High-Confinement Alumina Waveguides for Nonlinear Optics in the Visible and Ultraviolet

Elissa Grace McKay

Submitted in fulfilment of the requirements for the degree of
Doctor of Philosophy

School of Engineering

College of Science and Engineering
University of Glasgow

July 2024

Declaration of Authorship

I, Elissa McKay, declare that this thesis titled "High-Confinement Alumina Waveguides for Nonlinear Optics in the Visible and Ultraviolet", and the work presented in it are my own, other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified) and that the thesis has not been edited by a third party beyond what is permitted by the University's PGR Code of Practice. This thesis does not include work forming part of a thesis presented successfully for another degree. This thesis has been produced in accordance with the University of Glasgow's Code of Good Practice in Research. I acknowledge that if any issues are raised regarding good research practice based on review of the thesis, the examination may be postponed pending the outcome of any investigation of the issues.

Natale Pruiti carried out HSQ optimisation work described in Section 3.2, and Matteo Clerici performed dispersive wave phase-matching simulations described in Section 5.4.

This work was supported by the Defence Science and Technology Laboratory.

Work presented in this thesis has contributed to the following publications:

1. E. McKay, N. G. Pruiti, S. May, and M. Sorel, "Multipass lithography of HSQ etch masks for fabrication of low-loss alumina waveguides for blue light applications," in *Integrated Optics: Devices, Materials, and Technologies XXVII*, p. 124245, Jan. 2023.

2. E. McKay, N. G. Pruiti, S. May, and M. Sorel, "High-confinement alumina waveguides with sub-dB/cm propagation losses at 450 nm," *Scientific Reports*, vol. 13, p. 19917, Nov. 2023.

3. E. McKay, N. G. Pruiti, C. Suciu, M. Clerici, and M. Sorel, "Ultraviolet to Infrared Supercontinuum Generation in Alumina Waveguides", submitted to *Optics Letters*, 2024.

Abstract

There is increasing interest in photonic integrated circuits (PICs) operating in the blue and ultraviolet (UV), driven by applications such as sensing and optical timing. Novel material platforms are required to operate at these wavelengths; amorphous alumina has an exceptionally wide bandgap, and is therefore a promising candidate. High-confinement waveguide geometries allow for high component integration densities and are essential for nonlinear frequency conversion; however, they can be challenging to fabricate and are sensitive to material absorption losses. This work establishes a methodology for fabrication of high-confinement alumina waveguides, then demonstrates that fabricated devices are suitable for low-loss visible waveguiding and integrated nonlinear optics.

Challenges in developing a high-confinement alumina platform are twofold. Atomic layer deposition (ALD) is a deposition method which can produce pure and homogenous films with low absorption losses; however, depositing amorphous alumina using a conventional thermal ALD process is extremely slow, placing a practical limit on the maximum attainable film thickness. Using an oxygen plasma as the oxidant considerably reduces deposition time; in this work, I establish that alumina films deposited using a dioxygen plasma-enhanced atomic layer deposition (PEALD) process can be used to produce low-loss optical waveguides. Furthermore, alumina is difficult to etch selectively, especially whilst maintaining the low sidewall roughness required for operation in the UV-visible region. I address this by developing a waveguide definition process using an hydrogen silsesquioxane (HSQ) mask exposed using electron beam lithography (EBL) and a BCl_3 etch plasma, which allows etching of waveguides up to 800 nm thick.

Using this process, I fabricate 400 nm-thick amorphous alumina waveguides with confinement factors over 0.9 at a wavelength of 450 nm. With propagation losses of 0.8 dB/cm, these are the lowest-loss high-confinement waveguides reported for blue light. I present dispersion engineering calculations which demonstrate the necessity of high-confinement, unclad waveguides for stimulation of spectral broadening through nonlinear optical processes, and demonstrate the first supercontinuum generated in alumina waveguides, extending from 384 nm to 951 nm at a level of −50 dB, in 700 nm-thick waveguides. These results establish the viability of amorphous alumina deposited using dioxygen PEALD as a wide-bandgap material platform for UV-visible linear and nonlinear optics.

Acknowledgements

Thanks to Marc Sorel, for his supervision, support, and guidance; his feedback and encouragement were essential to completing this thesis. Stuart May provided early guidance and came up with the idea of using ALD alumina for nonlinear optics ("easy paper?" he said). Also, thanks to Matthew Aldous and Keith Newton at DSTL for their help and guidance in getting to the finish line.

Thanks to the JWNC staff, especially Linda Pollock; James Grant and Corrie Farmer, who were so generous in the time they spent providing etch advice and support; Donald Nicolson, who expedited many an etch run; and Paul Reynolds, who advised on EBL. Dilini Tania Hemakumara developed the ALD recipes used in this work.

Natale Giovanni Pruiti made a considerable contribution to this work: we collaborated to refine the Optoelectronics group's HSQ process. Natale also analysed the origin of waveguide losses, displayed in Chapter 4, and his work on optical setups was invaluable. Cosmin Suciu helped with risk assessment and measurement setup for supercontinuum generation—thanks, Coz, for staying in the lab until 8pm on a Friday to ensure I didn't torch my eyes out. I'm glad you got to see the first ever supercontinuum in an alumina waveguide as a result. Vera Biryukova knows everything there is to know about HSQ and shared lots of useful information. Martin McIntosh shared knowledge and resources on EBL and inorganic FTIR spectroscopy. William Peveler, of the School of Chemistry, provided access to a UV-visible spectrophotometer. Matteo Clerici provided expertise in nonlinear optics, as well as a really big laser. Stephen Strowes developed the LaTeX template used for this document, whilst all figure colour schemes were developed by ColorBrewer [1].

Most of the results reported in this thesis were produced in the final six months, and the most difficult part of this work was maintaining morale through years of failure. Thanks to those who literally and/or metaphorically held me when I was at the end of my rope: Cosmin Suciu, Di Gilbert, Adam McKenzie, Kirstie McAllister, Michael Sloan, and my parents.

Ciaran McCreesh had endless discussions about writing and experimental work, was the arbiter of good taste in all things LaTeX, very patiently performed rubber duck duties, cooked me huge amounts of delicious food, gave me a place to live for eight months during the Great Sewage Flood of 2023, and provided entirely unreasonable levels of emotional support. He now knows more about nonlinear optics than any layperson ever needs to know. Ciaran, I would still love you if you were a worm.

This work was funded by the Defence Science and Technology Laboratory.

Table of Contents

Abstract

Acknowledgements

1 Introduction **1**
- 1.1 Photonic Integrated Circuits . . . 2
- 1.2 Waveguiding Basics . . . 5
  - 1.2.1 Loss in Optical Materials . . . 6
- 1.3 Materials for Low-Loss Blue and UV Waveguiding . . . 10
- 1.4 Alumina Waveguides and PICs . . . 15
- 1.5 UV-Visible Supercontinuum Generation . . . 20
- 1.6 Aims and Objectives . . . 25

2 Atomic Layer Deposition of Optical-Quality Alumina **26**
- 2.1 Deposition . . . 28
  - 2.1.1 Film Deposition . . . 31
- 2.2 Ellipsometry and Optical Constants . . . 32
- 2.3 Material Bandgap . . . 35
- 2.4 Fourier Transform Infrared (FTIR) Spectroscopy: Chemical Composition . . . 38
  - 2.4.1 Method . . . 39
  - 2.4.2 Composition of As-Deposited Films . . . 39
- 2.5 Surface Roughness . . . 41
- 2.6 Film Properties Over Time . . . 42
- 2.7 Chemical Compatibility . . . 44
- 2.8 Annealing and Thermal Budget . . . 46
  - 2.8.1 Method . . . 47
  - 2.8.2 Microscopy: Cracking and Outgassing in Annealed Films . . . 48
  - 2.8.3 Ellipsometry: Optical Constants of Annealed Films . . . 48
  - 2.8.4 Bandgap Changes in Annealed Films . . . 50
  - 2.8.5 FTIR: Bonding in Annealed Films . . . 52
  - 2.8.6 AFM: Surface Roughness in Annealed Films . . . 53

2.8.7 Discussion . . . 53
2.9 Conclusion . . . 55

3 Alumina Processing Toolbox **57**
3.1 Process Flow . . . 58
3.2 Lithography . . . 59
3.2.1 Theory of EBL . . . 61
3.2.2 HSQ Chemistry and Application . . . 62
3.2.3 Base Lithography Process . . . 64
3.2.4 Maximum Mask Thickness . . . 65
3.2.5 Improving HSQ Etch Resistance . . . 66
3.2.6 Line-Edge Roughness Optimisation . . . 67
3.2.7 Rectification of Stitching Errors . . . 75
3.2.8 Finalised Process . . . 77
3.3 Etching . . . 79
3.3.1 Theory of Plasma Etching . . . 79
3.3.2 Etch Development . . . 80
3.4 Cladding . . . 85
3.5 Cleaving . . . 87
3.6 Conclusion . . . 89

4 Single-Mode Waveguides **91**
4.1 Theory . . . 92
4.2 Design . . . 97
4.2.1 Functional Requirements . . . 98
4.2.2 Fabrication Requirements . . . 99
4.2.3 Metrology Requirements . . . 99
4.2.4 Design Simulations . . . 100
4.3 Fabrication . . . 105
4.4 Test Method . . . 105
4.5 Results . . . 106
4.5.1 Revision 4 . . . 107
4.5.2 Revisions 5 and 6 . . . 108
4.5.3 Coupling Losses . . . 109
4.6 Low-Confinement Waveguides . . . 110
4.7 Discussion . . . 111
4.7.1 Wider Context . . . 111
4.7.2 Measurement Technique . . . 113
4.7.3 Future Improvements . . . 115
4.8 Conclusion . . . 116

5 Supercontinuum Generation in Alumina Waveguides **117**
5.1 Theory . . . 117
5.1.1 Optical Kerr Effect . . . 118
5.1.2 Self-phase modulation (SPM) . . . 119
5.1.3 Cross-Phase Modulation (XPM) . . . 120
5.1.4 Raman Scattering . . . 120
5.1.5 Soliton Fission . . . 121
5.1.6 Optical Wave Breaking (OWB) . . . 123
5.1.7 Four-Wave Mixing (FWM) . . . 123
5.2 Design and Fabrication . . . 124
5.2.1 Functional Requirements . . . 125
5.2.2 Fabrication Requirements . . . 126
5.2.3 Design Simulations . . . 126
5.2.4 Mask Design . . . 131
5.2.5 Fabrication . . . 131
5.3 Test Method . . . 132
5.4 Results and Discussion . . . 132
5.4.1 Fabrication . . . 132
5.4.2 Supercontinuum Generation . . . 134
5.4.3 Discussion . . . 141
5.5 Conclusion . . . 144

6 Conclusion **146**
6.1 Optical-Quality PEALD Alumina Films . . . 146
6.2 Fabrication of High-Confinement Ridge Waveguides . . . 147
6.3 Supercontinuum Generation . . . 147
6.4 Future Directions . . . 148

A Reviews **150**

B Process Flows **153**
B.1 Process Flows for Optical Devices . . . 154
B.2 Process Flows for Fabrication Development . . . 156

C Other Sources of Optical Loss **158**
C.1 Coupling Loss . . . 159
C.2 Radiative Loss . . . 159

References **161**

Abbreviations and Acronyms

AFM	Atomic force microscope
ALD	Atomic layer deposition
CAD	Computer-aided design
CD	Critical dimension
CE	Coupling efficiency
CMOS	Complementary metal-oxide semiconductor
CTE	Coefficient of thermal expansion
CVD	Chemical vapour deposition
DFT	Density functional theory
DW	Dispersive wave
EBL	Electron beam lithography
FDE	Finite difference eigenmode
FTIR	Fourier transform infrared
FWM	Four-wave mixing
GNLSE	Generalised nonlinear Schrödinger equation
GPC	Growth per cycle
GPS	Global Positioning System
GVD	Group velocity dispersion
HF	Hydrofluoric acid
HSQ	Hydrogen silsesquioxane
IC	Integrated circuit
ICP	Inductively coupled plasma
IPA	Isopropyl alcohol
IR	Infrared
JWNC	James Watt Nanofabrication Centre
LER	Line-edge roughness
LiDAR	Light detection and ranging
LPCVD	Low-pressure chemical vapour deposition
MIBK	Methyl isobutyl ketone

MOCVD	Metal-organic chemical vapour deposition
MRR	Micro-ring resonator
NEXAFS	Near-edge X-ray absorption fine structure spectroscopy
NMR	Nuclear magnetic resonance
NV	Nitrogen vacancy
OWB	Optical wave breaking
PEALD	Plasma-enhanced atomic layer deposition
PEB	Post-exposure bake
PEC	Proximity error correction
PECVD	Plasma-enhanced chemical vapour deposition
PIC	Photonic integrated circuit
PMMA	Polymethyl methacrylate
PVD	Physical vapour deposition
RF	Radio frequency
RIE	Reactive ion etch
RTA	Rapid thermal annealer
SCG	Supercontinuum generation
SEM	Scanning electron microscope
SHG	Second harmonic generation
SI	International System of Units
SOI	Silicon-on-insulator
SPM	Self-phase modulation
SRS	Stimulated Raman scattering
SSFS	Soliton self-frequency shift
TE	Transverse electric
TEOS	Tetraethoxysilane
TGA	Thermogravimetric analysis
THG	Third harmonic generation
TM	Transverse magnetic
TMA	Trimethylaluminium
TMAH	Tetramethylammonium hydroxide
TOD	Third-order dispersion
UV	Ultraviolet
VASE	Variable angle spectroscopic ellipsometry
XPM	Cross-phase modulation
XPS	X-ray photoelectron spectroscopy

Chapter 1

Introduction

Integrated photonics offers the capacity to perform complex optical processing on a single chip. Where electronic integrated circuits process data using transistors and electrons, photonic integrated circuits (PICs) modify and process light. Whilst most commercially available PIC process infrared (IR) light for high-bandwidth optical communications, many emerging applications of PICs require visible and UV light. Sensors can use broadband UV-visible light to probe biological molecules, whilst optical atomic clocks require UV-visible frequency conversion to stabilise timing lasers. Creating PIC-based broadband UV-visible sources would allow such devices to be made cheaply and in miniature.

The basic component of all PICs is the *waveguide*—a structure engineered to guide light through total internal reflection. Maximising signal transmission through a waveguide is critical; however, optical losses occur as a result of material absorption and light scattering, both of which increase dramatically as wavelength decreases [2, 3]. Consequently, developing material platforms which have low losses in the UV-visible region is an ongoing challenge. Producing waveguides with properties which can generate UV-visible light through nonlinear spectral broadening is particularly challenging, as design requirements demand that the waveguide cannot be engineered to minimise absorption or scattering losses.

Amorphous alumina waveguides have demonstrated the lowest UV-visible losses to date; however, prior work has used low-confinement waveguides, which minimise loss and are easy to fabricate, but have limited practical utility [4]. Consequently, there were no demonstrations of waveguide-based nonlinear broadening processes in etched alumina waveguides prior to this work. This thesis documents the development of an amorphous alumina material platform for visible and UV photonic integrated circuits (PICs), including characterisation of alumina films, development of processes for definition of high-confinement waveguides, testing of optical losses at visible wavelengths, and nonlinear generation of broadband light from UV to IR.

This chapter contextualises the work presented in this thesis. I begin by exploring applications driving a need for UV-visible material platforms, and considering the technological advances which could be achieved by miniaturisation of optical assemblies (Section 1.1). I introduce waveguides as the core building block of all PICs, and explain the origins of optical loss, demonstrating that UV-visible waveguiding is difficult because of the wavelength-dependence of both absorption and scattering losses (Section 1.2). I consider current benchmarks and challenges for waveguiding in the blue and UV region, justifying why alumina is a favourable material choice (Section 1.3), then focus more closely on prior work on alumina waveguides, with a particular view to applications in nonlinear optics (Section 1.4). There is very little prior work on nonlinear optics in alumina films, but Section 1.5 provides a brief review of spectral broadening in waveguides in the UV-visible region, considering challenges and figures of merit. With this context in place, Section 1.6 outlines the aims and objectives of this project.

1.1 Photonic Integrated Circuits

The term *integrated circuit* (IC) is most commonly applied to electronic circuits, describing the simultaneous planar fabrication of a large number of electronic components on a single substrate. Similar design and manufacturing principles can be applied to photonic devices, presenting similar advantages in terms of miniaturisation and mass production. These *photonic integrated circuits* (PICs) use optical waveguides to transport and modify light.

By the mid-20th century, the use of a solid structure to confine the direction of travel of electromagnetic radiation was well-understood, and radio waveguides—metal pipes—were a commercial technology [5]. The invention of ruby, then semiconductor, lasers in the 1960s [6, 7] enabled translation of this well-known phenomenon to the optical domain; the first PIC was proposed by Miller in 1969 [8] and demonstrated later that year [9]. Whilst developing laser repeaters, Miller noted the inefficiency and variability within his experimental setups. He envisioned a future in which optical table setups could be miniaturised onto a single chip: “If realized, this new art would facilitate isolating the laser circuit assembly from thermal, mechanical, and acoustic ambient changes through small overall size; economy should ultimately result.” [8] That economy, Miller noted, could take two forms: the first is the production efficiency of simultaneously creating many components using established high-volume semiconductor processing techniques. The second is within the system itself: the small size of the system should reduce its optical loss, allowing lower laser powers to be used without detriment to purpose. Miller even recognised that PICs would allow nonlinear optical processes to be harnessed for signal modulation and frequency conversion.

As Miller envisioned, PICs' first commercial applications were in the telecommunications industry, where they have have driven bandwith increases in optical fibre communications. Monolithically integrated sources, modulators, and amplifiers for near-IR light, made using III-V semiconductors, have been used in commercial telecoms systems since 2003; later developments such as multiplexers and demultiplexers have used Si waveguides to further increase data transmission bandwidths [10, 11]. Historically, the telecoms industry's needs have driven the focus of PIC development towards infrared platforms using Si and III-V semiconductors, which also integrate well into existing semiconductor manufacturing infrastructure. However, a range of emerging applications require waveguiding in the visible and even UV regimes, where these semiconductor materials are not transparent.

Optical atomic clocks are a driver of UV-visible PIC research, especially research into frequency conversion [15]. At present, the International System of Units (SI) definition of the second is based on emitted radiation from a ^{133}Cs hyperfine transition in the microwave region [16], and ^{133}Cs atomic clocks are used internationally as timing standards [17]. The upper limit of an atomic clock's timing stability is defined by the frequency of the oscillations used for timing, so atomic clocks using optical frequencies offer a significant timing stability improvement. Optical atomic clocks use a *clock laser* which stimulates an electronic transition in an atom or ion; many transitions of interest require a clock laser in the visible or UV (Figure 1.1). The Boulder Atomic Clock Optical Network (BACON) Collaboration, for example, has demonstrated ^{87}Sr, ^{171}Yb and ^{27}Al^{+} optical clocks which each offer a hundredfold increase of stability relative to the ^{133}Cs microwave standard. Comparing emissions from multiple species may further improve timing stability [18].

There are a range of applications in which we may want miniaturised but extremely precise timing, such as in gravimetry or Global Positioning System (GPS) navigation; PICs offer a route to producing robust, stable clocks with small form factors. However, clock lasers operating at optical frequencies cannot be processed by conventional electronics. Integrated devices offer a compact solution: in a microresonator, nonlinear processes can create a *frequency comb*, radiation consisting of resonances which are equally spaced in the frequency domain. Tooth spacing in a microresonator is typically in the GHz range and can be measured by electronics. A sufficiently broad comb, containing frequency multiples, can be stabilised using processes such as *f-2f self-referencing*, to ensure that timing stability is maintained [19].

Sensing and imaging technologies also benefit from integrated processing of UV-visible light. Underwater light detection and ranging (LiDAR) systems require blue or green light to maximise range [12, 20]. Biological molecules, such as proteins, have absorptions throughout the UV, blue, and green [21, 22], which can be exploited in multiple ways. Waveguide-based sensors operate by immersing a functionalised waveguide in an analyte solution, passing light through the waveguide, and measuring the change in refractive index relative to a standard.

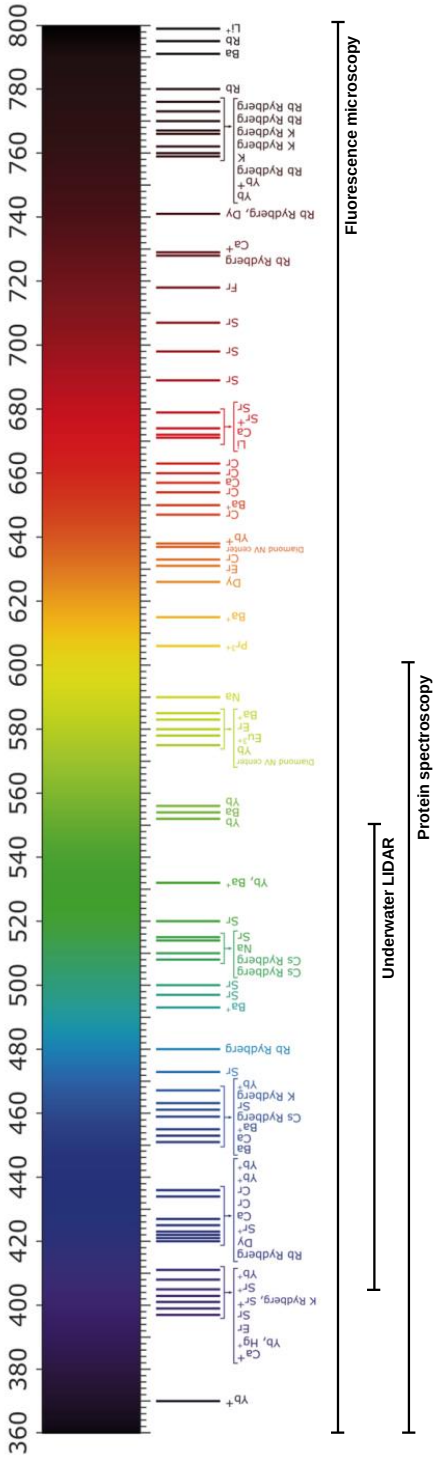

Figure 1.1: PIC applications in the UV-visible region. Vertical labels show wavelengths of interest in optical timing. Adapted with permission from Corato-Zanarella et al. [2]; additional data from Pershin et al. [12], Lin et al. [13], and Tinguely et al. [14].

This technique offers a simple alternative to time-consuming chemical analyses such as assays, which is especially significant for medical screening [23, 24]. Fluorescence microscopy, in which short-wavelength light is used to probe biological tissue, can similarly benefit from waveguides which can transmit blue or UV light [13, 14].

The integration of photon sources presents a bottleneck for these applications: light emitters are typically fabricated using different material platforms than other PIC components. Nonlinear optics can simplify photonic integration processes: a frequency comb or supercontinuum can produce a broad range of wavelengths for multiplexed use whilst only requiring integration of a single source. Sources based on nanotechnology, such as quantum dots, may offer alternative integration routes. There is considerable interest in photonic integration of diamond nitrogen vacancy (NV) centres, which, when excited, produce visible photons with long-lived—and therefore, detectable—quantum states [25].

Producing high-quality waveguides which have low losses in the blue and UV presents an ongoing challenge. At present, amorphous alumina is the only material platform to demonstrate sub-dB/cm losses in the blue [4]; technology demonstrators using the material platform have included autofluorescence microscopes [13] and waveguide-based biological sensors [23]. However, these demonstrations are limited in scope: they have used *low-confinement* waveguides, in which the waveguide itself has a small cross-section and the optical power is diffuse. Larger high-confinement waveguides, in which the majority of the optical power travels within the waveguide structure, offer several advantages over lower-confinement waveguides: they allow for higher component integration density, and are required for efficient nonlinear frequency conversion. However, their fabrication poses additional challenges, addressed in this work.

1.2 Waveguiding Basics

Optical waveguides are a building block for many of the optical components used in PICs, such as lasers and gain media, interferometers, and multiplexers [26]. Waveguides consist of a high-refractive-index core, which confines light, surrounded by a lower-index cladding material. Using a simple ray-optics description, light is described as a plane wave propagating through a medium in the direction of the 'ray'; when light is coupled into the core, it experiences total internal reflection at the core-cladding boundary and therefore propagates through the core material (Figure 1.2) [26].

Common waveguide geometries are shown in Figure 1.3. *Slab* waveguides consist of a film of high-refractive-index material on top of a lower-index cladding material, typically a dielectric such as SiO_2. Slab waveguides only confine light in one dimension, and therefore have few practical applications; however, they can be used to test whether materials are suitable for

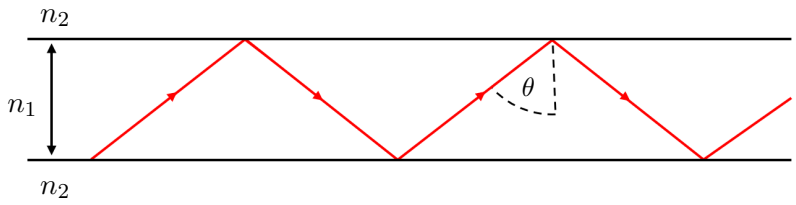

Figure 1.2: Total internal reflection within a slab waveguide. Here, a symmetrical waveguide is shown; the core refractive index n_1 is higher than the cladding refractive index n_2. If θ is above the critical angle, the light experiences total internal reflection.

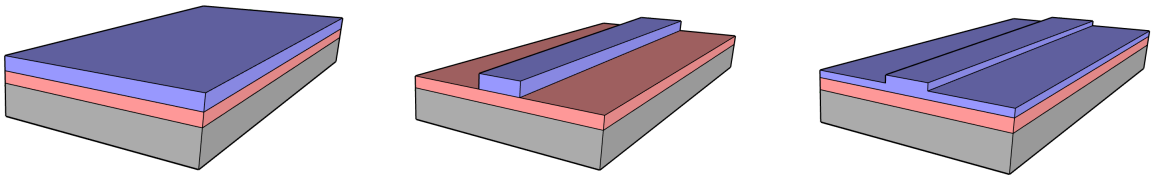

(a) Slab (b) Ridge (c) Rib

Figure 1.3: Slab, ridge and rib waveguide geometries. Purple areas represent the core material; red, the bottom cladding; grey, the carrier. For clarity, no upper cladding is shown.

use in optical waveguiding applications. *Ridge* waveguides consist of an isolated strip of core material on top of a cladding layer; this design confines light in two dimensions. This common design has numerous practical applications; in this thesis, 'waveguide' refers to a ridge waveguide unless otherwise specified. A *rib* waveguide consists of a ridge superimposed onto a slab. An upper cladding layer can be added to any of these designs; however, this is not essential [26].

Optical waveguides provide the basic building block for almost all of the optical components in PICs, such as micro-ring resonators (MRRs), mode converters, and beamsplitters. These may be fabricated using exactly the same process; the only difference is in the structure which is defined in the masking stage. Therefore, large and complex circuits can be fabricated simultaneously using the same steps: deposition of a core layer, defining the optical structures in a mask layer, and etching to transfer structures into the core layer. This thesis focuses on the development of a manufacturing process which can be used to create basic waveguides and which can be extended to fabricate a range of primarily passive components.

1.2.1 Loss in Optical Materials

The key purpose of a waveguide is to transmit as much light from source to receiver as possible; therefore, key figures of merit centre around loss. Loss can be measured in decibels

(dB), a logarithmic ratio of transmitted power P_1 to received power P_2,

$$Loss\ (dB) = 10\ \log \frac{P_2}{P_1} \tag{1.1}$$

The primary sources of loss in a well-designed waveguide are *coupling loss* and *propagation loss*. Coupling loss, measured in dB per facet, describes how much light is lost in the process of coupling light from a source into a waveguide; propagation loss, measured in dB/m or dB/cm, describes how much light is lost during travel through the waveguide. *Radiative loss*, the leakage of the mode out of the waveguide, can occur, but should be minimal in an optimised PIC design. The total (*insertion*) loss associated with passing light through an optical component consists of the sum of these effects. Coupling and radiative loss mechanisms are explored in Appendix C.

Minimisation of propagation loss is a core focus of this thesis. Reducing propagation loss allows an optical signal to retain more of its initial power over a transmission distance. Longer transmission lengths allow us to

- Create larger and more complex PICs
- Transmit lower signal levels, which may improve the limits of detection and quantification in sensor or single-photon applications
- Create MRR with high finesse (sharp resonances), which improves performance in optical timing and sensing applications [27]
- Engage in a wider range of nonlinear optical processes (described in more detail in Chapter 5; taken simplistically, we can say that nonlinear optical processes are intensity-dependent and therefore rely on retaining light within the waveguide)

Sources of propagation loss fall into two categories: absorption and scattering. *Intrinsic material absorption* defines the wavelength range in which a material is transparent and can therefore be used for low-loss waveguiding. The most commonly encountered form of intrinsic material absorption is *band edge absorption*, where a high-energy photon is absorbed to create an exciton; if there is sufficient energy, an electron is promoted to the conduction band. It is important to note the difference between the *electronic bandgap* (the photon energy required to create an electron-hole pair) and the *optical bandgap* (the photon energy required to produce an exciton, without necessarily overcoming the attractive forces required to separate the electron and hole) [28]. In the context of low-loss optics, it is the absorption of the photon, and the optical bandgap, which is relevant. The positions of the valence and conduction bands, and therefore the bandgap energy, are defined by the elemental and bonding composition of the material. As a result of quantisation, light absorption increases sharply at the bandgap energy; photons of energy higher than the bandgap can be absorbed.

However, a critical factor in integrated photonics, especially in the visible regime, is intrinsic absorption below the bandgap energy, a phenomenon known as the *Urbach tail* (Figure 1.4) [29]. Absorption in the Urbach tail region is substantially lower than at the bandgap, but nevertheless sets the lower bound for practical use of a material in waveguiding applications [30]. In amorphous solids, the Urbach tail is broader than in crystalline solids, as a result of the localised variation in bonding, and therefore energy levels, within the material [31]. *Free carrier absorption*, in which electrons already in the conduction band are excited to a higher energy state, is an additional source of absorption loss. However, free carrier absorption only produces appreciable effects in semiconductors [32], so will not be considered here.

Intrinsic material absorption would exist even in a perfect waveguide made of a perfectly homogenous film. However, real materials inevitably contain both elemental impurities and lattice defects, which may strongly absorb light at wavelengths where inherent material absorption is low. This is known as *extrinsic absorption*. A small concentration of impurities or defects can have a considerable impact on waveguide loss; for example, significant efforts have been made to reduce the concentration of OH^- impurities in silica optical fibres since these produce absorptions near telecoms wavelengths [33].

Therefore, a waveguide material must be selected based on the target wavelength of PIC operation. A waveguide core material optimised to minimise absorption losses should have a bandgap energy substantially higher than the energy of the propagating light, an Urbach tail which ends at a higher energy than the propagating light, and should be pure and defect-free. Guided modes also extend into the waveguide cladding (see Section 4.1), so these requirements also apply to cladding materials.

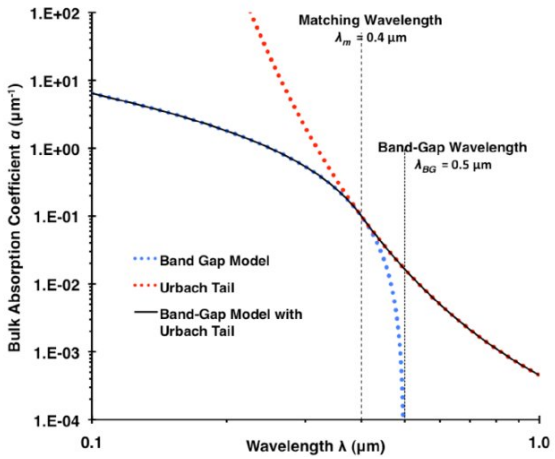

Figure 1.4: Impact of Urbach tail on absorption below bandgap energy. The blue dotted line shows a modelled bandgap without the influence of the Urbach tail, with an abrupt rise in absorption at $\lambda = 500\,\text{nm}$. The solid black line indicates a model incorporating the Urbach tail, which introduces a degree of absorption at energies lower than the bandgap. Reproduced from Moosmüller et al. [34] under Creative Commons Attribution 3.0 License.

Scattering losses occur when propagating light encounters a change in refractive index. *Volume scattering* occurs within the bulk of the core material, whilst *interface scattering* occurs at the interface of core and cladding.

Volume scattering occurs through different pathways—*Rayleigh scattering* or *Mie scattering*—depending on the size of the inhomogeneity relative to the propagating light. Rayleigh scattering occurs when the refractive index of a material varies on scales much shorter than the wavelength of light, producing a dipole which radiates light at the same frequency as the propagating wave, but in any direction. The intensity of Rayleigh scattering is proportional to $1/\lambda^4$, where λ is the wavelength of the propagating light; the shorter the wavelength of the propagating light, the greater the effect of Rayleigh scattering. A robust exploration of this phenomenon is provided in Boyd [35]. Rayleigh scattering is stimulated by a change in polarisability, and occurs in all materials as a result of oscillations of charge within the material. However, small-scale variations in refractive index, such as the presence of defects (point defects such as inclusions or vacancies, as well as larger defects such as nanovoids), will also stimulate Rayleigh scattering [36].

It would be intuitive to conclude that amorphous solids will stimulate higher levels of scattering than crystals, as a result of their disordered structure. However, low bulk scattering losses are attainable in amorphous materials, and recent research hypothesises that stress and elasticity correlations create a long-range "order" within amorphous solids [37]. Nevertheless, assuming the absence of nonlinear effects, Rayleigh scattering sets the lower bound for attainable loss in a transparent, unpatterned material [38].

When the size of the inhomogeneity reaches $\sim$10% of the wavelength of the light, the Rayleigh scattering model is no longer applicable and the Mie scattering model can instead be used to model the effects of inhomogeneities. Unlike Rayleigh scattering, Mie scattering is directional; the larger the inhomogeneity, the more light is scattered in the direction of propagation [39].

Scattering at waveguide sidewalls—wherever the waveguide interfaces with a lower-refractive-index surface—can be modelled differently. The Payne and Lacey model [3] of sidewall scattering considers sidewall scattering in two dimensions; more complex models have been developed, which consider the effects of the third dimension and more complex effects such as radiative loss [40]. However, a two-dimensional model is sufficient to provide a conceptual understanding of scattering in most dielectric waveguides. The Payne and Lacey model [3] describes loss in terms of the refractive indices of the core and cladding materials n_1 and n_2 in a waveguide of nominal width $2d$:

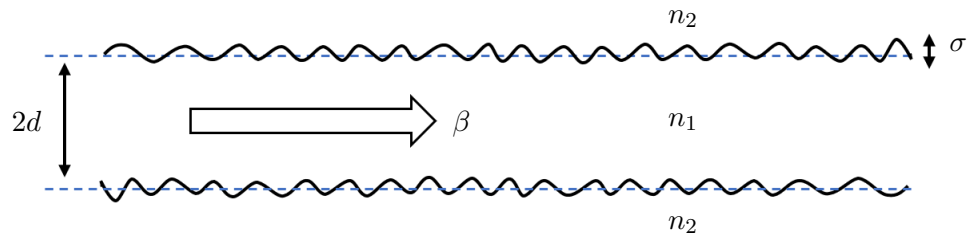

Figure 1.5: 2D slab waveguide described by Payne and Lacey scattering model. Redrawn from Payne and Lacey [3].

Given a modal propagation constant β (see Section 4.1), a modal field of $\varphi(d)$, scattering angle θ, and a free-space wavenumber k_0, Payne and Lacey [3] define propagation loss α as

$$\alpha = \varphi^2(d)\,(n_1^2 - n_2^2)^2\,\frac{k_0^3}{4\pi n_1}\int_0^{\pi} \tilde{R}(\beta - n_2 k_0 \cos\theta)\; d\theta \tag{1.2}$$

The autocorrelation function $\tilde{R}(\Omega)$ describes the sidewall roughness σ. Deriving a term

$$g(V) = \frac{d^4 k_0^2 (n_1^2 k_0^2 - \beta^2)(n_1^2 - n_2^2)}{1 + d\sqrt{\beta^2 - n_2^2 k_0^2}} \tag{1.3}$$

yields a further result [3]

$$\alpha \leq \frac{\sigma^2}{2 k_0 d^4 n_1}\, g(V) \tag{1.4}$$

which indicates that, for a waveguide of fixed geometry and composition (fixed n, k_0, β), the upper limit for scattering loss is set only by the waveguide sidewall roughness.

Consequently, interface scattering losses can be reduced by:

- Minimising difference in refractive index between core and cladding $(n_1^2 - n_2^2)$
- Using a core material with a low refractive index $({}^{(n_1^2-n_2^2)^2}/_{n_1})$
- Using the longest possible wavelength of light (k_0^3, where $k_0 = {}^1/_{\lambda}$)
- Reducing sidewall roughness (Equation 1.4)

1.3 Materials for Low-Loss Blue and UV Waveguiding

When selecting a PIC material platform, the material bandgap is often the primary consideration, as the intrinsic absorption loss sets the lower bound for propagation loss within a waveguide. However, other drivers of loss have a considerable influence on material choice. Material deposition processes must be capable of producing chemically pure films (to minimise extrinsic absorption and Rayleigh scattering) with a homogenous microstructure (to minimise both Rayleigh and Mie scattering).

Further losses are introduced during waveguide definition. Scattering losses arise at the boundary between waveguide core and cladding; whilst these are dependent on the refractive indices of the two materials, minimising waveguide sidewall roughness will reduce loss regardless of waveguide composition or geometry. It is also important to consider the effects of the fabrication process on bulk losses: heat or chemical processing may affect material structure and therefore intrinsic absorption or bulk scattering, as well as introducing external absorbers. Intrinsic absorption losses and bulk and interface scattering losses all increase as wavelength decreases, so low-loss waveguiding at blue and UV wavelengths presents a considerable challenge.

Early developments in PIC technology were driven by the telecoms industry [41], which has established standards in the infrared area of the spectrum [42]; consequently, there are well-established platforms which are transparent in this area, such as silicon-on-insulator (SOI) and tertiary and quaternary III-V platforms based on InP and GaAs. However, a range of emerging applications, such as optical atomic clocks [43] and waveguide-based sensors [23] would benefit from low-loss waveguides in the visible (400 nm to 700 nm) and UV ($<$400 nm). Mature IR platforms have bandgaps well above these wavelength ranges—alternative material plaforms must be considered.

Table 1.1 shows optical properties of material platforms for the visible region. Considering bandgap alone, there are a large number of candidate materials. However, a scalable material platform relies on the ability to consistently produce high-quality films, and to define high-quality (i.e., low-roughness) waveguides. Whilst there are a relatively large number of wide-bandgap materials, there are fewer in which optical-quality films can be grown, and fewer still in which high-quality films can be patterned to produce low-loss waveguides (Figure 1.6 reviews waveguide loss results in the UV and visible). The criteria used in Table 1.1 to define high-quality films and waveguides—losses below 2 dB/cm in a deposited slab and in an etched waveguide, respectively—provide only a crude measure of a material platform's maturity: even at the same propagation wavelength, design and fabrication choices have a considerable effect on waveguide losses and use cases.

When comparing waveguides and material platforms, it is often useful to consider a waveguide's *confinement factor*, the percentage of the propagating optical power which travels within the waveguide core (as opposed to the cladding). The confinement factor is related to the propagation wavelength, the refractive index of the material, and the waveguide's cross-sectional geometry. High-confinement structures maximise nonlinearity [44], can achieve low losses using smaller bending radii [45], and can produce more efficient grating couplers [46, 47, 48]. However, the higher power density within the waveguide core places higher demands on the material quality (absorption and bulk scattering loss). Increasing waveguide

Table 1.1: Bandgap wavelengths for common optical materials. The criterion for slab or waveguide loss is a published result demonstrating loss at any wavelength $\leq$700 nm. Values are indicative; manufacturing variations may yield different results. Bandgap wavelengths may refer to the material's optical or electronic bandgap; where a source specifies which, [o] denotes an optical bandgap measurement and [e] an electronic bandgap measurement.

Material	Bandgap wavelength (nm)	n at 633 nm	n_2 (m^2/W)	Slab loss <2 dB/cm	WG loss <2 dB/cm	Sources
Si	1107	3.88	4.5×10^{-18} at 1550 nm	No	No	[49, 50, 51]
GaN	359	2.38	2.5×10^{-18} at 724 nm	No	No	[52, 53, 54]
TiO_2	353 [e]	2.13	9.0×10^{-19} at 1000 nm	Yes	No	[55, 56, 57, 58]
Ta_2O_5	281	1.83	6.2×10^{-19} at unspecified	Yes	No	[59, 60, 61]
$LiNbO_3$	264 [o]	2.29	2.3×10^{-20} at 800 nm	Yes	Yes	[62, 63, 64, 65]
SiN	250 [o]	2.18	2.4×10^{-19} at 1558 nm	Yes	Yes	[66, 67, 68, 69]
Diamond	225	2.41	1.4×10^{-19} at 515 nm	No	No	[52, 70, 71]
AlN	200 [o]	2.16	1.8×10^{-17} at 532 nm	Yes	No	[67, 72, 73]
am-Al_2O_3	200	1.64	4.8×10^{-20} at 958 nm	Yes	Yes	[74, 75, 76, 77]
α-Al_2O_3	150	1.77	3.3×10^{-20} at 550 nm	No	No	[52, 78]
Hydex (SiO_2)	140	1.6-1.7	1.2×10^{-19} at 1560 nm			[79, 80, 81]

confinement also typically requires depositing and patterning thicker films, which can make film deposition and waveguide definition processes more challenging.

SiN is by far the most mature visible waveguiding platform, with multiple well-established commercial offerings such as those from Ligentec [66] and LioniX [82]. There are well-established methods of depositing films and sophisticated methods of minimising waveguide roughness, both through etching [83] and etchless (Damascene) processes [84]. The film composition can be altered to alter properties such as film stress, nonlinearity, and refractive index [85]. However, despite a measured bandgap of 250 nm [67], SiN has limited utility below 500 nm [86]: recently, Corato-Zanarella et al. confirmed that this is because the Urbach tail extends through the visible region such that intrinsic material absorption is about 1 dB/cm at 600 nm, rising to about 2.5 dB/cm at 450 nm [30]. The shortest wavelength at which SiN losses have been reported is 405 nm; Morin et al. [87] report losses of 0.93 dB/cm at this wavelength. However, these waveguides were 24 nm thick, indicating that the majority of the mode was travelling in the wide-bandgap silica cladding rather than the core (see Section 4.1), which limits the effect of material absorption. Such waveguides are of limited utility, as the diffuse mode will require a extremely large bending radius (millimetre-scale, in this paper). Smith et al., using a more realistic 150 nm-thick SiN platform from commercial supplier Ligentec, report sub-dB/cm losses at 635 nm, rising to 6.85 dB/cm at 450 nm, using 150 µm bends [66].

SiN provides a mature material platform for guiding red light. Whilst other materials may have lower losses in the red, SiN's maturity, demonstrated by the capacity to consistently deposit high-quality films and create high-quality waveguides, and the material's relatively high refractive index make it an attractive platform for longer wavelengths. However, materials

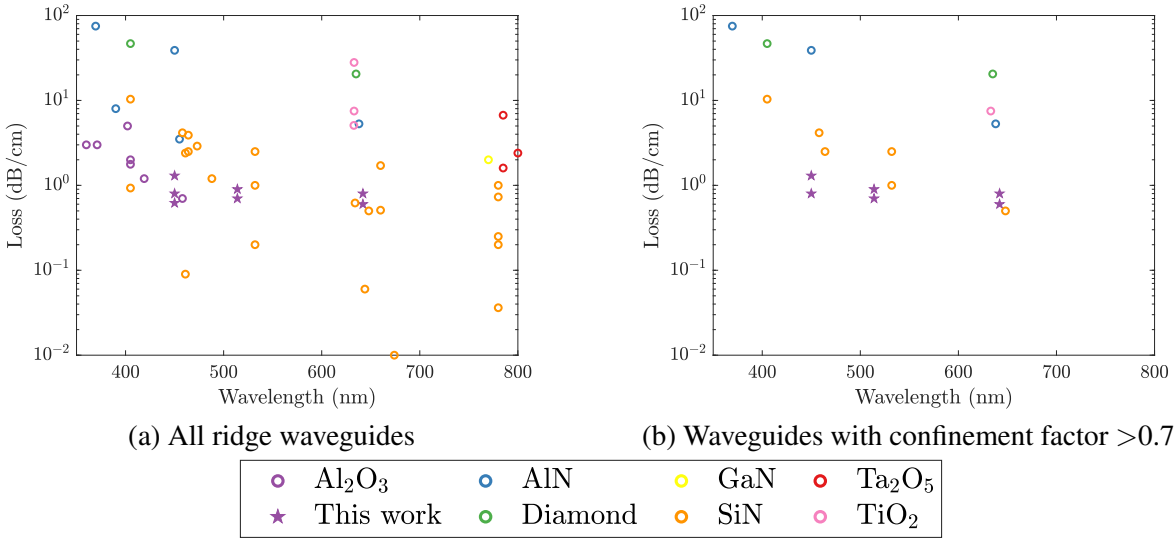

(a) All ridge waveguides (b) Waveguides with confinement factor >0.7

Figure 1.6: Review of propagation losses achieved in ridge waveguides at wavelengths below 800 nm. Results from this work demonstrate the lowest propagation losses in waveguides with confinement factor > 0.7 at 514 nm and below. Full details are provided in Table A.1.

with a wider bandgap—or a shorter absorption tail—must be considered for waveguiding in the blue and green.

A number of alternative materials have suitable bandgaps. Perhaps the most well-established is lithium niobate ($LiNbO_3$); as a non-centrosymmetric crystal, it has a high second-order nonlinear susceptibility ($\chi^{(2)}$), and bulk periodically poled $LiNbO_3$ has been used extensively for second harmonic generation. Whilst the material is reported to have a usable range down to about 400 nm [88], nonlinear generation has produced guided light down to 330 nm [89]. Now commercially available in thin-film form, the high film quality, linear electro-optic effects, and nonlinear properties have encouraged development of low-loss $LiNbO_3$ waveguide platforms in the infrared [88, 90]. However, the material poses fabrication challenges, and is notably difficult to etch: demonstrations using reactive ion etch (RIE) tools found in typical fabrication facilities suffer from angled sidewalls, poor selectivity, and creation of involatile byproducts which are difficult to remove from the chip [91]. Desiatov et al. report losses of 0.06 dB/cm at 637 nm in partially-etched rib waveguides, but with an etch depth of only 180 nm; to date there are no examples of fully-etched $LiNbO_3$ waveguides for the visible region.

III-V materials such as aluminium nitride (AlN) and gallium nitride (GaN) appear promising because of their capacity to form tertiary or quaternary compounds which could be used to produce integrated light sources, especially in the blue and UV [92]. For example, AlGaN can be produced with arbitrary relative compositions of the Group 3 elements, producing materials with a bandgap which varies between those of the two binary compounds [67]. However, growth processes currently limit the usage of these materials in PICs: Stassen et

al. demonstrate that film defect density is a limiting factor for losses in pure GaN (losses do not scale with degree of sidewall interaction), whilst Li et al. illustrate that increasing Al concentration drives defect growth in epitaxial processes [93]. Pure AlN may be deposited through sputtering or ALD as well as epitaxial growth. Historically, the polycrystalline forms deposited using sputtering [94, 95] have reported higher losses than epitaxial growth using metal-organic chemical vapour deposition (MOCVD), although recent sputtering optimisation by Mardani et al. has produced AlN films with slab losses of 1.5 dB/cm at 633 nm [73]. To date, single-crystal AlN has produced favourable etched waveguide loss results, with demonstrations of 8 dB/cm in the UV [96]. Notably, this result was also attained using a ring resonator with a small radius of 30 µm—radiative losses could likely have been reduced by increasing ring size.

Single-crystal diamond films are commercially available, but are variable in quality [97], and there are few demonstrations of the technology for visible waveguiding. Producing waveguides requires aggressive plasma etching, and sidewall roughness is presently a limiting factor [98].

Whilst optimised sputter deposition of titania (TiO_2) has produced films with slab losses of 1.0 dB/cm at 633 nm [55], waveguide losses remain high, above 5 dB/cm. Considering the material's relatively narrow bandgap (353 nm [56]), even optimised TiO_2 waveguides are unlikely to be useful for the blue and UV regions.

Tantala (Ta_2O_5) has a high nonlinear refractive index and is therefore an emerging nonlinear optics platform for the infrared; given its wide bandgap and extremely low slab losses in the visible (0.3 dB/cm at 633 nm) and low-loss operation of high-confinement waveguides in the near-infrared [59], both nonlinear frequency generation and waveguiding in the visible region are likely feasible with this platform.

Alumina (Al_2O_3) is a more mature platform than many of these options—as shown in Figure 1.6a, alumina is presently the only waveguiding platform which has demonstrated losses as low as 5 dB/cm in the UV. Optical-quality films can be deposited using ALD, a highly repeatable deposition method, which allows fabricators to sidestep the labour-intensive process of film optimisation [74]. Hendriks et al., at the University of Twente, have also optimised alumina sputtering processes, which can produce sub-dB/cm slab losses in the UV [99]. Optical-quality etching poses a difficulty (further described in Section 1.4)—whilst losses below 1 dB/cm in the blue have been attained [4] in fully-etched alumina waveguides, all demonstrations of fully-etched alumina waveguides for the visible region (outwith this work) use extremely thin films (on the order of 100 nm). Alumina has a low refractive index, measured at 1.64 at 633 nm in this work: this minimises index contrast with cladding, which is favourable for minimisation of scattering loss [3], but renders thin alumina waveguides

extremely low-confinement; they may have a prohibitively large footprint for high-density applications.

Figure 1.6 provides an overview of loss results attained in fully etched (ridge) waveguides in the visible and near-UV regions. Results may be from either straight waveguides or micro-ring resonators. Whilst SiN waveguides with extremely low propagation losses are listed in Figure 1.6a, when the analysis is restricted to waveguides with a confinement factor over 0.7 (Figure 1.6b), it becomes clear that achieving low-loss operation in high-confinement systems remains a challenge. In SiN, a well-developed platform, waveguide losses track well with the absorption limits derived by Corato-Zanarella et al. [30]; in all other materials, losses remain substantially higher, limited by both film quality and waveguide definition quality. Using alumina, a material with a wider bandgap, and an optimised etch process, this work demonstrates the capacity to produce high-confinement waveguides with loss values substantially below the SiN absorption limit.

1.4 Alumina Waveguides and PICs

Low propagation loss is a prerequisite for a UV-visible PIC platform; key requirements are a waveguide material with a wide bandgap, methods which allow consistent deposition of pure and homogenous films, and fabrication methods which allow components to be fabricated to specification and with low sidewall roughness. Ideally, fabrication processes are repeatable and use existing semiconductor manufacturing infrastructure. The ability to include active functions, such as light sources, amplifiers, and frequency conversion, is desirable; however, the majority of material platforms are likely to require integration of sources in a different material.

Whilst there is a substantial body of literature exploring the deposition of alumina films, the links between deposition parameters, structure, and optical properties remain unclear. Regardless, there are several deposition methods capable of producing high-quality amorphous alumina films, which have enabled demonstrations of alumina waveguides with low losses into the UV. However, the definition of high-confinement waveguides with the optical quality required for visible waveguiding remains a challenge.

Alumina polymorphs

Alumina, stoichiometrically represented as Al_2O_3, has a large number of polymorphs. α-Al_2O_3, known as *corundum*, is the most thermodynamically stable form, a hexagonal crystal structure in which aluminium atoms are 6-coordinate ($^{[6]}$Al). However, there are numerous long-lived metastable states, the most common being cubic-packed γ-Al_2O_3, in which aluminium atoms are primarily 4-coordinate ($^{[4]}$Al) [100]. Thin-film deposition methods such as

atomic layer deposition (ALD) and sputtering typically produce predominantly amorphous alumina films [101, 102], in which aluminium centres display a range of coordination numbers, mostly $^{[4]}$Al and $^{[5]}$Al [103, 104]. In addition, amorphous alumina may be non-stoichiometric [105].

Alumina's optical properties vary depending on its bonding composition. Consider the properties of the most and least ordered forms: sapphire (single-crystal corundum) and amorphous alumina. Sapphire has a higher refractive index than metastable forms of alumina [106]—at 633 nm, about 1.77 [52] compared to 1.64 for the amorphous alumina used in this project. Its reported bandgap of 150 nm is wider than that of the amorphous form, which appears to have a bandgap in the region of 200 nm [74, 75]. This difference in bandgap is not clearly understood. In 2006, Momida et al. presented an *ab initio* analysis suggesting that the presence of 4-coordinate aluminium reduces the material bandgap [107]; more recent work pinpoints alumina coordination as a driver of bandgap changes, affecting the position of the bottom of the conduction band [108]. Oxygen vacancy defects narrow the bandgap [109] and can produce absorption peaks in the near-UV [110]. Suitability for PIC applications is also decided by the amount of absorption in the Urbach tail region just below the bandgap; whilst this area has not been modelled in alumina, trends observed in other material systems suggest that the amorphous form may have greater extension of the Urbach tail, as the presence of a variety of bonding coordination within the material is likely to produce a wider range of feasible absorptions [31, 111].

Alumina Deposition Techniques

Whilst sapphire has excellent optical properties and can be produced as high-quality single crystals, it lacks a thin-film form, which prohibits its use in integrated optics. Instead, thin-film alumina is deposited top-down, using atomic layer deposition (ALD) or sputter deposition, and is typically amorphous.

An unpatterned film's suitability for optical waveguiding can be evaluated using two techniques. The material bandgap is measured through absorption or transmission, which does not provide any information about scattering losses. Slab loss tests measure propagation loss through an unpatterned film deposited on a lower cladding layer. These measurements combine losses from absorption, bulk scattering, and interface scattering at the top and bottom of the slab. It is typically assumed that the majority of losses will result from material quality, rather than from scattering from the top and bottom of the waveguide, although for rough films and at short wavelengths, topside scattering can have a considerable effect [99]. Table 1.2 summarises reported slab losses in alumina films.

Sputtering describes a range of physical vapour deposition (PVD) techniques in which plasma species—typically an inert gas such as Ar—are accelerated towards a target of the material to be deposited. Bombardment of the target releases particles of the material, which are deposited on the substrate. Alumina sputtering uses an aluminium target, and oxygen is added to the chamber in order to oxidise the metal. The first recorded use of sputtered alumina films in optics is in a demonstration of second harmonic generation by Chen, Tang and Telle in 1974 [112]; further development of sputtered alumina films for integrated optics applications first took place in the 1980s [113, 114, 115]. These early papers demonstrate the critical role of deposition parameters in defining optical characteristics such as refractive index and material loss: for example, Smit et al. [114] reported slab losses for as-deposited films ranging from 2 dB/cm to 40 dB/cm, whilst Este and Westwood reported slab losses from 0.5 dB/cm to 100 dB/cm [113]. However, no consistent reason for the high variability was established, and there were few further developments in alumina film deposition throughout the 1980s and 1990s. This lack of progress likely arises from a lack of applications for alumina waveguides. Widespread interest in UV-visible optics is relatively recent [116]; the telecoms market, which historically drove PIC development [41], has established standards using frequencies in the infrared area of the spectrum [42], where other low-loss waveguide materials (Si, SiN) already excel.

In recent years, extensive research at the University of Twente has yielded an ability to consistently deposit low-loss alumina films using radio frequency (RF) sputtering [99, 101, 117]: the film structure and optical properties are strongly dependent on the oxidation state of the sputtering target, with a higher oxidation state producing polycrystalline γ-alumina. Whilst the group has produced polycrystalline films with a higher refractive index (1.72 at 633 nm) and relatively low loss (1.8 dB/cm at 407 nm) [118], the lowest losses appear to arise from providing sufficient oxygen to produce amorphous alumina without facilitating nucleation of the crystalline phase [99].

Alternatively, alumina can be deposited using ALD, a process which builds up a film from sequential half-reactions using an organometallic precursor and an oxidant [119]. This process, in which films are deposited a sub-monolayer at a time and each half-reaction is driven to completion, inherently produces consistent results—namely, conformal films of even density, minimising bulk scattering losses, and with high purity, minimising extrinsic absorption. Most commonly, the oxidation is performed using water vapour, a process known as *thermal ALD*; however, the oxidation can also be performed using an oxygen or ozone plasma, known as *plasma-enhanced atomic layer deposition* (PEALD). There is only one study investigating slab loss in amorphous alumina deposited using PEALD: Caballero-Espitia et al. demonstrate lower losses using ozone as their oxidant instead of water vapour, and attribute the improvement in loss to a reduction in hydroxyl groups remaining in the bulk

Table 1.2: Review of as-deposited slab losses in alumina.

Authors	Institution	Year	Deposition method	Loss (dB/cm)	λ (nm)
Este/Westwood [113]	Bell, Ottawa	1984	DC magnetron sputter	0.5-100	633
Smit/van der Laan [114]	Delft	1986	RF sputter	2-40	633
Arnold/Cole [115]	Honeywell	1988	Ion beam sputter	0.4-2.9	633
Stadler/Oliveria [75]	MIT	1995	RF sputter	1	633
Bradley/Pollnau [122]	Twente	2007	RF sputter	0.29	633
León/Kobayashi [123]	CSEM	2019	DC magnetron sputter	1.6-7.7	633
Hendriks/García-Blanco [117]	Twente	2023	RF sputter	0.5	371
Kumar/Avrutsky [124]	Wayne State	2009	Thermal ALD	7	633
Aslan/Mendes [74]	Louisville	2010	Thermal ALD	<4 1.1	250 633
West/Sorace-Agaskar [4]	MIT	2019	Thermal ALD	0.3	405
Caballero-Espitia/ Marquez [120]	UNAM	2020	Thermal ALD Ozone plasma ALD	3.03 0.51	633

of the film [120]. Shestaeva et al. report the onset of absorption at 230 nm in alumina films deposited using dioxygen PEALD, but no other optical loss results using this process have been published.

Whilst sputtering and ALD can produce relatively similar losses, the nature of ALD processes—in which films are deposited a sub-monolayer at a time through reactions with strongly favourable thermodynamics—should produce replicable results, even across tools of different models [121]. Achieving consistency in sputtered alumina depositions is a recent phenomenon, and maintaining consistent film quality requires retuning of deposition parameters. However, ALD processes are extremely slow compared to sputtering, which is likely to contribute to the lack of high-confinement ALD alumina waveguides reported in the literature.

Etched Ridge Waveguides

Etching alumina presents a significant fabrication challenge. Whilst the material readily dry etches in halides, it has relatively low selectivity to most mask materials [122], which limits etch depth and quality.

The first dry-etched alumina ridge waveguides were produced by Bradley et al. in 2007 [122]. The process compensated for poor etch selectivity by using a thick photoresist etch mask, and has been used to etch high-confinement waveguides more than 1 µm thick [125]. The waveguides produced had losses of 0.21 dB/cm at 1550 nm, and has since been used to create a range of active components in rare earth-doped alumina, including amplifiers at 880 nm, 1060 nm, 1330 nm and 1550 nm [126, 127, 128] and lasers at 1550 nm [129]—valuable additions to platforms operating at telecoms wavelengths [130]. However, no results in the visible region have been reported using this fabrication method; scanning electron

micrographs [122] show levels of sidewall roughness which would likely produce excessively high losses.

Other groups have employed a different strategy to produce UV-visible waveguides, defining an intermediate *hardmask* in silica or SiN to reduce sidewall erosion during the aggressive alumina etch [4, 13]. Extremely low losses, even in the UV, have been reported using this method (Table 1.3); however, these loss results have been measured in thin, low-confinement waveguides. Since minimising the core thickness can minimise losses, it makes sense to follow this strategy where possible, and this approach has produced waveguides which were successfully used in practical applications such as blue light sensing [131] and autofluorescence microscopy [13].

However, given the extremely low refractive index of amorphous alumina, low-confinement devices are likely to have an extremely large footprint. Additionally, high-confinement waveguides are required to leverage the material's low UV-visible loss for nonlinear optics applications. There are two barriers to production of low-loss high-confinement alumina waveguides: deposition and waveguide definition. Confining more light within the waveguide core places higher demands on material quality, requiring minimisation of absorption loss and bulk scattering. However, reported slab losses indicate that high-quality alumina deposition processes exist, using both sputtering and ALD. Creating thick sputtered films is trivial, but process control remains challenging. ALD can produce consistent, high-purity conformal films, but thermal ALD processes are extremely slow to run and PEALD processes have almost never been studied with a view to optical applications. Furthermore, waveguide definition processes are more challenging in thick films; the strategy used by Bradley et al. produced waveguides which were too rough to produce loss results in the visible region [122]. The hardmask processes used to develop low-confinement alumina waveguides for the visible region may also be applicable to thicker waveguides; however, etch selectivity and critical dimension control pose a greater challenge.

Table 1.3: Review of UV-visible propagation losses in alumina ridge waveguides.

Authors	Institution	Deposition method	Year	Loss (dB/cm)	λ (nm)	Core thickness (nm)	Core width (nm)	Cladding
West/ Sorace-Agaskar [4]	MIT	Thermal ALD	2019	3 1.8 1.2 0.7	371 405 419 458	100	400-600	SiO_2
Lin/ le Thomas [13, 132]	Ghent	Thermal ALD	2021 2022	5 3	402 360	120 70	700 600	None SiO_2
Hendriks/ García-Blanco [117]	Twente/ Aluvia	Sputter	2023	2	405	70	600	SiO_2

1.5 UV-Visible Supercontinuum Generation

The current state-of-the-art in alumina waveguiding is low-loss (sub-dB/cm) propagation in the blue and UV, but only in low-confinement waveguides [4]. A considerable portion of the motivation to work in the UV-visible range arises from applications which require broadband light, such as for stabilised frequency combs for optical atomic clocks and spectroscopy (Section 1.1). Broadband light can be achieved through supercontinuum generation (SCG), the conversion of narrow-linewidth pump light into broad-bandwidth, spatially coherent light through a range of nonlinear processes such as harmonic generation, soliton fission, and self-phase modulation (SPM) [44].

Nonlinear optics

Nonlinear optical effects only occur at extremely high light intensity (irradiance). Light propagating with an electric field $\vec{E}$ induces a response in polarisation density $\vec{P}$ [35]:

$$\vec{P} = \epsilon_0 \sum_{n=1}^{\infty} \chi^{(n)} \vec{E}^n \tag{1.5}$$

where ϵ_0 is the vacuum permittivity. The $\chi^{(n)}$ terms are susceptibility terms defining the proportionality of the response. Higher-order χ terms are extremely small, so at lower light intensities, we may consider solely the $\chi^{(1)}$ term, which is a dimensionless constant. The linear optical response is therefore defined by

$$\vec{P} = \epsilon_0 \chi^{(1)} \vec{E} \tag{1.6}$$

However, at extremely high $\vec{E}$ intensity, higher-order (nonlinear) susceptibilty terms become relevant:

$$\vec{P} = \epsilon_0 [\chi^{(1)} \vec{E} + \chi^{(2)} \vec{E}^2 + \chi^{(3)} \vec{E}^3 + ...] \tag{1.7}$$

This nonlinear polarisation response gives rise to a range of effects, most saliently frequency conversion. Nonlinear optical phenomena have been extensively studied in bulk media and optical fibres [35, 133], but PIC-based nonlinear optics offer advantages over these systems (aside from small form factor). The small cross-sectional area of a waveguide allows light to be confined to a very small area, maximising irradiance and therefore the extent to which nonlinear effects can be stimulated given a specific pump power. In this context, having high-confinement waveguides is particularly important, since they minimise modal area, which increases irradiance and therefore increases nonlinear response [44]. Furthermore, waveguide geometries have far more degrees of freedom than fibres, and altering the waveguide design can affect the type of nonlinear effect which takes place [134]. This design process is known as *dispersion engineering*.

Ultrashort pulsed sources are usually used for nonlinear optics, because they can achieve an extremely high peak pulse power [135]. Refractive index is wavelength-dependent (*chromatic dispersion*); consequently, phase and group velocity are wavelength-dependent. If group velocity varies strongly with wavelength, then pulse broadening will have an adverse effect on parametric phase matching and pulse shape, and therefore nonlinear conversion efficiency [44]. The wavelength-dependence of group velocity, known as *group velocity dispersion* (GVD), therefore has a strong influence on nonlinear broadening processes.

Chromatic dispersion is determined by material dispersion, which is typically normal (i.e., refractive index increases as wavelength decreases), and by the waveguide geometry, which introduces a phase change during propagation. Dispersion engineering typically aims to define a waveguide cross-section which offsets the effects of material dispersion to produce a GVD near zero at the pump wavelength (Figure 1.7) [136]. Typically, waveguides must be relatively thick to achieve this—the amorphous alumina used in this work requires a film thickness of 800 nm to achieve a zero GVD (Section 5.2). We can also consider the fact that since broadening begins from the pump wavelength, a shorter-wavelength pump may allow generation of short-wavelength components. However, material dispersion becomes strongly normal towards the material bandgap (for example, see Figure 2.3), so in order to use a short-wavelength pump, a wide-bandgap waveguide material is required.

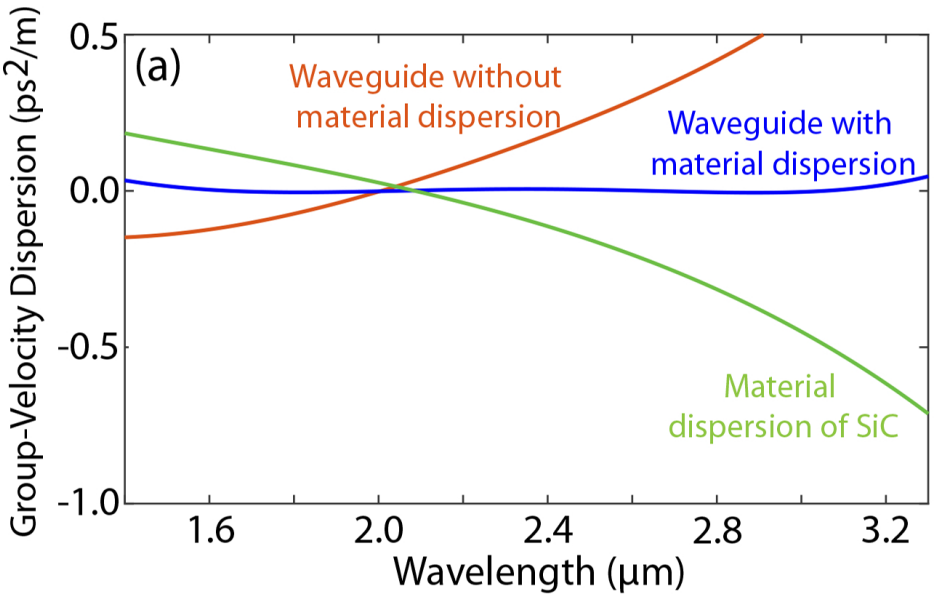

Figure 1.7: Dispersion engineering in an SiC waveguide. Waveguide dispersion (red) offsets material dispersion (green) to produce a flattened overall dispersion (blue). Adapted from Liang et al. [136].

Supercontinuum generation

SCG in optical fibres already finds a range of uses in spectroscopy, optical timing, and metrology [133]. Integrated photonic alternatives could substantially reduce the footprint and complexity of these devices. Supercontinua are typically described in terms of the wavelength range they cover without dropping below a specified attenuation level, measured in dB relative to the pump power. This implies that a favourable result is one which covers a wide range of wavelengths at a low level of attenuation. In fact, this is not strictly necessary for all applications: for example, access to frequency doubling is most desirable for optical clock applications [43]. Access to wavelengths in the blue and UV are desirable for spectroscopy and sensing applications, particularly in the life sciences [21, 22, 24].

Producing waveguides for supercontinuum generation poses a number of fabrication challenges. Dispersion-engineered waveguides are likely to be high-confinement, and potentially without an upper cladding layer (see Appendix A), conditions which make it challenging to achieve low propagation losses. Low-loss operation is not a prerequisite for *initiation* of nonlinear frequency conversion; however, nonlinear processes are intensity-dependent. Therefore, low losses increase the propagation distance over which nonlinear processes can occur. Additionally, light does not just need to be generated, it needs to be guided. For example, Hammani et al. produced unclad TiO_2 waveguides which generated blue light visible to the naked eye—but propagation losses were so high that no visible light was sensed at the detector [137]. The reliance on waveguide geometry to define dispersion characteristics may also place tighter tolerances on feature size control during fabrication.

Materials' third-order nonlinear response ($\chi^{(3)}$) is often quantified using the *nonlinear refractive index* n_2. (n_2 describes the magnitude of the *optical Kerr effect*, a third-order nonlinear effect in which refractive index changes with optical intensity.) Non-centrosymmetric crystals additionally possess $\chi^{(2)}$ nonlinearity, a second-order nonlinear response which can produce effects such as second harmonic generation (SHG) [44].

Figure 1.8 reviews UV and visible supercontinua achieved in optical waveguides at a level of at least −60 dB, whilst Table 1.1 lists n_2 for candidate materials for UV-visible supercontinuum generation. What produces good results—where, here, *good* means *generating broadband and/or short-wavelength light*, is complex: n_2, material bandgap, waveguide quality, and dispersion engineering all play a role.

SiN has a relatively high n_2 (2.4×10^{-19} m^2/W at 1558 nm [69]), which can be altered by altering the stoichiometry of the material. Si-rich SiN can produce the highest nonlinearities, but also has a narrower bandgap than stoichiometric SiN [138]. Exceptionally broad supercontinua achieved in SiN waveguides, extending well into the blue [139, 140], demonstrate what can be achieved with good fabrication processes, even given the deleterious effects of material

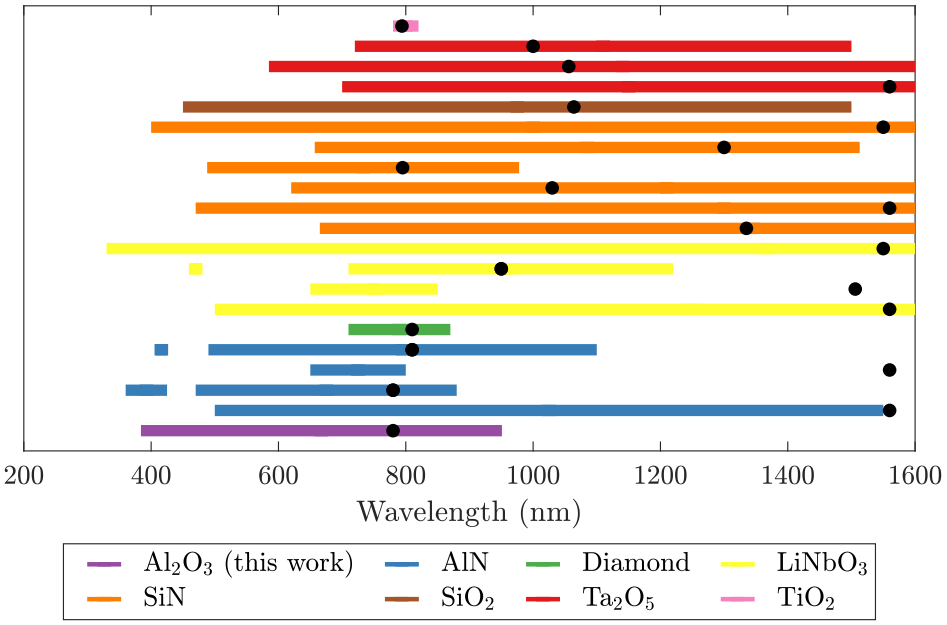

Figure 1.8: Review of supercontinuum generation within optical waveguides. In order to be included in this review, waveguides must have generated light below $\lambda = 800$ nm at a level of at least −60 dB. Black markers indicate pump wavelengths. Full details are provided in Table A.2.

absorption. Innovative dispersion engineering strategies can also offset losses: for example, Kou et al. produced a broad supercontinuum in an SiN waveguide (as low as 450 nm at a level of −50 dB) by laterally tapering the waveguide to stimulate coupling of fundamental and higher-order modes [139]. Nevertheless, material absorption is likely to limit the wavelength range attainable in SiN; alternative materials may be required for nonlinear broadening below 400 nm.

Ta_2O_5 demonstrates a considerably higher n_2 than SiN (6.2×10^{-19} m^2/W [59]), and extremely broad supercontinua have been demonstrated in Ta_2O_5 waveguides [141, 142]. However, the material platform lacks maturity, and low-loss linear operation at shorter than red wavelengths has not yet been demonstrated. UV generation has been reported as a result of SCG in AlN [143, 144]. $LiNbO_3$, another material with second-order nonlinear susceptibility, is a promising candidate for short-wavelength generation. Periodic poling of the material, which allows extended phase-matching for SHG, can be used to generate extremely broad spectra: a recent demonstration exploits this characteristic to produce a supercontinuum extending to 330 nm [89].

It is also worth noting that a high n_2 is not *required* for nonlinear generation. Broad supercontinua and tunable UV generation have been reported by Yoon Oh et al. using a doped silica waveguide platform [145]. High fabrication quality and relatively long waveguide lengths compensate for the material's low n_2 (1.2×10^{-19} m^2/W at 1560 nm [81]), although the input

pump power used is an order of magnitude higher than that used in leading SiN and AlN examples [140, 143].

Alumina in Nonlinear Optics

Despite the fact that the first mention of alumina films' use in optics is in a 1974 paper on SHG [112], there are no demonstrations of nonlinear optics in etched alumina waveguides. Four-wave mixing (FWM) in the IR has been demonstrated in alumina trench resonators co-integrated with SiN bus waveguides [146]: in this instance, a trench was etched into silica cladding and backfilled with sputtered alumina.

Amorphous alumina's n_2 has been measured as 4.8×10^{-20} m^2/W at 958 nm [77], an order of magnitude lower than that of SiN. As described in Section 1.4, deposition methods can have a considerable impact on the optical properties of amorphous alumina, and this value was derived from nonlinear measurements taken in aluminosilicate fibres with varying alumina content. Dragic et al. repeated their analysis using ellipsometric data from ALD films deposited by Kumar et al. [124] and arrived at a similar value. Sapphire's n_2—3.3×10^{-20} m^2/W at 550 nm [78]—is relatively similar, suggesting that the order of magnitude is known. Regardless, amorphous alumina's wide bandgap may offer the potential for anomalous dispersion conditions at relatively short wavelengths.

1.6 Aims and Objectives

Amorphous alumina has a wide bandgap and is therefore a promising material platform for UV and visible PICs. Technology demonstrators using low-confinement waveguides with low losses in the blue and UV deliver on this promise. However, high-confinement waveguides for the visible region have not yet been developed, limiting the utility of amorphous alumina waveguides for UV-visible nonlinear optics applications. Outstanding fabrication challenges include consistent deposition of thick optical-quality layers, and development of masking and etching processes for thick waveguides which are also smooth enough to produce acceptable scattering losses.

This project's overall aim was to establish an ALD alumina material platform for UV and visible linear and nonlinear optics. Objectives were:

- Deposit amorphous alumina and establish that the material's properties make it suitable for use in UV-visible waveguiding (Chapter 2)
- Develop a waveguide fabrication process which can be used to produce high-confinement amorphous alumina ridge waveguides with low losses in the visible and UV (Chapter 3)
- Design and fabricate alumina ridge waveguides which are suitable for low-loss single-mode propagation of visible light (Chapter 4)
- Design dispersion-engineered alumina waveguides which can generate broadband visible and UV light through nonlinear optical processes (Chapter 5)

Chapter 2

Atomic Layer Deposition of Optical-Quality Alumina

Amorphous alumina is an established wide-bandgap integrated optics platform, but a lack of high-confinement waveguides with low propagation losses has hindered its use in non-linear optics applications. Whilst ALD can be used to produce optical-quality amorphous alumina films, conventional thermal ALD processes, which use water vapour to oxidise an organometallic precursor, are too slow for deposition of thick waveguides. In this work, I use a dioxygen PEALD process, in which an oxygen plasma is used to perform the oxidation step, to reduce deposition time by 60%. This deposition process is relatively well-established, but no prior work has used PEALD alumina films in an integrated optics context. This chapter provides analysis of the properties of unpatterned dioxygen PEALD films with a view to their use in integrated optics.

As discussed in Section 1.2, an ideal dielectric material for waveguiding in the UV and visible region has

- A high bandgap energy—ideally far higher than the energy of the light to be guided, to avoid the absorption 'tail' caused by uncertainties in ionic energies
- A very low concentration of impurities, to minimise extrinsic absorption loss
- A homogenous microstructure free of voids, pinholes, and refractive index inhomogeneities, to minimise bulk scattering loss
- Low topside surface roughness after deposition, to minimise interface roughness

Successful experiments in deposition of amorphous alumina have demonstrated bandgaps in the range of 200 nm to 250 nm [74, 75]. Numerous groups have reported low surface roughness [105, 147], and atomic layer deposition (ALD) is known for its inherent tendency

to produce films which are pure, homogenous, conformal, and smooth [119, 121]. Amorphous alumina slab and ridge waveguide demonstrations indicate low losses in the UV and visible, indicating acceptable homogeneity and film purity [4, 74]. However, different deposition parameters can produce amorphous films with a wide range of optical properties (Section 1.4). The majority of work linking deposition parameters to optical properties, especially properties such as optical loss and optical bandgap, have focused on sputtered alumina; ALD alumina's application in waveguiding is secondary to its uses in catalysis [148], high-κ dielectrics [149], and optical coatings [150]. Nevertheless, it is clear that different ALD deposition parameters produce different film densities, stoichiometries, and optical properties [151, 152, 153]. Film reproducibility is, in theory, inherent to this deposition technique. In practice, whilst some groups have demonstrated high reproducibility, even across different deposition tools [121, 154] and substrate profiles [155], this cannot be taken for granted: Stønsteby et al. compare outputs from papers using similar deposition parameters, and highlight considerable disparities in growth per cycle (GPC) [151].

The ALD processes used in this project were originally developed for power electronics applications by Dilini Tania Hemakumara [149]. Prior to this work, the only optical characterisation was measurements of refractive index, used solely to measure deposition uniformity. However, the high breakdown voltage and lack of pinholes within the material suggested that the material may have favourable optical properties. Given the favourable optical loss results obtained with Hemakumara's recipes, I did not further optimise the deposition process. Consequently, the analysis presented is limited in depth: lacking a series of films deposited under different conditions, it is difficult to establish a strong link between deposition parameters, structural properties, and optical properties. However, this analysis remains necessary to characterise the optical quality of PEALD films deposited within the James Watt Nanofabrication Centre (JWNC).

I used a variety of material characterisation techniques to evaluate films' suitability for optical waveguiding. I used variable angle spectroscopic ellipsometry (VASE) to determine the optical constants of films (Section 2.2), directly measured film bandgap using UV-visible spectrometry (Section 2.3), used Fourier transform infrared (FTIR) spectroscopy to probe the bonding composition of deposited films and provide information on impurities (Section 2.4), and used atomic force microscopy to determine the surface roughness of deposited films (Section 2.5). I consider run-to-run reproducibility, although this dataset is limited: even using PEALD to reduce cycle time and increase GPC, ALD depositions are time- and resource-intensive, so data has been collected from different substrates and film thicknesses.

This chapter focuses solely on analysis of deposited alumina films; a process flow for production of optical devices is discussed in Chapter 3. However, it is critical to consider that processing may alter films and affect their capacity for low-loss waveguiding. Device

fabrication processes may involve subjecting materials to heat, liquid reagents, and plasma processing. Consequently, Sections 2.7 and 2.8 explore the effects of chemical reagents and temperature processing on PEALD alumina films.

2.1 Deposition

ALD is a form of chemical vapour deposition (CVD) process which is performed as sequential half-reactions, as illustrated in Figure 2.1. The substrate is placed in a reaction chamber, which is filled with a gaseous precursor of the metallic element of the film. The precursor *chemisorbs* onto the surface of the substrate. During chemisorption, the adsorbate chemically bonds to the surface and cannot readily be removed. The metal precursor is introduced to the chamber until the surface is saturated, then the reaction chamber is purged and the precursor of the non-metallic component of the film is introduced. The non-metallic precursor forms bonds with the adsorbed metal precursor, forming a layer of the target compound. Once enough time has passed to allow all of the available surface to react, the chamber is purged again, and the two half-reactions are alternated until a sufficiently thick layer of the target film has been deposited [119]. In a conventional ALD process, each half-reaction is thermodynamically driven to completion; therefore, as long as sufficient time is allowed for each step, the films produced should be high-purity and free of voids. Each cycle deposits no more than a monolayer of reaction product, so the film thickness is highly controllable.

ALD processes for depositing alumina typically use trimethylaluminium (TMA) ($Al(CH_3)_3$) as the metal precursor [119]. Oxidation, and displacement of the methyl groups, can be achieved by one of two routes. Conventional *thermal* ALD processes use water vapour to perform the oxidation, and heat to drive the reaction. Prior work on optical-quality alumina films has almost exclusively used thermal ALD processes [4, 74, 132, 156]. However, water vapour is notoriously difficult to purge from reaction chambers [154]; whilst the repeated purge steps required to perform ALD necessarily produce long deposition times, the use of water for the oxidation step results in exceptionally long purge and deposition times. It has also been suggested that different research groups report substantial variations in GPC as a result of poor or incomplete water purging [151]. As an alternative to water, an oxidant plasma—oxygen (O_2) or ozone (O_3)—can be used for the oxidation half-reaction.

The oxidation mechanisms within PEALD alumina processes are not yet well understood. *Ab initio* density functional theory (DFT) simulations and in-situ FTIR studies agree that for both thermal and plasma processes, TMA adsorbs onto the substrate, and the aluminium centre accepts a pair of electrons from a terminal oxygen atom. Each aluminium centre may bond to one or more surface oxygen atoms, stabilising when it reaches a tetrahedral coordination [157, 158]. The mechanism of the oxidation half-reaction varies depending on the oxidant.

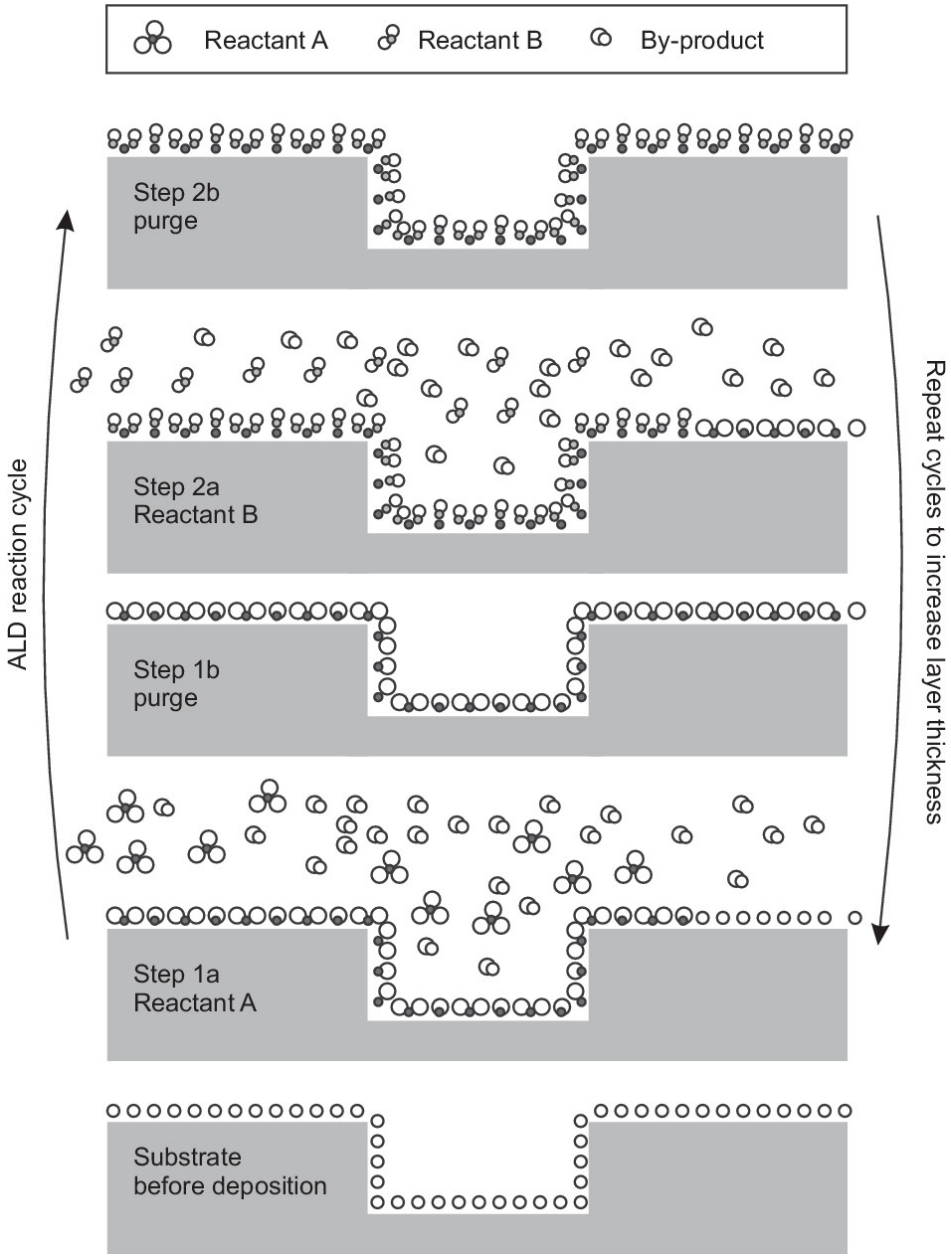

Figure 2.1: Schematic illustration of one ALD reaction cycle. Reprinted from Miikkulainen et al. [119], with the permission of AIP Publishing.

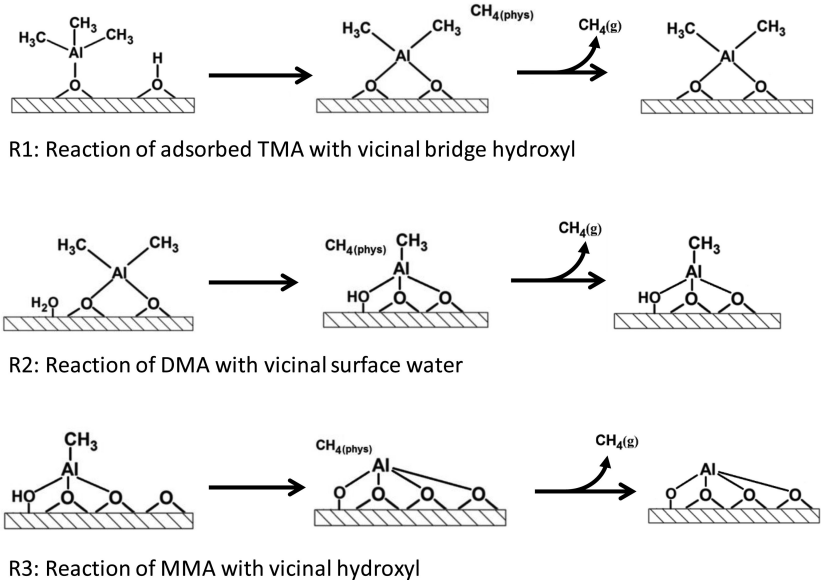

Figure 2.2: Schematic of selected surface reactions during thermal ALD process with TMA and water reactants. Adapted with permission from Sandupatla et al. [158]. ©2015 American Chemical Society.

Water vapour may bond with adsorbed aluminium centres, donating an electron pair, or it may hydrogen-bond to the surface, before undergoing a ligand exchange reaction in which methyl groups are removed as methane (CH_4) [158]. Both oxygen and ozone plasmas perform a combustion-type oxidation reaction; the use of oxygen plasma reduces the formation and lifetime of intermediate carbonate species [159, 160].

It is not possible to provide a single, clear reaction mechanism for either oxidation route: a variety of functional groups terminate the alumina surface after each half-reaction, producing a number of subsequent reactions [158, 159]. Figure 2.2 shows a selection of surface reactions which occur during thermal ALD processes—considering the variety of mechanisms within these sequences, it is clear how an amorphous structure, with varying aluminium coordination numbers, arises during the deposition process.

The link between deposition precursors and parameters and the optical quality of deposited films is not well understood. A near-edge X-ray absorption fine structure spectroscopy (NEXAFS) and X-ray photoelectron spectroscopy (XPS) study has demonstrated that aluminium coordination number plays a substantial role in defining the bandgap—although a specific ratio remains unspecified, as does the role of ALD parameters in producing any given aluminium coordination [108]. Furthermore, incomplete oxidation of TMA centres may result in oxygen vacancy defects, which produce a narrower bandgap [109] and absorption peaks in the blue and near-UV [110].

There is a stronger evidence base for the role of ALD parameters in defining film density and purity. Limiting the scope of this analysis to ALD processes using an oxygen plasma,

Verlaan et al. varied temperature between 50 °C and 400 °C, and found that films deposited at higher temperatures contained fewer hydroxyl, carbonate and methyl groups. Rai et al. have proposed and experimentally determined a combustion-like reaction mechanism which supports this finding [157]. Whilst it is unclear whether these impurities absorb light at wavelengths of interest, films with low impurity levels are typically preferable for optical waveguiding. Verlaan et al. also noted an increase in film density with increasing deposition temperature, which is likely to indicate an increase in refractive index [152]. Findings by Shestaeva et al. support this: they report an increase in refractive index with increasing deposition temperature [153]. The relationship between density and optical loss is unclear, but assuming that crystallisation does not occur at ALD deposition temperatures, a denser film may be less porous and therefore experience lower bulk scattering loss. In this work, I deposited films at 200 °C; because loss results from the initial deposition run were favourable, and there is little conclusive evidence linking recipe parameters with optical losses, I did not perform any further optimisation.

2.1.1 Film Deposition

All ALD processes were conducted using an Oxford Instruments FlexAL system. Thermal and plasma recipes were available, and are presented in Tables 2.1 and 2.2 for comparison; however, all deposition runs used the plasma process. Depositing a 400 nm-thick layer using the thermal ALD recipe, with TMA and water as reactants, would take 13.5 h. However, depositing the same thickness using an oxygen plasma for the oxidation half-reaction reduces this time to 5.2 h. Eliminating the need to purge water vapour from the chamber reduces purging time by 6 s, and reduces the time for a single cycle from 12.17 s to 6.02 s. The plasma process also has a high GPC, depositing 1.3 Å per cycle compared to 1.0 Å per cycle using the thermal process. For thermal ALD, higher GPC is frequently associated with less dense films [121], but the higher GPC of plasma processes—particularly when an oxygen plasma is used—originates from more efficient oxidation of the aluminium centres [157, 160, 161].

Table 2.1: Process parameters for thermal ALD of amorphous alumina at 200 °C using Oxford Instruments FlexAl system.

Step	Action	Time (s)	Other Parameters
1	TMA dose	0.02	
2	TMA purge	3	
3	H_2O dose	0.15	80 mT
4	H_2O purge	9	150 sccm, 80 mT
	Total	12.17	

Table 2.2: Process parameters for plasma ALD of amorphous alumina at 200 °C using Oxford Instruments FlexAl system.

Step	Action	Time (s)	Other Parameters
1	TMA dose	0.02	
2	TMA purge	2	
3	Plasma stabilisation	0.5	
4	O_2 plasma	2	60 sccm, 80 mT, 400 W
5	Post-plasma purge	1.5	100 sccm, 80 mT
	Total	6.02	

Table 2.3: Summary of alumina deposition runs performed during this project.

Date	Thickness (nm)	Substrates
January 2021	400	6” Si wafer with 5 µm SiO_2
May 2021	300	Si 4” Si wafer with 5 µm SiO_2 Si with 5 µm SiO_2 and 400 nm Al_2O_3
June 2023	100	Si 4” Si wafer with 5 µm SiO_2
September 2023	400	Si 4” Si wafer with 5 µm SiO_2 Fused SiO_2 slides

In total, four deposition runs were performed, as detailed in Table 2.3. Various thicknesses were deposited—300 nm and 400 nm-thick films for high-confinement waveguides and analysis work, and 100 nm-thick films for low-confinement waveguides. The films were deposited on different substrates to allow a range of applications and analysis. For waveguiding, a thick bottom oxide cladding layer is required; deposition runs included entire Si wafers with an optical-quality SiO_2 layer, which were cleaved into pieces and processed separately. Bare Si chips, suitable for ellipsometry and FTIR spectroscopy, were included in three runs. Fused silica slides, included in one deposition run, provide a wide-bandgap substrate, allowing transmission bandgap measurements to be performed. Three 20 mm × 20 mm chips cleaved from the January 2021 deposition run (400 nm of alumina on SiO_2, suitable for waveguiding) were included in the May 2021 run, depositing an additional 300 nm of alumina to produce 700 nm-thick films for nonlinear optics applications.

2.2 Ellipsometry and Optical Constants

Spectroscopic ellipsometry can be used to determine a film's optical constants (refractive index n and absorption coefficient k). Plane-polarised light is reflected from a film, producing

an elliptically polarised beam; the attenuation and phase shift can be fitted to a theoretical model to derive n and k, as well as film thickness. Multiple-angle ellipsometry allows derivation of a larger number of constants relating to the polarisation of the reflected light [162]. Information on chromatic dispersion, the variation of n with wavelength, is critical for optical component design (Section 4.2) and measurements over multiple films can provide information on deposition consistency. The variation of k with wavelength can provide information on the position of the optical bandgap. Note that there is no direct link between ellipsometric measurement and film quality: ellipsometry is a reflective technique, and does not directly provide information about bulk scattering.

I performed all measurements using a J. A. Woollam M-2000XI spectroscopic ellipsometer, over a wavelength range of 210 nm to 1700 nm and an angle range of 50° to 75°. I fitted data to a Tauc-Lorentz model, which includes terms to parameterise the material bandgap [163], achieving a typical mean squared error of 12. Only alumina deposited directly on Si can be measured using this technique; when alumina films are deposited on a thick bottom cladding layer, it is challenging to resolve interference oscillations, especially at the short wavelengths which are of interest [164].

Film thicknesses yielded by ellipsometric measurement are shown in Table 2.4. There is some variation in GPC across samples: this may indicate changes in the condition of the ALD chamber and/or variation between the deposited films' properties. However, substrate effects frequently produce a higher GPC during the first cycles of an ALD process [151]—further runs, ideally performed within a narrow timeframe, would be required to establish whether this variation is related to film thickness or process variability.

Figure 2.3 shows derived chromatic dispersion curves for the three deposition runs. There is no proven correlation between refractive index and optical loss, although alumina films with unusually low refractive index may be prone to porosity and therefore increased bulk scattering [75, 114]. Demirtaş et al. also noted that lower-refractive-index films deposited using ALD were less stoichiometric—they were lower in density because of a higher number of unreacted functional groups, and were therefore less pure and more likely to demonstrate extrinsic absorption [156]. Films with considerably higher refractive index (> 1.7) are likely to contain crystalline phases [117, 118]; however, crystalline phases are highly unlikely in ALD alumina films deposited using the ALD process described here. This process intrinsically produces an amorphous structure, and is performed considerably below the crystallisation temperature.

The *extinction coefficient* describes attenuation within the material, which rises as photon energy approaches the bandgap energy. An estimate of optical bandgap can be derived from spectroscopic ellipsometry, but this technique does not necessarily provide reliable

Table 2.4: Measured thickness of alumina deposited on Si substrates using plasma ALD process. Ellipsometric data has been fitted to a Tauc-Lorentz model to derive thickness.

Target thickness (nm)	Measured thickness (nm)	Number of cycles	Growth per cycle (nm)
300	295.2	2397	1.23
100	102.1	769	1.33
400	377.5	3100	1.22

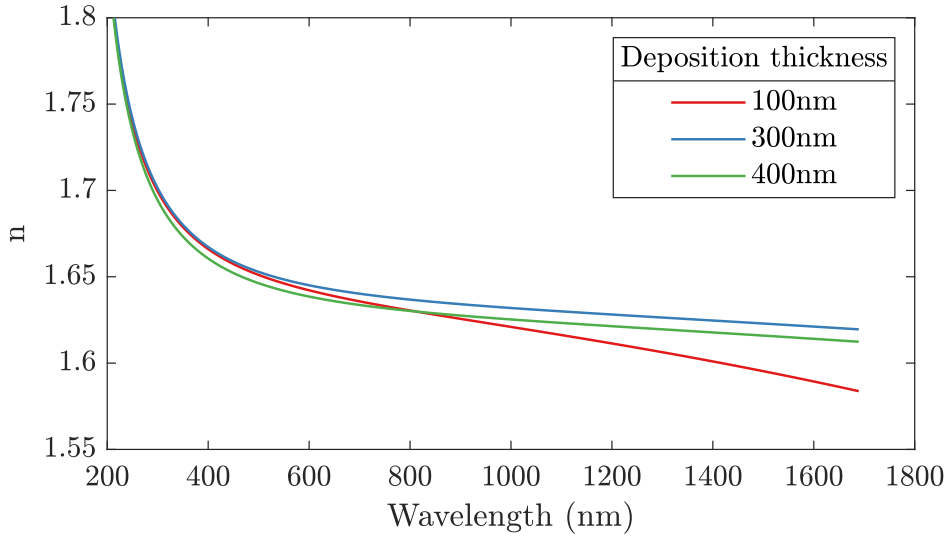

Figure 2.3: Chromatic dispersion of alumina films deposited using plasma ALD, derived using a Tauc-Lorentz model. Data from 300 nm-thick film reproduced from McKay et al. [165].

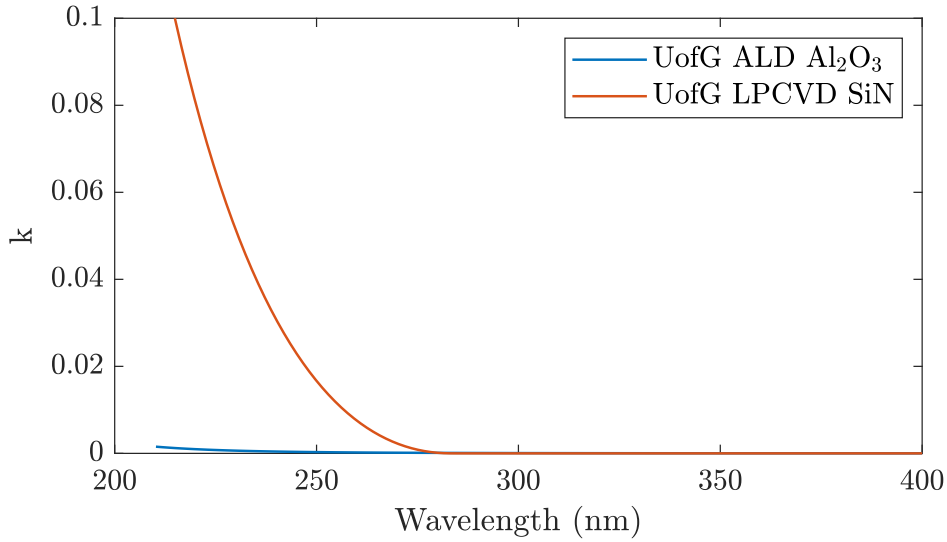

Figure 2.4: Extinction coefficient of 300 nm-thick alumina film deposited using plasma ALD, derived using a Tauc-Lorentz model. The extinction coefficient of a sample of low-pressure chemical vapour deposition (LPCVD) stoichiometric SiN, deposited in the JWNC, is provided to demonstrate that even where the alumina sample's extinction coefficient is nonzero, it is extremely low. SiN data provided by Natale Giovanni Pruiti.

information on bandgap—it is possible to derive varying values from well-fitted data [166]. Additionally, Corato-Zanarella et al. report that their ellipsometric process underestimated the extinction coefficient of SiN films where $k < 0.002$ [30]. None of the deposited amorphous alumina films had $k > 0.002$ within the range of the tool. A representative curve is shown in Figure 2.4, in which k becomes non-zero at about 300 nm and begins to rise markedly below 230 nm. However, the lowest wavelength available on the tool is 210 nm, and an abrupt rise in extinction coefficient is not evident in this range. An optical bandgap below 230 nm appears realistic, given the limited information available on the optical properties of amorphous alumina films in the far-UV region; extending the measurement range further into the UV, or performing transmission measurements, is required to quantify the position of the material bandgap.

2.3 Material Bandgap

VASE measurements could not quantify the optical bandgap of deposited alumina films; an abrupt rise in the extinction coefficient did not take place above the ellipsometer's minimum wavelength of 210 nm. UV-visible absorbance measurements offer an alternative, more direct way to measure a film's optical bandgap: monochromated light is passed through the sample and the transmitted light is measured by a photodetector.

I performed UV-visible absorbance measurements using a 378 nm-thick alumina film deposited on a microscope slide made from Heraeus Spectrosil 2000 fused silica, which is transparent to 180 nm [167]. I used a Horiba Duetta absorbance spectrometer to measure light absorbance through a sample mounted perpendicular to the light source over a wavelength range of 190 nm to 450 nm. I obtained background spectra using a blank microscope slide; spectra presented are the differential between the sample and background spectra.

Material absorption is not the sole contributor to absorbance (total attenuation) in this experiment. The effect of the transparent substrate on reflection is accounted for through collection of a background spectrum; however, the thin film itself can act as a Fabry-Pérot cavity, with reflections producing absorbances which do not indicate material absorption and which prevent accurate derivation of material bandgap. I calculated and removed Fabry-Pérot maxima and minima using the calculation procedure derived by Gough [168] from basic optical principles described in Hecht [169], Stenzel [170], Yu and Cardona [171], and Heebner et al. [172]. This methodology requires only chromatic dispersion data for the film and substrate, and the thickness of the measured film. I measured the thickness and chromatic dispersion of the alumina film using VASE (Section 2.2), and used silica chromatic dispersion data collected by Malitson [173]. Calculated maxima and minima showed a wavelength offset of 3 nm from the collected data—this disparity may be a result of fitting inaccuracies in the alumina

ellipsometry model, or because of disparities in the refractive index of the fused silica slide compared to the published silica ellipsometry data used in the calculation. I applied this 3 nm offset to the calculated reflection data to produce the corrected absorbance spectrum shown in Figure 2.5; there is little absorbance at wavelengths longer than 250 nm.

Optical bandgap can be derived from absorbance masurements using the Tauc method. The relationship between absorbance and bandgap is described by [174]

$$(\alpha \cdot h\nu)^{\frac{1}{\gamma}} = B(h\nu - E_g) \tag{2.1}$$

where h is Planck's constant, so $h\nu$ represents a photon of frequency ν, B is a constant, and E_g is the bandgap. For indirect electronic transitions, such as those which take place in amorphous solids, $\gamma = 2$. The absorption coefficient α is derived from absorbance measurements: $\alpha = {}^{ln(10)A}/_{t}$, where A is the measured absorbance and t is the film thickness. Photon energy $h\nu$ and $(\alpha \cdot h\nu)^{\frac{1}{2}}$ have a linear relationship near the bandgap; E_g can be estimated by performing a least-squares linear fitting to this region and extrapolating to the abcissa. Figure 2.6 shows a Tauc plot for the absorbance data presented in Figure 2.5; the derived bandgap position is dependent on the wavelength range selected for linear fitting (Table 2.5). Figure 2.6 shows a linear fitting to the region 190 nm to 200 nm, which avoids a noisy region above 200 nm; the derived bandgap is 5.75 eV (216 nm). Given that absorbance at the spectrometer's minimum wavelength of 190 nm remains relatively low (0.35), extending the measurement wavelength range further into the UV may provide a more definitive result.

The *electronic* bandgap of amorphous alumina is well-characterised, and has been measured as 7.0 eV (177 nm) in films deposited using thermal ALD [108]. Using UV-visible spectrometry, the *optical* bandgap, which is relevant to low-loss waveguiding, has been measured at 3.0 eV (413 nm) in thermal ALD films [175] and 3.6 eV (344 nm) in electrochemically synthesised films [176]. However, Aslan et al. measured slab losses in thermal ALD films and reported a rise in losses, indicating onset of the bandgap, at about 250 nm [74]; Shestaeva et al. reported the onset of absorbance at 230 nm in an oxygen PEALD film [153], but the sharp rise in absorbance which can be used to calculate optical bandgap was not within range of the spectrophotometer used to take measurements. This evidence demonstrates that there may be considerable variation in the bandgap of amorphous alumina films depending on deposition methods and parameters, and emphasises the necessity of characterising deposited films.

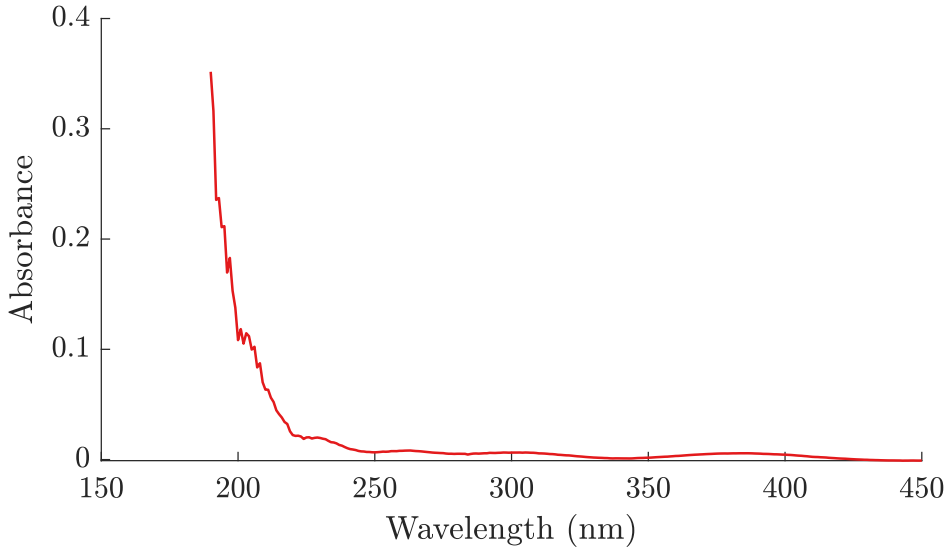

Figure 2.5: Absorbance spectrum of 378 nm-thick PEALD alumina film. Fabry-Pérot interference fringes have been reduced using the calculation procedure described by Gough [168].

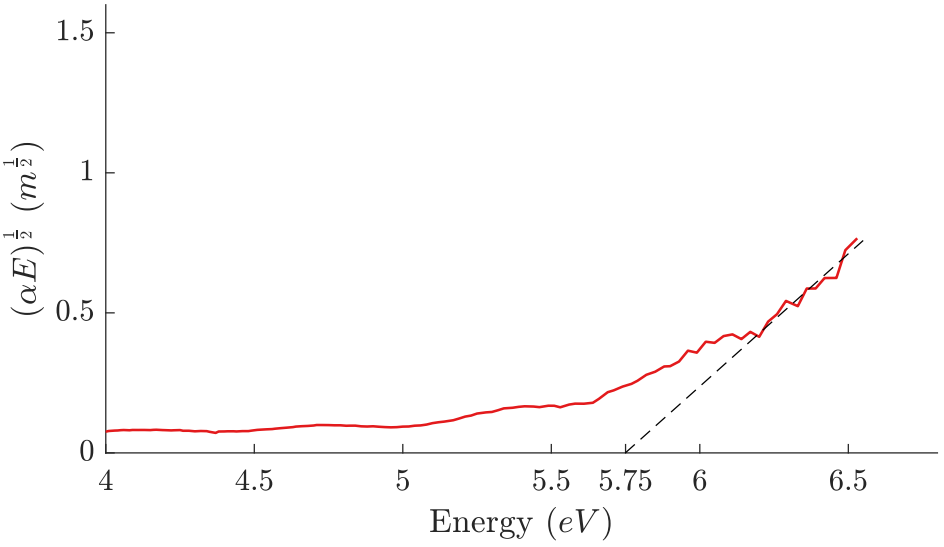

Figure 2.6: Tauc plot of 378 nm-thick PEALD alumina film. The dashed line is fitted to the wavelength range 190 nm to 200 nm; the optical bandgap is given by the abcissa of the fitted line (5.75 eV). Fabry-Pérot interference fringes have been reduced using the calculation procedure described by Gough [168].

Table 2.5: Variation of calculated bandgap value in PEALD alumina film with wavelength region used for linear fitting.

Wavelength range	E_g (eV)	E_g (nm)	R^2
190-195	5.86	212	0.91
190-200	5.75	216	0.95
190-210	5.45	227	0.93
190-220	5.35	232	0.96

2.4 Fourier Transform Infrared (FTIR) Spectroscopy: Chemical Composition

FTIR spectroscopy is a technique in which IR radiation is directed through a sample and the absorbance or transmission spectrum is measured. The IR radiation is absorbed to excite vibrational transitions within the material, with each absorption peak relating to a specific vibration mode for a bond. Therefore, FTIR spectroscopy provides information on the bonding composition of an alumina film, as well as on the presence of impurities. Radiation is absorbed according to the Beer-Lambert law—the signal obtained is proportional to the concentration of any given bond [177]. Therefore, whilst FTIR spectroscopy is most commonly used as a qualitative technique for identification of species, quantitative results are attainable if a tool is suitably calibrated. In this experiment, calibration was not possible, so peak heights can be considered generally indicative of the relative concentration of species.

IR radiation is only absorbed where an excitation alters the dipole moment of the affected bond. Fortunately, all of the species of interest in amorphous alumina films contain IR-active modes (Table 2.6). Aluminium centres can be coordinated to different numbers of oxygen atoms; in amorphous alumina, tetrahedral ($[AlO_4]$) and octahedral ($[AlO_6]$) configurations are found, as in crystalline forms; the additional presence of 5-coordinate alumina ($[AlO_5]$) is unique to the amorphous form [103, 178]. However, in amorphous alumina films, resolution of individual peaks associated with a specific aluminium coordination number is rarely feasible: the presence of multiple Al coordinations, and 2- and 3-coordinate oxygen within the structure, means that absorption energies vary over a wide range. The local environment of each Al centre defines the bandgap of the material, with the 6-coordinate crystalline form producing the widest bandgap [108].

The majority of FTIR analyses of amorphous alumina demonstrate a broad peak in the region of 400 cm^{-1} to 1000 cm^{-1}, but stop short of identifying specific coordination numbers [179, 180]. Those which attempt to elucidate aluminium coordination typically identify $[AlO_6]$ modes in the range of 480 nm to 600 nm and $[AlO_4]$ modes in the range of 700 nm to 1100 nm [120, 156, 181]. However, relatively recent work has consistently identified a large proportion of $[AlO_5]$ in amorphous alumina; the species has recently been described as “the fingerprint of the amorphous phase” [182]. This thermodynamically unlikely species was predicted through *ab initio* studies in 2009 [103] and identified through ^{27}Al 2D nuclear magnetic resonance (NMR) spectroscopy in 2010 [104]; more recent work corroborates this finding [182, 183]. Recent theoretical and practical work also cautions against explicitly assigning absorption wavelengths based solely on aluminium coordination number, as considerable shifts in peak position can occur as a result of oxygen coordination and local environment [182].

In addition to the Al-O bonds expected in the deposited material, precursors' unreacted functional groups may be present as impurities. For thermal processes, hydroxyl (-OH) groups, from incomplete reaction of adsorbed water molecules, are a common impurity [120]; they may also be adsorbed onto the material surface from atmospheric water vapour [148]. Hydroxyl groups introduce loss in the IR region—this has little impact on the visible optical applications described in this work, but will have an adverse effect at telecoms wavelengths [156].

Carbon-based contaminants may absorb light in the visible region [184]. ALD processes use TMA as a precursor, and if incomplete oxidation takes place then methyl groups (-CH_3, or -CH_2 if multiply bonded) may be incorporated into the film [185]. C-O bonds may be present, especially in PEALD films, as intermediate carbonate species form during combustion-like oxidation processes [159]. Positions of absorption peaks reported in the literature should be taken as indicative; localised variations in bonding may slightly alter absorption energies.

Table 2.6: IR-active modes in amorphous alumina. Wavenumbers are approximate, and will vary depending on the structure of the individual sample. Adapted from Demirtaş et al. [156].

Wavenumber (cm^{-1})	Mode
380-1065	AlO_x and Al-O-Al [182, 186]
1400-1600	CH_2 bend [152]
1423	CH_2 bend [187]
1604	OH bend [187]
1800-2100	C-O stretch [188]
2700-3100	CH_2 stretch [189]
3000-3700	OH stretch [189]

2.4.1 Method

I performed FTIR spectroscopy using a Bruker Vertex 80v vacuum FTIR system. I performed all measurements using alumina films deposited directly on Si, considering a wavenumber range from 4000 cm^{-1} to 400 cm^{-1} (2500 nm to 25 000 nm) at a resolution of 4 cm^{-1}. I obtained background spectra using blank Si; the displayed spectra are the differential between the sample and background spectra. Each line on the plot shows the average of 100 spectra. I analysed collected spectra using the Bruker OPUS software package, and applied a scattering baseline correction to each spectrum using default defined baseline points.

2.4.2 Composition of As-Deposited Films

I deposited alumina films directly onto Si during three deposition runs, with nominal thicknesses of 100, 300 and 400 nm. Figure 2.7 presents FTIR spectra from each film shortly after deposition. The relative absorbance intensity within the 400 cm^{-1} to 1100 cm^{-1} range is

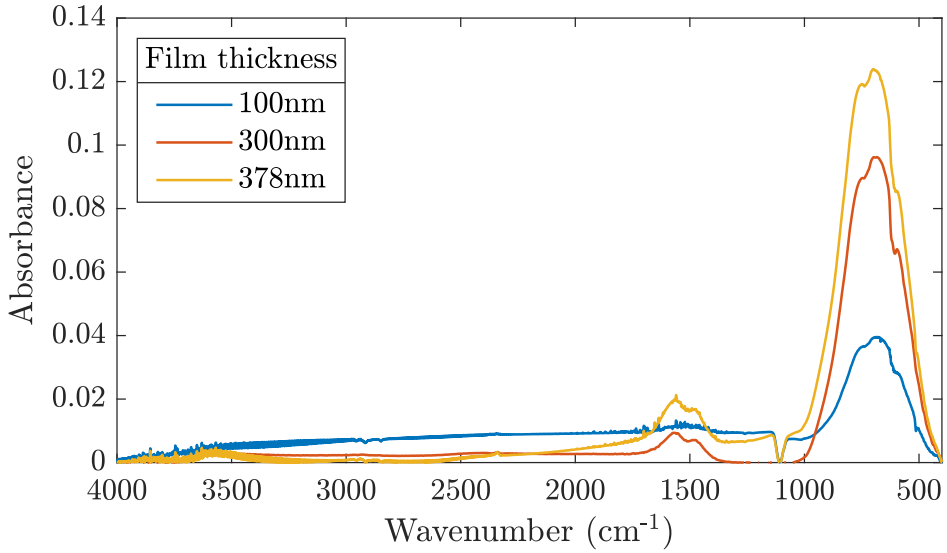

Figure 2.7: Mid-IR absorbance of amorphous alumina films directly after deposition.

similar for all samples, indicating a similar Al-O bonding composition; absolute absorbance scales approximately with sample thickness. All samples display peaks at 1400 cm^{-1} to 1600 cm^{-1}: these could be attributed to -CH_2 and -OH absorptions, although Verlaan et al. [152] and Kääriäinen and Cameron [180] have both identified a peak of similar shape and range in PEALD alumina films and attributed it to C-O bonding. A low-intensity absorbance at 3600 cm^{-1} indicates adsorbed water.

A particularly high level of broadband absorbance is observed for the 100 nm-thick film. Repeating the FTIR experiment produced the same result. To definitively assign a cause, further deposition runs would be required, which is outwith the scope of this work. One possible cause is a disparity between the Si wafer on which the alumina was deposited, and that which was used to collect a background spectrum; surface hydroxylation, for example, could somewhat affect the obtained spectrum. The difference between *absorbance* and *absorption* may also be considered; whilst the latter refers to attenuation as a result of a bond absorbing photons for excitation, a quantised process, the former simply means "any form of attenuation", including scattering or reflection losses. Scattering within the deposited film across a 100 nm path length is unlikely to produce appreciable scattering loss, especially given that the same deposition run produced low-loss waveguides in the blue (Section 4.6); however, carrier quality may also affect absorbance.

These observations are largely in line with the relatively large body of published FTIR spectra of ALD-deposited alumina. Katamreddy et al. [179] and Chowdhuri and Takoudis [190] report near-identical Al-O fingerprints from films deposited using thermal ALD and dioxygen PEALD, respectively. FTIR spectra of dioxygen PEALD films obtained by Verlaan et al. [152] and by Kääriäinen and Cameron [180] lack definition in the Al-O fingerprint region, but

are overall similarly shaped. Whilst Caballero-Espitia et al. do not identify carbon impurities in either thermal or ozone plasma ALD films, they identify considerably higher inclusion of -OH impuritities—this is likely because the oxidation step for both processes is unusually short [120].

2.5 Surface Roughness

Atomic force microscopy can be used to directly measure the morphology of a surface at the atomic scale. Atomic force microscopes (AFMs) use a sharp tip mounted on a cantilever with a weak spring constant to probe the sample surface; the tip is moved across the substrate in a raster scan and its displacement measured [191]. This provides a map of the surface being scanned. From a scan of the topside of a deposited film, it is simple to extract a numerical value for surface roughness, a contributor to scattering loss. Schematics showing the material morphology may also provide insight into the material's microstructure.

I used a Bruker Dimension Icon AFM to perform measurements on alumina which had been deposited on wafers with a 5 µm silica layer. I measured films from four deposition runs and scanned three randomly selected sites on each sample. I also measured a typical silica bottom cladding layer, since the roughness of a conformal film will be affected by the underlying layer. Each scan covered an area of 1 µm × 1 µm using a resolution of 256 points per scan. I used Gwyddion (version 2.59) [192] to analyse data, applying a levelling algorithm to each scan, then deriving the root mean square roughness R_{rms} of each sample [193] as

$$R_{rms} = \sqrt{\frac{1}{n}\sum_{i=1}^{n} y_i^2} \tag{2.2}$$

where n is the number of sample sites and y_n is the height measurement obtained for a given (x, z) coordinate.

Table 2.7 shows root mean square roughnesses of films deposited during four separate deposition runs. In all instances, the top surface of deposited films was smoother than that of the pre-deposition substrate, which is to be expected given the conformal nature of optimised ALD processes. Direct comparisons to the literature should be treated with caution, since the deposition conditions [105] and the underlying substrate [74] may affect roughness. However, Aslan et al. report topside roughness of 0.95 nm for a thermal process on a fused silica substrate (the relatively rough substrate—1.25 nm—may impact this value), West et al. report root mean square roughness of of 0.34 nm for a thermal process, and Li et al. report roughnesses ranging from 0.1 nm to 0.5 nm using a room-temperature plasma process. It should be noted that none of these papers mentions sampling over multiple sites or multiple

Table 2.7: Topside roughness of alumina films deposited using plasma ALD. Root mean square roughness analysis conducted using Gwyddion, version 2.59 [192].

Deposition date	Nominal thickness (nm)	R_{rms} (nm)			
		1	2	3	Avg
Typical pre-deposition (SiO_2) surface	–	1.04	1.08	1.05	*1.06*
January 2021	400	0.48	0.48	0.48	0.48
June 2021	300	0.80	0.82	0.91	0.85
June 2023	100	0.71	0.68	0.62	0.67
August 2023	400	0.64	0.63	0.63	0.64

runs with identical deposition conditions, so it is unclear whether the high run-to-run variance within these results is typical for the process.

However, even the highest of these results compares favourably with the surface roughness produced through other deposition methods. Whilst sputtering does not intrinsically produce atomically smooth films [194, 195], careful process optimisation and chemical-mechanical polishing can be used to produce films which are as smooth as those deposited using ALD [195, 196].

2.6 Film Properties Over Time

The literature contains numerous studies of the properties of as-deposited alumina [120, 148, 179]. However, alumina is somewhat hygroscopic. Immersion in water is known to corrode amorphous alumina (Section 2.7), and water adsorption—the formation of a hydroxylated layer on the material's surface, which may result in losses in the near-IR—has been demonstrated at relative humidities as low as 10% [197]. A platform centering a material which corrodes or becomes more lossy over time has little long-term, widespread viability.

Directly after deposition, I took VASE, FTIR and AFM measurements of a 300 nm-thick alumina film deposited directly onto an Si substrate. I stored this sample at uncontrolled ambient conditions (in a lab drawer) and repeated the same measurement protocols after 2 years.

VASE, using the same experimental and model parameters as for the original measurement (Section 2.2), produced a film thickness of 294.2 nm, similar to the original value of 295.2 nm. Re-fitting model parameters did not produce any appreciable change in n, k, or film thickness; considerable levels of water absorption would result in thickness and refractive index changes [198].

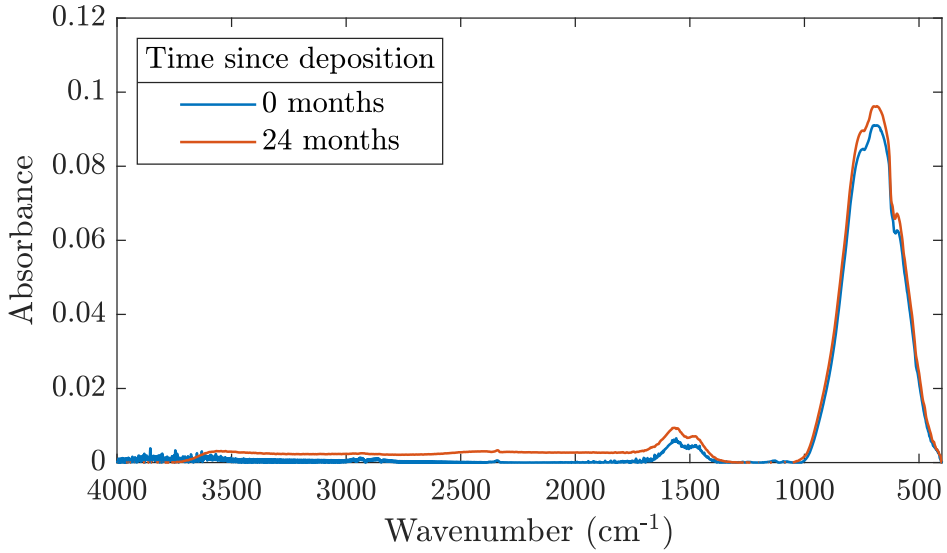

Figure 2.8: Mid-IR absorbance of 300 nm-thick amorphous alumina films at 0 and 24 months after deposition.

Repetition of the FTIR protocol demonstrated that film composition remained relatively consistent over time (Figure 2.8). Higher levels of absorption above 1400 cm^{-1} indicate a small increase in hydroxylation over time. This may affect film lossiness in the IR; there is little evidence to indicate the impact this could have on UV-visible absorption. Given that film thickness did not change appreciably in a two-year timespan, a more well-evidenced concern is an increase in film roughness caused by surface corrosion and restructuring.

Atomic force microscopy of this sample did not produce conclusive results. This film, which was deposited on Si, appeared particularly non-uniform, possibly a result of surface contamination prior to deposition. To average out the non-uniformity, I scanned a wide area of 2.5 µm × 5 µm using 512 points per scan. Data analysis was performed as described in Section 2.5. Results (Table 2.8) demonstrated high variance, so it is difficult to draw meaningful conclusions about aging-related roughening of alumina films. Performing this test on a smoother film—perhaps one deposited on a silica layer—is likely to produce more useful data.

Table 2.8: Topside roughness of alumina films deposited on Si using PEALD. Root mean square roughness analysis conducted using Gwyddion, version 2.59 [192].

Measurement date	R_{rms} (nm)			
	1	2	3	Avg
July 2021	2.57	1.04	1.18	1.60
July 2023	0.99	1.99	0.99	1.32

These experiments provide limited information on alumina aging. Over the course of two years, a small degree of film hydration appears to have occurred. Minimal change in film thickness implies that little corrosion has occurred. Atomic force microscopy results were intended to demonstrate the presence or absence of corrosion-related changes in film morphology [198]; however, poor repeatability renders the data unusable. In Section 4.5, I demonstrate low optical losses in waveguides made from films which were deposited more than two years prior (stored at 40-50% relative humidity), indicating that aging did not have a considerable effect on waveguide loss.

2.7 Chemical Compatibility

The ALD processes used in this project have produced films which, upon deposition, are consistently low-roughness (Section 2.4). However, during device fabrication, films are exposed to a range of chemical processes which may have undesirable effects.

Hydrofluoric acid (HF) is used to remove silica and related materials, such as HSQ; alumina etches in acids, although the etch rate is dependent on the deposition parameters [114, 199]. In this project, I used concentrated (25%) tetramethylammonium hydroxide (TMAH), which has a pH of > 13, as a developer; dissolution of ALD alumina in basic solutions has been reported [198], and etch rate increases with temperature and with increasing basicity.

Reagents used in rinsing and cleaning processes may also affect alumina films. Water can corrode or dissolve amorphous alumina: Correa et al. report hydration and roughening of alumina films deposited using thermal ALD following 10 days in deionised water [198], whilst Abdulagatov et al. report alumina dissolution at 90 °C [200]. Even relatively short immersions may result in restructuring of the material's surface [201], producing a rougher nanostructure which is likely to increase propagation losses. There is no information on the effect of heating alumina in organic solvents used in cleaning steps, such as acetone and isopropyl alcohol (IPA), or on the effect of oxygen plasma cleaning, known as *ashing*.

Furthermore, none of the studies cited in this section have studied the effect of chemical processing on alumina films produced using PEALD; given that deposition parameters can produce differences in chemical structure and reactivity [102], I studied the effects of chemical processing on the films used in this project.

I measured the surface roughness of 400 nm-thick alumina films deposited on a silica BOX layer before and after exposure to chemical processing, using the AFM protocol described in Section 2.5 (Table 2.9). Initial treatment durations were selected to resemble the maximum likely exposure to a reagent that a typical alumina sample would receive during processing. Therefore, I subjected samples to 10 min in each liquid reagent, with the exception of HF,

Table 2.9: Topside root mean square roughness of alumina films before and after chemical processing. Each value is the mean of three R_{rms} measurements. Root mean square roughness analysis conducted using Gwyddion, version 2.59 [192].

Process	Initial R_{rms} (nm)	Duration 1 (min)	Post-processing R_{rms} (nm)	% change from initial	Duration 2 (min)	Post-processing R_{rms} (nm)	% change from initial
Water 20 °C	0.48	10	0.64	34%	60	0.64	34%
Water 50 °C	0.56	10	0.85	50%	60	19.29	3316%
Water 80 °C	0.65	10	50.38	7678%			
TMAH 20 °C	0.60	10	0.42	-30%	60	0.41	-32%
TMAH 50 °C	0.60	10	0.42	-30%	60	0.75	25%
Acetone 20 °C	0.57	10	0.58	2%	60	0.57	0%
IPA 20 °C	0.58	10	0.69	17%	60	0.65	12%
O_2 ash 150 W	0.58	5	1.60	176%	(6x5)	0.70	21%
6:1 HF 20 °C	0.65	1	5.92	810%			
10:1 HF 20 °C	0.63	1	3.12	396%			

where the duration was limited to 1 min. I also measured surface roughness before and after a plasma ash of 5 min at a maximum tool power of 150 W.

Amongst the 30 initial measurements, one standard deviation is 0.08 nm, 12% of the mean value. (The complete dataset is included in the data repository associated with this work.) Where post-treatment root mean square roughness remained within approximately 3 standard deviations of the mean pre-treatment value, I performed a second treatment of extended duration, aiming to exacerbate any observed effect.

Deposited amorphous alumina films' reactivity is in accordance with the literature: amorphous alumina etches at high and low pH [202] and corrodes in hot water [198, 200], whilst organic solvents do not appear to affect the material. Process design rules based on material compatibility can be summarised as

- HF etches alumina leaving a rough surface (R_{rms} >3 nm), and must not be used during processing.
- TMAH corrodes alumina, leaving a smooth surface (R_{rms} <0.5 nm); at room temperature, etching over the exposure duration required during processing is negligible. At higher temperature, the effect of TMAH etching on feature size and sidewall angle must be considered.
- Room-temperature deionised water can be used for the durations required during processing (up to 10 min); however, extended contact with room-temperature water and any contact with hot water should be avoided.
- Organic solvents are unlikely to interact with amorphous alumina films and may be used freely.
- It is unclear whether oxygen plasma ashing will roughen films, and therefore it is unclear whether it should be used during processing.

An interesting upshot of this work is that films treated with TMAH became smoother than as-deposited films. This may have applications in low-loss waveguiding. During plasma etching processes, waveguide sidewalls are roughened by ion bombardment, increasing interface losses. TMAH etching could be used to etch and smooth waveguide sidewalls, reducing both sources of loss.

2.8 Annealing and Thermal Budget

Annealing—exposing a sample to heat to drive structural change—may improve a material's optical properties. Many fabrication processes also require sample heating regardless of whether it is desirable—such as when baking photoresist or curing cladding materials. Heat is also produced during plasma processing.

Amorphous alumina is thermodynamically unstable compared to crystalline phases of alumina; heating will produce structural changes [203]. At the atomic level, these changes consist of annihilation of $^{[5]}$Al and $^{[3]}$O in favour of more stable $^{[4]}$Al, $^{[6]}$Al, and $^{[2]}$O [100, 183]—predicted by Filatova and Konashuk to change the bandgap [108]. Afanas'ev et al. determined that annealing produces a wider electronic band gap in thermal ALD alumina films, but do not quantify the effect on optical band gap [204]. Several groups who have successfully produced low-loss waveguides using ALD alumina films report annealing at temperatures ranging between 400 °C to 600 °C, but none provide a comparison of waveguide losses before and after annealing [4, 205].

On a larger scale, annealing-induced bonding changes result in microstructural changes in the film, observed as a transition from an amorphous to a polycrystalline structure. Crystalline phases have a higher refractive index than the amorphous phase [105] and their nucleation will typically increase light scattering by producing refractive index inhomogeneities within the film. The phase (γ- or α-alumina) and the crystallisation temperature are strongly dependent on film thickness—or, specifically, the ratio of surface energy to configurational entropy [206]. Crystallisation temperatures ranging from 450 °C to 1000 °C are reported in the literature; thicker films crystallise at lower temperatures [203].

Annealing may facilitate the outgassing of impurities, such as unreacted methyl, carbonyl and hydroxyl groups from ALD precursors. Whilst reducing impurity levels could reduce extrinsic loss, in studies on thin catalytic alumina ($<$10 nm) and alumina/zinc oxide nanolaminates ($<$100 nm), the outgassing process resulted in pore formation [148, 207]—which will increase light scattering and surface roughness.

2.8.1 Method

In order to test the effect of temperature on the deposited ALD films, I developed annealing processes with a relatively long dwell time, aiming to maximise thermally-driven material restructuring. I used a Jipelec JetFirst 200 rapid thermal annealer (RTA) for all processes, which limited the dwell time to 20 min. I used a slow ramp rate of 5 °C s^{-1} to reduce the risk of film cracking due to coefficient of thermal expansion (CTE) mismatch between the film and the substrate [208], with the final 100 °C ramp conducted at 2.5 °C s^{-1} to avoid overshooting target temperature (Figure 2.9). All anneals were conducted in a nitrogen atmosphere to reduce the risk of film outgassing [209].

I annealed 300 nm-thick alumina films deposited directly on Si at 400 °C and 800 °C—below the lowest reported crystallisation temperature [210], and at the most commonly reported crystallisation temperature [75, 114]—and performed the ellipsometry and FTIR spectroscopy protocols detailed in Sections 2.2 and 2.4 before and after annealing.

I also annealed samples of 400 nm alumina on a silica BOX layer at 400 °C, 600 °C and 800 °C. These films are smoother, and have more consistent surface roughness, than those deposited directly on Si, and are more suitable for atomic force microscopy. I followed the protocol described in Section 2.5 to measure the surface roughness of each sample before and after annealing; I also performed optical microscopy on these samples.

Finally, I annealed a 378 nm-thick alumina film deposited on fused silica, and performed the bandgap measurement protocol described in Section 2.3.

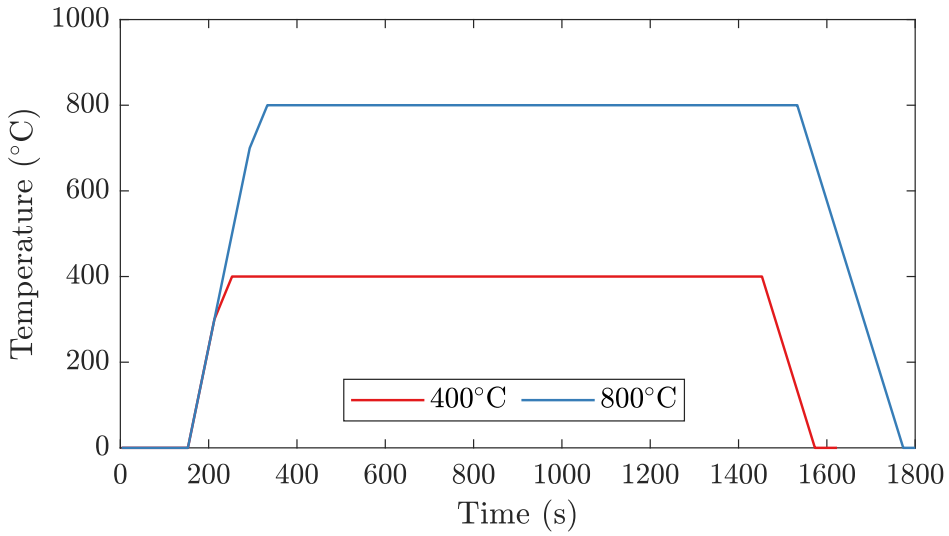

Figure 2.9: RTA profiles used for annealing of alumina films.

2.8.2 Microscopy: Cracking and Outgassing in Annealed Films

Figure 2.10 shows bright-field optical micrographs of 400 nm alumina-on-silica samples prior to annealing, and after annealing at 400 °C, 600 °C and 800 °C. Annealing at all temperatures produced a change in film colour, which may be indicative of a change in the films' optical properties in the visible region. Films annealed at 400 °C and 600 °C show little difference from their original state, whereas cracking is evident following annealing at 800 °C.

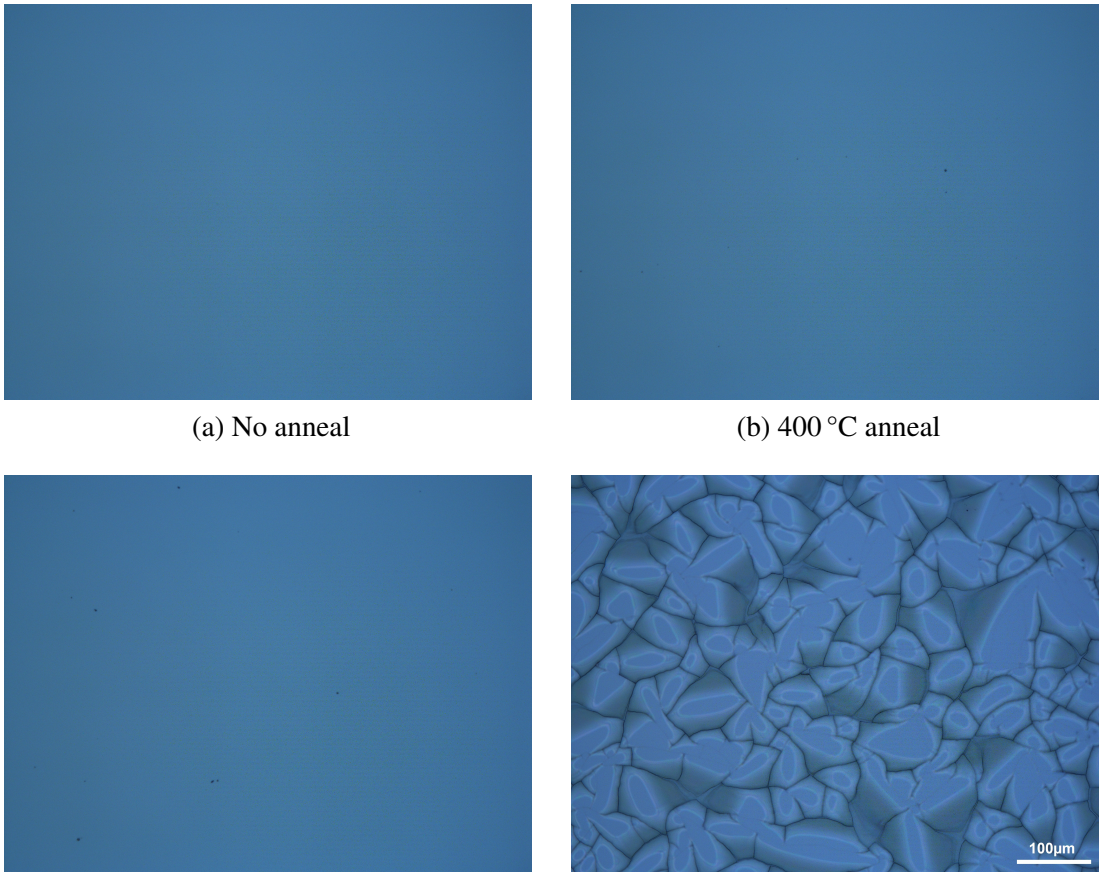

(a) No anneal (b) 400 °C anneal

(c) 600 °C anneal (d) 800 °C anneal

Figure 2.10: Bright-field optical micrographs of 400 nm-thick alumina films deposited on SiO_2. All images are taken at 200x magnification and using the same illumination settings.

2.8.3 Ellipsometry: Optical Constants of Annealed Films

I compared the thickness and optical constants of 300 nm-thick alumina films deposited on an Si substrate. Whilst data for the film annealed at 400 °C demonstrated a good fit to the model used for as-deposited films (mean square error of 12), the film annealed at 800 °C showed a poor fit to the original model (mean square error over 100), and a new model had to be developed using a decreased wavelength range of 210 nm to 1550 nm.

Table 2.10: Ellipsometer thickness measurements of alumina films before and after annealing. Deposited films had a nominal thickness of 300 nm prior to annealing.

Annealing temperature (°C)	As-deposited thickness (nm)	Thickness after anneal (nm)
400	294.0	293.7
800	295.3	255.7

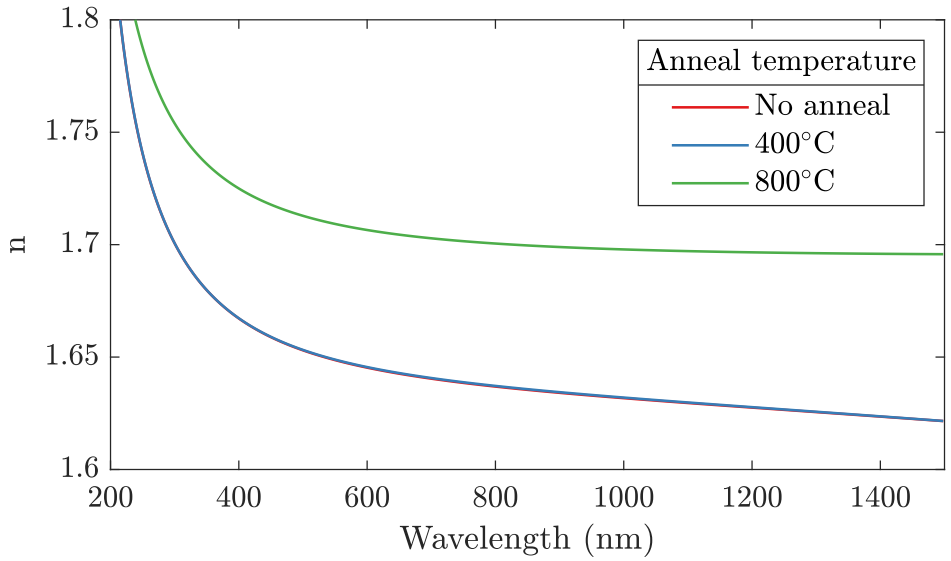

Figure 2.11: Refractive index of alumina films deposited using plasma ALD. Films annealed at 400 °C and 800 °C are compared with an as-deposited sample. Refractive index was derived using a Tauc-Lorentz model. The "no anneal" and "400 °C anneal" lines almost completely overlap.

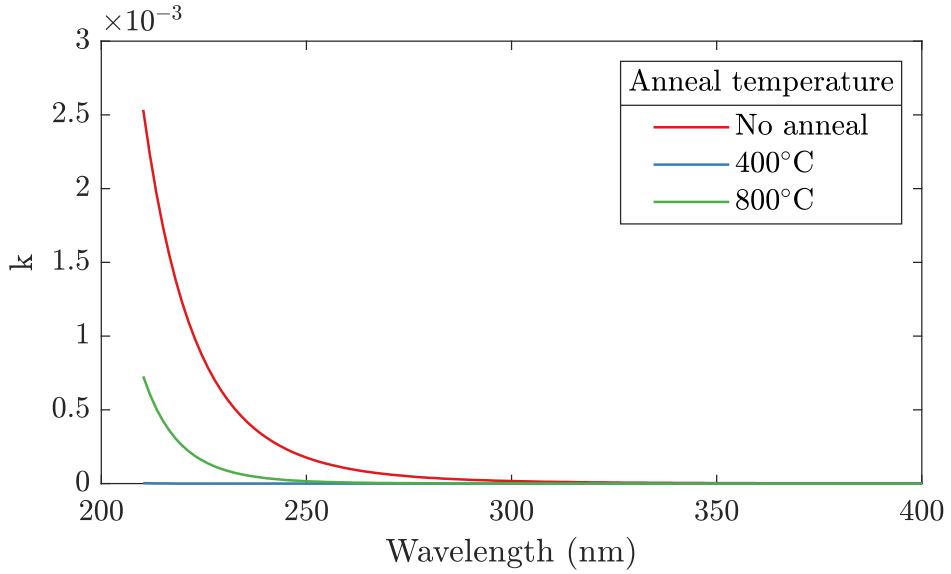

Figure 2.12: Extinction coefficient of alumina films deposited using plasma ALD, derived using a Tauc-Lorentz model. Films annealed at 400 °C and 800 °C are compared with an as-deposited sample. Note that the y-scale denotes extremely small changes in k.

The film sample annealed at 800 °C compacted considerably, showing a 13% decrease in thickness (Table 2.10) and an increase in refractive index (Figure 2.11). These observations are consistent with the onset of crystallisation [75, 204]. The film annealed at 400 °C showed very little change in either thickness or refractive index. Whilst annealing appears to affect the absorption coefficient of alumina films (Figure 2.12), it should be noted that material absorption is extremely low for all measured samples and that small changes in model fitting could produce the degree of variation observed between results [30, 166].

2.8.4 Bandgap Changes in Annealed Films

I used the UV-visible spectrophotometry procedure described in Section 2.3 to determine the bandgap of a 378 nm-thick PEALD alumina film deposited on a fused silica microscope slide. As for films which had not been annealed, I measured sample absorbance over a wavelength range from 190 nm to 450 nm and removed Fabry-Pérot reflections from the collected absorbance spectrum using the procedure outlined by Gough [168].

Figure 2.13 compares corrected absorbance spectra of as-deposited and annealed films. Noise in the spectrum of the as-deposited sample prevents a detailed comparison; however, there are no considerable differences between the two spectra at wavelengths above 200 nm. Below 200 nm, absorbance in the annealed film increases considerably more than in the as-deposited film. Figure 2.14 compares the bandgaps derived from the two samples. For the annealed sample, a linear fit to the region from 190 nm to 195 nm provides a high R^2 (Table 2.11); the calculated bandgap is 6.09 eV (204 nm). (Analogous results for the as-deposited film are provided in Table 2.5.)

Whilst Tauc plots based on spectra collected from 190 nm to 450 nm can be used to derive a wider bandgap for annealed films than as-deposited films, this result may be inaccurate: the area of the Tauc plot used for linear fitting is at a considerably shallower gradient for the as-deposited film. It may be necessary to measure absorbance at shorter wavelengths to achieve an appropriate linear fitting for as-deposited films.

As discussed in Section 2.3, these results should therefore be taken as indicative rather than quantitative. We can say that the PEALD alumina films deposited for this work are transparent at wavelengths of interest, and are, broadly, transparent at wavelengths above 250 nm; however, accurate derivation of the optical bandgap of deposited films may require collection of absorbance measurements at wavelengths below 190 nm.

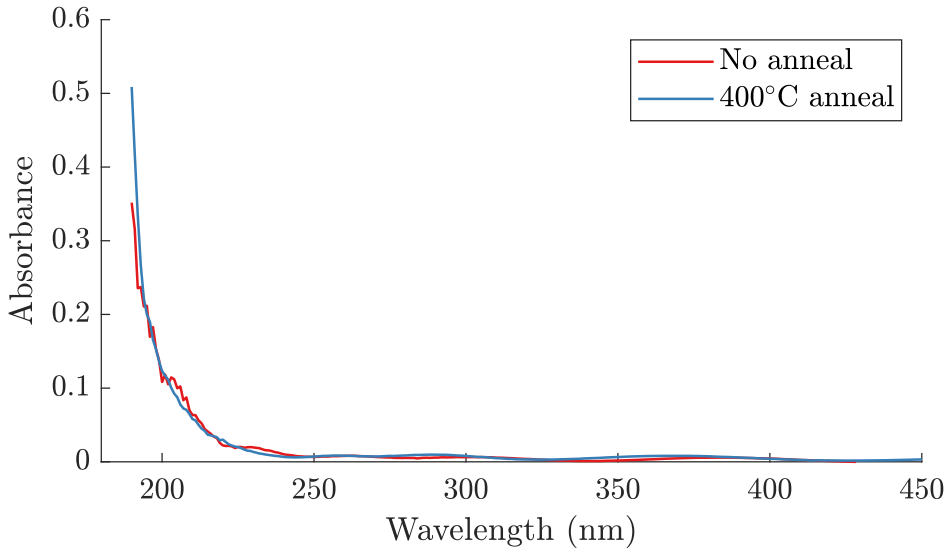

Figure 2.13: Absorbance spectra of 378 nm-thick PEALD alumina films with and without post-deposition annealing at 400 °C. Fabry-Pérot interference fringes have been reduced using the calculation procedure described by Gough [168].

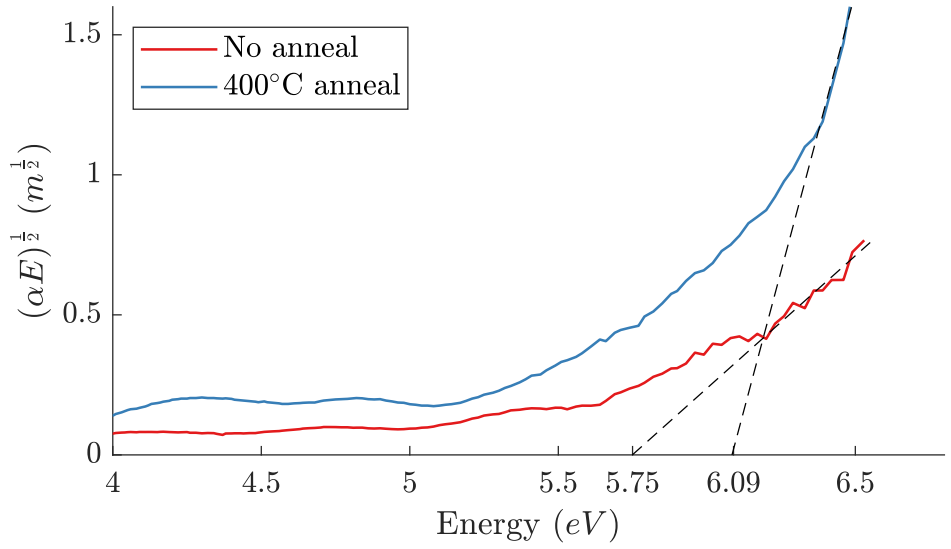

Figure 2.14: Tauc plot of 378 nm-thick PEALD alumina films with and without post-deposition annealing at 400 °C. The dashed lines are fitted to the wavelength ranges 190 nm to 200 nm for the non-annealed film and 190 nm to 195 nm for the annealed film; the optical bandgap for each film is given by the abcissa of the fitted line. Fabry-Pérot interference fringes have been reduced using the calculation procedure described by Gough [168].

Table 2.11: Variation of calculated bandgap value in 400 °C annealed PEALD alumina film with wavelength region used for linear fitting.

Wavelength range	E_g (eV)	E_g (nm)	R^2
190-195	6.09	204	0.97
190-200	5.87	211	0.93
190-210	5.58	222	0.92
190-220	5.40	230	0.92

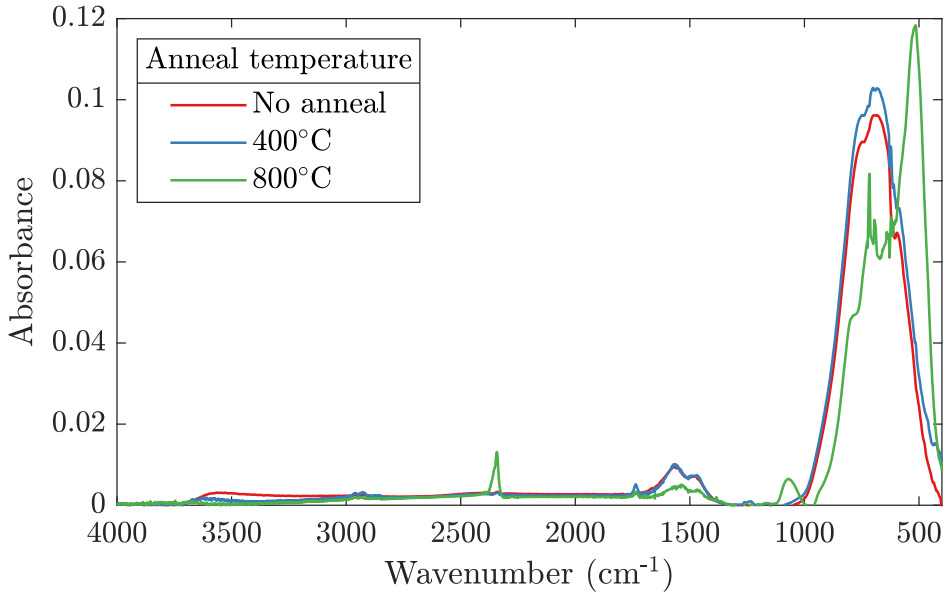

Figure 2.15: Mid-IR absorption of alumina films annealed at 400 °C and at 800 °C, compared to a non-annealed film.

2.8.5 FTIR: Bonding in Annealed Films

I used the FTIR protocol described in Section 2.2.5 to measure the absorption spectra—and probe the bonding composition—of alumina films annealed at 400 °C and 800 °C (Figure 2.15). Compared to an as-deposited sample, both annealed samples display a reduction in the -OH absorption at 2900 cm^{-1} to 3600 cm^{-1}, suggesting that water—likely surface-adsorbed—has been driven from the films. The film annealed at 800 °C also shows a decrease in the C-O mode at 1600 cm^{-1}. Results from Verlaan et al. demonstrate a similar reduction in this absorbance upon annealing [152]. High-temperature annealing also produces a striking change in alumina coordination. Whilst the sample annealed at 400 °C shows a similar bonding pattern to the as-deposited sample, the onset of crystallisation during annealing at 800 °C is evident, with a shift to a lower-energy absorption which is not prominent in the other samples. This may represent the transition of $^{[4]}$Al and $^{[5]}$Al to $^{[6]}$Al. The additional peak at 1100 cm^{-1}, only observed after annealing at 800 °C, indicates either the formation of an interfacial layer containing Al-O-Si [211] or thermally induced Si-O-Si bonding within the carrier wafer [179].

A spike at 2400 cm^{-1} after annealing at 800 °C is likely spurious; this absorption is associated with atmospheric CO_2, and may be present as a result of incomplete evacuation of the vacuum chamber [212].

2.8.6 AFM: Surface Roughness in Annealed Films

Reports on the degree of roughening of alumina films are varied. Crystallite formation may increase film roughness [204], as will the formation of nanopores caused by outgassing [148]. Outgassing of ALD films is reported at temperatures as low as 600 °C [148, 209]. Statistically significant increases in roughness are typically reported after crystallisation: for example, Correa et al. report no change in roughness following annealing at 450 °C (uncrystallised), but an increase at 900 °C (crystallised) [198]. Zhang et al. report no increase in surface roughness in 47 nm-thick alumina films grown using thermal ALD following a spike anneal at 1000 °C [213], whereas Broas et al. report a R_{rms} increase of an order of magnitude in 100 nm-thick films annealed using similar conditions [102]. This disparity may arise from the decrease in crystallisation temperature observed in thicker films [203].

Annealing produced an increase in R_{rms} across all samples, with roughness increasing with temperature (Table 2.12), although this increase can only be considered statistically significant for the sample annealed at 800 °C. In Figure 2.16, it is evident that roughness on the sample surface has increased; there is also evidence of a nanopore produced by film outgassing. Root mean square measurements are sensitive to large, isolated changes in height, so nanopores are likely to drive up the reported value.

2.8.7 Discussion

The annealing work I have performed is limited; I made no attempt to optimise annealing parameters and used the maximum allowed dwell time at each temperature. Even so, these experiments provide some useful information about allowable thermal processing of alumina films, as well as some avenues for future work.

Films annealed at 800 °C showed a clear change in alumina coordination, a decrease in thickness, and an increase in refractive index, in line with the changes expected during crystallisation. A decrease in both adsorbed hydroxyl groups and bulk carbonyl groups is apparent—however, this bulk desorption produces film outgassing which renders films unsuitable for optical applications.

Films annealed at 400 °C showed no statistically significant change in film roughness and little change in film thickness, refractive index, or alumina coordination. Neither ellipsometry nor UV-visible spectrometry measurements have quantitatively established that annealing at this temperature widens the optical bandgap of alumina films. The annealing process also reduces adsorbed water; this is likely to reduce absorption loss in the IR, but no effect on visible loss has been proven.

Table 2.12: Topside roughness of 400 nm-thick alumina-on-BOX films deposited using plasma ALD before and after 10 min annealing. Measurements reported are the mean of three 1 µm × 1 µm scans. Root mean square roughness analysis conducted using Gwyddion, version 2.59 [192].

Annealing temperature (°C)	R_{rms} pre-anneal (nm)	R_{rms} post-anneal (nm)	% change in roughness
400	0.45	0.49	9%
600	0.41	0.52	26%
800	0.41	1.14	178%

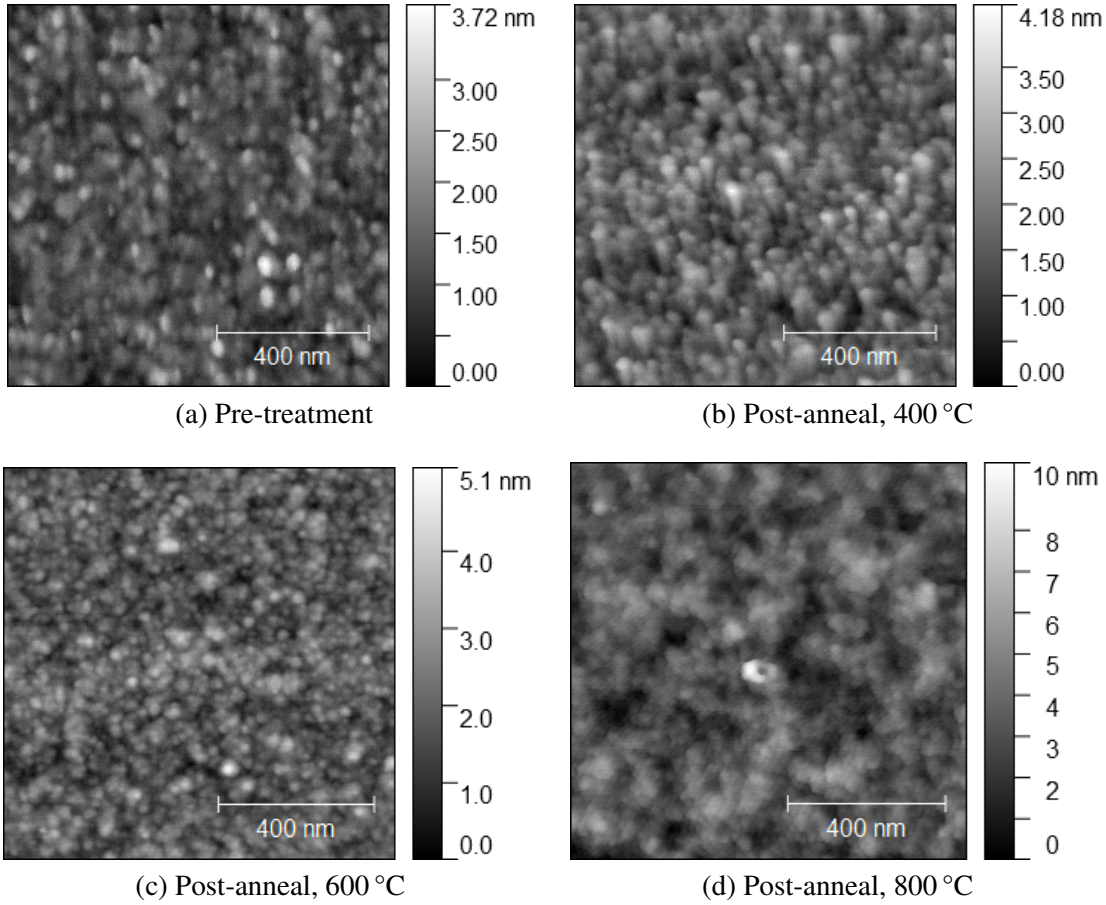

(a) Pre-treatment

(b) Post-anneal, 400 °C

(c) Post-anneal, 600 °C

(d) Post-anneal, 800 °C

Figure 2.16: Atomic force micrographs of alumina morphology before and after 10 min annealing at 400 °C, 600 °C and 800 °C under N_2 atmosphere.

In the basic process used to fabricate devices during this project, a capacity to withstand a temperature of 400 °C is sufficient. However, a higher temperature may be desirable for HSQ processing (Sections 3.2.4 and 3.4). We can say that the maximum allowable annealing temperature lies between 400 °C and 800 °C, and that temperatures must be kept below 800 °C to avoid cracking. The overall thermal load applied and the film thickness will also affect allowable annealing parameters [203].

2.9 Conclusion

ALD processes can reproducibly deposit high-purity, uniform conformal films—these characteristics point towards favourable optical properties. However, the thermal ALD processes used for all prior demonstrations of etched alumina waveguides suffer from extremely long deposition times, a result of the difficulty of purging water vapour from the reaction chamber. Films used in this project were deposited using an oxygen PEALD process, which addresses this concern by using an easily-purged oxygen plasma for the oxidation half-reaction. The PEALD process also produces a higher GPC than an analogous thermal ALD process.

The optical and structural properties of alumina films are highly dependent on the parameters used during deposition, such as the deposition temperature, the dwell time on each step, and the oxidising agent. Data on the chemical and optical properties of alumina deposited using dioxygen PEALD is extremely limited. Therefore, it is useful to design fabrication processes based on measured properties of films deposited using this specific process in this specific reactor, rather than using information from the literature.

The low number of deposition runs performed limits the conclusions which can be drawn from the data, especially in terms of the consistency of depositions. However characterisation of as-deposited alumina films suggests that they are likely to have favourable optical qualities. Spectroscopic ellipsometry and UV-visible spectrometry indicate that the films have a wide bandgap ($\sim$5.75 eV), with no evidence of intrinsic absorption at the target wavelengths for waveguiding (visible and near-UV). FTIR spectroscopy indicates surface hydroxylation and the presence of carbon within the bulk of the film. These impurities certainly absorb in the infrared; it is unclear whether they could produce extrinsic losses in the visible region. Optimisation of the deposition process—performing a longer oxidation step [159] or using a higher temperature [152]—could reduce carbon impurities. The surface roughness of deposited films—which will affect scattering losses—varied considerably, but R_{rms} remained below 1 nm. Films remained stable over 2 years at ambient conditions, albeit with a small increase in adsorbed water which may be removed upon heating.

Films must be able to maintain their optical quality during fabrication processes. Brief immersion in cold water and TMAH will not appreciably damage alumina films; immersion

in organic solvents is acceptable; hot water, HF, and hot TMAH will dissolve films. It is unclear whether oxygen ashing has an effect on film properties. Limited annealing tests suggest that heating to 400 °C in an inert atmosphere is unlikely to have a considerable adverse effect on film properties; establishing optimised annealing protocols, and an upper limit on acceptable annealing temperature, would require further work.

Prior work on optical-quality alumina has pointed to annealing a means of reducing losses in the visible region, although there are no published results of optical or waveguide losses before and after annealing [4, 205]. Optical bandgap measurements in this work do not point to considerable changes in bandgap following annealing at 400 °C, but results' reliability is limited because deposited films' bandgaps lie close to the spectrophotometer's cutoff wavelength.

TMAH etches alumina, leaving a smooth surface ($<$0.5 nm). Given that etch rates are relatively low, TMAH etching offers a highly controllable method for smoothing alumina surfaces. Top surfaces, which are variable in roughness, could be smoothed prior to further processing; finished waveguides, which are prone to interface scattering losses at etched sidewalls, could be smoothed by controlled immersion in TMAH. The processing techniques developed in this thesis allow etching of waveguides up to 800 nm thick whilst maintaining adequate sidewall quality; this value is limited by the thickness of HSQ masks (Section 3.2) and places a limit on etch power. A technique which allows controllable smoothing of sidewalls could allow the use of thicker, but rougher, masks, or more aggressive etching processes, allowing thicker alumina waveguides to be fabricated.

Chapter 3

Alumina Processing Toolbox

As described in Section 1.2, waveguides, along with most other PIC components, can be produced using the same basic steps:

- Deposit a lower cladding layer with a low refractive index
- Deposit a core layer with a higher refractive index
- Mask off areas of the core layer to define optical structures
- Etch away all (ridge) or part (rib) of the high-index layer, leaving the areas which were masked off
- Optionally, apply an upper cladding layer with a low refractive index

To minimise propagation loss within a fabricated waveguide, the core and cladding material must have low absorption and bulk scattering losses, and roughness at the core-cladding boundary must be minimised. Development of high-confinement alumina waveguides has been hampered by two main challenges: depositing thick layers of optical-quality amorphous alumina, and defining deep-etched waveguides with low sidewall roughness.

In Chapter 2, I evaluated the properties of amorphous alumina films deposited using a repeatable dioxygen PEALD process with a relatively short deposition time. This chapter presents a suite of processes which can be used to etch alumina films with acceptable sidewall roughness. The novel processes presented in this chapter are masking and etching processes; together, these define interface scattering at etched sidewalls, which is a key source of loss in the visible regime. A mask material must be sufficiently etch-resistant, but also produce smooth and vertical sidewalls when etched. Plasma etching offers a wide range of chemistries and processing parameters which can be varied to alter the selectivity, verticality and sidewall roughness achieved—but these parameters frequently trade off against each other.

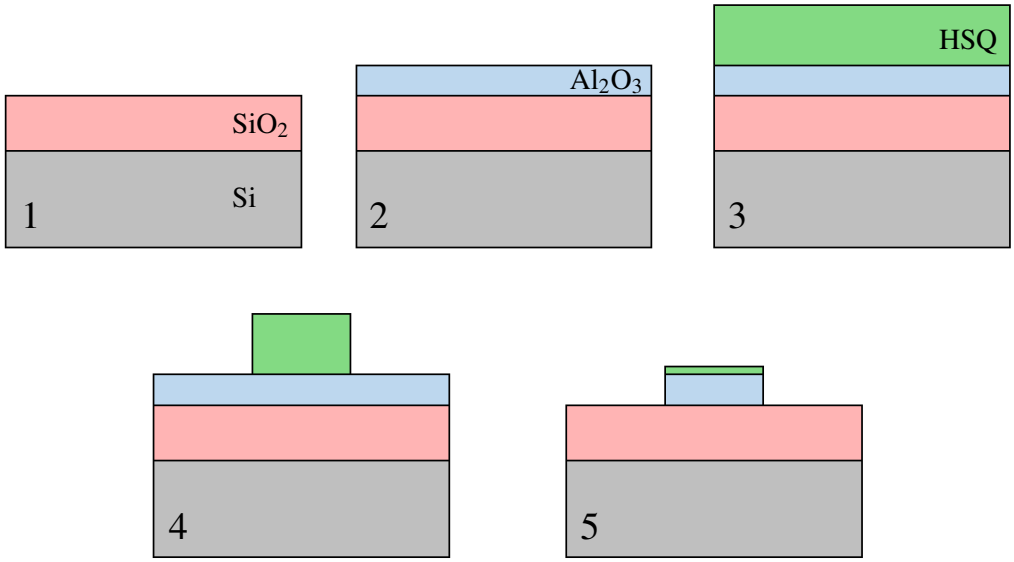

Figure 3.1: Graphical representation of waveguide process flow. Colours are consistent; new layers are labelled as they are introduced. Note that layer thicknesses have been selected for clarity and are not proportional to actual thicknesses.

This chapter is intended to act as a manual containing all the processing information required to fabricate fully-etched optical structures in thick alumina films deposited using dioxygen PEALD. Whilst there is a focus on masking and etching, I also address cladding and facet definition processes, with a view to their effects on optical loss. I begin by describing the full process flow in its final form (Section 3.1), then provide a deeper analysis of each step, describing the process optimisation performed to minimise waveguide losses and maximise manufacturability.

3.1 Process Flow

The generic alumina waveguide process flow, shown in Figure 3.1 is a ridge waveguide process flow similar to that used for SOI platforms.

- The devices are supported by an Si carrier wafer with a bottom cladding layer (5 µm of SiO_2 grown through thermal oxidation).
- Amorphous alumina is deposited using a dioxygen PEALD process.
- Wafers are cleaved into pieces, which are processed individually.
- PIC structures are patterned using HSQ resist, which is exposed using EBL. Following exposure, samples receive a post-exposure bake and are developed in TMAH.
- An inductively coupled plasma reactive ion etch (ICP-RIE) process is used to transfer the HSQ pattern into the alumina.

- Optionally, an upper cladding may be added: I use HSQ as a spin-on dielectric, baking it in a convection oven to crosslink the material.

In this project, I coupled light into the waveguides using end-fire coupling (Appendix C). As such, after the waveguides are defined, facets must be produced. The top surface of the waveguide is protected with polymethyl methacrylate (PMMA) resist, a scribe made which is aligned to the Si crystal plane, and the facet manually cleaved along the crystal plane.

Each step of this process flow has been designed for its capacity to produce components for low-loss PICs:

- The deposited alumina layers have a wide bandgap (low intrinsic absorption) and low surface roughness (low interface scattering). In Chapter 4, I demonstrate that deposited films do have low losses in the visible region. The PEALD deposition process has a relatively high growth rate, facilitating the deposition of thick layers. The films are relatively stable to atmospheric moisture over time, and are stable to the conditions (heat, solvents) required to define PIC structures (Sections 2.7 and 2.8).
- The optimised HSQ etch mask has vertical sidewalls with low sidewall roughness (Section 3.2).
- The optimised BCl_3 etch, when used in combination with the HSQ mask, can be used to etch waveguides up to 800 nm thick, and produces sidewalls of 80° with low sidewall roughness (Section 3.3).
- HSQ cladding is conformal, minimising scattering losses; it appears to have high transparency across the visible spectrum [214] and in this project was used to clad waveguides which had low losses in the blue (Section 3.4).
- The scribe and break process for producing end facets has a yield $> 90\%$ (Section 3.5).

All process parameters are provided in Table B.1.

3.2 Lithography

The majority of waveguides are fabricated using a *subtractive* process, in which a continuous film of the waveguide core material is deposited, waveguide shapes are defined using a lithography process, and the entire substrate is etched to remove the core layer in all but the masked areas.

Lithography processes involve spin-coating a uniform layer of a UV- or electron-sensitive *resist* over the whole sample, then selectively *exposing* part of the resist layer to *mask off* a waveguide geometry defined using computer-aided design (CAD). The resist is *developed* by

immersion in a solvent, which removes the resist in areas which will be etched. *Transfer* of the mask pattern into the substrate is most commonly performed using a plasma etch.

Together, masking and etching processes define waveguide geometry, which decides allowed light propagation modes (Section 4.2) and allowed nonlinear processes (Section 5.2). They also define sidewall roughness, and therefore interface scattering losses. The contributions of the two process steps cannot be clearly separated: during plasma etching, both mask and substrate are etched, a complex reaction in which outcomes are dependent on the relative reactivity of the mask and substrate, the area and geometry of the mask, the species generated in the plasma, and the energy of those species.

Desirable properties of a resist mask include

- Predictable variation of feature dimension with exposure dose and dimensions defined using CAD
- Vertical sidewalls: typically maximises dimension control during etching, and maximises sidewall verticality of etched structures
- Low sidewall roughness: typically minimises roughness of etched sidewalls
- Resistance to plasma etching: high selectivity allows etching of thicker layers, may reduce transfer of lithographic roughness from mask to substrate, and generally improves sidewall verticality
- The ability to remove the mask after etching; if mask material cannot be removed, it must not have an adverse impact on optical properties (low absorption and scattering loss)
- A lack of involatile etch byproducts which are difficult to remove or which adversely impact etching

Alumina is etched in halogen plasmas, so any mask material must be sufficiently stable to halogen plasma etching. The primary options for mask materials are displayed in Table 3.1, along with some of their key characteristics. The maximum thickness and plasma etch rate together define the depth which can be etched using a mask. *Hardmask* processes are processes where the mask must be defined in two steps: in other words, the mask itself is defined through a lithography-etch cycle. This can result in a reduction in sidewall quality [122], and can make dimension control more difficult.

Photoresist masks may be defined using EBL or conventional UV photolithography. EBL can define structures with more fidelity than UV lithography, especially in academic settings where the process development and financial costs of high-resolution UV lithography tools are prohibitive [215]. Masks made from photoresists, such as those in the novolac and polyacrylate families, suffer from low selectivity, but can compensate for this by being many

Table 3.1: Comparison of mask materials for alumina etching.

Mask material	Maximum thickness	Etch rate in halogen RIE	Hardmask?
Photoresist	Microns	High	No
HSQ	1 µm	Medium	No
PECVD silica	Microns	Medium	Yes
Metal (e.g. Ni, Cr)	Microns	Low	Yes

microns thick; however, they are prone to sidewall roughness and their high erosion rate limits sidewall verticality [122, 125]. Metal masks are highly etch-resistant [122]; however, they can be difficult to remove, especially given amorphous alumina's susceptibility to acid and base attack. Additionally, their use is forbidden in many of the JWNC's plasma processing tools, as metal redeposition on chamber walls may contaminate subsequent processing runs.

The Optoelectronics group at the University of Glasgow has an established HSQ masking process which has been used to develop high-quality waveguides in other materials such as SiN [216] and AlGaAs [217]. Exposed HSQ has a silica-like structure, but avoids the need for a hardmask process; the structures produced using the base process have smooth and vertical sidewalls. However, further optimisation was required to produce masks which are suitable for waveguiding at short wavelengths.

It is not possible to fully assess the effect of lithography without considering the effects of etching on the mask. All of the optimisation work described in this section used the optimised BCl_3 etching process described in Section 3.3. I make the assumption that, in all instances, the area masked off is sufficiently small, and structures are sufficiently far from each other, that loading (mask area dependence) and RIE lag (aspect ratio) effects do not need to be considered.

Line-edge roughness (LER) and stitching optimisation work described in this section was performed jointly with Natale Giovanni Pruiti, who performed optimisations of HSQ masks using multipass lithography reported in this work, and stitching tests not reported in this work.

3.2.1 Theory of EBL

All EBL was performed on a Raith EBPG5200 EBL tool. EBL exposes resists using a focused beam of high-voltage electrons, which is scanned over the sample. Compared to conventional UV lithography, EBL can produce substantially higher-resolution patterns; with suitable resists (such as HSQ) and optimised processes, extremely low lithographic roughness can be achieved. The direct writing mode allows for rapid prototyping, whereas conventional lithography requires that a new mask is fabricated for any pattern changes.

The electron source is a thermal field emitter. The electrons are accelerated across an electric field, then focused and shaped using a series of condenser lenses. The beam formed is round and Gaussian in shape. In order to write shapes, the beam is moved over the substrate using deflection coils. These coils are responsible for moving the beam from location to location, scanning the beam at a given location, and adjusting the focus and stigma [218].

Patterns are prepared using computer-aided drawing programs, and must be translated into a series of actions which can be performed on the EBL tool. Shapes are *fractured* into a series of rectangles, which can be written using a vector scan and described as a series of shots. Electron scattering within the substrate material also exposes the area surrounding the beam spot: Monte Carlo simulations are performed and the dosage across the pattern is altered to mitigate these loading effects (proximity correction) [219]. The pattern is then split into *fields* and *subfields*. The beam is scanned over the subfield using fine deflection coils, altering the dwell time in each location to alter the dose. Coarse deflection coils move the beam from subfield to subfield. Large movements from the central position cause distortion of the beam spot, so field-to-field movement is performed through stage movement [220]. Finally, a field writing order is assigned, ideally providing an efficient path through the write area.

Optimised EBL tool and job parameters are critical to producing low-loss waveguides. There are two key considerations: ensuring that the waveguides remain continuous, and ensuring that the sidewalls are smooth. Drift between the beam and the stage can cause variations in field positioning up to tens of nanometres per hour [218]. Consequently, field size and ordering should be chosen so that adjacent fields are written sequentially. Additionally, field size must be a multiple of subfield size to ensure that lithography remains continuous at field boundaries.

Lithography parameters may affect waveguide smoothness, since a pattern may be designed as a smooth line but is written as a series of (at best) circular shots. Selecting an appropriate exposure dose, step size, spot size, and filling pattern is essential, especially for curved structures, which are processed as many-sided polygons by EBL writing software [221].

3.2.2 HSQ Chemistry and Application

The term hydrogen silsesquioxane (HSQ) describes a range of monomers with the overall formula $[HSiO_{3/2}]_n$. When energy is applied, the monomer crosslinks, forming Si-O bonds and releasing hydrogen (Figure 3.2) [222].

Monomer crosslinking occurs upon application of heat, or energy in the form of an electron beam or UV radiation [223]; extensive crosslinking produces a *network* structure. Any uncrosslinked monomer readily dissolves upon immersion in basic developers. HSQ is

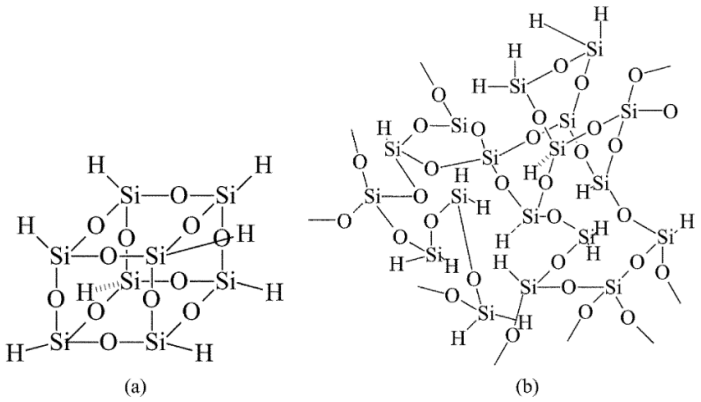

Figure 3.2: HSQ structure a) before and b) after crosslinking. Reproduced with permission from Yang et al. [222].

sensitive to exposure dose, particularly when using TMAH developer; a high degree of crosslinking is essential for the formation of a high-quality mask [224].

The line-edge roughness (LER) of HSQ masks is strongly dependent on processing conditions. Even if exposure conditions have been optimised to minimise roughness, nanometre-scale roughness arises from the granularity of the crosslinked resist [224]. Aggregates of different sizes form during the polymerisation process—aggregate size is related to the baking and/or exposure conditions to which the HSQ has been subjected [223, 225].

In order to be used as an EBL resist, the HSQ monomer is suspended in an organic solvent, allowing it to be distributed on substrates using spin coating. The resist is pipetted onto a clean substrate (0.1 mL to 0.2 mL is sufficient to coat typical samples of 2 cm^2 to 4 cm^2), then spun at several thousand revolutions per minute to produce a uniform film. Spin coating produces exceptionally consistent, uniform films: assuming that environmental conditions are consistent, the film thickness depends solely on the resist dilution and the spin speed. During the spinning process, the solvent evaporates off, leaving the active monomer evenly distributed on the substrate. Following spinning, a *softbake*—a brief, low-temperature bake—can be used to drive off more solvent [226].

HSQ is a negative-tone photoresist, in which electron beam exposure drives resist crosslinking. A *post-exposure bake* (PEB), typically performed at a higher temperature than the softbake, can be used to drive further crosslinking, albeit at the expense of contrast [227]. Development in basic media—in this instance, TMAH—removes uncrosslinked resist.

Minimisation of an HSQ mask's LER requires careful control of thermal and exposure parameters, to optimise for low-granularity line-edges. Furthermore, the resist's high sensitivity (to dose variation) means that consistent results require high uniformity and low run-to-run

variation across each step of the process, such as resist thickness, baking time and temperature, exposure dose, and develop time and temperature.

3.2.3 Base Lithography Process

The Optoelectronics group has well-established processes using Dow Corning FOx-16, which uses methyl isobutyl ketone (MIBK) as a solvent for HSQ monomer. However, these processes have been used primarily for the fabrication of low-loss waveguides in the IR, and for etching materials where selective etching processes are available [216, 217]. Initial lithography conditions were

- Spin HSQ 2000 rpm 1 min—resist thickness 700 nm
- Softbake on hotplate at 92 °C, 15 min
- Expose at dose 1400 μC/cm^2
- Post-exposure bake in convection oven at 180 °C, 1 h
- Develop in 25% TMAH 30 s, rinse in deionised water then IPA, blow dry using nitrogen gun

The base process required a number of modifications for use in this application. Firstly, a 700 nm-thick HSQ mask is insufficient for etching waveguides for nonlinear applications. Simulations (Section 5.2) indicate that, in order to achieve zero GVD, 800 nm-thick alumina films must be fully etched. When etching amorphous alumina with a silica mask, optimised etching processes yield a selectivity of approximately 1:1 [4]; the highest etch selectivity achieved in this project was 1:0.8. Increasing the thickness of the HSQ layer by decreasing the speed of the distribution spin is trivial; however, thicker films are prone to cracking and may not possess the uniformity required for high-quality EBL processing. Therefore, Section 3.2.4 describes work done to establish the maximum HSQ thickness which can be attained whilst maintaining film quality, and Section 3.2.5 describes the development of methods to increase the etch-resistance of HSQ masks.

400 nm-thick waveguides fabricated using the base process were too lossy to measure; scanning electron microscopy showed considerable sidewall roughness after etching (Figure 3.5). Whilst increasing exposure dose provided a considerable improvement, I also implemented changes to EBL settings and thermal processing to reduce sidewall roughness (Section 3.2.6). However, waveguide losses remained high as a result of *stitching errors*, where adjacent EBL fields are misaligned; Section 3.2.7 describes stitching mitigation protocols.

3.2.4 Maximum Mask Thickness

HSQ is applied to a sample using a spin coating technique, as described in Section 3.2.2. Spin coating produces highly reproducible film thicknesses; if all other parameters are kept the same, the film thickness depends solely on the spin speed.

The base process described in Section 3.2.3 was designed to produce HSQ masks of thicknesses up to 700 nm. The limiting factor is film uniformity: using Dow FOx-16 resist, this thickness is achieved at a spin speed of 2000 rpm. Given a maximum etch selectivity of 0.8, this resist thickness is insufficient for etching of thick waveguides.

At low spin speeds, the high variation of thickness with spin speed can compromise film uniformity [226]. However, the maximum feasible single-layer film thickness reported in the literature is 1.2 µm [228]; films thicker than this crack as a result of high tensile stress. High-quality lithography processes have been reported in films as thick as 1 µm [229]. Since alumina etch processes tend to have poor selectivity, at best 1:1 [4], I investigated the maximum mask thickness which can be achieved without compromising uniformity.

I characterised the spin behaviour of HSQ by spinning HSQ onto bare silicon at a variety of spin speeds and using the spectroscopic ellipsometry methodology documented in Section 2.2 to measure film uniformity. Results are reported in Table 3.2; a full process flow is provided in Table B.3.

The Si pieces used were 18 mm × 18 mm, the size of the substrate used for early generations of devices. I measured uniformity over 5 points at the corners and centre of a 10 mm × 10 mm square: at least 4 mm of the substrate is removed when each facet is cleaved, so the device footprint, and the area in which resist uniformity is critical, is only 10 mm × 10 mm. When spinning samples at speeds below below 2000 rpm, an HSQ *edge bead*—a band of thick resist at the edge of the chip—formed (Figure 3.3). However, uniformity in the central area remained in line with values attained at higher spin speeds (within process and measurement error). Therefore, lower spin speeds are acceptable for these samples—especially since HSQ is a negative resist, and areas which are not exposed are removed during the develop process.

Samples spun slower than 1000 rpm showed cracks throughout the film, the effect of tensile stress [228]. The same effect was apparent at the edges of samples spun at 1000 rpm, where

Table 3.2: Thickness and uniformity of HSQ spun at different spin speeds.

Spin speed (rpm)	Average HSQ thickness (nm)	HSQ uniformity (%)
1000	1018	2
2000	713	4
3000	584	3

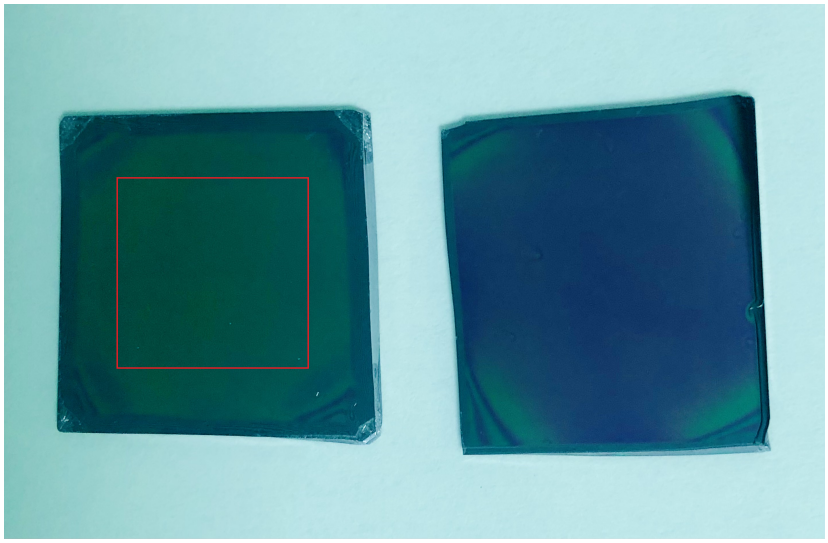

Figure 3.3: Photographs of HSQ films spun onto approx. 18 mm × 18 mm Si pieces. **Left:** Sample spun at 1000 rpm. Note the darker band around the edge of the sample. The area outlined in red is approx. 10 mm × 10 mm, showing the area in which waveguides will be patterned. **Right:** Sample spun at 2000 rpm.

the central HSQ thickness was 1 µm but the edge bead was substantially thicker. However, the cracking did not extend past the edge bead area, so posed no problems to the lithography.

Therefore, the maximum feasible thickness of an HSQ mask is 1 µm; given an etch selectivity of 0.8, alumina films up to 800 nm can be etched, and waveguides with zero GVD can be fabricated.

3.2.5 Improving HSQ Etch Resistance

HSQ is crosslinked during the EBL writing process, producing an amorphous structure similar to that of silica. Post-exposure thermal annealing further crosslinks HSQ films [230], which may improve etch resistance [223]. In order to test the effect of thermal annealing on etch mask resilience, I created HSQ monitor samples and annealed them for 20 min using the protocol described in Section 2.8. After treatment, I etched each sample in BCl_3 plasma for 10 min (full process flow provided in Table B.4). I used the spectroscopic ellipsometry methodology described in Section 2.2 to measure film thicknesses before annealing, after annealing, and after etching (Table 3.3). Etch resistance is reported as a proportion of film removed: whilst all samples had a similar HSQ thickness prior to annealing, the annealed samples compacted to different extents. Increasing temperature produced a reduction in etch rate, even when the effect of annealing on film compaction is taken into account. This is likely a result of thermally driven crosslinking within the HSQ structure.

Exposing HSQ to oxygen plasma has been used to create more etch-resistant masks [224]; uncrosslinked Si-H bonds are oxidised to form Si-O-Si crosslinks (a more silica-like structure)

Table 3.3: Etch rate of HSQ in BCl_3 plasma following 20 min annealing in N_2.

HSQ treatment	% HSQ eroded	HSQ erosion rate as % of erosion rate of untreated film
None	26	100
Asher 150 W 20 min	26	100
Anneal 400 °C	22	85
Anneal 500 °C	19	72
Anneal 600 °C	17	63

[231]. I exposed an etch test sample and a monitor piece to oxygen plasma in a PlasmaFab RF barrel asher, at a power of 150 W (maximum power) for a time of 20 min; however, there was no reduction in etch rate.

Using a mask annealing process certainly reduces the etch rate of HSQ in BCl_3 plasma; however, this does not provide a complete picture of the effects of HSQ annealing. I did not investigate the effect of annealing on HSQ roughness; whilst the formation of granules which affect line-edge roughness is primarily associated with development and the preceding processes [224], high temperatures may alter the mask's microstructure. In addition, the effect of annealing on the etch rate of amorphous alumina has not been established. However, this annealing process may prove valuable in fabricating a range of waveguides using HSQ masks, many of which have higher thermal budgets than this process (for example, SiN [216, 232]).

3.2.6 Line-Edge Roughness Optimisation

For a given waveguide geometry and propagation wavelength, interface scattering losses are defined by the waveguide's interface roughness (Section 1.2). According to the Payne-Lacey model, interface scattering losses scale with $1/\lambda^3$ [3]. When using EBL, the line-edge roughness (LER) of the mask is affected by both pattern preparation (shape fragmentation and shot filling patterns) and process factors which define resist sensitivity.

Firstly, consider a scenario in which lithographic shot filling patterns produce a perfectly straight edge (which is never the case in reality). Here, the LER is controlled solely by the *granularity* of the resist—in other words, the size of the resist molecules and the structures they make when they are crosslinked. When a resist film is developed, areas which are not crosslinked are dissolved; areas which are crosslinked are retained. Near the edge of the film, resist aggregates which are not well-bonded to the bulk of the mask will be removed. The incident electron beam does not provide a completely uniform energy distribution—it is approximately Gaussian in shape, and additional exposure occurs as an effect of electron scattering in the substrate. Therefore, areas at the edge of the resist are exposed to lower exposure doses than the bulk of the resist. Furthermore, *shot noise*—tiny fluctuations in the

number of incident electrons—can produce small-scale variations in exposure dose at the pattern's edge [224].

Resist *sensitivity* describes the strength of correlation between exposure dose and crosslinking; sensitive resists are prone to surface roughness as a result of incomplete crosslinking at pattern edges [227]. Low-sensitivity resists are also described as *high-contrast*. HSQ forms small aggregates, and therefore has the potential for very low sidewall roughness; however, its sensitivity is strongly dependent on processing parameters [224]. Strategies for improving contrast in HSQ resist masks include the use of a PEB to drive crosslinking to completion [233] and the use of a high-concentration and/or high-temperature developer solution to ensure complete removal of weakly bonded aggregates [225].

EBL beam scan patterns also have an effect on roughness. Using a sensitive resist, it is particularly important to design filling patterns which minimise dose variation at waveguide edges. Proximity error correction (PEC), which alters dosage to account for substrate scattering effects, plays a considerable role here [219]. The beam spot size, beam step size, and filling pattern can also have an effect, explored further in Section 3.2.6.

Directly quantifying LER is extremely challenging. Two primary techniques exist: directly measuring sidewall roughness using AFM, which requires challenging additional processing [40], and deriving sidewall roughness from waveguide loss measurements [30]. Fabricating enough waveguides to accurately quantify roughness consumes a substantial amount of material, as well as processing and testing time—and assumes that waveguide losses are measurable. The first set of waveguides fabricated during this project, made using an optimised etch and the standardised Optoelectronics group's established HSQ process, proved too lossy to measure. Given the difficulty of quantifying sidewall roughness, I used scanning electron microscopy to provide a qualitative indication of the effect of process changes on sidewall roughness.

Exposure Dose Optimisation

Exposure dose optimisation is required for two reasons. Typically, exposure dose tests are performed to control the *critical dimension* (CD), the linewidth of the smallest (or the most performance-sensitive) feature [226]. However, exposure dose also impacts LER: sensitive resists can produce high LER when underexposed. The use of a PEB and a strong developer, such as the 25% TMAH used in this project, can help to mitigate this effect—but only if the exposure dose is correctly optimised to ensure that the pattern is evenly exposed.

Initially, waveguides were written at an exposure dose of $1400\,\mu C/cm^2$: scanning electron microscope (SEM) inspection of an exposed sample showed features which were fully ex-

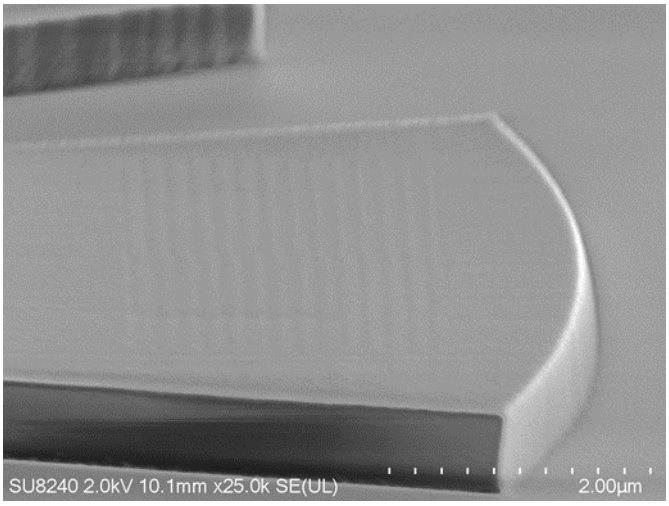

Figure 3.4: Scanning electron micrograph of wedge-shaped structure exposed in HSQ. Sidewalls are vertical and smooth.

posed and fully resolved at this dose. Figure 3.4 shows a typical mask structure following development, with vertical sidewalls and no visible sidewall roughness.

I fabricated revision 1 waveguides (see Table 4.1) using the Optoelectronics group's standard HSQ process flow; however, they proved too lossy to measure. Scanning electron microscopy of these waveguides showed significant sidewall roughness, especially in the curved areas of the waveguide (Figure 3.5a). This roughness had not been observed prior to etch. The roughness followed shape fragmentation patterns, suggesting that areas at the edges of shapes did not receive enough energy to fully crosslink. Sample height mismeasurement may also exacerbate this effect by affecting shape scaling (see Section 3.2.7).

For all subsequent lithography optimisations, I imaged the mask after development, subjected samples to a 5 min etch using a standard BCl_3 etch, and re-imaged. Following a dose test performed using this protocol, I increased the exposure dose to 1450 µC/cm^2—as shown in Figure 3.5b, the fragmentation pattern is no longer visible.

Post-Exposure Bake Optimisation

HSQ may be crosslinked using energy provided through EBL or heat [223]. EBL exposure at a dose which allows feature resolution may not provide sufficient energy to fully crosslink the resist. Post-exposure baking can improve contrast (reduce sensitivity) by increasing crosslinking in partially-crosslinked areas at the edge of the film [234]. However, thermal load may also adversely impact surface roughness [225].

Comparing the standardised HSQ lithography process, which performs a 1 h oven bake at 180 °C, to process flows reported in the literature [227, 233, 234, 235], substantial differences are evident. All of the flows cited used a short (1 min to 10 min) hotplate PEB, with tempera-

tures ranging from 95 °C to 200 °C. Under ideal conditions, hotplates provide more consistent and more efficient heat transfer than ovens [226].

To determine optimal PEB parameters, I spun HSQ onto an Si substrate and exposed a standard test pattern featuring a variety of lines, micro-rings, and pillars. Features patterned directly onto Si are easier to image with high resolution using scanning electron microscopy; omission of the dielectric layers used when fabricating waveguides reduces charge buildup on the sample surface, which creates image artifacts. Since the majority of the electron scattering is likely to happen within the Si layer, the impact on the lithography is expected to be negligible.

I prepared samples using different post-exposure bake conditions: the standard 1 h 180 °C oven bake under N_2 atmosphere, and 1 min hotplate bakes at 90 °C and 180 °C. I also fabricated a sample with no post-exposure bake. I used optical microscopy and high-magnification scanning electron microscopy to assess exposure and LER, etched the samples for 5 min, and re-imaged. A full process flow is provided in Table B.4.

Figure 3.6 shows optical micrographs of 1 µm-wide micro-rings with a bend radius of 10 µm. There is a noticeable difference in feature size when comparing the samples which received a PEB to the sample which did not. The colour gradient on the rings on the unbaked sample suggest that the sample has sloping sidewalls; it is interesting to note that there appears to be *less* resist removed from this sample, despite the fact that less crosslinking is likely to have occurred in this sample. Exposing samples without a PEB would therefore require an adjustment in exposure dose, and these results may not be representative of the best sidewall profile which could be achieved for this condition. However, given the precedent for use of a PEB and the favourable results produced using a PEB, I did not explore optimisation of an HSQ masking process without a PEB.

Scanning electron micrographs of 200 nm × 200 nm HSQ pillars (Figure 3.7) after resist development demonstrate thermally-driven differences in profile. Samples which have been processed with a higher thermal load show some resist bridging between pillar features, demonstrating the higher degree of crosslinking which has taken place. It appears that the original oven process may have a higher thermal load than the hotplate process at the same temperature. The higher thermal load reduces feature resolution, implying that sensitivity is increased.

Following a 5 min BCl_3 etch, differences between samples are more apparent. Square pillar features developed without a PEB or using a hotplate for PEB all resolved, showing similar size gaps between the features. Sidewall slope is apparent on samples which did not receive a PEB. Using an oven PEB, the 100 nm gap between pillars could not be resolved.

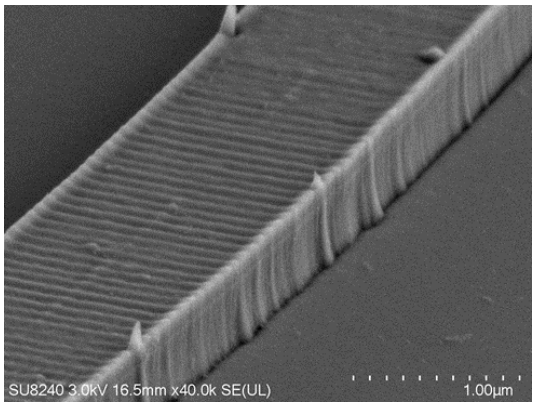

(a) Curved waveguide exposed using a dose of1400 µC/cm^2 (underexposed).

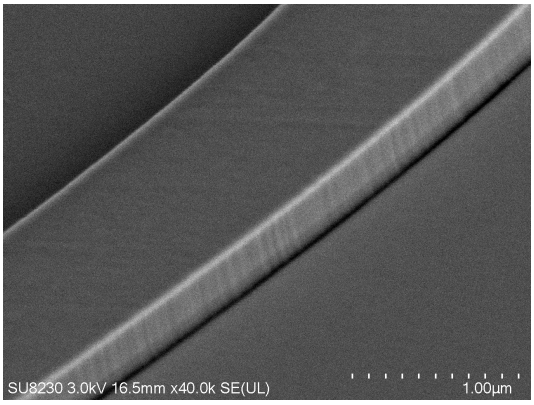

(b) Curved waveguide exposed using a dose of1450 µC/cm^2 (correctly exposed).

Figure 3.5: Scanning electron micrograph of underexposed and correctly exposed curved waveguides after etching. Underexposure produces uneven etching of the HSQ mask, and the same pattern is transferred into the underlying substrate. Reproduced from McKay et al. [165].

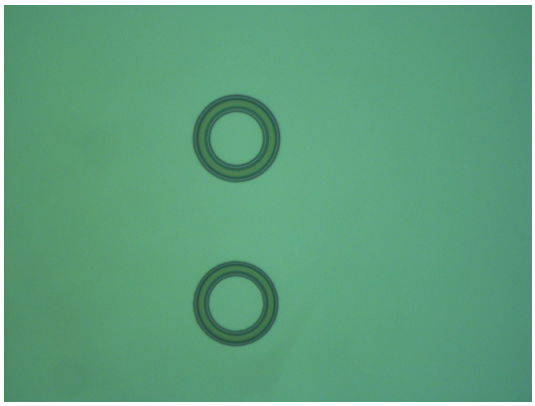

(a) No PEB

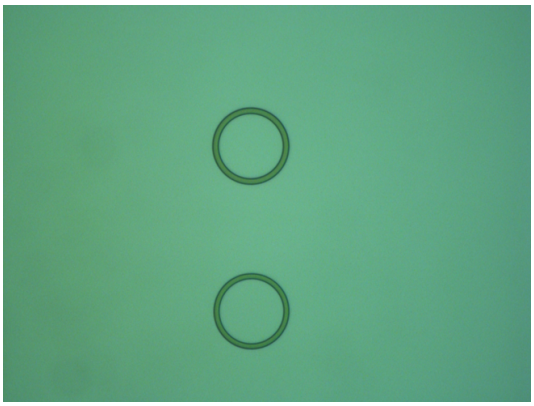

(b) 90 °C, hotplate 1 min

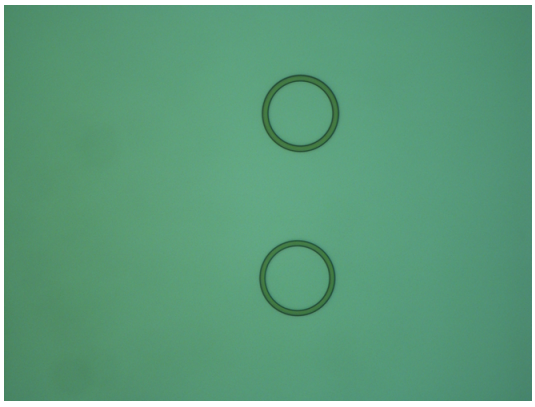

(c) 180 °C, oven 60 min

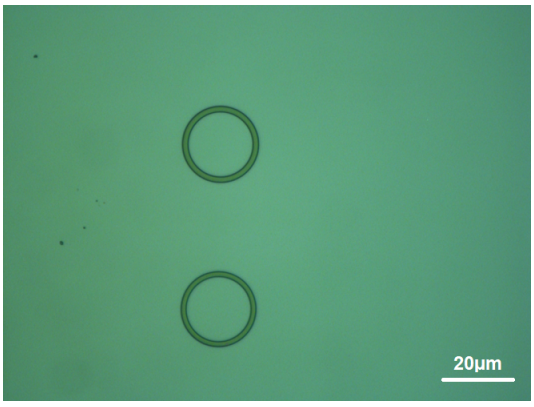

(d) 180 °C, hotplate 1 min

Figure 3.6: Optical micrographs of 10 µm-radius microrings fabricated in HSQ using different PEB conditions. All images taken at 1000x magnification and using identical bright-field illumination.

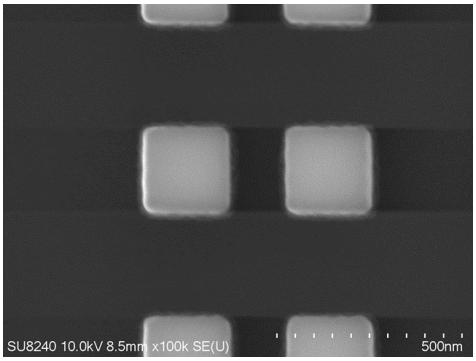

(a) No PEB

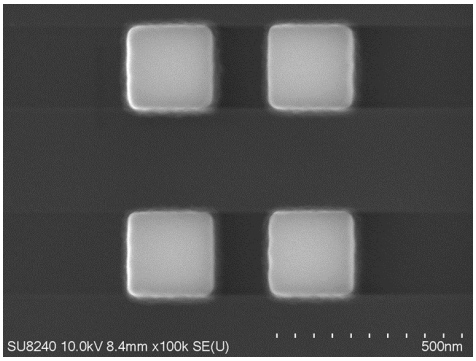

(b) 90 °C, hotplate 1 min

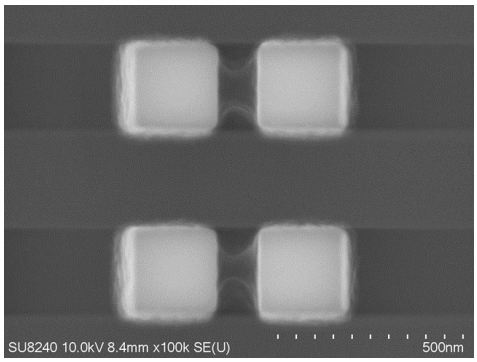

(c) 180 °C, oven 60 min

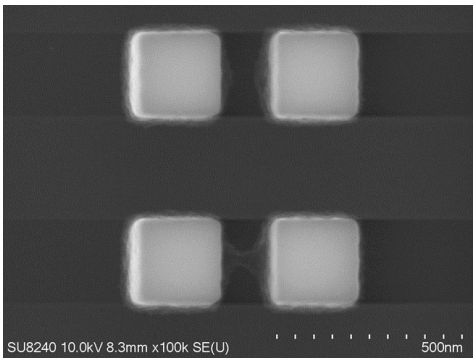

(d) 180 °C, hotplate 1 min

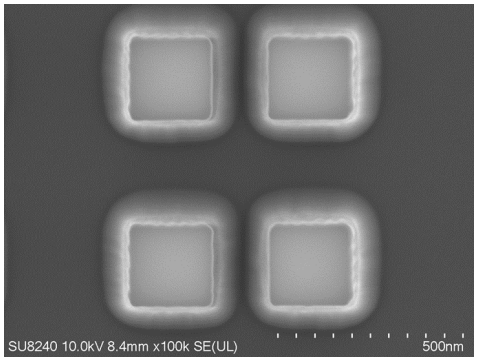

(e) No PEB, post etch

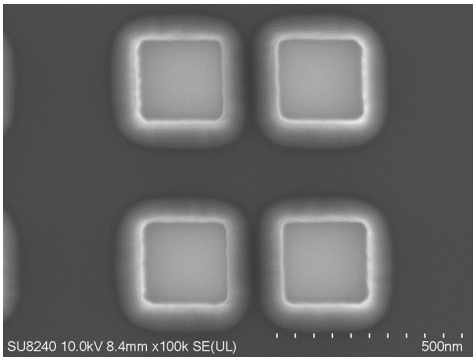

(f) 90 °C, hotplate 1 min, post etch

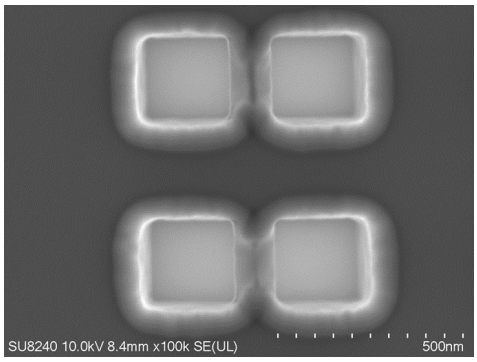

(g) 180 °C, oven 60 min, post etch

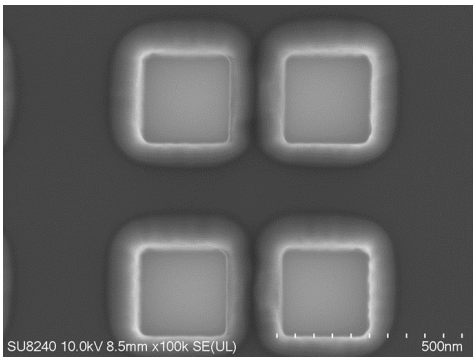

(h) 180 °C, hotplate 1 min, post etch

Figure 3.7: Scanning electron micrographs of 200 nm × 200 nm squares fabricated in HSQ using different PEB conditions. Images (e) – (h) show samples which have been etched for 5 min in BCl_3 plasma.

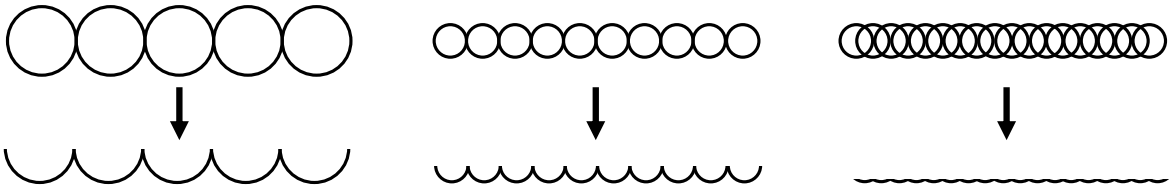

Figure 3.8: Effect of EBL spot size and step size on lithographic line-edge roughness. A large beam spot (a) may produce greater line-edge roughness than a small spot (b). Reducing step size can considerably reduce roughness (c).

Table 3.4: Beam currents and spot sizes on Raith EBPG5200 EBL tool.

Beam current (nA)	Theoretical spot size (nm)	Suggested beam step size (nm)	Measured spot size (nm)	Write time for standard test pattern (min)
1	4	2	38	11:16
2	6	4	36	8:33
4	9	5	38	5:16

Based on these results, I used a 180 °C, 1 min hotplate bake for subsequent fabrication runs. The choice to maintain a bake temperature of 180 °C, when a bake at 90 °C produced better feature resolution, was pragmatic as much as scientific. Feature resolution is not a concern when exposing waveguides, which must be well-spaced to avoid cross-coupling; given that etch selectivity results were based on processing of films baked at a higher temperature—with a higher degree of crosslinking, and therefore etch-resistance—and sidewall roughness after etching appeared similar for both films, I retained the higher bake temperature.

Beam Current and Step Size

Altering the electron beam current changes the spot size of the beam. Given that the beam shape is Gaussian and that shapes are exposed as a series of shots, the beam current can affect LER (Figure 3.8). In theory, the smallest possible spot size (lowest beam current) can produce the smoothest lithography; however, decreasing beam current can drastically increase exposure time [221].

Decreasing the beam step size can also produce smoother lines, reducing the effect of the Gaussian beam shape as shown in Figure 3.8. Whilst the standard HSQ process flow uses a beam step size of 4 nm, I used a beam step size of 1 nm in optimised lithography flows.

To determine the effect of beam current on LER, I exposed three versions of a test pattern on the same sample, using beam currents of 1 nA, 2 nA and 4 nA. All other parameters were identical (see Table B.4 for process flow) and the smallest possible beam step size (1 nm) was used when exposing all patterns. In this instance, I wrote the pattern on an alumina layer, and used scanning electron microscopy to image the patterns before and after 10 min BCl_3 etching. In the images taken following etching, approximately half of the structure is HSQ

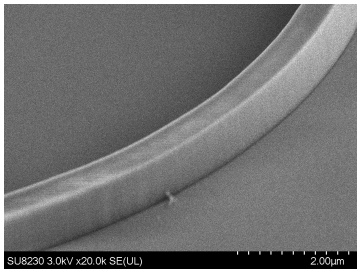

(a) 1nA beam spot, after development.

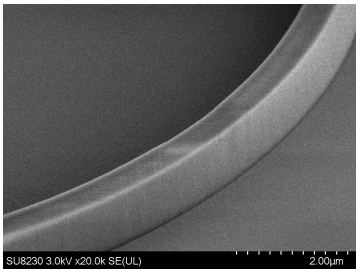

(b) 2nA beam spot, after development.

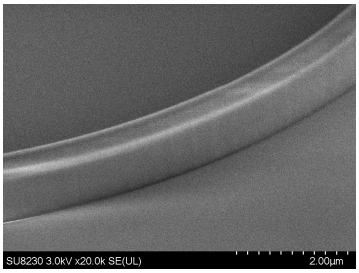

(c) 4nA beam spot, after development.

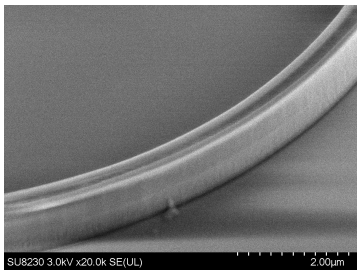

(d) 1nA beam spot, after etching.

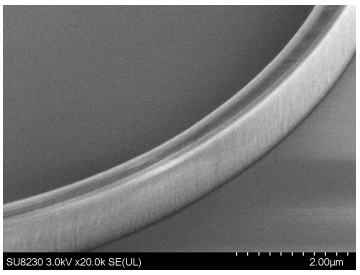

(e) 2nA beam spot, after etching.

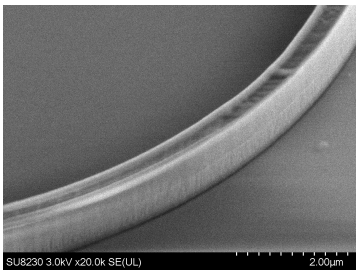

(f) 4nA beam spot, after etching.

Figure 3.9: Scanning electron micrographs of curved waveguides exposed using different beam spot sizes before and after etching. All waveguides are written on a single sample. All other parameters, including beam step size, remain identical.

and half is etched alumina. However, there is little discernible difference between structures written with different beam currents, likely because whilst the theoretical spot size increases with beam current, the measured spot size does not (Table 3.4). A further consideration is that, at higher beam currents, the minimum exposure dose is higher—therefore, the optimised process flow uses a 2 nA beam current and a 1 nm beam step size, balancing write time and the capacity to write structures using a low exposure dose.

Multipass Lithography

In a typical EBL process, the beam scan follows a pattern which allows it to move efficiently across the substrate. Movement between fields and subfields takes processing time and is generally minimised. *Multipass* EBL writing instead exposes each field multiple times, splitting the exposure dose equally between the passes. Whilst exposure time remains the same, the additional stage movement time does marginally increase overall EBL tool utilisation time; however, the technique can have substantial benefits in reducing LER and the impact of beam placement errors [236, 237].

Dose variability occurs as a result of *shot noise*—the variation in dose which arises from marginally different numbers of electrons incident on the sample with each beam step. Splitting the dose over multiple passes reduces the probability of any given shot receiving

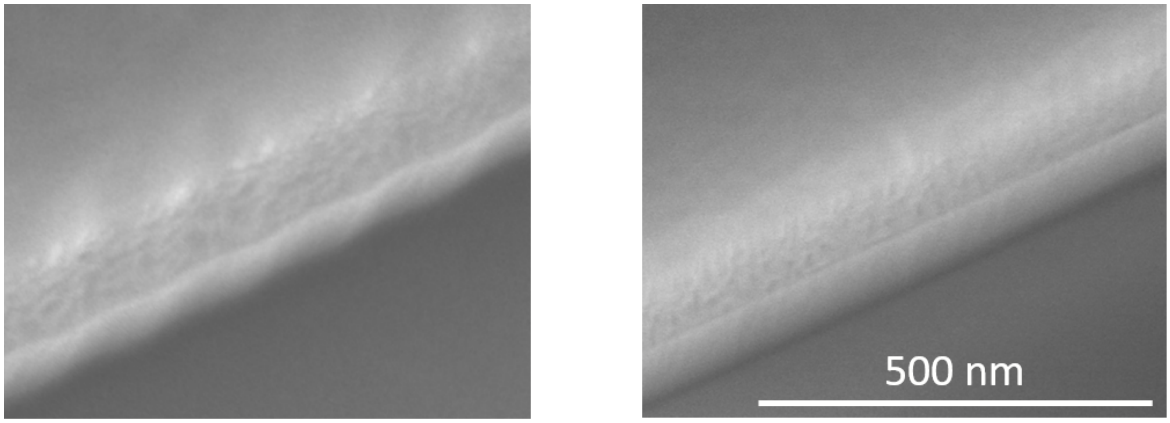

(a) Exposed using 1-pass EBL. (b) Exposed using 4-pass EBL.

Figure 3.10: Scanning electron micrographs showing sidewall roughness in SiN waveguides with HSQ masks exposed using single-pass and multipass exposure patterns. Reproduced from McKay et al. [165]; ©2023, Society of Photo-Optical Instrumentation Engineers (SPIE).

a substantially lower or higher dose than the target [236]. It also improves heat and charge dissipation, which reduces proximity effects [237].

Over the course of writing, there may be small (nanometre-level) variations in the position of the stage or the beam; this is frequently observed as *stitching errors*, where adjacent fields are misaligned. In this work, I applied a *subfield offset*: on each pass, subfield boundaries are shifted by half the width of a subfield. This reduces the impact of stitching errors [236].

Work performed by Natale Pruiti, using an identical HSQ masking process to etch SiN waveguides, demonstrates that implementing multipass lithography produces a qualitatively apparent improvement in sidewall roughness (Figure 3.10). SiN waveguides written using 4-pass EBL demonstrate a 46% reduction in optical loss at 635 nm, compared to waveguides written using single-pass lithography [165]. Therefore, the optimised exposure process uses 4-pass lithography.

3.2.7 Rectification of Stitching Errors

When testing waveguides fabricated using optimised lithography and etching processes, a further problem became apparent: fabricated waveguides (revisions 2 and 3—see Table 4.1) had high propagation losses. The lossy areas of the waveguide were visible to the naked eye, as bright spots at regular intervals of about 0.5 mm (Figure 3.11a). This corresponds to the size of an EBL write field, suggesting that the problem related to write field boundaries. Typically, problems occur at field boundaries as a result of small amounts of drift in stage positioning over time; if a long time elapses between writing of two adjacent fields, stage

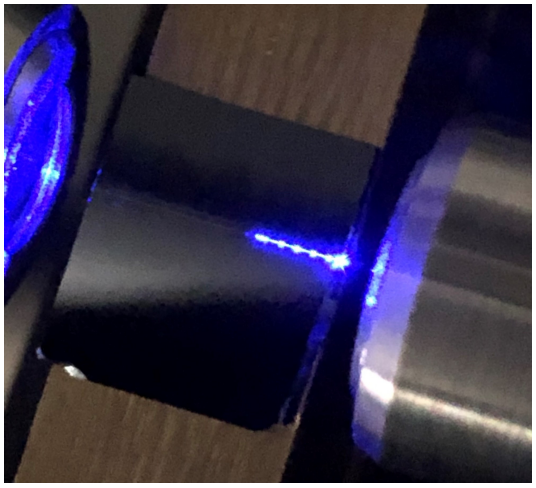

(a) Height offset not applied; light scatters at field boundaries. The abrupt transmission cut-off is caused by an unrelated mask defect.

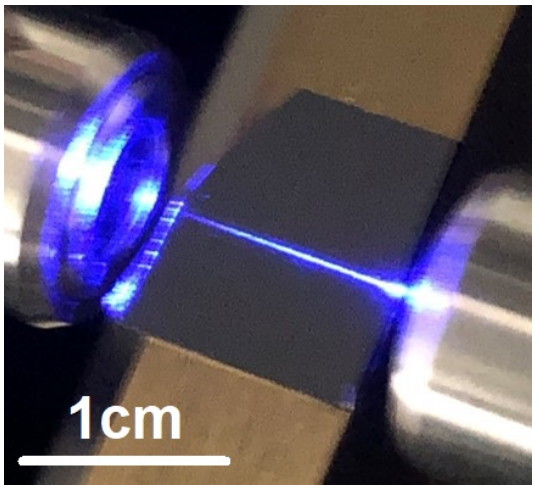

(b) Height offset applied; no scattering at field boundaries.

Figure 3.11: Photographs showing light scattering from etched alumina waveguides with and without application of a 1 µm height offset during EBL.

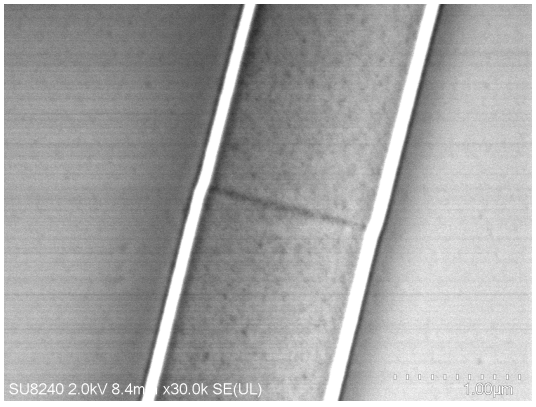

(a) Top view – scale bar edited to enhance contrast.

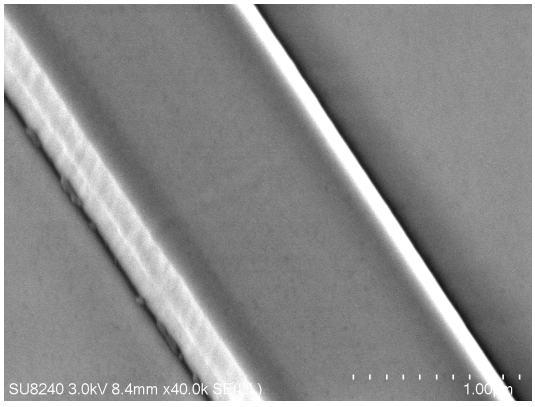

(b) Side view

Figure 3.12: Scanning electron micrographs of stitching defects in an etched waveguide. Mask was written without a height offset applied in EBL software.

drift may produce poor field-to-field alignment. However, even when all fields were written sequentially, similar amounts of scattering from all field boundaries could be observed.

SEM inspection of the waveguides showed a consistent but small gap between fields (Figure 3.12), suggesting that the field scaling was wrong. We hypothesised that the stage height was too high, scaling down the field size. The EBL system determines sample height using a glancing angle laser height meter with a wavelength of 700 nm, and adjusts stage height accordingly. Vernier-matching tests at a range of heights established that the height meter misreads sample heights by $\sim$1 µm—the combined height of the 400 nm alumina layer and 700 nm HSQ resist layer. This erroneous height calibration produced the 25 nm gap between fields seen in Figure 3.12; applying a 1 µm height offset eliminated the gap. I wrote test waveguides using the same exposure conditions, with and without the offset (Figure 3.11): no scattering at field boundaries (right) is apparent once the height offset is applied. The effect of the field scaling error is apparent when considering propagation loss results (Section 4.5)—waveguides with a height offset applied showed lower losses, especially for optical modes which have a high level of interaction with the areas most affected by the stitching defects.

3.2.8 Finalised Process

The finalised EBL process (Table 3.5) demonstrates lower LER than the initial conditions (Figure 3.13). Whilst the difference is evident considering the exposed mask (subfigures (a) and (b)), it is most evident after etching, where edges are smoother and the increased contrast is evident in the mask verticality.

Directly measuring the impact of any single process change on waveguide propagation loss (i.e., by fabricating waveguides using each process condition) would require a prohibitive amount of time and material. However, 400 nm-thick waveguides produced using the initial process were too lossy to measure, whereas 400 nm-thick waveguides produced using the final process demonstrated losses on the order of 1 dB/cm at 450 nm (Section 4.5).

This process has been optimised for use on low-loss waveguides—in other words, the process is designed to produce relatively high aspect ratio standalone structures with low sidewall roughness, where concerns regarding resolution can be ignored.

Table 3.5: Initial and final EBL conditions used to fabricate HSQ masks for low-loss waveguides.

Parameter	Initial	Final
Exposure dose ($\mu C\,cm^{-2}$)	1400	1450
Beam current (nA)	2	2
Beam step size (nm)	4	1
Number of passes	1	4
Height offset (μm)	0	1
Write order	Optimised	Sequential

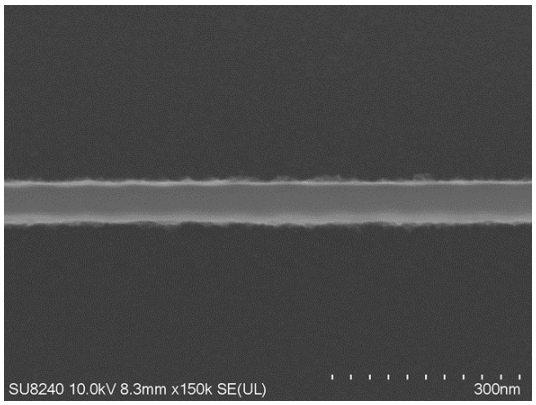

(a) Original lithography conditions: higher LER and 'foot' behaviour (residue at bottom edge of mask)

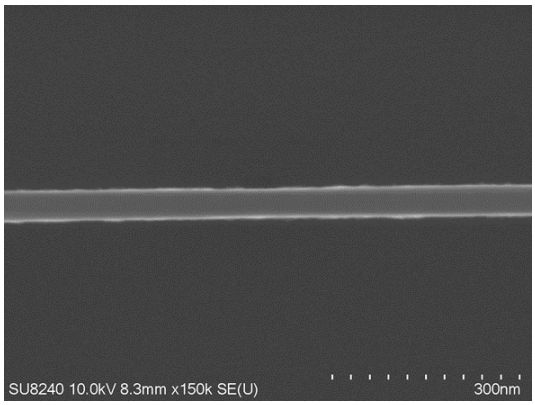

(b) Optimised lithography conditions: reduced LER and no 'foot'

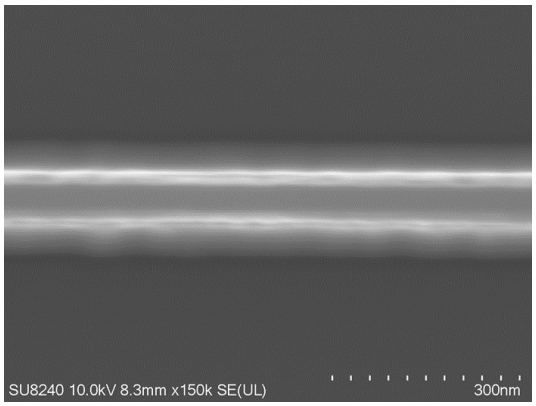

(c) Original lithography conditions, post etch: LER and 'foot' exacerbated

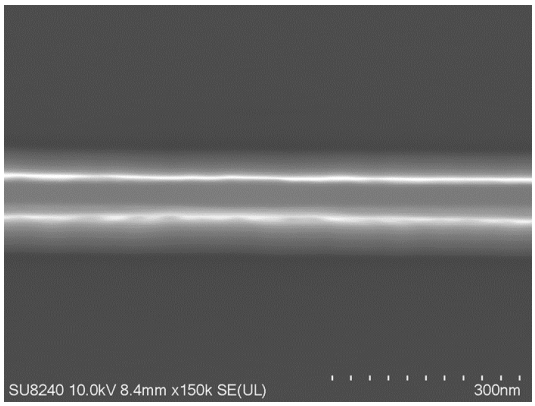

(d) Optimised lithography conditions, post etch: reduced LER

Figure 3.13: Top-down scanning electron micrographs showing 100 nm lines exposed in HSQ, using original (1400 $\mu C/cm^2$, 2 nA, 1-pass) and optimised (1450 $\mu C/cm^2$, 1 nA, 4-pass) lithography processes before and after etch.

3.3 Etching

Waveguides are defined by depositing a continuous layer of the core material, masking off the waveguide geometry, then etching away any material which is not protected by the mask. Typically, waveguide etching processes are performed using a *plasma etch*, in which gases in the etch chamber are dissociated to form a plasma containing highly energetic species which react with both mask and substrate. There are a wide variety of parameters which can be tuned to control the active species within the plasma and their directionality, thereby defining etch selectivity and etched structures' sidewall angle and roughness.

A high quality etching process for low-loss waveguides must

- Have adequate mask/substrate etch selectivity for etching to the target thickness
- Produce smooth sidewalls—not introducing roughness through mask erosion or direct damage to the substrate (minimising interface losses)
- Not produce any non-volatile byproducts which excessively inhibit the latter stages of the etch reaction
- Not produce any byproducts which are not removable, or which degrade the performance of the finished device (minimising extrinsic absorption)
- Some applications, such as MRRs, also require high sidewall verticality. For the devices described in this work, sidewall verticality is not critical, but this toolkit aims to provide a suite of processes which can be used to fabricate a range of devices.

Development of a plasma etch typically requires a compromise between requirements—for example, it may not be possible to achieve extremely smooth sidewalls at the same time as extremely vertical sidewalls, or high selectivity. In this work, HSQ was a favourable choice of mask material because of the Optoelectronics group's existing process capability and because of the resist's inherent capacity to create an extremely smooth and vertical mask using a single-step process. However, alumina has relatively poor selectivity to silica-like masks and HSQ cannot be spun thicker than 1 µm, so optimising for etch selectivity is critical. When fabricating waveguides which have low propagation losses in the visible region (Chapter 4), minimising sidewall roughness is most critical. For supercontinuum-generating waveguides (Chapter 5), dispersion engineering requires thick waveguides, and achieving an etch selectivity of $\sim$0.8 becomes critical.

3.3.1 Theory of Plasma Etching

A plasma etching reactor consists of a vacuum chamber with parallel upper and lower electrodes. A gas, or mixture of gases, is introduced to the chamber, and an RF field is applied across the electrodes, creating a weakly ionised plasma in which a small proportion

of the molecular gaseous species are ionised to produce ions, radicals, and free electrons. The substrate sits at the lower electrode within the chamber. The acceleration of highly mobile electrons towards the substrate generates a *DC bias* across the electrodes, and a plasma sheath—an electron-depleted area with no plasma—forms above the substrate [238].

Radicals are the primary reactive species within a weakly ionised plasma. Since they have no net charge, they are not attracted towards the substrate, but will reach the substrate through diffusion then adsorb and chemically react with it (*chemical etching* or *plasma etching*). This process is typically isotropic. Ionised species, present in lower concentrations, are attracted towards the substrate as a result of the DC bias (despite the oscillating RF field). Unreactive ions—such as Ar^+—etch the substrate through sputtering, or *physical etching*. Ion etching is typically anisotropic; the degree of anisotropy is related to the *mean free path*, the typical distance a species can travel before a collision. *Reactive ion etching* describes a system in which the balance between plasma and sputter etching is mediated to produce an etch with desirable properties. Given the low density of ions relative to radicals, ions are rarely the primary species reacting with the substrate; however, they can mediate the adsorption of radical species, the etching reaction, or the removal of etch byproducts, and therefore establish a particular etching profile [239].

In an inductively coupled plasma (ICP) RIE system, a coil around the chamber creates an additional oscillating RF field in which charged species are not induced to move towards either electrode. This allows creation of a high-density plasma at a low chamber pressure, increasing etch rate whilst allowing a long mean free path. The ICP power (ion density) and platen power (ion acceleration) can be controlled separately, providing additional degrees of freedom when designing etch processes [239].

In this instance, selectivity is the critical parameter requiring optimisation. Selectivity is defined by the etch rates of the mask and substrate materials, and is therefore controlled by modifying etch chemistry, as well as parameters which alter the balance between chemical and physical etching. Sidewall roughness has a lithographic component (Section 3.2.6): roughness from the mask can be transferred into the substrate. During etching, additional roughness arises from ion bombardment, and can also arise from *micromasking*, where etch byproducts do not desorb from the substrate and inhibit further etching. For this application, verticality is desirable, but must not be prioritised over sidewall roughness. The ratio of chemical to physical etching will affect sidewall verticality, as will the chamber pressure.

3.3.2 Etch Development

Alumina readily etches in halogen-containing plasmas (plasmas containing F, Cl, Br). Whilst fluorine chemistries produce high etch rates, non-volatile AlF_3 is produced as a byproduct;

therefore, chlorine or bromine chemistries are preferred despite lower etch rates [240]. As suggested by this variation in etch rate, which tracks with the relative chemical reactivity of these elements, halogen etching has a substantial chemical component, and much research has focused on how to maximise halogen radical concentration within the plasma [241, 242]. BCl_3 has emerged as a core component of etch chemistries, as Cl radicals etch alumina films [241], whilst BCl_xO_y etch byproducts are volatile and readily transported away from etched surfaces [122].

There is a large body of literature focusing on optimising the etch rate and selectivity of BCl_3-based alumina etches; however, alumina is more commonly used as a high-κ dielectric rather than as an optical material, and consequently optimisation is rarely performed with a focus on sidewall roughness or verticality [240, 241, 243].

As described in Section 1.4, there is a limited body of work on etched alumina waveguides. At the time of this optimisation work, only two published works discussed alumina etch optimisation for optics purposes: Bradley et al. (2007) used 5:2 BCl_3/HBr with a photoresist mask [122] and West et al. (2019) used 1:1 BCl_3/Ar with a plasma-enhanced chemical vapour deposition (PECVD) silica mask [4]. Recent work by Lin et al. [13, 132] has used an SiN hardmask with a BCl_3/Cl_2/Ar etch, although full process details have not been published. Ultimately, the only deep alumina etch remains the BCl_3/HBr protocol established by Kim et al. [244] and adapted by Bradley et al., which produces a shallow sidewall angle of 68° and relatively rough sidewalls which proved insufficient for visible optics [122].

I performed etch tests by exposing a standardised pattern featuring a variety of straight lines and curved shapes, designed to establish what sidewall roughness and resolution is attainable, in HSQ on a 400 nm-thick alumina film (see Table B.5 for full process flow). I used a Bruker Dektak XT stylus profilometer to measure the step height of the resist before etching and the combined step height of the resist and the etched substrate after etching. In each etch, I included an HSQ monitor piece (see Table B.6 for full process flow) consisting of cured HSQ on bare Si. This film thickness is easily measured using VASE, and in combination with step height data, can be used to calculate etch rate and selectivity.

I reproduced etches used in the literature to successfully fabricate alumina waveguides, 5:2 BCl_3/HBr [122], and 1:1 BCl_3/Ar [4] (Table 3.6). Upon exposure, HSQ forms a silica-like structure, so I anticipated that the HSQ mask would produce similar results to the silica mask used by West et al. Whilst the etch used by Bradley et al. showed high sidewall roughness with a photoresist mask, I anticipated that HSQ's higher etch resistance would improve both selectivity and roughness. I also included an etch recipe developed by Woo et al. [245], which included oxygen with the aim of increasing etch selectivity by binding reaction-inhibiting byproducts.

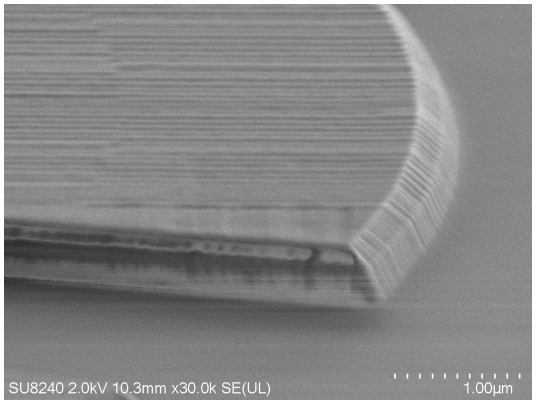

(a) 5:2 HBr/BCl_3 etch reproduced from Bradley et al. [122].

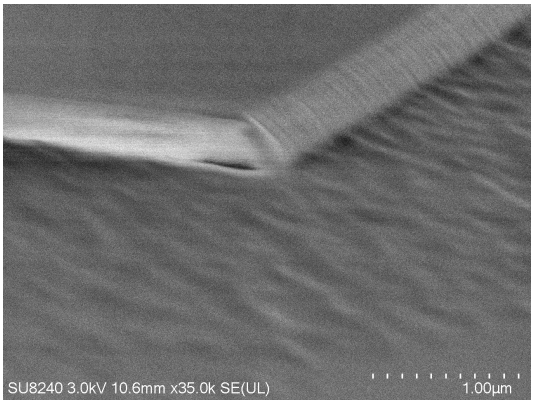

(b) 1:1 BCl_3/Ar etch reproduced from West et al. [4].

Figure 3.14: Scanning electron micrographs of alumina films processed using plasma etches which have been used for the fabrication of alumina waveguides. All structures used an HSQ mask.

The BCl_3/Ar etch demonstrated a selectivity of 0.71—whilst less than the rate associated with PECVD silica, this is to be expected given that HSQ is less dense and may not be fully crosslinked. Scanning electron micrographs (Figure 3.14b) show a shallow-angled sidewall. The HBr/BCl_3 also produced sidewalls with a relatively shallow angle, even after power optimisations to increase etch bias (Figure 3.14a). Sidewall roughness is clearly evident on both etches; however, this may be a result of underexposure of the mask (Section 3.2.6). The $O_2/Ar/BCl_3$ recipe etched the alumina extremely slowly, and also produced byproducts, likely involatile B_2O_3 [245], which visibly contaminated the sample and the carrier wafer. It is likely that the slow etch rate is a result of surface passivation with involatile B_2O_3.

Since none of these etches produced acceptable results, I developed an etch recipe based on an existing BCl_3 etch used by other (non-optics) JWNC users. Here, SEM imaging of the sidewall angle and roughness gave promising results. I maintained a low pressure of 2 mT throughout all etches, prioritising a long mean free path over a high etch rate. Limited data on varying the temperature of BCl_3-based alumina etching has not demonstrated a strong temperature dependence, so I maintained a temperature of 20 °C [240]. I varied ICP power and platen power to optimise selectivity and sidewall quality (Table 3.7). Increasing ICP

Table 3.6: Alumina etch selectivity tests using etches reproduced from published papers. Note that values used for HBr/BCl_3 etching are approximate as measurements were taken using a calibrated SEM, rather than direct measurement of sample and monitor piece.

Paper	Gases	ICP power (W)	Platen power (W)	HSQ etch rate (nm/s)	Al_2O_3 etch rate (nm/s)	Selectivity
Bradley et al. [122]	5:2 HBr/BCl_3	1000	25	∼2.00	∼4.00	∼0.50
Woo et al. [245]	3:6:14 $O_2/Ar/BCl_3$	700	50	0.44	0.15	0.34
West et al. [4]	1:1 BCl_3/Ar	1000	100	1.26	0.90	0.71

Table 3.7: Alumina etch selectivity tests using a BCl_3 etch recipe. Only the ICP and platen power are changed for each etch.

ICP power (W)	Platen power (W)	HSQ etch rate (nm/s)	Al_2O_3 etch rate (nm/s)	Selectivity	DC bias (-V)
1000	25	0.92	0.68	0.74	92
600	100	1.12	0.62	0.56	279
600	25	0.63	0.41	0.65	111
400	100	0.62	0.35	0.56	296
400	25	0.37	0.23	0.61	122
1400	15	1.02	0.80	0.79	71

power increases the density of reactive species, thereby increasing the etch rate of both alumina and HSQ. Increasing platen power increases the DC bias of the etch—increasing directionality, but with an adverse impact on selectivity.

SEM imaging of these samples showed a slight increase in sidewall roughness with increasing ICP power, with negligible difference in sidewall angle. I chose waveguide etch conditions to balance selectivity and roughness, and etched the first set of single-mode waveguides using pure BCl_3 with an ICP power of 600 W and a platen power of 25 W. This provided a sidewall angle of 80° (Figure 3.15) and a selectivity of 0.65—sufficient for etching of 400 nm-thick waveguides.

Waveguides for SCG require an etch depth of approximately 800 nm; given a maximum HSQ mask thickness of 1 µm, an etch selectivity of 0.8 is required. Using a high ICP power and low platen power minimises etch bias and increases selectivity, at the expense of sidewall angle and roughness; the highest selectivity attained was 0.79 (Figure 3.16). Etched sidewalls are vertical near the top and sloped closer to the bottom of the waveguide, which may indicate that charge buildup on the waveguide surface has inhibited physical etching, or that the etch mask demonstrated "foot" behaviour. The uneven etching of the mask may relate to a lithographic PEC error.

Figure 3.17 provides a comparison of the sidewall roughness generated by the two etching processes. Qualitatively, little difference is apparent; direct measurement, or extraction of roughness from optical loss measurements, would be required to establish a difference.

The etch selectivities attained in this work are certainly not the highest reported in the literature: approximately 1:1 selectivity, with comparable verticality, can be attained adding Ar to the gas mixture [4]. This addition enhances the physical component of the etching process—but this sputtering may adversely affect sidewall roughness. The addition of Cl_2 could increase the chemical component of etching and improve selectivity [241].

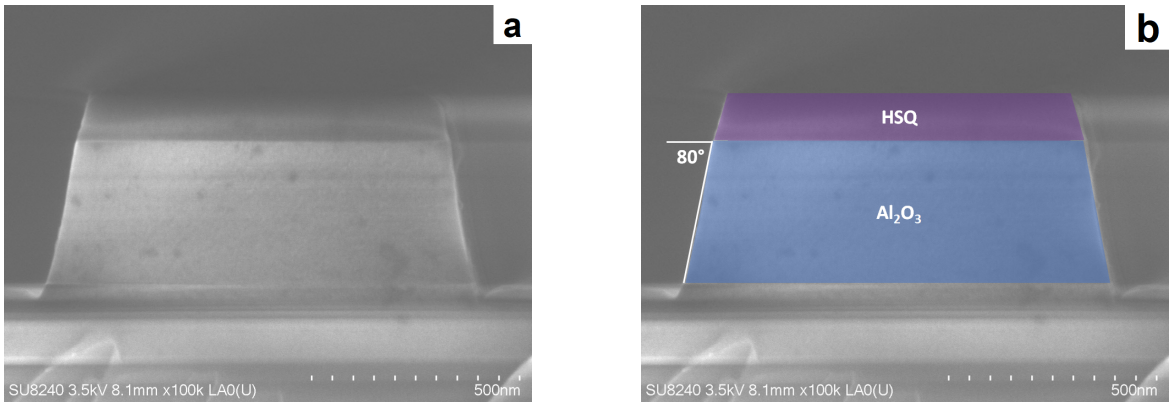

(a) (b)

Figure 3.15: Scanning electron micrographs showing cross-section of 400 nm-thick, 850 nm-wide alumina waveguide etched using optimised BCl_3 etch. Subfigure (b) provides a false-colour annotation to clarify layer thicknesses. Reproduced from McKay et al. [76].

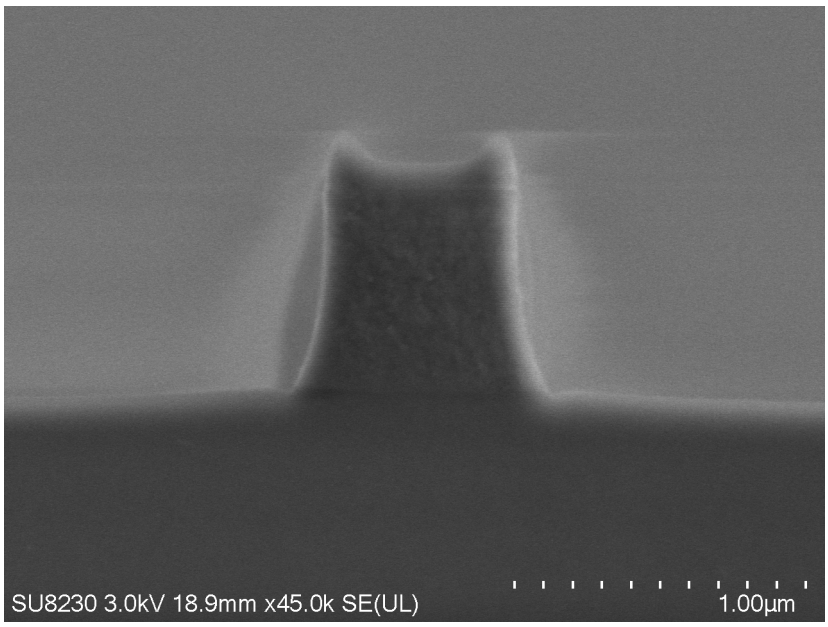

Figure 3.16: Scanning electron micrograph of the cross-section of a 700 nm-thick, 700 nm-wide supercontinuum-generating waveguide.

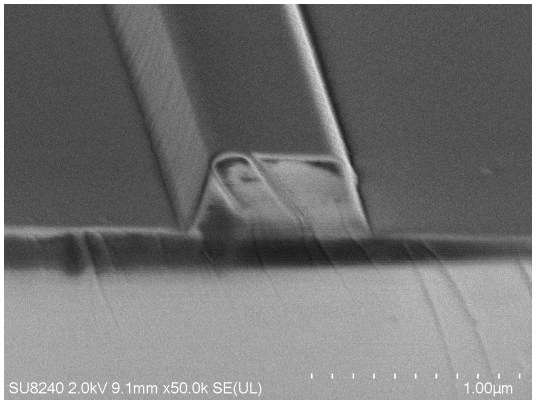

(a) Etch optimised for low-loss waveguides.

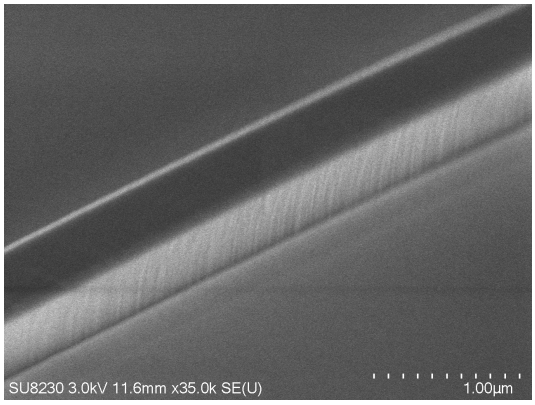

(b) Etch optimised for supercontinuum-generating waveguides.

Figure 3.17: Scanning electron micrographs showing sidewalls of alumina waveguides etched using BCl_3 plasma.

3.4 Cladding

Optical radiation remains in waveguides because the waveguide core has a higher refractive index than the surrounding cladding. Interface scattering is dependent on the refractive index difference between core and cladding (Section 1.2): a smaller refractive index difference reduces interface scattering. Therefore, an upper cladding layer is typically applied to minimise waveguide losses. Propagating light is not completely confined to the waveguide core; some of the light travels in the cladding (Section 4.1). Absorption and scattering losses can arise from cladding material, and desirable optical properties for a cladding layer are similar to those for the waveguide core.

In this work, I deposited alumina films on an Si carrier with a pre-grown silica cladding layer, and used the Optoelectronics group's standard HSQ cladding process to apply an upper cladding layer [246, 247]. The bottom oxide layer is grown using a wet thermal oxidation process, in which the silicon carrier is oxidised to produce an amorphous silica film. The bandgap of such films is considerably wider than that of the alumina itself; in fact, they are too wide to be measured directly. An estimate derived from photoconductivity measurements places the optical bandgap of amorphous silica at 133 nm [248].

I applied an upper cladding to waveguides by spinning an HSQ layer at 2000 rpm, producing a layer approximately 700 nm thick. The HSQ is cured by baking at 180 °C in a convection oven for 72 h. Thermally driven changes in HSQ films include solvent evaporation, film crosslinking, and (associated) film densification. Thermogravimetric analysis (TGA) by Siew et al. indicates that a baking temperature of 180 °C is sufficient to drive off the MIBK solvent in which monomers are dissolved, but is not sufficient to fully crosslink or densify films [249].

Thermal crosslinking of HSQ produces changes in bonding composition which are readily probed using FTIR spectroscopy. Fully-crosslinked HSQ is silica-like, consisting of Si-O-Si bonds in a *network* structure. Monomers contain Si-O-Si bonds in a *cage* structure; localised differences in bonding mean that cage and network bonds absorb at different wavenumbers and are easily distinguished. Films which are not fully bonded also contain Si-H and Si-OH bonds. Biryukova et al. cured films at 180 °C, and established that the cage structure predominated and that Si-H and Si-OH bonds were present [246]. Yang and Chen observed scission of Si-H bonds at temperatures above 300 °C [222], whilst Siew at al.'s TGA places the onset of Si-H scission at 350 °C [249]. Extended curing at 180 °C may elevate the Si-H scission temperature: Biryukova et al. required an anneal temperature of 550 °C to eliminate Si-H bonding when a low-temperature cure was used prior to anneal. They also investigated annealing at considerably higher temperatures than many other groups, and demonstrated ongoing structural changes at temperatures up to 1150 °C [246].

HSQ films also demonstrate temperature-driven densification, linked to bonding changes—refractive index is a good proxy for HSQ film density. Biryukova et al. observed changes in refractive index upon annealing. Films cured at 180 °C had a refractive index of 1.39 at a wavelength of 1550 nm; annealing continued to increase refractive index through the entire temperature range the authors explored, and films annealed at 1150 °C had a refractive index of 1.45 [246]. Several other groups report changes in film density at temperatures of 300 °C to 400 °C [222, 249]; Liou and Pretzer report thermally-driven changes in porosity and density in this range, continuing at temperatures as high at 800 °C [250].

There is limited data linking these structural changes to propagation loss: as for alumina, absorption in the IR does not necessarily mean that absorption in the visible region will occur. Changes in film density are a more concrete concern, since porosity in films will result in bulk scattering loss. Analysis of HSQ absorption loss poses the same challenges as analysis of alumina film losses: VASE lacks the sensitivity required to establish that films will allow low-loss propagation, so data on optical quality must be derived from transmission measurements or from demonstrating that waveguides function well. No published results using a 180 °C curing process demonstrate low-loss light propagation below 660 nm [247]. Chang et al. performed transmission measurement of HSQ films annealed to 300 °C which demonstrated $> 90\%$ transmission above 300 nm [214]—however, these films are not necessarily comparable to those cured at 180 °C.

The PEALD films deposited in this work cracked at 800 °C (Section 2.8), but not at 600 °C. I designed a relatively conservative annealing process at a temperature of $\sim$500 °C, using the same ramp rates and process parameters described in Section 2.8. I fabricated 400 nm-thick waveguides using the standard low-loss waveguide process flow (Table B.1), clad them using HSQ cured at 180 °C for 72 h, and tested losses. I then annealed the waveguides in a

Jipelec JetFirst RTA system, using a maximum ramp rate of 5 °C s^{-1} to reduce the risk of film cracking, and holding the maximum temperature of 500 °C for 1 min. However, the HSQ film cracked at the edges (Figure 3.18), rendering the devices unusable. This cracking pattern is similar to those observed when extremely thick HSQ films are spun, and therefore may relate to film stress. Given the limited extent of the cracking, it may be possible to anneal samples prior to facet cleaving, and then cleave off the damaged areas. However, given that waveguide losses in samples clad with HSQ baked at 180 °C demonstrated low losses across the visible region (Section 4.5), I did not pursue further work to reduce cladding losses.

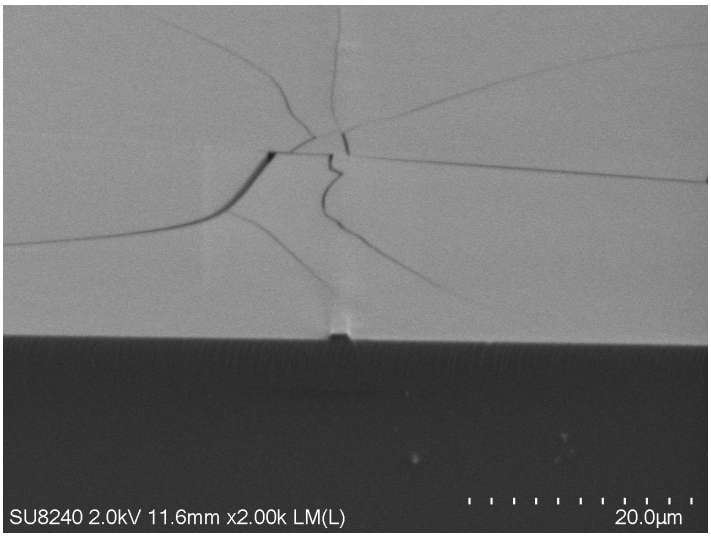

Figure 3.18: Scanning electron micrograph showing an etched alumina waveguide with an HSQ cladding. Following annealing at 500 °C, the cladding has cracked near the waveguide facet.

3.5 Cleaving

The devices fabricated in this work, like the majority of devices fabricated using wide-bandgap dielectrics, do not contain native light sources: light must be coupled in from an external source. The waveguides in this work are designed to be coupled to sources using end-fire coupling, in which a laser source is focused using a lens to form a tight spot which is aligned to the end of the waveguide (see Appendix C). End-fire coupling is an attractive coupling method because it presents few additional design or fabrication requirements, and any allowed propagation wavelength can be coupled into and out of the waveguide—this is critical for nonlinear optics experiments.

To facilitate end-fire coupling, the waveguide's cross-section must be exposed at the waveguide ends, known as waveguide *facets*. In order to minimise coupling loss, waveguide facets must be free from defects—surface roughness, scratches, or debris—which will cause light scattering.

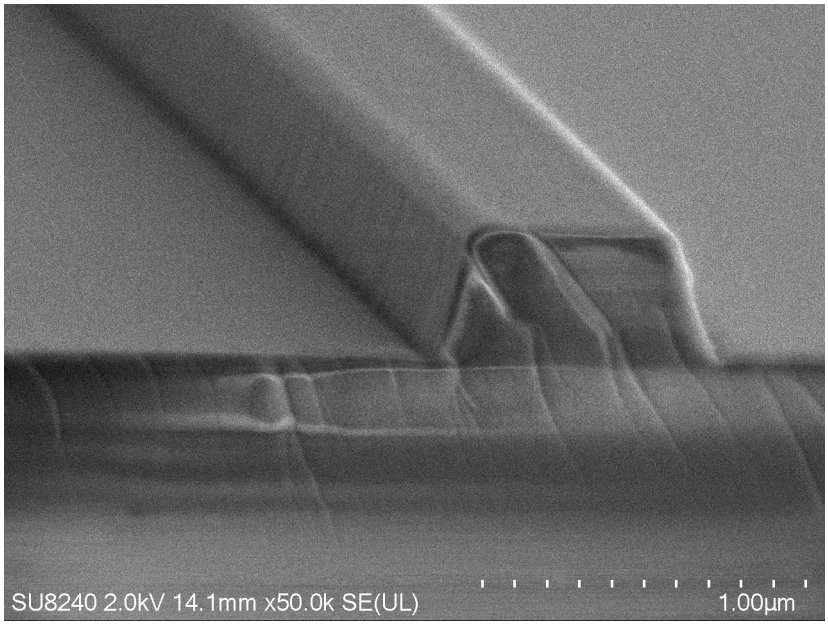

Figure 3.19: Scanning electron micrograph of a poorly cleaved waveguide facet. Striations, which adversely affect coupling losses, are evident.

Since the Si carrier wafer is a relatively brittle single-crystal structure, *cleaved* facets can be produced using a scribe-and-break technique, in which a diamond cleaving tip is aligned to the crystal plane of the substrate, a scribe (notch) is inscribed at one end of the sample, along the target break line, and the scribe is propagated along the crystal plane by applying pressure.

Initial waveguide designs (revisions 1-5) placed all waveguide widths on a single 18 mm × 18 mm chip. Of 14 facets cleaved using this chip geometry, only 7 cleaved correctly; the rest were affected by cleave *runoff*, where the cleave does not follow the crystal plane for its entire length. This affected waveguide functionality in different ways:

- Where the runoff was oriented towards the edge of the chip, the waveguides were not cleaved through, facets were not exposed, and light could not be coupled into the chip.
- Where the runoff was oriented towards the centre of the chip, waveguide bends were cleaved through, preventing light from reaching the end of the waveguide.
- In instances where there was severe runoff but facets were still cleaved in an appropriate location, the poor optical quality of the facets resulted in high coupling losses (Figure 3.19).

I redesigned the chip geometry to improve cleaving yield: revision 6 retained the revision 5 waveguide design, but split onto three 14 mm × 20 mm chips. When cleaving 18 mm × 18 mm chips, runout typically occurred after the cleave had propagated along the crystal plane for at least 12 mm; reducing the cleave distance to $<$12 mm prevents this occurrence. The addition of an extra 1 mm at each side of the chip increases the area to which pressure can be applied

during cleaving. The considerable increase in material and processing requirements are justified by the yield increase: of 26 facets cleaved using this chip geometry, 24 cleaved correctly. Typical losses are on the order of 5 dB/facet (Section 4.5.3).

Scribe-and-break cleaving is inherently unreliable, especially on Si substrates and when using manual processing techniques. Using grating couplers, which could be fabricated on the same mask layer as the alumina waveguides, could provide a more consistent coupling solution [46]; however, such couplers are only functional over a narrow wavelength range, and add design and fabrication complexity. Alternatively, using a dicing saw to cut facets, then polishing the facets to optical quality, could provide high-precision facet definition with a substantially lower chance of damaging the chip [251]. Etching facets, followed by a cleaving or dicing process, could also provide a route to consistent facet definition [252]. Given that coupling loss is also impacted by the roughness of the facet, the concerns associated with sidewall roughness also apply to facet roughness.

3.6 Conclusion

This chapter presents a suite of lithography and etching processes which can be used to fabricate low-loss alumina waveguides (demonstrated in Chapter 4). I have performed minimal optimisation of cleaving and cladding processes, but have considered their impact on optical loss and device yield. Critical optimisations include

- Establishing the maximum thickness of HSQ which can be spun without cracking (1 μm), which sets a lower bound on acceptable etch selectivity
- Ensuring HSQ exposure dose is sufficient for the mask to withstand etching
- Using multipass lithography to reduce mask sidewall roughness and associated propagation losses
- Applying a height offset to EBL exposures to eliminate field scaling errors
- Developing a BCl_3 etch and establishing the relationship between ICP and platen powers and selectivity
- Designing chip geometries which allow for more reliable manual cleaving

The optimisation work I have performed has considerable limitations in both scope and robustness. In terms of robustness, few individual experiments detailed in this chapter directly quantify any property that is linked to optical loss: lithography and etch roughness optimisations relied on the use of scanning electron micrographs which only offer a qualitative evaluation of sidewall roughness. A more convincing, but time-consuming, test would be to fabricate a set of waveguides using each process condition, and compare propagation losses.

Fabricating waveguides in a range of widths and modelling sidewall interactions would even provide a value for sidewall roughness.

Given the vast number of lithography and etch parameters which could be changed, many of which are interlinked, the aim was to produce a process which is *good enough* to produce low-loss waveguides for optics experiments; the finalised process flow certainly cannot be considered fully optimised.

Working with an HSQ mask and designing a process from scratch, optimisation (or even omission) of the softbake step may produce a reduction in the fundamental roughness of the mask layer. All subsequent steps in the lithography process would require re-evaluation. Quantifying optical losses and/or LER for each variable choice would allow for more robust and well-evidenced decision-making.

Etch optimisation was limited: whilst I adjusted ICP and platen powers to improve selectivity and sidewall roughness, I ignored parameters such as gas flow rate, chamber pressure, platen temperature, and the use of additional gases to enhance physical or chemical etching. An alternative route to improving selectivity is to strengthen the mask material. I explored annealing processes which may reduce mask erosion during etching, but did not explore the effect of annealing on HSQ roughness, the effect of annealing on alumina etch rate, or the effect of annealing on optical loss in alumina. All of these factors may preclude the use of mask-annealing processes.

Further opportunities for reducing optical loss lie in cladding and cleaving processes. The current HSQ cladding process, with a low-temperature cure, does not fully densify or crosslink the material. Whilst the effect of heat treatment on HSQ losses in the visible region is not apparent, thermally driven densification of HSQ cladding is likely to reduce scattering losses. Alternative cladding materials—such as silica deposited using low-pressure chemical vapour deposition (LPCVD) [253] or tetraethoxysilane (TEOS)-PECVD [254]—have demonstrated low losses in low-confinement waveguides for the visible region. Alternative facet definition processes, such as dicing or etching, may reduce coupling losses, although a primary consideration is the yield and consistency of the facet definition process.

Ultimately, the suitability of the process flow is decided by whether it can be used to create optical components with low losses, especially propagation losses. In Chapter 4, I describe the fabrication and testing of eight sets of waveguides using the lower-selectivity etching process at visible wavelengths. I achieve propagation losses as low as 0.8 dB/cm at 450 nm—the lowest losses attained in high-confinement ridge waveguides in the blue—whilst in Chapter 5 I demonstrate that UV light generated through nonlinear processes can be guided.

Chapter 4

Single-Mode Waveguides

Chapter 1 described the need for PICs operating in the visible region for a range of applications in optical timing, computing, and sensing. It also introduced waveguides as the base unit for a range of PIC components, and described the challenges of managing absorption and scattering loss when fabricating waveguides for the visible region. Chapter 2 analysed the suitability of PEALD alumina films for visible waveguiding, whilst Chapter 3 outlined a processing toolkit for creating high-quality waveguides using PEALD alumina. In Chapters 2 and 3, I related processing conditions to predictors of optical quality such as roughness and bonding composition, but did not present results which validate the effect of each optimisation on optical loss.

The time and material cost of fabricating waveguides to test the effect of every process modification on optical loss is prohibitive. However, six iterations of waveguides, exposed using different lithography parameters and mask designs, were required to establish processing conditions which produce low propagation losses in the visible region; in this chapter, I describe the evolution of the finalised process flow, and then evaluate its repeatability.

Section 4.1 provides sufficient theory to justify the design requirements set out in Section 4.2. Whilst fabrication details are covered in Chapter 3, Section 4.3 provides a brief overview of the fabrication differences between waveguide revisions. Section 4.4 describes the optical apparatus used to test waveguides' optical losses, whilst Section 4.5 provides results from loss testing of high-confinement waveguides. In Section 4.6 I discuss propagation losses in a limited run of low-confinement waveguides, which are more comparable to alumina waveguides previously reported in the literature. Section 4.7 situates the results in the wider context of visible waveguiding, as well as considering future improvements to mask, process, and metrology design.

4.1 Theory

A full picture of loss in waveguides requires an understanding of how the confinement of light produces optical modes, and how these modes interact with their surroundings.

Waves propagate within waveguides as *guided modes*[1], which have consistent electric and magnetic field profiles transverse to the propagation direction. Figure 4.1 illustrates the definitions used throughout this section: waves propagate in the z-direction and have consistent electric and magnetic field profiles in the (x, y) plane. Modal confinement is defined by the refractive index profile across a cross-section of the waveguide; assuming that the waveguide material is a homogenous, optically isotropic dielectric and that the cross-sectional dimensions are consistent, we can characterise a waveguide mode by considering the transverse field pattern of a single cross-section. The resulting amplitude and polarisation patterns are constant along the propagation direction.

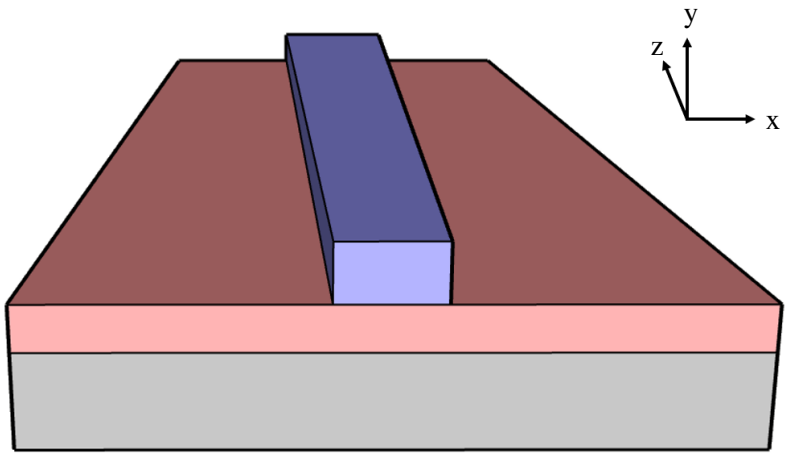

Figure 4.1: Ridge waveguide with no upper cladding, showing x and y (transverse) and z (propagation) directions. Purple areas represent the core material; red, the bottom cladding; grey, the carrier.

Using a ray-optic approximation, light can be visualised as a plane wave propagating through a homogeneous medium, with wave fronts moving in the direction of the 'ray'. Electric and magnetic fields are mutually perpendicular, and both are perpendicular to the propagation direction. For a mode with index v, with electric mode field profile $E_v(x, y)$, the electric field $E_v(r, t)$ is given by

$$E_v(r, t) = E_v(x, y) \exp -i(\beta_v z - \omega t) \tag{4.1}$$

Similarly, given a magnetic field mode profile of magnetic field profile $H_v(x, y)$, the magnetic field is given by

$$H_v(r, t) = H_v(x, y) \exp -i(\beta_v z - \omega t) \tag{4.2}$$

[1]Whilst other types of modes, such as radiation modes, exist, I use 'modes' alone to mean 'guided modes'.

ω gives the *angular frequency* of the mode, the rate of phase change over time. β is the modal propagation constant, the rate of phase change over distance. Altogether, the exponent defines the phase of the wave [255].

Transverse modes are quantised along the x and y axes (as defined in Figure 4.1): the number and types of modes which are allowed within a waveguide depend on the wavelength of the propagating light and the geometry and refractive index of the core and cladding. Figure 4.2 provides a simple example of a slab waveguide where the core has a higher refractive index (n_1) than the cladding (refractive index n_2). For clarity, only two dimensions are shown: in the following discussion, z defines the propagation direction and x the transverse direction.

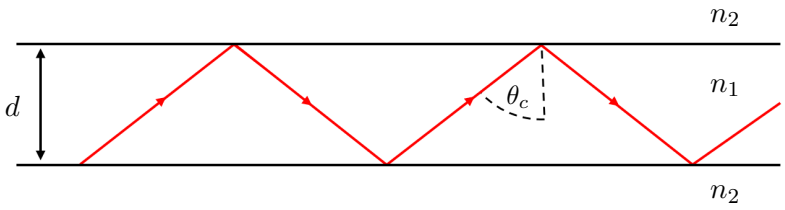

Figure 4.2: Ray-optic schema of light travelling through a symmetrical slab waveguide with core width d and refractive index n_1, cladding refractive index n_2, and critical angle θ_c.

A ray propagating through core material with a refractive index n_1 at an arbitrary angle θ_1 will eventually reach an interface with the cladding, which has refractive index n_2. At this interface, some proportion of the light will be refracted and some portion reflected. Snell's law gives the relationship between indices, the angle of incidence θ_1 and the angle of refraction θ_2 as

$$n_1 \sin\theta_1 = n_2 \sin\theta_2 \tag{4.3}$$

Above the critical angle θ_c, all light is reflected, none is refracted; this is known as *total internal reflection* [26]. The critical angle is given by

$$\sin\theta_c = \frac{n_2}{n_1} \tag{4.4}$$

A guided mode can only form if the ray direction satisfies the conditions for total internal reflection, i.e. the angle of reflection is above the critical angle. (If the cladding is not the same index on each side, then there will be different critical angles at the two interfaces.) However, light does not propagate at every angle of reflection above the critical angle: it only propagates at a discrete number of angles, producing discrete modes [26].

The 'ray' can be described as a wavevector k, which hits the wall of the core with angle θ. k can be separated into two orthogonal components, one travelling in the propagation (z) direction, the other in the transverse (x) direction:

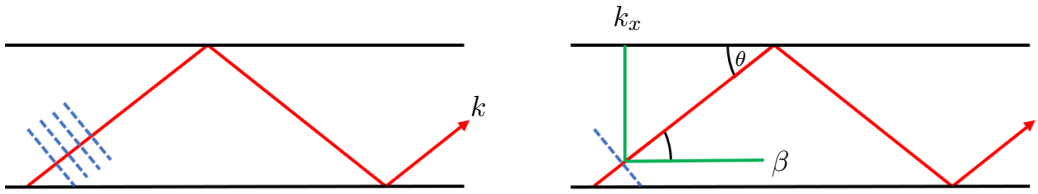

Figure 4.3: Light propagation through symmetrical slab waveguide, showing decomposition of wavevector k into x and z components. Hypothetical wavefronts for each component would travel in the directions of the green lines. The x component $k_x = k\cos\theta$ and the z component $\beta = kz = k\sin\theta$. Adapted from Boudrioua [256].

The transverse component of the plane wave experiences constructive and destructive interference with itself. A guided mode can only exist when the wave constructively interferes with itself, i.e. it satisfies the transverse resonance condition. This requires a specific phase shift within a transverse round-trip.

Within the core, $k_x = k\cos\theta$ and $\beta = k\sin\theta$; the phase shift in a full transverse round trip is $\phi = 2kd\cos\theta$ (Figure 4.3). To satisfy the transverse resonance condition, the reflected wave must be perfectly in phase, so ϕ must be a multiple of 2π: in other words, the resonance condition is $2kd\cos\theta = 2m\pi$, where m is an integer value. Here, the relationship between allowed transmission modes, waveguide dimensions, and refractive index profile becomes more apparent: the permitted angles, and therefore permitted modes, depend on the waveguide dimensions. The requirement that m is an integer indicates that only specific modes will propagate [257].

However, internal reflection at the core-cladding interface introduces an additional phase shift. In a symmetrical slab waveguide, the phase shift at each interface can be denoted ϕ_2, and the propagation condition becomes

$$2kd\cos\theta + 2\phi_2(\theta) = 2m\pi \tag{4.5}$$

If the cladding is asymmetrical, such that the core material with refractive index n_1 interfaces with two different claddings with refractive indices n_2 and n_3, then the $2\phi_2(\theta)$ term becomes $\phi_2(\theta) + \phi_3(\theta)$, where ϕ_2 indicates the change associated with the phase change at the $n_1 - n_2$ interface and ϕ_3 indicates the change associated with the phase change at the $n_1 - n_3$ interface [257]. This quantisation can be extended to 3 dimensions.

Whilst the ray-optic model provides an intuitive explanation for the evolution of guided modes, it fails to provide a complete picture of modal propagation. Modal confinement affects electric and magnetic field components differently. The electric and magnetic fields within a waveguide are mutually perpendicular and, in transverse modes, they are both perpendicular to the propagation direction. Therefore, in a rectangular waveguide such as the ridge waveguides

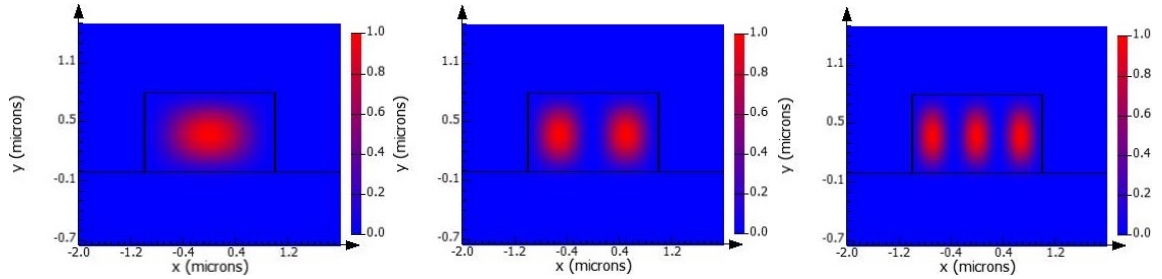

(a) TE00 (b) TE01 (c) TE02

Figure 4.4: FDE simulations showing amplitude of electric field at $\lambda = 450\,\text{nm}$ for three quasi-TE modes within an alumina ridge waveguide with cross-section $2000\,\text{nm} \times 800\,\text{nm}$. Simulations performed using Lumerical MODE finite difference eigenmode (FDE) solver, R1.2 [258].

considered in this thesis, a mode will consist of either electric or magnetic propagation—but not both[2]. For a transverse electric (TE) field, the electric field vector in the z direction, E_z, is 0, and for a transverse magnetic (TM) field, the magnetic field vector H_z is 0. Recalling the transverse resonance condition, whilst the critical angle (or angles, if the cladding is asymmetrical) for reflection are polarisation-insensitive, the phase shifts caused by internal reflection (ϕ_2 and ϕ_3) *are* polarisation-dependent, which means that TE and TM waves have different solutions for the transverse resonance condition; they have different propagation constants for the same mode number, and therefore have different characteristics [257].

Different mode orders are distinguished by adding two numbers to the mode descriptor—for example, TEab, where a indicates the number of nodes in the x direction and b indicates the number of nodes in the y direction. Figure 4.4 shows TE00, TE01, and TE02 modes in a ridge waveguide, which have 0, 1, and 2 nodes, respectively.

Figure 4.5, generated through solving an electromagnetic model, allows a qualitative evaluation of another crucial characteristic of guided modes: they extend into the cladding. Optical power travelling within the cladding is known as the *evanescent field*. The ratio of power propagating within the waveguide core to the total propagating power is known as the *confinement factor*, and is dependent on the relative refractive indices of the core and cladding, and the waveguide geometry [259]. When evaluating propagation loss within a waveguide, we must consider that light travelling within the cladding material is subject to material loss (Section 1.2)—if the confinement factor is low, most of the optical power is travelling within the cladding, and cladding material quality plays a critical role in minimising propagation loss.

[2]Numerical solutions of rectangular waveguide modes indicate that modes with similar effective indices couple to each other and produce hybrid modes which are not purely electric or magnetic; modes are more precisely quasi-TE or -TM. For low-order modes, one component dominates and it remains convention to describe modes as TE or TM.

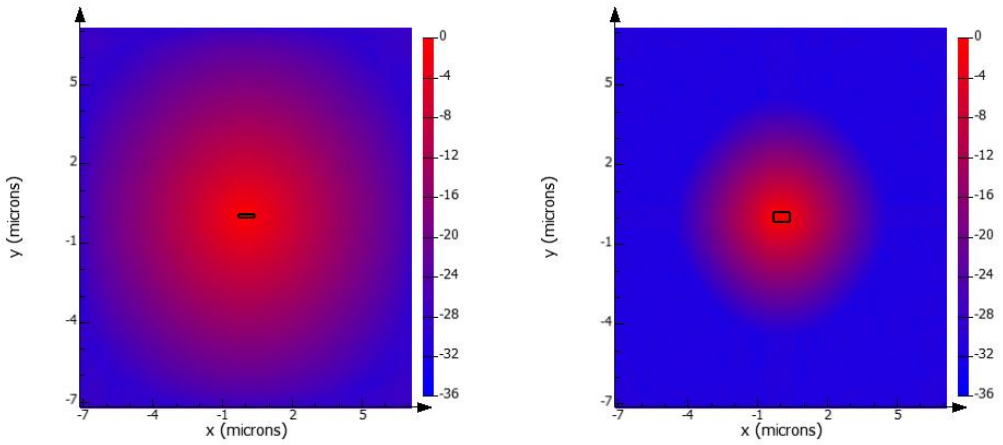

(a) Low-confinement: 800 nm × 100 nm. (b) High-confinement: 800 nm × 400 nm.

Figure 4.5: FDE simulations showing amplitude of quasi-TE00 electric field at $\lambda = 450$ nm for low- and high-confinement alumina waveguides fully clad in silica. Field amplitude is shown in log scale. Simulations performed using Lumerical MODE finite difference eigenmode (FDE) solver, R1.2 [258].

When designing a waveguide, it is essential to know the allowed transmission modes associated with a given geometry. This is typically visualised as the *effective refractive index* n_{eff} of the mode at a given wavelength. In a bulk medium, refractive index defines the change in radiation's phase velocity relative to its phase velocity in a vacuum; effective refractive index does the same for light confined within a waveguide.

A mode is considered to be 'guided' if $n_{eff} = {}^{\lambda\beta}/_{2\pi}$ is higher than the refractive index of the cladding material; *mode cutoff* is the point at which n_{eff} equals the refractive index of the cladding, at which radiation modes will predominate over guided modes. The higher the n_{eff}, the stronger the light confinement and the less evanescent field is present.

n_{eff} is defined by

- The waveguide geometry (d in Equation 4.5)
- The refractive indices of the core and cladding (θ in Equation 4.5, derived from Snell's Law, Equation 4.3)
- The mode being considered (which m is chosen to satisfy the transverse boundary condition in Equation 4.5)
- Whether the mode is TE or TM (from the polarisation-dependence of the boundary condition in Equation 4.5)
- The wavelength of the light (whilst this is made explicit in the right-hand side of Equation 4.1, a less obvious contribution arises from chromatic dispersion).

This simple explanation of modal confinement has limitations: ray optics cannot adequately explain the confinement of light when the size of the structure approaches the wavelength of the light itself. A more complete model can be derived from classical electrodynamic theory, starting with Maxwell's equations, a series of coupled vector partial differential equations which describe the behaviour of electric and magnetic fields and how they develop as a result of charges, currents, and the fields themselves. Waveguides are typically designed to support specific transmission modes; establishing the allowed modes requires solving an eigenvalue equation derived from Maxwell's equations. Therefore, classical electrodynamics are vital for PIC design. However, the eigenvalue equation is transcendental and cannot be solved analytically; commercial solvers employ numerical methods to to predict light propagation in any given waveguide design (Section 4.2). No intuitive design rules arise directly from an examination of classical electrodynamics [260, 261], so such a discussion is omitted here.

4.2 Design

This work aimed to create single-mode, high-confinement PEALD alumina ridge waveguides which have low propagation loss at wavelengths of 450 nm, 514 nm and 642 nm—selected based on available laser sources.

Total (insertion) loss for a waveguide comprises coupling loss and propagation loss; measuring the insertion loss of waveguides of different lengths allows extraction of the propagation loss. This technique is known as the *cutback method*. Waveguides of multiple lengths can be achieved by fabricating straight waveguides and cleaving them to different lengths, or by fabricating curved waveguides on the same chip—I chose the latter approach. At longer wavelengths, adiabatic tapers may be used to create straight waveguides which are highly confining over different propagation distances; in the UV and visible, material losses are a considerable contributor to insertion loss, so this approach is not viable.

Design requirements can be grouped into *functional requirements* and *non-functional requirements*. Functional requirements are those relating to the devices' performance in the domain for which they were designed: in other words, these requirements are essential in order to fulfil the aim of producing high-confinement ridge waveguides which have low propagation loss in the visible region. Non-functional requirements are design considerations which do not relate to the purpose of the devices, but which are required to achieve results. Non-functional requirements can be split further into *fabrication requirements* and *metrology requirements*.

4.2.1 Functional Requirements

Waveguides should be

- High-confinement ridge waveguides—the core must be completely etched, and waveguide dimensions must produce a confinement factor of at least 0.7
- Single-mode at each target wavelength—this requires a different waveguide width for each target wavelength
- Low-loss, clad and unclad, at each target wavelength

Recalling Section 1.2, propagation loss arises from absorption within the core and cladding, bulk scattering within the core and cladding, and scattering at the core-cladding interface. Chapter 2 considers the contributions of absorption and bulk scattering loss through indirect means (probing bonding composition and considering roughness of as-deposited films). Chapter 3 describes process choices designed to minimise sidewall roughness, and therefore interface scattering. None of this work directly links any given observation to propagation loss.

The waveguide geometry also affects the contributions of each type of loss. As discussed in Section 4.1, waveguide modes have evanescent fields which penetrate the cladding to different extents depending on n_{eff}. Both core and cladding materials experience absorption and bulk scattering, so waveguide losses are influenced by the proportion of the field travelling in each material. Since high-quality wide-bandgap cladding materials are relatively mature technologies, a low confinement factor—usually achieved with small cross-sectional dimensions, especially height (see Table A.1)—can be used to offset the effect of a lossy core material (see, for example, [254]). Conversely, increasing confinement (by increasing cross-sectional area) reduces modal interaction with sidewalls, and therefore reduces interface scattering [262].

Single-mode waveguides are useful in practical PIC applications; different modes have different effective refractive indices, so multimode waveguides have limited use in applications such as signal processing [41]. Since allowed transmission modes are defined by the wavelength of the light and the cross-sectional geometry, different waveguide dimensions are required for single-mode waveguiding of each different wavelength of light used for measurement.

I have also specified that waveguides should be ridge waveguides. This design maximises modal confinement, and also overall functionality; fully-etched waveguides are necessary for nonlinear optics (Section 5.2), can reduce bend radius and device footprint, and can increase design flexibility for PIC components such as gratings [46, 47, 48].

Fulfilling these functional design requirements requires the design of waveguides with cross-sections which are small enough to be single-mode at the propagation wavelength, but large enough that modal confinement is above 0.7.

4.2.2 Fabrication Requirements

Fabrication features which relate to the aim of fabricating low-loss waveguides are encapsulated in functional requirements. Requirements in this section relate to the practical limitations within the fabrication process.

- ALD deposition runs are long and costly; all waveguides should be the same height, so that a single deposition run can be performed.
- Waveguides must be thin enough that they can be fully etched using an HSQ layer with a maximum thickness of 1 µm—using the optimised etching protocol described in Section 3.3, waveguides must be no more than 600 nm thick.
- To ensure that a good cleave can be achieved, waveguides should be cleaved at least 4 mm from the edge of the chip.
- The final design revision included an additional requirement: to ensure that a good cleave can be achieved, each cleaved edge should be no longer than 14 mm.

4.2.3 Metrology Requirements

Many chip layout requirements relate to the cutback method of loss measurement. Propagation loss is measured per unit length, and can therefore be determined by measuring loss from waveguides of different lengths. If the waveguides are identical other than their lengths, then the loss per unit of length can be found by measuring input and output power for each waveguide and determining the linear relationship between waveguide loss and waveguide length. Further detail is provided in Section 4.4.

Consequently, metrology requirements are

- For waveguides designed for a given wavelength, there should be at least 4 different lengths on the chip, with the longest being at least double the shortest. Given that yield is rarely 100%, at least 3 of 4 waveguide lengths must be measurable to produce a result.
- To place multiple waveguide lengths on the same chip, the waveguides will necessarily contain bends. A bending radius and geometry which minimses bend losses must be selected.
- Each waveguide must contain the same number of bends, so that any bending losses are similar for all waveguides of a given geometry.

- To ensure that coupling losses are as similar as possible, all waveguides for a given wavelength must be on the same chip.
- Waveguides must be placed sufficiently far apart that no evanescent coupling between waveguides occurs.

4.2.4 Design Simulations

Simulating waveguide designs allows us to determine which waveguide geometries allow single-mode transmission: we calculate effective indices to establish allowed transmission modes. In this instance, we are seeking to maximise effective refractive index within the single-mode region, to maximise modal confinement and minimise modal interaction with the sidewalls.

Theory

Most useful waveguides, such as ridge waveguides, have no analytical solutions, but numerical methods can be used to determine electromagnetic field profiles and constants for any arbitrary cross-section.

I ran simulations using Ansys' Lumerical MODE software to determine the allowed transmission modes associated with cross-sectional dimensions. MODE uses a finite difference eigenmode (FDE) solver, in which the user creates a 3-dimensional model of a waveguide. The cross-section of this model is discretised using a 2-dimensional Yee mesh [263], which expresses the electromagnetic field at each mesh point. The software then uses the method described by Zhu *et al.* [264] to express Maxwell's equations as a matrix eigenvalue problem for each mesh point:

$$\mathbf{P}\begin{bmatrix}\mathbf{E}_x\\ \mathbf{E}_y\end{bmatrix} = \begin{bmatrix}\mathbf{P}_{xx} & \mathbf{P}_{xy}\\ \mathbf{P}_{yx} & \mathbf{P}_{yy}\end{bmatrix}\begin{bmatrix}\mathbf{E}_x\\ \mathbf{E}_y\end{bmatrix} = \beta^2\begin{bmatrix}\mathbf{E}_x\\ \mathbf{E}_y\end{bmatrix} \tag{4.6}$$

$$\mathbf{Q}\begin{bmatrix}\mathbf{H}_x\\ \mathbf{H}_y\end{bmatrix} = \begin{bmatrix}\mathbf{Q}_{xx} & \mathbf{Q}_{xy}\\ \mathbf{Q}_{yx} & \mathbf{Q}_{yy}\end{bmatrix}\begin{bmatrix}\mathbf{H}_x\\ \mathbf{H}_y\end{bmatrix} = \beta^2\begin{bmatrix}\mathbf{H}_x\\ \mathbf{H}_y\end{bmatrix} \tag{4.7}$$

where $\mathbf{E}$ and $\mathbf{H}$ denote the electric and magnetic fields in directions x and y (the cross-section dimensions), β is the propagation constant of the mode defined by these field distributions, and $\mathbf{P}$ and $\mathbf{Q}$ are sparse matrices which describe the mesh spacing and refractive index profile. Using numerical methods to solve these provides the eigenvalue β^2, which can be used to calculate the effective refractive index n_{eff}, as well as the cross-sectional mode profiles defined by $(\mathbf{E}_x, \mathbf{E}_y)$ and $(\mathbf{H}_x, \mathbf{H}_y)$.

Simulation of Allowed Transmission Modes

I ran FDE simulations using Lumerical MODE 2021 R1.2. Given a waveguide cross-section and a propagation wavelength, the solver provides a list of allowed transmission modes, their effective indices, and field distributions. Given the wide range of refractive indices that alumina and HSQ can present, I used sampled refractive index data derived from ellipsometer measurements (Section 2.1) in all simulations.

FDE simulations' accuracy is highly dependent on the spacing of the Yee mesh. Yee stated that "the space grid size must be such that over one increment the electromagnetic field does not change significantly" [263], whilst Zhu and Brown consider a grid size of $\lambda/15$ "usually enough to get fairly accurate results" [264]. A smaller mesh size produces higher resolution, but comes at significant cost in terms of computation time. I performed mesh size convergence testing to determine the largest possible mesh size which could be used: I altered the mesh size between 0.01 µm and 0.2 µm and considered the output n_{eff} for the TE00 mode. Results are consistent at mesh sizes below 0.03 µm. Given that the shortest wavelength considered for this project was 320 nm, I used a mesh size of 0.02 µm—approximately $\lambda/15$, as recommended by Zhu and Brown.

Given the functional requirements that waveguides are single-mode and high-confinement and the fabrication requirement that a single waveguide thickness is used, I varied waveguide width and tracked the calculated n_{eff} values for all modes. I performed these simulations at a range of heights and using wavelengths of 450 nm, 514 nm and 642 nm (Figure 4.6). Using this data, we selected a height of 400 nm, and waveguide widths of 700 nm, 850 nm and 1200 nm. At these dimensions, the FDE software finds higher-order modes—but where n_{eff} for those modes is below that of the cladding (approximately 1.47 [265]) the higher-order modes become radiation modes rather than propagating modes. In retrospect, a waveguide width of 600 nm should have been selected to ensure that waveguides designed for operation in the blue remained single-mode after cladding.

Further simulations were required to determine chip layout. It is vital to avoid *evanescent coupling*, in which the evanescent field propagating in the cladding cross-couples to another waveguide. The required separation between waveguides can be established by evaluating a log-scale plot of the electric field, such as those shown in Figure 4.5. Given the high modal confinement of the selected waveguide geometries, fundamental modes extinguish to −30 dB within 10 µm of the waveguide edge. However, given the potential for increased radiation at bends, I used a 100 µm separation between waveguides.

Within a mode, the phase front must be maintained; this condition poses problems when a mode encounters a bend [266]. As shown in Figure C.2, if a guided mode with a propagation constant β_z, in a medium where the maximum phase velocity (for unguided light) is β_0, is

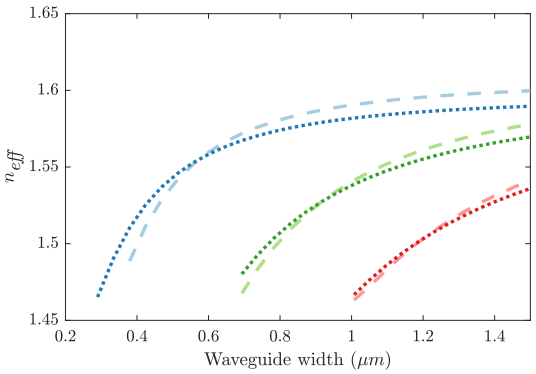

(a) 450nm, no upper cladding

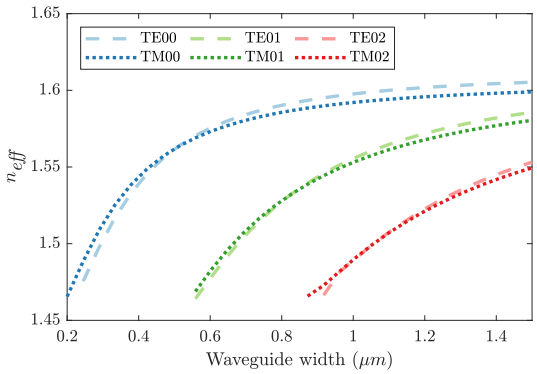

(b) 450nm, HSQ upper cladding

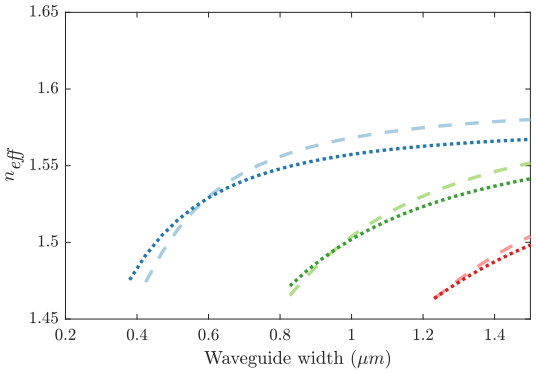

(c) 514nm, no upper cladding

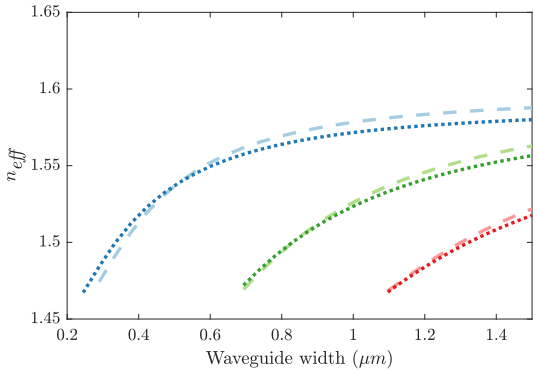

(d) 514nm, HSQ upper cladding

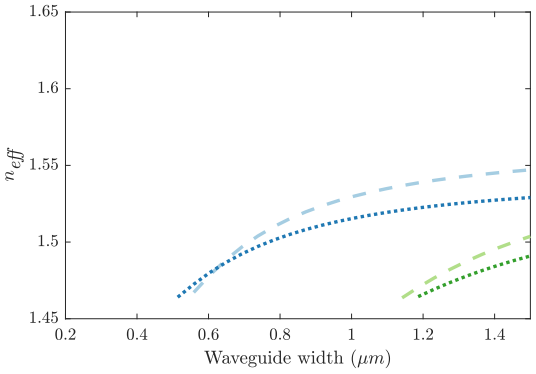

(e) 642nm, no upper cladding

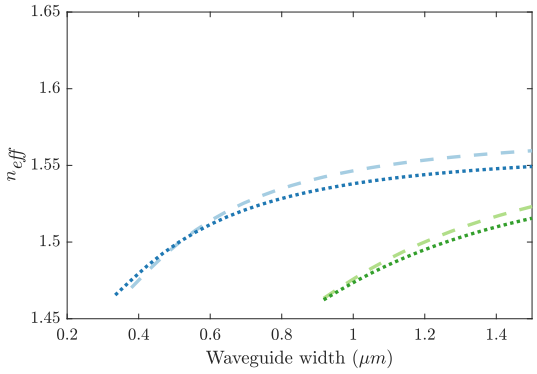

(f) 642nm, HSQ upper cladding

Figure 4.6: Simulated effective refractive indices (n_{eff}) for 400 nm-thick PEALD alumina waveguides without and with an HSQ upper cladding layer. Simulations performed using Lumerical MODE FDE solver, R1.2 [258].

travelling around a bend R, above a radius $(R + X_r)$ the light would have to travel faster than is possible in the medium (i.e., above β_0) in order to remain guided [32]. If this is not the case, radiative loss can occur. Lumerical simulations indicate that bending losses will be negligible at bend radii above 200 µm.

Mask Design

I produced GDS files for EBL writing using Tanner L-Edit v2016.2. Figure 4.7 shows the finalised mask design used for Revision 6 waveguides (see Table 4.1 for waveguide design revisions). Key features of this design are:

- Waveguides are written using a "paperclip" design.
- Two waveguides of each of four lengths (approximately, 1.5 cm to 3 cm in steps of 0.5 cm) are written.
- 20 µm-wide slab waveguides are written between each pair of waveguides. Coupling light into these waveguides is straightforward, since they have a large cross-section; aligning the input spot to a slab and moving it to a single-mode waveguide is considerably quicker and easier than initially aligning the spot to a single-mode waveguide.
- Markers at the edge of the chip provide an alignment reference during the cleaving process.
- Large features (green) are written with a large beam spot size to reduce EBL write time; only curved waveguides (blue) are written using optimised EBL conditions.
- No critical features are written within 2 mm of the edge of the chip: spinning and etching uniformity in this area is poor (Section 3.2.4).

When a waveguide bends, the mode distribution changes, shifting towards the outside of the bend (Figure 4.8). This creates a mode mismatch between the straight and curved waveguides, which can be a source of loss [267]. Given that each waveguide contains the same number of bends, bending loss is constant across all waveguides, and therefore will not contribute to measured propagation loss. However, in practical applications the overall insertion loss associated with a PIC is equally relevant. In Revision 1, all bends were circular, and optical loss at bends was visible to the naked eye. Scattering loss, caused by waveguide roughness, is likely the most significant contributor (Figure 3.5); however, I also re-evaluated bend geometry.

Bending geometries which feature *adiabatic* (slowly-varying) bends can reduce the degree of mode mismatch. Therefore, Revisions 2-6 used symmetrical quadratic *Bézier curves* generated using the CNST Nanolithography Toolbox [268]. Bézier curves are parametric curves generated through a linear interpolation between points; in the case of quadratic Bézier curves, three points are specified, giving the start and end of the bend, and a midpoint which

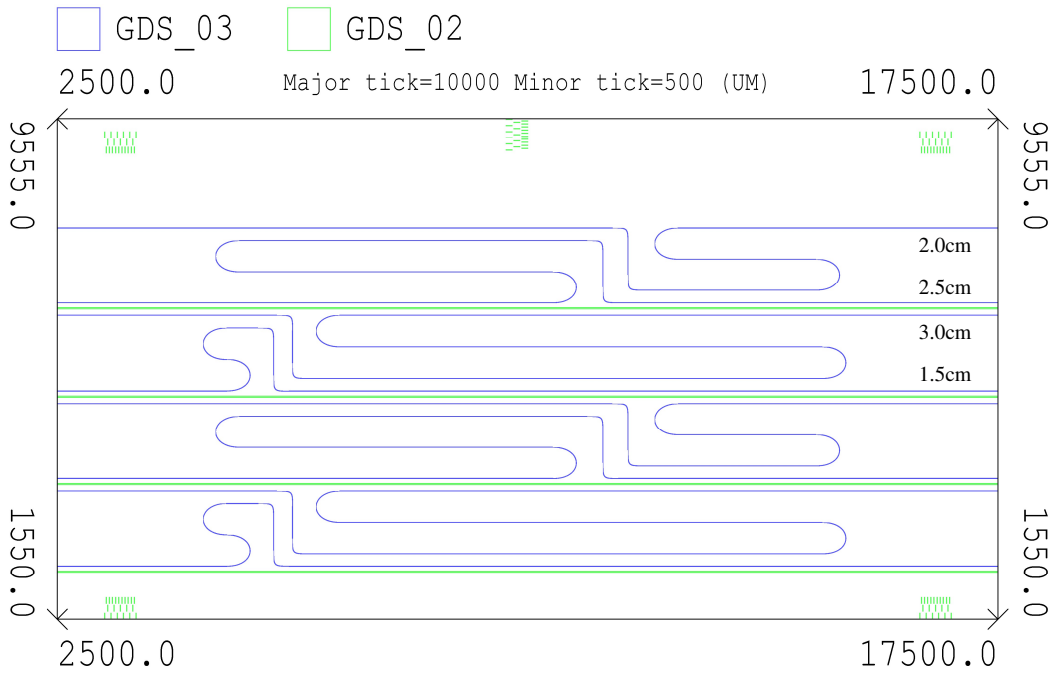

Figure 4.7: Finalised EBL mask design for high-confinement single-mode waveguides. Features written using a large beam spot size (alignment markers and slab waveguides) are shown in green; features using a small beam spot size (waveguides for measuring propagation loss) are shown in blue. Waveguide lengths are labelled.

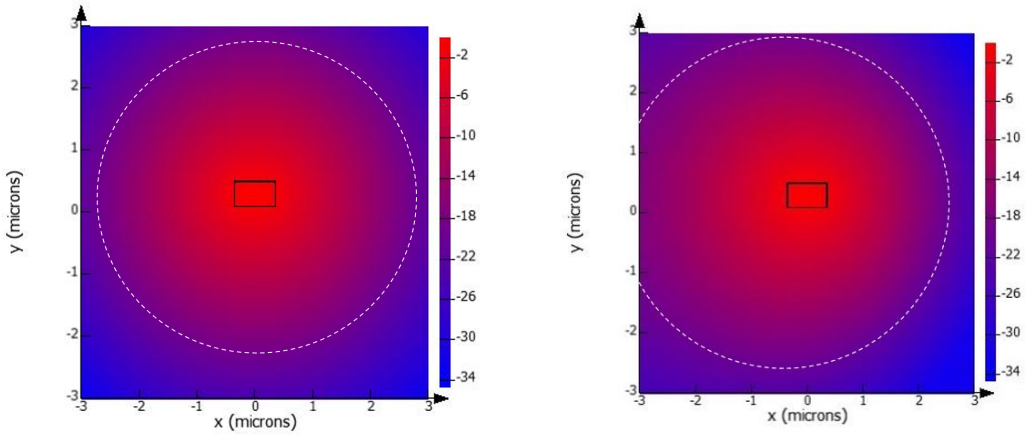

(a) Straight waveguide

(b) Waveguide with 50 µm radius semicircular bend.

Figure 4.8: FDE simulations showing amplitude of quasi-TE00 electric field at $\lambda = 450$ nm for curved and straight 700 nm × 400 nm alumina waveguides fully clad in silica. Field amplitude is shown in log scale; dotted lines mark −20 dB. Simulations performed using Lumerical MODE FDE solver, R1.2 [258].

defines the degree of curvature. I selected a *Bézier number*—which specifies the sharpness of the bend—of 0.5. In retrospect, it is likely that a lower Bézier number would be more effective in reducing losses, although even sharp Bézier curves have demonstrated lower losses than semicircular shapes [269].

4.3 Fabrication

The process flow presented in Section 3.1 is the finalised process flow for single-mode waveguides. All process parameters are provided in Table B.1.

I fabricated all single-mode waveguides using the same etching process; however, many of the lithography improvements arose from observations made during waveguide testing. Therefore, I fabricated 6 revisions of the single-mode waveguide design: process parameters for each revision are listed in Table 4.1. Revisions 5 and 6 used identical processing parameters and results from the two generations can be considered comparable; however, Revision 6 waveguides used smaller chips.

Table 4.1: Design and processing parameters for all single-mode waveguide revisions.

Revision	Bend design	Bend radius (μm)	Exposure dose (μC/cm^2)	No. of passes	Height offset (μm)	Beam current (nA)	Beam step size (nm)	Write order	Chip size (mm)
1	Circular	100	1400	1	0	2	4	Optimised	18×18
2	Bézier	200	1450	4	0	2	4	Optimised	18×18
3	Bézier	200	1450	4	0	1	2	Optimised	18×18
4	Bézier	200	1450	4	0	1	2	Sequential	18×18
5	Bézier	200	1450	4	1	1	2	Sequential	18×18
6	Bézier	200	1450	4	1	1	2	Sequential	14×20

4.4 Test Method

I measured waveguide losses using a *pseudo-cutback* method. Cutback loss measurements are so named because the insertion loss from a long waveguide is measured, the waveguide is cleaved into successively shorter pieces, and the insertion loss of each piece is measured. By plotting loss against waveguide length, loss per unit length—in other words, propagation loss—can be extracted [26]. The cutback method makes the assumption that all waveguide cleaves are identical: as discussed in Section 3.5, this is not the case. Pseudo-cutback measurements, in which waveguides of different lengths are written on a single chip, have several advantages over cutback measurements. Firstly, the measurement procedure is non-destructive. Secondly, all waveguide facets are defined using the same cleaving process—this certainly does not guarantee that all facets are identical in quality, but it does remove a source of variation from the process.

I designed an optical testing setup in conjunction with Natale Pruiti (Figure 4.9). End-fire coupling is used to couple light in and out of waveguides. TO-Can diode lasers provided light sources at 450 nm, 514 nm and 642 nm. Diode lasers have elliptical beam paths and

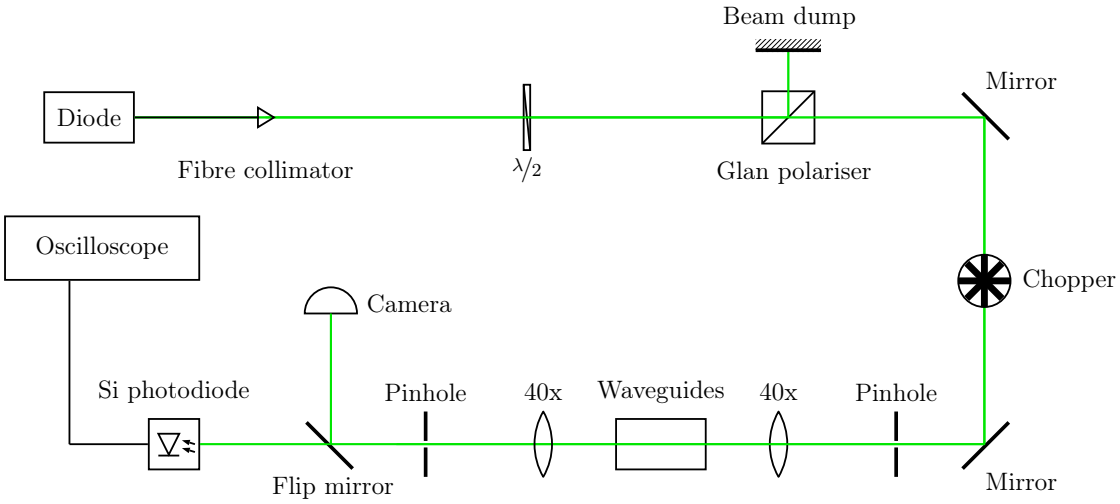

Figure 4.9: Schematic of optical setup used to test low-loss waveguides.

require collimation prior to coupling. To avoid re-collimating the beam each time the source is switched, we used diode lasers which had been pigtailed to optical fibres. These can be plugged into a fibre collimator to produce a collimated beam, with a collinear beam path for all test wavelengths. A half-wave plate and Glan polariser allow for selection of TE or TM polarisation; rotation of the half-wave plate allows for beam attenuation. Two steering mirrors allow the beam to be centred on an aspheric lens, producing a small spot size. The waveguide sample sits on a micropositioner-controlled stage, with the end facets positioned in line with the focused spot. A second aspheric lens focuses the light from the waveguide output, and a flip mirror allows output light to be directed to either a CMOS camera (a visual reference for fine-tuning alignment) or an amplified Si detector. An optical chopper periodically blocks the beam, allowing the signal intensity from the Si photodiode to be measured using an oscilloscope.

4.5 Results

Propagation loss testing is performed by measuring the insertion loss from multiple waveguides and fitting results to a line (Figure 4.10a). When working with low-power visible sources, the observation of light scattering patterns and intensities provides an additional source of information (Figure 4.10b).

In this work, six waveguide design revisions were required; visual inspection of light scattering during testing of the first three revisions indicated that they were not suitable candidates for low-loss waveguiding. Revision 1 waveguides demonstrated high levels of scattering at

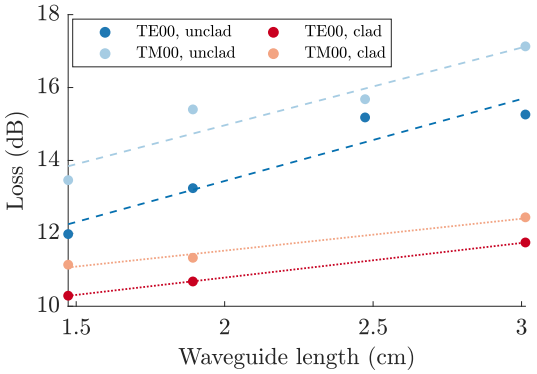

(a) Loss measurements of clad and unclad waveguides

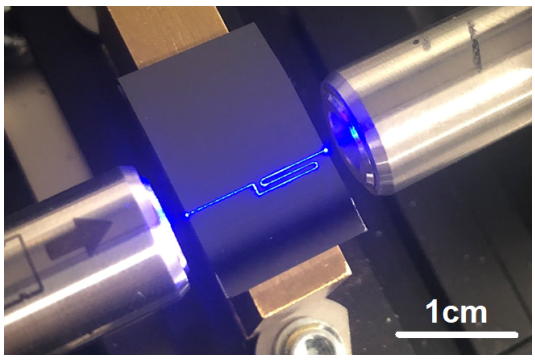

(b) Photograph of light propagation (right to left) in unclad waveguide

Figure 4.10: Observation and measurement of waveguide propgation losses at 450 nm. Reproduced from McKay et al. [76].

waveguide bends, visible to the naked eye. For subsequent revisions, I increased bend size, altered bending geometry, and increased EBL exposure dose. Revision 2 and 3 waveguides did not demonstrate high levels of light scattering at bends—however, stitching errors along the waveguides were apparent. Revision 4 waveguides used a sequential, rather than optimised, field writing order, aiming to reduce stitching errors, and produced acceptable losses.

For each chip of design revisions 4-6, I measured optical losses in the red, green, and blue without an upper cladding, then applied an HSQ upper cladding (Section 3.4) and remeasured the chip.

4.5.1 Revision 4

Table 4.2: Propagation losses in the visible region for single-mode waveguides, revision 4 (sequential writing, no height offset).

Wavelength (nm)	Unclad loss (dB/cm)		Clad loss (dB/cm)	
	TE00	TM00	TE00	TM00
642	4.8	3.7	0.8	0.5
514	6.8	4.9		
450	7.3	4.1	2.5	1.0

Revision 4 waveguide losses were relatively high (Table 4.2): I did not apply a height offset to rectify the EBL field scaling error (Section 3.2.7), and stitching errors at field boundaries produced obvious periodic scattering points (see Figure 3.11). This field scaling error has a clear impact on loss results. The effect is most profound for unclad waveguides: the (larger) difference in refractive index between core and its surroundings produces high interface losses. The TE00 mode is affected more than the TM00 mode, since TM modes interact more with the top and bottom of the waveguide (the interface with the BOX layer being least affected by

the refractive index discontinuity) and TE modes with the sides (which are both affected by the discontinuity). Consequently, TE00 losses were considerably higher than TM00 losses, even after application of an upper cladding layer.

Regardless, some of the trends which are apparent throughout all results are evident. Losses tend to increase as wavelength decreases, given the wavelength-dependence of both absorption and scattering losses (Section 1.2). Several 850 nm-wide waveguides, designed to be single-mode at 514 nm, were damaged, so I performed loss measurements at 514 nm on 700 nm-wide waveguides. Losses in the green are therefore higher than might be expected based on the values obtained in the red and blue, a result of greater levels of sidewall interaction.

4.5.2 Revisions 5 and 6

Table 4.3: Propagation losses in the visible region for single-mode waveguides, revision 5 (height offset applied). Results published in McKay et al. [165].

Wavelength (nm)	Unclad loss (dB/cm)		Clad loss (dB/cm)	
	TE00	TM00	TE00	TM00
642	1.6	1.4	0.6	0.6
514	4.0	3.6	1.4	1.3
450	6.5	6.5	1.3	1.1

Revision 5 waveguides were written with a height offset applied to the EBL job, rectifying the field scaling error observed in Revision 4. Table 4.3 presents propagation loss results, which are overall lower than those achieved from Revision 4 waveguides. The strong disparity between TE00 and TM00 losses is no longer evident. In several instances TM00 losses remain slightly lower than TE00 losses; this is to be expected, given that TM modes interact less with the etched sidewalls, which are likely the roughest interfaces.

When compared with later waveguide iterations, unclad losses in the blue and green are unusually high. However, the peculiarity of these results only became apparent after the waveguides had been cladded, preventing further investigation.

Table 4.4: Propagation losses in the visible region for single-mode waveguides, revision 6. Confinement factors are listed in brackets. Results published in McKay et al. [76].

Wavelength (nm)	Unclad loss (dB/cm)		Clad loss (dB/cm)	
	TE00	TM00	TE00	TM00
642	0.8 *(0.82)*	0.9 *(0.77)*	0.6 *(0.80)*	0.6 *(0.78)*
514	1.1 *(0.88)*	0.9 *(0.87)*	0.7 *(0.85)*	0.7 *(0.85)*
450	1.7 *(0.92)*	1.3 *(0.91)*	0.8 *(0.89)*	0.8 *(0.89)*

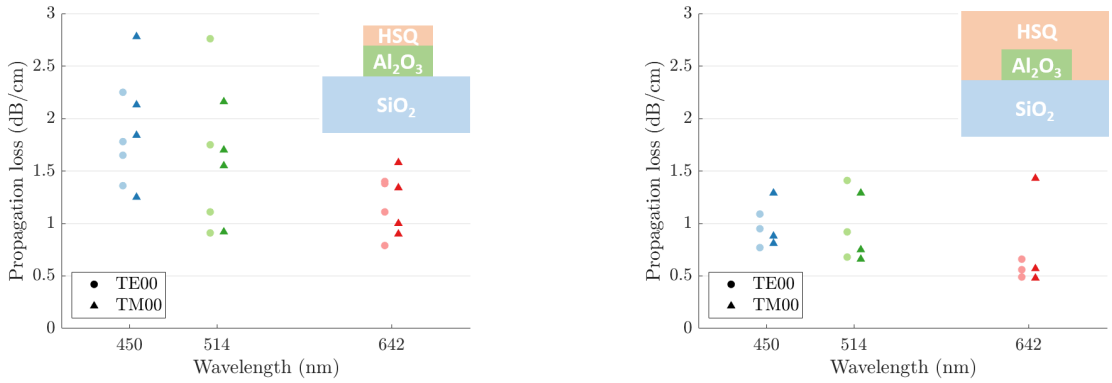

(a) Unclad waveguide losses

(b) Clad waveguide losses

Figure 4.11: Measured optical propagation losses in clad and unclad PEALD alumina waveguides, revisions 5 and 6. TE00 points are offset by −5 nm, TM00 points by 5 nm. Reproduced from McKay et al. [76].

Since Revision 5 waveguides demonstrated a fabrication process which could produce low-loss ridge waveguides using PEALD alumina films, no further changes were made to the waveguide shape or definition process in Revision 6; the only change was to the chip shape and the number of waveguides included on each chip. I fabricated and measured five unclad Revision 6 chips; contamination during the cladding process rendered two chips unusable after cladding. Table 4.4 presents the lowest loss results obtained across these samples, whilst Figure 4.11 presents Revision 5 and 6 loss results. There is some spread in measurements across samples; possible origins include

- Variations in waveguide quality, either chip-to-chip or across a single chip. Possible sources of chip-to-chip variation include deviations in the HSQ spinning process and variation in the condition of the ICP-RIE chamber; within-chip variations may be caused by particulate contanimation.
- Variations in cleave quality across a chip
- Variations in alignment of input beam to waveguide facet—the alignment process is manual, sensitive, and requires operator skill
- Variations in alumina film quality: Revision 6 chips used alumina from two separate deposition runs. No quality monitors were included in the first of these deposition runs, so there is no data to link lossiness to film properties such as bandgap or defectivity.

4.5.3 Coupling Losses

Each waveguide measurement measures a waveguide's insertion loss—a single figure which encompasses coupling, propagation, and radiation losses. Assuming that waveguide design does not result in appreciable levels of radiation loss and that results can be fitted to a straight

line where the slope gives the propagation loss, the constant term can be attributed to coupling losses.

Table 4.5 displays average coupling losses for Revision 5 and 6 waveguides. The high coefficients of variation indicate a high degree of variability in the limited dataset, but we can say that coupling loss typically decreases as wavelength increases, and that coupling losses are typically lower when waveguides are cladded. The latter condition makes sense, as clad waveguides experience slightly lower confinement; the larger modal area will improve overlap with the input spot. Edge coupling losses are considered relatively wavelength-insensitive, which is certainly true compared to alternative coupler designs such as grating couplers [270, 271]; the wavelength-dependence of loss apparent in this instance may be related to scattering losses or modal overlap.

Table 4.5: Coupling losses for Revision 5 and 6 waveguides. The coefficient of variation is listed in brackets.

Wavelength (nm)	Unclad loss (dB)		Clad loss (dB)	
	TE00	TM00	TE00	TM00
642	7.2 *(0.06)*	8.9 *(0.06)*	7.5 *(0.17)*	7.4 *(0.23)*
514	9.0 *(0.12)*	10.1 *(0.09)*	7.8 *(0.24)*	7.5 *(0.4)*
450	10.4 *(0.15)*	11.4 *(0.16)*	8.2 *(0.34)*	8.6 *(0.41)*

4.6 Low-Confinement Waveguides

Whilst this work focused on developing fabrication processes for high-confinement waveguides, the same processes can be used to fabricate low-confinement waveguides with minimal adaptation. I fabricated one set of waveguides using a 100 nm-thick alumina film. These results provide a comparison point with other alumina waveguides reported in the literature, which are typically on the order of 100 nm thick [4, 117, 132].

I repeated the simulations described in Section 4.2 and designed a mask containing four lengths of 700 nm-wide paperclip waveguides (single-mode at a propagation wavelength of 450 nm), producing a TE00 confinement factor of 0.31 and a TM00 confinement factor of 0.15. Only the fundamental mode can propagate in a waveguide with this geometry; with an upper cladding layer, the n_{eff} for both modes is $\sim$1.47. Simulations indicated that no modes could propagate in an unclad waveguide of this geometry.

Given the very low modal confinement, I designed bends with a radius of 400 µm and increased the separation between waveguides to 200 µm. Consequently, only one set of waveguides fitted on a chip. The only modification to the standard fabrication process (Table B.1) was to reduce the etching time to 30% of its usual value.

I measured the waveguides using the measurement protocol described in Section 4.4. TE00 loss was 1.1 dB/cm, whilst TM00 loss was 0.6 dB/cm. Here, the two polarisations have a considerable difference in confinement, which makes it difficult to speculate about the origin of the disparity in loss measurements. Whilst this dataset is limited, it is still worth reporting because it confirms the quality of a third alumina deposition run—in this work, three separate depositions have been used to produce low-loss waveguides. Furthermore, this result allows for benchmarking against prior work on alumina waveguides (Section 4.7).

4.7 Discussion

4.7.1 Wider Context

There are two major challenges in the fabrication of UV-visible PICs: selecting a wide-bandgap material platform with sufficient material quality, and developing high-quality processes which allow for repeatable fabrication of integrated optical components. Both absorption and scattering losses are wavelength-dependent, so the blue and UV present a considerably greater challenge than the red.

Figure 4.12 situates the results presented in Sections 4.5 and 4.6 within the wider context of low-loss waveguiding for the visible region. On consideration of subfigure (a), the results obtained in this work appear favourable, especially at shorter wavelengths—but SiN waveguides offer considerably lower losses. In subfigure (b) we consider waveguides which are high-confinement (here defined as confinement factor >0.7): in these waveguides, where the majority of the mode is travelling within the waveguide core, the distribution of SiN results becomes strikingly similar to the stoichiometric SiN absorption limits identified by Corato-Zanarella et al. [30]. At 450 nm, the alumina waveguides produced in this work offer losses below the SiN absorption limit.

Very few loss results are available for waveguides without an upper cladding, regardless of confinement factor. This may simply be because an upper cladding layer is a desirable practical feature: as well as reducing interface scattering, it ensures airborne or other particulate contamination does not touch the waveguide core, and makes devices more robust to mechanical damage. Few applications require waveguides without an upper cladding layer: sensing [13] and nonlinear optics are exceptions in this area. Therefore, whilst the loss results obtained in unclad waveguides can certainly be considered novel, they are not necessarily applicable to a wide range of applications.

Almost all prior or contemporary work on alumina waveguides has used low- to medium-confinement waveguides, with a maximum confinement factor of 0.4 [4]. TM00 losses from my 100 nm × 700 nm waveguides at 450 nm (0.62 dB/cm) are similar to those attained by

West et al. in 100 nm × 600 nm waveguides at 458 nm (0.7 dB/cm). Interestingly, the lowest losses attained in my high-confinement waveguides are only marginally higher (0.8 dB/cm).

In 2023, He et al. demonstrated alumina MRRs fabricated using an HSQ mask and BCl_3 etch. TE losses of 0.84 dB/cm at 390 nm demonstrate the first sub-dB/cm losses in the UV: as this result was obtained in shallow-etched rib waveguides, it is not included in Figure 4.12. By etching a 100 nm-deep, 4.5 µm-wide rib structure in a 400 nm-thick alumina layer, He et al. create a highly multimode waveguide which is high-confinement in the sense that most of the optical power travels within the alumina layer. This structure is minimally affected by sidewall scattering, and likely to demonstrate losses close to the absorption limit in the thermal ALD material they are using—however, it cannot support small bending radii and is unlikely to find use in nonlinear optics. Another interesting aspect of this work is that the alumina material the group used was commercially grown: the commercial availability of optical-quality alumina films represents a significant advance in the maturity of this material platform [205].

In short, SiN provides a high-quality, mature material platform in which high-confinement waveguides can offer low losses in the red and above, but alternative candidates are required for blue and UV waveguiding. There is a growing body of work in which amorphous alumina appears to be the frontrunner for this application. Both ALD and sputter processes have demonstrated sub-dB/cm material absorption in the near-UV [4, 205]. The commercial availability of optical-quality ALD alumina films represents a further step forward. This work is the first demonstration of the optical quality of dioxygen PEALD films, which are considerably quicker to deposit, and are therefore commercially attractive. Furthermore, we demonstrate low-loss waveguides fabricated from three different deposition runs (two 400 nm-thick depositions and one 100 nm-thick deposition), providing an indication of the repeatability of the deposition process. However, given the increasing availability of optical-quality alumina films, the primary contribution of this work lies in the waveguide definition processes: there are no other demonstrations of alumina waveguides for the visible region which are more than 120 nm thick.

In this work, the purpose of designing high-confinement waveguides was to achieve a low GVD, facilitating waveguide-based nonlinear frequency conversion processes. However, as PIC technologies mature, there may be further justification for the use of thicker films. Integration density can be increased: the most obvious example is in curved structures, where a higher n_{eff} reduces radiation loss (Appendix C). Thicker waveguide structures can also be advantageous when designing grating couplers, which often suffer from low coupling efficiency (CE) because of leakage of incident light to the substrate, as well as poor matching of the incident light to the mode diffracted by the grating [47]. Using a thicker waveguide does not inherently improve either condition, but provides an additional degree of freedom

which can allow a higher CE to be achieved. For example, Marchetti et al. demonstrate a higher conversion efficiency using a 260 nm-thick SOI layer compared to a standard 220 nm layer, although different apodised grating designs were required to maximise CE at each thickness [46].

It is worth evaluating the cost of using a high-confinement geometry in an application where a low-confinement geometry would suffice. Whilst the low-confinement loss data presented in this work is limited, it is approximately in line with other work presented in the literature [4, 205]. At a propagation wavelength of 450 nm, clad high-confinement waveguides demonstrated TM loss of 0.8 dB/cm, or 16.8% loss per centimetre, whereas low-confinement waveguides demonstrated TM loss of 0.6 dB/cm, or 12.9% loss per centimetre. In the example given here, consider MRRs fabricated using bend radii which produce negligible radiation losses for the high- and low-confinement platforms: 200 µm and 400 µm. To create a 180° bend, the distance travelled by light in the low-confinement platform is double that which is travelled using the high-confinement platform—so the low-confinement platform has both a higher footprint *and* higher losses associated with a longer propagation distance. Of course, the process flow was not optimised for low-confinement waveguides, and smoother etches and higher-quality claddings may further reduce losses; however, the reduction in path length using a higher-confinement platform could be a route to reducing total losses.

4.7.2 Measurement Technique

In this work, I used a pseudo-cutback method to derive propagation loss from the insertion losses of multiple waveguides. Using this methodology, we assume that coupling loss is identical for all waveguides, and that all waveguides are identical in quality—however, neither of these conditions are necessarily true, and a single unusually high insertion loss can skew the line fitting. This makes use of the cutback technique somewhat dubious in systems where propagation loss is not the dominant source of loss. As we develop sub-dB/cm loss capabilities, it becomes increasingly important to consider the suitability of loss measurement methodologies.

Propagation loss can also be measured by finding the Q-factor of a MRR. Measuring light transmission across a wavelength scan around the resonance wavelength provides the bandwidth of the ring's resonance, which is linked to the round-trip power loss. This allows a propagation loss measurement to be extracted from a single device, whilst removing assumptions about consistency of cleaving and waveguide quality. However, the method adds design complexity (designing a ring geometry and coupling gap), fabrication complexity (designing a process which attains sufficient resolution to expose coupling gaps, and where lateral bias is controlled well enough that coupling gaps are true to design), and measurement complexity (frequency-swept sources can be prohibitively expensive). Furthermore, a measurement

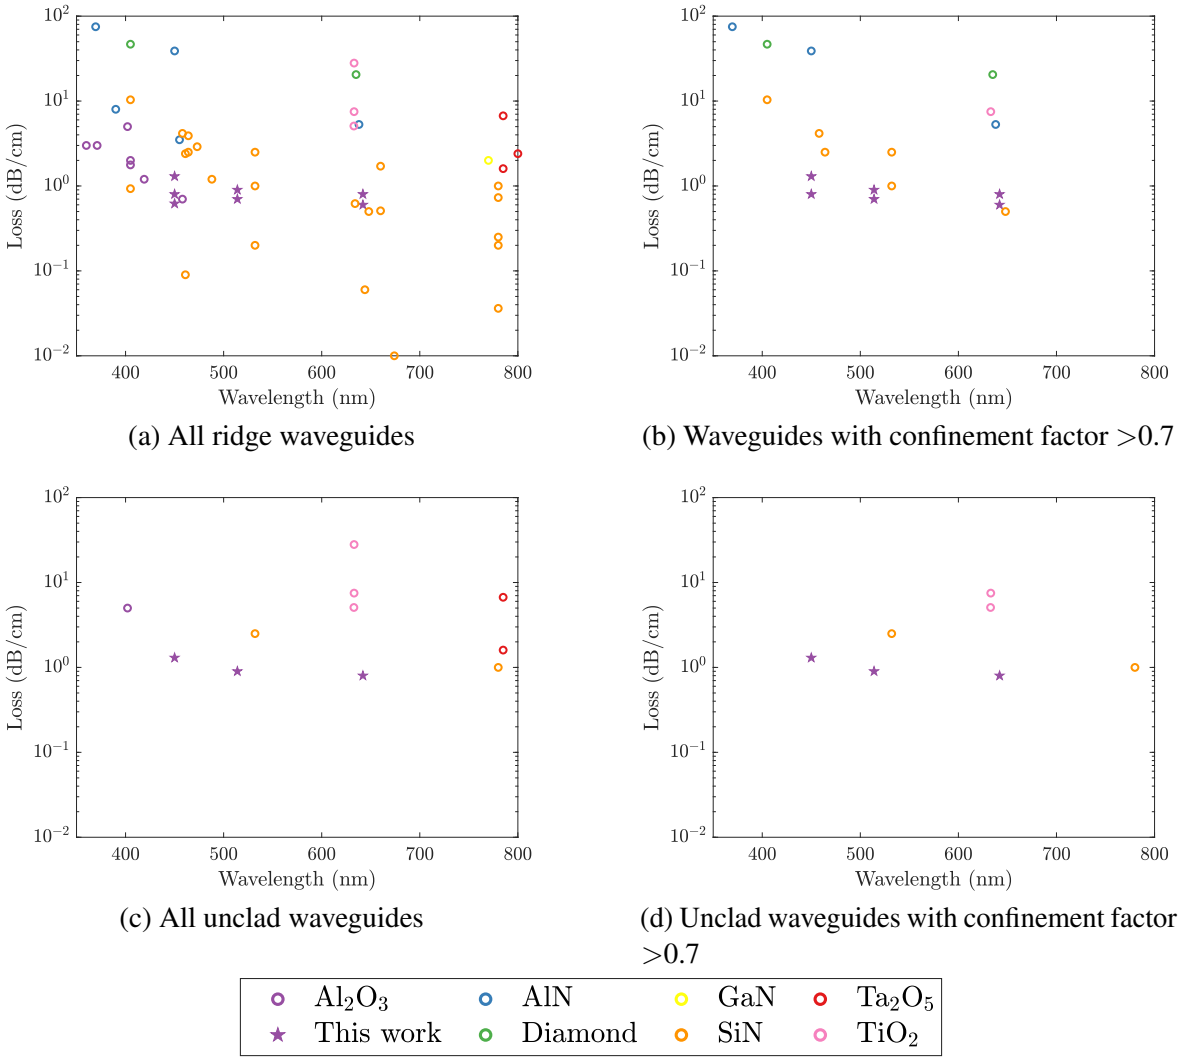

(a) All ridge waveguides

(b) Waveguides with confinement factor >0.7

(c) All unclad waveguides

(d) Unclad waveguides with confinement factor >0.7

Figure 4.12: Review of propagation losses achieved in ridge waveguides at wavelengths below 800 nm. Results from this work demonstrate the lowest propagation losses in waveguides with confinement factor > 0.7 at 514 nm and below. Full details are provided in Table A.1.

obtained from a single device does not demonstrate consistency and replicability in the fabrication process—and many papers do not discuss the total number of devices measured.

The robustness of the cutback method can be improved. Increasing the number and length of waveguides use may pose manufacturing difficulties, but certainly improves the chance of achieving a representative measurement—Chauhan et al. report losses in waveguides as long as 2 m [254]. Measurement of multiple samples is, of course, good scientific practice.

Using an end-fire coupling technique, measuring the same waveguides at multiple wavelengths is a possibility; low-cost diode lasers are suitable sources which are available in a range of wavelengths spanning the visible and near-UV. Losses should scale with wavelength, as was observed in this work. Combining measurement with simulation can be a powerful way of increasing results' robustness: for example, in recent work Corato-Zanarella et al. measured loss in MRRs of different widths and bend radii across a range of wavelength, and used this data to produce quantitative estimates of scattering and absorption loss. Our ongoing work uses a similar strategy: we have generated higher-density pseudo-cutback chip designs featuring a range of waveguide lengths and widths which are single-mode at 450 nm, 514 nm and 642 nm, with the aim of generating loss data at a range of confinement factors and deriving a (quantitative) relationship between propagation loss and sidewall roughness.

4.7.3 Future Improvements

The waveguides presented in this chapter are by no means perfectly optimised. I did not perform any optimisation of the alumina deposition process, nor relate it to optical losses. Some impurities derived from precursor molecules remain in the film (Section 2.7), although it is unclear whether these will produce absorptions at wavelengths of interest. West et al. annealed waveguides at 400 °C—with the stated aim of improving the quality of the PECVD silica cladding rather than the waveguide core. Their work also demonstrated exceptionally low slab losses (0.3 dB/cm at 405 nm), where films have presumably not been annealed [4]. He et al. annealed finished waveguides at 500 °C and 600 °C with the explicit aim of reducing losses in the alumina film, but do not provide comparison data on loss without the annealing step.

Other strategies discussed in Section 3.6 include annealing finished chips, which may reduce losses in the HSQ cladding, and quantifying sidewall roughness to allow for better decision-making when selecting lithography-etch parameters.

Finally, whilst bending losses are discounted from the cutback measurements, they are significant in PIC applications. In this work, I used Bézier curves to produce adiabatic variation of the mode; I selected this curve geometry because the CNST Nanolithography Toolbox makes them easy to generate. Other geometries such as Euler curves, or custom

optimisations based on waveguide simulations, could produce lower bending losses [272]. Whilst a reduction in bending loss with the same bend radius will produce lower insertion loss for a device, a further motivation for optimising bend geometry is to produce smaller bend radii and reduce device footprint and propagation lengths.

4.8 Conclusion

Amorphous alumina is emerging as the primary material platform for low-loss blue and UV waveguiding. In this work, we fabricated high-confinement PEALD alumina waveguides with unclad losses of 1.3 dB/cm and clad losses of 0.8 dB/cm at 450 nm. Low-confinement single-mode waveguides produced TM00 losses of 0.6 dB/cm at 450 nm.

Dioxygen PEALD films can be deposited considerably faster than the conventional thermal ALD films which have been used in other alumina waveguide demonstrations. This work demonstrates the first use of PEALD alumina in an optical waveguide application; furthermore, it demonstrates low-loss waveguides made from three deposition runs performed over a three-year period, indicating that the deposition process is relatively repeatable.

However, much of the contribution of this work arises from the process development work which allowed us to produce the first high-confinement fully-etched alumina waveguides for the visible region. Whilst low-confinement waveguides may be a practical choice which minimises loss, increasing film thickness and confinement factor allows us to create devices with smaller bending radii and smaller footprints, and overall improves PIC design flexibility. Whilst these waveguides are only 400 nm thick, allowing single-mode propagation, the fabrication processes developed in Chapter 3 can be used to etch waveguides up to 800 nm thick—thick enough to allow nonlinear spectral broadening (Chapter 5).

Chapter 5

Supercontinuum Generation in Alumina Waveguides

Waveguide-based nonlinear optics motivated the development of the high-confinement alumina waveguides demonstrated in Chapter 4. Waveguide-based nonlinear frequency conversion plays a critical role in miniaturising optical atomic clocks [43], whilst spectroscopy applications benefit from single-source broadband light in the blue and UV, where biological molecules have absorptions [14]. Optical coherence tomography, used to visualise biological tissues, has also demonstrated resolution improvements using broadband visible light generated through nonlinear optical processes [273].

The degree of nonlinearity in a waveguide can be altered through tailoring the waveguide cross-section (Section 5.1). Given access to a 780 nm pulsed source, I designed waveguides with anomalous GVD at 780 nm, aiming to produce spectral broadening through soliton fission processes (Section 5.2). However, variability in the fabrication process produced waveguides with GVD in the low normal regime. Waveguide testing using a high-power femtosecond pump laser (Section 5.3) produced broadening through self-phase modulation (SPM) and dispersive wave (DW) generation (Section 5.4), with the broadest supercontinuum spanning over an octave from near-UV to near-IR.

5.1 Theory

Nonlinear optical processes occur as a result of a material's nonlinear response to an extremely strong optical field. Given the general wave equation describing propagation of

electromagnetic radiation [35]

$$\nabla^2 \vec{E} - \frac{1}{c^2}\frac{\partial^2 \vec{E}}{\partial t^2} = \mu_0 \frac{\partial^2 \vec{P}}{\partial t^2} \tag{5.1}$$

where $\vec{E}$ is the electric field, μ_0 is the magnetic permeability of free space, and c the speed of light, the polarisation $\vec{P}$ consists of a series of terms

$$\vec{P} = \epsilon_0 \sum_{n=1}^{\infty} \chi^{(n)} \vec{E}^n \tag{5.2}$$

The size of each susceptibility $\chi^{(n)}$ is dependent on the bonding composition and structure of the material in which light is propagating. However, the higher-order $\chi^{(n)}$ terms, known as nonlinear susceptibilities, are small, and only have a detectable effect on optical processes where $\vec{E}$ is extremely high—in other words, where there is an extremely high light intensity I [35]. For power P and cross-sectional area A, I is defined as

$$I = \frac{P}{A} \tag{5.3}$$

Therefore, high input power, confined in a small area, is required to excite nonlinear effects. Waveguides offer the capacity to contain high optical intensities in a small cross-sectional area whilst also allowing tuning of dispersion properties. Here, I present a limited analysis of nonlinear effects which may arise in alumina waveguides, with a focus on developing design rules which produce a spectrum with a maximum blueshift and consistently high intensity over at least an octave. Considerable even-order nonlinear susceptibilities, such as $\chi^{(2)}$, are only present in non-centrosymmetric crystals [44]. The alumina used in this work is amorphous, and no second-order nonlinear processes were designed, simulated or observed in this work; therefore, second-order effects are ignored. Third-order nonlinear processes, associated with $\chi^{(3)}$, can occur in all materials, and are the focus of this work. The nonlinear processes which can be stimulated vary depending on the source; I used a pulsed laser with pulses of 90 fs, and therefore explore the effects of third-order nonlinear processes which produce spectral broadening in the femtosecond regime.

5.1.1 Optical Kerr Effect

The application of an intense optical field produces a shift in a material's refractive index, known as the Optical Kerr Effect. Where the intensity of the incident wave is $I(t)$, the linear refractive index n_0 is accompanied by an intensity-dependent component [35]:

$$n = n_0 + n_2 I(t) \tag{5.4}$$

In Section 1.5, I used the nonlinear refractive index n_2 as a measure of the strength of third-order nonlinear interactions in a given material. Its relationship to $\chi^{(3)}$ is given by [35]

$$n_2 = \frac{3\chi^{(3)}}{4n_0^2\epsilon_0 c} \tag{5.5}$$

Therefore, the magnitude of a material's nonlinear response is defined by its linear refractive index at the propagation wavelength and $\chi^{(3)}$, both properties which are defined by the material structure and composition.

5.1.2 Self-phase modulation (SPM)

If the refractive index within a medium is intensity-dependent, then phase velocity is also intensity-dependent, and propagating light's phase $\phi(t)$ will change in response to the electric field:

$$\phi(t) = \frac{\omega}{c} n_2 I(t) L \tag{5.6}$$

Here, $I(t)$ is the intensity of the pulse at time t and propagation distance L. The input pulse itself has a shape (a distribution in time) and therefore the refractive index has a distribution in time. This produces a time-dependent frequency change

$$\Delta\omega(t) = \frac{\omega}{c} n_2 L \frac{\partial I}{\partial t} \tag{5.7}$$

which can be approximated as

$$\Delta\omega = \frac{\omega}{c} n_2 L \frac{I_0}{\tau} \tag{5.8}$$

where I_0 is the peak intensity and τ the duration of the input pulse [274].

In the absence of additional nonlinear effects, SPM produces a linear, positive chirp—in the typical case where the input pulse already experiences positive chirp, SPM results in symmetrical broadening with the leading edge of the pulse redshifted and the trailing edge blueshifted [44]. Coherence is maintained, and pulses demonstrate a characteristic shape—a series of minima and maxima produced through constructive and destructive interference [275]. Spectral broadening through SPM is often desirable; as demonstrated in Equation 5.8, high levels of SPM can be achieved using short input pulses and high input powers. There is also a linear dependence on n_2, which is proportional to $\chi^{(3)}$ and inversely proportional to linear refractive index n_0.

A final contributor to the degree to which SPM takes place is the propagation length L. The nonlinear length L_{NL} is the propagation distance required to attain a phase shift of 1 radian through SPM; if this length is shorter than that of the waveguide, nonlinear effects will have

an appreciable effect on the propagating pulse.

$$L_{NL} = \frac{1}{\gamma P_0} \tag{5.9}$$

where P_0 is the pulse input power and γ is the effective nonlinear coefficient [44]:

$$\gamma = \frac{\omega_0 n_2(\omega_0)}{c A_{\mathit{eff}}(\omega_0)} \tag{5.10}$$

From this, we can observe that modal effective area A_{eff} also affects SPM.

5.1.3 Cross-Phase Modulation (XPM)

Where SPM is a nonlinear effect created by high optical field intensity within a single pulse, XPM is an analogous effect in which refractive index is affected by propagation of a different (i.e., different wavelength, polarisation, or spatial distribution/pulse) high-intensity optical field within the same medium. Where the two fields have angular frequencies ω_1 and ω_2 with intensities I_1 and I_2, the time-dependent phase change for the field with angular frequency ω_1 over a propagation distance L is given by [276]

$$\phi(t) = \frac{2\omega_1}{c} n_2 I_2(t) L \tag{5.11}$$

In the type of system discussed in this work—a train of high-intensity femtosecond pulses travelling in a dielectric waveguide—XPM can occur where pulses temporally overlap with each other. This may occur as a result of variation in group velocity, as well as through pulse broadening induced by GVD and SPM. Whilst these effects can produce the field overlap which allows XPM to take place, they also produce temporal *walk-off* which limits the propagation distance over which XPM can occur. For two pulses of duration $\Delta\tau$, dispersion D and a wavelength difference (between pulses) of $\Delta\lambda$, the walk-off length L_{WO} is given by [277]

$$L_{\mathit{WO}} = \frac{\Delta\tau}{D\Delta\lambda} \tag{5.12}$$

Unlike SPM, XPM results in asymmetrical broadening: the pulse with the higher group velocity is blue-shifted, and the pulse with lower group velocity is red-shifted [275].

5.1.4 Raman Scattering

Raman scattering is a process in which a photon is inelastically scattered by matter, producing an emitted photon with a different energy than the incident photon, and a bond in a different ro-vibrational state. *Stokes* scattering refers to Raman scattering processes in which the emitted photon is of lower energy, and the bond is excited: for a photon with angular frequency ω_0

interacting with a bond with resonance frequency Ω_R, the emitted (Stokes-shifted) photon has frequency $\omega_S = \omega_0 - \Omega_R$. *Anti-Stokes* scattering processes are those in which a bond is de-excited and a higher-energy photon emitted: $\omega_{AS} = \omega_0 + \Omega_R$. Stokes scattering occurs at a considerably higher rate than anti-Stokes scattering, since at any given time the vast majority of matter exists in the ground state [275].

Spontaneous Raman scattering is an incoherent process which occurs in at low intensity in all materials. Given a high incident field, this spontaneous process can seed *stimulated* Raman scattering (SRS), a third-order nonlinear process. Where the frequency difference between the pump and the Stokes wave (Ω_R) matches a vibrational mode, pump power is continuously transferred to a red-shifted, coherent *Stokes wave* [275]. The wavelengths generated are dependent on the wavelengths of the propagating light and the waveguide material, which together define which ro-vibrational energy levels are accessible. The degree to which SRS takes place can be defined by the Raman threshold, the input power at which, at the waveguide output, the Stokes and pump powers P_s and P_p are equal:

$$P_s(L) = P_p(L) \equiv I_0 A_{\mathit{eff}} \exp(-\alpha_p L) \tag{5.13}$$

Here, we can see that transfer of power to the Stokes wave also relies on the input pump power I_0, the mode effective area A_{eff}, optical loss α_p, and transmission length L [275].

5.1.5 Soliton Fission

Where a pulsed laser source is used, soliton dynamics can play a role in nonlinear broadening. Each pulse generated by the laser consists of temporally coherent radiation travelling at the phase velocity ($\nu = {}^{\omega}/_{\beta}$), whilst the pulse envelope travels at the group velocity ($\nu_g = {}^{d\beta}/_{d\omega}$). The wavelength-dependence of the group velocity is described by group velocity dispersion (GVD), β_2:

$$\beta_2 = \frac{\partial}{\partial\omega}\frac{1}{\nu_g} = \frac{\partial^2\beta}{\partial\omega^2} \tag{5.14}$$

If a pulse is launched into a waveguide in the normal dispersion regime, both GVD and SPM act to broaden the pulse. In the anomalous dispersion regime, however, GVD counteracts SPM, producing a structure known as a *temporal soliton*—a pulse which maintains its shape and frequency distribution during transmission. Solutions to the nonlinear Schrödinger equation, discussed further in Section 5.2.3, describe the nonlinear propagation of light; initial solutions for solitons take the form [278]

$$u(t, z = 0) = N\sqrt{P_0}\,\mathrm{sech}\left(\frac{t}{T_0}\right) \tag{5.15}$$

where T_0 is the soliton duration and P_0 is the peak power of the input pulse. The soliton order, N, is given by [44]

$$N = \sqrt{\frac{\gamma P_0 T_0^2}{|\beta_2|}} \tag{5.16}$$

Fundamental solitons ($N = 1$) travel as a temporally invariant shape; *higher-order* solitons ($N > 1$) are superpositions of fundamental solitons and have a periodically varying shape arising from interference [278]. From equation 5.16, it is apparent that soliton number is affected by γ, and is therefore proportional to $\sqrt{\omega_0}$ and $\sqrt{n_2}$ and inversely proportional to $\sqrt{A_{eff}}$. It is also apparent that for any given γ, P_0, and T_0, there is an upper limit on the value of β_2 which will produce soliton propagation ($N \geq 1$).

Soliton propagation does not result in considerable frequency broadening; however, soliton *fission* does. Higher-order solitons are unstable and may degenerate into lower-order solitons as a result of perturbations such as higher-order dispersion effects and SRS [135]. Solitons typically propagate, whilst being modified by nonlinear and dispersive effects, before fission occurs at the soliton fission length $L_{fiss} = {}^{L_D}/_{N}$, where dispersion length $L_D = {}^{T_0^2}/_{|\beta_2|}$ [275]. There are numerous mechanisms by which solitons can produce new frequency components: two common pathways are formation of *dispersive waves* (DWs), and *soliton self-frequency shift* (SSFS).

Intrapulse Raman scattering causes frequencies within a soliton to shift towards the red, known as SSFS. The high-frequency part of the pulse acts as a pump for the lower-frequency part [279]. Assuming the absence of chirp, the Raman-induced frequency shift Ω_p evolves linearly with propagation distance z [275]:

$$\Omega_p(z) = -\frac{8 T_R \gamma P_0}{15 T_0^2} z \tag{5.17}$$

Since $\Omega_p(z)$ is proportional to T_0^{-2}, SSFS becomes significant where pulse times are under 1 ps [275]. $\Omega_p(z)$ is also dependent on $\gamma \propto {}^{n_2}/_{A_{eff}}$ and peak pulse power P_0. The Raman response time T_R is on the order of 1 fs, and describes the delay associated with the Raman response.

DWs may be formed to dissipate the extra energy where a soliton number is non-integer [44]. Efficient transfer of energy to a dispersive wave during soliton fission requires phase matching: for a soliton pluse with central frequency ω_s and group velocity ν_g, a DW with frequency ω_d will form where [280]

$$\beta(\omega_d) = \beta(\omega_s) + (\omega_d - \omega_s)\nu_g^{-1} + \frac{\gamma P_0}{2} \tag{5.18}$$

DWs can be blue-shifted or red-shifted relative to the soliton frequency. Third-order dispersion (TOD), the change of GVD with angular frequency, is relevant here: if TOD is positive, DWs are blue-shifted, whilst if TOD is negative, DWs are red-shifted [281]. Conservation of momentum results in *spectral recoil* in which the soliton position moves in the opposite direction from the DW [282].

5.1.6 Optical Wave Breaking (OWB)

In the normal regime, positive GVD produces *chirp* in a pulse, in which lower frequencies travel faster than higher ones. SPM has an additional effect on chirp. The pulse has a temporal distribution, and since SPM is intensity-dependent, the pulse tails are less affected by SPM. Where frequency-shifted components temporally overlap with the pulse tails, a phenomenon known as *optical wave breaking* (OWB) can occur, in which new frequency components are created by nonlinear mixing of the frequencies. OWB also be induced by XPM [276].

An approximate expression for the propagation distance required for OWB to occur is given by [283]

$$z_{OWB} \approx \sqrt{\frac{L_D L_{NL}}{\varkappa^2 \sigma_2}} \tag{5.19}$$

where $\varkappa$ is a constant and σ_2 is the time integral of a normalised power distribution. The numerator $L_D L_{NL} = {T_0^2}/{|\beta_2|\gamma P_0}$, indicating that z_{OWB} is proportional to the square root of pulse duration, and inversely proportional to the square roots of peak power, effective area, and nonlinear refractive index.

OWB can precipitate the formation of DWs even where the pump operates in the normal-GVD regime, with no requirement for soliton formation or fission. Phase-matching conditions given by Equation 5.18 also apply in this case [284]. Transfer of power to a DW tends to be less efficient in the normal regime, and a short nonlinear length (inspection of Equation 5.9 demonstrates that high γ and high pump power minimise L_{NL}) and long dispersion length (low GVD) relative to the waveguide are required [285].

5.1.7 Four-Wave Mixing (FWM)

FWM describes parametric nonlinear processes in which four photons interact, creating new frequency components. This may involve the annihilation of three lower-frequency photons to produce one higher-energy photon or the annihilation of two photons to create two photons at different frequencies [275].

$$\omega_4 = \omega_1 \pm \omega_2 \pm \omega_3 \tag{5.20}$$

Total momentum and energy must be conserved—in other words, the wavevector mismatch Δk must be zero [275]:

$$\Delta k = \beta(\omega_1) + \beta(\omega_2) + \beta(\omega_3) - \beta(\omega_4) = 0 \tag{5.21}$$

$$\Delta k = \beta(\omega_1) + \beta(\omega_2) - \beta(\omega_3) - \beta(\omega_4) = 0 \tag{5.22}$$

Assuming no depletion of incident pump powers P_1 and P_2, and introducing the parameter $P_0 = P_1 + P_2$, the parametric gain g is given by [275]

$$g = \sqrt{(\gamma P_0 r)^2 - (\frac{\kappa}{2})^2} \tag{5.23}$$

where $\kappa = \Delta k + \gamma(P_1 + P_2)$ describes the *effective phase mismatch*, and $r = {}^{2(P_1 P_2)^{0.5}}/_{P_0}$.

Third harmonic generation (THG) is a specific case of FWM in which three degenerate photons of a frequency ω are annihilated to form one photon of frequency 3ω; in this instance the phase-matching condition is $\beta(3\omega) = \beta(\omega)$.

5.2 Design and Fabrication

Section 5.1 describes a selection of third-order nonlinear processes which contribute to spectral broadening. The magnitude of all of these phenomena increases monotonically with increasing nonlinear refractive index n_2 and with the power and intensity of the propagating light. Given that light intensity is area-dependent, a small effective area A_{eff} is generally favourable for nonlinear broadening. As discussed in Section 4.1, a large cross-sectional area increases confinement factor and reduces A_{eff}—however, an equally, if not more, compelling argument for producing unusually thick waveguides arises from dispersion engineering requirements.

GVD plays a key role in which third-order nonlinear processes are preferentially stimulated. As discussed in Section 1.5, material dispersion is typically normal, producing normal GVD; altering waveguide geometry to induce a phase change in the propagating light allows us to tailor the dispersion properties of the waveguide and counteract the effect of material chromatic dispersion on GVD. Achieving anomalous GVD is critical for formation of optical solitons. Other nonlinear optical processes are GVD-dependent in the sense that dispersion causes temporal broadening, and parts of a pulse, or multiple pulses, may have to overlap for the phenomenon to take place (SPM, XPM, OWB). It is typical to design waveguides so that they have low anomalous GVD at the pump wavelength, facilitating broadening through

soliton dynamics. Low dispersion over a wide bandwidth also facilitates further nonlinear interaction amongst newly generated frequency components [286].

A critical design question remains: what is the most useful spectral output? We could consider the brightness of the spectrum, its continuity, its coherence, the most redshifted or blueshifted component, or the presence of frequency doubling (for f-$2f$ self-referencing). Most supercontinua in the literature are reported by reporting the wavelength or frequency range which is continuously above a threshold (often −30 dB or −50 dB) relative to the output power at the pump wavelength. Given that a key driver of this work is atomic clocks, which require f-$2f$ self-referencing at short wavelengths, I aimed to generate at least an octave of broadening with a focus on generation of short-wavelength light. Pulsed lasers at 780 nm, 1064 nm and 1560 nm were available as pump sources; here, I present simulations using a pump laser at 780 nm, as zero dispersion is most easily attainable at this wavelength. This is a favourable outcome, as a short-wavelength pump reduces the degree of broadening required to generate wavelengths in the blue and UV.

There are very few design criteria compared to those for single-mode waveguides (Section 4.2). Comparison of multiple devices to derive a measurement is not required, releasing the chip from many of the metrology requirements associated with pseudo-cutback loss measurements. However, waveguide dimensions are critical: small variations in dimension can alter GVD and produce a considerable difference in output. Comparing measurements from devices with different cross-sectional dimensions remains desirable, since this can provide useful information about broadening dynamics.

5.2.1 Functional Requirements

- The waveguides must have sufficiently low insertion loss for generated light to be channeled to the detector.
- The waveguide cross-section must produce anomalous GVD at the pump wavelength (780 nm).
- Waveguide length (propagation distance) should be greater than L_{NL} and L_{fiss}.
- It is also desirable to include a range of waveguide cross-sections, some of which have low normal and strongly normal GVD, to demonstrate that spectral broadening is, in fact, produced by the waveguide geometry.

Given the large cross-sectional dimensions required to achieve anomalous GVD, all waveguides transmitted at least one mode throughout the generated frequency range. The process optimisations documented in Chapter 3 and tested in Chapter 4 were designed to minimise insertion loss. Therefore, design work focused on defining the effect of cross-section on GVD and propagation characteristics.

5.2.2 Fabrication Requirements

- The alumina core must be no more than 700 nm thick. This restriction is based on the alumina film thickness which was available.
- The maximum HSQ mask thickness is 1 µm (Section 3.2.4).
- Alumina films can be etched with a maximum selectivity of 0.8 (Section 3.3).
- To ensure that a good cleave can be achieved, waveguides should be cleaved at least 4 mm from the edge of the chip (Section 3.5).
- To ensure that a good cleave can be achieved, each cleaved edge should be no longer than 14 mm (Section 3.5). Given the available chip size (no larger than 20 mm × 20 mm), this limits propagation length to 12 mm.

These requirements placed considerable limitations on the capacity to fulfil the functional requirements—waveguide thickness has a strong effect on GVD, whilst a limitation on propagation distance may preclude the formation of frequency components far from the pump.

5.2.3 Design Simulations

Initially, I used FDE simulations to generate β values associated with a given waveguide cross-section; these values can be used to generate higher-order dispersions such as GVD, and can also provide information on phase-matching conditions for dispersive wave formation. However, analysis based solely on β values omits another critical factor: that broadening processes do not simply depend on the input pulse's wavelength, but on the frequency, temporal distribution, and intensity of the radiation, all of which vary during propagation. Therefore, I used the generalised nonlinear Schrödinger equation (GNLSE) to simulate the propagation of pulses based on simulated β values.

Group Velocity Dispersion (GVD)

In order to achieve nonlinear broadening caused by soliton fission, anomalous GVD at the pump wavelength is required. This is achieved by altering waveguide geometry until modal confinement introduces a phase change which counteracts (normal) material dispersion. FDE solvers use numerical methods to find the propagation constant β; GVD (β_2) is the second derivative of β with respect to angular frequency ($\beta_2 = {}^{\partial^2\beta}/_{\partial\omega^2}$). Lumerical MODE can derive β_2 for any given cross-section using a preset dispersion calculation function which derives β across a range of wavelengths.

Figure 5.1 demonstrates the effect of varying cross-section on the GVD of the fundamental modes at the pump wavelength of 780 nm. Here, we assume that all waveguides are fully-

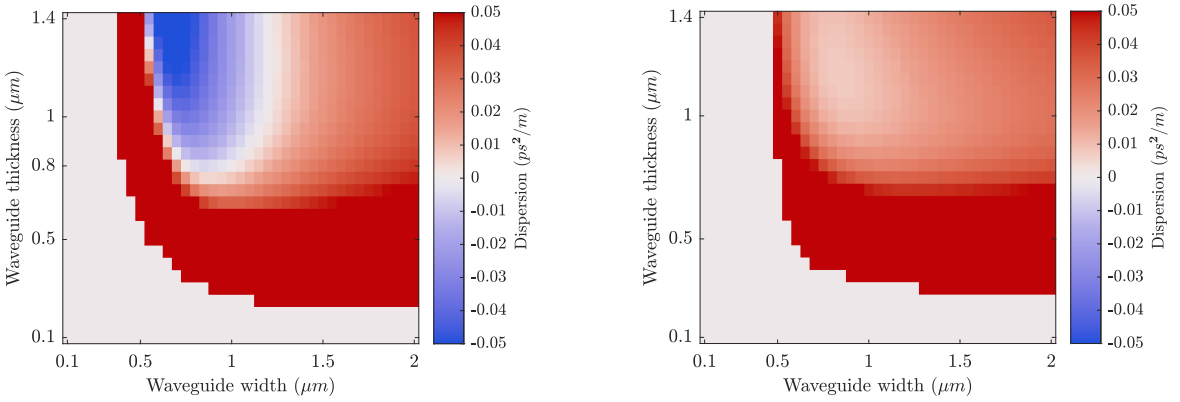

(a) TE00 (b) TM00

Figure 5.1: Maps showing variation of GVD with cross-sectional dimensions at $\lambda = 780$ nm for the fundamental modes of unclad alumina ridge waveguides. Values derived from simulations performed using Lumerical MODE FDE solver, R1.2 [258]. Subfigure (a) reproduced from McKay et al. [165].

etched alumina ridge waveguides without any residual HSQ. Within the wavelength range probed, only the TE00 mode achieves anomalous dispersion, requiring a minimum waveguide height of 800 nm. If an upper cladding is applied, the GVD becomes strongly normal in all instances—therefore, waveguides must remain unclad.

However, the thickest available alumina film was 700 nm thick—insufficient to achieve anomalous dispersion. The fact that applying an upper cladding has a dramatic effect on GVD points to an alternative strategy: β, and therefore GVD, are affected by the refractive index difference between the core and its surroundings. Overetching the waveguide increases the refractive index difference, and therefore alters GVD. Given that the maximum thickness of the HSQ mask is 1 µm, the depth of overetch is limited. If the silica bottom cladding layer has an etch rate similar to that of the alumina film, only 100 nm of cladding material can be etched. Fortunately, an overetch of 100 nm is sufficient to produce anomalous GVD in 700 nm-high waveguides (Figure 5.2), provided they are 800 nm to 1000 nm wide.

Soliton Fission and Dispersive Wave Formation

Table 5.1 provides values for A_{eff} and GVD (β_2) derived from FDE simulations. L_{NL} (Equation 5.9) defines the propagation length required for nonlinear effects to have a considerable effect on the pulse; it is much less than the propagation length of all waveguides. The dispersion length L_D defines the distance over which dispersion effects will produce considerable pulse broadening; given that the waveguide geometries have been selected for low GVD, L_D is consistently greater than the propagation length. For waveguides operating in the anomalous dispersion regime, L_{fiss} can be calculated; whilst there are three waveguide widths which

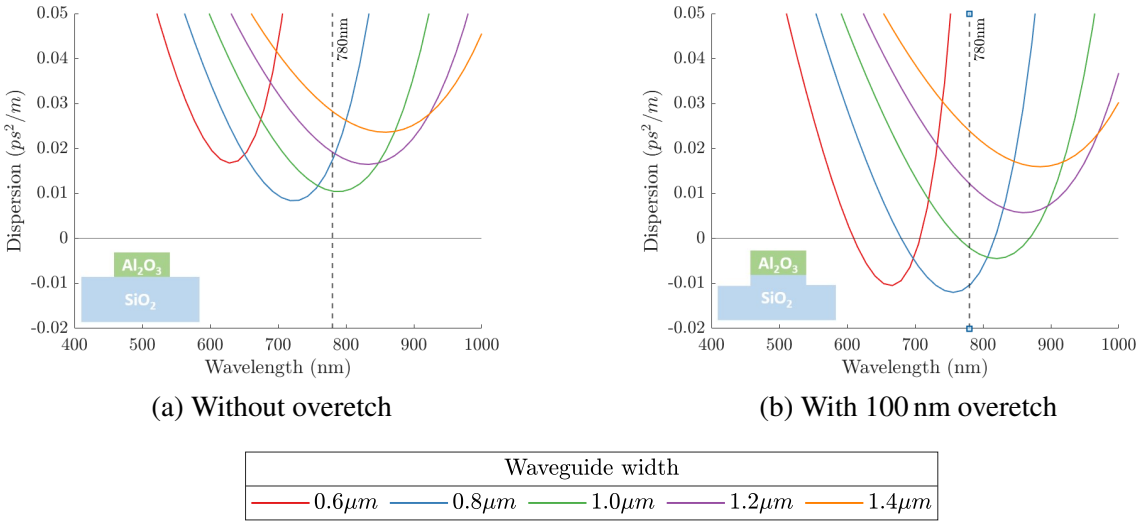

(a) Without overetch

(b) With 100 nm overetch

Figure 5.2: Effect of 100 nm overetch on GVD of TE00 mode in 700 nm-thick unclad alumina ridge waveguides. Values derived from simulations performed using Lumerical MODE FDE solver, R1.2 [258]. Inset schematics show representative waveguide cross-sections, not to scale.

produce anomalous GVD at 780 nm, only 800 nm and 900 nm-wide waveguides have soliton fission lengths which are shorter than the maximum allowable waveguide length (12 mm).

The phase matching condition for dispersive wave generation is achieved where integrated dispersion $\beta_{int} = 0$. A simplified expression for the phase mismatch between the dispersive wave and the soliton, which neglects a small nonlinear contribution, can be written as

$$\beta_{int} \approx \beta(\omega) - \beta(\omega_s) - (\omega - \omega_s)\nu_g^{-1} \tag{5.24}$$

Table 5.1: Parameters which govern nonlinear broadening processes, calculated for amorphous alumina waveguides with 700 nm core thickness, 100 nm overetch, and no upper cladding. Values derived from simulations performed using Lumerical MODE FDE solver, R1.2 [258].

Waveguide width (nm)	A_{eff} (μm^2) at 780 nm	β_2 (ps^2/m)	L_{NL} (μm)	L_D (m)	N	L_{fiss} (mm)
700	0.453	+0.006	64.1	1.31		
800	0.478	-0.011	67.6	0.77	107	7.2
900	0.510	-0.009	72.1	0.95	114	8.3
1000	0.546	-0.002	77.2	3.93	225	17.5
1100	0.584	+0.005	82.6	1.54		
1200	0.623	+0.012	88.1	0.66		
1300	0.664	+0.018	93.9	0.44		
1400	0.706	+0.024	99.8	0.34		

where ω is the angular frequency of a given wavelength, ω_s is the angular frequency at the centre of the soliton (the pump wavelength), and ν_g is the group velocity [280]. All of these parameters are readily calculated using Lumerical MODE FDE simulations.

Soliton fission is only expected where GVD at the pump wavelength is anomalous, so Figure 5.3 presents the results of integrated dispersion calculations (Equation 5.24) for the subset of waveguide widths where this condition is true. Because β_{int} relates to the phase mismatch with the pump wavelength, it is zero at 780 nm; dispersive waves may form to either side, and the wavelength at which they form defines the position of the dispersive wave. Of the range of 700 nm-thick waveguides with a 100 nm overetch which have anomalous GVD, only 900 nm and 1000 nm-wide waveguides are predicted to produce dispersive waves, with the shortest dispersive wave position (635 nm) occurring in 900 nm-wide waveguides. Note that for 1000 nm-wide waveguides, $L_{fiss} = 17.5$ mm—so emission of the dispersive wave is not predicted over the 12 mm propagation length.

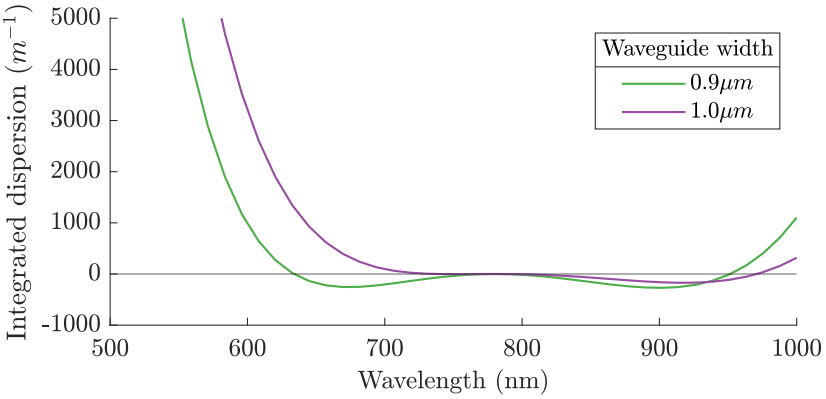

Figure 5.3: Integrated dispersion values attained in 700 nm-thick unclad alumina ridge waveguides with a 100 nm overetch. Values derived from simulations performed using Lumerical MODE FDE solver, R1.2 [258].

Generalised Nonlinear Schrödinger Equation (GNLSE)

The simulations provided in the previous two subsections describe the behaviour of light at a given wavelength without simulating the shape of a pulse or its evolution. However, nonlinear effects produce pulse evolution in both frequency and temporal domains—a complete picture of a pulse's behaviour as it propagates cannot be obtained from FDE simulations. The GNLSE can be used to model propagation: given an input pulse with a wavelength and temporal distribution, sequential transformations are used to model the evolution of the pulse envelope [287].

There are a number of assumptions used in the simplification: here, we assume that the propagation is unidirectional and single-mode. The *slowly varying envelope* approximation is also critical: we define time as existing within a coordinate frame moving at the group velocity

of the input pulse, so if there were no dispersion effects the centre of the pulse would remain at $t = 0$. Therefore, if we define $T = t - \beta_1 z$, where t denotes the time, z the propagation distance, and β_1 the propagation constant, then the evolution of a pulse envelope $A(z, T)$ can be described by

$$\frac{\partial A}{\partial z} + \frac{\alpha}{2}A - \sum_{k\geq 2} \frac{i^{k+1}}{k!}\beta_k \frac{\partial^k A}{\partial T^k} = i\gamma \left(1 + i\tau_{shock}\frac{\partial}{\partial T}\right)\left(A(z,t)\int_{-\infty}^{+\infty} R(T') \times |A(z, T - T')|^2 \, dT'\right) \quad (5.25)$$

The left-hand side of this equation models linear effects: loss (α) and the effects of dispersion (the Taylor expansion of β_k). The right-hand side models nonlinear effects. The first term describes the optical Kerr effect, with the nonlinear coefficient $\gamma = {}^{\omega_0 n_2(\omega_0)}/_{cA_{eff}(\omega_0)}$ describing the strength of nonlinear interactions. n_2 and A_{eff} are evaluated at the pump frequency ω_0. The time derivative term models broadening associated with self-steepening and optical shock, with $\tau_{shock} = {}^1/_{\omega_0}$. The response function $R(t) = (1 - f_R)\delta(t) + f_R h_R$ describes the instantaneous and delayed Raman contributions—f_R describes the fractional Raman response, whilst h_R describes the delayed (phonon-mediated) Raman response.

I used the GNLSE solver written by Dudley and Taylor [133] to simulate pulse evolution in the waveguides I designed. β and β_2 parameters are derived from Lumerical MODE simulations; I used polynomial fitting to generate higher-order dispersion terms. A_{eff} is calculated using Lumerical MODE. I used a value of n_2 derived by Dragic et al. (4.77×10^{20} m^2/W [77]). The fractional Raman contribution is typically determined through spectral broadening experiments [288] which have not been performed in alumina; however, Raman scattering appears to be very low in amorphous alumina [289] and is therefore ignored in these simulations. There remains a risk that these simulations do not accurately represent any redshift that will result from Raman scattering.

Figure 5.4 presents simulations of spectral broadening over 12 mm of propagation within a 900 nm-wide, 700 nm-high alumina waveguide with a 100 nm overetch into the bottom cladding material. Soliton fission occurs at ∼4 mm, shorter than predicted by calculations of L_{fiss} from A_{eff} and β_2, and is followed by considerable broadening across the visible and IR. Some discrepancies between GNLSE and FDE calculations can be expected, given that FDE simulations do not model the dynamic evolution of new frequency components, which can modify soliton dynamics through processes such as SPM, XPM, and SSFS, nor do they take into account the effect of higher-order dispersion (which can perturb soliton dynamics, leading to fission [135]). However, the shortcomings of the GNLSE must also be considered when evaluating these results: the wavelength-dependence of loss is not encoded in these calculations, so short-wavelength components, in particular, are unlikely to appear at

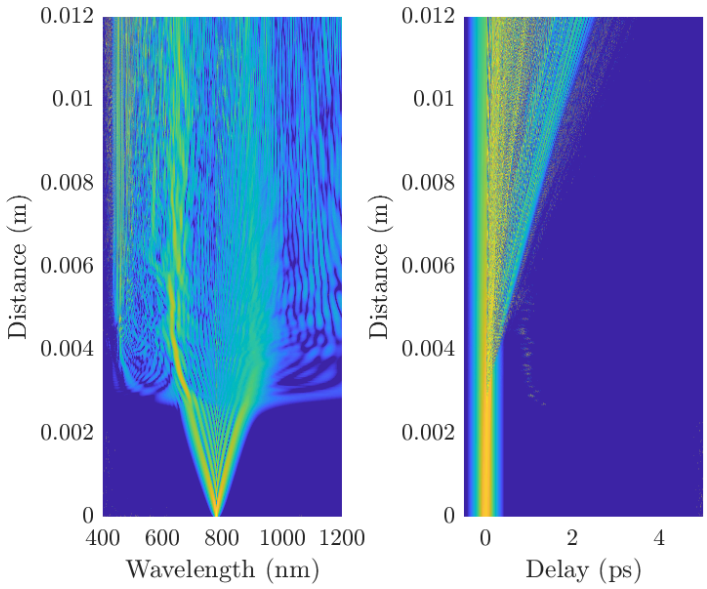

Figure 5.4: GNLSE simulation showing evolution of TE00 mode of a 780 nm, 90 fs input pulse with 18 kW peak power over 12 mm of propagation within a 900 nm-wide, 700 nm-high alumina waveguide with a 100 nm overetch into the bottom cladding material. Simulations performed using GNLSE solver written by Dudley and Taylor [133].

the amplitude shown. This may even preclude the generation of extreme short-wavelength components of the supercontinuum.

5.2.4 Mask Design

Mask design for these devices is straightforward: only straight waveguides are required. I used a chip size of 10 mm × 20 mm, with waveguides written parallel to the long edge. Excluding the edge zone leaves 6 mm of width in which waveguides can be written, and cleaving facets for end-fire coupling 4 mm from the edges of the chip produces a waveguide length of 12 mm.

I wrote waveguides in 8 widths, from 700 nm to 1400 nm in steps of 100 nm, with a 100 µm gap between waveguides. For redundancy, I wrote four sets of 10 waveguides, each set accompanied by a 20 µm-wide slab waveguide to be used as an alignment aid.

5.2.5 Fabrication

I fabricated devices using a variation of the standardised single-mode waveguide process flow. The 700 nm-thick alumina layer was deposited in two stages, by depositing a 300 nm-thick layer on top of a 400 nm-thick layer.

I spun a 1 µm-thick layer of HSQ to form the etch mask, and used the high-selectivity variant of the standard BCl_3 etch described in Section 3.3. All process parameters are provided in Table B.2.

5.3 Test Method

I used an end-fire coupling setup to direct light into the waveguides, shown in Figure 5.5. The pulsed source is a Menlo Systems C-Fiber 780 High Power laser, with a 90 fs pulse width and a measured maximum pulse energy of 1.6 nJ. Emitted radiation is linear and *s*-polarised (analogous to TE-polarised propagation within a waveguide). The beam is directed through a half-wave plate and a polarising beamsplitter, allowing selection of TE or TM polarisation; another half-wave plate can be used to control beam attenuation. Steering mirrors allow the beam to be centred on an aspheric lens, producing a small beam spot which can be focused on the waveguide facet. The waveguide sample sits on a micropositioner-controlled stage, with the end facets positioned in line with the focused spot. A second lens focuses the light from the waveguide output. A doublet lens with a focal distance of 15 cm is mounted on a rail system, allowing it to be moved back and forth to collimate the output beam. The beam is directed to a lensed fibre which provides input to a spectrometer. Post-holders before and after the waveguide stage allow for placement and movement of a power meter head, for measurement of in- and out-coupled optical power.

There are several alignment aids: a flip mirror prior to the stage allows integration of a 635 nm (visible) alignment laser, attached to a fibre collimator and aligned collinear with the 780 nm beam path. A second flip mirror after the waveguide stage allows output light to be directed to a CMOS camera, which provides a second visual alignment reference.

The spectrum is measured using an OceanOptics HR2000+ spectrometer, with a wavelength range of 200 nm to 1100 nm. High light intensities can saturate this spectrometer, adversely affecting results. A double-threaded mount on the pinhole directly after the stage allows filters to be added to the circuit: neutral density filters were required throughout, and a shortpass filter with a cut-off wavelength of 650 nm was required to collect light in the UV without detector saturation. Spectra were collected over an integration time of 50 ms.

5.4 Results and Discussion

5.4.1 Fabrication

Profilometry of etched waveguides indicated an etch depth of 806 nm, measured from the base of the waveguide to the centre of its top surface. Scanning electron microscopy of the

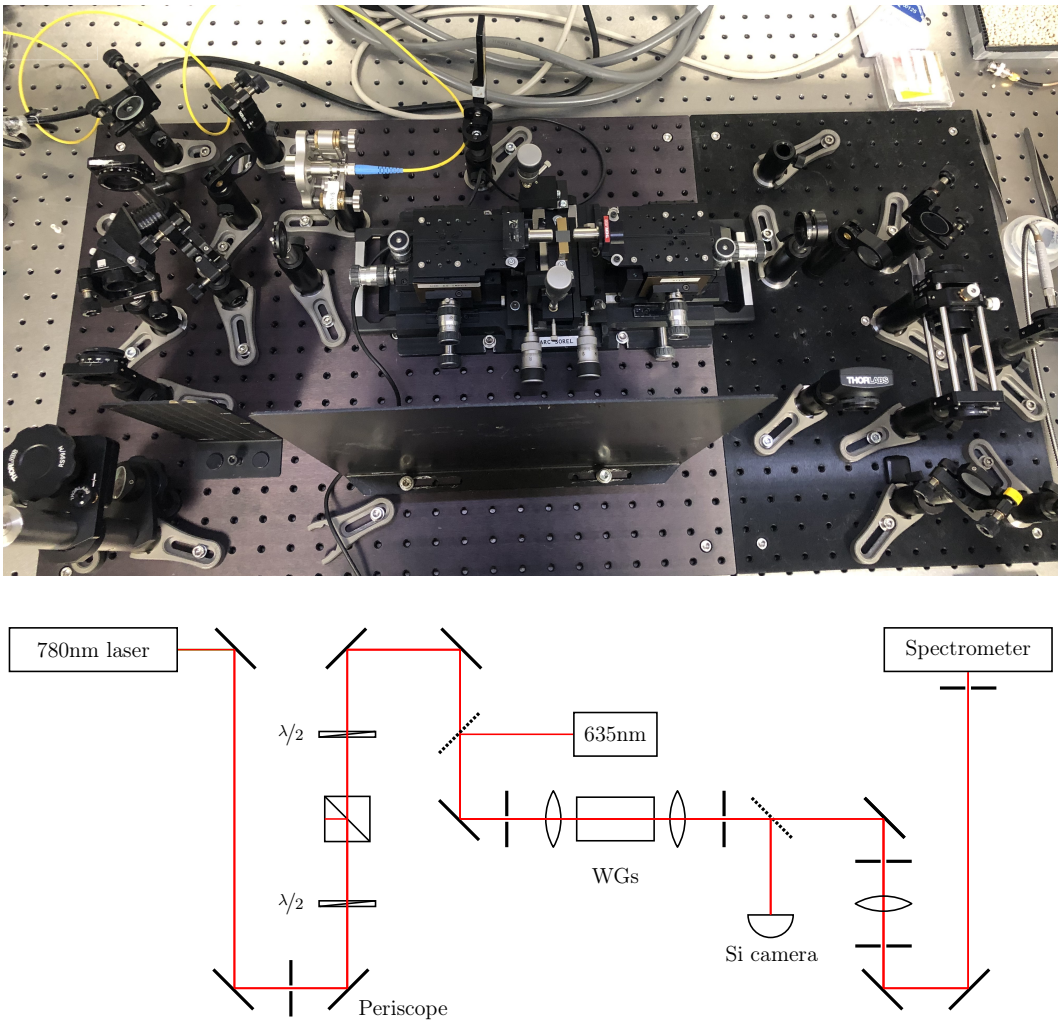

Figure 5.5: Schematic of optical setup used to generate supercontinua in alumina waveguides.

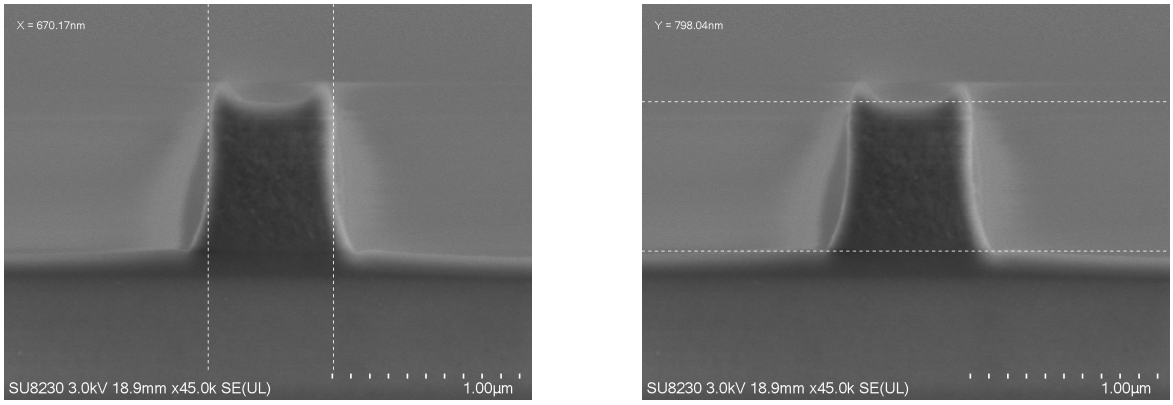

Figure 5.6: Dimensions of 700 nm-thick alumina waveguide with lithographic width of 700 nm. Measurements performed using calibrated SEM.

waveguide offcuts demonstrated that the etching time was too short: it was sufficient to etch through the alumina layer, but not sufficient to etch the required 100 nm into the silica layer below. Instead, approximately 100 nm of HSQ mask remained on top of the waveguide. The mask did not etch evenly, with more HSQ remaining at the sides of the waveguide. This may be a result of poorly-calibrated PEC.

Figure 5.7 presents updated GVD simulations, demonstrating that the waveguides operate in the low normal regime. Soliton dynamics cannot be evaluated; however, GNLSE simulations predict considerable spectral broadening through SPM, with further broadening through OWB and DW formation in waveguides where low GVD allows for temporal overlap of generated frequencies.

5.4.2 Supercontinuum Generation

I generated supercontinua in fabricated normal-GVD waveguides using the testing methodology described in Section 5.3. Figure 5.8 shows the spectrum produced by the pump laser, a 780 nm erbium laser with a 90 fs pulse width, when transmitted through a waveguide at low power (10 mW).

Typical insertion loss for a component was −11 dB (Table 5.2); insertion losses did not display a consistent relationship with waveguide width, so variations are likely to be caused by variations in waveguide quality or skill in manual coupling. Using coupling loss results from single-mode waveguides (Section 4.5.3), we may assume that there is ∼4 dB coupling loss per facet, and therefore that about 64 mW/637 pJ of power is coupled into each waveguide. The estimated propagation loss is therefore ∼3 dB/cm—this is higher than observed in single-mode waveguides, but supercontinuum-generating waveguides required a higher-energy, more selective etch and therefore had rougher sidewalls (Figure 3.17).

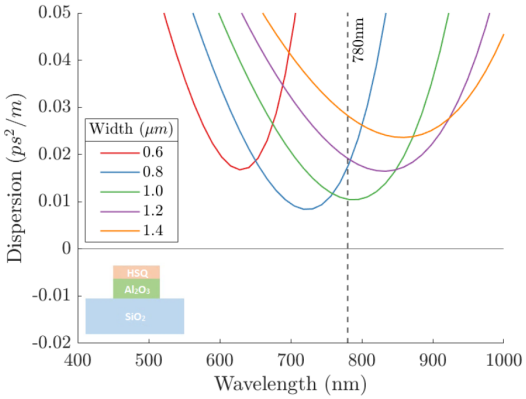

(a) GVD of TE00 mode, derived from simulations performed using Lumerical MODE FDE solver, R1.2 [258].

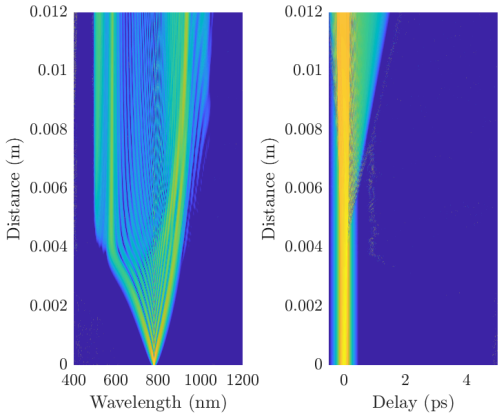

(b) Spectrum produced from 800 nm-wide waveguide pumped at 780 nm after 12 mm of propagation. Modelled using GNLSE solver written by Dudley and Taylor [133].

Figure 5.7: Simulations of dispersion and nonlinear properties of 700 nm-thick waveguides with 100 nm residual HSQ mask.

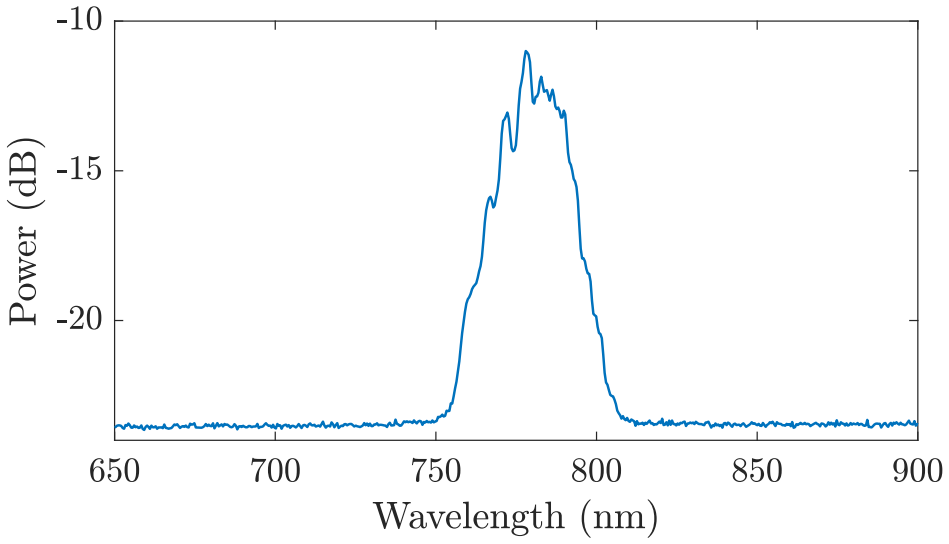

Figure 5.8: Measured output spectrum of pump laser used for supercontinuum generation, a Menlo Systems C-Fiber 780 High Power laser, at a power of 10 mW.

Effect of Waveguide Width on Generated Supercontinua

TE-polarised light produced considerably more spectral broadening than TM-polarised light; therefore, I only report TE00 results. GVD calculations (Figure 5.1) demonstrated that TE00 GVD is considerably lower than TM00 GVD throughout the range of waveguide geometries. The majority of broadening is expected to occur through SPM, which is not inherently GVD-dependent; however, it is intensity-dependent (Equation 5.6). Low GVD minimises temporal broadening and maximises localised intensity, increasing the phase change induced by SPM.

Figure 5.9 shows supercontinua obtained in waveguides ranging from 700 nm to 1400 nm wide. For wider waveguides (1100 nm to 1400 nm), SPM drives the majority of broadening, and the characteristic symmetrical maxima and minima produced by SPM are apparent. In narrower waveguides (700 nm to 1000 nm), DW formation precipitated by OWB produces additional high-frequency components.

Table 5.2 lists the spectral width of each supercontinuum at −30 dB and at the noise floor (approximately −32 dB). Whilst narrower waveguides produced spectra which were continuous from the blue to the IR, frequencies not produced by SPM are low-intensity. The noise floor of these measurements was limited by the OceanOptics HR2000+ spectrometer used to take measurements; a low integration time is required to avoid detector saturation. The noise floor can be lowered by using shortpass, longpass, and neutral density filters to capture multiple spectra which can be stitched together (see Section 5.4.2).

A photograph of light scattering from an 800 nm-wide waveguide during a full-power measurement (Figure 5.11) demonstrates the generation of red light after very little propagation, and the generation of blue light after ∼6 mm of propagation—similar to the propagation dynamics modelled by the GNLSE (Figure 5.7), which predicts 4 mm of SPM-induced broadening, followed by the generation of a DW in the blue. Optical loss within the waveguide, reducing the strength of SPM, may account for the difference in propagation distance between the model and the experiment.

Effect of Pump Power on Generated Supercontinua

To assess spectral broadening as a function of pump power, I collected spectra at different pump powers in a single 700 nm-wide waveguide. As in Section 5.4.2, I used a TE-polarised pump to maximise spectral broadening.

SPM-induced spectral broadening increases as pump power increases; the spectrum does not saturate within the range of pump power available. The short-wavelength component which has a peak at 475 nm is only apparent at pump powers above 100 mW—this feature is likely to occur as a result of OWB, and it may only occur once specific frequency components are

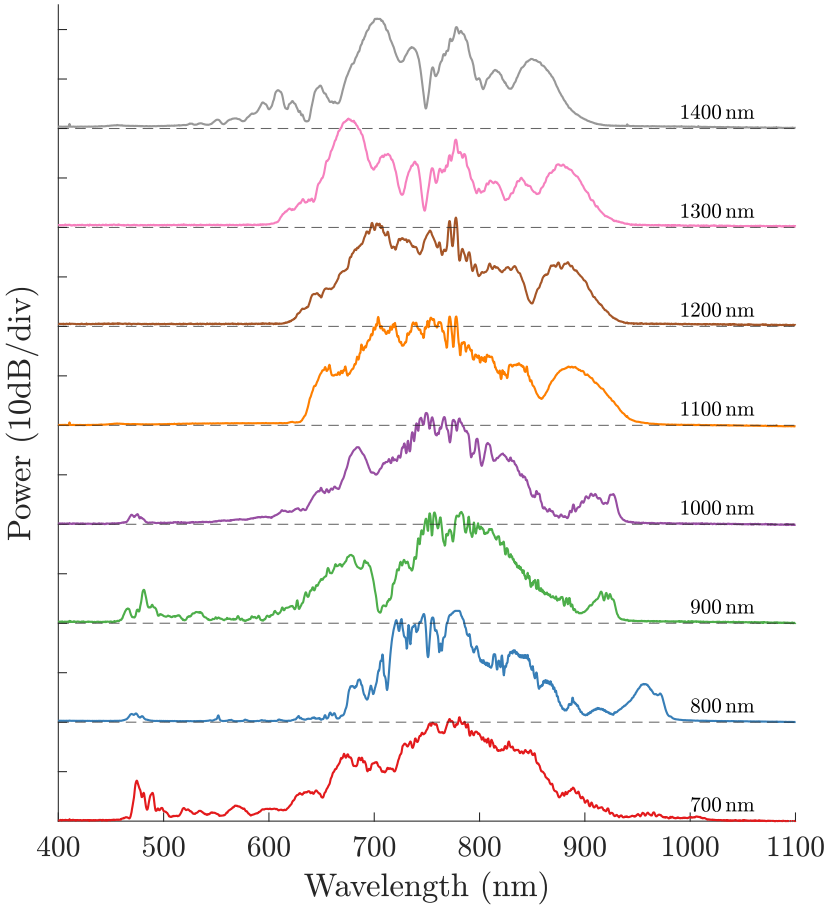

Figure 5.9: Measured supercontinua in 12 mm-long alumina waveguides with widths varying between 700 nm to 1400 nm, using a TE-polarised pump with a central wavelength of 780 nm, a pulse duration of 90 fs, and a power of 160 mW. Spectra are offset by 20 dB; dashed lines represent the noise floor of each spectrum.

Table 5.2: Width of generated supercontinua in alumina waveguides using a TE-polarised pump with a central wavelength of 780 nm, a pulse duration of 90 fs, and a power of 160 mW. Waveguides are 700 nm thick, with 100 nm of residual HSQ on top of the core layer; waveguide widths vary.

Waveguide width (nm)	Insertion loss (dB)	Spectral range at −30 dB (nm)	Spectral width at −30 dB (nm)	Spectral range above noise (nm)
700	11.8	630-924	294	459-1023
800	11.6	674-876	202	459-996
900	10.9	622-892	270	459-940
1000	11.8	647-858	211	463-942
1100	11.3	636-935	299	619-959
1200	10.9	625-927	302	615-940
1300	11.6	621-912	291	601-938
1400	10.6	587-895	308	488-922

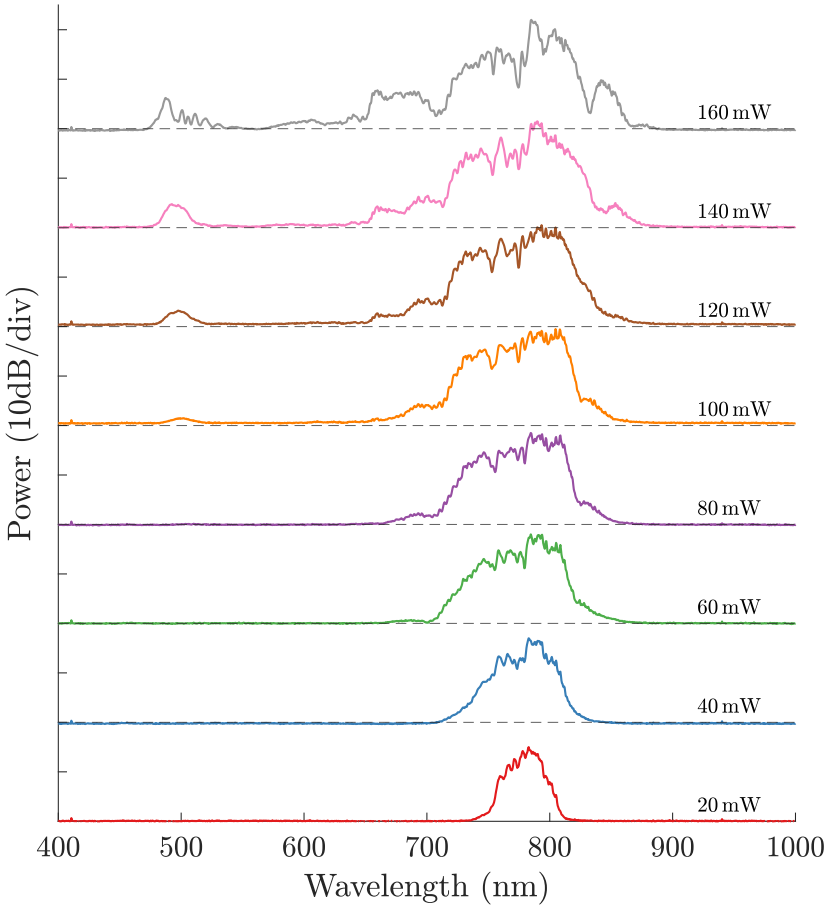

Figure 5.10: Measured supercontinua in 12 mm-long, 700 nm-wide alumina waveguides, using a TE-polarised pump with a central wavelength of 780 nm, 90 fs pulse duration, and varying pump power from 20 mW to 160 mW. Spectra are offset by 20 dB; dashed lines represent the noise floor of each spectrum.

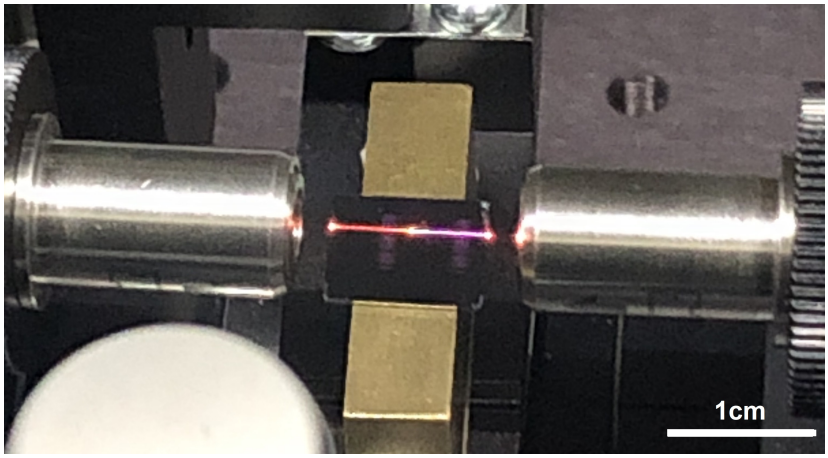

Figure 5.11: Photograph showing light scattering from 800 nm-wide, 12 mm-long alumina waveguide during SCG experiment using a pump power of 160 mW. Light travels left to right. Generation of red light is evident after $<$1 mm of light propagation; a blueshift is evident after $\sim$6 mm.

produced through SPM. Since SPM is intensity-dependent, lower pump powers likely fail to produce the required broadening.

Effect of Mode Hybridisation

A particularly broad spectrum was obtained by coupling a 5:2 ratio of TM to TE light into an 800 nm-wide waveguide, spanning from 384 nm to 951 nm at a level of −50 dB (Figure 5.12).

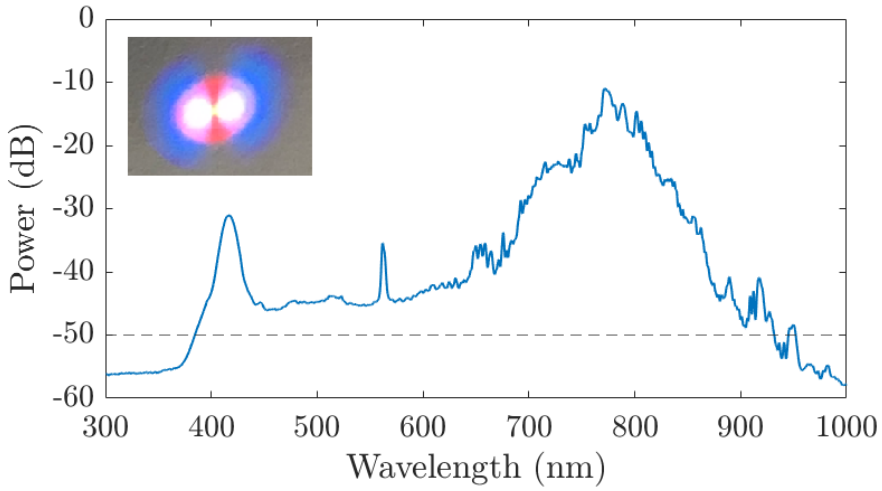

Figure 5.12: Measured supercontinuum in 12 mm-long, 800 nm-wide alumina waveguide, using a pump with a central wavelength of 780 nm, a pulse duration of 90 fs, and a power of 160 mW. The polarisation of the input light is 5:2 TM:TE. Inset photograph shows a far-field output spot.

The shape of the spectrum is similar to that obtained in the same waveguide using TE-polarised light (Figure 5.9); however, the peak typically observed at 475 nm is shifted to 418 nm. Deviating from the use of pure TE-polarised light did not produce additional spectral broadening in any other waveguide. 800 nm-wide waveguides are almost square when the 100 nm of residual HSQ mask is taken into account—therefore, the fundamental TE and TM modes have very similar n_{eff} ($\Delta < 0.001$), which facilitates mode hybridisation.

The integrated dispersion simulations described earlier in this chapter can be modified to account for coupling between modes. Rearranging β_{int} (Equation 5.24) gives

$$\beta_{DW} - \frac{\omega_{DW}}{\nu_g} \approx \beta_p - \frac{\omega_p}{\nu_g} \tag{5.26}$$

where ν_g is the group velocity of the pump β_{DW} and β_p are the propagation constants of the DW and pump, and ω_{DW} and ω_p are the angular frequencies of the DW and pump. Figure 5.13a shows the wavelength-dependence of $\beta_{DW} - \frac{\omega_{DW}}{\nu_g}$ for the TE00 mode and the phase-matching conditions for the TE00 and TM00 modes. The phase-matching condition for a TE00 DW with a TE00 pump is at 542 nm, whilst phase-matching of a TE00 DW with a TM00 pump is

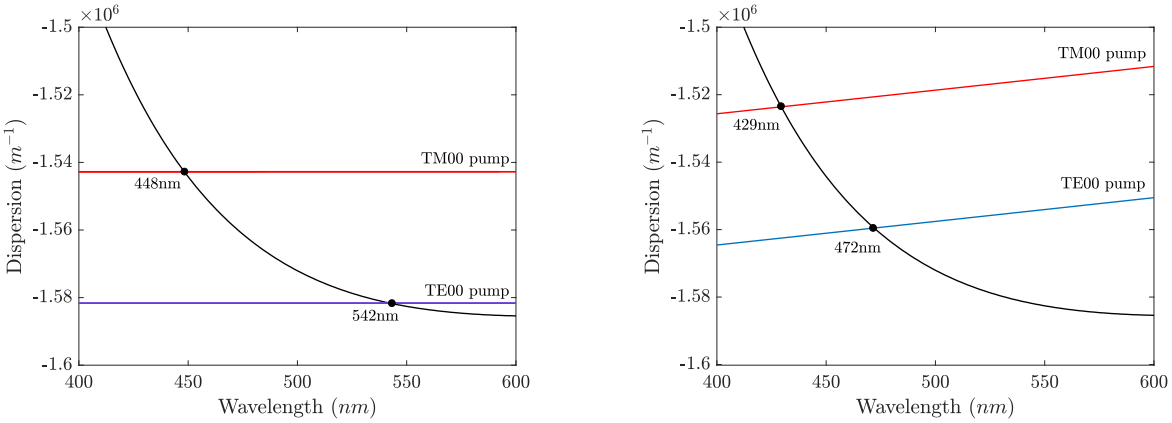

(a) Excluding nonlinear term (Equation 5.26).

(b) Including nonlinear term (Equation 5.27).

Figure 5.13: Dispersive wave phase-matching conditions. The black curve represents DW dispersion for the TE00 mode, whilst the red and blue lines represent the pump dispersion.

at 448 nm. TE00-TM00 mode hybridisation can therefore blueshift the DW phase-matching condition.

However, the phase-matching conditions derived from 5.26 do not match the DW positions observed during experiments. Further refinement of the simulation is required. Equation 5.24 and 5.26 remove the nonlinear term from 5.18. With this term included, Equation 5.26 becomes

$$\beta_{DW} - \frac{\omega_{DW}}{\nu_g} = \beta_p - \frac{\omega_p}{\nu_g} + \frac{\gamma P_0}{2} \tag{5.27}$$

where γ is the effective nonlinear coefficient (Equation 5.10) and P_0 is the peak pump power. Given the high peak power used in these experiments, this term becomes non-negligible. Since γ is wavelength-dependent, the phase-matching condition also becomes wavelength-dependent; in Figure 5.13b, the nonlinear contribution is added to the lines representing TE00 and TM00 pumps. Now, phase-matching of a TE00 DW with a TE00 pump occurs at 472 nm, and phase-matching of a TE00 DW with a TM00 pump occurs at 429 nm, where the measured values are 475 nm and 416 nm, respectively.

The noise floor for the majority of spectra reported in this chapter was approximately −32 dB, limited by detector saturation. In order to confirm the generation of UV light, I acquired the spectrum in three parts, using a shortpass filter to block light above 650 nm, a longpass filter to collect light above 900 nm, and a neutral density filter (optical density 2.0) to collect a full spectrum. I used characteristic peaks in each spectrum to scale the short- and long-wavelength components to the primary spectrum, producing a noise floor of −55 dB.

5.4.3 Discussion

This work is the first demonstration of SCG in alumina waveguides, and the first demonstration of any nonlinear optical process in an etched alumina waveguide. Efficient frequency conversion through third-order nonlinear processes requires a GVD near zero, and a film thickness of 800 nm is required to achieve zero GVD in ridge waveguides made using the alumina deposited in this work (Figure 5.1). Prior to this work, no alumina waveguides thicker than 120 nm had demonstrated low-loss operation in the visible region [132]. Chapters 2-4 of this work addressed two barriers to SCG in alumina waveguides, realising relatively repeatable deposition of thick optical-quality films and the capacity to etch low-loss (low-roughness) waveguides with selectivity sufficient to achieve anomalous GVD. Whilst the fabricated supercontinuum-generating waveguides were underetched, they demonstrated near-zero GVD; the thickness of the residual mask after etching demonstrates that the fabrication process is capable of producing waveguides with anomalous GVD, in which soliton dynamics may produce greater spectral broadening. All waveguides measured had relatively low insertion losses—maximising the light retained within the waveguide maximises the distance over which nonlinear effects can occur and produces a brighter output spectrum.

Supercontinua generated using both TE00 and mixed polarisation span an octave—they extend through a frequency f to its double $2f$. Octave-spanning spectral broadening is a common requirement in frequency combs for optical timing applications, which require extremely precise and stable outputs; f-$2f$ self-referencing can be used to used to quantify and account for intra-pulse temporal incoherence caused by chromatic dispersion and nonlinearities. Figure 1.1 shows a range of atomic emissions which are of interest for optical timing applications at wavelengths ranging from 370 nm to 800 nm, all of which are accessible using supercontinua generated in this work.

Whilst this work demonstrates frequency conversion in a small form factor, a high input pulse energy (1.6 nJ) was required, necessitating use of a bulky external pump source. Fully miniaturised PICs for applications such as optical timing or sensing will require integration or monolithic fabrication of sources with a small form factor, which is particularly challenging when short pulse times and high pulse energies are required. Recent work by Singh et al. demonstrates an on-chip Q-switched laser, generating nanojoule-level picosecond pulses using thulium-doped alumina as a gain medium [290]; on-chip amplification such as this may offer a future solution for integration of sources with small form factors, especially if nonlinear techniques can be used to compress pulses to the femtosecond range [291]. Supercontinuum generation could then offer a way to access a wide range of frequencies using a single source in the IR.

Wider Context

Figure 5.14 compares the supercontinua generated in this work to results reported in established material platforms for waveguide SCG, such as AlN, $LiNbO_3$, and SiN. SiN is a well-established material platform, deposition and fabrication processes are mature, and sophisticated dispersion engineering strategies have been employed to produce exceptionally broad supercontinua [139, 292]. However, generation of UV light has never been achieved in SiN waveguides. Approaching the material bandgap, altering waveguide geometry is insufficient to counteract strongly normal material dispersion, which limits the temporal coherence required for many nonlinear effects; furthermore, SiN's high propagation losses below 500 nm [30] may result in the loss of any short-wavelength light which is generated. Conversely, the alumina films used in this project are transparent above 250 nm (Figure 2.5) and do not demonstrate a strong variation in refractive index at wavelengths above 400 nm (Figure 2.3), allowing low GVD at considerably shorter wavelengths.

AlN and $LiNbO_3$ platforms benefit from $\chi^{(2)}$ nonlinearities; since $\chi^{(2)}$ is typically several orders of magnitude greater than $\chi^{(3)}$, waveguides made in these materials may require relatively low pump powers to achieve spectral broadening. At present, results in AlN may be limited by film quality [95, 96]; however, commercially available thin-film $LiNbO_3$ offers a high-quality material platform which has demonstrated guiding at wavelengths as low as 320 nm. Recent work by Wu et al. used periodically poled $LiNbO_3$ to produce cascaded SHG, resulting in a supercontinuum extending to 330 nm using a low pump energy (90 pJ) [293].

Silica is perhaps the most similar material platform to amorphous alumina. The material's wide bandgap has allowed generation of light at wavelengths as short as 265 nm, although wavelengths below 450 nm have not been quantified relative to pump or input power. However, the material's low n_2 means that a considerably higher-energy pump (2300 pJ) is required to stimulate broadening compared to other materials with higher nonlinearity. The alumina waveguides used in this work also required nanojoule-level pulse energies to generate appreciable levels of spectral broadening, and spectra did not saturate even at the maximum pulse power of 1.6 nJ. SiN, in which broadening takes place primarily through $\chi^{(3)}$ effects but n_2 is an order of magnitude higher, provides a comparison point: octave-spanning supercontinua can be generated at pulse energies below 0.5 nJ [139, 292].

Amorphous alumina's wide bandgap facilitates the nonlinear generation of short-wavelength light. Given the material's low n_2, high pump energies are required to stimulate broadening; other wide-bandgap materials, such as AlN and $LiNbO_3$, have higher n_2 and also offer efficient frequency conversion through $\chi^{(2)}$ processes. However, repeatable deposition of optical-quality AlN films remains challenging, and $LiNbO_3$ can pose fabrication cleanliness challenges. These concerns are mitigated in an amorphous alumina material platform:

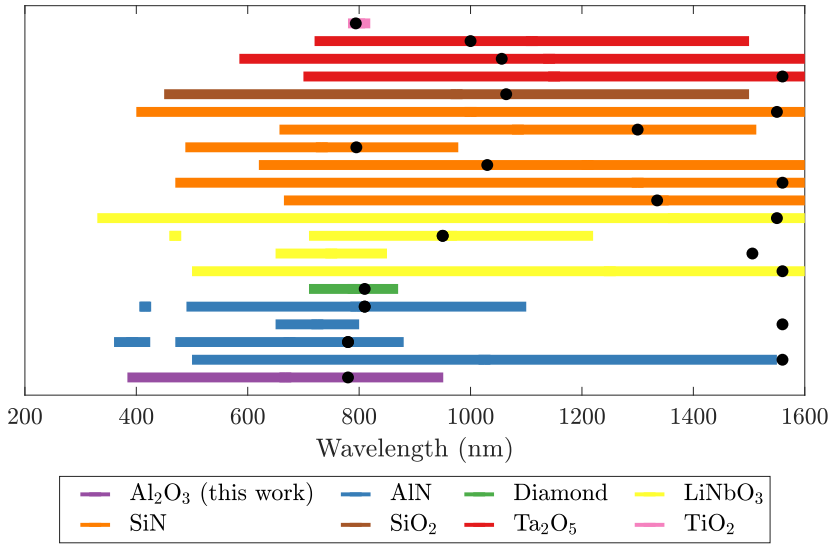

Figure 5.14: Review of supercontinuum generation within optical waveguides. In order to be included in this review, waveguides must have generated light below $\lambda = 800\,\text{nm}$ at a level of at least $-60\,\text{dB}$. Black markers indicate pump wavelengths. Full details are provided in Table A.2.

optical-quality films can be deposited repeatably using ALD, and high-confinement alumina waveguides can be fabricated using a complementary metal-oxide semiconductor (CMOS)-compatible process; waveguides' low insertion losses contribute to spectral brightness.

Future Improvements

This is a promising first demonstration of supercontinuum generation in alumina waveguides; there are numerous avenues to explore which could produce brighter supercontinua, broader supercontinua, or broadening with a lower pump power.

Firstly, the fabricated waveguides were underetched, and operated in the normal regime. Spectral broadening was driven by SPM, in which the induced phase change is directly proportional to light intensity (Equation 5.6). Simulations indicate that if the same waveguides were etched a further 100 nm, producing GVD in the anomalous regime, broadening mechanisms are dominated by soliton fission (Figure 5.4). This is predicted to produce a wider supercontinuum, and may also allow for reduced pump power, since soliton formation and fission is not directly linked to the intensity of the pulse.

Work in other material systems points to a number of ways to produce broader or brighter supercontinua. In this work, I created waveguides with a uniform cross-section, selected primarily on the basis of GVD; an alternative strategy is to optimise cross-section based on the output of the GNLSE. Altering a waveguide's width along its length will alter its dispersion properties and stimulate different broadening mechanisms along its length. Sophisticated

dispersion management techniques typically use optimisation algorithms to design a desirable output [294].

I optimised the dispersion properties of the TE00 mode in this work; however, a mixture of TE and TM-polarised light generated the greatest blueshift. A multimode GNLSE solver could identify the dynamics which increase spectral broadening in this instance, but implementation requires considerable computational power and adoption of parallel computing techniques [295]. The processes developed in this work offer a maximum etch depth of 800 nm, which is very close to the waveguide thickness required to attain zero GVD in the TE00 mode of an alumina ridge waveguide; excitation of higher-order modes can allow low GVD to be attained using a wider range of waveguide geometries [144]. Recent work by Kou et al. provides a powerful example of combining these techniques: using a tapered waveguide which stimulates broadening in both fundamental and higher-order modes, the authors demonstrate a two-octave supercontinuum in SiN [139].

5.5 Conclusion

In this chapter, I report a supercontinuum spanning 384 nm to 951 nm at a level of −50 dB in a PEALD alumina ridge waveguide. This is the first demonstration of nonlinear optical processes in etched alumina waveguides. This spectrum spans an octave and spans the entire visible region, allowing access to a wide range of frequencies of interest in optical timing and sensing applications.

Most nonlinear optical processes achieve high conversion efficiency when pumped near zero GVD. The amorphous alumina used in this work has a wide bandgap (∼5.75 eV/216 nm), producing GVD minima near 780 nm in submicron waveguides. This is a relatively short pump wavelength, which can readily broaden to span the visible and UV. An etch depth of 800 nm is required to produce anomalous GVD at the pump wavelength; however, prior to this work, no alumina waveguides thicker than 120 nm had been demonstrated in the visible region. I developed processes for etching high-confinement alumina waveguides (Chapter 3), allowing access to the anomalous dispersion regime in waveguides with low propagation losses. The supercontinuum-generating waveguides I fabricated were underetched, but pumping in the low normal regime produced symmetrical spectral broadening through SPM, with outlying frequencies generated through OWB. The waveguides had low insertion losses (−11 dB), which allows for long nonlinear interaction lengths and produces a relatively bright output spectrum.

This is a promising first demonstration of nonlinear broadening in alumina waveguides; future work will refine the waveguide design with the aim of extending the supercontinuum further into the UV. Pumping in the anomalous regime, as originally intended, may allow access to a

wider range of frequencies at a lower pump energy. Designing tapered waveguides based on output from the GNLSE can produce dispersion-managed waveguides which are optimised for brightness, flatness, or blueshift.

Chapter 6

Conclusion

Amorphous alumina is a noteworthy material in low-loss UV-visible integrated optics because of its exceptionally wide bandgap. In this work, I developed an amorphous alumina material platform for integrated linear and nonlinear photonics in the UV and visible regions. Previous demonstrations of low-loss alumina waveguides have focused on low-confinement waveguides; however, waveguide-based nonlinear optics requires high-confinement waveguides. ALD offers a repeatable method for depositing smooth and uniform amorphous alumina films; however, conventional thermal ALD processes are prohibitively slow, whilst poor etch selectivity has previously prevented smooth, optical-quality etching of thick waveguides. In this work, I established efficient and repeatable deposition and patterning processes for thick, optical-quality alumina films, and used these processes to fabricate low-loss single-mode waveguides and supercontinuum-generating waveguides. Fabricated single-mode waveguides are the lowest-loss high-confinement ridge waveguides in the blue and green; supercontinuum-generating waveguides demonstrate the first nonlinear optical processes in etched alumina waveguides. This is the first demonstration of a continuous UV-to-IR spectrum generated using only $\chi^{(3)}$ nonlinearities. This work therefore represents a progression in the state of the art for UV-visible integrated optics. Here, I consider the primary contributions of this work, routes to improvement of the devices I developed, and this work's implications for integrated optical devices and applications.

6.1 Optical-Quality PEALD Alumina Films

Whilst PEALD of amorphous alumina is well-established, this work provides some of the first data on dioxygen PEALD alumina films' suitability for optics applications. I demonstrate that deposited films are stable for at least 2 years at room temperature and humidity, that they can be heated to 400 °C without appreciable alteration in film properties, and that they

are sufficiently stable in cold water, organic solvents, and developers used in waveguide fabrication processes. To quantify the transparency of unpatterned films, I measured film absorption in the UV; deposited films' estimated bandgap is 5.75 eV (216 nm), which is in concordance with previously reported UV transmission-based measurements of ALD alumina films. Low propagation losses in single-mode waveguides fabricated using three different deposition runs offer the most decisive proof of deposited films' suitability for integrated optics applications.

6.2 Fabrication of High-Confinement Ridge Waveguides

Prior to this work, no alumina waveguides thicker than 120 nm had demonstrated low losses in the visible region. Given alumina's poor etch selectivity to common mask materials, aggressive etching, which produces rough sidewalls, is typically required to produce ridge waveguides. In this work, I used HSQ as an etch mask: whilst its etch rate in halide chemistries is typically higher than that of alumina, its low granularity and EBL compatibility allow production of extremely smooth etch masks. I established that uniform HSQ masks can be fabricated using film thicknesses up to 1 µm, adjusted EBL parameters to minimise lithographic roughness, and developed BCl_3 etching processes which balance selectivity and sidewall roughness. I designed a low-roughness etch process suitable for waveguides up to 650 nm thick, and a higher-selectivity process, which compromises on roughness, for waveguides up to 800 nm thick.

Fabricated high-confinement PEALD alumina waveguides had unclad losses of 1.3 dB/cm and clad losses of 0.8 dB/cm at 450 nm—this propagation loss value lies below the absorption limit of SiN, the current leader in the field, and is the lowest loss results reported in high-confinement waveguides in the blue. Low-confinement single-mode waveguides produced TM00 losses of 0.6 dB/cm at 450 nm, in line with previous demonstrations in thermal ALD alumina waveguides.

6.3 Supercontinuum Generation

I used my high-selectivity fabrication process to produce 700 nm-thick, unclad alumina waveguides. Deposited alumina films' wide bandgap produced a GVD minimum near the pump wavelength of 780 nm—this relatively short pump wavelength facilitated spectral broadening through the visible region and into the UV. Operating in the low normal regime, spectral broadening occurs through SPM, with outlying frequencies generated through OWB

and DWs once sufficient spectral and temporal broadening has occurred. This work is the first demonstration of nonlinear optical processes in an etched alumina waveguide.

The broadest supercontinuum produced in this work spanned 384 nm to 951 nm at a level of −50 dB—the first demonstration of continuous UV-to-IR light generation using only $\chi^{(3)}$ nonlinearities, which are typically less efficient than the $\chi^{(2)}$ processes which have been used to produce similarly broad and bright spectra. All waveguides had low linear losses, which maximises optical power retention, maintains nonlinear broadening over a considerable propagation distance, and produces a bright output spectrum.

6.4 Future Directions

Whilst I have produced alumina waveguides with low propagation losses, I performed limited optimisation work. At present, there is little information explicitly linking ALD parameters, structural properties such as bonding composition, impurity levels, and microstructure, and optical properties such as optical bandgap and other absorbances in the visible region. Many of the tools to perform this type of work are readily available—for example, ^{27}Al NMR and XPS studies have already been used to establish aluminium and oxygen coordination in amorphous films [108, 183]. There are only a few papers using amorphous alumina for low-loss integrated optics, which may explain the lack of research in this area: I hope that recent demonstrations, including this work and related publications [76, 165], will increase awareness of this application and stimulate further work linking deposition parameters, structure, and optical quality.

The waveguide definition work presented in this thesis also lacks a quantitative link between process parameters, waveguide roughness, and optical loss. Directly quantifying sidewall roughness is challenging, but recent work by Corato-Zanarella et al. demonstrates a convenient method of extracting quantitative sidewall roughness data from optical loss measurements [2]. Further process optimisation work could further reduce propagation losses in alumina waveguides—altering etch gas mixes and flow rates may produce smoother and more selective etches, whilst higher-quality cladding materials such as LPCVD or TEOS-PECVD may reduce cladding losses, especially in low-confinement waveguides.

The high-confinement waveguides demonstrated in this work do not have considerably higher losses than low-confinement alumina waveguides demonstrated in this and other work; it is worth investigating whether, for ALD alumina in the visible region, high-confinement waveguides should become a new standard. Considering solely linear applications, they offer a number of advantages—smaller bend radii for structures allow for minimal radiation losses in a smaller footprint, and the resulting reduction in propagation distance may actually reduce propagation loss compared to the ∼100 nm-thick waveguides which are presently most

common. Thick waveguides also allow for design flexibility when producing components such as grating couplers. Ongoing work aims to encapsulate nanodiamonds containing NV centres into circular Bragg grating resonators fabricated in alumina films: when excited at 532 nm, these point defects exhibit long-lived spin-dependent luminescence at 637 nm, and creating grating structures allows directional out-coupling and measurement of emitted photons [25, 296]. Alumina films' low propagation losses, including in high-confinement structures (gratings use 500 nm-thick films), facilitate this application. Longer-term, integration of nanodiamonds with defect centres into PICs may allow for long-lived spin states to be used for information processing.

Supercontinua generated in this work span the entire visible region and extend into the UV—they span a wide range of UV-visible frequencies which are of interest for optical timing and metrology (Figure 1.1). The pulse power required to produce broadening into the blue and UV is greater than in many other common nonlinear materials with higher n_2; tapered waveguides and stimulation of nonlinearities in higher-order modes could produce broadening at lower pump powers, or produce greater broadening at high pump power. Optical timing and metrology applications require frequency combs, produced by generating a supercontinuum in a tailored MRR; f-$2f$ self-referencing can be used to stabilise octave-spanning combs. The supercontinua generated in this work span an octave, and would be suitable for use with this stabilisation scheme.

Of course, the long-term aim of performing nonlinear optics in etched waveguides is the production of chip-scale systems. On-chip pulsed lasers in the IR do not yet offer the pulse width required to perform fully-integrated SCG, with recent demonstrations producing picosecond pulses [290]; however, waveguide-based nonlinear optical processes can be used to produce compressed pulses [291]. Supercontinuum generation could then offer a way to access a wide range of frequencies using a single source in the IR.

Appendix A

Reviews

Table A.1: Propagation losses in ridge waveguides at wavelengths below 800 nm.

Authors	Institution	Material	Deposition method	Year	Loss (dB/cm)	λ (nm)	Core thickness (nm)	Core width (nm)	Cladding
West/ Sorace-Agaskar [4]	MIT	Al_2O_3	ALD	2019	3 1.8 1.2 0.7	371 405 419 458	100	400-600	SiO_2
Lin/ le Thomas [13, 132]	Ghent	Al_2O_3	ALD	2021 2022	5 3	402 360	120 70	700 600	None SiO_2
Hendriks/ García-Blanco [117]	Twente/ Aluvia	Al_2O_3	Sputter	2023	2	405	70	600	SiO_2
McKay/Sorel [76]	Glasgow	Al_2O_3	ALD	2023	0.8 0.6 1.3 0.8	450 642 450 642	400	700 1200 700 1200	HSQ None
Stegmaier/ Pernice [94]	Karlsruhe	AlN (polycrystalline)	Sputter	2014	38.9	450	200	170	HSQ
Lu and Fanto/ Englund [95]	MIT	AlN (polycrystalline)	PVDNC	2018	75 5.3	370 638	200	250 450	SiON PMMA
Liu/Tang [96]	Yale	AlN (single crystal)	MOCVD	2018	8.0 3.5	390 455	500	400-800	SiO_2
Gao/ Vermeulen [98]	Brussels	Diamond	Microwave plasma CVD	2018	46.7 20.5	405 635	140 250	300	SiO_2
Chen/Zhao [297]	Arizona	GaN	MOCVD	2017	2	770	1500	1200-2000	SiO_2
Romero-Garcia/ Witzens [298]	Aachen	SiN	PECVD	2013	0.51	660	100	700	SiO_2
Subramanian/ Baets [262]	Ghent/imec	SiN	PECVD	2013	1 2.5 1	532 532 780	180 180 220	400 500 800	SiO_2 None None
Shim/Li [299]	Minnesota	SiN	LPCVD	2016	2.9	473	200		SiO_2
Sorace-Agaskar [300]	MIT	SiN	PECVD	2018	10.35 4.16 0.62	405 458 634	100		SiO_2
Sacher/Poon [301]	CalTech/ Planck Institute	SiN	LPCVD	2019	3.9-9.2 2.5-4.5 0.5-1.4	430-464 430-464 602-678	200 135 200	520	SiO_2
Sinclair/Paul [302]	Glasgow	SiN	LPCVD	2020	0.2	780	200	1100	SiO_2
Zhao/Gaeta [303]	Columbia	SiN	LPCVD	2020	0.73	780	730	1330	SiO_2
Domeniguetti/ Nussenzveig [304]	São Paulo	SiN		2021	0.25	780	730	2000	SiO_2
Morin/ Bowers [87]	UC Santa Barbara	SiN	LPCVD	2021	0.93	405	24	800	SiO_2
Chauhan/ Blumenthal [254]	UC Santa Barbara	SiN	LPCVD	2022	0.09 0.01	461 674	20	1300 2300 2300	SiO_2
Lelit/ Piramidowicz [305]	Warsaw	SiN	LPCVD	2022	1.71	660	320	1000	SiO_2
Smith/Strain [66]	Strathclyde	SiN	LPCVD (Ligentec)	2023	6.85 4.39 0.98 0.87	450 520 630 750	150		SiO_2
Lindecrantz/ Hellesø[306]	UiT Tromsø	Ta_2O_5	Sputter	2014	6.7	785	215	1500	None
Lin/Lee [307]	Sun Yat-Sen, Taiwan	Ta_2O_5	Sputter	2016	2.4	800	400	700	SiO_2
Sierra/ Carvalho [308]	São Paulo	Ta_2O_5	Sputter	2019	1.6	785	430	5000	None
Choy/Lončar [309]	Harvard	TiO_2	Sputter	2012	28	633	170	250	None
Evans/ Suntivich [310]	Cornell	TiO_2	Sputter	2015	7.5	633	250	250	None
Hegeman/ García-Blanco [58]	Twente	TiO_2	Sputter	2020	5.08	633	140	1200	None

Table A.2: Experimental results for supercontinuum generation in the visible region. To be listed, a paper must report a power of −50 dB or higher at a wavelength below 800 nm. If the paper reports the supercontinuum width at a power level lower than −30 dB, then two wavelength ranges are provided here: the reported supercontinuum width, and the supercontinuum width at −30 dB. If no result at −30 dB is listed, this indicates that no light with a power level of −30 dB was attained at any wavelength below 800 nm.

Authors	Institution	Material	Year	Wavelength range (nm) at spectral flatness (dB)			Pump wavelength (nm)	Pump power (pJ)	Waveguide type	Cladding	Waveguide height (nm)
Hickstein/ Papp [311]	NIST/Yale	AlN	2017	750-900 500-1550	at at	-30 -50	1560	800	Ridge	SiO_2	800
Liu/Tang [143]	Yale	AlN	2019	650-800 360-425 470-880	at at at	-30 -50 -50	780	237			1000
Chen/Zhao [144]	Arizona	AlN	2021	740-795 407-425 490-1100	at at at	-30 -60 -60	810	600	Ridge	SiO_2	1200
Shams-Ansari/ Lončar [312]	Harvard	Diamond	2019	710-870	at	-60	810	187	Pedestal	None	500
Lu/Tang [313]	Yale	$LiNbO_3$	2019	500-2000	at	-60	1560	800	Pedestal	None	600
Yu/Lončar [314]	Harvard	$LiNbO_3$	2019	650-850	at	-50	1506	185	Rib	None	800
Escalé/ Grange [315]	ETH Zurich	$LiNbO_3$	2020	460-480 710-1220	at at	-30 -30	950	67	Rib	None	400
Wu/ Diddams [293]	NIST	$LiNbO_3$	2024	330-2400	at	-50	1550	90	Rib	None	700
Halir/ Gaeta [292]	Columbia	SiN	2012	666-2025	at	-30	1335	160	Ridge	SiO_2	400
Epping/ Boller [140]	Twente/ LioniX	SiN	2015	470-2130	at	-30	1560	590	Ridge	SiO_2	1000
Johnson/ Gaeta [316]	Cornell	SiN	2015	670-700 620-1800	at at	-30 -50	1030	437	Ridge		900
Zhao/ Baets [317]	Ghent/imec	SiN	2015	488-978	at	-30	795	87	Ridge	SiO_2	300
Okawachi/ Gaeta [318]	Columbia	SiN	2017	657-1513	at	-30	1300	260			700
Kou/ Yamada [139]	AIST Tsukuba	SiN:D	2023	450-1600 400-1600	at at	-50 -60	1550	422	Ridge	SiO_2	1000
Yoon Oh/ Vahala [139]	Caltech/ NIST	SiO_2	2017	450-1500	at	-50	1064	2300	Pedestal	None	830
Fan/Lee [141]	Sun Yat-Sen, Taiwan	Ta_2O_5	2019	585-1697	at	-30	1056	49	Ridge	None	700
Lamee/ Papp [59]	NIST	Ta_2O_5	2020	700-2500	at	-50	1560	900	Ridge	None	570
Woods/ Apostolopoulos [142]	Southampton	Ta_2O_5	2020	720-1500	at	-50	1000	2190	Ridge	SiO_2	700
Evans/ Mazur [319]	Harvard	TiO_2	2013	780-820	at	unspecified	794	48	Ridge	Fluoro-polymer	250

Appendix B

Process Flows

B.1 Process Flows for Optical Devices

Table B.1: Process flow for fabrication of single-mode alumina waveguides.

Operation	Tool	Recipe	Notes
Inspect	Optical microscope		Check for defects, contamination
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Meas. roughness	AFM	ScanAsyst mode, 5 µm × 2.5 µm	Record r.m.s. roughness
Meas. film thickness	Ellipsometer	23 pts, multiple angle	
Al_2O_3 growth	FlexAL	Plasma ALD, Al_2O_3 200°C	Target thickness 400 nm
Inspect	Optical microscope		Check for film defects
Meas. roughness	AFM	ScanAsyst mode, 5 µm × 2.5 µm	Record r.m.s. roughness
Meas. film thickness	Ellipsometer	23 pts, multiple angle	
Cleave wafer	Manual		14 mm × 20 mm pieces
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Spin HSQ	Manual spinner 2	Neat HSQ, spin 2000 rpm, 1 min	Target thickness 700 nm
Bake	Hotplate	92°C, 15 min	
Expose – e-beam	EBPG 5200	1450 µC/cm^2	
Bake	Hotplate	180°C, 1 min	
Develop	Wet deck	25% TMAH 30 s, 2x DI water 30 s, IPA 30 s, N_2 dry	
Develop inspect	Optical microscope		
Alumina etch	ICP180	BCl_3	
Meas. profile	Dektak XT	100 µm, 30 sec	2 scans
Spin resist	Manual spinner 1	PMMA 200MW anisole, spin 4000 rpm, 1 min	
Bake	Hotplate	180°C, 5 min	
Cleave	Manual		Cleave waveguide facets
Resist strip	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	

Table B.2: Process flow for fabrication of supercontinuum-generating alumina waveguides.

Operation	Tool	Recipe	Notes
Inspect	Optical microscope		Check for defects, contamination
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Al_2O_3 growth	FlexAL	Plasma ALD, Al_2O_3 200°C	Target thickness 400 nm
Al_2O_3 growth	FlexAL	Plasma ALD, Al_2O_3 200°C	Target thickness 300 nm
Inspect	Optical microscope		Check for film defects
Meas. roughness	AFM	ScanAsyst mode, 5 µm × 2.5 µm	Record r.m.s. roughness
Meas. film thickness	Ellipsometer		
Cleave wafer	Manual		10 mm × 20 mm pieces
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Spin HSQ	Manual spinner 2	Neat HSQ, spin 1000 rpm, 1 min	Target thickness 1000 nm
Bake	Hotplate	92°C, 15 min	
Expose – e-beam	EBPG 5200	1600 µC/cm^2	
Bake	Hotplate	180°C, 1 min	
Develop	Wet deck	25% TMAH 30 s, 2x DI water 30 s, IPA 30 s, N_2 dry	
Develop inspect	Optical microscope		
Alumina etch	ICP180	BCl_3	Selectivity ∼0.8
Spin resist	Manual spinner 1	PMMA 200MW anisole, spin 4000 rpm, 1 min	
Bake	Hotplate	180°C, 5 min	
Cleave	Manual		Cleave waveguide facets
Resist strip	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	

B.2 Process Flows for Fabrication Development

Table B.3: Process flow used for HSQ spinning tests described in Section 3.2.4.

Operation	Tool	Recipe	Notes
Cleave Si wafer	Manual		18 mm × 18 mm pieces
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Spin HSQ	Manual spinner 2	Neat HSQ, various speeds, 1 min	
Bake	Hotplate	92°C, 15 min	
Meas. HSQ thickness	Ellipsometer	1 cm × 1 cm, 5 pts	

Table B.4: Process flow used for line-edge roughness tests described in Section 3.2.6.

Operation	Tool	Recipe	Notes
Cleave Si wafer	Manual		20 mm × 20 mm pieces
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Spin HSQ	Manual spinner 2	Neat HSQ, spin 2000 rpm, 1 min	Target thickness 700 nm
Bake	Hotplate	92°C, 15 min	
Expose – e-beam	EBPG 5200	Various parameters	
Bake	Hotplate	Various parameters	
Develop	Wet deck	25% TMAH 30 s, 2x DI water 30 s, IPA 30 s, N_2 dry	
Line-edge roughness inspection	SEM	10 kV, 15 nA	Inspect straight and curved structures at 100kx and 150kx magnification
Alumina etch	ICP180	BCl_3, 5 min	
Line-edge roughness inspection	SEM	10 kV, 15 nA	Inspect straight and curved structures at 100kx and 150kx magnification

Table B.5: Process flow for etching tests described in Section 3.3.

Operation	Tool	Recipe	Notes
Cleave wafer	Manual		18 mm × 18 mm pieces
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Spin HSQ	Manual spinner 2	Neat HSQ, spin 2000 rpm, 1 min	Target thickness 700 nm
Bake	Hotplate	92°C, 15 min	
Expose – e-beam	EBPG 5200	1450 µC/cm^2	Expose test pattern
Bake	Heraeus oven	180°C, 1 hour	
Develop	Wet deck	25% TMAH 30 s, 2x DI water 30 s, IPA 30 s, N_2 dry	
Develop inspect	Optical microscope		
Alumina etch	ICP180	BCl_3	Attach HSQ monitor to carrier
Meas. profile	Dektak XT	100 µm, 30 sec	2 scans
SEM	SU8240	3 kV, 5 nA	

Table B.6: Process flow used for HSQ monitor samples for use in etching tests described in Section 3.3.

Operation	Tool	Recipe	Notes
Cleave Si wafer	Manual		10 mm × 10 mm pieces
Solvent clean	Wet deck	Acetone 3 min, IPA 3 min, N_2 dry	
Plasma ash	Barrel Asher	O_2 ash, 2 min, 150 W	
Spin HSQ	Manual spinner 2	Neat HSQ, various speeds, 1 min	
Bake	Hotplate	92 °C, 15 min	
Bake	Oven	180 °C, 1 hr	
Meas. HSQ thickness	Ellipsometer	1 cm × 1 cm, 1 pt	
Alumina etch	ICP180	Various parameters	
Meas. HSQ thickness	Ellipsometer	1 cm × 1 cm, 1 pt	

Appendix C

Other Sources of Optical Loss

C.1 Coupling Loss

Coupling light from a source into a waveguide is challenging: in commercial PICs, heterogeneous and hybrid integration present a considerable impediment to scale-up. In academic settings, source integration is less common; use of an external, commercially available, well-controlled laser removes some complexity, variability, and technical challenge from experiments. Using an external source, there are two key methods of coupling light into a waveguide. Grating couplers, described above, are created by etching waveguide core material to form a periodic grating designed to direct light in and out of the waveguide. End-fire coupling is far simpler: a waveguide facet is defined through cleaving or dicing, and light from the source is focused on the facet.

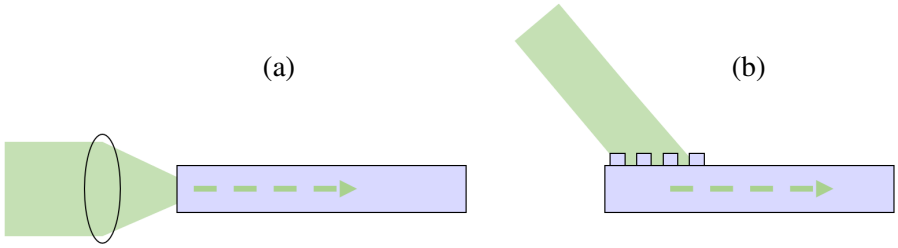

Figure C.1: Methods for coupling light into a waveguide. a) End-fire coupling. b) Grating coupling.

In any source integration process, coupling losses remain unavoidable—a result of light reflection or scattering off of the facet, poor alignment, or poor matching of the source's beam field to the waveguide's mode field [32]. Quix, a commercial photonic computer manufacturer, reports losses of 0.9 dB/facet, or 19% loss in power per facet, in their most recent product range [320].

C.2 Radiative Loss

Radiative losses refer to leakage of the light from waveguide modes into unguided modes. Many PICs are fabricated on a Si carrier wafer, which does not play any role in the optical processing. Si has a refractive index of > 3.5, higher than that of most waveguide materials; if the evanescent field from the propagating modes (see Section 4.1) extends to the Si substrate, then unguided *leaky modes* will dominate. In most optical waveguides, the presence of a suitably thick bottom cladding layer mitigates this kind of loss [26].

Waveguide bending losses, however, play a pivotal role in PIC design. A more robust treatment of waveguide modes is provided in Section 4.1—here, we may consider a mode to be guided light with the optical field necessarily remaining the same in the z-direction. Within a mode, the phase front must be maintained; this condition poses problems when a mode encounters a

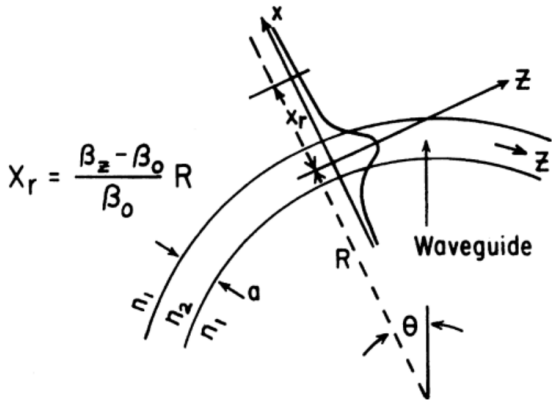

Figure C.2: Diagram illustrating the role of the phase front in bending loss. In order to maintain the phase front, light at the outside of the waveguide bend must travel faster than light near the inside bend; radiation loss will occur if the phase front cannot be maintained. Reproduced with permission from Hunsperger [32].

bend [266]. As shown in Figure C.2, if a guided mode with a propagation constant β_z, in a medium where the maximum phase velocity (for unguided light) is β_0, is travelling around a bend R, above a radius $(R + X_r)$ the light would have to travel faster than is possible in the medium (i.e., above β_0) in order to remain guided [32].

This result is critical in PIC design, because the propagation constant β is proportional to the effective refractive index of the waveguide. A higher β value allows a smaller bending radius to be used without introducing appreciable losses; this can have a considerable impact on device footprint and integration density. Therefore, it is desirable to maximise β; this can be achieved by using a high refractive-index waveguide material, or by increasing the waveguide cross-sectional area to maximise confinement.

References

[1] M. Harrower and C. A. Brewer, "ColorBrewer.org: An Online Tool for Selecting Colour Schemes for Maps," *The Cartographic Journal*, vol. 40, pp. 27–37, June 2003.

[2] M. Corato-Zanarella, A. Gil-Molina, X. Ji, M. C. Shin, A. Mohanty, and M. Lipson, "Widely tunable and narrow-linewidth chip-scale lasers from near-ultraviolet to near-infrared wavelengths," *Nature Photonics*, vol. 17, pp. 157–164, Feb. 2023.

[3] F. P. Payne and J. P. R. Lacey, "A theoretical analysis of scattering loss from planar optical waveguides," *Optical and Quantum Electronics*, vol. 26, pp. 977–986, Oct. 1994.

[4] G. N. West, W. Loh, D. Kharas, C. Sorace-Agaskar, K. K. Mehta, J. Sage, J. Chiaverini, and R. J. Ram, "Low-loss integrated photonics for the blue and ultraviolet regime," *APL Photonics*, vol. 4, p. 026101, Feb. 2019.

[5] H. Sobol, "Microwave Communications – An Historical Perspective," *IEEE Transactions on Microwave Theory and Techniques*, vol. 32, pp. 1170–1181, Sept. 1984.

[6] T. Maiman, "Stimulated Optical Radiation in Ruby," *Nature*, vol. 187, pp. 493–494, Aug. 1960.

[7] R. N. Hall, G. E. Fenner, J. D. Kingsley, T. J. Soltys, and R. O. Carlson, "Coherent Light Emission From GaAs Junctions," *Physical Review Letters*, vol. 9, pp. 366–368, Nov. 1962.

[8] S. E. Miller, "Integrated optics: An introduction," *The Bell System Technical Journal*, vol. 48, pp. 2059–2069, Sept. 1969.

[9] J. E. Goell and R. D. Standley, "B.S.T.J. briefs: Sputtered glass waveguide for integrated optical circuits," *The Bell System Technical Journal*, vol. 48, pp. 3445–3448, Dec. 1969.

[10] F. Kish, V. Lal, P. Evans, S. W. Corzine, M. Ziari, T. Butrie, M. Reffle, H. Tsai, A. Dentai, J. Pleumeekers, M. Missey, M. Fisher, S. Murthy, R. Salvatore, P. Samra,

S. Demars, N. Kim, A. James, A. Hosseini, P. Studenkov, M. Lauermann, R. Going, M. Lu, J. Zhang, J. Tang, J. Bostak, T. Vallaitis, M. Kuntz, D. Pavinski, A. Karanicolas, B. Behnia, D. Engel, O. Khayam, N. Modi, M. R. Chitgarha, P. Mertz, W. Ko, R. Maher, J. Osenbach, J. T. Rahn, H. Sun, K. Wu, M. Mitchell, and D. Welch, "System-on-Chip Photonic Integrated Circuits," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 24, pp. 1–20, Jan. 2018.

[11] Y. Su, Y. He, X. Guo, W. Xie, X. Ji, H. Wang, X. Cai, L. Tong, and S. Yu, "Scalability of Large-Scale Photonic Integrated Circuits," *ACS Photonics*, vol. 10, pp. 2020–2030, July 2023.

[12] S. M. Pershin, A. F. Bunkin, V. A. Zavozin, M. Y. Grishin, V. S. Makarov, P. A. Titovets, and M. O. Fedyuk, "Underwater Lidar: Remote Sensing in Strongly Scattering Media," *Physics of Wave Phenomena*, vol. 31, pp. 406–411, Dec. 2023.

[13] C. Lin, J. S. D. Peñaranda, J. Dendooven, C. Detavernier, D. Schaubroeck, N. Boon, R. Baets, and N. Le Thomas, "UV photonic integrated circuits for far-field structured illumination autofluorescence microscopy," *Nature Communications*, vol. 13, p. 4360, July 2022.

[14] J.-C. Tinguely, O. I. Helle, and B. S. Ahluwalia, "Silicon nitride waveguide platform for fluorescence microscopy of living cells," *Optics Express*, vol. 25, pp. 27678–27690, Oct. 2017.

[15] G. Moody, V. J. Sorger, D. J. Blumenthal, P. W. Juodawlkis, W. Loh, C. Sorace-Agaskar, A. E. Jones, K. C. Balram, J. C. F. Matthews, A. Laing, M. Davanco, L. Chang, J. E. Bowers, N. Quack, C. Galland, I. Aharonovich, M. A. Wolff, C. Schuck, N. Sinclair, M. Lončar, T. Komljenovic, D. Weld, S. Mookherjea, S. Buckley, M. Radulaski, S. Reitzenstein, B. Pingault, B. Machielse, D. Mukhopadhyay, A. Akimov, A. Zheltikov, G. S. Agarwal, K. Srinivasan, J. Lu, H. X. Tang, W. Jiang, T. P. McKenna, A. H. Safavi-Naeini, S. Steinhauer, A. W. Elshaari, V. Zwiller, P. S. Davids, N. Martinez, M. Gehl, J. Chiaverini, K. K. Mehta, J. Romero, N. B. Lingaraju, A. M. Weiner, D. Peace, R. Cernansky, M. Lobino, E. Diamanti, L. T. Vidarte, and R. M. Camacho, "2022 Roadmap on integrated quantum photonics," *Journal of Physics: Photonics*, vol. 4, p. 012501, Jan. 2022.

[16] Bureau International des Poids et Mesures, "The International System of Units," Dec. 2022.

[17] A. Bauch, "Caesium atomic clocks: function, performance and applications," *Measurement Science and Technology*, vol. 14, pp. 1159–1173, July 2003.

[18] K. Beloy, M. I. Bodine, T. Bothwell, S. M. Brewer, S. L. Bromley, J.-S. Chen, J.-D. Deschênes, S. A. Diddams, R. J. Fasano, T. M. Fortier, Y. S. Hassan, D. B. Hume, D. Kedar, C. J. Kennedy, I. Khader, A. Koepke, D. R. Leibrandt, H. Leopardi, A. D. Ludlow, W. F. McGrew, W. R. Milner, N. R. Newbury, D. Nicolodi, E. Oelker, T. E. Parker, J. M. Robinson, S. Romisch, S. A. Schäffer, J. A. Sherman, L. C. Sinclair, L. Sonderhouse, W. C. Swann, J. Yao, J. Ye, X. Zhang, and Boulder Atomic Clock Optical Network (BACON) Collaboration, "Frequency ratio measurements at 18-digit accuracy using an optical clock network," *Nature*, vol. 591, pp. 564–569, Mar. 2021.

[19] S. B. Papp, K. Beha, P. Del'Haye, F. Quinlan, H. Lee, K. J. Vahala, and S. A. Diddams, "Microresonator frequency comb optical clock," *Optica*, vol. 1, pp. 10–14, July 2014.

[20] P. M. Moser, "Spectral Transmission of Light through Seawater," Technical report, Defense Technical Information Center, 1992.

[21] A. P. Demchenko, *Ultraviolet Spectroscopy of Proteins*. Springer Science & Business Media, Nov. 2013.

[22] R. Antoine and P. Dugourd, "UV–Visible Absorption Spectroscopy of Protein Ions," in *Photophysics of Ionic Biochromophores* (S. Brondsted Nielsen and J. A. Wyer, eds.), Physical Chemistry in Action, pp. 141–153, Berlin, Heidelberg: Springer, 2013.

[23] M. de Goede, M. Dijkstra, R. Obregón, J. Ramón-Azcón, E. Martínez, L. Padilla, F. Mitjans, and S. M. García-Blanco, "Al_2O_3 microring resonators for the detection of a cancer biomarker in undiluted urine," *Optics Express*, vol. 27, pp. 18508–18521, June 2019.

[24] R. S. Wiederkehr and S. B. Mendes, "Extension of the broadband single-mode integrated optical waveguide technique to the ultraviolet spectral region and its applications," *Analyst*, vol. 139, no. 6, pp. 1396–1402, 2014.

[25] J. Smith, J. Monroy-Ruz, J. G. Rarity, and K. C. Balram, "Single photon emission and single spin coherence of a nitrogen vacancy center encapsulated in silicon nitride," *Applied Physics Letters*, vol. 116, p. 134001, Mar. 2020.

[26] G. Reed and A. Knights, *Silicon Photonics: An Introduction*. Wiley, 2004.

[27] K. Y. Yang, D. Y. Oh, S. H. Lee, Q.-F. Yang, X. Yi, B. Shen, H. Wang, and K. Vahala, "Bridging ultrahigh-Q devices and photonic circuits," *Nature Photonics*, vol. 12, pp. 297–302, May 2018.

[28] I. Vurgaftman, M. P. Lumb, and J. R. Meyer, *Bands and Photons in III-V Semiconductor Quantum Structures*. Oxford University Press, Dec. 2020.

[29] F. Urbach, "The Long-Wavelength Edge of Photographic Sensitivity and of the Electronic Absorption of Solids," *Physical Review*, vol. 92, pp. 1324–1324, Dec. 1953.

[30] M. Corato-Zanarella, X. Ji, A. Mohanty, and M. Lipson, "Absorption and scattering limits of silicon nitride integrated photonics in the visible spectrum," *Optics Express*, vol. 32, pp. 5718–5728, Feb. 2024.

[31] K. V. Shportko, "Disorder and compositional dependences in Urbach-Martienssen tails in amorphous $(GeTe)_x(Sb_2Te_3)_{1-x}$ alloys," *Scientific Reports*, vol. 9, p. 6030, Apr. 2019.

[32] R. G. Hunsperger, *Losses in Optical Waveguides*. New York, NY: Springer, 2009.

[33] V. Plotnichenko, G. Ivanov, E. Kryukova, V. Aksenov, V. Sokolov, and V. Isaev, "Influence of molecular hydrogen diffusion on concentration and distribution of hydroxyl groups in silica fibers," *Journal of Lightwave Technology*, vol. 23, pp. 341–347, Jan. 2005.

[34] H. Moosmüller, R. K. Chakrabarty, K. M. Ehlers, and W. P. Arnott, "Absorption Ångström coefficient, brown carbon, and aerosols: basic concepts, bulk matter, and spherical particles," *Atmospheric Chemistry and Physics*, vol. 11, pp. 1217–1225, Feb. 2011.

[35] R. W. Boyd, *Nonlinear Optics*. Burlington: Academic Press, third ed., Jan. 2008.

[36] A. G. Rojo and P. R. Berman, "Rayleigh scattering revisited: From gases to crystals," *American Journal of Physics*, vol. 78, pp. 94–101, Jan. 2010.

[37] S. Gelin, H. Tanaka, and A. Lemaître, "Anomalous phonon scattering and elastic correlations in amorphous solids," *Nature Materials*, vol. 15, pp. 1177–1181, Nov. 2016.

[38] F. Benabid, M. Notcutt, L. Ju, and D. G. Blair, "Rayleigh scattering in sapphire test mass for laser interferometric gravitational-wave detectors: I. Measurement of scattering attenuation coefficient," *Optics Communications*, vol. 167, pp. 7–13, Aug. 1999.

[39] C. F. Bohren and D. R. Huffman, *Absorption and Scattering of Light by Small Particles*. John Wiley & Sons, Sept. 2008.

[40] S. Roberts, X. Ji, J. Cardenas, M. Corato-Zanarella, and M. Lipson, "Measurements and Modeling of Atomic-Scale Sidewall Roughness and Losses in Integrated Photonic Devices," *Advanced Optical Materials*, p. 2102073, June 2022.

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

[42] International Telecommunication Union, "Recommendation G.652: Characteristics of a single-mode optical fibre and cable," Nov. 2016.

[43] Z. L. Newman, V. Maurice, T. Drake, J. R. Stone, T. C. Briles, D. T. Spencer, C. Fredrick, Q. Li, D. Westly, B. R. Ilic, B. Shen, M.-G. Suh, K. Y. Yang, C. Johnson, D. M. S. Johnson, L. Hollberg, K. J. Vahala, K. Srinivasan, S. A. Diddams, J. Kitching, S. B. Papp, and M. T. Hummon, "Architecture for the photonic integration of an optical atomic clock," *Optica*, vol. 6, pp. 680–685, May 2019.

[44] G. New, *Introduction to Nonlinear Optics*. Cambridge University Press, Apr. 2011.

[45] A. Melloni, R. Costa, G. Cusmai, and F. Morichetti, "The role of index contrast in dielectric optical waveguides," *International Journal of Materials and Product Technology*, vol. 34, pp. 421–437, Jan. 2009.

[46] R. Marchetti, C. Lacava, A. Khokhar, X. Chen, I. Cristiani, D. J. Richardson, G. T. Reed, P. Petropoulos, and P. Minzioni, "High-efficiency grating-couplers: demonstration of a new design strategy," *Scientific Reports*, vol. 7, pp. 1–8, Nov. 2017.

[47] X. Chen, C. Li, C. K. Y. Fung, S. M. G. Lo, and H. K. Tsang, "Apodized Waveguide Grating Couplers for Efficient Coupling to Optical Fibers," *IEEE Photonics Technology Letters*, vol. 22, pp. 1156–1158, Aug. 2010.

[48] S. Nambiar, A. Kumar, R. Kallega, P. Ranganath, and S. K. Selvaraja, "High-Efficiency Grating Coupler in 400 nm and 500 nm PECVD Silicon Nitride With Bottom Reflector," *IEEE Photonics Journal*, vol. 11, pp. 1–13, Oct. 2019.

[49] L. Tolbert, B. Ozpineci, S. Islam, and M. Chinthavali, "Wide bandgap semiconductors for utility applications," *Proceedings of the IASTED Multi-Conference — Power and Energy Systems*, vol. 7, Jan. 2003.

[50] H. K. Tsang and Y. Liu, "Nonlinear optical properties of silicon waveguides," *Semiconductor Science and Technology*, vol. 23, p. 064007, May 2008.

[51] D. E. Aspnes and A. A. Studna, "Dielectric functions and optical parameters of Si, Ge, GaP, GaAs, GaSb, InP, InAs, and InSb from 1.5 to 6.0 eV," *Physical Review B*, vol. 27, pp. 985–1009, Jan. 1983.

[52] S. Kasap, *Springer Handbook of Electronic and Photonic Materials*. Springer Science & Business Media, 2006.

[53] H. Chen, X. Huang, H. Fu, Z. Lu, X. Zhang, J. A. Montes, and Y. Zhao, "Characterizations of nonlinear optical properties on GaN crystals in polar, nonpolar, and semipolar orientations," *Applied Physics Letters*, vol. 110, p. 181110, May 2017.

[54] S. Adachi, *Optical Constants of Crystalline and Amorphous Semiconductors*. Boston, MA: Springer US, 1999.

[55] A. A. Fontenla, I. Hegeman, W. A. P. M. Hendriks, M. Dijkstra, and S. M. García-Blanco, "Low-loss reactive sputter deposited titanium oxide waveguides," in *2021 IEEE 17th International Conference on Group IV Photonics (GFP)*, pp. 1–2, Dec. 2021.

[56] S. A. Campbell, H.-S. Kim, D. C. Gilmer, B. He, T. Ma, and W. L. Gladfelter, "Titanium dioxide (TiO_2)-based gate insulators," *IBM Journal of Research and Development*, vol. 43, pp. 383–392, May 1999.

[57] R. Adair, L. L. Chase, and S. A. Payne, "Nonlinear refractive index of optical crystals," *Physical Review B*, vol. 39, pp. 3337–3350, Feb. 1989.

[58] I. Hegeman, M. Dijkstra, and S. M. García-Blanco, "Low loss TiO_2 channel waveguides," in *Integrated Optics: Devices, Materials, and Technologies XXIV*, vol. 11283, pp. 23–32, SPIE, Feb. 2020.

[59] K. F. Lamee, D. R. Carlson, Z. L. Newman, S.-P. Yu, and S. B. Papp, "Nanophotonic tantala waveguides for supercontinuum generation pumped at 1560nm," *Optics Letters*, vol. 45, pp. 4192–4195, Aug. 2020.

[60] M. Belt, M. L. Davenport, J. E. Bowers, and D. J. Blumenthal, "Ultra-low-loss Ta_2O_5-core/SiO_2-clad planar waveguides on Si substrates," *Optica*, vol. 4, pp. 532–536, May 2017.

[61] H. Jung, S.-P. Yu, D. R. Carlson, T. E. Drake, T. C. Briles, and S. B. Papp, "Tantala Kerr nonlinear integrated photonics," *Optica*, vol. 8, pp. 811–817, June 2021.

[62] S. M. Taleb, M. A. Fakhri, and S. A. Adnan, "Optical Investigations of Nanophotonic $LiNbO_3$ Films Deposited by Pulsed Laser Deposition Method," in *Defect and Diffusion Forum*, vol. 398, pp. 16–22, Jan. 2020.

[63] D. E. Zelmon, D. L. Small, and D. Jundt, "Infrared corrected Sellmeier coefficients for congruently grown lithium niobate and 5 mol. % magnesium oxide–doped lithium niobate," *JOSA B*, vol. 14, pp. 3319–3322, Dec. 1997.

[64] Y.-h. Wang, X.-x. Yu, F. Liu, and Y.-m. Wang, "Nonlinear refraction of lithium niobate crystal doped with different metal nanoparticles," *Materials Letters*, vol. 123, pp. 35–37, May 2014.

[65] B. Desiatov, A. Shams-Ansari, M. Zhang, C. Wang, and M. Lončar, "Ultra-low-loss integrated visible photonics using thin-film lithium niobate," *Optica*, vol. 6, pp. 380–384, Mar. 2019.

[66] J. A. Smith, H. Francis, G. Navickaite, and M. J. Strain, "SiN foundry platform for high performance visible light integrated photonics," *Optical Materials Express*, vol. 13, pp. 458–468, Feb. 2023.

[67] M. Soltani, R. Soref, T. Palacios, and D. Englund, "AlGaN/AlN integrated photonics platform for the ultraviolet and visible spectral range," *Optics Express*, vol. 24, pp. 25415–25423, Oct. 2016.

[68] 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, titanium dioxide, silicon dioxide, aluminum nitride, and silicon nitride," *Applied Optics*, vol. 51, pp. 6789–6798, Oct. 2012.

[69] K. Ikeda, R. E. Saperstein, N. Alic, and Y. Fainman, "Thermal and Kerr nonlinear properties of plasma-deposited silicon nitride/silicon dioxide waveguides," *Optics Express*, vol. 16, pp. 12987–12994, Aug. 2008.

[70] H. R. Phillip and E. A. Taft, "Kramers-Kronig Analysis of Reflectance Data for Diamond," *Physical Review*, vol. 136, pp. A1445–A1448, Nov. 1964.

[71] B. Zhang, S. Liu, X. Wu, T. Yi, Y. Fang, J. Zhang, Q. Zhong, X. Peng, X. Liu, and Y. Song, "Ultrafast dynamics of carriers and nonlinear refractive index in bulk polycrystalline diamond," *Optik*, vol. 130, pp. 1073–1079, Feb. 2017.

[72] M. Zhao, C.-H. Xu, W.-J. Hu, W.-J. Wang, L.-W. Guo, and X.-L. Chen, "Observation of Two-Photon Absorption and Nonlinear Refraction in AlN," *Chinese Physics Letters*, vol. 33, p. 104201, Oct. 2016.

[73] S. Mardani, W. a. P. M. Hendriks, M. Dijkstra, and S. M. García-Blanco, "Low-loss sputter coated aluminum nitride waveguides," in *Integrated Optics: Devices, Materials, and Technologies XXVII*, vol. 12424, pp. 5–9, SPIE, Mar. 2023.

[74] M. M. Aslan, N. A. Webster, C. L. Byard, M. B. Pereira, C. M. Hayes, R. S. Wiederkehr, and S. B. Mendes, "Low-loss optical waveguides for the near ultra-violet and visible

spectral regions with Al_2O_3 thin films from atomic layer deposition," *Thin Solid Films*, vol. 518, pp. 4935–4940, June 2010.

[75] B. J. H. Stadler, M. Oliveria, and L. O. Bouthillette, "Alumina Thin Films as Optical Waveguides," *Journal of the American Ceramic Society*, vol. 78, pp. 3336–3344, Dec. 1995.

[76] E. McKay, N. G. Pruiti, S. May, and M. Sorel, "High-confinement alumina waveguides with sub-dB/cm propagation losses at 450 nm," *Scientific Reports*, vol. 13, p. 19917, Nov. 2023.

[77] P. Dragic, M. Cavillon, and J. Ballato, "The linear and nonlinear refractive index of amorphous Al_2O_3 deduced from aluminosilicate optical fibers," *International Journal of Applied Glass Science*, vol. 9, no. 3, pp. 421–427, 2018.

[78] A. Major, F. Yoshino, I. Nikolakakos, J. S. Aitchison, and P. W. E. Smith, "Dispersion of the nonlinear refractive index in sapphire," *Optics Letters*, vol. 29, pp. 602–604, Mar. 2004.

[79] A. Tsaturyan, E. Kachan, R. Stoian, and J.-P. Colombier, "Ultrafast bandgap narrowing and cohesion loss of photoexcited fused silica," *The Journal of Chemical Physics*, vol. 156, p. 224301, June 2022.

[80] Y. Zhang, J. Wu, L. Jia, Y. Qu, Y. Yang, B. Jia, and D. J. Moss, "Graphene Oxide for Nonlinear Integrated Photonics," *Laser & Photonics Reviews*, vol. 17, no. 3, p. 2200512, 2023.

[81] D. Duchesne, M. Ferrera, L. Razzari, R. Morandotti, B. E. Little, S. T. Chu, and D. J. Moss, "Efficient self-phase modulation in low loss, high index doped silica glass integrated waveguides," *Optics Express*, vol. 17, pp. 1865–1870, Feb. 2009.

[82] K. Wörhoff, R. G. Heideman, A. Leinse, and M. Hoekman, "TriPleX: a versatile dielectric photonic platform," *Advanced Optical Technologies*, vol. 4, pp. 189–207, Apr. 2015.

[83] X. Ji, J. K. Jang, U. D. Dave, M. Corato-Zanarella, C. Joshi, A. L. Gaeta, and M. Lipson, "Exploiting Ultralow Loss Multimode Waveguides for Broadband Frequency Combs," *Laser & Photonics Reviews*, vol. 15, no. 1, p. 2000353, 2021.

[84] M. H. P. Pfeiffer, A. Kordts, V. Brasch, M. Zervas, M. Geiselmann, J. D. Jost, and T. J. Kippenberg, "Photonic Damascene process for integrated high-Q microresonator based nonlinear photonics," *Optica*, vol. 3, pp. 20–25, Jan. 2016.

[85] K. J. A. Ooi, D. K. T. Ng, T. Wang, A. K. L. Chee, S. K. Ng, Q. Wang, L. K. Ang, A. M. Agarwal, L. C. Kimerling, and D. T. H. Tan, "Pushing the limits of CMOS optical parametric amplifiers with USRN:Si_7N_3 above the two-photon absorption edge," *Nature Communications*, vol. 8, p. 13878, Jan. 2017.

[86] C. Sorace-Agaskar, D. Kharas, S. Yegnanarayanan, R. T. Maxson, G. N. West, W. Loh, S. Bramhavar, R. J. Ram, J. Chiaverini, J. Sage, and P. Juodawlkis, "Versatile Silicon Nitride and Alumina Integrated Photonic Platforms for the Ultraviolet to Short-Wave Infrared," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 25, pp. 1–15, Sept. 2019.

[87] T. J. Morin, L. Chang, W. Jin, C. Li, J. Guo, H. Park, M. A. Tran, T. Komljenovic, and J. E. Bowers, "CMOS-foundry-based blue and violet photonics," *Optica*, vol. 8, pp. 755–756, May 2021.

[88] A. Boes, L. Chang, C. Langrock, M. Yu, M. Zhang, Q. Lin, M. Lončar, M. Fejer, J. Bowers, and A. Mitchell, "Lithium niobate photonics: Unlocking the electromagnetic spectrum," *Science*, vol. 379, p. eabj4396, Jan. 2023.

[89] T.-H. Wu, L. Ledezma, C. Fredrick, P. Sekhar, R. Sekine, Q. Guo, R. M. Briggs, A. Marandi, and S. A. Diddams, "Visible-to-ultraviolet frequency comb generation in lithium niobate nanophotonic waveguides," *Nature Photonics*, pp. 1–6, Jan. 2024.

[90] Y. Qi and Y. Li, "Integrated lithium niobate photonics," *Nanophotonics*, vol. 9, pp. 1287–1320, June 2020.

[91] M. Churaev, R. N. Wang, A. Riedhauser, V. Snigirev, T. Blésin, C. Möhl, M. H. Anderson, A. Siddharth, Y. Popoff, U. Drechsler, D. Caimi, S. Hönl, J. Riemensberger, J. Liu, P. Seidler, and T. J. Kippenberg, "A heterogeneously integrated lithium niobate-on-silicon nitride photonic platform," *Nature Communications*, vol. 14, p. 3499, June 2023.

[92] C. Huang, H. Zhang, and H. Sun, "Ultraviolet optoelectronic devices based on AlGaN-SiC platform: Towards monolithic photonics integration system," *Nano Energy*, vol. 77, p. 105149, Nov. 2020.

[93] D. Li, K. Jiang, X. Sun, and C. Guo, "AlGaN photonics: recent advances in materials and ultraviolet devices," *Advances in Optics and Photonics*, vol. 10, pp. 43–110, Mar. 2018.

[94] M. Stegmaier, J. Ebert, J. M. Meckbach, K. Ilin, M. Siegel, and W. H. P. Pernice, "Aluminum nitride nanophotonic circuits operating at ultraviolet wavelengths," *Applied Physics Letters*, vol. 104, p. 091108, Mar. 2014.

[95] T.-J. Lu, M. Fanto, H. Choi, P. Thomas, J. Steidle, S. Mouradian, W. Kong, D. Zhu, H. Moon, K. Berggren, J. Kim, M. Soltani, S. Preble, and D. Englund, "Aluminum nitride integrated photonics platform for the ultraviolet to visible spectrum," *Optics Express*, vol. 26, pp. 11147–11160, Apr. 2018.

[96] X. Liu, A. W. Bruch, Z. Gong, J. Lu, J. B. Surya, L. Zhang, J. Wang, J. Yan, and H. X. Tang, "Ultra-high-Q UV microring resonators based on a single-crystalline AlN platform," *Optica*, vol. 5, pp. 1279–1282, Oct. 2018.

[97] D. Rani, O. R. Opaluch, and E. Neu, "Recent Advances in Single Crystal Diamond Device Fabrication for Photonics, Sensing and Nanomechanics," *Micromachines*, vol. 12, p. 36, Jan. 2021.

[98] F. Gao, J. V. Erps, Z. Huang, H. Thienpont, R. G. Beausoleil, and N. Vermeulen, "Down-scaling grating couplers and waveguides in single-crystal diamond for VIS-UV operation," *Journal of Physics: Photonics*, vol. 1, p. 015003, Dec. 2018.

[99] W. A. P. M. Hendriks, L. Chang, C. I. van Emmerik, J. Mu, M. de Goede, M. Dijkstra, and S. M. García-Blanco, "Rare-earth ion doped Al_2O_3 for active integrated photonics," *Advances in Physics: X*, vol. 6, p. 1833753, Jan. 2021.

[100] A. L. Dragoo and J. J. Diamond, "Transitions in Vapor-Deposited Alumina from 300° to 1200°C," *Journal of the American Ceramic Society*, vol. 50, no. 11, pp. 568–574, 1967.

[101] C. I. v. Emmerik, W. A. P. M. Hendriks, M. M. Stok, M. d. Goede, L. Chang, M. Dijkstra, F. Segerink, D. Post, E. G. Keim, M. J. Dikkers, and S. M. García-Blanco, "Relative oxidation state of the target as guideline for depositing optical quality RF reactive magnetron sputtered Al_2O_3 layers," *Optical Materials Express*, vol. 10, pp. 1451–1462, June 2020.

[102] M. Broas, O. Kanninen, V. Vuorinen, M. Tilli, and M. Paulasto-Kröckel, "Chemically Stable Atomic-Layer-Deposited Al_2O_3 Films for Processability," *ACS Omega*, vol. 2, pp. 3390–3398, July 2017.

[103] S. K. Lee, S. B. Lee, S. Y. Park, Y. S. Yi, and C. W. Ahn, "Structure of Amorphous Aluminum Oxide," *Physical Review Letters*, vol. 103, p. 095501, Aug. 2009.

[104] S. K. Lee, S. Y. Park, Y. S. Yi, and J. Moon, "Structure and Disorder in Amorphous Alumina Thin Films: Insights from High-Resolution Solid-State NMR," *The Journal of Physical Chemistry C*, vol. 114, pp. 13890–13894, Aug. 2010.

[105] M. T. Aguilar-Gama, E. Ramírez-Morales, Z. Montiel-González, A. Mendoza-Galván, M. Sotelo-Lerma, P. K. Nair, and H. Hu, "Structure and refractive index of thin alumina films grown by atomic layer deposition," *Journal of Materials Science: Materials in Electronics*, vol. 26, pp. 5546–5552, Aug. 2015.

[106] I. Levin and D. Brandon, "Metastable Alumina Polymorphs: Crystal Structures and Transition Sequences," *Journal of the American Ceramic Society*, vol. 81, no. 8, pp. 1995–2012, 1998.

[107] H. Momida, T. Hamada, Y. Takagi, T. Yamamoto, T. Uda, and T. Ohno, "Theoretical study on dielectric response of amorphous alumina," *Physical Review B*, vol. 73, p. 054108, Feb. 2006.

[108] E. O. Filatova and A. S. Konashuk, "Interpretation of the Changing the Band Gap of Al_2O_3 Depending on Its Crystalline Form: Connection with Different Local Symmetries," *The Journal of Physical Chemistry C*, vol. 119, pp. 20755–20761, Sept. 2015.

[109] T. V. Perevalov, O. E. Tereshenko, V. A. Gritsenko, V. A. Pustovarov, A. P. Yelisseyev, C. Park, J. H. Han, and C. Lee, "Oxygen deficiency defects in amorphous Al_2O_3," *Journal of Applied Physics*, vol. 108, p. 013501, July 2010.

[110] K. J. Caulfield, R. Cooper, and J. F. Boas, "Luminescence from electron-irradiated sapphire," *Physical Review B*, vol. 47, pp. 55–61, Jan. 1993.

[111] J. A. Guerra, A. Tejada, J. A. Töfflinger, R. Grieseler, and L. Korte, "Band-fluctuations model for the fundamental absorption of crystalline and amorphous semiconductors: a dimensionless joint density of states analysis," *Journal of Physics D: Applied Physics*, vol. 52, p. 105303, Jan. 2019.

[112] B.-U. Chen, C. L. Tang, and J. M. Telle, "CW harmonic generation in the uv using a thin-film waveguide on a nonlinear substrate," *Applied Physics Letters*, vol. 25, pp. 495–498, Nov. 1974.

[113] G. Este and W. D. Westwood, "Reactive deposition of low loss Al_2O_3 optical waveguides by modified DC planar magnetron sputtering," *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, vol. 2, pp. 1238–1247, July 1984.

[114] M. K. Smit, G. A. Acket, and C. J. van der Laan, "Al_2O_3 films for integrated optics," *Thin Solid Films*, vol. 138, pp. 171–181, Apr. 1986.

[115] S. M. Arnold and B. E. Cole, "Ion beam sputter deposition of low loss Al_2O_3 films for integrated optics," *Thin Solid Films*, vol. 165, pp. 1–9, Nov. 1988.

[116] D. J. Blumenthal, "Photonic integration for UV to IR applications," *APL Photonics*, vol. 5, p. 020903, Feb. 2020.

[117] W. Hendriks, M. Dijkstra, S. Mardani, I. Hegeman, and S. García-Blanco, "Low-loss photonic integrated circuits for UV applications," in *Integrated Optics: Devices, Materials, and Technologies XXVII*, p. 124242, Jan. 2023.

[118] W. A. P. M. Hendriks, M. Dijkstra, I. Hegeman, F. B. Segerink, E. G. Keim, M. J. Dikkers, and S. M. García-Blanco, "Poly-crystalline low-loss aluminium oxide waveguides," in *Integrated Optics: Devices, Materials, and Technologies XXV*, vol. 11689, pp. 110–117, SPIE, Mar. 2021.

[119] V. Miikkulainen, M. Leskelä, M. Ritala, and R. L. Puurunen, "Crystallinity of inorganic films grown by atomic layer deposition: Overview and general trends," *Journal of Applied Physics*, vol. 113, p. 021301, Jan. 2013.

[120] D. L. Caballero-Espitia, E. G. Lizarraga-Medina, H. A. Borbon-Nuñez, O. E. Contreras-Lopez, H. Tiznado, and H. Marquez, "Study of Al_2O_3 thin films by ALD using H_2O and O_3 as oxygen source for waveguide applications," *Optical Materials*, vol. 109, p. 110370, Nov. 2020.

[121] O. M. E. Ylivaara, X. Liu, L. Kilpi, J. Lyytinen, D. Schneider, M. Laitinen, J. Julin, S. Ali, S. Sintonen, M. Berdova, E. Haimi, T. Sajavaara, H. Ronkainen, H. Lipsanen, J. Koskinen, S.-P. Hannula, and R. L. Puurunen, "Aluminum oxide from trimethylaluminum and water by atomic layer deposition: The temperature dependence of residual stress, elastic modulus, hardness and adhesion," *Thin Solid Films*, vol. 552, pp. 124–135, Feb. 2014.

[122] J. Bradley, F. Ay, K. Wörhoff, and M. Pollnau, "Fabrication of low-loss channel waveguides in Al_2O_3 and Y_2O_3 layers by inductively coupled plasma reactive ion etching," *Applied Physics B*, vol. 89, pp. 311–318, Nov. 2007.

[123] J. J. D. León, D. M. Fryauf, J. E. Volk, and N. P. Kobayashi, "Aluminum Oxide Waveguides for Improved Transmission in the Visible and Near-Infrared Spectrum," *IEEE Photonics Technology Letters*, vol. 31, pp. 43–45, Jan. 2019.

[124] P. Kumar, M. K. Wiedmann, C. H. Winter, and I. Avrutsky, "Optical properties of Al_2O_3 thin films grown by atomic layer deposition," *Applied Optics*, vol. 48, pp. 5407–5412, Oct. 2009.

[125] S. M. Martinussen, *Grow thick, etch deep: novel waveguide geometries and material systems for integrated photonics*. PhD, University of Twente, Enschede, The Netherlands, Apr. 2021.

[126] J. Yang, K. van Dalfsen, K. Wörhoff, F. Ay, and M. Pollnau, "High-gain Al_2O_2:Nd^{3+} channel waveguide amplifiers at 880nm, 1060nm, and 1330nm," *Applied Physics B*, vol. 101, pp. 119–127, Oct. 2010.

[127] K. Wörhoff, J. D. B. Bradley, F. Ay, D. Geskus, T. P. Blauwendraat, and M. Pollnau, "Reliable Low-Cost Fabrication of Low-Loss Al_2O_3 Waveguides With 5.4-dB Optical Gain," *IEEE Journal of Quantum Electronics*, vol. 45, pp. 454–461, May 2009.

[128] S. A. Vázquez-Córdova, M. Dijkstra, E. H. Bernhardi, F. Ay, K. Wörhoff, J. L. Herek, S. M. García-Blanco, and M. Pollnau, "Erbium-doped spiral amplifiers with 20 dB of net gain on silicon," *Optics Express*, vol. 22, pp. 25993–26004, Oct. 2014.

[129] J. Notaros, N. Li, C. V. Poulton, Z. Su, M. J. Byrd, E. S. Magden, E. Timurdogan, C. Baiocco, N. M. Fahrenkopf, and M. R. Watts, "CMOS-Compatible Optical Phased Array Powered by a Monolithically-Integrated Erbium Laser," *Journal of Lightwave Technology*, vol. 37, pp. 5982–5987, Dec. 2019.

[130] L. Agazzi, J. D. B. Bradley, M. Dijkstra, F. Ay, G. Roelkens, R. Baets, K. Wörhoff, and M. Pollnau, "Monolithic integration of erbium-doped amplifiers with silicon-on-insulator waveguides," *Optics Express*, vol. 18, pp. 27703–27711, Dec. 2010.

[131] R. Morgan, D. Kharas, J. Knecht, P. Juodawlkis, K. Cahoy, and C. Sorace-Agaskar, "Waveguide-Integrated Blue Light Detector," in *2021 IEEE Photonics Conference (IPC)*, pp. 1–2, Oct. 2021.

[132] C. Lin, J. S. D. Peñaranda, J. Dendooven, C. Detavernier, D. Schaubroeck, R. Baets, and N. L. Thomas, "Photonic integrated alumina waveguide gratings for far-field structured illumination at UV wavelengths," in *Conference on Lasers and Electro-Optics (2021)*, p. STu1A.6, May 2021.

[133] J. M. Dudley and J. R. Taylor, *Supercontinuum Generation in Optical Fibers*. Cambridge University Press, Apr. 2010.

[134] J. M. C. Boggio, D. Bodenmüller, T. Fremberg, R. Haynes, M. M. Roth, R. Eisermann, M. Lisker, L. Zimmermann, and M. Böhm, "Dispersion engineered silicon nitride waveguides by geometrical and refractive-index optimization," *JOSA B*, vol. 31, pp. 2846–2857, Nov. 2014.

[135] J. M. Dudley, G. Genty, and S. Coen, "Supercontinuum generation in photonic crystal fiber," *Reviews of Modern Physics*, vol. 78, pp. 1135–1184, Oct. 2006.

[136] H. Liang, Y. He, R. Luo, and Q. Lin, "Ultra-broadband dispersion engineering of nanophotonic waveguides," *Optics Express*, vol. 24, pp. 29444–29451, Dec. 2016.

[137] K. Hammani, L. Markey, M. Lamy, B. Kibler, J. Arocas, J. Fatome, A. Dereux, J.-C. Weeber, and C. Finot, "Octave Spanning Supercontinuum in Titanium Dioxide Waveguides," *Applied Sciences*, vol. 8, p. 543, Apr. 2018.

[138] T. D. Bucio, C. Lacava, M. Clementi, J. Faneca, I. Skandalos, A. Baldycheva, M. Galli, K. Debnath, P. Petropoulos, and F. Gardes, "Silicon Nitride Photonics for the Near-Infrared," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 26, pp. 1–13, Mar. 2020.

[139] R. Kou, A. Ishizawa, K. Yoshida, N. Yamamoto, X. Xu, Y. Kikkawa, K. Kawashima, T. Aihara, T. Tsuchizawa, G. Cong, K. Hitachi, T. Nishikawa, K. Oguri, and K. Yamada, "Spatially resolved multimode excitation for smooth supercontinuum generation in a SiN waveguide," *Optics Express*, vol. 31, pp. 6088–6098, Feb. 2023.

[140] J. P. Epping, T. Hellwig, M. Hoekman, R. Mateman, A. Leinse, R. G. Heideman, A. v. Rees, P. J. M. v. d. Slot, C. J. Lee, C. Fallnich, and K.-J. Boller, "On-chip visible-to-infrared supercontinuum generation with more than 495 THz spectral bandwidth," *Optics Express*, vol. 23, pp. 19596–19604, July 2015.

[141] R. Fan, C.-L. Wu, Y.-Y. Lin, C.-Y. Liu, P.-S. Hwang, C.-W. Liu, J. Qiao, M.-H. Shih, Y.-J. Hung, Y.-J. Chiu, A.-K. Chu, and C.-K. Lee, "Visible to near-infrared octave spanning supercontinuum generation in tantalum pentoxide (Ta_2O_5) air-cladding waveguide," *Optics Letters*, vol. 44, pp. 1512–1515, Mar. 2019.

[142] J. R. C. Woods, J. Daykin, A. S. K. Tong, C. Lacava, P. Petropoulos, A. C. Tropper, P. Horak, J. S. Wilkinson, and V. Apostolopoulos, "Supercontinuum generation in tantalum pentoxide waveguides for pump wavelengths in the 900nm to 1500nm spectral region," *Optics Express*, vol. 28, pp. 32173–32184, Oct. 2020.

[143] X. Liu, A. W. Bruch, J. Lu, Z. Gong, J. B. Surya, L. Zhang, J. Wang, J. Yan, and H. X. Tang, "Beyond 100 THz-spanning ultraviolet frequency combs in a non-centrosymmetric crystalline waveguide," *Nature Communications*, vol. 10, p. 2971, July 2019.

[144] H. Chen, J. Zhou, D. Li, D. Chen, A. K. Vinod, H. Fu, X. Huang, T.-H. Yang, J. A. Montes, K. Fu, C. Yang, C.-Z. Ning, C. W. Wong, A. M. Armani, and Y. Zhao, "Supercontinuum Generation in High Order Waveguide Mode with near-Visible Pumping Using Aluminum Nitride Waveguides," *ACS Photonics*, vol. 8, pp. 1344–1352, May 2021.

[145] D. Yoon Oh, K. Y. Yang, C. Fredrick, G. Ycas, S. A. Diddams, and K. J. Vahala, "Coherent ultra-violet to near-infrared generation in silica ridge waveguides," *Nature Communications*, vol. 8, p. 13922, Jan. 2017.

[146] H. C. Frankis, Z. Su, N. Li, E. S. Magden, M. Ye, M. R. Watts, and J. D. B. Bradley, "Four-Wave Mixing in a High-Q Aluminum Oxide Microcavity on Silicon," in *2018 Conference on Lasers and Electro-Optics (CLEO)*, pp. 1–2, May 2018.

[147] X. Li, Q. Chen, L. Sang, L. Yang, Z. Liu, and Z. Wang, "Atomic Layer Deposition Al_2O_3 Thin Films in Magnetized Radio Frequency Plasma Source," *Physics Procedia*, vol. 18, pp. 100–106, Jan. 2011.

[148] C. George, P. Littlewood, and P. C. Stair, "Understanding Pore Formation in ALD Alumina Overcoats," *ACS Applied Materials & Interfaces*, vol. 12, pp. 20331–20343, May 2020.

[149] H. C. M. Knoops, K. Arts, J. W. Buiter, L. M. Martini, R. Engeln, D. T. Hemakumara, M. Powell, W. M. M. E. Kessels, C. J. Hodson, and A. O'Mahony, "Innovative remote plasma source for atomic layer deposition for GaN devices," *Journal of Vacuum Science & Technology A*, vol. 39, p. 062403, Dec. 2021.

[150] H. Liu, L. Jensen, J. Becker, M. C. Wurz, P. Ma, and D. Ristau, "Comparison of ALD and IBS Al_2O_3 films for high power lasers," in *Laser-Induced Damage in Optical Materials 2016*, vol. 10014, p. 1001421, Dec. 2016.

[151] H. H. Sønsteby, A. Yanguas-Gil, and J. W. Elam, "Consistency and reproducibility in atomic layer deposition," *Journal of Vacuum Science & Technology A*, vol. 38, p. 020804, Feb. 2020.

[152] V. Verlaan, L. R. J. G. v. d. Elzen, G. Dingemans, M. C. M. v. d. Sanden, and W. M. M. Kessels, "Composition and bonding structure of plasma-assisted ALD Al_2O_3 films," *physica status solidi c*, vol. 7, no. 3-4, pp. 976–979, 2010.

[153] S. Shestaeva, A. Bingel, P. Munzert, L. Ghazaryan, C. Patzig, A. Tünnermann, and A. Szeghalmi, "Mechanical, structural, and optical properties of PEALD metallic oxides for optical applications," *Applied Optics*, vol. 56, pp. C47–C59, Feb. 2017.

[154] J. L. v. Hemmen, S. B. S. Heil, J. H. Klootwijk, F. Roozeboom, C. J. Hodson, M. C. M. v. d. Sanden, and W. M. M. Kessels, "Plasma and Thermal ALD of Al_2O_3 in a Commercial 200 mm ALD Reactor," *Journal of The Electrochemical Society*, vol. 154, p. G165, May 2007.

[155] J. Yim, O. M. E. Ylivaara, M. Ylilammi, V. Korpelainen, E. Haimi, E. Verkama, M. Utriainen, and R. L. Puurunen, "Saturation profile based conformality analysis for atomic layer deposition: aluminum oxide in lateral high-aspect-ratio channels," *Physical Chemistry Chemical Physics*, vol. 22, no. 40, pp. 23107–23120, 2020.

[156] M. Demirtaş, C. Odaci, N. K. Perkgöz, C. Sevik, and F. Ay, "Low Loss Atomic Layer Deposited Al_2O_3 and Y_2O_3 Waveguides for Applications in On-Chip Optical Amplifiers," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 24, pp. 1–8, July 2018.

[157] V. R. Rai, V. Vandalon, and S. Agarwal, "Influence of Surface Temperature on the Mechanism of Atomic Layer Deposition of Aluminum Oxide Using an Oxygen Plasma and Ozone," *Langmuir*, vol. 28, pp. 350–357, Jan. 2012.

[158] A. S. Sandupatla, K. Alexopoulos, M.-F. Reyniers, and G. B. Marin, "Ab Initio Investigation of Surface Chemistry of Alumina ALD on Hydroxylated γ-Alumina Surface," *The Journal of Physical Chemistry C*, vol. 119, pp. 13050–13061, June 2015.

[159] V. R. Rai, V. Vandalon, and S. Agarwal, "Surface Reaction Mechanisms during Ozone and Oxygen Plasma Assisted Atomic Layer Deposition of Aluminum Oxide," *Langmuir*, vol. 26, pp. 13732–13735, Sept. 2010.

[160] S. B. S. Heil, P. Kudlacek, E. Langereis, R. Engeln, M. C. M. van de Sanden, and W. M. M. Kessels, "In situ reaction mechanism studies of plasma-assisted atomic layer deposition of Al_2O_3," *Applied Physics Letters*, vol. 89, p. 131505, Sept. 2006.

[161] G. Dingemans, M. C. M. van de Sanden, and W. M. M. Kessels, "Influence of the Deposition Temperature on the c-Si Surface Passivation by Al_2O_3 Films Synthesized by ALD and PECVD," *Electrochemical and Solid-State Letters*, vol. 13, p. H76, Dec. 2009.

[162] T. E. Jenkins, "Multiple-angle-of-incidence ellipsometry," *Journal of Physics D: Applied Physics*, vol. 32, p. R45, May 1999.

[163] G. E. Jellison, Jr. and F. A. Modine, "Parameterization of the optical functions of amorphous materials in the interband region," *Applied Physics Letters*, vol. 69, pp. 371–373, July 1996.

[164] "J. A Woollam Ellipsometry Tutorial." https://www.jawoollam.com/resources/ellipsometry-tutorial/thin-film-thickness.

[165] E. G. McKay, N. G. Pruiti, S. May, and M. Sorel, "Multipass lithography of HSQ etch masks for fabrication of low-loss alumina waveguides for blue light applications," in *Integrated Optics: Devices, Materials, and Technologies XXVII*, p. 124245, Jan. 2023.

[166] M. Di, E. Bersch, A. C. Diebold, S. Consiglio, R. D. Clark, G. J. Leusink, and T. Kaack, "Comparison of methods to determine bandgaps of ultrathin HfO_2 films using spectroscopic ellipsometry," *Journal of Vacuum Science & Technology A*, vol. 29, p. 041001, June 2011.

[167] "UQG Optics Heraeus Spectrosil 2000." https://www.uqgoptics.com/materials/heraeus-spectrosil-2000/.

[168] C. Gough, *Silicon-Rich Silicon Nitride: Tailoring Material Properties for Photonic Integrated Circuits*. PhD, University of Glasgow, Glasgow, 2024.

[169] E. Hecht, *Optics*. Pearson Education, 2017.

[170] O. Stenzel, *The Physics of Thin Film Optical Spectra: An Introduction*, vol. 44 of *Springer Series in Surface Sciences*. Cham: Springer International Publishing, 2016.

[171] P. Y. Yu and M. Cardona, *Fundamentals of Semiconductors: Physics and Materials Properties*. Graduate Texts in Physics, Berlin, Heidelberg: Springer, 2010.

[172] J. Heebner, R. Grover, and T. Ibrahim, *Optical Microresonators*, vol. 138 of *Optical Sciences*. New York, NY: Springer, 2008.

[173] I. H. Malitson, "Interspecimen Comparison of the Refractive Index of Fused Silica," *JOSA*, vol. 55, pp. 1205–1209, Oct. 1965.

[174] P. Makuła, M. Pacia, and W. Macyk, "How To Correctly Determine the Band Gap Energy of Modified Semiconductor Photocatalysts Based on UV–Vis Spectra," *The Journal of Physical Chemistry Letters*, vol. 9, pp. 6814–6817, Dec. 2018.

[175] S. Shi, S. Qian, X. Hou, J. Mu, J. He, and X. Chou, "Structural and Optical Properties of Amorphous Al_2O_3 Thin Film Deposited by Atomic Layer Deposition," *Advances in Condensed Matter Physics*, vol. 2018, p. e7598978, Mar. 2018.

[176] A. Mavrič, M. Fanetti, G. Mali, and M. Valant, "High-temperature stabilization of bulk amorphous Al_2O_3," *Journal of Non-Crystalline Solids*, vol. 499, pp. 363–370, Nov. 2018.

[177] A. Dutta, "Chapter 4 - Fourier Transform Infrared Spectroscopy," in *Spectroscopic Methods for Nanomaterials Characterization* (S. Thomas, R. Thomas, A. K. Zachariah, and R. K. Mishra, eds.), Micro and Nano Technologies, pp. 73–93, Elsevier, Jan. 2017.

[178] P. Nayar, A. Khanna, D. Kabiraj, S. R. Abhilash, B. D. Beake, Y. Losset, and B. Chen, "Structural, optical and mechanical properties of amorphous and crystalline alumina thin films," *Thin Solid Films*, vol. 568, pp. 19–24, Oct. 2014.

[179] R. Katamreddy, R. Inman, G. Jursich, A. Soulet, A. Nicholls, and C. Takoudis, "Post deposition annealing of aluminum oxide deposited by atomic layer deposition using tris(diethylamino)aluminum and water vapor on Si(100)," *Thin Solid Films*, vol. 515, pp. 6931–6937, June 2007.

[180] T. O. Kääriäinen and D. C. Cameron, "Plasma-Assisted Atomic Layer Deposition of Al_2O_3 at Room Temperature," *Plasma Processes and Polymers*, vol. 6, no. S1, pp. S237–S241, 2009.

[181] C. Morterra and G. Magnacca, "A case study: surface chemistry and surface structure of catalytic aluminas, as studied by vibrational spectroscopy of adsorbed species," *Catalysis Today*, vol. 27, pp. 497–532, Feb. 1996.

[182] L. Giacomazzi, N. S. Shcheblanov, M. E. Povarnitsyn, Y. Li, A. Mavrič, B. Zupančič, J. Grdadolnik, and A. Pasquarello, "Infrared spectra in amorphous alumina: A combined ab initio and experimental study," *Physical Review Materials*, vol. 7, p. 045604, Apr. 2023.

[183] S. K. Lee and S. Ryu, "Probing of Triply Coordinated Oxygen in Amorphous Al_2O_3," *The Journal of Physical Chemistry Letters*, vol. 9, pp. 150–156, Jan. 2018.

[184] C. Ramìrez-Espinoza, D. Salazar, R. Rangel-Rojo, J. L. Angel-Valenzuela, G. V. Vázquez, E. Flores-Romero, L. Rodríguez-Fernández, A. Oliver, D. A. Domínguez, and H. Márquez, "Design of Step-Index Optical Waveguides by Ion Implantation," *Journal of Lightwave Technology*, vol. 33, pp. 3052–3059, July 2015.

[185] A. C. Dillon, A. W. Ott, J. D. Way, and S. M. George, "Surface chemistry of Al_2O_3 deposition using $Al(CH_3)_3$ and H_2O in a binary reaction sequence," *Surface Science*, vol. 322, pp. 230–242, Jan. 1995.

[186] K. J. D. Mackenzie, "Infrared Frequency Calculations for Ideal Mullite ($3Al_2O_3{\cdot}SiO_2$)," *Journal of the American Ceramic Society*, vol. 55, no. 2, pp. 68–71, 1972.

[187] F. Carrillo, X. Colom, J. J. Suñol, and J. Saurina, "Structural FTIR analysis and thermal characterisation of lyocell and viscose-type fibres," *European Polymer Journal*, vol. 40, pp. 2229–2234, Sept. 2004.

[188] W. G. Golden, D. S. Dunn, C. E. Pavlik, and J. Overend, "The observation of the 1875 cm^{-1} band in the IR spectrum of CO on polycrystalline Pt," *The Journal of Chemical Physics*, vol. 70, pp. 4426–4427, May 1979.

[189] B. Gong and G. N. Parsons, "Quantitative in situ infrared analysis of reactions between trimethylaluminum and polymers during Al_2O_3 atomic layer deposition," *Journal of Materials Chemistry*, vol. 22, pp. 15672–15682, July 2012.

[190] A. Roy Chowdhuri and C. G. Takoudis, "Investigation of the aluminum oxide/Si(100) interface formed by chemical vapor deposition," *Thin Solid Films*, vol. 446, pp. 155–159, Jan. 2004.

[191] D. Rugar and P. Hansma, "Atomic Force Microscopy," *Physics Today*, vol. 43, pp. 23–30, Oct. 1990.

[192] D. Nečas and P. Klapetek, "Gwyddion: an open-source software for SPM data analysis," *Open Physics*, vol. 10, pp. 181–188, Feb. 2012.

[193] P. Eaton and P. West, "Chapter 5: AFM image processing and analysis," in *Atomic Force Microscopy*, Oxford University Press, Mar. 2010.

[194] I. O. Olarinoye and F. O. Ogundare, "Optical and microstructural properties of neutron irradiated RF-sputtered amorphous alumina thin films," *Optik*, vol. 134, pp. 66–77, Apr. 2017.

[195] S. Mardani, M. Dijkstra, W. A. P. Hendriks, M. P. Nijhuis-Groen, and S. M. García-Blanco, "Low-loss chemical mechanically polished Al_2O_3 thin films for UV integrated photonics," *Proceedings of European Conference on Integrated Optics (ECIO)*, vol. 23, pp. 390–392, May 2022.

[196] C. Hom-on, M. Horprathum, P. Eiamchai, S. Limwichean, V. Patthanasettakul, N. Nuntawong, C. Chananonnawathorn, T. Lertvanithphol, N. Triroj, and P. Jaroenapibal, "Surface roughness of aluminum oxide thin films deposited by DC and RF reactive magnetron sputtering," *Materials Today: Proceedings*, vol. 5, pp. 15228–15232, Jan. 2018.

[197] H. A. Al-Abadleh and V. H. Grassian, "FT-IR Study of Water Adsorption on Aluminum Oxide Surfaces," *Langmuir*, vol. 19, pp. 341–347, Jan. 2003.

[198] G. C. Correa, B. Bao, and N. C. Strandwitz, "Chemical Stability of Titania and Alumina Thin Films Formed by Atomic Layer Deposition," *ACS Applied Materials & Interfaces*, vol. 7, pp. 14816–14821, July 2015.

[199] V. Sammelselg, I. Netšipailo, A. Aidla, A. Tarre, L. Aarik, J. Asari, P. Ritslaid, and J. Aarik, "Chemical resistance of thin film materials based on metal oxides grown by atomic layer deposition," *Thin Solid Films*, vol. 542, pp. 219–224, Sept. 2013.

[200] A. I. Abdulagatov, Y. Yan, J. R. Cooper, Y. Zhang, Z. M. Gibbs, A. S. Cavanagh, R. G. Yang, Y. C. Lee, and S. M. George, "Al_2O_3 and TiO_2 Atomic Layer Deposition on Copper for Water Corrosion Resistance," *ACS Applied Materials & Interfaces*, vol. 3, pp. 4593–4601, Dec. 2011.

[201] M. M. Aslan, "Nanostructuring of alumina optical waveguides by hot water treatment for tuning sensor output," *Thin Solid Films*, vol. 520, pp. 2251–2256, Jan. 2012.

[202] P. P. Mardilovich, A. N. Govyadinoy, N. I. Mazurenko, and R. Paterson, "New and modified anodic alumina membranes part II. Comparison of solubility of amorphous (normal) and polycrystalline anodic alumina membranes," *Journal of Membrane Science*, vol. 98, pp. 143–155, Jan. 1995.

[203] A. Mavrič, M. Valant, C. Cui, and Z. M. Wang, "Advanced applications of amorphous alumina: From nano to bulk," *Journal of Non-Crystalline Solids*, vol. 521, p. 119493, Oct. 2019.

[204] V. V. Afanas'ev, A. Stesmans, B. J. Mrstik, and C. Zhao, "Impact of annealing-induced compaction on electronic properties of atomic-layer-deposited Al_2O_3," *Applied Physics Letters*, vol. 81, pp. 1678–1680, Aug. 2002.

[205] C. He, Y. Wang, C. Waldfried, G. Yang, J.-F. Zheng, S. Hu, and H. X. Tang, "Ultra-high Q alumina optical microresonators in the UV and blue bands," *Optics Express*, vol. 31, pp. 33923–33929, Oct. 2023.

[206] J. M. McHale, A. Auroux, A. J. Perrotta, and A. Navrotsky, "Surface Energies and Thermodynamic Phase Stability in Nanocrystalline Aluminas," *Science*, vol. 277, pp. 788–791, Aug. 1997.

[207] H. Liu, R. B. Yang, S. Guo, C. J. J. Lee, and N. L. Yakovlev, "Effect of annealing on structural and optical properties of ZnO/Al_2O_3 superlattice structures grown by atomic layer deposition at 150°C," *Journal of Alloys and Compounds*, vol. 703, pp. 225–231, May 2017.

[208] M. Broas, J. Lemettinen, T. Sajavaara, M. Tilli, V. Vuorinen, S. Suihkonen, and M. Paulasto-Kröckel, "In-situ annealing characterization of atomic-layer-deposited Al_2O_3 in N_2, H_2 and vacuum atmospheres," *Thin Solid Films*, vol. 682, pp. 147–155, July 2019.

[209] E. Beche, F. Fournel, V. Larrey, F. Rieutord, C. Morales, F. Madeira, and G. Audoit, "Direct Bonding Mechanism of ALD-Al_2O_3 Thin Films," *ECS Journal of Solid State Science and Technology*, vol. 4, pp. 171–175, Mar. 2015.

[210] I. N. Reddy, V. R. Reddy, N. Sridhara, V. S. Rao, M. Bhattacharya, P. Bandyopadhyay, S. Basavaraja, A. K. Mukhopadhyay, A. K. Sharma, and A. Dey, "Pulsed RF magnetron sputtered alumina thin films," *Ceramics International*, vol. 40, pp. 9571–9582, Aug. 2014.

[211] P. Vitanov, A. Harizanova, T. Ivanova, and T. Dimitrova, "Chemical deposition of Al_2O_3 thin films on Si substrates," *Thin Solid Films*, vol. 517, pp. 6327–6330, Oct. 2009.

[212] C. P. Rinsland, M. A. H. Smith, J. M. Russell, J. H. Park, and C. B. Farmer, "Stratospheric measurements of continuous absorption near 2400 cm^{-1}," *Applied Optics*, vol. 20, pp. 4167–4171, Dec. 1981.

[213] L. Zhang, H. C. Jiang, C. Liu, J. W. Dong, and P. Chow, "Annealing of Al_2O_3 thin films prepared by atomic layer deposition," *Journal of Physics D: Applied Physics*, vol. 40, pp. 3707–3713, June 2007.

[214] T.-S. Chang, T.-C. Chang, P.-T. Liu, T.-S. Chang, C.-H. Tu, and F.-S. Yeh, "Improvement of Hydrogenated Amorphous-Silicon TFT Performances With Low-k Siloxane-Based Hydrogen Silsesquioxane (HSQ) Passivation Layer," *IEEE Electron Device Letters*, vol. 27, pp. 902–904, Nov. 2006.

[215] K. Ronse, P. De Bisschop, G. Vandenberghe, E. Hendrickx, R. Gronheid, A. V. Pret, A. Mallik, D. Verkest, and A. Steegen, "Opportunities and challenges in device scaling by the introduction of EUV lithography," in *2012 International Electron Devices Meeting*, pp. 18.5.1–18.5.4, Dec. 2012.

[216] N. G. Pruiti, C. Klitis, C. Gough, S. May, and M. Sorel, "Thermo-optic coefficient of PECVD silicon-rich silicon nitride," *Optics Letters*, vol. 45, pp. 6242–6245, Nov. 2020.

[217] S. May, M. Clerici, and M. Sorel, "Supercontinuum generation in dispersion engineered AlGaAs-on-insulator waveguides," *Scientific Reports*, vol. 11, p. 2052, Jan. 2021.

[218] S. Thoms and P. Reynolds, "Ebpg 5200 introduction and workflow," Dec. 2020.

[219] T. H. P. Chang, "Proximity effect in electron-beam lithography," *Journal of Vacuum Science and Technology*, vol. 12, pp. 1271–1275, Nov. 1975.

[220] R. D. Moore, G. A. Caccoma, H. C. Pfeiffer, E. V. Weber, and O. C. Woodard, "EL-3: A high throughput, high resolution e-beam lithography tool," *Journal of Vacuum Science and Technology*, vol. 19, pp. 950–952, Nov. 1981.

[221] GeniSys GmBH, "BEAMER Applications: High Resolution and Fast Writing. App-Note Rev. 20200415," 2020.

[222] C.-C. Yang and W.-C. Chen, "The structures and properties of hydrogen silsesquioxane (HSQ) films produced by thermal curing," *Journal of Materials Chemistry*, vol. 12, no. 4, pp. 1138–1141, 2002.

[223] S. Choi, M. J. Word, V. Kumar, and I. Adesida, "Comparative study of thermally cured and electron-beam-exposed hydrogen silsesquioxane resists," *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena*, vol. 26, p. 1654, Sept. 2008.

[224] H. Namatsu, "Three-dimensional siloxane resist for the formation of nanopatterns with minimum linewidth fluctuations," *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures*, vol. 16, p. 69, Jan. 1998.

[225] Y. M. Georgiev, W. Henschel, A. Fuchs, and H. Kurz, "Surface roughness of hydrogen silsesquioxane as a negative tone electron beam resist," *Vacuum*, vol. 77, pp. 117–123, Jan. 2005.

[226] C. Mack, *Fundamental Principles of Optical Lithography: The Science of Microfabrication*. John Wiley & Sons, Dec. 2007.

[227] A. Jeyakumar and C. L. Henderson, "Enhancing the electron beam sensitivity of hydrogen silsesquioxane (HSQ)," in *Advances in Resist Technology and Processing XXI*, vol. 5376, pp. 490–501, SPIE, May 2004.

[228] J. N. Bremmer, D. Gray, Y. Liu, K. Gruszynski, and S. Marcus, "Rapid Thermal Processing of Hydrogen Silsesquioxane for Low Dielectric Constant Performance," *MRS Online Proceedings Library (OPL)*, vol. 565, p. 273, 1999.

[229] J. Vila-Comamala, S. Gorelick, V. A. Guzenko, and C. David, "3D Nanostructuring of hydrogen silsesquioxane resist by 100keV electron beam lithography," *Journal of Vacuum Science & Technology B*, vol. 29, p. 06F301, Nov. 2011.

[230] T. Nakamura, M. Sasaki, A. Kobayashi, K. Sawa, and K. Mine, "Oxidative Curing of Hydrogen Silsesquioxane Resin Films by Electron Beam Irradiation without Additional Heatings and Characterization of the Cured Films," *Japanese Journal of Applied Physics*, vol. 40, p. 6187, Nov. 2001.

[231] D. Lauvernier, S. Garidel, C. Legrand, and J.-P. Vilcot, "Realization of sub-micron patterns on GaAs using a HSQ etching mask," *Microelectronic Engineering*, vol. 77, pp. 210–216, Apr. 2005.

[232] W. Jiang, D. Xu, B. Xiong, and Y. Wang, "Effects of rapid thermal annealing on LPCVD silicon nitride," *Ceramics International*, vol. 42, pp. 1217–1224, Jan. 2016.

[233] D.-H. Kim and J.-H. Jang, "Effect of Post Exposure Bake of Hydrogen Silsesquioxane on the Nanoscale Grating Pattern," *ECS Meeting Abstracts*, vol. MA2009-02, p. 156, July 2009.

[234] D.-H. Kim, S.-K. Kang, G.-Y. Yeom, and J.-H. Jang, "Nanometer-Scale Fabrication of Hydrogen Silsesquioxane (HSQ) Films with Post Exposure Baking," *Journal of Nanoscience and Nanotechnology*, vol. 13, pp. 1918–1922, Mar. 2013.

[235] T. Mpatzaka, G. Zisis, I. Raptis, V. Vamvakas, C. Kaiser, T. Mai, M. Schirmer, M. Gerngroß, and G. Papageorgiou, "Process study and the lithographic performance of commercially available silsesquioxane based electron sensitive resist Medusa 82," *Micro and Nano Engineering*, vol. 8, p. 100065, Aug. 2020.

[236] J. Bolten, C. Manecke, T. Wahlbrink, M. Waldow, and H. Kurz, "At low costs: Study on optical propagation losses of silicon waveguides fabricated by electron beam lithography," *Microelectronic Engineering*, vol. 123, pp. 1–3, July 2014.

[237] S. Babin, H. Hartmann, and I. Y. Kuzmin, "Simulation and measurement of resist heating in multipass exposure using a 50 kV variably shaped beam system," *Microelectronic Engineering*, vol. 46, pp. 231–234, May 1999.

[238] K. Nojiri, *Dry Etching Technology for Semiconductors*. Cham: Springer International Publishing, 2015.

[239] S. A. Campbell, *Fabrication engineering at the micro- and nanoscale*. New York, NY: Oxford University Press, fourth ed., 2013.

[240] S. Tegen and P. Moll, "Etch Characteristics of Al_2O_3 in ICP and MERIE Plasma Etchers," *Journal of The Electrochemical Society*, vol. 152, p. G271, Mar. 2005.

[241] D.-P. Kim, J.-W. Yeo, and C.-I. Kim, "Etching properties of Al_2O_3 films in inductively coupled plasma," *Thin Solid Films*, vol. 459, pp. 122–126, July 2004.

[242] J. W. Lee, B. Pathangey, M. R. Davidson, P. H. Holloway, E. S. Lambers, B. Davydov, T. J. Anderson, and S. J. Pearton, "Comparison of plasma chemistries for dry etching thin film electroluminescent display materials," *Journal of Vacuum Science & Technology A: Vacuum, Surfaces, and Films*, vol. 16, pp. 2177–2186, July 1998.

[243] S.-M. Koo, D.-P. Kim, K.-T. Kim, and C.-I. Kim, "The etching properties of Al_2O_3 thin films in $N_2/Cl_2/BCl_3$ and $Ar/Cl_2/BCl_3$ gas chemistry," *Materials Science and Engineering: B*, vol. 118, pp. 201–204, Apr. 2005.

[244] D. W. Kim, C. H. Jeong, K. N. Kim, H. Y. Lee, H. S. Kim, Y. J. Sung, and G. Y. Yeom, "High rate sapphire (Al_2O_3) etching in inductively coupled plasmas using axial external magnetic field," *Thin Solid Films*, vol. 435, pp. 242–246, July 2003.

[245] J.-C. Woo, D.-S. Um, and C. I. Kim, "Dry Etching of Al_2O_3 Thin Films in $O_2/BCl_3/Ar$ Inductively Coupled Plasma," *Transactions on Electrical and Electronic Materials*, vol. 11, pp. 202–205, Oct. 2010.

[246] V. Biryukova, G. J. Sharp, C. Klitis, and M. Sorel, "Trimming of silicon-on-insulator ring-resonators via localized laser annealing," *Optics Express*, vol. 28, pp. 11156–11164, Apr. 2020.

[247] C. D. Vita, F. Toso, N. G. Pruiti, C. Klitis, G. Ferrari, M. Sorel, A. Melloni, and F. Morichetti, "Amorphous-silicon visible-light detector integrated on silicon nitride waveguides," *Optics Letters*, vol. 47, pp. 2598–2601, May 2022.

[248] Z. A. Weinberg, G. W. Rubloff, and E. Bassous, "Transmission, photoconductivity, and the experimental band gap of thermally grown SiO_2 films," *Physical Review B*, vol. 19, pp. 3107–3117, Mar. 1979.

[249] Y. K. Siew, G. Sarkar, X. Hu, J. Hui, A. See, and C. T. Chua, "Thermal Curing of Hydrogen Silsesquioxane," *Journal of The Electrochemical Society*, vol. 147, p. 335, Jan. 2000.

[250] H.-C. Liou and J. Pretzer, "Effect of curing temperature on the mechanical properties of hydrogen silsesquioxane thin films," *Thin Solid Films*, vol. 335, pp. 186–191, Nov. 1998.

[251] H. Baghsiahi, K. Wang, W. Kandulski, R. C. A. Pitwon, and D. R. Selviah, "Optical Waveguide End Facet Roughness and Optical Coupling Loss," *Journal of Lightwave Technology*, vol. 16, no. 31, pp. 2959–2968, 2013.

[252] J. Cardenas, C. B. Poitras, K. Luke, L.-W. Luo, P. A. Morton, and M. Lipson, "High Coupling Efficiency Etched Facet Tapers in Silicon Waveguides," *IEEE Photonics Technology Letters*, vol. 26, pp. 2380–2382, Dec. 2014.

[253] X. Ji, S. Roberts, M. Corato-Zanarella, and M. Lipson, "Methods to achieve ultra-high quality factor silicon nitride resonators," *APL Photonics*, vol. 6, p. 071101, July 2021.

[254] N. Chauhan, J. Wang, D. Bose, K. Liu, R. L. Compton, C. Fertig, C. W. Hoyt, and D. J. Blumenthal, "Ultra-low loss visible light waveguides for integrated atomic, molecular, and quantum photonics," *Optics Express*, vol. 30, pp. 6960–6969, Feb. 2022.

[255] J.-m. Liu, *Waveguides and couplers*, p. 71–72. Cambridge University Press, 2005.

[256] A. Boudrioua, *Optical Waveguide Theory*, ch. 1, pp. 1–37. John Wiley & Sons, Ltd, 2009.

[257] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*. John Wiley & Sons, third ed., Feb. 2019.

[258] Lumerical Inc., "Lumerical MODE, R1.2," 2021.

[259] P. Steglich, "Silicon-on-Insulator Slot Waveguides: Theory and Applications in Electro-Optics and Optical Sensing," in *Emerging Waveguide Technology*, IntechOpen, Aug. 2018.

[260] A. Banerjee, *Optical Waveguides Analysis and Design*. Cham: Springer International Publishing, 2022.

[261] D. Marcuse, *Theory of Dielectric Optical Waveguides*. Elsevier, Sept. 2013.

[262] A. Z. Subramanian, P. Neutens, A. Dhakal, R. Jansen, T. Claes, X. Rottenberg, F. Peyskens, S. Selvaraja, P. Helin, B. Du Bois, K. Leyssens, S. Severi, P. Deshpande, R. Baets, and P. Van Dorpe, "Low-Loss Singlemode PECVD Silicon Nitride Photonic Wire Waveguides for 532-900nm Wavelength Window Fabricated Within a CMOS Pilot Line," *IEEE Photonics Journal*, vol. 5, p. 2202809, Dec. 2013.

[263] K. Yee, "Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media," *IEEE Transactions on Antennas and Propagation*, vol. 14, pp. 302–307, May 1966.

[264] Z. Zhu and T. G. Brown, "Full-vectorial finite-difference analysis of microstructured optical fibers," *Optics Express*, vol. 10, pp. 853–864, Aug. 2002.

[265] S. Chongsawangvirod, E. A. Irene, A. Kalnitsky, S. P. Tay, and J. P. Ellul, "Refractive Index Profiles of Thermally Grown and Chemically Vapor Deposited Films on Silicon," *Journal of The Electrochemical Society*, vol. 137, p. 3536, Nov. 1990.

[266] E. a. J. Marcatili and S. E. Miller, "Improved Relations Describing Directional Control in Electromagnetic Wave Guidance," *Bell System Technical Journal*, vol. 48, no. 7, pp. 2161–2188, 1969.

[267] W. Bogaerts and S. K. Selvaraja, "Compact Single-Mode Silicon Hybrid Rib/Strip Waveguide With Adiabatic Bends," *IEEE Photonics Journal*, vol. 3, pp. 422–432, June 2011.

[268] K. Coimbatore Balram, D. A. Westly, M. I. Davanco, K. E. Grutter, Q. Li, T. Michels, C. H. Ray, R. J. Kasica, C. B. Wallin, I. J. Gilbert, B. A. Bryce, G. Simelgor, J. Topolancik, N. Lobontiu, Y. Liu, P. Neuzil, V. Svatos, K. A. Dill, N. A. Bertrand, M. Metzler, G. Lopez, D. Czaplewski, L. Ocola, K. A. Srinivasan, S. M. Stavis, V. A. Aksyuk, J. A. A. Liddle, S. Krylov, and R. R. Ilic, "The Nanolithography Toolbox," *Journal of Research of the National Institute of Standards and Technology*, vol. 121, p. 464, Oct. 2016.

[269] H. P. Bazargani, J. Flueckiger, L. Chrostowski, and J. Azaña, "Microring resonator design with improved quality factors using quarter Bezier curves," in *2015 Conference on Lasers and Electro-Optics (CLEO)*, pp. 1–2, May 2015.

[270] G. Son, S. Han, J. Park, K. Kwon, and K. Yu, "High-efficiency broadband light coupling between optical fibers and photonic integrated circuits," *Nanophotonics*, vol. 7, pp. 1845–1864, Dec. 2018.

[271] R. Marchetti, C. Lacava, L. Carroll, K. Gradkowski, and P. Minzioni, "Coupling strategies for silicon photonics integrated chips," *Photonics Research*, pp. 201–239, Jan. 2019.

[272] M. Bahadori, M. Nikdast, Q. Cheng, and K. Bergman, "Universal Design of Waveguide Bends in Silicon-on-Insulator Photonics Platform," *Journal of Lightwave Technology*, vol. 37, pp. 3044–3054, July 2019.

[273] X. Shu, L. J. Beckmann, and H. F. Zhang, "Visible-light optical coherence tomography: a review," *Journal of Biomedical Optics*, vol. 22, p. 121707, Dec. 2017.

[274] A. Zheltikov, A. L'Huillier, and F. Krausz, "Nonlinear Optics," in *Springer Handbook of Lasers and Optics* (F. Träger, ed.), Springer Handbooks, pp. 161–251, Berlin, Heidelberg: Springer, 2012.

[275] G. Agrawal, ed., *Nonlinear Fiber Optics*. Optics and Photonics, Boston: Academic Press, fifth ed., Jan. 2013.

[276] R. R. Alfano, *The Supercontinuum Laser Source*. New York, NY: Springer, 2016.

[277] K. Thyagarajan and A. K. Ghatak, "Fiber and Guided Wave Optics | Nonlinear Optics," in *Encyclopedia of Modern Optics* (R. D. Guenther, ed.), pp. 472–486, Oxford: Elsevier, Jan. 2005.

[278] L. F. Mollenauer and J. P. Gordon, *Solitons in Optical Fibers: Fundamentals and Applications*. Elsevier, Mar. 2006.

[279] J. H. Lee, J. van Howe, C. Xu, and X. Liu, "Soliton Self-Frequency Shift: Experimental Demonstrations and Applications," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 14, pp. 713–723, May 2008.

[280] H. Guo, C. Herkommer, A. Billat, D. Grassani, C. Zhang, M. H. P. Pfeiffer, W. Weng, C.-S. Brès, and T. J. Kippenberg, "Mid-infrared frequency comb via coherent dispersive wave generation in silicon nitride nanophotonic waveguides," *Nature Photonics*, vol. 12, pp. 330–335, June 2018.

[281] A. L. Gaeta, "Nonlinear propagation and continuum generation in microstructured optical fibers," *Optics Letters*, vol. 27, pp. 924–926, June 2002.

[282] D. V. Skryabin and A. V. Yulin, "Theory of generation of new frequencies by mixing of solitons and dispersive waves in optical fibers," *Physical Review E*, vol. 72, p. 016619, July 2005.

[283] D. Castelló-Lurbe and E. Silvestre, "Supercontinuum generation in silicon waveguides relying on wave-breaking," *Optics Express*, vol. 23, pp. 25462–25473, Oct. 2015.

[284] M. Conforti and S. Trillo, "Dispersive wave emission from wave breaking," *Optics Letters*, vol. 38, pp. 3815–3818, Oct. 2013.

[285] K. E. Webb, Y. Q. Xu, M. Erkintalo, and S. G. Murdoch, "Generalized dispersive wave emission in nonlinear fiber optics," *Optics Letters*, vol. 38, pp. 151–153, Jan. 2013.

[286] L. Zhang, Q. Lin, Y. Yue, Y. Yan, R. G. Beausoleil, and A. E. Willner, "Silicon waveguide with four zero-dispersion wavelengths and its application in on-chip octave-spanning supercontinuum generation," *Optics Express*, vol. 20, pp. 1685–1690, Jan. 2012.

[287] Y. Kodama and A. Hasegawa, "Nonlinear pulse propagation in a monomode dielectric guide," *IEEE Journal of Quantum Electronics*, vol. 23, pp. 510–524, May 1987.

[288] R. E. Slusher, G. Lenz, J. Hodelin, J. Sanghera, L. B. Shaw, and I. D. Aggarwal, "Large Raman gain and nonlinear phase shifts in high-purity As_2Se_3 chalcogenide fibers," *JOSA B*, vol. 21, pp. 1146–1155, June 2004.

[289] A. Raza, S. Clemmen, P. Wuytens, M. de Goede, A. S. K. Tong, N. L. Thomas, C. Liu, J. Suntivich, A. G. Skirtach, S. M. García-Blanco, D. J. Blumenthal, J. S. Wilkinson, and R. Baets, "High index contrast photonic platforms for on-chip Raman spectroscopy," *Optics Express*, vol. 27, pp. 23067–23079, Aug. 2019.

[290] N. Singh, J. Lorenzen, M. Sinobad, K. Wang, A. C. Liapis, H. C. Frankis, S. Haugg, H. Francis, J. Carreira, M. Geiselmann, M. A. Gaafar, T. Herr, J. D. B. Bradley, Z. Sun, S. M. García-Blanco, and F. X. Kärtner, "Silicon photonics-based high-energy passively Q-switched laser," *Nature Photonics*, pp. 1–7, Feb. 2024.

[291] D. T. H. Tan, A. M. Agarwal, and L. C. Kimerling, "Nonlinear photonic waveguides for on-chip optical pulse compression," *Laser & Photonics Reviews*, vol. 9, no. 3, pp. 294–308, 2015.

[292] R. Halir, Y. Okawachi, J. S. Levy, M. A. Foster, M. Lipson, and A. L. Gaeta, "Ultrabroadband supercontinuum generation in a CMOS-compatible platform," *Optics Letters*, vol. 37, pp. 1685–1687, May 2012.

[293] T.-H. Wu, L. Ledezma, C. Fredrick, P. Sekhar, R. Sekine, Q. Guo, R. M. Briggs, A. Marandi, and S. A. Diddams, "Visible-to-ultraviolet frequency comb generation in lithium niobate nanophotonic waveguides," *Nature Photonics*, vol. 18, pp. 218–223, Mar. 2024.

[294] C. Ciret and S.-P. Gorza, "Generation of ultra-broadband coherent supercontinua in tapered and dispersion-managed silicon nanophotonic waveguides," *JOSA B*, vol. 34, pp. 1156–1162, June 2017.

[295] WiseLabAEP, "GMMNLSE-Solver-FINAL," Jan. 2017.

[296] J. Monroy-Ruz, C. S. Kline, J. Smith, J. Rarity, and K. C. Balram, "Circular Bragg grating resonators for optical read-out of NV centres in nanodiamonds encapsulated in silicon nitride," in *2021 Conference on Lasers and Electro-Optics Europe & European Quantum Electronics Conference (CLEO/Europe-EQEC)*, p. 1, June 2021.

[297] H. Chen, H. Fu, X. Huang, X. Zhang, T.-H. Yang, J. A. Montes, I. Baranowski, and Y. Zhao, "Low loss GaN waveguides at the visible spectral wavelengths for integrated photonics applications," *Optics Express*, vol. 25, pp. 31758–31773, Dec. 2017.

[298] S. Romero-García, F. Merget, F. Zhong, H. Finkelstein, and J. Witzens, "Silicon nitride CMOS-compatible platform for integrated photonics applications at visible wavelengths," *Optics Express*, vol. 21, pp. 14036–14046, June 2013.

[299] E. Shim, Y. Chen, S. Masmanidis, and M. Li, "Multisite silicon neural probes with integrated silicon nitride waveguides and gratings for optogenetic applications," *Scientific Reports*, vol. 6, p. 22693, Mar. 2016.

[300] C. Sorace-Agaskar, S. Bramhavar, D. Kharas, W. Loh, P. W. Juodawlkis, J. Chiaverini, and J. M. Sage, "Multi-layer integrated photonics from the ultraviolet to the infrared,"

in *Frontiers in Biological Detection: From Nanosensors to Systems X*, vol. 10510, p. 105100D, International Society for Optics and Photonics, Feb. 2018.

[301] W. D. Sacher, X. Luo, Y. Yang, F.-D. Chen, T. Lordello, J. C. C. Mak, X. Liu, T. Hu, T. Xue, P. G.-Q. Lo, M. L. Roukes, and J. K. S. Poon, "Visible-light silicon nitride waveguide devices and implantable neurophotonic probes on thinned 200mm silicon wafers," *Optics Express*, vol. 27, pp. 37400–37418, Dec. 2019.

[302] M. Sinclair, K. Gallacher, M. Sorel, J. C. Bayley, E. McBrearty, R. W. Millar, S. Hild, and D. J. Paul, "1.4 Million Q-Factor 780 nm Wavelength Si_3N_4 Micro-Rings for Chip-Scale Atomic Systems," in *2020 Conference on Lasers and Electro-Optics (CLEO)*, pp. 1–2, May 2020.

[303] Y. Zhao, X. Ji, B. Y. Kim, P. S. Donvalkar, J. K. Jang, C. Joshi, M. Yu, C. Joshi, R. R. Domeneguetti, F. A. S. Barbosa, P. Nussenzveig, Y. Okawachi, M. Lipson, and A. L. Gaeta, "Visible nonlinear photonics via high-order-mode dispersion engineering," *Optica*, vol. 7, pp. 135–141, Feb. 2020.

[304] R. R. Domeneguetti, Y. Zhao, X. Ji, M. Martinelli, M. Lipson, A. L. Gaeta, and P. Nussenzveig, "Parametric sideband generation in CMOS-compatible oscillators from visible to telecom wavelengths," *Optica*, vol. 8, pp. 316–322, Mar. 2021.

[305] M. Lelit, M. Słowikowski, M. Filipiak, M. Juchniewicz, B. Stonio, B. Michalak, K. Pavłov, M. Myśliwiec, P. Wiśniewski, A. Kaźmierczak, K. Anders, S. Stopiński, R. B. Beck, and R. Piramidowicz, "Passive Photonic Integrated Circuits Elements Fabricated on a Silicon Nitride Platform," *Materials*, vol. 15, p. 1398, Jan. 2022.

[306] S. M. Lindecrantz and O. G. Hellesø, "Estimation of Propagation Losses for Narrow Strip and Rib Waveguides," *IEEE Photonics Technology Letters*, vol. 26, pp. 1836–1839, Sept. 2014.

[307] Y.-Y. Lin, C.-L. Wu, W.-C. Chi, Y.-J. Chiu, Y.-J. Hung, A.-K. Chu, and C.-K. Lee, "Self-phase modulation in highly confined submicron Ta_2O_5 channel waveguides," *Optics Express*, vol. 24, pp. 21633–21641, Sept. 2016.

[308] J. H. Sierra, R. C. Rangel, R. E. Samad, N. D. Vieira, M. I. Alayo, and D. O. Carvalho, "Low-loss pedestal Ta_2O_5 nonlinear optical waveguides," *Optics Express*, vol. 27, pp. 37516–37521, Dec. 2019.

[309] J. T. Choy, J. D. B. Bradley, P. B. Deotare, I. B. Burgess, C. C. Evans, E. Mazur, and M. Lončar, "Integrated TiO_2 resonators for visible photonics," *Optics Letters*, vol. 37, pp. 539–541, Feb. 2012.

[310] C. C. Evans, C. Liu, and J. Suntivich, "Low-loss titanium dioxide waveguides and resonators using a dielectric lift-off fabrication process," *Optics Express*, vol. 23, pp. 11160–11169, May 2015.

[311] D. D. Hickstein, H. Jung, D. R. Carlson, A. Lind, I. Coddington, K. Srinivasan, G. G. Ycas, D. C. Cole, A. Kowligy, C. Fredrick, S. Droste, E. S. Lamb, N. R. Newbury, H. X. Tang, S. A. Diddams, and S. B. Papp, "Ultrabroadband Supercontinuum Generation and Frequency-Comb Stabilization Using On-Chip Waveguides with Both Cubic and Quadratic Nonlinearities," *Physical Review Applied*, vol. 8, p. 014025, July 2017.

[312] A. Shams-Ansari, P. Latawiec, Y. Okawachi, V. Venkataraman, M. Yu, B. Desiatov, H. Atikian, G. L. Harris, N. Picqué, A. L. Gaeta, and M. Lončar, "Supercontinuum generation in angle-etched diamond waveguides," *Optics Letters*, vol. 44, pp. 4056–4059, Aug. 2019.

[313] J. Lu, J. B. Surya, X. Liu, Y. Xu, and H. X. Tang, "Octave-spanning supercontinuum generation in nanoscale lithium niobate waveguides," *Optics Letters*, vol. 44, pp. 1492–1495, Mar. 2019.

[314] M. Yu, B. Desiatov, Y. Okawachi, A. L. Gaeta, and M. Lončar, "Coherent two-octave-spanning supercontinuum generation in lithium-niobate waveguides," *Optics Letters*, vol. 44, pp. 1222–1225, Mar. 2019.

[315] M. Reig Escalé, F. Kaufmann, H. Jiang, D. Pohl, and R. Grange, "Generation of 280 THz-spanning near-ultraviolet light in lithium niobate-on-insulator waveguides with sub-100 pJ pulses," *APL Photonics*, vol. 5, p. 121301, Dec. 2020.

[316] A. R. Johnson, A. S. Mayer, A. Klenner, K. Luke, E. S. Lamb, M. R. E. Lamont, C. Joshi, Y. Okawachi, F. W. Wise, M. Lipson, U. Keller, and A. L. Gaeta, "Octave-spanning coherent supercontinuum generation in a silicon nitride waveguide," *Optics Letters*, vol. 40, pp. 5117–5120, Nov. 2015.

[317] H. Zhao, B. Kuyken, S. Clemmen, F. c. Leo, A. Subramanian, A. Dhakal, P. Helin, S. Severi, E. Brainis, G. Roelkens, and R. Baets, "Visible-to-near-infrared octave spanning supercontinuum generation in a silicon nitride waveguide," *Optics Letters*, vol. 40, pp. 2177–2180, May 2015.

[318] Y. Okawachi, M. Yu, J. Cardenas, X. Ji, M. Lipson, and A. L. Gaeta, "Coherent, directional supercontinuum generation," *Optics Letters*, vol. 42, pp. 4466–4469, Nov. 2017.

[319] C. C. Evans, K. Shtyrkova, J. D. B. Bradley, O. Reshef, E. Ippen, and E. Mazur, "Spectral broadening in anatase titanium dioxide waveguides at telecommunication and near-visible wavelengths," *Optics Express*, vol. 21, pp. 18582–18591, July 2013.

[320] C. Taballione, M. C. Anguita, M. de Goede, P. Venderbosch, B. Kassenberg, H. Snijders, N. Kannan, D. Smith, J. P. Epping, R. van der Meer, P. W. H. Pinkse, H. v. d. Vlekkert, and J. J. Renema, "20-Mode Universal Quantum Photonic Processor," June 2022. arXiv:2203.01801 [quant-ph].