Design, Fabrication and Measurement of Integrated Bragg Grating Optical Filters

by

Thomas E. Murphy

B.S. Electrical Engineering, Rice University (1994)
B.A. Physics, Rice University (1994)
S.M. Electrical Engineering, Massachusetts Institute of Technology (1997)

Submitted to the Department of Electrical Engineering and Computer Science in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February, 2001

Author . . . .
Department of Electrical Engineering and Computer Science
November 6, 2000

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

Accepted by . . . .
Arthur C. Smith
Chairman, Deparment Committee on Graduate Students

Design, Fabrication and Measurement of Integrated Bragg Grating Optical Filters

by

Thomas E. Murphy

Submitted to the Department of Electrical Engineering and Computer Science on
November 6, 2000, in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Abstract

The subject of this thesis is the design, analysis, fabrication, and characterization of integrated Bragg grating optical filters. We begin by describing the design and analysis of three essential building blocks needed for integrated Bragg grating filters: waveguides, directional couplers and Bragg gratings. Next, we describe and implement a flexible fabrication methodology for building integrated Bragg gratings filters in two important material systems: doped-glass channel waveguides and silicon-on-insulator ridge waveguides. Finally, we evaluate the performance of the fabricated directional couplers and Bragg gratings, comparing experimental results with theoretical predictions.

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

Acknowledgments

First and foremost, I would like to thank Prof. Hank Smith who supported and advised me throughout my time MIT. The goal of a Ph.D. program is not only to accomplish specific research goals but also to train one as a scientist. More than anyone else, Hank has contributed to this second goal. He has given me a tremendous amount of freedom to set my own goals and seek my own solutions to problems, but at the same time he has always been approachable and encouraging when I got stuck. Perhaps one of the most valuable skills I have acquired (hopefully) under Hank's tutelage is the ability to write clearly and present my work effectively. To the extent that the reader finds this work clear and understandable, I am deeply indebted to Hank.

For many professors, the role of academic adviser unfortunately entails little more than signing forms and reminding graduate students of upcoming deadlines. I have been extremely fortunate to have an academic adviser who has done much more than this. Terry Orlando has been a great source of guidance and advice, both professional and personal. He has gone out of his way to make himself available outside of the usual biannual registration-day meetings. I would like to thank Terry for always taking the time to listen to me, for giving me honest and unbiased advice, and for encouraging me at a time when I really needed encouragement.

In the spring of 1999, I had the unique opportunity to work as a teaching assistant with Prof. Greg Wornell. This semester proved to be one of the most rewarding and stimulating during my time as a graduate student. I'm grateful to Greg for giving me ample teaching responsibilities that I wouldn't have found in other courses.

I have benefited from many stimulating conversations with Prof. Hermann Haus, and I have great esteem for his unrivaled enthusiasm and physical insight. I would also like to acknowledge Dr. Brent Little, who provided the original underlying theory for my work on insensitive couplers. Brent has also been a valuable source of practical knowledge on optical physics and electromagnetics. I'm grateful to Prof. Erich Ippen, Prof. Lionel Kimerling, and Anu Agarwal who agreed to serve on my thesis committee.

MIT has been a wonderful place to conduct research, mostly because of the talented students who have worked alongside me. I am especially indebted to the members of the optics subgroup, who directly contributed to this work. Jay Damask was one of my first mentors at MIT, and the person who first introduced me to the world integrated optics. Much of my intuition and a great deal of my research style can be traced back to Jay. I couldn't possibly enumerate the ways in which Mike Lim has helped me out in the lab; many of the lithographic feats presented in this work could not have been accomplished without Mike's helpful advice and assistance. Juan Ferrera was for years the resident electron beam lithography expert in our lab, and he is responsible for developing the technique of adding alignment marks to interference-lithography-generated x-ray masks. He also provided invaluable help with the interference lithography system. Additionally, Juan taught me much of what I know about UNIX network administration. I'm grateful to Vince Wong, who was very generous with his time even when he was extremely busy.

Todd Hastings has been exceptionally helpful with electron-beam lithography and metrology. When I have new ideas, Todd is usually one of the first people I run them by. Jalal Khan has been a great officemate, and a valuable source of advice on optics modeling. Oliver Stolz worked very diligently with me for several months to investigate birefringence in silicon waveguides. I truly enjoyed working with him on this project, and my own research has benefited directly from his careful measurements and calculations.

Outside of the optics team, I would like to thank the other students and staff of the NanoStructures Lab. Maya Farhoud has been a wonderful officemate, and more importantly a great friend. The past five years would not have been nearly as much fun without her companionship. Mark Mondol performed much of the electron-beam lithography for me, and more generally he has also been an outstanding advocate for the students in our group. I have truly enjoyed working with Mark, and I will miss his sense of humor and his exceptional patience. Jim Daley deserves almost all the credit for keeping the lab up and running. Jim Carter has been a very valuable source of technical assistance in almost all areas of laboratory work. I would like to thank Mike Walsh, who developed the Lloyd's mirror lithography system, and spent the time to teach me how to use it. Mike has also contributed very valuable software which I used to model the reflection from multilayer stacks. Minghao Qi has greatly assisted me with the tasks of network administration, and Dario Gil has helped me out by taking over the job of maintaining our website. I would also like to thank Tim Savas, Keith Jackson, James Goodberlet, and Dave Carter for their helpful advice. I have also enjoyed many stimulating conversations about electromagnetic simulation with Rajesh Menon. Mark Schattenburg has been a tremendous source of advice on topics ranging from interference lithography to x-ray absorption. Cindy Lewis also deserves recognition for handling all of the administrative tasks of the group.

Several individuals outside of my research group have also helped to make my job at MIT easier. Desmond Lim deserves credit for teaching me the art of polishing waveguide facets. I've enjoyed sharing both lab space and ideas with J. P. Laine. Dave Foss, the local computer and networking guru, has been a tremendously helpful resource when we've had computer failures. Peggy Carney helped to straighten out my stipend on numerous occasions during my first three years as a fellowship student. Marilyn Pierce has helped with everything from selecting an oral exam comittee to registering for courses. I'd like to thank the folks over at the ATIC lab for all of their useful advice on more effectively using speech recognition software. I'm grateful to Alice White and Ed Laskowski at Bell Labs for all of the technical advice they've give me on waveguide technology, and for the dozens of glass waveguide substrates that they generously provided for this work.

Erik Thoen shared an apartment with me for five of my years at MIT, and he's been one of my most trusted friends. I'm grateful to Erik for helping me to always see the lighter side of life. Ii-Lun Chen has patiently stood by my side through some of the most difficult and strenuous times of my Ph.D. work. She has been a deep source of warmth and inspiration in my life, and a constant source of moral support.

I would like to thank the National Science Foundation for supporting me during my first three years as a graduate student, and DARPA for supporting my work as a research assistant in the latter years. I'm grateful to the MIT-Lemelson foundation and Unisphere

Corp., for recognizing and supporting our accomplishments.

Finally, I would like to acknowledge my family for their tireless love and support over the years. A special word of thanks goes to my aunt Roseanne Cook: having made it through her own Ph.D. program, her words of encouragement were especially understanding and empathetic. I would like to thank my parents, John and Josephine Murphy, for working hard to provide me with the best educational opportunities, for supporting all of my decisions, and for serving as the shining examples of how to live a rewarding and meaningful life.

Contents

Chapter 1

Introduction

During the past decade, the world has seen an explosive growth in optical telecommunications, fueled in part by the rapid expansion of the Internet. Not only are optical telecommunications systems constantly improving in their performance and capacity [1, 2], but the deployment of optical systems is spreading deeper into the consumer market. Ten years ago, optical systems were primarily used in point-to-point long distance links [3, 4]. In the future, fiber-optic networks will be routed directly into neighborhoods, households, and even to the back of each computer [5, 6]. In the more distant future, it is possible that even the signals bouncing between the different components inside the computer will be transmitted and received optically [7]. As optical fiber gradually replaces copper cables, it will become necessary for many of the electronic network components to be replaced by equivalent optical components: splitters, filters, routers, and switches.

In order for these optical components to be compact, manufacturable, low-cost, and integratable, it is highly desirable that they be fabricated on a planar surface. The integrated circuit revolution of the 1960s (which continues today) clearly demonstrated the tremendous potential afforded by planar lithographic techniques. Bragg gratings offer one possible solution for constructing integrated optical filters [8].

This thesis concerns the design, construction, and measurement of integrated optical filters based on Bragg gratings. The general objective of this thesis is to better understand how Bragg gratings work, to develop a flexible fabrication scheme for building Bragg gratings filters, and to evaluate the performance of fabricated devices.

Figure 1.1a depicts the structure of an integrated Bragg grating filter. A fine-period corrugation etched into the surface of an otherwise uniform waveguide creates a coupling between the forward- and backward-traveling light in the structure. The Bragg grating is analogous to the dielectric stack mirror depicted in Fig. 1.1b. The grating reflects light in a narrow wavelength range, centered at the so-called Bragg wavelength. As such, the Bragg grating forms a convenient implementation of an integrated optical bandpass filter.

Fiber Bragg gratings are widely used in optical telecommunications systems, for applications ranging from dispersion compensation to add / drop filtering. The integrated Bragg gratings considered here offer several advantages over their fiber counterparts. First, the integrated Bragg gratings described in this work are formed by physically corrugating a waveguide, and therefore they do not rely upon a photorefractive index change. This allows us to build Bragg gratings in materials which are not photorefractive (e.g. Si or InP), and it potentially allows stronger gratings to be constructed since the grating strength is not limited by the photorefractive effect. Second, the integrated Bragg gratings can be made smaller, and packed closer together than fiber-optic devices. Third, the planar fabrication process gives better control over the device dimensions. For example, the beginning and end of the Bragg grating can be sharply delineated rather than continuously tapered, abrupt phase shifts can be introduced at any point in the grating, and precise period control can be achieved – the integrated Bragg grating can be engineered on a tooth-by-tooth basis. Finally, multiple levels of lithography can be combined, with precise nano-alignment between them, allowing the Bragg gratings to be integrated with couplers, splitters, and other electronic or photonic components.

Figure 1.2 illustrates schematically the hierarchy of integrated optical devices which will be discussed in this thesis, beginning with the simplest structure: the integrated waveguide.

Figure 1.2b depicts a slightly more complicated device: the integrated directional coupler, which is designed to transfer power from one waveguide to another. Figure 1.2c depicts a more advanced version of the integrated directional coupler, which provides wavelength-, and polarization-insensitive performance.

Figure 1.2d depicts the simplest form of an integrated Bragg grating filter. In this structure, the filtered signal is reflected back into the input port of the device. Figure 1.2e depicts a more sophisticated Bragg grating filter, in which the grating is *apodized*, or windowed in order to provide a more optimal spectral response. One drawback of this topology is that further processing is required to separate the reflected filtered signal from the input signal. This can be accomplished using an optical circulator [9], but the circulator is an expensive device which cannot be easily integrated.

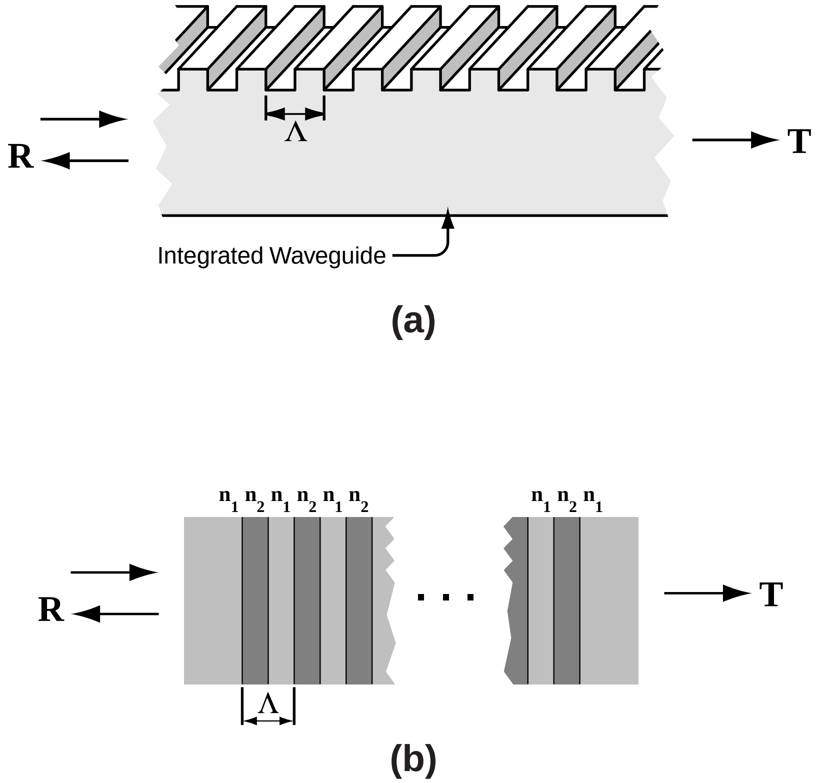

Figure 1.1: (a) Diagram of an integrated optical Bragg grating. The fine-period corrugation introduces a coupling between the forward and backward traveling modes of the waveguide. (b) The dielectric stack mirror depicted here is analogous to the integrated Bragg grating. (figs/1/grating-dielectric-stack.eps)

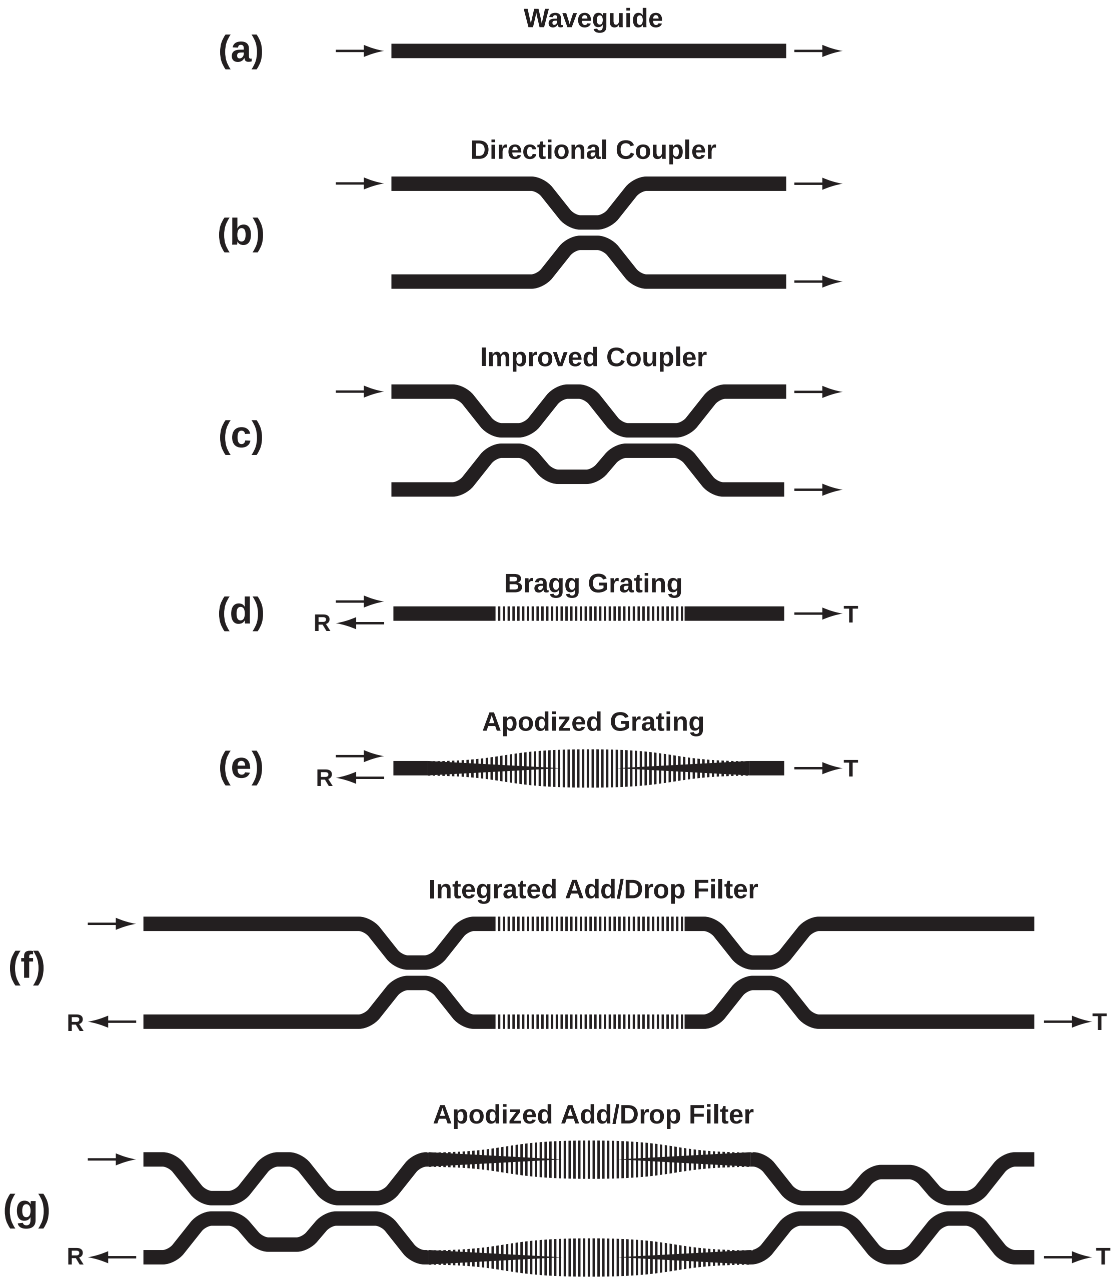

Figure 1.2: A hierarchy illustrating the type of devices which will be considered in this work, ranging from simple waveguides to complex combinations of directional couplers and gratings. (figs/1/device-types.eps)

By combining Bragg gratings and directional couplers, as depicted in Fig. 1.2f, it is possible to construct an integrated device which separates the reflected signal from the incident signal. This device can be further improved by replacing the gratings with higher-performance apodized gratings, and replacing the directional couplers with broadband polarization-insensitive couplers, as depicted in Fig. 1.2g.

Designing, building, and testing all of the structures depicted in Fig. 1.2 is an ambitious task, some of which will be carried out by future students. It is my hope that this thesis will lay the groundwork for this effort. Specifically, the analytical techniques described in this thesis will cover all of the devices depicted in Fig. 1.2. Additionally, the fabrication techniques described should provide the basic tools for constructing any of the devices depicted in Fig. 1.2, in more than one material system. Finally, we will describe completed measurements of waveguides, directional couplers, and integrated Bragg gratings of the type depicted in Fig. 1.2a-d.

This thesis is separated into three principal chapters describing respectively the design, fabrication, and measurement of Bragg gratings filters.

Chapter 2 will detail the design and analysis of waveguides, directional couplers and Bragg gratings filters. The purpose of this segment of the work will be to (1) illustrate the types of filters that can be constructed using Bragg gratings, and (2) describe how such grating filters can be designed, specifically how the waveguide and grating geometry should be selected in order to achieve a desired spectral response.

Chapter 3 will describe the development and implementation of a flexible fabrication technique for building Bragg gratings on integrated optical waveguides. The principal contribution from this portion of the work will be to identify the critical fabrication challenges presented by Bragg grating structures, and to develop fabrication techniques specifically designed to address these challenges.

Finally, in Chapter 4, measurements of integrated waveguides, directional couplers, and Bragg gratings are described. By comparing the spectral response with theoretical predictions, we will assess the device performance and the integrity of the fabrication techniques utilized.

Chapter 2

Theory and Analysis

This chapter is devoted to a theoretical analysis of integrated waveguides and Bragg gratings. There are entire textbooks written on this subject of integrated waveguides [10, 11, 12, 13], and this work is not intended to replace those excellent sources. Instead, this chapter is intended to provide a practical and relatively comprehensive summary of the theoretical and numerical techniques which are necessary for designing and building integrated Bragg grating devices.

We will begin in Section 2.1 by deriving the basic equations which describe the eigenmodes of dielectric waveguides. Since most integrated optical waveguides of interest have modal solutions which cannot be expressed analytically, we will describe flexible numerical techniques for computing the eigenmodes of integrated waveguides.

In Section 2.2, we turn to the topic of coupling between proximate waveguides. This section will describe how to accurately and efficiently model the transfer of power which occurs when two waveguides are brought close together.

In Section 2.3, we extend the coupled-mode theory to model the interaction between a forward-propagating mode and a backward-propagating mode in the presence of a Bragg grating. Included here is a description of a technique for modeling non-uniform gratings, including apodized gratings and chirped gratings.

The final portion of the chapter describes the transfer matrix method, a powerful technique which allows one to model arbitrary sequences of gratings and couplers by simply multiplying the transfer matrices of each constituent segment.

2.1 Modal Analysis of Waveguides

The dielectric waveguide is the most essential connective element in integrated optics: the waveguide is to optics what the wire is to electrical circuits. All of the theory presented in the remainder of this chapter is built up from an analysis of a simple dielectric waveguide. Therefore, this first portion of the chapter describes methods for computing the electromagnetic modes and propagation constants of dielectric waveguides.

2.1.1 Eigenmode Equations for Dielectric Waveguides

Loosely speaking, a dielectric waveguide is formed when a region with high index of refraction is embedded in (or surrounded by) a region of relatively lower index of refraction. Under these conditions, light can be confined in the central region by total internal reflection at the boundary between the high and low index materials.

Waveguides come in many shapes and sizes, but any dielectric waveguide can be mathematically described by a refractive index profile (often simply called the “index profile”), $n(x, y)$. The index profile is related to the dielectric constant ϵ by:

$$\epsilon(x, y) = \epsilon_0 n^2(x, y), \tag{2.1}$$

where ϵ_0 is the permittivity of free space. In this work, we will assume that the refractive index profile is real everywhere. Materials which have gain or loss can easily be modeled by adding an imaginary component to the refractive index profile. As indicated in Fig. 2.1, we have chosen to orient our coordinate axes such that the waveguide points in the z direction, and therefore the index profile depends only upon the two transverse coordinates x and y, or equivalently upon r and ϕ. The index profile $n(x, y)$ can be a piecewise-constant function, as depicted in Fig. 2.1, or it can be a smoothly-varying function in the x-y plane. Throughout this thesis we restrict our attention to dielectric waveguides, i.e. we assume that the materials comprising the waveguide are non-magnetic:

$$\mu(x, y) = \mu_0 \quad . \tag{2.2}$$

The eigenmodes of an optical waveguide are found by applying Maxwell’s Equations, with appropriate boundary conditions, to the index profile specified by Eq. 2.1.

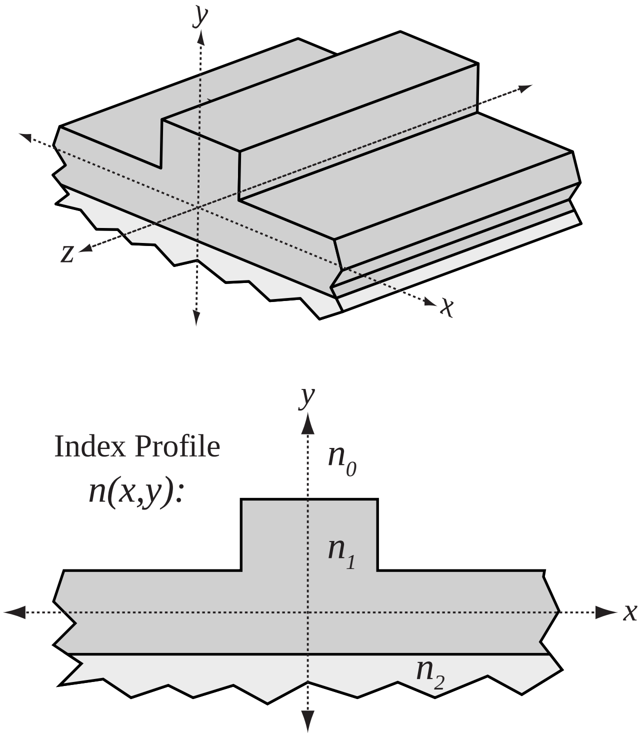

Figure 2.1: Schematic diagram of an optical waveguide. The waveguide is described by a refractive index profile $n(x, y)$. The coordinate axes have been oriented such that the waveguide points in the z-direction. In this example, the waveguide is comprised of homogeneous regions such that $n(x, y)$ is piecewise-constant. Other waveguides have smoothly-varying index profiles. (figs/2/waveguide-diagram.eps)

Maxwells equations for a dielectric waveguide are expressed as: [1]

$$\nabla \times \mathbf{E} = -\mu_0 \frac{\partial}{\partial t}\mathbf{H} \tag{2.3}$$

$$\nabla \times \mathbf{H} = n^2 \epsilon_0 \frac{\partial}{\partial t}\mathbf{E} \tag{2.4}$$

$$\nabla \cdot (n^2 \mathbf{E}) = 0 \tag{2.5}$$

$$\nabla \cdot (\mathbf{H}) = 0 \tag{2.6}$$

The above equations govern the electric and magnetic fields in an optical waveguide with no current sources or free charge ($\mathbf{J} = \sigma = 0$). The boundary conditions which must be satisfied at the interface between two different dielectric materials, designated 1 and 2, are summarized below:

$$\begin{gathered}(\mathbf{E}_1 - \mathbf{E}_2) \times \hat{\mathbf{n}} = \mathbf{0} \\ (\mathbf{H}_1 - \mathbf{H}_2) \times \hat{\mathbf{n}} = \mathbf{0} \\ (n_1^2 \mathbf{E}_1 - n_2^2 \mathbf{E}_2) \cdot \hat{\mathbf{n}} = 0 \\ (\mathbf{H}_1 - \mathbf{H}_2) \cdot \hat{\mathbf{n}} = 0\end{gathered} \tag{2.7}$$

In words, all components of the magnetic field are continuous across a dielectric interface, as are the tangential components of the electric field. The normal components of the electric field are discontinuous, in such a way that $(n^2 \mathbf{e} \cdot \hat{\mathbf{n}})$ is continuous.

Next, we assume that all field components have a time-dependence of $e^{j\omega t}$,

$$\mathbf{E}(x,y,z,t) = \mathbf{Re}\{\underline{\mathbf{E}}(x,y,z)e^{j\omega t}\} \tag{2.8}$$

$$\mathbf{H}(x,y,z,t) = \mathbf{Re}\{\underline{\mathbf{H}}(x,y,z)e^{j\omega t}\} \tag{2.9}$$

and we rewrite Eq. 2.3–2.6 in terms of the complex field quantities $\underline{\mathbf{E}}$ and $\underline{\mathbf{H}}$.

$$\nabla \times \underline{\mathbf{E}} = -jk\eta_0 \underline{\mathbf{H}} \tag{2.10}$$

$$\nabla \times \underline{\mathbf{H}} = jk \frac{1}{\eta_0} n^2 \underline{\mathbf{E}} \tag{2.11}$$

$$\nabla \cdot (n^2 \underline{\mathbf{E}}) = 0 \tag{2.12}$$

$$\nabla \cdot (\underline{\mathbf{H}}) = 0 \tag{2.13}$$

---

[1] Throughout this work, we shall use the following typeface conventions: (1) Real-valued electromagnetic vector fields are symbolized by bold, capital letters (e.g., $\mathbf{E}$). (2) Complex vector fields which have an assumed time dependence of $e^{j\omega t}$ are denoted with an underbar (e.g., $\underline{\mathbf{E}}$). (3) Complex vector fields which for which the spatial z-dependence has been factored out, are symbolized by lowercase boldfaces letters (e.g., $\mathbf{e}$). These conventions can be summarized by the following equation: $\mathbf{E}(x,y,z,t) = \mathbf{Re}\{\underline{\mathbf{E}}(x,y,z)e^{j\omega t}\} = \mathbf{Re}\{\mathbf{e}\, e^{j(\omega t - \beta z)}\}$.

In the above equations, k denotes the free-space wave vector, which is proportional to the optical frequency and has dimensions of inverse length,

$$k \equiv \frac{\omega}{c} \tag{2.14}$$

and η_0 is the free-space wave impedance,

$$\eta_0 \equiv \sqrt{\frac{\epsilon_0}{\mu_0}} \simeq 377\Omega \quad . \tag{2.15}$$

The full-vector eigenvalue equation can be derived from Maxwell's equations, as described in [14]. First, one computes the curl of Eq. 2.10

$$\nabla \times \nabla \times \underline{\mathbf{E}} = -jk\eta_0 \nabla \times \underline{\mathbf{H}} = k^2 n^2 \underline{\mathbf{E}} \quad , \tag{2.16}$$

which can be simplified via the vector identity,

$$\nabla \times \nabla \times \underline{\mathbf{E}} = \nabla(\nabla \cdot \underline{\mathbf{E}}) - \nabla^2 \underline{\mathbf{E}} \quad . \tag{2.17}$$

Then we rewrite the divergence equation as

$$\nabla \cdot (n^2 \underline{\mathbf{E}}) = \nabla(n^2) \cdot \underline{\mathbf{E}} + n^2 \nabla \cdot \underline{\mathbf{E}} = 0 \tag{2.18}$$

$$\nabla \cdot \underline{\mathbf{E}} = -\frac{1}{n^2} \nabla(n^2) \cdot \underline{\mathbf{E}} \tag{2.19}$$

Combining Eq. 2.17 and Eq. 2.19 yields the full vector wave equation for the complex electric field $\underline{\mathbf{E}}$:

$$\nabla^2 \underline{\mathbf{E}} + \nabla \left(\frac{1}{n^2} \nabla(n^2) \cdot \underline{\mathbf{E}} \right) + n^2 k^2 \underline{\mathbf{E}} = 0 \tag{2.20}$$

Note that only two components of the electric field are required. If the transverse components e_x and e_y are known, the longitudinal component may be calculated by applying Eq. 2.12. Therefore, it makes sense to separate the electric field into transverse and longitudinal components, and assume a z-dependence of $e^{-j\beta z}$.

$$\underline{\mathbf{E}}(x, y, z) = (\mathbf{e}_t + \hat{\mathbf{z}} e_z) e^{-j\beta z} \tag{2.21}$$

$$\underline{\mathbf{H}}(x, y, z) = (\mathbf{h}_t + \hat{\mathbf{z}} h_z) e^{-j\beta z} \tag{2.22}$$

with this substitution, the full-vector wave equation can be written in terms of the transverse components, $\mathbf{e}_t$,

$$\nabla^2 \mathbf{e}_t + \nabla \left(\frac{1}{n^2} \nabla (n^2) \cdot \mathbf{e}_t \right) + n^2 k^2 \mathbf{e}_t = \beta^2 \mathbf{e}_t \tag{2.23}$$

As mentioned above, the longitudinal component e_z can be computed from $\mathbf{e}_t$ using the divergence relation:

$$j\beta e_z = \nabla \cdot \mathbf{e}_t + \frac{1}{n^2} \nabla (n^2) \cdot \mathbf{e}_t \tag{2.24}$$

Eq. 2.23 can be written more succinctly if we express it in terms of the two transverse field components e_x and e_y. After some algebra, Eq. 2.23 becomes [15, 16],

$$\begin{bmatrix} P_{xx} & P_{xy} \\ P_{yx} & P_{yy} \end{bmatrix} \begin{bmatrix} e_x \\ e_y \end{bmatrix} = \beta^2 \begin{bmatrix} e_x \\ e_y \end{bmatrix} \tag{2.25}$$

where $P_{xx} \dots P_{yy}$ are differential operators defined as:

$$P_{xx} e_x = \frac{\partial}{\partial x} \left[\frac{1}{n^2} \frac{\partial (n^2 e_x)}{\partial x} \right] + \frac{\partial^2 e_x}{\partial y^2} + n^2 k^2 e_x \tag{2.26}$$

$$P_{yy} e_y = \frac{\partial^2 e_y}{\partial x^2} + \frac{\partial}{\partial y} \left[\frac{1}{n^2} \frac{\partial (n^2 e_y)}{\partial y} \right] + n^2 k^2 e_y \tag{2.27}$$

$$P_{xy} e_y = \frac{\partial}{\partial x} \left[\frac{1}{n^2} \frac{\partial n^2 e_y}{\partial y} \right] - \frac{\partial^2 e_y}{\partial x \partial y} \tag{2.28}$$

$$P_{yx} e_x = \frac{\partial}{\partial y} \left[\frac{1}{n^2} \frac{\partial n^2 e_x}{\partial x} \right] - \frac{\partial^2 e_x}{\partial y \partial x} \tag{2.29}$$

Notice that although the transverse components of the electric field need not be continuous across dielectric interfaces, each of the differentiated terms in Eq. 2.25 is continuous.

Eq. 2.25 is a full-vector eigenvalue equation which describes the modes of propagation for an integrated waveguide. The two coupled transverse field components e_x and e_y taken together are the eigenfunction, and the corresponding eigenvalue is β^2. The four remaining field components can be easily derived from these two transverse components by applying Maxwell's equations. The non-zero diagonal terms P_{xy} and P_{yx} reveal that the two field components e_x and e_y are coupled, that is, the eigenvalue equation cannot be divided into two independent eigenvalue equations which can be solved separately for e_x and e_y. Because of this coupling, the eigenmodes of an optical waveguide are usually not purely TE or TM in nature, and they are often referred to as *hybrid* modes [11]. Nevertheless, often one of the two transverse field components is much larger than the other, and the mode can be treated as approximately TE or TM in nature.

As with most eigenvalue equations, there can be more than one eigenpair which satisfies Eq. 2.25. For this reason, the eigenvalue equation is often written with subscripts on e_x, e_y, and β, but we have chosen to omit the subscripts here for clarity. Most integrated optical devices are designed to be "single-mode" waveguides, meaning that Eq. 2.25 has only one eigenmode for each polarization state.

Eq. 2.25 describes the eigenmodes of a waveguide in terms of the transverse electric field $\mathbf{e}_t$, however it is important to realize that equivalent eigenvalue relations can be derived for the other field components. In particular, some prefer to express the eigenvalue equations in terms of the two longitudinal components h_z and e_z [14]. Likewise, a set of equations similar to Eq. 2.25 can be derived for the transverse magnetic field $\mathbf{h}_t$[16]. In any case, only two components of the electromagnetic fields are required to completely specify the optical mode; the remaining components can be derived from Maxwell's equations.

2.1.2 Normalization and Orthogonality

One of the characteristics of eigenfunctions is that they can only be determined up to a scalar multiplicative constant, i.e. if the modal solutions e_x and e_y are scaled by any factor they will still satisfy the eigenvalue equation. To remove this ambiguity, it is often convenient to normalize the mode so that it it has unity power. The time-averaged electromagnetic power transmitted by a propagating mode of a waveguide is described by the Poynting vector, integrated over the x-y plane,

$$P = \frac{1}{4}\iint (\mathbf{e}\times\mathbf{h}^* + \mathbf{e}^*\times\mathbf{h})\cdot\hat{\mathbf{z}}\,dx\,dy \quad . \tag{2.30}$$

Thus, if we require that $P = 1$, we can easily determine the magnitude of the constant which multiplies e_x and e_y. Note however that the phase of this constant remains undetermined, because the field components enter Eq. 2.30 in complex-conjugate pairs. This ambiguity is resolved by arbitrarily choosing the transverse components $\mathbf{e}_t$ to be purely real quantities. It can be seen from Maxwell's equations that this choice of phase implies that the two longitudinal components h_z and e_z are purely imaginary, and the transverse magnetic field components are likewise real quantities. Of course, the complex nature of the electromagnetic fields is nothing more than a ramification of the $e^{j\omega t}$ time dependence assumed in Eq. 2.8. The fact that the transverse field components are real and the longitudinal components are imaginary should therefore be understood to mean that the longitudinal components lead (or lag) the transverse components in time by $\pi/2$.

One further ambiguity is that the propagation constant β is only determined up to a sign. This is to be expected, because the waveguide can support forward- and backward-traveling modes. We adopt the convention that positive values of β correspond to forward-traveling modes, while negative values correspond to backward-traveling modes.

Another characteristic of eigenvalue equations is that eigenfunctions which correspond to different eigenvalues are in orthogonal. For optical waveguides, the orthogonality condition between two discrete modes labelled m and n can be stated as:

$$\frac{1}{4}\iint(\mathbf{e}_n\times\mathbf{h}_m^* + \mathbf{e}_m^*\times\mathbf{h}_n)\cdot\hat{\mathbf{z}}\,dx\,dy = \delta_{mn} \qquad , \tag{2.31}$$

where δ_{mn} is the Kroneker delta function, and we have assumed unit-power normalization for the modes, as described above.

2.1.3 Weakly-Guiding Waveguides

For many waveguides, the refractive index profile varies by only a small fractional amount over the waveguide cross-section. That is, often the index of refraction for the central core region is only slightly higher than that of the surrounding cladding region. For example, in a standard optical fiber the difference in refractive index between the core and cladding is only about 0.3%. These types of waveguides are often referred to as *weakly-guiding* waveguides. The term weakly-guiding does not mean that the light leaks out of the waveguide (indeed, optical fiber has replaced copper as a transmission medium precisely because light does not leak out); rather it means only that the relative refractive index contrast is small.

For weakly-guiding waveguides, the modal analysis can be greatly simplified by replacing the full-vector eigenvalue equation by a simple scalar eigenvalue equation for a single field component. Examining the differential operators in Eq. 2.25, we see that when the refractive index profile is constant, the off-diagonal terms P_{xy} and P_{yx} vanish, leading to decoupled eigenvalue equations for e_x and e_y:

$$\begin{bmatrix} P_{xx} & 0 \\ 0 & P_{yy} \end{bmatrix}\begin{bmatrix} e_x \\ e_y \end{bmatrix} = \beta^2\begin{bmatrix} e_x \\ e_y \end{bmatrix} \tag{2.32}$$

Eq. 2.32 is known as the semivectorial eigenvalue equation. The polarization-dependent continuity relations for the two transverse field components are maintained in this equation, but the coupling between the two transverse components is ignored.

The semivectorial eigenmode equation can be further simplified by replacing the differential operators P_{xx} and P_{yy} with a simplified second-order Laplacian operator:

$$P_{xx} \simeq P_{yy} \simeq P = \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} + n^2 k^2 \tag{2.33}$$

$$P\phi(x,y) = \beta^2 \phi(x,y) \tag{2.34}$$

Where $\phi(x,y)$ can represent either transverse field component. Again, this simplification is only valid for weakly-guiding waveguides for which variations in the refractive index are small. The field $\phi(x,y)$ in Eq. 2.33 is assumed to be continuous at all points, even across dielectric interfaces. The scalar eigenmode equation described in Eq. 2.33 does not account for any polarization-dependence, and therefore cannot distinguish between TE and TM polarized modes. However, for many weakly-guiding waveguides the polarization-dependence arises primarily because of stress and strain in the material layers comprising the device, and not because of modal birefringence. The scalar mode equation is often sufficiently accurate for modeling weakly-guiding waveguides.

One way to think about the scalar approximation is to imagine light trapped inside of the waveguide core by total internal reflection. Recall that in order for light to be totally internally reflected the angle of incidence must exceed the critical angle. The critical angle for total internal reflection depends upon the index difference between the internal and external layers and when the index contrast is small, only light at grazing incidence will be totally internally reflected. Thus, for weakly-guiding waveguides, the confined light may be regarded as approximately TEM in nature.

2.1.4 Finite Difference Methods

Now that we have established the eigenmode equations for dielectric waveguides, we turn to the more practical matter of how to calculate the electromagnetic modes for an integrated waveguide. There are a few waveguides for which the eigenmodes can be computed analytically. For example, when the index profile consists of stratified layers of homogeneous dielectric materials, the eigenmode equations can be reduced to a simple one-dimensional problem which can be solved analytically by matching boundary conditions at all of the dielectric interfaces. The cylindrical optical fiber is another example of a problem which can be solved exactly; because of its cylindrical symmetry, the problem can be reduced to an equivalent one-dimensional eigenvalue equation.

Integrated waveguides, by contrast, are usually rectangular structures which confine the light in both transverse directions. Because they do not have planar or cylindrical symmetry, the eigenmodes of these structures cannot be computed analytically. Instead, numerical techniques must be used to solve the eigenvalue equations. There are many different numerical techniques for solving partial differential equations, including finite element methods [17, 18], finite difference methods [15, 19, 20, 16, 21], and boundary integral techniques [22]. Each of these techniques has its own advantages and disadvantages. In this section and in Appendix A, we will describe a finite-difference technique for discretizing the eigenvalue equation.

In the finite difference technique, differential operators are replaced by difference equations. As a simple example, the first derivative of a function $f(x)$ could be approximated as

$$f'(x) \simeq \frac{f(x+\Delta x) - f(x)}{\Delta x} \quad . \tag{2.35}$$

This is a very intuitive approximation. In fact, most elementary calculus textbooks define the first derivative of a function to be just such a finite-difference in the limit that $\Delta x \to 0$. As we will show later, similar finite-difference equations can be developed to approximate higher-order derivatives and mixed derivatives. Of course, Eq. 2.35 fails entirely if the function $f(x)$ is discontinuous in the interval $x \to x + \Delta x$. Moreover, Maxwells equations *predict* that the normal components of the electric field are discontinuous across abrupt dielectric interfaces. Therefore in order to develop an accurate model for the eigenmodes of an optical waveguide, we must construct a finite difference scheme which accounts for the discontinuities in the eigenmodes. We will later show how such a finite difference scheme can be derived.

Once the difference equations have been described, the partial-differential equation can be translated into an equivalent matrix equation. The functions $e_x(x, y)$ and $e_y(x, y)$ are replaced by vectors representing the value of the functions at discrete points. The differential operators P_{xx}, P_{yx}, P_{xy} and P_{yy}, are replaced by sparse, banded matrices which describe sums and differences between adjacent samples. With this substitution, Eq. 2.25 becomes a conventional matrix eigenvalue equation.

Figure 2.2 illustrates a typical finite difference mesh for a ridge waveguide. The refractive index profile has been broken up into small rectangular elements or pixels, of size $\Delta x \times \Delta y$. Over each of these elements, the refractive index is constant. Thus, discontinuities in the refractive index profile occur only at the boundaries between adjacent pixels. Because the index profile is symmetric about the y-axis, only half of the waveguide needs to be included in the computational domain. The computational window must extend far

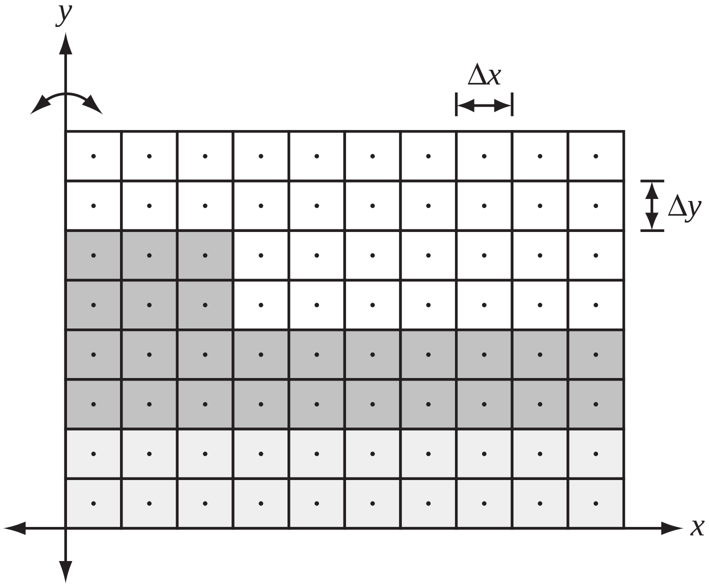

Figure 2.2: A typical finite difference mesh for an integrated waveguide. The refractive index profile $n(x, y)$ has been divided into small rectangular cells over which $n(x, y)$ is taken to be constant. For symmetric structures, such as this one, only half of the waveguide needs to be included in the computation window. (figs/2/fdmesh.eps)

enough outside of the waveguide core in order to completely encompass the optical mode.

The finite-difference grid points, i.e., the discrete points at which the fields are sampled, are located at the center of each cell. Some finite difference schemes instead choose to locate the grid points at the vertices of each cell rather than at the center. This approach works well for finite-difference schemes involving the magnetic field $\mathbf{h}$ which is continuous across all dielectric interfaces [20, 19]. However, the normal component of the electric field is discontinuous across an abrupt dielectric interface, which leads to an ambiguity if the grid points are placed at the cell vertices.

It is worth pointing out that the finite difference method described here can also be used to develop beam-propagation models. The structure of the sparse matrices remains unchanged, but the problem becomes one of repeatedly solving a sparse system of linear equations to simulate mode propagation, rather than computing eigenvalues [16, 23, 24, 25].

Scalar Finite Difference Equations

We will begin by deriving the finite difference equations for the scalar eigenmode approximation. Recall that in this approximation the coupled full-vector eigenmodes equations (Eq. 2.25) have been replaced by a single scalar eigenmode equation for one of the transverse field components denoted $\phi(x, y)$ (Eq. 2.33). This approximation is valid for so-called "weakly-guiding" waveguides in which the refractive index contrast is small. We repeat the scalar eigenvalue equation here for reference:

$$\left\{ \frac{\partial^2}{\partial x^2} + \frac{\partial^2}{\partial y^2} - n^2(x, y)k^2 \right\} \phi(x, y) = \beta^2 \phi(x, y) \tag{2.36}$$

In order to translate this partial differential equation into a set of finite difference equations, we must approximate the second derivatives in terms of the values of $\phi(x, y)$ at surrounding gridpoints. We shall use the subscripts N, S, E and W, to indicate the value of the field (or index profile) at grid-points immediately north, south, east and west of the point under consideration, P. This labeling scheme is illustrated in Fig. 2.3.

One of the most straightforward techniques for deriving finite difference approximations is Lagrange interpolation[26]. The Lagrange interpolant is simply the lowest order polynomial which goes through all of the sample points. The derivatives can then be easily computed from the polynomial coefficients of the interpolating function. For example, to approximate the second derivative of ϕ with respect to x at point P, $\frac{\partial^2 \phi}{\partial x^2}|_P$, we simply fit a quadratic equation to the three points ϕ_E, ϕ_P, and ϕ_W:

$$\tilde{\phi}(x) = A + Bx + Cx^2, \tag{2.37}$$

Where the three coefficients A, B and C are determined by the matching the function $\tilde{\phi}(x)$ at the three adjacent gridpoints, i.e.,

$$\tilde{\phi}(-\Delta x) = \phi_W, \quad \tilde{\phi}(0) = \phi_P, \quad \tilde{\phi}(+\Delta x) = \phi_E \qquad . \tag{2.38}$$

Notice that for convenience we have arbitrarily chosen to place the origin of our local coordinate system ($x = 0$) at point P. The second derivative is related to the x^2 coefficient,

$$\frac{\partial^2 \phi}{\partial x^2}|_P \simeq \frac{\partial^2 \tilde{\phi}}{\partial x^2} = 2C \qquad . \tag{2.39}$$

Solving the three equations of Eq. 2.38, we arrive at the following finite difference approx-

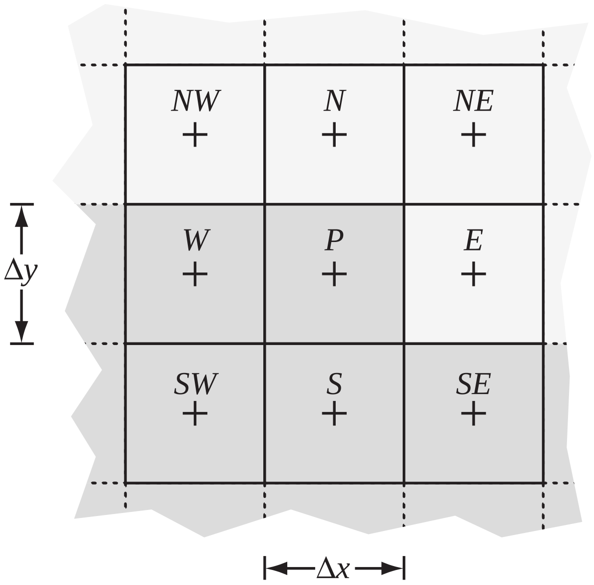

Figure 2.3: Labeling scheme used for the finite difference model. The subscripts P, N, S, E, W, NE, NW, SW and SE are used to label respectively the grid point under consideration, and its nearest neighbors to the north, south, east, west, north-east, north-west, south-west, and south-east. (figs/2/nsew-label.eps)

imations:

$$\frac{\partial^2\phi}{\partial x^2}|_P \simeq \frac{1}{(\Delta x)^2}(\phi_W - 2\phi_P + \phi_E) \tag{2.40}$$

$$\frac{\partial\phi}{\partial x}|_P \simeq \frac{1}{2\Delta x}(\phi_E - \phi_W) \qquad . \tag{2.41}$$

These approximations can also be derived by performing a second order Taylor expansion of the field about the point P. However, we have chosen to derive the finite-difference equations via polynomial interpolation because this approach is easily adaptable to non-uniform grid sizes, and more importantly it can be extended to account for predictable discontinuities in the field ϕ. Similar finite difference approximations apply in the vertical direction:

$$\frac{\partial^2\phi}{\partial y^2}|_P \simeq \frac{1}{(\Delta y)^2}(\phi_S - 2\phi_P + \phi_N) \tag{2.42}$$

$$\frac{\partial\phi}{\partial y}|_P \simeq \frac{1}{2\Delta y}(\phi_N - \phi_S) \qquad . \tag{2.43}$$

With these approximations, the differential operator P may be replaced with its finite difference representation to arrive at the following discretized difference equation:

$$\frac{\phi_W}{(\Delta x)^2} + \frac{\phi_E}{(\Delta x)^2} + \frac{\phi_N}{(\Delta y)^2} + \frac{\phi_S}{(\Delta y)^2} + \left(n_P^2 k^2 - \frac{2}{(\Delta y)^2} - \frac{2}{(\Delta x)^2}\right)\phi_P = \beta^2\phi_P \tag{2.44}$$

The finite difference operator, which we shall denote $\hat{P}$, can be more conveniently represented by the following diagram which illustrates the coefficients which multiply each of the adjacent sample-points.

$$\hat{P} : \quad \begin{array}{|c|c|c|} \hline 0 & \dfrac{1}{(\Delta y)^2} & 0 \\ \hline \dfrac{1}{(\Delta x)^2} & n_P^2 k^2 - \dfrac{2}{(\Delta y)^2} - \dfrac{2}{(\Delta x)^2} & \dfrac{1}{(\Delta x)^2} \\ \hline 0 & \dfrac{1}{(\Delta y)^2} & 0 \\ \hline \end{array} \tag{2.45}$$

As we will describe in Appendix A, the finite difference equations must be slightly modified for points which lie on the boundary of the computation window. If we apply Eq. 2.45 for each point in the computation window, we obtain an M-dimensional eigenvalue problem, where M is the total number of grid-points, i.e., $M = n_x n_y$. Figure 2.4 illustrates the structure of the eigenvalue equation for a simple grid with $n_x = 4$ and $n_y = 3$. As is customary, we have chosen to number of grid points from left to right and bottom to top, which results in a block tridiagonal matrix structure as shown in the lower portion of Fig. 2.4.

Vector Finite Difference Equations

As noted earlier, the scalar eigenvalue equation is only applicable in cases where the refractive index contrast is small. The polynomial interpolation process seems reasonable for describing continuous functions, but not all components of the electric field are continuous at dielectric interfaces. When the refractive index differences are small, the fields may be treated as continuous without significantly affecting the accuracy of the solution. For problems which do not meet this criterion, a more accurate finite difference model is required. Appendix A describes how the finite-difference equations can be modified to account for such index discontinuities.

Computation of Eigenvalues

As described above, the finite difference method essentially translates a partial differential eigenvalue equation into a conventional matrix eigenvalue equation. The partial differential operators have been replaced by large sparse matrices, and the eigenfunctions have been replaced by long vectors representing a sampling of the eigenfunctions at discrete grid-points.

Once we have set up this matrix equation, we must solve for the eigenvalues and eigenvectors. Naturally, since the matrix is of dimension $M = n_x n_y$, there should be M eigenpairs. However, we are only interested in computing the largest few eigenvalues. The smaller eigenvalues correspond to unphysical eigenmodes.

There are many routines available for computing a few selected eigenvalues of large sparse matrices. The most common technique is the shifted inverse power method [27]. Unfortunately, this technique proves to be relatively slow and it is only capable of computing one eigenfunction at a time. A complete review of the available routines for computing eigenvalues of sparse matrices is given in [28]. One of the most promising algorithms is the implicitly restarted Arnoldi method [29]. This method allows one to simultaneously compute a few of the largest eigenvalues of the sparse matrix. For this work, we used the built-in Matlab function `eigs`, which implements a variant of the Arnoldi method.

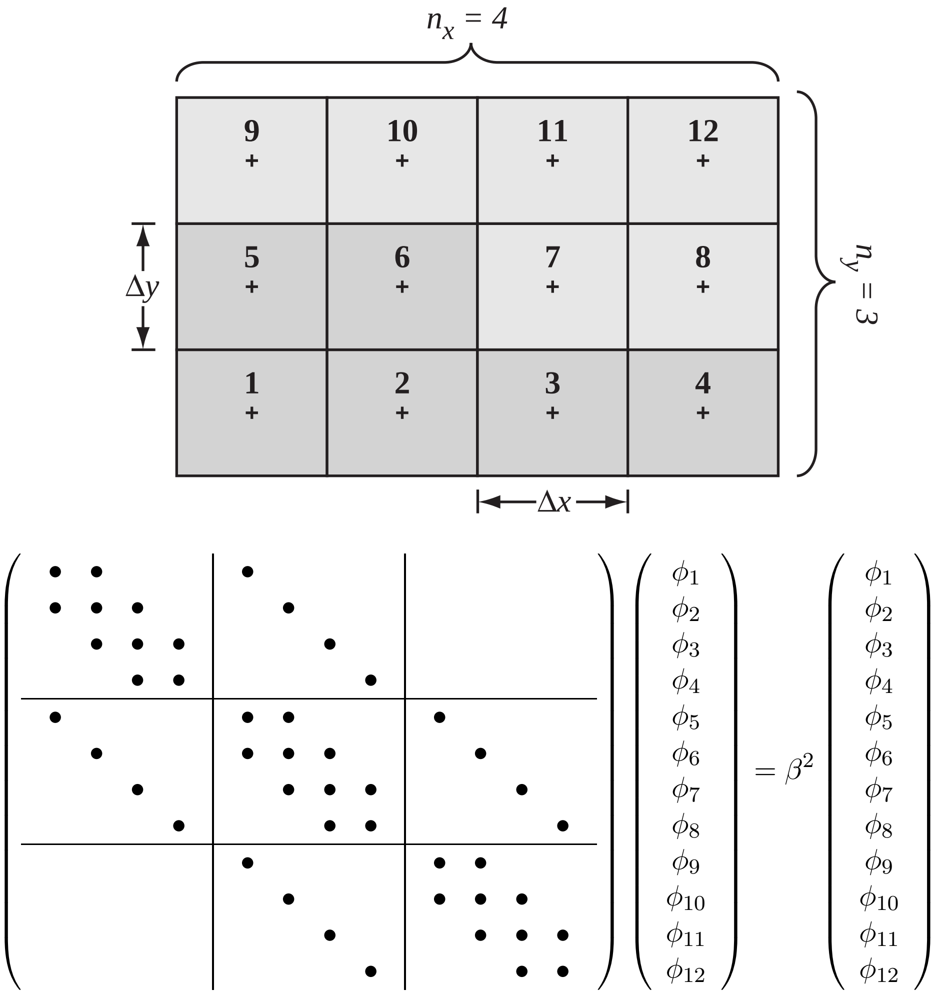

Figure 2.4: Structure of the eigenvalue equation for the finite difference problem, applied to a simple 4×3 index mesh. The partial differential operator P has been replaced by its finite difference matrix equivalent, and the eigenfunction $\phi(x, y)$ is replaced by samples at discrete grid points. The grid points are numbered sequentially from left to right, and bottom to top, which results in the block tridiagonal matrix structure shown. The •'s in this equation represent nonzero elements of the matrix. (figs/2/matrix-shape.eps)

2.1.5 Computed Eigenmodes for Optical Waveguides

The two types of dielectric waveguides considered in this work are illustrated in Fig. 2.5. The first is a doped-glass channel waveguides, whose index profile and mode shape are designed to match well with that of an optical fiber [30, 31]. The second type of waveguide that will be analyzed is a silicon-on-insulator (SOI) ridge waveguide [32, 33].

The doped-glass waveguide consists of a rectangular core region surrounded by a cladding region with slightly lower index of refraction. The lower cladding layer is fused silica (SiO_2). By doping the core region with phosphorus or germanium, the index of refraction can be raised slightly with respect to the underlying silica. The top cladding layer is co-doped with both boron and phosphorus in order to match the refractive index of the lower cladding. The index contrast for this type of waveguide typically ranges from 0.3% to 0.8%, and can be adjusted by varying the dopant concentration in the core layer. Figure 2.6 plots the calculated mode profile for a Ge-doped glass channel waveguide with an index contrast of 0.8%. The waveguides measure six microns on each side, which insures that the structure only supports one bound mode for each polarization state. Notice that because of the rectangular symmetry of the structure, only one quadrant of the mode needs to be included in the calculation.

In the silicon-on-insulator ridge waveguide, the light is confined in the silicon ridge structure which sits on top of an oxide separation layer. Provided the oxide layer is thick enough, the light will remain confined in the silicon layer without escaping into the silicon substrate. By choosing the ridge height appropriately, the structure can be made to have only one bound mode per polarization state, even for relatively large mode sizes [34, 35]. Because the structure requires no top cladding layer (the air above the waveguide forms the top cladding), this structure avoids some of the challenging problems of material overgrowth. Figure 2.7 plots the calculated mode profiles for an SOI ridge waveguide. For this structure, the core height is 3 µm, and the ridge width is 4 µm.

2.2 Coupled Waveguides

In the preceding section, we described techniques for describing the modes of propagation for an optical waveguide. In analyzing the waveguide, we assumed that the structure can be described by a z-invariant refractive index profile $n(x, y)$ which extends infinitely in the transverse directions. Such a waveguide performs no real optical function except to transmit a light signal from one point to another. In this section, we investigate a slightly

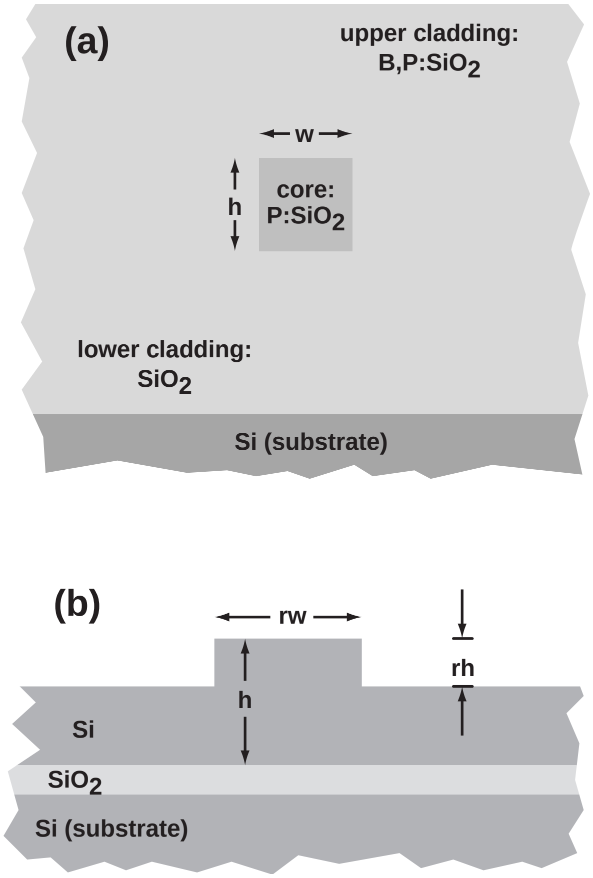

Figure 2.5: The two types of waveguides considered in this work. (a) A glass channel waveguide, of the type described in references [30, 31]. The core is doped with phosphorus (P) or germanium (Ge) to increase the refractive index relative to that of the underlying undoped silica. The upper cladding is co-doped with boron and phosphorus to match the refractive index of the lower core. (b) Silicon-on-insulator (SOI) ridge waveguide. The optical mode is guided in the silica ridge, and confined by the oxide layer below and air above [32]. (figs/2/waveguide-types.eps)

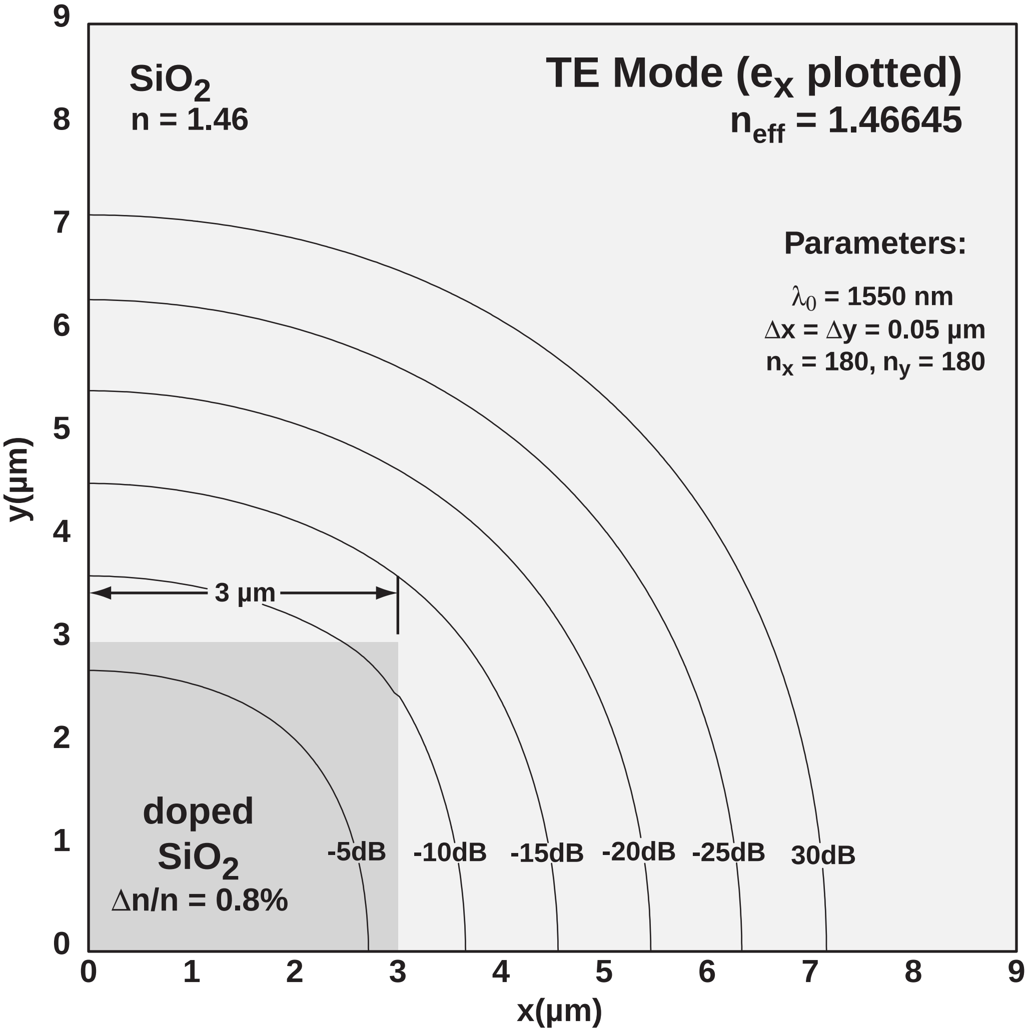

Figure 2.6: Mode profile for an integrated channel waveguide in silica. For the simulation, the index contrast was taken to be 0.8%, which could be achieved by doping the core of region with germanium. The dimensions of the waveguide are 6 μm × 6 μm. Plotted here is the transverse electric field component e_x for the fundamental TE mode. The orthogonal field component e_y (not plotted) is approximately 30-40 dB lower than e_x. Note that because of the rectangular symmetry of the waveguide, only one quadrant was included in the computational window. The fundamental TE and TM modes are degenerate, because the waveguide is perfectly square. (figs/2/glass-wg-modeprofile.eps)

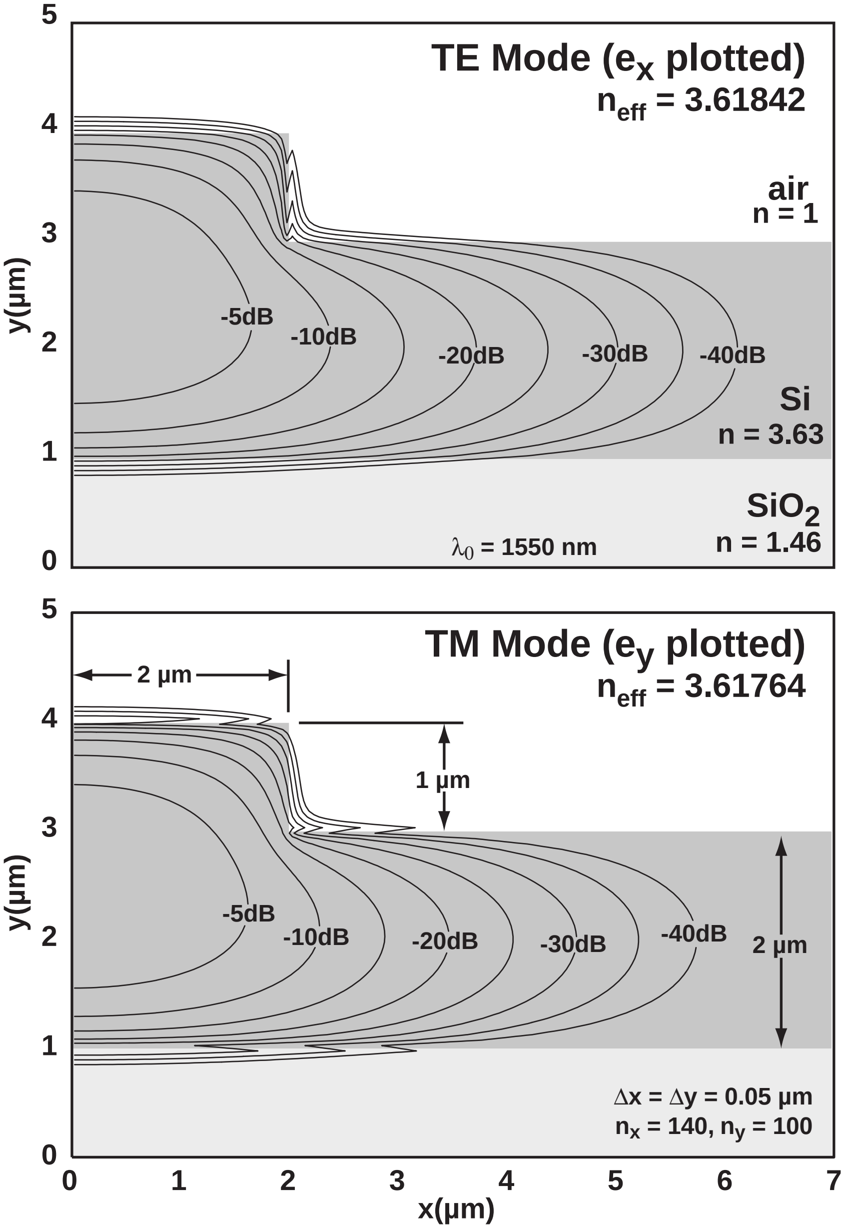

Figure 2.7: Mode profile for an integrated silicon-on-insulator (SOI) ridge waveguide. The upper portion of this plot depicts the principal field component e_x for the fundamental TE mode of the structure, and the lower portion of the plot depicts the principal field component e_y for the TM mode of the structure. (figs/2/soi-wg-modeprofile.eps)

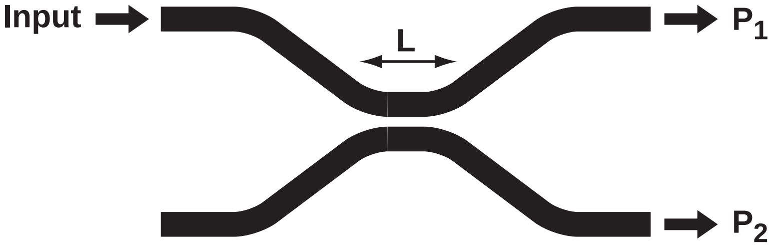

Figure 2.8: The structure of a typical integrated directional coupler. Two waveguides, initially separated, are brought together so that power may transfer from one to the other. The path of approach and length of interaction must be carefully engineered to achieve the desired amount of power transfer. (figs/2/simple-coupler-schematic.eps)

more complicated structure consisting of two (or more) interacting waveguides in close proximity. Such a structure, which we call a waveguide coupler, performs the important task of transferring light from one waveguide to another. Waveguide couplers are important components in Mach-Zehnder interferometers, power splitters, and a variety of other integrated optical devices.

Figure 2.8 illustrates schematically the type of structure which we wish to describe. Two waveguides, labeled 1 and 2, which are initially separated are slowly brought close to one another over some interaction length L. The waveguide separation and interaction length must be selected in order to achieve the desired amount of power transfer.

2.2.1 Variational Approach

We begin by simplifying the problem to the analysis of two parallel waveguides, as depicted in Fig. 2.9, ignoring for the moment the gradual approach and separation at either end of the device. We shall denote the electromagnetic modes of waveguides 1 by $\mathbf{e}_1$ and $\mathbf{h}_1$, and those of waveguide 2 by $\mathbf{e}_2$ and $\mathbf{h}_2$.

$$\begin{aligned} n_1^2(x,y) &\to \mathbf{e}_1(x,y), \quad \mathbf{h}_1(x,y) \\ n_2^2(x,y) &\to \mathbf{e}_2(x,y), \quad \mathbf{h}_2(x,y) \qquad , \end{aligned} \tag{2.46}$$

where $n_1(x,y)$ denotes the index profile of waveguide 1 in the absence of waveguide 2, and $n_2(x,y)$ is the index profile of waveguide 2 in the absence of waveguide 1.

Next, we attempt to describe the electromagnetic fields of the coupled waveguide system as a superposition of the modes of waveguides 1 and 2.

$$\begin{aligned} \underline{\mathbf{E}}(x,y,z) &= a_1(z)\mathbf{e}_1(x,y) + a_2(z)\mathbf{e}_2(x,y) \\ \underline{\mathbf{H}}(x,y,z) &= a_1(z)\mathbf{h}_1(x,y) + a_2(z)\mathbf{h}_2(x,y) \end{aligned} \tag{2.47}$$

In the above equation, $a_1(z)$ and $a_2(z)$ are scalar functions of z which represent respectively the mode amplitude in waveguide 1 and the mode amplitude in waveguide 2. When the two waveguides are very far apart, i.e. when d is large compared to the mode size, the two optical modes should propagate independently without interaction as described in Section 2.1. In this case, the solution for $a_i(z)$ is:

$$\begin{aligned} a_1(z) &= a_1(0)\exp(-j\beta_1 z) \\ a_2(z) &= a_2(0)\exp(-j\beta_2 z) \end{aligned} \tag{2.48}$$

Eq. 2.48 can be viewed as the solution to the following differential equation:

$$\frac{d}{dz}\begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} = \begin{bmatrix} -j\beta_1 & \\ & -j\beta_2 \end{bmatrix}\begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} \tag{2.49}$$

The goal of this section is to derive a new differential equation which describes the evolution of $a_1(z)$ and $a_2(z)$ when the two waveguides are brought close together. Essentially, we seek to replace Maxwells equations by a system of two coupled differential equations for the scalar mode amplitudes. This analysis comes under the rubric of coupled mode theory[36].

Before proceeding, we should point out that the mode expansion of Eq. 2.47 is only a convenient approximation. While for a single isolated waveguide, the electromagnetic fields may be accurately described as a superposition of the orthogonal modes, the modes of the two constituent waveguides in a coupler do not comprise an orthonormal basis set. Nevertheless, it seems reasonable to use the expansion of Eq. 2.47 as a trial function. A more rigorous analysis of the waveguide coupler will be presented in Section 2.2.2.

One of the most complete methods for analyzing coupled waveguides is the variational approach described in reference [37], which we summarize here. The variational method

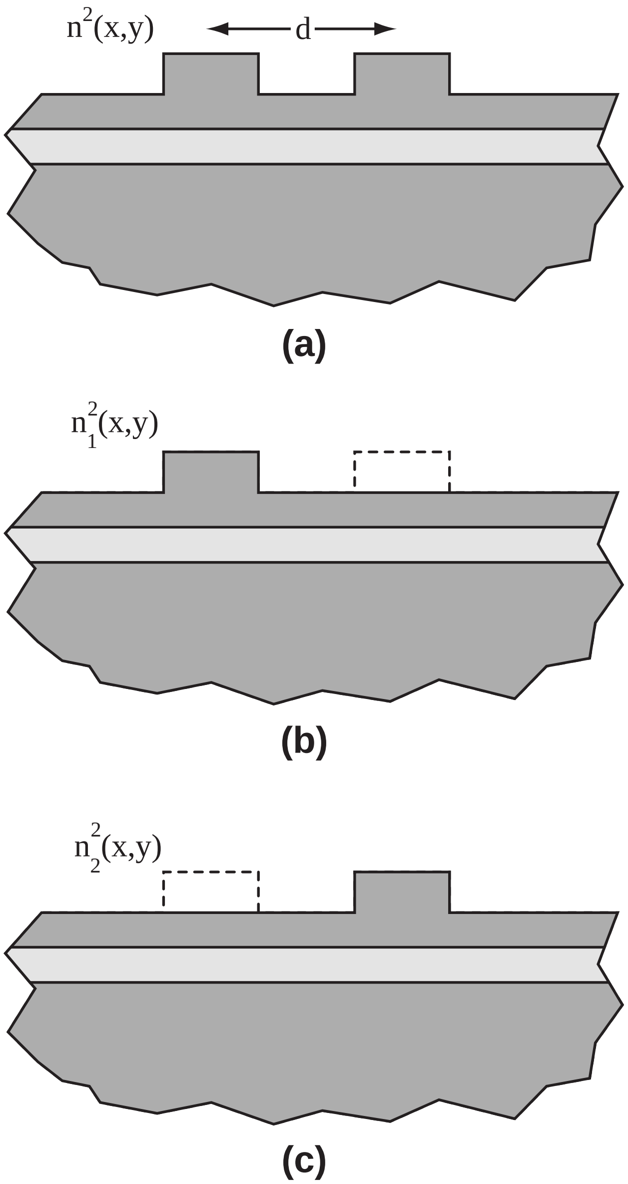

Figure 2.9: (a) Cross-sectional diagram of two parallel waveguides (ridge waveguides in this example) separated by a center-to-center distance d. (b) Refractive index profile for waveguide 1 in the absence of waveguide 2. (c) Refractive index profile for waveguide 2 in the absence of waveguide 1. (figs/2/parallel-waveguides.eps)

begins with an integral expression for the propagation constant β.

$$\beta = \frac{\frac{j}{4}\iint \left\{(\nabla_t \times \mathbf{h} - \frac{jk}{\eta_0} n^2 \mathbf{e}) \cdot \mathbf{e}^* - (\nabla_t \times \mathbf{e} - jk\eta_0 \mathbf{h}) \cdot \mathbf{h}^*\right\} dx\, dy}{\frac{1}{4}\iint (\mathbf{e}_t \times \mathbf{h}_t^* + \mathbf{e}_t^* \times \mathbf{h}_t) \cdot \hat{\mathbf{z}}\, dx\, dy} \tag{2.50}$$

In the above equation, the integration is performed of the entire x-y plane, and n^2 refers to the complete refractive index profile, including both waveguides. The fields $\mathbf{e}$ and $\mathbf{h}$ and propagation constant β likewise refer to the eigenmodes of the coupled waveguide system considered as a whole. (Note that we have assumed a z-dependence of $e^{-j\beta z}$ for the field quantities of Eq. 2.50.) Of course, we could apply the techniques of Section 2.1 to rigorously solve for the modes of the aggregate structure, but for now, we will treat the quantities $\mathbf{e}$, $\mathbf{h}$ and β as unknown eigenmodes which we wish to approximate in terms of the modes of the constituent waveguides considered separately.

As shown in [37], Eq. 2.50 is a variational expression for the propagation constant β. This means that if the electromagnetic fields $\mathbf{e}$ and $\mathbf{h}$ are perturbed slightly, i.e.,

$$\begin{aligned} \mathbf{e} &\rightarrow \mathbf{e} + \delta\mathbf{e} \\ \mathbf{h} &\rightarrow \mathbf{h} + \delta\mathbf{h} \end{aligned} \quad , \tag{2.51}$$

the value of the integral of Eq. 2.50 changes only to second-order in $\delta\mathbf{e}$ and $\delta\mathbf{h}$.

$$\beta \rightarrow \beta + \iint \left[\alpha^2 |\delta\mathbf{e}|^2 + \gamma^2 |\delta\mathbf{h}|^2\right] dx\, dy \tag{2.52}$$

Therefore, one way to determine the two fields $\mathbf{e}$ and $\mathbf{h}$ is to find the two vector functions which minimize the value of the integral given in Eq. 2.50. [2]

Clearly, it is unreasonable to minimize Eq. 2.50 over all continuous vector functions of x and y. In the variational approach, we instead perform a constrained minimization in which we assume that the fields $\mathbf{e}$ and $\mathbf{h}$ are described by the superposition of the isolated waveguide modes.

$$\begin{aligned} \tilde{\mathbf{e}}(x,y) &= a_1 \mathbf{e}_1(x,y) + a_2 \mathbf{e}_2(x,y) \\ \tilde{\mathbf{h}}(x,y) &= a_1 \mathbf{h}_1(x,y) + a_2 \mathbf{h}_2(x,y) \end{aligned} \tag{2.53}$$

The above equation is idential to Eq. 2.47 except that we have factored out the $e^{-j\beta z}$ dependence from $\underline{\mathbf{E}}$, $\underline{\mathbf{H}}$, and $a_i(z)$. We have added a tilde to the fields $\tilde{\mathbf{e}}$ and $\tilde{\mathbf{h}}$ to distinguish

---

[2] In fact, a similar minimization principle forms the basis for all finite-element mode solvers [17, 38].

them from the *exact* solutions $\mathbf{e}$ and $\mathbf{h}$. Rather than minimizing Eq. 2.50 over the space of all continuous functions, we instead minimize over a two-dimensional subspace consisting of all possible linear combinations of the two isolated waveguide modes.

We shall use $\tilde{\beta}$ to denote the minimal value achieved by Eq. 2.50 under the linear superposition constraint. This optimal value should closely approximate the real propagation constant β of the coupled waveguide structure. Likewise, the fields which minimize Eq. 2.50 should closely approximate the actual electromagnetic mode of the aggregate structure:

$$\tilde{\beta} \simeq \beta, \qquad \tilde{\mathbf{e}} \simeq \mathbf{e}, \qquad \tilde{\mathbf{h}} \simeq \mathbf{h} \tag{2.54}$$

In fact, the error between the actual fields $\mathbf{e}$ and $\mathbf{h}$ and the optimal linear superposition $\tilde{\mathbf{e}}$ and $\tilde{\mathbf{h}}$ can easily be shown to be orthogonal to any function in the subspace over which the minimization is performed. (By "orthogonal", we refer to the inner product described in Eq. 2.31.)

If the trial functions of Eq. 2.53 are substituted into Eq. 2.50, the resulting integral expression for β can be cast into the following form:

$$\beta = \frac{\mathbf{a}^\dagger \mathbf{H} \mathbf{a}}{\mathbf{a}^\dagger \mathbf{P} \mathbf{a}} \qquad , \tag{2.55}$$

where $\mathbf{a}$ is a two-dimensional vector of mode amplitudes and $\mathbf{H}$ and $\mathbf{P}$ are Hermitian matrices defined below.

$$\mathbf{a} = \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \tag{2.56}$$

$$P_{ij} = \frac{1}{4} \iint (\mathbf{e}_j \times \mathbf{h}_i^* + \mathbf{e}_i^* \times \mathbf{h}_j) \cdot \hat{\mathbf{z}}\, dx\, dy \tag{2.57}$$

$$H_{ij} = P_{ij}\beta_j + \frac{k}{4\eta_0} \iint \left[(n^2 - n_j^2)\mathbf{e}_j \cdot \mathbf{e}_i^*\right] dx\, dy \tag{2.58}$$

All of the integrals in the above equations are taken over the entire x-y plane. However, notice that the integrand involved in H_{ij} vanishes for points outside of waveguide i and therefore the range of this integral may be restricted to waveguide i.

Minimizing Eq. 2.55 with respect to the two components of $\mathbf{a}$ leads to the following equation:

$$\beta \mathbf{P} \mathbf{a} = \mathbf{H} \mathbf{a} \qquad . \tag{2.59}$$

The above equation can be seen to be a generalized eigenvalue equation, with eigenvectors $\mathbf{a}$. Because we know that β describes the propagation constant of the fields, we can obtain the coupled mode equations by simply replacing β with $j\frac{d}{dz}$ in Eq. 2.59:

$$\begin{bmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{bmatrix} \frac{d}{dz} \begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} = -j \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} \begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} \quad . \tag{2.60}$$

Eq. 2.60 is a coupled two-dimensional linear differential equation which describes the evolution of the coefficients $a_1(z)$ and $a_2(z)$ for the parallel waveguide system. The diagonal elements of $\mathbf{P}$ represent the modal power carried by each of the two isolated waveguide modes separately. The off-diagonal elements of $\mathbf{P}$ describe the extent to which the two isolated waveguide modes are not orthogonal to each other.

We will now examine the solution to this system of equations, in the case where the two waveguides under consideration are identical (we shall let $\beta_1 = \beta_2 \equiv \beta_0$). For identical waveguides, Eq. 2.60 can be cast into the following form:

$$\begin{bmatrix} 1 & x \\ x & 1 \end{bmatrix} \frac{d}{dz} \begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} = -j \left\{ \begin{bmatrix} 1 & x \\ x & 1 \end{bmatrix} \begin{bmatrix} \beta_0 & 0 \\ 0 & \beta_0 \end{bmatrix} + \begin{bmatrix} \Delta & \mu \\ \mu & \Delta \end{bmatrix} \right\} \begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} \tag{2.61}$$

where the constants x, Δ, and μ are defined by,

$$x \equiv \frac{1}{4P} \iint (\mathbf{e}_1 \times \mathbf{h}_2^* + \mathbf{e}_2^* \times \mathbf{h}_1) \cdot \hat{\mathbf{z}} \, dx \, dy \tag{2.62}$$

$$\Delta \equiv \frac{1}{4P} \frac{k}{\eta_0} \iint_{\text{guide 2}} \mathbf{e}_1 \cdot \mathbf{e}_1^* \, dx \, dy \tag{2.63}$$

$$\mu \equiv \frac{1}{4P} \frac{k}{\eta_0} (n_{\text{core}}^2 - n_{\text{clad}}^2) \iint_{\text{guide 1/2}} \mathbf{e}_1 \cdot \mathbf{e}_2^* \, dx \, dy \tag{2.64}$$

By inverting the matrix on a left-hand side of Eq. 2.61, the coupled mode equations simplify to:

$$\frac{d}{dz} \begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} = -j \begin{bmatrix} \beta_0' & \mu' \\ \mu' & \beta_0' \end{bmatrix} \begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} \quad , \tag{2.65}$$

where the quantities β_0' and μ' are given by

$$\beta_0' \equiv \beta_0 + \frac{\Delta - \mu x}{1 - x^2} \tag{2.66}$$

$$\mu' \equiv \frac{\mu - x\Delta}{1 - x^2} \tag{2.67}$$

Notice that the diagonal elements of the Jacobian matrix in Eq. 2.65 are not equal to β_0, the propagation constant of the isolated waveguides. This means that the presence of the nearby waveguide in fact changes the propagation constant slightly.

Eq. 2.65 can easily be solved by eigenvalue decomposition. The solution for $a_1(z)$ and $a_2(z)$ can be written as a transfer matrix:

$$\begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} = e^{j\beta_0' z} \begin{bmatrix} \cos(\mu' z) & -j\sin(\mu' z) \\ -j\sin(\mu' z) & \cos(\mu' z) \end{bmatrix} \begin{bmatrix} a_1(0) \\ a_2(0) \end{bmatrix} \tag{2.68}$$

If, at $z = 0$ light is launched into waveguide 1, the relative power in the two waveguides as a function of z is:

$$\left| \frac{a_2(z)}{a_1(0)} \right|^2 = \sin^2(\mu' z), \quad \left| \frac{a_1(z)}{a_1(0)} \right|^2 = \cos^2(\mu' z) \quad . \tag{2.69}$$

The above equation illustrates the fact that for two coupled identical waveguides the power slowly sloshes back and forth between them at a rate described by μ'. Interestingly, full power transfer from waveguide 1 to waveguide 2 can be achieved even for weakly coupled waveguides (with arbitrarily small μ'), provided the interaction length z is sufficiently long.

Henceforth, we will drop the prime from the quantities β_0' and μ' in Eq. 2.68. We refer to μ as the “coupling constant” for the structure. It has dimensions of inverse length, and describes the spatial rate at which power transfers between the two waveguides.

2.2.2 Exact Modal Analysis

In Section 2.2.1, we described a variational technique in which we approximated the electromagnetic fields of the parallel waveguide system in terms of a linear superposition of the modes of each constituent waveguide. In principle, it is possible to rigorously compute the electromagnetic modes for the coupled waveguide system using the techniques described in Section 2.1.

The variational technique gives valuable insight into the structure of the eigenmodes for the coupled waveguide system. If we consider the case of identical waveguides, the eigenvalue equation (Eq. 2.59) simplifies to:

$$\begin{bmatrix} \beta_0 & \mu \\ \mu & \beta_0 \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} = \tilde{\beta} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix} \quad , \tag{2.70}$$

where the quantities β_0 and μ are defined in Eqs. 2.66 - 2.67 and $\tilde{\beta}$ represents an eigenvalue of the system of equations. The eigenvalues and corresponding normalized eigenvectors of this equation are,

$$\beta_s = \beta_0 + \mu, \quad \mathbf{a}_s = \frac{1}{\sqrt{2}} \begin{bmatrix} +1 \\ +1 \end{bmatrix} \tag{2.71}$$

$$\beta_a = \beta_0 - \mu, \quad \mathbf{a}_a = \frac{1}{\sqrt{2}} \begin{bmatrix} +1 \\ -1 \end{bmatrix} \tag{2.72}$$

Thus, the approximate eigenmodes of the coupled waveguide system are symmetric and antisymmetric linear combinations of the isolated waveguide modes. [3] The symmetric mode has a propagation constant which is slightly higher than the antisymmetric mode. This gives rise to another physical interpretation of the coupling constant μ: μ describes the splitting between the symmetric and antisymmetric modes of propagation for the parallel waveguide system.

$$\mu = \frac{1}{2}(\beta_s - \beta_a) \tag{2.73}$$

In fact, a more rigorous way to analyze the coupling between parallel waveguides is to directly compute the symmetric and antisymmetric modes. Figure 2.10 illustrates the symmetric and antisymmetric TE modes for two coupled SOI ridge waveguides, of the type depicted in Fig. 2.7.

The coupled mode theory presented earlier agrees very well with the more rigorous direct solution method described here, especially when the waveguide separation becomes large [39]. Moreover, the coupled mode approach has a few advantages over directly solving for symmetric and antisymmetric modes. One limitation of the direct solution method is that the simulations must be repeated if the waveguide separation changes. For the

---

[3] If the two waveguides comprising the coupler are not identical, then the eigenmodes of the coupled system will not have definite symmetry. Nevertheless, the two lowest order modes may be approximated by linear combinations of the isolated waveguide modes according to the eigenvalue equation.

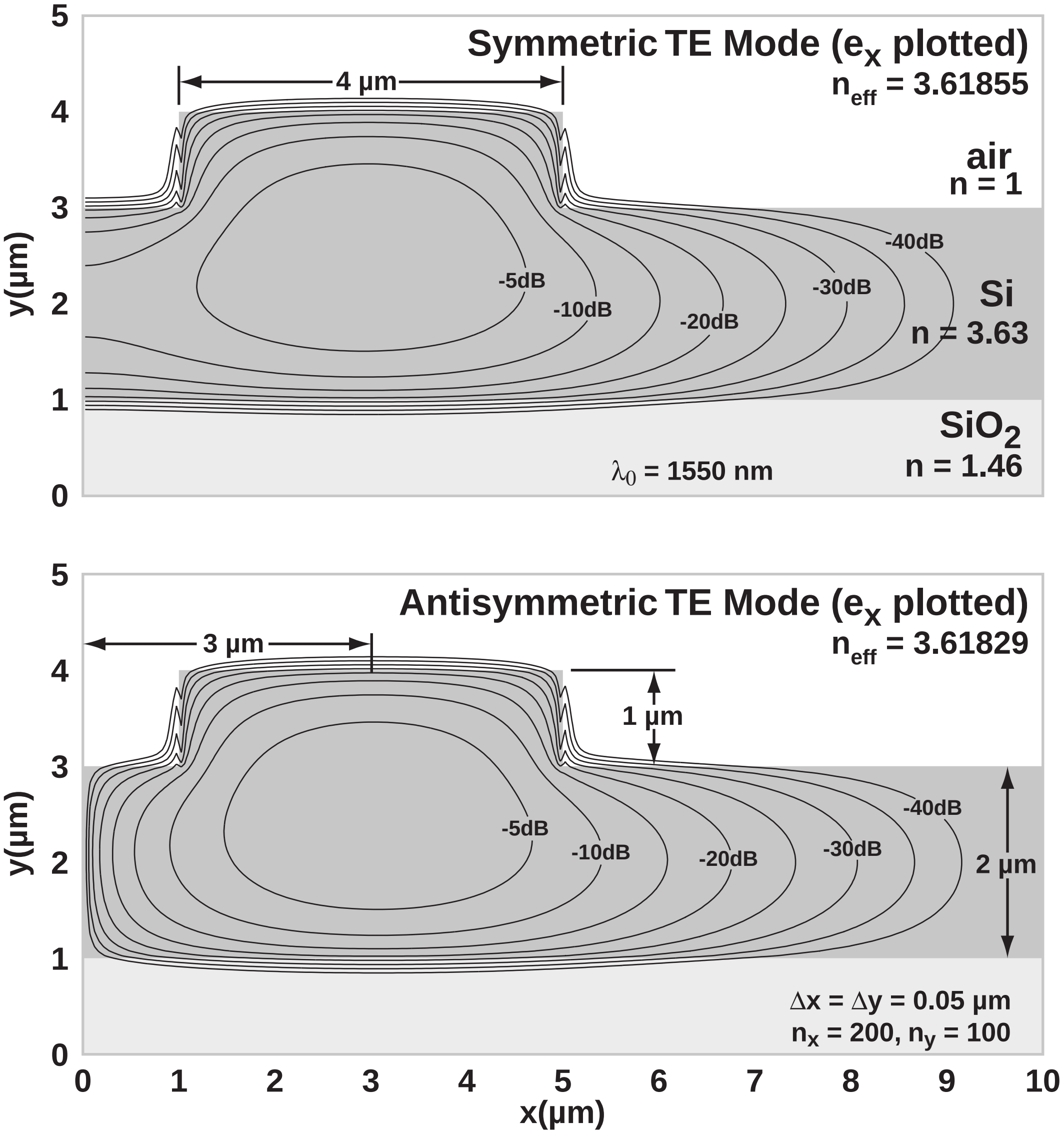

Figure 2.10: Symmetric and antisymmetric TE modes for an SOI ridge waveguide. Note that because of the symmetry, only half of the structure needs to be included in the computation. (figs/2/symmetric-antisymmetric.eps)

coupled mode approach, the isolated waveguide mode only needs to be computed once – different waveguide separations can be analyzed by simply changing the regions of integration for the calculation. Also, the direct calculation of symmetric and antisymmetric modes proves to be numerically challenging as the waveguide separation increases. The method relies on accurately computing the difference between two similar propagation constants (β_s and β_a), which leads to numerical inaccuracies when the two numbers are subtracted. By contrast, the coupled mode approach computes the coupling constant μ by way of an overlap integral which is not susceptible to these problems.

2.2.3 Real Waveguide Couplers

Thus far, we have discussed the problem of parallel waveguides without considering the gradual approach and separation of the two waveguides at the input and output of the coupler. Because most waveguides cannot be bent at a sharp angle without incurring significant loss, any realistic coupler must include a gradual approach and separation. In order to design a coupler with the desired splitting ratio, one must account for the power transfer which occurs in these curved regions.

If the bending of the waveguide is very gradual in comparison to the coupling rate, i.e., if the waveguide separation changes slowly over a length scale of $1/\mu$, the effects of bending may be modeled by simply replacing μ by $\mu(z)$ in Eq. 2.68. This condition is known as the adiabatic condition. In this case, the solution to the coupled system of differential equations is:

$$\begin{bmatrix} a_1(z) \\ a_2(z) \end{bmatrix} = e^{j\beta_0 z} \begin{bmatrix} \cos\phi(z) & -j\sin\phi(z) \\ -j\sin\phi(z) & \cos\phi(z) \end{bmatrix} \begin{bmatrix} a_1(0) \\ a_2(0) \end{bmatrix} \tag{2.74}$$

$$\phi(z) \equiv \int_0^z \mu(z)\,dz \tag{2.75}$$

The above equation is identical to Eq. 2.68, with the exception that μz has been replaced by an integral over the length of the coupler.

A more rigorous analysis of the problem of nonparallel waveguides is given in references [40, 41, 42]. The authors show that when the waveguide dimensions or separation change with position, the couple mode equations of Eq. 2.60 should be replaced by the following modified differential equation:

$$\mathbf{P}\frac{d\mathbf{a}}{dz} + \frac{1}{2}\frac{d\mathbf{P}}{dz}\mathbf{a} = -j\mathbf{H}\mathbf{a} \tag{2.76}$$

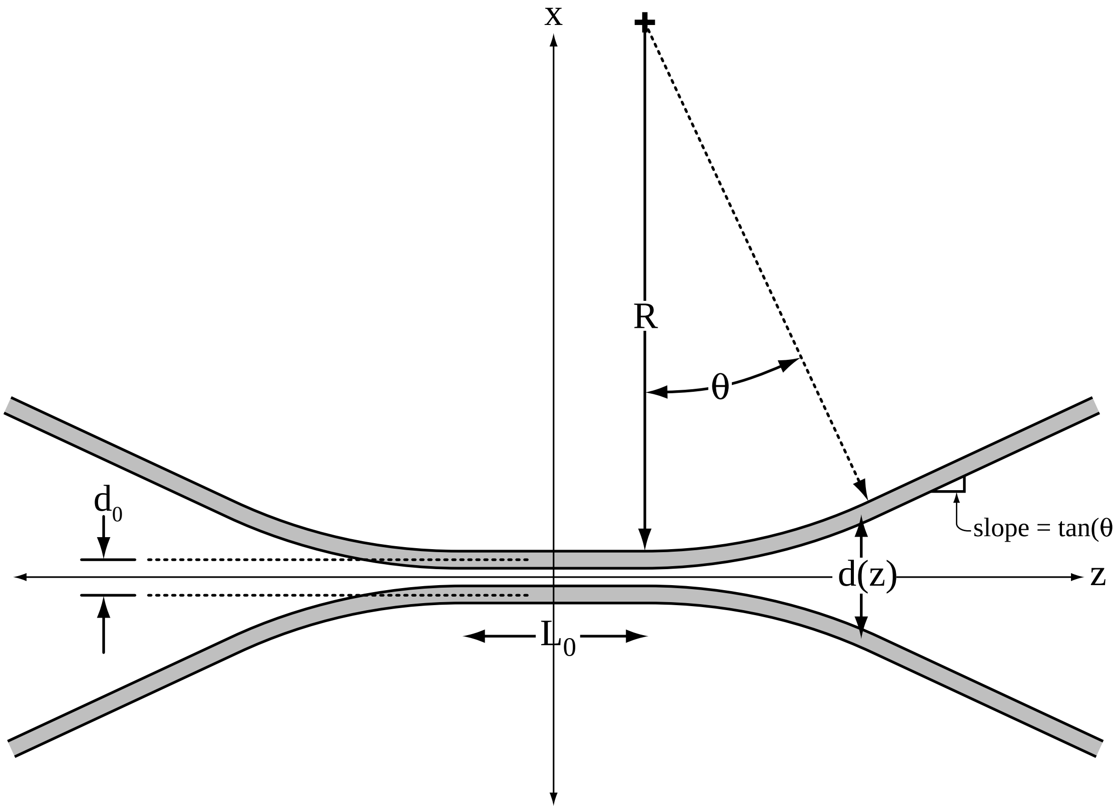

Figure 2.11: Schematic of a more realistic directional coupler, in which two identical waveguides approach each other over a curved path with radius R. (figs/2/real-coupler-schematic.eps)

where the matrices $\mathbf{P}$ and $\mathbf{H}$ are the same as those described earlier, but they are now understood to be slowly varying functions of z. Most couplers have a plane of symmetry such that with the proper choice of origin the elements of $\mathbf{P}$ and $\mathbf{H}$ are even functions of z. In this case, the $\frac{d\mathbf{P}}{dz}$ terms are odd functions of z and therefore integrate to zero when the entire coupler is considered. Therefore, provided the structure is symmetric, Eq. 2.74 can be used to describe the coupling.

Figure 2.11 depicts the geometry of a typical integrated waveguide coupler. Two waveguides initially approach each other along a sloped path. The waveguides gradually become parallel as they traverse an arc of radius R and angle θ, and remain parallel for a length L_0. Then, the waveguides gradually separate, following a symmetric path. Fig-

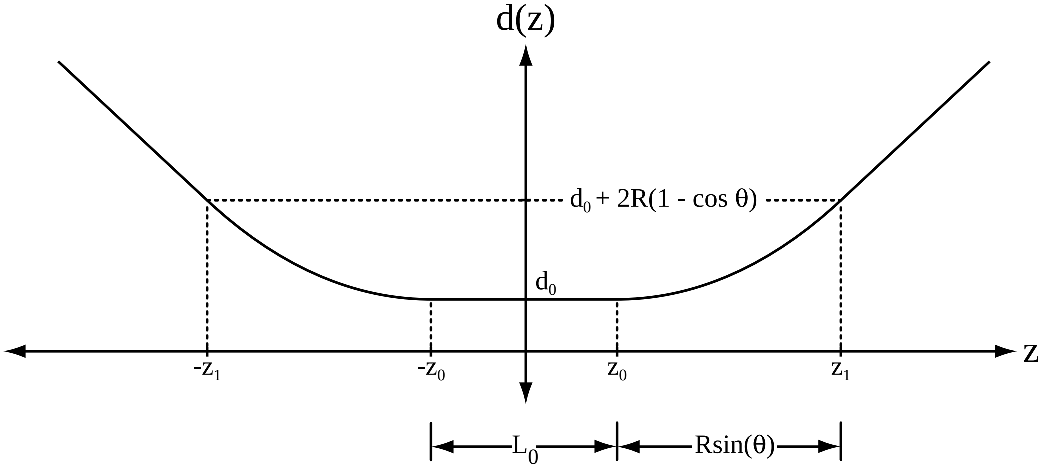

Figure 2.12: Waveguide center-to-center separation as a function of z for the coupler depicted in Fig. 2.11. (figs/2/d-vs-z.eps)

ure 2.12 plots the waveguide separation as a function of separation d for this structure. The waveguide separation can be described mathematically by the following piecewise function:

$$d(z) = \begin{cases} d_0 & |z| < \frac{L_0}{2} \\ d_0 + 2R - 2\sqrt{R^2 - \left(|z| - \frac{L_0}{2}\right)^2} & \frac{L_0}{2} < |z| < \frac{L_0}{2} + R\sin\theta \\ d_0 + 2R(1-\cos\theta) + 2\left(|z| - \frac{L_0}{2} - R\sin\theta\right)\tan\theta & \frac{L_0}{2} + R\sin\theta < |z| \end{cases} \tag{2.77}$$

For the purposes of this analysis, we assume that $d(z)$ continues to grow linearly to ∞ outside of the coupling region. In practice, the coupling constant μ becomes negligibly small as the waveguide separation increases and therefore the exact functional form of $d(z)$ outside of the principal coupling region is unimportant.

One way to treat the structure depicted in Fig. 2.11 is to simply compute the integrated coupling numerically, using Eq. 2.75 along with the calculated coupling constants as a function of d. For example, the coupler may be divided into short segments over which the coupling is approximately constant, and then the integral of Eq. 2.75 could be approximated by a discrete summation.

However, as described below, by making a few simplifying assumptions, one can derive a simple analytical approximation for the total integrated coupling.

In most cases, the coupling constant μ falls exponentially with increasing waveguide separation d. This occurs because the electromagnetic mode decays exponentially outside of the core, and therefore so does the overlap integral of Eq. 2.64. It is often very accurate to model this relationship with the following simple functional form [39, 30]:

$$\mu(d) = A \exp\left(\frac{d}{\tilde{d}}\right) \quad , \tag{2.78}$$

where the coefficient A represents the extrapolated coupling at $d = 0$, and $\tilde{d}$ represents the $1/e$ decay length.

Next, we assume that the waveguide separation in the non-parallel regions can be modeled as a quadratic function:

$$d(z) \simeq \begin{cases} d_0 & |z| < \frac{L_0}{2} \\ d_0 + \frac{1}{R}(z - \frac{L_0}{2})^2 & \frac{L_0}{2} < |z| \end{cases} \tag{2.79}$$

The above equation can be derived by simply performing a Taylor expansion of Eq. 2.77. When this approximation is combined with the exponential decay approximation of Eq. 2.78, the total integrated coupling can be computed analytically:

$$\int_{-\infty}^{+\infty} \mu(z)\, dz = \mu(d_0) \left(L_0 + \sqrt{\pi \tilde{d} R}\right) \tag{2.80}$$

Thus, under certain conditions, the coupling between the two non-parallel waveguides can be treated as if the coupler were comprised of two parallel waveguides separated by d_0 with an effective coupling length given by

$$L_{\text{eff}} = L_0 + \sqrt{\pi \tilde{d} R} \tag{2.81}$$

as described in [39], this approximation is valid under the following condition:

$$\frac{R}{\tilde{d}} \sin^2 \theta \gg 1 \tag{2.82}$$

2.2.4 Results of Coupled Mode Analysis

Figure 2.13 plots the coupling constant μ as a function of waveguide separation d for the doped-glass channel waveguide structure considered in this work. Notice that the relationship between μ and d is approximately exponentially decreasing, as evidenced by the linear slope of the curves when plotted on semilogarithmic axes. There is a slight deviation from the exponential decay when the waveguides are very close together.

Also, it is worth pointing out that there is a slight discrepancy between the TE and TM coupling constants, but the difference is too small to be visible in Fig. 2.13.

Figure 2.14 plots the coupling constant μ as a function of waveguide separation for the silicon-on-insulator ridge waveguide considered in this work. Notice that because of the stronger index contrast between the core and cladding, the SOI waveguide has a much more severe polarization dependence than the doped-glass channel waveguide.

2.2.5 Design of Insensitive Couplers

So far, we have described the coupling between parallel waveguides, and explained how to accurately treat the coupling between nonparallel waveguides. Together with the eigenmode analysis of Section 2.1, these techniques allow one to design an integrated coupler with the desired power splitting ratio. However, these calculations can only accurately predict the coupling at one specific wavelength and for one polarization state. If the wavelength and polarization change, the eigenmodes of the waveguides also change, and therefore so does the coupling constant μ. It is often impossible to achieve broadband polarization-insensitive performance, especially in planar integrated devices.

Figure 2.15 depicts the calculated power splitting ratio for an integrated glass waveguide coupler, as a function of wavelength. This particular coupler was designed to have a power splitting ratio of 50% at a free-space wavelength of 1550 nm, but over a span of 100 nm the calculated coupling varies from 40 to 60%.

Even in applications where the wavelength and polarization are well-defined and carefully controlled, small deviations or nonuniformities in the fabrication process, material properties, or device operating conditions can significantly alter the power splitting ratio.

We now describe an improved design for a directional coupler which yields a power splitting ratio that is insensitive to wavelength, polarization, and fabrication parameters.

Others have achieved wavelength-insensitive performance by using tapered waveguide couplers [43, 44] or asymmetric waveguide couplers [45, 46, 47]. Although these approaches can greatly reduce the wavelength dependence, they do not explicitly account for polarization and fabrication uncertainties. Another approach to achieving insensitive

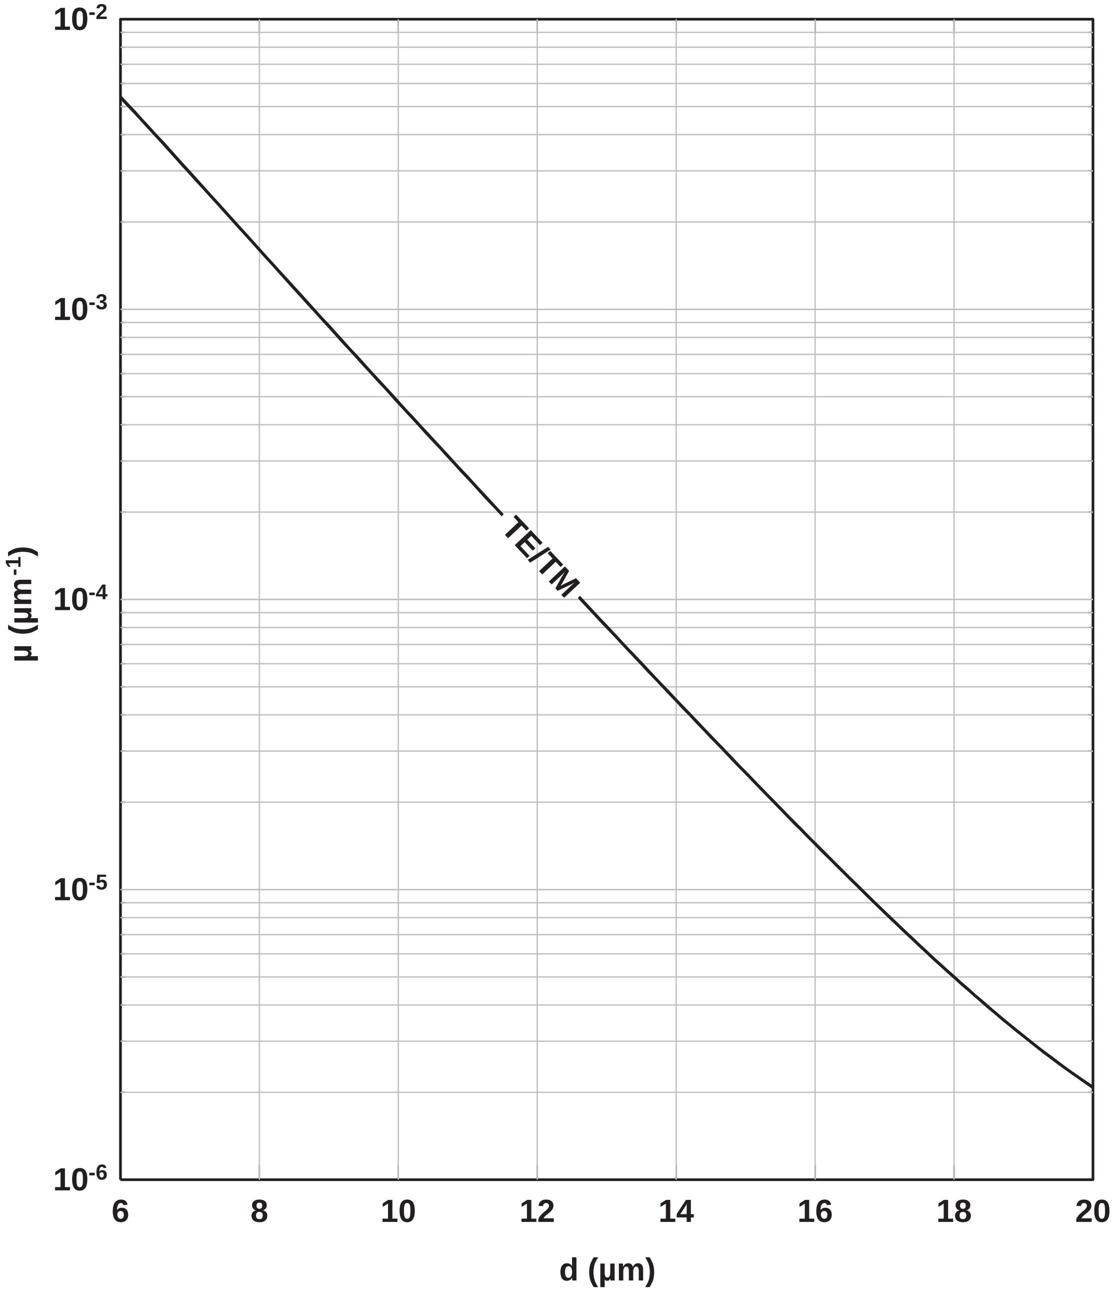

Figure 2.13: Coupling constant $\mu(z)$ as a function of d for a doped-glass channel waveguide. These data were calculated using the variational theory describe earlier.

(figs/2/mu-vs-d-glass.eps)

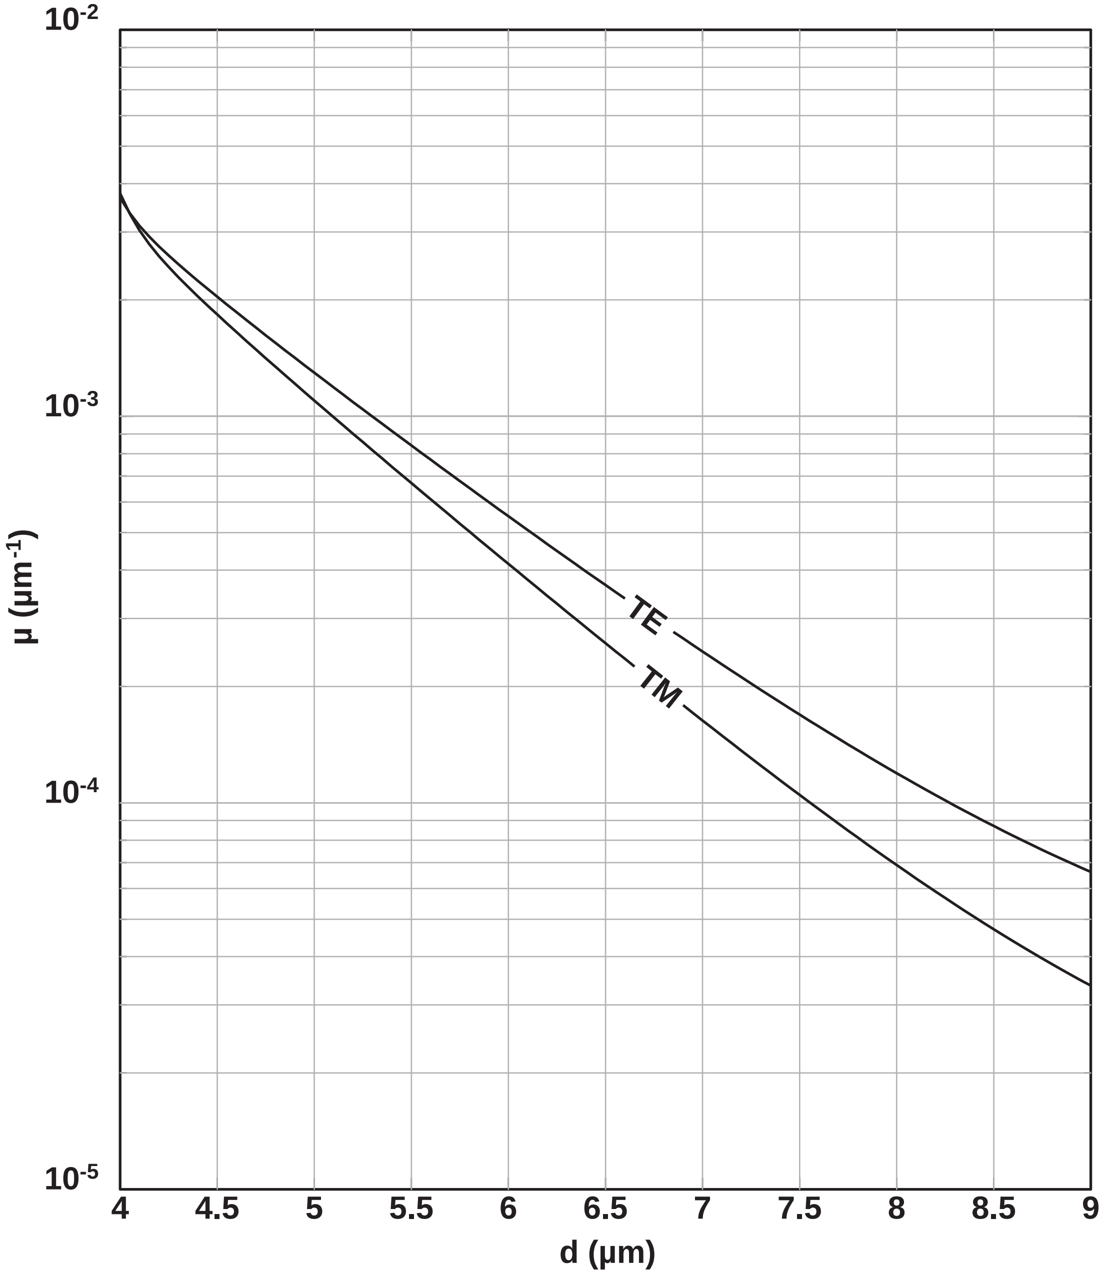

Figure 2.14: Coupling constant $\mu(z)$ as a function of d for an SOI ridge waveguide

(figs/2/mu-vs-d-soi.eps)

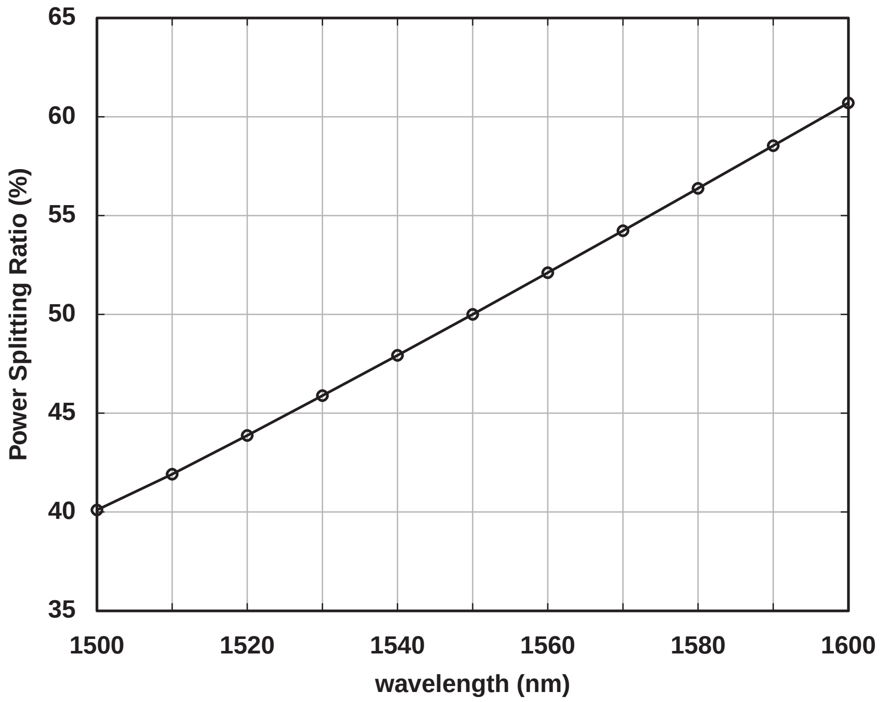

Figure 2.15: Calculated splitting ratio as a function of wavelength for an integrated doped-glass waveguide coupler. The coupler was designed to have a power splitting ratio of 50% at a free-space wavelength of 1550 nm. (figs/2/coupler-dispersion.eps)

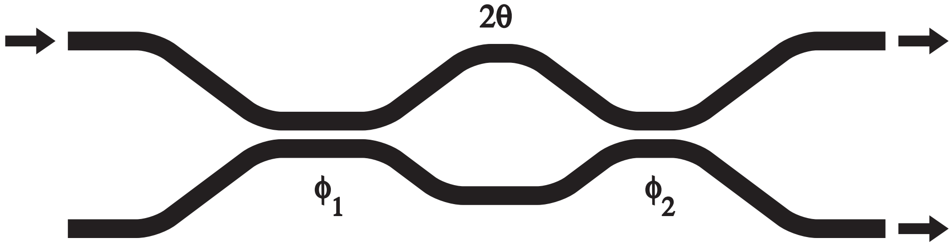

Figure 2.16: Structure of a parameter-insensitive directional coupler: two conventional couplers are cascaded in a Mach-Zehnder configuration, with a relative phase shift introduced between them. (figs/2/mz-coupler-structure.eps)

performance is to use adiabatic couplers [48], but this approach can require very long structures, and the application to couplers with arbitrary splitting ratio is not obvious.

Another method for achieving wavelength insensitive performance is to use a cascade of two similar couplers in a Mach-Zehnder configuration [49, 50, 51]. Using numerical optimization techniques in combination with empirically determined wavelength-dependence data, researchers at NTT have employed this technique to build broadband 20% couplers [51]. By contrast, we describe here a simple set of universal design rules for constructing parameter-insensitive cascaded couplers with any desired splitting ratio.

Figure 2.16 depicts the structure of the parameter-insensitive coupler. The device consists of a cascade of two conventional couplers, with an intervening phase shift. The dimensionless quantities ϕ_1 and ϕ_2 represent the total integrated coupling of the two constituent directional couplers,

$$\phi_1 \equiv \int_{(1)} \mu(z)\, dz, \quad \phi_2 \equiv \int_{(2)} \mu(z)\, dz \tag{2.83}$$

and 2θ is the relative phase difference between the two arms of the structure. This phase shift can be achieved by making one arm slightly longer than the other, or through some other thermo-optic or electro-optic effect.

The cascaded coupler depicted in Fig. 2.16, can be analyzed using a transfer matrix approach, which will be described completely in Section 2.4. For now, we will simply give the result for the power splitting ratio. The splitting ratio, defined as the fractional amount of power transferred from guide 1 to guide 2, is given by [52, 51]:

$$S = \cos^2\theta \sin^2(\phi_1 + \phi_2) + \sin^2\theta \sin^2(\phi_1 - \phi_2) \tag{2.84}$$

each of the three parameters used in Eq. 2.84 implicitly depends on the wavelength, polarization, material properties and dimensions of the device. However, the deviations in ϕ_1 and ϕ_2 are usually highly correlated with each other. By connecting the couplers with an appropriate phase shift, these variations can be made to counteract each other such that the resulting power splitting ratio remains unchanged.

Compared with other methods of achieving wavelength-insensitive performance, this approach has the advantage that it is easily adaptable to any waveguide geometry. Rather than attempting to numerically model the coupling dispersion or polarization dependence, as was done in reference [51], this technique instead treats all potential variations as fractional changes in ϕ_1 and ϕ_2. Accordingly, the device should be capable of compensating for any uncontrolled variation (structural or optical) which affects both constituent couplers proportionately. The variations in ϕ_1 and ϕ_2 tend to be more severe than variations in θ because of the evanescent nature of the waveguide coupling.

As described in a previous work [52], insensitive performance can be achieved by choosing the three parameters in the following way:

$$\phi_1 = \frac{3\pi}{8}\left(1 + \frac{1}{N}\right), \quad \phi_2 = \frac{3\pi}{8}\left(1 - \frac{1}{N}\right) \tag{2.85}$$

$$\cos^2\theta = \sin\left(\frac{3\pi}{2N}\right)\left[N + \sin\left(\frac{3\pi}{2N}\right)\right]^{-1} \tag{2.86}$$

where N is a dimensionless real parameter larger than 3 which can be chosen to give any desired splitting ratio as described in (2.84). Figure 2.17 shows how to select the three parameters in order to achieve a prescribed splitting ratio. In the particularly important case where 50% splitting is desired, we choose $\phi_1 = \pi/2$, $\phi_2 = \pi/4$, and $\theta = \pi/3$.

Figure 2.18 illustrates how the coupler compensates for fractional changes in ϕ_1 and ϕ_2. For these calculations, ϕ_1, ϕ_2 and θ were nominally chosen so as to give 50% splitting. Along the y-axis, we plot the actual power splitting ratio as a function of the fractional change $\Delta\phi/\phi$. Clearly, for the cascaded Mach-Zehnder structure, the first order dependence of splitting ratio on $\Delta\phi/\phi$ has been eliminated.

Having specified ϕ_1, ϕ_2 and θ, the device structure must be selected to achieve the desired accumulated coupling and phase shift. This task is no more complicated than the

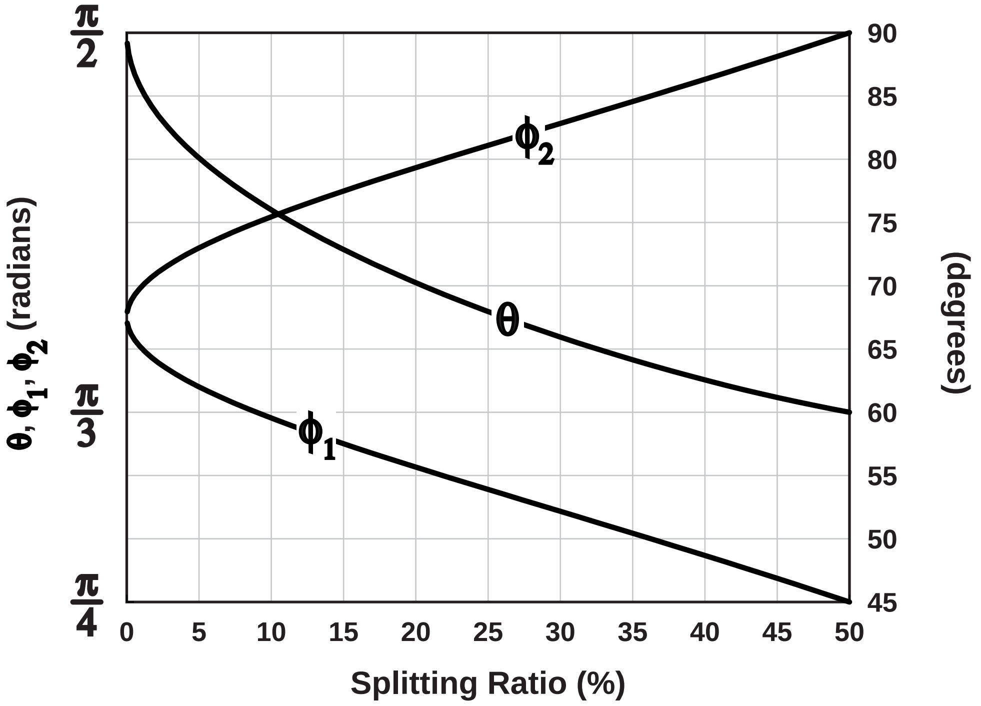

Figure 2.17: Plot describing how the three dimensionless parameters ϕ_1, ϕ_2 and θ of the coupler should be selected in order to achieve insensitive performance for any desired splitting ratio. (figs/2/insensitive-coupler-parameters.eps)

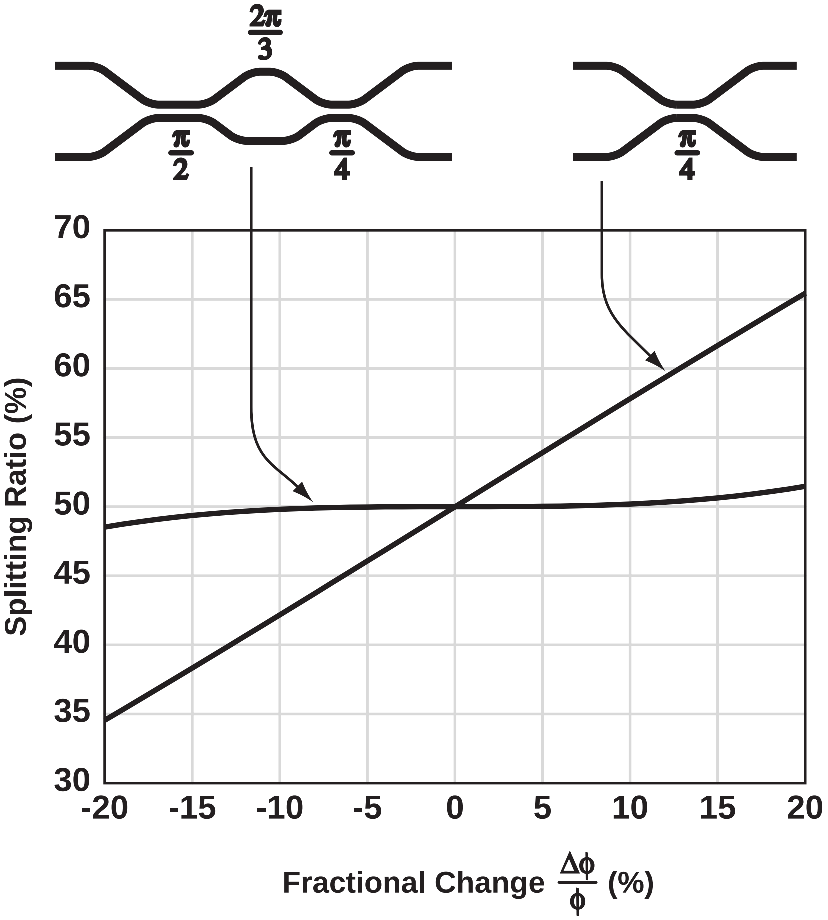

Figure 2.18: Plot comparing a conventional directional coupler to the improved Mach-Zehnder structure. Along the vertical axis, we plot the power splitting ratio, and horizontal axis represents a fractional change in the parameters ϕ_1 and ϕ_2. (figs/2/insensitive-coupler-theory.eps)

design of a conventional directional coupler.

2.3 Bragg Gratings

In the previous section, we developed a theory to describe the interaction between two copropagating modes of adjacent waveguides. We now turn to the problem of contradirectional coupling, the coupling between forward-traveling and backward-traveling modes. By introducing a periodic modulation in a waveguide, it is possible to create an interaction between the otherwise decoupled forward-traveling and backward-traveling modes of an optical waveguide.

Wave propagation in a periodic medium is a well studied phenomenon in the field of solid-state physics. The electronic bandgap which exists in semiconductors arises because the electrons move in a periodic medium defined by the crystal lattice. The diffraction of x-rays by a crystal lattice is another example of wave propagation in a periodic structure [53]. The problem of electromagnetic propagation in a periodic waveguide is a similar phenomenon, with the exception that the light is confined in the transverse directions so that the interactions occur only in one dimension.

There are many ways to induce a periodic modulation in an integrated waveguide. In some cases, it is possible to directly modulate the refractive index of the guiding layer, by illuminating the waveguide with ultraviolet light[54, 55, 56], or by utilizing acousto-optical or electro-optical effects. These techniques, however, can only be used in materials which have photorefractive, acousto-optic, or electro-optic properties. A more flexible way of introducing a periodic modulation is to physically corrugate the waveguide structure using the techniques of lithography and pattern transfer. This is the approach considered in this thesis.

Before proceeding, let us clarify the relationship between forward-traveling modes and backward-traveling modes. Recall that the eigenmode equation for an optical waveguide, derived in Section 2.1, has eigenvalues which are β^2, the propagation constant squared. Positive values of β correspond to forward-traveling modes, while negative values correspond to backward-traveling modes. That the eigenvalue equation depends only upon β^2 seems to indicate that the electromagnetic mode profiles are identical for forward- and backward-traveling modes. However, the eigenmode equations derived earlier only govern two components of the transverse electromagnetic field. The remaining four field components can be derived from these two transverse components, but the resulting expressions depend upon the sign of β. The relationship between the n^{th} forward-traveling mode and corresponding backward-traveling mode can be expressed as:

$$\mathbf{e}_{-nt} = \mathbf{e}_{nt} \qquad \mathbf{h}_{-nt} = -\mathbf{h}_{nt} \tag{2.87}$$

$$e_{-nz} = -e_{nz} \qquad h_{-nz} = h_{nz} \qquad . \tag{2.88}$$

We conventionally use negative subscripts to indicate backward-traveling modes and positive subscripts to indicate forward-traveling modes. It is easy to verify that although the transverse electric field components are identical for the forward and backward modes, the two modes are orthogonal,

$$\iint (\mathbf{e}_+^* \times \mathbf{h}_- + \mathbf{e}_- \times \mathbf{h}_+^*) \cdot \hat{\mathbf{z}}\, dx\, dy = 0 \quad . \tag{2.89}$$

2.3.1 Bragg Grating Basics

Figure 2.19 depicts the type of structure which we seek to analyze. For most of the examples presented in this chapter, we will consider rectangular grooves etched into the surface of an optical waveguide [4], although the analytical techniques can be applied to any periodic modulation.

The Bragg grating may be thought of as a one-dimensional diffraction grating which diffracts light from the forward-traveling mode into the backward-traveling mode. The condition for diffraction into the reverse traveling mode is called the Bragg condition. In order for light to be efficiently diffracted in the opposite direction, the reflections from subsequent periods of the grating must interfere constructively. This means that the Bragg period Λ must be related to the free space wavelength λ_0 by:

$$\Lambda = \frac{\lambda_0}{2n_{\text{eff}}} \quad , \tag{2.90}$$

where n_{eff} is the effective index of refraction of the structure, which depends upon the materials comprising the waveguide. Generally speaking, for fiber-optic networks the desired operating wavelength is approximately 1550 nm. Table 2.1 lists the appropriate Bragg grating periods for a few common waveguide materials.

Eq. 2.90 assumes that the phase accumulation between reflections from adjacent grat-

---

[4] Most lithographic techniques produce gratings with such a rectangular profile, because of the directional etching processes used.

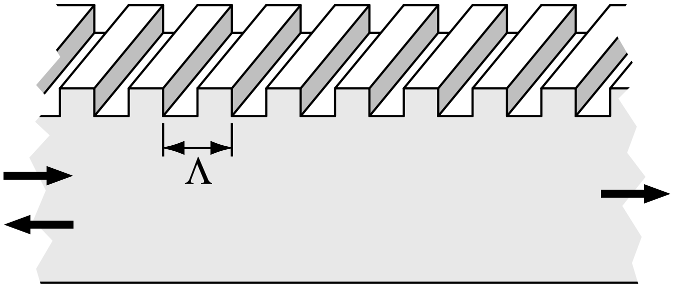

Figure 2.19: A diagram illustrating the type of Bragg grating considered in this work. A grating, with spatial period Λ is formed by physically corrugating one surface of the waveguide. The Bragg grating acts as a one-dimensional diffraction grating which reflects light from forward-traveling mode into the backward-traveling mode. (figs/2/rectangular-grating.eps)

Material	Si	InP / InGaAsP	SiO_2
n **(refractive index)**	3.5	3.17	1.46
Bragg grating period Λ	220 nm	245 nm	535 nm

Table 2.1: Bragg periods Λ for some common waveguide materials, assuming a free-space operating wavelength of 1550 nm.

ing teeth is precisely one wavelength. This is the condition for a *first-order* Bragg grating. It is also possible to utilize higher-order diffraction to couple to the forward and backward modes. The more general condition for constructive interference is that the phase accumulation between subsequent reflections must be an integral number of wavelengths. The Bragg condition for an m^{th}-order Bragg grating is:

$$\Lambda = \frac{m\lambda_0}{2n_{\text{eff}}} \quad . \tag{2.91}$$

In this work, we consider only first-order ($m = 1$) Bragg gratings, because the diffraction efficiency (and hence the contradirectional coupling strength) is generally strongest for the first diffracted order. Nevertheless, some people build higher-order Bragg gratings simply because the required grating pitch is substantially larger, which simplifies the fabrication process.

2.3.2 Contradirectional Coupled Mode Theory

We now turn to the practical question of how to theoretically model a Bragg grating. The most common way of analyzing Bragg gratings is to use a coupled mode formulation [57, 8, 36, 58]. In this approach, the grating is modeled as a small perturbation which produces a coupling between the forward and backward modes of an unperturbed waveguide. It is important to realize, however, that Bragg grating structures do not have characteristic eigenmodes in the same way that a waveguide does. A better description of the physics of Bragg gratings is to describe the electromagnetic fields as *Bloch waves* or *Floquet modes*, a technique similar to that used to describe electron propagation in a crystal lattice. A few authors have used this approach to describe wave propagation in a Bragg grating [59, 60, 61]. The analysis of B. E. Little et al. clarifies the relationship between the Floquet analysis and the more simple coupled mode analysis [60]. There are others who have analyzed Bragg gratings using a rigorous full-wave numerical analysis.

Often the periodic modulation in the waveguide can be accurately treated as a small perturbation of an otherwise z-invariant structure. Figure 2.20 illustrates how a corrugated waveguide can be separated into a z-invariant structure with well-defined modes plus a periodic perturbation. The goal of coupled mode theory is to express the electromagnetic fields of the complete structure as a superposition of the fields of the unperturbed waveguide.

To begin, we shall denote the refractive index profile for the complete structure by $n^2(x, y, z)$, and the refractive index profile of the unperturbed waveguide by $n_0^2(x, y)$.

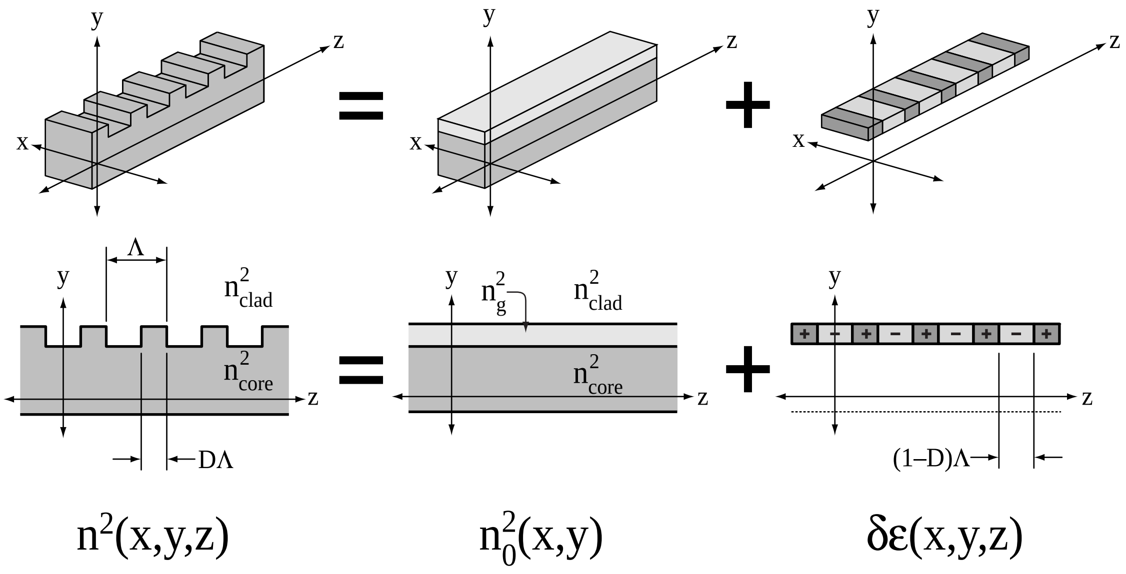

Figure 2.20: The Bragg grating corrugation may be modeled as a perturbation of an otherwise z-invariant waveguide, as shown in this diagram. $n^2(x,y,z)$ represents the full refractive index profile of the structure, $n_0^2(x,y)$ represents the unperturbed z-invariant waveguide, and $\delta\epsilon(x,y,z)$ is the perturbation function. Notice that the perturbation $\delta\epsilon(x,y,z)$ is non-zero only within the grating region. (figs/2/grating-perturbation.eps)

$n^2(x,y,z)$ and $n_0^2(x,y)$ are related by

$$n^2(x,y,z) = n_0^2(x,y) + \delta\epsilon(x,y,z) \quad , \tag{2.92}$$

where the quantity $\delta\epsilon(x,y,z)$ is a periodic function of z which represents the corrugation. This function is nonzero only within the grating cross-section.

Likewise, we denote the electromagnetic fields of the complete structure by $\underline{\mathbf{E}}(x,y,z)$ and $\underline{\mathbf{H}}(x,y,z)$, and the electromagnetic fields of the unperturbed waveguide by $\underline{\mathbf{E}}_0\,(x,y,z)$ and $\underline{\mathbf{H}}_0(x,y,z)$:

$$n_0^2(x,y,z) \quad \longrightarrow \quad \underline{\mathbf{E}}_0(x,y,z),\ \underline{\mathbf{H}}_0(x,y,z) \tag{2.93}$$

$$n^2(x,y,z) \quad \longrightarrow \quad \underline{\mathbf{E}}(x,y,z),\ \underline{\mathbf{H}}(x,y,z) \quad . \tag{2.94}$$

The fields $\underline{\mathbf{E}}$, $\underline{\mathbf{H}}$, $\underline{\mathbf{E}}_0$, and $\underline{\mathbf{H}}_0$, satisfy Maxwells equations in their respective index profiles, i.e.,

$$\nabla \times \underline{\mathbf{E}} = -jk\eta_0 \underline{\mathbf{H}} \qquad \nabla \times \underline{\mathbf{E}}_0 = -jk\eta_0 \underline{\mathbf{H}}_0 \tag{2.95}$$

$$\nabla \times \underline{\mathbf{H}} = jk\frac{1}{\eta_0} n^2 \underline{\mathbf{E}} \qquad \nabla \times \underline{\mathbf{H}}_0 = jk\frac{1}{\eta_0} n_0^2 \underline{\mathbf{E}}_0 \quad . \tag{2.96}$$

Now consider the vector quantity $\mathbf{S}$, defined as

$$\mathbf{S} \equiv \underline{\mathbf{E}}_0^* \times \underline{\mathbf{H}} + \underline{\mathbf{E}} \times \underline{\mathbf{H}}_0^* \quad . \tag{2.97}$$

Computing the divergence of $\mathbf{S}$, using Eqs. 2.95-2.96 gives

$$\nabla \cdot \mathbf{S} = -\frac{jk}{\eta_0} \delta\epsilon \, \underline{\mathbf{E}}_0^* \cdot \underline{\mathbf{E}} \quad . \tag{2.98}$$

If we integrate Eq. 2.98 over a volume V, and apply the divergence theorem to the left-hand side, we obtain,

$$\iint_{\partial V} \mathbf{S} \cdot \mathbf{dA} = -\frac{jk}{\eta_0} \iiint_V \delta\epsilon \, \underline{\mathbf{E}}_0^* \cdot \underline{\mathbf{E}} \, dV \quad . \tag{2.99}$$

This relationship is known as the *Lorentz reciprocity theorem*. The integral on the left-hand side is a surface integral, where $\mathbf{dA}$ points in the outward direction. If we take the volume V to be an infinitesimally thin slab of width dz which spans the entire x-y plane, Eq. 2.99

becomes:

$$\frac{d}{dz}\iint \mathbf{S}\cdot\hat{\mathbf{z}}\,dA = -\frac{jk}{\eta_0}\iint \delta\epsilon\,\underline{\mathbf{E}}_0^*\cdot\underline{\mathbf{E}}\,dA \quad , \tag{2.100}$$

where now both integrals are taken over the x-y plane. To simplify the subsequent analysis, we shall use brackets to indicate integration over the x-y plane:

$$\langle f(x,y)\rangle \equiv \iint f(x,y)\,dx\,dy \quad . \tag{2.101}$$

At this point, no approximations have been made. The Lorentz reciprocity theorem described in Eq. 2.100 forms the basis for coupled-mode theory, as we shall describe below.

Now, we attempt to express the electromagnetic fields of the complete structure as a linear superposition of the modes of the unperturbed waveguide.

$$\underline{\mathbf{E}}(x,y,z) = A_+(z)\mathbf{e}_+(x,y) + A_-(z)\mathbf{e}_-(x,y) \tag{2.102}$$

$$\underline{\mathbf{H}}(x,y,z) = A_+(z)\mathbf{h}_+(x,y) + A_-(z)\mathbf{h}_-(x,y) \tag{2.103}$$

The electromagnetic fields $\mathbf{e}_\pm(x,y)$ and $\mathbf{h}_\pm(x,y)$ represent the modes of the unperturbed waveguide at the optical frequency of interest. The goal of coupled mode theory is to replace Maxwells equations for $\underline{\mathbf{E}}$ and $\underline{\mathbf{H}}$ by a set of coupled ordinary linear differential equations which describe the evolution of the scalar coefficients $A_\pm(z)$. The Lorentz reciprocity relation of Eq. 2.100 provides the necessary link between the electromagnetic fields and the expansion coefficients.

First, let us take $\underline{\mathbf{E}}_0(x,y,z)$ and $\underline{\mathbf{H}}_0(x,y,z)$ in Eq. 2.100 to be the forward-traveling mode of the unperturbed waveguide, i.e.,

$$\underline{\mathbf{E}}_0(x,y,z) = \mathbf{e}_+(x,y)e^{-j\beta_0 z}, \quad \underline{\mathbf{H}}_0(x,y,z) = \mathbf{h}_+(x,y)e^{-j\beta_0 z} \quad . \tag{2.104}$$

Substituting these fields into Eqs. 2.100, along with the mode expansion of Eqs. 2.102-2.103, leads to the following differential equation for $A_+(z)$:

$$\frac{d}{dz}A_+(z) = -j\left[\beta_0 + \frac{k}{4P\eta_0}\langle\delta\epsilon\,\mathbf{e}_+^*\cdot\mathbf{e}_+\rangle\right]A_+(z) - \frac{jk}{4P\eta_0}\langle\delta\epsilon\,\mathbf{e}_+^*\cdot\mathbf{e}_-\rangle A_-(z) \quad , \tag{2.105}$$

where P denotes the time-averaged power carried by the forward-traveling mode,

$$P \equiv \frac{1}{4}\langle(\mathbf{e}_+\times\mathbf{h}_+^* + \mathbf{e}_+^*\times\mathbf{h}_+)\cdot\hat{\mathbf{z}}\rangle \quad . \tag{2.106}$$

A similar equation may be derived for $A_-(z)$ by taking $\underline{\mathbf{E}}_0(x,y,z)$ and $\underline{\mathbf{H}}_0(x,y,z)$ to be the reverse-traveling mode. We omit the details and merely present the result:

$$\frac{d}{dz}A_-(z) = +j\left[\beta_0 + \frac{k}{4P\eta_0}\langle\delta\epsilon\,\mathbf{e}_-^*\cdot\mathbf{e}_-\rangle\right]A_-(z) + \frac{jk}{4P\eta_0}\langle\delta\epsilon\,\mathbf{e}_-^*\cdot\mathbf{e}_+\rangle A_+(z) \quad . \tag{2.107}$$

Eq. 2.105 and Eq. 2.107 constitute a pair of coupled linear differential equations which relate the forward- and backward- mode amplitudes. The terms inside angle brackets are simple overlap integrals of the mode with the grating perturbation $\delta\epsilon$. Note, however, that even though these integrals are taken over the x-y plane, the integrand is proportional to $\delta\epsilon$, which is a periodic function of z.

The coupled-mode equations can be simplified by expanding $\delta\epsilon(x,y,z)$ as a Fourier series in z. The Fourier series expansion for $\delta\epsilon$ is:

$$\delta\epsilon(x,y,z) = \sum_m \delta\epsilon_m(x,y)e^{-j2m\pi z/\Lambda} \quad . \tag{2.108}$$

Each Fourier coefficient in Eq. 2.108 is a function of x and y only. Eq. 2.108 can be inverted to obtain an expression for the Fourier coefficients in terms of $\delta\epsilon(x,y,z)$:

$$\delta\epsilon_m(x,y) = \frac{1}{\Lambda}\int_{z_0}^{z_0+\Lambda}\delta\epsilon(x,y,z)e^{j2m\pi z/\Lambda}\,dz \quad . \tag{2.109}$$

Provided the periodic perturbation is real-valued, the Fourier coefficients are related by $\delta\epsilon_{-m} = \delta\epsilon_m^*$.

Substituting the expansion of Eq. 2.108 into the coupled mode equations yields:

$$\begin{aligned}\frac{d}{dz}A_+(z) = -j\left[\beta_0 + \frac{k}{4P\eta_0}\langle\delta\epsilon_0\,\mathbf{e}_+^*\cdot\mathbf{e}_+\rangle\right]A_+(z) \\ - \frac{jk}{4P\eta_0}\langle\delta\epsilon_1\,\mathbf{e}_+^*\cdot\mathbf{e}_-\rangle e^{-j2\pi z/\Lambda}A_-(z) + \ldots\end{aligned} \tag{2.110}$$

$$\begin{aligned}\frac{d}{dz}A_-(z) = +j\left[\beta_0 + \frac{k}{4P\eta_0}\langle\delta\epsilon_0\,\mathbf{e}_-^*\cdot\mathbf{e}_-\rangle\right]A_-(z) \\ + \frac{jk}{4P\eta_0}\langle\delta\epsilon_{-1}\,\mathbf{e}_-^*\cdot\mathbf{e}_+\rangle e^{+j2\pi z/\Lambda}A_+(z) + \ldots\end{aligned} \tag{2.111}$$

In the above equation, we have only retained those terms which are *phase-matched*. Specifically, we have omitted terms from the Fourier expansion which have a spatial periodicity much different from the accompanying terms in the equations. For example, when the function $A_-(z)$ is multiplied by the $m = 1$ Fourier coefficient, $\delta\epsilon_1$, the resulting product has roughly the same periodicity as $A_+(z)$. Likewise, when the function $A_+(z)$ is multiplied by the $m = -1$ Fourier coefficient the resulting function has the same spatial periodicity as $A_-(z)$. The omitted Fourier expansion terms do not meet the phase-matching condition, meaning they oscillate with spatial frequencies which are much different from $e^{\pm j\pi z/\Lambda}$. When the coupled-mode equations are solved by integrating over z, these rapidly-oscillating terms quickly average to zero, whereas the retained terms do not.

As shown in Eq. 2.110, for a first-order Bragg grating only the $m = 0$ and $m = \pm 1$ Fourier coefficients enter into the coupled-mode equations. The $m = 0$ term induces a change in the propagation constant, i.e. it modifies the diagonal elements of the coupling matrix, while the $m = \pm 1$ terms couple the forward and backward modes. We may cast the coupled mode equations into a simpler form by defining two parameters β and κ, each involving overlap integrals between the unperturbed waveguide modes and the grating perturbation:

$$\beta \equiv \beta_0 + \frac{k}{4P\eta_0}\langle \delta\epsilon_0 \, \mathbf{e}_\pm^* \cdot \mathbf{e}_\pm \rangle \tag{2.112}$$

$$\kappa \equiv -\frac{jk}{4P\eta_0}\langle \delta\epsilon_1 \, \mathbf{e}_+^* \cdot \mathbf{e}_- \rangle \quad . \tag{2.113}$$

With these definitions, the coupled mode equations can be summarized in matrix form:

$$\frac{d}{dz}\begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} = \begin{bmatrix} -j\beta & \kappa e^{-j2\pi z/\Lambda} \\ \kappa^* e^{+j2\pi z/\Lambda} & j\beta \end{bmatrix} \begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} \quad . \tag{2.114}$$

The magnitude of the complex coupling constant κ describes the rate at which power transfers between the forward and backward modes. κ is often described as the “reflectivity per unit length” of the grating. The phase of κ is directly related to the phase of $\delta\epsilon_1$. For a grating which is symmetric about $z = 0$, all of the Fourier coefficients will be real, and therefore κ will be purely imaginary. Even for gratings which do not have this symmetry, the coupled mode equations are of the form given in Eq. 2.114, provided the refractive index perturbation is real. [5]

By choosing the unperturbed waveguide appropriately, the $m = 0$ Fourier coefficient can be eliminated. This is accomplished by choosing $n_0^2(x, y)$ to be the z-averaged value of

---

[5] If the grating includes a modulation in the optical gain or loss, the perturbation can no longer be modeled by a real-valued function, and the coupled mode equations will be different from those given in Eq. 2.114. [62, 63, 64]

$n^2(x,y,z)$:

$$n_0^2(x,y) = \frac{1}{\Lambda}\int_0^{\Lambda} n^2(x,y,z)\,dz \quad . \tag{2.115}$$

It is easy to verify that with this choice, $\delta\epsilon_0(x,y) = 0$, and there is no change in the propagation constant as a result of the perturbation.

If we consider the limit as $\kappa \to 0$, the equations decouple and the solutions for $A_\pm(z)$ are of the form $e^{\mp j\beta z}$, as expected. Often, it is more convenient to factor out these rapid oscillations from $A_\pm(z)$ to obtain coupled-mode equations for the slowly-varying envelope functions. The rapid oscillations are removed by making the following change of variables:

$$A_\pm(z) = a_\pm(z) e^{\mp j\pi z/\Lambda} \quad . \tag{2.116}$$

With this substitution, the coupled mode equations simplify to

$$\frac{d}{dz}\begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} = \begin{bmatrix} -j\delta & \kappa \\ \kappa^* & j\delta \end{bmatrix}\begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} \quad , \tag{2.117}$$

where δ is a measure of the deviation from the Bragg condition, given by

$$\delta \equiv \beta - \frac{\pi}{\Lambda} \quad . \tag{2.118}$$

2.3.3 Spectral Response of Bragg Grating

Bragg gratings are useful because of their frequency-dependent nature. We will now examine this frequency dependence by solving the coupled mode equations which were developed in the previous section.

In the analysis presented earlier, we tacitly assumed that all field quantities were oscillating at the same optical frequency $\omega = kc$. The quantity β_0 and the associated fields $\mathbf{e}_\pm$, and $\mathbf{h}_\pm$ were understood to be the propagation constant and electromagnetic modes for the unperturbed waveguide *at the optical frequency* ω. We now asked the question, what happens when the optical frequency changes?

Strictly speaking, the eigenfunctions and propagation constants must be recomputed using the new value for k. Then, the overlap integrals in Eqs. 2.112 - 2.113 must be recomputed using the new modes. In order to numerically compute the spectral response for a Bragg grating, this process must be repeated for each optical frequency to be considered. While this is possible for simple one-dimensional problems, it becomes impractical for more realistic waveguides. Fortunately, in most cases the mode profiles do not change substantially over the bandwidth of interest. Thus, it is usually sufficient to neglect the small change in the electromagnetic modes and simply treat the change in frequency as a change in the propagation constant β. If we Taylor-expand the propagation constant $\beta(\omega)$ to first order about a central frequency ω_c, we obtain [6]:

$$\begin{aligned}\beta_0(\omega) &= \beta_0(\omega_c) + \frac{d\beta_0}{d\omega}|_{\omega_c}(\omega - \omega_c) \\ &= \beta_0(\omega_c) + \frac{1}{v_g}(\omega - \omega_c)\end{aligned} \quad , \tag{2.119}$$

where v_g is the group-velocity of the unperturbed waveguide. Using this relationship, the quantity δ may be expanded to first-order in terms of the deviation from the central frequency ω_c:

$$\delta(\omega) = \beta(w_c) - \frac{\pi}{\Lambda} + \frac{1}{v_g}(\omega - \omega_c) \quad . \tag{2.120}$$

In the above equation, we have assumed that the unperturbed waveguide geometry has been selected as described in Eq. 2.115 so that there is no first-order change in the propagation constant. Usually, the center frequency ω_c is chosen to coincide with the Bragg frequency, i.e. $\beta(\omega_c) = \pi/\Lambda$, which leads to

$$\delta(\omega) = \frac{1}{v_g}(\omega - \omega_c) \quad . \tag{2.121}$$

The expression for the coupling coefficient κ also depends upon ω, both through the electromagnetic fields which appear in the overlap integrals and through a proportionality constant which depends on $k(= \omega/c)$. However, these variations are small and can usually be neglected over the bandwidth of interest. Therefore, we shall take κ to be a frequency-independent parameter which is calculated at the center frequency ω_c.

Thus, we can calculate the spectral response of the Bragg grating by solving the coupled mode equations, treating δ as a variable which measures the deviation from the Bragg condition as described in Eq. 2.121.

The pair of coupled mode equations in Eq. 2.117 comprise a 2 × 2 system of linear

---

[6] Notice that to avoid confusion, we had used the subscript c to refer to the central frequency (the optical frequency at which the electromagnetic modes were computed.) In contrast, the subscript 0 refers to the unperturbed waveguide geometry.

homogeneous differential equations. The standard technique used to solve such a system of differential equations is eigenvector decomposition [65]. Applying this technique, the solution to the coupled mode equations is:

$$\begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} = [\mathbf{v}_1 \quad \mathbf{v}_2] \begin{bmatrix} e^{\gamma_1 z} & 0 \\ 0 & e^{\gamma_2 z} \end{bmatrix} [\mathbf{v}_1 \quad \mathbf{v}_2]^{-1} \begin{bmatrix} a_+(0) \\ a_-(0) \end{bmatrix} \quad , \tag{2.122}$$

where $\mathbf{v}_1$ and $\mathbf{v}_2$ are the eigenvectors of the coupling matrix and γ_1 and γ_2 are the corresponding eigenvalues. The eigenvalues and associated eigenvectors are given by:

$$\gamma_1 = +\gamma, \quad \mathbf{v}_1 = \begin{bmatrix} \gamma - j\delta \\ \kappa^* \end{bmatrix} \tag{2.123}$$

$$\gamma_2 = -\gamma, \quad \mathbf{v}_2 = \begin{bmatrix} -\kappa \\ \gamma - j\delta \end{bmatrix} \quad , \tag{2.124}$$

where the quantity γ is defined as:

$$\gamma \equiv \sqrt{|\kappa|^2 - \delta^2} \quad . \tag{2.125}$$

Substituting the eigenvalues and eigenvectors from Eq. 2.123 into Eq. 2.122, we obtain an analytical solution to the coupled mode equations:

$$\begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} = \begin{bmatrix} \cosh(\gamma z) - j\frac{\delta}{\gamma}\sinh(\gamma z) & \frac{\kappa}{\gamma}\sinh(\gamma z) \\ \frac{\kappa^*}{\gamma}\sinh(\gamma z) & \cosh(\gamma z) + j\frac{\delta}{\gamma}\sinh(\gamma z) \end{bmatrix} \begin{bmatrix} a_+(0) \\ a_-(0) \end{bmatrix} . \tag{2.126}$$

The above equation gives the solution to the coupled mode equations in transfer matrix form. If the forward and backward mode amplitudes are known at $z = 0$, they can be calculated at any other location z by simply multiplying by the transfer matrix given an Eq. 2.126.

Often, the boundary conditions on a_+ and a_- at $z = 0$ are not simultaneously known. Consider, for example, a grating which begins at $z = 0$ and extends to a length L. Now imagine injecting light into the left-hand side of the grating. We wish to calculate the reflection and transmission of the grating. We further assume that no light enters the grating from the right hand side. Substituting these boundary conditions into the transfer-matrix equation gives:

$$\begin{bmatrix} t \\ 0 \end{bmatrix} = \begin{bmatrix} \cosh(\gamma L) - j\frac{\delta}{\gamma}\sinh(\gamma L) & \frac{\kappa}{\gamma}\sinh(\gamma L) \\ \frac{\kappa^*}{\gamma}\sinh(\gamma L) & \cosh(\gamma L) + j\frac{\delta}{\gamma}\sinh(\gamma L) \end{bmatrix} \begin{bmatrix} 1 \\ r \end{bmatrix} \quad , \tag{2.127}$$

where r and t represent the reflection and transmission coefficients of the Bragg grating. The second equation in Eq. 2.127 can be directly solved for r, giving:

$$r(\delta) = \frac{\dfrac{\kappa^*}{\gamma}\tanh(\gamma L)}{1 + j\dfrac{\delta}{\gamma}\tanh(\gamma L)} \quad , \tag{2.128}$$

where γ is a function of δ as described in Eq. 2.125. Similarly, the amplitude transmission function t can be shown to be[7]:

$$t(\delta) = \frac{\mathrm{sech}(\gamma L)}{1 + j\dfrac{\delta}{\gamma}\tanh(\gamma L)} \quad . \tag{2.129}$$

The peak reflectivity of a uniform Bragg grating occurs when $\delta = 0$,

$$R_{\mathsf{max}} = |r(\delta = 0)|^2 = \tanh^2(\kappa L) \quad . \tag{2.130}$$

Figure 2.21 plots the calculated reflection spectral response for five different Bragg gratings with increasing lengths. From this figure, we can identify two different operating regimes. When the product κL is small compared to 1, the spectral response of the Bragg gratings can be accurately approximated as the Fourier transform of the grating shape. Since the grating begins and ends abruptly and extends for a length L, the spectral response has a characteristic "sinc" shape whose bandwidth is inversely proportional to the grating length L. In this limit, the spectral response can be shown to approach:

$$r(\delta) = j\frac{\kappa^*}{\delta}\sin(\delta L) \quad . \tag{2.131}$$

We refer to gratings with $\kappa L < 1$ as *weak Bragg gratings*, because in general they only reflect a fraction of the incident light. The peak reflectivity in this regime depends upon the value of κ, but the overall spectral shape and bandwidth is determined only by the grating length.

A weak Bragg grating does not make a suitable add/drop filter, because it only partially reflects the input signal. However, there are cases where the "sinc"-shaped spectral response of a weak grating is desirable. In many binary communications systems, the encoded signal has precisely at the same "sinc"-shaped spectral response. Thus, the weak

---

[7] If the transmission function is defined in terms of $A_\pm$ instead of the slowly-varying envelope functions $a_\pm$, there is an additional phase term in the expression for $t(\delta)$

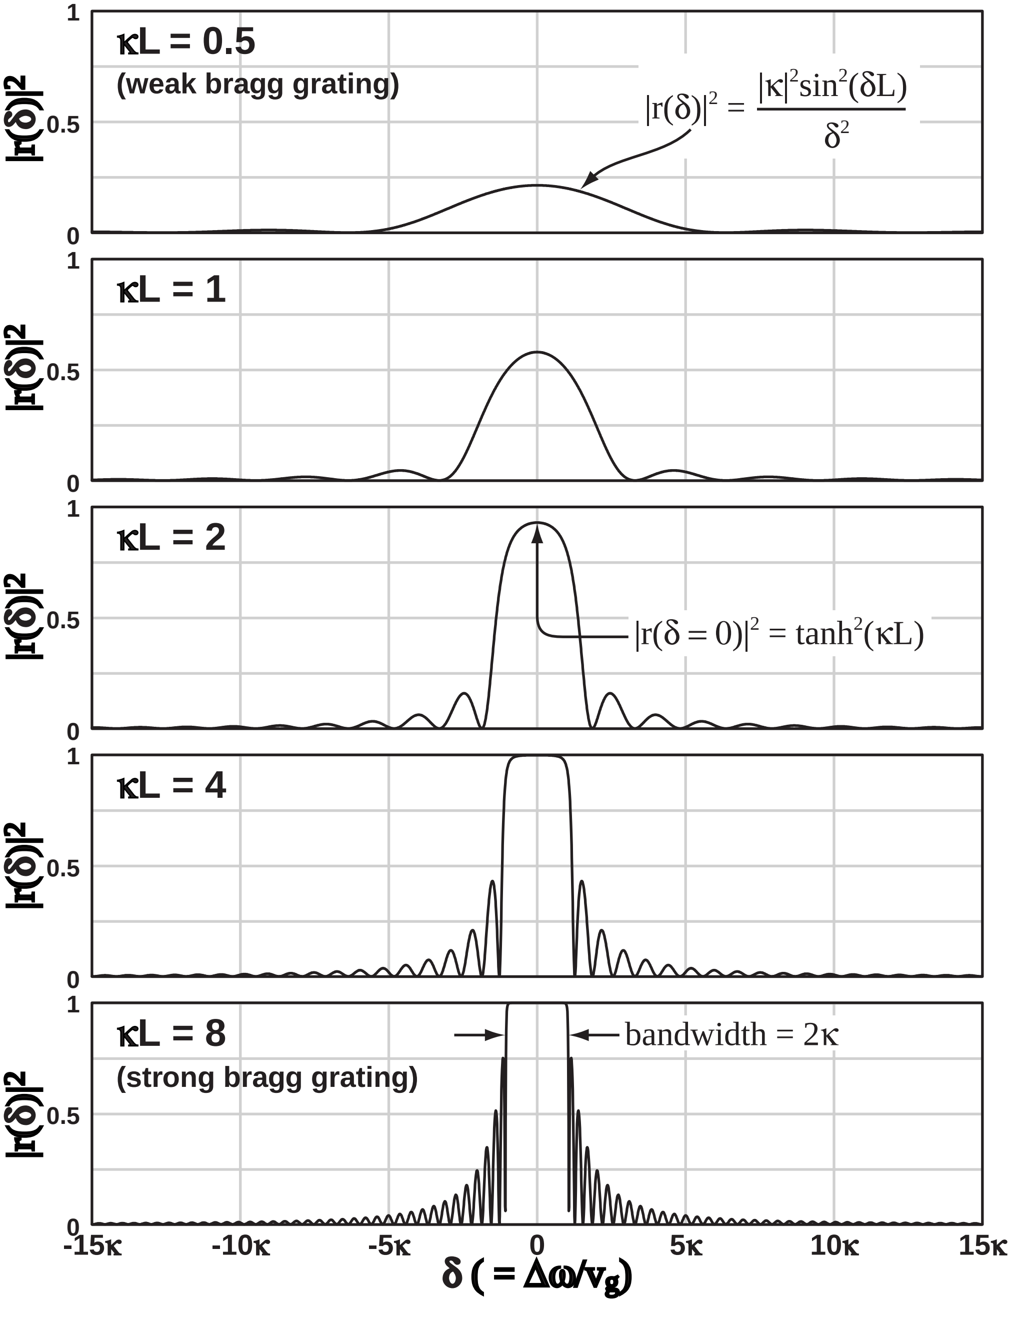

Figure 2.21: Calculated reflection spectral response from five different Bragg gratings, with increasing lengths. (figs/2/kappa-various-l.eps)

Bragg grating provides a convenient implementation of an all-optical *matched filter*, which should provide the optimal signal-to-noise ratio for detecting the signal in the presence of white background noise [66, 67, 68].

For gratings where the product κL exceeds 1, the spectral response has a plateau-like shape, as shown in the lower portion of Fig. 2.21. In this regime, the grating has a very high reflectivity within a band of frequencies called the *stopband*. The spectral width of the stopband is proportional to the grating strength κ:

$$\Delta f = \frac{\kappa c}{\pi n_g} \quad , \tag{2.132}$$

where Δf is the bandwidth and c/n_g is the group velocity. Outside of the stopband, the spectral response shows a series of ripples or sidelobes. The sidelobes quickly decay as we move away from the Bragg condition until the structure is effectively transparent [8]. If the grating is made longer without changing the value of κ, the bandwidth remains unchanged, but the peak reflectivity rises closer to 1, the spectrum becomes more square, and the sidelobes get closer together.

2.3.4 Calculation of Grating Strength

In the previous section, we solved the coupled mode equations to obtain the spectral response for a Bragg grating. The grating strength κ defines the so-called *stopband* of the Bragg grating, which (for strong gratings) is directly related to the bandwidth of the spectral response. We now turn to a more detailed discussion of how this grating strength is calculated for typical waveguide structures.

As described earlier, we choose the unperturbed waveguide structure such that there is no 0^{th} Fourier coefficient in the grating perturbation. For a grating comprised of rectangular teeth with a fixed duty cycle, this means that when constructing the unperturbed waveguide, the grating should be modeled as a homogeneous slab whose index lies somewhere in between the core and cladding indices. We define the duty cycle D to be the fraction of the grating period which is occupied by a tooth of core material. This geometry is illustrated in Fig. 2.20. With this definition, the refractive index profile of the unper-

---

[8] In fact, it can be shown that for points well outside of the stopband, the reflectivity of the grating can again be approximated by Eq. 2.131.

turbed waveguide is:

$$n_0^2(x,y) = \begin{cases} Dn_{\text{core}}^2 + (1-D)n_{\text{clad}}^2 & \text{in the grating region} \\ n^2(x,y) & \text{elsewhere} \end{cases} \quad . \tag{2.133}$$

With this choice, it can easily be verified that $\delta\epsilon_0(x,y) = 0$, as desired. The remaining Fourier coefficients can be computed using Eq. 2.109,

$$\delta\epsilon_m(x,y) = \begin{cases} (n_{\text{core}}^2 - n_{\text{clad}}^2)\dfrac{\sin(m\pi D)}{m\pi} & \text{in the grating region} \\ 0 & \text{elsewhere} \end{cases} \quad . \tag{2.134}$$

Note that we have arbitrarily chosen the position of the $z = 0$ origin to coincide with the center of a grating tooth, and consequently all of the Fourier coefficients are real. For first-order Bragg gratings, we are only interested in the first Fourier coefficients. Substituting the $m = 1$ coefficients from Eq. 2.134 into Eq. 2.113, we obtain following equation for the coupling constant κ:

$$\kappa = \frac{k}{4\pi P\eta_0}(n_{\text{core}}^2 - n_{\text{clad}}^2)\sin(\pi D) \iint\limits_{\text{(grating)}} \mathbf{e}_+^* \cdot \mathbf{e}_- \, dx\, dy \quad . \tag{2.135}$$

In the above equation, we have omitted the arbitrary phase of κ, recognizing that the phase is determined by the relative position of the $z = 0$ origin.

The integrand in Eq. 2.135 can be expanded using the relationships between the forward and backward traveling modes given in Eq. 2.87:

$$\mathbf{e}_+^* \cdot \mathbf{e}_- = |\mathbf{e}_t|^2 - |e_z|^2 \tag{2.136}$$

The finite difference techniques described in Section 2.1.4 allow one to calculate the transverse electric fields $\mathbf{e}_t$. The longitudinal field component e_z and magnetic field $\mathbf{h}$ may be computed from the two transverse components, using Maxwells equations. In most cases of interest, the longitudinal field component is much smaller than the transverse components, and the equation for κ can be simplified by neglecting the e_z terms. In many cases it is also possible to neglect one of the transverse electric field components and simply use a semivectorial technique to calculate the principal field component. In the context of the semi-vectorial method, the overlap integrals can be re-expressed in terms of one scalar field quantity representing the principal transverse electric field component (either $\mathbf{e}_x$ or

$\mathbf{e}_y$),

$$\iint\limits_{\text{(grating)}} \mathbf{e}_+^* \cdot \mathbf{e}_- \, dx\, dy \simeq \iint\limits_{\text{(grating)}} |\phi(x,y)|^2 \, dx\, dy \tag{2.137}$$

$$4P\eta_0 \simeq \frac{2\beta}{k} \iint |\phi(x,y)|^2 \, dx\, dy \quad , \tag{2.138}$$

where we have used the symbol $\phi(x, y)$ to represent $\mathbf{e}_x$ for TE modes and $\mathbf{e}_y$ for TM modes. With this simplification, the expression for the grating strength κ becomes:

$$\kappa = \frac{k^2}{2\pi\beta}(n_{\text{core}}^2 - n_{\text{clad}}^2)\sin(\pi D)\Gamma \quad . \tag{2.139}$$

The quantity Γ in Eq. 2.139 is a dimensionless number less than 1 which describes the extent to which the electromagnetic mode overlaps with the grating region.

$$\Gamma \equiv \frac{\displaystyle\iint\limits_{\text{(grating)}} |\phi(x,y)|^2 \, dx\, dy}{\displaystyle\iint |\phi(x,y)|^2 \, dx\, dy} \quad . \tag{2.140}$$

Figure 2.22 plots the calculated coupling constant κ as a function of duty cycle D, for various grating depths, for the doped glass channel waveguide.

In Fig. 2.23, we plot the coupling constant as a function of duty cycle for various grating depths for the silicon-on-insulator ridge waveguide. One interesting effect which is seen in this plot is that the strongest grating strength occurs when the duty cycle is larger than 50%. Even though the first Fourier coefficient is maximized by choosing $D = 0.5$, the mode overlap integral increases substantially with duty cycle, leading to the skewed plots presented in Fig. 2.23. Notice that there is also a significant polarization dependence in the data presented in Fig. 2.23 – the calculated grating strength for TM-polarized light is substantially smaller than for TE-polarized light for the same grating edge depth and duty cycle.

Fiber Bragg Gratings

It is instructive to evaluate the strength of a Bragg grating in a waveguide in comparison to that of a fiber Bragg grating. In a fiber Bragg grating, the periodic modulation is created by illuminating a photosensitive fiber with a periodic UV standing-wave. Typically,

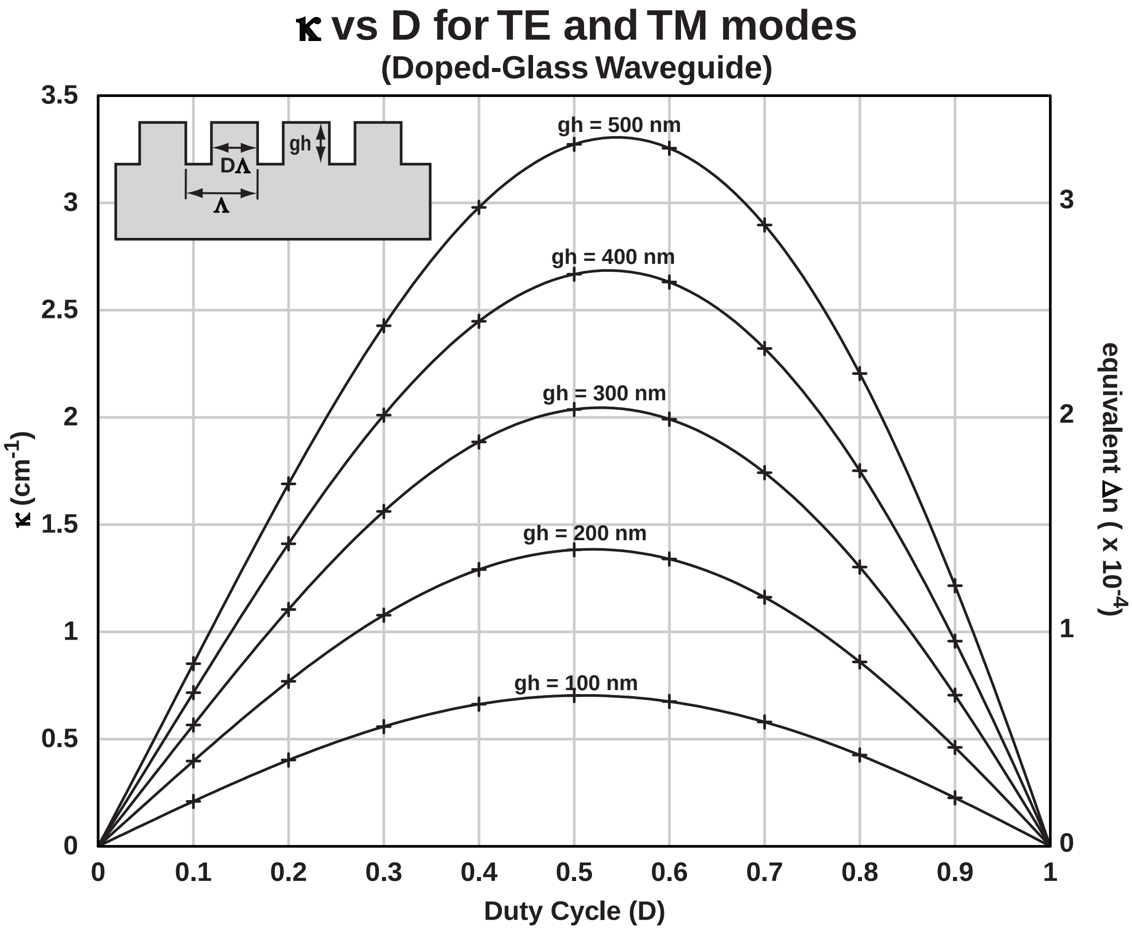

Figure 2.22: Calculated grating strength κ vs. duty cycle D for a doped-glass channel waveguide. The results are indistinguishable for TE and TM polarizations. The grating strength is calculated for five different grating etch depths ranging from 100 to 500 nm.

(figs/2/kappa-glass.eps)

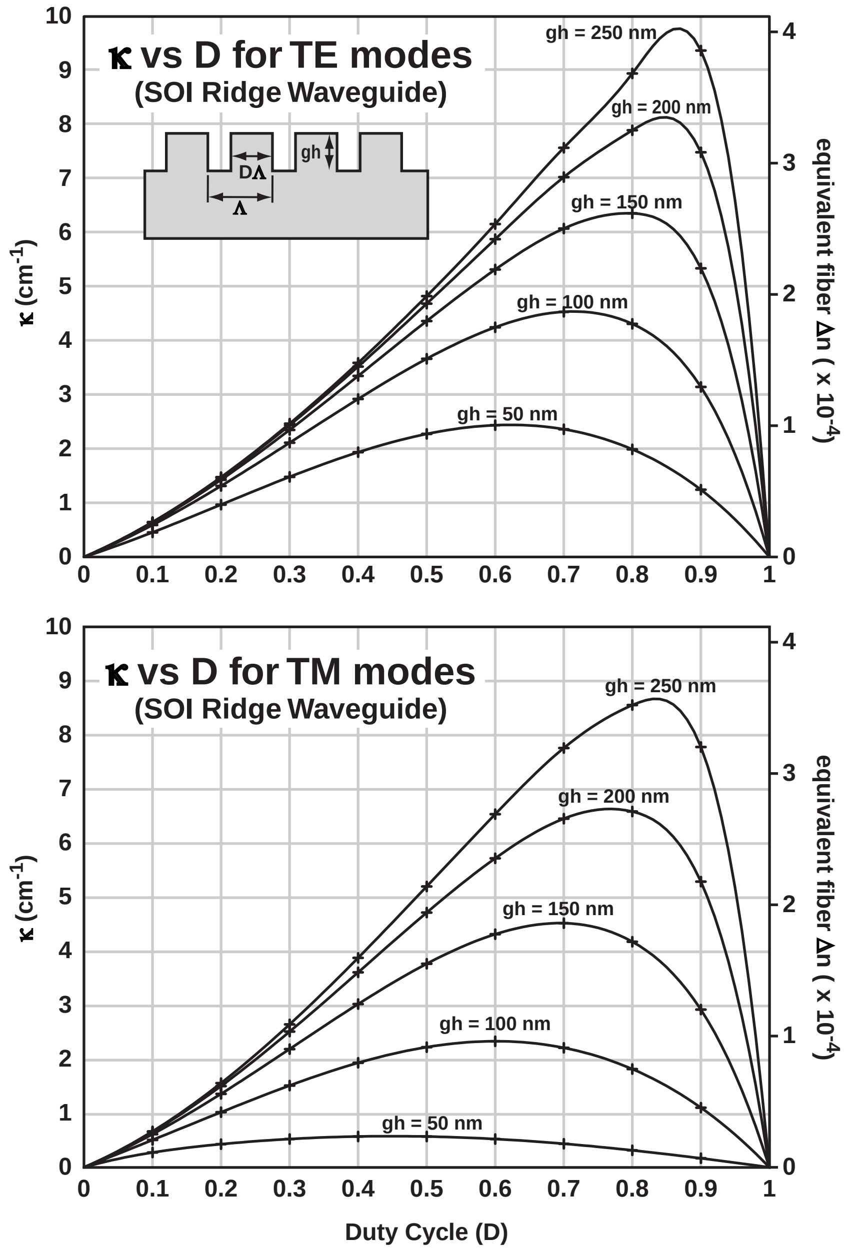

Figure 2.23: Calculated grating strength κ vs. duty cycle D for an SOI ridge waveguide. The upper plot applies for the TE polarization, while the lower plot corresponds to TM polarization. In each case, the grating strength is calculated for five different grating edges ranging from 50 to 250 nm. (figs/2/kappa-soi.eps)

the fiber core is photorefractive, meaning that its index of refraction can be permanently changed by exposing the fiber to UV radiation. In contrast to the integrated Bragg gratings discussed so far, a fiber Bragg grating has no physical corrugation – there is simply a small modulation in the refractive index of the core. The amount by which the index of refraction can be changed via UV illumination depends upon the composition of the core and the wavelength of exposure. Many techniques have been developed for enhancing the photosensitivity of fibers[69], but generally the maximum attainable refractive index change in an untreated fiber seldom exceeds 1×10^{-3} [70].

The coupled mode analysis described earlier can also be applied to fiber Bragg gratings. The periodic perturbation $\delta\epsilon(x,y,z)$ for a fiber Bragg grating with a peak-to-peak refractive index modulation of Δn is:

$$\delta\epsilon(x,y,z) = \begin{cases} n\Delta n\cos(2\pi z/\Lambda), & \text{inside core} \\ 0, & \text{outside core} \end{cases} \quad , \tag{2.141}$$

where we have assumed that $\Delta n/n$ is small. For the fiber Bragg grating, the periodic perturbation is a pure sinusoid[9] which means that the Fourier decomposition of $\delta\epsilon(x,y,z)$ only has two nonzero coefficients:

$$\delta\epsilon_{\pm1} = \begin{cases} \dfrac{n\Delta n}{2}, & \text{inside core} \\ 0, & \text{outside core} \end{cases} \quad . \tag{2.142}$$

We can now apply the same procedure described earlier for corrugated Bragg gratings, using Eq. 2.142 for the first Fourier coefficient of the grating. The resulting expression for κ is:

$$\kappa = \frac{k^2}{4\beta} n\Delta n\Gamma \quad , \tag{2.143}$$

where Γ is, as before, a dimensionless number less than 1, but it now describes the fraction of the modal power which resides in the waveguide core. (In comparison with the corrugated Bragg grating, the fiber Bragg can achieve a larger value for Γ because the index modulation covers the entire core region.)

For the sake of comparison, we shall evaluate Eq. 2.143 using some values which are

---

[9] The sinusoidal index approximation assumes that the UV-induced refractive index change is proportional to the intensity pattern formed by a superposition of two plane waves of equal amplitude. This is only a convenient approximation of what really occurs when a fiber Bragg grating is fabricated. Reference [71] describes a technique for measuring the actual refractive index modulation in a fiber Bragg grating.

typical for the strong fiber Bragg gratings. Specifically, we assume that the peak-to-peak index modulation amplitude is Δn = 1×10^{-3}, the free-space wavelength is taken to be 1550 nm, and the overlap factor Γ is taken to be 90%. With these numbers, the equivalent grating strength is approximately κ = 18 cm^{-1}. If we assume that the group index is approximately 1.46 (typical for fiber), this corresponds to a stopband spectral width of approximately 120 GHz.

2.3.5 Apodized and Chirped Bragg Gratings

So far, we have assumed that the Bragg grating is a perfectly periodic structure which has a sharply-defined beginning and endpoint. Often, either by design or because of the fabrication techniques used, the Bragg grating will deviate somewhat from this perfect periodic structure. If the pitch of the Bragg grating changes slowly along the length of the grating, we say that there is a *chirp*. More generally, even when the grating is perfectly periodic, if the waveguide itself has some slow variation in its propagation constant across the length of the device, the structure will exhibit chirp-like properties. When the amplitude of the grating constant κ changes with position, we say that the grating is *apodized*. In both of these cases, the coupled mode equations must be modified to account for the nonuniformity in the periodic structure. We now describe how to generalize the coupled mode equations to treat apodized and chirped gratings.

First, we recall the coupled mode equations for $A_{\pm}(z)$ derived earlier for uniform gratings. The relevant equations are repeated here for reference:

$$\frac{d}{dz}\begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} = \begin{bmatrix} -j\beta & \kappa e^{-jk_\mathrm{g}z} \\ \kappa^* e^{+jk_\mathrm{g}z} & j\beta \end{bmatrix} \begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} \quad , \tag{2.144}$$

where k_g is the grating spatial frequency, defined by:

$$k_\mathrm{g} \equiv \frac{2\pi}{\Lambda} \quad , \tag{2.145}$$

Now imagine that the grating slowly changes with z. Strictly speaking, the Fourier analysis of the grating perturbation no longer applies when the grating is aperiodic. Nevertheless, provided the grating changes slowly with z, the coupled-mode analysis can still be applied locally.

The effect of grating apodization can be included by simply letting κ in Eq. 2.144 be a slowly varying function of z. The way to account for grating chirp is to replace the product

$k_g z$ by a more general integral $\int k_g(z)\,dz$. [10] Thus, the generalized coupled mode equations for nonuniform gratings are:

$$\frac{d}{dz}\begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} = \begin{bmatrix} -j\beta(z) & \kappa_0(z)e^{-j\phi(z)} \\ \kappa_0(z)^* e^{+j\phi(z)} & j\beta(z) \end{bmatrix} \begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} \tag{2.146}$$

$$\text{where } \phi(z) \equiv \int_0^z k_g(z')\,dz' \quad . \tag{2.147}$$

Note: in the above equation, we have also allowed the propagation constant β to vary with z, to account for any slow change in the structure or composition of the waveguide itself, which is indistintuishable from grating chirp. We shall assume that the function κ_0 is a real-valued function which describes the magnitude of the grating perturbation. The grating phase $\phi(z)$ includes all of the information about the chirp and phase of the Bragg grating. For a uniform Bragg grating, this term is a linear function of z, but for nonuniform gratings it not a strictly linear function. The first derivative of $\phi(z)$ is the local k-vector of the grating.

As before, we often wish to rewrite Eq. 2.146 in terms of the slowly-varying amplitude functions $a_\pm(z)$. But now, since there is no well-defined grating periodicity to factor out of the equations, we simply choose an arbitrary reference point, β_r,

$$A_\pm(z) = a_\pm(z)e^{\mp j\beta_r z} \quad , \tag{2.148}$$

Substituting the above equation into the coupled mode equations leads to the following set of differential equations for the slowly-varying amplitude functions:

$$\frac{d}{dz}\begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} = \begin{bmatrix} -j\delta(z) & \kappa(z) \\ \kappa(z)^* & j\delta(z) \end{bmatrix} \begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} \quad , \tag{2.149}$$

where the functions $\kappa(z)$ and $\delta(z)$ are defined as:

$$\delta(z) = \beta(z) - \beta_r \tag{2.150}$$

$$\kappa(z) = \kappa_0(z)\exp\left[j(2\beta_r z - \phi(z))\right] \quad . \tag{2.151}$$

Eq. 2.149 is identical to Eq. 2.117, with the exception that the matrix elements are slowly-varying functions of z. The effect of grating chirp it is to introduce a slow variation in the phase of $\kappa(z)$. And the effect of apodization is to introduce a slow variation in the

---

[10] One might be tempted to account for grating chirp by simply replacing k_g in Eq. 2.144 by a slowly-varying function of z, but this leads to an inconsistent interpretation of k_g as the local spatial frequency of the grating.

magnitude of $\kappa(z)$.

Because the coefficients of the coupling matrix are no longer constants, the coupled mode equations for nonuniform gratings can no longer be solved in closed form. The differential equation can be solved numerically by simple integration. Again, we are faced with the problem that the boundary conditions are not completely specified at either end of the structure. For light incident from the left-hand side of the grating, the easiest way to compute the reflection and transmission coefficients is integrate backwards from right to left. At the right hand side of the grating, we assume that there is no light impinging from the right, and we arbitrarily choose the output amplitude to be 1. Therefore, the boundary conditions at the end of the grating are:

$$\mathbf{a}(L) = \begin{bmatrix} a_+(L) \\ a_-(L) \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \end{bmatrix} \quad . \tag{2.152}$$

With this as our starting point, we integrate the differential equation from $z = L$ back to $z = 0$,

$$\mathbf{a}(0) = \mathbf{a}(L) + \int_L^0 \mathbf{M}(\mathbf{z})\, \mathbf{a}(\mathbf{z})\, dz \quad , \tag{2.153}$$

where $\mathbf{M}(z)$ is the 2 × 2 coupling matrix in Eq. 2.149. The reflection and transmission coefficients are then given by:

$$r = \frac{a_-(0)}{a_+(0)}, \qquad t = \frac{1}{a_+(0)} \quad . \tag{2.154}$$

In Section 2.4, we describe how the numerically computed reflection and transmission coefficients can be used to construct the full 2 × 2 transfer matrix for the nonuniform grating.

If only the reflection coefficient is sought, the problem may be simplified by deriving a new differential equation for the local amplitude ratio $r(z)$, defined as

$$r(z) \equiv \frac{a_-(z)}{a_+(z)} \quad . \tag{2.155}$$

Computing the derivative of Eq. 2.155 with respect to z and using the coupled mode equations, we obtain the so-called Ricatti differential equation:

$$\begin{aligned}\frac{d}{dz}r(z) &= \frac{1}{a_+}\frac{da_-}{dz} - \frac{a_-}{a_+^2}\frac{da_+}{dz} \\ &= \kappa^*(z) + 2j\delta(z)r(z) - \kappa(z)r^2(z)\end{aligned} \quad . \tag{2.156}$$

As before, this equation can be numerically integrated from right to left to obtain the reflection coefficient of the Bragg grating.

A simpler approach to the analysis of nonuniform gratings is to simply divide the grating into short segments, over which δ, κ, and Λ are approximately constant. The transfer matrix for the entire structure can then be obtained by simply multiplying the transfer matrices for each short segment. Essentially, this technique amounts to performing a primitive numerical integration of the differential equations. Transfer matrix methods will be further discussed in the following section.

Example 1: Apodized Bragg Grating

As an example of how these techniques can be applied, we demonstrate how the spectrum of an apodized grating is computed. As illustrated in Fig. 2.21, the reflection spectrum from a Bragg grating typically has a series of sidelobes on either side of the stopband. Were we to use such a grating to filter one wavelength channel from a multi-wavelength system, the sidelobes could lead to crosstalk if they overlap with the adjacent wavelength channels. The sidelobes can be substantially reduced by apodizing the grating [72, 73]. As a simple example, we consider a Bragg grating where the grating strength κ has a raised-cosine profile:

$$\kappa(z) = \kappa_0 \cos^2(\pi\frac{z}{L}) \quad , \tag{2.157}$$

where the grating extends from $-L/2$ to $+L/2$. Figure 2.24 plots the computed spectral response for such an apodized Bragg grating. In computing this spectrum, the indices of refraction were taken to be 1.46, the peak value of κ was chosen to be 2 cm^{-1}, and the total length of the structure is 5 cm. For comparison, we also plot the reflection spectrum from an equivalent uniform Bragg grating with the same length and peak reflectivity. Notice that the apodized structure exhibits greatly reduced sidelobes levels, and a broader bandwidth.

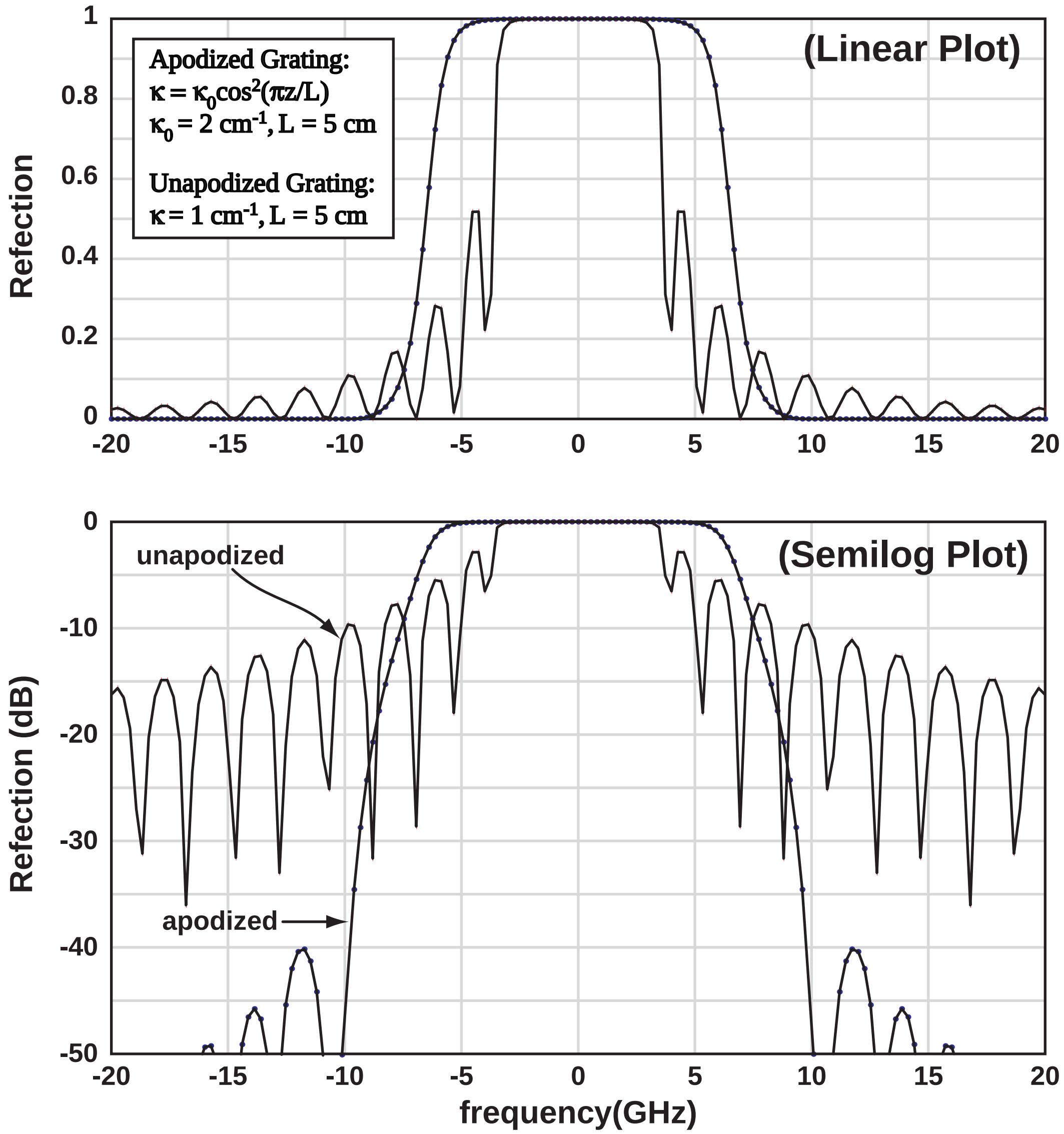

Figure 2.24: Calculated reflection spectrum from an apodized Bragg grating, with a raised-cosine apodization profile, in comparison to a uniform Bragg grating with the same peak reflectivity. The following parameters were used in calculating the spectra presented here: $n_{\text{eff}} = n_g = 1.46$, $\int \kappa(z)\,dz = 5$, $\kappa_0 = 2\ \text{cm}^{-1}$ (peak κ for the apodized structure), and $\kappa_{\text{uniform}} = 1\ \text{cm}^{-1}$. (figs/2/apodized-spectrum.eps)

Example 2: Linearly-Chirped Bragg Grating

In the previous example, we investigated the effect apodization. We now provide an example which illustrates how to calculate the spectral response of a linearly chirped grating. To begin, we assume that the grating period Λ is a linear function of z:

$$\Lambda(z) = \Lambda_0 + Az \quad , \tag{2.158}$$

where A is a dimensionless constant (typically given in units of pm / cm) describing the rate of chirp. The grating phase $\phi(z)$ is computed by integrating the k-vector over z:

$$\begin{aligned}\phi(z) &= \int_0^z \frac{2\pi}{\Lambda_0 + Az'}\, dz' &(2.159)\\ &= \frac{2\pi}{A} \ln\left(1 + \frac{Az}{\Lambda_0}\right) \quad . &(2.160)\end{aligned}$$

Provided the term Az/Λ_0 is small compared to one, we can accurately approximate Eq. 2.159 with a Taylor series:

$$\phi(z) = \frac{2\pi}{\Lambda_0} z - \frac{\pi A}{\Lambda_0^2} z^2 + O\left(\frac{Az}{\Lambda_0}\right)^3 \quad . \tag{2.161}$$

The linear term in the above equation corresponds to the linear phase progression of a uniform grating with period Λ_0. The quadratic term arises because of the grating chirp. If we choose $\beta_r = \pi/\Lambda_0$ in Eq. 2.150, the linear term drops out and we obtain the following coupled mode parameters:

$$\begin{aligned}\delta &\equiv \beta - \frac{\pi}{\Lambda_0} &(2.162)\\ \kappa(z) &\equiv \kappa_0 \exp\left[+j\frac{\pi A}{\Lambda_0^2} z^2\right] \quad . &(2.163)\end{aligned}$$

Notice that a linear chirp in the grating period gives rise to a quadratic term in the grating phase $\phi(z)$.

One application of chirped Bragg gratings is to compensate for signal dispersion (pulse spreading) which occurs during transmission over a long distance of fiber [74, 75]. However, for the types of filters discussed in this work, grating chirp adversely affects the device performance. Figure 2.25 plots the calculated reflection and transmission functions for a chirped Bragg grating with $\kappa = 1\,\text{cm}^{-1}$, $L = 5\,\text{cm}$, and $A = 10\,\text{pm}/\text{nm}$. Notice that the chirp significantly broadens the spectral response, raises the sidelobe levels, and changes the peak rejection ratio from -37 dB to -15 dB.

2.4 Transfer Matrix Methods

The previous two sections describe how Maxwells equations can be simplified to a set of coupled differential equations which govern the mode amplitudes. When solving coupled mode equations, the solutions have always be expressed in the form of a *transfer matrix*. That is, the mode amplitudes at some position z are written as a linear combination of the mode amplitudes at $z = 0$ or some other reference location. Although it might seem coincidental that the solutions can always be expressed in this form, it is a direct consequence of the linearity of the differential equations.

Eq. 2.68 gives the transfer matrix for a codirectional coupler, and Eq. 2.126 the transfer matrix for a Bragg grating. The transfer matrix for a more complex structure can easily be computed by simply multiplying the transfer matrices for its constituent parts. Section 2.4.4 gives an example of how this is done.

2.4.1 Multi-Waveguide Systems

There is one important difference between the transfer matrix of a directional coupler and the transfer matrix of a Bragg grating. In the case of a directional coupler, the transfer matrix given in Eq. 2.68 only applies to the forward-propagating mode amplitudes of the two waveguides. In contrast, the transfer matrix of a Bragg grating describes the forward and backward mode amplitudes of a single waveguide.

We must recognize, however, that if the directional coupler is operated in the reverse direction, it also provides coupling between the two backward-traveling modes. Because there are two waveguides in a directional coupler, a simple 2 × 2 coupling matrix cannot completely describe the system. A more complete transfer matrix solution for the directional coupler would be:

$$\begin{bmatrix} a_{1+}(z) \\ a_{2+}(z) \\ a_{1-}(z) \\ a_{2-}(z) \end{bmatrix} = \begin{bmatrix} \cos(\phi) & -j\sin(\phi) & 0 & 0 \\ -j\sin(\phi) & \cos(\phi) & 0 & 0 \\ 0 & 0 & \cos(\phi) & +j\sin(\phi) \\ 0 & 0 & +j\sin(\phi) & \cos(\phi) \end{bmatrix} \begin{bmatrix} a_{1+}(0) \\ a_{2+}(0) \\ a_{1-}(0) \\ a_{2-}(0) \end{bmatrix} \quad , \tag{2.164}$$

$$\phi \equiv \mu z \quad . \tag{2.165}$$

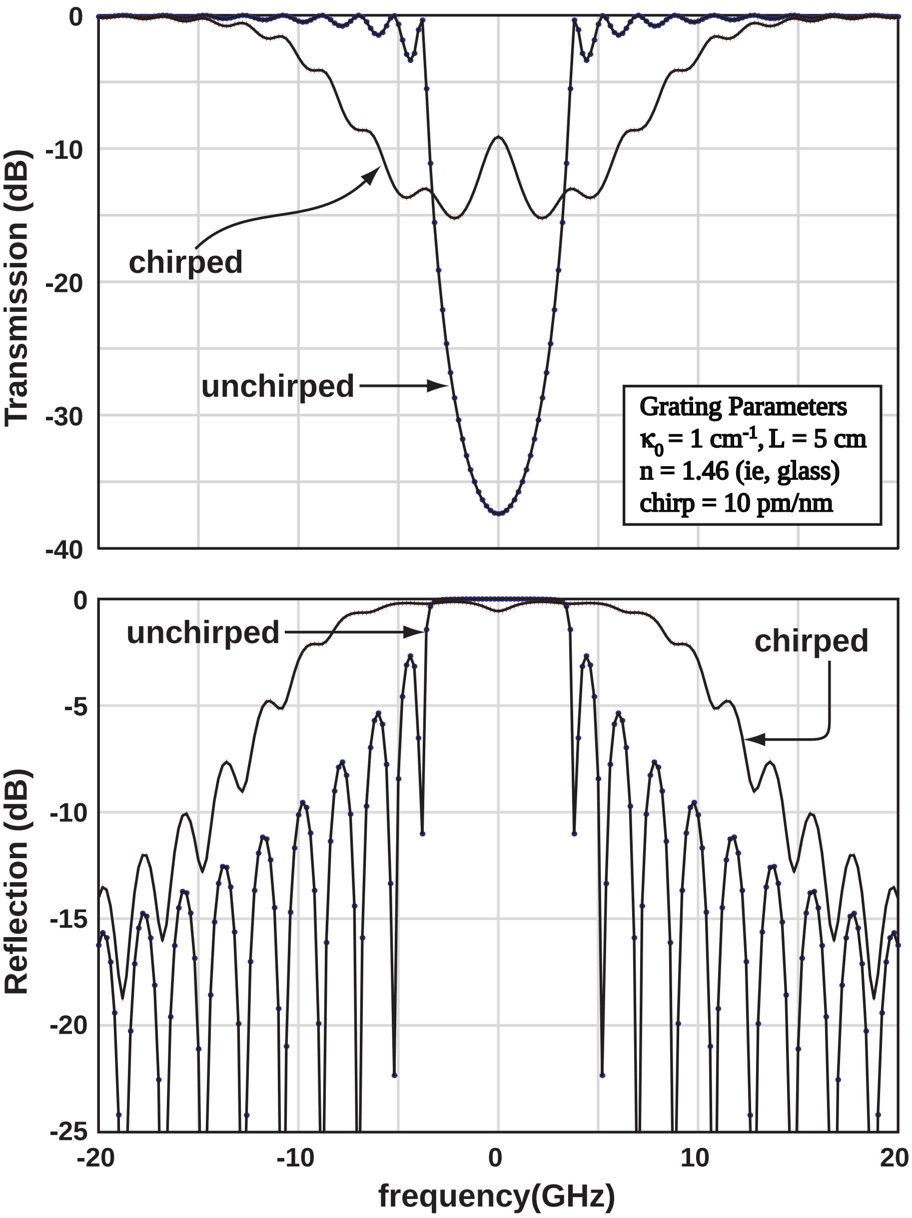

Figure 2.25: Calculated spectral response from a chirped Bragg grating with the following parameters: $\kappa = 1\,\mathrm{cm}^{-1}$, $L = 1\,\mathrm{cm}$, and $A = 10\,\mathrm{pm}\,/\,\mathrm{cm}$. (figs/2/linear-chirp-spectrum.eps)

In the above equation, we have stacked all 4 mode amplitudes into a single vector. The transfer matrix consists of block 2 × 2 matrices, and there is no interaction between the forward-traveling modes and backward-traveling modes. [11]

Generally, for a system consisting of N waveguides, the mode amplitudes will form a vector of length $2N$, and the transfer matrices will be $2N \times 2N$. Conventionally, the first N elements of the amplitude vector are taken to be the forward-traveling mode amplitudes, while the last N elements represent the corresponding backward-traveling mode amplitudes. The general form of a transfer matrix is:

$$\begin{bmatrix} \mathbf{a}_+(z) \\ \mathbf{a}_-(z) \end{bmatrix} = \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} \\ \mathbf{T}_{21} & \mathbf{T}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{a}_+(0) \\ \mathbf{a}_-(0) \end{bmatrix} \quad , \tag{2.166}$$

where $\mathbf{a}_\pm$ are each length-N vectors, and $\mathbf{T}_{11} \ldots \mathbf{T}_{22}$ are $N \times N$ matrices.

2.4.2 Transfer Matrices and Scattering Matrices

When we solved for the spectral response of a Bragg grating in Section 2.3.2, we pointed out that in most cases the forward and backward mode amplitudes are not simultaneously known at $z = 0$. In treating this kind of problem, it is often more useful to express the solution in terms of a *scattering matrix*. Whereas a transfer matrix describes the mode amplitudes at $z > 0$ in terms of the mode amplitudes at $z = 0$, a scattering matrix relates the light emerging from a structure to the light impinging on the structure. The general form of a scattering matrix is given below:

$$\begin{bmatrix} \mathbf{a}_-(0) \\ \mathbf{a}_+(z) \end{bmatrix} = \begin{bmatrix} \mathbf{S}_{11} & \mathbf{S}_{12} \\ \mathbf{S}_{21} & \mathbf{S}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{a}_+(0) \\ \mathbf{a}_-(z) \end{bmatrix} \quad , \tag{2.167}$$

where the quantities $\mathbf{a}_\pm$ are, as before, length-N vectors, and the $\mathbf{S}_{ij}$ are $N \times N$ matrices.

With a little algebraic manipulation, it is possible to relate the elements of the scattering matrix to the elements of the corresponding transfer matrix:

$$\begin{bmatrix} \mathbf{S}_{11} & \mathbf{S}_{12} \\ \mathbf{S}_{21} & \mathbf{S}_{22} \end{bmatrix} = \begin{bmatrix} -\mathbf{T}_{22}^{-1}\mathbf{T}_{21} & \mathbf{T}_{22}^{-1} \\ (\mathbf{T}_{11} - \mathbf{T}_{12}\mathbf{T}_{22}^{-1}\mathbf{T}_{21}) & \mathbf{T}_{12}\mathbf{T}_{22}^{-1} \end{bmatrix} \quad . \tag{2.168}$$

Similarly, it is possible to write the transfer matrix elements in terms of the scattering ma-

---

[11] is also possible to expand the transfer matrix approach to include both orthogonal polarization states and higher-order modes (if they exist).

trix elements:

$$\begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} \\ \mathbf{T}_{21} & \mathbf{T}_{22} \end{bmatrix} = \begin{bmatrix} (\mathbf{S}_{21} - \mathbf{S}_{22}\mathbf{S}_{12}^{-1}\mathbf{S}_{11}) & \mathbf{S}_{22}\mathbf{S}_{12}^{-1} \\ -\mathbf{S}_{12}^{-1}\mathbf{S}_{11} & \mathbf{S}_{12}^{-1} \end{bmatrix} \quad . \tag{2.169}$$

Because the scattering matrix describes the outputs in terms of the inputs, one might ask why transfer matrices are used at all. The benefit of the transfer matrix is that complex structures can be analyzed by simply multiplying the transfer matrices of each segment. Needless to say, scattering matrices cannot be cascaded in this manner.

Lossless Systems

For structures which do not have any optical gain or loss, there are additional constraints placed on the transfer matrix and scattering matrix by the requirements of power conservation and time reversal symmetry. Power conservation simply means that the net amount of power flowing into the device must be 0. Time reversal symmetry means that for every valid solution to Maxwell's equations, another solution may be obtained by replacing t by $-t$, β by $-\beta$, $\mathbf{E}$ by $\mathbf{E}^*$ and $\mathbf{H}$ by $-\mathbf{H}^*$[76]. This time-reversed solution is similar to playing a movie of the system backwards. Clearly, systems with loss do not have the same time reversal symmetry since the backwards-running solutions would exhibit gain.

Consider first the scattering matrix for a lossless linear system with inputs $\mathbf{x}$ and outputs $\mathbf{y}$:

$$\mathbf{y} = \mathbf{S}\mathbf{x} \quad . \tag{2.170}$$

Power conservation means that the magnitude of the vector on the right hand side is the same as the magnitude of the vector on the left-hand side. Thus, for lossless systems, multiplication by a scattering matrix is a norm-preserving operation. Mathematically, the requirement of power conservation can be stated as:

$$\mathbf{x}^\dagger\mathbf{x} = \mathbf{y}^\dagger\mathbf{y} = \mathbf{x}^\dagger\mathbf{S}^\dagger\mathbf{S}\mathbf{x} \quad . \tag{2.171}$$

This equation can only be satisfied if:

$$\mathbf{S}^\dagger\mathbf{S} = \mathbf{I} \quad , \tag{2.172}$$

where $\mathbf{S}^\dagger$ is the complex-conjugate transpose of the scattering matrix. Matrices with this property are referred to as *unitary* matrices.

For lossless systems, time reversal symmetry provides a similar relationship between the matrix elements. The time reversed solution to a problem is obtained from the forward-running solution by making the following substitutions:

$$\mathbf{x} \to \mathbf{y}^* \qquad , \quad \mathbf{y} \to \mathbf{x}^* \tag{2.173}$$

$$\begin{bmatrix} \mathbf{a}_+(0) \\ \mathbf{a}_-(z) \end{bmatrix} \to \begin{bmatrix} \mathbf{a}_-^*(0) \\ \mathbf{a}_+^*(z) \end{bmatrix}, \quad \begin{bmatrix} \mathbf{a}_-(0) \\ \mathbf{a}_+(z) \end{bmatrix} \to \begin{bmatrix} \mathbf{a}_+^*(0) \\ \mathbf{a}_-^*(z) \end{bmatrix} \tag{2.174}$$

By taking the complex conjugate of Eq. 2.170, and substituting a the time-reversed solutions of Eq. 2.173, we arrive at the following additional constraint on the scattering matrix:

$$\mathbf{S}^*\mathbf{S} = \mathbf{I} \quad . \tag{2.175}$$

Eq. 2.175, when combined with Eq. 2.172, can only be satisfied if the scattering matrix $\mathbf{S}$ is symmetric:

$$\mathbf{S} = \mathbf{S}^T \quad . \tag{2.176}$$

Similarly, by taking the complex conjugate of Eq. 2.166, and substituting the time-reversed solutions of Eq. 2.174, we obtain the following relationships between the the elements of the transfer matrix:

$$\mathbf{T}_{22} = \mathbf{T}_{11}^*, \quad \mathbf{T}_{21} = \mathbf{T}_{12}^* \quad . \tag{2.177}$$

Therefore, only two elements of the transfer matrix are needed to completely describe a lossless system. It is easy to verify that the transfer matrices given earlier for Bragg gratings and directional couplers satisfy Eq. 2.177 [12].

2.4.3 Useful Transfer Matrices

We now present a summary of useful transfer matrices for integrated optical devices.

---

[12] For directional couplers, Eq. 2.177 applies to the complete 4 × 4 transfer matrix of the system, i.e. that given in Eq. 2.164. It does not apply to the simple 2 × 2 submatrix describing only forward-traveling modes.

Uniform Waveguide (with Loss)

The transfer matrix for a lossless waveguide with propagation constant β is simply:

$$\mathbf{T} = \begin{bmatrix} e^{-j\beta z} & 0 \\ 0 & e^{+j\beta z} \end{bmatrix} \quad . \tag{2.178}$$

The above equation expresses what we learned earlier in Section 2.1, forward traveling waves propagate with wavevector β and backward traveling waves propagate with wavevector $-\beta$. If the waveguide has loss or gain, the appropriate transfer matrix is:

$$\mathbf{T} = \begin{bmatrix} e^{-(\alpha+j\beta)z} & 0 \\ 0 & e^{+(\alpha+j\beta)z} \end{bmatrix} \quad , \tag{2.179}$$

where positive values of α correspond to loss and negative values correspond to gain.

Partial Reflector

When an integrated waveguide reaches the end of a chip, and the light emerges into free space, there will be some reflection at the chip facet simply because of the refractive index change. More generally, anytime there is an abrupt change in the waveguide, some reflection can occur. If the electromagnetic modes are reasonably well matched on both sides of the discontinuity, the effect can often be modeled using the simple Fresnel reflection equations. The transfer matrix describing Fresnel reflection at a dielectric interface between two materials with indices n_1 and n_2 is:

$$\mathbf{T} = \frac{1}{2\sqrt{n_1 n_2}} \begin{bmatrix} n_1+n_2 & n_2-n_1 \\ n_2-n_1 & n_1+n_2 \end{bmatrix} \quad . \tag{2.180}$$

This transfer matrix relates the mode amplitudes immediately to the right of the interface to those immediately to the left.

More generally, the transfer matrix for a partially reflecting surface (such as a beamsplitter) can be always be cast in the following form:

$$\mathbf{T} = \frac{j}{\sqrt{1-r^2}} \begin{bmatrix} 1 & r \\ r & 1 \end{bmatrix} \quad , \tag{2.181}$$

where r a real number representing the amplitude reflection coefficient.

Directional Coupler

The complete transfer matrix for a symmetric directional coupler is given by:

$$\mathbf{T} = \begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^* \end{bmatrix} \quad , \tag{2.182}$$

where the submatrix $\mathbf{M}$ is a simple 2×2 matrix described by:

$$\mathbf{M} \equiv e^{-j\beta z} \begin{bmatrix} \cos\phi & -j\sin\phi \\ -j\sin\phi & \cos\phi \end{bmatrix}, \quad \phi \equiv \int_0^z \mu(z')\,dz' \quad . \tag{2.183}$$

Note: to simplify the mathematics, the $e^{j\beta z}$ phase terms are often omitted, as was done in Eq. 2.164. This simplification is valid as long as the two waveguides which make up the coupler are mirror images of each other. If the two waveguides differ, the transfer matrix must be modified to keep track of the phase difference between the two guides. Also, if the coupler is placed within a Fabry-Perot cavity (for example, a cavity formed by the polished facets of the chip), the spectral response of the complete structure cannot be calculated without including this phase term.

Bragg Grating

The transfer matrix for a Bragg grating of length L and grating strength κ is:

$$\mathbf{T} = \begin{bmatrix} e^{-jk_g z} & 0 \\ 0 & e^{+jk_g z} \end{bmatrix} \begin{bmatrix} \cosh(\gamma L) - j\frac{\delta}{\gamma}\sinh(\gamma L) & \frac{\kappa}{\gamma}\sinh(\gamma L) \\ \frac{\kappa^*}{\gamma}\sinh(\gamma L) & \cosh(\gamma L) + j\frac{\delta}{\gamma}\sinh(\gamma L) \end{bmatrix} \quad , \tag{2.184}$$

where the quantities κ, δ, and γ are defined in Section 2.3.2. Again,the phase term in front of Eq. 2.184 can often be omitted.

In Section 2.3.5, we also describe how to calculate the reflection and transmission from a nonuniform Bragg grating, that is, a grating which has chirp and/or apodization. The complete 2×2 transfer matrix for such a nonuniform grating can be expressed in terms of the numerically calculated reflection and transmission coefficients r and t:

$$\mathbf{T} = \frac{1}{t^* t} \begin{bmatrix} t & -tr^* \\ -t^* r & t^* \end{bmatrix} \quad , \tag{2.185}$$

The above equation can be derived from the time-reversal symmetry properties of transfer

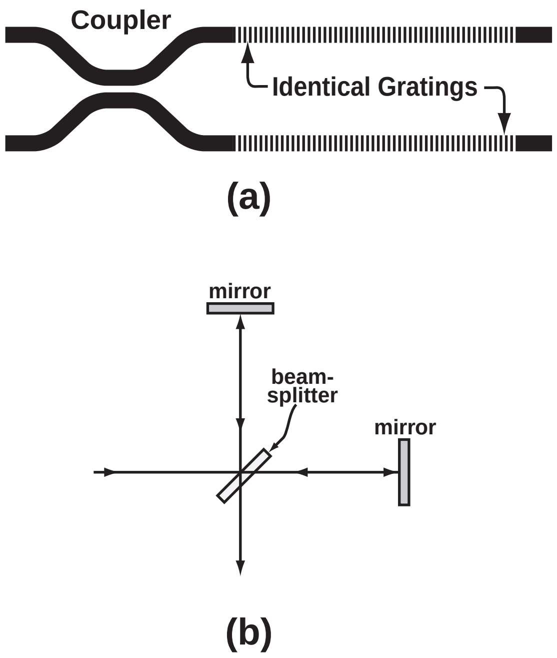

Figure 2.26: A diagram of the Michelson interferometer, as an integrated optical device (a) and as a free-space structure (b). The integrated device depicted in (a) is functionally equivalent to the free-space device depicted in (b). The integrated directional coupler takes the place of the beam splitter, and the integrated Bragg gratings act as wavelength-selective mirrors. (figs/2/michelson-schematic.eps)

matrices given in Eq. 2.177.

2.4.4 The Integrated Interferometer

The transfer matrix techniques described above provide all of the tools needed to analyze almost arbitrary combinations of couplers and gratings. We now describe how to use these techniques to analyze an integrated form of the Michelson interferometer.

Figure 2.26a depicts schematically the structure which we wish to consider. On the left-hand side is a conventional directional coupler which is nominally designed to provide 50% power splitting between the upper and lower waveguides. On the right hand side, there are two identical Bragg gratings formed in opposite arms of the interferometer. Fig. 2.26b depicts the equivalent free-space Michelson interferometer. In the integrated

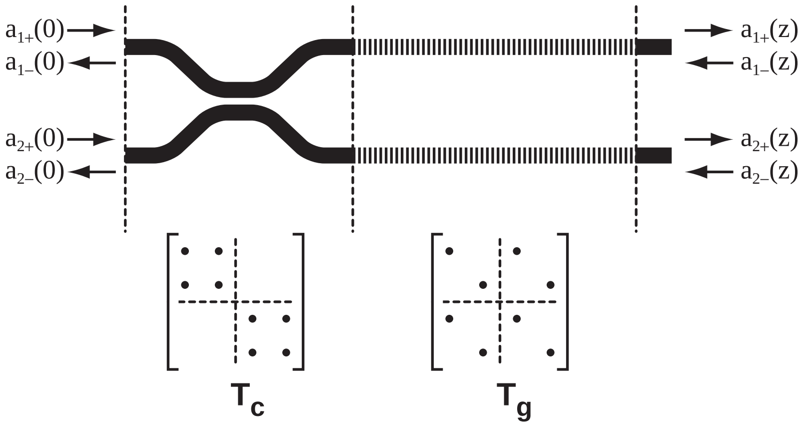

Figure 2.27: The transfer matrix method can be applied to the Michelson interferometer by simply multiplying the transfer matrices for each segment of the device. We depict here the shape of the transfer matrices for the directional coupler and for the identical Bragg gratings. (figs/2/michelson-matrices.eps)

device, the directional coupler serves the same purpose as the beam splitter, and the Bragg gratings serve as wavelength-selective mirrors.

Figure 2.27 illustrates how the transfer matrix approach can be applied to the integrated Michelson interferometer. Because there are two waveguides in the device, each of which can support a forward and backward mode, four mode amplitudes are needed to describe the system. As described earlier, the mode amplitudes are collected as a vector of four quantities:

$$\mathbf{a}(z) = \begin{bmatrix} a_{1+}(z) \\ a_{2+}(z) \\ a_{1-}(z) \\ a_{2-}(z) \end{bmatrix} . \tag{2.186}$$

The 4 × 4 transfer matrix representing the directional coupler is given in Eq. 2.182, but repeated here for reference:

$$\mathbf{T}_c = \begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^* \end{bmatrix}, \quad \mathbf{M} = e^{-j\beta L_c} \begin{bmatrix} \cos\phi & -j\sin\phi \\ -j\sin\phi & \cos\phi \end{bmatrix} . \tag{2.187}$$

where L_c is the total path length of either arm of the directional coupler. The 2 × 2 matrix $\mathbf{M}$ is a symmetric, unitary matrix which has the following properties:

$$\mathbf{M}^* = \mathbf{M}^\dagger = \mathbf{M}^{-1} \quad . \tag{2.188}$$

The 4 × 4 transfer matrix representing the identical Bragg gratings can be expressed in the following general form:

$$\mathbf{T}_g = \frac{1}{t^*t}\begin{bmatrix} t\mathbf{I} & -tr^*\mathbf{I} \\ -t^*r\mathbf{I} & t^*\mathbf{I} \end{bmatrix} \quad , \tag{2.189}$$

where $\mathbf{I}$ is the 2 × 2 identity matrix and the quantities r and t represent respectively the reflection and transmission coefficients of the Bragg gratings when light is incident from the left-hand side. For a uniform grating, these quantities would be given by Eqs. 2.128 and 2.129. For a nonuniform grating, they can be numerically computed using the techniques described in Section 2.3.5.

The complete transfer matrix for the integrated Michelson interferometer is found by simply multiplying the 4 × 4 transfer matrices for each portion of the device, in the reverse order:

$$\begin{aligned} \mathbf{T} &= \mathbf{T}_g\mathbf{T}_c \qquad (2.190)\\ &= \frac{1}{t^*t}\begin{bmatrix} t\mathbf{I} & -tr^*\mathbf{I} \\ -t^*r\mathbf{I} & t^*\mathbf{I} \end{bmatrix}\begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^* \end{bmatrix} \qquad (2.191)\\ &= \frac{1}{t^*t}\begin{bmatrix} t\mathbf{M} & -tr^*\mathbf{M}^* \\ -t^*r\mathbf{M} & t^*\mathbf{M}^* \end{bmatrix} \quad . \qquad (2.192) \end{aligned}$$

Using Eq. 2.168, this transfer matrix can be converted into an equivalent scattering matrix. The resulting scattering matrix for the structure depicted in Fig. 2.27 is:

$$\mathbf{S} = \begin{bmatrix} r\mathbf{M}^2 & t\mathbf{M} \\ t\mathbf{M} & -\frac{t}{t^*}r^*\mathbf{I} \end{bmatrix} \quad . \tag{2.193}$$

Of particular interest is the case where no light is launched into the waveguides from the right hand side. Using the scattering matrix given in Eq. 2.193, it is easy to solve for the reflected and transmitted mode amplitudes. For example, the signal emerging from the left edge of the device is given by:

$$\mathbf{a}_-(0) = r\mathbf{M}^2\mathbf{a}_+(0) \tag{2.194}$$

$$\begin{bmatrix} a_{1-}(0) \\ a_{2-}(0) \end{bmatrix} = re^{-j2\beta L_c} \begin{bmatrix} \cos 2\phi & -j\sin 2\phi \\ -j\sin 2\phi & \cos 2\phi \end{bmatrix} \begin{bmatrix} a_{1+}(0) \\ a_{2+}(0) \end{bmatrix} \quad . \tag{2.195}$$

The factor r in the above equation is no more than the reflection spectral response of the identical Bragg gratings. The 2×2 matrix (with the phase term) is equivalent to the transfer matrix for a directional coupler with twice the effective length of the original coupler. Signals which enter the device from the left pass once through the directional coupler and are then reflected by the Bragg gratings. The reflected signals once again pass through the coupler. Thus, the result of Eq. 2.194 could almost be predicted without using the transfer matrix methods: the signal passes through the directional coupler twice (which accounts for the factor of 2 in the matrix terms of Eq. 2.194), and the filtering action of the Bragg grating is contained entirely in the first factor.

Usually, the directional coupler is designed so that after two passes the light is completely transferred from one waveguide to the other. (This is accomplished by selecting $\phi = \pi/4$.) With this choice, if light enters the device in the upper waveguide, the reflected signal will emerge from the lower waveguide. In this way, the Michelson interferometer provides a way of spatially separating the incident input light from the reflected output light.

Any light which is not reflected by the Bragg grating emerges in the two waveguides on the right side of the device:

$$\mathbf{a}_+(z) = t\mathbf{M}\mathbf{a}_+(0), \tag{2.196}$$

$$\begin{bmatrix} a_{1+}(z) \\ a_{2+}(z) \end{bmatrix} = te^{-j\beta L_c} \begin{bmatrix} \cos \phi & -j\sin \phi \\ -j\sin \phi & \cos \phi \end{bmatrix} \begin{bmatrix} a_{1+}(0) \\ a_{2+}(0) \end{bmatrix} \quad . \tag{2.197}$$

Notice that in transmission, the light only passes through the coupler once, and therefore the emerging light is typically divided between the two waveguides.

Mach-Zehnder Interferometer

The transmitted light can be recombined by adding an additional coupler to the structure, as depicted in Fig. 2.28. The structure is the integrated form of a Mach-Zehnder interferometer. The analysis described above can be extended to obtain the transfer matrix for the

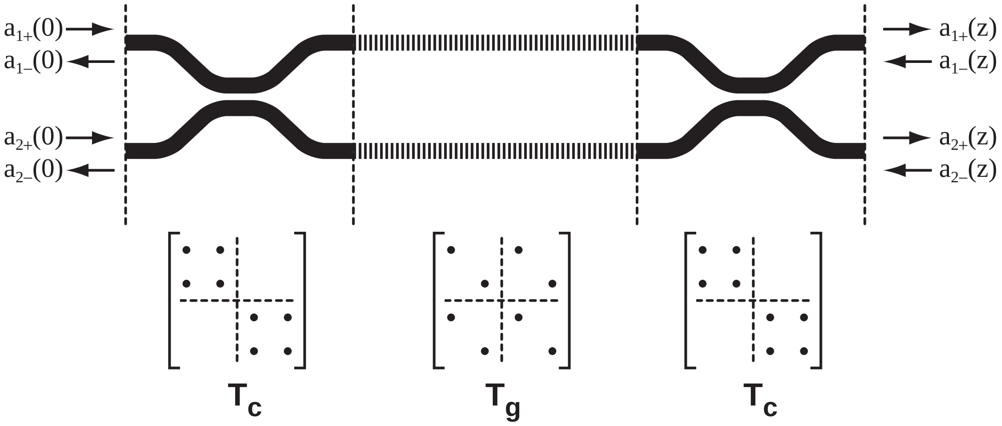

Figure 2.28: Diagram of an integrated Mach-Zehnder interferometer. This device is similar to the Michelson interferometer, but a second coupler has been added on the right hand side to recombine the signals transmitted by the Bragg gratings.
(figs/2/mach-zehnder-matrices.eps)

Mach-Zehnder interferometer. The transfer matrix for the device is now the product of the 3 transfer matrices representing the right coupler, the gratings, and the left coupler:

$$\mathbf{T} = \frac{1}{t^*t}\begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^* \end{bmatrix}\begin{bmatrix} t\mathbf{I} & -tr^*\mathbf{I} \\ -t^*r\mathbf{I} & t^*\mathbf{I} \end{bmatrix}\begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^* \end{bmatrix} \quad . \tag{2.198}$$

which simplifies to:

$$\mathbf{T} = \frac{1}{t^*t}\begin{bmatrix} t\mathbf{M}^2 & -tr^*\mathbf{I} \\ -t^*r\mathbf{I} & t^*[\mathbf{M}^*]^2 \end{bmatrix} \quad . \tag{2.199}$$

As before, we can use Eq. 2.169 to compute the corresponding scattering matrix:

$$\mathbf{S} = \begin{bmatrix} r\mathbf{M}^2 & t\mathbf{M}^2 \\ t\mathbf{M}^2 & -\frac{t}{t^*}r^*\mathbf{M}^2 \end{bmatrix} \quad . \tag{2.200}$$

When there is no light entering the device from the right hand side, the signals reflected out the left-hand side can be written as,

$$\mathbf{a}_{-}(0) = r\mathbf{M}^2\mathbf{a}_{+}(0), \tag{2.201}$$

$$\begin{bmatrix} a_{1-}(0) \\ a_{2-}(0) \end{bmatrix} = re^{-j2\beta L_c} \begin{bmatrix} \cos 2\phi & -j\sin 2\phi \\ -j\sin 2\phi & \cos 2\phi \end{bmatrix} \quad . \tag{2.202}$$

and signals transmitted to the right hand side are given by,

$$\mathbf{a}_{+}(L) = t\mathbf{M}^2\mathbf{a}_{+}(0), \tag{2.203}$$

$$\begin{bmatrix} a_{1+}(L) \\ a_{2+}(L) \end{bmatrix} = te^{-j2\beta L_c} \begin{bmatrix} \cos 2\phi & -j\sin 2\phi \\ -j\sin 2\phi & \cos 2\phi \end{bmatrix} \quad . \tag{2.204}$$

If the directional couplers are designed appropriately (with $\phi = \pi/4$), the Mach-Zehnder interferometer provides full power transfer from one waveguide to the other in both reflection and transmission.

Interferometer with Imbalanced Arms

The interferometer devices considered in this section are, like all interferometers, phase-sensitive. That is, they require that the optical path lengths be equal in the upper and lower arms of the device. If there is some difference in path length, the signals which are reflected by the Bragg gratings will not properly recombine in the desired port. We will now show how the effect of unequal path lengths can be easily calculated using transfer matrices.

Figure 2.29 depicts the structure of an imbalanced Michelson interferometer. The structure is identical to that of Fig. 2.27, with the exception that the upper arm of the device has been made longer by an amount ΔL. (Of course, a phase imbalance could also be achieved by changing the refractive indices in one arm of the device.) The phase imbalance between the upper and lower arms can be modeled with the following transfer matrix:

$$\mathbf{T}_\theta = \begin{bmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^* \end{bmatrix}, \quad \mathbf{D} \equiv \begin{bmatrix} e^{-j\theta} & 0 \\ 0 & 1 \end{bmatrix} \quad , \tag{2.205}$$

where θ represents the total phase difference between the two arms. If the phase imbalance is achieved by making one arm longer than the other, as illustrated in Fig. 2.29, θ is given

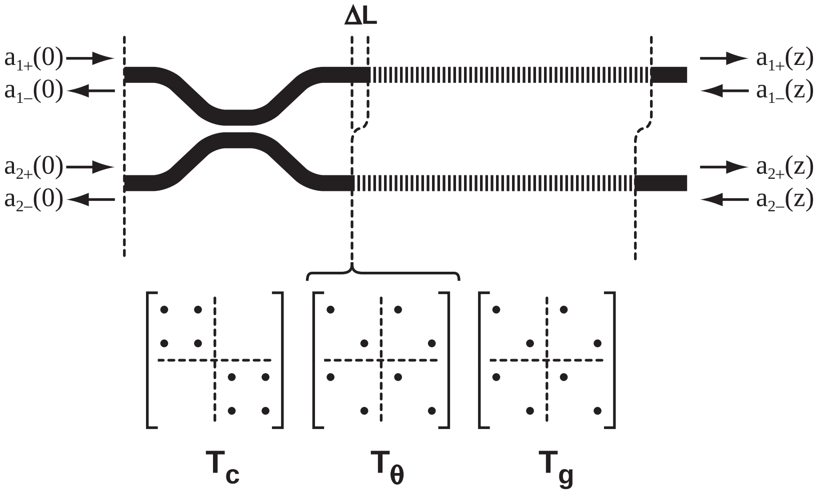

Figure 2.29: A diagram of a Michelson interferometer with mismatched arms. The effect of differing arm lengths can be included in the model by inserting a third transfer matrix which represents the phase difference between the upper and lower arms. (figs/2/mismatched-michelson-matrices.eps)

by:

$$\theta = \beta \Delta L \quad . \tag{2.206}$$

The transfer matrix for the imbalanced Michelson interferometer is found by multiplying three transfer matrices for the grating, phase discrepancy, and coupler:

$$\begin{aligned} \mathbf{T} &= \mathbf{T}_g \mathbf{T}_\theta \mathbf{T}_c & (2.207)\\ &= \frac{1}{t^* t} \begin{bmatrix} t\mathbf{I} & -tr^*\mathbf{I} \\ -t^* r\mathbf{I} & t^*\mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^* \end{bmatrix} \begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{M}^* \end{bmatrix} & (2.208)\\ &= \frac{1}{t^* t} \begin{bmatrix} t\mathbf{DM} & -tr^*\mathbf{D}^*\mathbf{M}^* \\ -t^* r\mathbf{DM} & t^*\mathbf{D}^*\mathbf{M}^* \end{bmatrix} \quad . & (2.209) \end{aligned}$$

The corresonponding scattering matrix is,

$$\mathbf{S} = \begin{bmatrix} r\mathbf{MDDM} & t\mathbf{DM} \\ t\mathbf{MD} & -\frac{t}{t^*} r^* \mathbf{I} \end{bmatrix} \quad . \tag{2.210}$$

As before, we can use the scattering matrix to compute the reflected mode amplitudes in terms of the incident mode amplitudes (again, assuming no light enters the structure from the right hand side.)

$$\mathbf{a}_-(0) = r\mathbf{MDDM}\mathbf{a}_+(0) \tag{2.211}$$

Once again, this result could almost be deduced by inspection. The incident signal passes through the coupler once, and then through the phase shift twice (once forward and once backward), and then goes back through the coupler again. This accounts for the matrix factor **MDDM** in Eq. 2.211. The reflectivity of the Bragg grating is entirely encompassed in the first factor r. Expanding Eq. 2.211 gives:

$$\begin{bmatrix} a_{1-}(0) \\ a_{2-}(0) \end{bmatrix} = re^{-j(2\beta L_c + \theta)} \begin{bmatrix} \cos 2\phi \cos\theta - j\sin\theta & -j\sin 2\phi\cos\theta \\ -j\sin 2\phi \cos\theta & \cos 2\phi\cos\theta + j\sin\theta \end{bmatrix} \begin{bmatrix} a_{1+}(0) \\ a_{2+}(0) \end{bmatrix} \tag{2.212}$$

If we further assume that light is launched into the upper waveguide (waveguide 1), then the relative power reflected in the lower waveguide is described by:

$$\left| \frac{a_{2-}(0)}{a_{1+}(0)} \right|^2 = |r|^2 \sin^2 2\phi \cos^2\theta \quad . \tag{2.213}$$

The above equation shows that in order to completely divert the reflected signal to the lower waveguide, two conditions must be satisfied: (1) the directional coupler must provide 50% coupling, i.e., $\phi = \pi/4$, and (2) the optical path lengths must be matched, or mismatched by an integral number of half-wavelengths, i.e., $\theta = n\pi$.

2.5 Summary

The design of integrated optical devices is a hierarchical process, and the organization of this chapter was meant to follow this hierarchy. We began with Maxwells equations, and derived some of the fundamental properties of propagating electromagnetic modes. From there, we proceeded to describe how Maxwells equations can be simplified using coupled mode theory. Coupled mode theory replaces Maxwells equations by a set of coupled differential equations for the mode amplitudes. The electromagnetic analysis of waveguide modes is critical to coupled mode theory, because the coupling coefficients depend upon the electromagnetic modes through overlap integrals. The solution to the coupled mode equations can always be expressed in terms of transfer matrices. In the final portion of this chapter, we described how to analyze interconnected waveguides, gratings, and couplers using transfer matrices.

Chapter 3

Fabrication

The previous chapter provided all of the theoretical background needed to understand, design, and analyze waveguides, couplers, and Bragg gratings. This chapter marks a turning point from theory to practice. The following sections describe the fabrication techniques which we have developed to build integrated Bragg grating devices described in Chapter 2.

Many of the fabrication techniques commonly used in building integrated optical devices are borrowed, or adapted from the semiconductor industry. Indeed, many integrated optics research group specifically select materials, waveguides, and processing techniques which are compatible with existing semiconductor-processing equipment, and many groups cite compatibility with current processing techniques as a selling point for their technology [30, 31][1]. As we will show, the fabrication of integrated optical devices presents several unique challenges that are quite different from those encountered in electronic devices. Throughout this chapter, we hope to illuminate these challenges and describe some innovative solutions.

Section 3.1 begins by describing waveguide technology. We will summarize prior work on passive integrated optical devices and describe the processing techniques developed to build waveguides and couplers.

In Section 3.2, we will shift our focus from waveguides to Bragg gratings, for which the feature sizes are generally about one order of magnitude finer. We will describe and demonstrate a host of lithographic techniques which can be used to pattern the fine-period

---

[1] Often this equipment is handed down from semiconductor facilities after upgrades, as is the case with many of the tools used to build the devices described in this work.

Bragg gratings required for this work, and discuss the benefits and shortcomings of each method.

Section 3.3 addresses the question of how waveguides and gratings can be combined to build integrated Bragg gratings. Specifically, we will describe the implementation of a flexible lithographic technique for patterning gratings atop relatively tall waveguides. Additionally, we will describe techniques for ensuring accurate alignment of the waveguides and gratings. Finally, we will discuss the critical problem of maintaining the grating integrity during subsequent overgrowth steps.

3.1 Fabrication of Waveguides and Couplers

This section is devoted to the fabrication techniques for waveguides and directional couplers. We have grouped waveguides and couplers together because the lithographic techniques used to build them are almost identical. Generally, a waveguide or coupler can be built using a single patterning step, i.e., only one layer of photolithography is needed. Alignment, if it is needed at all, can easily be accomplished using standard lithographic alignment techniques. We will first provide a summary of the many different waveguide compositions, structures, and fabrication techniques reported in the literature. Next, we will specifically describe the techniques which were used to build integrated waveguides and couplers in doped-SiO_2, and silicon-on-insulator (SOI) ridge waveguides.

3.1.1 Integrated Waveguide Technology

There are a number of candidate material systems for building integrated optical waveguides, each with advantages and disadvantages. One might be led to believe, from the analysis of Section 2.1, that the only required feature of an integrated waveguide is a high-index core material surrounded by low-index cladding. In practice, there are many additional criteria which determine how useful a waveguide technology is. Some of the desirable properties of optical waveguides are:

- Single-mode operation
- Efficient coupling to optical fiber
- Low propagation loss

- Low bending loss (or equivalently, tight bending radius)
- Low birefringence and low polarization-dependent loss
- Easy to integrate

Single-mode operation means that there is only one electromagnetic mode for the waveguide. In a multi-mode waveguide, each electromagnetic mode propagates with a different speed, which leads to undesirable signal dispersion if the waveguide is used to transmit data. For this reason, most integrated waveguides of interest support only one bound mode for each polarization state.

Optical fiber is currently the medium of choice for transmitting optical signals. One might argue that it is only because of the dominance of optical fiber that integrated optics is such an active field of research. Although many people envision futuristic integrated optical circuits consisting of complex interconnected devices, at present most integrated optical devices are single elements inserted between optical fiber. As such, one of the most important criteria in selecting an optical waveguide is how well one can couple light to and from a standard optical fiber. There are two requirements for efficient coupling to an optical fiber: (1) the optical mode of the integrated waveguide must have a similar size and shape to that of the optical fiber, and (2) the effective refractive index of the integrated waveguide must be close to that of the optical fiber. The coupling efficiency can be estimated using a simple overlap integral between the mode of the optical fiber and that of the integrated waveguide [11, 30, 39]:

$$\eta = \frac{\left| \iint \phi_0(x,y)\phi_1(x,y)\,dx\,dy \right|^2}{\iint |\phi_0(x,y)|^2\,dx\,dy \cdot \iint |\phi_1(x,y)|^2\,dx\,dy} \tag{3.1}$$

where ϕ_0 represents the electromagnetic mode of the fiber and ϕ_1 is that of the waveguide. (For an optical fiber, the mode ϕ_0 can be well represented by a single scalar quantity as described in Section 2.1. This is often the case for integrated waveguides, too, especially for integrated waveguides which are designed to match to an optical fiber.) Additionally, many also insert a factor in front of Eq. 3.1 to account for Fresnel reflection when the refractive indices are not matched. Although many researchers have developed novel techniques for improving coupling efficiencies [77, 78, 79], the problem of getting light from a fiber into waveguide remains one of the critical challenges to integrate optical devices.

Clearly, one wants to minimize the amount of propagation loss for an optical waveguide. Propagation loss has several sources. If the materials comprising the waveguide are not optically transparent, the waveguide will exhibit propagation loss. Also, even if the materials are completely transparent, any roughness or inhomogeneity incurred during lithographic processing can scatter light out of the waveguide, leading to additional loss. The latter form of loss generally increases in proportion to the core-clad index difference. Of course, the amount of loss which is tolerable depends upon the overall size of the structure. For long waveguides, the propagation loss must be kept low whereas for tightly-confined compact microoptic devices higher propagation loss can often be tolerated.

Bending loss refers to the leakage of light from the waveguide as it bends in a curved trajectory. There are many excellent models for calculating bending loss in optical waveguides [80, 81, 82, 83]. Generally, the amount of bending loss decreases exponentially with increasing radius of curvature. Also, waveguides with high index contrast (core-clad index difference) exhibit lower bending loss than waveguides with low index contrast. Thus, the desire for low bending loss often competes with the desire for high coupling efficiency to an optical fiber.

Birefringence and *polarization-dependent loss* simply refer to the difference in propagation constant and loss between the fundamental TE and TM modes. Again, the desire for low polarization dependence stems from the prevalence of optical fiber as a transmission medium. Because the optical fiber has no preferred polarization axis, an integrated optical device should ideally provide polarization independent performance as well.

Finally, *integrability* simply refers to the ease with which a waveguide can be integrated with other optical or electronic components.

Depending upon the application, there are other desirable characteristics in addition to those discussed here including, for example, photosensitivity and acoustooptic sensitivity.

Commonly Used Waveguide Systems

One of the earliest systems used for integrated optical devices is a combination of SiO_2 and Si_3N_4. One advantage of this material system is that it can be easily integrated on a silica substrate. The structure typically begins with a lower buffer (cladding) layer of SiO_2, grown by high-pressure thermal oxidization, wet oxidization, or plasma enhanced chemical vapor deposition (PECVD). The core layer of Si_3N_4 is typically deposited using low-pressure chemical vapor deposition (LPCVD) or via nitridation of the oxide in an ammonia atmosphere [84, 85, 86, 87]. A protective top layer of SiO_2 can be deposited, although it is not essential. Because of the high refractive index difference ($n \simeq 2$ for Si_3N_4, vs n = 1.46 for SiO_2), the core layer must be relatively thin in order to ensure single-mode operation. Typical core thicknesses range from 100-200 nm. Longitudinal confinement can be achieved either by directly etching a shallow ridge directly on the nitride layer, or by forming a ridge in the upper oxide cap layer. The thickness of the underlying oxide buffer layer must be large enough to ensure that the optical mode remains well confined in the oxide-nitride layers and does not leak into the high-index silicon substrate. Waveguides based upon this geometry have been reported with propagation losses as low as 0.3 dB/cm [85]. Unfortunately, because of the small core-thickness and high index contrast these devices are not well matched to an optical fiber. Moreover, they exhibit significant polarization dependence.

Many researchers have sought to build integrated waveguides in polymers [88]. Although there have been reports of low-loss propagation in polymers, the minimum loss attainable in polymers is typically much higher than that of dielectric waveguides. Another significant problem with polymer waveguides is the difficulty in patterning them. Most lithographic techniques involve coating the substrate with a photosensitive resist, and often the solvents used in the resist are incompatible with polymeric layers. To circumvent this problem, some have actually used the photoresist itself as a waveguide which can be conveniently patterned via lithography[89]. Other approaches include using lithographic techniques to directly modulate the refractive index of the polymer [90].

Compound semiconductors are often used in integrated optical components, because they can be conveniently integrated with active elements such as lasers, and photodetectors. The materials most commonly used for constructing such integrated waveguides are $Al_{1-x}Ga_xAs$ and $In_{1-x}Ga_xAs_{1-y}P_y$. The index of refraction can be controlled by adjusting the mole fraction (x,y) of the atomic constituents. For most III-V compound semiconductors, the refractive indices range from 3.1–3.5. Because of this, it is difficult to build III-V waveguides which couple efficiently to an optical fiber.

Lithium niobate is an attractive material for fabrication of optical waveguides because of its acoustooptic and electrooptic properties [91]. These properties make $LiNbO_3$ ideal for constructing optical modulators and switches. Unfortunately, $LiNbO_3$ loses these valuable properties when it is not grown in single crystal form, which makes it difficult to integrate $LiNbO_3$ on a Si substrate. Waveguides can be formed in crystaline $LiNbO_3$ by a process of proton exchange, or metal diffusion, but these methods typically yield a graded-index profile which cannot be controlled to the same degree as an etched waveguide structure.

Glass or silica (SiO_2) are often used as the raw material for fabricating passive optical waveguides. One advantage of these materials is that they are very similar to the materials which make up optical fibers. Because of the close match in index of refraction, it is possible to couple efficiently from fiber into a glass waveguide. As evidenced by the very low propagation loss of optical fibers, glass is intrinsically low loss compared with other waveguide materials. Although the attenuation in planar glass waveguides is still orders of magnitude higher than for optical fiber, it is still much lower than in most other dielectric, semiconductor, or polymer waveguides. Glass waveguides are typically fabricated on Si substrates, which are readily available in large sizes and are less expensive than III-V semiconductor substrates.

Researchers at LETI, in France, have fabricated SiO_2 channel waveguides using plasma-enhanced chemical vapor deposition (PECVD) and reactive-ion etching [84, 92]. They begin with an 8-15 µm thick lower cladding layer of SiO_2. A core layer, doped with phosphorus, is formed by including phosphine in the gas mixture during deposition. The refractive index of the core can be controlled by adjusting the amount of phosphorus doping. The maximum reported refractive index for phosphorus-doped glass was 1.53 (compared with 1.46 for undoped SiO_2.) However, in order to achieve optimal coupling to a fiber, the core should have an index of refraction that is only 0.3-0.8% higher than the surrounding cladding, which corresponds to a phosphorus content of less than 1% by weight. One of the challenges to PECVD deposition is controlling the phosphorus concentration and uniformity in this low Δn regime. Using these techniques, the LETI group achieved propagation losses ranging from 0.1–0.2 dB/cm.

Researchers at Bell Labs, employ a similar procedure to construct glass optical waveguides [30, 93]. In their work, the lower cladding of SiO_2 is deposited using a high-pressure thermal oxidization (HiPOX). This process is limited to an oxide thickness of approximately 15 µm, because the oxidation rate decreases rapidly as the SiO_2 layer gets thicker. The core layer is formed using CVD, again with phosphorus doping to increase the refractive index relative to the underlying oxide. The typical core composition is approximately 7% P_2O_5 by weight, which corresponds to a refractive index contrast of 0.6%. The core is then patterned and etched using photolithography with a tri-layer resist and reactive-ion etching. A top cladding layer is deposited using LPCVD (TEOS) with both boron and phosphorus dopants. This borophosphosilicate glass can be made to have a refractive index which is matched to SiO_2, but a lower glass transition temperature than the underlying layers. This enables the top cladding to be deposited without softening the core layer. The composition of the top BPTEOS letter is approximately 2.5% P_2O_5 by weight and 5% B_2O_3 by weight.

Researchers at NTT and elsewhere developed a process for constructing glass waveguides using flame hydrolysis deposition [31, 94, 95, 96, 97, 98]. The process of flame hydrolysis is also used to generate the glass preforms from which optical fiber is drawn. In flame hydrolysis deposition, gaseous $SiCl_4$ is burned in an oxygen / hydrogen torch, while the flame is traced over a substrate. This produces a porous layer of "soot" comprised of fine SiO_2 particles on the substrate. When this layer is heated in a furnace, it flows to form a solid, high-quality optical layer. The index of refraction and flow temperature of the deposited layer can be adjusted by introducing $TiCl_4$ or $GeCl_4$ into the gas mixture. Using this technique, researchers at NTT demonstrated channel waveguides with propagation losses as low as 0.01 dB / cm [99, 97], and fiber-coupling losses of only 0.05 dB [96].

Doped glass waveguides typically require relatively thick oxide cladding layers in order to isolate the mode from the high-index Si substrate. One way to alleviate this problem is to use an antiresonant structure consisting of a Si/SiO_2 multilayer stack as a lower cladding. Duguay et al. describe successful fabrication of such an antiresonant reflecting optical waveguide (ARROW)[100]. Unfortunately, because this approach relies on the polarization-dependent reflection from the Si/SiO_2 interface, the waveguides are only transparent for one polarization.

Another useful platform for integrated optical waveguides is silicon-on-insulator, or SOI. Silicon, although reflective at visible wavelengths, is transparent to infrared light. Moreover, it is clearly well-suited to integration with electronic devices. However, in order to build a silicon waveguide, a low-index cladding material must be positioned between the waveguide and the substrate in order to prevent the optical mode from leaking out. Fortunately, the semiconductor industry has developed many techniques for forming a buried oxide insulator on a silicon substrate. The simplest technique for forming silicon-on-insulator (SOI) involves implanting oxygen ions directly into the crystalline silicon. This technique can be used to form relatively thin silicon layers. For thicker layers, a bond-and etchback technique can be used [101], or the silicon layer thickness can be augmented by epitaxy after the oxygen implant [102, 32]. Although the high index contrast between silicon and oxide / air seems to imply a small mode size, relatively large optical modes can be confined in silicon ridge waveguides without sacrificing single-mode operation [34]. Single-mode waveguides with propagation loss as low as 0.1 dB / cm have been achieved using this method [33]. Although silicon-on-insulator is a relatively new platform for integrated optics, there are now commercial companies which base their product line on this promising technology [103].

3.1.2 Pattern Generation and Waveguide CAD

The waveguides and couplers described in this work typically consist of long, gently curved structures which gradually approach and separate as they traverse a chip. This type of pattern contrasts sharply with the dense patterns of horizontal and vertical lines which define, for example, a silicon microchip. Moreover, in order for a device to achieve the proper coupling and phase, the waveguide separations and path lengths must be carefully controlled over the entire length of the device. Because of this requirement, most lithographic CAD packages are unsuitable for generating waveguide patterns. One of the challenges in constructing waveguide devices is to translate these complex geometrical shapes into a sequence of simple shapes that can be written onto a photomask.

Specifically, the electron-beam lithography systems which are used to write photomasks require that the pattern be organized into rectangular shapes which lie on a Cartesian grid. The minimum addressable size is called the grid-size or address unit. In designing a photomask for optical waveguides, the complicated curved and sloped waveguides must be approximated by a series of rectangular shapes lying on this Cartesian grid.

To simplify this task, we have developed specialized flexible software routines which will construct the quantized representation of a device based upon a high-level description of the device geometry. We will outline the organization of the software to illustrate how it can be used (and extended) to construct relatively complex waveguide geometries.

Figure 3.1a depicts the structure of a simple integrated coupler, with all of the relevant dimension specified. The 7 labeled parameters, d_1, d_0, w, L_c, R, θ, and L, are sufficient to completely specify the geometry of the structure. These parameters are provided to the software through an input text file, along with the few general parameters such as the desired address unit. Dimensions of length can be specified in any of the standard SI units; when the input file is processed, the software automatically converts all dimensions into address units. Appendix B gives an example of an input file used by the layout program.

The first level of processing that the program performs is to break down the device into a series of primitive segments, as shown in Fig. 3.1b. Although customized segment types can be easily designed by the user, most patterns can be adequately described in terms of just two segment types: lines and arcs. The segments which make up each waveguide are ordered into an array which describes where they are placed and how they are connected. A device is described by one or more of these segment arrays.

In the final level of processing, the program sequentially steps through each segment

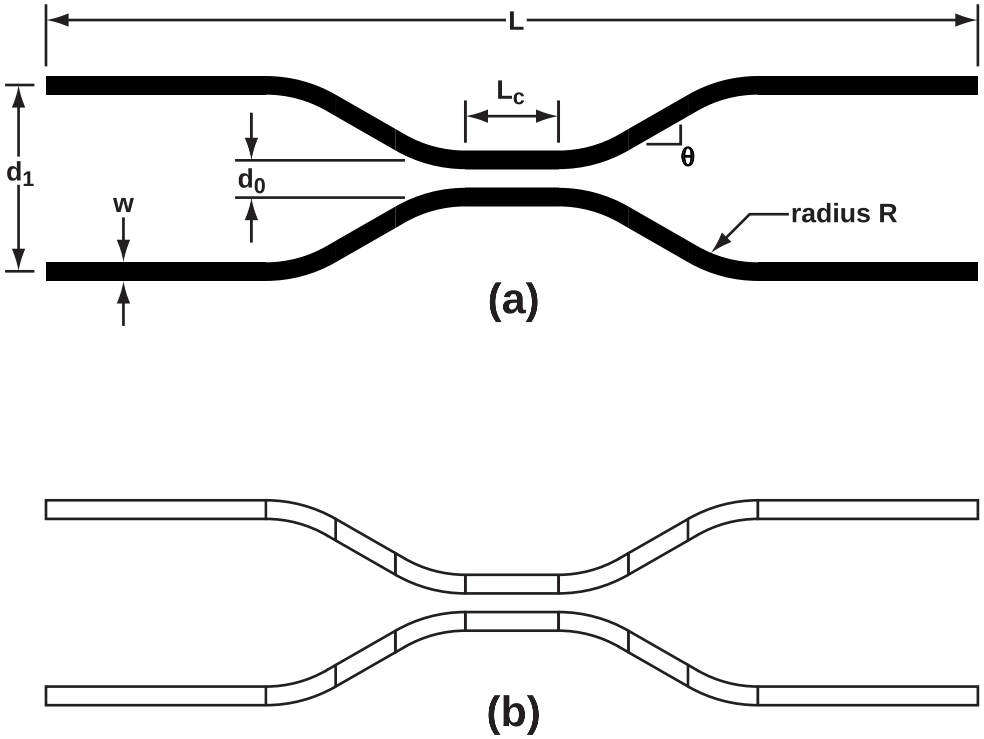

Figure 3.1: (a) Diagram of an integrated directional coupler, with all of the necessary dimensions labeled. (b) The coupler may be broken down into a sequence of primitive “segments”. In this example, the segments are either straight waveguide pieces or arcs. (figs/3/coupler-segments.eps)

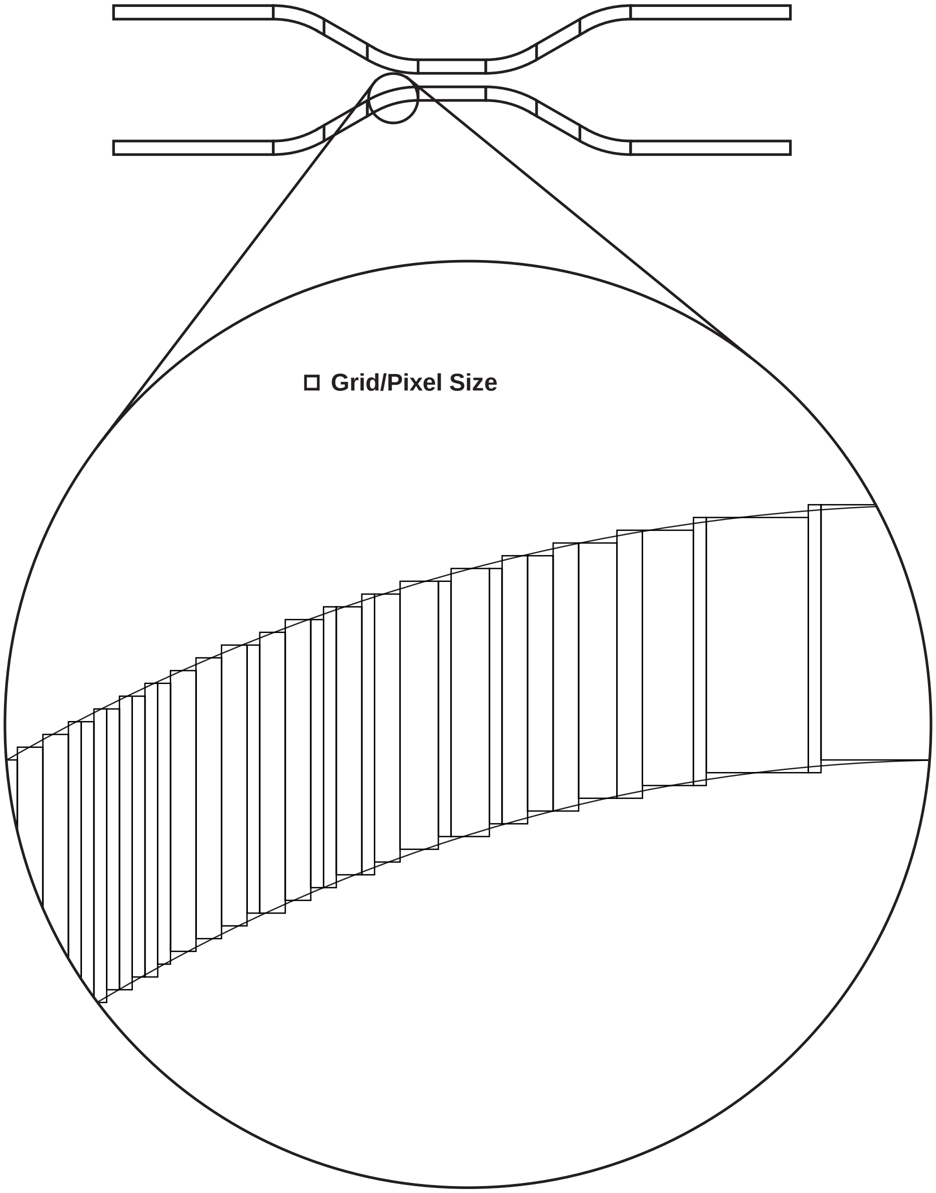

Figure 3.2: The quantized representation of a curved waveguide segment. The mathematically defined arc is approximated by a series of rectangular boxes which lie on a fixed Cartesian grid. Where possible, adjacent boxes are combined in order to most efficiently describe the pattern with the minimum number of boxes. (figs/3/quantized-coupler.eps)

and generates a series of Cartesian boxes to approximate the mathematically-defined structure. Figure 3.2 illustrates the pattern of quantized boxes used to approximate a curved waveguide with a width of 20 address units. Notice that where possible, adjacent boxes are combined in order to most efficiently describe the structure with the smallest number of boxes. The output file is formatted as a KIC / CIF file, which can be viewed and edited by a freely-available CAD program called KIC, or converted to GDSII or MEBES format for mask generation. To assist with the layout, the program can also optionally generate an encapsulated PostScript (EPS) representation of the segments, which can be printed or viewed with a number of common software products.

One benefit of this organization is that it allows the waveguide designer to easily construct new device types without rewriting the core software routines. The only portion of code which must be provided is a high-level routine which instructs the program how to generate the appropriate series of arcs and lines based upon a few user-specified geometrical parameters. It also makes it relatively easy to refine or adjust waveguide parameters such as waveguide width or grid size without starting over.

3.1.3 Glass Channel Waveguides and Couplers

In our initial experiments with glass waveguides, we used germanium-doped SiO_2 structures deposited via flame hydrolysis. These materials are now commercially available in a variety of core thicknesses and compositions. For our devices, the wafers provided by the manufacturer consist of a 20 μm-thick lower cladding layer of undoped SiO_2, topped with a 6.6 μm thick core layer of germanium doped SiO_2, as shown in Fig. 3.3a. The nominal index difference was 0.3%, which corresponds to a Ge mole fraction of approximately 3%. The process of flame hydrolysis only deposits glass on one side of the wafer, which can lead to significant stress and bowing when the wafer is cooled after annealing. To minimize this effect, the suppliers use 1 mm thick 4" silicon wafers. (Quartz substrate wafers can also be used.)

The waveguides were formed by selectively etching away the core material using reactive-ion etching, as depicted in Fig. 3.3. In order to fully removes the core layer, a relatively deep (6-7 μm) edge is needed. Although plasma etching of oxides is a well studied process in semiconductor processing, the oxide layer's used in electronic devices are typically an order of magnitude thinner than those used for optical waveguides.

The most suitable plasma chemistries for etching SiO_2 and other doped glasses are CHF_3 or a mixture of CF_4 and H_2 [104, 105, 106]. The masking material can be photoresist,

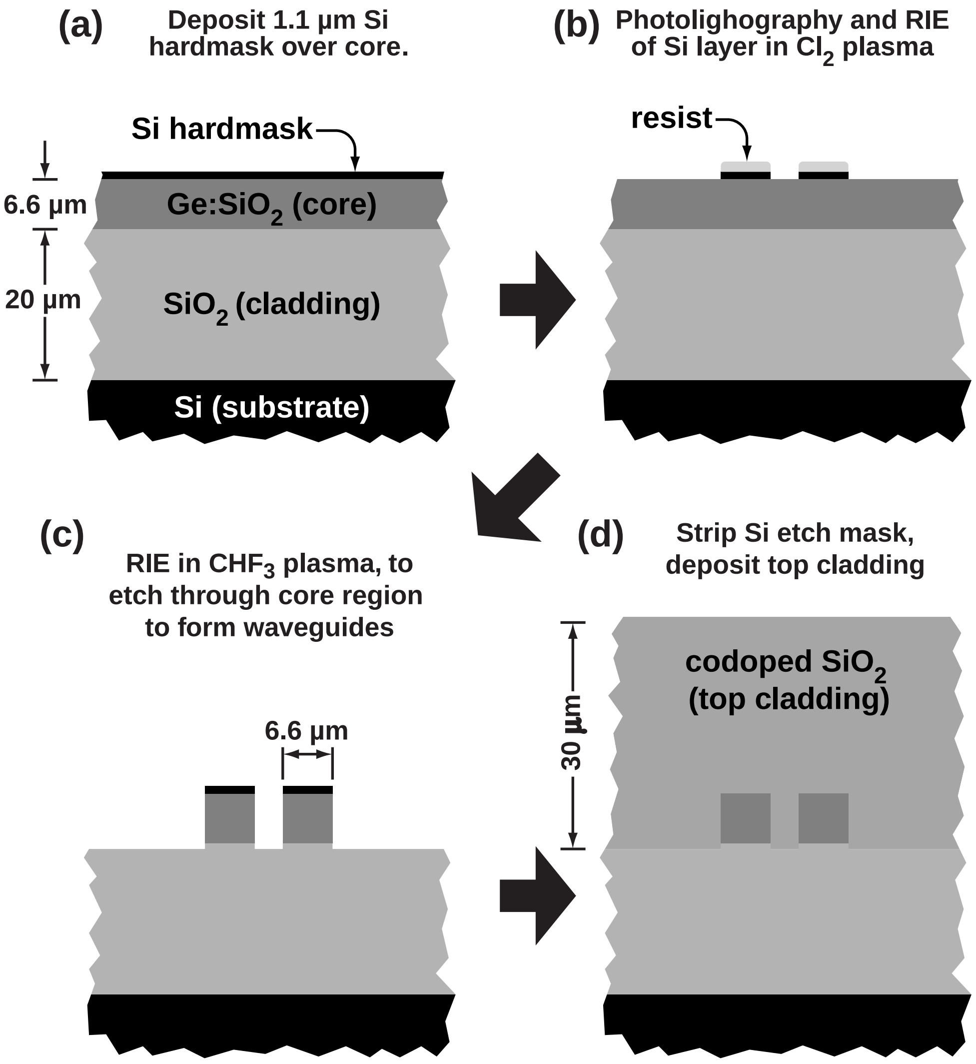

Figure 3.3: Outline of the process used to pattern channel waveguides and couplers in doped-glass. (figs/3/glass-wg-process.eps)

metal, or a material such as silicon [106, 30, 31].

Because of the deep etch required to form the waveguides, even thick photoresist masks typically do not hold up well, resulting in sloped sidewalls and increased edge roughness [106]. Researchers at Lucent use a tri-layer process to pattern the etch mask. The initial layer consist of about 2-3 μm of photoresist which is hardened by a high-temperature bake. Onto this a thin oxide interlayer is deposited, followed by a normal photoresist imaging layer. After exposing and developing the top imaging layer, the pattern is transferred, first to the oxide interlayer and then through the thick photoresist layer.

Metals can also be used for the etch mask, but they often prove problematic because of the difficulty in patterning thick metallic layers. In a conventional RIE system, it is difficult to achieve SiO_2:metal etch selectivities higher than about 20:1. Thus, the metal thickness must be at least 300 nm in order to withstand a 7 μm SiO_2 etch. Such thick metallic layers often have high stress which can lead to problems with adhesion. In our initial attempts at waveguide fabrication, we used a tungsten hard mask formed via sputter deposition. The waveguide patterns were transferred to this tungsten mask using photolithography and wet etching. The etching selectivity was found to be 15:1, but the process of wet etching led to poor control of the feature size and unacceptable roughness in the waveguide edges [39].

Researchers at NTT use a silicon layer as the hard mask. Silicon has the advantage that relatively thick silicon layers can be patterns using dry etching in Cl_2 or SF_6 plasmas. Also, reasonably high SiO_2:Si selectivities can be obtained by etching in CHF_3. For these reasons, we chose to use silicon as a hard mask when etching the waveguides and couplers.

Returning to Fig. 3.3, a uniform layer of amorphous silicon deposited initially using RF sputter deposition in an argon plasma. We used an MRC sputtering system with 10 mT pressure and 200 W of RF power for a duration of 90 min., which gives a silicon film thickness of approximately 1.2 μm at the center of a 4 inch wafer, and 0.9 μm near the wafer perimeter. Next, we spin on a layer of AZ5214 photoresist to a thickness of 1 μm and expose the waveguide patterns using ultraviolet contact lithography. (Note, for a positive-working resist this step requires a clear-field photomask.)

The waveguide pattern is then transferred to the silicon by reactive-ion etching in a chlorine (Cl_2) plasma. For this process, we used a modified Perkin Elmer system with a chamber pressure of 20 mT and an RF power of 75 W, which results in a silicon etch rate of 50-60 nm / min.. The chlorine plasma exhibits extremely high Si:SiO_2 selectivity, hence the patterns can be substantially over etched without significantly etching the underlying glass. This degree of tolerance is important because the silicon film thickness is not completely uniform across the wafer, nor is the plasma etch rate. It is relatively easy to determine whether the silicon has completely cleared using optical microscopy: when the silicon film is etched entirely it is possible, by decreasing the working distance, to see a second image of the silicon hard mask reflected by the silicon substrate 26 μm below.

After patterning the silicon hard mask, the sample is ready for etching the waveguides. If desired, the residual photoresist can be removed, to obtain an unambiguous measurement of the silicon film thickness. The waveguides are etched in a PlasmaTherm RIE system, with a flowrate of 20 sccm of CHF_3, a chamber pressure of 20 mT, and an RF power of 250 W. With these parameters, the DC bias voltage is observed to fall slowly over the duration of the etch. As long as the bias voltage remains between 100 and 300 V, the etch rate remains approximately 50-55 nm / min., and the SiO_2:Si selectivity is about 9-12. It is believed that the changing bias voltage is associated with the formation of a polymer coating on the walls of the etching chamber. If the bias voltage falls below 100 V, the etch profile is observed to degrade substantially, and the chamber should be cleaned before proceeding. Figure 3.4 illustrates a 7 μm deep waveguide etched into doped glass using the process described above.

After etching the waveguides, the samples are cleaned in a hot solution of 1:3 H_2O_2: H_2SO_4 (piranha) to remove the thin polymer layer deposited on the silicon. A standard RCA clean will also remove this polymer, but it slowly etches the SiO_2 and should therefore be avoided if possible. The silicon hard mask is removed in a hot (80 C) solution of tetramethyl ammonium hydroxide (20% by weight in water.) At elevated temperatures, TMAH will etch silicon but not SiO_2. Typically, the remaining hard mask can be removed in about 3-5 minutes, depending upon the temperature. While the silicon substrate is slightly etched during this process, the effect is not noticeable.

After the waveguides are formed, the wafers are returned to the vendor where the final glass cladding layer is deposited, again via flame hydrolysis deposition. Figure 3.5 is an optical micrograph showing the cross-section of the completed waveguide. In obtaining this micrograph, the chip was illuminated from the rear which causes the waveguides to appear brighter as a result of the guided light from back illumination.

3.1.4 SOI Waveguides

The fabrication of ridge waveguides in silicon-on-insulator (SOI) is much more straightforward than the fabrication of glass channel waveguides. The etching techniques for silicon

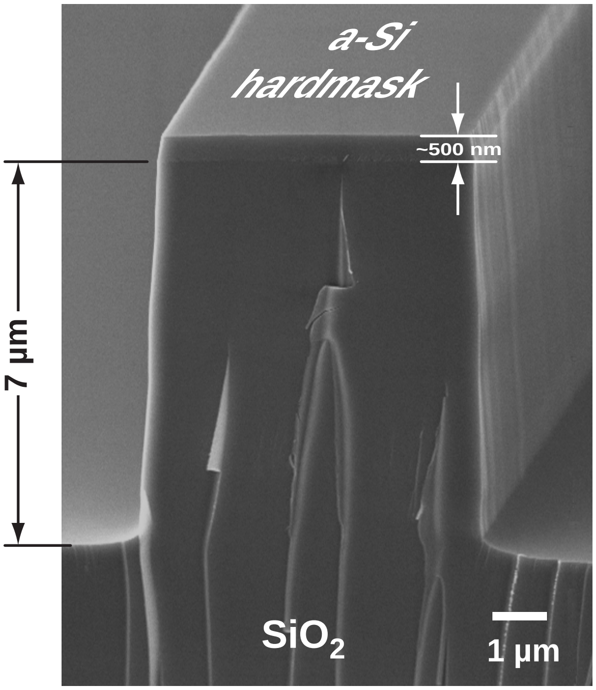

Figure 3.4: Scanning electron micrograph showing a doped glass waveguide etched to a depth of 7 µm. The silicon hard mask has not yet been removed. (figs/3/glass-wg-etch-sem.eps)

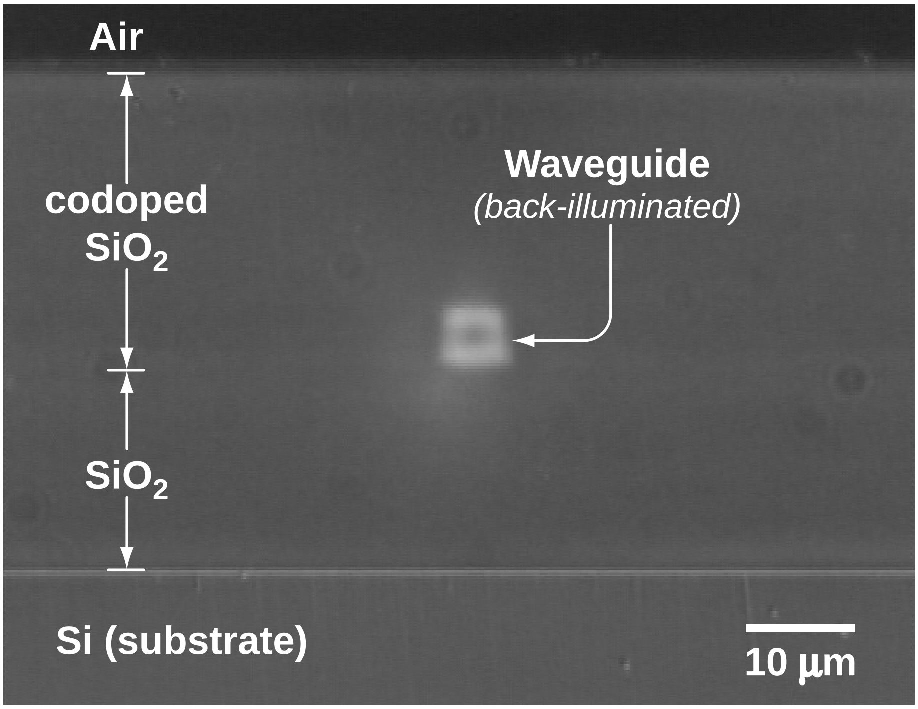

Figure 3.5: An optical micrograph showing the edge of a chip containing integrated glass channel waveguides. The chip is illuminated from the rear, which causes the waveguides to glow brightly when viewed in transmission. (figs/3/glass-wg-facet.eps)

are very well-established in the semiconductor industry, the typical ridge heights required are not prohibitively deep (0.5-1.5 μm), and no top cladding layer is needed.

There are several reported techniques for fabricating SOI ridge waveguides. Schmidtchen et al. start with SIMOX wafers with an epitaxially grown silicon layer and form the waveguides by anisotropic chemical etching using an SiO_2 mask [32]. Rickman et. al. likewise use epitaxially regrown SIMOX, but they instead form the waveguides with reactive-ion etching using CH_3Br with a photoresist mask[102]. Trinh at all start with bonded and etched-back SOI (BESOI) and form the waveguides using reactive-ion etching in SF_6/O_2 with an SiO_2 mask [107, 108]. Fischer et. al. likewise use BESOI, but describe an isotropic wet etching procedure which uses HNO_3, HF, and CH_3COOH. They also report using an evaporated Si_3N_4 antireflective coating on the end facets to reduce Fresnel reflection [33].

For the SOI waveguides, which are relatively new to our repertory, we have not yet designed an optical photomask with couplers and waveguides. Instead, we patterned a series of prototype waveguides using a VS2A electron-beam lithography system.

The raw materials used in this work were commercially available ELTRAN SOI wafers provided by Canon. They use a novel combination of epitaxy and selective back etching to achieve smooth, high-quality silicon on insulator with excellent uniformity. The ELTRAN fabrication sequence is illustrated in Fig. 3.6[109, 110]. The process begins by depositing porous silicon and, followed by an epitaxial silicon where which will ultimately become our waveguide core. The top surface of the silicon and is then oxidized to the desired depth, forming the lower cladding. This structure is then flip-bonded to another wafer and separated at the porous silicon layer. Separation is achieved either by grinding away the top wafer or by blasting the interface with a high-pressure water jet. The remaining porous silicon and is removed via a selective with chemical etch which stops abruptly on the epitaxial silicon. Finally, the structure is annealed in a hydrogen environment which has the effect of polishing the silicon to a smooth surface. This process has several advantages for integrated optical devices. First, in comparison to SIMOX (implanted) wafers, the oxide layer can be grown to almost any desired thickness of which allows for a high degree of optical isolation between the waveguide and the substrate. Second, the Si/SiO_2 interface is very sharply defined. Third, the film thicknesses (both Si and SiO_2) are controlled via epitaxy and selective wet etching, rather than by chemical-mechanical polishing, which leads to superior thickness uniformity. For the devices reported in this work, the oxide thickness was 1 μm, and silicon thickness was 3 μm.

Figure 3.7 outlines the process steps used to pattern SOI ridge waveguides. First, a 400 nm thick layer of polymethyl methacrylate (PMMA) is deposited by a spin-coating. A thin

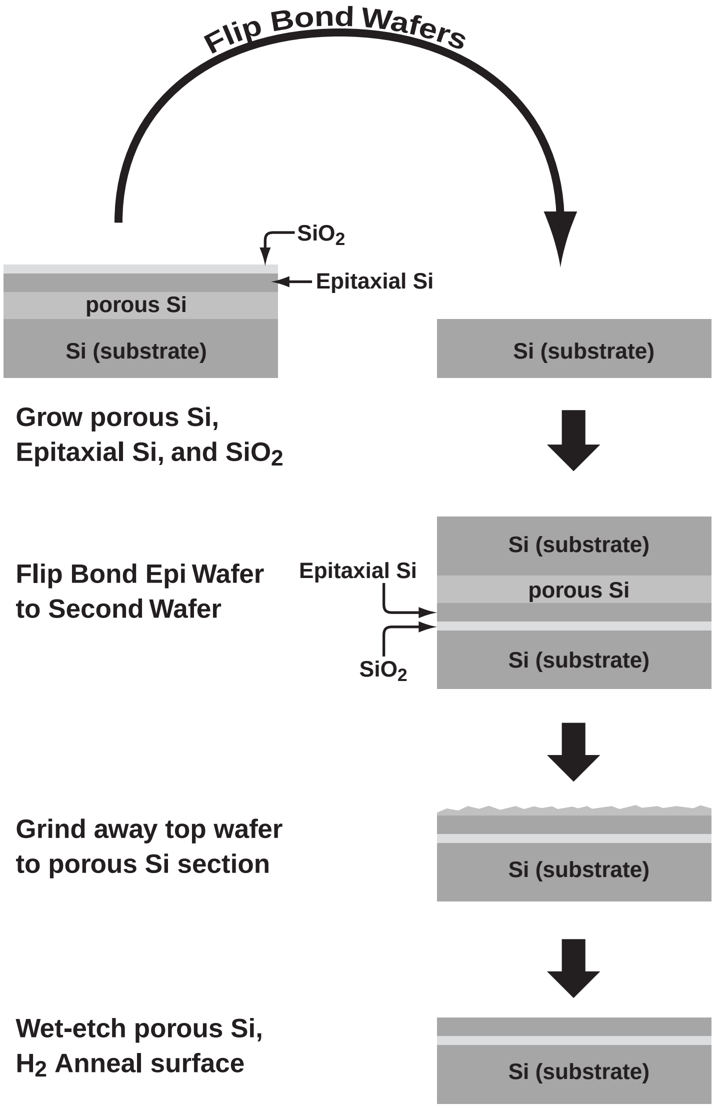

Figure 3.6: Diagram of process used by Canon to produce ELTRAN SOI wafers[109, 110]. (figs/3/eltran-process.eps)

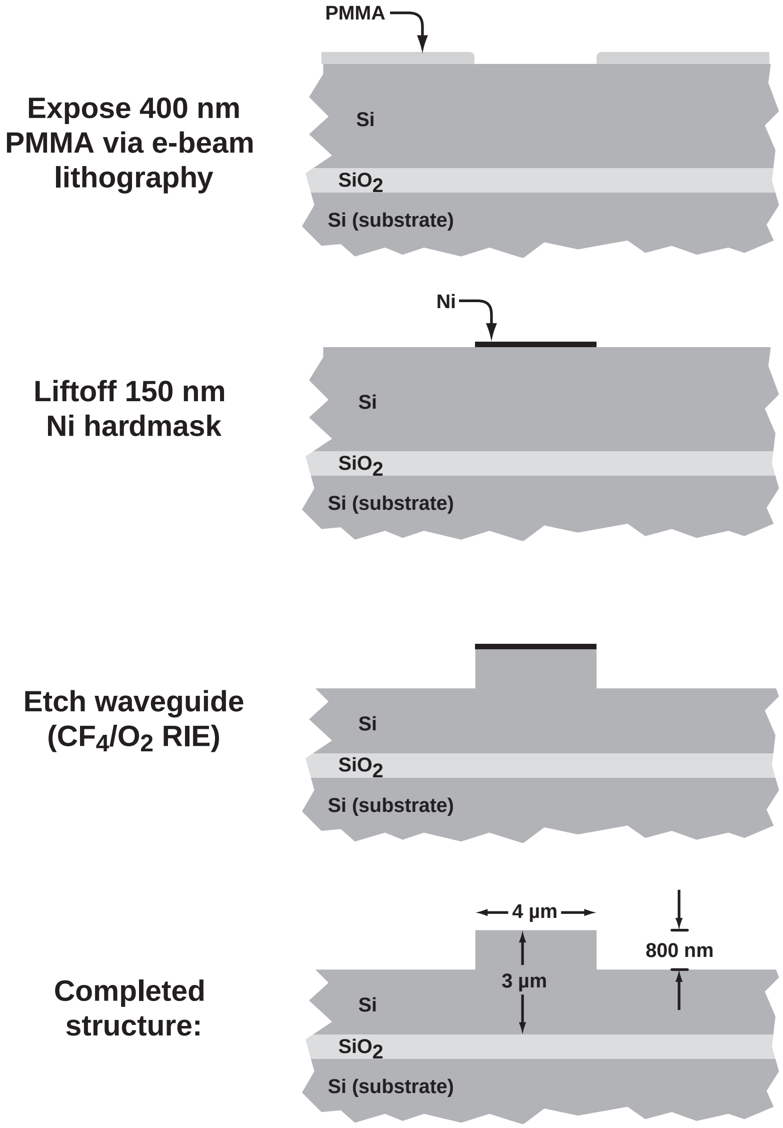

Figure 3.7: Outline of process used to form ridge waveguides in SOI. (figs/3/soi-wg-process.eps)

(5 nm) layer of chromium is evaporated over the surface to prevent charging during the e-beam exposure. The waveguides were written with a beam currents of 520 pA and an areal dose of xxx $\mu C / cm^2$. We biased the waveguides patterns down by about 200 nm to account for anticipated feature broadening due to the finite beam size.

After removing the chromium layer via wet etching, we developed samples in a solution of 60:40 IPA:MIBK for 60 seconds. A 150 nm thick layer of nickel was next deposited and lifted off in a heated solution of NMP. The final nickel feature width was 4 μm.

The waveguides were etched in a PlasmaTherm reactive-ion etching system, using the nickel as a hard mask. The etching parameters were 18 and 6 sccm of CF_4 and O_2, respectively, at a chamber pressure of 15 mT, an RF power of 250 W and a corresponding DC bias voltage of 260 V. The approximate etch rate achieved with these parameters is 21 nm / min. The waveguides could also the etched using Cl_2, as described earlier.

Following the waveguide etch the nickel is stripped in a hot solution of 1:3 H_2O_2:H_2SO_4 (piranha), completing the process. Figure 3.8 is an electron micrograph illustrating a completed SOI ridge waveguides. This sample was cleaved in cross-section for microscopy, but typically the facets are die-sawed and polished in order to more accurately define the facet orientation.

3.2 Lithographic Techniques for Bragg Gratings

The previous section described the methods for constructing integrated waveguides and couplers. We now turn to the more challenging task of fabricating Bragg gratings. As shown in Table 2.1, the periodicity of Bragg gratings ranges from 220 – 530 nm, which places them just beyond the reach of conventional photolithography. The purpose of this section is to summarize high-resolution lithography techniques which are suitable for patterning Bragg gratings.

3.2.1 Interference Lithography

Interference lithography is the cornerstone for most integrated Bragg gratings devices. The principle behind interference lithography is simple and appealing: two coherent plane waves which overlap in space form a standing wave which is recorded in a photoresist or other photosensitive material, as depicted in Fig. 3.9. It is relatively easy to compute the

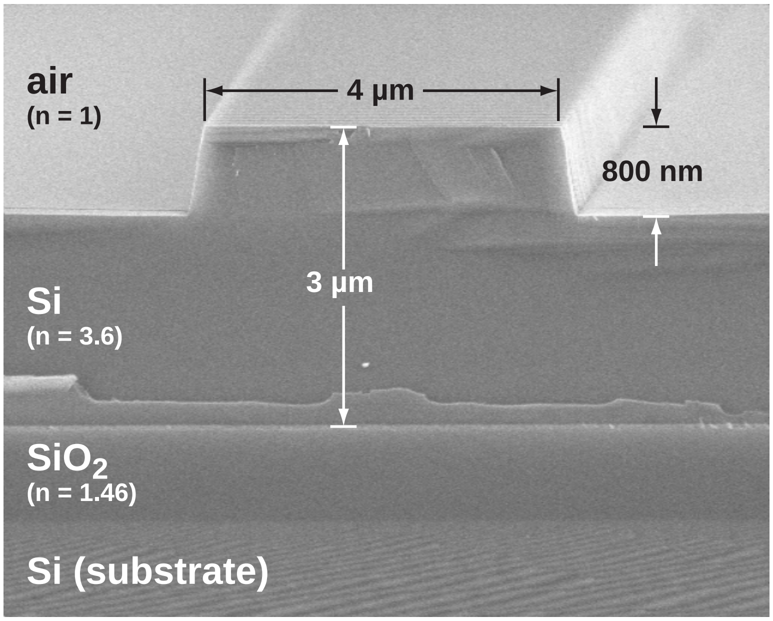

Figure 3.8: Scanning electron micrograph showing the cross-section of a ridge waveguide in silicon-on-insulator (SOI). (figs/3/soi-wg-sem.eps)

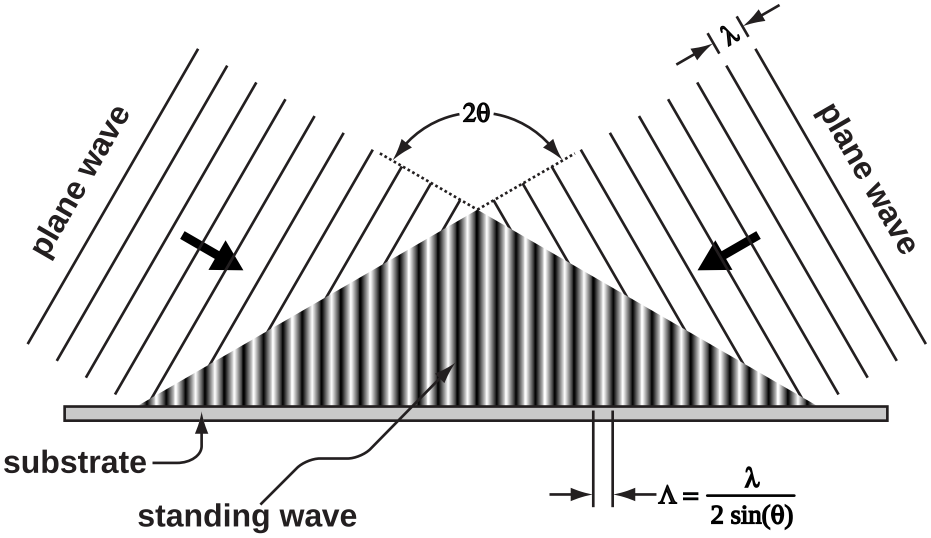

Figure 3.9: In interference lithography, two plane waves (or more accurately large-radius spherical waves) intersect forming a region of standing wave. This standing wave can be used to expose a photosensitive material on the substrate. The period of the resultant standing wave is related to the illuminating wavelength and the angle of intersection. (figs/3/interfering-plane-waves.eps)

periodicity of the standing wave:

$$\Lambda = \frac{\lambda_{\mathrm{IL}}}{2\sin\theta} \quad , \tag{3.2}$$

where λ_{IL} is the lithography wavelength of the illuminating source (not to be confused with the infrared free-space wavelength at which the device operates) and θ is the half-angle between the two interfering plane waves. In principle, the standing wave should form a perfectly periodic pattern which is *spatially coherent*. By spatial coherence we mean that by knowing the position of one local maximum in the standing wave, one could predict the positions of all the other peaks. The spatial coherence of the standing wave is of course related to the spatial and temporal coherence of the interfering beams, which is in turn related to the bandwidth and spatial character of the illuminating source.

Figure 3.10 illustrates the structure of one of the interference lithography systems used at MIT[111, 112]. The illumination source is an argon ion laser operating at a wavelength of 351 nm. The beam is split into two arms, focused through a pinhole spatial filter, and

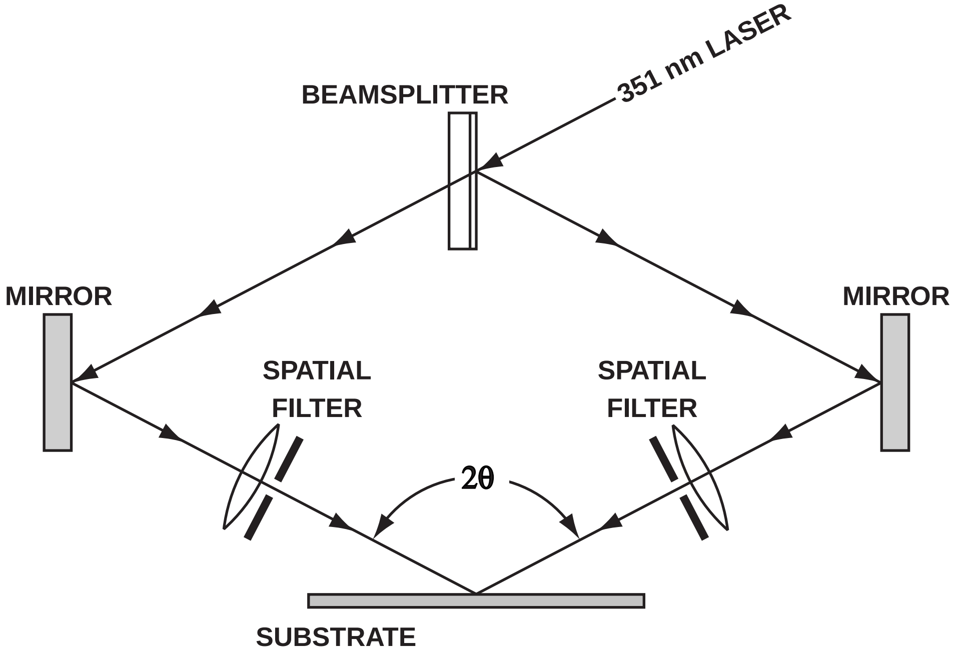

Figure 3.10: Diagram of the conventional interference lithography system used at MIT. Two beams from an Ar ion laser are interfered on the substrate to create a standing wave interference pattern. The lens and spatial filter assembly create a diverging beam which can be approximated as a spherical wave. Not depicted: there is a fringe-locking feedback mechanism which guarantees that the interference pattern remains stationary even over relatively long exposure times. (figs/3/il-conventional.eps)

interfered on the surface of the sample to be exposed. Similar systems have been used to expose Bragg gratings in photosensitized optical fiber [113, 114]. When exposing optical fiber, a series of lenses are typically used to expand and collimate the beam before it reaches the fiber. A cylindrical lens is also used to concentrate the illumination on the fiber core.

Figure 3.11 depicts a variant of this system, which is based upon a Lloyd’s mirror. In this system there is only one arm, but a mirror mounted perpendicular to the sample folds half of the beam back onto itself creating a region of interference. For the Lloyd’s mirror system, the illuminating source is a He:Cd laser with wavelength of 325 nm. Similar Lloyd’s mirror systems are used for exposing Bragg gratings in optical fiber [115, 116].

Under certain assumptions, the Lloyd’s mirror system can be thought of as functionally

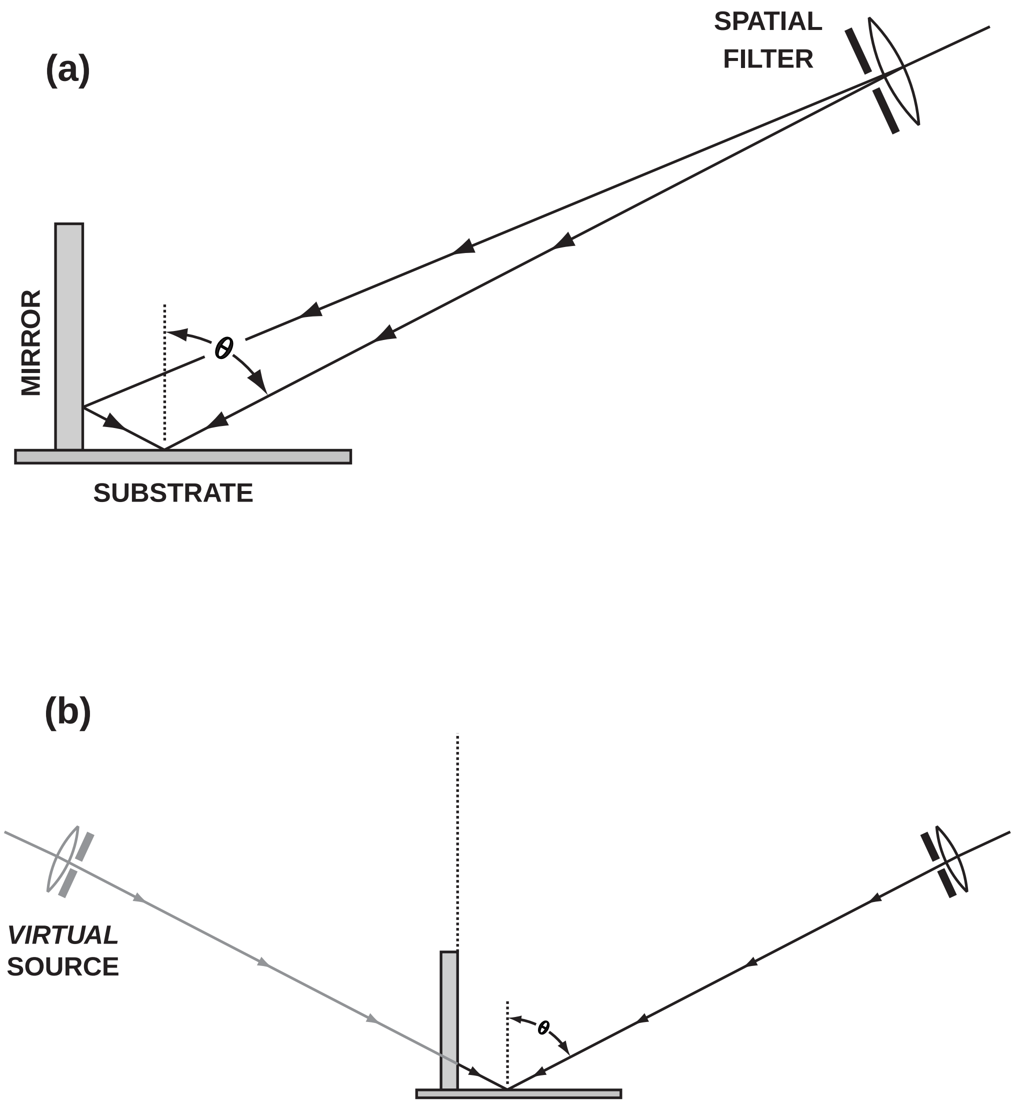

Figure 3.11: (a) A variant of the conventional interference lithography system, based on a Lloyd's mirror. The mirror is mounted perpendicular to the surface of the substrate, which causes half of the expanded beam to fold back on itself creating a standing wave at the substrate. (b) The Lloyd's mirror system can be thought of as equivalent to the conventional two-beam interference lithography system. The effect of the mirror can be modeled by placing a virtual point source opposite to the pinhole behind the mirror. (figs/3/il-lloyds-mirror.eps)

equivalent to the conventional two-beam interference lithography system, as shown in Fig. 3.11b. If we regard the spatial filter as an idealized point source and we further assume that the mirror is perfectly flat, then the effect of the mirror can be likened to having to point source placed symmetrically with respect to the mirror plane.

One important limitation faced by both of these systems is that the interfering waves are not plane waves; they are more accurately described as spherical waves [117]. Although interfering spherical waves produce a predictable (and hence spatially coherent) interference pattern, the resultant grating will not have a perfectly periodic structure. As we will discuss in Chapter 4, these distortions can easily affect the performance of a Bragg grating device. One way to fix this problem is to use lenses to collimate the illuminating beams to form plane waves (or more specifically, extremely wide beams with virtually flat wavefronts.) [113, 114] Another promising approach is to instead focus the beam to a small spot on the substrate while scanning the substrate on a precision controlled stage.

One of the potential problems with interference lithography (and with all other forms of photolithography) is that the light reflected by the wafer from the interfering beams can create another standing wave in the orthogonal direction. This can lead to undesirable ripples in the resist edges after development, and poor linewidth control. To minimize this effect, we typically use an antireflective coating beneath the resist to suppress the unwanted reflection from the substrate.

Figure 3.12 illustrates an example of a grating exposed in resist using interference lithography. The underlying antireflective coating layer is designed to reduce substrate reflection, and the thin oxide interlayer allows one to reliably transfer the grating pattern through the ARC using a simple two-step reactive-ion etching process. The complete process for transferring this pattern to the substrate will be discussed in Section 3.3.

3.2.2 X-ray Lithography

Although interference lithography is in many ways an ideal system for making Bragg gratings, it has some limitations. One disadvantage is that it is difficult to achieve alignment between the interference fringes and existing patterns on the substrate. Moreover, the presence of existing features on the substrate (e.g. metal layers under the resist) could interfere with the exposure. Perhaps most seriously, the distortion map of an interference lithography exposure depends sensitively upon how all of the optics in the system are aligned with respect to the surface. Consequently it is difficult to achieve exactly the same distortion map between two samples.

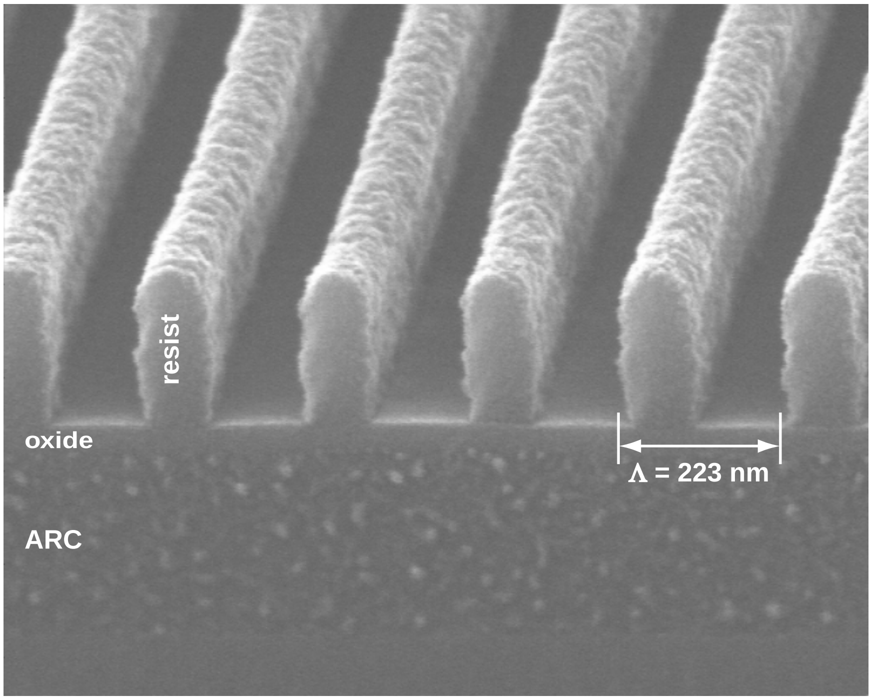

Figure 3.12: A 220 nm period grating exposed in resist using interference lithography. The underlying antireflective coating (ARC) is needed to suppress undesirable reflections from the substrate. (figs/3/grating-from-il.eps)

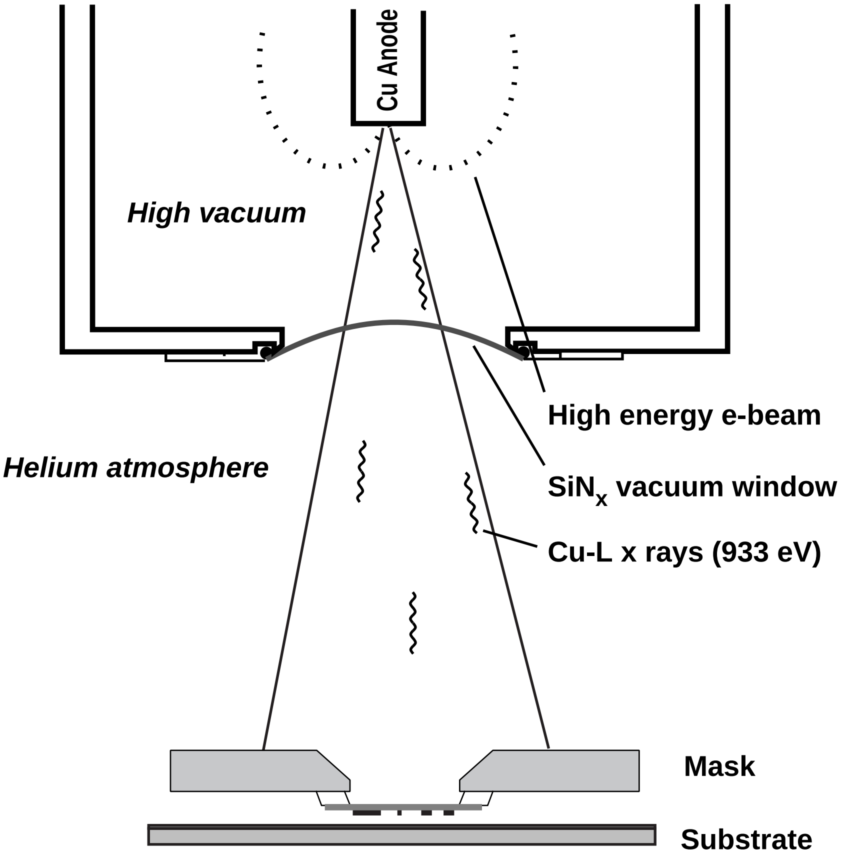

Figure 3.13: Diagram of an electron-bombardment x-ray source. (figs/3/x-ray-schematic.eps)

These problems can be mitigated by combining interference lithography with x-ray lithography. In this approach, interference lithography is used once to pattern an x-ray mask, which can be used repeatedly to transfer the patterns to substrates.

Figure 3.13 depicts the x-ray lithography system used at MIT [118]. The x-rays are generated by focusing a beam of electrons onto a copper target. The electrons are boiled off of a tungsten filament and deflected to a high-voltage copper anode (which is held at 8 kV with respect to the filament and chamber walls.) The impinging electrons generate x-rays characteristic of the copper-L transition, with energy of 933 eV. Also, there is a broadband background of bremstrahlung x-ray emission associated with the deceleration of the electrons. The emitted x-rays typically emerge from the vacuum system through a thin nitride

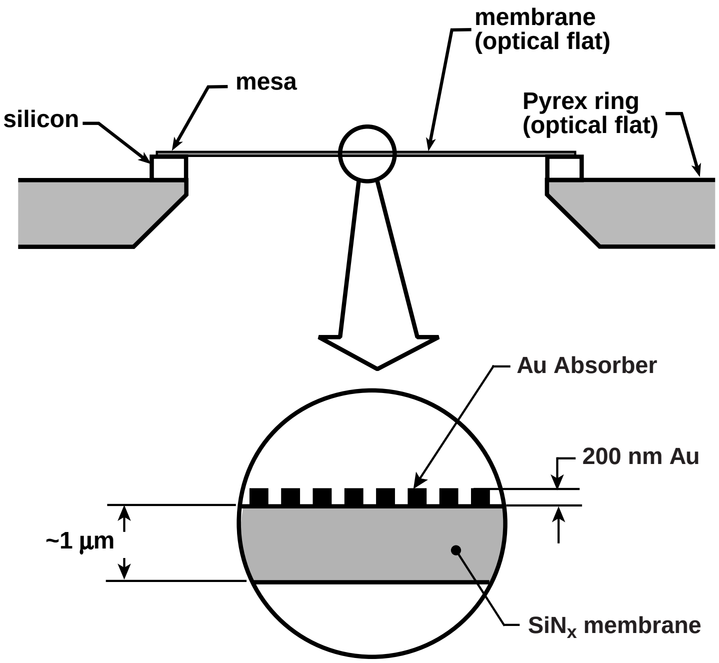

Figure 3.14: Diagram of an x-ray mask. A pattern of gold lines on a silicon nitride membrane defines a Bragg grating which can be transferred to the substrate via x-ray lithography. (figs/3/x-ray-mask-detail.eps)

membrane window and illuminate the mask and sample.

Figure 3.14 depicts the structure of an x-ray mask. A 1 µm thick membrane (which is 50% transparent to the 933 eV x-rays) supports a pattern of 200 nm thick gold lines which define the Bragg grating (or other pattern.)

In addition to its superior resolution, x-ray lithography benefits from the very different way that x-rays interact with materials. In contrast to optical and ultraviolet light, which is easily reflected by metal and dielectric surfaces, almost all materials are absorbing at x-ray wavelengths. For the x-ray wavelengths typically used in lithography, the indices of refraction for most materials lie slightly below 1. Consequently, x-rays are not reflected, except at grazing incidence. This greatly simplifies the lithographic processing, because antireflective layers are not needed, and the process of exposing gratings is not sensitive to the presence of underlying patterns. As an example, Fig. 3.15 shows a grating exposed in PMMA using x-ray lithography. Notice that in comparison with the structure shown in

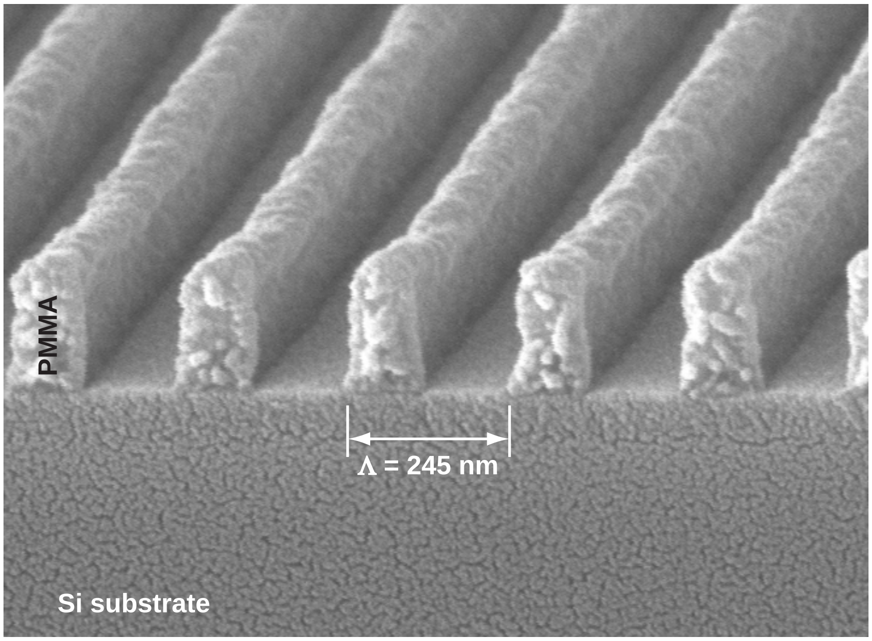

Figure 3.15: Scanning electron micrograph of a 244 nm period grating exposed in PMMA using x-ray lithography. (figs/3/grating-from-xray.eps)

Fig. 3.12, this exposure did not require the use of an antireflective coating.

3.2.3 Electron-Beam Lithography

One of the most flexible high-resolution lithography tools is scanning-electron-beam lithography, which we alluded to earlier when we described pattern generation for photomasks. In an electron-beam lithography system, a focused e-beam is traced over a surface to define almost any prescribed pattern. A combination of electromagnetic beam deflection and interferometrically controlled stage motion allows one to write patterns over a relatively large area.

One of the limitations of e-beam lithography is throughput. Because each pixel comprising the pattern must be exposed separately and sequentially, direct e-beam exposures are typically time-consuming. For this reason, e-beam lithography is primarily used for building prototype devices and for generating masks which can be used repeatedly in a higher throughput system such as optical photolithography or x-ray lithography.

A more serious limitation of e-beam lithography is the difficulty in achieving accurate pattern placement. In most e-beam systems, large area patterns are formed by stitching together a mosaic of small fields or stripes. The area within each field is accessed by deflecting the focused beam, while successive fields are written by moving the substrate. Accurate placement of patterns requires very well controlled stage motion in tandem with well-calibrated field deflection. Many factors, including for example thermal expansion, charging, beam current and focus drift, conspire to make this a difficult task. Often, there are small abrupt discontinuities, called stitching errors, at the boundary between adjacent fields. Sometimes the field deflection itself acquires distortion, especially at the edges and corners of the field.

Distortion and stitching errors are especially problematic for Bragg grating devices. Even relatively short Bragg grating devices typically span several electron-beam fields, and therefore it is impossible to avoid the problem of field stitching. Reflection from a Bragg grating is a coherent effect in which the small reflections from each period interfere constructively to create a high reflectivity over a narrow bandwidth. Stitching errors and distortion in the Bragg grating spoil the coherence of the device. Even relatively small random stitching errors can quickly degrade the optical performance of a device [119].

One approach to solving this problem is to somehow place a pattern on the substrate prior to writing which will act as a guide for pattern placement. An ongoing project in the NanoStructures Lab called spatial-phase-locked electron-beam lithography (SPLEBL) seeks to implement this technique using interference-lithography-generated gratings or grids on the substrate [120]. In SPLEBL, the sample to be exposed is first exposed in interference lithography, which generates a spatially-coherent reference pattern that can be viewed or sampled by the e-beam system. Then, as the e-beam system writes the desired patterns, the reference grating is periodically (or even continuously) viewed in order to assess the beam position. By “locking” the electron beam deflection to the reference pattern, stitching errors and field distortion can be significantly reduced [121, 122].

3.2.4 Phase Mask Interference Lithography

One of the most common techniques for patterning Bragg gratings is to use a variant of interference lithography based upon a diffractive phase mask[123, 124]. Figure 3.16 depicts the most common configuration for phase mask photolithography. Light impinging normally from above the phase mask is diffracted into three (or possibly more) orders.

The phase mask is carefully designed to suppress the 0th diffracted order. The $+1$ and -1 diffracted orders interfere, leading to a standing wave pattern. It is relatively easy to verify that the periodicity of the resulting standing wave is half the period of phase grating on the mask.

$$\Lambda = \frac{P}{2} \quad . \tag{3.3}$$

This result applies regardless of the wavelength of illumination (provided the illuminating wavelength is short enough to produce ±1 orders.) This means that, in contrast to the interference lithography systems discussed earlier, phase mask lithography does not require narrowband conherent illumination to produce a grating.

Furthermore, phase mask lithography systems are typically more stable than conventional interference lithography systems. In a conventional interference lithography system, the split beams traverse two relatively long paths before interfering on the substrate, whereas in a typical phase mask system the mask is usually held rigidly against the substrate (e.g. by vacuum contact.) Provided the phase mask and substrate are rigidly connected, the illuminating source can be scanned over the phase mask to expose the desired area. For these reasons, phase mask lithography is the method of choice for patterning fiber Bragg gratings.

One disadvantage of phase mask lithography is that the period of the exposed grating cannot be changed without replacing the phase mask. That is, the system is not tunable. Furthermore, phase mask lithography does not solve the problem of how to construct the phase mask itself. Typically, phase masks are either made using e-beam lithography (which is often prone to stitching errors) or one of the interference lithography techniques described in Section 3.2.1. Another limitation is that it is difficult to completely eliminate the 0th order.

3.2.5 Other Techniques for Producing Bragg Gratings

The methods described above are the most common ways of patterning Bragg gratings. However, there are many other nanolithography techniques which could be used, two of which we briefly mention here.

Deep UV contact photolithography is a technique in which a 1× flexible mask is pulled into intimate contact with a surface and exposed with UV illumination. If the mask and resist structures are designed appropriately, near field optical effects enable one to pattern

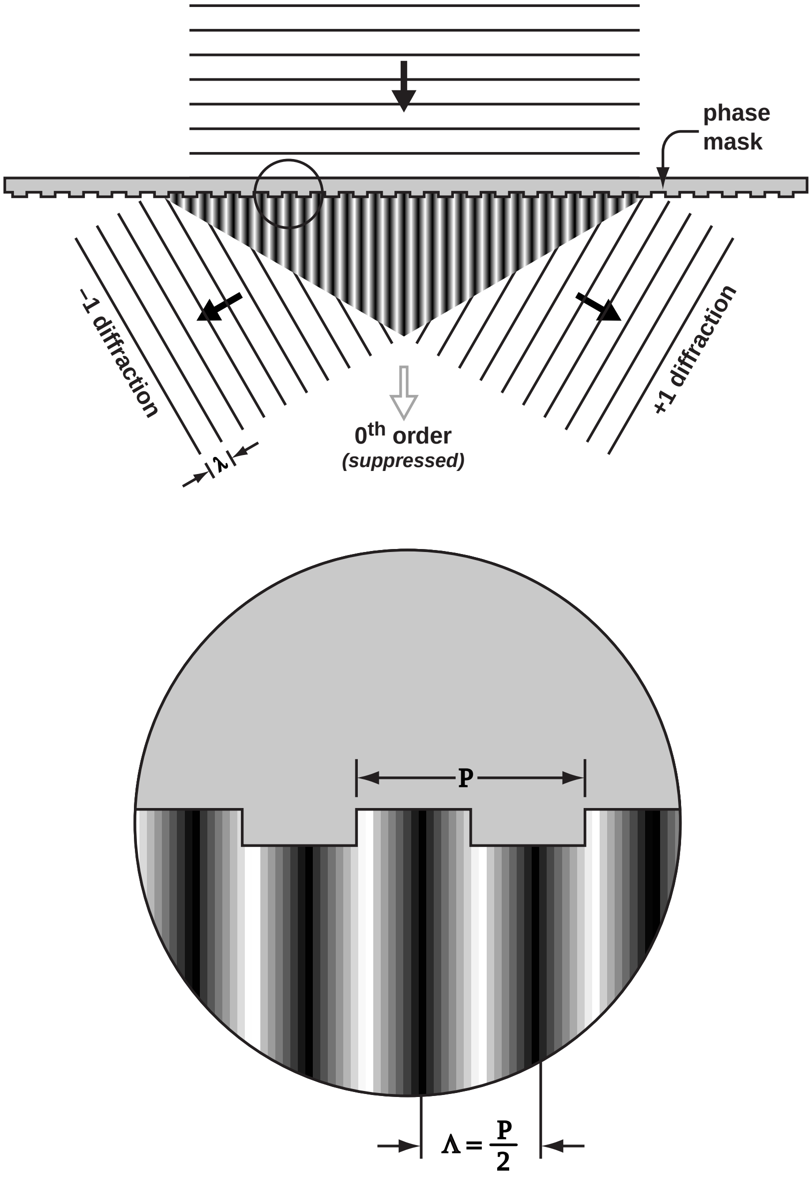

Figure 3.16: Diagram of a phase-mask interference lithography system (sometimes also called near-field holography.) Light incident on the phase mask is diffracted into at least three orders ($0, \pm 1$). The phase mask is designed to minimize (as much as possible) the 0th diffracted order. The ± 1 orders interfere, creating a standing wave whose period is half the period the phase mask grating. The standing wave can be used to expose a resist or other photosensitive material (not depicted.)

(figs/3/phase-mask-il.eps)

features which would normally lie beyond the resolution limit of the illuminating optics [125].

Nanoimprint lithography is another novel technique which can be used pattern Bragg gratings. In nanoimprint lithography, the grating patterns would be pressed or stamped into a polymer coating on the substrate, forming a relief structure [126]. There are variants of this method which use UV curable polymers and various combinations of flexible substrates and flexible stamps.

Both of these techniques require a 1× grating mask, which must be patterned with one of the other methods described in this section.

3.3 Multilevel Lithography: Waveguides with Gratings

In Section 3.1 we described the techniques of making integrated waveguides and in Section 3.2 we described the lithographic techniques for patterning Bragg gratings. In this final section, we focus on the problem of combining waveguides and gratings to build a functional Bragg grating filter.

3.3.1 Challenges to Building Integrated Bragg Gratings

One of the principal challenges to building the integrated Bragg gratings required for this work is that the devices call for fine-period Bragg gratings patterned over relatively tall waveguides. For the doped-glass waveguides, the Bragg grating period is 10× smaller than the desired waveguide height. For the silicon-on-insulator ridge waveguides, the grating period is 2-5× smaller then the ridge height. The nanolithography techniques described in Section 3.2 often rely upon a planar substrate. Figure 3.17 illustrates schematically the difficulty in patterning resist over a waveguide structure; conventional spin-coating techniques lead to extremely uneven resist coverage. Although it is sometimes possible to pattern gratings in such uneven resist [87], it becomes more difficult for taller waveguides.

Another challenge to building integrated Bragg gratings is ensuring that the Bragg grating is properly aligned to the underlying waveguide. Microelectronic devices often have dozens of lithographic layers, all of which must be aligned to one another. However, microelectronic devices are typically manufactured entirely with projection photolithography; alignment can be easily achieved by placing complementary alignment mark sets

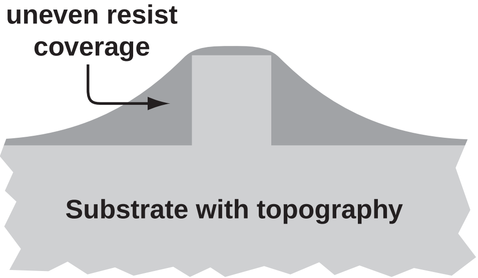

Figure 3.17: Schematic diagram illustrating the problem with performing lithography over topographic features. When the features are substantially taller than the desired thickness of the photoresist, achieving uniform resist coverage is problematic. Moreover, when patterning fine-period structures the resist cannot be made arbitrarily thick. (figs/3/uneven-resist.eps)

on each of the photomasks. The specialized high-resolution lithography techniques described in Section 3.2 are not all amenable to such simple alignment schemes. Moreover, the feature size of Bragg gratings are far too small to be seen with the alignment optics in an optical lithography system.

Figure 3.18a depicts a single waveguide with a misaligned grating. In this figure the grating lines are not oriented perpendicular to the waveguide. The first effect of misalignment is that the grating period seen by the devices increased by a trigonometric factor:

$$\Lambda = \frac{\Lambda_0}{\cos\theta} \simeq \Lambda_0(1 + \frac{1}{2}\theta^2) \quad , \tag{3.4}$$

where θ is the angular misalignment. This is a second-order effect which can slightly modify the Bragg frequency. If the misalignment is too severe, the grating may actually diffract light out of the waveguide core leading to radiation.

The requirements of angular alignment are more stringent for the Michelson interferometer device illustrated in Fig. 3.18b. Recall from the analysis of Section 2.4.4, in order for this device to completely transfer the reflected signal from one waveguide to the other, the optical path length must be matched in the opposing arms of the device. More specifically, the optical path length mismatch must be an integral number of half-wavelengths. Since the Bragg grating period is precisely 1/2 wavelength, this condition can be conveniently restated that the arms must be matched to within an integral number of Bragg

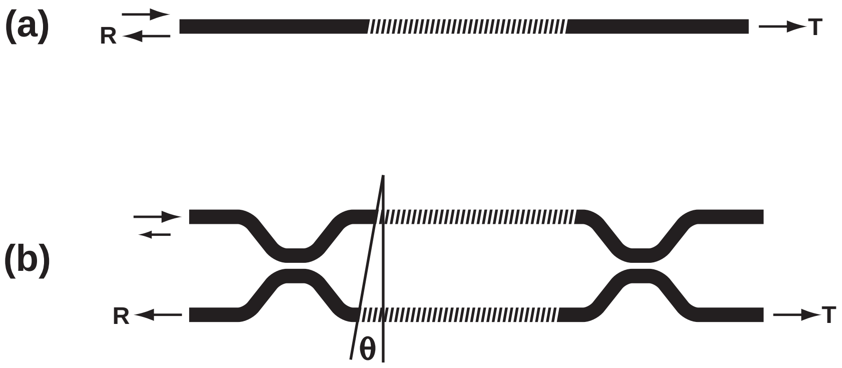

Figure 3.18: (a) Diagram of an integrated waveguide with a misaligned grating. The angle of the grating is not properly aligned perpendicular to the waveguide axis, which leads to a lengthening of the effective period seen by the waveguide. (b) Diagram of an integrated Michelsen interferometer with angularly misaligned gratings. In this case, the angular misalignment can cause the reflected light to be directed to the wrong output port. (figs/3/angular-misalignment.eps)

periods [2]. This will always be satisfied if the gratings are angularly aligned with respect to the waveguides. Even if one arm has a few more grating teeth then the other, the device should function correctly provided the k-vector of the grating is angularly aligned to the waveguide axis.

3.3.2 Prior Work on Integrated Bragg Gratings

One of the earliest places where integrated gratings have been used is in distributed feedback (DFB) and distributed Bragg reflector (DBR) semiconductor lasers [127, 128, 129]. Loosely speaking, in DFB and DBR lasers the reflective facets or mirrors which provide feedback into the laser cavity have been replaced with Bragg gratings. A conventional diode laser with a Fabry-Perot cavity typically lases over several longitudinal modes, whereas in a DFB/DBR laser the Bragg gratings have a wavelength-dependent reflectivity which favors one mode. The true benefit of a Bragg grating in semiconductor lasers is that it allows for single-mode oscillation [3]. The linewidth of the light emerging from a DFB/DBR laser is governed more by the peak reflectivity of the Bragg grating than its spectral width. Consequently, DFB lasers typically use a much higher grating strength κ than one would find in a passive filter. The Bragg gratings in lasers are often patterned over a large area of substrate and overgrown with guiding active regions, as in reference [129]. Sometimes the Bragg corrugations are chemically etched to form smooth sawtooth-shaped gratings which are better suited to overgrowth [127].

Researchers at Bell laboratories have fabricated Bragg gratings on SiO_2/Si_3N_4 rib waveguides [57, 131]. In this work, the waveguide was formed by etching a 8 nm ridge into the Si_3N_4 core prior to depositing the top cladding layer of SiO_2. Bragg gratings were then etched to a depth of about 60 nm into the top air-SiO_2 interface. Because of the almost planar geometry of the waveguide, they could use conventional interference lithography techniques to pattern the grating. Members of our group at MIT later developed a process for patterning Bragg gratings on similar rib waveguides [87]. In this work, the layer structure is identical but the waveguides were formed by a rib etched into the top SiO_2 layer. We described a self-alignment procedure that automatically confines the Bragg gratings to the top of the rib waveguides. Because of the relatively shallow waveguide features, the gratings could be patterned over the waveguides via x-ray lithography.

---

[2] This argument applies only when the waveguides comprising the device are mirror-images of one another.

[3] For uniform grating structures, there are actually two degenerate lasing modes placed symmetrically with respect to the Bragg stopband. This degeneracy can be lifted by placing a $\lambda/4$ phase shift in the Bragg grating, whereupon the lasing mode will be shifted to the center of the stopband [130].

Motivated by the tremendous success of fiber Bragg gratings, many groups have sought to construct similar photorefractive gratings in integrated glass channel waveguides. Maxwell et al. report one of the earliest demonstrations of a photo-induced Bragg grating in an integrated Ge:SiO_2 Channel waveguide[132]. Since then, several groups have reported similar results, and much research has been devoted to achieving increased photosensitivity in doped-SiO_2 and other types of glass waveguides [133, 55, 134, 135, 136]. In all of these cases, the process of forming the Bragg grating is essentially identical to that used to construct fiber Bragg gratings: typically phase mask interference lithography is used to expose a grating through the transparent cladding of the completed waveguide.

Kashyap et al. were the first to apply this technique to form matched gratings in opposite arms of an integrated Mach-Zehnder interferometer [137]. Again, the gratings were photo imprinted in the core using phase mask lithography. Because of the difficulty in aligning to gratings they used a post-exposure laser trimming procedure to balance the arms of the interferometer. Hibino et al. from NTT described a similar device which uses heating elements placed above the waveguides to balance or imbalance the path lines in the two arms [56, 138]. Researchers at Bell Labs were the first to construct an integrated Mach-Zehnder interferometer with identical Bragg gratings in the arms without any post-exposure trimming or adjustment [54, 139]. In part because of the difficulty in achieving angular alignment, they were unable to completely transfer power from one waveguide to the other. As pointed out in this work, another challenge to successfully implementing the integrated Mach-Zehnder interferometer is achieving perfectly matched gratings.

Our group in MIT has developed a dual-hardmask process for constructing integrated quarter-wave-shifted Bragg gratings on InP/InGaAsP channel waveguides [140]. In the following sections, we will describe a similar technique which has been developed for glass waveguides and silicon waveguides.

3.3.3 Bragg Gratings on Glass Waveguides

Figure 3.19 outlines the process we have developed to form Bragg gratings on glass waveguides. This sequence, which we call the dual-hard-mask process, effectively solve the problem of patterning gratings over relatively tall waveguide features. We will later describe each of these steps in detail, but first we will summarize the process in general terms in order to clearly identify the advantages of the process.

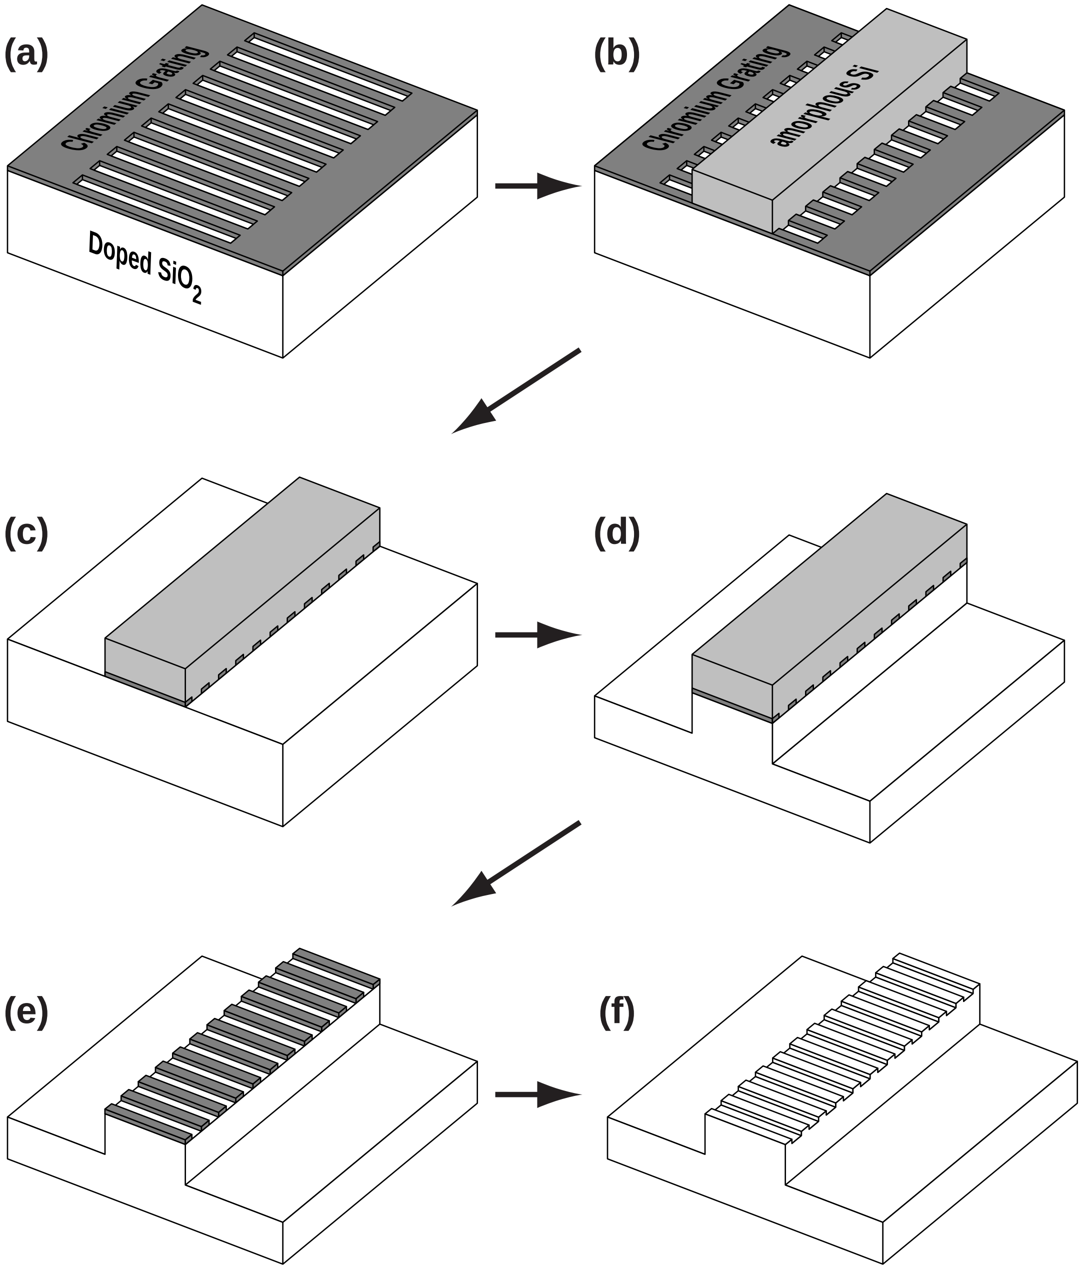

Figure 3.19: Overview of the process used to form Bragg gratings on integrated glass waveguides. The gratings are patterned first, but etched last. The sequence ensures that all lithography steps are performed over a surface free of significant topography. (figs/3/dual-hardmask-glass.eps)

Overview of Process

We begin with an optical substrate consisting of a lower cladding and core layer. Onto the core, we pattern the fine-period Bragg gratings using one of the lithography techniques described in Section 3.2. The gratings are not yet etched into the core, rather they are patterned in a thin metallic masking layer which sits on top of the core layer. The second masking layer defining the waveguides is then formed on top of the gratings, using the techniques described in Section 3.1.3. After the waveguide hard mask is patterned, the device is formed by a series of etching steps. First, the excess grating is removed from the regions outside the waveguides. Next, the waveguides are etched to the desired depth, and the waveguide hard mask is stripped, revealing the underlying grating mask. Finally, the gratings are etched and the grating mask removed. The device is completed by conformally depositing a top layer of cladding over the structure.

The first benefit of this process is that the most critical nanolithography step (that in which the Bragg gratings are formed) is performed over an essentially planar surface. Because the grating hard mask can be made relatively thin, the second level of lithography likewise has very little surface topography to contend with. A second benefit is that the gratings can be made wider than the waveguides, allowing for relaxed placement tolerance when patterning the waveguides: when the excess grating is stripped, the remaining grating lines are automatically confined to the tops of the waveguides.

Lithographic Alignment

So far, we have not specifically described the lithographic sequence used to pattern the Bragg gratings and waveguides. Three lithographic exposures are required to pattern the two hard masks. The three lithographic masks are depicted in Fig. 3.20. The first mask is an optical photomask which defines large rectangular regions where the gratings should be placed. The second mask is an x-ray mask covered with fine-period Bragg gratings, and the third mask is an optical photomask defining the waveguides. All of these layers must be aligned with respect to one another.

It is relatively straightforward to include alignment marks on the two optical masks in order to allow for alignment of the grating windows with the waveguides. The x-ray mask is fabricated using a combination of interference lithography and e-beam lithography, as shown in Fig. 3.21 [117]. Interference lithography is used to pattern a 535 nm period grating on an x-ray mask. However, before the grating pattern is developed a second optical

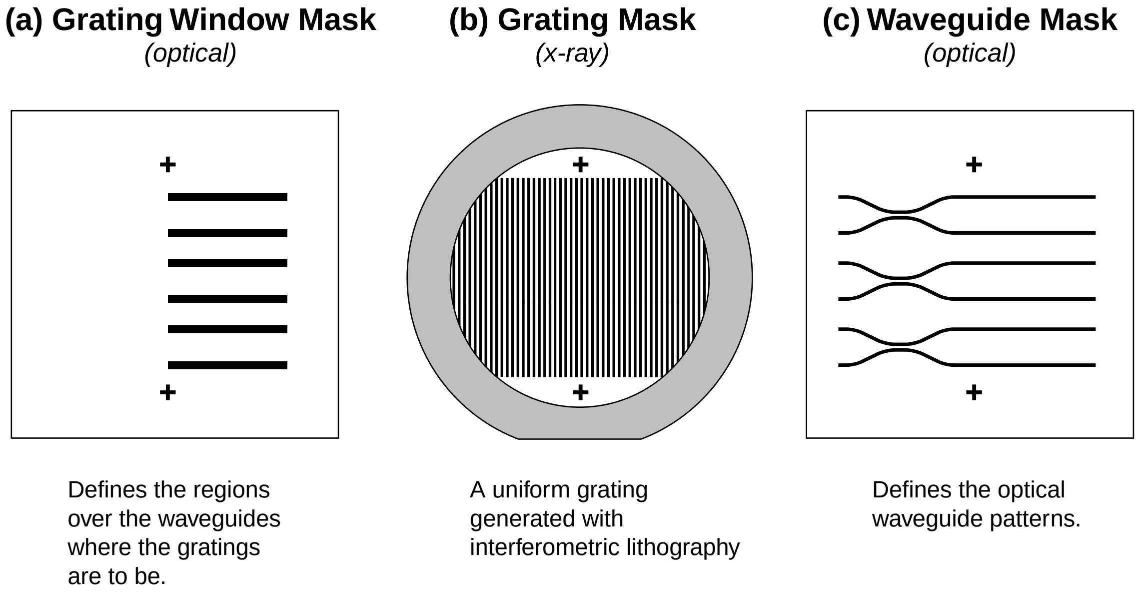

Figure 3.20: An overview of the three masks used to pattern the Bragg gratings and waveguides. The first mask defines rectangular regions where the gratings will reside. The second mask contains fine-period Bragg gratings, and the third mask contains the waveguide features. (figs/3/mask-overview.eps)

exposure is performed, exposing two large areas at the top and bottom of the mask. After the x-ray mask is developed and the pattern transferred through the antireflective coating, a gold absorber is electroplated. In this way, the areas which were exposed optically become opaque regions on the mask. The x-ray mask is then replicated (or daughtered) onto a second x-ray mask, a process that reverses the polarity of the pattern. After the daughter mask is developed and electroplated, it has two transparent regions at the top and bottom. It is in these regions that the alignment marks are to be written. The daughter mask is recoated with PMMA and loaded into an electron-beam lithography system. The e-beam system is then used to examine the grating pattern and write a pair of alignment marks at the top and bottom of the mask which are precisely aligned with the grating. It is possible to follow a single grating line over several centimeters on the mask in order to accurately achieve angular alignment between the alignment marks and the grating. The PMMA is then developed and electroplated. It should be pointed out that in order for the e-beam system to write alignment marks which are lined up with the grating, it is necessary to scan the e-beam over a portion of the gratings, which will cause the PMMA in these regions to be exposed. We have developed a process which allows us to protect these exposed regions during the subsequent electroplating step so that the gratings are not affected.

The procedure described above places a pair of conventional alignment marks on the x-ray mask which are angularly aligned to the grating. Later, when the waveguide mask is used it can be aligned with these marks, ensuring that the grating lines are perpendicular to the waveguide direction. Based upon the distance between the alignment marks and the accuracy with which we can align lithography layers, we estimate that the waveguides are angularly aligned to the gratings to within 0.006 deg.. If the arms of the interferometer are separated by 75 μm, this means that the optical path lengths in the opposite arms of the interferometer should be matched to within about 7 nm. This figure is more than five times better than the best reported results in the literature for lithographically aligned Mach-Zehnder interferometers[139]. Moreover, this figure could be further improved by increasing the separation between the alignment marks or by decreasing the separation between the waveguides.

Patterning the Two Hard Masks

The gratings are transferred to the substrate with a double-exposure lithography process. First, the substrate is coated with 250 nm of PMMA. The grating windows are then exposed and developed using UV contact photolithography. The grating window mask is transpar-

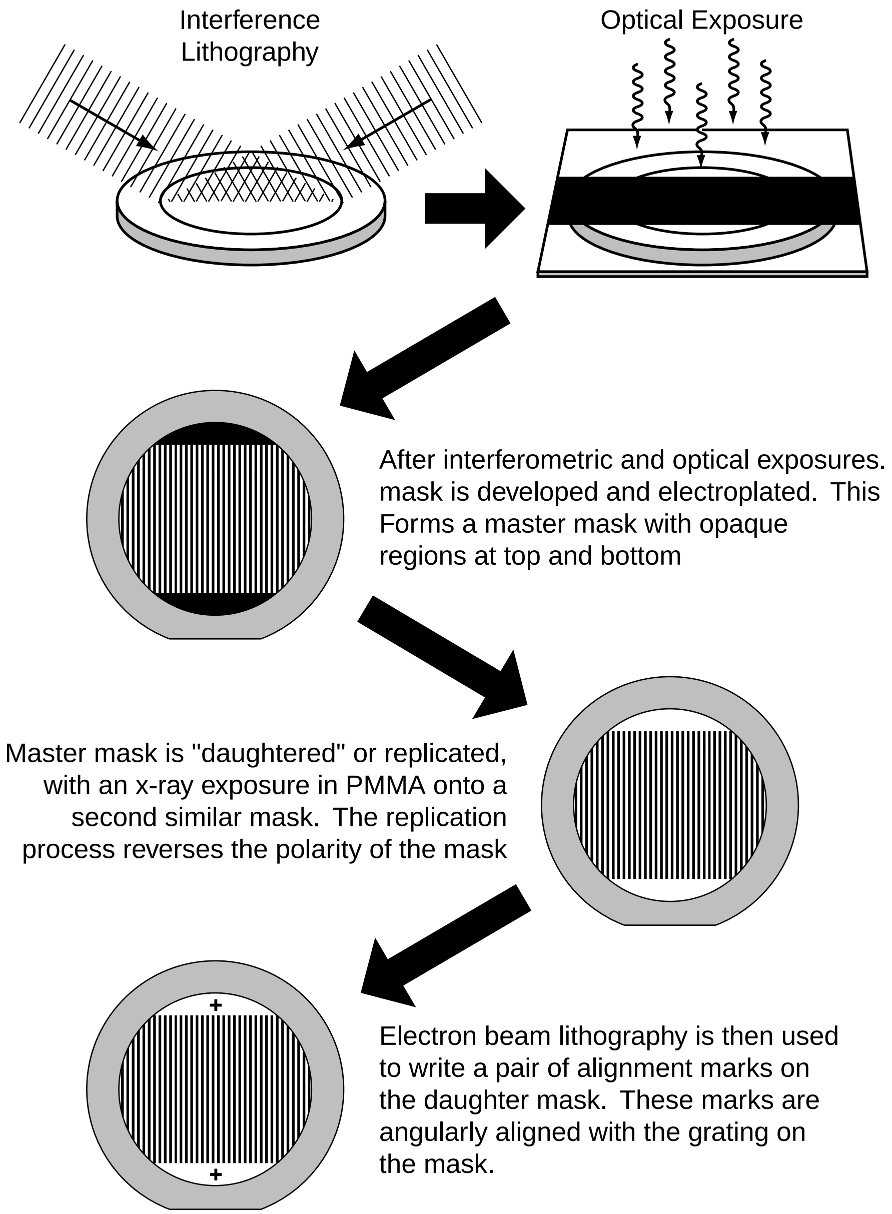

After interferometric and optical exposures. mask is developed and electroplated. This Forms a master mask with opaque regions at top and bottom

Master mask is "daughtered" or replicated, with an x-ray exposure in PMMA onto a second similar mask. The replication process reverses the polarity of the mask

Electron beam lithography is then used to write a pair of alignment marks on the daughter mask. These marks are angularly aligned with the grating on the mask.

Figure 3.21: Procedure for adding alignment marks to an interferometrically generated x-ray mask. This process combines the advantages of interference lithography with the flexibility of electron-beam lithography. The resulting mask has an interferometrically generated grating which is free of stitching errors, with precisely-placed alignment marks that are angularly aligned to the gratings.

(figs/3/add-marks-to-mask.eps)

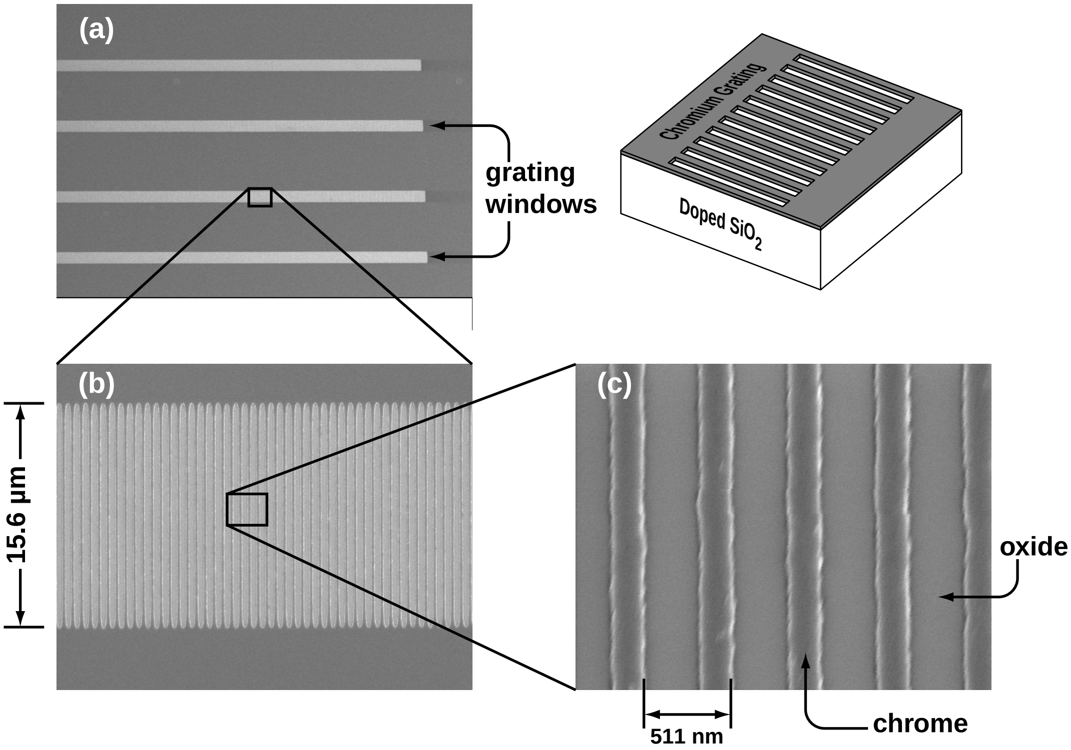

Figure 3.22: A series a scanning electron micrographs illustrating the results of the dual exposure process. We used two sequential exposures (one optical, one x-ray) to confine Bragg gratings to well-defined rectangular regions. (figs/3/grating-window-sem.eps)

ent outside of the grating region, so that after development the PMMA only remains in the regions designated for gratings. The x-ray mask is then used to expose gratings in the remaining PMMA. The grating window mask and x-ray mask are designed with a pair of complementary alignment marks which allow them to be aligned with respect to each other. After the double exposure and development, 40 nm of chromium is evaporated and lifted off in a heated NMP solvent. Figure 3.22 is a series of top-down micrographs showing the completed grating patterns. Since the resist must be developed twice, this process cannot be used with chemically amplified photoresists that require post-exposure baking or flood exposure prior to development. However, the process can be easily modified by simply performing two liftoff steps rather than just one.

After the gratings are formed, the waveguide mask is patterned using the techniques described in Section 3.1.3. First, an amorphous silicon hard mask layer is deposited via RF sputtering, and this layer is patterned via conventional contact photolithography and reactive-ion etching. A second pair of complementary alignment marks are included on the x-ray mask and the waveguide mask allowing for alignment between the waveguides and gratings. We have found that it is not necessary to protect the alignment marks on the substrate when the amorphous silicon layer is deposited. Because the sputter deposition process does not produce a completely conforming layer, it is usually possible to see the alignment marks as relief patterns in the top surface of the deposited silicon. Figure 3.23 is a micrographs showing the waveguide hard mask pattern on top of the grating hard mask.

Etching Sequence

Next, the excess grating outside of the waveguid is removed. Initially, we used a wet chemical etch to remove the chromium, but we found that this approach had problems with undercutting. In order to entirely remove the large areas of unpatterned chrome, the gratings had to be severely overetched, leading to undercutting. We solved this problem by instead removing the chrome by reactive ion etching. The chrome was etched in the same systems used to etch the silicon layer. We used a mixture of 34 sccm of Cl_2 and 6 sccm of O_2, with a chamber pressure of 20 mT, at an RF power of 100 W. With this process, the chromium layer can be removed in approximately 3 min.. Unfortunately, this dry etching process is not very efficient, and it partially transfers the grating pattern into the underlying oxide via physical sputtering. However, the undesired grating features are smoothed out during the subsequent deep waveguide etch.

The waveguides were next etched using RIE as described in Section 3.1.3. After etching the waveguides, the silicon mask is removed in hot TMAH. Figure 3.24 is a scanning electron micrograph showing an etched waveguide after the silicon hard mask has been removed. As shown in this figure, the chromium grating lines are not affected by the TMAH. Finally, the gratings are etched to the desired depth using the chrome as a mask, and the chrome is removed using a wet etch. Figure 3.25 depicts the completed structure with gratings over waveguides.

Overgrowth

One of the most challenging tasks to forming Bragg gratings on channel waveguides is the deposition of the top cladding layer. There are three requirements for this deposition. First, the top layer must conform to the underlying structure, completely filling the gaps

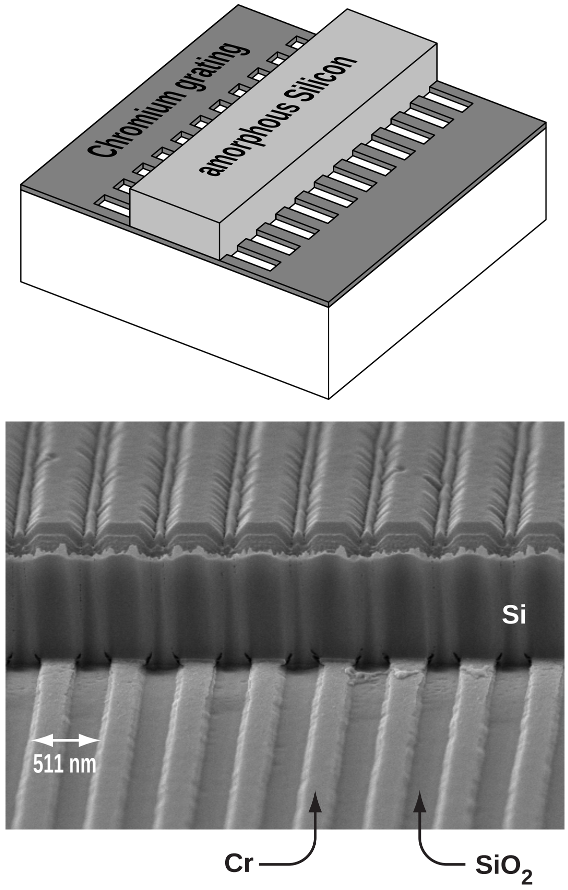

Figure 3.23: Scanning electron micrographs showing a patterned amorphous silicon hard mask on top of a thin chromium grating pattern. (figs/3/si-over-cr-sem.eps)

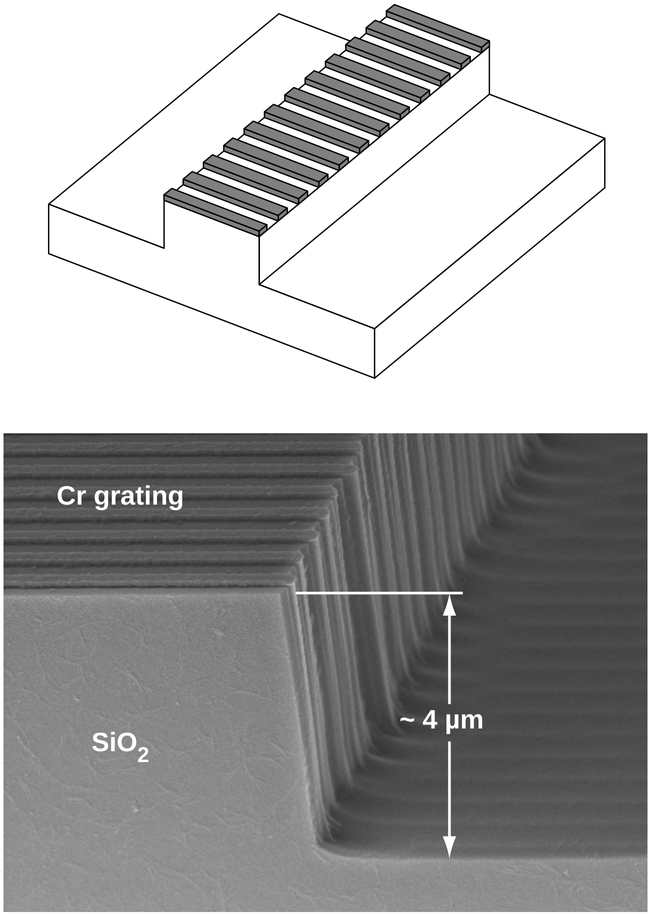

Figure 3.24: Scanning electron micrograph showing an etched waveguide after the amorphous silicon masking layer is removed in TMAH. Notice that the fine-period chromium lines remain intact under the silicon layer. (figs/3/cr-over-wg-sem.eps)

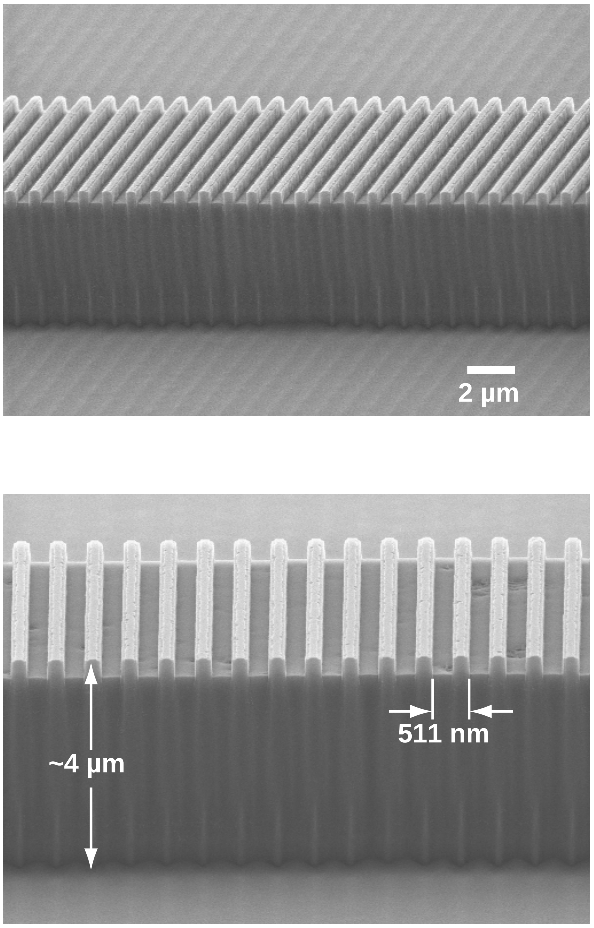

Figure 3.25: Micrographs depicting an almost-completed structure, immediately prior to overgrowth. The dual-hard-mask process described in this section allows us to pattern the fine-period gratings on top of relatively tall waveguides, as shown here. (figs/3/gr-on-glass-wg-sem.eps)

between grating teeth and the spaces between waveguides without forming voids (often called *keyholes* because of their distinctive shape.) Second, the shape of the underlying features must be preserved through the deposition and annealing of the top layer. Third, the index of refraction (at infrared wavelengths) of this top layer must be very well matched to that of the lower cladding. If the index of the top cladding layer is too large, even by a fraction of a percent, the waveguide will lose its ability to confine the light. Conversely, if the index is slightly too low, the waveguide can easily become multimode.

In our first attempt at overgrowth, we sent our samples to a commercial supplier (PIRI Inc.), who deposited the top layer by flame hydrolysis. From prior work with integrated waveguides and couplers, we knew that the deposition process used by the supplier could conformally cover large waveguide patterns and that the index of the top layer was well matched to the lower cladding (cf. Fig. 3.5.) In order to test the overgrowth process on fine features, we patterned the wafer with a large area of uniform grating. Specifically, we etched a 511 nm-period grating into the core layer to a depth of 250 nm (such that the teeth had an aspect ratio of approximately 1:1.) To allow for easier inspection, we did not include waveguides beneath the gratings. After overgrowth, we inspected the overgrown gratings using several techniques.

First, we visually inspected the wafer with an intense white light source, looking for the rainbow diffraction pattern usually produced by gratings of this period. This test failed to confirm the presence of a grating. We postulated that because of the small refractive index difference, the diffraction might be too difficult to see with the naked eye. Therefore, we next tried to detect a grating by carefully measuring the diffraction of a He:Ne laser. In this experiment we mounted an un-cladded grating next to the overgrown grating and measured the relative intensity of the diffracted beam with a lock-in detector. Unfortunately, we were unable to detect any diffracted beam from the overgrown sample. As a further test, we polished the samples to a smooth facet parallel to the grating lines and coupled light into the edge of the sample directly from a cleaved optical fiber. The opposite edge was observed on an infrared CCD camera while the wavelength of the input light was slowly tuned over the accessible 1.55 μm band. In this test, we were looking for a dark spot in transmission indicating reflection of the incident light at certain well-defined angles, but again we were unable to see any evidence of a Bragg grating. Finally, we cleaved and polished the sample in the direction perpendicular to the gratings and examined the cross-section using atomic force microscopy and scanning electron microscopy. A short HF etch was performed in an attempt to differentiate between the three layers. After this test, we concluded that the Bragg grating had been destroyed during the overgrowth or annealing process.

The flame hydrolysis process used by PIRI in the final cladding layer uses at least two dopants (Ti and Ge) to control the refractive index and glass transition temperature T_g of the layer. Our experiments indicate that although the deposition process does not destroy the waveguides, the fine-period structures are smoothed out. For most single-level waveguide devices, this smoothing of the underlying features is beneficial, because it removes any undesired roughness from the edges of the waveguide which would otherwise lead to scattering loss.

Another possible explanation of our observations is that during the deposition and anneal, some of the dopant species interdiffuse at the junction between the two layers. Depending upon the temperature and duration of the anneal, the diffusion depth could be significantly larger than the Bragg grating features. Finally, we point out that the high-temperature anneal used in flame hydrolysis deposition serves the dual purpose of consolidating the deposited glass “soot” and annealing the consolidated film. It is possible that this process requires a higher temperature and longer duration than would be required for films deposited by other means (e.g. sputter deposition or CVD.)

For our second attempt at achieving overgrowth on top of Bragg gratings, we sought the assistance of the silica waveguide fabrication team at Bell Labs. As discussed earlier in Section 3.1.1, the Bell Labs group uses CVD techniques rather than flame hydrolysis to deposit the waveguide layers. The bottom layer is a thick SiO_2 deposited by high-pressure steam oxidization (HiPOX), and the core layer is P-doped SiO_2 deposited to a thickness of 6.6 μm by CVD. The top cladding layer is co-doped with B and P to reduce the flow temperature and match the index of the bottom HiPOX layer. As before, we patterned a series of 535 nm period gratings into the core layer in order to test the final overgrowth step. Figure 3.26 shows the structure of the grating immediately before and after overgrowth. (The scanning electron micrograph of the overgrown sample was taken in cross-section after a brief stain-etch in diluted HF.) As shown in Fig. 3.26, there is some residual grating remaining after the final cladding layer is deposited, but the grating has been greatly diminished and smoothed. For this experiment, the top cladding layer was deposited in two steps, with an intermediate anneal, in order to avoid formation of voids in between the teeth. We are uncertain whether the distortion of the grating profile occurs during the first or second anneal. Additionally, because of transient uncertainties in the flow parameters used in the cladding deposition, the precise composition of the initial 100 nm of cladding is not well-known.

Changing the composition and deposition parameters for the top cladding layer is a potentially time-consuming process, because there are two separate dopant species which must be controlled, and there is a tight constraint on the refractive index which must be

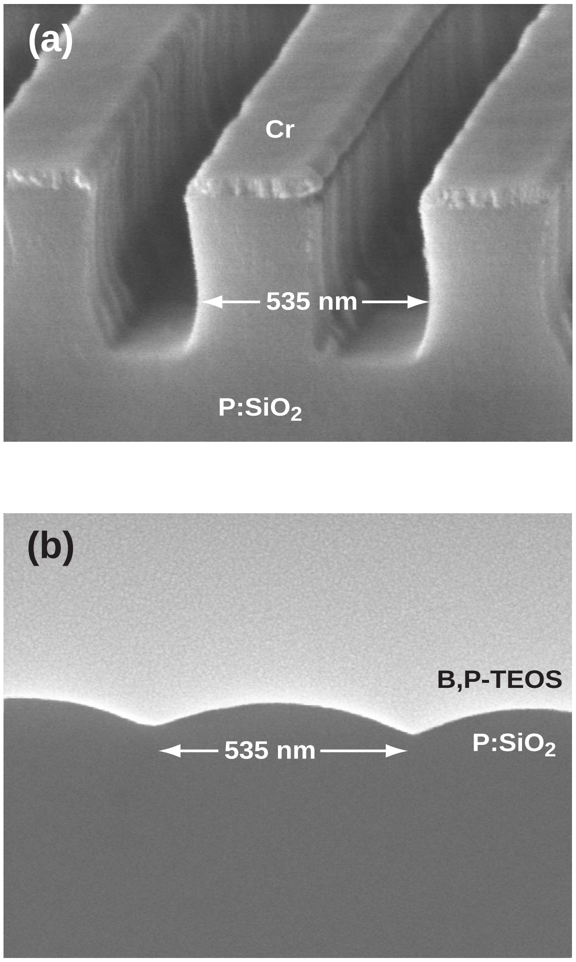

Figure 3.26: (a) Cross-section of test gratings prior to overgrowth. The gratings were etched to a depth of approximately 250 nm using the reactive-ion etching process described in Section 3.1.3. (b) Cross-section of gratings after overgrowth.

(figs/3/p-glass-overgrowth.eps)

maintained. Therefore, we initially focused our efforts on increasing the softening temperature of the core layer. For our next experiment, we changed the core composition from P-doped SiO_2 to Ge-doped SiO_2, because the glass transition temperature for Ge-doped glass is significantly higher than that of P-doped glass. The top cladding layer remains B-P-TEOS glass. The goal of this experiment is therefore to increase the T_g differential between the core and upper cladding layer. This experiment has not been completed at the time of writing.

Another possible solution to this problem is to pattern the grating in the lower cladding layer before the core is deposited. In this scheme, the core would conformally fill the grating during deposition. In principle, the pure SiO_2 should have a significantly higher softening point than all of the subsequent layers.

3.3.4 Bragg Gratings on Silicon Waveguides

The dual-hard-mask process depicted in Fig. 3.19 can also be adapted for use with silicon-on-insulator ridge waveguides. The masking layer materials must, of course, be modified to provide suitable masks for etching silicon. Figure 3.27 outlines the slightly modified procedure which we implemented for SOI ridge waveguides. As with the glass waveguide process, all lithography steps are performed over almost-planar surfaces.

Patterning the Two Hard Masks

The gratings were patterned in a thin (40 nm) layer of oxide, using interference lithography and reactive-ion etching. The samples were first coated with a uniform layer of 40 nm SiO_x, using e-beam evaporation. Next, a 165 nm-thick layer of ARC was spun on, to eliminate back-reflection. A second oxide layer, 30 nm thick, was e-beam evaporated on top of the ARC, and finally the imaging photoresist layer was spun on to a thickness of 200 nm. Figure 3.28 depicts the complete layer structure before exposure. The top layer of resist is where the grating pattern will be exposed. The ARC thickness is carefully chosen to minimize the back-reflection into the resist. The top oxide layer is used as a mask for transferring the grating to the ARC. The bottom oxide layer will be used as a mask for transferring the grating to the silicon.

Figure 3.29 illustrates how the patterns are transferred from the exposed resist into the oxide masking layer. First, the sample was exposed with interference lithography and developed. The grating is then transferred to the oxide interlayer using a CHF_3 plasma.

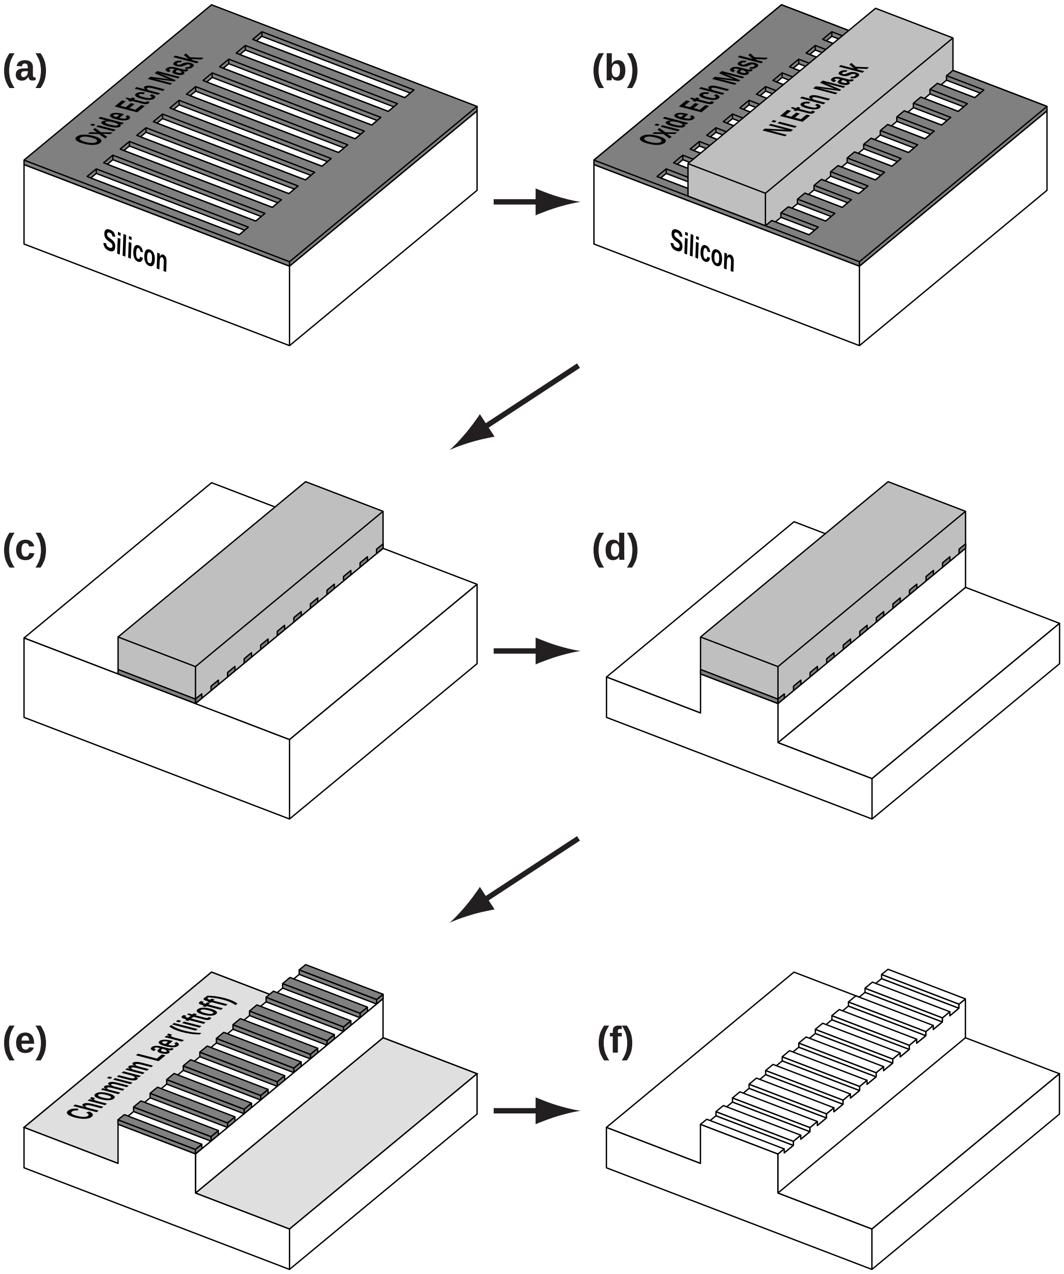

Figure 3.27: The dual hardmask process used to pattern fine-period Bragg gratings on top of silicon ridge waveguides. The process is similar to that described earlier for glass waveguides, with smaller feature sizes and different masking materials.
(figs/3/dual-hardmask-si.eps)

Figure 3.28: Layer structure used for patterning Bragg gratings in Si with interference lithography. The top photoresist layer is used to record the standing wave interference pattern. The thin interlayer is used as a mask for transferring the pattern through the ARC. The ARC is designed to minimize reflections from the substrate, and the bottom oxide layer will ultimately form the mask for etching the silicon. (figs/3/holo-layers-si.eps)

With the oxide interlayer as a mask, the pattern is then easily transferred through the ARC using an O_2 plasma. Finally, the bottom oxide layer is etched, again using reactive-ion etching in CHF_3, which stops at the underlying silicon. Generally, each step of this process is monitored immediately before processing a batch of samples, to determine the etch rates. Table 3.1 lists the etching processes used for each step, along with representative etch durations.

Process Step	**Etch Depth**	**Recipe**	**Duration**
Etch oxide interlayer	30 nm	20 sccm CHF_3 at 10 mT, 175 W	2:30
Etch ARC	165 nm	10 / 5 sccm O_2 / He at 10 mT, 250 W	5:00
Etch bottom oxide layer	40 nm	20 sccm CHF_3 at 10 mT, 175 W	2:30

Table 3.1: Etching processes used to transfer Bragg gratings from a interferometrically-patterned resist layer into the underlying oxide and ARC layers. Each of the steps was performed in a PlasmaTherm reactive-ion etching system. The etch times indicated here are representative figures; actual etch times must be adjusted based upon monitor runs.

Once the gratings were etched all the way through the oxide, all of the other layers are stripped in a hot solution of 1:3 H_2O_2:H_2SO_4. Because this step does not always fully remove the ARC, the samples were then cleaned in an O_2 plasma asher prior to further processing.

For the SOI devices, we didn't confine the gratings to windows using a double exposure as we did for the glass waveguides. Instead, we let the gratings fill a substrate, choosing to define the grating regions at a later point.

Section 3.1.4 completely describes the process used to pattern the waveguide hard mask, and therefore we will not repeat the details here. The waveguides were written directly into PMMA using e-beam lithography, and a nickel hard mask was then lifted off. The one critical modification is that the patterns written by e-beam had to be aligned with the existing gratings on the substrate. We initially attempted to use the e-beam to scan the existing grating prior to writing the waveguides, as described earlier. However, we found that it was too difficult to image a shallow grating in oxide through 200 nm of PMMA. One solution to this problem is to add a photolithography step which would selectively remove PMMA from a portion of the substrate, allowing the e-beam an unobstructed view of the grating. For our initial experiments with SOI waveguides, we adopted a simpler scheme.

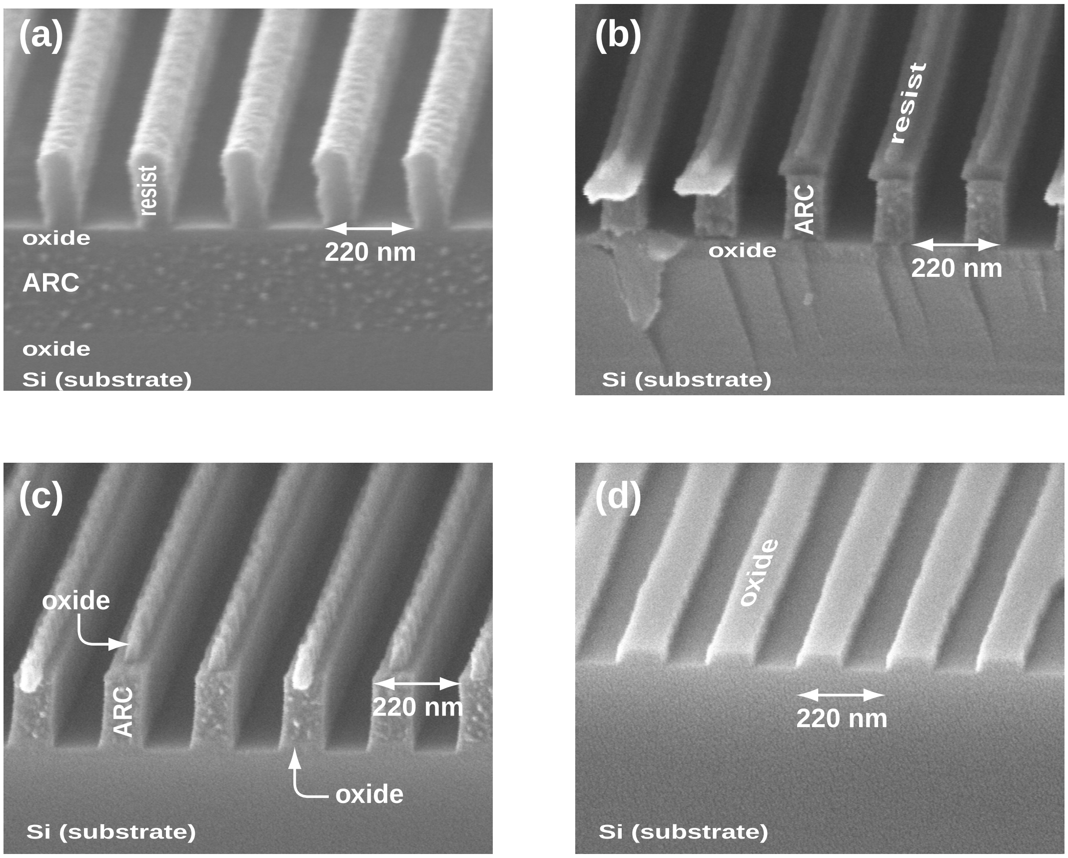

Figure 3.29: Trilayer process for patterning Bragg gratings. (a) The grating patterns exposed into the top layer of resist and developed. (b) After etching through the oxide interlayer, the patterns transferred through the ARC. (c) The pattern is then transferred into the final oxide layer using the ARC as a mask. (d) All of the imaging layers are removed leaving only the patterned oxide grating. (figs/3/trilayer-process.eps)

When exposing the gratings, we used the Lloyds mirror system, which leaves a clearly visible edge on the sample marking the point of contact between the mirror and the substrate. Furthermore, this edge should be well aligned with the direction of the grating lines. By marking two points on the edge with a scribe, we created a pair of approximately aligned features to which the e-beam could register. Although this process is not as accurate as the technique described earlier for glass devices, we calculated that it is sufficiently accurate to achieve our initial goal of building simple Bragg gratings on waveguides. In order to build integrated an Michelsen interferometer, a more accurate alignments process would be needed.

Figure 3.30 is a micrograph showing a nickel waveguide mask patterned via liftoff directly on top of a SiO_x grating pattern.

Etching Sequence

After the waveguide mask was patterned, the excess grating was removed in a short CHF_3 etch. Prior to etching the waveguides, the samples were cleaned in an oxygen plasma in order to remove any polymer which may have formed on the exposed silicon. Figure 3.31 depicts the nickel masking layer after the excess oxide has been removed. The waveguides were next etched to the desired depth, using the process described in Section 3.1.4. Before removing the nickel hard mask, we deposited a 40 nm thick protective layer of chromium, by e-beam evaporation. The nickel was then stripped in 1:3 H_2O_2:H_2SO_4, which lifts the chrome off the tops of the waveguides but leaves a thin layer outside of the waveguides to protect these areas from further etching. Figure 3.32 is a scanning electron micrograph showing an oxide grating on top of a silicon ridge.

Before etching the gratings, we performed a simple photolithography step to define the grating regions. We coated the sample with a 1.5 μm thick layer of photoresist. The grating windows were then exposed into this resist and developed, revealing the oxide grating underneath. We intentionally used a resist coating which was somewhat thicker than the ridge height, in order to avoid problems with incomplete resist coverage. Although this process does not yield a completely uniform resist thickness, the nonuniformity does not appreciably affect the exposure.

The gratings were etched into the top of the waveguide ridge using the oxide as a hard mask. The regions outside of the waveguides were protected by the chromium liftoff layer, and the areas outside of the grating windows are protected by the thick photoresist layer. To etch the gratings, we used a plasma with 40 sccm of Cl_2 at a chamber pressure of 20

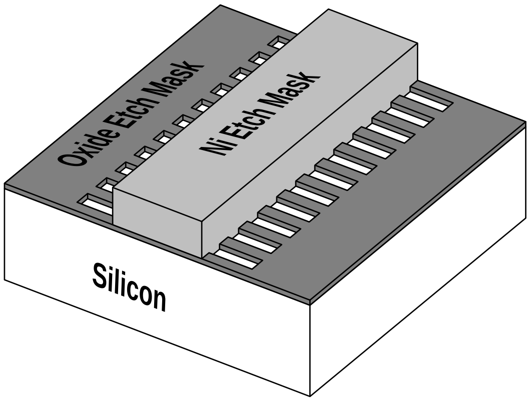

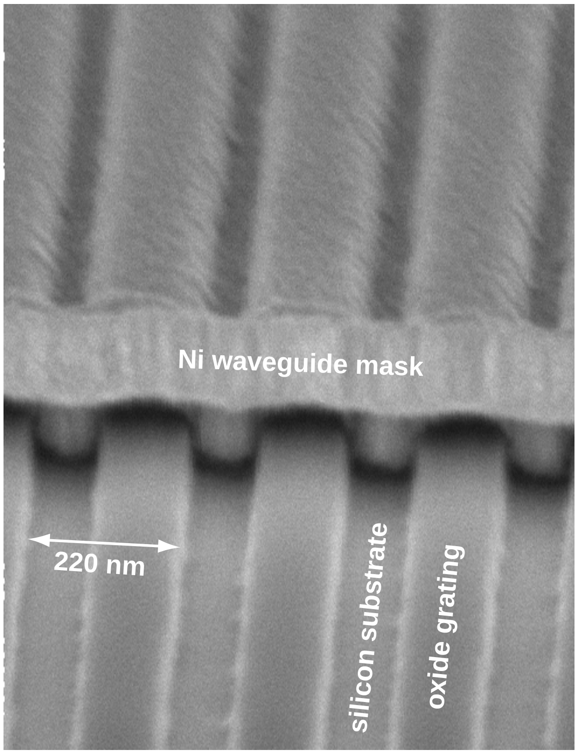

Figure 3.30: Scanning electron micrograph depicting a nickel waveguide pattern directly on top of an oxide grating. (In this picture, the nickel line is lifting slightly at the edge because of a nonoptimal resist profile during liftoff.) (figs/3/ni-on-oxide-grating.eps)

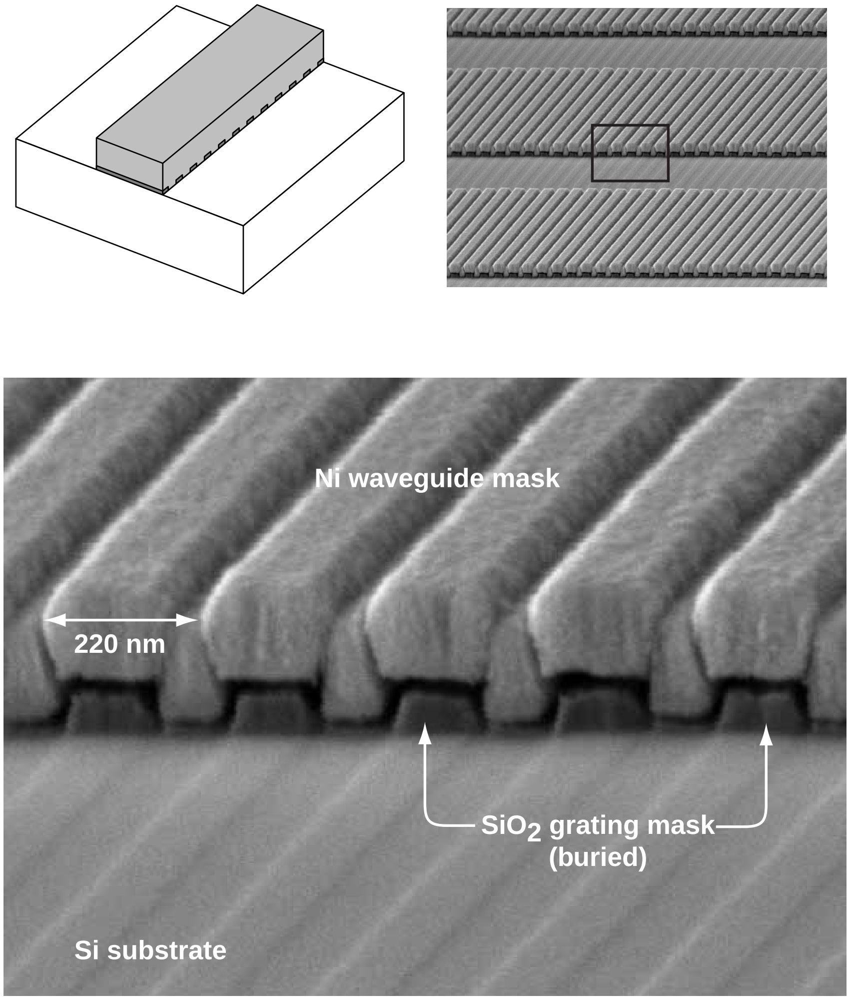

Figure 3.31: Scanning electron micrographs showing the waveguide mask with the grating buried underneath. The excess grating has been removed from the regions outside of the waveguide. There is a very shallow grating modulation in the silicon which was formed when the oxide mask is etched. (figs/3/ni-after-ox-removal.eps)

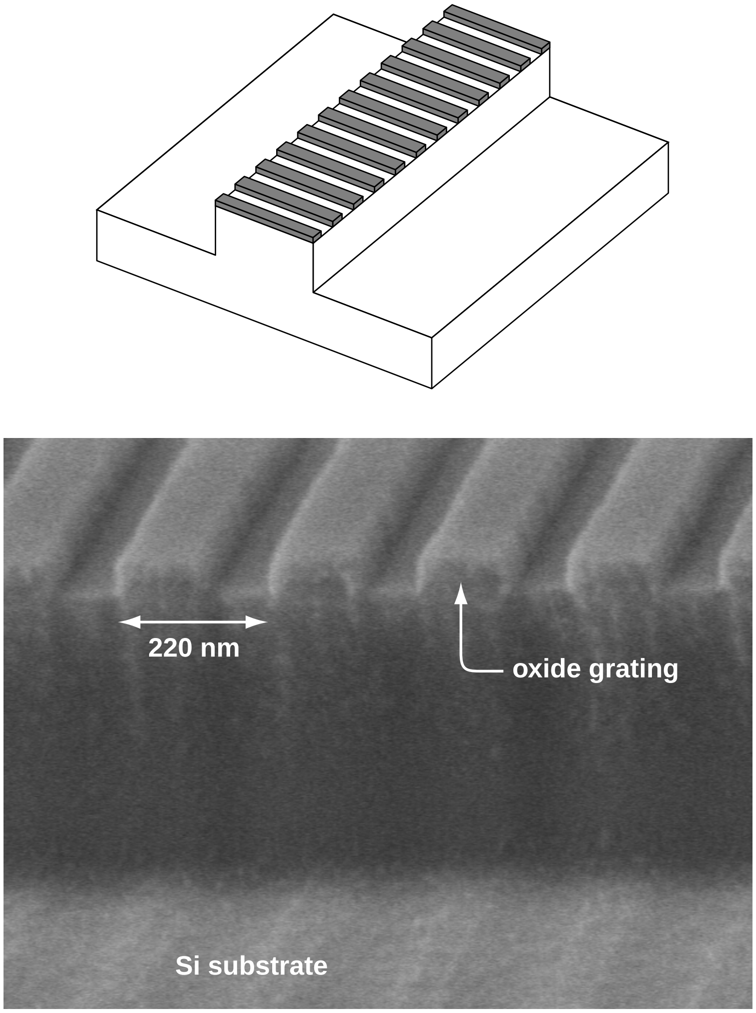

Figure 3.32: Scanning electron micrograph showing an oxide grating pattern on top of a silicon ridge waveguide. This micrograph was taken before we had added the chrome liftoff step to the procedure. (figs/3/ox-on-si-wg.eps)

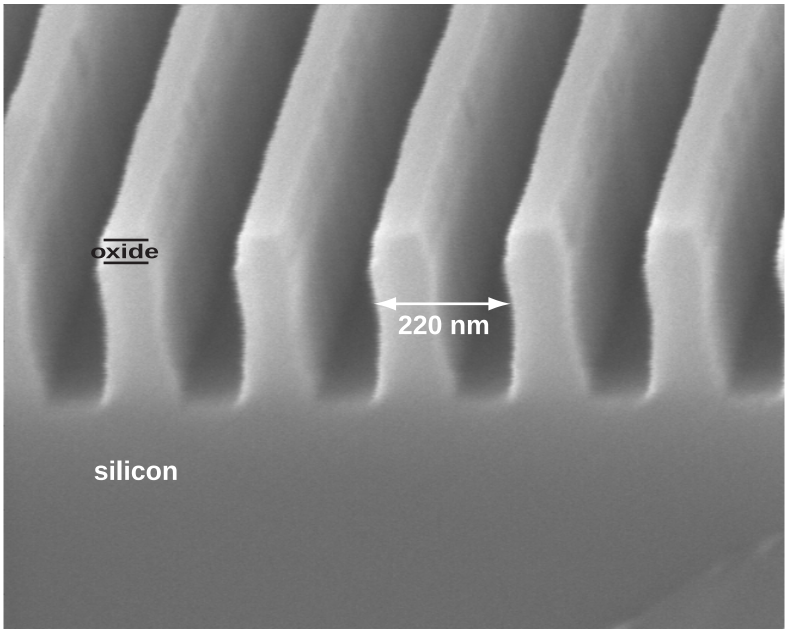

Figure 3.33: Deep grating etched into silicon using chlorine reactive-ion etching with an oxide mask. Although there is some rounding of the oxide features, the mask remains mostly intact. (figs/3/deep-grating-in-si.eps)

mT and a power of 75 W. Figure 3.33 shows the profile of a deep grating etched with these parameters, illustrating the excellent selectivity to SiO_2. [4]

Finally, the oxide grating mask is removed using reactive-ion etching in CHF_3. The sample is once again cleaned in an oxygen plasma, and lastly the protective chrome layer is removed in a wet chemical etch. Figure 3.34 shows a completed structure with gratings patterned on top of the waveguides.

---

[4] This high selectivity can often lead to “grass” formation in the large exposed areas in between waveguides. The grass is formed by small amounts of SiO_2 which are sputtered away from the mask and redeposited in the open areas. The protective chrome layer described earlier prevents this grass from forming.

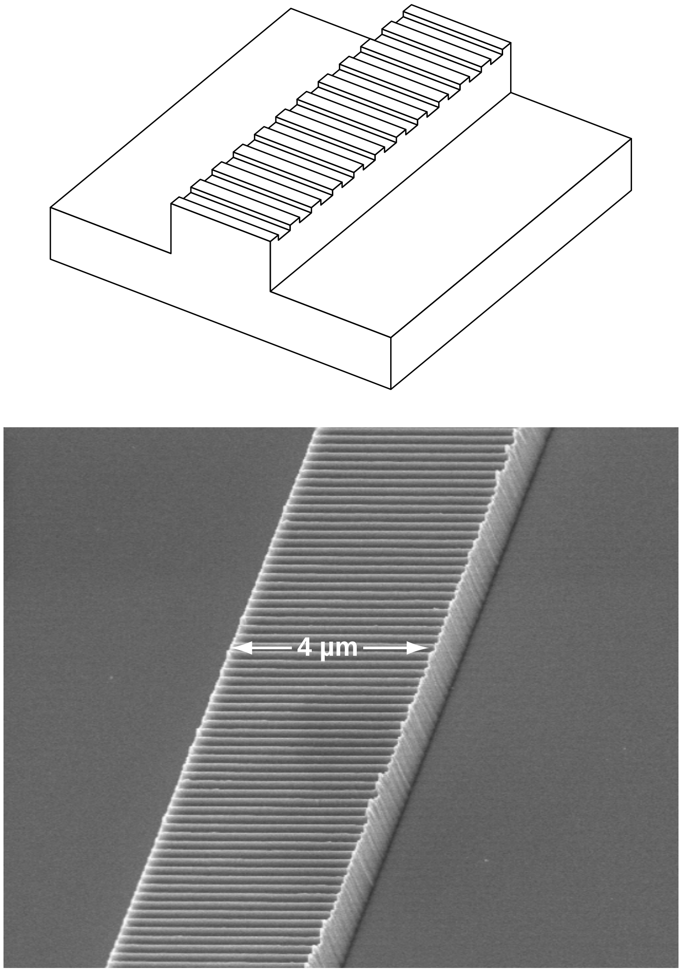

Figure 3.34: Scanning electron micrograph of a silicon ridge waveguide with integrated Bragg grating. (figs/3/si-ridge-with-gr.eps)

3.4 Summary

In this chapter, we described the fabrication techniques for building integrated waveguides and Bragg gratings. Although we have specifically concentrated on two types of waveguides, doped-glass channel waveguides and silicon-on-insulator ridge waveguides, the processes describe here can be easily adapted to other systems.

The fabrication of integrated Bragg gratings involves two lithographic steps: one which defines the relatively coarse waveguides features, and one which defines the fine-period gratings features. This chapter began by describing the techniques for forming waveguides and couplers. While the waveguides can be patterned using conventional optical photolithography, high-resolution nanolithography must be used to print the Bragg gratings. Therefore, we next described several high-resolution lithographic techniques which are suitable for patterning Bragg gratings. Finally, we addressed the problem of combining these two lithographic techniques to form gratings over waveguides.

Chapter 4

Measurement

Thus far, we have described how to design and construct integrated waveguides, couplers and Bragg gratings. In this chapter we discuss the measurement of the devices we have built. The goal of this segment is to (1) provide a practical description of the measurement techniques used and (2) present the results of our work. First, we will focus on integrated glass waveguides and couplers, describing the results of our work on an improved directional coupler which has a power-splitting ratio that is insensitive to wavelength, polarization, and fabrication parameters. Next, we will describe our measurements of integrated Bragg gratings on silicon-on-insulator ridge waveguides.

4.1 Glass Waveguide Measurements

As mentioned in Section 3.3.3, we have not yet perfected the deposition of glass cladding over fine-period Bragg gratings. Nevertheless, glass waveguides provide an ideal platform for investigating directional couplers. We describe here the measurement techniques and performance of a set of directional couplers made with doped-glass waveguides.

4.1.1 Summary of Device Design

We designed and fabricated a set of directional couplers, using the analytical and lithographic techniques described earlier. Figure 4.1 depicts schematically the structure of the directional couplers, with all of the relevant dimensions labeled.

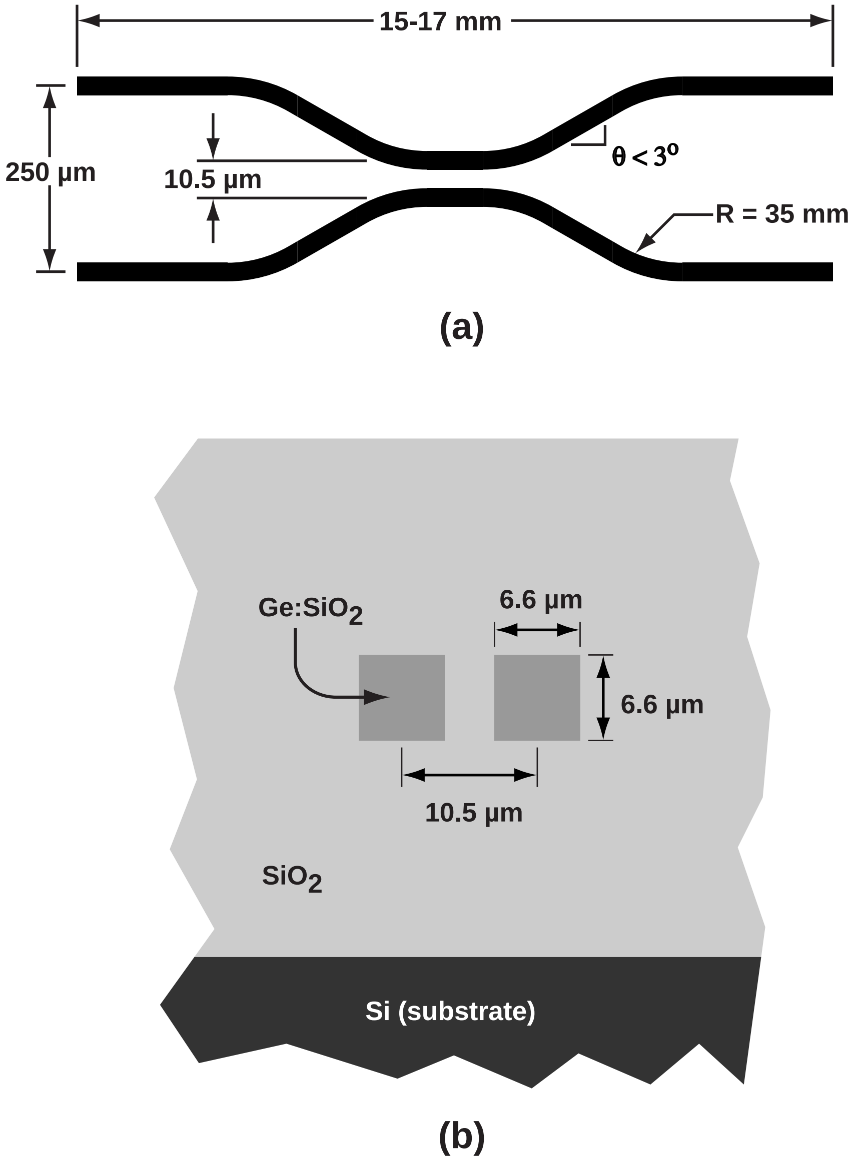

Figure 4.1: (a) Top-down diagram (not to scale) of an integrated directional coupler, with all of the dimensions labeled. (b) Cross-sectional view of the waveguides at the point of minimum separation. The couplers were fabricated in integrated germanium-doped glass channel waveguides, with a nominal refractive index contrast of 0.3%. (figs/4/coupler-geometry.eps)

The coupler is comprised to two identical square-cross-section channel waveguides, with a nominal width and height of 6.6 µm. These dimensions were chosen to yield single mode operation and efficient coupling to a standard optical fiber. The materials used were Ge-doped SiO_2 waveguides with an index contrast of 0.3%, as specified by the supplier. Based upon the calculated mode profile of this waveguide and that of a standard optical fiber, we estimate that the transition loss due to mode mismatch at the fiber-waveguide interface should be less than 0.15 dB.

The radius of curvature was chosen to the 35 mm for all of the devices constructed, and maximum bend angle was approximately 3 degrees. The bending loss for these waveguides was modeled using a conformal transformation method described elsewhere [83, 39]. This bending radius is predicted to yield negligible loss ($<$ 0.5 dB) for the structures considered. Other researchers who have worked with similar waveguides report using similar or slightly larger bending radii [30, 96].

The minimum waveguide separation (center-to-center) was fixed at 10.5 µm, which leads to an inner-wall separation of 3.9 µm. The on-chip and off-chip waveguide separation was designed to be 250 µm, which allows us to potentially couple light into or out of both waveguides with commercially available fiber v-groove arrays.

In addition to conventional directional couplers, we also fabricated a set of improved insensitive couplers, using the design rules described in Section 2.2.5. Figure 4.2 illustrates the structure of the insensitive directional coupler. The relevant device dimensions are the same as those specified in Fig. 4.1. The target power splitting ratio for these devices was 50%, although the design could easily be modified to accommodate any desired power-splitting ratio. The phase delay was achieved by making the upper arm of the device slightly longer than the lower arm. The path length difference between the two arms was approximately 350-360 nm. We did not include a range of coupling lengths (ϕ_1 and ϕ_2) in our design, rather we simply used our best estimate for these parameters based upon coupled-mode calculations. We did, however, include a range of phase shifts (θ) above and below the nominal design.

4.1.2 Measurement Technique

After the devices were fabricated and over-cladded, the samples were cut into small chips using a die saw. The facets were then polished using a Bueler Ecomet tabletop polishing system. Initially, we polished with a minimum grit-size of 0.1 µm, but we later found that generally a 0.3 µm grit pad provides a sufficiently smooth surface. The input facets were

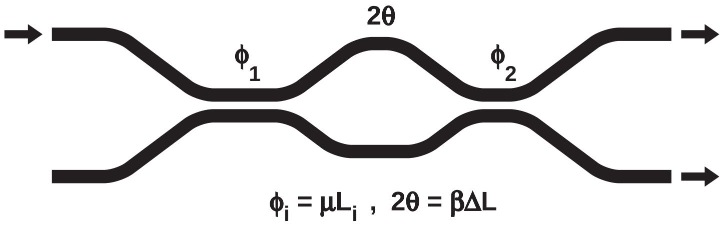

Figure 4.2: Schematic diagram of an improved directional coupler, formed by cascading two conventional couplers in a Mach-Zehnder configuration, with an intervening phase shift. The three parameters ϕ_1, ϕ_2 and θ were chosen to yield a coupler with 50% splitting ratio, as described in Section 2.2.5. Specifically, we chose $\phi_1 = \pi/2$, $\phi_2 = \pi/4$, and $\theta = \pi/3$. (figs/4/insensitive-coupler-diagram.eps)

polished normal to the waveguide direction to enable coupling to a cleaved optical fiber, and the output facets were polished at an angle of 10 degrees from the vertical to eliminate unwanted reflections into the waveguide. This polishing scheme causes the light to emerge from the facet at a small downward angle (out of the plane of the chip).

Figure 4.3 illustrates the setup of our measurement system for characterizing directional couplers. We measured the performance of the couplers using an external cavity tunable diode laser. The attainable wavelength range for the source is approximately 1480–1590 nm.

Light from the laser is first directed through a polarization controller, which allows us to polarize the output light in any desired direction. (The method of polarization control will be described later.) Next, a fused-fiber coupler is used to tap off a fraction of the input light. This tapped-off signal provides a monitor power against which other measurements are compared. This allows one to eliminate the effects of drift in the output laser power.

The light is butt-coupled into the waveguide directly from the cleaved facet of the optical fiber, using an index-matching gel to suppress reflections at the facet. If the index-matching gel is not used, there is generally a small air gap between the input fiber and the polished waveguide facet which leads to undesirable Fabry-Perot effects. Even though the free spectral range of such a thin air gap is typically very large, the measured power will change substantially if the air gap increases or decreases by a fraction of the optical wavelength. When using an index-matching gel, the input coupling efficiency is insensitive to

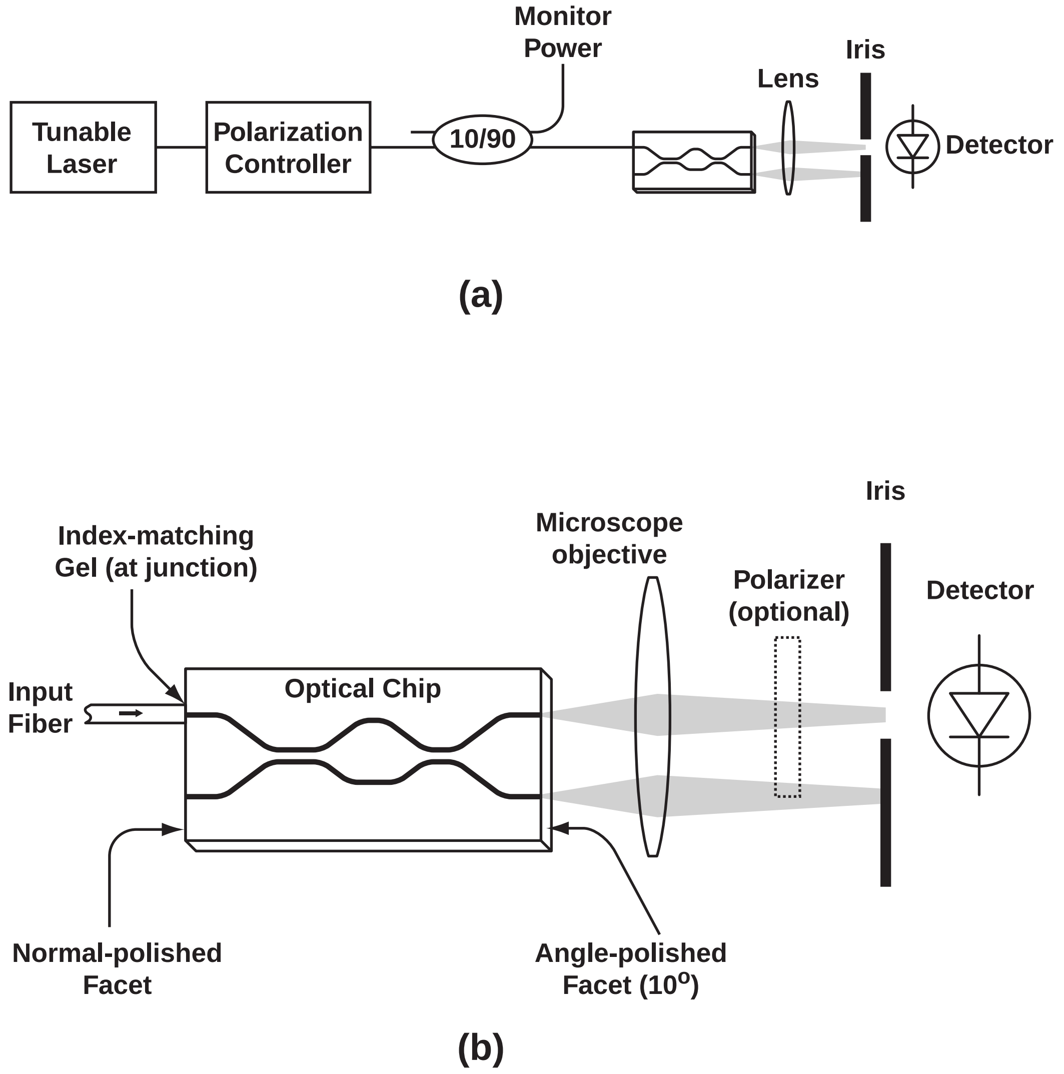

Figure 4.3: Measurement system used for characterizing integrated couplers. (a) Light from a tunable external cavity diode laser is passed through a polarization controller, 10/90 splitter, and then launched into the device. (b) Light is coupled directly from the fiber into the waveguide facet, using an index matching gel. The signal emerging from the output facet is imaged onto a photodetector (through a linear polarizer.) (figs/4/coupler-measurement-diagram.eps)

the size of the gap between the waveguide and fiber.

At the output facet, the light emerges from the chip into free space, where it is captured by a microscope objective which forms an image of the chip facet on a detector. An iris is used to selectively measured the light emerging from one waveguide.

The input fiber and a microscope objective are both mounted on 3-axis translation stages, allowing accurate, stable alignment of the input fiber and imaging optics. The optical chip is likewise mounted on a translation stage, but this stage is typically kept stationary except when moving from one device to the next.

Initially, the detector is replaced with an infrared camera in order to align the input fiber and focus the microscope objective on the facet. After the system is aligned, the infrared camera is replaced with a detector, and each of the alignment stages is adjusted to maximize the detected power.

As shown in Fig. 4.3b, a linear polarizer is inserted immediately before the iris in order to determine the polarization state of the light emerging from the waveguide. The input fiber used in this experiment is not polarization-maintaining, which means that using the polarization controller alone we cannot accurately predict the polarization state of the light entering the waveguide. However, by placing a linear polarizer before the detector, it is possible to fix the polarization state using an empirical approach. First, the linear polarizer is oriented in the vertical (TM) direction. Then, the polarization controller is adjusted to minimize the detected power. This guarantees that the light emerging from the waveguide is TE-polarized [1]. Wherever possible, the fiber path between the polarization controller and the input waveguide is clamped to the table to prevent any twisting or bending which might alter the polarization state. After the polarization is adjusted in this manner, the linear polarizer can be removed (or rotated by 90 degrees). Provided the input fibers are relatively stationary, this technique provides fairly stable polarization control.

There is unfortunately no easy way to insert a linear polarizer between the input optical fiber and the chip. We therefore make the assumption that when the light emerging from the output facet is TE-polarized, the input light is likewise TE-polarized, and similarly for TM. This assumption is quite reasonable for most integrated waveguides. (In fact, if one could build an integrated device which rotates the polarization state from TE to TM, it might solve many of the problems which currently plague integrated optical devices.)

---

[1] It is easier in practice to zero the detected power with polarizer in the TM position that it is to maximize the power with the polarizer in the TE position.

The directional couplers are measured by scanning the input laser wavelength while measuring the output power. More specifically, we measure the ratio of the output power to the monitor power in order to account for any drift or variation in the laser power. The microscope objective is then translated in order to image the second output port on the detector, and wavelength scan is repeated.

4.1.3 Measurement Results

Before measuring the devices, the output of the optical fiber was focused directly onto the detector in order to quantify the power entering the device. When the device was inserted into the light path, the total power transmission decreases by approximately 1.5 dB. This total insertion loss figure includes fiber-coupling loss at the input facet, bending loss, and intrinsic waveguide loss.

Figure 4.4a shows the measured splitting ratio for our initial set of directional couplers. These devices were designed to have a power-splitting ratio of 50% at λ = 1550 nm. As shown, the coupler achieves a splitting ratio which is higher than the desired value of 50% for all wavelengths measured and for both polarization states. We attribute this to uncertainty in the indices of refraction used in our calculations when designing devices, and structural and material deviations from the nominal design, which may occur during the various fabrication steps. In principle, this bias can be removed by performing a more exhaustive empirical investigation of the material properties and by more carefully controlling and characterizing each step of the fabrication process. However, in addition to this offset, the splitting ratio for the conventional coupler exhibits a sloped wavelength dependence which is characteristic of directional couplers and cannot be flattened without using a different design.

More importantly, the conventional coupler also exhibits significant polarization dependence: the splitting ratio is consistently about 5% higher for TM- than for TE-polarized light at all of the wavelengths we measured. This discrepancy cannot be accounted for unless we assume that the index profile of the waveguide is different for TE and TM polarizations. This can occur if there is birefringence in the deposited glass films, especially the top cladding layer which fills the region between the waveguide.

Figure 4.4b plots the measured power-splitting ratio for an improved directional coupler, designed using the techniques described in Section 2.2.5. In spite of the polarization dependence and the apparent uncertainty in material and fabrication parameters, the improved coupler achieves a splitting ratio between 45 and 50% (within 5% of the targeted

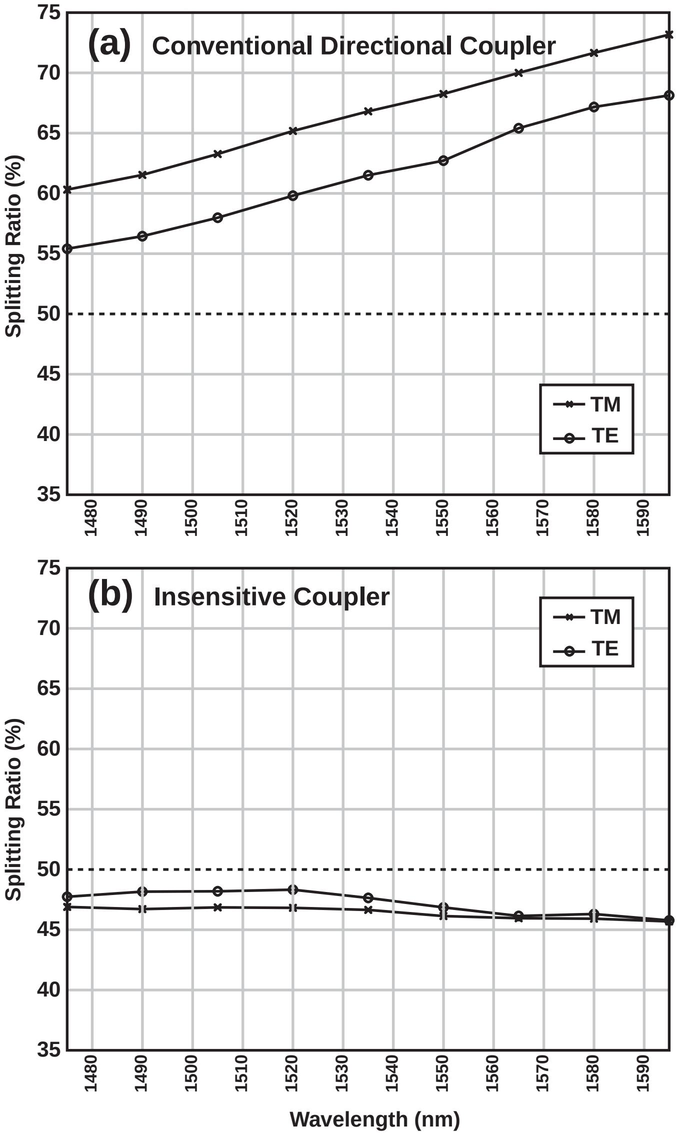

Figure 4.4: (a) Measured splitting ratio for the initial set of integrated directional couplers. The achieved power splitting ratio is uniformly higher than the target value of 50%, for both polarization states. (b) The improved directional coupler exhibits a wavelength-flattened response which is much closer to the desired splitting ratio. (figs/4/coupler-comparison-orig.eps)

value) over the entire wavelength range from 1475 to 1595 nm, with very little polarization dependence.

We fabricated approximately 12 such cascaded Mach-Zehnder devices, all of them with identical couplers but with different phase shifts. In principle, it should be possible to precisely control the phase shift (ie, θ) through the photolithographic process. One of the most prevalent integrated optical devices available today, the waveguide-grating-router [141], relies for its operation on a path-length difference which must be lithographically defined to the sub-wavelength level. In order to fabricate such a lithographically-defined phase difference in the opposing arms of a Mach-Zehnder interferometer, one must: (1) accurately know the effective index of propagation for the waveguide, and (2) use a photomask that is generated using a grid size that is sufficiently small to properly describe the waveguides.

In order to help reduce the cost of the photomask set, and to minimize the number of electron-beam field boundaries spanned by the device, we used a relatively large grid size of 1/8 µm for these initial devices. Moreover, when designing the photomask set, we had performed no prior experimental measurements of the effective index for the waveguides. Thus, we simply relied on our best estimate for the effective index, based upon vendor-supplied data. We included a range of phase shifts in the design in order to allow for the uncertainties.

Unfortunately, our measurements indicate that the range of phase shifts (θ) which we actually obtained with these devices was significantly lower than the target value of $\pi/3$. Only one of the 12 devices was relatively close to $\pi/3$. The data plotted in Fig. 4.4b was obtained from this device.

In order to address this problem, we fabricated a second set of couplers using our own e-beam lithography to directly write the patterns. For these devices, the grid size was approximately 24 nm, which allows us to more accurately describe the pathlength difference. Furthermore, we adjusted the directional couplers in an effort to center the splitting ratio at the desired value of 50%.

Figure 4.5 presents the measured splitting ratio for the second set of devices. As shown in Fig. 4.5a, the power splitting ratio for the conventional directional coupler is now centered on 50%, but, as expected, the sloped wavelength dependence has not been removed, nor has the polarization-dependence been eliminated. Figure 4.5b presents the measured power-splitting ratio for the insensitive directional couplers. As before, the insensitive couplers achieve a power-splitting ratio that is close to the desired value of 50% for both

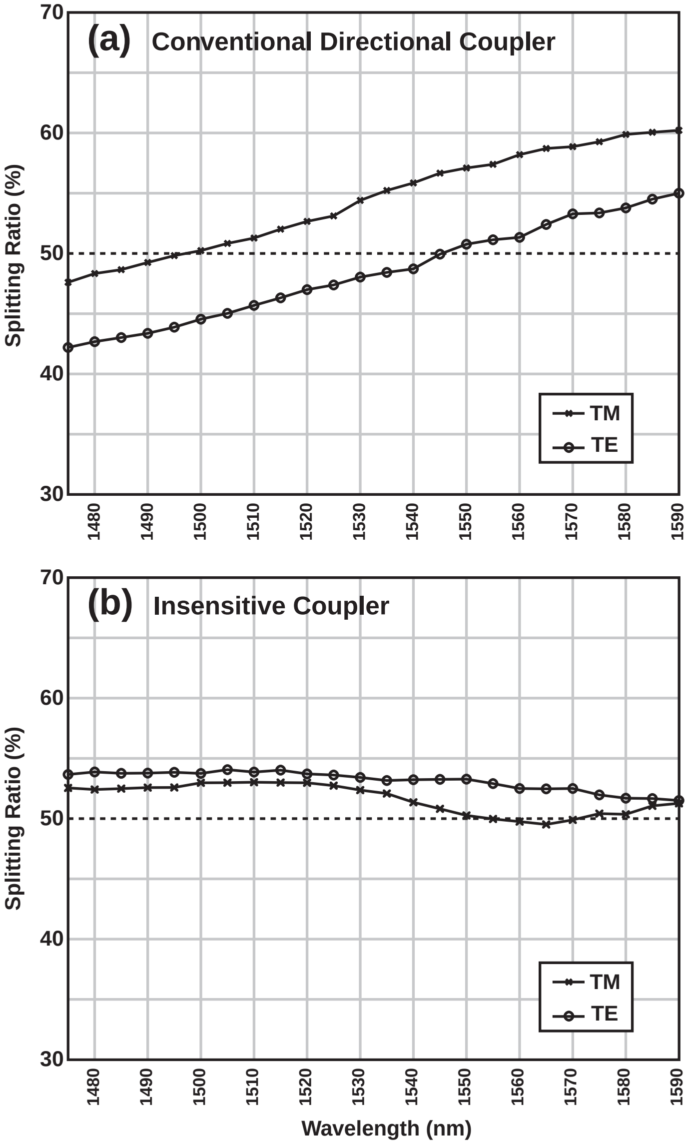

Figure 4.5: (a) Measured splitting ratio for the second set of integrated directional couplers. The achieved power splitting ratio is now centered on 50%, but there is still a significant polarization-dependence as well as the characteristic sloped wavelength dependence. (b) The measured splitting ratio for the improved directional coupler exhibits a wavelength-flattened response that is much closer to the desired splitting ratio of 50% (figs/4/coupler-comparison-data.eps)

polarization states.

Unfortunately, although the insensitive coupler provides a reduced wavelength-dependence, the achieved power-splitting ratio seems to be slightly higher than the desired value of 50%. Recall that the analysis presented in Fig. 2.18 predicts that the insensitive directional coupler should be able to correct for relatively large fractional changes which proportionately affect each of the constituent directional couplers. The reason for the apparent discrepancy is that the two constituent directional couplers in the device are not centered at the same wavelength. That is, the two directional couplers which make up the device achieve their nominal values for ϕ_1 and ϕ_2 at two different wavelengths. Our measurements indicate that the first directional coupler of the structure, which should nominally provide full-power transfer, achieves its target value of $\phi_1 = \pi/2$ at $\lambda = 1518$ nm for the TE polarization, whereas the second directional coupler achieves its nominal value of $\phi_2 = \pi/4$ at $\lambda = 1547$ nm. Because of this, the power-splitting ratio achieved is not as insensitive as it would otherwise be. Although the structure can compensate for fractional changes in ϕ_1 and ϕ_2, it cannot correct for a miscalibration or design flaw which only affects one of the directional couplers and not the other. In spite of this limitation, the device does provide improved performance over the conventional directional coupler. Moreover, we believe that this problem can be solved by performing more exhaustive experimental measurements of the waveguide coupling rate.

For the second set of devices, we also investigated the effect of changing the intervening phase shift θ above and below the nominal value of $\pi/3$. Figure 4.6 presents the measured power-splitting ratio for the insensitive couplers when the phase shift θ is changed by $\pm 2^\circ$and $\pm 4^\circ$. These plots demonstrate that the device performance is relatively insensitive to the phase shift θ.

Finally, we looked at how the device response changes when the waveguide width is increased or decreased with respect to its nominal value. We intentionally fabricated couplers which were identical to those described earlier, but with the waveguides made uniformly wider or narrower in order to simulate the effect of an unintended feature size bias in the fabrication process. Note that the center-to-center waveguide separations were held constant while the waveguide width was changed. Figure 4.7 plots the achieved power-splitting ratio for the conventional and improved directional couplers when the waveguide width is increased and decreased by 0.2 µm. As shown in Fig. 4.7b, the insensitive directional couplers retain their wavelength-flattened response when the waveguide width is changed. There is, however, a small change in the power-splitting ratio as a result of the change in waveguide width. Again, this change can be ascribed to the miscalibration of the constituent directional couplers described earlier.

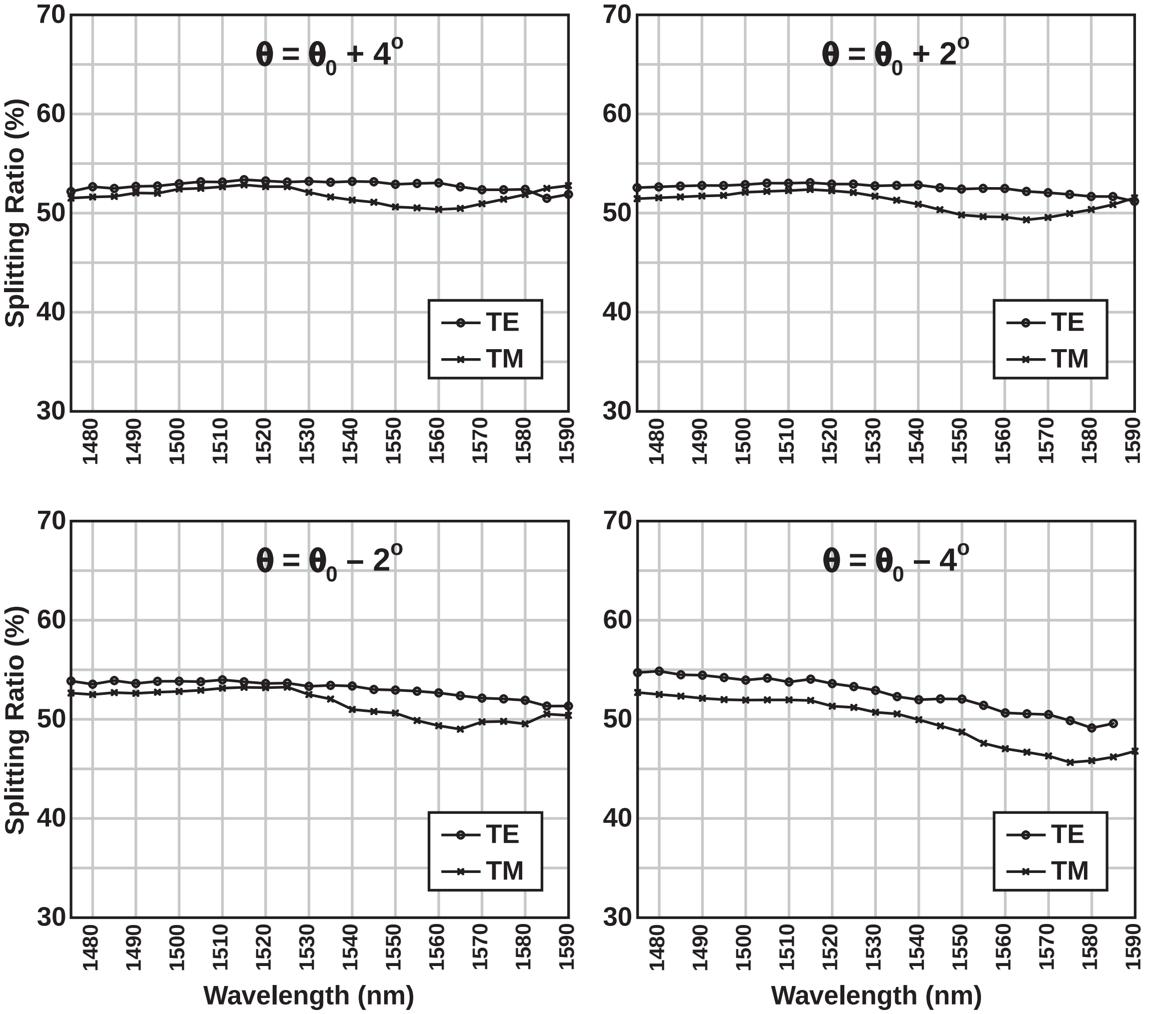

Figure 4.6: Measured performance of the insensitive directional coupler, when the intervening phase shift θ is increased or decreased with respect to its nominal value of $\pi/3$. (figs/4/coupler-vary-theta.eps)

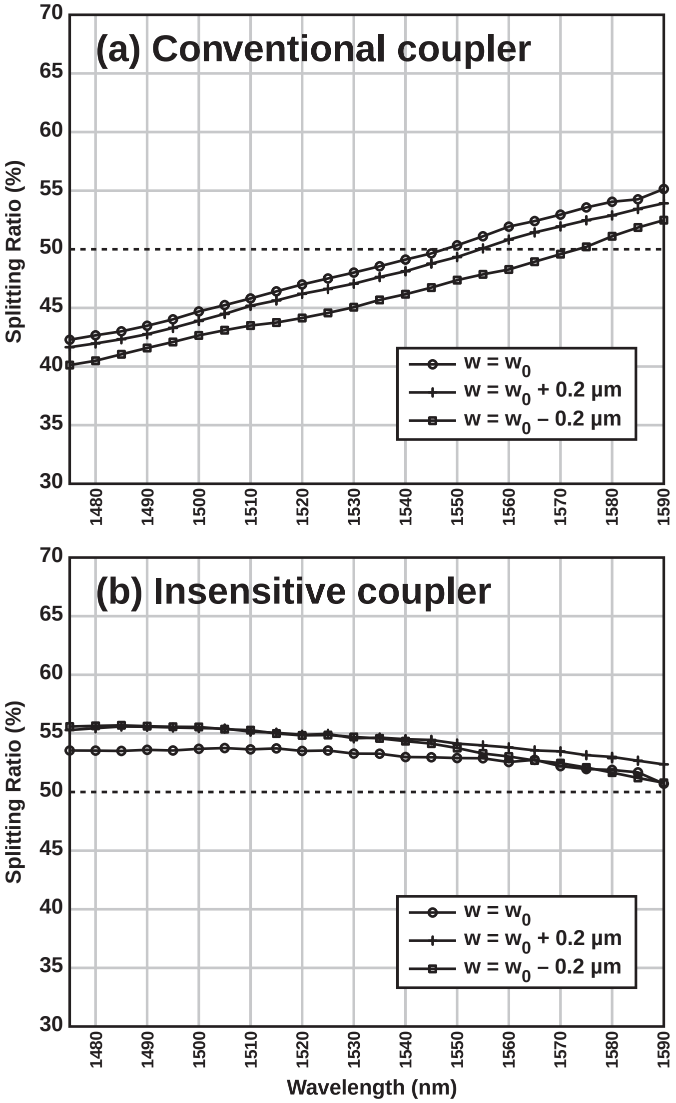

Figure 4.7: Measured performance of the insensitive coupler, when the waveguide width is biased above and below its nominal value of 6.6 µm. (figs/4/wider-narrower.eps)

4.2 Silicon-on-Insulator Waveguide Measurements

4.2.1 Loss Characterization

Before fabricating integrated Bragg gratings, we first constructed a set of conventional waveguides in order to evaluate the propagation loss for the structure. In order for a Bragg grating to function as desired, the linear loss coefficient (denoted by α) must be significantly smaller than the grating strength κ. In other words, the rate of reflection into the counter-propagating mode must be greater than attenuation rate in order to build working Bragg grating devices. Moreover, the propagation loss is an important figure-of-merit for any integrated waveguide. In order to characterize the loss, we fabricated a set of simple straight ridge waveguides of approximately 5 mm length using the techniques described in Section 3.1.4.

The samples were cut with a die saw and polished using the technique described earlier for glass waveguides. However, because of the smaller mode size, and higher refractive index relative to glass, the samples were polished with a minimum grit-size of 0.05 µm. Prior to dicing, a thick photoresist layer was spun over the waveguides to protect them from damage during subsequent sawing and polishing. Both chip facets were polished normal to the waveguide to facilitate loss measurement via a Fabry-Perot technique[142], which we describe next.

When the chip facets are polished normal to the waveguides, the waveguide forms a Fabry-Perot cavity, because of the internal reflection at the air-silicon interface. By measuring the spectral response of this cavity, it is possible to infer the waveguide propagation loss. The spectral transmission of a Fabry Perot cavity can be calculated using transfer matrices as described in Section 2.4 [76]. We shall omit the derivation and simply provide the result:

$$T(\beta) = \frac{T_0}{\left|1 - r_1 r_2 e^{-2(\alpha + j\beta)L}\right|^2} \quad , \tag{4.1}$$

where T_0 is a constant which represents what the power transmission would be in the absence of any internal reflections, r_1 and r_2 are the internal amplitude reflectivities at the two opposing chip facets, L is the facet-to-facet waveguide length, β is the propagation constant, and α is the amplitude loss coefficient. This function can be rewritten in the

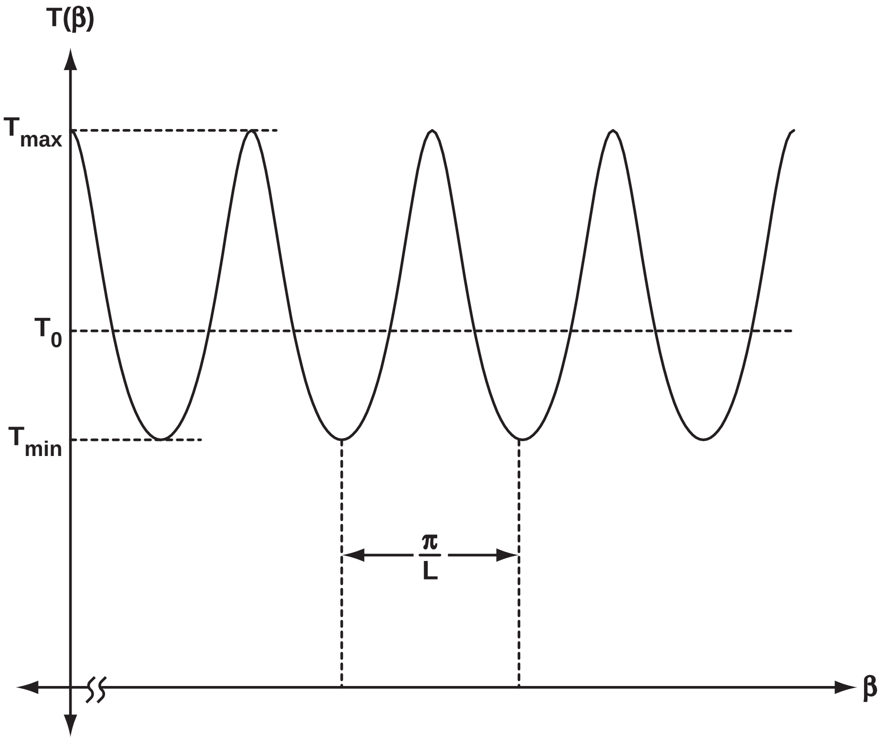

Figure 4.8: Diagram of theoretical Fabry-Perot transmission spectrum, calculated for the case where $A = |r_1 r_2| \exp(-2\alpha L) = 0.2$. (figs/4/fabry-calculated.eps)

following simple form:

$$T(\beta) = \frac{T_0}{|1 - A\exp(-2j\beta L + \phi)|^2} \quad , \tag{4.2}$$

where A is a scalar constant, less than 1, which is related to the facet reflectivities and loss coefficient by,

$$A \equiv |r_1 r_2| e^{-2\alpha L} \quad , \tag{4.3}$$

and ϕ is a phase constant which allows for the possible nonzero phase of the complex amplitude reflection coefficients r_i. Figure 4.8 plots the calculated power transmission for the case where $A = 0.2$.

From Eq. 4.2, the ratio of the maximum to the minimum transmission is:

$$\frac{T_{\text{max}}}{T_{\text{min}}} = \left(\frac{1+A}{1-A}\right)^2 \quad . \tag{4.4}$$

This equation can be readily inverted to obtain an expression for A in terms of the maximum and minimum transmission.

$$A = \frac{\sqrt{T_{\text{max}}} - \sqrt{T_{\text{min}}}}{\sqrt{T_{\text{max}}} + \sqrt{T_{\text{min}}}} \quad . \tag{4.5}$$

Thus, A may be viewed as a measure of the *contrast* in the Fabry-Perot spectrum.

The internal reflection coefficient at the chip facet can be approximated using the standard Fresnel reflection equation for normal incidence:

$$|r| = \frac{n-1}{n+1} \quad . \tag{4.6}$$

For the waveguides discussed here, because the mode is almost entirely confined in the silicon core region, one can use $n = n_{\text{Si}} \simeq 3.5$ in the above equation.

After measuring the contrast in the Fabry-Perot fringes and inferring the internal reflection from Eq. 4.6, it is possible to estimate the propagation loss using Eq. 4.3 and Eq. 4.5.

Figure 4.9 illustrates the experimental set up used to perform the Fabry-Perot loss measurement. As for the glass waveguide measurements described in Section 4.1, the illuminating source was a tunable external cavity diode laser connected through a polarization controller and 10/90 power splitter. Light was coupled into the waveguides using a conical-shaped lensed fiber, which has a minimum spot size of approximately 5 μm and a focal length of approximately 15 μm. Light emerging from the waveguides was focused onto a detector through a linear polarizer. Figure 4.10 depicts a representative measured spectrum for one of the devices considered. Using the equations given above, the propagation loss for the sample is calculated to be about 4 dB/cm.

One of the limitations of this technique is that the loss estimate relies upon an accurate estimate of the internal reflectivity at the two facets. For perfect facets, this reflectivity should be reasonably described by Eq. 4.6. However, if the facet is scratched, damaged, or misoriented, the internal reflectivity might deviate significantly from this estimate. In practice, we have noticed that there can be significant device-to-device variation in the loss estimate, even for adjacent waveguides on the same chip. In comparison to the buried the glass channel waveguides described earlier, the facets of these air-cladded ridge wave-

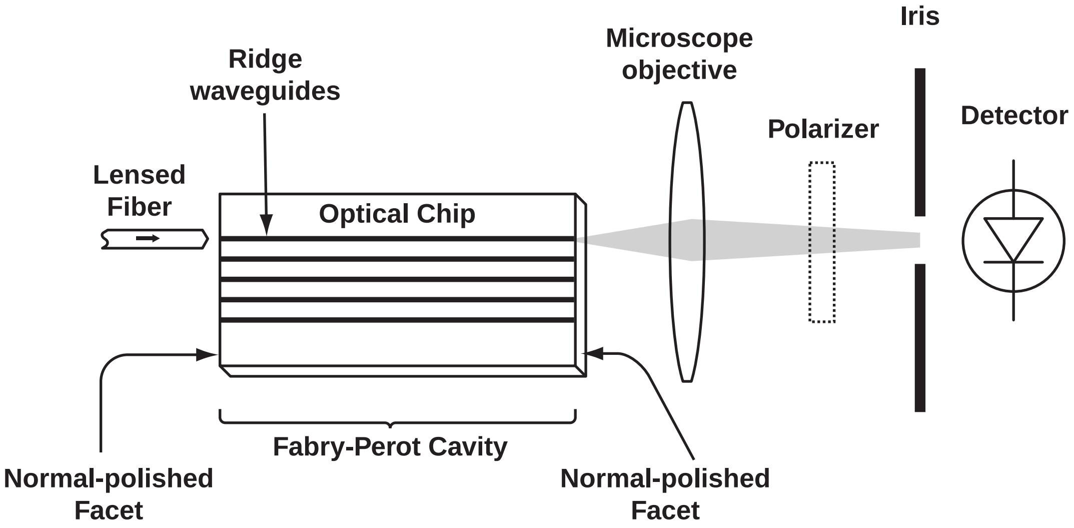

Figure 4.9: Experimental setup used to perform Fabry-Perot measurements of ridge waveguides. Light is launched into the waveguide facet using a lensed fiber, while the output facet is imaged onto a photodetector with a microscope objective. The illuminating source (not depicted) was a tunable infrared laser. (figs/4/loss-measurement-diagram.eps)

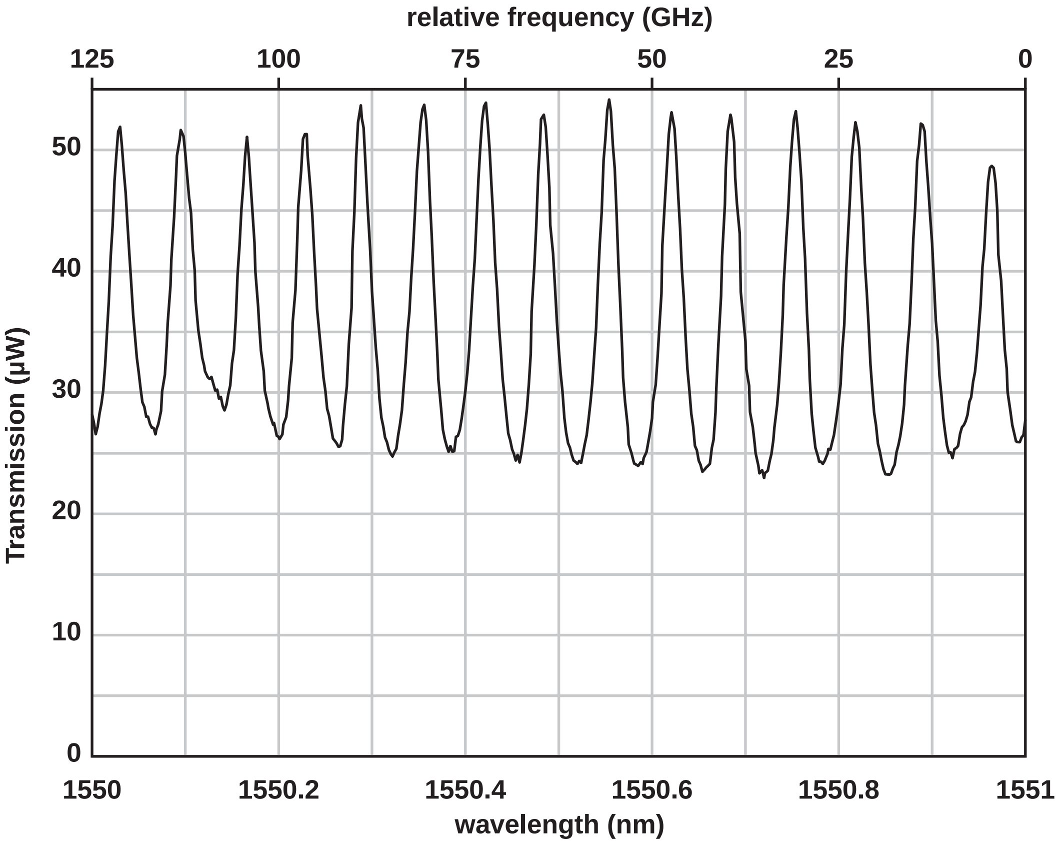

Figure 4.10: A representative Fabry-Perot transmission spectrum for an integrated silicon-on-insulator ridge waveguide. The total waveguide length for the sample was 4.88 mm. The propagation loss inferred from these measurements is approximately 4 dB / cm. (figs/4/fabryloss.eps)

guides are more susceptible to damage during the process of polishing because the guiding layer is not as well protected. This uncertainty limits the accuracy with which we can estimate the propagation loss. We point out that this uncertainty would also apply for more conventional cutback loss measurement techniques.

Based upon our measurements, we estimate a propagation loss of approximately 4 ± 2 dB/cm. This corresponds to an amplitude loss coefficient [2] of $\alpha = 0.5 \pm 0.2$ cm^{-1}, which is significantly lower than the anticipated grating strengths κ of 3-4 cm^{-1}. The measured propagation loss is approximately equivalent (within the measurement uncertainty) for TE and TM polarizations.

4.2.2 Measurement of Integrated Bragg Gratings

After establishing that the propagation loss for the SOI ridge waveguides was sufficiently low, we next constructed a set of integrated Bragg gratings. Figure 4.11 summarizes the relevant geometrical parameters for these devices.

The waveguide width was 4 µm, and the silicon ridge layer was 3 µm thick at its center. The ridge was etched to a depth of 800 nm. The Bragg gratings were etched to a depth of about 150 nm, with a duty cycle of 45-50%, and a period of approximately 223 nm. The grating length for the initial set of devices was approximately 4 mm, and the complete chip length was 17 mm.

In contrast to the couplers described earlier, these devices must be angle-polished on *both* facets in order to avoid forming Fabry-Perot cavities between the grating and facets. Figure 4.12 illustrates three possible ways to prepare the chip facets to eliminate internal reflection. The first approach, depicted in Fig. 4.12a refracts the light out of the plane of the chip, which makes measurement of the device difficult[3]. The second method, depicted in Fig. 4.12b corrects this problem but the devices in the set must be staggered. We chose to use the scheme depicted in Fig. 4.12c, in which the waveguides are bent to an angle of six degrees at the edges of the chip. With this approach, the facets can be cut and polished normally. The bending radius was chosen to be 15 mm, which is predicted to yield a negligible bending loss for the waveguides considered.

Figure 4.13 depicts the measurement set up for the integrated Bragg gratings. Notice

---

[2] The amplitude loss coefficient α is related to the loss in dB per unit length by $\alpha = \frac{\ln(10)}{20} \times$ (dB/length)

[3] We were able to use this technique for the glass channel waveguides because the refraction is significantly weaker for lower-index glass waveguides.

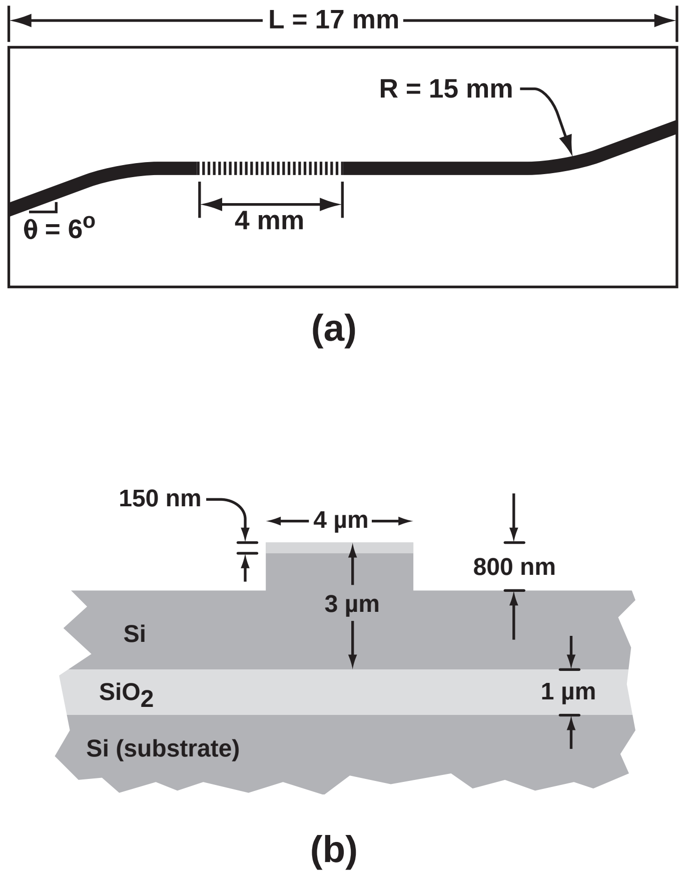

Figure 4.11: Geometry of the silicon-on-insulator ridge waveguides with Bragg gratings. (a) Top-down diagram illustrating overall device dimensions. (b) Waveguide and grating cross-section. (figs/4/soi-grating-geometry.eps)

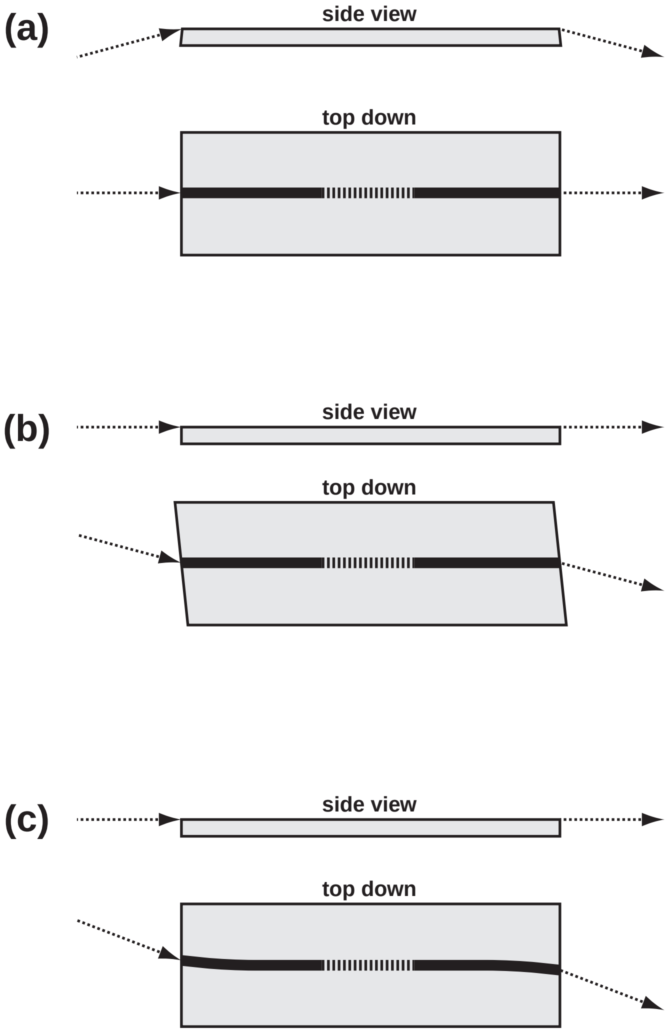

Figure 4.12: When measuring integrated Bragg gratings, it is important to eliminate internal reflections at the chip facets. This can be achieved by orienting the facets at an angle with respect to the waveguides. This figure illustrates three different ways of angle-polishing the waveguides facets. We have chosen to adopt the scheme depicted in (c), in which the waveguides are bent slightly at either end of the chip. (figs/4/angle-polishing-schemes.eps)

that the chip must be rotated by 21 degrees in order to account for the refraction at the input and output facets. Light was coupled into the device using a lensed fiber, and imaged at the output facets using a microscope objective.

Prior to measuring the devices, the output of the lensed fiber was focused directly onto the detector in order to quantify the relative amount of light entering the device. After the device was inserted and all of the elements aligned, the detected power decreased by 12.3 dB for the TE polarization and 13.6 dB for the TM polarization. These insertion loss figures include mode mismatch at the input facet, waveguide propagation loss, possible grating induced loss, Fresnel reflection loss at both facets, and bending loss. The theoretically calculated Fresnel reflection loss is 2.9 dB for TE and 3.5 dB for TM[4]. From our prior measurements of the waveguide propagation loss for similar structures without gratings, the waveguide themselves account for 3-10 dB of loss. Because of this uncertainty in our estimate of the waveguide propagation loss, it is difficult to quantify the excess loss associated with the grating. However, as we will describe in Appendix C, we also measured 8 mm-long Bragg gratings in addition to the 4 mm gratings described here. The measured insertion loss did not increase measurably when the grating was made longer, indicating that the Bragg gratings introduce very little additional loss.

Figure 4.14 depicts the measured TE transmission spectrum for a 4 mm-long integrated Bragg grating. A similar spectrum is seen for the TM polarization. There are four distinct dips in the transmission spectrum (and there might be more at shorter wavelengths which our laser could not reach.) The first three dips at 1510, 1525, and 1537 nm are caused by coupling of the forward traveling mode to radiation modes. The final dip at 1543 nm corresponds to the anticipated coupling from the forward to the backward guided mode. We associated the short-wavelength transmission dips with radiation loss after measuring the device in reflection using a circulator. Whereas there is no measured reflected signal at the wavelengths corresponding to the first three transmission dips, we saw a peak in the reflected signal at wavelengths corresponding to the transmission dip at 1543 nm.

The radiation loss on the blue side of the spectrum is a well studied phenomenon for fiber Bragg gratings [143]. As in the case of fiber Bragg gratings, this loss can be identified with grating-assisted coupling between the forward-traveling mode and backward-traveling cladding modes. The spectrum of the cladding modes can be computed using the beam propagation method, as described in reference [143].

To illustrate the physical mechanism of this radiation mode coupling, consider the sim-

---

[4] Because the waveguide are illuminated at an angle to the facet, the reflection is different for TE and TM, which is why the computed Fresnel reflection loss differs between the two polarization states.

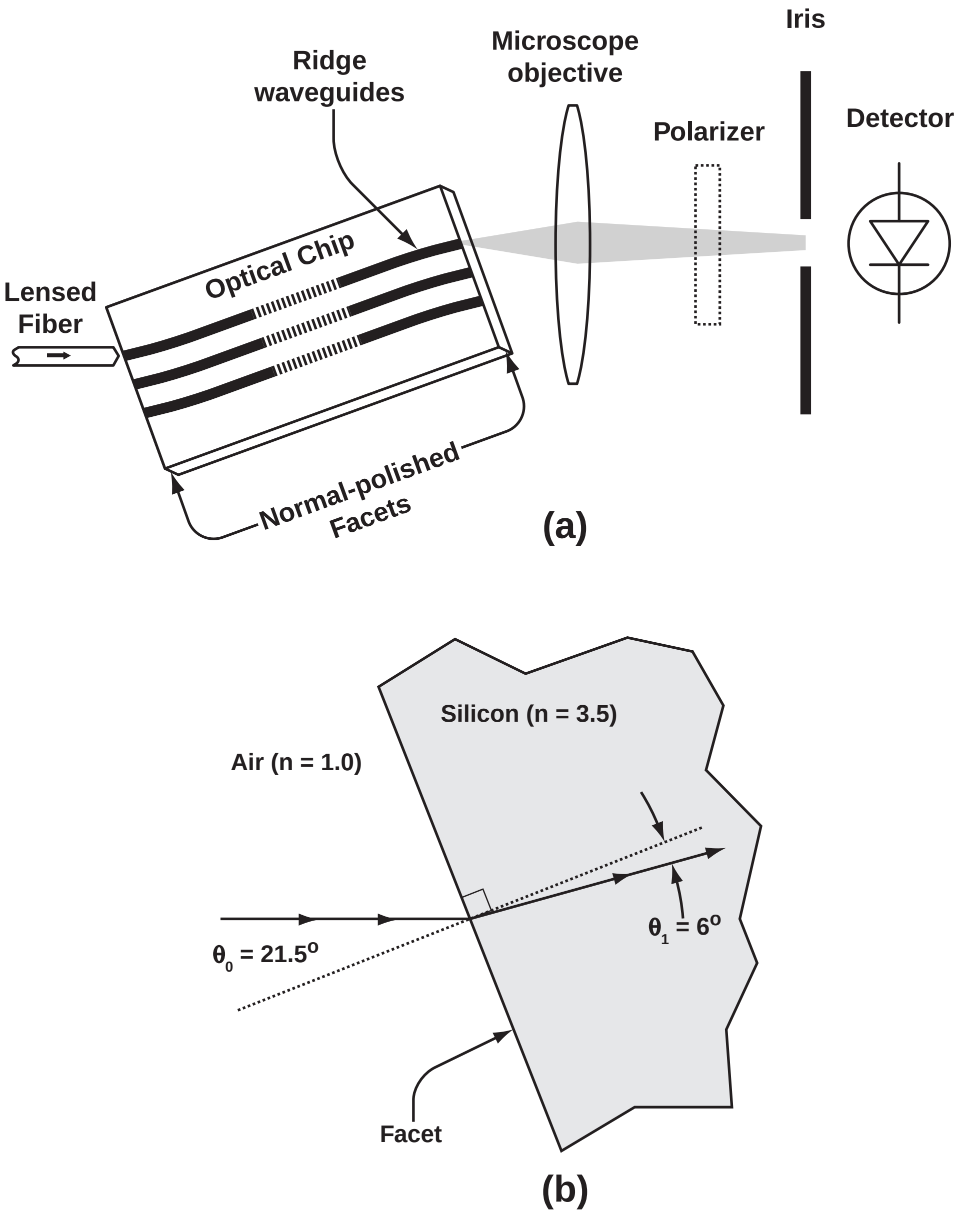

Figure 4.13: (a) Experimental setup used to measure the transmission spectral response for integrated Bragg gratings in silicon-on-insulator ridge waveguides. (b) Enlarged diagram depicting the light path at the input facet. Because the waveguides approach the facet at an angle of six degrees from the normal, the entire chip must be rotated by 21 degrees to allow for refraction as the light enters the chip.

(figs/4/grating-measurement-diagram.eps)

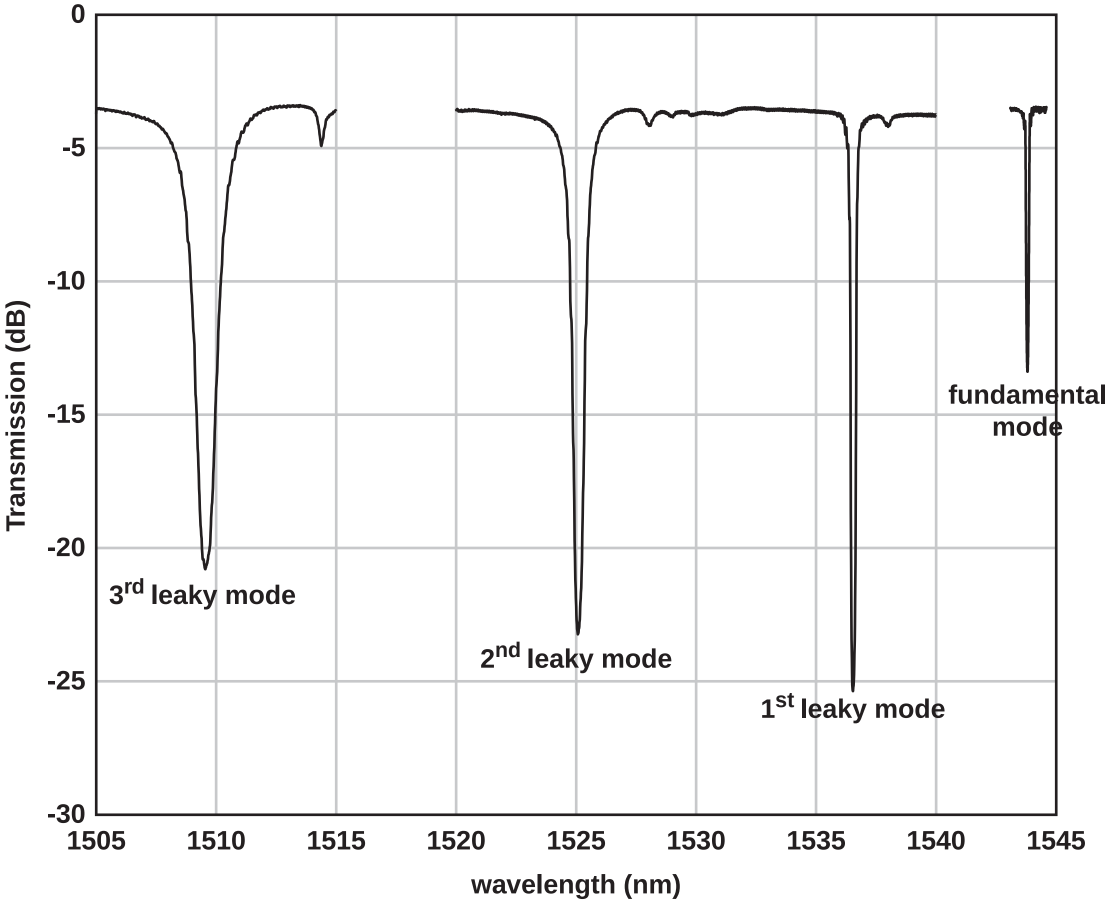

Figure 4.14: Measured transmission spectral response for a 4 mm long Bragg grating on a silicon-on-insulator ridge waveguide. The spectral response shows four distinct dips in transmission. The three peaks on the blue side of the spectrum are attributed to radiation loss, whereas the narrow dip at 1543 nm corresponds to Bragg reflection into the backward-traveling bound mode. (figs/4/radiationmodes.eps)

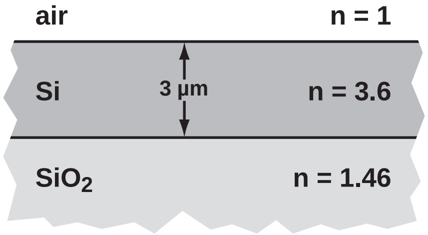

Figure 4.15: The nature of the radiation losses seen in Fig. 4.14 can be explained by considering this simplified structure consisting of a simple SiO_2:Si:air slab waveguide whose dimensions are chosen to be similar to the ridge waveguide which we actually measured. (figs/4/simple-slab-model.eps)

plified problem of a silicon slab waveguide, whose dimensions are chosen to be similar to the ridge waveguide which we actually measured. Specifically, we let the silicon guiding layer be 3 µm thick and we assume that the lower cladding is SiO_2 ($n = 1.46$) and the upper cladding is air ($n = 1$), as depicted in Fig. 4.15. The propagation constants, β_i, can be computed exactly for each of the modes of this structure[14]. Figure 4.16 depicts the calculated dispersion relation (β vs. λ) for the first four TE modes of the structure. Notice that we have also plotted the negative propagation constants which correspond to the lowest four backward-travelling modes. The Bragg grating creates a coupling between the modes wherever the difference in propagation constants is equal to the k-vector of the grating:

$$\beta_n(\lambda) - \beta_m(\lambda) = \frac{2\pi}{\Lambda} \quad . \tag{4.7}$$

The four arrows drawn in Fig. 4.16 illustrate the four wavelengths at which light from the forward-traveling fundamental mode can couple to a backward traveling mode. The position and relative spacing of these four wavelengths agrees qualitatively with the measured data presented in Fig. 4.14, in spite of the simplified approximations made (i.e., replacing the ridge with a slab.)

Of course, the ridge waveguide was designed to have only one bound mode (for each polarization state), which raises the question of why we measure these extraneous dips in the transmission. The explanation is that although the waveguide only supports one bound mode, there are several *leaky modes* associated with the higher-order slab modes of the structure. For these leaky modes, the mechanism of power loss is that the light escapes

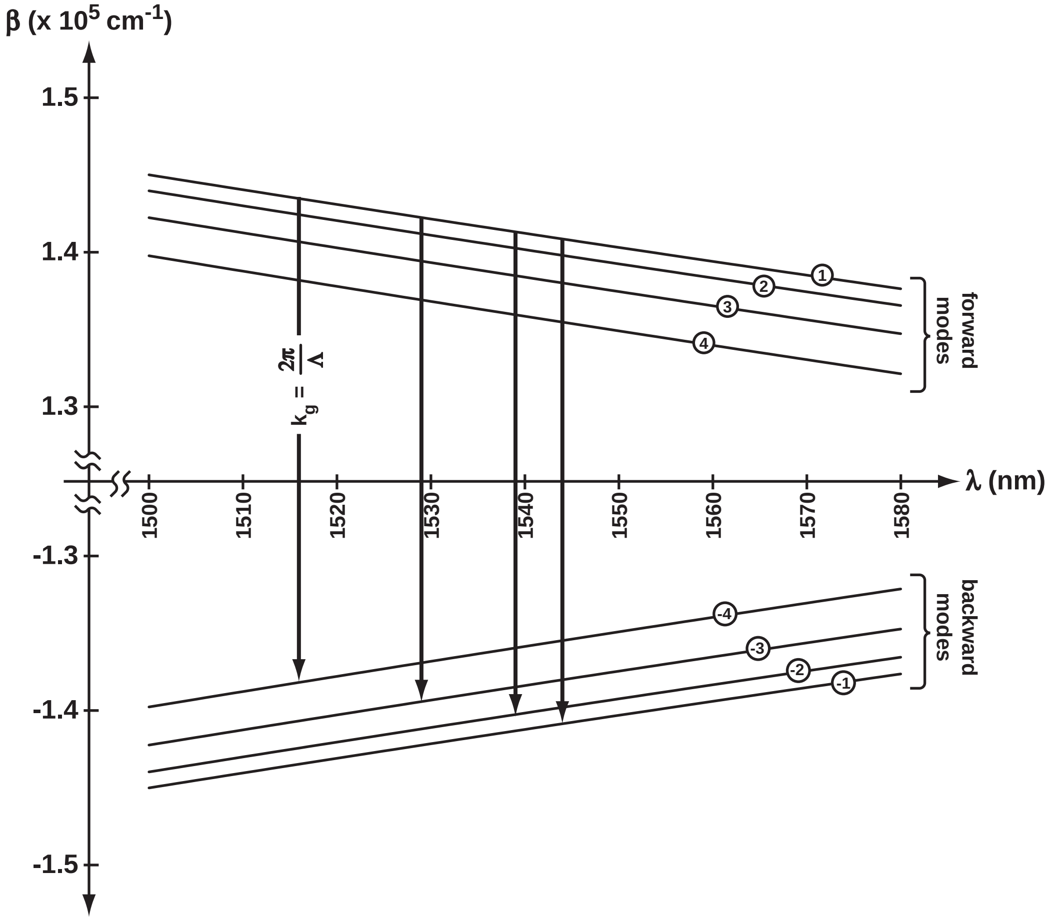

Figure 4.16: Calculated dispersion relation (β vs. λ) for the first four modes of the slab waveguide structure depicted in Figure 4.15. The corresponding backward-traveling modes (with negative values for β) are also plotted. Coupling can occur between modes wherever the difference in propagation constants matches the spatial frequency (k-vector) of the Bragg grating. The four vertical lines in this figure indicate the four wavelengths at which light from the forward-traveling fundamental mode can couple to one of the four backward-traveling modes. (figs/4/leaky-slab-simulation.eps)

by leaking out into the silicon slab on either side of the ridge. By contrast, the fundamental mode is completely confined by the ridge, meaning that the fields decay evanescantly to either side of the ridge.

Figure 4.17 is an expanded plot of the fundamental Bragg transmission spectrum. We have included both TE and TM spectra for comparison. The separation between the TE and TM peaks is approximately 50 GHz, or 0.4 nm. The theoretically predicted modal birefringence for the structure is 43 GHz, or 0.35 nm. Another effect seen in Fig. 4.17 is that the TE-polarized mode experiences a stronger Bragg reflection than the TM mode. Based upon the depth of the transmission response, we estimate that the grating strength (κ) is 4.5 cm^{-1} for the TE mode and 2.1 cm^{-1} for the TM mode. The TE grating strength agrees well with calculations based on coupled mode theory. (cf Fig. 2.23), while the TM result is somewhat lower than the value of 3.7 cm^{-1} predicted by coupled mode theory.

Figure 4.18 plots a comparison between the measured theoretical spectral response (for TE polarization) and the calculated spectral response, assuming $\kappa = 4.5$ cm^{-1}. The inferred peak reflectivity of this structure is approximately 90%, and bandwidth (at half-maximum) is approximately 15 GHz. Although the reflectivity of this structure is not high enough to compete with fiber Bragg gratings or other currently used high-performance WDM filters, this represents to our knowledge the first experimental demonstration of an integrated Bragg grating in an SOI ridge waveguide.

As discussed in Section 2.3.3, when extracting a modulated data signal from a noisy background, the optimal filter is not a flat bandpass filter but a matched filter. The length of the gratings considered here were chosen to yield a spectral response which is approximately matched to a 10 Gb/s on-off encoded optical signal.

4.3 Summary

We have described in this chapter the measurement and characterization of integrated waveguides, couplers, and Bragg gratings. Specifically, we investigated a series of directional couplers, fabricated in glass channel waveguides, and we demonstrated an integrated Bragg grating filter fabricated on a silicon-on-insulator ridge waveguide.

For the glass waveguide devices, we presented a comparison between conventional single-stage directional couplers and a more sophisticated cascaded coupler. We demonstrated that by connecting two dissimilar directional couplers with an intervening phase shift, one can obtain a power-splitting ratio that is insensitive to wavelength, polariza-

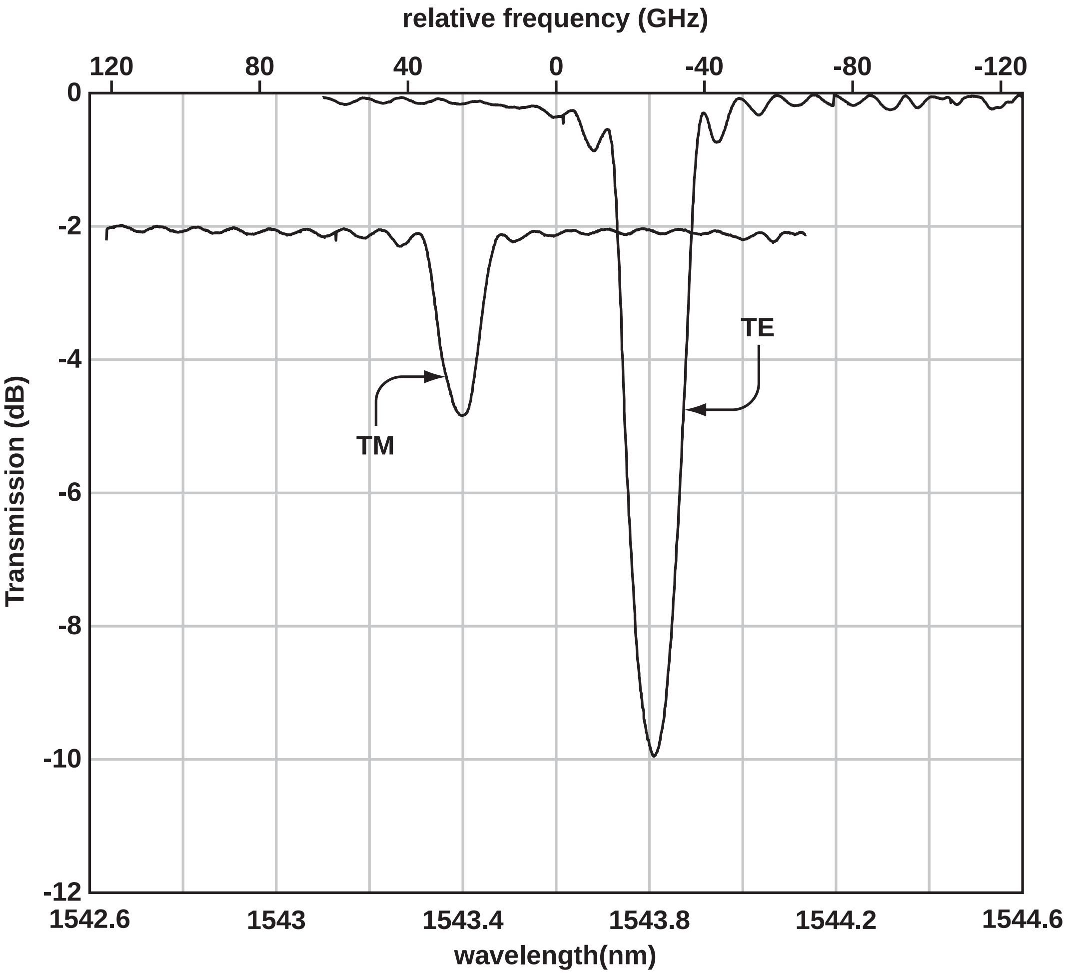

Figure 4.17: Transmission spectrum for the 4 mm-long Bragg grating. For comparison, both TE and TM polarizations are plotted. The TM data has been offset by -2 dB in order to more clearly differentiate the two spectra. The separation between the TE and TM peaks indicates that the birefringence for the structure is approximately 50 GHz, or 0.4 nm. The estimated grating strength for the TE polarization approximately 2× larger than for the TM polarization. (figs/4/tetmcompare4.eps)

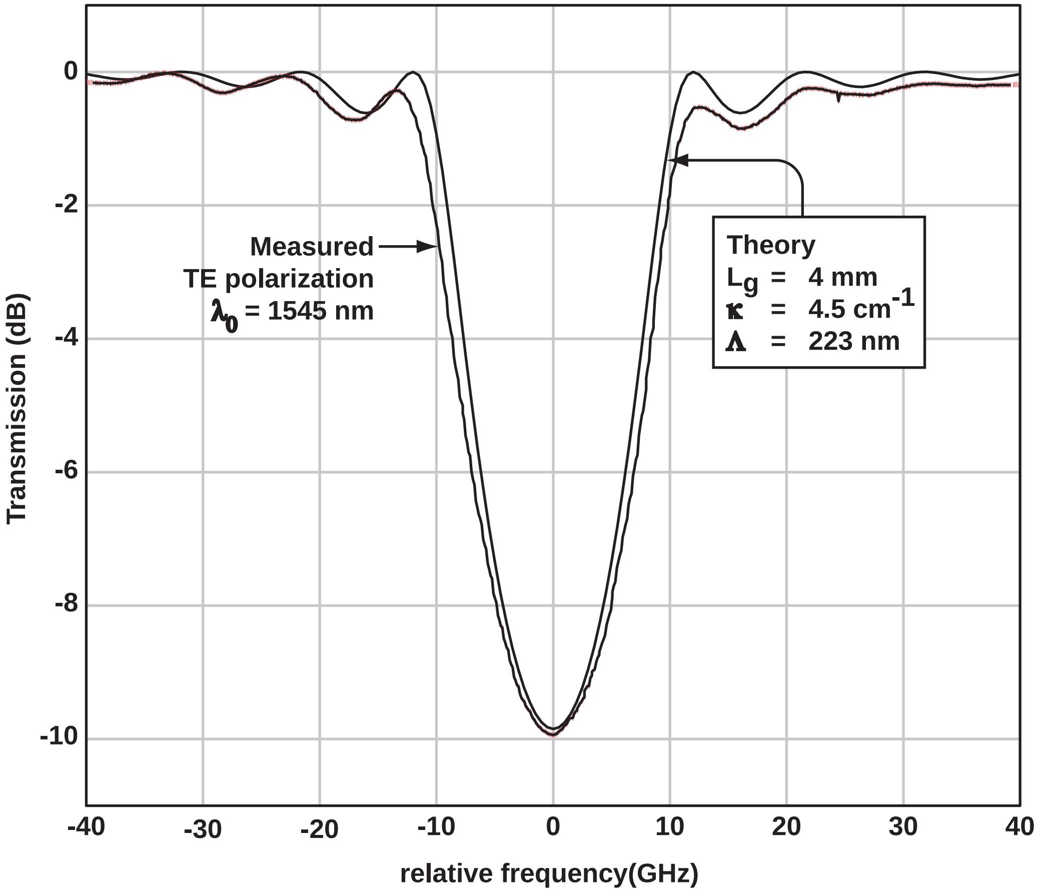

Figure 4.18: A comparison between the measured spectral response for the 4 mm-long Bragg grating, and that predicted from coupled mode theory. The grating strength used in the theoretical model was 4.5 cm^{-1}. (figs/4/grating-spectrum-fit-4.eps)

tion, and fabrication parameters. Others have recognized that wavelength insensitivity can be achieved using such a Mach-Zehnder structure, but their design technique involves numerical optimization based upon measured wavelength-dependence data [51]. This approach does not explicitly account for polarization dependence or other fabrication uncertainties, and it cannot be easily extended to other waveguide systems. Our measurements show, by direct comparison with a similarly fabricated directional coupler, that wavelength-insensitive and polarization-insensitive performance can be achieved using the relatively simple analytical design rules described in Section 2.2.5. The couplers which we constructed achieved close to 50 percent power splitting over the entire wavelength range from 1475-1590 nm, for both polarization states.

For the silicon-on-insulator devices, we measured and analyzed the spectral response of a set of integrated Bragg gratings. Our measured spectral response shows reasonable agreement with that predicted by the coupled mode theory described in Section 2.3. The filters achieve a bandwidth of 15 GHz, for a 4 mm-long grating, with a peak reflectivity of 90% for the TE polarization and 50% for the TM polarization. Furthermore, we described a simple model which explains the origin of the radiation loss which we observed on the blue side of the spectrum. To date, most of the integrated filters fabricated in SOI are resonator-based structures, based on sub-micron size silicon channel waveguides[144]. The Bragg grating structures demonstrated here operate on a completely different principle. Although it is somewhat misleading to compute a Q (quality factor) for these devices, we can point out that for a resonator-based device to have the same bandwidth as the Bragg grating filter described here, would require a Q of 13,000.

Chapter 5

Conclusions

5.1 Summary

This thesis has sought to provide a thorough description of my research in integrated Bragg grating filters. The work was grouped into three principal chapters describing, respectively, the design, fabrication and measurement of devices. This ordering, while logical, doesn't necessarily describe the actual or even the optimal sequence of events: most successful device development projects involve many cycles of iteration and prototyping.

The design of integrated optical devices begins with an electromagnetic analysis of dielectric structures. To this end, I developed and implemented a set of finite difference routines which accurately predict the electromagnetic modes of integrated waveguides. These routines are crucial not only in designing such structures, but in understanding their observed behavior. Once the electromagnetic modes are well-characterized, the behavior of more complicated systems consisting of multiple waveguides and gratings can be simulated using coupled mode theory. In coupled mode theory, Maxwells equations are replaced by a set of coupled linear differential equations in which the coupling coefficients are related to overlap integrals involving the electromagnetic modes of the waveguide. I describe how the coupled mode equations can be applied to directional couplers and Bragg gratings, including the case where the gratings are apodized or chirped. Finally, the solution to the coupled differential equations can be expressed in the form of a transfer matrix. Complicated aggregate devices can be easily modeled by simply multiplying their respective transfer matrices.

The fabrication of integrated photonic devices presents several unique problems which have not been adequately addressed by the microfabrication community. I implemented specialized software routines for generating the long, slowly curved waveguide patterns required in waveguides and directional couplers. I described a flexible fabrication methodology for patterning integrated waveguides, couplers, and Bragg gratings in two different material systems: integrated glass channel waveguides and silicon-on-insulator ridge waveguides. The relevant lithographic tools used to build these devices include conventional optical lithography, scanning-electron-beam lithography, interference lithography, and x-ray nanolithography. The goal of our fabrication efforts has been to bring together all of these techniques in an efficient and robust process.

Finally, the measurements of fabricated devices were presented, along with a description of the measurement techniques. For the glass waveguide system, I demonstrated an improved directional coupler that shows reduced sensitivity to wavelength and polarization. In the silicon-on-insulator system, I measured and characterized a series of integrated Bragg gratings. A simple physical model was presented to explain the anomalous loss in transmission which was observed in the blue side of the spectrum. The measured spectral response agrees well with theoretical predictions based upon coupled mode analysis.

5.2 Future Work

No one enjoys writing a document about the things they didn't do. Nevertheless, integrated optics is an exciting research field and like every exciting field there is an endless supply of challenging problems to consider. Although my work in this field has allowed me to design, fabricate and measure integrated photonic devices, there is much room for improvement, and there are many stones which have been left unturned.

The finite difference routines which I described are adequate for many devices, but there exist more advanced numerical techniques for simulating waveguides. In particular, the routines which I describe have a numerical accuracy (grid quantization error) which is proportional to the grid size $\Delta x \Delta y$. Most commercially available numerical mode solvers use a similar set of routines with the same accuracy order. The design and analysis of integrated photonic structures could benefit significantly from routines which converge more rapidly with decreasing grid size. For example, if the numerical accuracy were proportional to $\Delta x^2 \Delta y^2$, one wouldn't need to use as fine a grid size to accurately compute the mode structure[145, 146]. Another improvement which could be made would be to modify the existing program to implement the beam propagation method. This would enable one to numerically model structures like tapers, bends, and y-branches.

The coupled mode analysis for integrated Bragg gratings seems to accurately predict the coupling coefficient κ for the TE polarization, but overestimates the coupling coefficient for the TM polarization. This discrepancy might be caused by the fact that we neglected the longitudinal field component e_z when computing the grating strength overlap integral. More attention should be devoted to the problem of calculating the coupling coefficient κ, especially for high index contrast systems where there can be significant polarization dependence.

Interference lithography is an extremely powerful tool for building integrated Bragg gratings, but more work is needed to characterize and eliminate the phase distortion produced by these interference lithography systems. In particular, because the systems currently used in the NanoStructuresLab employ spherical waves, the gratings produced have an inherent chirp which limits the usable grating length in practical applications. In order to construct longer Bragg grating filters, it will be necessary to either use some other lithographic means, or renovate the interference lithography systems to eliminate this chirp.

For the glass waveguide devices, we have not yet solved the problem of maintaining the grating structure during the cladding overgrowth process. One potential solution would be to etch the grating into the lower cladding layer before depositing the core. Since the bottom SiO_2 base layer should have the highest glass transition temperature in the system, presumably this approach would not suffer from the same problems of reflow. As the integrated optics community becomes more interested in smaller, high-index-contrast waveguides, the problem of maintaining the structural fidelity of the structure during overgrowth will become even more important.

The silicon-on-insulator devices could be improved in several ways. First, we have not been careful in our initial experiments to minimize the birefringence of the devices. More experimental and theoretical investigation of these devices could lead to a design which has little or no birefringence. Another significant limitation of this material systems is the polarization-dependence in the grating strength κ. One potential solution to this problem would be to corrugate the sidewalls of the structure in addition to the top surface, in order to balance the TE and TM feedback. The devices which we measured had a mode size which was slightly too small to allow efficient butt-coupling to an optical fiber. While the structure could be made larger without sacrificing single-mode operation, it becomes hard to achieve a sufficiently high grating strength κ for a large waveguide. Again, there might be a way to improve the grating strength by corrugating the sidewalls and the other silicon regions outside of the ridge.

The coupling to leaky modes on the blue side of the spectrum is another effect which deserves further investigation. Although radiation loss is an undesirable effect for add/drop filters, it might be possible to exploit this effect for other applications such as gain equalization.

Having demonstrated integrated directional couplers and integrated Bragg gratings separately, the next task will be to assemble these two building blocks into an integrated Mach-Zehnder interferometer, in order to separate the reflected signal from the input signal. The ability to integrate directional couplers and gratings in this way without the need for any post-fabrication trimming or adjustment could prove to be the strongest selling point for integrated Bragg gratings filters.

Finally, much can be done to optimize the spectral response of a Bragg grating by intentionally introducing chirp and apodization. The theoretical framework presented in this thesis describes how to analyze and design such a structure, but the devices which we have built to date have not taken advantage of these degrees of freedom. One way to achieve apodization is to adjust the lateral extent of the grating, by making the grating only cover part of the waveguide rather than the entire waveguide. Another way would be to adjust the depth of the grating across the device. For Bragg gratings formed by physical corrugations, the process of apodization invariably changes the effective index of the waveguide. For this reason, apodization and chirp are inextricably tied together. There are two ways to construct a filter with pure apodization. One way is to modify the waveguide width in tandem with the apodization profile such that the effective index of the waveguide remains constant. Another approach would be to intentionally chirp the grating in addition to apodization in order to compensate for the change in effective index induced by the apodization.

I have mentioned here a few of the possible directions for future research, but the possible applications of integrated Bragg gratings are limited only by the imagination. It is my hope that this work provides enough theoretical and experimental information to encourage others to develop new creative ideas and applications.

Appendix A

Finite Difference Modesolver

Section 2.1.4 introduced the topic of finite difference methods by describing an approximate solution method for the scalar eigenvalue equation. The scalar eigenvalue equation is valid whenever the index contrast of the waveguide is small. The derivation presented in Section 2.1.4 assumes that the scalar field and its derivative are continuous at all points and the transverse planes. However, Maxwells equations clearly state that some of the transverse field components are discontinuous across abrupt dielectric interfaces. This appendix details how the finite difference equations can be modified to account for these predictable discontinuities in the electromagnetic fields. The reader should refer to Section 2.1.4, which defines the geometry of the finite difference mesh and introduces the notation used to label grid points.

A.1 Vector Finite Difference Method

First, consider the differential operator P_{xx} which operates on e_x in the full-vector eigenmode equation (Eq. 2.25):

$$P_{xx}e_x = \frac{\partial}{\partial x}\left[\frac{1}{n^2}\frac{\partial(n^2 e_x)}{\partial x}\right] + \frac{\partial^2 e_x}{\partial y^2} + n^2k^2 e_x \qquad . \tag{A.1}$$

Now suppose that we wish to approximate this differential operator with a finite difference scheme. As before, we shall use the subscripts P, N, S, E and W to label the point under consideration and its closest neighbors to the north, south, east and west respectively. If the refractive index is the same for all five of these points, we may simply use the scalar finite difference scheme described in Eq. 2.45. Also, because e_x is continuous across horizontal interfaces, the $\frac{\partial^2 e_x}{\partial y^2}$ term may be approximated with the simple three point difference of Eq. 2.42. However, when there is an index discontinuity between W and P or between P and E, the finite difference equations must be modified to account for the discontinuity in e_x. To accomplish this, we use a modified interpolation procedure in which we fit the points e_{xW}, e_{xP} and e_{xE} with a piecewise quadratic polynomial.

$$\hat{\phi}(x) = \begin{cases} A_W + Bx + Cx^2 & \text{in cell W} \\ A_P + Bx + Cx^2 & \text{in cell P} \\ A_E + Bx + Cx^2 & \text{in cell E} \end{cases} \tag{A.2}$$

The interpolating function $\hat{\phi}(x)$ is identical to Eq. 2.37 except that it allows for a possible discontinuity at the edge of each cell. We have further assumed that the first and second derivatives (or equivalently the polynomial coefficients B and C) are continuous across cell boundaries.

As before, we require that the interpolating function passes through the three points e_{xW}, e_{xP} and e_{xE}:

$$\begin{aligned} e_{xW} &= A_W - B\Delta x + C(\Delta x)^2 \\ e_{xP} &= A_P \\ e_{xE} &= A_E + B\Delta x + C(\Delta x)^2 \end{aligned} \tag{A.3}$$

But we now add additional constraints that describe the continuity of $n^2 e_x$ at $x = \pm\frac{\Delta x}{2}$:

$$\begin{aligned} n_W^2 \left\{ A_W - \frac{1}{2}B\Delta x + \frac{1}{4}C(\Delta x)^2 \right\} &= n_P^2 \left\{ A_P - \frac{1}{2}B\Delta x + \frac{1}{4}C(\Delta x)^2 \right\} \\ n_P^2 \left\{ A_P + \frac{1}{2}B\Delta x + \frac{1}{4}C(\Delta x)^2 \right\} &= n_E^2 \left\{ A_E + \frac{1}{2}B\Delta x + \frac{1}{4}C(\Delta x)^2 \right\} \end{aligned} \tag{A.4}$$

Eq. A.3 and Eq. A.4 together give five linear equations which can be solved for the five unknown polynomial coefficients. (Actually, B and C are the only two coefficients of interest because they describe respectively the first and second derivatives of e_x.) For brevity, we shall omit the algebraic steps, and simply summarize the results.

$$\begin{aligned}\frac{\partial\hat{\phi}}{\partial x} = B = \frac{1}{\Delta x}\Bigg\{ & \frac{-(n_W^2 n_P^2 + 3n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} e_{xW} + \\ & \frac{3(n_E^2 - n_W^2)n_P^2}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} e_{xP} + \\ & \frac{(n_E^2 n_P^2 + 3n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} e_{xE} \Bigg\}\end{aligned} \tag{A.5}$$

$$\begin{aligned}\frac{\partial^2\hat{\phi}}{\partial x^2} = 2C = \frac{1}{(\Delta x)^2}\Bigg\{ & \frac{4(n_W^2 n_P^2 + n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} e_{xW} + \\ & \frac{4(n_E^2 n_P^2 + n_W^2 n_P^2 + 2n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} e_{xP} + \\ & \frac{4(n_E^2 n_P^2 + n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} e_{xE} \Bigg\}\end{aligned} \tag{A.6}$$

Eq. 2.42 and Eq. A.6 can be combined to yield the finite difference representation of the operator P_{xx}.

$\hat{P}_{xx}$: (A.7)

0	$\frac{1}{(\Delta y)^2}$	0
$\frac{\alpha_W}{(\Delta x)^2}$	$n_P^2 k^2 - \frac{2}{(\Delta y)^2} - \frac{2\alpha_P}{(\Delta x)^2}$	$\frac{\alpha_E}{(\Delta x)^2}$
0	$\frac{1}{(\Delta y)^2}$	0

Where the constants α_W, α_P and α_E are dimensionless ratios defined by:

$$\alpha_W \equiv \frac{4(n_W^2 n_P^2 + n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} \tag{A.8}$$

$$\alpha_P \equiv \frac{2(2n_P^4 + n_E^2 n_P^2 + n_W^2 n_P^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} \tag{A.9}$$

$$\alpha_E \equiv \frac{4(n_E^2 n_P^2 + n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} \tag{A.10}$$

A similar analysis applied to e_y leads to the following finite difference representation of P_{yy}:

$$\hat{P}_{yy} : \begin{array}{|c|c|c|} \hline 0 & \dfrac{\gamma_N}{(\Delta y)^2} & 0 \\ \hline \dfrac{1}{(\Delta x)^2} & n_P^2 k^2 - \dfrac{2\gamma_P}{(\Delta y)^2} - \dfrac{2}{(\Delta x)^2} & \dfrac{1}{(\Delta x)^2} \\ \hline 0 & \dfrac{\gamma_S}{(\Delta y)^2} & 0 \\ \hline \end{array} \tag{A.11}$$

Where the constants γ_S and γ_N are dimensionless ratios defined by:

$$\gamma_S \equiv \frac{4(n_S^2 n_P^2 + n_N^2 n_S^2)}{n_P^4 + 2n_N^2 n_P^2 + 2n_S^2 n_P^2 + 3n_N^2 n_S^2} \tag{A.12}$$

$$\gamma_P \equiv \frac{2(2n_P^4 + n_N^2 n_P^2 + n_S^2 n_P^2)}{n_P^4 + 2n_N^2 n_P^2 + 2n_S^2 n_P^2 + 3n_N^2 n_S^2} \tag{A.13}$$

$$\gamma_N \equiv \frac{4(n_N^2 n_P^2 + n_N^2 n_S^2)}{n_P^4 + 2n_N^2 n_P^2 + 2n_S^2 n_P^2 + 3n_N^2 n_S^2} \tag{A.14}$$

The finite difference approximations for P_{xy} and P_{yx} are slightly more complicated because these differential operators involve mixed partial derivatives in both directions. However, the framework described above can be used to compute a mixed partial derivative by simply applying the finite difference approximations for $\frac{\partial}{\partial x}$ and $\frac{\partial}{\partial y}$ separately. We will illustrate this technique for the simple case where we wish to approximate $\frac{\partial^2 f}{\partial x \partial y}$ for a continuous function $f(x,y)$. To approximate $\frac{\partial^2 f}{\partial x \partial y}|_P$ we first use the three point finite difference approximations to obtain $\frac{\partial f}{\partial y}|_W$ and $\frac{\partial f}{\partial y}|_E$,

$$\frac{\partial f}{\partial y}|_W \simeq \frac{1}{2\Delta y}(f_{NW} - f_{SW}) \tag{A.15}$$

$$\frac{\partial f}{\partial y}|_E \simeq \frac{1}{2\Delta y}(f_{NE} - f_{SE}) \tag{A.16}$$

$$\tag{A.17}$$

Notice that we have now introduced the new subscripts NW, SW, NE and SE to label the points northwest, southwest, northeast and southeast respectively. Now, we simply apply a similar finite difference equation in horizontal direction, giving,

$$\frac{\partial^2 f}{\partial x \partial y}|_P \simeq \frac{1}{2\Delta x}\left\{\frac{\partial f}{\partial y}|_E - \frac{\partial f}{\partial y}|_W\right\} \tag{A.18}$$

$$\frac{\partial^2 f}{\partial x \partial y}|_P \simeq \frac{1}{4\Delta x \Delta y}\left\{f_{NE} - f_{SE} - f_{NW} + f_{SW}\right\} \tag{A.19}$$

Thus, the finite difference representation of the differential operator $\frac{\partial^2}{\partial x \partial y}$ is

$$\frac{\partial^2}{\partial x \partial y} : \begin{array}{|c|c|c|} \hline -\dfrac{1}{4\Delta x \Delta y} & 0 & +\dfrac{1}{4\Delta x \Delta y} \\ \hline 0 & 0 & 0 \\ \hline +\dfrac{1}{4\Delta x \Delta y} & 0 & -\dfrac{1}{4\Delta x \Delta y} \\ \hline \end{array} \tag{A.20}$$

It is easy to to verify that for continuous functions, the finite difference representation of $\frac{\partial^2}{\partial y \partial x}$ is the same as for $\frac{\partial^2}{\partial x \partial y}$. The discretized operators corresponding to P_{yx} can be derived by applying a similar analysis with the modified finite-difference model described in Eq. A.6. This yields the following finite difference representation of P_{yx}:

$$\hat{P}_{yx} : \begin{array}{|c|c|c|} \hline \dfrac{1-\alpha_{NW}}{4\Delta x \Delta y} & \dfrac{3(\alpha_{NE}-\alpha_{NW})}{4\Delta x \Delta y} & \dfrac{\alpha_{NE}-1}{4\Delta x \Delta y} \\ \hline 0 & 0 & 0 \\ \hline \dfrac{\alpha_{SW}-1}{4\Delta x \Delta y} & \dfrac{3(\alpha_{SW}-\alpha_{SE})}{4\Delta x \Delta y} & \dfrac{1-\alpha_{SE}}{4\Delta x \Delta y} \\ \hline \end{array} \tag{A.21}$$

Where the constants α_{NE}, α_{NW}, α_{SE} and α_{SW}, are again dimensionless ratios which become 1 in the absence of dielectric interfaces.

$$\alpha_{NE} \equiv \frac{2(n_{NE}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2)}{n_N^4 + 2n_{NE}^2 n_N^2 + 2n_{NW}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2} \tag{A.22}$$

$$\alpha_{NW} \equiv \frac{2(n_{NW}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2)}{n_N^4 + 2n_{NE}^2 n_N^2 + 2n_{NW}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2} \tag{A.23}$$

$$\alpha_{SE} \equiv \frac{2(n_{SE}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2)}{n_S^4 + 2n_{SE}^2 n_S^2 + 2n_{SW}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2} \tag{A.24}$$

$$\alpha_{SW} \equiv \frac{2(n_{SW}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2)}{n_S^4 + 2n_{SE}^2 n_S^2 + 2n_{SW}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2} \tag{A.25}$$

Likewise, the finite difference representation of P_{xy} is:

$$\hat{P}_{xy} : \begin{array}{|c|c|c|} \hline \dfrac{1-\gamma_{NW}}{4\Delta x \Delta y} & 0 & \dfrac{\gamma_{NE}-1}{4\Delta x \Delta y} \\ \hline \dfrac{3(\gamma_{SW}-\gamma_{NW})}{4\Delta x \Delta y} & 0 & \dfrac{3(\gamma_{NE}-\gamma_{SE})}{4\Delta x \Delta y} \\ \hline \dfrac{\gamma_{SW}-1}{4\Delta x \Delta y} & 0 & \dfrac{1-\gamma_{SE}}{4\Delta x \Delta y} \\ \hline \end{array} \tag{A.26}$$

where the four constants γ_{NE}, γ_{NW}, γ_{SE} and γ_{SW} are given by:

$$\gamma_{NE} \equiv \frac{2(n_{NE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2)}{n_E^4 + 2n_{NE}^2 n_E^2 + 2n_{SE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2} \tag{A.27}$$

$$\gamma_{SE} \equiv \frac{2(n_{SE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2)}{n_E^4 + 2n_{NE}^2 n_E^2 + 2n_{SE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2} \tag{A.28}$$

$$\gamma_{NW} \equiv \frac{2(n_{NW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2)}{n_W^4 + 2n_{NW}^2 n_W^2 + 2n_{SW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2} \tag{A.29}$$

$$\gamma_{SW} \equiv \frac{2(n_{SW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2)}{n_W^4 + 2n_{NW}^2 n_W^2 + 2n_{SW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2} \tag{A.30}$$

Examination of Eq. A.21 and Eq. A.26 reveals that the matrix elements of the finite difference operators $\hat{P}_{xy}$ and $\hat{P}_{xy}$ vanish except for points which are located next to dielectric interfaces. Thus, it is the presence of dielectric interfaces which couples the two transverse field components.

For reference, I have summarized the finite difference implementation of the full-vector eigenmode problem in Tables A.1 and A.2.

To summarize, we have described a technique for developing finite difference approximations which can be applied even in the presence of abrupt dielectric interfaces where the field is discontinuous. The technique described here involves fitting a piecewise-quadratic polynomial to three adjacent grid points. Other authors have developed similar techniques for accounting for the field discontinuity. Most notably, Stern [21] describes a technique for modeling field discontinuities using only two adjacent grid points. For simple one-dimensional problems which can be solved exactly, I have found that the discretization presented here yields a solution which is 3-4× more accurate for the same grid size.

A.2 Finite Difference Boundary Conditions

Thus far, we have concentrated on developing finite difference representations of partial differential equations, without regard to the edges of the computation window. The finite difference approximations given in the previous section can be applied without modification to all internal points in the computation window. But for points which lie on the edge of the computation window, boundary conditions must be considered. We will describe here three different boundary conditions: symmetric, antisymmetric, and absorbing.

By "absorbing", we mean that the field is assumed to be zero at grid points immediately

$$P_{xx}e_x \equiv \frac{\partial}{\partial x}\left[\frac{1}{n^2}\frac{\partial(n^2 e_x)}{\partial x}\right] + \frac{\partial^2 e_y}{\partial y^2} + n^2k^2 e_x$$

0	$\frac{1}{(\Delta y)^2}$	0
$\frac{\alpha_W}{(\Delta x)^2}$	$n_P^2k^2 - \frac{2}{(\Delta y)^2} - \frac{2\alpha_P}{(\Delta x)^2}$	$\frac{\alpha_E}{(\Delta x)^2}$
0	$\frac{1}{(\Delta y)^2}$	0

$$P_{xy}e_y \equiv \frac{\partial}{\partial x}\left[\frac{1}{n^2}\frac{\partial n^2 e_y}{\partial y}\right] - \frac{\partial^2 e_y}{\partial x \partial y}$$

$\frac{1-\gamma_{NW}}{4\Delta x\Delta y}$	0	$\frac{\gamma_{NE}-1}{4\Delta x\Delta y}$
$\frac{3(\gamma_{SW}-\gamma_{NW})}{4\Delta x\Delta y}$	0	$\frac{3(\gamma_{NE}-\gamma_{SE})}{4\Delta x\Delta y}$
$\frac{\gamma_{SW}-1}{4\Delta x\Delta y}$	0	$\frac{1-\gamma_{SE}}{4\Delta x\Delta y}$

$$P_{yx}e_x \equiv \frac{\partial}{\partial y}\left[\frac{1}{n^2}\frac{\partial n^2 e_x}{\partial x}\right] - \frac{\partial^2 e_x}{\partial y \partial x}$$

$\frac{1-\alpha_{NW}}{4\Delta x\Delta y}$	$\frac{3(\alpha_{NE}-\alpha_{NW})}{4\Delta x\Delta y}$	$\frac{\alpha_{NE}-1}{4\Delta x\Delta y}$
0	0	0
$\frac{\alpha_{SW}-1}{4\Delta x\Delta y}$	$\frac{3(\alpha_{SW}-\alpha_{SE})}{4\Delta x\Delta y}$	$\frac{1-\alpha_{SE}}{4\Delta x\Delta y}$

$$P_{yy}e_y \equiv \frac{\partial^2 e_y}{\partial x^2} + \frac{\partial}{\partial y}\left[\frac{1}{n^2}\frac{\partial(n^2 e_y)}{\partial y}\right] + n^2k^2 e_y$$

0	$\frac{\gamma_N}{(\Delta y)^2}$	0
$\frac{1}{(\Delta x)^2}$	$n_P^2k^2 - \frac{2\gamma_P}{(\Delta y)^2} - \frac{2}{(\Delta x)^2}$	$\frac{1}{(\Delta x)^2}$
0	$\frac{\gamma_S}{(\Delta y)^2}$	0

Table A.1: A summary of full-vector finite difference equations for the transverse $\mathbf{e}$ fields. Each of the finite difference operators is represented by a 3 by 3 table which describes the finite difference coefficients for the point under consideration and the neighboring eight points. For example, the upper left-hand element of the P_{yx} table represents the coefficient which multiplies e_{xNW} in the finite difference approximation of $P_{yx}e_x$. The constants $\{\alpha_\nu\}$ and $\{\gamma_\nu\}$ are dimensionless ratios, tabulated in Table A.2. Each of these constants evaluates to 1 in the absence of dielectric interfaces.

$$
\begin{aligned}
\alpha_W &\equiv \frac{4(n_W^2 n_P^2 + n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} & \qquad \gamma_S &\equiv \frac{4(n_S^2 n_P^2 + n_N^2 n_S^2)}{n_P^4 + 2n_N^2 n_P^2 + 2n_S^2 n_P^2 + 3n_N^2 n_S^2} \\
\alpha_P &\equiv \frac{2(2n_P^4 + n_E^2 n_P^2 + n_W^2 n_P^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} & \qquad \gamma_P &\equiv \frac{2(2n_P^4 + n_N^2 n_P^2 + n_S^2 n_P^2)}{n_P^4 + 2n_N^2 n_P^2 + 2n_S^2 n_P^2 + 3n_N^2 n_S^2} \\
\alpha_E &\equiv \frac{4(n_E^2 n_P^2 + n_E^2 n_W^2)}{n_P^4 + 2n_E^2 n_P^2 + 2n_W^2 n_P^2 + 3n_E^2 n_W^2} & \qquad \gamma_N &\equiv \frac{4(n_N^2 n_P^2 + n_N^2 n_S^2)}{n_P^4 + 2n_N^2 n_P^2 + 2n_S^2 n_P^2 + 3n_N^2 n_S^2} \\
\alpha_{NE} &\equiv \frac{2(n_{NE}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2)}{n_N^4 + 2n_{NE}^2 n_N^2 + 2n_{NW}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2} & \qquad \gamma_{NE} &\equiv \frac{2(n_{NE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2)}{n_E^4 + 2n_{NE}^2 n_E^2 + 2n_{SE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2} \\
\alpha_{NW} &\equiv \frac{2(n_{NW}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2)}{n_N^4 + 2n_{NE}^2 n_N^2 + 2n_{NW}^2 n_N^2 + 3n_{NE}^2 n_{NW}^2} & \qquad \gamma_{NW} &\equiv \frac{2(n_{NW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2)}{n_W^4 + 2n_{NW}^2 n_W^2 + 2n_{SW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2} \\
\alpha_{SE} &\equiv \frac{2(n_{SE}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2)}{n_S^4 + 2n_{SE}^2 n_S^2 + 2n_{SW}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2} & \qquad \gamma_{SE} &\equiv \frac{2(n_{SE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2)}{n_E^4 + 2n_{NE}^2 n_E^2 + 2n_{SE}^2 n_E^2 + 3n_{NE}^2 n_{SE}^2} \\
\alpha_{SW} &\equiv \frac{2(n_{SW}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2)}{n_S^4 + 2n_{SE}^2 n_S^2 + 2n_{SW}^2 n_S^2 + 3n_{SE}^2 n_{SW}^2} & \qquad \gamma_{SW} &\equiv \frac{2(n_{SW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2)}{n_W^4 + 2n_{NW}^2 n_W^2 + 2n_{SW}^2 n_W^2 + 3n_{NW}^2 n_{SW}^2}
\end{aligned}
$$

Table A.2: A summary of the dimensionless constants used in the finite difference scheme given in Table A.1. Notice that each of the constants defined here becomes 1 when there are no index discontinuities.

outside of the computation window. In this case, the finite difference approximations only require a slight modification: we simply omit those matrix elements which would otherwise refer to points outside of the computation window. For example, suppose we wish to approximate the scalar eigenvalue equation for a point P located on the leftmost edge of the computation window with absorbing boundary conditions. Eq. 2.45 describes the appropriate finite-difference discretization for internal grid points. The modified finite-difference discretization for a point on the leftmost absorbing boundary would be:

0	$\frac{1}{(\Delta y)^2}$	0
0	$n_P^2 k^2 - \frac{2}{(\Delta y)^2} - \frac{2}{(\Delta x)^2}$	$\frac{1}{(\Delta x)^2}$
0	$\frac{1}{(\Delta y)^2}$	0

(A.31)

Now consider the case where instead we assume that the field is symmetric at the left boundary of the computation window. For a point P located on the left boundary, ϕ_W can be related to ϕ_P by:

$$\phi_W = \phi_P \tag{A.32}$$

When this relation is combined with the previously developed finite difference representations of Eq. 2.45, the effect is to simply add the ϕ_W coefficient to the ϕ_P coefficient as illustrated below:

0	$\frac{1}{(\Delta y)^2}$	0
0	$n_P^2 k^2 - \frac{2}{(\Delta y)^2} - \frac{1}{(\Delta x)^2}$	$\frac{1}{(\Delta x)^2}$
0	$\frac{1}{(\Delta y)^2}$	0

(A.33)

Likewise, an antisymmetric boundary condition on the left edge can be accommodated by simply subtracting the ϕ_W coefficient from the ϕ_P coefficient.

We have described how the three different boundary conditions can be treated in the finite difference method. Although we have used the scalar finite difference equation for illustration purposes, these techniques work just as well for the full-vector finite difference equations. In the full-vector treatment, it is important to realize that the two transverse electric field components have opposite symmetry. That is, if e_x is symmetric at the left boundary, e_y should be antisymmetric and vice versa. Also, even though we have focused on the left / west boundary, these techniques can be easily applied to the other three edges of the computation window.

The boundary conditions described here assume that the fields are either symmetric, antisymmetric, or zero immediately outside of the computation window. In reality, the electromagnetic fields usually fall exponentially at points outside of the core. Therefore, the assumption that the fields reach zero at the edges of the computation window is merely a convenient approximation. In practice, one must ensure that the computation window is large enough that the fields are sufficiently small at the edges. Another solution is to modify the boundary conditions to account for the exponential falloff of the mode away from the core. However, one must know how rapidly the fields are decaying in order to implement this boundary condition. For this reason, the exponentially decaying boundary condition is usually implemented iteratively [147, 19]: first, one solves for the modes with absorbing boundary conditions, then the boundary conditions are modified according to the computed decay rate near the edges of the computation window, and the solution is then refined using the modified boundary conditions. After several iterations, the solution converges to an exponentially decaying mode. One limitation of this procedure, is that the exponential falloff rates can be different for each mode of the structure. Therefore, this technique cannot be used to calculate several eigenmodes simultaneously.

A.3 Finite Difference Equations for Transverse H Fields

From the preceding discussion, it appears that the finite difference equations would be simple were it not for the discontinuity of the field components across dielectric interfaces. Based upon this observation, one may ask why we do not formulate the eigenmodes equation in terms of the transverse magnetic field $\mathbf{h}_t$, rather than the transverse electric field $\mathbf{e}_t$. After all, the magnetic field should be continuous across dielectric interfaces even if the electric field is not.

In fact, it is possible to express the eigenmode equations in terms of the transverse magnetic fields, which are continuous everywhere. However, the slope of the magnetic field is not necessarily continuous across dielectric interfaces. Because of this limitation, the simple scalar finite difference equations cannot be used with the magnetic field. As with the transverse electric field, a modified set of finite difference equations must be derived for the transverse magnetic field. The analysis is almost identical to that presented above, with the exception that the quadratic interpolating function has a discontinuous slope rather than a discontinuous offset. Rather than presenting the complete analysis (which is quite similar to the one presented above), the resulting finite difference discretization for transverse magnetic fields is summarized in Table A.3 and Table A.4.

One argument in favor of using magnetic fields is that the magnetic field should be well behaved at dielectric sharp corners. By contrast, consider the continuity of the electric field component e_x at a corner. e_x should be continuous across the horizontal interface, but discontinuous across the vertical interface; these two requirements seem to be at odds in the vicinity of a sharp corner. Rigorous analyses show that the electric field has a singularity at such sharp corners [148]. In practice, the corner is never perfectly sharp, and this singularity can usually be neglected.

A.4 Semivectorial Finite Difference Method

In many cases when the full vector finite difference equations are used, one of the transverse field components is found to be orders of magnitude smaller than the orthogonal component. In these cases, the finite difference method can be simplified by neglecting the small transverse field component and instead solving the finite difference equations for the one remaining field component. This approach is equivalent to neglecting P_{xy}, P_{yx}, and either P_{xx} or P_{yy} in the full vector treatment described earlier. With this simplification, the finite difference problem requires the same storage and computational time as the scalar problem. However, some of the vectorial properties of the electromagnetic fields are retained in the solution through the boundary conditions which are applied in the finite difference discretization. For this reason, this approach is called the semivectorial finite difference method.

The following three Matlab scripts implement a finite-difference semivectorial mode solver. The first program, `rib.m` is used to construct a discretized matrix representation of the refractive index profile. The second program, `svbuildmtx.m`, constructs the finite difference matrix for the semivectorial mode solver, and the last program, `sveigenmodes` simply calculates and formats the eigenmodes of the finite difference matrix. There are a few additional routines (not presented here) which can be used to plot the calculated eigenmodes.

$$P_{xx}h_x \equiv \frac{\partial^2 h_x}{\partial x^2} + n^2\left[\frac{\partial}{\partial y}\left(\frac{1}{n^2}\frac{\partial h_x}{\partial y}\right)\right] + n^2k^2h_x$$

0	$\frac{\sigma_N}{(\Delta y)^2}$	0
$\frac{1}{(\Delta x)^2}$	$n_P^2k^2 - \frac{2}{(\Delta y)^2} - \frac{2\sigma_P}{(\Delta x)^2}$	$\frac{1}{(\Delta x)^2}$
0	$\frac{\sigma_S}{(\Delta y)^2}$	0

$$P_{xy}h_y \equiv \frac{\partial^2 h_y}{\partial y \partial x} - n^2\left[\frac{\partial}{\partial y}\left(\frac{1}{n^2}\frac{\partial h_y}{\partial x}\right)\right]$$

$\frac{\zeta_{NW}-1}{4\Delta x\Delta y}$	$\frac{\zeta_{NE}-\zeta_{NW}}{4\Delta x\Delta y}$	$\frac{1-\zeta_{NE}}{4\Delta x\Delta y}$
0	0	0
$\frac{1-\zeta_{SW}}{4\Delta x\Delta y}$	$\frac{\zeta_{SW}-\zeta_{SE}}{4\Delta x\Delta y}$	$\frac{\zeta_{SE}-1}{4\Delta x\Delta y}$

$$P_{yx}h_x \equiv \frac{\partial^2 h_x}{\partial x \partial y} - n^2\left[\frac{\partial}{\partial x}\left(\frac{1}{n^2}\frac{\partial h_x}{\partial y}\right)\right]$$

$\frac{\sigma_{NW}-1}{4\Delta x\Delta y}$	0	$\frac{1-\sigma_{NE}}{4\Delta x\Delta y}$
$\frac{\sigma_{SW}-\sigma_{NW}}{4\Delta x\Delta y}$	0	$\frac{\sigma_{NE}-\sigma_{SE}}{4\Delta x\Delta y}$
$\frac{1-\sigma_{SW}}{4\Delta x\Delta y}$	0	$\frac{\sigma_{SE}-1}{4\Delta x\Delta y}$

$$P_{yy}h_y \equiv \frac{\partial^2 h_y}{\partial y^2} + n^2\left[\frac{\partial}{\partial x}\left(\frac{1}{n^2}\frac{\partial h_y}{\partial x}\right)\right] + n^2k^2h_y$$

0	$\frac{1}{(\Delta y)^2}$	0
$\frac{\zeta_W}{(\Delta x)^2}$	$n_P^2k^2 - \frac{2\zeta_P}{(\Delta y)^2} - \frac{2}{(\Delta x)^2}$	$\frac{\zeta_E}{(\Delta x)^2}$
0	$\frac{1}{(\Delta y)^2}$	0

Table A.3: A summary of full-vector finite difference equations for the transverse $\mathbf{h}$ fields. Each of the finite difference operators is represented by a 3 by 3 table which describes the finite difference coefficients for the point under consideration and the neighboring eight points. For example, the upper left-hand element of the P_{yx} table represents the coefficient which multiplies h_{xNW} in the finite difference approximation of $P_{yx}h_x$. The constants $\{\sigma_\nu\}$ and $\{\zeta_\nu\}$ are dimensionless ratios, tabulated in Table A.4. Each of these constants evaluates to 1 in the absence of dielectric interfaces.

$$
\begin{aligned}
\zeta_W &\equiv \frac{2n_P^2}{n_P^2+n_W^2} & \sigma_S &\equiv \frac{2n_P^2}{n_P^2+n_S^2} \\
\zeta_E &\equiv \frac{2n_P^2}{n_P^2+n_E^2} & \sigma_N &\equiv \frac{2n_P^2}{n_P^2+n_N^2} \\
\zeta_P &\equiv \tfrac{1}{2}(\zeta_E+\zeta_W) & \sigma_P &\equiv \tfrac{1}{2}(\sigma_N+\sigma_S) \\
\zeta_{NE} &\equiv \frac{2n_N^2}{n_N^2+n_{NE}^2} & \sigma_{NE} &\equiv \frac{2n_E^2}{n_E^2+n_{NE}^2} \\
\zeta_{NW} &\equiv \frac{2n_N^2}{n_N^2+n_{NW}^2} & \sigma_{NW} &\equiv \frac{2n_W^2}{n_W^2+n_{NW}^2} \\
\zeta_{SE} &\equiv \frac{2n_S^2}{n_S^2+n_{SE}^2} & \sigma_{SE} &\equiv \frac{2n_E^2}{n_E^2+n_{SE}^2} \\
\zeta_{SW} &\equiv \frac{2n_S^2}{n_S^2+n_{SW}^2} & \sigma_{SW} &\equiv \frac{2n_W^2}{n_W^2+n_{SW}^2}
\end{aligned}
$$

Table A.4: A summary of the dimensionless constants used in the finite difference scheme given in Table A.3. Notice that each of the constants defined here becomes 1 when there are no index discontinuities.

File: rib.m

This Matlab program constructs an index mesh for a generalized rib waveguide. The script is presented here only to give an example of how to construct the index mesh for a finite difference problem. This routine may be customized or replaced to accommodate any desired waveguide geometry; the only required output parameters are the two-dimensional matrix `eps`, and the quantities specifying the grid size `dx` and `dy`.

```
function [x,y,xc,yc,nx,ny,eps] = rib(n1,n2,n3,h1,h2,h3,rh,rw,side,dx,dy);

% This function creates an index mesh for the finite-difference
% mode solver.  The function will accommodate a generalized three
% layer rib waveguide structure.  (Note: channel waveguides can
% also be treated by selecting the parameters appropriately.)
%
% USAGE:
%
% [x,y,xc,yc,nx,ny,eps] = rib(n1,n2,n3,h1,h2,h3,rh,rw,side,dx,dy)
%
% INPUT
%
% n1 - index of refraction for substrate
% n2 - index of refraction for core
% n3 - index of refraction for top cladding
% h1 - height of substrate region
% h2 - height of core region
% h3 - height of top cladding region
% rh - height of rib ( <= h2)
% rw - half-width of rib
% side - excess space to the right of waveguide
% dx - horizontal grid spacing
% dy - vertical grid spacing
%
% OUTPUT
%
% x,y - vectors specifying mesh coordinates
% xc,yc - vectors specifying grid-center coordinates
% nx,ny - size of index mesh
% eps - index mesh (n^2)

ih1 = round (h1/dy);
ih2 = round (h2/dy);
ih3 = round (h3/dy);
irh = round (rh/dy);
irw = round (rw/dx);
iside = round (side/dx);

nx = irw+iside+1;
ny = ih1+ih2+ih3+1;
```

```
xc = (1:(nx-1))'*dx - dx/2;
yc = (1:(ny-1))*dy - dy/2;
x = (0:(nx-1))'*dx;
y = (0:(ny-1))*dy;

eps = zeros(nx-1,ny-1);

iy = 1;

for i = 1:ih1,
   eps(:,iy) = n1^2*ones(nx-1,1);
   iy = iy+1;
end

for i = 1:(ih2-irh),
   eps(:,iy) = n2^2*ones(nx-1,1);
   iy = iy+1;
end

for i = 1:irh,
   eps(:,iy) = [[n2^2*ones(irw,1)];[n3^2*ones(iside,1)]];
   iy = iy+1;
end

for i = 1:ih3,
   eps(:,iy) = n3^2*ones(nx-1,1);
   iy = iy+1;
end

nx = length(xc);
ny = length(yc);
```

File: svbuildmtx.m

This Matlab function lies at the heart of the semivectorial modesolver. Based upon the specified index mesh (obtained, for example, from `rib.m`), the program builds the sparse matrix **A** representing the finite-difference equations. Notice that the matrix elements are calculated in vectorized form rather than with nested "for" loops – this greatly increases the speed of the routine by making use of Matlab's optimized vector algebra routines. The matrix elements are first calculated in lines 58-103, ignorning boundary conditions at the edges of the computation window. The sparse matrix is actually assembled in line 120, and boundary conditions are implemented in lines 124-164. We have also developed a similar routine for building the matrix associated with the full-vector problem. However, the semivectorial routines are usually sufficient for most problems, and the extension to full-vector routines is straightforward.

```
function A = svbuildmtx (lambda, dx, dy, eps, boundary, field);

% This function constructs the finite difference matrix for
% the semivectorial mode solver.  Using the index mesh specified in
% the matrix eps (= nˆ2) and the boundary conditions specified in
% ’boundary’, this function generates a sparse matrix A
% representing the differential operator for the eigenmodes.
%
% USAGE:
%
% A = buildmtx (lambda, dx, dy, eps, boundary, field);
%
% INPUT:
%
% lambda - optical wavelength
% dx - horizontal grid spacing
% dy - vertical grid spacing
% eps - index mesh (= nˆ2(x,y))
% boundary - 4 letter string specifying boundary conditions to be
% applied at the edges of the computation window.
%   boundary(1) = North boundary condition
%   boundary(2) = South boundary condition
%   boundary(3) = East boundary condition
%   boundary(4) = West boundary condition
% The following boundary conditions are supported:
%   ’A’ - field is antisymmetric
%   ’S’ - field is symmetric
%   ’0’ - field is zero immediately outside of the
%         boundary.
% field - can be ’EX’, ’EY’, ’HX’, ’HY’, or ’scalar’
%
% OUTPUT:
%
% A - sparse matrix representing differential operator for the
%     eigenvalue problem.

[nx,ny] = size(eps);

% now we pad eps on all sides by one grid point
eps = [eps(:,1),eps,eps(:,ny)];
eps = [eps(1,:); eps ; eps(nx,:)];

% compute free-space wavevector
k = 2*pi/lambda;

en = ones(1,nx*ny);
es = ones(1,nx*ny);
ee = ones(1,nx*ny);
ew = ones(1,nx*ny);
ep = ones(1,nx*ny);

en(:)  = eps(2:nx+1 ,3:ny+2);
es(:)  = eps(2:nx+1 ,1:ny);
```

```
ee(:)  = eps(3:nx+2 ,2:ny+1);
ew(:)  = eps(1:nx   ,2:ny+1);
ep(:)  = eps(2:nx+1 ,2:ny+1);

switch lower(field)
  case 'ex'
    alphaw = 4*(1 + ep./ee)./(3 + 2*ep./ee + 2*ep./ew + ep.^2./(ew.*ee));
    alphap = 2*(ep./ee + ep./ew + 2*ep.^2./(ee.*ew))./ ...
        (3 + 2*ep./ee + 2*ep./ew + ep.^2./(ew.*ee));
    alphae = 4*(1 + ep./ew)./(3 + 2*ep./ee + 2*ep./ew + ep.^2./(ew.*ee));

    an = ones(1,nx*ny)/dy^2;
    as = ones(1,nx*ny)/dy^2;
    ae = alphae/dx^2;
    aw = alphaw/dx^2;
    ap = ep*k^2 - 2*alphap/dx^2 - 2*ones(1,nx*ny)/dy^2;

  case 'ey'
    gammas = 4*(1 + ep./en)./(3 + 2*ep./en + 2*ep./es + ep.^2./(en.*es));
    gammap = 2*(ep./en + ep./es + 2*ep.^2./(en.*es))./ ...
        (3 + 2*ep./en + 2*ep./es + ep.^2./(en.*es));
    gamman = 4*(1 + ep./es)./(3 + 2*ep./en + 2*ep./es + ep.^2./(en.*es));

    an = gamman/dy^2;
    as = gammas/dy^2;
    ae = ones(1,nx*ny)/dx^2;
    aw = ones(1,nx*ny)/dx^2;
    ap = ep*k^2 - 2*gammap/dy^2 - 2*ones(1,nx*ny)/dx^2;

  case 'hx'
    an = 2*ep./((ep+en)*dy^2);
    as = 2*ep./((ep+es)*dy^2);
    ae = ones(1,nx*ny)/dx^2;
    aw = ones(1,nx*ny)/dx^2;
    ap = ep*k^2 - an - as - 2*ones(1,nx*ny)/dx^2;

  case 'hy'
    an = ones(1,nx*ny)/dy^2;
    as = ones(1,nx*ny)/dy^2;
    ae = 2*ep./((ep+ee)*dx^2);
    aw = 2*ep./((ep+ew)*dx^2);
    ap = ep*k^2 - ae - aw - 2*ones(1,nx*ny)/dy^2;

  case 'scalar'
    an = ones(1,nx*ny)/dy^2;
    as = ones(1,nx*ny)/dy^2;
    ae = ones(1,nx*ny)/dx^2;
    aw = ones(1,nx*ny)/dx^2;
    ap = ep*k^2 - 2*ones(1,nx*ny)/dx^2 - 2*ones(1,nx*ny)/dy^2;
end

ii = zeros(nx,ny);
ii(:) = (1:nx*ny);
```

```
iall = zeros(1,nx*ny);
is = zeros(1,nx*(ny-1));
in = zeros(1,nx*(ny-1));
ie = zeros(1,(nx-1)*ny);
iw = zeros(1,(nx-1)*ny);

iall(:) = ii;
is(:) = ii(1:nx,1:(ny-1));
in(:) = ii(1:nx,2:ny);
iw(:) = ii(1:(nx-1),1:ny);
ie(:) = ii(2:nx,1:ny);

A = sparse ([iall,iw,ie,is,in], ...
    [iall,ie,iw,in,is], ...
    [ap(iall),ae(iw),aw(ie),an(is),as(in)]);

% now we must account for the boundary conditions.

% north boundary
ib = zeros(1,nx); b = boundary(1); ib(:) = ii(1:nx,ny);
if     (b == ’S’) sign = +1;
elseif (b == ’A’) sign = -1;
elseif (b == ’0’) sign =  0;
end
for i = ib,
  A(i,i) = A(i,i) + sign*an(i);
end

% south boundary
ib = zeros(1,nx); b = boundary(2); ib(:) = ii(1:nx,1);
if     (b == ’S’) sign = +1;
elseif (b == ’A’) sign = -1;
elseif (b == ’0’) sign =  0;
end
for i = ib,
  A(i,i) = A(i,i) + sign*as(i);
end

% east boundary
ib = zeros(1,ny); b = boundary(3); ib(:) = ii(nx,1:ny);
if     (b == ’S’) sign = +1;
elseif (b == ’A’) sign = -1;
elseif (b == ’0’) sign =  0;
end
for i = ib,
  A(i,i) = A(i,i) + sign*ae(i);
end

% west boundary
ib = zeros(1,ny); b = boundary(4); ib(:) = ii(1,1:ny);
if     (b == ’S’) sign = +1;
elseif (b == ’A’) sign = -1;
```

```
elseif (b == '0') sign =  0;
end
for i = ib,
  A(i,i) = A(i,i) + sign*aw(i);
end
```

File: sveigenmodes.m

This last routine simply calls the Matlab function `eigs` (in line 33) to compute the eigenvalues and associated eigenvectors of the finite difference matrix $\mathbf{A}$. The remainder of the script simply formats the $n_x n_y \times 1$ eigenvectors into two-dimensional arrays of size $n_x \times n_y$ so that they can easily be plotted.

```
 1   function [phi,neff] = sveigenmodes(A,guess,nmodes,lambda,nx,ny);

     % This function calculates and formats the eigenmodes for the
     % semivectorial finite difference mode solver.  It uses the MATLAB
 5   % function eigs to calculate a few eigenvalues of the matrix A.
     %
     % USAGE:
     %
     % [phi,neff] = sveigenmodes(A,guess,nmodes,lambda,nx,ny)
10   %
     % INPUT:
     %
     % A - sparse matrix containing the finite-difference
     %     representation of the differential operator for the
15   %     full-vector mode solver, generated by svbuildmtx.m
     % guess - scalar shift to apply when calculating the eigenvalues.
     %     This routine will return the eigenpairs which are closest
     %     to this guess in magnitude
     % nmodes - number of modes to calculate
20   % lambda - wavelength
     % nx,ny - dimensions of finite difference mesh
     %
     % OUTPUT:
     %
25   % phi - three-dimensional vector containing the field for each
     %     calculated mode, phi(:,:,k) = kˆth eigenmode
     % neff - vector of modal effective indices,
     %     neff(k) = effective index of kˆth eigenmode

30   shift = (2*pi*guess/lambda)ˆ2;
     options.tol = 1e-8;
     options.disp = 0;                       % suppress output
     [v,d] = eigs(A,speye(size(A)),nmodes,shift,options);
     neff = lambda*sqrt(diag(d))/(2*pi);
35
     phi = zeros(nx,ny,nmodes);
```

```
temp = zeros(nx,ny);

for k = 1:nmodes;
  temp(:) = v(:,k)/max(abs(v(:,k)));
  phi(:,:,k) = temp;
end;
```

Appendix B

Waveguide Pattern Generation Software

In this appendix, we provide an example which illustrates how to use the waveguide pattern generation software. The file listed below contains a description of a series of coupler devices, of the type described in Section 2.2.5. When the pound character # is encountered in the input file, the remainder of the line is treated as a comment. Dimensions of length can be specified in almost any of the standard SI units, and angles may be specified in either degrees or radians. The settings for each device type are “sticky”, meaning that unless they are changed they will automatically apply to subsequent devices of the same type. This makes it easy to construct an array of similar, but slightly different devices. Figure B.1 illustrates the corresponding patterns generated by the layout program.

```
kicfilename = "testchip.kic";    # kic output filename
epsfilename = "testchip.eps";    # eps output filename
rptfilename = "testchip.rpt";    # report filename

epsScaleX = 5;                   # x scale factor for eps output
epsScaleY = 5;                   # y scale factor for eps output

lambda = 0.1 um;                 # grid size / address unit

DeviceSeparation = 400 um;       # optional device separation

device CouplerPair {
    Name = "mzc1";               # Name of device
    Ltotal = 20 mm;              # total length of device
    DeviceLayer = "WG";          # layer on which device resides
    DrawText = false;            # do not draw text label
    x0 = 10 um;                  # starting position (left)
```

```
    y0 = 0;                      # starting position (bottom)
    Width = 6.0 um;              # waveguide width
    RadiusA = 20 mm;             # radius of curvature A
    RadiusB = 20 mm;             # radius of curvature B
    RadiusC = 20 mm;             # radius of curvature C
    Theta1 = 2.5 deg;            # bend angle (on left side)
    Theta2 = 2.5 deg;            # bend angle (on right side)
    WidthA = 6.0 um;             # width of bend section A
    WidthB = 6.0 um;             # width of bend section B
    WidthC = 6.0 um;             # width of bend section C
    OffsetA = 0;                 # offset of bend section A
    OffsetB = 0;                 # offset of bend section B
    OffsetC = 0;                 # offset of bend section C
    Dmin1 = 10.0 um;             # minimum waveguide c-c separation
    Dmin2 = 10.0 um;             # minimum waveguide c-c separation
    Dmax1 = 250 um;              # maximum waveguide c-c separation
    Dmax2 = 120 um;              # maximum (arm) waveguide c-c separation
    Lc1 = 1213 um;               # length of first coupler segment
    Lc2 = 2757 um;               # length of second coupler segment
    LArmUpper = 100 um;          # length of grating arm segment
    dyLower = 7.7 um;            # vertical path difference in arms
    Center = true;               # whether to center device horizontally
}

# Note: multiple devices may be defined in a single input file.
# When constructing a set of similar devices, you only need to
# specify which parameters are different from the most recently
# defined device.

device CouplerPair { Name = "mzc2"; dyLower = 7.8 um; }
device CouplerPair { Name = "mzc3"; dyLower = 7.9 um; }
device CouplerPair { Name = "mzc4"; dyLower = 8.0 um; }
device CouplerPair { Name = "mzc5"; dyLower = 8.1 um; }
device CouplerPair { Name = "mzc6"; dyLower = 8.2 um; }
device CouplerPair { Name = "mzc7"; dyLower = 8.3 um; }
device CouplerPair { Name = "mzc8"; dyLower = 8.4 um; }
device CouplerPair { Name = "mzc9"; dyLower = 8.5 um; }
```

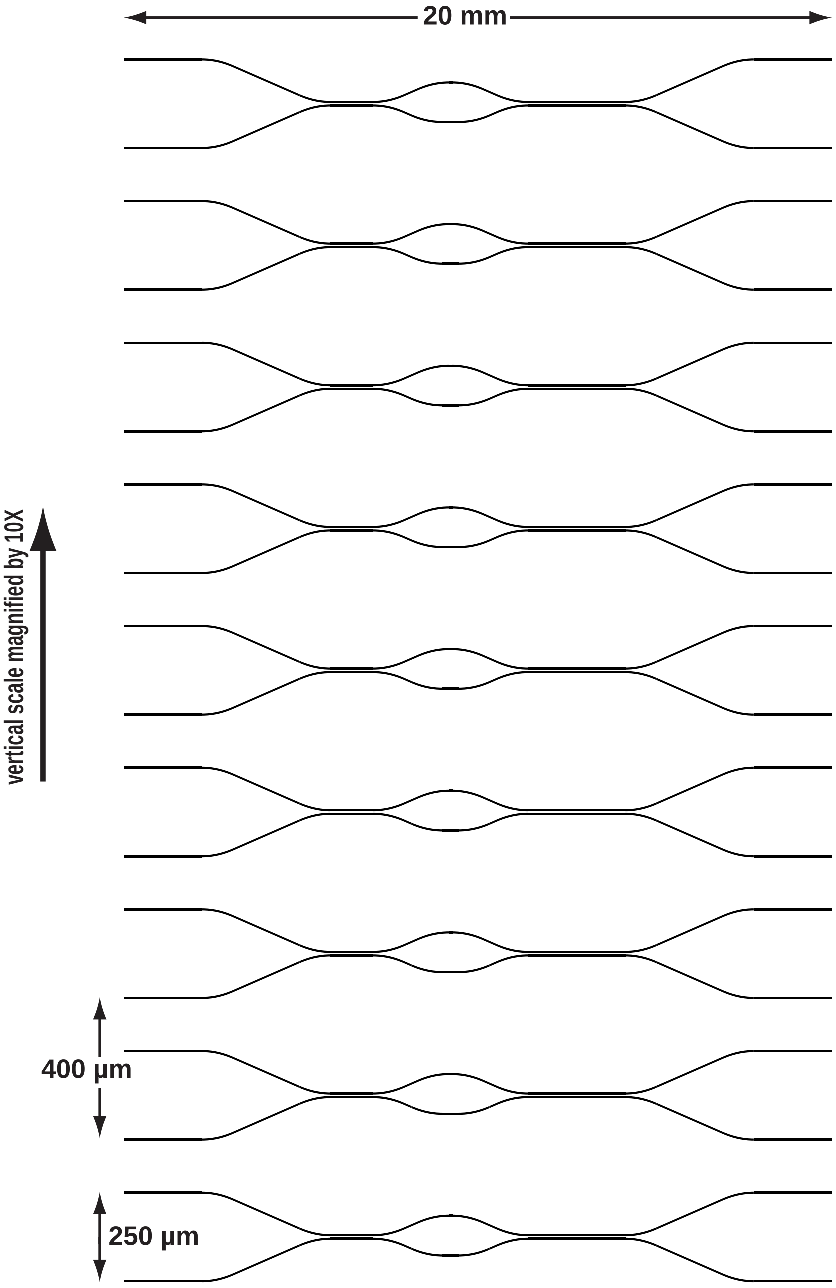

Figure B.1: Waveguide patterns generated by the layout program, with the input file given in this appendix. In this figure, the vertical scale has been expanded by 10 × in order to better illustrate the devices. (figs/B/test-chip-pattern.eps)

Appendix C

Phase Distortion in Bragg Gratings

C.1 Evidence of Chirp in Bragg Gratings

In Section 4.2.2, we presented the measured spectral response for 4 mm-long Bragg gratings, and compareed the results to theoretical predictions. Although the measured response agrees well with a theoretical model, a closer inspection of the data plotted in Fig. 4.18 reveals that the measured spectral response is slightly broader than that predicted. One possible explanation for this discrepancy is simply that the grating length is somewhat shorter than 4 mm, or that the group index is slightly smaller than our estimate. However, the grating length can be measured unambiguously to a high degree of accuracy, and the discrepancy in group velocity required to explain the observed bandwidth change is too large to be supported by theory.

The theoretical spectral response plotted in Fig. 4.18 assumes that the Bragg grating is a uniform, perfectly periodic structure which begins and ends abruptly. Under these conditions, the transmission spectrum is given by Eq. 2.129. When these criteria are not satisfied, the spectral response cannot be expressed in closed form, and it must instead be computed by numerically solving the coupled mode equations, as described in Section 2.3.5. Because of the fabrication techniques used to form the gratings, we know that the grating regions begin and end abruptly, but we cannot rule out a distortion in the grating pattern. To first-order, such a distortion can be described by simple linear chirp:

$$\Lambda(z) = \Lambda_0 + Az \quad , \tag{C.1}$$

Figure C.1 compares the measured and calculated transmission spectrum, when we allow for a linear chirp in the grating. The chirp rate A used in this model was ± 27 pm / nm[1], and the grating strength (κ) was 4.5 cm^{-1}. Because the spectral response of the chirped grating must be numerically calculated for each frequency point plotted, is not easy to extract the theoretical parameters by a least squares fitting algorithm. (Each spectral response takes approximately two minutes to compute on a 400 MHz computer.) Therefore, the chirp rate of ± 27 pm / cm was determined by simple trial and error method. Nevertheless, we believe this technique can determine the chirp rate to within 2 pm / nm. The excellent agreement between theory and experiment shown in Fig. C.1 supports our hypothesis that there is a chirp in the system.

The effect of grating chirp can be seen more clearly for longer Bragg gratings. Figure C.2 shows the measured spectral response for an 8 mm-long Bragg grating, for both polarization states. For these devices, the spectral response shows two distinct dips in transmission. Figure C.3 presents a more detailed picture of the measured TE spectral response, along with a theoretically computed spectrum. For this device, the inferred chirp rate is ± 56 pm / cm, and the grating strength is 4.0 cm^{-1}. Also shown for comparison in Fig. C.3 is the theoretically calculated spectral response that the Bragg grating would have in the absence of chirp. As for the 4 mm devices, the measured spectral response is very accurately modeled by including a linear chirp in the grating. The deviation between theory and measurements could arise if the grating distortion has a more complicated form than a simple linear chirp.

When the linear chirp rate is 56 pm / cm over 8 mm (as it is for the data plotted in Fig. C.3), the grating period is 45 pm longer or shorter at the end of the grating than it is at the beginning. For a 223 nm period grating, this corresponds to a change in period of only 200 parts per million over the length of the grating. Although this seems like a small change, it corresponds to a change in the free-space Bragg wavelength of about 0.3 nm, which is equivalent to a 40 GHz frequency shift.

Another way to quantify the distortion is in terms of the equivalent in-plane distortion of the grating. Imagine comparing the chirped grating to a uniform unchirped grating, whose period is chosen to correspond to the that of the chirped grating at $z = 0$, (i.e., Λ_0.) The two gratings would be identical at the $z = 0$, but the distorted grating would slowly walk off with respect to the uniform grating as one moves away from $z = 0$. This walk-off

---

[1] Because we have measured only the amplitude of the spectral response (and not the phase), we can determine the magnitude of the chirp rate, but not the sign.

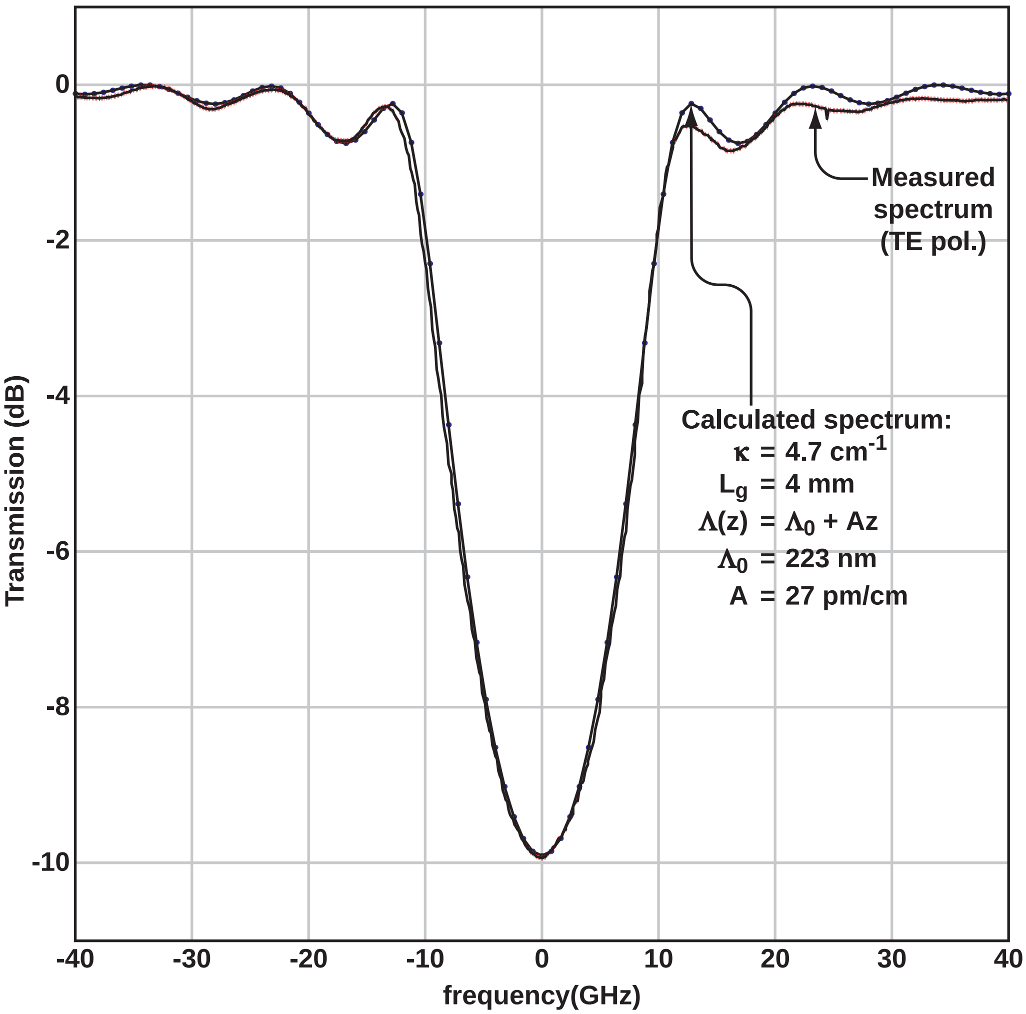

Figure C.1: A comparison between the measured transmission spectral response for a 4 mm Bragg grating and the theoretically calculated response, when we allow for a possible linear chirp in the grating. The measured spectrum is well described by a linear chirp of 27 pm / cm. (figs/C/grating-chirp-fit-4.eps)

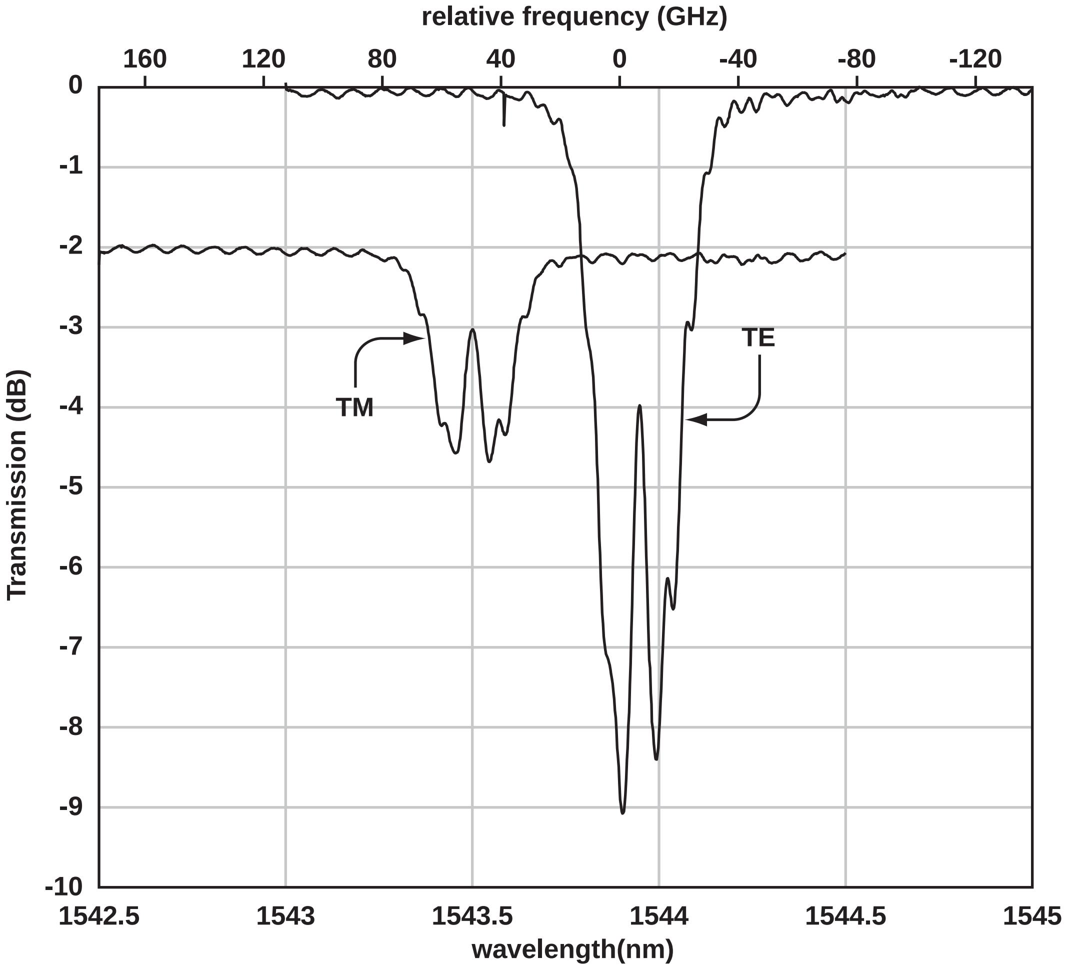

Figure C.2: Transmission spectrum for an 8 mm-long integrated Bragg grating in SOI. For comparison, both TE and TM polarizations are plotted. The TM data has been offset by -2 dB in order to more clearly differentiate the two spectra. (figs/C/tetmcompare8.eps)

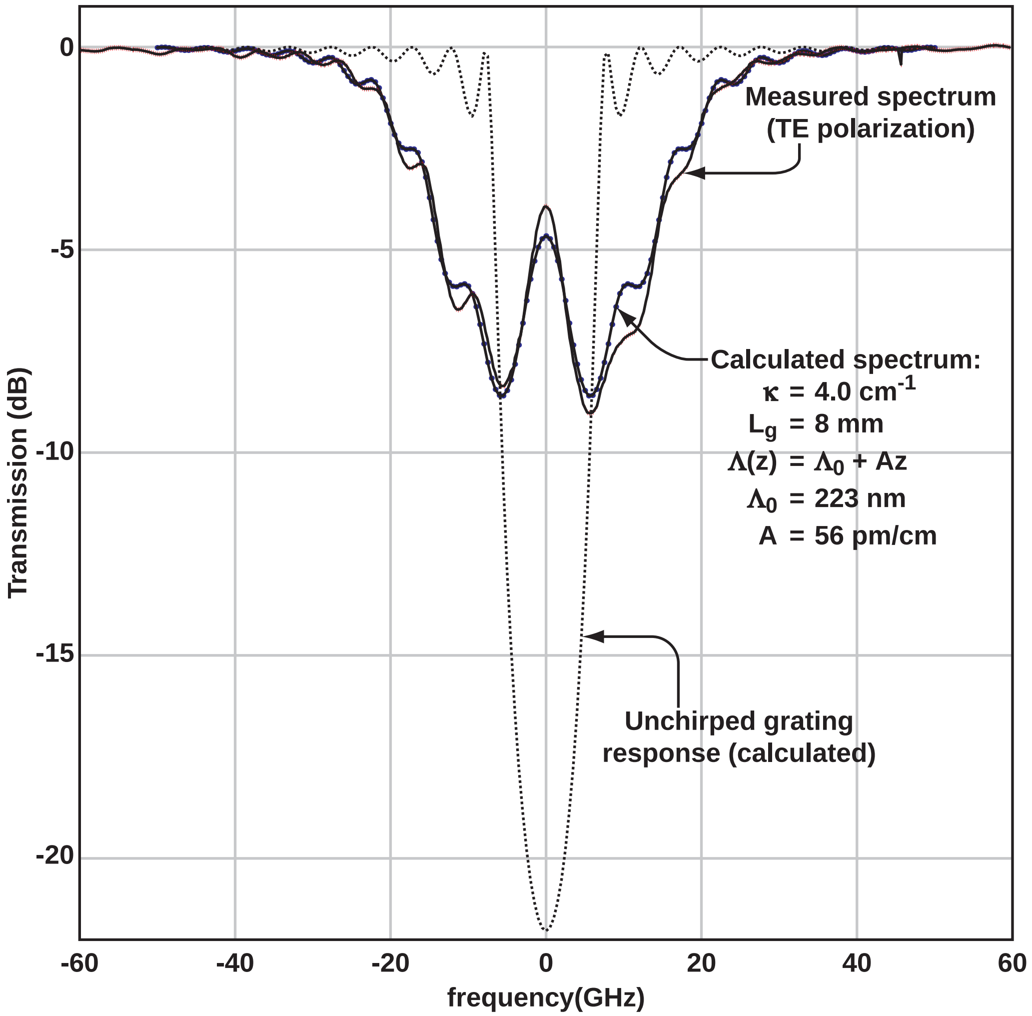

Figure C.3: A comparison between the measured and calculated transmission spectra for an 8 mm-long Bragg grating. The measured spectrum is well described by a linear chirp of 56 pm / cm. For reference, we have also plotted the predicted spectral response for a similar grating without chirp (dashed line.) (figs/C/grating-chirp-fit-8.eps) can be simply related to the grating phase $\phi(z)$, which was introduced in Section 2.3.5.

$$\Delta z = \frac{\Lambda_0}{2\pi}\left(\phi(z) - \frac{2\pi}{\Lambda_0}z\right) \quad , \tag{C.2}$$

where $\phi(z)$ represents the phase of the chirped grating and $2\pi z/\Lambda_0$ is the phase of the unchirped grating. Using the Taylor series expansion given in Eq. 2.161 for $\phi(z)$, we obtain following expression for the in-plane distortion:

$$\Delta z = -\frac{A}{2\Lambda_0}z^2 \tag{C.3}$$

Thus, for a linearly chirped grating, the in-plane distortion grows quadratically. For the data plotted in Fig. C.3, if Λ_0 is chosen to be the period of the grating at its midpoint, the ends of the grating will be displaced from their unchirped positions by 200 nm, almost one full period.

C.2 Sources of Distortion

The analysis of the previous section provides fairly strong evidence that there is a chirp in the devices, but it does not unambiguously identify the source or sign of this chirp. There are many effects which could explain the observed chirp, and they can be broadly classified in two categories: either the grating itself is chirped, or there is some equivalent change in the medium of propagation. Even for a uniform grating, if the effective index (or equivalently, the propagation constant, β) of the guiding medium changes over the length of the device, the structure will behave as if the grating were chirped.

One potential source of grating distortion is the interference lithography technique used to generate gratings. As described in Section 3.2.1, the interfering waves in the interference lithography systems used at MIT are not plane waves but diverging spherical waves formed by focusing a laser beam through a small pinhole. The standing wave pattern formed by interfering spherical waves can be accurately described as a family of hyperbolas, as described in reference [117]. Figure C.4 depicts the calculated grating distortion for the Lloyd's mirror interference lithography system used to form the gratings on the devices which we measured. Figure C.4a shows contours of constant grating period and Fig. C.4b depicts the calculated local grating chirp[2]. The calculated chirp grows to either side of the plane of symmetry of the system, and depends weakly on the vertical

---

[2] The local grating chirp is proportional to the second derivative of the grating phase at each point.

position y. In the Lloyd's mirror system, the plane of symmetry ($z = 0$) corresponds to the plane of the mirror. When lithography is completed, the mirror edge makes a clearly visible boundary in the grating which marks the plane of symmetry of the system. For the device reported in this work, we were careful to position the gratings within 1 cm of the plane of symmetry in order to minimize the hyperbolic distortion. Based upon the theory, the observed grating chirp is too large by a factor of 4-5 to be explained entirely by the natural hyperbolic phase distortion of the grating. For comparison, we have also calculated the theoretical distortion map for the conventional two-beam interference lithography system, also used at MIT. In this system, the arm length is slightly smaller and therefore the distortion is predicted to be more severe.

Another effect which could lead to a physical chirp in the grating is if the wafer or mirror were not flat. Most wafers are not completely flat, even when held on a vacuum chuck. However, it is important point out that the flatness only matters over a small, 4-8 mm portion of the wafer where the gratings reside. The mirror must likewise be flat over a similar length scale. Furthermore, because the gratings are essentially one-dimensional structures that are periodic in the z direction, the critical figure is the flatness in the z direction and not the flatness in the y direction. To lowest order, the nonflatness can be described by a simple linear height variation in the z direction. This type of distortion is equivalent to making the wafer and mirror nonperpendicular. The diverging spherical wave approximation mentioned earlier can easily be extended to account for this type of angular misalignment. For the system used to expose the gratings, the mirror and wafer would have to be nonperpendicular by about 2 degrees in order to explain the observed chirp. This corresponds to a wafer height variation of 280 µm over 8 mm. This level of nonflatness would have been visible to the naked eye during the setup of the system.

Another possibility is that after the grating is formed on the substrate, subsequent processing of the sample causes the surface to distort. For example, perhaps when the sample is released from the vacuum chuck which holds it during the exposure, the surface relaxes to some distorted shape. Another place where distortion may occur is when the sample is polished and mounted prior to measurement. The samples are mounted onto metal tabs using a heated epoxy prior to measurement. It is possible that the sample distorts during cooling because of the potentially different thermal expansion coefficients of the epoxy and the sample. In order for these effects to explain the observed chirp, the type of induced distortion must be more complex than a simple linear expansion or contraction: a linear chirp in the grating requires that the in-plane surface distortion has a quadratic profile.

As mentioned earlier, it is also possible for the chirp to arise because of a variation in the guiding medium, even if the grating itself is perfectly periodic. To understand this

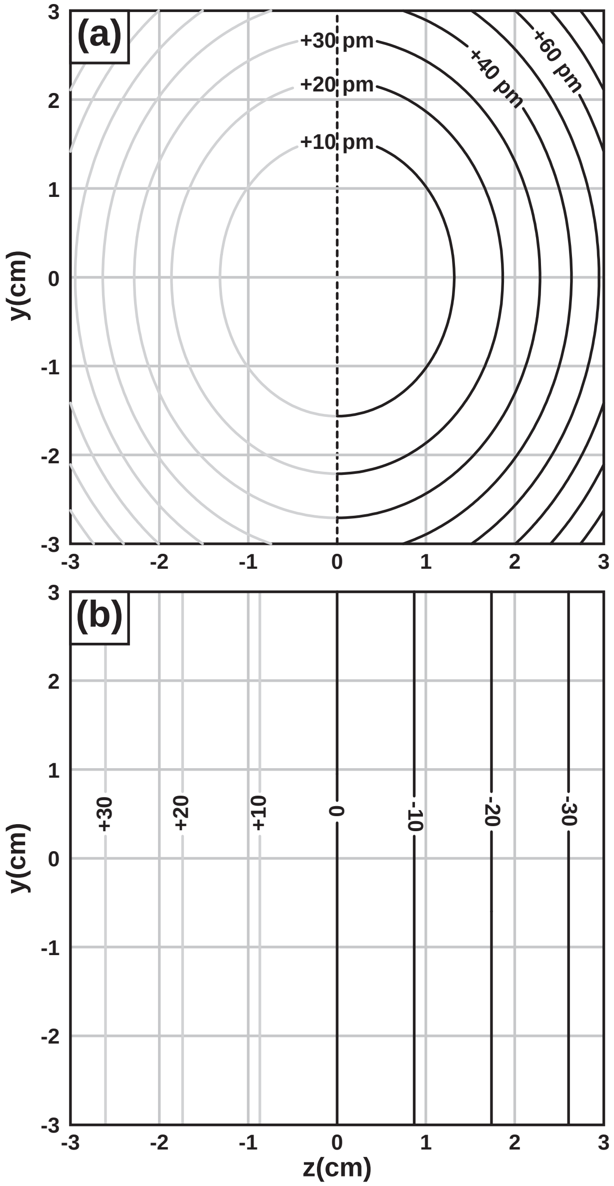

Figure C.4: Calculated distortion map for the Lloyds mirror interference lithography system used to generate the Bragg gratings which we measured. The following parameters were used in these calculations: the illuminating wavelength was 325 nm, the central period of the grating was 223 nm, and the distance from the center of the sample to the pinhole was 165 cm. (a) Contours showing the change in grating period (the period at the center is 223 nm.) (b) Contour plot of the local grating chirp. The contour labels have dimensions of pm / cm. (figs/C/hypermap-325.eps)

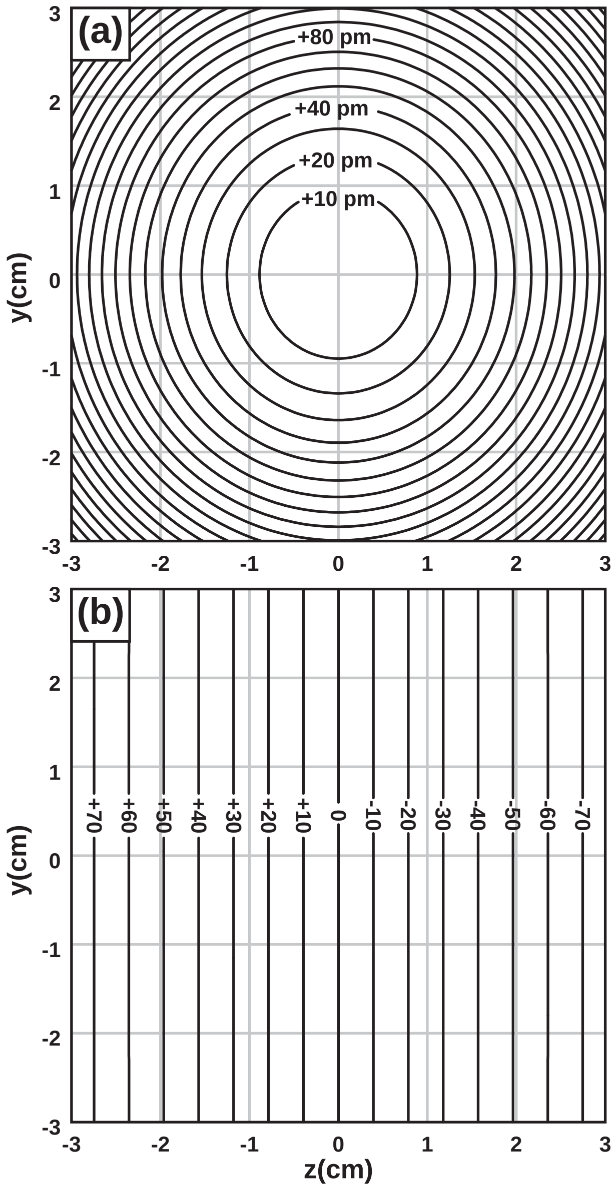

Figure C.5: Calculated distortion map for the conventional two-beam interference lithography system used at MIT. The following parameters were used in these calculations: the illuminating wavelength was 351 nm, the period of the grating at the center was 223 nm, and the distance from the center of the sample to the pinhole was 100 cm. (figs/C/hypermap-351.eps) effect, it is necessary to return to the coupled mode equations for a nonuniform grating:

$$\frac{d}{dz}\begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} = \begin{bmatrix} -j\beta(z) & \kappa_0 e^{-j\phi(z)} \\ \kappa_0^* e^{+j\phi(z)} & j\beta(z) \end{bmatrix} \begin{bmatrix} A_+(z) \\ A_-(z) \end{bmatrix} \tag{C.4}$$

$$\text{where } \phi(z) \equiv \int_0^z k_{\mathrm{g}}(z')\,dz' \quad . \tag{C.5}$$

In this equation, we assume that the grating strength κ has a constant magnitude, but we allow for both grating chirp (through the function $\phi(z)$) and waveguide chirp (through the function $\beta(z)$). Of course, the waveguide chirp function $\beta(z)$ will change with frequency ω. When this problem was treated in Section 2.3.5, we made a slowly-varying-amplitude approximation by factoring out the rapid oscillations which have spatial frequency near some arbitary reference point β_r. We now take a slightly different approach in order to clearly show the equivalence of waveguide chirp and grating chirp. Imagine evaluating the waveguide chirp function $\beta(z)$ at some reference frequency ω_r:

$$\beta_r(z) = \beta(z)|_{\omega=\omega_r} \quad . \tag{C.6}$$

Next, we define a slowly-varying envelope functions $a_\pm(z)$ in terms of this reference chirp function $\beta_r(z)$,

$$A_\pm(z) = a_\pm(z)\exp\left[\mp j\int_0^z \beta_r(z')\,dz'\right] \quad . \tag{C.7}$$

By substituting Eq. C.7 into Eq. C.4, we arrive at the following coupled mode equations for the envelope functions $a_\pm(z)$:

$$\frac{d}{dz}\begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} = \begin{bmatrix} -j\delta & \kappa(z) \\ \kappa(z)^* & j\delta \end{bmatrix} \begin{bmatrix} a_+(z) \\ a_-(z) \end{bmatrix} \quad , \tag{C.8}$$

where the $\kappa(z)$ and δ are defined as:

$$\delta(z) = \beta(z) - \beta_r(z) \tag{C.9}$$

$$\kappa(z) = \kappa_0 \exp\left\{-j\int_0^z \left[k_g(z') - 2\beta_r(z')\right]\,dz'\right\} \quad . \tag{C.10}$$

Notice that δ is now the difference between two functions of z, however the difference between the two should be relatively independent of z over the bandwidths of interest. To first-order, the difference between the two functions is simply proportional to the deviation from the reference frequency ω_r:

$$\delta = \frac{1}{v_g}(\omega - \omega_r) \quad , \tag{C.11}$$

where v_g is the group velocity. Notice that when the coupled mode equations are cast in this way, grating distortion and waveguide distortion are indistinguishable through the relation described in Eq. C.10.

There are several geometrical and optical parameters that could give rise to a variation in the waveguide over the length of the device. One possibility is that the silicon core thickness changes slightly from one end to the other. Another possibility is that the duty cycle, or grating etch depth changes across the device. Also, the ridge height or ridge width could vary systematically across the device. We have investigated each of these effects separately, using metrological tools where possible to characterize the samples and analytical tools to compute the effect of small changes in these parameters. Based upon our observations and calculations, the observed chirp cannot be explained by a variation in any of these geometrical parameters. Furthermore, it is unlikely that these parameters could change substantially without also affecting the magnitude of the grating strength $|\kappa|$, which would lead to additional effects in the transmission spectrum. Another possibility is that the refractive index of the silicon guiding layer changes slightly from one end of the device to the other. This could occur either because of compositional nonuniformity or because of nonuniform stress / strain in the sample. The required change in refractive index is approximately 0.001 over 8 mm in order to account for the observed chirp. Currently, our metrological tools are not able to measure the refractive index with sufficient spatial resolution and measurement accuracy to diagnose this type of change.

C.3 Measurements of Distortion

In the preceding section, we discussed various things which could give rise to a chirp in the Bragg grating. Based upon spectral measurements alone, it is impossible to determine whether there is a chirp in the grating, or some other systematic variation in the waveguide composition or structure. This ambiguity could be removed if there were some metrological tool for accurately measuring the grating distortion.

One way of measuring the grating distortion is to measure the diffraction angle when light is incident on the grating from above. By viewing the position of the diffracted beam on a distant screen, it should be possible to experimentally measure the distortion map similar to that presented in Fig. C.4a. This technique requires that the grating area be larger than the size of the incident beam, and the spatial resolution of this technique is similarly limited by the beam size.

Another way to characterize the distortion in a grating is to examine the grating in an electron-beam-lithography system. In an e-beam lithography system, the sample sits on a movable stage whose position is measured and controlled by an interferometer feedback system. The electron-beam-lithography system can be used as a scanning-electron-microscope to view the grating as the stage moves to precisely defined locations.

Figure C.6 depicts the technique for characterizing the grating distortion using e-beam metrology. A scanning e-beam system is used to capture a series of images of the grating at evenly spaced point on the sample. The technique relies upon a feedback-controlled interferometer system to accurately position the stage prior to capturing each image. The total image area is typically chosen to be 2-4 µm, in order to avoid undesirable image distortion that can occur for larger field deflections.

After the images are captured, they can be analyzed using off-line image processing to determine the relative position of the grating lines in each of the images. Consider first the case of a perfectly periodic, undistorted Bragg grating. If the displacement d is precisely an integral number of periods, then all of the images should appear almost identical to one another. In the more likely case that the displacement d is not an integral number of periods, the captured images will “walk off” as one proceeds across the sample. It is relatively easy to show that for an undistorted grating, the images will walk off in a linear progression. The rate of this progression is an arbitrary figure which is related to the precise value of the displacement d in relation to the grating period Λ. If the grating is not uniform, e.g. if there’s a systematic chirp, the relative position of the grating lines within each image will follow a non-linear progression as shown in Fig. C.6b. For a grating with a simple linear chirp, the relative position should have a quadratic shape, where the second derivative is proportional to the chirp rate A.

Figure C.7 plots the measured relative grating position vs. z for a grating produced by the Lloyd’s mirror interference lithography system. The approximate grating period for the sample was 223 nm, and the adjacent sample points were spaced at 250 µm intervals over a total length of about 2.6 cm. There is a pronounced curvature in the data plotted in Fig. C.7, which is indicative of the chirp in the grating. The data were fitted with a quadratic polynomial to extract the chirp rate. The quadratic polynomial coefficient indicates this grating is chirped by about 53 pm/cm, a figure which is consistent with the observed spectra from other samples. The discrepancy at the left edge of the plot is likely

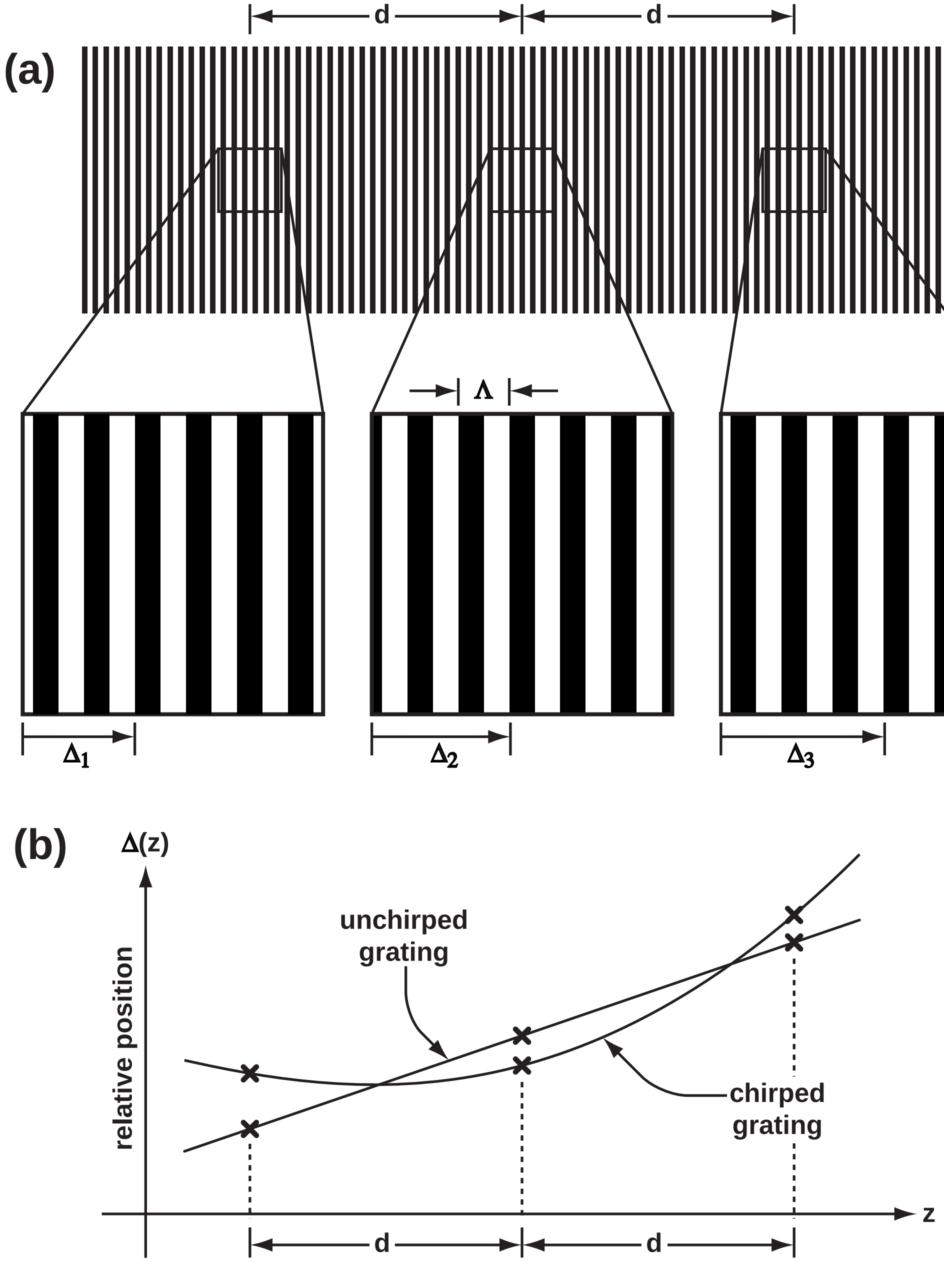

Figure C.6: Diagram illustrating the technique for characterizing the distortion in a Bragg grating using an e-beam metrology technique. (a) The grating is imaged at a series of evenly spaced places on the sample. (b) The relative position of the grating lines in each image can be used to infer the distortion in the grating. (figs/C/chirp-measure-method.eps) caused by scattered light from the edge of the Lloyd's mirror (the left edge of the data plotted in Fig. C.7 is very close to the boundary of the grating region, where the Lloyds mirror contacts the wafer.)

These preliminary measurements suggest that the distorted spectra of the integrated Bragg gratings is caused by a real chirp in the grating pattern and not by a nonuniformity in the waveguide. However, the level of chirp which we measure is inconsistent with a simple spherical-wave model of the interference lithography system. Further investigation will be needed to conclusively identify (and hopefully eliminate) the source of this chirp.

More measurements are needed to quantify the accuracy of this technique. A more comprehensive study of the distortion in grating's produced by both interference lithography systems could provide valuable information about the fidelity and spatial coherence of interference-lithography-produced gratings.

C.4 Summary

The measured spectra of the integrated Bragg gratings shows strong evidence of a linear chirp in the grating pattern. This effect is more pronounced for longer Bragg gratings than for short Bragg gratings. The chirp could be explained by a number of different phenomena both in a lithographic system used to produce the gratings and in the waveguide itself. We describe the experimental technique, based on e-beam metrology for measuring the grating distortion. Our measurements point towards a real chirp in the grating, but further work is needed to determine the exact source of this chirp.

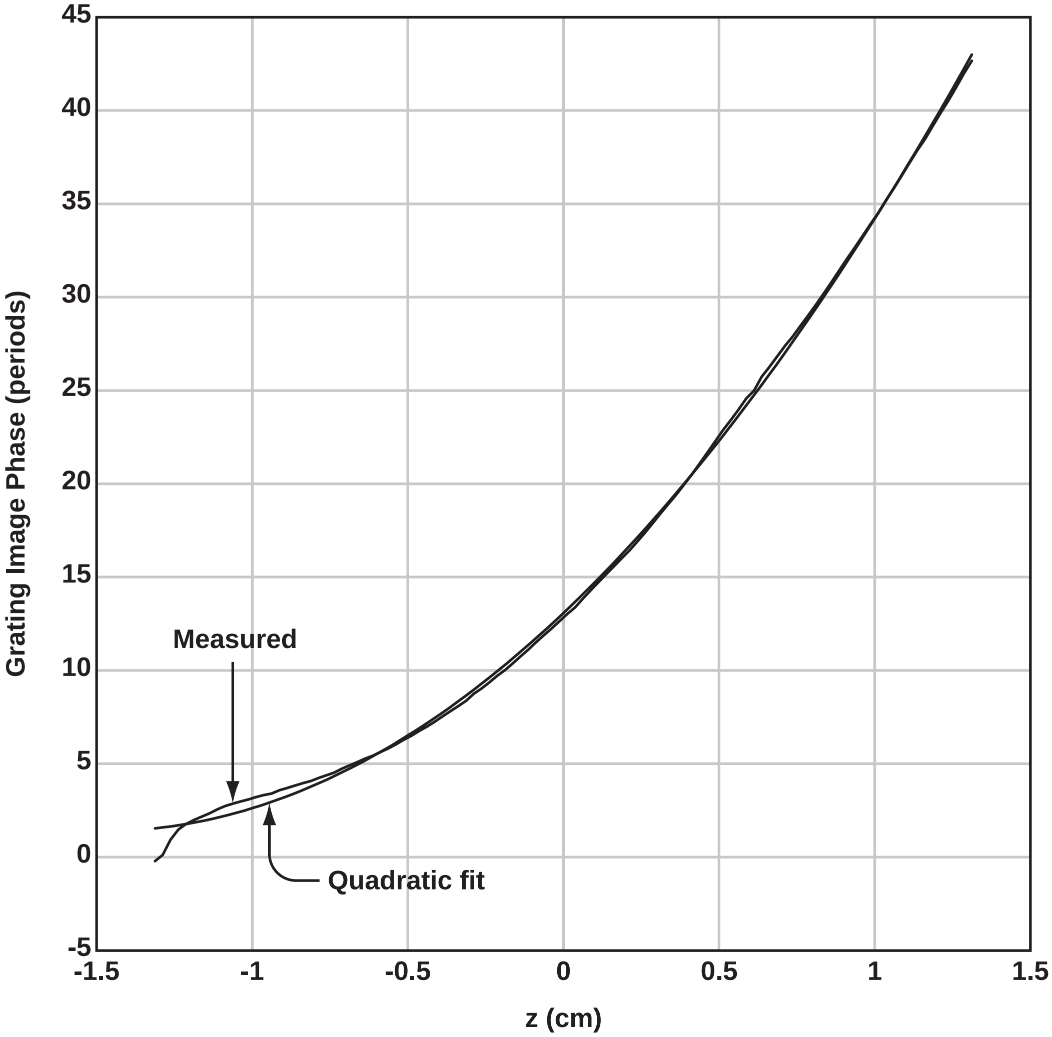

Figure C.7: Measured phase distortion for a grating produced in the Lloyds mirror interference lithography system. The grating is sampled at 250 μm intervals and the relative grating position is plotted as a function of z. The curvature in the data presented here indicates that the grating is chirped by approximately 53 pm / cm. The discrepancy at the left edge of the plot is believed to be caused by scattered light from the edge of the Lloyds mirror. (figs/C/chirp-measure.eps)

Bibliography

[1] Y. Yano, T. Ono, K. Fukuchi, T. Ito, H. Yamazaki, Y. M., and K. Emura. 2.6 terabit/s WDM transmission experiment using optical duobinary coding. In *22nd European Conference on Optical Communication*, volume 5, pages 3–6. 1996.

[2] M. Murakami, T. Matsuda, and T. Imai. Quarter terabit (25*10 Gb/s) over 9288 km WDM transmission experiment using nonlinear supported RZ pulse in higher order fiber dispersion managed line. In *24th European Independent Conference on Optical Communication*, volume 3, pages 77–81. 1998.

[3] P. S. Henry. Lightwave Primer. *J. Quantum Electron.*, **QE-12**(12), 1862 – 1879, 1985.

[4] J. M. Senior. *Optical Fiber Communications*. Prentice-Hall, 1985.

[5] P. W. Shumate. The broadest broadband fiber to the home. *Scientific American*, **281**(4), 104 – 105, 1999.

[6] T. Wood, G. C. Wilson, R. D. Feldman, and J. A. Stiles. (FTTH) system providing broad-band data over cable modems along with analog and digital video. *IEEE Photon. Technol. Lett.*, **11**(4), 475 – 477, 1999.

[7] T. Woodward and A. Krishnamoorthy. 1-Gb/s integrated optical detectors and receivers in commercial CMOS technologies. *J. Select. Topics in Quantum. Electron.*, **5**(2), 146 – 156, 1999. (and other articles from this volume).

[8] A. Yariv and M. Nakamura. Periodic Structures for Integrated Optics. *J. Quantum Electron.*, **QE-13**(4), 233–253, 1977.

[9] J. S. V. Delden. Optical Circulators Improve Bidirectional Fiber Systems. *Laser Focus World*, pages 109 – 112, November 1995.

[10] D. L. Lee. *Electromagnetic principles of integrated optics*. Wiley, 1986.

[11] A. W. Snyder and J. D. Love. *Optical Waveguide Theory*. Chapman and Hall, 1983.

[12] D. K. Cheng. *Field and Wave Electromagnetics*. Addison-Wesley, 2nd edition, 1990.

[13] J. D. Jackson. *Classical Electrodynamics*. Wiley, 2nd edition, 1975.

[14] T. Tamir (editor). *Guided-wave optoelectronics*, volume 26 of *Springer-Verlag series in electronics and photonics*. Springer-Verlag, 2nd edition, 1990. See section 2, Theory of Optical Waveguides, by H. Kogelnik.

[15] C. L. Xu, W. P. Huang, M. S. Stern, and S. K. Chaudhuri. Full-vectorial Mode Calculations by Finite Difference Method. *IEE Proc. Opt.*, **141**(5), 281 – 286, 1994.

[16] W. P. Huang and C. L. Xu. Simulation of Three-Dimensional Optical Waveguides by a Full-Vector Beam Propagation Method. *J. Quantum Electron.*, **29**(10), 2639–2649, 1993.

[17] Z.-E. Abid, K. L. Johnson, and A. Gopinath. Analysis of Dielectric Guides by Vector Transverse Magnetic Finite Elements. *J. Lightwave Technol.*, **11**(10), 1545 – 1549, 1993.

[18] C. Yeh, K. Ha, S. B. Dong, and W. P. Brown. Single-mode optical waveguides. *Appl. Opt.*, **18**(10), 1490 – 1504, 1979.

[19] G. R. Hadley and R. E. Smith. Full-Vector Waveguide Modeling Using an Iterative Finite-Difference Method with Transparent Boundary Conditions. *J. Lightwave Technol.*, **13**(3), 465 – 469, 1995.

[20] P. Lüsse, P. Stuwe, J. Schüle, and H.-G. Unger. Analysis of Vectorial Mode Fields in Optical Waveguides by a New Finite Difference Method. *J. Lightwave Technol.*, **12**(3), 487 – 493, 1994.

[21] M. S. Stern. Semivectorial Polarized Finite Difference Method for Optical Waveguides with Arbitrary Index Profiles. *IEE Proc. Opt.*, **135**(1), 56 – 63, 1988.

[22] J. E. Goell. A Circular-Harmonic Computer Analysis of Rectangular Dielectric Waveguides. *Bell Sys. Tech. J.*, pages 2133 – 2160, September 1969.

[23] W. P. Huang, C. L. Xu, S.-T. Chu, and S. K. Chaudhuri. The Finite Difference Vector Beam Propagation Method: Analysis and Assessment. *J. Lightwave Technol.*, **10**(3), 295 – 305, 1992.

[24] W. P. Huang, C. L. Xu, and S. K. Chaudhuri. A Finite-Difference Vector Beam Propagation Method for Three-Dimensional Waveguide Structures. *IEEE Photon. Technol. Lett.*, **4**(2), 148 – 151, 1992.

[25] C. L. Xu, W. P. Huang, and S. K. Chaudhuri. Efficient and Accurate Vector Mode Calculations by Beam Propagation Method. *J. Lightwave Technol.*, **11**(7), 1209 – 1215, 1993.

[26] J. Peraire. Finite Difference Discretization of Elliptic Equations, 1999. Course notes from 16.920 (Numerical Methods for Partial Differential Equations).

[27] W. H. Press, W. Vetterling, S. A. Teukolsky, and B. P. Flannery. *Numerical Recipes in C.* Cambridge University Press, 2nd edition, 1992.

[28] R. B. Lehoucq and J. A. Scott. An Evaluation of Software for Computing Eigenvalues of Sparse Nonsymmetric Matrices. Technical Report Preprint MCS-P547-1195, Argonne National Laboratories, lehoucq@mcs.anl.gov, 1996. Available at http://www.cs.sandia.gov/ rlehoucq/papers.html.

[29] D. C. Sorensen. Implicit Application of Polynomial Filters in a k-step Arnoldi method. *SIAM J. Matrix Analysis and Applications*, **13**(1), 357 – 385, 1992.

[30] Y. P. Li and C. H. Henry. Silica-based optical integrated circuits. *IEE Proc. Opt.*, **143**(5), 263–280, 1996.

[31] M. Kawachi. Silica Waveguides on Silicon and Their Application to Integrated-Optic Components. *Optics and Quantum Electron.*, **22**, 391–416, 1990.

[32] J. Schmidtchen, A. Splett, B. Schüppert, and K. Petermann. Low Loss Single Mode Optical Waveguides with Large Cross-section in Silicon-on-Insulator. *Electron. Lett.*, **27**(16), 1486 – 1488, 1991.

[33] U. Fischer, T. Zinke, J.-R. Kropp, F. Arndt, and K. Peterman. 0.1 dB/cm Waveguide Losses in Single-Mode SOI Rib Waveguides. *IEEE Photon. Technol. Lett.*, **8**(5), 647 – 648, 1996.

[34] R. A. Soref, J. Schmidtchen, and K. Peterman. Large Single-Mode Rib Waveguides in GeSi-Si and Si-on-SiO_2. *J. Quantum Electron.*, **27**(8), 1971 – 1973, 1991.

[35] S. P. Pogossian, L. Vescan, and A. Vonsovici. The Single-Mode Condition for Semiconductor Rib Waveguides with Large Cross-section. *J. Lightwave Technol.*, **16**(10), 1851 – 1853, 1998.

[36] H. A. Haus and W. P. Huang. Coupled-Mode Theory. *Proc. IEEE*, **79**(10), 1505 – 1518, 1991.

[37] H. A. Haus, W. P. Huang, S. Kawakami, and N. A. Whitaker. Coupled-Mode Theory of Optical Waveguides. *J. Lightwave Technol.*, **5**(1), 16 – 23, 1987.

[38] E. Schweig and W. B. Bridges. Computer Analysis of Dielectric Waveguides: a Finite-Difference Method. *IEEE Trans. Microwave Theory Tech.*, **MTT-32**(5), 531 – 541, 1984.

[39] T. E. Murphy. *Integrated Optical Grating-Based Matched Filters for Fiber-Optic Communications*. Master's thesis, Massachusetts Institute of Technology, 1996.

[40] Z. Weissman, A. Hardy, and E. Marom. On the Applicability of the Coupled Mode Theory to Non-Parallel Waveguide Systems. *Optics Comm.*, **71**(6), 341 – 344, 1989.

[41] H. A. Haus and W. P. Huang. Mode Coupling in Tapered Structures. *J. Lightwave Technol.*, **7**(4), 729 – 730, 1989.

[42] W. P. Huang and H. A. Haus. Self-Consistent Vector Coupled-Mode Theory for Tapered Optical Waveguides. *J. Lightwave Technol.*, **8**(6), 922– 926, 1990.

[43] A. Takagi, N. Takato, Y. Hida, T. Oguchi, and T. Nozawa. Silica-Based Line-Symmetric Series-Tapered (LSST) Broadband Directional Coupler with Asymmetric Guides Only in Their Parallel Coupling Regions. *Electron. Lett.*, **32**(18), 1700 – 1703, 1996.

[44] A. Takagi, K. Jinguji, and M. Kawachi. Silica-Based Waveguide-Type Wavelength-Insensitive Couplers (WINC's) with Series-Tapered Coupling Structure. *J. Lightwave Technol.*, **10**(12), 1814 – 1824, 1992.

[45] A. Takagi, K. Jinguji, and M. Kawachi. Broadband Silica-Based Optical Waveguide Coupler with Asymmetric Structure. *Electron. Lett.*, **26**(2), 132 – 133, 1990.

[46] A. Takagi, K. Jinguji, and M. Kawachi. Design and Fabrication of Broad-Band Silica-Based Optical Waveguide Couplers with Asymmetric Structure. *J. Quantum Electron.*, **28**(4), 848 – 855, 1992.

[47] A. Takagi, K. Jinguji, and M. Kawachi. Wavelength Characteristics of (2×2) Optical Channel-Type Directional Couplers with Symmetric or Nonsymmetric Coupling Structures. *J. Lightwave Technol.*, **10**(6), 735 – 746, 1992.

[48] R. Adar, C. H. Henry, R. F. Kazarinov, R. c Kistler, and G. R. Weber. Adiabatic 3-DB Couplers, Filters and Multiplexers Made with Silica Waveguides on Silicon. *J. Lightwave Technol.*, **10**(1), 46 – 50, 1992.

[49] F. Gonthier, D. Ricard, S. Lacroix, and J. Bures. Wavelength-Flattened $2{\times}2$ Splitters Made of Identical Single-Mode Fibers. *Optics Lett.*, **16**(15), 1201 – 1203, 1991.

[50] K. Jinguji, N. Takato, Y. Hida, T. Kitoh, and M. Kawachi. Two-Port Optical Wavelength Circuits Composed of Cascaded Mach-Zehnder Interferometers with Point Symmetrical Configurations. *J. Lightwave Technol.*, **14**(10), 2301 – 2310, 1996.

[51] K. Jinguji, N. Takato, A. Sugita, and M. Kawachi. Mach-Zehnder Interferometer Type Optical Waveguide Coupler with Wavelength-Flattened Coupling Ratio. *Electron. Lett.*, **26**(17), 1326 – 1327, 1990.

[52] B. E. Little and T. E. Murphy. Design Rules for Maximally Flat Wavelength-Insensitive Optical Power Dividers Using Mach-Zehnder Structures. *IEEE Photon. Technol. Lett.*, **9**(12), 1607–1609, 1997.

[53] C. Kittel. *Introduction to Solid State Physics.* Wiley, 6th edition, 1986.

[54] G. E. Kohnke, C. H. Henry, E. J. Laskowski, M. A. Capuzzo, T. A. Strasser, and A. E. White. Silica Based Mach-Zehnder Add-Drop Filter Fabricated with UV Induced Gratings. *Electron. Lett.*, **32**(17), 1579 – 1580, 1996.

[55] R. Kashyap, G. D. Maxwell, and D. L. Williams. Photoconduction and Germanium and Phosphorus Doped Silica Waveguides. *Appl. Phys. Lett.*, **62**(3), 214 – 216, 1993.

[56] Y. Hibino, T. Kitagawa, K. O. Hill, F. Bilodeau, B. Malo, J. Albert, and D. C. Johnson. Wavelength Division Multiplexer with Photoinduced Bragg Gratings Fabricated in a Planar-Lightwave-Circuit-Type Asymmetric Mach-Zehnder Interferometer on Si. *IEEE Photon. Technol. Lett.*, **8**(1), 84 – 86, 1996.

[57] H. J. Lee, C. H. Henry, R. F. Kazarinov, and K. J. Orlowsky. Low Loss Bragg Reflectors on SiO_2-Si_3N_4-SiO_2 Rib Waveguides. *Appl. Opt.*, **26**(13), 2618 – 2620, 1987.

[58] Y. Yamamoto, T. Kamiya, and H. Yanai. Improved Coupled Mode Analysis of Corrugated Waveguides and Lasers. *J. Quantum Electron.*, **QE-14**(4), 245 – 258, 1978.

[59] B. E. Little and H. A. Haus. A Variational Coupled-Mode Theory for Periodic Waveguides. *J. Quantum Electron.*, **31**(12), 2258 – 2264, 1995.

[60] B. E. Little. A Variational Coupled-Mode Theory Including Radiation Loss for Grating-Assisted Couplers. *J. Quantum Electron.*, **12**(2), 188 – 195, 1996.

[61] D. L. Jaggard and C. Elachi. Floquet and Coupled-Waves Analysis of Higher-Order Bragg Coupling in a Periodic Medium. *J. Opt. Soc. America*, **66**(7), 674 – 681, 1976.

[62] W. Streifer, D. R. Scifres, and R. D. Burnham. Coupled Wave Analysis of DFB and DBR Lasers. *J. Quantum Electron.*, **QE-13**(4), 134 – 141, 1977.

[63] W. Streifer, D. R. Scifres, and R. D. Burnham. Analysis of Grating-Coupled Radiation in GaAs:GaAlAs Lasers and Waveguides. *J. Quantum Electron.*, **QE-12**(7), 422 – 428, 1976.

[64] E. Kapon and A. Katzir. The Effect of Complex Coupling Coefficients on Distributed Feedback Lasers. *J. Quantum Electron.*, **QE-18**(1), 66 – 71, 1982.

[65] M. Braun. *Differential Equations and Their Applications*. Springer-Verlag, 3rd edition, 1986.

[66] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. McGraw-Hill, 3rd edition, 1991.

[67] C. W. Helstrom. *Probability and Stochastic Processes for Engineers*. Prentice-Hall, 2nd edition, 1991.

[68] G. Einarsson. *Principles of Lightwave Communications*. Wiley, 1996.

[69] P. J. Lamaire, A. M. Vengsarkar, W. A. Reed, and D. J. DiGiovanni. Thermally Enhanced Ultraviolet Photosensitivity in GeO_2 and P_2O_5 Doped Optical Fibers. *Appl. Phys. Lett.*, **66**(16), 2034 – 2036, 1995.

[70] A. Othonos. Fiber Bragg Gratings. *Rev. Sci. Instrum.*, **68**(12), 4309–4341, 1997.

[71] J. D. Mills and C. W. J. H. nd W. S. Brocklesby. Evanescent Field Imaging of an Optical Fiber Bragg Grating. *Appl. Phys. Lett.*, **75**(26), 4058 – 4060, 1999.

[72] B. E. Little and C. Wu. Window Functions for Ideal Response in Distributed Feedback Reflection Filters. *IEEE Photon. Technol. Lett.*, **9**(1), 76 – 78, 1997.

[73] P. S. Cross and H. Kogelnik. Sidelobes Suppression in Corrugated-Waveguide Filters. *Optics Lett.*, **1**(1), 43 – 45, 1977.

[74] S. Mihailov, F. Bilodeau, O. Hill, D. Johnson, J. Albert, D. Stryckman, and C. Shu. Comparison of fiber Bragg grating dispersion-compensators made with holographic and e-beam written phase masks. *IEEE Photon. Technol. Lett.*, **11**(5), 572 – 574, 1999.

[75] R. Tiberio, D. Carr, M. Rooks, S. Mihailov, F. Bilodeau, J. Albert, D. Stryckman, D. Johnson, K. Hill, A. McClelland, and B. Hughes. Fabrication of electron beam generated, chirped, phase masks for fiber Bragg grating dispersion compensator. *J. Vac. Sci. Technol. B*, **16**(6), 3237 – 3240, 1998.

[76] H. A. Haus. *Waves and Fields in Optoelectronics*. Prentice-Hall, 1984.

[77] C. M. Weinert. Design of Fiber-Matched Uncladded Rib Waveguides on InP With Polarization-Independent Mode Matching Loss of 1 dB. *IEEE Photon. Technol. Lett.*, **8**(8), 1049 – 1051, 1996.

[78] O. Mitomi, K. Kasaya, Y. Tohmori, Y. Suzaki, H. Fukano, Y. Sakai, M. Okamoto, and S. Matsumoto. Optical Spot-Size Converters for Low-Loss Coupling between Fibers and Optoelectronics Semiconductor Devices. *J. Lightwave Technol.*, **14**(7), 1714 – 1719, 1996.

[79] R. E. Smith, C. T. Sullivan, G. A. Vawter, G. R. Hadley, J. R. Wendt, M. B. Snipes, and J. F. Klein. Reduced Coupling Loss Using a Tapered-Rib Adiabatic-Following Fiber Coupler. *IEEE Photon. Technol. Lett.*, **8**(8), 1052 – 1054, 1996.

[80] E. A. J. Marcatili. Bends in Optical Dielectric Guides. *Bell Sys. Tech. J.*, pages 2103 – 2132, September 1969.

[81] M. K. Smit, E. C. M. Pennings, and H. Blok. A Normalized Approach to the Design of Low-Loss Optical Waveguide Bends. *J. Lightwave Technol.*, **11**(11), 1737 – 1742, 1993.

[82] K. Thyagarajan, M. R. Shenoy, and A. K. Ghatak. Accurate Numerical Method for the Calculation of Bending Loss and Optical Waveguides Using a Matrix Approach. *Optics Lett.*, **12**(4), 296 – 298, 1987.

[83] I. C. Goyal, R. L. Gallawa, and A. K. Ghatak. Bent Planar Waveguides and Whispering Gallery Modes: a New Method of Analysis. *J. Lightwave Technol.*, **8**(5), 768 – 773, 1990.

[84] S. Valette, S. Renard, H. Denis, J. P. Jadot, A. Fournier, P. Philippe, P. Gidon, A. M. Grouillet, and E. Desgranges. Si-Based Integrated Optics Technologies. *Solid State Technology*, pages 69 – 75, February 1989.

[85] C. H. Henry, R. F. Kazarinov, H. J. Lee, K. J. Orlowsky, and L. E. Katz. Low Loss Si_3N_4-SiO_2 Optical Waveguides on Si. *Appl. Opt.*, **26**(13), 2621 – 2624, 1987.

[86] D. E. Zelmon, J. T. Boyd, and H. E. Jackson. Low Loss Optical Waveguides Fabricated by Thermal Nitridation of Oxidized Silicon. *Appl. Phys. Lett.*, **47**(4), 353 – 355, 1985.

[87] V. V. Wong, J. Ferrera, J. N. Damask, T. E. Murphy, and H. I. Smith. Distributed Bragg Grating Integrated Optical Filters; Synthesis and Fabrication. *J. Vac. Sci. Technol. B*, **13**, 2859 – 2864, 1995.

[88] B. L. Booth. Low Loss Channel Waveguides in Polymers. *J. Lightwave Technol.*, **7**(10), 1445 – 1453, 1989.

[89] A. Mukherjee, B. J. Eapen, and S. K. Baral. Very Low Loss Channel Waveguides in Polymethylmethacrylate. *Appl. Phys. Lett.*, **65**(25), 3179 – 3181, 1994.

[90] L. Eldada, C. Xu, K. M. T. Stengel, L. W. Shacklette, and J. T. Yardley. Laser-Fabricated Low-Loss Single-Mode Raised-Rib Waveguiding Devices in Polymers. *J. Lightwave Technol.*, **14**(7), 1704 – 1712, 1996.

[91] E. Wooten, K. Kissa, A. Yi-Yan, E. Murphy, D. Lafaw, P. Hallemeier, D. Maack, D. Attanasio, D. Fritz, G. McBrien, and D. Bossi. A review of lithium niobate modulators for fiber-optic communications systems. *J. Select. Topics in Quantum. Electron.*, **6**(1), 69 – 82, 2000.

[92] G. Grand, J. P. Jadot, H. Denis, S. Valette, A. Fournier, and A. M. Grouillet. Low-Loss PECVD Silica Channel Waveguides for Optical Communications. *Electron. Lett.*, **26**(25), 2135 – 2137, 1990.

[93] C. H. Henry, G. E. Blonder, and R. F. Kazarinov. Glass Waveguides on Silicon for Highbrid Optical Packaging. *J. Lightwave Technol.*, **7**(10), 1530 – 1539, 1989.

[94] T. Izawa, H. Mori, Y. Murakami, and N. Shimizu. Deposited Silica Waveguide for Integrated Optical Circuits. *Appl. Phys. Lett.*, **38**(7), 483 – 485, 1981.

[95] M. Kawachi, M. Yasu, and T. Edahiro. Fabrication of SiO_2:TiO_2 Glass Planar Optical Waveguides by Flame Hydrolysis Deposition. *Electron. Lett.*, **19**(15), 583 – 584, 1983.

[96] N. Takato, K. Jinguji, M. Yasu, H. Toba, and M. Kawachi. Silica-Based Single-Mode Waveguides on Silicon and Their Application To Guided-Wave Optical Interferometers. *J. Lightwave Technol.*, **6**(6), 1003 – 1010, 1988.

[97] N. Takato, M. Yasu, and M. Kawachi. Low-Loss High-Silica Single-Mode Channel Waveguides. *Electron. Lett.*, **22**(6), 321 – 322, 1986.

[98] N. Nourshargh, E. M. Starr, and T. M. Ong. Integrated Optic 1×4 Splitter in SiO_2:GeO_2. *Electron. Lett.*, **25**(15), 981 – 982, 1989.

[99] T. Kominato, Y. Ohmori, H. Okazaki, and M. Yasu. Very Low-Loss GeO_2-doped Silica Waveguides Fabricated by Flame Hydrolysis Deposition Method. *Electron. Lett.*, **26**, 327 – 328, 1990.

[100] M. A. Duguay, T. L. Koch, and L. Pfeiffer. Antiresonant Reflecting Optical Waveguides in SiO_2-Si Multilayer Structures. *Appl. Phys. Lett.*, **49**(1), 13 – 15, 1986.

[101] B. N. Kurdi and D. G. Hall. Optical Waveguides in Oxygen-Implanted Buried-Oxide Silicon-on-Insulator Structures. *Optics Lett.*, **13**(2), 175 – 177, 1988.

[102] A. G. Rickman, G. T. Reed, and F. Namavar. Silicon-on-Insulator Optical Rib Waveguide Loss and Mode Characteristics. *J. Lightwave Technol.*, **12**(10), 1771 – 1776, 1994.

[103] Bookham Technology, http://www.bookham.com/.

[104] H. I. Smith. *Submicron- and Nanometer-Structures Technology*. NanoStructures Press, 2nd edition, 1994. (notes for MIT course 6.781).

[105] S. Wolf and R. N. Tauber. *Silicon Processing for the VLSI Era*, volume 1. Lattice Press, 1986.

[106] M. V. Bazylenko and M. Gross. Reactive Ion Etching of Silica Structures for Integrated Optics Applications. *J. Vac. Sci. Technol. A*, **14**(6), 2994 –3003, 1996.

[107] P. D. Trinh, S. Yegnanarayanan, and B. Jalali. Integrated Optical Directional Couplers in Silicon-on-Insulator. *Electron. Lett.*, **31**(24), 2097 – 2098, 1995.

[108] P. D. Trinh, S. Yegnanarayanan, F. Coppinger, and B. Jalali. Silicon-on-Insulator (SOI) Phased-Array Wavelength Multi/Demultiplexer with Extremely Low Polarization Sensitivity. *IEEE Photon. Technol. Lett.*, **9**(7), 940 – 942, 1997.

[109] T. Yonehara, K. Sakaguchi, and N. Sato. Epitaxial layer transfer by bond and etch back of porous Si. *Appl. Phys. Lett.*, **64**(16), 2108 – 2110, 1994.

[110] N. Sato and T. Yonehara. Hydrogen annealed silicon-on-insulator. *Appl. Phys. Lett.*, **65**(15), 1924 – 1926, 1994.

[111] E. H. Anderson. *Fabrication and Electromagnetic Applications of Periodic Nanostructures*. Ph.D. thesis, Massachusetts Institute of Technology, May 1988.

[112] M. E. Walsh. *Nanostructuring Magnetic Thin Films Using Interference Lithography*. Master's thesis, Massachusetts Institute of Technology, 2000.

[113] G. Meltz, W. W. Morey, and W. H. Glenn. Formation of Bragg Gratings and Optical Fibers by a Transverse Holographic Method. *Optics Lett.*, **14**(10), 823 – 826, 1989.

[114] C. G. Askins, T.-E. Tsai, G. M. Williams, M. A. Putnam, M. Bashkansky, and E. J. Friebele. Fiber Bragg Reflectors Prepared by a Single Excimer Pulse. *Optics Lett.*, **17**(11), 833 – 835, 1992.

[115] H. G. Limberger, P. Y. Fonjallaz, and R. P. Salathe. Spectral Characterization of Photoinduced High Efficient Bragg Gratings in Standard Telecommunication Fibers. *Electron. Lett.*, **29**(1), 47 – 49, 1993.

[116] P. Lambelet, p. Y. Fonjallaz, H. G. Limberger, R. P. Salathe, C. Zimmer, and H. H. Gilgen. Bragg Grating Characterization by Optical Low-Coherence Reflectometry. *IEEE Photon. Technol. Lett.*, **5**(5), 565 – 567, 1993.

[117] J. Ferrera. *Nanometer-Scale Placement in Electron-Beam Lithography*. Ph.D. thesis, Massachusetts Institute of Technology, June 2000.

[118] R. Ghanbari. *Physics and Fabrication of Quasi-one-Dimensional Conductors*. Ph.D. thesis, Massachusetts Institute of Technology, February 1993.

[119] J. N. Damask, V. V. Wong, J. Ferrera, H. I. Smith, and H. A. Haus. High-Coherence QWS Gratings for Optoelectronic Devices: Why Spatial-Phase-Locked E-Beam Lithography is Necessary. In *Optical Fiber Communications Conference*. February 20–25 1994.

[120] J. Ferrera. *Highly Coherent Gratings for Optoelectronics: An Application of Spatial-Phase-Locked Electron Beam Lithography*. Master's thesis, Massachusetts Institute of Technology, May 1994.

[121] J. Ferrera, V. V. Wong, S. Rishton, V. Boegli, E. H. Anderson, D. P. Kern, and H. I. Smith. Spatial-Phase-Locked Electron-Beam Lithography: Initial Test Results. *J. Vac. Sci. Technol. B*, **11**, 2342 – 2347, 1993.

[122] V. V. Wong, J. Ferrera, J. Damask, J. Carter, E. Moon, H. A. Haus, H. I. Smith, and S. Rishton. Spatial-Phase Locked E-Beam Lithography and X-ray Lithography for Fabricating First-Order Gratings on Rib Waveguides. *J. Vac. Sci. Technol. B*, **12**, 3741 – 3745, 1994.

[123] A. M. Hawryluk. *Transmission Diffraction Gratings for Soft X-ray Spectroscopy and Spatial Period Division*. Ph.D. thesis, Massachusetts Institute of Technology, October 1981.

[124] K. O. Hill, B. Malo, F. Bilodeau, D. C. Johnson, and J. Albert. Bragg Gratings Fabricated in Monomode Photosensitive Optical Fiber by UV Exposure through a Phase Mask. *Appl. Phys. Lett.*, **62**(10), 1035 – 1037, 1993.

[125] J. G. Goodberlet. Patterning 100 nm features using deep-ultraviolet contact photolithography. *Appl. Phys. Lett.*, **76**(6), 667 – 669, 2000.

[126] S. Y. Chou. Nanoscale silicon field effect transistors fabricated using imprint lithography. *Appl. Phys. Lett.*, **71**(21), 1881 – 1883, 1997.

[127] K. Aiki, M. Nakamura, and J. Umeda. Lasing Characteristics of Distributed-Feedback GaAs-GaAlAs Diode Lasers with Separate Optical and Carrier Confinement. *J. Quantum Electron.*, **QE-12**(10), 597 – 603, 1976.

[128] K. Aiki, M. Nakamura, and J. Umeda. A Frequency-Multiplexing Light Source with Monolithically Integrated Distributed-Feedback Diode Lasers. *J. Quantum Electron.*, **QE-13**(4), 220 – 223, 1977.

[129] H. Soda, Y. Kotaki, H. Sudo, H. Ishikawa, S. Yamakoshi, and H. Imai. Stability in Single Longitudinal Mode Operation in GaInAsp / InP Phase-Adjusted DFB Lasers. *J. Quantum Electron.*, **QE-23**(6), 804 – 814, 1987.

[130] H. A. Haus and C. V. Shank. Antisymmetric Taper of Distributed Feedback Lasers. *J. Quantum Electron.*, **QE-12**(9), 532 – 539, 1976.

[131] C. H. Henry, R. F. Kazarinov, H. J. Lee, N. A. Olsson, and K. J. Orlowsky. A Narrow-Band Si_3N_4-SiO_2 Resident Optical Reflector. *J. Quantum Electron.*, **QE-23**(9), 1426 – 1428, 1987.

[132] G. D. Maxwell, R. Kashyap, B. J. Ainslie, D. L. Williams, and J. R. Armitage. UV Written 1.5 μm Reflection Filters in Single Mode Planar Silica Guides. *Electron. Lett.*, **28**(22), 2106 – 2107, 1992.

[133] F. Bilodeau, B. Malo, J. Albert, D. C. Johnson, K. O. Hill, Y. Hibino, M. Abe, and M. Kawachi. Photosensitization of Optical Fiber and Silica-on-Silicon / Silica Waveguides. *Optics Lett.*, **18**(12), 953 – 955, 1993.

[134] T. A. Strasser, T. Erdogan, A. E. White, V. Mizrahi, and P. J. Lamaire. Ultraviolet Laser Fabrication of Strong, Nearly Polarization-Independent Bragg Reflectors in Germanium-Doped Silica Waveguides On Silica Substrates. *Appl. Phys. Lett.*, **65**(26), 3308 – 3310, 1994.

[135] J. Albert, F. Bilodeau, D. C. Johnson, K. O. Hill, S. J. Mihailov, D. Strykman, T. Kitagawa, and Y. Hibino. Polarization-Independent Strong Bragg Gratings in Planar Lightwave Circuits. *Electron. Lett.*, **34**(5), 485 – 486, 1998.

[136] A. Saliminia, A. Villeneuve, T. V. Galstyan, S. LaRochelle, and K. Richardson. First- and Second-Order Bragg Gratings in Single-Mode Planar Waveguides of Chalcogenide Glasses. *J. Lightwave Technol.*, **17**(5), 837 – 842, 1999.

[137] R. Kashyap, G. D. Maxwell, and B. J. Ainslie. Laser-Trimmed Four-Port Bandpass Filter Fabricated in Single-Mode Photosensitive Germanium-Doped Planar Waveguide. *IEEE Photon. Technol. Lett.*, **5**(2), 191 – 194, 1993.

[138] T. Kitagawa, F. Bilodeau, B. Malo, S. Theriault, J. Albert, D. C. Johnson, K. O. Hill, K. Hattori, and Y. Hibino. Single-Frequency Er^{3+}-doped Silica-Based Planar Waveguide Laser with Integrated Photo-Imprinted Bragg Reflectors. *Electron. Lett.*, **30**(16), 1311 – 1312, 1994.

[139] T. Erdogan, T. A. Strasser, M. A. Milbrodt, E. J. Laskowski, C. H. Henry, and G. E. Kohnke. Integrated-Optical Mach-Zehnder Add-Drop Filter Fabricated by a Single UV-Induced Grating Exposure. *Appl. Opt.*, **36**(30), 7838 – 7845, 1997.

[140] M. H. Lim, T. E. Murphy, J. Ferrera, J. N. Damask, and H. I. Smith. Fabrication Techniques for Grating-Based Optical Devices. *J. Vac. Sci. Technol. B,* **17**(6), 3208–3211, 1999.

[141] C. Dragone. An $N \times N$ optical multiplexer using a planar arrangement of two star couplers. *IEEE Photon. Technol. Lett.*, **3**(9), 812 – 815, 1991.

[142] R. G. Walker. Simple and Accurate Loss Measurement Technique for Semiconductor Optical Waveguides. *Electron. Lett.*, **21**(13), 581 – 583, 1985.

[143] B. J. Eggleton, P. S. Westbrook, C. A. White, C. Kerbage, R. S. Windeler, and G. L. Burdge. Cladding-Mode Resonances and Error-Silica Microscope Optical Fibers. *J. Lightwave Technol.*, **18**(8), 1084 – 1100, 2000.

[144] J. Foresi, P. Villeneuve, J. Ferrera, E. Thoen, G. Steinmeyer, S. Fan, J. Joannopoulos, L. Kimerling, H. Smith, and E. Ippen. Photonic-Band-Gap Waveguide Microcavities. *Nature*, **390**, 143–145, 1997.

[145] C. Vassallo. Improvement of finite difference methods for step-index optical waveguides. *IEE Proc. Opt.*, **139**(2), 137–142, 1992.

[146] J. Yamauchi, M. Sekiguchi, O. Uchiyama, J. Shibayama, and H. Nakano. Modified Finite-Difference Formula for the Analysis of Semivectorial Modes in Step-Index Optical Waveguides. *IEEE Photon. Technol. Lett.*, **9**(7), 961 – 963, 1997.

[147] G. R. Hadley. Transparent Boundary Conditions for Beam Propagation. *Optics Lett.*, **16**(9), 624 – 626, 1991.

[148] A. S. Sudbo. Why are accurate computations of mode fields in rectangular dialectic waveguides difficult? *J. Lightwave Technol.*, **10**(4), 418 – 419, 1992.