VERTICALLY-INTEGRATED PHOTONIC DEVICES IN SOI

Vertically-Integrated Photonic Devices in Silicon-on-Insulator

by

CHRISTOPHER BROOKS

B.Sc. (McMaster) 2002, M.A.Sc. (McMaster) 2005

A Thesis
Submitted to the School of Graduate Studies
in Partial Fulfilment of the Requirements
for the Degree
Doctor of Philosophy

McMaster University

DOCTOR OF PHILOSOPHY (2010)
(Engineering Physics)

McMaster University
Hamilton, Ontario

TITLE: Vertically-Integrated Photonic Devices in Silicon-on-Insulator

AUTHOR: Christopher Brooks

SUPERVISORS: Dr. P.E. Jessop, Dr. A.P. Knights

NUMBER OF PAGES: xiii, 210

Abstract

The functional density of photonic integrated circuits can be significantly increased by stacking multiple waveguide layers. These vertically-integrated devices require optical couplers to switch light signals between their layers. In this thesis, optical coupling between two stacked silicon-on-insulator slab waveguides has been demonstrated with a coupling efficiency of 68±4 %, obtained with a coupler length of 3535 μm. The main advantage of using a silicon-based material system for photonic integrated circuits is its compatibility with existing electronics manufacturing processes, facilitating cost-effective fabrication and the monolithic integration of both photonics and electronics on a single device.

Coupling between more complex silicon-on-insulator waveguide structures with lateral confinement was then demonstrated. The coupling ratio between stacked silicon rib waveguides was measured to be 54±4 %, while ratios of 71±4 % and 93±4 % were obtained for stacked channel waveguide and multimode interferometer-based couplers respectively. The corresponding coupler lengths for these three designs were 572 μm, 690 μm and 241 μm respectively. The sensitivity of these couplers to the input wavelength and polarization state has also been evaluated. These vertically-integrated couplers, along with other structures, have been thoroughly simulated, including their tolerance to fabrication errors. Novel fabrication processes used to demonstrate coupling in proof-of-concept devices have been developed, including an in-house wafer bonding procedure.

Acknowledgements

I would like to thank my supervisors, Dr. Paul Jessop and Dr. Andy Knights, for their knowledge and guidance that was invaluable to the success of this project.

I would like to thank the technicians and staff of the Centre for Emerging Device Technologies at McMaster University for their help in fabricating and characterizing my devices, especially: Doug Bruce, Doris Stevanovic, Graham Pearson, Zhilin Peng, Peter Jonasson, Brad Robinson, Jacek Wojcik and Evgueni Chelomentsev. I would like to thank Todd Simpson at the University of Western Ontario Nanofabrication Laboratory for his help running their plasma etcher. I would also like to thank everyone in the Jessop and Knights groups, especially: Mike Bulk, Dylan Logan, Edgar Huante-Ceron, Jonathan Doylend and David Walters for their advice and help throughout the project. A special thanks goes to Jonathan and Edgar for helping with device simulations.

Last, and most importantly, I would like to thank my parents, my brother, and my sister, for their love and support.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

While reductions in feature scale have continued to improve the performance of both electronic and photonic integrated circuits, vertical integration achieved through the stacking of multiple device layers will eventually be required for further performance gains. In this chapter, the benefits of vertically-integrated photonic devices will be discussed along with the advantages of using silicon-based materials to manufacture them. A summary of the published research in the area of vertically-integrated optical couplers will be given, with the fabrication techniques used to realize them being presented as well.

1.1 Advantages of Vertically-Integrated Photonic Devices

The functional density of traditional single-layered photonic integrated circuits is limited by the amount of in-plane surface area available. While technological advances in materials and feature size scaling have steadily increased the number of components that can be manufactured on a typically-sized die, the density and overall performance of optoelectronic devices can be further enhanced if vertical integration is employed. The number of external optical and electrical connections that can be made can also be increased if multiple device layers are stacked together. Three-dimensional routing of optical signals can also prevent undesired optical cross-talk between components since they can be separated onto individual layers. This cross-talk is unavoidable if the components are too closely packed onto a single device layer in an attempt to maximize available device real estate. Since individual optical components generally require more space than transistor-based electronics, vertical integration is even more crucial for enhancing the functional density of photonic integrated circuits.

Many types of optoelectronic devices are often made with different materials, such as lasers, detectors and switches, along with the electronics that control them and their connecting waveguides. Stacking multiple layers of different materials and components would further enable the realization of a truly integrated system-on-a-chip, while maximizing device surface area. Certain vertically-integrated devices are also more compact than their in-plane counterparts. Vertical integration also facilitates the fabrication of novel devices, as well as novel signal routing schemes, that will be presented in this chapter. Stacking multiple layers of optoelectronic devices on a single chip is an efficient way to fabricate densely integrated components, thereby enhancing functionality and performance.

1.1.1 Analogues With Vertically-Integrated Electronics

The progress of photonic device research has generally followed behind that of electronics. This has also been the case with vertical integration, as many of the technologies employed, such as wafer bonding, were initially developed to realize 3-D electronics. The performance of electronic integrated circuits has steadily improved as progresses in semiconductor manufacturing technology have resulted in vanishingly small transistor sizes. The time its takes for the transistor to turn on or off, or the gate delay, decreases with shorter transistor gate lengths, thus increasing chip performance. While transistor delays have continued to decrease, the bandwidth of the interconnects joining them has remained basically the same. These interconnects can only handle so much speed and will soon become the limiting factor in chip performance.

Interconnect bandwidth can be improved by replacing conventional aluminum wires with a more conductive material such as copper. Yet, this does not address the sheer number and length of the interconnects used to connect the ever-increasing number of transistors fabricated on a single die. These long interconnects also consume power as unwanted capacitance, caused by all of these wires being packed together in such a small space. Vertical integration, achieved by stacking multiple layers of transistors, greatly reduces the overall length of these interconnects as many of them are now used to connect between the closely spaced layers. As is the case with photonic integrated circuits, electronics will also benefit from increased functional density and available external connections achieved through vertical integration. The current pace of performance gains can be continued even if device scaling reaches its limits.

Vertically-integrated field effect transistors have been demonstrated in numerous publications, including [1]. IBM has demonstrated electronics on multiple layers that were assembled using a wafer bonding and layer transfer technique that will be detailed later in this chapter. Once the layers are stacked, vias are made between the layers and filled with either aluminum or copper to form vertical interconnects [2].

Along with wafer bonding techniques, the stacking of multiple transistor layers can also be achieved using successive growths of silicon from crystalline seeds located in the bottom layer [3]. Vias are etched through to access these seeds that provide a crystal lattice template for each growth process. The resulting islands of silicon grown in each layer are electrically isolated by oxide layers in between. They form the substrate onto which complementary metal-oxide-semiconductor (CMOS) transistors can be fabricated before growing another layer of silicon islands on top.

Optical interconnects can also be used to address the interconnect delay issue of electronics as their bandwidth is significantly larger than conventional metal wires. Optical interconnects between stacked CMOS transistor layers have been demonstrated using vertical-cavity surface-emitting lasers (VCSELs) and photodetectors bonded in between each layer [4]. This also demonstrates the integration of heterogeneous materials on the same substrate since the lasers and photodetectors were fabricated separately with III-V semiconductor materials while the CMOS electronics were fabricated using conventional silicon-based materials.

1.2 Silicon-based Optoelectronics

Silicon has long been the preferred material for electronics and has more recently become a material option for photonic devices. The band gap of silicon is 1.1 eV, meaning light with a wavelength greater than 1.1 μm will not be absorbed. This makes silicon suitable for the transmission of optical telecommunications wavelengths, which are typically centered around 1550 nm. These infrared wavelengths were chosen since the absorption of silica-based fibers is particulary low for the wavelength band of 1500-1600 nm.

The material system that is most commonly used for silicon photonic devices is silicon-on-insulator (SOI). This material system has been used to demonstrate high-speed optical modulators [5], photodetectors [6] and laser emission using the Raman effect [7]. Many other optoelectronic devices have also been realized in SOI, including: splitters, couplers, switches and (de)multiplexers, among others [8, 9]. SOI wafers consist of a top layer of crystalline silicon, referred to as the device layer, a buried insulating layer of silicon dioxide (SiO_2), often referred to as simply the buried oxide, and a base layer of silicon to provide mechanical support, commonly referred to as the handle. This structure is shown in Figure 1.1, with the corresponding index of refraction for the 1550 nm wavelength provided for each material layer.

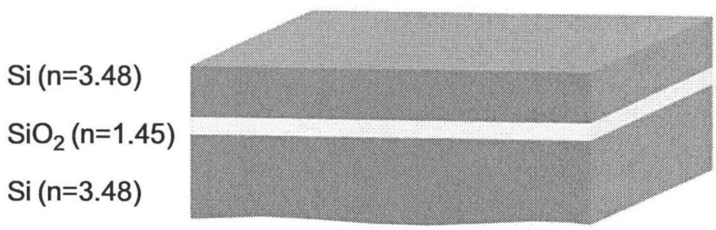

Figure 1.1: SOI wafer structure

SOI was originally developed to improve the performance of electronics by eliminating current leakage into the silicon substrate. The buried insulator layer confines carriers from transistors to the upper device layer of silicon. The reduction in current leakage results in faster transistor switching and reduced power consumption [10]. For silicon photonics, the buried oxide layer provides optical confinement resulting in waveguiding in the upper layer of silicon.

1.2.1 Advantages of Silicon for Optoelectronics

The main advantage of using silicon-based materials for photonic devices is their compatibility with existing CMOS technology, enabling the monolithic integration of photonics and electronics on the same chip. This compatibility also allows for cost-effective manufacturing of silicon optoelectronic devices by leveraging the existing CMOS foundry infrastructure to which significant investments have been made for several decades. While SOI wafers are more expensive than conventional silicon wafers, they are relatively inexpensive when compared to more exotic materials such as III-V semiconductors and lithium niobate. Silica-on-silicon wafers are also inexpensive but they provide limited options for high-speed active devices due to the insulating properties of the silica guide layer.

The waveguiding layer of SOI can be easily made several microns thick, unlike III-V semiconductor materials that are generally limited to around 1 μm since they must be completely grown using molecular beam epitaxy. These larger dimensions make SOI waveguides easier to fabricate, though low-loss nanowire waveguiding structures with dimensions on the order of ~200 nm have also been demonstrated in SOI [11]. As well, waveguides with larger cross-sections have greater optical coupling to single-mode fibers, which have core diameters of around 8-10 μm, resulting in lower device insertion losses. An important consideration is the refractive index difference between the fiber and the silicon guide layer, which is not a concern with silica waveguides. This results in a Fresnel reflection of 31 % (or a loss of 1.6 dB) for each fiber-to-waveguide interface. These losses can be nearly eliminated by using an anti-reflection coating on the waveguide facets.

SOI has a much higher refractive index contrast ($\Delta n \sim 2$) compared to other material systems used for waveguide devices, such as silica ($\Delta n \sim 0.5 - 0.75\,\%$)

and III-V semiconductors ($\Delta n \sim 0.1 - 0.2$). The high-index contrast of SOI reduces the amount of evanescent field propagating outside the waveguide core, when comparing the different materials using the same waveguide dimensions. This greater optical confinement allows for smaller waveguide bend radii and consequently more compact devices [9].

The most limiting disadvantage of using SOI for optoelectronics is that an all silicon-based laser has yet to be developed due to the indirect band gap of silicon. The Raman lasing that has been demonstrated still requires an external pump laser based on a III-V diode [7]. This inhibits the complete integration of optoelectronic devices on a single silicon substrate using standard CMOS processing exclusively. While III-V laser diodes can be incorporated into silicon devices using hybrid integration techniques [12], an all-silicon current-injection laser would be the ideal solution. Silicon-based light sources using dislocation engineering, nanocrystals and rare-earth dopants such as erbium are currently being investigated [13].

1.2.2 SOI Wafer Fabrication

While the thermal growth of an oxide layer onto silicon is straightforward, it is challenging to grow another layer of crystalline silicon on top due to the crystal lattice mismatch between the silicon and the amorphous structure of SiO_2. The resulting silicon layer would be either amorphous or polycrystalline, with both having unacceptably low mobilities for electronics and high absorption losses for optical devices. Therefore unique processes for the manufacturing of SOI wafers have been developed to ensure that the upper layer of silicon is in fact crystalline. These processes include: separation by implanted oxygen (SIMOX), bond and etch-back (BESOI) and Smart-Cut®. These processes are illustrated in Figure 1.2, which has been adapted from [14]. A comprehensive review of these techniques can be found in [10].

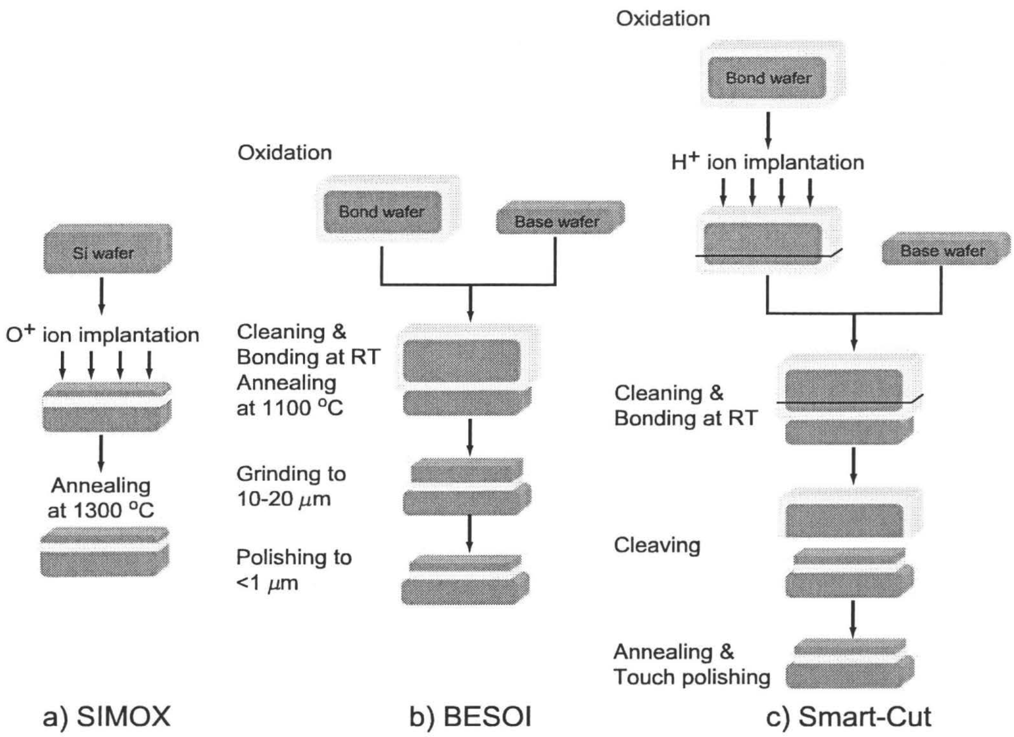

Figure 1.2: SOI wafer fabrication processes
a) SIMOX process using oxygen ion implantation b) BESOI process using grinding/polishing c) Smart-Cut® process using an atomic cleave plane produced by hydrogen ion implantation

The SIMOX process involves the implantation of oxygen ions below the surface of a silicon wafer. Damage to the lattice of the crystalline silicon caused by the implantation process is then repaired using high temperature annealing at around 1300 °C. This anneal also sharpens the boundaries of the buried oxide layer. The depth and thickness of the buried oxide layer is determined by the implantation energy used. For a 200 keV implant, the buried oxide layer thickness is approximately 0.5 μm with an upper crystalline silicon layer of 0.3 μm [13]. Since a device layer with a thickness of several microns is often desired for SOI-based optical devices, additional silicon can be epitaxially grown using chemical vapor deposition. The thickness of both the buried oxide and the silicon device layer have uniformities that are typically within a few percent.

BESOI wafers are produced by physically bonding a silicon wafer with another silicon wafer that has been thermally oxidized. An initial bond is formed at room temperature by simply bringing the two clean hydrophilic surfaces (Si and SiO_2) into contact. Thermal annealing at temperatures of up to 1100 °C strengthens the bond to that of bulk material. The upper silicon wafer is then thinned to the desired device layer thickness (> 1 μm) using chemical mechanical polishing (CMP) techniques.

The Smart-Cut® process, which has been patented by SOITEC Inc. of France, uses implanted hydrogen ions as an atomic cleave plane to separate the upper layer of silicon. This process results in improved uniformity of the device layer ($< \pm 5$ %) in comparison to the grinding and polishing techniques used for BESOI ($\pm$0.3-0.5 μm) [14]. The use of H^+ ions instead of heavier O^+ ions also reduces the probability of lattice defects in the device layer, in comparison to SIMOX wafers.

Following the hydrogen implant, the oxidized wafer is bonded to another silicon wafer, as it would be with the bond and etch-back process. Subsequent thermal annealing causes the implanted wafer to split along the plane where the silicon lattice bonds have been weakened by the implanted hydrogen ions. This annealing also strengthens the bond between the base wafer and the remaining portion of the implanted wafer. A fine CMP is then used to reduce the roughness of the SOI surface. The device layer thickness of Smart-Cut® wafers is controlled by the implantation energy used and is typically limited to around 1.5 μm. As with the BESOI process, the thickness and uniformity of the buried oxide layer of Smart-Cut® wafers can be more easily controlled compared to the SIMOX process since it is thermally grown in a separate step.

The choice of SOI wafer type depends on the application requirements and manufacturing budget. The advantage of using BESOI wafers over those fabricated via the SIMOX and Smart-Cut® processes is that device layers that are several microns thick are possible without the additional costs of epitaxial growth. Yet, BESOI is not the ideal choice if the application requires strict uniformity specifications. Another advantage of BESOI wafers is that the device layer will have a very low defect concentration since they have not been subjected to high-dose ion implantation. For applications requiring a device layer that is both thin (< 1.5 μm) and highly uniform, Smart-Cut® wafers are most likely the ideal choice, although they are generally more costly than SIMOX and BESOI. If a thicker device layer is required a choice must be made between the higher costs of epitaxial growth with SIMOX and Smart-Cut® wafers and the compromised uniformity of lower-cost BESOI wafers.

1.2.3 Optical Properties of Silicon Photonic Devices

SOI can be used for optoelectronic devices since the buried oxide layer provides optical confinement in the upper layer of silicon. This vertical confinement is due to the difference in the indices of refraction between the silicon device layer and the buried oxide ($\Delta n = 2.03$) and between the silicon device layer and the cover layer of air ($\Delta n = 2.48$). The cover can also be another oxide layer. The buried oxide layer is often referred to as the substrate when discussing the waveguiding properties of SOI since a negligible amount of light penetrates into the base silicon layer. The buried oxide layer need only be 0.4 μm thick for guide layers exceeding

2 μm, while thinner guide layers require thicker buried oxides to prevent coupling into the substrate [13]. Horizontal optical confinement in an SOI waveguide can be achieved by etching a structure into the device layer, such as the rib waveguide shown in Figure 1.3.

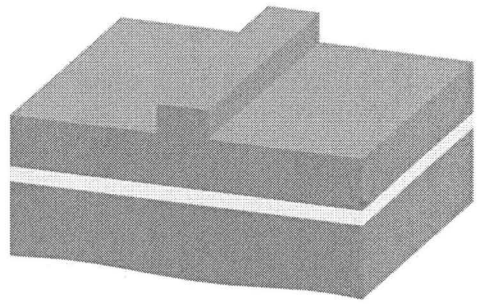

Figure 1.3: SOI rib waveguide

Propagation losses for SOI waveguides are relatively low at 0.1-0.5 dB/cm for the telecommunications wavelength range of 1.5 to 1.6 μm [13]. These losses are primarily due to sidewall roughness since silicon itself is transparent at these wavelengths. Some scattering at crystal defect sites can also occur, but this is generally insignificant due to the low defect concentration of commercial SOI. The presence of carriers from wafer doping will also increase propagation losses due to absorption. Low-loss SOI waveguides require dopant concentrations to be less than 1×10^{16} cm^{-3}, resulting in a resistivity greater than 10 Ω·cm [13].

Control of the refractive index and absorption coefficient of the silicon device layer is required for active optoelectronic devices such as modulators and switches, among others. A change in the refractive index is usually achieved using the free carrier plasma dispersion effect with a current-injection diode or the thermo-optic effect with a metal strip heater. The electro-optical effects available in silicon will be thoroughly explained in Section 2.8 in the next chapter.

1.3 Published Vertically-Integrated Photonic Devices

Vertically-integrated optical couplers that transfer light between stacked waveguide layers are the basic building blocks required for the three-dimensional routing of optical signals. These devices have been demonstrated using different structures and in various material systems including: III-V semiconductors, polymers, nitride, silica and SOI. This vertical signal routing can also be accomplished using 3-D photonic crystal structures, as detailed in [15] and demonstrated in [16]. These structures represent a separate area of research and will not be included in the following summary of published vertically-integrated photonic devices.

1.3.1 III-V Semiconductor Vertical Couplers

Coupling between stacked waveguides in a III-V semiconductor material system consisting of InGaAsP and InP was first demonstrated in [17]. Optical power was transferred between two device layers using a vertically-integrated directional coupler. A comparison between a conventional directional coupler that couples along the horizontal axis only and its vertically-integrated counterpart is illustrated in Figure 1.4.

Power is transferred between the waveguides of the conventional coupler when the separation between them is reduced such that their optical fields overlap. The coupler length at which the power is completely transferred depends on the amount

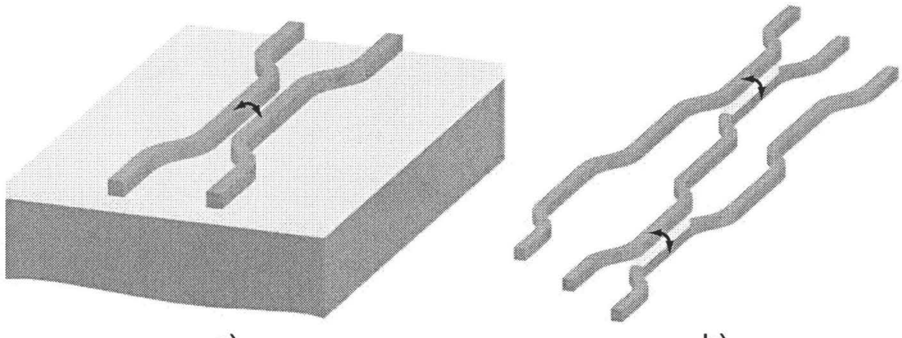

Figure 1.4: Comparison of horizontal and vertical directional couplers in III-V semiconductors
a) Horizontally-arranged directional coupler b) Vertically-stacked directional couplers in a three-layered photonic device

of overlap and how well matched the two waveguides are. Its vertical counterpart operates using the same principle, but with its fields coupling along the vertical axis wherever the waveguides overlap. Additional couplers can be fabricated in multiple layers, as well as along the horizontal plane of each layer, increasing the overall functional density of the device. A three-layered photonic devices with two vertically-integrated couplers is shown in Figure 1.4 b).

The minimum spacing between waveguides in the horizontal plane is limited by fabrication techniques. Vertically-integrated directional couplers can be made with a very thin separating layer between the two stacked waveguides, resulting in greater field overlap and hence a shorter coupler length at which maximum power transfer occurs. Coupler lengths of these stacked devices can be reduced to hundreds of microns, instead of millimeters (1-3 mm) required for conventional horizontally-orientated couplers.

The above coupler was fabricated using a backside processing technique. The layers for both waveguides were first deposited using metal-organic chemical vapor deposition (MOCVD). The waveguide for one of the layers is patterned and etched and then the wafer is bonded facedown to a support wafer. The backside substrate of the wafer is removed to access the layers for the opposite waveguide, which is then finally patterned and etched such that it is aligned directly above the other guide. This same processing technique was also used to fabricate vertically-integrated micro-resonator disks stacked on top of waveguides that had been etched into the lower layer [18]. These micro-resonator disks are used to filter wavelengths from an optical signal. Wavelengths that satisfy the resonant condition of the disk will vertically couple into the disk from the bus waveguide below, leaving the other wavelengths to continue propagating along their existing path.

A novel three-layer InP-InGaAsP waveguide vertical beam splitter has also been demonstrated [19]. This vertically-integrated device uses multiple directional couplers to equally split the optical power from one input waveguide into eight output guides distributed among the three layers. The enhanced overlap of these strongly coupled waveguides allowed for a 583 μm device length, more than 100 times shorter than that of the equivalent horizontal coupler needed to achieve a 1:8 splitting of optical power.

While multiple waveguide layers could have been deposited using MOCVD, these devices were fabricated using a wafer bonding process. The lowermost waveguides are first patterned and etched into the base substrate before bonding another MOCVD wafer facedown. Its backside substrate was then removed, allowing the waveguides in the middle layer to be patterned and etched. The same process is repeated for the upper layer waveguides. To achieve fusion bonding, the wafers were cleaned to remove any contaminants and then placed together in contact under a pressure of 2-3 MPa. Heating the wafers to near the growth temperature of one of the materials allowed atomic redistribution to occur at the bonding interface. This filled in any deviation between the bonded surfaces resulting in an atomically smooth interface between them. Alignment during the photolithographic patterning of the middle and upper waveguides was achieved using infrared light to backside illuminate the alignment marks of the lower waveguides.

The majority of the vertically-integrated devices developed using III-V semiconductors has been demonstrated by the optoelectronics research group at the University of California, Santa Barbara, headed by Dr. John E. Bowers. A review of the group's work in this area can be found in [20]. Here, the electro-optic effect of III-V semiconductors has also been used to actively-control the amount of vertical coupling between two stacked waveguides by varying the reverse bias voltage applied between them. This group has more recently demonstrated the vertical integration of a III-V diode laser above a silicon waveguide, making optical interconnects in CMOS electronics a viable solution to current interconnect bandwidth limitations [12].

1.3.2 Polymer-Based Vertical Couplers

Vertically-integrated photonic devices have also been demonstrated in polymer-based material systems. A power splitter used to separate input light into two optically isolated waveguide layers is illustrated in Figure 1.5, along with the gray-scale photolithography process used to fabricate them.

The photoresist mask used to define the waveguide structure is first exposed to ultraviolet light through a photomask pattern whose transmission is gradually reduced. The resulting pattern in the developed photoresist mask is sloped, which translates into a sloped structure as the thinner resist regions are eroded first during

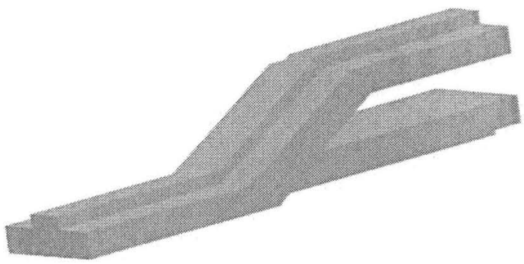

a)

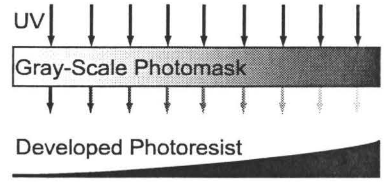

b)

Figure 1.5: Gray-scale photolithography for vertically-integrated polymer-based photonic devices
a) Polymer-based power splitter b) Gray-scale photoresist exposure and development

the subsequent plasma etching of the polymer layer. The coupler length required to equally split the power between the two stacked waveguides was typically 1 mm [21]. The polymer employed also exhibits an electro-optic effect when a voltage is applied. This change in refractive index was used to actively-control the power splitting ratio.

The input light can also be separated into its TE and TM polarization components by making each output waveguide of the splitter with different polymer materials [22]. Each material will have its own birefringence properties resulting in a splitting of the two polarization components between each stacked waveguide layer.

Reflectors angled at 45 ° have also been used to direct light between vertically-integrated polymer waveguides [23]. These reflectors were fabricated using imprinting techniques, where an etched silicon wafer relief template is pressed into a softened polymer film.

A vertically-integrated coupler based on a multimode mode interferometer (MMI) structure has also been demonstrated using polymers. The general operating principles of MMI couplers can be found in [24]. An illustration comparing horizontally-arranged and vertically-integrated MMI-based couplers is shown in Figure 1.6. These devices support the propagation of multiple optical modes and the beating between these modes results in behavior known as self-imaging. The field from a single-mode input waveguide is first coupled into a multimode slab region. Its modes are then excited resulting in a splitting of the input field into multiple images as the light propagates along the MMI section. At certain MMI coupler lengths the image of the input field will become two-fold. The input light can then be equally split into two output waveguides placed at this length. The vertically-integrated MMI coupler operates using the same principle, with the exception that the image launched from either the lower or upper input waveguide is split along the vertical axis.

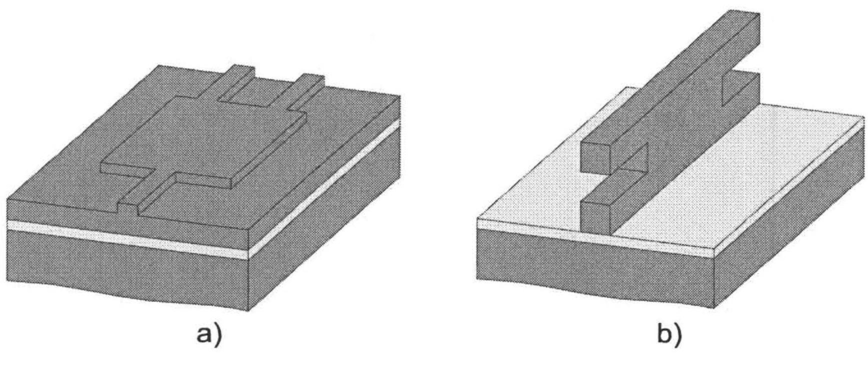

Figure 1.6: Comparison of horizontally and vertically-integrated MMI-based couplers
a) Horizontally-arranged MMI coupler b) Vertically-stacked MMI coupler

Polymer-based MMI couplers have been demonstrated in [25], where the best coupling of about 84 % was predicted to occur at a coupler length of 5.65 mm. This length can be reduced to 1.1 mm if a stepped MMI coupler structure is employed. This shorter length can be attributed to additional reflections off the upper material interface placed closer to the propagating optical field. These two types of polymer-based MMI couplers are illustrated in Figure 1.7.

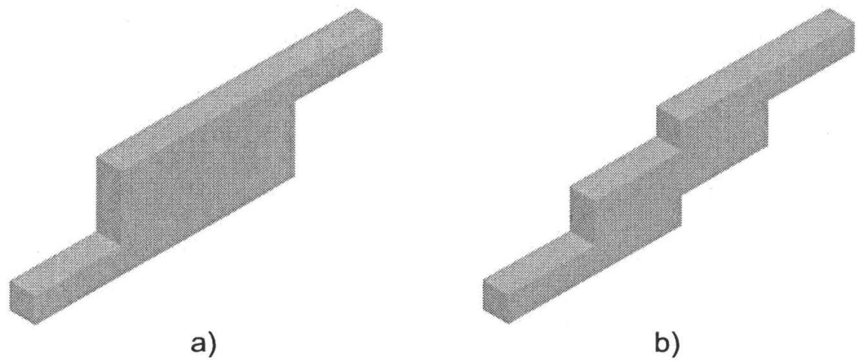

Figure 1.7: Vertically-integrated polymer-based MMI couplers
a) Single-section polymer-based MMI coupler b) Stepped polymer-based MMI coupler

1.3.3 Silica-Based Vertical Couplers

Vertical optical coupling has also be demonstrated in the silica glass material system [26]. A high coupling efficiency of 96 % between stacked silica waveguides was experimentally demonstrated over a coupling length of 800 μm. Wavelength-selective coupling between vertically-integrated silica waveguides via radiation modes has also been demonstrated using grating couplers [27]. These 0.5 mm long grating structures reflect, and hence couple, only the wavelengths that satisfy its Bragg condition.

1.3.4 MEMS-Based Vertical Couplers

Micro-electro-mechanical systems (MEMS) technology has been used to demonstrate active switching of light between stacked waveguides in a silicon oxynitride

(SiON) material system. This device is illustrated in Figure 1.8 and has been presented in [28]. The upper SiON waveguide is suspended above the lower guide with an air gap in between. Coupling between the two guides is introduced by electrostatically closing the air gap such that their evanescent fields overlap. As is common in many MEMS-based devices, a voltage is applied to the bridge structure causing an accumulation of charge on its lower surface. This charge build-up results in an electrostatic force that pulls the bridge structure into the layer below.

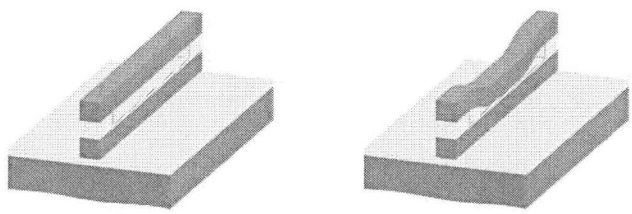

a) b)

Figure 1.8: MEMS-based vertically-integrated directional coupler
a) Open bridge with zero applied voltage b) Closed bridge with applied voltage greater than pull-in voltage

The bridge structure can also be made into a cantilever arm by etching through one end of the bridge during the fabrication process. The released end of the cantilever can be electrostatically bent to direct light between multiple waveguide layers, as demonstrated using nitride waveguides in [29]. Highly efficient switching can be achieved since this device directly couples light from its end facet, instead of relying on the coupling of evanescent fields.

1.3.5 SOI-Based Vertical Couplers

Soref was the first to propose and demonstrate vertical optical coupling between two silicon guide layers [30]. Coupling between stacked slab waveguide structures was achieved using a SIMOX wafer with two silicon device layers created with successive oxygen ion implantations and epitaxial silicon growths. Soref was also the first to propose using a patterned SIMOX process to define waveguides buried into lower layers. A suitable material would be patterned onto the top surface to mask certain regions during the ion implantation step, resulting in a buried oxide layer of varying thickness and depth around the waveguide structure.

This technique has been used by Dr. Jalali's group at UCLA to pattern buried micro-resonator disks into SOI wafers [31]-[34]. This process will be further explained in Section 1.5. A bus waveguide etched into the top surface overlaps with the resonator disk below, resulting in the filtering of wavelengths from the input signal that satisfy its resonance condition. This condition is primarily determined by the radius of the disk, which was typically $\sim$20 μm.

This group has also fabricated a field-effect transistor on top of the resonator disk to demonstrate that electronics can also be integrated with a multi-layered silicon photonic device [35]. They have also demonstrated coupling between an additional layer by using another ion implantation step to bury a third waveguide [36].

Gray-scale photolithography has been used to fabricate efficient input couplers for conventional single-layer SOI optoelectronic devices. Here, the silicon guide layer is vertically tapered from a 10 μm thick SOI waveguide, matching the core diameter of an optical fiber, to a smaller 0.25-1 μm thick guide [37]. This technology can also be extended to fabricating vertically-integrated power splitters similar to the devices demonstrated using polymer materials. Coupling between two vertically-stacked SOI waveguides using grating structures has also been simulated [38].

An alternative to creating stacked silicon device layers using either ion implantation or wafer bonding is the deposition of polysilicon thin films. This technique has been used to demonstrate vertical coupling between an upper polycrystalline ring resonator deposited and patterned above a lower crystalline silicon waveguide [39]. Ring resonators, which resemble race tracks, also filter narrow bands of wavelengths from an optical signal. Equivalent to micro-resonator disks, the filtered wavelengths are those that satisfy the resonance condition of the ring, which was fabricated with a radius of 45 μm. The main disadvantage of using polysilicon as a photonic device layer is its unacceptably large optical absorption coefficient.

1.4 Vertically-Integrated Photonic Devices Presented in This Thesis

Vertically-integrated optical couplers in SOI have been demonstrated in this project using various stacked structures. Directional couplers based on stacked rib waveguides have been presented in [40, 41]. These devices are similar to the couplers demonstrated using III-V materials but now leverage the advantages of the SOI material system (*i.e.* cost-effective manufacturing, integration with CMOS electronics and compact device size).

These structures are also suitable for fabrication using wafer bonding techniques. This is advantageous over patterned SIMOX-fabricated devices since electronic devices can now be integrated into each photonic layer before they are stacked with wafer bonding. While the patterned SIMOX process can be used to define multiple layers of buried waveguides, only the top surface is accessible for the metallization processes required to manufacture electronics. While vias can be etched through to the lower layers, this consumes device real estate, negating the main advantage of vertical integration.

As well, vertical coupling in SOI photonic devices has so far been demonstrated using micro-resonator disks and rings that are inherently wavelength sensitive. This can be advantageous if individual wavelength signals are to be routed to different layers. Here, resonators of varying radii fabricated on different layers will filter their corresponding resonant wavelength through vertical coupling from an overlapping waveguide. While this facilitates a novel signal routing scheme, many other applications would require broadband vertical couplers capable of routing a wide range of wavelengths simultaneously between layers. The couplers fabricated for this project have minimal wavelength sensitivity.

All vertically-integrated directional couplers require that their stacked waveguides be sufficiently close such that their optical fields overlap. This also results in undesirable optical cross-talk between the two layers since coupling will occur wherever the waveguides in each layer overlap. Positioning the waveguides on each layer in a such way to minimize this undesirable overlapping would be an inefficient use of device real estate. While gray-scale photolithography can be used to gradually taper the separating material between them from being thick enough for optical isolation to being thin enough for efficient coupling, this process adds complexity to the device fabrication.

An alternative coupler based on a vertically-integrated MMI structure in SOI has been demonstrated and is presented in [42, 43]. Though independently proposed, it was latter found to be a silicon-based version of the polymer device shown previously in Figure 1.7 a). The advantage of using an MMI structure is that coupling can be easily restricted to only certain areas of the device since the oxide between the stacked waveguides is made sufficiently thick for complete optical isolation. A multimode slab of silicon can then be either grown or deposited into a trench that has been etched into the separating oxide before bonding another waveguide layer directly above the coupler region.

1.5 Techniques for Manufacturing Vertically-Integrated Silicon Optoelectronics

Multiple stacked waveguide layers can be fabricated using successive epitaxial growths, or amorphous depositions, of silicon and oxide layers. Silicon layers are typically grown/deposited using silane (SiH_4) gas-based chemistries. If this deposition occurs on top of an amorphous oxide layer, the resulting silicon layer will not have grown using a crystal lattice template and its resulting crystal structure will be either polycrystalline or amorphous itself. The grain boundaries of these films cause unacceptable optical propagation losses. While these losses can be reduced to 9 dB/cm using high temperature anneals and special hydrogen plasma passivation steps [44], waveguides fabricated in crystalline silicon can have losses that are below 1 dB/cm. Deposited polysilicon layers can also be laser recrystallized but the resulting material is not typically defect-free [45].

Crystalline silicon can be epitaxially grown on top of an oxide layer by initiating the growth from a silicon seed below [3], accessible through a via etched into the oxide. This process is known as epitaxial lateral overgrowth (ELO), as the selective growth of the silicon seed fills the via and then laterally spreads to cover the oxide. The resulting layer of silicon is typically non-uniform and must be planarized using CMP. Optoelectronic devices are then fabricated in the uppermost layer of silicon before depositing another layer on top. This is a complex fabrication process, while patterned ion implantation and wafer bonding are more likely to be adaptable to existing electronics manufacturing.

1.5.1 Patterned Ion Implantation Process

Ion implantation can also be used to define waveguides buried below the wafer surface. The smooth upper surface of silicon is then regrown epitaxially before defining the next layer of waveguides. Buried waveguides in silicon have been produced using a focused beam of 250 keV protons that selectively slows down the rate of porous silicon formation during subsequent anodization. This results in a buried silicon core surrounded by a porous silicon cladding [46]. As discussed previously, a patterned SIMOX process can be used to fabricate buried SOI waveguides. This process is illustrated in Figure 1.9.

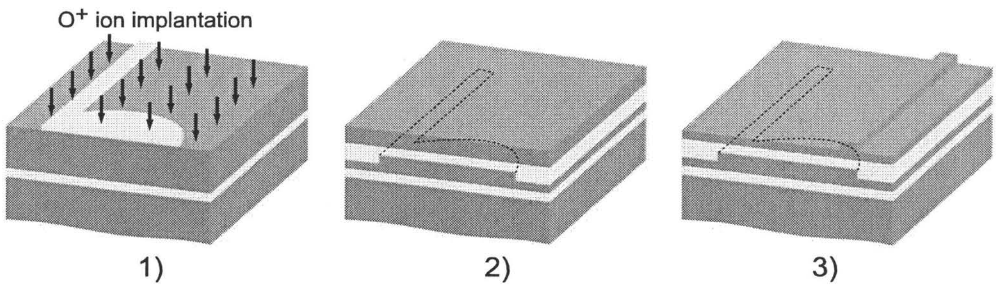

Figure 1.9: Patterned SIMOX process for fabricating vertically-integrated SOI micro-resonators

1) Patterned oxygen ion implantation 2) High-temperature anneal to form buried resonator disk 3) Patterning and etching of upper bus waveguide

Soref was first to propose this technique for fabricating buried SOI waveguides, and demonstrated coupling between two silicon slab waveguide layers defined by successive oxygen ion implantations and epitaxial growths [30]. It has since been applied using masked ion implantations to demonstrate vertically-integrated micro-resonator disks [34]. Here, the mask defining the resonator is patterned onto the initial SOI wafer, followed by an oxygen ion implantation, as illustrated in Figure 1.9 1). The mask pattern has now been transferred into a buried layer of oxide defining the resonator disk after a high-temperature anneal at 1300 °C , as shown in step 2). A bus waveguide that will vertically couple into the resonator below is then patterned and etched onto the upper silicon surface in step 3).

Successive patterned implantations and epitaxial growths can be used to define multiple layers of buried waveguides and resonators before etching the uppermost waveguide. Other implant species could also be used to locally change the index of refraction to create buried waveguides. Yet, the formation of buried oxide layers using oxygen ion implantation provides the large index change required for compact devices.

1.5.2 Wafer Bonding Techniques

The main disadvantage with using ion implantation techniques to fabricate buried waveguides in multiple layers is that only the uppermost device layer can be integrated with electronics. This would restrict the active control of optical coupling to between the top two waveguide layers only. Another limitation of the patterned SIMOX process is its required high-temperature annealing that will damage any existing electronic structures elsewhere on the chip. In other words, it is not a suitable back-end process for conventional CMOS manufacturing.

Ideally, each device layer would be fabricated onto separate SOI wafers and then bonded together. The backside handle of one of the SOI wafers would have to be removed before bonding another wafer. Both electronics and photonics can be monolithically integrated onto each device layer, and planarized using conventional CMOS-compatible processes, before they are stacked.

Another advantage with using wafer bonding is the improvement in overall yield for multi-layered devices, since each layer can be inspected for fabrication errors before bonding them together. Fabrication errors in subsequent ion implanted layers would render the overall device useless regardless of the successful fabrication of previous layers. This yield issue is also a problem with multilayered devices fabricated using either successive epitaxial growths or depositions of silicon. Wafer bonding does have its own yield considerations, along with the added complication of precise alignment being required during bonding as it is generally easier to accomplish accurate alignment with photolithographic processes.

An overview of the silicon wafer bonding processes currently available is given in [47]. One of these techniques is to use intermediate layers deposited between the two wafers to form a bonding interface. The bonding materials that have been used include: thin gold films, polymers, solders, low melting temperature glass frits and spin-on-glasses (SOGs). Many of these wafer-level techniques are simply modified versions of those applied at the die level in electronics packaging. The majority of these techniques involve an annealing step to reflow the intermediate layer to solidify the bond upon cooling.

This project evaluated SOG as a potential candidate for wafer bonding. These liquids are composed of siloxane molecules (Si-O) or silicates dissolved in a solvent, such as n-Butanol, and are frequently used for the planarization of electronics. They form near glass-like thin films after being spun onto a wafer and cured with an anneal at 800 °C. The thickness of the film can range from hundreds of nanometers to microns depending on the spin rate used and how many layers have been spun.

SOI wafers have been manufactured using an intermediate layer of SOG doped with 2 wt-% boron and 2 wt-% phosphorus [48]. These additional dopants cause the glass layer to reflow when annealed at 1050 °C, further solidifying the bond. The SOG is first spun onto oxidized wafers and then the pair is bonded in a vacuum fixture and annealed in an O_2 furnace. SOG has also been used to bond III-V diode lasers to a silicon substrate in [49].

Another bonding technique frequently used in MEMS processing is anodic bonding. Here, silicon wafers are bonded to glass wafers containing alkali ions such as sodium or potassium. The two wafers are placed in contact under pressure and then a high voltage field is applied across the wafer pair. The mobile sodium ions migrate away from the bonded interface, leaving behind fixed charge in the glass that creates a high electric field across the bond interface. This forms a chemical bond that fuses the wafers together. Existing electronics on the wafers to be bonded must be suitably shielded to avoid damage from the high voltage field. SOG doped with potassium oxide has been used to create an intermediate glass layer suitable for anodic bonding [50].

Direct bonding methods form the most robust bonding interface between two wafers, whose strength approaches that of bulk silicon. This technique is commonly used for manufacturing SOI wafers with either the bond and etch back or the Smart-Cut® process. Multiple device layers of silicon have been directly bonded together using the Smart-Cut® process, with each layer separated by a thin oxide film [51].

Direct bonding involves contacting wafers without the assistance of any intermediate layers, of significant pressure or applied voltage fields. These bonding schemes rely on the tendency for smooth surfaces to adhere. First, both wafers are cleaned in either sulphuric or nitric acid to create hydrophilic surfaces with dangling silanol (Si-OH) bonds. The wafers are brought into contact and adhere through hydrogen bonding. The bond is then strengthened by annealing the wafer pair to polymerize the silanol bonds into siloxane bonds (Si-O) bonds. The bonding strength approaches that of bulk material at temperatures greater than 1000 °C [47].

This bonding technique is not suitable as a back-end process since its high temperature may damage existing devices. Therefore a lower-temperature bonding procedure is required to stack wafers that have already been processed. The use of plasma-activation has proven to be a robust low-temperature bonding technique and has been demonstrated in [52]-[55]. Here, the wafers to be bonded are first exposed to an O_2 gas plasma. High bonding strength can be achieved at lower annealing temperatures since more silanol groups are created by exposure to the O_2 plasma than by wet chemical activation alone. It has also been suggested that plasma treatment creates a porous surface structure which enhances the diffusivity of water molecules from the bonded interface [54]. Published results using this bonding technique have demonstrated that bulk bonding strength can be achieved with a low-temperature anneal of only 400 °C [52]. Further details for the bonding processes presented above will be provided in Chapter 4: Device Fabrication.

IBM has developed a novel wafer bonding technique for stacking multiple layers of integrated circuits, and is illustrated in Figure 1.10. The processed SOI wafer is first bonded facedown to a quartz support wafer using polymeric adhesives, as shown in step 1). The backside handle of the SOI wafer is removed by griding and etching in step 2), with the buried oxide acting as an etch stop, before oxide fusion bonding is used to transfer the thin device layer onto another processed SOI wafer in step 3). The transparency of the quartz support wafer allows for the two wafers to be visually aligned before bonding. The bonded pair was annealed at a temperature of

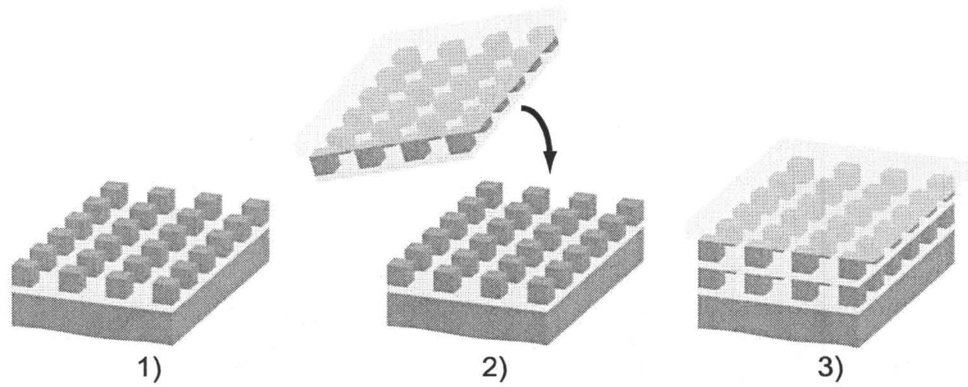

Figure 1.10: IBM wafer bonding technique
1) Initial processed SOI wafer 2) Bonding of the quartz wafer-held device layer, with handle removed, to another processed SOI wafer 3) Completed bonding with quartz support wafer to be removed later

only 400 °C, though bonding strength was not reported. Finally, the support wafer is released using laser ablation and the remaining adhesive is stripped. This process can be repeated to stack multiple layers. Between each bond, vias are etched through to the layer below and filled with a conductive material to form electrical interconnects [56].

1.5.3 Wafer Bonding With Precise Alignment

The most significant impediment of using wafer bonding for vertical integration is the precise alignment tolerances required. Stacked waveguides will efficiently couple only if they are placed directly above each other. Currently available commercial wafer bonders have typical alignment errors of $\pm$1-2 μm. A backside infrared source is commonly used to illuminate through the wafers to be bonded and alignment is performed visually using an infrared camera before placing them in contact. This specification is inadequate if smaller structures are to be vertically-integrated.

An improved alignment technique inspired by interlocking LEGO® blocks has been recently demonstrated [57]. Interlocking convex pyramids and concave V-pits etched into silicon wafers were used to achieve alignment accuracy of better than 200 nm between a bonded pair of wafers. The smooth sidewalls of the etched features are sloped at a precise angle of 54.76 ° since they were fabricated using an anisotropic chemical that etches preferentially along one crystallographic plane. This significant improvement in alignment tolerance facilitates the vertical-integration of photonic devices with feature sizes smaller than 1 μm.

IBM has reduced misalignment errors between bonded wafer pairs to only 180 nm, without using interlocking structures [56]. This was accomplished using wafer bow compensation techniques, which is the main cause of misalignment during bonding. This technique was applied in conjunction with their concept of using a transparent quartz support wafer to transfer thin layers of SOI integrated circuits onto another substrate.

A summary of the different techniques that can be used to manufacture vertically-integrated photonic devices in silicon is provided in Table 1.1. The main advantages of using a wafer bonding based process are: the possibility for integration of electronics and photonics onto each layer, higher overall device yield by individual inspection of each layer before stacking and the absence of high temperatures that could damage existing fabricated electronics.

Table 1.1: Summary of techniques for manufacturing vertically-integrated photonic devices in silicon

	Polysilicon Deposition	Ion Implantation	Wafer Bonding
Sub-100 nm alignment	✓	✓	
Low optical absorption		✓	✓
Integration of electronics onto each layer	✓		✓
Inspection of each layer before stacking			✓
Low-temperature post-processing			✓

1.6 Summary of Accomplished Work

The following accomplishments have been achieved during this project:

1. Thorough simulation of five proposed vertically-integrated optical couplers yet to be demonstrated in SOI, along with their tolerance to fabrication errors
2. Development of an in-house wafer bonding procedure based on existing plasma-activation techniques
3. Development of novel fabrication processes involving wafer bonding and backside processing to create proof-of-concept devices
4. Successful demonstration of vertical coupling between stacked slab waveguides using double-bonded SOI wafers
5. Novel demonstration of vertical coupling between stacked SOI rib waveguides, as well as between stacked SOI channel waveguides
6. Novel demonstration of vertical coupling within a vertically-integrated MMI structure fabricated in SOI

Chapter 2

Optoelectronic Device Theory

In this chapter, the electromagnetic theory of dielectric waveguides will be presented beginning with Maxwell's equations. The solutions of a simple slab waveguide will then be analytically derived before describing the beam propagation method used for propagating optical modes in waveguides with arbitrary cross-sections. The transfer of optical power between waveguides with overlapping modes will be explained using coupled mode theory. Finally, the material properties of silicon that can be employed to actively control optical switching will be described.

2.1 Electromagnetic Wave Equations

Maxwell's equations for the propagation of light in an isotropic, non-conducting, and non-magnetic dielectric medium are:

$$\nabla\times\mathbf{E} = -\mu_0\frac{\partial\mathbf{H}}{\partial t}, \qquad \nabla\cdot\mathbf{E} = 0 \tag{2.1}$$

$$\nabla\times\mathbf{H} = \epsilon_0 n^2\frac{\partial\mathbf{E}}{\partial t}, \qquad \nabla\cdot\mathbf{H} = 0 \tag{2.2}$$

where $\mathbf{E}$ and $\mathbf{H}$ are the time-dependent vectors of the electric and magnetic fields respectively, $\nabla = (\partial/\partial x, \partial/\partial y, \partial/\partial z)$ is the del operator, t is time, ϵ_0 and μ_0 are the dielectric permittivity and magnetic permeability of free space respectively, and n is the refractive index of the medium. It is common notation to express the $\epsilon_0 n^2$ term as simply the dielectric constant, ϵ, and to use μ in place of μ_0.

Combining the above equations to eliminate the magnetic field yields:

$$\nabla\times(\nabla\times\mathbf{E}) = -\mu\epsilon\frac{\partial^2\mathbf{E}}{\partial t^2} \tag{2.3}$$

Applying the vector identity of $\nabla\times\nabla\times\mathbf{A} = \nabla(\nabla\cdot\mathbf{A}) - \nabla^2\mathbf{A}$ and the assumption that the medium is non-conducting ($\nabla\cdot\mathbf{E} = 0$), yields the familiar wave equation:

$$\nabla^2\mathbf{E} = -\mu\epsilon\frac{\partial^2\mathbf{E}}{\partial t^2} \tag{2.4}$$

which supports the notion that light is an electromagnetic wave. A similar equation for the magnetic field can also be derived.

The solutions to the wave equation are electromagnetic waves of the form:

$$\mathbf{E}(x, y, z, t) = \mathbf{E}_0(r, t)e^{i(\mathbf{k}\cdot\mathbf{r}-\omega t)} \tag{2.5}$$

where $\mathbf{E}_0(r, t)$ is the slowly varying electric field amplitude of the wave in time and space, $\mathbf{k}$ is the wave vector, and ω is the angular frequency. The amplitude of the wave vector is nk_0, which is equal to $2\pi/\lambda_0$, where λ_0 is the vacuum wavelength. Note that the exponential component of these solutions is in the form of a harmonic wave. If $\mathbf{E}_0$ is constant, the solutions simplify to plane waves and the wave equation in 2.4 becomes:

$$-\mathbf{k}^2\mathbf{E} = -\omega^2\mu\epsilon\mathbf{E} \tag{2.6}$$

For plane wave solutions, the time and space dependence of the electric field become $\partial\mathbf{E}/\partial t = -i\omega\mathbf{E}$ and $\nabla \times \mathbf{E} = i\mathbf{k} \times \mathbf{E}$. Similar expressions exist for the magnetic field of a plane wave solution [58]. Maxwell's equations in 2.1 and 2.2 can now be written as:

$$\mathbf{k}\times\mathbf{E} = \omega\mu\mathbf{H} \tag{2.7}$$

$$\mathbf{k}\times\mathbf{H} = -\omega\epsilon\mathbf{E} \tag{2.8}$$

which shows that light is composed of electric and magnetic fields propagating in the direction of **k** and that the three vectors form an orthogonal set. In other words, **E** and **H** are always perpendicular to each other. This vector pair can also be oriented in any direction, which is referred to as the state of polarization. Linearly polarized light, for which the direction of the vibrating **E**-vector is fixed in time, is shown in Figure 2.1.

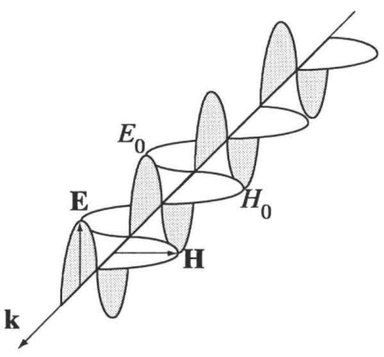

Figure 2.1: Electromagnetic wave propagating in free space
The vectors $\mathbf{E}$, $\mathbf{H}$, and $\mathbf{k}$ form an orthogonal set with the wave propagating in the direction of $\mathbf{k}$ given by $\mathbf{E} \times \mathbf{H}$.

2.2 Wave Propagation in Slab Waveguides

The most basic optical waveguiding structure is the dielectric slab waveguide, shown in Figure 2.2. The refractive indices of the cover layer, guiding layer, and substrate are n_c, n_g, and n_s respectively. Light cannot be guided unless $n_g > n_s$ and $n_g > n_c$ and the thickness, T, of the guiding layer is above a critical thickness. In such a two-dimensional waveguide, light is confined in the y-direction resulting in a beam that only expands along the x-axis while propagating. The coordinate system is chosen such that the y-direction is perpendicular to the layers and the direction of light propagation is along the z-axis.

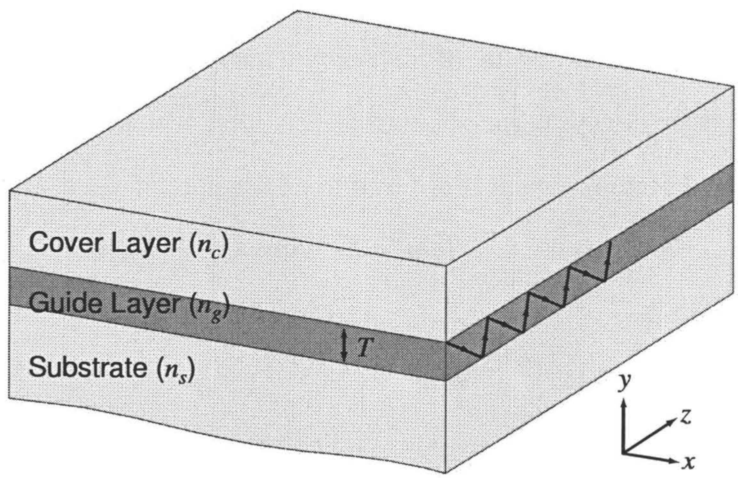

Figure 2.2: Slab waveguide structure

A simple ray optics model can be used to illustrate basic slab waveguide concepts. A guided mode will exist provided that the incident light ray makes an angle between the wave normal and dielectric interfaces such that it satisfies the condition for total internal reflection at both the upper and lower interfaces of the guiding layer. For now, a mode can be considered as simply a path along which the light propagates. In this case, the light ray will propagate in the guiding layer without attenuation along a zig-zag path, provided the waveguide material is lossless. The critical angles at each interface depend on the relative refractive indices of each layer. If the conditions for total internal reflection are not satisfied, light will be coupled into either the substrate or the cladding, or into both layers, resulting in attenuation.

Electromagnetic theory can be used for a more complete characterization of the modes for a slab waveguide. The solutions that will be given in this section follow the derivation given by Nishihara in [59]. A similar derivation can be found in [60]. For a plane wave propagating along the z-direction, the electromagnetic fields vary as:

$$\mathbf{E} = \mathbf{E}(x, y)e^{i(\omega t - \beta z)} \tag{2.9}$$

$$\mathbf{H} = \mathbf{H}(x, y)e^{i(\omega t - \beta z)} \tag{2.10}$$

where β is the propagation constant representing the z component of the wave vector, k_z. The effective index of a mode, N, is defined by:

$$\beta = k_0 N \tag{2.11}$$

and can be thought of as the index of refraction that governs the mode as it propagates along the z-direction. It will be lower than the material refractive index of the guide layer since a portion of the optical field will be propagating in the cover and substrate layers. Guided modes exist when $n_s < N < n_g$ provided that the index of the cover is lower than, or equal to, the index of the substrate layer. It is usually lower as the cladding layer is most often air.

Since the index profile of a slab waveguide is invariant along the x-direction, its electromagnetic fields are also independent of x. This results in the following conditions: $\partial/\partial t = i\omega$, $\partial/\partial z = -i\beta$, and $\partial/\partial x = 0$. Accordingly, Maxwell's equations in 2.1 and 2.2 will yield two different modes with mutually orthogonal polarization states. One is the TE mode, which consists of field components E_x, H_y, and H_z. The other is the TM mode with E_y, H_x, and E_z field components.

The wave equations for the TE modes of a slab waveguide are:

$$\frac{\partial^2 E_x}{\partial y^2} + (k_0^2 n^2 - \beta^2)E_x = 0, \tag{2.12}$$

$$H_y = -\frac{\beta}{\omega\mu}E_x, \;\; H_z = -\frac{1}{i\omega\mu}\frac{\partial E_x}{\partial y} \tag{2.13}$$

and for the TM modes they are:

$$\frac{\partial^2 H_x}{\partial y^2} + (k_0^2 n^2 - \beta^2)H_x = 0, \tag{2.14}$$

$$E_y = \frac{\beta}{\omega\epsilon}H_x, \;\; E_z = \frac{1}{i\omega\epsilon}\frac{\partial H_x}{\partial y} \tag{2.15}$$

The propagation characteristics of the TE and TM modes are determined by the eigenvalue equations arising from the field solutions and the boundary conditions at the interfaces $y = -T$ and $y = 0$. From Eq. 2.12, the field solutions of the TE mode are written in the form:

$$\begin{aligned} E_x &= E_c e^{-\gamma_c y}, \; y > 0 \text{ (in the cover)} \\ &= E_g \cos(k_y y + \phi), \; -T < y < 0 \text{ (in the guiding layer)} \\ &= E_s e^{-\gamma_s (y+T)}, \; y < -T \text{ (in the substrate)} \end{aligned} \tag{2.16}$$

where the propagation constants in the y-direction are expressed in terms of the effective index such that:

$$\gamma_c = k_0\sqrt{N^2 - n_c^2}, \;\; k_y = k_0\sqrt{n_g^2 - N^2}, \;\; \gamma_s = k_0\sqrt{N^2 - n_s^2} \tag{2.17}$$

Physically, these field solutions assume a guided wave (with a cosine or sine form depending on the value of the phase shift ϕ) whose energy exponentially decays into the substrate and cladding layers. Only the TE modes will be considered here as a similar analysis is used for the TM modes.

The tangential field components E_x and H_z must be continuous at the interfaces. Consequently, the corresponding boundary conditions of E_x and $\partial E_x/\partial y$ both being continuous across the interfaces yield at $y = 0$:

$$\begin{aligned} E_c &= E_g \cos\phi \\ -\gamma_c E_c &= -k_y Eg \sin\phi \end{aligned} \tag{2.18}$$

which combine to give $\tan\phi = \gamma_c/k_y$, and at $y = -T$:

$$\begin{aligned} E_s &= E_g \cos(k_y T - \phi) \\ \gamma_s E_s &= k_y E_g \sin(k_y T - \phi) \end{aligned} \tag{2.19}$$

which then combine to give $\tan(k_y T - \phi) = \gamma_s/k_y$. Combining all the above relations results in the following eigenvalue equation:

$$k_y T = (m+1)\pi - \tan^{-1}\left(\frac{k_y}{\gamma_s}\right) - \tan^{-1}\left(\frac{k_y}{\gamma_c}\right) \tag{2.20}$$

where $m = 0, 1, 2, \ldots$ denotes the mode number of each solution. This transcedental equation can be numerically solved for k_y, and consequently N, if the refractive indices of the waveguide layers and the guide layer thickness T are given. Before, we considered modes as the confined path along which light could propagate without radiating into the substrate or cladding. Now, we can consider modes to be the eigenvector solutions to the above equation for a specific index of refraction profile, with the eigenvalues representing the propagation constants, or effective indices, of the modes. The resulting electric field distributions $E_x(y)$ for the first three TE-polarized modes, $\mathrm{TE}_{m=0,1,2}$, of an arbitrary slab waveguide are illustrated in Figure 2.3. The mode subscript indicates the number of zero crossings, or nodes, of the field in the guiding layer.

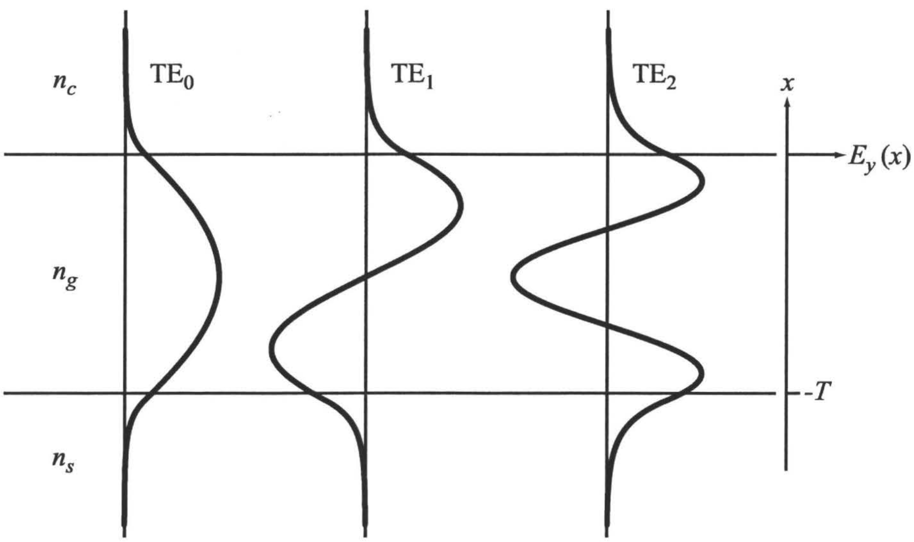

Figure 2.3: Electric field distributions of TE slab guided modes

It is convenient to apply the numerical evaluation of Eq. 2.20 using the following normalizations. The normalized frequency V and the normalized guide index

b_E can be defined as:

$$V = k_0 T \sqrt{n_g^2 - n_s^2} \tag{2.21}$$

and

$$b_E = (N^2 - n_s^2)/(n_g^2 - n_s^2) \tag{2.22}$$

The asymmetry of the waveguide is also defined by:

$$a_E = (n_s^2 - n_c^2)/(n_g^2 - n_s^2) \tag{2.23}$$

Using these definitions, the normalized form of Eq. 2.20 becomes:

$$V\sqrt{1 - b_E} = (m+1)\pi - \tan^{-1}\sqrt{\frac{1 - b_E}{b_E}} - \tan^{-1}\sqrt{\frac{1 - b_E}{b_E + a_E}} \tag{2.24}$$

The numerically derived normalized dispersion curves based on the above equation are shown in Figure 2.4. The effective index can then be obtained graphically if the waveguide parameters are known.

Light will couple into the substrate layer if the effective index of the mode is equal to the refractive index of the substrate ($N = n_s \rightarrow b_E = 0$). The cut-off value of the normalized frequency required for supporting guided modes is given by:

$$V_m = V_0 + m\pi, \quad V_0 = \tan^{-1}\sqrt{a_E} \tag{2.25}$$

where V_0 is the cut-off value for the fundamental mode (TE_0). If the normalized frequency V of the waveguide ranges over $V_m < V < V_{m+1}$, the TE_0, TE_1,..., and TE_m modes are supported and the number of guided modes is $m + 1$.

For the TM modes of a slab waveguide, the analysis is similar to the preceding,

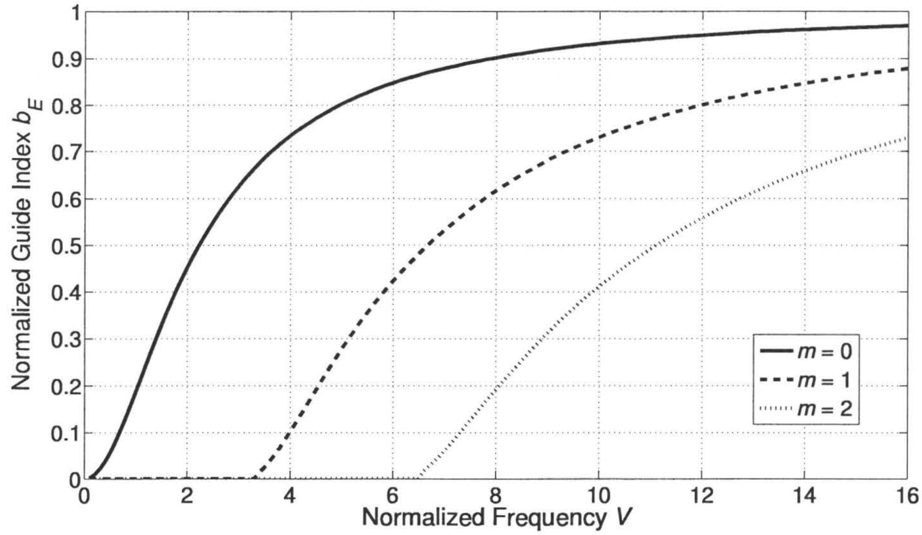

Figure 2.4: Dispersion curves of TE slab guided modes
For symmetric slab waveguides with $n_s = n_c \rightarrow a_E = 0$.

but with H_x and E_x continuous at the interfaces, yielding the following eigenvalue equation:

$$k_y T = (m+1)\pi - \tan^{-1}\left(\frac{n_s}{n_g}\right)^2\left(\frac{k_y}{\gamma_s}\right) - \tan^{-1}\left(\frac{n_c}{n_g}\right)^2\left(\frac{k_y}{\gamma_c}\right) \tag{2.26}$$

The normalized form of this equation is more complicated than that of the TE modes and is therefore not shown, but it can be found in [59].

2.3 Effective Index Method for Rib Waveguides

Most waveguide devices require both lateral and vertical optical confinement, which is not possible with a simple slab waveguide structure. This additional horizontal confinement can be achieved by using a three-dimensional structure such as the rib waveguide, which consists of sidewalls fabricated by removing material from the guide layer. The effective index method (EIM) can be used to obtain the approximate modes of a rib waveguide, as no analytical solution for this waveguide geometry exists. This method is illustrated in Figure 2.5 and has been adapted from [59].

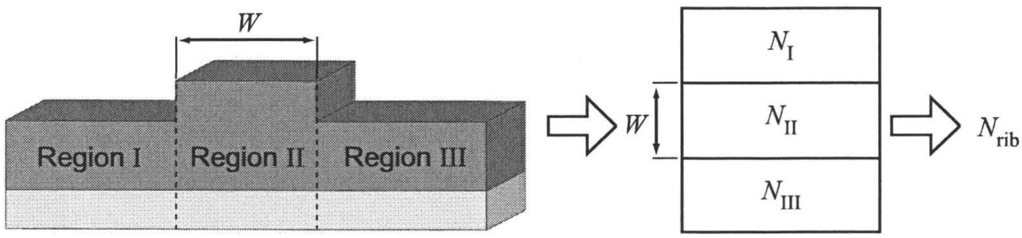

Figure 2.5: Analytical model for the effective index method

It is important to note that the modes of a rib waveguide are not purely TE or TM, as is the case with the slab waveguide. They represent two families of hybrid modes: TE-like modes, E^x_{pq}, with transverse field components E_x and H_y and TM-like modes, E^y_{pq}, with components E_y and H_x. These hybrid modes are essentially TEM modes ($E_z \doteq H_z \doteq 0$) that are polarized along either the x- or y-directions. Their notation has an additional subscript to count the nodes that will now occur along both the x- and y-axes.

The effective index method reduces the three-dimensional problem of a rib waveguide into two separable slab waveguide problems. Each region (I, II and III) are first treated as individual slab waveguides with optical confinement in the y-direction before solving for the overall effective indices of the rib waveguide modes.

To solve for the E^x_{pq} modes, we first solve for the effective indices of the TE_m

modes in each region. Recall that the effective index of a slab waveguide is found by numerically solving Eq. 2.20. Now the overall effective index can be solved for by considering the TM_m solutions of a slab waveguide whose material refractive indices are replaced with the effective indices that were just solved for in each region along the y-direction. To solve for the E^y_{pq} modes, the effective indices of the TM_m modes are first obtained in each region. Then the effective indices of the TE_m modes of the slab waveguide based on the TM_m effective indices are found.

The wave equation can also be explicitly solved for a given set of boundary conditions using a partial differential equation solver based on the finite element or finite difference methods. Basically, the system of equations characterizing the problem are typically expressed in matrix form before applying a numerical eigenvalue solution algorithm. The propagation constants of each mode can then be simply derived from the eigenvalues of the resulting solutions with their fields being defined by the corresponding eigenvectors. This method can be applied to any waveguide of arbitrary cross-section to solve for both its guided and radiation modes.

2.4 Beam Propagation Method

Solving for the optical modes of a given waveguide cross-section does not take into account variations in the refractive index profile along the propagation direction. A more versatile algorithm known as the beam propagation method (BPM) can solve for the varying optical field as it propagates by using a step-wise discretization. This technique is particularly useful when simulating waveguides with curved or tapered sections, but it can also be used to simulate the optical power transferred between waveguides whose modes overlap. The BPM algorithm used for this project was implemented by the BeamPROP simulation package from RSOFT Corporation. The derivation of this algorithm will be described here and is based on the work found in [61]-[64].

The BPM technique uses finite difference methods to solve the paraxial approximation of the well-known Helmholtz equation. By assuming a scalar field by neglecting polarization effects, and restricting the propagation to a narrow range of angles, the wave equation can be written in the form of the Helmholtz equation for monochromatic waves:

$$\frac{\partial^2 \psi}{\partial x^2} + \frac{\partial^2 \psi}{\partial y^2} + \frac{\partial^2 \psi}{\partial z^2} + k(x,y,z)^2 \psi = 0 \tag{2.27}$$

Here the scalar electric field has been written as $E(x,y,z,t) = \psi(x,y,z)e^{-i\omega t}$ using $k(x,y,z) = k_0 n(x,y,z)$, with $k_0 = 2\pi/\lambda_0$ being the wavenumber in free space. The geometry of the problem is defined entirely by the refractive index profile, $n(x,y,z)$.

The rapidly varying phase of the field ψ can be factored out by introducing a slow-varying field u such that: $\psi(x,y,z) = u(x,y,z)e^{i\bar{k}z}$. Here, $\bar{k}$ is a constant representing the average phase variation. This reference wavenumber is often expressed in terms of a reference refractive index, $\bar{n}$, such that $\bar{k} = k_0\bar{n}$. This changes the above Helmholtz equation to:

$$\frac{\partial u}{\partial z} = \frac{i}{2\bar{k}} \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + (k^2 - \bar{k}^2)u \right) \tag{2.28}$$

An implicit finite-difference (FD) approach based on the well-known Crank-Nicholson scheme can used to numerically solve this modified Helmholtz equation along the z-propagation direction. This FD-BPM algorithm represents the field in the transverse (xy) plane by discrete points on a grid, and at discrete planes along the propagation direction. It interpolates the known field solution from the previous plane to obtain the solution of the next unknown field using a step-wise discretization. The initial field should ideally be the solution of the $z = 0$ refractive index profile, which can be found using various mode solving algorithms. If we consider a simplified scalar field in 2-D by neglecting the y field component, we can apply the Crank-Nicholson scheme such that Eq. 2.28 is represented at the midplane between the known plane n and the unknown plane $n + 1$ by:

$$\frac{u_i^{n+1} - u_i^n}{\Delta z} = \frac{i}{2\bar{k}} \left(\frac{\delta^2}{\Delta x^2} + \left[k(x_i, z_{n+1/2})^2 - \bar{k}^2 \right] \right) \frac{u_i^{n+1} + u_i^n}{2} \tag{2.29}$$

where u_i^n is the field at a transverse grid point i and longitudinal plane n, with a grid spacing of Δx and the planes spaced apart by Δz. Here, the δ^2 second order difference operator is defined by: $\delta^2 u_i = (u_{i+1} + u_{i-1} - 2u_i)$, and $z_{n+1/2} \equiv z_n + \Delta z/2$. This equation can then be expressed in a standard tridiagonal matrix form and solved numerically. The transparent boundary condition (TBC), which allows radiation to freely escape the computational domain, would be typically used for this solution.

Since this simple 2-D scalar algorithm does not describe propagation behavior in all directions, the more complex semi-vectorial FD-BPM algorithm was applied in all three dimensions to accurately predict coupling behavior between waveguides. Its lengthly derivation can be found in [61]. The full-vectorial method was not used as it is only necessary when coupling between polarizations is expected, as would be the case in a polarization converter for example. The semi-vectorial method is adequate to accurately simulate each polarization state separately.

This BPM technique can also be modified into a mode solving algorithm, as proposed in [65] and [66]. This so-called imaginary distance-beam propagation method (ID-BPM) replaces the longitudinal coordinate z with $z' = jz$, resulting in an exponential growth of the individual modes. The growth rate of each mode is equal to its real propagation constant, β_m. An arbitrary field, such as a Gaussian, is launched into a structure with a constant refractive index profile along the z-axis and is then propagated along the imaginary axis. The field will evolve into the fundamental mode ($m = 0$) because it has the highest propagation constant resulting in the highest growth rate, dominating all other modes. Higher-order modes can then be obtained using an orthogonalization procedure to subtract contributions from lower order modes while performing subsequent propagations.

2.5 Single-Mode SOI Rib Waveguides

Most waveguide-based devices require only one mode to be guided since device operation depends on the value of its effective index, which varies for each mode. Such waveguides are said to be single-mode. In the case of our stacked waveguide couplers, the presence of higher-order modes will decrease the overall coupling efficiency since each mode will completely transfer its optical power at a different propagation distance. It is important to note that the term single-mode actually means that two modes will be guided, one for each orthogonal polarization state.

As was shown with the slab waveguide, the number of guided modes depends on the normalized frequency, which is a function of slab thickness and the refractive index contrast between the guide and substrate layers. Referring back to Figure 2.4, higher-order modes will be prohibited from propagating by sufficiently lowering the thickness and index contrast. The same is true for optical fibers, which are designed to be single-mode by using a small core diameter and a low-index contrast between the core and cladding.

This poses a problem for SOI rib waveguides, which have a high-index contrast and relatively large cross-sections. It was often assumed that the cross-sectional dimensions of a three-dimensional waveguide must be similar to the thickness of a single-mode slab waveguide, for the same material system, in order to propagate only the lowest-order mode. This would restrict the guide layer of an SOI rib waveguide to a thickness of 0.3 μm, which would prevent low-loss coupling of light from single-mode fibers that have core diameters of $\sim$8 μm. Soref was the first to propose that single-mode SOI rib waveguides could be fabricated with cross-sectional dimensions of several microns [67].

The principle behind large cross-section single-mode SOI waveguides is that for certain rib dimensions, higher-order modes will couple into the external slab regions where they become radiation modes. If the rib width is below a certain cut-off value, relative to the slab thickness, higher-order modes will be forced out of the rib leaving only the fundamental modes to be guided. The electric fields for both the TE and TM-like fundamental modes for a single-mode SOI rib waveguide are shown in Figure 2.6. These fields were obtained using the ID-BPM algorithm as implemented by the BeamPROP simulation package. Note that the field discontinuities are horizontal along the material interfaces for the TM-polarized mode, but perpendicular to them for the orthogonal TE-polarized mode.

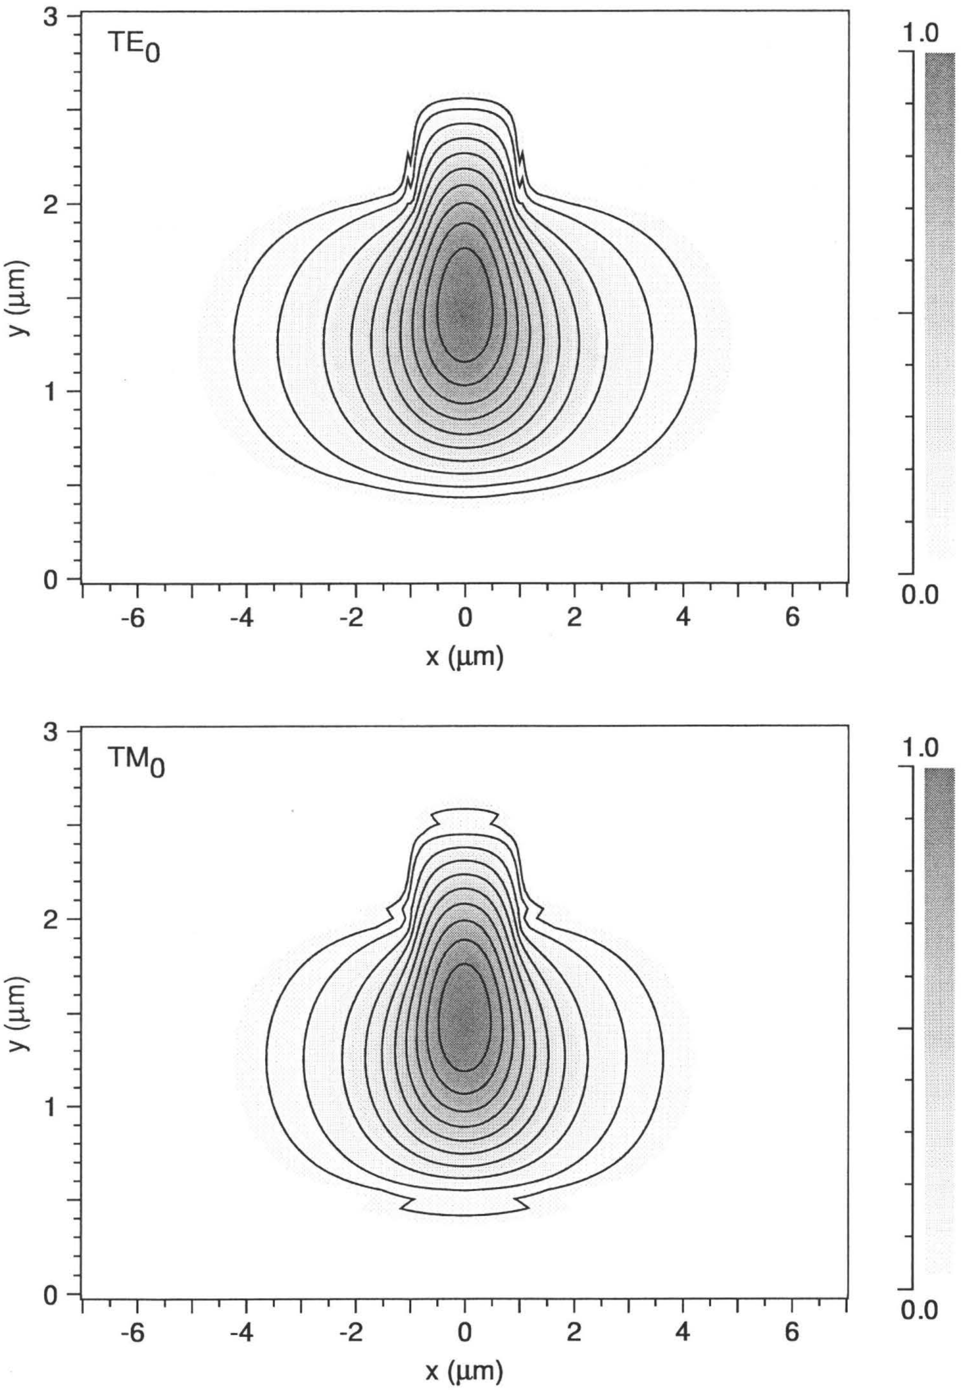

Figure 2.6: Fundamental modes of a single-mode SOI rib waveguide for the TE and TM polarizations

Referring to the parameters illustrated in Figure 2.7, the single-mode condition was analytically derived by Soref to be:

$$\begin{aligned}\frac{W}{H} &\leq \left(\frac{q+4\pi b}{4\pi b}\right)\frac{1+c\sqrt{\left(\frac{q+4\pi b}{q+4\pi r b}\right)^2-1}}{\sqrt{\left(\frac{q+4\pi b}{q+4\pi r b}\right)^2-1}} \\ q &= \frac{\gamma_0}{\sqrt{n_1^2-n_0^2}}+\frac{\gamma_2}{\sqrt{n_1^2-n_2^2}}\end{aligned} \tag{2.30}$$

with $\gamma_{0,2} = 1$ for the E^x_{pq} modes, $\gamma_{0,2} = (n_{0,2}/n_1)^2$ for the E^y_{pq} modes and c is a constant which was numerically found to be 0.3.

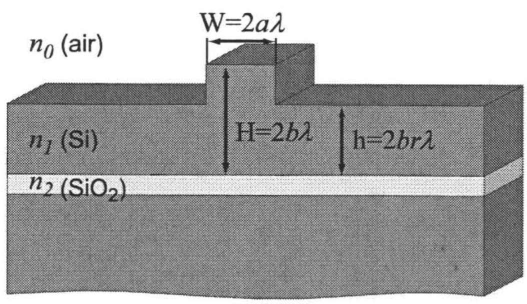

Figure 2.7: SOI waveguide parameters for single-mode condition

Equation 2.30 reduces to:

$$\frac{W}{H} \leq c + \frac{r}{\sqrt{1-r^2}} \tag{2.31}$$

in the limit of large b ($4\pi b \gg q$). The variable r is the ratio of the slab thickness to the guide layer thickness (h/H). If this single-mode condition is not satisfied, then higher-order modes, such as the one shown in Figure 2.8, will also be guided.

Soref restricted this single-mode condition to rib waveguides with $r > 0.5$,

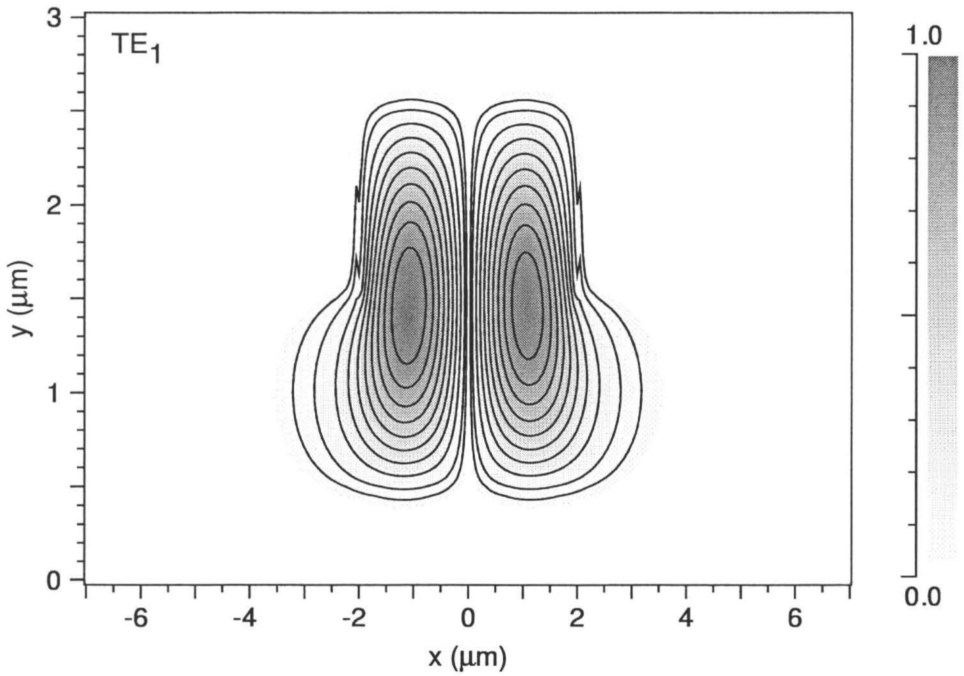

Figure 2.8: First higher-order mode of a multimode SOI rib waveguide

meaning the guide layer is not etched beyond half of its thickness. To explain this additional requirement, we will consider the second-order modes which have a double-lobed intensity distribution along the vertical direction. For $r > 0.5$, the lower lobe will radiate out into the fundamental slab modes of the outer regions. If the waveguide is etched beyond $r = 0.5$, the external slab modes may be too thin for sufficient radiation to occur. Therefore, waveguides with $r > 0.5$ are guaranteed to be single-mode in the vertical direction. The single-mode condition is satisfied in the horizontal direction if the rib width is sufficiently narrow such that the lobes of higher-order modes are cut-off horizontally.

The single-mode condition of SOI rib waveguides was studied experimentally by Rickman [68] and using the effective index method by Pogossian [69]. Pogossian found that the constant c of Soref's formula (Eq. 2.31) should be lowered to a more restrictive value of zero based on the EIM results or to -0.05 based on Rickman's experimental results. The single-mode condition for both $c = 0$ and $c = 0.3$ are plotted in Figure 2.9.

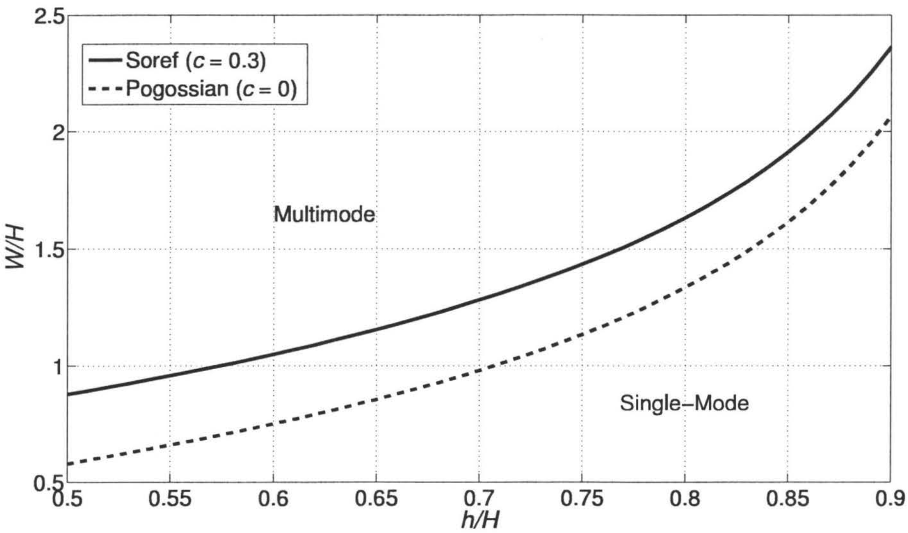

Figure 2.9: Single-mode conditions for SOI rib waveguides

A mode solver can be used to confirm the single-mode condition for a particular rib waveguide geometry, as was done for this project. If the effective indices of the higher-order modes fall below the effective index of the fundamental slab mode of the side regions they will radiate out of the rib [63]. Only the fundamental mode will propagate since its effective index is higher than the effective index of the external fundamental slab mode. Visual inspection of the optical field profiles will also indicate which rib geometries will confine only the single lobe of the fundamental mode.

2.6 Coupled Mode Theory

We have shown that various guided modes can exist in a waveguide. These are normal modes defined by the waveguide structure and its boundary condition. The orthogonal relationship between these modes dictates that they each propagate without mutual coupling, carrying power independently. Yet, if these modes are slightly perturbed by the modes of another nearby waveguide they are no longer independent, resulting in coupling between the two waveguides. Coupled mode theory can be used to describe the coupling behavior in a perturbed waveguide system by first knowing the normal modes of each unperturbed structure.

2.6.1 Coupled Mode Equations

Consider two single-mode waveguides, I and II, each with normal modes a and b respectively. If these waveguides are sufficiently separated, by some distance x_s, the fields of each mode, ψ_a and ψ_b, will propagate independently with constants β_a and β_b. The waveguides become a coupled system when this separation is reduced such that the two fields overlap. Both scenarios are illustrated in Figure 2.10.

Here, the light is exclusively launched into waveguide I. This light will propagate uncoupled from waveguide II if there is sufficient separation between them, as is the case in Figure 2.10 a). The perturbed scenario in Figure 2.10 b) results in optical power being transferred between the two waveguides as coupling occurs along the propagation direction. It is possible for the power to be completely transferred to waveguide II after propagating a certain distance. This perturbed system no longer has individual waveguide modes, but rather a superposition of them resulting in normal odd and even supermodes with fields ψ_o and ψ_e. These supermodes are mu-

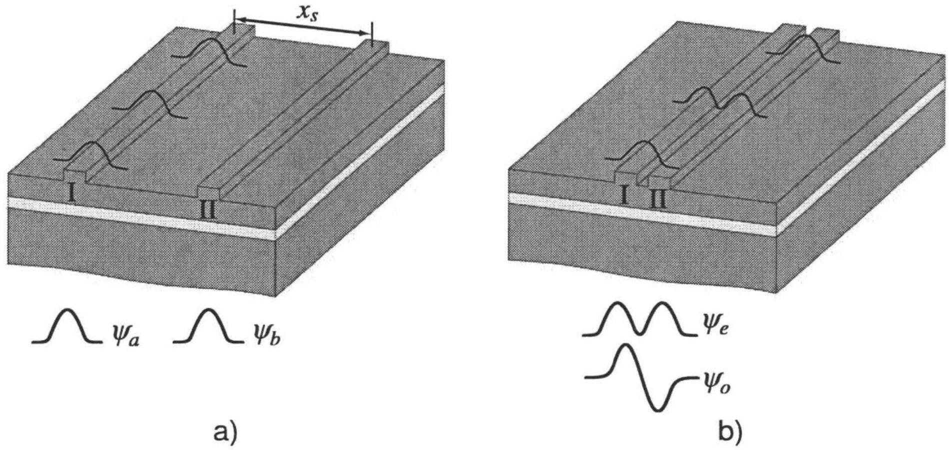

Figure 2.10: Directional rib waveguide optical coupler
a) Optically isolated waveguides without cross-coupling b) Coupling between waveguides with sufficient optical field overlap

tually coupled, with propagation constants of β_o and β_e. Coupled mode theory can be used to obtain these supermodes, and their corresponding beat period, and will be described following the derivation in [59].

We begin by expressing the individual optical fields of each coupled waveguide by:

$$\begin{cases} \psi_a(x,y,z,t) = A(z)e^{-j\beta_a z}f_a(x,y)e^{i\omega t} \\ \psi_b(x,y,z,t) = B(z)e^{-j\beta_b z}f_b(x,y)e^{i\omega t} \end{cases} \tag{2.32}$$

where f_a and f_b are the normalized field distribution functions. If no coupling is present, $A(z)$ and $B(z)$ will reduce to constants. With coupling present along the propagation z-axis they are no longer independent, and the relationship between them is given by the coupled mode equations:

$$\begin{cases} \pm\frac{dA(z)}{dz} = -j\kappa_{ab}B(z)e^{-j(\beta_b-\beta_a)z} & (\beta_a \neq 0) \\ \pm\frac{dB(z)}{dz} = -j\kappa_{ba}A(z)e^{+j(\beta_b-\beta_a)z} & (\beta_b \neq 0) \end{cases} \tag{2.33}$$

where κ_{ab} and κ_{ba} are the coupling coefficients between the two modes. The exponential terms correspond to the phase-constant mismatching of the individual modes. The coupling coefficient is found by integrating the spatial overlap of the field distribution functions over the dielectric constant increments $\Delta\varepsilon$ in waveguide II such that:

$$\kappa_{ab} = c \int_{II} f_a^* \Delta\varepsilon f_b \, dx \, dy \tag{2.34}$$

where c is a constant related to the normalization of ψ_a and ψ_b. This integration procedure is illustrated in Figure 2.11.

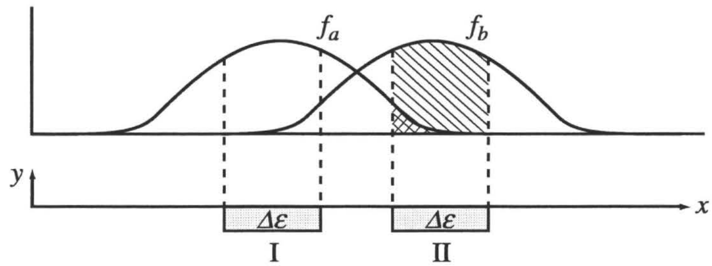

Figure 2.11: Optical field overlap for coupled mode theory

2.6.2 Coupling Length and Power Transfer

In the case where both modes are propagating in the same direction ($\beta_a > 0$, $\beta_b > 0$), the coupled mode equations in 2.33 reduce to:

$$\begin{cases} A(z) = Ae^{-\gamma z}e^{-j\Delta z} \\ B(z) = Be^{-\gamma z}e^{+j\Delta z} \end{cases} \tag{2.35}$$

where we have defined: $2\Delta \equiv \beta_b - \beta_a$, $\gamma = \pm\sqrt{\kappa^2 + \Delta^2}$, and $B/A = \kappa/(\gamma - \Delta)$. It can shown that $\kappa_{ab} = \kappa^*_{ba}$, both of which are set to a positive real value of κ for simplicity.

The solutions are then rewritten as:

$$\begin{cases} A(z) = \left[A_e e^{-j\sqrt{\kappa^2+\Delta^2}z} + A_o e^{+j\sqrt{\kappa^2+\Delta^2}z}\right] e^{-j\Delta z} \\ B(z) = \left[\frac{\kappa A_e}{\sqrt{\kappa^2+\Delta^2}-\Delta} e^{-j\sqrt{\kappa^2+\Delta^2}z} + \frac{-\kappa A_o}{\sqrt{\kappa^2+\Delta^2}+\Delta} e^{+j\sqrt{\kappa^2+\Delta^2}z}\right] e^{+j\Delta z} \end{cases} \tag{2.36}$$

By inserting Eq. 2.36 into Eq. 2.32, the following is derived:

$$\begin{aligned} \psi_a(x,y,z,t) &= [A_e e^{-j\beta_e z} + A_o e^{-j\beta_o z}] f_a(x,y) e^{j\omega t} \\ \psi_b(x,y,z,t) &= [B_e e^{-j\beta_e z} + B_o e^{-j\beta_o z}] f_b(x,y) e^{j\omega t} \end{aligned} \tag{2.37}$$

using the following relations:

$$\beta_e,\ \beta_o = \beta_m \pm \beta_c \tag{2.38}$$

$$\beta_m \equiv (\beta_a + \beta_b)/2, \ \ \beta_c \equiv \sqrt{\kappa^2 + \Delta^2} \tag{2.39}$$

The fields of the uncoupled system, ψ_a and ψ_b, are now expressed as a linear combination of the two modes with propagation constants β_o and β_e.

If light is exclusively launched into waveguide I, the boundary conditions are $A(0) = 1$ and $B(0) = 0$ at $z = 0$, and $A(z)$ and $B(z)$ become:

$$\begin{cases} A(z) = e^{-j\Delta z}\left(\cos\beta_c z + j\frac{\Delta}{\beta_c}\sin\beta_c z\right) \\ B(z) = e^{+j\Delta z}\frac{-j\kappa}{\beta_c}\sin\beta_c z \end{cases} \tag{2.40}$$

which can be expressed in the normalized power forms as:

$$\begin{cases} \frac{|A(z)|^2}{|A(0)|^2} = 1 - F\sin^2\beta_c z \\ \frac{|B(z)|^2}{|A(0)|^2} = F\sin^2\beta_c z \end{cases} \tag{2.41}$$

where

$$F \equiv \left(\frac{\kappa}{\beta_c}\right)^2 = \frac{1}{1+(\Delta/\kappa)^2} \tag{2.42}$$

which is a measure of the maximum coupling efficiency achievable. The power transferred between the two coupled waveguides will beat periodically, as plotted in Figure 2.12.

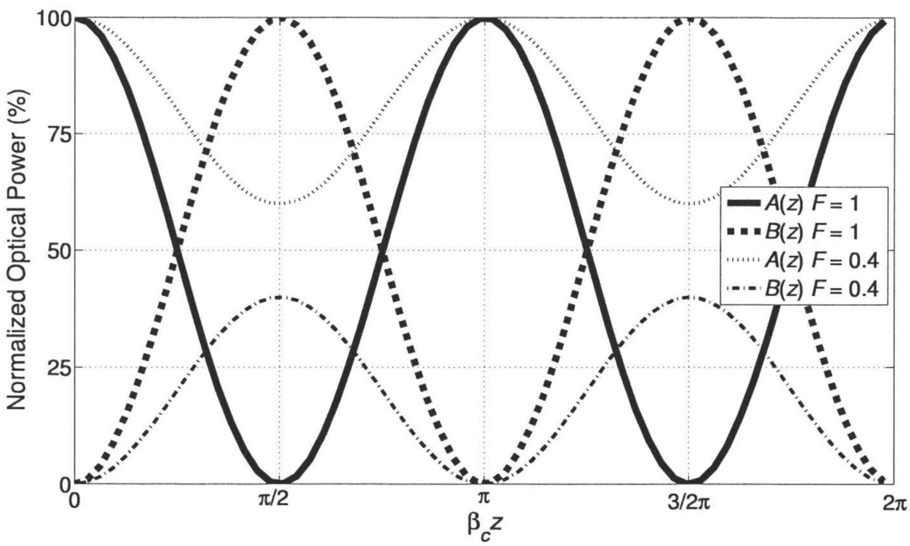

Figure 2.12: Power transfer between coupled waveguides

The location at which maximum power is transferred is given by the beat length:

$$L_\pi = \frac{\pi}{2\beta_c} \tag{2.43}$$

which reduces to:

$$L_\pi = \frac{\pi}{2\kappa} \tag{2.44}$$

when the two modes are synchronized such that $\beta_a = \beta_b \rightarrow F = 1$. This means that complete power transfer between coupled waveguides will only occur if both individual waveguides have the same propagation constants, usually achieved through matching refractive index profiles. The contradirectional coupling case, where the two fields are travelling in opposite directions, will not be shown here. It is usually reserved for describing the behavior of Bragg grating structures that produce backward travelling reflected waves.

The coupling coefficient will not be directly calculated in this project. Rather, the coupler length required to achieve a π-phase shift between the supermodes is found using:

$$L_\pi = \frac{\pi}{|\beta_0 - \beta_1|} = \frac{\lambda_0}{2(|N_0 - N_1|)} \tag{2.45}$$

where N_0 and N_1 are the effective indices of the first two supermodes as calculated by an ID-BPM mode solver. This coupling length will also be found using conventional BPM. The input field launched into one of the waveguides will initially excite the supermodes, that will then couple as they propagate. The normalized power in each waveguide will be plotted to graphically determine the optimal coupler length.

2.6.3 General Theory of Mode Coupling

The preceding derivation only considered coupling between two modes, whose fields were described by a scalar. The fields are actually vectors and all modes, including radiation modes, should be considered as well. A more rigorous general theory of mode coupling is provided in [59]. Its derivation will not be provided here, but its applicable results will be highlighted.

This general theory has been used to derive the coupling coefficient expression in Eq. 2.34, and for the justification that $\kappa_{ab} = \kappa_{ba}^*$ as well. It also proves that the linear combination of perturbed modes, referred to here as supermodes, form an orthonormal set. This orthogonality is important in understanding how the supermodes interfere given a certain phase shift accumulated between them.

Referring back to Figure 2.10, the orthogonality of the supermodes is illustrated by the even field having two positive lobes, while the odd field is comprised of lobes with opposite signs. The superposition of these modes at $z = 0$ results in the constructive interference of the lobes in waveguide I, while the lobes destructively interfere in waveguide II. This means that all of the power will be located in waveguide I. The signs for the field lobes of the odd supermode will invert after propagating a distance, L_π, such that a π-phase has accumulated between the two supermodes. Now, the superposition of these modes will result in all of the power being confined in waveguide II. If the same amount of power is launched into both waveguides, the signal in each will be completely exchanged at $z = L_\pi$.

2.7 Multimode Interferometers

While it was previously stated that waveguides supporting multiple modes were usually undesirable, they can be exploited to realize power splitting devices, such as the one illustrated in Figure 2.13. Here, a single-mode rib waveguide launches light into a wide multimode rib, known as a multimode interferometer (MMI). The number of modes supported by the MMI depends on its width, W_{MMI}. Beating between the numerous modes that have been excited results in different interference conditions along the propagation axis. The input optical field will be equally split into two separate lobes after propagating a certain distance, where it can then be collected by two single-mode rib waveguides. The theory behind this behavior will now be explained, following the derivation given in [24].

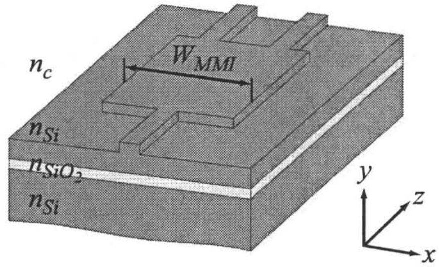

Figure 2.13: MMI-based power splitter in SOI

Consider a multimode slab section of width W_{MMI}, with a guide index of n_g. This multimode waveguide supports ν lateral modes along the x-axis that are identified by the mode numbers $m = 0,\ 1,\ \ldots\ (\nu - 1)$. The y-axis will be ignored in this discussion for simplicity. The lateral wavenumber k_{xm} and the propagation constant β_m of each mode are related to the guide index by the following dispersion equation:

$$k_{xm}^2 + \beta_m^2 = k_0^2 n_g^2 \tag{2.46}$$

with $k_0 = 2\pi/\lambda_0$ and $k_{xm} = (m+1)\pi/W_{em}$. Here, W_{em} is the effective width of the MMI section, taking into consideration the penetration of the modal fields into the substrate and cover layers. For high-index contrast waveguides, such as those fabricated in SOI, the penetration depth is small, such that this effective width is approximated as the width of the MMI itself ($W_{em} \simeq W_{MMI}$).

Using a binomial expansion with $k_{xm}^2 \ll k_0^2 n_g^2$, the propagation constant β_m can be found from Eq. 2.46 to be:

$$\beta_m \simeq k_0 n_g - \frac{(m+1)^2 \pi \lambda_0}{4 n_g W_{MMI}^2} \tag{2.47}$$

We can then define L_π as the beat length of the two lowest-order modes such that:

$$L_\pi \doteq \frac{\pi}{\beta_0 - \beta_1} \simeq \frac{4 n_g W_{MMI}^2}{3 \lambda_0} \tag{2.48}$$

allowing the spacing of the propagation constants to be written as:

$$(\beta_0 - \beta_m) \simeq \frac{m(m+2)\pi}{3 L_\pi} \tag{2.49}$$

Consider an input field profile $\Psi(x,0)$ launched into the MMI slab at $z = 0$. We can decompose this field into the modal field distributions $\psi_m(x)$ of all the modes using:

$$\Psi(x,0) = \sum_{m=0}^{\nu-1} c_m \psi_m(x) \tag{2.50}$$

where c_m are the field excitation coefficients, representing the amount of input light that is coupled into each mode of the MMI section. The field profile at a distance z can then be written as a superposition of all the guided mode field distributions such that:

$$\Psi(x,z) = \sum_{m=0}^{\nu-1} c_m \psi_m(x) e^{j(\beta_0 - \beta_m)z} \tag{2.51}$$

with the phase of the fundamental mode as a common factor taken out of the sum. We can find the expression of the field at a distance $z = L$ by substituting Eq. 2.49

into Eq. 2.51 to obtain:

$$\Psi(x, L) = \sum_{m=0}^{\nu-1} c_m \psi_m(x) \exp\left[j\frac{m(m+2)\pi}{3L_\pi}L\right] \tag{2.52}$$

By inspecting Eq. 2.52, it can been seen that $\Psi(x, L)$ will be an image of $\Psi(x, 0)$, if

$$\exp\left[j\frac{m(m+2)\pi}{3L_\pi}L\right] = 1 \quad \text{or} \quad (-1)^m \tag{2.53}$$

This is commonly referred to as self-imaging. The first condition means that the phase changes of all the modes along L must differ by integer multiples of 2π. In this case, all of the guided modes will combine with the same relative phases as they did at $z = 0$, resulting in a direct replica of the input field. The second condition means that the phase changes must be alternatively even and odd multiples of π. In this case, the even modes will be in phase and the odd modes in antiphase. This type of interference will also produce a replica of the input field, but it will be mirrored with respect to the $x = 0$ plane. The first and second conditions are satisfied at:

$$L = p(3L_\pi) \text{ with } p = 0, \ , 1, \ 2, \ \ldots \tag{2.54}$$

for p even and p odd, respectively. This self-imaging behavior is illustrated in Figure 2.14.

Now consider the images obtained half-way between the direct and mirrored image positions, i.e., at distances:

$$L = \frac{p}{2}(3L_\pi) \text{ with } p = 1, \ 3, \ 5, \ \ldots \tag{2.55}$$

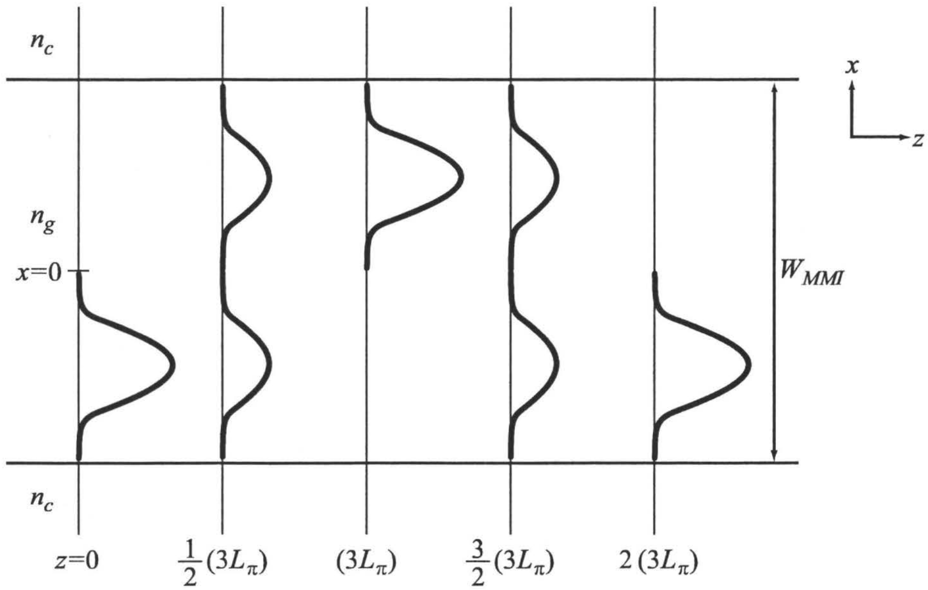

Figure 2.14: Optical field self-imaging behavior in an MMI

We will find that the images are two-fold, with each half mirrored with respect to the $x = 0$ plane. This distance provides the length required for the MMI coupler to split the input optical power equally into two output images. This two-fold imaging can be used to realize 2×2 3-dB couplers, which are typically used for power splitting and switching applications. In general, multi-fold images are formed at intermediate distances between those given in equations 2.54 and 2.55. This can be exploited to create $N \times M$ couplers.

In this project, we will fabricate the MMI slab such that the x and y-axes of the above derivation are inverted. Self-imaging will now occur along the vertical direction. These vertically-integrated MMI couplers will either completely transfer, or split, incoming light between the upper and lower regions of the multimode section.

2.8 Active Silicon Photonic Device Theory

Silicon photonic devices, such as modulators and switches, can be actively-controlled by changing the refractive index of the guide layer. Consider a modulator based on a Mach-Zehnder interferometer that splits the signal light into two separate waveguide arms. A phase shift is induced between the two waveguide modes by changing the refractive index in one of the arms. These modes will constructively or destructively interfere upon recombining in a second splitter, depending on the amount of phase shift induced. Light can therefore be modulated or switched by varying the refractive index in one of the waveguide arms. The phase shift of an interferometer is found using:

$$\Delta\phi = \frac{2\pi}{\lambda_0}\Delta n L \tag{2.56}$$

where L is the length of the interferometer arm over which the refractive index change Δn is applied.

Change in the refractive index of any material can be achieved using various physical effects, but not all of these are available in silicon. The Pockels effect is commonly used in III-V semiconductor and lithium niobate devices. Here, change in the refractive index has a linear dependence on the strength of an applied electric field. This effect relies on the asymmetry of the crystal structure and is therefore absent in bulk, unstrained silicon due to its centrosymmetric crystal structure. Recently, a linear electro-optic effect has been demonstrated in silicon material that had been asymmetrically strained by silicon nitride layers [70]. The Kerr effect, where the change in the refractive index has a quadratic dependence on the strength of the applied electric field, is available in any material but is generally a weak effect. In silicon, the refractive index changes by only 10^{-4} orders of magnitude even with a very large electric field of $\sim 10^6$ V/cm, which is above the breakdown voltage of lightly doped silicon [71].

2.8.1 Current-Injection P-I-N Diodes

The change in refractive index for active SOI devices is commonly achieved using the free carrier plasma dispersion effect. This is a significantly slower process than the Pockels effect since it involves the diffusion of carriers. Carrier injection into an SOI rib waveguide is typically realized using a forward-biased p-i-n diode, as illustrated in Figure 2.15.

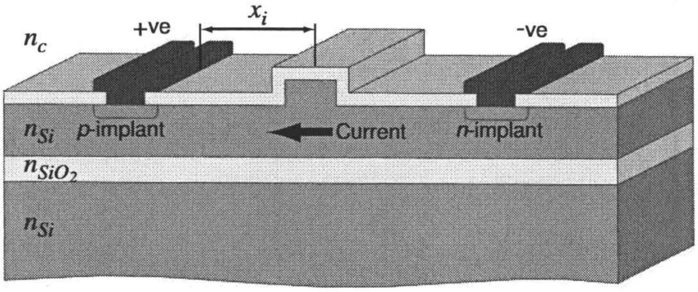

Figure 2.15: Current injection diode for the tuning refractive index of an SOI waveguide

A voltage-induced change in the free carrier concentration of electrons, ΔN_e, or holes, ΔN_p, will lead to a change in the refractive index of silicon, Δn, given by:

$$\Delta n_e = -8.8\times10^{-22}\Delta N_e \tag{2.57}$$

$$\Delta n_p = -8.5\times10^{-18}(\Delta N_p)^{0.8} \tag{2.58}$$

where the electron and hole concentration changes are in units of carriers/cm^3.

These expressions were derived for the telecommunications wavelength of 1550 nm using experimental absorption spectra and Kramers-Kronig analysis by Soref in [72]. Since the electro-refraction due to free carriers is inseparable from electro-absorption, the change in the index of refraction is usually limited to $\sim 10^{-3}$ in order to have a waveguide with a reasonable loss ($\sim$5 dB/mm).

The dopant regions and metal contacts of the p-i-n diode must be placed sufficiently far away from the optical mode to avoid optical absorption. This distance, shown here as x_i, is generally >10 μm for 5 μm thick SOI rib waveguides. The maximum switching speed achievable will decrease the further the dopant regions are separated, therefore it is desirable to use highly-confined waveguide structures to minimize this separation, as proposed in [73].

These p-i-n diode-based devices are generally limited to switching speeds of $\sim$20 MHz due to the slow carrier generation and/or recombination processes involved in accumulating free carriers [74]. A metal-oxide-semiconductor (MOS) capacitor-based device can accumulate carriers along its gate oxide with far greater switching speeds since it relies on field effects induced by the applied voltage. Modulation in SOI waveguide devices using MOS capacitor phase shifters can exceed frequencies of 1 GHz, as demonstrated by Intel Corporation [5].

2.8.2 Thermo-optic Devices

The refractive index can also be changed using the thermo-optic effect, without the absorption associated with carrier injection. This process is relatively slow, typically on the order of milliseconds, as it relies on the slower diffusion of thermal energy. The change in the refractive index of silicon is a linear function of temperature given by [13]:

$$\left(\frac{\mathrm{d}n}{\mathrm{d}T}\right)_{\mathrm{Si}} = +1.86\times10^{-4}\ \mathrm{K}^{-1} \tag{2.59}$$

Note that the change in the refractive index is positive for the thermo-optic effect while it was negative for the free carrier effect. These processes can compete with each other in a diode-based device if thermal effects are not carefully considered.

This temperature change can be achieved by driving current through a thin metal strip heater placed above the silicon waveguide. A buffer layer of oxide is placed in between to avoid optical absorption by the metal layer. Relatively fast thermo-optic switching in SOI waveguides has been demonstrated in [75]. A simplistic 1-D heat flow model was used that predicted the switching power required for a π-phase shift to be:

$$P_{\pi} = \lambda_0 \kappa_{\mathrm{SiO_2}} \left(\frac{W_H}{t_{\mathrm{SiO_2}}} + 0.88\right) \left(\frac{\mathrm{d}n}{\mathrm{d}T}\right)_{\mathrm{Si}}^{-1} \tag{2.60}$$

where $\kappa_{\mathrm{SiO_2}} = 1.4$ W/(m·K) is the thermal conductivity of the buffer oxide layer, $t_{\mathrm{SiO_2}}$ is the thickness of this oxide, and W_H is the width of the heater. A cut-off frequency of 49 kHz was observed when a heater power of 46 mW was used to induce a π-phase shift in one of the modulator arms. This phase shift was achieved by changing the temperature of the SOI waveguide by 8 K over 700 μm of length. Submicrosecond thermo-optic switching has also been demonstrated in [76].

2.8.3 Laser-Induced Free Carrier Generation

The free carrier plasma dispersion effect can also be achieved by generating carriers through laser excitation. Light will be absorbed by the silicon if the photon energy of the incident laser beam is greater than the band gap, generating electron-hole pairs through interband absorption. Since silicon has an indirect band gap of 1.1

eV, any light with a wavelength below 1.1 μm will be absorbed, along with the emission of phonons. The free carriers have a finite lifetime before recombining, releasing energy in the form of heat. During this lifetime, the free carriers will diffuse away from the region where they were created.

An all-optical switch using a nitride bus waveguide to couple 632 nm-HeNe laser light into a silicon waveguide below has been proposed in [77]. This light is absorbed by the silicon waveguide, changing its index of refraction through free carrier generation. The laser-induced free carrier concentration, N_c [cm^{-3}], as a function of space and time is expressed as:

$$\frac{\partial N_c(x,z,t)}{\partial t} = G_n + D\left[\frac{\partial N_c(x,z,t)}{\partial x^2} + \frac{\partial N_c(x,z,t)}{\partial z^2}\right] - \frac{N_c(x,z,t) - N_{c0}}{\tau_R} \quad (2.61)$$

where G_n is the free carrier generation rate [cm^{-3}s^{-1}] produced by optical excitation, D is the ambipolar carrier diffusion coefficient, N_{c0} is the initial carrier concentration, and τ_R is the recombination rate. The derivation of this carrier concentration equation can be found in [78]. Note that this equation only solves for the carriers generated along the surface xz-plane, as we assume that the concentration will be uniform over the shallow depth of the silicon device layer. This is justified by the diffusion length of >100 μm for bulk silicon being much larger than the thickness of the SOI device layer [79]. We also assume that there is equal generation of both electrons and holes. This partial differential equation will be solved later using MATLAB.

An Ar-ion laser beam at a wavelength of 514 nm will be used to generate free carriers in the upper waveguide layer of devices with two vertically-integrated silicon guide layers separated by a thin oxide layer in between. The heat generated from recombination processes is assumed to uniformly diffuse between the two guide layers at steady-state, while the remaining free carriers will be confined to the upper layer by the electrically insulating oxide in between. This means that the change in the refractive index will depend only on the excess of free carriers in the upper layer, allowing thermal effects to be neglected.

Various processes can occur in silicon that contribute to the overall recombination rate, including: Shockley-Read-Hall (SRH), radiative, Auger, and surface recombination. SRH recombination, involving deep-level traps, is a relatively slow process on the order of 10 ms in pure crystalline silicon and can be neglected. Radiative recombination in silicon is a minor process because of its indirect band gap and Auger recombination, involving energy transfer to a third carrier, plays only a significant role at carrier densities $>1\times10^{18}$ cm^{-3} [80]. However, surface recombination is a significant effect in thin semiconductor layers, dominating the determination of carrier lifetime. The lifetime due to surface recombination is given by:

$$\tau_{sur} = \frac{1}{D}\frac{d^2}{\pi} + \frac{d}{2S} \tag{2.62}$$

where d is the film thickness, S is the surface recombination velocity and D is the carrier diffusion coefficient. The typical recombination rates in bulk silicon and SOI wafers will discussed in the following chapter.

The free carrier generation rate used in Eq. 2.61 is defined by:

$$G_n = \frac{\alpha_T}{h\nu}\frac{P_{abs}}{\pi\omega_0^2}\exp\left(-2\frac{r^2}{\omega_0^2}\right) \tag{2.63}$$

where α_T is the total absorption coefficient, $h\nu$ is the photon energy, P_{abs} is the power absorbed by the silicon device layer, ω_0 is the laser beam radius for a $1/e^2$

intensity, and $r = \sqrt{x^2 + z^2}$ [78]. The total absorption coefficient depends on the inherent material absorption coefficient, α, and the absorption cross-section for free carriers, σ, and is defined by: $\alpha_T = \alpha + \sigma N_c(x, z, t)$. The free carrier cross-section is expressed by a semi-empirical formula: $\sigma = 1 \times 10^{-12} T \lambda^2$, where T is the material temperature and λ is the wavelength of the incident laser beam.

The amount of laser light absorbed by the silicon device layer depends on the fraction of light reflected by the upper air-silicon interface and by the buried silicon-oxide interface, along with the absorption coefficient of the silicon layer. The reflectivity at each interface depends on the real and imaginary parts of the refractive indices of the materials, n and κ, and is found using the Fresnel equation:

$$R_1 = \frac{(n_{Si} - 1)^2 + \kappa^2}{(n_{Si} + 1)^2 + \kappa^2}, \quad R_2 = \frac{(n_{SiO_2} - n_{Si})^2 + \kappa^2}{(n_{SiO_2} + n_{Si})^2 + \kappa^2} \tag{2.64}$$

The light transmitted through the upper silicon layer into the one below is found using: $T_{Trans} = (1 - R_1)e^{-\alpha_T H}(1 - R_2)$, where H is the guide layer thickness. The remaining light is assumed to be absorbed completely by the silicon layer such that: $P_{abs} = (1 - R_1 - T_{Trans})P_0$, where P_0 is the incident power of the laser beam. The carrier generation equation will be solved for a varying absorbed optical power, and these results can then be correlated to a specific guide layer thickness and incident laser power.

The refractive index change with respect to the 1550 nm light propagating in the upper waveguide is found using the free carrier plasma dispersion equations 2.57 and 2.58. The changes in the free electron and hole concentrations, ΔN_e and ΔN_p, are assumed to be equal, with each found using: $N_c - N_{c0}$.

Chapter 3

Device Design and Simulation

Various vertically-integrated optical coupler designs have been simulated to determine their lengths required for optimal coupling efficiency. In this chapter, the simulation results for these proposed designs will be presented along with their tolerance to fabrication errors. The feasibility of actively controlling the coupling to achieve a vertically-integrated optical switch will also be investigated.

3.1 Stacked Slab Waveguide Coupling

A pair of stacked slab waveguides is the most basic example of a vertically-integrated optical coupler. This device was demonstrated before advancing to more complicated coupler designs. The substrate structure and simulation parameters for coupling between slab waveguides are shown in Figure 3.1. The devices for this project were fabricated using two different double SOI substrates: one with silicon layers of thickness $H = 5$ μm, separated by a $t_{BOX} = 0.05$ μm thick oxide layer, and the other with 2.5 μm silicon layers separated by a 0.1 μm oxide. The buried oxide thickness of these wafers is 0.5 μm and 0.375 μm respectively. Vertical coupling between the slab layers of these substrates was simulated using a 3-D beam propagation method (BPM), implemented by the BeamPROP software package from RSOFT Corporation.

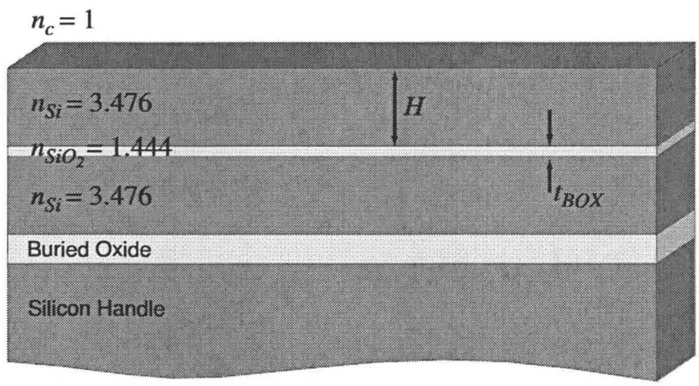

Figure 3.1: Parameters for stacked slab waveguide coupling simulations

Experimentally, it will be desirable to only launch light into the lower slab layer to accurately determine the amount of coupling between the stacked guide layers. The fabricated device will consist of multiple single-mode rib waveguides in the bottom layer, with the layer above having been removed. These input waveguides will excite the lower slab mode of the double slab coupler section, which will be fabricated with varying lengths.

This proposed device structure is illustrated in Figure 3.2, with two of the input waveguides shown. Either one can be used to launch light into the double slab sections of differing length. If a single layer slab waveguide had been instead used for the input, its varying length would have resulted in a different excitation condition for each double slab section. This is due to the inherent lateral spreading of slab modes. The lateral confinement of the rib waveguides ensures that the input mode dimensions remain consistent.

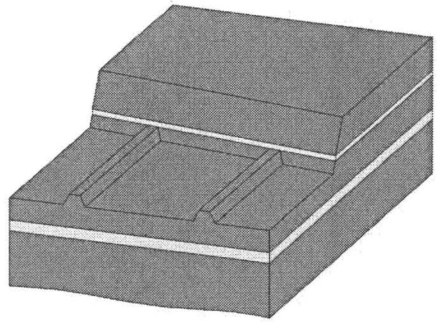

Figure 3.2: Device structure for slab waveguide coupling demonstration

The fundamental mode of the input rib waveguide was computed and then launched into the lower guide of the double slab coupler for the following BPM simulations. The resulting coupling of optical power between the slab layers as a function of coupler length is shown in Figure 3.3, for the substrate with $H = 5$ μm and $t_{BOX} = 0.05$ μm.

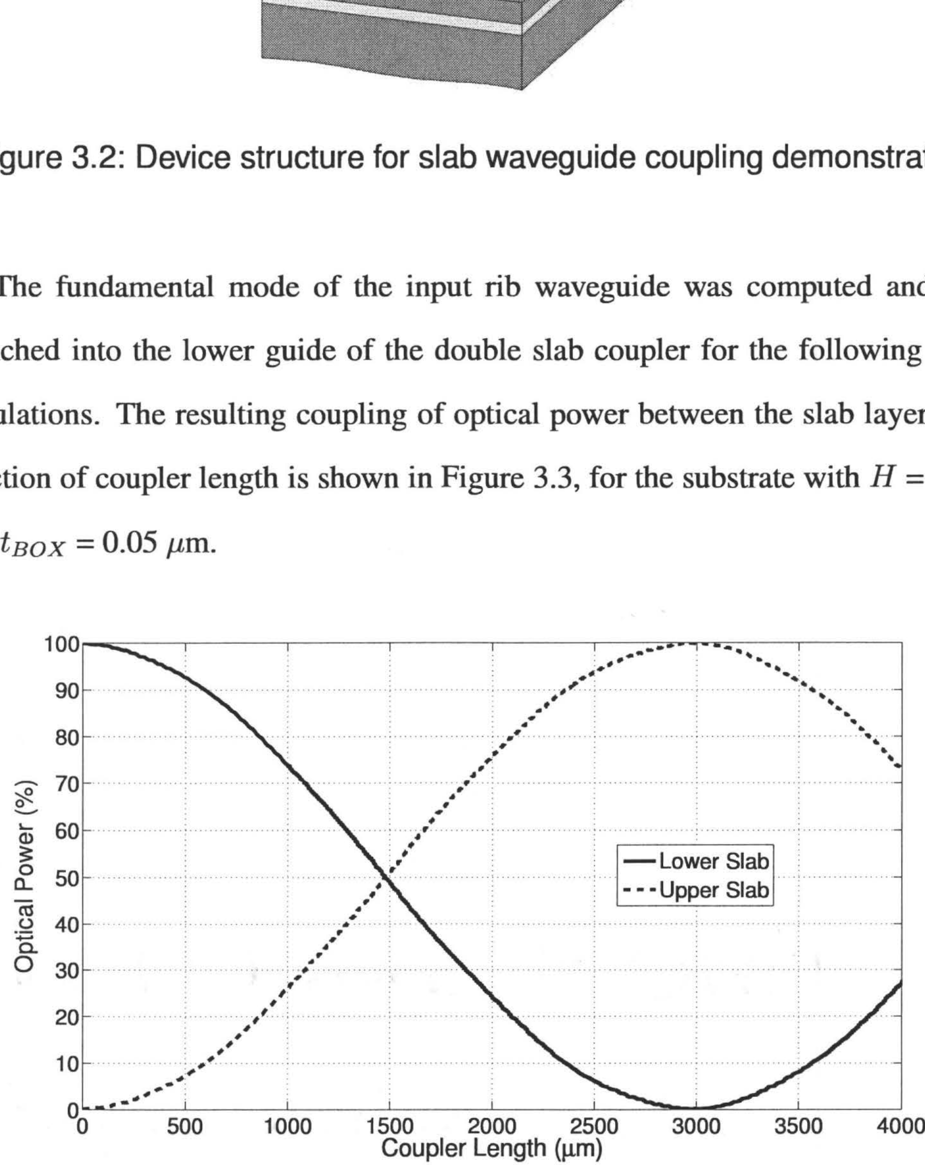

Figure 3.3: Coupling of optical power between stacked slab waveguides
For slab waveguides of thickness $H = 5$ μm with a separating oxide of thickness $t_{BOX} = 0.05$ μm.

This plot reveals that the input optical power is completely transferred to the upper waveguide after propagating a length of 2985 μm. This coupling length was simulated using TE-polarized input light with a wavelength of 1550 nm and an air cover layer ($n_c = 1$). The coupler length reduces to 834 μm for the substrate with H = 2.5 μm and $t_{BOX} = 0.1$ μm. The optical modes of thinner guide layers have more light propagating outside the silicon slab, resulting in a larger optical field overlap between the modes of the two stacked slab waveguides.

The simulated coupling follows the expected sinusoidal dependence on propagation distance as predicted by coupled mode theory. The normalized transferred power into the upper waveguide as a function of propagation length, L, is given by: $P(L) = \sin^2(\pi L/2L_\pi)$, where L_π is the coupler length required for complete power transfer.

The coupling length was found to be significantly polarization dependent. This arises from the different orientations of the optical field discontinuities for the TE and TM polarization modes. The discontinuities of the TM modes run parallel to the material layers, intersecting the plane between the overlapping modes from each waveguide. This decreases the overall optical field overlap for TM-polarized input light, thus increasing the coupling length to 15996 μm for the H = 5 μm substrate.

3.2 Proposed Vertically-Integrated Optical Coupler Designs

Coupling between stacked waveguide layers of silicon can be achieved with various types of vertically-integrated devices. While there are several types of waveguides that can be employed, the rib waveguide is the most common optical confinement structure used in silicon photonics. These guides can be stacked either slab-to-slab, rib-to-rib, or rib-to-slab. Integrating multiple layers would require alternating between the first two options, while the rib-to-slab configuration can be stacked consecutively. These coupler designs are illustrated in Figure 3.4 along with two other designs that will also be simulated.

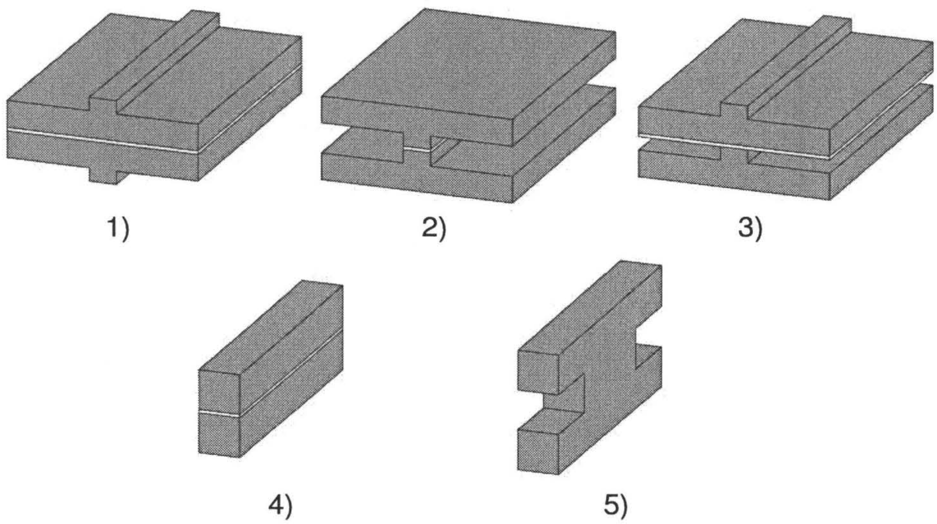

Figure 3.4: Proposed vertically-integrated coupler designs
Stacked rib waveguide couplers: 1) slab-to-slab 2) rib-to-rib 3) rib-to-slab 4) Stacked channel waveguide coupler 5) Vertically-integrated multimode interferometer coupler

The parameters used in simulating the stacked rib waveguide coupler are illustrated in Figure 3.5. The refractive index of the silicon guide layers is $n_{Si} = 3.476$, while the indices of the cover and separating oxide layers are both $n_{SiO_2} = n_c = 1.444$, for the input wavelength of $\lambda_0 = 1550$ nm. The rib waveguide dimensions for all simulations were chosen to ensure single-mode propagation, as predicted by [67].

The supermodes of an example slab-to-slab stacked rib waveguide coupler were computed using the imaginary distance-beam propagation method (ID-BPM) algorithm of the BeamPROP software package. The effective indices of these modes

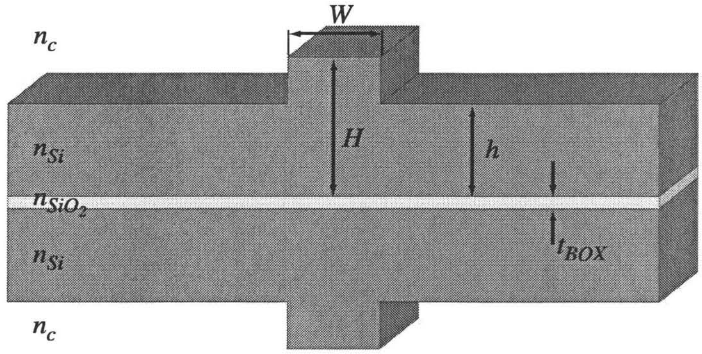

Figure 3.5: Simulation parameters for stacked rib waveguide couplers

can be used to calculate the required coupling length to accumulate a π-phase shift between them, as predicted by coupled mode theory. This length can then be compared to that found using conventional BPM for the same structure. The dimensions for this simulated coupler are $H = W = 2$ μm, $h = 1.5$ μm, and $t_{BOX} = 0.05$ μm.

The electric fields of the two computed supermodes are plotted in Figure 3.6, with their orthogonality illustrated by their respective three-dimensional contour plots. The zero-order supermode has two positive magnitude lobes, while the first-order supermode has two lobes of opposite sign. The coupling of these supermodes will result in optical power being transferred between the upper and lower waveguides.

Consider optical power launched exclusively into the upper rib waveguide. The overlapping of the supermodes at the zero propagation distance results in constructive interference of the upper lobes of each mode, while the lower lobes will destructively interfere. The opposite will occur after propagating a distance that introduces a π-phase shift between the supermodes, which inverts the magnitudes of the lobes for the first-order supermode. This phase shift occurs since each mode propagates

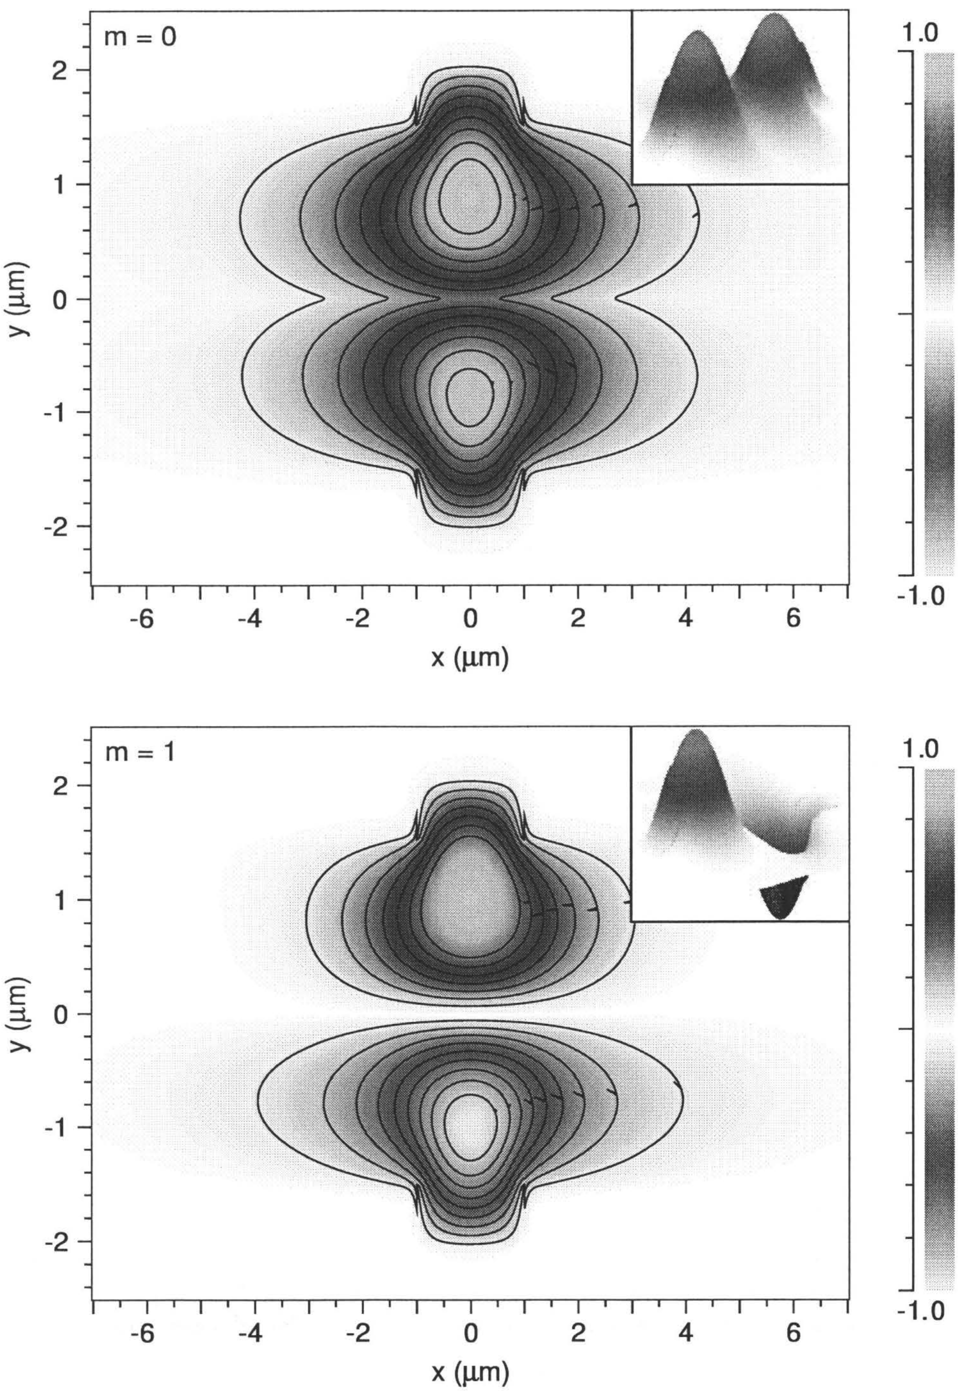

Figure 3.6: Supermodes of a stacked rib waveguide coupler
For stacked rib waveguides with $H = 2\ \mu\text{m}$, $h = 1.5\ \mu\text{m}$, $W = 2\ \mu\text{m}$ and $t_{BOX} = 0.05\ \mu\text{m}$

with its own effective index of refraction. At this length the input light will have completed its coupling from the upper to the lower waveguide. Intermediate lengths will result in a distribution of the optical power between the two guides, with the ratio corresponding to the accumulated phase shift.

The fundamental mode of the upper rib waveguide alone was first computed using ID-BPM and then launched into the stacked waveguide structure to solve for its supermodes using the same algorithm. The effective index of the zero-order supermode was found to be N_0 = 3.45323, while the index of the first-order supermode is N_1 = 3.44811. The number of decimal places shown is indicative of the typical accuracy of the mode solving simulation. The coupler length required to achieve a π-phase shift between these supermodes is $L_\pi = \lambda_0/2(|N_0 - N_1|) = 152 \pm 2\ \mu$m, for a TE-polarized input with a wavelength of λ_0 = 1.55 μm. This value matches the length found using BPM, confirming that either simulation technique can be used.

The zero-order supermode for the TM input polarization is shown in Figure 3.7. This plot illustrates the discontinuities of the electric field that run along the material boundaries. As with the slab waveguide couplers, these discontinuities result in a reduced optical field overlap, increasing the coupling length for the TM polarization to 701 $\pm$ 9 μm. If a fixed coupler length of 152 μm is chosen based on the TE polarization results, a TM-polarized input will only couple 9.5 % into the lower waveguide. This polarization sensitivity is a detriment if the input polarization state is random, but can also be exploited to control optical switching between layers.

The remaining simulations were completed using BPM since the mode solving algorithm did not always converge when simulating stacked waveguide structures. Here, the fundamental mode profile of the upper rib waveguide alone is first obtained using the mode solver, which was unconditionally stable with only one

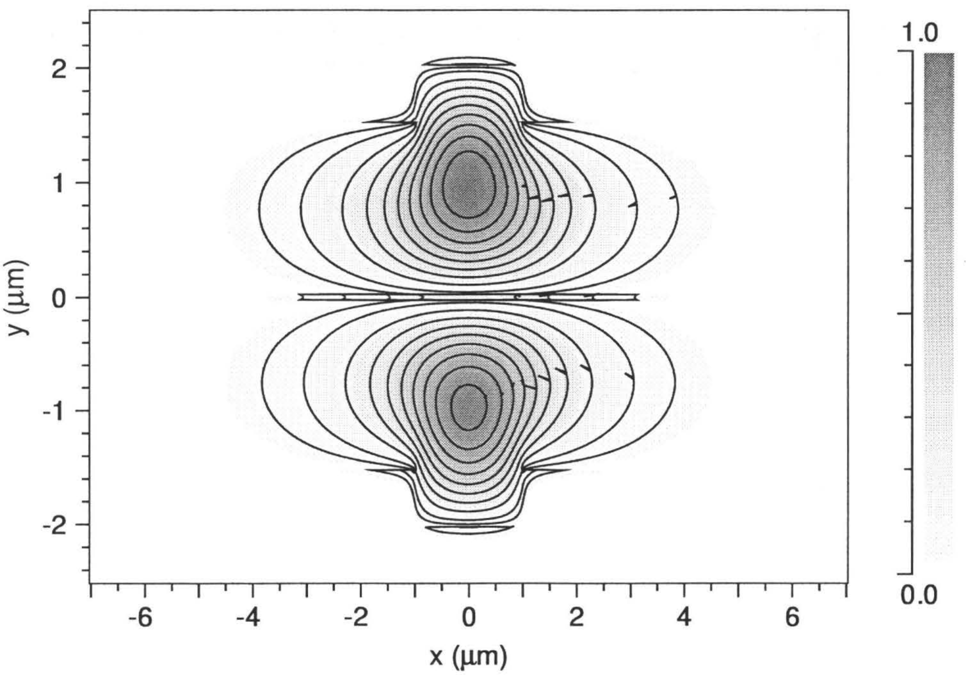

Figure 3.7: Zero-order supermode of a stacked rib waveguide coupler for a TM-polarized input

waveguide present. This mode profile is then launched into the stacked waveguide structure and propagated over a distance sufficiently long to determine the coupler length, L_π, required for a π-phase shift. The resulting plots of optical coupling as a function of coupler length will all resemble that of the slab coupling results shown in Figure 3.3. The reader is referred to this plot since only the location of maximum power transfer will vary.

A non-uniform simulation grid was used for the 3-D semi-vectorial BPM algorithm to ensure accurate results in a reasonable computation time. The grid density was gradually increased until the computed coupling length converged to the same result. The grid lines were ideally spaced by 0.1 μm along the x-axis, with 0.01 μm spacing near material interfaces, and by 0.005 μm along the y-axis, with 0.002 μm near material interfaces. A finer grid was chosen along the y-axis since the coupler modes will evolve primarily in the vertical direction. As well, a minimum of 20 vertical grid lines were fit within the separating oxide layer. A propagation step size of 0.05 μm along the z-axis was sufficient to obtain an accurate result.

The coupling lengths for various waveguide geometries of each proposed coupler design are shown in Figure 3.8. Here, the thickness of the device layer and separating oxide are fixed at $H = 2$ μm and $t_{BOX} = 0.05$ μm, while the slab thickness h and the rib width W are varied. The input is TE-polarized with a wavelength of 1550 nm and the cladding layers of both waveguides are oxide.

To compare the coupling length required for each structure, consider the same rib waveguide parameters of $H = 2$ μm, $h = 1.5$ μm, $W = 2$ μm and $t_{BOX} = 0.05$ μm. The coupling lengths for the slab-to-slab, rib-to-rib and rib-to-slab stacked rib waveguide couplers were found to be 154 μm, 466 μm, and 339 μm, respectively. The slab-to-slab coupler requires a shorter length than the rib-to-rib structure due to a larger optical field overlap. These coupling lengths increase to 695 μm, 1647 μm, and 1237 μm for a TM-polarized input.

It is important to note that the rib-to-slab coupler did not achieve complete power transfer at $L_\pi = 339$ μm. It was limited to 97.5 % due to the asymmetry of the cladding for each rib waveguide. While the rib waveguides of the slab-to-slab and rib-to-rib couplers are mirrored about the separating oxide, the cladding of the rib-to-slab coupler differs for each waveguide. Coupled mode theory dictates that the refractive index profile of both waveguides be identical to completely couple light at L_π. The maximum optical coupling achievable was shown to increase with increasing rib width, up to 98.8 % for $W = 2.5$ μm, but falls below 50 % for widths less than 1 μm. The practicality of the rib-to-slab design is questionable due

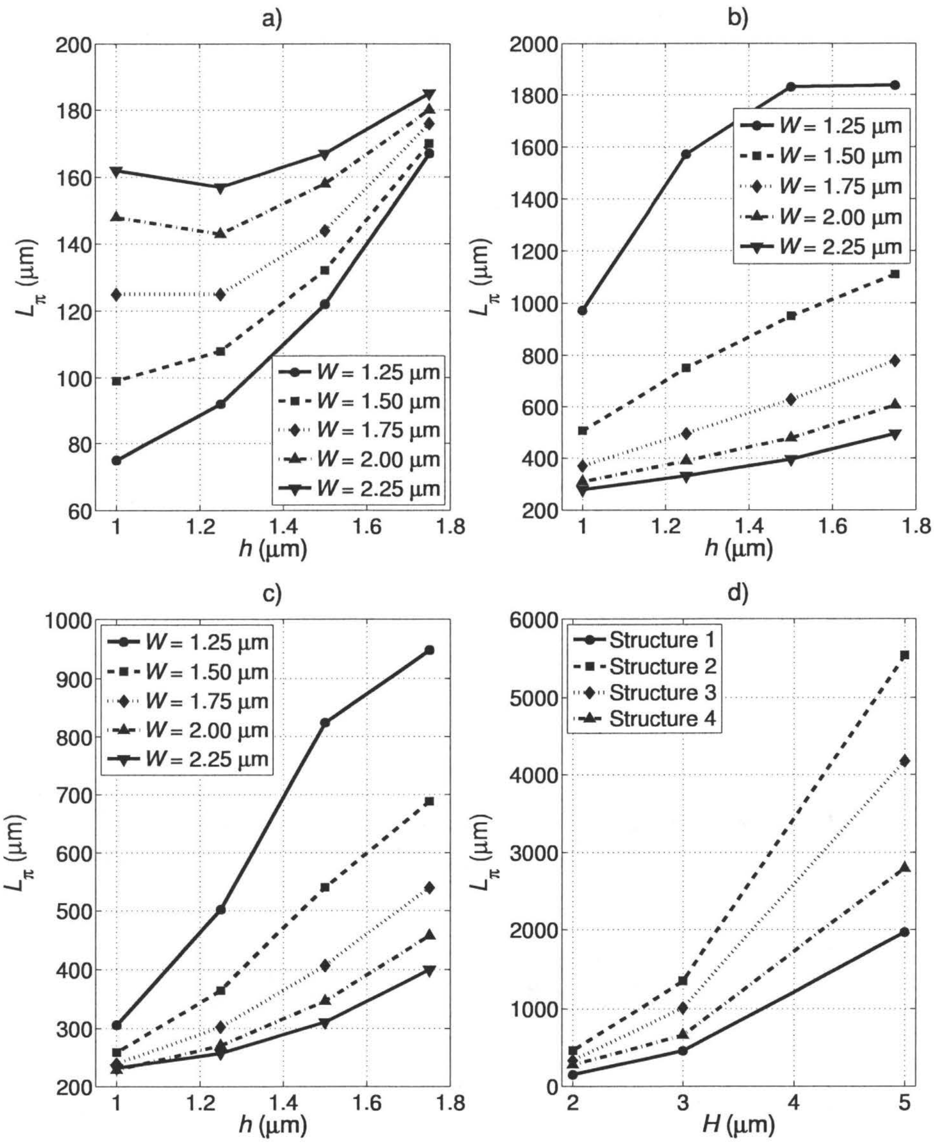

Figure 3.8: Coupling length for varying waveguide geometry

Stacked rib waveguide couplers: a) slab-to-slab b) rib-to-rib c) rib-to-slab d) Varying device layer thickness for structures 1-4 to this inherent limitation.

Another type of vertically-integrated coupler is the design consisting of stacked square channel waveguides, shown in Figure 3.4 4). While these channels are inherently multimode for the simulated dimensions of H = 2 μm, h = 0 μm, W = 2 μm and t_{BOX} = 0.05 μm, it is possible to use tapered mode converters to only excite their fundamental modes. The coupling lengths were found to be 221 μm for a TE-polarized input and 946 μm for the TM polarization. This coupler design does not require alternating structures for multiple guide layers, as is the case with the stacked rib waveguide couplers, thereby maintaining a constant coupling length for each layer.

The dimensions of these channel structures can be reduced to those of nanowires ($\sim$200 nm), to ensure both single mode operation and a very short coupling length on the order of tens of microns. Such dimensions are impractical if current aligned bonding techniques are used in their manufacturing process, since the width of the guide itself is smaller than the best case alignment tolerance of ± 1 μm.

The required coupling length for each design increases with increasing guide layer thickness for a fixed separating oxide layer, as plotted in Figure 3.8 d). The thicker guide layers reduce the evanescent field propagating outside the silicon, resulting in a decrease in the optical field overlap. Here, the width and slab thickness of each guide is scaled to the guide layer thickness, maintaining the same ratios of the H = 2 μm designs.

3.2.1 Tapered Mode Converter for Multimode Waveguide Couplers

As mentioned previously, a tapered mode converter will be necessary in order to excite only the fundamental mode of one of the stacked multimode channel waveguides. If higher order modes are excited, each will have its own coupling length required for a π-phase shift, reducing the overall optical coupling at a fixed coupler length. The proposed mode converter, which has been demonstrated in [81], linearly tapers the slab region of a single-mode rib waveguide from a starting width, W_{slab}, to the width of the channel waveguide, W, and is illustrated in Figure 3.9.

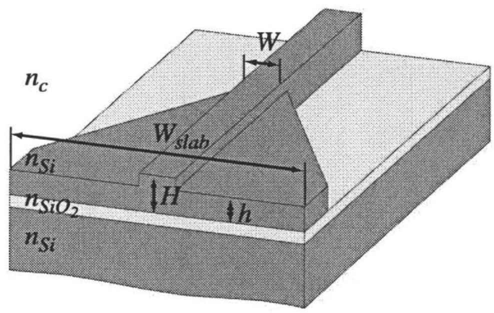

Figure 3.9: Mode converter design for stacked channel couplers

The dimensions for the simulated mode converter were: $H = 2$ μm, $h = 1.5$ μm, $W_{slab} = 15$ μm and $W = 2$ μm, with an oxide cladding and a TE-polarized input. The fundamental mode of the input rib waveguide was launched and gradually tapered into the channel waveguide over various lengths. The converted mode was allowed to propagate in the channel waveguide for an additional 50 μm to visually observe if higher order modes had been excited. This was the case until the converter length was increased above 150 μm. The transmitted power remained constant at 99.95 %, even at the relatively short converter length of 50 μm. This required converter length for single-mode excitation was also confirmed for a TM-polarized input.

3.2.2 Fabrication Tolerance Analysis

The coupling efficiency of an optical coupler will decrease if fabrication errors result in different dimensions for the two coupling waveguides, as predicted by coupled mode theory. The tolerance of each coupler design to fabrication errors was evaluated using BPM. Here, TE-polarized input light is launched into the upper waveguide, then coupled into a lower waveguide whose dimensions are varied slightly from the optimal design. The coupling efficiency is then measured at the coupler length previously determined at which complete power transfer between perfectly matching waveguides is obtained.

Device tolerance to deviations in the lower waveguide is shown in Figure 3.10 for the proposed coupler designs. Here, Designs 1-4 refer to the slab-to-slab, rib-to-rib, rib-to-slab and stacked channel waveguide couplers, respectively. The waveguide parameters chosen for each design were: $H = 2\ \mu$m, $h = 1.5\ \mu$m, $W = 2\ \mu$m and $t_{BOX} = 0.05\ \mu$m. The stacked channel coupler has a zero slab thickness. Their respective optimal coupler lengths are: 154 μm, 466 μm, 339 μm and 221 μm, for a TE-polarized input.

The decrease in coupling efficiency can be attributed to both the difference in the effective indices of the individual waveguide modes and to the change in coupling length. Coupled mode theory predicts that this difference in effective indices will reduce the maximum power that can be coupled due to the modes being unsynchronized, while the different coupling length means that a π-phase shift will not be obtained after propagating the designed coupler length. This is illustrated in Figure 3.11, with the optical power coupled between the upper and lower waveguides plot-

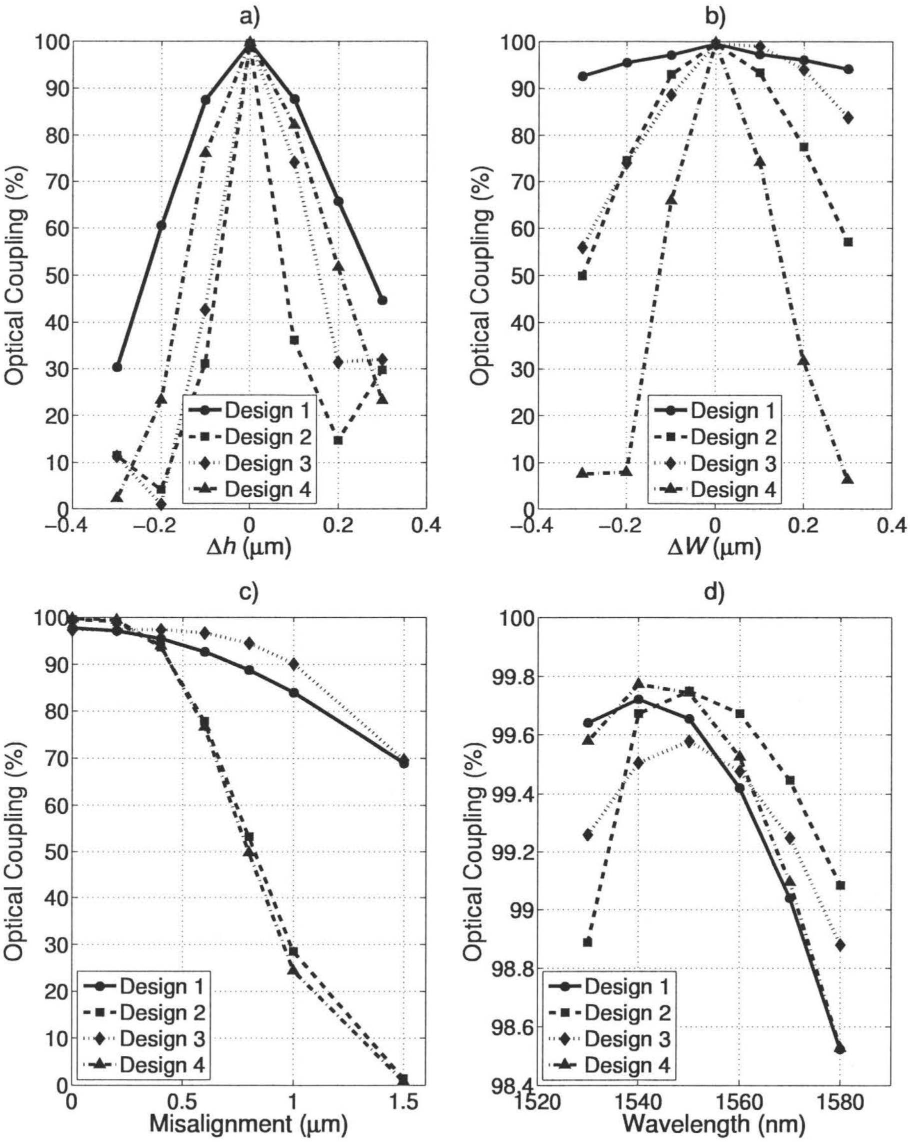

Figure 3.10: Optical coupling tolerance to deviations between stacked waveguides

a) Deviation in slab thickness b) Deviation in waveguide width c) Misalignment of waveguides d) Deviation in input wavelength

ted as a function of coupler length for the slab-to-slab stacked rib waveguide design with variations in slab thickness of $\Delta h = 0.1$ μm and 0.3 μm.

Deviations in the slab thickness can arise from both the non-uniformity in the

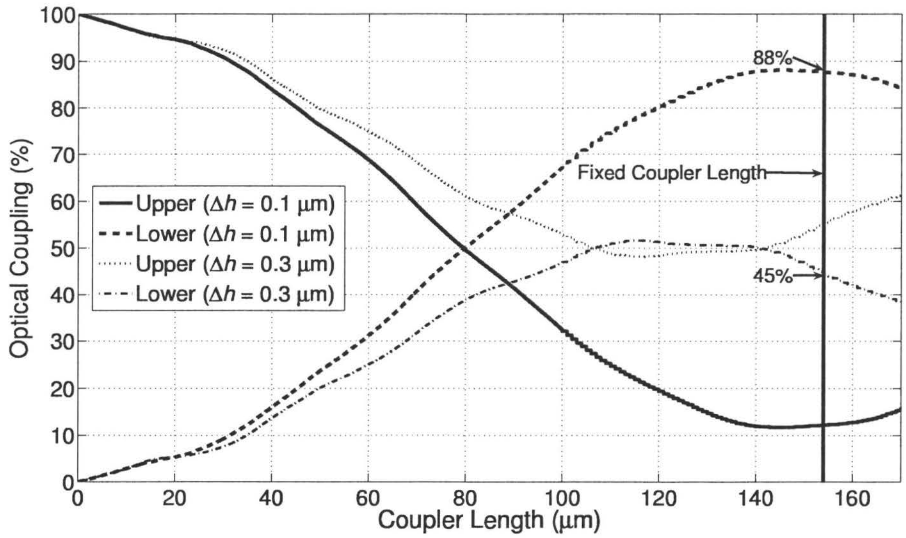

Figure 3.11: Optical coupling in a slab-to-slab stacked rib coupler with variations in slab thickness

silicon guide layer thickness of the starting material and from etch depth imprecision during fabrication. It is recommended that these devices be fabricated using SIMOX SOI wafers, which have device layers that are uniform within 2 %. The device layers of bond and etch-back SOI wafers have typical non-uniformities of ± 0.5 μm, which would result in an unacceptable coupling efficiency for these $H =$ 2 μm thick designs.

The waveguide width can differ between waveguides due to variations in the photolithographic process used to pattern the ribs. Our contact UV-photolithography process can pattern 2 μm lines within $\pm$ 0.2 μm, while modern CMOS fabrication facilities would have precision to tens of nanometers.

Each waveguide must be transversely aligned in a precise manner to ensure optimal coupling efficiency. The designs with larger optical field overlap areas involving slab regions exhibited the greatest tolerance to misalignment of the lower waveguide. The fabricated devices for this project will have multiple sets of couplers with varying intentional misalignments to compensate for shifts of up to 2 μm between the stacked waveguides. All of the devices exhibited a slight sensitivity to changes in the input wavelength over the telecommunications L band covering 1530-1565 nm.

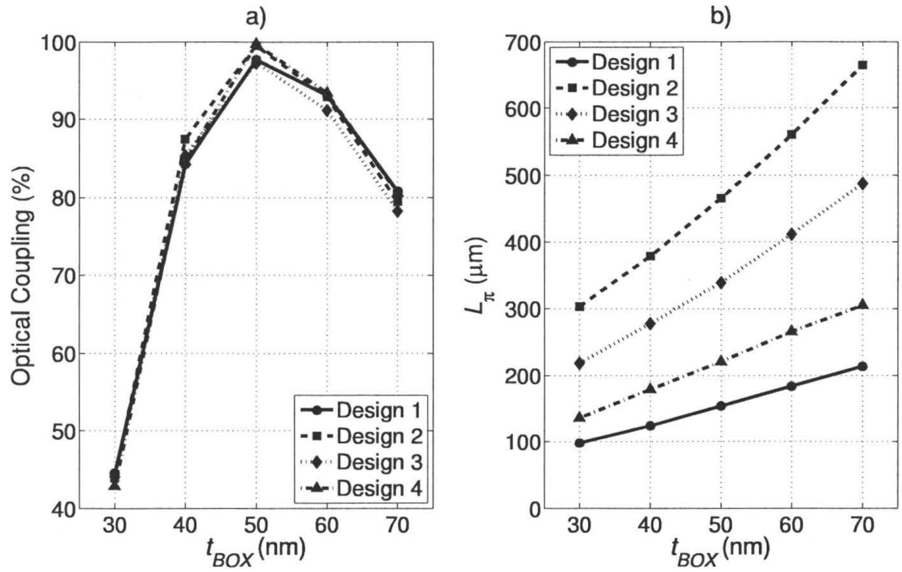

Figure 3.12: Device tolerance to deviations in separating oxide thickness
a) Optical coupling tolerance b) Change in coupling length

Device tolerance to variations in the separating oxide thickness is shown in Figure 3.12. This thickness was chosen to be 50 nm for each of the proposed designs, but errors during oxide growth or deposition can result in a different or nonuniform thickness over an entire wafer. Here, the separating oxide thickness is assumed to be uniform over the entire coupler length. A reduction of the separating oxide thickness results in a shorter coupling length due to enhanced optical field overlap. This coupling length will not match the fabricated coupler length, reducing optical coupling efficiency. The wafers used in this project should have oxide thickness non-uniformities of less than 5 nm.

3.2.3 Simulation Results for Fabricated Devices

The device layer thickness of 2 μm was chosen to reduce the computational time for the above simulations. The actual device layer thickness used for fabricating the slab-to-slab stacked rib coupler was chosen to be H = 5 μm, which improves device tolerance to fabrication errors. The coupling length changes to 1104 μm for the dimensions of $H = 5$ μm, $h = 3.5$ μm, $W = 2$ μm and $t_{BOX} = 0.05$ μm, with TE-polarized light, and to 4979 μm for the TM polarization. This width was chosen since more narrow features would be challenging to pattern with the photolithographic process used for this project, while the etch depth was chosen to ensure single-mode propagation. These lengths are for a coupler with oxide cladding layers, while the fabricated device will have air as the cladding for both the upper and lower waveguides. This reduces the TE-polarization coupling length to 782 μm. The coupling length for various rib geometries is shown in Figure 3.13 for H = 5 μm, t_{BOX} = 0.05 μm, air cladding layers and a TE-polarized input with a wavelength of 1550 nm.

The fabrication of the rib-to-rib stacked waveguide coupler would require an alignment accuracy between the overlapping waveguides that is beyond our capabilities. Therefore, a demonstration of its optical coupling was not attempted, nor was the rib-to-slab design due to its inherent limitation of incomplete power transfer. The stacked channels device with a tapered mode converter input was fabricated by patterning and etching both waveguides simultaneously, eliminating any misalignment between them. The material used for these devices had a guide layer thickness of H = 2.5 μm and a thicker separating oxide of t_{BOX} = 0.1 μm. The

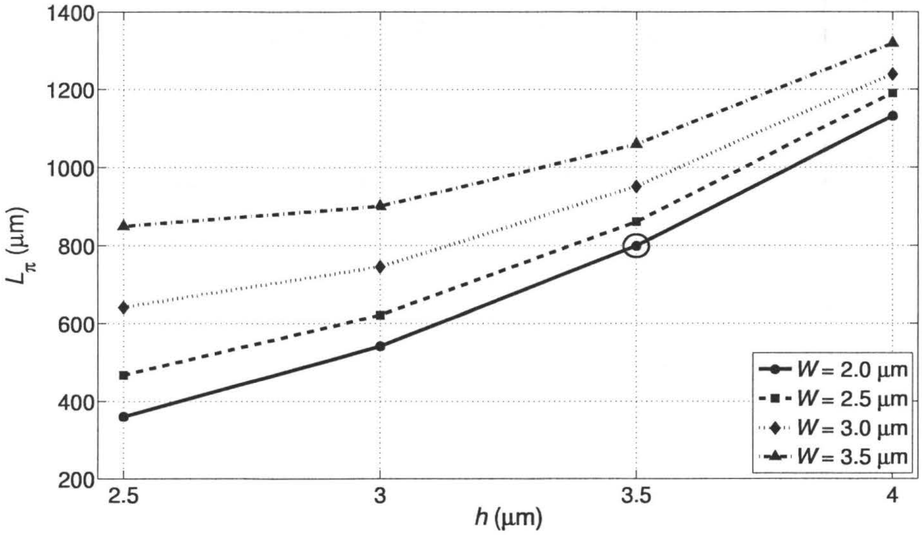

Figure 3.13: Optical coupling lengths for the fabricated stacked rib waveguide coupler

required coupling length was determined to be 640 μm for a TE-polarized input and 1840 μm for the TM-polarization. Here, the cladding layer is chosen to be air.

A tolerance analysis of the two fabricated stacked waveguide devices was completed and the results are shown in Figures 3.14 and 3.15. Note, the separating oxide thickness is different for both devices, referred to here as Device 1 for the slab-to-slab stacked rib waveguide coupler and Device 4 for the stacked channels coupler. An air cladding was used for all simulations, as will be the case with the fabricated devices.

The H = 5 μm BESOI wafers used in this project had a quoted non-uniformity of $\pm$ 0.5 μm, while the H = 2.5 μm SIMOX wafers had a non-uniformity of only $\pm$0.05 μm. While $\pm$ 0.5 μm over an entire wafer is an unacceptably large deviation, only small pieces were used for fabrication. Measuring the local thickness of each guide layer allows for the etch depth of each rib waveguide to be chosen such that

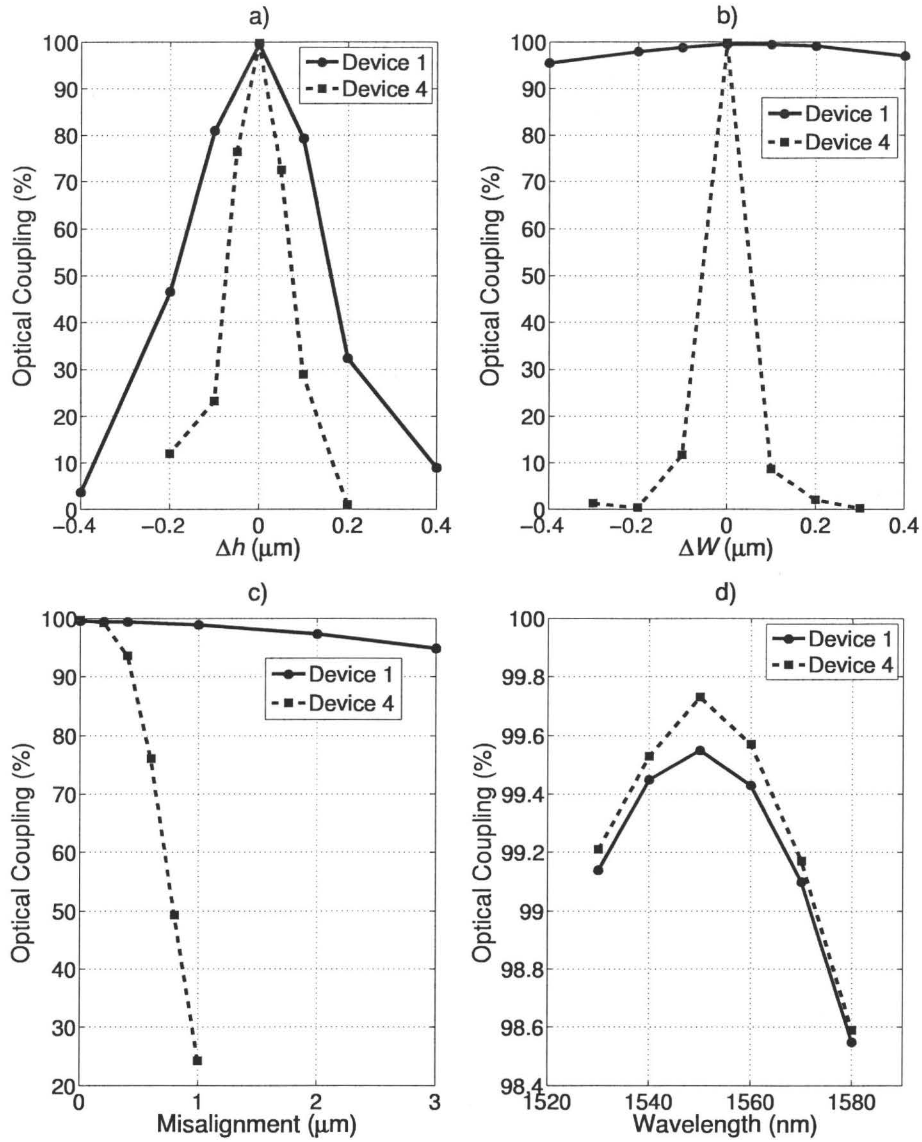

Figure 3.14: Tolerance to deviations between waveguide dimensions of fabricated devices

a) Deviation in slab thickness b) Deviation in waveguide width c) Misalignment of waveguides d) Deviation in input wavelength

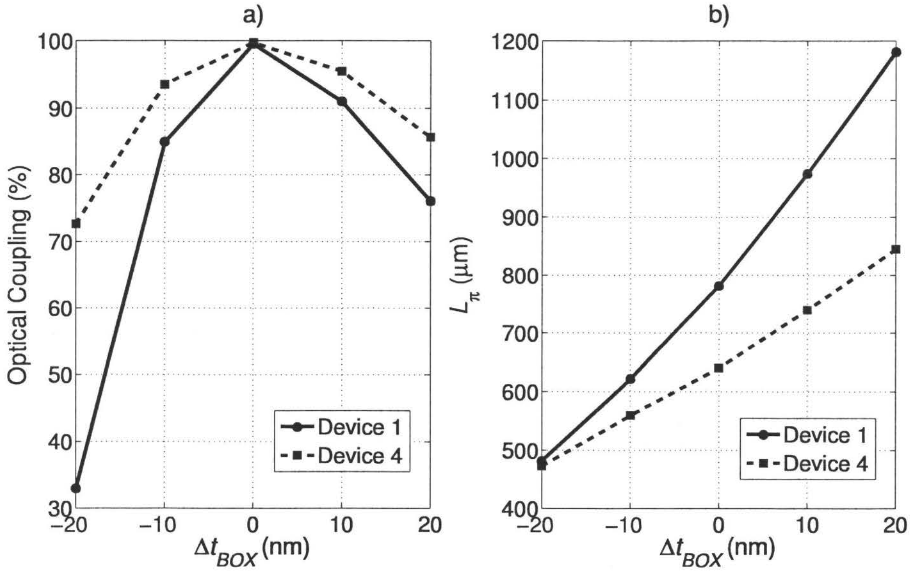

Figure 3.15: Tolerance to deviations in separating oxide thickness for fabricated devices

a) Optical coupling tolerance b) Change in coupling length

the slab regions are the same height. Deviations in slab thickness have a greater impact on device performance than the overall guide layer height since the optical mode profile is primarily defined by the lower corners of the rib.

3.3 Vertically-Integrated Multimode Interferometer Coupler

Another structure that could be used as a vertically-integrated optical coupler is the multimode interferometer (MMI), which is illustrated in Figure 3.16. With this structure, the coupling between waveguides occurs within a selectively placed thick silicon section. One disadvantage of the other proposed structures is that a very thin separating oxide layer is required to obtain sufficient mode overlap. Unless this oxide is made thin in only the selective areas where coupling is to occur, undesirable cross-talk between the guide layers will result wherever the waveguides overlap. This diminishes the overall functional density that can be achieved through vertical integration. With a vertically-integrated MMI coupler, the waveguides of each device layer are sufficiently separated to prevent optical cross-talk. Coupling will only occur in the regions where the thick separating silicon section is placed.

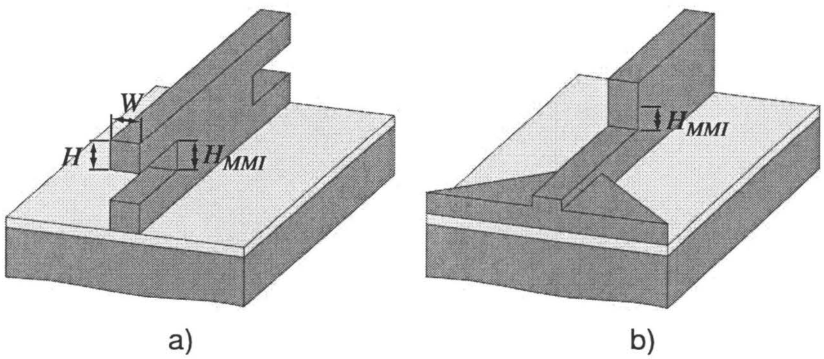

Figure 3.16: Vertically-integrated MMI coupler design
a) With upper and lower input/output waveguides b) Simulated structure with mode converter and no output waveguides

The coupling between stacked channel waveguides with an MMI section in between was simulated using BPM. A tapered mode converter is used to excite only the fundamental mode of the lower input channel waveguide, as illustrated in Figure 3.16 b).

Consider channel waveguides of height $H = 2$ μm and width $W = 2$ μm, and an MMI section of thickness $H_{MMI} = 1$ μm in between. The optical power will beat between the multiple modes of the silicon channel, which has an overall height of 5 μm, resulting in the expected self-imaging behavior of a typical MMI coupler. This 2×5 μm silicon channel supports six modes, as determined by their convergence using the ID-BPM mode solver. The optical power measured in the upper and lower 2 μm squares of the silicon channel is plotted versus MMI coupler length in Figure

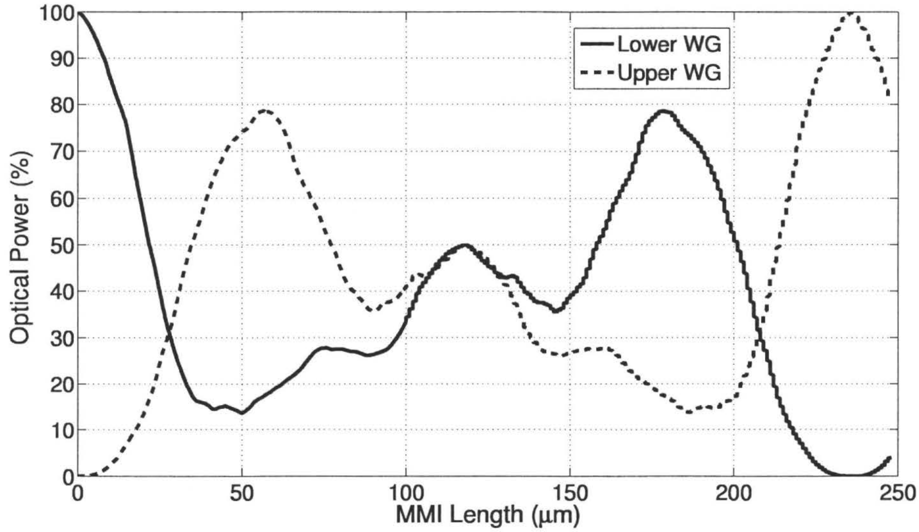

Figure 3.17: Optical power coupled between upper and lower waveguides of an MMI coupler

3.17. This represents the power that would be coupled into the output channel waveguides if they were placed at that particular coupler length. Note that these two plots do not always add up to 100 % since some of the light will be in the 1 μm of silicon in between. The coupling behavior observed is not sinusoidal as with the other proposed designs since the coupling depends on the beating between multiple modes as opposed to only two supermodes.

The coupler length should be set to 118 μm to achieve 50 % power splitting or to 236 μm for complete power transfer from the lower to the upper waveguide. These TE-polarized results are for an oxide cladding, but changing this layer to air has a negligible effect, reducing the coupling length to 235 μm. This coupler also exhibits less polarization sensitivity than the other proposed designs since the optical field discontinuities associated with the separating oxide are no longer present. Changing to a TM-polarized input reduces the coupling length to 224 μm, reducing

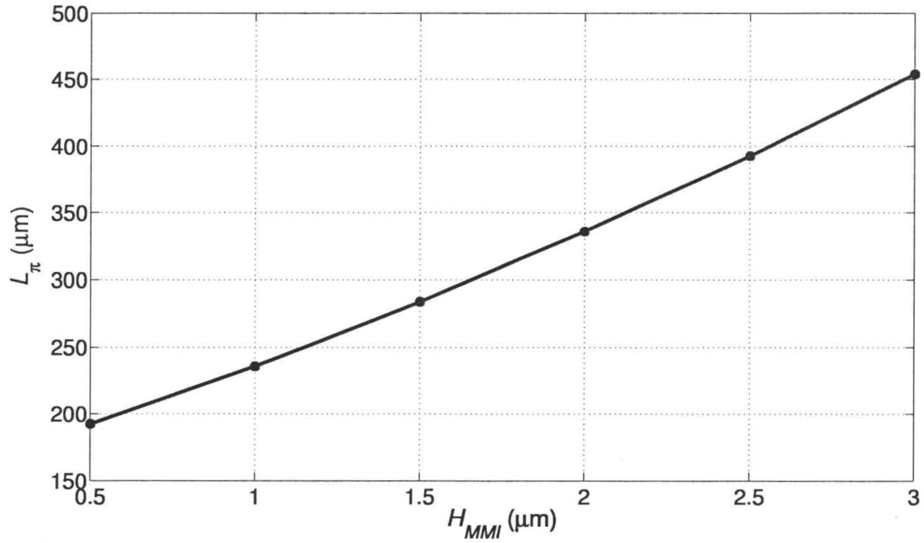

Figure 3.18: Coupling length dependence on multimode section thickness

the coupling efficiency to 79.8 % if the coupler length is fixed to the optimal length for a TE-polarized input. If the thickness of the MMI section is increased to H_{MMI} = 2 μm, the coupling lengths become 336 μm and 322 μm for the TE and TM polarizations, respectively. The dependence of the coupling length on the thickness of the MMI section is shown in Figure 3.18.

The other proposed coupler designs employing stacked rib waveguides could also be separated by a thick silicon slab to avoid optical cross-talk. Yet, it was impossible to obtain complete power transfer between the stacked waveguides since these devices now have a combination of rib and slab modes. Simulations of designs 1) through 3) with 0.5 μm of silicon in between resulted in shorter coupling lengths of 40 μm, 135 μm and 78 μm, but the maximum coupling efficiencies obtained were only 82 %, 85 % and 54 %, respectively.

The tolerance of the MMI coupler to fabrication errors was evaluated and the results are presented in Figure 3.19. The coupling efficiency is measured at the fixed length of 236 μm for the coupler with $H_{MMI} = 1$ μm and at 336 μm for $H_{MMI} = 2$ μm. The changes in width and misalignment were kept the same for both the upper waveguide and the MMI section, while the lower waveguide remained fixed.

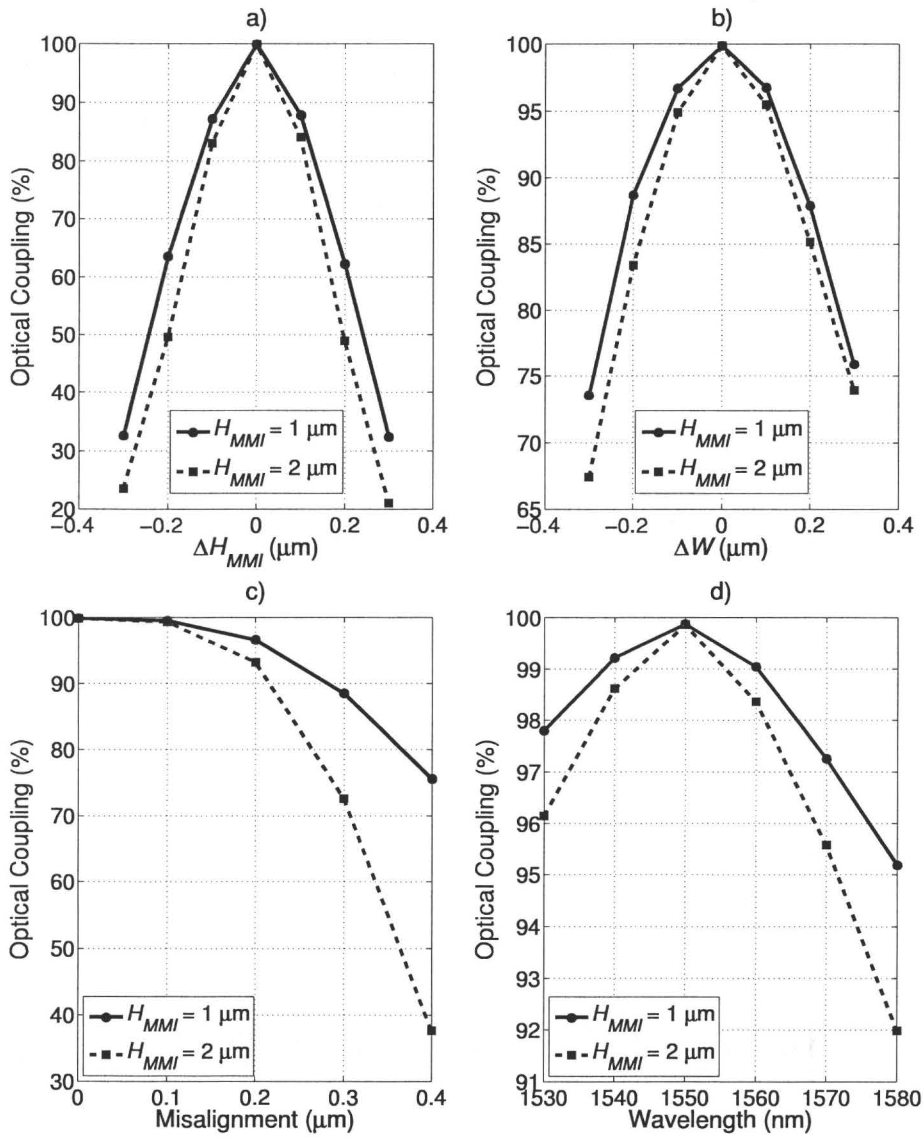

Figure 3.19: Fabrication tolerance for deviations in MMI coupler dimensions a) Deviation in MMI section thickness b) Deviation in MMI section width c) Misalignment of MMI section d) Deviation in input wavelength

3.4 Active Control of Optical Coupling

The above simulations provide the required coupling length to achieve complete optical power transfer between stacked waveguides. Ideally, the amount of power coupled at this fixed length could be actively controlled by changing the refractive index of one of the waveguides. This results in a mismatch of the effective indices of each waveguide mode, disrupting the coupling between them since they are no longer synchronized. As discussed in the theory chapter, this refractive index change can be obtained with the free carrier plasma dispersion effect using a current-injection diode, the thermo-optic effect using a metal strip heater or by generating free carriers using laser excitation. The dependence of optical coupling on the refractive index change in one waveguide is plotted in Figure 3.20 for the designs with $H = 2\ \mu$m. This dependence is also plotted for the devices to be fabricated later, the slab-to-slab stacked rib waveguide coupler and the stacked channels coupler.

The coupling efficiency of all the designs can be reduced to nearly zero with a 0.005 to 0.01 refractive index change along the entire coupler length. The required index change varies for each design due to their different coupling lengths and the optical confinement of the waveguide structure used. The index change would have to be precisely tuned to completely eliminate coupling due to its sinusoidal dependence on coupling length. This range of index change corresponds to a change in free carrier concentration, for both electrons and holes, of $\sim 5 \times 10^{18}$ carriers/cm^3 or a temperature change of 27 to 54 K. Both of these conditions can be obtained with a conventional current-injection diode or a metal strip heater. It was demonstrated in [76] that 6 mW of electrical power was required to raise the temperature of a 110

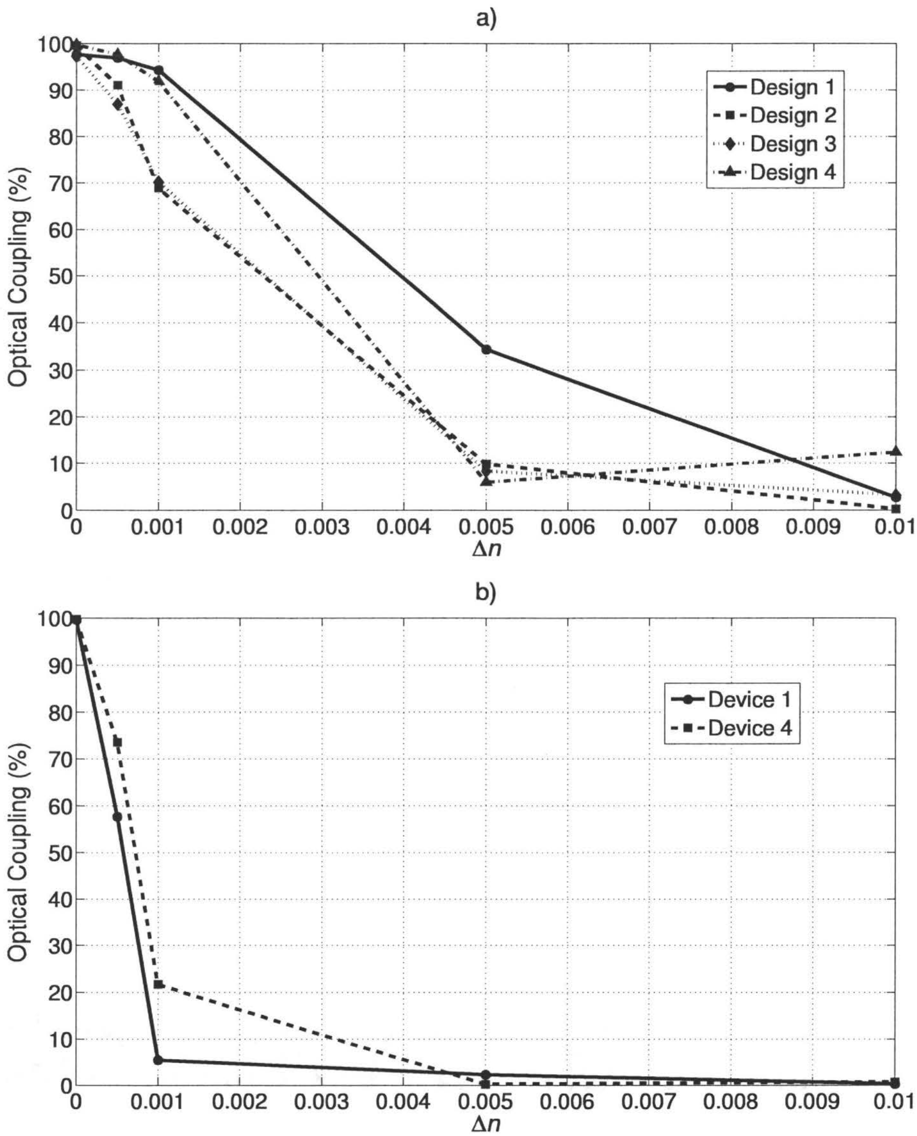

Figure 3.20: Optical coupling tuning with refractive index change

a) Designs with $H = 2$ μm b) Fabricated devices

μm long interferometer arm by 36 K.

Small fabrications errors can also be compensated for by actively controlling the coupler. Complete power transfer will not be achievable at the designed coupler length if both stacked waveguides are not identical, due to a mismatch in their effective mode indices. If this mismatch is small, it can be compensated for by either increasing or decreasing the refractive index of one or both of the waveguides. A diode can either inject or remove carriers, decreasing or increasing the index of refraction, by operating in either forward or reverse bias.

Active control of the vertically-integrated MMI coupler design can not be achieved by changing the refractive index of the entire multimode section of silicon. Unlike the stacked waveguide devices, each mode propagating in this silicon section will be perturbed by the same index variation since there is no separating oxide, thus maintaining their coupling behavior. A refractive index change of 0.01 within the entire multimode section results in only a 0.2 % reduction of the coupling efficiency.

It is possible to actively control the coupling of an MMI coupler if a different switching technique is used. Consider two MMI couplers of length $L_\pi/2$ separated by upper and lower channel waveguides, as illustrated in Figure 3.21. Propagating this length in each coupler alone will result in a 50 % splitting of the optical power through multimodal self-imaging. Light launched into the lower input waveguide of the first 50/50 splitter will equally couple into each of the channel waveguides. If these channel waveguides are unperturbed, each mode will then couple into the other 50/50 splitter and result in complete power transfer to the upper output waveguide, as if the coupling had occurred in a single MMI section of length L_π.

Alternatively, if the refractive index of one of the channel waveguides is changed, the resulting phase shift between the modes before they couple into the second splitter will disrupt their coupling. The coupling into the upper output waveguide can be completely eliminated if a π-phase shift is introduced into one of the channel waveguides. The latter coupling scenario has been simulated for an MMI coupler with $H = 2\ \mu$m, $W = 2\ \mu$m, $H_{MMI} = 0.5\ \mu$m and $L_\pi = 192\ \mu$m, and is presented in Figure 3.21.

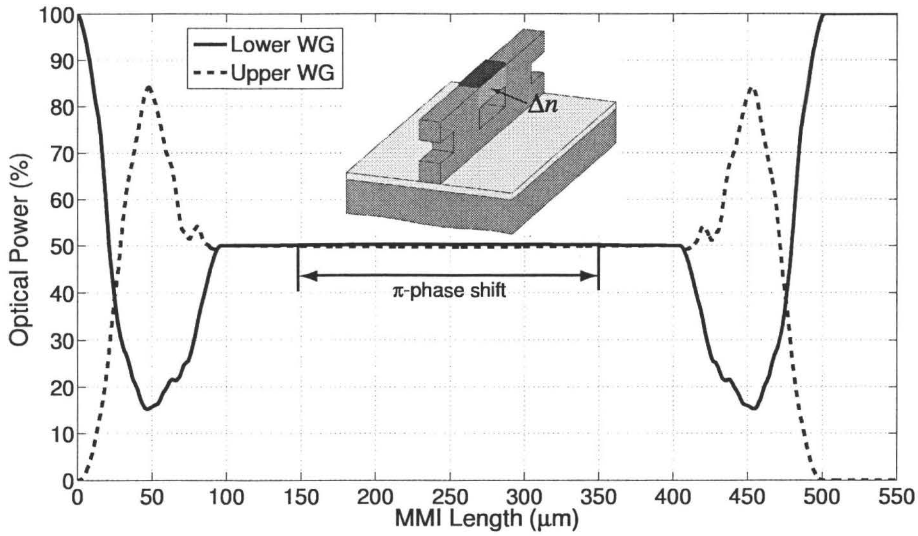

Figure 3.21: Optical power coupling with an actively-controlled vertically-integrated MMI coupler

This proposed switching structure is essentially a Mach-Zehnder interferometer, whose modulation can be controlled by changing the refractive index in one arm. Here, the separating channel waveguides are 309 μm long and an index change of 0.00368 over 209 μm will be required to obtain a π-phase shift between their modes. This change in material index does not correspond to the same change in the effective index of the mode since a portion of the optical field is propagating outside the silicon guide in the substrate and cladding. This index change corresponds to a change in both the free electron and hole concentrations of $\sim 1 \times 10^{18}$ carriers/cm^3

or a temperature change of 20 K. A 2 $\times$ 2 MMI switch employing the thermo-optic effect has been demonstrated using a similar layout as the one proposed above, only without vertical integration [82].

3.4.1 Thermally-Controlled Optical Coupling

The feasibility of a metal strip heater to actively control the refractive index difference between two stacked waveguides was also investigated. This thermal control was simulated with the heat transfer module of COMSOL corporation's Multiphysics software package, which uses the finite element method to solve partial differential equations.

Thermal simulations were completed for three of the fabricated devices: the stacked rib waveguide coupler, the stacked channel waveguide coupler and the MMI coupler. Aluminum strip heaters are placed either on top of the coupler or off to the sides on top of the slab regions. The heaters are typically 10 μm in width and the current applied to them ranges from 100 to 500 mA. During fabrication of the stacked channels device, a thin slab region was left on the lower waveguide to allow for active control of the coupling using either a diode, heater or laser-induced free carrier generation. BPM simulations predict that the presence of a 0.2 μm thick slab will only reduce the optical coupling by 0.4 % at the designed coupler length. A 0.3 μm thick oxide buffer layer is placed between the metal heater and the waveguide below to prevent optical absorption.

As expected, the relatively thin oxide layer between the stacked waveguides does not provide sufficient thermal insulation. Any heat transferred into one waveguide would rapidly diffuse into the other, resulting in a negligible temperature difference between them. This was the case regardless of the location of the heater

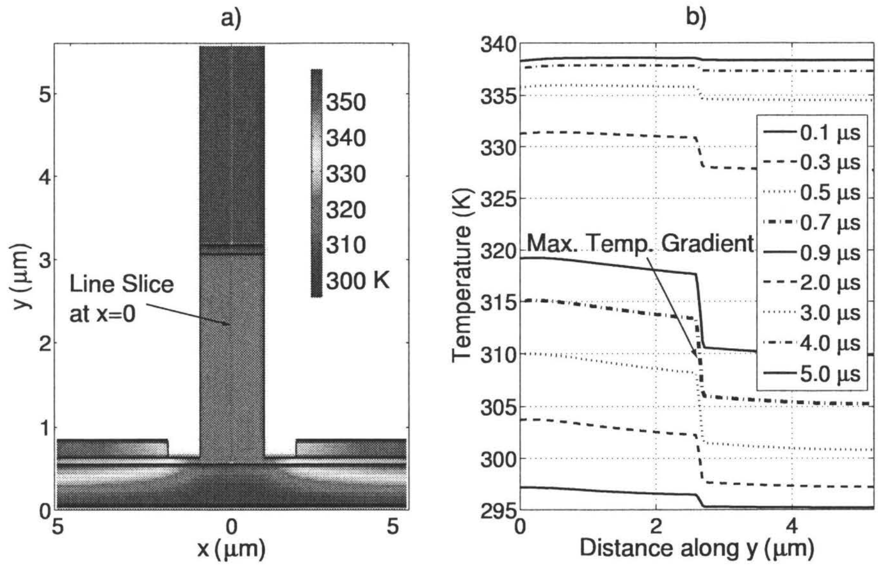

Figure 3.22: Time-dependent thermal diffusion in stacked channel waveguides

a) Temperature profile at 0.7 μs b) Temperature gradient at various times

strip. A time-dependent model was then used since steady-state solutions do not completely characterize the heat transfer behavior.

The time-dependent thermal diffusion in a stacked channel coupler is shown in Figure 3.22. The heaters are placed 1 μm from the sidewalls of the channels, with no buffer oxide required since an insignificant amount of light will be propagating in the thin slab regions. The slice plots taken through the center of the waveguides at various times after 300 mA of current is applied to the heater reveal that the largest temperature difference of $\sim$8 K between the two waveguides occurs at 0.7 μs. A contour plot of this temperature gradient is shown in Figure 3.22 a).

These time-dependent results were consistent over a range of applied current, leading to an investigation of modulating the current to confine heat in only one of the waveguides. Both square and sine waves of varying frequency and amplitude were used to modulate the applied current. An obvious frequency to use would be 1.4 GHz, corresponding to the switching period of 0.7 μs. Yet, this frequency and several other attempts did not result in a stable temperature gradient. In fact, at certain times the heat would be localized in the waveguide farthest from the heater. This current modulation technique was attempted with all three fabricated designs, with no success.

The observed temperature instability can be attributed to the fact that the air cladding and the silicon substrate below both act as heat sinks. Once the current is turned off the heat diffuses in all directions, it does not just return back to the heater. This contrasts the behavior of a p-n diode that can be operated in forward bias and then reversed. These simulations preclude the use of a metal strip heater to change the index of refraction between stacked waveguides. Active coupling control can only be achieved using free carriers generated with a current-injection diode or by laser excitation, as they will be confined to only one of the waveguides by the electrically insulating oxide layer.

3.4.2 Laser-Induced Free Carrier Generation

The refractive index of the upper guide layer can also be varied by generating free carriers through pump laser excitation. Light will be absorbed if the photon energy of the incident laser beam is greater than the band gap, generating electron-hole pairs through interband absorption. Since silicon has an indirect band gap of 1.1 eV, any light with a wavelength below 1.1 μm will be absorbed, along with the emission of phonons. The free carriers have a finite lifetime before recombining, releasing energy in the form of heat. During this lifetime, the free carriers will diffuse away from the region where they were created.

A time-dependent partial differential equation was used to estimate the steady-state concentration of free carriers generated for a given set of laser parameters, and was solved using the PDE toolbox in MATLAB. The model was chosen to be two-dimensional along the surface of the device layer since the carrier diffusion length in silicon is much larger than the depth of the device layer. Therefore, it was assumed that the carrier concentration will be uniform along the depth of the device layer, although there will be significantly more recombination at the upper surface and at the buried oxide.

Another assumption is that the changes in both the free electron and hole concentrations are equal. As well, any heat generated through phonon emission and recombination will be neglected since it can be assumed from the metal strip heater simulations that the heat will diffuse uniformly between the stacked waveguides, resulting in the same refractive index change for both guides. Yet, the free carriers will be confined to the upper waveguide by the buried electrically insulating oxide layer. This model and its parameters are thoroughly described in Section 2.8.

The argon-ion laser chosen for the laser excitation experiments has a wavelength of 514 nm and a maximum output power of 130 mW. The model parameters corresponding to this wavelength are: n_{Si} = 4.209, n_{SiO_2} = 1.468 and an absorption coefficient of α = 14718 cm^{-1}, which are used to determine the fraction of light reflected off the top surface and the amount of light absorbed by the upper silicon layer [83]. For a device layer thicker than 2 μm, 38 % of the incident light will be reflected, while the remaining 62 % is absorbed in the upper silicon layer. Virtually no light will be transmitted into the lower waveguide below.

An optical fiber with a spot size radius of ω_0 = 2 μm was used to excite the silicon surface at normal incidence. The 2 μm width of the stacked channel coupler is half the spot size diameter, meaning only 63 % of the beam will actually illuminate the upper waveguide. This is not an issue with the slab coupling device or the rib-on-rib stacked waveguide coupler since they have slab regions that will absorb the entire beam. The fiber delivered a maximum power of 42 mW due to losses from coupling the laser into a bare fiber end. The initial carrier concentration is chosen to be $N_{c0} = 1 \times 10^{15}$ cm^{-3}, which is the phosphorous impurity concentration for a silicon wafer with a resistivity of $\sim$5 Ω-cm. The diffusion coefficient is chosen to be $D = 19$ cm^{2}s^{-1} [79].

Recombination can occur through various processes, with surface recombination being the dominant process. Since the recombination time will have to be assumed, a range of reasonable values will be simulated to determine an approximate carrier concentration change and diffusion length. The dependence of the steady-state change in carrier concentration and diffusion length as a function of recombination time for an absorbed power of 25 mW is shown in Figure 3.23. The carrier concentration is measured at the peak of its profile along the z-axis, while the diffusion length is the distance from the center of the profile to the base level carrier concentration of $N_{c0} = 1 \times 10^{15}$ cm^{-3}.

The peak change in carrier concentration is not strongly dependent on the chosen carrier lifetime, and can be approximated as $N_c = 1 \times 10^{19}$ cm^{-3}. While the peak change is consistent, the total concentration of free carriers generated is larger for the longer recombination times that result in longer diffusion lengths. The net phase shift caused by the free carriers depends on both the carrier concentration change and the length over which this change occurs.

The free carriers will diffuse away in all directions from the generation region before they recombine. The concentration profile along the z-axis is presented in

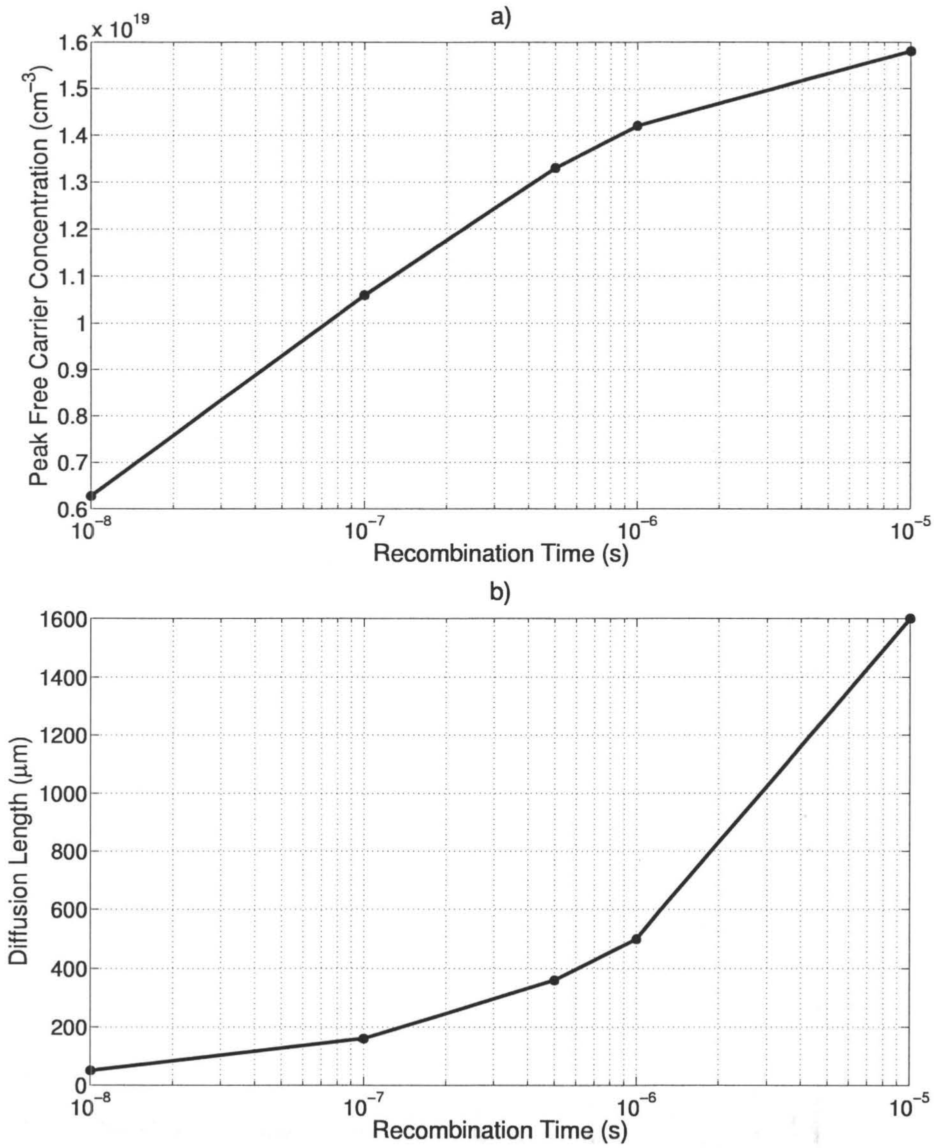

Figure 3.23: Change in laser-induced free carrier concentration and diffusion length for varying recombination time

a) Maximum increase of free carriers b) Distance of free carrier diffusion from the center of the laser beam

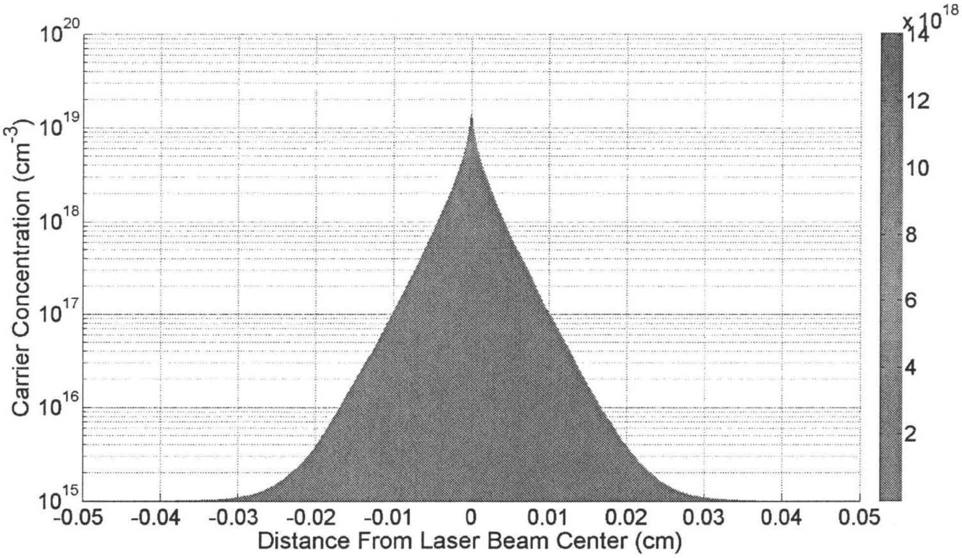

Figure 3.24: Laser-induced free carrier concentration profile

Figure 3.24 for an absorbed power of 25 mW and a recombination time of 500 ns. The diffusion length here is ~0.035 cm or 350 μm. By integrating this profile we find an average net change of $\Delta N_c = 1.1 \times 10^{17}$ carriers/cm^3 over 700 μm.

The peak change in free carrier concentration as a function of absorbed optical power is shown in Figure 3.25 for three different laser beam radii and a recombination time of 500 ns. These results can be scaled to the incident laser power, knowing the reflection and absorption coefficients for a specific silicon layer thickness. The larger beam radii did not result in a significant increase of the diffusion length, as it is more dependent on the recombination time. They did result in a decrease in the peak change in carrier concentration, with beam radii of 5 and 10 μm resulting in a peak change that was 75 % and 56 % of the change obtained with $\omega_0 = 2$ μm, respectively. This arises from the spreading of the optical power absorbed away from the center of the carrier generation region.

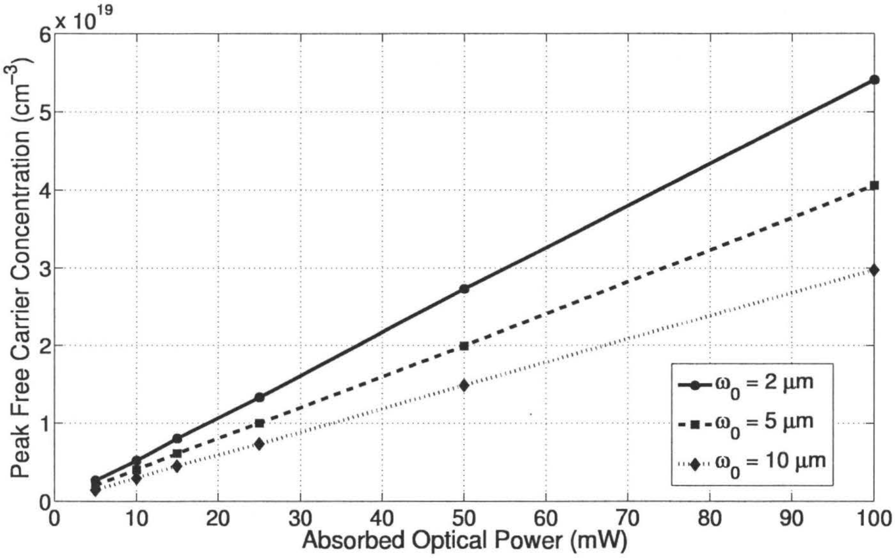

Figure 3.25: Peak free carrier concentration with varying absorbed optical power

To determine the phase shift induced by the free carriers, we estimate a concentration change of $\Delta N_c = 1 \times 10^{17}$ cm^{-3} over a length of $L = 700$ μm. This change in carrier concentration for both electrons and holes results in a refractive index change of $\Delta n = -4.3 \times 10^{-4}$ in the upper guide layer. The phase shift introduced between the modes of the stacked waveguides is therefore: $2\pi\Delta nL/\lambda_0 =$ -1.2 radians. If a π-phase shift results in 100 % of the light remaining in the lower waveguide, then we can approximate that a phase shift of -1.2 rad will result in 39 % of the light remaining. Complete switching would require this refractive index change to be induced over a length of 1800 μm, corresponding to a relatively long recombination time of $\sim$5 μs.

This phase shift is strongly dependent on the length over which the refractive index change is applied. Free carrier lifetime, and hence diffusion length, is shorter for SOI waveguide geometries, compared to bulk silicon. Carrier lifetime also decreases as the dimensions of the waveguide become smaller. This is due to increased surface recombination occurring at both the upper interface and the buried oxide interface. Typical recombination times between 10 ns and 100 ns have been reported for SOI rib waveguides with dimensions of 2-5 μm [80]. While longer recombination times on the order of microseconds have been observed in bulk SOI wafers, these times are only applicable if laser excitation is used to generate carriers in the stacked slab waveguide coupler. Therefore, a compromised value of 500 ns was used for the above calculation.

The phase shift introduced between the upper and lower waveguide modes will decrease if the laser-induced refractive index change is applied over a shorter length. Given the same laser excitation conditions as above, the amount of uncoupled light reduces to 18 % and 6 % for lengths of 320 μm and 100 μm respectively. These lengths correspond to recombination times of 100 ns and 10 ns, respectively. The observability of a change in coupling efficiency will have to be experimentally verified, since the recombination time of the devices is unknown.

It can be concluded from the above simulations that the most practical method of actively controlling the optical coupling is a current-injection diode. The laser excitation method will be attempted to demonstrate proof-of-concept active control since the the fabrication of a diode across one of the waveguides requires a complex process, in addition to the challenges encountered in creating stacked waveguide devices. Thermal control can only be used in the MMI coupler switching configuration, between two 50/50 splitters, as heat uniformly diffuses between the stacked waveguides of the other devices. Since all of the coupler designs exhibit polarization sensitivity, the coupling efficiency can also be controlled by varying the input polarization, which can be easily accomplished experimentally.

3.5 Fabricated Device Designs

Various vertically-integrated coupler designs were fabricated to demonstrate optical coupling between stacked guide layers. A summary of the dimensions chosen for these devices is summarized in Table 3.1 (all dimensions in μm). A preliminary demonstration of coupling between slab waveguides was attempted before advancing to more complex structures: stacked rib waveguides (slab-to-slab), stacked channel waveguides and a vertically-integrated MMI coupler.

Table 3.1: Dimensions of fabricated vertically-integrated couplers

Structure	H	W	h	$t_{BOX}\|H_{MMI}$	TE L_π	TM L_π
Stacked slabs	5.0	∞	5.0	0.05	2985	15996
1 (Ribs Slab-to-Slab)	5.0	2.0	3.5	0.05	782	3546
4 (Channels)	2.5	2.0	0	0.1	640	1840
5 (MMI)	2.0	2.0	0	1.0	235	224

Thicker waveguide layers were chosen for these proof-of-concept designs to improve device tolerance to fabrication errors, at the expense of longer coupling lengths. Coupling lengths can also be reduced by fabricating devices with more narrow waveguide widths, but the minimum achievable feature size of the photolithography process used for this project limited them to 2 μm.

Chapter 4

Device Fabrication

The vertically-integrated coupler designs proposed in Chapter 3 were fabricated using a range of novel processing techniques. In this chapter, detailed fabrication processes for these devices will be presented, along with images of the resulting structures. The development of an in-house wafer bonding technique will also be discussed.

4.1 Wafer Bonding Process Developed at McMaster University

Each waveguide layer would be ideally fabricated separately, and then stacked using an aligned wafer bonding technique. This allows each layer to be inspected for fabrication errors before bonding them together. This approach increases device yield, which is not only important in demonstrating proof-of-concept but is crucial for the feasibility of the device as a commercial product. A commercial wafer bonder with precise alignment was not available at McMaster University, or within easy access, during this project.

A novel fabrication technique avoiding aligned bonding was therefore developed, starting with commercially-supplied double-layer SOI wafers. Wafer bonding is a requirement mid-process, because the double-layer SOI is bonded facedown to a support wafer and then etched back to access the second silicon device layer. Conventional photolithography is used to align and pattern the upper waveguide to match the lower. This section describes in detail the development of the in-house wafer bonding process used in fabricating the vertically-integrated optical couplers.

4.1.1 Initial Attempts at Wafer Bonding Utilizing Spin-On-Glass

The direct bonding process of two silicon wafers (used in manufacturing SOI wafers) involves cleaning both wafers in either sulphuric or nitric acid to create hydrophilic surfaces. Hydroxyl (O-H) groups form on these surfaces resulting in silanol (Si-OH) bonds. The two wafers are brought into contact and will bond with minimal pressure since the smooth and flat surfaces of each wafer have a tendency to adhere through hydrogen bonding. This bond is strengthened by annealing the wafer pair to polymerize the silanol bonds into siloxane bonds (Si-O) bonds in the following reaction:

$$\text{Si-OH} + \text{HO-Si} \rightarrow \text{H}_2\text{O} + \text{Si-O-Si} \tag{4.1}$$

The water molecules formed at the bonding interface create intrinsic voids below annealing temperatures of 300 °C. Above this temperature, these voids disappear as more siloxane bonds are formed and the water is liberated. At temperatures greater than 800 °C, the bond strength is increased as the more easily deformed surfaces oxidize in the presence of trapped water, bringing them into better contact. The bonding strength approaches that of bulk material at temperatures greater than 1000 °C [47].

This chemical surface activation technique requires the bonding surfaces to be very smooth and flat. For the purposes of the current project, modified techniques were investigated since the proposed fabrication processes involve bonding patterned samples that have been exposed to numerous process steps beforehand. Initial attempts of bonding a patterned silicon sample to an unpatterned support wafer involved the use of spin-on-glass (SOG) as an intermediate layer. SOG is an organic liquid consisting of siloxane molecules dissolved in a solvent, such as n-Butanol. Here, the bonding mechanism also involves hydroxyl molecules absorbed at the bonding interface resulting in attractive forces between the two wafers. It was hypothesized that the cured SOG would reflow, or at least soften, around the edges of the patterned features during a 1100 °C furnace anneal. This (in theory) would strengthen the bond between the patterned sample and the support wafer as the SOG solidified during cooling.

The SOG used for this project, "Dielectric Coating DC4-500", was supplied by Futturex, Inc. The first step of this modified bonding process was to spin the liquid SOG onto a 2"-diameter silicon support wafer. The resulting film thickness was 500 nm after being spun at 3000 rpm for 40 s and baked on a 200 °C hotplate for 2 min. This initial hotplate bake evaporated the excess solvent. The resulting film needed to be cured at high temperatures in a rapid thermal annealer (RTA) to remove silanol groups and form a stable glass layer. This annealing step was performed in nitrogen using temperatures of 300, 500, 800, and 900 °C with rise and dwell times of 1 min for each temperature step, followed by a gradual cooling over several minutes.

The patterned sample (i.e. that to be bonded) consisted of 2 μm wide rib waveguides spaced by 200 μm, with 100 μm wide relief features placed in between.

These relief features increased the surface area over which bonding could occur. This pattern was also surrounded by 2-3 mm of unetched silicon measured from the sample edge. These edges encapsulate the sample once it is bonded facedown. The patterned samples were thermally oxidized with 50 nm of oxide and then cleaned in sulphuric acid heated to 80 °C, with hydrogen peroxide slowly added to create a bubbling solution. The samples were rinsed in deionized (DI) water and blown dry with a nitrogen gun before being brought into contact with the SOG-covered support wafer that had also been cleaned beforehand. Bonding was initialized when the patterned sample was placed facedown onto the support wafer using tweezers.

The bonded pair was then sandwiched between 2" × 2" ceramic blocks weighing approximately 500 g each. A larger 1 kg ceramic block was placed on top to provide pressure while the bond was annealed in a box furnace. The temperature of this box furnace was gradually increased to 1100 °C with the sample and ceramic blocks in place. The bonded pair was then annealed overnight in an ambient environment for 16 hr. The furnace was gradually cooled before removing the bonded samples.

The resulting bond could have been inspected by transmitting infrared light from an incandescent source through the bonded pair and capturing the resulting image using an IR camera. Any voids in the bonded interface that may later cause the sample to delaminate should be visible in these conditions. The strength of the bond can also be measured using either a pull or knife-edge test procedure. Since the final fabrication process involved mechanically and chemically thinning the double-layer SOI sample, this process step was used as the ultimate test of the bond interface quality.

The thinning procedure begins with melting crystal bond onto a brass block that has been heated to 160 °C on a hotplate. The bonded pair was then mounted onto the block, with the backside of the patterned double-SOI sample facing up, followed by cooling the block to solidify the melted crystal bond. The patterned sample was thinned to a thickness of $\sim$100 μm, as measured by a micrometer depth gauge, using a rotating silicon carbide sandpaper disc and water. The bonded pair was then removed from the block by heating it on a 175 °C hotplate to liquify the crystal bond. Residual crystal bond was removed in acetone and methanol before continuing.

The remaining 100 μm of backside silicon was etched in tetramethyl ammonium hydroxide (TMAH) heated to 90 °C. This chemical etch stops on the buried oxide of the double-layer SOI due to its selectivity, which would not have been the case had the mechanical thinning technique been used exclusively. It was usually during this process step that the patterned double-SOI sample would delaminate from the support wafer if processing failure did occur. While the bond was sufficiently strong to withstand the shearing forces of the rotating silicon carbide disc, any voids present in the bonding interface would blister when only tens of microns of silicon remained during the chemical etch. These blisters would break allowing the chemical etchant to seep into encapsulated areas, causing the entire sample to delaminate from the support wafer.

Voids in the bonding interface are caused by trapped contaminants, such as hydrocarbon particles or metal ions from the tweezers, or by voids in the SOG layer itself even before bonding. Spinning the SOG onto the support wafer without forming any comets or voids was close to impossible in the McMaster cleanroom (limited to a class 10,000 environment). Only SOG-spun wafers without noticeable comets or large voids were used for bonding, but even these wafers did not result in acceptable results. Ideally, the samples would be cleaned and bonded in a sealed environment to virtually eliminate any trapped contaminants.

It is a possibility that the cured SOG layer did not reflow or even soften in the box furnace. Uncured SOG was tested but this also resulted in sample delamination. Samples were also brought into contact with liquid SOG in between, as investigated in [49], but without success. It was suspected that the trapped layer of liquid SOG cracked during annealing as the solvent vapor had nowhere to escape as it evaporated. SOG was also spun onto the patterned samples for some of the process attempts, but the presence of etched features caused the SOG to bead along their edges resulting in an uneven surface that did not successfully bond.

Bonding with another type of SOG that had been doped with phosphorus was then attempted. This phosphosilicate polymer SOG contained 4 %wt phosphorous bonded to O atoms in the Si-O backbone. The "Accuglass P-114A" SOG used was supplied by Honeywell International, Inc. The same spinning and annealing procedures were used for this SOG type, with the exception that its resulting thickness was only 100 nm. This phosphosilicate SOG was known to reflow at the temperature used to anneal the bonded wafer pair due to the presence of dopant atoms. Despite this reflow, the bonded samples continued to delaminate during the chemical etch. Therefore, using SOG as an intermediate bonding layer was abandoned, primarily due to the unavoidable comets and voids observed in the spun SOG films.

4.1.2 Plasma-Activated Bonding

Plasma-activated bonding between oxidized silicon surfaces was the next technique attempted. It was assumed that thermally grown oxide layers would contain significantly less contaminants than SOG layers and would be relatively smooth. The challenge was developing a process that would create a durable bond between the patterned sample and an oxidized support wafer.

To ensure a high-quality, contaminant-free thermal oxide surface for the support wafer, square samples were cleaved from larger silicon wafers that had been thermally oxidized by the manufacturer with a 2.4 μm thick oxide film. These support wafers were made to be larger than the square double-SOI samples by 2-3 mm around each edge. Any remaining oxide that had been used for patterning the double-SOI sample was removed using buffered hydrofluoric acid (BHF) before thermally growing a fresh oxide layer that was 50 nm thick. Both the sample and the support wafer were then cleaned using a solution of DI water, ammonium hydroxide and hydrogen peroxide mixed in a 5:1:1 ratio and heated to 70 °C before submerging the samples for 10 min. The samples are rinsed in DI water, dried using a nitrogen gun, then placed in BHF for a short 10 sec dip, followed by a final DI rinse and drying. The BHF dip was added to the cleaning procedure to etch away any contaminants that may have remained on the oxide surface after the first clean.

Both the sample and support wafer were then immediately placed in the chamber of a Technics parallel-plate plasma etcher. Their surfaces were exposed to an oxygen plasma for 1 min using the following chamber conditions: O_2 gas flow rate of 12 sccm, a chamber pressure of 200 mTorr and a plasma-generating RF power of 100 W at a frequency of 13.56 MHz. A wide range of plasma conditions and times have been proposed in [52, 53, 54], with all resulting in successful bonding with low annealing temperatures. Our conditions and time were chosen based on the results of [55], which suggested using plasma conditions that were already typically used for the etcher in our cleanroom.

Exposure to the O_2 plasma creates OH groups on the oxide surfaces of both samples. These groups result in spontaneous bonding of the samples when they are brought into contact. High bonding strength can be achieved at lower annealing temperatures since more silanol groups are created by exposure to O_2 plasma than by wet chemical activation alone. Published research has also suggested that plasma treatment creates a porous surface structure which enhances the diffusivity of water molecules from the bonded interface [54]. As well, oxygen plasma is typically used to remove hydrocarbons from silicon wafers, therefore it should also reduce the amount of surface contaminants that would later cause voids.

Contact was initiated by placing the patterned double-SOI sample facedown onto the oxidized support wafer using tweezers. The samples were thoroughly blown with a nitrogen gun before bonding, but no chemicals or water were used to avoid alteration of the plasma-activated surfaces. Published results using this bonding technique have demonstrated that bulk bonding strength can be achieved with a low-temperature anneal of only 400 °C [52]. A nitrogen tube furnace was used for this annealing step, with the bonded pair sandwiched between the two 500 g ceramic blocks. The larger 1 kg block could not be placed on top due to the limited clearance of the tube furnace.

Subsequent mechanical and chemical thinning of these plasma-activated bonded samples also resulted in delamination in the TMAH solution. An alternative thinning method employing reactive ion etching (RIE) was then attempted using a plasma-based silicon etcher located at the University of Western. The details of this etch process will be discussed later as this etcher was also used to etch device features. The bonded pair was initially thinned using silicon carbide sandpaper before using the plasma etcher, since a maximum etch rate of only 4-6 μm of silicon per minute was achievable. This etch process also selectively etches silicon and will stop once the buried oxide is reached.

This thinning procedure did result in some improvement as the bonded pair did not delaminate, but there were still many large voids that blistered. The sample did not delaminate since unlike the chemical etchant, the plasma does not fill the encapsulated areas through the blistered voids. The inconvenience of using a remotely located etcher led to the exploration of improving the bonding procedure such that it would be suitable for chemical thinning.

An alternative method that could improve the contact technique used to bring the two samples together was then proposed. It was suspected that the uneven bonding pressure applied to the sample using tweezers caused large air pockets to be trapped in certain areas. These large voids resulted in local blistering of the silicon, which was a major concern if they were located along the edges that encapsulated the patterned sample. A modified mask aligner-based process was proposed to roughly align and bring the two plasma-activated samples into contact, and is pictured in Figure 4.1.

A conventional mask aligner uses a vacuum plate to hold a photomask above a photoresist-covered wafer, which is held by a lower vacuum chuck. This chuck is moved laterally by micrometers to align the wafer before bringing it into contact with the photomask for subsequent ultraviolet exposure. To bond two plasma-activated samples, the mask aligner was modified such that the patterned sample is held by the plate normally used for the photomask. The opening and surrounding vacuum ring of this plate was too wide to hold our smaller samples, but it was used to hold a glass photomask blank with a small vacuum hole drilled into its center. Another vacuum line from a separate pump is attached to this hole and sealed with epoxy. The upper sample was held by the small vacuum hole, while the glass pho-

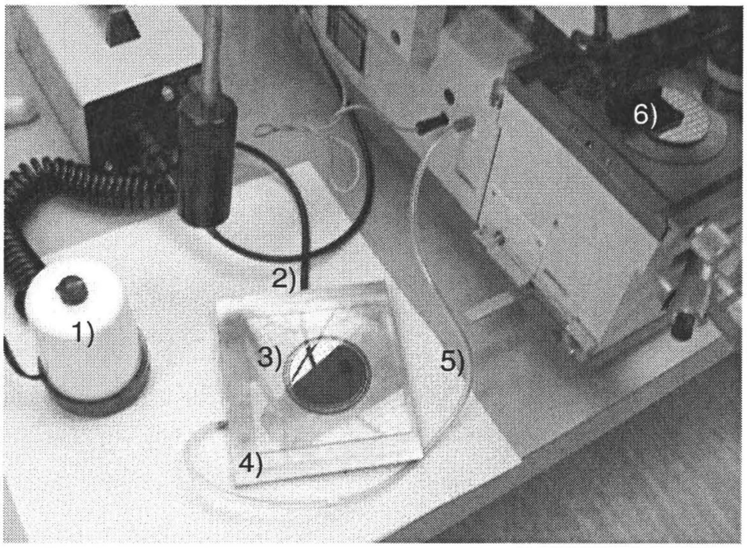

Figure 4.1: In-house wafer bonding system using a modified mask aligner
1) Mini vacuum pump 2) Vacuum line for holding upper wafer 3) Glass plate with vacuum hole for holding upper wafer 4) Photomask holder with vacuum ring for holding glass plate 5) Vacuum line to photomask holder 6) Vacuum chuck of mask aligner holding lower wafer

tomask blank was held onto the metal plate by the vacuum of the mask aligner. The oxidized support wafer was held by the lower vacuum chuck of the mask aligner.

The cleaved flats of both samples were first aligned by rotating the lower chuck. The chuck was also positioned such that the smaller patterned sample above was roughly centered relative to the larger support wafer below. The chuck was then raised to bring the two samples together while applying a uniform contact force. The vacuum pump that held the upper sample was turned off and the bonded pair was lowered and removed from the mask aligner. The resulting bond was then annealed at a temperature of 1200 °C in a nitrogen-filled RTA for 30 min, without the use of ceramic blocks.

Subsequent mechanical and chemical thinning of these mask aligner-bonded samples did *not* result in delamination in the chemical etchant. Smaller voids that eventually blistered were observed, but they would only affect a small number of the many devices fabricated on each sample. These smaller voids can be attributed to contaminants trapped between the samples. Although great care was taken to thoroughly blow off each sample using a nitrogen gun several times before the actual bonding occurred, this procedure could still be vastly improved by cleaning, plasma-activating and bonding the samples in a vacuum environment. Commercial wafer bonders that complete all of these processes under vacuum are available, but one was not accessible during this project. Regardless, using the modified mask aligner in conjunction with plasma-activation resulted in satisfactory results with acceptable consistency. This technique was also tried using doped SOG as an intermediate bonding layer, but resulted in delamination in the chemical etchant due to the voids in the spun SOG film. A summary of all the attempted bonding procedures is shown in Table 4.1.

Table 4.1: Summary of attempted wafer bonding processes

Intermediate Layer	Bonding Technique	Bonding Pressure	Annealing Conditions	Backside Thinning	Process Result
Cured SOG	Tweezers placement	Ceramic blocks	1100 °C ambient box furnace 16 hr	Mechanical/ TMAH	Delamination in TMAH
Uncured SOG	Tweezers placement	Ceramic blocks	1100 °C ambient box furnace 16 hr	Mechanical/ TMAH	Delamination in TMAH
Wet SOG	Tweezers placement	Ceramic blocks	1100 °C ambient box furnace 16 hr	Mechanical/ TMAH	Delamination in TMAH
Doped SOG	Tweezers placement	Ceramic blocks	1100 °C ambient box furnace 16 hr	Mechanical/ TMAH	Delamination in TMAH
Oxidized wafer	O_2 plasma + tweezers	Ceramic blocks	400 °C N_2 tube furnace 2 hr	Mechanical/ TMAH	Delamination in TMAH
Oxidized wafer	O_2 plasma + tweezers	Ceramic blocks	400 °C N_2 tube furnace 2 hr	RIE at Western	Bubbling around voids
Oxidized wafer	O_2 plasma + aligner	Mask aligner	1200 °C N_2 RTA 30 min	Mechanical/ TMAH	Satisfactory

4.1.3 Attempt at Fabricating Double-SOI Wafers

This project used double-layered SOI wafers as the starting material for fabricating vertically-integrated optical couplers. The specifications of these custom-fabricated wafers is presented in the next section. When additional material was needed, a recently acquired commercial wafer bonder was tested to avoid outsourcing the fabrication of the double-SOI wafers. This wafer bonder was purchased by Dr. Matiar Howlader from the Electrical Engineering department at McMaster and was installed in the Engineering Physics cleanroom at McMaster. Whole 100 mm-diameter silicon wafers were initially cleaned using the same previously mentioned procedure before thermally oxidizing them with 50 nm of oxide. SOI wafers would be used later if the test results with silicon wafers were acceptable.

The wafers were then exposed to an O_2 plasma in the chamber of the commercial bonder. The plasma conditions were: 200 W of RF power for 30 s, with a 60 Pa chamber pressure and an O_2 flow rate of 20 sccm. They were then exposed to nitrogen ion radicals generated by 2500 W of microwave power at 2.45 GHz for 30 s, with a 60 Pa chamber pressure and a N_2 flow rate of 20 sccm. This technique is known as sequential plasma-activated bonding (SPAB), since two different types of plasma are used, and is presented in [84, 85]. The nitrogen radicals allow for increased bonding strength at room temperature by creating unstable oxinitrides on each surface that stabilize into silicon-oxinitride bonds when the wafers are brought into contact. The bonding strength is increased by annealing at higher temperatures since additional Si-O-Si bonds are formed from the O_2 plasma-generated OH groups also present on the surfaces.

This particular commercial bonder does not bring the wafers into contact within the plasma chamber. While contact can be made in a separate anodic bonder module, the plasma-activated wafers must be removed from the chamber first. It was standard procedure in Dr. Howlader's group to place the wafers together using tweezers outside the bonder, resulting in bonding at room temperature. Finally, the bonded pair is cold-rolled under 0.2 MPa pressure at room temperature to remove trapped air. The modified mask aligner could not be used since the wafers are too large to be held by the small vacuum hole. As well, vision of the wafer flats would be obstructed preventing the wafers from being properly aligned.

Infrared transmission images for two examples of the bonded oxidized silicon wafer pairs are shown in Figure 4.2 a), b). An incandescent bulb was used to back-side illuminate the bonded wafers and the resulting transmitted pattern is captured using an IR camera. Large voids were observed with all of the bonding attempts, before and after annealing at 1200 °C. These voids would later blister with the removal of one of the wafer handles, resulting in an large proportion of unusable material area. Images of the commercially-bonded double-SOI wafers used for this project, starting with either single-layer SIMOX or bond and etch back SOI wafers, are also shown for comparison. While some voids were observed around the edges of these commercially-bonded wafers, the majority of the material was usable and was cleaved into smaller square samples for subsequent processing.

What may appear to be voids in the BESOI wafer pair are actually interference fringes from the non-uniformity of the silicon device layers due to uneven wafer thinning. The voids around the edges of these larger wafers were outside the imaging area. These interference fringes are not visible with the SIMOX wafers due to their highly uniform device layers defined through oxygen ion implantation. Due to the unacceptable amount of usable material area that would be obtained with the in-house wafer bonding processes, the decision was made to acquire additional

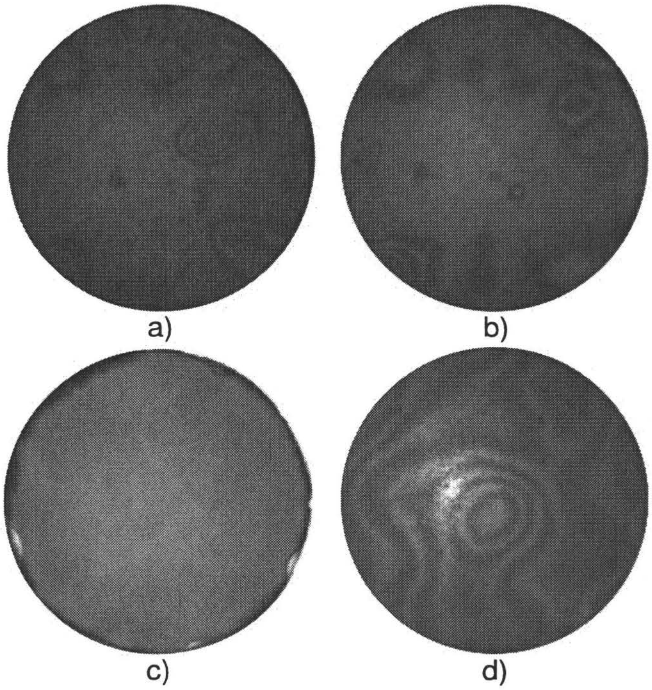

Figure 4.2: IR transmission images of bonded oxidized silicon wafer pairs
a), b) In-house bonded oxidized silicon wafer pairs c) Commercially-bonded SIMOX wafer pair d) Commercially-bonded BESOI wafer pair

custom-bonded material from the same industrial supplier that had been used previously. The results using the commercial wafer bonder at McMaster were most likely limited by contaminants since the bonding occurs outside the chamber. Additionally, the use of tweezers to initially place the wafers in contact may have resulted in trapped air pockets arising from applying uneven bonding pressure.

4.2 Double-SOI Wafer Specifications

SOI wafers with two silicon device layers were custom-fabricated for this project by IceMos Technology Limited located in Belfast, Northern Ireland. Two single-layer SOI wafers were each oxidized with half of the thin separating oxide layer before being bonded together using a standard contact and annealing process. The handle wafer of one of the SOI wafers was then removed using chemical mechanical polishing, revealing its thick buried oxide layer. This oxide was later removed at McMaster using hydrofluoric acid, revealing the upper silicon device layer.

Two types of wafers with different silicon device layer and separating oxide thicknesses were manufactured by IceMos. The first type consisted of silicon device layers of thickness $H = 5 \pm 0.5$ μm separated by an oxide layer of thickness $t_{BOX} = 50 \pm 1$ nm, a 1 μm thick buried oxide and a 550 μm thick silicon handle. The initial 150 mm-diameter, $\langle 100 \rangle$-orientated BESOI wafers used for bonding were supplied by IceMos, and their device layers were lightly P-doped with a resistivity of $\sim$10 Ω-cm.

The other wafer type was fabricated by bonding two customer-supplied 100 mm $\langle 100 \rangle$ SIMOX wafers that had been manufactured by IBIS Technology Corporation of MA, USA. These wafers were also lightly P-doped to a resistivity of $\sim$10 Ω-cm, and have a 0.375 μm thick buried oxide and a 550 μm thick silicon handle. The resulting double-SOI wafers have device layers of thickness $H = 2.50 \pm 0.05$ μm separated by an oxide of thickness $t_{BOX} = 100 \pm 2$ nm. These SIMOX wafers were chosen for thinner devices to ensure a highly-uniform and precise thickness for each silicon device layer. Scanning electron microscope (SEM) images of cleaved samples from both types of double-SOI wafers are shown in Figure 4.3. Note that the commercially-bonded SIMOX wafer pair does not have an optical-quality facet, while the BESOI wafer pair image is taken from a slab waveguide coupling device that has been properly thinned and cleaved. This illustrates the importance of thinning the sample before cleaving the facet.

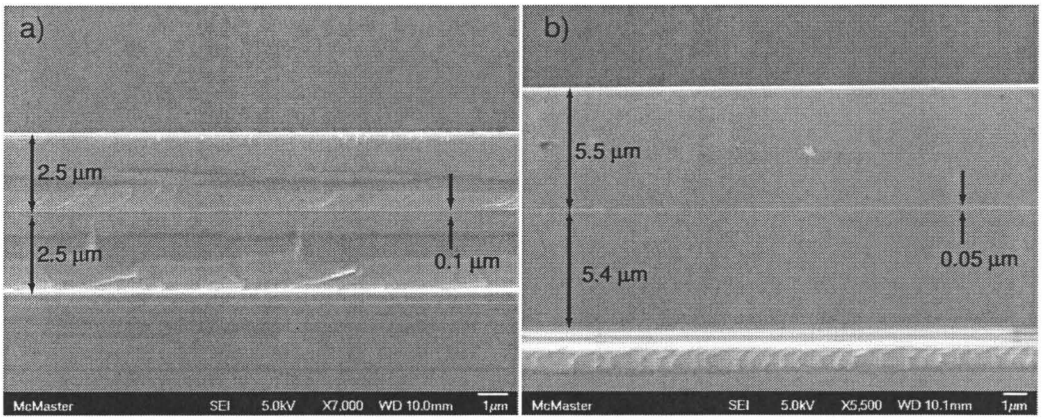

Figure 4.3: SEM images of cleaved double-SOI wafers
a) Commercially-bonded SIMOX wafer pair b) Commercially-bonded BESOI wafer pair

The vertically-integrated MMI devices described in Section 3.3 were fabricated using single-layer SIMOX wafers with a 4.95 ± 0.05 μm thick device layer, a 0.3875 μm thick buried oxide and a 700 μm thick silicon handle. This device layer was created using ion-implantation followed by epitaxial growth to reach the desired silicon thickness. These 150 mm $\langle 100 \rangle$ wafers were lightly P-doped to a resistivity of $\sim$10 Ω-cm and were manufactured by IBIS Technology Corporation.

4.3 Detailed Fabrication Process Steps

The following process steps are common to all of the devices fabricated for this project. They will be explained here in detail before presenting the complete fabrication process for each coupler design. Each process was completed in the class 10000 cleanroom located at McMaster University, unless otherwise noted.

Wafer Cleaning: Smaller 2.5 to 3 cm^2 samples were first scribed and cleaved from whole 100 mm or 150 mm-diameter wafers. They were then cleaned in a piranha solution consisting of sulfuric acid (H_2SO_4) and 30 % hydrogen peroxide (H_2O_2) in a 3:1 ratio. The H_2SO_4 was heated to 80 °C on a hotplate, the samples were submerged, and then the H_2O_2 was slowly added to create a bubbling solution. The samples were removed after 5-10 min and were rinsed in deionized (DI) water for 3 min before being dried with a nitrogen gun. Buffered hydrofluoric acid (10:1 BHF) was then used to remove the upper oxide layer left behind by the manufacturer, if applicable.

Mask Layer Growth/Deposition: Oxide mask layers were either thermally grown in dry oxygen in a JetFirst RTA at 1200 °C or deposited in a plasma-enhanced chemical vapor deposition (PE-CVD) chamber. The conditions for the CVD process were: silane (SiH_4) gas flow rate of 90 sccm, nitrous oxide (N_2O) gas flow rate of 70 sccm, a chamber pressure of 650 mTorr and a plasma-generating RF power of 100 W. The deposition rate for these conditions was ~70 nm/min.

Aluminum mask layers were deposited using a Bell Jar evaporator pumped below 1.5×10^{-3} mTorr. The current flow through the 99.999 % pure aluminum wire was adjusted to give a deposition rate of 1 nm/s as measured by a quartz crystal thickness monitor.

Photolithography: Device features were patterned using ultra-violet (UV) contact photolithography (PL) and Shipley S-series positive photoresists. Shipley S1808 photoresist was used for patterning narrow features such as waveguides. The hexamethyldisilazane (HMDS) primer used to promote better resist adhesion was first spun onto the sample at 4000 rpm for 30 s. The photoresist was then spun on at 4000 rpm for 30 s, resulting in a thickness of

0.8 μm. The sample was soft baked on a 110 °C hotplate for 2 min. A Karl Suss MJB3 mask aligner was used to expose the sample to UV light through a chrome/quartz photomask brought into vacuum contact. Further photomask details can be found in Appendix A. The exposure time was chosen such that the total UV exposure was 32 mJ/cm^2, as measured by the CI-1 photodetector channel. The sample was then developed in a solution of DI water and 351 developer in a 5:1 ratio for ~30 s. The sample was rinsed in DI water for 2 min and then hard baked at 130 °C for 2 min.

Thicker Shipley photoresists were also used for improved coverage over existing features. The process conditions used for each resist type are summarized in Table 4.2. The same spin rate was used for both the primer and photoresist and a time of 2 min was used for all hotplate baking.

Table 4.2: Photoresist process conditions

Resist Type	Spin Rate (rpm)	Thickness (μm)	Soft Bake (°C)	Exposure (mJ/cm^2)	Develop (s)	Hard Bake (°C)
S1808	4000	0.8	110	32	30	130
S1818	4000	1.8	110	70	60	135
S1827	4000	2.7	100	116	40	120
S1827	2200	3.7	100	190	30	120

Chemical Etching: Oxide mask layers selectively protected by patterned photoresist were etched using 10:1 BHF, with a room temperature etch rate of 50 nm/min for thermally-grown oxides and 100 nm/min for CVD-deposited oxides. The photoresist was then usually removed in an acetone bath, followed by a DI rinse, a methanol bath and then a final DI rinse. The remaining oxide mask was used to protect features during silicon etching in either tetramethyl ammonium hydroxide (TMAH) or by reactive ion etching (RIE).

TMAH preferentially etches in the $\langle 100 \rangle$ direction exposing the $\langle 111 \rangle$ crystallographic plane of silicon, creating sloped features with near-perfect smoothness. The resulting sidewalls are angled by $\sim$55 ° from the wafer surface. TMAH is also highly selective in terms of etching silicon instead of oxide, making the buried oxide layer a convenient etch stop. This high selectivity also means that samples should undergo a 1-2 s BHF dip to remove any native oxide growth before being placed in the TMAH solution. The 25 % TMAH in water solution used for this project had an etch rate of $\sim$0.2 μm/min when heated to 60 °C. This etch rate can be increased to $\sim$80 μm/hr for deeper etches by heating the solution to 90 °C.

Aluminum mask layers patterned with photoresist were etched using the common aluminum etchant "type A" solution containing 80 % phosphoric acid, 5 % nitric acid, 5 % Acetic Acid, and 10 % distilled Water. This solution etches at a rate of 10 nm/s when heated to 50 °C.

Plasma Etching: Reactive Ion Etching (RIE) was used to etch features with vertical sidewalls into the silicon device layers. The plasma etcher used for this project was the Alcatel 601E inductively-coupled plasma reactive ion etcher (ICP-RIE), located at the University of Western Ontario. This etcher uses a modified Bosch process that alternates between silicon etching with sulfur hexafluoride (SF_6) for 2 s and sidewall passivation with octafluorocyclobutane (C_4F_8) for 1 s. The C_4F_8 plasma creates a teflon-like layer on the sidewalls that prevents undercutting by the SF_6 plasma, resulting in features with a high-aspect ratio. The chamber conditions are as follows: 100 sccm of SF_6

for 2 s, 100 sccm of C_4F_8 for 1 s, a plasma-generating RF power of 1000 W, a substrate bias of 80 W RF power, a 7.5 mTorr chamber pressure and a substrate plate temperature of 20 °C achieved using backside liquid nitrogen-cooled helium. The resulting etch rate is 1.5-2 μm/min with a selectivity of etching silicon to oxide of ~30:1 and a selectivity to photoresist of ~10:1. If a photoresist mask is used, exposure to the plasma will cause it to harden, making it insoluble in acetone. The piranha solution was found to be effective in removing hardened photoresist without leaving behind any residue.

4.4 Slab Waveguide Coupler Fabrication

The step-by-step process used to fabricate the stacked slab waveguide coupler is illustrated in Figure 4.4. This coupler design was the least complex to fabricate since its process did not require any wafer bonding or accurate alignment during photolithography. These devices were fabricated using the double-SOI wafers with 5 μm thick silicon layers. As shown in step 1), the 0.13 μm oxide mask used to define the transition from the input waveguides to the stacked slab coupler was thermally grown at 1200 °C for 35 min. It was then patterned using S1808 photoresist and BHF. The photomask used for this step had a long straight edge separating a large clear area from a solid area of chrome. This edge was angled from the cleaved facet of the sample to vary the coupler length by ~1500 μm along the width of the device. The photoresist was then removed in acetone and methanol.

As shown in step 2), the 5 μm thick silicon layer above the input waveguides was etched using TMAH at 60 °C until the 0.05 μm thick separating oxide layer is reached. This oxide layer acts as a convenient etch stop. The 2 μm wide input

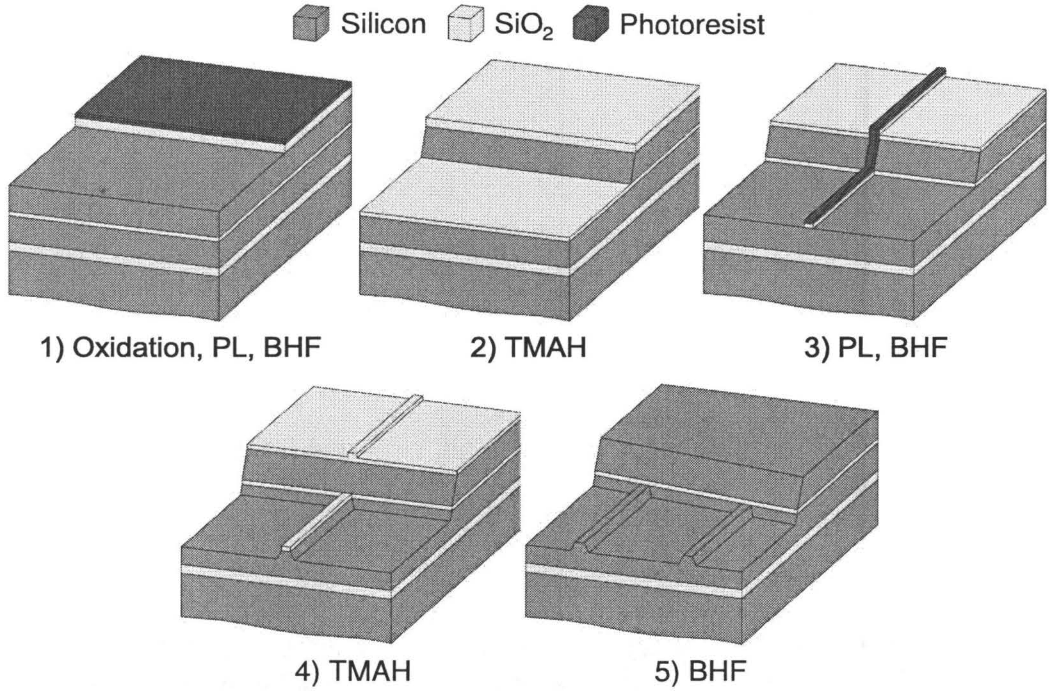

Figure 4.4: Fabrication process for vertically coupling slab waveguides

waveguides were then patterned into both the thermal oxide mask and the separating oxide layer using S1808 photoresist and BHF in step 3). The ribs on the photomask were made to be slightly wider at 2.2 μm to compensate for the undercutting that occurs during the isotropic etch of the oxide mask. Approximately 0.05 μm of oxide remains above the double slab coupler section for protection during the subsequent silicon etch since the thickness of the separating oxide layer is less than half of the thermal oxide mask. The photoresist was then removed in acetone and methanol.

As shown in step 4), the input waveguides were etched to a depth of 2 μm using TMAH at 60 °C, creating ribs with sloped sidewalls and a bottom width of ~5 μm. The remaining oxide above the coupler section and the input waveguides was removed in BHF in step 5). The last illustration shows two of the input waveguides coupling into a double slab section whose transition edge has been angled to vary the coupler length. The final device has 99 input waveguides spaced by 200 μm. Each input rib will excite a double slab coupler with a length varying by $\sim$1500 μm between the first and last input waveguide. The actual length for each coupler is determined by the location of the cleaved output facet.

SEM images of the resulting device are shown in Figure 4.5. Note that the transition from the single-layer input waveguide to the double slab coupler section does not appear to be ideally smooth. This was not expected to affect device performance as the optical mode propagating in the rib is not perturbed by these external irregularities as it excites the lower multimode slab waveguide.

Light is coupled into the device using an optical fiber, while the output light is collected using a microscope objective lens. Optical-quality facets are required at both ends of the device to ensure sufficient transmission of light. Both the input and output facets were prepared using the mechanical thinning and cleaving technique. Here, crystal bond is melted onto a brass block that has been heated to 160 °C on a hotplate and then the sample was mounted facedown onto the block. This encapsulates the etched features to protect them during thinning. The sample was thinned to a thickness of $\sim$250 μm, as measured by a micrometer depth gauge, using a rotating silicon carbide sandpaper disc [200 grit] and water. The sample was then carefully removed from the block by heating it on a 175 °C hotplate to liquify the crystal bond. Residual crystal bond was removed in acetone and methanol, followed by a final piranha clean.

Thinning the sample assists the cleaving process by reducing the force required to propagate the cleave along a crystallographic plane, starting from an initial scribe. The sample was scribed on its top surface using a Karl Suss wafer scriber that has

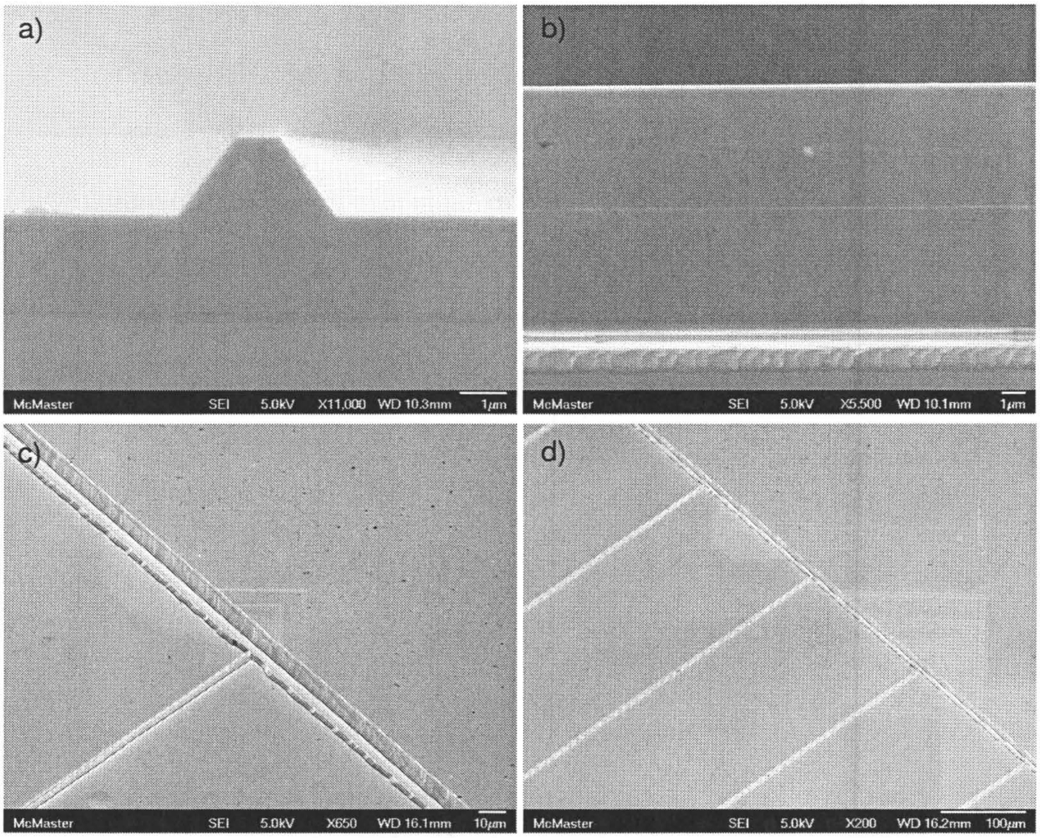

Figure 4.5: SEM images of a fabricated slab waveguide coupler
a) Input waveguide facet b) Coupler output facet c) Transition from single-mode rib waveguide to stacked slab waveguides d) Multiple waveguides coupling into varying double slab coupler lengths

a micrometer-controlled rotational stage, a weighted diamond tip and a microscope with cross-hairs to accurately place the scribe mark. This 1-2 mm long mark was made from the edge of the sample using a few passes of the diamond tip placed in contact with the sample surface. Its location for the output facet was positioned such that the double slab coupler length was $\sim$4000 μm in the middle of the sample. The location of the input facet is not critical.

The facets were then cleaved by pressing down on the backside of the sample, above the scribe mark, using a needle mounted on a 3-axis stage. This cleaves the sample along the crystallographic plane perpendicular to the input waveguides. The sample was placed facedown on a rubber pad covered with tissue, with a thin aluminum block on top to keep it in place. The location of the scribe mark was found using a microscope that was tilted to view the edge of the sample. Pressing down on the backside of the sample caused the cleaved ends to separate away from each other, preventing any damage that would be caused if the ends came in contact. The resulting facet was perfectly smooth and was expected to efficiently couple light from an optical fiber into the input rib waveguides.

4.5 Stacked Rib Waveguide Coupler Fabrication

Fabrication of the stacked rib waveguide coupler required both wafer bonding and accurate alignment during photolithography. These devices were fabricated using the double-SOI wafers with 5 μm thick silicon layers and the processing steps used are illustrated in Figure 4.6. The perspective of these illustrations has the output facet facing the reader.

As shown in step 1), the 0.13 μm thick oxide mask used to define the first set of waveguides and alignment marks was thermally grown and patterned using S1808 photoresist and BHF. The 2 μm wide rib waveguides were spaced by 200 μm with 100 μm wide relief regions placed in between. These relief features later acted as bonding regions providing mechanical support for both waveguide layers. Once again, the ribs on the photomask were 2.2 μm wide to compensate for the undercutting of the oxide mask, as was the case for each of the fabricated coupler designs. This pattern was then etched using ICP-RIE to a depth of 2 μm in step 2). The photoresist was not removed before this etch to protect the oxide mask

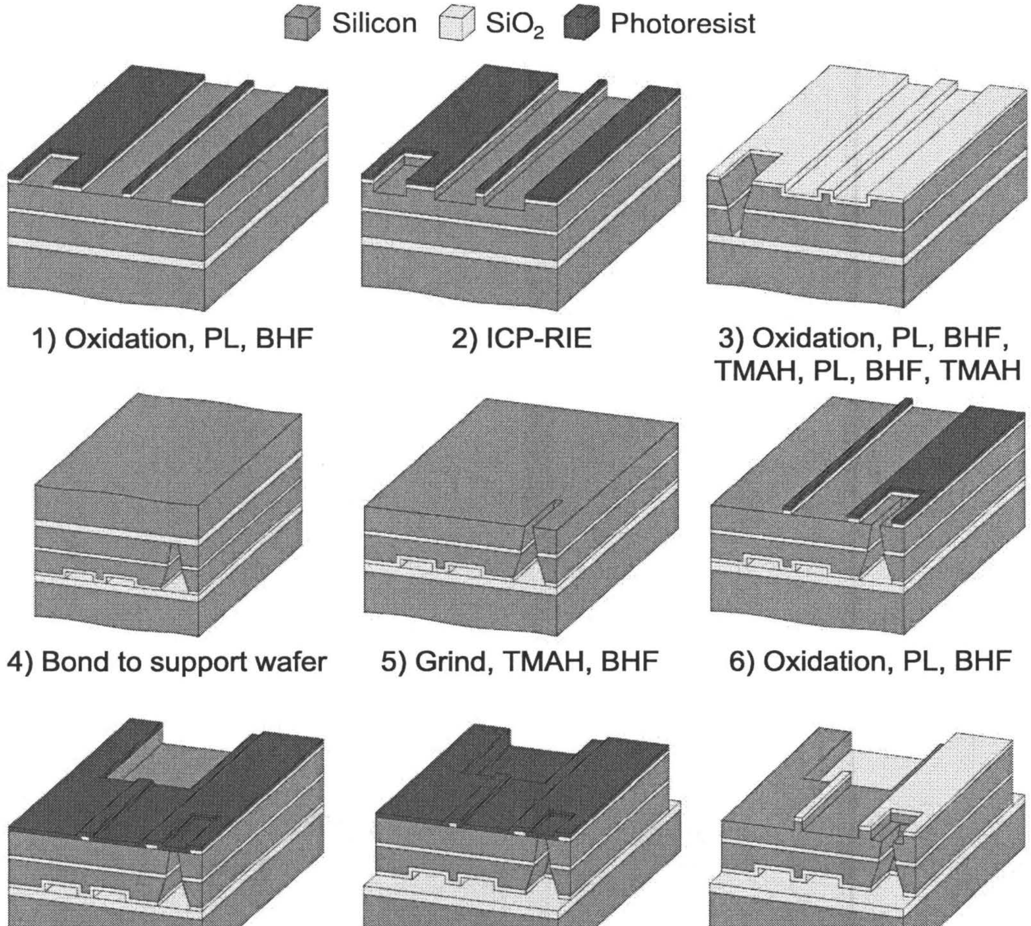

Figure 4.6: Proof-of-concept fabrication process for stacked rib waveguide couplers

from plasma-induced damage, as it would later serve as a bonding surface. The plasma-hardened photoresist was subsequently removed in a piranha bath.

As shown in step 3), a thin 0.05 μm oxide mask was thermally grown at 1200 °C for 8 min to protect the waveguides during subsequent etching steps. The alignment marks have to be etched through both silicon layers to ensure their visibility when the second set of waveguides is patterned. This was accomplished using photolithography and BHF to open the areas just above the alignment marks. Thicker S1827 photoresist was used for improved coverage over the existing features. The alignment marks were etched through to the separating oxide layer using TMAH at 60 °C, and then the photolithography, BHF and TMAH process steps were repeated to etch them through to the buried oxide.

TMAH etching was used in this step to avoid having to use the ICP-RIE at the University of Western, as the sidewall profile of the alignment marks is not important. The alignment features were widened during the drawing of the photomask to account for the slopping of their sidewalls during the TMAH etches. As well, the thickness of the two silicon device layers were each measured at nine areas spaced around the sample. These measurements were subsequently used to determine the etch depth required for the second set of rib waveguides in order to obtain equally matched slab regions for each layer.

The double-SOI sample was then bonded facedown to an oxidized support wafer in step 4), using the previously described plasma-activated wafer bonding technique. The handle of the double-SOI sample was mechanically thinned using silicon carbide sandpaper until $\sim$100 μm remained, as shown in step 5). The remaining silicon was etched in TMAH at 90 °C until the buried oxide layer is reached. This oxide layer was removed in BHF, revealing the second silicon guide layer and the alignment marks that had been etched through.

These alignment marks were used to register the second set of waveguides to the first for the patterning of the 0.13 μm thick thermal oxide mask in step 6) using S1808 photoresist and BHF. This pattern of waveguides consisted of nine sets of ribs that have been intentionally staggered from the lower waveguides to compensate for misalignment during this crucial photolithography step. Each set of 11 waveguides was staggered in increments of 1 μm to compensate for up to ± 4 μm of lateral misalignment. The photoresist was then removed in acetone and methanol.

As shown in step 7), the length of the coupler section was defined using thicker S1827 photoresist and BHF. The oxide mask defining the upper rib waveguide was etched back to the transition where the single rib waveguide in the lower layer meets the two stacked rib waveguides. The location of this transition differed for each of the 11 waveguides in the nine misalignment sets, allowing the coupling efficiency to be later measured as a function of coupler length. The upper silicon layer was then etched using ICP-RIE to a depth of 3.5 μm, with the coupler section protected by photoresist. The remaining silicon would be removed during the etch of the upper rib waveguide. The hardened photoresist was then removed in a piranha bath.

The input and output facets of the device were etched using ICP-RIE in step 8). The existing features were protected by a thick layer of S1827 photoresist, with each facet patterned using a photomask with a long straight chrome edge. The location of the output facet determined the length of the coupler sections. The plasma etching was carried through the separating oxide layer and the lower silicon layer until the oxidized support wafer was reached. A piranha bath was used to remove the hardened photoresist mask. The final ICP-RIE etch of the upper rib waveguides was completed in step 9), using the previously patterned oxide mask as

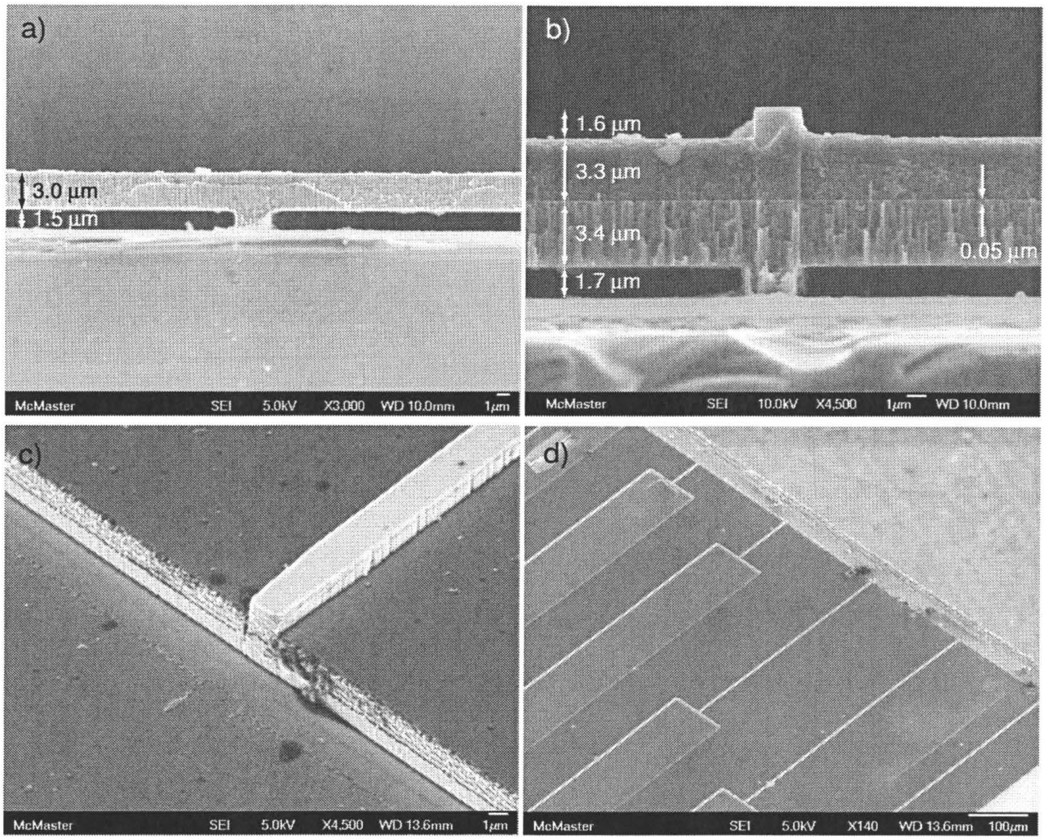

Figure 4.7: SEM images of a fabricated stacked rib waveguide coupler
a) Input waveguide facet b) Coupler output facet c) Transition from single-mode rib waveguide to stacked rib waveguides d) Various coupler lengths

protection. The etch depth was chosen such that the slab regions of the upper and lower waveguides were the same thickness.

SEM images of the resulting device are shown in Figure 4.7. Note that there is some roughness on the input and output facets, and on the waveguide sidewalls as well. This roughness is due to the scalloped etching of the ICP-RIE process, since the etching and passivation gases are alternated. The resulting scallops repeat every 50-60 nm, corresponding to a 2 s etch time at 1.5 μm/min followed by 1 s of passivation. There had also been some random roughness along the edges of the photoresist mask, which is inherent in the contact photolithography process used. This mask roughness translated into the silicon during the plasma etch.

The coupler shown in these SEM images is the best aligned device on the sample, with only 0.4 μm of lateral misalignment between the upper and lower rib waveguides. The slab thickness for each guide differs by only 0.1 μm, which is the approximate error for these measurements. The separating oxide thickness cannot be accurately measured using these images, but was measured by the manufacturer to be nominally 49 ± 1 nm. Both rib widths were measured to be slightly larger at 2.3 μm, instead of the expected 2 μm.

While the facets of the two waveguide layers have been completely etched with RIE, the remaining support wafer prevents access to these facets with an optical fiber. A dicing saw was used to trim the support wafer $\sim$20 μm away from each facet surface. The diamond resin blade used for dicing was 100 μm thick with a diamond grit size of 9 μm. The spin rate of the water cooled blade was 30000 rpm and the sample feed rate was 2 mm/s. The dicing saw stage can be moved in 2 μm increments, allowing accurate placement of the blade.

4.5.1 Complications During Fabrication of the Stacked Rib Waveguide Coupler

Aside from the challenge of developing a reliable wafer bonding technique, other problems were encountered during the fabrication of the stacked rib waveguide coupler. The preparation of optical-quality facets proved to be the greatest challenge in fabricating these devices. Ideally, the waveguides would be planarized before bonding to the support wafer using either a deposited oxide or a spin-on-glass, followed by chemical mechanical polishing. These planarization techniques were not accessible during this project, resulting in air gaps surrounding the encapsulated rib waveguides following wafer bonding. These air gaps caused severe chipping of the facets when either mechanical cleaving or a dicing saw was used to directly cleave or cut through the waveguide layers. This chipping was most likely caused by insufficient mechanical support of the waveguide layers.

The alternative was to use plasma etching to create vertical input and output facets. While the roughness of these facets would result in additional optical loss, this technique was the only remaining option. This was initially attempted using photoresist to pattern the facets onto waveguides whose ribs had already been etched. Thicker S1827 photoresist was chosen to better planarize the etched waveguides and the photomask with a straight chrome edge was used for patterning. The resulting resist mask had large notches in front of the ribs, which translated into the silicon during the following plasma etch. SEM images of the resulting etched facet is shown in Figure 4.8. Note that the sidewalls of the rib waveguides of these initial devices were sloped since TMAH was used for the waveguide etches to minimize usage of the ICP-RIE at Western. Partial delamination of the photoresist along the facet edge also occurred during etching, resulting in undesirable etching of the upper rib waveguide in certain areas. Optical characterization of these devices was attempted, but the notches prevented any coupling of light into the rib waveguides, although some slab propagation was observed.

The notches in what should have been a straight facet edge were caused by the air gap between the photomask and the photoresist during UV exposure. This air gap is due to the photoresist not adequately planarizing the rib waveguides that had been etched to a depth of 2 μm. The solution to this problem was to first pattern and etch the oxide mask for the rib waveguides, remove its photoresist, then define

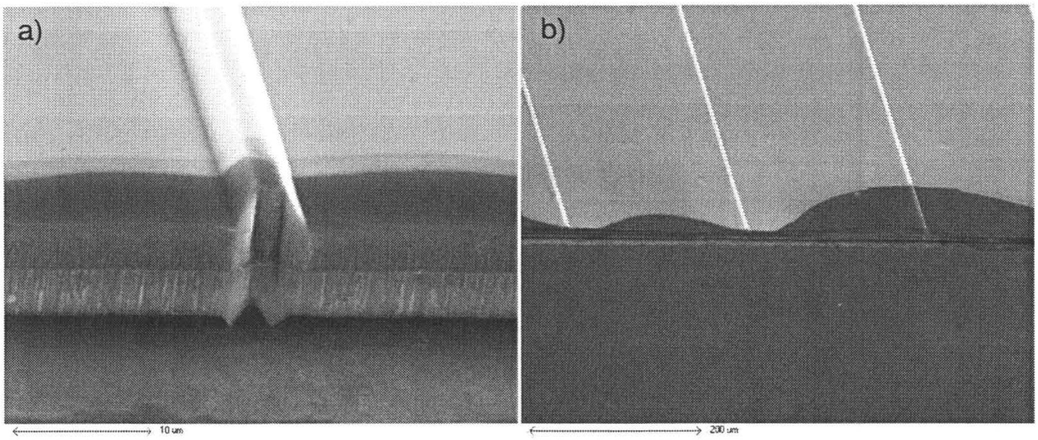

Figure 4.8: SEM images of the initial attempt at plasma-etched facets
a) Notch in the output facet of the stacked rib waveguide coupler b) Undesirable facet etching due to delamination of photoresist from the facet edge

and etch the facets *before* etching the rib waveguides in the final step. The thick photoresist for the facet mask easily planarized the ribs of the oxide mask since they were only 0.13 μm thick. This eliminated the air gap between the photomask and the photoresist during facet patterning. Once the facets have been etched to the support wafer, the photoresist was removed and the waveguides were etched to a depth of 2 μm using the previously patterned oxide mask that protects the top of their ribs.

Ideally, the entire process would have been completed using only chemical etching, resulting in perfectly smooth sidewalls for the rib waveguides. But, etching the facets before the final waveguide etch exposes the lower waveguides, which were previously encapsulated. If a chemical etch had been used for the final step, it would not only etch the lower waveguide but also cause slopping of the already etched facets. This meant that plasma etching had to be used for both the upper and lower waveguides, resulting in vertical sidewalls with some roughness. These devices also have a slightly shorter coupling length than the originally planned devices with sloped sidewalls. Plasma etching is the preferred process for the commercial manufacturing of photonic devices since it can be used to create curved features along any crystallographic plane.

4.6 Vertically-Integrated MMI Coupler Fabrication

A proof-of-concept fabrication process was used to demonstrate vertical optical coupling using an MMI structure, and is illustrated in Figure 4.9. This process did not require any wafer bonding and uses only single-layer SIMOX wafers with a 4.95 μm thick silicon device layer. Samples from these wafers were first thermally oxidized with a 0.05 μm thick oxide layer. This oxide layer would be used to define the transition from the input waveguide into the MMI coupler section and was patterned using S1808 photoresist and BHF, in step 1). The photomask used for this step contained 11 different coupler lengths repeated over 9 sets of rib waveguides. The photoresist was then removed in acetone and methanol baths before proceeding with the TMAH etch.

As shown in step 2), the silicon device layer was etched to a depth of 3 μm in TMAH at 60 °C. The remaining silicon was used to fabricate the input waveguide. Chemical etching was chosen for this step to ensure a smooth surface on top of the rib waveguide to be etched later. This was at the expense of creating a sloped facet leading into the MMI coupler section.

The initial oxide mask was removed in BHF before thermally oxidizing the samples at 1200 °C for 82 min, resulting in a 0.22 μm thick oxide mask. This oxide thickness was chosen by assuming a conservative selectivity of 20:1 for the ICP-

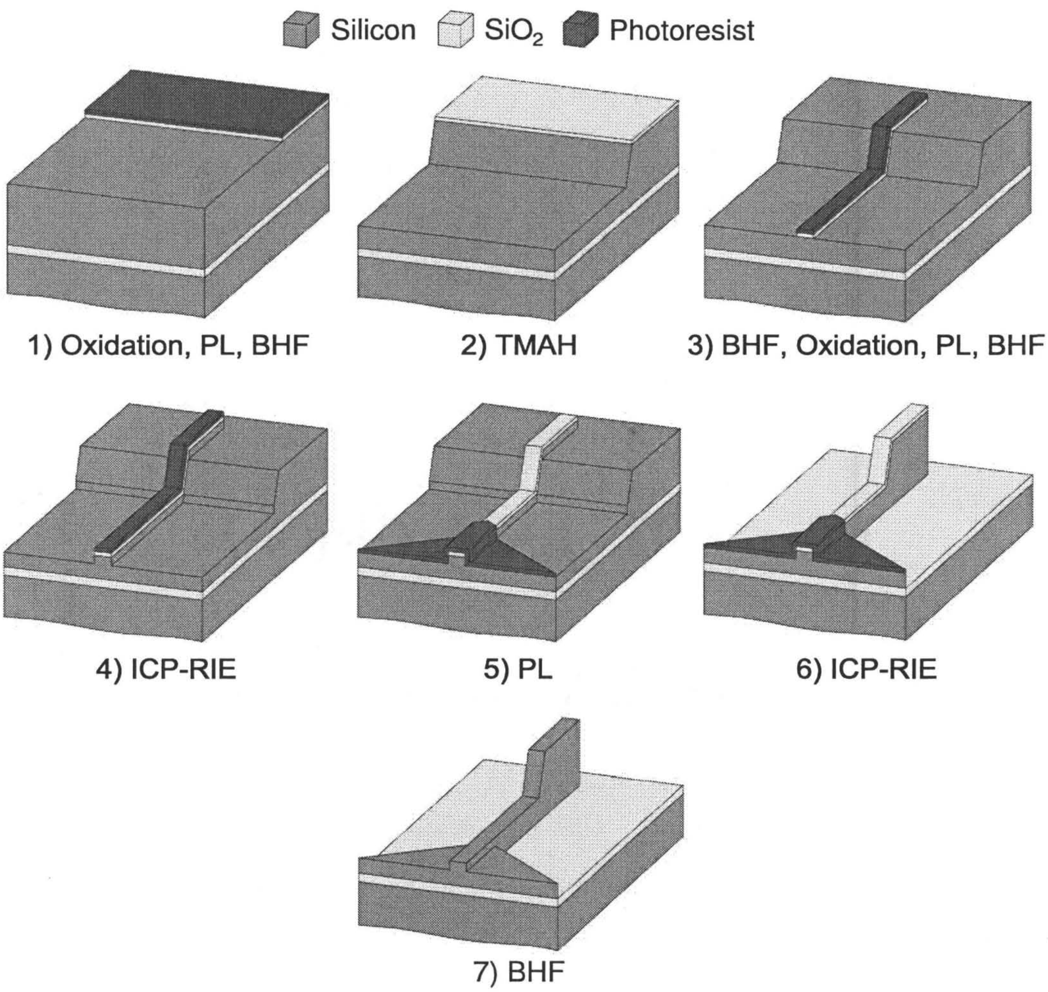

Figure 4.9: Proof-of-concept fabrication process for vertically-integrated MMI couplers

RIE process, as it was left exposed to the plasma during a 4.5 μm deep silicon etch. A 0.15 μm thick oxide would be the minimum thickness required if we assume a more typical selectivity of 30:1. While the thermal growth of the oxide mask consumes $\sim$0.1 μm of the silicon device layer, this type of mask was chosen as the selectivity of the etcher was less than 15:1 for unannealed CVD oxides. This would necessitate a very thick CVD oxide mask that would result in unacceptable undercutting during the BHF etch used to define the rib waveguide width.

Rib waveguides were patterned into the oxide mask using S1808 photoresist and BHF, as shown in step 3). The output facet of the MMI coupler section was also patterned in this step, defining its overall length. The photoresist was first exposed to UV light through the waveguide photomask and then through a photomask with a straight chrome edge to define both the waveguide and the output facet before developing.

ICP-RIE was used to etch 0.5 μm of silicon, creating the rib of the input waveguide, as shown in step 4). The hardened photoresist was then removed in a piranha bath. As shown in step 5), triangular sections were patterned above the rib waveguides using S1808 photoresist. This photoresist mask defines the tapered mode converters used to ensure that only the fundamental mode of the channel waveguide that couples into the MMI section is excited. Lateral misalignment during this process step was compensated for by using staggered waveguides and uniformly spaced triangular sections. The previously patterned waveguides were staggered in increments of 0.2 μm for each of the nine misalignment sets, compensating for up to $\pm$0.8 μm of misalignment. The photoresist mask protected the tapered mode converter regions during the ICP-RIE etch of the MMI coupler section to the buried oxide in step 6). A piranha solution was used to remove the hardened photoresist and then the oxide mask was removed using BHF to reveal the completed device in step 7). SEM images of the resulting device are shown in Figure 4.10.

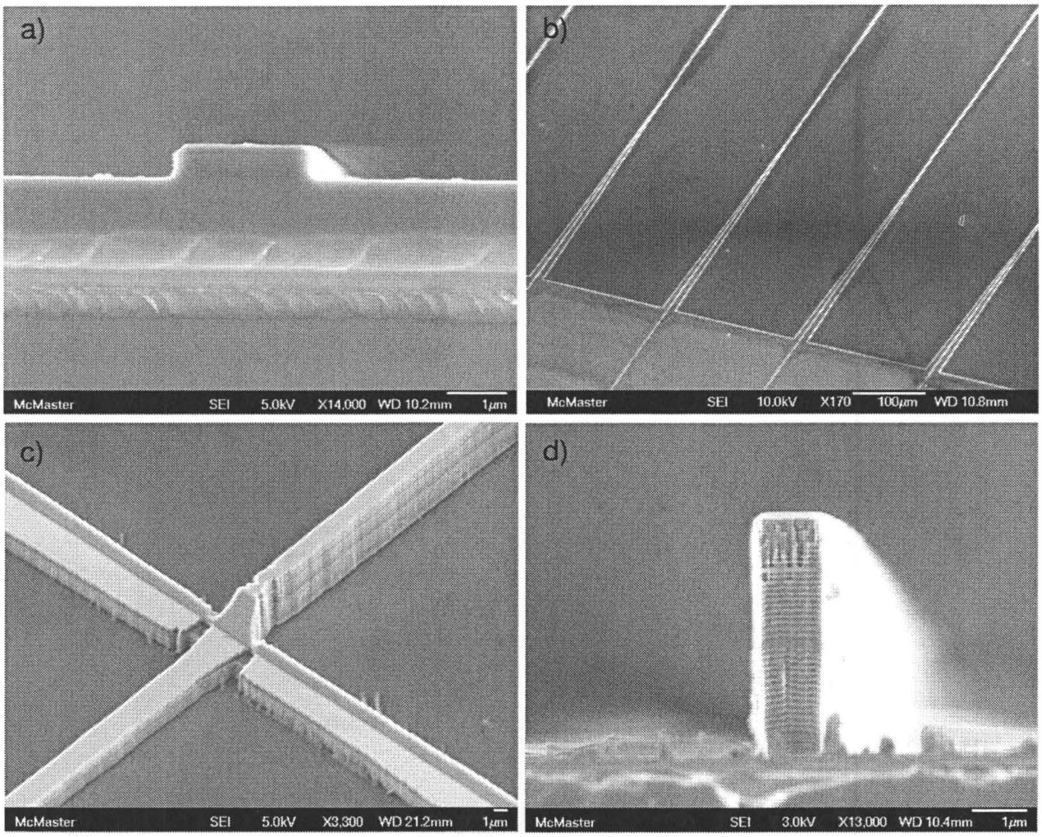

Figure 4.10: SEM images of the fabricated vertically-integrated MMI coupler
a) Facet of input rib waveguide b) Tapered mode converters c) Transition from channel waveguide to MMI section d) Output facet of MMI section

Note the less than optimal patterning of the transition from the lower input channel waveguide into the multimode coupler section. This irregular pattern arose from patterning the waveguide photoresist mask over a previously etched ledge. The resulting transition decreases the amount of light transmitted into the MMI section, but was not expected to significantly affect its coupling behavior since it only perturbs the optical modes over a short length. Also, the MMI section width at the output was only $\sim$1 μm. There was a slight tapering of the coupler section width from its nominal 2 μm over the last 10 μm of coupler length. This was apparent even before etching during the patterning of the facet photoresist mask, which was completed along with the waveguides. Once again, this was expected to only effect the coupling over a short distance. The roughness of the plasma-etched RIE facet was apparent, in contrast to the input facet that had been prepared using sample thinning and edge cleaving. The decision of using different facet preparation techniques for each facet is explained in the following section.

4.6.1 Complications During Fabrication of the Vertically Integrated MMI Coupler

Facet preparation proved to be the greatest challenge in fabricating the MMI coupler devices. Mechanical cleaving could not be used as the plane of the cleave does not always follow the initial scribe. The length of the relatively short MMI coupler section is crucial to device performance, thus a dicing saw was initially used to create the output facet of the device. While the 2 μm placement precision of the dicing saw should have resulted in an accurate coupler length, the vibration of the blade caused the MMI sections to chip at various lengths, as shown in Figure 4.11. The dicing saw had been successfully used for the facet preparation of conventional rib waveguides, but those structures have slab regions that provide mechanical support. The MMI coupler sections do not have surrounding slab regions, and this is the most probable cause for the chipping of these tall and narrow waveguides.

To overcome this problem, plasma etching was used to etch the output facet of the device to the buried oxide, at the expense of additional optical loss due to waveguide edge roughness. The output facet was patterned during the waveguide

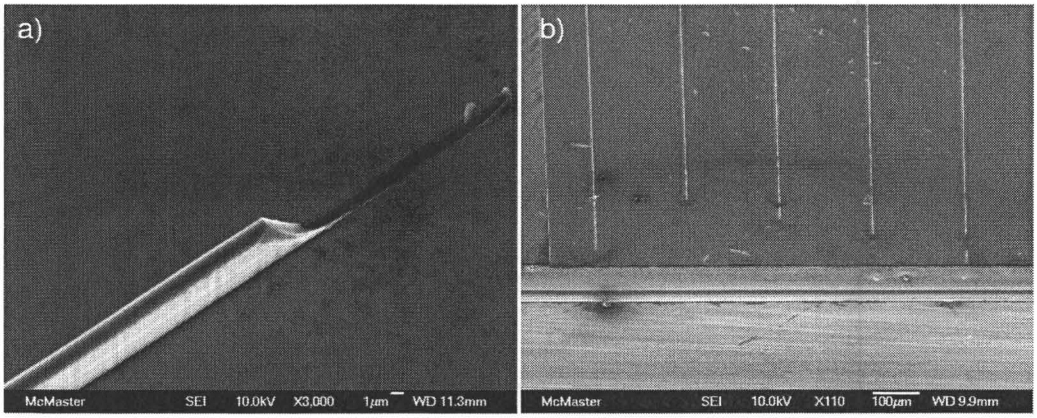

Figure 4.11: SEM images of sawed MMI coupler facets
a) Close-up of chipped output facet b) Multiple chipped waveguides

photolithography step, such that the waveguide sidewalls and the facet edge would be etched simultaneously. The remaining handle wafer was then trimmed $\sim$20 μm away from the output facet edge using the dicing saw. The blade was sufficiently far away from the MMI sections such that its vibration did not cause any chipping. The input facet was prepared using the mechanical thinning and cleaving technique, since its position was not critical. The advantage of cleaving the input facet was that it would be atomically smooth, ensuring efficient coupling of light from an optical fiber.

The first attempted fabrication process used RIE to thin the silicon above the input waveguide, as opposed to using chemical etching. This process step was completed after the waveguide sidewalls had been etched. These initial devices did not efficiently guide light, and the roughness on top of the input waveguide ribs created by the plasma etching was suspected as the main loss mechanism. An alternative process using TMAH to thin the silicon above the input waveguides was then attempted. This step had to be completed before patterning the waveguides, which meant that the lower rib section would be patterned with an air gap between the photomask and the photoresist. Despite this potential issue, the resulting devices had consistent rib widths for both the input waveguide and the MMI coupler section and were found to efficiently guide light.

This proof-of-concept vertically-integrated MMI coupler does not have the upper waveguides required to demonstrate actively-controlled optical switching. A fabrication process to create such a device, but that does not require wafer bonding, is described in Appendix B. The fabrication process that could have been used if an aligned wafer bonder and planarization techniques were available is also presented in Appendix B. This process would be applicable for the manufacturing of commercial devices employing the vertically-integrated MMI coupler.

4.7 Stacked Channel Waveguide Coupler Fabrication

A proof-of-concept stacked channel waveguide coupler was fabricated using the process illustrated in Figure 4.12. This process is similar to the one used for the MMI coupler, with the exceptions of using double-layered SOI wafers as the starting material and an aluminum mask to pattern the waveguides. The stacked channel waveguide coupler for this project was fabricated using the double SIMOX wafers with 2.5 μm thick silicon device layers, separated by a 0.1 μm thick oxide.

In the first step, a 0.05 μm thick CVD oxide was deposited and patterned using S1808 photoresist and BHF to define the transition from a single channel waveguide to the stacked channel coupler. The position of this transition was varied on the photomask to define multiple coupler lengths on the same device. A CVD oxide was chosen to avoid consuming any silicon from the upper layer to ensure that both

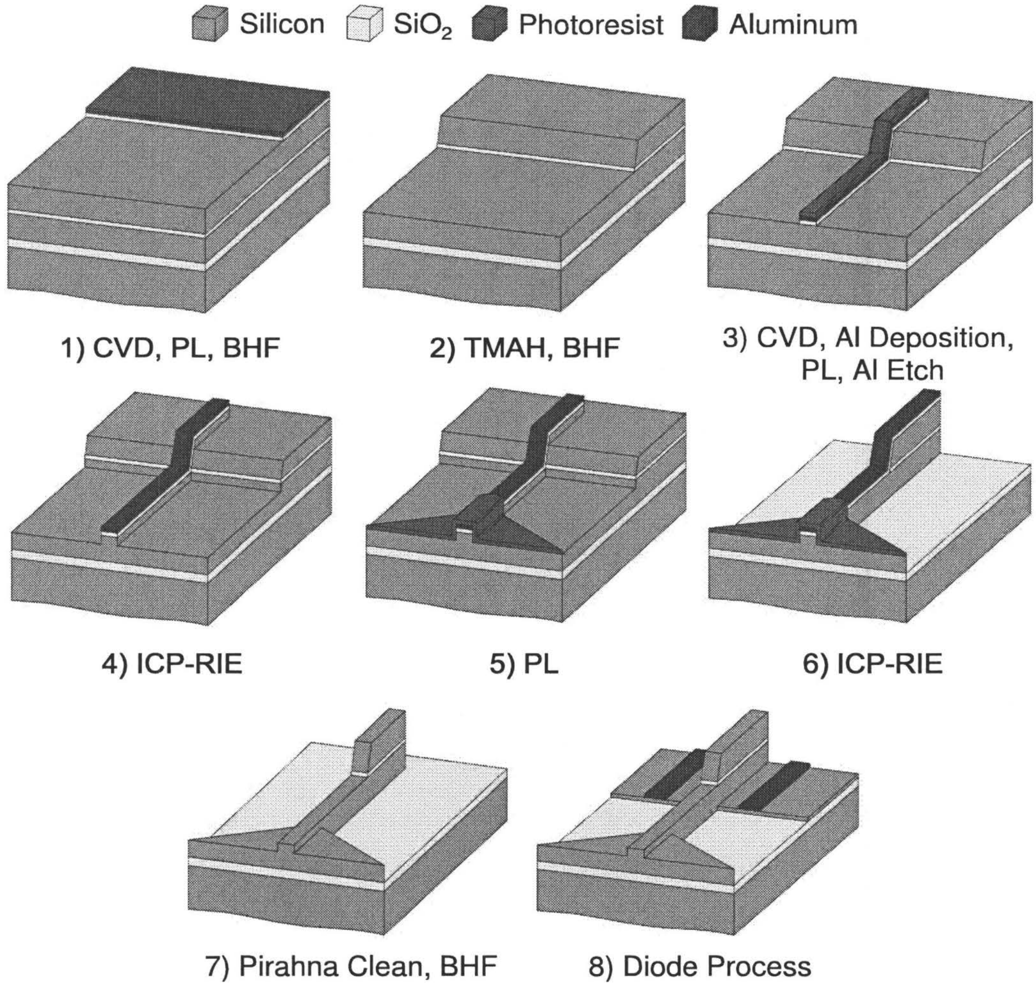

Figure 4.12: Proof-of-concept fabrication process for stacked channel waveguide couplers

layers remained perfectly matched. The photoresist was removed in acetone and methanol before etching the upper silicon layer in TMAH at 60 °C, in step 2). The separating oxide layer acted as an etch stop for this process step. This oxide was then selectively removed above the input waveguide region using BHF, along with the CVD oxide mask above the coupler section.

As shown in step 3), a 0.05 μm thick buffer layer of CVD oxide was deposited before depositing 0.1 μm of aluminum using a Bell Jar evaporator. This buffer layer prevented the aluminum from forming an alloy with the upper silicon layer. An aluminum mask was chosen here as it will not be significantly etched during the relatively long plasma etch process. If a CVD oxide mask had been chosen it would have to be unreasonably thick ($\sim$0.5 μm) to survive plasma etching through both silicon layers and the 0.1 μm of oxide in between. The channel waveguides, and their output facets, were patterned into the aluminum mask using a slightly thicker S1818 photoresist and aluminum etchant type A. This photoresist was removed in acetone and methanol before continuing with the ICP-RIE etch in step 4). This plasma etch was to a depth of 0.5 μm to create the rib of the input waveguides.

In step 5), the tapered mode converters have been patterned into S1808 photoresist with the same triangular sections used previously for the MMI coupler. The photoresist in this step was developed using CD-30 developer as opposed to 351 developer. This silicate-based developer does not attack the aluminum, which was observed using the sodium hydroxide-based 351 developer. When developing with CD-30, it was found through experimentation that the UV exposure time had to be doubled from the optimal time used with 351 to ensure complete developing.

This photoresist mask protected the input rib waveguides during the plasma etch of step 6). Part of the channel waveguide mask was exposed to the plasma as it etched through both silicon layers and the separating oxide in between. This deep etch required a robust mask, such as the chosen aluminum. In step 7), the photoresist and aluminum masks have been removed in a piranha bath, and the oxide buffer layer removed in BHF, to reveal the final device. Step 8) illustrates a device with a current-injection diode fabricated into $\sim$0.2 μm of silicon that could have been left behind during the plasma etch in step 6). This process allows either a passive device to be fabricated by etching to the buried oxide, or an active device by leaving behind a layer of silicon that is sufficiently thin to not affect optical coupling. The process steps that would be used for diode fabrication can be found in Appendix B.

SEM images of the resulting device are shown in Figure 4.13. Note the irregular transition from the lower input channel to the stacked channel coupler, the cause of which will be explained in the following section. The facets of the stacked channel devices were prepared with the same techniques used for the MMI coupler. The handle wafer was trimmed back from the output facet using a dicing saw and the input facet was prepared by mechanically thinning and cleaving the device.

4.7.1 Complications During Fabrication of the Stacked Channel Waveguide Coupler

The main challenge with the fabrication of the stacked channel coupler was developing a reliable process for the robust mask required during the etching of the waveguides. An oxide mask would have to made very thick given the selectivity of the plasma etcher and would have to be deposited using CVD to avoid consuming any silicon. Further, the etcher has a lower selectivity to CVD oxides than to thermally grown layers, eliminating oxide as a suitable mask candidate.

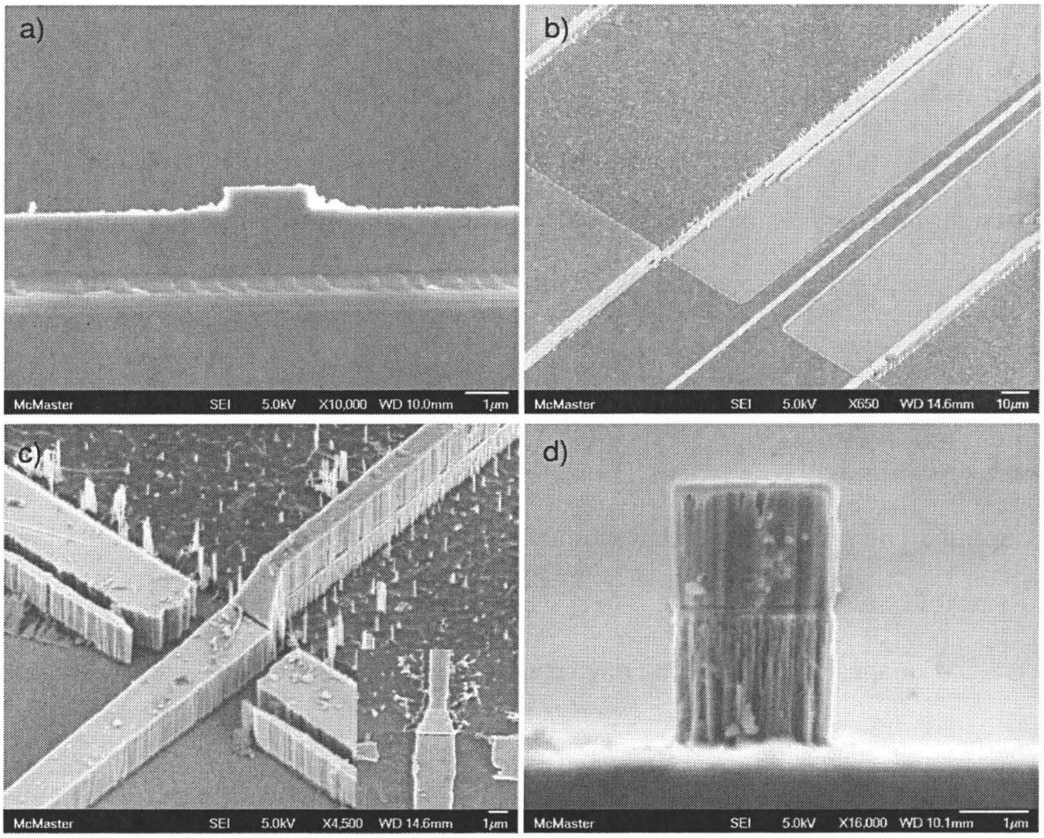

Figure 4.13: SEM images of the fabricated stacked channel waveguide coupler
a) Facet of input rib waveguide b) Input mode converter c) Transition from input channel waveguide to coupler section (top view shown in inset) d) Output facet of coupler section

Using hardened photoresist as a waveguide mask was tested without success. Here, the photoresist is left on after the first ICP-RIE step, then additional photoresist is spun on, exposed and developed to pattern the tapered mode converters. While the hardened photoresist did not appear to be altered during this second photolithography step, subsequent plasma etching resulted in sloped sidewalls. This meant that the edges of the hardened photoresist had eroded and continued to do so during the deep plasma etch.

A metal mask layer, such as evaporated aluminum, was a successful solution despite the new challenges that came with it. With the aluminum layer below the photoresist being highly reflective, the conditions required for optimal photolithography significantly changed. Patterning waveguides over the previously etched transition between the single input waveguide and the stacked channel coupler proved to be problematic. Using the standard recipe for S1808 photoresist resulted in a coupler section that did not join with the input waveguide over the transition edge. The coupler waveguides tapered to a point that was tens of microns away from this edge. Varying the UV exposure time improved the pattern near the transition but resulted in waveguides that where either under or over-developed. Thicker S1827 photoresist was also tried since it improves resist coverage over the transition edge, but the resulting waveguides were always too narrow.

A compromised thickness was then attempted using S1818 photoresist. The exposure time had to be varied to find the best case conditions for this photolithography process. These process conditions can be found in Table 4.2. The resulting waveguide width of 2 μm was correct for both the input waveguide and the coupler section but varied near the transition edge. The input channel waveguide widens to $\sim$2.5 μm at this edge, while the joining coupler section waveguides narrow to $\sim$1.5 μm, as shown in the inset of Figure 4.13 c). This mismatch in waveguide width will result in optical loss as well as the possibility of exciting higher order modes in the coupler channel waveguides.

Also note the unetched silicon surrounding the input channel waveguide at the transition ledge. This is due to the waveguide photoresist pattern being slightly underdeveloped leaving behind thicker resist that had accumulated along edges and in corners during spinning. This unetched silicon should not perturb the propagating mode since it is sufficiently far away from the channel waveguide. These problems with the transition are only a issue with our proof-of-concept fabrication process since the waveguide patterning is done over an etched ledge. A manufacturing process suitable for commercial devices employing the stacked channel waveguide coupler structure is presented in Appendix B.

Another issue with using an aluminum mask was that it caused micromasking during the plasma etching process, resulting in pillars of unetched silicon commonly referred to as "grass". This is due to the alloy that the aluminum forms with the silicon, and that remains even after chemically etching of the aluminum. To avoid this issue, a buffer layer of CVD oxide was used to form a protective barrier between the aluminum and the silicon below. The original process involved etching this buffer oxide using BHF with the aluminum and photoresist masks still in place. This sometimes resulted in the waveguide lines being lifted off. Hardening the photoresist before etching the oxide using a higher hotplate baking temperature and exposing it to a light oxygen plasma descum process resulted in some success, but these results were not consistent. A decision was made to etch through the buffer oxide using ICP-RIE, and then allow the process to continue to etch the silicon below. Despite these efforts, some grass was still observed on the samples after plasma etching, and is visible in Figure 4.13 c).

Chapter 5

Device Characterization

The optical coupling efficiency of the fabricated vertically-integrated devices was evaluated using an infrared laser and camera setup. The measurement techniques that were commonly used for each device will be described before presenting their individual results.

5.1 Measurement Techniques

5.1.1 Experimental Setup for Device Characterization

The optical coupling efficiency of each device was measured using the setup shown in Figure 5.1. Light from a 1550 nm diode laser is coupled into the input waveguide of the device under test using a tapered optical fiber with a 2.5 μm spot size and a working distance of 14 μm. The laser output power is typically set to 1 mW (0 dBm). The fiber and device are mounted on 3-axis linear stages with manually controlled differential micrometers, with the fiber stage having optional piezoelectric control as well. A ring-illuminated microscope is placed above the samples in order to visually position the tapered optical fiber.

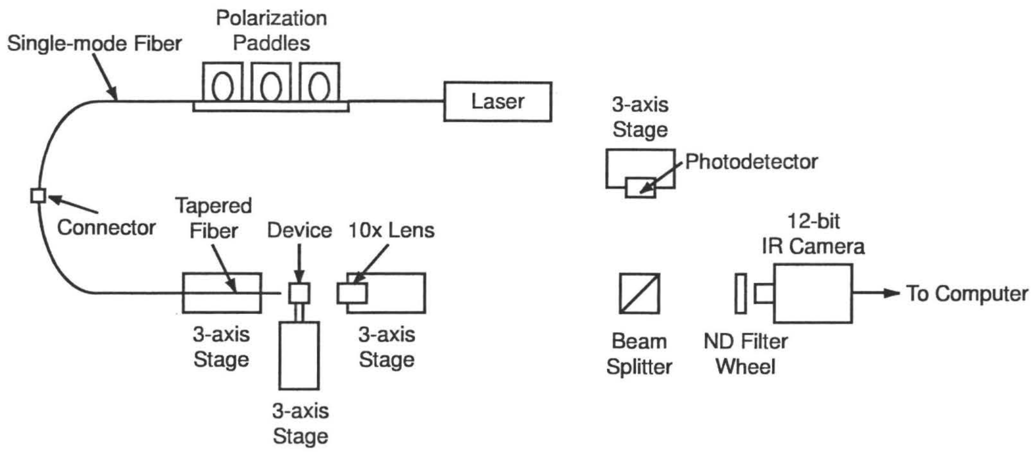

Figure 5.1: Setup used for device characterization

Polarization paddles connected in-line between the diode laser and the tapered fiber are used to control the input state of polarization. This type of polarization controller utilizes stress-induced birefringence to achieve the desired polarization state by tilting three independent spools of looped fiber. Light from the stacked waveguide coupler section is then collected using a 10× objective lens mounted on a 3-axis stage. The output optical field is imaged using a high-resolution (320 × 256 pixels), 12-bit cooled-CCD IR camera, model name Merlin from Indigo Systems. National Instruments Vision software is used for image capture and to perform image operations. A 50/50 beam splitter cube is routinely placed between the objective lens and camera to reflect half of the output light into an InGaAs photodetector with a sensor diameter of 3 mm. This photodetector provides accurate total output power measurements that are used to optimize the alignment of the input fiber. A polarizing beam splitter cube is also temporarily placed in front of the camera to set the input polarization state by tilting the individual polarization paddles. A neutral density (ND) filter wheel is routinely used to attenuate image brightness to prevent camera saturation.

5.1.2 Optical Field Image Analysis

It would be very challenging to distinguish the output light in each stacked waveguide using a slit or other aperture shape to physically block the output from one layer while capturing light from the other with a large-area photodetector. Another option would be to use a second tapered fiber to raster scan the output optical field to generate its profile using a fiber-coupled photodetector. Yet, the most efficient and highest resolution technique is to use an IR camera to capture the entire output field instantaneously. These images can be analyzed using software to generate line slice plots of the optical field and to integrate the image intensity over a defined area to determine the optical power in each output waveguide. The optical coupling ratio of two vertically-integrated waveguides was determined using:

$$P_c = \frac{P_{Upper}}{P_{Upper} + P_{Lower}} \times 100\,\% \tag{5.1}$$

where P_{Upper} is the optical power in the upper output waveguide and P_{Lower} is the optical power in the lower waveguide.

The 12-bit camera used in these experiments has a maximum number of $2^{12} = 4096$ possible brightness levels. In practice, the background field was around 1000 units and the maximum output signal was kept around 3500 units to avoid camera saturation. This still yields an experimental range of nearly 2500 levels, with each pixel carrying a corresponding relative error of less than one percent $(1/2500 = 0.04\,\%)$.

The optical power in each defined region was calculated by multiplying the

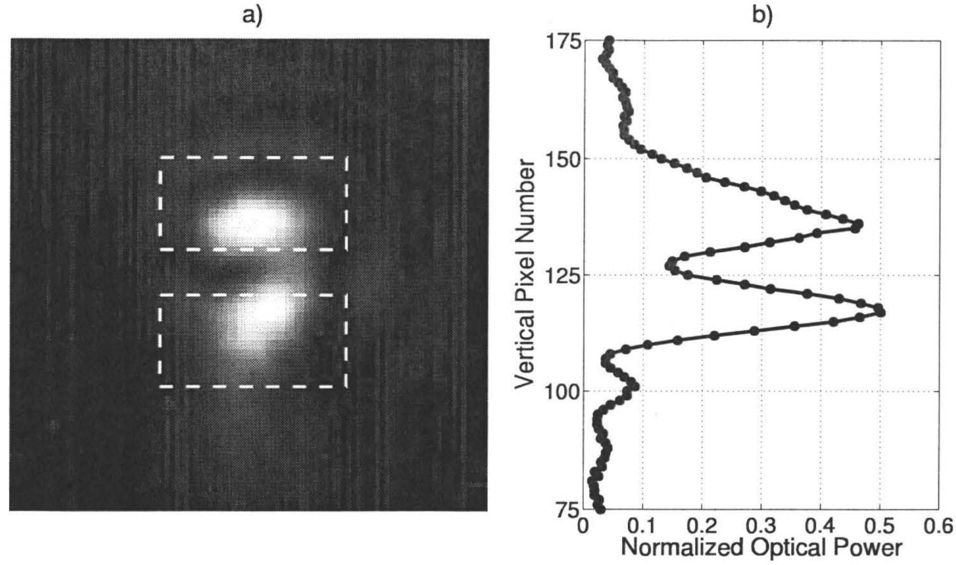

Figure 5.2: Image analysis of an MMI coupler output
a) Output image with integration boxes b) Vertical line slice plot through center of lobes

mean intensity within the integration box with its area. A uniform average background intensity was first subtracted from each pixel before integration. The background image was acquired with the input laser off, and represents the dark current of the camera and any stray infrared light. An example of this image analysis is provided in Figure 5.2, using the output of an MMI coupler whose length results in approximately a 50/50 split of the input signal.

Here, the image has been cropped to 100 $\times$ 100 pixels from the original size of 320 $\times$ 256. An objective lens with a higher magnification could have been used to better fill the acquisition area, but difficulties with focusing both output lobes simultaneously were encountered with a 20$\times$ objective. The image resolution of this 5 μm tall device is approximately 10 pixels per μm vertically. The resulting output field is nearly a two-fold image of the input field launched into the lower region of the MMI coupler section, and the integration boxes are chosen to be 2 μm in height with 1 μm in between. The line slice plot shown has been taken through the center of the lobes along the vertical axis and is normalized to the coupling ratio of 51$\pm$3 %, as measured by the integration boxes.

The position of the output image can shift vertically for each coupler length measured due to non-uniform sample thickness and small angles in the stage translation. This necessitates adjusting the placement of the integration boxes for each captured image. For devices with a separating oxide between the waveguides, the lobes of the output field can be easily distinguished. This allows the integration boxes to be placed visually using the midpoint between the two lobes.

The output field of the MMI coupler does not have easily distinguishable lobes since there is no separating oxide. In this case, the vertical position of the image was kept constant for each measurement to ensure consistent placement of the integration boxes. To adjust the image position for each coupler length, the input optical fiber was first lowered to illuminate the substrate below the waveguide. This revealed the upper plane of the oxide and silicon just below the multimode coupler. The location of this plane was used to shift the image vertically, by moving the objective lens, to the exact center of the image acquisition window.

5.1.3 Additional Optical Measurements

A large-area photodetector was used during alignment of the input fiber to maximize the transmitted optical power. These measurements can also be used to quantify the excess optical loss of the coupler devices as a function of their length, in comparison to the power transmitted through a waveguide without a coupler section. Unfortunately, these excess losses were sporadic when measured for various coupler lengths, instead of the expected linear increase with length. In fact some shorter couplers would have output images that were barely detectable, while a longer adjacent coupler would transmit enough light such that the neutral density filters had to be used to avoid camera saturation.

These sporadic loss measurements arise from the variations in facet quality between different coupler lengths, from disrupted waveguides caused by voids introduced during photolithography, from inconsistent sidewall roughness and finally, from variations in the transition from the single-layer input waveguide to the stacked waveguide coupler section. Patterning waveguides over the transition edge proved to be problematic as the waveguide width tapered near this edge, and random variations between adjacent coupler lengths were observable even with an optical microscope. Even if the facet quality and coupler transitions were consistent, the length of the input waveguide was varied to have different coupler lengths on the same sample. Therefore, the optical loss accrued along the input waveguide changes and cannot be easily factored out of the excess loss measurement for each coupler length.

The ideal technique to measure excess optical loss in this case would be the cut-back method. Here, the transmitted optical power is repeatedly measured for the same waveguide with the output facet being trimmed between each measurement. This reduces the coupler length while maintaining a constant length for the input waveguide. Polishing techniques should be used to obtain consistent facet quality for each measurement. Yet, the cut-back method is impractical for the devices in this project given all of the challenges encountered with facet preparation, as detailed in the previous chapter. Therefore, optical loss measurements will not be reported for the fabricated devices. The main goal of this project was to demonstrate proof-of-concept optical coupling between vertically-integrated SOI waveguides.

Our waveguides are expected to be very lossy due to the non-optimized patterning and etching processes used during this project.

The sensitivity of the coupling ratio to the wavelength of the input light was also quantified. Since all of the coupler devices were designed for the 1550 nm telecommunications wavelength, the majority of the coupling ratio measurements were obtained using a diode laser fixed at this wavelength. Wavelength sensitivity was evaluated using a separate tunable diode laser, whose wavelength can vary from 1530 nm to 1570 nm using a rotatable diffraction grating. The tapered input fiber was first positioned to optimize the transmitted power through a coupler length that provided the highest coupling ratio for that particular sample. Then the input wavelength was varied in increments of 10 nm, with the output image acquired for each wavelength. It is important to note that the amount of birefringence induced in the polarization paddles is wavelength dependent. These paddles were adjusted, with a cube polarizer placed in front of the camera, to maintain a constant state of polarization for each wavelength measurement.

All of the measurements were repeated over several trials to determine their uncertainty. When measuring the coupling ratio, the output image from each coupler length was collected for a particular set of 11 lengths, and then the measurements were repeated for the same set. This ensured that the input fiber and objective lens were realigned and refocused for each measurement. At least five separate images were collected for each coupler, with the uncertainty in the averaged coupling ratio determined by the standard deviation of the trials. This measurement procedure was the same for either a TE or TM-polarized input. The polarization state was adjusted between each set of measurements.

There were four main sources of error for the measured coupling ratios. Variations in the focus of the output image were most likely the largest contributor to this error. It was at times challenging to determine the optimal objective lens focus for both lobes of the output image and it seemed that a compromised focus had to be chosen. All of the proposed designs are also inherently polarization sensitive, with the performance of the MMI coupler being the least dependent. Even slight variations in the input polarization state can lead to measurable uncertainty in the measured coupling ratios. The excitation conditions for the input waveguide mode also varied as the fiber was realigned for each measurement. This was especially an issue with the stacked slab coupler and great care was taken to maintain consistent coupler excitation conditions with optimal fiber alignment. Finally, some of the error can be attributed to inconsistencies in the placement of the integration boxes during image analysis. This was especially the case with the MMI couplers since they do not have a separating oxide dividing the upper and lower waveguide regions, making them less distinguishable.

5.1.4 Laser-Induced Free Carrier Measurements

Actively-controlling the optical coupling for a fixed coupler length was attempted using laser-induced free carrier generation. An argon-ion laser with a wavelength of 514 nm and a maximum power of 130 mW was used for these experiments. The output light of the laser was coupled into an optical fiber using two gimbal-mounted mirrors and a gimbal-mounted objective lens and fiber holder pair. The coupling losses due to misalignment and numerical aperture mismatch resulted in a maximum transmitted power of 42 mW at the end of the cleaved fiber. The transmitted optical power was measured for various laser voltage settings using a silicon-based photodetector capable of measuring this wavelength. The optical fiber used to guide the 514 nm light had a core diameter of $\sim$4 μm (NA = 0.12) and a single-mode cut-off wavelength of $\sim$430 nm.

This optical fiber was mounted in a fiber holder that is positioned perpendicular to the device surface using a 3-axis linear stage. This setup is pictured in Figure 5.3 and was the same setup used for measuring the coupling ratio for each device. The tapered optical fiber delivering 1550 nm light was first aligned to the sample to obtain optical coupling as seen with the IR camera. A camera sensitive to the visible range was was then placed in the ocular of the microscope above the device to image its surface and the 514 nm spot from the Ar-ion laser. This laser spot was accurately positioned above the coupler section using the 3-axis stage and any change in the coupling ratio would be observed by the IR camera. The vertical translation of the stage allowed for control of the laser spot size at the device surface.

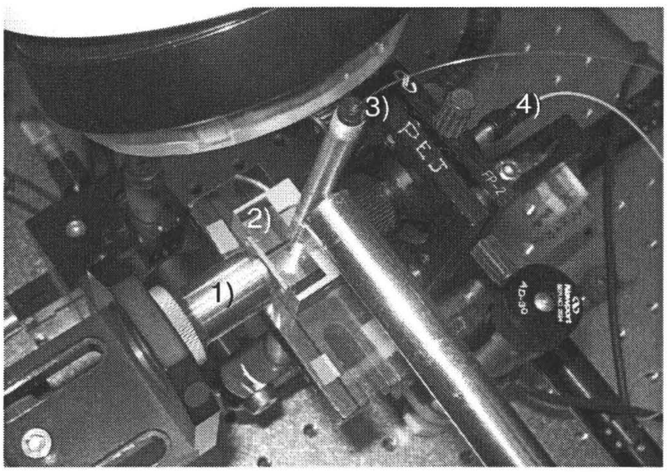

Figure 5.3: Setup used for laser excitation experiment

1) Objective lens for output light 2) Device under test 3) Optical fiber for 514 nm laser light 4) Tapered optical fiber for input 1550 nm laser light

5.2 Characterization Results

The following characterization results are reported for the best-case sample (the so-called 'hero' device) for each type of coupler fabricated. When intentional misalignments have been used, the set of 11 coupler lengths that provided the highest coupling ratios for a given sample are the devices that have been fully characterized.

5.2.1 Stacked Slab Waveguide Coupler

Example images from the output of a stacked slab waveguide coupler are shown in Figure 5.4. These images were acquired for the coupler lengths of 3535 μm and 4270 μm, yielding coupling ratios of 68±4 % and 51±4 % respectively, as measured by the integration boxes placed along the entire width of the acquisition window. The images here have been cropped to 200 $\times$ 100 pixels.

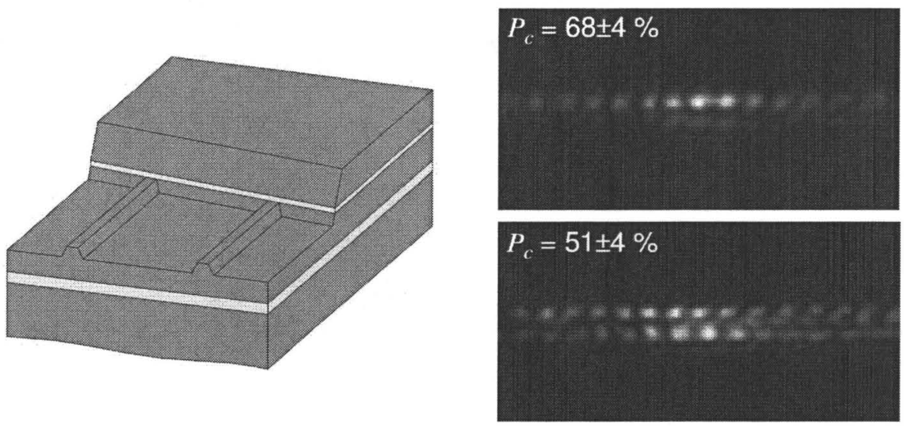

Figure 5.4: Example IR output images from the stacked slab coupler

An illustration of the device structure is also provided, showing two of the 99 possible input rib waveguides. Any of these input guides can be used to excite the lower slab mode of the coupler section whose length varies over the width of the sample. Since only the modes of the stacked slab waveguides are visible, and not the input rib waveguide mode itself, the input fiber is positioned such that the optical power collected at the center of the slab modes is maximized. This ensures that the mode of the lower slab waveguide is being excited by the rib mode alone, and not by light propagating in the side slab regions of the rib waveguide. Data was collected for every tenth input waveguide, resulting in an incremental change of 146 μm for the coupler length over a range of 3094 to 4541 μm.

The coupling ratios for both the TE and TM input polarizations are plotted in Figure 5.5 for various coupler lengths. A sinusoidal function expressing the dependence of the coupling ratio on coupler length has been fit to the TE-polarized experimental data and is in the form: $P_c(L) = F\sin^2(\pi L/2L_\pi)$, where L_π is the coupler length at which the maximum power transfer occurs and F is the maximum coupling ratio obtained, representing the degree of synchronization between the coupled modes.

Complete optical power transfer was not achieved, as the value of F was 68$\pm$4 % for the TE input polarization at a coupler length of 3535 μm. For the TM polarization, this peak would occur for coupler lengths longer than those fabricated and therefore could not be accurately determined. TE-polarized input light coupled more strongly than the TM state due to the orientation of their respective mode discontinuities, as discussed in the theory and simulation chapters. This was also experimentally observed by Soref, who was the first to publish coupling between the stacked layers of a double-SIMOX wafer [67].

This incomplete transfer of power can most likely be attributed to the non-uniform thickness of the two silicon slab waveguides. SEM measurements of the output facet found that the slab heights were 5.5$\pm$0.1 μm and 5.4$\pm$0.1 μm for the

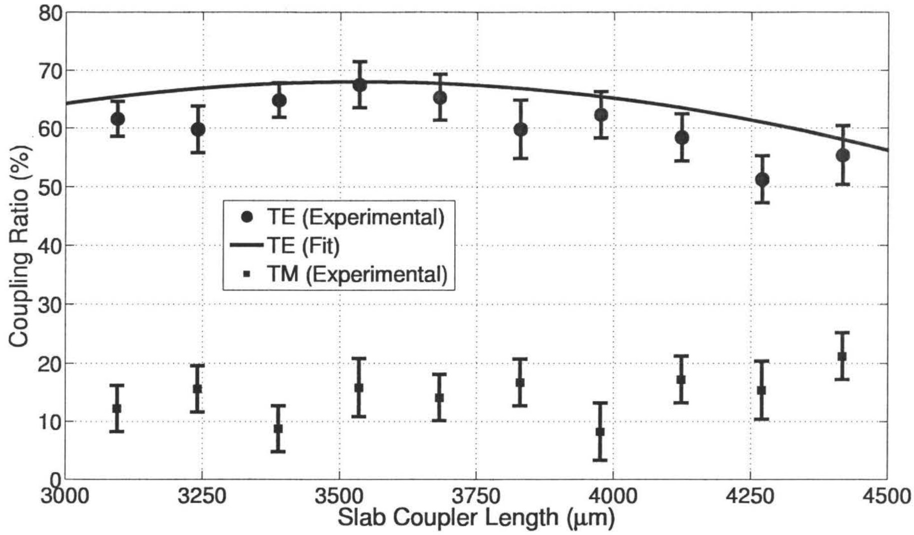

Figure 5.5: Coupling ratios measured for the stacked slab coupler

upper and lower layers respectively. This measurement was made near the center of the sample, so this thickness non-uniformity could have also varied along the width of the sample. While the thickness could have varied by up to 0.5 μm along the length of the coupler section according to wafer specifications, such a large variation would be unlikely given the short coupler length of ~4 mm.

This non-uniformity explains not only why complete power transfer was not achieved but also why the coupling ratios for each coupler length do not completely agree with the sinusoidal fit. As well, some of the slab regions of the input waveguide would have been excited by the light from the tapered fiber despite our best efforts, resulting in a slightly different excitation condition for each coupler length. The variability in the excitation conditions between measurement trials was most likely the largest contributor to the typical experimental error of ± 4 % for the measured coupling ratios.

Simulations for a slab coupler device with matching slab thicknesses of 5.5 μm predicted coupling lengths of 3970 μm and 21279 μm for the TE and TM polarizations, respectively. Further simulations showed that even a small variation of 0.1 μm in the slab thicknesses can reduce the maximum achievable coupling ratio to only 70 %, occurring at a coupler length of 3122 μm.

The wavelength sensitivity of the maximum coupling ratio obtained, occurring at the coupler length of 3535 μm, is plotted in Figure 5.6 for the TE input polarization. The coupling ratio optimized at 1550 nm was found to vary by 4 % over a wavelength range of 1530-1570 nm. This is more than the simulated variation of 2 %, but agrees within the experimental error of 2 % for these measurements. This error can most likely be attributed to small variations in the input polarization state between each trial. This error is less than that of the above coupler ratio measurements since the alignment of the input fiber and the focus of the objective lens remained fixed while the wavelength was varied.

The Ar-ion laser excitation setup was then used to attempt active-control of the optical power coupled into the upper waveguide. The 514 nm laser spot was first positioned above the input waveguide that was being used to excite the stacked slab coupler and was then translated along the length of the coupler. Unfortunately, no change in the coupling ratio was observed even for the largest laser output of 42 mW. Not even a change in the shape and location of the lobes of the two slab modes was observed. Changes to the lobes were easily observed when the input wavelength or polarization state was changed. Different fiber focal distances were attempted to vary the beam diameter, but no changes were observed with this or with lateral translations of the fiber position. As well, no free carrier-induced attenuation was observed when the laser beam was positioned either above the coupler section

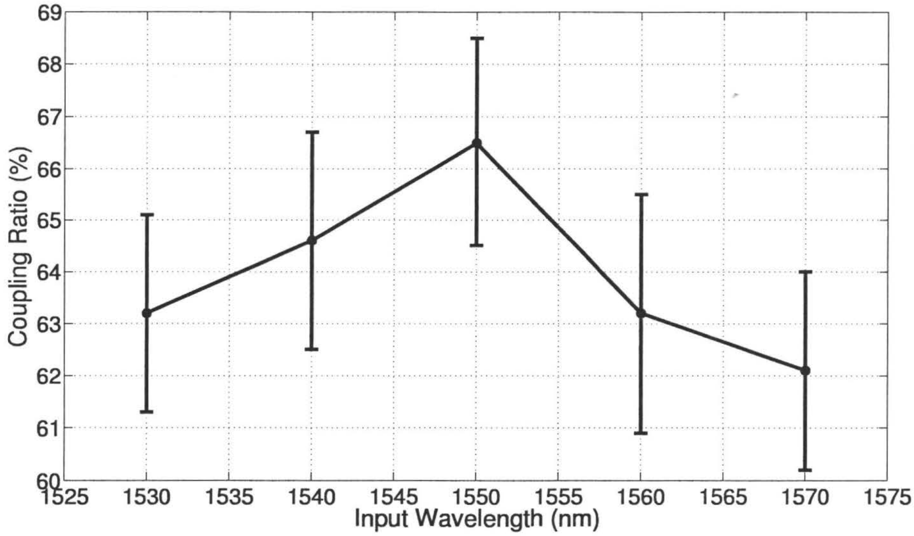

Figure 5.6: Wavelength sensitivity of the stacked slab coupler

or above the input waveguide.

It was suspected that the carrier lifetime in these devices was on the order of nanoseconds and therefore was not long enough to allow for significant diffusion of the carriers away from their generation site. The diffusion length for the generated carriers was probably at most tens of microns, which is not large enough to induce a significant change in the material index of refraction or absorption. This shorter than expected carrier lifetime is most likely due to a large amount of surface recombination. The total light that is absorbed in the upper 5 μm slab should be 26 mW for the maximum Ar-ion laser output of 42 mW. Since 80 % of this light will be absorbed in the first micron, the majority of carriers may have recombined at the surface interface before diffusing over any significant length.

The simulation model that was employed does not account for the dependence of the generation and recombination processes on the depth of the device layer. We also neglected Auger recombination, which may have also contributed to the shorter than expected lifetime, especially near the generation site. A 980 nm diode pump laser with a maximum output of 40 mW was also attempted, but it also resulted in no observable change. These laser parameters result in even less total power absorbed (6 mW) with 90 % of the absorption occurring in the first micron. A 632 nm HeNe laser could have also been used, but the lasers available were limited to a maximum output of only 7-10 mW.

It was surprising that a refractive index change was not observed in the slab couplers, since they are basically bulk SOI wafers. Typically, these wafers have relatively long carrier lifetimes on the order of microseconds. Not surprisingly, no effect was observed when these various laser excitation conditions were attempted for the other fabricated waveguide-based couplers. Waveguides with etched side-walls have even shorter carrier lifetimes due to increased surface recombination [80]. As well, while the entire laser beam is incident on the slab section, the channel waveguide couplers will only be excited by a portion of the beam since their widths are smaller than the spot diameter.

Ideally, a bus waveguide made in material that guides visible light, such as silicon nitride, would be placed above to uniformly couple visible light along the entire length of the coupler section, as proposed in [77]. Concentrating all of the optical power in a highly localized spot results in increased carrier recombination and hence insufficient diffusion length. Experimentally, the visible light used to generate carriers could have been delivered along the coupler length using a linear array of optical fibers, but this was not attempted.

5.2.2 Stacked Rib Waveguide Coupler

Optical coupling between stacked rib waveguides was then demonstrated, despite continued issues with facet quality. The other types of vertically-integrated couplers were fabricated with mechanically cleaved input facets, facilitating efficient coupling of light from the tapered input fiber. The output facets of the stacked channel waveguide and MMI coupler devices were plasma-etched, resulting in additional scattering losses due to their roughness. This was generally not a problem since collecting light from an output guide with an objective lens is more easily accomplished than exciting the rib mode of the input waveguide, due to the small spot size of the tapered fiber. Unfortunately, the stacked rib waveguide couplers had to be fabricated using plasma etching for both the input and output facets.

The roughness of the input facet, as seen in SEM images presented in the Figure 4.7 in Chapter 4, prevented excitation of the input rib mode. Only slab waveguiding was observed for all of the waveguides of the two successively fabricated samples. While coupling between the stacked slab regions did occur, this was not the desired result for these devices.

The input facets were then prepared for a second time using a dicing saw, as a final attempt to demonstrate vertical coupling between confined waveguide structures. This resulted in most of the input facet being severely chipped due to the air gaps surrounding the suspended lower rib waveguides. It was expected that this chipping would not have occurred had the rib waveguides been planarized, filling these gaps. Rib mode propagation was achieved for only three waveguides on one of the samples. Of these three waveguides, only one demonstrated coupling between stacked rib modes. The two others are suspected of having unacceptable thickness variations or lateral misalignments between their stacked structures.

Rib mode propagation was visually distinguished from slab waveguiding by slightly translating the input fiber from its optimal position and observing the resulting image. The tapered input fiber excited modes in both the single-mode rib and its surrounding multimode slab regions since its output spot does not perfectly overlap the field of the rib waveguide. The lobes of the slab modes were identified as they changed in both shape and location, along with their intensity, with small movements of the input fiber. The two lobes of the stacked rib waveguide supermode remained fixed, with only their intensity decreasing as the fiber was moved. Similar behavior was observed when the input polarization state was varied, as the slab lobes would shift in position and change shape, while the position of the stacked rib lobes remained fixed as the coupling ratio between them varied.

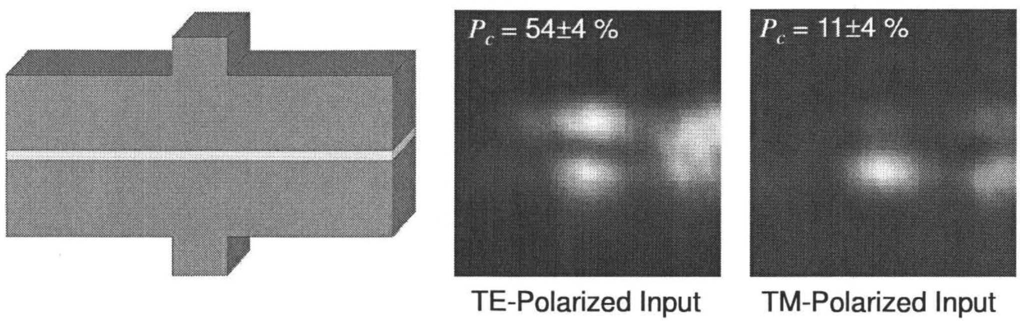

Figure 5.7: IR output images from the stacked rib waveguide coupler

Infrared output images captured from the only stacked rib waveguide device that exhibited coupling are shown in Figure 5.7, for both the TE and TM input polarizations. These images have been cropped to 75 $\times$ 75 pixels, with some slab guiding still visible to the right even with the fiber alignment optimized. Their respective coupling ratios were measured to be 54$\pm$4 % and 11$\pm$4 %, occurring at a coupler length of 572 μm. More efficient vertical coupling may have been achieved if other nearby waveguides had exhibited rib mode propagation. Simulations using the SEM-measured dimensions of this device predicted a coupling length of $L_\pi = 780\ \mu$m and 3354 μm for the TE and TM polarizations respectively, with a maximum power transfer of 81 %. The coupling ratio at the measured length of 572 μm was expected to be 67 % for the TE polarization and 6 % for a TM-polarized input.

The wavelength sensitivity of the coupling ratio is plotted in Figure 5.8 for the TE input polarization. The coupling ratio was found to vary by less than 3 % over a wavelength range of 1530-1570 nm. This is slightly more than the simulated wavelength sensitivity of ± 1 %, but within the 2 % experimental error attributed to small variations in the input polarization state.

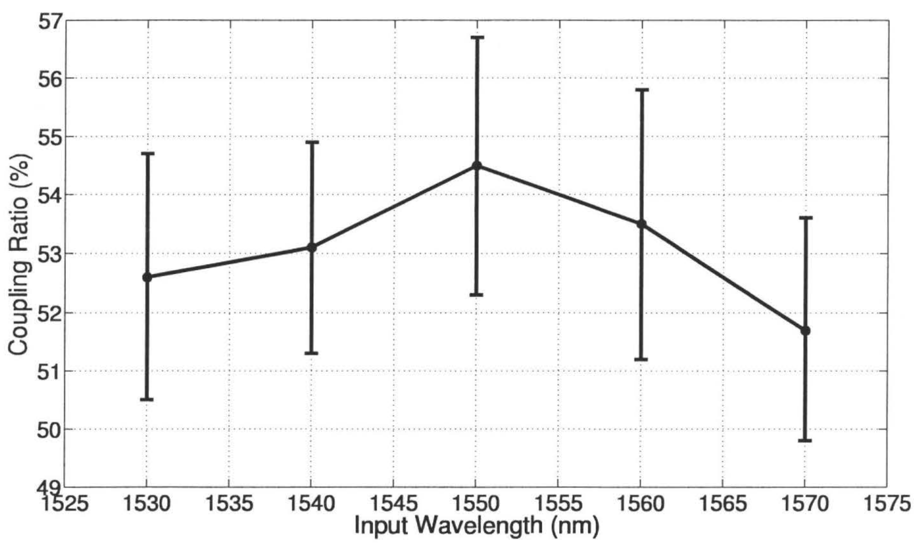

Figure 5.8: Wavelength sensitivity of the stacked rib waveguide coupler

5.2.3 Stacked Channel Waveguide Coupler

Images of the output field for the stacked channel waveguide coupler are shown in Figure 5.9 for various coupler lengths. An illustration of the device structure is also provided. The $L_c = 0$ μm length corresponds to the device with only a lower 2 $\times$ 2.5 μm channel, without a stacked channel coupler section above, and its output image can be thought of as the input field for the other coupler lengths. This device had $\sim$0.2 μm of silicon remaining on its lower channel.

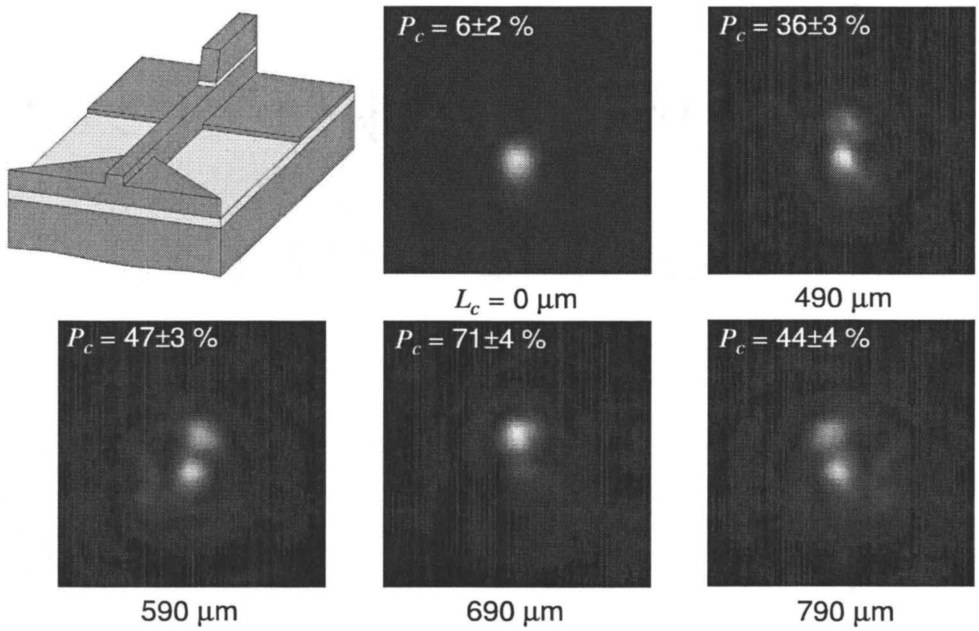

Figure 5.9: Output IR images from a stacked channel waveguide coupler

The best-case sample was fabricated with the following coupler lengths: **0**, *82*, **132**, *182*, **490**, **590**, *640*, **690**, **790**, *890*, *990*, and *1090* μm. The italicized lengths are those that did not have sufficient optical transmission to obtain accurate measurements. While the transmission losses of the other waveguides were sporadic, the corresponding measurements should still be accurate since they are based on the ratio of the power in the upper and lower waveguide regions.

The large number of excluded lengths can be attributed to the non-optimal, and inconsistent transition from the input channel into the stacked channel coupler. The width of the channels were tapered at this transition due to problems with patterning the aluminum mask used to define them. This tapering was far worse for these devices than what was observed for the MMI couplers, which will be presented in the following section. As expected, the stacked channel couplers were significantly more lossy than the MMI couplers and it was challenging to find a set of coupler lengths whose output field could be clearly imaged.

The tapered channel widths also resulted in the excitation of higher-order modes in the coupler sections. This can be observed in the output images of Figure 5.9, with the lobes of the output fields being horizontally staggered. The upper lobe was either to the left or right of the lower lobe, meaning that the excitation condition was different for each coupler length. Each higher-order mode that is excited will have its own optimal coupling length, reducing the overall coupling ratio at a given coupler length. This limited the maximum measured coupling ratio to 71 ± 4 % for a TE-polarized input. This coupling ratio was obtained for the 690 μm coupler length, while the simulated optimal length was 640 μm. The coupling ratios for both the TE and TM input polarizations are plotted for various coupler lengths in Figure 5.10. The coupling ratio for the input waveguide without a coupler section is not zero since a portion of its evanescent field will be within the upper integration box as well as some haloing from the setup optical components. The simulated results for a device with optimal coupling are also plotted for comparison.

The wavelength sensitivity of the highest coupling ratio, measured at a length of 690 μm, is plotted in Figure 5.11 for a TE-polarized input. The coupling ratio

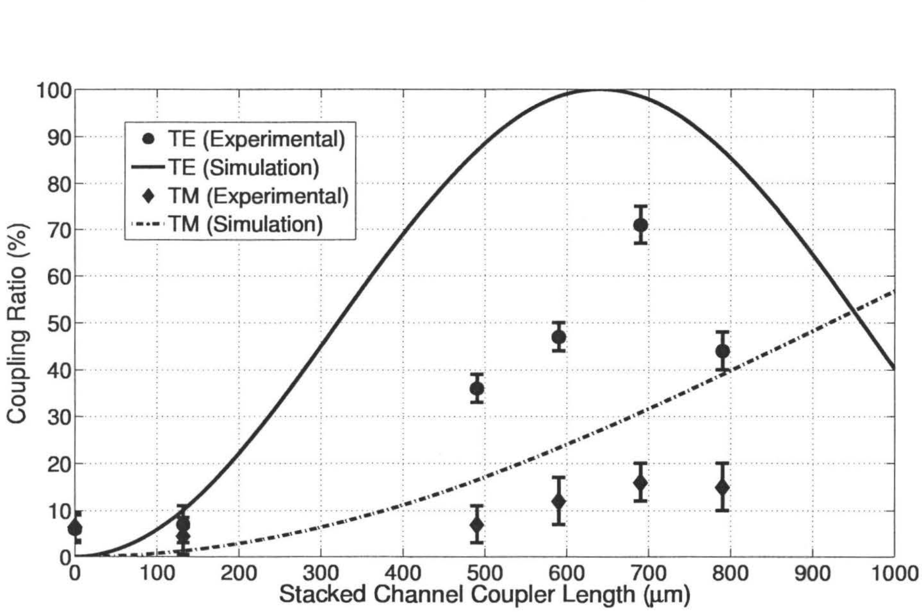

Figure 5.10: Coupling ratio measured for the stacked channel coupler

was found to vary by less than 2 % over a wavelength range of 1530-1570 nm. This is slightly more than the simulated wavelength sensitivity of ± 1 %.

The polarization sensitivity of the stacked channel coupler can be exploited to control the amount of coupling between the guides by changing the input polarization state. Either a passive or active control of the polarization state can be employed. For this project, we simply changed this state by tilting the polarization paddles, whose internal optical fiber was connected in-line between the diode laser and the input tapered fiber. The paddles were adjusted with a polarizing beam splitter cube placed in front of the camera until the light was extinguished. The polarizer had been rotated to the desired extinction angle, resulting in linearly polarized input light with an angle perpendicular to that of the extinction angle. The polarizer was then removed between each measurement and the resulting output IR image was acquired. The coupling ratio as a function of the input polarization state for the 690

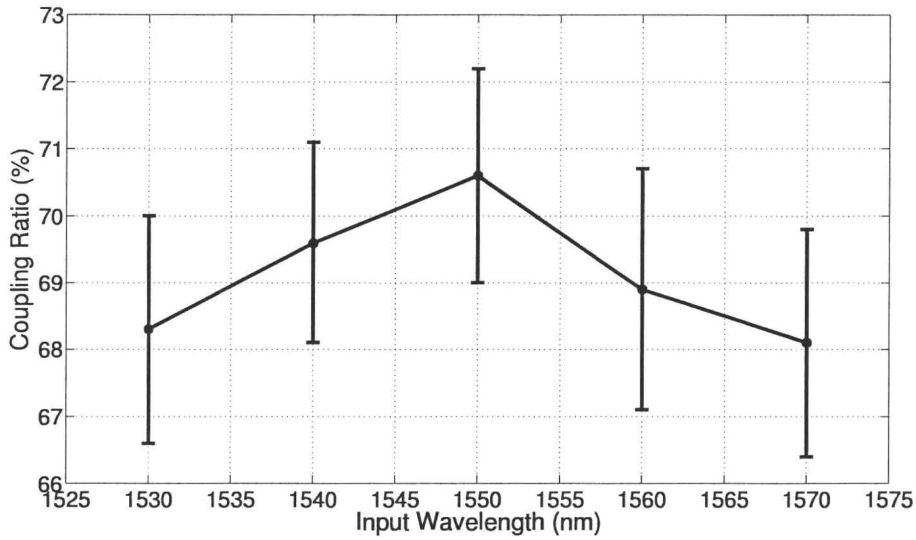

Figure 5.11: Wavelength sensitivity of the stacked channel coupler

μm long stacked channel coupler is plotted in Figure 5.12. Here, 0 ° corresponds to the TE polarization state, while 90 ° corresponds to a TM-polarized input.

The ~0.2 μm silicon slab remaining on the lower channel waveguide could be used as the substrate for a p-i-n diode; although the current set of fabricated devices, for which the facets have already been prepared, are too thin and short to be considered robust enough for the fabrication process needed for the diode. As well, the diode would have to be patterned over the coupler sections that are located at the edge of the sample. Once the photoresist is spun onto the sample, its non-uniform edge bead would cover the coupler sections, resulting in a very challenging photolithography process.

Ideally, the diode fabrication process would be completed in the center of the sample before the final facet preparation step. This could not be attempted on the devices reported here because of the problematic fabrication of the stacked channel

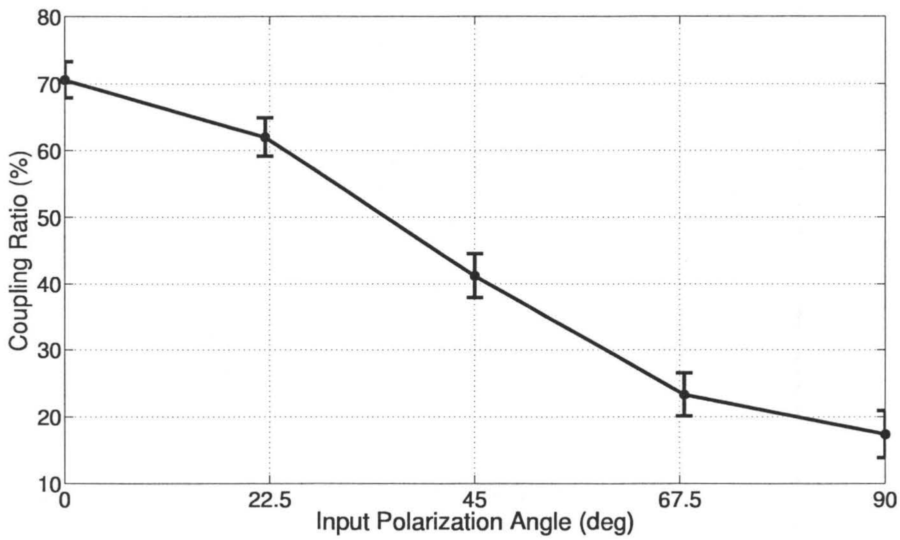

Figure 5.12: Stacked channel coupling ratio tuned by varying the input polarization state

coupler, as detailed in the previous chapter, consumed all of the double-SIMOX material that had been acquired. The fabrication process was however successful enough to demonstrate proof-of-concept optical coupling. Of the three samples that visually appeared to be successfully fabricated, only two of them had sufficient optical transmission to demonstrate passive coupling. Laser excitation was also attempted for these devices, with either the beam placed directly above the coupler or above its side slabs, but no observable change occurred.

5.2.4 Multimode Interferometer Coupler

Images of the output field for the vertically-integrated MMI coupler are shown in Figure 5.13. These images show the self-imaging evolution of the input mode at various coupler lengths for a TE-polarized input. These fields would then excite

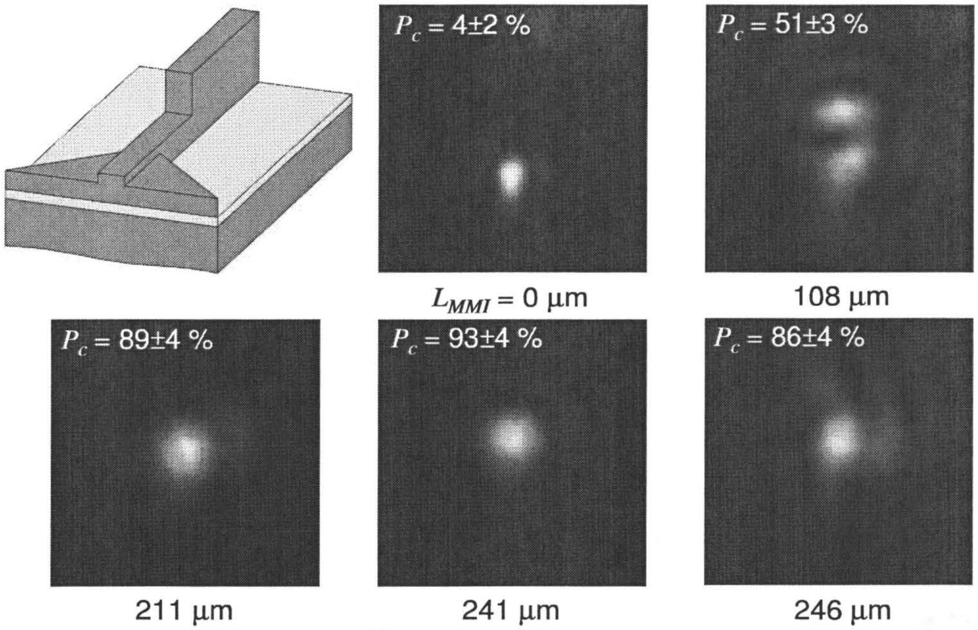

Figure 5.13: Output IR images from a vertically-integrated MMI coupler

the modes of the upper and lower output waveguides had they been placed where the coupler section terminated. The structure of the MMI coupler with its linearly tapered single-mode converter has also been illustrated.

This input image becomes nearly two-fold after propagating approximately half the designed coupling length at $L_{MMI} = 108$ μm. It evolves into what is nearly a mirrored self-image of the input at $L_{MMI} = 241$ μm, which is 5 μm longer than the designed coupling length. Images for two other intermediate coupler lengths are also shown. Vertical line slice plots that have been taken through the center of these images are provided in Figure 5.14. They have been normalized to their corresponding coupling ratio found through integration of the image intensity.

The coupling ratio for both the TE and TM input polarizations are plotted in Figure 5.15. A typical error of $\pm 4\,\%$ was found for both polarizations, but the

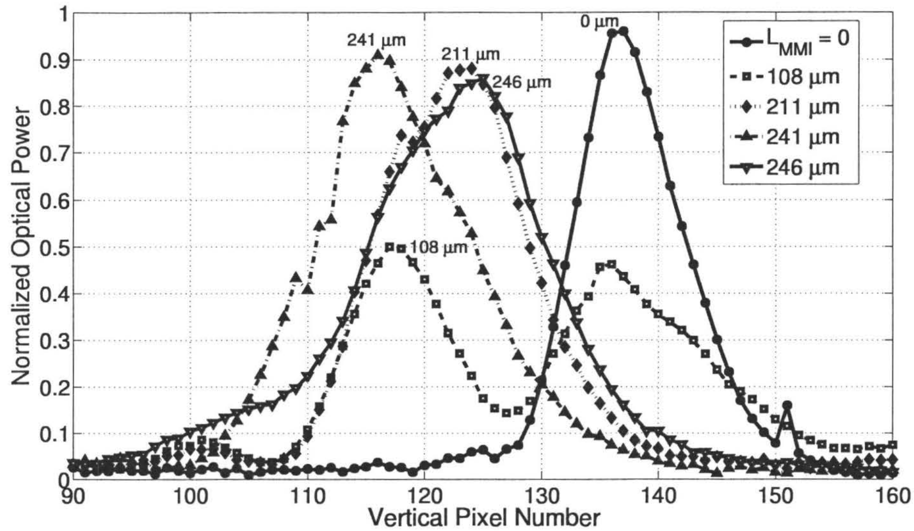

Figure 5.14: Vertical line slice plots of the MMI coupler output images

error bars have been omitted for the TM-polarized results for plot clarity. The simulated coupling ratios are also plotted for comparison. The best-case sample was fabricated with the following coupler lengths: **0**, *103*, **108**, **113**, **211**, *216*, *221*, **226**, **231**, *236*, **241**, and **246** μm. The italicized lengths are those that did not have sufficient optical transmission to obtain accurate measurements.

The maximum coupling ratio was found to be 93±4 % at the coupler length of 241 μm for the TE polarization and was 88±4 % at the same length for the TM polarization. The averaged coupling ratios were typically 5 % lower for the TM polarization for all of the measured coupler lengths. It is important to note that the percentage of the total optical power that was present in the separating 1 μm silicon slab in between was 16±4 % and 32±4 % for the TE and TM polarizations respectively. This power would be lost if the field was coupled into upper and lower output channel waveguides since the image has yet to completely evolve into

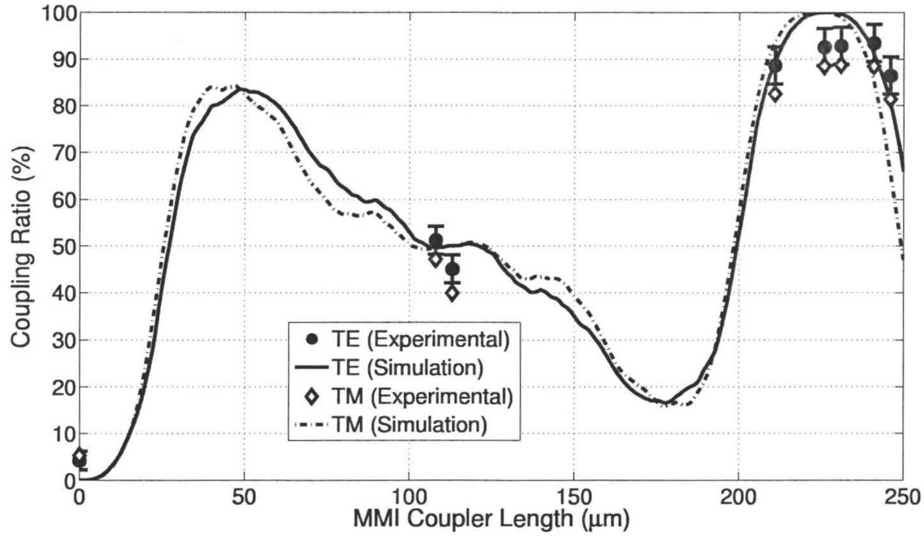

Figure 5.15: Coupling ratio measured for the MMI coupler

a mirrored self-image in the upper waveguide region. While the coupling ratio between the upper and lower waveguides is less polarization-dependent than the other devices, the input polarization state will affect the insertion loss of the coupler. This is known as polarization-dependent loss, or PDL, in the telecommunications industry.

Complete power transfer was not achieved and this may be due to the narrowing of the waveguide width at both the coupler transition and the output facet. Yet, a relatively high coupling ratio was achieved since this tapering was over a short length of the coupler section. It was also suspected that the input image launched into the coupler section was slightly multimode due to the non-optimal transition from the input channel waveguide into the MMI coupler section. This was observed as a small horizontal shift between the lobes of the two-fold image at $L_{MMI} = 108$ μm.

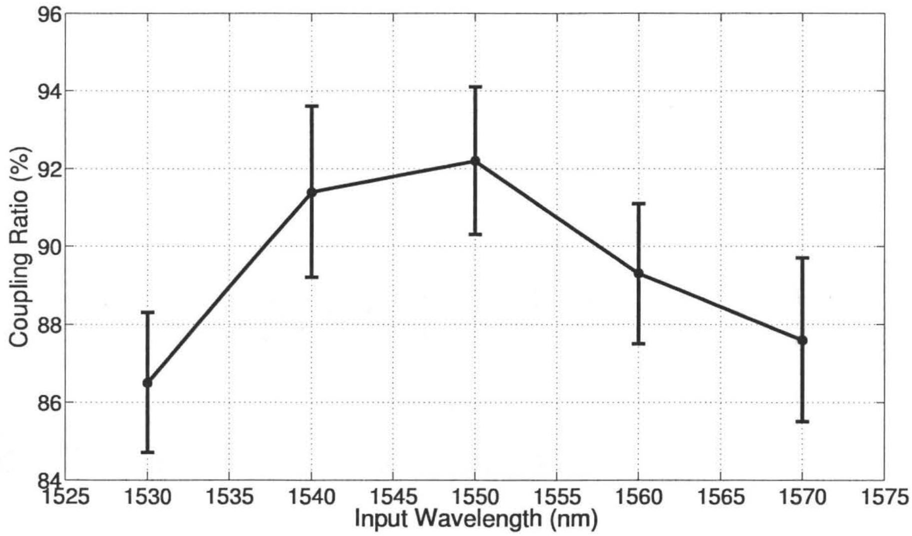

Figure 5.16: Wavelength sensitivity of the MMI coupler

The wavelength sensitivity of the coupling ratio at a length of 241 μm is plotted in Figure 5.16 for the TE input polarization. The coupling ratio was found to vary by less than 6 % over a wavelength range of 1530-1570 nm. This is slightly more than the simulated wavelength sensitivity of $\pm$4 %, but again is within the 2 % error of these measurements.

Chapter 6

Conclusions

The development of efficient vertically-integrated optical couplers is crucial to the realization of multi-layered photonic devices. This vertical integration will continue the improvements made in the performance and functional density of photonic integrated circuits once the scaling reduction of individual components has reached its limitation. Vertical coupling between various stacked waveguide structures has been successfully demonstrated in this thesis. These devices were realized using the silicon-on-insulator material system, which is the only high-index contrast material compatible with existing CMOS electronics manufacturing. This compatibility will enable the cost-effective fabrication of the proposed vertically-integrated photonic devices by leveraging the ongoing developments being made with 3-D electronics and their associated manufacturing processes such as wafer bonding. It will also facilitate the monolithic integration of both photonic and electronic components onto each layer of the proposed multi-layered device, realizing a highly-dense system-on-a-chip.

Previous demonstrations of vertically-integrated photonic devices in SOI have relied on patterned implantation techniques, while the coupling structures demonstrated in this thesis are compatible with wafer bonding processes. Wafer bonding is the preferred technique for the commercial realization of multi-layered optoelectronic devices since each device layer can be fabricated and inspected for defects separately before stacking them together, increasing overall device yield. As well, the patterned implantation process restricts the integration of electronics to the upper layer only since the layers below are not easily accessible for the required metallization processes. Wafer bonded structures enable the integration of photonics and electronics onto each layer, both manufactured using the same processes, before they are stacked. Other vertically-integrated photonic devices demonstrated using SOI have relied on polysilicon waveguide structures but their optical losses arising from material grain boundaries would be unacceptable for the telecommunications industry.

All of the fabricated proof-of-concept devices exhibited coupling efficiencies greater than 50 %, while coupling ratios above 90 % were achieved using the vertically-integrated MMI structure. The MMI-based coupler is the most likely candidate for commercial applications as it is the least polarization dependent, due to the absence of a separating oxide layer, and has the shortest coupler length of all of the structures proposed in this thesis. It is also the only structure that prevents undesirable optical cross-talk between overlapping waveguides by optically isolating each device layer everywhere except for the coupler section. The other designs would require more complex vertical tapering of their coupler sections to optically isolate each waveguide layer, as illustrated in Figure B.3 of Appendix B. While the polarization sensitivity of these devices is generally undesirable, it can be exploited to control the amount of coupling between stacked device layers in a polarization diversity switching scheme.

Numerous accomplishments have been achieved during the progress of this project. Various novel vertically-integrated optical coupler designs yet to be demonstrated in SOI have been proposed and thoroughly simulated along with their tolerance to fabrication errors. In order to experimentally demonstrate optical coupling between stacked SOI waveguides, many challenges encountered with their fabrication had to be overcome. A robust and reliable wafer bonding procedure was developed, along with many other novel fabrication processes and facet preparation techniques. The development of these processes was essential to the realization of the proof-of-concept devices that exhibited vertical coupling. For the first time, coupling between stacked SOI rib waveguides has been successively demonstrated, along with coupling between stacked SOI channel waveguides. Novel demonstration of vertical coupling within a vertically-integrated MMI structure fabricated in SOI was also achieved. The addition of these proposed devices to the research and development of vertically-integrated optoelectronics is a significant contribution to their future commercial realization.

6.1 Future Work

The development of a planarization process would be the best option for addressing the challenges encountered with facet preparation during this project. The air gaps surrounding the suspended lower rib waveguide complicated the preparation of optical-quality facets as the lack of structural support for the stacked waveguide structure resulted in severe chipping when a dicing saw was used. This chipping problem was also encountered with the tall and narrow structures of the stacked channel and MMI-based coupler designs. In the end, the facets were prepared using a plasma etching process whose inherent roughness resulted in high-loss facets, though vertical coupling was still successively demonstrated. If a planarizing material, such as spin-on-glass, had been used to encapsulate the waveguide structures there would have been adequate structural support during facet preparation with a dicing saw. This would be followed by a polishing procedure to create optical-quality facets.

While planarization would have addressed facet preparation issues, the addition of an extra material layer to the device stack may have also complicated wafer bonding since the surfaces to be bonded must be very smooth and level for adhesion to occur. In addition to determining the appropriate planarization material that would be compatible with our plasma-activated wafer bonding process, a wafer surface polishing procedure would also have to be developed to smooth out the non-uniformities of the planarizing material before bonding. The development of such a process would probably not be an issue in an industrial environment as it commonly used for planarizing the interconnect layer of electronic integrated circuits.

This project could be furthered if industrial-level wafer bonding and planarization processes become accessible at McMaster. Much of this project was dedicated to developing a reliable bonding procedure, which would have been expedited had a commercial wafer bonder been available. Ideally, this wafer bonder would have alignment capabilities and would complete all of the required process steps (wafer cleaning, plasma-activation, alignment and contact) in an enclosed, contaminant-free environment to minimize voids formed between the bonded wafer pair. Such a bonder will be installed in the near-future for Dr. Matiar Howlader's research group in the Electrical Engineering department of McMaster.

The combination of aligned wafer bonding and planarization techniques would facilitate the fabrication of the proposed devices in a near-industrial manner. The MMI couplers could be fabricated with both upper and lower input/output channel waveguides, making a 2×2 switching structure possible. This process and others requiring aligned wafer bonding and planarization techniques are detailed in Appendix B.

Future work could also involve the fabrication of current-injection diodes across the coupler structures to actively-control the transfer of optical power between stacked layers, creating a vertically-integrated optical switch. Other photonic components such as photodetectors and modulators, along with their controlling electronics, could also be integrated with the proposed optical couplers to demonstrate a highly-dense system-on-a-chip. The coupling length required for complete power transfer between stacked waveguides can also be shortened by reducing the overall dimensions of the waveguide structures. This would also decrease the tolerance of the device to misalignments between its stacked waveguides, induced during the bonding process step. Further research could involve using more advanced bonding techniques, such as interlocking structures and wafer bow compensation, to improve alignment precision.

Appendix A

Photomasks Used for Device Fabrication

The first 4" $\times$ 4" photomask used for this project consisted of a laser-patterned chrome on quartz substrate and was manufactured by Adtek Photomask in Montréal, Québec. The minimum feature size was 1$\pm$0.1 μm, written with an address size of 0.025 μm. The CAD drawing provided to Adtek is shown in Figure A.1.

The pattern for the lower waveguides is located in tile 1). These 2.2 μm wide lines are separated by 200 μm. There are 100 μm wide relief ledges, that serve as bonding surfaces, placed in between. Tile 2) consists of 9 boxes used to reveal the alignment marks of the lower waveguides. Tile 3) contains the pattern for the upper waveguides, which have been slightly staggered from the uniformly spaced lines defined in tile 1). There are 9 sets of misalignments in increments of 0.2 μm, compensating for up to $\pm$0.8 μm of misalignment between the upper and lower waveguides. Each set contains 11 waveguides, for a total of 99 waveguides per device. Tile 4) consists of 100 μm wide open channels above the waveguides, allowing the option of removing the oxide mask just above them. Tile 5) was used to define 11 different lengths for the stacked waveguide coupler section and were repeated for each misalignment set.

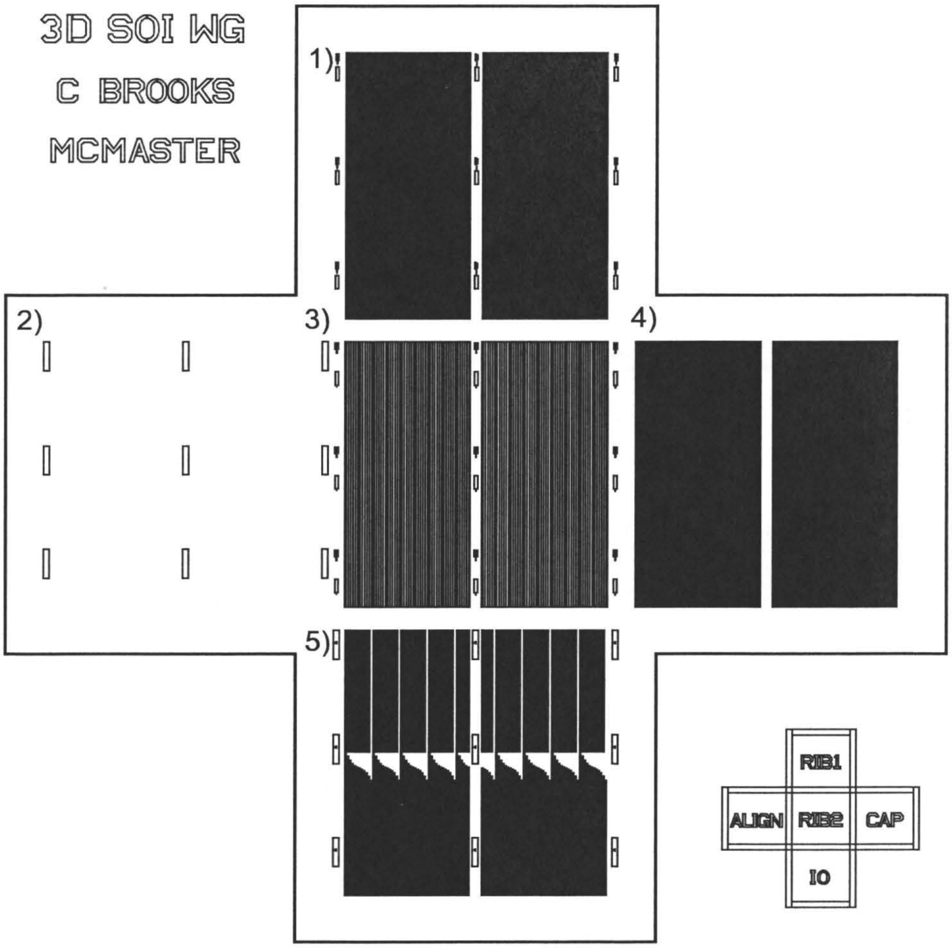

Figure A.1: CAD drawing for laser-written chrome photomask from Adtek

Tolerance simulations completed after this mask was made revealed that the coupler performance was less sensitive to misalignment than previously estimated. To ease alignment tolerances during fabrication, a second set of upper waveguides was placed on another mask with misalignments ranging from ± 4 μm in increments of 1 μm. This mask was named "VDC20" and the tile containing these waveguides was labelled "CB".

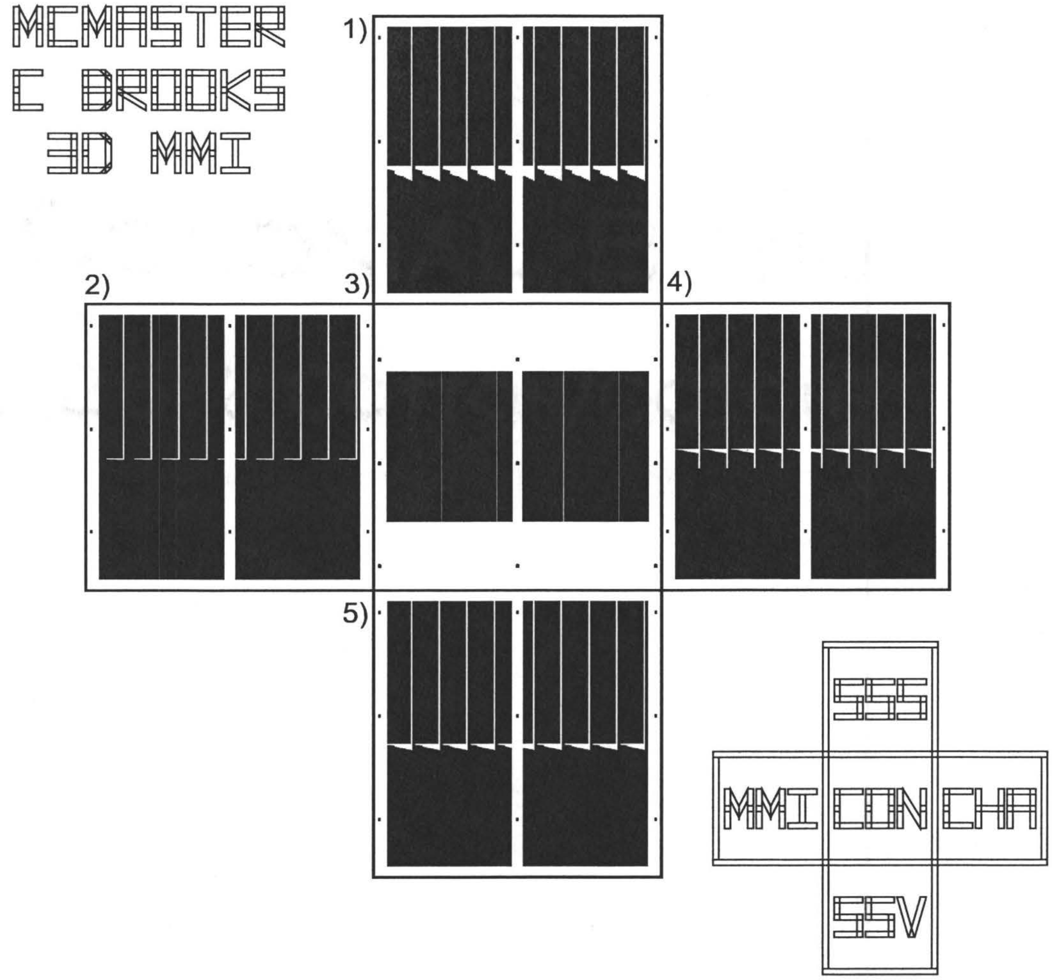

Figure A.2: CAD drawing for photomask from the University of Alberta

The third photomask used for this project consisted of a laser-patterned chrome on soda lime glass substrate and was fabricated at the University of Alberta NanoFab. The minimum feature size was 1 μm, but line widths had to be made 0.7 μm larger than desired due to the undercutting of the isotropic etchant used to pattern the chrome. This 5" $\times$ 5" photomask was later diced to a 4" $\times$ 4" square at the University of Western, which is the largest size that the mask aligner at McMaster can accommodate. The CAD drawing provided to the University of Alberta is shown in Figure A.2.

The center tile contains the pattern for the tapered mode converters. These 500 μm-long triangular sections are uniformly spaced and were placed at both the input and output of the device. These converter sections register to the waveguides located in tile 3) of the Adtek photomask. The 0.2 μm misalignment increments of these lines will compensate for up to ± 0.8 μm of misalignment when patterning the converter sections on top of the waveguides, which are only 2 μm wide themselves. The other tiles contain the coupler lengths for various coupler designs. Tile 1) would have been used for the stacked rib waveguide coupler if chemically-etched sloped sidewalls were chosen, tile 2) was used for the MMI coupler lengths, tile 4) was used for the stacked channels device and tile 5) was used for the stacked rib waveguide coupler with vertical sidewalls.

The tiles for both masks had 9 sets of alignment marks placed around the edges and along the center of the tile. The alignment marks for two overlapping tiles, as well as two resolution patterns, are shown in Figure A.3. The alignment marks consisted of either a cross within a box or a small box within a larger box. A series of long lines within open channels was placed above to compensate for rotational misalignment. Two open squares measuring 300 μm $\times$ 1150 μm were also included below the alignment marks to measure the depth for the various etching process steps.

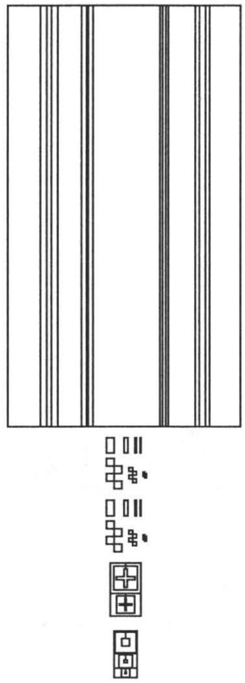

Figure A.3: Alignment marks and resolution patterns

The output facet for the majority of the devices was etched using the ICP-RIE at the University of Western. A photoresist mask was used to protect the rest of the device during this etch and was patterned using a photomask that contained a long straight edge separating a large clear area of quartz from a solid area of chrome. This mask was labelled "DOSE", but its origin was unknown.

Appendix B

Alternative Fabrication Processes

B.1 Alternative MMI Coupler Process

An alternative proof-of-concept process to fabricate vertically-integrated MMI couplers with both upper and lower waveguides is illustrated in Figure B.1. This process was not used since it requires selective epitaxial growth (SEG) of silicon along with chemical mechanical polishing (CMP) for planarization, and these processes were not available at McMaster. These illustrations omit the patterning and etching of the tapered mode converters required for single-mode device operation.

In step 1), the coupler length is defined using an aluminum mask patterned with photoresist and chemically etched. In step 2), plasma etching is used to etch through both silicon layers of the double-SOI wafer until the buried oxide is reached. The separating oxide is sufficiently thick to optically isolate each device layer. In step 3), silicon is selectively grown to fill in the etched trench, forming the MMI coupler section. In step 4), the aluminum is removed and the wafer is planarized using CMP. Steps 5) through 8) involve patterning and etching both the upper and

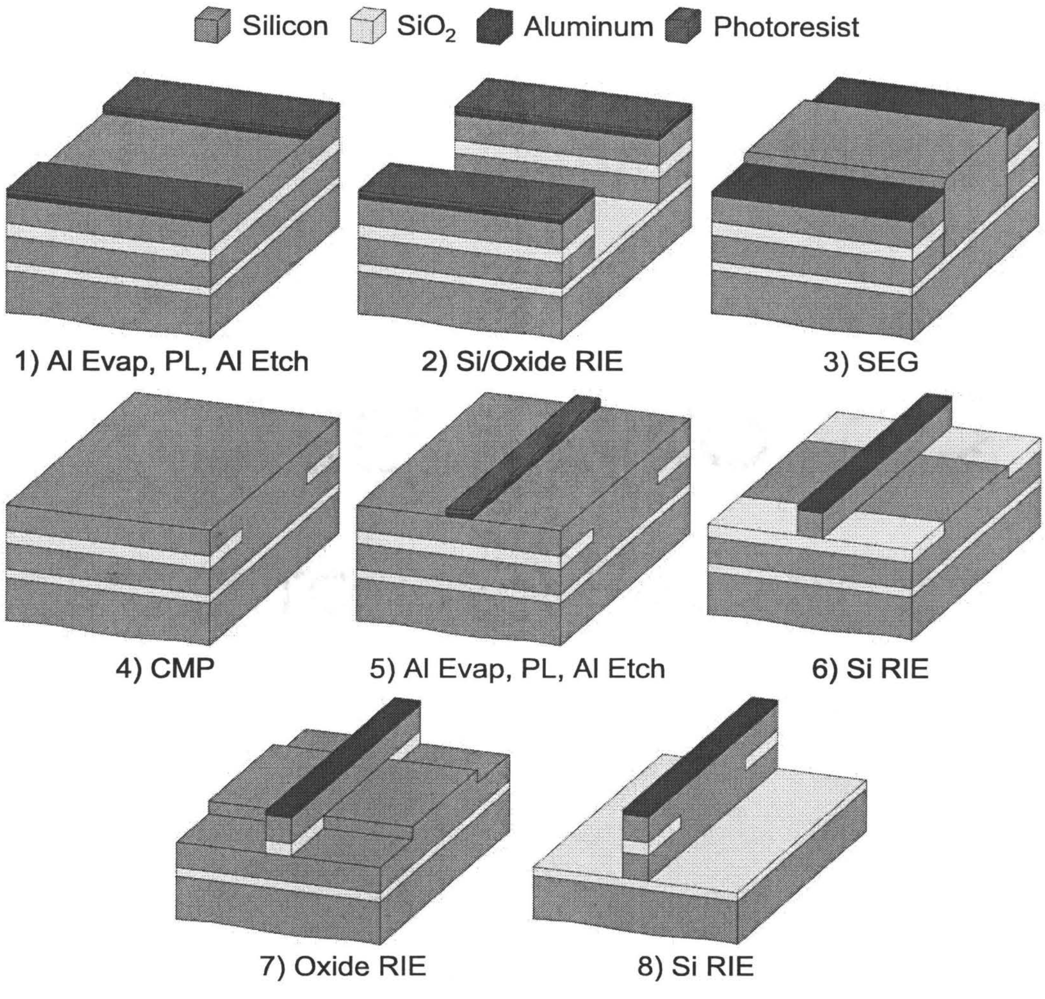

Figure B.1: Alternative proof-of-concept fabrication process for vertically-integrated MMI couplers lower input/output waveguides simultaneously, resulting in perfect alignment between them.

B.2 Commercial MMI Coupler Process

The fabrication process that could be used in industry to commercially manufacture vertically-integrated MMI couplers is illustrated in Figure B.2. This process requires precise aligned bonding, but allows each device layer to be inspected before stacking them together. A cut-away perspective is used to reveal the fabrication of the MMI coupler section that is completely encapsulated by oxide. This oxide can also planarize electronic components used for actively-controlled optical switching and will later aide facet preparation.

In step 1), a conventional single-layer SOI wafer with an etched and planarized channel waveguide is bonded facedown to an oxidized support wafer. The handle of the SOI wafer is then removed using CMP, revealing its buried oxide. In step 2), the buried oxide is patterned and etched just above the lower channel waveguide to define the length of the MMI coupler section. Some lateral misalignment can occur in this photolithography step. This trench is filled using SEG of silicon in step 3), and is then planarized in step 4). In step 5), a second single-layer SOI wafer with an etched and planarized channel is bonded facedown resulting in stacked channel waveguides separated by the MMI coupler. This process step requires precisely aligned wafer bonding. Finally, the handle of the second SOI wafer is removed using CMP in step 6). These process steps can be repeated to stack multiple optically isolated device layers that are joined by selectively placed vertically-integrated couplers.

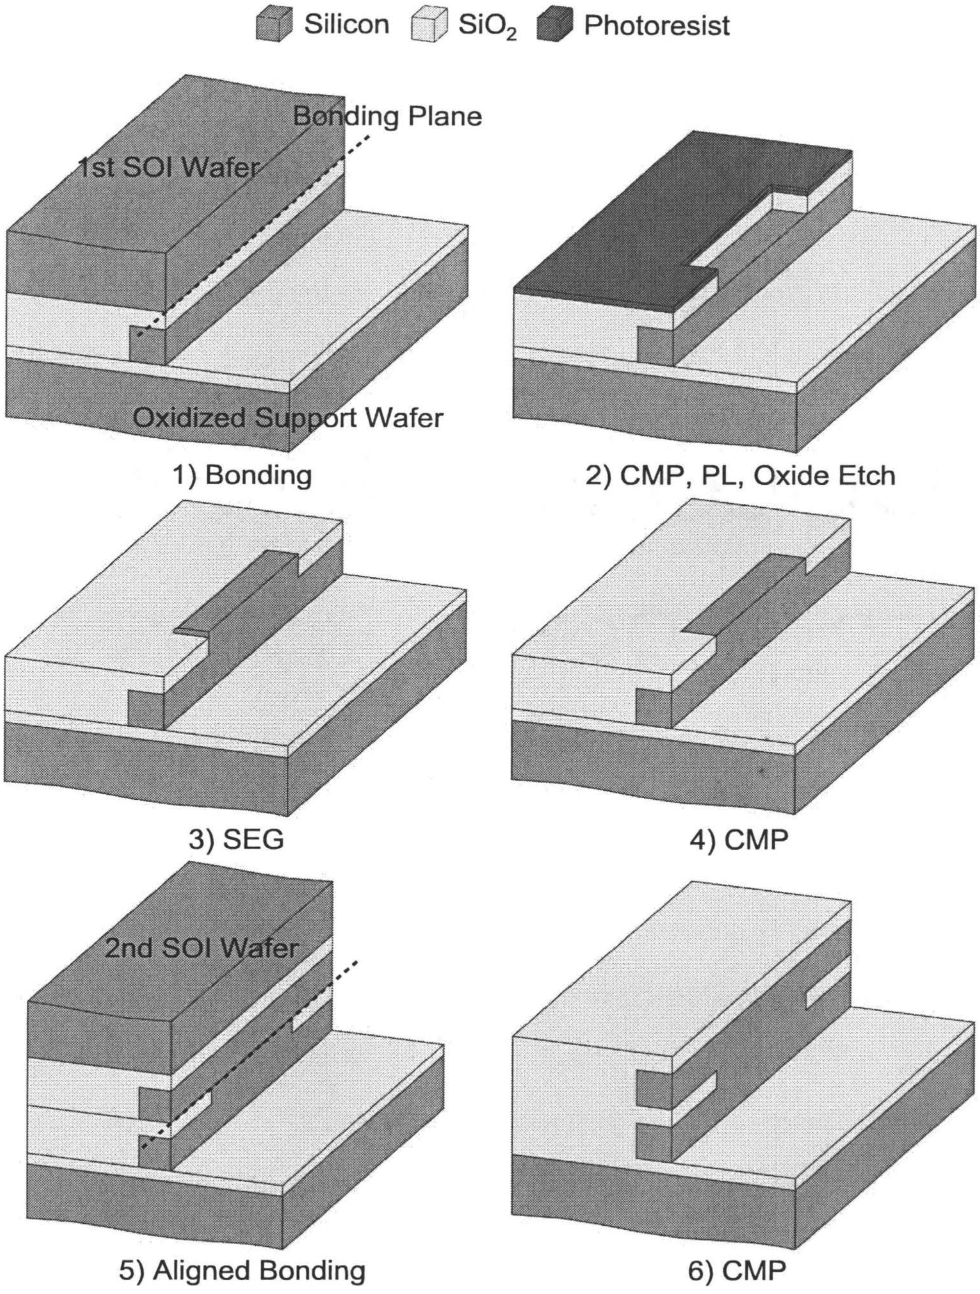

Figure B.2: Industrial fabrication process for vertically-integrated MMI couplers

B.3 Commercial Stacked Channel Coupler Process

The fabrication process that could be used in industry to commercially manufacture stacked channel waveguide couplers is illustrated in Figure B.3. These waveguide layers can be optically isolated by using gray-scale photolithography to vertically taper the height of the waveguides, as demonstrated in [37]. The encapsulating oxide is excluded from these illustrations to better reveal the fabrication of the coupler section.

In step 1), the vertical taper of the lower channel waveguide is patterned using gray-scale photolithography, defining the length of the coupler as well. The subsequent plasma etching of silicon with this tapered photoresist in place will translate into the waveguide height being tapered as well. This tapered waveguide will be optically isolated from any waveguides placed above, except along the coupler length. In step 2), the photoresist is removed and the waveguide is planarized using a deposited oxide and CMP, leaving a thin layer of oxide above the coupler section. In step 3), a second SOI wafer that has followed the same process is bonded facedown with precise alignment. The handle of the second SOI wafer is removed using CMP in step 4). Once again, this process can be repeated to stack multiple device layers.

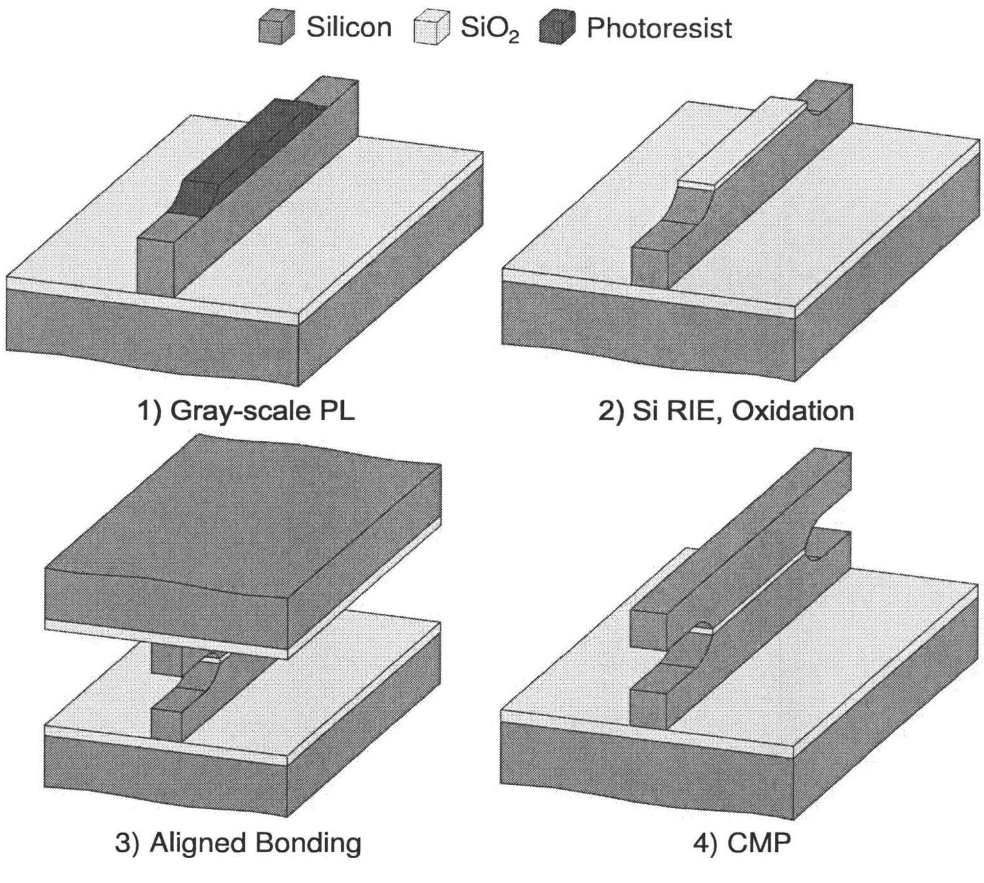

Figure B.3: Industrial fabrication process for stacked channel waveguide couplers

B.4 Diode Fabrication Process

While current-injection diodes were not fabricated onto the passive coupling devices made during this project, their process is illustrated in Figure B.4 to demonstrate the compatibility of these couplers with electrical integration. The process is shown for a stacked rib waveguide coupler as an example, but would be similar for the stacked channel coupler as well. In step 1), thick photoresist is patterned to open a window above the p-type dopant region and then an ion such as boron is implanted through this opening. The photoresist is removed and the process in repeated for the adjacent n-type dopant region in step 2). A common n-type dopant would be phosphorus.

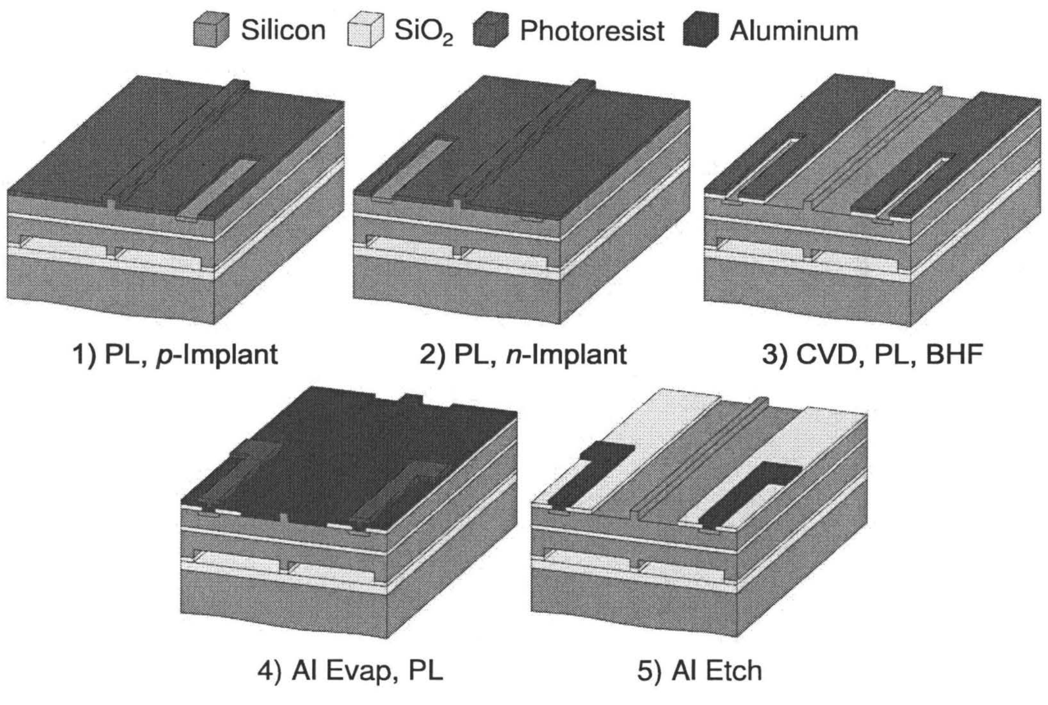

Figure B.4: Fabrication process for diode-controlled stacked rib waveguide couplers

In step 3), an insulating CVD oxide is deposited and removed above the rib waveguide using photoresist and BHF. This is to ensure that the two rib waveguides remain matched. Contact windows above the doped regions are opened in the remaining oxide and then the photoresist is removed. In step 4), an aluminum layer is evaporated over the entire sample for the contacts that are then patterned using photolithography. Finally, the aluminum is chemically etched in step 5) to fabricate the contacts. A lift-off process could also be used to create the contacts.

References

[1] L. Wei, R. Zhang, K. Roy, Z. Chen, and D. B. Janes, "Vertically integrated SOI circuits for low-power and high-performance applications," *IEEE Transactions on Very Large Scale Integration Systems*, vol. 10, no. 3, pp. 351–362, 2002.

[2] K. W. Guarini, A. W. Topol, M. Leong, R. Yu, and L. Shi, "Electrical integrity of state-of-the-art 0.13 μm SOI CMOS devices and circuits transferred for three-dimensional (3D) integrated circuit (IC) fabrication," in *IEEE International Electron Devices Meeting*, pp. 943–945, 2002.

[3] S. Pae, T. Su, J. P. Denton, and G. W. Neudeck, "Multiple layers of silicon-on-insulator islands fabrication by selective epitaxial growth," *IEEE Electron. Device Lett.*, vol. 20, no. 5, pp. 194–199, 1999.

[4] S. W. Bond, O. Vendier, M. Lee, S. Jung, M. Vrazel, A. Lopez-Lagunas, S. Chai, G. Dagnall, M. Brooke, N. M. Jokerst, D. S. Wills, and A. Brown, "A three-layer 3-D silicon system using through-Si vertical optical interconnections and Si CMOS hybrid building blocks," *IEEE J. Sel. Top. Quant. Electron.*, vol. 5, no. 2, pp. 276–286, 1999.

[5] A. Liu, R. Jones, L. Liao, D. Samara-Rubio, D. Rubin, O. Cohen, R. Nicolaescu, and M. Paniccia, "A high-speed silicon optical modulator based on a metal-oxide-semiconductor capacitor," *Nature*, vol. 427, pp. 615–618, 2004.

[6] A. P. Knights, J. D. B. Bradley, S. H. Gou, and P. E. Jessop, "Silicon-on-insulator waveguide photodetector with self-ion-implantation-engineered-enhanced infrared response," *J. Vac. Sci. Technol. A*, vol. 24, no. 3, pp. 783–786, 2006.

[7] H. Rong, R. Jones, A. Liu, O. Cohen, D. Hak, A. Fang, and M. Paniccia, "A continuous-wave Raman silicon laser," *Nature*, vol. 433, pp. 725–728, 2005.

[8] R. A. Soref, "Silicon-based optoelectronics," *Proc. IEEE*, vol. 81, no. 12, pp. 1687–1706, 1993.

[9] B. Jalali, S. Yegnanarayanan, T. Yoon, T. Yoshimoto, I. Rendina, and F. Coppinger, "Advances in silicon-on-insulator optoelectronics," *IEEE J. Sel. Top. Quant. Electron.*, vol. 4, no. 6, pp. 938–947, 1998.

[10] G. Celler and S. Cristoloveanu, "Frontiers of silicon-on-insulator," *J. Appl. Phys.*, vol. 93, no. 9, pp. 4955–4978, 2003.

[11] F. Grillot, L. Vivien, S. Laval, and E. Cassan, "Propagation loss in single-mode ultrasmall square silicon-on-insulator optical waveguides," *J. Lightwave. Tech.*, vol. 24, no. 2, pp. 891–896, 2006.

[12] J. E. Bowers, H. Park, A. W. Fang, R. Jones, O. Cohen, and M. J. Paniccia, "A technology for integrating active photonic devices on SOI wafers," in *Indium Phosphide and Related Materials Conference Proceedings, 2006 International Conference on*, pp. 218–221, 2006.

[13] G. T. Reed and A. P. Knights, *Silicon Photonics*. John Wiley & Sons Ltd., 2004.

[14] "SOI wafer: Solutions for high speed/low power consumption," tech. rep., Shin-Etsu Handotai Co., Ltd., Tokyo, Japan, 2001.

[15] J. D. Joannopoulos, S. G. Johnson, J. N. Winn, and R. D. Meade, *Photonic Crystals: Molding the Flow of Light*. Princeton University Press, 2008.

[16] S. Noda, M. Imada, M. Okano, S. Ogawa, M. Mochizuki, and A. Chutinan, "Semiconductor three-dimensional and two-dimensional photonic crystals and devices," *IEEE J. Quantum Electron.*, vol. 38, no. 7, pp. 726–735, 2002.

[17] B. Liu, A. Shakouri, P. Abraham, and J. E. Bowers, "Vertical coupler with separated inputs and outputs fabricated using double-sided process," *Electron. Lett.*, vol. 35, no. 18, pp. 1552–1554, 1999.

[18] D. V. Tishinin, P. D. Dapkus, A. E. Bond, I. Kim, C. K. Lin, and J. O'Brien, "Vertical resonant couplers with precise coupling efficiency control fabricated by wafer bonding," *IEEE Photon. Tech. Lett.*, vol. 11, no. 8, pp. 1003–1005, 1999.

[19] M. Raburn, B. Liu, P. Abraham, and J. E. Bowers, "Double-bonded InP-InGaAsP vertical coupler 1:8 beam splitter," *IEEE Photon. Tech. Lett.*, vol. 12, no. 12, pp. 1639–1641, 2000.

[20] M. Raburn, B. Liu, K. Rauscher, Y. Okuno, N. Dagli, and J. E. Bowers, "3-D photonic circuit technology," *IEEE J. Sel. Top. Quant. Electron.*, vol. 8, no. 4, pp. 935–942, 2002.

[21] S. M. Garner, S.-S. Lee, V. Chuyanov, A. Chen, A. Yacoubian, W. H. Steier, and L. R. Dalton, "Three-dimensional integrated optics using polymers," *IEEE J. Quant. Electron.*, vol. 35, no. 8, pp. 1146–1155, 1999.

[22] S. M. Garner, V. Chuyanov, S.-S. Lee, A. Chen, W. H. Steier, and L. R. Dalton, "Vertically integrated waveguide polarization splitters using polymers," *IEEE Photon. Tech. Lett.*, vol. 11, no. 7, pp. 842–844, 1999.

[23] H.-S. Lee, S. An, and E.-H. Lee, "A new method of optical interconnection for multiply stacked planar optical circuit boards using 45 ° reflection coupling," *Microelectron. Eng.*, vol. 87, no. 5-8, pp. 1336–1339, 2010.

[24] L. B. Soldano and E. C. M. Pennings, "Optical multi-mode interference devices based on self-imaging: Principles and applications," *J. Lightwave. Tech.*, vol. 13, no. 4, pp. 615–627, 1995.

[25] J.-M. Lee, J. T. Ahn, D. H. Cho, J. J. Ju, M.-H. Lee, and K. H. Kim, "Vertical coupling of polymeric double-layered waveguides using a stepped MMI coupler," *Electronics and Telecommunications Research Institute (ETRI) Journal*, vol. 25, no. 2, pp. 81–88, 2003.

[26] S. Ikuta, T. Sekimoto, W. Pan, S. T. Chu, and Y. Kokubun, "High-coupling efficiency vertical ARROW coupler with large tolerance and short coupling length for three-dimensional optical interconnects," *IEEE Photon. Tech. Lett.*, vol. 11, no. 8, pp. 1006–1008, 1999.

[27] S. Ura, R. Nishida, T. Suhara, , and H. Nishihara, "Wavelength-selective coupling among three vertically integrated optical waveguides by grating couplers," *IEEE Photon. Tech. Lett.*, vol. 13, no. 2, pp. 133–135, 2001.

[28] F. Chollet, M. de Labachelerie, and H. Fujita, "Compact evanescent optical switch and attenuator with electromechanical actuation," *IEEE J. Sel. Top. Quant. Electron.*, vol. 5, no. 1, pp. 52–59, 1999.

[29] I. Shubin and P. Likamwa, "A guided-wave optical switch controlled by a micro-electro-mechanical cantilever," in *Proc. 13th Annual LEOS Meeting*, 2000.

[30] R. A. Soref, E. Cortesi, F. Namavar, and L. Friedman, "Vertically integrated silicon-on-insulator waveguides," *IEEE Photon. Tech. Lett.*, vol. 3, no. 1, pp. 22–24, 1991.

[31] P. Koonath, K. Kishima, T. Indukuri, and B. Jalali, "Sculpting of three-dimensional nano-optical structures in silicon," *Appl. Phys. Lett.*, vol. 83, no. 24, pp. 4909–4911, 2003.

[32] P. Koonath, T. Indukuri, and B. Jalali, "Vertically-coupled micro-resonators realized using three-dimensional sculpting in silicon," *Appl. Phys. Lett.*, vol. 85, no. 6, pp. 1018–1020, 2004.

[33] T. Indukuri, P. Koonath, and B. Jalali, "Subterranean silicon photonics: Demonstration of buried waveguide-coupled microresonators," *Appl. Phys. Lett.*, vol. 87, no. 081114, 2005.

[34] P. Koonath, T. Indukuri, and B. Jalali, "Monolithic 3-D silicon photonics," *J. Lightwave. Tech.*, vol. 24, no. 4, pp. 1796–1804, 2006.

[35] T. Indukuri, P. Koonath, and B. Jalali, "Three-dimensional integration of metal-oxide-semiconductor transistor with subterranean photonics in silicon," *Appl. Phys. Lett.*, vol. 88, no. 121108, 2006.

[36] P. Koonath and B. Jalali, "Multilayer 3-D photonics in silicon," *Opt. Express*, vol. 15, no. 20, pp. 12687–12691, 2007.

[37] A. Sure, T. Dillon, J. Murakowski, C. Lin, D. Pustai, and D. W. Prather, "Fabrication and characterization of three-dimensional silicon tapers," *Opt. Express*, vol. 11, no. 26, pp. 3555–3561, 2003.

[38] P. Dong and A. G. Kirk, "Compact double-grating coupler between vertically stacked silicon-on-insulator waveguides," *Appl. Optic.*, vol. 44, no. 35, pp. 7540–7547, 2005.

[39] K. Preston, B. Schmidt, and M. Lipson, "Polysilicon photonic resonators for large-scale 3D integration of optical networks," *Opt. Express*, vol. 15, no. 25, pp. 17283–17290, 2007.

[40] J. K. Doylend, A. P. Knights, C. Brooks, and P. E. Jessop, "CMOS compatible vertical directional coupler for 3D optical circuits," in *Proc. Photonics North*, 2005.

[41] C. J. Brooks, J. K. Doylend, A. P. Knights, and P. E. Jessop, "Vertically-stacked SOI waveguides for 3-D photonic circuits," in *Group IV Photonics, 2006. 3rd IEEE International Conference on*, pp. 72–74, 2006.

[42] C. Brooks, A. P. Knights, , and P. E. Jessop, "Vertically-integrated multimode interferometers for 3-D photonic circuits in SOI," in *Proc. Photonics West*, 2008.

[43] C. Brooks, A. P. Knights, , and P. E. Jessop, "Vertically-integrated multimode interferometer coupler for 3D photonic circuits in SOI," *Submitted to Opt. Express*.

[44] L. Liao, D. R. Lim, A. M. Agarwal, X. Duan, K. K. Lee, and L. C. Kimerling, "Optical transmission losses in polycrystalline silicon strip waveguides: Effects of waveguide dimensions, thermal treatment, hydrogen passivation, and wavelength," *J. Electron. Mat.*, vol. 29, no. 12, pp. 1380–1386, 2000.

[45] J. F. Gibbons and K. F. Lee, "One-gate-wide CMOS inverter on laser-recrystallized polysilicon," *IEEE Electron Device Lett.*, vol. EDL-1, p. 117118, 1980.

[46] E. J. Teo, A. Bettiol, M. B. H. Breese, P. Yang, G. Z. Mashanovich, W. R. H. G. T. Reed, and D. J. Blackwood, "Three-dimensional control of optical waveguide fabrication in silicon," *Opt. Express*, vol. 16, no. 2, pp. 573–578, 2008.

[47] M. A. Schmidt, "Wafer-to-wafer bonding for microstructure formation," *Proc. IEEE*, vol. 86, no. 8, pp. 1575–1585, 1998.

[48] R. C. Cole, R. E. Robertson, R. M. Garrett, and D. C. Mayer, "Formation of a bonded SOI interface using a spin-on-dielectric," in *IEEE International SOI Conference*, pp. 44–45, 1996.

[49] H. C. Lin, K. L. Chang, G. W. Pickrell, K. C. Hsieh, and K. Y. Cheng, "Low temperature wafer bonding by spin on glass," *J. Vac. Sci. Tech. B*, vol. 20, no. 2, pp. 752–754, 2002.

[50] H. Quenzer, A. V. Schulz, T. Kinkopf, and T. Helm, "Anodic-bonding on glass layers prepared by a spin-on glass process: preparation process and experimental results," in *The 11th International Conference on Solid-State Sensors and Actuators*, pp. 230–233, 2001.

[51] C. Maleville, T. Barge, B. Ghyselen, and A. Auberton, "Multiple SOI layers by multiple Smart-Cut transfers," in *IEEE International SOI Conference*, pp. 134–135, 2000.

[52] M. Wiegand, M. Reiche, and U. Gosele, "Time-dependent surface properties and wafer bonding of O_2-plasma-treated silicon (100) surfaces," *J. Electrochem. Soc.*, vol. 147, no. 7, pp. 2734–2740, 2000.

[53] A. Sanz-Velasco, P. Amirfeiz, S. Bengtsson, and C. Colinge, "Room temperature wafer bonding using oxygen plasma treatment in reactive ion etchers

with and without inductively coupled plasma," *J. Electrochem. Soc.*, vol. 150, no. 2, pp. G155–G162, 2003.

[54] T. Suni, K. Henttinen, I. Suni, and J. Makinen, "Effects of plasma activation on hydrophilic bonding of Si and SiO_2," *J. Electrochem. Soc.*, vol. 149, no. 6, pp. G348–G351, 2002.

[55] Y. Zhang, S.-W. Myung, H.-S. Choi, I.-H. Kim, and J.-H. Choi, "Optimum conditions for the surface modification of polyurethane by oxygen plasma treatment," *J. Ind. Eng. Chem.*, vol. 8, no. 3, pp. 236–240, 2002.

[56] A. W. Topol, D. C. L. Tulipe, L. Shi, S. M. Alam, D. J. Frank, S. E. Steen, and J. Vichiconti, "Enabling SOI-based assembly technology for three-dimensional (3D) integrated circuits (ICs)," in *Electron Devices Meeting, 2005. IEDM Technical Digest. IEEE International*, pp. 352–355, 2005.

[57] L. Jiang, G. Pandraud, P. J. French, S. M. Spearing, and M. Kraft, "A novel method for nanoprecision alignment in wafer bonding applications," *J. Micromech. Microeng.*, vol. 17, pp. S61–S67, 2007.

[58] E. Hecht, *Optics*. Pearson Education, Inc., 4 ed., 2002.

[59] H. Nishihara, M. Haruna, and T. Suhara, *Optical Integrated Circuits*. McGraw-Hill Book Company, 1989.

[60] A. Yariv, *Introduction to Optical Electronics*. Holt, Rinehart and Winston, 1976.

[61] *BeamPROP manual*. RSoft Corporation, 2005.

[62] M. Feit and J. Fleck, "Light propagation in graded-index optical fibers," *Appl. Opt.*, vol. 17, no. 24, pp. 3990–3998, 1978.

[63] D. Yevick, "A guide to electric field propagation techniques for guided wave optics," *Opt. and Quant. Elec.*, vol. 26, no. 3, pp. S185–S197, 1994.

[64] G. Hadley, "Wide-angle beam propagation using Pade approximant operators," *Opt. Lett.*, vol. 17, no. 20, pp. 1426–1428, 1992.

[65] D. Yevick and W. Bardyszewski, "Correspondence of variational finite-difference (relaxation) and imaginary-distance propagation methods for modal analysis," *Opt. Lett.*, vol. 17, no. 5, pp. 329–330, 1992.

[66] J. Chen and S. Jungling, "Computation of higher-order waveguide modes by the imaginary-distance beam propagation method," *Opt. and Quant. Elec.*, vol. 26, no. 3, pp. S199–S205, 1994.

[67] R. A. Soref, J. Schmidtchen, and K. Petermann, "Large single-mode rib waveguides in GeSi-Si and Si-on-SiO_2," *IEEE J. Quantum Electron.*, vol. 27, no. 8, pp. 1971–1974, 1991.

[68] A. G. Rickman, G. Reed, and F. Namavar, "Silicon-on-insulator optical rib waveguide loss and mode characteristics," *J. Lightwave Technol.*, vol. 12, no. 10, pp. 1771–1776, 1994.

[69] S. P. Pogossian, L. Vescan, and A. Vonsovici, "The single-mode condition for semiconductor rib waveguides with large cross section," *J. Lightwave Technol.*, vol. 16, no. 10, pp. 1851–1853, 1998.

[70] R. S. Jacobsen, K. N. Andersen, P. I. Borel, J. Fage-Pedersen, L. H. Frandsen, O. Hansen, and M. Kristensen, "Strained silicon as a new electro-optic material," *Nature*, vol. 441, pp. 199–202, 2006.

[71] R. A. Soref and B. R. Bennett, "Electrooptical effects in silicon," *IEEE J. Quantum Electron.*, vol. QE-23, no. 1, pp. 123–129, 1987.

[72] R. A. Soref and B. R. Bennett, "Kramers-Kronig analysis of electro-optical switching in silicon," *Proc. SPIE*, no. 704, pp. 32–37, 1986.

[73] C. A. Barrios, V. R. Almeida, R. Panepucci, and M. Lipson, "Electrooptic modulation of silicon-on-insulator submicrometer-size waveguide devices," *J. Lightwave Technol.*, vol. 21, no. 10, pp. 2332–2340, 2003.

[74] C. K. Tang and G. T. Reed, "Highly efficient optical phase modulator in SOI waveguides," *Electron. Lett.*, vol. 31, p. 451452, 1995.

[75] R. L. Espinola, M.-C. Tsai, J. T. Yardley, and R. M. Osgood, "Fast and low-power thermooptic switch on thin silicon-on-insulator," *IEEE Photon. Tech. Lett.*, vol. 15, no. 10, pp. 1366–1368, 2003.

[76] M. W. Geis, S. J. Spector, R. C. Williamson, and T. M. Lyszczarz, "Submicrosecond submilliwatt silicon-on-insulator thermooptic switch," *IEEE Photon. Tech. Lett.*, vol. 16, no. 11, pp. 2514–2516, 2004.

[77] F. Gan, F. J. Grawert, J.-M. Schley, S. Akiyama, J. Michel, K. Wada, L. C. Kimerling, and F. X. Kartner, "Design of all-optical switches based on carrier injection in Si/SiO_2 split-ridge waveguides (SRWs)," *J. Lightwave Technol.*, vol. 24, no. 9, pp. 3454–3463, 2006.

[78] M. v. Allmen and A. Blatter, *Laser-Beam Interactions with Materials*. Springer Science+Business Media, 1995.

[79] J. Batista, A. Mandelis, and D. Shaughnessy, "Temperature dependence of carrier mobility in Si wafers measured by infrared photocarrier radiometry," *Appl. Phys. Lett.*, vol. 82, no. 23, pp. 4077–4079, 2003.

[80] D. Dimitropoulos, R. Jhaveri, R. Claps, J. C. S. Wo, and B. Jalali, "Lifetime of photogenerated carriers in silicon-on-insulator rib waveguides," *Appl. Phys. Lett.*, vol. 86, no. 071115, 2005.

[81] D. Dai, J.-J. He, and S. He, "Compact silicon-on-insulator-based multimode interference coupler with bilevel taper structure," *Appl. Optic.*, vol. 44, no. 24, pp. 5036–5041, 2005.

[82] J. Liu, J. Yu, S. Chen, and J. Xia, "Fabrication and analysis of 2×2 thermo-optic SOI waveguide switch with low power consumption and fast response by anisotropy chemical etching," *Opt. Commun.*, vol. 245, pp. 137–144, 2005.

[83] E. D. Palik, *Handbook of Optical Constants of Solids*. Academic Press, 1985.

[84] M. Howlader, S. Suehara, and T. Suga, "Room temperature wafer level glass/glass bonding," *Sens. Acutators, A*, vol. 127, no. 1, pp. 31–36, 2006.

[85] M. Howlader, F. Zhang, and M. Kibria, "Void nucleation at a sequentially plasma-activated silicon/silicon bonded interface," *J. Micromech. Microeng.*, vol. 20, no. 6, pp. 1–10, 2010.