UNIVERSITY OF SOUTHAMPTON

Faculty of Engineering and Physical Sciences
Zepler Institute for Photonics and Nanoelectronics
Optoelectronics Research Centre

Wideband mid-infrared group IV photonic devices and platforms

by

Callum John Stirling

MSci (Hons) MInstP

ORCiD: 0000-0001-5274-2100

A thesis for the degree of
Doctor of Philosophy

February 2022

Mid-infrared group IV photonics is a field which, by adapting techniques from silicon photonics at visible and near-infrared wavelengths and using mature semiconductor fabrication processes, could establish an enabling technology for a diverse range of applications in numerous areas. In particular, integrated photonic sensors could take advantage of the characteristic absorptions of many chemicals at mid-infrared wavelengths, due to strong fundamental molecular vibrations in this region. Such "lab-on-a-chip" devices would be applied to areas like medical diagnostics and environmental monitoring.

To develop complex mid-infrared photonic integrated circuits, a set of core building-block components are required. Currently, the components in mid-infrared group IV photonics are limited to operating at relatively narrow wavelength ranges. This lack of spectral bandwidth is not an issue for some applications, but to unlock the full potential of the field, it is essential to develop wideband devices.

The operating wavelength range of mid-infrared devices may be limited by absorptions of the material platform or the geometry of the component; this work considers both to increase the available spectral bandwidth. Silicon-on-insulator waveguides with propagation losses ~ 1.5 dB/cm are shown to only support the fundamental mode over an octave of frequency, as an experimental demonstration of a technique that in principle will be applicable to much of the mid-infrared range. Beam splitters were fabricated on silicon-on-insulator platforms with low insertion losses and high performance over a bandwidth of $3.1 - 3.7$ µm: multimode interferometers are shown with an insertion loss of <1 dB and imbalance of <0.5 dB; and an insertion loss of ~ 0.2 dB was achieved for 50/50 Y-splitters. Further, considering material platforms, silicon membrane devices have been successfully transfer printed onto a high-transparency zinc selenide substrates, to develop waveguides without substrate absorption losses.

Contents

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;

2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;

3. Where I have consulted the published work of others, this is always clearly attributed;

4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;

5. I have acknowledged all main sources of help;

6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;

7. Parts of this work have been published as:
   Journal publications

   - B. Troia, J. Soler Penades, Z. Qu, A. Z. Khokhar, A. Osman, Y. Wu, **C. Stirling**, M. Nedeljkovic, V. M. N. Passaro, and G. Z. Mashanovich, "Silicon ring resonator-coupled Mach-Zehnder interferometers for the Fano resonance in the mid-IR", *Applied Optics* 56(31):8769–8776, 2017

   - G. Z. Mashanovich, M. Nedeljkovic, J. Soler-Penadés, Z. Qu, W. Cao, A. Osman, Y. Wu, **C. J. Stirling**, Y. Qi, Y. X. Cheng, L. Reid, C. G. Littlejohns, J. Kang, Z. Zhao, M. Takenaka, T. Li, Z. Zhou, F. Y. Gardes, D. J. Thomson, and G. T. Reed, "Group IV mid-infrared photonics [Invited]", *Optical Materials Express* 8(8):2276–2286, 2018

- **C. J. Stirling**, R. Halir, A. Sánchez-Postigo, Z. Qu, J. D. Reynolds, J. Soler Penadés, G. S. Murugan, A. Ortega-Moñux, J. G. Wangüemert-Pérez, Í. Molina-Fernández, G. Z. Mashanovich, and M. Nedeljkovic, "Broadband 2×2 multimode interference coupler for mid-infrared wavelengths", *Optics Letters* 46(21):5300–5303, 2021

- **C. J. Stirling**, J. D. Reynolds, Z. Qu, T. D. Bradley, L. Mastronardi, F. Y. Gardes, and M. Nedeljkovic, "Mid-infrared silicon-on-insulator waveguides with single-mode propagation over an octave of frequency", *Optics Express* 30(5) 2022

Conference publications

- J. Soler Penadés, A. Osman, A. Sánchez-Postigo, Z. Qu, Y. Wu, **C. J. Stirling**, D. P. Cheben, A. Ortega-Moñux, J. G. Wangüemert-Pérez, M. Nedeljkovic, and G. Z. Mashanovich, "Silicon and germanium mid-infrared platforms", *Semiconductor and Integrated Opto-Electronics (SIOE)*, Cardiff, UK, Apr. 2019

- G. Z. Mashanovich, J. Soler Penadés, A. Osman, A. Sánchez-Postigo, Z. Qu, Y. Wu, **C. J. Stirling**, D. P. Cheben, A. Ortega-Moñux, J. G. Wangüemert-Pérez, I. Molina-Fernández, and M. Nedeljkovic, "Silicon and germanium-based mid-infrared platforms", *Photonics & Electromagnetics Research Symposium (PIERS)*, Rome, Italy, June 2019

- G. Z. Mashanovich, J. Soler Penades, Y. Wu, V. Mittal, A. Osman, W. Cao, Z. Qu, Y. Qi, **C. Stirling**, M. Milosevic, L. Reid, D. J. Rowe, D. J. Thomson, and M. Nedeljkovic, "Mid-IR silicon and germanium photonic circuits for communications and sensing", *IEEE Summer Topicals Meeting*, Fort Lauderdale, USA, July 2019

- **C. J. Stirling**, B. Guilhabert, W. Cao, K. S. Kiang, A. Z. Khokhar, M. Nedeljkovic, M. J. Strain, and G. Z. Mashanovich, "Investigations into group IV photonic waveguides with a wide working optical bandwidth", *Silicon Photonics XV (SPIE Photonics West)*, ed. by G. T. Reed and A. P. Knights, vol. 11285, San Francisco, USA, Feb. 2020, pp. 224–233

- **C. J. Stirling**, J. D. Reynolds, Z. Qu, T. D. Bradley, L. Mastronardi, F. Y. Gardes, G. Z. Mashanovich, and M. Nedeljkovic, "Mid-infrared SOI waveguides with broadband single-mode propagation", *15th International Conference on Mid-Infrared Optoelectronic Materials and Devices (MIOMD)*, Sept. 2021

- M. Nedeljkovic, C. Wei, **C. J. Stirling**, Y. Qi, L. Reid, Y. Wu, W. Cao, M. Tittle, and G. Z. Mashanovich, "Silicon photonic photodetectors, sensors, and spectrometers for the mid-infrared", *15th International Conference on Mid-Infrared Optoelectronic Materials and Devices (MIOMD)*, Sept. 2021

- G. Mashanovich, **C. Stirling**, L. Reid, C. Wei, Y. Qi, Y. Wu, Z. Qu, G. Georgiev, W. Cao, C. Mitchell, and M. Nedeljkovic, "Broadband silicon photonic mid-infrared devices", *The European Optical Society Annual Meeting (EOSAM)*, Sept. 2021

- G. Mashanovich, C. Wei, L. Reid, G. Georgiev, Y. Qi, **C. Stirling**, A. Osman, Y. Wu, Z. Qu, D. Thomson, K. Li, J. Heffernan, K. Groom, W. Cao, C. Mitchell, and M. Nedeljkovic, "Silicon and germanium photonic devices for sensing and communications", *IEEE Photonics Conference (IPC)*, Oct. 2021

Outreach publications

- **C. J. Stirling**, A. L. Donko, N. Baktash, K. Grabska, A. L. Camacho Rosales, D. T. Clarke, J. J. Prentice, and M. T. Posner, "Student-led outreach and public engagement activities at the University of Southampton to celebrate the inaugural International Day of Light", *Optics Education and Outreach V (SPIE Optics + Photonics)*, ed. by G. Groot Gregory, vol. 10741, San Diego, USA, Aug. 2018, pp. 67–77

- C. Holmes, P. John, and **C. Stirling**, "A Nobel cause: public engagement and outreach", *Europhysics News* 50(2):19–22, 2019

Signed:...   Date:..................

Acknowledgements

First and foremost, this work would not have been possible with the great support and guidance of my supervisors, Prof. Goran Mashanovich and Dr. Milos Nedeljkovic. Their insight and tireless encouragement have kept my research focussed in the right direction and have provided me with the platform to be successful.

I must thank all my past and present colleagues in the Silicon Photonics Group for their help on countless occasions, but particularly Lauren Reid, Dr. Zhibo Qu, Dr. Jordi Soler Penadés, Dr. Jamie Reynolds, Dr. Dave Rowe and Dr. Lorenzo Mastronardi for their help within and without the laboratories and cleanroom.

The results of this project are due in part to fruitful collaborations with colleagues at other universities. I would like to extend special thanks to my collaborators at the University of Málaga, for the lengthy discussions over the subwavelength-structured MMIs and their simulations: Dr. Robert Halir, Prof. Gonzalo Wangüemert-Pérez, Dr. Alejandro Sánchez-Postigo and Dr. Alejandro Ortega-Moñux. I would also like to thank Prof. Michael Strain and Dr Benoit Guilhabert at the University of Strathclyde, without whom the membrane transfer printing would not have been possible.

I would like to thank all of the technical staff of the Zepler Institute and in particular the cleanroom staff of the Southampton Nanofabrication Centre. Without you, I would not have made a single device.

The great opportunity I have had to work in outreach and public engagement during my PhD has been of tremendous importance. On many occasions, it has often helped remind of the motivation behind my work and served to renew my enthusiasm for research. I would like to thank everyone with whom I have worked alongside to communicate research over the years, but in particular Dr. Pearl John for sharing her great expertise in this area, and Dr. Andrei Donko for working through the silly ideas (we won't mention *Laser Sharks*).

I am indebted to all my friends, near and far, for their comfort and support throughout the last four years. I must give special recognition to the League of Average Gentlemen for reminding me that, no matter how confident or successful you may be, a humbling defeat is only a gameweek away.

I must thank the Engineering and Physical Sciences Research Council for providing the funding that has allowed me explore the last four years in research.

Last, but not least, I want to thank those closest to me: my parents, Gary and Julie, and my sister Ellen for their unconditional love and support. I cannot forget my partner Lucy, whose understanding and reassurance has been indispensable throughout my PhD studies.

"Nothing travels faster than the speed of light with the possible exception of bad news, which obeys its own special laws."
- Douglas Adams, Mostly Harmless

List of Figures

List of Tables

List of Symbols

a	Feature size of a grating
A, B	Output fields of a 2×2 MMI
$\mathbf{B}$	Magnetic flux density
c	Speed of light in a vacuum
$C_{c,s}$	Cognitive/social rate, for PSO
d	Etch depth
$\mathbf{E}$	Electric field
$\mathcal{F}[f(x)]$	Fourier transform of the function $f(x)$
g	Gap size
G	Global best position, for PSO
h	Waveguide height
i	Imaginary unit, $\sqrt{-1}$
$\hat{\mathbf{i}}, \hat{\mathbf{j}}, \hat{\mathbf{k}}$	Unit vector in the x, y or z-direction
I	Pixel intensity
$\mathbf{k}$ (or k)	Wavevector (or wavenumber)
$\mathbf{K}$	Grating vector
L	Waveguide length
M	An integer number
m	Index for sequential terms (e.g. the m^{th} term)
N	An integer number
n	Real refractive index of a material
n_{eff}	Effective index of a mode
P	Optical power
r	Position
R	Particle best position, for PSO
s	Slab thickness
T	Bottom cladding thickness
$\mathbf{T}$	Transfer matrix
t	Time
v_{ph}	Phase velocity of a wave
w	Waveguide width
W	Inertial weight for PSO

x, y, z	Cartesian co-ordinates
α	Propagation loss
β	Propagation constant of a mode
γ	Constant for the transcendental waveguide equation
δ	Figure of merit term for the PSO
ε	Permittivity of a material (ε_0 for free-space)
θ	Geometric angle
Θ	Figure of merit term for the PSO
κ	Constant for the transcendental waveguide equation
Λ	Period of a grating
λ	Wavelength
μ	Permeability of material (μ_0 for free-space)
ν	Frequency
π	Ratio of a circle's circumference to its diameter
ρ	Charge density
σ	Conductivity
τ	Random number
ϕ	Optical phase
ζ	Indicator for TE/TM polarisation
ω	Angular frequency
∇	The vector differential operator, del

List of Acronyms, Initialisms and Abbreviations

Al$_2$O$_3$	Aluminium oxide (also sapphire)
BOX	Buried oxide
CdTe	Cadmium telluride
CMOS	Complementary metal-oxide-semiconductor
CTE	Coefficient of thermal expansion
DC	Direct current
EME	Eigenmode expansion
ESM	Endlessly single mode
FIB	Focussed ion beam
FDTD	Finite-difference time-domain
FEXEN	Fourier expansion simulation enviroment
Ge	Germanium
GON	Germanium-on-silicon nitride
GOS	Germanium-on-silicon
HF	Hydrofluoric acid
ICL	Interband cascade laser
ICP	Inductively coupled plasma
IPA	Isopropyl alcohol
MIR	Mid-infrared
MMI	Multimode interferometer
MZI	Mach-Zehnder interferometer
NIR	Near-infrared
NOI	Silicon nitride-on-insulator
NMP	N-methyl-2-pyrrolidone
OPO	Optical parametric oscillator
PIC	Photonic integrated circuit
PCF	Photonic crystal fibre
QCL	Quantum cascade laser
RIE	Reactive-ion etching
RF	Radio frequency

SEM	Scanning electron microscope
Si	Silicon
SiC	Silicon carbide
SiGe	Silicon-germanium alloy
SiN$_x$	Silicon nitride
SOI	Silicon-on-insulator
SON	Silicon-on-silicon nitride
SOS	Silicon-on-sapphire
SWG	Subwavelength grating
TE	Transverse electric
TM	Transverse magnetic
ZnSe	Zinc selenide

Data Access Statement

All data supporting the work presented in this thesis are openly available from the University of Southampton repository at https://doi.org/10.5258/SOTON/D2012. The data includes all simulation and experimental required to obtain the figures in this thesis.

In addition, two sets of MATLAB scripts are included. The first can be used to obtain the heatmaps of the intensity profiles of the outputs of the endlessly single-mode, single-mode and multimode waveguides, as in Figure 4.15 of Chapter 4. The second includes all functions written for the fitting of an analytical model of the transmission spectrum through a 2×2 Mach-Zehnder interferometer to the experimentally measured spectrum, by varying the insertion loss, imbalance and phase error of embedded 2×2 multimode interferometers, as in Chapter 5.

Chapter 1

Introduction

At a very general level, group IV photonics is the utilisation of a planar stack of thin films (made from elements in group IV of the periodic table) on a silicon wafer to create devices that use visible, near-infrared (NIR) or mid-infrared (MIR) light, depending on the materials used. By processing the thin films to make micrometre or nanometre-scale structures, the light can be confined, guided and manipulated to enable a wide range of applications. Devices that use these types of structures are known as photonic integrated circuits (PICs), which are analogous to the integrated circuits on planar semiconductors of the microelectronics industry.

1.1 Silicon photonics

Integrated photonics in silicon has attracted increasingly great research and industrial interest over the past few decades for a range of applications at NIR wavelengths ($780 - 2000$ nm, but particularly the telecommunications wavelengths of 870 nm, 1310 nm and 1550 nm), such as: short-reach optical interconnects [1], LiDAR and 3D imaging [2]; signal processing [3]; computing [4]; gyroscopes [5]; wireless optical communications [6]; and on-chip refractometers for gas sensing [7].

   The popularity of silicon photonics, despite it not being an ideal photonic material for modulators, detectors or sources, has endured because it is poised to use the mature production lines of the microelectronics industry, that are used for making complementary metal-on-semiconductor (CMOS) devices [8, 9]. This will enable high-volume production at low cost, for a cost efficiency that is hard to rival with other technologies.

   By extending silicon (and other group IV) photonic devices out to MIR wavelengths, additional applications can be unlocked that are specific to this range (discussed in the next section). However, the same advantages in terms of

cost-efficient fabrication hold true for group IV photonic devices at MIR wavelengths, provided that they can remain compatible with CMOS processes.

1.2   Mid-infrared wavelength range
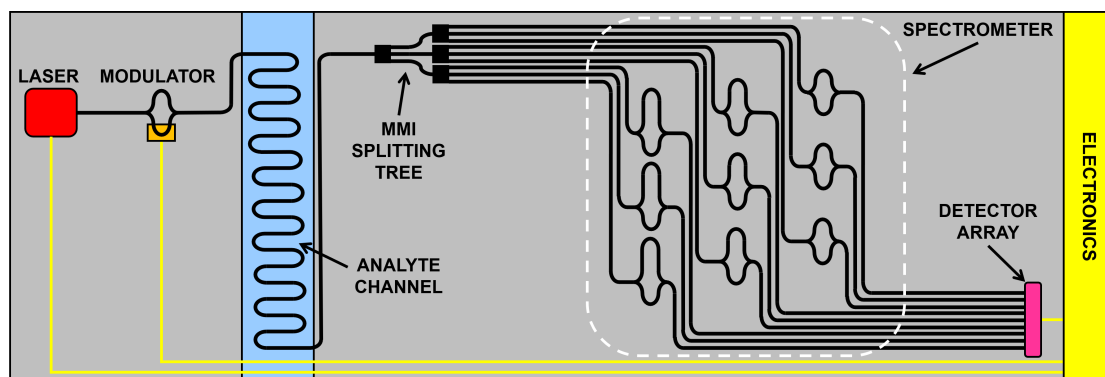
FIGURE 1.1: Illustration of an integrated photonics chemical sensing device. The black lines indicate waveguides, while the yellow lines are electrical connections.

One of the largest attractions to implementing integrated photonics in the MIR wavelength range is the opportunity to miniaturise MIR spectroscopy setups (that typically require bench-top scale apparatus) into a photonics circuit for a "lab-on-a-chip" [10, 11]. An illustration of the type of device that could be made is shown in Figure 1.1, where all of the functionality of a bench-top optical setup is available on a single chip. Light is provided by an broadband integrated laser source that could be an array of discrete lasers, a tunable source [12] or utilise supercontinuum generation [13]. The signal is then modulated to improve the signal-to-noise ratio and sensitivity [14], acting similarly to the chopper in the experimental setups in Chapter 3. The waveguides are exposed to an analyte over a particular channel and the spectral analysis performed using the spectrometer section of the circuit. The signal is then recovered using an on-chip detector array at the output and passed to the electronics, that are integrated on the same chip.

The control and analysis of a liquid-phase analyte can be implemented by combining the PIC with a microfludic chip, as shown in Figure 1.2(a). Taking measurements of liquids is a key method, for example, for enabling direct analysis of biofluids in healthcare applications. To illustrate how this might work, consider the identification of a drug in a blood sample. The blood is dropped onto the sample input, which consists of a paper pad to act as a filter to remove any biological cells in the sample and leave only blood plasma and any solutes, including the analyte. The sample is drawn along the microfluidic channel by capillary forces until it meets a region where the dissolved chemicals can diffuse through an analyte-permeable membrane into a different solvent in an underlying channel. The water in blood is highly absorbing in the MIR range so an organic solvent with a greater transparency is

FIGURE 1.2: Prospective devices using combined PICs and microfluidics: (a) Diagram of a potential method to combine a PIC with microfluidics; (b) PIC-based therapeutic drug monitoring device, that uses an insertable disposable microfluidics/PIC device (reproduced with permission from David J. Rowe).

used to increase the signal-to-noise ratio of the device. This solvent is then passed over the PIC to interact with the waveguide for spectroscopic analysis. In this example, everything required is integrated in a single device, but alternatively the chip could be made to be disposable. In this case, external electronics and light sources would be used, which may be more suitable to some applications where contamination or potentially hazardous biological material may be used (for example, in medicine). An example of such a device is shown in Figure 1.2(b).

In principle, such sensors could be made at high-volume and low-cost by leveraging silicon photonics for a wide variety of biological and chemical sensing applications, such as medical diagnostics or environmental monitoring. Lightweight, portable and cheap sensors could enable rapid and reliable chemical testing independent of the location and access to facilities, potentially revolutionising decentralised point-of-need testing. In context of the time in which this thesis is being written (during the COVID-19 global pandemic), the exigency for point-of-need testing with widespread availability and fast throughput has only been further emphasised.

Spectroscopy is a powerful tool at MIR wavelengths because many chemicals have uniquely identifiable absorption spectra in this region due to molecular vibrations [15]. At shorter wavelengths, MIR spectroscopy can provide information on the functional groups in a molecule, as shown in Figure 1.3. In the "fingerprint region" of $7 - 20\,\mu$m, there are many absorption features that are difficult to assign to a particular functional group, but the absorption spectrum in this region is unique to the given chemical, so it can be identified with comparison to reference data.

From Figure 1.3, it is evident that a considerable wavelength range of the MIR region must be covered to distinguish the bond absorptions of different functional groups, identify broad absorptions or be able to investigate the important fingerprint region. However, many recent sensing demonstrations in integrated photonics have
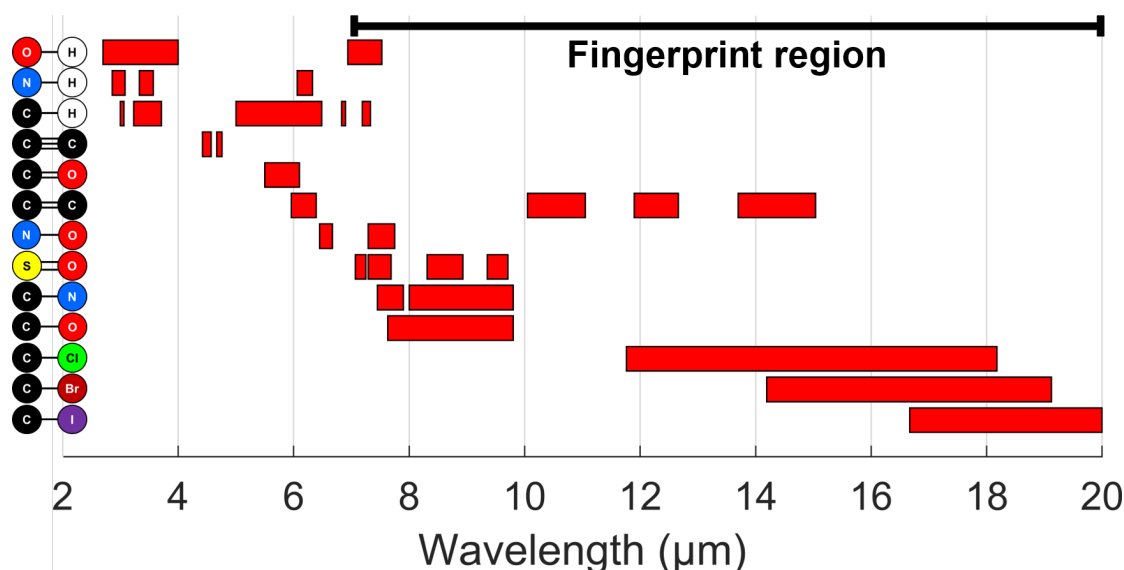
FIGURE 1.3: Absorptions in the mid-infrared wavelength range due to common molecular bonds. The multiple absorptions of each bond type will correspond to different types of vibration and the functional group, which is not specified here. The fingerprint region shown contains absorption features that are unique to a given compound. All data obtained from [16].

focused on small wavelength ranges for specific absorptions [17–20] or have used multimode waveguides [21], both of which have issues to address. For the latter, the multimode waveguides are problematic to use in a practical PIC as they prevent predictable optical behaviour (this will be discussed further in Chapter 4). For the former, focusing on a narrow wavelength range can only be used for a very controlled environment where the outcome of the test is binary; the chemical, which has an absorption at this wavelength, will either be there or not be there, and there will be no other chemical that could have an overlapping absorption feature.

To illustrate, consider an example: an integrated photonics test has been developed for cocaine using narrowband MIR spectroscopy, similar to the system in [24]. In this hypothetical scenario, the absorption peak at a wavelength of 5.8 μm, Figure 1.4(a), is being used to identify the presence of cocaine; it is a reasonably strong peak with high transmission on either side in the spectrum, so easily identifiable with a narrow bandwidth. In a laboratory environment where the chemicals present are known and quantifiable, the test is simple and the positive or negative outcome of the test for cocaine obtained from whether or not there is an absorption at 5.8 μm. However, in most situations where identifying the presence of cocaine would be desired (for example, in accident and emergency department or roadside police testing), the testing conditions cannot be controlled. There may be multiple substances present including the common and widely available painkiller paracetamol, which has a similar absorption at 5.8 μm, which can be seen in Figure 1.4(b). Evidently, it would be fairly likely that a false positive could be obtained on the test, which makes it at best unreliable and useless at worst. However, if the hypothetical testing device

FIGURE 1.4: MIR transmission spectra for (a) cocaine [22] and (b) paracetamol [23]. The inserts show the molecular strucutre of each chemical.

could work across a wider wavelength range (such as $\lambda = 5 - 10\,\mu m$), then the two absorption spectra are dramatically different and are easily distinguishable.

Although this exact hypothetical scenario is very specific, the underlying reasons for the problems in the test would arise in many applications. In the cocaine and paracetamol example, both have $C\!=\!\!=\!O$ functional groups and it is the stretching vibration of this bond that causes the absorption discussed above (the cocaine peak has a slightly split peak because it has two $C\!=\!\!=\!O$ groups with different adjacent atoms in the molecule). The absorptions in the MIR range correspond to specific bending and stretching vibrations of the molecular bonds, chemicals that have functional groups in common will have absorptions at similar points in the transmission spectrum. Therefore, multiple absorption peaks (and therefore a wider operating wavelength) need to be measured to distinguish similar molecules. Furthermore, as stated above, the fingerprint region holds little information at specific peaks (in narrow wavelength ranges) but is incredibly useful when the entire range is assessed. Expanding the working bandwidth of integrated photonics devices is essential for practical MIR spectroscopy applications.

Sensing capability is not the only driver to have a larger accessible wavelength range in MIR silicon photonics. In particular, there are opportunities for integrated photonics to have a large impact in application areas such as astronomy [25, 26], free-space communications [27] and infrared countermeasures for defence [28]. For each of these applications, MIR photonic devices need to have a working spectral bandwidth that can cover one or both of the the atmospheric transmission windows at $3 - 5\,\mu m$ and $8 - 14\,\mu m$, shown in Figure 1.5, to maximise their usefulness.

FIGURE 1.5: Atmospheric absorption spectrum calculated with data from the *HITRAN* database [29], using *HITRAN on the Web* [30], over 1 nautical mile (1.85 km). The red lines indicate the atmospheric transmission windows at $3 - 5\,\mu\text{m}$ and $8 - 14\,\mu\text{m}$.

1.3   Thesis outline

The work detailed in this thesis is intended to form part of a library of core building blocks for group IV integrated photonic circuits that could be used over a broad wavelength range in the MIR. The various components are developed and demonstrated, potentially providing more functionality. In particular, this work focuses on the passive components in the device, such as the waveguides and power splitters critical for analyte channel, splitting tree and spectrometer parts of the sensing device shown in Figure 1.1.

Chapter 2 is a review of existing literature, which consists of two parts: a discussion of recent results for the various passive device components at MIR wavelengths; and an evaluation of materials suitable for use in a MIR photonic device. In both cases, the background theory and working principle for each component is explained to contextualise the discussion.

Chapter 3 covers simulation, design, fabrication and characterisation techniques that were used within this work and were common to all devices.

Chapter 4 details the development of waveguides that are single-moded over an octave of frequency. Single-mode waveguides are critical for predictable circuit behaviour and are therefore needed for wideband PICs.

Chapter 5 shows the work on a multimode interferometer that can act as a power splitter or coupler within a PIC. The device was a subwavelength structure, that enables low imbalance, insertion loss and phase error over a 700 nm bandwidth centred on the MIR wavelength of 3.4 μm.

Chapter 6 shows an alternative type of broadband power splitter: a Y-junction with an insertion loss of 0.2 dB over a $2.8 - 3.7\,\mu$m wavelength range and a fabrication-tolerant design.

Chapter 7 discusses the work undertaken to develop a new material platform, by combining group IV waveguides with chalcogenides with a CMOS-compatible back-end process. Chalcogenides have high transparency at MIR wavelengths, but also offer sufficient refractive index contrast with group IV materials that could potentially realise compact devices with low propagation loss.

Chapter 8 concludes this thesis and projects possible future research directions for the work detailed in previous chapters.

Chapter 2

Literature Review

2.1 Passive photonic integrated circuit components

In this section, a selection of components commonly used in the passive integrated photonics is presented, to give the theoretical background on the components and to provide context that will be relied upon in later chapters.

2.1.1 Waveguides

Waveguides are the core underlying component of all photonic integrated circuits. To understand discussions of more complex waveguide designs and other components in general, it is important to discuss the underlying physics and behaviour of simple waveguiding structures first.

All forms of electromagnetic radiation propagating through some medium are subject to Maxwell's equations for electromagnetism. These equations describe the relationship between electric field $\mathbf{E}$ and the magnetic flux density $\mathbf{B}$ in the medium with charge density ρ, conductivity σ, permittivity ε and permeability μ:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon} \qquad \text{Gauss' law;} \qquad (2.1a)$$

$$\nabla \cdot \mathbf{B} = 0 \qquad \text{Gauss' law for magnetism;} \qquad (2.1b)$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t} \qquad \text{Faraday's law of induction;} \qquad (2.1c)$$

$$\nabla \times \mathbf{B} = \mu \left(\sigma \mathbf{E} + \varepsilon \frac{\partial \mathbf{E}}{\partial t} \right) \qquad \text{Ampère's law.} \qquad (2.1d)$$

Note that the change in **E** and **B** with respect to both time t and space are considered, since the operator ∇ is given by

$$\nabla = \left(\hat{\mathbf{i}} \frac{\partial}{\partial x} + \hat{\mathbf{j}} \frac{\partial}{\partial y} + \hat{\mathbf{k}} \frac{\partial}{\partial z} \right) \tag{2.2}$$

with $\hat{\mathbf{i}}, \hat{\mathbf{j}}$ and $\hat{\mathbf{k}}$ as the unit vectors in x, y and z respectively.

Assuming a charge-free, non-magnetic and dielectric material (such that $\rho = 0$, $\mu = \mu_0$, and $\sigma = 0$), it is relatively straightforward to derive the equation for a propagation of a light wave in terms of **E** or **B**:

$$\nabla^2 \mathbf{E} + \nabla \left(\frac{1}{n^2} \nabla n^2 \mathbf{E} \right) = \frac{1}{v_{\text{ph}}^2} \frac{\partial^2 \mathbf{E}}{\partial t^2}; \tag{2.3}$$

$$\nabla^2 \mathbf{B} + \frac{1}{n^2} \nabla n^2 \times (\nabla \times \mathbf{B}) = \frac{1}{v_{\text{ph}}^2} \frac{\partial^2 \mathbf{B}}{\partial t^2}, \tag{2.4}$$

where $v_{\text{ph}} = 1/\sqrt{\mu \varepsilon}$ is the phase velocity of the wave and $n = c/v_{ph}$ is the refractive index (c is the speed of light in free-space, $c = 1/\sqrt{\mu_0 \varepsilon_0}$). In homogeneous media, $\nabla n^2 = 0$ and so Equations 2.3 and 2.4 reduce to the well-known wave equations:

$$\nabla^2 \mathbf{E} = \frac{1}{v_{\text{ph}}^2} \frac{\partial^2 \mathbf{E}}{\partial t^2}; \tag{2.5}$$

$$\nabla^2 \mathbf{B} = \frac{1}{v_{\text{ph}}^2} \frac{\partial^2 \mathbf{B}}{\partial t^2}. \tag{2.6}$$

However, to manipulate light in useful way, an inhomogeneous structure is required to enable the guiding of a wave: this type of structure is known as a waveguide. The simplest waveguide structure is the configuration of materials show in in Figure 2.1, the planar waveguide. In this case, the refractive index only varies in the y-direction, $n = n(y)$, the structure extends infinitely in the x-direction and the

FIGURE 2.1: Planar Waveguide

light propagates in the z-direction. The three layers have refractive indices n_t, n_w and n_b for the top cladding layer, waveguiding layer and bottom cladding layer respectively (for a symmetric waveguide, the cladding layers will be the same material, with $n_t = n_b$). Light can be confined in the central waveguiding layer if it has a larger index than the cladding layer, so $n_w > n_t, n_b$. The confinement of light within the waveguiding layer will depend on the difference between the refractive indices of the waveguide and cladding layers, Δn. For larger values of Δn, the light will be more tightly confined, while the reverse is true for smaller values of Δn.

Light within the waveguide is considered as having transverse electric (TE) or transverse magnetic (TM) polarisation, or a combination thereof. For TE-polarisation, the electric field is parallel to the layers in the planar waveguide and zero in the other directions (so $\mathbf{E} = \mathbf{E_x}, \mathbf{E_y} = \mathbf{E_z} = 0, \mathbf{B_x} = 0$), while for TM-polarisation the magnetic field is parallel to the layers (so $\mathbf{B} = \mathbf{B_x}, \mathbf{B_y} = \mathbf{B_z} = 0, \mathbf{E_x} = 0$). Considering the TE-polarisation, the electric field is then given by

$$\mathbf{E}(\mathbf{r}, t) = E_x(y) \exp[\mathrm{i}(\omega t - \beta z)]\hat{\mathbf{i}}, \tag{2.7}$$

with angular frequency ω and propagation constant β.

The behaviour of guided light within a waveguide can be derived using Equations 2.3 and 2.7; the full derivation is well-known so will not be shown here, but the reader is directed to [31] for the full treatment of the mathematics. By combining these equations the TE wave equation can be found (likewise for TM-polarisation):

$$\frac{\mathrm{d}^2 E_x(y)}{\mathrm{d}y^2} + \left[k_0^2 n^2(y) - \beta^2\right] E_x(y) = 0; \tag{2.8}$$

where k_0 is the free-space wavenumber and is defined by $k_0 = 2\pi/\lambda_0$, for free-space wavelength λ_0. The propagation constant β is related to k_0 by $\beta = k_0 n_{\mathrm{eff}}$, where n_{eff} is the effective refractive index of the guided mode (the distribution of electric and magnetic fields for light guided by the waveguide) and $n_w > n_{\mathrm{eff}} > n_t, n_b$. Modes for which $n_{\mathrm{eff}} < n_t, n_b$ can exist, but they will radiate out of the waveguide as they propagate.

For a waveguiding layer of thickness h, the solutions of Equation 2.8 for each $n(y)$ can be assumed to have the form

$$E_x(y) = \begin{cases} A \exp[-\gamma_t y] & y \geq 0 \\ B \exp[\mathrm{i}\kappa_w y] + C \exp[-\mathrm{i}\kappa_w y] & 0 > y > -h \\ D \exp[\gamma_b(y + h)] & y \leq -h \end{cases} \tag{2.9}$$

where A, B, C and D are determined by the boundary conditions of the structure (namely that E_x and $\mathrm{d}E_x/\mathrm{d}y$ should be continuous over each of the interfaces between the different materials to satisfy Gauss' and Faraday's respective laws) and γ_t, κ_w and

γ_b are positive, real constants:

$$\gamma_t^2 = \beta^2 - k_0^2 n_t^2; \tag{2.10a}$$

$$\kappa_w^2 = k_0^2 n_w^2 - \beta^2; \tag{2.10b}$$

$$\gamma_b^2 = \beta^2 - k_0^2 n_b^2. \tag{2.10c}$$

By substituting Equation 2.9 into Equation 2.8 for each of the three waveguide regions and, using the boundary conditions of the structure, a transcendental equation for the TE-polarisation can be obtained:

$$\tan(\kappa_w h) = \left[\frac{\gamma_t}{\kappa_w} + \frac{\gamma_b}{\kappa_w} \right] \Big/ \left[1 - \frac{\gamma_t}{\kappa_w}\frac{\gamma_b}{\kappa_w} \right]. \tag{2.11}$$

Likewise, a transcendental equation can also be found for TM-polarisation:

$$\tan(\kappa_w h) = \left[\frac{\gamma_t}{\kappa_w}\left(\frac{n_w}{n_t}\right)^2 + \frac{\gamma_b}{\kappa_w}\left(\frac{n_w}{n_b}\right)^2 \right] \Big/ \left[1 - \frac{\gamma_t}{\kappa_w}\frac{\gamma_b}{\kappa_w}\left(\frac{n_w^2}{n_t n_b}\right)^2 \right]. \tag{2.12}$$

For a given polarisation, the solutions of the corresponding equation will yield a value β and thus the field amplitudes can be calculated in all regions of the waveguide. Since there are only a finite number of solutions to these eigenvalue equations, there will be a finite number of values for β.

Important physical conclusions can be extracted from the mathematics. Firstly, each value of β corresponds to a unique distribution of the electric and magnetic fields: a mode of propagation. Since there are a finite number of values for β, only discrete modes are allowed, each with a different effective refractive index n_{eff}. Where there is only one value for β, the waveguide is referred to as being single-moded; likewise, a waveguide with multiple values of β is multimoded. If there are no solutions to the transcendental equation, then there are no guided modes in the waveguide.

Secondly, from Equation 2.9, in addition to the electric field (and therefore optical power) confined within the waveguide, the field also continues outside of the waveguide and into the cladding layers. The field decays exponentially with the distance from the waveguide, so it is known as the evanescent field, and can be both an advantage and disadvantage depending on the application of the device. If the cladding absorbs light at the wavelength of the waveguide mode, this overlap of the modal field with cladding can lead to optical loss and attenuation of the signal in the waveguide. Generally, this is undesired but it is also the mechanism by which MIR absorption spectroscopy can be implemented in a waveguide; the overlap of the evanescent field with an analyte (shown in Figure 2.2) can lead to a detectable change in intensity through the waveguide, which can be used to identify the analyte.

In practice, the planar waveguide is not useful in an integrated photonic circuit, as the light is only confined in one direction. By adding another dimension of confinement (such as a rib waveguide, Figure 2.3), the waveguide can realise practical

FIGURE 2.2: Illustration of the evanescent field sensing of an analyte (the field shown is for $n_b > n_t$, which is expected for a waveguide exposed to an analyte).

FIGURE 2.3: Cross-section of a rib waveguide.

devices but the modes can no longer be found analytically, with the exception of certain two-dimensional geometries (such as circular optical fibres, where the radial symmetry allows this). Consequently, the waveguide must be modelled numerically, which underlies the need for the simulation tools discussed in Chapter 3. Further, the modes within the waveguide will not be solely TE- or TM-polarised, but have a combination of the two polarisations. However, in integrated photonics, the waveguides are typically designed so that the TE and TM components of polarisation are weighted to be close to 100% or zero respectively (or vice-versa), making it a quasi-TE (or quasi-TM) mode. For the avoidance of confusion, note that in literature it is common to drop the "quasi" term; this will be followed through the remainder of this work.

Rib and strip waveguides are by far the most ubiquitous design used in group IV photonics; a cross-section of a rib waveguide can be seen in Figure 2.3. The underlying substrate for a typical wafer material stack is omitted from this diagram, but it is usually silicon of thickness $0.5 - 1$ mm to prevent the wafer bending and bowing; the

exact thickness depends on the size of the wafer. The various parameters concerned in rectangular waveguide design are the width of the waveguide core w, the core height h, etch depth d and slab thickness $s = h - d$; strip waveguides differ only from rib waveguides in that $s = 0$. The thickness of the lower cladding layer T is typically $2 - 5\,\mu\text{m}$ at MIR wavelengths, sufficient to optically isolate the waveguide and prevent leakage of power into the wafer substrate (note that this is dependent on wafer materials and the specific wavelength of operation).

For any given waveguide design, single mode propagation will only be supported over a limited range. This is evident when it is recognised that all of the variables in Equations 2.11 and 2.12 have a wavelength dependence and therefore the number of solutions (i.e. discrete values of β, or modes) will also be wavelength-dependent. Although Equations 2.11 and 2.12 are for planar waveguides, this also holds true for the rib and strip waveguides.

Figure 2.4 shows the effective indices of the first three waveguide modes for a silicon-on-insulator, $1.2\,\mu\text{m}$-wide and $500\,\text{nm}$-tall strip waveguide; when the effective index of each mode is less than the refractive index of the cladding (shown by the dotted red line), the mode is no longer guided. Since the value of β for each mode varies with wavelength (and therefore also its effective index n_{eff}), there is a
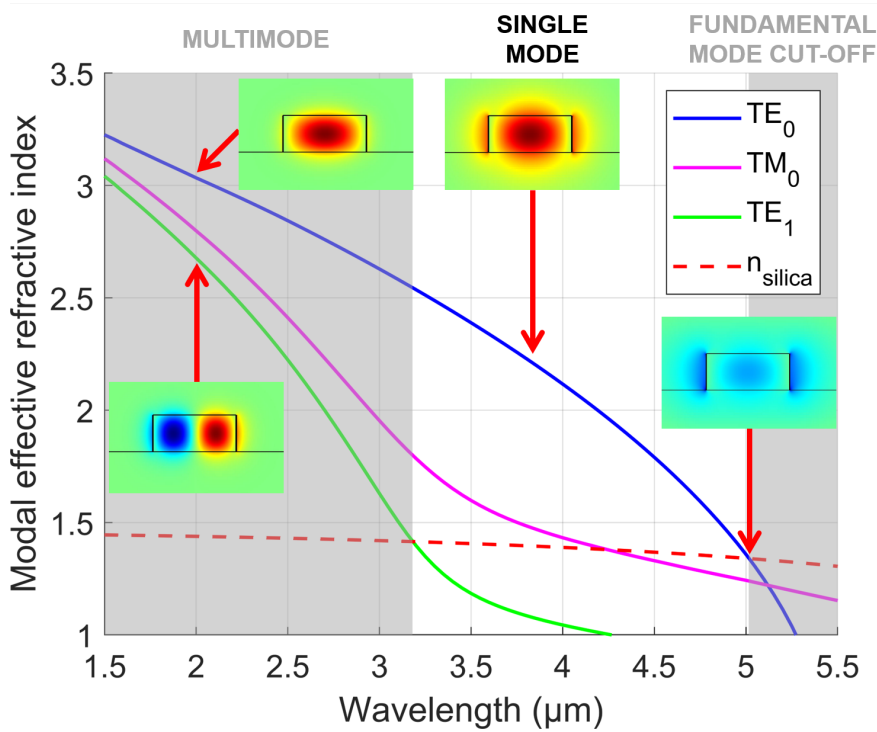
FIGURE 2.4: Effective indices of the first three supported modes in a silicon-on-insulator strip waveguide, with $1.2\,\mu\text{m}$ width and $500\,\text{nm}$ height, simulated in Lumerical MODE Solutions. Inserted images show the amplitude of the TE mode fields at the wavelengths indicated by the red arrows (red indicates positive amplitude, while blue shows negative). For effective index values below the refractive index of the SiO_2, the waveguide modes are radiative. The grey-shaded areas indicate wavelength regions where the waveguide is not single-mode.

wavelength range ($\lambda = 3.2 - 5\,\mu$m for this example) where only the fundamental modes are guided and outside this range the waveguide is either multimoded ($\lambda < 3.2\,\mu$m) or the fundamental mode is cut-off ($\lambda > 5\,\mu$m). Note that the support of both the fundamental TE/TM modes is still classed as the waveguide having single-mode propagation, as coupling into only the desired mode can be achieved either by polarisation control on-chip [32, 33] or by using external optics.

2.1.2   Propagation loss

When characterising a waveguide, one of the most important considerations is the propagation loss or attenuation; that is, the proportion of the input optical power that is absorbed, scattered or otherwise lost while propagating through a unit distance in a waveguide, α. Generally, it is useful to express the propagation loss of light through a waveguide of length L on a logarithmic-scale, so

$$\alpha L = -10 \log_{10} \left(\frac{P_{\text{out}}}{P_{\text{in}}} \right), \tag{2.13}$$

where P_{in} and P_{out} are the optical powers input into and output from the waveguide respectively. The propagation loss effectively limits the usefulness of a device. For example, considering a chemical sensing device, with a larger propagation loss it becomes more difficult (if not impossible) to distinguish an absorption peak, thereby reducing the sensitivity of a device. Specifically, the fractional change in optical power ΔP through a waveguide where there is also an analyte-induced loss α' is given by

$$\Delta P = \exp(-\alpha L) \cdot [1 - \exp(-\Gamma \alpha' L)] \approx \exp(-\alpha L) \cdot \Gamma \alpha' L \tag{2.14}$$

where Γ is the modal confinement factor in the analyte [11]. It is easy to see that, due to the exponential dependence on α in Equation 2.14, the waveguide propagation loss can quickly dominate and mask any absorption due to the analyte.

The propagation loss in a waveguide can be influenced by several mechanisms, but the primary ones considered here are absorptions from the waveguide and cladding materials, and the scattering induced by the roughness of the waveguide sidewalls. To illustrate this discussion consider Table 2.1, which shows a selective group of various waveguides used for MIR wavelengths, with differing group IV material combinations and geometries. The discussed material platforms are silicon-on-silicon dioxide (commonly known as silicon-on-insulator, or SOI), silicon-on-silicon nitride (SON), silion-on-sapphire (SOS), silicon nitride-on-insulator (NOI), germanium-on-insulator (GOI), germanium-on-silicon (GOS), germanium-on-silicon nitride (GON) and silicon-germanium (SiGe) graded index waveguides. Most of these waveguides had no upper cladding (i.e. the upper cladding was air), so the phrase "X-on-Y" denotes the X as the waveguide material

#	Waveguide	λ (µm)	α (dB/cm)	w (µm)	h (µm)	d (µm)	Ref.
A1	SOI	2	1.00 ± 0.08	0.6	0.24	0.1	[34]
A2		$2.12 - 2.28$	$0.45 - 0.65$	0.9	0.22	0.22	[35]
A3		3.39	0.6 ± 0.2	2	2	1.2	[36]
A4		$3.68 - 3.88$	$2 - 3$	1.2	0.4	0.4	[37]
A5		3.75	1.75 ± 0.22	1.2	0.4	0.24	[38]
A6		3.75	2.65 ± 0.08	1.2	0.4	0.4	[38]
A7		3.77	1.46 ± 0.20	1.35	0.4	0.22	[39]
A8		3.8	1.28 ± 0.65	1.3	0.5	0.5	[40]
A9		3.8	2.72 ± 0.57	1.1	0.5	0.5	[40]
B	SOI slot	3.8	1.4 ± 0.2	1.378	0.5	0.5	[41]
C	SON	3.39	5.2 ± 0.6	2	2	0.8	[42]
D1	SOS	2.08	$1.09 - 1.44$	$0.4 - 1$	0.28	0.28	[43]
D2		4.5	4.3 ± 0.6	1.8	0.6	0.6	[44]
D3		5.18	$1.69 - 1.92$	$0.4 - 1$	0.28	0	[43]
E	SOI pedestal	3.7	2.7 ± 0.15	8	5	n/a	[45]
F1	Suspended Si	2.75	3.0 ± 0.7	1.0	0.34	0.24	[46]
F2		3.8	3.4	1.1	0.5	n/a	[47]
F3		3.8	0.82	1.3	0.5	n/a	[48]
F4		7.67	3.1 ± 0.3	1.3	0.5	n/a	[49]
G1	NOI	2.65	0.16	4	2.5	2.5	[50]
G2		3.7	2.1	4	2.5	2.5	[51]
H1	SiGe	$4 - 5$	< 0.4	5.0	2.7	2.7	[52]
H2		4.5	< 1	3.3	3	3	[53]
H3		7.4	< 2	7.0	3	3	[53]
H4		$5.5 - 8.5$	$2 - 3$	4	13	4	[54]
H5		$5 - 11$	< 4.6	6.2	13	6.1	[55]
I	GOI	3.8	4.5 ± 0.5	1.1	0.515	0.265	[56]
J1	GON	3.8	3.35 ± 0.5	2	1	1	[57]
J2		3.73	7.86 ± 0.7	1	1.5	1.5	[58]
K1	GOS	3.8	0.58 ± 0.12	2.7	2.9	1.7	[59]
K2		5.8	2.5	2.9	2	2	[60]
K3		7.575	2.5	4.3	3	1.8	[61]
K4		$7.5 - 11.0$	< 5	4	2	1	[62]
K5		$9.0 - 11.5$	< 10	6	1.5	0.9	[63]
L1	Suspended Ge	2.155	5.4	0.9	0.3	0.15	[64]
L2		3.8	2.9	1.1	0.4	0.25	[56]
L3		7.67	2.6 ± 0.3	3.5	1	0.3	[65]
L4		7.7	5.3 ± 1.0	2.7	1	n/a	[66]

TABLE 2.1: Propagation losses of some waveguide geometries based on group IV materials at different wavelengths in the MIR wavelength range, using the geometrical parameters shown in Figure 2.3.

and Y as the lower cladding, but the reader is referred to the given references for each specific waveguide design.

The roughening of the waveguide sidewalls is introduced when the waveguide is fabricated by vertical etching [67], as described in Chapter 3, and it can cause light to be scattered out of the waveguide. This scattering can be reduced by increasing the

confinement of the mode within the waveguide, as there is less overlap of the mode with the sidewalls [68, 69]. This is evidenced by considering waveguides A8 and A9 in Table 2.1, which only differ in terms of the waveguide width; the mode is more confined in the silicon waveguide for the larger design, resulting in ~ 1.5 dB/cm less propagation loss compared with the smaller waveguide. Alternatively, the scattering can also be reduced by simply decreasing the height of the sidewalls. This explains the difference in propagation loss between waveguides A5 (a fully etched strip waveguide) and A6 (a partially etched rib waveguide), as the sidewalls are smaller if some of the slab remains. Light can be scattered of the waveguide surface as well; the propagation loss of waveguide A3 was reduced to 0.6 ± 0.2 dB/cm from 1.9 ± 0.2 dB/cm when the surface was oxidised with 30 nm of thermal oxide to reduce the surface roughness. It should be noted that the scattering is proportional to λ^{-4} [70], as expected from Rayleigh scattering, so the sidewall roughness becomes less dominant as a loss mechanism for longer wavelengths.

The material absorption is not just determined by the absorption from the waveguide material, but from the cladding layers as well. It was shown above that an evanescent field exists within the waveguide cladding (Equation 2.9). This can mean even if the waveguide core is transparent at a particular wavelength, absorption from the cladding can lead to a considerable propagation loss. For example, for a SOI waveguide, the silicon dioxide cladding has significant absorption for wavelengths longer than 4 µm. Using different cladding materials (such as SON or SOS, waveguides C-D3 in Table 2.1) can extend the low-loss propagation to longer wavelengths, but ultimately the cladding absorption will limit the operable wavelength range. Likewise, germanium-based waveguides (transparent up to ~ 14 µm [71]) also suffer increased propagation losses due to cladding absorption. GOS waveguides (K1-K5 in Table 2.1) can be low loss up to ~ 8 µm until multiphonon absorption begins in silicon [72]. Waveguide K4 in particular has a region of low loss $(9.5 - 11$ µm) but considerably larger losses outside, following the absorption spectrum for silicon.

Alternatively, appropriate engineering can minimise the overlap of the waveguide mode and the cladding by increasing the confinement. GOI and GON waveguides (I and J1-J2 in Table 2.1) have claddings that absorb at shorter wavelengths in comparison to silicon, but offer a considerably higher refractive index contrast with germanium than silicon does. This will increase the confinement and thereby decrease the loss, although so far better losses than GOS have not been achieved. In spite of this, higher index contrast layers may be useful for increasing the modal overlap with an analyte; the waveguide J2 (air-clad GON) was predicted as having a threefold increase in the evanscent field fraction present in air compared to a GOS waveguide with a similar design.

Graded-index SiGe waveguides (H1-H5 in Table 2.1) use an alloy that changes composition gradually from high silicon concentration to high germanium, as

opposed to the step index change between the cladding and waveguide layers in a conventional rib or strip waveguide. These comparatively large waveguides can have a tailored mode confinement, in principle minimising silicon absorption by preventing overlap with the mode and a pure silicon layer. However, the alloy means that there will inherently always be some overlap with silicon and the waveguides in [55] had silicon-related absorptions, particularly around 9 μm. Further, the SiGe waveguides are very large compared the other waveguides in the table, which brings other issues into consideration, such as the large footprint of a potential circuit made from these waveguides.

Other attempts at reducing the cladding absorption have been to simply reduce or remove it. Underetching of silicon and germanium waveguides for pedestal (waveguide E in Table 2.1) or suspended structures (waveguides F1-F4, L1-L4) have been successful at reducing the propagation losses and in particular have extended low-loss propagation to the end of the silicon transparency range. This is advantageous because the significant overlap of the evanescent field and the cladding is desirable for MIR absorption spectroscopy [18], where an analyte is introduced to a waveguide with no upper cladding layer and is identified by identifying its absorption peaks in the waveguide transmission spectrum. Suspended waveguides could potentially offer increased sensitivity, as an analyte could also be present beneath the waveguide, increasing the absorption from the analyte [20]. For designs that require high mode confinement with an analyte, a slot waveguide maybe more appropriate. For these waveguides, a gap is introduced in the centre of the waveguide, resulting in the light being tightly confined in this slot. Since the analyte will be present in this slot, this could greatly increase the device sensitivity due to significantly increased overlap of the waveguide mode and the analyte [73, 74], although at shorter wavelengths the increased sidewall-roughness-induced scattering from the slot might cancel out any benefit [75].

Other mechanisms can have a significant impact on the propagation loss and therefore must be considered in waveguide design. Combining crystalline materials with dissimilar lattice constants can cause defects at the material interface, which in turn can contribute to scattering loss [61, 76]. For example, for germanium deposited on silicon (i.e. GOS), typical densities of threading dislocation defects are 10^7 cm^{-2} [59, 77, 78]. As with other waveguide scattering, this scales by λ^{-4} so it is a smaller concern at longer wavelengths; in-depth discussion of this can be found in Section 2.2. Free-carrier absorption is caused by the presence of electrons and holes in a semiconductor, in addition to usual material absorption. It is generally considered negligible unless the material is highly doped (i.e. with high levels of free carriers), but for certain materials it is worth considering. For example, in germanium the free-carrier absorption is much higher than it is in silicon [79, 80] and it also scales significantly with wavelength. Consequently, for germanium waveguides at longer wavelengths, the free carrier absorption may result in a large or dominant loss

contribution, even at background doping concentrations.

2.1.3 Splitters and couplers

In a PIC-based sensing device, power splitting is essential to make realistic devices such as the one suggested in Chapter 1. As a simple example, a Mach-Zehnder inteferometer (a key compoment for many sensors, see Section 2.1.4) requires a splitter and a coupler to operate. There are several ways to approach power splitting, the most common and relevant to this work (Y-junctions, directional couplers, and multimode interferometers) are discussed in this section.

2.1.3.1 Y-junctions

Y-junctions (Figure 2.5) are the simplest splitter design, where the input waveguide is split into the two output waveguides with the same width. The output waveguides are then separated by a sufficient distance to minimise any coupling between them.

Their simplicity means that they are straightforward to design and their symmetry means that there fundamentally will not be any imbalance between the outputs. However, their significant drawback is their basic design requires the tip of the "Y" to be infinitely sharp to minimise the insertion loss, as the tip interacts with the centre of the waveguide mode where the optical power is greatest. For example, some optical loss at the tip can be seen in the simulation of a sub-optimal Y-junction shown in Figure 2.6. However, the resolution of waveguide fabrication limits the tip sharpness that is realistically achievable, which can lead to considerable insertion losses: for example, a loss per junction of 2.6 dB have been recorded in the NIR [81]. Additionally, small fabrication imperfections can also cause asymmetry in the junction, yielding imbalance in the outputs.

The fabrication limitations have been addressed in the NIR range. By setting the tip size to a minimum size allowed by fabrication (200 nm), the optimal geometry of the splitting region around the tip was found using a particle swarm optimisation, to

FIGURE 2.5: Illustration of a Y-junction as a waveguide power splitter.

FIGURE 2.6: Optical power through a Y-junction. Red areas indicate high intensity and blue areas indicate low intensity.

minimise the insertion loss to $0.28 \pm 0.02\,\text{dB}$ at $\lambda = 1.50\,\text{μm}$ [82]. Furthermore, the design has a higher tolerance to fabrication overall, meaning that it is less susceptible to the effects of any imperfections. A similar design has been used for a polarising beam splitter with an insertion loss less than $1.4\,\text{dB}$ over $\lambda = 2.45 - 2.55\,\text{μm}$ [83]. This approach is discussed further in Chapter 6.

A different approach to relax this fabrication requirement is to use very gradual changes in the waveguide geometry, giving a smoothly changing refractive index profile and not exciting any higher order modes that can cause wavelength dependence or loss. This technique is called adiabatic mode evolution and has been used for various waveguide configurations in the NIR [84–86]. Although different in principle to the Y-junction, the layout of the device can be arranged in a similar configuration for an identical effect (see Figure 2.7). In the MIR, adiabatic power splitters in a Y-junction configuration have been demonstrated by the input and output waveguides tapering down to a small tip, separated laterally by a small gap. The mode in the input waveguide is squeezed out by the decreasing width and couples into the output waveguides via intermediate modes propagating within the gaps. This splitter showed an insertion loss of $< 0.6\,\text{dB}$ over $3.66 - 3.89\,\text{μm}$ [87];

FIGURE 2.7: Illustration of an adiabatic power splitter, in a Y-junction configuration

importantly, the splitter was fairly tolerant to the smallest tip sizes of the taper over $30 - 120\,\text{nm}$ (the upper half of this range is achievable with *e*-beam lithography).

2.1.3.2 Directional couplers

A directional coupler (Figure 2.8) is a device in which two parallel waveguides are separated by a small gap g. Optical power can couple across from one waveguide to another in an arbitrary ratio $P_{\text{out,1}} : P_{\text{out,2}}$ that can be tuned with the design of the coupler. Typically, this is either 50:50 for a power splitter, or 0:100 for a cross-coupler.

FIGURE 2.8: Illustration of a directional coupler, with coupling length L_C and waveguide separation g.

The working principle behind the directional coupler is based on the sum of two supermodes (Figure 2.9) that propagate in the two waveguides of the coupling section with constants $\beta_{1,2}$. As the modes travel, they will have a relative phase difference; if the input is into the top waveguide only (so $P_{\text{in,2}} = 0$), when the modes are in phase the optical power will be localised in the top waveguide, while when they are in anti-phase the power will be in the bottom waveguide. The phase difference between the two modes, after a coupler length L_C, is

$$
\begin{aligned}
\beta_1 L_C - \beta_2 L_C &= \frac{2\pi n_{\text{eff,1}}}{\lambda} L_C - \frac{2\pi n_{\text{eff,2}}}{\lambda} L_C \\
&= \frac{2\pi \Delta n}{\lambda} L_C = 2C L_C
\end{aligned}
\tag{2.15}
$$

where $\Delta n = n_{\text{eff,1}} - n_{\text{eff,2}}$ is the difference between the two effective indices of the supermodes and the coupling coefficient $C = \pi \Delta n / \lambda$. The coupling coefficient (and also Δn) will decrease exponentially as the gap size g increases [88]. Δn can be found using numerical eigenmode solvers, as detailed further in Chapter 3.

FIGURE 2.9: The modal field for the supermodes in a directional coupler. Red indicates positive electric field amplitude, while the blue areas indicate negative amplitude.

FIGURE 2.10: Location of optical power in directional couplers of different lengths: (a) long cross coupler showing multiple oscillations of the optical power within the directional couplers; (b) cross coupling; (c) 3 dB power splitting. All couplers shown have the same waveguide separation, g. Red areas indicate high intensity and blue areas indicate low intensity.

Consequently, for a given geometry, selecting the desired power ratio can be achieved by designing the length of the coupler to give the correct phase difference at the output of the coupler. This is illustrated in Figure 2.10, where it can be seen that different coupler lengths will give different power ratios, for the same waveguides and value of g.

Importantly, the coupling coefficient C has an inherent wavelength-dependence. This means that, for the same device operating at a different wavelength, the phase difference between the supermodes will change and the power ratio will differ as a consequence. This limits the tolerable performance range with respect to wavelength of the standard directional coupler for broadband PICs. At $\lambda = 3.8\,\mu$m, this means that one reported 50:50 splitter had a bandwidth that is tens of nanometres (assuming a normalised power of 0.5 ± 0.05 in each output) [89].

The bandwidth of directional couplers can be improved by manipulating the waveguide dispersion to compensate for the wavelength dependence of $n_{\text{eff},1}$ and $n_{\text{eff},2}$, yielding an 85 nm bandwidth around 3.8 µm wavelength for 50:50 splitting and 145 nm bandwidth for 0:100 cross-coupling [90]. However, this necessitated a more complicated fabrication, with multiple etch steps on the same waveguide, and it still does not provide an operable bandwidth sufficient for sensing applications. Alternatively, inserting a subwavelength grating metamaterial into the gap of a conventional directional coupler has been demonstrated to provide a bandwidth of

175 nm at a wavelength of 3.80 μm (for 0:100 cross-coupling) [91]. This device uses the dispersion of the subwavelength metamaterial to compensate for the wavelength dependence of $n_{\text{eff},1}$ and $n_{\text{eff},2}$, but this can be achieved using the same etch step that is used to define the waveguides and therefore the fabrication is less complicated. Subwavelength structures are discussed further in Section 2.1.5.

2.1.3.3   Multimode interferometers

FIGURE 2.11: Illustration of a $M \times N$ MMI (in this case, $M = 3$ and $N = 4$). The multimode region of the MMI has a width w_{MMI} and length L_{MMI}.

A multimode interferometer (MMI) consists of a multimode waveguide, of width W_{MMI} and length L_{MMI}, with M input waveguides and N outputs (Figure 2.11). The multimode waveguide is wide enough to support a large number of modes, which are excited by the modal field launched from one or more of the input waveguides. The higher-order modes will interfere, resulting in the input mode being replicated along the length of the multimode region. As an illustration, consider the single-input MMIs shown in Figure 2.12, where the points at which a single replication of the input mode (the self-imaging point) and two replications (the two-fold imaging point) occur are highlighted. If L_{MMI} is equal to the distance from the input to the two-fold imaging point, the MMI can be used as a 1×2 splitter for example (or, in reverse, as a combiner for two inputs). It is also evident that the MMI could be used as a splitter or combiner for more than two waveguides choosing a suitable length to coincide with, for example, the three-fold imaging point, or as a 2×2 cross-coupler by using the self-imaging point with an off-centre input waveguide.

The position along the multimode region at which N replications of the input occur is determined by the beat length, L_π, between the two lowest order modes in the multimode region, given by [92]:

$$L_\pi = \frac{\pi}{\beta_0 - \beta_1} \approx \frac{4 n_{\text{eff,w}} w_{\text{e}}^2}{3\lambda},\tag{2.16}$$

where $\beta_{0,1}$ are the propagation constants of the two lowest order modes, $n_{\text{eff,w}}$ is the effective index of the waveguide and w_{e} is the effective width of the multimode

FIGURE 2.12: Intensity plots of the field in the multimode region of: (a) a symmetric MMI with a centred input; (b) an asymmetric MMIs with an offset input, with the self-imaging and two-fold imaging points shown by the dotted white line. Both MMIs have the same beat length L_π. The red areas indicate high intensity, and the blue areas show low intensity.

region. Since the waveguide mode has a non-negligible penetration into the waveguide cladding (see Section 2.1.1) w_e will be slightly larger than w_{MMI} and, for an effective cladding index $n_{eff,c}$, is given by [92]:

$$w_e = w_{MMI} + \left(\frac{\lambda}{\pi}\right) \left(\frac{n_{eff,c}}{n_{eff,w}}\right)^{2\zeta} (n_{eff,w}^2 - n_{eff,c}^2)^{-\frac{1}{2}} \tag{2.17}$$

where $\zeta = 0$ for TE-polarisation and $\zeta = 1$ for TM-polarisation.

Generally, for an $M \times N$ MMI and for propagation in the z-direction, N images will first occur when

$$z = \frac{3L_\pi}{N}, \tag{2.18}$$

so the self-imaging ($N = 1$) and two-fold imaging points ($N = 2$) will be $z = 3L_\pi$ and $z = 3L_\pi/2$ respectively. In the specific case of a symmetric MMI with a centred input

waveguide, the N images will first occur at

$$z = \frac{3L_\pi}{4N},$$
(2.19)

meaning that the self-imaging point will be at $z = 3L_\pi/4$ and two-fold imaging point $z = 3L_\pi/8$.

The quadratic dependence of the MMI beat length, and therefore also the entire MMI length, on the width of the multimode region is a significant consequence for the MMI design, as it can be made with a small footprint. In principle, this is only limited by having a sufficiently wide multimode region to support enough higher order modes. Moreover, the smallest feature size in the device is that of the input and output waveguides, meaning that it is fairly tolerant to fabrication processes. The modal mismatch between the thin single-mode input and output waveguides and the wider multimode region will cause optical loss, but this can be reduced by tapering the input and output waveguides [93], as illustrated in Figure 2.11.

The inherent wavelength dependence of the MMI can be exploited for the (de)multiplexing of signals (i.e. separating or combining a signal into different wavelength components), such as in an angled MMI [94] shown in Figure 2.13. This can also be useful for some sensing applications, for example, the different wavelengths of a broadband source can be split into separate channels for sensing different gases on the same device [19]. However, for broadband PICs, the wavelength dependence considerably limits its utility. Conventional silicon MMIs have been reported with losses as low as 0.1 dB/MMI and imbalance as low as 0.15 dB [39, 48, 95] at 3.8 μm, and 0.24 dB/MMI at 7.9 μm on GOS [61], but these all have a wavelength-dependent performance. MMIs based on the graded-index SiGe alloys can offer bandwidths up to an octave of frequency [96, 97], but these rely on the dispersion on the material itself and do not provide an approach to improve the wavelength range of an MMI more generally. Tunable graphene-based plasmonic MMIs can offer low imbalance over a 2 μm bandwidth and with a compact footprint,

FIGURE 2.13: Illustration of an angled MMI.

but they suffer from very high insertion loss (up to 10 dB per device) and a much more complicated fabrication process [98, 99].

Alternatively, engineering of the multimode region itself can yield devices with a broader operating bandwidth. In the NIR, MMIs in which a strip of a sub-wavelength metamaterial has been inserted has been shown to halve the length of a conventional MMI [100]. Significantly, an MMI in which the entire multimode region is replaced with this sub-wavelength metamaterial has been demonstrated with a 300 nm bandwidth around 1.55 μm [101]. In this design, the refractive index varies with wavelength to compensate for the change in beat length with wavelength; importantly, the dispersion is due to etched features and not the waveguide material, so this approach is applicable to different waveguide platforms (unlike the ultra-broadband SiGe MMIs). Subwavelength engineering is discussed further in Section 2.1.5, and the subwavelength grating MMI will be revisited in Chapter 5.

2.1.4   Mach-Zehnder interferometers

Mach-Zehnder interferometers (MZI) are useful components for PICs and are used for a variety of applications, such as spectrometers [102–104], refractive index sensors [7, 105, 106], modulators [107–109], switches [110, 111] or (as discussed in Chapter 5) for device characterisation [112].

In an MZI, input optical power is divided into two separate optical paths, called the arms of the MZI. The phase change of the light by propagating in each arm of length $L_{1,2}$ will be given by $\phi_{1,2} = \beta_{1,2}L_{1,2}$, where $\beta_{1,2}$ is the propagation constant in each arm. The light in both arms is then recombined and will interfere, with the output power depending on the phase difference $\Delta\phi = \phi_2 - \phi_1$ between the two arms. If we consider the transfer matrix of a 1×1 MZI (such as those in Figure 2.14), the electric field at the input and output ports, $\mathbf{E}_{\text{in}}$ and $\mathbf{E}_{\text{out}}$, of the MZI are related by its total transfer matrix $\mathbf{T}_{\text{MZI}}$, which in turn is the product of the transfer matrices of its constituent components:

$$\mathbf{E}_{\text{out}} = \mathbf{T}_{\text{MZI}}\mathbf{E}_{\text{in}} = \mathbf{T}_{50:50}\mathbf{T}_{\text{arms}}\mathbf{T}_{50:50}\mathbf{E}_{\text{in}} \tag{2.20}$$

where $\mathbf{T}_{50:50}$ is the well-known transfer matrix for an ideal 50:50 beamsplitter (any splitter may be used)

$$\mathbf{T}_{50:50} = \frac{1}{\sqrt{2}}\begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \tag{2.21}$$

and $\mathbf{T}_{\text{arms}}$ is the transfer matrix for arms of the MZI, with respective phases $\phi_1 = \phi$ and $\phi_2 = \phi + \Delta\phi$ and assuming no propagation loss

$$\mathbf{T}_{\text{arms}} = \begin{bmatrix} \exp[i(\phi + \Delta\phi)] & 0 \\ 0 & \exp(i\phi) \end{bmatrix} = \exp(i\phi)\begin{bmatrix} \exp(i\Delta\phi) & 0 \\ 0 & 1 \end{bmatrix}. \tag{2.22}$$

With a single input to the MZI, and multiplying through the transfer matrices:

$$\mathbf{E}_{\text{in}} = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \Rightarrow \mathbf{E}_{\text{out}} = \frac{\exp(i\phi)}{2} \begin{bmatrix} i[\exp(i\Delta\phi) + 1] \\ 1 - \exp(i\Delta\phi) \end{bmatrix}. \tag{2.23}$$

Since the optical output power is related to the output field by $P_{\text{out}} \propto |E_{\text{out}}|^2$, it is then straightforward to show that the normalised output power of the 1×1 MZI is

$$\frac{P_{\text{out}}}{P_{\text{in}}} = \cos^2\left(\frac{\Delta\phi}{2}\right). \tag{2.24}$$

It is this sinusoidal-like behaviour of MZIs that makes them, for example, useful for modulators and switches (although the latter requires more output ports) because small changes in the phase difference can make large changes in the output power.

The phase difference between the MZI arm can be induced in two different ways, depending on whether symmetric or asymmetric MZIs are used. In the case of the symmetric MZI (where $L_1 = L_2$), the phase is due to a change in the waveguide effective index Δn_{eff}, so that $\beta_2 = \beta_1 + \Delta\beta$. As shown in Figure 2.14(a), this is caused by changing the local refractive index by Δn over a section of the waveguide, which in turn will change the effective index of the waveguide mode. In silicon, the localised refractive index shift can be achieved by using thermo-optic or plasma dispersion effects [113], or by exposing a section of the waveguide to a different substance other than the top cladding (such as an analyte, for refractive index sensing techniques). In an asymmetric MZI, Figure 2.14(b), the phase difference is directly due to the

FIGURE 2.14: Illustration of 1×1 MZIs, using Y-junctions as power splitters. The pink box for the symmetric MZI indicates an area with a localised refractive index shift.

difference in the length of the optical path ΔL. In this case, the waveguides in each arm are kept the same, so $\beta_1 = \beta_2$. In either case, the phase difference between the two arms can be described similarly and is given by

$$\Delta\phi = \begin{cases} \frac{2\pi L}{\lambda}\Delta n_{\text{eff}} & \text{for a symmetric MZI} \\ \frac{2\pi n_{\text{eff}}}{\lambda}\Delta L & \text{for an asymmetric MZI.} \end{cases} \tag{2.25}$$

The free spectral range (FSR), $\Delta\lambda_{\text{FSR}}$, of a resonator such as an MZI is defined as the separation in wavelength between two extinctions (where $P_{\text{out}} = 0$) in the transmission spectrum in an MZI. From Equation 2.24, it is evident this occurs when $\Delta\phi = m2\pi$, where $m = 0, 1, 2...$ is an integer multiple. Considering the phase change with wavelength:

$$\left|\frac{\partial\phi}{\partial\lambda}\right| = \left|\frac{\partial}{\partial\lambda}\left[\frac{2\pi n_{\text{eff}}(\lambda)L}{\lambda}\right]\right| = \frac{2\pi L}{\lambda^2}\left[n_{\text{eff}}(\lambda) - \lambda\frac{\partial n_{\text{eff}}}{\partial\lambda}\right] = \frac{2\pi n_{\text{g}}L}{\lambda^2} \tag{2.26}$$

where n_{g} is the group index of the waveguide. Assuming an asymmetric MZI, then a relative phase change $\Delta\phi = 2\pi$ will be induced for light travelling a distance $L = \Delta L$, which yields

$$\Delta\lambda_{\text{FSR}} \approx \frac{\lambda^2}{n_{\text{g}}L}. \tag{2.27}$$

2.1.5   Subwavelength structures

Historically, the phenomenon of electromagnetic radiation interacting with periodic structures of different materials as if they were a single homogenous material has been known since the 19[th] century, with Hertz's experiments with radio wave and wire grids [114, 115]. For these structures, their periodicity is much smaller than the wavelength of light and so diffraction effects are suppressed, and the structure can be considered as an effective medium. In recent years, the engineering of silicon photonic devices on a subwavelength scale has been widely considered, as it allows selective manipulation of the device refractive index without the need of additional materials, and can be implemented in the same fabrication steps that are used to define conventional devices. Subwavelength structured devices have been used in a variety of passive components, such as waveguides [47, 116, 117], directional couplers [91], MMIs [100, 101], Y-junctions [118], polarisation splitters [119, 120], edge couplers [121] and grating couplers [122–126].

Figure 2.15(a) shows a periodic one-dimensional grating with period Λ and alternating features of size a and $(\Lambda - a)$ with refractive indices n_1 and n_2 respectively. The bandstructure of this structure is shown in Figure 2.15(b) where the bandgap occurs when $k = \pi/\Lambda$ (i.e. the edge of first Brillouin zone) with allowed bands above and below the bandgap. When the frequency ω is above the bandgap, the structure carries radiative modes (i.e. the light is diffracted out of the waveguide), while at
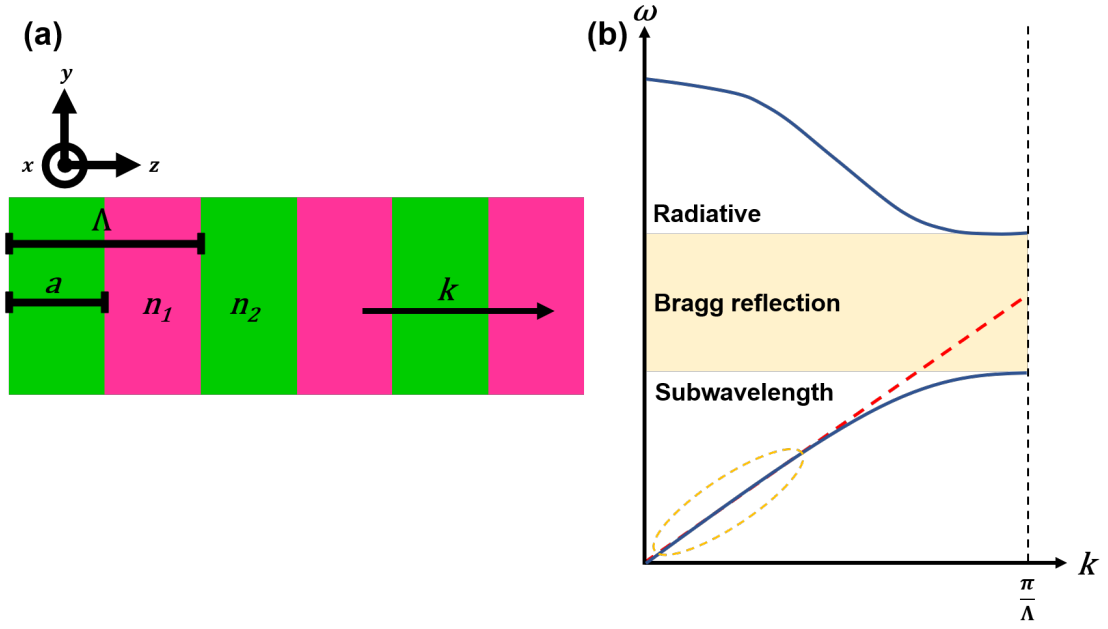
FIGURE 2.15: Illustration of (a) a one-dimensional dielectric periodic structure and (b) its resulting photonic bandstructure. For light propagating with wavevector k, propagating modes at frequency ω are shown in the bandstructure by blue lines. The yellow shaded region indicates the bandgap, with no propagating modes. The red dashed line shows the dispersion for a homogenous waveguide (with $n_1 = n_2$), and the region in which the periodic structure can be approximated as homogenous is shown by the dashed orange circle.

frequencies below the bandgap, the structure has guided modes. If $n_1 = n_2$, the lattice structure becomes a homogenous material and the bandgap closes, and we have the case of the conventional strip waveguide [127].

At frequencies within the photonic bandgap, no propagating mode exists (the field will exponentially decay within the structure) and the light is reflected back opposite to the direction of propagation. This is the Bragg reflection and, for a given structure, the wavelength at which it occurs is $\lambda = 2n_{\text{eff,BF}}\Lambda$ (note that n_{eff} is the effective index of a mode within the periodic structure, known as the Bloch-Floquet mode [114]). Therefore, the condition for the above structure to be subwavelength is

$$n_{\text{eff,BF}} < \frac{\lambda}{2\Lambda}. \tag{2.28}$$

At smaller values of k (equivalently at long wavelengths, $\lambda \gg \Lambda$) within the dashed orange circle in Figure 2.15(b), the grating can be considered as a homogeneous waveguide. The refractive index of this material will be polarisation-dependent (i.e. birefringent), with the index for light polarised parallel ($n_\parallel$) or perpendicular ($n_\perp$) to the periodic grating interfaces approximated by Rytov's

equations [128]:

$$n_\parallel^2 \approx \frac{a}{\Lambda} n_1^2 + \left(1 - \frac{a}{\Lambda}\right) n_2^2; \tag{2.29a}$$

$$\frac{1}{n_\perp^2} \approx \frac{a}{\Lambda} \frac{1}{n_1^2} + \left(1 - \frac{a}{\Lambda}\right) \frac{1}{n_2^2}. \tag{2.29b}$$

When considering the approximated homogeneous material as a whole, it is sometimes more convenient to describe its anistropic refractive index with a single tensor $\boldsymbol{n}_{SWG}$ [129]. By denoting the refractive index in each of the x-, y- and z-directions as n_{xx}, n_{yy} and n_{zz} respectively (using the same directions as in Figure 2.15), $\boldsymbol{n}_{SWG}$ is given by:

$$\boldsymbol{n}_{SWG}^2 = \begin{bmatrix} n_{xx}^2 & 0 & 0 \\ 0 & n_{yy}^2 & 0 \\ 0 & 0 & n_{zz}^2 \end{bmatrix} = \begin{bmatrix} n_\parallel^2 & 0 & 0 \\ 0 & n_\parallel^2 & 0 \\ 0 & 0 & n_\perp^2 \end{bmatrix}. \tag{2.30}$$

The power of subwavelength structuring of a dielectric waveguide is evident when considering Equations 2.29a and b; by simply designing a and Λ appropriately, the refractive index of the waveguide can be tuned between n_1 and n_2, provided the condition of Equation 2.28 is met.

2.1.6   Coupling light from an external source

Where on-chip sources have not yet been realised, or where it is more useful to have a device to be connected to a larger optical system (for example, for use with optical fibres in telecommunications or for experimental characterisation in a laboratory), the approaches for coupling light into and out of an integrated photonics device are important. Generally, these strategies can be broken down into two groups: in-plane (where the input or output light propagates parallel to the surface of the chip) or out-of-plane (where light is, or is close to, normal to the surface of the chip) [130]. In general, in-plane coupling is broadband and polarisation-independent but requires additional post-fabrication processing and consistent coupling losses can be hard to achieve due to difficulty in alignment, while in contrast out-of-plane coupling is narrowband and polarisation-dependent and has higher losses but has a relaxed alignment approach, requires no post-fabrication processing and can be implemented anywhere on a device.

2.1.6.1   In-plane coupling

For in-plane coupling, the end facet of the waveguide (perpendicular to the waveguide direction) needs to be first exposed with post-fabrication processing, such

FIGURE 2.16: Illustration of butt coupling, with light propagation indicated by yellow arrows (not to scale).

as by polishing or cleaving (see Chapter 3). Light propagating in the same plane is coupled into the waveguide by aligning an optical fibre with the waveguide (Figure 2.16) or a lens focussed on the end facet.

While simple in principle, these approaches have no inherent bandwidth limitations or polarisation dependence beyond that of the waveguide and fibre. However, the simplicity does mean that it is heavily reliant on external controls and requires precise alignment in three directions whether using a lens or fibre. This is a particular issue when reliable and reproducible alignment is required, for example assuming consistent insertion losses during optical characterisation. Further, low coupling loss also necessitates high-quality facets, which can be problematic to achieve consistently with cleaving or polishing.

When considering alignment with a fibre, the cross-sectional area of the fibre core is approximately 100 times larger than the cross-sectional area of a waveguide into which it is coupling. This means that there is a significant mismatch between the optical modes in the fibre and the modes in the waveguide and this modal mismatch leads to considerable losses in the coupling. To mitigate this, the waveguide can be tapered to the output facet, as in Figure 2.17(a), to increase the waveguide width and therefore the mode area; this approach also eases the alignment in the lateral direction. Alternatively, and somewhat counter-intuitively, a waveguide that tapers in the opposite direction, to decrease the waveguide width towards the facet as in Figure 2.17(b), have also been used as the mode becomes less confined as the width decreases
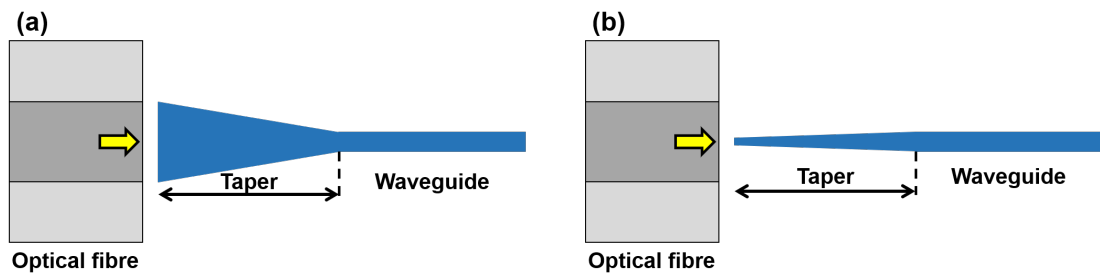
FIGURE 2.17: Illustration of waveguide tapering schemes for edge coupling: (a) wide waveguides at the chip facet; (b) narrow waveguides at the chip facet. The yellow arrows indicate the direction of light propagation.

and therefore the waveguide mode area increases [131–133]. This effect can be further enhanced using the sub-wavelength patterning to tune the waveguide index (and thus modal effective index) [121] or by considering the mode evolution in the vertical direction as well, with two-dimensional tapers [134] or multilayer structures [135].

2.1.6.2   Out-of-plane coupling

By far the most popular approach for the out-of-plane coupling is to use waveguide gratings, which can be achieved by a straightforward vertical etch into the waveguide. In the simplest implementation, consider an infinite one-dimensional grating in a waveguide, such as the one shown in Figure 2.18, with etch depth d, period Λ and unetched feature (or tooth) dimension a (the dimension of the etched features, or groove, is correspondingly $\Lambda - a$). Note that the etch depth of the grating is not necessarily the same as the etch depth used to define the waveguide.

FIGURE 2.18: Illustration of a waveguide grating coupling, with light propagation indicated by yellow arrows. Only the first diffraction order into the top and bottom cladding are shown.

When an optical mode propagating in the z-direction with constant β is present in the waveguide grating, the light will diffract according to the Bragg condition, here expressed as [130]:

$$\mathbf{k_{m,z}} = \beta - m\mathbf{K}, \tag{2.31}$$

where $\mathbf{k_{m,z}}$ is the z-component of the wavevector $\mathbf{k_m}$ for the m^{th} diffraction order. $\mathbf{K}$ is a vector used to describe the grating and is solely in the z-direction, with magnitude $|\mathbf{K}| = 2\pi/\lambda$. The wavevector of the diffracted light will have a magnitude

$|\mathbf{k}_m| = 2\pi n_{t,b}/\lambda$, depending on whether the diffraction orders are in the top or bottom cladding respectively. Therefore, the diffraction orders will only occur when Equation 2.31 can be satisfied, i.e. for $\beta - m\mathbf{K} < \mathbf{k}_{m,z}$ the orders are not permitted.

Assuming that $\mathbf{k_{m,z}}$ is non-zero and that β is only in the z-direction, then Equation 2.31 can be re-written for the first diffraction order directed upwards into the upper cladding (with index n_t):

$$\sin\theta = \frac{\beta - K}{k_1} \Rightarrow \frac{\lambda}{\Lambda} = n_{\text{eff}} - n_t \sin\theta \tag{2.32}$$

where θ is the angle of $\mathbf{k_1}$ from the y-axis. Consequently, by placing an optical fibre at the same angle θ, light can be coupled out of the waveguide. By reciprocity, the reverse will also hold true, meaning that the same grating design and fibre arrangement can be used to couple light into the chip as well.

There are a few important things to note by approaching input and output coupling in this way. Firstly, although Equation 2.32 gives the design condition for the upwards diffracted order, there will also be light in the higher diffraction orders and in the diffraction orders that propagate into the bottom cladding, leading to inherent optical loss at each grating. For material platforms that have an underlying substrate (for example, SOI) there will be some reflection at the interface of the bottom cladding and substrate (shown in Figure 2.18), which can couple back into the waveguide, which can be useful for increasing the coupling efficiency of the grating coupler. In fact, the inclusion of metal layers beneath the grating coupler to increase this reflection has been suggested [136], but this approach is not CMOS-compatible. Secondly, while it is possible to have vertical propagation ($\theta = 0$ and $\mathbf{k}_{m=1,z} = 0$), this will also give a second diffraction order that is reflected back into the waveguide, which is undesirable and so this specific design is generally avoided. Thirdly, the dependence of Λ on n_{eff} means that the grating couplers have a strong polarisation-dependence; in silicon photonics, waveguides have TE and TM modes with significantly different values of n_{eff} and will therefore require different optimal grating periods for coupling each polarisation. Finally, although the Bragg condition (Equation 2.31) implies that the grating coupler will only work for a specific λ, this is for an infinite grating. For a finite grating in a real device, there will be a range of k-vectors that are allowed by the Bragg condition, meaning that the grating coupler has some, albeit narrow, spectral bandwidth around the specific design value of λ.

Practically, the grating coupler has advantages over in-plane coupling. Alignment effectively only takes place in two directions, as the fibre is lowered very close to the surface of the chip and then its position is only varied over the area of the grating (the optimised one-dimensional grating design will be applied over $\sim 10\,\mu\text{m}$-wide waveguide). This makes the alignment simpler, faster and more reproducible, as the optimal coupling position only has to be found in two axes rather than three. Furthermore, as the grating couplers are made using nanoscale fabrication

methods, they typically have more consistent insertion losses as any defects tend to be smaller than macroscale processes (like polishing or cleaving). Gratings are also not limited to be placed at the edge of a chip, meaning that they can be used for testing at multiple points on a device or across a wafer, particularly as they can be made as erasable [137].

In the MIR range, the best grating coupler designs have bandwidths of $\sim 200\,\text{nm}$ and coupling efficiencies of less than $-3\,\text{dB}$ [35, 126, 138]. Material platforms that have a low refractive index contrast, such as GOS [139], tend to have grating couplers with low coupling efficiencies without more sophisticated design strategies [140].

2.2   Materials for the mid-infrared wavelength range

When selecting appropriate materials to use for group IV photonics in the MIR wavelength range, there are two important questions to answer:

1. can the materials enable compact devices with low-loss propagation?

2. can the materials be integrated easily together and the devices be realised using CMOS-compatible fabrication techniques?

This section will discuss the material properties and integration techniques to aid the answering of these questions and to explain the choices of material platforms in this thesis.

2.2.1   Transparency range

Before discussing individual materials, it is important to highlight the absorption mechanisms that dictate the transparency range of semiconductors. The electronic bandgap of the semiconductor determines the shortest wavelength that can be transmitted with low loss; at shorter wavelengths, photons have energies larger than the bandgap and will therefore be absorbed, exciting electrons from the valence to conduction band. Typically, the bandgap-related absorptions are not relevant in the mid-infrared range as the photon energy is too small; for example, silicon has a bandgap of $1.1\,\text{eV}$ which is equivalent to a photon wavelength $\lambda \sim 1.1\,\mu\text{m}$ [141].

At the other edge of the transparency range, the transmittance is determined by multiphonon absorptions [71, 72]. Phonons (the quantisation of crystal lattice vibrations) are termed either acoustic or optical, where adjacent atoms in the lattice move in or out of phase respectively. Individually, optical phonons will generally have lower frequencies and lower energies than are required for the absorption of photons (acoustic phonons have even lower frequencies than optical). However, multiple
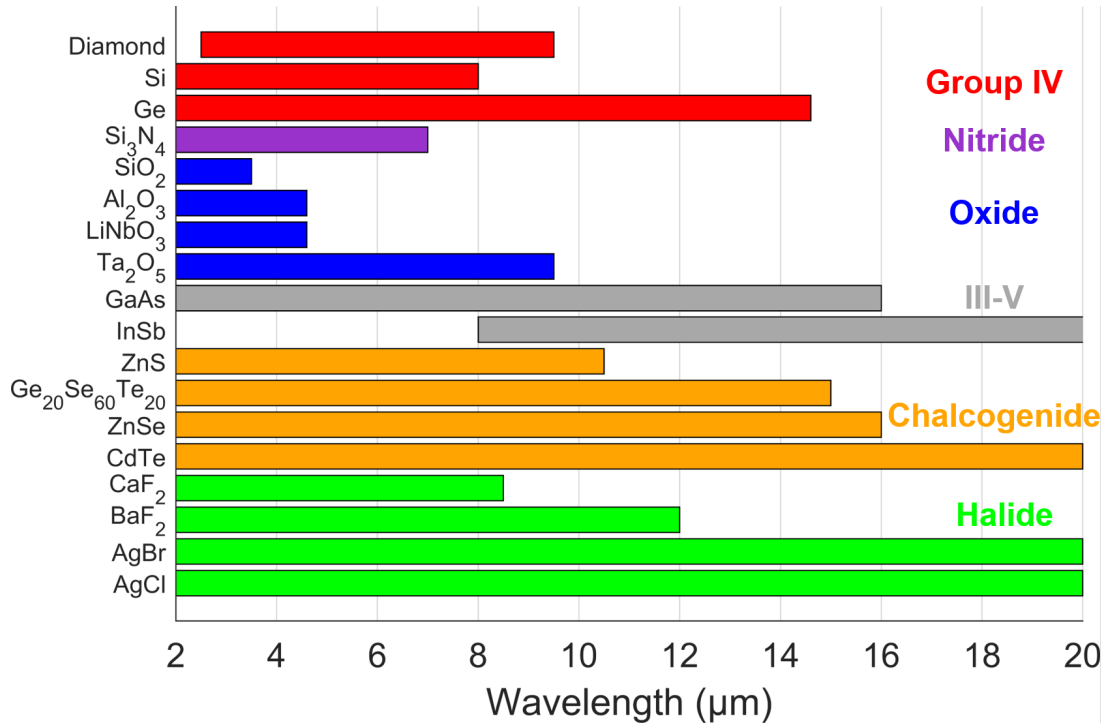
FIGURE 2.19: Transparency ranges of various crystalline materials commonly used at MIR wavelengths, approximated from the extinction coefficients. Data from [142–158].

optical phonons can have a total energy equivalent to a single photon, allowing it to be absorbed; the multiphonon absorption occurs when the photon energy becomes low enough (or photon wavelength becomes long enough) for this to happen. Crystal lattices with heavier elements will have the multiphonon absorption edge at a longer wavelength, as the phonon energies will be lower. This can be seen by comparing the upper edge of transparency ranges of crystalline materials with their counterparts below them in the periodic table: for example, contrasting silicon with germanium or oxides with chalcogenides (see Figure 2.19).

Figure 2.19 shows the transparency ranges (with minimal material absorption) of materials commonly used at MIR wavelengths. The choice of appropriate waveguide materials is determined by the part of the MIR wavelength range for which a device would be used. As discussed in Section 2.1.2 silicon waveguides are typically used, while at longer wavelengths germanium waveguides have been considered. The limitations of using silicon dioxide as a cladding longer wavelengths is also evident, so different cladding materials must be used for $\lambda > 3.5\,\mu m$.

Looking beyond group IV materials however, there are a range of materials that at first glance seem more suitable for MIR waveguides. The transparency of silver halides (AgBr and AgCl) cover the entire MIR range so could in principle allow low-loss propagation at any wavelength. Such waveguides have been demonstrated at $\lambda = 10\,\mu m$, albeit with high loss attributed to the large sidewall roughness [159]. However, it has long been known that silver halides will decompose when exposed to visible light [160], a fact that made them very useful for early photographic plates

[161]. This would likely make them unsuitable for devices that need to be placed *in situ*, greatly limiting the number of applications that they could be used for.

Chalcogenide glasses (containing sulphur, selenium and tellurium atoms instead of oxygen) have also been explored for low-loss waveguides in the MIR [162–165], and are particularly suitable for applications at longer wavelengths. However, many chalcogenides also contain toxic elements such as arsenic and antimony, leading to their suitability for biological sensing being widely discussed within the fibre community [166, 167]. Furthermore, chalcogenide glasses generally lack mechanical robustness when contrasted with comparable materials. For example, characteristic measures such as the Knoop hardness, fracture energy and Young's modulus all have values for chalcogenides that are small fractions of the values for oxide glasses [168, 169].

GaAs waveguides have also been demonstrated at longer wavelengths [170]. However, III-Vs have a shared issue with chalcogenides and silver halides: by conventional wisdom, they are not compatible with CMOS foundries. As discussed in Chapter 1, this is critical to have devices that can be produced at high-volume and low cost and maximise the potential of MIR integrated photonics. Therefore, the reliance continues to be on the group IV materials (particularly silicon, germanium and the nitride and oxide of silicon), unless non-CMOS materials can be integrated into the material platforms.

2.2.2   Other material considerations

Beyond the transparency range, other material properties must be considered to assess suitability. Table 2.2 shows the crystalline lattice constant, the coefficient of thermal expansion (CTE) and refractive index of a selection of the crystals from the previous section.

Crystal	n @ $\lambda = 6\,\mu m$	Lattice constant @ 300 K (Å)	CTE @ 300 K $(10^{-6}\,K^{-1})$
Si	3.4	5.43108	2.6
Ge	4.0	5.6576	5.7
GaAs	3.3	5.65325	5.7
InSb	4.0	6.4794	5.19
ZnS	2.2	5.4105	7.1
ZnSe	2.4	5.6687	7.1
CdTe	2.7	6.481	5.0

TABLE 2.2: Crystal properties of select cubic lattice semiconductors in Figure 2.19. Data from [143, 144, 150–152, 154, 155, 171].

As discussed in Section 2.1.1, there must be a contrast between the refractive index of the waveguide core and the cladding. For example, for the SOI platform typically used in the NIR and short wavelength MIR, Si has a refractive index of $n \sim 3.4$ while the SiO_2 cladding has a refractive index of $n \sim 1.4$. While the refractive

index contrast does not necessarily have to be this marked, a greater contrast enables strong confinement of the light within the waveguide core and enable more compact PICs to be realised. For example, for the GOS platform the difference in refractive index between Ge ($n \sim 4.0$) and Si is only $\Delta n = 0.6$, meaning the GOS waveguides have larger waveguide cross-sections, bend radii, etc. in comparison to SOI. It is evident from Table 2.2 that the chalcogenides could offer a good contrast to group IV or III-V waveguides.

For certain heterogenous bonding techniques or epitaxial growth of dissimilar materials, the size of the crystal lattice is very significant. Mismatch in the lattice constants can cause threading dislocations at the interface of the two layers, causing optical scattering from these defects if the waveguide mode overlaps with them [59, 76]. Likewise, significantly differing thermal expansions of the lattice will contribute to defects at the interface of the layers, which is an issue as many fabrication processes take place at elevated temperatures [171]. To avoid inducing additional propagation losses, the lattice constants and CTEs must be matched closely.

If the need for CMOS-compatibility is temporarily ignored (this will be addressed in the following section), the choice of core and cladding materials can be decided based on the transparency range, refractive index constrast and crystal structure parameters. If group IV waveguides are assumed (as they have the highest refractive index, excluding InSb), then GaAs and ZnSe would allow for the entire transparency range of Ge and ZnS to cover the full transparency of Si. CdTe is transparent over the full MIR range but it has a large lattice mismatch with both Si and Ge (16.20% and 12.70% respectively), while InSb is not suitable for shorter wavelengths.

At this point, trade-offs between the materials have to be considered. ZnS offers a large index contrast with Si ($\Delta n = 1.2$) and has a small lattice mismatch (0.38%), but the difference in the CTE would likely cause signifcant strain at the Si-ZnS interface during fabrication. GaAs and ZnSe both advantages as cladding layers for Ge waveguides, but drawbacks as well. GaAs has a small lattice mismatch (0.08%) and a similar CTE, but a lower index contrast ($\Delta n = 0.7$). ZnSe, on the other hand, gives a large refractive index constrast ($\Delta n = 1.6$) and also has a small lattice mismatch with Ge (0.20%), but with a significant difference in CTE (although smaller than between Si and Ge, which is an existing platform).

2.2.3   Heterogeneous integration

A material platform with a non-group IV bottom cladding offers a challenge because wafers containing layers of both group IV materials and, for example, chalcogenides are not currently commercially available. This means that material layers must be added during fabrication.

A typical approach would be to deposit layers of each material onto a substrate, with the cladding first and then the waveguide layer. However, although the crystallinity of the films can be controlled by selecting the deposition conditions, deposition of group IV materials at CMOS-compatible temperatures (lower than $400°$ C) will lead to an amorphous structure. For Ge, this is an issue as an amorphous Ge waveguide will have higher losses than its single-crystal counterpart [172, 173]. The additional loss from the material absorption can also occur at specific bands due to different defects in the material, which could potentially mask analyte absorptions; amorphous Ge has absorption peaks at wavelengths of 5.4 µm and 13.8 µm [174], within the region of interest for chemical sensing applications with Ge waveguides. Further, as noted in the previous section, epitaxial growth of these thin films can create defects at the layer interfaces for materials with different CTEs and a lattice mismatch, particularly at the elevated temperatures required for deposition. Additionally, once a non-CMOS buried cladding layer is deposited, the resulting material stack would be incompatible with further CMOS processing.

Alternatively, group IV waveguides could be first made on a separate wafer and then later transplanted to a transparent substrate. Using an approach such as this could in principle have most of the fabrication workflow be CMOS-compatible, as the integration would be a back-end process. There are several ways to approach developing a waveguide platform in this way. Direct bonding is a process for adhering two semiconductor wafers or samples together solely with the adhesive forces between the two surfaces (van der Waals and hydrogen bonding). The advantages of this technique mean that it requires no intermediate adhesive layers like other methods [175], which could cause additional waveguide loss, and can be used to bond materials with dissimilar material properties without introducing strain into the material [176]. Further, the bonding will often spontaneously occur at low temperatures and the strength of the bond can be improved by plasma treatment of the surfaces before bonding [177, 178]. This latter property means that it could relax the criteria for the choice of materials in the previous section so all of GaAs, ZnS, ZnSe and CdTe could be used as waveguide cladding. Further, it is used on an industrial scale as part of the process to produce SOI wafers [179, 180], although this requires annealing for stronger covalent bonding between the surfaces and reintroduces issues with lattice mismatch and dissimilar CTE [181].

However, the surface conditions for direct bonding are very strict, requiring a surface roughness of less than 1 nm and the surface to be free of all potential contamination [177]. Further, there is also the question of defining the devices in the platform before integration with the substrate, as the donor substrate must be removed. The SOI wafer production flow follows a process similar to Figure 2.20(a); an implantation of H^+ ions is used to create a layer of microscopic defects beneath the wafer surface, which will result in a crack along this layer and parallel to the surface when the wafer is annealed [179, 180]. However, this is not suitable as the heating will

also lead to chemical bonding at the interfaces and therefore defects due to lattice mismatch. This approach also results in the presence of a non-CMOS compatible layer before fabrication.

Alternatively, Si substrates can be removed by first thinning the substrate by grinding and then a selective acid etch [182]. In this case, the process flow would appear similar to Figure 2.20(b), where the waveguides of the PIC have already been

(a)

Step 1: H⁺ ion implantation   **Step 2: Direct bonding**

Step 3: Anneal, split and separate   **Step 4: Polish**

(b)

Step 1: Device fabrication   **Step 2: Thin substrate**

Step 3: Direct bonding   **Step 4: Two wet etches**

FIGURE 2.20: Different approaches for substrate removal after direct bonding of different wafers: (a) substrate removal by H$^+$ ion implantation; (b) substrate removal by grinding and wet etching.

defined. The features do not necessarily mean that the wafer no longer meets the minimal roughness requirement for direct bonding, as direct bonding of crystals to pre-etched waveguides has been shown previously, for use as optical isolators [183] and for III-V layers to Si waveguides for light sources [184]. Note that a buffer layer is required between the device layer and substrate to prevent etching of the waveguides. The simplest case for this would be for the buffer layer to be SiO_2 (for example, if Wafer A was SOI) that could be removed by a second selective acid etch, to etch only the SiO_2 and not the underlying Si.

Although this approach would yield the full material platform required, questions remain over the quality of the bond between the two wafer surfaces when one already has a waveguide defined, as [183] only shows this for small samples and [184] required pre-etched outgassing channels to prevent void formation at the bonding interface [185]. Furthermore, this process would be costly in terms of resources, the entire donor substrate is effectively being destroyed each time.

A more resource-efficient approach might be then for the transfer of complete circuits to be integrated with a transparent substrate. In recent years, transfer printing has been used to achieve the integration of heterogeneous materials for multiple applications [186]: placing micrometre-scale light emitting diodes on flexible substrates with nanoscale-accuracy for displays [187], integrating III-V lasers into photonic circuits [188] and transferring membranes devices without damage [189]. Typical areas of the transferred devices are on the order of $100\,\mu m \times 100\,\mu m$, although membranes up to $2\,mm \times 2\,mm$ have been successfully printed [190].

Although the individual technique can differ, all transfer printing generally follows the same process. The devices to be transferred are fabricated on a donor wafer or substrate and "released" from it, such as by underetching and suspension. An elastomeric stamp is then used to take advantage of reversible adhesion, so that the released devices can both be picked up and deposited by manipulating the speed of the stamp [191]. The placed devices are held in place on the receiving substrate by adhesive forces (i.e. van der Waals) between the thin device and the surface of the substrate. Since the technique can be used for thin membranes, it could work a feasible technique for integrating group IV waveguides with dissimilar but transparent substrates, by fabricating an entire PIC on a membrane and then transferring it to a new substrate. Furthermore, since the membranes can be placed with high accuracy, it may be possible to place multiple membranes adjacent to one another to realise a larger device, with edge coupling between the membranes or with a technique such as photonic wire bonding [192] to facilitate inter-membrane coupling.

2.3 Summary of literature review

In this chapter, the existing relevant literature on mid-infrared integrated photonic components, materials and platforms has been explored. In particular, the theoretical concepts have been introduced and working principle of each component has been discussed to, in most cases, illustrate the wavelength dependence that must be overcome for a fully wideband circuit. Other components that were not addressed directly in this work but are useful to later discussions (such as MZIs or directional couplers) were described. Furthermore, the opportunities and challenges of integrating materials outside of group IV into silicon photonics has been reviewed.

Chapter 3

Tools and techniques

3.1 Introduction

This chapter is concerned with the methods and techniques used to develop the devices and platforms detailed in this thesis. Like the previous chapter, the discussions here are important for providing context for the presentation and analysis of the results in later chapters.

3.2 Simulation

Various commercial software packages exist for the simulation of photonic devices, such as Photon Design, RSoft and the Lumerical DEVICE suite of programs. The choice of software is largely determined by the personal preference of the user because each package offers similar optical solvers: in this work, the Lumerical DEVICE suite was solely used. Details of the individual simulations and their results will be given in the relevant chapters.

In particular, the finite-difference time-domain (FDTD) and waveguide simulator programs were used. The FDTD program is a two- and three-dimensional electromagnetic simulator and is useful for comprehensive simulation of a structure or device. However, for larger structures the full simulation can often be costly in terms of computing resources. In this work, it was particularly used for designing grating couplers. The waveguide simulator, MODE, can be used for: performing modal analysis of the waveguides (such as for determining the single-moded geometry, or calculating bending losses); bidirectional eigenmode expansion (useful for designing MMIs or waveguide tapers); and variational FDTD (varFDTD). The varFDTD simulator is referred to as a "2.5D solver" as it collapses a three-dimensional geometry into a two-dimensional problem, making the simulation faster and using less

FIGURE 3.1: Exemplar simulation setup for modal analysis of a waveguide: (a) Cross-sectional structure of a waveguide, with simulation region highlighted by orange line and additional mesh over orange region; (b) The meshed structure as a map of the refractive index, with structure superimposed; (c) Simulated *E*-field intensity of the waveguide mode.

FIGURE 3.2: Exemplar simulation setup for determining coupling efficiency of a grating coupler, with a fibre above (in blue). The red and green arrows show the input and output ports of the simulation, with the yellow monitors used to track the propagation of the optical power.

computer memory. The varFDTD solver is only applicable for problems where light is confined to one plane of the device, i.e. where there is negligible vertical coupling between different slab modes in the material stack, or conversion from TE polarisation to TM. This makes it useful for designing devices like Y-junctions or directional couplers.

Examples of simulation setups for the modal analysis of a waveguide and the coupling efficiency of a grating coupler are shown in Figures 3.1 and 3.2 respectively.

Where possible, to check the simulation results, it is good practice to first use the faster solvers in MODE to sweep many design variations to the optimal geometry. The optimum can then be simulated again in FDTD to confirm the result.

3.3   Fabrication

The fabrication processes used in this work generally followed the same approach and it is therefore useful to detail them all in this section. Specific parameters (such as resist thicknesses and etch depths) and where any fabrication steps have deviated from those outlined below will be stated in the relevant following chapters.

3.3.1   Layout design

After the requisite components or the circuit as a whole have been designed and simulated, the layout of the device on the wafer or chip is drawn using a computer-aided design tool, which was the Tanner L-Edit software package in this work. In this software, the masks are prepared from a top view of the PIC, as is required by the lithographic and vertical etching fabrication processes to fabricate the device. By default, the layout can be drawn manually with a selection of geometry tools, but this software additionally has the advantage that external C++ code can be used as macros for drawing features or large designs. In particular this facilitates quick designing of the layouts, as a library of macros can be established for similar components that only need the design parameters (for example, waveguide width) to be updated for the new layout. All layout designs were output from L-Edit in the industry-standard GDSII file format and can then be used to create lithographic masks or transferred directly to the chip with maskless lithography.

The form of maskless lithography used in this project was electron beam lithography. In this case, the GDSII layouts designs were converted using the GenISys BEAMER software. Unlike optical lithography techniques, the pattern data from the GDSII must be first optimised for electron beam lithography, to include information about the beam (such as dose size) and to account for various process effects, such as electron scattering in different material stacks.

In either case (mask or maskless), some features are common for both layout designs. For designs that require multiple lithography and etch steps, alignment marks are essential. These are pre-defined marks that are placed at the edge of the design area on the wafer or chip, to provide a reference point for the lithography tool to ensure that the features defined in later steps have the correct placement respective to the earlier ones. Furthermore, all dry etch steps need to have metrology boxes. These are squares or rectangles, often $100\,\mu\text{m} \times 100\,\mu\text{m}$ in size, that are etched at the same time as smaller features. Unlike the smaller features, their large size means that a light beam in an ellipsometry tool can be positioned to reflect off them and the remaining material thickness therefore be determined. Consequently, they are crucial for determining accurate etch depths.

3.3.2   Lithography

Lithography typically follows the process as in Figure 3.3, which shows an SOI wafer that will be processed. First, a layer of a photosensitive chemical, known as a photoresist, is applied evenly to the surface of the wafer. This is achieved using a technique known as spin coating, where the wafer or sample is placed on a turntable and the resist is dispensed onto the surface. The turntable is then rapidly rotated, causing the resist to spread across the surface; the thickness of the resist (with a given

FIGURE 3.3: Process flow for lithography, before etching.

viscosity) can then be controlled by the rotation speed of the turntable. The coated wafer or sample is then baked to remove the solvent from the resist.

In the exposure step, a mask is used to selectively expose the photoresist to ultraviolet (UV) light. Features on the mask are transferred to the resist by blocking areas from the light, while the exposed areas will chemically react to the light. With a positive resist, the exposed areas will be softened and for a negative resist, they will be hardened (see Figure 3.3). The softer areas of the resist can then be removed by developing the resist (by cleaning it in a solvent), leaving areas of the Si surface exposed to be etched.

Note that there are different approaches for the exposure step. For projection lithography as shown in Figure 3.4(a), the mask has macroscopic features and the light transmitted through the mask is focused onto the surface using a lens, resulting in the micrometre- or nanometre-scale features in the resist. As a rule of thumb, the critical dimension (i.e. the size of the smallest feature in the transferred pattern) is around 200 nm for this style of lithography, although this depends on the wavelength of the

FIGURE 3.4: Diagram of lithography tools: (a) projection lithography, where light is passed through a lens and projected onto the wafer; (b) *e*-beam lithography, where the layout is written directly onto the wafer using a beam of electrons.

source. In contrast, in contact lithography the mask is placed in contact with resist, similar to the setup as in Figure 3.4(a) except without the projection optics. Consequently, the features in the mask need to be the same size as the transferred pattern. Mask resolutions below 1 μm are achievable, but this process results in etched features with comparatively rough edges, which can lead to scattering in waveguides. As a consequence, contact lithography is generally not used for fabrication of PICs, to avoid large waveguide propagation losses.

Alternatively, the designed pattern can be transferred without the use of a mask, in processes like electron beam (*e*-beam) lithography that is shown in Figure 3.4(b). In this case, under vacuum electrons are accelerated by an anode and focused into a fine beam, before being used to write the pattern directly into the resist (controllable by the deflector). The resists used for *e*-beam lithography behave in a similar way to the photoresists for UV lithography, with the chemical or physical structure changing on exposure.

Maskless lithography techniques take much longer than light-based lithography, as it is much quicker to expose a wafer to light than it is to write individual features and, since the same mask can be used repeatedly for many exposures, the use of masks facilitates scalable, high-volume fabrication. On the other hand, masks also require a significant initial cost and will have a non-negligible lead time when ordered from an external supplier, making *e*-beam far more practical for rapid prototyping. The resolution of *e*-beam lithography is also better than can be achieved by UV lithography as the wavelength of an acclerated electron is smaller than the wavelength of a UV photon. A resolution of $50 - 100$ nm can be routinely achieved, with better resolution if the acceleration voltage is increased. For these reasons, *e*-beam lithography was almost exclusively used in this work, unless otherwise noted.

A positive resist ZEP520A was used for *e*-beam lithography with the corresponding developer ZED-N50. Before spinning, the wafer surface is cleaned with oxygen-based plasma for 2 minutes to remove any organic contaminants and then placed in an oven at $200\,^\circ\mathrm{C}$ to dehydrate for $1 - 2$ hours. As with all resists, the thickness of the ZEP520A can be controlled by adjusting the rotational speed of the spinning turntable; the spin-speed curve is given in the manufacturer documentation. The wafer or sample is then baked on a hotplate to harden the resist by removing some of the solvent, for $2 - 5$ minutes at $180\,^\circ\mathrm{C}$ (depending on the applied thickness). To avoid charge build-up while patterning in the *e*-beam, and the resulting beam deflection and pattern distortion, the conductive polymer ESpacer 300Z was applied on top of the resist. The ESpacer can be rinsed off in water post-exposure and before development. After etching, the ZEP520A can be removed using an oxygen-based plasma asher or by dissolving in anisole.

3.3.3   Etching

With the designed pattern now transferred to the wafer, the exposed surface layer can be removed or etched to fabricate waveguides and other devices. Two types of etching process were used in this work, simply referred to as "wet" or "dry" etching to denote the presence or lack of liquid chemicals during the process.

Note that the etching processes will also etch the resists, although if appropriately selected it will be at a slower rate than for the wafer. Consequently, when applying the resist for lithography, a sufficiently thick layer must be used so that the resist is not etched completely through. However, for deeper etches this can be an issue when attempting to define smaller features in lithography.

3.3.3.1   "Dry" etching

"Dry" etching is a general term but, in this work, refers to the bombardment of the wafer or sample surface with ions to remove material through transfer of kinetic energy and chemical reaction. Specifically, this process is called reactive-ion etching (RIE).

An example of a RIE tool is shown in Figure 3.5. An electrode is situated inside and is electrically isolated from a vacuum chamber, on which the wafer or sample is placed. A radio frequency (RF) voltage is applied to the electrode, while the chamber is grounded; the electric field removes the electrons from the gas (input from the top of the chamber) and creates a plasma consisting of electrons and positive molecular

FIGURE 3.5: Diagram of a RIE tool. The electrode has a RF voltage applied to it, while the chamber is grounded.

ions. Low pressure is maintained in the chamber using vacuum pumps at the bottom of the chamber.

As the oscillating electric field cycles, the mobile electrons are accelerated up or down within the chamber, while the positive ions (which are much more massive) move comparatively little. Electrons that strike the walls of the chamber are removed to the ground, while electrons that enter the electrode will cause it to have an electric charge, as it is isolated in terms of direct current (DC). Likewise, electrons being removed from the plasma will cause the plasma to be positive overall, leading to a DC bias between the plasma and the electrode.

The resulting DC electric field means that the positive ions are accelerated towards the sample and collide with it. If only inert chemicals (such as argon) are present in the gas inserted to the chamber, bonds in the sample will only be broken through the transfer of kinetic energy. However, the etch rate can be increased by using oxidising chemicals such fluorine, to chemically react with the sample. Since the ions are accelerated in one direction, the subsequent removal of the wafer material is anisotropic (i.e. the wafer is only etched in one direction), as shown in Figure 3.6. Furthermore, since the etch rate will be more dependent on the chemical reaction than the etching from kinetic bombardment, the etch is selective (i.e. for an SOI wafer, the etch rate for Si is much higher than for SiO_2 with the appropriate chemistry).

FIGURE 3.6: Illustrative diagram of the anisotropic and selective etch of an SOI wafer.

Increasing the RF power applied to the electrode will increase the DC bias and in turn the energy of the ions, therefore increasing the etch rate. However, some ions will then have too much energy, causing additional etching in undesired directions or to other materials, thereby decreasing the selectivity and anisotropy of the etch. Consequently, an induction coil is added to the upper part of the chamber with a different RF power applied to it. This coil can increase the plasma density, known as an inductively coupled plasma (ICP), meaning the number of positive ions increases and therefore so too does the etch rate without increasing the DC bias.

This type of ICP-RIE system was used in this work: an Oxford Instruments ICP 380 plasma system. As all dry etching processes in this work were on SOI wafers or samples of varying thicknesses, the same recipe was used but with different lengths of time exposed to the plasma, corresponding to the desired etch depth and material thickness. In all cases, a C_4F_8/SF_6 chemistry was used, at a table temperature of 15°C and at a pressure of 15 mTorr. The applied DC and RF voltages were set to 800 V and

50 V respectively. This chemistry is used as it is highly selective for Si over SiO_2 and will therefore allow overetching of Si, which is important to ensure complete definition of etched features with large depth-width aspect ratios.

Generally, the areas to be etched on a design have dimensions that are on the order of micrometres or larger and the etching process will behave with a reasonably consistent etch rate across the wafer. In contrast, etching features that have small dimensions in-plane with the material layers, such as small holes, can lead to a loss of the etch anisotropy. The small features limits the access of the plasma as the feature is etched and, as illustrated in Figure 3.7, this can lead to a rounding at the bottom of the hole or for a thin layer of material to remain at the bottom of the hole (or both); it is a more notable problem as the etch depth increases (for the same size hole). The geometry of the feature will then deviate from the simulation or design, which can lead to further issues in both the device performance and later fabrication steps, as will be discussed later. One way to mitigate this is to "overetch" the features, or running the etch process beyond its endpoint, to ensure that the features are completely defined. When etching features completely through the Si layer, this relies on the selectivity of the plasma chemistry as described above for there to be no significant etching of the underlying SiO_2.

FIGURE 3.7: Potential incomplete etching in small etched features; the black dotted line shows desired cross-section of the etched feature.

3.3.3.2   "Wet" etching

"Wet" etching consists of any techniques that use liquid chemicals to remove material through chemical reaction. In this work, the sole chemical used was hydrofluoric (HF) acid at a ratio 7:1, for H_2O:HF. HF etches are highly selective, in that it reacts with Si at approximately $1/100$ of the rate at which it reacts with SiO_2. For SiO_2, the reaction that occurs is

$$SiO_2 + 6HF \rightarrow H_2SiF_6 + 2H_2O \tag{3.1}$$

with an etch rate of $\sim 100\,\text{nm/min}$.

HF etches are isotropic so the etch rate is the same in every direction, in contrast to dry etching. In particular, this is useful for underetching layers and can be used for

suspending structures. This is shown in Figure 3.8, where a hole through the top Si layer in an SOI wafer has been made using a dry etch to expose the SiO_2 beneath. The wafer is then immersed in aqueous HF which isotropically etches the SiO_2 layer, appearing as a rounded etched region. If left for a sufficiently long time, the lateral etching will mean that the SiO_2 layer will be removed under the Si layer, meaning that areas of the Si layer can be suspended.

FIGURE 3.8: Isotropic HF etch of underlying SiO_2 in an SOI platform to suspend the Si layer.

3.3.4 Waveguide facet preparation

Typically, fabricated samples are designed so that the waveguides are not close to the edge of a chip or wafer, to both allow for easier handling and to prevent causing lithography errors. For use in a butt-coupled setups (see Section 3.4.2.4), the end facets of waveguides have to be exposed in order to couple light in and out of a sample. To do this, the excess material of the sample is removed using one of three methods: cleaving, dicing or grinding. Each of these have their own advantages and disadvantages. If appropriate, the sample edges are then polished to improve the facet quality.

3.3.4.1 Facet exposure

Cleaving, the simplest method of the three, is the technique of deliberately breaking the sample in one direction to separate different parts of the sample (in this case, separating the wanted and unwanted parts of the chip). Making a small scratch at the edge of the sample and applying physical stress is all that is required to cleave a sample with a crystalline substrate. As well as being simple and quick to employ, this technique has the advantage that the end facets should not require any further processing (i.e. polishing), as proper usage will mean that the break is along a crystal plane. However, the efficacy can heavily depend on the skill of the user; improper implementation can lead to the break passing through features in the worst case, or an uneven and irregular break in the best. In the latter case, this would likely require additional polishing, as otherwise it becomes difficult to ensure consistent fibre-facet alignment and would introduce an unsystematic insertion loss, thereby losing one of the inherent advantages for cleaving.

When dicing a sample, a circular saw makes cuts in a sample at programmed locations and orientations. As this is often used to divide a larger wafer into smaller chips, it is most useful when the design has been fabricated on a wafer (so multiple chips can have the waveguide facets exposed at the same time), but it is can also be used for individual chips. Due to the nature of the tool, the placement of the cuts (and therefore end facets) is very reproducible for this technique but the dicing saw leaves the facets with a comparatively rough surface, meaning fine polishing is needed for any usable coupling into the chip. Dicing also requires the resist S1813 to be spun onto the surface to act as a protective layer during the process.

However, it should be noted that there is a form of dicing (not used in this work) that can provide optical quality facets. Ductile dicing is an optimised form of dicing that relies on machining material via plastic deformation, rather than the brittle machining that results in the rough edges [193]. This technique can yield nanometre-scale roughness on the cuts and has been used to make optical quality end facets in GOS [61] and chalcogenide waveguides [194], and even to produce waveguides in lithium niobate [195].

Grinding is in principle an almost identical process to polishing, that will be outlined below in Section 3.3.4.2. The only differences are that the grit size of the silicon carbide (SiC) abrasive discs are larger (15 µm diameter first, then 8.4 µm) to remove more material with each disc revolution. An advantage of grinding is that, like dicing, the placement of the chip edge (and waveguide facet) can be precisely controlled and is reproducible. Further, as grinding uses the same tool as for polishing, it is straightforward to carry out both processing steps subsequently. Conversely, grinding is much slower than both cleaving and dicing, as the length of the processing depends on the amount of material to be removed. As a consequence, grinding can be fairly laborious, taking at least one hour or possibly multiple (when re-mounting to do the other side), while both cleaving and dicing will take less than 30 minutes (cleaving is effectively instantaneous, while dicing also requires the wafer or chip to be mounted into a tool).

In this work, where consistent and reproducible fibre-facet alignment was required, dicing or grinding were used to create a flat facet at the correct location within a waveguide, before further polishing with a finer grit size to ensure a high-quality facet. However, in some cases (such as the imaging of the waveguide facets for the endlessly single-mode waveguides in Chapter 4), cleaving was sufficient for the facet exposure as no alignment with the fibres was required.

3.3.4.2   Polishing

Ahead of polishing (or grinding), a protective layer needs to be present on top of the chip to protect the surface and any fabricated devices. For the Y-splitters (Chapter 6), a thick layer of resist AZ2070 was spun onto the chip, achieved by spinning

comparatively slowly at only 1000 rpm. The resist was hardened by baking it for longer than typical (for 7 minutes at 100°C) to remove as much solvent as possible from the resist. For the wideband MMIs (Chapter 5), hard-baking the AZ2070 caused small features to be damaged and they would be removed with the resist. In this case, the wax that is used for mounting the samples (see below) was applied to the surface; it can act as a protective layer as it is solid at room temperature. As it is applied by hand and not spun, it is generally much thicker and less uniform compared to hard-baked AZ2070. As the wax layer will also melt upon mounting the chip, additional care had to be taken to ensure sufficient surface coverage.

The polishing apparatus can be seen in Figure 3.9(a). Before placing it in the the polishing tool (Metaserv 2000 Grinder/Polisher), the sample was fixed to the mount to be safely held in the tool. A small amount of molten wax was applied to the backside of the chip using a toothpick, before placing on the mount. Likewise, a "dummy chip" (a spare or scrap sample of similar size) was attached to the other side of mount, to balance it while in the polishing tool. A sample ready to be polished after being attached to the mount is shown in Figure 3.9(b).
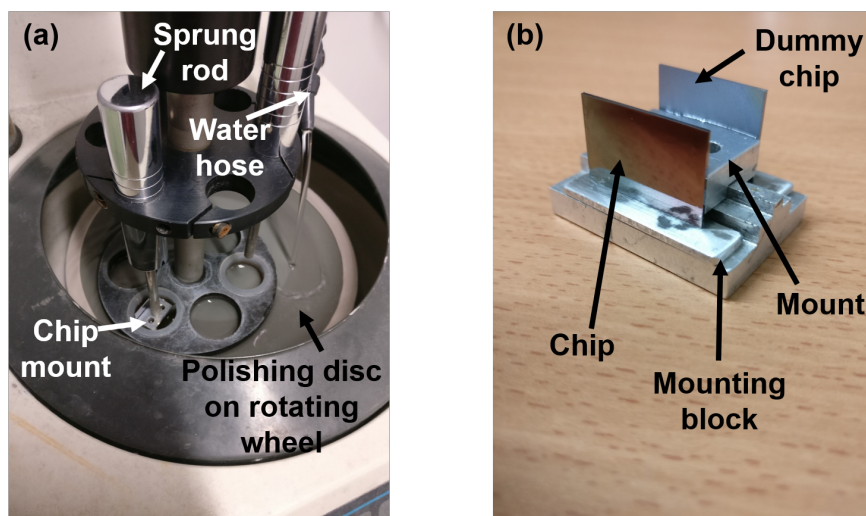
FIGURE 3.9: Photographs of the apparatus used for the polishing of waveguide end facets: (a) sample in the polishing tool; (b) sample to be polished on the chip mount.

An aluminium oxide (Al_2O_3) abrasive disc is attached to the surface of the rotating wheel and then the sample holder is positioned such that the edge of the sample is in the contact with the disc. The sample is orientated so that its edge is parallel to the radius of the polishing wheel; this way, the abrasive action is normal to the substrate and the end facet polishing can be in a consistent direction. The sprung rod is then lowered to apply pressure to the chip mount and ensure the chip is firmly in contact with the polishing surface. The water directed onto the disc acts as both a lubricant and to remove excess debris.

The sample is polished for a predetermined interval and rotational speed (dependent on the sample), before the end facets are inspected under optical

microscope. If the facets were sufficiently visible for that polishing step, the abrasive disc was swapped for one with a smaller grit size, starting at a 3 μm diameter, then 1 μm and finally 0.3 μm. When imaged under an optical microscope, if the waveguide facets were sufficiently contrasted with the dark adjacent etched areas and appeared bright (indicative of reflectivity and therefore a smooth surface) then the polishing was determined to be completed.

After polishing all sides of the samples on which there are waveguide facets, the chip is then detached from the mount by melting the wax on the backside of the chip. Both the protective layers (either wax or AZ2070) can be removed by submerging in a solvent, either acetone or N-methyl-2-pyrrolidone (NMP), before rinsing in isopropyl alcohol (IPA) and water separately.

3.3.5   Post-fabrication inspection

During and following fabrication steps, various microscopy techniques can be used to inspect chips and ensure the critical dimensions agree with the designs. Optical microscopy, which uses visible light, use objective lenses to magnify a sample and produce a real image for the user's eye. The resolution of optical microscopes is restricted by the diffraction limit (or Abbe limit), which for visible light corresponds to several hundred micrometres. However, they are still useful for device or chip-level inspection and in particular for ensuring features have developed properly and are correctly aligned during lithography.

Scanning electron microscopes (SEMs) uses an electron beam rather than light to image a sample. Conceptually very similar to the *e*-beam lithography tool (Section 3.3.2), in SEMs the electron beam is scanned across the surface of the sample to yield its topography and spectroscopic analysis of its composition. In particular relevance to this work (although not the sole imaging technique used), electrons that escape from the surface of the sample as a result of inelastic collision between the beam and the surface (termed secondary electrons) are used to build the image. The low energy of the secondary electrons means that they are localised to a small area on the surface and can therefore give nanometre-scale resolutions.

A focussed ion beam (FIB) tool uses, instead of a beam of electrons, a beam of ions that can be used for both imaging and milling of the sample. The imaging principle is similar to that of the SEM, where the collisions of the ions with the surface generates secondary electrons, but much more massive ions will cause damage to the surface even at low currents and, for example, can result in roughening of waveguide surfaces. As a consequence, it is preferable to only use FIB imaging for samples in which destructive testing is acceptable (normally after optical characterisation has finished). However, the destructive nature of the FIB beam is useful to mill samples, in which higher currents are used so that the FIB beam cuts away material. This is

commonly used to provide a cross-section of materials that lie beneath the surface of the sample.

3.4 Optical characterisation

3.4.1 Normalising transmission

Normalisation waveguides are important to include in any chip layout for optical characterisation, as they can provide a reference for the losses in the system. For example, consider the simple implementation of the cut-back method for determining propagation loss in a waveguide, as shown in Figure 3.10.

FIGURE 3.10: Schematic diagram of the grating-coupled implementation of the cut-back method.

In this layout, there is an array of waveguides where each successive waveguide increases by a length ΔL, so the length of the m^{th} waveguide is

$$L_m = L + m\Delta L \tag{3.2}$$

and, from Equation 2.13 in Chapter 2, it will have a loss of αL_m, assuming that the waveguide loss is constant across the array.

However, this is not the only loss present across the system: the waveguide tapering and the coupling efficiency of the grating couplers will also contribute losses. By dividing the power of each of the longer waveguides by the shortest waveguide (or normalising the transmission to the shortest waveguide), the losses that arise from elements common to all devices (in this case, the tapers, gratings and propagation along the initial length L) can be factored out and so the only loss being measured is due to the differences between the devices, which is $\alpha\Delta L$ in this case.

Importantly, this works across more complicated measurement layouts and is used to factor out losses caused by aspects such as waveguide bending, attenuation or scattering for example.

3.4.2   Measurement setups

Multiple experimental setups were used for the optical characterisation of the devices presented in this thesis. These were predominantly arranged so that the coupling of light in and out of the chips would be achieved using gratings (at $\lambda = 1.95\,\mu$m, $\lambda = 3.80\,\mu$m and $\lambda = 7.67\,\mu$m) but, for broadband characterisation over $\lambda = 2.5 - 3.7\,\mu$m, light was butt-coupled into the chips.

3.4.2.1   Characterisation at $\lambda = 1.95\,\mu$m
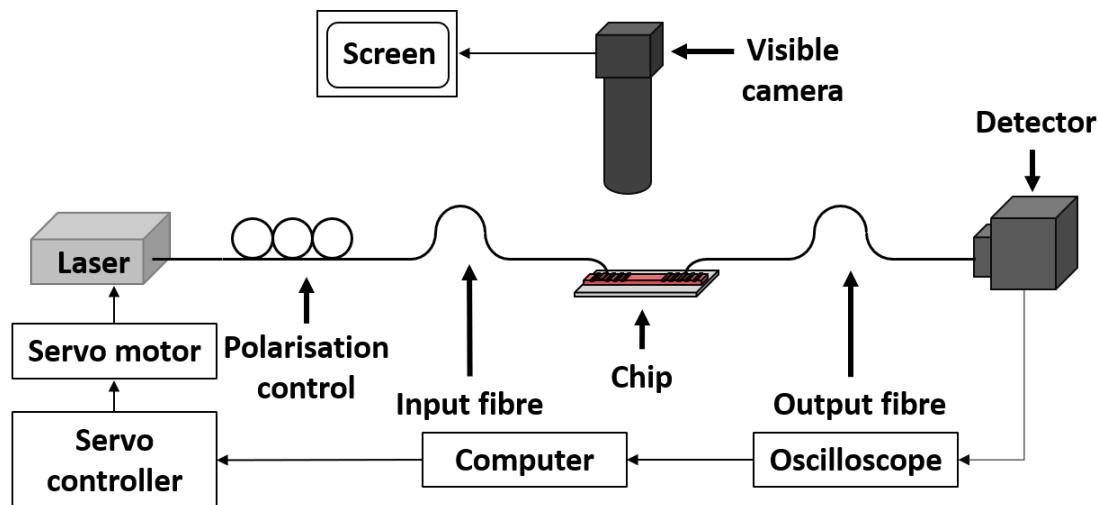
FIGURE 3.11: Experimental setup used for characterising devices at 1.95 µm.

A schematic for the experimental setup used for characterisation at $\lambda = 1.95$ µm is shown in Figure 3.11. A Thorlabs TLK-L1950R tunable laser (centred at $\lambda = 1.95$ µm, tunable over $\lambda = 1.89$ - 2.02 µm) was used as a source and coupled into a silica optical fibre (Thorlabs SMF-28e). This laser has an external cavity and tuning is provided by driving a servo motor to adjust the position of a grating that forms part of the cavity. The fibre is wound through a 3-paddle polarisation controller, that uses stress-induced birefringence within the fibre for control of the polarisation. The fibre is attached to an angled mount, that allows for control of the fibre relative to the sample stage. The fibre mounts are supported on 3-axis stages for position control during alignment; the output fibre (also Thorlabs SMF-28e) is mounted in a similar way. The chip under test is placed on the sample stage, and fibre-to-chip coupling is achieved using gratings patterned onto the chip during fabrication. The output fibre is coupled onto a Thorlabs PDA10DT-EC InGaAs photodetector (sensitive over $\lambda = 0.9$ - 2.57 µm), with a PicoScope 2000 oscilloscope used to recover the detector signal.

Both the oscilloscope and the controller for the servo motor were connected to a computer and run by a Python program. This provided automation of the sweep of

the laser wavelength and synchronisation between tuned wavelength and the recorded detector signal.

The end of the fibre was initially aligned with the grating coupler by observing both the end of the fibre and the sample with a visible camera. After initial alignment, the laser was tuned to the central wavelength (λ = 1.95 µm) and the detector signal monitored. The positions of both the input and output fibres were then adjusted systematically to optimise the magnitude of the detector signal.
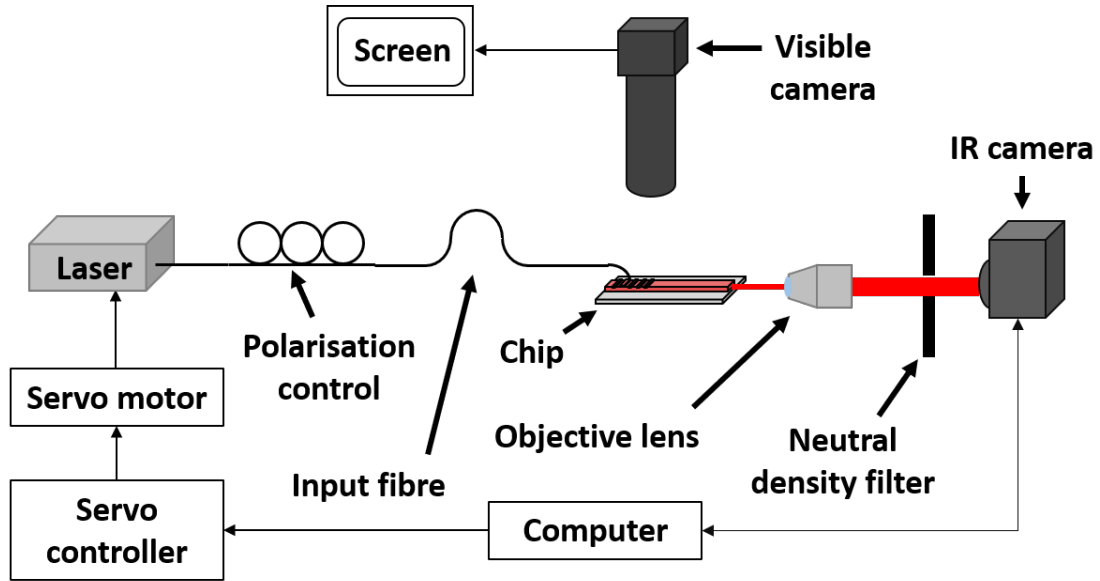
FIGURE 3.12: Experimental setup used for imaging waveguide outputs at 1.95 µm.

The setup of this apparatus was adapted for the imaging of the waveguide output as in Figure 3.12, for the confirmation of single-mode waveguide propagation (see Chapter 4 for further discussion). In this case, the chip was cleaved to expose the waveguide facet and remove the output grating coupler (the input was left intact). The output fibre and photodetector were then replaced with a $63\times$ objective lens, to expand the output beam, and a Xenics Xeva 1.7 320 infrared camera, to image the waveguide output. A neutral density filter was placed in the beam path between the lens and the camera to prevent damage to the infrared camera from high optical power. As with the standard arrangement for this setup, the camera and the tuned laser wavelength were synchronised by connecting them to a computer and running a Python script.

3.4.2.2   Characterisation at $\lambda = 3.8\,\mu\text{m}$

The experimental setup used for optical characterisation around $\lambda = 3.80\,\mu\text{m}$ is shown in Figure 3.13. In this setup, a Daylight Solutions quantum cascade laser (QCL), tunable over $3.72 - 3.90\,\mu\text{m}$, was used as a source and coupled into the input fibre (Thorlabs InF$_3$ single-mode fibres) using a ZnSe lens. The QCL was orientated such

FIGURE 3.13: Experimental setup used for characterising devices at 3.8 μm.

that its output would be TE-polarised with respect to the waveguides, but the propagation through the fibre would cause some non-negligible polarisation rotation. Consequently, an adjustable half-wave plate was placed in the beam path prior to coupling into the fibre, to compensate for this polarisation rotation and to ensure that the output of the fibre was TE-polarised. Alignment of the input and output fibres with the on-chip grating couplers was facilitated placing the fibre holders in three-axis stages and imaging overhead using a visible camera. The output light, after couplig into a fibre using the grating, was collected using a Vigo Systems PVI-4TE fan-cooled HgCdTe detector (sensitive over $3.0 - 6.5\,\mu m$ or, for measurements requiring greater sensitivity, an InfraRed Associates Inc. IS-1.0 liquid nitrogen-cooled InSb detector (sensitive over $1.0 - 5.5\,\mu m$).

To recover the comparatively low-amplitude voltage signal from the detector, a lock-in amplifier was used. The input beam was modulated using a chopper, driven by the chopper controller with the same frequency that was passed to the lock-in as a reference.

3.4.2.3   Characterisation at $\lambda = 7.67\,\mu m$

For longer wavelength measurements, a similar experimental setup to that shown in Figure 3.13 was employed, but instead using a single-mode continuous-wave distributed-feedback QCL (Thorlabs QD7500CM1) as the source. Unlike the other sources in this work, this laser was not tunable and had a 106 mW maximum output power at a wavelength of 7.67 μm. A black diamond-2 lens with a focal length of 1.9mm was used to collimate the output laser beam before modulation using a chopper wheel, as in the setup for $\lambda = 3.8\,\mu m$. A second black diamond-2 lens was used to focus the beam and enable coupling into a single-mode As_2Se_3 fibre (Coractive IRT-SE-28/170). Fibre-to-chip coupling was achieved using on-chip grating couplers and, after coupling into another single-mode fibre, the output light was collected with

a liquid nitrogen-cooled HgCdTe detector (Infrared Associates Inc. MCT-13-1.00). As above, the signal-to-noise ratio was improved using a lock-in amplifier.

3.4.2.4 Broadband characterisation across $\lambda = 2.5 - 3.7\,\mu$m
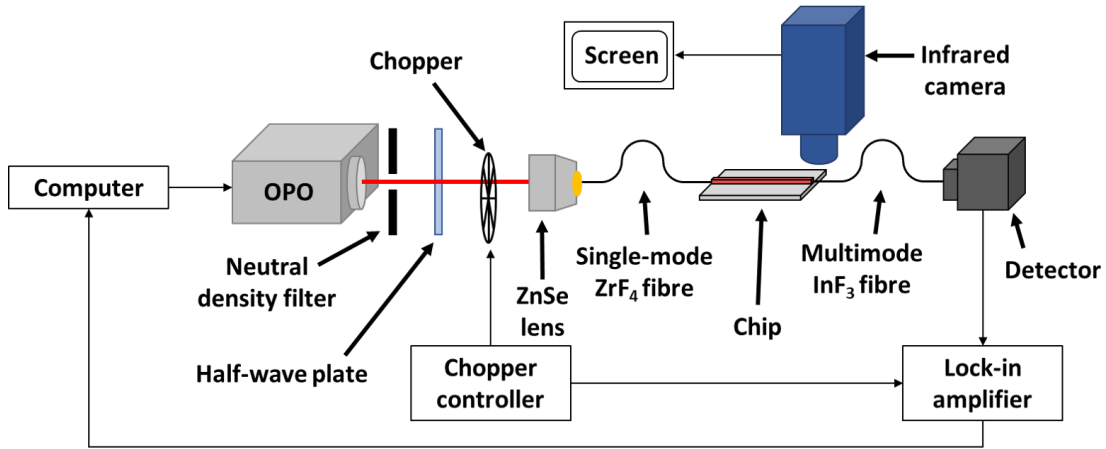
FIGURE 3.14: Experimental setup used for broadband characterisation over $2.5 - 3.7\,\mu$m.

For large optical bandwidth measurements, a butt-coupled setup was required to avoid the wavelength-dependence of grating coupling. In this setup, an M Squared Firefly-IR optical parametric oscillator (OPO) was used as a source (with an manufacturer-specified upper bound on the linewidth as $< 10\,\text{cm}^{-1}$), as it is tunable over $\lambda = 2.5 - 3.7\,\mu$m. Similar to the $3.8\,\mu$m grating, the OPO output beam was coupled into a single-mode fibre using a ZnSe lens and an adjustable half-wave plate was used to compensate for polarisation rotation within the fibre. The same detector used in the $3.8\,\mu$m setup was used to measure the optical signal, and it was recovered using a similar chopper and lock-in amplifier arrangement.

Unlike the $3.8\,\mu$m setup, a large-core multimode fibre was used to collect the waveguide output power, to facilitate easy chip-to-fibre coupling. Further, the alignment of both the input and output fibres with the waveguide facets was achieved using an overhead FLIR infrared camera.

Observing the waveguide facet with the OPO emitting, a bright spot would appear on the chip edge if the fibre had an approximately correct height alignment with the chip. Then by observing the output facet, the effective coupling into the waveguide can be seen as a small bright spot will appear at the output if the input is aligned correctly. The output fibre can then be positioned so it is in focus with the chip (indicating the height alignment is approximately correct), before the measured signal on the detector is used to optimise the alignment.

To prevent damage to the camera when observing the intense spots at the input and output waveguide facets, one or more neutral density filters were placed in the

OPO beam path before coupling into the fibre. These were then removed once observation with the camera has been completed.

3.5   Summary

This chapter detailed the various techniques used throughout the development cycle of the devices in this work. In particular, the techniques shown here are common throughout the design, fabrication and characterisation procedures for the majority of the devices, where deviations from these procedures will be detailed in the relevant chapters.

The simulation tools are introduced and described, with preference in this work being for the Ansys Lumerical DEVICE suite of modelling tools. The procedures for translating the simulated designs into fully fabricated devices are presented, with discussion around areas where there is risk the fabricated devices differing from the intended design to provide context for later chapters. The experimental setups used for characterising the devices are also detailed, as they must be considered in device design and in the discussion of the experimental results.

Chapter 4

Endlessly single mode waveguides

4.1   Introduction

For most photonic integrated circuits (PICs), waveguides that support multiple modes are undesirable because, if light is coupled into more than one of these supported modes, they will interfere and the output will depend on the phase relationship between the modes. This phase relationship will vary along the length of the waveguide as each mode has a different effective index (the phase constant $\beta = 2\pi n_{\text{eff}}/\lambda$). Consequently, the optical behaviour of any PIC made of such waveguides becomes complicated to predict and, once other circuit components are considered for a PIC of realistic complexity, the behaviour is near impossible to predict and reproduce consistently. Therefore, waveguides that only support the fundamental mode, termed "single-moded", are essential for any real PIC.

For PICs that only require a small working optical bandwidth, single-mode propagation at or around a specific wavelength is easily achievable with an analysis of the waveguide modes, as discussed in Chapter 2. However, conventional rib or strip waveguides have a limited wavelength range over which they support only the fundamental waveguide mode: towards shorter wavelengths it will become multimoded, while at longer wavelengths the waveguide will cease to be guiding. For a device that needs to cover a much larger wavelength range (for example, a MIR spectroscopy device), an alternative waveguide design needs to replace the conventional rib or strip.

Over the last few decades, photonic crystal fibres (PCFs) have been used for single-mode propagation seemingly independent of wavelength, otherwise known as endlessly single-mode (ESM) guidance [196, 197]. In the PCF, a two-dimensional photonic crystal runs along the length of the fibre to act as a cladding, as shown in Figure 4.1, while light is guided along the defect (the missing hole) in the centre of the fibre. This structure causes, in theory, the higher-order modes to be cut-off at all wavelengths, while the fundamental mode can still expand in size with wavelength

FIGURE 4.1: Illustration of the cross-section of a PCF. The white circles indicate air holes that run the length of the fibre, while the blue area is the core material of the fibre.

and is therefore supported at all wavelengths (this mechanism is explained in more detail in Section 4.2.1).

The application of the ESM guidance concept to integrated photonics has been previously proposed for SOI and for a SOI/III-V hybrid platform [198]. For this design, the cladding on either side of the waveguide is a one-dimensional grating in a direction transverse to the propagating wave. This is comparable to taking only the middle row of holes, including the defect, in the PCF; the PCF defect is equivalent to the waveguide core. The etched features of the grating (i.e. the "grooves") are filled with an effective material (instead of air for the PCF). This material has a refractive index that can be controlled by nanopatterning of the silicon layer in the longitudinal direction (i.e. parallel to the direction of propagation).

Although these waveguides were predicted to be single-moded across 1.1 - 5 μm in simulations [198] (i.e. the whole transparency range of SOI), they have not been experimentally demonstrated in silicon photonics. This chapter will detail the successful design, fabrication and characterisation of these waveguides to support single-mode waveguiding over an octave of frequency.

4.2   Modelling and design

4.2.1   Geometry and working principle

The geometry of the ESM waveguide is shown in Figure 4.2, for an SOI platform. A two-dimensional grating is patterned into the silicon layer of thickness h, with a strip of width w left unpatterned to act as the waveguide core. In each of the transverse and

FIGURE 4.2: Three-dimensional illustration of a fabricated ESM waveguide, with the labelled design parameters. [top inset] Top-view of a section of the sub-wavelength grating cladding. [bottom inset] The sub-wavelength regions of the wavelength cladding, approximated as a homogenous material with index n_{SWG} (in yellow).

longitudinal directions (the x- and z-directions in Figure 4.2), the grating had a period Λ_{T} and Λ_{L} respectively. Consequently, the grating is formed by a patterned array of holes with size $\Lambda_{\text{T}} - a_{\text{T}}$ and $\Lambda_{\text{L}} - a_{\text{L}}$ in the transverse and longitudinal directions respectively. The width of the waveguide cladding w_{clad} formed by the grating with m transverse periods is then

$$w_{\text{clad}} = \left(m - \frac{a_{\text{T}}}{\Lambda_{\text{T}}} \right) \Lambda_{\text{T}} \tag{4.1}$$

as the external part of the last grating period (with width a_{T}) is part of the slab.

The waveguide cladding can be approximated as a one-dimensional cladding filled with an effective index material if Λ_{L} is small enough so that the grating is subwavelength in the z-direction (i.e. the period is smaller than the Bragg period [114]). As discussed in Chapter 2, the subwavelength regions of the grating can be modelled as an homogeneous anisotropic material (see bottom insert of Figure 4.2) with refractive index tensor n_{SWG}:

$$n_{\text{SWG}}^2 = \begin{bmatrix} n_{xx}^2 & 0 & 0 \\ 0 & n_{yy}^2 & 0 \\ 0 & 0 & n_{zz}^2 \end{bmatrix} \tag{4.2}$$

where

$$n_{xx}^2 \sim \frac{a_{\text{L}}}{\Lambda_{\text{L}}} \cdot n_{\text{Si}}^2 + \left(1 - \frac{a_{\text{L}}}{\Lambda_{\text{L}}} \right) \cdot n_{\text{air}}^2; \tag{4.3a}$$

$$n_{yy}^2 = n_{xx}^2; \tag{4.3b}$$

$$\frac{1}{n_{zz}^2} \sim \frac{a_{\text{L}}}{\Lambda_{\text{L}}} \cdot \frac{1}{n_{\text{Si}}^2} + \left(1 - \frac{a_{\text{L}}}{\Lambda_{\text{L}}} \right) \cdot \frac{1}{n_{\text{air}}^2}. \tag{4.3c}$$

Note that as well as the refractive index of silicon, n_{Si}, the refractive index of air, n_{air}, is also used because there is no deposited upper cladding and the grating holes are therefore filled with air.

As discussed in Chapter 2, modes within a waveguide can be considered guided if their effective index is greater than the refractive index of the cladding. Below this point, the modes can couple into the cladding modes and will be radiative. Therefore, for single-moded guidance, the higher order modes must have an effective mode below the cladding index, with only the fundamental mode having an effective index above this cut-off value. In material platforms with high-index contrast, this can be more difficult to design as the cladding index (for example, the refractive index of silicon dioxide $n_{SiO_2} \sim 1.4$) is fairly low while any higher order modes concentrated in the high-index waveguide (such as silicon, $n_{Si} \sim 3.4$) will have a higher index as well.

The ESM waveguide overcomes this limitation by using cladding with a generally higher index, by a similar mechanism as the PCF [197]. The propagation constant, and consequently effective index, of the cladding mode for this waveguide is comparatively large over the entire wavelength range, preventing higher order modes from being supported. In particular, at shorter wavelengths, the cladding mode becomes more concentrated in the non-etched features and its effective index increases as a consequence. This avoids the emergence of the higher order modes at short wavelengths, where they have a higher effective index than typical waveguide cladding (but a smaller effective index than the fundamental mode). Also notably, this design leads to the field of the fundamental mode in the waveguide being more highly concentrated within the waveguide core across the wavelength range in comparison to the conventional waveguide. This is discussed further in the following section.

4.2.2   Waveguide modes

The waveguide parameters in [198] were used as a starting point, with $n_{zz} = 1.5$ (by Equation 4.3c, achievable with a longitudinal duty cycle $a_L / \Lambda_L = 0.6$ for $n_{Si} \sim 3.4$) and a transverse duty cycle $a_T / \Lambda_T = 0.5$. The waveguides were designed for an SOI platform with a 0.5 µm-thick silicon layer, with the grating holes etched halfway through this layer (etch depth $d = 250$ nm). The core was fixed with $w = 1.2$ µm; from Figure 2.4 in Chapter 2, a 1.2 µm-wide, 0.5 µm-thick SOI conventional strip waveguide will support the TE_0 mode over $1.5 - 5$ µm, so it was expected it would also sufficiently support the fundamental mode of the ESM waveguide across the desired range as well.

The modes in the waveguide were simulated using Lumerical MODE, as shown in Figure 4.3. To determine which modes were guided, perfectly matched layer (PML) boundary conditions (shown with thick orange lines in Figure 4.3) were applied to the simulation region. PML boundary conditions allow power to leave the simulation region, so leaky modes (such as those that radiate out of the waveguide) are not found

by the simulator. The critical parameters (the duty cycles, transverse period and the etch depth) were varied over different wavelengths to find the optimal solution with the largest effective index of the fundamental mode (indicating confinement within the waveguide) and with no other modes supported by the waveguide. The optimal solution found had a similar design as in [198], with equal transverse and longitudinal duty cycles, $a_\text{T}/\Lambda_\text{T} = a_\text{L}/\Lambda_\text{L} = 0.6$, a transverse period $\Lambda_\text{T} = 600$nm and an etch depth $d = 250$ nm. Only the fundamental TE mode, TE_0, was found over the wavelength range $\lambda = 1.5\,\mu\text{m} - 5.5\,\mu\text{m}$ when this optimal design was simulated.

To confirm the single-moded guidance principle, the cladding mode was also simulated by removing the waveguide core from the structure with the optimised parameter values. As the cladding mode would be leaky, metal boundary conditions were used instead of PML to prevent power leaving the simulation region. The cladding mode was then tracked over the wavelength range $\lambda = 1.5\,\mu\text{m} - 5.5\,\mu\text{m}$; the

FIGURE 4.3: Screenshots of the simulation of the ESM waveguide mode in Lumerical MODE. (a) The three-dimensional view of the simulation. (b) The cross-sectional view, in which the PML boundary condition (the thick orange line) shows the edge of the eigenmode solver region. A mesh was applied across the entire simulation region to account for the large simulation region, with a finer mesh over the waveguide core (to enable modes to be found).

effective indices for the TE_0 mode and the cladding mode are shown in Figure 4.4. For comparison, the effective indices of the first two guided TE modes in a strip waveguide with the same cross-section ($1.2\,\mu m \times 0.5\,\mu m$) are also shown.

Evidently from Figure 4.4, the effective index of the cladding mode is sufficiently large across the simulated wavelength range to only allow the TE_0 mode to be guided (the TM_0 mode could not be found by the eigenmode solver with PML boundaries) and the TE_0 modes of the waveguides behave differently over the wavelength range. Despite both TE_0 modes being guided over approximately $\lambda = 1.5\,\mu m - 5.0\,\mu m$ (and for the ESM up to $5.5\,\mu m$), for the ESM waveguide the mode is concentrated within the waveguide as evidenced by the large effective index. On the other hand, for the strip waveguide the effective index gradually decreases, showing the mode is becoming less confined within the waveguide.

As there was no source available that covered the full wavelength range over which the ESM waveguide was simulated, specific wavelengths were chosen for characterisation: $1.55\,\mu m$, $1.95\,\mu m$ and $3.80\,\mu m$. For these wavelengths, the refractive index values of the homogenised subwavelength regions of the waveguide cladding are given in Table 4.1. The simulated mode profiles in the waveguide at each of these wavelengths are shown as the insets in Figure 4.4.

FIGURE 4.4: Modal analysis of an ESM waveguide, showing the modal effective indices of the ESM fundamental waveguide TE mode and the ESM cladding mode. The first two TE modes of a conventional SOI strip waveguide with the same waveguide core size are also shown for contrast. [inset] the TE_0 mode profiles (E-field) at $\lambda = 1.55\,\mu m$, $\lambda = 1.95\,\mu m$ and $\lambda = 3.80\,\mu m$.

Wavelength, λ (μm)	n_{Si}	n_{xx} (= n_{yy})	n_{zz}
1.55	3.4757	2.7656	1.4913
1.95	3.4510	2.7469	1.4901
3.80	3.4244	2.7269	1.4888

TABLE 4.1: Refractive index values for the homogenous anisotropic sub-wavelength regions of the cladding grating at the characterisation wavelengths. It was assumed that $n_{air} = 1$ at all wavelengths.

4.2.3   Bend radius

Exploiting a large cladding index for ESM guidance brings its disadvantages as well. Due to the small index contrast between the TE_0 and cladding modes, the lateral confinement of the fundamental mode in the ESM waveguide is minimal, leading to significant bending losses at even large bend radii.

The setup for the simulation of the bending loss is shown in Figure 4.5, carried out in the full 3D-FDTD. Although FDTD is not the fastest simulation tool, unlike for the modal analysis, MODE could not be used for various practial reasons. In MODE, the modes found by the solver (with PML boundaries) depends on the size of the simulation region, but at shorter radii (below 125 μm) where the mode is very weakly confined, the modal field would overlap with the boundaries (undesirable for PML boundaries) and give inaccurate results. Conversely, enlarging the simulation region meant that a finer mesh had to be used to avoid averaging of the cladding features, meaning that the computing resources required significantly increased and became prohibitive for shorter bend radii. Alternatively, not using the finer mesh meant that the solver found non-physical modes at all radii and, since the solver only finds an input number of modes, it would sometimes not find the actual waveguide mode. This lead to massive and unrealistic losses when considering modal mismatch in the bend and the solution was to increase the number of modes the solver looks for, which would also increase the required computing resources.

The simulation was setup with curved strips of the homogenised anisotropic material as the waveguide cladding. The software treats anisotropic refractive indices fixed to the dimensions of the simulation, whereas for the waveguide the anisotropy is relative to the direction of propagation. Therefore, a rotation attribute (the blue line in Figure 4.5 that extends rightward from the left side of the bend) was applied to the index of the homogenised strips such that the index rotated with direction of propagation. The rotation attribute is not available for varFDTD so, despite being a faster simulation tool, it could not be used. The fundamental mode was excited in the straight waveguide immediately ahead of the bend and the bending loss was calculated by considering the light passing through the input and output monitors (the yellow lines in Figure 4.5). The bending loss was evaluated at a wavelength of 3.8 μm, where the modal confinement was the weakest of the characterisation wavelengths (i.e. the effective index of the waveguide mode was closest to the

FIGURE 4.5: Top-view of the 3D-FDTD simulation of the waveguide bend radius (pictured for a 25 μm radius).

effective of the cladding mode, see Figure 4.4) and so the bending loss would be the limiting case. Note that the structure does not extend through the PML boundaries, which increased reflections at the simulation edges but avoided diverging simulations.

The results of the bending loss simulation are shown in Figure 4.6, in which the bending loss converge to approximately 0.6 dB/90° for radii above 300 μm. As the simulation appeared to converge, and because the FDTD computation time became prohibitive, larger bend radii were not investigated. It should be noted that the small lateral confinement does not solely determine the bending loss, as the simulation predicts larger propagation losses due to the overlap of the waveguide mode with the SiO_2 lower cladding at 3.8 μm because the absorption is considerably large at this wavelength and the overlap is larger as wavelength increases (see Figure 4.4). Nevertheless, the bending radius and loss is impractical for many PICs and is a known issue for ESM guidance in the photonic crystal fibre [196]. However, unlike the photonic crystal fibre, in integrated photonics there may be alternate design approaches to minimise the radii of waveguide bends. For example, one approach could be to use multimode strip or rib waveguides for the bends, thereby increasing the lateral mode confinement and decreasing the bending loss. To avoid exciting the higher order modes in the multimode waveguides, adiabatic transitions between the ESM and multimode sections can be used with combination with bends that have a gradually changing curvature, like Euler bends [199]. Overcoming this limitation is highlighted as an area of work for the future.

FIGURE 4.6: Simulated bending losses for the ESM waveguide at $\lambda = 3.8\,\mu m$, calculated in 3D-FDTD.

4.2.4 Subwavelength design

The period of the subwavelength features is determined by considering both the wavelengths at which the waveguide will be characterised and realistic fabrication requirements. Recalling from Chapter 2 that the condition for single-moded guidance is

$$n_{\text{eff,BF}} < \frac{\lambda}{2\Lambda_{\text{L}}} \tag{4.4}$$

and using a fixed duty cycle $a_{\text{L}}/\Lambda_{\text{L}} = 0.6$, a design criterion for the grating can be found.

The size of the etched holes in the longitudinal direction are $\Lambda_{\text{L}} - a_{\text{L}}$, or equivalently $0.4\Lambda_{\text{L}}$ for the given duty cycle. To realise these small features in a device, the resolution of the lithography techniques have to be considered: *e*-beam lithography allows reliable fabrication of features down to $\sim$50 nm. However, the design of the ESM waveguides assumes that all the features in the subwavelength regions are defined as a perfectly right-angled cross-sections with a uniform etch depth. From the dry etching discussion in Chapter 3, the small etched features will likely have some rounding at the bottom of the hole and this effect will be more significant for smaller ratios of $d : 0.4\Lambda_{\text{L}}$. For the subwavelength regions, this is a problem as the values of n_{SWG} will be sensitive to variations in the shape of the holes; incomplete etching will give a different cross-sectional profile of the refractive index. Consequently, for a given etch depth, the hole size $0.4\Lambda_{\text{L}}$ must be maximised to give the best etch profile, therefore meaning that the longitudinal period must also be maximised, while still fulfilling the subwavelength condition in Equation 4.4.

This design criteria can be observed in Figure 4.7. The effective index of the Bloch-Floquet mode $n_{\text{eff,BF}}$ will vary with grating design. However, if the subwavelength condition is met, the grating will change from behaving as a periodic

FIGURE 4.7: The subwavelength threshold, $\lambda/2\Lambda_L$, for different infrared wavelengths, for fixed duty cycle $a_L/\Lambda_L = 0.6$.

medium to an homogeneous medium, i.e. the Bloch-Floquet wave will become the usual guided wave. Consequently, the effective index of the Bloch-Floquet for the design criteria can be roughly approximated as the effective index of the fundamental mode in a waveguide with homogeneous index n_{yy} (as this will only depend on the fixed duty cycle of the grating, not its period directly), with the same upper and lower cladding layers as the subwavelength waveguide. From Equation 4.4, for a given period the wavelength of 1.55 μm will give the lowest threshold value so the waveguide mode was evaluated at this wavelength, yielding $n_{\text{eff,BF}} \sim 2.55$ (shown by the red dotted lie in Figure 4.7). Using this criteria, the largest possible grating period while still meeting the subwavelength criterion is $\Lambda_L = 300$ nm, for a corresponding hole size of 120 nm.

4.2.5   Grating coupler design

Lacking a continuous source available across $\lambda = 1.55 - 3.8$ μm, the waveguides were to be characterised separately at three different wavelengths in this range: 1.55 μm, 1.95 μm and 3.80 μm. Consequently, multiple sets of waveguides with the same design were required for each wavelength, with a corresponding different grating coupler design to optimise fibre-to-chip coupling for that wavelength.

The grating couplers were simulated using a 2D simulation within FDTD, with the fibre angle fixed at 10° from the normal to the waveguide (such as the simulation pictured in the simulation section of Chapter 3). A 2D simulation was used to

minimise the computing time required; an accurate simulation was achieved by reducing the geometry in the in-plane direction perpendicular to the waveguide to a constant small value ($\sim$ 1 µm). In each case, the etch depth was fixed equal to the etch depth of the waveguide cladding (250 nm) or to the etch depths of the reference waveguides (500 nm, see Section 4.3.2) for ease of fabrication, while the duty cycle and period of the grating were varied. The designs used for the fabricated gratings are shown in Table 4.2 and illustrated in Figure 4.8; note that the grating coupler for 3.8 µm has two periods and duty cycles because a subwavelength design was used (this last design was not new and had been reliably used in previous work).

   The grating couplers were implemented on-chip using a waveguide that had a core width of 10 µm for easy fibre alignment (see Figure 4.8), before tapering down to the designed core width of 1.2 µm over 1 mm. In the cladding of the taper, the holes were orientated to be parallel to the edges of the waveguide core. However, in the eigenmode expansion (EME) simulations in MODE, Figure 4.9(a), the holes were replaced with strips with anisotropic index n_{SWG} for faster simulations. The loss is simulated by calculating the scattering matrix for the taper input and output porrts, based on the fundamental TE modes in the straight waveguides at either end of the taper. The EME simulation predicted that the taper loss was negligible for tapers longer than 200 µm at $\lambda = 3.8$ µm, Figure 4.9(b), so the insertion loss due to the taper was minimal for this length. It is unclear what is causing the "rippling" in the loss for shorter taper lengths, but it is perhaps caused by interference between the fundamental mode and radiating modes, resulting in mode beating and fluctuations in intensity over the output port. This would explain why the effect is more

FIGURE 4.8: Implementation of the grating couplers over 10 µm-wide waveguides, for (a) $\lambda = 1.55$ µm, (b) $\lambda = 1.95$ µm and (c) $\lambda = 3.80$ µm. The red features show the ESM grating cladding, the purple features show the 250 nm etched gratings and the blue features show the 500 nm etched gratings.

Wavelength, λ (µm)	Period (µm)	Duty cycle (%)	Etch depth (nm)	Simulated coupling efficiency (%)
1.55	0.78	75	250	54
1.95	0.8	40	250	35
3.80	2.0 (1.0)	50 (75)	500	27

TABLE 4.2: Optimised grating coupler parameters and simulated efficiencies. Bracketed values indicate the duty cycle and period for the subwavelength region of the grating coupler design for $\lambda = 3.8$ µm.

pronounced at shorter taper lengths where the loss is higher and more power is
coupled into the radiating modes.

4.3   Results and discussion

The waveguides were fabricated using the techniques described in Chapter 3 on an
SOI sample with 500 nm Si and 2 µm SiO$_2$ layers. Two different lithography and etch
steps were used: a 250 nm etch to define the cladding holes and the grating couplers
for 1.55 µm and 1.95 µm, and a 500 nm etch to define the reference waveguides (see
Section 4.3.2) and the grating coupler at 3.80 µm. A captured image of the fabricated
ESM waveguide in an SEM is shown in Figure 4.10. The measured dimensions of the
cladding from SEM micrographs showed good agreement with the design, as shown
in Table 4.3, meaning that the waveguide was expected to behave as simulated.
    To confirm single-moded propagation in the waveguides across the wavelength
range, evidence was required to answer the following questions:

1. Do the waveguides guide at each measured wavelength?
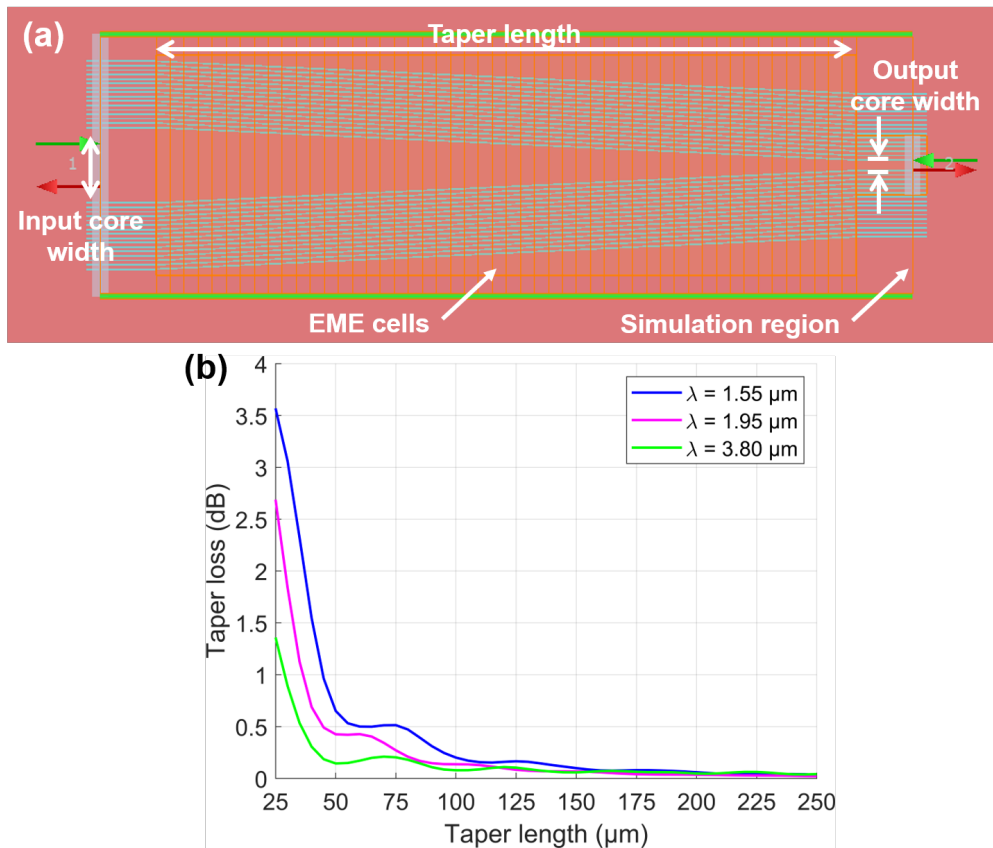
FIGURE 4.9: Simulation of the ESM taper in MODE: (a) top-view screenshot of the
EME simulation; (b) the taper loss with respect to taper length.

FIGURE 4.10: Top-view SEM micrograph of the fabricated ESM waveguide.

Dimension	Designed value (nm)	Measured value (nm)
Waveguide width, w	1200	1180
Transverse period, Λ_T	600	580
Transverse hole size, $\Lambda_T - a_T$	240	230
Longitudinal period, Λ_L	300	310
Longitudinal hole size, $\Lambda_L - a_L$	120	130

TABLE 4.3: Comparison of the designed and measured (from SEM micrographs) ESM cladding dimensions. Note that the values were measured to the nearest 10 nm.

2. At the shortest wavelength, are the waveguides single-moded or multimoded?

For the latter question, it was sufficient to show only single-moded propagation at the shortest wavelength because, for a waveguide with a given geometry, it expected that a greater number of higher order modes would be supported when a shorter wavelength is propagated than for a longer wavelength. Therefore, if only the fundamental mode is supported at the shortest wavelength, then only the fundamental mode will be supported at the other end of the measured range, provided that the longest is still guided. Consequently, by answering the two questions above, there is strong evidence that the waveguide behaves as expected from the ESM working principle.

4.3.1 Propagation loss

The propagation loss of the waveguides was determined using the cut-back method (see Chapter 3) with waveguide lengths spanning 0.15 - 1.5 cm in 0.15 cm intervals. Multiple sets of cut-back waveguides were used for measurements at each wavelength of 1.55 μm, 1.95 μm and 3.80 μm, with the corresponding grating coupler design. The waveguide transmissions in each set were normalised to the corresponding shortest waveguide in that set, as the different grating coupler designs would have different insertion losses.

The propagation loss can be a good proxy for demonstrating guidance in the waveguide, as high propagation loss would suggest that light is leaking from the waveguide. In this case, to investigate the effect of the width of the cladding on the leakage of the guided mode into the slab, two variations of the waveguide cladding thickness were fabricated. The first had the grating repeated for 10 periods (resulting in $w_{\text{clad}} = 5.64\,\mu m$, see Equation 4.1) and the second was repeated for 16 periods (as pictured in Figure 4.10, for $w_{\text{clad}} = 9.24\,\mu m$). For the 10-period cladding, the propagation loss at $3.80\,\mu m$ was found to be large and with significant variance between the waveguides, as shown in Figure 4.11. In addition, coupling was found between adjacent waveguides (i.e. power was measured at the output of one waveguide, with the light input into a different waveguide). Since the centre of the waveguides were separated by $50\,\mu m$, this coupling was likely due to intermediate coupling of the light into the slab and then into the next waveguide, rather than due to a directional coupler-like effect. Consequently, it was concluded that light was leaking out of the waveguide and that the 10-period cladding was determined to not be wide enough to ensure guidance.

From this point onwards, all references to the ESM waveguides are to those with width $w_{\text{clad}} = 9.24\,\mu m$. For this cladding thickness, four sets of waveguides across two different chips were fabricated for the determination of the propagation loss. However, for $\lambda = 1.95\,\mu m$, an error in the lithography step to define the waveguide cladding meant that two sets of the waveguides were unusable. Despite this, from the remaining waveguides the best propagation loss was determined to be $1.46 \pm 0.13\,\text{dB/cm}$, as showing in Figure 4.12 at $\lambda = 1.95\,\mu m$.

From the four sets of waveguides measured at $\lambda = 3.80\,\mu m$, the best propagation loss was $1.55 \pm 0.35\,\text{dB/cm}$ if both the fitted loss and corresponding uncertainty are

FIGURE 4.11: Waveguide cut-back propagation loss measurements for ESM waveguides with $w_{\text{clad}} = 6\,\mu m$ at $\lambda = 3.80\,\mu m$. The transmission measurements are normalised to the transmission through the shortest waveguide (0.15 cm). The quoted uncertainty is the 95% confidence interval of the linear fit.

considered. In both cases at $\lambda = 1.95\,\mu\text{m}$ and $\lambda = 3.80\,\mu\text{m}$, not only does this demonstrate that the waveguides support propagation at both wavelengths, but also that the propagation loss is comparable to that of other SOI waveguides at these wavelengths (see Table 2.1 in Chapter 2), indicating that the cladding design does not introduce significant additional loss.

At $\lambda = 1.55\,\mu\text{m}$, the propagation loss could not be determined. Beyond the shortest waveguides, measurable transmission through the waveguides was not observed and as a consequence there were insufficient data points to calculate the propagation loss. Coupling loss into the waveguides was not a limiting factor, as the coupling efficiency for the grating couplers at 1.55 µm was simulated as larger than for the grating couplers used at the longer wavelengths and the taper loss was simulated as minimal, as discussed in Section 4.2.5. When investigated by using both optical microscopy and an SEM, no fabrication defects were found that would affect transmission. However, similar to waveguides with the thinner cladding, coupling between adjacent waveguides was observed, suggesting that the light was leaking into the slab. Given that the confinement of the waveguide mode was simulated as being the greatest at this wavelength (see Figure 4.4) and that the longer wavelengths were guided with low propagation loss, the cladding width was not likely to be the cause of the leaking waveguide.

The cause becomes more evident when the subwavelength threshold is reconsidered. From Section 4.2.4, the longitudinal period chosen for the cladding design lies close to the limit of the subwavelength threshold. However, the Rytov equations that are used to calculate the refractive index of the subwavelength regions as an homogeneous material are only applicable for the deep subwavelength regime,

FIGURE 4.12: Waveguide cut-back propagation loss measurements for ESM waveguides with $w_{\text{clad}} = 9.6\,\mu\text{m}$ at $\lambda = 1.95\,\mu\text{m}$. The transmission measurements are normalised to the transmission through the shortest waveguide (0.15 cm) in each set. The quoted uncertainties in the propagation losses are the 95% confidence interval of the linear fits.
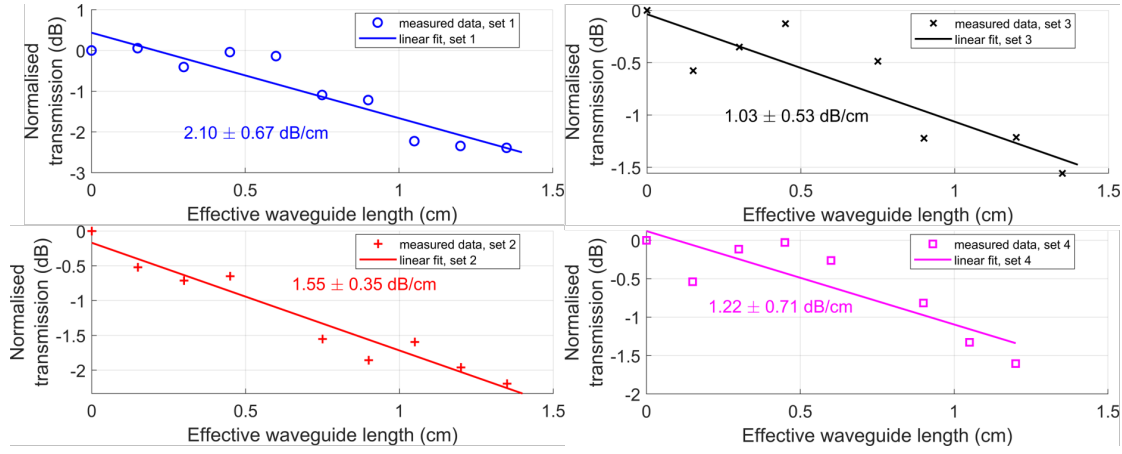
FIGURE 4.13: Waveguide cut-back propagation loss measurements for ESM waveguides with $w_{\text{clad}} = 9.6\,\mu\text{m}$ at $\lambda = 3.80\,\mu\text{m}$. The transmission measurements are normalised to the transmission through the shortest waveguide (0.15 cm) in each set. The quoted uncertainties in the propagation losses are the 95% confidence interval of the linear fits.

i.e. for $\lambda \gg \Lambda$. Consequently, when the value $n_{\text{eff,BF}}$ is close to the subwavelength threshold of $\lambda/2\Lambda_{\text{L}}$, the Rytov equations are no longer good approximations. Therefore, the mode simulations that use the approximations of the Rytov equations are likely less accurate at 1.55 µm and, while it was predicted as being guided at this wavelength, the TE_0 mode may not be guided in reality. To confirm this, future iterations of the design will have to be adapted to balance the hole size (and duty cycle) for fabrication purposes and the longitudinal period for the approximation of the Rytov equations to be applicable.

For example, if it is assumed that the duty cycle is kept the same and that the *e*-beam lithography limit is 50 nm, then $\Lambda_{\text{L}} = 125\,\text{nm}$ and $\lambda/2\Lambda_{\text{L}} = 6.2$, which is well above the effective index of the grating Bloch-Floquet mode ($n_{\text{eff,BF}} \sim 2.55$) and would therefore be in the deep subwavelength regime. However, in this case, then aspect ratio of the holes (for the 250 nm etch depth) will likely mean that the holes are ill-defined after etching and the fabricated device will not match the simulation as a consequence. To adjust for this, the duty cycle can be altered to increase the hole size without changing the period (effectively lowering $n_{\text{eff,BF}}$) but this will also decrease n_{zz} and will require new simulations of the waveguide to ensure broadband single-moded guidance. This will effectively impose another limit on the hole size. Alternatively, for original duty cycle, $\Lambda_{\text{L}} = 200\,\text{nm}$ is also above the subwavelength limit ($\lambda/2\Lambda_{\text{L}} \sim 3.9$) and gives a 80 nm hole size, for a hole aspect ratio of $\sim 3:1$ (depth-width) that is more realistic fabricate.

4.3.2 Test for single-moded behaviour

In a waveguide in which multiple modes have been excited, the individual modes a propagate with a different propagation constant β. Consequently, the relative phase between the waveguide modes will be wavelength-dependent and, as the modes within the waveguide will interfere with each other, then output of the waveguide will form an interference pattern that varies with the wavelength. In contrast, single-mode waveguides will not form an interference pattern, as there will be no higher order modes supported to interfere with the fundamental mode.

In order to obtain the interference pattern (or lack of), the ESM waveguides were cleaved to remove the output grating coupler and the associated taper to expose the waveguide facet; the output facets were then imaged in the setup shown in Figure 3.12 of Chapter 3. As a reference for outputs of single-mode and multimode waveguides, strip waveguides (with widths 0.45 μm and 1.2 μm for the single-mode and multimode waveguides respectively) were also designed and fabricated in the same material platform. The guided TE modes predicted for these waveguides at $\lambda = 1.95$ μm are shown in Figure 4.14. It should be noted that the multimode waveguide has the same core width as the ESM waveguide and is the same waveguide for which the modal analysis is shown in Figure 2.4 in Chapter 2 and the strip waveguide in Figure 4.4. The same grating coupler design was used for the input and the waveguides were also cleaved to expose their output facets. The reference waveguides were designed for 1.95 μm only because, as discussed previously, it is sufficient to show only single-mode propagation at the lower end of a wavelength range and the waveguides (under the current design) were shown not to work at 1.55 μm. Furthermore, the outputs could not be imaged at 3.80 μm because this was outside of the sensitivity range of the camera that was used.
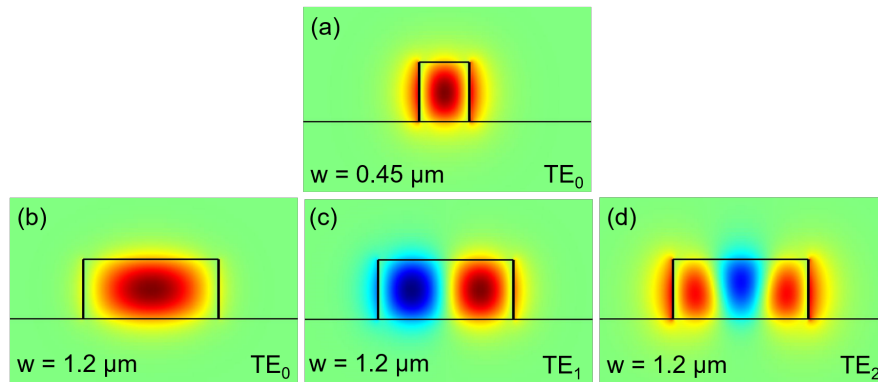
FIGURE 4.14: The electric field (red and blue areas indicate positive and negative amplitudes respectively) of (a) the guided fundamental TE mode of the single-mode strip waveguide (0.45 μm width) and (b)-(d) the first three guided modes for the multimode strip waveguide (1.2 μm width), at $\lambda = 1.95$ μm. Both waveguides had a height of 500 nm.

To observe the evolution of the waveguide output with wavelength, the wavelength was swept in 0.5 nm steps over $\lambda = 1.95 - 2.01\,\mu$m and an image of the output was recorded at each step. Examples of these images are shown in Figures 4.15(a-c), all of which were taken at $\lambda = 1.95\,\mu$m. The output intensity was measured in this way in an attempt to perform spatial and spectral (S^2) imaging as used in optical fibres [200], which has previously been achieved using a similar IR camera for real time measurements [201]. Recording images at each step of the wavelength sweep and plotting the subsequently results an intensity map of the waveguide output, with pixel intensities $I(x, y, \lambda)$. Equivalently, each pixel at position (x, y) can be considered as having its own intensity spectrum $I(\lambda)$. In the S^2 imaging technique, the Fourier transform of the intensity spectrum $\mathcal{F}[I(x, y, \nu)]$ is used to yield the relative intensities of each guided mode against their differential group delay, with respect to the fundamental mode. Note that $I(x, y, \nu)$ is equivalent to $I(x, y, \lambda)$, but defined in terms of frequency ν. In this way, each mode in the waveguide (or fibre) appears as a strong peak, providing clear evidence for either single-moded or multimoded guidance.

In fibres, S^2 imaging requires fibres that are metres in length, due to comparatively small differences in the modal effective indices and consequently small group indices and group delays. In waveguides, the refractive index of the waveguiding materials are typically much larger than in fibres so the group index of each waveguide mode will also be larger. Consequently, it was expected that these waveguides would have greater differential group delays (per unit length) that should have been resolvable in short centimetre-long waveguides. However, in [201], the wavelength resolution was ~ 1 pm in a NIR system, roughly two orders of magnitude better than was achievable with the MIR laser used here. Since the wavelength resolution directly determines the resolution of the Fourier transform, the limited wavelength step combined with the comparatively small group delays (due to short waveguide lengths) meant that no mode peaks could be resolved even in the designed multimode waveguide.

However, while the S^2 analysis was unsuccessful, the recorded intensities for the multimode waveguide showed variance with wavelength, suggesting an interference pattern. To show this, the intensities were summed in the y-direction to yield, for a given wavelength step, the intensity profile in the x-direction only, $I(x, \lambda)$. The plot of $I(x, \lambda)$ then shows a "heat map" of the intensity profile evolving with wavelength. This analysis was performed on the recorded intensities of each of the ESM, single-mode and multimode waveguides, as shown in Figures 4.15(d-f).

For the ESM waveguide, Figure 4.15(d), the centre of the intensity profile (outlined by the white-dotted 3 dB contour) at the output remains centrally located throughout the wavelength sweep, with some intensity fluctuations at the output edges attributed to background noise. Similar behaviour of the output intensity with wavelength can be seen for the single-mode waveguide in Figure 4.15(e), where the output intensity also has no movement in the x-direction (the noise is less evident

FIGURE 4.15: Imaged outputs of the ESM, single-mode strip, and multimode strip waveguides. The output intensity (normalised to the largest pixel intensity in the image) at each camera pixel at position (x, y) captured at $\lambda = 1.95$ μm is shown for each of the (a) ESM, (b) single-mode and (c) multimode waveguides; the low intensity area in the southwest of (b) is attributed to the use of multiple filters for imaging this waveguide. The heat maps of the output intensity, as a function of pixel position and wavelength, are shown for the (d) ESM, (e) single-mode and (f) multimode waveguides. The intensity values are normalised to the largest pixel intensity of each heat map. The white dotted line shows the -3 dB contour for the heat map.

here, as the signal-to-noise ratio was larger for this waveguide). However, for the
output of the multimode waveguide in Figure 4.15(f), the intensity has a significant
lateral shift (in the *x*-direction), much larger than any noise fluctuations, indicative of
a multimode interference pattern having been formed. Conversely, the lack of
significant lateral shift of the intensity for the ESM and single-mode waveguides
indicates that no interference pattern has been formed and therefore that no higher
modes are supported in this wavelength range. Thus, it can be concluded that the
ESM waveguide does indeed only support the fundamental mode at $1.95\,\mu\text{m}$ and,
since the same waveguide has been shown as guiding at $3.80\,\mu\text{m}$ from the propagation
loss measurements, it therefore follows that the ESM waveguide will also only allow
single-mode propagation at $3.80\,\mu\text{m}$. Further, it is notable that since the ESM and
multimode waveguide cores have the same cross-sectional area, the lack of higher
order modes provides evidence that they are being suppressed by the cladding
design, rather than simply being avoided with a suitable cross-sectional geometry.

4.4   Summary

In this chapter, the well-known endlessly single-mode guidance principle of photonic
crystal fibres has been successfully experimentally demonstrated in integrated
photonics. Building on previous theoretical work, a simulated modal analysis of the
waveguide is presented to illustrate the working principle of the waveguide design,
predicting single-mode propagation over a wavelength range of $1.5 - 5.5\,\mu\text{m}$. The
results of the optical characterisation of the waveguides are shown, yielding
propagation losses of $\sim 1.5\,\text{dB/cm}$ and confirmation of single-mode propagation over
the wavelength range $1.95 - 3.80\,\mu\text{m}$, equivalent to an octave of frequency. For further
development, the design of the cladding grating must be adjusted to match the
simulation and to allow for single-mode propagation at shorter wavelengths.
Furthermore, different approaches for designing the waveguide bends are required to
prevent the waveguide bending loss from being prohibitive.

Chapter 5

Wideband multimode interferometers

5.1 Introduction

As discussed in Chapter 2, power splitters and couplers are essential for making MIR photonic sensing devices, particularly in realising on-chip spectrometers or splitting light between sensing and reference arms. However, many existing power splitters have a significantly narrow bandwidth over which they can effectively operate or have other fabrication-related drawbacks (such as the requirement for an infinitely small tip on a Y-junction), meaning that this is an important area to address for the development of MIR PICs.

In the NIR wavelength range, MMIs have been demonstrated with a subwavelength-patterned multimode region for an extended bandwidth compared to conventional designs [101]. In this chapter, this concept is extended to the MIR, where demand for greater optical bandwidth is required. These MMIs were characterised using a method of embedding the MMIs in MZIs and are demonstrated with a twofold bandwidth improvement over conventional MMI designs.

5.2 Design

From Chapter 2, recall that the two-fold imaging point of an MMI occurs at the point $z = 3L_\pi/2$ (for propagation in the z-direction), where L_π is the beat length of the MMI. Therefore, to create a 3 dB power splitter based on an MMI, the length of the MMI needs to be $L_{\mathrm{MMI}} = 3L_\pi/2$. Since the length of the MMI is fixed in a given design, any variance in L_π will result in a loss of performance.

However, further recall that the beat length L_π of an MMI is given by

$$L_\pi \approx \frac{4w_e^2}{3\lambda} n_{\text{eff,w}}$$ (5.1)

where $n_{\text{eff,w}}$ is the effective index of the waveguide and w_e is the effective width of the multimode region (slightly larger than its actual width, w_{MMI}). Examining each factor in Equation 5.1 individually, evidently L_π has an explicit wavelength dependence from the factor $1/\lambda$. From Equation 2.17 in Chapter 2, w_e has some wavelength-dependence but for TE-polarised waveguides and to a first-order approximation, $w_e \approx w_{\text{MMI}}$ and is thus determined by the device design. As a consequence, it can be considered as approximately invariant with respect to wavelength. The effective index of the waveguide $n_{\text{eff,w}}$ is determined by waveguide materials and their dimensions (such as thickness). As such, for an SOI waveguide, its value is $n_{\text{Si}} > n_{\text{eff,w}} > n_{\text{silica}}$ to be guided (typically $n_{\text{eff,w}} = 2.5 - 3$). As such, the change in the value $n_{\text{eff,w}}$ is comparatively small to $1/\lambda$. Therefore, L_π has strong dependence on λ and accordingly it has a peak performance at the central wavelength for which it is designed. This means that, as the wavelength changes from its optimal value, the MMI is effectively detuned and its performance degrades.

To overcome this limitation, the multimode region of the MMI can be replaced with a periodic grating, as shown in Figure 5.1. In this design, the multimode region of the MMI now consists of an array of silicon strips that span the width of the MMI. These strips have a thickness h, equal to the thickness of the silion layer, and a length in the z-direction of a. The silicon strips alternate with regions that are fully etched to the SiO$_2$ layer and are filled with air. Recalling from Section 2.1.5 in Chapter 2, if the period of this grating Λ is sufficiently small with respect to λ, it can be considered as a
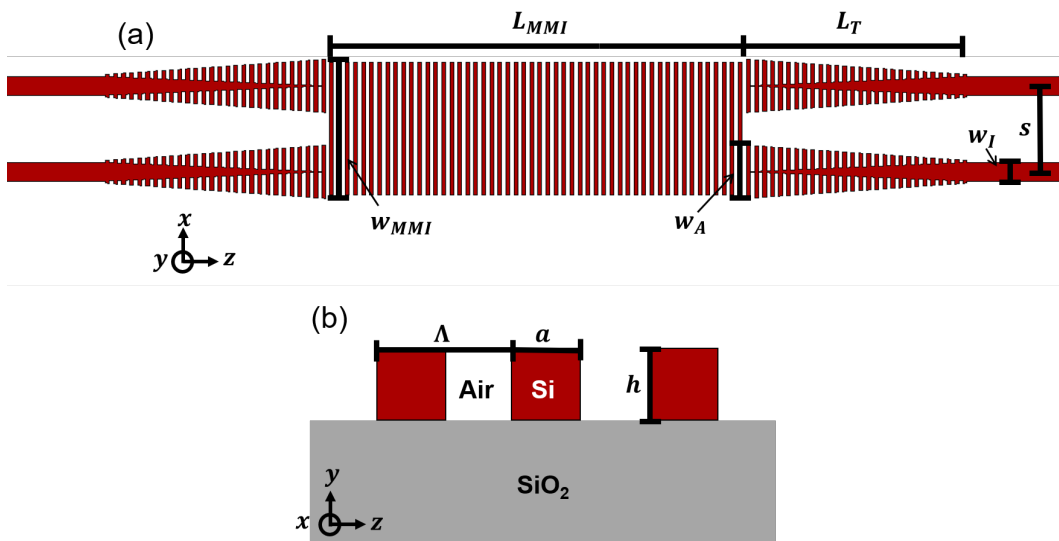
FIGURE 5.1: Structure of the SWG-MMI: (a) top-view diagram; (b) cross-sectional view of the multimode region.

subwavelength grating (SWG) with a refractive index tensor n_{SWG}:

$$n_{\text{SWG}}^2 = \begin{bmatrix} n_{xx}^2 & 0 \\ 0 & n_{zz}^2 \end{bmatrix} \tag{5.2}$$

where the grating is considered in two-dimensions only. It is this anisotropy of the metamaterial that gives the SWG-MMI its advantageous properties. The beat length of the MMI is now given by [101]:

$$L_{\pi,\text{SWG}} \approx \frac{4w_e^2}{3\lambda} \frac{n_{zz}^2}{n_{xx}}. \tag{5.3}$$

In contrast to the beat length of the conventional MMI given in Equation 5.1 (where $n_{zz} = n_{xx}$), the factor of n_{zz}^2/n_{xx} is roughly proportional to variation of L_π with wavelength, thereby cancelling out the wavelength dependence. Furthermore, the value of n_{zz}^2/n_{xx} is two or three times smaller than $n_{\text{eff,w}}$ and this reduces the beat length and the length of the MMI correspondingly. This is shown in Figure 5.2, in which the beat lengths of a conventional MMI and a SWG-MMI with the same w_{MMI} have been calculated by Equations 5.1 and 5.3, with n_{xx} and n_{zz} approximated by the Rytov equations, Equations 2.29(a-b) in Chapter 2, for illustrative purposes. Figure 5.2 also shows the effect of the duty cycle a/Λ on the beat length, as it changes the values of n_{xx} and n_{zz}. By designing the pitch and duty cycle, the bandwidth of the MMI is enhanced by tuning the dispersion of the SWG metamaterial to compensate for $1/\lambda$.

While each of n_{xx} and n_{zz} can be approximated by the Rytov equations as in the

FIGURE 5.2: Beat length of conventional and SWG-MMIs, for a silicon or silicon-air multimode region with $w_{\text{MMI}} = 7.7\,\mu\text{m}$, using n_{xx} and n_{zz} approximated by the Rytov equations.

previous chapter, in this chapter it was more appropriate to utilise the full Floquet analysis of the grating. In particular, this provides a more accurate value for n_{xx} and consequently also provides greater accuracy for the calculation of the beat length and the MMI performance. The simulations of the MMI structure were carried out by collaborators at the University of Málaga using their in-house Floquet mode simulation tool, Fourier expansion simulation environment (FEXEN). FEXEN is based on the Fourier eigenmode expansion method and is optimised for the simulation of periodic structures [202].

To design the SWG-MMI, the initial geometry was first defined based on the parameters defined in Figure 5.1. The duty cycle of the grating was first fixed to 50%, as this maximised both the size of the silicon strips and the etched regions between them, which was important to facilitate the etch step in the fabrication. Single-mode strip waveguides ($w_I = 1.1\,\mu\text{m}$, $h = 500\,\text{nm}$) were used at the input and outputs of the MMI. These waveguides were linearly tapered with a SWG cladding (with the same a and Λ as the multimode region) to provide an adiabatic transistion from the strip waveguide to the multimode region, ending at a width $w_A = 3.1\,\mu\text{m}$. The taper structure was "inversed" over the taper length $L_T \sim 13\,\mu\text{m}$, such that the strip waveguide decreased in width as the SWG cladding increased in width. To ensure that there was negligible coupling between the two inputs or the two outputs, both inputs and both outputs were separated by a distance $s = 5.1\,\mu\text{m}$.

Using the Floquet mode simulations, the grating period that resulted in the smallest variation in beat length of the design wavelength range of $\lambda = 3 - 4\,\mu\text{m}$ was $\Lambda \sim 465\,\text{nm}$. Using this period and a fixed duty cycle, the critical design parameters are then the length and width of the multimode region itself. These dimensions were iteratively optimised, yielding an optimal design with $w_{\text{MMI}} = 7.7\mu\text{m}$ and a grating length of 52 periods (equivalent to $L_{\text{MMI}} \sim 24.2\,\mu\text{m}$).

The performance of an MMI can be defined by three metrics: the insertion loss IL, the imbalance IB, and the phase error ϕ_{err}. The insertion loss is the amount of optical power lost by the inclusion of the MMI in the device, such that if the overlap of the input field with the fundamental mode of the MMI output waveguides is A and B respectively, then the insertion loss is

$$IL\,(\text{in dB}) = -10\log_{10}(|A|^2 + |B|^2). \tag{5.4}$$

For a perfect 3 dB (50:50) power splitter, the input power is split equally between the outputs. Therefore, the imbalance is then the ratio of the output powers:

$$IB\,(\text{in dB}) = -10\log_{10}\left(\frac{|A|^2}{|B|^2}\right). \tag{5.5}$$

For an ideal 50:50 beamsplitter (with $IL = IB = 0\,\text{dB}$), $A = B = 1/\sqrt{2}$. The phase error is simply the deviation of the phase difference between the MMI outputs $\Delta\phi$ from its ideal value of $\pi/2$, i.e. $\phi_{\text{err}} = \Delta\phi - \pi/2$.

Based on these metrics, the simulated performance of the optimal SWG-MMI design is evaluated in Figure 5.3. Evidently, the SWG-MMI shows excellent performance, with all three performance metrics close to zero over most of the wavelength range.

FIGURE 5.3: Simulated performance of the SWG-MMI with $W_{MMI} = 7.7\,\mu m$, 50% duty cycle and $L_{MMI} = 52\Lambda$, where $\Lambda \sim 465\,nm$.

5.3 Fabrication

The designed MMI was fabricated on an SOI sample with a 500 nm-thick Si layer using the *e*-beam lithography and dry etching techniques outlined in Chapter 3. In particular, the samples were overetched (i.e. etched beyond the end of the silicon layer), in an attempt to avoid the rounding at the bottom of grating features and to ensure that no silicon remained in the etched areas. As it was expected that this may cause some lateral etching in addition to the vertical etch, designs with variations in the critical parameters (i.e. the number of grating periods and the duty cycle of the SWG) around the simulated optimal values were included. Specifically, the duty cycle was biased by $+5\%$ to account for lateral etching, then the duty cycle was varied in different structures by $\pm 10\%$. The number of periods in the grating was varied over $42 - 55$, to investigate the tolerance of the MMI performance to its length. Figure 5.4 shows the best-performing measured device (discussed later), after fabrication. The measured duty cycle of the grating in this device in the SEM was $\sim 54\%$, suggesting that the effect of the lateral etching was overestimated.

The fabricated MMIs were embedded in asymmetric MZIs to evaluate their performance [112], with a path difference of $\Delta L = 20\,\mu m$ between the arms. Reference waveguides were used to obtain a reference for the coupling, propagation and bending losses in the MZIs. For comparison, conventional MMIs designed for $\lambda = 3.4\,\mu m$ (with $w_{MMI} = 6.85\,\mu m$ and $L_{MMI} = 76\,\mu m$) were also fabricated in asymmetric MZIs with the same path difference.

FIGURE 5.4: SEM micrograph of the best-performing fabricated SWG-MMI.

In addition, for broadband measurements of the devices, the input and output waveguide facets were exposed and polished to enable fibre-to-chip butt coupling. However, the preparation of the chip ahead of polishing introduced a challenge into the fabrication. The previously used standard process was to apply a layer of protective resist (AZ2070) via spin coating. Generally, this layer would be comparatively thick (several micrometres) and hard-baked (7 minutes at $110°C$) to fully protect the surface of the chip. This protective layer could then be removed by cleaning with acetone, IPA, and water.

For the first batch of devices, after the removal of the resist, it was found that the subwavelength features in the MMIs were partially broken or even completely removed (for example, a SWG-MMI that received comparatively minor damage is shown in Figure 5.5). No damage had been noticed in inspections of the chip surface under optical microscope during the polishing process (i.e. before removal of the resist) and acetone would not cause damage to the silicon layer in this way. Therefore, it was concluded that the gratings were breaking at some point during the processing to polish the waveguide facets and were held in place by the resist, only to come away when the resist itself was removed. As the polishing process uses multiple heating steps, the hard-baked resist (that had filled the gaps within the grating during the spin coating) was likely expanding during heating and applying stress to grating features.

FIGURE 5.5: SEM micrograph of a SWG-MMI that received damage to the subwave-
length features, after polishing with AZ2070 as a protective layer.

Since these features are much smaller than other structures in the silicon layer (such as the strip waveguides), they would be much weaker and therefore easier to break, which provided an explanation for only damage to the grating features being observed.

To fix this issue, the protective layer was no longer spin coated onto the surface of the chip. Instead, heated liquid wax was manually applied; even heated, the wax was likely too viscous to fill the gaps in the grating so any expansion after solidification would not apply significant stress to the grating. No further breakages were observed after switching to the wax. The wax was simply removed by cleaning with acetone, IPA and water.

5.4 Characterisation using MZIs

As mentioned above, the characterisation of the MMIs (both conventional and subwavelength-structured) involved including them in asymmetric MZIs with an arm length difference $\Delta L = 20\,\mu\text{m}$, for an easily measurable FSR $\Delta\lambda_{\text{FSR}} \approx 140\,\text{nm}$ (the spectrum was measured at $2\,\text{nm}$ wavelength intervals and the maximum manufacturer-specified laser linewidth was $< 10\,\text{cm}^{-1}$, equivalent to $\sim 11\,\text{nm}$ at $\lambda = 3.4\,\mu\text{m}$). The normalised transmission spectra of the MZIs can be modelled analytically, including accounting for wavelength-dependent errors in A, B and $\Delta\phi$ where their value differs from their theoretical ideal to represent a real device:

$$A(\lambda) = \frac{1}{\sqrt{2}} + A_{\text{err}}(\lambda); \tag{5.6a}$$

$$B(\lambda) = \frac{1}{\sqrt{2}} + B_{\text{err}}(\lambda); \tag{5.6b}$$

$$\Delta\phi(\lambda) = \frac{\pi}{2} + \phi_{\text{err}}(\lambda). \tag{5.6c}$$

By fitting these analytical model to measured spectra, the performance metrics can be extracted. Note that, compared to other common methods such as cascading MMIs to find the insertion loss, using the MZIs is the only way to determine the phase error of the MMIs. Furthermore, cascading the MMIs would require multiple measurements to determine the insertion loss of each MMI design, whereas embedding the MMIs in MZIs required only two measurements (one for each MZI output), facilitating much faster characterisation.

The analytical model for the 2×2 MZI used, as shown in Figure 5.6, is detailed below. As shown in Chapter 2, the electric field at the input and output ports of the MZI are related by the transfer matrix of the MZI $\mathbf{T}_{\text{MZI}}$, where $\mathbf{T}_{\text{MZI}}$ is the product of

FIGURE 5.6: Schematic of the asymmetric MZIs used for the characterisation of the MMIs, where $\Delta L = 20\,\mu m$. The bend radius in the MZI arms is $30\,\mu m$.

the individual transfer matrices of each component in the MZI:

$$\begin{bmatrix} E_{\text{out},1} \\ E_{\text{out},2} \end{bmatrix} = \mathbf{T_{MZI}} \begin{bmatrix} E_{\text{in},1} \\ E_{\text{in},2} \end{bmatrix} \tag{5.7}$$

with

$$\mathbf{T_{MZI}} = \mathbf{T_{MMI}}\mathbf{T_{arms}}\mathbf{T_{MMI}}. \tag{5.8}$$

The MMI transfer matrix, $\mathbf{T_{MMI}}$, is expressed in terms of A, B and $\Delta\phi$ and is applicable for any 2×2 MMI (including both conventional and subwavelength-structured):

$$\mathbf{T_{MMI}} = \begin{bmatrix} A & B\exp(i\Delta\phi) \\ B\exp(i\Delta\phi) & A \end{bmatrix}. \tag{5.9}$$

The transfer matrix for the arms of the MZI, $\mathbf{T_{arms}}$, is determined by the difference in phase by travelling path lengths L and $L + \Delta L$ through identical waveguides with modal propagation constants β:

$$\mathbf{T_{arms}} = \begin{bmatrix} \exp[i\beta(L+\Delta L)] & 0 \\ 0 & \exp(i\beta L) \end{bmatrix} = \exp(i\beta L)\begin{bmatrix} \exp(i\beta\Delta L) & 0 \\ 0 & 1 \end{bmatrix}. \tag{5.10}$$

For simplicity, assuming light is launched into only one MZI input (as was the case in the experimental characterisation), then

$$\begin{bmatrix} E_{\text{in},1} \\ E_{\text{in},2} \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \tag{5.11}$$

and so Equation 5.7 can be reduced to

$$\begin{bmatrix} E_{\text{out},1} \\ E_{\text{out},2} \end{bmatrix} = \exp(i\beta L)\begin{bmatrix} A^2\exp(i\beta\Delta L) + B^2\exp(2i\Delta\phi) \\ AB\exp(i\Delta\phi)[\exp(i\beta\Delta L) + 1] \end{bmatrix}. \tag{5.12}$$

Since the output power $P \propto |E|^2$, the normalised powers at the output ports of the MZI are then:

$$\begin{bmatrix} P_{\text{out},1} \\ P_{\text{out},2} \end{bmatrix} = \begin{bmatrix} A^4 + B^4 + 2A^2B^2 \cos(\beta \Delta L - 2\Delta\phi) \\ 2A^2B^2[1 + \cos(\beta \Delta L)] \end{bmatrix}. \tag{5.13}$$

Consequently, the output powers of the 2×2 MZI is expressible in terms of A, B and $\Delta\phi$. By varying A, B, and $\Delta\phi$, the analytical model can be fitted to the experimentally measured MZI transmission spectra; the imbalance, insertion loss and phase error can be directly calculated from these resulting values of A, B and $\Delta\phi$.

5.5 Results and discussion

The MZIs were characterised using the butt-coupled OPO setup detailed in Chapter 3, over a wavelength range of $\lambda = 3.0 - 3.7\,\mu\text{m}$ (limited by the upper wavelength of the source). Note that, for the avoidance of ambiguity in this section, the experimentally measured transmissions are labelled $P_{\text{exp},1}$ and $P_{\text{exp},2}$, while the fitted analytical models (equivalent to $P_{\text{out},1}$ and $P_{\text{out},2}$ on a dB-scale) are labelled $P_{\text{mod},1}$ and $P_{\text{mod},2}$ for the corresponding outputs.

To show the effect of the noise in the measured spectra, the raw transmission spectra for the best-performing SWG-MMI (identified later) and the reference waveguide used for normalisation are shown in Figure 5.7 over the wavelength range of the source, $\lambda = 2.5 - 3.7\,\mu\text{m}$. The reference waveguides had a similar layout to the MZIs to remove the coupling, propagation and bending losses (although there was some slight additional loss for the MZI, due to slightly longer path length, but this loss contribution was expected to be minimal). The dips in transmission around

FIGURE 5.7: Raw transmission spectra for two outputs of the MZI based on the best performing SWG-MMI, and the raw transmission spectrum of the reference waveguide used for normalisation.

$\lambda = 3.4\,\mu\text{m}$ are attributed to organic contamination, likely due to residues leftover from the polishing process. These absorptions were present in all spectra and were cancelled out in the normalisation. For wavelengths below $3.2\,\mu\text{m}$, the detector noise floor relative to the laser power limited the measurement of the MZI extinction. To avoid overestimating the value of the insertion loss and imbalance in this region during the fitting, the noise floor was subtracted from the spectra; this would make negligible difference to the rest of the spectrum, but yield a better fit for the noise-limited region. To estimate the noise floor, the output of the lock-in amplifier was measured with no input power (i.e. with the laser turned off) and then dividing that value by the output of the lock-in amplifier by the measured transmission spectrum of the straight normalisation waveguide. The rest of the measurements were carried out over the range with large measurable transmission, $\lambda = 3.0 - 3.7\,\mu\text{m}$.

The experimentally measured transmission spectra for a selection of MZIs that used the different SWG-MMI designs are shown in Figure 5.8. All of the MZIs in Figure 5.8 used embedded MMIs with the same duty cycle (55%), but with different lengths; across the measured devices, the 55% duty cycle devices were generally better.

As expected, Figure 5.8 shows that increasing the number of periods in the grating (and therefore the length) changes the MMI performance, the beat length will be the same for the same grating design. The best-performing MMIs will result in interferograms with the aligned alternating peaks (showing low phase error), peaks with equal depths (for low imbalance) and a minimal zero-point (i.e. the top of interferogram close to 0 dB, showing for insertion loss). From these criteria, the MZI for the SWG-MMI with $L_{\text{MMI}} = 55$ periods has the best performance.

The experimentally measured normalised transmission spectra for the MZIs using the best-performing SWG-MMI and the comparison conventional MMIs are shown in Figures 5.9(a) and (b) respectively. To fit $P_{\text{mod},1}$ and $P_{\text{mod},2}$ to $P_{\text{exp},1}$ and $P_{\text{exp},2}$ respectively, each of $A(\lambda)_{\text{err}}$, $B_{\text{err}}(\lambda)$ and $\phi_{\text{err}}(\lambda)$ were varied and the modelled MZI spectra were calculated accordingly. In order to balance accuracy without overfitting, each error parameter was approximated as a 6$^{\text{th}}$-order polynomial with respect to wavelength (giving 7 coefficients to optimise per parameter or 21 in total for the entire problem). The fitting of the spectra was then treated as a minimisation problem (i.e. minimising the difference between the experimental and modelled spectra), solved by varying the values of the coefficients using a genetic algorithm in MATLAB [203].

Additionally, the finite linewidth of the measured spectra also limited the extinction of the MZIs, while the analytical model would extend to negative infinity and the subsequent fit would significantly overestimate the performance metrics. As a consequence, the laser linewidth had to be accounted for in the model. This was achieved by using a Lorentzian function to represent the linewidth:

$$f(\nu) = \frac{1}{1 + g^2(\nu)} \quad \text{with} \quad g(\nu) = \frac{\nu - \nu_0}{\Delta \nu} \tag{5.14}$$

where ν is the frequency, ν_0 is the central frequency of the distribution and $\Delta\nu$ is the full-width at half maximum equal to the linewidth, fitted to $\Delta\nu = 0.7\,\mathrm{cm}^{-1}$ (this is less than the linewidth specified by the manufacturer, but that value was only an upper limit). This function was then convoluted with the outputs of the model, for each extinction at different values of ν_0 (or equivalently, the wavelength at which each extincition occurs). This correction was specifically applied to the fitted analytical MZI spectra, so it is only present $P_{\mathrm{mod},1}$ and $P_{\mathrm{mod},2}$ in Figure 5.9.

From Figure 5.9, adjusting the model in this way resulted in a very good fit for the measured transmissions. The exception to this is that at wavelengths below
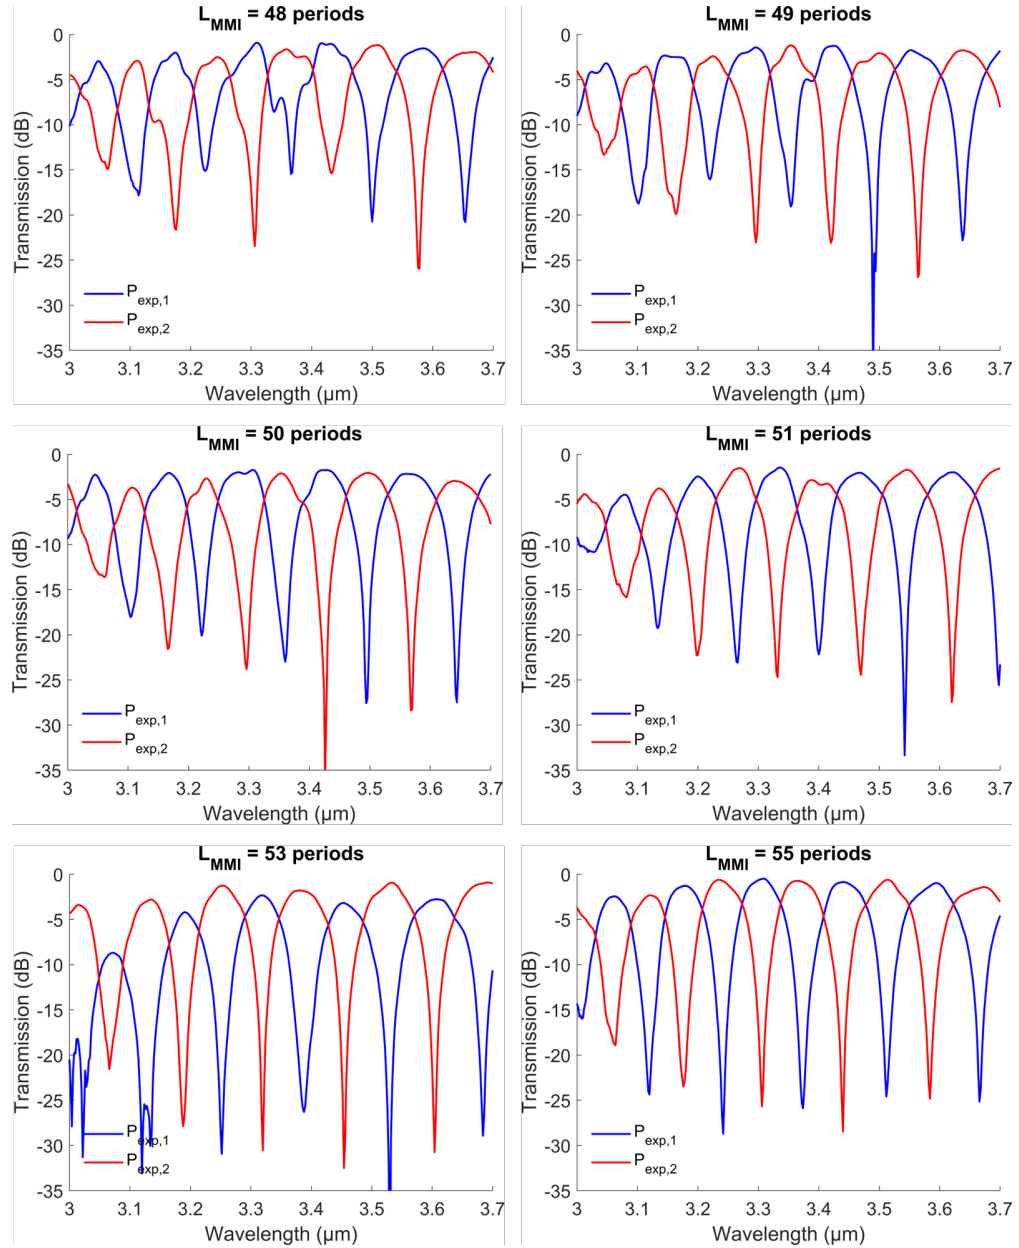
FIGURE 5.8: Measured transmission spectra for MZIs that used MMIs of different lengths, but all with a 55% duty cycle. All spectra have the noise floor subtracted from the experimental data.
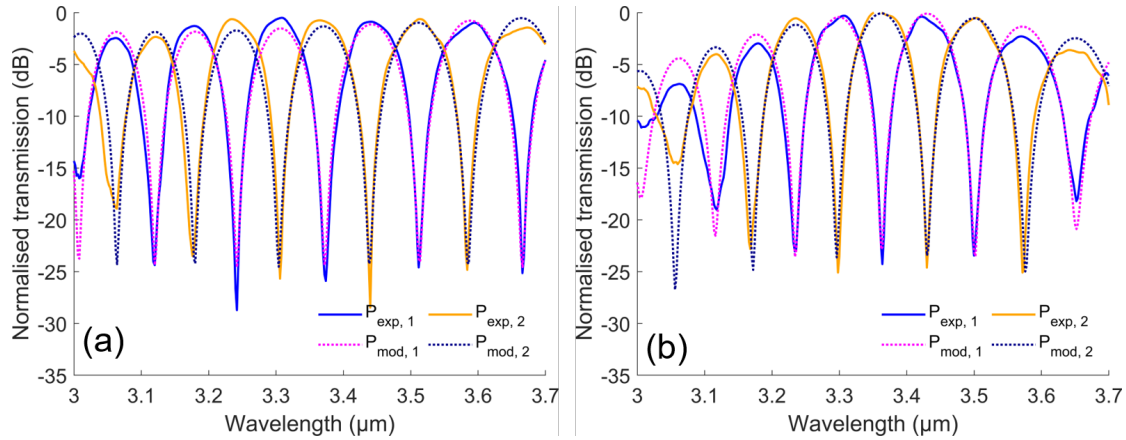
FIGURE 5.9: Transmission spectrum and the fitted model for MZIs with embedded (a) SWG-MMIs and (b) conventional MMIs. $P_{exp,1}$ and $P_{exp,2}$ are the experimentally measured outputs of the MZI; $P_{mod,1}$ and $P_{mod,2}$ are the fitted model for each output. For both (a) and (b), the noise floor has been subtracted from the experimental data.

3.1 μm, the signal-to-noise ratio was low and this limited close matching of the fit as in the rest of the spectrum. Consequently, the extracted insertion loss and imbalance are overestimated in this part of the spectrum. The cause for this is not known as it should be sufficiently above the SiO_2 near $\lambda = 2.7$ μm, but one likely possibility is that it is caused by sidelobes of the SWG metamaterial photonic bandgap extending further into the spectrum, as has been noted as an issue for other subwavelength devices in literature [204].

The extracted performance metrics (insertion loss, imbalance, and phase error) are shown in Figure 5.10 for the best performing fabricated device. This MMI differs from the device with the optimal simulated performance: duty cycle of 55% instead of 50% (due to the overestimated biasing); 55 period length rather than 52 ($L_{MMI} \sim 25.6$ μm rather than $L_{MMI} \sim 24.2$ μm); and $W_{MMI} = 7.6$ μm, not $W_{MMI} = 7.7$ μm. The reason that this design yielded a better performance is that the simulated devices assumed that the SWG was fully etched, i.e. that no silicon remained between the grating strips. However, as discussed in Chapter 3, the ICP-RIE etching of small holes can lead to the ideally square-bottomed holes being rounded and even incompletely etched. Although this was identified as a potential issue prior to the fabrication and an attempt was made to mitigate this (the devices were intentionally over-etched), it does not appear to be the case here. A residual layer of this thickness within the grating is very difficult (if not impossible) to image in the SEM and a FIB cut along the grating would smooth the edges of the residual layer, meaning that it cannot be observed directly. Rather, the presence of the residual layer was inferred from the inclusion of such a layer in the simulation, to match the experimental results. The lateral etching that was expected but not observed (see Section 5.3) further suggests that perhaps the MMIs should have been etched longer to fully remove the silicon in the etched areas.

Despite this, by accounting for a 15 nm layer of residual silicon between the
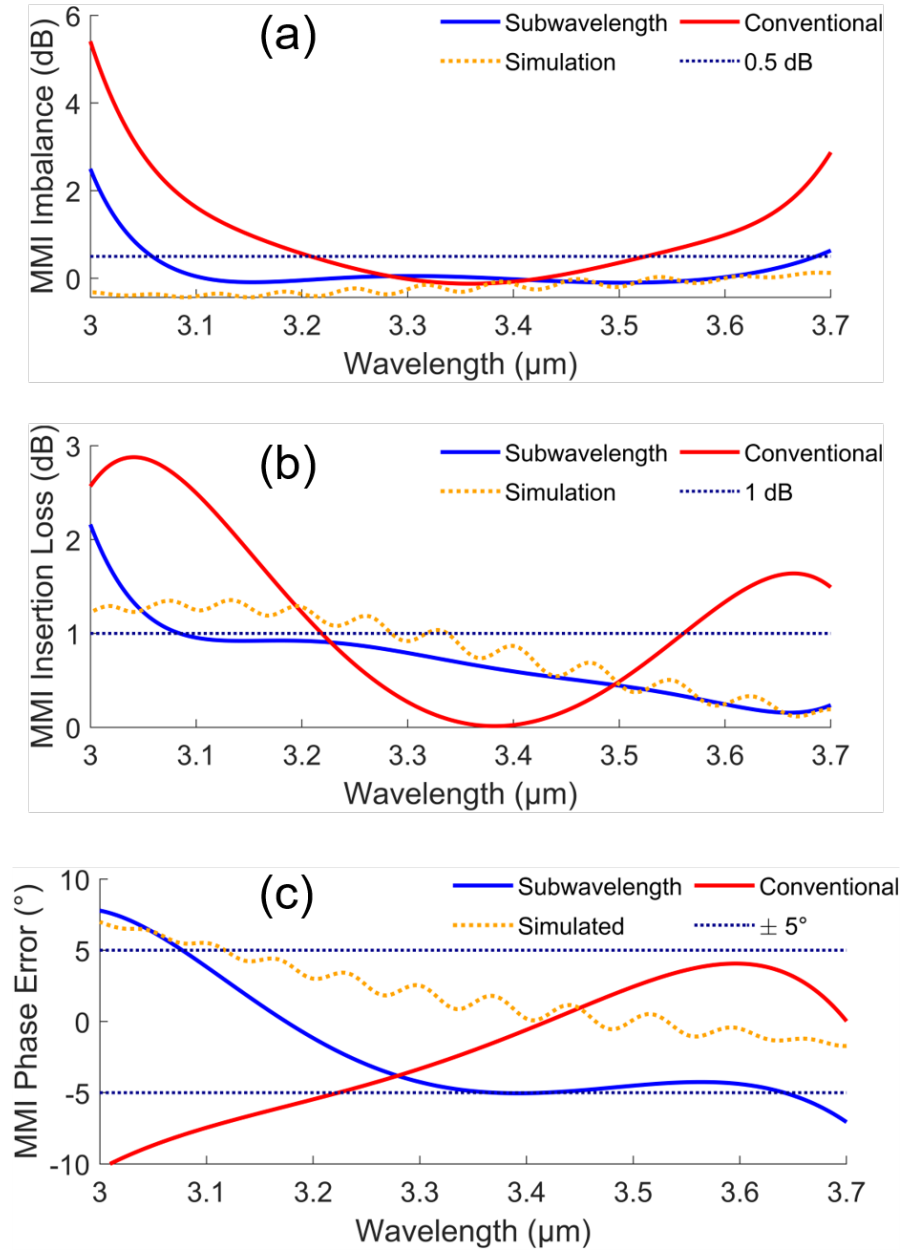
FIGURE 5.10: (a) The imbalance; (b) insertion loss; and (c) phase error of the SWG and conventional MMIs calculated from the measured MZI transmission spectra. The simulated MMI uses the estimated fabricated dimensions: a 15 nm residual silicon layer, WMMI = 7.6 μm, 55% duty cycle and 55 periods.

grating strips, the simulated performance matches closely with the measured performance, as shown in Figure 5.10. Considering the bandwidth for each of the performance metrics individually, in terms of the imbalance - Figure 5.10(a) - the usuable wavelength range for the SWG-MMI is almost double that of the conventional MMI. Defining $IB = 0.5$ dB as a acceptable threshold for the imbalance, the conventional MMI can cover 316 nm of the spectrum, while the SWG-MMI covers 628 nm. Likewise, for the insertion loss, the 1 dB-bandwidth of the SWG-MMI (616 nm) is approximately twice that of the conventional MMI (343 nm).

The difference between the two MMIs is less noticeable when considering the phase error, but the SWG-MMI outperforms the conventional MMI here as well. For $|\phi_{\text{err}}| < 5°$, the bandwidth of the SWG-MMI is 564 nm, in contrast to 476 nm for the conventional MMI. Furthermore, the phase error for the SWG-MMI is roughly uniform for $\lambda > 3.4\,\mu\text{m}$, whereas the phase error of the conventional MMI changes consistently over the spectrum.

5.6   Summary

By applying a subwavelength-structure to the multimode region of an MMI, a twofold bandwidth enhancement can be achieved over a conventional MMI around the MIR wavelength of 3.4 µm. Importantly, the device has a compact footprint (the multimode region is approximately a third of the length of the comparable conventional MMI) and has a simple fabrication of a single etch step in SOI. However, this design is intolerant to fabrication errors; in this work, MMI length, etch profile and duty cycle have all been shown to adversely affect performance. Further work is needed to ensure to improve the fabrication tolerance.

Despite this, the straightforward fabrication process to achieve this bandwidth is notable as it does not require additional materials and complex fabrication (like the tunable MMIs in Chapter 2), nor is it reliant on the inherent dispersion of the material to achieve broadband performance (like the SiGe MMIs in Chapter 2). This means that the same approach could be applied to different material platforms to be extended out to longer MIR wavelengths. In particular, the relative size of the subwavelength features can increase with longer wavelengths, relaxing the fabricaton tolerances and thereby avoiding issues such as the residual silicon layer found here (although thicker layers could be limited by the aspect ratios of the etched features) or offering the flexibility to choose a period to ensure the sidelobes of the SWG bandgap does not overlap with the operating wavelength range.

Chapter 6

Low-loss broadband Y-junction power splitters

6.1   Introduction

Y-junctions are introduced in Section 2.1.3.1 of Chapter 2. As power splitters, they are in principle relatively insensitive to wavelength when compared to other devices such as MMIs or direction couplers. However, in practice the utility of a Y-junction is limited by the size of the tip of the "Y" where the input waveguide separates into two output waveguides. Ideally, this tip would be infinitely small but it has a finite size due to the limited resolution of lithography. Zhang *et al.* optimised the geometry of a Y-junction to have a low loss ($0.28 \pm 0.02\,\text{dB}$) despite this limitation for a NIR wavelength range of $\lambda = 1500 - 1580\,\text{nm}$ [82] by using a particle-swarm optimisation (PSO). In this chapter, this approach is extended to develop a Y-junction with low-loss over a 900 nm wavelength range in the MIR, a significant improvement over other 1×2 power splitters in the MIR (see Section 2.1.3).

6.2   Particle-swarm optimisation

Before considering the optimisation problem for the Y-junction, it is important to detail the algorithm used. The PSO was conceptually developed from the movement and social behaviour of creatures that exhibit swarming, such as flocks of birds or schools of fish [205]. The algorithm for the PSO has subsequently been applied to optimisation problems in electromagnetic design [206], as described below.

Initially, particles representing solutions to the optimisation problem are generated in a user-defined parameter space of M dimensions. M is chosen to be equal to the total number of parameters which can be varied in the optimisation and a

reasonable range of values for each parameter is input by the user. A figure-of-merit (FOM) function must be defined to evaluate the "goodness" of a particle's position in the parameter space against the goal of the optimisation. For each particle, the algorithm remembers the position at which it had the greatest value of the FOM, the personal best position in the m^{th} dimension R_m, where $m = 1, 2, ..., M$. For the swarm as a whole, the algorithm remembers the position with the highest FOM ever achieved by any particle, the global best position G_m. At the beginning of the optimisation, the particles are initialised with either random positions and velocities or user-input first guesses so, since these will be the only recorded positions, R_m and G_m are obtained from these starting positions. The algorithm then "flies" each particle for a time step Δt, treating the particles as a swarm of bees for example, updating both the particle's position and velocity after each step, and R_m and G_m if they have improved.

The velocity in the m^{th} dimension, v_m, of a particle depends on its current position, r_m, relative to R_m and G_m:

$$v_{m,\text{new}} = W v_{m,\text{old}} + C_{\text{c}} \tau_1 (R_m - r_m) + C_{\text{s}} \tau_2 (G_m - r_m). \tag{6.1}$$

The other terms in Equation 6.1 are essentially scaling factors for calculating the new velocity: W is termed the inertial weight, which dictates by how much a particle will maintain its previous velocity; C_{c} is the degree to which the particle is influenced by the memory of its best location (the "cognitive rate"), C_{s} is the degree to which the particle is influenced by the rest of the swarm (the "social rate"); and $\tau_{1,2}$ are two separate random numbers between 0 and 1. The position of the particle is simply updated using this velocity by

$$r_{m,\text{new}} = r_{m,\text{old}} + v_{m,\text{new}} \Delta t. \tag{6.2}$$

Evidently from Equation 6.1, the particles that are furthest from R_m or G_m will travel towards them fastest, gaining speed until they overshoot these positions. The velocity of the particle will then swap directions back towards the best positions, effectively gradually circling the optimal solution. With many particles, this is in concept similar to a random swarm of creatures approaching a single location. The rate at which the particle swarm converges on the solution will depend on the number of particles in the swarm, which is usually a balance with required computing resources. The PSO will end after a completion criteria is met, generally either after it has run for a given number of generations (the number of times the particle's position and velocity are updated) or after the FOM converges to a value with a specified tolerance to the change between generations.

6.3 Geometry optimisation

The Y-junction was to be designed for an SOI platform, with a 500 nm-thick Si layer that is fully etched to the underlying oxide layer. To demonstrate broadband low-loss in the MIR range, the Y-junction was designed to work across $\lambda = 2.8 - 3.7\,\mu\text{m}$ to correspond to the tunable wavelength range of the OPO in the broadband characterisation setup (see Chapter 3) and single-moded bandwidth of the waveguide.

6.3.1 Input and output waveguide

Before the PSO could be implemented for the Y-junction geometry, the input and output waveguides for the junction (used to calculate the input and output powers, and therefore insertion loss) had to be defined. The height and etch depth of the waveguides were fixed due to the design restrictions of the Y-juntion (in a 500 nm-thick strip waveguide).

The modal analysis of the waveguide was completed in MODE, using the approach described in Chapter 3, and it was simulated over the full range of the source in the experimental setup $\lambda = 2.5 - 3.7\,\mu\text{m}$. To determine the single-mode condition, the effective index of the TE_0 and TE_1 modes were calculated for varying waveguide widths and at different wavelengths; where the effective index of a mode was less than the refractive index of the cladding, the mode is not be guided. By using TE-polarised input light, only the TE modes need to be considered, so the single-mode condition is when the TE_0 mode is supported but the TE_1 mode is not.

Figure 6.1 shows the result of this analysis, with Figures 6.1(a) and (b) for the TE_0 and TE_1 modes respectively. The cladding refractive index is shown by the dashed white line, which is the contour of the index equal to the smallest value of the index of the cladding, $n_{\text{silica}} = 1.399$ at $\lambda = 3.7\,\mu\text{m}$. Evidently, there is no waveguide width at
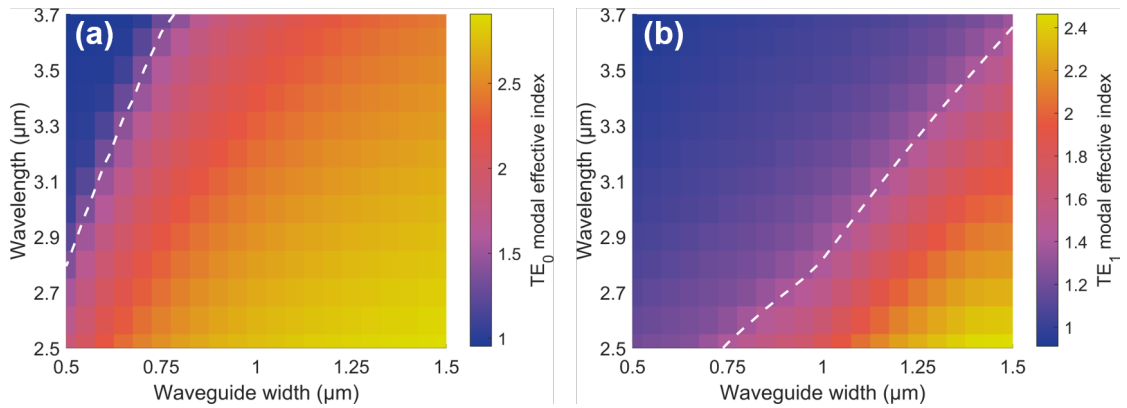
FIGURE 6.1: Modal effective indices of the TE_0 and TE_1 modes in the input and output Y-junction waveguides, for a fully-etched SOI waveguide with $h = d = 500\,\text{nm}$. The white dashed lines are the contours for $n_{\text{silica}}(\lambda = 3.7\,\mu\text{m} = 1.399)$, where it is smallest.

which the TE_0 mode is supported and the TE_1 is not across the entire wavelength range. However, a waveguide width $w = 1\,\mu m$ meets the single-mode condition across the wavelength range $\lambda = 2.8 - 3.7\,\mu m$, while allowing for the greatest confinement of the fundamental mode.

6.3.2   Y-junction geometry

Rather than use an random geometry for as a starting point for the optimisation, the initial Y-junction was pre-defined to allow comparison before and after the optimisation. The initial geometry is shown in Figure 6.2, in which the shape of the junction depends on the width at 13 equally-spaced intervals along a fixed $4\,\mu m$ length: it is only within this length the shape is optimised. The number of widths was equal to the number of dimensions in the parameter space. Therefore, this number of parameters was selected to balance the degrees of freedom in shaping the geometry with the overall optimisation time. Between these intervals, the edges of the junction are found by interpolation between the width values using a spline curve. For the initial geometry, a sigmoid function or "S"-curve was used to calculate each width as a function along the length and are given in Table 6.1. The separation between the two output waveguides is fixed as $0.2\,\mu m$, which is the approximate resolution of UV lithography.

FIGURE 6.2: Geometry of the Y-junction before the PSO.

w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8	w_9	w_{10}	w_{11}	w_{12}	w_{13}
1000	1008	1022	1057	1143	1323	1600	1877	2057	2143	2178	2192	2200

TABLE 6.1: The widths w_m of the initial Y-junction geometry in nanometres, for $m = 1, 2, ..., 13$.

    The PSO was implemented using the Lumerical varFDTD solver, using functions in-built to the software. For the 13 different widths, the PSO was then solved in a 13-dimension parameter space. In each generation, a swarm of 10 particles (each representing a trial geometry) was generated and the fractional transmitted optical power $P_{\text{trans}}(\lambda)$ through the Y-junction was monitored. For a faster simulations, the

FIGURE 6.3: Screenshot of the simulation of the Y-junction in varFDTD. The orange rectangle shows the solver, the yellow lines and rectangles are the simulation monitors and the arrow shows the source.

simulation was setup as symmetric and therefore $P_{\text{trans}}(\lambda)$ only had to be calculated for a single output (this setting is shown by the green rectangle in Figure 6.3).

The FOM for the optimisation was defined based on the simulated transmission $P_{\text{trans}}(\lambda)$ and was to be maximised over the PSO. The aim of the PSO was twofold: to decrease the insertion loss of the Y-junction over all λ, but also to increase the uniformity of this loss. By defining

$$\Theta = -10 \log_{10}(\overline{P}_{\text{trans}}) \tag{6.3}$$

where $\overline{P}_{\text{trans}}$ is the arithmetic mean of $P_{\text{trans}}(\lambda)$ and

$$\delta = 10 \log_{10}(\max[P_{\text{trans}}(\lambda)] - \min[P_{\text{trans}}(\lambda)]) \tag{6.4}$$

then the FOM used for the PSO was

$$\text{FOM} = \frac{3}{5}\delta + \frac{1}{\Theta - 3\,\text{dB}}. \tag{6.5}$$

The term Θ is equivalent to the average insertion loss of the Y-junction and $3\,\text{dB}$ was subtracted from it to account for P_{trans} being only for one output; even in an ideal lossless case, the best value would be $\Theta = 3\,\text{dB}$. The factor $3/5$ is used to scale the terms and balance their contribution to the FOM in later generations. The FOM was designed such at, for early generations where there are large differences in P_{trans}, the optimisation algorithm will reduce the average insertion loss first; the easiest way to increase the FOM is for the value of Θ to approach $3\,\text{dB}$. However, in later generations once the average insertion loss has decreased, the maximum and minimum values of P_{trans} will be much closer together and correspondingly the value of δ becomes large, which increases the uniformity. This prevents the optimisation algorithm from ending

too early and only reducing the overall loss, but not the uniformity. Equally, this also avoids a solution where δ has been maximised but the average insertion loss is high.

Alternate FOMs can be used for other optimisation problems, or to weight this optimisation problem differently. For example, it might be acceptable for the final design to have a higher loss if the uniformity in the loss is minimal. Furthermore, it should be noted that the reproducibility of the result is affected by the end conditions of the optimisation (such the tolerance on the FOM, or the preset number of generations) and the limits imposed on the parameter space. These limits, input before the algorithm is run, define boundaries on the acceptable positions (in each parameter space) the particles can take and it can reduce the time taken for the optimisation. However, harsh limits can also cause the simulation to converge to the limit, which may not represent the optimal value. The limits for each parameter are listed in Table 6.2; no width was allowed to be thinner than the single-mode input waveguide and that the first and last widths had fairly constrained limits. The final consideration for reproducibility is that, given the random nature of the initial particle positions within a generation, optimisations run using random number generators that have different seeds may lead runs that converge to different results, particularly for a short number of generations or relaxed tolerances on the FOM.

Width	w_1	w_2	w_3	w_4	w_5	w_6	w_7
Lower	1000	1000	1000	1100	1100	1100	1400
Upper	1100	1400	1500	1800	2000	2500	2700

Width	w_8	w_9	w_{10}	w_{11}	w_{12}	w_{13}
Lower	1600	1800	1800	2000	2000	2200
Upper	3000	3000	3000	2700	2400	2300

TABLE 6.2: The limits of the widths w_m input into the PSO in nanometres, for $m = 1, 2, ..., 13$.

The PSO was run for 41 generations of 10 particles each, before the algorithm was concluded because it had reached the specified tolerance, as shown in Figure 6.4(a). The zeroth generation to corresponds the FOM for the initial geometry of the waveguide. The improvement in Y-junction performance is shown in Figure 6.4(b) with the insertion loss now $\sim 0.2\,\text{dB}$ over the wavelength range of the simulations, in constrast to the inital geometry where the insertion loss is strongly wavelength-dependent. The resultant Y-junction geometry is illustrated in Figure 6.5, with the widths specified in Table 6.3.

FIGURE 6.4: The outcome of the PSO of the Y-junction geometry: (a) the FOM of the PSO over 41 generations; (b) simulated insertion loss of the Y-juction before and after optimisation.

FIGURE 6.5: Geometry of the Y-junction after the PSO.

w_1	w_2	w_3	w_4	w_5	w_6	w_7	w_8	w_9	w_{10}	w_{11}	w_{12}	w_{13}
1000	1236	1140	1785	1705	2370	2631	2864	2621	2633	2544	2334	2200

TABLE 6.3: The widths w_m of the optimised Y-junction geometry in nanometres, for $m = 1, 2, ..., 13$.

The underlying reason for the improvement in performance is demonstrable by the field within the Y-junction, shown in Figure 6.6. From Figure 6.6(a), the initial Y-junction has a higher insertion loss because the electric field is concentrated in the gap between the output waveguides. This light cannot couple into the waveguides and thus leaks into the slab. Conversely, from Figure 6.6(b) the Y-junction with the optimised geometry this is substantially reduced, demonstrating that the improvement in insertion loss arises from overcoming the finite resolution of the fabrication.

FIGURE 6.6: $|E|$ along the length of the Y-junctions at $\lambda = 2.8\,\mu m$:
(a) the initial geometry; (b) the optimised geometry.

6.4   Results and discussion

To determine the characterisation loss, the Y-junctions were fabricated onto the chip using e-beam lithography and dry etching as described in Chapter 3, before the waveguide facets were prepared by cleaving and polishing. The layout of the chip design used "chains", which consisted of arrangements of the Y-junction and straight waveguides as illustrated in Figure 6.7.

FIGURE 6.7: Schematic illustration of the chain layout for determining the insertion loss of the Y-junctions.

In this layout, a series of separate waveguides were arrayed edge-to-edge across the chip, for fibre-waveguide butt coupling. In each subsequent waveguide, an increasing number of closed loops were added, each of which consisted of two Y-junctions used to split and re-couple the light (in effect, these were small symmetric MZIs with no phase difference between the arms): the result waveguides resemble chains. Consequently, the loss will increase between each subsequent waveguide by twice the insertion loss of the Y-junctions and the first waveguide with no chain loop can be used to obtain a reference for the coupling loss and some measure of the waveguide loss. In practice, as each loop is added some of the straight waveguide

length is lost, equivalent to the length of the loop, as the waveguide length was fixed by the width of the chip. Therefore, the calculated loss of the Y-junctions will be marginally underestimated as some of the loss is normalised as propagation loss. However, the removed lengths are short so unless the propagation loss of the fabricated waveguides is very high then this underestimate should be negligible. The advantage this method has for measuring the insertion loss is that the transmission through the waveguides will remain fairly high throughout the characterisation. This is not the case for alternate layouts such as a splitting tree, where the transmission will be halved at each subsequent measurement even with lossless splitters. This is important for ensuring consistent insertion losses between waveguides, as it is much easier to measure transmission with a large signal-to-noise ratio.

The measured insertion loss for the Y-junctions are shown in Figure 6.8. For transmission through each waveguide in the chain layout to determine the total insertion loss for the corresponding number of Y-junctions and then a linear fit was used to determine the loss per junction at each measured wavelength point. For example, Figure 6.8(a) shows the linear fit at $\lambda = 3.175\,\mu\text{m}$, for a fitted loss of $0.18\,\text{dB}$ per Y-junction. Note that facets of the waveguides corresponding to 0, 2 and 6
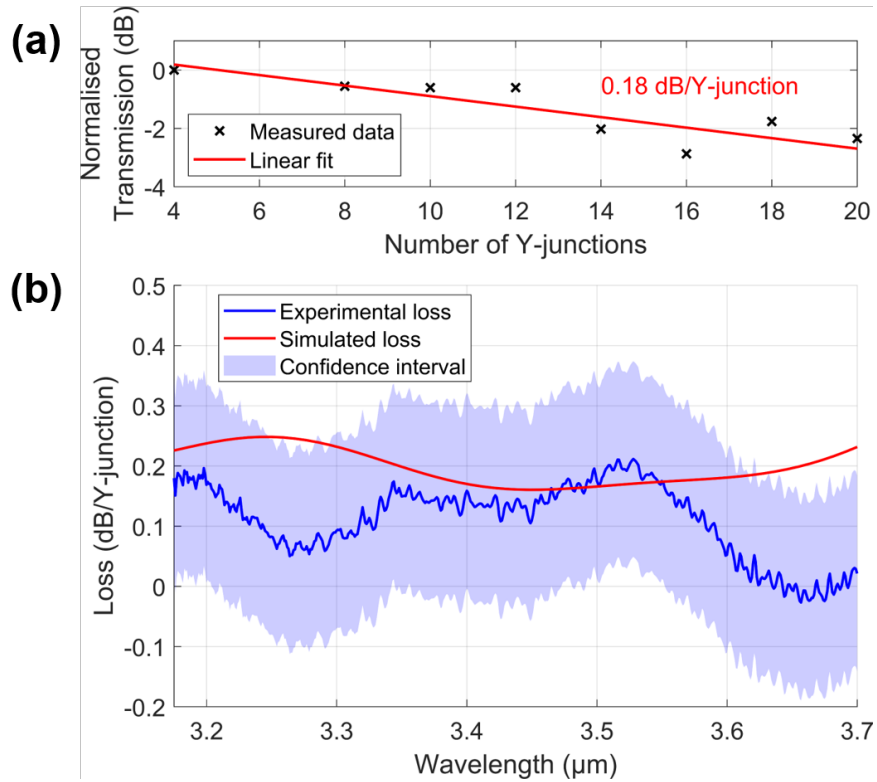
FIGURE 6.8: Measured insertion loss of the wideband Y-splitter over $2.8 - 3.7\,\mu\text{m}$. (a) Transmission as a function of the number of Y-junctions at $\lambda = 3175\,\text{nm}$. The transmissions for waveguides with 0, 2 and 6 Y-junctions are not included due fabrication defects causing substantially lower transmission. (b) The simulated and experimental loss per Y-junction as a function of wavelength. Experimental measurements for $\lambda < 3150\text{nm}$ are not shown.

Y-junctions sustained damage during polishing and consequently had poor measurable transmission, so these data points are not included in the fit and the measurement data is instead normalised to the waveguide with 4 Y-junctions.

Although the transmission through each waveguide was measured over the wavelength range of $2.8 - 3.7\,\mu m$, the transmission through the waveguides was very low for $\lambda < 3.15\,\mu m$. The reason for this is unclear but it may be in part caused by an absorption related to the SiO_2 layer, as it has a strong absorption around $\lambda = 2.6 - 2.9\,\mu m$ [207] if hydroxyl groups (O——H) are present [208]. Therefore, the loss spectrum in Figure 6.8(b) only shows the loss obtained for $\lambda = 3.15 - 3.7\,\mu m$. For this wavelength, the measured insertion loss was very low and less than $0.2\,dB$ over most of the spectrum. The confidence level of this measured loss is given by the largest error across the spectrum ($\pm 0.16\,dB$); as can be seen in Figure 6.8(b), the simulated loss agrees with the measured insertion loss within the uncertainty level across the majority of the spectrum. Note that, although the areas of negative loss within the error range imply that the Y-junction could result in gain in the chip, this is not physically realistic - Figure 6.8(b) is presented in this way solely to illustrate the large experimental error.

The comparatively large error is due to the small insertion losses themselves. Even for a large number of consecutive junctions, the total loss was fairly small and could be dominated by other losses in the experimental setup. This is particularly true for a butt-coupled chip such as the one used in this chapter, where the insertion losses are generally harder to control (as opposed to grating-coupled chips) due to potential inconsistencies between the quality of the waveguide facets after polishing or repeatable fibre-chip alignment. These additional introduced losses do not have to be large to introduce considerable error and possibly prevent the characterisation of the device, when the device losses are so small. Some strategies to improve the accuracy of the loss characterisation could have been achieved a layout that only required a single input (thereby only the output alignment would change) or by using chains with a greater number of Y-junctions included. Furthermore, a reference waveguide that accounted from the small extra loss from the loop S-bends (rather than a straight waveguide) must be used.

6.5   Summary

In this chapter, the geometry of a Y-junction was optimised to overcome practical fabrication limitations, to yield a wideband low-loss power splitter. The better geometry was the result of a particle-swarm optimisation, resulting in a fairly constant simulated insertion loss of $\sim 0.2\,dB$ over a MIR wavelength range of $\lambda = 2.8 - 3.7\,\mu m$. By fabricating the Y-junction on a SOI wafer, the insertion loss was experimentally determined to be less than $0.2\,dB$ over $\lambda = 3.15 - 3.7\,\mu m$, which agreed with the

simulation within the experimental error. These power splitters are important as basic component of many PICs, but in particular for mid-infrared spectrometers.

Chapter 7

Group IV membrane waveguides on transparent substrates

7.1 Introduction

The heterogeneous integration of dissimilar materials was discussed in Chapter 2 and, in particular, the use of transfer printing group IV materials onto new substrates was explored. Transfer printing offers a potential direction for providing a lower cladding with a transparency that spans a greater wavelength range than is possible with solely group IV materials. This could facilitate waveguides with a low-loss over a considerable wavelength range in the MIR, by minimising absorption losses from the waveguide mode overlapping with the cladding. Furthermore, zinc selenide (ZnSe) was suggested as a potential substrate because it has a transparency that extends beyond that of both Si and Ge, it can offer a large refractive index contrast to both materials and, for Ge, a small mismatch in the crystal lattice size.

This chapter details the transfer printing of group IV membranes onto crystalline ZnSe substrates for the reasons outlined above. Initially, Si membranes were used as a proof-of-concept despite Ge waveguides in principle offering the greater transparent bandwidth: firstly, Si waveguides are well-established at both NIR and MIR wavelengths, meaning that it is easier to control the fabrication processes and also determine the source of any potential losses; secondly, Si-based wafers (i.e. SOI) are more widely available commercially, enabling quicker progression in the development of the integration technique.

The transfer printing process used is outlined in Figure 7.1. The process operates by taking advantage of the reversible adhesion of the stamp [191] which is made from an elastomeric polymer, in this case polydimethylsiloxane (PDMS). For this stamp, the adhesive force between the stamp and the membrane depends on the speed at which
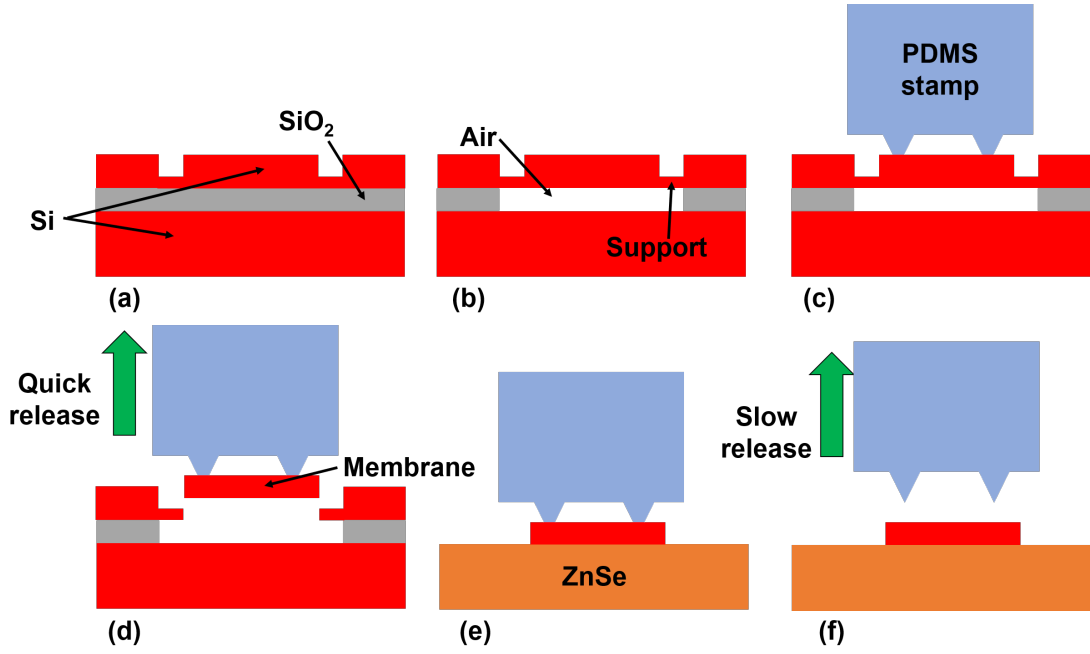
FIGURE 7.1: Outline of the transfer printing of Si membranes onto a ZnSe substrate: (a) The membranes and structures on the membrane are defined using lithography and dry etching; (b) the membrane is suspended by wet etching the SiO_2 beneath the membrane; (c) a PDMS stamp is brought into contact with the suspended membrane; (d) the stamp is retreated quickly to break the supports and pick-up the membrane; (e) the stamp is lowered to bring the membrane into contact with the substrate; (f) the stamp is retreated slowly to release the membrane onto the substrate.

the stamp is moved, meaning that it can be used to both pick-up and place the membrane.

The membranes are first defined by using separate lithography and etch steps: a full etch to create the membranes, followed by a shallower etch for the waveguides and grating couplers (the grating coupler etch depth was fixed the same as the waveguide etch depth to minimise the number of fabrication steps). The membranes are then suspended so that they can be released from the SOI wafer, before being picked up and then placed onto the ZnSe substrate. In order to encourage adhesion of the membrane to the ZnSe surface, the backside of the membrane can be wetted (it will be dried during the deposition of the membrane).

7.2   Blank membrane iterations

The membranes were transfer printed by the University of Strathclyde, using their in-house custom tool that was adapted from a dip-pen lithography tool, NanoInk Inc. NLP2000. Their process is detailed in [189], for similar but much smaller and thinner membranes. The typical size of previously printed membranes are roughly $100\,\mu m \times 100\,\mu m$, which is too small for any feasible PIC. Much larger membranes

were therefore required but increasing the size is non-trivial for the process and was expected to be challenging. $2\,\mathrm{mm} \times 2\,\mathrm{mm}$ membranes have been previously been demonstrated for transfer printing [190] so the use of larger membranes was in principle feasible, but this work used a different style of tool. Consequently, an initial run was made to determine the maximum size of the membrane that could be successfully printed. In this iteration, no waveguides or other features were patterned onto the membranes.

7.2.1 Design

A series of square membranes were designed with edge sizes ranging $100\,\mathrm{\mu m} - 1000\,\mathrm{\mu m}$. In order to ensure that the membranes were strong enough to not break under their own weight, the membranes were fabricated on an SOI wafer with a $1.5\,\mathrm{\mu m}$-thick Si layer. Each membrane was surrounded by a $20\,\mathrm{\mu m}$-wide trench that would be fully etched through the Si layer, to provide access to the underlying SiO_2. In addition, the initial designs had the entire area of the membranes patterned with an array of $1\,\mathrm{\mu m} \times 1\,\mathrm{\mu m}$ holes (see Figure7.2) that would also be fully etched to act as vias that provide access for the acid during the underetching for the area directly beneath the membrane. This avoided the need for long wet etching times which would roughen the Si layer and meant that all membranes could be suspended using the same wet etch. The via array used a $10\,\mathrm{\mu m}$ separation between the vias and the vias were inset by $20\,\mathrm{\mu m}$ from the edge of the membrane, allowing suspension to be achieved with 100 minutes of submersion in HF.

To prevent the membranes detaching from the SOI wafer upon suspension, supports extended across the trench and between membrane and wafer surface. The supports started with a width of $16\,\mathrm{\mu m}$ on the wafer side of the trench, tapering down to a width of $6\,\mathrm{\mu m}$ on the membrane. These dimensions were chosen so that the supports were more likely to break on the membrane side when they were picked up by the stamp, but also to balance with the supports having sufficient strength to hold the membranes in place. The number of supports varied depending on the size of the membranes: a support was placed alternately every $100\,\mathrm{\mu m}$ along parallel edges of the membrane, starting at $10\,\mathrm{\mu m}$ from the corner of the membrane, with additional supports added on the perpendicular edges for the membranes with dimensions larger than $500\,\mathrm{\mu m}$.

7.2.2 Membrane fabrication

An example of a fabricated membrane (with a size of $200\,\mathrm{\mu m} \times 200\,\mathrm{\mu m}$) is shown in Figure 7.2. The membranes were fabricated using *e*-beam lithography and dry etching as described in Chapter 3. Comparatively thick resist layers ($\sim 800 - 900\,\mathrm{nm}$) had to

FIGURE 7.2: Optical microscope image of a 200 µm × 200µm membrane after fabrication, after a failed suspension.

FIGURE 7.3: Optical microscope image showing a non-suspended membrane. The white dashed lines shows the edge of the area where there is underetching. [insert] Zoomed image of the underetched area around a via.

be used to prevent completely etching through the resist during the deep dry etches. Following the dry etching, the membranes were then wet etched in HF to suspend them before transfer printing.

However, initial attempts to release the membranes from the wafer using the stamp failed. Under inspection using an optical microscope, the similar colouring between the membrane and the rest of the wafer except for short areas adjacent to the trenches suggested that the underetching had failed (Figure 7.3). In Figure 7.3, the white lines highlight the edge of these regions; next to the trench, the underetching extends approximately halfway to the vias ($\sim$ 5 µm), while around the vias it extends only $\sim$ 1 µm. The lack of underetching was confirmed with inspection in an SEM (Figure 7.4), where the SiO_2 has clearly not been removed.

The change in colour around the vias indicated that small amount of underetching had occurred around the vias. In turn, this suggested that either the

FIGURE 7.4: SEM images showing the unetched SiO$_2$. (a) Angled SEM showing some underetching at the edge of the membrane. (b) Cross-sectional SEM of the membrane (after cleaving) showing no apparent underetching through the HF access holes.

holes were not sufficiently large to allow flow of the HF, or that the vias had not been fully defined in the dry etching step. The latter would be caused by the effect discussed in Chapter 3, where etched features with a small aspect ratio (i.e. a deep etch but a small area) can cause rounding at the bottom of the holes. In this case, the holes would be etched enough to cut through the Si layer, but be much smaller at the bottom of the hole than their designed area. Alternatively, since the change in colour was very close to the vias, this could mean that the dry etch did not go through the entire Si layer at any point in the via and the discolouration could be the result of another part of the fabrication process.

The simplest solution to this issue was to dry etch the membranes further and remove the remaining material from within the vias. However, the patterned resist had been removed ahead of the transfer printing so additional etching meant that, in addition to the material within the vias, the entire surface of the membranes would be etched. For membranes that included waveguides this would have been an issue, but for the blank membranes the only problem that would arise would be if the membranes were excessively thinned such that they collapsed under their own weight.

Following the additional dry etching, the membranes were submerged in HF for a second wet etch; this time, the suspension of the waveguides was successful. This was confirmed with inspection using both optical microscopy and imaging and milling in a FIB tool (see Figure 7.5). After suspension, the flexible membranes sagged in the middle under their own weight, which was more significant for the largest membranes: Figure 7.5(a) shows the sagging for the largest fabricated membranes, which were 1 mm^2. Milling through the centre of this sagging with a FIB beam

revealed that the membrane was coming into contact with the Si substrate. In addition to the weight of the membrane, this was likely caused by the drying step after the wet etching, in which the removal of the water would result in capillary forces that act to pull the bottom surface of the membrane close the surface of the substrate. While the bending of the membrane could potentially lead to issues in the transfer printing through non-uniform contact with the PDMS stamp, the strong contact between the two Si surfaces is only possible if both were very smooth and consequently it was concluded that the entirety of the SiO_2 layer was removed by the second wet etch. Furthermore, this meant that the limiting factor with respect to the vias was that their aspect ratio was too low and they were incompletely etched as a result, rather than their designed size being insufficient to allow HF flow.

FIGURE 7.5: Blank membrane with underlying SiO_2 layer removed. (a) Optical micrograph of the suspended $1\,mm^2$ blank membrane before transfer printing, with the location of the FIB cut shown. (b) Angled SEM micrograph of the FIB cut through the centre of the membrane, showing membrane in contact with the Si substrate.

7.2.3 Printing result

The now-suspended membranes were then transfer printed, starting with the largest size and progressing to decreasingly small membranes until the transfer print did not fail, thereby yielding the largest size that could be printed. Note that, for this technique, a new size of stamp is required for each membrane size for good stamp-membrane contact: the same area as the membrane but $50\,\mu m$ smaller in each dimension to facilitate the membrane deposition. Furthermore, the membranes were first printed onto a PDMS substrate to investigate the pick-up aspect of the transfer, in particular to see the effect (if any) that the sagging of the membrane and consequent non-uniform contact between the stamp and membrane would have. The use of the PDMS substrate acted as a control for the pick-up as the stamp-membrane adhesive force would be similar to the membrane-substrate force (as they were the same

material), until the stamp was retreated (recall the stamp-membrane force depends on the speed of movement of the stamp). Therefore, the membrane-substrate force would be guaranteed to be larger than the stamp-membrane force and would make the membrane easy to deposit, meaning that any issues would be a result of the pick-up only.

Figure 7.6 shows the successful transfer print of blank $1000\,\mu m \times 1000\,\mu m$ (or $1\,mm^2$) Si membranes onto the PDMS substrate. Aside from a minor breakage in one corner, the membranes were printed intact and had good contact with the substrate (the darker areas in the membrane show places with worse contact), that could be optimised by small changes in the pick-up conditions. This demonstrated two things: the comparatively large and thick membranes could be transfer printed with this technique, and that the sagging of the membranes had little effect on the pick-up.

FIGURE 7.6: Optical micrographs of multiple blank $1\,mm^2$ Si membranes transfer printed onto PDMS.

Having shown that the largest membranes could be printed, transfer printing was then attempted onto the transparent ZnSe substrate. The substrates used for this were commercially available ZnSe optical windows (Crystran), with a manufacturer-specified polished flatness of $\lambda/10$ at 633 nm and a root-mean-square surface roughness of $< 10\,nm$. The high surface quality was necessary for good contact and adhesion of the membrane to the surface, as any localised defects could prevent it from adhering. As with the PDMS substrate, transfer printing was attempted starting with the largest size and, as before, the $1\,mm^2$ membranes were successfully printed, as can be seen in Figure 7.7. Evidently, the quality of the initial print onto the ZnSe was worse, in comparison to the print onto the PDMS. From the analysis above, the issue was caused by the release and deposition step, rather than the pick-up by the stamp.

The most significant issue in the quality of the print are the large dark areas of bulging in the membrane, where the membrane has not properly adhered to the ZnSe and is bending upwards. This is caused when there is a non-uniform release from the stamp, if the membrane-surface adhesive force is insufficient to overcome the stamp-membrane force. If this occurs, the membrane is only partially released and can be pulled up by the stamp as it is moved away from the surface. The difficulty in

FIGURE 7.7: Optical micrographs of multiple blank $1\,\text{mm}^2$ Si membranes transfer printed onto ZnSe.

ensuring uniform contact between the membrane and the accepting surface is a result of the larger stamp size; the membrane must be parallel to the surface for deposition, but any small deviation in the angle of the membrane will be magnified as a larger vertical displacement for larger membranes. Likewise, larger membrane areas also means uniform wetting of the backside of the membrane is more difficult to control and, since it can affect the adhesion of the membrane to the surface, result in the bulging.

The breakages in the membranes during the transfer printing are also more noticeable for the printing onto ZnSe. One possible cause is that the stamps ($950\,\mu\text{m} \times 950\,\mu\text{m}$ for the $1\,\text{mm}^2$ membranes) were not large enough to effectively lift the membranes during pick-up, which would also explain the minor breakages for the printing onto the PDMS but does not explain the severity of the breakages for the ZnSe-printed membranes. Notably, the breakages appeared to be adjacent to the bulged areas so, similar to the bulging, the breakages could too be caused by parts of the membrane remaining attached to the stamp. Alternatively, since the membranes were taken from different places in the wafer than for the PDMS samples, it could simply be that the local mechanical properties in a particular region of the wafer made it more difficult to break the supports during pick-up and lead to the membranes breaking.

The final noticeable issue is the dirty areas under the ZnSe-printed membranes. The presence of this dirt is a by-product of the wetting process and while it would cause some scattering loss for the waveguide, it was not expected that the loss would be limiting. Furthermore, it could be mitigated by better control of the backside wetting process.

While the transfer printing of the membranes onto ZnSe was not as successful as onto PDMS, PDMS was a known substrate to the collaborators that carried out the printing. As a direct consequence, the process for the release of the membranes onto the substrate was already considerably optimal for PDMS, while ZnSe had not been used before. Furthermore, considering the discussion above, the issues of bulging, breakages and dirt were mainly results of either a non-uniform contact, sub-optimal

wetting process, or both. Therefore, it was concluded that transfer print of the $1\,\text{mm}^2$ membranes onto ZnSe should be achievable, provided that the wetting process was optimised for this substrate. The development of the membrane devices then progressed to include waveguides.

7.3 Device design

The waveguides, bends and grating couplers were simulated using MODE and FDTD to determine the optimal geometries for the membrane devices. These designs would then be used to build the layouts discussed later.

7.3.1 Waveguide geometry

The membrane waveguides were designed to operate at a wavelength of $\lambda = 7.67\,\mu\text{m}$, which is towards the upper end of the Si transparency range and could be compared to other long-wavelength Si waveguides, such as the suspended Si [49]. As the membranes had to remain intact, a waveguide slab was required with sufficient thickness to prevent the membrane from breaking, meaning that the etch depth had to be minimised.

The modal analysis of the waveguide was simulated in MODE, as described in Chapter 3. With the height of the waveguide fixed to the thickness of the top Si layer (i.e. $h = 1.5\,\mu\text{m}$), the waveguide width w and etch depth d of the waveguide were swept. The single-mode condition for the waveguide (considering TE-polarisation) is when the effective index of the TE_1 is larger than the refractive index of the cladding

FIGURE 7.8: TE_1 modal effective index in the Si on ZnSe waveguide by varying geometry for $\lambda = 7.67\,\mu\text{m}$. The white dashed line is the contour for $n_{\text{ZnSe}} = 2.4080$, below which the mode leaks into the cladding.

FIGURE 7.9: Simulated modal field intensity, $|E|$, of the TE_0 mode for the Si membrane waveguide on ZnSe.

(in this case, n_{ZnSe}) and can therefore be considered guiding. Consequently, by investigating the TE_1 mode, the single-mode waveguide geometry is determined.

Figure 7.8 shows the variance of the effective index of the TE_1 mode with the waveguide geometry at $\lambda = 7.67\,\mu m$. The fundamental TE mode, TE_0 was guided across the entire shown geometry range. The condition for single-moded propagation is given by the white line in Figure 7.8, which is equal to the contour for the refractive index of ZnSe, $n_{ZnSe} = 2.4080$. A suitable geometry therefore gives an effective index below this line, but it must also be reasonably close to it as smaller waveguide geometries will also restrict the confinement of the TE_0 mode, as well as to minimise the etch depth for the membrane strength: in this case, $w = 3\,\mu m$ and $d = 1.1\,\mu m$.

The simulated TE_0 mode for these waveguide dimensions is shown in Figure 7.9. Note the non-negligible overlap of the modal field with the underlying cladding, which illustrates the need for transparent cladding.

7.3.2   Bending loss

The bending loss for this waveguide geometry was simulated in MODE as well. This simulation, using the eigenmode solver, accounted for both the predicted propagation

FIGURE 7.10: Bending radius for Si membrane waveguide on ZnSe, with $w = 3\,\mu m$, $h = 1.5\,\mu m$ and $d = 1.1\,\mu m$.

loss in travelling through the bend and the mismatch between the modes in the straight and bent waveguides.

The results of this simulation are shown in Figure 7.10. A bending loss below $< 0.1\,\mathrm{dB}/90°$ was deemed acceptable for this design (although this is entirely device-dependent), so a bending radius of $35\,\mathrm{\mu m}$ was selected. Note that while a bend radius of $30\,\mathrm{\mu m}$ also meets this criterion, it was very close to $0.1\,\mathrm{dB}/90°$ so a slightly larger radius was used.

7.3.3   Grating coupler design

The grating couplers were simulated using a two-dimensional simulation within FDTD, with the fibre angle fixed at $10°$ from the normal to the waveguide, as shown in Figure 7.11(a). The simulation was reduced from three to two dimensions to reduce the computing time required and was achieved by shrinking the in-plane geometry perpendicular to the waveguide to a constant small size. The light was launched from the fibre into the grating and the coupling efficiency determined by the light coupled into the waveguide; by reciprocity, the reverse coupling efficiency (launching from the waveguide) will have the same coupling efficiency. The etch depth was fixed equal to the etch depth of the waveguide for ease of fabrication, with the performance optimised by varying the duty cycle (the ratio and the grating feature to its period, a/Λ) and grating period. The results of this simulation are shown in Figure 7.11(b). In this case, the best simulated coupling efficiency was $\sim 22\%$, for a 80% duty cycle and period $\Lambda = 3.2\,\mathrm{\mu m}$. The grating couplers were then defined over $10\,\mathrm{\mu m}$-wide waveguides (the approximate diameter of the fibre cores) and approximately $25\,\mathrm{\mu m}$ in length.

As described in Chapter 2, the coupling efficiency of a grating coupling is often enhanced by reflections from layers beneath the waveguide; in the case of SOI, this is the reflection from the interface of the underlying SiO_2 cladding layer and the Si wafer substrate. However, for the Si membrane waveguides that have been transfer printed
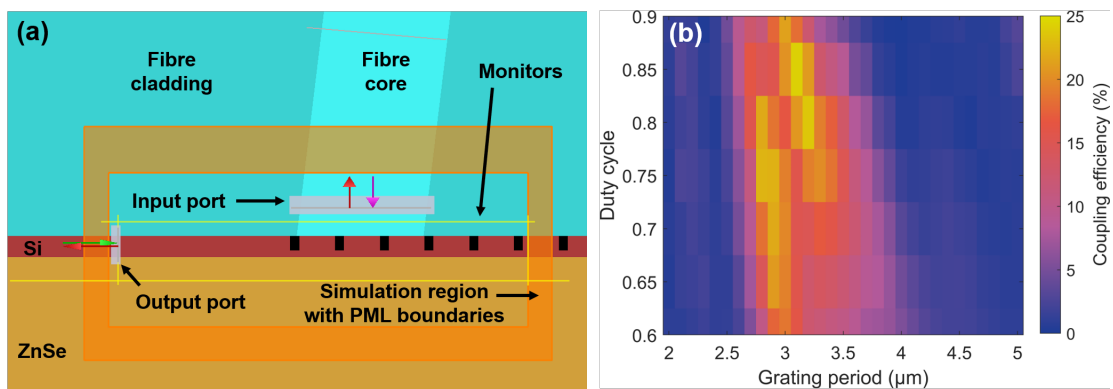
FIGURE 7.11: Simulation of a Si membrane grating coupler on ZnSe, with $d = 1.1\,\mathrm{\mu m}$ at $\lambda = 7.67\,\mathrm{\mu m}$: (a) simulation setup; (b) simulated coupling efficiency.

onto bulk ZnSe, there are no layers to cause the reflections beyond the interface of the waveguide and bulk ZnSe. Furthermore, since the refractive index of ZnSe ($\sim$ 2.4) is closer to that of Si ($\sim$ 3.4) than the index of SiO_2 ($\sim$1.4) is, the index contrast for the Si-ZnSe waveguide is smaller than for SOI, so even the initial reflection at the bottom waveguide-cladding interface will be smaller. As a consequence, the simulated coupling efficiency is lower than other grating couplers, but the equivalent coupling loss ($\sim$ 6.6 dB per coupler) is not so limiting as to prevent measurement of the waveguides and therefore the design was sufficient for this purpose.

7.4   First run with waveguides

For the first run that used membranes with fabricated waveguides, the size of the membranes was chosen to be the same as the largest membranes in Section 7.2: $1\,mm^2$. Using the largest size membranes as the basis for the layout designs meant that the available area for characterisation devices was maximised.

7.4.1   Layout design

The layout of the waveguides on the membranes are shown in Figure 7.12. Two different designs were used, for determining the propagation loss and bending loss of the waveguides respectively. In both designs, light was coupled in and out of the membranes using the grating couplers as described above. These were fabricated across 10 μm-wide waveguides to ease the fibre-grating alignment, which were then tapered down to the single-mode waveguides over a length of 100 μm.

To determine the propagation loss, the cut-back method was implemented with effectively straight waveguides with increasing lengths, as all other losses (from the

FIGURE 7.12: Waveguide layouts on $1\,mm^2$ membranes for (a) propagation loss and (b) bending loss.

bends, tapers and grating couplers) were the same in each waveguide and could therefore be removed with normalisation. The waveguides were designed in the quasi-spiral pattern shown in Figure 7.12(a) to maximise the length of each waveguide in the pattern. Each input grating coupler was separated by 50 μm vertically down the waveguide. Using the layout arranged in this way meant that 12 waveguides could fit onto the $1\,mm^2$ membrane, with effective lengths (ignoring taper, grating couplers and bends) ranging $1.35\,mm - 5.20\,mm$ in $350\,μm$ intervals.

For the bending loss layout, 7 waveguides were designed with $0 - 24\,90°$ bends, in intervals of 4 bends. Note that this does not include the bends in the shortest waveguide used from normalisation, the waveguide on the bottom left in Figure 7.12(b).

The design of the via array was changed from the array used for the blank membranes. To avoid the issues that resulted from the incomplete etching of the vias, the area of the HF vias was quadrupled from $1\,μm \times 1\,μm$ to $2\,μm \times 2\,μm$ while the separation between the centre of the vias was maintained as $10\,μm$. Furthermore, the layout of the vias was changed from a square array to a triangular array, shown in Figure 7.13. The triangular array meant that the via pattern was slightly more dense, as the rows would now be separated by $\sim 8.66\,μm$ rather than by $10\,μm$ for the square array, which would provide more coverage for HF access and therefore facilitate consistent underetching. The lateral offset between vias in different rows also meant that where the via array was overwritten by the waveguides, shown in Figure 7.13(c), it was easier to ensure there were no areas without vias.

FIGURE 7.13: Membrane via arrays: (a) square array, used for blank membranes; (b) triangular array, used for membranes with waveguides; (c) illustration of the overwriting of the via array adjacent to the waveguide, in which the overwritten region is highlighted in orange.

The overwriting of the via array, Figure 7.13(c), was necessary because vias that overlapped with waveguides would result in significant and prohibitive losses. Any via that was within $4.5\,μm$ of the waveguide was removed, either entirely or partially. Additional vias were then placed on either side of the waveguide and parallel to it with a $10\,μm$ separation (consistent with the rest of the via array), to allow the wet etch duration to be maintained. This meant that the edge-to-edge separation of the vias and the waveguide was $2.5\,μm$, which was predicted to cause negligible perturbation of the waveguide mode in the eigenmode simulations.

7.4.2   Fabrication

The membranes with waveguides were fabricated as described for the blank membranes. An example of the fabricated membrane is shown in Figure 7.14, for one of the fabricated bending loss membranes: note how the combination of the triangular via array and the overwriting of it with additional vias results in consistent coverage of the membrane area with HF vias. Figure 7.14(b) shows that the vias slightly expanded into the lateral waveguide cladding area, which is a result of the imperfect anisotropy of dry etching in practice (i.e. there will be some lateral etching) and is more pronounced for the deeper etches (which have longer etch durations) for the vias. Additionally, the flared features at the corners of the vias are artefacts of errors in the lithography, expanding the vias closer to the waveguide. Neither of these fabrication errors were expected to result in significant losses, but could contribute to some loss from scattering or perturbation of the waveguide mode.

FIGURE 7.14: SEM micrographs of the waveguides on Si membranes after suspension. (a) The shortest bending loss waveguide, used for normalisation. (b) The vias adjacent to the waveguide, showing the overwritten via array.

Following the incomplete suspension of the blank membranes, some membranes underwent destructive testing to ensure that they had been completely suspended. The membranes were cleaved to allow angled and cross-sectional analysis before being inspected in an SEM: the results of this are shown in Figure 7.15. Figure 7.15(a) shows that the underlying SiO_2 had been fully removed in the wet etching step and, as was seen for the blank membranes, the membranes sagged and came into contact with the Si wafer substrate (aside from around the supports). Interestingly, some membranes did not break during the cleave and instead were left overhanging the edge of the chip, as can be seen in Figure 7.15(b). The structural stability of this membrane suggests the sagging is not caused by the large membrane size, as the membrane remains consistently in the focus of the SEM and is thus not significantly bending under its own weight. Rather, this sagging is more likely caused by the drying process that follows the wet etching and the resultant capillary forces described earlier. In further iterations, this could be avoided by using a HF vapour etch (as opposed to the liquid HF etch used here), which has an order of magnitude
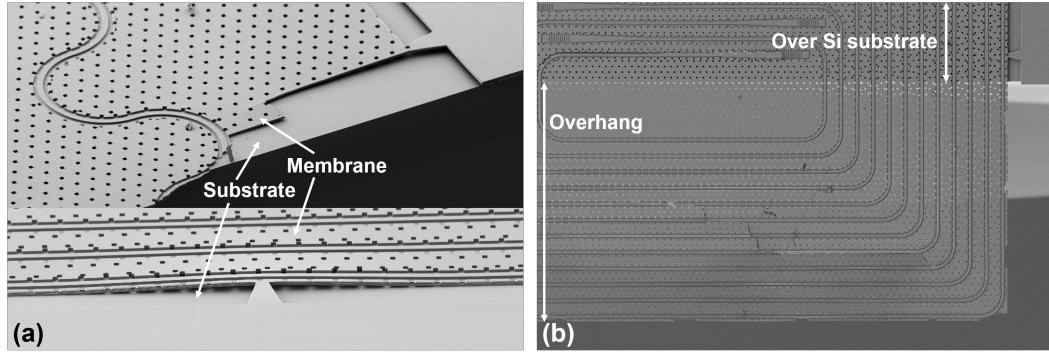
FIGURE 7.15: SEM micrographs of suspended membranes after cleaving. (a) Lifting of the membrane around the supports. (b) Overhanging membrane still attached after cleaving.

slower etch rates but is commonly used in the fabrication of microelectromechanical systems for this very reason [209, 210].

The membranes following the transfer printing are shown in Figure 7.16. The optimisation of the uniform pick-up and backside wetting are evident, as the quality of the adhesion onto the ZnSe surface was remarkably better than the result of the blank membranes printing. The membranes survived the printing intact and without any breakages, there was no visible residual dirt and, aside from very localised deformations, there was no considerable bulging as seen previously.

FIGURE 7.16:   Optical micrographs of the $1\,mm^2$ Si membranes with waveguides printed onto a ZnSe optical window, for (a) determining propagation loss and (b) determining bending loss.

Furthermore, considering these local deformations in detail, while less than ideal this out-of-plane bending would only lead to negligible losses. The largest of deformation was measured in an optical profilometer to determine its radius of curvature, as shown in Figure 7.17. The hemispherical fit of this measurement yielded a bending radius of 1.12 mm, which is orders of magnitude larger than the radius that

would be needed for significant radiative bending losses. Simulating a waveguide bent in this way (using MODE and a bent waveguide analysis similar to the method used to determine the waveguide bending loss) confirmed that this had a minimal loss contribution, predicting only an additional loss of $\sim 0.05\,\mathrm{dB/cm}$ due to this deformation.

FIGURE 7.17: Out-of-plane waveguide bending caused by membrane deformation, with a 1.12 mm radius of curvature. [inset] direction of measurement along a waveguide defect using an optical profilometer.

7.4.3   Characterisation and discussion

The membrane waveguides were characterised at $\lambda = 7.67\,\mu\mathrm{m}$ by placing the ZnSe optical window with the printed substrates, shown in Figure 7.18(a), in the experimental setup described in Chapter 3. In the propagation loss measurements described in Chapter 4, the waveguides used in the cut-back method were all straight waveguides at least 1.5 mm in length. As a consequence, the input and output grating coupler pairs were also separated by a distance of at least 1.5 mm, as they are placed at either end of the waveguides.

However, for layout of membrane waveguides shown in Figure 7.16(a), the lateral separation between the grating couplers was much smaller: 360 μm for the shortest waveguide, up to 910 μm for the longest. Furthermore, the vertical offset between the grating coupler pairs was only 25 μm and the fibres were approximately in-line with each other when placed over the gratings. This meant the input and output fibres had to be placed very close together for the correct positioning on the membrane, as can be seen in Figure 7.18(b).

As a result, light that was not coupled into the waveguides would pass through the ZnSe substrate and reflect off the metal stage beneath the sample, directed back up towards the output grating coupler and output fibre. In effect, this created a large

FIGURE 7.18: Photographs of the printed Si membranes on the ZnSe optical window. (a) A size comparison of the membranes alongside a ballpoint pen for reference. (b) Aligning the optical fibres with the membranes for characterisation.

background noise signal that was dependent on the separation between the input and output fibres. Due to the variance of the background noise with the fibre position, it was difficult to determine when aligning whether, for example, the larger signal at the shortest waveguide was due to good transmission at optimal fibre alignment or simply the fibres being placed closer together. This was further compounded by the comparatively low grating coupler efficiency (Section 7.3.3), which meant less light was coupling into the waveguides in the first place. This effect had not been considered prior to the very small membrane sizes because, as mentioned above, typically the input and output grating couplers are sufficiently spatially separated such that this effect is negligible compared to the noise floor on the detector. For the propagation loss membranes, this effect was so significant that it prevented the distinguishing of the waveguide signal when aligning, meaning that characterisation measurements of the propagation loss membrane did not yield any usable results.

For the bending loss membrane, however, the input and output grating couplers had a much larger vertical offset between them and the fibres were consequently out-of-line when placed over the grating couplers. This reduced the impact of the background noise and meant that measured signal from the detector appeared to depend on fibre-grating alignment as expected. To confirm that this was from light guided through the waveguide and not the stage reflection, in addition to measuring the signal through the waveguides to determine the propagation loss, the effect of the fibre position on the magnitude of the reflection was also determined. To do this, the fibres were aligned with the gratings in the same way as the propagation loss measurements and then the stage (with the sample) was moved so that the fibres would be over the ZnSe substrate with no membrane. The results of these measurements are shown in Figure 7.19.

Evidently from Figure 7.19, the background noise decreases with fibre position at

FIGURE 7.19: Propagation loss for Si membranes transfer printed onto ZnSe substrate. Change in detector signal when light is coupled into a waveguide and background signal when the fibres are in corresponding positions without aligning to a waveguide. Both data sets are normalised to the largest signal of each data set, and the loss was calculated as the gradient of the linear fit. The last two data points for the background are not included in the fit, as this is the thermal noise floor of the detector (this is approximated as an average of the last two data points).

a greater rate than the signal through the waveguides, indicating that there is indeed measureable waveguide transmission. This is shown by calculating an effective "loss" of $297.91 \pm 48.66\,\mathrm{dB/cm}$ for the waveguide lengths for the corresponding fibre positions; note that this is not a true measured loss, but rather it is just the change in background signal presented in a way for easy comparison to the waveguide transmission.

The measured waveguide propagation loss of $82.12 \pm 11.25\,\mathrm{dB/cm}$, while smaller than the effect of decreasing background signal, is very large compared to other waveguide losses. Firstly, it is important to note that effect of the fibre positions on the measured loss has not been removed because it was determined without the presence of the membrane (i.e. over the ZnSe surface), leading to a non-constant loss contribution for the waveguides. The background signal is likely larger than it would be with the membrane as the membrane would introduce losses for the free-space beam, such as scattering and reflections, that then will not be reflected to the output fibre. Consequently, the measured background cannot be simply subtracted from the waveguide signals as this would be overestimating its contribution and lead to an underestimate of the propagation loss.

Secondly, the propagation loss was determined using bent waveguides, rather than straight. In particular, since the number of bends in each waveguide is not

constant and increases with waveguide length, the bending losses are not normalised and therefore the fitted propagation loss includes these bending losses.

Considering other potential loss contributions, the quality of the membrane-substrate adhesion could be an important factor. While the bending loss caused by localised deformations was expected to be negligible, if the deformation is caused by a particle or contaminant on the ZnSe surface, light could scatter off this particle. Additionally, modal mismatch between waveguides in areas of good adhesion (with a ZnSe lower cladding) and areas of poor adhesion (with effective air lower cladding) could introduce other losses. To mitigate these, the quality of the membrane-substrate bonding must be improved.

Other waveguide losses could be introduced from the design itself. As noted in Section 7.4.2, fabrication errors in the lithography and etching of the vias meant that they expanded closer to the waveguide than intended, which could introduce loss. The waveguide cladding itself was thin (only 2.5 µm was etched either side of the waveguide core) so there is also possibility that the light was leaking into the membrane slab, which could cause considerable loss.

Despite the waveguide loss being considerable, the transfer printing of a waveguide device of this size is notable and therefore confirming that the waveguides allow measurable transmission at all is an achievement. To more accurately determine the propagation loss, the devices were to be redesigned to minimise the background noise and demonstrate the full use of a transparent substrate.

7.5 Second waveguide iteration

The second fabrication run with the Si membranes onto ZnSe substrates required a considerable re-design to ensure minimal losses and to avoid the large background noise of the experimental setup affecting the characterisation again. While the $1\,\mathrm{mm}^2$ membranes were shown to be effective, smaller membranes were also designed to facilitate suspension and high-quality bonding, while the waveguide geometries were reconsidered to minimise the contributions of any loss mechanisms.

To potentially reduce the sagging of the membranes observed in both the blank membrane and waveguide membrane iterations, some membranes can be underetched with the slower HF vapour. As suggested previously, this will eliminate the capillary forces associated with drying after liquid HF etches and avoid the sagging. However, since the liquid HF etches have been shown to remove the entirety of the SiO_2 layer beneath the membranes for effective release, it is important to also have another set of membranes etched in liquid HF in the event there are issues with the vapour HF.

7.5.1   Re-designed devices

When considering the effect of the device design, several design parameters were reconsidered for the waveguides, namely: the waveguide cross-section; the cladding width; and the placement of the vias adjacent to the waveguide. Factors that would not affect the propagation loss as they are factored out in the normalisation but will impact waveguide transmission were also re-considered, such as the bending losses, grating coupler efficiency, and taper lengths. To facilitate a new membrane layout, an S-bend was also designed.

To provide a reference for the losses that could be caused by factors such as coupling into the membrane slab or scattering from the via fabrication errors, suspended waveguides (that would not be transfer printed) were also designed. These waveguides would have the same dimension as the membrane waveguides, with the exception of the grating couplers which needed a different design for lower air cladding.

7.5.1.1   Waveguide cross-section

To prevent breakages to the membranes, the thickness of the waveguide slab can be increased to strengthen it. However, considering the effective index of the first-order higher mode TE_1, increasing the slab thickness can mean that it is guided (Figure 7.8 in Section 7.3.1). To avoid this, decreasing the waveguide width will squeeze the TE_1 mode out and prevent the waveguide from being multimoded. For this design, the waveguide geometry was changed to a width $w = 2.5\,\mu m$ and etch depth $d = 1.05\,\mu m$, compared to $w = 3\,\mu m$ and $d = 1.1\,\mu m$ in the previous run. Note that this design is also more fabrication tolerant to the etch depth in meeting the single-mode condition.

7.5.1.2   Cladding width

The membrane slab (the $1.5\,\mu m$-thick Si adjacent to the etched waveguide cladding) can, if it is sufficiently large, be approximated as an infinitely large waveguide. It will therefore have an infinite number of modes, one of which will have an effective index equivalent to that of the waveguide mode and consequently allow coupling between them. The strength of the coupling will exponentially depend on the thickness of the cladding, meaning that if the cladding is too thin then the waveguide will be leaky. Determining the minimal width of the cladding is therefore critical to minimising the propagation loss.

Attempting to simulate a sufficiently large number of modes to represent an infinite slab would be impractical in terms of computing resources. To simplify this, the waveguide modes of two identical waveguides will allow coupling between them

so problem can be treated as a directional coupler (Section 2.1.3.2 in Chapter 2). In this case, the waveguide separation g is to be minimised under the condition that the length over which 100% coupling occurs, L_C, is infinitely long, which is equivalent to no coupling. If the cladding width is then equal to this value of g, it can be reasonably expected that there will therefore be no coupling between the waveguide mode and a corresponding mode in the waveguide slab.

Recalling from Chapter 2 that 100% cross-coupling occurs when the directional coupler supermodes are in anti-phase (i.e. they have a π-phase difference):

$$\frac{2\pi\Delta n}{\lambda} L_C = \pi \Rightarrow L_C = \frac{\lambda}{2\Delta n} \tag{7.1}$$

where Δn is determined using the MODE eigenmode solver to find the two supermodes of the adjacent waveguides. The result of this simulation is shown in Figure 7.20.
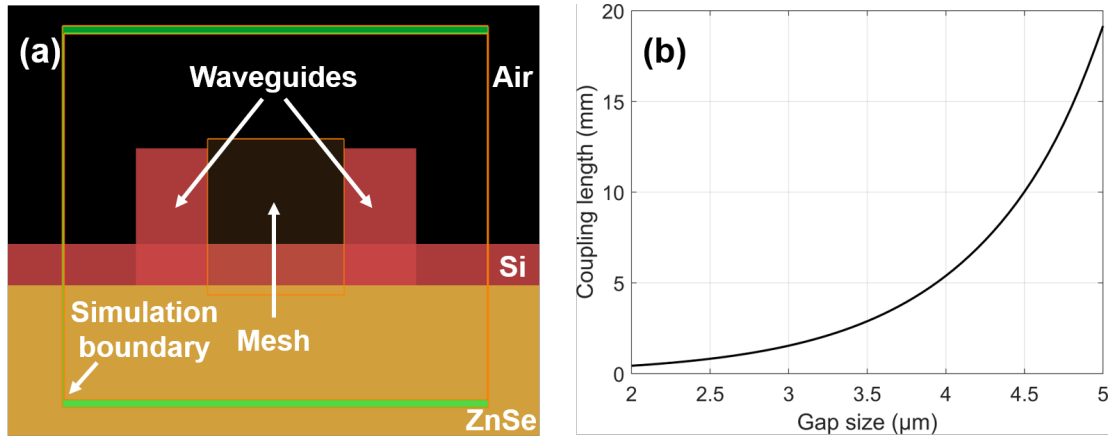
FIGURE 7.20: Simulation of a Si membrane on ZnSe waveguide to estimate slab coupling. (a) a screenshot of the simulation setup in MODE. (b) the coupling length as a function of the size of the gap between the waveguides for $\lambda = 7.67\,\mu\text{m}$.

Choosing the design condition as $L_C > 10\,\text{mm}$, which is more than triple the length of the longest waveguide in the re-designed propagation loss layout (see Section 7.5.2), $w_{\text{clad}} = 4.5\,\mu\text{m}$ is sufficiently large to ensure that there is no coupling between the waveguide and membrane slab modes.

7.5.1.3 Via placement

To reduce any effect the vias adjacent to the waveguide may have on the waveguide mode (they could, for example, act as a grating in the radiative regime to cause scattering loss), the position of the vias relative to the waveguide was varied. If the vias are close enough to perturb the fundamental waveguide mode, then the effective index of that mode will change. Therefore, by monitoring the modal effective index,

the minimal distance from the waveguide the vias must be placed to not affect the mode can be found.

The via size was kept the same as before ($2\,\mu m \times 2\,\mu m$) and the fundamental mode was monitored in MODE's eigenmode solver. The results of this simulation are shown in Figure 7.21, in which a distance of $4.75\,\mu m$ between the via edge and the edge of the waveguide results in a change in the effective index $\Delta n_{\mathrm{eff}} \sim 10^{-6}$. In order for the vias to align with the edge of the waveguide cladding as before, the cladding width is also adjusted to $w_{\mathrm{clad}} = 4.75\mu m$ as well (note that this increases the coupling length $L_C \sim 15\,\mathrm{mm}$). Equivalently, this distance also means that the separation between the centre of the vias and the centre of the waveguide is $7\,\mu m$. Note that this means that the via-to-via separation is $14\,\mu m$, meaning that the HF etch times must also be increased accordingly. The spacing across the rest of the via array was then also increased to $14\,\mu m$ to match, as this would decrease the number of vias and therefore decrease the amount of material etched from the membrane, thereby increasing the membrane strength.

FIGURE 7.21: The simulated change in the effective index of the fundamental mode of the Si membrane waveguide caused by placing $2\,\mu m \times 2\,\mu m$ vias adjacent to the waveguide.

7.5.1.4   Bending loss

The waveguide bending loss simulations were also repeated in MODE for the adjusted waveguide dimensions. However, in this case, the bends were also simulated both with and without the vias presented, to evaluate if their presence would increase the loss. The results of this simulation are shown in Figure 7.22. Evidently, compared to the previous bending loss simulations, the new waveguide design made little difference to the bending loss (see Section 7.3.2), particularly for larger radii: the bending loss is still $< 0.1\,\mathrm{dB}/90°$ for a radius of $\sim 35\,\mu m$.

The simulation was also run with the vias present to observe any effect that they would have on the bending loss. Note implementing the simulation in this way treats the via as if it is a part of the structure along the entire waveguide length, as MODE

FIGURE 7.22: Bending radius for Si membrane waveguide on ZnSe, with $w = 2.5\,\mu\text{m}$, $h = 1.5\,\mu\text{m}$ and $d = 1.05\,\mu\text{m}$.

only calculates modes for a cross-section of the waveguide. This is not physically represented in the device design, as the bends have vias spaced at $10\,\mu\text{m}$ intervals as with the rest of the waveguide. The inclusion of the vias predicts a decrease in bending loss at shorter radii (due to increased lateral confinement of the mode), which indicates that the vias are close enough to the waveguide to affect the mode in the bent waveguide for shorter radii. To avoid this, the bend radius was increased to $60\,\mu\text{m}$ where the predicted loss was very similar both with and without the vias, and therefore the vias had little effect on the mode. This also meant that the bending loss was reduced to $\sim 0.05\,\text{dB}/90°$.

7.5.1.5 Grating coupler design

The waveguide and grating couplers use the same etch depth so that only a single lithography and etch step can be used for both. As the etch depth was changed in Section 7.5.1.1, the grating coupler analysis must be repeated. The simulation setup of the grating couplers was the same as in Section 7.3.3, with the simulation source inside the fibre and the fibre angle set to $10°$. Figure 7.23(a) shows the coupling efficiency as a function of the period and duty cycle, as simulated in FDTD. For this design, the optimal design is for $\Lambda = 2.9\,\mu\text{m}$ and a duty cyle of 75%, which yields a coupling efficiency of 23.8%. Note that, while the coupling efficiency is not significantly better than for the previous design, this design has a greater fabrication tolerance in terms of the duty cycle (only the duty cycle changes due to fabrication errors, the period does not).

Additionally, for the suspended reference waveguides, a new grating coupler design is needed. An almost identical simulation setup was used, except the ZnSe substrate was removed and replaced with air. Figure 7.23(b) shows the simulated (in FDTD) coupling efficiency for a suspended Si membrane grating coupler, with a peak coupling efficiency of 48.6% for a 75% duty cycle and $\Lambda = 4.1\,\mu\text{m}$. As would be
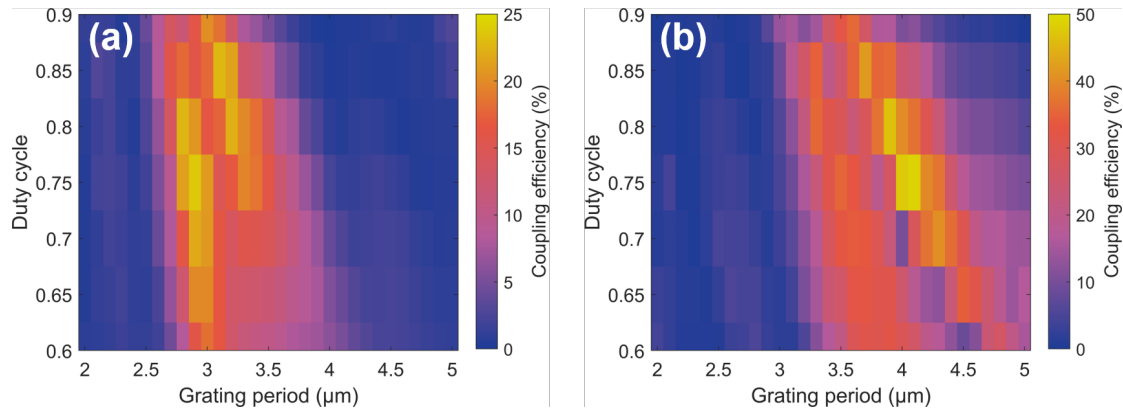
FIGURE 7.23: Simulated coupling efficiency, with $d = 1.05\,\mu\text{m}$ at $\lambda = 7.67\,\mu\text{m}$, for grating couplers on: (a) the Si membrane on ZnSe; and (b) the suspended Si membrane.

expected, the coupling efficiency is larger due to the greater refractive index contrast between the Si waveguide and the air undercladding.

7.5.1.6   Taper length

The linear tapers used in the previous iteration were $100\,\mu\text{m}$ in length due to the limitations of the space on the membrane. However, tapers use an adiabatic changing in the waveguide width to avoid losses in moving from a multimode waveguide over the grating coupler to a single-mode waveguide. To that end, if the taper is too short that change in width will be too rapid and this can cause insertion losses in the structure. Therefore, the length of the liner taper was investigated as a potential loss source to avoid in the new design.
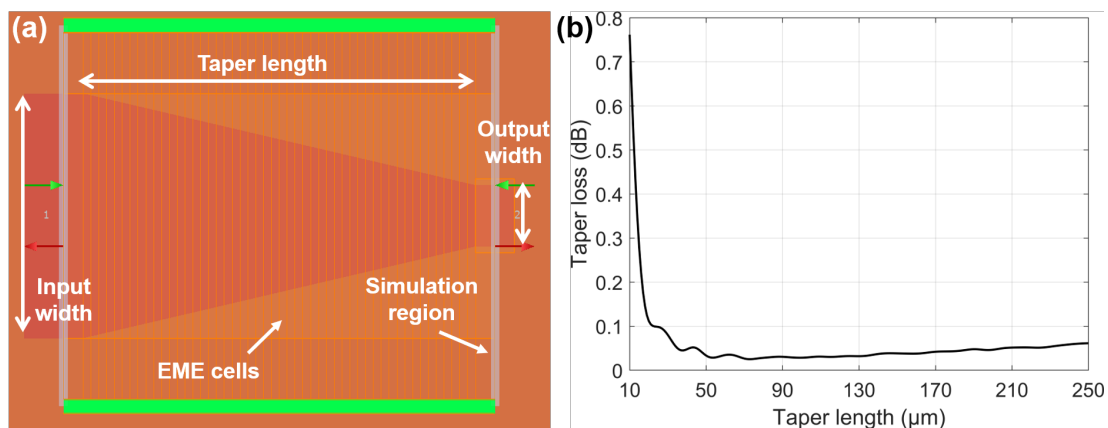
FIGURE 7.24: Simulated insertion loss for the tapering from $w = 10\,\mu\text{m}$ to $w = 2.5\,\text{m}$ for a Si membrane on ZnSe waveguide, with $h = 1.5\,\mu\text{m}$ and $d = 1.05\,\mu\text{m}$: (a) a screenshot of the simulation setup in MODE EME; (b) the simulated loss in the taper with respect to length.

Using the waveguide dimensions already outlined (10 μm over the grating coupler, down to a single-mode width of 2.5 μm), the taper was simulated in MODE's eigenmode expansion (EME) solver. For simplicity the vias were not included in the simulation, as they will be a negligible source of loss if placed the correct distance from the waveguide. The loss is simulated by calculating the scattering matrix for the taper input and output ports, based on the fundamental TE modes in the straight waveguides at either end of the taper. The result of this simulation is shown in Figure 7.24. The taper losses are at a minimum value of $\sim 0.03\,\mathrm{dB}$ for length of $\sim 75\,\mu\mathrm{m}$. At taper lengths longer than the minimum length, the taper loss increases roughly linearly due to the propagation loss in the waveguide, predominantly caused by the Si material absorption at this wavelength; the eigenmode solver predicted losses of $\sim 2\,\mathrm{dB/cm}$ for the fundamental TE mode in both the input and output straight waveguides. As with the tapers for the ESM waveguides (Section 4.2.5 in Chapter 4), the fluctuations in the loss for lengths of $25 - 50\,\mu\mathrm{m}$ are probably caused by mode beating between the radiating and fundamental modes.

7.5.1.7   S-bend

The smaller membrane design (see Section 7.5.2) was considerably larger in one dimension than it was the other, which meant that $90°$ bends could not be used because the start and end points in the bend have are separated in each of x and y by a distance equal to the radius. To make an "S" shape from $90°$ bends will result in double this separation; for the 60 μm bends, this requires a footprint of 120 μm $\times$ 120 μm to fit the bends in, start to finish. S-bends, on the other hand, are ideal for this shape because they can be used for small offsets in one direction with low loss, provided the other direction is large. As with the other waveguide components, it was simulated to minimise any loss contribution it would add.
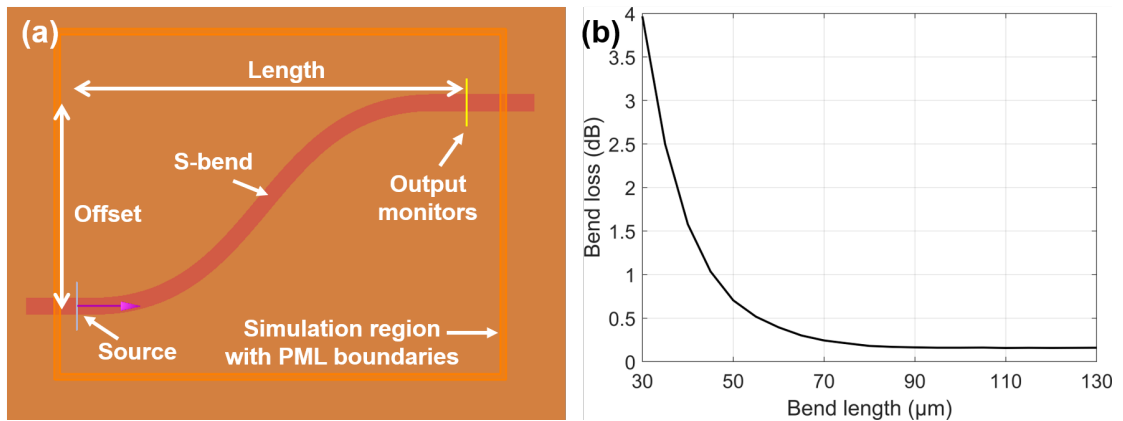
FIGURE 7.25: Simulation of a S-bend in a Si membrane on ZnSe with dimensions $w = 2.5\,\mu\mathrm{m}$, $h = 1.5\,\mu\mathrm{m}$ and $d = 1.05\,\mu\mathrm{m}$: (a) a screenshot of the simulation setup in FDTD for a 50 μm-long bend with 30 μm offset; (b) the simulated loss in the S-bend.

Figure 7.25(a) shows the simulation of the S-bend as it was set up in FDTD, in which the length of the S-bend was varied for a given offset and the transmission through (and therefore loss) was monitored. For the offset of 30 μm required for the small membrane design, Figure 7.25(b) shows that a bend length of 100 μm results in a loss of ~ 0.16 dB. For comparison to the 90° bends, this design then requires a footprint of 100 μm × 30 μm to fit the bends onto the membranes start-to-finish, which is more suitable for the membranes with differing edge lengths.

7.5.2   Re-designed layout

Two different types of membranes were designed for second iteration. The first remained as 1 mm^2 in size to maximise the length of the waveguides for the cut-back method and is pictured in Figure 7.26(a). For this membrane, the input and output grating couplers on each waveguide are separated by ~ 800 μm horizontally and ~ 715 μm vertically. This separation was kept constant for each waveguide so that the fibres would have the same relative positioning for each measurement and therefore the background reflection would remain the same, so it could be removed in normalisation. The effective lengths (not including bends, tapers and grating couplers) of the four waveguides were $1.3 - 2.2$ mm in 300 μm intervals. Only four waveguides were used (rather than 12 in the last iteration) so that the adjacent waveguides and grating couplers could have a greater separation (adjacent gratings were separated by 75 μm to ease identification of a waveguide signal.

Additionally, smaller membranes (150 μm × 1000 μm in area) were designed including two waveguides, as shown in Figure 7.26(b). These smaller membranes were used in attempt to both reduce the chance of membrane sagging and to facilitate a better quality adhesive bond between the membrane and the ZnSe surface (it is easier to ensure uniform contact and avoid any contamination with smaller

FIGURE 7.26: Redesigned membrane waveguide layouts: (a) the cut-back method for propagation loss on 1 mm^2 membranes; (b) the two-waveguide method for propagation loss on 150 μm × 1000 μm membranes.

membranes). By conventional wisdom, two data points is an insufficient number for determining the propagation loss with the cut-back method. However, repeated measurements of the same two data points (i.e. the same waveguide lengths but on many different membranes) and then using a linear fit of between the groups of points can also give an accurate value of the propagation loss. The effective lengths of each of waveguide on the membranes were 200 μm and 450 μm. Unlike the larger membrane however, the input and output grating couplers were only separated vertically by 30 μm and have different horizontal spacings (650 μm and 900 μm for the shorter and longer waveguides respectively), so the background signal cannot be normalised out if it is significant.

In addition, a second set of membranes was also designed (including both the large and small membranes). These only included the vias immediately adjacent to the waveguides so that the waveguides would only be locally suspended, with only the SiO_2 directly beneath the waveguide removed. These were used to act as a reference for the loss contributions of the vias, bends and tapers; if the propagation loss for these waveguides was low but was large for the transfer printed membrane waveguides, then the loss would be due to the transfer printing process. As mentioned previously, the only difference in design for these waveguides (beyond not being on membranes) was that they had a different grating coupler design, appropriate for air undercladding rather than ZnSe.

7.5.3   Status and future iterations

Due to the impact of unforeseen events, this area of work has not advanced to complete this iteration. The causes for this are outlined below and were predominantly caused by factors outside of the control of the author impacting critical facilities and pieces of equipment.

The COVID-19 global pandemic resulted in major restrictions on civil liberties and the majority of the world spent 2020 and significant parts of 2021 under some form of lockdown. Crucial facilities, in particular the nanofabrication cleanroom, were closed in accordance to governmental and University-level regulations and a phased re-opening only allowed for a severely restricted number of users to follow infection control rules (for example, social distancing). This directly resulted in this aspect of the work having no progress for over a year and, even when access was partially restored, reduced throughput in terms of fabrication as a consequence of maintained social distancing rules. Understandably, as this was a global pandemic, the restricted access to facilities and decreased productivity was mirrored in the experiences of collaborators who were due to carry out the transfer printing, which added delays.

In addition, these issues were compounded by setbacks related to key pieces of equipment. Tool down-times were anticipated and were built into the project planning schedules. Typically, these down-times might last a few weeks or, at worse, a

couple of months. However, the *e*-beam lithography tool spent much of late 2020 and early 2021 in a "down" state and, at the time of writing, is now in an indefinite period of down-time. As a tool being unusable for this amount of time had not been expected by either this author or the cleanroom staff, alternate provision for *e*-beam lithography was not available for several months. Efforts were made to adapt the work to use alternative lithography tools (namely, a Nikon NSR-S204B Scanner), which came with associated time costs such as the re-design of existing layouts and lead times for reticles. However, this tool also experienced periods of down-time (albeit much less severe). The cumulative impact of these delays pushed this work beyond what was allowed by its time-restricted funding.

The current iteration of the membranes must be completed to isolate the propagation loss from the other loss sources in the experimental setup. Combined with an improved adhesion of the membrane to the ZnSe surface, if the re-design of the membrane devices can yield a reasonable propagation loss then this approach can show promise for heterogeneous integration of group IV waveguides with further transparent substrates. If higher losses were instead obtained, then the propagation loss suspended comparison structures should allow determination of whether the losses come from the waveguide design or the transfer printing. If it is the former, the combined material platform may still offer promise, while the latter would suggest that the approach is not feasible for practical devices.

7.6   Summary

The transfer printing of $1\,\mathrm{mm}^2$ Si membrane PICs onto a ZnSe substrate has been demonstrated. The ZnSe substrate can offer a much wider transparency range as a cladding than other materials used in group IV, which could be ideal for waveguides with a broadband low-loss propagation. As a back-end process, transfer printing could enable the heterogeneous integration of ZnSe into group IV photonics in a way that is CMOS-compatible. However, the most recent propagation loss of these waveguides was estimated as $82.12 \pm 11.25\,\mathrm{dB/cm}$, a value that is likely an overestimate as it includes other sources of loss that have not been factored out with normalisation. To determine a propagation loss figure with a greater accuracy, re-designed membranes are discussed in detail, which should demonstrate the viability of the material platform and integration technique for use with PICs. Unfortunately, this work was halted by the COVID-19 pandemic and equipment failures, so further iterations are required.

Chapter 8

Conclusions and future work

8.1 Conclusions

Group IV materials offer great opportunity for the implementation of mid-infrared (MIR) photonic devices for a variety of applications including, in particular, the realisation of "lab-on-a-chip" devices for chemical and biological sensing applications. For devices to cover the MIR fingerprint region and to distinguish similar chemicals, both the active and passive parts of a photonic integrated circuit need to be implemented for wideband operation of the circuit. The work in this project covers passive integrated components and material platforms that can cover a wide spectral bandwidth in the MIR.

In the literature review (Chapter 2), previous demonstrations of MIR integrated photonic components, materials and platforms are detailed. From this review, it is clear that despite efforts to extend the operable wavelength, waveguides in traditional group IV material platforms (SiGe alloys notwithstanding) cannot support low-loss propagation over the entire MIR wavelength range. As a consequence, the heterogeneous materials outside of group IV should be considered. Furthermore, many device components have an inherent wavelength dependence to their performance so new design approaches must be used for a wideband circuit. From this basis the work in this project has been developed to expand their usable bandwidth. The techniques used for the simulation, fabrication and characterisation of these devices are detailed in Chapter 3.

Wideband single-moded waveguide propagation is required for predictable behaviour of complex optical circuits, which can otherwise limit the sensitivity or resolution of devices such as spectrometers. The first practical demonstration of the application of the endlessly single-mode guidance principle in photonic crystal fibres to integrated photonics is shown in Chapter 4. In this waveguide, the lateral waveguide cladding is replaced with a linear grating (in practice, a two-dimensional grating that is subwavelength in one direction). The effective index of the waveguide

cladding mode thus varies with wavelength such to prevent any higher modes being supported by the waveguide and to ensure good confinement of the fundamental mode in the waveguide core. The simulated modal analysis of the waveguides predicts that the single-mode bandwidth of these waveguides covers the wavelength range of $1.5 - 5.5 \, \mu$m, which covers the mid-infrared range accessible with low loss to SOI waveguides and some of the near-infrared range as well. These waveguides were fabricated in a SOI platform with a 500 nm-thick silicon layer, with waveguide width $w = 1.2 \, \mu$m, etch depth $d = 250$ nm and cladding width $w_{\text{clad}} = 9.6 \, \mu$m. Optical characterisation of these waveguides resulted in an experimentally determined propagation loss of ~ 1.5 dB/cm at wavelengths of $1.95 \, \mu$m and $3.80 \, \mu$m. Furthermore, single-mode propagation was confirmed by imaging the output waveguide facet over $1.95 - 2.01 \, \mu$m. Consequently, it can be concluded that the waveguides were single-moded over a wavelength range of $1.95 - 3.80 \, \mu$m, equivalent to an octave of frequency. In principle, these waveguides could also extend down to near-infrared (NIR) wavelengths (i.e. $\lambda = 1.55 \, \mu$m) also, but in this case ceased to be guiding as the cladding grating could no longer be considered subwavelength at this wavelength. A straightforward change in the cladding period should enable this.

Efficient power splitting in integrated photonics is required for any optical circuit beyond a minimal complexity and is essential for the realisation of on-chip spectrometers. MMIs can be useful as splitters as they can facilitate multi-channel power splitting and the implementation of switches. Chapter 5 details the extension of a design approach for MMIs, in which the multimode region is replaced with a subwavelength grating, from NIR wavelengths to the MIR. This grating acts like a homogeneous anisotropic material and the anisotropy of its refractive index yields a more compact device with a simulated wideband performance over the $3 - 4 \, \mu$m wavelength range. These MMIs were fabricated on an SOI platform with a 500 nm-thick silicon layer, with a full etch to the underlying SiO_2 layer, and were embedded in MZIs to facilitate characterisation alongside conventional MMIs for comparison. The fabrication process was also developed to prevent breakage of the subwavelength features during the application of thick resists as protective layers for polishing. Optical characterisation of the MZIs revealed that the subwavelength-structured MMIs had a high-perfomance bandwidth approximately double that of the conventional MMIs, despite residual silcion left in the etched features (contrary to the design). Agreeing closely with the simulation around a central characterisation wavelength of $3.4 \, \mu$m, the 0.5 dB-bandwidth of the imbalance was 628 nm, the 1 dB-bandwidth for the insertion loss was 616 nm and the $5°$-bandwidth of the phase error was 564 nm. Since this design approach is not dependent on the materials used, it could be extended to longer wavelengths in the MIR by using a different material platform than SOI.

Efficient power splitting is also the focus of Chapter 6. In this chapter, a waveguide Y-junction was implemented as power splitter. The optimal geometry of

the Y-junction was found using a particle-swarm optimisation, for a fabrication-tolerant, low-loss and broadband splitter over a wavelength range of $2.8 - 3.7\,\mu m$. The devices were fabricated in a fully-etched SOI platform with a 500 nm-thick silicon layer, with the Y-junctions forming a "chain" to allow characterisation over many splitters. The loss of these devices were experimentally measured as $\sim 0.2\,dB$ over a wavelength range of $3.15 - 3.7\,\mu m$, agreeing with the simulation results within the experimental error. The lower end of the spectrum ($\lambda = 2.8 - 3.15\,\mu m$) could not be measured due to limiting losses in the device, likely due to SiO_2 absorption peaks.

Chapter 7 showed the heterogeneous integration of Si-based membrane waveguides with a ZnSe substrate, as a proof-of-concept for using transparent substrates. This will reduce the propagation loss that would otherwise be caused by material absorption from underlying cladding layers, that limits the practically usable range for many platforms. The proposed integration method was to use high-accuracy transfer printing to place pre-fabricated Si-membranes, with all waveguides and grating couplers already defined, onto polished ZnSe optical windows; in principle, this process could be CMOS-compatible as the ZnSe is introduced at a back-end stage only. Various iterations of the membranes and printing process were discussed, starting with blank waveguides (i.e. with no waveguides) to determine the maximal membrane dimensions and progressing to membranes for characterisation of propagation and bending losses. Membrane devices, $1000\,\mu m \times 1000\,\mu m$ in size, were transfer printed onto the ZnSe substrate with high-quality adhesion and an upper limit to the propagation loss of the waveguides was determined as $82.12 \pm 11.25\,dB/cm$, with the accuracy limited by the experimental setup. To determine and minimise any propagation losses resulting from the membrane design, a complete loss analysis of the waveguide components was performed for a new iteration of the design. This next iteration would allow the evaluation of the integration technique because, as the losses from the membrane design and experimental setup were now minimised, the determined propagation loss would result solely from the propagation loss of the waveguides and the quality of the adhesion of the membrane to the ZnSe substrate. However, due to unforeseen circumstances, this work could not be feasibly completed within the constraints of the project, so this next iteration is left as future work.

MIR group IV photonics has been developed through this thesis by revisiting the design of components that are fundamental building blocks of photonic integrated circuits. While the endlessly single-mode waveguides, subwavelength-structured MMIs, and broadband Y-junctions are all demonstrated on an SOI platform and at the shorter end of the MIR spectrum, the design approaches detailed could be in principle extended to longer wavelengths and different platforms as they are not dependent on the materials. This is in contrast to other broadband MIR devices, in particular the SiGe alloy platform, where the performance relies on the dispersion of the material

platform itself. Furthermore, the initial steps taken to explore ZnSe as a lower cladding for group IV waveguides is an important one, given that the need for the integration of non-group IV materials into CMOS-compatible photonics has been highlighted previously [11].

8.2 Future work

As discussed in the previous section, the immediate focus for future research is further investigations on the heterogeneous integration of group IV waveguides with transparent substrates, to determine whether the propagation loss of the waveguides is limited by the printing technique. However, beyond this, the obvious area to address is how the different components will be combined for a PIC with low loss and single-mode operation over a significant wavelength range. In reference to overarching goal for MIR absorption spectroscopy, in which an entire bench-top laboratory setup can be implemented on a single chip, the proposed chip design for MIR absorption spectroscopy is shown again here, in Figure 8.1.

FIGURE 8.1: Repeated illustration of an integrated photonics chemical sensing device from Chapter 1. The black lines indicate waveguides, while the yellow lines are electrical connections.

This thesis considers the passive components in the on-chip spectrometer so this will be the primary focus in this section. However, it should be noted that examples of integrated lasers [12, 211–213] and detectors [214–218] at various shorter MIR wavelengths, as have modulators [14, 219] have been demonstrated at MIR wavelengths. For broadband MIR photonic sensing devices, these active components need to cover longer wavelengths or have a tunable operation, or both. Design approaches around this can reduce these requirements however, for example, by using a laser array and multiplexing different wavelength signals into a broadband waveguide to remove the need for a tunable source. In the interim, or for applications that require a disposable chip, broadband light could be coupled in and out of the chip for use with external sources and detectors using the edge couplers, that allow

low-loss transition between fibre and waveguide modes (as discussed in Chapter 2). However, these will need to be first optimised for MIR wavelengths and for a broad wavelength range, as most current designs have been for the NIR where efficient fibre-chip coupling is required for silicon photonics packaging.

Consider then the other components in the spectrometer chip. Looking first at the waveguides, the ESM waveguides demonstrated in this thesis are single-moded over a broad wavelength range but suffer from a large bending loss due to low lateral confinement over the waveguide mode. In Chapter 4, this was highlighted as an area of further research; one potential solution could be replace the ESM waveguide in the bends with multimode strip or rib waveguides. This would increase the confinement of the fundamental mode and, with careful design (i.e. adiabatic transitions between ESM and multimode waveguides, and using Euler bends), the excitement of higher order modes could be avoided.

Furthermore, currently the wavelength range of operation for the ESM waveguides are limited by the material absorption in the SOI platform, it would be useful to extend this to longer wavelengths that cover the MIR fingerprint region. It is worth nothing that the ESM waveguides have been demonstrated over an octave of frequency, which at longer wavelengths translates to a much wider coverage of wavelength (for example, $\lambda = 8 - 16 \, \mu$m is also an octave of frequency). Migrating the design to longer wavelengths would likely mean using Ge-based waveguides rather than Si, which would require a different tuned refractive index of the homogeneous effective material regions. In principle, at longer wavelengths, the fabrication in terms of the hole size is less restricted, as the changing subwavelength threshold means that the grating period and hole size can be increased. However, the height of the waveguide will need to be increased for longer wavelengths (i.e. thicker guiding layers will have to be used) and correspondingly deeper etches will have to be used, which could lead to improperly defined holes that will in turn affect the design. Effective tuning of the design for use at longer wavelengths, and to account for different materials and fabrication tolerances, is therefore highlighted as an area for further work.

The other area of focus related to waveguides concerns that section of the waveguide that is used in the analyte channel. The sensitivity of the photonics sensor will depend on the overlap of the waveguide mode with any analyte and the length of the waveguide over which this interaction occurs (see Equation 2.14 in Chapter 2). However, while the overlap with the analyte was not explicitly calculated, the modal analysis of the ESM waveguide showed that the mode was considerably confined within the Si waveguide throughout the simulated wavelength range, suggesting that the interaction length would have to very long for significant analyte-induced absorptions. Typically, rather than the layout shown in Figure 8.1, the long interaction length might be achieved using a spiral waveguide but the aforementioned large bending losses for the ESM waveguides would likely make this unfeasible. The ESM

waveguides could be used in the rest of the chip but then the question is: what waveguide design can be used across the analyte channel? Waveguides that have a large mode-analyte overlap, such as slot waveguides, are unlikely to support broadband operation to work with the ESM waveguides, as is the case with conventional single-mode strip or rib waveguides. Could conventional multimode waveguides be designed with a sufficient modal confinement in the analyte and, if so, could they be used in a way that would avoid exciting higher order modes over a long interaction length? Or alternatively, if coupling back into the ESM waveguides effectively filters out the higher order modes, can this be achieved without significant system loss? Careful design and optimisation are required for waveguides over the analyte channel.

The other fundamental components essential for creating the spectrometer chip in Figure 8.1 are the power splitters and combiners used: the MMIs and the Y-junctions. The power splitters are used within the chip for separating the input light into different channels and for making MZIs. The most immediately obvious area for future research is to adapt the same design approaches used for subwavelength-structured MMIs and optimised to work at longer wavelengths and on different material platforms. In particular, both of these devices are defined by a full etch, but a redesign that used a slab for the waveguide would allow them to be included in suspended devices or transfer printed membranes.

However, considering the existing devices, in this thesis both of these devices have been optimised for broadband operation, but in both cases the devices were designed with single-mode waveguides at the input and outputs of the splitters. This means that then the obvious question for future research is then: how are these designs implemented into a low-loss PICs with the ESM waveguides? Considering first the Y-junction, what happens if the waveguide cladding is replaced from a step-index cladding to the ESM cladding? Unlike an unoptimised Y-junction, where the sharp etched tip splits the power in the fundamental mode, the optimised splitter geometry reduces the power that overlaps with this point to increase the tolerance to fabrication. However, this also implies that the optimised geometry is not reliant on a large index contrast at the tip, which could be useful for use with the ESM cladding. Repeating the optimisation of the Y-junction structure with the ESM cladding, potentially with more freedom (i.e. allowing the length of the Y-junction to be variable rather than fixed, as was the case in this work), should yield a design compatible with the ESM waveguides.

Conversely, the subwavelength-structured MMIs will inherently not work using the same cladding as the ESM waveguides, as they rely on fully etched claddings to work. Despite this, it should actually be simpler in principle to combine the MMIs and ESM couplers. For most MMI designs, light is launched into the multimode region from tapered waveguides, which themselves generally have a cross-sectional width that would (at the output of the taper) support multimode mode propagation if it

were a full-length waveguide, rather than a short taper. In that case then, the multimode transitions highlighted for the ESM waveguide bends or for the analyte channel above could also be used here as an intermediary between the ESM waveguide and the MMI taper. Furthermore, as the adiabatic transition from ESM waveguide to multimode strip and from strip waveguide to subwavelength grating will have already been optimised by other areas of future work, then this should be straightforward to implement. For certainty, this multi-transition approach should be experimentally verified to be low-loss.

The final area for future work assumes that the heterogeneous integration of membrane waveguides with a transparent substrate through transfer printing is confirmed to be viable (i.e. with no limiting losses) with Si membranes and a ZnSe substrate. If this is the case, then the process must be immediately extended to use Ge membranes, for a platform that is fully transparent over $\lambda = 2 - 16 \, \mu$m. If this too can be achieved, and the other components detailed in this thesis are suitably adapted for the platform, then the demonstration of the ESM waveguides and power splitters on this platform is the next logical progression, as this could finally yield a PIC with a working wavelength range that spans much of the MIR, albeit only using passive devices. Then, using an external optical setup and combining with microfluidics, a sensing demonstration over the fingerprint region region to detect a chemical (preferably using biological material, if possible) should be attempted to showcase the utility of the MIR group IV photonic sensing devices.

References

[1] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, "Silicon optical modulators", *Nature Photonics* 4:518–526, 2010.

[2] C. Rogers, A. Y. Piggott, D. J. Thomson, R. F. Wiser, I. E. Opris, S. A. Fortune, A. J. Compston, A. Gondarenko, F. Meng, X. Chen, G. T. Reed, and R. Nicolaescu, "A universal 3D imaging sensor on a silicon photonics platform", *Nature* 590:256–261, 2021.

[3] M. Teng, A. Honardoost, Y. Alahmahdi, S. S. Polkoo, K. Kojima, H. Wen, C. K. Renshaw, P. LiKamWa, G. Li, S. Fathpour, R. Safian, and L. Zhuang, "Miniaturized Silicon Photonics Devices for Integrated Optical Signal Processors", *Journal of Lightwave Technology* 38(1):6–17, 2020.

[4] A. Katumba, M. Freiberger, F. Laporte, A. Lugnan, S. Sackesyn, C. Ma, J. Dambre, and P. Bienstman, "Neuromorphic Computing Based on Silicon Photonics and Reservoir Computing", *IEEE Journal of Selected Topics in Quantum Electronics* 24(6):1–10, 2018.

[5] B. Wu, Y. Yu, J. Xiong, and X. Zhang, "Silicon Integrated Interferometric Optical Gyroscope", *Scientific Reports* 8:8766, 2018.

[6] K. Wang, Z. Yuan, E. Wong, K. Alameh, H. Li, K. Sithamparanathan, and E. Skafidas, "Experimental Demonstration of Indoor Infrared Optical Wireless Communications With a Silicon Photonic Integrated Circuit", *Journal of Lightwave Technology* 37(2):619–626, 2019.

[7] G. Antonacci, J. Goyvaerts, H. Zhao, B. Baumgartner, B. Lendl, and R. Baets, "Ultra-sensitive refractive index gas sensor with functionlized silicon nitride photonic circuits", *APL Photonics* 5(8):081301, 2020.

[8] A. Rickman, "The commercialization of silicon photonics", *Nature Photonics* 8:579–582, 2014.

[9] D.Thomson, A. Zilkie, J. E. Bowers, T. Komljenovic, G. T. Reed, L. Vivien, D. Marris-Morini, E. Cassan, L. Virot, J.-M. Fédéli, J.-M. Hartmann, J. H. Schmid, D.-X. Xu, F. Boeuf, P. O'Brien, G. Z. Mashanovich, and M. Nedeljkovic, "Roadmap on silicon photonics", *Journal of Optics* 18(7) 2016.

[10] R. Soref, "Mid-infrared photonics in silicon and germanium", *Nature Photonics* 4:495–497, 2010.

[11] H. T. Lin, Z. Q. Luo, T. Gu, L. C. Kimerling, K. Wada, A. Agarwal, and J. J. Hu, "Mid-infrared integrated photonics on silicon: a perspective", *Nanophotonics* 7(2):393–420, 2018.

[12] W. Zhou, N. Bandyopadhyay, D. Wu, R. McClintock, and M. Razeghi, "Monolithically, widely tunable quantum cascade lasers based on a heterogeneous active region design", *Scientific Reports* 6:25213, 2016.

[13] M. R. Karim, H. Ahmad, S. Ghosh, and B. M. A. Rahman, "Design of dispersion-engineered As_2Se_3 channel waveguide for mid-infrared region supercontinuum generation", *Journal of Applied Physics* 123(21):213101, 2018.

[14] M. Nedeljkovic, S. Stanković, C. J. Mitchell, A. Z. Khokhar, S. A. Reynolds, D. J. Thomson, F. Y. Gardes, C. G. Littlejohns, G. T. Reed, and G. Z. Mashanovich, "Mid-Infrared Thermo-Optic Modulators in SoI", *IEEE Photonics Technology Letters* 26(13):1352–1355, 2014.

[15] R. W. Waynant, I. K. Ilev, and I. Gannot, "Mid-infrared laser applications in medicine and biology", *Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences* 359(1780):635–644, 2001.

[16] Merck KGaA, *IR Spectrum Table*, https://www.sigmaaldrich.com/GB/en/technical-documents/technical-article/analytical-chemistry/photometry-and-reflectometry/ir-spectrum-table, Accessed: 27th October 2021.

[17] P. T. Lin, V. Singh, J. Hu, K. Richardson, J. D. Musgraves, I. Luzinov, J. Hensley, L. C. Kimerling, and A. Agarwal, "Chip-scale Mid-Infrared chemical sensors using air-clad pedestal silicon waveguides", *Lab Chip* 13(11):2161–2166, 2013.

[18] C. Ranacher, C. Consani, A. Tortschanoff, R. Jannesari, M. Bergmeister, T. Grille, and B. Jakoby, "Mid-infrared absorption gas sensing using a silicon strip waveguide", *Sensors and Actuators A: Physical* 277:117–123, 2018.

[19] L. Bodiou, Y. Dumeige, S. Normani, G. Louvet, P. Němec, V. Nazabal, and J. Charrier, "Design of a Multimode Interferometer-Based Mid-Infrared Multispecies Gas Sensor", *IEEE Sensors Journal* 20:13426–13435, 22 2020.

[20] W. Liu, Y. Ma, Y. Chang, B. Dong, J. Wei, Z. Ren, and C. Lee, "Suspended silicon waveguide platform with subwavelength grating metamaterial cladding for long-wave infrared sensing applications", *Nanophotonics* 10(7) 2021.

[21] V. Mittal, G. Devitt, M. Nedeljkovic, L. G. Carpenter, H. M. H. Chong, J. S. Wilkinson, S. Mahajan, and G. Z. Mashanovich, "Ge on Si waveguide mid-infrared absorption spectroscopy of proteins and their aggregates", *Biomedical Optics Express* 11(8):4714–4722, 2020.

[22] *Cocaine*, NIST Chemistry WebBook, SRD 69,
`https://webbook.nist.gov/cgi/cbook.cgi?ID=C50362&Units=SI`, Accessed:
27th October 2021.

[23] *Acetaminophen*, NIST Chemistry WebBook, SRD 69,
`https://webbook.nist.gov/cgi/cbook.cgi?ID=C103902&Type=THZ-IR-SPEC&Index=0`, Accessed: 27th October 2021.

[24] Y.-C. Chang, P. Wägli, V. Paeder, A. Homsy, L. Hvozdara, P. van der Wal,
J. Di Francesco, N. F. de Rooij, and H. P. Herzig, "Cocaine detection by a
mid-infrared waveguide integrated with a microfluidic chip", *Lab Chip*
12:3020–3023, 2012.

[25] J. Bland-Hawthorn and S. G. Leon-Saval, "Astrophotonics: molding the flow of
light in astronomical instruments [Invited]", *Optics Express* 25(13):15549–15557,
2017.

[26] R. Diener, J. Tepper, L. Labadie, T. Pertsch, S. Nolte, and S. Minardi, "Towards
3D-photonic, multi-telescope beam combiners for mid-infrared
astrointerferometry", *Optics Express* 25(16):19262–19274, 2017.

[27] J. J. Liu, B. L. Stann, K. K. Klett, P. S. Cho, and P. M. Pellegrino, "Mid and
long-wave infrared free-space optical communication", *Laser Communication
and Propagation through the Atmosphere and Oceans VIII*, ed. by J. P. Bos,
A. M. J. van Eijk, and S. Hammel, vol. 11133, SPIE, 2019, pp. 1–18.

[28] D. H. Titterton, "Development of Infrared Countermeasure Technology and
Systems", *Mid-infrared Semiconductor Optoelectronics*, ed. by A. Krier, vol. 118,
Springer Series in Optical Sciences, Springer, 2006.

[29] I. E. Gordon, L. S. Rothman, C. Hill, R. V. Kochanov, Y. Tan, P. F. Bernath,
M. Birk, V. Boudon, A. Campargue, K. V. Chance, B. J. Drouin, J.-M. Flaud,
R. R. Gamache, J. T. Hodges, D. Jacquemart, V. I. Perevalov, A. Perrin,
K. P. Shine, M.-A. H. Smith, J. Tennyson, G. C. Toon, H. Tran, V. G. Tyuterev,
A. Barbe, A. G. Császár, V. M. Devi, T. Furtenbacher, J. J. Harrison,
J.-M. Hartmann, A. Jolly, T. J. Johnson, T. Karman, I. Kleiner, A. A. Kyuberis,
J. Loos, O. M. Lyulin, S. T. Massie, S. N. Mikhailenko, N. Moazzen-Ahmadi,
H. S. P. Müller, O. V. Naumenko, A. V. Nikitin, O. L. Polyansky, M. Rey,
M. Rotger, S. W. Sharpe, K. Sung, E. Starikova, S. A. Tashkun,
J. Vander Auwera, G. Wagner, J. Wilzewski, P. Wcisło, S. Yu, and E. J. Zak, "The
HITRAN2016 molecular spectroscopic database", *Journal of Quantitative
Spectroscopy and Radiative Transfer* 203:3–69, 2017.

[30] *HITRAN on the Web*, `https://hitran.iao.ru/`, Accessed: 27th October 2021.

[31] G. Z. Mashanovich, G. T. Reed, B. D. Timotijevic, and S. P. Chan, "Silicon
Photonic Waveguides", *Silicon Photonics: The State of the Art*, ed. by G. T. Reed,
John Wiley & Sons, 2008, chap. 2, pp. 15–46.

[32] Y. Xiong, D.-X. Xu, J. H. Schmid, P. Cheben, and W. N. Ye, "High Extinction Ratio and Broadband Silicon TE-Pass Polarizer Using Subwavelength Grating Index Engineering", *IEEE Photonics Journal* 7(5):1–7, 2015.

[33] H. Zafar, R. Flores, R. Janeiro, A. Khilo, M. S. Dahlem, and J. Viegas, "High-extinction ratio polarization splitter based on an asymmetric directional coupler and on-chip polarizers on a silicon photonics platform", *Optics Express* 28(15):22899–22907, 2020.

[34] M.-S. Rouifed, C. G. Littlejohns, G. X. Tina, Q. Haodong, T. Hu, Z. Zhang, C. Liu, G. T. Reed, and H. Wang, "Low Loss SOI Waveguides and MMIs at the MIR Wavelength of 2µm", *IEEE Photonics Technology Letters* 28(24):2827–2829, 2016.

[35] N. Hattasan, B. Kuyken, F. Leo, E. M. P. Ryckeboer, D. Vermeulen, and G. Roelkens, "High-Efficiency SOI Fiber-to-Chip Grating Couplers and Low-Loss Waveguides for the Short-Wave Infrared", *IEEE Photonics Technology Letters* 24(17):1536–1538, 2012.

[36] G. Z. Mashanovich, M. M. Milošević, M. Nedeljkovic, N. Owens, B. Xiong, E. J. Teo, and Y. Hu, "Low loss silicon waveguides for the mid-infrared", *Optics Express* 19(8):7112–7119, 2011.

[37] B. Dong, X. Guo, C. P. Ho, B. Li, H. Wang, C. Lee, X. Luo, and G. Lo, "Silicon-on-Insulator Waveguide Devices for Broadband Mid-Infrared Photonics", *IEEE Photonics Journal* 9(3):1–10, 2017.

[38] T. Hu, B. Dong, X. Luo, T.-Y. Liow, J. Song, C. Lee, and G.-Q. Lo, "Silicon photonic platforms for mid-infrared applications [Invited]", *Photonics Research* 5(5):417–430, 2017.

[39] M. Nedeljkovic, A. Z. Khokhar, Y. Hu, X. Chen, J. Soler. Penadés, S. Stankovic, H. M. H. Chong, D. J. Thomson, F. Y. Gardes, G. T. Reed, and G. Z. Mashanovich, "Silicon photonic devices and platforms for the mid-infrared", *Optical Materials Express* 3(9):1205–1214, 2013.

[40] G. Z. Mashanovich, F. Y. Gardes, D. J. Thomson, Y. Hu, K. Li, M. Nedeljkovic, J. Soler Penadés, A. Z. Khokhar, C. J. Mitchell, S. Stankovic, R. Topley, S. A. Reynolds, Y. Wang, B. Troia, V. M. N. Passaro, C. G. Littlejohns, T. Dominguez Bucio, P. R. Wilson, and G. T. Reed, "Silicon Photonic Waveguides and Devices for Near- and Mid-IR Applications", *IEEE Journal of Selected Topics in Quantum Electronics* 21(4):407–418, 2015.

[41] J. Soler Penadés, A. Z. Khokhar, M. Nedeljkovic, and G. Z. Mashanovich, "Low-Loss Mid-Infrared SOI Slot Waveguides", *IEEE Photonics Technology Letters* 27(11):1197–1199, 2015.

[42] S. Khan, J. Chiles, J. Ma, and S. Fathpour, "Silicon-on-nitride waveguides for mid- and near-infrared integrated photonics", *Applied Physics Letters* 102(12):121104, 2013.

[43] F. Li, S. D. Jackson, C. Grillet, E. Magi, D. Hudson, S. J. Madden, Y. Moghe, C. O'Brien, A. Read, S. G. Duvall, P. Atanackovic, B. J. Eggleton, and D. J. Moss, "Low propagation loss silicon-on-sapphire waveguides for the mid-infrared", *Optics Express* 19(16):15212–15220, 2011.

[44] T. Baehr-Jones, A. Spott, R. Ilic, A. Spott, B. Penkov, W. Asher, and M. Hochberg, "Silicon-on-sapphire integrated waveguides for the mid-infrared", *Optics Express* 18(12):12127–12135, 2010.

[45] P. T. Lin, V. Singh, Y. Cai, L. C. Kimerling, and A. Agarwal, "Air-clad silicon pedestal structures for broadband mid-infrared microphotonics", *Optics Letters* 38(7):1031–1033, 2013.

[46] Z. Cheng, X. Chen, C. Y. Wong, K. Xu, and H. K. Tsang, "Mid-infrared Suspended Membrane Waveguide and Ring Resonator on Silicon-on-Insulator", *IEEE Photonics Journal* 4(5):1510–1519, 2012.

[47] J. Soler Penadés, C. Alonso-Ramos, A. Z. Khokhar, M. Nedeljkovic, L. A. Boodhoo, A. Ortega-Moñux, Í. Molina-Fernández, P. Cheben, and G. Z. Mashanovich, "Suspended SOI waveguide with sub-wavelength grating cladding for mid-infrared", *Optics Letters* 39(19):5661–5664, 2014.

[48] J. Soler Penadés, A. Ortega-Moñux, M. Nedeljkovic, J. G. Wangüemert-Pérez, R. Halir, A. Z. Khokhar, C. Alonso-Ramos, Z. Qu, Í. Molina-Fernández, P. Cheben, and G. Z. Mashanovich, "Suspended silicon mid-infrared waveguide devices with subwavelength grating metamaterial cladding", *Optics Express* 24(20):22908–22916, 2016.

[49] J. Soler Penadés, A. Sánchez-Postigo, M. Nedeljkovic, A. Ortega-Moñux, J. G. Wangüemert-Pérez, Y. Xu, R. Halir, Z. Qu, A. Z. Khokhar, A. Osman, W. Cao, C. G. Littlejohns, P. Cheben, I. Molina-Fernandez, and G. Z. Mashanovich, "Suspended silicon waveguides for long-wave infrared wavelengths", *Optics Letters* 43(4):795–798, 2018.

[50] P. T. Lin, V. Singh, H.-Y. G. Lin, T. Tiwald, L. C. Kimerling, and A. M. Agarwal, "Low-Stress Silicon Nitride Platform for Mid-Infrared Broadband and Monolithically Integrated Microphotonics", *Advanced Optical Materials* 1(10):732–739, 2013.

[51] P. T. Lin, V. Singh, L. Kimerling, and A. M. Agarwal, "Planar silicon nitride mid-infrared devices", *Applied Physics Letters* 102(25):251121, 2013.

[52] M. Sinobad, A. DellaTorre, R. Armand, B Luther-Davies, P. Ma, S. Madden, A. Mitchell, D. J. Moss, J.-M. Hartmann, J.-M. Fedeli, C. Monat, and C. Grillet, "Mid-infrared supercontinuum generation in silicon-germanium all-normal dispersion waveguides", *Optics Letters* 45(18):5008–5011, 2020.

[53] M. Brun, P. Labeye, G. Grand, J.-M. Hartmann, F. Boulila, M. Carras, and S. Nicoletti, "Low loss SiGe graded index waveguides for mid-IR applications", *Optics Express* 22(1):508–518, 2014.

[54] J. M. Ramirez, Q. Liu, V. Vakarin, J. Frigerio, A. Ballabio, X. Le Roux, D. Bouville, L. Vivien, G. Isella, and D. Marris-Morini, "Graded SiGe waveguides with broadband low-loss propagation in the mid infrared", *Optics Express* 26(2):870–877, 2018.

[55] M. Montesinos-Ballester, V. Vakarin, Q. Liu, X. Le Roux, J. Frigerio, A. Ballabio, A. Barzaghi, C. Alonso-Ramos, L. Vivien, G. Isella, and D. Marris-Morini, "Ge-rich graded SiGe waveguides and interferometers from 5 to 11 μm wavelength range", *Optics Express* 28(9):12771–12779, 2020.

[56] G. Z. Mashanovich, M. Nedeljkovic, J. Soler-Penadés, Z. Qu, W. Cao, A. Osman, Y. Wu, C. J. Stirling, Y. Qi, Y. X. Cheng, L. Reid, C. G. Littlejohns, J. Kang, Z. Zhao, M. Takenaka, T. Li, Z. Zhou, F. Y. Gardes, D. J. Thomson, and G. T. Reed, "Group IV mid-infrared photonics [Invited]", *Optical Materials Express* 8(8):2276–2286, 2018.

[57] W. Li, P. Anantha, S. Y. Bao, K. H. Lee, X. Guo, T. Hu, L. Zhang, H. Wang, R. Soref, and C. S. Tan, "Germanium-on-silicon nitride waveguides for mid-infrared integrated photonics", *Applied Physics Letters* 109(24):241101, 2016.

[58] W. Li, P. Anantha, K. H. Lee, H. D. Qiu, X. Guo, S. C. K. Goh, L. Zhang, H. Wang, R. A. Soref, and C. S. Tan, "Spiral Waveguides on Germanium-on-Silicon Nitride Platform for Mid-IR Sensing Applications", *IEEE Photonics Journal* 10(3):2201107, 2018.

[59] M. Nedeljkovic, J. Soler Penadés, C. J. Mitchell, A. Z. Khokhar, S. Stankovic, T. D. Bucio, C. G. Littlejohns, F. Y. Gardes, and G. Z. Mashanovich, "Surface-Grating-Coupled Low-Loss Ge-on-Si Rib Waveguides and Multimode Interferometers", *IEEE Photonics Technology Letters* 27(10):1040–1043, 2015.

[60] Y.-C. Chang, V. Paeder, L. Hvozdara, J.-M. Hartmann, and H. P. Herzig, "Low-loss germanium strip waveguides on silicon for the mid-infrared", *Optics Letters* 37(14):2883–2885, 2012.

[61] M. Nedeljkovic, J. Soler Penadés, V. Mittal, G. S. Murugan, A. Z. Khokhar, C. Littlejohns, L. G. Carpenter, C. B. E. Gawith, J. S. Wilkinson, and G. Z. Mashanovich, "Germanium-on-silicon waveguides operating at mid-infrared wavelengths up to 8.5 μm", *Optics Express* 25(22):27431–27441, 2017.

[62] K. Gallacher, R. W. Millar, U. Griškevičiūte, L. Baldassarre, M. Sorel, M. Ortolani, and D. J. Paul, "Low loss Ge-on-Si waveguides operating in the 8-14 μm atmospheric transmission window", *Optics Express* 26(20):25667–25675, 2018.

[63] D. A. Kozak, T. H. Stievater, R. Mahon, and W. S. Rabinovich, "Germanium-on-Silicon Waveguides at Wavelengths From 6.85 to 11.25 microns", *IEEE Journal of Selected Topics in Quantum Electronics* 24(6):1–4, 2018.

[64] T.-H. Xiao, Z. Zhao, W. Zhou, C.-Y. Chang, S. Y. Set, M. Takenaka, H. K. Tsang, Z. Cheng, and K. Goda, "Mid-infrared high-Q germanium microring resonator", *Optics Letters* 43(12):2885–2888, 2018.

[65] A. Osman, M. Nedeljkovic, J. Soler Penadés, Y. Wu, Z. Qu, A. Z. Khokhar, K. Debnath, and G. Z. Mashanovich, "Suspended low-loss germanium waveguides for the longwave infrared", *Optics Letters* 43(24):5997–6000, 2018.

[66] A. Sánchez-Postigo, A. Ortega-Moñux, J. Soler Penadés, A. Osman, M. Nedeljkovic, Z. Qu, Y. Wu, Í. Molina-Fernández, P. Cheben, G. Z. Mashanovich, and J. G. Wangüemert-Pérez, "Suspended germanium waveguides with subwavelength-grating metamaterial cladding for the mid-infrared band", *Optics Express* 29(11):16867–16878, 2021.

[67] D. E. Hagan and A. P. Knights, "Mechanisms for optical loss in SOI waveguides for mid-infrared wavelengths around 2 μm", *Journal of Optics* 19(2):025801, 2016.

[68] K. K. Lee, D. R. Lim, H.-C. Luan, A. Agarwal, J. Foresi, and L. C. Kimerling, "Effect of size and roughness on light transmission in a Si/SiO2 waveguide: Experiments and model", *Applied Physics Letters* 77(11):1617–1619, 2000.

[69] E. Jaberansary, T. M. B. Masaud, M. M. Milosevic, M. Nedeljkovic, G. Z. Mashanovich, and H. M. H. Chong, "Scattering Loss Estimation Using 2-D Fourier Analysis and Modeling of Sidewall Roughness on Optical Waveguides", *IEEE Photonics Journal* 5(3):6601010–6601010, 2013.

[70] F. Grillot, L. Vivien, S. Laval, and E. Cassan, "Propagation loss in single-mode ultrasmall square silicon-on-insulator optical waveguides", *Journal of Lightwave Technology* 24(2):891–896, 2006.

[71] S. J. Fray, F. A. Johnson, J. E. Quarrington, and N. Williams, "Lattice bands in germanium", *Proceedings of the Physical Society* 85(1):153, 1965.

[72] F. A. Johnson, "Lattice Absorption Bands in Silicon", *Proceedings of the Physical Society* 73(2):265, 1959.

[73] F. Dell'Olio and V. M. N. Passaro, "Optical sensing by optimized silicon slot waveguides", *Optics Express* 15(8):4977–4993, 2007.

[74] P. T. Lin, S. W. Kwok, H.-Y. G. Lin, V. Singh, L. C. Kimerling, G. M. Whitesides, and A. Agarwal, "Mid-Infrared Spectrometer Using Opto-Nanofluidic Slot-Waveguide for Label-Free On-Chip Chemical Sensing", *Nano Letters* 14(1):231–238, 2014.

[75] D. M. Kita, J. Michon, S. G. Johnson, and J. Hu, "Are slot and sub-wavelength grating waveguides better than strip waveguides for sensing?", *Optica* 5(9):1046–1054, 2018.

[76] H. Chen, H. Fu, J. Zhou, X. Huang, T.-H. Yang, K. Fu, C. Yang, J. A. Montes, and Y. Zhao, "Study of crystalline defect induced optical scattering loss inside photonic waveguides in UV-visible spectral wavelengths using volume current method", *Optics Express* 27(12):17262–17273, 2019.

[77] Y. Yamamoto, P. Zaumseil, T. Arguirov, M. Kittler, and B. Tillack, "Low threading dislocation density Ge deposited on Si (1 0 0) using RPCVD", *Solid-State Electronics* 60(1):2–6, 2011.

[78] Y. Wang, B. Wang, D. F. S. Eow, J. Michel, K. E. K. Lee, S. F. Yoon, E. A. Fitzgerald, C. S. Tan, and K. H. Lee, "Performance of AlGaInP and LEDs on silicon substrates through low threading dislocation density (TDD) germanium buffer layer", *Semiconductor Science and Technology* 33(10):104004, 2018.

[79] M. Nedeljkovic, R. Soref, and G. Z. Mashanovich, "Free-Carrier Electrorefraction and Electroabsorption Modulation Predictions for Silicon Over the $1 - 14\,\mu m$ Infrared Wavelength Range", *IEEE Photonics Journal* 3(6):1171–1180, 2011.

[80] M. Nedeljkovic, R. Soref, and G. Z. Mashanovich, "Predictions of Free-Carrier Electroabsorption and Electrorefraction in Germanium", *IEEE Photonics Journal* 7(3):2600214, 2015.

[81] T. Baehr-Jones, R. Ding, Y. Liu, A. Ayazi, T. Pinguet, N. C. Harris, M. Streshinsky, P. Lee, Y. Zhang, A. E.-J. Lim, T.-Y. Liow, S. H.-G. Teo, G.-Q. Lo, and M. Hochberg, "Ultralow drive voltage silicon traveling-wave modulator", *Optics Express* 20(11):12014–12020, 2012.

[82] Y. Zhang, S. Yang, A. E.-J. Lim, G.-Q. Lo, C. G., T. Baehr-Jones, and M. Hochberg, "A compact and low loss Y-junction for submicron silicon waveguide", *Optics Express* 21(1):1310–1316, 2013.

[83] B. Zhang, W. Chen, P. Wang, S. Dai, H. Li, H. Lu, J. Ding, J. Li, Y. Li, Q. Fu, T. Dai, Y. Wang, and J. Yang, "Particle swarm optimized polarization beam splitter using metasurface-assisted silicon nitride Y-junction for mid-infrared wavelengths", *Optics Communications* 451:186–191, 2019.

[84] J. Xing, Z. Li, Y. Yu, and J. Yu, "Design of polarization-independent adiabatic splitters fabricated on silicon-on-insulator substrates", *Optics Express* 21(22):26729–26734, 2013.

[85] Y. Wang, S. Gao, K. Wang, and E. Skafidas, "Ultra-broadband and low-loss 3 dB optical power splitter based on adiabatic tapered silicon waveguides", *Optics Letters* 41(9):2053–2056, 2016.

[86] D. Mao, Y. Wang, E. El-Fiky, L. Xu, A. Kumar, M. Jaques, A. Samani, O. Carpentier, S. Bernal, M. S. Alam, J. Zhang, and M. Zhu, "Adiabatic Coupler With Design-Intended Splitting Ratio", *Journal of Lightwave Technology* 37(24):6147–6155, 2019.

[87] J. X. B. Sia, W. Wang, X. Guo, J. Zhou, Z. Zhang, M. S. Rouifed, X. Li, Z. L. Qiao, C. Y. Liu, C. Littlejohns, G. T. Reed, and H. Wang, "Mid-Infrared, Ultra-Broadband, Low-Loss, Compact Arbitrary Power Splitter Based on Adiabatic Mode Evolution", *IEEE Photonics Journal* 11(2):1–11, 2019.

[88] L. Chrostowski and M. Hochberg, "Fundamental building blocks", *Silicon Photonics Design: From Devices to Systems*, Cambridge University Press, 2015, pp. 92–110.

[89] B. Dong, X. Guo, C. P. Ho, H. Wang, C. Lee, X. Luo, and G.-Q. Lo, "Silicon-on-Insulator Waveguide Devices for Broadband Mid-Infrared Photonics", *IEEE Photonics Journal* 9(3):1–10, 2017.

[90] B. Dong, X. Luo, T. Hu, T. X. Guo, H. Wang, D.-L. Kwong, P. G.-Q. Lo, and C. Lee, "Compact Low Loss Mid-Infrared Wavelength-Flattened Directional Coupler (WFDC) for Arbitrary Power Splitting Ratio Enabled by Rib Waveguide Dispersion Engineering", *IEEE Journal of Selected Topics in Quantum Electronics* 24(4):1–8, 2018.

[91] B. Dong, T. Hu, X. Luo, Y. Chang, X. Guo, H. Wang, D.-L. Kwong, G.-Q. Lo, and C. Lee, "Wavelength-Flattened Directional Coupler Based Mid-Infrared Chemical Sensor Using Bragg Wavelength in Subwavelength Grating Structure", *Nanomaterials* 8(11):893, 2018.

[92] L.B. Soldano and E.C.M. Pennings, "Optical multi-mode interference devices based on self-imaging: principles and applications", *Journal of Lightwave Technology* 13(4):615–627, 1995.

[93] D. J. Thomson, Y. Hu, G. T. Reed, and J.-M. Fedeli, "Low Loss MMI Couplers for High Performance MZI Modulators", *IEEE Photonics Technology Letters* 22(20):1485–1487, 2010.

[94] Y. Hu, T. Li, D. J. Thomson, X. Chen, J. Soler Penades, A. Z. Khokhar, C. J. Mitchell, G. T. Reed, and G. Z. Mashanovich, "Mid-infrared wavelength division (de)multiplexer using an interleaved angle multimode interferometer on the silicon-on-insulator platform", *Optics Letters* 39(6):1406–1409, 2014.

[95] A. Sánchez-Postigo, J. G. Wangüemert-Pérez, J. Soler Penadés,
      A. Ortega-Moñux, M. Nedeljkovic, R. Halir, F. El Mokhtari Mimun,
      Y. Xu Cheng, Z. Qu, A. Z. Khokhar, A. Osman, W. Cao, C. G. Littlejohns,
      P. Cheben, G. Z. Mashanovich, and Í. Molina-Fernández, "Mid-infrared
      suspended waveguide platform and building blocks", *IET Optoelectronics*
      13(2):55–61, 2019.

[96] V. Vakarin, J. M. Ramírez, J. Frigerio, A. Ballabio, X. Le Roux, Q. Liu,
      D. Bouville, L. Vivien, G. Isella, and D. Marris-Morini, "Ultra-wideband
      Ge-rich silicon germanium integrated Mach-Zehnder interferometer for
      mid-infrared spectroscopy", *Optics Letters* 42(17):3482–3485, 2017.

[97] V. Vakarin, J. M. Ramírez, J. Frigerio, Q. Liu, A. Ballabio, X. Le Roux,
      C. Alonso-Ramos, G. Isella, P. Cheben, W. N. Ye, L. Vivien, and
      D. Marris-Morini, "Wideband Ge-Rich SiGe Polarization-Insensitive
      Waveguides for Mid-Infrared Free-Space Communications", *Applied Sciences*
      8(7) 2018.

[98] R. Zheng, D. Gao, and J. Dong, "Ultra-Compact Broadband Tunable Graphene
      Plasmonic Multimode Interferometer", *IEEE Photonics Technology Letters*
      28(6):645–648, 2016.

[99] C.-C. Huang and T.-C. Sun, "Numerical simulations of tunable ultrashort
      power splitters based on slotted multimode interference couplers", *Scientific
      Reports* 9:12756, 2019.

[100] A. Ortega-Moñux, C. Alonso-Ramos, A. Maese-Novo, R. Halir,
      L. Zavargo-Peche, D. Pérez-Galacho, I. Molina-Fernández,
      J. G. Wangüemert-Pérez, P. Cheben, J. H. Schmid, J. Lapointe, D. Xu, and
      S. Janz, "An ultra-compact multimode interference coupler with a
      subwavelength grating slot", *Laser & Photonics Reviews* 7(2):L12–L15, 2013.

[101] R. Halir, P. Cheben, J. M. Luque-González, J. D. Sarmiento-Merenguel,
      J. H. Schmid, G. Wangüemert-Pérez, D.-X. Xu, S. Wang, A. Ortega-Moñux, and
      Í. Molina-Fernández, "Ultra-broadband nanophotonic beamsplitter using an
      anisotropic sub-wavelength metamaterial", *Laser & Photonics Reviews*
      10:1039–1046, 6 2016.

[102] M. Nedeljkovic, A. V. Velasco, A. Z. Khokhar, A. Delâge, P. Cheben, and
      G. Z. Mashanovich, "Mid-Infrared Silicon-on-Insulator Fourier-Transform
      Spectrometer Chip", *IEEE Photonics Technology Letters* 28(4):528–531, 2016.

[103] D. M. Kita, B. Miranda, D. Favela, D. Bono, J. Michon, H. Lin, T. Gu, and J. Hu,
      "High-performance and scalable on-chip digital Fourier transform
      spectroscopy", *Nature Communications* 9(4405) 2018.

[104] Q. Liu, J. M. Ramirez, V. Vakarin, X. Le Roux, C. Alonso-Ramos, J. Frigerio, A. Ballabio, E. Talamas Simola, D. Bouville, L. Vivien, G. Isella, and D. Marris-Morini, "Integrated broadband dual-polarization Ge-rich SiGe mid-infrared Fourier-transform spectrometer", *Optics Letters* 43(20):5021–5024, 2018.

[105] K. Misiakos, I. Raptis, A. Salapatas, E. Makarona, A. Botsialas, M. Hoekman, R. Stoffer, and G. Jobst, "Broad-band Mach-Zehnder interferometers as high performance refractive index sensors: Theory and monolithic implementation", *Optics Express* 22(8):8856–8870, 2014.

[106] D. Yuan, Y. Dong, Y. Liu, and T. Li, "Mach-Zehnder Interferometer Biochemical Sensor Based on Silicon-on-Insulator Rib Waveguide with Large Cross Section", *Sensors* 15(9):21500–21517, 2015.

[107] M. A. Van Camp, S. Assefa, D. M. Gill, T. Barwicz, S. M. Shank, P. M. Rice, T. Topuria, and W. M. J. Green, "Demonstration of electrooptic modulation at 2165nm using a Mach-Zehnder interferometer", *Optics Express* 20(27):28009–28016, 2012.

[108] M. Nedeljkovic, S. Stanković, C. J. Mitchell, A. Z. Khokhar, S. A. Reynolds, D. J. Thomson, F. Y. Gardes, C. G. Littlejohns, G. T. Reed, and G. Z. Mashanovich, "Mid-Infrared Thermo-Optic Modulators in SoI", *IEEE Photonics Technology Letters* 26(13):1352–1355, 2014.

[109] T. Li, M. Nedeljkovic, N. Hattasan, W. Cao, Z. Qu, C. G. Littlejohns, J. Soler Penades, L. Mastronardi, V. Mittal, D. Benedikovic, D. J. Thomson, F. Y. Gardes, H. Wu, Z. Zhou, and G. Z. Mashanovich, "Ge-on-Si modulators operating at mid-infrared wavelengths up to 8 μm", *Photonics Research* 7(8):828–836, 2019.

[110] Y. Zou, S. Chakravarty, C.-J. Chung, and R. T. Chen, "Miniature mid-infrared thermooptic switch with photonic crystal waveguide based silicon-on-sapphire Mach–Zehnder interferometers", *Optical Interconnects XVI*, ed. by H. Schröder and R. T. Chen, vol. 9753, SPIE, 2016, pp. 112–117.

[111] S. Wang and D. Dai, "Polarization-insensitive 2×2 thermo-optic Mach-Zehnder switch on silicon", *Optics Letters* 43(11):2531–2534, 2018.

[112] M. A. Tran, T. Komljenovic, J. C. Hulme, M. L. Davenport, and J. E. Bowers, "A Robust Method for Characterization of Optical Waveguides and Couplers", *IEEE Photonics Technology Letters* 28(14):1517–1520, 2016.

[113] G. T. Reed and C. E. J. Png, "Silicon optical modulators", *Materials Today* 8(1):40–50, 2005.

[114]   R. Halir, P. J. Bock, P. Cheben, A. Ortega-Moñux, C. Alonso-Ramos,
        J. H. Schmid, J. Lapointe, D.-X. Xu, J. G. Wangüemert-Pérez,
        Í. Molina-Fernández, and S. Janz, "Waveguide sub-wavlength structures: a
        review of principles and applications", *Laser & Photonics Reviews* 9(1):25–49,
        2015.

[115]   P. Cheben, R. Halir, J. H. Schmid, H. A. Atwater, and D. R. Smith,
        "Subwavelength integrated photonics", *Nature* 560:565–572, 2018.

[116]   J. G. Wangüemert-Pérez, P. Cheben, A. Ortega-Moñux, C. Alonso-Ramos,
        D. Pérez-Galacho, R. Halir, I. Molina-Fernández, D.-X. Xu, and J. H. Schmid,
        "Evanescent field waveguide sensing with subwavelength grating structures in
        silicon-on-insulator", *Optics Letters* 39(15):4442–4445, 2014.

[117]   J. D. Sarmiento-Merenguel, A. Ortega-Moñux, J.-M. Fédéli,
        J. G. Wangüemert-Pérez, C. Alonso-Ramos, E. Durán-Valdeiglesias, P. Cheben,
        Í. Molina-Fernández, and R. Halir, "Controlling leakage losses in
        subwavelength grating silicon metamaterial waveguides", *Optics Letters*
        41(15):3443–3446, 2016.

[118]   R. Fernández de Cabo, D. González-Andrade, P. Cheben, and A. V. Velasco,
        "High-Performance On-Chip Silicon Beamsplitter Based on Subwavelength
        Metamaterials for Enhanced Fabrication Tolerance", *Nanomaterials* 11(5):1304,
        2021.

[119]   T. Hu, M. S. Rouifed, H. Qiu, X. Guo, C. G. Littlejohns, C. Liu, and H. Wang, "A
        Polarization Splitter and Rotator Based on a Partially Etched Grating-Assisted
        Coupler", *IEEE Photonics Technology Letters* 28(8):911–914, 2016.

[120]   A. Herrero-Bermello, J. M. Luque-González, A. V. Velasco, A. Ortega-Moñux,
        P. Cheben, and R. Halir, "Design of a Broadband Polarization Splitter Based on
        Anisotropy-Engineered Tilted Subwavelength Gratings", *IEEE Photonics
        Journal* 11(3):1–8, 2019.

[121]   P. Cheben, J. H. Schmid, S. Wang, D.-X. Xu, M. Vachon, S. Janz, J. Lapointe,
        Y. Painchaud, and M.-J. Picard, "Broadband polarization independent
        nanophotonic coupler for silicon waveguides with ultra-high efficiency",
        Optics Express 23(17):22553–22563, 2015.

[122]   R. Halir, P. Cheben, S. Janz, D.-X. Xu, Í. Molina-Fernández, and
        J. G. Wangüemert-Pérez, "Waveguide grating coupler with subwavelength
        microstructures", *Optics Letters* 34(9) 2009.

[123]   Z. Cheng, X. Chen, C. Y. Wong, K. Xu, C. K. Y. Fung, Y. M. Chen, and
        H. K. Tsang, "Focusing subwavelength grating coupler for mid-infrared
        suspended membrane waveguide", *Optics Letters* 37(7):1217–1219, 2012.

[124] D. Benedikovic, P. Cheben, J. H. Schmid, D.-X. Xu, B. Lamontagne, S. Wang, J. Lapointe, R. Halir, A. Ortega-Moñux, S. Janz, and M. Dado, "Subwavelength index engineered surface grating coupler with sub-decibel efficiency for 220-nm silicon-on-insulator waveguides", *Optics Express* 23(17):22628–22635, 2015.

[125] J. Kang, Z. Cheng, W. Zhou, T.-H. Xiao, K.-L. Gopalakrisna, M. Takenaka, H. K. Tsang, and K. Goda, "Focusing subwavelength grating coupler for mid-infrared suspended membrane germanium waveguides", *Optics Letters* 42(11):2094–2097, 2017.

[126] N.Chen, B. Dong, X. Luo, H. Wang, N. Singh, G.-Q. Lo, and C. Lee, "Efficient and broadband subwavelength grating coupler for 3.7 µm mid-infrared silicon photonics integration", *Optics Express* 26(20):26242–26256, 2018.

[127] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light*, 2nd, Princeton University Press, 2008.

[128] S. M. Rytov, "Electromagnetic Properties of a Finely Stratified Medium", *Soviet Physics JETP* 2(3):466–475, 1956.

[129] J. M. Luque-González, A. Herrero-Bermello, A. Ortega-Moñux, Í. Molina-Fernández, A. V. Velasco, P. Cheben, J. H. Schmid, S. Wang, and R. Halir, "Tilted subwavelength gratings: controlling anisotropy in metamaterial nanophotonic waveguides", *Optics Letters* 43(19):4691–4694, 2018.

[130] R. Marchetti, C. Lacava, L. Carroll, K. Gradkowski, and P. Minzioni, "Coupling strategies for silicon photonics integrated chips [Invited]", *Photonics Research* 7(2):201–239, 2019.

[131] K. Kasaya, O. Mitomi, M. Naganuma, Y. Kondo, and Y. Noguchi, "A simple laterally tapered waveguide for low-loss coupling to single-mode fibers", *IEEE Photonics Technology Letters* 5(3):345–347, 1993.

[132] K. K. Lee, D. R. Lim, D. Pan, C. Hoepfner, W.-Y. Oh, K. Wada, L. C. Kimerling, K. P. Yap, and M. T. Doan, "Mode transformer for minituarized optical circuits", *Optics Letters* 30(5):498–500, 2005.

[133] G. Roelkens, P. Dumon, W. Bogaerts, D. Van Thourhout, and R. Baets, "Efficient silicon-on-insulator fiber coupler fabricated using 248-nm-deep UV lithography", *IEEE Photonics Technology Letters* 17(12):2613–2615, 2005.

[134] L. Vivien, S. Laval, E. Cassan, X. Le Roux, and D. Pascal, "2-D taper for low-loss coupling between polarization-insensitive microwaveguides and single-mode optical fibers", *Journal of Lightwave Technology* 21(10):2429–2433, 2003.

[135] Q. Fang, J. Song, X. Luo, X. Tu, L. Jia, M. Yu, and G. Lo, "Low Loss Fiber-to-Waveguide Converter With a 3-D Functional Taper for Silicon Photonics", *IEEE Photonics Technology Letters* 28(22):2533–2536, 2016.

[136] F. Van Laere, G. Roelkens, M. Ayre, J. Schrauwen, D. Taillaert, D. Van Thourhout, T. F. Krauss, and R. Baets, "Compact and Highly Efficient Grating Couplers Between Optical Fiber and Nanophotonic Waveguides", *Journal of Lightwave Technology* 25(1):151–156, 2007.

[137] R. Topley, G. Martinez-Jimenez, L. O'Faolain, N. Healy, S. Mailis, D. J. Thomson, F. Y. Gardes, A. C. Peacock, D. N. R. Payne, G. Z. Mashanovich, and G. T. Reed, "Locally Erasable Couplers for Optical Device Testing in Silicon on Insulator", *Journal of Lightwave Technology* 32(12):2248–2253, 2014.

[138] Z. Cheng, X. Chen, C. Y. Wong, K. Xu, C. K. Y. Fung, Y. M. Chen, and H. K. Tsang, "Mid-Infrared Grating Couplers for Silicon-on-Sapphire Waveguides", *IEEE Photonics Journal* 4(1):104–113, 2012.

[139] C. Alonso-Ramos, M. Nedeljkovic, D. Benedikovic, J. Soler Penadés, C. G. Littlejohns, A. Z. Khokar, D. Pérez-Galacho, L. Vivien, P. Cheben, and G. Z. Mashanovich, "Germanium-on-silicon mid-infrared grating couplers with low-reflectivity inverse taper excitation", *Optics Letters* 41(18):4324–4327, 2016.

[140] S. Radosavljevic, B. Kuyken, and G. Roelkens, "Efficient 5.2 µm wavelength fiber-to-chip grating couplers for the Ge-on-Si and Ge-on-SOI mid-infrared waveguide platform", *Optics Express* 25(16):19034–19042, 2017.

[141] M. A. Green, "Self-consistent optical parameters of intrinsic silicon at 300 K including temperature coefficients", *Solar Energy Materials and Solar Cells* 92(11):1305–1310, 2008.

[142] S. Webster, Y. Chen, G. Turri, A. Bennett, B. Wickham, and M. Bass, "Intrinsic and extrinsic absorption of chemical vapor deposition single-crystal diamond from the middle ultraviolet to the far infrared", *Journal of the Optical Society of America B* 32(3):479–484, 2015.

[143] D. Chandler-Horowitz and P. M. Amirtharaj, "High-accuracy, midinfrared $(450\,\mathrm{cm}^{-1} \leq \omega \leq 4000\,\mathrm{cm}^{-1})$ refractive index values of silicon", *Journal of Applied Physics* 97(12):123526, 2005.

[144] J. H. Burnett, E. C. Benck, S. G. Kaplan, E. Stover, and A. Phenis, "Index of refraction of germanium", *Applied Optics* 59(13):3985–3991, 2020.

[145] K. Luke, Y. Okawachi, M. R. E. Lamont, A. L. Gaeta, and M. Lipson, "Broadband mid-infrared frequency comb generation in a Si_3N_4 microresonator", *Optics Letters* 40(21):4823–4826, 2015.

[146] R. Kitamure, L. Pilon, and M. Jonasz, "Optical constants of silica glass from extreme ultraviolet to far infrared at near room temperature", *Applied Optics* 46(33):8118–8133, 2007.

[147] J. Kischkat, S. Peters, B. Gruska, M. Semtsiv, M. Chashnikova, M. Klinkmüller, O. Fedosenko, S. Machulik, A. Aleksandrova, G. Monastyrskyi, Y. Flores, and W. T. Masselink, "Mid-infrared optical properties of thin films of aluminum oxide, titanum dioxide, silicon dioxide, aluminum nitride, and silicon nitride", *Applied Optics* 51(28):6789–6798, 2012.

[148] M. Leidinger, S. Fieberg, N. Waasem, F. ühnemann, K. Buse, and I. Breunig, "Comparative study on three highly sensitive absorption measurement techniques characterizing lithium niobate over its entire transparent spectral range", *Optics Express* 23(17):21690–21705, 2015.

[149] E. Franke, C. L. Trimble, M. J. DeVries, J. A. Woollam, M. Schubert, and F. Frost, "Dielectric function of amorphous tantalum oxide from the far infrared to the deep ultraviolet spectral region measured by spectroscopic ellipsometry", *Journal of Applied Physics* 88(9):5166–5174, 2000.

[150] E. D. Palik, "Gallium Arsenide (GaAs)", *Handbook of Optical Constants of Solids*, ed. by E. D. Palik, vol. I, Academic Press, 1997, pp. 429–443.

[151] R. T. Holm, "Indium Antimonide (InSb)", *Handbook of Optical Constants of Solids*, ed. by E. D. Palik, vol. I, Academic Press, 1997, pp. 491–502.

[152] H. Liu, S. Li, X. Yang, D. Liu, Y. Ji, F. Zhang, and D. Chen, "Physical model of the infrared dielectric function of zinc sulfide", *Journal of Applied Physics* 122(22):223109, 2017.

[153] A. Lemière, F. Désévédavy, P. Mathey, P. Froidevaux, G. Gadret, J.-C. Jules, C. Aquilina, B. Kibler, P. Béjot, F. Billard, O. Faucher, and F. Smektala, "Mid-infrared supercontinuum generation from 2 to 14 µm in arsenic- and antimony-free chalcogenide glass fibers", *Journal of the Optical Society of America B* 36(2):A183–A192, 2019.

[154] H. Qi, X. Zhang, M. Jiang, and D. Li, "Optical Constants of Zinc Selenide in Visible and Infrared Spectral Ranges", *Journal of Applied Spectroscopy* 84(4):679–682, 2017.

[155] E. D. Palik, "Cadmium Telluride (CdTe)", *Handbook of Optical Constants of Solids*, ed. by E. D. Palik, vol. I, Academic Press, 1997, pp. 409–427.

[156] D. F. Bezuidenhout, "Calcium Fluoride (CaF$_2$)", *Handbook of Optical Constants of Solids*, ed. by E. D. Palik, vol. II, Academic Press, 1997, pp. 815–835.

[157] M. E. Thomas and W. J. Tropf, "Barium Fluoride (BaF$_2$)", *Handbook of Optical Constants of Solids*, ed. by E. D. Palik, vol. III, Academic Press, 1997, pp. 683–699.

[158] M. Tazawa and A. Hadni, "Silver Chloride (AgCl), Silver Bromide (AgBr), Silver Iodide (AgI)", *Handbook of Optical Constants of Solids*, ed. by E. D. Palik, vol. III, Academic Press, 1997, pp. 553–572.

[159] T. Lewi and A. Katzir, "Silver halide single-mode strip waveguides for the mid-infrared", *Optics Letters* 37(13):2733–2735, 2012.

[160] A. P. H. Trivelli and S. E. Sheppard, "On the Visible Decomposition of Silver Halide Grains by Light", *The Journal of Physical Chemistry* 29(12):1568–1582, 1925.

[161] H. Baines, *The Science of Photography*, ed. by E. S. Bomback, 3rd, Fountain Press, 1970.

[162] A. Gutierrez-Arroyo, E. Baudet, L. Bodiou, J. Lemaitre, I. Hardy, F. Faijan, B. Bureau, V. Nazabal, and J. Charrier, "Optical characterization at 7.7 μm of an integrated platform based on chalcogenide waveguides for sensing applications in the mid-infrared", *Optics Express* 24(20):23109–23117, 2016.

[163] V. Mittal, N. P. Sessions, J. S. Wilkinson, and G. S. Murugan, "Optical quality ZnSe films and low loss waveguides on Si substrates for mid-infrared applications", *Optical Materials Express* 7(3):712–725, 2017.

[164] E. Baudet, A. Gutierrez-Arroyo, M. Baillieul, J. Charrier, P. Němec, L. Bodiou, J. Lemaitre, E. Rinnert, K. Michel, B. Bureau, J. L. Adam, and V. Nazabal, "Development of an evanescent optical integrated sensor in the mid-infrared for detection of pollution in groundwater or seawater", *Advanced Device Materials* 3(2):23–29, 2017.

[165] C. Xin, H. Wu, Y. Xie, S. Yu, N. Zhou, Z. Shi, X. Guo, and L. Tong, "CdTe microwires as mid-infrared optical waveguides", *Optics Express* 26(8):10944–10952, 2018.

[166] Y. Matsuura, "Optical fibers for medical applications", *Lasers for Medical Applications*, ed. by H. Jelínková, Woodhead Publishing, 2013, chap. 4, pp. 110–124.

[167] P. Lucas, G. J. Coleman, S. Jiang, T. Luo, and Z. Yang, "Chalcogenide glass fibers: Optical window tailoring and suitability for bio-chemical sensing", *Optical Materials* 47:530–536, 2015.

[168] E. Hartouni and J. J. Mecholsky, "Mechanical Properties of Chalcogenide Glasses", *Infrared and Optical Transmitting Materials*, ed. by R. W. Schwartz, vol. 0683, SPIE, 1986, pp. 92–97.

[169] J. S. Sanghera, L. B. Shaw, and I. D. Aggarwal, "Applications of chalcogenide glass optical fibers", *Comptes Rendus Chimie* 15(12):873–883, 2002.

[170] C. Charlton, M. Giovannini, J. Faist, and B. Mizaikoff, "Fabrication and Characterization of Molecular Beam Epitaxy Grown Thin-Film GaAs Waveguides for Mid-Infrared Evanescent Field Chemical Sensing", *Analytical Chemistry* 78(12):4224–4227, 2006.

[171] J. E. Ayers, T. Kujofsa, P. Rago, and J. Raphael, *Heteroepitaxy of Semiconductors: Theory, Growth and Characterization*, 2nd, CRC Press, 2016.

[172] T. M. Donovan, W. E. Spicer, J. M. Bennett, and E. J. Ashley, "Optical Properties of Amorphous Germanium Films", *Physical Review B* 2(2):397–413, 1970.

[173] G. Z. Mashanovich, C. J. Mitchell, J. Soler Penades, A. Z. Khokar, C. G. Littlejohns, W. Cao, Z. Qu, S. Stanković, F. Y. Gardes, T. B. Masaud, H. M. H. Chong, V. Mittal, G. S. Murugan, J. S. Wilkinson, A. C. Peacock, and M. Nedeljkovic, "Germanium Mid-Infrared Photonic Devices", *Journal of Lightwave Technology* 35(4):624–630, 2017.

[174] J. Tauc, A. Abrahám, R. Zallen, and M. Slade, "Infrared absorption in amorphous germanium", *Journal of Non-Crystalline Solids* 4:279–288.

[175] F. Niklaus, P. Enoksson, E. Kälvestan, and G. Stemme, "Low-temperature full wafer adhesive bonding", *Journal of Micromechanics and Microengineering* 11(2):100–107, 2001.

[176] J. Haisma, B. A. C. M. Spierings, U. K. P. Biermann, and A. A. van Gorkum, "Diversity and feasibility of direct bonding: a survey of a dedicated optical technology", *Applied Optics* 33(7):1154–1169, 1994.

[177] D. Pasquariello and K. Hjort, "Plasma-assisted InP-to-Si low temperature wafer bonding", *IEEE Journal of Selected Topics in Quantum Electronics* 8(1):118–131, 2002.

[178] H. G. Stenhouse, S. J. Beecher, and J. I. Mackenzie, "Direct bonding diamond to zinc selenide", *Optical Materials Express* 7(8):2922–2927, 2017.

[179] M. Bruel, "Silicon on insulator material technology", *Electronics Letters* 31(14):1201–1202, 1995.

[180] M. Bruel, B. Aspar, and A.-J. Auberton-Hervé, "Smart-Cut: A New Silicon On Insulator Material Technology Based On Hydrogen Implantation and Wafer Bonding", *Japanese Journal of Applied Physics* 36(3B):1636–1641, 1997.

[181] P. Gueguen, C. Ventosa, L. Di Cioccio, H. Moriceau, F. Grossi, M. Rivoire, P. Leduc, and L. Clavelier, "Physics of direct bonding: Applications to 3D heterogenous or monolithic integration", *Microelectronic Engineering* 87(3):477–484, 2010.

[182] P. Srivastava, J. Das, D. Visalli, J. Derluyn, M. Van Hove, P. E. Malinowski, D. Marcon, K. Geens, K. Cheng, S. Degroote, M. Leys, M. Germain, S. Decoutere, R. P. Mertens, and G. Borghs, "Silicon Substrate Removal of GaN DHFETs for Enhanced (<1100 V) Breakdown Voltage", *IEEE Electron Device Letters* 31(8):851–853, 2010.

[183] Y. Shoji, T. Mizumoto, H. Yokoi, I-W. Hsieh, and R. M. Osgood Jr., "Magneto-optical isolator with silicon waveguides fabricated by direct bonding", *Applied Physics Letters* 92(7):071117, 2008.

[184] Y. Hu, D. Liang, K. Mukherjee, Y. Li, C. Zhang, G. Kurczveil, X. Huang, and R. G. Beausoleil, "III/V-on-Si MQW lasers by using a novel photonic integration method of regrowth on a bonding template", *Light: Science & Applications* 8(93):1–9, 2019.

[185] D. Liang and J. E. Bowers, "Highly efficient vertial outgassing channels for low-temperature InP-to-silicon direct wafer bonding on the silicon-on-insulator substrate", *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena* 26(4):1560–1567, 2008.

[186] B. Corbett, R. Loi, W. Zhou, D. Liu, and Z. Ma, "Transfer print techniques for heterogeneous integration of photonic components", *Progress in Quantum Electronics* 52:1–17, 2017.

[187] A. J. Trindade, B. Guilhabert, D. Massoubre, D. Zhu, N. Laurand, E. Gu, I. M. Watson, C. J. Humphreys, and M. D. Dawson, "Nanoscale-accuracy transfer printing of ultra-thin AlInGaN light-emitting diodes onto mechanically flexible substrates", *Applied Physics Letters* 103(25):253302, 2013.

[188] A. De Groote, P. Cardile, A. Z. Subramanian, A. M. Fecioru, C. Bower, D. Delbeke, R. Baets, and G. Roelkens, "Transfer-printing-based integration of single-mode waveguide-coupled III-V-on-silicon broadband light emitters", *Optics Express* 24(13):13754–13762, 2016.

[189] J. McPhillimy, B. Guilhabert, C. Klitis, M. D. Dawson, M. Sorel, and M. J. Strain, "High accuracy transfer printing of single-mode membrane silicon photonic devices", *Optics Express* 26(13):16679–16688, 2018.

[190] H. Yang, D. Zhao, J.-H. Seo, S. Chuwongin, S. Kim, J. A. Rogers, Z. Ma, and W. Zhou, "Broadband Membrane Reflectors on Glass", *IEEE Photonics Technology Letters* 24(6):476–478, 2012.

[191] J. Wu, S. Kim, W. Chen, A. Carlson, K.-C. Hwang, Y. Huang, and J. A. Rogers, "Mechanics of reversible adhesion", *Soft Matter* 7(18):8657–8662, 2011.

[192] M. R. Billah, M. Blaicher, T. Hoose, P.-I. Dietrich, P. Marin-Palomo, N. Lindenmann, A. Nesic, A. Hofmann, U. Troppenz, M. Moehrle, S. Randel, W. Freude, and C. Koos, "Hybrid integration of silicon photonic circuits and InP lasers by photonic wire bonding", *Optica* 5(7):876–883, 2018.

[193] T. G. Bifano, T. A. Dow, and R. O. Scattergood, "Ductile-regime grinding: a new technology for machining brittle materials", *Journal of Engineering for Industry* 113(2):184–189, 1991.

[194] V. Mittal, A. Aghajani, L. G. Carpenter, J. C. Gates, J. Butement, P. G. R. Smith, J. S. Wilkinson, and G. S. Murugan, "Fabrication and characterization of high-contrast mid-infrared $GeTe_4$ channel waveguides", *Optics Letters* 40(9):2016–2019, 2015.

[195] L. G. Carpenter, S. A. Berry, and C. B. E. Gawith, "Ductile dicing of $LiNbO_3$ ridge waveguide facets to achieve 0.29 nm surface roughness in single process step", *Electronics Letters* 53(25):1672–1674, 2017.

[196] T. A. Birks, J. C. Knight, and P. St J. Russell, "Endlessly single-mode photonic crystal fiber", *Optics Letters* 22(13):961–963, 1997.

[197] J. C. Knight, T. A. Birks, P. St. J. Russell, and J. P. de Sandro, "Properties of photonic crystal fiber and the effective index model", *Journal of the Optical Society of America A* 15(3):748–752, 1998.

[198] K. Bougot-Robin, J. P. Hugonin, M. Besbes, and H. Benisty, "Broad working bandwidth and "endlessly" single-mode guidance within hybrid silicon photonics", *Optics Letters* 40(15):3512–3515, 2015.

[199] M. Cherchi, S. Ylinen, M. Harjanne, M. Kapulainen, and T. Aalto, "Dramatic size reduction of waveguide bends on a micron-scale silicon photonic platform", *Optics Express* 21(15):17814–17823, 2013.

[200] J. W. Nicholson, A. D. Yablon, S. Ramachandran, and S. Ghalmi, "Spatially and spectrally resolved imaging of modal content in large-mode-area fibers", *Optics Express* 16(10):7233–7243, 2008.

[201] D. R. Gray, M. N. Petrovich, S. R. Sandoghchi, N. V. Wheeler, N. K. Baddela, G. T. Jasion, T. Bradley, D. J. Richardson, and F. Poletti, "Real-Time Modal Analysis via Wavelength-Swept Spatial and Spectral (S^2) Imaging", *IEEE Photonics Technology Letters* 28(9):1034–1037, 2016.

[202] L. Zavargo-Peche, A. Ortega-Moñux, J. G. Wangüemert-Pérez, and Í. Molina-Fernández, "Fourier based combined techniques to design novel sub-wavelength optical integrated devices", *Progress In Electromagnetics Research* 123:447–465, 2012.

[203] The MathWorks, Inc., *MATLAB: Genetic Algorithm*, https://uk.mathworks.com/help/gads/genetic-algorithm.html, Accessed: 27th October 2021.

[204] T. Barwicz, B. Peng, R. Leidy, A. Janta-Polczynski, T. Houghton, M. Khater, S. Kamapurkar, S. Engelmann, P. Fortier, N. Boyer, and W. M. J. Green, "Integrated metamaterial interfaces for self-aligned fibre-to-chip coupling in volume manufacturing", *IEEE Journal of Selected Topics in Quantum Electronics* 25(3) 2019.

[205] J. Kennedy and R. Eberhart, "Particle swarm optimization", *Proceedings of ICNN'95 - International Conference on Neural Networks*, vol. 4, 1995, pp. 1942–1948.

[206] J. Robinson and Y. Rahmat-Samii, "Particle swarm optimization in electromagnetics", *IEEE Transactions on Antennas and Propagation* 52(2):397–407, 2004.

[207] R. A. Soref, S. J. Emelett, and W. R. Buchwald, "Silicon waveguided components for the long-wave infrared region", *Journal of Optics: Pure and Applied Optics* 8(10):840–848, 2006.

[208]   J. Stone and G. E. Walrafen, "Overtone vibrations of OH groups in fused silica optical fibers", *The Journal of Chemical Physics* 76(4):1712–1722, 1982.

[209]   A. Witvrouw, B. Du Bois, P. De Moor, A. Verbist, C. A. Van Hoof, H. Bender, and C. Baert, "Comparison between wet HF etching and vapor HF etching for sacrificial oxide removal", *Micromachining and Microfabrication Process Technology VI*, ed. by J. M. Karam and J. A. Yasaitis, vol. 4174, SPIE, 2000, pp. 130–141.

[210]   Juan Valle, Daniel Fernández, and Jordi Madrenas, "Experimental analysis of vapor HF etch rate and its wafer level uniformity on a CMOS-MEMS process", *Journal of Microelectromechanical Systems* 25(2):401–412, 2016.

[211]   R. Wang, S. Sprengel, G. Boehm, M. Muneeb, R. Baets, M.-C. Amann, and G. Roelkens, "2.3 µm range InP-based type-II quantum well Fabry-Perot lasers heterogeneously integrated on a silicon photonic integrated circuit", *Optics Express* 24(18):21081–21089, 2016.

[212]   A. Spott, J. Peters, M. L. Davenport, E. J. Stanton, C. D. Merritt, W. W. Bewley, I. Vurgaftman, C. S. Kim, J. R. Meyer, J. Kirch, L. J. Mawst, D. Botez, and J. E. Bowers, "Quantum cascade laser on silicon", *Optica* 3(5):545–551, 2016.

[213]   A. Spott, E. J. Stanton, A. Torres, M. L. Davenport, C. L. Candey, I. Vurgaftman, M. Kim, C. S. Kim, C. D. Merritt, W. W. Bewley, J. R. Meyer, and J. E. Bowers, "Interband cascade laser on silicon", *Optica* 5(8):996–1005, 2018.

[214]   R. Wang, S. Sprengel, M. Muneeb, G. Boehm, R. Baets, M.-C. Amann, and G. Roelkens, "2 µm wavelength range InP-based type-II quantum well photodiodes heterogeneously integrated on silicon photonic integrated circuits", *Optics Express* 23(20):26834–26841, 2015.

[215]   J. J. Ackert, D. J. Thomson, L. Shen, A. C. Peacock, P. E. Jessop, G. T. Reed, G. Z. Mashanovich, and A. P. Knights, "High-speed detection at two micrometres with monolithic silicon photonics", *Nature Photonics* 9:393–396, 2015.

[216]   M. Muneeb, A. Vasiliev, A. Ruoco, A. Malik, H. Chen, M. Nedeljkovic, J. S. Penades, L. Cerutti, J. B. Rodriguez, G. Z. Mashanovich, M. K. Smit, E. Tourni, and G. Roelkens, "III-V-on-silicon integrate micro-spectrometer for the 3 µm wavelength range", *Optics Express* 24(9):9465–9472, 2016.

[217]   R. Fain, S. A. Miller, M. Yu, A. G. Griffith, J. Cardenas, and M. Lipson, "CMOS-compatible Mid-Infrared Silicon Detector", *Conference on Lasers and Electro-Optics*, Optical Society of America, 2017, STu1N.4.

[218]   B. W. Jia, K. H. Tan, W. K. Loke, S. Wicaksono, K. H. Lee, and S. F. Yoon, "Monolothic integration of InSb photodetector on silicon for mid-infrared silicon photonics", *ACS Photonics* 5(4):1512–1520, 2018.

[219]   D. J. Thomson, L. Shen, J. J. Ackert, E. Huante-Ceron, A. P. Knights, M. Nedeljkovic, A. C. Peacock, and G. Z. Mashanovich, "Optical detection and modulation at $2\,\mu m - 2.5\,\mu m$ in silicon", *Optics Express* 22(9):10825–10830, 2014.