UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND PHYSICAL SCIENCES

Optoelectronics Research Centre

Simulation, Fabrication and Assembly Techniques for Passive Alignment of Silicon Photonic Integrated Circuits

---

By

Nathan Soper

Thesis for the

degree of Doctor of Philosophy

November 2019

Dedicated to Nana

26/12/1933 - 25/11/2017

“Who can find a virtuous woman? for her price is far above rubies.”

Proverbs 31:10

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL SCIENCES AND ENGINEERING
Optoelectronics Research Centre

DOCTOR OF PHILOSOPHY

SIMULATION, FABRICATION AND ASSEMBLY TECHNIQUES FOR PASSIVE ALIGNMENT OF SILICON PHOTONIC INTEGRATED CIRCUITS

by Nathan Soper

This thesis demonstrates a number of simulation, fabrication and assembly techniques that have been designed, developed and tested in order to facilitate the passive alignment of silicon photonic integrated circuits to allow a packaging method that can decrease the cost and time penalties associated with active alignment methods. The requirement for an affordable packaging method is highlighted and key achievements are identified in the literature.

Novel and unique surface grating couplers are demonstrated to be compatible with an in-plane passively aligned packaging method which exhibit insertion losses of 1.5 dB in simulation and 2.59dB in physical characterisation with 1- and 3-dB bandwidths of 35nm and 59nm respectively, as well as having improved misalignment tolerance.

Additionally, a novel means of deep-silicon etching is developed which allows customised, packaging-specific structures to be fabricated allowing passively aligned assembly within the packaging method described in this thesis by improving vision and access of a fibre, whilst decreasing the free-space propagation distance. This process has been optimised through dynamic etch parameters.

Finally, this thesis also outlines a number of key innovations that allow the passive assembly of a hybrid, two-chip packaging method by using silica microspheres and epoxy, to bond and align a silicon chip to an optical fibre.

Contents

List of figures

Declaration of authorship

I, Nathan Soper, declare that this thesis entitled “Simulation, Fabrication and Assembly Techniques for Passive Alignment of Silicon Photonic Integrated Circuits” and the work presented in it are my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published in the following publications, as indicated in the relevant chapters:

Signed: .....

Date: .....

Acknowledgements

There are many people who need to be thanked and acknowledged due to their influence during this project. Of course, first mention must go to my supervisors: Graham, Goran and Dave who have always provided help, advice, wisdom and encouragement. Next the rest of the Silicon Photonics Group at Southampton which has now grown to be too numerous to mention everyone but Dr Scott Reynolds deserves a special mention and it's been a pleasure to have worked alongside him and his hard work and dedication has been priceless. Dr Ali Khokar's knowledge and expertise with e-beam lithography and cleanroom processing in general has been invaluable over the last few years. In addition, I would like to thank the hard work of the Zepler Institute technical staff who have worked tirelessly to keep the cleanroom functional allowing this research to have been possible.

Chapter 1 Introduction

1.1 Introduction to Silicon Photonics in General

In the 21st century we find ourselves on the "cusp of revolutionary change"[1]. The availability of data and communications has drastically altered the way in which we live, work and conduct our lives. More data is being produced than ever before and more of it is also being transmitted and received than ever before. The basic human need to communicate has migrated into cyberspace. In fact, many have proposed that we are now moving into a "knowledge-based economy"[2], [3], with the collapse of previous economic systems based on services or industry. The technological infrastructure that underpins the storage and transfer of knowledge (information) then becomes vitally important.

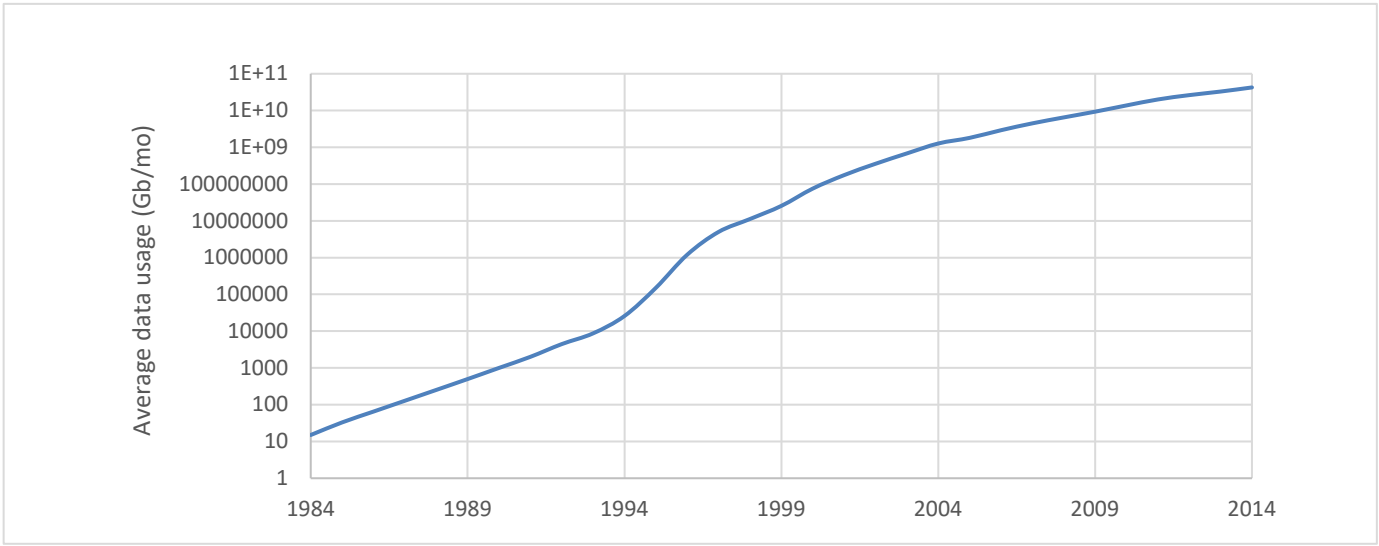

Figure 1.1 Graph showing average internet usage across the world from 1984 to 2014 [4]

Figure 1.1 shows the average amount of internet data that has been transferred over a 30 year period from 1984 to 2014, as provided by Cisco[4]. As can been seen, there is a general upwards trend over that period. However, the reader is pointed

towards the scale of the vertical axis, which is indeed plotted logarithmically. In fact, to plot this chart on a linear axis is virtually impossible, such is the colossal growth in data usage and transfer in that period. This chart shows an increase of internet traffic of over 2.8 billion times in that 30-year period. Growths of this magnitude are very rare in technology and industry and can only really be compared to that of monetary hyperinflation[5]–[8].

This then presents a problem for the engineer and for society. How can the predicted "capacity crunch"[9] be avoided? Since the invention of the optical fibre for telecoms by Charles Kao[10], photonics has been suggested, developed and applied to scenarios where very high capacity data streams are needed. Traditionally this has been for long haul networks across oceans and continents[11]. The advent of Erbium Doped Fibre Amplifiers (EDFA) has only made their use more ubiquitous[12]–[14].

Silicon Photonics has been proposed as a means of maintaining the future growth of both long and short reach communications for over a decade[15]. Silicon Photonics is seen as an answer to many of the challenges surrounding the current expansion, such as how to reduce the costs of transmitting large volumes of data and how to provide greater capacity without scaling up complexity, hence avoiding a situation of technological overextension.

Silicon Photonics seeks to piggyback on existing CMOS technology and infrastructure that has been developed from the electronics industry. Process flows have been created that can recycle CMOS tools [16]. Most of the exact same tools and equipment are used at the design, fabrication and packaging stages of Silicon Photonic processing.

Silicon Photonics has gained a great deal of interest in the last couple decades. Figure 1.2 shows a graph, modified from one presented by Silicon Photonics pioneer Roel Baets. The graph depicts the number of publications with key words "silicon photonics" or similar that have been produced over a 20-year period, from 1996 to 2016, around the beginning of this project. There is a sharp rise in the total number of publications, indicating increased academic interest and research spending in that time.

A range of devices and circuits have been demonstrated within the field of Silicon Photonics including modulators [17], [18], filters[19], [20], transceivers [21] and more. These devices have found broad application, including in data and communications as mentioned previously as well as sensing, defence, medicine and others[22].

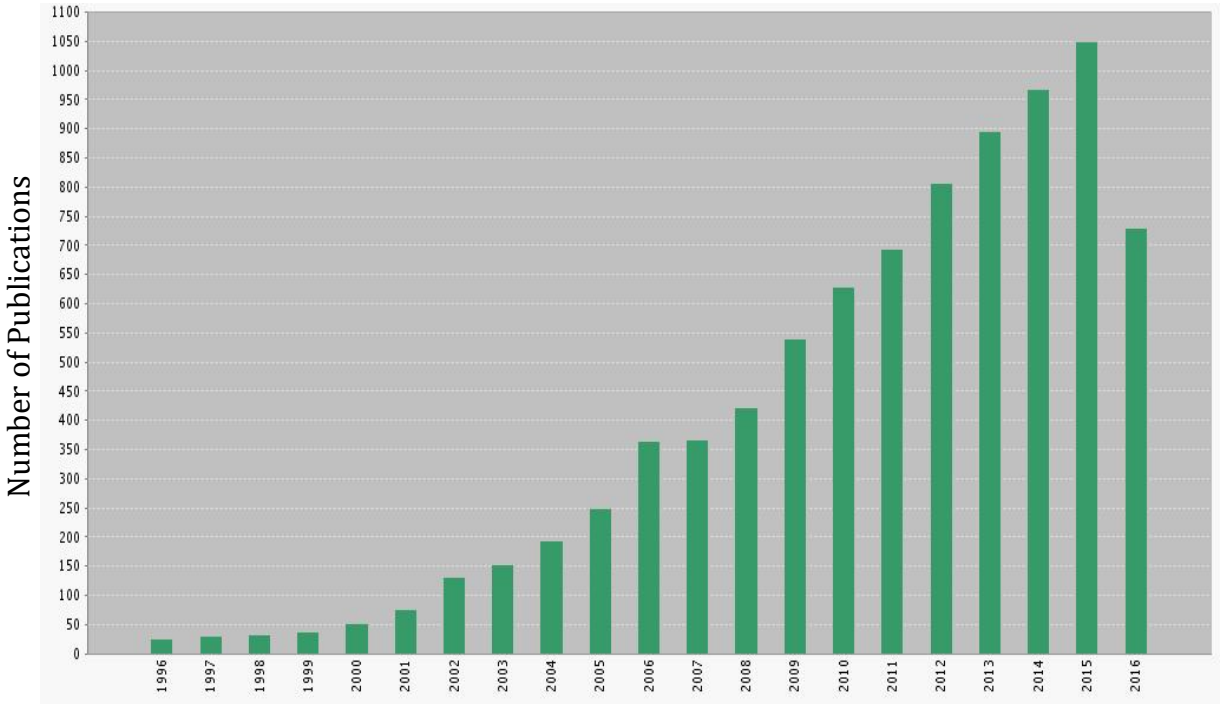

Figure 1.2 Graph of increase in the number of publications in Silicon Photonics from 1996-2016

Silicon photonics now has been established as a platform for +25Gb/s data coms[23], with comprehensive roadmaps already outlined[24].

The interest in Silicon Photonic research has been buttressed with newer multi-project wafer services[25] and fabless silicon photonics[26][27][28][29][30]. Indeed, commercial companies such as Luxtera (Cisco), Intel, Elenion and Global Foundries are already offering commercial Silicon Photonic integrated circuits[31]. Programmable 'FPGA style' integrated photonics[32] is also becoming available, giving the testing and development flexibility that is on offer in electronics.

1.1.1 Additional Silicon Photonics Applications

The need for Silicon Photonics is not simply due to the potential for cheaper data and hence bigger profits. It is also increasingly becoming a vital part of many

different sectors. This section will outline some of the key ways that Silicon Photonics is being adopted in a wide variety of industries and applications.

- **LIDAR**
  Lidar (light ranging and detection or light-radar) has been proposed by many as a means of short-range sensing and imaging. This is being most heavily employed by companies seeking to build automated self-driving vehicles. Silicon Photonics is being extensively researched as means of providing the low-cost integrated lidar capability needed for these autonomous vehicles[33]–[35].
- **Microwave photonics, MWP**
  Silicon Photonics has also been researched as a way of creating planar microwave devices, typically used for mobile communications. These could dramatically reduce the size and complexity of current microwave infrastructure supporting cellular networks[36], [37],[38].
- **Biochemical sensing**
  There has been a heavy amount of interest in Silicon Photonics as a way of providing biological and chemical sensing. This could reduce the cost and time penalty for testing diseases in humans or animals; or testing for harmful chemicals in pipelines for instance[39]–[43].
- **Gyroscopes**
  Silicon Photonics has been demonstrated as a platform for developing very compact optical gyroscopes for extremely low-powered flight[44], [45]. Many modern aircraft use fibre-based gyros for navigation but these are impractical for drones and other low-power aircraft which could instead use a lighter and lower powered silicon chip instead.
- **Wireless communications**
  There has been a number of demonstrations of the use of Silicon Photonics to implement traditional wireless communication networks[46]–[49].
- **Non-linear optics**
  Silicon photonics has also been used as means of harnessing the properties of non-linear optics[50]–[52]. These could open up a wide variety of new areas of research and development such as frequency doubling devices.
- **Coherent transceivers**
  Many people have demonstrated coherent transceivers in silicon, paving the way for high-bandwidth all-optical communications[53]–[56].

1.2 Need for Cost Effective Packaging Solutions

One of the biggest driving forces behind the need for packaging is that of cost. Indeed, one of the basic rationales for photonics based in silicon is cost as well. Somewhat paradoxically, despite cost being so often claimed to be the reason for pushing Silicon Photonics and hence a means of supporting cheaper datacoms systems, it seems there has been a curiously small amount of attention to economising Silicon Photonics production itself. To highlight this, Figure 1.3 shows a pie-chart adapted from Barwicz et al. [57].

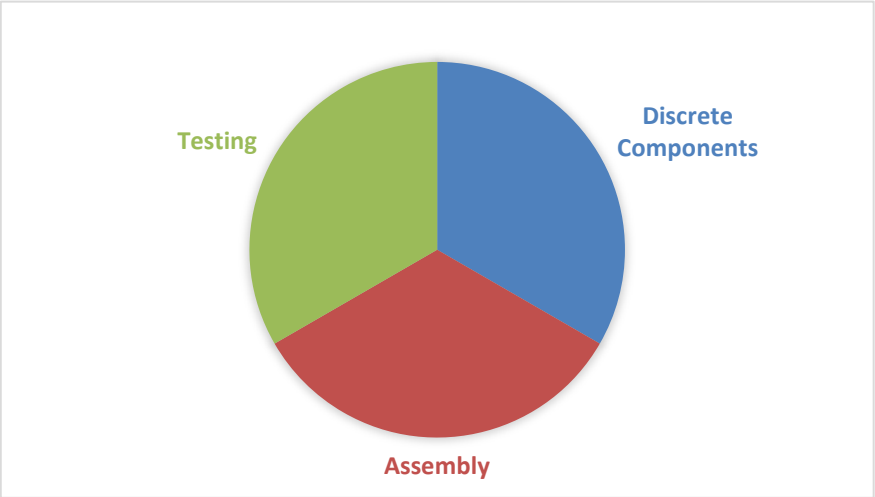

Figure 1.3 Pie-chart showing the distribution of costs associated with traditional optics devices, modified from[62]

This pie-chart shows the relationship between the different costs associated with developing traditional optical devices. As can be seen there is an even split between the various aspects of testing, assembly and the discrete components respectively. Since the costs are spread evenly between each sector, there aren't any specific bottlenecks to be concerned with, nor are there any obvious areas which should receive further research or investment to be optimized for cost.

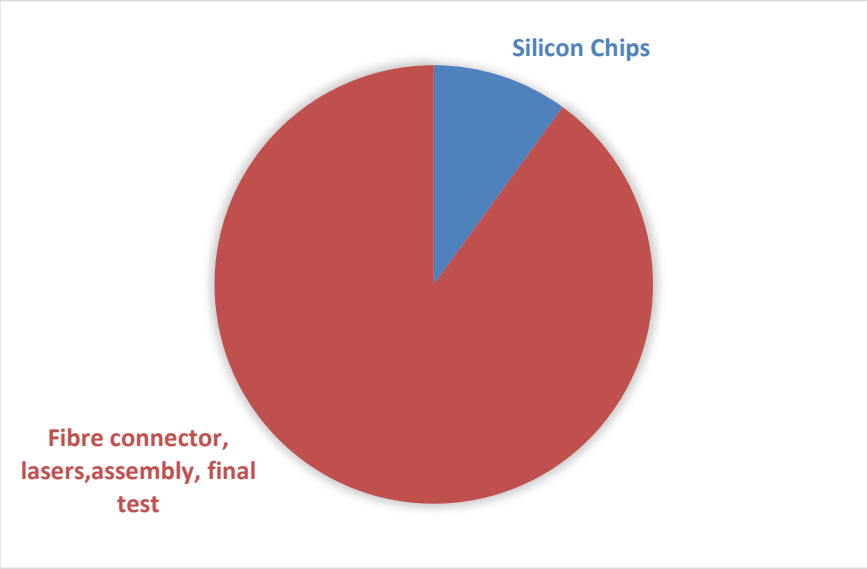

Figure 1.4 Pie Chart showing the distribution of costs with developing silicon photonic devices, modified from [62]

However, when a similar analysis is made for Silicon Photonics, a rather different assessment is to be made. This can be seen from Figure 1.4, also adapted from Barwicz et al. In the case of Silicon Photonics, the pie is divided very differently, with up to 90% of the cost coming from packaging and assembly. This phenomena has been reported by others too – *"large cost and yield loss emerges from the photonic packaging process"*[1]. The cost of silicon and CMOS fabrication is, relatively speaking, very cheap. The cost of everything between the foundry and the consumer however is higher. One wonders why despite taking up to 90% of the cost of production, packaging has received so little attention at an academic level. It appears very much as though the argument for an economical approach has been applied to photonics in general and hence spawned the birth of Silicon Photonics. Yet ironically the same incentives have yet to be applied to Silicon Photonics itself as highlighted by the figures above. Even very basic calculations show that if the costs associated with packaging and assembly of Silicon Photonics can be reduced even by a modest amount, this would leave a huge impression in the total cost of developing silicon photonic products. It therefore becomes a crucial 'battle' across the landscape of Silicon Photonics, where victory could see a swing in the overall 'war'.

Lockwood and Pavesi[58], give some basic rules that must be followed for a monolithic integration of photonics and CMOS to be possible.

Consequently, in the final part of this section, a short excerpt from Lockwood and Pavesi is reproduced. This excerpt sets out some design rules that will help to ensure that Silicon Photonics which is CMOS compatible can be produced for an economic price.

"…

1) *Waveguides should be constructed on silicon substrates or be constructed together with silicon electronic devices.*
2) *Waveguide fabrication processes should not damage electronic devices.*
3) *Waveguides should not be damaged by the processes for fabricating electronic devices.*
4) *Waveguide materials must not be hazardous to silicon electronics*
5) *Geometrical criteria, such as the layout of photonic devices, should not interfere with the electronic circuit layout."* [58]

1.3 Goals and Motivation of this Project

At the start of this project, and still almost universally, most fabricated silicon photonic integrated circuits (PICs) use(d) active alignment for packaging. This is where light is coupled into a circuit and simultaneously measured at the output. Then by using some kind of feedback control system the position of the input and output fibres can be adjusted in order to achieve efficient coupling into the PIC, whether by a grating coupling or butt coupling. Figure 1.5 shows a feedback diagram illustrating this point.

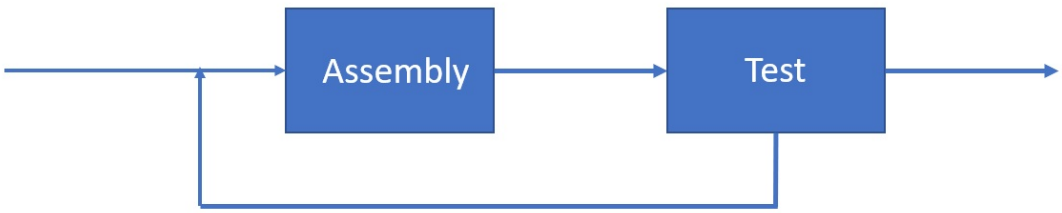

Figure 1.5 Diagram showing the feedback loop needed for active alignment

This then has a number of drawbacks, most obviously that the cost of production will increase as will the time required to produce packaged devices. The volume that can be fabricated will be negatively affected and so for Silicon Photonics to be used widely and in a mainstream, commercial way it is vital that a simpler, low-cost alternative is found to overcome the "packaging bottleneck"[59].

A goal for Silicon Photonics is to produce a method for passively aligning an input/output fibre to a PIC. There would be no need for any feedback loops or extra control systems, thus lowering the overheads and complexity associated with current methods of assembly and alignment. This project seeks to use and build upon ideas for proposed passive alignment towards achieving the desired result cheaply, simply and effectively.

This project itself has a few key aims. These were to help take an initial patented idea and to realise it experimentally. That is to develop the necessary processes and methods that can take the concept from paper into the real world. This involves developing new techniques for design, simulation, fabrication and assembly in order to create an economic and sustainable way of passively aligning optical fibres to silicon chips, bringing together the long-haul and short-haul optical networks that are already established. This involved developing a method that is robust and repeatable and hence would be viable in a mass-market scenario. In addition, the techniques developed need to be streamlined such that they can be reproduced with as little time and effort as possible. This project seeks to do for integrated Silicon Photonics what packaging has achieved in electrical CMOS circuits. This end-goal would be to have a working demonstrator that shows and exhibits the principle of passive alignment.

1.4 Thesis Layout

This short section will outline each of the subsequent chapters in this thesis. Chapter 2 will go through the preceding work that has been established to make this project possible. This includes the basic scientific fundamentals and a review of previous packaging work in the literature. Chapter 3 discusses the design principals that were

used and the key ideas that were employed to refine the design and simulations. Chapter 4 details the evolution of the masks that were used in fabrication. Chapter 5 describes the laboratory equipment that was used for characterisation. Chapter 6 details the fabrication process that was used and the innovation that was made for each step of fabrication to make it possible. Chapter 7 goes through the assembly process, highlighting a number of novel techniques that were developed for assembly. This chapter also give the results from practical work done. The final chapter is a summary and conclusion for the project.

Chapter 2 Background Theory and Review of Literature

2.1 Introduction

This chapter will discuss the key background that must necessarily be appreciated in order to make sense of the work done for this project. This is by no means an exhaustive assessment; however, it will cover all the principles that are required to set this work in its proper scientific context.

This section covers both the scientific fundamentals that are being built upon and also the recent developments that have been made and published in the field. This chapter will then frame the rest of the work discussed in the project.

2.2 Basic Scientific Fundamentals

2.2.1 Ray Propagation in Waveguides

This section outlines some of the scientific fundamentals that are key to understanding this project as well as Silicon Photonics in general. For a long time, it has been known that electromagnetic radiation can be confined within a material that can then be used to transmit a pulse or continuous wave over a distance through space. This was first demonstrated experimentally in 1841 by Colladon[60] and also in 1842 by Babinet[60], done using a stream of water into which a beam of visible light was coupled. In 1869 Tyndall explained this effect in terms of reflection and

refraction[61], describing the phenomena now commonly referred to as *total internal reflection*[62] or TIR as it is often abbreviated. Whilst the use of a stream is a far from practical way of transmitting light over a given distance, the principal can be applied to other materials just as easily, as will now be explained

An intermediate overview of the scientific principles of total internal reflection and guided modes shall now be discussed, showing the justification for the project as well as giving the reader an insight into the physical behaviour of light that is necessary to design and create photonic devices.

An electromagnetic wave will propagate in a straight line through a material (or medium) at a constant velocity as described in Eq. 2.1, where v is the velocity of the wave, c is the speed of light and n is a dimensionless property known as the refractive index which is specific to the material of the medium in which the electromagnetic wave is propagating and is also specific to the wavelength of the wave.

$$\boldsymbol{v = \frac{c}{n}} \qquad \text{Eq. 2.1}$$

From this equation it can be inferred that a particular wave will travel at a different velocity in materials with different refractive indices, for a given wavelength. A vacuum (i.e. no material) has a refractive index defined as unity or numerically equal to 1. Therefore, an electromagnetic wave will travel at a velocity equal to c as can be seen in Eq. 2.1 which is a value of 299792458 ms^{-1} to nine significant figures, often acceptably approximated to $3x10^8$ ms^{-1} or roughly 1000 million km/h. Every physical material in bulk quantities will have a refractive index higher than this due the presence of matter which is absent in a vacuum. Table 2.1[63] gives some common examples of refractive indices as well as showing how it can vary due to factors such as physical state and concentration when in solution. It also shows how different materials can have a near identical refractive index and for many photonic applications can be considered, at least optically, as the same.

Table 2.1 Comparison of different refractive indices [63]

Material	Index
Vacuum	1.00000
Air at STP	1.00029
Ice	1.31
Water at 20 C	1.33
Acetone	1.36
Ethyl alcohol	1.36
Sugar solution (30%)	1.38
Sugar solution (80%)	1.49
Sodium chloride	1.54
Polystyrene	1.55-1.59
Sapphire	1.77
Diamond	2.417

To have an understanding of how light can propagate we can begin with the *ray theory* as a model. Whilst other models of light propagation certainly exist such as particle or wave models, ray theory gives a more intuitive understanding of propagation. The ray model provides a reasonable, albeit not exhaustive, understanding of the propagation of light within different media which have differing refractive indices. It is also useful to describe and explain coupling and loss mechanisms which are present when a ray passes from one media to another, such as from a glass fibre into air.

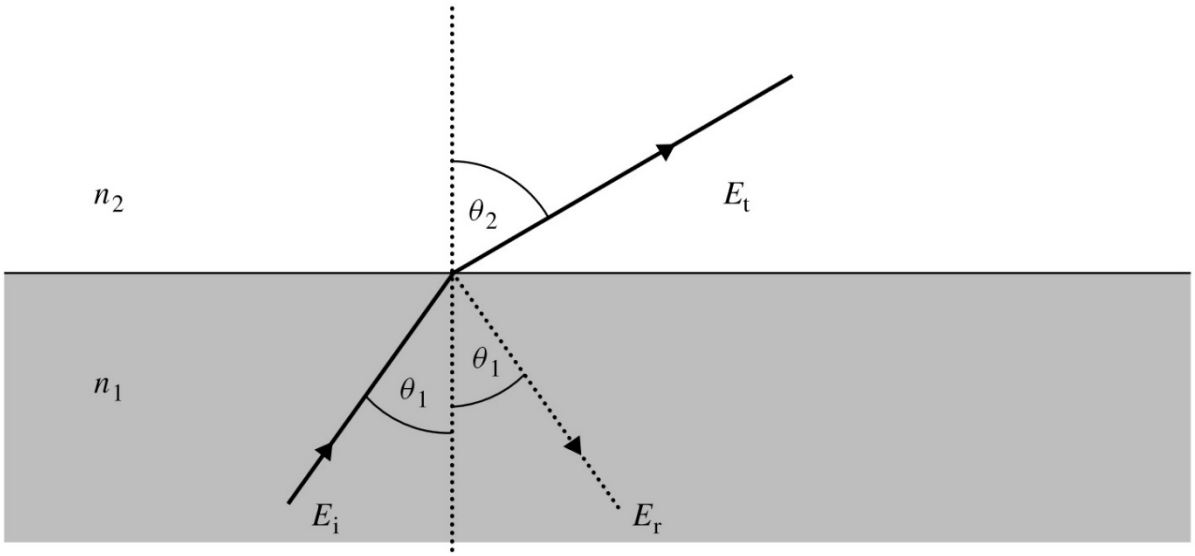

Figure 2.1 A diagram of a refracted electromagnetic wave [64]

The ray model of light can be seen depicted in Figure 2.1 [64] where n_1 and n_2 are the refractive indices of the two different media, E_i, E_r and E_t are the incident, reflected and transmitted rays respectively. θ_1 and θ_2 are the angles of the

corresponding rays to the normal. Whilst the reflected ray E_r is at the same angle of θ_1 as is the incident ray E_i, E_t is propagating at a different angle. Optically speaking, only the refractive index of each medium is different and so this must factor into the change of the direction which the propagating ray experiences.

The governing physical principles of the conservation of energy and the conservation of momentum also apply in this scenario. Since the energy and momentum must always be conserved when the phase velocity of a ray is changed, direction of propagation is hence affected in order for conservation to hold true, whilst the frequency will remain unchanged.

The relationship between the angle of incidence and the refractive index is described by 'Snell's Law' named after the 17th century scientist Willebrord Snellius, though it had been observed prior to Snell. The effect of a ray of light changing direction due to a change in refractive index is aptly called '*refraction'*. Snell's law is mathematically described according to Eq. 2.2 where it can be seen that the ratio between the different refractive indices is equal to the inverse ratio of the sines of the corresponding angles of propagation from the normal.

$$\frac{\boldsymbol{n_1}}{\boldsymbol{n_2}} = \frac{\boldsymbol{\sin\theta_2}}{\boldsymbol{\sin\theta_1}} \qquad \text{Eq. 2.2}$$

From this relationship a ray can be guided to a specific angle by engineering the ratios of the refractive indices of the two materials in which it propagates. This forms the basic premise of all photonics and guided optics since using this simple relationship fosters the possibility of creating optical networks and pathways for photonic transmission.

This theory can also be extended however to allow for optical confinement, where a ray is confined within only one medium and then able to propagate. In order for a ray to be guided within a single medium then the angle of transmittance needs to be greater than 90° from the normal such that it will be confined. Eq. 2.3 shows this quantitatively.

$$\boldsymbol{n_1 \sin \theta_1 = n_2 \sin 90^\circ = n_2} \qquad \text{Eq. 2.3}$$

Therefore, for any two given materials there will be a '*critical angle*' over which the incident wave will be totally reflected within the medium and hence will solely propagate onwards back into the original medium along the line of refraction. None of the incident light will be transmitted through to the surrounding lower index cladding. This critical angle θ_c is represented mathematically in Eq. 2.4

$$\boldsymbol{sin\theta_c = \frac{n_2}{n_1}} \qquad \text{Eq. 2.4}$$

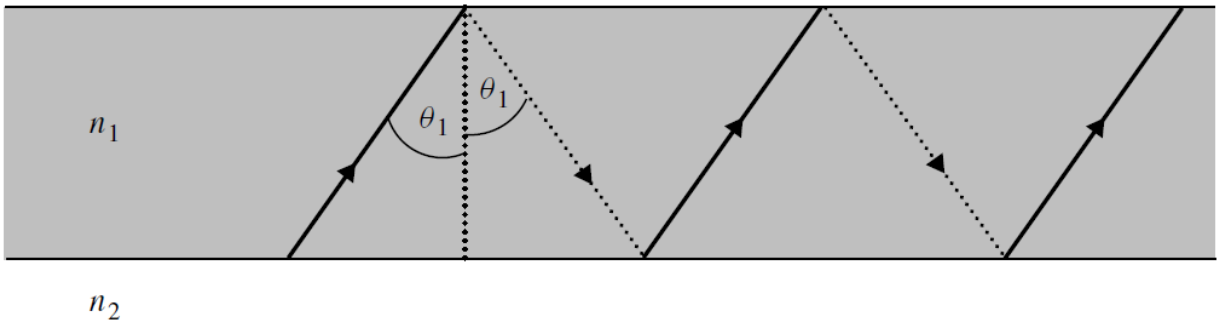

Figure 2.2 Diagram of Total Internal Reflection depicted using the ray model[64]

When this medium is surrounded on all sides then the wave will be able to propagate down the length of the medium and if this material is then extended to a specific length then it is possible for waves to be confined and guided to a desired location in space.

This effect is what is has already been referred to as total internal reflection (TIR) and is illustrated in Figure 2.2[64]. Here the ray of light is 'zig-zagging' up and down across the length of the medium with refractive index of n_1 since it is surrounded by the medium with index n_2 which is depicted as being white in the image. This diagram shows the principle well in two dimensions, but the very same principle can easily and intuitively be extended to three dimensions though this is more difficult and cumbersome to represent adequately.

TIR has been achieved with silica fibres which typically have a refractive index of 1.47[65]. This was first demonstrated for practical communications purposes by Charles Kao in 1965 for which he received the 2007 Nobel prize[10].

Over the last five decades the whole field of the study of fibre optics has arisen out of this technology that is now at work within the modern telecommunications networks which have now become the heartbeat of modern society, circulating media, knowledge and entertainment to the masses.

One final aspect of ray theory will be discussed before moving on- *numerical aperture* or NA. This is important when considering the situation when a guided ray of light is ejected or launched into waveguides such as an optical fibre. This is simply an application of Snell's Law defined as above. Figure 2.3[64] is a diagram illustrating just such a situation.

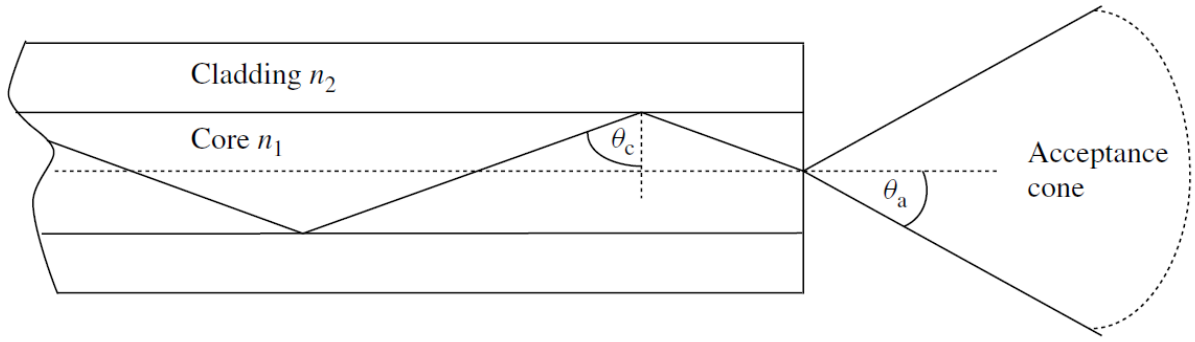

Figure 2.3 Diagram showing the acceptance cone [64]

Figure 2.3 shows a ray propagating down the length of the waveguide, in this case an optical fibre, at the critical angle, θ_c. θ_a is the angle of acceptance which in the case for a symmetrical and circular optical fibre will subtend to form a region called the acceptance cone. The transmitted ray can only be found within the acceptance cone, and conversely only an incident beam can be launched into the optical fibre from within the acceptance cone due to reciprocity. Rays that do not meet this condition will not be totally internally reflected rendering them unsuitable. By applying Snell's Law mathematically, it is possible to derive an expression which includes the acceptance angle (θ_a) which is shown in Eq. 2.5.

$$n_a \sin\theta_a = n_1 \sin[90° - \theta_c]$$

$$= \boldsymbol{n_1 \cos\theta_c} \qquad \text{Eq. 2.5}$$

Previously it has already been shown in Eq. 2.4 that $sin\theta_c = \frac{n_2}{n_1}$ and so it is possible to substitute this expression into Eq. 2.5 with the use of a commonly known trigonometric identity; $sin^2\theta + cos^2\theta = 1$.

Hence, by substitution it can be shown:

$$\boldsymbol{n_a \sin\theta_a = n_1\sqrt{1 - sin^2\theta_c}}$$

$$= \boldsymbol{n_1\sqrt{1 - \frac{(n_2)^2}{(n_1)^2}}}$$

$$= \boldsymbol{n_1\sqrt{(n_1)^2 - (n_2)^2}} \qquad \text{Eq. 2.6}$$

This gives the mathematical expression, which is defined as the numerical aperture, NA, which is summarised below in Eq. 2.7

$$\boldsymbol{NA = n_a \sin\theta_a = \sqrt{(n_1)^2 - (n_2)^2}} \qquad \text{Eq. 2.7}$$

NA is usually defined in air but does not have to be, especially when light is being launched into a fibre from a special epoxy or polymer waveguide that are often used in packaging. NA can also be used to assess the coupling efficiency of a given coupling scheme.

2.2.2 Waveguides and guided modes

TIR can also be used in a planar structure as shown in Figure 2.4 and is known as a waveguide. Waveguides have been demonstrated in several different materials for different applications, but all work on the same principal of guided

electromagnetic waves – or often simply called '***light***' even at wavelengths into the non-visible infrared region of the spectrum.

At this stage it is necessary to derive from first principles, the theory of waveguides. Optical waveguides are a necessary component of any PIC, in an analogous way to a metallic wire in an electrical circuit. However, despite the similarities, the physical phenomena are quite different. Whereas a metallic wire will behave according to Ohm's law, light in a waveguide will propagate in discrete modes[66]. The energy confined within a waveguide cannot take any value in a spectrum, but rather will come in specific modes of propagation.

In order to derive the theory of waveguides it is convenient to consider the following situation. Figure 2.4 shows three planes, which have been cropped for illustration purposes. In this scenario $n_2 > n_3 > n_1$., though the exact values need not concern us. Each n corresponds to the refractive index of that material. This configuration is very similar to a silicon-on-insulator waveguide with air cladding at the top.

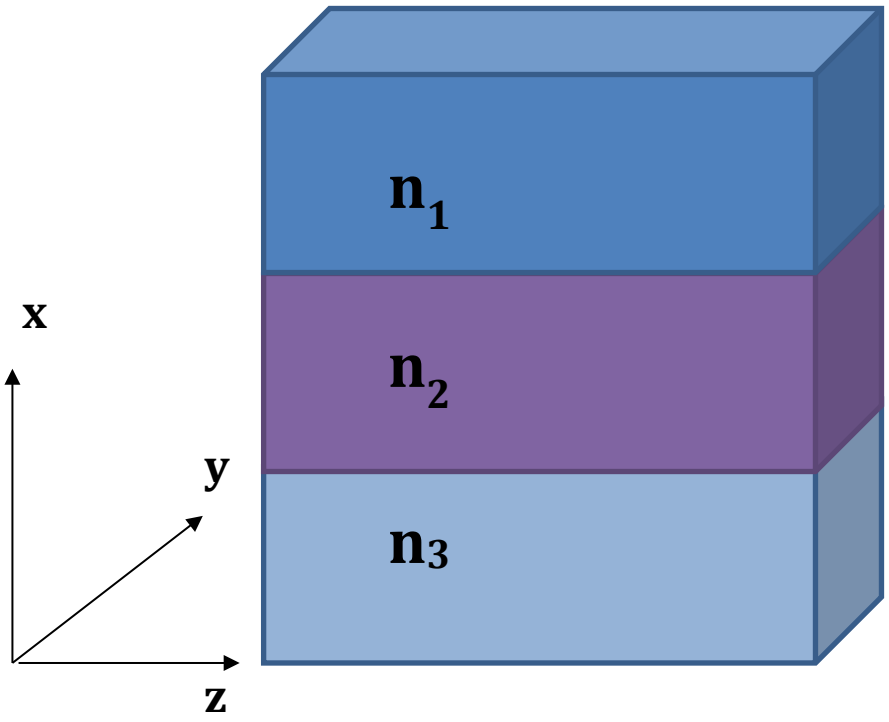

Figure 2.4 Schematic diagram of a planar waveguide

In this basic derivation, each plane is considered to extend infinitely in the y and z direction, though all layers are a non-infinite value in the x direction. Usually, convention will have the light propagating in the z direction.

In order to compute the mode within the waveguide we must consider the energy confined within the middle layer since the theory that has already been explained tells us that the optical energy will be confined to the higher index medium. This will involve finding a spatially confined solution to Maxwell's wave equation[66], which is shown below.

$$\nabla^2 \boldsymbol{E}(z,y,z,t) = \left(\frac{n^2}{c^2}\right)\frac{\partial^2 \boldsymbol{E}(x,y,z,t)}{\partial^2 t^2} \qquad Eq.2.8$$

In this version of the equation, E is the electric field in vector form (indicated by the bold font), n is the refractive index and c is of course the speed of light. The electromagnetic wave will be described by the expression:

$$\boldsymbol{E}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z},t) = \boldsymbol{E}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z})e^{i\omega t} \qquad Eq.2.9$$

Where the electric field has a frequency of ω. Since we know the solution of the electric field will take this form, it can be substituted back into the previous equation and with a little manipulation the following is obtained:

$$\nabla^2 \boldsymbol{E}(\boldsymbol{r}) + k^2 n^2 \boldsymbol{E}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}) = 0 \qquad Eq.2.10$$

In this expression, the wavevector k = ω/c makes the expression easier to understand. Since we are considering the situation when the wave is propagating in the z direction, there is no confinement along that axis. Hence the electric field becomes thus: $\boldsymbol{E}(\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}) = \boldsymbol{E}(x,y)e^{-i\beta z}$. There is a propagation constant β which governs the propagating mode; it is a measure of the change of amplitude and phase in a given direction per unit length. Again, by substituting this expression into the previous equation we obtain the following:

$$\frac{\partial^2 \boldsymbol{E}(x,y)}{\partial^2 x^2} + \frac{\partial^2 \boldsymbol{E}(x,y)}{\partial^2 y^2} + (k^2 n^2 - \beta^2)\boldsymbol{E}(x,y) = 0 \qquad Eq.2.11$$

This partial differential equation will have solution that is either a real exponential or a complex one which will be oscillatory. The equation now takes the form of an elliptic partial differential equation. The crucial part of this equation then becomes $k^2n^2 - \beta^2$, since the form of the solution depends on whether this expression is greater than or less than zero. Recalling that $n_2 > n_1, n_3$, we can deduce that if the value of β is between kn_2 and kn_3 then the solution will be a mode that is confined within the waveguide[66]. However, if $\beta < n_3$ the radiation will begin to 'leak' out of the waveguide, since there is no confinement at the boundary and the wave will propagate out. This is a radiative mode which is not guided.

2.2.3 Discrete Modes of Propagation.

In order to understand how the mode will propagate through the waveguide we need to obtain the form of the solutions for β. For a full solution of β see [67]. The important result is that the solution(s) are discrete, arising from the fact that it is the solution of a periodic tangent function. This means that though β can take multiple values which are less than kn_3 it can only take certain discrete values. Therefore, only certain discrete modes of propagation are allowed. The exact value will depend on what the refractive indices, waveguide thickness and frequency are. This theory has been deduced considering a 2-dimensional waveguide, however the same process applies when considering more complex waveguide topologies, the total waveguide mode solutions will be a superposition of modes in each direction of the waveguide cross section.

2.2.4 Waveguide mode Cut-off

Since the propagation constant depends up on both wavelength and waveguide thickness, one can define a cut-off wavelength or cut-off thickness. This means that for a given waveguide, if one increases the wavelength, eventually it will no longer be able to guide any modes – there will be no real solution to the wave equation. In the same way, if a given wavelength is selected, there is a minimum thickness that can support the mode.

2.2.5 Waveguide theory conclusion

This theory underpins the work that was carried out in this project. It was used to design the waveguides that were fabricated. Rather than solving these equations analytically, numerical computations were used to design the waveguides, but it is necessary for the photonic engineer to know these equations and results as well as having an intuitive understanding in order to begin to design any waveguide circuit.

Once waveguides have been designed, they can be used to direct light around a photonic circuit on a PIC. This way they can route signals between different devices such as modulators or filters and additionally, the signal can be coupled out of the waveguide and into an optical fibre allowing the signal to be used for either long or short haul communications.

2.2.6 Types of Waveguide

There are many different examples of waveguides that have been fabricated such as strip, rib, slot, suspended and photonic crystal [68]–[71], however for most purposes the rib waveguide is to be preferred since in exhibits lower loss [64] and can be easily fabricated alongside other devices such as free carrier depletion

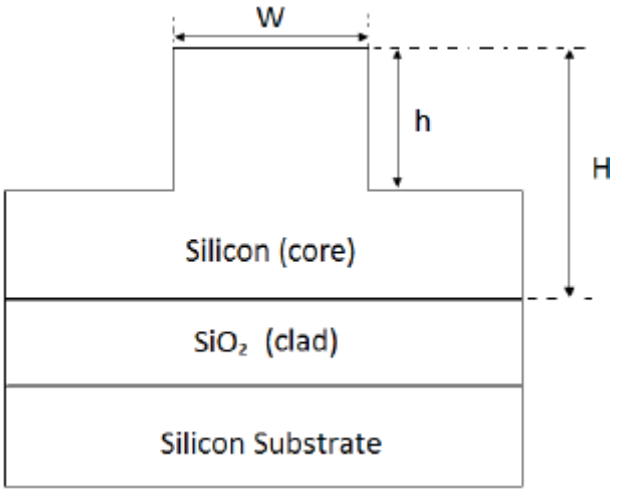

Figure 2.5 Cross-sectional schematic of SOI rib waveguide [73]

modulators[72] but also simpler passive devices such as grating couples (explained in detail later). Figure 2.5 shows a cross-sectional schematic of a typical rib waveguide [73].

This is an SOI based rib waveguide, where only a section of the silicon is etched. Single mode rib waveguides are possible according to design rules that were originally formulated in pioneering work by Soref et al.[74]. For this project they are used due to the reasons mentioned above.

2.2.7 Choice of Silicon

Silicon waveguides are capable of confining and transmitting light with relatively low loss, typically lower than 5 dB/cm [58], [75] and as low as 0.1dB/cm for some designs[76]. Silicon has very high transparency at normal communication wavelengths around 1550nm making it a material very well suited for fabricating waveguides designed to operate in this region commonly referred to as the C band in optical communications which is defined as the region of 1530-1570 nm[77]. It is also popular for silica fibre communications due to silica's transparency at these wavelengths as well as the availability of erbium fibre doped amplifiers in this region[12]–[14].

Silicon is also benefited by having a high refractive index of approximately 3.48 [78], at a wavelength of 1550nm and at a temperature of 295K (note: refractive index is also dependent on temperature). This makes it possible to confine light in a silicon waveguide when it is implemented with common cladding materials such as silicon dioxide (SiO_2 or simply 'oxide') and air with refractive indices of 1.44 and 1.00 respectively at similar wavelengths and temperatures. Whilst silicon is a useful material for fabricating waveguides, coupling light into such a waveguide is more challenging. If silicon is to be used as a waveguide material it must be justifiable. Since the world's long-haul networks predominantly use silica glass waveguides (optical fibres) it seems that the best candidate for planar waveguides would be the same material. Silica waveguides would have comparable dimensions to an optical fibre and so wouldn't require much in the way of special coupling. In fact, silica waveguides are a popular choice and much research has been done into them[79]. There are of course also drawbacks to using silica. Firstly as Kawachi rightly points out in [79], Silica waveguides can be used in LSI chips with little difficulty. However, when considering VLSI waveguides of the size of silica, the waveguides become very large, circa tens of microns in width whereas a silicon waveguide is sub-micron. Given that the current state-of-the-art transistors can be as small as 5nm[80], having

waveguides three orders of magnitude larger obviously can be problematic since on an optoelectronic chip, both transistors and waveguides could be present and connected together. Any gains in terms of footprint that may have been made by shrinking to 5nm will then be lost with waveguides of such a large size. The other consideration that must be made is that of CMOS compatibility. Silica waveguides are typically made through flame-hydrolysis. Whereas a silicon waveguide can, in principle, be fabricated on-chip at the same time as any electronic circuitry, but also there are far more CMOS standard fabrication facilities which are already prepared for silicon PIC fabrication.

2.3 Fibre to PIC Coupling Methods

2.3.1 Coupling Challenges

In order for Silicon PICs to be useful it must be possible to interface them with optical fibres as seamlessly as possible. This allows them to be connected to long-haul fibre networks and then also to be connected to other silicon PICs in a network. Therefore, light must be coupled between the PIC and an optical fibre with as little loss as possible and vice versa; light coupled in a fibre must be coupled into a PIC with minimal loss. Whilst both fibres and waveguides on PICs are guided modes there is a big difference between the respective guided modes.

The most typical optical fibre is the Corning SMF-28[65]. These fibres are almost ubiquitous within photonics and so can be used as a good bench-mark example to highlight the challenge in coupling from a fibre into a PIC waveguide. The first important characteristic to realise about the SMF (single-mode fibre) is the mode-field diameter or MFD. The MFD is the diameter of the propagating mode within the optical fibre. In the case of the SMF-28 at 1550nm this is 10.4±0.5μm[65], which is very typical of all SMF's. This mode will also be circularly symmetrical since the geometry of the optical fibre is circular. However, the waveguide dimensions of a typical silicon PIC are rather different. Usually they will have dimensions of roughly 0.5 x 0.22μm[81] and are usually an oval shape from the rectangular geometry of a typical waveguide, or possibly an irregular shape for rib waveguides when the mode is wider at the bottom than the top. Therefore, when the area that is associated with each MFD is considered it can easily be found that the total area of

the fibre mode is roughly 500 times larger than the mode of a silicon waveguide. Figure 2.6 shows a diagram of what the difference in MFD looks like. This diagram is adapted from the version in Zimmerman et al. [82].

In addition to this, polarization sensitivity must also be accounted for, otherwise polarization dependant losses (PDL) will be incurred when coupling from one into the other. Within an SMF, the polarization of the electric field is usually unknown and always unstable and so matching the polarization with that of the PIC waveguide is difficult. (The transverse-electric TE polarisation is usually used for communications purposes while the transverse-magnetic field is more popular for sensing). Therefore, in order to alleviate the effect of PDL the polarization of the

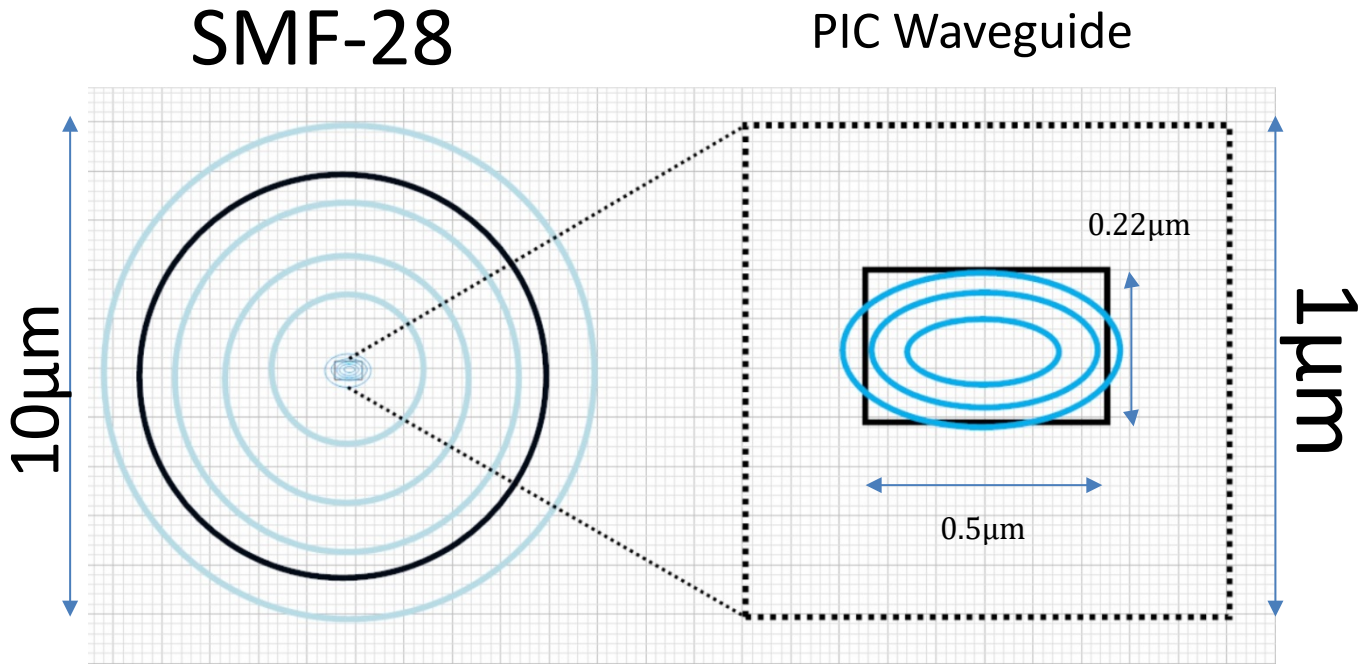

Figure 2.6 Diagram of the MFD difference between an SMF and PIC waveguide

mode must be managed adequately. This can either be done in the fibre or later in on the PIC [83]. Within Silicon Photonics there are three main ways of coupling from fibre to PIC[28]. These three can be described as different 'coupling philosophies' where each approach has his positives and negatives. They are as follows: edge-coupling, vertical-coupling and evanescent-coupling.

Grating couplers were chosen for this project to allow for efficient coupling into silicon waveguides. Here a short discussion on the fundamentals of grating

couplers shall be given. Grating couplers have been used in optics for a number of decades [84], [85], and are commonly used in silicon photonics as a means of coupling light from an optical fibre into a silicon nano-photonic waveguide[81], [86]. Grating couplers are able to accomplish effective light coupling by satisfying the phase-match condition. This is where the z directed propagation constant of the light in the waveguides are matched. In practice this can never be achieved without a grating coupler since a silicon waveguide's propagation constant is greater than the incident propagation constant resolved parallel to the waveguide (see Reed and Knights 2004 [64]for full derivation); usually defined as the z direction. Figure 2.7

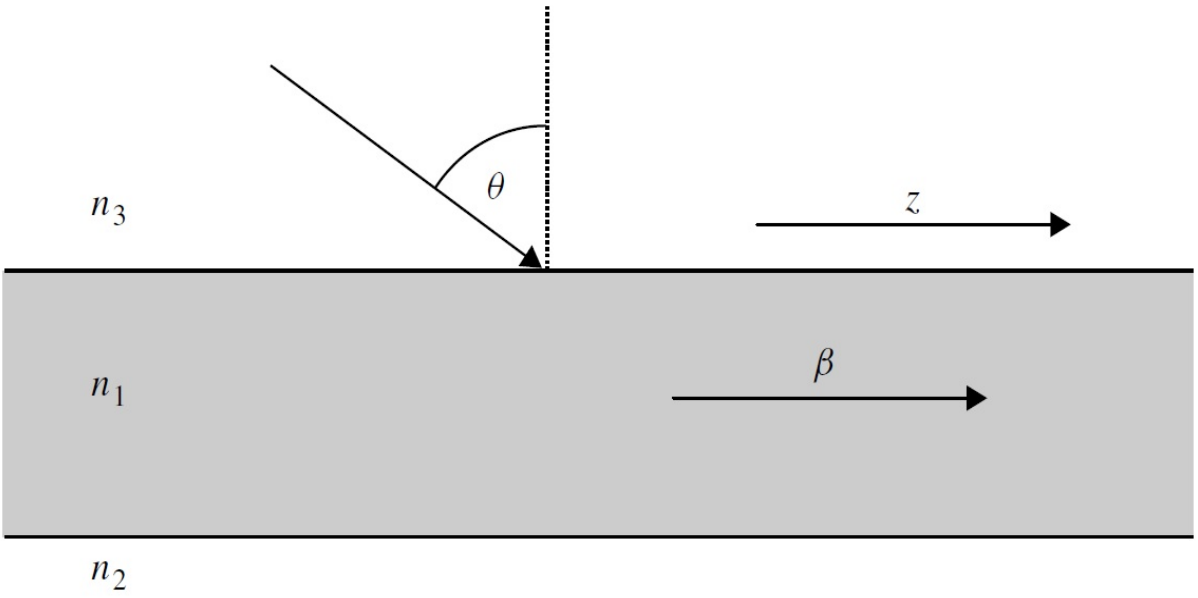

Figure 2.7 A diagram of incident light upon a waveguide [64]

shows a diagram of an incident light beam upon a waveguide[64], with the required propagation constant of $zn_3\sin\theta$.

Despite this problem, a periodic grating structure can be used to allow light to be coupled into the waveguide in a satisfactory manner. With a grating structure etched onto the waveguide the propagation constant becomes a function of the grating period and then it becomes possible to engineer a grating that is capable of coupling light as required.

2.4 Review of Literature

2.4.1 Overview of Silicon Photonic Packaging in the Literature

This section will outline a brief but comprehensive survey of the literature that has been published in recent years and is relevant to this project. It is necessary to do this for two reasons. Firstly, as it will provide inspiration for the work here, set it in context and prevent unnecessary repetition of work carried out previously. Secondly, it is necessary to correctly acknowledge the work that has gone before in order to make this project possible.

Publications in the area of Silicon Photonics packaging have been somewhat sparse compared to other areas in Silicon Photonics in recent years. This is understandable since there is no point considering how to package devices that have not yet been developed! However, now that a wide range of silicon photonic devices and processes have been formalised, packaging is very much a necessity and one can only predict that publications will increase.

The approaches that have been demonstrated in this area up until now have been quite varied, as can be seen from a very cursory reading of the sub-titles in this section. This suggests a couple of things. Firstly, there has been no definitive solution that meets all requirements. Though many different competing methods have been conceived and demonstrated, none has risen above the rest to common acceptance. Secondly and more excitingly, it suggests that there is plenty of scope for innovation and new ideas since no formula has been found, fuelling the basis of the work presented here. The different packaging methods here are arranged in terms of ‘packaging philosophy’, meaning that there are certain principles to packaging that have been adopted by different groups and this has been used to organise and group them in this text. As such not every single example of packaging is included but rather each different general approach or ‘philosophy’ is considered with the most notable work in each being expounded further.

2.4.2 Active, Out-of-Plane

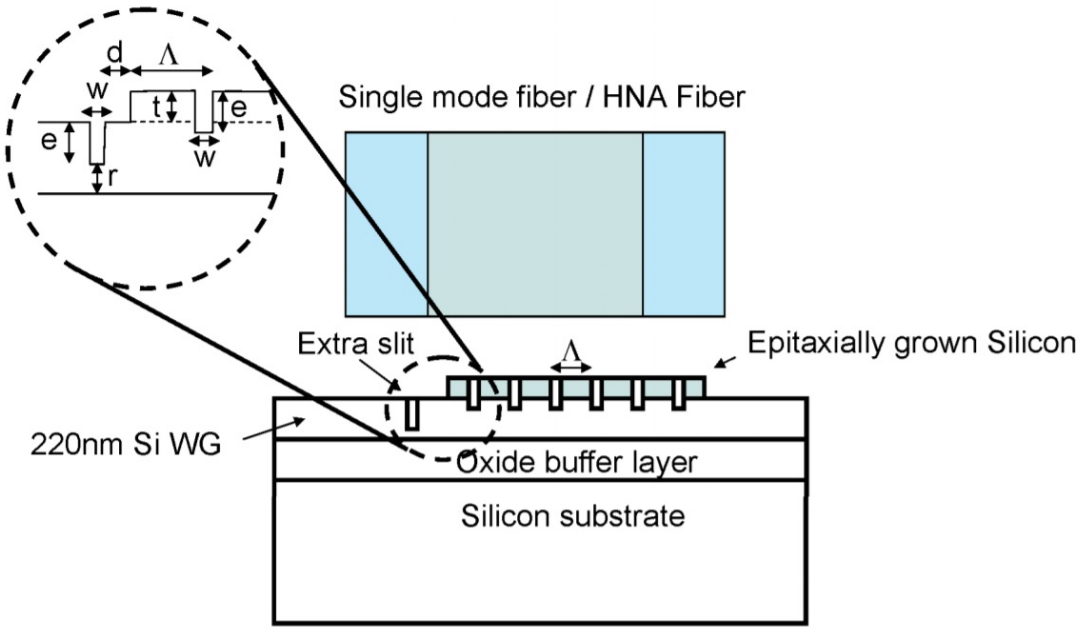

Figure 2.8 Diagram of the perfectly vertical grating coupler by Roelkens et al.[87]

2.4.2.1 Vertical Coupling

The challenge of out of plane coupling can be seen quite clearly by simple inspection. Grating couplers typically require a slight angle of incidence in order to couple light effectively in a desired direction. This poses a problem from the point of view of packaging. A fibre must be positioned exactly in the correct place in all 6 axes typically to a tolerance of ±1μm. In addition to this, the angle of incidence makes the packaged chip bulky and likely fragile, both factors which should be minimised for effective packaging. With this in mind, an ingenious solution to this problem was presented by Roelkens et al. They propose a novel grating design which allows for "perfectly vertical fiber coupling"[87]. This method uses a non-uniform grating to achieve perfect 0° coupling from the normal. Figure 2.8 shows a schematic of the design. An extra slit is etched into the grating to create extra Bragg reflections in the coupler which directs the beam vertically. By having a perfectly vertical coupling mechanism this effectively eliminates 3 of the axes of rotation and so only translational movement need be considered. The authors suggest that this is a method to reduce the cost of packaging, which is true if an active approach is assumed. Actively aligning a fibre to 3 axes of translation alone and not considering the

additional 3 axes of rotation would minimise the time and complexity required for active alignment. This paper reports coupling efficiency of between 50-65% depending on fibre used, a higher coupling efficiency is possible with a narrower lensed fibre and a lower efficiency is possible with a larger conventional SMF. This was achieved using Gent's in-house mode solver CAMFR. When repeated on more advanced commercially available software the author of this thesis noticed slightly higher losses, achieving coupling efficiencies of 35-50%. This approach does still suffer from a few weaknesses, however. Firstly, the approach is still an "out-of-plane" solution which will add bulkiness to the finished package. For fibres this would be problematic, even with the use of a low-profile ferule. The technique however would be very suited to the use of VCSEL lasers as the authors rightly point out in their conclusion. Also, this paper uses epitaxial growth of silicon on an SOI wafer in order to create the thicker silicon layer needed for the grating coupler. This could also be achieved by using a thicker SOI wafer, 400nm SOI wafers are now readily available which were less available at the time of publication (2007). The extra step of epitaxial growth of silicon is not desirable but also not impossible. Even with a thicker SOI wafer used, multiple etch steps would be required in order to produce this specialised grating coupler. However, given that that majority of production cost lies in packaging, increasing fabrication cost and complexity to save packaging costs could be a sensible approach. Ultimately however, this method is not a packaging solution as such, but rather a helpful contribution that could be adopted into a packaging scheme. This technique is only really a partial solution but is included in this review of other methods for completeness. It would still need to be used in conjunction with an alignment technique and most likely an active one. If a passive alignment technique can be developed, then it would be an alternative to the work in this project since the packaging process would be lower-cost and quick anyway, so specially fabricated grating couplers would be a bonus. Alternatively, the special gratings could be incorporated into the work here to produce an excellent passive solution, subject to grating coupler efficiency being improved. The theoretical maximum coupling efficiencies are probably unacceptable by today's standards. All that said, this paper is a helpful contribution to the study of packaging and must be considered within its context

which is now over ten years old. Now that Silicon Photonics has gained much more commercial and industrial support, lower insertion losses are demanded.

2.4.2.2 G-Pack

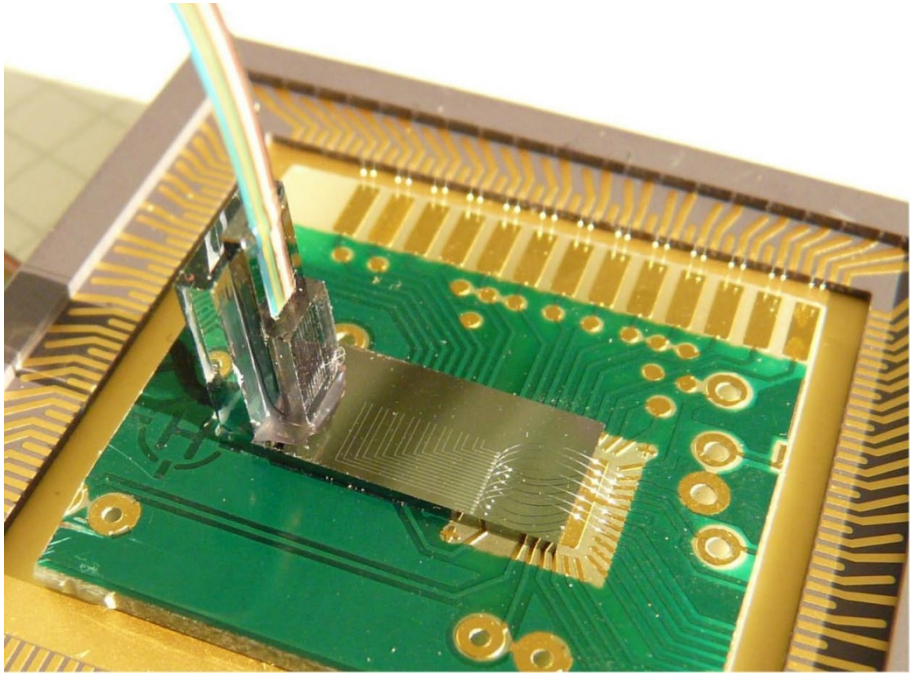

Figure 2.9 Image of G-Pack by Zimmerman et al. [88]

Another alternative out-of-plane solution is the g-Pack as presented by Zimmermann et al. which is depicted in Figures 2.9-2.10[88]–[90].

This involves using an epoxy to adhere the fibre to an optical chip at the desired angle in order to achieve efficient coupling into a grating coupler. The fibres

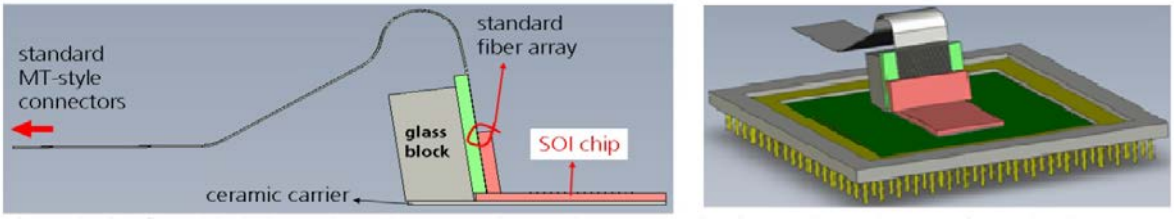

Figure 2.10 Diagram of the g-pack method [88]

are held between two chips, one with v-grooves in order to hold the fibres and the other chip (usually a silica block) is used to hold the fibres into the v-grooves and also give a surface that can be mated to the chip. The end facets are then all polished to give the desired angle before being bonded to the chip. In their standard method the fibre assembly is actually bonded to a ceramic carrier which the optical chip sits on. The grating couplers used for this approach gave a 30% efficiency with 40nm

1dB bandwidth which should be considered relatively poor given the state of the art which is around 83%[91]. There is a 1dB loss penalty for a misalignment of ±2μm[88]. This approach has many benefits – a whole array of fibres can be easily coupled into an optical chip with this method and at a conventional angle which permits this to be used with previously manufactured chips and systems. Also it must be pointed out that this method had been physically demonstrated quite effectively by commercial companies such as Luxtera who have used this method for a quad-channel 40Gb/s transceiver(Figure 2.11)[92]. Ultimately market acceptance does show the usefulness of a design.

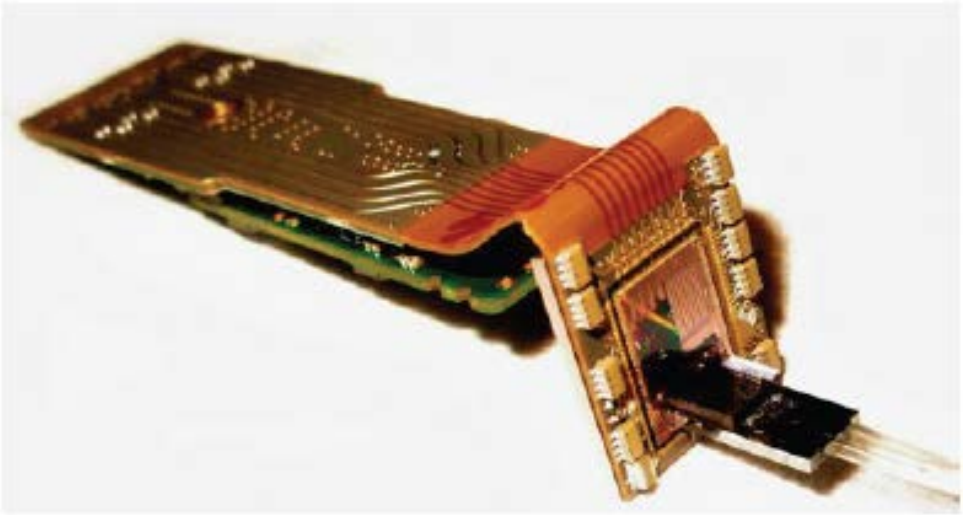

Figure 2.11 Image of the "WDM package" from Luxtera[92]

There are however some points of critique which must be made. Firstly, the out-of-plane nature of this approach will always be limiting. As can be seen from the figures above, the fibres must be located in a block, usually silica, which is then adhered to the optical chip. The extra added bulk here is quite considerable. This is a big price to pay in terms of footprint for what can really only be considered marginal gains in terms of performance. The whole packaging procedure incurs a loss penalty of 2dB which is quite expensive given that it is actively aligned. Furthermore, the active alignment must be considered in many orientations since there are effectively 5 pieces that much be aligned together, the fibres, optical chip, glass block, v-groove chip and the carrier. None of these can be considered passive in alignment and each will impose a slight alignment tolerance of ~0.5dB[89]. Also,

the use of polishing is not a standard procedure and adds to the complexity and cost of design.

Whilst this approach has many commendable qualities, it shows that there are still hurdles which must be overcome before it could be considered on a mass scale. Altogether the g-pack approach is bulky, somewhat lossy and pricey (active alignment) so therefore should be considered as a generous contribution towards an effective packaging solution rather than a real solution in and of itself.

2.4.3 Active, In-plane

2.4.3.1 Overview

This section will outline some of the in-plane solutions that have been suggested. These solutions sidestep the problem outlined above – the bulkiness of an out-of-plane solution. Without reducing the bulkiness of a packaged PIC the fundamental benefits of using a high index low footprint material such as silicon become undermined. To this end, an in-plane solution is preferable.

2.4.3.2 Polished Fibres

Another general approach to packaging has been to use polished fibres. This was elegantly presented by Snyder and O'Brien in 2012[93], though had been

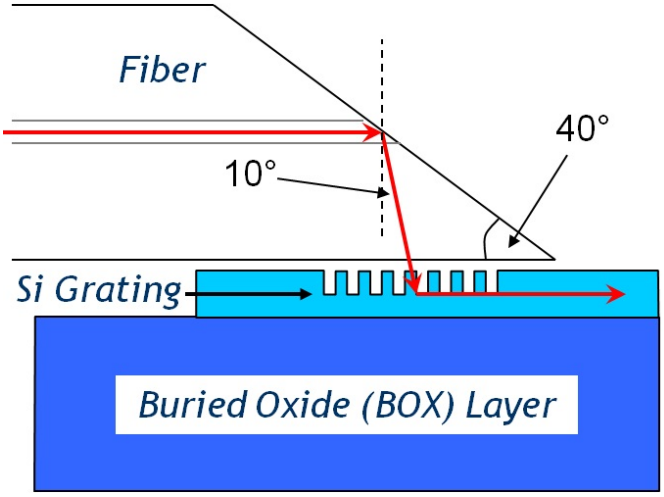

Figure 2.12 Diagram of polished fibre packaging method, taken from [93]

previously suggested by Luxtera[94]. The fibre is polished at a specific angle

as can be seen in Figure 2.12. In this instance the fibre is polished at an angle of 40° to the normal.

This angled facet then gives the guided mode a surface to reflect off by means of TIR. As is shown in this image, this gives the incoming light an incident angle of 10° which is very typical for standard grating couplers. Losses are reported to be less than 1dB with a misalignment less than 5µm, less than 1dB for a 5° roll and a vertical extension of up to 13.5µm with less than 1dB as is shown in Figure 2.13 [95].

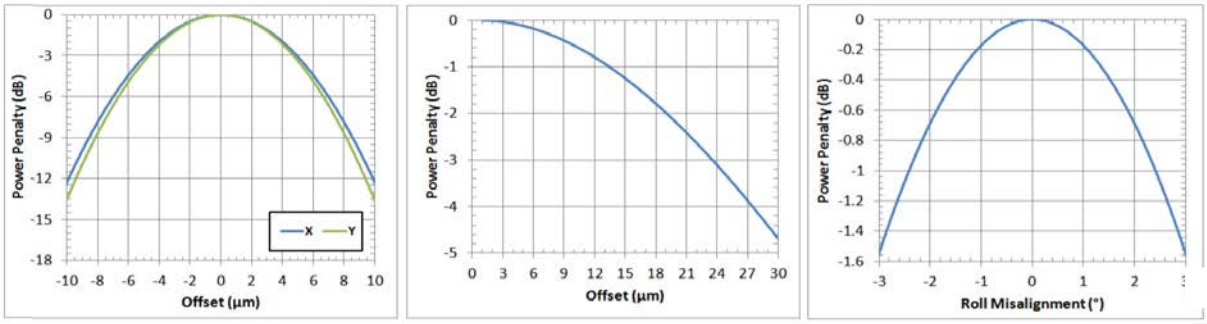

Figure 2.13 Optical Performance of the Polished Fibre packaging method, adapted from [95]

To its credit this method does have many attractive qualities, namely its simplicity and versatility. This method could be implemented with essentially any grating coupler based PIC and the low number of essential parts reduces the complexity of fabrication and assembly. In addition to this, it is a low-profile in-plane method. Very little extra bulk would be added to the packaged chip. A multi-fibre array could be used to allow for coupling multiple channels into a chip with no extra complexity.

However, this approach is not without its limitations. Firstly, polishing fibres is not a standard procedure in standard CMOS fabrication, nor in Silicon Photonic fabrication in general. The question remains as to whether fibres could be polished accurately and repeatably to a specific desired angle for production purposes. Whilst it is shown that this can be achieved in small quantities there is little suggestion that this could be replicated in larger quantities in a cost-effective manner. Also, the durability of such an approach has yet to be determined. A UV-cure epoxy is used to attach the fibres but due to the size and orientation of the fibre, there is very

little mating surface for the adhesive to secure to, raising doubts as to whether optical alignment could be ensured over a longer time period. Of course, a protective outer casing could be applied but this hasn't been explored by the authors. The biggest weakness of this method is the alignment procedure. Whilst the authors claim that it is "*suited to passive alignment*" they provide no method by which this could actually happen. So whilst this technique may well be suited to passive alignment it requires additional features to make passive alignment a possibility. With this in mind it can only really be considered an active approach at present. In order for this to be a passive method of alignment, then straight away some of its key features (versatility and simplicity) are threatened since the additional hardware needed would more than likely impact both of these aspects.

2.4.3.3 In-plane Block

Another approach which deserves mention is that of Galan et al[96]. It follows in very much the same spirit of Zimmerman et al. by using a block to hold an array of fibres. However, they have rightly identified the limitations that arise from an out-of-plane block-based system. They highlight a few advantages with using grating couplers as Zimmerman does, namely a planar fabrication process, wafer level testing and the spatial alignment tolerance; which is relatively good. They then highlight that this will involve devising a new packaging scheme since grating couplers are by default an out-of-plane device. Figure 2.14 shows a schematic of how

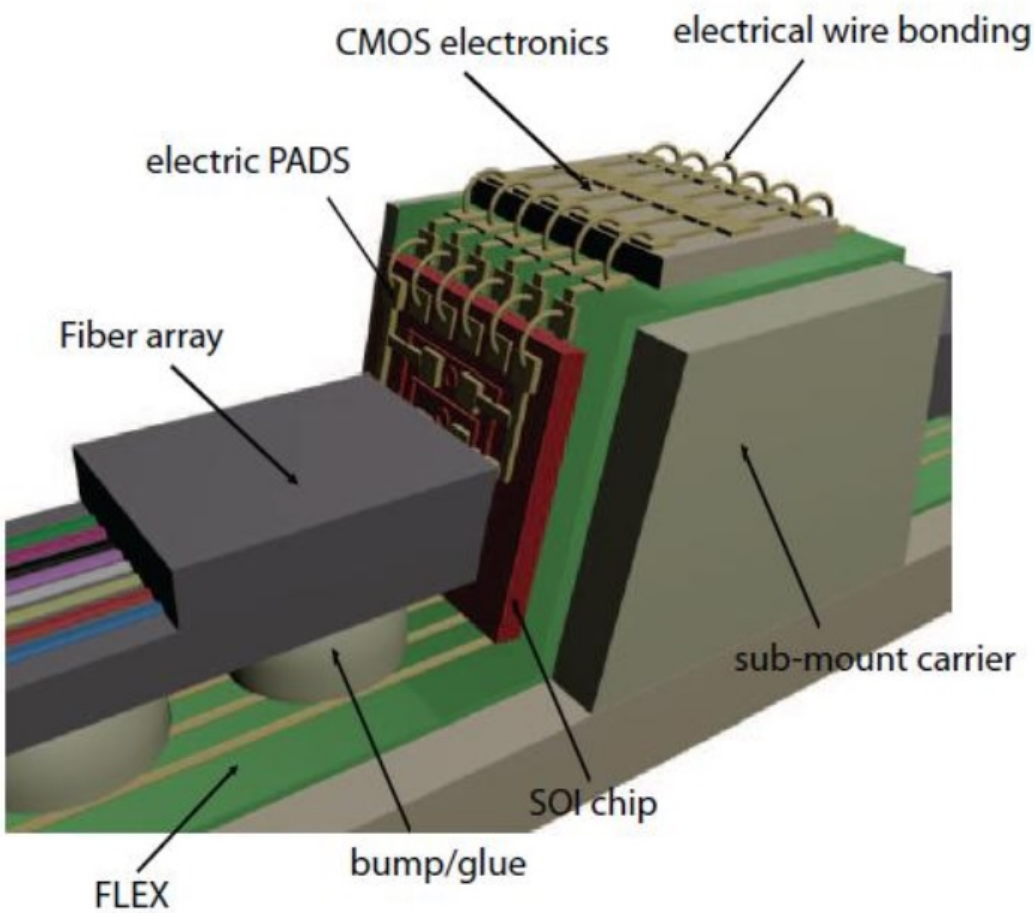

Figure 2.14 Diagram of low-profile block packaging by Galan et al. [96]

this would take place, by using a special carrier and then by gluing the fibre array in place. Electrical connections can then be added by wire bonding to incorporate CMOS chips with the optical chips. Another important feature of this design is the "FLEX" (flexible wiring layer) which is used to allow wire bonding in conjunction with the chip carrier.

This packaging method has demonstrated good results as well. They report that each grating coupler exhibits a coupling efficiency of 24% (6.2dB). These losses are quite high but there is nothing intrinsic to the packaging design which suggests that the grating efficiency couldn't be increased to be closer to current state-of-the-art efficiencies. They use a rather large footprint in their grating design of 12 x 12μm which is presumably to allow for a small amount of misalignment. The 'square' nature of the grating coupler also will help in this regard whereas a conventional one will be more rectangular. They also report misalignment tolerances of being ±2μm for a 1 dB penalty in the x axis and slightly better in the y axis with ±3μm. They then use an index matching epoxy, as is normal in packaging, and report 1dB extra loss as a result of gluing the fibres in place. What isn't discussed in depth is a total loss attributed to the packaging, only individual elements. Whilst it could be assumed that by summing these together one could calculate the total insertion losses, most likely each of these figures is *best case scenario* and there could be other factors not reported. As such it is curious that they neglect to report the total losses attributed. In either case, the total insertion losses will be at least 7.2dB just from the figures that are given. Also, it must be stressed that this is still an active method of alignment. No mention of any form of passive alignment is given in this method and so it must be assumed that the alignment is carried out in an active scheme.

Overall this method does aptly highlight many of the issues that have been already mentioned and does seek to alleviate them, primarily by offering an in-plane solution. It also allows up to 8 optical channels to be packaged in an array alongside CMOS electronics and so could find wide reaching applications. However, the complexity, insertion losses and active alignment procedure do limit its appeal.

2.4.4 In-plane, passive approach

From the discussion that has been made so far, it is clear that ultimately any packaging scheme that is to be adopted would ideally be both passive and in-plane. There have been a number of impressive demonstrations of packaging schemes which apply this philosophy.

2.4.4.1 V-groove Passive Alignment

Galan et al. have also demonstrated a V-groove based coupling method which is compatible with passive alignment[97]. It is interesting to note that whilst this method in the author's opinion is a better approach than the previous method mentioned by Galan et al.[96], the method detailed here was actually devised and published earlier and therefore one must presume that the authors of that work felt there were still other aspects of packaging that needed to be addressed.

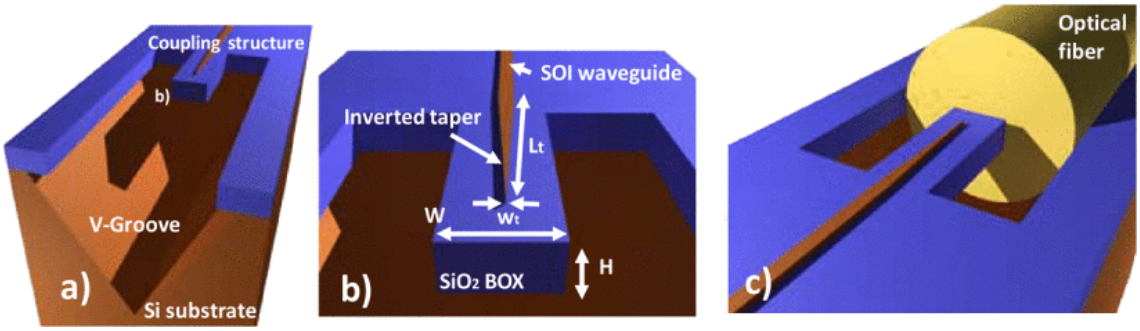

Figure 2.15 Diagram of V-groove inverted taper coupler, taken from [97]

This approach involves the use of inverted tapers rather than grating couplers which is probably why a new packaging scheme was devised two years later. Whilst using inverted tapers is a fundamental departure from grating couplers, it does offer several advantages. As the authors rightly point out, the biggest potential advantage is the effective elimination of polarisation dependant loss (PDL). Another advantage is that tapers typically offer much higher bandwidth since they aren't inherently bandwidth restricted as grating couplers are. In some unique instances this is less desirable where channels are multiplexed either through polarisation or wavelength but in most scenarios the reduction of PDL and increase in bandwidth is favourable.

This method uses a V-groove which is created by undercutting the substrate using an anisotropic etch (this process will be discussed in detail later). The V-groove is then used to hold the cylindrical fibre. Figure 2.15 shows a schematic of the structure that is used to couple light from the fibre into the waveguide. The inverted taper is used to gradually adjust the mode from the large core of the optical fibre into the silicon waveguide. There is a small oxide waveguide that is used such that the mode in the fibre will seamlessly transition into the oxide waveguide since they

share the same refractive index and so optically speaking the light 'sees' no difference between the waveguides. Then the mode is gradually drawn into the higher index silicon waveguide due to the taper, until it is confined totally within the silicon. Theoretically this should be an adiabatic process with no loss of energy.

The reported performance of this packaging method is also impressive. At a standard wavelength of 1.55μm they report coupling losses of 7.5dB with a flat, broad band of 70nm with only ±0.5dB PDL. This is 1.5dB higher than the 6dB suggested by their simulated predictions, but for a taper of only 400μm this is still impressive. The authors also rightly point out that an improvement could be made by increasing the dimensions of the oxide waveguide which is used. The fibre core is circular and has a diameter of 10μm, whilst their oxide waveguide is rectangular and is 8x2μm, an obvious mode mismatch. By applying "6-7μm" of extra oxide cladding the authors predict a 1-2 dB improvement in their total coupling efficiency which is not insignificant.

This method should be praised for its ingenuity and results, especially considering that this is a passive method of alignment. The fabrication tolerances seem to be one of the big limiting factors and so substantial gains could be made by improving the fabrication methods and making the amendments which the authors suggest.

A more recent development using V-grooves and tapers has been demonstrated by IBM, who have put in a substantial amount of work into Silicon Photonic packaging in the last few years [98], based upon the work of Cheban at the National Research Council (NRC) of Canada[99]. IBM rightly observes the potential of Silicon Photonics but also correctly note the bottleneck which prevents mass market adoption – packaging. They mention that in many cases the use of "legacy telecom practises" can result in the packaging process alone costing up to an order of magnitude more than chip itself. This hugely limits the scalability in terms of manufacturing throughput and the number of ports on a single chip is limited as well. In this author's opinion the team from IBM best captures the correct design philosophy behind finding an economically viable packing solution, though their

approach is vastly different. This is a testament to the breadth and depth of creativity and ingenuity that is present currently within the field.

Whilst strictly speaking IBM have developed several methods, the basic concept is very similar in all cases. In order to hold the fibres a V-groove is used as in many other similar projects. Figure 2.16 [100] shows schematic of how the fibres are aligned to the metamaterial converter which is used to couple light into the silicon waveguide. Standard optical fibres can be used for this method, avoiding the need for special lensed fibres or polished ones.

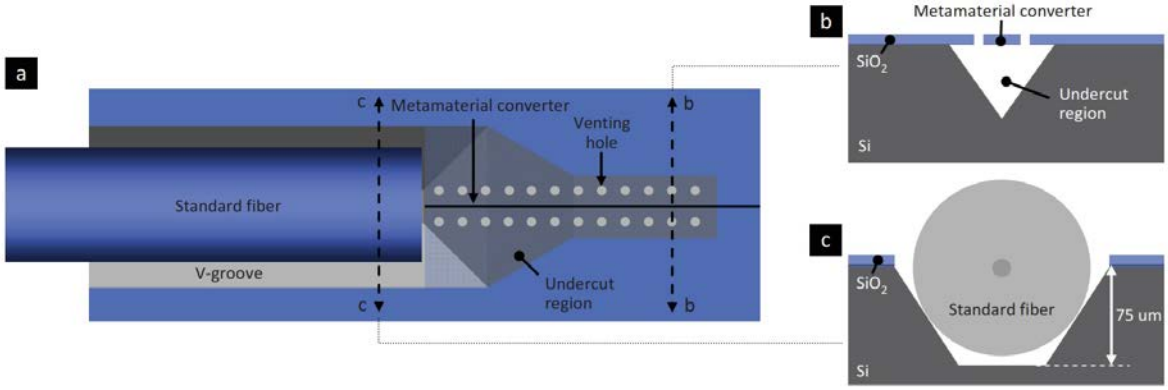

Figure 2.16 Diagram of the IBM packaging method[100]

Of special mention is the subwavelength metamaterial converter which is used rather than a conventional inverted taper as would be typical for this situation. Their reasons for doing so are explained in [100]. Firstly, an inverse taper is fundamentally a 2D solution to a 3D problem. That is, ideally the taper would taper in the vertical direction as well as the horizontal, but this is not normally achievable and hence losses are introduced. Secondly an inverse taper is extremely sensitive to its dimensions and as such, the fabrication tolerances are very important to control, which in practice becomes difficult for mass production. Creating large volumes of identical tapers through standard processing techniques is very challenging; typically, in mass production the repeatability and uniformity of fabricated devices is just as, if not more important than the peak theoretical performance.

So then, they have developed a different method for coupling – using a metamaterial converter. The use of metamaterials for mode converting is relatively recent[101] and has been extended here by IBM for coupling to a conventional optical fibre rather than a lensed fibre. They use a periodically etched, sub-wavelength waveguide to gradually couple into the silicon waveguide as can be seen in Figure 2.17. The venting holes that are pictured in Figure 2.17 above are used to introduce a fluid in which the metamaterial converter sits. Therefore, the cladding

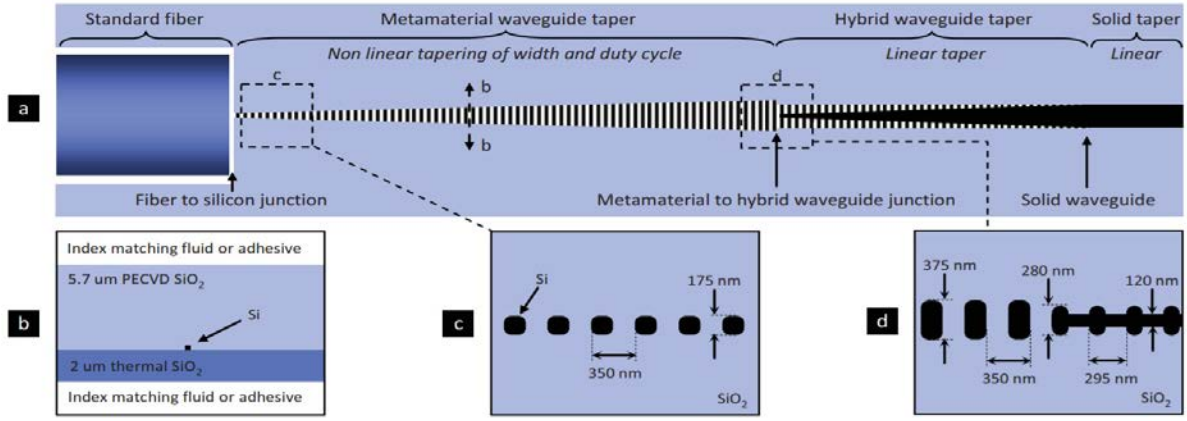

Figure 2.17 Diagram of the Metamaterial Converter used by IBM [100]

material can easily be adjusted for different wavelengths or devices, making this method very compatible with a range of different applications. One assumes speculatively that this approach could also find use within a sensing context, though IBM's focus has traditionally been very much in datacoms.

The performance of this Metamaterial Converter is impressive to say the least. They report excellent coupling losses of up to 1.3dB at an O-band wavelength of 1310nm, with a 0.8dB bandwidth of up to 100nm. Both these figures are

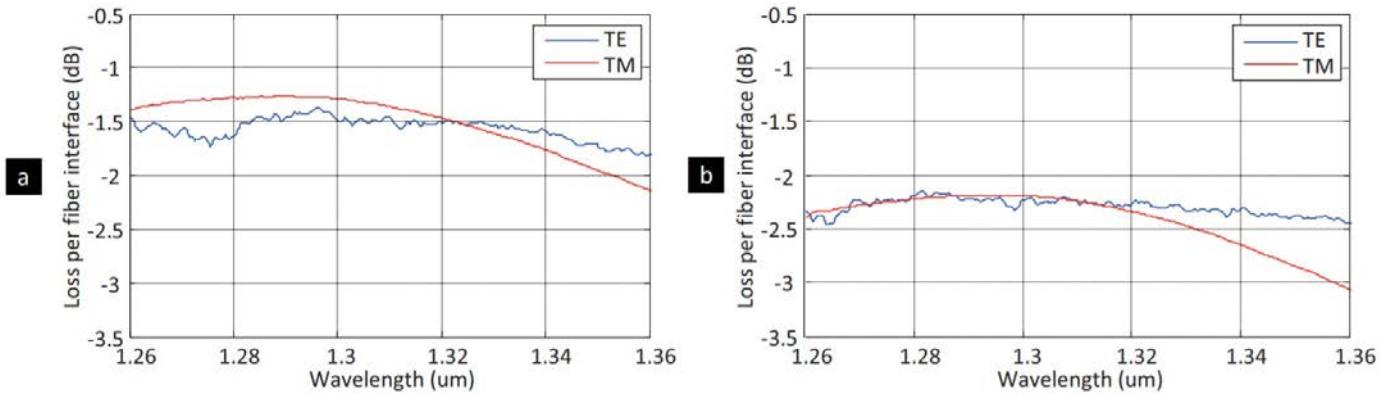

Figure 2.18 Graph of the results of IBM's packaging method [100]

extremely good, far exceeding most competition. Figure 2.18 shows a graph of the performance of the metamaterial converter. As can be seen by graph a, an almost flat band is shown, with almost identical profile for both TE and TM. Graph b shows the losses when the V-groove was etched too deep resulting in a slight fibre misalignment and additional loss of ~0.9dB.

Though this method deserves much applause, there are still a few points of critique to be made. Firstly, an adequate assessment of the tolerance to misalignment is not made in this work, though this is acknowledged by the authors at the conclusion. Hence it is difficult to ascertain how useful the method would be, even though it is clear that given good assembly, this method could be adopted in a widescale way. Also, they report only the use of O-band wavelengths and not C-band, but one assumes there is nothing fundamental to their design that would prevent it also being used for the C-band. One other thing to mention is the use of the immersive fluid. In the tests done by IBM either water or isopropanol was used in order to provide the desired index. Obviously for a packaged product a material which is liquid at operating temperature is problematic. Whilst it wouldn't be conceptually or practically difficult to replace the liquid with an index matching epoxy which could be UV cured, this is yet to be demonstrated. In conclusion then, this method deserves commendation for both the approach and execution.

More recently IBM have also demonstrated the same principle using an array of 12 fibres[102]. In this work they use Monte Carlo simulations to show that the maximum expected assembly misalignment is 1.3μm. This is an impressive number, but no mention is given to the loss penalty that would be introduced by such a misalignment. Indeed, this work shows simply that fibres can be aligned to tight tolerances but does not actually show conclusively that the fibres and chip can be coupled with low loss, even though one could assume that this would be the case.

A very similar piece of work has been demonstrated by Snyder et al., from Imec[103] and Cardenas et al. [104] . The concept is the same in essence, using an inverted taper from a different material to couple into a silicon waveguide. Instead of a metamaterial taper, they rather use a nitride inverse taper to guide the mode from a typical SMF-28 fibre into the silicon waveguide. This is done by creating a

nitride layer on top of the silicon layer with an oxide cladding. They remove the silicon substrate underneath the buried oxide layer, as they rightly point out that the mode radius of the fibre is over 4.6μm and therefore this mode will leak into the substrate unless it is removed. A Bosch etch is used to remove silicon substrate from the underside until the buried oxide is reached which acts like a membrane to etch up to. Figure 2.19 shows a diagram of the coupling structure. Much like IBM's method, they immerse the chip in water in order to provide an appropriate index to match the modes.

The results from this method are impressive as well, with the clearest advantage being bandwidth which exceeds 100nm for a 1dB penalty which is extremely wide. PDL is less than 1dB as well making this coupling method extremely low loss. The total package is reported as having an insertion loss of 4dB for both TE and TM modes across the whole C and L bands.

The strengths and weaknesses are essentially all the same as that of the example above, except perhaps for the fact that a nitride layer is generally simpler to fabricate than the metamaterial coupler that was developed by IBM.

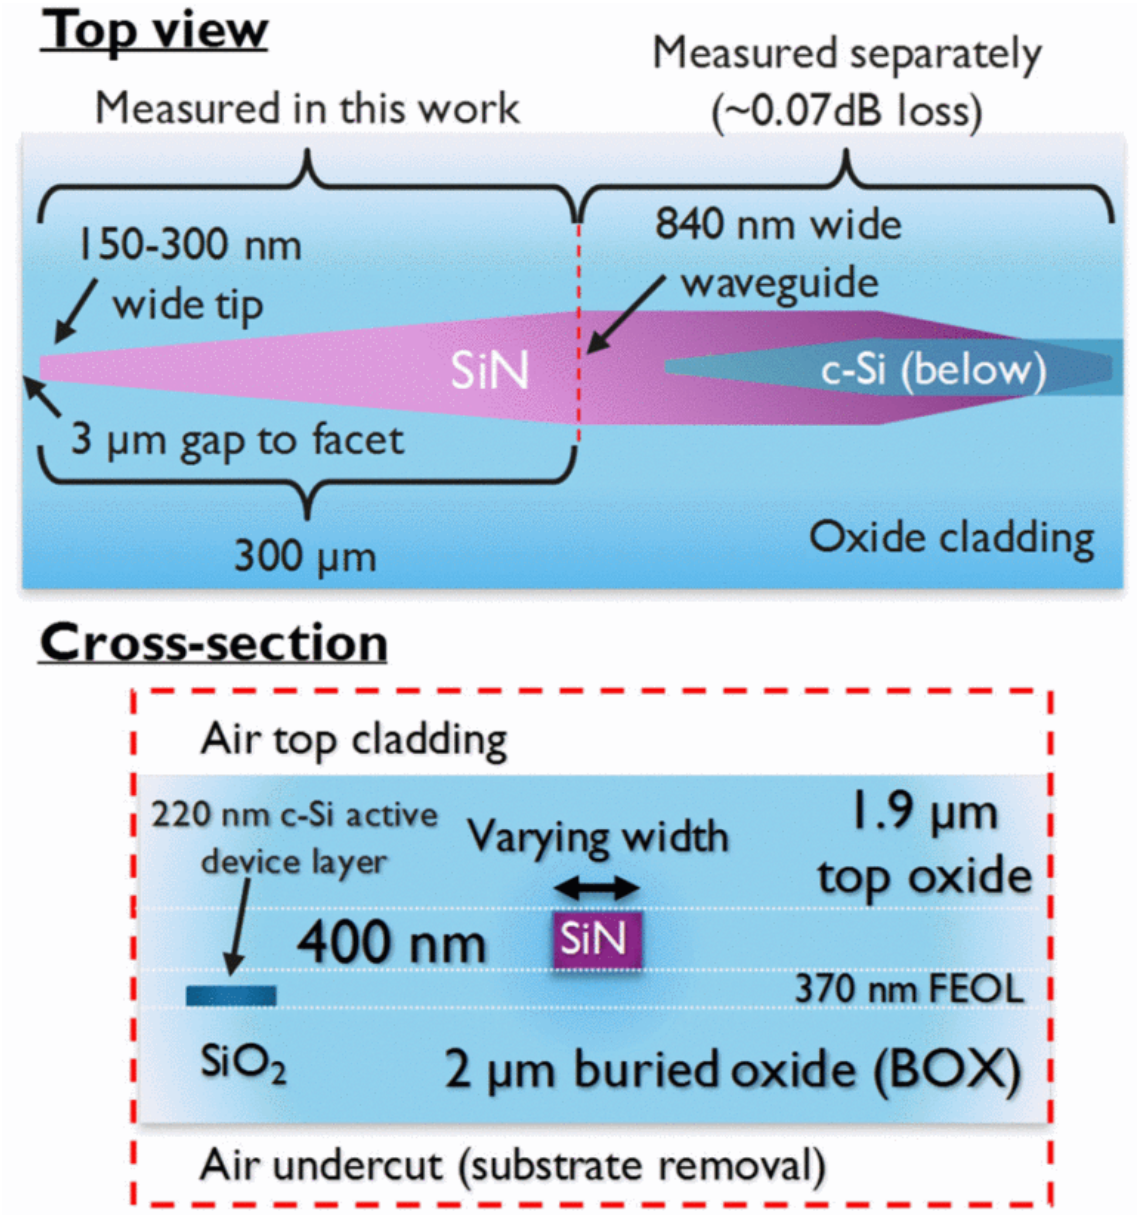

Figure 2.19 Diagram of Imec's SiN inverse taper coupler[103]

2.4.4.2 Passively Aligned Laser Package

Cohen et al. demonstrated a method for passively aligning a laser as far back as 1992[105]. Even though the field has progressed and evolved dramatically in the last few decades there are still many important contributions from this work that deserve mention. Whilst this work doesn't feature light coupled from a silicon waveguide into a fibre there are still many important principals for photonic packaging that must be discussed.

A chip with four AlGaAs lasers is coupled into a fibre array. No specific coupling scheme is used, the lasers are simply butt coupled into the fibre array. The alignment comes from a separate "alignment plate" which is used to fix both the laser chip and the fibres as can be seen in Figure 2.20. Hence the alignment accuracy for this method is very much dependant on the accuracy of photolithography which is used to define the fiducials used for aligning the three structures. The fiducial marks are defined at the same stage as the laser ridges and v-grooves allowing precise alignment with respect to each other. During assembly the fiducials are all aligned visually and when the cross shaped fiducials are aligned the chips are soldered to each other to provide the necessary adhesion.

The end results of this packaging method are impressive too. They report that the passively aligned packaged laser achieves 80% of the coupling efficiency of an identical actively packaged laser. The chip's tolerance to lateral misalignment is also good with a 1dB loss of ±5μm and a 3dB loss of ±18μm, however this is with multimode operation. The absolute coupling efficiency isn't specified, all losses are present with respect a normalised figure. At first glance this is very successful but

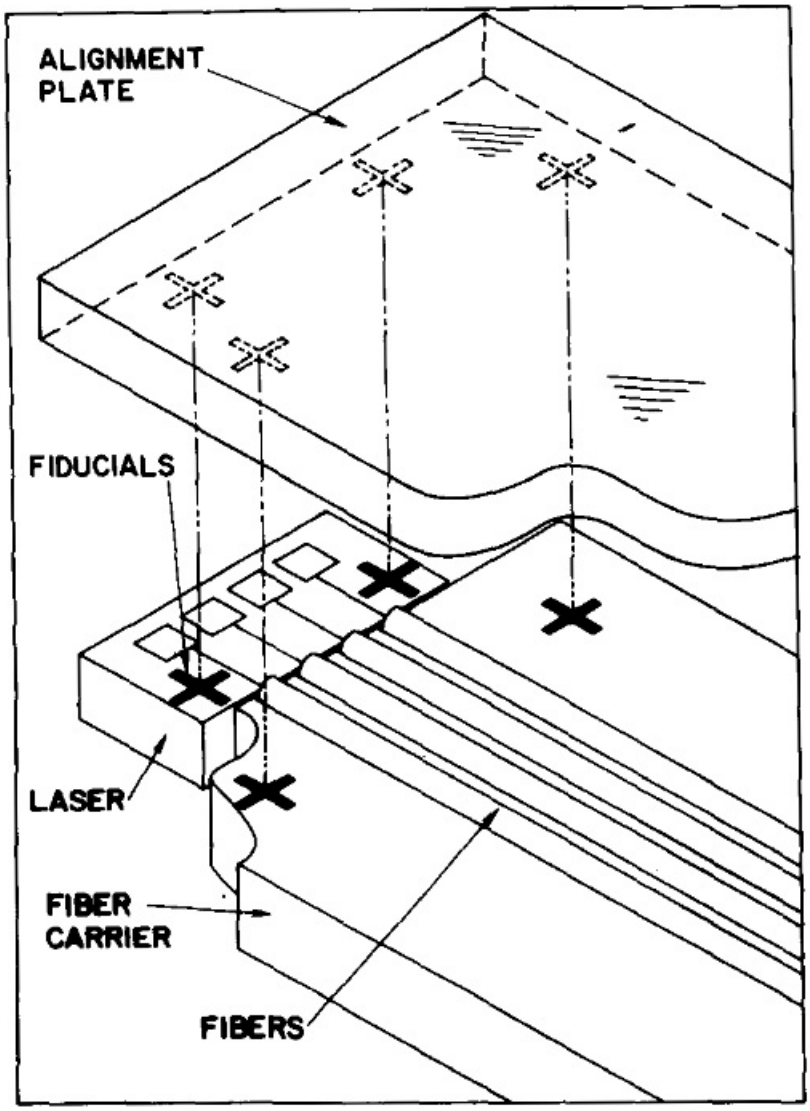

Figure 2.20 Diagram of a laser to fibre package by Cohen et al. [105]

must be qualified in a few important ways. With regards to alignment this method is technically passive; no optical feedback is used to align the chips to each other. However, so many extra elements are added that the gains that are made by passive alignment begin to be undermined by all the added complexity. Due to the extra complexity this method becomes extremely slow and so this severely limits its usefulness for large scale adoption. That said, this work does provide many important contributions, namely the use of photolithographically aligned chips which are passively aligned.

2.4.4.3 Capping Chip Method

An interesting piece of work from Bernabé et al. [106]demonstrates a method that is here summarised as being a capping chip method. This uses an additional chip both to hold and locate the fibres as well as reflect the light propagating out of the fibre onto the optical chip which contains the waveguides – an SOI wafer. This is depicted in Figure 2.21. Here the capping chip is called a

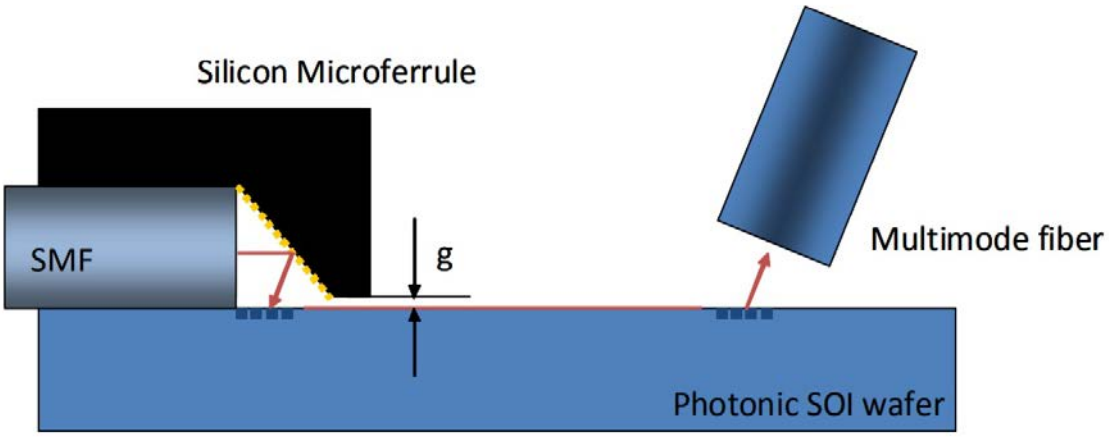

Figure 2.21 Diagram of Capping Chip method by Bernabe et al. [106]

micro ferrule rather than a capping chip. This is what the authors of that work have designated it since it is in essence working as a ferrule, however for consistency with the work in this project it shall be called a capping chip. This work also clearly lays out some '*silicon photonics packaging laws*' which is should be applied to create a functional packaging solution. They are as follows,

1. Multiple channel compatibility (at least 12 channels) with single mode fibers. Tight alignment tolerances must be addressed [107]
2. Final assembled devices should be low profile ("Inplane") and compact in order to fit the next generations transceivers footprints [108]
3. High throughput process and low assembly cost [109]

In many ways these three statements succinctly summarise the conclusions that have already been drawn out this far in this literature review. One must consider each of these three '*laws*' in order to produce a workable packaging solution. The authors of this paper have sought to incorporate these elements into their packaging philosophy. Whilst one must always be cautious around laying down 'laws' such as these for fear of enslavement to dogmatic principles over and against creative and original ideas, in this case these maxims do helpfully focus and guide us towards a packaging solution that will be of use.

In this work a silicon chip is wet-etched to create the fibre V-grooves, but the end of the V-groove can be used to create a mirror due to the angled sidewall that is formed. This is then coated in gold in order to be reflective enough. This process allows a whole array of fibre v-grooves to be etched simultaneously that will all be aligned respective to each other. In this case the pitch between v-grooves being 250μm.

After fabrication and assembly, the authors report that this packaging method introduces an extra packaging penalty of 4dB on average. This meets their predicted values very nicely. These chips were packaged without the use of an index matching material in the gap between fibre and optical chip and the authors predict that if such a material was included that the packaging penalty would be reduced by roughly 1-2dB, depending on the success of the assembly. Another important element of this team's analysis was a study of the thermally dependant behaviour of this method. They report a more or less stable performance over the temperature range of -40 to 85°C, representing the full range of reasonable operational temperatures (for now![110]).

This method has many attractive qualities and the packaging philosophy in general should certainly be commended. As always however there are a few points of critique which are to be made. The authors claim this to be a "semi-passive" method. That is, that computer vision is used to locate the two chips together, without an optical feedback circuit. Whilst this is technically not an active alignment the authors are correct to dub this a semi-passive alignment method, using their optimistic naming scheme. One also notices that according to Figure 2.21 the capping chip will sit proud of the optical chip by an amount g. This would make attaching the two chips challenging and it is not described in the assembly section of the paper how this was overcome, since the fibres were inserted after the chip attachment. No assessment is made to the potential misalignment from this method, nor the impending losses that would entail. Therefore, whilst this work contains much to be embraced, it still is not the required fully passive alignment solution that is being sought.

2.4.5 Pluggable Packaging

One recent piece of work that deserves mention is the pluggable coupling method that has recently been demonstrated by Scarcella et al. from Tyndall[111]. By pluggable, this means that the optical connections to the chip can be plugged in and out much like a commercially available connector, such as is popular in consumer audio visual equipment for instance. Figure 2.22 shows a both an image of what the method looks like when used but also giving a graph of the performance.

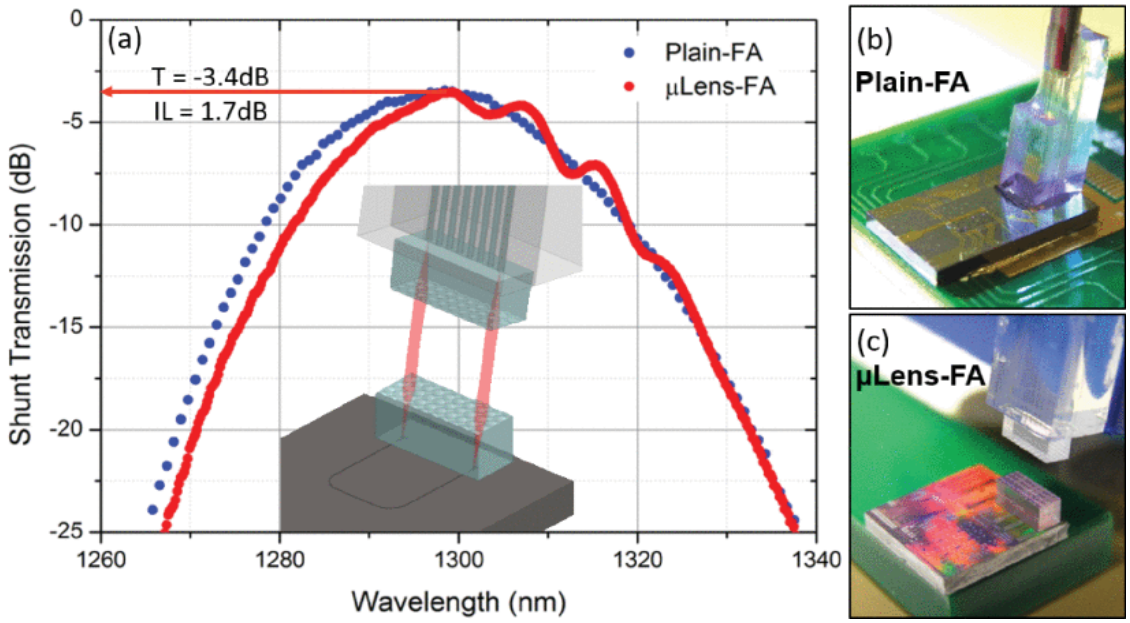

Figure 2.22 Results and images of Pluggable Packaging from Tyndall [111]

The key ingredient that is necessary to this coupling method is the presence of a micro lens which is used to focus light onto the chip surface. The light is then coupled into the waveguide by using a conventional surface grating coupler. In the instance where a micro lens is incorporated into the fibre array then the array will sit roughly 1mm above the surface of the chip. Using this array of micro lenses, the

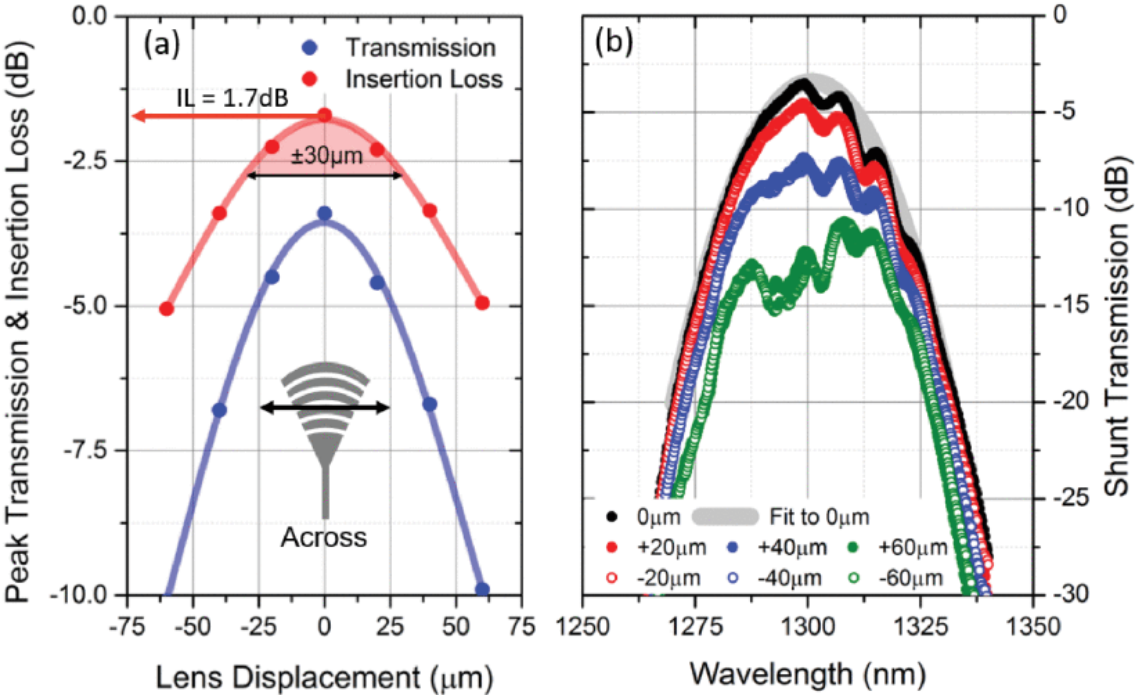

Figure 2.23 Graph showing the misalignment effects of the Pluggable Coupler from [111]

tolerable misalignment that can be allowed is increased. This is demonstrated by Figure 2.23 which shows a graph of insertion losses against misalignment. As can be seen from the graph, the peak insertion losses are 1.7dB when there is no misalignment. This figure then rolls off parabolically as the lens is misplaced in any direction, with a ±30μm misalignment penalty only costing 1dB of decreased performance. This is an incredible achievement by any estimation, where most packaging methods are considering alignment tolerances of ±1-3μm, so essentially this allows an order of magnitude of freedom over most methods, with still extremely low losses.

So far, the method that has been described can hardly be called a packaging method, even if it is an ingenious method of coupling. However, the extremely relaxed alignment tolerances have vastly reduced the precision and complexity of any alignment method that is used. Hence it can be considered “pluggable” as the name suggests. This is demonstrated very elegantly by their use of Lego bricks as alignment structures. This is depicted in Figure 2.24. As can be seen nothing more than commercially available toy Lego bricks can be used as an alignment structure.

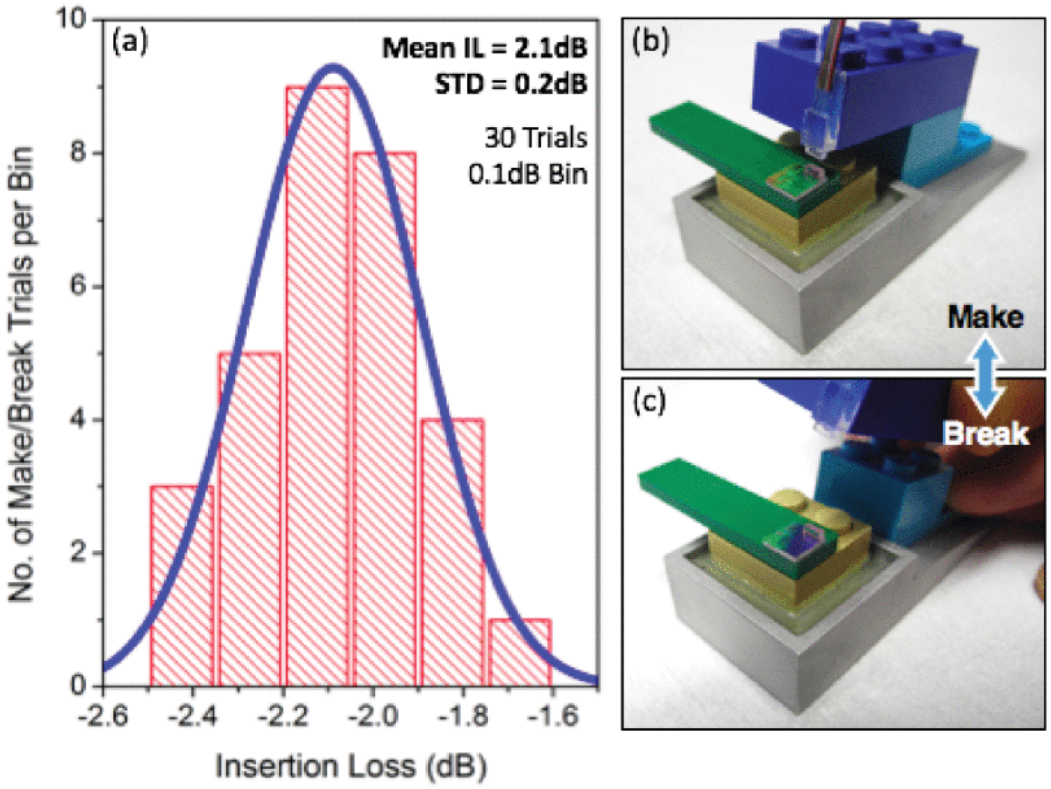

Figure 2.24Further results and images of Lego Bricks used for Pluggable Coupling [111]

This is an extremely novel and amusing concept but what it shows most importantly is that efficient coupling (2.1dB) can be achieved using very cheaply mass manufactured parts. This must also be commended as a demonstration of the precision and manufacture standards of Lego which must be in no small part integral to its decades of success as a children's toy.

Whilst it could be argued that this method fits squarely within the section on "out-of-plane" solutions, it has been placed in its own section here since it should better be considered only a semi-package as it currently stands. This method could potentially find use in applications where an optical channel must be plugged in and out several times but perhaps less so in situations where the channel is to be left alone for an extended period of time, as in most near-future applications such as optical channels in data centres. Since it is still in essence an out-of-plane solution it will likely be too cumbersome for many uses but no-doubt this will find niche applications for which it will be extremely useful. Most likely in-situ sensing applications, such as biosensing could find this method very appealing since a biosensor could be plugged onto or off of a patient or test subject and then removed

when enough data has been acquired. The work demonstrated here was carried out using active alignment as well, but the authors point out that this is only as a proof of concept and as such with the use of direct writing of micro lenses or wafer bonding that this could be transformed into a passive alignment packaging method. In summary this is a very attractive proposition but probably shouldn't be seen as a direct competitor to conventional packaging solutions (yet) and rather as something that could be complimentary, allowing semi-permanent packaging.

2.4.6 Packaging Design Rules and Future Standardisation

In order for silicon photonics to be accepted into the mainstream of data and telecoms applications it will likely be necessary for certain industry standards to be made and held otherwise progress and cohesion will be stifled. It is clear to see from the example of electronics and before that landline telecoms networks, that advancement and acceptance of new technologies will only be realised once manufacturers and researchers can agree on the standards and crucially investors will typically only inject the necessary funds once there is a clear standard, such that their investment isn't usurped by a new competing one. To that end design rules are always something of a compromise, but one that is very much necessary.

Already design rules for Silicon Photonics are being proposed and as such it is necessary in this work to be aware of what these could be and especially in packaging to devise methods that would be compatible. In a paper by Morrissey et al. [112] the authors highlight the issues of not having a standard. They state quite bluntly that "*Currently, the lack of readily available design rules and the fragmentation of both assembly and packaging capabilities, are the most significant roadblocks limiting the commercialization of PICs*". Whilst this is quite a bold statement, this author is in agreement and hence it is right that until design rules are instituted, all research output must be sensitive to any potential standards that are formalised. The authors also rightly point out that the rules for Silicon Photonic packaging must cover the whole scope of variables, taking into account not just optical properties but also thermal, electrical and mechanical aspects of any given packaging method.

Many research institutions and commercial entities are working together to form PDKs (process design kits) and PDRs (process design rules) in order to lubricate the development of Silicon Photonics technology and this will in turn provide a helpful boost for newcomers to the field such that the same foundations need not be laid again. This also limits the disastrous situation where large amounts of time and money end up being wasted by developing devices that cannot be easily or economically packaged.

Figure 2.25 shows a thoughtfully devised set of design rules for silicon photonic packaging. As can be seen, these rules relate specifically to their methods

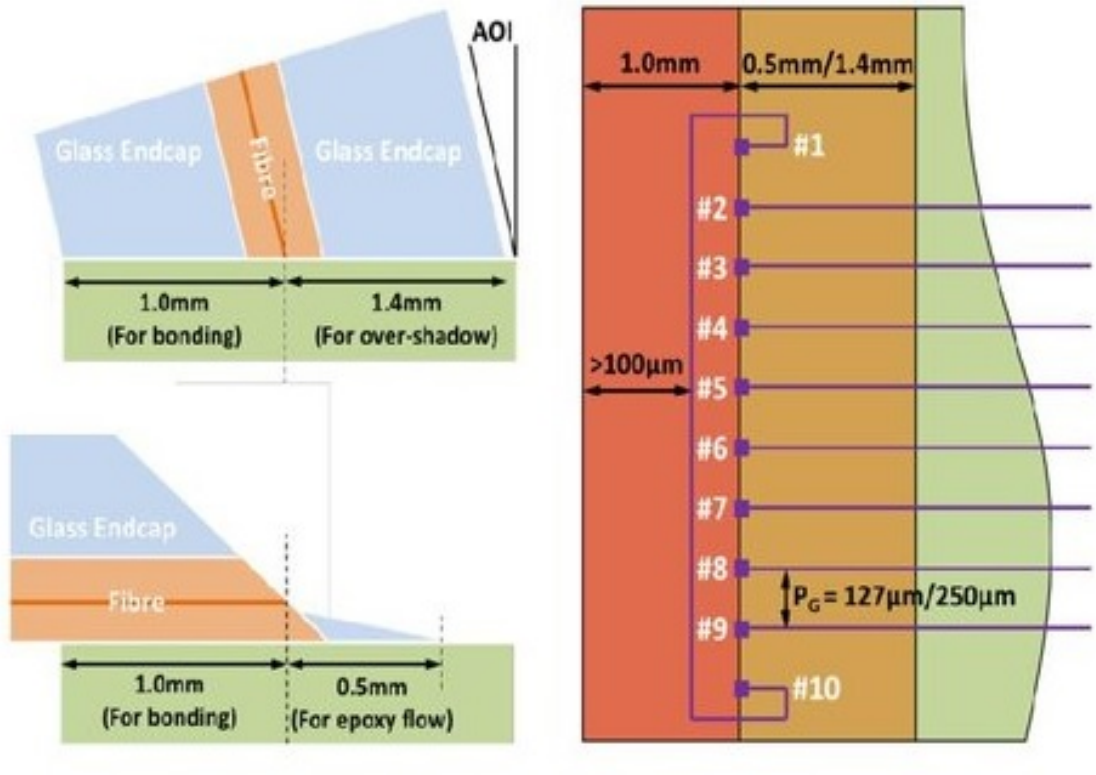

Figure 2.25 Diagram of design rules devised by Tyndall for silicon photonic packaging [112]

of packaging as they have been outlined previously in this chapter, polished fibre and pluggable ferrule base package. It remains to be seen as to whether these particular methods will be adopted in the mainstream, yet despite that this is a valuable exercise, having one eye on the future and outlining rules such as these. Of more interest is perhaps the image to the right, which shows the dimensions and

spacing between optical channels in a fibre array. Here provision for two redundant channels are made (#1-10) in order to facilitate active alignment. Of course, the packaging method outlined in this work seeks to avoid this altogether but if non-bespoke optical chips are to be fabricated with such extra channels the work presented in this thesis needs to be compatible (i.e. leave these arrays directly adjacent to the optical channels free). Without the two redundant channels, this represents an 8-fibre array being packaged to a single optical chip, with standards of pitch and orientation being shown.

In addition to optical packaging, close attention must be made to electrical packaging as well since in all likelihood both photonics and electronics will remain close bedfellows of the foreseeable future. For the most part it is proposed that electrical contacts will be positioned at the 'north' and 'south' sides of the PIC, as in [112], and then the optical channels will be positioned at the 'west' and 'east' edges of the pic. Therefore, a method must be devised to fit neatly to this regime. For example, a capping chip must not extend too far toward the north or south edges of the chip which would occlude any potential electrical packaging.

Until a firm PDR is synthesised a certain degree of flexibility must always be retained when designing and fabricating a packaging method, and it was with this in mind that the work in this project was undertaken.

Chapter 3 Theoretical Design and Simulation

3.1 The approach used in this project

This section outlines the approach that was used for this project, as proposed and patented by Thomson et al. in 2013 [113], discussing the basic principles as well as the different technologies that are required to fabricate it. The method has changed significantly but not fundamentally since the patent was filed. This section will highlight how the concept began and how it has evolved to meet both fabrication and assembly demands.

3.1.1 Original Method

By combining different technologies and processes, it is possible to construct a method for passive alignment that is compatible with Silicon Photonics fabrication and packaging. Figure 3.1 shows a diagram of the original proposed solution. This involves using a "***Capping Chip***" in conjunction with a "***Device Chip***" and a conventional optical fibre, hence this can be considered a 'hybrid two-chip' packaging solution. The Capping Chip undergoes an anisotropic wet etch to form the required V-groove as well as the optical mirror. By combining these three elements as depicted in Figure 3.1, light can be coupled into the optical chip without the need for active alignment. The device chip contains grating couplers as well as

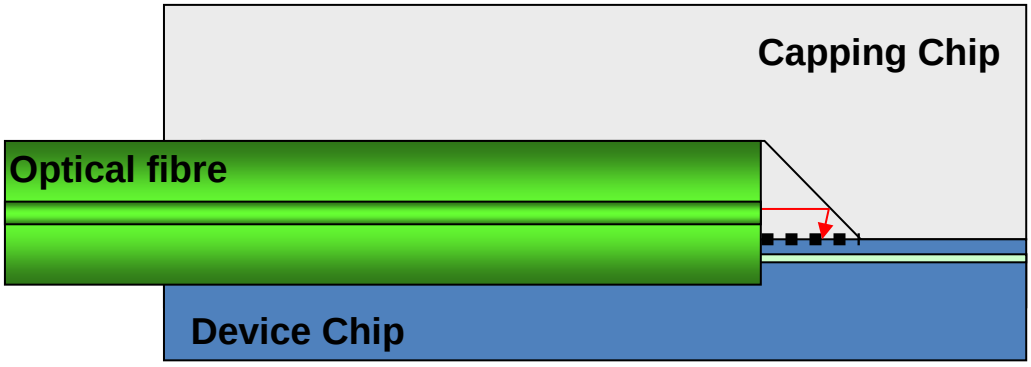

Figure 3.1 Diagram of a passively aligned optical fibre and chip

waveguides and any other devices such as modulators etc. as is suggested by its name. By their very nature grating couplers restrict the optical bandwidth that can be effectively coupled into waveguides but with 3dB bandwidths of 40+nm generally achievable[114], [115] this is more than adequate for many applications at 1550nm C band wavelengths. The optical chip also featured a dry etch which creates a channel to lower the optical fibre down by 55.5µm. Since a typical SMF-28 has a diameter of 125µm, this considerably shortens the propagation distance between the end of the fibre and the grating coupler from 111µm to 52µm as depicted in Figure 3.1 [116]. This in turn will dramatically reduce the diameter of the diverging light beam spot from an unacceptably large 24µm to a much more reasonable 14µm. (See Kowalevicz and Bucholtz for further discussion on this [117]).

This solution provides an effective yet conceptually straightforward method of passive alignment, with the optical fibres being held in the same plane as both the capping and device chips which greatly reduces the effective size compared with previous packaging solutions incorporating grating couplers[118]–[120]. The proposed solution offers the potential for a more robust, less complicated and cheaper alternative to previous designs.

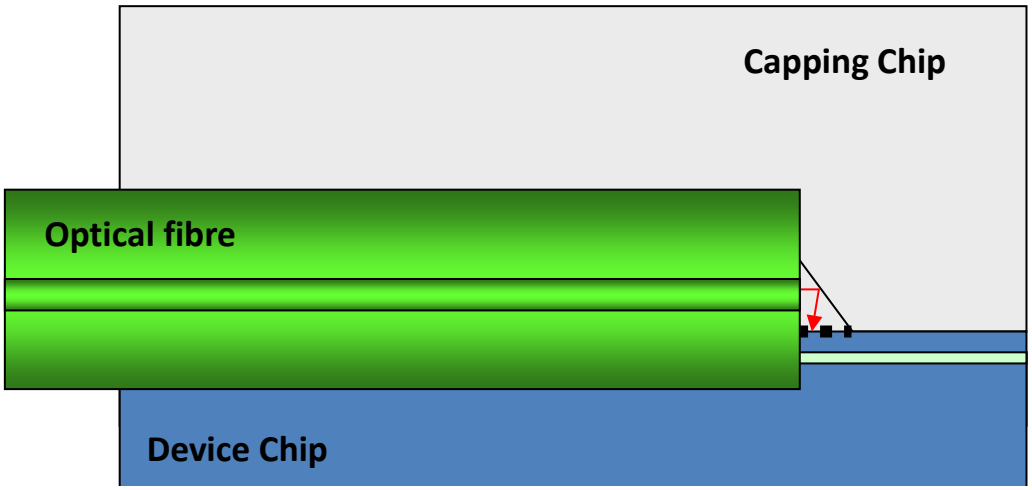

Figure 3.2 Diagram of a Passively aligned fibre and over-etched capping chip.

One simple way that this method can be improved yet further is with a simple "*over-etch*" of the Capping chip. As can be seen from Figure 3.1, only a fraction of the mirror that is present is actually required in order to reflect the beam towards the

grating coupler. All of the sloped edge above this critical point is unused but still adds to the propagation distance of the beam.

Figure 3.2 displays what such a configuration wold look like. A conservative over-etch of length 20μm would reduce the propagation distance yet again down to 30-35μm. These modifications bring the total propagation distance to be comfortably below the Rayleigh length of a light beam ejected from a fibre[117]; where a propagating Gaussian beam has doubled in radius. Such an etch could also be extended to be through the entire wafer in order to give a means of visually inspecting the fibre as well as providing bleed space for any potential epoxies which may be deposited inside the cavity (see Section 3.2 and Chapter 4). Whilst this propagation distance is greater than for conventional vertical coupling, it is only so by ~20 μm. At a typical angle of 10° for vertical coupling, when the width of the fibre

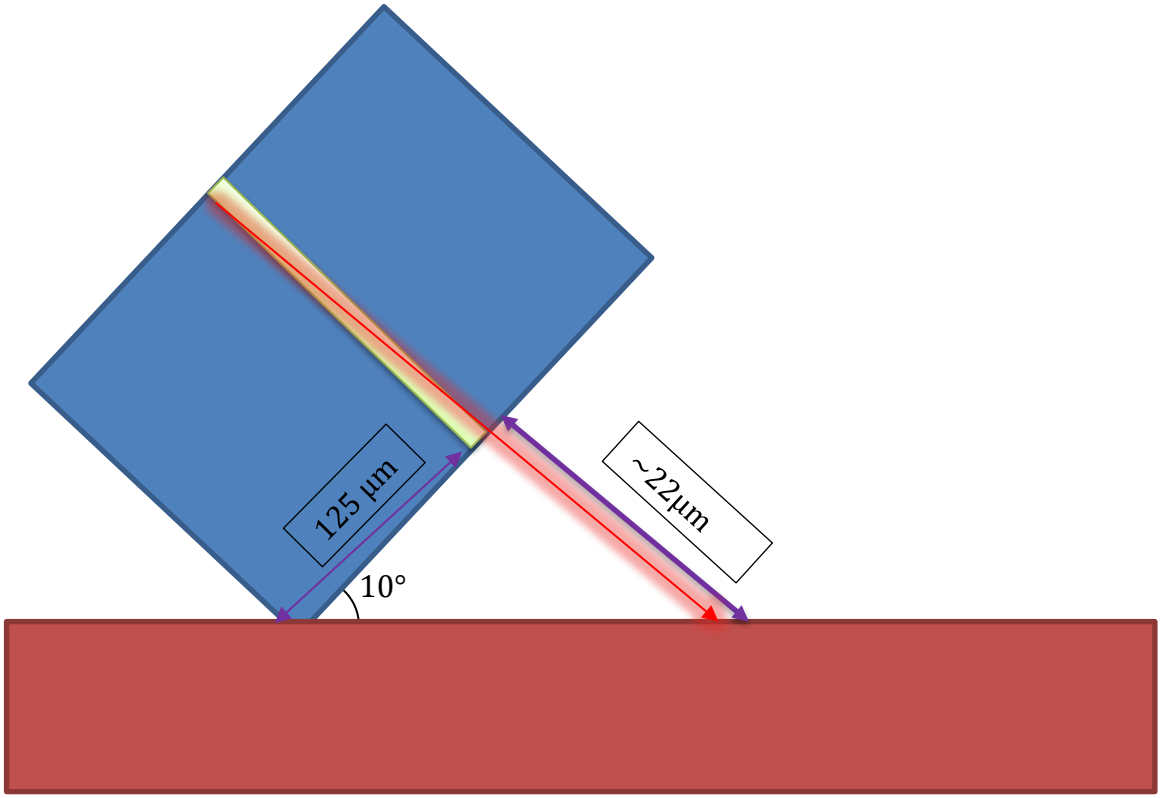

Figure 3.3 Diagram showing the propagation distance in conventional vertical coupling at 10° using a standard SMF-28 fibre.

cladding is considered there is still a propagation distance of approximately 22μm as can be seen in Figure 3.3. So, despite the fact that there is an increased propagation distance in this packaging scheme, the extra distance is less than one

might think, and the extra difference has been accounted for in this design of specific features of the devices.

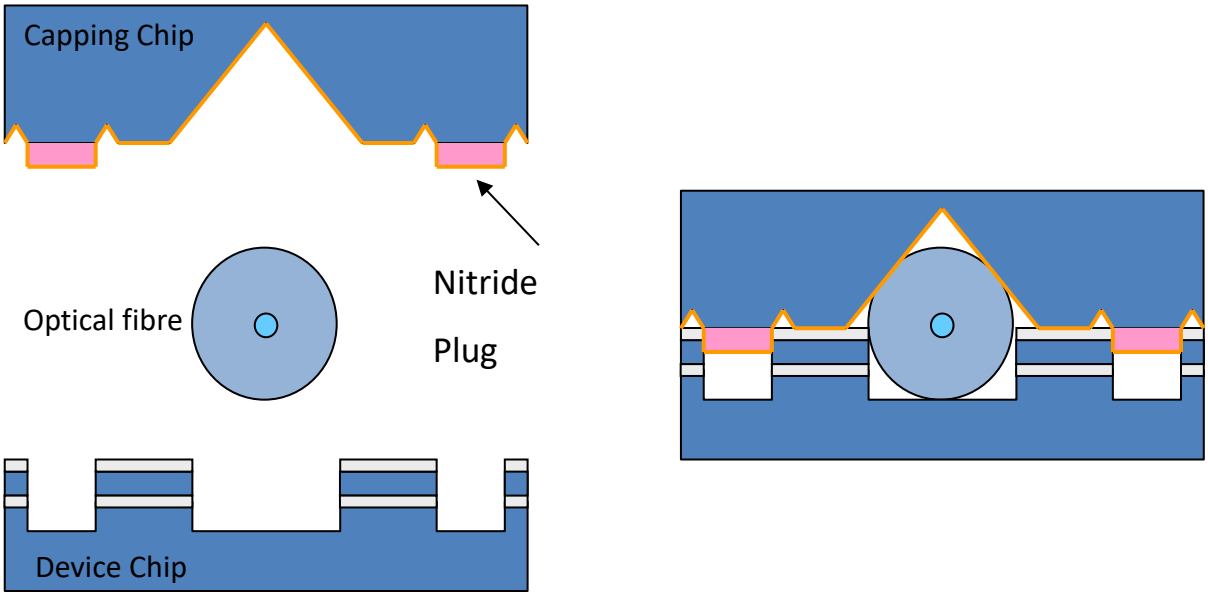

Figure 3.4 Side-on angel of the original packaging design

In order to bond both the chips together a number of '***plugs***' were designed into the capping chip. These were to be formed from a Si_3N_4 (or simply '**nitride**') on Si wafer, avoiding the need to grow or deposit nitride onto a Si wafer. Oxide is deposited on the surface and used as a hard mask which is then patterned through lithography in order to outline the plugs. By etching the remaining oxide with hydrofluoric acid (HF) a hard mask remains over a region of the nitride which will form the plugs. The substrate is then etched in orthophosphoric acid in order to remove the nitride surrounding the plug area, leaving a well-defined structure. Once the plugs have been formed, they can be placed into the holes in the device chip during assembly. Figure 3.4 shows a diagram of what the finished plugs would have looked like.

3.1.2 Evolution of the packaging method

Despite the method outlined above being promising and having many good ideas, it became apparent after the first design and fabrication cycle that improvements needed to be made. Like all research, there are always unforeseen issues which only become apparent after creating what seemed like a very viable

design when on paper. This section will compare the original design to the final design and highlight the key changes that were made along the way.

The first issue was the nitride plugs that were intended to bond the two chips together simply were not able to provide adequate structural rigidity at a size that was reasonable. The original dimension for the nitride plugs was to have a depth of 4μm. This 4μm plug would then be glued into a corresponding hole using an epoxy. This depth simply was not enough to align and bond two chips that were of the order of several hundred microns in width and length. Increasing the thickness of nitride wasn't really an option either for two reasons. Firstly, 4μm is the maximum commercially available on a silicon wafer, which would mean having to grow over 4μm onto a wafer which is an extremely slow and tedious process. The depths required would need to be of the order of 100μm which is just not feasible in terms of process time. The second factor was that of stress. A nitride layer 100μm thick would cause a great deal of stress[121], [122]. This in turn leads to the wafer bowing, making accurate etching of V-grooves impossible. Hence it become obvious that an alternative solution was needed.

Three separate solutions were thought of. These are:

1. Use SU-8 to create micro pyramids which would act as the plug
2. Etch the whole surface of the capping chip except a section of the plugs to leave pyramids that would act as plugs
3. Use a silica microsphere to sit in the hole and act as the plug

Each of these approaches will now be discussed and reasons given for accepting or rejecting each one.

3.1.2.1 SU-8 Pyramids

SU-8 has been used to create a number of different shapes and geometries for microfabrication and MEMS[123]–[125]. SU-8 is a photoresist that has been designed to be able to make very thick films. Once it has been cured it then becomes extremely hard and sturdy. As shown in references in the previous chapter this allows SU-8 to be used for making microscopic 3D structures that would be ideal for use as plugs in this packaging method. Unfortunately however, it was not feasible to

use SU-8 in the Southampton Nanofabrication Cleanroom (SNC) due to the regulations imposed there. Talks were held with external clean-room facilities to find a way of using SU-8 but were sadly unfruitful in the end and this approach had to be temporarily abandoned. It would be an interesting option to explore in the future, however.

3.1.2.2 Two-step KOH etched pyramids

The second option was to use an extra KOH step to create pyramids on the surface. This would involve etching away the entire surface of the chip, except for a few specific points which would become raised pyramids. This method has been achieved before[126], [127] and so is not without precedent.

Figure 3.5 indicates how this would be achieved, the square mask is able to create a pyramid but there is a predictable undercut that comes from the KOH etch

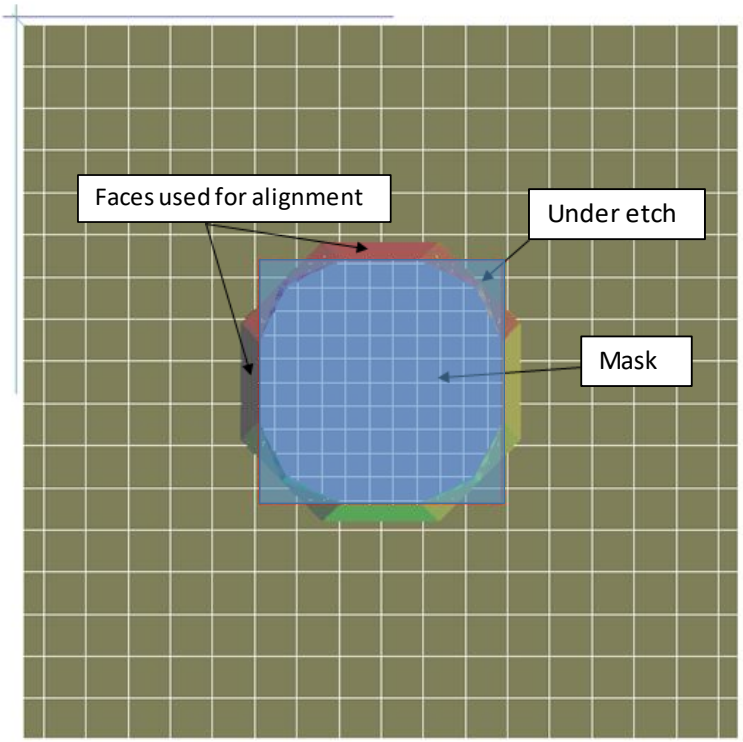

Figure 3.5 3D diagram showing how the pyramids would be etched

process. Correspondingly, Figure 3.6 shows a diagram of the side-on profile of the suggested method.

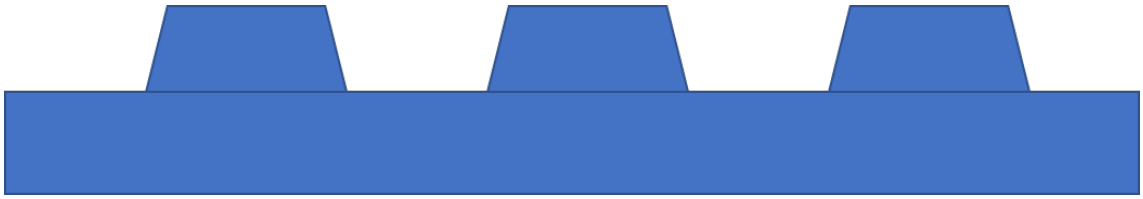

Figure 3.7 Diagram of Chip with pyramids viewed side-on

This method has many attractive qualities. There is no new processing required, merely additional processing of the same kind. This would decrease the time taken overall to develop the necessary recipes. It requires few constituent parts since the same wafer is being used for both the v-grooves and the pyramids. However, there are some issues. This method required a very large etch area which can affect the etch chemistry and make it less predictable than originally suggested. It is also heavily reliant on e-beam alignment, which may become problematic for a large-scale scenario.

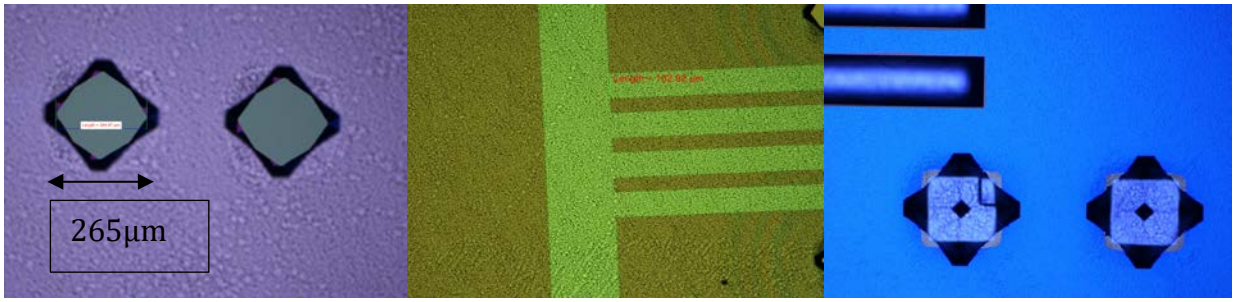

(a) **(b)** **(c)**

Figure 3.6 Microscope images of double KOH etch pyramids

Figure 3.7 shows a series of images from testing the KOH pyramid method. Image (a) shows that the initial lithography was good and well defined according to the required dimensions. (B) shows that the V-grooves were able to correctly be defined in relation to the pyramids. However, c shows how in many instances the protection offered by the nitride layer was insufficient. This was with a 1µm thick layer. At this point it becomes clear that in order for the pyramids to be adequately protected a thicker nitride layer is needed. Hence, very quickly it begins to suffer from the same issue as the original plugs and so becomes problematic.

3.1.2.3 Silica Microspheres

The final proposed option was to use a silica microsphere to act as a plug and provide the necessary alignment. This already has precedent for aligning chips together[128] as well as passively aligning MEMS structures[129]. Different materials of spheres are available such as sapphire and ceramics, but silica was chosen due to cost and reduced risk of contamination since silica (oxide) is already incorporated into the chip. Using a sphere also improves the design by allowing for easier 'self-alignment' as Figure 3.8 shows. Since the sphere is round, it helps both

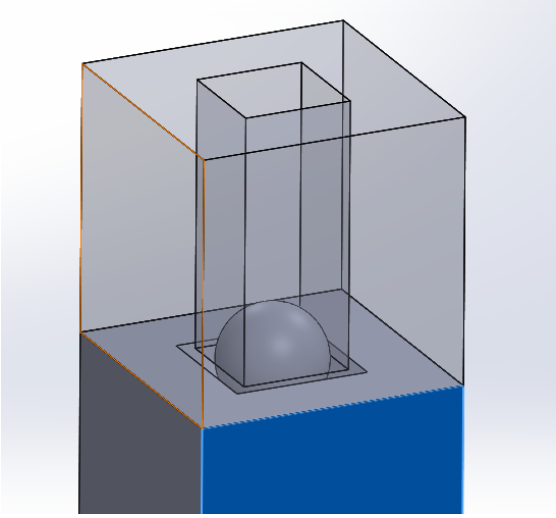

Figure 3.8 Diagram of silica microsphere used to align the Capping and Device Chip

chips to locate themselves together. It also reduced the requirement for an extra alignment step since the hole which the sphere sits in can be patterned at the same time as the v-grooves and etched simultaneously as well. This reduces maximum misalignment.

This does of course rely on the accuracy of the size of the spheres which is typically ±1.5µm for 200µm diameter spheres.

3.1.2.4 Through-wafer etch

The other big change that was made to the original design was the etch depth of the fibre trench in the device chip. As Figure 3.1 shows, it was originally intended that

the device chip would have the fibre trench etched by a few tens of microns to allow for the fibre to sit lower in the trench and reduce the propagation distance of the light beam. However, this was further extended so that this trench was etched through the entire wafer, so as to allow vision of the fibre during assembly and for easier access of adhesive epoxy.

Figure 3.9 shows the final design that was pursued in this project, using both the microsphere and a through-wafer etch.

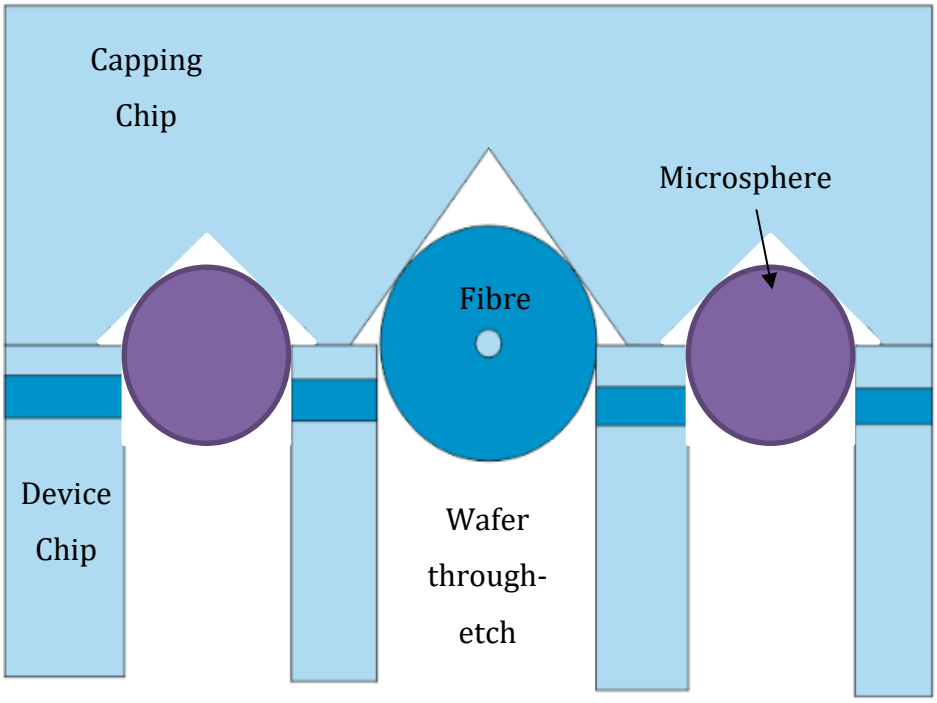

Figure 3.9 Diagram of final method used, with microspheres and wafer through-etch

3.2 Alternative Approaches

There have been a few alternative approaches that have been explored in slightly less depth throughout the project which shall be mentioned here. These alternatives are better thought of as complimentary methods since they could easily be employed alongside the method explained above.

3.2.1 Inverse Taper

Whilst the bulk of the work of this project focused on using grating couplers, the possibility of using the same passive alignment technique but with an inverse taper was also explored. The basic function and design of an inverse taper was inspired by Cardenas et al.[104]. Figure 3.10 depicts how this could be designed. The capping chip would still be used to support and align the fibres to a taper rather than a grating coupler. This method is similar to what was proposed by Wood et al. in [130], but rather than using grown oxide, an index matching epoxy can be used; which is used for the packaging process anyway. No mirror is required either for this method and hence no metal deposition. The inverse taper has a few advantages over grating couplers. There is reduced polarization dependant loss associated with inverse tapers as mentioned in the previous chapter. In addition, an increased bandwidth is also possible. This could be very advantageous for data and telecoms due to the increased option for wavelength division multiplexing (WDM). It can also be relatively easily integrated alongside vertical coupling methods for whichever application is required. Simulations show that inverse tapers can be used with relatively low loss ~1dB, however taper length can become an issue since this must be relatively long.

Since the wafer is being etched all the way through, this could extend to the region directly under the silicon optical layer. Index-matching epoxy would then be flooded into the remaining gaps, creating an epoxy/oxide waveguide which the fibre would couple into. Then, along the length of the taper the mode would slowly be drawn into the silicon waveguide. More details about how this was designed can be found in the subsequent chapter on simulations.

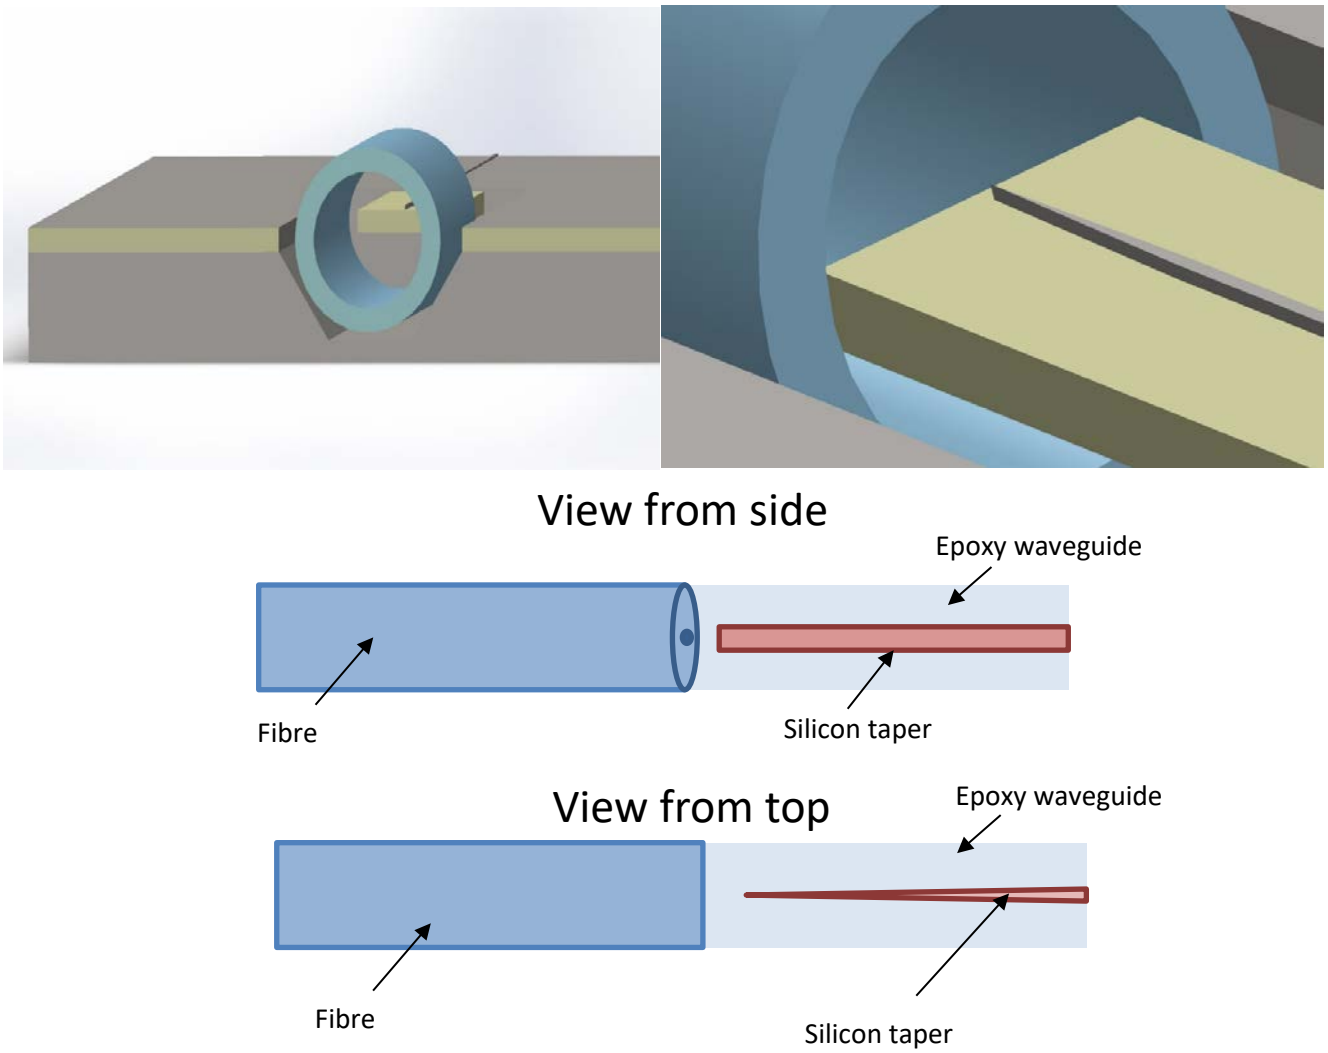

Figure 3.10 Diagram and design proposal for an inverted taper

3.2.2 Holey Fibre Taper

An extension of the inverse taper approach would be to adjust it for use with a holey fibre[131]. The use of hollow-core "holey" fibres is being researched for extremely high data speeds. In such a fibre, the mode is guided through airgaps in the core rather than silica. In this case, the taper could be suspended on the oxide layer and then inserted into the gap in the hollow fibre. Figure 3.11 shows how this could be done.

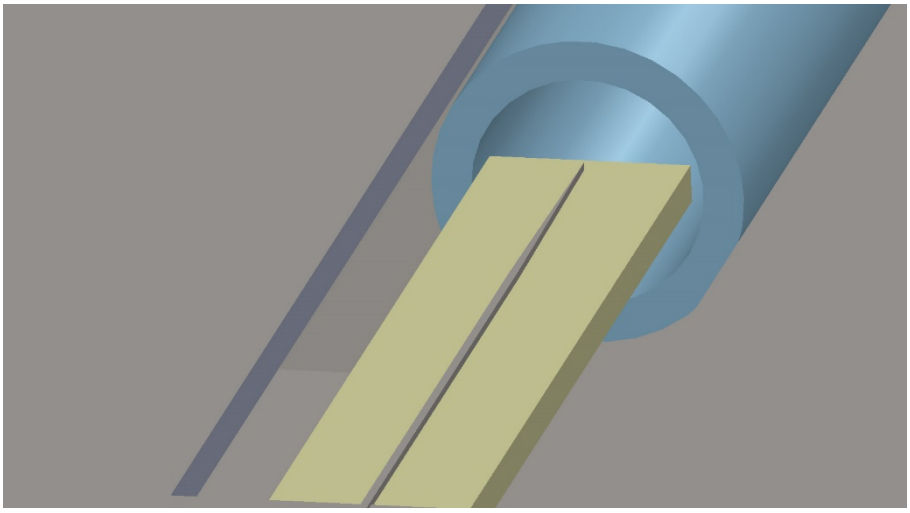

Figure 3.11 Diagram showing the potential Hollow Core fibre-to-chip inverted taper

Unfortunately, this approach became difficult to model successfully and so has not been explored in great depth. However, the concept remains interesting despite the fact that as yet the need for an interface between a hollow-core fibre and silicon photonic chip has not been substantial and hence this has not been pursued further into fabrication.

3.2.3 Laser Passive Alignment

One other hurdle for silicon photonics is the integration of a compatible light source or laser. If a suitable source can be integrated on chip, then this avoids many of the issues associated with packaging silicon photonic devices.

The most obvious barrier for useful source integration is the lack of an appropriate silicon based laser source[132], [133]. Whilst silicon is extremely useful in optoelectronics and photonics, especially considering CMOS compatibility, due to the nature of its indirect bandgap, it has proved extremely difficult to realise a

functioning electrically pumped silicon laser diode capable of fulfilling the requirements of most silicon photonic devices. The indirect bandgap leads to light emission by way of phonon interaction which greatly limits radiative electron-hole recombination. The low quantum efficiency of silicon ($\eta_i \approx 10^{-6}$ [133]) makes creating a laser or even LED source extremely challenging (see Paniccia et al. and Fang et al. for detailed discussion [132], [134]).

However the possibility of using III/V lasers presents a more viable alternative for the near future of Silicon Photonics[135]. Whilst III/V lasers offer a strong platform to fabricate useful sources, they introduce an extra engineering problem to tackle – *how can a III/V source be coupled with a silicon chip?* Again there have been many efforts to marry together the flexibility and functionality of silicon with a convenient and capable III/V source, a so called hybrid laser[136]–[138].

A variety of methods are possible, for example a vertical-cavity surface-emitting lasers (VCSEL) could be positioned vertically above a grating coupler, perhaps within some sort of etched capping chip cavity, however typical VCSEL's

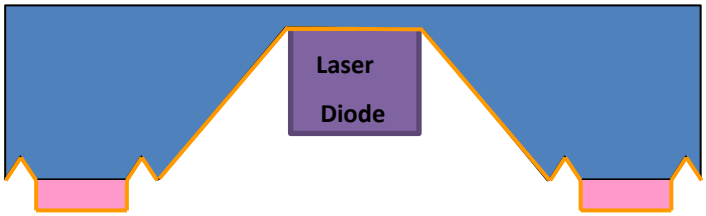

Figure 3.12 Diagram of a laser diode bonded to capping chip

work at wavelengths unsuitable for most forms of silicon photonics. Conventional butt coupling could be used to couple light from a laser emitting in plane with the waveguide.

There should be no reason however, why a similar approach cannot be employed within the packaging solution already discussed here, in fact there would be definite advantages if a passively aligned laser could be demonstrated. Figure 3.12 shows how this could work. Within the field of packaging more generally, passive alignment remains an issue when considering integrated sources in much the same way as it does with optical fibres. Therefore, if a III/V source can be

successfully passively aligned to a silicon chip then this would give great scope for a multitude of different device applications.

One method that has been successfully demonstrated is to use an adhesive to bond together an SOI wafer with a III/V die[139]. Divinylsiloxane-benzocyclobutene (DVS-BCB or simply BCB) is commonly used to achieve such a goal[140]–[143]. This would allow for heterogeneous integration of both materials. However, there would need to be extra steps introduced to the fabrication to allow for this to work. If the laser was to be positioned on the capping chip, then whatever metal that is deposited on the capping chip as a reflective mirror would almost certainly have to be removed from the bonding site in order to give a good adhesion between the die and wafer and avoid the risk of an electrical short circuit. Whilst this is certainly feasible, it adds an extra step meaning more complexity and would lead to greater costs and time for an industrial process. A laser could also be placed on the optical chip and this is another possibility to consider. With any arrangement, precision of alignment also becomes an issue. With specialist equipment a precision of under 3μm is possible[144] which should allow for acceptable alignment to a grating coupler although a smaller alignment error of ±1μm is far more desirable.

Flip-chip bonding has also been used successfully to bond III/V laser diodes onto silicon and SOI wafers[145]–[147]. This involves the use of particular solders which can make the process "self-aligning"[148], [149]. This technique could also provide the necessary requirements to allow a III/V laser diode to be bonded onto an SOI substrate in order to be packaged. This could be a useful method for assembling both chips but also could be useful for bonding any III/V laser chips as well. Also, once this packaging system in this work has been fully demonstrated, then in future it will need to be integrated with electronics in order make modulators and such things. Flip-chip bonding also provides a way of adding the necessary electronics chips into the packaging system.

One other possibility that is of interest is to use a eutectic metal, with surface activated bonding (SAB) in order securely mate both substrates. This has already been demonstrated using a VCSEL [150]. In this work a laser is bonded to a silicon substrate by plasma activation of an Au film on the substrate. This method could

potentially offer a way of bonding a laser to the capping chip, given the fact that an Au (or similar metal) film is to be deposited on the capping chip anyway to provide the mirror needed to reflect the light beam onto the grating coupler. This could also potentially be useful for electrical contacts given the high conductivity of Au. The trench which was originally conceived to house a circular fibre could be widened to create an angled trench with a flat bottom rather than the V-grooves seen previously. Everything else beyond this point would be done the same way as has already been shown. An extremely flat and smooth surface is required to achieve an acceptable alignment as was achieved by Higurashi et al. vis use of an Ar RF plasma[150].

3.3 Numerical Simulations

3.3.1 Simulation overview

Simulations have become a vital and common part of the design and testing of photonics devices[151], [152]. Simulations allow the fundamental physical equations which govern the behaviour of light to be computed numerically. Given Maxwell's equations, the photonic designer may predict how light will propagate through any given medium. With this they can then design waveguides and more complicated devices. Software packages that are available "off the shelf" can be used which are essentially Maxwell solvers with a user-friendly interface. By using a simulation software package all the photonic devices that are needed can be developed and numerically tested before even setting foot into a foundry.

This project used various different simulation techniques to design the different devices used. This chapter will narrate and explain the journey from an initial concept to creating the final device. Simulations are an important part of the design cycle and the important aspects of modelling will be highlighted and explained in the following sections.

3.4 Waveguide Simulations

The goal of this project was not to design and create a silicon waveguide – this has of course been done before. But silicon waveguides form the basis of virtually all silicon photonics circuits and consequentially are important in this project. Therefore, basic choices about waveguides had to be made and then a simulation could be implemented to confirm that they would be suitable for the project.

Rib waveguides were selected rather than strip waveguides. Rib waveguides are less susceptible to issues with side-wall roughness[153], are very low-loss[154] and are very useful for modulators[72], a key building block of all telecoms and datacoms.

The basic design method for rib waveguides was outlined by Soref et al. [74] back in the early 90s. These principles were then adopted when designing waveguides for this project. A commercially available mode solver, FIMMWAVE, was used for this process.

Figure 3.13 gives a graphical cross-section of the waveguide, showing how the fundamental mode is confined in the silicon rib, when oxide is used as a cladding material within an SOI wafer. Since silicon has a comparably large index, small sized waveguides can be used. These simulations also showed that this waveguide design would have a loss of 0dBcm^{-1} to two decimal places, so are theoretically lossless for TE polarized light. In practice, sidewall roughness from the imperfections in both the lithography and etching processes will introduce intrinsic losses, as can impurities in the substrate. Tight bends can be formed with this

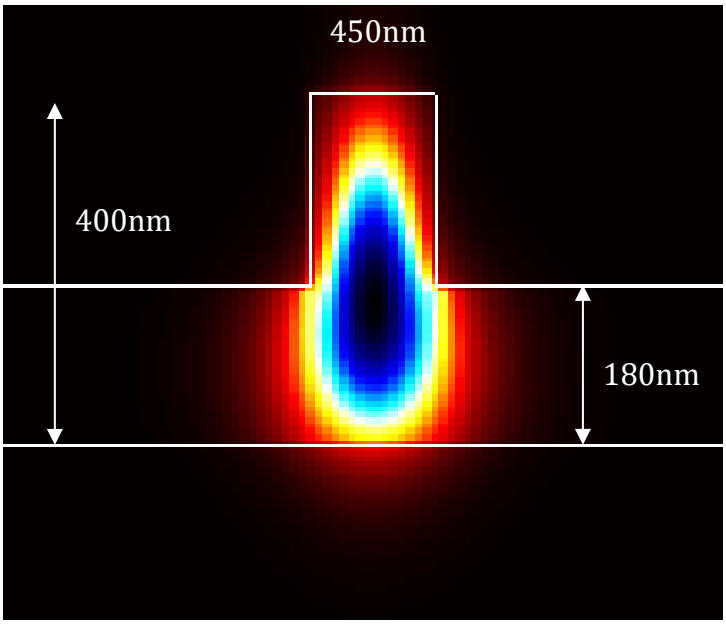

Figure 3.13 Image of a Waveguide Simulation showing the mode easily confined

waveguide, with mode-loss being negligible until a bend radius of <30µm is reached, as shown in Figure 3.14, which is much smaller than is needed for this project anyway.

These waveguides can also be easily tapered as well with a 500µm tapering length to a 10µm wide waveguide or 1000µm which can be used for a 14µm wide waveguide as used for this project. Tapering is necessary to couple waveguides into modes with a larger mode-profile, such as for an optical fibre. For this project it was

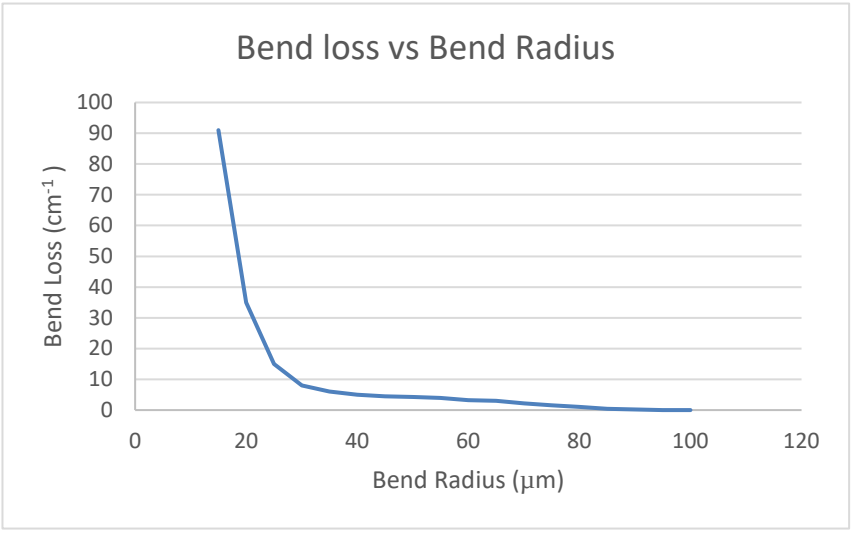

Figure 3.14 Graph showing the bend loss of the waveguide against radius

necessary for the grating couplers to adequately match the input fibre mode. A taper of these dimensions shows no loss due to tapering when modelled in this way. Of course, fabrication imperfections can induce a minimal loss due to tapering which is discussed in the results section.

3.5 Grating coupler design

The grating couplers for this project were a lot more complex than the waveguides due to the unique design constraints. The basic design of grating couplers can be seen in [155], with advantages and disadvantages given in [107][156][157].

Grating couplers are used as the basis of fibre to chip coupling in this project for the following reasons. Firstly, it finds standard use in industry and R&D meaning there is already a strong platform to work on. In addition, grating couplers are more resistant than edge coupling methods to partial misalignment and so will be better suited to a passive alignment approach, as evidenced in the previous chapter from the literature. Finally, surface coupling allows for far more in the way of wafer scale testing which can dramatically improve fabrication time, yield and costs. These three reasons make grating couplers the primary choice for this project.

Most grating couplers in the past have been designed to work in an out of plane configuration rather than in the scheme being proposed here, so there are a few peculiarities associated with designing customised gratings that are applicable to this packaging solution.

Firstly, due to the constraints of the wet etched mirrors and channels, the beam incident upon the grating coupler is required to fall upon the grating coupler at an angle of ~20° to normal, which is determined by the angle of the silicon crystal planes used to form the mirror etch. This is significantly different to the 7-10° which is more commonly used with grating couplers in silicon photonics. This then requires a totally different grating coupler that is designed from scratch, since the condition for a grating coupler to function correctly is dependent on several things; the wavelength of the light, the period of the grating, the refractive index of the cladding medium and the effective index of the grating coupler and finally the angle of the incident light. The relationship is shown in Eq. 3.1 [64].

$$\Lambda = \frac{\lambda}{N - n\,sin\theta} \qquad Eq.\,3.1$$

Here Λ is the period of the grating, λ is the wavelength of the light, N is the effective index of the grating, n is the refractive index of the cladding media (usually air so is often unity) and θ is the angle of the incident light to the normal. An approximate calculation shows that for typical values for a grating coupler, a change in angle of an additional 10° will amount to a difference in ideal period of around 40nm, meaning that it is necessary to design new gratings for specific intended angles of incidence. Eq. 3.1 was used to determine an initial estimate for the grating period which was then improved upon with finite difference time domain (FDTD) simulations[1], [158], [159]. Using a typical wavelength of 1550nm (transverse electric field light polarisation or TE polarisation) gives a grating period of 0.403μm as a good starting point which can then be refined with simulations.

The next variation was that the proposed coupling scheme uses backwards or reverse grating couplers. This is again due the constraints from the wet etched mirrors. This means that the incident light will be at an angle of -20° to the normal given the conventional way of defining such an incident beam. This also requires customised grating couplers. The efficiency of the grating couplers for this project are also important since there is a long propagation length between the end of the fibre and the surface of the grating. Initial estimates were of the order of 50μm for this distance. However, due to design improvements there is now scope to decrease this to around 30 or even 25μm as discussed previously. Over this distance however there is a considerable change to the profile of the light beam. This was ascertained by using FDTD simulations with proprietary software from Lumerical™ and then plotted for comparison. This comparison can be seen clearly by examining Figure 3.15 and Figure 3.16. Here there is a considerable change in both the peak electric field intensity but also of the power density. The field profile is spread over a much larger total area after 50μm of propagation in free space due to the beam divergence. Hence, all this needs to be accounted for when designing suitable grating couplers.

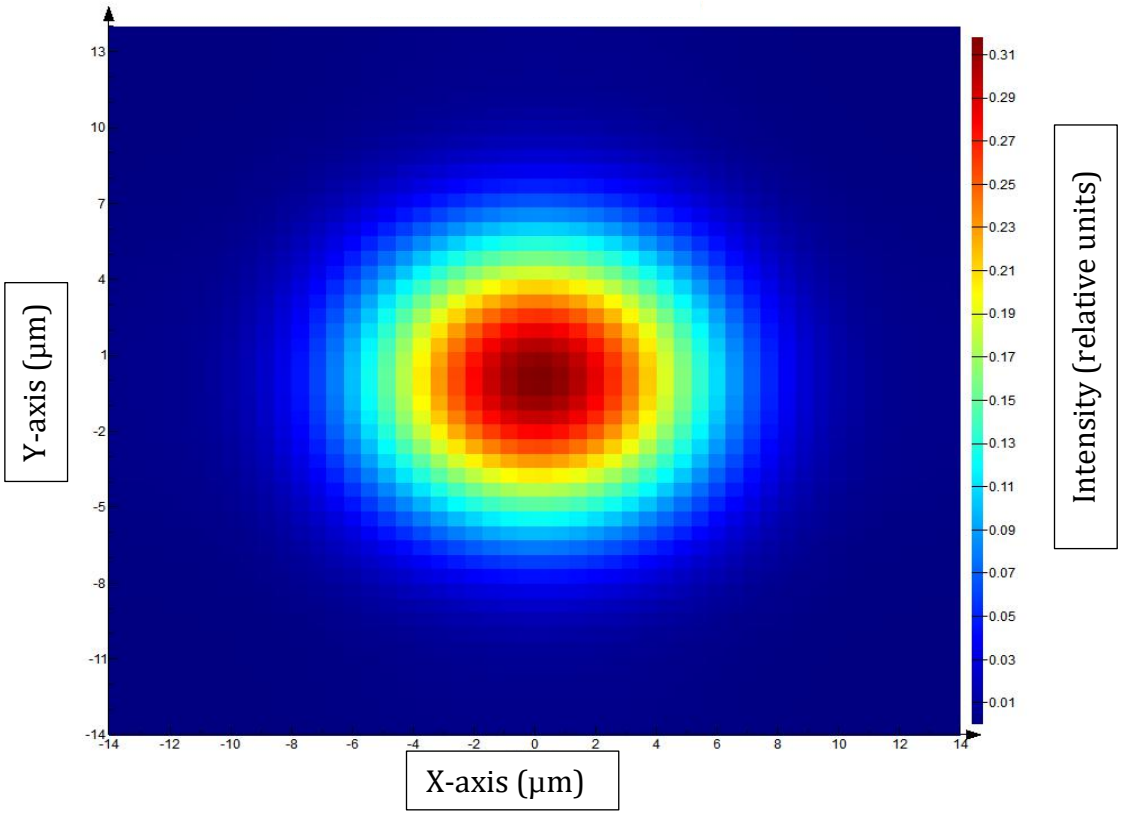

Fig. 3.15 Simulated Field profile after light beam has propagated 50µm in free space

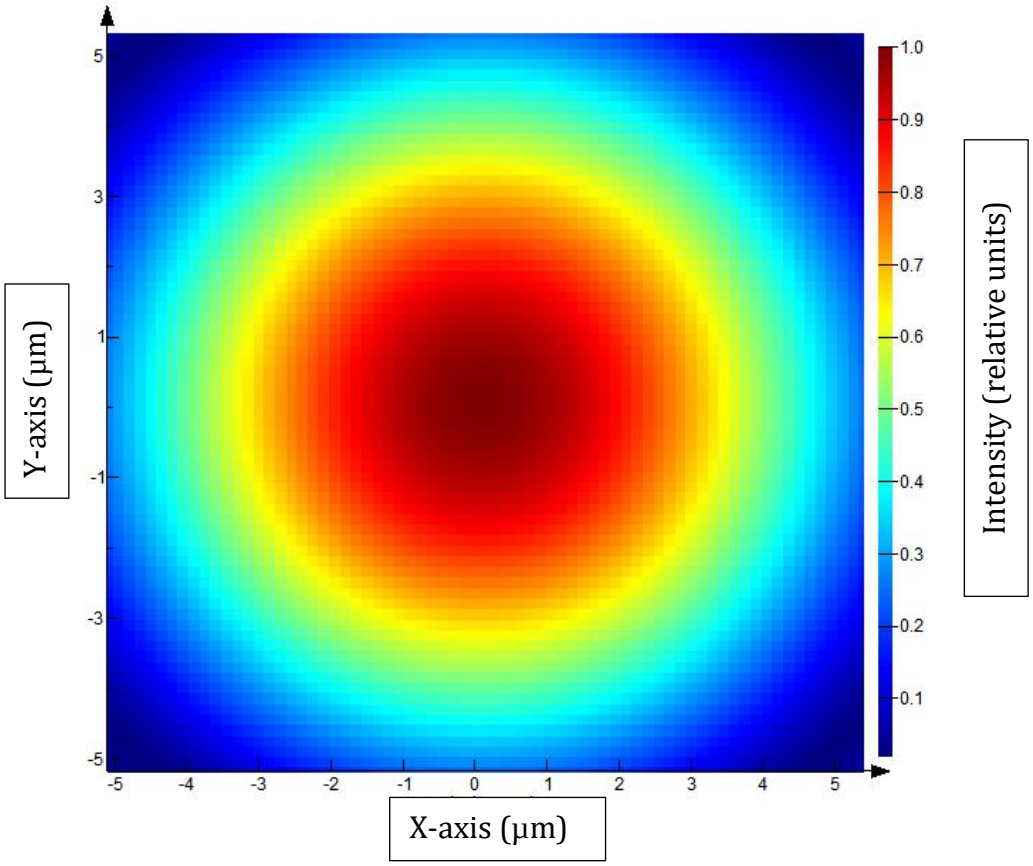

Fig. 3.16 Simulated field profile of emitted light at fibre end

The characteristics of the propagating beam were used in the simulation to give an accurate representation of the physical system.

Another consideration was the application of apodization to increase the efficiency of the grating couplers. The use of apodized grating couplers has been well documented [160]–[162]but has yet to be incorporated with the current packaging scheme as proposed here.

The model which is depicted in Figure 3.17 was used to incorporate all of the design features that have been mentioned so far, again within a Lumerical FDTD Solutions© environment; the source and chip can be clearly seen, showing the relative dimensions of the waveguide and grating. The period and fill factor of the grating was varied for the first five 'teeth' to give a non-uniform grating which will match as closely as possible the Gaussian profile of the incident light beam. The fill factor, which is the fraction of each period that is '*filled*', for the rest of the grating

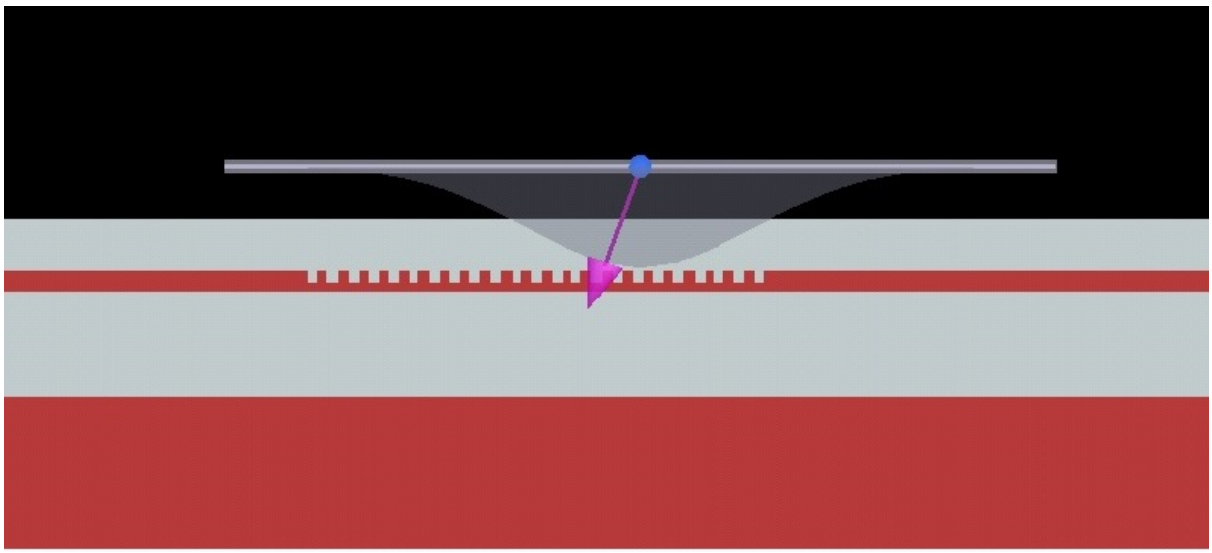

Figure 3.17 Side on view of grating coupler model within simulation environment

was optimised to the central wavelength of 1550nm in order to give the greatest transmission of light into the waveguide. The width of grating was increased slightly to 14μm so as to accommodate the increased area of the incident light beam, the justification for which can be seen in Figure 3.18, where the grating efficiency is only increased extremely minimally above this width, but device footprint would be negatively impacted. A waveguide height of 400nm was selected as this is easily

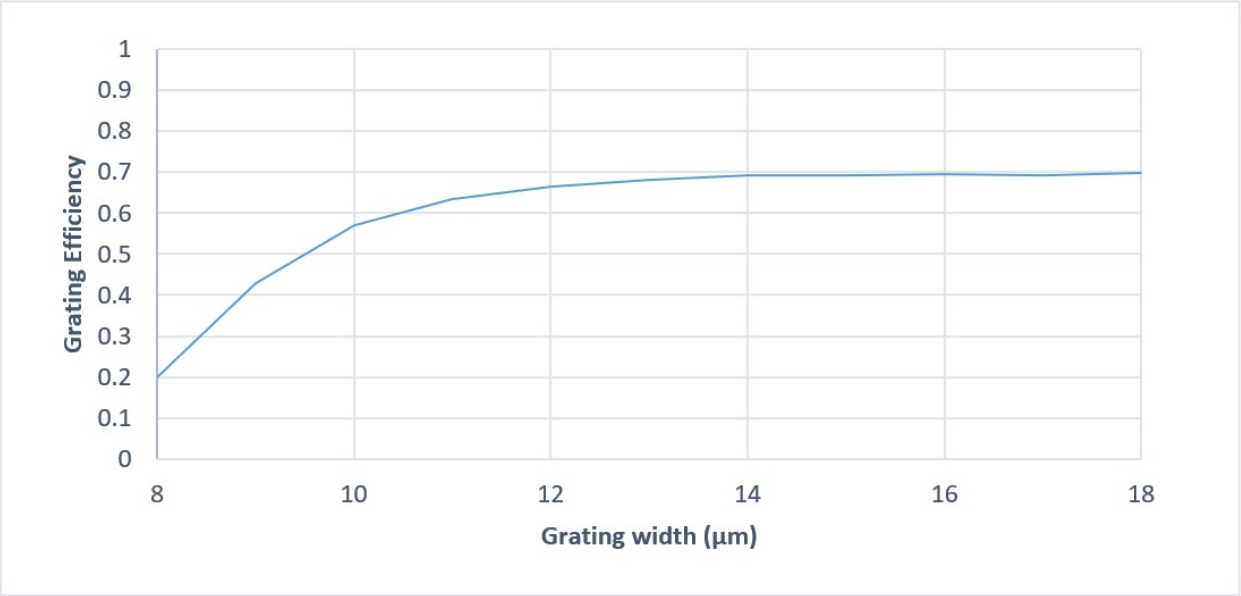

Figure 3.18 Graph showing grating width against efficiency

fabricated using conventional silicon-on-insulator (SOI) wafers. The waveguides were 450nm in width with an etch depth of 220nm.

A base grating period of 0.490µm and duty cycle of 0.418 was used following investigations with computer simulations. This period is shifted significantly from the original estimate of 0.403µm, showing that the estimate for effective index needed to be refined by utilising the specific parameters of the packaging system. These parameters were paired with the first five teeth of the grating being, adjusted to match the Gaussian profile of the incident light beam and so this incorporates a graded change in wavelength and fill factor to manipulate the grating efficiency. This was deduced by using a "*particle swarm*" optimisation algorithm, which is an iterative approach to optimising a given parameter (in this case the transmitted power). A brief introduction to this method shall now be given. Particle swarm optimisation was first inspired by observing patterns of behaviour with wildlife and was initially developed by Kennedy and Eberhart[163], [164]. Particles are used to model the characteristics and are put onto a 'space'. This space will differ depending on the constraints of the particular optimisation. In this case it was the grating period and fill factor. The maximum value will be located on the search space and the particles will "swarm" to the maximum value. This form of machine learning enables an algorithmic code to find a maximum solution to any given problem. For a deeper explanation see Poli et al.[163]

A central wavelength of 1550nm was used as is commonly employed within silicon photonics and communications. At the time of writing the total transmission efficiency of the grating coupler in simulation stands at 65% or a coupling loss of 1.87dB; as seen in Figure 3.19. The 1- and 3-dB bandwidths can be seen to be 40 and 64nm respectively. These results are comparable to others such as Antelius et al. [162], Chen et al. [160] and Li et al. [165].

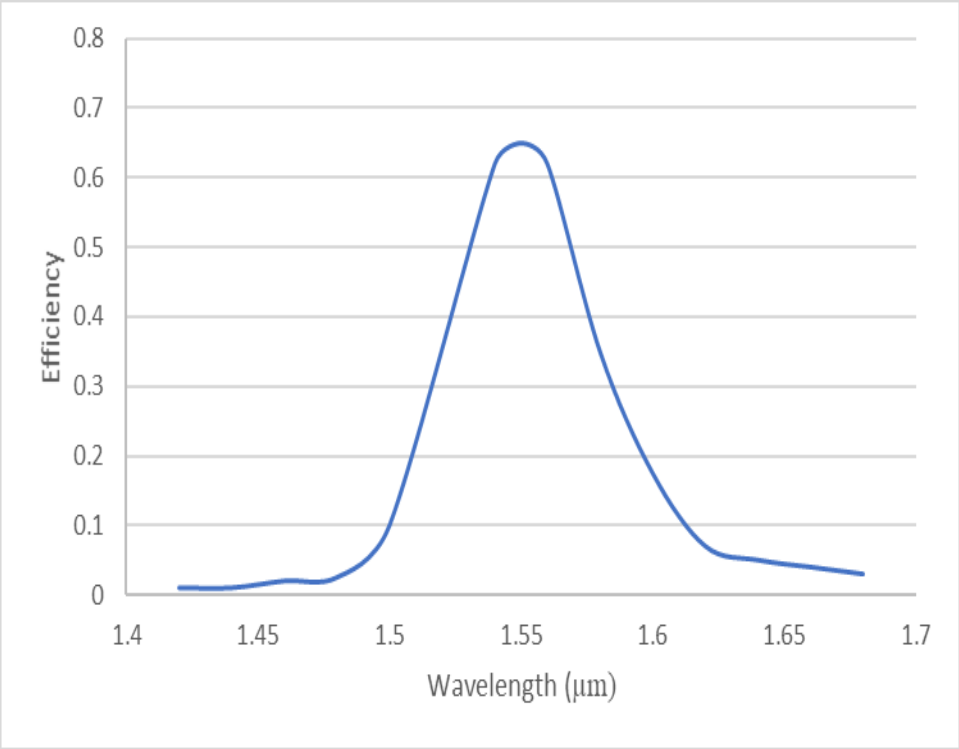

Figure 3.19 Transmission spectrum of simulated grating coupler

Potentially there is the possibility to vary the buried oxide thickness as well as the etch depth which should allow for further gains in efficiency[162], however this then affects the cost and complexity of fabrication negatively. At the time of writing it seems that these improvements are unnecessary complications to this project, but could be implemented in future.

Since initially developing this grating, the design has been modified in order to account for the potential use of an index matching epoxy to fill the void between the fibre and the grating. An epoxy can be chosen which has a refractive index equal (or close) to that of the glass fibre. Therefore, propagating light will only be effectively moving between two media rather than three. This will of course create the need for a different grating and hence all the principles stated so far have been

applied to this new case. Figure 3.20 shows how the relationship between the grating period and efficiency after adding in an index matching epoxy is changed. This graph shows how the peak efficiency has both slightly increased but also the

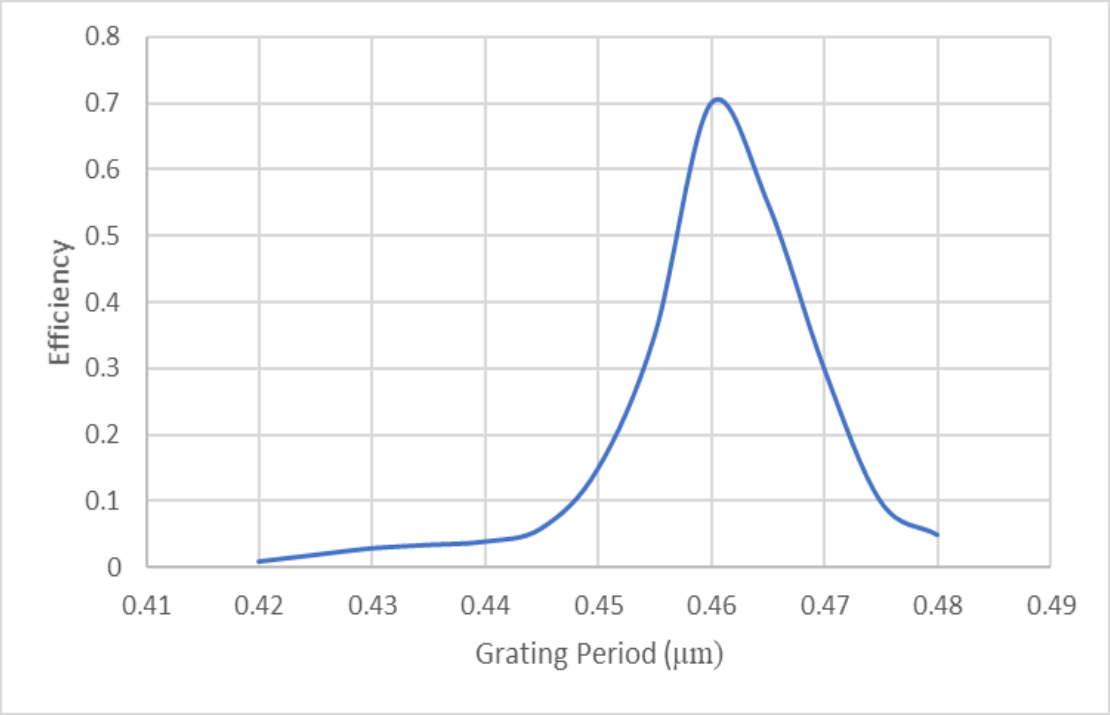

Figure 3.20 Graph showing the relationship between grating period to peak efficiency

corresponding period at which this peak is obtained has shifted significantly. This graph is plotted for a wavelength of 1550nm.

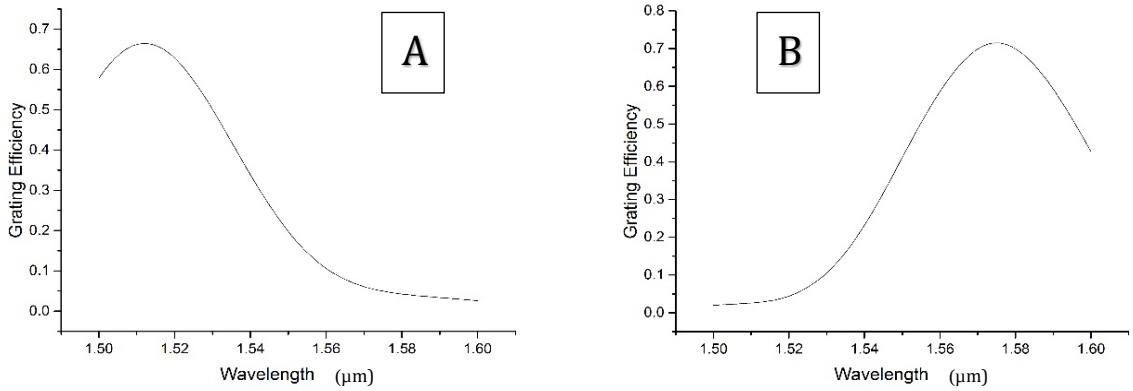

Figure 3.21 Two graphs showing a comparison between the same grating with a fill-factor of 0.3, A and a fill-factor of 0.6, B

Figure 3.21 demonstrates the dominant effect that the fill-factor has on this particular grating. By comparing Graph A with Graph B, the effect of differing fill-factor can be seen because Graph A and B both are taken from the same grating coupler, differing only in the fill-factor with Graph A being 0.3 and Graph B being 0.6. Therefore, it can be seen that the dominating effect of varying fill-factor is a shift in the wavelength of peak efficiency, while the efficiency itself is only affected very slightly. Figure 3.20 is a graph which shows how the peak efficiency varies with the grating period and keeping other variables constant. By amalgamating Figures. 3.20 and 3.21 together for a grating coupler with an index matching epoxy then the parameters which yield the greatest efficiency grating are a period of 0.460μm with a fill-factor of 0.515, in order to be centred on 1550nm. The simulated performance of such a grating can be seen in Figure 3.22. From this it can be seen that the addition of the index matching epoxy actually increases the peak efficiency of the grating coupler as well. The peak efficiency has increased to 71% which is an increase of 6% from the figure quoted previously. The 1- and 3-dB bandwidths are slightly reduced in this configuration falling to 35 and 59nm respectively. Whilst this is lower than before it is still comparable to other results in the literature, (see previous pages).

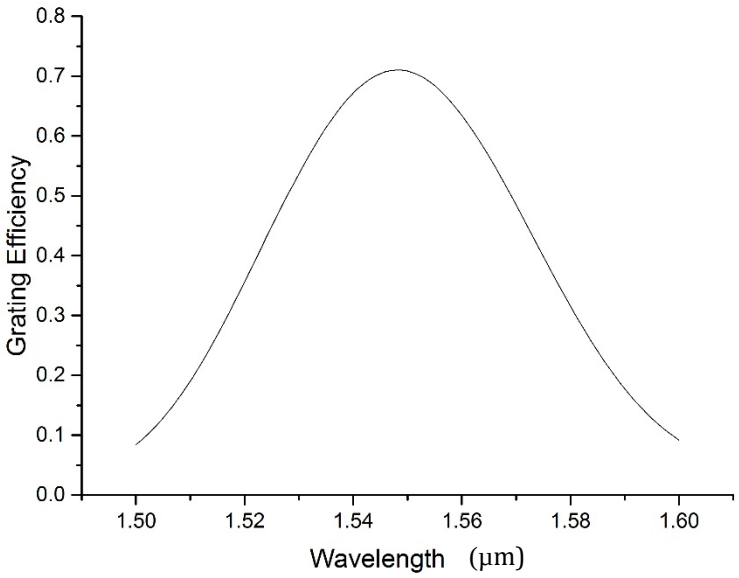

Figure 3.22 Graph showing the simulated performance of grating with index matching epoxy

Computer simulations can also be used to estimate how the grating coupler will perform when the fibre is misaligned in space. Figure 3.23 shows a heat-map displaying the performance of an apodized grating coupler with a range of x/y values around the centre point of the grating. The grating and waveguide would be positioned parallel to the y-axis in this diagram. Whilst the efficiency of the grating remains surprisingly high in a lateral direction, it is comparatively poor when there is a misalignment in the y direction by more than 4μm. It is likely that the apodization is responsible for this behaviour. Whilst apodizing the grating will allow it to achieve a greater peak efficiency it will also create a positional "sweet-spot" to which the fibre must be precisely aligned to in order for the grating to operate at maximum efficiency. Consequentially, for the sake of this project, apodization was postponed so as to give more flexibility to misalignment, at the cost of peak efficiency.

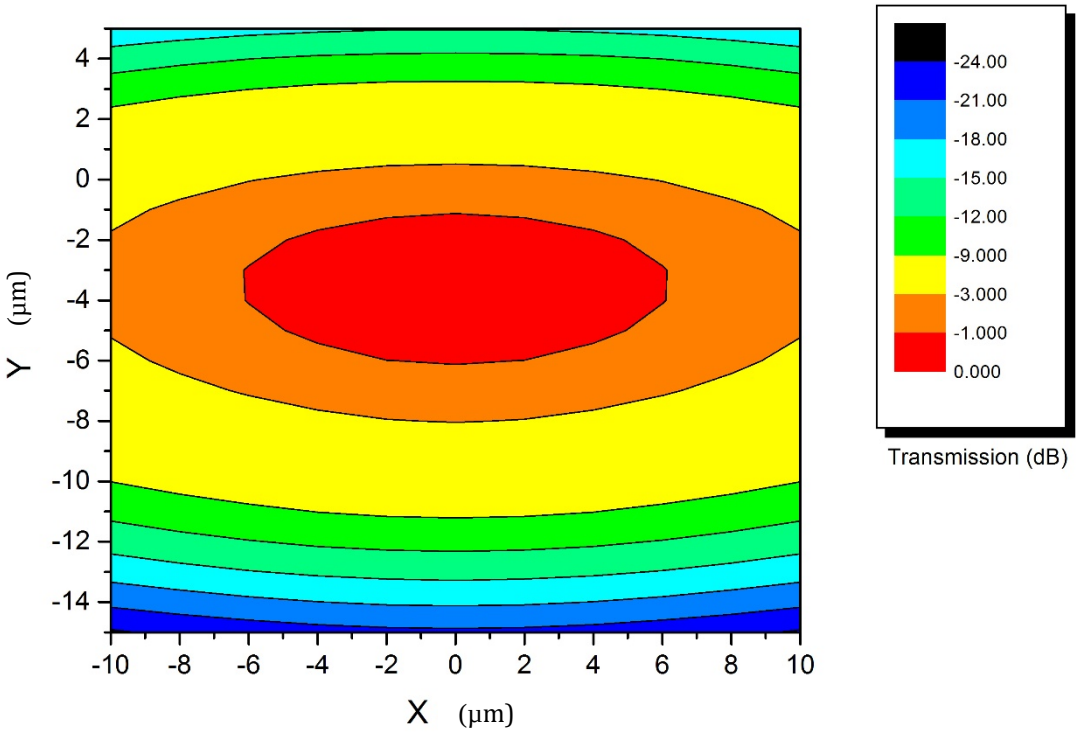

Figure 3.23 2D plot of the transmission of the designed grating coupler with respect to displacement in x and y axes

One further modification to the current grating design would be to incorporate polarization independent grating couplers. The current proposed grating design is tailored for TE polarised light as is typically used in silicon photonics, but this could potentially be expanded to allow for TM modes to be coupled into waveguides. This could create extra data channels which could increase the volume of data which a silicon photonic device can transmit or receive.

3.6 Inverted Taper Simulations

In addition to the grating couplers, another proposed coupling solution was to use an inverted taper. The reasons and rationale for this is explained elsewhere in section 3.2.1. Here a short discussion of the simulation process will be given.

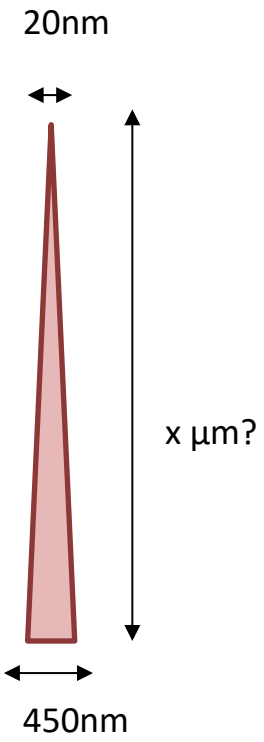

Figure 3.24 Diagram of the problem of making an inverse taper

Figure 3.24 shows the problem that is to be solved. A taper tip of 20nm is chosen since it is very small but still feasible with respect to fabrication. Ideally an infinitesimally small tip would be chosen but this is not practical so 20nm is used as a compromise, which gives minimal reflections. The length of the taper then needs to be computed. Non-linear tapers are also possible but were not considered here due to complexity but once designed, they could be incorporated into a photomask and fabricated relatively simply. This taper is used to couple light into a strip waveguide rather than a rib, but a subsequent taper section can be used to couple

into the typical rib. Both FDTD and EME (eigenmode expansion) approaches were used to optimize this taper with their results agreeing.

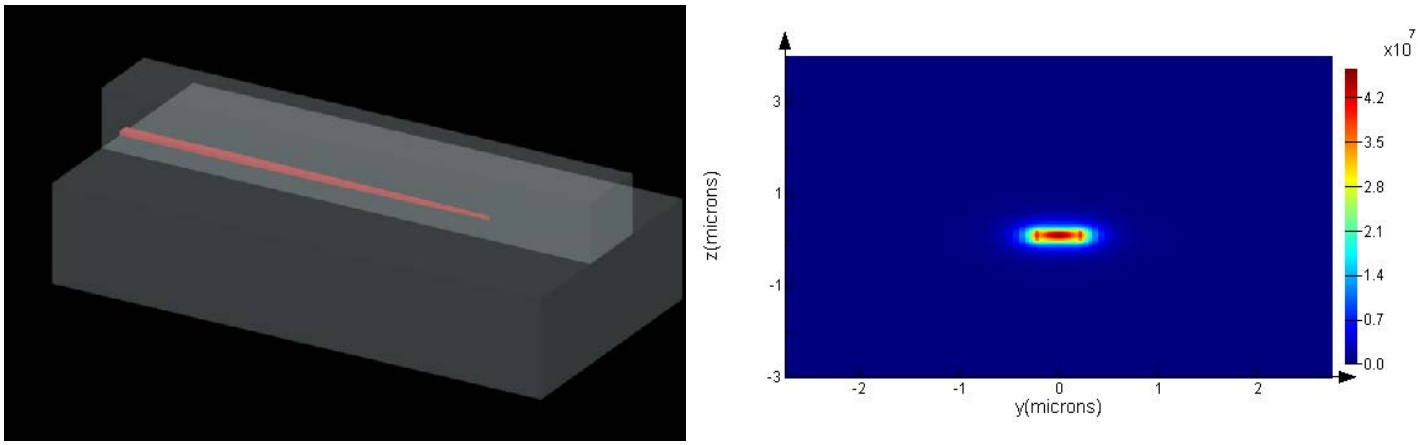

Figure 3.25 Image of Inverse Taper within Lumerical FDTD. The image on the righthand side shows the mode confined within the waveguide after tapering

Figure 3.25 shows the inverse tape within the Lumerical design environment and also the resultant mode within the waveguide, showing that the mode can be squeezed down to the required size.

In addition to this, Figure 3.26 shows a plot of taper efficiency against length. As can be seen the taper efficiency increases up to around 400μm, where the curve

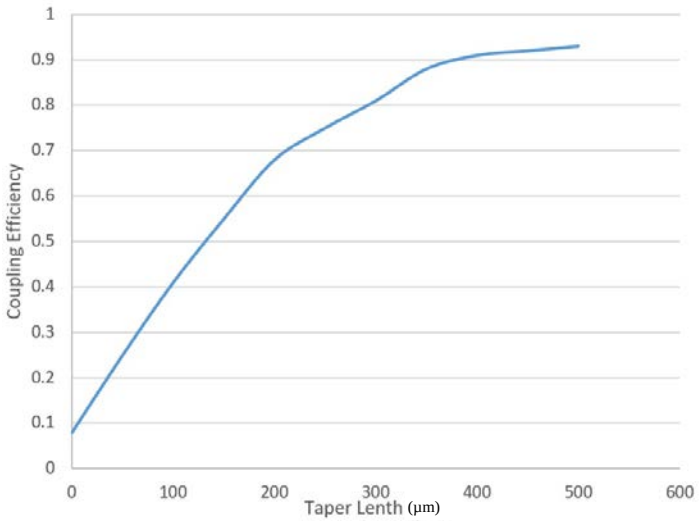

Figure 3.26 Graph showing coupling efficiency against taper length for the simulated inverse taper

begins to flatten. At the other end of the curve the minimum taper length of 0μm (butt coupling) shows an efficiency below 10%. This simulation shows that a

theoretical efficiency of over 90% or 0.45dB coupling efficiency is possible with a taper length of 500μm or above. Of course, the critical parameter here is alignment accuracy. These values will fall very sharply when any amount of misalignment is introduced. This simulation assumes a perfect alignment in all six axes.

Chapter 4 Mask Design Overview

This chapter details the process of how the mask was designed for devices that were used for this project. The design of the mask underwent several changes as the project developed. These are highlighted and explained in this section.

All the device masks that were written were done using a software package called L-edit[166], from Mentor Graphics. This program was chosen because it allows devices to be built in a hierarchical manner. So individual elements (such as a waveguide) can be drawn in a 'cell'. Then several cells can be instanced alongside each other to create a device (grating couplers, tapers and waveguides for example). Indeed, in this way, a whole chip or even a whole wafer can be written as one cell, made up of many smaller cells – in very much the same way a biological organism is made up of specific cells which make up organs.

Each element was scripted using C++ such that individual parameters could be adjusted from a script rather than having to manually change each one. For example, if the period of a grating was to be changed, this would be extremely cumbersome to do manually, since each etched area of the grating would need to have its dimensions adjusted. With a script this can be changed in an instant across a whole cell. The script for each section was stored in an online library and can be seen in the appendix of this document.

4.1 Device chip mask design

The mask for the device chip needed to be altered at a number of different points for a variety of reasons. This section will take each of these changes in turn.

Figure 4.1 shows the original mask design for the device chip that was designed in conjunction with Dr Reynolds. The left-hand side of the diagram shows the fibre trenches, and this is where the fibres would sit when assembled. The light from the fibres is coupled into a grating coupler and a short U-bend is used to return the

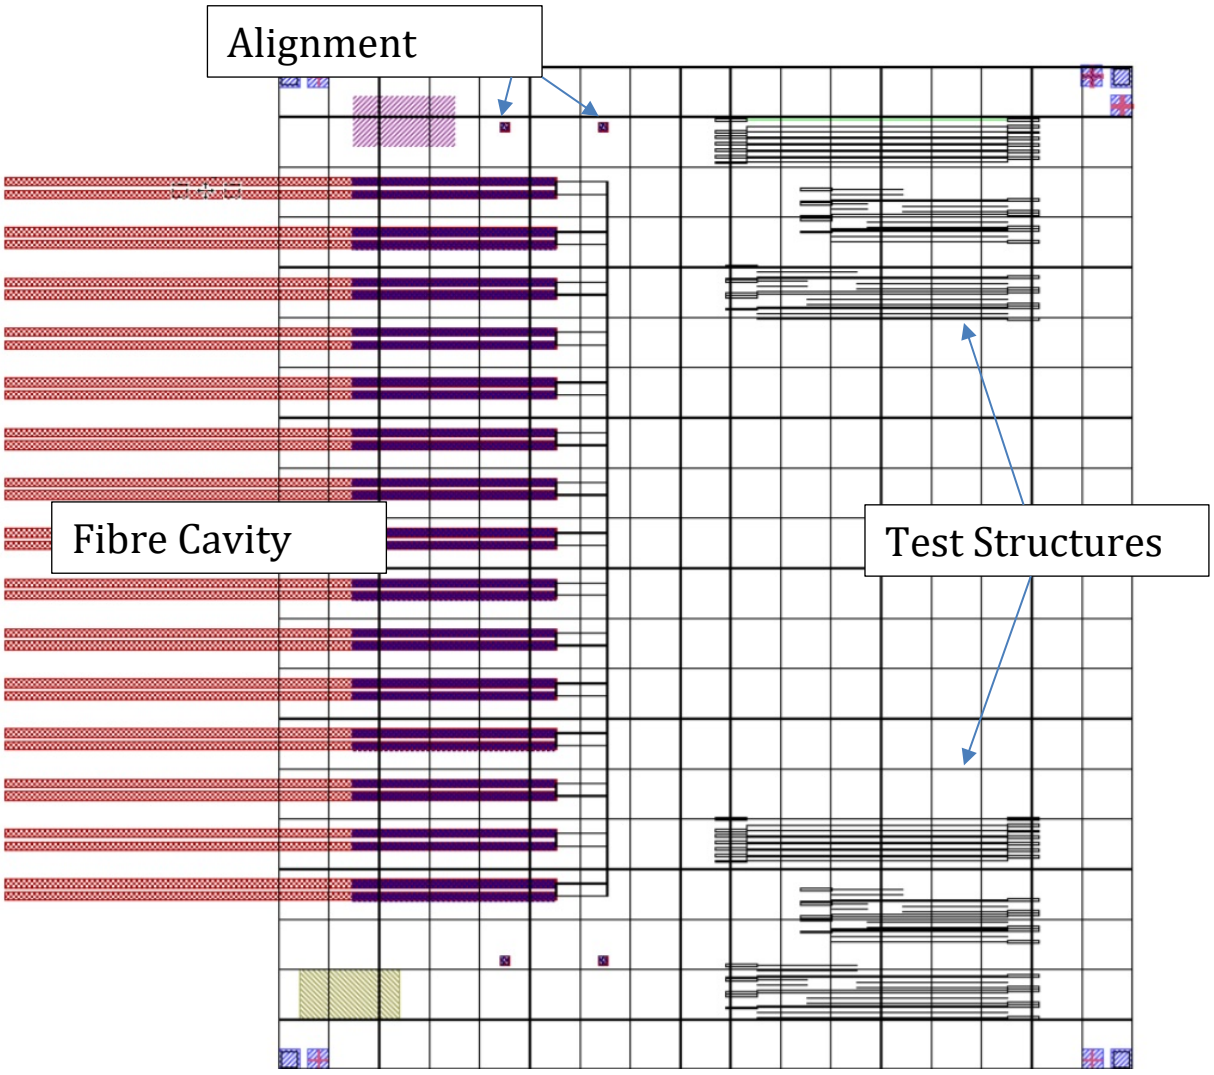

Figure 4.1 Mask layout of the original Device Chip

signal into the adjacent fibre. Thus, the array is made up of a series of fibre couplets. The black lines on the right-hand side of the diagram show the measurement structures. These are a series of test structures that can be used to measure the intrinsic, bend and taper losses of the devices and can be used as a reference point when assessing the performance of the chip. More detail on these is given in 7.2.1 The very small purple squares around the fibre cavities are the initial plugs that were intended to fasten the capping chip to the device chip.

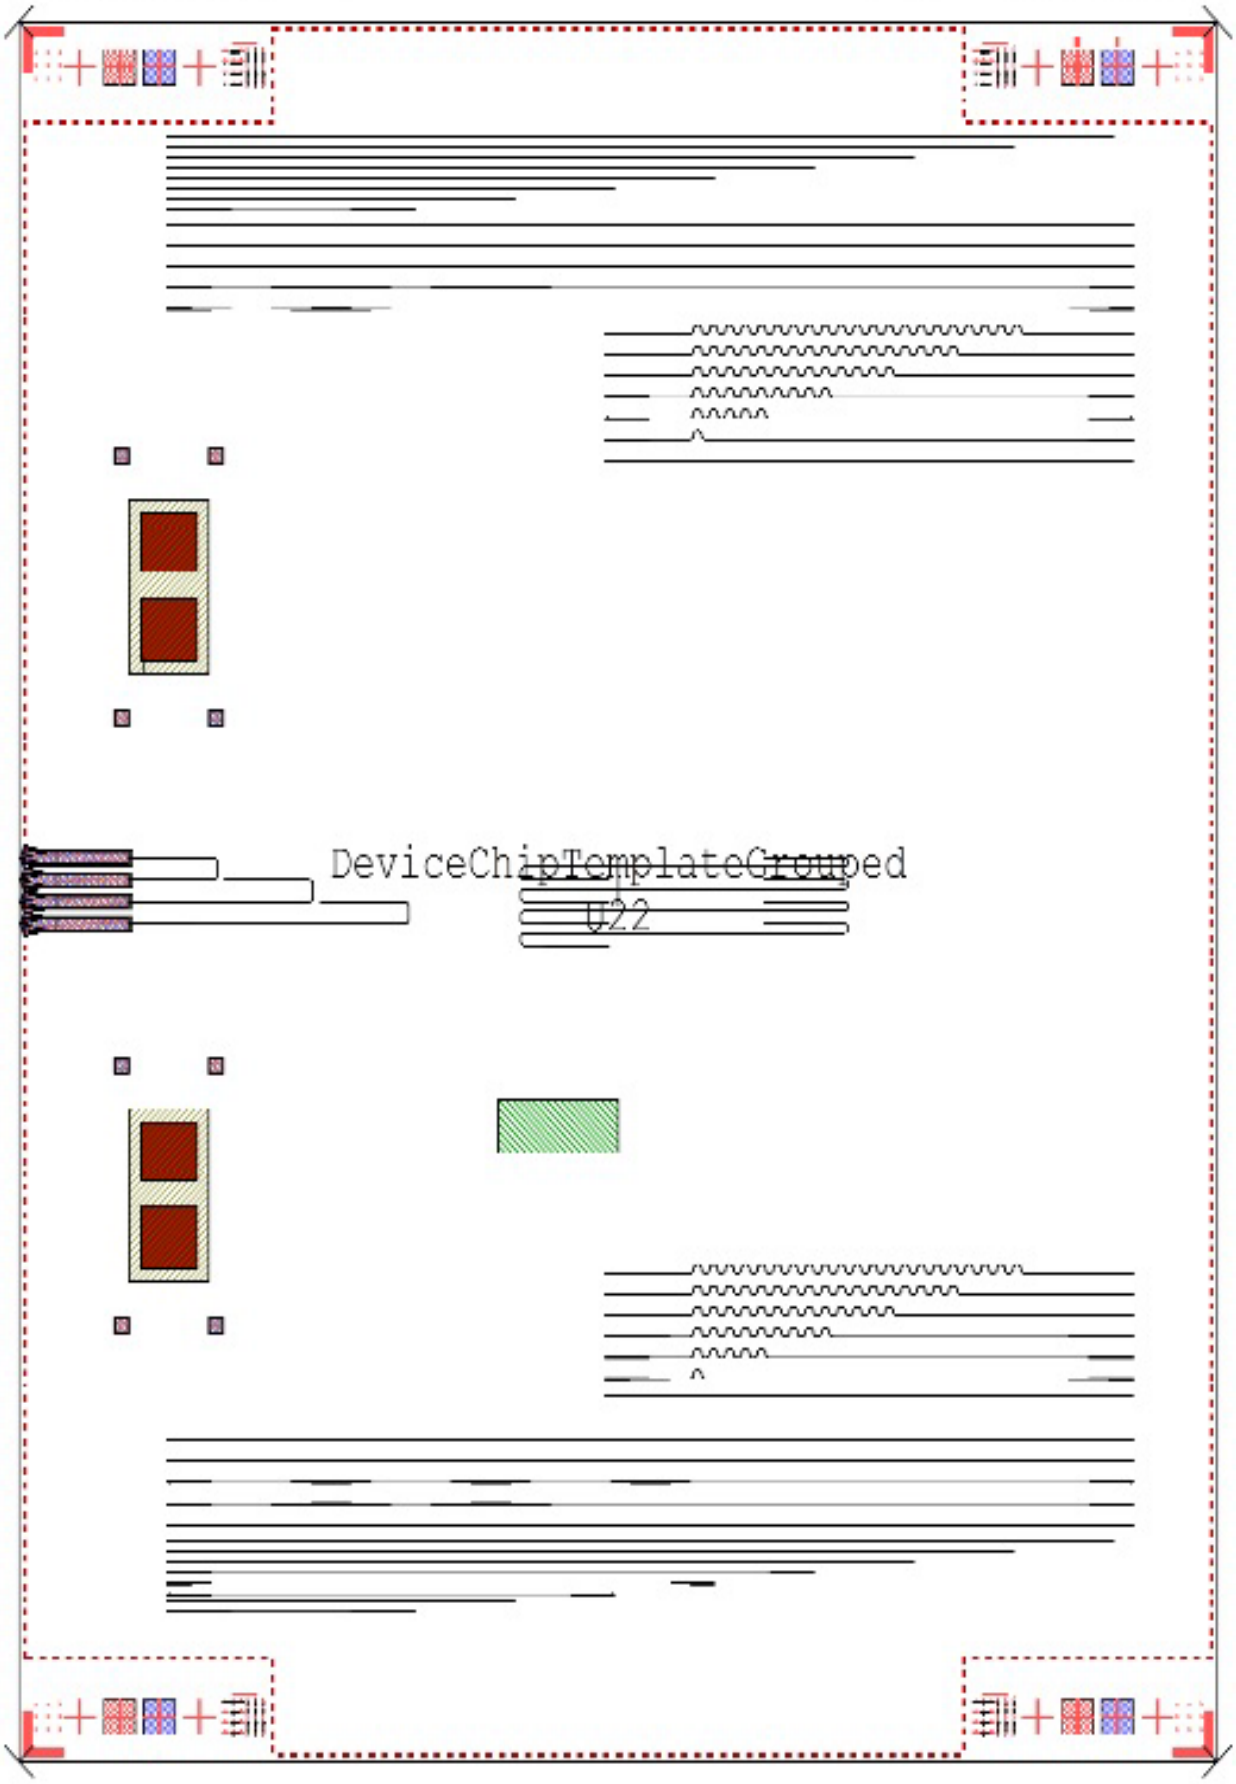

Figure 4.2 Mask of final Device Chip

There is a lot of redundant space on each chip. Of course, when this packaging method is used on real chips, the chip designer may fill this space with whatever devices are needed. The other thing to notice is just how small the plugs are in relation to the rest of the chip in Figure 4.1. The stresses and strains that would build up around these plugs would be enormous in proportion to their size. Another thing to point out is the size of the fibre array – 30 fibres. This is an unnecessary quantity for this project, and only increases the complexity of alignment. This mask design was altered 8 times during the course of the project so in order to make this discussion as a straightforward to understand as possible, only the final mask layout will be compared to the original version, which shows all of the modifications that were made along the way, during the evolution of the mask design.

Figure 4.2 shows the final mask design. A few things to point out include:

1. The plugs have been removed and instead larger adhesive holes have been drawn, with smaller holes that are used for alignment spheres.
2. The number of fibres has been reduced to just four – simplifying the alignment process, but allowing proof of principle, with subsequent scale up to more fibres.
3. Two of the fibres are now joined to non-packaged surface grating couplers such that the loss of one packaged fibre can be measured against a control.
4. The fibre cavity geometry has been significantly changed and now exhibits a ‘fluted’ end.

Each of these changes will now be discussed.

4.1.1 Adhesive holes

Figure 4.3 shows a close up of the adhesive holes and the sphere holes. The cross-hatched squares around the outside is where the spheres from the capping chip sit, and the brown larger squares are holes which are etched through the wafer and can be used to apply the adhesive epoxy. The hatched rectangle which these sit inside, is the outline of the corresponding hole from the capping chip which is given for reference- this hole is larger to allow easier access of epoxy but also to give a larger surface to bond to, forming a plug. A series of tests were carried out to determine

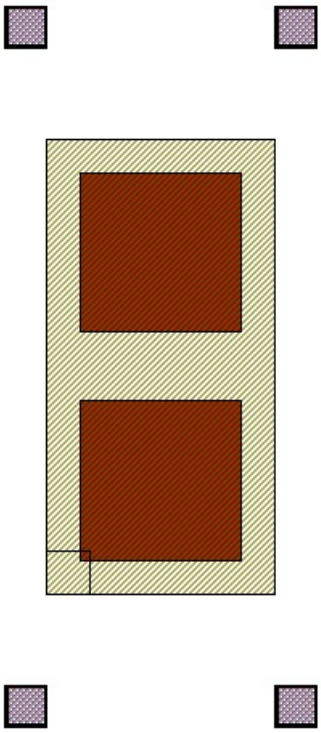

Figure 4.3 Close up mask of alignment structure

this configuration for the project. When deciding the size and shape of these structures, a trade off must be made. There is a trade between strength of adhesion and chip footprint. The other factor to consider is alignment. At least two spheres are needed to fix the capping chip in the correct orientation, otherwise the chip could rotate around the sphere. Of course, the more spheres the better, but again this comes at the cost of footprint and complexity. Whilst the sphere itself is not large, it would interrupt the course of waveguides on the chip such that they would need to bend around them, and this would make device design very difficult and so

this is something to avoid. In the final design, the capping chip was designed to only take up 1/3 of the area of the device chip, such that as much of the device chip could be left free for the devices as possible.

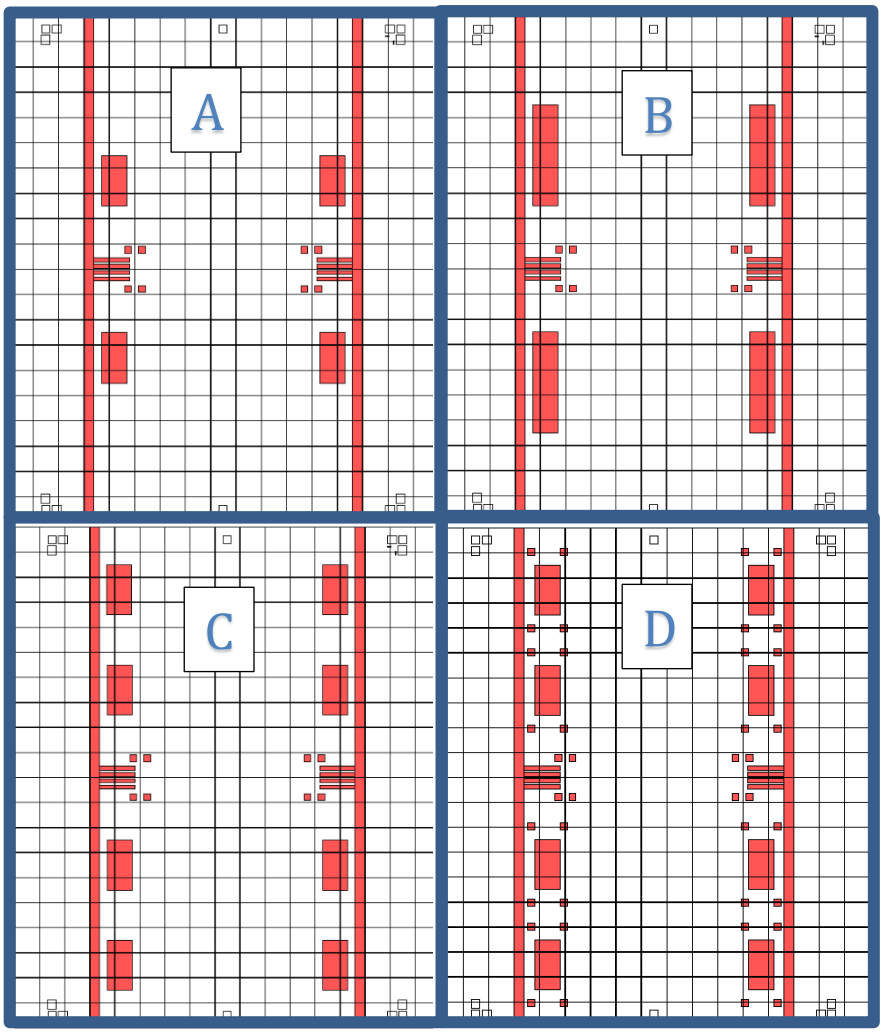

Figure 4.4 Masks of different sphere and hole varieties

Different sizes of adhesive plugs were used, however when tested to failure, it was determined that above these dimensions, the increases in bond strength were negligible. Figure 4.4 shows the different proposed shapes and sizes that were suggested. Cell A, B and C had different lengths and as such different adhesive areas, with the hypothesis being that increased bonding area would allow for a greater chip bonding strength. The first 3 chips only had spheres around the crucial fibre trenches. Cell D shows the addition of more spheres, around the adhesive holes. The sheer tests (explained in 7.2.2) showed that a combination of cell D and A were the best compromise, since the overall adhesive strength of A was still more than

enough for the purposes of this project, whilst the extra spheres that were available in Cell D aided in alignment. The spheres around the trenches were then removed when the fluting was introduced. Quantitative results of these test are given in the results chapter.

4.1.2 Reduction of Fibre Cavities

This was quite straightforward; the number of fibre cavities was reduced to four. The principle of passive alignment can be demonstrated with a reduced number of fibres in the array. Of course, in production, a larger array is to be preferred but for now only 4 can be used as a proof of concept and vastly simplifies the alignment procedure.

4.1.3 Addition of surface grating couplers for testing

This design also incorporates a packaged fibre that leads to a non-packaged grating coupler which allows the performance of the package to be better assessed. This allows for easier characterisation since a known grating coupler can be used with it as a control, but also one which is easier to use with standard lab equipment.

4.1.4 Fluted Fibre cavity end

The end of the fibre cavity has been changed significantly to incorporate a ‘fluted’ end since it resembles the flared end of a horn or flute. This was decided alongside consultation with Optocap and improves the design from an assembly point of view. The wider opening allows fibres to be inserted much more easily and would allow them to self-locate into the fibre cavity. This modification is also carried out in conjunction with the Capping Chip chamfered edge (explained in the next section).

4.2 Capping Chip mask design

The Capping Chip underwent a series of changes in much the same way that the device chip did. Of course, many of these changes were to compliment changes that were made on the device chip (such as hole number and location), but there were also a series of innovative adjustments made to the design of the capping chip which were unique.

As before, rather than go through each and every mask design, the original and the final design will be compared and a few key stages in between also described, which best illustrates the evolution of the capping chip mask.

Figure 4.5 shows the key features of the design. It has the trenches for the V-grooves, holes for the plugs and little additional detail. Several key features have changed. Figure 4.5 shows both the wafer layout, and an array of 4 chips side by side.

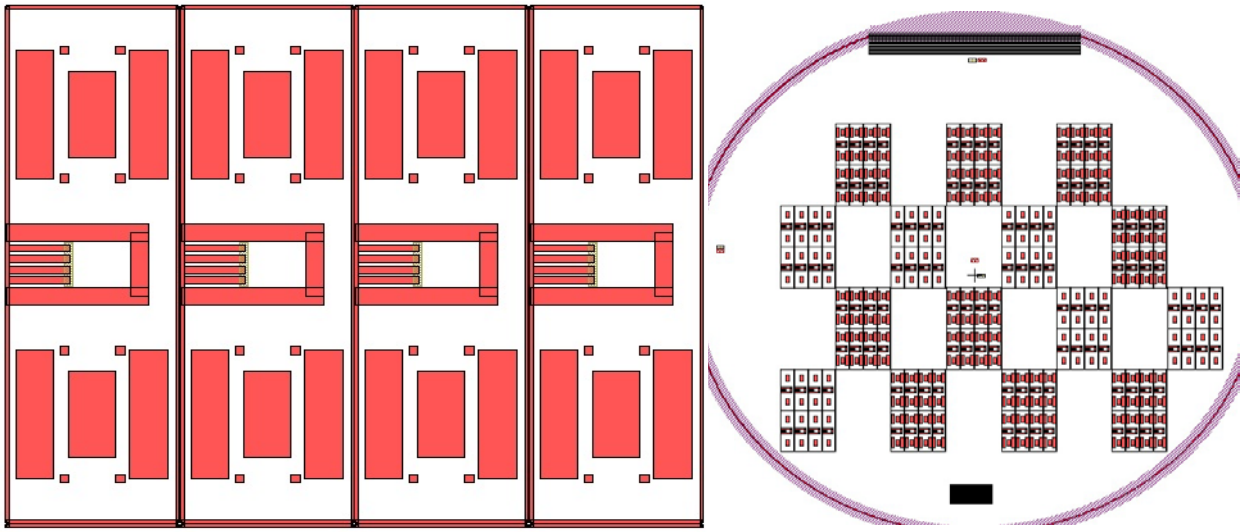

Figure 4.5 Mask of final capping chip layout at both chip and wafer level

There are several features that are different from the original design that will be discussed here. These are:

1. Chamfer around edge
2. Epoxy well around fibre cavities
3. Gold area reduced to critical area
4. Surface area reduction areas or recesses
5. Chip size reduction.

Each of these features will be described and an explanation given as to what benefits they bring to the design.

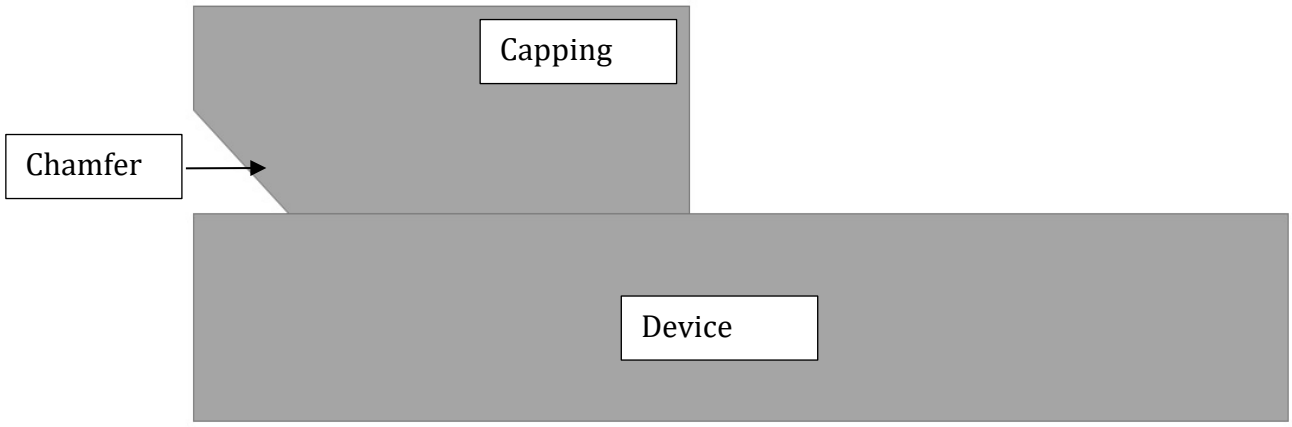

Figure 4.6 Diagram of chamfered edge of capping chip

4.2.1 Chamfer around edge

As can be seen, there is an outline around each chip, which is etched in conjunction with fibre cavity V-grooves. As this is a KOH etch, it will form the angled sidewalls in the exact same fashion as the V-grooves. Its purpose is not appearance, but rather it aids the assembly process. The decision to add the chamfer was based on consultation with Optocap and benefits the design in two main ways. Firstly, around the fibre V-groove it opens the channels out further in the same manner as the fluting. But it will also aid the adhesion process, giving a lip which can be sealed with an epoxy around the chip, ensuring a hermetically sealed environment inside the chip and a strong bond around the edge. Figure 4.6 shows a side-on image of how this works.

It is worth showing the intermediate design of chamfer that was rejected. This can be seen in Figure 4.7. The features are coloured blue here to illustrate the difference between this chip layout and the previous. In this version, the chamfer was elongated down the entire length of the wafer. Each chip was aligned in a column as can be seen in Figure 4.8. This introduced a serious problem when etching the chamfer. The groove that would run down almost the entire length of the wafer

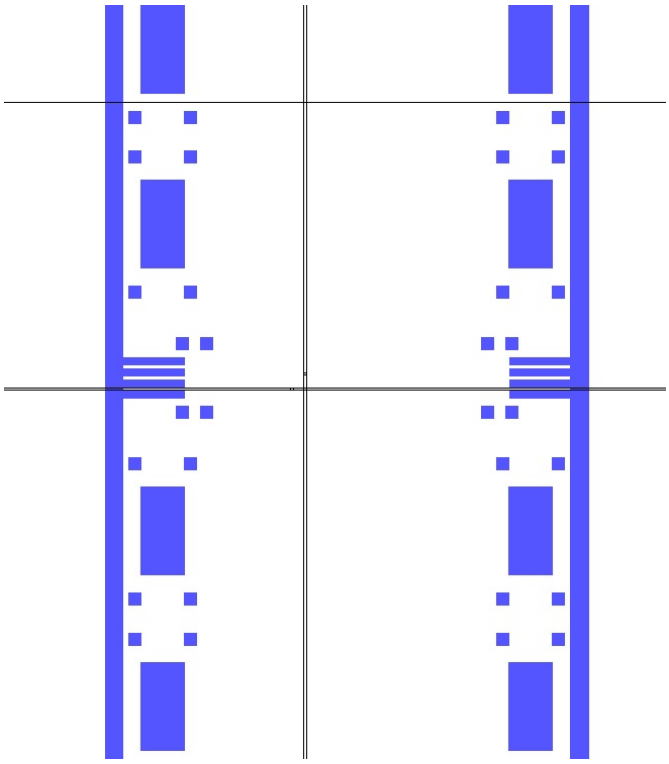

Figure 4.7 Intermediate mask, with elongated chamfer

would make it extremely fragile and was liable to crack. This happened on a number of occasions, hampering development. The chamfer was also only located along one edge of the chip in this design. It can now be seen why the wafer layout in Figure 4.5 was chosen. This newer design maximises wafer space whilst leaving the wafer strong enough to endure the processing that is necessary.

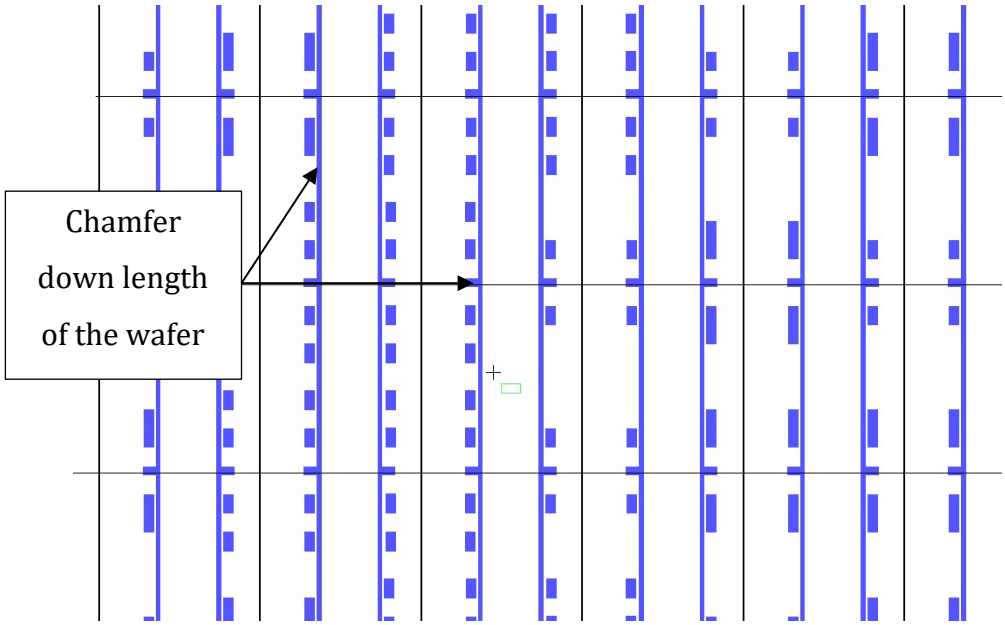

Figure 4.8 Wafer layout of intermediate capping chip

4.2.2 Epoxy Well Around Fibres

This new addition is very similar in appearance to the chamfer but performs a different purpose. Figure 4.9 shows this. The four channels for the V-grooves can be seen, the chamfer for the edge of the chip is also shown on the left-hand side of the image. The new additions are the large and wide channels around the V-grooves. These wide channels dubbed 'wells' are there to prevent any excess adhesive from leaking into areas where it ought not to be. Therefore, it functions much like a storm

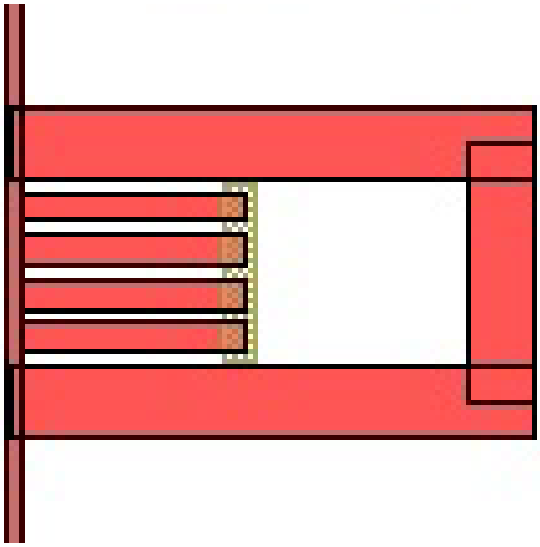

Figure 4.9 Mask close up of the V-grooves showing the well and reduced Gold mirror

drain, giving any excess liquid somewhere to flow to. This is especially important around the fibre area which is most critical. It also to a limited degree, catches debris that may be on the surface of the chip. Any dust or dirt particles that are on the chip during fabrication may move around. The wells will catch them and then prevent them getting caught in the fibre trench which would be the alternative. In this way they are sacrificial features on the chip, but since they are etched and patterned at the same time as the V-grooves, no extra work is needed to include them from a fabrication point of view.

4.2.3 Gold Area Reduction

Figure 4.9 also shows a gold hatched area which are the mirrors at the end of the V-grooves. Initially it was intended to simply deposit gold across the entire chip surface, but this was rejected for assembly reasons. The adhesive would often not bond well to the gold on the chip surface giving a very poor mating strength. So instead the critical area was outlined as shown in the figure, but then the rest was etched away using a simple metal etchant. This then enabled the rest of the chip to achieve a much higher bonding strength, both to silicon and to the silica microspheres.

4.2.4 Surface Area Recesses

These features are similar to that in section 4.2.2. But they also are designed to reduce the total surface area of the chip which ends up in contact with the Device

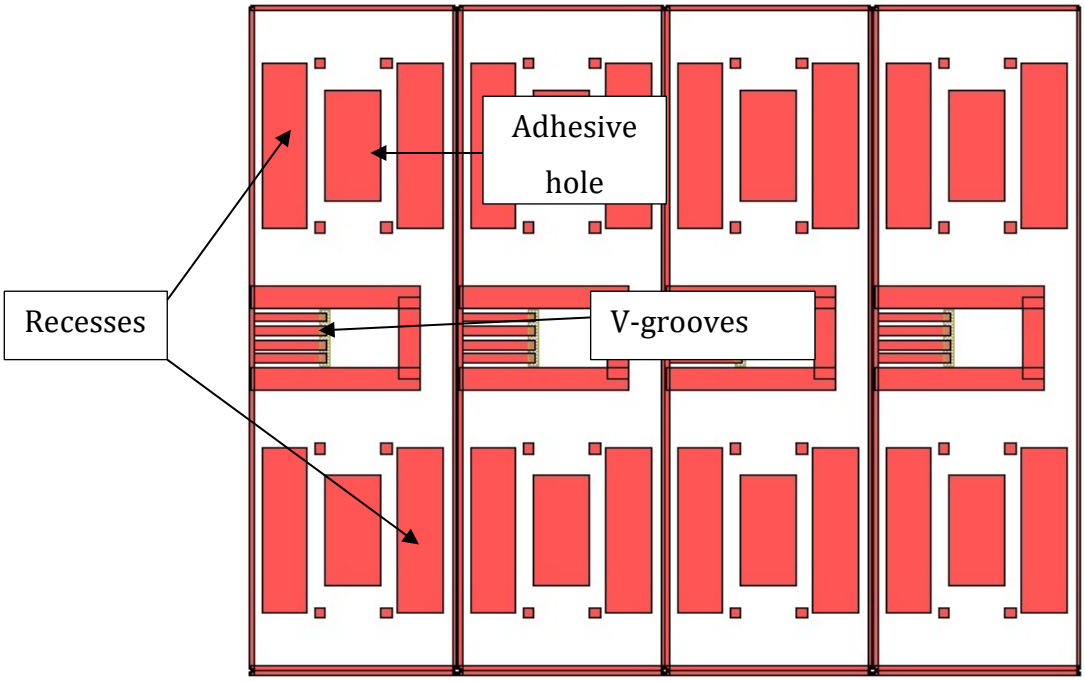

Figure 4.10 Mask showing Capping Chip features

Chip. Figure 4.10 points out where the recesses are. After having repeated problems with chip contamination, these were introduced as a means to mitigate the problem. They are simply KOH etched regions of the chip that mean that the mating surface of the capping chip is reduced in area, while the overall chip size remains the same. Any dirt or debris that is left on the chip can fall into these areas and also it provides a place for excess epoxy to flow during adhesion.

4.2.5 Chip Area Size reduction

The final change is the overall size of the chip that has decreased. The original Capping Chip was the same size as the Device Chip. However, is soon became apparent that this was in fact not at all necessary since there were less critical components of the Capping Chip. Thus, the width of the chip was significantly reduced. A trade-off occurs here, however. The larger a chip is, the harder it is to align accurately due to there being more possibility of unevenness due to contamination or other defects. However, conversely the smaller the chip is, the

harder it is to actually use during assembly without very specialist equipment. Eventually a comprise was found, where the chip was smaller than before such that there was less redundant space, but also large enough to be easily handled and aligned. The recesses could potentially have been removed altogether and the chip made shorter, but due to the timescale of the project this was rejected since it would require redesign of the device chip as well.

Chapter 5 Laboratory Equipment Used

This chapter gives clear indication as to the methods of experimentation used in this project.

For detailed discussion of the tools used for fabrication, see the Appendix. Figure 5.1 shows a photograph of the lab equipment used.

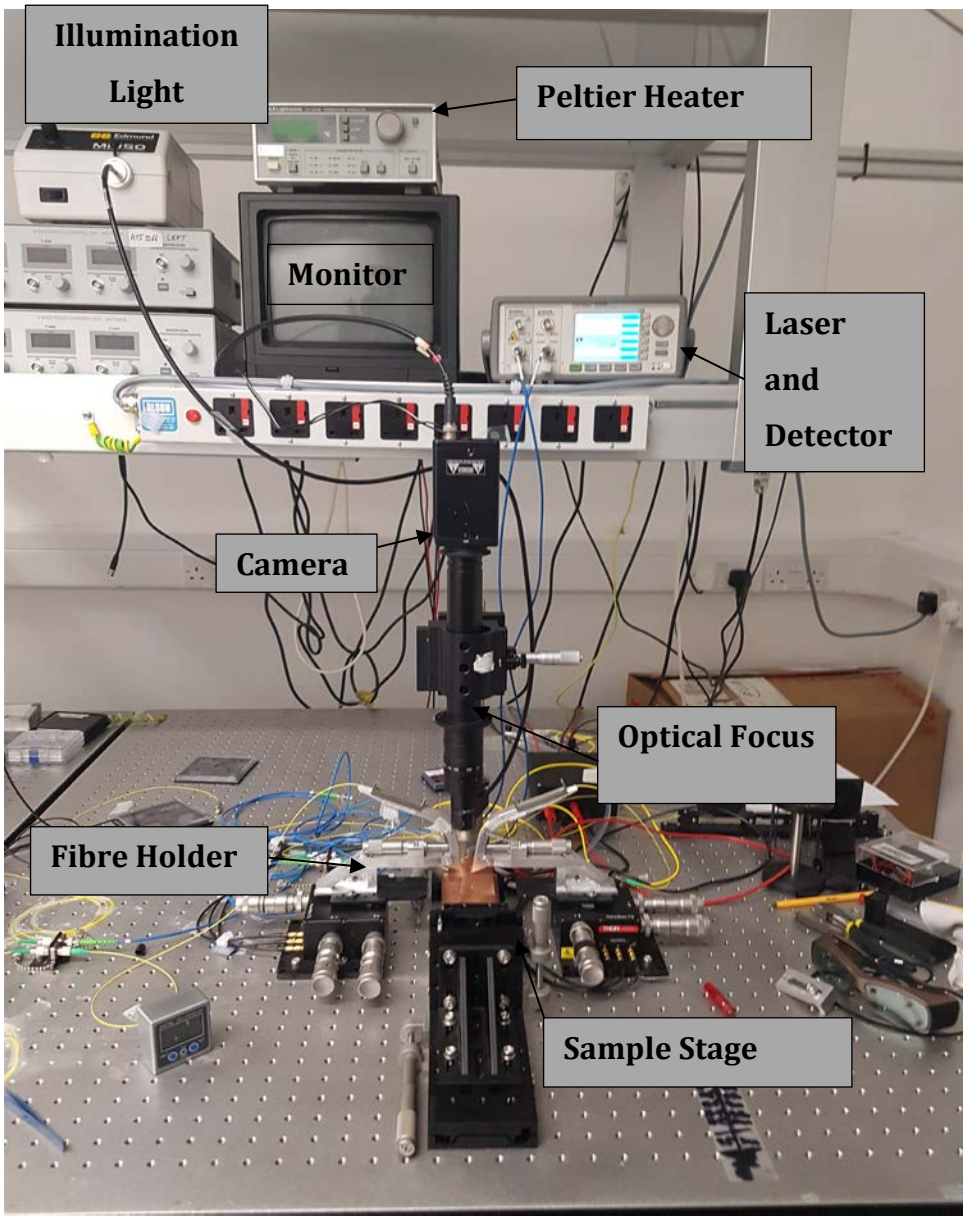

Figure 5.1 Photograph of the lab setup for characterisation

5.1.1.1 Sample Stage

In order to have a high degree of precision over the placement of the sample chips, a specific sample stage is used. It uses a solid base that can be moved in one axis by use of a micrometre such that precise lateral translations can be made with as little as 1μm adjustments. The stage used was a Newport M-426 which is moved by a Newport Vernier micrometre with 1μm resolution. The stage has M6 bolt holes giving a high degree of precision when being secured to a bench and conforming to ISO standards.

The sample itself sits upon a copper plate. This is to allow good thermal contact between the sample and the stage. If an insulating material was used, then localised hotspots could occur which will affect optical characteristics. Copper is highly thermally conductive with a thermal conductivity of 385 $Wm^{-1}K^{-1}$[167]. The vast majority of an SOI chip is bulk silicon by mass, which is also a reasonably good thermal conductor, 142 $Wm^{-1}K^{-1}$[168], and hence it can be assumed with a good deal of confidence that the temperature will be constant across the chip.

Not only must the stage be a constant temperature, but it must also be a controllable temperature as well. The lab is air conditioned and has no windows keeping ambient temperature as constant as possible, however the presence of people and other equipment will give off heat which could cause local heating of the chip stage. In fact. without a temperature controlling mechanism, the high thermal conductivity of the stage could actually make things worse since temperature fluctuations would be more severe and swift. To alleviate this problem a thermoelectric Peltier heater is used. The Peltier is sandwiched between the moveable micrometre-controlled stage and the copper sample stage. Thermal compound is placed in between the Peltier and the copper to ensure a good thermal contact. The Peltier simply uses the thermoelectric effect to provide either cooling or heating based on the magnitude and direction of current flow through it. A separate thermocouple is used to regulate temperature through a feedback mechanism.

5.1.1.2 Temperature Controller

In order to control the Peltier heater, an LDT-5910B was used which can regulate the temperature to ±0.1°C. A photograph is shown in Figure 5.2. This device takes a feedback loop from the thermocouple sensor and then outputs through the Peltier device in order to regulate the temperature precisely. Temperature control is vital

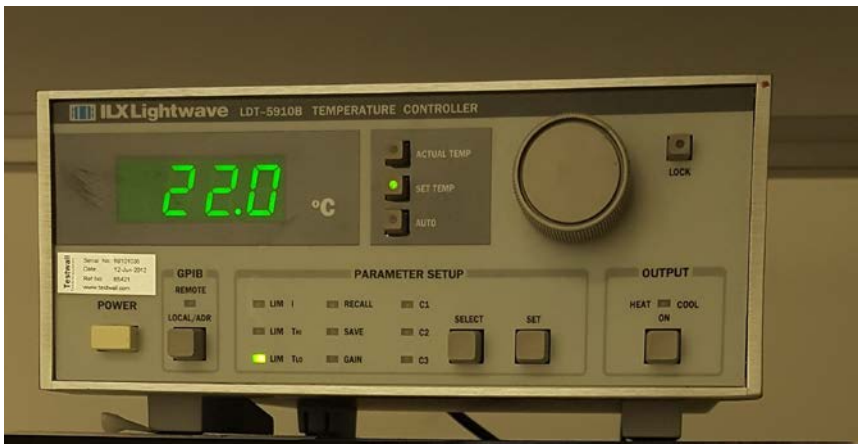

Figure 5.2 Photograph of the Newport Temperature Controller used for accurate optical measurements

due to the thermo-optic effect. That is, that refractive index is temperature dependant and so the optical characteristics of any given photonic device will vary with temperature. For example, bandwidth and wavelength dependant devices such as grating couplers will change their behaviour due to temperature as the Bragg condition will shift with a change in temperature and subsequently the insertion loss of that device will also change. The thermo-optic coefficient($\frac{dn}{dT}$) of silicon is 1.8 x10^{-4} at 300k and with a wavelength of 1550nm[169] [78], [170], which can result in peak wavelengths shifting by a matter of tens of nanometres. This is illustrated elegantly in work by Klimov et al. where a temperature sensor is created with Bragg gratings in silicon over a range of 5-160°C indicating a wavelength shift of approximately 0.082nm/°C [171].

5.1.1.3 Fibre Holders

In this project bespoke fibre holders were needed to precisely control the position of input fibres in three translational axes and one axis of rotation. Dr Scott Reynolds designed the fibre holders used for this project. The fibre holders which he designed and had manufactured are able to securely hold fibres and are attached

to a stage giving adjustments in 3 translational axes and importantly can adjust accurately the angle of incidence of the fibre, which is especially important when using grating couplers and especially in this project where grating couplers with an unusual desired angle of incidence were used.

5.1.1.4 Fibre Stage

Precision fibre stages where used to mount the fibre holders. A Thorlabs Nanomax 3-axis stage was used for both the input and output fibres. These feature micrometres much like the sample stage and are capable of providing up to 4mm of translational movements in the x,y, and z axes. The stage micrometres feature both coarse and fine adjustments making alignments easier to make and reproduce. For example, coarse adjustments are used for moving the fibre from one device to the next and fine adjustments are used to then locate the fibre in the optimum position above the device, e.g. above a grating coupler. The coarse adjuster uses 10μm markings and the fine uses 1μm[172]. In addition to this, piezo controllers can be used to give a theoretical resolution of 20nm[172] which is more than adequate for the purposes of this project. In practise acceptable results are possible without the use of the piezo controllers, the behaviour of a device can still be easily seen with manual adjustment. For the best results though the extra precision of the piezo controllers is used.

5.1.1.5 Optical Bench

The fibre stages are bolted to an optical bench which gives a number of advantages when carrying out optical characterisation. Firstly, they allow the stage to be secured firmly to eliminate the risk of the fibre being moved during measurements. Secondly they reduce vibrations by using pressurised air such that the table 'floats' on a cushion of air. This means that there is no mechanical connection between the fibre stage and the ground, only a fluid connection which will drastically minimise vibrations coming from both the user (moving, breathing, heartbeat etc.) or any other piece of electrical equipment e.g. cooling fans which will vibrate at certain frequencies. All must be considered when alignments of a fraction of a micron are being made. In this project a Thorlabs optical bench was used with a compressed air pressure of 1.5 bar for all measurements. The table's unloaded weight of 423kg and stiffness means that it is extremely resistant to external

vibrations and forces, with a deflection of <1.7μm for a load of 150kg[173] which is vastly larger than any potential external force that is reasonable in a lab setting.

5.1.1.6 Vision

In order to align fibres to such a high degree of precision within a lab environment, accurate vision of the target and the fibres is needed, though for passive alignment in a production scenario, this is not needed. Several pieces of equipment work in conjunction to provide adequate vision of the target site. Since each piece is only useful in conjunction with the other pieces and as such, they are considered jointly in one section.

Firstly, an Edmund MI-150 light box is used to illuminate the sample, with a variable light output that can be controlled by the user. The lightbox emits visible light with an internal IR filter. The light is also unpolarized. Both of these factors eliminate the chance of the light from the lightbox interfering with the optical measurements. The light from this box is then reflected onto the sample using a mirror with a camera also attached in the path of the beam of light. The camera used was a Panasonic digital camera normally used for CCTV applications. This is then paired with an optical focussing tube from Newport which enlarges the image enough such that larger devices on a PIC can be viewed. The camera and the focussing tube are both bolted to the optical bench. This way the camera, sample stage and the fibre stage are all bolted to the same reference plane. Therefore, if one were to move for whatever reason, the other two should also move with it such that

the frame of reference for the camera should remain constant. Finally, the feed from the camera is connected to an analogue display via a coaxial BNC connector.

Figure 5.3 Photograph of the lab monitor showing what the user sees when performing optical characterisation

Figure 5.3 shows what the user sees when performing optical characterisaiton. The horizontal lines are waveguides etched onto the PIC and the larger black object to the left is the outline of the input optical fibre. The other coloured lines on the screen are Moiré fringes and should be ignored.

5.1.1.7 Laser and Dector

In order to generate light for measurements, a tuneable light source is needed in order to have a practical means for measuring the behaviour of the devices at different wavelengths. A detector is also needed to calculate the loss associated with the measurements. The detected power can simply be subtracted from the insertion power and the insertion loss can then be recorded.

In this project an Agilent (Keysight) 81940A Laser source was used and an 81634B low polarization dependence optical power sensor was used as a detector. The laser has a power output up to 13dBm or 20mW, with a wide range of wavelengths available, 1520-1630nm, which adequately covers the telecoms C band which is of interest for this project as well as the adjacent L band. It also offers exceptional wavelength accuracy of ±20pm, again more than sufficient for this work[174]. The detector has an extremely low power and polarization uncertainty of ±2.5% and <±0.005dB respectively. This means that the detector gives extremely accurate results and the user need not worry about polarization changes in the fibres between the laser and detector. The InGaAs sensor gives a wavelength range of 800-1650nm ensuring that all the desired wavelengths can be detected sufficiently[175].

The detector is then output to a PC where the spectrum can be saved and analysed. The software on the PC automates the process such that a wide band is emitted and detected. By comparing inserted and received power, the insertion losses of the device and equipment can be displayed in dB. This data is then exported into a spreadsheet for further analysis and plotting.

5.1.1.8 Polarisation

Polarisation is an important factor when characterising the effect of these devices. Since grating couplers are included as the means of coupling, the polarisation of the incoming light is critical. Conventional grating couplers are essentially a 2D structure and so they must be designed for a specific polarisation. In this case the TE modes were considered as is conventional in telecoms applications. A Thorlabs Fiber Paddle Controller was used in this instance which is depicted in Figure 5.4. This device uses many fibres wrapped in a coil. Stress can be induced in the fibre by mechanical bending and twisting which will affect the birefringence of the fibre in the paddle[176][177]. The result of this is that the polarisation of the beam will be adjusted at the output[178]. By both twisting and bending the fibre, the desired polarisation can be achieved.

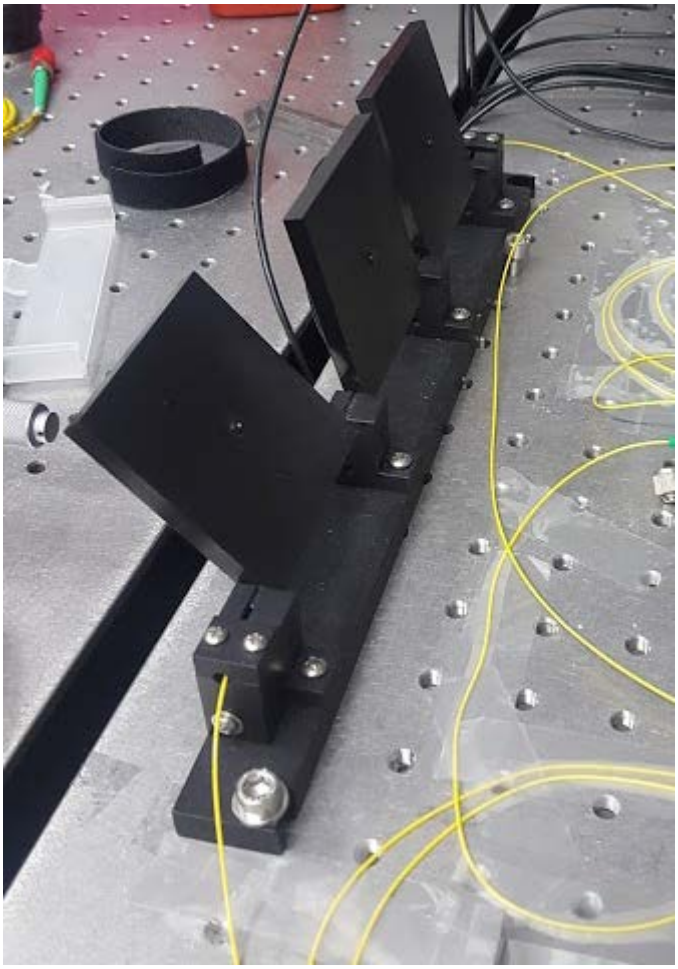

Figure 5.4 Photograph of a polarisation controller

This process does not introduce extra losses from back relfections, though there will be a minimal extra loss from the extra length of fibre needed. Without the use of a polarimeter it is still possible to calibrate the the correct polarisation by considering the power. When a maximum is found at the output with minimal distortion then the user can be confident that the correct polarisation has been achieved since the input and output grating couplers are polarisation dependant.

Chapter 6 Fabrication

6.1 Fabrication Overview

This chapter deals with all the various aspects of fabrication that were needed for the project. The chapter is split into two parts, each focussing on one of the two chips that were developed as part of the hybrid two-chip method, namely the Device Chip and the Capping Chip. A wide variety of different techniques were used and developed over the course of this project which will be discussed in detail. For technical discussion on the tools used, see the Appendix. Detail will be given on the parameters used for each process, but the physical, chemical and technological information is contained in the Appendix. For a comprehensive understanding of the fabrication for this project, read the Appendix in conjunction with this chapter.

Process development has been an integral part of this project. This chapter documents how each process has been developed and has evolved throughout the course of the project. In some cases, a particular process was rejected altogether after it was determined to be ineffective or in some cases incompatible with later processes. Each chip undergoes a series of processes and so in some instances, compromises must be made in the interest of compatibility; when the process flow is long and complex, it is not always possible to select the ideal process for each individual process step.

Also, care was taken to reduce the total number of steps in the entire process flow. Each step will increase the chance of contamination being introduced since despite the best efforts of the fabricator and using the proper equipment, foreign contaminants are always a risk. The main source of contamination is from handling wafers, which of course can never be eliminated in a research environment. The other reason to reduce the number of steps is one of cost. It is important to recall at this point that one of the main objectives of this project is to reduce the cost of a reliable packaging method. Whilst typically fabrication is not an overwhelmingly expensive part of the whole development cycle, ingenuity must still be utilised to reduce processing where possible and therefore reduce costs. With respect to

fabrication, this usually involves increased volume and the ability to process multiple wafers in parallel. For example, if a process is only possible at a chip scale, it will never be possible to implement this in a mass market situation since the low volume will drive up costs. Therefore, wherever possible, processes which allow for a high volume and with as little human time taken as can be reasonably achieved. By bearing in mind these factors it should become more apparent as to why certain processes where favoured over other ones.

6.2 Device Chip

The device chip contains all the optical devices that are used for this project. It is in essence a typical silicon PIC with a few very important modifications that are necessary to make it compatible with the capping chip and thus also compatible with the two-chip hybrid packaging method. The extra processing which is discussed here is designed such that it could be incorporated with any given silicon PIC, with any devices, passive or active. The perfect solution would be one which requires no extra processing to the device chip where the capping chip can just be attached to any silicon PIC which is built according to a standardised set of design rules. Until this is achieved a minimal amount of processing must also be applied to the device chip to make it 'package-able'.

The device chip starts as a basic SOI (silicon on insulator) wafer which was sourced in bulk from Simgui, China. The basic dimensions are shown in Figure 6.1. The top layer is here called the "silicon layer" which will contain the waveguides and other optical devices such as grating couplers, tapers, MMI's etc. It is 400nm thick in

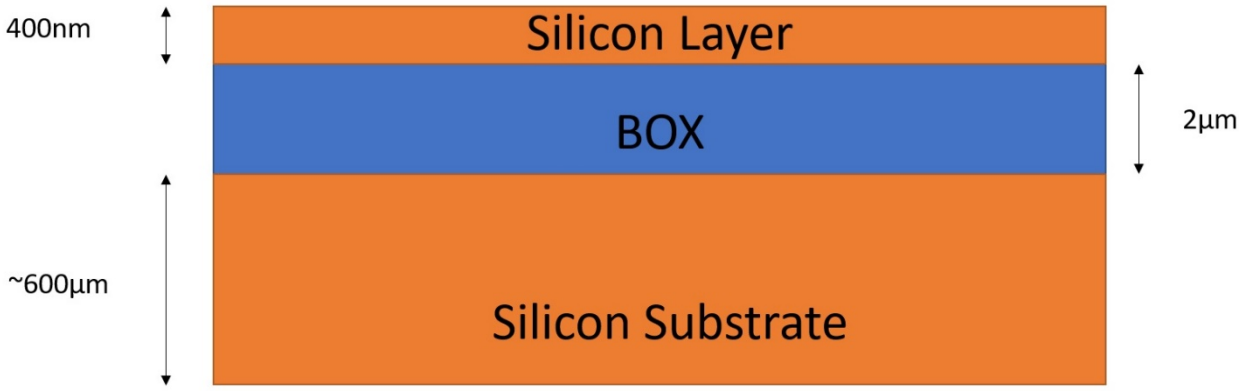

Figure 6.1 A diagram of an SOI wafer cross section

this configuration but is also available in thinner layers. This begins as a whole layer but can be etched into, to form the necessary devices. It sits on top of the "buried oxide layer" which is typically acronymised, being called the BOX. This provides the necessary underlayer of lower index cladding. In this case it is 2μm but thicknesses of 1-3μm are common. The first two layers are supported by the silicon substrate layer which is roughly 600μm though this can vary somewhat. In fact, all these values for thickness can vary by up to 1% but this should not cause any issues for this project either from a performance or fabrication point of view, however some resonant devices can require trimming due to these inaccuracies. These wafers are specialised for photonics purposes, rather than electronics, such that the silicon is grown with photonic properties in mind, rather than electronic. In short this means that optical losses will be reduced but perhaps resistances and capacitances will not be optimised. For a passive photonics project such as this one, it is no problem.

The need for an SOI wafer rather than a traditional silicon wafer is obvious when considering the silicon waveguide. The high-index silicon waveguide needs to be surrounded by a lower index cladding material (for further details see Chapter 2). Fortunately, silicon forms silicon dioxide (SiO_2) when oxidised which is a lower index material. Therefore, oxide can be used conveniently as a cladding material for silicon waveguides and is also a material which is widely used across all photonics, both fibre and planar.

The method of creating SOI wafers is quite long and complex by industry standards but is outlined briefly here for completeness. For a fuller explanation see [179], but since these wafers were purchased rather than made in-house only a brief description is required here. A standard silicon wafer is created using the usual Czochralski process[180]. This wafer is then oxidised at high pressure and temperature so that a 2μm layer is formed around the wafer. Ion implantation follows prior to precise bonding to a second silicon wafer. This wafer is then split or cut in half and undergoes chemical-mechanical polishing (CMP) down to a thin 400nm layer and to smooth the final surface which is used as the optical layer. Thus, using this method the optical silicon layer is formed with an insulating layer to prevent the optical mode in a silicon waveguide from leaking into the substrate. This is the basic 'building block' for the device chip.

6.2.1 Process flow

At this point it is necessary to show the process flow for each step that is applied to the SOI wafer. Each step was carried out at the Southampton Nanofabrication Centre (SNC) as previously mentioned. The process flow changed in a number of significant ways during the length of the project and the process flow shown here is the final form it took. Where relevant, changes to the process flow will be discussed. Whilst it would be possible to show each process flow that was made, this would be cumbersome and confusing and as such only the final version is displayed; along with the knowledge that this was a working document and so was changed where needed throughout the project. Indeed, for future work this document is liable to change yet again in future, it is very much a pragmatically governed process flow.

Identifier	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31	32	33	34	35	36
	Deposit bs oxide	Alignemnt Marks						Devices								Fibre Cavity Front													Fibre Cavity Back							
		S1813 S/B	Expose	Develop	Etch	Ash		Zep S/B	Espacer	Ebeam	Develop	Etch	Ash	Oxide Dep		AZ2070 S/B	Espacer	Ebeam	Rinse	UV Flood	Bake	Develop	Oxide etch	Deep etch	Ash	Si3n4 dep	Oxide Dep		AZ9260 s/b	Expose	Develop	Oxide etch	DRIE	Oxide strip	Si3N4 Strip	
Wafer 1	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔	✔

Figure 6.2 Process flow for the Device Chip

Figure 6.2 shows a table of the final process flow for the device chip. For reasons of clarity, a larger version can be found in the Appendix. This process flow obviously can extend downwards such that each wafer that is processed can be recorded with its progress through the flow; ticked off at each stage.

6.2.2 Backside Oxide

The first step that was required was to deposit additional oxide to the reverse of the SOI wafer. When the oxide layer for the BOX is formed on a plain silicon wafer it is done using LPCVD oxide growth in a furnace. Therefore, every surface of the wafer will receive an oxide layer. Hence the reverse of the SOI wafer also contains a 2μm layer of oxide. This layer is usually unpolished and so is inconvenient for devices, either photonic of electrical and so is usually ignored. However, in this project there is a significant amount of wafer processing that occurs on the backside. Estimating the thickness of this reverse layer is simple, even though it is very

challenging to measure it directly using ellipsometry. The thickness will be approximately the same as the BOX thickness which can be easily measured from the front side. Though the BOX specification for these wafers is 2μm, it was usually found to be slightly less than that by around 1%, so typically would range from 1970-1990nm. This oxide layer is only being used as a hard mask and as such does not need to be extremely precise, so long as it is thick enough to adequately mask the silicon substrate beneath. Therefore, an extra 2μm of oxide was deposited onto the back of the wafer using plasma-enhanced chemical vapour deposition (PECVD). PECVD was chosen here since it can be deposited relatively quickly (~40 mins per wafer) and will only deposit on the backside of the wafer. This process can be carried out at pretty much any place on the process flow before the hard mask is created but it was decided that it should be carried out first since it is a simple process and many wafers can be created in a day, then stored and processed thereafter. Such is the nature of sharing a fabrication facility that it is more efficient to block book a tool for one or even a few days rather than for an hour here or there. For this reason, this was the first process in the flow.

The recipe used to deposit the oxide layer is shown beneath in Table 6.1.

Table 6.1 Oxide Recipe

Material	Gasses	Gas Flow rate (sccm)	Temperature (°C)
SiO_2	$SiH_4/N_2/N_2O$	4.2/80/350	350

The deposition rate for this recipe is 1.16nm/s but this will vary by up to ±0.01 nm/s depending on the condition of the chamber. Therefore, this process was run with a well-conditioned chamber which had only run this process since being thoroughly cleaned. For the desired thickness of 2μm the process was run for 1700s or 29 minutes. After the plasma had ceased running, Argon was flushed through the chamber to remove any reactive species and to ensure that the deposition reaction had finished.

A higher power with higher gas flows would have ideally been used but limitations in the cleanroom prevented the standard oxide recipe from being changed. Deposition rates of almost 5x greater than 1.16nm/s are possible [181] but this recipe is kept slow to enable very thin films to be deposited with ease. This recipe gives a reliable deposition rate with very acceptable uniformity, both will usually only vary about 1% across a 6-inch (150mm) wafer.

6.2.3 Alignment Marks

Alignment marks are necessary in this process flow because there are multiple etch steps that all must be aligned to each other; most importantly the optical devices must be precisely aligned to the fibre trenches on both sides of the wafer. Therefore, alignment marks must be etched into the wafer to give a reference point that all other structures are aligned to. Now one may rightly suggest that the alignment marks can be patterned and etched during the first device etch step, and often this is the case. However, the devices for this project are only etched to a shallow depth (relatively speaking) and this can prove very difficult to detect using standard e-beam lithography when a thicker resist is used. Hence on this occasion the alignment marks were etched first, prior to any devices.

In order to pattern the wafers, a 1.3μm layer of S1813 photoresist [182] was spun onto each wafer using a resist spinner after the wafers had been dehydrated in a 200°C oven for 20 minutes. S1813 was chosen because it is a readily available resist which is inexpensive (comparatively) and can be quickly and easily spun on and exposed. S1813 is a positive resist which in most cases makes exposure and development simpler since over-exposure is less of a risk and with S1813 the risk of over-development is very low anyway due to its chemical design. With a positive resist, the exposed portion will be developed away and so the quality of exposure only needs to be good enough such that the developer will remove it (for further details on this see the Appendix). Added to this, S1813 was found to be particularly robust with respect to exposure times, allowing a generous exposure time and increasing yields. Also, development is a very clear-cut process with this resist, whereas with other resists it can be far more time consuming and ambiguous as to when development has finished. After spinning, the wafer was baked at 110°C for 60s using a hotplate to harden the resist and remove any stickiness from the surface

as well as prepare the resist chemically for exposure. The thickness of the resist is usually around 1.3μm but can vary slightly. Fortunately, this is not a problem since the tolerance to thickness is good, both for exposure and development

Alignment marks were patterned using a photolithographic mask aligner and the pattern was transferred from a chrome photomask and exposed for 2 seconds. Post exposure, the wafers were developed using MF-319 developer. Each wafer was dipped in a developer bath for 60s and then transferred immediately into a water bath and then rinsed with water for 60 seconds. Wafers were then spun-dried and

inspected under a microscope to assess the development. In all cases this lithographic recipe gave good results with all the unexposed resist being removed easily and no bleeding into the photomask. All the processing up to this point was carried out in a 'yellow-room' to remove the possibility of any exposure due to ambient high-energy photons.

Table 6.2 Silicon Etch Parameters

Etch Material	Etching tool	Gasses used	Flow rate (sccm)	Temperature (°C)
Silicon	Oxford Instruments ICP	SF_6/C_4F_8	24/45	15

The wafers were then etched using an Inductively coupled plasma (ICP) etcher to remove enough of the substrate such that clearly visible alignment marks are formed. Firstly, a silicon etch was used to remove the first 400nm of silicon and then an oxide etch was used to remove the box. The etch parameters used are shown in Table 6.2. The depth of this etch is not too critical but must only be deep enough to be detected by the e-beam. It was found that etching down to the bottom substrate layer was sufficiently deep. This meant that this etch was not time critical, but the silicon etch recipe is very selective to silicon alone and therefore for this stage the BOX acts as an etch stop. The oxide etch recipe is not especially selective to anything and will etch silicon at almost the same rate as the oxide. After 20 minutes of etching the substrate had been reached and could be verified by ellipsometry. Over time it

is very easy to visually determine a change in the colour of the etched trenches such that total BOX removal can be ascertained by eye or at low magnification.

Table 6.3 Oxide Etch Parameters

Etch Material	Etching tool	Gasses used	Flow rate (sccm)	Temperature (°C)
Oxide	Oxford Instruments ICP	$O_2/C_4F_8/CHF3$	8.5/34/37.4	15

For oxide etching a different recipe was used which is shown in Table 6.3, however the same tool was used. Finally, the photoresist was removed using the plasma asher. Ashing in an 1800W plasma for 15 minutes totally removed the resist from the surface of the wafer and no traces of the resist were left, leaving the alignment marks clearly visible both via the e-beam lithography tool and the naked eye

6.2.4 Device Fabrication

The next stage of fabrication involved fabricating the optical devices that would be present on each chip: grating couplers, tapers and waveguides for this project. E-beam lithography was used for this stage rather than photolithography since it offers far greater precision and so is able define the structures adequately for optical purposes. Of course, this would be limiting in a mass-market situation, but the same procedure could be done using a deep UV (DUV) scanner. For the purposes of research however it is more practical and economical to use e-beam lithography.

Firstly, the resist must be applied in the same way as the previous section. ZEP 520A[183] was used as a resist since it allows for high resolution, especially for right-angles, and has good resistance to dry-etching which will be used for the next step. The resist is spun onto the wafer at high rpm to give an even coverage of 400nm. This is then baked at 180°C for 120s to be adequately soft-baked. An additional layer of E-spacer is pipetted onto the wafer surface and then spun-on. This layer is extremely thin and is conductive. Without this layer the performance

of the e-beam is significantly reduced since the charged electron beam will also charge the insulating BOX layer which will then produce adverse effects on the writing process.

The wafers were then written using the e-beam with a dose of 190 $\mu C/cm^2$ which provided good exposure and was advised by Dr Ali Khokar. The e-beam 'wrote' the structures onto the wafer but in order to save time only a portion of the

Figure 6.3 Diagram showing the e-beam etch outline

whole wafer was etched surrounding the waveguides. Figure 6.3 demonstrates how this is done. The surrounding silicon walls are sufficiently far enough (approximately 20 microns) such that no light is coupled into them, but this will save many hours of e-beam write time.

After the write is complete (several hours) the wafers were given a post-exposure bake of 100°C for 120s which finished the cross-linking process within the resist. The E-spacer was then be removed by rinsing, which is simple since the conductive chemical is easily soluble in a polar solvent like water.

Development was then carried out using ZED-N50 developer. The same process as before is used to develop the wafers, being submerged for 1 minute and then thoroughly rinsed and dried. Then the wafers were ready for etching.

The etching process is identical to that descried previously. The same recipe and conditions are used. However, in the case of the devices, the etch depth is very critical. If the trenches around the waveguide are etched either too deep or too shallow this will have an adverse effect on the performance. The waveguide mode will not be adequately guided within the waveguide unless the correct etch depth is used. Furthermore, the etch depth is even more critical for the grating couplers.

Small variations in the etch depth can drastically alter the effective index of the device and thus the coupling characteristics.

To ensure accurate etching, the etch chamber was 'conditioned' for 20 minutes before the wafers were etched. This involved etching a dummy wafer with some resist on it with the silicon etch recipe for all that time. This process helps to bring the etch chemistry of the chamber to a stable point. Etch rates will tend to vary over time until a steady state is found, which in this case was found to be at 20 minutes. In this case, the stability of the etch rate rather than the absolute speed is what is crucial, to ensure that the wafer is correctly etched. After this conditioning procedure a stable etch rate of 5nm/s was attained, meaning that the etch depth of 220nm was reached in a total etch time of 44s. Each wafer was etched in total for this amount of time. After etching, ellipsometry was used to measure the final thickness of the etched silicon layer. Any samples that were found to be ±10nm more than the desired etch depth were rejected. Thankfully after fine tuning the conditioning and etch recipes, only 1 wafer was ever rejected in this way. Bearing in mind that 10nm here represents only 2s of etching and less than 5% of the total etch depth, this was a successful procedure.

The remaining e-beam resist can then be removed using plasma ashing for 15 mins, as previously described. Then the wafer is 'passivated' by depositing 1μm of oxide onto its surface. In electronics the term 'passivation' means to provide an insulating layer from any metallic contacts on the surface of the chip, and this prevents short circuits. In photonics the process is identical but the rational a little different. Of course, air can be used as a cladding material around the silicon waveguide and is also a lower index material than oxide. However, the oxide layer protects the silicon waveguides and adds a shield which will stop any dust or particles from being present on the surface of the waveguide. 1μm of oxide will totally encapsulate the propagating mode at this wavelength and waveguide dimensions. Again, the same PECVD oxide recipe was used for this process. LPCVD oxide would not be appropriate since this would alter the silicon layer whereas PECVD oxide will simply deposit the extra oxide required, as explained previously.

At this stage a traditional SOI wafer would be considered complete. However, the packaging method used here requires a few extra steps. Before moving on, it is useful to see a cross sectional diagram of the wafer as it would be at this stage.

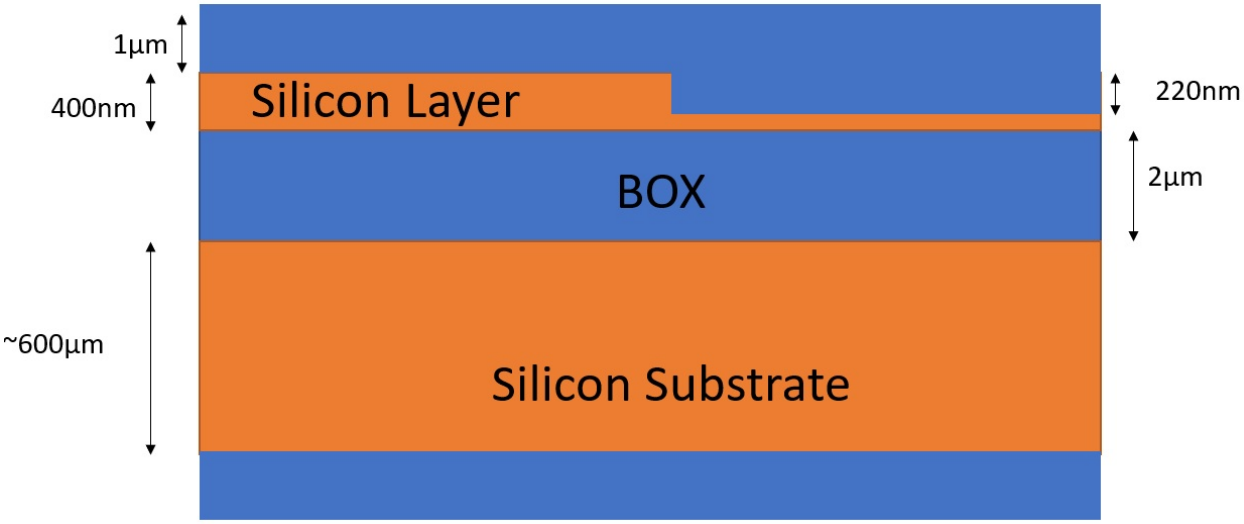

Figure 6.4 Diagram of the SOI wafer after devices have been etched

Figure 6.4 shows as a cross-sectional schematic diagram of how the wafer looks at that stage.

6.2.5 Fibre Cavity Front

The next step was to create the cavity or trench for the optical fibres to sit in. This has become a rather complex procedure since the decision was made to drop the fibres to a lower position with respect to the waveguide, which would decrease the propagation distance, but then also to etch through the entire wafer for improved vision and adhesion. The processing from the front and backside will be considered individually here since although they are a combined effort to create one structure, the processing steps are essentially distinct and as such it is more appropriate to deal with them in turn.

The first stage as before is to lithographically pattern the wafer. This step must be done with careful consideration of the previously defined structures so must be accurately aligned to the alignment marks that were made. For the benefit of the reader, the relevant section from the process flow has been reprinted here to make reading easier.

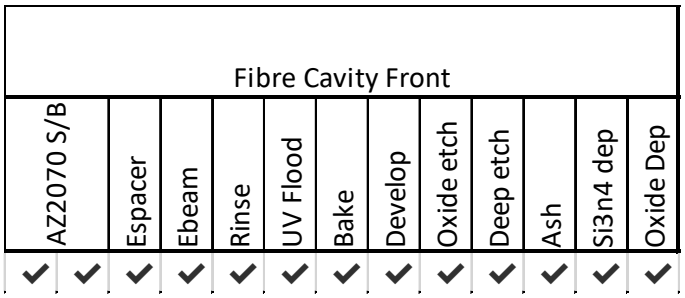

Figure 6.5 Image of the process flow for the Front Fibre Cavity

This lithographic process is called a 'hybrid-process' and was pioneered by Dr Ali Khokar and Dr Scott Reynolds. The hybrid process uses both e-beam and photolithography in order to adequately pattern the wafer.

AZ2070 is used as the resist for this etch step since a thick resist is needed for the etching process. It is spun onto the wafer to a thickness of 7µm (2500rpm [184]) and baked at 110°C for up to 7 minutes. The data sheet recommends 1 minute per micron of resist. However, in practice the resist was usually sufficiently baked after 240s. The variation is presumed to be due to the amount of solvent that had evaporated out of the resist over time, such that an older bottle of resist would bake quicker. E-spacer was then applied to the wafer as described previously and then the wafer is ready to be written in the e-beam.

AZ2070 is a negative resist and as indicated previously, this makes the process slightly more challenging. However, it is also compatible with the e-beam, so this is the ideal resist for this process. The e-beam is required for its accuracy which is crucial for alignment to the devices. However, writing a whole wafer (negative resist) with e-beam is simply not feasible and the write times would be exceedingly long, and thus the cost would be astronomically high. Therefore, the hybrid process allows the critical outline of the structures to be written accurately and then the remaining area is exposed using the UV flood on an EVG mask aligner. Despite needing two steps, the overall time is reduced and costs too. This process would of course be compatible with a DUV scanner for industrial applications

High quality alignment was achieved between each lithographic step. As explained before this technique was already devised by Dr Khokar and Dr Reynolds but was used for the purposes of this project.

This resist also requires a post-exposure bake to finish the cross-linking of the resist, in addition to the UV. This took 60s at 110°C[184]. Finally, the wafers were developed using the same process as mentioned before (developer -> water -> spin-dry -> microscope examination) but the appropriate developer was used. In this case it is AZ-726 MIF. Development time was very difficult to predict. The developer loses its potency over time and due to the thickness of the resist it can be difficult to ensure that the development has finished, and the trench become clear. Therefore, after an initial development time of 3 mins and checking, the wafers were developed for subsequent 30s periods until the trench was clear. AZ400k is also recommended by the data sheet at a ratio of 4:1 with water. However, after many trials this developer always performed very poorly with these conditions, in contrast the AZ-726 worked well. Therefore, AZ-726 was the chosen developer for this project. For different structures it could be that AZ400K performs better but unfortunately it consistently underperformed in this project.

Next, the oxide layer must be etched away to reveal the bare silicon beneath. As before, the oxide recipe would very happily etch through the 400nm silicon layer as well, meaning that etching need only be done in one sitting. The same recipe and conditions were used as before for oxide etching. The oxide etch recipe will etch the resist at an almost comparable 1:1 ratio as the oxide, hence the thickness of the resist required. Once the silicon substrate beneath is reached this etch process can stop. As before, ellipsometry is used to confirm this, but after practise, visual inspection can be used to determine that the silicon layer has been reached.

The final stage for the front side of the fibre cavity is to deep etch the front side. After several trials it was decided that 100μm would be etched from the front side and then capped. In order to etch this, the Plasmatherm Versaline etcher was selected, using the Bosch deep silicon etch as outlined in the previous chapter. The rules that are applied in the SNC dictate that only soft masks of <2μm may be used to prevent resist from reflowing under temperature and pressure and causing the

wafer to stick to the clamp. This had been a significant problem in the past. The resist thickness was estimated at this stage by using profilometry. In most cases the thickness was greater than 2μm and hence the resist was removed using plasma ashing as before. If the resist thickens was roughly 2μm then it was kept. Thankfully oxide works very well as a hard mask and so can be used for the deep etch process, given that it is a relatively shallow deep etch. The following recipe shown in Table 6.4 was used.

This etch recipe gave an etch rate of almost exactly 10μm/min. From the front side, so 10 minutes of etching was done to reach a depth of ~100μm. The exact depth here isn't too critical, so long as it is over 100μm and ideally less that 160μm since this the maximum depth that can be measured using the profilometer. At this stage each wafer was measured using the profilometer and the depth was recorded, such that the correct amount of etching could be repeated onto the backside of the wafer. The etch steps refer to standard steps of the Bosch Etch which are explained subsequently.

Table 6.4 DRIE Bosch Etch recipe

Etch Step	Gases	Flow Rates (sccm)	ICP Power (W)	LF Bias (V)	Pressure (mTorr)	Time (s)
Dep	SF_6/Ar	150/30	2000	10	25	2
A	SF_6/Ar	350/30	2500	~300	80	1.5
B	C_4F_8/ Ar	150/30	2000	10	40	2

Finally, the wafer was capped with nitride to 400nm and then an etch stop of 2μm of oxide was also applied. Both were done using PECVD, using an Oxford Instruments tool.

For nitride deposition, the following recipe was used as shown in Table 6.5.

Table 6.5 Nitride PECVD recipe

Material	Gasses	Flow (sccm)	Temp (°C)
Nitride	$SiH_4/N_2/NH_3$	12.5/500/20	350

The thicknesses of either of these steps is not too critical but excessive thickness will introduce unnecessary stress. 400 and 2000nm respectively was found to be adequate after some tests, such that the nitride would sufficiently cap the devices from the oxide HF etch and the oxide would provide an adequate etch stop for the deep etching from the backside. At this stage the front-side fibre cavity can be considered complete.

6.2.6 Back-side Cavity

This stage will complete the work that is carried out on the frontside to create the all-important fibre cavity that makes this packaging method viable for assembly. Despite only accounting for a small portion of the over-all processing that is applied to the Device chip, the process for the backside cavity took a great deal of time to properly develop and much thought was put into this step. Once again, the relevant

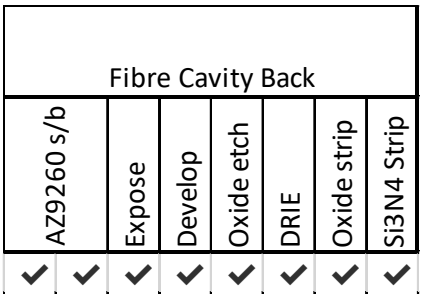

Figure 6.6 Process flow for the Backside Cavity

section from the total process flow has been reprinted here to allow for easier reading, seen in Figure 6.6.

It is important to stress at this stage how important this step was for the overall project. In order for the packaging method to be properly applied to a silicon wafer it was necessary to etch through the entirety of the wafer as mentioned before.

Whilst etching through a silicon wafer is not a new process, a unique recipe must be designed for this packaging method. The following section outlines the engineering problem in detail, the test carried out by the author and then final recipe that was ascertained in order to achieve the desired etch. This specific section of work required several months and is a pivotal part of the project.

Photolithography is sufficient for this process since the shapes have already been defined on the front side. Backside alignment requires the use of an EVG mask aligner which has cameras on the top and bottom, such that when the wafer is flipped over, the alignment marks on the frontside can be aligned to it. The precision alignment step is already carried out on the frontside with the e-beam, which can then be aligned to.

AZ9260 is used as the photoresist for this step since it allows for a thicker resist that is sufficient for the later etching process. AZ9260 can achieve similar thicknesses to AZ2070 but is a positive resist and doesn't require a post exposure bake. A 7µm thick soft mask is required which is spun onto the wafer using a spinner and then baked for ~120s at 110°C, again as before, sometimes an extra 30s of baking is required for the resist to be adequately baked. Exposure was carried out on the appropriate EVG mask aligner with backside capability. Due to the thickness of the resist, an 11 second exposure was required to expose the resist fully. A number of test wafers were needed to ascertain this length of exposure with shorter exposure times producing insufficient clearing when developed. At longer exposures there is a risk of over-exposure so 11s was chosen as the best exposure time.

Due to the poor performance of AZ400k with AZ2070, it was never used for developing AZ9260 at this stage, though the data sheet recommends it. Again, it has not been decisively deduced exactly why this developer has performed so poorly in this project but instead AZ-726 has always performed to a high standard and so was used again for developing the AZ9260. Development of such a thick resist could be time consuming. Initially the wafers were submerged in an AZ-726 bath for 3 minutes, then rinsed, spun-dried and examined under a microscope. From this point the development procedure was repeated in 30s bursts until the trenches on the

backside were clear. The unpolished backside of the wafer could sometimes make visual inspection troublesome, so great care needed to be taken at this stage. Fortunately, AZ9260 is typically clearer under a microscope than AZ2070 is, which again makes it a better choice for this stage.

After development, the wafer is then ready for etching. It is worth mentioning now that alignment between the front and backside was found to be exceptionally accurate (<2μm), as can be demonstrated by cross sectional SEM images that will be shown later.

Deep reactive ion etching (DRIE) is an extension to RIE (see Appendix for more details) that allows it to be used to etch greater depths. Fundamentally, the process is very similar to RIE. It is a dry etch which uses a plasma to etch a substrate. There are however some important differences. Typically RIE rates are of the order of 100-1000nm/min, whereas DRIE rates are much higher, circa 2-20μm/min [185]. This allows a great deal more capability than standard RIE, with etch depths of hundreds of microns being possible, which have been used to great success particularly in the study and research of MEMS [186]–[188]. Sidewalls of etched trenches can be almost completely vertical which allows for etched holes to be extremely deep but also extremely narrow, with ratios of up to 100:1 possible, whereas conventional RIE struggles to maintain comparable aspect ratios. The selectivity to the mask becomes problematic when etching for longer periods during deeper etches. Both hard and soft masks are very difficult to make thicker than a few microns and so this then imposes a boundary on etching since the entire mask will be etched away over longer etches unless an extremely high selectivity to Si can be achieved[189]. Etching is also independent of wafer orientation with respect to its crystal structure which is not the case in wet chemical etches. Any shape that can be patterned onto a wafer can be etched into the substrate almost as much as is desired. DRIE is usually used in conjunction with Si substrates as it was in this project so for this discussion DRIE will centre on Si etching.

There are two main methods of DRIE: Bosch and cryogenic. For this particular project the Bosch etch (also called Bosch advanced silicon etch, ASE) was used. The Bosch etch is named after the Robert Bosch GmbH company which

patented it[190]. This process uses a combination of SF_6 and C_4F_8 to etch an Si substrate[191]. However, unlike RIE etching there is a series of individual steps that are used in order to produce the desired etch. High aspect ratios of 20-30:1 are possible as well as fast etch rates; upwards of 3μm/min are standard for the Bosch process[190]. The Bosch process consists of a cycle of different etch and passivation stages. During passivation a thin layer (~10nm) of fluorocarbon polymer is deposited across the wafer[192]. Usually C_4F_8 is used for this purpose. Then during the etch step this is removed from the bottom of trench which is etched by the addition of SF_6. The cycle then is repeated as many times as necessary in order to obtain the etch depth required. Each step of the cycle will only run for a few seconds, so hundreds of cycles are sometimes required for deeper etches.

Figure 6.7 shows a diagram of the Bosch process. The steps of the process are as follows; Step 1: this is the wafer with a mask and an exposed area which is to be etched. Step 2: Ions accelerated by bias voltage, bombard the exposed area and

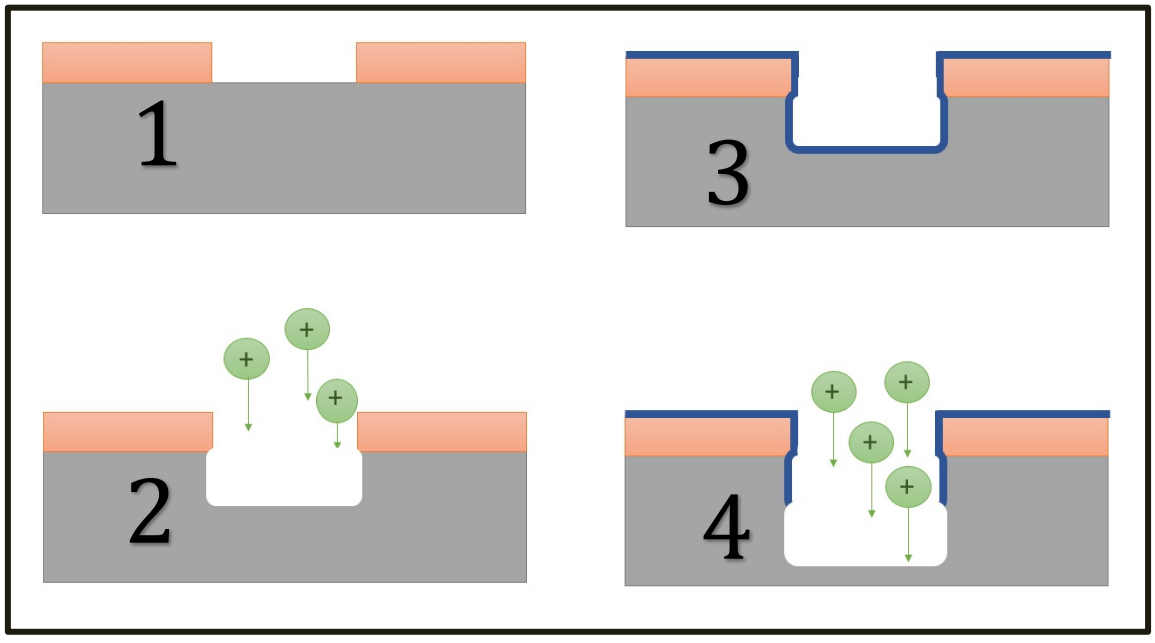

Figure 6.7 Diagram showing the steps of the Bosch Process

create a trench that is a few microns deep. Step 3: The polymer is deposited onto the surface of the wafer to passivate it. Step 4: The directional ion bombardment etch is used to etch further down into the substrate. As can be seen the Bosch etch will usually result in a slight undercut due to slight isotropy when etching. The diagram is of course exaggerated hugely in order to highlight the problem. The effect is small

yet noticeable. Figure 6.8 shows a close-up scanning electron microscope image of a trench that has been etched using the Bosch Process. The wavy sidewall profile is the result of this undercut. This is often called scalloping because the shape resembles the silhouette of a scallop shell. As can be seen from the image with less zoom, the affect is very small when considered on a larger scale.

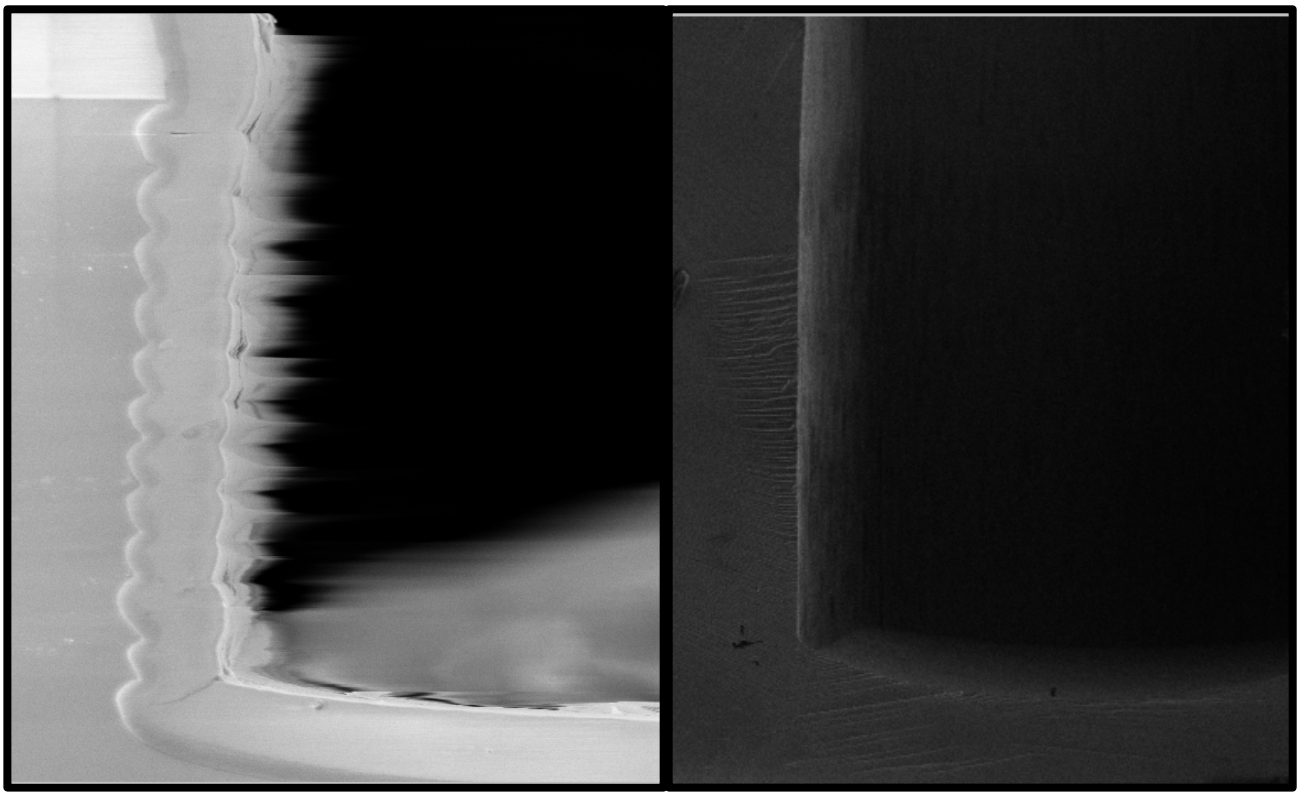

Figure 6.8 Scanning Electron Microscope image of a trench etched with the Bosch Process. A zoomed-out image is given next to it for sake of reference.

One of the downsides to using the Bosch process is that the etch rate is considerably dependent upon the feature size that is being etched. This effect is called DRIE lag. This is typically caused by charging effects and capacitance being different around different sized features. The distribution of charges will cause the ions to etch in a non-uniform way. Other than creating a mask with equal feature sizes there is little that can be done to compensate for this effect. Process parameters can be varied slightly to alleviate some effects of DRIE lag, but this usually provides only marginal gains.

The Bosch Process does however have a very good selectivity to oxide and so this can be used as an etch stop buffer. This is not a total solution however as ion bombardment will cause the etch species to bounce off the oxide layer and will etch the sidewalls creating a 'notch' at the base of the trench. Figure 6.9 shows how notching can affect the etch profile. The trench will flare out at the bottom due to the etch species bouncing off of an oxide etch stop.

DRIE using the Bosch Process is an extremely versatile and important tool within nanofabrication, allowing extremely deep trenches with very high aspect rations, resulting in almost totally vertical sidewalls. Yet despite its usefulness, the Bosch Process must be tailored to each individual mask shape in order to minimize

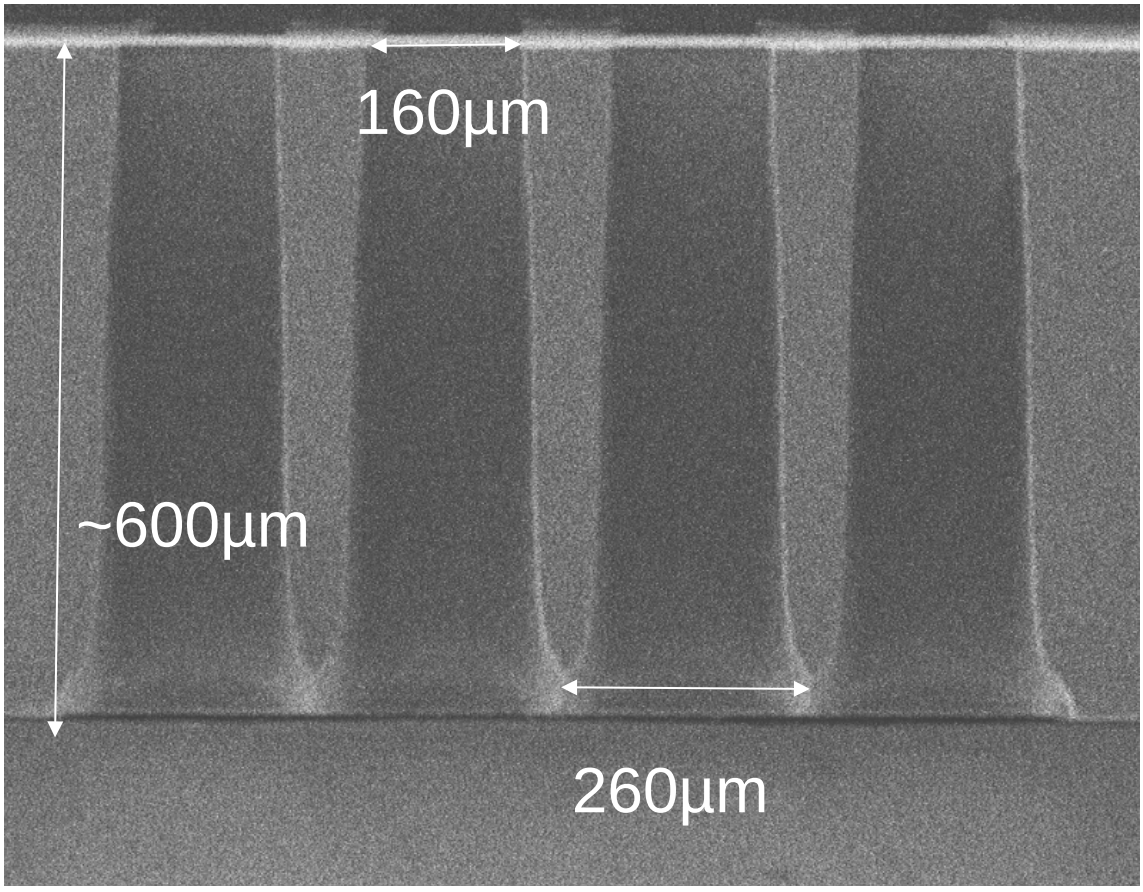

Figure 6.9 SEM image of notching on a DRIE trench

adverse effects that arise due to DRIE lag, scalloping and notching. Process development for DRIE will be explained later.

For this project a Plasma-therm Versaline was used for DRIE as depicted in Figure 6.10. This model offers very advanced etching methods to produce high

aspect ratios, smooth walls, high uniformity and high selectivity[193]. This model also features endpoint detection where it can detect that a material has been etched up to an etch stop layer, though it was never successfully calibrated for this project.

Figure 6.10 Image of a Plasma-Therm Versaline DRIE

As mentioned before, a maximum of 2μm of resist was permitted to be used with the Plasmatherm deep etcher and so this made etching the full depth of the wafer challenging. AZ9260 is very resistant to deep etching, with a selectivity of around 500:1, making it ideal for such a process. Unfortunately, early attempts to do so frequently resulted in the resist reflowing due to the temperature and pressures that it was subjected to for extended periods of time. This would lead to the wafer getting stuck and the machine needing repair. Whilst it was somewhat frustrating to have to redesign the process to work around this, it was most likely beneficial in the end.

Due to the constraints mentioned above, an alternative was needed. Therefore, the process was redesigned to use an oxide hard mask instead of the soft mask. This required extra oxide to be deposited to the backside of the wafer, which in this process flow happened first. The extra deposited oxide provided and

adequate layer of hard masking. Due to the change in masking material, the parameters of the etch needed to be adjusted slightly to compensate. To illustrate this point, a quick discussion must be made of difficulties of deep etching using the Bosch process.

When etching to depths of >300μm, there becomes the danger of a phenomena known as “grass” or “grassing”[194]–[196]. This is caused by lack of power that the bombardment ions possess as they have to travel the whole length of the trench, losing energy as they travel. As the power decreases, their ability to etch effectively is diminished which will lead to areas which aren’t etched properly. As the trench etches its way down, these areas remain as pillars or blades of “grass” which then become extremely difficult to remove[195]. The best way to mitigate against this is to increase the bias voltage and hence accelerate the ions to a greater velocity during the bombardment stage. Figure 6.11 shows an image of an early attempt that was made to etch through the wafer where insufficient bias was used. Over bombardment though can become an issue as well, with sidewall profiles suffering and scalloping becoming a bigger issue. The other negative effect of high LF bias voltage is that it will cause greater bombardment to the masking material, in this case oxide. Whilst the oxide is very resistant at lower bias voltages, it does begin to be etched away to a non-negligible degree once the bias voltage is increased

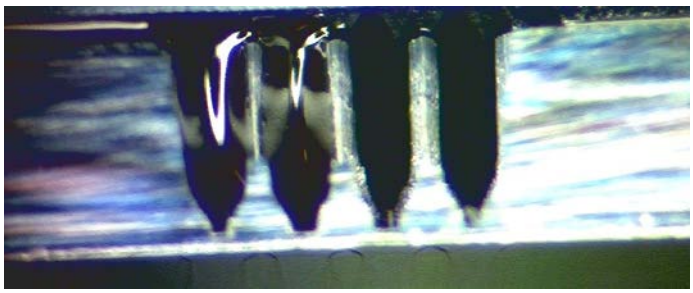

Figure 6.11 Picture showing and early attempt to etch through the wafer

The other consideration to make is that the process above is not constant across the whole wafer. There will be different charging effects, and hence etch rates on different sized features. Generally smaller size features will etch slower since they will usually receive a lower concentration of ions during the bombardment stage. Larger features will etch quicker, up to a certain point at which other factors

begin to dominate. For this reason, when deep Bosch etching is used, it is typically used with uniform circular shapes. This attempts to ensure that the etch rate will be constant across the wafer and the round shape prevents any undue charging effects that can cause the edges and corners to etch more slowly, since charge accumulates at a point.

In this project however, the geometry of the features to be etched are non-uniform in both shape and size. This presented a unique engineering challenge that is bespoke to this particular mask.

Consequently, many tests needed to be conducted to obtain the correct etch parameters. A series of tests were undertaken to tackle the problem at hand. A dynamic method of etching needed to be found that would adjust the etching parameters in step with the changing etch environment and topology during etching. In order to do this the recipe is required to "morph" dynamically in time. Different parameters can be morphed over time. Figure 6.12 shows the example of bias voltage being morphed in different ways over time. The temperature of the

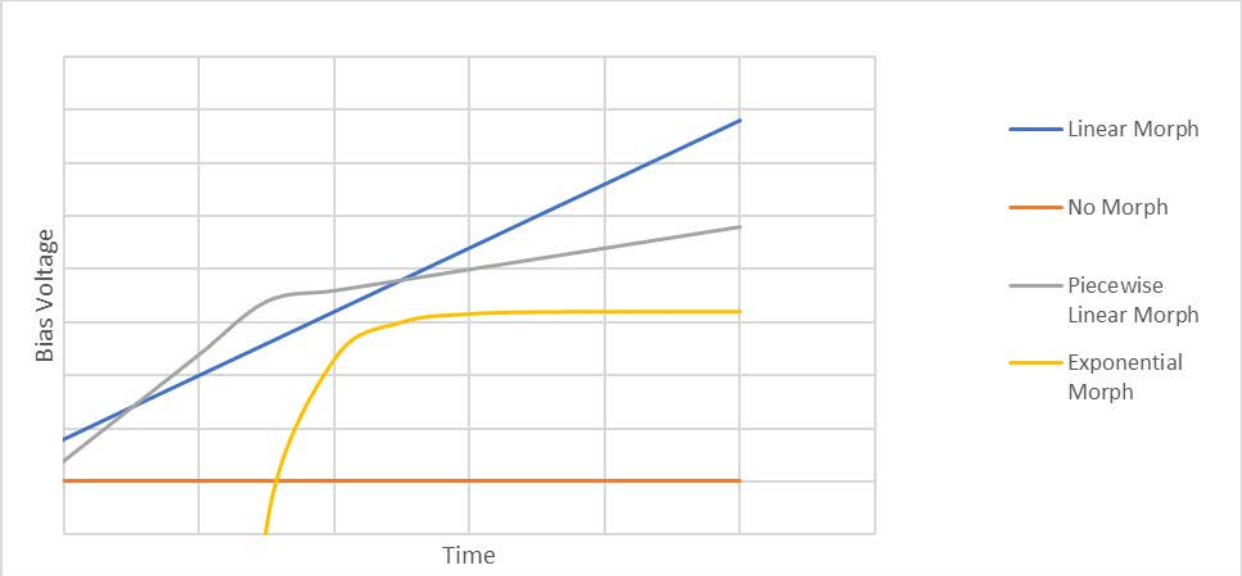

Figure 6.12 Graph showing different types of etch morph

chamber was found to have little impact in this regard, and it was determined early on that the LF bias was the dominating feature in this process. Alternatively, the etch chemistry could have been altered with differing flow rates of each gas. This was avoided due to the complexity of understanding the etch chemistry accurately

enough to be able to manipulate it effectively. Fortunately, it was found that varying the chamber bias voltage was sufficient.

With regards to Figure 6.12 a basic “No Morphing” would correspond to a constant voltage throughout the etch, shown as the red line. This is acceptable for very shallow etches only. A linear morph is where the morphing changes linearly with a constant of proportionality. This is shown in Figure 6.12 as the blue line. This means that the increase voltage each unit it of time will increase by a constant amount. If this is insufficient then a piecewise linear morph can be used, shown above with the grey line. This is simply the addition of several linear morphs at specific intervals. Finally, for very complex etches an exponential morph can be used where bias will vary according to an exponential rate.

A series of tests were conducted to empirically find a morphing parameter that could adequately etch the trenches for this project. At the end of each test, the wafers were measured for depth to ensure they could adequately etch the desired depth and then examined under an SEM to ensure the mitigation of scalloping, notching and grass.

The bias can be adjusted as the process is performed and this is written into the recipe. A beginning and end voltage can be set if a linear ramp is required or more complex functions can be written into the recipe. A base voltage of 350V was used since this was the default for any effective etching to take place.

To begin with, it was unknown what kind of morphing would be required for this topology. So initially a series of linear morphs were trialled.

Table 6.6 Bosch Etching Tests

Voltage step rate (V/min)	Max etch depth (μm)	Mask etch depth (μm)
0	~300	<1
0.2	~330	~1
0.4	~380	~1.5
0.6	~430	~2
0.8	~500	~3
1	~550	~3.5

From Table 6.6 it can be seen that at about 0.6V/min the thickness of the hard mask becomes problematic. Max etch depth here indicates the maximum depth that could be reached before grassing and other undesired effects became an issue. For example, square features became rounded off, such that they became round holes on the other side. The other issue is that of 'landing' the etch. If a high LF bias is used

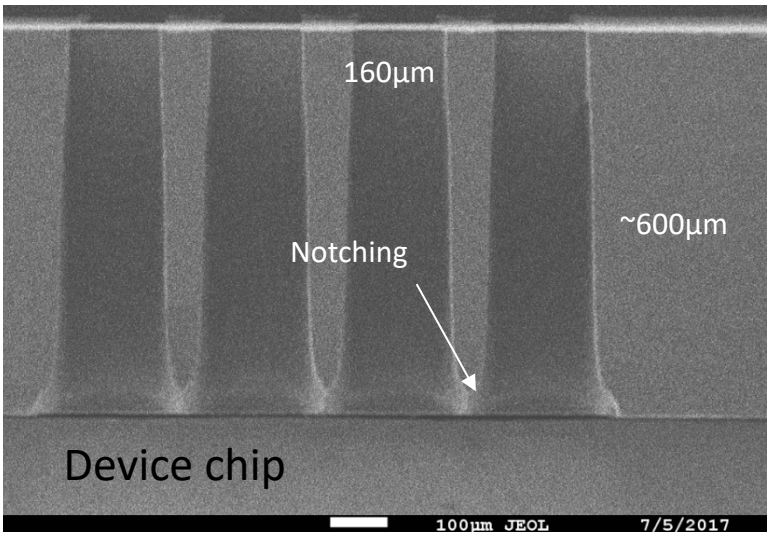

Figure 6.13 Notching at the base of the through wafer etch.

for too long, then the ion bombardment off of the oxide membrane becomes a problem and the trenches become notched. Figure 6.13 shows an SEM image of when this has been a problem.

The depth cannot be ascertained to a high degree of accuracy and would vary somewhat across a wafer and across different features as mentioned above, since once the limit is reached (i.e. maximum etch depth) the etching process begins to break down and uniformity likewise. Depths were measured using a microscope and focussing to the bottom of the trench which gives an accuracy usually to around ±10μm. Taking this data, plus plenty of other qualitative data, it was found after more tests that 0.9V/min gave good results, given that just over 100μm had been etched from the front side.

At this stage, both sides of the cavity have been etched and the only remaining material in the cavity is the oxide membrane and nitride cap used to protect the devices. Visible inspection is possible, through these thin membranes such that daylight can pass through it, showing that the etch has been successful. A microscope was then used to assess how successful the etch had been.

Some sacrificial samples were used with the SEM and the results are displayed below. Figure 6.14 shows how much the notching had been improved by the final recipe. The top of the image shows the 100µm etch that was performed from the frontside and a much smaller amount of notching can be seen. Given that the total etch depth is 500µm, this amount of represents around 5% of the total etch depth and is small enough in relation to the fibre to not be problem. In addition, Figure 6.15 shows an example of the whole depth of the etch. The bottom 212µm show a strange 'bowing' affect but this was due to tool malfunction on this particular run (hence the sacrificial use with the SEM). The remaining 282µm show very good uniformity and reaches the oxide membrane with a soft landing. This recipe was then used after it was shown to produce the desired etch.

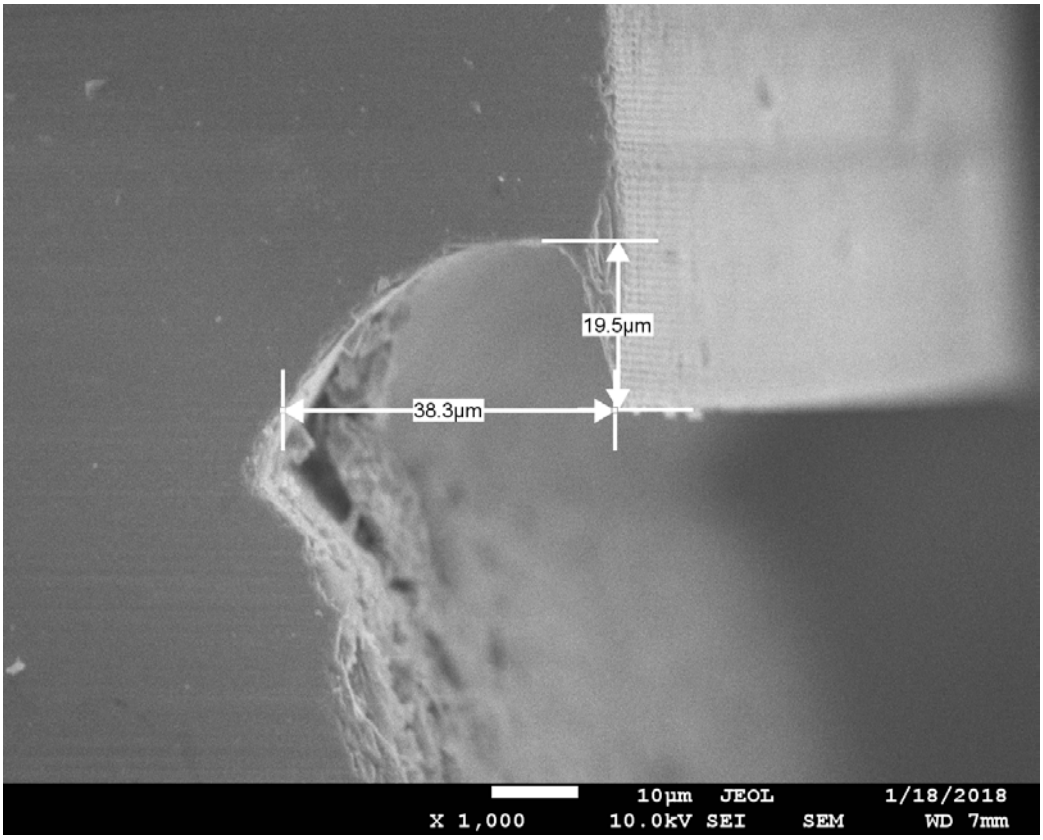

Figure 6.14 SEM image showing much smaller notching on the final recipe

Finally, hot phosphoric acid at 140°C was used to remove the nitride layer and HF 7:1 was used to remove any excess oxide, such that the final oxide thickness was as close to 1µm as possible. The wafer at this stage is complete and can be diced in preparation for assembly. Dicing will be considered part of the assembly procedure in this project. An image of the final device chip can be seen in Figure 6.16, displaying the different structures that have been etched into it for packaging purposes.

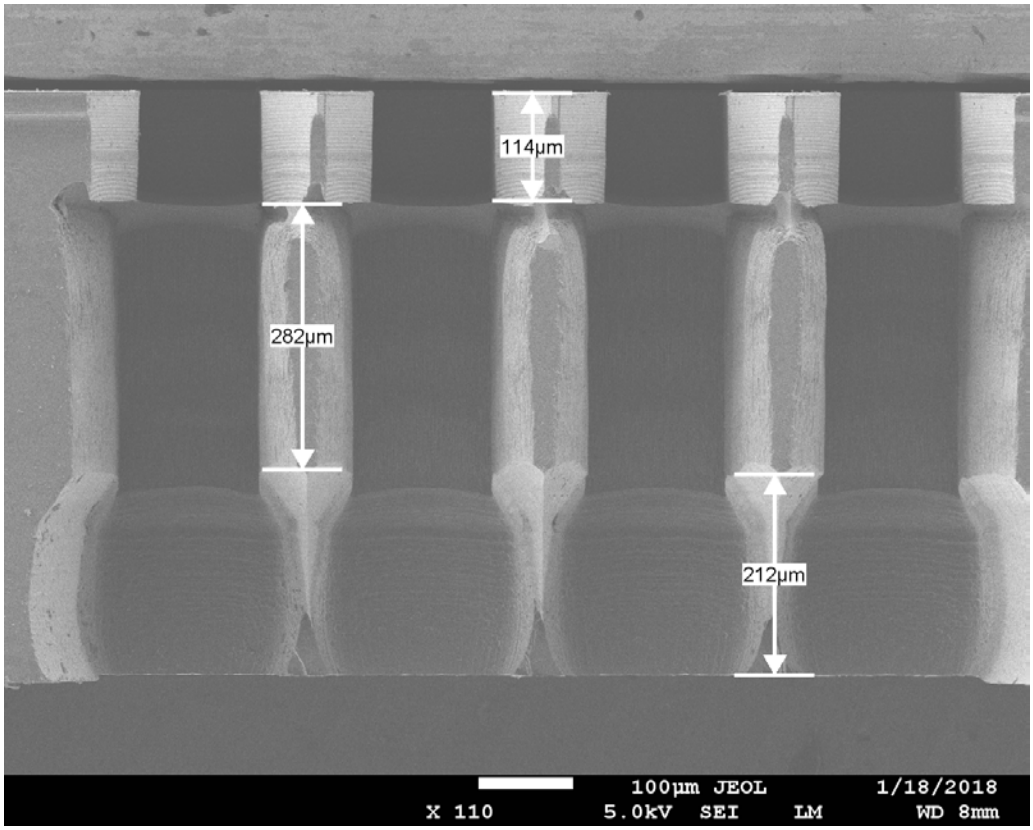

Figure 6.15 SEM image of the side-on profile of the through wafer etch

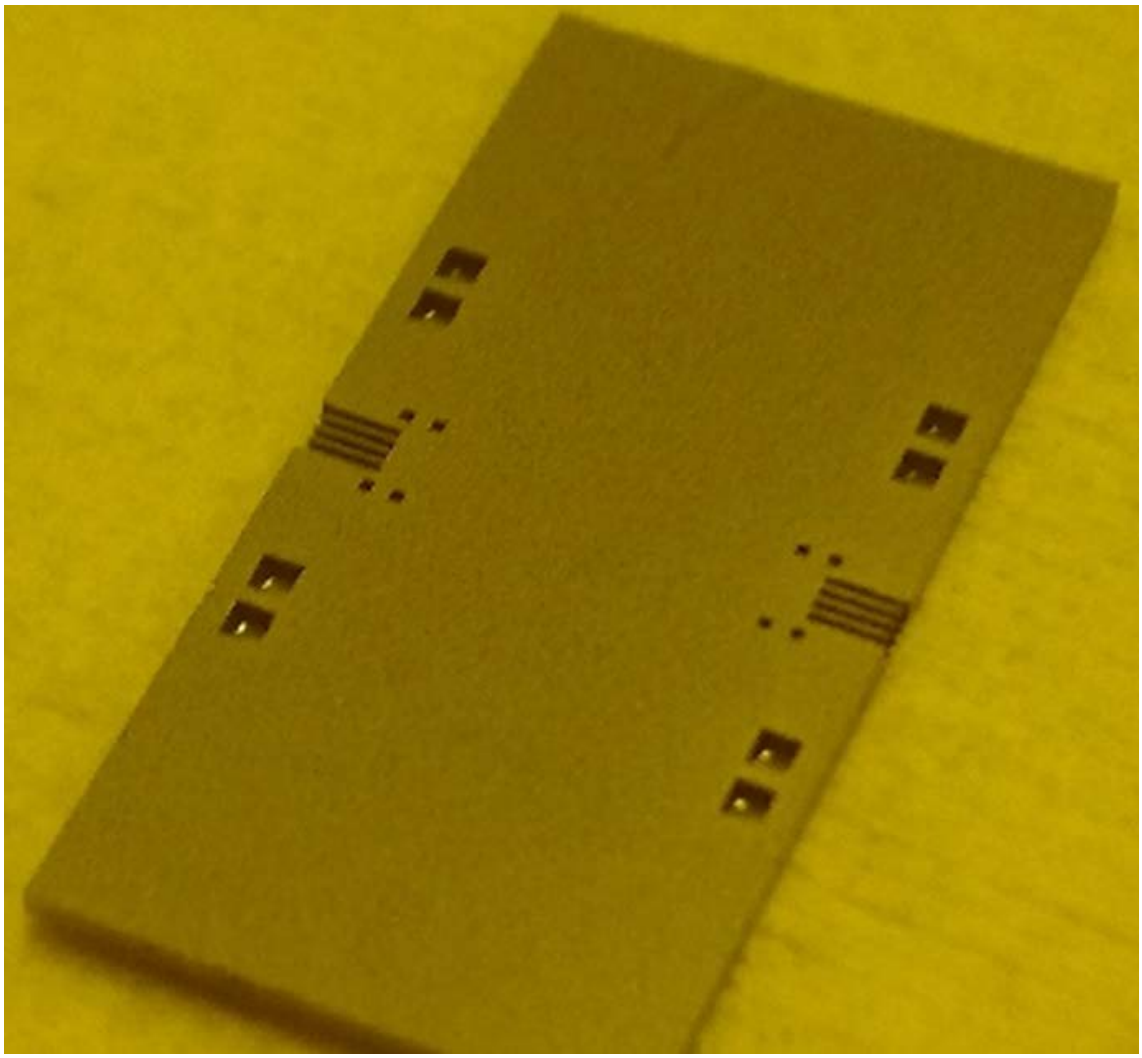

Figure 6.16 Photograph of the final Device Chip

6.3 Capping Chip Fabrication

In comparison to the Device Chip, the Capping Chip is somewhat less intensive from a fabrication point of view. The mask layout and designs for the Capping Chip have altered quite drastically from the beginning of the project, but the processing steps haven't changed extensively, in fact they have become less complex. Figure 6.17 is a table which shows the process flow for the capping chip in the same way as that of the Device chip.

Wafer	Location	Identifier	Photolith CGA marks								Photolith Trenches +Spheres								Nitride Strip		Deposition		
			AZ9260 Spin	Bake	Rest	Expose	Develop	Nitride Etch	Ash	KOH Etch	AZ9260 Spin	Bake	Rest	Expose	Develop	Nitride Etch	Ash	KOH Etch	Orthophos		Au Deposition	Lithography	Gold etch
1			✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓		✓	✓	✓

Figure 6.17 Capping Chip Process Flow

The processing for the Capping Chip mainly revolves around the KOH etching of silicon, as described earlier in the previous chapter. In order to create the hard mask for etching, a nitride layer is used. This was achieved in two ways. A factory made 400nm nitride-on-silicon wafers were used for much of the processing, but for tests "homemade" nitride wafers were used by using LPCVD nitride. For these test wafers, 400nm of nitride was grown onto plain silicon wafers using a furnace for 80 hours to create a suitable thickness. All other wafers were used "out-of-the-box".

6.3.1 Crystallographic Alignment marks

As discussed previously, KOH etching of silicon can be used to create V-grooves which can then be used to hold fibres. However careful care must be taken to ensure that the structures which are to be etched are aligned to the crystal planes.

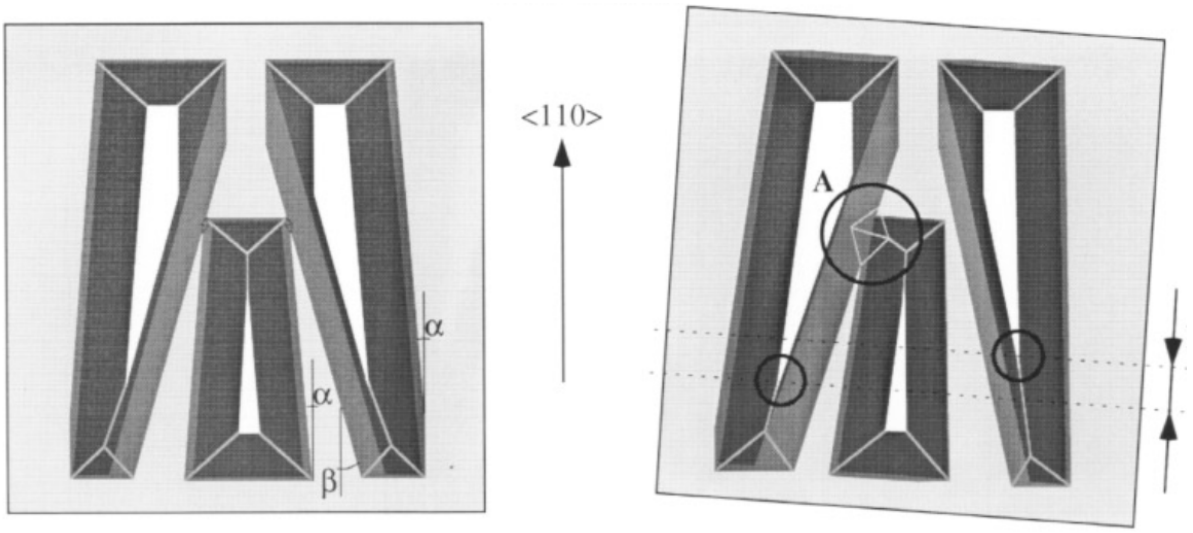

Figure 6.18 Diagram of crystallographic alignment forks [197]

The crystallographic plane needs to be accurately known in order to create v-grooves that etch in a straight line. In order to do this, a technique first published by

Vangbo and Bäcklund[197] can be used where they demonstrated a way to find the exact orientation of the crystal plane. By creating a series of "*alignment forks*" the plane can be seen clearly after these structures have been etched in KOH. Figure 6.18 [197] shows how this can be achieved. When the fork shaped structures are etched then it can be seen where there is asymmetry between each structure. If any given alignment fork is exhibiting asymmetry, then it can be deduced that this particular structure does not line up with the crystallographic plane as can be seen from the right-hand picture in Figure 6.18. However, in the left-hand picture it can be seen that this alignment fork does line up (or at least is very close to) the desired {110} crystallographic orientation. Hence, by using a series of structures like this it is possible to work out where to align a V-groove mask to. As long as each structure is aligned parallel to this fork then it will be etched straight. Over a length of many tens of microns the etched trench can vary in width and orientation by a degree which would render it useless for optical alignment, so this step is very critical for the whole process flow.

Photolith CGA marks							
AZ9260 Spin	Bake	Rest	Expose	Develop	Nitride Etch	Ash	KOH Etch
✓	✓	✓	✓	✓	✓	✓	✓

Figure 6.19 Process flow for alignment forks

Figure 6.19 shows again the relevant points from this section of the process flow. To begin with, AZ9260 is used to create a soft mask. The steps used create the soft mask are identical to that which was described in the previous section for AZ9260 (dehydration, spinning, baking, rest, exposure and development). Of course, the only difference is that the photomask used is different since it corresponds to

the alignment forks. By using the same developer, the development was clear and repeatable.

Next, the nitride layer needs to be etched to form the hard mask. Again, the same ICP nitride etch that was mentioned before was used. This process is not too critical, so long as the silicon substrate beneath the nitride is reached. A metrology box was used to confirm that the silicon substrate had been reached; this is simply a large rectangular portion of the wafer that is etched at the same time as the structures. This can be examined using ellipsometry to find its thickness. Once this was confirmed the soft mask needed to be removed since it is not compatible with KOH etching – the next step. As before, O_2 plasma ashing was used to remove any trace of the soft mask. After 15 minutes of ashing the resist is totally removed, and it can be etched using KOH.

It is important to have good control over the KOH process to ensure that accurate depth of the forks and V-grooves is obtainable. The KOH solution was used at 40% concentration and was heated to 80°C. Initial tests were done at 70°C but etching was considerably slower. Using 80°C at 40% concentration, the KOH etch had the useful property of an etch rate that was very close to 1μm/min, making etch times simple to track. The exact etch time would vary slightly (±0.1 μm/min) depending on feature shape but using 1μm/min as an etch rate proved successful. This was aided by a MATLAB script that was written to estimate the etch depth of different features based on a specific KOH temperature and concentration, which can be found in the Appendices. After etching, the wafers were washed in a weir for over 10 minutes until the resistivity indicated that any ions from the etching process were flushed away. Then they were blow dried and then finally spin dried.

In the case of the alignment forks, the depth is not critical, so long as the shape of the forks is visible. An etch depth of ~40μm was sufficient. If etched too much then the fork would lose its shape and no longer work, but too little and the fork cannot be seen clearly enough, especially when located under several microns of resist.

6.3.2 Fibre trench Fabrication

The next step was to create the V-grooves and other associated structures onto the surface of the wafer. The cropped process flow in Figure 6.20 indicates how this was done.

Photolith Trenches +Spheres							
AZ9260 Spin	Bake	Rest	Expose	Develop	Nitride Etch	Ash	KOH Etch
✓	✓	✓	✓	✓	✓	✓	✓

Figure 6.20 Process flow for V-grooves

All the processes for this stage are identical to those discussed above, except that the etch time differs and of course the photomask is different also.

Once the resist had been spun onto the wafer, the mask aligner is used to align the photomask, not to the wafer as such, but to the correct alignment mark which shows the most symmetry. Using this method yielded good results in this project and the KOH-etched V-grooves were aligned accurately.

Once the hard mask had been created the wafers were etched in KOH for 80 minutes to produce an etch depth required of 80μm.

6.3.3 Gold Deposition and Lithography

The samples were cleaned thoroughly before gold deposition. Initial tests showed week adhesion between the gold and wafer surface which was remedied by first placing a thin ~10nm layer of titanium on the chip, and then depositing gold onto the surface. This was carried out using BAK700 deposition chamber at extremely low pressures. A final lithography step was used to remove the unnecessary gold from the rest of the chip, such that only the mirror area was left remaining. This was

simply done with photolithography (AZ9270) and then gold etchant is used to remove the remaining area. The gold etchant is a chemical etch that comes in a premade solution. Finally, the samples are rinsed and dried. In the same manner as before.

6.4 Final Capping Chip

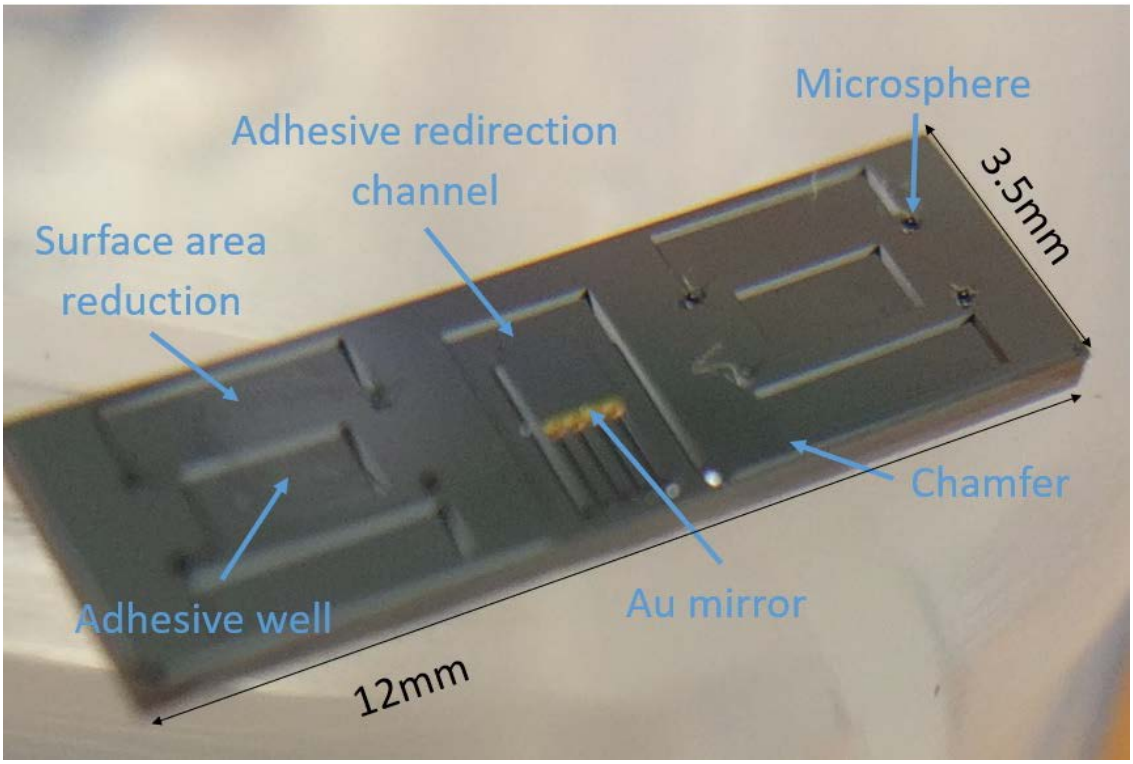

Figure 6.21 Photograph of the final Capping Chip

Figure 6.21 shows the final Capping Chip after all processing is finished. Notice the different sections of the chip: the adhesive well, surface area reduction, adhesive redirection channel, chamfer and gold mirror. The dimensions of the chip are given for reference as well.

Chapter 7 Assembly and Characterisation

7.1 Assembly Procedure

Once the two chips of had been fabricated they needed to be assembled along with the input fibres. The assembly procedure was carried out in part by Optocap and in part internally. Optocap were valuable consultants on this procedure as it falls far outside the scope of typical university-based silicon photonic research but was necessary for this project. This first section will outline the assembly procedure as it was performed internally.

Once the chips had been diced, cleanliness was vitally important. As soon as the samples were removed from the cleanroom environment, they became susceptible to foreign contamination. Even the dicing process is considered “dirty”. For a standard silicon PIC, the only real area of concern is the coupling area. Surface grating couplers or facets must of course be free from debris or contamination which would cripple coupling efficiency. However, once the light is coupled into the waveguide, the cladding layer of oxide should protect the propagating optical mode from any interference from a contaminant. When the waveguide is being used for sensing applications, this is more problematic since the surface of the chip will use be coated in a substance for sensing. But even in this case, the majority of the chip surface is unimportant, only the waveguide area. Whilst it is always a good idea to keep the chip free from contamination, typical dirtier processes such as dicing are not normally as issue.

However, in this project since the chip will be bonded to another, any amount of debris will prevent the two chip surfaces being brought flush to each other. Therefore, cleanliness was of the utmost importance. The best way to limit contamination is to reduce the number of times the chips are handled. Steps were taken in order to reduce the total amount of handling required for all the assembly. Before leaving the cleanroom, all samples were thoroughly washed and blow dried

to remove any particles that may have come through the dicing process. However, despite this, some samples still showed signs of contamination. Therefore, Optocap were able to provide dicing services for this project which used a special "snow-cleaning" method which involved an extremely high volume of water and air pressure during dicing to flush and remove any debris that is formed during dicing. The samples were then sent back to Southampton for the next stage of the assembly.

Ideally, assembly would take place within a cleanroom environment as well but unfortunately it was difficult to gain permission to do this because of the extra equipment needed and relevant safety forms that needed to be completed to allow the epoxy to be brought into the cleanroom. Because of this the chips were

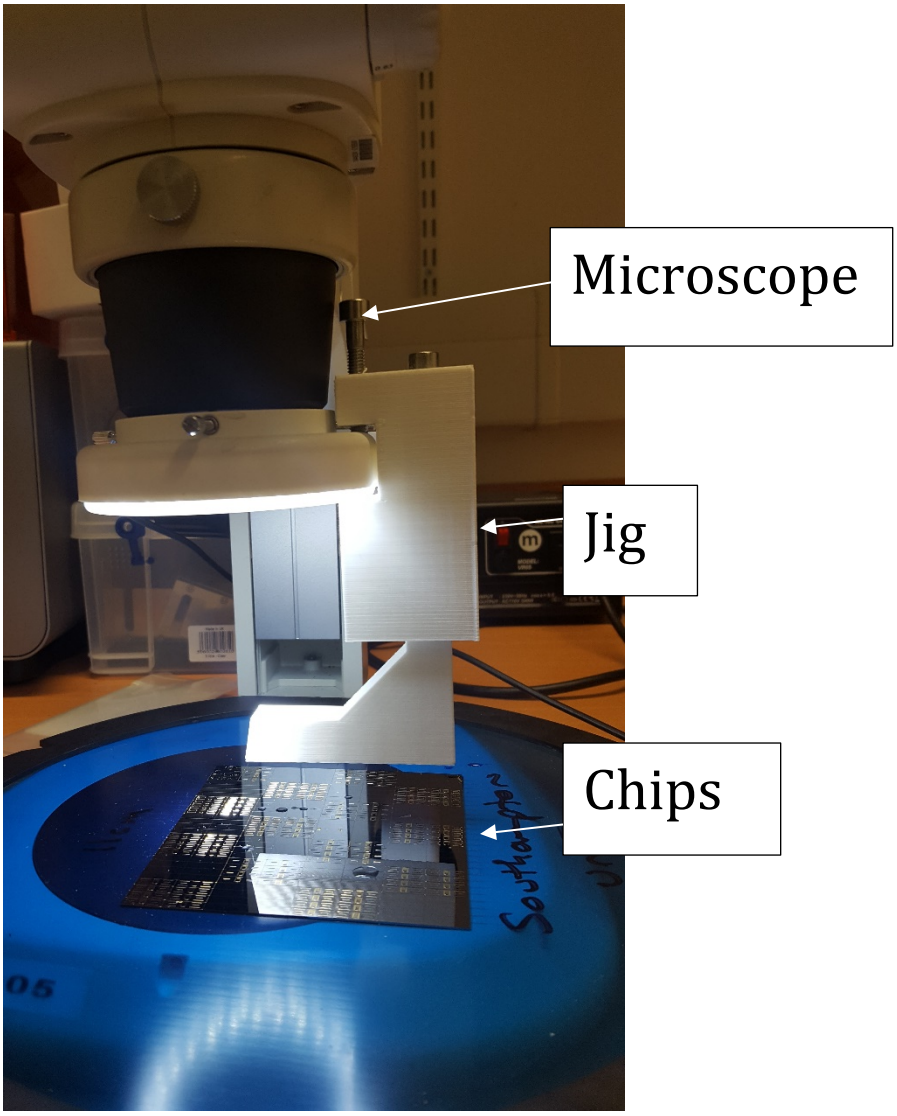

Figure 7.1 Photograph of the setup used for sphere population

processed within an optical lab, which still ensured stable temperature, pressure and humidity.

Before the chip attachment could take place, the capping chip holes needed to be populated with the microspheres. Without any special equipment for this, some customised tools needed to be fashioned. A custom jig was designed in conjunction with Dr Reynolds, which he then used to 3D print it.

Figure 7.1 shows how the jig was used. It attached to a conventional microscope with a bolt and then would hang down to the focal point of the microscope. At the focal point, a small piece of scrap optical fibre was attached using epoxy as an adhesive. In this way, the microscope focussed onto the tip of the optical fibre. The microscope could be raised or lowered such that the jig and the fibre tip would be raised by the same amount. This setup was used in order to populate the holes with spheres.

The first step of sphere population was to put a small amount of adhesive epoxy into the hole. This took some practise as one can imagine. Too little epoxy and the sphere would not stick, but too much and the sphere would not sit flush to the edge of the hole. Figure 7.2 demonstrates this problem. The image to the right shows the instance when there is too much epoxy. The sphere will not sit flush and will float up out of the hole. Also, excess epoxy spills out onto the chip surface. This is simply not acceptable. It was found that the correct amount of epoxy could be deposited in the following manner. First, the fibre was dipped into a pool of wet

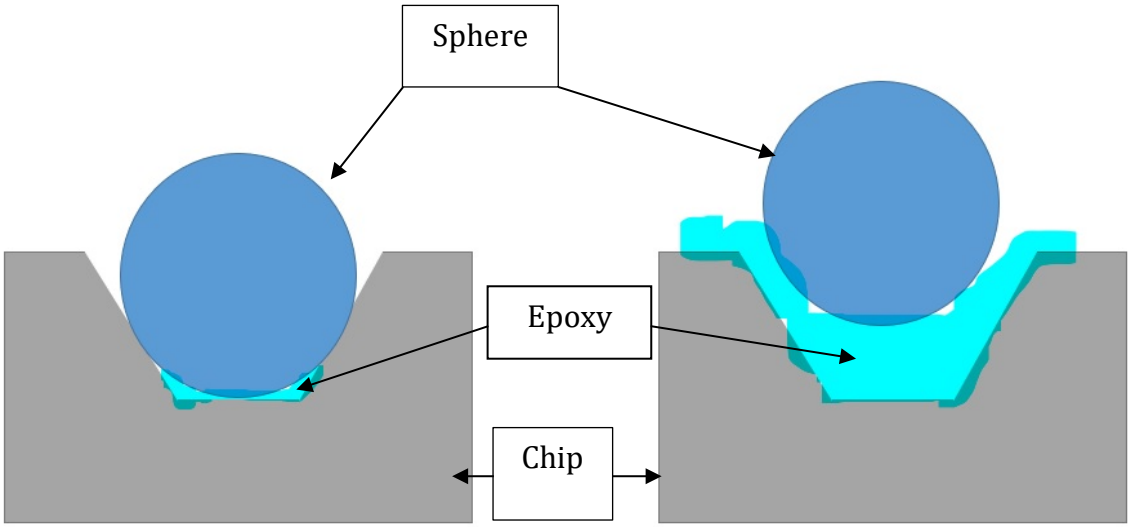

Figure 7.2 Comparison of epoxy levels

epoxy for a brief moment, ~0.5s. It was done in such a way that the surface tension of the epoxy pool was not properly broken. In just this brief moment, a small quantity of epoxy is wicked up the fibre. By leaving the fibre in the epoxy pool for only a very brief time and by not breaking surface tension, a roughly repeatable quantity of epoxy would be drawn onto the fibre each time, though the exact quantity cannot be accurately known. To deposit the epoxy into the hole, the fibre was placed down into the hole. Once the fibre made contact with the bottom of the hole, the epoxy would then be deposited in a droplet that would sit in the hole. Almost all of the epoxy would be removed from the fibre such that this process could be repeated roughly 20x before cleaning was required. When necessary, the fibre was cleaned with a microfibre wipe and some IPA solvent.

The same jig could be used to deposit the sphere as well. A number of spheres were poured onto the chip surface of an unused chip at the edge of the array of diced chips. The fibre on the jig was then lowered onto a sphere. Visual inspection at this point was used to confirm the sphere was uniform in shape and size. Any deformed spheres (of which there were surprisingly many) were rejected at this stage. From the batch of spheres used, they were either perfect in shape and size, or very obviously incorrectly formed (with a bulge on the side for example). Once a correctly shaped sphere was located, the fibre was used to make contact with the sphere. Static charge on the dielectric spheres enabled the fibre to pick up the sphere. At this stage the sphere is held relatively securely and can be placed into the hole. Once the

sphere sits in the hole, the adhesive properties of the epoxy in the hole are stronger than the static charge and after the sphere is dropped into the hole and it remains there. This is then repeated for all holes which need sphere population.

Finally, after all the spheres are located, UV light is used to cure the epoxy, which is essentially a non-reversable process and must only be done once all spheres are in place. A simple UV diode soldered onto a PCB with a battery and switch was used for this. This whole procedure is somewhat unrefined, but it could all be easily replicated with industrial pick and place tools and a precision UV lamp. An alternative would be to use a mask aligner to flood the chips with UV but cleanroom regulations prevented this from happening. During the assembly procedure, lower light levels were used, and the samples were kept in the dark when not being worked on, so as to prevent any undesired curing. However, it was found that the energy needed to actually cure the epoxy was much higher than was present in ambient light, such that yellow-room conditions were not strictly necessary.

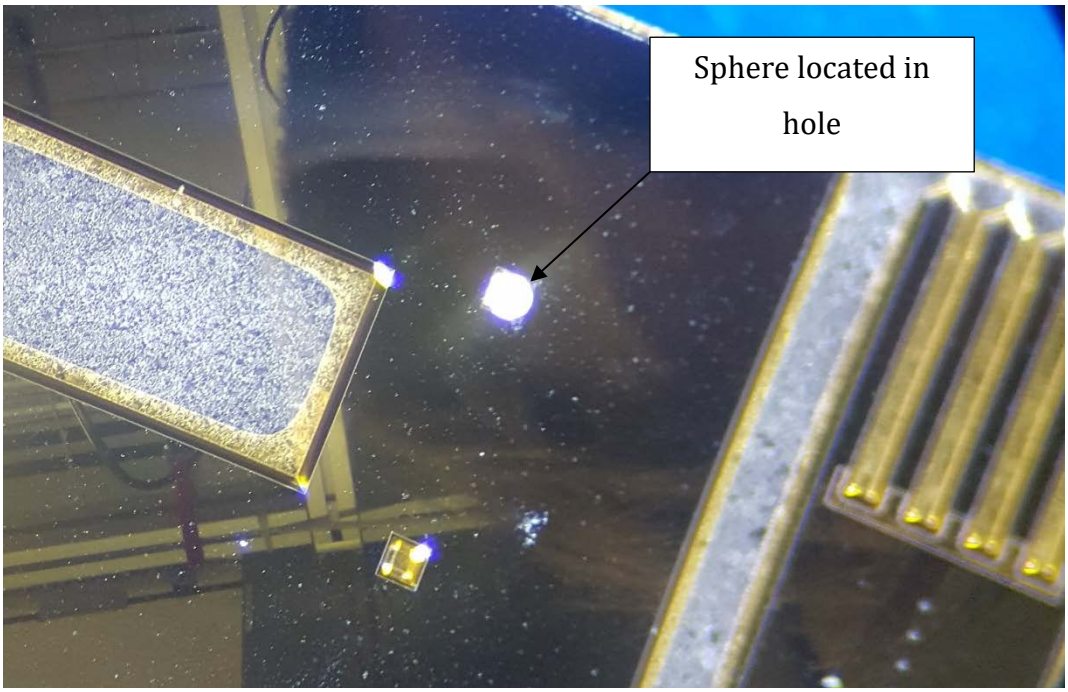

Figure 7.3 Photograph of a sphere in a hole next to an empty hole

Figure 7.3 shows an image of a sphere once it had been placed within a hole. The reader must appreciate that no proper imaging equipment was available for this picture and this is simply the product of photographing the microscope lens. Even

so, it is possible to see the result. Once all the spheres had been populated, the chips could be sent for bonding and fibre attachment.

7.2 Characterisation of Individual Parts

7.2.1 Optical Characterisation

In order to properly characterise the grating couplers for this packaging method, several other metrics are necessary. The propagation, bending and tapering losses must also be ascertained, such that remaining loss can be attributed to the grating couplers.

Each device chip had several test structures that were used to obtain these details. Each chip had two banks of test structures, both at the top and the bottom. In this way any variations across an individual chip can be accounted for and averaged.

7.2.1.1 Propagation Loss

A very simple structure can be used to calculate the propagation losses of a waveguide. By having an array of waveguides that vary only in length, the propagation losses can be measured by using the cutback method[198]. This is calculated on the assumption that losses are only introduced due to scattering in the waveguides and that roughness can be treated as an averaged constant. There are no sources of gain in any of these circuits so only losses need to be considered.

Figure 7.4 shows an image of the mask used. The reader should appreciate

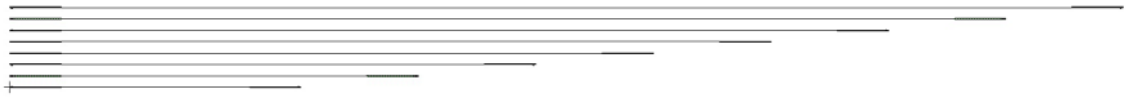

Figure 7.4 Mask image of propagation testing structures

that it is difficult to accurately display microscopic structures such as these without blowing the image up very significantly. However, the pattern can be seen that each waveguide varies only in total length. The same tapers and grating couplers were used for each waveguide so as not to impact the results.

This data was taken from both the top and bottom structures from each chip and across four different wafers.

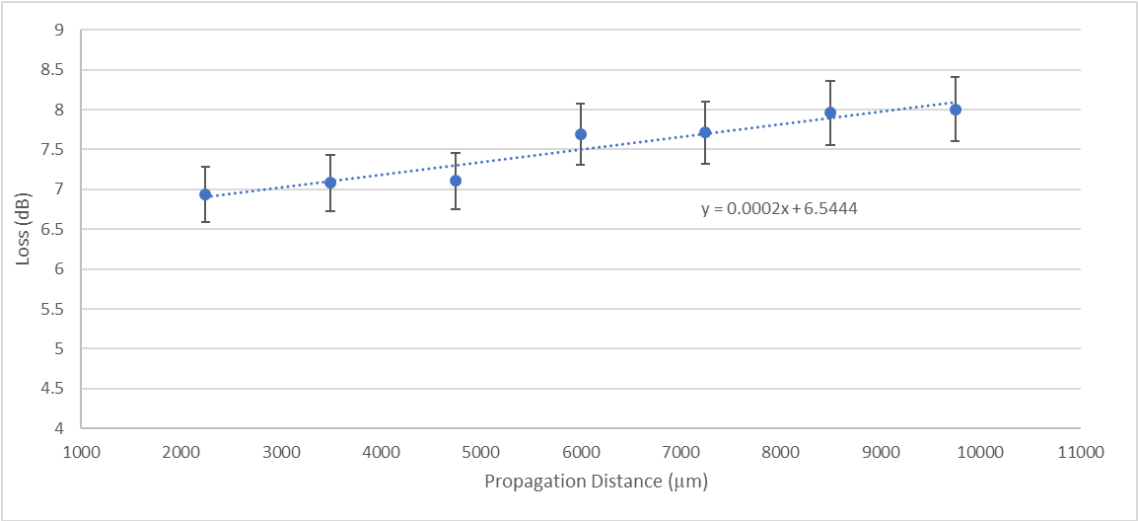

Figure 7.5 Graph of Propagation Losses

Figure 7.5 shows a graph of the propagation loss. The mean was from each section on each chip. This formed a normal distribution. Error bars of are shown of ±σ. By drawing a trend line and taking the gradient, the propagation loss can be calculated via a linear regression model. The trendline has a gradient of 0.0002±0.000019 or a propagation loss of 2.0±0.19dB/cm. The y-intercept would give the additional losses attributed to the gratings and tapers; however, these were not immediately of interest since they were stock forward grating couplers that were designed for a previous project. A full spectrum was taken but the data points shown here pertain to a wavelength of 1550nm. Before taking measurements of insertion losses, the system was normalized against a single piece of fibre running through the polarization controller and back to the detector. That way, these losses are not included in the measurements. All measurements were taken in lab conditions at 25°C.

7.2.1.2 Bend Losses

In a likewise manner the bend losses were calculated using an array of different waveguides. Figure 7.6 shows the mask layout. Each successive waveguide in the array has additional bends, with the radius of each bend being 50μm. The total

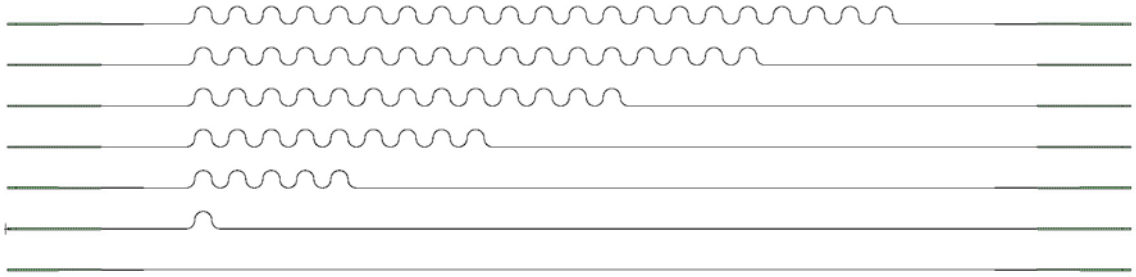

Figure 7.6 Mask image of bend loss structures

propagation distance is not constant, so the losses must be offset by the difference in distance. It is possible to create structures that have increasing numbers of bends that remain constant in length by means of more elaborate shapes. However, this increases chip footprint and so was not used here. Figure 7.7 shows a graph of the bending losses.

Again, the gradient will give the loss per bend after being adjusted for propagation losses, which in this case is 0.029±0.00192dB/bend, which agrees well with the predicted value, which suggested that there would be negligible loss for

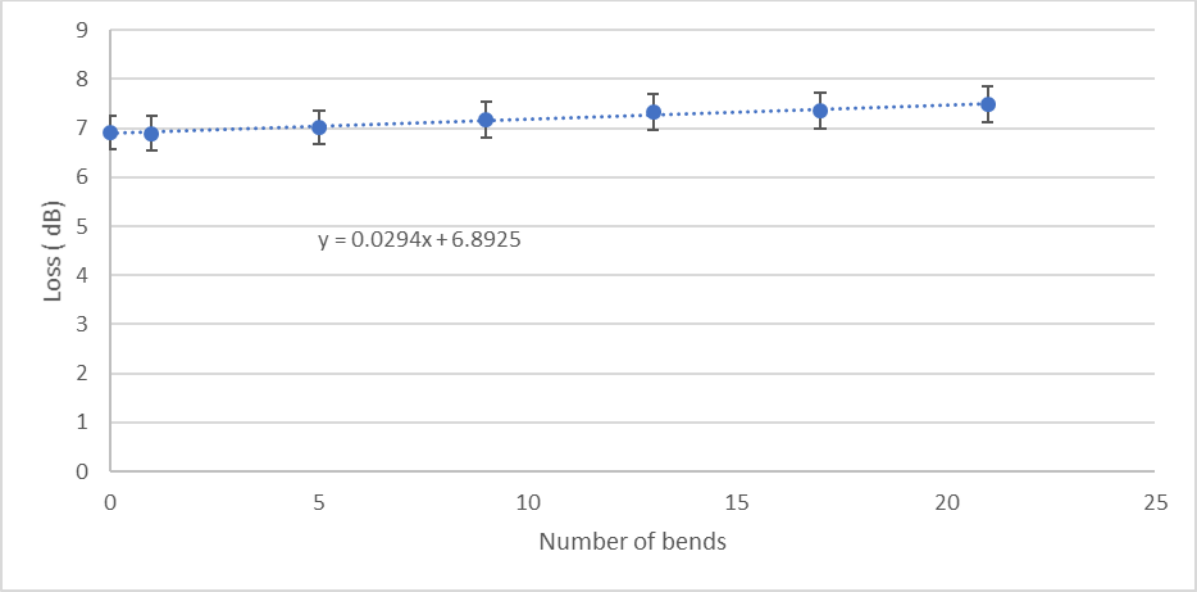

Figure 7.7 Graph showing losses per bend

bends of this radius. The same averaging method was used. The σ bounds are shown again for completeness.

7.2.1.3 Taper losses

Once again, the process to measure the taper loss is the same. An array of waveguides is made, each with identical length. However, they are made up of increasing numbers of tapering sections along their length, such that each successive waveguide section is made up of an increasing number of tapers.

Figure 7.8 shows an image of the mask that is a little unclear, however the wider sections of the end of the tapers can be seen, as they increase in number for each successive waveguide.

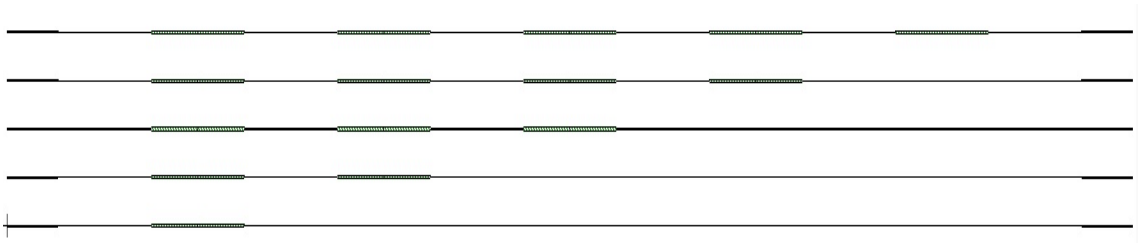

Figure 7.8 Mask image of taper loss structures

Figure 7.9 shows the loss attributed to each taper section. This graph shows more variation than previous graphs, however the pattern can be seen. Again σ uncertainty bounds are shown here. The same averaging technique and linear

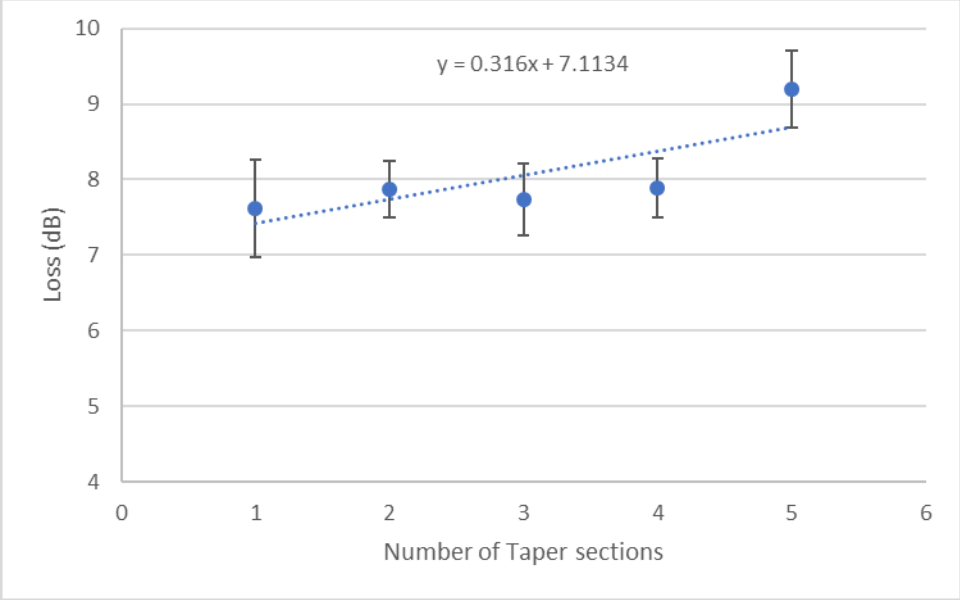

Figure 7.9 Graph of the Taper losses

regression was used as previously. The gradient gives the loss per taper section which in this graph is 0.31±0.14dB per taper section (tapering down to waveguide width and then back up to full width of grating coupler). The relative uncertainty in this calculation suggests that the fabrication process for these tapers was not as repeatable as other devices.

7.2.1.4 Grating Coupler Losses

Figure 7.10 Mask image of the grating coupler sections

The final optical device to be characterised were the specially designed reverse grating couplers for the packaging method. Since they were reverse grating couplers, it was a little more difficult to measure them with a conventional setup, so the mask design was adjusted to compensate. Figure 7.10 shows the mask design, which includes 50μm bends to allow the gratings to be characterised the normal way. The tapers, bends and propagation losses have already been characterised such that they can be subtracted off the measured loss of these waveguides, giving the grating coupler loss remaining. Figure 7.11 shows the spectrum of the measured grating coupler, with an identical simulated grating coupler plotted alongside for comparison. Peak efficiency is 51.3% (2.89dB) which indicates additional loss with the measured example, but the general shape and bandwidth is very comparable. The bandwidth remains within 5% of the predicted values of 35 and 59nm for 1 and 3dB points respectively.

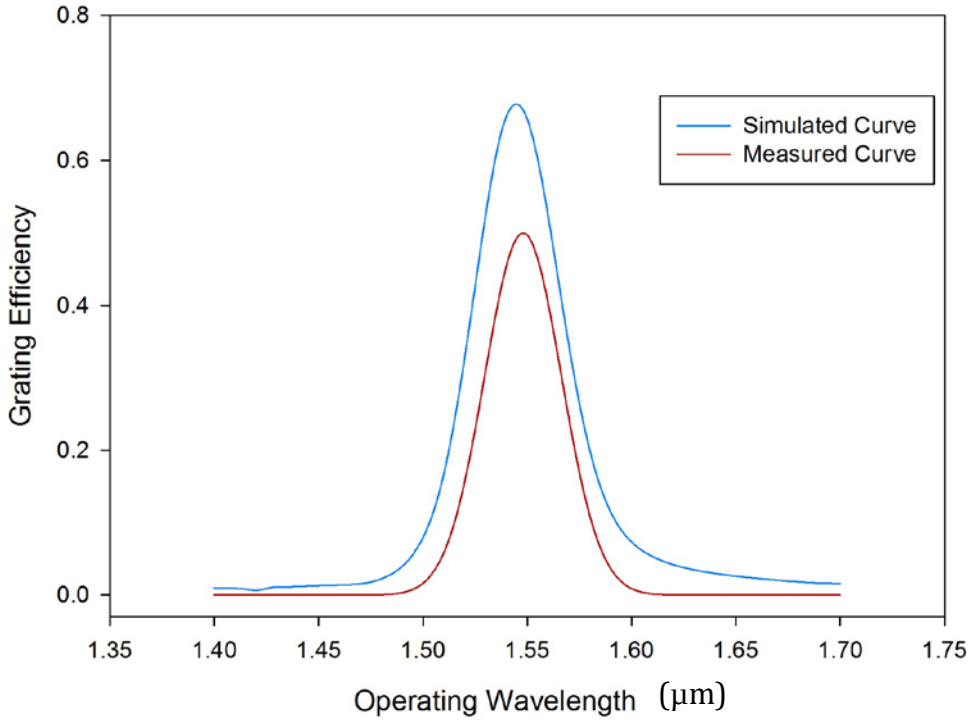

Figure 7.11 Graph of characterised grating coupler compared with the simulated and measured example

7.2.2 Assembly Tests

This section outlines how the assembly method was tested. Due to the specialist equipment that was needed, this work was carried out on request by Optocap. The goal of these tests was to ascertain the mechanical stability of the proposed packaging method, and the feasibility of using the proposed design, ensuring that the parts involved could be mated together effectively and then not move. It also was used to find the best candidate for the epoxy, both in terms of chip adhesion and for optical properties.

Two main trials were conducted to devise a way of assembling both chips together. If it was determined impossible to effectively bond the two respective chips, then this would severely limit the feasibility of the project. Figure 7.12 shows how it was initially conceived to happen.

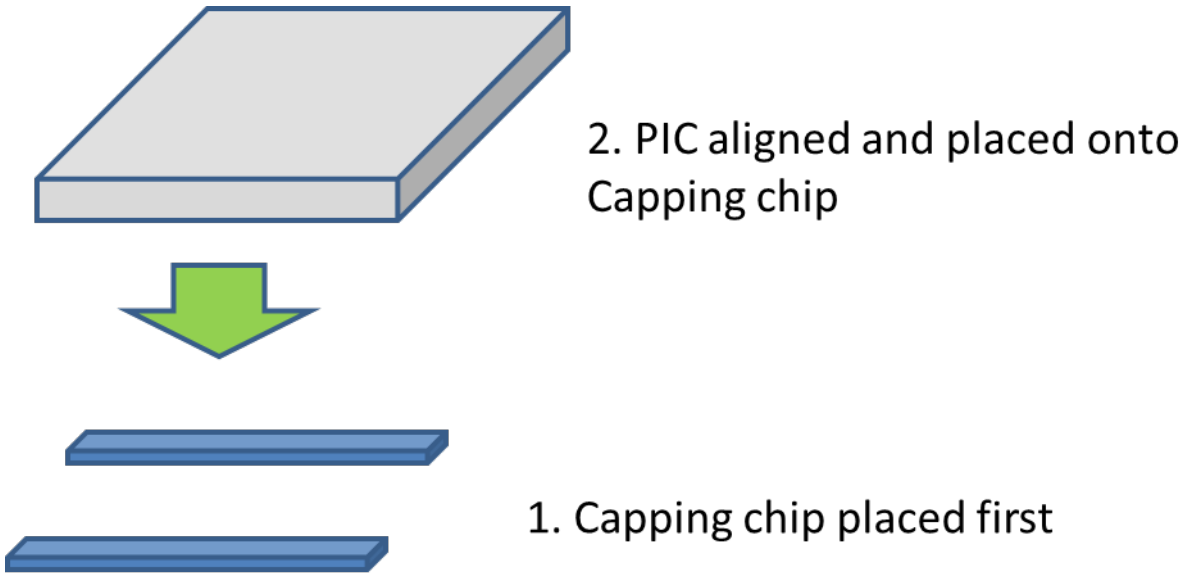

Figure 7.12 Diagram showing the initial assembly attempt

This approach however needed to be abandoned due to the constraints with the machinery that was available. The small size of the holes proved to be too challenging to use for the insertion of the epoxy.

After the initial trial was unsuccessful, the chips were redesigned to allow for the assembly to be possible. These changes were previously outlined in the mask design section.

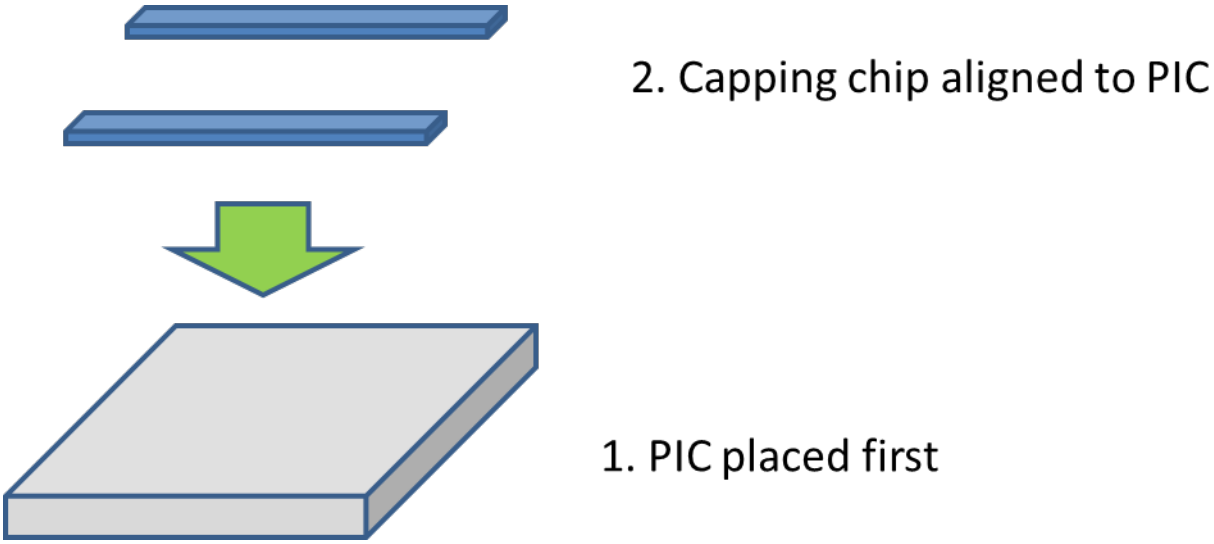

Figure 7.13 Diagram of the second attempt to assemble both chips

Figure 7.13 shows how the second attempt to attach both chips was made. By reversing the order in which the chips are attached this provided greater access for the insertion of the epoxy. The addition of the chamfer that was included in the mask design then allowed easier application of epoxy. Figure 7.14 shows how the epoxy could now be applied after designing the chamfer with the right dimensions with respect to the viscosity of the epoxy such that it would form an effective bond. In addition to this, the epoxy that was placed on the chamfer could easily be UV cured to provide an initial adhesion, which was then used as the foundation for thermal curing.

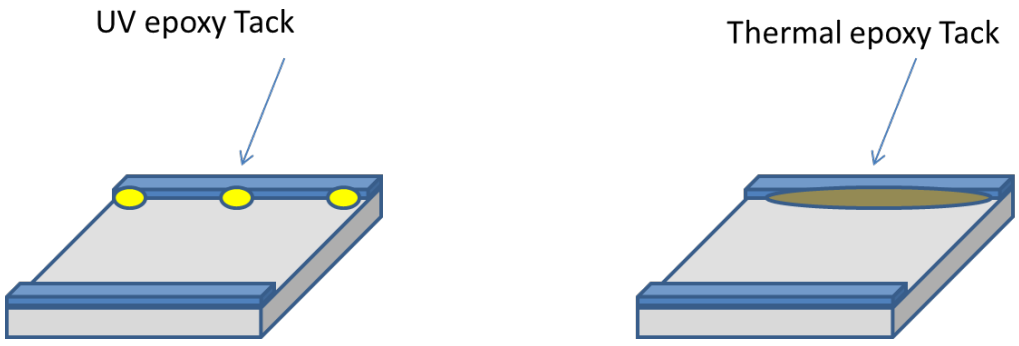

Figure 7.14 Diagram showing the application of epoxy to the chamfer

The adhesion holes on the chip were accessed with a low-viscosity epoxy, that would not be suitable for the chamfer. Figure 7.15 shows how this was accomplished.

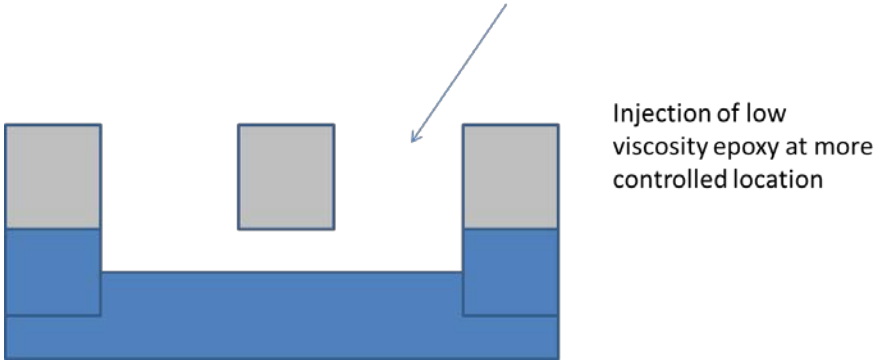

Figure 7.15 Diagram showing how lower viscosity epoxy was used for the alignment holes

The chips were able to be successfully aligned using this method. Figure 7.16 shows the alignment that was possible. The two chips were able to be placed within the tolerance that was expected (±1.5µm). Specialist cameras were used by Optocap to show obtain these pictures and measurements.

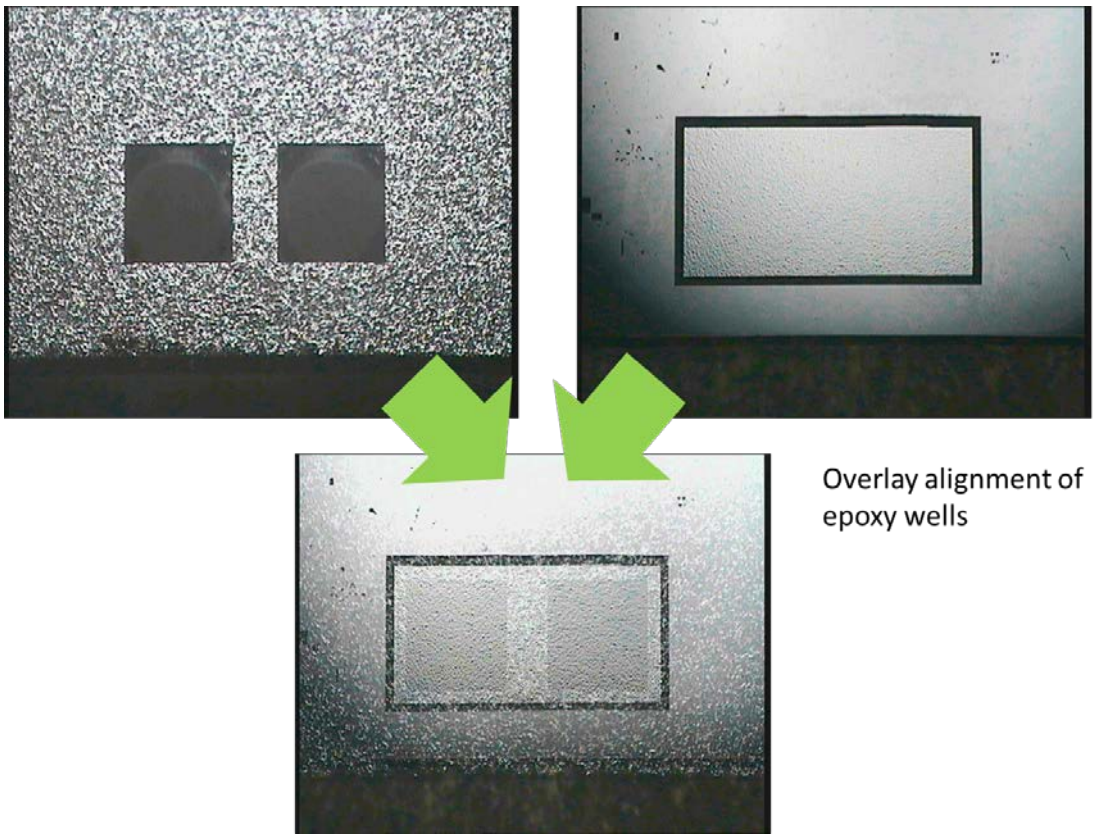

Figure 7.16 Images showing alignment of the two chips

When both chips had successfully been aligned, an array of fibres in a block are inserted into the channels as shown in Figure 7.17. The addition of the fluted fibre channels allowed the fibres to be inserted with only basic alignment equipment, done by inspection. Before the widening of the fibre channels this was found to be difficult or not possible by Optocap during initial assembly trials. This test also shows the need for the adhesive well around the fibre trenches. As it can be seen, there is some overflow of epoxy around where the fibres sit. The well was designed to drain some of this excess and prevent a non-flat surface to mate the other chip to.

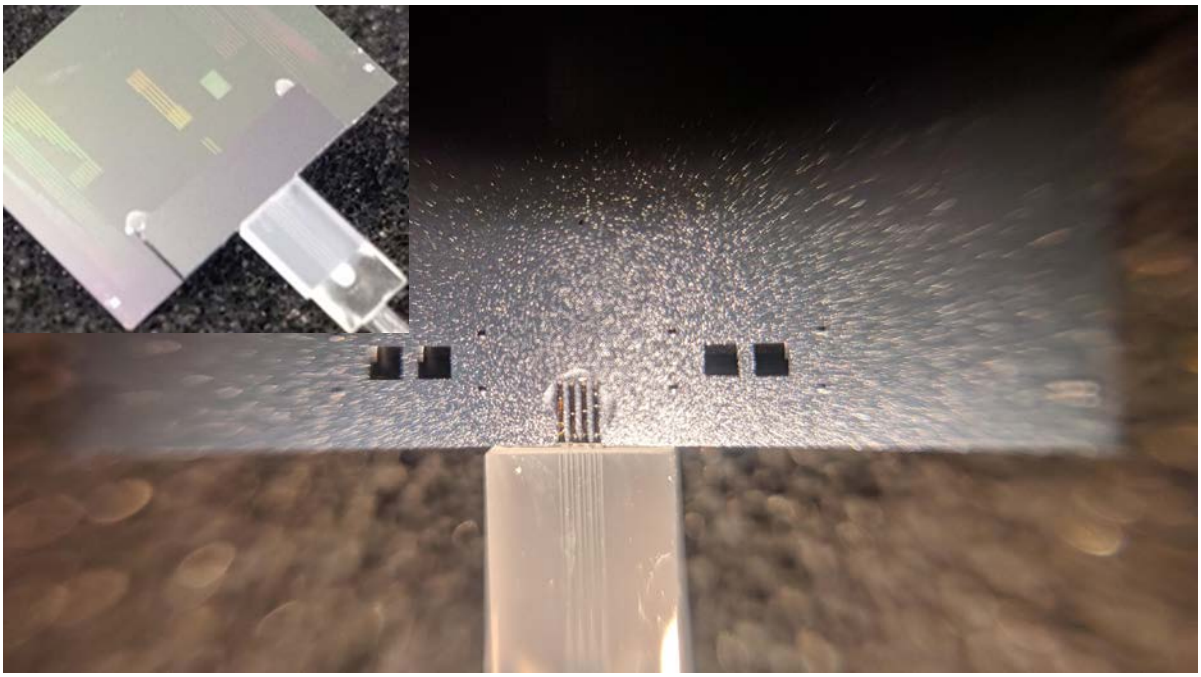

Figure 7.17 Image of a fibre array inserted into the assembled chip

Once successful alignment and bonding was demonstrated, tests were conducted to show the rigidity of bonded chips.

Figure 7.18 shows some of the sampled chips that were used for testing. Stress tests were used to determine the strength of the bonded chips. Vertical stress was applied to the Capping Chip, in such a way as to try and separate the chips by lifting them apart. The chips were tested up to an effective weight of 36kg without moving, at which point the chips shattered as shown in Figure 7.18. This is significantly higher than 5kg which is the industry standard. Therefore, the weak point of this assembly method is in fact the fragility of the chips themselves, not the bonding. Indeed, the bonds remained intact even when the chips themselves shattered as can be seen in the image below. To obtain these results UV Snap Cure was used for the UV cured epoxy, 353NDT was used as a thermal epoxy and Dymax OP21 was used for the optical adhesive.

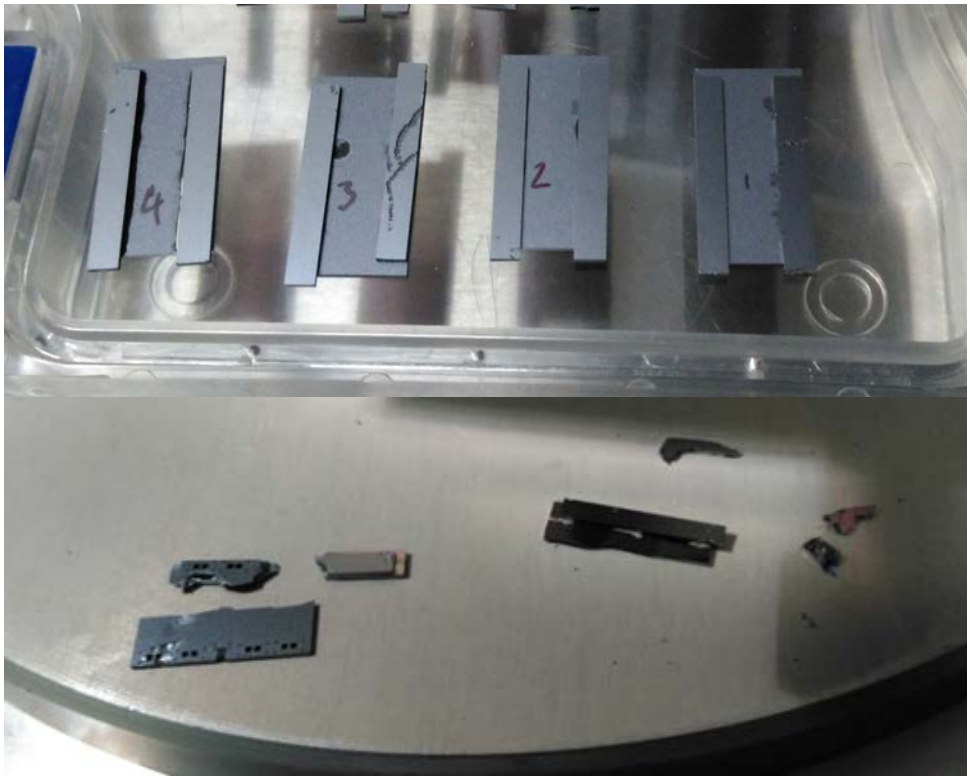

Figure 7.18 Image of Bonded Samples of both the Capping chip and Device Chip after testing

7.3 Assembly Failure Analysis

Unfortunately, due to several constraints, it was not feasible to complete the assembly to a satisfactory level during the project. As such a working demonstrator was never produced. This section will analyse the root cause for this, and highlight some key ways in which it can be improved to allow for assembly in the future.

Table 7.1 Comparisons of different causes of contamination

Tool	Level of Contamination Risk	Number of times used in processing	Comments
Plasma Tools	Low	9	Low pressure sealed environment
Lithography	Low	8	Non-contact
Development	Medium	7	Cleanliness of dishes hard to know
Wet Etching	High	5	Contamination builds up in tank
Dicing	Very High	2	Having two chips doubles the amount of dicing
Assembly	Very High	2	Each chip needed to be assembled by hand in a non-clean room environment

The main issue stemmed from contamination of the samples during the many stages of fabrication. Fabrication will inevitably introduce some level of contamination. The Appendix highlights a number of ways in which the cleanroom environment is designed to mitigate against contamination. However, despite best efforts, samples which were satisfactorily clean were not realised.

Each of the sources of contamination are shown in Table 7.1, along with an estimate for how likely they are to introduce contamination. As can be seen, there are two important factors to consider. Firstly, the sheer number of processes that were necessary in order to fabricate both the chips. Certain forms of contamination can be removed along the way by cleaning. However, there is a small but non-negligible amount of contamination that could not be removed. This contamination accumulates along each stage. Whilst it is small, it would have an effect. Secondly, there are several exceptionally dirty processes. Particularly the KOH etching, dicing and assembly stages. KOH etching has been shown to cause contamination from iron oxide particles which can be very difficult to remove[199]. Dicing is a dirty process for obvious reasons since debris can be scattered by the dicing saw. Assembly in this project was a dirty process since it was undertaken outside of cleanroom conditions.

Whilst contamination is present in all fabrication, it is of particular concern in this project. This is for a few reasons. Firstly, since there are two chips, the risk of contamination is essentially doubled. In any given batch of wafers, some will be found to be unsatisfactory due to contamination. However, this is usually not a problem since yields are typically still high enough. Secondly, since both chips must be mated together perfectly flush, there cannot be any excess contamination that builds up over the surface of the wafer. For a normal chip, the user is only interested in a very small portion of the chip, there rest of which can be contaminated with very little impact on the performance of the circuit. However, for this project the entire surface must be kept as clean as possible. This is especially apparent due to the V-grooves which begin to act like a sink for debris and dirt. Then when the fibre is pushed down the trench during assembly, it will pick up all of that dirt and move it towards the vital mirror area. All of these effects in conjunction, create a real risk of contamination. Figure 7.19 shows an example of the typical levels of contamination

on a chip. With levels of contamination this high, it would not feasible to assemble both chips together for a working prototype. The contamination particles were all

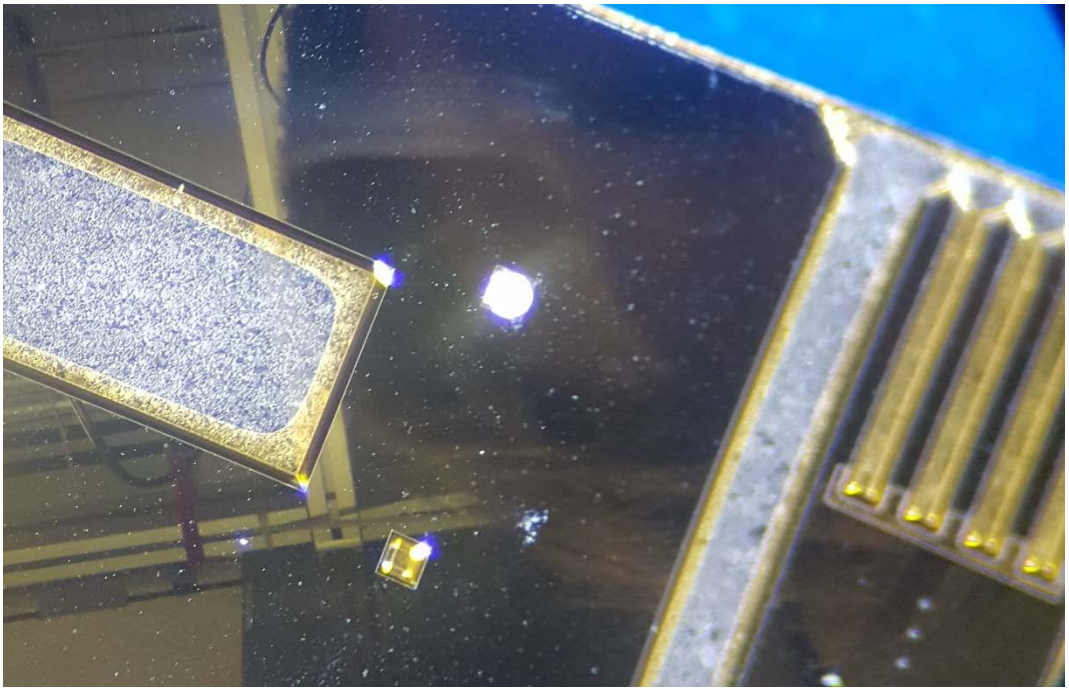

Figure 7.19 Image of Capping Chip showing the accumulation of contamination

of the order of 1µm. Particles larger than this could be removed and were not present in the cleanroom anyway. Particles smaller that this were not such a concern. Given that packaging method requires an alignment of ±3-4µm, only a small number of particles would cause catastrophic consequences. These particles become adsorbed onto the surface of the silicon by means of electrostatic forces[200].

The presence of such particles made attempts to produce fine alignment futile. Until a suitable method to remove or prevent the extra contamination can be found then assembly cannot be completed.

There were several attempts to mitigate the magnitude of the contamination. Firstly, the recesses in the mask design were introduced into the mask design. This did improve the situation somewhat. However, a qualitative assessment of the capping chip after this introduction still showed that an excess of contamination had built up in the important fibre channels. An extra attempt to mitigate the contamination was made by the suggestion of Optocap. This involved completing the dicing stage (which is most likely the dirtiest) at their premises and using a CO_2

'snow' clean[201]. In this process, high pressure CO_2 is directed towards the wafer and freezes the contaminants which then break off and are removed. This was done to the wafer before dicing. However, despite this procedure, the sum total of contamination was largely unchanged. Inspection before and after indicated that the snow clean was mostly ineffective at removing the contamination on the wafer. The subsequent dicing process then also created more contamination particles. An extra factor to consider is that a protective resist layer could not be applied in the case of the Device chip. Typically, a resist layer can be spun onto the wafer before dicing which would protect the surface. After dicing, this layer can be dissolved away, taking the contamination and debris with it. However, the Device chip needed holes to be etched through the whole wafer. This meant that it could not be used with the conventional spinner since there would be a vacuum leak when the wafer was placed onto the spinner chuck. Therefore, it needed to be diced bare. This further caused this wafer to accumulate debris.

Because of these factors, it was not feasible to attempt a final assembly and test, though it is necessary to point out that with more time, the effects of contamination could be substantially reduced. This is discussed in detail in the Future Work section of the next chapter.

Chapter 8 Summary, Conclusions and Future Work

8.1 Summary

Silicon Photonics will no-doubt become increasingly adopted as a means of meeting the data demands of the future. Photonics is being utilised in shorter reach networks such as data centres and potentially into homes and businesses. Innovative and adaptable silicon photonic packaging will accordingly become increasingly necessary in order to match demand in an affordable and efficient manner.

Published work in the area of packaging has yet to synchronise with other work in Silicon Photonics. The literature review showed that a range of different approaches to Silicon Photonic packaging have been proposed and demonstrated without a clear victor. Despite this, passive alignment has certainly become a necessary component of any viable packaging method due to the reductions that can be made in terms of time and equipment needed during the assembly. Additionally, it was shown from the literature review that an in-plane packaging methodology is to be preferred such that the final package is robust and compact. The field of Silicon Photonic packaging still requires exploration and as such, much of the work done for this project can be considered 'frontier work', setting standards and proving concepts.

Though a final demonstration of this packaging method was not obtained during the project, a number of key milestones were reach which did provide confidence that it would be feasible, given a solution to the contamination that built up on the chips. Despite the disappointment, there are still many original and crucial results that have been obtained from this work that move this project towards its final goal. This section will outline in summary what the work here has

demonstrated and what impact this has had on the overall goal of a method of passive alignment for silicon photonics.

8.1.1 Packageable Grating Couplers

This work has both explained the need for and provided a solution of how to couple light into a silicon PIC that is also compatible with a passive means of alignment. The final method used was a customised grating coupler that was optimised both for the unusual angle of incidence, but also for the addition of an index matching epoxy and the increased beam divergence that is introduced when using a mirror. This grating coupler was used in a reverse configuration with an effective angle of -19.7° to the normal. This specific design facilitates passive alignment for the first time. It has been demonstrated that grating couplers with an insertion loss of 2.89dB are possible this way, with scope for improvement to 1.59dB based upon simulated results. This represents a comparable result to other packaging methods but also is compatible with passive alignment and crucially is the lowest reported insertion losses of this kind of grating coupler.

8.1.2 Demonstration of Unique Process Flow

This work has overcome and demonstrated a number of challenges with regards to fabrication, showing the feasibility of such chips to be fabricated within a typical CMOS fabrication facility. Both chips can be fabricated using mostly standard CMOS techniques and hence are scalable to large throughputs. The particular highlight in this regard was the total design from scratch of a novel recipe for etching through the entirety of the silicon wafer, for a non-standard layout. A starting bias voltage of 350V was used for the Bosch etch, with a ramp of 9V/min. This ensured excellent quality and depth of etch with minimal sacrificial damage to the substrate. Without this crucial step it would not be possible to have adequate vision of the target site, nor would it be possible to insert the required epoxy and it would also make fibre insertion far less effective. This work neatly shows a method to etch the required fibre channel, but this method could also be easily adjusted, with further calibration, in order to allow an evolved packaging method in the future to be possible.

A novel 3D printed jig was used to facilitate assembly. This methodology can be easily transposed onto an industrial pick and place tool. The jig was used to deposit an adequate amount of epoxy into the relevant holes on the chip surface. The same method was used to securely drop the silica microsphere into the hole.

8.1.3 Novel Chip Layout and Features

This work has also demonstrated a novel design for chip layout that has been adapted with packaging in mind. The use if silica microspheres allows the two chips to self-locate and remain in the correct location in a secure manner. Particular attention in this regard has been applied to the Capping Chip to keep the Device Chip as standard as possible and hence makes it simply to apply to almost all Silicon PICs. Key features are shown in Figure 8.1. These include:

1. A chamfered edge around chip to allow easier fibre location and for a bead of epoxy to be run around the edge of the chip for better adhesion.
2. Microspheres have been strategically located around the Capping Chip in order to allow the precise and effective bonding to the Device Chip. The tests in the previous chapter show that this method was able to align both chips to a satisfactory alignment tolerance.
3. Adhesive wells have been etched into the surface to provide a vessel large enough for the injected epoxy to fix itself to both chips and create the adhesive plugs needed. The results from the previous chapter show the effectiveness of these.

4. It was demonstrated that an Au layer could be deposited onto the mirror section at the end of the V-groove with the introduction of an intermediate metal bonding layer.

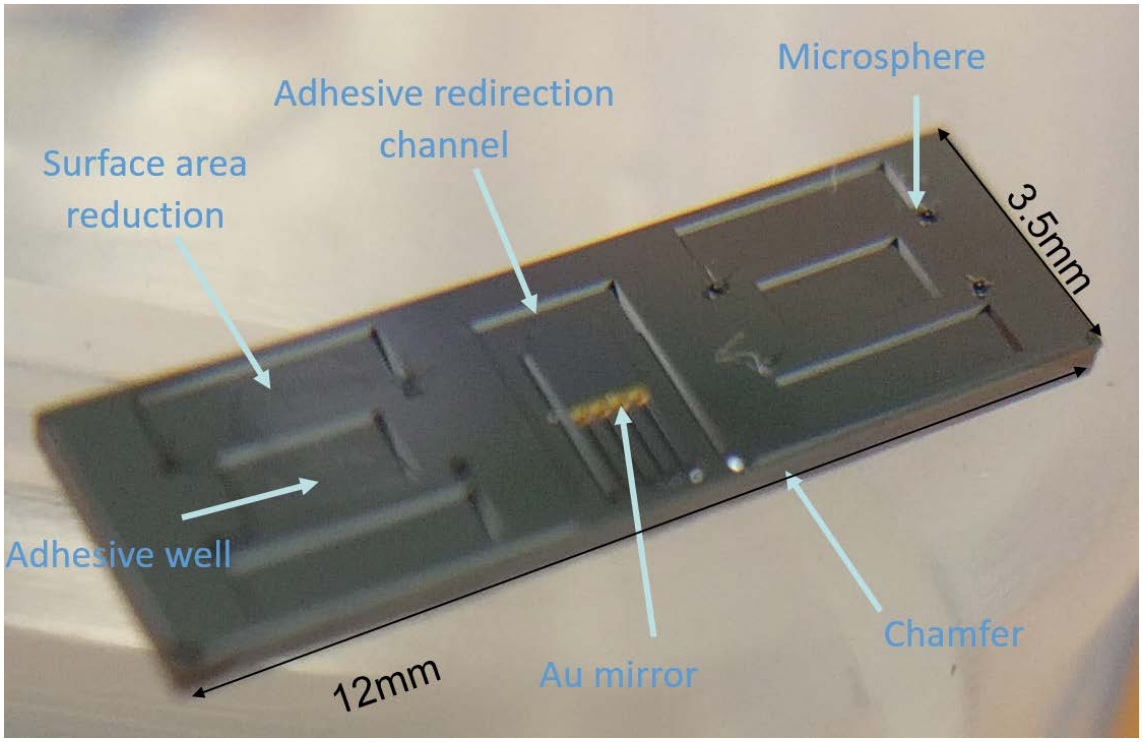

Figure 8.1 Photograph of final fabricated Capping Chip

5. A Surface Reduction Area or recess has been applied to help mitigate against stray dust and particles but also to reduce the total amount of area that needed to be in contact with the Device Chip. Since the total amount of area has been reduced, it means that there is a greater chance for a flat mating surface, and this allows for higher probability of successful bonding.

By employing these different ideas together, it was demonstrated that the two chips could be bonded to an accuracy of ±1.5μm and provide a bonding strength which was actually stronger than the tensile strength of the silicon itself, shattering at 36kg effective sheer stress. This represents a demonstrably robust yet repeatable method for chip alignment and bonding.

8.2 Conclusion

Whilst it is regrettable that working demonstrator was not realised in the timescale of this project, the work shown here has made vital and important steps towards an economic packaging method that can be scaled to industry standards.

Key results have been obtained which demonstrate the feasibility of this method, given extra time. The following section on future work outlines what improvements can be made to make this a reality.

Figure 8.2 shows a photograph of the finished and bonded samples, indicating that all the necessary fabrication processes and functions that needed to be developed have been completed. The development of the fabrication process took

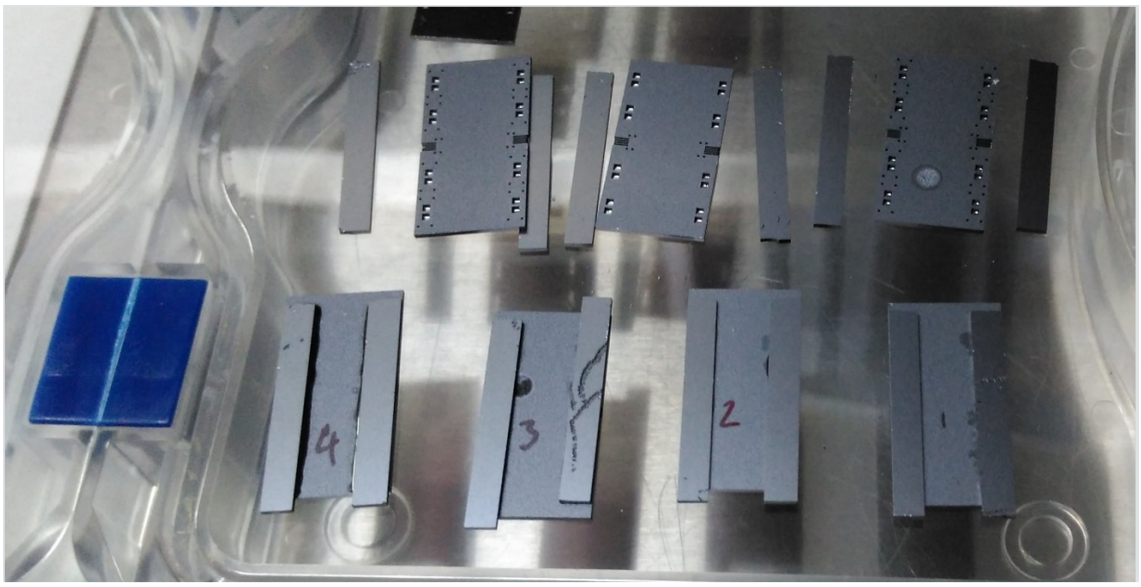

Figure 8.2 Photograph of fabricated final samples

many months, but eventually a prototype which could be packaged using passive alignment was produced- if aided by superior cleanliness. The specific structures which have been designed and fabricated have improved the passive alignment method and has overcome a number of hurdles in the initial design. A repeatable and CMOS compatible fabrication method had been established, and a repeatable assembly method has also been demonstrated. Therefore, with a little extra improvement, this method can be scaled to industry specification.

This project has also demonstrated the design and characterisation of a novel coupling mechanism that is especially sensitive to passive alignment packaging constraints. Grating coupler losses have been minimised to levels which are lower than those reported elsewhere for the unique set of constraints that were required for passive alignment. Without this innovation, it would be impossible to feasibly package silicon chips with any utility. Therefore, this is an important and positive

result from this work which forms the foundation of the subsequent work on fabrication and assembly.

Overall this project has achieved a number of key milestones along the way to a passively aligned packaging method. It provides a number of key contributions that have brought passively aligned silicon photonics closer to a working reality and the work shown here can act as a springboard to enable passive alignment in future.

8.3 Future Work

8.3.1 Reduction to contamination

In order to allow assembly to be undertaken, the contamination must be reduced. This section will outline a few ways that the processes in this project could be adapted to reduce the risk of contamination.

Firstly, the layout of the device chip could be altered such that the holes are only present on the edges of the wafer. Conventionally a 4-inch chuck is used with a 6-inch wafer. This means that the inner 4-inch diameter of the wafer cannot contain holes. The remaining outside edge could be used but this would significantly reduce wafer footprint. A better compromise would be to use a smaller chuck, either 2 or 3 inches, and then reserve only a small area of the chip to be kept free of holes. This would still negatively affect the yield of the chip but would allow resist coating and hence make the dicing procedure a much cleaner one. Alternatively, an edge mounted chuck could be used. This uses clips at the edge of the wafer to hold it in place whilst it spins. Only much lower spin speeds are possible however and so this would need to be tested in order to ascertain if a suitable thickness could be achieved.

The second way to reduce contamination would be to use cleaner tanks for wet etching. For both the KOH and hot phosphoric acid etching, the wafers were etched in a bath which can contain debris from other processes. As wafers are etched, contamination can break off the surface of the wafer. If these contaminants do not strongly react with the etchant, then they can contaminate the wafer yet again. The more times the same tank is used, the greater the likelihood is of contamination. Therefore, an improvement could be to use fresh etchant each time.

Though potentially more expensive, it could be tested to see what effect his has on the cleanliness of the wafers after wet etching.

Another way to reduce the contamination would to keep the wafers in the cleanroom during the initial assembly phase. Unfortunately, this was not possible for this project since the silica spheres required numerous safety assessments before being allowed in a cleanroom environment. This is very understandable since if they were mishandled then the microspheres would end up scattered throughout a clean environment and could cause widescale contamination themselves. A compromise was sought during this project with the cleanroom staff but could not ultimately be achieved. Given more time, it could be possible to conduct this phase in a clean room environment. Then, when the microspheres had all been located, the wafer could be sealed in cleanroom conditions, shipped for assembly and snow cleaned before dicing and final assembly.

The final way that would significantly improve the levels of contamination would be to do wafer scale assembly. This would follow the previous method, but the dicing would take place after the two wafers were bonded. This would of course require the two wafer layouts to correspond to each other making overall yield decrease since the Capping Chip cannot be so densely filled with chips due to its structural integrity. However, it should significantly reduce the contamination present in the fibre channels since the two chips would both protect each other during the dicing process. It could even be possible to do wafer scale testing of the packaged chips before dicing, since a temporary fibre could be inserted into the fibre channels to test each chip. Theoretically this could be done with passive alignment. Then the wafers would be diced after being bonded together and the final fibres could be inserted and glued in.

8.3.2 Evolution of the Silica Sphere Alignment Method

The issue of which alignment structures to use is still something of an open question. The use of microspheres has been demonstrated here but this is not a perfect solution. The tolerance of sphere radii is difficult to improve upon. A way of separating different sized spheres would be a bonus. A primitive means of doing so has been trialled with Dr Reynolds' assistance. Figure 8.3 shows how this was done. The glass tube was filled with a viscous liquid and the spheres were then inserted in the top. The

Figure 8.3 Image showing a method to separate different silica spheres

differences in size will result in differing buoyancies and thus they will sink at different rates. The tap at the bottom can be opened and closed at intervals such that the spheres collected will be closer in size. However, this proved difficult as removing the liquid after the process cleanly was an issue and this method only increased the amount of grease and particles present on the spheres making the problem worse. It is an interesting concept however and could be improved. The fact remains that a means to separate the differing sized spheres would be very beneficial and would improve the alignment tolerances.

8.3.3 Packaging of non-standard fibres

This packaging method could be extended to be used for non-conventional fibres (lensed, holey fibres and multi-core/mode fibres for instance). This would be an attractive area for future research. The state of the art in fibre research is also rapidly moving to keep pace with data demands. Multi-core fibres such as the one depicted in Figure 8.4[202] are being extensively worked on. An interesting approach would be to use an array of grating couplers to correspond to the multiple cores on such a fibre to allow packaging into a PIC.

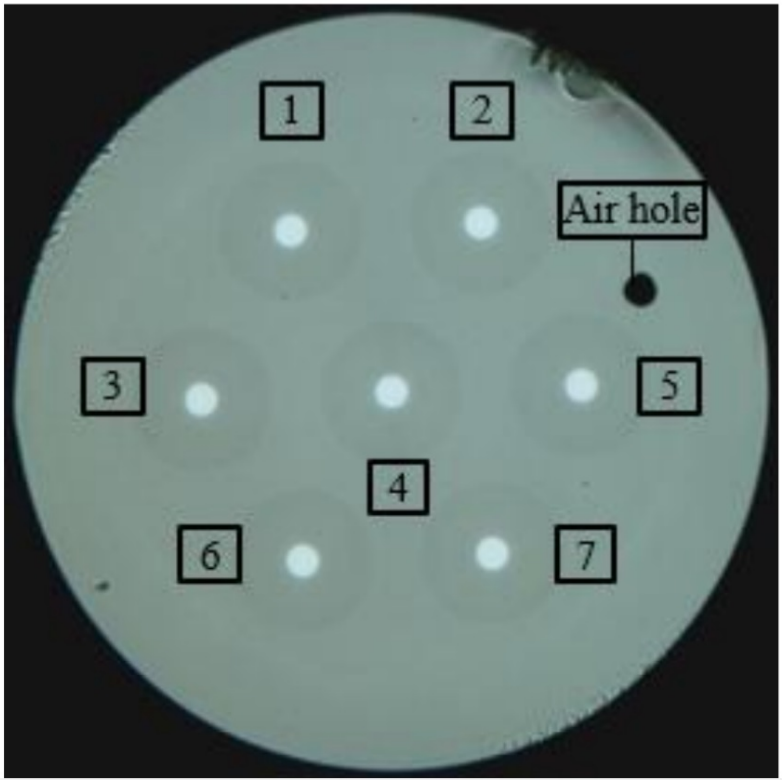

Figure 8.4 Multicore fibre by Donko et al. [202]

8.3.4 Demonstration of Lasers and Inverted Tapers

It has been suggested that this packaging system could be seamlessly integrated alongside other coupling methods, namely lasers and inverted tapers. However, more work is still required to demonstrate this in practise. An ideal demonstrator would be packaging a laser and a fibre onto the same chip, using both coupling methods side-by-side which would be easily possible given the flexibility that has been designed into this method.

Appendices

A.1 Clean Room Equipment

a. The Cleanroom Itself

The cleanroom itself is perhaps not typically considered a tool in its own right but probably should be considered as such given the important role it plays in nano-fabrication. In order to be classed as “cleanroom” a space must meet the criteria as made clear in documentation[203] by the International Organization for Standardization. The metrics used to measure cleanliness are particle size and number in a given volume of space.

The biggest cause of contamination when fabricating nanoscale devices is the fabricator themselves[204]. In most controlled laboratory working spaces, the equipment poses a risk to the person conducting the work. However, in nanofabrication the opposite is also true. The worker must protect their work from contamination, not the other way around as may be appropriate in a chemistry lab for example. A stationary person can generate up to 100,000 particles ≥0.5μm every minute. That number can increase tenfold to 1,000,000 particles every minute when moving and working[205]. This poses an obvious problem when fabricating chips with feature sizes of comparable dimensions.

In order to prevent any contamination, cleanrooms filter any incoming air to remove any dust or other foreign contaminants, as well as maintaining positive pressure throughout causing air to flow away from the cleanroom. Air locks are used to prevent dirty air seeping in. Strict regulations on what is permissible inside the cleanroom also prevent contamination.

Cleanrooms are usually divided into different classes by their different cleanliness. A class 100 cleanroom refers to the number of particles sized 0.5mm of larger per cubic foot of air. A class 1000 clean room will contain up to 1000 such particles in each cubic foot of air.

In this project classes 100, 1000 and 10k where used for processing, with the process beginning in the higher class and then moving to the lower classes when appropriate, (i.e. 100 to 1000 to 10k).

All cleanroom processing for this project was carried out at the Southampton Nanofabrication Centre unless explicitly mentioned and hereafter will be referred to as the SNC

b. Lithography

Lithography is the process by which a wafer can be '*patterned*'. Literally speaking, lithography means to write onto a stone, but within nanofabrication, substrates are often referred to as a stone, such as 'monolithic' chips where all processing and transmission can be achieved on only one chip. Structures such as waveguides and couplers must be written onto a wafer before they can be etched into the surface.

When '*writing'* a wafer a suitable *'ink'* must be used. The ink that is usually used is called a resist, or photoresist. This is a chemical polymer material which is deposited across the entire wafer to a specified thickness (200nm up to 10μm in extreme cases). This is the ink, which is then written into. Portions of the resist will be removed by lithography and then development which will leave exposed regions of the wafer which can then be etched away. The resist will usually either be unselective to etching or thicker than the etch depth in order to properly protect the substrate which is being etched. Once the etching has been completed then the resist can be removed leaving the pattern etched onto the substrate surface.

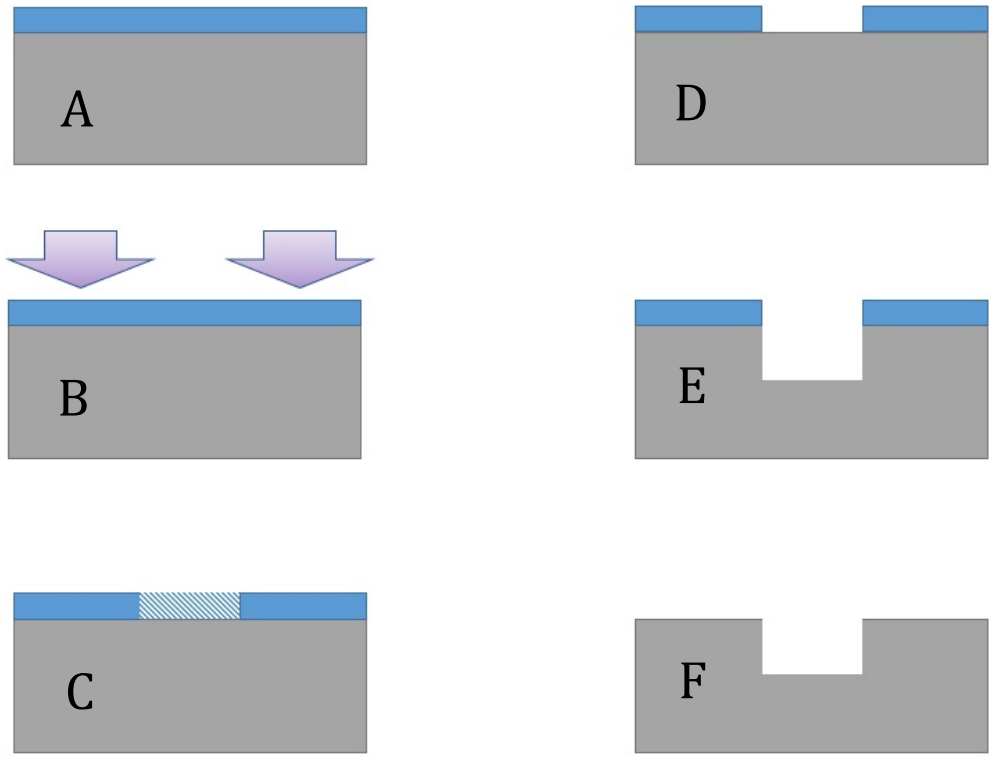

Figure A.1 Diagram showing steps for lithography

Figure A.1 shows the steps in lithography. Figure A.1A shows a wafer with a resist deposited on the top of it. B shows a pattern being written onto the resist. This then leaves a written region on the wafer, which is signified by the cross-hatched part of C. D shows development where the written section of the wafer is removed by developing the resist through a different chemical process. Strictly speaking after step D, lithography has been completed but step E and F, etching and resist removal, are shown for completeness. The process flow shown in Figure A.1 can be repeated as many times as necessary to create all the necessary structures on a substrate.

There are two main methods of '*writing*': photolithography and e-beam lithography. Both have been used in this project in order to pattern wafers appropriately. Each method will be discussed in turn with advantages and disadvantages discussed and giving the rationale for using each type.

1. Photolithography

The first main branch of lithography is photolithography. As the name suggests this method, uses light of a specific wavelength to write onto the surface of the wafer. In the case of photolithography, the resist used is a called a photoresist since it will be reactive under exposure to light. Typically, UV light is used for this procedure for two main reason. Firstly, it means that photolithography can be done in the presence of lower energy, higher wavelength light that will not be reactive as UV light will. Therefore, it doesn't have to be done in low-light conditions. Secondly, using UV light means the optical wavelength is smaller allowing smaller structures to be written to the wafer. Usually this is called '*deep UV*', using wavelengths from 200-400nm, meaning that minimum feature size is roughly the same. Equation A.1 shows the relationship between feature size and optical wavelength in photolithography[206].

$$\boldsymbol{w = k\frac{\lambda}{NA}} \qquad \textbf{Eq. A.1}$$

Here w is the width of the smallest feature, k is a constant (also called k-factor) which covers different factors associated with the process. λ and NA are the wavelength and numerical aperture respectively. From this it can be seen that the using a smaller wavelength is a simple way of decreasing feature size. In practice, a resolution of as small as 1μm down to a few tens of nm is possible using photolithography.

Since UV light is used to activate the patterning on the photoresist, the whole room where this happens must be free from any lower wavelengths of visible light which could activate lithography. Therefore, photolithography is usually carried out in a '*yellow-room'* where light filters are used such that the ambient light is mainly yellow in colour; long enough wavelengths such that they won't interfere with the lithography process. Figure A.2 shows what such a yellow-room looks like, this

Figure A.2 Picture of the yellow room

particular one being located in the SNC and was used for the majority of lithography in this project.

The light source for UV lithography is usually a mercury lamp, which emits light at a number of specific wavelengths in the UV part of the spectrum. Filters can be applied if a specific region of the UV spectrum is required.

In order for a specific pattern to be etched onto the wafer surface a '*mask*' must be used. Whilst it would be possible to have a very fine UV beam and move this beam as appropriate, a mask can be used as a stencil and the UV light will '*flood*' the entire mask area. The mask will then shield certain parts of the wafer from receiving UV light.

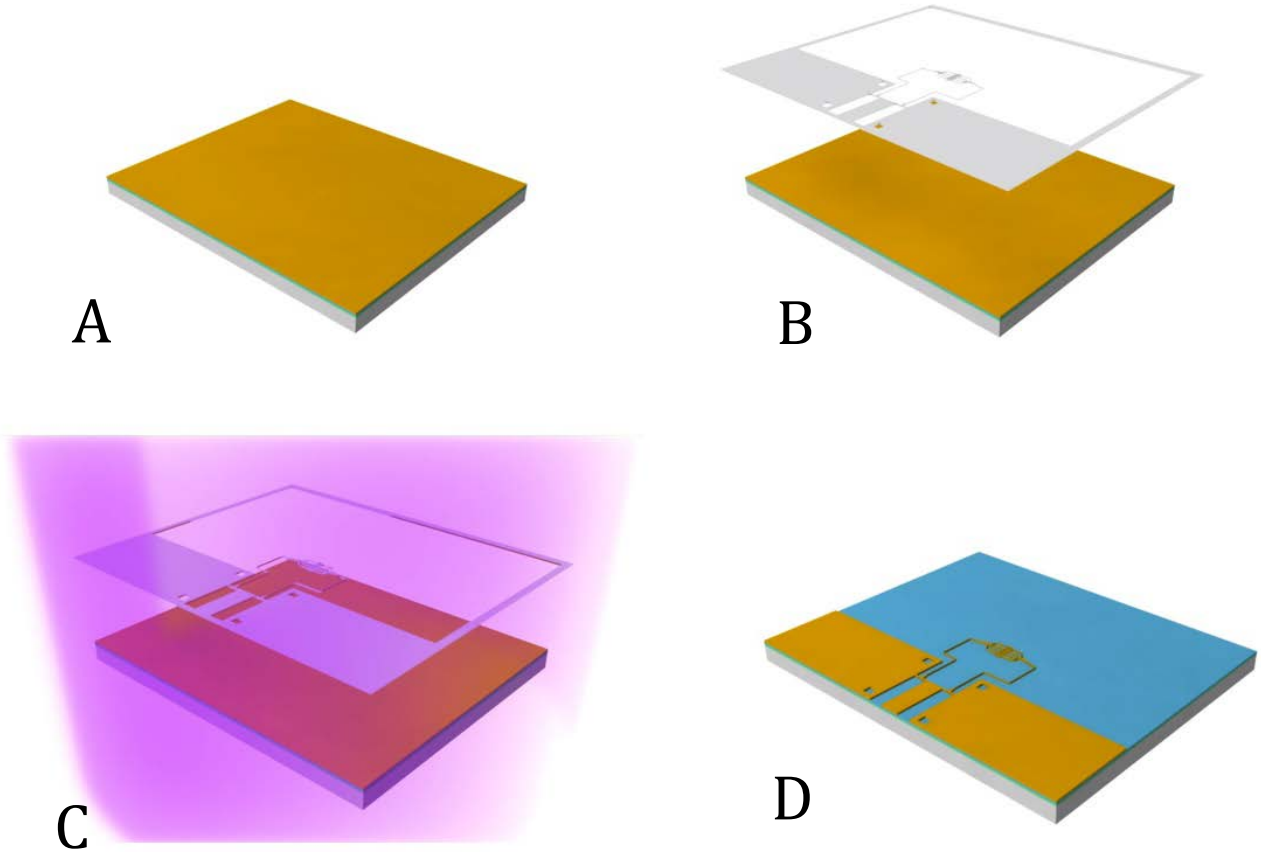

Figure A.3 Diagram showing steps of Photolithography

Figure A.3 shows the same process as Figure A.1, yet this is from a bird's eye view rather than an adjacent view. From this angle it is clearer how a mask works. Figure A.3A shows the substrate with resist applied. Figure A.3B shows a mask applied over the top of the substrate. This is usually made from chrome on glass, chrome being opaque to the light and glass transparent. Figure A.3C illustrates the mask and substrate under UV exposure and Figure A.3D the finished substrate after development.

Both positive and negative photoresists were used for this project. A positive photoresist means that the area of UV exposure will be removed after development. Conversely, with a negative photoresist, the exposed area will then remain after

development and the unexposed area will be removed by the development chemicals. Figure A.3 shows a positive resist being used.

When a positive resist is used then the UV light will cause polymer scission, where the polymer chains are broken up. This then allows the developing chemical to dissolve away the regions of the resist that were exposed. Conversely, if it is a negative resist then the UV light will cause cross-linking which binds together long

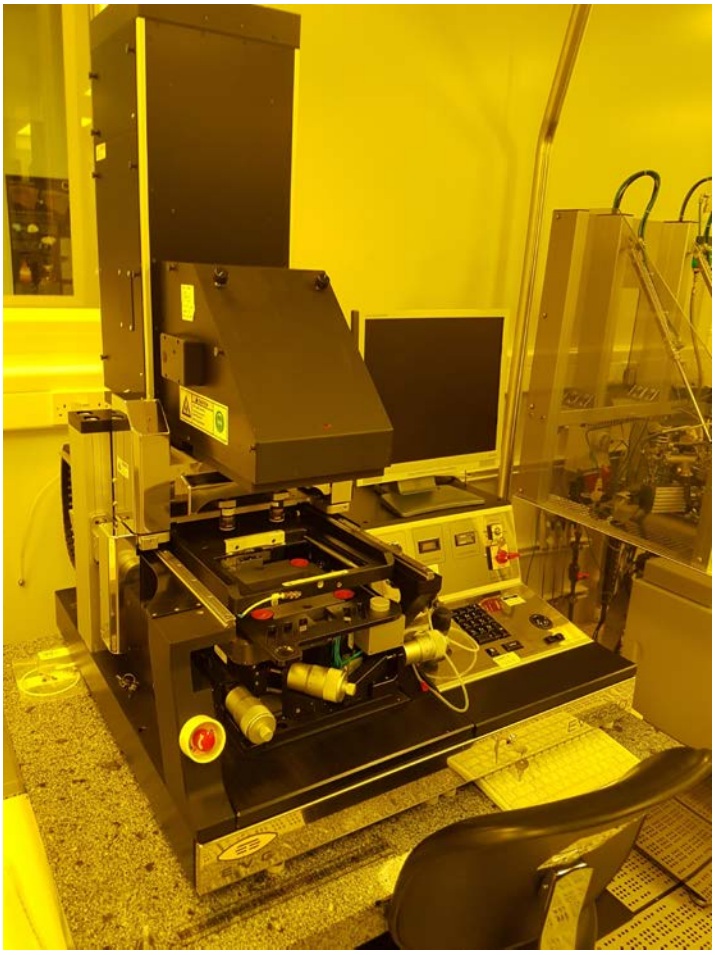

Figure A.4 EVG620 mask aligner used in the SNC

polymer chains. This is often further activated thermally through a post exposure bake. The developing chemical will then dissolve away the unexposed regions of resist which were not cross-linked.

For precision alignment of the mask to the substrate a mask aligner must be used. For this project an EVG620 was used for photolithography. Figure A.4 shows an image of the mask aligner used for this project. This model allows for vacuum contact exposure which uses a vacuum to secure the mask in place which his

brought into contact or close proximity to the wafer substrate. Another useful feature of this model is backside alignment. This allows a mask on the front side of the wafer to be aligned to a mask on the backside meaning that both sides of the wafer can be processed. The EVG620 uses precision microscope cameras and a micron-scale stage to allow front-side alignment of ±1.0μm 3σ and backside alignment ±1.25μm 3σ[207]. One limiting factor of this model is the maximum size of 150mm wafers, whereas usually industry processes use 200mm wafers.

2. E-beam Lithography

E-beam lithography (EBL or electron-beam lithography) is an alternative to using photolithography. As the name suggests UV photons are replaced with an electron beam. Because of this the fundamental physical limitation of diffraction effects is mitigated. E-beam spot sizes can be reduced to as small as 5nm[185]. The de Broglie wavelength for a typical 25keV electron will be only 8pm. The limitation with e-beam lithography is more to do with the equipment available. Typically, e-beams of 10-100keV are used with very small current of the order a few nA.

With e-beam lithography the resists must be sensitive to electrons rather than photons, though often resists are manufactured for either type of lithography.

E-beam lithography can be used in flood exposure like with photolithography, however more often a direct writing procedure is used due to the finer line-widths that are possible.

An immediate difference between electrons and photons is the discrepancy in mass. This means that mass related scattering effects can happen as the electrons meet the resist. Figure A.5, adapted from Franssila[185], demonstrates these effects. When the electron acceleration voltage is low (10kV) and the resist thickness is large then there can be significant scattering as in Figure A.5A. This is reduced with decreased resist thickness as in Figure A.5B. As an alternative the electron acceleration voltage can be increased as in Figure A.5C to reduce scattering. The situation of decreased resist thickness and increased electron acceleration is shown in Figure A.5D for completeness.

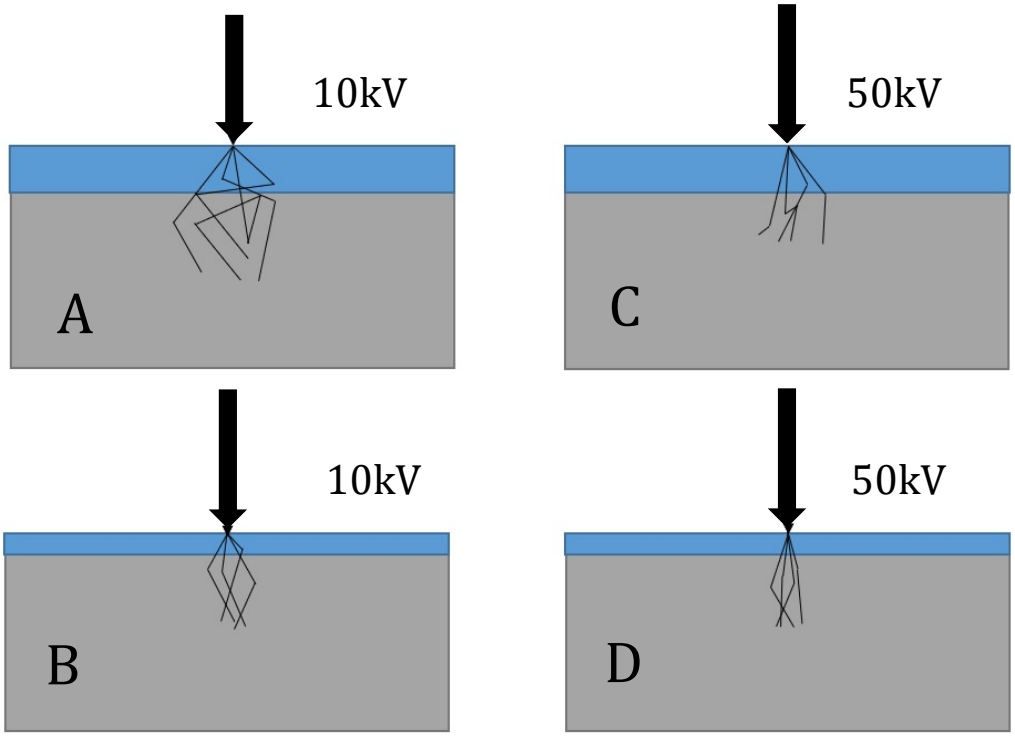

Figure A.5 Comparison of different e-beam acceleration voltages and resist thicknesses

When using direct-write, consideration must be made for the time this takes. When writing across an entire wafer this must be done in a raster scan. Therefore, if the write density is large and the wafer size is large (up to 200mm) then the entire job can take many hours.

Usually an e-beam will use '*stitching*' to increase area size. In this method the e-beam will only move a very small distance in the x and y directions and when a

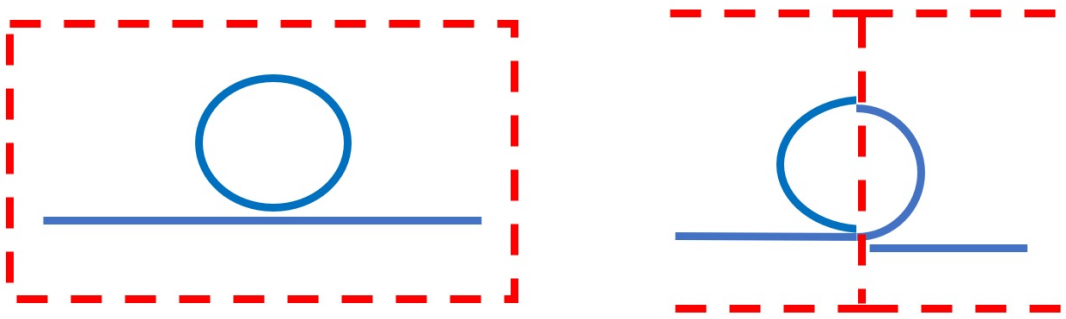

Figure A.6 Stitching error from e-beam lithography

section is completed, the wafer will move on a stage such that the next cell can be written. This can lead to stitching-errors where the adjacent cells meet together. To mitigate the effects of this, any critical devices should not be written over two cells. For example, grating couplers which require a very precise period should be kept away from the stiches where possible. The stitching errors are usually only a maximum of 10nm and so less critical components such as metrology structures can be placed in the stitching gaps. Figure A.6 shows how stitching errors can affect features that are written. This is an extreme case which is used to highlight the problem, normally it wouldn't be anything nearly this bad.

No physical mask is used for e-beam lithography however. Instead, a software mask is used. In this case it is still referred to as a mask but no physical mask is ever made. Instead, the e-beam is programmed using the mask structure

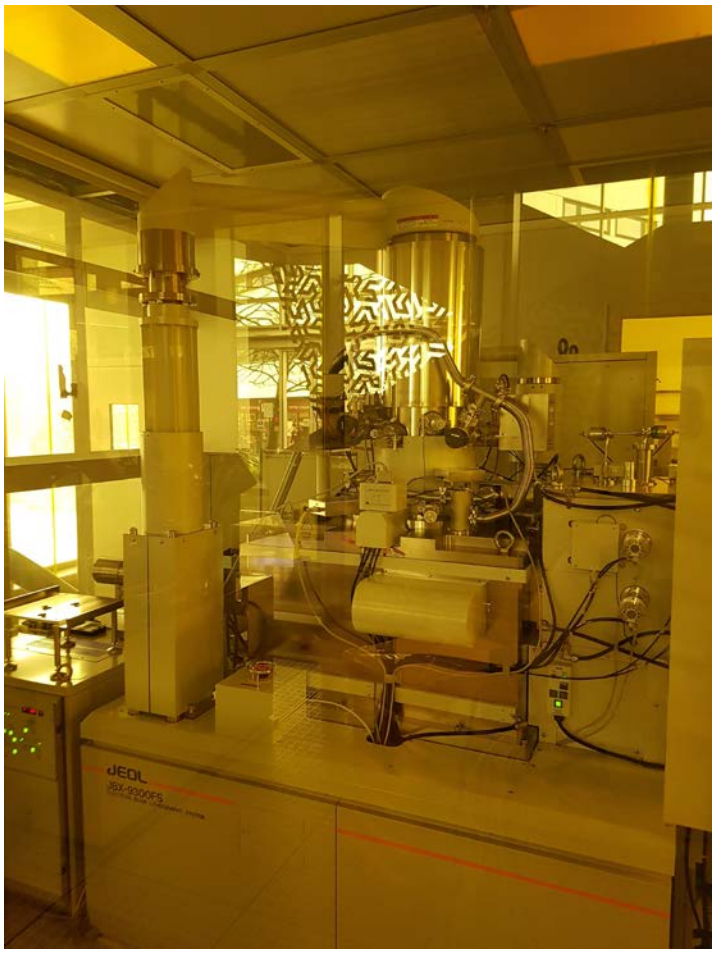

Figure A.7 JEOL JBX 9300FS e-beam lithography system used in the SNC

and it is written directly onto the wafer. For this project a JEOL JBX-9300FS was used to write wafers with the help of Dr Ali Khokar. This model can write down to widths of 10s of nm using electron energies of 50-100keV and a 4nm spot size.

3. Lithography Summary

Both forms of lithography that have been discussed have been used in this project. After explaining their mechanisms, a comparison can be made between the two different technologies.

Table A.1 shows a comparison of both methods. From this side-by-side comparison it can be seen that both have their respective niches. For prototyping in a research context e-beam lithography has several specific advantages. The flexibility it offers allows mask designs to be modified with ease. Single wafers can be written simply and the higher resolution on offer means that there are less restraints on mask

design, providing a quicker turn around for research. Indeed, if a single wafer is to be fabricated then a chrome mask must be made using an e-beam so the time taken for fabrication will always be slower. However, the cost and throughput (speed of multiple wafers) makes e-beam lithography simply unsustainable for any commercial or production use.

Table A.1 Table of Lithography methods by comparison

	UV Photolithography	E-beam Lithography
Resolution	Lower (50nm max but typically much lower ~1μm)	High (up to 15nm)[208]
Resolution Potential	Diffraction limited to 50nm	Not diffraction limited (technological limit rather than physical)3
Write complexity	No time penalty for higher density writes	Write time increases dramatically with higher density features
Single Wafer Speed	Slow- mask must be made using e-beam	Fast – direct write
Multiple Wafer Speed	Fast – dozens of wafers per hour	Very Slow – several hours per wafer
Cost	Cheap – speed and equipment reduce costs	Very Expensive – resists alone can cost many thousands £

c. Wafer Preparation and Resist Removal

This section will outline how a wafer is prepared for either form of lithography and then how the resist can be removed after etching has been completed. Whilst strictly speaking these parts aren't part of the lithography process they have been included here for lack of better alternative.

Before a resist is applied to a wafer it is dehydrated. The presence of moisture on the wafer can have an adverse effect on the resist thickness once it is applied and its adhesion. In order to dehydrate a wafer to prepare it for resist application it should be placed in an oven to evaporate any moisture that may be present. For this project a Thermo Fisher Scientific Heraeus drying oven was used to dehydrate all wafers prior to resist application, as depicted in Figure A.8

Figure A.8 Picture of a Thermo Fisher Scientific Heraeus drying oven in the SNC

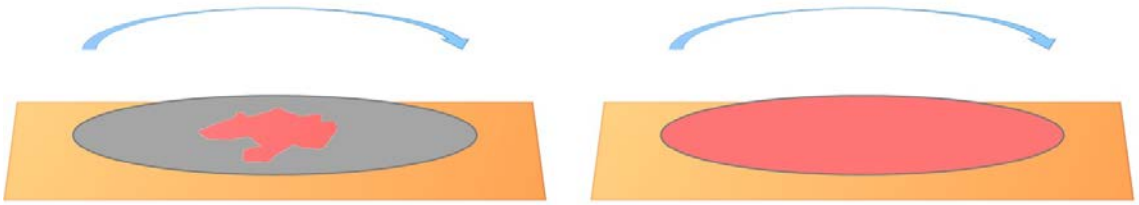

Figure A.9 Diagram showing a resist spreading to cover a wafer through spinning.

Figure A.9 shows how a photoresist or e-beam resist, is usually '*spun on'* using a resist spinner. This is a device that can spin the wafer or substrate at several thousand rpm. A vacuum is used to hold the sample in place to prevent it flying off, being damaged and most likely destroyed. Once the sample is secure then a resist can be applied to the surface in liquid form. This is either done through a pipette or poured over the surface from a bottle. Figure A.9 shows how this happens. Once the resist has spread across the entire sample the spin speed can either increase, remain constant or stop depending on what thickness is required. Thickness will be a function of spin-speed and resist viscosity.

For this project a Brewer Science CEE-200 was used to apply resists to wafers. This model features the ability to program specific spin recipes so that a

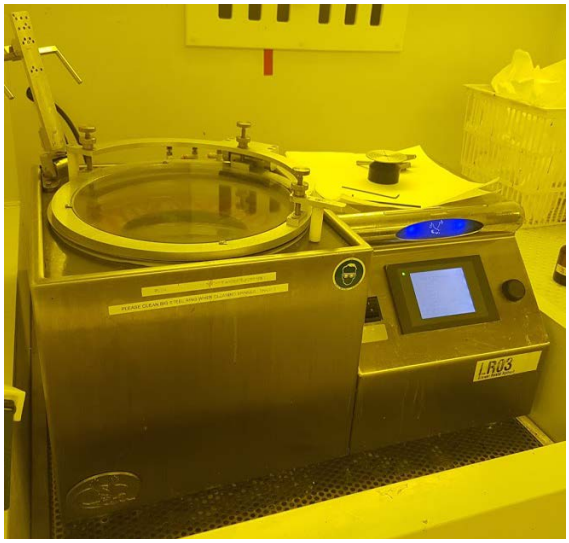

Figure A.10 Brewer Science CEE-200 resist spinner

particular resist can be applied repeatability many times. It is able to reach spin

speeds up to 6000±1 rpm allowing very thin resists to be deposited. Figure A.10 shows this particular spinner located in the SNC yellow room.

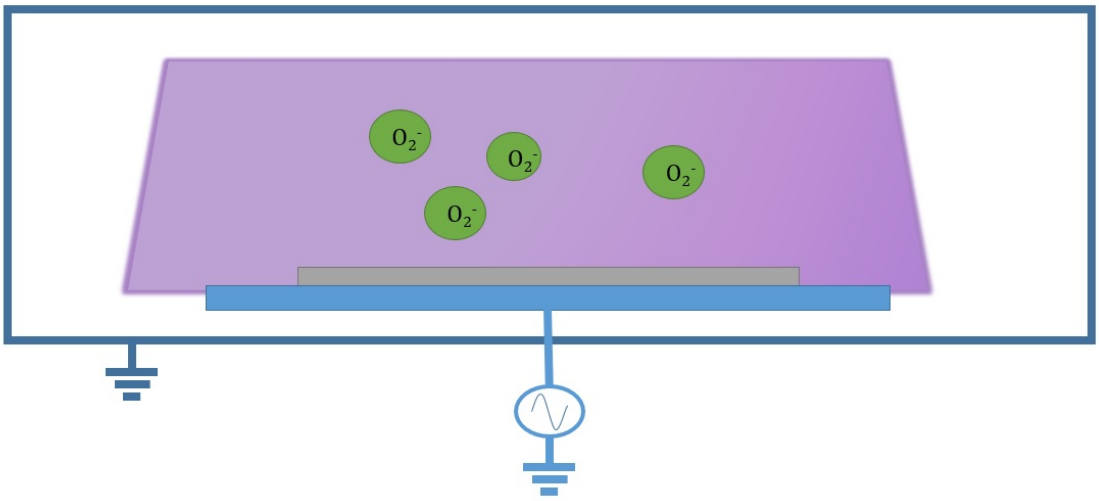

Figure A. 11 Diagram of Plasma Ashing using O_2

Once a resist has served its purpose it must be removed so as not to change the optical properties of any devices created or to allow a new resist to be applied or other think film. In order to do this a plasma asher can be used. A plasma asher uses a plasma of free radicals (usually O_2) to remove the organic resist polymer

Figure A.12 Image of PVA Tepla 300 Plasma Asher in the SNC

compounds on the wafer surface. This way resists can be removed dry and with no residue. Figure A.11 shows a diagram of how this is achieved. All ashing done for this project was done using a PVA Tepla 300 Plasma System located in the SNC as shown in Figure A.12.

d. Dry Etching

1. Etching Overview

In order to create waveguides and the other associated structures needed for this project the wafer substrate must be etched in order to remove unwanted material. However, there are many different etching methods and so this section will outline what all these different methods are and how they can be used together to create complex structures and systems. Etching can broadly fit into either of two categories: dry and wet. Dry etching uses etch species in a plasma, which will both chemically and physically remove unwanted parts of the substrate. Wet etching uses chemicals in a liquid phase, which will etch the substrate by means of a chemical reaction.

2. Reactive Ion Etching

This section will discuss the mechanism of Reactive Ion Etching (RIE) also simply called plasma etching. Despite being called Reactive Ion Etching , this is sometimes considered an improper name since the majority of etching is due to excited neutral particles[185]. RIE is done at very low pressures of 1-150 mTorr[209], with specific reactive gasses injected which are then excited by an RF field within the etching chamber. Figure A.13 is a diagram of an RIE chamber. The RF bias is applied to a platen that the wafer sits on. The gases are injected at the top electrode. Within the chamber there is a mixture of excited molecules such as CF_4^* and ions such as CF_3^+. These are then accelerated towards the platen due to the RF bias which is applied. These two different etch species will then impart energy onto the wafer substrate. If enough energy is delivered, then the Si bonds can be broken, and material is

etched away. There is a double effect of both "chemical (reactive) and physical(bombardment)" etching[185].

Typical gases that are used for RIE are halogen compounds such as fluorides (CF_4, SF_6, CHF_3) or chlorides($SiCl_4$). Other gases are used as stabilizers (He or Ar) and scavengers (O_2) in order for the etching process to proceed smoothly. He is used to cool the wafer by applying the gas to the back of the wafer. Therefore, to allow adequate cooling the wafer's integrity must be maintained such that excessive helium doesn't leak into the chamber which could disrupt the etching process.

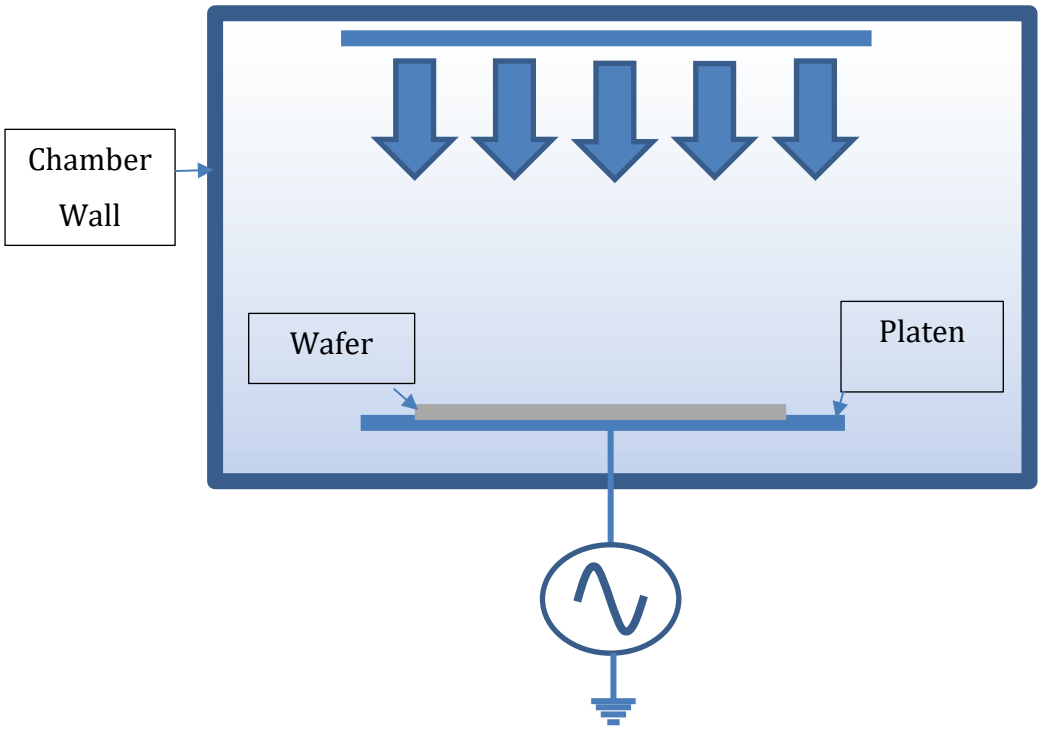

Figure A.13 Diagram of RIE

Many factors will influence etching such as RF power, temperature, pressure and the abundance of each gas. By controlling each of these parameters and creating a 'recipe', repeatable etches can be made which are then used to fabricate whatever structures are necessary. Other factors that can affect etching are the masking material, the total etched area on the substrate surface and the proximity to the centre of the platen. Therefore, these must all be taken into consideration when using RIE as an etch mechanism and the recipe must be fine-tuned to provide optimal etching results.

3. Inductively Coupled Plasma RIE

Inductively Coupled Plasma RIE (ICPRIE or simply ICP) is a very similar tool to standard RIE that was explained previously. There is however the addition of a large inductor to the chamber. Figure A.14 shows a diagram of an ICP etcher. At first glance, it seems very similar to the previous figure of

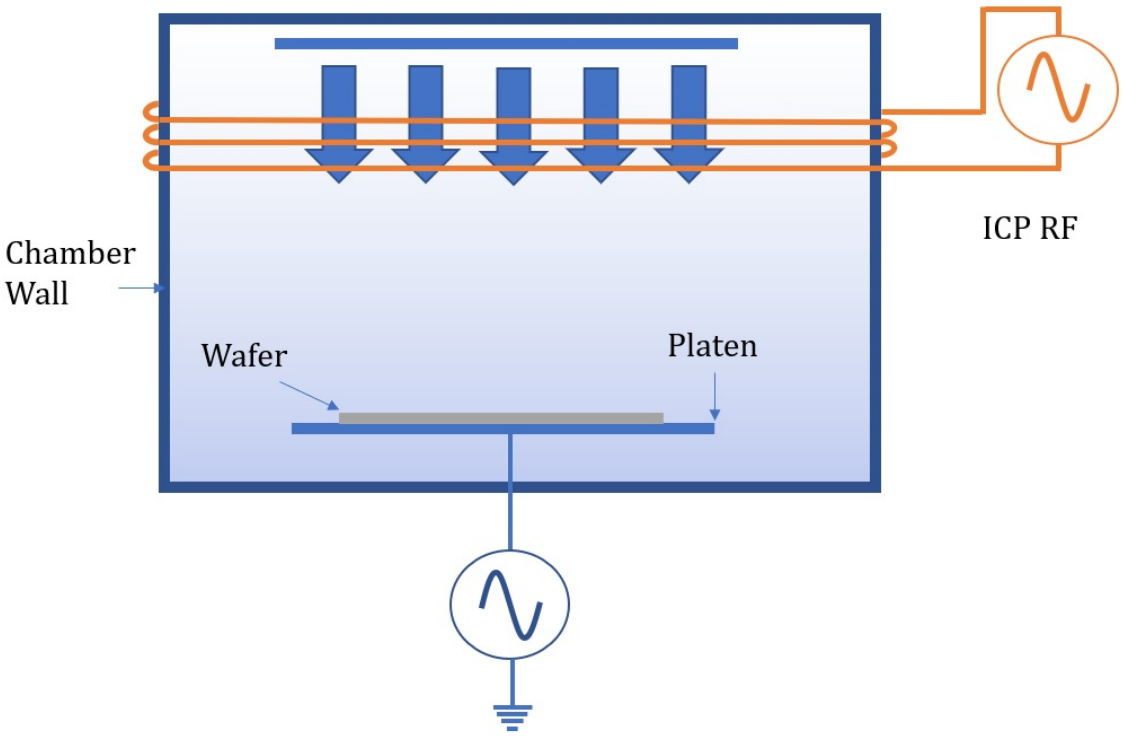

Figure A.14 Diagram of an ICP RIE

the RIE. There is the addition however of an additional RF source which is put through an inductor which surrounds the etch chamber. The inductor which is coiled around the chamber can act similarly to the focussing coils in an old-fashioned cathode ray tube television. 13.56MHz is the standard frequency that is used for this. Using inductive coupling the density of the ionic etch species in the plasma can be significantly increased. Etch rates can then be increased with a high density plasma, greater than $10^{11}/cm^3$[192]. This is in contrast to a typical RIE which will normally only achieve densities of the order of $10^9/cm^3$. By reducing the pressure of the chamber (of the order of mTorr), the etch species experience less hindrance and so directionality is improved. Under normal circumstances, lowering the pressure would obviously have an adverse effect on the density of the plasma

but the introduction of inductive coupling allows plasma density to remain high. There are now two separate RF frequencies that are usually referred to as the RF (platen) power and the ICP power. The ICP power is used create and maintain the low-pressure high-density plasma and the RF power is used to accelerate ions to energies sufficient for bombardment such that the substrate can be etched.

For this project an Oxford Instruments *Plasmalab System 100* was used, as depicted in Figure A.15. This particular model is capable of both RIE and ICP etching. It can maintain plasma densities of up to 5 x $10^{11}/cm^3$ at pressures as low as 1mTorr[209] which provides perfect conditions for high precision etching.

Figure A. 15 Picture of an Oxford Instruments Plasmalab System 100 within the SNC

4. Deep Reactive Ion Etching

Deep reactive ion etching (DRIE) is an extension to RIE that allows it to be used to etch to greater depths. Fundamentally, the process is very similar to RIE. It

is a dry etch which uses a plasma to etch a substrate. There are however some important differences. Typically RIE rates are of the order of 100-1000nm/min, whereas DRIE rates are much higher, circa 2-20μm/min [185]. This opens up a great deal more capability than standard RIE, with etch depths of hundreds of microns being possible, which have been used to great success particularly in the study and research of MEMS [186]–[188]. Sidewalls of etched trenches can be almost completely vertical which allows for etched holes to be extremely deep but also extremely narrow, with ratios of up to 100:1 possible, whereas conventional RIE struggles to maintain comparable aspect ratios. The selectivity to the mask also becomes problematic when etching for longer periods during deeper etches. Both hard and soft masks are very difficult to make thicker than a few microns and so this then imposes a boundary on etching since the entire mask will be etched away over longer etches unless an extremely high selectivity to Si can be achieved[189] . Etching is also independent of wafer orientation with respect to its crystal structure which is not the case in wet chemical etches. Any shape that can be patterned onto a wafer can be etched into the substrate almost as much as is desired. DRIE is usually used in conjunction with Si substrates as it was in this project so for this discussion DRIE will centre on Si etching.

There are two main methods of DRIE: Bosch and cryogenic. For this particular project the Bosch etch (also Bosch advanced silicon etch, ASE) was used. The Bosch etched is named after the Robert Bosch GmbH company which patented it[190]. In this process uses a combination of SF_6 and C_4F_8 to etch an Si substrate[191]. However, unlike RIE etching there is a series of individual steps that are used in order to produce the desired etch. High aspect ratios of 20-30:1 are possible as well as fast etch rates; upwards of 3μm/min are standard for the Bosch process[190]. The Bosch process consists of a cycle of different etch and passivation stages. During passivation a thin layer (~10nm) of fluorocarbon polymer is deposited across the wafer[192]. Usually C_4F_8 is used for this purpose. Then during the etch step this is removed from the bottom of trench which is etched by the addition of SF_6. The cycle then is repeated as many times as necessary in order to obtain the etch depth required. Each step of the cycle will only run for a few seconds so hundreds of cycles are sometimes required for deeper etches.

Figure A.16 shows a diagram of the Bosch process. The steps of the process are as follows; Step 1: this is the wafer with a mask and an exposed area which is to be etched. Step 2: Ions bombard the exposed area and create a trench that is a few microns deep. Step 3: The polymer is deposited onto the surface of the wafer to passivate it. Step 4: The directional ion bombardment etch is used to etch further

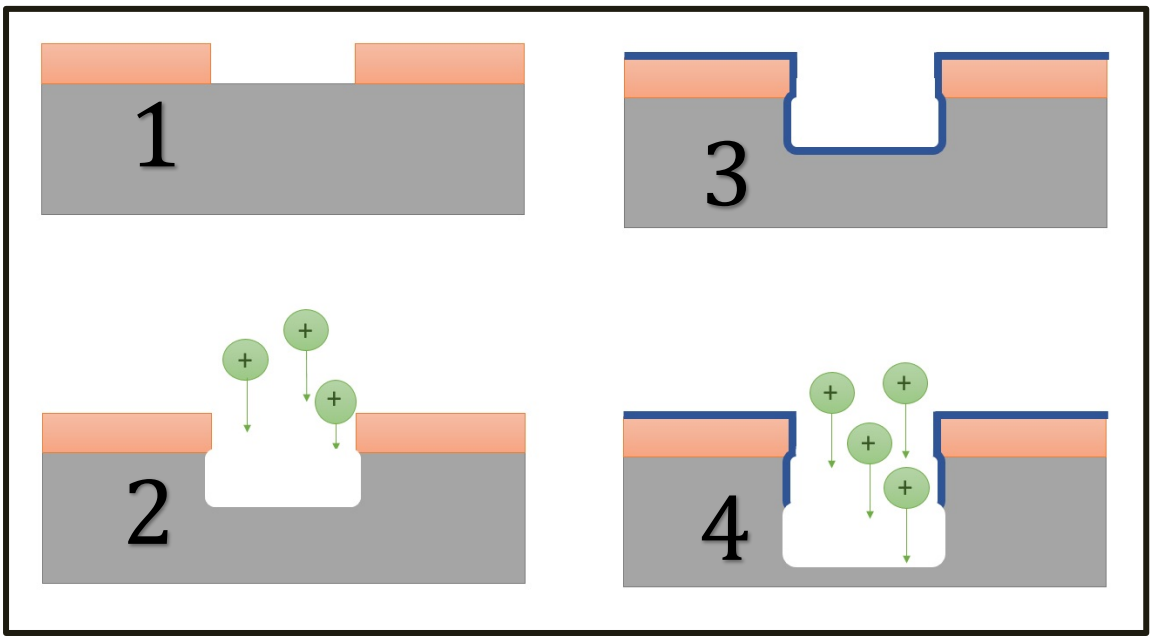

Figure A.16 Diagram showing the steps of the Bosch Process

down into the substrate. As can be seen the Bosch etch will usually result in a slight undercut due to slight isotropy when etching. The diagram is of course exaggerated hugely in order to highlight the problem. In reality, the effect is small yet noticeable. Figure A.17 shows a close-up scanning electron microscope image of a trench that has been etched using the Bosch Process. The wavy sidewall profile is the result of this undercut. This is often called scalloping because the shape resembles the silhouette of a scallop shell. As can be seen from the image with less zoom, the affect is very small when considered on a larger scale.

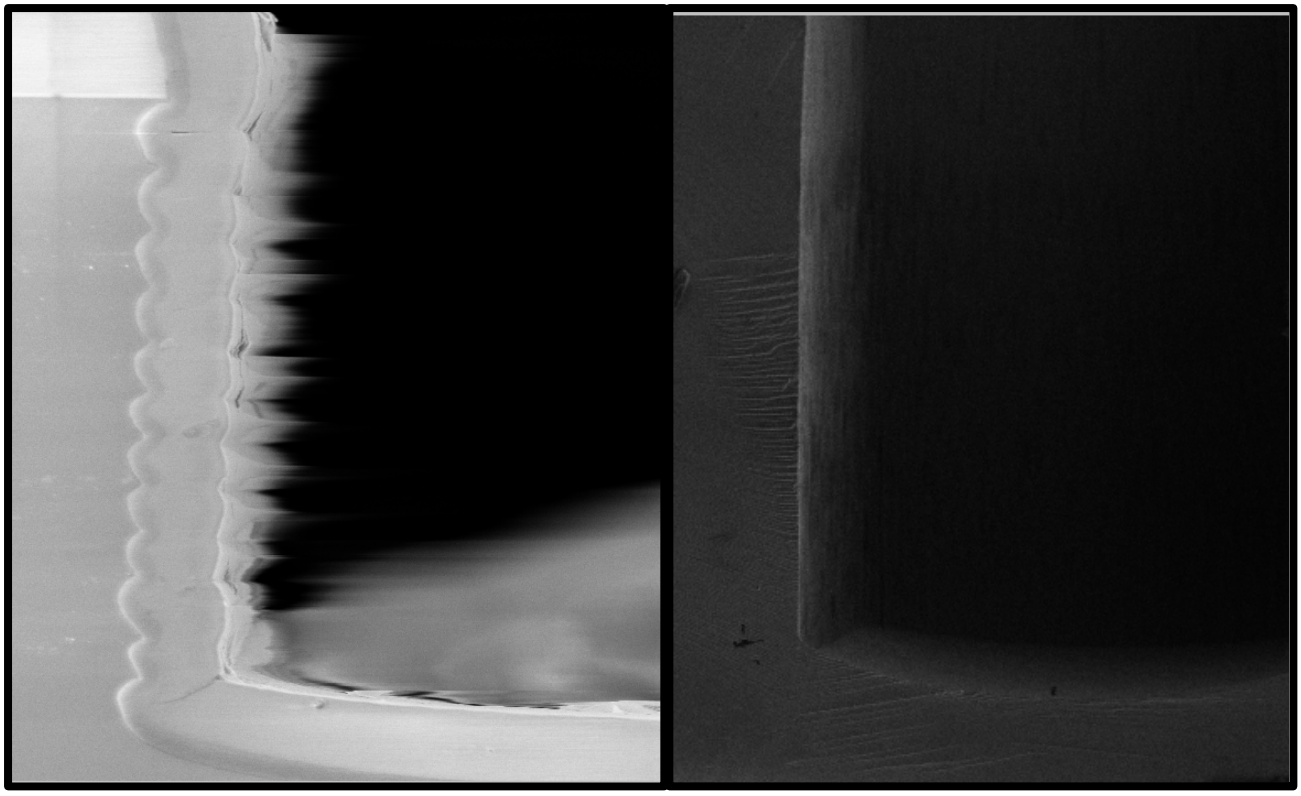

Figure A.17 Scanning Electron Microscope image of a trench etched with the Bosch Process. A zoomed-out image is given next to it for sake of reference.

One of the downsides to using the Bosch process is that the etch rate is considerably dependent upon the feature size that is being etched. This is called DRIE lag. Other than creating a mask with equal feature sizes there is little that can be done to compensate for this effect. Process parameters can be varied slightly to alleviate some effects of DRIE lag, but this usually has only provided marginal gains.

The Bosch Process does however have a very good selectivity to oxide and so this can be used as an etch stop buffer. This is not a total solution however as ion bombardment will cause the etch species to bounce off of the oxide layer and will etch the sidewalls creating a 'notch' at the base of the trench. Figure A.18 shows how notching can affect the etch profile. The trench will flare out at the bottom due to the etch species bouncing off of an oxide etch stop.

DRIE using the Bosch Process is an extremely versatile and important tool within nanofabrication, allowing extremely deep trenches with very high aspect rations, resulting in almost totally vertical sidewalls. Yet despite its usefulness,

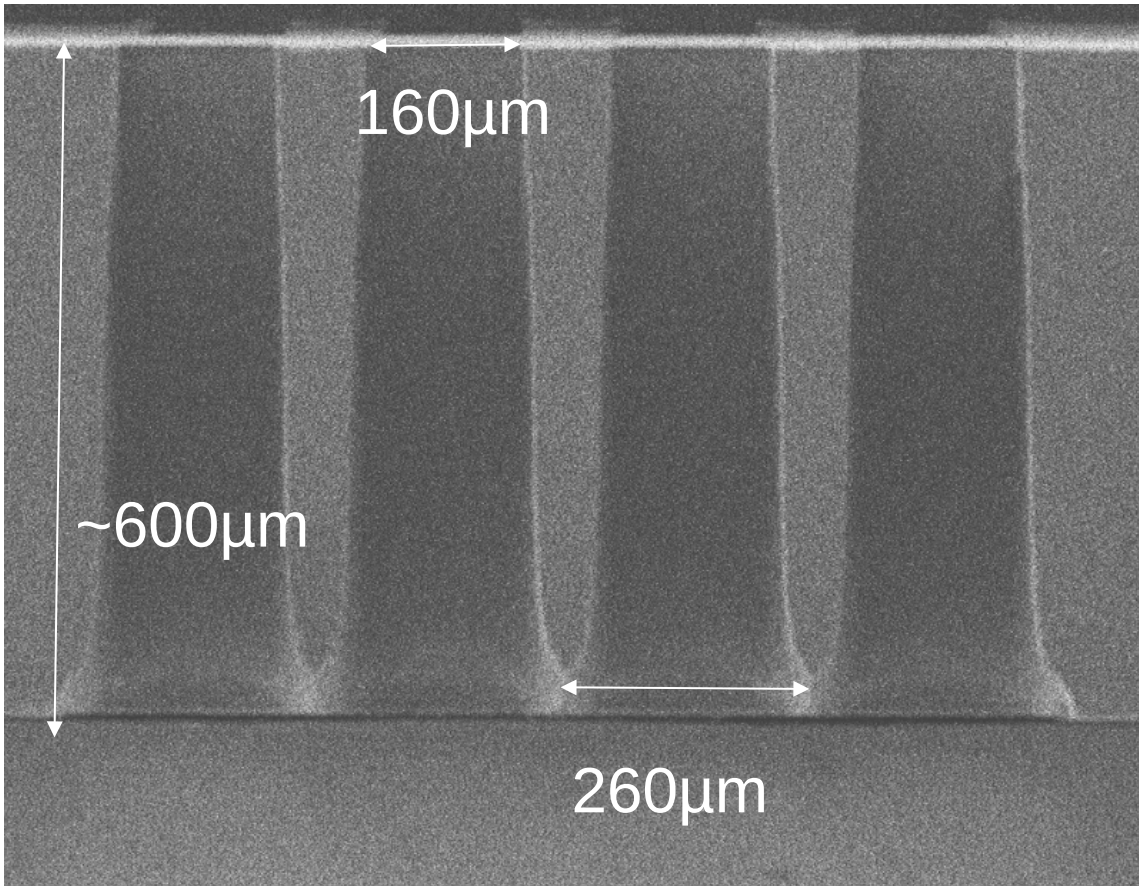

Figure A.18 Notching on a DRIE trench

Bosch Process must be tailored to each individual mask shape in order to minimize adverse effects that arise due to DRIE lag, scalloping and notching. Process development for DRIE will be explained in a subsequent chapter.

For this project a Plasma-therm Versaline was used for DRIE as depicted in Figure A.19. This model offers very advanced etching methods to produce high aspect ratios, smooth walls, high uniformity and high selectivity[193]. This model also features endpoint detection where it can detect that a material has been etched up to an etch stop layer.

Figure A.19 Plasma-Therm Versaline DRIE

e. Wet Etching

1. Wet Etching Overview

As has been already mentioned, one of the two main methods of etching is wet etching. There are a variety of different wet etches that are available to etch different substrates. Three particular etches will be considered in this chapter, all of which were used extensively in the course of this project.

Wet etching has a rich history, with evidence of metals being etched with organic acids even into prehistoric times. But even the use of a masking material to etch a metal substrate beneath dates back to at least the 16th and possibly 15th century AD [210]

Within the electronics industry wet etching was first employed in the manufacture of printed circuit boards, used mostly for the production of radio equipment needed during World War II [211]. By the 1960s the principle was being used in microelectronics to etch silicon for integrated circuits. The basic idea behind

wet etching for nanofabrication is exactly the same as that of photochemical milling which is often used in aerospace or automotive production[212]. There are four basic steps that are as follows:

1. Cleaning : The substrate to be etched must be clean in order for the etching chemical to make a good contact with the substrate which is being etched.

2. Masking: Unless a 'global' etch is required in which the entire substrate is to be etched, then a mask is required. The pattern to be etched will be incorporated into the mask and the exposed areas of the mask will then be etched by the chemical.

3. Etching: The Substrate is exposed to the etchant. This usually means that it is submerged entirely in a solution containing the etch chemical. In some instances, a continuous spray is applied the etch site which gives the same effect. To bring the etch to an end the substrate is rinsed with de-ionized water.

4. Stripping: Once the etch has been completed the masking material is removed and often cleaned again to remove any lasting residue.

These steps are in essence exactly the same that are used in nanofabrication of electronic, photonic or MEMS devices, however the scales are of course

Table A.2 Comparison of different etching methods and their associated isotropies.

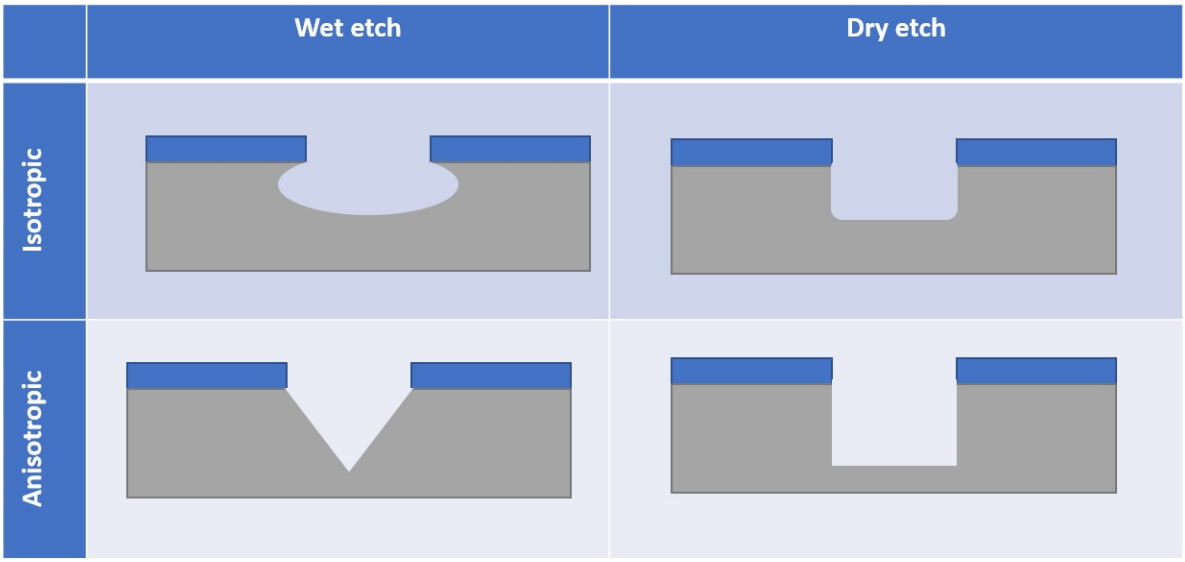

considerably smaller. This project used a mixture of fabrication techniques drawn both from IC fabrication and MEMS fabrication. IC fabrication usually demands much smaller feature sizes but MEMS fabrication often requires much greater aspect ratios to create the necessary structures [211].

Wet etching offers a few advantages over dry etching. Firstly, there are economic benefits to wet etching. This is especially the case when bulk etching. One etching tank can often hold 25-50 100mm wafers at a time[191]. So, whilst the speed of the etch is typically slower, the throughput can be considerably higher, and at lower cost as well due to the lack of expensive plasma etching tools that are required. Secondly, the rates of chemical etching are much easier to control and are less prone to fluctuating. Finally, wet etching is a purely chemical process and so selectivity can be extremely high.

Hence as long as a masking material is not reactive with the etchant then theoretically the selectivity to the substrate will be extremely high.

Another factor to consider when wet etching is isotropy which until now has been assumed to be unimportant in dry etching. In truth all etching will display isotropy to some extent but the effects in wet etching are often much more apparent. Table A.2 illustrates the effects of isotropy using different etching methods.

As can be seen, anisotropy is present in both wet etching and dry etching. Anisotropy can be extremely useful, so long as it fits the requirements needed for a specific etch. For example, the anisotropic etch shown in Table A.2 would be totally unsuitable for most waveguides yet can be useful for other structures such as a holder for fibres.

When etching a substrate in a solution a few other parameters are important. Nearly all chemical reactions are temperature dependant, so the temperature of the reactants is sometimes increased to speed up the rate of the reaction. The concentration of the etchant in solution must also be considered. Increasing the concentration of the etchant will increase the rate of reaction up to a point. Concentration and temperature can be adjusted in measure to produce an etch rate that is needed but which can be carried out safely. High temperature etchants that

are also highly acidic can pose a significant risk to the fabricator, so a compromise of concentration and temperature can be negotiated to minimize risks whilst maximising etch rates and results.

Generally speaking most wet etches fall into one of two categories, fast or slow[185]. If the reaction is fast, then the rate will be determined by the availability of the etchant which is dependent upon concentration and the speed at which etchants can be transported to the etch site and reaction products can be transported away. If the etch is slow then the reaction itself determines the rate.

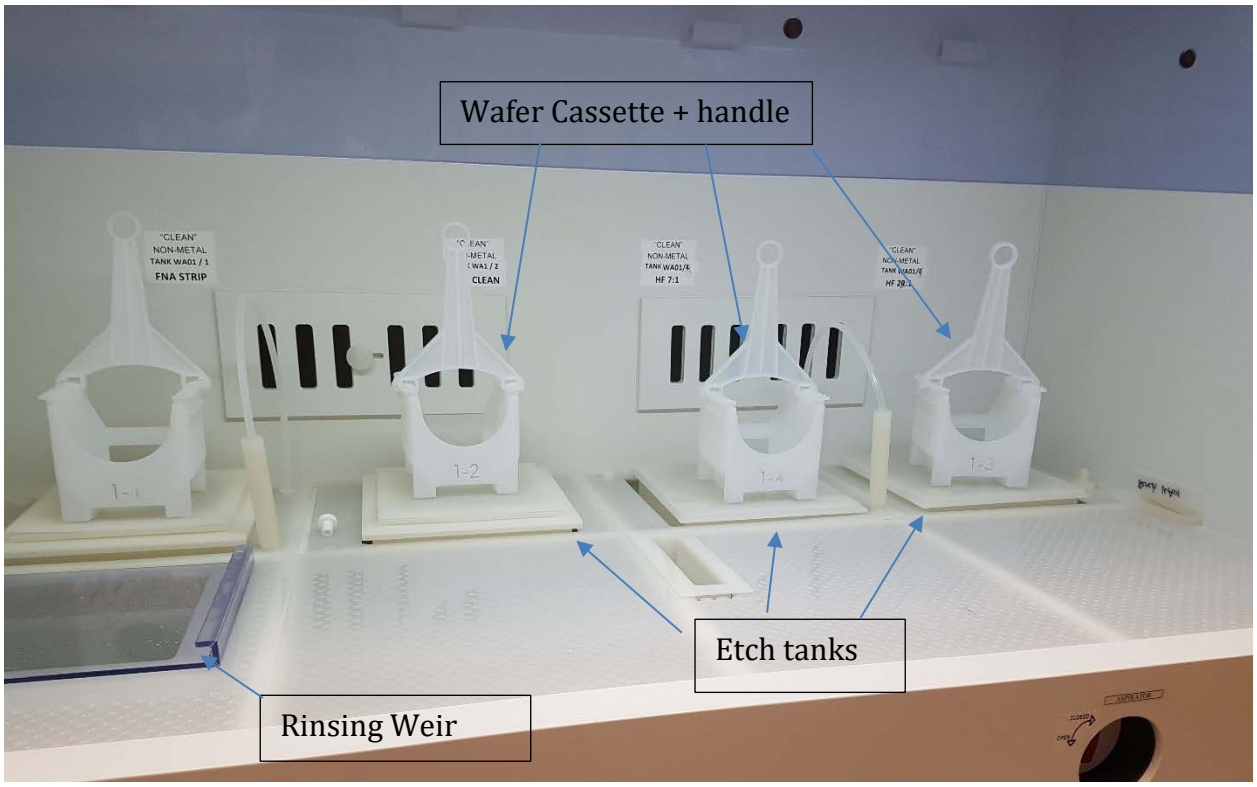

Figure A.20 Image of a wet bench for wet etching in the SNC

Figure A.20 shows a typical wet bench, which is used for wet etching, this particular one is located in the SNC and was used for fabrication in this project. There are several tanks that are each filled with a specific etching solution. Each tank is also temperature and concentration controlled. This wet bench is configured for etching entire wafers, but the principle is identical if smaller substrates are being etched. Wafers are loaded onto cassettes which are unreactive with respect to the etchants.

They are then manually loaded into the etching tank using the handle for the duration of the etch. Since a cassette can hold around 20 wafers at a time, wet etching can have a much higher wafer throughput than dry etching which is usually is restricted to a single wafer at a time and parallel processing is only possible with multiple tools and people to run them, making it extremely expensive in high volumes compared to wet etching where one user can simultaneously etch a whole cassette.

Once etching is complete the wafers and cassette must be rinsed through the adjacent weir. The weir takes its name from those which are built into rivers since they flush the tank with water in the same manner and thus remove and dilute the etchants from the wafer and so finish the etching procedure. All etching takes place in the presence of an extract to remove any gases that are produced, with a protective barrier also present.

The etchants used for wet etching are usually very strong acids or alkalis which pose a significant risk to the fabricator and hence personal protective equipment (PPE) is a must at all times, as is proper working practises.

All wet etching for this project was conducted at the SNC.

Several different wet etches will now be discussed in more detail and explaining their relevance to the project.

2. Potassium Hydroxide etching

Potassium hydroxide (KOH using chemical formula notation) is normally used to etch silicon substrates according to the following reaction[191]:

$$Si + 2OH^- \rightarrow Si(OH)_2^{++} + 4e^- \ (oxidation) \qquad \textbf{Eq. A.2}$$

$$Si(OH)_2^{++} + 4e^- + 4H_2O \rightarrow Si(OH)_6^{--} + 2H_2 \ (reduction) \qquad \textbf{Eq. A.3}$$

As can be seen this reaction includes both an oxidation and reduction step. Aqueous hydroxide ions are required to react with bulk silicon, so KOH in an aqueous solution provides an abundance of necessary etch ions. This reaction makes two products, $Si(OH)_2^{++}$ and four extra electrons and so this is an oxidation reaction. These electrons then allow the reduction to proceed spontaneously leaving H_2 gas which will escape the tank as bubbles and must be extracted and silicic acid which is a by-product.

Other etchants such as Tetramethylammonium hydroxide (TMAH) can be used to provide hydroxide ions for the same reaction but KOH was used in this project. KOH solutions for wet etching usually have a concentration of 40-50% and are carried out at temperatures of 70-90°C.

KOH etching is also interesting for fabrication since it is a strongly anisotropic etch. This comes from the crystal structure of silicon. Figure A.21 shows the different crystal planes of a silicon lattice. Silicon is a group 4 material and so

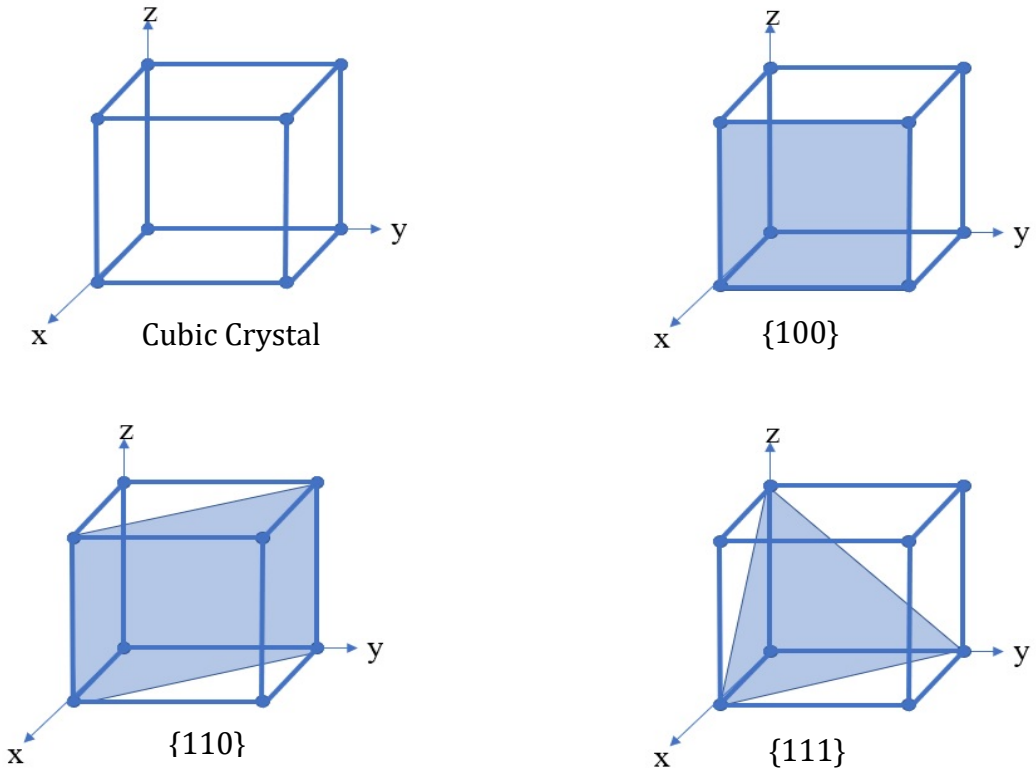

Figure A.21 Diagram showing the crystal planes of silicon

each silicon atom forms four covalent bonds with surrounding silicon atoms in a lattice. This creates the crystal structure of silicon. The bond energies between each plane are not equal and so when KOH is used to etch silicon the etch rate will vary with respect to each plane. The {100} plane is etched 200 times faster than the {111} plane[185]. This will cause the silicon trench that is formed to etch anisotropically. Since etching in the {111} plane is so much slower than the other planes it can effectively be considered an etch stop for many applications. The angle between the {111} and {100} is 54.7° to one decimal place. In a lattice structure, each coordinate can be assigned using a unit vector. Inspection of the angle between {100} and {110} shows that this is 45°. Therefore, the angle between {111} and {100} will can be calculated using basic trigonometry as follows:

$$\mathbf{1} = \sqrt{\mathbf{3}}\cos\boldsymbol{\theta} \qquad \textbf{Eq. A.4}$$

Solving this problem gives that θ = 54.7° and means that a trench forms with sidewalls that will be precisely 54.7°.

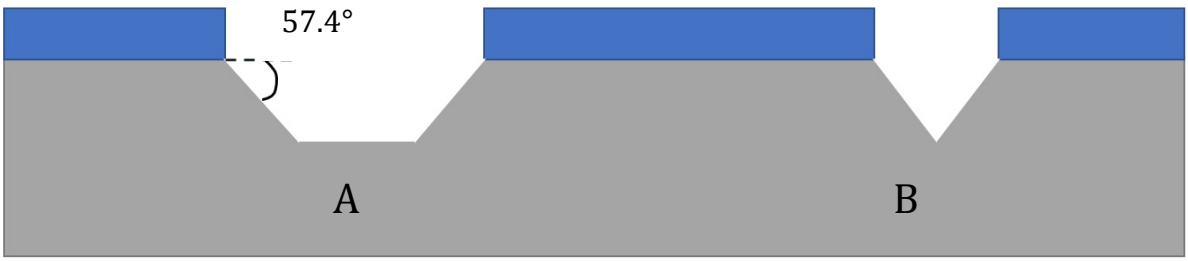

Figure A.22 Sidewall profile of Silicon etched in KOH

As can be seen from Figure A.22 the exact profile of the etched trench will depend on the width of the exposed area in the mask. In this figure, trench A is wider than trench B but both have been etched for the same time. Trench A is a trench with an angled sidewall whereas trench B has a triangular cross section, typically called a V groove due to its resemblance to the letter V. Trench B also shows an effective etch

stop. Since the {100} plane etches so much slower, the etch will effectively cease once it has reached a depth as defined by the width of the opening and from the angle subtended at 54.7°. Given a longer etch time, trench A would also eventually form a perfect V groove but with a greater depth. This can also be extended from two to three dimensions and inverted pyramids can be made from a square shaped mask, or a V grooved trench can be made from a long rectangular shaped mask.

3. Hydrofluoric Acid Etching

The next wet etch that should be discussed is hydrofluoric acid etching. Hydrofluoric acid (HF) is an extremely strong acid, with a disassociation constant of 3.17[213] meaning that very low pH solutions can be made which are very highly acidic. Unlike KOH, HF is highly selective to oxide rather than to silicon. HF and oxide is an acid base reaction rather than electron transfer reaction as with KOH and silicon. The basic equation can be seen in Equations A.5 and A.6[185]

$$SiO_2 + 6HF \rightarrow H_2SiOF_6 + 2H_2O \qquad \textbf{Eq. A.5}$$

$$SiO_2 + 4HF \rightarrow SiF_4 + 2H_2O \qquad \textbf{Eq. A.6}$$

This reaction is often used to remove '*native*' oxide that will develop on a silicon substrate that is exposed to air. This is typically only a few nm at most and can be etched away very quickly to leave the bare silicon underneath. It can also be used for etching thicker films as has been used in this project.

Typically, HF based wet etching is used with a buffer. This will slow down the etch rate making the whole process far more controllable. Table A.3 gives a comparison of different buffering solutions, adapted from Kim et al. [214].

Table A.3 Comparison of etch rates with different HF buffering solutions (nm/min)

	Wet Oxide	TEOS	Nitride	Al
HF (49%)	1763	3969	15	38
NH_4F:HF (7:1)(BHF)	133	107	1	3
HF:H_2O 1:10	48	157	1.5	320

As can be seen from the table, the etch rate can be hugely changed by the addition of a buffering agent. In the case of Aluminium, the etch rate can actually increase after buffering but this is due to the presence of water which can catalyse the reaction. Temperature can be controlled to affect etch rate, but this is usually unnecessary, since the etch rate is fast enough for most processes.

4. Phosphoric Acid Etching

Phosphoric acid etching (also called hot-phosphoric etching or orthophosphoric acid etching) is typically used to etch silicon nitride substrates. The reaction is as follows[215]:

$$3Si_3N_4 + 4HPO_4 + 27H_2O \rightarrow 4(NH_4)_3 + 9H_2SiO_3 \qquad \textbf{Eq. A.7}$$

Silicon nitride is usually very chemically resistant to etching but can be etched according to the process outlined in Eq. A.7. This reaction can be extremely slow and so to increase the rate the temperature is often raised upwards of 150°C. Phosphoric acid has a boiling point of 158°C so will begin to evaporate and so this procedure must be done under reflux, where water is dripped into the bath

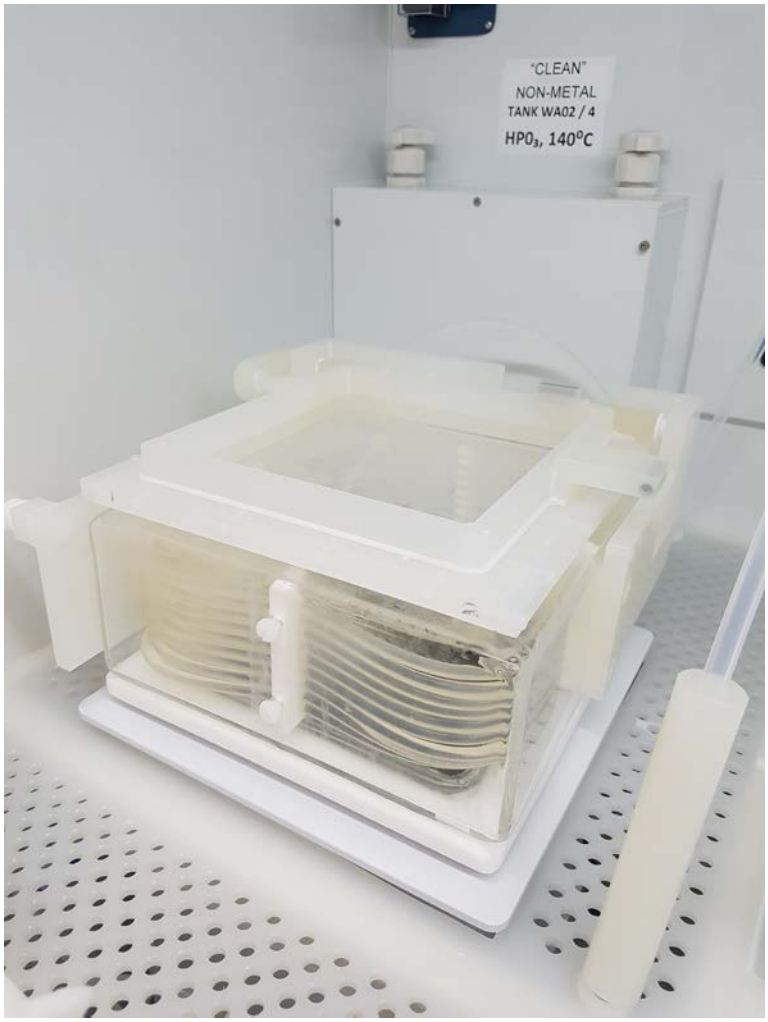

Figure A.23 Image of a Phosphoric Acid Etching Bath in SNC

allowing the concentration of phosphoric acid to remain constant with a temperature above the boiling point. Phosphoric acid has good selectivity to both silicon and oxide. This means that both materials can be used for masking or etch stops when etching a nitride film. Figure A.23 shows an image of a hot phosphoric etching bath. Unlike a typical etching bath, it is temperature controlled by PID

(proportional integral derivative) control. It also has a controlled drip mechanism to keep the concentration of phosphoric acid constant. There are also water-cooling lines, which surround the air immediately above the bath that help to condense any vapours that escape during etching.

f. Deposition

1. Deposition Overview

Deposition is necessary for a number of reasons. The most basic need for deposition is for passivation. Passivation is the process by which the surface of a substrate is covered with an inert substance (usually SiO_2) in order to protect it from reacting with other elements. In microelectronics and photovoltaics this is extremely useful to prevent the electrical properties of the device from shifting over time, but it is also useful in photonics to prevent the device's optical properties changing over time. In addition to passivation, deposition is necessary for creating hard masks. A hard mask works on the very same principle as a soft mask but is made out of a 'hard' material such as oxide or perhaps nitride. Once a film has been deposited then it can be patterned and then etched, leaving the substrate beneath exposed to be etched further. Deposition also can be used to form an etch stop on the back of a wafer as explained previously. The final need for passivation in this project was for protection. In order to protect devices that have been fabricated from being damaged during further processes, a protective layer can be deposited onto the wafer, which forms a defensive shield.

Two main forms of deposition are considered for this project, plasma enhanced chemical vapour deposition and low-pressure chemical vapour deposition. Each of these methods will now be explained in turn giving their relative strengths and weaknesses.

2. Plasma Enhanced Chemical Vapour Deposition

Plasma enhanced chemical vapour deposition (PECVD) has a number of key advantages. It is a low temperature process and is usually carried out below 400°C. It is a quick process meaning that films can be deposited at high rates. Adhesion is usually very good and the step coverage is also good [216].

The chemical reaction required for deposition is activated by a high-energy plasma rather than using thermal energy at very high temperatures (up to 1000°C in some cases). The design of a PECVD chamber is very similar to that of RIE but of course the gases present are different, and the RF power is also different. Since they are usually the same tool, this has the added bonus that etching can be used to clean the chamber. As PECVD happens it will of course deposit a film on the chamber and any other exposed areas. This needs to be cleaned away periodically but can simply be done with an etch recipe.

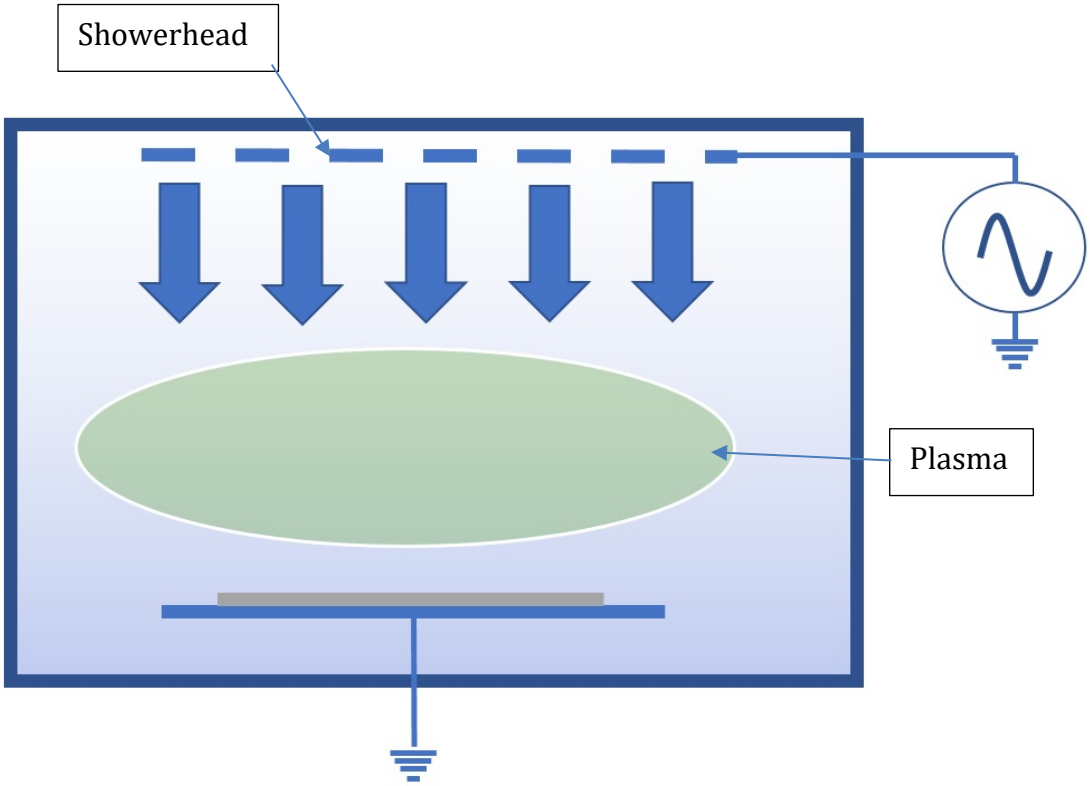

Figure A.24 Diagram of PECVD chamber

Figure A.24 shows a diagram of a PECVD chamber. Note the similarities between the PECVD chamber and the RIE. Indeed often, the production tools are capable of both processes with only a change in a few parts, software and recipes needed. There are a couple important differences, however. Firstly, the top electrode is now a showerhead, as symbolised by the dotted line. Reaction gases are funnelled through the showerhead and just like a conventional shower for washing, the gases are spread and distributed

across the chamber and hence across the wafer. The plasma is symbolised by the green cloud, which energises the reactive chemicals in order to allow spontaneous deposition.

In this project an Oxford Instruments Plasmalab System 100 was used for PECVD. As can be seen this is the same model as is used for ICP etching but with only a few slight different parts and running a different process.

Figure A.25 Oxford Instruments PlasmalabSystem 100 in the SNC used in this project.

3. Low Pressure Chemical Vapour Deposition

Low pressure chemical vapour deposition (LPCVD) will now be discussed and explained. As Boyle's law predicts, if the pressure is reduced then by contrast the reciprocal constant of proportionality (temperature) must increase. So, with LPCVD, whilst the pressure is decreased the temperature is greatly increased. Typical temperatures range from 700-1000°C and pressures are usually at ~1 Torr [185].

Due to the high temperature requirements for LPCVD it is usually carried out within a furnace which can accommodate such extremely high temperatures.

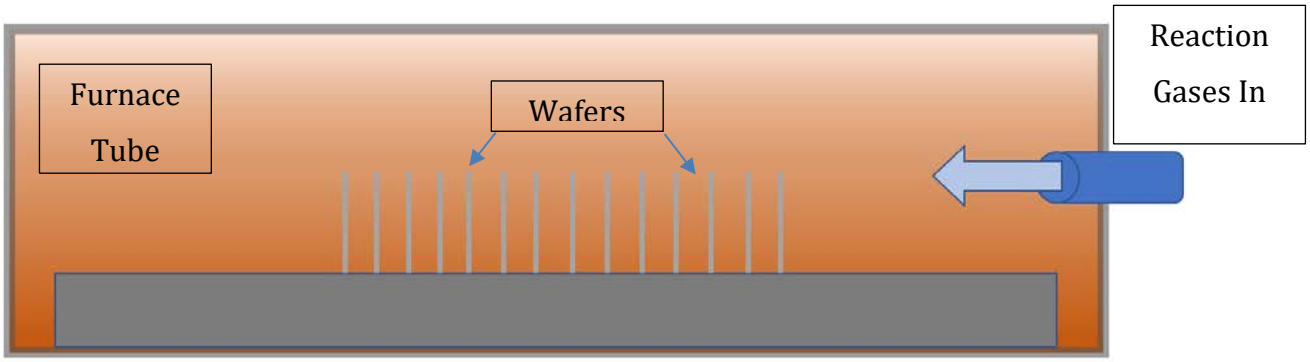

Figure A. 26 Diagram of LPCVD furnace

Both oxide and nitride films were grown onto silicon wafers for this project for different uses. The following reactions show the chemical change that is used to grow a film.

$$\textit{Nitride Deposition} \quad 3\,SiH_2Cl_2 + 4\,NH_3 \rightarrow Si_3N_4 + 6HCl + 6H_2 \qquad \textbf{Eq. A.8}$$

$$\textit{Oxide Deposition} \quad SiH_4 + 0_2 \rightarrow SiO_2 + 2H_2 \qquad \textbf{Eq. A.9}$$

LPCVD has a few benefits over PECVD. Firstly, films can be deposited on both sides of the wafer simultaneously and will be of identical thickness. Secondly the throughput will be much higher since many wafers can be placed in the same LPCVD tube at the same time. So then, even though the process is normally slower than PECVD, the throughput can be much higher when a large quantity of wafers is

needed. Indeed, a slower deposition rate can sometimes be considered an advantage when done in a batch since it makes the deposition thickness more controllable. LPCVD films typically etch much slower than PECVD ones. This can be seen either as an advantage or disadvantage but films grown onto a substrate are often used for hard-masking and in that case a slower etch rate is advantageous.

g. Metrology

When fabricating PICs, or any devices for that matter, it is necessary to have a way of quantitatively measuring the processes that have been carried out. With respect to nanofabrication as it has been outlined throughout this chapter, there are a few specific metrics that must be used. These are etch depth, film thickness and dimensions of the features that have been created. Material quality is a more qualitative metric that can also be ascertained using metrology tools. This section will describe the different metrology tools that have been used in this project, and hence all metrics that are given in this thesis are derived using such tools, including their accuracies and tolerances.

1. Ellipsometry

Ellipsometry is a method of obtaining the thickness of a thin film deposited upon a bulk substrate, first described by Rothen in 1945[217]. This method takes its name from the fact that it uses light polarized elliptically and is used for metrology. When light reflects off a surface it can be detected and there will be a change in polarization, which is represented as both an amplitude ratio Ψ and a phase change Δ[218]. This polarization difference will be a function of the sample's thickness and refractive index. Therefore, if the material is known then its thickness can be obtained from the information provided by the ellipsometer. Equally if the thickness of a film in known precisely then the identification of the material can be obtained given from the refractive index. In some cases, such as with nitride, the stoichiometric make-up of the material can vary and therefore its refractive index too. So then, by comparing the ellipsometer reading to that of a known Si3N4 sample, then a qualitative assessment can be made of a given nitride sample's stoichiometry.

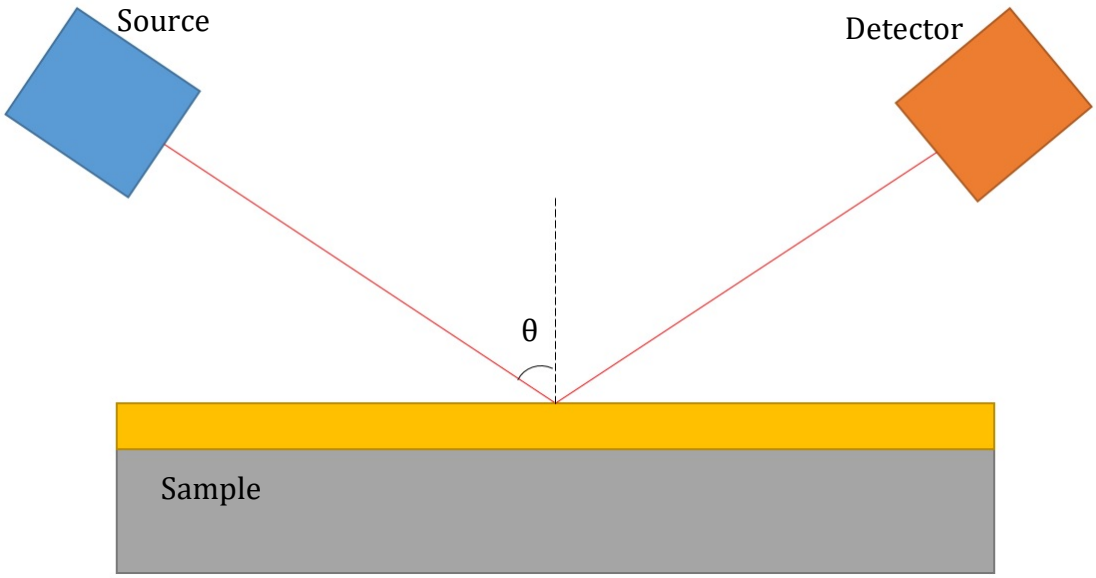

Figure A.27 Diagram of ellipsometry

Figure A.27 shows a diagram of how ellipsometry works. A source provides a beam with a known wavelength and polarization that is projected towards the sample. In this example the sample is an oxide film on a silicon wafer. The reflected beam is received by the detector which can then compares the beam with what was at the source. Angle θ, of the both the source and detector must also be known. Then in this way the thickness of the oxide sample can be obtained by using a lookup table with known refractive indices of an oxide film on silicon.

It is usual practice to create metrology boxes in a wafer that are large enough to be used specifically for the purposes of ellipsometry. A metrology box is essentially just a large box, roughly 1mm x 1mm which is also etched at the same time as any devices on the wafer. This box then can be used with an ellipsometer to calculate the etched depth when etching into a substrate.

For this project a J. A. Woollam ellipsometer was used within the SNC. This machine has a large library of materials but also has the option to create new materials, for example a mixture of different films on top of one another. Figure A.28 shows an image of the machine used. This model features a motorized stage with digital input so that different regions of a sample or wafer can be measured. It also features very small spot size of 25 x 60μm so that very small features can be measured. High-speed CCD detection allows results to be obtained very quickly.

Figure A.28 Image of a J. A. Woollam ellipsometer in the SNC

2. Profilometery

Once a substrate has been etched, after lithography and features have been made then a suitable method needs to be obtained to measure the etch depth as well as the profile of the etch. A profilometer can do such a job. Usually speaking a contact profilometer is used. A contact profilometer uses a small needle on a very carefully calibrated arm to measure the profile of an etched trench. Microscopic grooves and trenches on the surface of a wafer can be translated into electrical signals which are then used to plot a profile.

Figure A.29 shows a diagram of how a profilometer uses a small needle to plot a profile of a trench which has been etched into a substrate. From this data, the etch depth can be calculated by the total displacement of the needle as it travels over the surface. The needle can also pick up any imperfections in the surface of the substrate. The other useful feature is that the plotted graph will give an accurate side-on profile of the etched trench. The angle of the sidewall can also be calculated from this data which cannot be easily obtained otherwise.

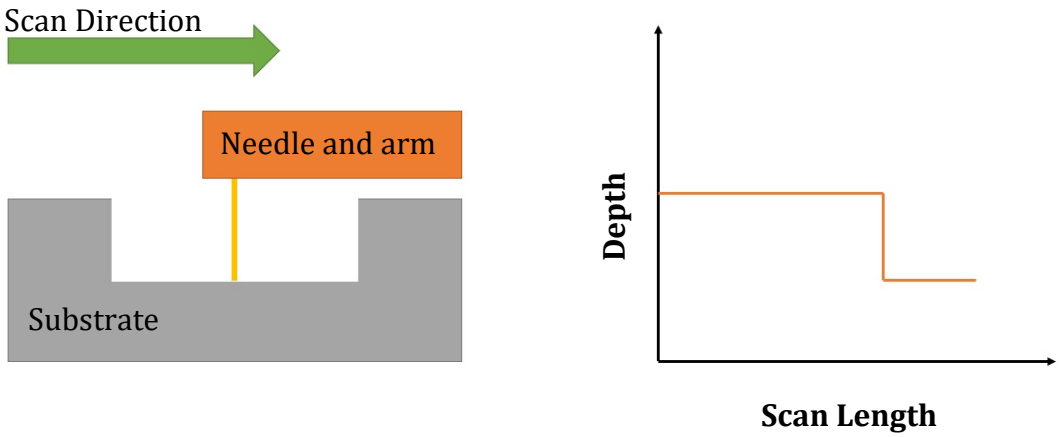

Figure A. 29 Diagram of profilometer and example of a plotted profile

Profilometry is a very powerful tool when used in conjunction with ellipsometry. When etching a substrate with a resist as a mask then both these tools can be used to calculate the etch depth.

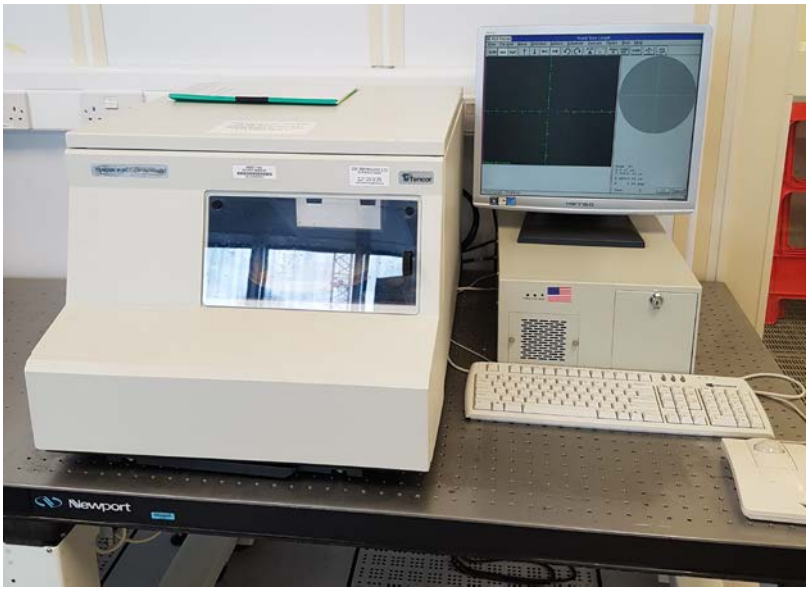

Figure A.30 Image of KLA-Tencor P-11 Profilometer in the SNC

For this project a KLA-Tencor P-11 Profiler was used for profilometry, as shown in Figure A.30. This model can measure a total step of up to 327µm vertically. This includes both directions so normally speaking an etch depth of half that can be measured. Since a silicon wafer is typically around 600µm this is a significant depth and more than adequate for most etching purposes. This particular model is quite dated as of 2018, with the P-17 now available. Because of this, accurate resolution statistics are hard to obtain. It does appear however to be close to 0.8nm. Over a range of over 100µm.

There are limitations to this tool however. The thickness of the needle will determine the smallest size of feature that can be measured. This coupled with a low sampling rate can lead to erroneous profile plots. The maximum sampling frequency of 200Hz is more than adequate for most structures but can occasionally approach the pitch frequency of some devices such as grating couplers meaning this tool can become inaccurate.

3. Optical Microscopy

Optical microscopes are useful for both qualitative and limited quantitative analysis. A microscope is necessary when determining the quality of a photoresist or a developed trench which is simply not possible by eye. Very often when developing a wafer for etching, a wafer can appear to be fully developed but under

a microscope it can be seen that some resist residue is still present at the bottom of the trench which cannot be seen otherwise.

Optical microscopes can also be used to moderate success to determine extremely deep etches, such as DRIE trenches. The maximum thickness of a profilometer is usually only a few hundred microns. In the occasion that a deeper etch is required then optical microscopy can be used to determine the etch depth. This is however not an 'exact science' as the method makes clear. In order to estimate the etch depth the microscope is first focused onto the surface of the substrate and then on the bottom of the trench by moving the microscope stage vertically. The total displacement required to bring the bottom of the trench into focus will give a good estimate for the etch depth. This is of course dependant on what the microscope user deems to be "in focus" and hence the accuracy is usually only ±5μm. Whilst this would be totally unacceptable for most optical devices it can

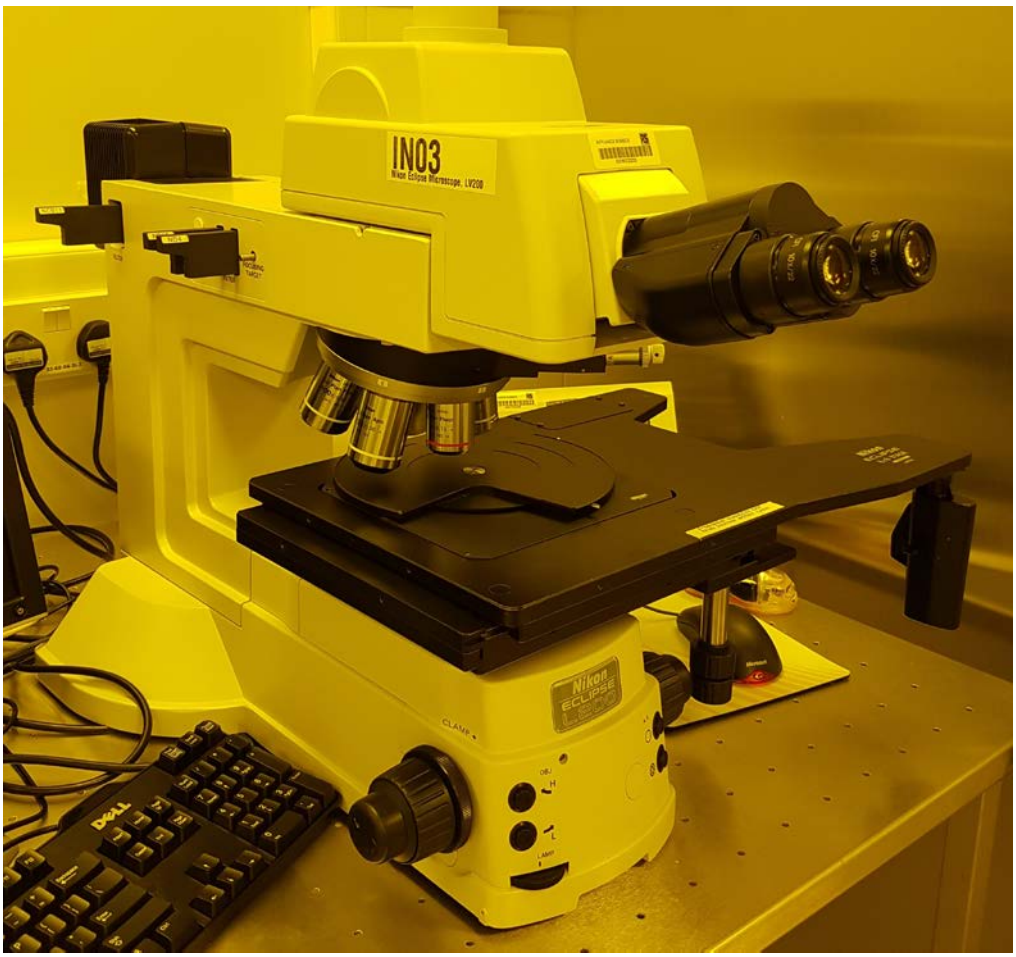

Figure A.31 Image of an optical microscope used for this project in the SNC yellow-room

give a good enough estimate for some larger structures, as was used in this thesis. The other limitation with optical microscopy is that very deep trenches can be difficult to illuminate sufficiently such that they can be properly analysed.

4. Scanning Electron Microscope

In the same way that an electron beam could be used to write smaller structures than a photonic one in lithography, the same concept can be applied to microscopy. Scanning electron microscopy (SEM) uses an electron beam to 'observe' a sample. Since the wavelength of electrons is so small an SEM can resolve objects down to 5 nm[185]. Figure A.32 shows an SEM image of a photonic device on a silicon PIC to give an idea of the imaging capability of an SEM.

Electrons are accelerated across a voltage of 1-50kV into a beam of spot size 1-10nm and a current of 0.1-0.001nA[219]. Larger currents can be used to if a greater spot size is needed. Deflection coils are used to raster scan the beam across the surface of the sample.

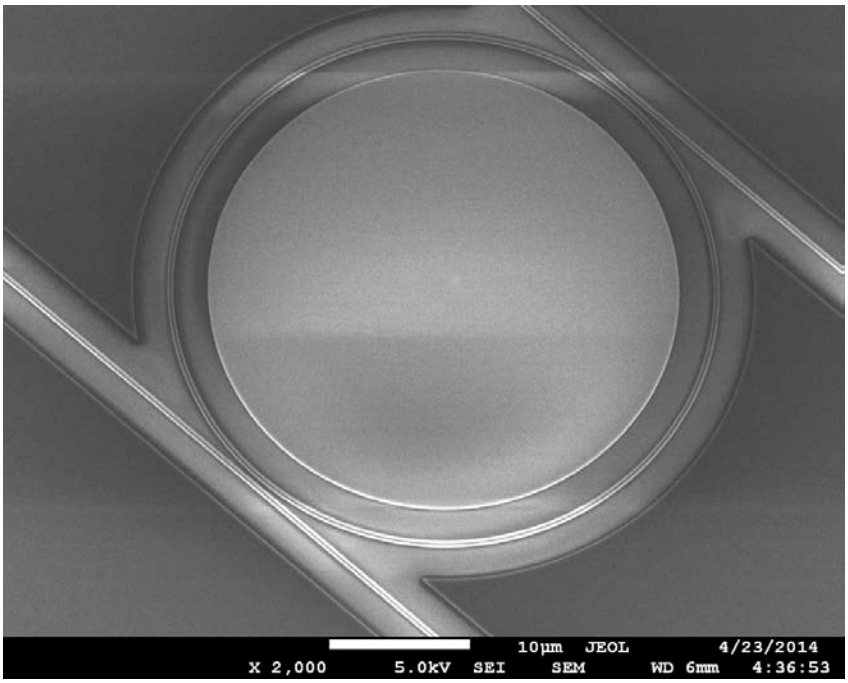

Figure A.32 SEM image of a photonic device fabricated in silicon

In order to increase visibility, dielectric substrates are usually first coated in a thin layer of a conductor, usually gold. Semiconductors such as silicon however usually are conductive enough to be seen clearly. Whilst SEM is a very useful tool for metrology, it is also usually a sacrificial process. Samples must usually be cleaved or diced before they can be viewed since they must be small enough to fit in the SEM chamber. In addition, metal contamination will occur, either from metal deposition or from other samples that have been used in the same SEM. Therefore, this technique, though powerful, can only be used sparingly.

For this project a JEOL JSM-7500F SEM was used for imaging samples in order to give both a qualitative and quantitative of etch profiles where profilometry was insufficient. This tool can image a sample in three dimensions giving a much greater understanding of the etched trench topology.

Figure A.33 Image of a JEOL JSM-7500F SEM used in the SNC

A.2 Full Process Flows

Identifier			wafer 1	wafer 2	wafer 3	wafer 4	wafer 5	wafer 6
Deposit B/S Sio2		1	✓	✓	✓	✓	✓	✓
S1813 S/B	Alignemnt Marks	2	✓	✓	✓	✓	✓	✓
		3	✓	✓	✓	✓	✓	✓
Expose		4	✓	✓	✓	✓	✓	✓
Develop		5	✓	✓	✓	✓	✓	✓
Etch		6	✓	✓	✓	✓	✓	✓
Ash		7	✓	✓	✓	✓	✓	✓
Zep S/B	Devices	8	✓	✓	✓	✓	✓	✓
		9	✓	✓	✓	✓	✓	✓
Espacer		10	✓	✓	✓	✓	✓	✓
Ebeam		11	✓	✓	✓	✓	✓	✓
Develop		12	✓	✓	✓	✓	✓	✓
Etch		13	✓	✓	✓	✓	✓	✓
Ash		14	✓	✓	✓	✓	✓	✓
Sio2 Dep		15	✓	✓	✓	✓	✓	✓
AZ2070 S/B	Fibre Cavity Front	16	✓	✓	✓	✓	✓	✓
		17	✓	✓	✓	✓	✓	✓
Espacer		18	✓	✓	✓	✓	✓	✓
Ebeam		19	✓	✓	✓	✓	✓	✓
Rinse		20	✓	✓	✓	✓	✓	✓
UV Flood		21	✓	✓	✓	✓	✓	✓
Bake		22	✓	✓	✓	✓	✓	✓
Develop		23	✓	✓	✓	✓	✓	✓
SiO2 etch		24	✓	✓	✓	✓	✓	✓
Deep etch		25	✓	✓	✓	✓	✓	✓
Ash		26	✓	✓	✓	✓	✓	✓
Si3n4 dep		27	✓	✓	✓	✓	✓	✓
Sio2 Dep		28	✓	✓	✓	✓	✓	✓
AZ9260 s/b	Fibre Cavity Back	29	✓	✓	✓	✓	✓	✓
		30	✓	✓	✓	✓	✓	✓
Expose		31	✓	✓	✓	✓	✓	✓
Develop		32	✓	✓	✓	✓	✓	✓
Sio2 etch		33	✓	✓	✓	✓	✓	✓
DRIE		34	✓	✓	✓	✓	✓	✓
Sio2 strip		35	✓	✓	✓	✓	✓	✓
Si3N4 Strip		36	✓	✓	✓	✓	✓	✓

Wafer		1	2	3	4	5	6
Location							
Identifier							
AZ9260 Spin	Photolith CGA marks	✓	✓	✓	✓	✓	✓
Bake		✓	✓	✓	✓	✓	✓
Rest		✓	✓	✓	✓	✓	✓
Expose		✓	✓	✓	✓	✓	✓
Develop		✓	✓	✓	✓	✓	✓
Nitride Etch		✓	✓	✓	✓	✓	✓
Ash		✓	✓	✓	✓	✓	✓
KOH Etch		✓	✓	✓	✓	✓	✓
AZ9260 Spin	Photolith Trenches +Spheres	✓	✓	✓	✓	✓	✓
Bake		✓	✓	✓	✓	✓	✓
Rest		✓	✓	✓	✓	✓	✓
Expose		✓	✓	✓	✓	✓	✓
Develop		✓	✓	✓	✓	✓	✓
Nitride Etch		✓	✓	✓	✓	✓	✓
Ash		✓	✓	✓	✓	✓	✓
KOH Etch		✓	✓	✓	✓	✓	✓
Orthophos	Nitride Strip	✓	✓	✓	✓	✓	✓
Au Deposition	Deposition	✓	✓	✓	✓	✓	✓
Lithography		✓	✓	✓	✓	✓	✓
Gold etch		✓	✓	✓	✓	✓	✓

A.3 Device Chip Layout

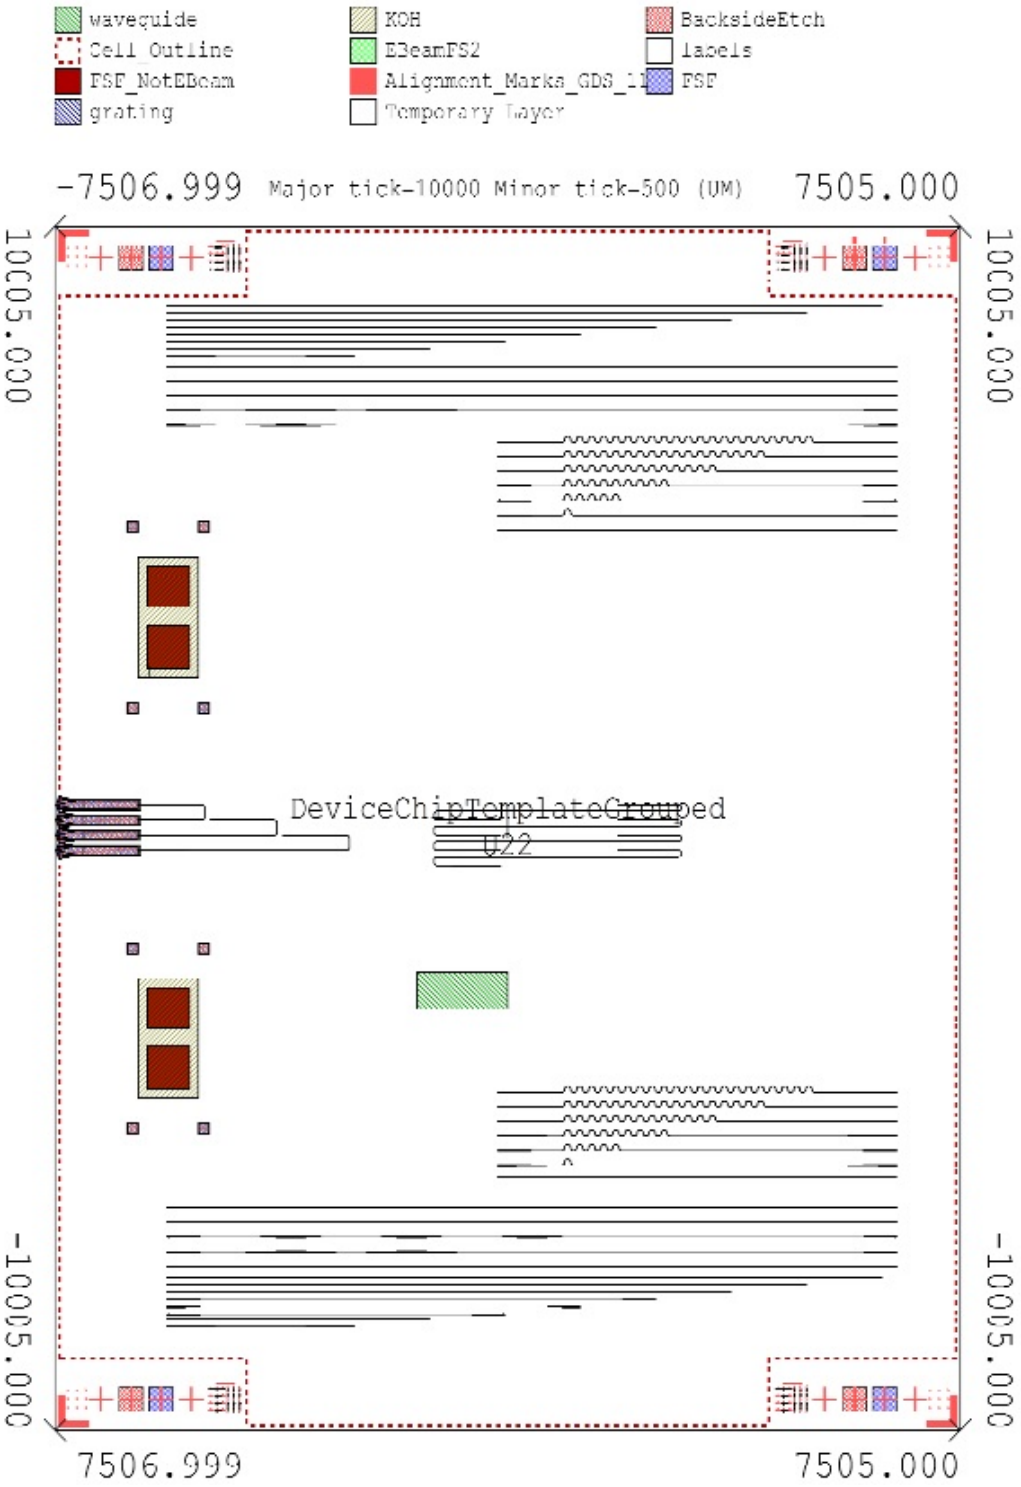

A.4 KOH etch MATLAB script

```
%% etch rates
ER_40pc = [ 2, 5, 7, 12, 23, 37.5, 62.5, 112.5, 187.5];
```

```
%% temperatures
T = [ 20, 30, 40, 50, 60, 70, 80, 90, 100] ;
Temp = input('Please select the temperature of the etch from the
following list by entering the associated number. \n 1 -- 20 degrees C
\n 2 -- 30 degrees C \n 3 -- 40 degrees C \n 4 -- 50 degrees C \n 5 --
60 degrees C \n 6 -- 70 degrees C \n 7 -- 80 degrees C \n 8 -- 90 degrees
C\n 9 -- 100 degrees C \n --->')
w = input('Please enter the width of the V-groove in microns:'); % width
of groove

d = w * tan( pi * (54.7/180)); %% depth to etch to

t = ( d / ER_40pc(Temp));

formatSpec = 'This etch will take %3.2f hours \n';
fprintf(formatSpec,t)
```

A.5 C++ Mask Script

```
/*****************************************************************************
* Macro Name: PropLossGrat
File Name: PropLossGrat.cpp
* Creator  : Nathan Soper
* Revision History: v1
* 18/03/17
*****************************************************************************/


#include <cstdlib>
#include <cstdarg>
#include <cstdio>
#include <cstring>
#include <cctype>
#include <cmath>
#include <iostream>
#include <string>
#include <sstream>
#include <stdio.h>
#include <math.h>
#include <cmath>

#define EXCLUDE_LEDIT_LEGACY_UPI

#include <ldata.h>

//Include the user component function files
#include "components.h"


using namespace std;
using std::string;
using std::cout;
using std::endl;

extern "C" {
       void PropLossGrat(void);
       int UPI_Entry_Point(void);
}
```

```
//////////////////////// main function ////////////////////////////////////
void PropLossGrat(void)
{

	//////////////////////Set up layers////////////////
	LCell  pCell = LCell_GetVisible();
	LFile  pFile = LCell_GetFile(pCell);

	LPoint centre = LPoint_Set(0, 0);
	//LLayer layer_waveguide[] = {LLayer_Find(pFile, "New Layer")};
	//"layer" variable is used to transfer control of which layer is being
written to.

	//SetLayersAll(pFile,  LAY_waveguide,  LAY_P_dop,  LAY_Ppp_dop,  LAY_N_dop,
LAY_Npp_dop, LAY_metal, LAY_gratings, LAY_vias);  //Create all layers.
	//LLayer layer[]= {LLayer_Find(pFile, "Waveguides")};
	SetLayersAll(pFile,  LAY_waveguide,  LAY_P_dop,  LAY_Ppp_dop,  LAY_N_dop,
LAY_Npp_dop, LAY_metal, LAY_gratings, LAY_vias);  //Create all layers.
	LLayer layer[] = {
		LLayer_Find(pFile, LAY_waveguide), //0
		LLayer_Find(pFile, LAY_P_dop),           //1
		LLayer_Find(pFile, LAY_Ppp_dop),  //2
		LLayer_Find(pFile, LAY_N_dop),           //3
		LLayer_Find(pFile, LAY_Npp_dop),  //4
		LLayer_Find(pFile, LAY_metal),           //5
		LLayer_Find(pFile, LAY_gratings), //6
		LLayer_Find(pFile, LAY_vias)             //7
	};

	//////  1. declare all component types, and set all common variables:

	//Declare grating coupler struct and set variables:
	GratingCouplerSpecial_TYPE GC;
	GC.Period = 0.533 * 1000;
	GC.W_grat = 17 * 1000;
	GC.L_grat = 25 * 1000;
	GC.W_wg = 14 * 1000;
	GC.angle = 0;
	GC.InPort = centre;
	GC.Duty = 0.4;
	GC.Teeth = 40;
	GC.PeriodApp1 = 0.49 * 1000;  // The periods for the appodized section of
the grating coupler
	GC.PeriodApp2 = 0.49 * 1000;
	GC.PeriodApp3 = 0.49 * 1000;
	GC.PeriodApp4 = 0.49 * 1000;
	GC.PeriodApp5 = 0.49 * 1000;
	GC.Duty1 = 0.4;
	GC.Duty2 = 0.4;
	GC.Duty3 = 0.4;
	GC.Duty4 = 0.4;
	GC.Duty5 = 0.4;

	//Declare waveguide struct and variables:
	WG_TYPE WG;
	WG.W = 0.45 * 1000;
	WG.angle = 180;
	WG.L = 5000 * 1000;

	//Declare waveguide struct and variables:
```

```
	WG_TYPE WG_cl;
	WG_cl.W = 0.45 * 1000;
	WG_cl.angle = 180;
	WG_cl.L = 3500 * 1000;

	//Declare input taper struct and set variables;
	Taper_TYPE TAPin;
	TAPin.W1 = GC.W_wg;
	TAPin.W2 = WG.W;
	TAPin.angle = 0;
	TAPin.L = 1000 * 1000;

	//Declare output taper struct and set variables;
	Taper_TYPE TAPout;
	TAPout.W1 = WG.W;
	TAPout.W2 = GC.W_wg;
	TAPout.angle = 0;
	TAPout.L = TAPin.L;

	//Declare Bend and set variables
	bend_TYPE Bend;
	Bend.W = WG.W;
	Bend.radius = 50 * 1000;


	/////// 2. Draw waveguides of different lengths, with with grating couplers
both LHS and RHS:
	int i, j;
	int n_lengths = 7;                //Set number of length variations
	long d_length = 2500 * 1000;      //Set difference in length between each
variation
	long L0 = 2500 * 1000;            //Set   the   length   of   the   shortest
waveguide
	long offset = 300 * 1000;         //Set offset between adjacent waveguides.
	long L_total = 4000 * 1000; //Set total waveguide length (direct), including
bends.ie taper to taper length
	long BendSep = 200 * 1000;  //Set the separation between the starts of the
two bends in the middle of the Prop structure

	for (i = 0; i<n_lengths; i++)
	{
		//First set the changing length for each iteration
		long PL_in = (L0 + (i - 1)*d_length) / 4;
				//length of prop loss section input part
		long PL_mid = 2 * Bend.radius + BendSep + (L0 + (i - 1)*d_length) /
2;		//length of prop loss section middle part
		long PL_out = (L0 + (i - 1)*d_length) / 4;
					//length of prop loss section output part


		GC.InPort.x = centre.x;     //Set RHS grating start point.
		GC.InPort.y = centre.y + i*(offset + 4 * Bend.radius); //Increment
y-coord by "offset" value.
		GratingCouplerSpecial(pFile,   pCell,   layer,   GC,   GC.OutPort);
	//Draw LHS grating.

		TAPin.InPort = GC.InPort;   //Set LHS taper staring point
		TAPin.angle = 0;
		Taper(pCell, layer[0], TAPin, TAPin.OutPort);    //Draw LHS Taper

		Bend.InPort = TAPin.OutPort;
		Bend.start_angle = 0;
```

```
Bend.end_angle = 90;
bend(pCell, layer[0], Bend, Bend.OutPort);

Bend.InPort = Bend.OutPort;
Bend.start_angle = 90;
Bend.end_angle = 180;
bend(pCell, layer[0], Bend, Bend.OutPort);

//Draw length of waveguide to the first bend:
WG.InPort = Bend.OutPort;
WG.L = (L_total - BendSep) / 2;
waveguide(pCell, layer[0], WG, WG.OutPort);


//Draw first two input bends
Bend.InPort = WG.OutPort;
Bend.start_angle = 180;
Bend.end_angle = 90;
bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 1

Bend.InPort = Bend.OutPort;
Bend.start_angle = 90;
Bend.end_angle = 0;
bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 2

//Draw Prop loss section waveguide input part
WG.angle = 0;
WG.L = PL_in;
WG.InPort = Bend.OutPort;
waveguide(pCell, layer[0], WG, WG.OutPort);

//Draw two bends to middle part
Bend.InPort = WG.OutPort;
Bend.start_angle = 0;
Bend.end_angle = 90;
bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 1

Bend.InPort = Bend.OutPort;
Bend.start_angle = 90;
Bend.end_angle = 180;
bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 2

//Draw Prop loss section waveguide middle part
WG.angle = 180;
WG.L = PL_mid;
WG.InPort = Bend.OutPort;
waveguide(pCell, layer[0], WG, WG.OutPort);

//Draw two bends to output waveguide part
Bend.InPort = WG.OutPort;
Bend.start_angle = 180;
Bend.end_angle = 270;
bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 1

Bend.InPort = Bend.OutPort;
Bend.start_angle = 270;
Bend.end_angle = 0;
bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 2

//Draw Prop loss section waveguide output part
WG.angle = 0;
WG.L = PL_out;
```

```
                WG.InPort = Bend.OutPort;
                waveguide(pCell, layer[0], WG, WG.OutPort);

                //Draw final two bends to chip output
                Bend.InPort = WG.OutPort;
                Bend.start_angle = 0;
                Bend.end_angle = 270;
                bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 1

                Bend.InPort = Bend.OutPort;
                Bend.start_angle = 270;
                Bend.end_angle = 180;
                bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 2

                //Draw length of waveguide to output taper:
                WG.L = (L_total - BendSep) / 2;
                WG.InPort = Bend.OutPort;
                WG.angle = 180;
                waveguide(pCell, layer[0], WG, WG.OutPort);

                //Draw Butt coupling output section:
                //WG.InPort = Bend.OutPort; //Set waveguide start point
                //WG.L = L_total - (Bend.OutPort.x - TAPin.OutPort.x); //Change   wg
length
                //waveguide(pCell, layer[0], WG, WG.OutPort);    //Draw waveguide


                //Draw final two bends to chip output
                Bend.InPort = WG.OutPort;
                Bend.start_angle = 180;
                Bend.end_angle = 270;
                bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 1

                Bend.InPort = Bend.OutPort;
                Bend.start_angle = 270;
                Bend.end_angle = 0;
                bend(pCell, layer[0], Bend, Bend.OutPort);      //bend 2



                TAPout.angle = 0;
                TAPout.InPort = Bend.OutPort;     //Set LHS taper staring point
                Taper(pCell, layer[0], TAPout, TAPout.OutPort); //Draw LHS Taper


                //Draw output grating coupler
                GC.InPort = TAPout.OutPort; //Set RHS grating start point.
                GC.angle = 180;
                GratingCouplerSpecial(pFile,   pCell,   layer,   GC,   GC.OutPort);
        //Draw RHS grating.


        }//end for



}//end BendLossGrat (end main)

int UPI_Entry_Point(void)
{
        LMacro_BindToMenuAndHotKey_v9_30("Tools", NULL /*hotkey*/,
```

```
            "PropLossGrat", "PropLossGrat", NULL /*hotkey category*/);
    return 1;
}
```

References

[1] L. Chrostowski and M. E. Hochberg, *Silicon photonics design*. .

[2] M. Boden and I. Miles, *Services and the knowledge-based economy*. Continuum, 2000.

[3] Burton-Jones and Alan, "Knowledge Capitalism: Business, Work, and Learning in the New Economy," *OUP Cat.*, 2001.

[4] "The History and Future of Internet Traffic." [Online]. Available: https://blogs.cisco.com/sp/the-history-and-future-of-internet-traffic. [Accessed: 30-Aug-2018].

[5] J. A. P. Clément and International Monetary Fund, *Postconflict economics in Sub-Saharan Africa : lessons from the Democratic Republic of the Congo*. .

[6] J. Coomer and T. Gstraunthaler, "The hyperinflaTion in Zimbabwe," 2011.

[7] G. E. Makinen, "The Greek Hyperinflation and Stabilization of 1943–1946," *J. Econ. Hist.*, vol. 46, no. 03, pp. 795–805, Sep. 1986.

[8] C. D. Campbell and G. C. Tullock, "Hyperinflation in China, 1937-49," *https://doi.org/10.1086/257516*, Oct. 2015.

[9] D. J. Richardson, "Applied physics. Filling the light pipe.," *Science*, vol. 330, no. 6002, pp. 327–8, Oct. 2010.

[10] M. A. Shampo, R. A. Kyle, and D. P. Steensma, "Charles K. Kao—father of fiber optics.," *Mayo Clin. Proc.*, vol. 86, no. 8, p. e45, Aug. 2011.

[11] J. Hecht, *Understanding fiber optics* . *Prentice Hall; 5 edition , 2005*

[12] R. J. Mears, L. Reekie, I. M. Jauncey, and D. N. Payne, "Low-noise erbium-doped fibre amplifier operating at 1.54μm," *Electronics Letters*, vol. 23, no. 19. IET Digital Library, pp. 1026–1028, 10-Sep-1987.

[13] C. G. Atkins, J. F. Massicott, J. R. Armitage, R. Wyatt, B. J. Ajnslie, and S. P. Craig-Ryan, "High-gain, broad spectral bandwidth erbium-doped fibre amplifier pumped near 1.5μm," *Electron. Lett.*, vol. 25, no. 14, p. 910, Jul. 1989.

[14] Y. Sun *et al.*, "80 nm ultra-wideband erbium-doped silica fibre amplifier," *Electron. Lett.*, vol. 33, no. 23, p. 1965, Nov. 1997.

[15] G. T. Reed, *Silicon Photonics: The State of the Art*. John Wiley & Sons, 2008.

[16] M. Hochberg *et al.*, "Silicon Photonics: The Next Fabless Semiconductor Industry," *IEEE Solid-State Circuits Mag.*, vol. 5, no. 1, pp. 48–58, 2013.

[17] M. Ziebell *et al.*, "40 Gbit/s low-loss silicon optical modulator based on a pipin diode.," *Opt. Express*, vol. 20, no. 10, pp. 10591–6, May 2012.

[18] F. Y. Gardes, D. J. Thomson, N. G. Emerson, and G. T. Reed, "40 Gb/s silicon photonics modulator for TE and TM polarisations," *Opt. Express*, vol. 19, no. 12, p. 11804, Jun. 2011.

[19] S. Xiao, M. H. Khan, H. Shen, and M. Qi, "Multiple-channel silicon micro-resonator based filters for WDM applications," *Opt. Express*, vol. 15, no. 12, p. 7489, 2007.

[20] F. Xia, M. Rooks, L. Sekaric, and Y. Vlasov, "Ultra-compact high order ring resonator filters using submicron silicon photonic wires for on-chip optical interconnects," *Opt. Express*, vol. 15, no. 19, p. 11934, 2007.

[21] C. Gunn, A. Narasimha, B. Analui, Y. Liang, and T. J. Sleboda, "A 40Gb silicon photonics tranceiver," in *Integrated Optoelectronic Devices 2007*, 2007, pp. 64770N-64770N–8.

[22] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, "Silicon optical modulators," *Nat. Photonics*, vol. 4, no. 8, pp. 518–526, Jul. 2010.

[23] T. Baehr-Jones *et al.*, "A 25 Gb/s Silicon Photonics Platform," Mar. 2012.

[24] D. Thomson *et al.*, "Roadmap on silicon photonics," *J. Opt.*, vol. 18, no. 7, p. 073003, Jul. 2016.

[25] A. Khanna *et al.*, "ePIXfab: the silicon photonics platform," 2013, p. 87670H.

[26] M. Hochberg and T. Baehr-Jones, "Towards fabless silicon photonics," *Nat. Photonics 2010 48*, Aug. 2010.

[27] "CORNERSTONE Home." [Online]. Available: http://www.cornerstone.sotonfab.co.uk/. [Accessed: 04-Apr-2018].

[28] L. Carroll *et al.*, "Photonic Packaging: Transforming Silicon Photonic Integrated Circuits into Photonic Devices," *Appl. Sci.*, vol. 6, no. 12, p. 426, 2016.

[29] "Institute of Microelectronics (IME) > SERVICES > MULTI-PROJECT WAFER (MPW) SERVICES." [Online]. Available: https://www.a-star.edu.sg/ime/SERVICES/MULTI-PROJECT-WAFER-MPW-SERVICES. [Accessed: 04-Apr-2018].

[30] "The MOSIS Service: Vendors : AIM : VIEW." [Online]. Available: https://www.mosis.com/vendors/view/AIM. [Accessed: 04-Apr-2018].

[31] A. Mekis *et al.*, "A Grating-Coupler-Enabled CMOS Photonics Platform," *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 3, pp. 597–608, May 2011.

[32] G. T. Reed, M. M. Milosevic, X. Chen, and D. J. Thomson, "Trimming of ring resonators via ion implantation in silicon," 2017, vol. 10242, p. 102420Q.

[33] P. J. M. Suni, J. Bowers, L. Coldren, S. J. Ben Yoo, and L. Author, "Photonic Integrated Circuits for Coherent Lidar."

[34] C. V. Poulton *et al.*, "Coherent solid-state LIDAR with silicon photonic optical phased arrays," *Opt. Lett.*, vol. 42, no. 20, p. 4091, Oct. 2017.

[35] J. C. Hulme *et al.*, "Fully integrated hybrid silicon two dimensional beam scanner," *Opt. Express*, vol. 23, no. 5, p. 5861, Mar. 2015.

[36] M. Burla, L. R. Cortés, M. Li, X. Wang, L. Chrostowski, and J. Azaña, "Integrated waveguide Bragg gratings for microwave photonics signal processing," *Opt. Express*, vol. 21, no. 21, p. 25120, Oct. 2013.

[37] J. Wu *et al.*, “Passive silicon photonic devices for microwave photonic signal processing,” *Opt. Commun.*, vol. 373, pp. 44–52, Aug. 2016.

[38] L. R. Chen, “Silicon Photonics for Microwave Photonics Applications,” *J. Light. Technol. Vol. 35, Issue 4, pp. 824-835*, vol. 35, no. 4, pp. 824–835, Feb. 2017.

[39] D. Hill, “Ultrahigh sensitivity slot-waveguide biosensor on a highly integrated chip for simultaneous diagnosis of multiple diseases,” in *2008 IEEE/LEOS International Conference on Optical MEMs and Nanophotonics*, 2008, pp. 52–53.

[40] T. Claes, J. G. Molera, K. De Vos, E. Schacht, R. Baets, and P. Bienstman, “Label-Free Biosensing With a Slot-Waveguide-Based Ring Resonator in Silicon on Insulator,” *IEEE Photonics J.*, vol. 1, no. 3, pp. 197–204, Sep. 2009.

[41] C. L. Arce, E. Hallynck, S. Werquin, J. W. Hoste, D. Martens, and P. Bienstman, “Silicon photonics biosensing: different packaging platforms and applications,” in *SPIE BiOS*, 2015, p. 932006.

[42] M. Iqbal *et al.*, “Label-Free Biosensor Arrays Based on Silicon Ring Resonators and High-Speed Optical Scanning Instrumentation,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 16, no. 3, pp. 654–661, 2010.

[43] J. Hu, X. Sun, A. Agarwal, and L. C. Kimerling, “Design guidelines for optical resonator biochemical sensors,” *J. Opt. Soc. Am. B*, vol. 26, no. 5, p. 1032, May 2009.

[44] M. Á. Guillen-Torres, E. Cretu, N. A. F. Jaeger, and L. Chrostowski, “Ring Resonator Optical Gyroscopes—Parameter Optimization and Robustness Analysis,” *J. Light. Technol.*, vol. 30, no. 12, pp. 1802–1817, Jun. 2012.

[45] Y. Chaojun, W. Junli, Li Haijun, and J. Binghua, “Design of Passive Components in Integrated Optic Chip on SOI Applied in Fiber Optic Gyroscope System,” in *2007 8th International Conference on Electronic Measurement and Instruments*, 2007, pp. 2-96-2–99.

[46] M. S. Eggleston *et al.*, “Silicon photonics enabled hyper-wideband wireless

communication link," in *2017 IEEE MTT-S International Microwave Symposium (IMS)*, 2017, pp. 431–434.

[47] M. Ko, J.-S. Youn, M.-J. Lee, K.-C. Choi, H. Rucker, and W.-Y. Choi, "Silicon Photonics-Wireless Interface IC for 60-GHz Wireless Link," *IEEE Photonics Technol. Lett.*, vol. 24, no. 13, pp. 1112–1114, Jul. 2012.

[48] M. Ko, M.-J. Lee, H. Rücker, and W.-Y. Choi, "Silicon photonics-wireless interface ICs for micro-/millimeter-wave fiber-wireless networks," *Opt. Express*, vol. 21, no. 19, p. 22962, Sep. 2013.

[49] E. Lacombe *et al.*, "10-Gb/s Indoor THz Communications Using Industrial Si Photonics Technology," *IEEE Microw. Wirel. Components Lett.*, pp. 1–3, 2018.

[50] M. A. Foster, A. C. Turner, M. Lipson, and A. L. Gaeta, "Nonlinear optics in photonic nanowires," *Opt. Express*, vol. 16, no. 2, p. 1300, 2008.

[51] Y. Okawachi, A. L. Gaeta, and M. Lipson, "Breakthroughs in Nonlinear Silicon Photonics 2011," *IEEE Photonics J.*, vol. 4, no. 2, pp. 601–606, Apr. 2012.

[52] S. Lefrancois, C. A. Husko, A. Blanco-Redondo, and B. J. Eggleton, "Nonlinear Silicon Photonics and the Moment Method," in *CLEO: 2015*, 2015, p. FW1D.4.

[53] C. R. Doerr *et al.*, "Silicon Photonics Coherent Transceiver in a Ball-Grid Array Package," in *Optical Fiber Communication Conference Postdeadline Papers*, 2017, p. Th5D.5.

[54] K. Kikuchi *et al.*, "Silicon-photonics-based coherent optical subassembly (COSA) for ultra-compact coherent transceiver," in *2017 IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS)*, 2017, pp. 1–3.

[55] C. R. Doerr *et al.*, "O, E, S, C, and L Band Silicon Photonics Coherent Modulator/Receiver," in *Optical Fiber Communication Conference Postdeadline Papers*, 2016, p. Th5C.4.

[56] C. R. Doerr *et al.*, "Packaged Monolithic Silicon 112-Gb/s Coherent Receiver," *IEEE Photonics Technol. Lett.*, vol. 23, no. 12, pp. 762–764, Jun. 2011.

[57] T. Barwicz *et al.*, "A Novel Approach to Photonic Packaging Leveraging Existing High-Throughput Microelectronic Facilities," *IEEE J. Sel. Top. Quantum Electron.*, vol. 22, no. 6, pp. 455–466, Nov. 2016.

[58] D. J. Lockwood and L. Pavesi, *Silicon Photonics II: Components and Integration*. Springer Science & Business Media, 2010.

[59] B. W. Snyder and P. A. O'Brien, "Developments in packaging and integration for silicon photonics," in *SPIE MOEMS-MEMS*, 2013, p. 86140D.

[60] A. Ghatak, *Optics*. Tata McGraw-Hill Education, 2005.

[61] J. Tyndall, *Notes of a Course of Nine Lectures on Light Delivered at the Royal Institution of Great Britain April 8-June 3, 1869*. Longmans, 1870.

[62] M. H. Weik, *Fiber Optics Standard Dictionary*. Springer, 2013.

[63] "Index of Refraction." [Online]. Available: http://hyperphysics.phy-astr.gsu.edu/hbase/Tables/indrf.html. [Accessed: 23-Mar-2018].

[64] G. T. Reed and A. P. Kights, *Silicon Photonics: An Introduction*. 2004.

[65] C. Incorporated, "Corning ® SMF- 28 ® Ultra Optical Fiber," no. November, 2014.

[66] R. G. Hunsperger, *Integrated optics, theory and technology*. Springer-Verlag, 1982.

[67] R. G. Hunsperger, "Theory of Optical Waveguides," 1982, pp. 31–46.

[68] J. Soler Penadés *et al.*, "Suspended SOI waveguide with sub-wavelength grating cladding for mid-infrared," *Opt. Lett.*, vol. 39, no. 19, p. 5661, Oct. 2014.

[69] J. Soler Penades, A. Khokhar, M. Nedeljkovic, and G. Mashanovich, "Low Loss Mid-Infrared SOI Slot Waveguides," *IEEE Photonics Technol. Lett.*, pp. 1–1, 2015.

[70] B. D. Timotijevic *et al.*, "Tailoring the response and temperature characteristics of multiple serial-coupled resonators in silicon on insulator,"

2007, vol. 6477, p. 64770B.

[71] W. Song, M. Chahal, G. K. Celler, Y. Jaluria, G. T. Reed, and W. Jiang, "The influence of substrate on SOI photonic crystal thermo-optic devices," *Opt. Express*, vol. 21, no. 4, p. 4235, Feb. 2013.

[72] D. J. Thomson *et al.*, "50-Gb/s Silicon Optical Modulator," *IEEE Photonics Technol. Lett.*, vol. 24, no. 4, pp. 234–236, Feb. 2012.

[73] R. V. R. Choudhury, "Simulation of varied Si waveguides to study FCA effect in TPA based optical logic gates," in *2015 International Conference on Control Communication & Computing India (ICCC)*, 2015, pp. 442–447.

[74] R. A. Soref, J. Schmidtchen, and K. Petermann, "Large single-mode rib waveguides in GeSi-Si and Si-on-SiO/sub 2/," *IEEE J. Quantum Electron.*, vol. 27, no. 8, pp. 1971–1974, 1991.

[75] Y. Vlasov and S. McNab, "Losses in single-mode silicon-on-insulator strip waveguides and bends.," *Opt. Express*, vol. 12, no. 8, pp. 1622–31, Apr. 2004.

[76] U. Fischer, T. Zinke, J.-R. Kropp, F. Arndt, and K. Petermann, "0.1 dB/cm waveguide losses in single-mode SOI rib waveguides," *IEEE Photonics Technol. Lett.*, vol. 8, no. 5, pp. 647–648, May 1996.

[77] R. L. Espinola, J. I. Dadap, R. M. Osgood, Jr., S. J. McNab, and Y. A. Vlasov, "C-band wavelength conversion in silicon photonic wire waveguides," *Opt. Express*, vol. 13, no. 11, p. 4341, May 2005.

[78] B. J. Frey, D. B. Leviton, and T. J. Madison, "Temperature dependent refractive index of silicon and germanium," 2006, pp. 62732J-62732J–10.

[79] M. Kawachi, "Silica waveguides on silicon and their application to integrated-optic components," *Opt. Quantum Electron.*, vol. 22, no. 5, pp. 391–416, Sep. 1990.

[80] H. Bu, "5 nanometer transistors inching their way into chips," 2017. [Online]. Available: https://www.ibm.com/blogs/think/2017/06/5-nanometer-

transistors/.

[81] D. Taillaert *et al.*, “Grating Couplers for Coupling between Optical Fibers and Nanophotonic Waveguides,” *Jpn. J. Appl. Phys.*, vol. 45, no. 8A, pp. 6071–6077, Aug. 2006.

[82] L. Zimmermann, T. Tekin, H. Schroeder, P. Dumon, and W. Bogaerts, “How to bring nanophotonics to application – silicon photonics packaging,” *IEEE LEOS Newsl.*, no. December, pp. 4–14, 2008.

[83] H. Fukuda, K. Yamada, T. Tsuchizawa, T. Watanabe, H. Shinojima, and S. Itabashi, “Silicon photonic circuit with polarization diversity,” *Opt. Express*, vol. 16, no. 7, p. 4872, Mar. 2008.

[84] S. Ura, T. Suhara, and H. Nishihara, “Aberration characterizations of a focusing grating coupler in an integrated-optic disk pickup device.,” *Appl. Opt.*, vol. 26, no. 22, pp. 4777–82, Nov. 1987.

[85] M. Miler and M. Skalsky, “Chirped and curved grating coupler focusing both outgoing beam and guided wave,” *Opt. Commun.*, vol. 33, no. 1, pp. 13–16, Apr. 1980.

[86] M. T. W. Ang *et al.*, “Grating couplers using silicon-on-insulator,” in *Optoelectronics ’99 - Integrated Optoelectronic Devices*, 1999, pp. 79–86.

[87] “SOI grating structure for perfectly vertical fiber coupling,” *Günther Roelkens Dries Van Thourhout Roel Baets*, no. 123.

[88] L. Zimmermann and T. Tekin, “G-Pack—A generic testbed package for Silicon photonics devices,” *Proc. Gr. IV …*, 2008.

[89] L. Zimmermann, L., “ePIXpack - advanced smart packaging solutions for silicon photonics,” in *14th European Conference on Integrated Optics, ECIO 08 Eindhoven. Proceedings. Contributed and and invited papers*, 2008, pp. 33–36.

[90] L. Zimmermann, G. B. Preve, T. Tekin, T. Rosin, and K. Landles, “Packaging and Assembly for Integrated Photonics—A Review of the ePIXpack Photonics

Packaging Platform," *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 3, pp. 645–651, May 2011.

[91] R. Marchetti *et al.*, "High-efficiency grating-couplers: demonstration of a new design strategy," *Sci. Rep.*, vol. 7, no. 1, p. 16670, Dec. 2017.

[92] P. De Dobbelaere *et al.*, "Demonstration of first WDM CMOS photonics transceiver with monolithically integrated photo-detectors," in *2008 34th European Conference on Optical Communication*, 2008, pp. 1–2.

[93] B. Snyder and P. O'Brien, "Planar Fiber Packaging Method for Silicon Photonic Integrated Circuits," *Opt. Fiber Commun. Conf.*, vol. 1, p. OM2E.5, 2012.

[94] Iiit. J. P. J. R.-L. J. W. Lawrence C. Gunn, "Fiber to chip coupler," Mar. 2004.

[95] B. W. Snyder and P. A. O'Brien, "Developments in packaging and integration for silicon photonics," in *SPIE MOEMS-MEMS*, 2013, p. 86140D.

[96] J. V. Galan, T. Tekin, G. B. Preve, A. Brimont, M. Llopis, and P. Sanchis, "Low profile silicon photonics packaging approach featuring configurable multiple electrical and optical connectivity," in *8th IEEE International Conference on Group IV Photonics*, 2011, pp. 377–379.

[97] J. V. Galan *et al.*, "CMOS compatible silicon etched V-grooves integrated with a SOI fiber coupling technique for enhancing fiber-to-chip alignment," in *2009 6th IEEE International Conference on Group IV Photonics*, 2009, pp. 148–150.

[98] T. Barwicz *et al.*, "Photonic Packaging in High-Throughput Microelectronic Assembly Lines for Cost-Efficiency and Scalability," no. c, pp. 4–6, 2015.

[99] P. Cheben *et al.*, "Broadband polarization independent nanophotonic coupler for silicon waveguides with ultra-high efficiency," *Opt. Express*, vol. 23, no. 17, p. 22553, Aug. 2015.

[100] T. Barwicz *et al.*, "An O-band Metamaterial Converter Interfacing Standard Optical Fibers to Silicon Nanophotonic Waveguides," in *Optical Fiber Communication Conference*, 2015, p. Th3F.3.

[101] P. Cheben *et al.*, "Refractive index engineering with subwavelength gratings for efficient microphotonic couplers and planar waveguide multiplexers," *Opt. Lett.*, vol. 35, no. 15, p. 2526, Aug. 2010.

[102] T. Barwicz *et al.*, "Automated, self-aligned assembly of 12 fibers per nanophotonic chip with standard microelectronics assembly tooling," in *2015 IEEE 65th Electronic Components and Technology Conference (ECTC)*, 2015, pp. 775–782.

[103] B. Snyder *et al.*, "Ultra-broadband, polarization-insensitive SMF-28 fiber edge couplers for silicon photonics," in *2017 IEEE CPMT Symposium Japan (ICSJ)*, 2017, pp. 55–58.

[104] J. Cardenas, C. B. Poitras, K. Luke, L. Luo, P. A. Morton, and M. Lipson, "High Coupling Efficiency Etched Facet Tapers in Silicon Waveguides," vol. 26, no. 23, pp. 2380–2382, 2014.

[105] M. S. Cohen, M. F. Cina, E. Bassous, M. M. Oprysko, and J. L. Speidell, "Passive laser-fiber alignment by index method," *IEEE Photonics Technol. Lett.*, vol. 3, no. 11, pp. 985–987, Nov. 1991.

[106] S. Bernabe *et al.*, "In-plane pigtailing of silicon photonics device using 'semi-passive' strategies," in *2012 4th Electronic System-Integration Technology Conference*, 2012, pp. 1–6.

[107] C. Kopp *et al.*, "Silicon Photonic Circuits: On-CMOS Integration, Fiber Optical Coupling, and Packaging," *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 3, pp. 498–509, May 2011.

[108] C. R. Cole, "100-Gb/s and beyond transceiver technologies," *Opt. Fiber Technol.*, vol. 17, no. 5, pp. 472–479, Oct. 2011.

[109] T. Baehr-Jones, T. Pinguet, P. Lo Guo-Qiang, S. Danziger, D. Prather, and M. Hochberg, "Myths and rumours of silicon photonics," *Nat. Photonics*, vol. 6, no. 4, pp. 206–208, Apr. 2012.

[110] R. Rahmani, I. Moser, and M. Seyedmahmoudian, "A Complete Model for

Modular Simulation of Data Centre Power Load," Mar. 2018.

[111] C. Scarcella *et al.*, "Pluggable Single-Mode Fiber-Array-to-PIC Coupling Using Micro-Lenses," *IEEE Photonics Technol. Lett.*, vol. 29, no. 22, pp. 1943–1946, Nov. 2017.

[112] P. E. Morrissey, P. O'Brien, L. Carroll, and K. Gradkowski, "Packaging of silicon photonic devices: from prototypes to production," in *Silicon Photonics XIII*, 2018, vol. 10537, p. 20.

[113] D. J. Thomson, "Apparatus comprising at least one optical device optically coupled to at least one waveguide on an optical chip." 12-Jun-2014.

[114] C. Zhang *et al.*, "High Efficiency Grating Coupler for Coupling between Single-Mode Fiber and SOI Waveguides," *Chinese Phys. Lett.*, vol. 30, no. 1, p. 014207, Jan. 2013.

[115] D. Vermeulen *et al.*, "High-efficiency fiber-to-chip grating couplers realized using an advanced platform," vol. 18, no. 17, pp. 18278–18283, 2010.

[116] D. D. Thomson, "Passive Alignment Silicon Photonics." .

[117] J. . A. M. Kowalevicz and F. Bucholtz, "Beam Divergence from an SMF-28 Optical Fiber," Oct. 2006.

[118] S. Bernabe, C. Kopp, L. Lombard, and J.-M. Fedeli, "Microelectronic-like packaging for silicon photonics : A 10 Gbps multi-chip-module optical receiver based on Ge-on-Si photodiode," in *3rd Electronics System Integration Technology Conference ESTC*, 2010, pp. 1–5.

[119] Haoshuo Chen *et al.*, "Packaged mode multiplexer based on silicon photonics." pp. 1–3, 2012.

[120] L. Zimmermann, G. B. Preve, T. Tekin, and T. Rosin, "Packaging and assembly for integrated photonics - the ePIXpack photonics packaging service," in *2010 IEEE Photinic Society's 23rd Annual Meeting*, 2010, pp. 167–168.

[121] S. Reboh, P. Morin, M. J. Hÿtch, F. Houdellier, and A. Claverie, "Mechanics of

silicon nitride thin-film stressors on a transistor-like geometry," *APL Mater.*, vol. 1, no. 4, p. 042117, Oct. 2013.

[122] E. Cianci, F. Pirola, and V. Foglietti, "Analysis of stress and composition of silicon nitride thin films deposited by electron cyclotron resonance plasma-enhanced chemical vapor deposition for microfabrication processes," *J. Vac. Sci. Technol. B Microelectron. Nanom. Struct.*, vol. 23, no. 1, p. 168, Jan. 2005.

[123] E. H. Conradie and D. F. Moore, "SU-8 thick photoresist processing as a functional material for MEMS applications," *J. Micromechanics Microengineering*, vol. 12, no. 4, p. 303, Jul. 2002.

[124] A. Schneider, R. H. Ibbotson, R. J. Dunn, and E. Huq, "Arrays of SU-8 microcantilevers with integrated piezoresistive sensors for parallel AFM applications," *Microelectron. Eng.*, vol. 88, no. 8, pp. 2390–2393, Aug. 2011.

[125] J. Hammacher, A. Fuelle, J. Flaemig, J. Saupe, B. Loechel, and J. Grimm, "Stress engineering and mechanical properties of SU-8-layers for mechanical applications," *Microsyst. Technol.*, vol. 14, no. 9–11, pp. 1515–1523, Oct. 2008.

[126] I. Bajwa, "KOH etching of (100) Si wafer, No 1 KOH etching of (100) Si wafer, No 1 Keywords KOH etching Silicon Technical Report (Graduate Student Fellow Program)," Membrane Science Commons, 2016.

[127] D. Resnik *et al.*, "Micropyramidal hillocks on KOH etched {100} silicon surfaces: formation, prevention and removal Related content Effective roughness reduction of {100} and {311} planes in anisotropic etching of {100} silicon in 5% TMAH Micropyramidal hillocks on KOH etched {100} silicon surfaces: formation, prevention and removal," 1999.

[128] A. V Krishnamoorthy *et al.*, "Optical Proximity Communication With Passively Aligned Silicon Photonic Chips," vol. 45, no. 4, pp. 409–414, 2009.

[129] Q. X. Zhang *et al.*, "A Lens Holder in Conjunction With a MEMS Platform for On-Chip Aligning and Fixing of a Ball Lens in Silicon Photonics Packaging," *IEEE Photonics Technol. Lett.*, vol. 23, no. 9, pp. 588–590, May 2011.

[130] M. Wood, P. Sun, and R. M. Reano, “Compact cantilever couplers for low-loss fiber coupling to silicon photonic integrated circuits,” *Opt. Express*, vol. 20, no. 1, p. 164, Jan. 2012.

[131] D. J. Richardson *et al.*, “Hollow Core Fibres and their Applications,” in *Optical Fiber Communication Conference*, 2017, p. Tu3H.1.

[132] Z. Fang, Q. Y. Chen, and C. Z. Zhao, “A review of recent progress in lasers on silicon,” *Opt. Laser Technol.*, vol. 46, pp. 103–110, Mar. 2013.

[133] L. Pavesi, “A Review of the various approaches to a silicon laser,” in *Integrated Optoelectronics Devices*, 2003, pp. 206–220.

[134] H. Rong *et al.*, “An all-silicon Raman laser.,” *Nature*, vol. 433, no. 7023, pp. 292–4, Jan. 2005.

[135] G. Roelkens *et al.*, “III-V/silicon photonics for on-chip and intra-chip optical interconnects,” *Laser Photon. Rev.*, vol. 4, no. 6, pp. 751–779, Nov. 2010.

[136] K. Tanabe, K. Watanabe, and Y. Arakawa, “III-V/Si hybrid photonic devices by direct fusion bonding.,” *Sci. Rep.*, vol. 2, p. 349, Jan. 2012.

[137] R. Jones *et al.*, “Integrated Hybrid Lasers and Amplifiers on a Silicon Platform,” *OFC/NFOEC 2008 - 2008 Conf. Opt. Fiber Commun. Fiber Opt. Eng. Conf.*, vol. 1, pp. 1–3, Feb. 2008.

[138] Y. De Koninck, G. Roelkens, and R. Baets, “Design of a Hybrid III–V-on-Silicon Microlaser With Resonant Cavity Mirrors,” *IEEE Photonics J.*, vol. 5, no. 2, pp. 2700413–2700413, Apr. 2013.

[139] G. Roelkens *et al.*, “III-V/Si photonics by die-to-wafer bonding,” *Mater. Today*, vol. 10, no. 7–8, pp. 36–43, Jul. 2007.

[140] G. Roelkens, D. Van Thourhout, and R. Baets, “Ultra-thin benzocyclobutene bonding of III–V dies onto SOI substrate,” *Electron. Lett.*, vol. 41, no. 9, pp. 4–5, 2006.

[141] S. Keyvaninia, M. Muneeb, S. Stanković, P. J. Van Veldhoven, D. Van Thourhout,

and G. Roelkens, “Ultra-thin DVS-BCB adhesive bonding of III-V wafers, dies and multiple dies to a patterned silicon-on-insulator substrate,” *Opt. Mater. Express*, vol. 3, no. 1, p. 35, Dec. 2012.

[142] S. Keyvaninia *et al.*, “III-V-on-silicon multi-frequency lasers.,” *Opt. Express*, vol. 21, no. 11, pp. 13675–83, Jun. 2013.

[143] I. Christiaens and G. Roelkens, “Adhesive wafer bonding with Benzocyclobutene,” *J. Light. …*, 2004.

[144] Z. Song, Z. Tan, L. Liu, and Z. Wang, “Void-free BCB adhesive wafer bonding with high alignment accuracy,” *Microsyst. Technol.*, Jun. 2014.

[145] Chien-Chih Liu, Yen-Kuang Lin, Mau-Phon Houng, and Yeong-Her Wang, “The microstructure investigation of flip-chip laser diode bonding on silicon substrate by using indium-gold solder,” *IEEE Trans. Components Packag. Technol.*, vol. 26, no. 3, pp. 635–641, Sep. 2003.

[146] K. Ohira *et al.*, “On-chip optical interconnection by using integrated III-V laser diode and photodetector with silicon waveguide.,” *Opt. Express*, vol. 18, no. 15, pp. 15440–7, Jul. 2010.

[147] A. V. Krishnamoorthy *et al.*, “16 x 16 VCSEL array flip-chip bonded to CMOS VLSI circuit,” *IEEE Photonics Technol. Lett.*, vol. 12, no. 8, pp. 1073–1075, Aug. 2000.

[148] S. K. Patra and Y. C. Lee, “Design of solder joints for self-aligned optoelectronic assemblies,” *IEEE Trans. Components, Packag. Manuf. Technol. Part B*, vol. 18, no. 3, pp. 543–551, 1995.

[149] Y. C. Lee, “Soldering technology for optoelectronic packaging,” in *1996 Proceedings 46th Electronic Components and Technology Conference*, 1996, pp. 26–36.

[150] E. Higurashi, T. Imamura, T. Suga, and R. Sawada, “Low-Temperature Bonding of Laser Diode Chips on Silicon Substrates Using Plasma Activation of Au Films,” *IEEE Photonics Technol. Lett.*, vol. 19, no. 24, pp. 1994–1996, Dec. 2007.

[151] D. F. G. Gallagher and T. P. Felici, “Eigenmode expansion methods for simulation of optical propagation in photonics: pros and cons,” 2003, vol. 4987, p. 69.

[152] W. Bogaerts, M. Fiers, and P. Dumon, “Design Challenges in Silicon Photonics,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 20, no. 4, pp. 1–8, Jul. 2014.

[153] L. Vivien *et al.*, “Comparison between strip and rib SOI microwaveguides for intra-chip light distribution,” 2004.

[154] A. G. Rickman and G. T. Reed, “Silicon-on-insulator optical rib waveguides: loss, mode characteristics, bends and y-junctions,” *IEE Proc. - Optoelectron.*, vol. 141, no. 6, pp. 391–393, Dec. 1994.

[155] G. Roelkens, D. Vermeulen, S. Selvaraja, R. Halir, W. Bogaerts, and D. Van Thourhout, “Grating-Based Optical Fiber Interfaces for Silicon-on-Insulator Photonic Integrated Circuits,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 3, pp. 571–580, May 2011.

[156] C. Kopp *et al.*, “Silicon photonic circuits: On-CMOS integration, fiber optical coupling, and packaging,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, pp. 498–509, 2011.

[157] C. Kopp, K. Gilbert, P. Grosse, and J.-M. Fedeli, “Silicon photonics packaging : characterization of a waveguide grating coupler and modeling of the fiber coupling ratio,” 2007, vol. 6478, p. 64780N.

[158] S. Xiao, R. Vahldieck, and H. Jin, “Full-wave analysis of guided wave structures using a novel 2-D FDTD,” *IEEE Microw. Guid. Wave Lett.*, vol. 2, no. 5, pp. 165–167, May 1992.

[159] A. Scherer, J. Vučković, M. Lončar, and T. Doll, “Design and Fabrication of Silicon Photonic Crystal Optical Waveguides,” *J. Light. Technol. Vol. 18, Issue 10, pp. 1402-*, vol. 18, no. 10, p. 1402, Oct. 2000.

[160] X. Chen, C. Li, and C. Fung, “Apodized waveguide grating couplers for efficient coupling to optical fibers,” *Photonics Technol. …*, vol. 22, no. 15, pp. 1156–

1158, 2010.

[161] C. Li, H. Zhang, M. Yu, and G. Q. Lo, “CMOS-compatible high efficiency double-etched apodized waveguide grating coupler.,” *Opt. Express*, vol. 21, no. 7, pp. 7868–74, Apr. 2013.

[162] M. Antelius, K. B. Gylfason, and H. Sohlström, “An apodized SOI waveguide-to-fiber surface grating coupler for single lithography silicon photonics.,” *Opt. Express*, vol. 19, no. 4, pp. 3592–8, Feb. 2011.

[163] R. Poli, J. Kennedy, and T. Blackwell, “Particle swarm optimization,” *Swarm Intell.*, vol. 1, no. 1, pp. 33–57, Aug. 2007.

[164] C. Sammut and G. I. Webb, *Encyclopedia of Machine Learning*, vol. 33. Boston, MA: Springer US, 2010.

[165] C. Li, H. Zhang, M. Yu, and G. Lo, “CMOS-Compatible Silicon Double-etched Apodized Waveguide Grating Couplers for High Efficient Coupling,” *Opt. Fiber Commun. Conf. Fiber Opt. Eng. Conf. 2013*, p. OW3F.2, 2013.

[166] J. Miller and N. Williams, “Tanner L-Edit IC A Complete IC Physical Design Environment.”

[167] H. D. Young, *Fundamentals of mechanics and heat*. McGraw-Hill, 1973.

[168] H. R. Shanks, P. D. Maycock, P. H. Sidles, and G. C. Danielson, “Thermal Conductivity of Silicon from 300 to 1400°K,” *Phys. Rev.*, vol. 130, no. 5, pp. 1743–1748, Jun. 1963.

[169] J. Komma, C. Schwarz, G. Hofmann, D. Heinert, and R. Nawrodt, “Thermo-optic coefficient of silicon at 1550 nm and cryogenic temperatures,” *Appl. Phys. Lett.*, vol. 101, no. 4, p. 041905, Jul. 2012.

[170] M. Han and A. Wang, “Temperature compensation of optical microresonators using a surface layer with negative thermo-optic coefficient.,” *Opt. Lett.*, vol. 32, no. 13, pp. 1800–2, Jul. 2007.

[171] N. N. Klimov, S. Mittal, M. Berger, and Z. Ahmed, “On-chip silicon waveguide

Bragg grating photonic temperature sensor," *Opt. Lett.*, vol. 40, no. 17, p. 3934, Sep. 2015.

[172] "3-Axis NanoMax Flexure Stages." [Online]. Available: https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=2386. [Accessed: 18-May-2018].

[173] "Optical Table and Active Isolator Leg Bundles." [Online]. Available: https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=5930. [Accessed: 18-May-2018].

[174] "81940A Compact Tunable Laser Source with Continuous Sweep Mode, 1520nm to 1630nm [Discontinued] | Keysight (formerly Agilent's Electronic Measurement)." [Online]. Available: https://www.keysight.com/en/pd-1281761-pn-81940A/compact-tunable-laser-source-with-continuous-sweep-mode-1520nm-to-1630nm?cc=GB&lc=eng. [Accessed: 18-May-2018].

[175] "81634B Low Polarization Dependence Optical Power Sensor | Keysight (formerly Agilent's Electronic Measurement)." [Online]. Available: https://www.keysight.com/en/pd-1279686-pn-81634B/low-polarization-dependence-optical-power-sensor?cc=GB&lc=eng&lsrch=true&searchT=81634B. [Accessed: 18-May-2018].

[176] R. Ulrich, S. C. Rashleigh, and W. Eickhoff, "Bending-induced birefringence in single-mode fibers," *Opt. Lett.*, vol. 5, no. 6, p. 273, Jun. 1980.

[177] R. Ulrich and A. Simon, "Polarization optics of twisted single-mode fibers," *Appl. Opt.*, vol. 18, no. 13, p. 2241, Jul. 1979.

[178] Thorlabs, "Fiber Paddle Controllers: Achieving Distinct Polarization States."

[179] "Smart Cut™ technology, Smart Choice - Soitec." [Online]. Available: https://www.soitec.com/en/products/smart-cut. [Accessed: 06-Jul-2018].

[180] *Zeitschrift für physikalische Chemie.* .

[181] M. Metzler and R. Patel, “Plasma Enhanced Chemical Vapor Deposition (PECVD) of Silicon Dioxide (SiO2) Using Oxford Instruments System 100 PECVD Plasma Enhanced Chemical Vapor Deposition (PECVD) of Silicon Dioxide (SiO2) Using Oxford Instruments System 100 PECVD Plasma Enhance Chemical Vapor Deposition of Silicon Dioxide (SiO2) Oxford PlasmaLab 100 PECVD.”

[182] “MICROPOSIT S1800 SERIES PHOTORESISTS.” [Online] Available: https://amolf.nl/wp-content/uploads/2016/09/datasheets_S1800.pdf [Accessed: 22-Mar-2018]

[183] D. Center Furukawa Sogo Bldg and K. Kawasaki-shi, “ZEON CORPORATION Specialty Materials Division Technical Report Technical Report ZEP520A ZEP520A High Resolution Positive Electron Beam Resist.”

[184] "AZ®nLOF™2000" [Online] https://www.microchemicals.com/micro/tds_az_nlof2000_series.pdf [Accessed: 22-Mar-2018]

[185] S. Franssila, *Introduction to Microfabrication*. Chichester, UK: John Wiley & Sons, Ltd, 2010.

[186] H. Jansen, H. Gardeniers, M. de Boer, M. Elwenspoek, and J. Fluitman, “A survey on the reactive ion etching of silicon in microtechnology,” *J. Micromechanics Microengineering*, vol. 6, no. 1, pp. 14–28, Mar. 1996.

[187] M. Wu and W. Fang, “Design and fabrication of MEMS devices using the integration of MUMPs, trench-refilled molding, DRIE and bulk silicon etching processes,” *J. Micromechanics Microengineering*, vol. 15, no. 3, pp. 535–542, Mar. 2005.

[188] C. M. Waits, B. Morgan, M. Kastantin, and R. Ghodssi, “Microfabrication of 3D silicon MEMS structures using gray-scale lithography and deep reactive ion etching,” *Sensors Actuators A Phys.*, vol. 119, no. 1, pp. 245–253, Mar. 2005.

[189] E. H. Klaassen *et al.*, “Silicon fusion bonding and deep reactive ion etching: a

new technology for microstructures," *Sensors Actuators A Phys.*, vol. 52, no. 1–3, pp. 132–139, Mar. 1996.

[190] F. Laermer and A. Schilp, "Method of anisotropically etching silicon," Mar. 1996.

[191] N. Maluf and K. Williams, *Introduction to microelectromechanical systems engineering*. Artech House, 2004.

[192] M. J. Madou, *Fundamentals of microfabrication : the science of miniaturization*. CRC Press, 2002.

[193] "Plasma-Therm: VERSALINE DSE." [Online]. Available: http://www.plasma-therm.com/versaline-dse.html. [Accessed: 22-Mar-2018].

[194] S. Leopold, C. Kremin, A. Ulbrich, S. Krischok, and M. Hoffmann, "Formation of silicon grass: Nanomasking by carbon clusters in cyclic deep reactive ion etching," *J. Vac. Sci. Technol. B, Nanotechnol. Microelectron. Mater. Process. Meas. Phenom.*, vol. 29, no. 1, p. 011002, Jan. 2011.

[195] P. Dixit and J. Miao, "Effect of SF_6 flow rate on the etched surface profile and bottom grass formation in deep reactive ion etching process," *J. Phys. Conf. Ser.*, vol. 34, no. 1, pp. 577–582, Apr. 2006.

[196] A. J. Knobloch, M. Wasilik, C. Fernandez-Pello, and A. P. Pisano, "Micro, Internal-Combustion Engine Fabrication With 900 [micro sign]m Deep Features Via DRIE," in *Microelectromechanical Systems*, 2003, vol. 2003, pp. 115–123.

[197] M. Vangbo and Y. Bäcklund, "Precise mask alignment to the crystallographic orientation of silicon wafers using wet anisotropic etching," *J. Micromechanics Microengineering*, vol. 6, no. 2, pp. 279–284, Jun. 1996.

[198] A. M. Agarwal, L. Liao, J. S. Foresi, M. R. Black, X. Duan, and L. C. Kimerling, "Low-loss polycrystalline silicon waveguides for silicon photonics," *J. Appl. Phys.*, vol. 80, no. 11, p. 6120, Jun. 1998.

[199] C. B. Nielsen, C. Christensen, C. Pedersen, and E. V. Thomsen, "Particle precipitation in connection with KOH etching of silicon," *J. Electrochem. Soc.*, vol. 151, no. 5, 2004.

[200] K. A. Reinhardt and W. Kern, *Handbook of silicon wafer cleaning technology*. .

[201] J. Johnson *et al.*, "Removing Surface Contaminates from Silicon Wafers to Facilitate EUV Optical Characterization," *Am. Phys. Soc. Four Corners Meet. 2003, Oct. 24-25, 2003 Arizona State Univ. Tempe, Arizona, Meet. ID 4CF03, Abstr. id. S2.008*, 2003.

[202] A. Donko, M. Beresna, Y. Jung, J. Hayes, D. Richardson, and G. Brambilla, "Point-by-point inscription of Bragg gratings in a multicore fibre."

[203] "ISO 14644-1:2015." [Online]. Available: https://www.iso.org/standard/53394.html.

[204] W. Whyte, *Cleanroom technology : fundamentals of design, testing and operation*. John Wiley & Sons, 2010.

[205] I. Anteney, "GP-002 – ECS Cleanrooms Induction Table of Contents," no. January 2014, pp. 1–16, 2015.

[206] B. LaFontaine, "Lasers and Moore's Law," *SPIE Prof.*, 2010.

[207] EVG, "EVG®620 Automated Mask Alignment System." [Online]. Available: https://www.evgroup.com/en/products/lithography/photolithography/mask_aligners/evg620semiauto/?SelectedTab=1. [Accessed: 21-Mar-2018].

[208] E. Anderson, "Double exposure makes dense high-resolution diffractive optics," *SPIE Newsroom*, 2007.

[209] Oxford Instruments, "Plasmalab System 100." [Online]. Available: http://www.oxfordplasma.de/systems/100ll.htm.

[210] W. Durant and A. Durant, *The story of civilization*. MJF Books, 1992.

[211] M. J. Madou, *Fundamentals of microfabrication : the science of miniaturization*.

CRC Press, 2002.

[212] E. P. (Ernest P. DeGarmo, J. T. Black, and R. A. Kohser, *Degarmo's materials and processes in manufacturing*. .

[213] D. C. Harris, *Quantitative chemical analysis*. W.H. Freeman and Co, 2010.

[214] Bong-Hwan Kim *et al.*, "MEMS fabrication of high aspect ratio track-following microactuator for hard disk drive using silicon on insulator," in *Technical Digest. IEEE International MEMS 99 Conference. Twelfth IEEE International Conference on Micro Electro Mechanical Systems (Cat. No.99CH36291)*, 1999, pp. 53–56.

[215] D. Seo, J. S. Bae, E. Oh, S. Kim, and S. Lim, "Selective wet etching of Si3N4/SiO2 in phosphoric acid with the addition of fluoride and silicic compounds," *Microelectron. Eng.*, vol. 118, pp. 66–71, Apr. 2014.

[216] S. M. Sze, *VLSI technology*. McGraw-Hill, 1988.

[217] A. Rothen, "The Ellipsometer, an Apparatus to Measure Thicknesses of Thin Surface Films," *Rev. Sci. Instrum.*, vol. 16, no. 2, pp. 26–30, Feb. 1945.

[218] "What is Ellipsometry? - J.A. Woollam." [Online]. Available: https://www.jawoollam.com/resources/ellipsometry-tutorial/what-is-ellipsometry. [Accessed: 22-Mar-2018].

[219] L. Reimer, *Scanning Electron Microscopy*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1985.

Publications

-Matthew T. Posner, Pearl V. John, Deanna Standen, Natalie V. Wheeler, Lieke D. van Putten, Nathan Soper, Tina L. Parsonage, Nicholas H. L. Wong, Gilberto Brambilla, "Reflecting photonics: reaching new audiences through new partnerships – IYL 2015 and the Royal Horticultural Society Flower Show," Proc. SPIE 9946, Optics Education and Outreach IV, 994603 (27 September 2016);

-G. T. Reed, R. Topley, A. Z. Khokhar, D. J. Thomson, S. Stanković, X. Chen, S. Reynolds, N. Soper, C. J. Mitchell, Y. Hu, G. Martinez-Jimenez, N. Healy, S. Mailis, A. C. Peacock, M. Nedeljkovic, J. Soler Penades, F. Y. Gardes and G. Z. Mashanovich, "Silicon photonics: some remaining challenges," Photonics West 2016, San Francisco, California, USA, 13-18 February 2016 (invited)

-N. Soper, S. Reynolds, D. Thomson, G. Reed, G. Mashanovich,"*Packaging Solutions for Silicon Photonics: Past attempts, current progress and beyond"*. IEEE CRALT 2016, Bangalore, India, 23rd August 2016

-M. T. Posner, A. Jantzen, L. D. van Putten, A. Ravagli, A. L. Donko, N. Soper, N. H. L. Wong, and P. V. John, "Cathedral outreach: student-led workshops for school curriculum enhancement in non-traditional environments," in *ETOP 2017 Proceedings*, X. Liu and X. Zhang, eds., (Optical Society of America, 2017), paper 1045207.

-N. Soper, S. Reynolds, D. J. Thomson, I. Sari, G. Z. Mashanovich, G. T. Reed, "Fabrication techniques for a Passively Aligned Silicon Photonics Package" Photonex 2017, Coventry, UK, 11th October