· UNIVERSITY OF OXFORD ·

· DEPARTMENT OF ENGINEERING SCIENCE ·

Manufacturing of Integrated Optical Devices Assisted by Ultra-Short Pulse Lasers

Zhi Kai Pong

Keble College

A thesis submitted for the degree of

Doctor of Philosophy

Supervised by Patrick S. Salter and Martin J. Booth

October 2025

Declaration

I declare that this thesis is entirely my own work, and except where otherwise stated, describes my own research.

Zhi Kai Pong

Keble College

To my parents, for everything

Abstract

Integrated photonics is a key enabling technology for communications, sensing, and quantum information processing. Femtosecond laser direct writing (FLDW) offers unique advantages for this field, including three-dimensional structuring, rapid prototyping, and compatibility with diverse transparent substrates. This thesis investigates the design, fabrication, and characterisation of laser-written waveguides and devices, with particular emphasis on understanding how fabrication strategies and device geometry influence optical performance.

This thesis starts with an overview of the research field, theoretical background, and experimental setup for the studies. The research progresses are then reported through three main experimental studies. First, a fabrication method is introduced that compensates for edge aberrations without post-processing, reducing failure rates and enabling in-situ device writing. Second, multiscan waveguides fabricated with spherical phase-induced modulation (SPIM-WG) are investigated through simulations using COMSOL Multiphysics software, showing how geometry, separation, and symmetry govern supported waveguide modes. Agreement with experimental results demonstrate how these SPIM-WG structures can be tailored for coupling to fibre or to devices with asymmetric spatial modes such as ppKTP waveguides. Third, directional couplers are analysed theoretically and experimentally, demonstrating how their splitting ratios vary with geometry and input conditions, and extending this understanding to three-dimensional, out-of-plane configurations.

Together, these results show how geometry, refractive index modification, and characterisation can be harnessed as design parameters for integrated photonics. Beyond fundamental insight, they establish design strategies with direct relevance to industrial applications in optical communications, co-packaged optics, data-centre interconnects, precision sensing, and quantum photonics. Looking forward, the methods developed here point toward heterogeneous integration with other platforms and the use of additional degrees of freedom of light, such as spatial modes and polarisation, to realise scalable, application-ready photonic systems.

Acknowledgements

I would like to express my sincere gratitude to my supervisors, Professor Patrick Salter and Professor Martin Booth, for their invaluable guidance and patience throughout my doctoral studies. Their insightful feedback and support have been instrumental in shaping this research. I am truly thankful for their time, wisdom, and generosity.

I am also indebted to Dr. Bangshan Sun for his mentorship during the early stages of this journey. His thoughtful advice and willingness to share his experience provided me with a strong foundation and confidence to move forward. I extend my heartfelt thanks to Dr. Mohan Wang for her kind assistance with experimental work, whose practical insights were crucial to the completion of my DPhil research. I would also like to thank the other members of the lab for their assistance and constructive discussions over the years.

My sincere appreciation goes to Professor Chao He for providing opportunities for collaboration and for the interactions with his new and energetic research team. The experience was both valuable and enriching. I am especially grateful to Dr. Zimo Zhao for the countless conversations we shared over lunches and dinners, which broadened my perspectives and made the research journey far more enjoyable.

To my long-time friends Alfred Lam and Tony Wong, thank you for your steady friendship, encouragement, and for reminding me to keep a sense of balance through the challenges of the DPhil years. I am also extremely appreciative of all those else who showed their support along the way, and I truly could not thank you all enough.

Last but not least, I owe my deepest thanks to my family for their unconditional love, patience, and belief in me. Their support has been my greatest source of strength and motivation throughout this journey.

Zhi Kai Pong

October 2025

Publications and Presentations

Thesis-related publications

1. **Chapter 4: Correcting for edge aberrations**

   Z. K. Pong, M. Wang, M. J. Booth and P. S. Salter, "Laser written waveguides to the sample edge." (In preparation)

2. **Chapter 5: Multicore waveguide design for spatial mode converters**

   B. Sun, F. Morozko, P. S. Salter, S. Moser, **Z. K. Pong**, R. B. Patel, I. A. Walmsley, M. Wang, A. Hazan, N. Barré, A. Jesacher, J. Fells, C. He, A, Katiyi, Z.-N. Tian, A. Karabchevsky and M. J. Booth, "On-chip beam rotators, adiabatic mode converters, and waveplates through low-loss waveguides with variable cross-sections," Light: Science & Applications 11 (2022).

3. **Chapter 6: Polarisation and angle-dependent modulation of coupling ratio in directional couplers**

   Z. K. Pong, B. Sun, Z. Li, P. S. Salter and M. J. Booth, "Polarization based modulation of splitting ratio in femtosecond laser direct written directional couplers," Optics Communications 575, 131303 (2025).

Other publications

1. A. A. Wang, Z. Zhao, Y. Ma, Y. Cai, R. Zhang, X. Shang, Y. Zhang, J. Qin, **Z. K. Pong**, T. Marozsák, B. Chen, H. He, L. Luo, M. J. Booth, S. J. Elston, S. M. Morris and C. He, "Topological protection of optical skyrmions through complex media," Light: Science & Applications 13 (2024).

2. Y. Zhang, A. A. Wang, Z. Zhao, Y. Ma, R. Liu, R. Zhang, **Z. K. Pong**, Y. Cai and C. He, "Skyrmions based on optical anisotropy for topological encoding," arXiv:2508.16483 (2025 preprint).

3. J. Qin, Y. Montelongo, A. Pradeep, **Z. K. Pong**, S. J. Elston, P. S. Salter and M.

J. Booth, “Gap-Controlled Photonic Coupling for Intrachip Interconnections” (In preparation).

Conference presentations

1. **Z. K. Pong**, B. Sun, P. S. Salter, et al., “Polarization effects in laser-written directional couplers,” in 2024 Conference on Lasers and Electro-Optics (CLEO) (2024)

Other co-authored presentations

1. P. S. Salter, B. Sun, **Z. K. Pong**, M. J. Booth, “Laser written waveguides with engineered cross-section,” in SPIE LASE (2023).
2. P. S. Salter, **Z. K. Pong**, and M. J. Booth, ”Laser writing waveguides up to the sample edge,” in SPIE LASE (2025).

Poster presentations

Z. K. Pong, B. Sun, P. S. Salter and M. J. Booth, “Femtosecond laser-written Quantum Optical devices with Adaptive Optics”

University of Oxford Quantum Day, Oxford, UK (July 2022)
University of Oxford Photonics Day, Oxford, UK (September 2022)
International Summer School in Quantum Technologies, Birmingham, UK (August 2023)
University of Oxford Quantum Day, Oxford, UK (September 2023)
University of Oxford Photonics Day, Oxford, UK (October 2023)

Patents

1. P. S. Salter, M. J. Booth and **Z. K. Pong**, “Method for laser writing waveguides to the sample edge” (Under review)

Contents

List of Figures

CHAPTER 1

Overview

1.1 Introduction

This thesis investigates and develops the technology of femtosecond laser direct writing (FLDW) for fabricating waveguide-based components in integrated optical devices. This introductory chapter outlines the research motivation and provides a structured overview of the whole thesis, preparing further discussions in subsequent chapters. Various aspects of the technology are explored, including novel fabrication methods and characterisation, which lay the foundation for further developing sophisticated techniques for industrial applications.

1.2 Motivation

Optical computing and optical interconnects have received rapidly increasing attention in recent years as electronic implementations are reaching bottlenecks in information density and energy efficiency [1], as exemplified by the slowing down of Moore's and Dennard's scaling laws [2–4]. Optical-based implementations are promising candidates to alleviate these issues due to high parallelism and energy efficiency [5, 6]. However, it encounters challenges of its own which have yet to be solved and proposes significant engineering difficulties.

Current implementations are mostly silicon based due to its compatibility with mature CMOS technology [7]. However, CMOS technology is planar in nature, and it is very challenging to utilise all the spatial dimensions. Moreover, getting light in and out

of silicon devices poses enormous difficulties due to mode-size mismatches with other components such as optical fibres [8].

FLDW has garnered interest as a candidate to overcome these challenges. Using ultra-short pulse lasers, three-dimensional structures can be created inside a transparent material such as glass. The intense ultra-short pulses at the laser focus are sufficient to permanently modify the material, creating regions with increased refractive index to form waveguide structures. The geometry of these structures can be controlled by moving the laser focus inside the material, hence the technology is known as direct laser writing [9].

Glass substrates in particular have numerous advantages as the material of choice. Compared to other commonly used materials for FLDW such as polymers, glass has stronger mechanical strength and is more resistant to environmental factors such as temperature, humidity or chemical changes [10]. It is also easier to prepare larger substrates with improved flatness to enable higher component densities and finer patterning. Moreover, since glass has a similar thermal expansion coefficient to silicon, it offers the potential of further integrating with silicon photonic components for more efficient co-packaged optics with lower thermal stress [11].

Nevertheless, there are still a number of challenges that the FLDW technology faces, and this thesis aims to develop the methods for FLDW from three different directions. The first aims to tackle the need for post-processing of samples to reduce time-consumption and minimise errors, at the same time opening up potentials for in-situ fabrication where external components are already bound onto the substrate before fabrication. The second method looks at creating waveguides with arbitrary cross-sections to improve mode-matching and reduce coupling loss, which enables further possibilities of integrating with a much larger range of sources, connectors and receivers, which have implications in a wide range of applications in computing, communication and extending to quantum optical implementations. The final part investigates the polarisation effects of evanescent coupling of out-of-plane geometry, which fills in a gap in the literature where past inves-

tigations are either planar or did not fully account for polarisation effects. This work lays the foundation for further understanding the behaviour of polarisation in integrated optical devices and harnessing polarisation as an extra degree of freedom of light for enhanced parallelism.

1.3 Thesis outline

This thesis outlines the motivation for developing and better understanding FLDW as a manufacturing technology.

Chapter 2 provides an in-depth discourse on the background of FLDW technology, including the basic principles of pulsed-lasers and material modification mechanism which allows the creation of waveguide structures in the first place. It will then discuss the role of adaptive optics in assisting device fabrication, explaining how a spatial light modulator helps to compensate for depth-dependent aberrations when combined with a 4-f system in the optical path inside the fabrication system, which allows the consistent fabrication of identical waveguide structures at different depths under the surface. The basics of waveguide theory will then be introduced and discuss various properties important for understanding their principle of operation and characterisation.

Chapter 3 presents the experimental setup and procedures. The fabrication setup and characterisation setup will be described in detail. In addition, the procedure for device preparation and processing together with methods for data processing and analysis will also be explained.

Chapter 4 explores a new fabrication method for processing sample in-situ without the need for post-processing. In the current fabrication workflow, the fabricated sample requires post-processing to polish back on the waveguide interface due to edge aberration effects. This is very time-consuming and a major point of failure during the fabrication process in the form of undesired mechanical damage to the sample edge. Moreover, this creates a barrier for fabrication in-situ since the surface has to be exposed for this polishing

step. The new method compensates for the edge aberration by only using half of the pupil for fabrication near the edge, and it is shown that the modes of the waveguide created using this new method is similar to that of the normal waveguide, which verifies the feasibility of this method.

Chapter 5 presents another fabrication method for creating waveguides that has an arbitrary mode-section, enabling potential applications for mode-converters and more exotic waveguide engineering capabilities. Simulation results are provided to show that it is possible to build up a larger refractive index structure from smaller individual waveguides which has specific preference for certain spatial modes. Combined with a method that involves artificially inducing a spherical aberration, the method known as spherical phase-induced multiscan waveguide (SPIM-WG) is shown to be able to create devices that can transform between arbitrary mode shapes with a variable cross-section adiabatically. This enables devices that can act as on-chip beam rotators, adiabatic mode converters, and waveplates, which opens up many possibilities for integrated optical processing and connections.

Chapter 6 presents a characterisation of a polarisation-dependent phenomenon in directional couplers. The coupling between two arms of directional couplers exhibits a polarisation dependence which results in a modulation of the splitting ratio depending on the input polarisation. By creating multiple sets of directional couplers, it is possible to investigate how the coupling ratio varies as a function of interaction region lengths. Moreover, the effect of geometrically offsetting the two arms in the vertical direction relative to the sample surface is also looked at which verifies the experimental results agree with the proposed general way of understanding such configurations.

Finally, Chapter 7 summarises the work and highlights potential applications. The chapter will conclude with potential directions for future work, exploring possibilities for further developing the methods for advancing the technology.

CHAPTER 2

Background

2.1 Ultra-short pulse lasers

Lasers are devices that generate and amplify light [12]. The output beam of a laser is coherent, directional and highly collimated, which can be focused to a tiny focal region of the size on the order of its wavelength. Most of the power of the beam is delivered to that region, inducing material modifications which make lasers a suitable candidate for fabricating structures inside a substrate.

Laser operation can be classified into continuous wave operation or pulsed operation. Continuous wave lasers have a constant amplitude in time, whereas pulsed lasers have a limited time duration for each pulse. In this section, I will define what it means by ultra-short pulse, and explain the main parameters related to pulsed operation. I will then describe the physical process of how ultra-short pulse lasers permanently modify material properties when focused inside a transparent dielectric.

2.1.1 Mode-locking

Ultra-short pulses are pulses of extremely short duration, on the order of femtoseconds (10^{-15} s). For this reason, ultra-short pulse lasers are often also referred to as ultrafast lasers or femtosecond lasers in literature.

Ultra-short pulses are generated through a process known as “mode-locking”. A free-running laser without any control to the laser spectrum has longitudinal modes that

oscillate independently with random phases, resulting in noisy and incoherent output [13]. Through mode-locking, it is possible to produce well-defined pulses with all modes oscillating in phase given the gain medium has sufficient bandwidth. There are many methods to implement mode-locking in a laser, for example utilising nonlinear optical effects such as the optical Kerr effect to modulate the laser for obtaining the desired output [14].

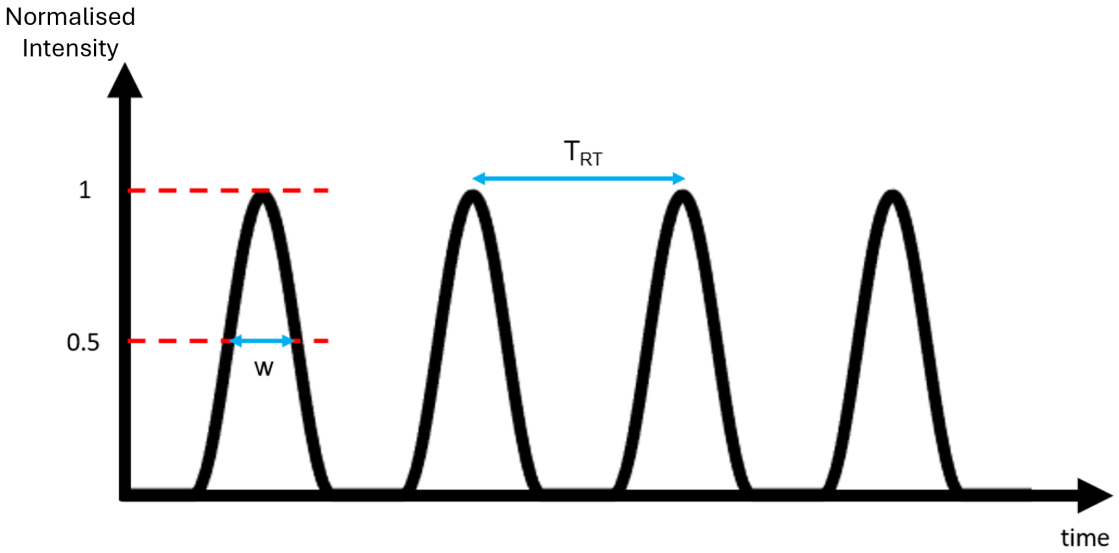

Figure 2.1: Schematic of mode-locked laser output. T_{RT}: repetition period. w: Pulse duration.

Fig. 2.1 shows a schematic of the temporal profile of a mode-locked laser output. Some important parameters of the laser are defined: The pulse duration w is usually defined as the FWHM (full-width at half maximum) of a pulse. The repetition period T_{RT} is the time between two consecutive pulses. The parameter repetition rate is more commonly used over the repetition period, which is given by $1/T_{RT}$. The lasers used in FLDW typically have maximum pulse energies from a few μJ to a few mJ, and the repetition rate can range from kHz to MHz. Typical pulse durations are around 50-300 fs.

2.1.2 Chirped-pulse amplification

Higher pulse energies provide greater flexibility in laser writing of glass, enabling controlled access to different modification regimes. It is difficult to amplify ultrashort laser pulses

directly without damaging the laser gain medium and optical components. In modern ultrashort pulse laser amplifier systems, Chirped Pulse Amplification (CPA) is a technique developed to circumvent this issue while enable such systems to remain relatively compact. CPA consists of three major steps: the pulses are first stretched temporally using dispersive elements such as a pair of gratings to separate the pulses' frequency components in a step known as "chirping", which reduces the peak power and prevents subsequent damage to amplification components. The stretched pulse is then safely amplified through the gain media. Finally, the pulse is recompressed into its original form through reversing the stretching process, obtaining the desired high-power pulses that would otherwise not be possible [15]. Through this process, ultrafast laser pulses of sufficiently high intensity can be generated for FLDW [16].

2.2 Femtosecond direct laser writing

FLDW in glass has become one of the most important approaches for realising integrated photonic devices, owing to the unique ability of femtosecond pulses to induce localised, permanent refractive index modifications within transparent dielectrics. This enables the direct inscription of buried optical waveguides with three-dimensional freedom, without the need for masks, chemical processing, or multiple fabrication steps. This section describes the material modification process and how the fabrication parameters affect the resulting modification.

2.2.1 Material modification

When focused inside a transparent dielectric, the laser delivers energy to the focal volume. When this energy density is high enough, it will initiate absorption through nonlinear field ionisation and avalanche ionisation, as illustrated in Fig. 2.2 [17]. This process creates an electron-ion plasma, which induces permanent structural changes in the material when the energy of the plasma is transferred to the lattice. A high numerical aperture (NA $\geq$ 0.5) objective lens is usually employed to focus the light in space, ensuring a tightly confined

focal volume and enhanced spatial resolution, while short pulse durations temporally confine the energy delivery.

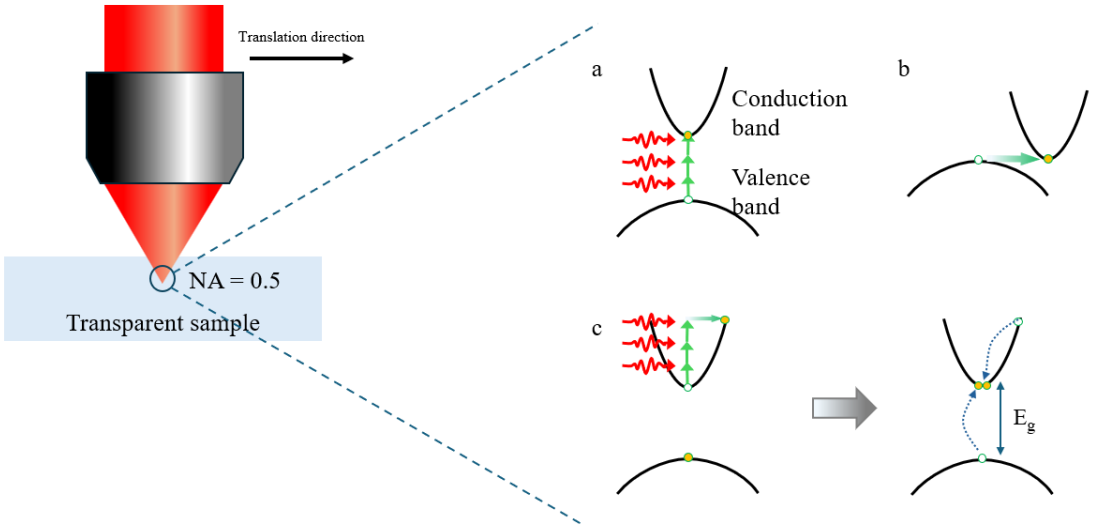

Figure 2.2: A focused femtosecond laser beam is tightly focused into a transparent sample and translated laterally during irradiation. The high intensity at the focal region induces nonlinear absorption processes that enable material modification even within optically transparent media. (a) Multiphoton ionisation: Several photons are simultaneously absorbed to excite an electron from the valence band to the conduction band. (b) Tunneling ionisation: Under a strong optical field, electrons tunnel through the potential barrier from the valence band into the conduction band. (c) Avalanche ionisation: Free carriers absorption followed by impact ionisation. Adapted from [18].

Femtosecond pulses have the advantage that the structure modification is localised [19] since the time duration of laser pulses is shorter than the time for heat to diffuse out of the focal volume to the surrounding area of the laser-irradiated region [20], and this fabrication regime is known as the heat accumulation regime. In the case of longer-pulsed lasers such as nanosecond lasers, the modified region could extend outside of the focal volume which could lead to undesirable effects such as reducing the spatial resolution of fabricated structures [21].

2.2.2 Choice of material

The modification behaviour of materials under laser irradiation is intrinsically linked to their physical properties, making material choice a critical factor in determining the out-

comes of laser writing processes. Key parameters such as the electronic bandgap, linear and compositional makeup collectively dictate how deposited laser energy is coupled, confined, and relaxed within the medium. Certain wide bandgap dielectrics such as fused silica are more sensitive to fabrication parameters, where the material modification can be categorised into 3 types depending on the pulse energy delivered. Type I modification occurs when the fabrication energy exceeds the breakdown threshold of the material, resulting in an isotropic refractive index change at low energies, suitable for fabricating waveguides and other low-loss devices [22]. Type II modification occurs when the pulse intensity is higher, resulting in nanograting formation [23] which are self-assembled nanostructures with sub-wavelength period. Those structures are characterised for their anisotropy and higher loss, which hinders performance for low loss applications but could be useful for other applications such as femtosecond laser-assisted etching to produce micro-fluidic channels [24] or controlled birefringent structures [25]. At even higher energies, Type III modification results in void formation [26], inducing disruptive damage to the material from micro-explosions. On the other hand, borosilicate glasses such as Corning Eagle 2000 have a lower bandgap due to impurities [27], therefore they demonstrate a different behaviour where the heating regime avoids the formation of nanogratings even at higher energies, allowing a larger window for parameter adjustment which in some cases enables easier fabrication.

2.2.3 Types of lasers and parameters

Different types of lasers have been employed for waveguide fabrication [28]. For example, Ti:sapphire amplified lasers typically have 50-200 fs pulse duration and 1-200 kHz repetition rate. They can have a broad range of wavelengths (although some are fixed, e.g. at 800 nm), and can be used to process a wide range of materials. However, they provide refractive changes of 1-2$\times 10^{-3}$ typically, which is relatively low compared to using other lasers. Ytterbium-based (Yb) lasers have been used to provide better fabrication results. They operate at 1030 nm wavelength which can be frequency-doubled to 515 nm. Their pulse duration is around 300 fs with repetition rates up to 1 MHz. They are also com-

pact and cost-effective since they are directly pumped by a laser diode and do not require amplification stages.

For the applications relevant to this report, a positive refractive index change is required for waveguide operations, and with these parameters at suitable pulse energies femtosecond lasers can induce a refractive index change 5-10$\times 10^{-3}$ in glass which is suitable for creating waveguides [29].

2.3 Adaptive optics for laser processing

Optical aberrations refer to deviations of light from its ideal form. Aberrations can arise from a number of reasons, including imperfections in components of the optical system, misalignment or refractive index mismatch between air and another material.

In ultra-short pulse laser processing, researchers have noticed early on that spherical aberration caused by focusing inside glass displaces and elongates the focus [30]. Other researchers have noted that the threshold pulse energy for material modification is increased when spherical aberration is more significant [31]. However, it is only in recent years that the more flexible adaptive optics techniques are implemented in ultra-short pulse fabrication systems to improve the quality of fabrication.

Adaptive optics (AO) was originally envisioned for improving astronomical observation [32]. It has found widespread applications in different areas since then, including biological microscopy [33] and retinal imaging [34].

In this section, I will begin by introducing Fourier optics which is an important tool for understanding propagation of light through lenses and calculating the light intensity distribution at the focal region [35]. This is important for focusing laser near diffraction limit with high NA and very relevant for laser processing in general, and heavily impacted by optical aberrations. I will then give an overview of Zernike polynomials, which are widely used in representing aberration modes in optical systems. Lastly, I will describe the choice and operation of adaptive elements.

2.3.1 Phase aberrations

First proposed by Frits Zernike [36], the Zernike modes are a set of polynomials that are continuous and orthogonal over a unit disk [37]. They are commonly used in optics and image processing for modelling the phase distortion to a wavefront passing through circular apertures. The normalised Zernike polynomials have the desirable property that the coefficient of each mode represents the root mean square wavefront error attributed to that mode. Since a significant portion of optical systems employs circular components and pupils, the Zernike modes have been widely adopted for describing wavefronts propagating through these systems.

The mathematical representation for an arbitrary function on a disk in polar coordinates (r, θ) can be written as a sum of orthonormal polynomials:

$$W(r,\theta) = \sum_{m,n} C_n^m Z_n^m(r,\theta) \tag{2.1}$$

where C_n^m are the Zernike coefficients and $Z_n^m(r,\theta)$ are the Zernike polynomials. The two indices m and n are the respective degrees of the azimuthal component and the radial component. m and n are integers subject to $n \geq |m| \geq 0$ and $n - |m| = even$.

The Zernike polynomials have the form

$$Z_n^m(r,\theta) = \begin{cases} \sqrt{2n+2} R_n^m(r) \cos(|m|\theta) \text{ for } m \geq 0 \\ \sqrt{2n+2} R_n^m(r) \sin(|m|\theta) \text{ for } m < 0 \\ \sqrt{n+1} R_n^m(r) \text{ for } m = 0 \end{cases} \tag{2.2}$$

where the radial term $R_n^m(r)$ is given by

$$R_n^m(r) = \sum_{s=0}^{(n-m)/2} \frac{(-1)^s (n-s)!}{s![(n+|m|)/2 - s]![(n-|m|)/2 - s]!} r^{n-2s}. \tag{2.3}$$

There are various ways of labelling Zernike polynomials with one index k by ordering

m and n differently, and the convention used in this report is from Neil's paper [38].

In ultra-short pulse laser fabrication, the lower order modes (piston, tip, tilt, defocus) are often ignored as they do not affect the fabrication process. Piston ($k = 1$) adds a constant phase which does not affect the shape of the focus, and tip ($k = 2$) and tilt ($k = 3$) are equivalent to a lateral translation of the laser focus in the plane perpendicular to the beam direction. The defocus ($k = 4$) component translates the focus in the axial direction.

Table 2.1 shows the polynomials that are corrected in the fabrication system used throughout this study. Higher order modes have little contribution to the shape of the focus, therefore correcting for these 5 modes is sufficient to significantly improve the quality of fabrication. These modes are plotted in Fig. 2.3.

Table 2.1: Zernike polynomials for k = 5, 6, 7, 8, 11

k	n	m	$Z_k(r,\theta)$	Aberration term
5	2	2	$\sqrt{6}r^2\cos(2\theta)$	Vertical astigmatism
6	2	-2	$\sqrt{6}r^2\sin(2\theta)$	Oblique astigmatism
7	3	1	$2\sqrt{2}\,(3r^3-2r)\cos(\theta)$	Horizontal coma
8	3	-1	$2\sqrt{2}\,(3r^3-2r)\sin(\theta)$	Vertical coma
11	4	0	$\sqrt{5}\,(6r^4-6r^2+1)$	Primary spherical

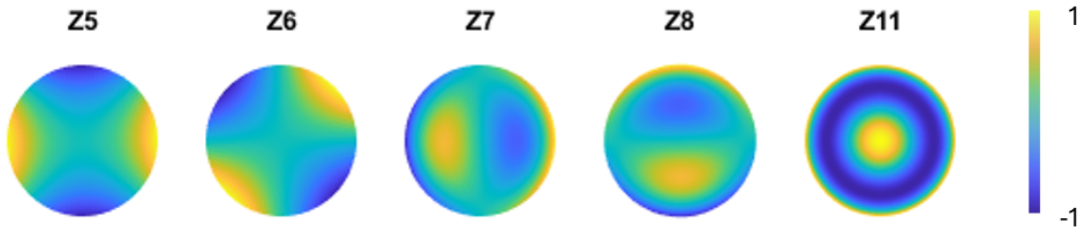

Figure 2.3: The Zernike modes in Table 2.1 over the unit disk.

Spherical aberration arises when focusing inside a material of different refractive index from another medium and is depth dependent, therefore it warrants extra discussion. The nominal depth d_{nom} is defined as the depth to which we are focusing within the sample.

The analytic expression for spherical aberration at d_{nom} is [39]:

$$\phi_{SA}(\rho) = \frac{-2\pi d_{nom}}{\lambda}\left(\sqrt{n_2^2 - (\mathrm{NA}\rho)^2} - \sqrt{n_1^2 - (\mathrm{NA}\rho)^2}\right) \tag{2.4}$$

where λ is the wavelength of light, n_1 and n_2 are the refractive indices of the two media, NA is the numerical aperture and ρ is the normalised pupil radius. This expression contains a defocus phase which leads to the actual focal position being deeper if $n_2 > n_1$. This results in d_{act} being greater than d_{nom} as illustrated in Fig. 2.4.

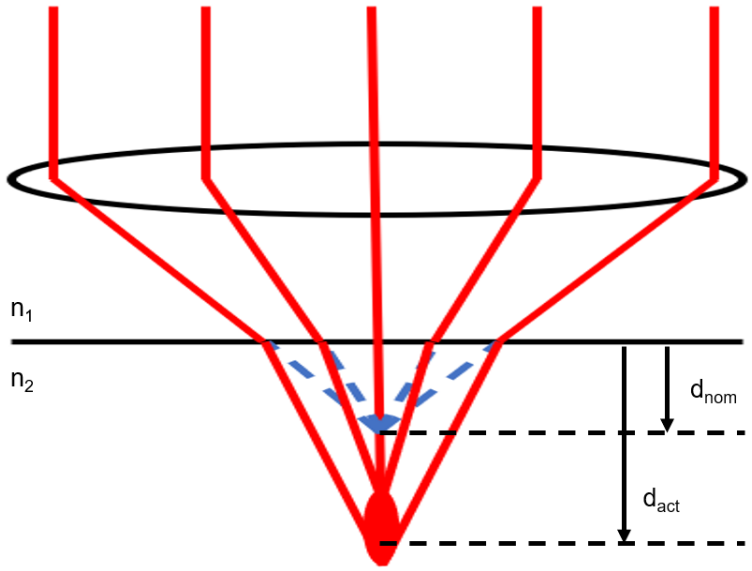

Figure 2.4: Actual and nominal depths of the laser focus. Since $n_1 < n_2$, the beam refracts and the actual focus (where red solid lines converge) is deeper than the nominal depth (where blue dashed lines converge).

Since defocus is just an equivalent translation of the sample stage, it can be removed from the total abberation ϕ_{SA} in order to reduce the magnitude of the correction needed from the adaptive element. The defocus-free spherical aberration function denoted by $\hat{\phi}_{SA}(\rho)$ is [40]:

$$\hat{\phi}_{SA}(\rho) = \phi_{SA} - \frac{\langle \phi'_{SA}, D'_{n2} \rangle}{\langle D'_{n2}, D'_{n2} \rangle} D_{n2} \tag{2.5}$$

where $\langle , \rangle$ is the inner product. The defocus phase $D_{n2}(\rho)$ is given by

$$D_{n2}(\rho) = \frac{2\pi d_{nom}}{\lambda}\sqrt{n_2^2 - (\mathrm{NA}\rho)^2} \tag{2.6}$$

and $\phi'_{SA} = \phi_{SA} - mean(\phi_{SA})$ and $D'_{n2} = D_{n2} - mean(D_{n2})$ correspond to their non-prime counterparts subtracting their respective mean values.

The actual fabrication depth d_{act} is therefore related to the nominal focal depth d_{act} by

$$d_{act} = d_{nom}\left(1 + \frac{\langle \phi'_{SA}, D'_{n2} \rangle}{\langle D'_{n2}, D'_{n2} \rangle}\right) \tag{2.7}$$

which is a function of NA. A plot of the relationship between the ratio of d_{act} and d_{nom} is shown in Fig. 2.5. Assuming an air-material interface, n_1 is set to 1, and three curves are plotted for different n_2 to observe the effect of refractive index on the scaling factor. At low NA, which is the part of the graph nearer NA = 0, the paraxial approximation is applicable. The ratio of d_{act} and d_{nom} is approximately equal to the refractive index of the substrate. As NA increases, this scaling factor becomes significant at around $NA > 0.5$. Moreover, we can see that with higher refractive index, the ratio is more sensitive to the NA used in fabrication and begins to increase at lower NA, therefore extra care should be taken when fabricating inside high refractive index material.

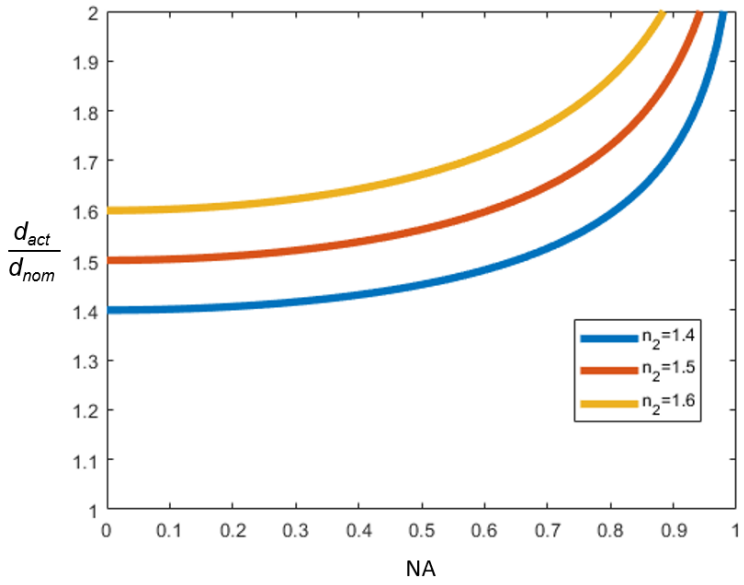

Figure 2.5: The ratio of nominal to actual focusing depth as a function of the numerical aperture calculated using Eq. 2.7.

In summary, Zernike polynomials provide a powerful mathematical basis for quanti-

fying and correcting optical aberrations. By identifying and compensating the dominant low-order modes, particularly astigmatic and spherical aberrations, adaptive optics enables precise control of the focal field, thereby enhancing fabrication accuracy and consistency in FLDW.

2.3.2 Adaptive optical elements

To compensate for the aberrations described in the previous section, adaptive optical elements are employed to dynamically modify the phase of the optical wavefront. This process is known as aberration correction, whereby a compensating phase pattern equal in magnitude but opposite in sign to the measured or estimated aberration is applied to the beam. The goal is to restore the wavefront to its ideal form and recover a diffraction-limited focus at the desired position within the sample as much as possible. The corrective phase can be determined using various methods, including direct wavefront sensing, interferometric techniques, or iterative optimisation based on image or signal metrics [41]. Implementing aberration correction not only restores beam symmetry and focal confinement but also enhances energy delivery efficiency and fabrication precision—both of which are critical for high-quality ultra-short pulse laser processing. An adaptive optical element is a device which dynamically modifies the phase of the wavefront across the beam. The two most commonly used adaptive elements are deformable mirrors (DMs) and spatial light modulators (SLMs).

DMs are usually electrically controlled by an array of actuators, which modifies the angle and shape of a reflective surface. There are different implementations of DMs, for example they can have a continuous membrane or are made of individual segments. They are actuated electro-mechanically by electrostatic, electromagnetic or piezoelectric means [42]. Being mirrors, they have high reflectivity and are insensitive to polarisation. They can operate over large bandwidths, but usually have much less spatial resolution than a SLM. They can operate at rates up to several kHz.

SLMs also come in numerous forms, with the vast majority operating with a liquid

crystal layer on silicon back plane. The liquid crystal is controlled by an applied voltage. They are generally slower than DMs, and they tend to have lower optical efficiency since the light has to pass through a liquid crystal layer in a SLM but is reflected directly in a DM. The number of drivable pixels is typically much higher than that of DMs, but the modulation range is limited. This can be overcome by phase wrapping 2π to 0 radians. Since SLMs are able to have discontinuities between adjacent pixels this allows for interesting applications such as parallelisation [43].

For FLDW, SLMs are suitable because laser light is narrowband and single polarisation. High speed operation is also not essential for the majority of fabrication processes since the aberrations induced by the system and refractive index mismatch of the sample is relatively constant over the course of fabrication.

2.3.3 Applications of adaptive optics in laser processing

Although correcting for system and depth-dependent aberrations are crucial, the functionality of the SLM is certainly not limited to aberration correction. There has been substantial work in the direction of utilising adaptive elements to enhance parallelisation, significantly reducing the total laser processing time. This is achieved by splitting the beam into an array of foci through displaying the appropriate phase pattern hologram on the SLM, usually generated via advanced computational techniques [44]. Demonstrations creating arrays with more than 1000 individual fabrication foci have been realised [45].

Beam shaping can also be utilised to adjust the shape and dimensions of the laser focus, hence allowing additional control over the material modification process. Applications include creating circular waveguides by stretching the beam in the horizontal direction [46], creating a flat top intensity profile for large area processing [47], and creating special optical vortex beams carrying optical angular momentum for creating complex material modifications [48].

Specifically for the work in this thesis, in Chapter 4, I will utilise an SLM to apply a phase mask for spatial filtering, and in Chapter 5 intentionally apply a spherical phase

pattern to control the focus shape for which allows for a novel way of creating waveguides in a different fabrication regime.

2.3.4 4-f lens relay system

The 4-f imaging system plays an important role in many optical systems, and it is also useful for adaptive optics in laser fabrication. Fig. 2.6 shows the geometry of a 4-f system. The system is constructed by having the back focal plane of the first lens coinciding with the front focal plane of the second. The overlapping plane is known as the Fourier plane. The 4-f system effectively performs one Fourier transform through each lens. The image at the back focal plane of the second lens will therefore be an inverted copy of the object, to a good approximation, at the front focal plane of the first lens, preserving both amplitude and phase information.

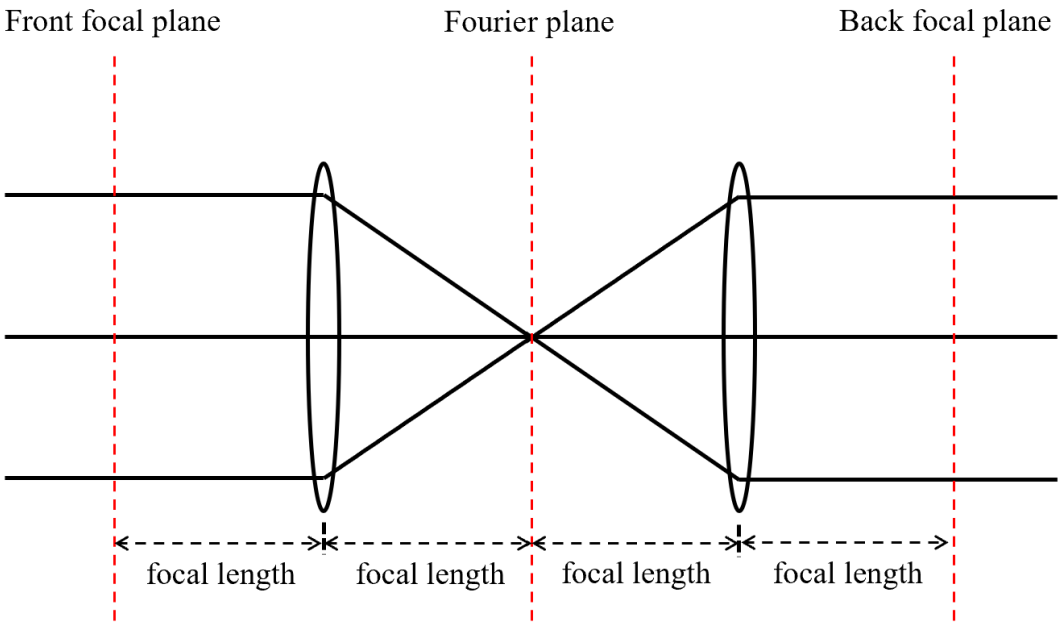

Figure 2.6: Geometry of a 4-f system. Black solid lines are rays through the system, and red dashed lines indicated the positions of the important planes.

The inexactness of this imaging process arises from the finite aperture of the lenses, which limits the range of spatial frequencies that can pass through the system. High spatial frequency components corresponding to fine details or sharp edges in the object

are clipped by the finite aperture and thus not fully transmitted to the image plane. As a result, the reconstructed image is a band-limited approximation of the original object rather than an exact replica. This approximation holds well near the optical axis under the paraxial approximation where the angles of propagation are small, but deviations become more significant for rays far from the optical axis or when the numerical aperture is large.

This particular combination of optical components has the advantage that the phase information will be conserved through the system in a way that would not be possible using a single lens. Although phase information may not always be important for typical imaging, it is important for our application since we modulate the wavefront using an adaptive element and we do not want this wavefront to be further distorted. This is crucial for the effective utilisation of the spatial light modulator and relaying the modulated wavefront to the objective pupil plane which we will implement in our fabrication system described in section 3.1.1.

Moreover, having access to the Fourier plane also allows us to perform spatial filtering, which will be important for the application described in Chapter 4. A grating pattern can easily be produced on the SLM by simply displaying a Z2 or Z3 Zernike mode depending on the direction of writing, which effectively shifts the beam by a fixed amount. By inserting a pinhole at the Fourier plane, we can then block the zeroth order which allows for obtaining the desired beam shaping. The exact setup will be described in more detail in the relevant chapter.

The working principle underlying the 4-f system is explained by Fourier optics, where Fourier analysis is applied to analyse what happens spatially during light propagation by considering wavefronts as being made up of a superposition of plane waves. The Rayleigh-Sommerfeld diffraction integral describes the propagation of light between two planes by considering Huygens' principle [49]:

$$E\left(x_f, y_f\right) = \frac{1}{i\lambda} \int \int E\left(\omega, x_i, y_i\right) \frac{\exp(ikR)}{R} \cos\theta dx_i dy_i \tag{2.8}$$

where x_i, y_i are the spatial coordinates on the input plane, and x_f, y_f are the spatial coordinates on the final plane where the field propagates towards. The distance R between the beginning and end points is given by

$$R = \sqrt{z^2 + (x_f - x_i)^2 + (y_f - y_i)^2} \tag{2.9}$$

Assuming that x and y are small compared to z, we can use the paraxial approximation which assumes θ is a small-angle and obtain

$$R \approx z + \left[(x_f - x_i)^2 + (y_f - y_i)^2\right]/2z \tag{2.10}$$

Using this expression equation 2.8 is simplified to the well-known Fresnel diffraction integral:

$$E(x_f, y_f) \approx \frac{e^{ikz}}{i\lambda z} \int\int E(x_i, y_i)\, e^{i\frac{k}{2z}\left[(x_f - x_i)^2 + (y_f - y_i)^2\right]} dx_i dy_i \tag{2.11}$$

If we expand the exponent inside the integral and factorise the component, we obtain a Fourier transform expression which is equivalent to the Fresnel integral:

$$E(x_f, y_f) \approx \frac{e^{ikz} e^{ik\frac{x_f^2 + y_f^2}{2z}}}{i\lambda z} \mathcal{F}\left\{E(x_i, y_i)\, e^{ik\frac{x_i^2 + y_i^2}{2z}}\right\} \tag{2.12}$$

where $\mathcal{F}\{\}$ represents the Fourier transform. If z is sufficiently large and x_i and y_i sufficiently small that the $e^{ik\frac{x_i^2 + y_i^2}{2z}}$ term is small, we obtain the Fraunhofer diffraction integral:

$$E(x_f, y_f) \approx \frac{e^{ikz} e^{ik\frac{x_f^2 + y_f^2}{2z}}}{i\lambda z} \mathcal{F}\{E(x_i, y_i)\} \tag{2.13}$$

This usually means that there is an aperture in the input plane that has a finite extend. If there is a single lens, assuming the Fraunhofer approximation is valid, the field amplitude and phase at the back and front focal planes are given by a direct Fourier transform with a quadratic phase factor that varies across the output plane.

A rule of thumb for the validity of Fresnel and Fraunhofer approximations are given by the Fresnel number

$$\mathrm{F} = \frac{(x_i^2 + y_i^2)_{\max}}{z\lambda} \tag{2.14}$$

which depends on the wavelength λ.

The Fresnel approximation is valid when $F \approx 1$ and the Fraunhofer approximation is valid when $F << 1$. In the far field scenario, where the Fraunhofer approximation is applicable, the propagation paths of incoming rays are almost parallel. When considering the focusing of light by a positive lens, parallel rays are focused into a point on the focal plane, therefore the Fraunhofer diffraction can be used to analyse positive lens systems. This framework allows us to understand the 4-f system through Fourier transforms, and how it indeed preserves the relative amplitude and phase of the wavefront.

In summary, adaptive optics has emerged as an indispensable technique in ultra-short pulse laser processing for compensating optical aberrations and enhancing fabrication precision. By combining the theoretical framework of Fourier optics, the analytical description of aberrations using Zernike polynomials, and the implementation of adaptive elements such as spatial light modulators and deformable mirrors, it is possible to achieve diffraction-limited focusing even within refractive index-mismatched materials. Beyond simple aberration correction, adaptive optics enables flexible beam shaping, dynamic focus control, and large-scale parallelisation, thereby significantly extending the capabilities of FLDW fabrication. The integration of adaptive optics with Fourier-based 4-f relay systems ensures accurate phase transfer and spatial filtering, forming a robust foundation for the advanced laser processing methodologies developed in the subsequent chapters.

2.4 Waveguides

Waveguides are fundamental components in modern optical devices, serving as structures that confine and direct light along a desired path. They operate on the principle of total internal reflection, where light is guided within a high-refractive-index core surrounded

by a lower-refractive-index cladding [50]. This confinement allows efficient transmission of optical signals over extended distances with minimal loss, forming the basis of technologies such as optical communication, integrated photonics, and laser-based sensing systems. Therefore, it is important to develop an understanding of waveguide structures and their underlying operating principles to enable their effective design, optimisation, and characterisation for specific applications.

2.4.1 Waveguide structure

The simplest waveguide is a structure where a higher refractive index material (n_2) is surrounded by a lower refractive index material (n_1) such as the one shown in Fig. 2.7. The higher refractive index region is referred to as the core and the surrounding region the cladding. Since FLDW induces a positive refractive index change as explained in section 2.2.1, this effectively creates a waveguide structure in the glass sample. This was first demonstrated by Davis et al. in 1996 [9].

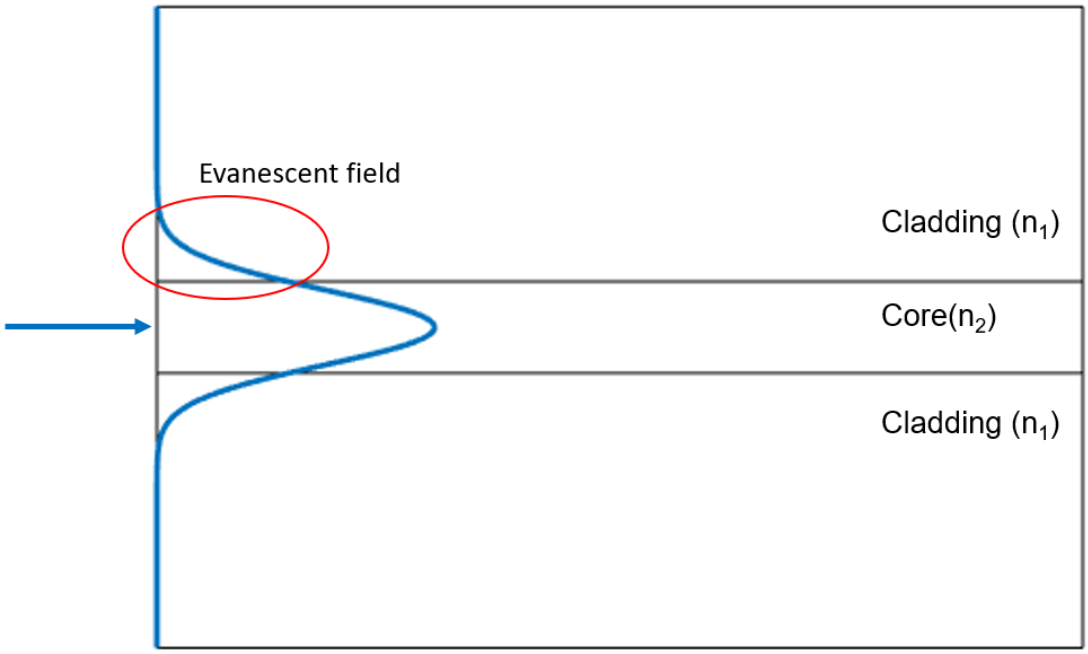

Figure 2.7: Schematic of a waveguide structure with $n_1 < n_2$. Blue arrow indicates input light and the direction of propagation. Blue curve represents the electric field amplitude of a guided mode inside the waveguide. Red ellipse identifies the evanescent field which is the part of the guided mode that extends out of the core and decays exponentially.

When the waveguide dimensions are much larger than the wavelength, light is guided through the structure through total internal reflection as long as the incident angle at each reflection is larger than the critical angle $\theta = \arcsin(n_1/n_2)$. However, this model is no longer accurate when the dimensions are comparable with the wavelength, which is the case of our fabrications since the structures are of the size of a few microns. We therefore need to solve for the field distribution using Maxwell equations. By solving Maxwell equations at the boundaries of a waveguide, it is a well-known result that standing waves known as modes will be formed in the waveguide and there is an exponentially decaying field just outside the core, also known as the evanescent field. When two waveguides are in close proximity such that the evanescent field of one waveguide overlaps with the core of the other, evanescent coupling occurs where the energy propagating in the first waveguide will be coupled into the second waveguide. This is the principle behind waveguide interactions in directional couplers and more sophisticated structures [51].

2.4.2 Single mode waveguides

The number of modes that are permitted in a waveguide is dependent on several parameters, including the size, refractive index contrast between the core and cladding, and the wavelength of light being transmitted. A useful parameter that captures those information is the V-number which is defined as

$$V = \frac{2\pi a}{\lambda}\sqrt{n_{core}^2 - n_{clad}^2} \tag{2.15}$$

for circular waveguides, where a is the radius, λ is the wavelength, and n_{core} and n_{clad} the core and cladding refractive indices respectively.

When V is less than 2.405, the waveguide is in single mode operation, which is often desired to avoid mode dispersion [52]. This refers to the phenomenon where different modes have different propagating speeds and leads to the spreading out of optical signals over long distances causing the signal quality to degrade.

The fundamental mode in a circular waveguide is neither purely transverse electric (TE) nor transverse magnetic (TM), but a hybrid mode known as the HE mode. To be precise, there are actually two degenerate HE modes with orthogonal polarisation. An exact description of the mode is complicated, but the physical model can be simplified by assuming the weakly guiding approximation holds [53].

The weakly guiding approximation is valid when the profile height parameter Δ defined as

$$\Delta = \frac{1}{2}\left\{1 - \frac{n_{clad}^2}{n_{core}^2}\right\} \tag{2.16}$$

is much less than one.

Physically, this suggests the refractive index of the core of the waveguide is only slightly higher than that of the cladding. In this case, the modes can be approximated by transverse electromagnetic (TEM) modes, where the polarisation is almost completely transverse to the propagation direction along a single direction. By identifying a pair of principal axes and denoting the direction of propagation by z, the two fundamental modes can then be considered as purely x-polarised and y-polarised respectively. For this reason, in a weakly guiding waveguide the fundamental modes are also known as linearly polarised (LP) modes.

2.4.3 Waveguide characterisation

The mode profile refers to the spatial distribution of the electromagnetic field of a specific mode of propagation within the waveguide. It depends on the structure of the waveguide, such as the size of the core and refractive index distribution. In most cases, the exact mode profile can only be calculated by numerical methods from solving the wave equation. However, the mode profile is exactly a Gaussian when the waveguide has an infinite parabolic refractive index distribution profile. When the weakly guiding approximation is valid, the fundamental mode-field distribution of an arbitrary profile can be approximated by some Gaussian function to the first order [54]. This approximation has the advantage

that the physical behaviour of a fabricated waveguide is obvious from just calculating the dimensions of the fitted Gaussian.

The most straightforward characterisation of a waveguide is its dimensions in the x- and y- direction transverse to the waveguide length (z-direction is taken to be the direction of waveguide along which light propagates). For ideal single mode operation, the waveguide core cannot exceed a certain size, since the V-number is directly proportional to the radius. Differences in x- and y- dimensions also lead to shape birefringence [55]. In an ideal circular waveguide, the two linearly polarised fundamental mode should be degenerate and propagate at the same speed. However, once shape birefringence is introduced, this will lead to polarisation effects such as the length-dependent polarisation modulation during propagation in a waveguide as a result of different propagation velocities of the fundamental x- and y- modes. The laser focus of the fabrication beam is naturally elongated in the transverse direction, which causes the fabricated waveguide to be stretched vertically. Various methods have been proposed to correct for this shape birefringence, such as using a cylindrical lens to create an elliptical focus through astigmatic beam shaping [56], or utilising alternative advanced beam-shaping methods to create a circular cross-section in order to avoid shape birefringence [57]. Chapter 6 is devoted to characterising a polarisation phenomenon that arises as a result of asymmetry in the geometry.

The coupling efficiency can be estimated via an overlap integral calculation between the waveguide mode and the incoming fibre mode. The intensity and phase profiles should match to achieve maximal coupling between a single mode optical fibre and the fabricated waveguide. The overlap integral η is given by the following expression where the integrals are integrating over the coupling cross-section area:

$$\eta = \frac{|\int E_1^* E_2 dA|^2}{\int |E_1| dA^2 \int |E_2| dA^2} \tag{2.17}$$

where E_1 is the complex electric field of the fibre mode in the transverse plane, and E_2

is the supported complex mode profile in the fabricated waveguide [58]. If the phase information is not available and only the modulus of the fields are known, Eq. 2.17 gives an upper bound on the coupling efficiency instead.

Waveguide losses arise from several mechanisms that reduce the transmitted optical power along the propagation length. The intrinsic losses are due to material absorption and scattering from microscopic inhomogeneities within the substrate, which are often determined by the choice and purity of the material. Bending losses become significant when the waveguide path deviates from a straight line, as light near the outer curvature may exceed the critical angle for total internal reflection and leak from the core. Mode mismatch losses can occur at coupling interfaces where the waveguide mode does not perfectly overlap with the fibre mode or with another photonic component. Additionally, radiation losses may arise in multimode or weakly confined structures where the refractive index contrast is insufficient to fully guide higher-order modes.

In summary, accurate characterisation of waveguides, including their mode profiles, dimensions, coupling behaviour, and propagation losses, is essential for optimising fabrication parameters and achieving high-performance photonic devices.

2.4.4 Applications of direct laser written waveguides

Waveguides fabricated in fused silica and borosilicate glasses exhibit low propagation losses, long-term durability and excellent thermal and mechanical stability, making them well-suited for telecommunications, sensing, and laser delivery systems. Furthermore, FLDW allows for the integration of complex waveguide geometries, such as splitters, directional couplers, interferometers, and three-dimensional waveguide crossings, which are difficult or impossible to achieve with conventional planar lithographic techniques. Owing to these benefits, FLDW has found applications in a wide range of fields, including astronomical photonics, optical communications, quantum photonics, optofluidics and sensing [59]. In this section, I will focus on three primary application areas: optical interconnects, on-chip information processing, and topological photonics.

The first application area, optical interconnects, exploits the three-dimensional fabrication flexibility of FLDW to produce low-loss and compact routing structures that bridge fibres and photonic chips. These interconnects address key packaging challenges in integrated photonics by enabling scalable and efficient fibre-to-chip coupling and multilayer routing.

Djogo et al. reported an 84-channel fused silica interposer for high-density edge coupling of multicore fibres to a silicon photonic chip, combining FLDW with chemical etching techniques to enable scalable optical input/output architectures [60]. Similarly, Desmet et al. demonstrated a fused silica interposer fabricated via FLDW that integrates both waveguides and alignment V-grooves in a monolithic platform, achieving propagation losses of 0.88 dB/cm and coupling losses of 1.24 dB [61]. This design simplifies packaging by allowing direct and robust fibre-to-chip alignment, reducing assembly complexity.

Advances in fabrication strategies, including multiscan techniques, have further minimised propagation and coupling losses in laser-written interconnects. For example, recent work reported coupling losses as low as 0.2 dB/facet and propagation losses of 0.07 dB/cm in a 25-channel interferometer array, yielding overall insertion losses below 1 dB [62]. These results highlight the capability of FLDW to implement high-performance interconnect networks, where precise control of waveguide cross-section and refractive index profiles is essential to reduce scattering and mode mismatch.

Beyond planar waveguide arrays, FLDW facilitates the creation of true three-dimensional routing geometries that planar fabrication cannot achieve. For example, hybrid 3D optical interconnects have been realised with low-crosstalk waveguide crossings through optimised crossing angles in fused silica [63]. Such designs enable compact layouts with minimal footprint, particularly advantageous for densely integrated photonic circuits.

Secondly, the technology also enables the fabrication of integrated photonic circuits for on-chip information processing. Just as interconnects benefit from precise control over waveguide geometry and spatial arrangement, quantum and classical photonic devices

require complex three-dimensional architectures to manipulate light at the single-photon level, implement multiport interference, and perform logic operations.

Marshall et al. [64] demonstrated that directional couplers fabricated using FLDW can perform comparably to their lithographically produced counterparts. In their experiments, photons generated via spontaneous parametric down-conversion (SPDC) were used to create three-photon entangled states. The quantum behavior was verified by measuring coincidence detection rates at the waveguide outputs as a function of the relative delay between input photons, revealing the characteristic Hong-Ou-Mandel dip consistent with quantum theoretical predictions.

Building on these capabilities, Spagnolo et al. [65] fabricated a three-port device known as a tritter, in which three waveguides interact through overlapping evanescent fields. This configuration enabled three-photon interference via the bosonic coalescence effect, where all events with two photons exiting the same port were suppressed. By injecting photons into separate input ports, genuine quantum interference was observed without decomposing the process into cascaded two-mode interactions, highlighting the versatility of FLDW for fabricating complex three-dimensional quantum circuits.

Further advancing integrated quantum photonics, Atzeni et al. [66] demonstrated the combination of an entangled photon source and directional couplers on a single chip. Their approach employed a balanced directional coupler to evenly split the pump beam between two periodically poled lithium niobate waveguides, generating photon pairs through Type-0 SPDC. Interchangeable modules allowed the recombination of photons to prepare different entangled states, including configurations with a balanced coupler alone or with half-wave plates at $\pm 22.5^\circ$ in each arm. Verification via interference fringes across multiple polarisation bases confirmed the prepared quantum states. This modular, integrated approach overcomes the complexity of bulk optical setups and paves the way for scalable, flexible on-chip quantum information processing.

Collectively, these studies demonstrate that FLDW provides a versatile platform for

high-fidelity, three-dimensional quantum photonic devices, enabling scalable architectures for both classical and quantum on-chip information processing.

Finally, beyond interconnects and on-chip information processing, FLDW has emerged as a powerful tool for investigating and implementing topological photonic structures. Drawing parallels from the propagation of light in photonic lattices under the paraxial approximation and the Schrödinger equation, the evolution of light in a photonic lattice is similar to the evolution of a wave function in a quantum system, a coupled waveguide array is originally envisaged a powerful tool to study quantum phenomenon that is difficult to observe through traditional methods. Nevertheless, topological photonics has established its own identity, exploiting robust, symmetry-protected light transport that is immune to defects and disorder, offering potential advantages for photonic circuits in harsh environments or at high integration densities. The three-dimensional fabrication freedom of FLDW enables the realisation of complex lattice geometries, helical waveguides and other topological structures that are difficult or impossible to achieve with planar fabrication techniques. An extensive review of the research area by Yan et al. is given in [67].

For instance, Rechtsman et al. [68] demonstrated that arrays of helical waveguides in fused silica, fabricated via FLDW, can emulate the physics of Floquet topological insulators. Light injected into these waveguide lattices propagates along the edges without backscattering, even in the presence of defects. These lattices exhibits topologically protected edge modes and defect-immune transport, highlighting the potential of this approach for robust optical routing and signal processing.

Menssen et al. [69] introduced a vortex distortion to a hexagonal lattice and observed a mode associated with the topological defect. The experiments demonstrated a stationary bulk zero mode. A vortex distortion was introduced at the centre of the lattice. The mode was excited using beams that were individually phase and amplitude tuned . They observed that the light in zero mode was tightly confined to the centre as predicted from

theory. Adaptive optics was utilised in this example to correct for the depth-dependent aberrations and ensure all the waveguides fabricated at different depths were uniform. The laser written waveguides were near-identical over a 400 μm depth. They also demonstrated the adiabatic translation of the zero mode and that the zero mode was topologically protected against random errors of the lattice, which was deliberately introduced as errors in waveguide position.

The idea was also applied to quantum photonics, and there is rapid development in the new research field known as quantum topological photonics. Tambasco et al. proposed a quantum topological interferometer and beam splitter fabricated by FLDW, demonstrating quantum interference of topological states of light by bringing boundary states into close proximity and interfere [70]. More recently, Ehrhardt et al. for the first time demonstrated topological Hong-Ou-Mandel interference equipped with intrinsic topological protection against imperfections. By tailoring the birefringence of the waveguides within the coupler, they demonstrated that the coincidences of indistinguishable photons are fully suppressed regardless of the length of the directional coupler, showing that the topological properties are propagation-invariant [71].

In summary, by providing access to fully three-dimensional architectures, precise control over refractive index modification and integration with conventional photonic components, FLDW is uniquely positioned to advance the field of topological photonics. These developments hold promise for next-generation photonic circuits that leverage topological protection for low-loss, high-density, and defect-tolerant optical information processing.

2.5 Conclusion

This chapter covers the background for fabricating optical devices using femtosecond ultra-short pulse lasers for FLDW. Laser writing has emerged as a highly versatile and powerful technology for material modification, offering unique advantages over conventional fabrication methods. Its ability to deliver tightly focused energy with femtosecond

to picosecond precision enables localised processing across a wide range of materials. The technology benefits from its high degree of tunability, where control over parameters such as wavelength, pulse energy, repetition rate and beam shaping allows the process to be tailored to specific material responses. Furthermore, recent advances in adaptive optics and high-power ultrafast laser sources continue to expand the scope of laser writing, enabling higher throughput, improved resolution, and greater flexibility in structuring complex geometries.

Overall, laser writing represents not only a mature research tool for exploring light-matter interactions but also a transformative platform for industrial-scale manufacturing. Its capability to combine precision, scalability and versatility positions it as a cornerstone technology in fields ranging from integrated photonics and quantum technologies to integration with biomedical engineering and microelectronics. As laser sources, optical components and computational control continue to advance, laser writing is expected to play an increasingly central role in next-generation fabrication, bridging the gap between laboratory-scale demonstrations and practical applications.

The work in subsequent chapters in this thesis contributes to this area by investigating new fabrication methods and characterising device behaviour. This includes the exploration of alternative writing strategies, the use of beam control techniques to refine the modification process, and systematic evaluation of optical properties such as propagation loss, mode structure, and polarisation dependence. Through this combined improved approaches of fabrication and characterisation, the thesis provides a detailed study of the factors governing device quality and functionality, offering practical insights for the development of laser-written photonic components.

CHAPTER 3

Experimental methods

This chapter describes the optical systems used for fabrication and characterisation of waveguides, and experimental methods used throughout this thesis.

3.1 Waveguide fabrication

A femtosecond laser writing system was employed for the experiments described in this thesis. This section first introduces the system components and then explains how it is calibrated and aligned. The following subsections then describe the fabrication procedure which can be roughly divided into pre-processing, fabrication, and post-processing steps.

3.1.1 Fabrication setup

The laser used in waveguide writing was Light Conversion PHAROS SP-06-1000-pp. The PHAROS laser system was a high-repetition-rate femtosecond laser based on the chirped pulse amplification (CPA) technique, which employed Yb:KGW as the active medium. A schematic of the system and its auxiliary components is shown in Fig. 3.1. The main laser compartment (green dashed box) contains the Oscillator (OSC), Regenerative Amplifier (RA), Stretcher-Compressor (S-C), and Pulse Picker (PP). The oscillator generates seed pulses at a high repetition rate, which are temporally stretched in the S-C unit to reduce peak intensity before being amplified in the RA. The amplified pulses are then recompressed to their original femtosecond duration within the same S-C unit. The Pulse Picker, based on an electro-optic modulator, selects individual pulses from the oscillator pulse train, thereby controlling output repetition rate of the emitted beam.

Peripheral subsystems provided essential operational support. The Power Supply Unit (PS) delivered electrical current to the laser diodes that pump the active medium and power other electronic components. The chiller circulated coolant to maintain thermal stability, while the Timing Electronics Module (TEM) automatically synchronised the operation of all active components, including the oscillator, amplifier, and pulse picker.

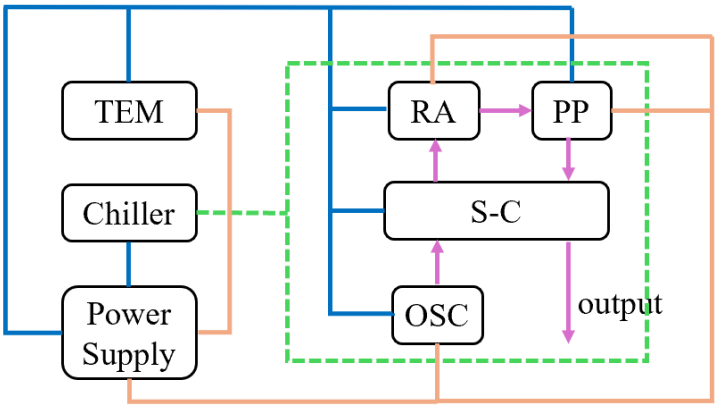

Figure 3.1: Schematic diagram of the PHAROS femtosecond laser system and auxiliary components. The main laser compartment (green dashed box) contains the Oscillator (OSC), Regenerative Amplifier (RA), Stretcher–Compressor (S–C), and Pulse Picker (PP). The Timing Electronics Module (TEM), Chiller, and Power Supply provide synchronisation, cooling, and electrical power respectively. Blue lines indicate electrical power connections, orange lines indicate timing synchronisation connections, and purple arrows represent the optical beam path. Adapted from [72].

The PHAROS had a pulse duration of t = 168 fs and a maximum average power of up to 6 W. The fundamental wavelength is at 1030 nm, and the second harmonic which was frequency-doubled to 515 nm was used for fabrication. The shorter wavelength created smaller fabrication feature sizes which enabled better mode matching with fibre for optimal coupling. The repetition frequency was set to 1 MHz in this work for fabrication in the heat accumulation regime. A combination of a motorised rotating half-wave plate and a polarisation beam splitter was used to adjust the average laser power which selectively filtered out a proportion of the incoming beam depending on the waveplate rotation angle. The pulse picker installed in the laser was an electrically controlled optical switch, which extracts pulses from a pulse train and controls the output of the laser.

Fig. 3.2 shows a photo of the system and Fig. 3.3 is a simplified schematic. Lenses L1 and L2 expanded the beam onto a liquid-crystal on silicon spatial light modulator (SLM) (Hamamatsu Photonics X10468-09(X)) to correct for aberrations. Since the SLM operated in reflection, mirrors M2 and M3 ensured that the angle of incidence was small for optimal performance. L3 and L4 formed a 4-f system which imaged the SLM onto the pupil plane of the objective (0.5 NA; 20x; Zeiss Plan Neofluar), which then focused the laser into the specimen placed on the sample stage.

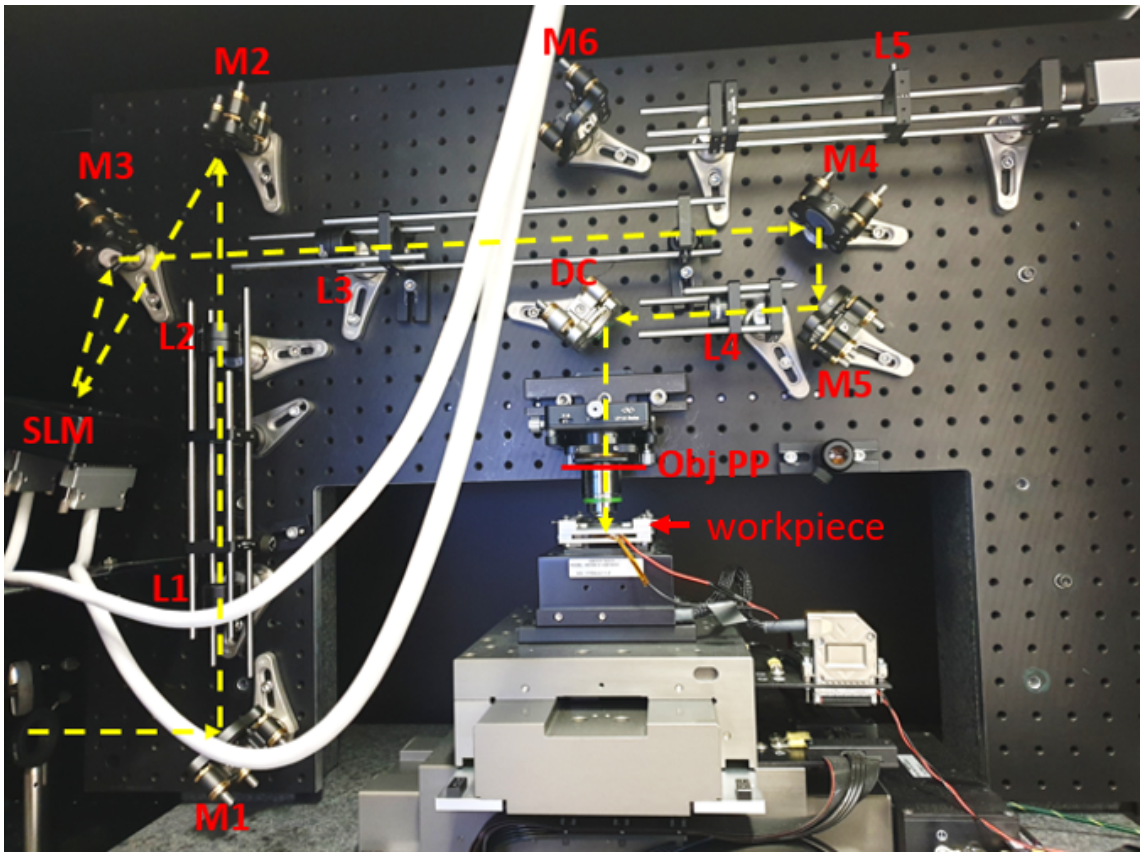

Figure 3.2: Photo of fabrication system. The laser was located at the bottom left of the photo. M: Mirror; L: Lens; SLM: Spatial Light Modulator; DC: Dichroic Mirror; Obj PP labels the pupil plane of the objective. The CCD camera for imaging the sample was located at the top right.

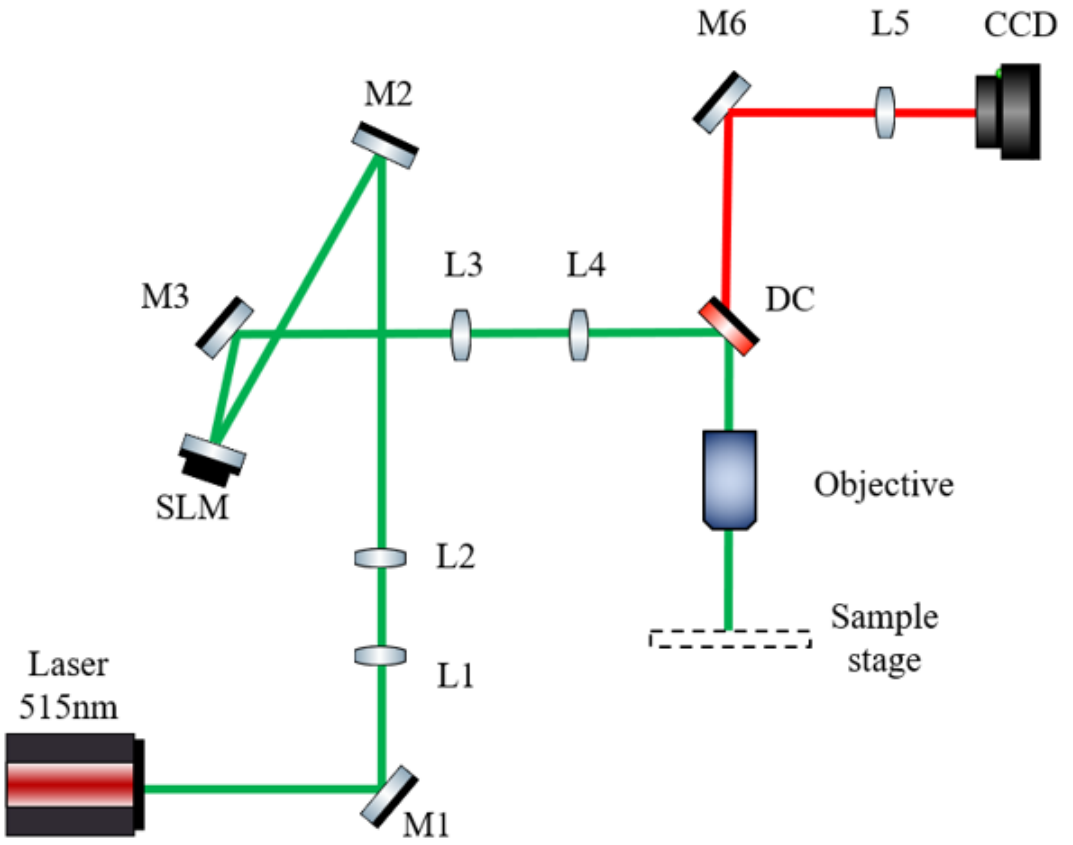

Figure 3.3: Simplified schematic of the fabrication system. L1 and L2 expanded the beam onto the SLM; L3 and L4 formed a 4-f imaging system that preserved the modulated wavefront delivered to the sample. The reflection from the sample passed through the Dichroic Mirror (DC) and was captured by the CCD camera.

The three-axis air bearing translation stage (AerotechABL10100L, xy-motion; ANT95-3-V, z-motion) was controlled with Aerotech G-code. The stage was back-illuminated by a red LED which allowed imaging of the sample with a CCD camera (Baumer VLU-12M) with the aid of a Dichroic mirror (Thorlabs DMLP550). The focal plane of the camera was adjusted such that it coincided with the laser focus to allow for easier alignment and observation of the fabrication process. A motorised rotating half-wave plate was placed at the output of the laser to control the energy delivered to the sample. These were all connected to an external an external desktop computer and controlled by computer software. The SLM was controlled by a LabView program, and the laser, Aerotech stage, CCD camera and waveplate were controlled by dedicated software packages that came

with the products.

3.1.2 System alignment and calibration

It was important to calibrate the system at regular intervals, as the alignment of components could gradually drift over a few weeks or months, due to environmental changes such as temperature changes or mechanical vibration. The critical part in the system was the alignment of the 4-f system and ensuring the pattern displayed on the SLM filled the back focal plane of the objective.

To achieve this, L5 was adjusted such that the image plane of the camera coincided with the back focal plane of the objective by observing when the sharp edge of the objective aperture comes into focus. Once the sharp edge was imaged, L4 was adjusted to bring the SLM pattern into focus. Typically, a simple pattern such as a Zernike mode Z5 is displayed which was a simple cross-shape to identify the centre of the pattern. The SLM pattern was then centred on the back focal plane by a combination of adjusting the physical position of the objective using the X/Y positioners of the objective holder, and software which controlled the offset on the SLM display. The size of the pattern can then be adjusted until it completely fills the objective aperture. After completing these steps, L5 was then reset to the original position where it coincided with the objective focal plane instead to image the sample.

The laser power was another parameter that could change over time. To calibrate for the laser power, the objective was removed, and a power meter was placed on the stage which was equivalent to measuring the power at the back focal plane objective. By changing the angle of the half-wave plate for power control, the correspondence between the waveplate angle and power was recorded. Unless otherwise specified, the value of fabrication power mentioned throughout the thesis were determined from this measurement.

To find out and calibrate the optimal parameters for waveguide fabrication, at regular time intervals (e.g. once every month) multiple straight waveguides were fabricated while

varying the fabrication parameters for each and inspected for the mode profile on the characterisation setup. Optimisation of fabrication was necessary since the material modification process was highly sensitive to fabrication parameters, which greatly impacted the guiding performance of waveguides. In Eagle glass, the optimal fabricating parameters were found to be an average power of P_{avg} = 110 mW after the objective, and fabrication speed of 6 mm/s with the appropriate aberration correction applied. In fused silica, the optical average power was around 50 mW and optimal writing speed at 4 mm/s.

3.1.3 Sample Preparation

The glasses used in the experiments were either Corning Eagle glass or UV-grade Quartz, commonly referred to as fused silica. The Eagle glass samples were available as 20×40 mm substrates, and fused silica samples were available as 10×20 mm. To facilitate post-processing and accommodate for the size of the sample holders, the samples were usually cut into smaller pieces with length <20 mm using a mechanical rotating blade glass cutter.

Before mounting the sample, the surface was thoroughly cleaned with acetone or isopropyl alcohol to remove any debris or stains which could potentially obstruct light and distort the laser focus.

3.1.4 Fabrication Procedure

After the sample was mounted, it was usually useful to identify a corner as a reference point for the coordinates and positions of fabricated structures. The tilt of the sample could be inferred from the position feedback from the stage, which was then corrected accordingly during fabrication.

The shape of the fabrication beam was calibrated by observing the reflection of the focus cross-section on the sample surface, and any phase aberrations were corrected by adjusting the coefficients of Zernike modes displayed on the SLM using LabView software. After obtaining the desired focus, the stage was translated to the starting position for waveguide writing. The depth of fabrication was also inputted into the

software which automatically calculated the corresponding depth-dependent aberration to ensure consistent fabrication at different depths.

A Thorlabs waveplate controller was used to control the rotation of a half-wave plate which effectively acted as a power controller. A polarising beam-splitter was placed after the rotation mount, which separated the incoming beam into vertical and horizontal components, and only the vertical component passed straight through. The splitted beam is monitored by a power meter which allowed checking whether the power was set to the correct value.

Aerotech Basic script was used to control the stage movement in the Aerotech Controller. The stage movement was either generated using MATLAB code or typed in manually. The script could be synchronised with the laser pulse-picker which controlled when the laser output turned on and off, and a high degree of control over the fabrication could therefore be obtained by a combination of stage movement and power switching.

Depending on the experimental design, multiple scripts might be required and the fabrication parameters such as the laser writing power might require adjustment.

3.1.5 Post-processing

After the fabrication was completed, both ends of the sample was polished using a polishing machine (Krelltech FLex Waveguide Polisher) since there were aberrations at the edges of the sample which degraded the quality of waveguide at both ends (see Chapter 4). Polishing pads of 30 μm, 9 μm, 3 μm and 1 μm were used consecutively for around 5-10 minutes each to ensure best results. The end facets of the sample were then checked on the waveguide characterisation setup under LED illumination to determine whether the surfaces were of good quality and ready to be characterised, otherwise further polishing might be required.

3.2 Waveguide characterisation

3.2.1 Waveguide measurement setup

To characterise and measure the performance of the fabricated devices, a testing rig was used to couple light from fibre into the waveguides. A fibre-coupled laser source (Thorlabs S1FC780PM) was used for the characterisation purpose. The laser was coupled to a single mode fibre (Thorlabs PM630-HP (PANDA)) and the light emitted had a wavelength of 785 nm and a maximum output power of up to 6.25 mW. The other end of the fibre was mounted onto a fibre chuck (Thorlabs HFR001) on a 6-axis stage (Thorlabs MAX600/M Series) for coupling into the waveguide sample. The waveguide sample was mounted onto a 3-axis translation stage (Newport M-562-XYZ-LH). The other end of the waveguides was imaged onto a CCD camera (Baumer TXD-14) using an objective (Olympus ULWD MS Plan 80x/0.75NA) and an achromatic lens (Thorlabs AC254-100-A-ML). A photo and a simplified schematic of the setup were shown in Fig. 3.4 and Fig. 3.5 respectively.

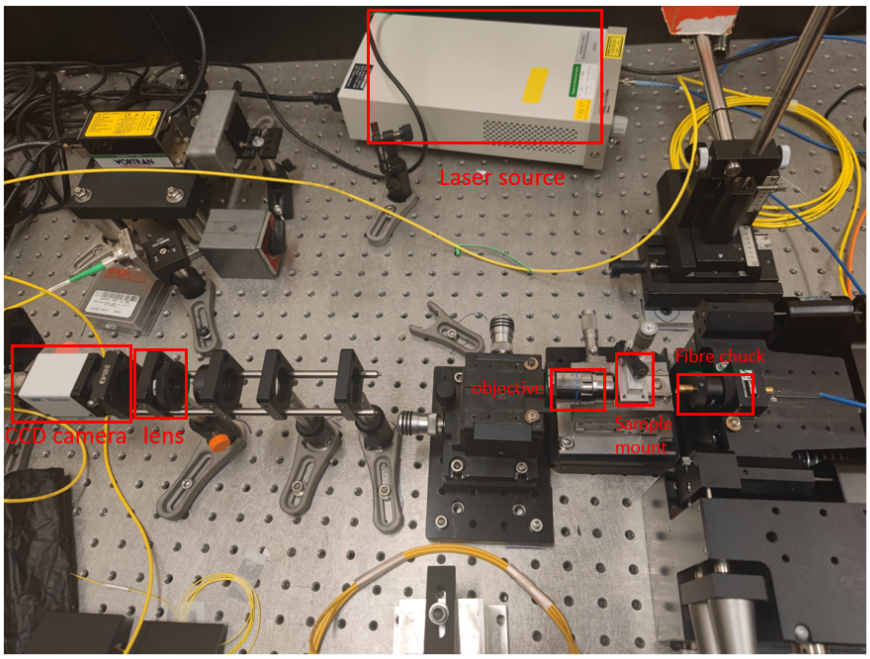

Figure 3.4: Photo of the characterisation setup. The camera was connected to a desktop computer to capture images and adjust camera settings.

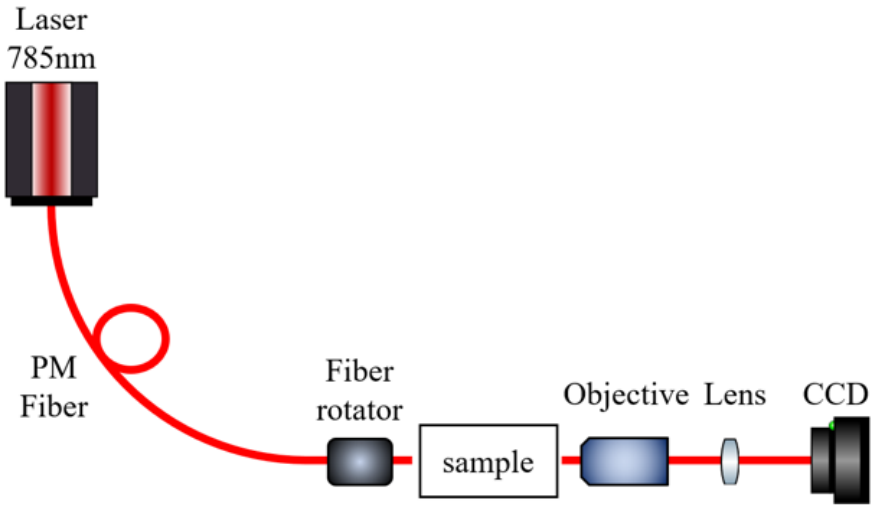

Figure 3.5: Simplified schematic of the characterisation setup.

3.2.2 Characterisation Procedure

To perform a characterisation routine, the waveguide sample was first cleaned and mounted on a 3-axis stage. The sample was then brought into focus of the imaging arm by translating along the optical axis with the assistance of LED illumination while also ensuring the end facet was of good quality and free from any debris which might scatter or block light.

The LED was then swapped for a fibre holder mounted on a 6-axis stage. The fibre was first brought as close as possible to the sample manually, then using translation knobs to ensure the laser was coupled into the correct waveguide by observing the output intensity change on the camera. Finally, the fine adjustment knobs were used to ensure maximal coupling.

Two types of measurements were taken: images and power readings. For images it was essential to ensure the images are not over-saturated by adjusting the camera shutter speed, and using the built-in capture function from the camera software. For power readings, a power meter was inserted in place of the camera, where the power meter should first be zeroed with the laser turn off. Since the fibre was extended from the fibre holder, air disturbances could cause fluctuation in the reading, so it was preferable to wait for the

reading to first stabilise and take a time-average of the data.

3.2.3 Data processing

After obtaining the images of the waveguide modes, the images were processed in MATLAB. Since the waveguide modes were approximately Gaussian, the centres of each waveguide mode were identified as the maximum value. A Gaussian profile was then fitted to the waveguide modes, and the dimensions and symmetry of the waveguide modes were then obtained from the Gaussian fit. If the data were in the form of power measurements, they were also inputted into MATLAB for post-processing and data visualisation.

3.2.4 High dynamic range measurement

On the current characterisation setup, we only had access to an 8-bit CCD camera, which did not have a high enough fidelity when High Dynamic Range (HDR) measurement was required. This was particularly important for measurements where crosstalk was a major concern, such as in quantum applications. Crosstalk was caused by coupling between waveguides and must be kept to a minimum to enable applications such as precise individual qubit addressing. Any unwanted noise from other waveguides in the system might completely destroy encoded information in the photons, therefore a method to characterise waveguides with high dynamic range was necessary to measure and quantify this undesirable effect.

To tackle this problem, a protocol using Neutral Density (ND) filters was devised to implement a HDR measurement on the characterisation setup using the same camera. Since ND filters did not have a flat frequency response, it was crucial to calibrate the ND filters at the operating wavelength in order to obtain accurate ratings of the filters. This was done by replacing the CCD camera with a power meter at the image plane of the system and measuring the power reading through each ND filter while keeping the input power from the laser source constant. From the measurements, the optical density (OD)

of each ND filter was then given by:

$$OD = -\log \frac{P}{P_0} \tag{3.1}$$

where P was the power after the filter and P_0 was the incident power.

After calibrating the ND filters, the procedure for performing high dynamic range measurement was as follows:

1. Couple the fibre to waveguide input as usual, and place ND filters in position.
2. Turn up the input power such that the centre of waveguide output is saturated for all images taken through ND filters except the last one (with highest OD).
3. Save the images and run them through the MATLAB code.

Fig. 3.6 shows an example of a series of images I_N. By passing these images through the algorithm, an HDR measurement was obtained.

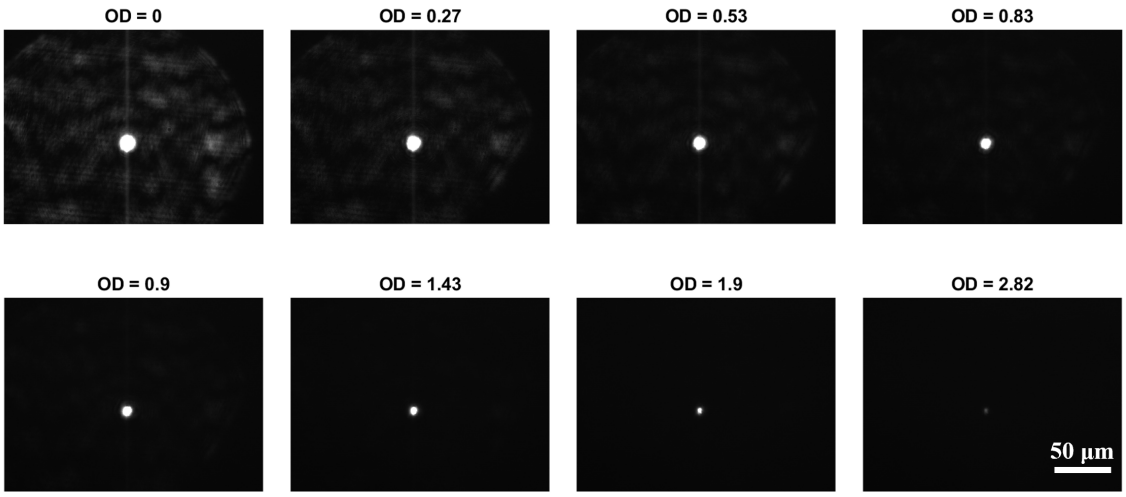

Figure 3.6: Images taken with successive ND filters. Top left image is taken without a ND filter, and bottom right image with the highest optical density.

The most attenuated image was set to have a peak intensity below saturation (255 for 8-bit image). The other images had different degrees of saturation. We refer to the sequence of images as $I_N(x, y)$ for n=0 to N and the attenuation of each of the ND

filters as R_n, for n=1 to N, so that for a non-saturated pixel, we should have $I_N(x_1, y_1) = R_N I_0(x_1, y_1)$. The images were then put through a MATLAB script which performed the following routine:

1. We start with the last (most saturated) image $I_N(x, y)$.
2. Create a temporary image $T(x, y)$ of the same size and set all pixels to equal $I_N(x, y)$.
3. Take the next image $I_{N+1}(x, y)$. For all saturated pixels in $T(x, y)$, replace the pixel with the corresponding pixel value from $I_{N+1}(x, y)$ multiplied by R_{N+1}/R_N. Note that some of these replaced pixels will also be saturated.
4. Continue through all images: For all saturated pixels in $T(x, y)$, replace the pixel with the corresponding pixel value from the next image as $I_{N+1}(x, y)R_{N+1}/R_N$
5. The output is then equal to $T(x, y)$.

Fig. 3.7(a) plots the pixel values of a row passing through the centre of the waveguide output for the most saturated image $I_8(x, y)$. Fig. 3.7(b) plots the same row of the combined image $T(x, y)$ for comparison, and Fig. 3.7(c) plots the log-scaled combined image for better visualisation of the results. Fig. 3.8 shows a 3D plot of the log-scaled total combined image.

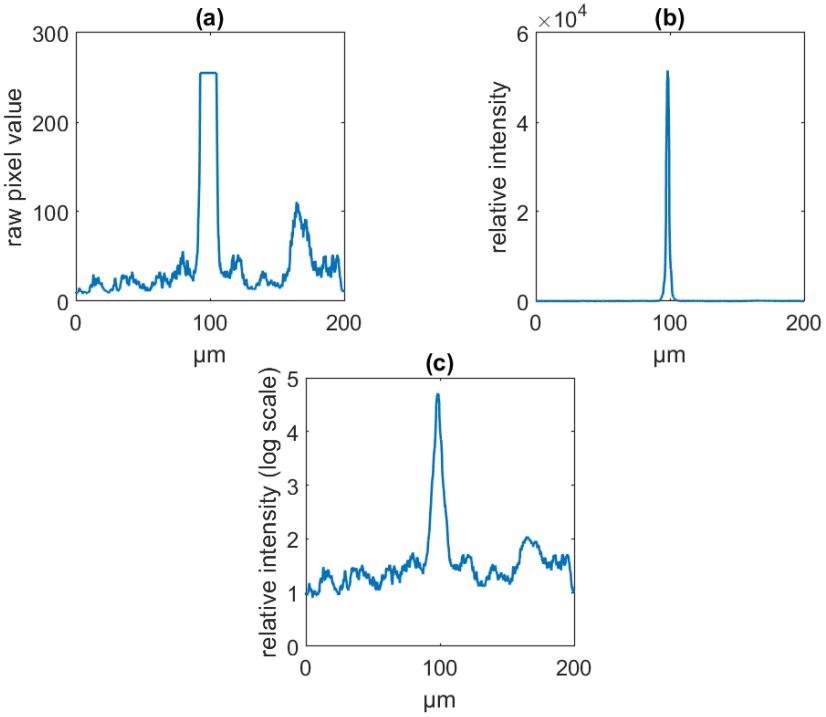

Figure 3.7: (a) Raw pixel values of a saturated image. (b) Relative intensity obtained after processing HDR measurement data. (c) Log-scale plot of relative intensity obtained from HDR measurement.

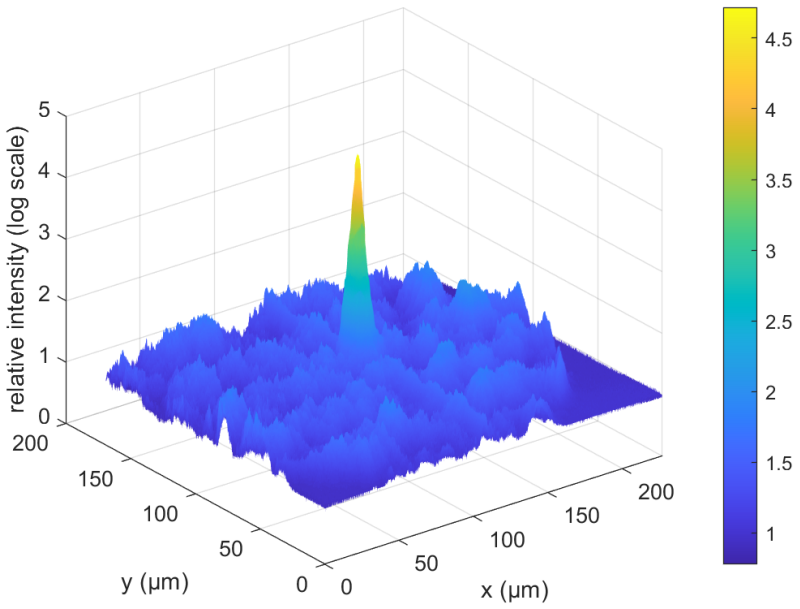

Figure 3.8: 3D plot of relative intensity obtained from HDR measurement on log scale.

The range of the HDR measurement depends on the maximum OD of the filters. With an OD of 3, the theoretical minimum power that can be measured is 10^{-3}.

In summary, this chapter has detailed the experimental setups for the fabrication and characterisation of femtosecond-laser-written waveguides. The integrated PHAROS laser system provided a stable and tuneable source of ultrashort pulses for high-precision material modification, while adaptive optics ensured accurate aberration correction and beam shaping. The alignment and calibration routines maintained fabrication consistency over extended experimental periods. Waveguide performance was subsequently characterised using a dedicated optical testing rig, supported by custom data processing and high dynamic range imaging techniques to enable precise quantitative analysis. Together, these methods established a robust experimental platform for investigating and optimising FLDW waveguide structures, providing the foundation for the results presented in the following chapters.

CHAPTER 4

Correcting for edge aberrations

4.1 Introduction

This chapter investigates a new method that can maintain comparable performance to polished waveguides without the need of polishing. FLDW excels at creating waveguides for integrated photonic applications. A major scenario would be to create optical interconnects that connect different parts of an integrated photonic circuit, where the different components including sources, processors, or detectors may all have different configurations, where laser-written waveguides with the appropriate geometry can serve as connectors. The laser-written devices themselves may also perform data computing routines as shown by various examples in Chapter 2. Regardless of the application, the ability of laser-written samples to interface with other devices is very important [73].

In the usual fabrication routine, the edge of the sample is polished back to expose the waveguides before aligning with other devices to ensure the full waveguide structure is exposed for external coupling. This hinders the possibilities of creating efficient laser-written devices and is a major obstacle to scaling up laser-writing in pre-assembled configurations. Mechanical polishing is also a source of error since incomplete polishing or a tilted facet affects the final quality of the sample, and there is also the possibility of damaging the facet which incurs more time resources to repolish or may even require the sample to be remade.

If there is a way to eliminate the need for the polishing step, this could open many new

avenues for integration with other optical devices, at the same time reducing the need for polishing which requires substantial processing time and effort. It also removes the need for active alignment between the fabricated sample and other components, which is yet another time-consuming and error-prone step. Fig. 4.1 shows a schematic showing how in-situ writing may look like if the new technique described in this chapter is applied, where the glass sample block to be fabricated is first mounted between two other components, after which the waveguides are written. The correct positioning of fabrication can be accurately determined by imaging techniques from the surface of the sample, thus avoiding the need for active alignment. This is in contrast to the standard approach, where the waveguides are written first and then the sample is aligned with other components.

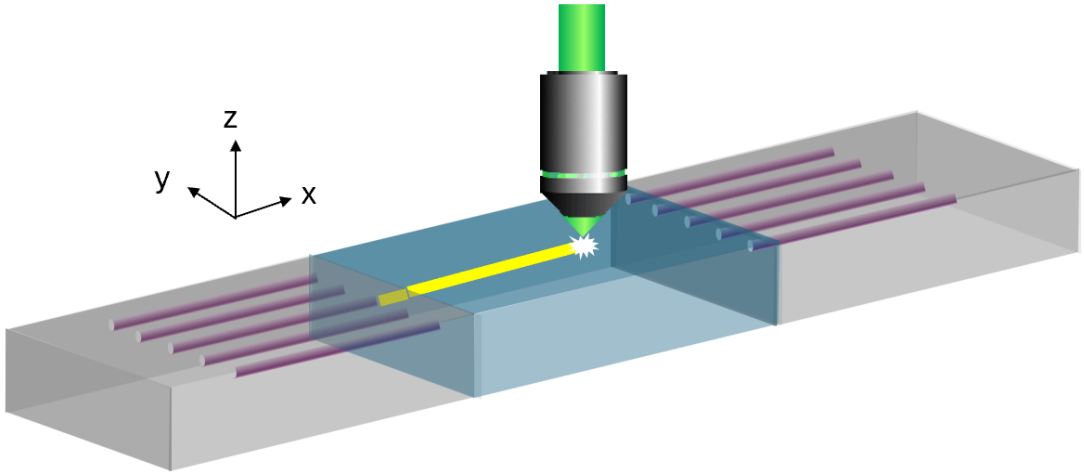

Figure 4.1: Schematic of in-situ fabrication of waveguides where the glass block is sandwiched between other components.

This chapter is structured as follows. The mechanism of edge aberrations will first be explained, and the proposed method for overcoming these challenges will be introduced. The experimental result will then be presented and followed by detailed discussion. Finally, a summary of results will be given, and potential applications and follow-up work will be discussed in section 4.5 “Conclusion and outlook for future work”.

4.2 Edge aberrations

Edge aberrations occur when the focusing cone of the laser extends beyond the sample edge such that a portion of the light rays passes through the side of the sample. Fig. 4.2 (a) shows a simplified schematic of how the focusing cone extends beyond the edge and a portion of it is incident on the side of the sample. This leads to a focal splitting as shown by the ray diagram in Fig. 4.2 (b), thus effectively decreasing the energy delivered to the focus and impairing the quality of fabrication. The polishing step essentially polishes away this section of the sample where the edge aberrations occur, but this is a wasteful, error-prone and time-consuming step as explained earlier.

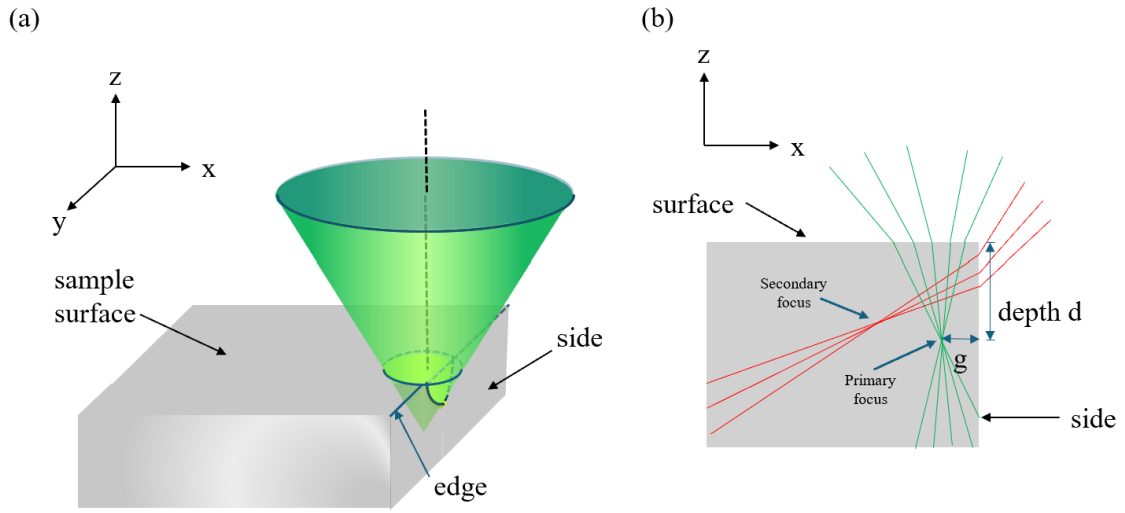

Figure 4.2: (a) Schematic showing the focusing cone of the fabrication laser beam intersecting with the sample edge. (b) Ray diagram showing focal splitting that occurs at sample edge.

An analytic expression for the phase aberration of the portion of the beam passing through the side facet is given as follows (the portion which is still incident on the top surface has the same phase aberration as earlier expressions given in Chapter 2) [74]:

$$\Psi(\rho,\theta) = \frac{2\pi}{\lambda} g n_1 \left(\sqrt{\frac{n_2}{n_1} + \rho^2 \sin\alpha_1^2 \cos^2\theta - 1} - \rho \sin\alpha_1 \cos\theta \right) \tag{4.1}$$

where g is the distance from the centre of the focusing cone on the sample surface to the sample edge, n_1 and n_2 are the refractive indices of the surrounding medium and

sample respectively, and α_1 is the maximum angle of the focusing cone that relates to the numerical aperture of the objective lens NA which is given by $NA = n_1 \sin \alpha_1$.

It has been proposed that the edge aberration can be fully corrected by displaying the corresponding phase pattern for different gaps g from the sample edge at different depths d as the writing approaches the sample edge [74]. However, this method is limited by the refresh rate of the spatial light modulator (SLM) which is typically of the order of ≈100 Hz, meaning that the translation speed has to be on the order of a few μm/s for the phase correction to be appropriately applied. This is impractical for industrial applications and hinders the possibility for scaling up to more complex device designs.

The effect of edge aberrations on the focal intensity can be quantitatively evaluated by calculating the normalised focal intensity (I/I_0) as a function of the normalised distance from the sample edge (g/d). The simulation results are shown in Fig. 4.3. As the focus approaches the sample edge ($g/d \to 0$), the focal intensity decreases due to partial truncation of the focusing cone by the sample boundary, which causes a portion of the incident beam to refract or scatter out of the focal region. This effect is more pronounced for higher numerical apertures, where the focusing cone subtends a larger angle and a greater fraction of rays interact with the edge. As g/d increases, the entire focusing cone is eventually contained within the sample, and the normalised intensity approaches unity, corresponding to the undistorted focal condition. The inset shows an example plot of the objective lens pupil plane, where the yellow region represents the portion of the beam contribute rays that are incident on the top surface of the sample and included in the focal intensity calculation.

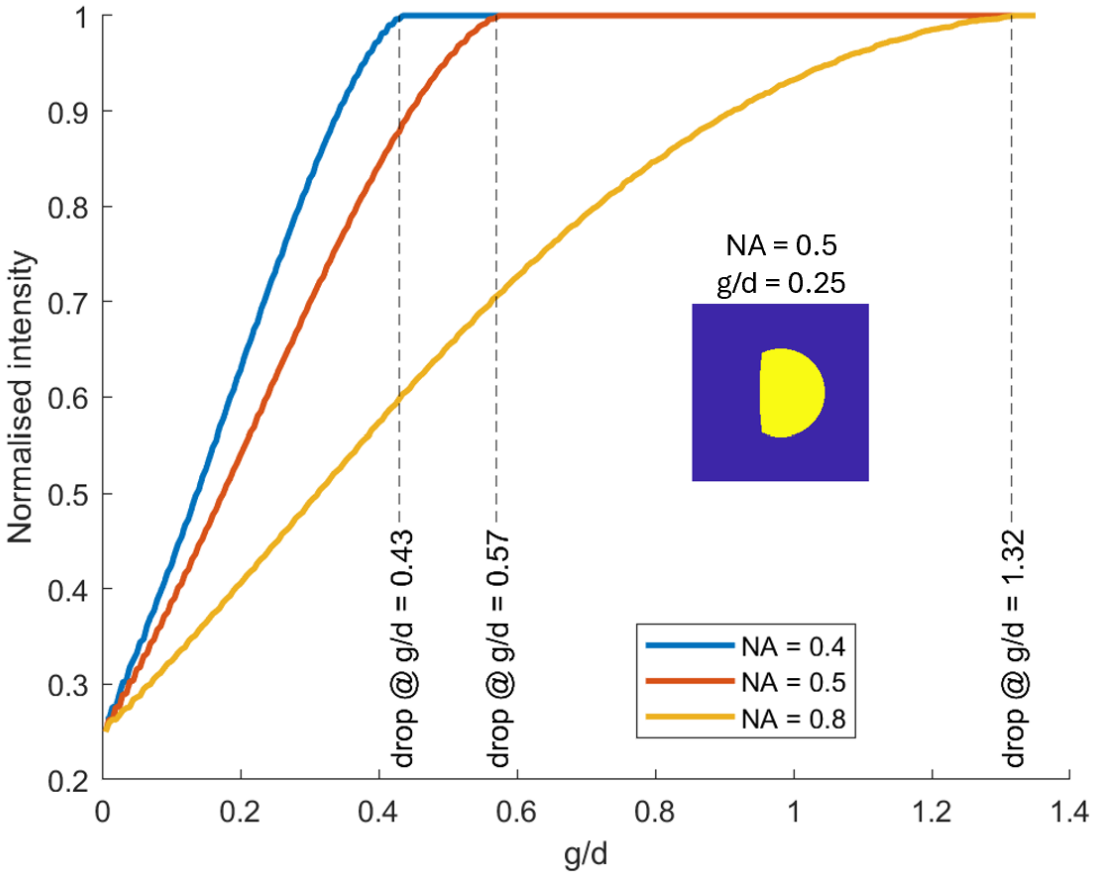

Figure 4.3: Simulated focal intensity as a function of the normalised distance from the sample edge (g/d) for different numerical apertures (NAs). The normalised focal intensity (I/I_0) decreases as the laser focus approaches the sample edge due to partial truncation of the focusing cone by the sample boundary at different NA. The dashed lines indicate where the intensities start dropping off, meaning the cone begins to intersect with the edge and edge aberrations take effect. Higher-NA systems exhibit stronger sensitivity to edge proximity, with a more gradual recovery of focal intensity as g/d increases. The inset shows an example plot of the objective lens pupil plane at NA = 0.5 and $g/d = 0.25$, where the yellow region indicates the portion of the beam contributing rays incident on the top surface of the sample and included in the focal intensity calculation.

4.3 Method for writing up to the sample edge

As discussed above, correcting exactly for the phase aberrations is impractical, and ideally a new scheme should not affect the fabrication speed by a significant amount and should be easy to implement while achieving a comparable result. To resolve this issue, a new method is proposed which will be explored in the following subsections.

4.3.1 Half pupil method

To begin with, we start with the simple idea that we investigate what happens when we block out half of the pupil at the fabrication objective pupil. This means there is only half of the fabrication beam going through the objective, which guarantees that the focusing cone of the laser stays within the boundaries on top of the sample surface up to the sample edge and avoids any possible edge aberrations. This will be referred to as the "half mask" or "half pupil" method throughout the rest of this chapter.

Fig. 4.4 shows the simulation of the laser focus calculated using the point spread function using the scalar diffraction integral of the pupil, similar to that described in Chapter 2 in relation the 4-f system [75]. The top row shows that there was a slight tilt of the focus, but the longitudinal dimension remains roughly the same. With the partial pupil amplitude blocked, the focus tilted in the direction normal to the substrate edge. The lower row shows that the focus was stretched in the x-direction, which was expected since blocking half of the pupil effectively reduces the NA of the objective by a factor of two.

This behaviour can be understood by noting that blocking half of the objective pupil effectively reduces the NA of the focusing system by approximately a factor of two. Since the NA determines the maximum convergence angle of the rays that form the focal spot, reducing it decreases the lateral confinement of the focus while leaving the axial extent largely unchanged. In other words, the range of zenith angles relative to the optical axis of the transmitted rays remains similar, so the focal depth is not strongly affected.

However, because only half of the pupil area contributes to the focus, the total optical power reaching the focal volume is reduced by a factor of two. Moreover, the intensity, which depends on both the power and the area over which it is distributed, is therefore further reduced. The lateral broadening of the focal spot (approximately twice as wide) spreads the available energy over a larger area, leading to an overall reduction in peak intensity by roughly a factor of four. To maintain the same energy density or fluence at

the focus as in the full-pupil case, the total incident power must thus be increased by approximately fourfold. This requirement is still practical, as the laser powers used for FLDW are typically on the lower end of the full power range of femtosecond lasers.

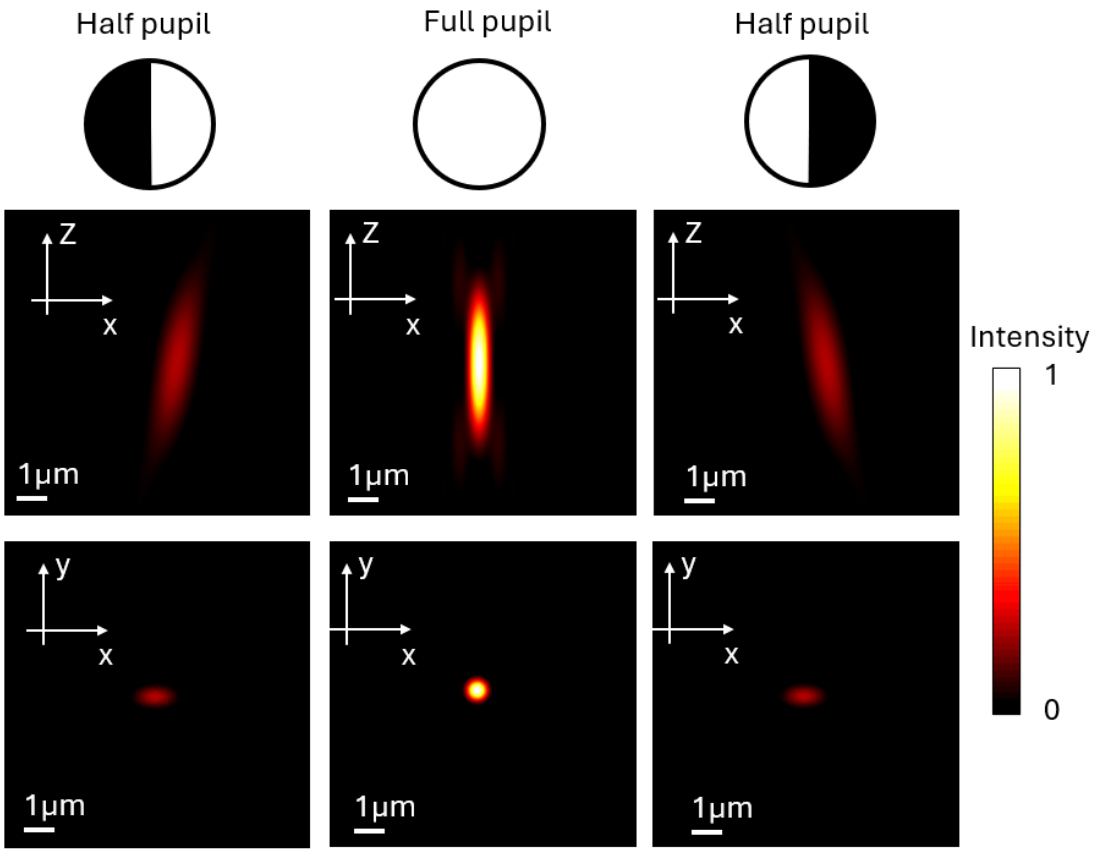

Figure 4.4: Simulation of laser focus when respective amplitude masks are applied. Upper row: x-z view from the side of the sample showing how the focus is tilted when half mask is applied. Lower row: x-y view that shows the top-down view of the cross section of the laser focus inside the glass sample.

In the actual fabrication system, the amplitude mask was applied via spatial filtering. The top row of Fig. 4.5 shows the corresponding phase pattern displayed on the actual SLM (without additional systematic and depth-dependent corrections for clarity). The grating like pattern was obtained by setting Zernike mode Z2 = 60 in order to shift the beam to its first order. By displaying the grating on half of the SLM, the other half which remains in the zeroth order beam was blocked by a pinhole inserted at the Fourier plane, thus only half of the beam reached the pupil plane of the objective.

The lower row of Fig. 4.5 shows the actual shape of the reconstructed laser focus. We observed the reflection of the laser focus on the sample surface, and by translating the sample in the z-direction longitudinally parallel to the beam direction, we effectively obtained cross-sections of the laser focus in the x-y plane. By stitching together the cross-sections of the laser focus, we could then reconstruct the shape of the laser focus. The tilt matched the simulated focus in Fig. 4.4, which suggests that the modelled focus and experimentally created focus agree with one another. Note that the relative intensity was not depicted accurately in order to avoid camera saturation since the energy is more concentrated for the full pupil case.

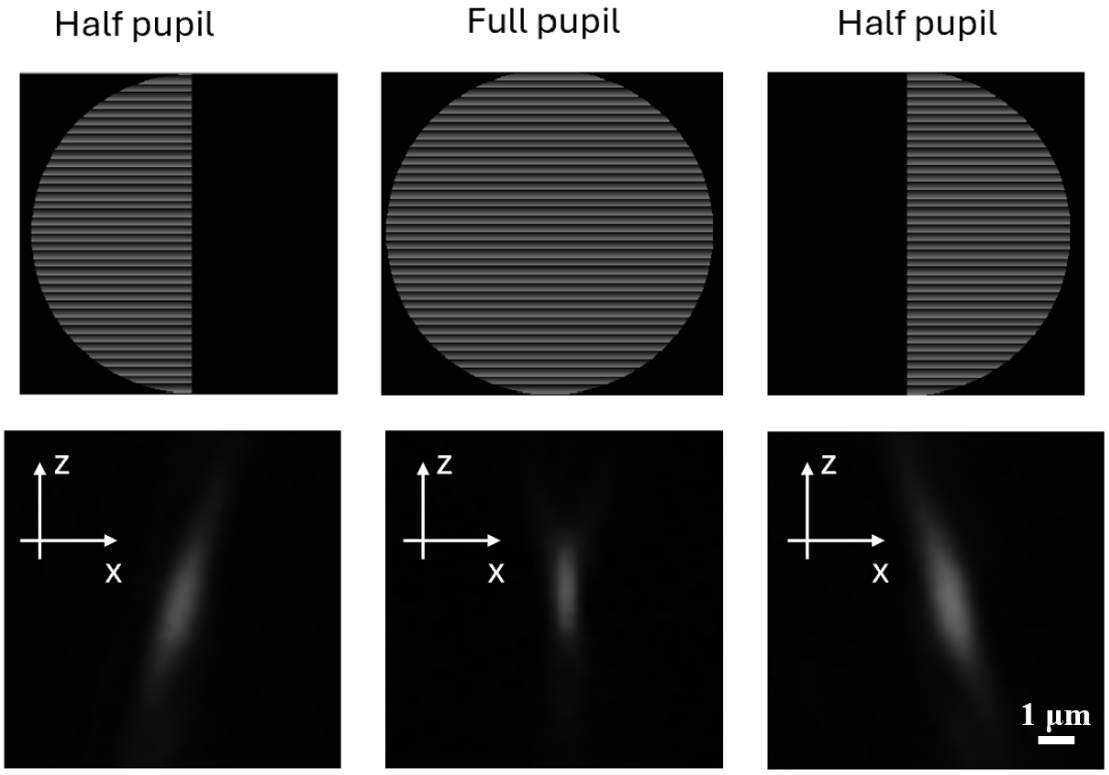

Figure 4.5: Experimental images of the respective amplitude masks and focus. Upper row: amplitude mask displayed on the SLM. Lower row: x-z images of the focus obtained by stitching together the cross-sections of laser reflection from the sample surface. Laser power has not been modified between acquisition but camera exposure times are different to avoid camera saturation.

Spatial filtering is not the only method to implement the amplitude mask, and using a physical shutter that blocks half of the beam at the objective pupil plane should achieve the same effect in principle.

4.3.2 Writing scheme

Given the above, we can now propose a writing scheme which is outlined in the schematic shown in Fig. 4.6. The waveguide is first written with the femtosecond laser using the standard fabrication method with full pupil. However, instead of writing all the way through the sample, it terminates before the focusing cone reaches the sample edge, avoiding edge aberrations. The half mask is then applied, and the laser writes the remaining section of waveguide up to the sample edge.

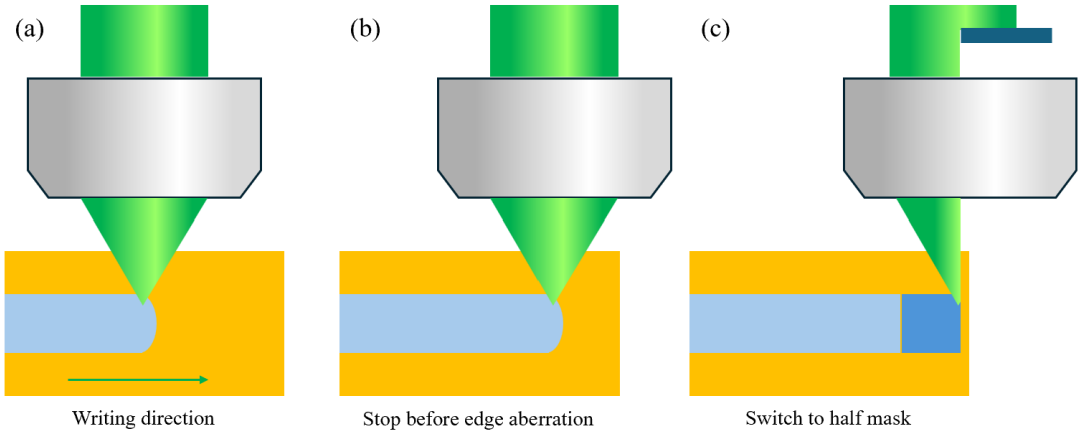

Figure 4.6: Writing scheme of the proposed method. (a) The conventional laser writing method with full pupil. (b) The writing terminates before edge aberration occurs. (c) An amplitude mask is applied and the half pupil method is used to fabricate the waveguide segment up to the sample edge.

From the geometry of the setup, a good estimation for the distance to the sample edge g is given by the relation $NA = n \sin \alpha_1$ and $\tan \alpha_1 = \frac{g}{d}$, where the outermost rays of the light cone just tangentially intersects the edge, and d is the nominal depth for fabrication (see Fig. 4.2). The objective used is a 0.5 NA objective lens and the maximal depth investigated is 750 µm actual depth corresponding to approximately 500 µm nominal depth, and $n \approx 1.5$, solving for g we obtain $g \approx 177$ µm.

The full writing scheme is therefore as follows: the waveguide is first written by the femtosecond laser using the standard method with the full pupil until 200 µm from the sample edge. The half mask is then applied, and the laser writes the remaining section

from 200 μm to 10 μm from the edge, where the laser power cuts off in order to avoid damage to the end facet caused by surface ablation.

4.3.3 Preparation of sample and fabrication implementation details

This subsection covers some implementation specifics for sample preparation and fabrication process. These are easily overlooked steps yet they may have a significant impact on the quality of fabrication and measurement.

Firstly, to write up to the edge of the sample, it is crucial that the sample surface is well formed and does not have any pre-existing damage. Fig. 4.7 (a) shows that the surface of a glass sample can be damaged if it is mechanically polished. The damaged parts on the surface will create aberrations for fabrication near the edge that usually go unnoticed since they are polished off anyway, but in this particular application it is a huge source of error. Compared to Fig. 4.7 (b) where the edge of a sample is obtained directly from the manufacturer who has specialised tools for cutting, the edge is significantly more regular which avoids the unwanted aberrations.

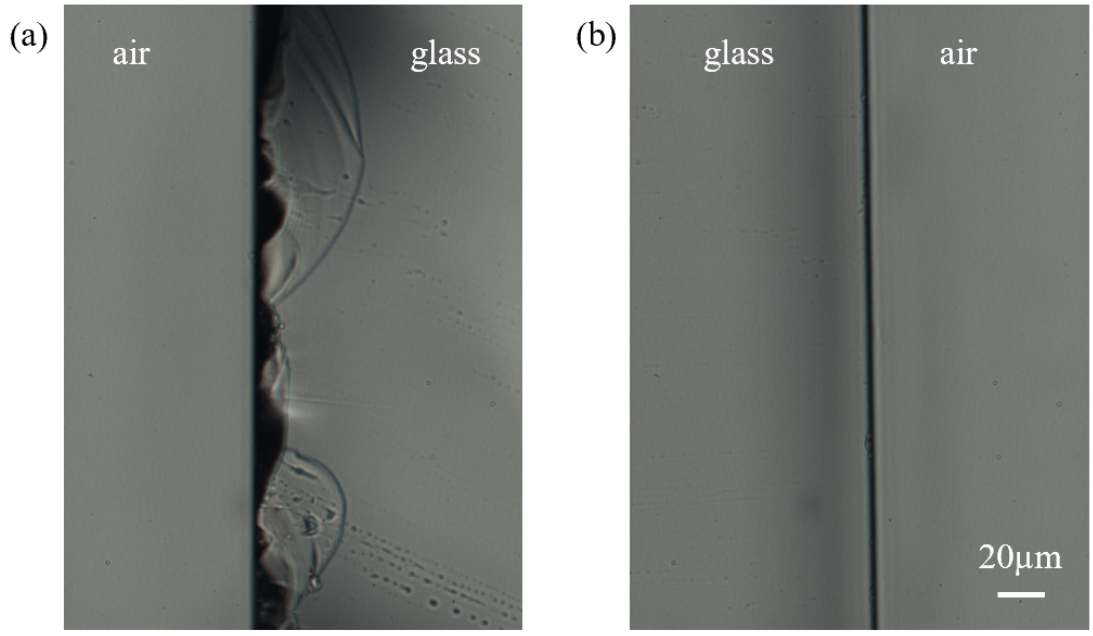

Figure 4.7: Images of the surface at sample edge under microscope (a) with damage from polishing (b) without damage.

Secondly, since the writing consists of connecting two distinct waveguide segments, the usual **LINEAR** command of the AeroBasic code used to control stage movement is independent of the pulse picker, which acts as a switch controlling whether the laser emits light. This means that the pulse picker is first turned on, then the stage moves until it reaches the end position, then the pulse picker is turned off. This means that there is a short time frame which the laser was stationary at the beginning and end of the motion where it needs to accelerate and decelerate respectively. This is shown in Fig. 4.8 (a)(i), where the motion only starts after the pulse picker is turned on. The result of this fabrication can be seen in Fig. 4.8 (b)(i), where a 'blob' of refractive index change is observed at the transition because of the energy accumulation during the time frame where the stage is stationary. This phenomenon introduces additional loss that is attributed to the irregular ending of the waveguide segment which leads to scattering and mode mismatch. To overcome this issue, the **PSOWINDOW** command is used instead. This command accepts position feedback from the stage, and synchronises the pulse picker to trigger within a specified window. This allows the stage to pick up velocity first by starting its movement earlier, then only trigger the pulse picker when it reaches the intended starting position. Similarly, this avoids the deceleration at the end of movement since the pulse picker is turned off first before the stage decelerates. The relationship between the stage movement and pulse picker is shown in Fig. 4.8 (a)(ii), and Fig. 4.8 (b)(ii) shows that the transition is smoother between the two waveguide segments

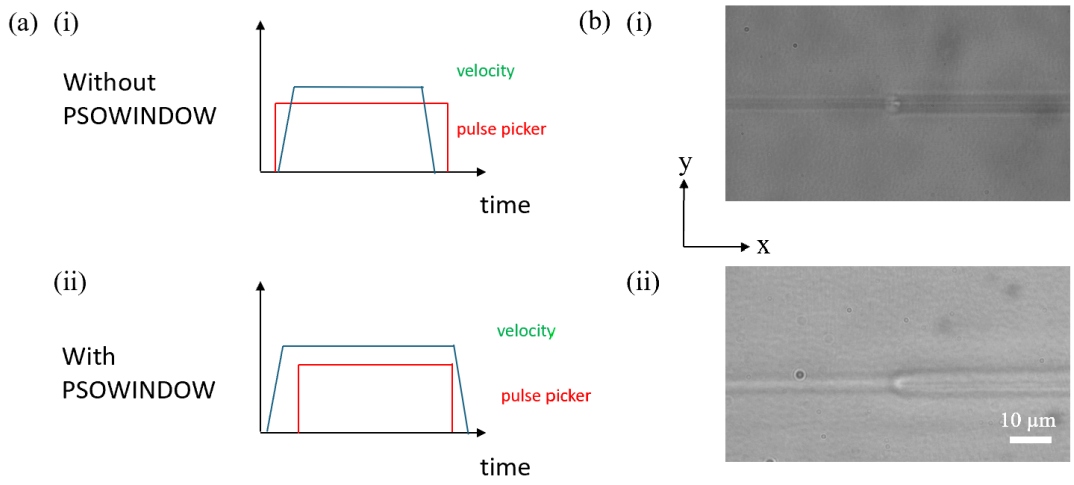

Figure 4.8: (a) Relationship between the stage velocity and pulse picker of the laser with time (i) without using **PSOWINDOW** command and (ii) using **PSOWINDOW** command. (b) Corresponding microscope images of the transition region of the two waveguide segments. (Contrast of images adjusted for visual clarity.)

Finally, as mentioned above the waveguides should stop before reaching the edge to avoid surface damage. Fig. 4.9 shows a comparison between waveguides with surface damage if not terminated appropriately. At the edge of the sample, in Fig. 4.9 (a) the surface damage can be observed as structures which scatter light, whereas in Fig. 4.9 (b) they are not observable for waveguides without surface damage.

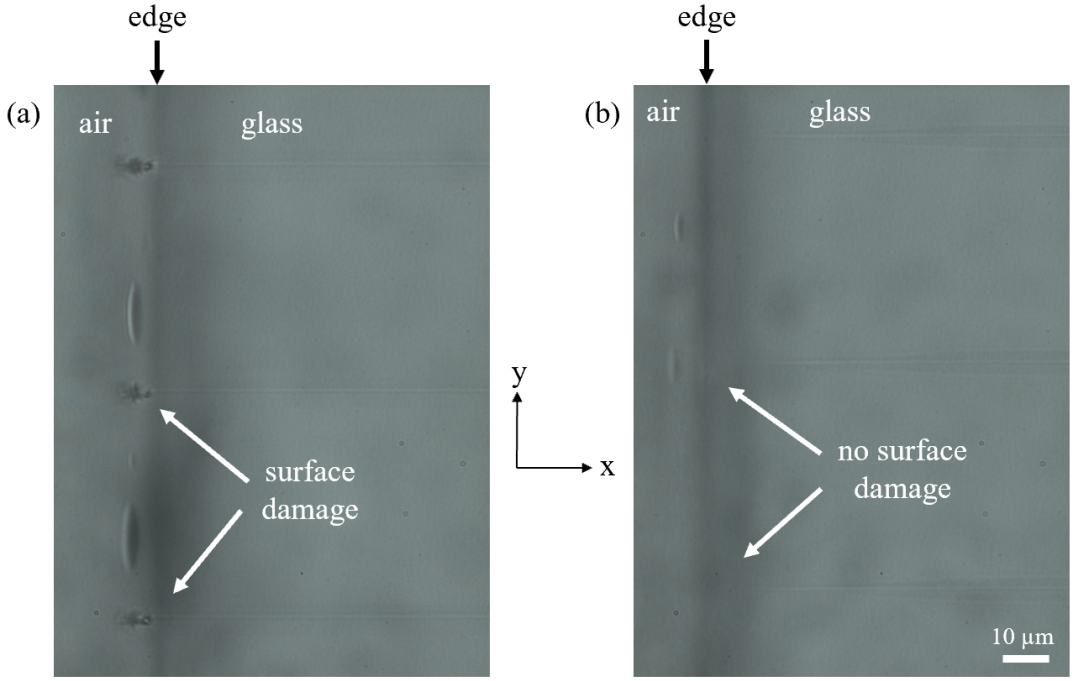

Figure 4.9: Images of waveguide outputs under microscope (a) with surface damage and (b) without surface damage.

Fig. 4.10 shows an alternative pair of images of the end facets of the waveguide. The surface damage scatters a significant portion of light and greatly decreases the coupling efficiency. It is therefore important for this writing scheme that the waveguide terminates just before reaching the sample edge, for which 10 μm is found to be a good compromise to both ensure the surface is not damaged and minimise coupling loss.

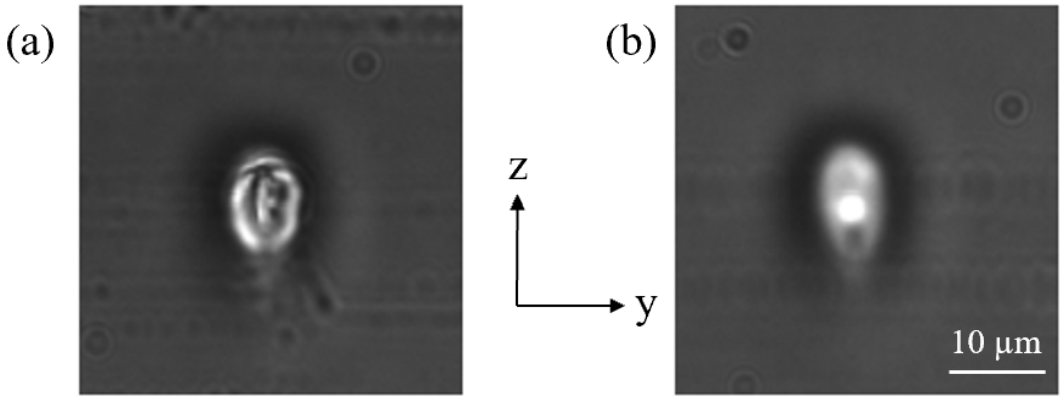

Figure 4.10: Images of waveguide outputs under LED illumination (a) with surface damage (b) without surface damage.

4.4 Experimental results

To verify how generally the proposed half pupil method applies to different materials, fabrication is performed in both fused silica and Eagle glass and their results are compared. The experiments can be divided into four main parts: the first part investigates how the waveguide terminates with a physical mask on the sample surface, which gives insight as to how edge aberrations affect the termination of waveguides and act as a model of in-situ fabrication where other components may have been pre-mounted to the glass sample. The second part compares the mode profiles of waveguides written with full pupil and half pupil respectively as an indication of how well the mode matches to a single mode fibre which acts as a standard for measuring coupling loss. The third part attempts to calibrate the difference between the z-positions of the centres of the modes of the waveguide segments written with full pupil and half pupil (z-offset) to ensure the two waveguide segments are properly aligned to reduce coupling loss, and the final fourth part measures the transmission loss of the full writing scheme combining the two segments.

If we want to fully eliminate polishing, the full writing scheme should start from the opposite end of the sample with the other half mask applied in a symmetric manner. However, to investigate the effects of this new writing scheme, it will only be used on one end of the sample in our experiments, where the other end is still polished in the conventional manner for consistency and better comparison.

4.4.1 Effect of edge aberrations on waveguide termination

To understand the effect of edge aberrations on fabricated waveguides, it would be instructive to observe how they terminate close to the edge. However, directly observing the edge termination is difficult using a standard microscope, since the edge obstructs the viewing and the imaging also suffers from the same edge aberrations. To work around this, it is possible to artificially create an "edge" by simply attaching a piece of foil on the surface of the glass sample. The edge of the foil effectively acts as an edge, with the added benefit that this also models the situation of in-situ fabrication where there are other

externally coupled components, and we want to avoid the fabrication beam extending beyond the edge and damaging the end facet where the components connect. By writing the waveguide over the edge of the foil, it is also possible to gain an initial understanding of how a physical mask on the sample surface may affect fabrication.

Fig. 4.11 shows the schematic of how the experiment is set up. A piece of tin foil is placed on top of the glass sample, so the foil edge is used to model the effect of the edge of the glass sample. Two sets of four different waveguides are fabricated in fused silica and Eagle glass respectively at 375 μm and 750 μm depth, each including the four cases of full pupil writing through the edge, full pupil stopping 10 μm before the edge, half pupil writing through the edge, and half pupil stopping 10 μm before the edge. The waveguides are written in the x-direction. The dashed red line denotes the position of the foil on the surface which acts as an indicator for the corresponding positions of the foil in Fig. 4.12 and Fig. 4.14. The middle segment is not drawn to avoid cluttering the figures.

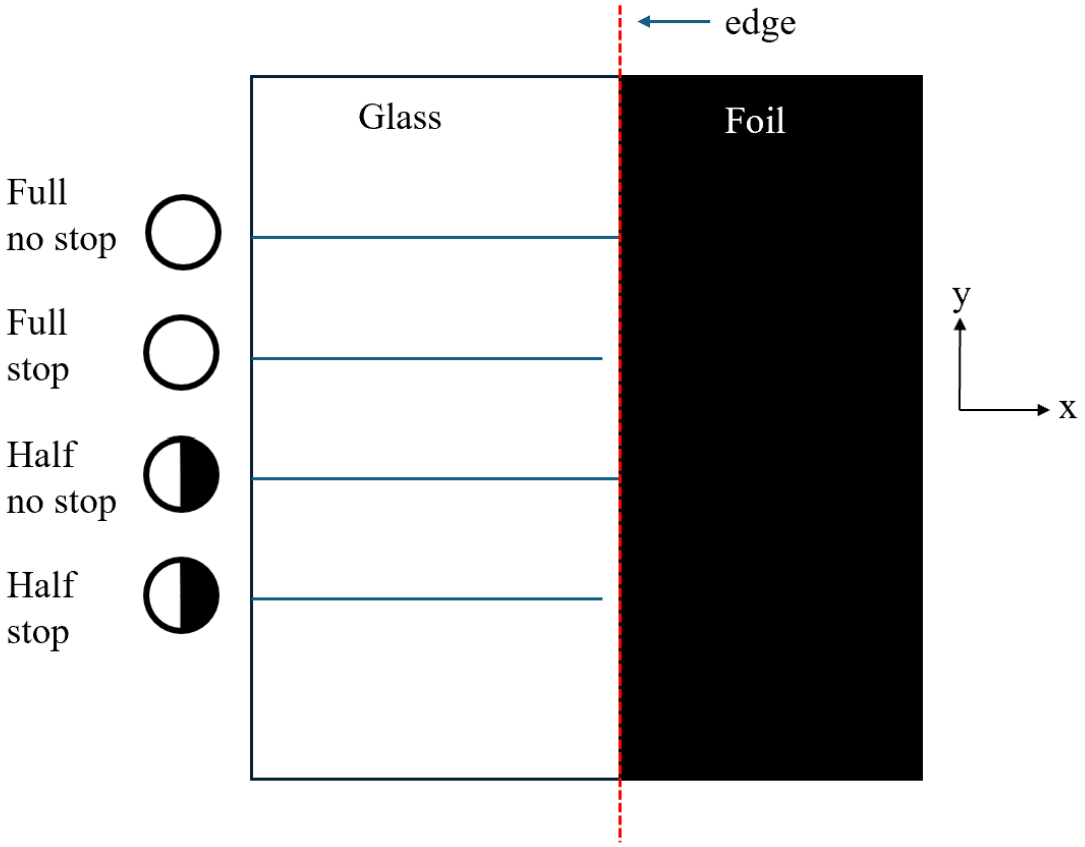

Figure 4.11: Schematic showing how the foil is placed on the glass sample relative to the written waveguides. "Full" and "Half" refer to the pupil used, "no stop" indicates the waveguide is written across the edge and "stop" indicates the waveguide terminates 10 μm before the edge of the foil.

Fig. 4.12 shows the results in fused silica at 375 μm and 750 μm fabrication depth. The dashed red line indicates where the position of the foil edge is on the surface of the sample.

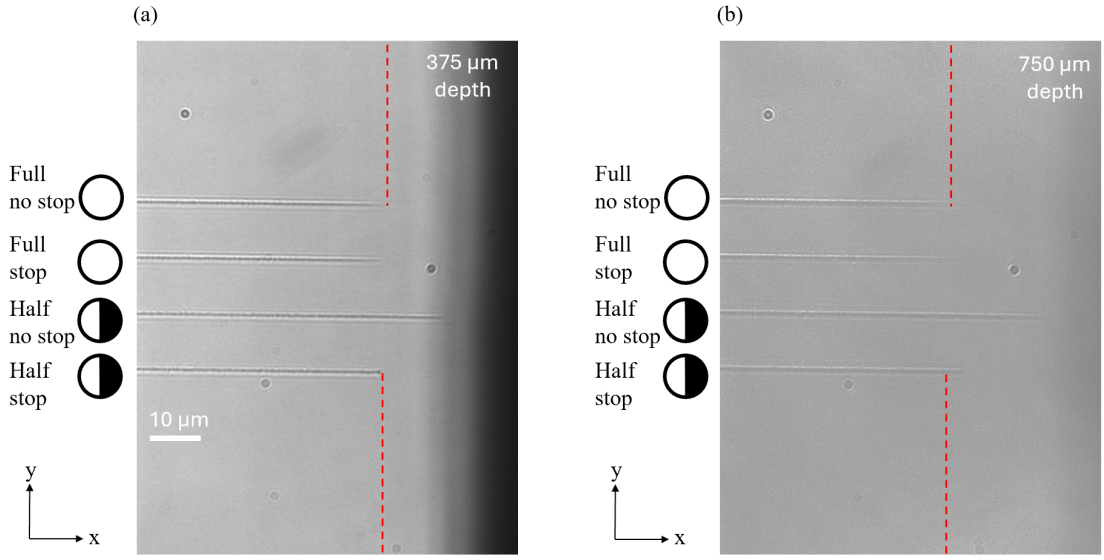

Figure 4.12: Experimental image of the waveguides written in fused silica. Dashed red line indicates the position of the foil edge on the surface, and labels are same as that of Fig. 4.11. (Contrast of images adjusted for visual clarity.)

Firstly, comparing the two waveguides written with full pupil and half pupil without stopping, the waveguide written with full pupil tapers off when it approaches the edge. This is due to the edge aberrations reducing the laser fluence delivered to the laser focus, thus weakening the material absorption effects. The effect is more apparent at the greater depth of 750 µm.

On the other hand, the waveguide written with half pupil extends beyond the edge. This is a result of how the fabrication beam cone focuses even when the centre of the beam has moved beyond the foil edge on the surface. Fig. 4.13 shows a schematic illustrating the effect. In Fig. 4.13 (a), the case where the waveguide is fabricated with full pupil is shown. At position 1, the focusing cone does not intersect with the foil, and the waveguide maintains its full structure. When moving to position 2 (where the centre of the beam is right at the edge), part of the focusing cone is blocked by the foil, thus reducing the power delivered to the focus, causing the waveguide to start tapering off. As the focusing cone reaches position 3, some rays still reaches the focus due to refraction, thus explaining why the waveguide extends beyond the foil edge on the surface. The situation is slightly different for the half pupil case shown in Fig. 4.13 (b), where the

waveguide retains its structure up to position 2 as the half-cone does not get blocked by the foil, and only beyond this point does the tapering effect begin. This is also strong evidence that the half pupil can write up to the sample edge avoiding edge aberrations. This explains why the extension beyond the edge is much longer for the waveguide written with half pupil.

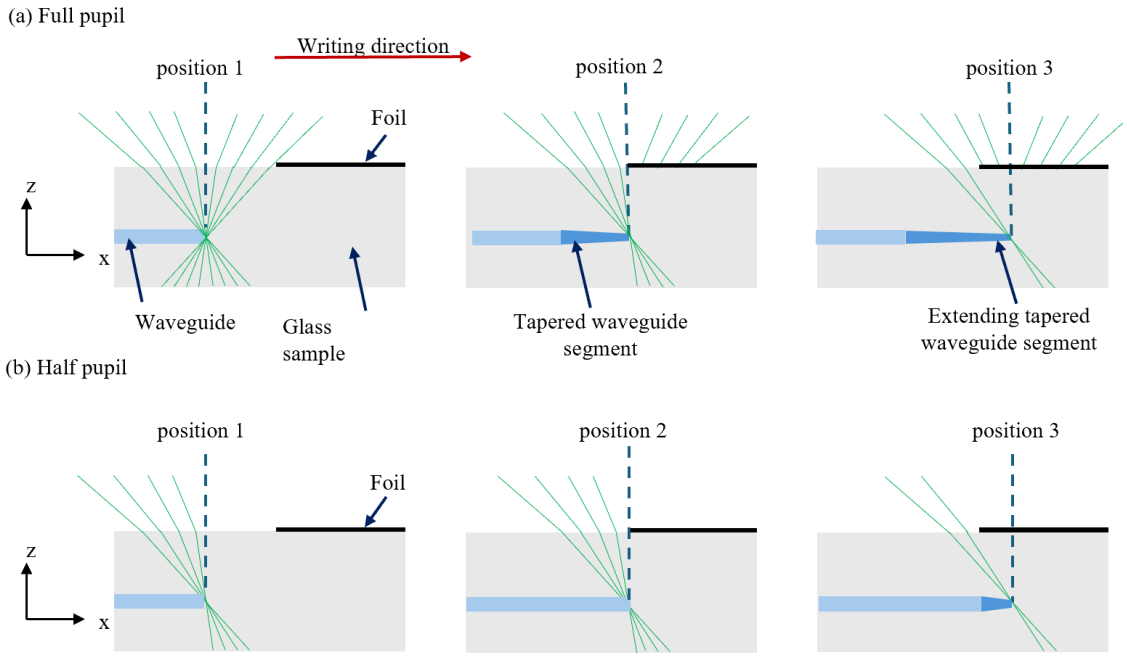

Figure 4.13: Schematic showing how rays forming the focusing cone extend under the foil even when the centre of the beam has moved past the edge. (a) Fabrication with full pupil, where the waveguide starts tapering off when the edge of the cone starts intersecting with the edge. (b) Fabrication with half pupil, where the waveguide maintains its full structure until position 2, where the centre of the beam reaches the edge, and the tapering only happens beyond that.

Secondly, comparing the two pairs of waveguides that stop before the edge at 375 μm and 750 μm depths also illustrates this effect. At 375 μm, the waveguide written with half pupil stop right at the edge maintaining the full waveguide structure, while the waveguide for the full pupil case tapers off. Furthermore, it can be observed that at the greater depth of 750 μm, the waveguide written with the full pupil tapers off more significantly. Again, this is a result of the edge aberration, which occurs earlier at greater depth hence the tapering off is more observable.

The corresponding fabrication in Eagle glass is shown in Fig. 4.14. It is apparent

that the width of the waveguides written with half pupil of both 375 μm or 750 μm fabrication depth in Eagle glass is bigger than those for full pupil, indicating the structure of the waveguide is larger. The increased structure size may be a result of the elongated focus, which may create additional structures thus increasing the apparent size of the waveguides due to the lower bandgap in Eagle glass (see Chapter 2). Eagle glass enters a cumulative heating regime whereas fused silica does not [76]. The material modification process is therefore different in the two substrates, resulting in different mode sizes which is critical for this application. Moreover, Fused silica has a higher fabrication threshold, therefore when power decreases the fabrication cuts off more abruptly whereas in Eagle glass there is a smoother transition.

For the same reason, the waveguides in Eagle glass extend much further beyond the edge for both the full pupil and half pupil cases compared to those fabricated in fused silica. Since the refractive indices of fused silica and Eagle glass is quite close, this phenomenon can only be attributed to the difference in material modification process. An additional observation that can be made is that the tapering off for the waveguides written with full pupil is less significant compared to the fused silica counterpart, again due to Eagle glass being more susceptible to material modification.

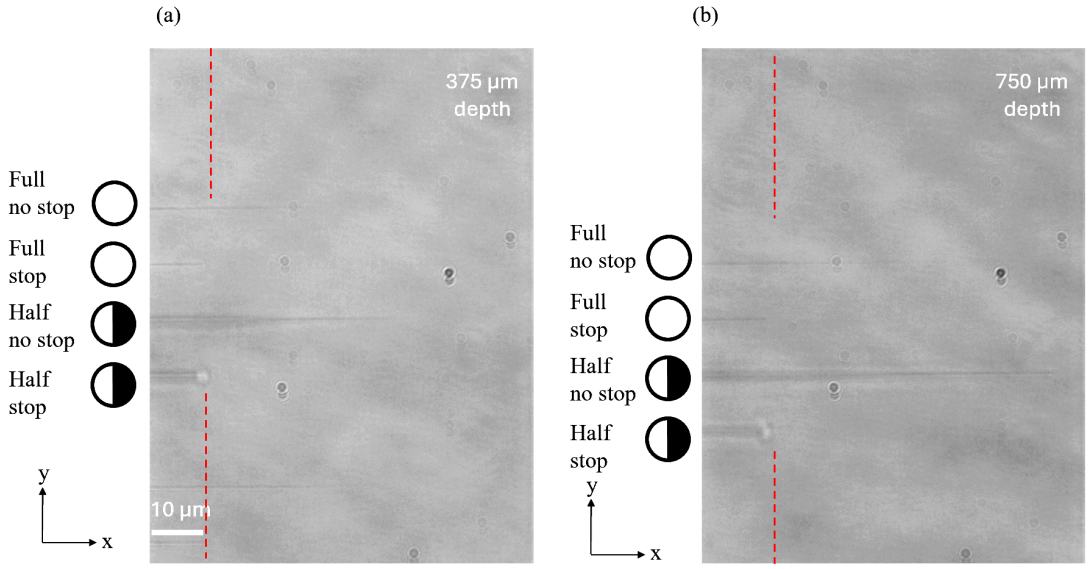

Figure 4.14: Experimental image of the waveguides written in Eagle glass. Dashed red line indicates the position of the foil edge on the surface, and labels are same as that of Fig. 4.11. (Contrast of images adjusted for visual clarity.)

In summary, it is possible to observe the behaviour of waveguide termination under the influence of an edge on the surface. The fabrication extending beyond the edge itself suggests that it is insufficient to simply place a physical mask directly on top of the sample, and writing over the edge may cause ablation on the surface or damage components beyond the edge (if any).

4.4.2 Mode size measurements

In this subsection, for ease of discussion, I shall use x- and y- axes of the cross-section of a fabricated waveguide to refer to the waveguide mode profiles (discussed in chapter 2).

The waveguide mode size gives quantitative information of the properties of the waveguides which allows the calculation of their overlap integrals with the mode of a single mode fibre. The overlap integral is a useful indicator of the coupling loss (see section 2.4.3). For this subsection, only the mode profiles were of interest, therefore the waveguides were written uniformly either with the full pupil or half pupil through the whole sample, with the opposite end cut and polished to expose the waveguides for

coupling from fibre. A schematic of the setup is shown in Fig. 4.15. Multiple waveguides were written at different depths with the correct depth aberration compensation applied, where each of their transmitted modes were imaged and the imaging system was focused on the surface of the sample end facet.

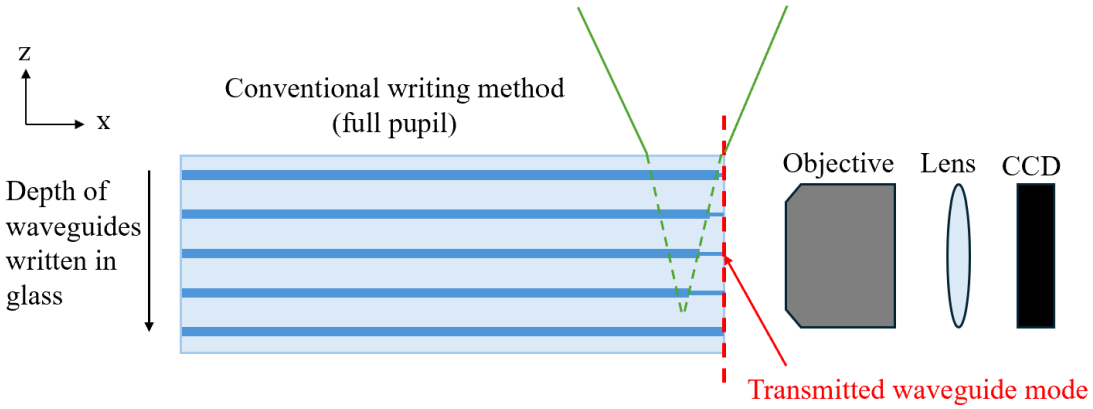

Figure 4.15: Schematic for waveguide geometry and how the modes are imaged.

Fig. 4.16 shows how the waveguides look like at the output facet under LED illumination. It is evident that the waveguides written with full pupil shown in (a) seem to increase in size as depth increases, while those written with half pupil as shown in (b) are relatively consistent in size. As will be discussed later, the apparent increase in size is a combination of the main waveguide structure tapering off leading to poorer confinement and mode expansion, together with the existence of additional secondary structures in Eagle glass.

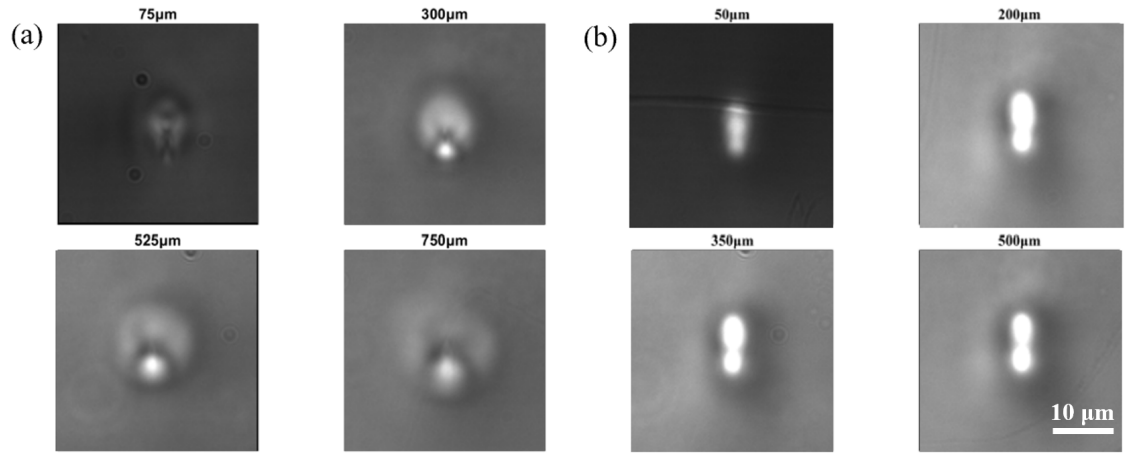

Figure 4.16: Images of waveguide outputs under LED illumination fabricated in Eagle glass at different depths. (a) Waveguides written with full pupil. (b) Waveguides written with half pupil.

To measure the modes, an optical fibre is butt-coupled to each waveguide from the other end of the sample as described in chapter 3. Fig. 4.17 shows example mode profiles for waveguides written in fused silica imaged for different depths, where (a) shows the modes of waveguides written with full pupil, and (b) shows the modes of waveguides written with half pupil. It can be immediately observed qualitatively that the mode size increases with increasing depth for the full pupil, and the waveguides written with half pupil remains approximately the same size.

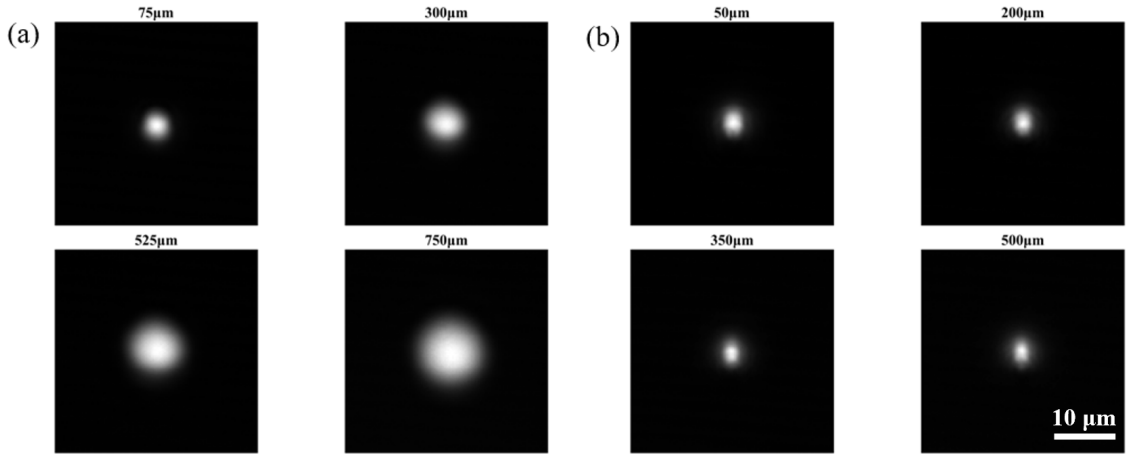

Figure 4.17: Modes of waveguides written in fused silica imaged at output facet without polishing. (a) Waveguides written with full pupil. (b) Waveguides written with half pupil.

Fig. 4.18 shows the processed data by fitting a Gaussian curve along the x- and y-directions through the centre of the mode measured at the end facet and estimating the $\frac{1}{e^2}$ mode field diameter. The results indicate that the waveguides written with half pupil (data points represented by circles) show consistent mode size at different depths, whereas the waveguides written with full mask (data points represented by crosses) show mode sizes increases almost linearly with increasing depth as a result of diffraction due to poorer confinement of the aberration-affected segment. Note that the y-direction for mode profile is actually the z-direction of the sample when considering the whole sample geometry. The inconsistency in nomenclature here arises from the fact that the mode taken at a cross-section orthogonal to the waveguide direction is considered as a two-dimensional object, hence it makes sense to use x- and y- mode profiles to describe the modes at the waveguide end facet. On the other hand, when considering the sample geometry, it is customary to refer to the direction of depth from surface (along the direction of the writing beam) as the z-direction in a transverse writing geometry. It should be clear from the context which coordinate system is being used, in particular only in the context of modes will the y-direction of the mode profile be along the direction of beam.

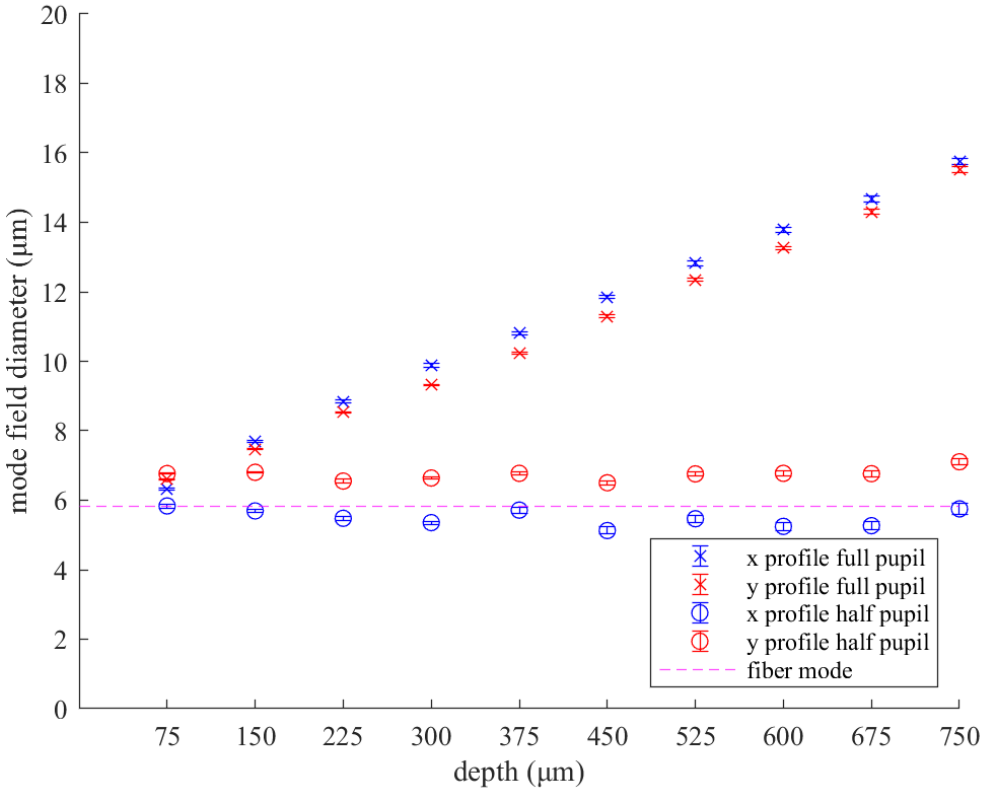

Figure 4.18: Estimated sizes of the diffracted waveguide modes measured at the end facet in fused silica.

The linearly increasing trend for full pupil can be explained in the following manner. The weaker fabrication due to edge aberrations means the mode is no longer properly confined inside the waveguide, and it spreads as it propagates from the end of the proper waveguide before aberration takes place towards the end facet. At greater depths, the edge aberration comes into effect earlier, which means that near the edge of the sample the waveguides start to taper off earlier as well. This results in more mode spreading, thus explaining the linearly increasing trend of the mode size with increasing depth.

On the contrary, the half pupil method shows that the mode field diameters of the waveguides can be maintained at a consistent size at different depths. In particular, a waveguide mode size comparable to that of a single mode fibre can be produced. This suggests that the half pupil method is promising to bridge the gap between waveguides written with full pupil and the fused silica sample edge.

Turning now to Eagle glass, the same sets of experiments are performed. Fig. 4.19 shows the example mode profiles for waveguides written with full pupil and half pupil respectively as in Fig. 4.17. The overall trend is still similar in that the mode sizes of waveguides written with full pupil increases with increasing depth, and the mode sizes of waveguides written with half pupil stay relatively consistent in size. However, the increase in size for the waveguides written with full pupil is not as uniform compared to that of those written in fused silica. Like the previous experiment, the aberrated focus due to edge aberration in Eagle glass is still able to create additional structures, which interacts with the propagation of the mode and finally leads to this distorted mode shape at the output.

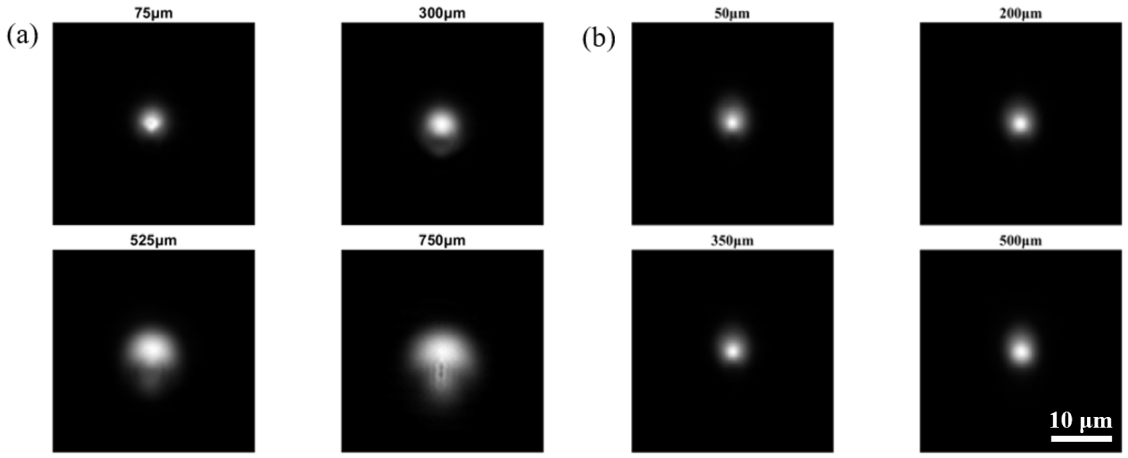

Figure 4.19: Modes of waveguides written in Eagle glass imaged at output facet without polishing. (a) Waveguides written with full pupil. (b) Waveguides written with half pupil.

Fig. 4.20 shows the corresponding $\frac{1}{e^2}$ mode field diameters of the waveguides in the same manner as Fig. 4.18. The increase in mode size with increasing depth starts at about 225 μm in Fig. 4.20, which is later than that of waveguides fabricated in fused silica which has close to a linear trend for all depths. This suggests that the edge aberration has less of an effect in Eagle glass. This is likely due to the lower fabrication energy threshold in Eagle glass such that the aberrated beam is still within the optimal fabrication window at small depths and thus only has a substantial difference at greater depths.

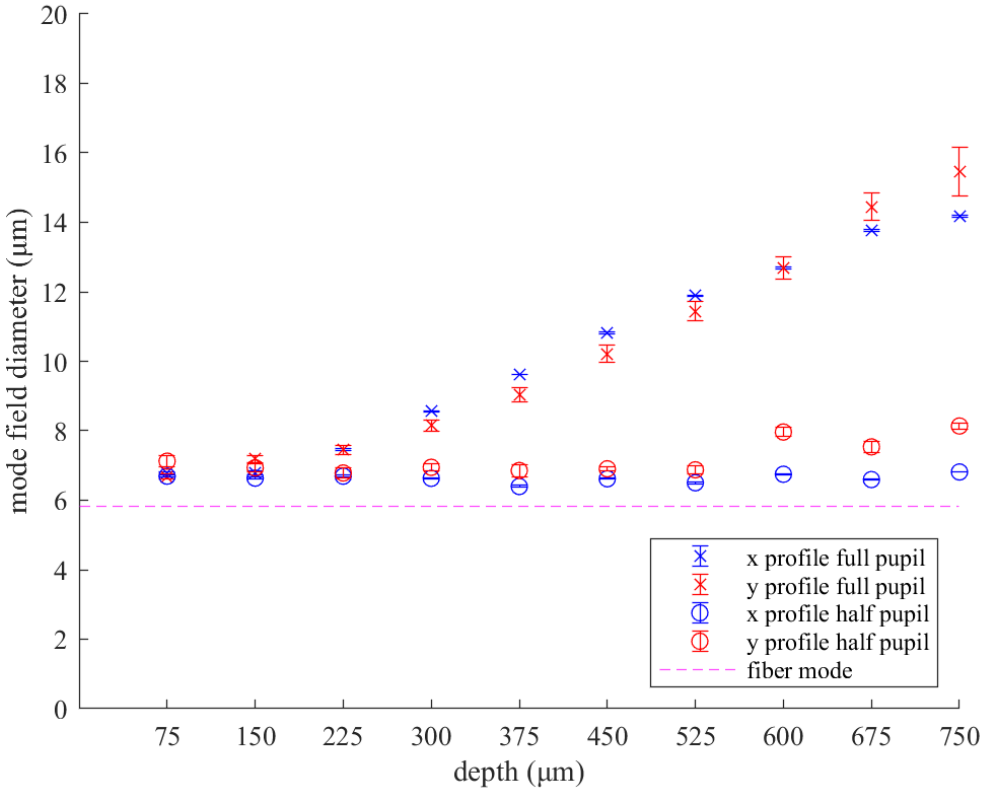

Figure 4.20: Estimated sizes of the diffracted waveguide modes measured at the end facet in Eagle glass.

To further understand this effect, the mode profiles of waveguides fabricated with full pupil in Eagle glass at 75 μm and 750 μm under the sample surface are plotted in Fig. 4.21. The measured data points are the optical intensities of the mode profile obtained by taking a section along the y-direction through the centre of the mode of the waveguides (see Chapter 2). The measured intensities are divided by the largest measured intensity to obtain the normalised intensities. In Fig. 4.21 (a), the mode for 75 μm is more or less symmetric and regular, as indicated by the Gaussian fit matching quite closely with the measured data. On the other hand, in Fig. 4.21 (b), at 750 μm the mode is much more irregular, and the fitted Gaussian does not accurately represent the measured mode shape. Along the y-mode profile the plot is not smooth and there is a secondary lobe (centred around 25 μm on the x-axis of the plot). This suggests that secondary structures are indeed formed close to the surface of the sample due to edge

aberrations, which may be related to the commonly observed formation of side lobes at lower fabrication energy [77]. Fitting a simple Gaussian no longer matches the mode well, which is the reason why the errorbar of Fig. 4.20 increases at greater depths for the y-profile of the waveguides written with full pupil. The errorbar here is proportional to the sum of squares difference between the best-fit Gaussian curve and the data points, indicating that the Gaussian curve is no longer a good model at greater depths and the mode no longer has a Gaussian profile.

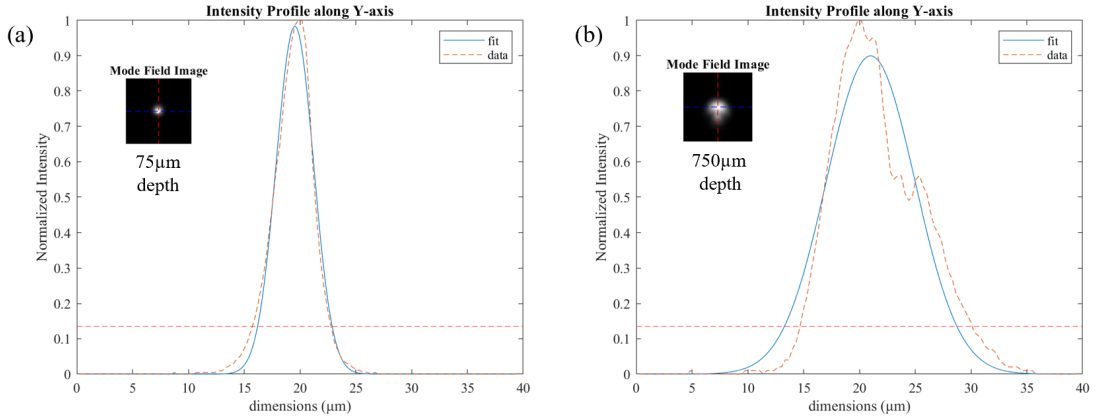

Figure 4.21: Mode intensity profile of waveguides fabricated with full pupil in Eagle glass at (a) 75 µm (b) 750 µm under the sample surface. The horizontal dashed orange line indicates the $\frac{1}{e^2}$ intensity, from which the mode diameter is determined. Insets are the same mode images from Fig. 4.19 and the dashed red line indicates where the cross-section is taken.

Finally, the overlap integral between the measured waveguide modes and the mode of a single mode fibre is calculated. The results are shown in Fig. 4.22. This shows more clearly that transmission of the waveguides written with full pupil drops off more or less linearly with increasing depth, indicating that edge aberrations does significantly affect coupling losses and explains why polishing is often required in standard fabrication. On the other hand, those written with half pupil matches is almost constant, showing that their mode profiles match well with the single-mode fibre and can maintain 80% to 90% throughput even at greater depths. This suggests that fabricating with the half pupil method near the sample edge significantly improves mode-matching and reduces coupling losses.

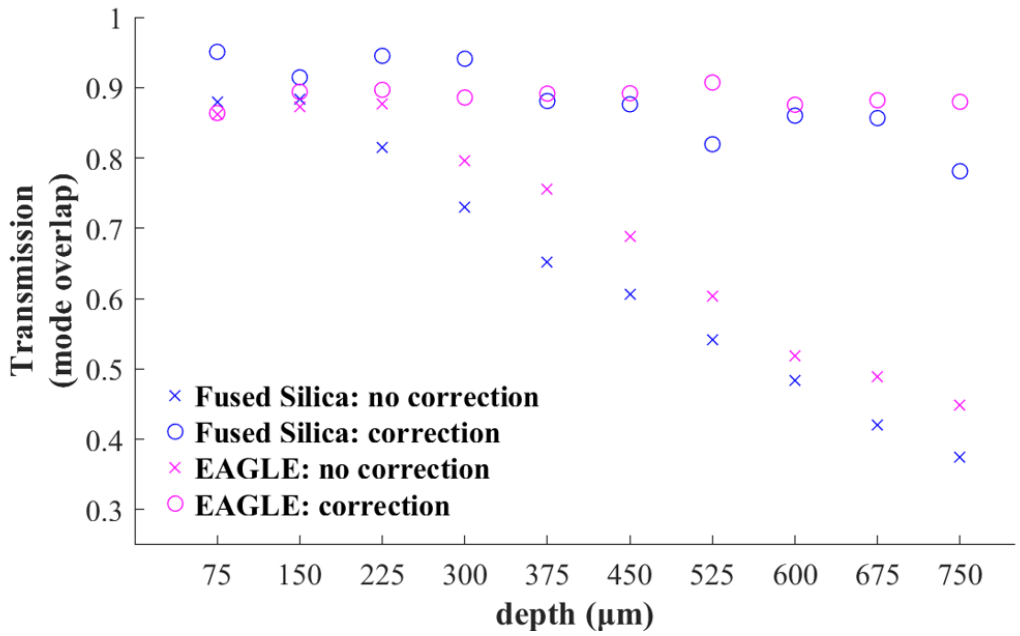

Figure 4.22: Estimated mode overlap losses of waveguides written in Eagle glass with respect to a single mode fibre.

To summarise the findings above, edge aberrations lead to poorer confinement of the waveguides written with full pupil, and the effect is more pronounced at greater depths. On the other hand, the half pupil method retains a relatively consistent mode across different depths in both fused silica and Eagle glass. Moreover, it was found that edge aberrations affect the modes of fabricated waveguides differently in fused silica and Eagle glass due to the difference in their material properties. The calculation of overlap integrals also suggests that the half pupil method is a promising strategy to match the fibre mode by fabricating up to the sample edge without the need to polish the sample and also minimise coupling losses.

4.4.3 z-offset calibration

Due to the different distribution of energy at the focus of the fabrication laser beam as confirmed in Fig. 4.5, the respective non-linear energy absorption process is expected to differ between waveguide writing using the full pupil and the half pupil. Moreover, since only half of the pupil is applied for phase correction, the depth-dependent aberrations

also cannot be fully corrected since only half of the azimuthal angles are present in the corresponding Zernike modes. As a combination of these two effects, it is observed that there is an offset in the z-direction between the two waveguide segments, implying that the centres of the waveguide modes are misaligned. The difference between the z-positions of the centre of the mode of the waveguid segments written with full pupil and half pupil will be referred to as the z-offset.

Fig. 4.23 shows a schematic how the z-offset manifests during fabrication. When switched to the half pupil, if all other depth-related settings remain the same, there could be a z-offset between the two segments leading to misalignment and increases the overall loss of the combined waveguide.

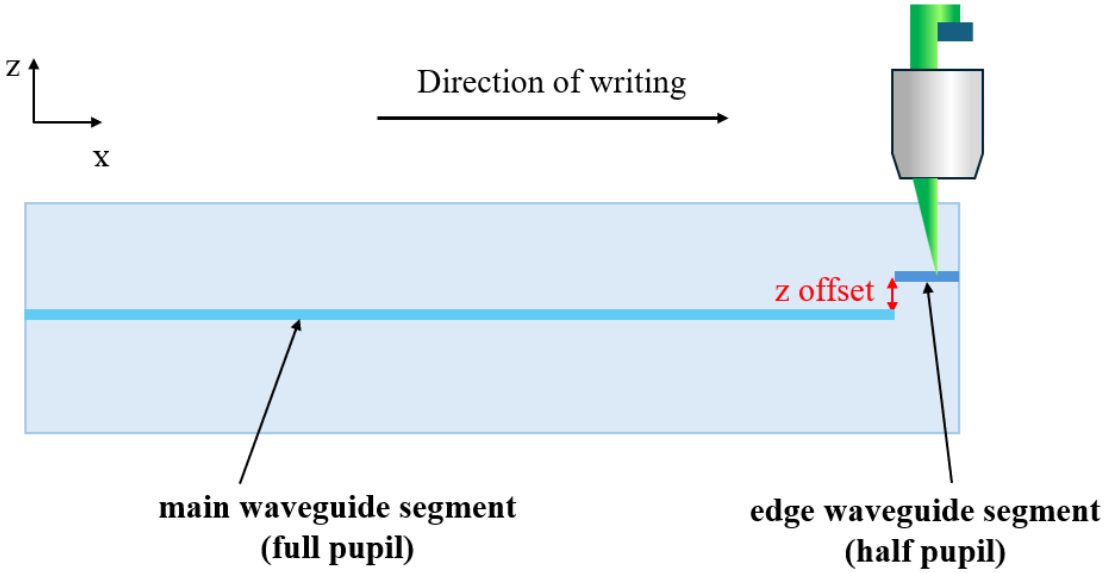

Figure 4.23: Schematic of z-offset between the two waveguide segments written with full pupil and half pupil.

Fig. 4.24 shows how the z-offset is characterised in experiments. Waveguides written with full pupil and half pupil are fabricated at the same depth settings, which is then polished back to expose the waveguides since the goal is to obtain the position of the waveguide modes. The red dashed line indicates the z-position of the centre of the mode written with half pupil. It is used as a reference from which the z-position of the waveguide modes written with full pupil are determined and the z-offset is then recorded.

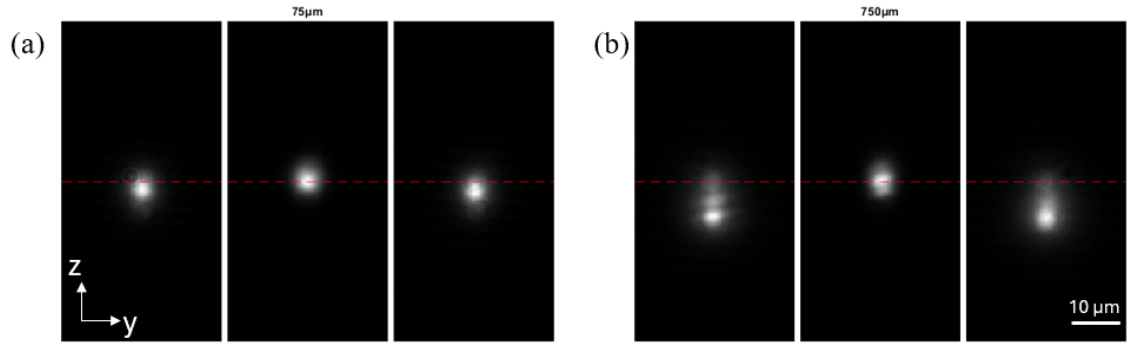

Figure 4.24: Waveguide modes for calibrating z-offset at different depths below the sample surface: (a) 75 µm (b) 750 µm. The three images from left to right in (a) and (b) are written with: full pupil, half pupil and full pupil, respectively. The red dashed line indicates the z-position of the centre of the mode written with half pupil.

Fig. 4.25 shows the z-offset as a function of fabrication depth. For this set of data, the z-offset is about 2 µm to 3 µm for shallow depths, then increases significantly from 600 µm fabrication depth. This set of data is used as a reference for the fabrication in section 4.4.4.

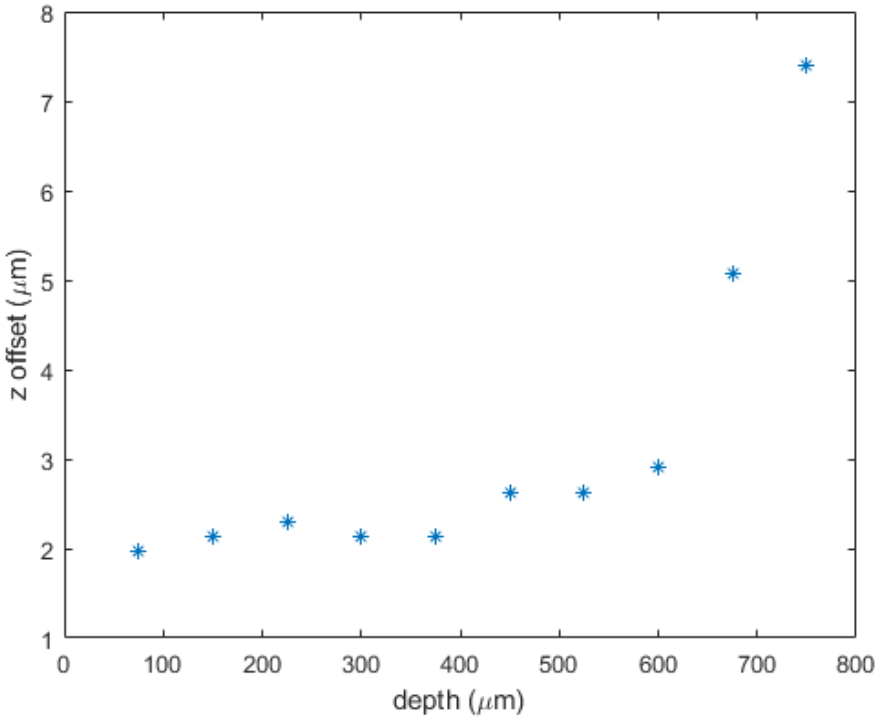

Figure 4.25: Difference between the z-positions of waveguides written with full pupil and half pupil (z-offset) as a function of fabrication depth.

However, it is later noticed that the waveguide modes in Fig. 4.24 may indicate there is some issue with the fabrication. Since the sample is polished, the waveguides written with full pupil should exhibit similar mode shapes even at greater depths, whereas in Fig. 4.24 (b) the fabrication at 750 μm clearly shows an elongated mode where the waveguiding region is actually the lower part of the waveguide structure. This secondary side lobe in the images is an indication of insufficient energy delivered to the focus of fabrication, indicating there may be an error in setting the fabrication power or incorrect aberration correction.

The fabrication at greater depths should be studied further to ensure the accuracy and reliability of results. Nevertheless, it is mostly unlikely in most practical situation to require fabrication as such depths. The calibration procedure developed here can be applied in future studies to give insights to the transition region between the two waveguide segments.

4.4.4 Throughput estimation and measurements

Equipped with the above knowledge, the transmission throughput of the waveguides using the new writing scheme can now be measured. Sets of waveguides written with the new writing scheme were fabricated with different depths, with z-offsets ranging from 1.5 μm to 4.5 μm. Fig. 4.26 shows the improvement in transmission relative to a waveguide written only with the full pupil at the same fabrication depth.It is noted that the experiments do not show a conclusive trend of how the z-offset value that gives the highest throughput depends on depth. The variations of transmission enhancement with respect to the z-offset could be due to experimental errors and this could be a direction for future studies.

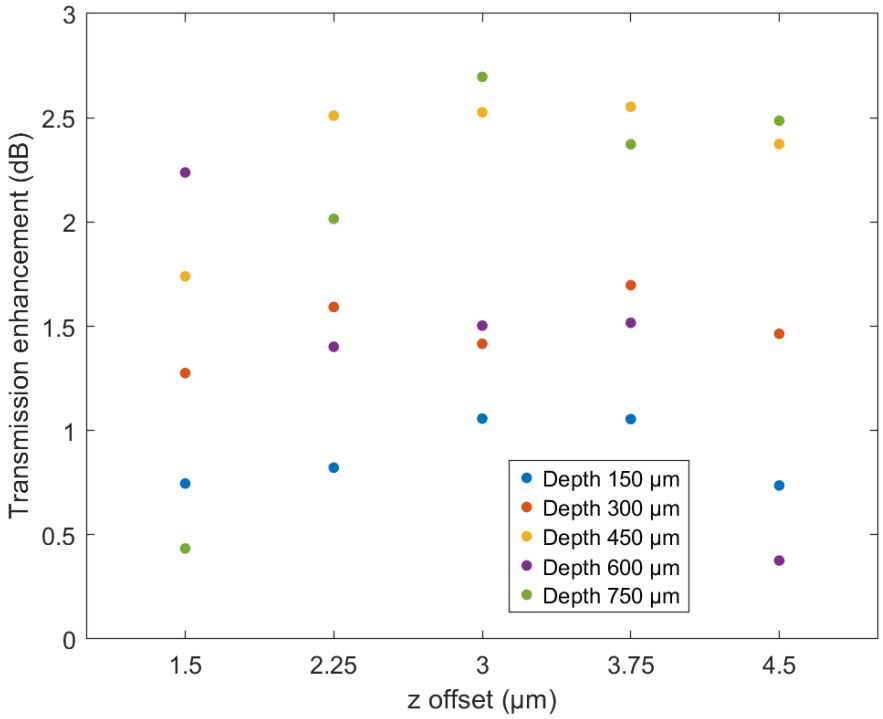

Figure 4.26: Throughput measurements with varying z-offsets at different fabrication depths.

Finally, the throughput enhancement of the combined waveguide with highest transmission throughput over the full pupil waveguide is shown in Fig. 4.27. The results shows that the new method improves the overall transmission of waveguides compared to those without corrections. The improvement increases with depth, which is expected since the effect of edge aberrations also increases with depth. However, it is noted that the error bars are quite large, meaning that the same measurement procedure cannot consistently produce reliable results. This is due to the current measurement setup cannot align the fibre with the waveguides in a consistent manner, thus leading to high variance in power measurements. Improving the measurement setup should lead to better precision in the results.

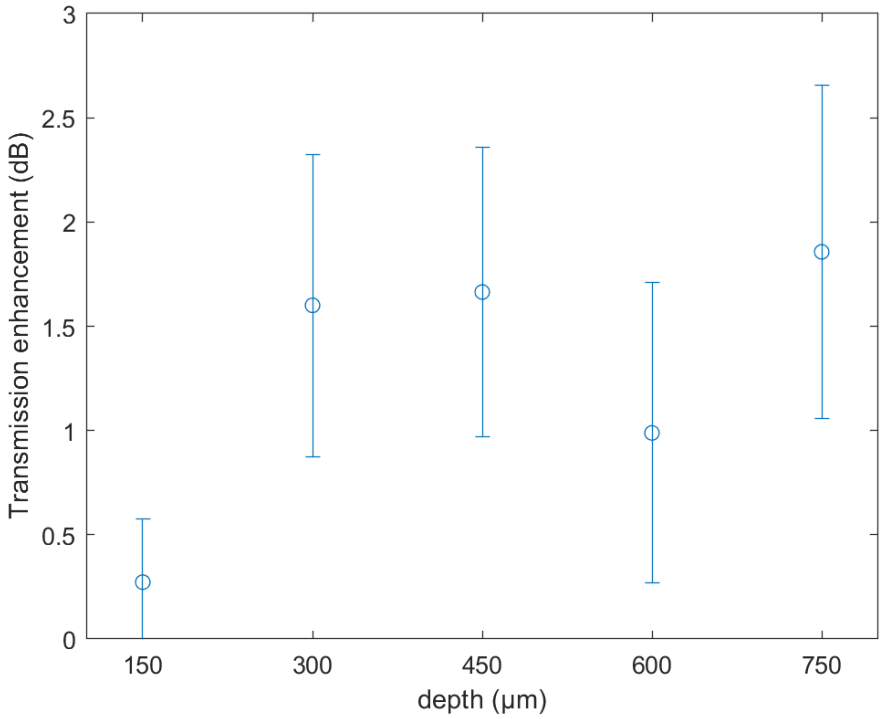

Figure 4.27: Throughput measurements with varying z-offsets at different fabrication depths.

4.5 Conclusion and outlook for future work

In conclusion, a new fabrication method for compensating for edge aberrations is presented and investigated. The novelty is in writing the segment of a waveguide up to a sample edge with half of the original pupil and increasing the energy instead of using a full pupil all the way which suffers from edge aberrations. The results showed that this method is a promising proposal to allow for fabrication of waveguides up to the sample edge, potentially eliminating the need of polishing after fabrication and enabling in-situ fabrication in the future.

The difference between the material modification characteristics in fused silica and Eagle glass is compared. It is noted that in Eagle glass, edge aberrations gives rise to the formation of secondary structures near the sample edge, leading to an irregular mode at the output facet, whereas in fused silica the expansion of the waveguide mode due to poorer confinement is more regular. Regardless of the actual shape, in both materials the expansion of the waveguide modes leads to a similar loss in the overlap integrals

between the waveguide modes and the optical fibre mode, indicating that the coupling loss increases with increasing fabrication depth. On the other hand, using the half pupil fabrication method can ensure that the dimensions of the fabricated waveguide modes remain relatively consistent at different fabrication depths, giving an estimated loss of less than 20% from the overlap integral calculations.

It is also observed that without additional compensation, there is a vertical offset in the z-direction between the waveguide segments written with full pupil and half pupil due to difference in the focal energy distribution and depth-dependent aberration correction. A procedure is described to calibrate the amount of offset and compensate accordingly during fabrication. Finally, measurements confirm that the waveguides written with half pupil have propagation loss comparable to those written with the full pupil without incurring additional loss, which verifies the feasibility of the new writing scheme.

The half-pupil technique thus provides a robust and adaptable foundation for future developments in in-situ fabrication. A logical next step will be to demonstrate the method in an integrated configuration, such as writing directly between pre-aligned fibre arrays mounted in v-groove substrates and the glass block, thereby realising coupling without polishing or manual alignment. While this would require highly precise spatial matching between input and output components, techniques such as tomographic refractive index profiling [78] could enable the accurate localisation of guiding structures within the substrate, paving the way for heterogeneous photonic integration.

On the fabrication side, the current workflow involves coordination between several independent control systems (for example, the Aerotech stage controller, LabVIEW-based SLM control, and motorised waveplate driver), each requiring manual reconfiguration when switching between full-pupil and half-pupil operation. Developing a unified control interface to automate these processes would not only streamline fabrication but also reduce alignment time and the potential for human error. Such an integrated control framework would represent an important step toward scaling the technique for larger and

more complex photonic devices.

Finally, future work could focus on enhancing fabrication stability through online monitoring and feedback. Incorporating an online monitoring module into the fabrication system would provide valuable insights into the quality of fabricated waveguides, allowing real-time correction without lengthy recalibration routines. Similarly, improving the characterisation setup with a dedicated coupling platform and vibration-isolated mounts would reduce measurement variance and ensure more reliable assessment of propagation and coupling losses. Together, these developments would enable the half-pupil method to evolve from a proof-of-concept demonstration into a practical and scalable fabrication tool for next-generation integrated photonic systems.

4.6 Acknowledgements

I would like to thank my supervisors Professor Patrick S. Salter and Professor Martin J. Booth for their preliminary investigations on correcting edge aberration [74]. Professor Patrick S. Salter also suggested the idea of using a foil to model the effects of an edge as outlined in section 4.4.1. Additionally, I would also like to thank Dr. Mohan Wang who kindly assisted in fabrication and data collection for waveguides in Fused Silica, and for valuable discussions on the implementation of the writing scheme. All other experiments and data processing were performed by me.

CHAPTER 5

Multiscan waveguide design for spatial mode converters

5.1 Introduction

Integrated photonics has witnessed rapid advances in recent years, driven by the demand for compact, low-loss, and reconfigurable optical components. Among the various functionalities required, the ability to manipulate spatial modes is of fundamental importance. The orthogonality of spatial modes implies that they can encode information as independent channels, which enlarges the capacities of classical optical communication and quantum communication networks [79]. Such control underpins applications in optical communication, mode-division multiplexing, quantum photonic circuits, and emerging areas such as neuromorphic computing [80, 81]. In addition to information processing, some optical systems may require interfacing with components with unique spatial mode profiles, therefore the ability to tailor the mode shapes can enable new demonstrations and applications [82].

Traditionally, devices such as beam rotators, mode converters, and waveplates have been realised with bulk optical components, which demand precise alignment and occupy a significant footprint. The integration of these functionalities onto a chip platform offers dramatic improvements in scalability and stability. Recent developments have shown that by tailoring the cross-sectional geometry and refractive index distribution of waveguides, entirely new classes of integrated devices can be realised [83].

A central challenge in designing such structures lies in predicting how fabrication

schemes influence the guided mode distribution. Conventional approaches often require multiple fabrication iterations and extensive experimental optimisation, making development costly and time consuming. Numerical modelling therefore plays an indispensable role: it allows systematic exploration of geometrical parameters and enables the prediction of modal properties prior to physical realisation.

Among recent fabrication strategies, Spherical Phase Induced Multiscan Waveguide (SPIM-WG) inscription has emerged as a particularly promising approach [77]. By introducing controlled phase modulation to the writing beam during femtosecond laser inscription, SPIM-WG enables the direct creation of multiscan structures by carefully arranging multiple sequential scans. This capability provides access to complex cross-sections not readily achievable with conventional single-focus writing or lithographic techniques.

This chapter investigates the design, simulation, and fabrication of SPIM-WGs for spatial mode conversion. We begin with the principles of phase-modulated writing, then present eigenmode simulations performed with COMSOL Multiphysics to quantify how separation between waveguide scans, wavelength, and geometry determine the supported modes. Both symmetric and asymmetric cross-sections are explored, including circular arrays optimised for fibre coupling and tailored converters for matching to the skewed modes of ppKTP waveguides. Finally, experimental results in Eagle glass are compared with the numerical models, demonstrating strong agreement and highlighting the synergy between phase-modulated fabrication and simulation-driven design.

5.2 Principle of phase-modulated writing

SPIM-WG is a fabrication method developed with the aim of achieving high-precision, high-contrast refractive index modifications in glass by deliberately reshaping the focal volume of a femtosecond laser. Rather than simply relying on the diffraction-limited focus produced by the objective, SPIM-WG employs an additional phase modulation to engineer the intensity distribution within the focal region so that the resulting material modification more closely follows a designed spatial profile.

The key step is the introduction of a controlled spherical phase, imposed by a spatial light modulator, to redistribute energy across the focal volume. In practice, the dominant contribution arises from the Zernike spherical aberration mode Z_{11} with an amplitude near -1, which drives the formation of an "energy-relocated" focus. In this regime, the effective pulse energy is preferentially concentrated toward the lower part of the focal volume, suppressing upper-lobe structures that otherwise affect the modification process. As a consequence, SPIM-WG produces refractive index changes that matches well with the engineered focal shape while simultaneously achieving higher contrast and reduced lateral dimensions compared with unmodified focusing.

The mechanism underlying the enhanced refractive index contrast can be understood from the altered pulse intensity created by the aberrated focus. Because the imposed phase reduces the local peak intensity, higher pulse energies are required to reach the nonlinear ionisation threshold. Once plasma is generated, it becomes opaque to the writing wavelength, and the increased deposited energy produces a correspondingly larger refractive index change [84]. In addition, the modified plasma distribution inverts the longitudinal intensity profile which leads to a different refractive distribution compared to conventional fabrication parameters [85].

To refine the focal profile further, SPIM-WG incorporates higher-order spherical modes in conjunction with Z_{11}. These were explored using a Nelder–Mead optimisation routine [86], in which a target 3D focal intensity distribution representative of an ideal energy-relocated regime was specified. At each iteration, the algorithm simulated the resulting focal field, evaluated the mean-squared error relative to the target distribution, and updated the phase mask accordingly. The result from simulations confirmed that the converged solution remained dominated by $Z_{11} = -1$, with only minor contributions from higher-order terms such as Z_{37}. These additional modes serve primarily to suppress residual side lobes and fine-tune the axial extent of the modification region. Together, these results confirm that engineered spherical phase modulation improves the precision and control of SPIM-WG over conventional laser-written waveguide fabrication schemes

[87].

The effectiveness of this approach is illustrated in Fig. 5.1, which shows tomographic microscopy images of refractive index modifications in Eagle glass written with SPIM-WG. Optical tomography provides a non-destructive means of reconstructing the three-dimensional refractive index distribution within transparent materials. By capturing multiple intensity images of a sample from different angles and applying inverse reconstruction algorithms, each image encodes information about the optical phase delay accumulated along the light path, which is directly related to the local refractive index variations within the material. Using an proximal gradient descent algorithm, these phase projections are computationally combined to yield a three-dimensional map of the refractive index distribution with sub-micron spatial resolution [88]. In addition to the uniform positive index regions that form the waveguide cores, the images reveal the characteristic negative index lobes above the waveguide cores, a distinctive feature of this method. Rather than being undesirable, these regions provide an additional degree of freedom for engineering mode profiles, as later demonstrated in converter designs (see section 5.3.4).

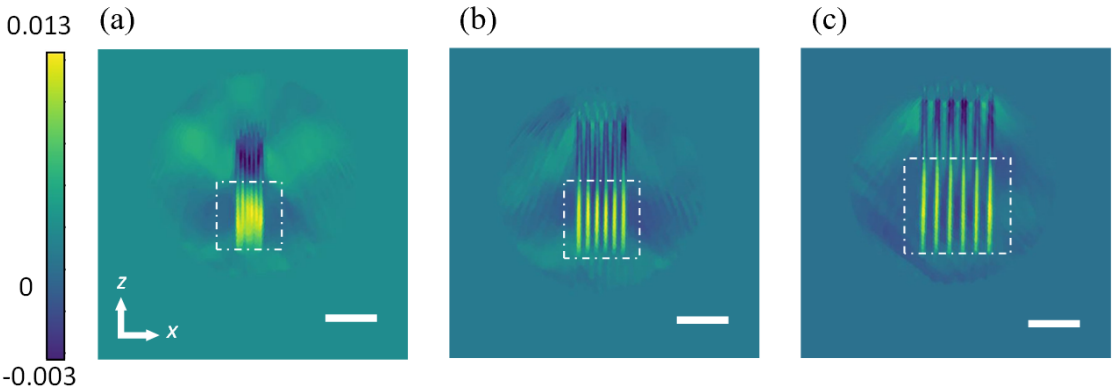

Figure 5.1: Adapted from [77]. Tomographic microscopy images of refractive index modifications in Eagle glass fabricated with the SPIM-WG method using a 6×6 multiscan scheme. Panels correspond to lateral and axial scan separations of (a) 0.5 µm, (b) 1.0 µm, and (c) 1.5 µm, respectively. The dashed boxes highlight the central multiscan region, while the colour scale represents the refractive index change. Scale bars: 5 µm.

This approach makes it possible to inscribe multiple closely spaced guiding regions by multiple sequential scans with the phase-modulated focus to build up complex geometries.

The resulting refractive index modification therefore reflects the designed arrangement of waveguide scans that overlap to form collective structures.

Overall, SPIM-WG provides a sufficiently high-precision method for inscribing complex multiscan geometries. By leveraging spherical phase modulation, it overcomes the limitations of conventional thermal-regime writing and offers a robust route to fabricating tailored cross-sections for spatial mode conversion.

5.3 Simulation studies of cross-sectional geometry

5.3.1 Modelling principles and procedure using COMSOL

Numerical simulations were carried out using the Wave Optics Module of COMSOL Multiphysics, which determines the eigenmodes of arbitrary waveguide cross-sections through a finite element method (FEM) framework. The workflow proceeds through a series of stages that transform an abstract refractive index distribution into concrete predictions of effective indices and field profiles.

The starting point of each simulation is the definition of material and structural parameters. Fig. 5.2 summarises the key quantities: refractive index contrast, operating wavelength, scan width, scan separation, and structural length. Together, these values specify the geometry and material composition of the system and set the conditions for the modes that the solver can return.

Parameters

Name	Expression	Value	Description
n_max	1.514	1.514	maximum refractive index
n_min	1.5	1.5	minimum refractive index
lda0	785[nm]	7.85E-7 m	wavelength
w	0.6[um]	6E-7 m	core width
sep	0.8[um]	8E-7 m	core separation
l	5[um]	5E-6 m	core length
f0	c_const/lda0	3.819E14 1/s	frequency

Figure 5.2: Key parameters defined in COMSOL for the multiscan waveguide simulations reported in this chapter, including refractive index contrast, operating wavelength, scan width, scan separation, and structural length. These values define the geometry and material composition of the model and directly determine the modal properties obtained from the COMSOL solver.

The next stage is geometry construction. In SPIM-WG devices, this is modelled by embedding Gaussian-shaped refractive index modifications in the substrate to represent the laser-written structures. The Gaussian index profiles are specified inside the material properties section, and by arranging these waveguide structures in different geometries, the intended multiscan cross-sections can be reproduced. Fig. 5.3 shows one such example, where rectangular outlines indicate the positions of the individual waveguide scans within the computational window. The operating wavelength is specified at this stage, as it strongly influences the effective mode indices and confinement.

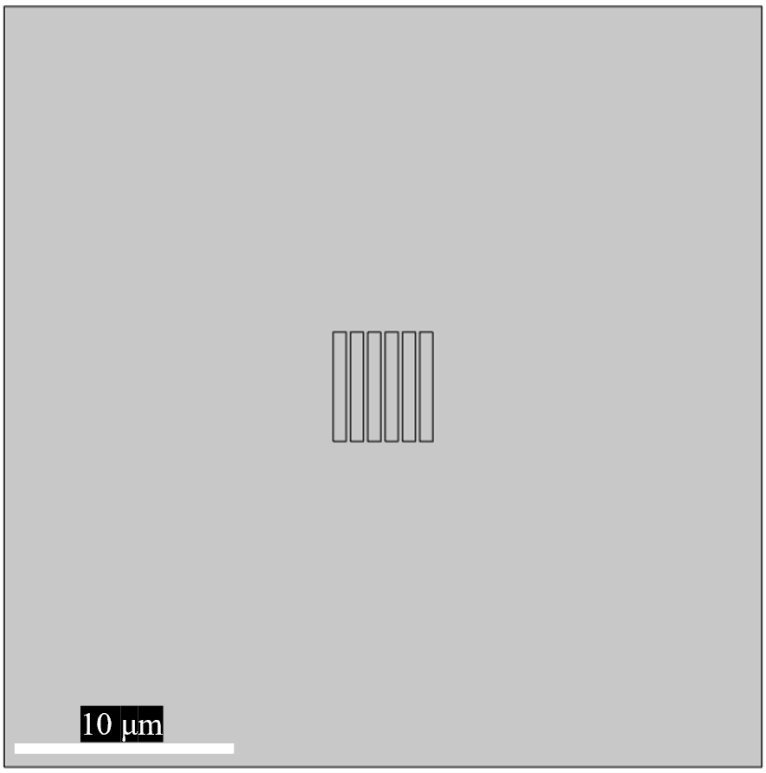

Figure 5.3: Geometry of a simulated multiscan waveguide cross-section. Rectangular outlines mark the positions of the individual laser-written waveguide scans within the computational window.

Once the geometry is established, the computational domain is discretised into a finite element mesh. The mesh must be fine enough around the structures to resolve the confined fields, while a coarser mesh can be used in the cladding to reduce computational effort. Fig. 5.4 illustrates this strategy, with dense elements around the guiding region and progressively larger elements at the periphery. The COMSOL solver has presets that automatically generates meshes, and for the simulations presented here a “Fine” mesh is used.

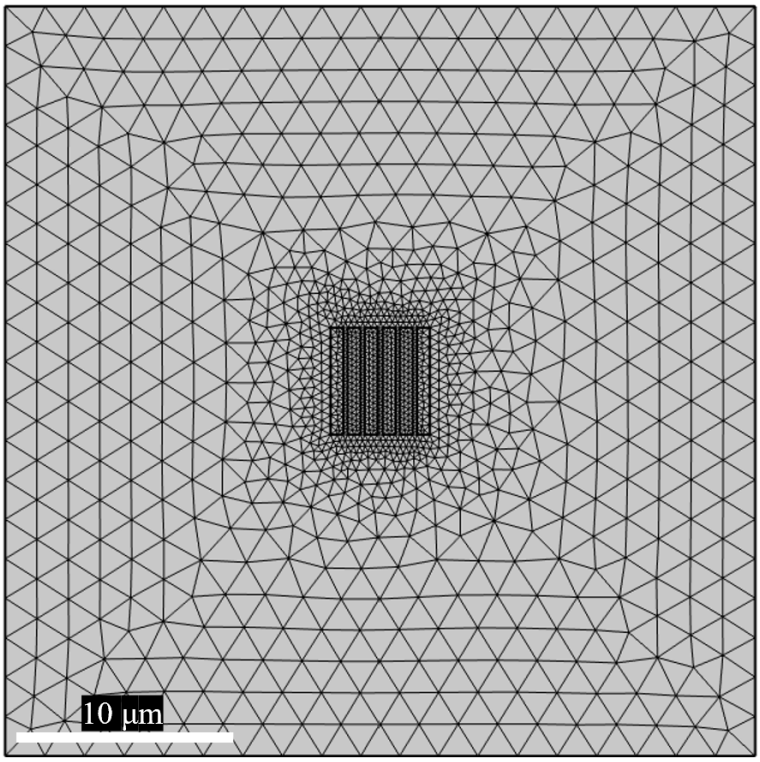

Figure 5.4: Finite element mesh used in COMSOL simulations. The mesh is refined in the guiding region to capture the confined fields accurately, while coarser elements are used in the cladding to reduce computational cost.

With geometry and mesh in place, COMSOL formulates the eigenvalue problem for the frequency-domain Helmholtz equation. Fig. 5.5 shows the refractive index distribution of a horizontal multiscan array used in the simulations, where waveguide structures are embedded in a lower-index substrate, matching the measured refractive index distribution of Fig. 5.1.

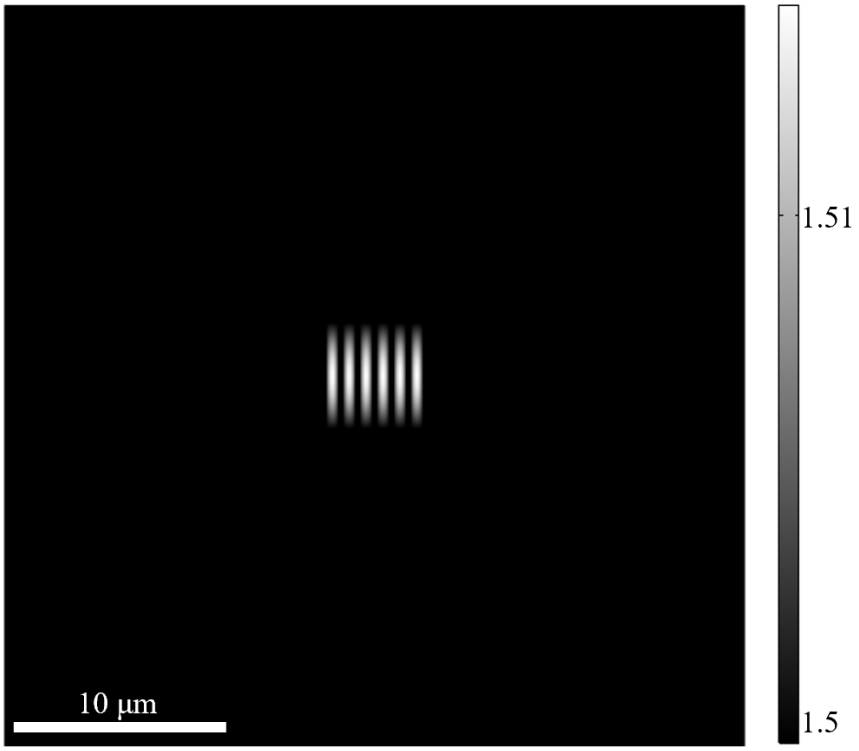

Figure 5.5: Refractive index distribution of a horizontal multiscan waveguide. Each structure is modelled with a Gaussian refractive index profile, with a peak contrast of $n = 1.514$ relative to the substrate index $n = 1.50$.

The COMSOL solver then computes the eigenvalues and eigenvectors of the system. The eigenvalues correspond to the effective indices of the supported modes, while the eigenvectors give the spatial field distributions. These outputs are inspected for confinement, symmetry, and agreement with the expected modal behaviour to be consistent with established waveguide theory. In structures with multiple modes, comparisons of effective indices reveal whether degeneracies are preserved in symmetric arrays or lifted in asymmetric ones.

Finally, the results can be post-processed to extract parameters such as mode-field diameter, effective area, or overlap integrals, providing deeper insight into the guiding properties. A central strength of the FEM approach is its ability to handle smooth Gaussian-like profiles typical of femtosecond-laser-written waveguides, as well as distorted or asymmetric geometries that cannot be described analytically.

In summary, the COMSOL workflow proceeds through parameter definition, geometry construction, meshing, eigenvalue solution, and post-processing. Each stage is essential for accurate prediction of supported modes and effective indices, establishing a direct link between structural design and optical behaviour.

5.3.2 Multiscan cross-sections

When multiple laser-written structures are placed within close proximity, their evanescent fields overlap and the system no longer supports modes confined to individual guides. Instead, the structure behaves as a single composite cross-section with collective modal properties, often referred to as supermodes. These supermodes form the basis for designing spatial mode converters in multiscan arrays.

The precise form of a supermode is determined by the arrangement, separation, and symmetry of the scans, together with the operating wavelength. With suitable design, these synthetic structures can reproduce the modal families of step-index fibres while being directly inscribed into bulk substrates. The discussion in this section indicate that a physical intuition about the mode profile behaviour can be established. Sub-micron separations produce strong overlap, so the array acts effectively as a single enlarged guide with a smooth central peak. As separation increases, the overlap weakens and the field distribution fragments into lobes centred on the individual scans. At sufficiently large distances, the overlap vanishes and each structure guides independently. Wavelength plays an analogous role: long wavelengths with broad mode fields favour collective guiding, while shorter wavelengths accentuate localisation within the individual structures.

The behaviour of supermodes is also highly sensitive to structural asymmetry. Perfectly symmetric arrays support evenly distributed fields, but even small variations in core size, position, or index contrast bias the distribution towards particular regions. Such asymmetry may arise from fabrication tolerances, yet it can also be deliberately engineered to produce tilted or shifted mode patterns. This introduces a valuable design parameter: multiscan cross-sections can be tailored not only for collective guiding, but also for

shaping or steering modes in controlled ways.

Overall, multiscan cross-sections act as synthetic waveguides whose modal properties are governed by geometry. Their design determines not only the number of supported modes, but also their extent, symmetry, and robustness. Numerical modelling, presented in the following subsections, builds on this qualitative framework to quantitatively confirm how specific cross-sectional choices translate into concrete field distributions. This predictive capability is critical for engineering spatial mode converters with precise functional outcomes.

5.3.3 Effect of separation and wavelength

One of the most critical parameters in multiscan arrays is the separation between neighbouring waveguides. Together with the operating wavelength, it determines whether the cross-section supports a smooth collective supermode or a set of independent, localised modes. To examine this systematically, simulations were carried out for three representative horizontal separations: 0.8 μm, 1.5 μm, and 3.0 μm, each at wavelengths of 1550 nm, 785 nm, and 532 nm. These values illustrate the transition from strong overlapping, to partial overlapping, to negligible overlapping of the individual modes.

At a separation of 0.8 μm, the evanescent tails of adjacent guides overlap almost completely. The simulated fundamental mode extends broadly across the array, forming a single collective distribution. At 1550 nm wavelength, the mode envelope fills nearly the entire cross-section, while at 785 nm wavelength the confinement tightens but the field remains continuous. Even at 532 nm wavelength, the field does not break into discrete lobes, confirming that such small separations produce a robust supermode spanning the whole array, as shown in Fig. 5.6.

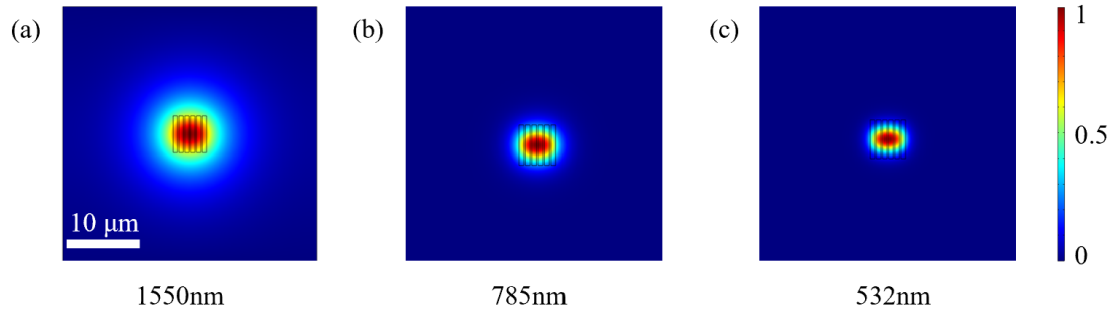

Figure 5.6: Simulated mode fields for a multiscan waveguide with a separation of 0.8 µm at three wavelengths. (a) 1550 nm, (b) 785 nm, (c) 532 nm. At all wavelengths the array supports a smooth collective fundamental mode, with tighter confinement at shorter wavelengths.

At a separation of 1.5 µm, the array enters a transitional regime. At 1550 nm wavelength the mode distribution still resembles that of a single enlarged guide, but at shorter wavelengths the influence of the individual modes becomes visible. At 785 nm wavelength, slight undulations appear in the field envelope, while at 532 nm wavelength, the lobes become clearly resolved with peaks aligned to the scan positions. The gradual transition from smooth to multi-lobed distributions can be seen in Fig. 5.7. Although not shown in the figure, the multiscan waveguide also became multimodal at shorter wavelengths.

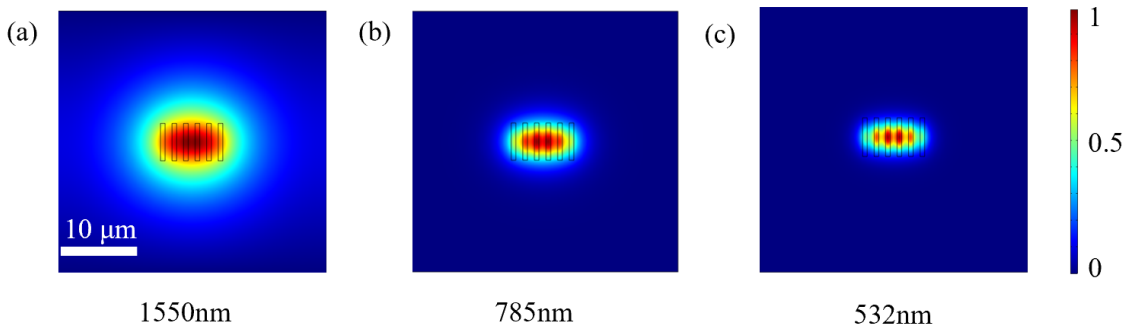

Figure 5.7: Simulated mode fields for a multiscan waveguide with a separation of 1.5 µm at three wavelengths. (a) 1550 nm, (b) 785 nm, (c) 532 nm. The field evolves from a broad collective envelope at longer wavelengths to distinct lobes at shorter wavelengths.

At a separation of 3.0 µm, the overlap between neighbouring guides is further decreased. The simulated fields remain confined to individual structures at all wavelengths. At 1550 nm wavelength, there is still faint extension into neighbouring regions, whereas

at 785 nm and 532 nm wavelength the modes are highly localised and centred on each scan, as illustrated in Fig. 5.8.

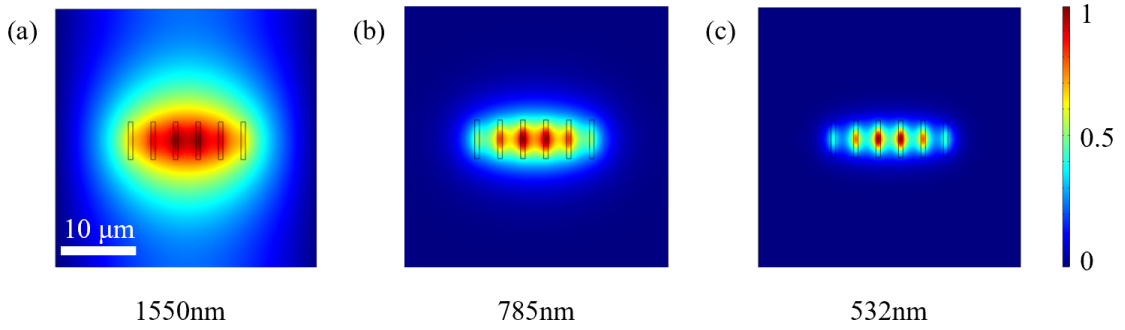

Figure 5.8: Simulated mode fields for a multiscan waveguide with a separation of 3.0 µm at three wavelengths. (a) 1550 nm, (b) 785 nm, (c) 532 nm. The fields are localised within individual structures and the collective supermode behaviour significantly diminishes.

While the field maps capture these trends qualitatively, cutline profiles provide a clearer quantitative comparison. Fig. 5.9 shows normalised intensity profiles extracted along a central horizontal line. At 0.8 µm, the profiles collapse into a single broad peak across all wavelengths, illustrating complete overlap. At 1.5 µm, the profiles show shoulders at 1550 nm and splitting at 785 nm and 532 nm, consistent with partial localisation. At 3.0 µm, the profiles display well-separated peaks, confirming independent single-core guiding.

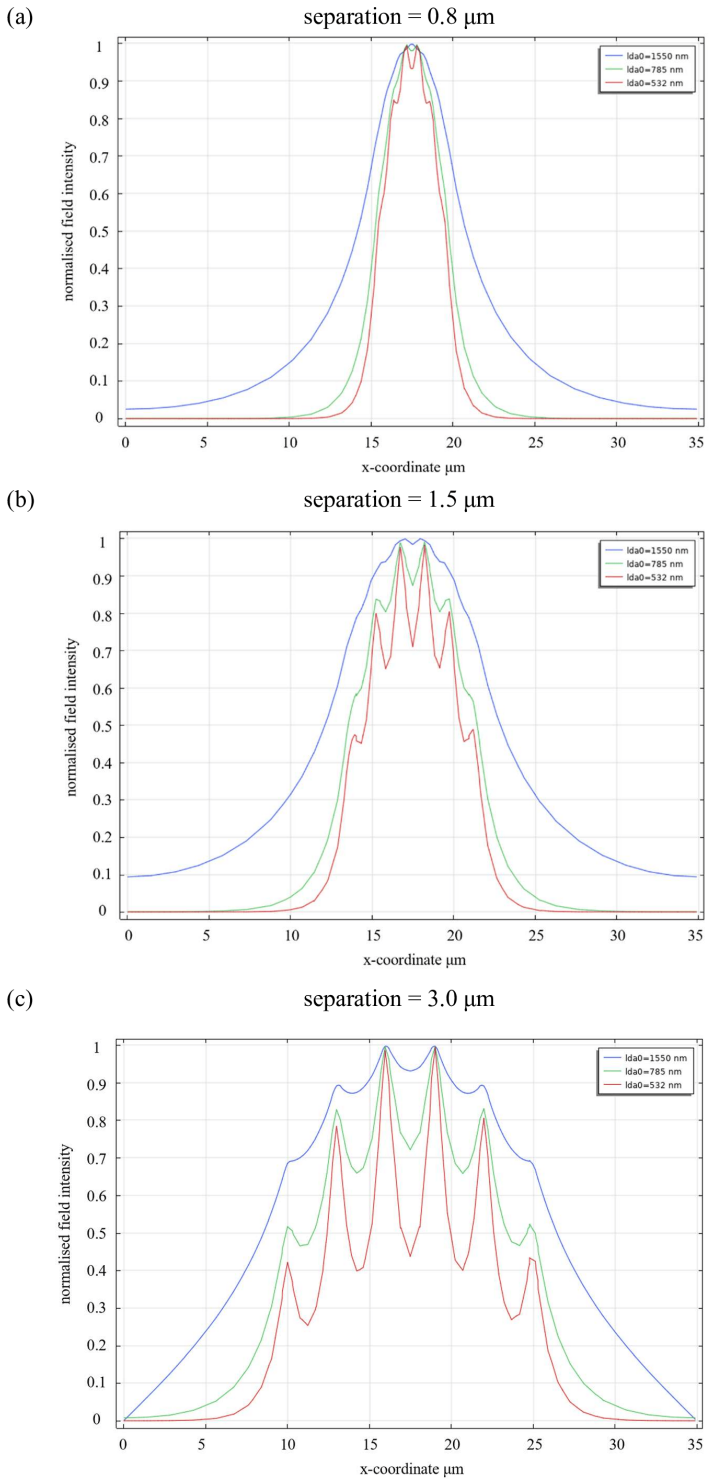

Figure 5.9: Normalised intensity profiles along a horizontal cutline for separations of (a) 0.8 µm, (b) 1.5 µm, and (c) 3.0 µm. Curves correspond to wavelengths of 1550 nm, 785 nm, and 532 nm. The transition from a broad single peak to distinct separated lobes is evident as the separation increases and the wavelength decreases.

To test whether the same principles apply in other orientations, a vertical two-column

arrangement was simulated. The refractive index profile in Fig. 5.10 show two vertical strips with higher refractive indices. The corresponding fundamental modes, shown in Fig. 5.11, demonstrate that at 1550 nm the field fills most of the elongated region, while at shorter wavelengths the confinement strengthens and the mode separates to becomes more distinctly confined in each region. This confirms that separation and wavelength interplay governs collective versus localised guiding in both horizontal and vertical geometries.

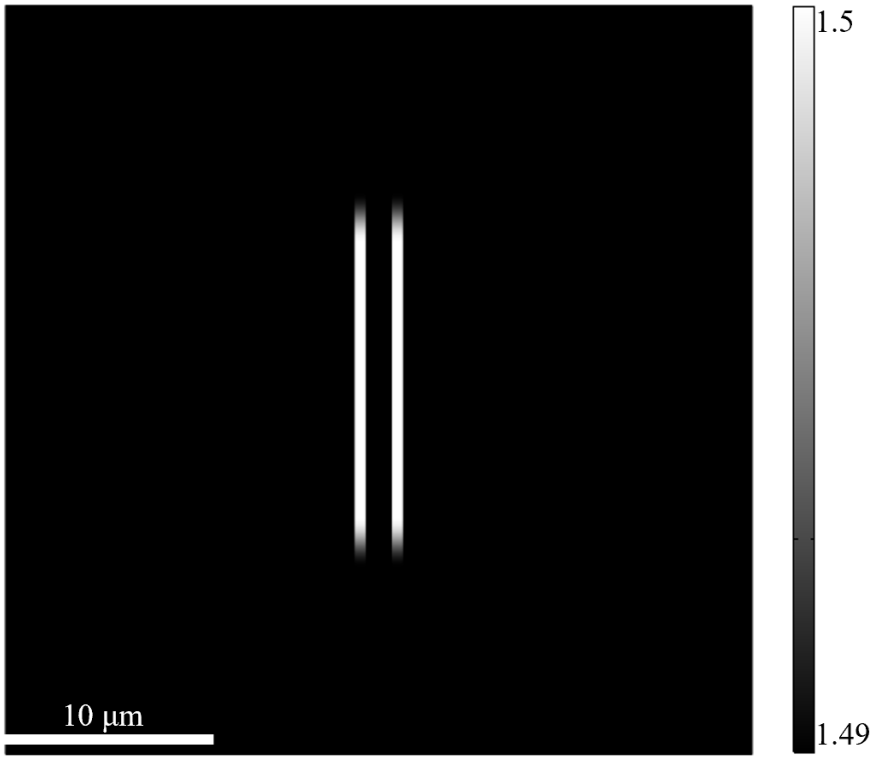

Figure 5.10: Simulated refractive index profile of a vertical two-column arrangement.

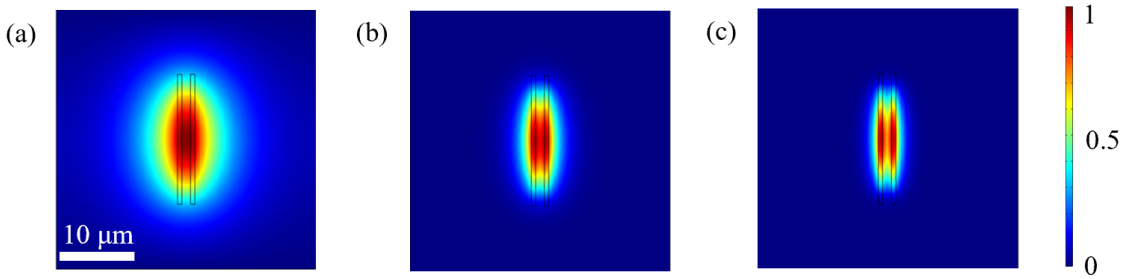

Figure 5.11: Simulated fundamental mode profiles for a vertical two-column waveguide arrangement at three wavelengths: (a) 1550 nm, (b) 785 nm, and (c) 532 nm. The vertical geometry leads to elongated modes with stronger confinement at shorter wavelengths.

These results establish a clear picture of how separation and wavelength jointly govern the transition between collective and localised guiding. Small separations favour strongly overlapping fields that combine into smooth supermodes, whereas large separations isolate the cores and suppress collective behaviour. Also, longer wavelengths enhance overlap and promote shared guiding, while shorter wavelengths emphasise confinement and localisation. Together, these factors provide a straightforward framework for anticipating the modal properties of multiscan structures and for tailoring cross-sections to achieve the desired balance between collective behaviour and independent guiding in spatial mode converter design.

5.3.4 Symmetric and asymmetric geometries

In addition to linear or rectangular arrays, multiscan structures can also be designed in circular configurations. Such arrangements are particularly attractive because they naturally support symmetric guided modes that closely resemble the fundamental mode of a standard single-mode fibre. This makes circular multiscan designs well-suited for fibre coupling, as the overlap between the guided field and the fibre mode can be maximised without additional mode-shaping elements. Fig. 5.12 shows the simulated refractive index profile of a circular SPIM-WG cross-section. It is observed that negative index region appears above the guiding structure as an artefact of the phase-modulated inscription, yet it does not disturb the mode formation within the array. The simulated field in Fig. 5.13 confirms that the structure supports a symmetric, centrally peaked fundamental mode that remains circular in shape since it is well-confined in the overall structure. The specifics and limitations of the interplay between the guiding region and negative index region can be a direction for future studies.

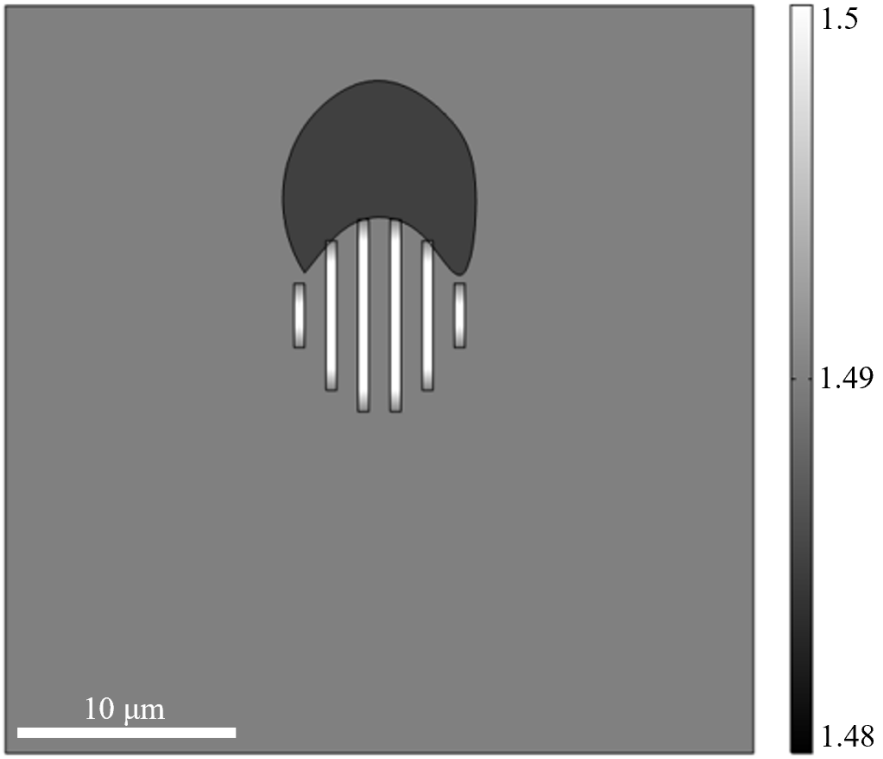

Figure 5.12: Simulated refractive index profile of a circular multiscan waveguide cross-section designed with the SPIM-WG method. A negative index region appears above the guiding structure due to artefacts of the phase-modulated inscription, but it does not perturb the formation of the symmetric mode supported by the circular array.

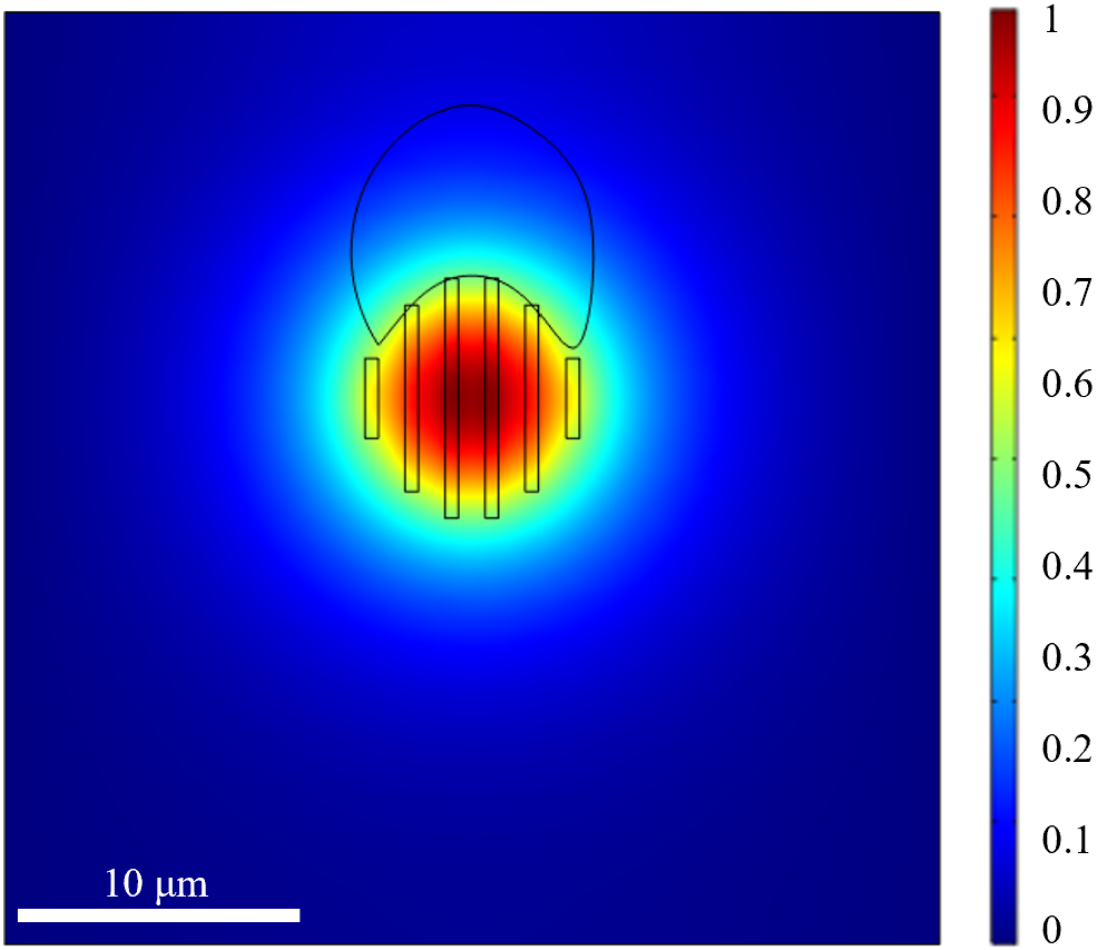

Figure 5.13: Simulated fundamental mode field of the circular multiscan waveguide corresponding to the refractive index profile shown in Fig. 5.12. The guided mode remains symmetric and well confined within the circular array of scans, despite the presence of a negative index region above the structure. The colour bar shows the normalised field intensity.

While circular symmetry provides efficient fibre coupling, many practical waveguides exhibit asymmetric mode profiles. A prominent example is the case of periodically poled potassium titanyl phosphate (ppKTP) waveguides, widely employed in nonlinear optics for frequency conversion and quantum light generation. In these devices, the guiding region is defined by a rubidium ion-exchange process that penetrates into the crystal from the surface with a gradually decaying concentration. The resulting refractive index distribution produces a skewed Gaussian-like mode that overlaps poorly with the circular mode of single-mode fibres, leading to substantial coupling losses around 70% from previous measurements. Simulations further suggest that improving the efficiency to above 80%

would already suffice to surpass the shot-noise limit in phase-sensing experiments without post-selection [89].

SPIM-WG devices fabricated in Eagle glass can be employed as intermediate converters to bridge this mismatch. By arranging the SPIM-WG scans appropriately, the refractive index profile can be engineered to emulate the asymmetry of the ppKTP guiding region. Fig. 5.14 shows one such design, where the SPIM-WG cross-section is tailored to match the skewed Gaussian-like profile of a ppKTP mode.

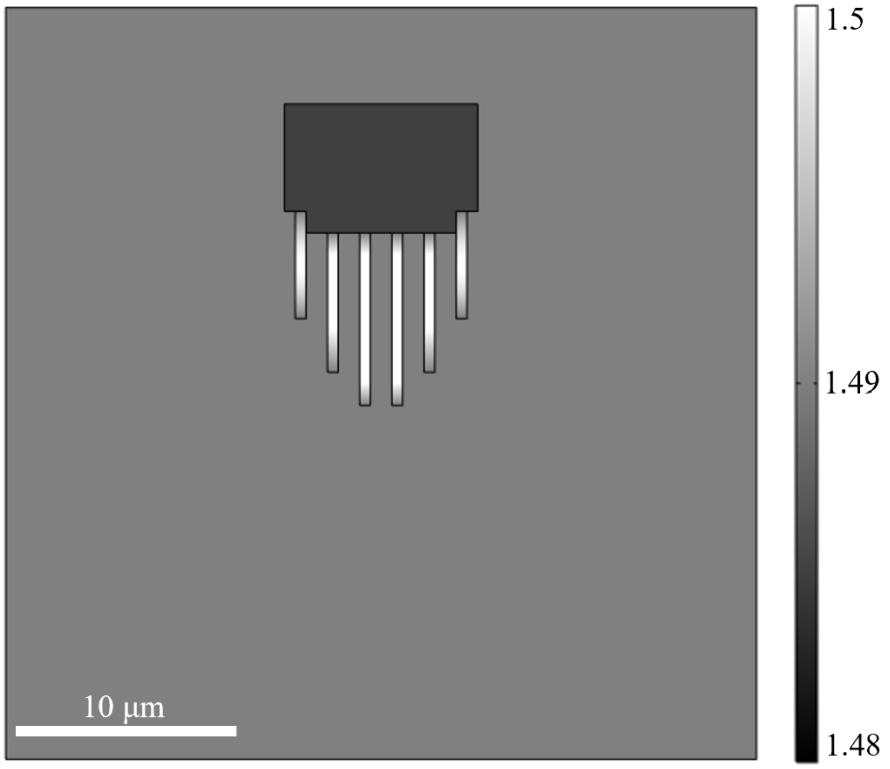

Figure 5.14: Simulated refractive index profile of a SPIM-WG cross-section designed for interfacing with ppKTP waveguides. The rectangular region at the top represents the ion-exchanged guiding layer of a ppKTP device, characterised by a skewed Gaussian-like mode. The SPIM-WG scans written in Eagle glass are arranged to transform a circular fibre mode into this asymmetric distribution. A negative index region, intrinsic to the SPIM-WG inscription process, is exploited to control the asymmetry through geometry design.

The resulting guided field is shown in Fig. 5.15. By raising the inscribed regions on the side and having the negative index region slightly extending from above in the centre, the negative index region that accompanies SPIM-WG inscription is deliberately exploited

to fine-tune the asymmetry, offering additional control over the guided field, resulting in profile where the top of the mode is flat. The simulated fundamental mode reproduces the skewed intensity distribution typical of ppKTP waveguides [90], demonstrating that SPIM-WG cross-sections can faithfully match the desired asymmetry.

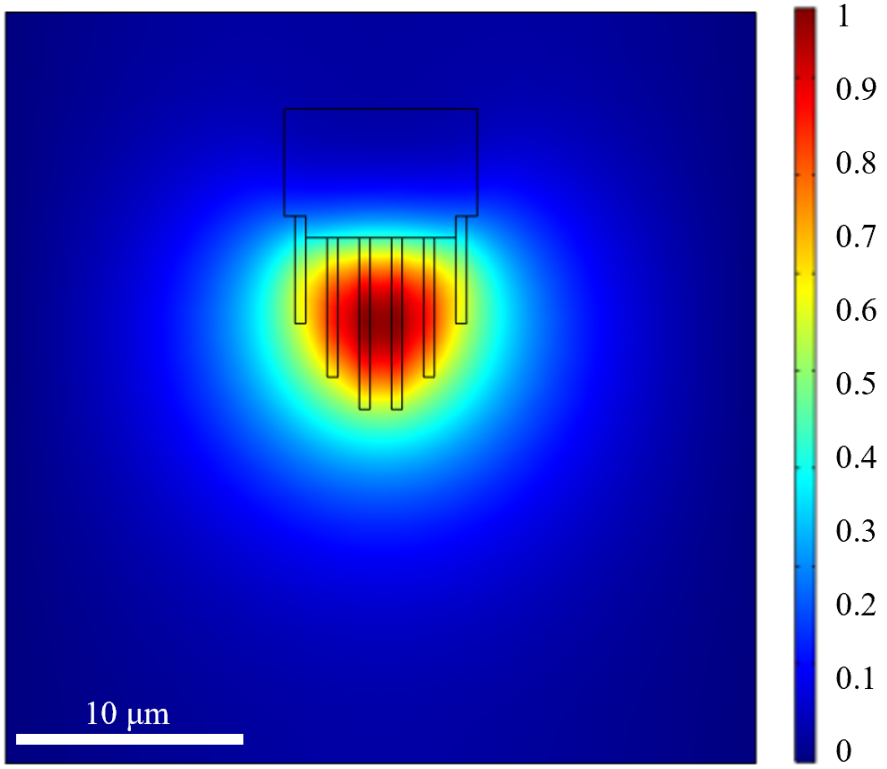

Figure 5.15: Simulated fundamental mode of the SPIM-WG cross-section shown in Fig. 5.14. The asymmetric refractive index distribution, defined by the ppKTP ion-exchanged region together with the SPIM-WG scans, guides a skewed mode profile that closely matches that of a ppKTP waveguide. The negative index region formed during SPIM-WG inscription is used to fine-tune the asymmetry, thereby improving mode matching to the ppKTP structure.

Overall, symmetric circular SPIM-WG arrays provide efficient fibre-compatible modes, while tailored asymmetric designs enable coupling to specialised waveguide structures such as ppKTP. These two cases highlight the versatility of SPIM-WG fabrication: by modifying the cross-sectional geometry, one can either preserve symmetry for fibre interfacing or deliberately introduce asymmetry for efficient coupling to nonlinear photonic devices.

5.3.5 Comparison with fabrication results

The SPIM-WG method was experimentally implemented in Eagle glass, and the fabricated structures were examined through both microscope imaging and guided-mode characterisation. Fig. 5.16 presents two representative examples: a circular array designed for efficient fibre coupling, and an asymmetric array tailored to match the mode profile of ppKTP waveguides.

In each case, the microscope image reveals the inscribed cross-section, the measured near-field pattern confirms the guided mode distribution, and the extracted intensity profiles quantify the confinement. The circular design produces a symmetric mode centred within the array, in close agreement with the simulated field shown in Fig. 5.13. By contrast, the asymmetric design yields a skewed distribution, as captured in the cutline profiles, which reproduces the characteristic ppKTP mode predicted in Fig. 5.15. Both cases also exhibit the negative index lobes intrinsic to SPIM-WG inscription, providing an additional degree of control over the modal shape.

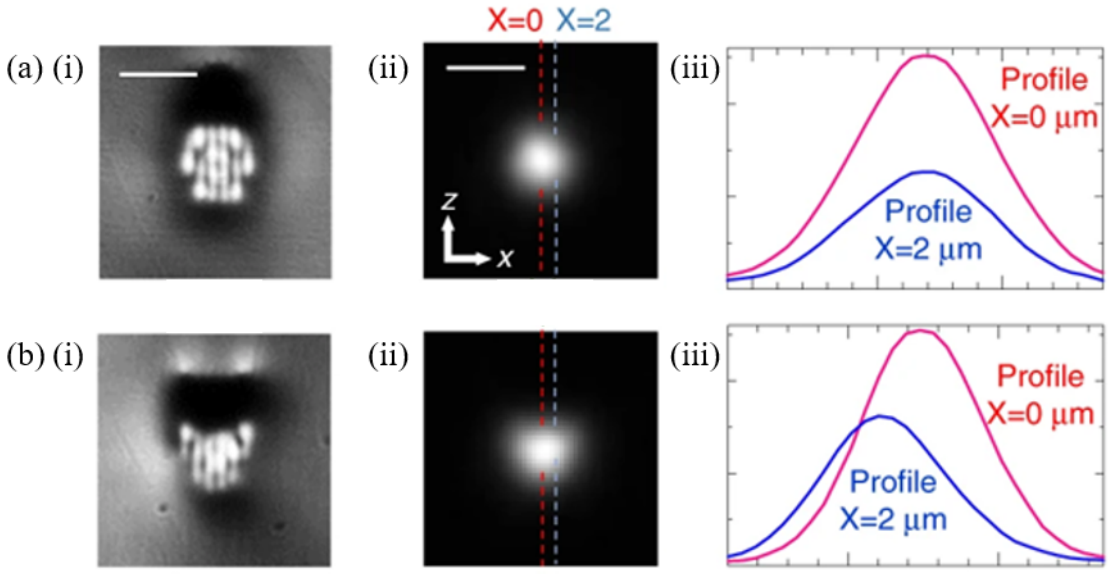

Figure 5.16: Adapted from [87]. Fabricated SPIM-WG structures in Eagle glass. For each design: (i) microscope image of the cross-section, (ii) measured near-field mode, and (iii) extracted intensity profiles along $x = 0$ and $x = 2$ μm. (a) Circular array optimised for fibre coupling, showing symmetric confinement. (b) Asymmetric array tailored for ppKTP mode matching, showing the expected skewed distribution in the cutline profiles. Scale bars: 10 µm.

By carefully shaping the core geometry and utilising the negative index region created during inscription, the mode profiles can be made to match that of external components. A spatial mode converter can then enable efficient transformation between circular fibre modes and ppKTP modes. Fig. 5.17 gives an example of using 18 waveguide scans to build a mode converter between fibre mode and ppKTP mode. The mode converter is formed by 3 regions: from the end coupling with fibre, a straight region first extends the waveguide from the end facet with the arrangement in Fig. 5.17(a), then over a sufficiently long region (usually a few mm) transitions into the arrangement in Fig. 5.17(b), where another straight region then extends to the opposite end facet and couples with a ppKTP device. Each waveguide is written with the SPIM-WG method, and the multiscan approach builds up the overall mode converter [87].

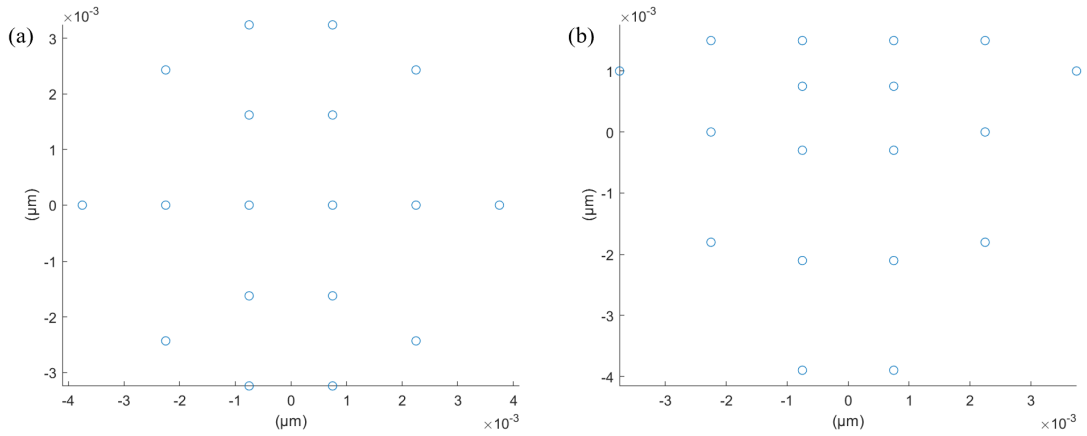

Figure 5.17: Positioning of waveguide scans at the end facet of a mode converter between fibre mode and ppKTP mode. (a) Fibre mode. (b) ppKTP mode.

The close correspondence between simulated and fabricated mode profiles highlights the reliability of the SPIM-WG approach. These results confirm that designs developed through numerical modelling can be directly realised in practice, enabling precise engineering of spatial mode converters.

5.4 Conclusion

The investigation of waveguide cross-sections for spatial mode converters highlights the influence of geometry on optical function. By adopting a supermode perspective, this chapter has shown that the guided fields of multiscan structures are inherently collective, shaped directly by the arrangement, symmetry, and separation of the scans. Numerical simulations provided a systematic framework for predicting these behaviours, while fabrication experiments in Eagle glass confirmed that SPIM-WG technology can reliably translate designed cross-sections into physical devices. Together, these results affirm that cross-sectional engineering is not merely supportive but foundational to the realisation of spatial mode converters.

Several avenues for future exploration emerge from these findings. One direction is the study of increasingly complex cross-sections, which could enable the generation of higher-order spatial modes such as those carrying orbital angular momentum, of interest for high-capacity communications and high-dimensional quantum encoding. Another important direction lies in circuit-level integration, where SPIM-WG cross-sections designed for specific transformations—such as mode rotation, splitting, or reshaping—are cascaded to form compact multifunctional photonic systems.

Equally important is the question of robustness. While this chapter has noted that asymmetries can often be harnessed as design tools, a quantitative understanding of tolerances to fabrication imperfections—such as misalignment, variation in index contrast, or beam-shaping artefacts—would be highly beneficial. Such studies would improve reproducibility and help identify which designs are best suited for scalable manufacturing.

Finally, the extension of these concepts into nonlinear and quantum photonics represents a particularly promising frontier. Tailored spatial modes could be used to optimise phase-matching conditions for nonlinear processes, or to control the spatial properties of entangled photons for quantum information processing. In this way, simulation-driven design of SPIM-WG cross-sections could become a powerful enabler of next-generation

integrated quantum technologies.

Overall, this chapter has established a clear link between cross-sectional geometry and optical functionality in SPIM-WG devices. The convergence of simulation, fabrication, and physical intuition provides a versatile framework for spatial mode engineering, with potential applications spanning communications, nonlinear optics, and quantum photonics.

5.5 Acknowledgements

I would like to acknowledge Dr. Bangshan Sun for conceiving the idea of SPIM-WGs, leading the project, and carrying out the fabrication of waveguides and devices. I also thank Dr. Alexander Jeasacher for developing the tomographic microscopy and phase optimisation routines, and for performing the corresponding measurements and analyses in Fig. 5.1, which is used as a basis for the refractive index profiles in modelling. I am grateful to the wider team of collaborators for their contributions to phase control, refractive index characterisation, and valuable discussions on fabrication regimes and device performance. The results have been published in [87].

My own contribution focused on carrying out the COMSOL simulations, which formed a central component of this work. These simulations established the link between cross-sectional design and optical behaviour, guided the fabrication strategy, and enabled direct comparison between theoretical predictions and experimental results. In this way, the modelling tied together the various strands of fabrication, optimisation, and characterisation into a coherent framework for engineering SPIM-WG devices.

CHAPTER 6

Polarisation and angle-dependent modulation of coupling ratio in directional couplers

6.1 Introduction

Manipulation and control of the polarisation of light is a fundamental aspect of photonics and has a wide range of applications, including photonic computing and communications [91]. Since waveguide-based structures are important components for building integrated optical devices and circuits, many applications in topological photonics utilise waveguide arrays and evanescent coupling to demonstrate physical phenomena [92]. On-chip quantum and data processing applications are also highly dependent on polarisation properties [93]. Recent efforts have been directed towards polarisation insensitive photonic devices [94]. While the impressive work demonstrates the ability to minimise polarisation-dependent effects, in general polarisation effects could be harnessed for further functionality. It is therefore necessary to understand how such optical devices respond to different polarisations in order to have full control on manipulating different states and preserving the quality of encoded information.

Laser-written directional couplers are investigated in this chapter as they are fundamental building blocks of optical circuits [95]. A waveguide directional coupler splits incident light into different output ports depending on the interaction length of the coupling region. It is formed by bringing two waveguides close to each other. Their function relies on evanescent coupling, which occurs when two waveguides are in close proximity

and their evanescent fields overlap [96]. In this paper, the waveguide arm coupled to the input light will be referred to as the primary waveguide, and the other arm referred to as the secondary waveguide. The splitting ratio is the proportion of transmitted power in the primary and secondary waveguide outputs with respect to the input light of a directional coupler. The ratio varies with respect to the interaction region length of a directional coupler.

Previous work on femtosecond laser direct written photonic circuits has demonstrated different approaches to utilise polarisation-dependent properties of integrated optical devices to let them act in lieu of bulk components, such as polarisation beam splitters [97], waveplates [98] and retarders [99]. However, the characterisation typically only considers vertical and horizontal linear polarisation inputs, but not for arbitrary linear polarisations. Moreover, the devices were mostly coplanar, which gives an incomplete description of device behavior when considering three-dimensional structures. Szameit et al. investigated how out of plane geometries affect coupling constants, but did not take input polarisation into account [100]. Sansoni et al. fabricated out of plane directional couplers specifically so that the coupling was the same for horizontal and vertically polarised light [101]. Nevertheless, it was focused on finding a specific operating point for a fixed interaction region length. These works have shown that the x- and y- polarised light give rise to different periods of power exchange, but characterisations of other polarisation states were unclear. Other recent work has been undertaken to utilise this capability to demonstrate more sophisticated topological and quantum phenomena but the reported results encounter similar limitations in the sense where only specific input polarisation states are considered [102, 103]. However, to fully exploit the 3D capabilities of FLDW, it is crucial to understand how these out-of-plane geometries change the properties of the devices and the transmission for different polarisation states.

Fig. 6.1 shows two configurations in which one might imagine when considering polarisation and angle-dependent coupling. There could be (a) different coupling angles with a fixed input polarisation, or (b) fixing the coupling angle and changing the polarisa-

tion instead. On first inspection one may think the two scenarios are equivalent up to a rotation, but on closer inspection one realises that the configuration actually is not symmetric for different polarisations. Consider Fig. 6.1(b), where the direction of coupling is parallel to x-polarised light but orthogonal to y-polarised light. The difference in field direction leads to a slight variation in the effective refractive indices of the fundamental modes, leading to polarisation-dependent effects. This suggests that a more careful analysis is needed to properly examine the different configurations.

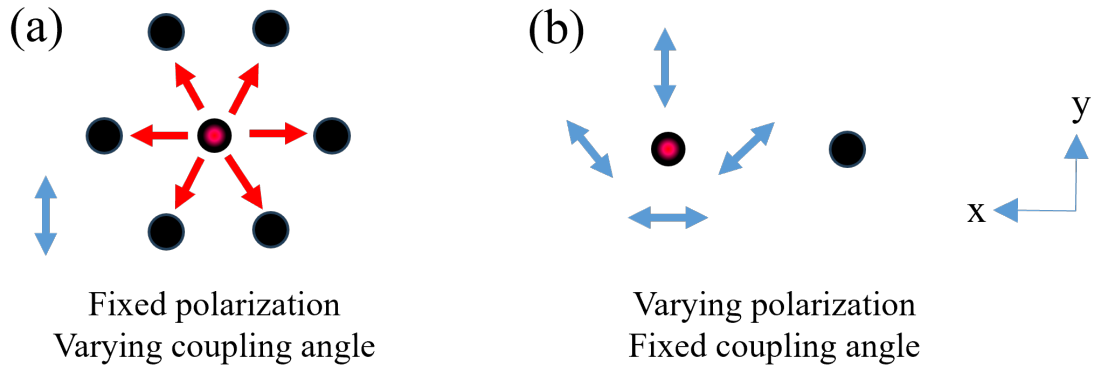

Figure 6.1: (a) Geometrical configuration where light is guided in the centre waveguide couples into different waveguides at various angles. (b) Configuration where the coupling angle is fixed but the guided polarisation is varying. (Red arrows denote the coupling directions and blue arrows indicate the polarisation field directions.)

This chapter aims to present a general polarisation-dependent effect where the orientation of linearly polarised incident light dictates the splitting ratio for each output port of a directional coupler and can modulate the maximum value attainable. Fig. 6.2 shows the general setup of the orientation of the input polarisation angle θ with respect to the y-direction (vertical) of the directional couplers. Note that the red dashed line indicates the interaction region for evanescent coupling of the two arms of a directional coupler.

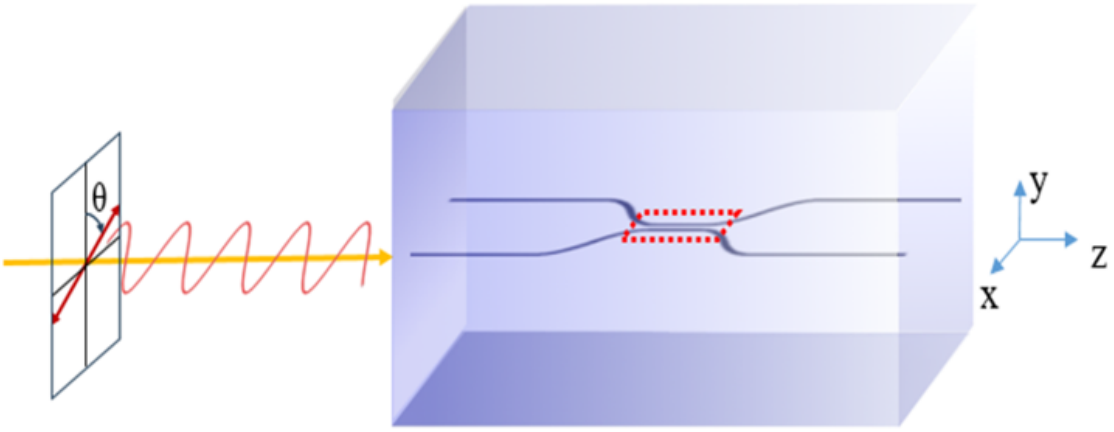

Figure 6.2: Schematic showing the input light polarisation angle θ with respect to the vertical y-direction of the directional couplers.

The polarisation effects on splitting ratio as a function of the interaction region length is experimentally characterised. This enables further understanding of how the performance of directional coupler devices depends on polarisation. The dependence of the splitting ratio on the polarisation orientation angle and relative position of the waveguide arms is studied theoretically, predicting an additional modulation effect which is in agreement with the experimental data. This gives insights into how the fabrication and circuit design can be controlled in order to improve the performance of devices by better understanding the polarisation effects.

Moreover, it is further demonstrated that the effect in directional couplers extended to out-of-plane geometries with three-dimensional structures, where the two arms are not at the same depth from the substrate surface. The experimental results are in agreement with the theoretical description of the coupling for monochromatic light. Fig. 6.3 shows the waveguides being offset at an angle which is denoted as ϕ, which is used when studying the effect on structures that are not coplanar.

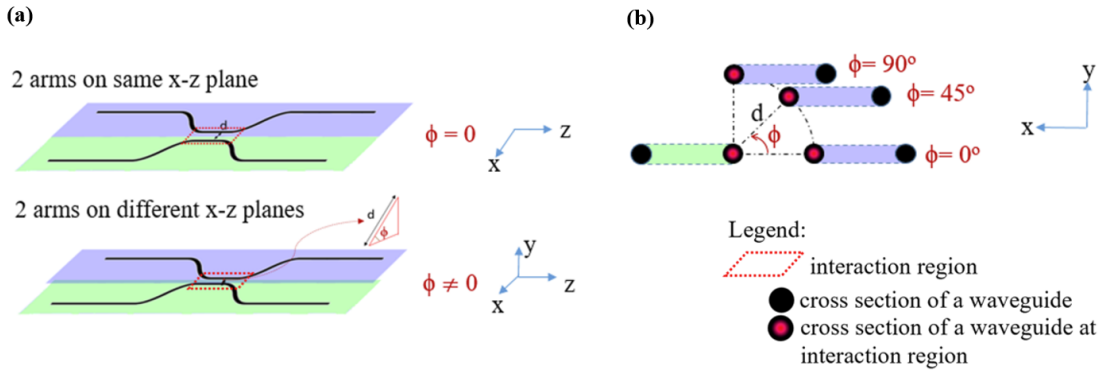

Figure 6.3: (a) 3D view of two waveguides lying on same or different x-z planes with a constant separation distance d slightly offset at an angle ϕ at the interaction region – the upper diagram shows when $\phi = 0$ and the lower when $\phi \neq 0$. (b) 2D cross-section view on x-y plane of two waveguides at different ϕ at the interaction region.

Below, Section 6.2 describes the background and theoretical study. Section 6.3 describes details of the experimental design and the directional coupler fabrication parameters. Section 6.4 reports and discusses the experimental results and the results of extension to 3D structures. Section 6.5 gives the concluding remarks.

6.2 Properties of directional couplers

The splitting ratio r indicates the proportions of transmitted power in the primary and secondary waveguide outputs with respect to the input light of a directional coupler. According to coupled-mode theory [104], the expression for r is given by

$$r = \frac{P_2}{P_1 + P_2} = \sigma^2 \sin^2\left(\frac{C}{\sigma} L_0\right) \tag{6.1}$$

where P_1 and P_2 are the output power of the primary and secondary waveguides respectively, C is the coupling coefficient between the two waveguides, σ is a dephasing term that depends on the waveguide asymmetry and L_0 is the length of the interaction region.

The dephasing term σ is given by

$$\sigma = \frac{1}{\sqrt{1 + \left(\frac{\Delta\beta}{2C}\right)^2}} \tag{6.2}$$

where $\Delta\beta = |\beta_1 - \beta_2|$ is the difference of the propagation constants of the two waveguides [105]. Therefore, if there is some general asymmetry in the waveguide refractive index distribution, the maximum power coupling ratio will be reduced to a fraction of the original. Nevertheless, the expression predicts a sinusoidal behavior with respect to the interaction region length. However, Eq. 6.1 does not fully consider polarisation effects and asymmetry for different input polarisation states.

When the input polarisation is aligned with the principal axes of symmetry of the cross-section geometry (along x- and y-axis for $\phi = 0°$ coplanar directional couplers as illustrated in Fig. 6.2), the light propagates through the arms with a specific mode, which can be viewed as a superposition of symmetric and anti-symmetric modes as shown in Fig. 6.4 [106]. Power exchange between the primary and secondary arms can then be explained in terms of the beating between the symmetric and anti-symmetric modes. However, when the input polarisation is not aligned along either of these principal axes, both x-polarised and y-polarised modes are present, and the situation is more complicated. The x/y symmetric/anti-symmetric modes all have different propagation constants, therefore these four modes will all contribute differently to the total propagating wave.

Assume the two waveguides lie in the xz-plane, where z is the direction of light propagation through the waveguides (Fig. 6.2). Birefringence caused by non-ideal device fabrication by FLDW, which often produces waveguides with non-circular cross-sections [107], leads to non-identical effective refractive indices of the primary and secondary waveguides for different input polarisation states. Furthermore, the geometric configuration of the two waveguides leads to further asymmetry, particularly in the strain field, such that the propagation constants of the symmetric mode and the anti-symmetric mode of the x- and y-polarised light are therefore also different [108]. This difference leads to

a change in the splitting ratio when the input polarisation changes and is inherent to the geometry of the directional couplers.

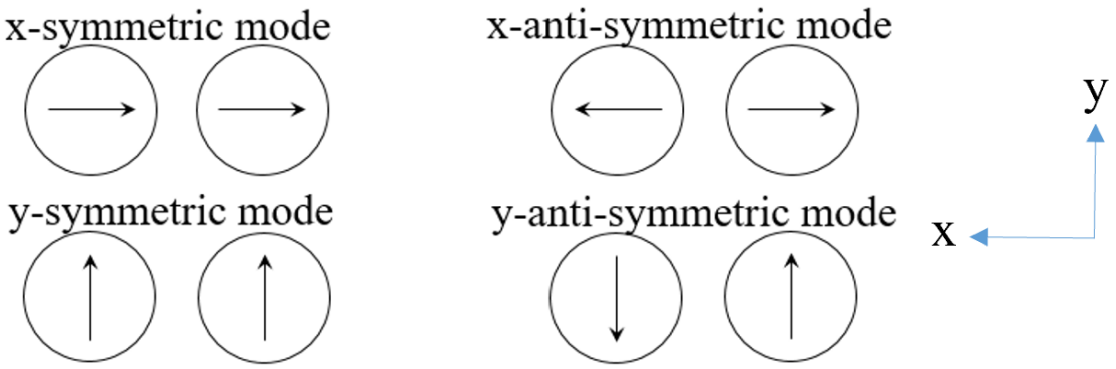

Figure 6.4: Fundamental supermodes of a directional coupler in the interaction region. Each circle represents the cross-section of a waveguide. Arrows indicate the orientation and relative phase of the transverse electric field of each supermode. Horizontal and vertical arrows correspond to x- and y-polarized supermodes, respectively. Symmetric supermodes have fields in phase across the two waveguides, while antisymmetric supermodes have fields out of phase.

Let the propagation constants of the x-symmetric (even) mode be denoted as k_{xe} and that of the x-anti-symmetric (odd) mode be denoted as k_{xo}. Similarly, the propagation constants for symmetric mode and anti-symmetric mode of the y-polarised modes are k_{ye} and k_{yo} respectively. Denote C/σ in Eq. 6.1 as κ, where the effective coupling coefficients κ_x and κ_y for the respective polarisations are related to the rate of power exchange between two waveguides.

Theoretically, the input light can be expressed as a sum of the symmetric and anti-symmetric modes with equal amplitude, and the beating between the modes creates a modulated envelope which manifests as the power exchange between the two arms of the directional coupler. The rate of power exchange is slightly different for polarised input light along the two principal axes, hence there will be an additional modulation to the splitting ratio for other polarisations. For linearly polarised light with the angle θ between the polarisation axis and the vertical y-direction (Fig. 6.2), the modulation is most pronounced at $\theta = \pi/4$. The derivation of this result is shown below.

6.2.1 Derivation for pure linearly x-polarised and y-polarised input light

Let $E_{xe}(z)$ represent the electric field amplitude for the propagating electric field corresponding to the x-symmetric mode as a function of z, omitting the time-varying component which is assumed to be the same for all terms. Similarly, for the other modes

$$\begin{aligned} E_{xe}(z) &= \cos(k_{xe}z), \quad E_{xo}(z) = \cos(k_{xo}z), \\ E_{ye}(z) &= \cos(k_{ye}z), \quad E_{yo}(z) = \cos(k_{yo}z) \end{aligned} \tag{6.3}$$

where z is along the propagation direction. Fig. 6.5 shows an example schematic of the field amplitudes of the symmetric and anti-symmetric modes in the two waveguides where the input light is coupled to WG1.

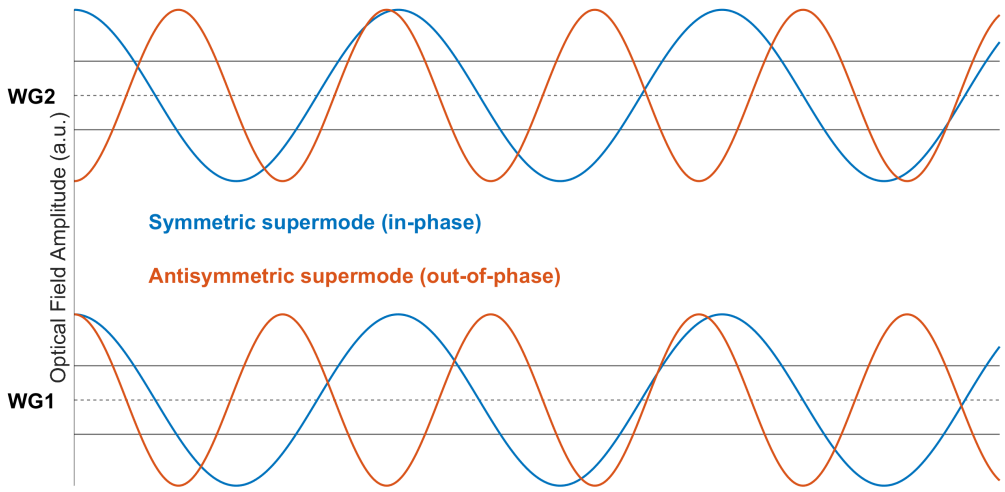

Figure 6.5: Symmetric (in-phase) and antisymmetric (out-of-phase) optical supermodes shown along the propagation direction. Horizontal lines indicate the positions of waveguides WG1 and WG2, while the curves illustrate the field amplitudes and their relative phase and sign in the respective waveguides; vertical scaling is schematic.

Consider input light polarised along the y-direction coupled to the primary waveguide arm. This input light can be expressed as a sum of the symmetric and anti-symmetric modes with equal amplitude. For unit amplitude input, at $z = 0$ we have $P_1 = 1$ and $P_2 = 0$. Assuming losses by absorption and scattering are negligible, we have $P_1 + P_2 = 1$

for all z since power is conserved. For simplicity, in the following, the power expressions P_x, P_y and P_θ refer to the output power of the x-polarised, y-polarised and linearly polarised input light with polarisation axis at an angle θ to the vertical y-axis. They represent power in the primary waveguide arm i.e. P_1 in Eq. 6.1. Note that P_2 in the secondary waveguide arm is therefore simply $1–P_1$. Fig. 6.6 shows an example plot of power beating in a directional coupler. The power variation follows a sinusoid as described in Eq. 6.1.

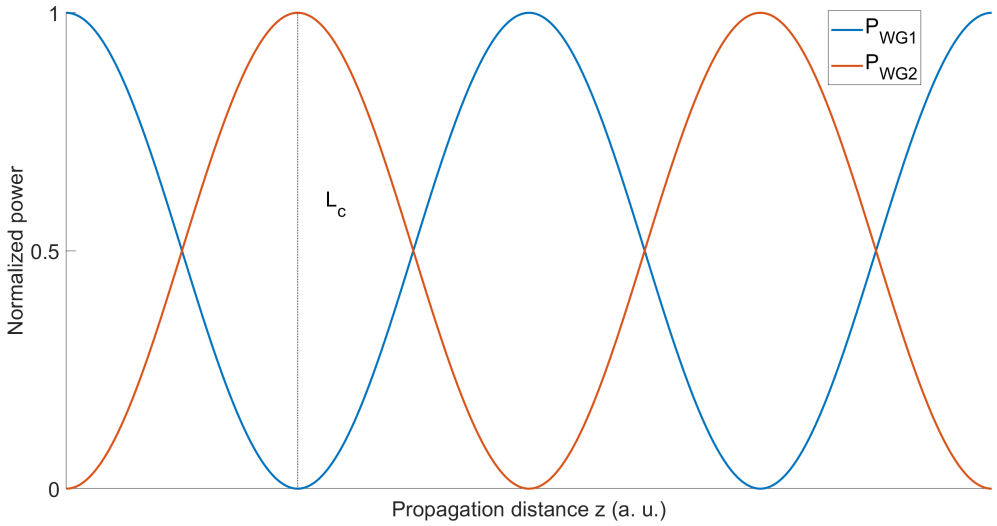

Figure 6.6: Normalized optical power in waveguides WG1 and WG2 versus propagation distance, showing periodic energy exchange due to interference between the two supermodes. The coupling length L_c marks the length required for the first complete power transfer from WG1 to WG2.

Let the wave amplitude in the primary waveguide corresponding to the y-polarised input light be E_y. Then

$$\begin{aligned} E_y(z) &= \frac{1}{2}\left(E_{ye}(z) + E_{yo}(z)\right) \\ &= \frac{1}{2}\left[\cos\left(k_{ye}z\right) + \cos\left(k_{yo}z\right)\right] \\ &= \cos\left(\frac{k_{ye}+k_{yo}}{2}z\right)\cos\left(\frac{k_{ye}-k_{yo}}{2}z\right). \end{aligned} \tag{6.4}$$

As k_{ye} and k_{yo} have similar but unequal values, the second cosine term is of much

lower frequency than the first one. The expression therefore represents a modulated sinusoid with an envelope corresponding to the second term. The power flow in the waveguide is obtained by integrating the wave amplitude, therefore it is proportional to the square of the envelope. The power P_y corresponding to the y-polarised input with unit amplitude is given by:

$$P_y(z) = \cos^2\left(\frac{k_{ye} - k_{yo}}{2} z\right). \tag{6.5}$$

Similarly,

$$E_x(z) = \cos\left(\frac{k_{xe} + k_{xo}}{2} z\right) \cos\left(\frac{k_{xe} - k_{xo}}{2} z\right) \tag{6.6}$$

and

$$P_x(z) = \cos^2\left(\frac{k_{xe} - k_{xo}}{2} z\right). \tag{6.7}$$

The envelope shows how the power flow in the primary arm changes with distance. Define quantities κ_x and κ_y as follows:

$$\begin{aligned} \kappa_x &= k_{xe} - k_{xo} \\ \kappa_y &= k_{ye} - k_{yo} \end{aligned} \tag{6.8}$$

which is simply the inverse of the beat length between the two modes [109]. They can be interpreted as the effective coupling coefficients for the respective polarisations which manifests as the power exchange between the two arms of the directional coupler.

6.2.2 Derivation for arbitrary linearly polarised input light

Now consider an arbitrary linearly polarised light with polarisation axis along θ coupled to the primary waveguide arm. Interference effects between modes (e.g. between k_{xe} and k_{ye}) are insignificant compared to the power exchange in directional couplers [109]. The overall power can therefore be found by decomposing the polarisation vector into x- and y- components. Since the wave vector corresponding to this input is a vector sum of E_x and E_y components in the amplitude domain, for input light polarised along θ with unit

amplitude we have:

$$\begin{aligned} P_\theta\left(z=L_0\right) &= P_x\left(z=L_0\right)\cos^2\theta + P_y\left(z=L_0\right)\sin^2\theta \\ &= \cos^2\left(\frac{\kappa_x}{2}L_0\right)\cos^2\theta + \cos^2\left(\frac{\kappa_y}{2}L_0\right)\sin^2\theta \end{aligned} \tag{6.9}$$

where the output power is measured at the output of a directional coupler with interaction distance L_0.

Since the effective coupling coefficients of E_x and E_y are different, there will also be beating between the envelopes of the two waves. It is possible to express this envelope analytically as follows: starting from Eq. 6.9,

$$\begin{aligned} & P_\theta\left(z=L_0\right) \\ =& \cos^2\left(\frac{\kappa_x}{2}L_0\right)\cos^2\theta + \cos^2\left(\frac{\kappa_y}{2}L_0\right)\sin^2\theta \\ =& \frac{1+\cos\left(\kappa_x L_0\right)}{2}\cos^2\theta + \frac{1+\cos\left(\kappa_y L_0\right)}{2}\sin^2\theta \\ =& \frac{1}{2}\left[1+\cos\left(\kappa_x L_0\right)\cos^2\theta + \cos\left(\kappa_y L_0\right)\sin^2\theta\right] \\ =& \frac{1}{2}\left[1+\frac{1}{2}\left(\cos\left(\kappa_x L_0\right)+\cos\left(\kappa_y L_0\right)\right)+\left(\cos^2\theta-\frac{1}{2}\right)\left(\cos\left(\kappa_x L_0\right)-\cos\left(\kappa_y L_0\right)\right)\right] \\ =& \frac{1}{2}\left[1+\cos\left(\frac{\kappa_x+\kappa_y}{2}L_0\right)\cos\left(\frac{\kappa_x-\kappa_y}{2}L_0\right)-\cos 2\theta\sin\left(\frac{\kappa_x+\kappa_y}{2}L_0\right)\sin\left(\frac{\kappa_x-\kappa_y}{2}L_0\right)\right] \end{aligned} \tag{6.10}$$

which is of the form $p\cos(t)+q\sin(t)$. Invoking the auxiliary angle relation

$$p\cos(t)-q\sin(t)=\sqrt{p^2+q^2}\cos(t+\phi),$$

and letting

$$\begin{aligned} p &= \cos\left(\frac{\kappa_x-\kappa_y}{2}L_0\right), \\ q &= \cos 2\theta\sin\left(\frac{\kappa_x-\kappa_y}{2}L_0\right), \\ t &= \frac{\kappa_x+\kappa_y}{2}L_0, \end{aligned}$$

Eq. 6.10 can be expressed as

$$\begin{aligned} &P_\theta\left(z=L_0\right) \\ =&\frac{1}{2}\left[1+\sqrt{\cos^2\left(\frac{\kappa_x-\kappa_y}{2}L_0\right)+\cos^2 2\theta\sin^2\left(\frac{\kappa_x-\kappa_y}{2}L_0\right)}\cos\left(\frac{\kappa_x+\kappa_y}{2}L_0+\phi\right)\right] \end{aligned} \tag{6.11}$$

where

$$\phi=\begin{cases}\arccos\frac{p}{\sqrt{p^2+q^2}}, & q\geq 0\\ -\arccos\frac{p}{\sqrt{p^2+q^2}}, & q<0\end{cases}.$$

Although the expression is complicated, the square-root term corresponding to the slow-varying component represents the modulation envelope, whereas the fast-varying component, which corresponds to the power exchange in the directional coupler, has a varying phase component which suggests that its period is also not constant.

To give a simpler example demonstrating this behavior, consider $\theta=\pi/4$ where the beating is most pronounced when E_x and E_y have equal amplitudes:

$$\begin{aligned} P_{\frac{\pi}{4}}\left(z=L_0\right)&=P_x\cos^2\left(\frac{\pi}{4}\right)+P_y\sin^2\left(\frac{\pi}{4}\right)\\ &=\frac{1}{2}\left(P_x+P_y\right)\\ &=\frac{1}{2}\left[\cos^2\left(\frac{\kappa_x}{2}L_0\right)+\cos^2\left(\frac{\kappa_y}{2}L_0\right)\right]\\ &=\frac{1}{2}\left[\frac{1+\cos\left(\kappa_x L_0\right)}{2}+\frac{1+\cos\left(\kappa_y L_0\right)}{2}\right]\\ &=\frac{1}{2}\left[1+\cos\left(\frac{\kappa_x+\kappa_y}{2}L_0\right)\cos\left(\frac{\kappa_x-\kappa_y}{2}L_0\right)\right] \end{aligned} \tag{6.12}$$

The second cosine term shows an additional modulation to the sinusoidal variation of the splitting ratio with increasing interaction region length. This sinusoidal variation with wavenumber $(\kappa_x+\kappa_y)/2$ has a period between that of pure x- and y- polarised input, and the modulation envelope has a wavenumber proportional to $(\kappa_x-\kappa_y)/2$. This additional modulation is important to consider when using integrated photonic circuits comprising

directional couplers with light of variable input polarisation state.

Fig. 6.7 shows an example plot of normalised power in waveguide 2 as a function of interaction length z for different input polarisations. The blue and orange curves correspond to purely x- and y-polarised inputs, respectively, illustrating that the two polarisations have different coupling lengths $L_{c,x}$ and $L_{c,y}$ (indicated by the vertical dotted lines). The green dashed curve shows the response for a $45°$ linearly polarised input, which is a superposition of the contributions from the two polarisations and therefore exhibits a beating envelope that reflects the difference between the coupling strengths for x- and y- polarisations.

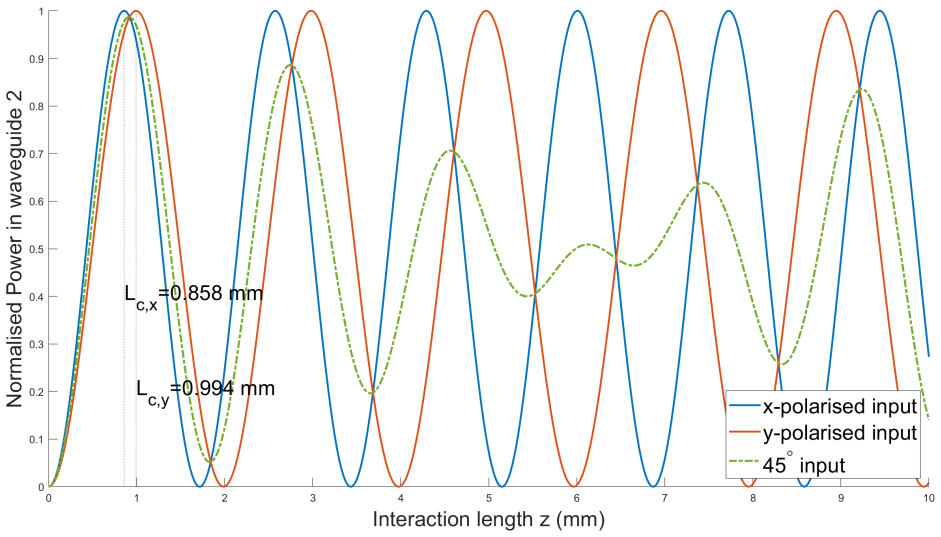

Figure 6.7: Normalised power in waveguide 2 as a function of interaction length z for x-polarised, y-polarised and $45°$ input states, illustrating the different coupling lengths $L_{c,x}$ and $L_{c,y}$ and the resulting beating for the mixed-polarisation case. (κ_x = 1.83 rad/mm, κ_y = 1.58 rad/mm.)

6.2.3 Output polarisation state

Beyond the redistribution of optical power between waveguides, the directional coupler also provides a platform in which the polarisation state itself can evolve in a non-trivial manner. Polarisation-dependent coupling and phase accumulation gives rise to changes in the output polarisation state [110]. At a high level, this behaviour can be understood as the combined effect of amplitude evolution driven by evanescent coupling and phase evo-

lution arising from waveguide birefringence. While the coupling interaction governs how optical power is exchanged between the two waveguides, birefringence controls the relative phase between orthogonal polarisation components and therefore determines whether linear input states remain linear or evolve into elliptical ones.

The polarisation behaviour can be represented using trajectories on the Poincaré sphere, which provides a geometric visualisation of the polarisation state. For a fixed input polarisation (linear at 45°), the polarisation state at each propagation position is evaluated as a function of propagation through the interaction region, accounting for polarisation-dependent coupling and waveguide birefringence. The corresponding Stokes parameters are calculated and normalised, yielding a point on the surface of the Poincaré sphere. Successive points trace out a trajectory that represents the evolution of the polarisation state along the coupler up to the output of the coupler. In this representation, trajectories confined to the equatorial plane correspond to purely linear polarisation states, while deviations away from the equator indicate the generation of elliptical polarisation.

As shown in Fig. 6.8, the birefringence-free case results in a trajectory confined to the equator of the Poincaré sphere, demonstrating that polarisation-dependent coupling alone modifies only the relative amplitudes of the orthogonal components and does not generate ellipticity. Introducing a finite birefringence leads to a progressive deviation of the trajectory away from the equator, reflecting the accumulation of a relative phase between the polarisation components and the resulting evolution toward elliptical polarisation. Increasing the magnitude of the birefringence enhances this effect, while the overall shape of the trajectory remains governed by the coupling interaction.

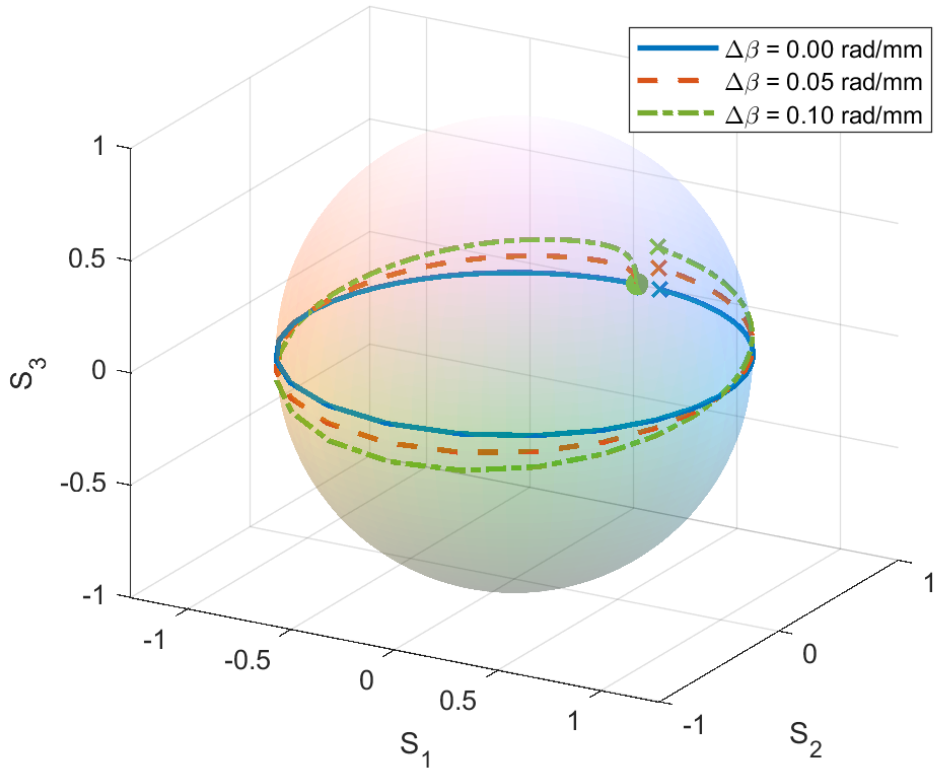

Figure 6.8: Poincaré-sphere trajectories illustrating the effect of waveguide birefringence on the output polarisation of a directional coupler for an input linear 45° polarisation. The solid curve corresponds to the birefringence-free case ($\Delta\beta = 0$ rad/mm), for which the polarisation evolution remains confined to the equatorial plane. Dashed curves show increasing birefringence ($\Delta\beta = 0.05$ and 0.10 rad/mm), which introduce a relative phase between the orthogonal polarisation components and lead to the generation of elliptical polarisation. Markers indicate the output polarisation at the end of the interaction region ($z = 2$ mm).

A full characterisation of the polarisation response of such couplers would require independent measurements of the birefringence of the individual waveguides, together with a detailed analysis of its interplay with polarisation-dependent coupling. The figure should therefore be regarded as an illustrative example that highlights the qualitative role of birefringence in the coupler, rather than as a complete description of a specific device, which lies beyond the scope of the present study. Nevertheless, the results indicate that directional couplers can act as versatile elements for manipulating polarisation on chip. By appropriately engineering both coupling geometry and waveguide birefringence, such structures could be exploited as compact polarisation transformers, enabling controlled

polarisation rotation or the generation of elliptical and circular polarisation states. This capability is of particular interest for polarisation-encoded optical communications and integrated quantum photonics, where precise control of multiple degrees of freedom of light is essential.

Finally, there exist interaction lengths for which the splitting ratio is independent of the input polarisation. Polarisation independence occurs when P_θ is independent of θ. From Eq. 6.11, it can be observed that this occurs when

$$L_0 = \frac{2n\pi}{|\kappa_x - \kappa_y|} \tag{6.13}$$

where n is any positive integer. This suggests that there is always some L_0 for which the splitting ratio is polarisation independent. For practical considerations, L_0 should not be exceedingly large, therefore the difference between κ_x and κ_y cannot be too small.

6.3 Experimental setup

6.3.1 Fabrication design and parameters

The waveguides were written with a single scan at 150 μm depth from the surface, 8 mm/s scanning speed, and 110 nJ pulse energy measured at the objective pupil. These parameters were found by optimising for the best transmission performance (minimising propagation and coupling loss) as detailed in previous chapters.

After completion of the waveguide fabrication, the samples were polished by using a sequence of 30 μm, 9 μm, 3 μm and 1 μm polishing films. A layer of at least 150 μm glass was polished off both the input and output facets of the chip.

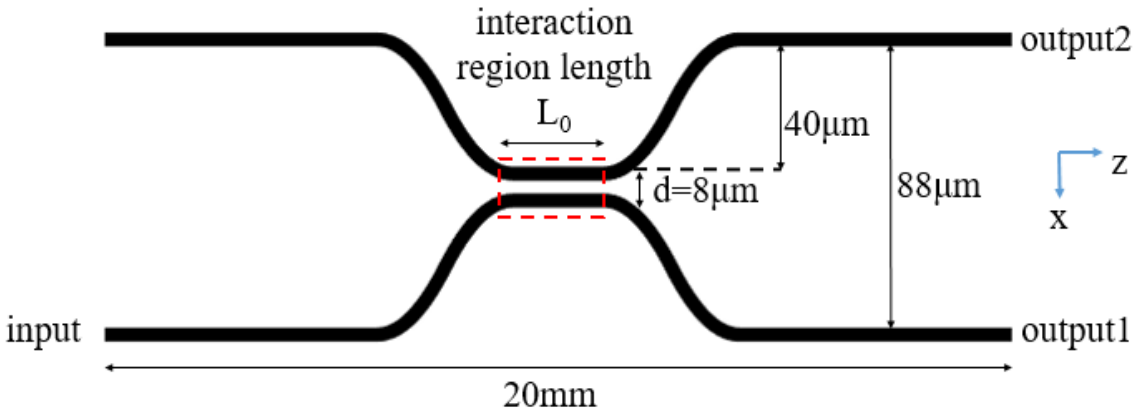

Figure 6.9: Schematic of the design of a fabricated directional coupler.

Fig. 6.9 shows a schematic for the design of fabricated directional couplers. Each directional coupler was created by fabricating two waveguides separated by 88 μm at the input and output facets, then brought together in the middle of the sample with 8 μm separation in the interaction region. The S-Bends at either side of the interaction region yield a 0.04 mm shift in the x-direction over a 4 mm length. This corresponds to each section having a radius of curvature of 100 mm, which is chosen as sufficiently large to minimise bend losses.

Multiple sets of directional couplers were fabricated with different interaction region lengths from 2 mm to 10 mm in 0.5 mm steps to investigate variation of splitting ratio between two arms when polarisation properties of input light were changed.

6.3.2 Characterisation

For directional coupler characterisation, the fibre was coupled to one of the input waveguide arms. The splitting ratio of each directional coupler was then measured by integrating over each output using pixel values obtained from the images captured with the CCD camera. Fig. 6.10(a) shows the output facet of the primary and secondary waveguides under LED illumination. Fig. 6.10(b) shows a sample image of the mode field for light output from the directional coupler device captured with the CCD camera. The mode field is slightly elliptical and can be approximated by a Gaussian. The $1/e^2$ diameter is

6.12 μm in the x-direction and 6.82 μm in the y-direction.

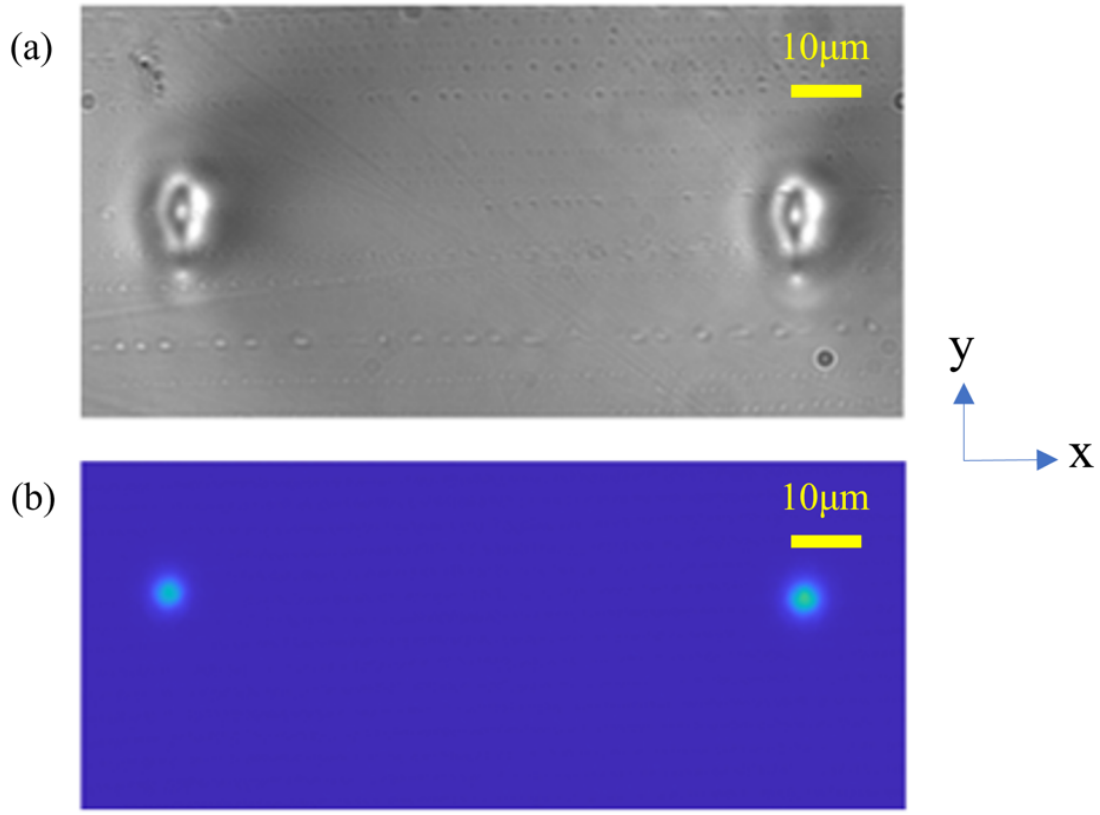

Figure 6.10: View of experiment measurements at the output end of a directional coupler: (a) Output facet of primary and secondary waveguides under widefield microscope showing the cross-section of fabricated waveguides which are non-circular and elongated along the y-direction. (b) Microscopic image of laser light guiding modes at the two outputs.

6.4 Experimental results

Measurements of the variation of splitting ratio against interaction region length when the input light had a linear input polarisation at different angles θ to the vertical y-axis were obtained by integrating over the pixel intensities of the CCD images captured. It is observed that the coupling variation had a maximum period when the input polarisation θ was at 0° and a minimum at 90°. The difference in period indicates a difference in the coupling coefficient C, which is in agreement with the theoretical analysis presented in Section 6.2. This observed polarisation dependence is common to all directional coupler configurations provided there is a difference between κ_x and κ_y.

From the derivation in Section 6.2, it was shown that only κ_x and κ_y is needed to know the variation for all other values of θ. The values for κ_x and κ_y are thus inferred from the experimental data measured at θ = 0° and 90° by fitting the measured data to a sinusoid using Eq. 6.1 and minimising least-square errors, with an additional phase that takes into account the coupling in the bend region that connects the input and output to the interaction region, which corrects for the phase difference accumulated outside of the interaction region and is constant regardless of interaction region lengths [111]. Let the period of the fitted variation be T, then κ_x and κ_y are given by

$$\kappa_x = \frac{2\pi}{T_x} \quad \text{and} \quad \kappa_y = \frac{2\pi}{T_y}. \tag{6.14}$$

The following figures presents the experimental data points and corresponding fitted curves. The main source of measurement error arises from the coupling between the fibre and the waveguide input. Measurements were performed on the set of directional couplers for at least five times, and the measured data points are shown with error bars which indicate the range of measured values.

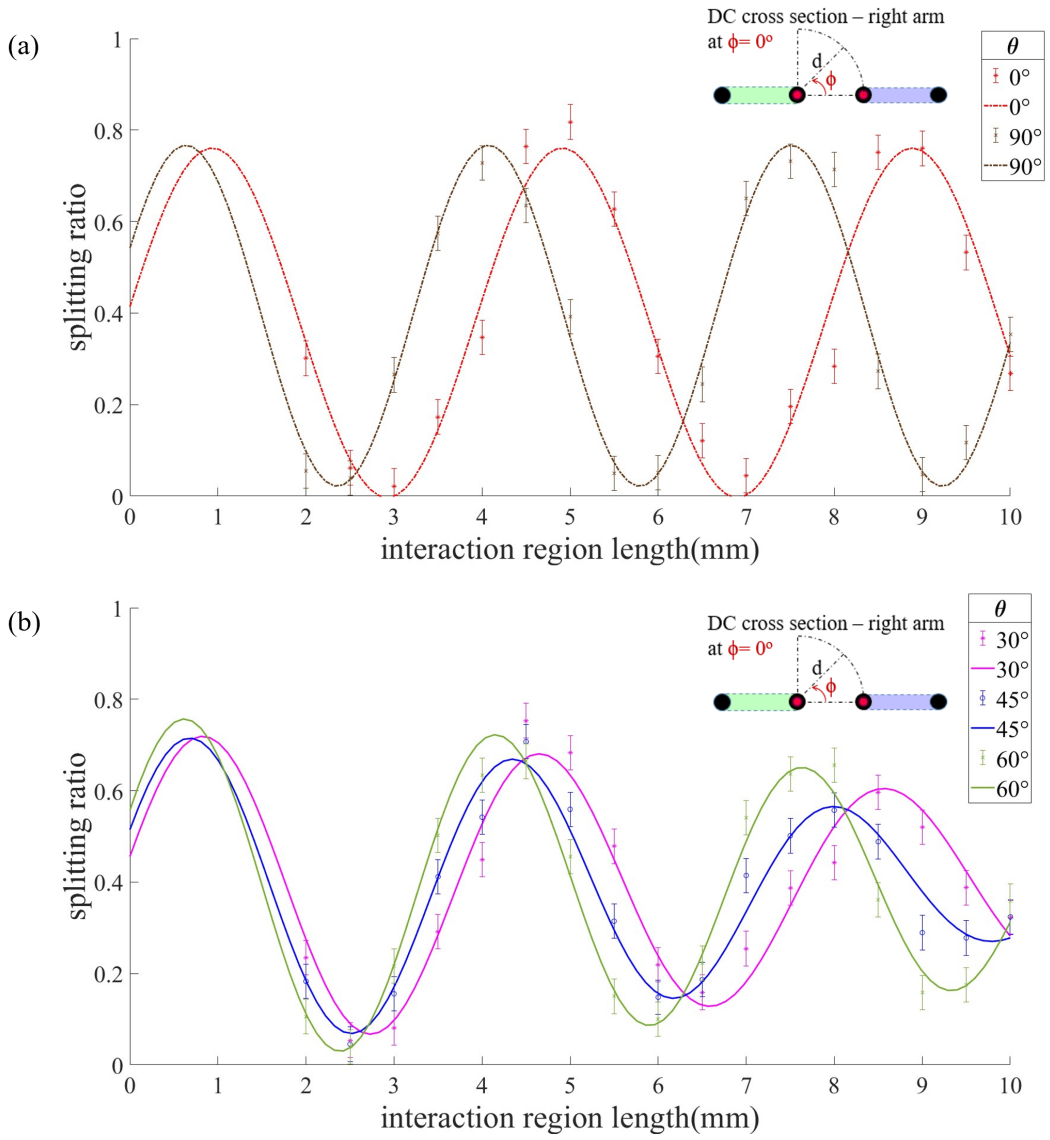

Figure 6.11: Splitting ratio against interaction region length of laser-written directional coupler (DC) with input polarisation angle θ (relative to the vertical symmetry axis), and ϕ at 0°. (a) Measured splitting ratios and least-squares fitted sinusoidal curves for θ at 0° and 90°. (b) Measured splitting ratios and theoretical predicted curves for θ at 30°, 45° and 60°. Error bars indicate measurement errors.

Fig. 6.11(a) show the fitted sinusoidal curves for polarisation angles θ = 0° and θ = 90°, which are used to infer the values of κ_x and κ_y using Eq. 6.14. The values are found to be κ_x = 1.83 rad/mm and κ_y = 1.58 rad/mm, which are typical of similar devices in literature [112]. Fig. 6.11(b) shows the predicted curves for polarisation angles θ = 30°, 45° and 60° using Eq. 6.9. The results show that the curves predicted by theory are in good agreement with the measured data. They are no longer sinusoidal at a single

frequency. Over several cycles a reduction in the maximum modulation is observed and the effect is most pronounced for a polarisation angle of 45°, which is exactly the case of $\theta = \pi/4$ as shown previously in Eq. 6.12. The sum of least-squares error for sinusoidal curve fitting in Fig. 6.11(a) is 0.070, whereas the errors for the theoretical predictions in Fig. 6.11(b) are 0.038, 0.029 and 0.020 respectively for θ = 30°, 45° and 60°. This shows that the theory matches the measured result with high accuracy.

To further verify the correctness of this understanding, another set of waveguides were fabricated where the waveguide arms of the directional couplers that were slightly offset such that they have the same separation distance d but are on different xz-planes and at an angle ϕ from each other on the cross-section view (Fig. 6.3). By compensating for the depth-dependent aberrations, the cross-sections of the waveguides is ensured to remain uniform at different depths. Comparing results for in-plane and out-of-plane directional couplers enables further understanding of the dependence on polarisation.

Fig. 6.12 presents the data for directional couplers with waveguide arms at $\phi = 90°$ with the same separation distance d = 8 μm. Similar to Fig. 6.11, the variation of splitting ratio with respect to interaction length is sinusoidal at a single frequency when the polarisation angles match the symmetry of the system, whilst in Fig. 6.12(b) the maximum splitting ratio is modulated over several cycles, with the effect greatest for an input linear polarisation at angle of θ = 45°. This confirms that the observed polarisation phenomenon is generalisable to the out-of-plane geometrical configurations of directional couplers.

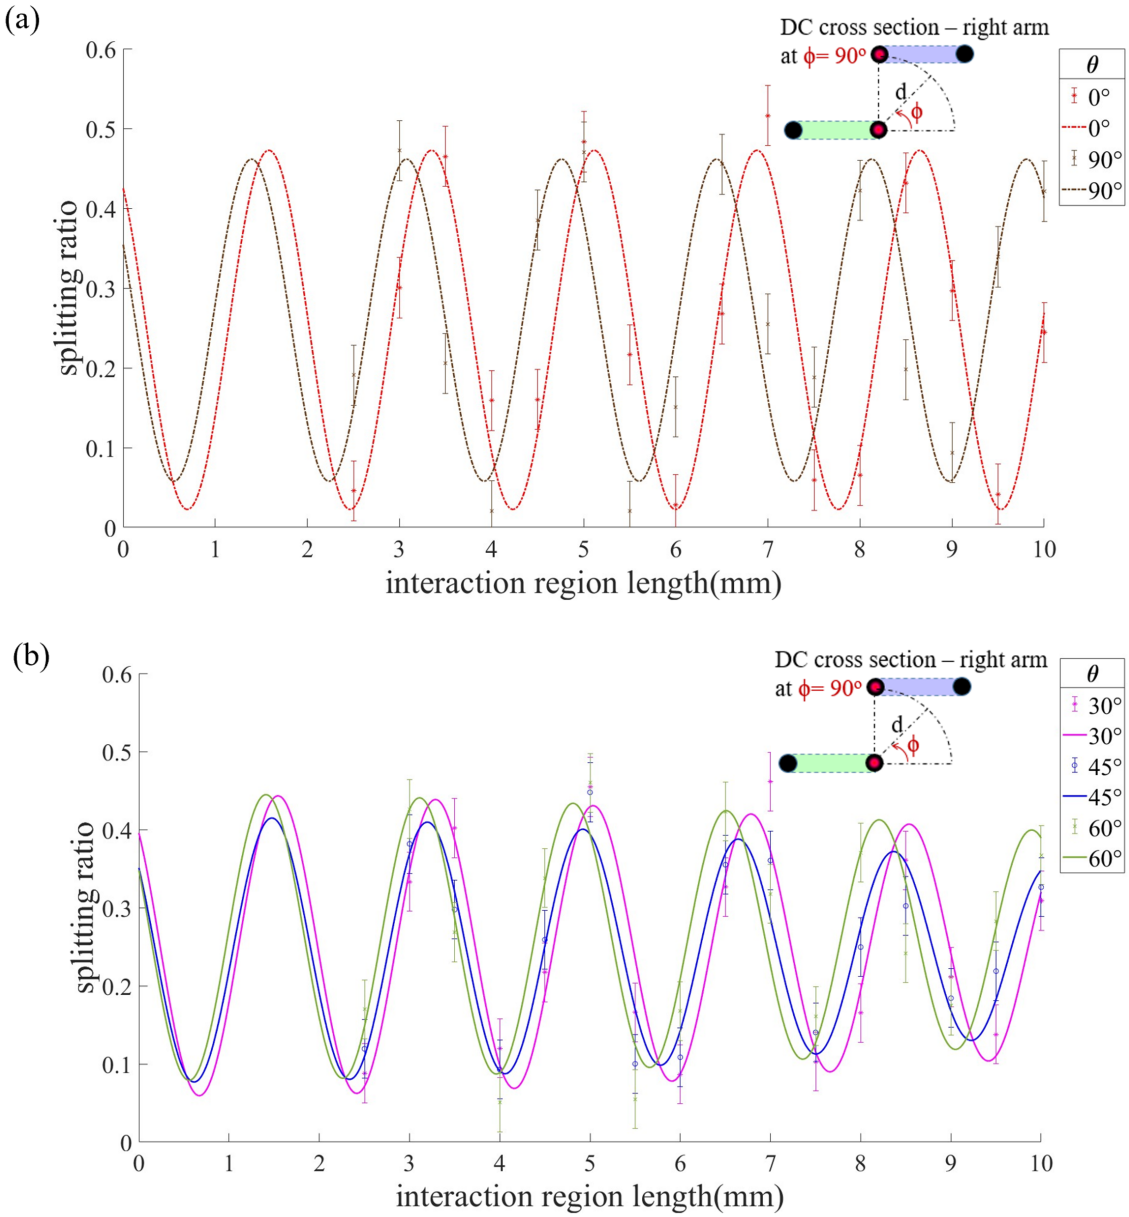

Figure 6.12: Splitting ratio against interaction region length of laser-written directional coupler with input polarisation angle θ (relative to the vertical symmetry axis), and ϕ at 90°. (a) Measured splitting ratios and least-squares fitted sinusoidal curves for θ at 0° and 90°. (b) Measured splitting ratios and theoretical predicted curves for θ at 30°, 45° and 60°. Error bars indicate measurement errors. Fig. 6.13 shows the case for $\phi = 45°$, where the geometry of the interaction region is effectively rotated by 45°. It is observed that the splitting ratios for polarisations at θ = 45° and 135° now show sinusoidal variation with a single frequency, matching the principal axes of the system. The values of κ are found to be $\kappa_{45°}$ = 1.80 rad/mm and $\kappa_{135°}$ = 1.77 rad/mm. The amplitude of the splitting ratio for angles not along the principal axes displays a modulation similar to the case of $\phi = 0°$ in Fig. 6.11.

It is also observed that when ϕ = 90°, the splitting ratio has a limited upper bound of around 0.5. From Fig. 6.12(a), the values of κ are found to be κ_x = 3.73 rad/mm

and κ_y = 3.56 rad/mm, approximately doubled from the case of $\phi = 0°$ in Fig. 6.11. This is partially due to the asymmetry of the waveguides causing dephasing between the coupling of the two waveguides. This effect is reflected by the σ term in Eq. 6.1. As a result, the amplitude of splitting ratio is reduced by a factor of σ^2, and the period of the variation is also reduced by a factor of σ. The value of σ can be found by finding the ratio between the amplitudes for the respective splitting ratios of $\phi = 0°$ and $\phi = 90°$. This gives $\sigma_x = 0.635$ and $\sigma_y = 0.671$. From Fig. 6.10(a), noting that the modified refractive index region of the fabricated waveguide cross-section was elongated along the y-direction. Since the separation of the two waveguides from center to center is fixed at 8 μm, the modified refractive index region of the upper waveguide overlaps partially with the lower one. This changes the refractive index profile in the interaction region between the two waveguides and affects the mode shape, hence the coupling between the two arms is also different. The measured values of κ resulted from a combination of both effects.

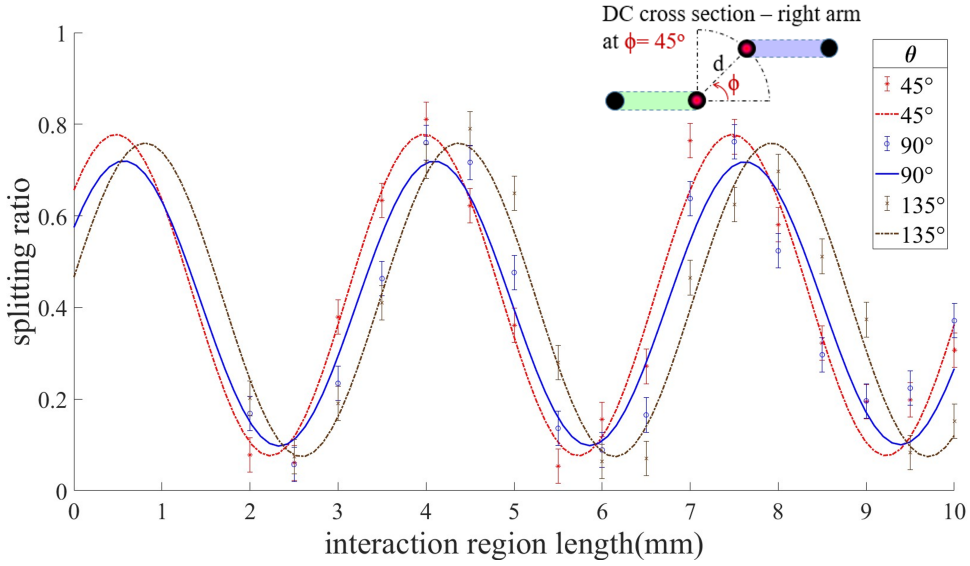

Figure 6.13: Splitting ratio against interaction region length of laser-written directional coupler with input polarisation angle θ (relative to the vertical y-axis), and ϕ at 45°: dashed curves - least-squares fitted sinusoidal curves for θ at 45° and 135°; solid curve - theoretical predicted curves for θ at 90°. Error bars indicate measurement errors.

Given the experimental results as shown above, the proposed theory is able to correctly predict the variation for arbitrary linear input polarisation given the geometry of the directional couplers. A simple prediction procedure can be performed as follows: Firstly

measure the κ_x and κ_y for any directional coupler geometry from finding the period of variation for polarisation along the principal axes by minimising least-square errors of data points for fitted sinusoidal curves. Eq. 6.9 is then used to infer the variation for all other values of angle θ. This enables much quicker characterisation and verification of directional coupler devices for more efficient prototyping and design processes.

Notably, the off-axis linearly polarised light gave the same power exchange for $+\theta$ and $-\theta$ states of polarisation. This is because the coupling is insensitive to the relative phase information and is solely determined by the magnitude of principal components. It is therefore predicted that for non-linearly polarised light, the variation can be predicted by the same procedure, i.e. projecting to the principal components and calculating the combined effect using Eq. 6.9. The phenomenon is also expected to showcase wavelength dependency. In particular, as wavelength increases, the asymmetry from the device geometry should be less significant, resulting in decreasing κ_x–κ_y, and the overall modulation will also become less apparent. Conversely, the modulation should be more pronounced at shorter wavelengths.

6.5 Conclusion and outlook for future work

In summary, a polarisation effect in directional couplers was analyzed theoretically and observed in experiments where there is a modulation of the power splitting ratio of monochromatic light in a directional coupler when the input polarisation does not align with the principal axes of the geometrical configuration. It is also showed how different geometrical configurations of relative angular offset affect this variation. Supported by experimental results, this modulation phenomenon is generalisable to out-of-plane geometrical configurations. The understanding is an important contribution to further optimisations of fabrication parameters in 3D photonic circuits. Combined with ways of controlling the birefringence of individual waveguides [113–115], this knowledge potentially enables new designs of applications with additional polarisation related multiplexing or integrated polarisation state readout, fully utilising the three-dimensional capabilities

of FLDW.

There are several potential future directions that may be further explored. Firstly, it would be illuminating to study and characterise the effect of shape birefringence on single waveguides, and how this relates to the directional coupler geometry. There has been substantial effort in creating circular cross-section waveguides for eliminating polarisation dependence in waveguides, but not much has been reported on tuning and utilising shape birefringence. This may open up new ways to control directional coupler behaviour instead of relying solely on device geometry for more sophisticated polarisation control and processing in integrated optical devices.

CHAPTER 7

Conclusion and Outlook

This thesis has investigated femtosecond laser direct writing (FLDW) as a platform for fabricating integrated photonic devices, with an emphasis on methods that expand the accessible design space of waveguides and provide greater control over their optical behaviour. The work addressed three complementary challenges: how to improve the robustness of fabrication close to device edges, how to design and realise complex cross-sections of multiscan waveguide arrays for their collective spatial mode control, and how to understand and characterise polarisation- and angle-dependent effects in directional couplers. Together, these studies highlight practical strategies that build on the established strengths of FLDW and expand its potential for realising complex three-dimensional integrated photonic devices.

The first contribution of the thesis was the demonstration of an in-situ fabrication method to overcome the limitations imposed by edge aberrations in FLDW. Conventionally, fabricated samples required polishing back the facet to expose usable interfaces, which was time-consuming, error-prone, and hindered integration with external devices. By applying an appropriate phase mask that restricted the fabrication beam to one half of the objective aperture, waveguides could be fabricated up to the sample edges without polishing. Characterisation showed that these waveguides exhibited guiding properties consistent with those produced by conventional procedures, confirming that the new fabrication scheme effectively compensated for the edge aberration. This method opens the door to fabrication in contexts where post-processing is undesirable, such as in

pre-assembled packages or hybrid photonic systems, and thus addresses a key barrier to practical deployment.

The second focus was the use of spherical phase-induced multiscan waveguides (SPIM-WGs) as a route to engineering arbitrary cross-sections of waveguide arrays for spatial mode control. Controlled spherical phase modulation enabled multiscan geometries to be written with high precision, and numerical simulations established how separation, wavelength, and geometry determine the supported modes. Both fibre-compatible symmetric arrays and deliberately asymmetric arrays designed to couple to ppKTP modes were studied, with the fabricated devices closely matching the predicted behaviour. This combination of modelling and experiment demonstrates how cross-sectional design can be used as a tool for creating devices that go beyond simple confinement, providing new opportunities for mode conversion and spatial tailoring in integrated optics.

The third line of investigation explored polarisation and angle-dependent behaviour in directional couplers, particularly when extended into three-dimensional out-of-plane geometries. A theoretical description was developed to show how splitting ratios of directional couplers vary with input polarisation, and experiments confirmed the predicted modulation effects. Devices fabricated with angular offsets demonstrated that the same behaviour extends to out-of-plane geometries, showing that the effect is not limited to planar structures. These results highlight the importance of accounting for polarisation in circuit design and point to ways in which it can be treated as a useful design parameter, rather than an unwanted side effect, thereby adding an extra degree of control in integrated photonics.

Viewed together, the three areas of investigation highlight different but interconnected aspects of FLDW: removing practical constraints in device fabrication, expanding flexibility in waveguide design, and deepening understanding of the behaviour of complex photonic circuits. In combination, they demonstrate how laser-written devices can evolve from simple waveguides towards complex, reliable, and functionally rich photonic

systems. This progression fits within the wider trajectory of FLDW research, where three-dimensional laser writing is increasingly being applied to reconfigurable and multi-functional devices [116–118]. State-of-the-art demonstrations of high-fidelity processors fabricated by femtosecond inscription now approach the performance of lithographic platforms while retaining unique three-dimensional flexibility [119]. The work presented here contributes to this broader progress by providing specific strategies for extending the capabilities of the platform.

Looking forward, several promising directions for future research emerge. A natural next step is to scale device-level methods into larger integrated systems. Real-world applications require multiple elements to be combined, such as converters, splitters, interferometers, and phase shifters. The techniques demonstrated in this thesis can be integrated to realise compact circuits that combine edge-optimised access, tailored cross-sections of multiscan waveguide arrays, and polarisation-aware directional couplers. Such developments would complement current efforts to build reconfigurable circuits with low power consumption and reduced crosstalk [120], and point towards laser-written systems that are not only functional but also scalable. In parallel, hybrid integration with other photonic platforms remains essential for industrial deployment. FLDW is attractive for prototyping and small-scale production because it is maskless, rapid, and inherently three-dimensional. Scaling to industry, however, will require compatibility with existing ecosystems. Recent demonstrations of hybrid devices, where laser-written waveguides directly interface with silicon photonics, illustrate how these approaches can be combined [121]. Pursuing such hybrid strategies could enable co-packaged optics for data centres, neuromorphic processors, or quantum networks, where three-dimensional geometries and tailored modal control can provide functionality that complements lithographic platforms [122].

Further opportunities lie in the continued development of SPIM-WGs as a flexible tool for spatial mode engineering. The results presented here have shown that carefully designed multiscan structures can reproduce both symmetric fibre-like modes and deliber-

ately asymmetric profiles. Extending this approach to more intricate cross-sections could enable devices that support higher-order spatial modes, including orbital angular momentum states [123] and other structured light fields of interest such as optical skyrmions [124], which are being explored for high-capacity communications and high-dimensional optical processing [125]. Developing converters and mode shapers that can reliably access these states would substantially broaden the scope of integrated photonics and provide a route to compact devices capable of generating or manipulating structured light directly on-chip.

Benchmarking and optimisation will also be important for translating these advances toward practical deployment. The present work established qualitative agreement between modelling and experiment, but systematic quantitative studies of propagation loss, coupling efficiency, and fabrication tolerance are needed to provide a stronger basis for reproducibility and scalability. Establishing such benchmarks would allow rigorous comparison across designs and identify structures most robust to variation, as an essential step in moving from laboratory demonstrations toward industrially relevant systems. At the same time, comparing alternative fabrication methods will sharpen understanding of where different approaches are best applied. Techniques such as the overlap-controlled multi-scan method [126] provide another route to producing complex geometries and examining them alongside SPIM-WGs can clarify their respective strengths. In some cases, hybrid strategies may combine the precision of multi-scan inscription with the design flexibility of phase-modulated writing, opening further possibilities for structured waveguides and customised mode profiles.

Finally, polarisation itself offers a distinct degree of freedom for future integrated photonics. The characterisation carried out in this thesis has shown how directional couplers exhibit systematic modulation when the input polarisation is varied, and that these effects extend into out-of-plane geometries. Rather than treating such behaviour as a limitation, it can be harnessed to provide additional channels for multiplexing, to prepare and analyse polarisation states, and to enable polarisation-sensitive routing within

larger circuits. Beyond communications [127] and quantum information processing [128], polarisation control also has relevance for emerging areas such as topological photonics, where the interplay between polarisation and spatial structure can be used to realise robust, disorder-resistant transport [129]. Incorporating polarisation in this way would complement the advances in waveguide cross-section design, expanding the set of controllable parameters for sophisticated three-dimensional photonic systems.

In conclusion, this thesis has demonstrated how FLDW can be advanced through edge-aberration compensation, simulation-driven multiscan waveguide cross-sectional design, and the characterisation of polarisation-dependent coupling. Together these contributions extend the design space of integrated photonics and illustrate the versatility of FLDW as a three-dimensional fabrication platform. By combining modelling, fabrication, and characterisation, the work provides a basis for future developments. With continued progress in device design, fabrication strategies, and system-level integration, femtosecond laser writing is well positioned to play a central role in the next generation of optical technologies.

Bibliography

1. H. H. Radamson *et al.* State of the Art and Future Perspectives in Advanced CMOS Technology. *Nanomaterials* **10,** 1555 (2020).
2. L. Johnsson & G. Netzer. The impact of Moore's Law and loss of Dennard scaling: Are DSP SoCs an energy efficient alternative to x86 SoCs? *Journal of Physics: Conference Series* **762,** 012022 (2016).
3. J. Shalf. The future of computing beyond Moore's Law. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **378,** 20190061 (2020).
4. C. E. Leiserson *et al.* There's plenty of room at the Top: What will drive computer performance after Moore's law? *Science* **368,** eaam9744 (2020).
5. P. L. McMahon. The physics of optical computing. *Nature Reviews Physics* **5,** 717–734 (2023).
6. C. G. Kibebe, Y. Liu & J. Tang. Harnessing optical advantages in computing: a review of current and future trends. *Frontiers in Physics* **12** (2024).
7. S. Ning *et al.* Photonic-Electronic Integrated Circuits for High-Performance Computing and AI Accelerators. *Journal of Lightwave Technology* **42,** 7834–7859 (2024).
8. R. Marchetti, C. Lacava, L. Carroll, K. Gradkowski & P. Minzioni. Coupling strategies for silicon photonics integrated chips [Invited]. *Photonics Research* **7,** 201 (2019).
9. K. M. Davis, K. Miura, N. Sugimoto & K. Hirao. Writing waveguides in glass with a femtosecond laser. *Optics Letters* **21,** 1729–1731 (1996).
10. V. Sukumaran, T. Bandyopadhyay, V. Sundaram & R. Tummala. Low-Cost Thin Glass Interposers as a Superior Alternative to Silicon and Organic Interposers for Packaging of 3-D ICs. *IEEE Transactions on Components, Packaging and Manufacturing Technology* **2,** 1426–1433 (2012).
11. N. Psaila, Z. Chaboyer, J. Macdonald & P. Tadayon. Passively Aligned Glass Micro-Optic Bridge for Expanded-Beam Vertical Coupling and Pluggable Silicon Photonics. *Journal of Lightwave Technology* **42,** 5223–5230 (2024).
12. A. E. Siegman. *Lasers* (Anthony Edward) (Sausalito, Calif. : University Science, Sausalito, Calif., 1986).
13. P. M. W. French. The generation of ultrashort laser pulses. *Reports on Progress in Physics* **58,** 169 (1995).
14. M. Piché. Beam reshaping and self-mode-locking in nonlinear laser resonators. *Optics Communications* **86,** 156–160 (1991).
15. D. Strickland & G. Mourou. Compression of amplified chirped optical pulses. *Optics Communications* **56,** 219–221 (1985).
16. A. Dubietis, G. Jonušauskas & A. Piskarskas. Powerful femtosecond pulse generation by chirped and stretched pulse parametric amplification in BBO crystal. *Optics Communications* **88,** 437–440 (1992).
17. M. Ams *et al.* Investigation of Ultrafast Laser–Photonic Material Interactions: Challenges for Directly Written Glass Photonics. *IEEE Journal of Selected Topics in Quantum Electronics* **14,** 1370–1381 (2008).
18. R. Osellame, G. Cerullo & R. Ramponi. *Femtosecond laser micromachining : photonic and microfluidic devices in transparent materiels* 1st 2012. Includes bibliographical references and index. (Springer, Berlin, 2012).
19. S. M. Eaton *et al.* Transition from thermal diffusion to heat accumulation in high repetition rate femtosecond laser writing of buried optical waveguides. *Optics Express* **16,** 9443 (2008).
20. X. Liu, D. Du & G. Mourou. Laser ablation and micromachining with ultrashort laser pulses. *IEEE Journal of Quantum Electronics* **33,** 1706–1716 (1997).
21. K. Sugioka & Y. Cheng. Ultrafast lasers—reliable tools for advanced materials processing. *Light: Science & Applications* **3,** e149–e149 (2014).
22. J. W. Chan, T. Huser, S. Risbud & D. M. Krol. Structural changes in fused silica after exposure to focused femtosecond laser pulses. *Optics Letters* **26,** 1726–1728 (2001).

23. L. Sudrie, M. Franco, B. Prade & A. Mysyrowicz. Writing of permanent birefringent microlayers in bulk fused silica with femtosecond laser pulses. *Optics Communications* **171,** 279–284 (1999).
24. B. Zhang, X. Liu & J. Qiu. Single femtosecond laser beam induced nanogratings in transparent media - Mechanisms and applications. *Journal of Materiomics* **5,** 1–14 (2019).
25. M. Sakakura, Y. Lei, L. Wang, Y.-H. Yu & P. G. Kazansky. Ultralow-loss geometric phase and polarization shaping by ultrafast laser writing in silica glass. *Light: Science & Applications* **9** (2020).
26. M. K. Bhuyan *et al.* Ultrafast laser nanostructuring in bulk silica, a "slow" microexplosion. *Optica* **4,** 951 (2017).
27. D. J. Little, M. Ams & M. J. Withford. Influence of bandgap and polarization on photo-ionization: guidelines for ultrafast laser inscription [Invited]. *Optical Materials Express* **1,** 670 (2011).
28. G. Della Valle, R. Osellame & P. Laporta. Micromachining of photonic devices by femtosecond laser pulses. *Journal of Optics A: Pure and Applied Optics* **11,** 013001 (2008).
29. A. Fuerbach *et al. Refractive index change mechanisms in different glasses induced by femtosecond laser irradiation* (SPIE, 2016).
30. Q. Sun *et al.* Effect of spherical aberration on the propagation of a tightly focused femtosecond laser pulse inside fused silica. *Journal of Optics A: Pure and Applied Optics* **7,** 655 (2005).
31. C. Hnatovsky *et al.* High-resolution study of photoinduced modification in fused silica produced by a tightly focused femtosecond laser beam in the presence of aberrations. *Journal of Applied Physics* **98** (2005).
32. H. W. Babcock. The Possibility of Compensating Astronomical Seeing. *Publications of the Astronomical Society of the Pacific* **65,** 229 (1953).
33. M. J. Booth. Adaptive optics in microscopy. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **365,** 2829–2843 (2007).
34. I. Kozak. Retinal imaging using adaptive optics technology. *Saudi Journal of Ophthalmology* **28,** 117–122 (2014).
35. J. W. Goodman. *Introduction to Fourier optics* Fourth edition. Includes bibliographical references and index. (W.H. Freeman, New York, 2017).
36. v. F. Zernike. Beugungstheorie des schneidenver-fahrens und seiner verbesserten form, der phasenkontrastmethode. *Physica* **1,** 689–704 (1934).
37. V. Lakshminarayanan & A. Fleck. Zernike polynomials: a guide. *Journal of Modern Optics* **58.** doi: 10.1080/09500340.2011.554896, 545–561 (2011).
38. M. A. A. Neil, M. J. Booth & T. Wilson. New modal wave-front sensor: a theoretical analysis. *Journal of the Optical Society of America A* **17,** 1098–1107 (2000).
39. M. J. Booth, M. A. A. Neil & T. Wilson. Aberration correction for confocal imaging in refractive-index-mismatched media. *Journal of Microscopy* **192,** 90–98 (1998).
40. P. S. Salter, M. Baum, I. Alexeev, M. Schmidt & M. J. Booth. Exploring the depth range for three-dimensional laser machining with aberration correction. *Optics Express* **22,** 17644 (2014).
41. M. J. Booth. Adaptive optical microscopy: the ongoing quest for a perfect image. *Light: Science & Applications* **3,** e165–e165 (2014).
42. P.-Y. Madec. *Overview of deformable mirror technologies for adaptive optics and astronomy* (SPIE, 2012).
43. P. S. Salter & M. J. Booth. Adaptive optics in laser processing. *Light: Science & Applications* **8** (2019).
44. R. W. Gerchberg. A practical algorithm for the determination of phase from image and diffraction plane pictures. *Optik* **35,** 237–246 (1972).
45. S. Hasegawa, H. Ito, H. Toyoda & Y. Hayasaki. Massively parallel femtosecond laser processing. *Optics Express* **24,** 18513–18524 (2016).
46. P. S. Salter *et al.* Adaptive slit beam shaping for direct laser written waveguides. *Optics Letters* **37,** 470–472 (2012).
47. A. Hendriks, D. Naidoo, F. Roux, C. López-Mariscal & A. Forbes. *The generation of flat-top beams by complex amplitude modulation with a phase-only spatial light modulator* (SPIE, 2012).
48. L. Yang *et al.* Direct laser writing of complex microtubes using femtosecond vortex beams. *Applied Physics Letters* **110** (2017).

49. M. Born & E. Wolf. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* 7th ed. (Cambridge University Press, Cambridge, 1999).
50. W. S. C. Chang. in *Principles of Optics for Engineers: Diffraction and Modal Analysis* 109–147 (Cambridge University Press, Cambridge, 2015).
51. W.-P. Huang. Coupled-mode theory for optical waveguides: an overview. *Journal of the Optical Society of America A* **11,** 963–983 (1994).
52. R. R. A. Syms & J. Cozens. *Optical guided waves and devices* Includes bibliographical references and index (McGraw-Hill, London, 1992).
53. A. W. Snyder & J. D. Love. in *Optical Waveguide Theory* (eds Snyder, A. W. & Love, J. D.) 280–300 (Springer US, Boston, MA, 1983).
54. A. W. Snyder. Understanding monomode optical fibers. *Proceedings of the IEEE* **69,** 6–13 (1981).
55. N. Imoto, N. Yoshizawa, J. Sakai & H. Tsuchiya. Birefringence in single-mode optical fiber due to elliptical core deformation and stress anisotropy. *IEEE Journal of Quantum Electronics* **16,** 1267–1271 (1980).
56. R. Osellame *et al.* Femtosecond writing of active optical waveguides with astigmatically shaped beams. *Journal of the Optical Society of America B* **20,** 1559–1567 (2003).
57. M. Ams, G. D. Marshall, D. J. Spence & M. J. Withford. Slit beam shaping method for femtosecond laser direct-write fabrication of symmetric waveguides in bulk glasses. *Optics Express* **13,** 5676 (2005).
58. A. Panda, P. Sarkar & G. Palai. Studies on coupling of optical power in fiber to semiconductor waveguide at wavelength 1550 nm for photonics integrated circuits. *Optik* **157,** 944–950 (2018).
59. S. Gross & M. J. Withford. Ultrafast-laser-inscribed 3D integrated photonics: challenges and emerging applications. *Nanophotonics* **4,** 332–352 (2015).
60. G. Djogo *et al.* Femtosecond laser additive and subtractive micro-processing: enabling a high-channel-density silica interposer for multicore fibre to silicon-photonic packaging. *International Journal of Extreme Manufacturing* **1,** 045002 (2019).
61. A. Desmet, A. Radosavljevic, J. Missinne, D. Van Thourhout & G. Van Steenberge. Laser Written Glass Interposer for Fiber Coupling to Silicon Photonic Integrated Circuits. *IEEE Photonics Journal* **13,** 1–12 (2021).
62. N. N. Skryabin *et al.* Femtosecond-laser-written low-loss multiscan waveguides in fused silica. *Physical Review Applied* **22** (2024).
63. X. Li *et al.* Three-dimensional broadband 90° optical hybrid with low crosstalk and phase error processed via femtosecond laser direct writing. *Optics Communications* **579,** 131537 (2025).
64. G. D. Marshall *et al.* Laser written waveguide photonic quantum circuits. *Optics Express* **17,** 12546 (2009).
65. N. Spagnolo *et al.* Three-photon bosonic coalescence in an integrated tritter. *Nature Communications* **4,** 1606 (2013).
66. S. Atzeni *et al.* Integrated sources of entangled photons at the telecom wavelength in femtosecond-laser-written circuits. *Optica* **5,** 311 (2018).
67. W. Yan, B. Zhang & F. Chen. Photonic topological insulators in femtosecond laser direct-written waveguides. *npj Nanophotonics* **1** (2024).
68. M. C. Rechtsman *et al.* Photonic Floquet topological insulators. *Nature* **496,** 196–200 (2013).
69. A. J. Menssen, J. Guan, D. Felce, M. J. Booth & I. A. Walmsley. Photonic Topological Mode Bound to a Vortex. *Physical Review Letters* **125** (2020).
70. J.-L. Tambasco *et al.* Quantum interference of topological states of light. *Science Advances* **4,** eaat3187 (2018).
71. M. Ehrhardt, C. Dittel, M. Heinrich & A. Szameit. Topological Hong-Ou-Mandel interference. *Science* **384,** 1340–1344 (2024).
72. *Pharos User's Manual* Light Conversion (2018).
73. A. J. Littlefield *et al.* Low loss fiber-coupled volumetric interconnects fabricated via direct laser writing. *Optica* **11,** 995–1007 (2024).
74. P. S. Salter & M. J. Booth. Focussing over the edge: adaptive subsurface laser fabrication up to the sample face. *Optics Express* **20,** 19978 (2012).
75. R. Holinirina Dina Miora *et al.* Calculating point spread functions: methods, pitfalls, and solutions. *Optics Express* **32,** 27278 (2024).

76. S. M. Eaton *et al.* Heat accumulation effects in femtosecond laser-written waveguides with variable repetition rate. *Optics Express* **13,** 4708 (2005).
77. B. Sun *et al.* Fast, Precise, High Contrast Laser Writing for Photonic Chips with Phase Aberrations. *Laser & Photonics Reviews* **18** (2024).
78. A. Jesacher, P. S. Salter & M. J. Booth. Refractive index profiling of direct laser written waveguides: tomographic phase imaging. *Optical Materials Express* **3,** 1223 (2013).
79. T. Mizuno, H. Takara, A. Sano & Y. Miyamoto. Dense Space-Division Multiplexed Transmission Systems Using Multi-Core and Multi-Mode Fiber. *Journal of Lightwave Technology* **34,** 582–592 (2016).
80. Q. Ling *et al.* On-chip optical matrix-vector multiplier based on mode division multiplexing. *Chip* **2,** 100061 (2023).
81. K. R. Mojaver, S. M. R. Safaee, S. S. Morrison & O. Liboiron-Ladouceur. Recent Advancements in Mode Division Multiplexing for Communication and Computation in Silicon Photonics. *Journal of Lightwave Technology* **42,** 7860–7870 (2024).
82. G. S. Thekkadath, B. A. Bell, R. B. Patel, M. S. Kim & I. A. Walmsley. Measuring the Joint Spectral Mode of Photon Pairs Using Intensity Interferometry. *Physical Review Letters* **128.** PRL, 023601 (2022).
83. Z.-Z. Li *et al.* Three-dimensional on-chip mode converter. *Optics Letters* **48,** 1140–1143 (2023).
84. P. H. D. Ferreira, G. F. B. Almeida & C. R. Mendonça. A simple strategy for increasing optical waveguide performance using spherical aberration. *Optics Laser Technology* **142,** 107235 (2021).
85. T. T. Fernandez *et al.* Controlling plasma distributions as driving forces for ion migration during fs laser writing. *Journal of Physics D: Applied Physics* **48,** 155101 (2015).
86. J. A. Nelder & R. Mead. A Simplex Method for Function Minimization. *The Computer Journal* **7,** 308–313 (1965).
87. B. Sun *et al.* On-chip beam rotators, adiabatic mode converters, and waveplates through low-loss waveguides with variable cross-sections. *Light: Science & Applications* **11** (2022).
88. N. Barré *et al.* Tomographic refractive index profiling of direct laser written waveguides. *Optics Express* **29,** 35414 (2021).
89. G. S. Thekkadath *et al.* Quantum-enhanced interferometry with large heralded photon-number states. *npj Quantum Information* **6** (2020).
90. C. Eigner *et al.* Periodically poled ridge waveguides in KTP for second harmonic generation in the UV regime. *Optics Express* **26,** 28827–28833 (2018).
91. J. Carolan *et al.* Universal linear optics. *Science* **349,** 711–716 (2015).
92. W. Cheng, W. Liu, Q. Liu & F. Chen. Observation of topological Anderson phase in laser-written quasi-periodic waveguide arrays. *Optics Letters* **47,** 2883–2886 (2022).
93. D. González-Andrade *et al.* Spatial and Polarization Division Multiplexing Harnessing On-Chip Optical Beam Forming. *Laser & Photonics Reviews* **17** (2023).
94. C. Pentangelo *et al.* High-fidelity and polarization-insensitive universal photonic processors fabricated by femtosecond laser writing. *Nanophotonics* **13,** 2259–2270 (2024).
95. T. Meany *et al.* Laser written circuits for quantum photonics. *Laser Photonics Reviews* **9,** 363–384 (2015).
96. A. M. Streltsov & N. F. Borrelli. Fabrication and analysis of a directional coupler written in glass by nanojoule femtosecond laser pulses. *Optics Letters* **26,** 42–43 (2001).
97. L. A. Fernandes, J. R. Grenier, P. R. Herman, J. S. Aitchison & P. V. S. Marques. Femtosecond laser fabrication of birefringent directional couplers as polarization beam splitters in fused silica. *Optics Express* **19,** 11992 (2011).
98. G. Corrielli *et al.* Rotated waveplates in integrated waveguide optics. *Nature Communications* **5** (2014).
99. L. A. Fernandes, J. R. Grenier, P. R. Herman, J. S. Aitchison & P. V. S. Marques. Femtosecond laser writing of waveguide retarders in fused silica for polarization control in optical circuits. *Optics Express* **19,** 18294 (2011).
100. A. Szameit, F. Dreisow, T. Pertsch, S. Nolte & A. Tünnermann. Control of directional evanescent coupling in fs laser written waveguides. *Optics Express* **15,** 1579 (2007).

101. L. Sansoni *et al.* Two-Particle Bosonic-Fermionic Quantum Walk via Integrated Photonics. *Physical Review Letters* **108** (2012).
102. X.-L. Zhang *et al.* Non-Abelian braiding on photonic chips. *Nature Photonics* **16,** 390–395 (2022).
103. X.-Y. Xu, X.-W. Wang, D.-Y. Chen, C. M. Smith & X.-M. Jin. Quantum transport in fractal networks. *Nature Photonics* **15,** 703–710 (2021).
104. A. Hardy & W. Streifer. Coupled mode theory of parallel waveguides. *Journal of Lightwave Technology* **3,** 1135–1146 (1985).
105. W.-J. Chen, S. M. Eaton, H. Zhang & P. R. Herman. Broadband directional couplers fabricated in bulk glass with high repetition rate femtosecond laser pulses. *Optics Express* **16,** 11470 (2008).
106. A. W. Snyder & W. R. Young. Modes of optical waveguides. *Journal of the Optical Society of America* **68,** 297–309 (1978).
107. F. Kapron, N. Borrelli & D. Keck. Birefringence in dielectric optical waveguides. *IEEE Journal of Quantum Electronics* **8,** 222–225 (1972).
108. K. Thygarajan, A. Kumar & A. K. Ghatak. Perturbation analysis of the coupling between two cylindrical parabolic index waveguides. *Optics Communications* **39,** 26–30 (1981).
109. A. W. Snyder & J. D. Love. *Optical waveguide theory* Includes bibliographical references and indexes. (Chapman and Hall, London, 1983).
110. G. Corrielli *et al.* Symmetric polarization-insensitive directional couplers fabricated by femtosecond laser writing. *Optics Express* **26,** 15101 (2018).
111. P. L. Auger & S. Iraj Najafi. New method to design directional coupler dual wavelength multi/demultiplexer with bends at both extremities. *Optics Communications* **111,** 43–50 (1994).
112. A. Crespi *et al.* Integrated photonic quantum gates for polarization qubits. *Nature Communications* **2,** 566 (2011).
113. F. Yu *et al.* Polarization Independent Quantum Devices With Ultra-Low Birefringence Glass Waveguides. *Journal of Lightwave Technology* **39,** 1451–1457 (2021).
114. J. Lapointe, J.-P. Bérubé, A. Dupont, M. Bellec & R. Vallée. Modified astigmatic beam technique for laser writing. *Applied Optics* **61,** 2333–2337 (2022).
115. Y.-D. Wang *et al.* Ultralow birefringent glass waveguide fabricated by femtosecond laser direct writing. *Optics Letters* **48,** 554–557 (2023).
116. J. Zhu, M. Dexheimer & H. Cheng. Reconfigurable systems for multifunctional electronics. *npj Flexible Electronics* **1** (2017).
117. Y. Zhang *et al.* A Novel Multifunctional Material for Constructing 3D Multi-Response Structures Using Programmable Two-Photon Laser Fabrication. *Advanced Functional Materials* **34,** 2313922 (2024).
118. A. Wang *et al.* In-chip critical plasma seeds for laser writing of reconfigurable silicon photonics systems. *Nature Communications* **16** (2025).
119. L. Zhong, Y. Wang, X. Han, J. Hu & J. Qiu. Laser-direct-lithography of large-area 3D integrated photonics: technological challenges and advances. *Chip*, 100157 (2025).
120. F. Ceccarelli *et al.* Low Power Reconfigurability and Reduced Crosstalk in Integrated Photonic Circuits Fabricated by Femtosecond Laser Micromachining. *Laser & Photonics Reviews* **14,** 2000024 (2020).
121. G. Gualandi *et al.* Laser-written reconfigurable photonic integrated circuit directly coupled to a single-photon avalanche diode array. *Light: Science & Applications* **14** (2025).
122. J. Li, C. Cao, Y. Qiu, C. Kuang & X. Liu. Optical Waveguides Fabricated via Femtosecond Direct Laser Writing: Processes, Materials, and Devices. *Advanced Materials Technologies* **8** (2023).
123. Y. Chen *et al.* Mapping Twisted Light into and out of a Photonic Chip. *Physical Review Letters* **121** (2018).
124. A. A. Wang *et al.* Optical Skyrmions in Waveguides. *arXiv:2505.06735* (2025).
125. I. Bente *et al.* The potential of multidimensional photonic computing. *Nature Reviews Physics* **7,** 439–450 (2025).
126. Y. Wang *et al.* Precise mode control of laser-written waveguides for broadband, low-dispersion 3D integrated optics. *Light: Science & Applications* **13** (2024).
127. Z. Liu *et al.* Versatile on-chip polarization-sensitive detection system for optical communication and artificial vision. *Light: Science & Applications* **14** (2025).

128. Z. Zhang *et al.* Entanglement-based quantum information technology: a tutorial. *Advances in Optics and Photonics* **16,** 60–162 (2024).
129. A. A. Wang *et al.* Topological protection of optical skyrmions through complex media. *Light: Science & Applications* **13** (2024).