Fabrication of Optical-Mode Converters for Efficient Fiber-to-Silicon-Waveguide Couplers

by

Raúl E. Barreto

B.S., Electrical Science and Engineering
Massachusetts Institute of Technology, 2005

Submitted to the Department of Electrical Engineering and Computer Science
In partial Fulfillment of the Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

at the

Massachusetts Institute of Technology

December 2006

Signature of Author.....
Department of Electrical Engineering and Computer Science
December 21, 2006

Certified by.....
.....
Henry I. Smith
Professor of Electrical Engineering
Thesis Supervisor

Accepted by.....
.....
Arthur C. Smith
Chairman, Department Committee on Graduate Theses

Fabrication of Optical-Mode Converters for Efficient Fiber-to-Silicon-Waveguide Couplers

by

Raúl E. Barreto

Submitted to the Department of Electrical Engineering and Computer Science
on December 21, 2006 in partial Fulfillment of the
Requirements for the Degree of
Master of Engineering in Electrical Engineering and Computer Science

ABSTRACT

Optical-mode converters are needed to efficiently couple light from an optical fiber to a photonic circuit by matching and transforming the propagating modes. This work is based on a horizontally-tapered coupler, in which light from an optical fiber is coupled into a large polymer waveguide and then gradually transferred to a smaller silicon waveguide whose width increases with distance along the guide. Several devices were designed and fabricated to measure the efficiency of the coupler. E-beam exposure doses and writing strategies were optimized to create the tapered silicon waveguides. A fabrication process was developed to form the polymer waveguides without etching the underlying silicon, and a set of marks was created to achieve sub-micron alignment between the two waveguides. Fabrication results showed that the coupler successfully transfers light between the two waveguides and that there is low loss in the polymer. A more accurate characterization of the coupler's efficiency was delayed due to fabrication problems not related to the developed process.

Thesis Supervisor: Henry I. Smith
Title: Professor of Electrical Engineering

Acknowledgments

There are many people who have greatly helped me and, directly or indirectly, contributed to the completion of this Thesis. First, I would like to thank my advisor, Prof. Henry Smith, who gave me the opportunity to join his group and become involved with this work. When the project ran into financial difficulties he always found the resources necessary so that I could continue my work, and provided guidance and insight when technical difficulties arose.

I want to recognize the invaluable contributions of everyone who worked in the Pirelli project. In particular, I am very thankful to Tymon Barwicz whose experience and dedication significantly eased my learning curve. I thank Anatoly Khilo and Marcus Dehlam for the design and characterization of the coupler. Charles Holzwarth and Milos Popovic were also always willing to help and very knowledgeable; thank you for getting involved when we most needed it.

Among the many people who trained me in the equipment I needed and helped me troubleshoot it multiple times, I especially want to thank Jim Daley for his help in the NSL, Mark Mondol for having the patience to answer all my Raith questions, and Dave Terry who was always eager to help in TRL.

On a personal note, I would like to thank my family and friends for their support and motivation. My parents have always encouraged and advised me; to them I owe my education. My friends were also vital in my survival at MIT, in particular those in APR who were like family away from home. Last but not least, I want to thank Kelly, who for the last year kept me going when times were rough and celebrated with me when they were bright.

Table of Contents

List of Figures

List of Tables

Chapter 1

Introduction

There is a constant need to improve performance in communication systems to increase data rates and transmission distances. So far, electronic communications has been very successful mainly because of advances in the integration of microelectronic devices. The need to decrease device sizes and increase the number of transistors integrated into a single chip is driving electron-based communications to its limit as data rates now exceed 1 Gb/s. Photonics provides one alternative to microelectronic circuits.

Optical systems are currently used in applications where the transmission length exceeds 100 m. Telecommunications, Ethernet and Storage-Area Networks are among the applications where optics is widely used. The ability to fabricate small devices and integrate them into optical circuits is essential to the development of photonic communications. In terms of size, photonic devices are limited by the wavelength of light used. Monolithic integration, on the other hand, reduces fabrication cost, eliminates alignment problems and reduces coupling loss.

Coupling is an issue of particular importance in photonics since optical fibers are used to carry data. That data, in the form of light-pulses, needs to be transmitted (coupled) from the fiber to the chip. Fiber-to-chip coupling is difficult because the large difference in size and refractive index between the fiber core and the chip's waveguide

core creates a modal mismatch that results in power loss. A special type of optical coupler is thus needed between the fiber and the chip to reduce coupling loss.

Several coupler designs have been proposed, including tapered waveguides and grating waveguides. In this work we present a coupler, called the horizontal coupler, based on two sets of waveguides: a thin, horizontally-tapered waveguide and a thick, straight waveguide. The emphasis of the present thesis is to address the challenges of fabricating a multi-mode horizontal coupler using silicon for the tapered waveguides and a polymer, called Cyclotene, for the straight waveguides.

The thesis is divided into six additional chapters. A brief overview of the theory of high refractive-index-contrast (HIC) photonic devices and optical couplers is presented in Chapter 2. This chapter, combined with the introductory chapter, allows the reader to understand where the need for optical couplers arises, what the biggest challenges are in designing the type of horizontally tapered coupler that is presented, and some of the more general fabrication challenges. Chapter 3 explains the particular design to be fabricated. The design specifications that must be met during fabrication are listed in the first section of the chapter, along with the aspects of the design that are most susceptible to fabrication errors. The second section of Chapter 3 discusses and illustrates the devices that will be fabricated to measure and characterize the efficiency of the coupler.

Chapter 4 presents a detailed description of the fabrication process, including the sets of experiments that were conducted, and the challenges faced during the process. This chapter is divided into four sections: one section for the fabrication process of each of the two sets of waveguides, one section that describes the sets of marks that were used for lithographic alignment, and another section that details the cleaving process used to

cut the chips. Chapter 5 then presents the measurements and fabrication results obtained to date. Problems in the fabrication are addressed in this chapter. Finally, Chapter 6 summarizes the work performed and discusses future work. The discussion of future works includes preliminary considerations for the fabrication process of a single-mode horizontal coupler.

Chapter 2

Background

This chapter presents the reader with an overview of the most common devices used in photonic circuits as well as a basic explanation of the theory behind them. Some of the optical devices that can be integrated in a photonic circuit include waveguides, filters, polarization splitters, and switches. The first section of the chapter introduces devices required for optical-network communications. The second section elaborates on the theory and behavior of waveguides; the building blocks of optical-mode converters. Section three discusses the key attributes of Silicon-on-Insulator (SOI) substrates for the fabrication of photonic devices.

2.1 Devices in Optical Networks

The capacity of a single optical fiber, which at present is about 1 Gb/s, could soon exceed 10^{14} bits/s. This capacity exceeds the total current traffic of voice and data in the United States (aproximately 10^{12} bits/s) and could make optical fiber the dominant form of communication in the near future [1]. Nodes in an optical network communicate with one another on specific wavelengths, which are all multiplexed into a single optical fiber in a narrow spectral band centered at 1.55 μm. This technique is called Wavelength-Division Multiplexing (WDM) and requires an optical add-drop multiplexer (OADM) at every node [2].

Before considering the operation of OADMs, it is important to present the most basic optical component: the waveguide. Waveguides direct the optical signal by confining light in a core layer surrounded by cladding materials of lower refractive-index. Electromagnetic modes propagate through a waveguide, according to the principle of Total Internal Reflection (TIR), when light is coupled into it. If only the fundamental mode propagates, the waveguide is called single-mode. Optical fibers are an example of cylindrical waveguides with low refractive-index-contrast.

Waveguides are used to route light between optical devices within a substrate, and between a substrate and an optical fiber. If the propagating modes of the connected devices are different, power is lost at the boundary due to the modal mismatch. This problem is particularly important when connecting an optical fiber, which has low refractive-index-contrast and dimensions of over 5 μm, to a photonic chip with high-index contrast (HIC) and dimensions of approximately 100 nm. An efficient optical-mode converter is needed to match and transform the modes inside the optical fiber to propagation modes that will not leak into the under-cladding of the photonic circuit. Additionally, it must match the effective indices of the optical fiber and the connecting photonic device to prevent power loss due to Fresnel reflection at the boundary [3]. The reflection coefficient is given by

$$R = \left(\frac{n_2 - n_1}{n_2 + n_1} \right)^2 \tag{2.1}$$

Optical add-drop filters are more complex photonic devices and form the basis for OADMs. These devices must drop a data stream at a given wavelength and replace it by a new stream at the same wavelength. There are various designs for the fabrication of

optical add-drop filters, including microring-resonators and racetrack-resonators as shown in Fig. 2.1. For a detailed explanation of their behavior refer to [2]. Waveguides, couplers, and microring-resonators can also be combined to create frequency-selective switches and wavelength converters which are needed for all-optical network routers. An in-depth description of the requirements, advantages and weaknesses of all-optical networks can be found in [4] and [5].

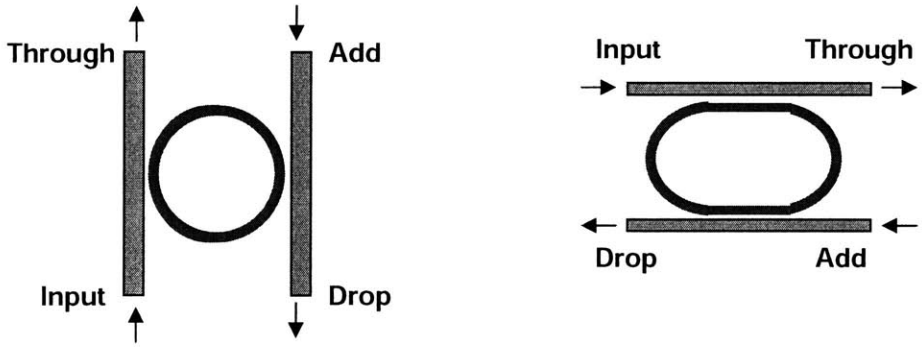

Fig. 2.1: Schematic of microring-resonator (left) and racetrack-resonator (right). Both devices can be used as optical add-drop filters.

2.2 Waveguides

Waveguides in photonic circuits consist of a core layer of material, where light is confined, sandwiched between two other materials. A schematic representation of a waveguide is show in Fig. 2.2. When light propagating through the core reaches the boundary of one of the cladding layers it can either reflect back into the core or transmit through the cladding. If the angle of incidence exceeds the critical angle given by

$$\theta_c = \arcsin \frac{n_3}{n_2} \tag{2.2}$$

all the light is reflected back into the core. Based on this principle, the refractive index of the core material in a waveguide must be higher than that of the claddings, and larger refractive-index contrasts result in better light confinement.

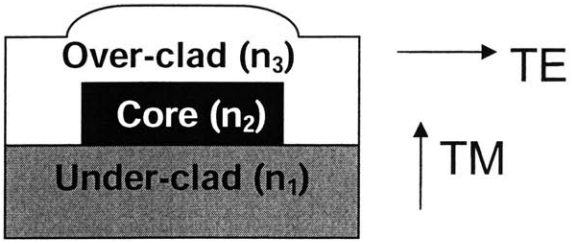

Fig. 2.2: Schematic cross-section of a rectangular waveguide. Light is confined in the core which has a higher refractive index than both cladding materials. The over-clad can be air, as is the case of the fabricated waveguides reported in this thesis, which has a refractive index of 1.

The modes that can propagate in a waveguide are determined by its dimensions and material. Light is an Electro-Magnetic (EM) wave that behaves according to Helmholtz equation:

$$(\nabla^2 - \mu\varepsilon\frac{\partial^2}{\partial t^2})E = 0 \tag{2.3}$$

For the following analysis let us assume that light propagates through the waveguide in the z direction and that the EM wave is TE polarized (electric field perpendicular to the plane of incidence). A similar analysis can be extended to a TM polarized waveguide. A solution to (2.3) under these assumptions is of the form:

$$E_x(y, z) = e^{-jk_z z}(Ae^{-jk_y y} + Be^{jk_y y}) \tag{2.4}$$

To get a basic understanding of how the dimensions and material of the waveguide determine the propagation modes, consider first the simplest structure: parallel plate waveguide. The distance between the parallel plates is d, and the parameters μ and ε are the permeability and permittivity intrinsic in the core material as shown in Fig. 2.3 (a).

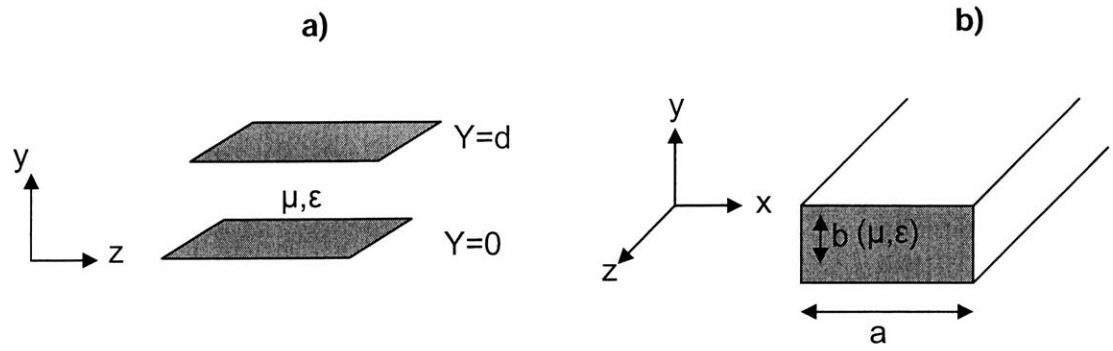

Fig. 2.3: Illustration of parallel-plate waveguide (a) and rectangular waveguide (b).

If no light is to propagate outside the waveguide, the boundary conditions dictate that $E_x = 0$ at $y = 0$ and $y = d$. We can then get an expression for the y and z components of the wave vector, which lead to the guidance condition.

$$k_y = \frac{m\pi}{d} \tag{2.5}$$

$$k_z = \sqrt{\omega^2 \mu\varepsilon - k_y} \tag{2.6}$$

$$\omega > \frac{m\pi}{d\sqrt{\mu\varepsilon}} \tag{2.7}$$

Equation (2.5) indicates that propagating modes constructively interfere every two reflections. The guidance condition dictates that only modes (m values, where m is an integer) that satisfy equation (2.7) can propagate in the waveguide. A mode is evanescent if the corresponding z component of the wave vector has an imaginary part. Evanescent modes decay exponentially from the boundary of the waveguide. If the same analysis is applied to a rectangular waveguide, such as the one shown in Fig. 2.3 (b), the expressions for wave vector and guidance condition are as follow:

$$k = \sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2} \tag{2.8}$$

$$\omega > \frac{\pi}{\sqrt{\varepsilon\mu}}\sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2} \tag{2.9}$$

The integers m and n in equations (2.8) and (2.9) determine the modes in rectangular waveguides.

The most important metric of a waveguide's performance is power loss. Loss in a waveguide is a result of optical absorption of the material and light scattering. Optical absorption can be caused by impurities in the material and distortions in the lattice that trap photons. Scattering is mostly due to surface roughness. These sources of loss lead to an exponential attenuation as light propagates and should be minimized.

2.3 SOI Selection for HIC Devices

SOI substrates offer various benefits to the fabrication of photonic devices. The refractive-index contrast between the Si layer and a buried oxide is 3.5 to 1.45. This high contrast provides strong light confinement if Si is used as a waveguide core and oxide is used for the under-clad, and reduces the waveguide dimensions needed to achieve single-mode propagation. The strong confinement also allows the fabrication of devices with sharp bends and the geometries needed for microring-resonators [6]. Another favorable characteristic of SOI substrates is that both layers are optically transparent at the 1.55μm-wavelength commonly used in optical networks. Several techniques are currently used to produce SOI substrates, including Smart-Cut®, Separation by IMplantation of OXygen,

and Bond and Etch back Silicon-On-Insulator. More information on the different techniques can be found in [7].

Chapter 3

Design

This chapter discusses the specifications of the horizontal coupler's design. The dimensions, choice of materials and optimization parameters are presented in the first section. Also included in the first section are the aspects of the design most susceptible to fabrication errors. The design and simulations presented in this section were done by A.M Khilo. The second section of the chapter presents the devices that will be fabricated to characterize the performance of the coupler. These devices were designed to compensate for probable errors in the fabrication and to measure their effect on the performance of the coupler.

3.1 Design Specifications

The design of the horizontal coupler is based on a straight polymer waveguide and a horizontally-tapered silicon waveguide. An optical fiber is connected to the polymer waveguide and power from the fiber is coupled into the waveguide. The Si waveguide starts inside the polymer and its width is increased over a certain length. Power from the fundamental mode inside the polymer is gradually transferred to the Si waveguide as its width increases. All the power should be transferred by the time the Si waveguide reaches its final width. The width of the Si waveguide is constant as it exits the polymer and carries the power into a photonic circuit. We omit the photonic circuit in our work

and simply connect the Si waveguide to another coupler that carries the power to an output optical fiber. A schematic of the input coupler is shown in Fig. 3.1.

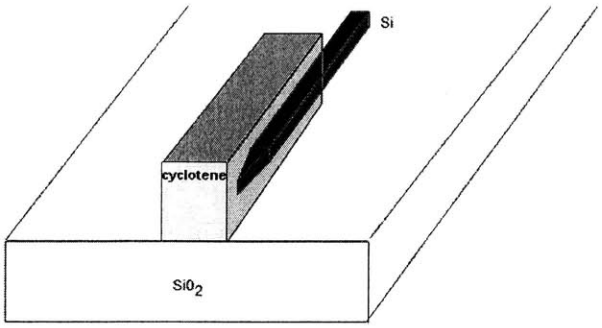

Fig. 3.1: Schematic of the horizontal coupler showing the silicon oxide under-cladding, the straight polymer waveguide and the tapered Si waveguide

The substrate used to fabricate the coupler is Silicon-on-Insulator (SOI). The buried oxide layer provides the necessary under-cladding because the refractive index of silicon dioxide is 1.45, where as that of silicon is 3.5. It also prevents optical modes from reaching into the substrate as long as its thickness is larger than the length of the evanescent fields of the optical modes. The main four contributors to power loss on the structure in Fig. 3.1 that are addressed in the design are: the boundary between the optical fiber and polymer waveguide, the boundary between the Si tip and polymer waveguide, the taper region, and the end of the polymer waveguide.

The parameters that can be optimized to minimize power loss at the fiber-polymer waveguide boundary include the polymer's refractive index, and the width and height of the polymer waveguide. Simulations show minimum coupling loss when the refractive index of the polymer is 1.5 (at a wavelength of 1.55 μm). Cyclotene was the polymer selected because its refractive index at the wavelength of interest is 1.53, it has high

optical transparency and was readily available at MIT. The mode field diameter, radius and refractive index of the optical fiber have to be considered to optimize the width and height of the polymer waveguide. For a multi-mode fiber with a mode-field diameter of 4.35μm and an outer radius of 5 μm, the optimal width and height for the polymer are 5.2 μm and 4.3 μm respectively.

Loss in the boundary between the Si tip and the polymer waveguide depends on the horizontal offset between the two waveguides and on the size of the Si tip. Signal deterioration due to destructive interference occurs when power is coupled into higher-order modes and then back into the fundamental mode. A horizontal shift between the tip of the Si waveguide and center of the polymer waveguide can cause the destructive coupling back into the fundamental mode. To avoid a large coupling this offset should be kept below 1 μm. The size of the Si tip should be kept below 50 nm to keep the optical mode in the tip small. These two constraints require accurate lithographic alignment between the two waveguides and dimensional control.

Loss in the taper region of the Si waveguide is minimized by slowing the modal transformation between the polymer and the Si. This is done by increasing the length of the taper region. The length cannot be made arbitrarily large because an inherit problem of the fabrication process used to create the Si waveguide creates an offset every 100 μm. This fabrication problem, called stitching error and explained in Chapter 4, causes additional power loss. Because there is a tradeoff between minimizing the number of stitching errors and slowing down the modal transformation, three different lengths of the taper region will be fabricated. The thickness of the buried oxide layer and the top Si layer in the available SOI substrate are 2 μm and 240 nm respectively. The final width of

the Si waveguide is optimized for those given values. Simulations show that the optimal width for the straight section of the Si waveguide is 400 nm.

The fourth component of loss occurs at the interface where the polymer waveguide ends. At this interface the Si is no longer covered by the polymer and there is a modal mismatch. If the transition between completely covered by polymer and completely uncovered is gradual then the loss due to the modal mismatch is lower than for an abrupt transition. With that in mind, the end of the polymer waveguide was designed to go off at an angle. The width of the angled section was chosen at 4.6 μm and the length, at 50 μm. These dimensions result in an angle of 5.25°. A schematic with the dimensions of the Si and polymer waveguides is shown in Fig. 3.2. In addition to the specifications already mentioned, the fabrication process must also keep surface roughness in both waveguides at a minimum. Roughness in a waveguide produces scattering loss, which is dependent on the overlap between the optical mode and the rough surface.

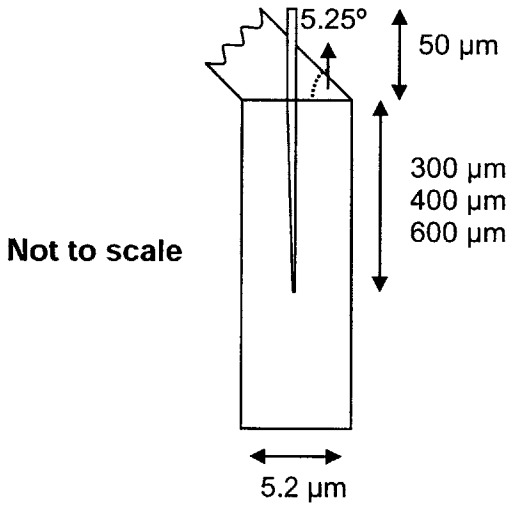

Fig. 3.2: Top-view schematic of the horizontal coupler. The dimensions not shown are the size of the Si tip (50 nm), the final width of the Si waveguide (400 nm), and the length and width of the slanted polymer section (50 μm and 4.6 μm).

3.2 Devices for Optical Characterization

A total of 44 devices were be fabricated in each chip to characterize the performance of the coupler. Marks for lithographic alignment, lines used to cleave the chips, and lines to test lithographic resolution were also designed. A CAD drawing of one of the chips is shown in Fig. 3.3. The alignment marks are discussed and illustrated in Chapter 4. All devices are spaced apart by 100 μm and have a 4mm length.

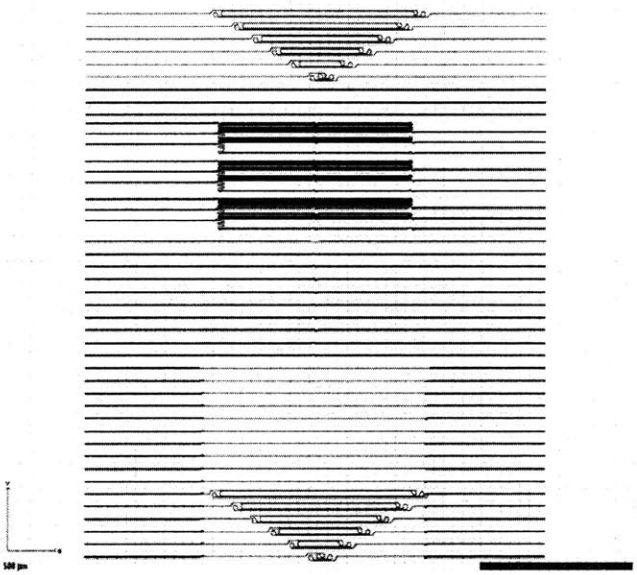

Fig. 3.3: CAD drawing of a horizontal coupler chip, showing all 44 devices.

The first six devices (bottom six in Fig. 3.3) form a set of Si "paperclips" connected to input and output horizontal couplers. Each paperclip consists of straight and curved Si sections. While the length of the curves is the same for each paperclip, the length of the straight sections increases by 1 cm between adjacent paperclips. This allows each paperclip to have different total lengths while fitting the device between the same input/output coordinates. The purpose of these devices is to determine the loss/cm

of the Si waveguides in the coupler, which can be affected by stitching errors, sidewall roughness and material absorption. A closer view of one of the paperclips is shown in the CAD drawing of Fig. 3.4.

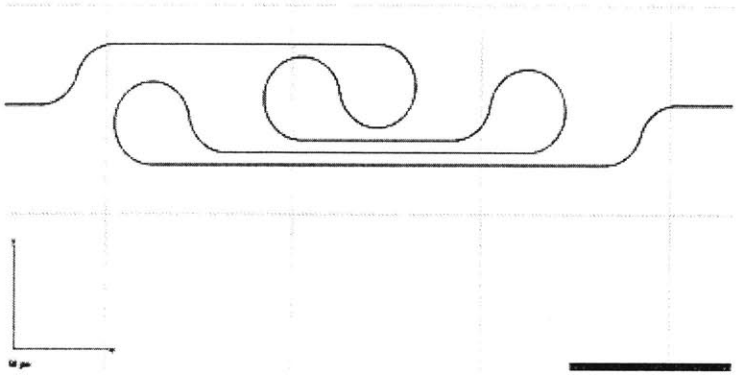

Fig. 3.4.: Close-up CAD drawing of a Si paperclip.

The next 10 devices consist of 2400μm-long straight Si waveguides connecting two horizontal couplers. The taper region of all the couplers in the chip was designed by creating a series of rectangles with linearly increasing width. The length of the taper region was divided into 32 rectangles and starting with the tip-size of 50 nm, each consecutive rectangle increased in width by 12 nm until reaching the final width. This writing strategy was used to minimize bugs in the e-beam write, as explained in Chapter 4. The dimensions were chosen to be multiples of the e-beam step size. In addition to trying the three different taper lengths already discussed, the devices were designed to account for possible misalignment errors during photolithography. Three horizontal offsets were designed for the polymer waveguides (-2/3μm, 0, 2/3 μm). Combined with the 3 different taper lengths, this gives us nine combinations. One additional device was

designed in which the polymer does not have the slanted end section, to test if there really is an advantage to using it. A second set of 10 devices was designed identically, but the length of the connecting Si waveguide was reduced to 100 μm. If there is significant loss in the Si layer, these devices will give us better results.

Couplers cascaded front-to-back were designed to more accurately determine the total loss of the horizontal coupler by obtaining average measurements. Three input/output-coupler-pairs are cascaded in one device, nine in another device, and fifteen in another one. Some of the couplers are connected through straight Si waveguides, while others are connected through U-turn waveguides. The three different horizontal offsets were used in these devices, giving a total of 9 devices, and the taper length was fixed at 400 μm. Fig. 3.5 shows the cascaded couplers.

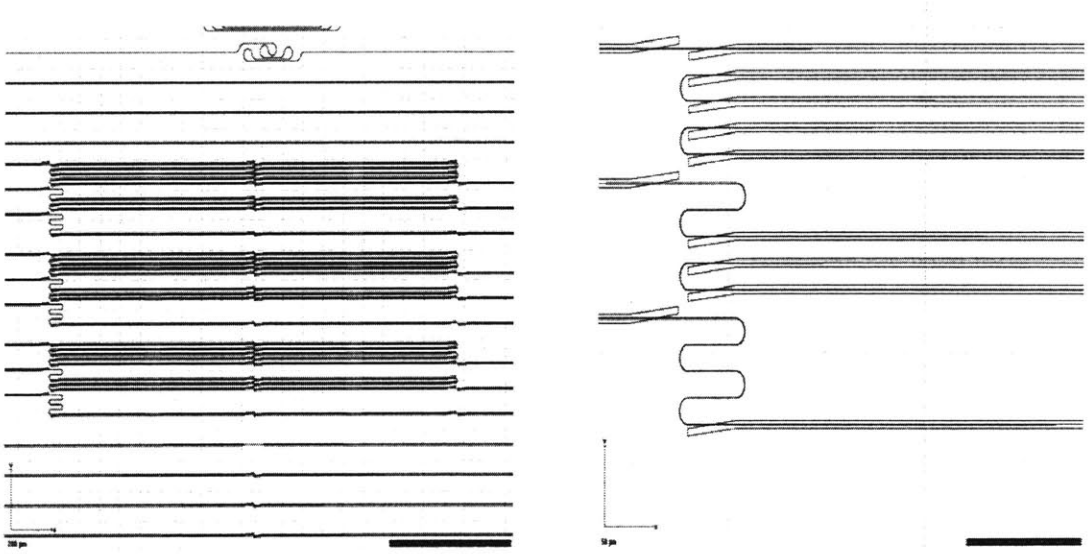

Fig. 3.5.: CAD drawings of the cascaded horizontal couplers. The image on the left shows all nine devices (3 sets of 3, 9 and 15 cascaded couplers). The right image shows a closer look of the cascaded couplers at the output terminal.

The remaining 9 devices are either pure polymer or pure Si. The first three are straight polymer waveguides used to test if there is high loss associated with the polymer. The last 6 devices are another set of Si paperclips, but not connected to any couplers. These can be used to measure the loss in the Si waveguides in case the polymer is too lossy, or to determine if there is a problem purely related to the Si layer.

Chapter 4

Fabrication

Chapter 4 presents a detailed description of the fabrication process needed to create the couplers. The major challenges addressed are dimensional control of the waveguides, lithographic alignment and etching of the polymer waveguides. Fabrication can be divided into two main processes, one to form the silicon waveguides out of the initial SOI substrate, and a second process to create the polymer waveguides above the silicon layer. It should be noted that these two processes cannot be treated independently, and some of the greatest fabrication difficulties arise from making them compatible. Precise alignment between the silicon and polymer waveguides is explained separately in a sub-section that describes the alignment marks created on each layer.

After the fabrication of both sets of waveguides is completed, the chips must be carefully cleaved to allow for good optical characterization. Obtaining a clean and smooth cleave is not trivial and the process used, called assisted-cleave, will also be explained in this chapter.

4.1 Silicon-Waveguide Fabrication

Fabrication of the Si waveguides is based on scanning-electron-beam lithography (SEBL). This part of the overall fabrication process requires the most complex shapes, smallest feature sizes and highest resolution. All design specifications must be met while

maintaining low roughness on the waveguides' sidewalls. A diagram of the process is shown in Fig. 4.1.

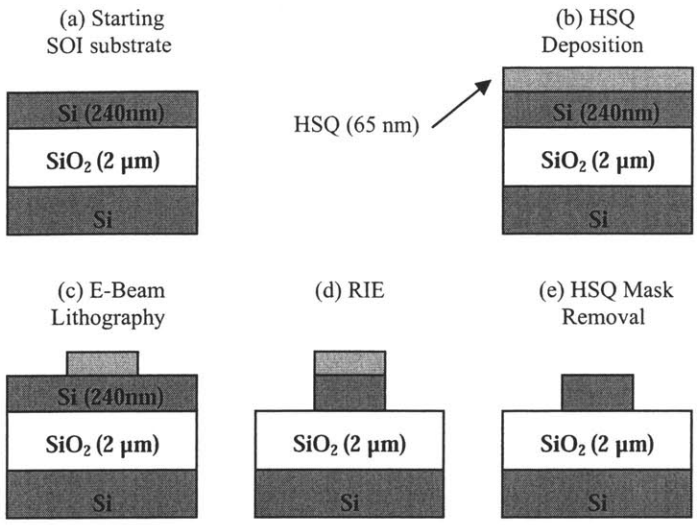

Fig. 4.1: Si waveguides fabrication-process overview (a) Initial SOI substrate formed 240nm of silicon and 2 μm of buried silicon oxide. (b) 65 nm of HSQ are spun on the substrate. (c) Scanning-electron-beam lithography and HSQ development. (d) Reactive-ion etching of top silicon layer. (e) HF removal of remaining HSQ mask

The starting substrate for the fabrication consists of 240 nm of silicon, out of which the waveguides are fabricated, and 2 μm of silicon oxide undercladding. The first step of the fabrication process consists of depositing 65nm of hydrogen silsesquioxane (HSQ). HSQ, commonly used as low-k dielectric, acts as a non-chemically-amplified high-resolution negative resist. Negative resists can be directly written using SEBL, because areas of the resist that are exposed to the electron beam are chemically modified to be less soluble to a developer than non-exposed areas. HSQ is a good choice because sub-50 nm resolution is achievable, it can be easily developed, and when cured it becomes a porous oxide that acts as an etching mask with good selectivity over silicon.

A mask thickness of 65nm is enough to withstand a 240nm Si etch and can be obtained by spinning XR-1541 HSQ at 2500 rpm for 60 seconds. Immediately after the spin, a hotplate bake at 90 C for 60 seconds is used to harden the film, and improve adhesion and uniformity.

The second step in the fabrication, e-beam lithography, is done with a Raith 150 SEBL system and several exposure parameters need to be optimized to obtain the desired features. The first parameters to be considered are e-beam energy, aperture and step size. Higher e-beam energies improve resolution and, thus, in order to achieve the sub-50 nm feature sizes required the maximum energy of 30 KeV was chosen. In selecting a good step size there is a tradeoff between increasing the writing speed at larger step-sizes and increasing deflection errors, which affect the position and size of the exposed structures, at smaller step-sized. A size of 6nm was found to be a good compromise between the two [2]. It is important to note that once a step size is selected, the dimensions of the devices must be selected to be multiples of that size in order to maximize accuracy when the devices are discretized. Finally, an aperture of 30 μm was found to work well for HSQ exposures at 30 KeV.

The next key parameter that needs to be optimized is exposure dose, but before discussing that topic it is worth mentioning one of the limitations inherited with the Raith 150, the stitching error. Structures to be written in the Raith are divided into 100x100 μm fields. Exposures are done by scanning one field and then moving the stage where the sample is mounted to the next field. At the boundary between two fields there is a field-stitching error that creates an offset in any structure that crosses the boundary. The stitching error has a mean of zero and a standard deviation of 20nm if the fields are

perfectly calibrated prior to the exposure. Any rotational error in the calibration would increase the mean of the stitching error. These stitching errors can create considerable loss in the signal at every boundary and, thus, can quickly add up over the length of large structures.

With respect to the parameters that need to be optimized, it is now essential to optimize exposure dose. The optimal dose is inversely proportional to the width of the structure to be written. The minimum dose required to generate each desired line width was chosen as the optimal dose, as higher doses increase sidewall roughness. Fig. 4.2 shows two 402-nm-lines exposed with different doses; it can be seen that the line with higher dose has significantly worse sidewall roughness. Additionally, all doses used must exceed the clearing dose, defined in [2] as the minimum dose required for the developer to dissolve the entire resist in a given time.

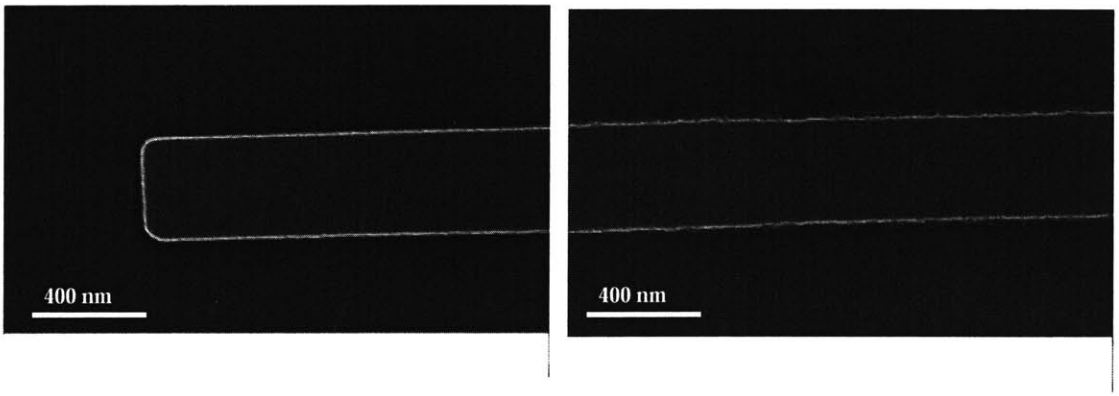

Fig. 4.2: Impact of dose-increase on sidewall roughness. Both lines have the same width, but the right line has a higher dose and the resulting sidewall roughness is worse.

Two calibration runs were performed to determine the optimal doses to be used for each line-width and analyze any errors in the writing strategy. One calibration was done on a silicon substrate and consisted of a dose matrix, three 4-mm lines, and four 500-μm trapezoids. Each element in the dose matrix consisted of five lines with widths

ranging from 450 nm to single-pixel lines. The purpose of the matrix was to determine the minimum doses required along the Si taper and the straight Si waveguides. Doses used for single-pixel lines must de defined in the Raith as line doses, whereas the other doses are defined as area doses. Only one area and one line dose can be specified at the time of exposure, but each structure can be designed to have a different percentage of the specified corresponding dose. The specified percentage cannot exceed 3000%, as this will lead to errors in the exposure.

The 4-mm lines were used to determine the average magnitude of the stitching error. The trapezoids were created to determine if a simple trapezoid figure divided into different segments of different doses was good enough to produce the desired taper. From this run it was determined that defining the tapered waveguides as trapezoids was not adequate as it was hard to produce a smooth and linearly decreasing transition between the areas with different doses, especially across field boundaries. This was the main reason why the writing strategy detailed in Chapter 3 had to be used. Because this calibration run was done on a Si substrate, not SOI, it was not ideal to obtain the optimal dose values from this run but rather an estimate.

The second calibration run was done on an SOI substrate and incorporated the writing strategy for the tapers that was mentioned in Chapter 3. In addition to the tapers, this run also included a dose matrix, as explained earlier, and 4-mm lines. From this calibration run it was determined that the new writing strategy for tapered waveguides worked better than the one based on trapezoids. It was found that using an exponential profile for the doses along the tapers was the best approach to obtain the desired dimensions and slope. Abrupt jumps in the size of the tapers can be avoided by properly

selecting the slope of dose's exponential profile; this is particularly important for the smaller-sized sections of the taper. It was also determined that the base doses used were too large to achieve the smallest target dimensions. Finally, measurements of the stitching error were on the order of 44 nm. Fig. 4.3 summarizes these findings.

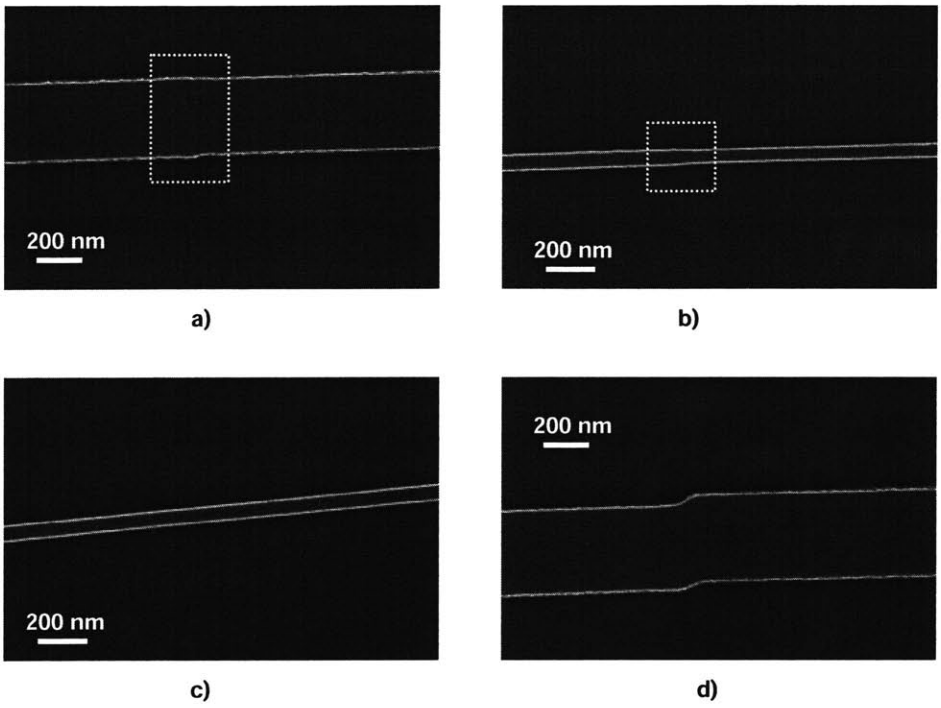

Fig. 4.3: SEM results of SOI calibration run after exposure and development. a) Tapered waveguide showing a smooth and small transition between the two largest-sized rectangles. b) Tapered waveguide showing a steeper size change between the two smallest-sized rectangles. c) The lowest-dose single pixel line had a measured width of 70nm. d) The average stitching error in the 4-mm boxes was 44 nm.

Because of time constraints and to avoid the cost of another calibration run, it was decided to use the information obtained from the previous two calibrations to extrapolate the optimal doses. SEM micrographs were taken of the exposed and developed Si and SOI samples and of the SOI sample after etching and mask removal. SEM measurements have a 5% error, but this number is acceptable when analyzing the most critical features which have a width of 50 nm.

The un-etched Si and SOI samples were compared to find structures with overlapping doses and determine the bias between the two. The un-etched and etched SOI structures were then compared to quantify the amount the size changed after etching. Adding this number to the previous bias gives us the total width bias that must be added to a structure in the Si calibration to find the corresponding size in an etched SOI sample. Development and etching conditions are discussed later this chapter.

The strategy used to extrapolate the doses consists of four steps. The first step is to find the dose in the Si calibration that would create a 402-nm-wide rectangle after adding the total bias described above. The second step is to find the dose in the Si calibration that would create 62-nm-wide rectangle after taken into account the total bias. Note that the design specification for the tip of the tapered waveguide is a width of 50 nm. As described in Chapter 3, the writing strategy for the taper is to create a set of 31 such that the width difference between adjacent boxes is 12 nm. Therefore, the required width of the second smaller rectangle in the taper is 62nm. Once the required dose for the 402-nm-wide and 62-nm-wide rectangles are known, the doses for the 29 rectangles between can be extrapolated by fitting the data to an exponential function of the form $(A)e^{Bx}$. In order to reach the 50-nm-width of the smallest rectangle we use a single-pixel line, and in the last step of the extrapolation we find the required line dose in the Si calibration that would result in a 50-nm-wide rectangle after adding the total bias. This dose represents the 32^{nd} and final data point on the exponential curve.

The bias measurements are summarized in table 4.1. Using a base area-dose of 398 $\mu C/cm^2$ the resulting parameters of the exponential fit are A = 151.52 and B = 0.0462. These values correspond to an exponential fit between the area doses and the

corresponding rectangle number (1-32). Each rectangle corresponds to a specific width (402nm – 62nm), thus the positive-exponential fit above can be converted to a negative-exponential plot, shown in Fig. 4.4, which indicates the percentage of the base area-dose required for each width. Additionally, the line dose required for the 50-nm-wide rectangle is 3,659 μC/cm.

Table 4.1

Bias Measurements from the Calibration Runs

Intended Width	Average Bias #1	Average Bias #2	Average Total Bias
402 nm	30 nm	16 nm	46 nm
30 nm	16 nm	4 nm	20 nm
single pixel	20 nm	10 nm	30 nm

Average bias #1 corresponds to the width difference between identical structures in the developed SOI sample and the developed Si sample. Average bias #2 corresponds to the increase in width after the structures in the SOI sample are etched and the mask is removed.

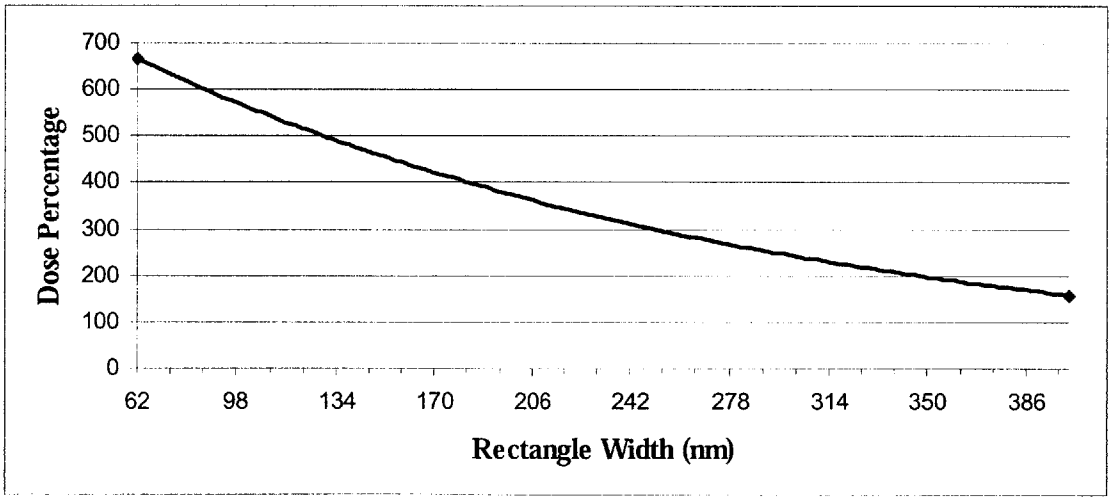

Fig 4.4: Exponential fit for the percentage of the 398μC/cm^2 area-dose used in each of the segments of the tapered waveguide, not including the 50-nm-wide tip.

After the sample is exposed with the correct doses, the following step in the fabrication process is development. The sample is developed in 25%

Ttetramethylammonium hydroxide (TMAH) for 2 minutes at room temperature. TMAH is an ammonium-salt solution that dissolves the HSQ that was not e-beam exposed. After this step the sample has the desired patterned-etch-mask.

Etching of the silicon waveguides needs to be highly anisotropic and have good selectivity over the oxide undercladding. A highly anisotropic etch is one in which almost no horizontal etch occurs and thus the sidewalls have a vertical profile. Selectivity, on the other hand, is defined as the ratio of the etch rate of two films. Selectivity between the silicon and the etch mask also need to be high, but if selectivity over the oxide underclad is high then it should also be high over the mask because cured HSQ is essentially porous silicon oxide. Our etching requirements can be met by using reactive ion etching (RIE) with either a chlorine based chemistry or hydrogen bromine (HBr). In RIE, the etchant gas is flowed into a chamber where an RF plasma is created. Ions and free radicals create a combination of physical and chemical etching, where the ratio of the two is affected by the voltage drop and gas composition. HBr was chosen for our application because it produces more vertical sidewalls and has higher selectivity over oxide than a Cl-based chemistry [7].

The etching step was divided in two to obtain the desired results. The first step requires a high-voltage bias for a short period of time. A high voltage is required initially in order to break through the native oxide present on top of the silicon. A thin native-oxide-layer of a few Armstrongs always forms whenever a silicon layer is exposed to air. Because HBr does not chemically etch silicon oxide, the physical component of RIE must be increased so that the ions can sputter away the oxide. However, this high-voltage step cannot last for long because both mask erosion and sidewall roughness are higher as the

voltage is increased. An initial step of 10 seconds at 350 V was found to be enough to remove the native oxide without significantly etching the mask or increasing roughness. A second step of 18 minutes at 150 V was then used to etch the 240 nm of silicon. Both steps are done at a pressure of 10 mT and with an HBr flow of 20 sccm. It was found that 40 nm of the HSQ mask were removed during etching. For improved accuracy, all height measurements were taken with a Dektak depth profiler. A Dektak scans a stylus tip across a surface and the height of a structure can be determined by a plot of the stylus's vertical position. Fig. 4.5 shows the resulting silicon waveguide after RIE. The remaining HSQ is removed in the following fabrication step.

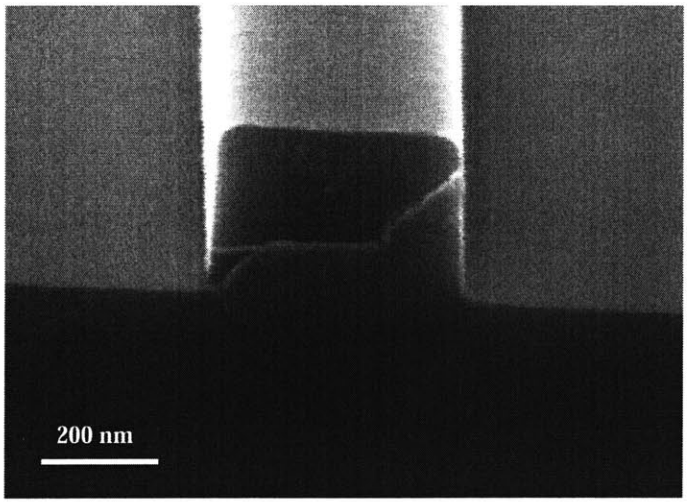

Fig. 4.5: Straight Si waveguide after RIE. Width was measured at 400 nm and height at 265 nm, including 25 nm of the HSQ mask.

Once the RIE is completed, the sample is cleaned in a solution water:ammonium-hydroxide:hydrogen-peroxide (5:1:1) at 85 C. This cleaning step ("RCA") removes any metal contaminants that may be present in the sample and that could lead to silicide formation during subsequent fabrication steps. Silicide formation, which occurs when

metallic residue reacts with Si at high temperatures, prevents light from propagating through the waveguides via optical absorption. The RCA clean damages the remaining HSQ mask, increasing the roughness at the top of the Si waveguides, and requires it to be removed. The remaining HSQ can be removed with a hydrofluoric (HF) acid wet-etch. Etch of the oxide under-cladding is not a major concern for the multi-mode coupler because the Si waveguides do not have an over-clad. If significant etch of the under-clad occurs it can lead to an undercut of the Si waveguides and, in turn, poor coverage of the over-clad. Although creating and undercut is not a critical concern in this application, we still want to minimize etching of the oxide and develop the HF etch accordingly. A highly-diluted 49%-HF:water solution (1:400) was found to have the highest selectivity of HSQ over oxide. Immersing the sample in the solution at room temperature for 4 minutes completely removes the remaining HSQ while etching only 5-10nm of oxide. This step completes the fabrication process for the Si waveguides.

4.2 Lithographic Alignment

Precise alignment between the Si and polymer waveguides is critical in minimizing coupling loss, as described in Chapter 3. One set of alignment marks, based on Moiré patterns, was created on the Si substrate and a complimentary set was defined in the photomask for the polymer waveguides such that the misalignment between the two would be less than 1 μm. A schematic of the first set of marks is shown in figure 4.6.

The alignment mark shown in Fig. 4.6 is based on six sets of gratings and a cross for coarse alignment. The cross on the right side has four 50x50μm quadrants and is complimented on the photomask by a set of four 50x50μm boxes. The two

complimentary sets are used to obtain an initial alignment. The six gratings, used twice in the alignment mark, are then used to obtain a more precise alignment both vertically and horizontally. A second set of alignment marks is defined in another position of the Si substrate and the photomask, and used to prevent rotational misalignment. The gratings with periods P1 and P2 are positioned side by side, with P1 on the Si substrate facing P2 on the photomask and vice-versa. Interference fringe patterns, also called moiré patterns, form when the sample and mask are imaged. Alignment is done by moving the sample with respect to the mask, which causes the fringe patterns to move in opposite directions. Alignment is obtained when the spatial phase of both sets of fringes match [8]. This is better illustrated in Fig. 4.7.

P1 and P2 were chosen to produce the required alignment of 1 μm. P3 and P4 work in the same way, but were designed with smaller values in order to achieve a stricter misalignment of 500 nm. Two additional sets of gratings, Pf and Pf*, are needed because alignment as defined above is periodic. When the substrate is moved by a certain distance, called capture range in [8], the fringes will re-align. The values of Pf, which faces Pf* on the photomask, and Pf*, which faces Pf on the photomask, are designed such that if the resulting fringe patterns are also matched then we have reached the desired alignment. Finally, it should be noted that all sets of fringes fit within a 100x100 μm window in order to prevent stitching errors during SEBL.

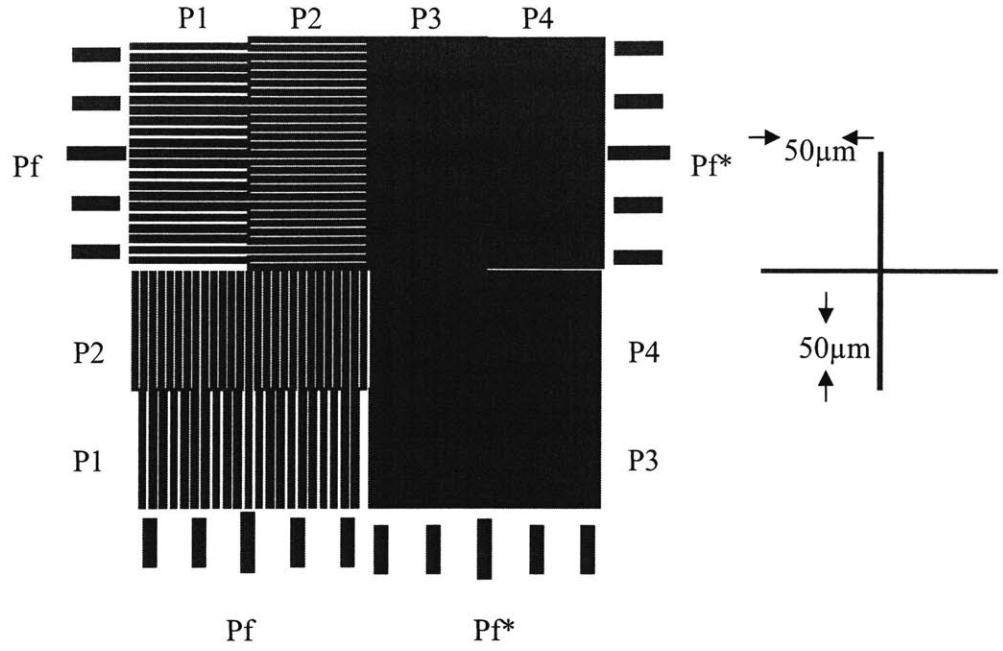

Fig. 4.6: Schematic of the alignment marks created on the Si layer. The cross on the far right provides coarse alignment. P1 and P2 are 4.5μm and 3.7 μm, and provide alignment accuracy of 1 μm. P3 and P4 are 2.2μm and 2.0 μm, and provide alignment accuracy of 500nm. Pf and Pf* (20.83 μm and 22.0 μm) ensure that all sets of fringes align only once within the 100 μm window. Gratings on the top half provide vertical alignment, and gratings on the bottom half provide horizontal alignment.

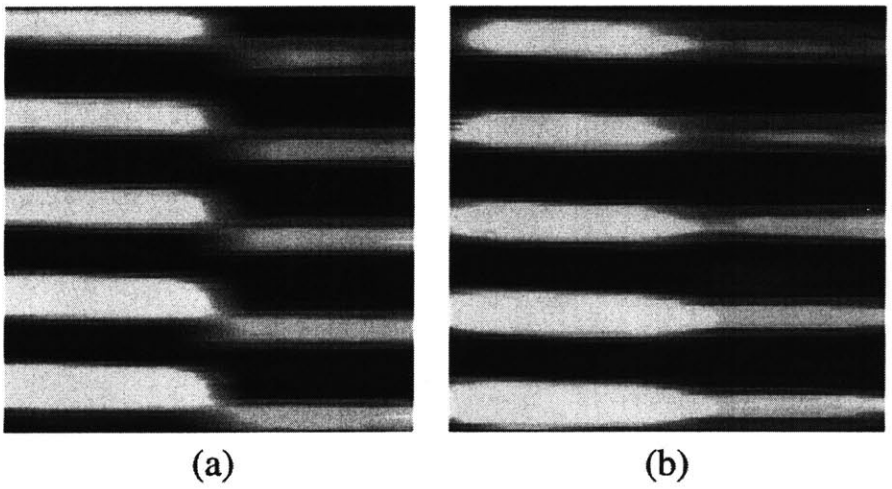

Fig. 4.7: Images of interference fringe patterns. (a) The spatial-phase difference on the left indicates misalignment. (b) The spatial-phase match on the right corresponds to proper alignment. Figure taken from [8].

4.3 Polymer-Waveguide Fabrication

This section of Chapter 4 presents the process flow used to create the Cyclotene polymer-waveguides above the Si layer. Discussion of the fabrication process includes: uniform deposition of the polymer on top of the sample, proper thermal solidification of the polymer, selection of the etch masks used to define the waveguides, and Si-compatible etching of the polymer. An overview of the process is illustrated in Fig. 4.8.

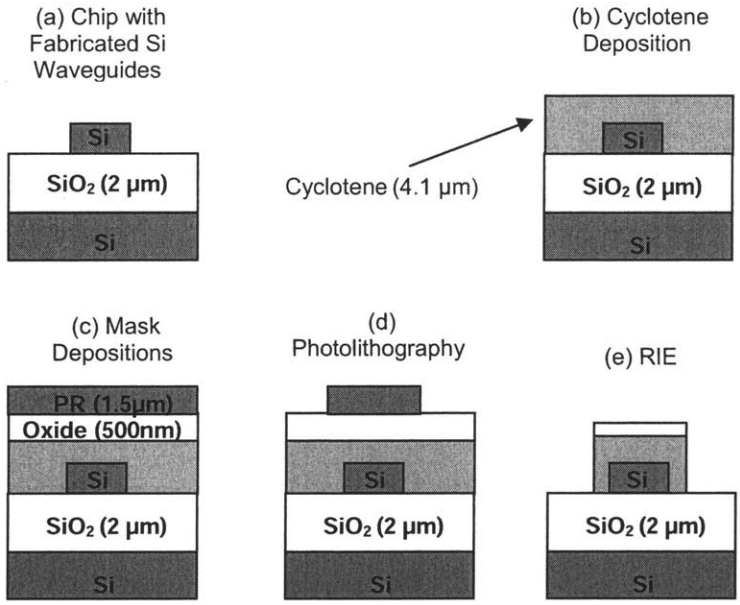

Fig. 4.8: Polymer waveguides fabrication-process overview (a) The sample initially contains 240-nm-tall Si waveguides and a silicon oxide underclad of 2µm. (b) 4.1 µm of the Cyclotene polymer are spun on the sample. (c) 500 nm of silicon oxide used as a hard mask are CVD-deposited, followed by 1.5µm of photoresist. (d) The photoresist is selectively-exposed using contact lithography and then developed to define an etch mask for the oxide. (e) RIE of the oxide hard mask followed by RIE of the Cyclotene complete the process.

Cyclotene is available in two types, photosensitive and dry-etch. The first type can be directly patterned by photolithography, but has a very short shell-life which makes it impractical for research applications. There are also different versions of dry-etch

Cyclotene, which can be spun up to different thicknesses. The dry-etch Cyclotene already available for this work was model 3022-35. The first step of the fabrication, spinning the polymer unto the sample, requires great care. The surface of the sample must be free of particles and residue, which create comet-style defects in the film and could also lead to poor adhesion. Contamination in the dispense mechanism will also result in film defects. Prior to Cyclotene coating, the surface of the sample should be covered with an adhesion promoter. The Dow Chemical Company recommends the AP3000, a promoter composed mostly of propylene glycol monomethyl ether, to prevent lamination of the polymer. The adhesion promoter was statically dispensed, spread at 600 rpm for 7 seconds and then at 3000 rpm for 30 seconds. It is important that the adhesion promoter is uniformly and cleanly deposited, because problems in this step will translate into a poor Cyclotene coating. Pre baking the sample at 90 C for 5 minutes was found to help in the quality and uniformity of the film.

In order to obtain the desired polymer thickness, two separate spins are required because the maximum thickness achieved by the 3022-35 model is approximately 2.2 μm. To achieve high thickness, each spin must be done at low speeds, which make it difficult for the polymer to spread evenly across the sample. The process found to minimize this problem is to statically dispense the Cyclotene in the center of the sample and use an amount that covers most of its area. The polymer should then be initially spread at 500 rpm for 7 seconds and then increase the speed to the desired value and spread for 30 more seconds. A final speed of 1000 rpm for the first spin and 1300 rpm for the second spin resulted in an average total film thickness of 4.1 μm. The thickness was found to change by up to +/-100 nm across a quarter-six-inch wafer. 4.1 μm was the

target thickness instead of the original 4.3 μm for reasons that are explained later in this section. If Cyclotene is deposited on the backside of the sample during the spin it is important to clean it with a solvent before proceeding to bake the sample, because the high temperature will cause the backside to adhere to the surface it is placed on.

A 5-min bake at 90 C was done immediately after each spin to stabilize the film prior to any following processing. After the first Cyclotene spin the sample needs to be baked to achieve a partial thermal cure. This cure should reach about 80% polymerization of the Cyclotene and improves adhesion between the two polymer layers. The extent of polymerization depends on the temperature and duration of the cure, and the recommended partial polymerization is obtained with a 40-min cure at 240 C. Adequate curing is of key importance in the fabrication of the polymer waveguides. If Cyclotene is exposed to oxygen at temperatures above 150 C the cyclobutene group can be oxidized; this would be detrimental to the waveguides' performance. It is important that the cure is done in an inert ambient. Furthermore, a slow temperature ramp is preferred in order to improve planarization of the film. The partial cure was done by placing the sample in an oven at 90 C and then flowing nitrogen. The oven was left at 90 C for 30 minutes before raising the temperature to 150 C for another 30 minutes. After one hour of curing, the temperature was slowly ramped to 240 C and left at that value for 40 minutes. The process for the final cure, which is done after the second Cyclotene spin and achieves a polymerization of over 95%, only differs in that the final temperature is set to 250 C for 60 minutes. It is important, for both cures, that the oven be allowed to cool down below 150 C before removing the sample and that the samples be placed horizontally to minimize any polymer flow that could affect uniformity. Cracking of the

polymer film was observed during subsequent processing if the final cure was not done properly.

The following fabrication step, etch-mask deposition, depends strongly on the RIE process that is used on the final step. For that reason, both steps will be discussed simultaneously. Because the type of Cyclotene used is not photosensitive, a layer of photoresist is needed to transfer the desired patter from a photomask through contact lithography. Photoresist has a poor etch selectivity of almost 1:1 over the Cyclotene, and cannot be used as an effective etch mask if vertical sidewalls are required. A hard mask, either metallic or inorganic, needs to be deposited before spinning the photoresist. If a metallic mask is used it must be removed after etching without damaging the Cyclotene waveguides. This would add great complexity to the fabrication process that can be avoided by using an inorganic mask. Silicon oxide was chosen as the hard mask because of its ease of deposition and compatibility with the other materials.

A series of experiments were performed to optimize plasma etching of the Cyclotene. Oxygen is required in the plasma to attack the benzocyclobutene (BCB) and benzene rings from which Cyclotene is derived. Cyclotene is different from most polymers in that silicon is present in its backbone, and fluorine is generally required to complement the oxygen plasma. Although fluorine is recommended to etch Cyclotene, etching of BCB with pure oxygen plasma has been reported [9] and was investigated in the first set of experiments. The presence of fluorine in the plasma etch is undesired because it causes mask erosion. The rate at which the hard mask is eroded has to be considered, not only to avoid its total etch, but also because it affects the sidewall profile

of the waveguides. Erosion at the edge of the mask is both vertical and lateral, and as illustrated in Fig. 4.9 as the etch progresses this results in slanting of the sidewall.

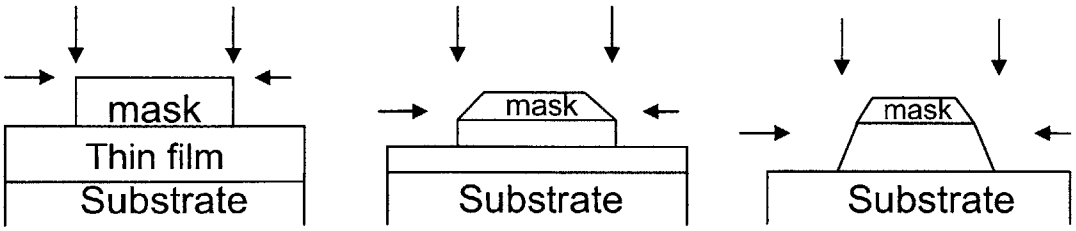

Fig. 4.9: Effects of mask erosion on sidewall profile. A facet forms as the mask erodes horizontally, which leads to a slanted sidewall-profile on the film.

To test etching with pure oxygen plasma, 80 nm of silicon oxide were CVD-deposited to serve as hard mask and 1 μm of OCG-835 positive photoresist was spun on top of the oxide. After spinning and pre-baking the photoresist, the sample was aligned to the photomask as explained in section 4.2 and then exposed. The exposure system used has a wavelength of 320 nm and delivers 4mW/cm^2. Good contrast in the resist was obtained with an exposure time of 45 seconds and development in OCG-934 for 1 minute. The photomask pattern is then transferred to the hard mask by etching the oxide in those areas not covered by the remaining photoresist. This was done in an RIE chamber using CHF_3 plasma at 300 V for 3 minutes. A cross-section of a test sample after the oxide RIE is shown in Fig. 4.10.

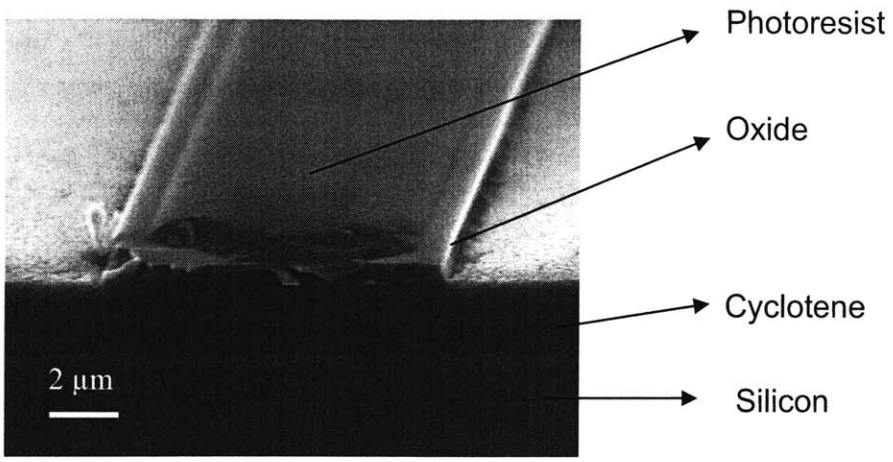

Fig. 4.10: Cross-section showing the different layers in a polymer waveguide after the oxide hard mask is selectively-etched. Only 2 μm of Cyclotene were deposited in the sample shown.

The first RIE conditions tried were 20 sccm of pure oxygen plasma at 300 V and 500 V for 5 minutes. The total etch at 300 V was 0.4 μm, while at 500 V it was 0.45 μm. Experiments with longer etch times showed that in the absence of fluorine only 0.5 μm of Cyclotene were cleanly etched. After this initial etch, a "grassy" layer forms in the surface and cannot be sputtered. This occurs because in the absence of fluorine amorphous silicon oxide forms on the surface of a sample and slows the etch rate [10].

A plasma containing 10% CF_4 (18sccm O_2 : 2sccm CF_4) was investigated next. At 300 V this plasma could not sputter away the "grassy" layer around the waveguides. The layer started to be sputtered at 500 V, but the fluorine concentration was still too low. Before further increasing the ratio of fluorine in the plasma, a thicker mask of silicon oxide is needed to withstand the etch. Fig. 4.11 shows the result of this last etch as well as some of the previous tests.

From the experiments conducted with the 80-nm oxide mask, it was determined that in a O_2/CF_4 (18:2) plasma the selectivity between Cyclotene and silicon oxide was

roughly 11:1. Selectivity is a function of the percentage of CF_4 in the plasma as well as bias, but the 11:1 ratio provides us with a good idea of what thickness to use for the new hard mask (> 400 nm). The next set of experiments was done using 500 nm of silicon oxide and a thicker photoresist layer of 1.5 μm. The selective etch of the hard mask was done in three seven-minute intervals using CHF_3 plasma at 300 V. The etch was split in three because as it progresses the chamber gets dirty and this causes fluctuations in the power. Stopping the etch every 7 minutes and cleaning the chamber prevented the power from increasing by more than 50%.

A new test was done by etching the Cyclotene for 48 minutes in a16:4 (O_2/CF_4) plasma at 500 V. This etch showed significant improvement in etch rate and in removing the "grassy" layer at the surface, but was still not enough. By increasing the ratio of CF_4 in the plasma to 16:8 (O_2:CF_4) we obtain a clean etch of the Cyclotene where the entire residue in the surface is finally removed. Cross-sections showing the result of these two etches are shown in Fig. 4.12.

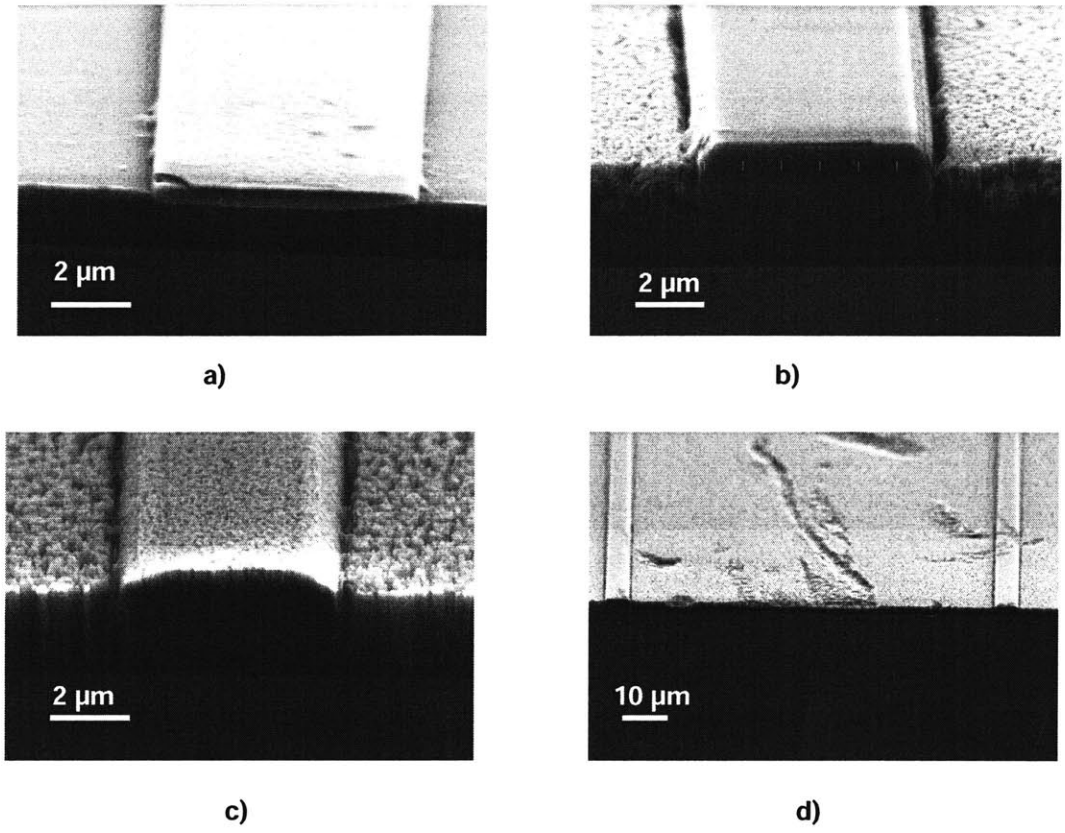

Fig. 4.11: Initial RIE results for polymer waveguides with an 80-nm hard mask. (a) Pure oxygen plasma at 500 V for 5 minutes. (b) A "grassy" layer forms in the surface after an initial clean etch and cannot be sputtered in the absence of fluorine. (c} A plasma containing 10% CF_4 at 500V completely removes the oxide hard mask. (d) The "grass" starts getting sputtered under the last etch conditions, but is still not enough. Notice that in all of these images the width of the waveguides exceeds the desired 5.2 µm, but as etching is optimized the width will shrink to the correct value.

Several key observations can be made from the results shown in Figure 4.12. First, the etch rate of Cyclotene increases as the percentage of CF_4 in the plasma goes from 0 to 33%. This finding agrees with what is reported in [11], where it was found that maximum etch rates occur with about 60% CF_4. Second, mask erosion resulted in a non-perfectly-vertical sidewall profile. The width of the waveguides was measured at 5.2 µm (the designed value) at the base and 4.8 µm at the top. This represents a change in width of 8% and a sidewall angle of 87°. It was also found that, on average, 200 nm of the oxide mask remained un-etched. Simulations performed to account for the 8% difference

in width indicate no significant change in the performance of the coupler, while simulations performed to account for the remaining oxide mask indicate a negligible effect if the height of the polymer is reduced by the amount of the remaining mask [simulations performed by A.M. Khilo]. In order to compensate for the additional 200 nm of oxide, Cyclotene was spun to a total thickness of to 4.1 μm instead of the original 4.3 μm as indicated previously.

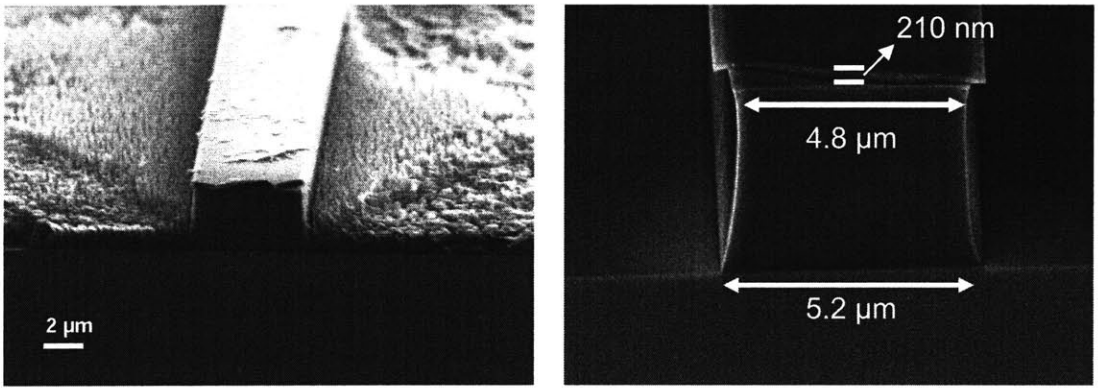

Fig. 4.12: RIE results of polymer waveguides with a 500-nm hard mask. (a) An increase in the percentage of CF_4 in the plasma from 10% to 20% shows significant improvement in Cyclotene etch. (b) Etching with a 33%-CF_4 plasma at 500V resulted in waveguides with a 5.3-μm-wide base, a sidewall slope of approximately 87°, and 210 nm of remaining oxide mask.

Once we know that the 2:1 ((O_2:CF_4) plasma at 500 V etches the Cyclotene cleanly we do not want to further increase the percentage of CF_4 and need to determine if the Si waveguides will be partially etched under those conditions. Fluorine atoms etch silicon with a rate that is linearly dependent on their concentration. The etch mechanism is mostly chemical and based on the formation of SiF_2 groups onto the silicon surface [12]. An initial test to determine the extent of the Si etch was done based on the

perceived color change in an SOI sample. Changes in the thickness of a thin Si film affect its reflectance, which in terms changes the perceived color of the film. A reflected wavelength spectrum can be calculated by considering the thickness and refractive index of the substrate and the Si layer. This spectrum can be converted into color-scheme (such as RGB) parameters [13]. The result is a plot that shows the power reflectivity percentage of each color of the scheme over a range of Si thicknesses. An SOI sample can be visually compared to the plot to determine where it falls. The curve that each color traces in the plot is periodic and each color has a different period. Thus, the color seen in the SOI sample may belong to a set of different Si thicknesses. The plot used was generated by T. Barwicz.

An un-patterned SOI sample was etched in the 2:1 (O_2:CF_4) plasma at 500V for 6 minutes. The resulting color of the sample was marked on the color chart at 2 different locations. To resolve which was the correct thickness the etch was repeated for another 6 minutes and the new color was marked on the chart. Knowing the location of this new data point allows us to determine which of the first two points was correct. It should be noted that this test based on the color chart only provides a rough estimate of the Si etch rate, but it allows for an estimate without the need to e-beam write the sample which is both time and money consuming. Test results showed a decrease in Si thickness from 240 nm to approximately 215 nm. Although the number is just an estimate, it was clear from the test that there was a finite etch rate of Si.

Knowing that a 33%-CF_4 plasma will partially etch the Si waveguides, it is necessary to switch to a different plasma at the end of the etch. Five new plasmas were tested based on the SOI color chart to determine which had the lowest Si etch-rates. The

first three have 0%, 10% and 20% CF_4, respectively, at 500 V. The other two are based on alternating CHF_3 and O_2 intervals. In the first, the CHF_3 intervals lasted 15 seconds and the O_2 intervals 45 seconds; both are done at 500 V. For the other recipe the intervals last 10 seconds and 50 seconds respectively and the CHF_3 bias is only 100V. The idea behind the alternating intervals is that if the pure oxygen plasma leads to a formation of an amorphous silicon oxide film on the surface, that could be removed by the CHF_3. The test results are shown in table 4.2.

Table 4.2

Approximate Total Si-Etch Under Different Fluorine-Containing Plasmas for 6 mins

Recipe Description	Si Etch
0 % CF_4 at 500V	0 nm
10 % CF_4 at 500V	0 nm
20 % CF_4 at 500V	5 nm
33% CF_4 at 500V	25 nm
CHF_3 (15 secs, 500 V) O2 (45 secs, 500 V)	30 nm
CHF_3 (10 secs, 100 V) O2 (50 secs, 500 V)	0 nm

These values are rough estimates based on the SOI color chart explained in this section and are mainly used to determine if there is finite Si etch or not.

Based on these results we can narrow the list of acceptable plasma for the end of the Cyclotene etch. The CHF_3-O_2 interval-etch was tested first to determine how well it attacks Cyclotene. As shown in Fig. 4.13, the etch rate is extremely slow at only 70 nm of Cyclotene in 10 minutes. Because the other recipes etch Cyclotene better, this one was discarded. To test which one of the remaining two recipes is better, they were both used

to etch a sample that had 500 nm of Cyclotene left. The 10 %-CF_4 plasma successfully removed the Cyclotene and left the surface clean, while the pure oxygen plasma could not etch cleanly all the way down. The results are also shown in Fig. 4.13. An SOI sample with etched Si waveguides was exposed to a 10 %-CF_4 plasma at 500 V for 8 minutes to corroborate the SOI color chart test. As expected, no etch of the Si waveguides was observed or measured.

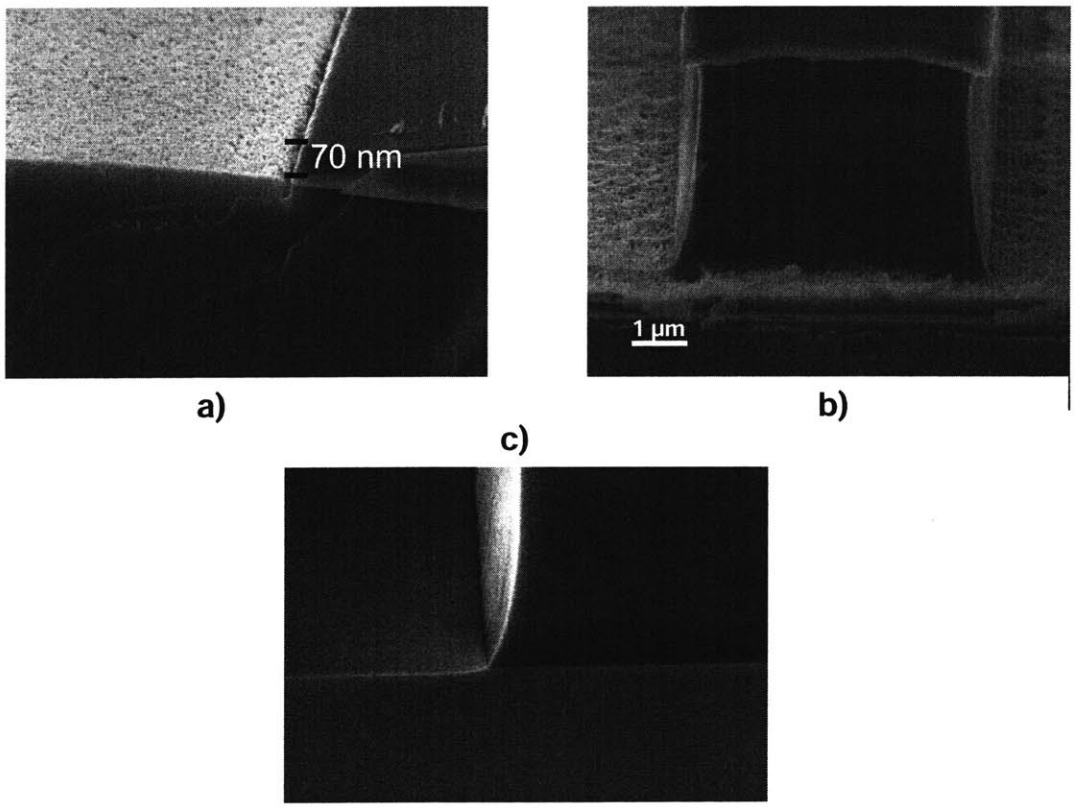

Fig. 4.13: Test results of different etch recipes considered for the second step of the Cyclotene etch. (a) The CHF_3/O_2 intervals had the worst etch rate, with only 70 nm in 10 minutes. (b) The pure oxygen plasma could not etch 500 nm cleanly. (c) Zoom-in image showing a clean surface after a 10%-CF_4 plasma etched 500 nm of Cyclotene.

We can use these results to design an etch process for the polymer waveguides. The first step uses a 33%-CF_4 plasma at 500V. It was found experimentally that the

actual etch rate of Cyclotene depends strongly on the power of the plasma. Power fluctuations are hard to control because even though O_2/CF_4 plasmas do not allow contamination of the chamber during the etch, variations in the initial condition of the chamber can result in power differences of up to 20% between etches conducted on different days. For this reason, every few minutes the etch is paused and the height of the waveguides is measured in a Dektak depth profiler. This has the added benefit that the sample is allowed to cool down between etches, preventing the polymer from reaching a high temperature. When the measured height is 3.9 μm (4.3 μm – 400 nm) we stop and switch to the lower CF_4-concentration plasma. There must be more than 240 nm of Cyclotene left at the end of the first step so that the Si waveguides are not etched. Differences in Cyclotene thickness of up to100 nm were measured across the sample area, so the process is designed to end the first step when there are 400 nm +/- 100 nm of Cyclotene left. The lower bound of 300 nm is still high enough to completely cover the Si waveguides and the upper bound of 500 nm is the maximum thickness that a 10%-CF_4 plasma was observed to remove cleanly. The second etch is stopped when the Cyclotene is completely removed, which occurs when the height of the Si waveguides is measured at 240 nm (a few minutes of over-etch are added to ensure a clean surface). This is the final step in the fabrication process for the polymer waveguides.

4.4 Assisted-Cleave

In order to accurately and uniformly cut each chip to the desired length, a diesaw is used to make the cuts along the cleave-marks. A standard process in which the diesaw generates a cut through the entire sample from the top cannot be used because the diesaw's blade would damage the waveguides and not allow good optical characterization. A process called assisted-cleave, and illustrated in Fig. 4.14, is needed.

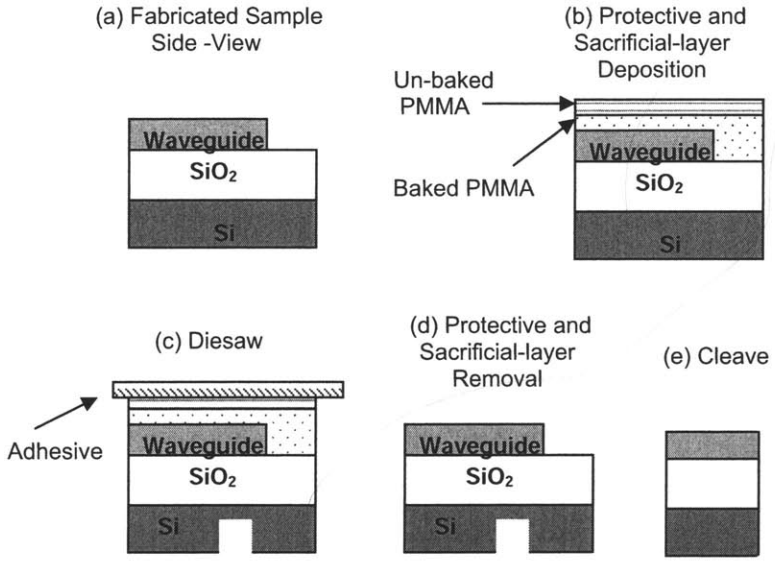

Fig. 4.14: Assisted-cleave process overview. (a) Side-view of sample after fabrication is completed. (b) A PMMA-protective-layer is spun and baked, and a PMMA-sacrificial-layer is spun on top. (c) The sample is partially cut from the back, while an adhesive tape holds the sample to the diesaw . (d) The sacrificial layer is removed, along with the adhesive, and the protective layer is removed next. (e) The sample is cleaved along the groove.

The first step of the process is to deposit two layers on top of the fabricated sample. One strong layer is needed to cover and protect the waveguides. A second layer is needed because an adhesive tape will be used to hold the top of the sample to the diesaw while cutting it from the back. A weaker layer is needed so that when dissolved,

the adhesive will be lifted off the sample. The layers have to be chosen so that the protective layer can withstand the removal of the sacrificial layer, and so that both layers can be removed in solvents that do not attack Cyclotene. PMMA was chosen for the protective layer because it can be spun to a thickness of over 4.3 μm, can be solidified by baking, and removed using N-Methyl-2-Pyrrolidone (NMP). NMP at room temperature was found not to attack cured Cyclotene. To solidify the PMMA, an oven bake at 180C for two hours was found sufficient. PMMA can also be spun on top of the solidified layer and, if left unbaked, removed by immersing in acetone. Because acetone will not dissolve the baked PMMA layer, un-baked PMMA can be used as the sacrificial layer.

Once both layers are deposited, the sample can be diesawed. First, a reference cut is made on the top of the sample at a known distance from the cleave-marks. The sample is then flipped and the reference cut is used to find the cleave-marks. Finally, a cut is made 80% through the back of the sample along the cleave-marks at the lowest possible speed. The cut is partial to prevent damage to the front of the wafer, and speed is minimized to decrease roughness along the cut. Once all the necessary cuts are made, the sacrificial layer is removed in acetone. As the sacrificial layer dissolves, the adhesive is lifted off the top of the sample. The protective layer is then removed by placing the sample in an NMP bath at room temperature until all the PMMA dissolves. The final fabrication step is to carefully cleave the sample along the grooves of the cuts.

Chapter 5

Results

Results of the first fabrication that integrates all the silicon and polymer devices is presented in Chapter 5. Unexpected problems in the fabrication occurred, which limited the extent of performance characterization. The results are presented in two sections, one that includes optical and Infra Red (IR) measurements for the devices, and one that shows micrographs of the devices after fabrication. Significant detective work was done to identify possible causes for the fabrication problems; the findings and conclusions are presented at the end of section 2.

5.1 Optical and IR Measurements

An initial set of optical measurements was done by coupling a tunable laser to the input facet of each device from an optical fiber. For a detailed description of each device please refer to Chapter 3. A photodetector at the output facet showed a strong signal only for the 10 input-output couplers connected with a short Si waveguide and for the three straight polymer waveguides. Furthermore, it showed a stronger output signal for the devices with the shortest Si-waveguide taper-length. No light was detected at the output facet of any of the remaining devices. We can arrive at three conclusions from these preliminary results: the polymer waveguides successfully confine light with small loss,

the horizontal coupler design succeeded in transferring light from the polymer to the Si and then back to the polymer, there is a higher-than-expected loss associated with the Si structures that does not allow accurate characterization of the effectiveness of the coupler.

The three straight polymer waveguides were of the same length, so it is not possible to accurately determine their loss per unit length. A GaAs IR camera, which can detect light scattering, was used to look at the waveguides from input to output facet. No scattering loss was evident in the IR images, which combined with the high-power signal picked up by the photodetector at the output facet suggests good light confinement and low attenuation in the polymer waveguides.

The light intensity picked up by the photodetector at the output of the couplers connected by a 100μm-long Si waveguide was similar to that of the straight polymer waveguides. Contrary to theoretical expectations, the intensity decreased as the length of the couplers' taper region increased. This suggests a problem either with the tapered region itself or with the Si waveguides in general. The average difference in output intensity compared to the polymer waveguides was approximately 3 dB. Given the slanted end-section of the polymer waveguides, as shown in Fig. 5.1, it is highly unlikely that light can be coupled from one polymer waveguide to the other without going through the Si. Thus, the light shown at the output facet indicates that the horizontal coupler indeed transfers the propagating modes from the polymer to the Si and back to the polymer.

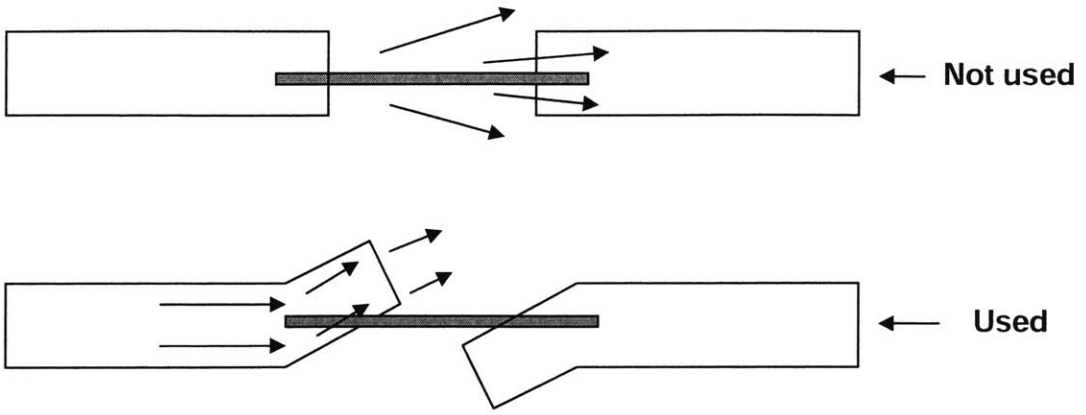

Fig. 5.1: Top view illustration of polymer waveguides connected with a short Si waveguide. In the top figure the polymer is not slanted at the end and light can be radiated from one polymer waveguide to the next without going through the Si. The slanted end-section in the bottom figure prevents this effect.

The fact that no light was detected at the output of the other coupler devices which have longer Si sections suggests a high loss associated with this layer. This was confirmed by measurements in the pure-Si "paperclips" where no light was detected at the output facet. This also confirms that the problem is with the Si waveguides in general, and not specific to the tapered region. IR images were taken for all devices and bright spots corresponding to light scattering where identified every 100 μm. This pattern corresponds to locations where stitching errors occur, and showed decreasing intensity at each stitching site. Observing light scattering at stitching sites is common, and the intensity decrease indicates that there is loss in the Si waveguides due to either scattering caused by the stitches or light absorption between the stitches. As shown in Fig. 5.2, the bright spots can be isolated and the loss can be estimated by integrating the intensity in each spot and subtracting the noise level. Measurements by M. Dahlem and calculations by A.M. Khilo give a rough estimate of 170 dB/cm power loss in the Si waveguides. This number is not consistent with the measurements from the couplers

connected with short Si sections where most of the Si is covered by the polymer. This suggests a higher loss for the exposed Si. Although the number is a rough approximation, it indicates an abnormally high loss that prevents accurate characterization of the horizontal coupler's performance (with theoretical loss of less than 1 dB/cm). It is then necessary to obtain images of the fabrication results to identify possible causes for the loss.

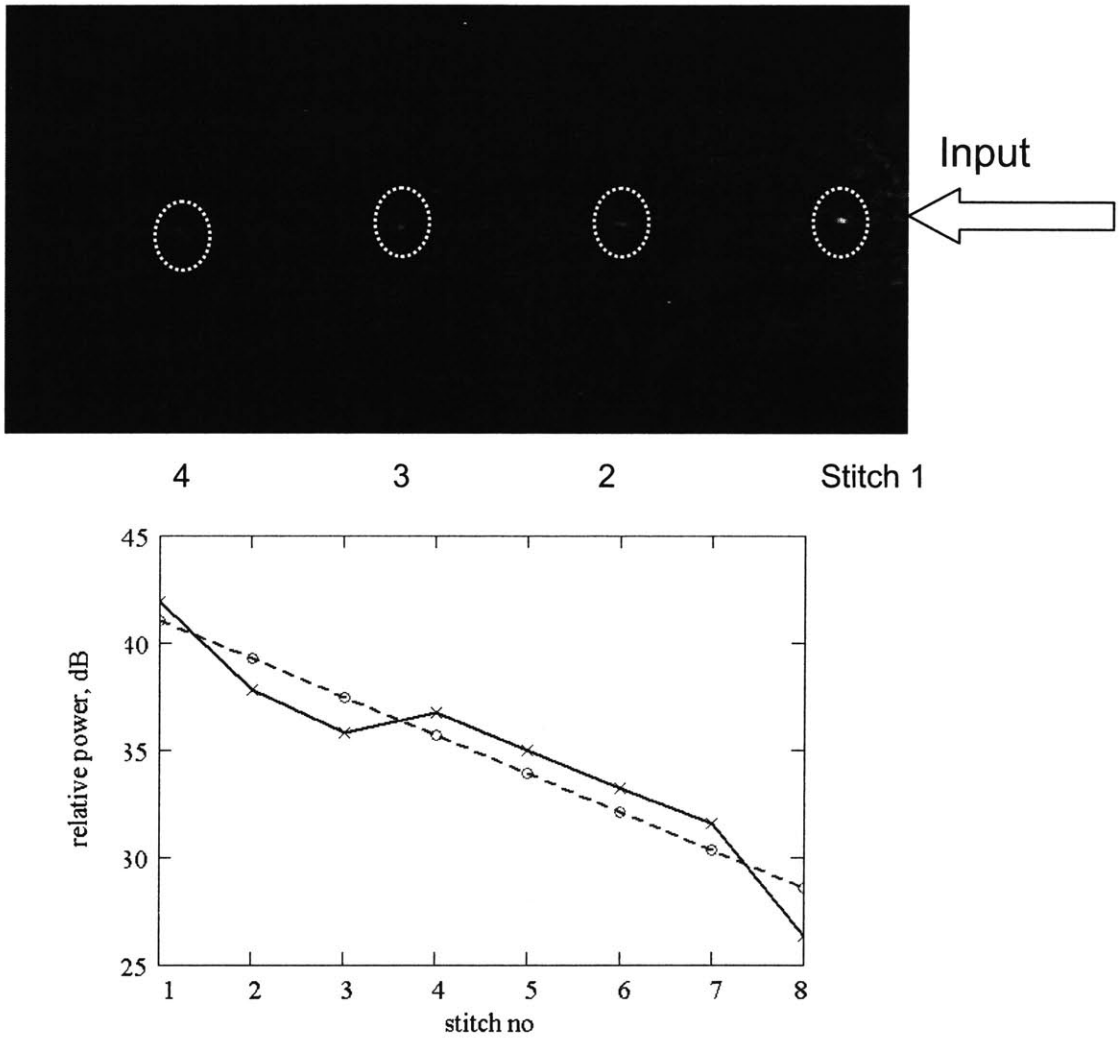

Fig. 5.2: Estimated loss in Si waveguides based on IR images. The top IR image shows the bright spots that correspond to light scattering at the stitching sites. The scattered power attenuation was estimated by isolating each spot, integrating its intensity and subtracting the noise level. The bottom graph plots the decrease in intensity across the stitching sites and gives a rough estimate of 170 dB/cm loss. [Measurements by M. Dahlem, calculations by A.M. Khilo]

5.2 Fabrication Results

Scanning-electron micrographs were taken to show the fabrication results and look for three probable causes of loss: horizontal misalignment between the polymer and Si waveguides, stitching errors, and residue present on the surface of the chip. Some of the devices are shown in Fig. 5.3 from a top view, including the connection between two horizontal couplers, the cascaded couplers and a section of the Si paperclips. The tapered region in the Si waveguides cannot be seen in these images because it is fully covered by the polymer waveguides.

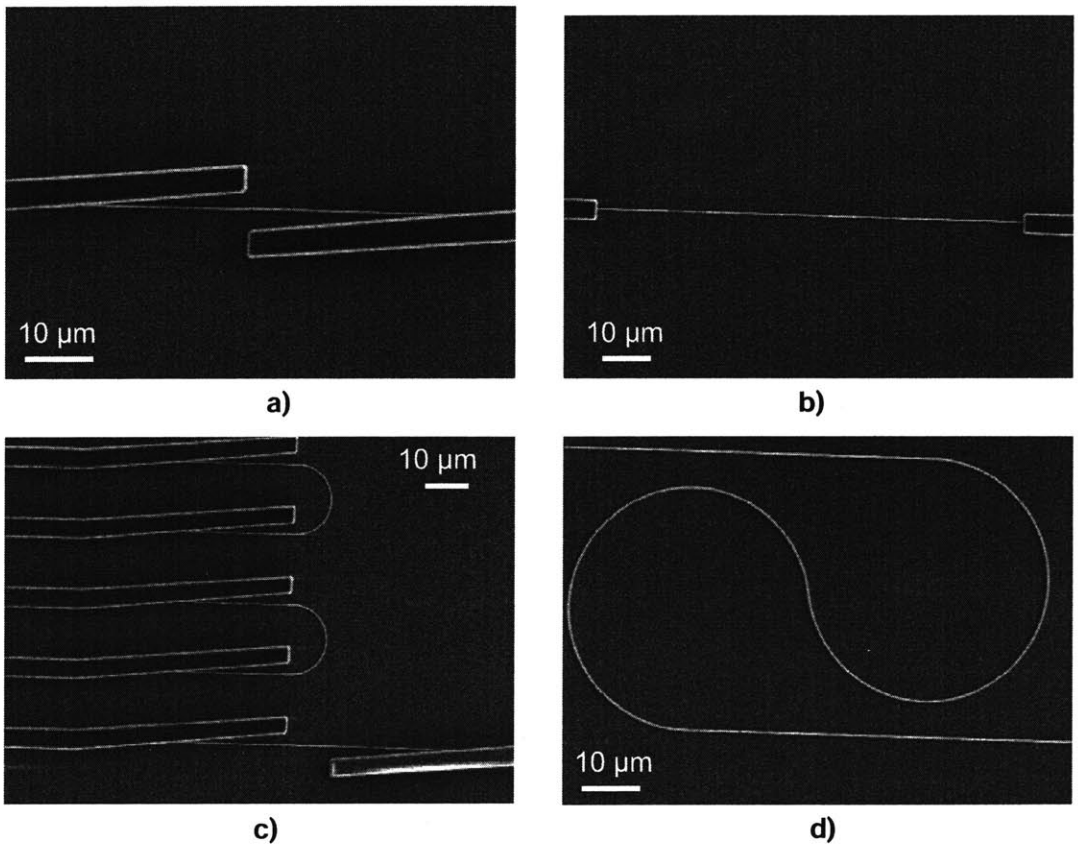

Fig. 5.3: Top-view micrographs of some of the fabricated devices. (a) Short Si waveguide connecting two horizontal couplers where the polymer end-section slants. (b) Short Si waveguide connecting two horizontal couplers in the absence of the polymer slanted section. (c) Five horizontal couplers cascaded through Si "u-turns" and connected to an output waveguide. (d) Section of the Si paperclips showing the "u-turns" and "s-bends" used in the device.

A closer look at the devices shows that the dimensions for the straight Si and polymer waveguides are, as expected, 400 nm and 5.2 μm respectively. The horizontal offset between the two sets of waveguides was measured at 0.6 μm, which meets the design specification of submicron alignment. The micrographs also show two probable causes of high loss: a large stitching error and a non-uniform residue across the chip. The SEM images are shown in Fig. 5.4. The average stitching error was 70 +/- 2 nm, and measurements ranged from 65 nm to 75 nm. The large stitchings are caused by a field calibration error prior to the e-beam exposure.

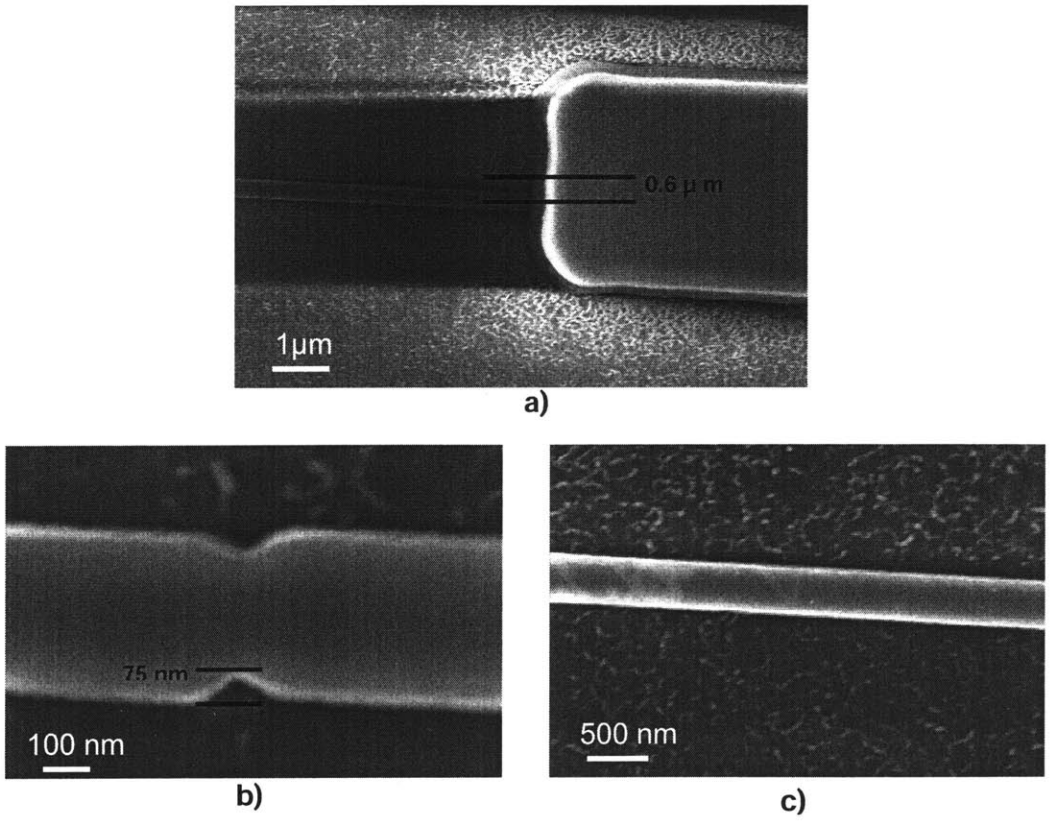

Fig. 5.4: High-magnification top-view micrographs of the fabricated devices. (a) The horizontal misalignment between the center of the polymer waveguides and the Si waveguides is 0.6 μm. An over-hang of the remaining oxide mask can be observed in this image (b) The average stitching error is 70 nm. (c) As can be seen in this image, along with the previous two, there is a non-uniform residue across the surface and on top of some of the waveguides.

Mode-matching (MM) simulations were performed to estimate how much loss can be caused by stitching error. In MM simulations the stitching is modeled as an offset between two sections of the waveguide, as shown in Fig. 5.5, and does not take into account the actual "notch" shape that was observed and which affects loss. Fig. 5.5 shows that for the average stitching error measured, MM simulation estimate a loss of less than 5 dB/cm. However, the plot also shows that larger stitching errors can account for losses of over 200 dB/cm [simulation by A.M. Khilo]. It should be noted that the performed simulation is a conservative estimate, since it does not take into account the shape of the stitch. Furthermore, the simulation estimates the effect of the stitching error on the straight section of the Si waveguides (400 nm) and it is likely that the same stitching causes much higher loss on the tapered section. Taking this into consideration, it is possible that stitching error accounts for a large part of the measured loss in the Si waveguides. One way to get a better estimate would be to perform a three-dimensional finite-difference time-domain (FDTD) simulation.

For the residue, there are two possible origins: remaining PMMA from the protective layer and unetched Cyclotene. The non-uniform appearance of the residue and the fact that it shows up on top of some of the waveguides suggests that it is not Cyclotene. A fabricated chip was immersed in NMP at 80 C to remove the residue if it was in fact PMMA. For a more detailed discussion of the process please refer to Chapter 4. After the prolonged NMP bath the residue was still present on the chip and optical measurements were similar to the previous ones. The failed NMP clean prompted the use of a more aggressive cleaning solution, piranha etch. Piranha is a sulfuric-acid:hydrogen-peroxide (3:1) wet-etch solution that removes most organic contaminants. It is also one

of the few wet etches that successfully removes fully-cured Cyclotene [14]. The idea behind this process was to remove everything from the sample except the Si waveguides and the oxide undercladding and measure again the loss in the Si. In order to completely remove the Cyclotene, four 12-minute piranha solutions were needed followed by CHF_3 RIE to remove the remaining oxide hard mask.

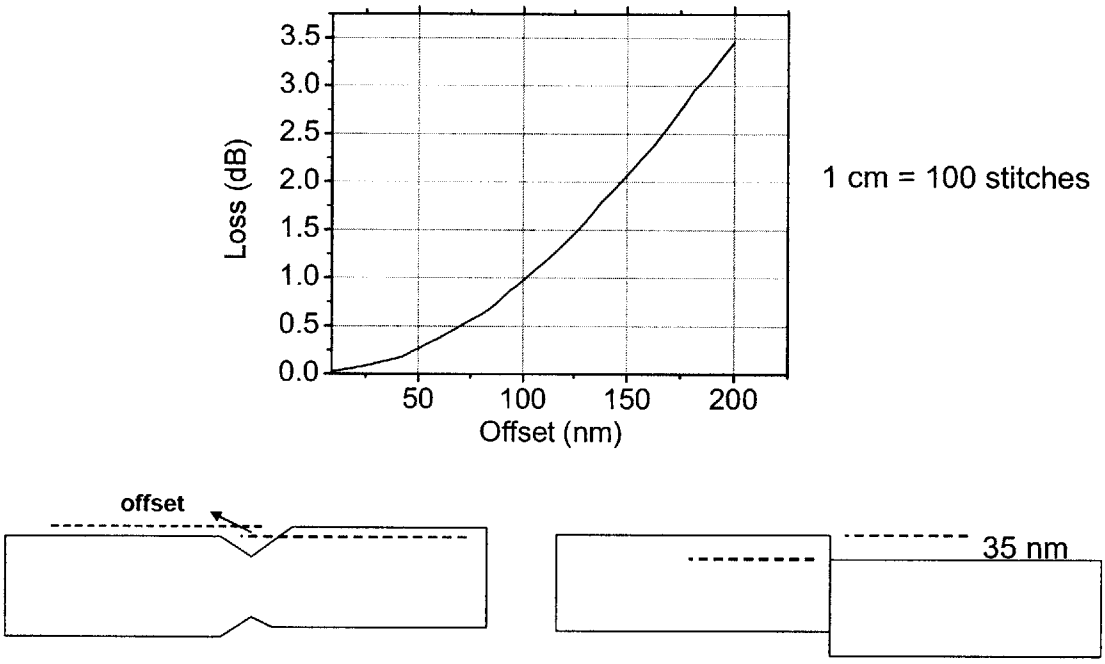

Fig. 5.5: Mode-matching simulation of stitching-error-induced power loss per stitch. As shown in the bottom illustrations, mode-matching models the stitching error as a horizontal offset. Although the simulation estimates a loss of less than 5 dB/cm for the average stitching error observed, it shows that losses of more than 200 dB/cm can be attributed to larger stitching errors. Because the performed MM simulation is a conservative estimate, it is possible that stitching error accounts for a large amount of the measured loss in the Si waveguides. [Simulation by A.M. Khilo]

Optical measurements once again showed no light in the output facet of the Si waveguides and approximations based on IR images showed only a slight reduction in power attenuation (160 dB/cm compared to an initial 170 dB/cm). These measurements, although complicated by breaks in the Si waveguides that resulted from the chip's bouncing in the piranha baths, still show high loss in the Si. Cross-sectional and top-view micrographs from the chip before and after the piranha etch are shown in Figs. 5.6 and 5.7, while Fig. 5.8 shows the breaks in the Si waveguides.

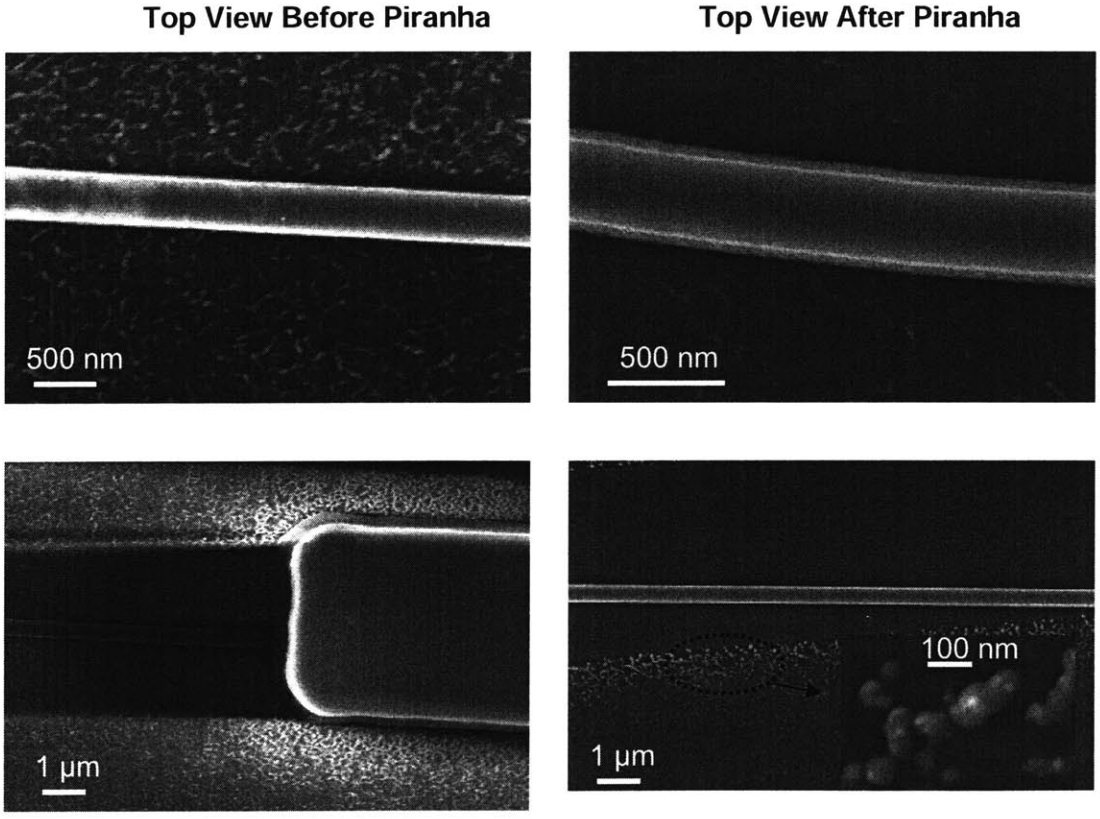

Fig. 5.6: Top-view micrographs comparing a fabricated chip before and after piranha-etch and CHF_3 RIE. The images on the right show significant reduction of residue, but some is still present. There appears to be an increase in the roughness of the Si waveguides, most likely as a consequence of the aggressive wet etch. The bottom-right micrograph includes a close-up of a narrow strip of residue that was seen at the edges of the polymer waveguides. Notice that other than this strip, the surface is clean where the polymer used to be.

Fig. 5.6 shows a significant reduction in the residue, but not full removal. The cross-sections shown in 5.7 reveal that the reduction corresponds to the removal of large material deposits on the surface of the chip. Besides material deposits, there is a rough layer on the surface prior to piranha and CHF_3 RIE which seems to correspond to a thin layer of un-etched Cyclotene. Both the deposits and the roughness can contribute to the loss measured in the Si waveguides. Images after the etches show a rough surface even after the polymer was fully removed. The new rough surface no longer appears to be a thin Cyclotene layer, but rather roughness on the oxide undercladding. It is possible that the thin layer that was previously on the surface acted as a mask during the CHF_3 etch and transferred the rough pattern to the oxide, but if the layer was Cyclotene then it should have been removed by the piranha prior to the CHF_3.

The post-fabrication work done on the chips suggests, aside from the stitching error, a problem during fabrication not associated with the actual process. Results from the micrographs are not consistent with fabrication steps discussed in Chapter 4. It is likely that that chips got contaminated during one of the fabrication steps (contamination in the HBr RIE chamber for example), and the best solution is to repeat the fabrication. If in fact there was contamination not related to the fabrication process it is highly unlikely to reappear, and a better field-calibration during e-beam exposure will reduce the stitching error. Thus, a second fabrication would allow for much more accurate characterization of the horizontal couplers.

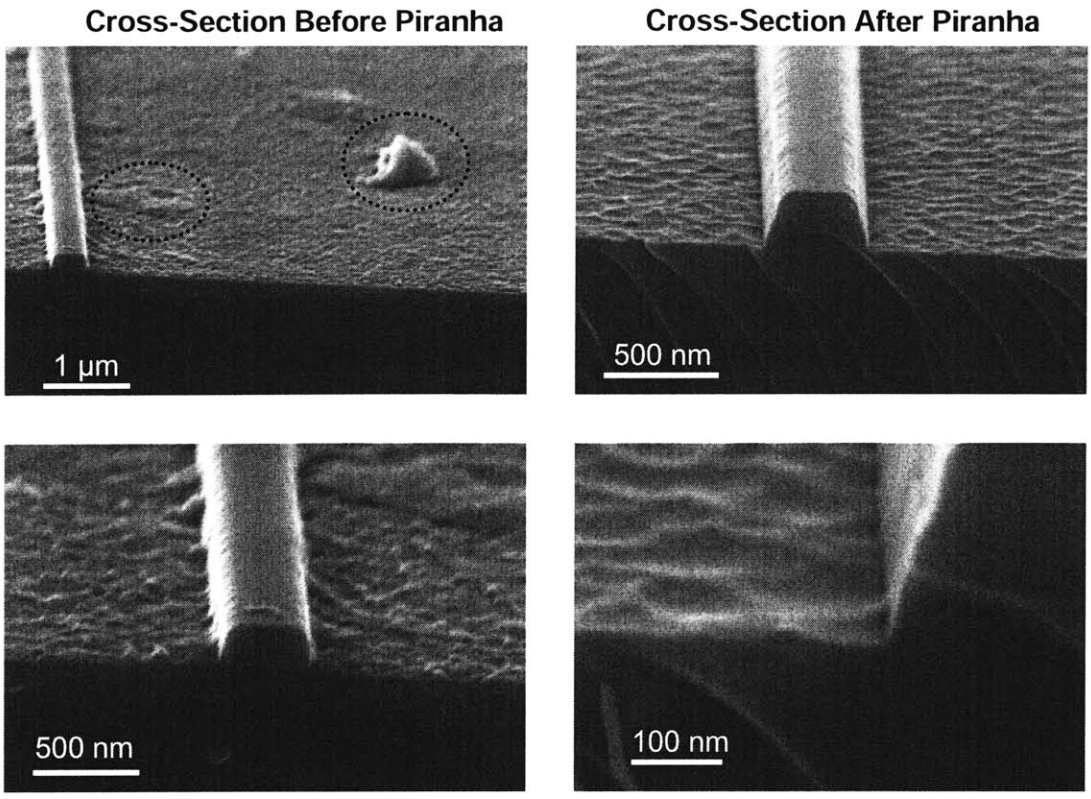

Fig. 5.7: Cross-sectional micrographs comparing a fabricated chip before and after piranha etch and CHF_3 RIE. Notice in the left images the presence of large deposits of material, and the roughness of the surface. The surface looks similar to that of a sample in which the Cyclotene has not been fully etched (see Fig. 4.13), and a closer look in the bottom-right micrograph points to the presence of a thin layer of un-etched Cyclotene. As shown in the right images, the large material deposits are no longer presents after piranha and RIE. Although the Si waveguides appear taller, indicating that there is no longer a thin Cyclotene layer on the surface, roughness is still evident.

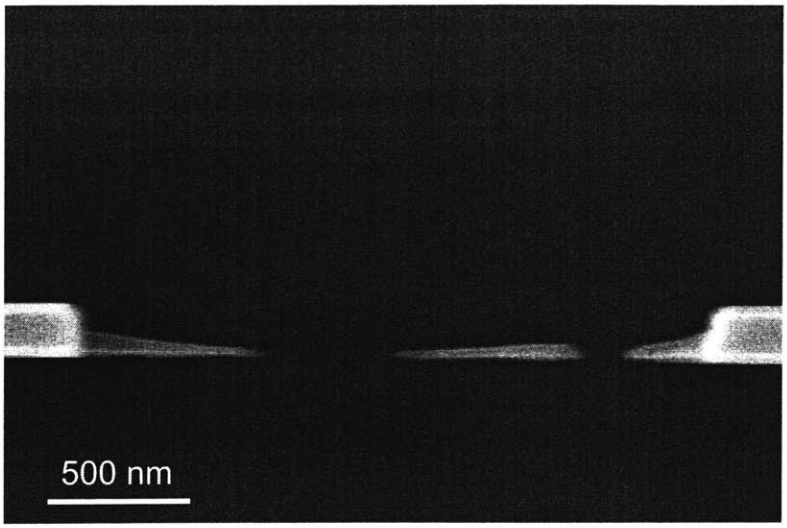

Fig. 5.8: Top-view micrograph of breaks in Si waveguides after piranha etch. The breaks become more frequent as the width of the waveguide tapers down to 50 nm. This is consistent with the theory that the breaks were caused by the bouncing of the chip in the piranha bath because thinner waveguides are more fragile.

Chapter 6

Conclusions

6.1 Summary

A process for the fabrication of optical-mode converters based on a horizontal coupler design was developed and tested. The coupler presented is multi-mode consisting of straight polymer waveguides and tapered Si waveguides, fabricated in SOI substrates. The main challenges addressed were: dimensional accuracy of Si waveguides, deposition and patterning of polymer waveguides, compatibility of polymer-waveguide fabrication process with underlying SOI substrate, and post-fabrication cleaving process.

SEBL exposure doses and writing strategies were optimized to create horizontally tapered Si waveguides with the specified dimensions and tapered profile. The deposition and thermal cure processes found to produce the best quality Cyclotene films were discussed in detail. The material selection and thickness considerations for the two etch masks needed to pattern the polymer are also presented, along with the photolithography exposure and development conditions.

The lithographic alignment requirements were discussed, and a set of marks was created to achieve sub-micron alignment. The main fabrication challenge addresses in the work was the implementation of an etch process for Cyclotene that is compatible with the underlying Si waveguides. The process is based on two O_2:CF_4 plasmas of different ratios, combined with frequent thickness monitoring through Dektak measurements. The

process was found to create low-loss polymer waveguides and negligible etching of the Si.

Finally, a cleaving process that allows good optical characterization was presented. The process uses two protective layers of PMMA and partial diesaw cuts through the backside of the sample to cleave the chips without damaging the structures. One fabrication was performed to integrate all the Si and polymer devices. Results from the fabrication showed that the horizontal coupler transfers the propagating modes from an optical fiber to the polymer waveguides and then to the Si waveguides. An unexpected high loss in the Si waveguides prevented accurate characterization of the coupling efficiency. It is believed that sample contamination during the fabrication process and large stitching errors are responsible for the unexpected loss.

6.2 Future Work

The first issue to be addressed in the future is the completion of a second fabrication. Because the problems encountered in the first one are not believed to be related to the actual fabrication process, it is highly unlikely that they will appear a second time. Before starting that fabrication it would be worthwhile to rule-out the SOI substrate as the root of the problem. A simple test consisting of straight Si waveguides can be performed to confirm that there are no problems with the Si prior to the fabrication. A second test should be performed to fabricate a few sets of straight Si and polymer waveguides to make sure the integration of the two does not affect the optical transparency of the Si.

If a second fabrication is successful, the next step is to develop a process for the fabrication of a single-mode optical converter that can couple light from a standard

optical fiber. The single-mode coupler will also be fabricated on an SOI substrates and consist of Si and polymer waveguides. While most of the fabrication techniques presented in the current work can be used for the new coupler, several key differences complicate the process.

The single-mode horizontal coupler requires an over-cladding above the Si waveguides. In the first version of the coupler, the section of the Si waveguides that did not overlap with the polymer was surrounded by air. This air-clad situation is impractical for real world applications as the performance of air-clad devices may strongly deteriorate with time due to environmental contamination. Adding an over-cladding layer makes the coupler more useful for real-world applications and compatible with other photonic devices such as OADMs also fabricated at MIT. As discussed in Chapter 4, the use of an overcladding makes the removal of the HSQ mask more difficult. It is possible to use HSQ as an overcladding, in which case the etch mask does not have to removed. In this case, the coupler's performance would strongly depend on the refractive index of HSQ which is affected by the annealing conditions used to harden it.

A second major difference in the single-mode coupler is the step-like patterning of the polymer waveguides. A cross-section and top-view schematic of the single-mode coupler is shown in Fig. 6.1. In this geometry, called a rib waveguide, higher-order modes leak out over short distances and only the fundamental mode propagates through the entire length. This improves coupling from the large standard fiber to the polymer waveguide.

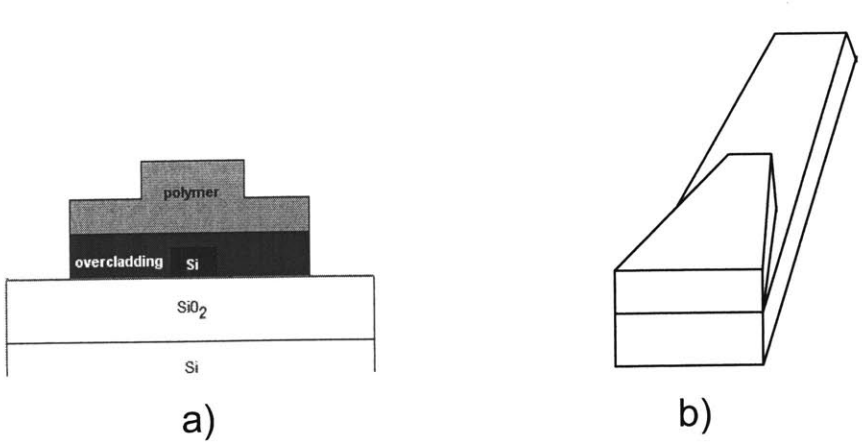

Fig. 6.1: Cross-section and top-view schematics of the single-mode horizontal coupler. (a) The cross-section shows the overcladding around the Si waveguides. (b) This top view shows the shape of the polymer waveguide, which requires two lithographic steps.

As shown in Fig. 6.1 (b), the single-mode coupler requires two lithographic steps for the patterning of the polymer waveguides. An initial layer of polymer has to be deposited above the SOI and the desired pattern must be aligned to the underlying Si structures. That first polymer layer must be etched completely before depositing the second layer. A second photo mask has to be designed to pattern the second layer of polymer and align it to the first. It is important to determine the required alignment accuracy between both layers before designing the second photo mask. It is also important to have good step-coverage in the deposition of the second polymer layer.

The dimensions of the polymer waveguides also represent a challenge in the fabrication of the single-mode coupler. In particular, a height of approximately 10 μm is needed for the polymer to match the dimensions of the single-mode optical fibers. Cyclotene can be used for in thick layers, but it would be recommended to use a version thicker than 3022-35 so that fewer coats are needed. Two coats of 3022-46 or a single coat of 3022-57 can be used to achieve the desired height. Decreasing the number of

coats not only saves fabrication time and cost, but also reduces adhesion problems. Although Cyclotene can still be used in the fabrication of the single-mode coupler, there are two main reasons why another polymer would work best. First, the required thickness, and multiple lithographic and etch steps makes compatibility with the Si layer much harder. A polymer than can be etched in a pure oxygen plasma would significantly simplify the fabrication process. Second, a lower-refractive-index polymer would better match the refractive index of the optical fiber and thus improve performance.

A list of polymers suitable for optical applications can be found in [15]. The selected polymer must have a refractive index between 1.45 and 1.55, have low absorption loss, be dry etchable or photosensitive, and reach a thickness of 10μm. Halogenated acrylates, deuterated acrylates, and fluorinated polymides are some of the polymer types available.

Bibliography

[1] D. Bertsekas, R. Gallager, *Data Networks*, Prentice Hall, Upper Saddle River NJ, (1991).

[2] T. Barwicz, "Accurate nanofabrication techniques for high-index-contrast microphotonic devices," *Ph.D. Thesis*, Massachusetts Institute of Technology, Cambridge MA, (2005).

[3] T. Montalbo, "Fiber to waveguide couplers for silicon photonics," *S.M. Thesis*, Massachusetts Institute of Technology, Cambridge MA, (2004).

[4] R. Ramaswami, K.N. Sivarajan, *Optical Networks: a practical perspective*, Morgan Kaufmann Publishers, San Francisco CA, (2001).

[5] B. Mukherjee, *Optical Communication Networks*, McGraw-Hill, New York NY, (1997).

[6] W. Bogaerts, P. Dumon, P. Jaenen, J. Wouters, S. Beckx, V. Wiaux, D.V. Thourhout, D. Taillaert, B. Luyssaert and R. Baets, "Silicon-on-insulator nanophotonics," *Proceedings of SPIE*, International Society for Optical Engineering, vol. 5956, (2005).

[7] J.D. Plummer, M.D. Deal, P.B. Griffin, *Silicon VLSI Technology: Fundamentals, Practice and Modeling*, Prentice Hall, Upper Saddle River NJ, (2000).

[8] E.E. Moon, P.N. Everett, and H.I. Smith, "Immunity to signal degradation by overlayers using a novel spatial-phase-matching alignment system," *J. Vac. Sci. Technol. B* **13**, pp. 2648-2652, (1995).

[9] M. Yu, H. Shan, A. Taylor, "Low k Film Etch in Applied Materials eMxP+ Chamber," *Proceedings of Plasma Deposition and Treatment of Polymers*, MRS, vol. 544, pp. 203-207, (1998).

[10] *Processing Procedures for CYCLOTENE 3000 Series Dry Etch Resins* (Dow Chemical Company, Midland, MI, 2005); http://www.dow.com/cyclotene/prod/302235.htm

[11] M. Schier, "Reactive ion etching of benzocyclobutene using a silicon nitride dielectric etch mask," *J. Electrochem. Soc.*, vol. 142, no. 9, pp. 3238-3240, (1995).

[12] D.L. Flamm, V.M. Donnelly, and J.A. Mucha, "The reaction of fluorine atoms with silicon," *J. Appl. Phys.* **52**, pp.3633-3639, (1981).

[13] J. Henrie, S. Kellis, S. Schultz, and A. Hawkins, "Electronic color charts for dielectric films on silicon," *Optics Express*, vol. 12, no. 7, pp. 1464-1469, (2004).

[14] *Rework Procedures for CYCLOTENE 3000 Series and 4000 Series Resins* (Dow Chemical Company, Midland, MI, 2005); http://www.dow.com/cyclotene/prod/302235.htm

[15] H. Ma, A.K.-Y. Yen, L.R. Dalton, "Polymer-Based Optical Waveguides: Materials, Processing, and Devices," *Advanced Materials*, vol. 14, no. 19, pp. 1339-1365, (2002).

[16] L.Liao, "Low loss polysilicon waveguides for silicon photonics," *M.S. Thesis*, Massachusetts Institute of Technology, Cambridge MA, (1995).

[17] J.A. Kong, *Electromagnetic Wave Theory*, EMW Publishing, Cambridge MA, (2005).

[18] M. B. Frish, J. A. Fijol, E. E. Fike, S. A. Jacobson, P. B. Keating, W. J. Kessler, C. Bozler, M. Fritze, C. Keast, J. Knecht, R. Williamson, and C. Manolatou, "Coupling of single-mode fibers to planar Si waveguides using vertically tapered mode converters," in *Integrated Photonics Res., 2002 Tech. Dig.*, Vancouver, BC, Canada, pp. IFB2-1–IFB2-3, (2002).

[19] S. Itabashi, H. Fukuda, T. Tsuchizawa, T. Watanabe, and K. Yamada, "Silicon wire waveguides and silicon microphotonic devices," *NTT Technical Review*, vol. 4, no. 3, pp. 48-56, (2006).

[20] D.P. Mancini, K.A. Gehoski, E. Ainley, K.J. Nordquist, D.J. Resnick, T.C. Bailey, S.V. Sreenivasan, J.G. Ekerdt, and C.G. Willson, "Hydrogen silsesquioxane for direct electron-beam patterning of step and flash imprint lithography templates," *J. Vac. Sci. Technol. B* **20**, pp. 2896-2901, (2002).

[21] M. Morita, T. Ohmi, E. Hasegawa, M. Kawakami, and M. Ohwada, "Growth of native oxide on a silicon surface," *J. Appl. Phys.* **68**, pp.1272-1281, (1990).

[22] S. Meltzer, D. Mandler, "Study of silicon etching in HBr solutions using a scanning electrochemical microscope," *J. Chem. Soc., Faraday Trans.* **91**, pp. 1019-1024, (1995).

[23] *Cure and Oxidation Measurements for Cyclotene Advances Electronic Resins* (Dow Chemical Company, Midland, MI, 2005); http://www.dow.com/cyclotene/prod/302235.htm

[24] P. Chinoy, "Reactive Ion Etching of Benzocyclobutene Polymer Films," *IEEE Transactions on Components, Packaging, and Manufacturing Technology*, vol. 20, no. 3, pp. 197-206, (1997).

[25] W.R. Buchwald, K. Vaccaro, "Sidewall profile of thick benzocyclobutene reactively ion etched in CF_4O_2 plasmas," *J. Vac. Sci. Technol. B* **23**, pp. 51-56, (2005).

[26] *NANO™ PMMA and Copolymer* (MicroChem, Newton, MA, 2001); http://www.microchem.com/products/pmma.htm

[27] T. Barwicz, M.A. Popovic, P.T. Rakich, M.R. Watts, H.A. Haus, E.P. Ippen, and H.I. Smith, "Microring-resonator-based add-drop filters in SiN: fabrication and analysis," *Optics Express*, vol. 12, no. 7, pp. 1437-1442, (2004).