MITIGATION OF LOSS, CROSSTALK, AND RESONANCE-SHIFT FOR SCALABLE SILICON PHOTONIC INTEGRATED CIRCUITS

By

ROBERT DIAZ GATDULA

A dissertation submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Doctor of Philosophy

Graduate Program in

Electrical and Computer Engineering

Written under the direction of

Professor Wei Jiang

And approved by

_____

_____

_____

_____

_____

New Brunswick, New Jersey

October 2019

ABSTRACT OF THE DISSERTATION

Mitigation of Loss, Crosstalk, and Resonance-Shift for Scalable Silicon Photonic Integrated Circuits

by ROBERT DIAZ GATDULA

Dissertation Director:

Professor Wei Jiang

Integrated optical interconnect technology has reached a point in which academia and industry research are dedicating many of its resources to maturing the technology and developing it for large scalability. Silicon photonics and silicon photonic-related technologies are the main contenders in driving scalability due to silicon's advantage in being a successful mature material in other large-scale developments such as in achieving highly dense transistors. Silicon's past success has jumpstarted silicon photonics. However, there are still many hurdles to scaling up.

One of the biggest issues with silicon waveguides is that they can be quite lossy. Despite fabrication processes becoming more advanced, the losses of single-mode silicon waveguides are still several orders worse than that in silica-based optical fibers. Typically, to reach lower losses, designers can widen or thicken silicon waveguides. However, there are various multimode-related caveats that can make this non-trivial. As such, we explore waveguide losses with a multimode perspective and provide some insight to avoid intermodal waveguide scattering that might inadvertently increase loss rather than decrease it in wider waveguides.

Another issue with waveguides is the need to bend them back and forth when routing light to various parts of a photonic integrated circuit. Dense waveguides have been proven to demonstrate low crosstalk in straight sections by creatively engineering the relative width geometry between neighboring waveguides in the form of a waveguide superlattice. However, the bending of waveguide superlattices introduces bending-related physics that can increase the crosstalk. We explore the bending regime and demonstrate dense waveguide superlattice bends of small footprint and relative crosstalk no greater than -19.6 dB for a waveguide superlattice with a minimum bending radius of 5 μm.

We also explore the sensitivities of microring-based transceiver circuits, which can drastically be reduced in performance by fabrication deviations. Furthermore, with ambient conditions such as temperature constantly changing, the performance of microrings can be unstable. More than likely, photonic integrated circuits will be packaged in systems including electronic circuits that can tune the performance of their photonic counterparts. Thus, we provide various gradient-based algorithms to enable automated tuning and thermal adaptivity for multi-microring photonic integrated circuits.

Dedication

For my family

Acknowledgements

This dissertation would not have been possible if it were not for the myriad of people in my life that supported me technically and personally as I went through my academic journey. Great science and engineering are not only the products of novel ideas and gritty work, but also through the many behind-the-scene coincidences that would otherwise not have given that extra push. This includes but is not limited to the inspiration working with others, the scaffolding provided by those who have provided to help jumpstart the work, and the support of those who will guide you down the right path regardless of the trials, tribulations, and diversions.

I would first and foremost thank my dissertation advisor Dr. Wei Jiang for opening my eyes to the perspective and disruptiveness that silicon photonics can provide to the scientific community. He introduced me to the world of photonics when I was a third-year undergraduate student and brought me in as an undergraduate research assistant in his lab. I appreciate his continued support for me as a doctoral student since then and the unique perspective he gave me for each project that I have worked on. In addition, I am grateful for the various connections he has made with other experts in the field. If it were not for those initial connections, I may not have met or worked with other great scientists or engineers.

I also would like to thank Dr. Yicheng Lu, Dr. Mehdi Javanmard, Dr. Ming Lu, and Dr. Michael Caggiano for their support in the completion of my dissertation. Because of their useful insights for science and silicon photonics, I have been inspired to continue my learnings regardless of where my future career may go. Their

questioning and healthy skepticisms of my research are always appreciated and, franky, necessary for me to continue developing impactful ideas.

My experience as a researcher in Dr. Wei Jiang's group would not have been amazing if it were not for the people that I have worked within the group, which include Siamak Abbaslou, Minning Zhu, Jun Tan, Weiwei Song, Amrita Banerjee, Yating Zhou. It has been an unforgettable, unique experience to do science with them, which made each day in the office or lab – whether it was a good day or bad day – worth coming in.

I would specifically like to thank my colleague Siamak Abbaslou, who has been through many of the of the same situations I have at the same time during graduate school. He and I always bounced ideas back and forth, whether or not they were truly good ideas. While I was always busy in the clean room, he was in the lab assisting me with getting simulation and measurement work done when I did not have the time to do it myself. He has also been there to talk about personal matters as well, and it was nice to be able to have time to do that outside the demands of our research work.

While I have not always discussed my research work in detail with them, I am thankful for the opinions of Dr. Michael Caggiano and Dr. Sigrid McAfee. Hearing about their many years of experience in academia and industry research has been eye opening and intriguing.

In times where I was particularly worried about finishing my doctoral degree, I would like to amply thank Dr. Zoran Gajic and Dr. Narayan Mandayam for their support and sympathy in my times of trouble. Their laying out of the necessary steps to progressing in a clear, understandable manner gave me solace that I have not

diverged too far from each stepping stone. I would also like to give another set of thanks to Dr. Yicheng Lu for his likewise support.

I would like to thank the current and former administrative and technical staff in both the Electrical & Computer Engineering Department and the Institute for Advanced Materials, Devices, and Nanotechnology, who include John Scafidi, John McCarthy, Noraida Martinez, Steve Orbine, Robert Lorber, Arletta Hoscilowicz, Christy Lafferty, Tea Akins, Mayra Howell, and Nancy Pamula. I am sure that there would be times that would otherwise been troublesome without any of their expertise or knowledge.

Out of all the other current and former graduate students I had the pleasure to interact with, I thank Wen-Chiang Hong, Kliti Kodra, Yanbiao Pan, Sourav Sagar, Manonit Kumar, and Maria Peifer. It was always great to discuss our various research interests as well as know what is going on in each other's lives.

From Brookhaven National Laboratory, I would like to thank Dr. Ming Lu and Dr. Aaron Stein for their constant support in nanofabrication. None of the devices that I have made in my research have not passed through the various lithography, etching, and deposition instruments that these two have overseen in their clean room. If there was a process, equipment, or chemical that I needed, they would figure out how to best assist me.

From Sandia National Laboratory, I would like to thank my summer internship advisors Christopher DeRose and Patrick Chu. They have introduced me to so many interesting and applicable silicon photonic research in their labs that I probably would not have experienced in any other group.

During my co-op internship at Nokia Bell Labs, I have had to pleasure to work with so many different people working on various photonic devices and system research. This includes Dr. Po Dong, Dr. Kwangwoong Kim, Dr. Argishti Melikyan, Dr. Chandrasekhar Sethumadhavan, Dr. Brian Eggleston, Yishen Huang, Dr. Guilhem DeValicourt, Dr. Jeffrey Sinsky, Dr. Chia-Ming Chang, and Dr. Young-Kai Chen. Working with these people with their various expertise have been an undoubtedly unique experience. I have made both great research colleagues and friends from there, and during the entirety of my internship there, I found it to be my second academic home away from Rutgers University.

For all the times I had the chance to spend time with them, I would like to thank my friends Shoham Bhadra, Christopher Chen, Danielle Walsh, David Walsh, Michael Zlotkowski, Grittney Tam, Dashmeet Singh, Kelly Lee, Grant Junno, Haley Gittleman, and Calvin Woo. Whether it is going out swing dancing, rock climbing, having holiday dinners, playing board games, going out to dinner or drinks, having potluck lunches or dinners, singing karaoke, or playing video games, I appreciate spending the time with them and experiencing a little bit of life outside of my graduate research.

I would like to dedicate a very special thanks to my dearest friend Julia Burmistrova. There are not enough words to describe how much she has been there for me, making sure that I did not drown in my work, helping me provide balance to my busy life – pretty much on a day-to-day basis – and entertaining my engineering aptitude (She was an engineering student as well). I not only had a lot of self-discovery during the efforts of my research, but I also found out a lot about myself in the many times she and I have interacted, whether it was talking about our personal

lives and beliefs, eating food and cooking, having empathy and caring for others, raising pets, traveling to other countries, swing dancing and being part of the dance community, teaching and competing in dancing, taking care of our physical and mental health, discussing politics and policy making – the list goes on and on. It was always enlightening to have Julia around. Sometimes it was quite frustrating when we viewed things differently. However, there is no doubt that she helped me understand myself and inspired me to become a better person in more ways than one.

Lastly, but not least, I would like to thank the never-ending and loving support from my parents and sisters. I was lucky to be within a half-hour drive from them and always have welcoming arms whenever I entered the house, whether it was just things I needed for my apartment or during times when I wanted to spend quality time with them. They have seen my life progress since the very beginning, helping me grow as an adult and encouraging me in whatever unique or unorthodox way to continue exploring my curiosity in science and engineering.

Much of this work is supported in part by DARPA Young Faculty Award under Grant No. N6600112-1-4246. This research is carried out in part at the Center for Functional Nanomaterials, which is a U.S. Department of Energy Office of Science Facility, at Brookhaven National Laboratory under Contract No. DE-SC0012704, in part by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-08-1-0394, and in part by the U.S. Department of Energy within the Center for Functional Nanomaterials, Brookhaven National Laboratory, Office of Basic Energy Sciences under Grant DE-AC02-98CH10886.

Table of Contents

List of Abbreviations

3D-FDTD	Three-dimensional Finite Difference Time Domain
ADAGRAD	Adaptive Gradient
ADC	Analog to Digital Converter
AWG	Arrayed Waveguide Grating
BER	Bit Error Rate/Ratio
CMOS	Complementary-Metal-Oxide-Semiconductor
CMP	Chemical-Mechanical Polishing
CPU	Central Processing Unit
CVD	Chemical Vapor Deposition
DAC	Digital to Analog Converter
DC	Direct Current
DIW	Deionized Water
DMT	Discrete Multi-Tone
DRC	Design Rule Check
DSL	Digital Subscriber Line
DSP	Digital Signal Processing
DWDM	Dense Wavelength Division Multiplexing
DUT	Device Under Test
EDFA	Erbium-Doped Fiber Amplifier
ESW	Equivalent Straight Waveguide
FDTD	Finite Difference Time Domain
FEC	Forward Error Correction
FPGA	Field Programmable Gate Array
FSR	Free Spectral Range
FTTB	Fiber-To-The-Building
FTTH	Fiber-To-The-Home
FTTN	Fiber-To-The-Node/Neighborhood
FTTX	Fiber-To-The-X
GS	Gerchberg-Saxton
HIO	Hybrid Input-Output
HSQ	Hydrogen-Silsesquioxane
IPA	Isopropyl Alcohol
LiDAR	Light Detection and Ranging
LOQC	Linear Optical Quantum Computing
MPD	Monitoring Photodiode/Photodetector
MZI	Mach-Zehnder Interferometer
NRZ	Non-Return-to-Zero
OEIC	Optoelectronic Integrated Circuit

OMA	Optical Modulation Amplitude
ONN	Optical Neural Network
OOK	On-Off Keying
OPA	Optical Phased Array
OVA	Optical Vector Analyzer
PAM	Pulse Amplitude Modulation
PBC	Polarization Beam Combiner
PBS	Polarization Beam Splitter
PBSR	Polarization Beam Splitter-Rotator
PD	Photodiode/Photodetector
PDK	Photonics Design Kit
PECVD	Plasma-Enhanced Chemical Vapor Deposition
PIC	Photonic Integrated Circuit
PID	Proportional-Integral-Derivative
PRBS	Pseudorandom binary sequence
PS	Phase Shifter
PSO	Particle Swarm Optimization
QAM	Quadrature Amplitude Modulation
RF	Radio Frequency
RIE	Reactive Ion Etching
RMS	Root-Mean-Squared
SC	Supercell
SDM	Space Division Multiplexing
SEM	Scanning Electron Microscope
SNR	Signal To Noise Ratio
SOI	Silicon-On-Insulator
SPSO	Standard Particle Swarm Optimization
SWI	Swept Wavelength Interferometry
TE	Transverse Electric
TEC	Thermoelectric Cooler
TIA	Transimpedance Amplifier
TM	Transverse Magnetic
VOA	Variable Optical Attenuator
WDM	Wavelength Division Multiplexing
WG	Waveguide
WGSL	Waveguide Superlattice

List of Tables

List of Illustrations

Chapter 1: Introduction

1.1 Optical interconnects in telecommunication and datacenters

Optical interconnects have often played a crucial role in the rapid development of state-of-the-art applications. One such prime example are the optical interconnects which form the high-speed backbone of telecommunications: optical fibers. By sending data in the form of short light pulses in low loss glass fibers, telecommunication technologies with fiber-optics have helped enable modern society to broadcast live video, stream multimedia, and store terabytes upon terabytes of data. This is primarily due to the higher potential data throughput of optical fibers compared to more traditional communication mediums such as copper cables. As illustrated in **Figure 1.1**, optical fiber deployment for broadband network architecture is generically described as fiber-to-the-X (FTTX), where X corresponds to the distance of the optical fibers to the end user. As years of optical fiber development have gone by, deployment has gone from the node/neighborhood (FTTN) to as close as the home/building (FTTH/FTTB).

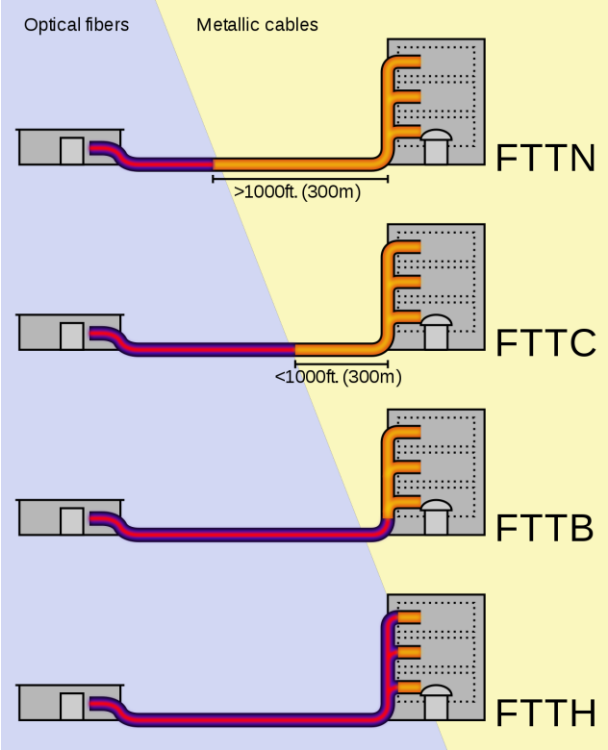

Figure 1.1 Illustration of various FTTX deployments for telecommunication applications.

In conjunction with FTTX deployment, optical transceivers – which convert light pulses into electrical pulses and vice versa – have become smaller. This is due to the crucial role of integrated optical interconnects (or integrated optics). These integrated optics systems have scaled up mainly due to the utilization of readily-available materials that can provide small form factors for fiber optic coupling, modulation, and photodetection. Datacenters, which to consumers are described as "the cloud," have scaled up substantially because of optical fibers and transceivers. In a sense like FTTX, optical communication deployment in datacenters can be described as going from rack-to-rack communication distances to shorter-reach regimes such as board-to-board, chip-to-chip, and ultimately viable intra-chip optics.

Each new deployment can potentially provide higher data throughputs and larger scalability than its predecessors.

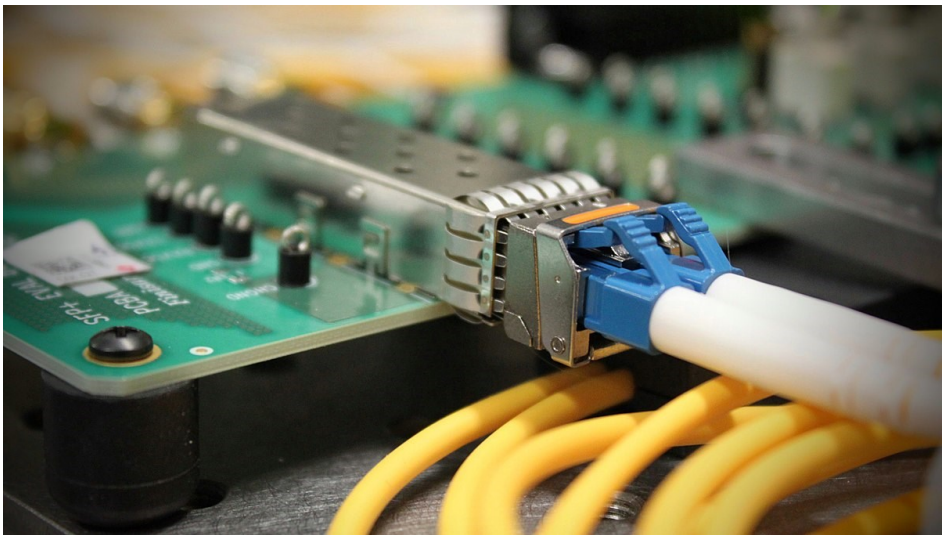

Figure 1.2 An optical transceiver with fiber optic cables connected to its TX (transmitter) and RX (receiver) ports. This module converts high-speed light communication to electrical communication and vice versa. The small form factor of it is attributed to the implementation of compact integrated optics.

1.2 Increasing data capacity through optical multiplexing

Optical communications have gone through many evolutions to become a competitive networking platform compared to traditional metal cabling, from overcoming the losses caused by imperfectly drawn glass fibers to overcoming spectral deficiencies via quadrature modulation. Carrying optical data can be in the form of five physical dimensions as illustrated in **Figure 1.3**: frequency (or wavelength), time, quadrature, polarization, and space [1].

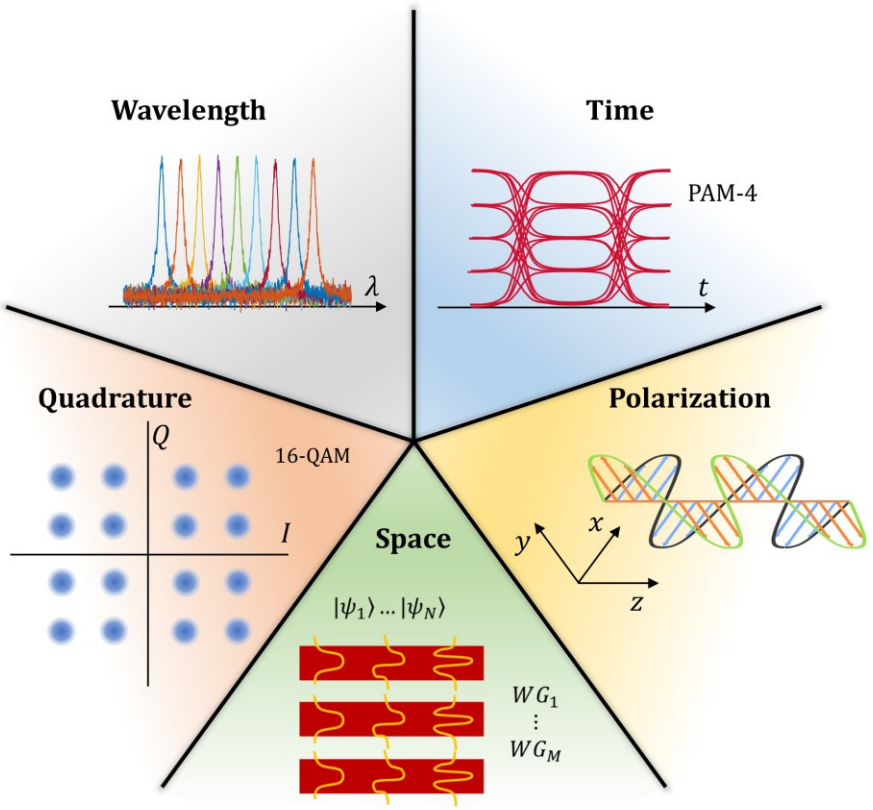

Figure 1.3 The five physical dimensions of optical multiplexing.

Optical communication is driven by a need to decrease signal losses – maintain a high signal-to-noise ratio (SNR) – while transferring as much data as possible. Each technology evolution has pushed the so-called Shannon limit for optical fibers, which describes the highest theoretical capacity limit [2]. At the time of writing this dissertation, academic and commercial research is pushing the Shannon limit via photonic integrated circuits (PICs). For example, researchers are designing short-reach coherent optical communication transceivers with dense wavelength division multiplexing (DWDM), utilizing carrier frequency spacing on the order of 100 GHz or lower. In terms of optical wavelength, this is less than or equal to 0.8 nm channel spacing. As these signals do not interfere with each other, multiple signals can be

propagated through the same communication bus, drastically increasing the potential data throughput. This is also potentially achievable through space division multiplexing (SDM), which exploits orthogonal optical modes in the communication bus. Together with DWDM and advanced modulation formats, terabit-per-second data rates are even more possible. By increasing the density of channels, we can increase the data throughput of optical interconnects.

1.3 **Dense photonic integrated circuits in other applications**

In a similar light, dense PICs can also benefit other applications. One is in free-space communication and sensing using optical phased arrays (OPAs) and their specific application in light detection and ranging (LiDAR). Inspired by radio-based phased array and radar, researchers aim to achieve small form factor systems with optical wavelength versions. This is possible because OPAs are essentially a form of light diffraction, allowing solid-state beamforming and beam steering. With denser apertures, beam quality increases, maximum beam steering angle increases, and device efficiency increases. Furthermore, OPAs can allow arbitrary beam forming and sensing with applications such as holography and optical fingerprinting in mind. [3]–[6]

Another interesting application desiring dense PICs is deep learning, more specifically, optical neural networks (ONN). ONNs are a physical form of artificial neural networks which can provide solutions to difficult problems via machine learning algorithms. One way to incorporate ONNs is through digital or analog optical switch fabrics. By increasing the number of devices, we can increase the number of

neural network nodes, and therefore increase the optical computing power of the ONN. [7]

Dense PICs can also benefit quantum computing systems in the form of linear optical quantum computing (LOQC). LOQC is the idea of using linear optical elements such as beam splitters, phase shifters, and mirrors to process quantum information. Quantum bits (qubits) can be carried via photons with probabilities of being in a certain optical mode, frequency, or both. Increasing the number of processed qubits via more and denser optical elements can drastically provide larger quantum computing power. [8], [9]

1.4 **Silicon photonics: a disruptive technology**

1.4.1 Why silicon photonics?

A promising material for these photonic applications is silicon. It has been widely used as the material of choice in the microelectronic industry, especially in transistors utilizing complementary-metal-oxide-semiconductor (CMOS) technology. Similarly, silicon has shown much potential as the choice material in photonics. Due to crystalline silicon's high refractive index (n = 3.45) in comparison to both air (n = 1.00) and silicon dioxide (a.k.a. silica, n = 1.45), light in the infrared regime can be easily confined to make PICs. However, silicon photonics incorporates more materials than just crystalline silicon and silica. The technology covers several other forms of silicon, including amorphous silicon, silicon nitride, and silicon oxynitride.

Silicon photonics also has comparatively low signal degradation and high bandwidth to fiber optics, making the implementation of these two technologies

extremely easy. In addition to its optical properties, silicon is shown to be the most economically viable material, allowing for development in sophisticated technologies and largescale integration at an affordable price. With the ease of interfacing with fiber optic networks and the compatibility of the material to standard microelectronic processes, silicon photonics can potentially be in lieu of microelectronics or be integrated to have hybrid optoelectronic integrated circuits (OEICs) – either with back-end or front-end implementation [10].

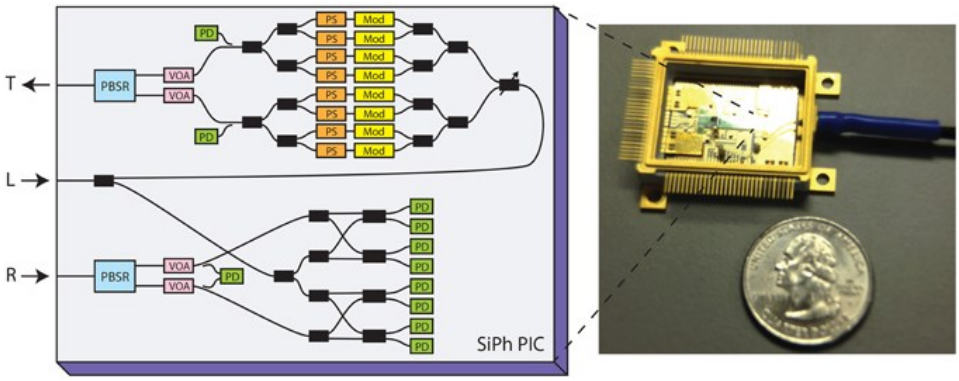

Figure 1.4 A silicon photonics transceiver consisting of polarization beam splitter-rotators, variable optical attenuators, phase shifters, modulators, Y-splitters/combiners, and photodetectors. "L" corresponds to the light source input, "T" corresponds to the transmitted signal, and "R" corresponds to the received signal.

1.4.2 Academic research

Academically, silicon photonics has established its strong feasibility in wavelength division multiplexing (WDM) with various implementations of on-chip optical modulators, filters, and routers [11]–[14]. Furthermore, there is already plenty of development in SDM, specifically mode division multiplexing, polarization division multiplexing, and orbital angular momentum coupling [15]–[23].

1.4.3 Industry research & development

Silicon photonic systems have also already been proven to be a commercial success from the likes of Intel, Luxtera, and Acacia Communications – to name a few – as they are a part of the transceiver (transmitter/receiver) package, reaching data rates on the order of 100 Gb/s. It has the potential to go to shorter reach communication regimes such as intra-chip (in the form of interposers via chip stacking) or inter-chip distances. Ultimately, silicon photonics will enable the high-speed data-rates shown in fiber-optics to be realized in PICs with much room for improvement. In addition, due to its potential to be densely integrated, silicon photonics can provide compact electronic controls and optical antennas for the large beam steering angles needed in OPAs. OPA development with silicon has gained much traction in recent years with its own handful of end-fire and surface-emitting implementations [3], [5], [24]–[30].

1.4.4 Scaling up through wafer-level fabrication and the foundry model

Commercial success and growth of silicon photonics would not be possible if it were not for the foundry model, meaning that PIC fabrication is not done internally in a research group or company but instead done in an external foundry [31]. By leveraging state-of-the-art fabrication facilities, silicon photonic engineers can focus on designing optimal devices and systems. Often, these foundries provide a photonics design kit (PDK) and design rule check (DRC) tables, which allow engineers to shorten the time between system conception and actual physical realization. A handful of

these foundries include AIM Photonics, the Institute of Microelectronics (IME), IMEC, Applied Nanotools (ANT), Leti, TowerJazz, Ligentec, and Taiwan Semiconductor Manufacturing Company (TSMC), ranging from singular die made patterned with e-beam lithography to more large-scale 300 mm wafers patterned with optical lithography. Device manufacturing capabilities range from passive devices (e.g., waveguides, heaters, optical filters) to active devices (e.g., optical modulators, detectors, switching networks).

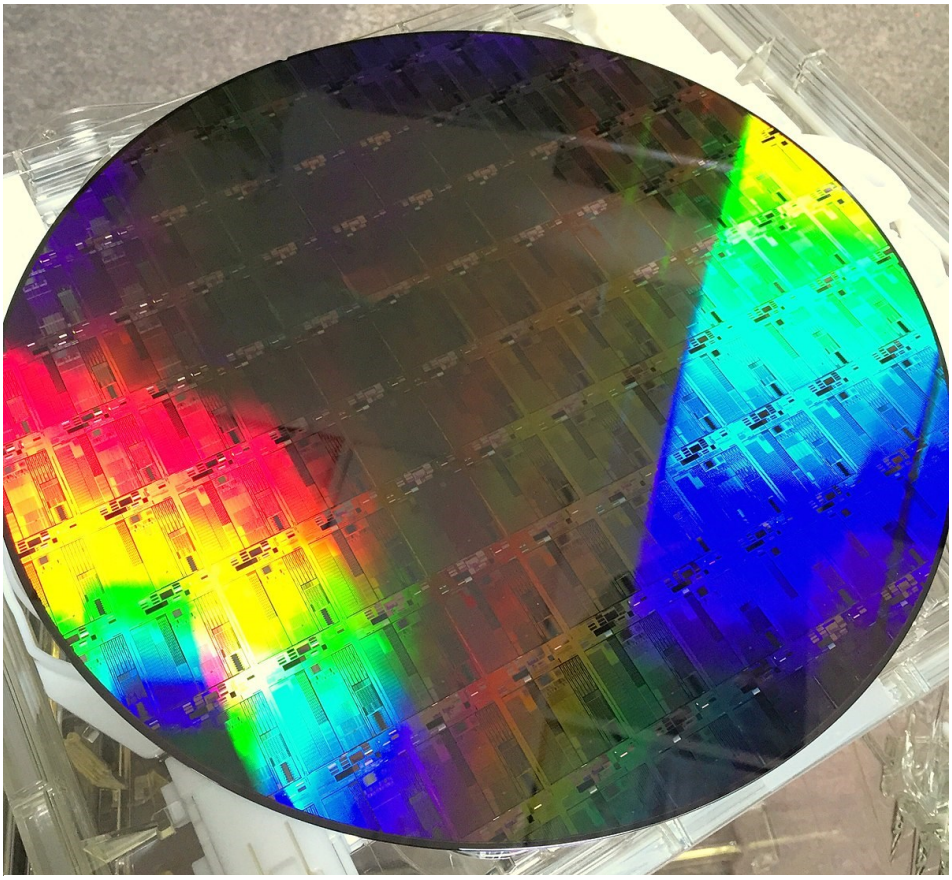

Figure 1.5 Silicon photonics manufactured on a 300 mm wafer platform.

1.5 **Organization of the dissertation**

Despite the current state of silicon photonics, there is still much room for improvement towards large scale implementations of silicon PICs. In this proposal, we show possible avenues for silicon photonics towards scalability. Specifically, we focus on studying the micro-scale edge roughness of real photonic waveguides and what we can exploit from such information, improving the density of on-chip optical interconnects through the idea of a waveguide superlattice and devising closed-loop algorithms for adaptable silicon photonic devices.

Chapter 2: Waveguide Loss from Sidewall Roughness

2.1 Background and motivation

In theory and design, waveguides are perfect. However, in real photonic waveguides, imperfections are inevitable, resulting in optical power loss. In fact, if it were not for the ability to manufacture low-loss, single mode optical fibers, fiber optic communications would not have been able to achieve its competitive state in the telecommunications industry. Overcoming the loss obstacle was the first step in optical fibers to becoming a scalable technology platform. Bringing that level of low loss to an integrated optical platform is a challenge.

We should understand why there is loss in the first place to effectively reduce it through state-of-the-art engineering design and manufacturing. Most PICs utilize semiconductor materials for the sake of monolithic optical waveguiding and modulation via free-carrier injection (also known as the plasma dispersion effect) or carrier depletion. Thus, losses due to free-carrier absorption is evident, especially in highly-doped structures such as optical modulators and filters – but not so much in passive structures with low doping concentrations. The high carrier-concentration in metals is also a concern for waveguide loss, especially if metal contacts or vias are near the waveguiding structures. In general, avoiding this is quite trivial, especially when using planar processes supporting multi-layer fabrication. In addition, surface-states – electronic states found out the interface of materials – can provide additional loss without proper fabrication considerations. Due to the nature of multi-layer

fabrication, substrate leakage is a source of optical loss, as it may support multiple planar waveguide modes that may overlap with our intended guided modes.

Lastly, but certainly not least, are losses due to the surface roughness of the waveguiding materials. This source of optical loss will be the focus of this chapter: knowing the causes of roughness, modelling it appropriately, and proposing a methodology to reduce loss induced by roughness. In our examination of optical loss, we focus on silicon waveguides and go under the assumption that surface roughness is a ubiquitous source of loss.

2.2 **Causes of sidewall roughness in planar waveguides**

Surface roughness in planar waveguides can be caused by a multitude of processes as PICs go through their fabrication journey – see **Figure 2.1**. Most of silicon photonic devices are fabricated on such planar processes, which means there are steps to minimize top and bottom surface roughness such as chemical-mechanical polishing (CMP). In addition, the blank SOI wafers most silicon photonics have been manufactured to have atomic-scale top and bottom roughness so, from the get-go, losses in the vertical direction are minimal for most silicon photonic applications [32]. Minimizing sidewall roughness is much more challenging, however, despite the capabilities of high-resolution lithography and smooth dry etching. The micro-dimensions of structures can be controlled only up to specified tolerance levels.

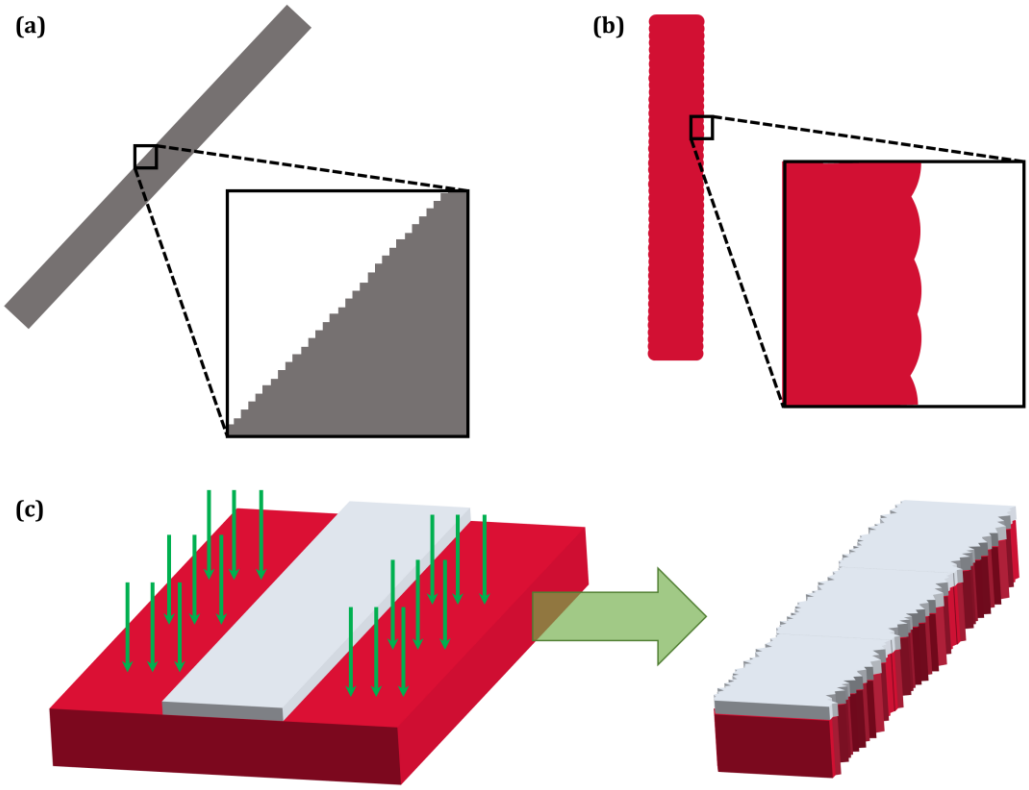

Figure 2.1 Roughness induced by (a) optical lithography mask manufacturing, (b) electron-beam lithography, and (c) dry etching.

Two common lithography techniques for microphotonic structures are optical lithography and electron beam (e-beam) lithography. Optical lithography masks may be patterned by e-beam lithography (sometimes direct laser writing lithography), and such masks are manufactured within a certain roughness tolerance. Thus, patterning would result in those imperfections transferring to the photoresist. E-beam lithography has issues of electron forward scattering and backscattering, resulting in overdosing or underdosing of e-beam resists which would entail a rough pattern, and likewise would be transferred to the substrate when etched. In dry etching, the bombardment of high-energy ions can be slightly random, too. Thus, even if one has a perfect lithography patterning step, sidewall roughness can appear after the etching step.

2.3 **Theoretical approaches to waveguide loss**

Waveguide scattering theory with dielectric materials was originally defined as a relation of optical coupling between the fundamental mode to high-order radiation modes [33]; however, it is an extensive theory and ultimately may not be useful for intelligently engineering low-loss waveguide structures. Since then, there have been several alternative approaches to theoretically describe waveguide loss in slab and rectangular dielectric waveguides. Often-referenced theories are closed-form solutions of radiation loss and backscattering loss [34]–[36]. These theories are based on correlated surface roughness and the waveguide parameters obtained from the weakly-guiding approximation. Other theories include three-dimensional analyses based on the volume current method [37], expanding radiation modes as a series of cylindrical harmonics for high-index contrast waveguides [38], a combined analysis of radiation loss and light leakage towards the substrate in silicon-on-insulator (SOI) square waveguides [39], and a unified theory for radiation and backscattering loss based on the gradient of the waveguide effective index with respect to its width [40].

2.4 **Single mode perspective of waveguide loss**

The theoretical approaches assume a single mode waveguide design. Thus, all other modes are either radiation modes or the backwards propagating counterparts. If we categorize these optical waveguides as weakly guiding, meaning that part of the guided optical modes is present in the low-index cladding material, the surface variation in the core-cladding interface will result in overlap between evanescent tails

of the optical modes and radiation modes. As a result, there would be loss optical loss. Payne and Lacey in [35] propose a closed-form solution, modeling the waveguide sidewall roughness as depicted in **Figure 2.2(a)**.

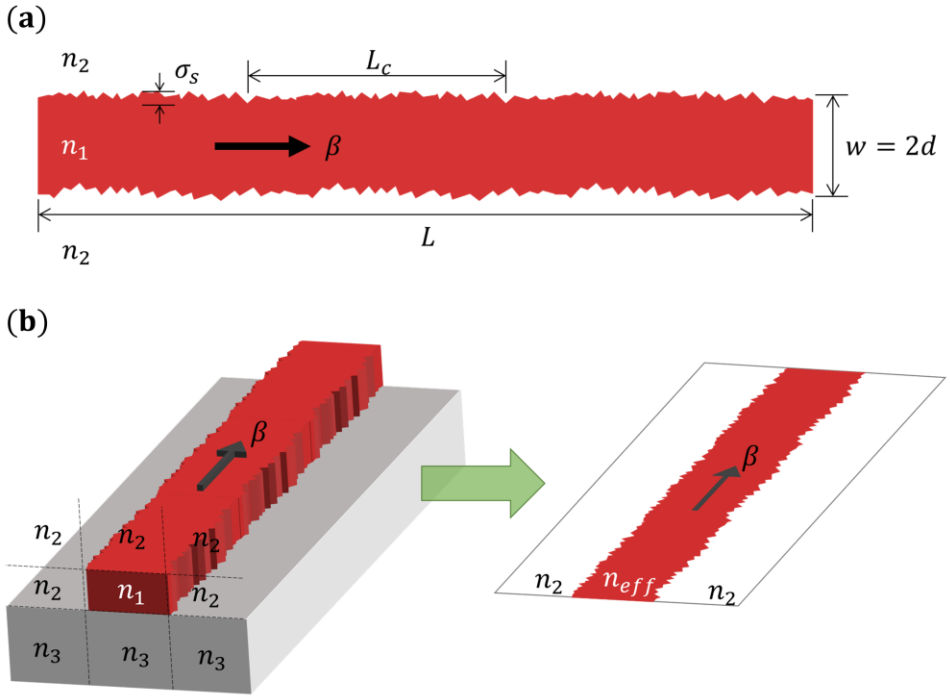

Figure 2.2 (a) Two-dimensional schematic of a waveguide with sidewall roughness having core index n_1 and cladding index n_2. The thick black arrow indicates the propagation direction. (b) A three-dimensional channel waveguide can be approximated into an effective two-dimensional waveguide using the effective index method. The refractive index of the substrate n_3 is taken into consideration when the waveguide is collapsed to an effective index n_{eff}.

Sidewall surfaces are described to have two roughness parameters. The standard deviation, denoted as σ_s, describes the parallel surface variations with respect to the normal vector of the mean surface plane. This is occasionally described as the root-mean-squared (RMS) roughness and will be mentioned as such from here on. The correlation length L_c describes the perpendicular variations. It is equal to the lag length of the surface where its autocorrelation function is e^{-1}. More than likely, the

random variations between both sides of the waveguide will have the same roughness parameters but would be decorrelated, thus the radiation loss α_r is calculated as

$$\alpha_r = \frac{\sigma_s^2}{\sqrt{2}k_0 d^4 n_1} g(V) f_e(x, \gamma), \tag{2.1}$$

where constant k_0 is the free-space wavenumber. The function $g(V)$ is determined by the waveguide geometry and is defined as

$$g(V) = \frac{U^2 V^2}{(1+W)}, \tag{2.2}$$

where the dimensionless waveguide parameters U, V, and W are defined as

$$U = d\sqrt{n_1^2 k_0^2 - \beta^2}, \tag{2.3}$$

$$V = k_0 d\sqrt{n_1^2 - n_2^2}, \tag{2.4}$$

and

$$W = d\sqrt{\beta^2 - n_2^2 k_0^2}. \tag{2.5}$$

The constant β is the propagation constant of the waveguide mode. Function $f_e(x,\gamma)$ is defined as

$$f_e(x,\gamma) = \frac{x\left\{[(1+\chi^2)^2 + 2\chi^2\gamma^2]^{1/2} + 1 - \chi^2\right\}^{1/2}}{[(1+\chi^2)^2 + 2\chi^2\gamma^2]^{1/2}}, \tag{2.6}$$

where dimensionless parameters x and γ are defined as

$$\chi = W\frac{L_c}{d} \tag{2.7}$$

and

$$\gamma = \frac{n_2 V}{n_1 W\sqrt{\frac{n_1^2 - n_2^2}{2n_1^2}}}. \tag{2.8}$$

Waveguide loss also comes from roughness-induced backscattering. In essence, a waveguide with sidewall roughness can be reduced to small sections of butt-coupled waveguides with different effective indices, resulting in multiple points of reflection and therefore backscattering loss. Ladouceur and Paladian in [36] calculate this backscattering loss α_b as

$$\alpha_b = \left[\frac{U^2 W}{2d^3\beta(1+W)}\right]^2 \frac{\sigma_s^2 L_c}{\pi}\frac{1}{1+4\beta^2 L_c^2}. \tag{2.9}$$

Therefore, the total propagation loss for a waveguide with sidewall roughness can be estimated as a summation of (2.1) and (2.9). For some waveguide of length L, insertion loss in dB is defined as

$$IL = 10\,log_{10}\{exp[-(\alpha_r + \alpha_b)L]\}. \tag{2.10}$$

These closed-form solutions are based on a two-dimensional model of a waveguide, but they can represent three-dimensional waveguides via the effective index method – see **Figure 2.1(b)** – collapsing the vertical dimension into an effectively horizontal two-dimensional waveguide. When calculating (2.1) for a three-dimensional waveguide, we would be setting $n_1 = n_{eff}$. The effective index in the collapsed model would be representative of a specific slab mode and not all the possible modes. **Figure 2.3** depicts a surface plot of the propagation loss from a waveguide with varying RMS roughness and correlation lengths at $\lambda_0 = 1550$ nm using the effective index method. Clearly, if the RMS roughness is significantly small, we can reduce the propagation loss. Interestingly, if the correlation length small or large, the propagation loss can also be reduced, and there is a specific correlation length in which we have maximum loss for any arbitrary RMS roughness.

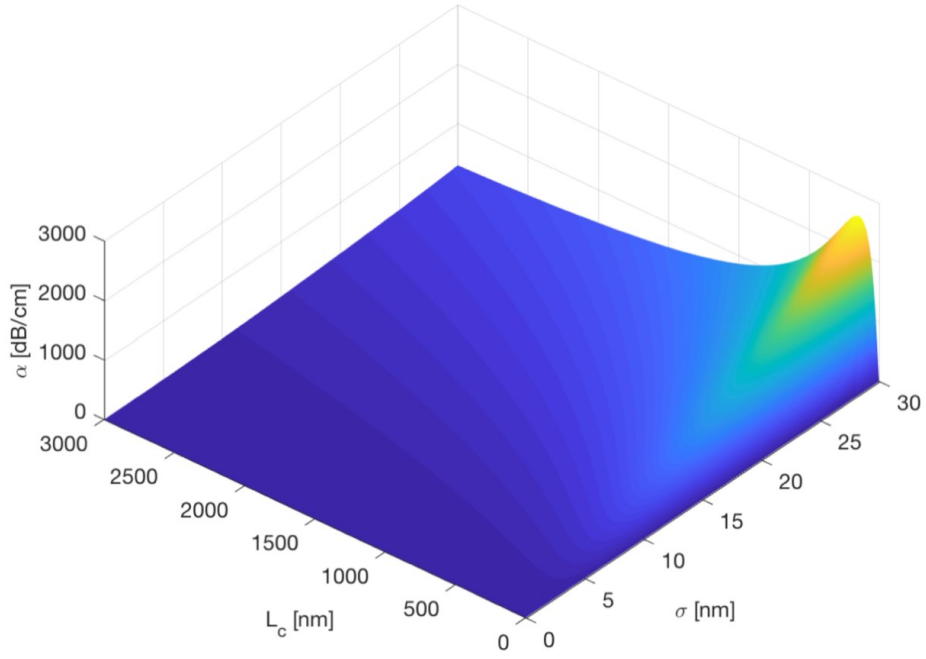

Figure 2.3 Propagation loss of a 260 nm × 420 nm (height × width) silicon waveguide at various sidewall roughness parameters. The waveguide is cladded with silica.

To further simplify the analysis of waveguide loss, we define a correlation length $L_{c,bound}$, which would mark the largest amount of waveguide loss. The bounds of (2.1) and (2.9) are defined at different correlation lengths. However, radiation loss is the primary source of loss, so the largest amount of loss would occur at correlation length

$$L_{c,bound} = \frac{1}{\sqrt{3}(\beta - n_2 k_0)}. \tag{2.11}$$

This correlation length is comparable to the beat length between guided modes and radiation modes.

Up until this point, we are assuming that only a single mode is propagating through the waveguide. However, a rectangular waveguide supports at least one TE mode and one TM mode regardless of its size. Furthermore, if we are working with wide or tall waveguides, we need to consider the loss for each excited guided mode.

2.5 **Multimode perspective of waveguide loss**

2.5.1 Modeling multimode waveguide loss

Loss from a multimode perspective can be represented by a linear combination of intermodal coupling and scattering loss. Suppose we take a slice of a roughened waveguide that might support *M* guided modes from length *L* to *L+dL* as visualized in **Figure 2.4**.

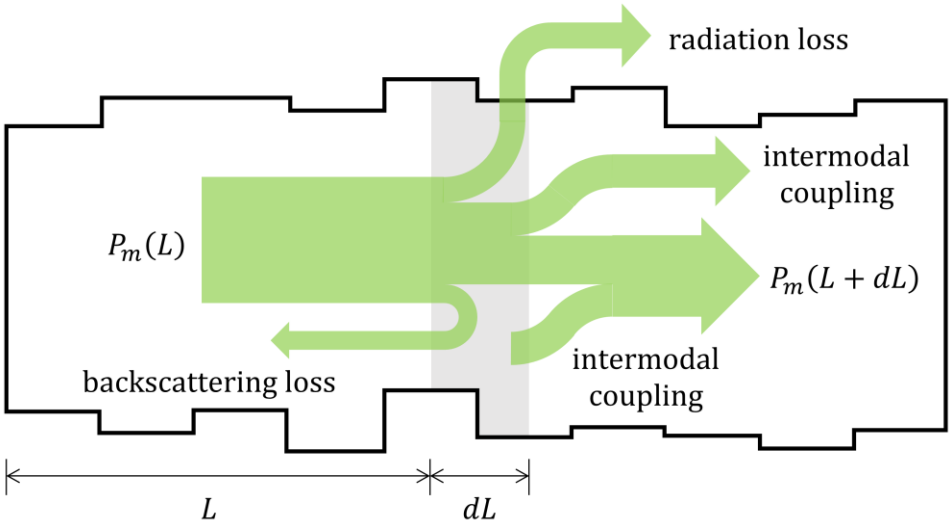

Figure 2.4 Various coupling mechanisms in a multimode waveguide for the mth guided mode at some slice in a roughened waveguide. Intermodal coupling accounts for coupling to and from other guided modes. If the waveguide is smooth and there is no material absorption, we would expect no loss or intermodal coupling, i.e. $P_m(L+dL) = P_m(L)$.

The output power P_m for the $(m-1)^{\text{th}}$ mode, where $m = \{1, \ldots, M\}$, and the total loss P_α, which represents both radiation and backscattering loss, can be summarized as

$$\begin{bmatrix} P_0(L+dL) \\ \vdots \\ P_M(L+dL) \\ P_\alpha(L+dL) \end{bmatrix} = \begin{bmatrix} |\kappa_{00}(L+dL)|^2 & \cdots & |\kappa_{0N}(L+dL)|^2 \\ \vdots & \ddots & \vdots \\ |\kappa_{M0}(L+dL)|^2 & \cdots & |\kappa_{MN}(L+dL)|^2 \\ \alpha_0(L+dL) & \cdots & \alpha_N(L+dL) \end{bmatrix} \begin{bmatrix} P_0(L) \\ \vdots \\ P_M(L) \end{bmatrix}, \tag{2.12}$$

where $N = M$, $n = \{1, \ldots, N\}$, $|\kappa_{mn}|^2$ represents the intermodal coupling constant from the $(n-1)^{\text{th}}$ mode to the $(m-1)^{\text{th}}$ mode and α_n represents the loss constant for the $(n-1)^{\text{th}}$ mode. To maintain energy conservation, the sum of each column in the coupling matrix should be

$$|\kappa_{0n}|^2 + \cdots + |\kappa_{Mn}|^2 + \alpha_n = 1. \tag{2.13}$$

Clearly, there could be coupling between guided modes, where each mode has its own waveguide loss constant. Most waveguides are designed for the fundamental mode. However, to avoid significant fundamental mode loss, the waveguides are made to be slightly wider, thus increasing the possibility for other modes to be excited. These other modes can also have higher loss, and so even slight coupling to such modes may result in higher loss than expected. Predicting the exact coupling constants in (2.12) and (2.13) is unfortunately not analytically trivial, but they provide the idea that the total coupling loss should lie within a certain set of bounds. We can perhaps assume that the total power loss should not be worse than the waveguide mode with the highest loss. Thus, by calculating the maximum potential loss that can be contributed by each mode, we can ensure that we are looking at a more predictable range of loss in a waveguide.

2.5.2 Evidence of multimode-induced loss in a slightly multimode silicon waveguide

To get an idea of the the multimode perspective of loss, we look at a high-index contrast SOI channel waveguide with an air cladding. The silicon layer has a thickness of 260 nm, which at a wavelength λ = 1550 nm supports two slab modes – the fundamental TE mode and the fundamental TM mode. The waveguide width of such a waveguide is chosen to be at 420 nm. At this width, the waveguide supports a channel mode each for the TE and TM slab modes, denoted as TE_0 and TM_0. However, at this width, we are near the cutoff for single mode waveguiding. In fact, this design also supports the TE_1 mode, totalling the number of possible guided modes to three.

Suppose we look at the maximum possible loss for each guided mode for a waveguide length of 10 μm as illustrated in **Figure 2.5**. We had chosen this length specifically to avoid the noise floor of our photodetectors in experiments when we measure very lossy waveguides. If we assume no intermodal coupling could occur, we would expect a mode's loss to not go below its respective bound. Any noticeable fluctuation may be due to varying of the sidewalls' correlation length or due to other sources of loss from devices attached to these waveguides. However, one should be wary of such assumptions.

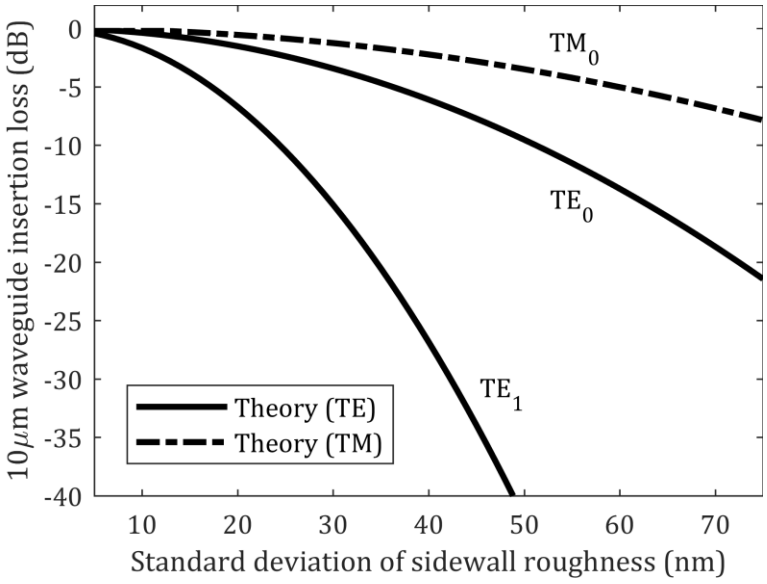

Figure 2.5 Theoretical maximum loss curves for various supported modes in a 260 nm x 420 nm silicon waveguide after 10 μm of propagation. The TE_1 mode depicts more loss since its mode lies more in the cladding compared to the fundamental TE_0 mode. The TM_0 mode depicts the least loss due to the evanescent tails of its mode lying largely in the top and bottom of the waveguide, which we assume has no surface roughness.

We simulate rough waveguides using a three-dimensional finite-difference-time-domain (3D-FDTD) method with a mesh size of 5 nm to ensure as much accuracy as our simulation computer can handle. We generated 5 rough waveguides for each

combination of the following set of parameters: σ_s = {10, 20, 30, 40, 50, 60, 70} nm, L_c = {25, 50, 75, 100, 125, 150} nm. This results in a total sample size of 210. We set the source wavelength to 1550 nm, excite the fundamental TE_0 mode, launch the mode source through each of these 10 μm long rough waveguides, and record the TE_0 transmission via the overlap integral. These results are depicted in **Figure 2.6** along with the theoretical loss bound for the TE_0 mode. While the results do seem to follow the theoretical curve fairly well, there are some things to consider. Immediately, it is evident that there are some rough waveguides showing loss below the theoretical bound. These discrepancies can be attributed to intermodal coupling.

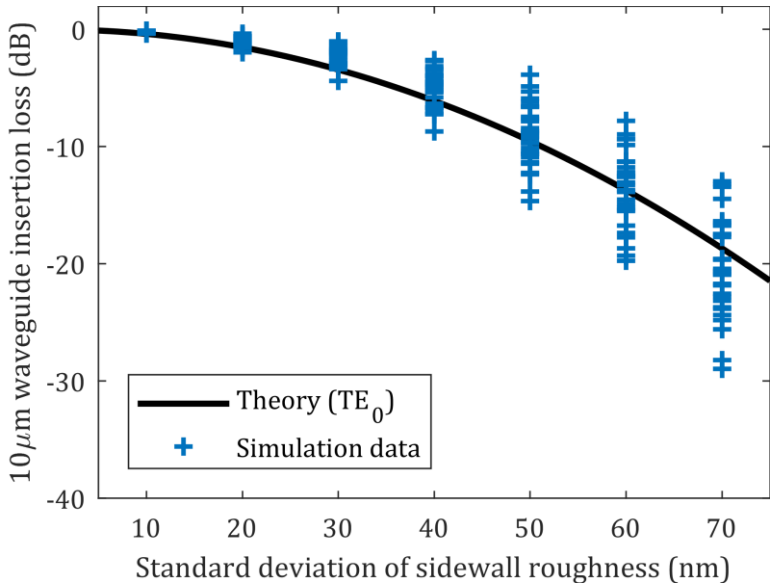

Figure 2.6 Simulation results for TE_0 input light with theoretical maximum loss for reference. As the actual loss may deviate as much as 10 dB from the maximum. This can be attributed to coupling to a higher-order guided mode, such as TE_1, that has significantly larger loss.

To further verify this claim, we show the output waveguide mode of a rough waveguide and compare it to the original source mode – see **Figure 2.7(a)**. Qualitatively, one can see that the mode has been polarization rotated. In other words,

the TM_0 mode was coupled. We look at another perspective in **Figure 2.7(b)**, depicting the propagating mode after the light passes the rough region of the waveguide and passes through a smooth straight waveguide. The light appears to be bending back and forth between each side of the smooth waveguide. This is in fact evidence of mode beating between the TE_0 and TE_1 modes. We also calculate the overlap integral of the output mode to each of the other two guided modes to determine the fraction of each mode being excited after propagating through the rough waveguide region. These results are respresented **Figure 2.7(c)**, thus quantitatively establishing evidence that there is loss induced by several modes coupled to each other.

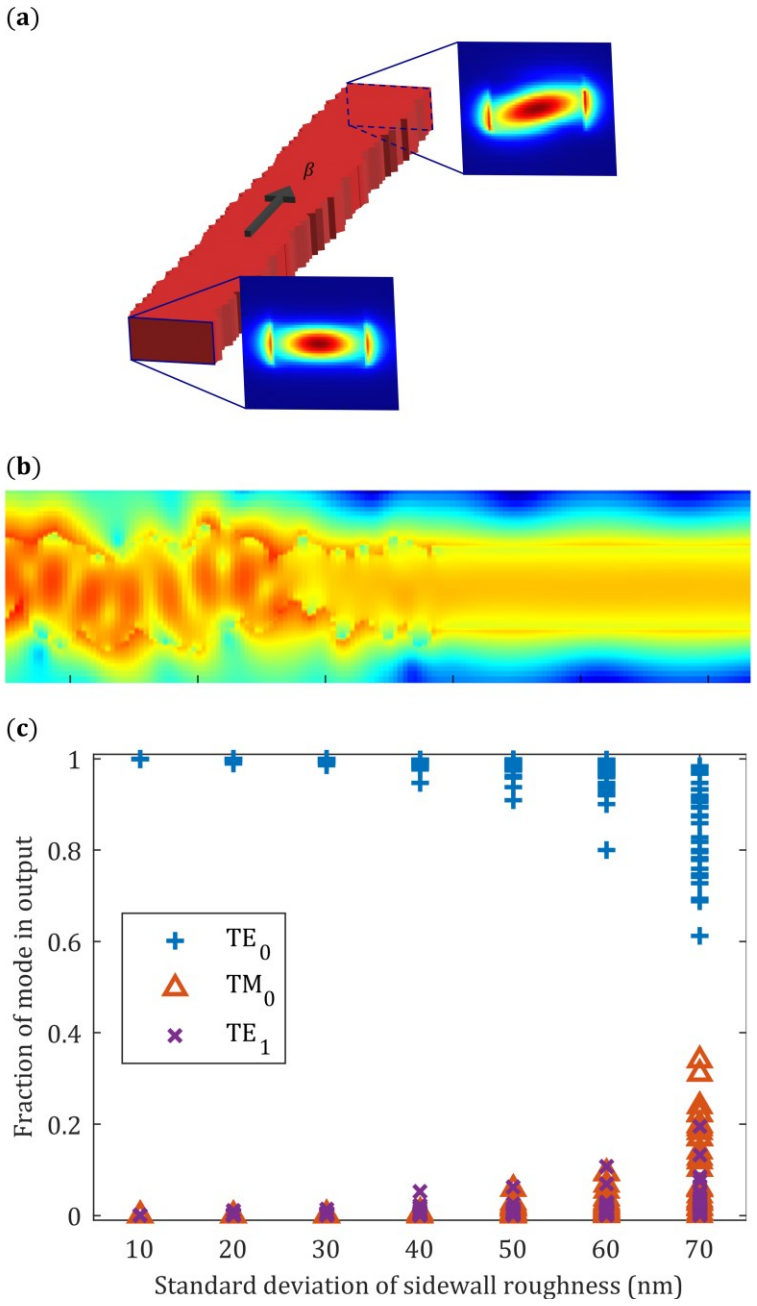

Figure 2.7 Simulations depicting (a) polarization rotation (plots are in linear scale) and (b) mode beating (plot is on a logarithmic scale) at the output of a roughened waveguide. The electric field intensities shown are from different rough waveguides. Polarization rotation is evidence of the TM_0 mode being excited, while mode beating is evidence of the TE_1 mode being excited. (c) Calculated data of the fraction of each mode at the output of various simulated roughened waveguides. Larger amplitude variations of the sidewalls increase the probability of having thin and wide waveguide sections, thus increasing the chances of the TM_0 and TE_1 modes being excited respectively.

We fabricated a set of rough waveguides with the same parameters defined in our simulations and used an SOI wafer with a 2 μm buried oxide layer. We patterned the waveguides with electron beam lithography with a resolution of approximately 4 nm when using hydrogen-silsesquioxane (HSQ) as a negative resist. To reduce sidewall roughness induced by ion bombardment, we utilized an inductively coupled plasma etcher with a HBr-Cl^2 recipe. No cladding was deposited over the waveguides, but a thin layer of HSQ – approximately 100 nm – remained. To obtain a high-resolution image of the waveguides' sidewalls, we used a scanning electron microscope (SEM). Some waveguides fabricated are shown in **Figure 2.8**. From these SEM images, we ran an edge detection algorithm and obtained the width along the total length of each rough waveguide. From here, we then calculated the RMS roughness and correlation length, which can be calculated from the surface distribution function and autocorrelation of each sidewall respectively. For RMS roughness, we have found that our fabricated devices are within 5 nm of their intended design. However, for correlation length, we have found that it is more difficult to control. Our analysis shows that our correlation length can vary between 60 nm to 240 nm.

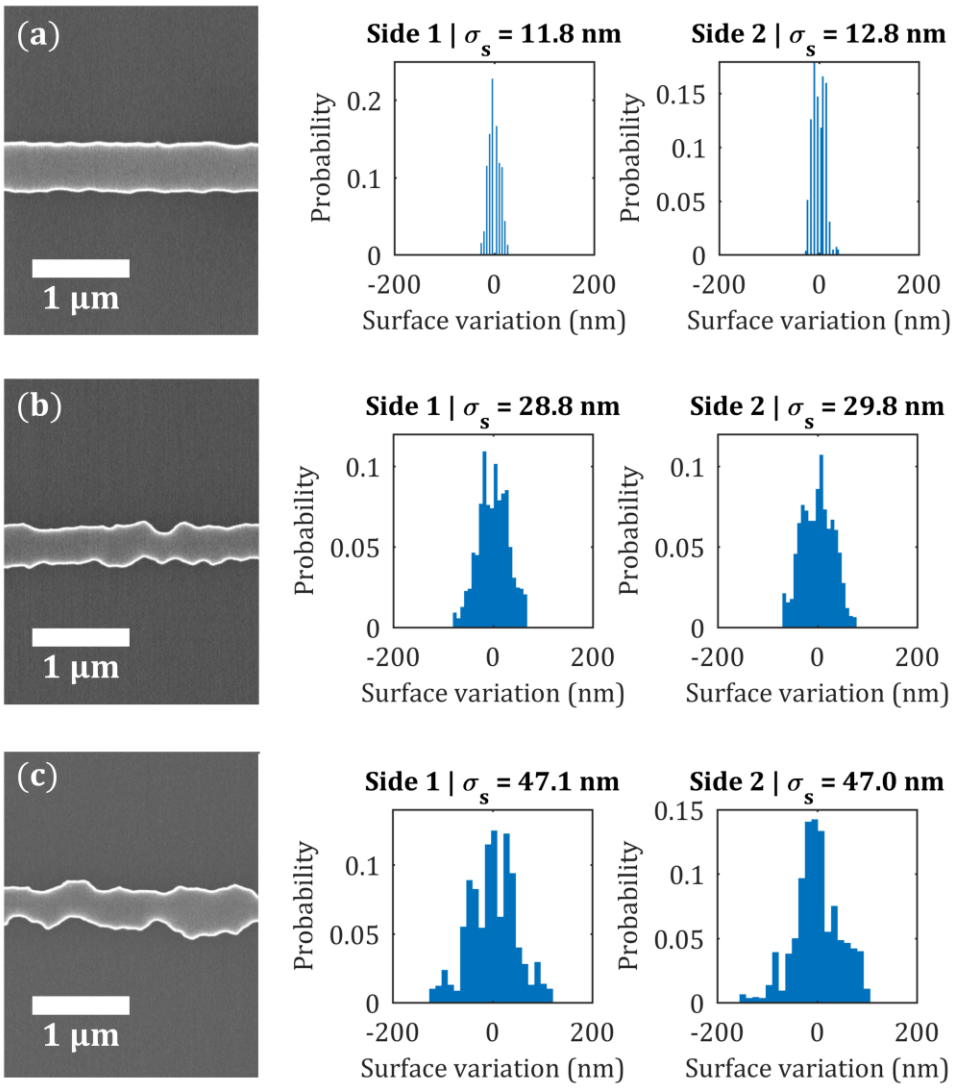

Figure 2.8 Scanning electron microscrope (SEM) images and corresponding surface distribution functions for fabricated waveguides designed with (a) σ_s = 10 nm, (b) σ_s = 30 nm, and (c) σ_s = 50 nm. Excitation to other guided modes can be a result of a waveguide abruptly thinning, widening, or shifting in the lateral direction. These physical features can be seen more often in waveguides with larger amplitude variations.

We measured each waveguide by coupling a tunable laser through single-mode fibers to on-chip grating couplers, which were designed to maximize the coupling efficiency for the TE_0 waveguide mode. We also placed smooth sidewall reference waveguides around the chip to normalize its nearest rough waveguide

measurements. This normalization considers the insertion losses from the measurement setup, grating couplers, waveguide bends, and tapers. The insertion loss for all 210 measurements at source wavelength of 1550 nm are depicted in **Figure 2.9** as a scatter plot with the theoretical TE_0 insertion loss curve for reference.

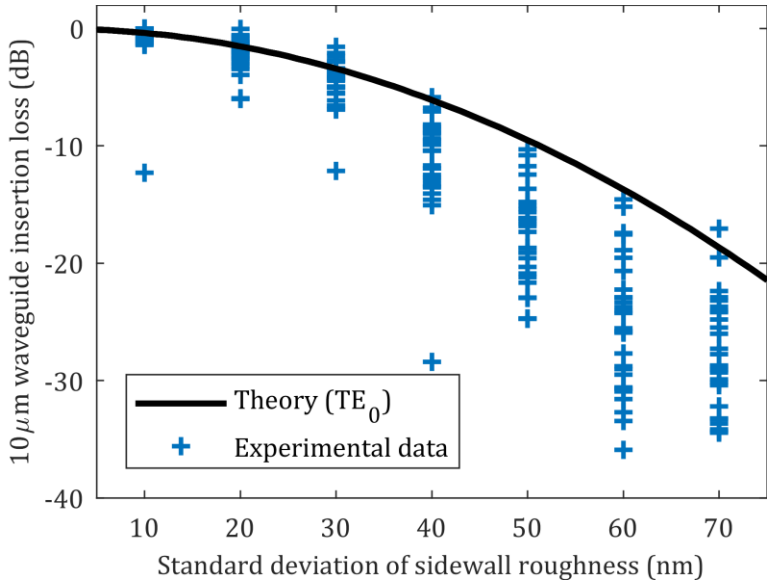

Figure 2.9 Experimental results for TE_0 input light with theoretical maximum loss for reference. It is also clear that as the actual loss does deviate from the theoretical maximum. Some outliers may even suggest that significant coupling to TE_1 coupling may be the result of extraordinary loss.

As expected, the insertion loss significantly deviates more from the theory as the RMS roughness increases in value. For this set of roughness parameters, the insertion loss should not go below the curve as it represents the maximum loss regardless of the correlation length. The experimental results also show more loss compared to the simulations. This increased loss may be attributed to several factors such as leakage into the substrate, change in the vertical mode confinement due to the thin film of remaining HSQ resist, or defects either in the HSQ film or at the silicon-HSQ interface. The larger spread of loss is a result of unintended variations in the actual RMS

roughness and correlation length values. Some outliers also suggest that there are specific ranges of correlation lengths that result in strong coupling to the TE_1 mode, thus causing a large increase in insertion loss.

2.6 **Discussion**

In the simulation and experimental results depicted in **Figure 2.6** and **Figure 2.9**, we are under the assumption that we can only measure the fundamental TE_0 mode. The total insertion loss in the waveguide should be higher if we also measure the light coupled to the TM_0 and TE_1 modes. As suggested by **Figure 2.7**, as much as 40 percent of the output light can be coupled in the higher order guided modes. Regardless, it should be expected that the insertion loss is bounded by the highest-order guided mode – in this case, TE_1.

Traditionally, waveguide loss is minimized by widening the waveguide and predicting the loss for the intended guided modes. However, the design should consider even the slightest excitation of the highest-order guided mode as even slight coupling can result in detrimental loss. For example, suppose we were to design for a three TE-mode waveguide. To minimize the loss, we should expect to widen the waveguide to a point where the sidewall variations – which should be known from the foundry – would not result in a waveguide segment that supports a higher order lossy mode. Suppose a foundry can guarantee that their sidewall roughness is typically on the order of σ_s = 3 nm and L_c = 100 nm. We simulated waveguides of various widths with input source modes – TE_0, TE_1, TE_2, and TE_3 – and calculated their insertion losses as shown in **Figure 2.10(a)**.

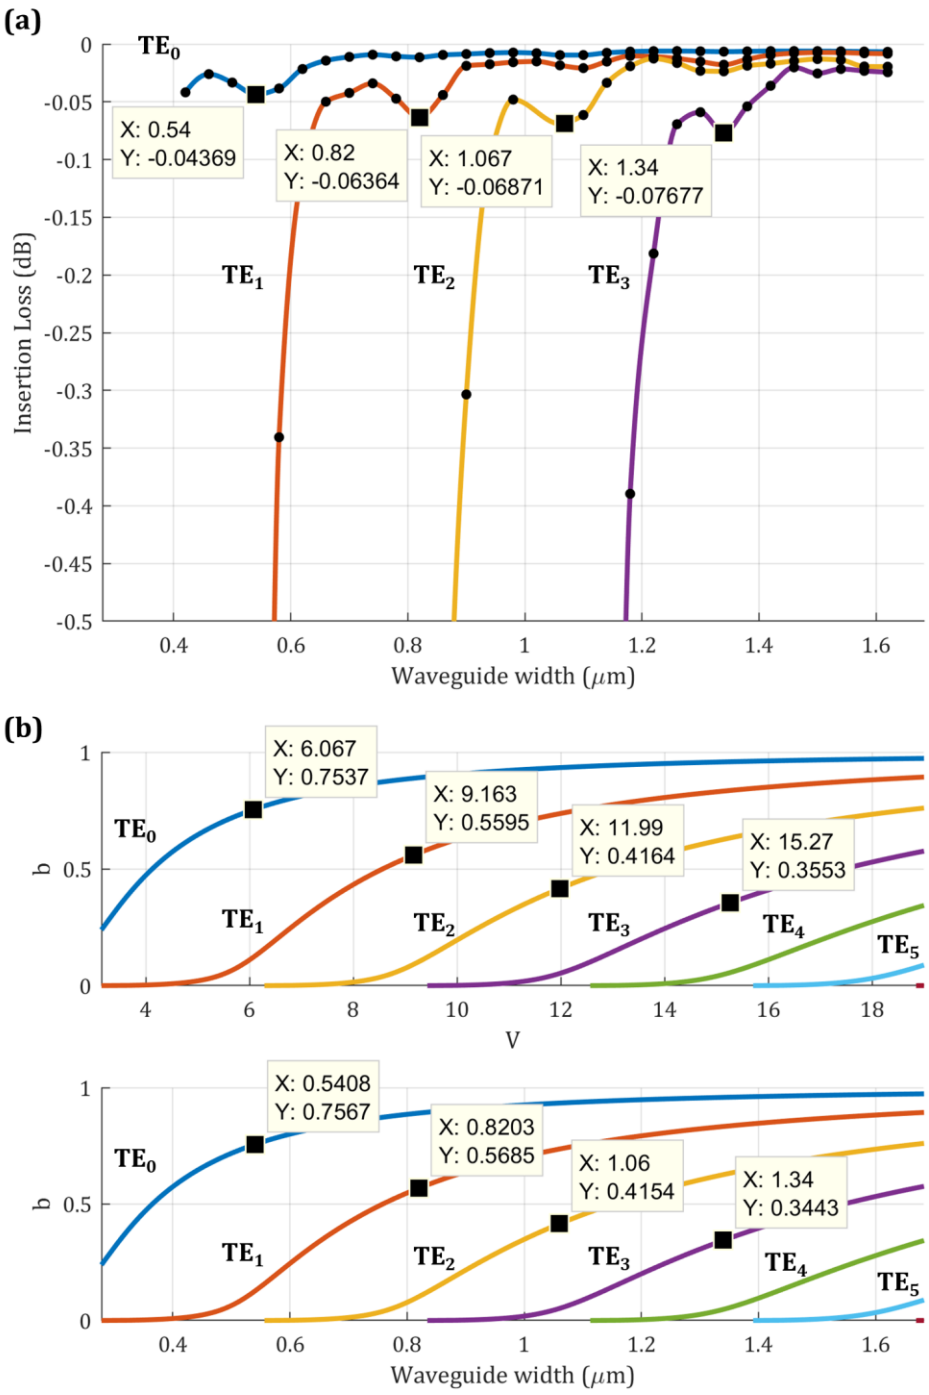

Figure 2.10 (a) Simulated insertion loss results of the first 4 TE modes as a waveguide's width increases. The sidewall roughness – σ_s = 3 nm and L_c = 100 nm – is the same in all simulations and we assume there is no top or bottom side roughness. The waveguide length is 10 μm. (b) Mode solutions calculated from equations (2.3) to (2.5) noting the points where the dip in insertion loss occurs for each mode in a rough waveguide.

Interestingly, the largest dips in the insertion loss occur at the point of waveguide widening where a new mode is supported, albeit incredibly lossy. For the fundamental TE_0 mode this dip occurs when the TE_2 mode can begin to be guided. Likewise, a dip in insertion loss for the TE_1 mode occurs when the TE_3 mode can begin to be guided. We suspect that there is some relation with the TE_0 and TE_2 modes both being symmetric modes and the TE_1 and TE_3 modes both being anti-symmetric modes. As the modes are perturbed while propagating through a rough waveguide, it is more likely for the symmetric modes to overlap with one another as well as the antisymmetric modes with one another. Thus, intermodal coupling may occur and potentially higher loss. Using simplified models of waveguide loss only provide a first-order insight on how to best engineer low-loss waveguides. With the additional understanding of mode-coupling induced loss, foundries would be able to provide better design rules that would benefit multimode photonic integrated circuits. In our case, if we want a low-loss waveguide for the first three TE modes, we may want to design the waveguides to be around 1.22 μm wide based on what we observe in **Figure 2.10(a).** Additionally, since higher order guided modes tend to have higher loss than the lower order modes, intentionally designed sidewall roughness can potentially be a means towards high order mode clean-up filters. In other words, we can perhaps control the sidewall roughness to maximize the higher-order loss with the least amount of intermodal coupling and low-order loss.

It is possible to achieve low-loss waveguides for great distances using doped-silica waveguides or even polymer waveguides. However, while lower-index contrast waveguides such as these tend to have lower loss from sidewall roughness given that

there is an appropriate way to manufacture them, we lose the capability to pack waveguides densely. Waveguide modes are not as tightly confined in such materials, and thus to avoid loss from tight bends and achieve loss crosstalk between adjacent waveguides, the footprint must be larger and ultimately less scalable.

Sidewall roughness is the primary source of propagation loss in silicon-based waveguides and widening the waveguides can help reduce loss significantly. However, to reach losses on the in sub-dB/m regime, the top and bottom side roughness should be considered. As the main provider of SOI wafers at the time of this writing, SOITEC has proven they have reached atomic-level smoothness in their device-level silicon layer. However, as the SOI wafer goes through its fabrication journey, various processes can inadvertently roughen the top of the waveguides. Sometimes over-etching can occur on the bottom corners of the waveguides, resulting in rougher corners an additional difficulty to fill the over-etched trenches during upper cladding deposition – both of which can be sources of extra loss.

Roughness reduction techniques exist to reduce the loss, which typically include various rounds of annealing or wet etching [41]–[43]. However, there are caveats. The most obvious is that, through several rounds of roughness reduction, waveguides will be become thinner or shorter. An inherent bias exists; however, it would be simple to include such bias in the design process. Another but more serious issue is the shape of the waveguide cross-sections changing from a more typical rectangular shape that can be trivially modelled to something more amorphous or irregularly shaped – in this case, more difficult to model. To this end, waveguide loss

should not be the only figure of merit in PIC design. Other scalability perspectives, such as the simplification of device modelling should be considered.

2.7 Conclusion

In summary, we defined a model which describes the combined power coupling among guided modes, radiation modes, and counterpropagating modes (backscattering) under the assumption that we have no free-carrier absorption loss, no surface-state absorption, and negligent substrate leakage. We verified in both simulation and experiment that the design of border-line single mode waveguides can have sidewall roughness that induces intermodal coupling and larger losses than expected [44]. We emphasized and explored this concept by specifically looking at waveguides with large sidewall variations. While real waveguides may be optimized to have minimal sidewall roughness, intermodal coupling and higher-order loss can still add up as multimode waveguides get longer. Therefore, the idea of intermodal-coupling-induced loss should be considered in the design process.

Chapter 3: Waveguide Superlattice Bends

3.1 Background and motivation

Chip-scale optical interconnects have been envisioned for future chips comprising more than 100 cores. Parallel signal transmission over many waveguide channels can meet the high bandwidth demands of future technology, but the area occupied by tens of thousands of waveguides may be a concern [45]. In other words, we will need to increase the capacity of optical interconnect buses to utilize the space and cost of the system effectively just as we need to increase the capacity in fiber optics. In addition, OPA applications utilizing integrated optics can benefit from compacting their waveguide array apertures, demanding a decrease in aperture pitches to increase beam steering angles. However, the high signal crosstalk is a bottleneck. There have been many approaches and structures to avoid this issue [17], [20], [46]–[50]. In particular, the idea of a waveguide superlattice (WGSL) has shown to be promising [48].

However, as rigorous as the design process is for the WGSL is, it only assumes that the waveguides are in a straight line. When the superlattice bends, a trivial solution would be to temporarily increase the inter-waveguide pitch and squeeze the waveguides together after the bend, but with the layout of a chip being increasingly precious, there should be every effort to decrease the footprint of each component in the interconnect system. In this chapter, we look at the design rules of a WGSL and provide avenues to consider when designing for the bending regime.

3.2 **Achieving low crosstalk among dense waveguide arrays**

Crosstalk in waveguides, specifically between a pair of waveguides, can be affected by three main parameters: the propagation constant difference (or phase mismatch) $\Delta\beta = \beta_1 - \beta_2$, the coupling strength (or coupling coefficient) κ, and the propagation length L. For light propagating in the z-direction, the coupling coefficient can be calculated as

$$\kappa = \kappa_{12} = \frac{\omega\epsilon_0 \iint_{-\infty}^{\infty} (\epsilon_r - \epsilon_{r,2}) \overrightarrow{E_1^*} \cdot \overrightarrow{E_2} dxdy}{\iint_{-\infty}^{\infty} \hat{z} \cdot (\overrightarrow{E_1^*} \times \overrightarrow{H_1} + \overrightarrow{E_1} \times \overrightarrow{H_1^*}) \overrightarrow{E_1^*} \cdot \overrightarrow{E_2} dxdy}, \tag{3.1}$$

where ϵ_r is the dielectric profile containing both waveguides and $\epsilon_{r,2}$ is the dielectric profile containing only the 2nd waveguide [51], [52]. The crosstalk between a pair of waveguides in [53] is described as

$$\frac{P_{1\to2}}{P_1} = \frac{1}{(\Delta\beta/2\kappa)^2 + 1} sin\left(\sqrt{(\Delta\beta/2\kappa)^2 + \kappa^2} L\right). \tag{3.2}$$

Across any arbitrary propagation length, the maximum crosstalk is given by

$$max\left[\frac{P_{1\to2}}{P_1}\right] = \frac{1}{(\Delta\beta/2\kappa)^2 + 1}. \tag{3.3}$$

To achieve a low crosstalk condition, $\Delta\beta >> \kappa$, trivially speaking, we can make the pair of waveguides asymmetric in width. Doing this by means of making one waveguide smaller or wider than the other results in both a larger phase mismatch and potentially a smaller coupling coefficient. In [48], we denote this as a supercell (SC)-2 configuration. However, as illustrated in **Figure 3.1**, we can have low intra-cell

crosstalk (between an asymmetric pair), but we may have substantial inter-cell crosstalk (between the same type of waveguide from one supercell to the next). Even if these waveguides are farther than the intra-cell pitch, they are completely phase-matched, so their coupling is considerably strong.

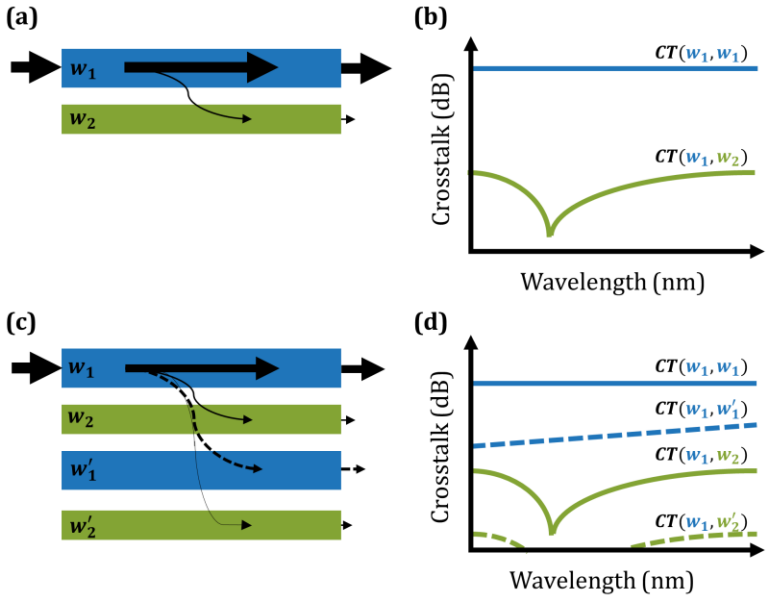

Figure 3.1 (a) A trivial pair of asymmetric waveguides with widths w_1 and $w_2 \neq w_1$ and (b) a depiction of its expected crosstalk. (c) If we consider laying this pair in a periodic array, we may see (d) considerable intra-cell crosstalk – particularly between w_1 and w_1'.

One could wrongfully assume that we can achieve a densely packed array – with a constant pitch between each waveguide – by stacking the SC-2's side-by-side. Unfortunately, with SC-2, inter-cell crosstalk is substantial at sub-wavelength pitch (center-to-center distance between waveguides). Therefore, the design for a large dense array of waveguides using an SC-2 configuration is limiting. However, the idea of a WGSL provides an effective solution.

Based on heuristic guidelines defined in [48], crosstalk is generally mitigated significantly by maximizing the phase mismatch across all waveguide channels. We

defined the superlattice modes as Λ_n ($n = 1 \ldots n_{super}$), where n_{super} denotes the number of waveguides in an SC-n_{super}. We want to maximize the phase mismatch between superlattice modes ($\Lambda_m - \Lambda_n$) to achieve low crosstalk. Analytically, this means

$$|K_{mm} - K_{nn}| >> \sum_{k \neq m} |K_{mk}| + \sum_{k \neq n} |K_{nk}|, \tag{3.4}$$

where $[K] = [B]^{-1}[\Delta A] + [\beta]$ such that $[\Delta A]$ is a matrix describing the index perturbation due to the presence of surrounding waveguides, $[B]$ is a metric matrix, and $[\beta]$ is a diagonal matrix containing the propagation constant β_n of each waveguide mode. This theory does not assume the elements of $[\Delta A]$ are small and, therefore, the results are not based on small-perturbation approximations. We can also look at an overestimation of the upper bound of crosstalk between input waveguides and output waveguides, which is defined as

$$|u_m(L)| < \sum_l \left| \sum_n B_{mn} \left[Q_{nl} Q_{lq}^{-1} \right] \right|, \tag{3.5}$$

where $u_m(L)$ is the output amplitude and $[Q]$ is obtained from the eigenvalue decomposition of $[K]$. Note that due to the utilization of eigenvalue decomposition, using (3.5) for WGSL design tends to be computationally demanding. It turns out that an SC-5 WGSL – as depicted in **Figure 3.2** – is an effective dense waveguide array design to achieve sub-wavelength pitch with low crosstalk.

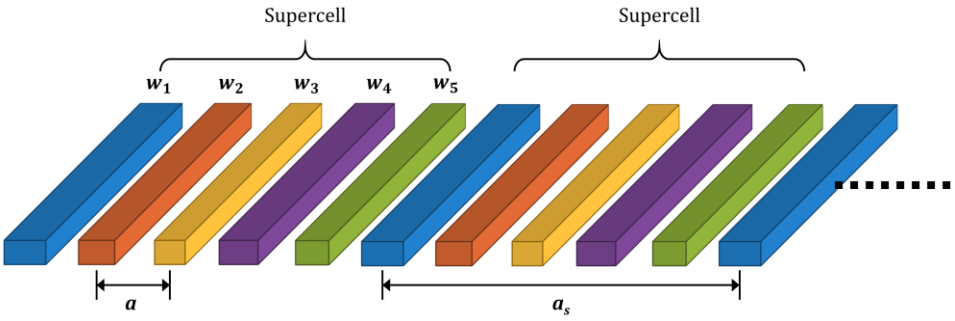

Figure 3.2 Schematic of a waveguide superlattice.

In an SC-5 WGSL, the waveguides in one supercell all have different widths and laid out in an interlacing-recombination fashion. Specifically, for a silicon waveguide with thickness of 260 nm, the widths in one supercell are 450, 390, 330, 420, and 360 nm, in that order. The waveguides are cladded with silicon dioxide. In [48], this design demonstrated inter-supercell and intra-supercell crosstalk below -20 dB across a light source wavelength span from 1500 nm to 1580 nm. The inter-waveguide pitch was designed at a = 0.78 μm.

3.3 The case for waveguide superlattice bends

There may be cases that the WGSL will have to bend, such as in an optical interconnect system with multiple CPU cores, and the optical communication bus may have to bend around a complex network. This is illustrated in **Figure 3.3**.

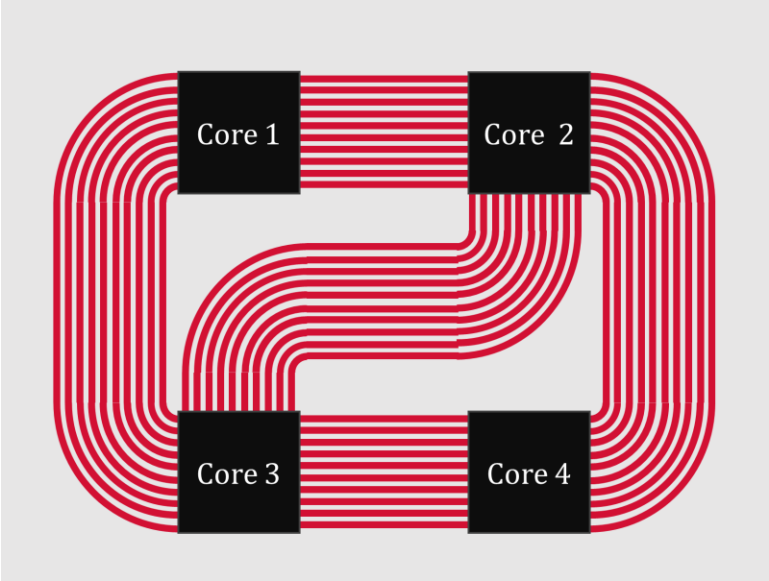

Figure 3.3 An illustration of how waveguide arrays may be routed amongst various components of a PIC interposer for CPU cores.

One may argue that a WGSL will just have to route and take the resultant route to a neighboring CPU core. The overall bus length will even be shorter and, thus, the amount of optical power loss would be less. However, a fabrication-rooted concept should be considered. Typically, lithography masks have a finite resolution when made, and are typically digitized in a rectangular coordinate system – see **Figure 3.4**. What this means is that a diagonal structure will be susceptible to what is known as a staircasing effect. In other words, the mask edges for a diagonal structure will be systematically rough and, therefore, a diagonal waveguide bus may be more lossy than one that bends back and forth.

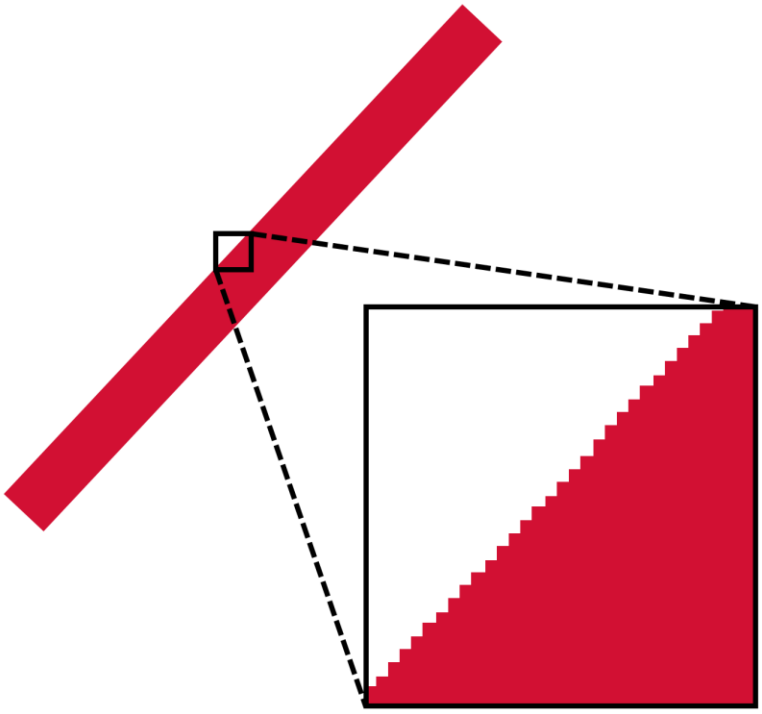

Figure 3.4 Lithography mask digitization in diagonal waveguides emphasizing the staircasing effect.

There are other systems that may also be beneficial with a bending WGSL. Take for example a thermo-optic switch based on dense waveguide arrays [47] – an illustration shown in **Figure 3.5(a)**. Despite being a compact structure in the straight waveguide array regions, there is a substantial amount of layout space dedicated to routing the waveguide arms curving back and forth. This can be potentially reduced as depicted in **Figure 3.5(b)** – perhaps as a dense spiral waveguide-like structure.

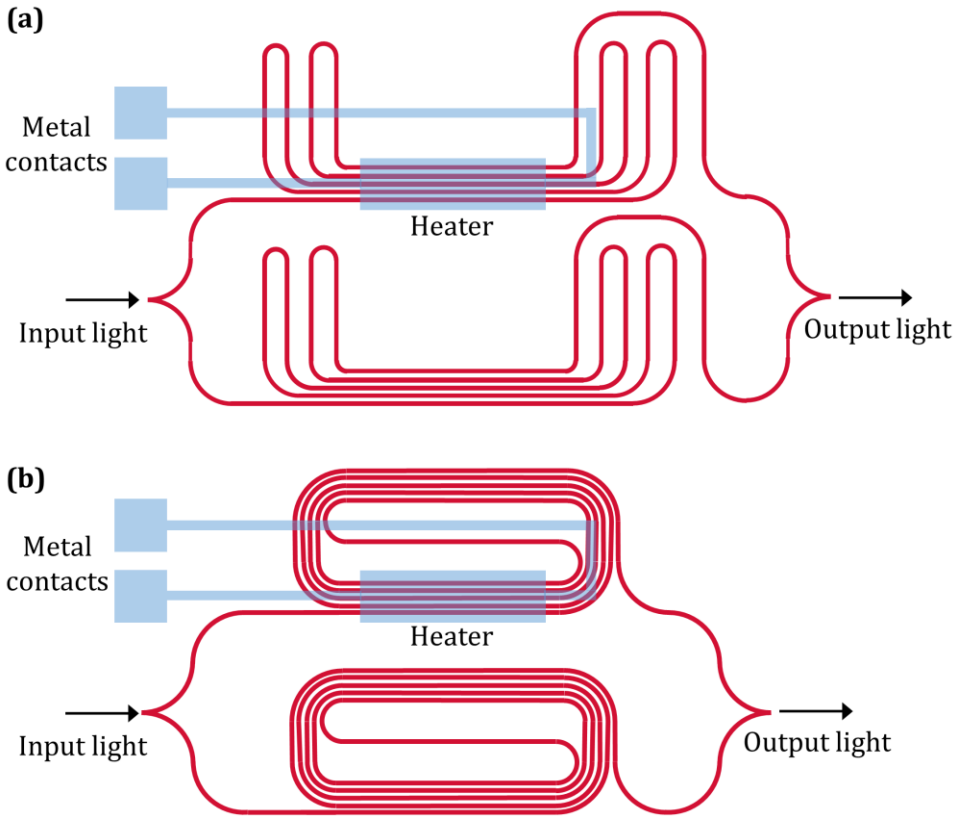

Figure 3.5 (a) An efficient thermo-optic switch using a dense waveguide array as demonstrated by [47] and (b) an alternative space-efficient switch with utilizing waveguide superlattice bends.

3.4 Coupling with a pair of concentric waveguide bends

As with straight waveguide case, we first simplify the problem for waveguide bends with a modification to (3.3), considering the bend in a cylindrical coordinate system:

$$max\left[\frac{P_{1\to 2}}{P_1}\right] = \frac{1}{\left(\Delta\beta_\phi/2\kappa\right)^2 + 1}, \tag{3.6}$$

where angular phase mismatch $\Delta\beta_\varphi$ is defined as

$$\Delta\beta_\phi = \beta_1\left(\frac{R_1}{R_2}\right) - \beta_2, \tag{3.7}$$

and R_1 and R_2 correspond to the bending radii [54], [55]. As one can clearly see, the ratio of the radii can inadvertently provide a phase matching condition even if the straight waveguide pair does not. We illustrate this further in **Figure 3.6**.

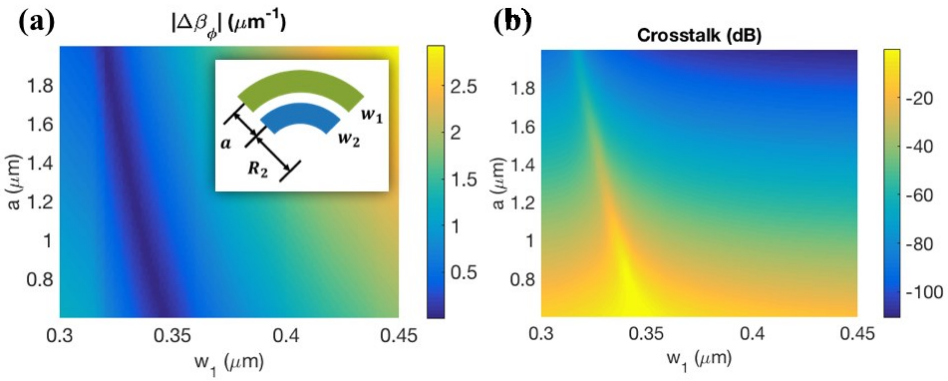

Figure 3.6 (inset) Schematic of a waveguide pair in a bend. (a) Phase mismatch at λ = 1550 nm with w_1 varying and R_1 = (R_2 + a). w_2 and R_2 are set at a constant 360 nm and 19.68 μm respectively. (b) Simulated crosstalk results show that phase mismatch is a significant factor even at large pitch dimensions.

One can see that $\Delta\beta_\varphi$ trends with crosstalk even when the waveguides are supposedly a fair distance apart. Conversely, if waveguides are phase matched ($\Delta\beta_\varphi = 0$), they will certainly not be phase matched when the waveguides bend ($\Delta\beta_\varphi \neq 0$). To this end, it is crucial to understand these ideas to push the limits of a WGSL bend.

3.5 Dielectric constant profile for a waveguide superlattice bend

For an array of curved silicon waveguides, the mode equation can be developed by considering the fields in each homogeneous region (along one of the waveguide cores, or in one cladding region). In each region, the dielectric function is a constant, therefore one can readily show that the following equation holds in each homogeneous region.

$$\nabla^2 H_z + \epsilon k_0^2 H_z = 0, \tag{3.8}$$

where z is normal to the plane of the waveguide array and k_0 is the free space wavevector. The effective dielectric constant approach can be applied [56]. Letting $H_z(\rho, z) = \rho^{1/2} u(\rho)$, one can readily obtain

$$\frac{d^2u}{dz^2} + \frac{d^2u}{d\rho^2} + [k_0^2 \epsilon_b(\rho) - \beta^2]u = 0, \tag{3.9}$$

where β is the propagation constant of the mode, and an effective dielectric function of the bend is introduced as

$$\epsilon_b(\rho) = \epsilon(\rho) + \frac{1}{4k_0^2\rho^2} + \frac{\beta^2}{k_0^2}\left(1 - \frac{R^2}{\rho^2}\right). \tag{3.10}$$

With the substitution of $\rho \to \tilde{x}$, $z \to \tilde{y}$, (3.9) can be treated as the mode equation of a straight waveguide

$$\frac{d^2u}{d\tilde{x}^2} + \frac{d^2u}{d\tilde{y}^2} + [k_0^2 \epsilon_b(\tilde{x}) - \beta^2]u = 0. \tag{3.11}$$

Here, we are interested in the property of this dielectric function and its implication rather than directly solving the mode equation. For this purpose, we can assume the modal propagation constant is obtained through other means – such as FDTD simulations – or otherwise assume it is very close that of the straight waveguide, $\beta_0 = n_{eff,0} k_0$, which is typically accurate (within less than 1 percent) for silicon waveguides with $R > 3$ μm. Then one can plot this dielectric function as shown in **Figure 3.7**. For a single waveguide, it can be shown that the cladding region will be

raised to an effective dielectric constant equal to β/k_0, which breaks the guiding condition, at a radius of

$$\rho_{bend,rad} \approx \frac{R\beta}{(n_{ox}k_0)}. \tag{3.12}$$

The distance between the center axis of the waveguide and the ρ_{rad} is usually fairly large for a silicon waveguide, usually $\rho_{bend,rad} - R > 0.5R$ for typical values of modal effective indices. For an array of curved silicon waveguides, the breaking of guiding conditions occurs at $\rho_{array,rad} - R = a$. This location is much closer to the original waveguide and the modal field is substantially stronger.

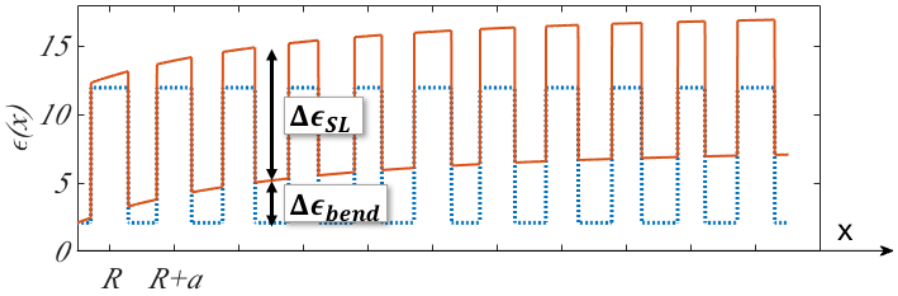

Figure 3.7 Dielectric constant profile of a waveguide superlattice for straight and bend cases. For the bend, the bending effect is accounted by the effective index method (R = 10 μm). Straight superlattice: blue dotted line. Bend: orange line.

3.6 **Crosstalk performance of waveguide superlattice bends**

Typically, in a straight WGSL, the worst-case crosstalk occurs in the first nearest neighbors. But at certain bending conditions when the radii cause for a low phase mismatch, worst case crosstalk could potentially occur in farther neighbors. A trivial way to solve this issue is to increase the inter-waveguide pitch, say to 1 μm. However,

we may not desire to change the pitch to maximize space savings. Another way to slightly increase the widths of some waveguides, specifically all waveguides in a supercell, except the first one set to 450 nm. This in turn could increase $\Delta\beta_\varphi$ for some waveguide pairs with others decreasing as a tradeoff. In addition, we could be increasing κ due to the edges of each waveguide being closer to their corresponding neighbors, so the waveguide widths should be modified such that this is not a significant factor.

In our prior WGSL, we use an SC-5 design in an interlacing-recombination configuration such that the waveguide width difference between each interlacing channel ($\Delta w = w_1 - w_4 = w_4 - w_2 = w_2 - w_5 = w_5 - w_3$) is set to 30 nm and $w_1 = 450$ nm. A schematic of this SC-5 WGSL bend is shown in **Figure 3.2**. From here, we explore what width differences can reduce the overall crosstalk in the WGSL when it bends.

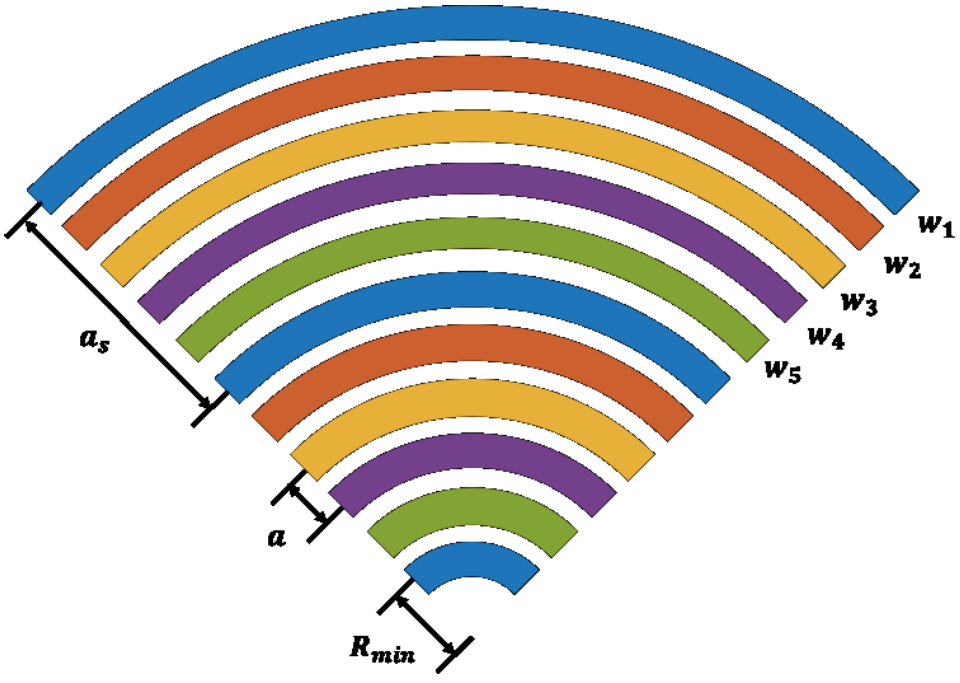

Figure 3.8 Schematic of a WGSL bend with two periods of SC-5 and an extra waveguide that would mark the start of a 3rd period of an SC-5, totalling 11 channels in the array.

This is a simple starting point; however, by simplifying the problem to changing just the waveguide width difference, we should expect there to be some tradeoff. Considering (3.6), we expect crosstalk between waveguide *m* and waveguide *n* in decibels to be approximately

$$CT_{dB} \approx 10\,log_{10}\left(\frac{|\kappa_{mn}|^2}{|\Delta\beta|^2}\right). \tag{3.13}$$

From **Figure 3.6**, we can deduce that at larger waveguide pitches, the phase mismatch is the dominating factor in crosstalk. Thus, we can estimate the new crosstalk when the phase mismatch between two waveguides change as

$$CT_{dB,new} \approx 10\,log_{10}\left(\frac{|\kappa_{mn}|^2}{\left|\left[1+\frac{\Delta\beta-\Delta\beta_0}{\Delta\beta_0}\right]\Delta\beta_0\right|^2}\right), \tag{3.14}$$

where $\Delta\beta$ is the new phase mismatch and $\Delta\beta_0$ is the previous phase mismatch. We can also separate (3.14) into two terms as

$$CT_{dB,new} \approx 10\,log_{10}\left(\frac{|\kappa_{mn}|^2}{|\Delta\beta_0|^2}\right) + 10\,log_{10}\left(\frac{1}{\left|1+\frac{\Delta\beta-\Delta\beta_0}{\Delta\beta_0}\right|^2}\right), \tag{3.15}$$

where the first term is the previous crosstalk CT_{dB}, and the second term is the change in crosstalk ΔCT_{dB}. Due to the assumption that the coupling constant κ_{mn} does not change, there will be some error in these calculations, but they provide an intuitive approach to estimate the change in crosstalk when one changes the phase mismatch. In **Figure 3.9**, we show some calculations predicting the expected change in crosstalk

in various waveguide designs when changing the width difference Δ*w* from 30 nm to 25 nm for various waveguide superlattice bend configurations at inter-waveguide pitch *a* = 0.78 μm. In these calculations, we are using the angular phase mismatch defined in (3.7).

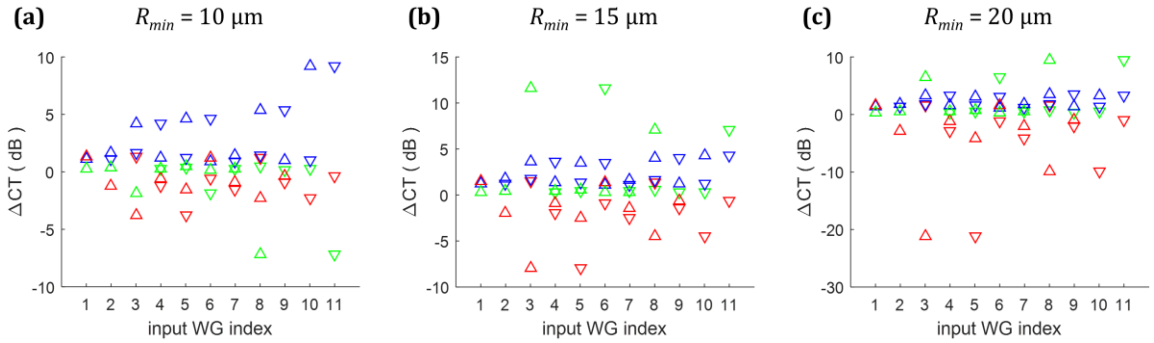

Figure 3.9 Calculated change in crosstalk for waveguide superlattice bends with pitch *a* = 0.78 μm and configured with a minimum bending radius of (a) 10 μm, (b) 15 μm, and (c) 20 μm. Changing the supercell width difference from 30 nm to 25 nm may reduce the crosstalk between some waveguide pairs but may also increase the crosstalk between others. (blue) 1st nearest neighbors, (red) 2nd nearest neighbors, (green) 3rd nearest neighbors

Considering the case with *a* = 0.78 μm and R_{min} = 15 μm, we can see that decreasing Δ*w* would potentially decrease the crosstalk of most 2nd nearest neighbors with most of the impact evident between channels 3 and 5 – as well as some significant impact to channels 8 and 10. On the other hand, the crosstalk increases for 1st nearest neighbors. However, if the crosstalk was originally low to these channels, then there should still be an overall benefit to the design modification – as is the case for 3rd nearest neighbors and those farther, too. Interestingly, channels 3 and 8 are an SC-5 period away as well as channels 5 and 10, which suggests that the crosstalk impact is also somewhat periodic.

We simulated both cases using a 3D-FDTD method over a full 180-degree bend and recorded the maximum crosstalk that could occur for each wavelength from 1500 nm to 1580 nm in the span of that bend. In other words, regardless of the amount of bending – between 0 and 180 degrees – the ideal crosstalk should not go above these upper bounds.

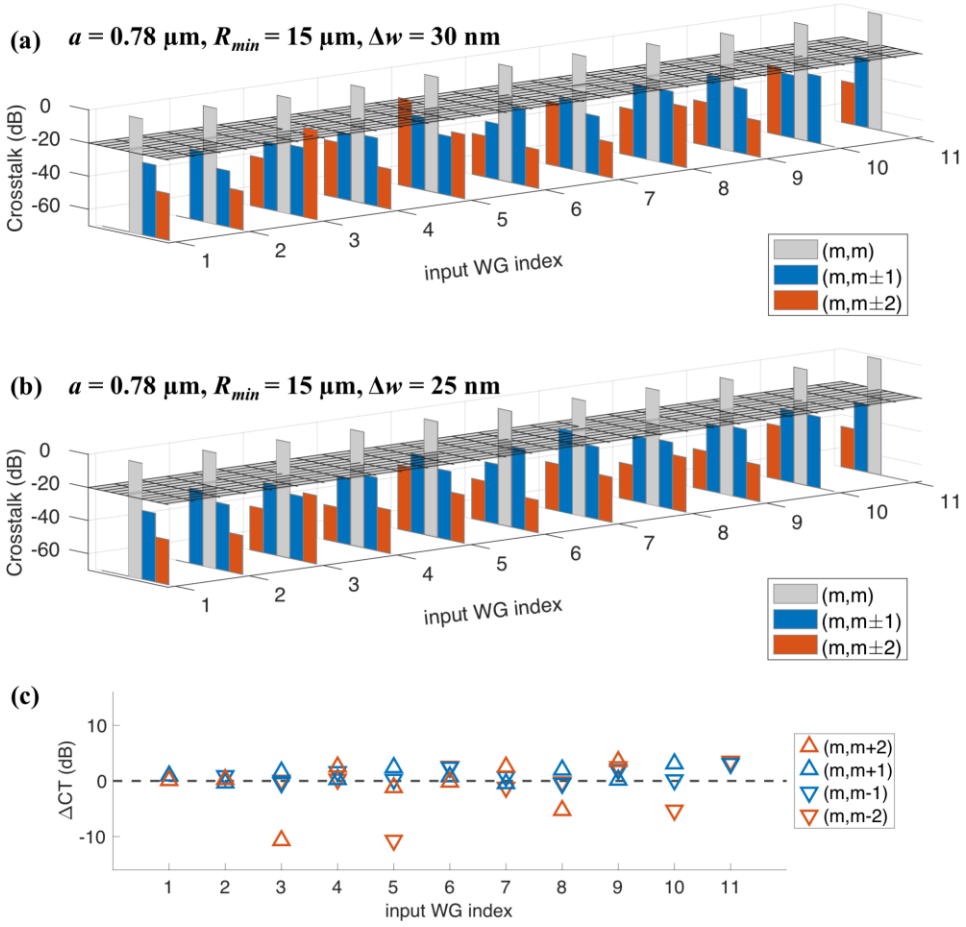

Figure 3.10 Simulated results for inter-waveguide pitch 0.78 µm at R_{min} = 15 µm for (a) Δw = 30 nm and (b) Δw = 25 nm, showing maximum relative crosstalk each channel and its 1st and 2nd nearest neighbors with gray plane marking -20 dB. (c) The expected change in crosstalk when Δw is reduced to 25 nm. As the waveguide width difference is decreased, there is potential benefit in reducing crosstalk from 2nd nearest neighbors by increasing the phase mismatch.

At these set of parameters, the maximum crosstalk is borderline -20 dB, and some of these results in **Figure 3.10(a)** show that 2nd nearest neighbor crosstalk is higher than that in the straight superlattice case – some even higher than that of the 1st nearest neighbors. In a realistic setting, some roughness-induced crosstalk could make this even worse, and so we take preventative measures via design by using the Δw = 25 nm case. A comparison is shown in **Figure 3.10(c)**. This is close to our calculated predictions. The 3rd nearest neighbors and farther are not shown for simplicity.

We fabricated these WGSL bend structures on an SOI wafer and characterized it with a repeatable waveguide coupling and measurement setup. We illustrated experimental data statistically for each input-output channel combination of a fabricated WGSL bend with an inter-waveguide pitch a = 0.78 μm and minimum bending radius R_{min} = 15 μm in **Figure 3.11(a)** and **(c)**. At each wavelength, we show the averages and standard deviations (statistics over index m) for the direct transmission $T(m,m)$, the first nearest neighbor crosstalk $T(m,m\pm1)$, and the second nearest neighbor crosstalk $T(m,m\pm2)$. In addition, **Figure 3.11(a)** and **(c)** show the worst crosstalk (thick black curve) for each particular wavelength $T(m,n)_{worst}$ and summarize the maximum overall relative crosstalk for key waveguide channels. The relative crosstalk is determined by the difference of their transmissions to the direct transmission of the input channel, $CT(m,n) = T(m,n) - T(m,m)$, as shown in **Figure 3.11(b)** and **(d)**.

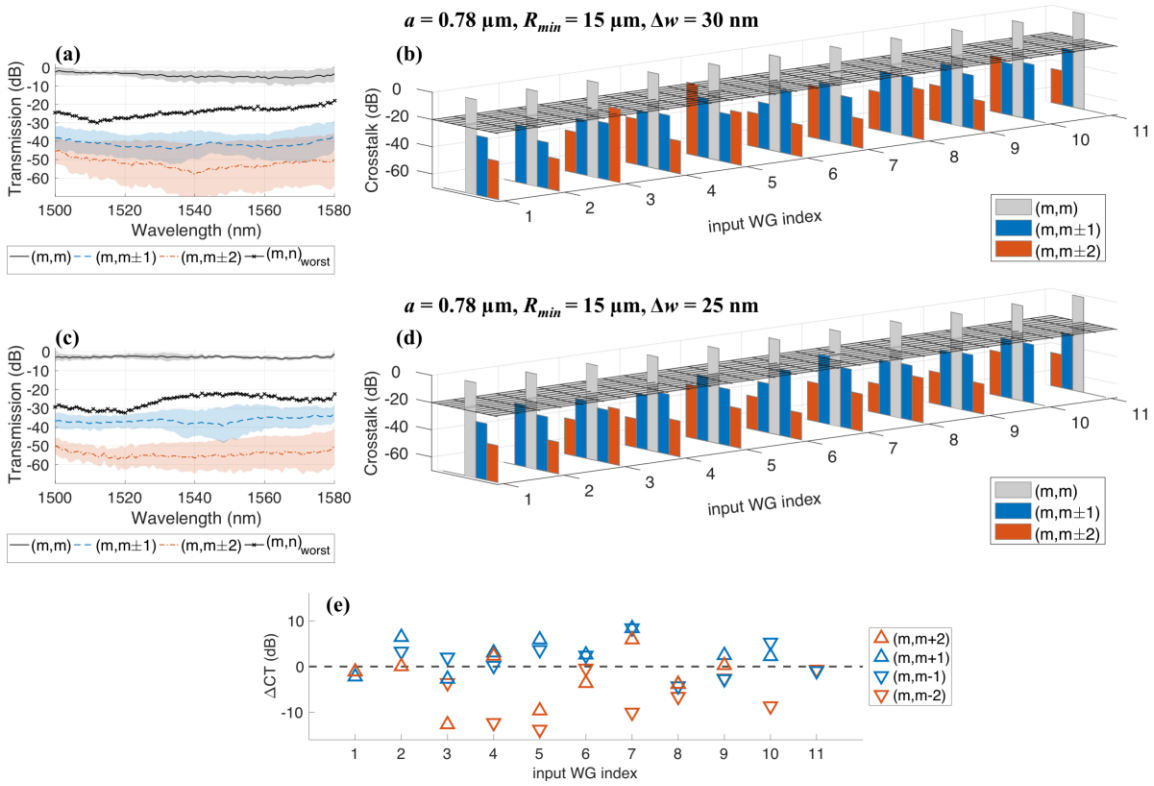

Figure 3.11 Experimental results for inter-waveguide pitch 0.78 µm at R_{min} = 15 µm when Δw = 30 nm and Δw = 25 nm. (a, c) statistics over channel number m for all direct channels and their 1st and 2nd nearest neighbors representing the average for each wavelength and standard deviation as the shaded regions, (b, d) maximum relative crosstalk for each channel and its 1st and 2nd nearest neighbors with gray plane marking -20 dB, and (e) the actual change in crosstalk when Δw is reduced to 25 nm.

The effectiveness of changing Δw =30 nm to 25 nm can be seen in **Figure 3.11(e)**. As expected, the crosstalk to second nearest neighbors are high for Δw = 30 nm; and it decreases significantly as Δw changes to 25 nm, with a small increase in crosstalk to first nearest neighbors. The crosstalk to third nearest neighbors $T(m,m\pm3)$ are very low and approach the noise floor of our measurement setup, and the crosstalk of even farther neighbors is sheer noise. Overall, the modified WGSL design pulls down the maximum relative crosstalk from -16.0 dB in the original

design to -19.9 dB in the reduced Δw design, which is significant improvement for many applications.

We further experimentally explored the crosstalk performance of various other WGSL bend configurations. We obtained complete results for 1.0 μm pitch at minimum bending radii of 10 μm (see **Figure 3.12**) and 5 μm (see **Figure 3.13**) with Δw = 30 nm.

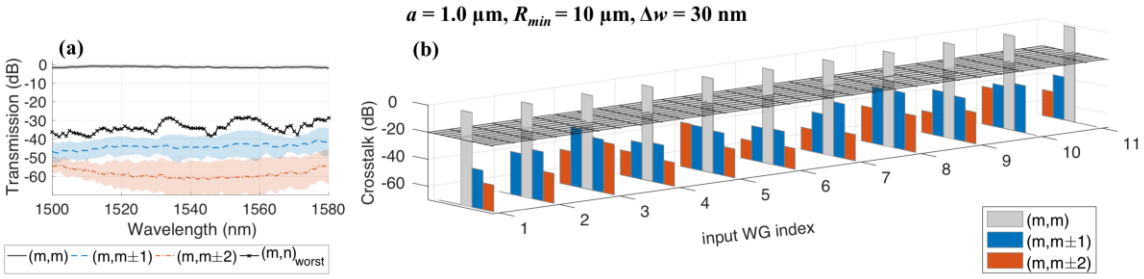

Figure 3.12 Experimental results for inter-waveguide pitch 1.0 μm at R_{min} = 10 μm when Δw = 30 nm. (a) statistics over channel number m for all direct channels and their 1st and 2nd nearest neighbors representing the average for each wavelength and standard deviation as the shaded regions, (b) maximum relative crosstalk for all 11 channels and their 1st and 2nd nearest neighbors with gray plane marking -20 dB.

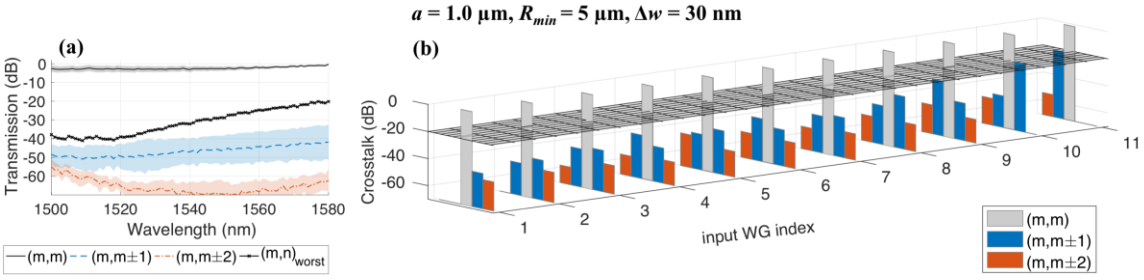

Figure 3.13 Experimental results for inter-waveguide pitch 1.0 μm at R_{min} = 5 μm when Δw = 30 nm. (a) statistics over channel number m for all direct channels and their 1st and 2nd nearest neighbors representing the average for each wavelength and standard deviation as the shaded regions, (b) maximum relative crosstalk for all 11 channels and their 1st and 2nd nearest neighbors with gray plane marking -20 dB.

We also obtained maximum relative crosstalk results for 0.78 μm pitch at minimum bending radii of 20 μm, 10 μm, and 5 μm with Δw = 30 nm. The maximum relative crosstalk for each WGSL bend configuration is summarized in **Figure 3.14(a)** along with a scanning electron microscope (SEM) image of a fabricated structure. Evidently, the maximum relative crosstalk decreases with increasing bending radius. In general, a maximum relative crosstalk of -20 dB is sufficient for most applications. Based on this value, the minimum bending radius for a 0.78 μm pitch is around 15 μm. Many applications can have relaxed pitch requirements, e.g., allowing for a 1 μm pitch. Under such conditions, **Figure 3.14(a)** shows that we can achieve a maximum relative crosstalk as low as -26.4 dB at a 10 μm minimum bending radius. Furthermore, for 1 μm pitch, we have demonstrated good performance for bending radii at 5 μm where we achieved respective maximum relative crosstalk of -19.6 dB.

One may wonder at what scale the WGSL reduces footprint compared to typical design rules in a general waveguide array. For straight waveguide arrays, the footprint scales with inter-waveguide pitch linearly. In bends, however, both the inter-waveguide pitch and minimum bending radius contribute to the footprint. The footprint of a bend can be estimated as

$$A_{bend} \approx \left(\frac{\theta}{2}\right)[(Na)^2 + 2(Na)R_{min}], \tag{3.16}$$

for any array bending at $\theta = (0, \pi]$ radians with $(N+1)$ waveguides of inter-waveguide pitch a, and minimum bending radius R_{min}. Interestingly, pitch contributes both quadratic and linear terms to the footprint, while the minimum bendings radius contributes a linear term. The footprint of an 11-channel array of various pitches and

bending radii is depicted in **Figure 3.14(b)**. Due to the quadratic dependence on the pitch, in the regime of large pitches, primary space savings come from minimizing the pitch, while secondary space savings come from minimizing the bending radius. For example, suppose we take an 11-channel waveguide array with a minimum bending radius of 15 μm and compare the overall footprint reduction of a 0.78 μm pitch WGSL to a 3.0 μm pitch uniform waveguide array. The area reduction is more than 6 times as shown in the first row of **Table 3.1**. If the pitch is already small, the area reduction due to minimizing bending radius can also be important. For example, we consider a 1.0 μm pitch WGSL, which allows the minimum bending radius to reach ~5 μm and results in a smaller area than a 0.78 μm pitch WGSL with $R_{\min}$ = 15 μm, as shown in **Table 3.1**. Based on these trends, it can be easily observed that a WGSL can provide significant saving of on-chip estate in bends. The several cases presented here offer a variety of options to be picked from based on system-level considerations in specific applications.

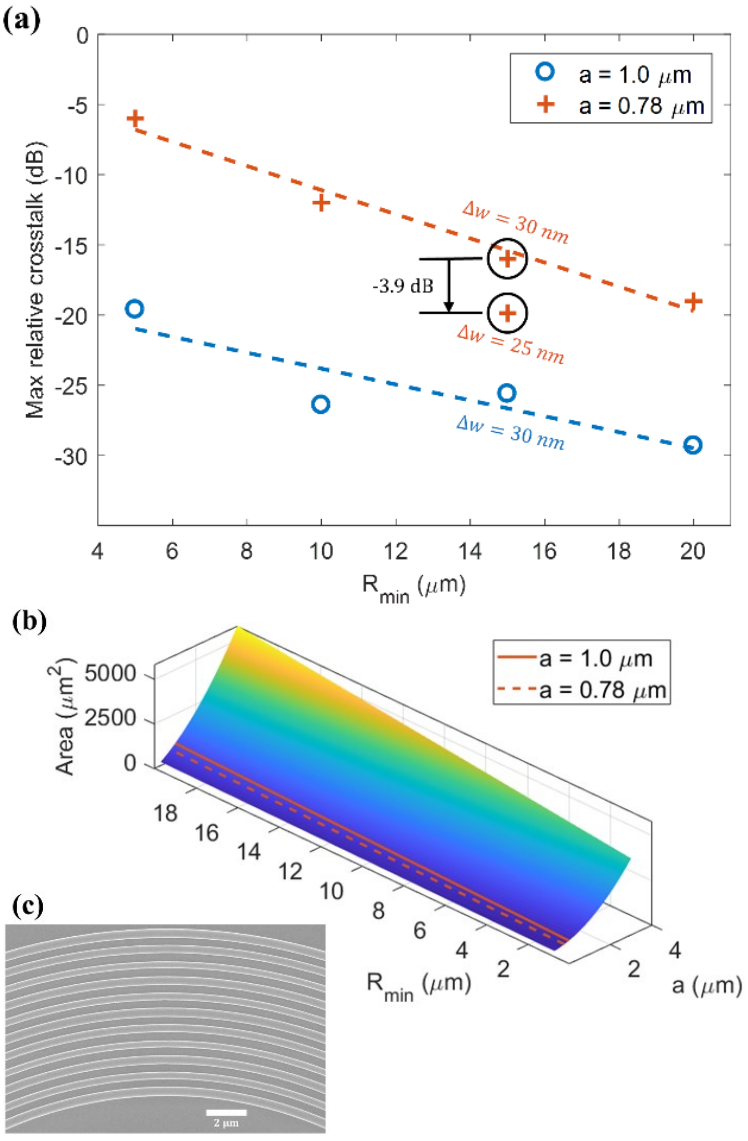

Figure 3.14 (a) Measured trends of maximum relative crosstalk with minimum bending radius for inter-waveguide pitch a = 0.78 μma = 0.78 μ m and a = 1.0 μma = 1.0 μ m. (b) Footprint of various 11-channel array U-bends. (c) SEM image of a fabricated WGSL bend. White bar represents 2 μm.

Table 3.1 Select designs of a WGSL U-bend

a [μm]	R_{min} [μm]	CT_{max} [dB]	A [μm²]a)	A_0 [μm²]b)	A/A_0
0.78	15	-19.9	519.9	3266	0.1592
1.0	5	-19.6	362.9	2229	0.1628

a) Area for a WGSL bend
b) Area for 3 μm pitch waveguide array bend (as a reference) with same R_{min}

3.7 Conclusion

In conclusion, we have explored the physics of guiding light through a waveguide superlattice bend and demonstrated that WGSLs can be implemented in bends with low crosstalk [57]. In the half-wavelength pitch regime (a = 0.78 μm), we have shown that the previous WGSL design can reach a minimum bending radius as tight as R_{min} = 15 μm with a maximum crosstalk of -16.0 dB. We observe that sometimes the second nearest neighbor can have higher crosstalk than the nearest neighbors. Such an effect stems from bend-induced skew of propagation constant, which destroys the phase mismatch structure in the WGSL for second nearest neighbors under certain circumstances. Interestingly, reducing the width difference Δw in the WGSL can help to overcome this issue. By changing Δw from 30 nm to 25 nm, we can push down the maximum crosstalk to -19.9 dB for a=0.78 μm and R_{min} = 15 μm. For applications that allow for 1 μm waveguide pitches, we can substantially reduce the minimum bending radius to 5 μm with -19.6 dB crosstalk. Finally, both cases of WGSL bends (a = 0.78 μm, R_{min} = 15 μm and a = 1.0 μm, R_{min} = 5 μm) have proven that they can achieve low crosstalk at small bending radii, thus potentially enabling high-density, flexible signal routing for a multitude of chip-scale applications.

Chapter 4: Closed-loop Control of Silicon Photonic Microring Transceivers

4.1 Background and motivation

Wavelength division multiplexing is a key multi-channel integration technique employing the idea of non-interacting light sources of different wavelengths as different parallel channels in a single waveguide bus. It scales up for capacity-crunching telecommunication networks via DWDM, which corresponds to a channel spacing of 50 GHz (0.4 nm) and even 25 GHz (0.2 nm). However, most of these WDM PICs are based on arrayed waveguide grating (AWG) filters and Mach-Zehnder interferometer (MZI) filters/modulators, both of which are relatively massive devices. For scaling up, the number of devices on a single PIC is limited to such devices.

Microring-based devices are an attractive alternative to AWGs and MZIs as they can function as a modulators or filters and they can be made to be incredibly small. However, there comes a tradeoff in such microring-based WDM PICs [11], thus requiring some robust techniques to overcome microrings' sensitivities to its fabrication and ambient environment.

4.2 Sensitivities of microring-based photonic integrated circuits

4.2.1 Fabrication sensitivities

The effective index of silicon is particularly sensitive to fabrication deviations. Specifically, the width of silicon waveguides can only be maintained within a certain tolerance level of fabrication, thus resulting in slightly shorter or wider waveguides.

Unfortunately, microrings in their very nature are highly sensitive to these deviations. The peak resonance can drastically shift – among other characteristics – and thus, they will not work as intended. The resonances of a microring can be estimated as

$$\lambda_{res} = \frac{n_{eff} L}{m}, \qquad m = 1,2,3, \ldots \tag{4.1}$$

where n_{eff} is the effective index of the waveguide, $L = 2\pi r_{ring}$ is the round-trip length of the microring, and m is the resonance mode number. A change in the effective index of the waveguide – such as when the waveguide width changes – will change the optical roundtrip length ($L_{opt} = n_{eff}L$) of the microring, and thus change λ_{res}. An estimation of this is depicted in **Figure 4.1**.

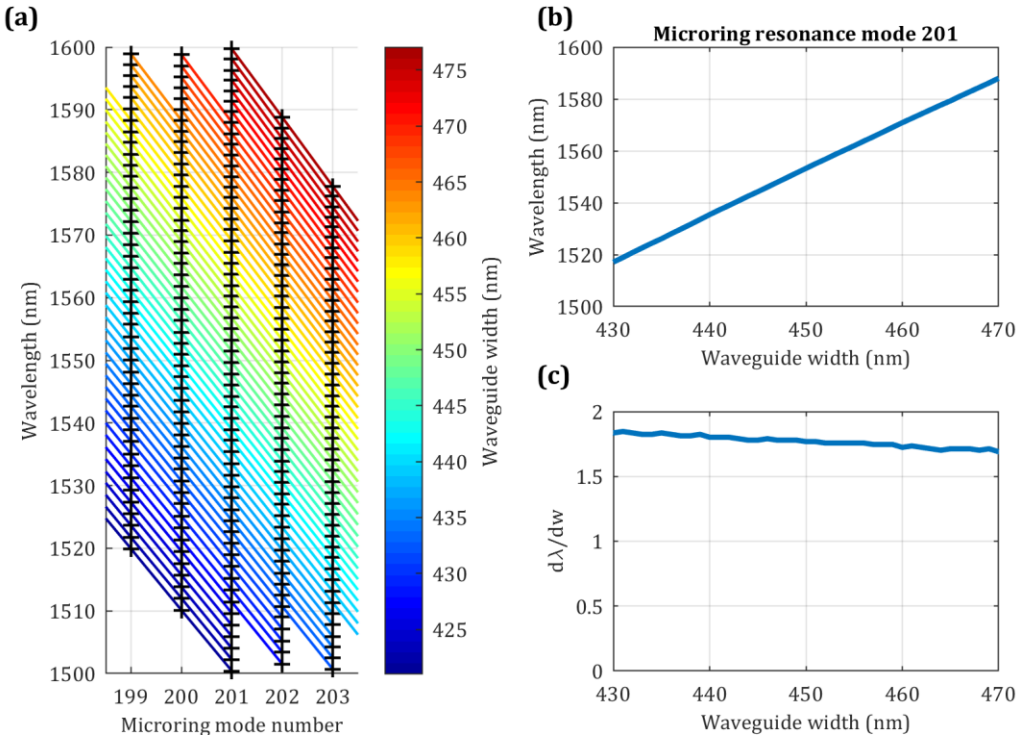

Figure 4.1 (a) Calculations of the resonant modes of a microring of radius r_{ring} = 20 µm, waveguide thickness of 260 nm, and varying waveguide widths. (b) Resonant wavelength of mode number 201. (c) Resonant wavelength shift per change in waveguide width.

In addition, this also changes the coupling efficiency between the bus waveguide and the microring resulting in a change in the extinction ratio and quality factor.

One may argue that post-fabrication trimming can compensate for the resonance shift, and it certainly does. However, trimming does little to mitigate for ambient effects such as temperature changes.

4.2.2 Temperature variation

On top of being sensitive to fabrication deviations, silicon has shown to have a high thermo-optic coefficient. In other words, a change in temperature will change the refractive index of silicon, therefore change the effective index of a silicon waveguide, and ultimately shift the wavelength resonance of a microring resonator. This temperature shift can come from several sources, including the global environment that the silicon PIC is in and heat coming from local sources such as electrical circuitry found in active PICs. An illustration of the sources of heat are shown in **Figure 4.2(a)** and **(b)**. It may be possible to thermally isolate such sensitive structures, but in a typical planar CMOS-based process, this could ultimately be impractical. Microrings are also particularly susceptible to self-heating effects [58], [59]. The accumulation of energy in a microring near its resonant modes can heat up the microring and thus change the refractive index – see **Figure 4.2(c)**. This results in a folding in the transmission spectrum and unwanted hysteresis when using it as a modulator. A typical way to avoid this is lowering the input power of the laser source.

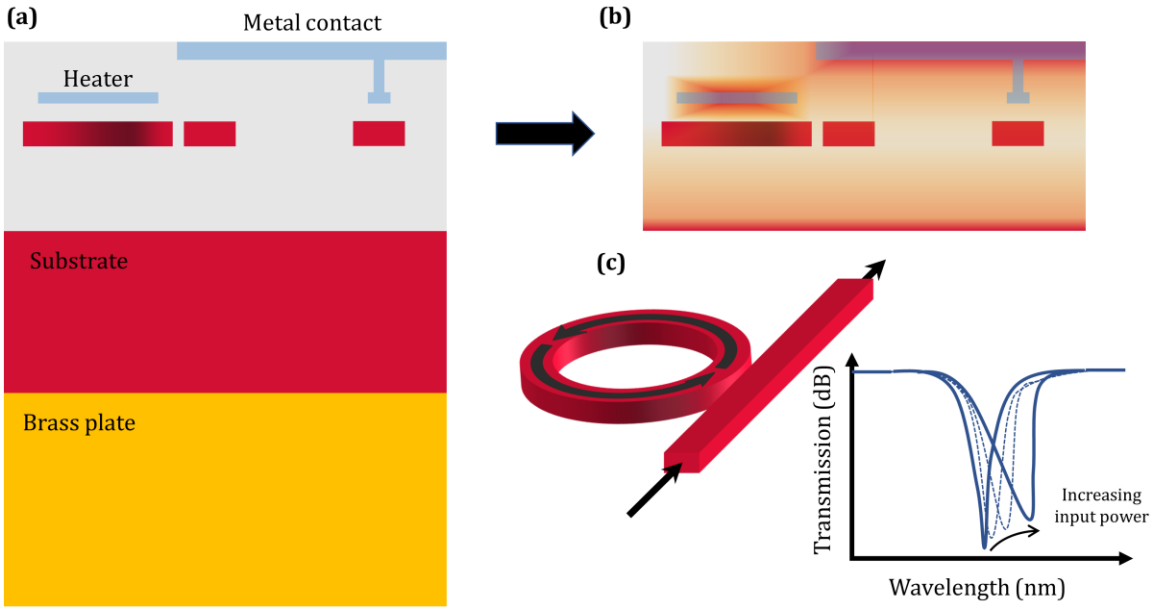

Figure 4.2 (a) Sources of heat around a silicon PIC, which could include refractive index tuning heaters, upper layer electrical routing, and substrate temperature which can transfer heat from typical packaging such as a brass plate. Ambient temperature will also change the global temperature of the PIC. (b) An illustration of the temperature gradients around the silicon waveguide layer. (c) Self-heating from resonant microring modes can fold the optical spectrum.

4.3 System-on-a-chip with photonic integrated circuits

Despite its sensitivities, microring-based PICs can still be employed in future scalable systems, such as transceivers. Silicon PICs will more than likely be packaged with electronics either via chip stacking or flip-chip bonding, thus allowing electronic-based control systems to mitigate for the performance degradations of silicon photonic systems. One approach is through the use of ohmic heaters to thermally tune microrings [11], [60]. This, in turn, changes the negative characteristic of a microring with temperature into a feature, resulting in highly efficient resonance shifting for microring-based modulators and filters [61], and wavelength tuning in WDM

transmitters and receivers have been demonstrated [12], [62]–[69]. Furthermore, automated microring filter tuning has also been explored [64]–[67], [69]–[73].

Unsurprisingly, designs of future microring-based WDM transceivers will require compact, simple, and low power closed-loop control systems for thermal adaptation. Compact, simple PICs can increase the bandwidth per footprint and minimize the energy requirement per bit. Low-power systems are necessary for short-reach optical communications, and silicon photonics integrated with efficient electronics and control algorithms can achieve that.

In this chapter, we aim to use a single-monitoring signal that can be used for multi-microring tuning to simplify the closed-loop control loop. Furthermore, we implement simple gradient-based algorithms to minimize the computation overhead needed to tune and track the resonances of microring-based devices. This is also known as wavelength locking. We argue that such approaches can be used to scale up microring-based transceivers (systems employing a transmitter and receiver) easily.

4.4 **Closed-loop control of a silicon photonic integrated circuit**

Closed-loop control is method in which an output (or outputs) of a system is fed back into the system to modify or maintain the output. It has advantages over open-loop control since it uses real-time information rather than pre-configured settings, thus maintaining high accuracy and repeatability in the control. Moreover, feedback reduces sensitivity to internal and external disturbances, and hence makes the system or device more robust. **Figure 4.3** is an illustration of how a closed-loop control system would look like for a PIC system.

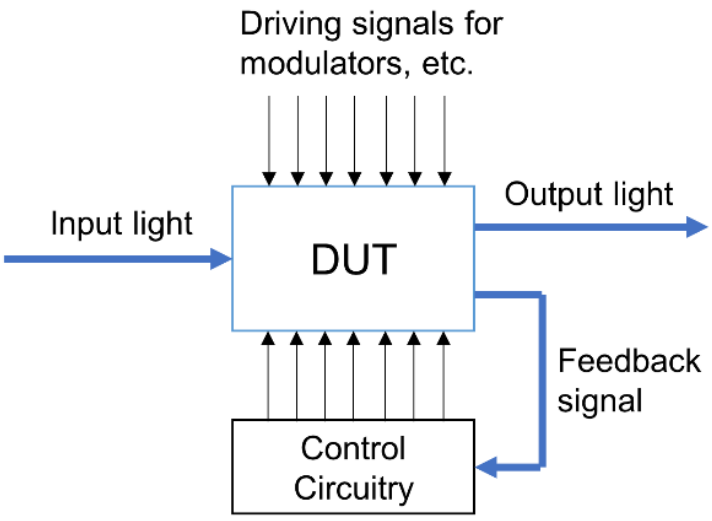

Figure 4.3 Closed-loop control of a device under test (DUT).

While the illustration shows the input and output as light, these signals can be electrical. The electrical-optical interfaces can be somewhere on the DUT. Likewise, the feedback signal can be electrical or optical. The device or devices that convert such feedback signals into a readable format for the control circuitry are often known as encoders. Control circuitry can take many forms. In the simplest sense, it can be an external photodetector that is monitored by an ammeter that is read by a person; and that person uses various control knobs/buttons to modify the DUT output – i.e., manual closed-loop control. For automatic closed-loop control, the control circuitry can be a microcontroller, field programmable gate array (FPGA), proportional-integral-derivative (PID) controller, or some custom circuitry. Often, this control circuitry outputs electrical signals to either drive on-chip ohmic heaters or other electro-optic interfaces such as PN junctions. It should be noted that these closed-loop control systems should be designed with low power and bandwidth density in mind, especially when scalability is considered.

In the following sections, we will be using a microcontroller and its peripherals to close the control loop in a couple of our WDM devices – see **Figure 4.4**. The microcontroller (an Analog Devices ADuC7122) has a clock speed of 41.78 MHz and contains multiple analog-to-digital converters (ADCs) for reading voltage signals and multiple digital-to-analog converters (DACs) for driving voltage signals. All of them have 12-bit resolution. To reach appropriate dynamic range, we use electrical amplifiers to increase the reading and driving signals.

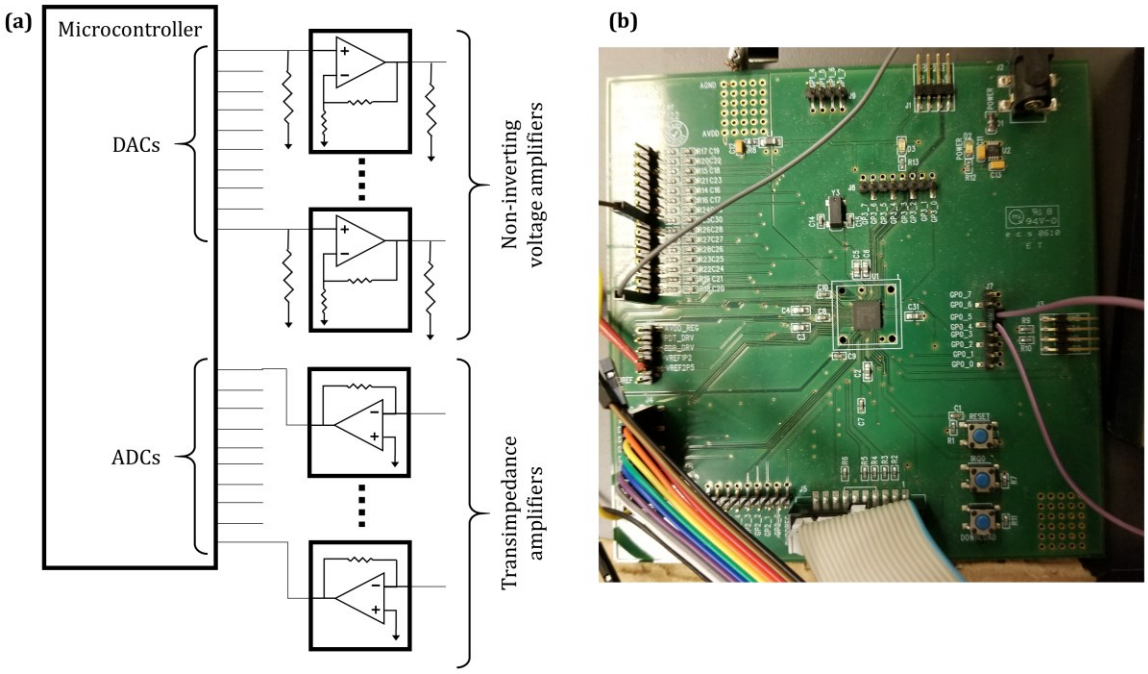

Figure 4.4 (a) A schematic of a microcontroller, its peripherals, and optional amplifiers. (b) An image of the ADuC7122 microcontroller evaluation board.

4.5 **Wavelength locking of a WDM transmitter**

One of many critical devices on a silicon PIC is the optical modulator, which in our case is a microring-based modulator. Using microring modulators is particularly attractive in WDM transmitters, especially if integrated with multi-wavelength

sources, simplifying the WDM transmitter architecture [74]–[77]. Previous reports have demonstrated wavelength locking in microring-based WDM transmitters using optical power detection in either the through or drop ports and adjusting the power to a prescribed level, or using more complicated detection schemes including homodyne detection, bit error ratio (BER) optimization, and photocurrent-induced by two-photon absorption (TPA) [71], [78]–[85]. These reported techniques, however, require at least one monitor for each microring and only realize individual microring control. To this end, scaling up would require complicated on-chip electronic routing to access each microring's monitor. We instead proposed a more simple and universal control method [86], providing simultaneous wavelength locking of a microring modulator array with a single monitoring signal. In this section, we present and demonstrate our proposed closed-loop control system, which detects the radio frequency (RF) components of a modulated optical signal in the through port bus of the microring modulator array, simultaneously locks multiple microring modulators by automatically finding the best heater power bias to achieve maximum modulation.

The single monitoring signal is essentially an encoder which utilizes a monitoring photodetector (MPD) and an RF power detector – see **Figure 4.5**. The key idea is to use the RF power detector to measure the non-DC components of the optical signal in the through port and use the encoded signal to control the microrings' corresponding heaters. If all the microring modulators are off resonance to a multi-wavelength source, no optical modulation would be evident. This would mean the RF power detector would measure zero power. On the other hand, if all the microring

modulators are on resonance to the multi-wavelength source – particularly, each microring is modulating its own wavelength channel – maximum optical modulation would be evident. The RF power detector would then return a global maximum signal, and, thus, the WDM transmitter would be wavelength locked.

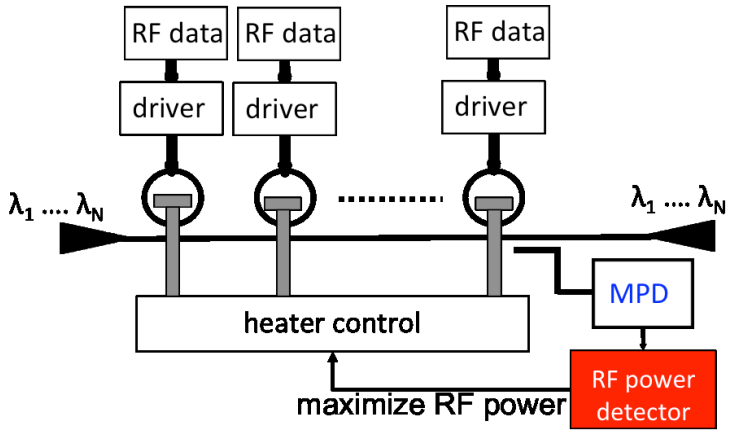

Figure 4.5 Closed-loop control of a microring modulator array with RF power detection in the through port.

However, this is only guaranteed under certain conditions of the WDM transmitter: (1) the encoder cannot detect the high beating frequency caused by heterodyning of the multi-wavelength light source, (2) the data applied to different microring modulators are decorrelated, and (3) the RF power drops from modulation loss caused by two microrings simultaneously modulating the same wavelength. The two conditions are trivially satisfied as the channel spacing between two WDM channels is typically more than 50 GHz and the data in different channels are typically independent in real applications. The third condition can be satisfied with a properly programmed control algorithm.

This proposed locking mechanism provides multiple advantages. First, the largest RF power corresponds to the largest optical modulation amplitude (OMA) in on-off keying (OOK) modulation. The RF power of an on-off keying (OOK) signal is

proportional to the square of the peak-to-peak voltage of the signal; and the peak-to-peak voltage of a transimpedance amplified photodetector is linearly proportional to the OMA. This by-product advantage aids in determining the best heating power for tuning the microring modulators. Secondly, the MPD, its amplifier, and the RF power detector of the encoder do not need to have high bandwidth. The RF power detector only needs to provide a monitoring signal output rather than full data decoding. In a non-return-to-zero (NRZ) OOK signal, the RF power below 20 percent of the fundamental frequency would occupy around 40 percent of the total RF power. A bandwidth in the range of 10 percent to 50 percent of the fundamental frequency is sufficient for most cases. Thus, the RF power detector can be made in with simple components such as rectifiers, which would only require a microwatt-order power budget. Thirdly, this mechanism gives the ability to simultaneously wavelength lock multiple microrings. This is due to the incoherent summation of the RF power from different wavelength channels. To this end, the use of a single MPD for multiple microrings can significantly reduce the power budget of transmission systems.

We experimentally demonstrated this proposed locking method using a silicon photonic chip with three microring modulators coupled to a bus waveguide – the total packaged microring modulator count was twenty. Each microring contains a PN junction along the curvature of the microring for high-speed modulation and a resistive heater over microring cladding for thermo-optic tuning [77]. We first determined the microrings' sensitivity to temperature as depicted in **Figure 4.6**. The packaging is mounted on a thermoelectric cooler (TEC) which uses the Peltier effect to control temperature. A thermocouple is placed near the silicon chip to feedback the

actual temperature to the TEC driver and ensure the set and actual temperatures are the same. When the temperature is changed from 25 °C to 70 °C, we measure a resonance red shift of approximately 3.8 nm. This corresponds well with the theoretical red shift of 0.08 nm/°C.

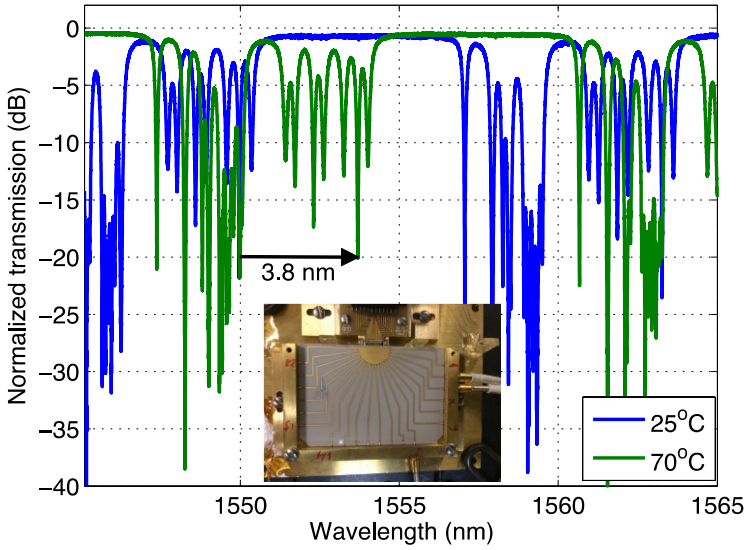

Figure 4.6 Spectra of the silicon photonic chip under different temperatures. The inset shows a picture of the device packaged with RF and DC boards.

For high-speed testing, three CW wavelength-tunable laser sources are multiplexed and launched into the chip each with an approximate power of 3 dBm. The output of the chip is power-divided into two branches, one branch tapping 20 percent of the output to be launched into the MPD. This branch, which contains a power of about -15 dBm when three microrings are modulating, is measured by an external 10-Gb/s optical receiver with an optical bandwidth from 1200 nm to 1600 nm. The electrical output of the receiver is measured by a RF power detector (Mini-Circuits, Zx47-60-S+) with a detection frequency range of 10 MHz to 8 GHz. The RF power detector converts the RF power into a voltage signal which is read by an ADC

in our microcontroller. Three DACs of the microcontroller are used to drive the microrings' heaters. The other branch of the output optical signal is amplified by an erbium-doped fiber amplifier (EDFA) and filtered by a tunable filter with a 3dB-bandwidth of ~0.5 nm to select the channel for optical eye diagram acquisition or BER testing. The drive signal is a 20-Gb/s pseudorandom binary sequence (PRBS) with a length of 2^{15}-1, a voltage swing of 5 V and a DC bias of 3 V.

Before actual wavelength locking, however, we first verified that the RF power signal can accurately indicate the OMA magnitude and incoherent summation of RF power occurring within WDM channels. **Figure 4.7(a)-(d)** illustrate the optical eye diagrams for one of the channels at different RF power levels. With increasing RF power, the eye opening widens the maximum RF power indicates the microring is modulated at the best heater bias. The incoherent RF power summation is depicted in **Figure 4.7(e)**, where the RF power is measured versus the number of microrings tuned to their best heater bias. The trend clearly shows that RF power is linearly proportional to the number of microrings locked. **Figure 4.7(f)** depicts the corresponding change in RF power versus each microring's heater bias. Troughs in the RF power indicate the points when several microrings are simultaneously modulating the same wavelength channel, resulting in modulation loss. Note that for a microring that utilizes wavelength shifting as a switching mechanism, modulation can occur on either side of its spectral resonance, hence the existence of another peak close to the maximum RF power.

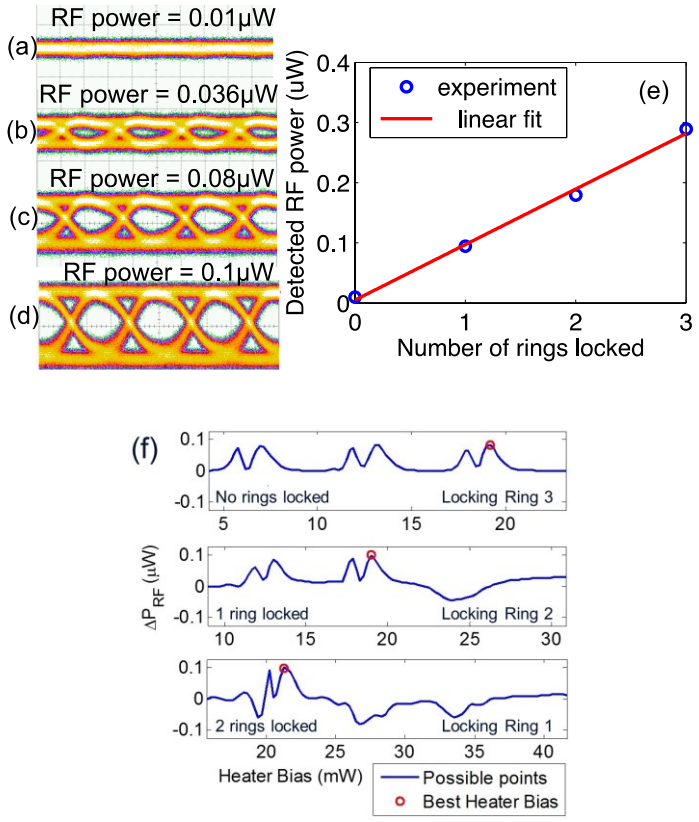

Figure 4.7 (a-d) Optical eye diagrams with increasing monitor RF power. (e) Detected RF power as a function of number of rings locked to show incoherent summation of RF power among different channels. (f) Change in RF power for each ring as a function of their respective heater power biases.

We demonstrated active wavelength locking of the microrings while the temperature of the packaged silicon photonics device cycles between nominal temperatures of 27 °C and 70 °C. The microcontroller is programmed to maximize the incoherent RF power via the gradient method. However, blindly doing so may result in local maximums. To avoid this, we initialize each microring at the same wavelength, specifically a wavelength not being used to transmit data – this would also minimize the RF power detected by the encoder. Then, each microring is consecutively tuned first by an exhaustive search until there is a notable gradient in the RF power. At that point, the gradient method takes over for a quick convergence to the best heater bias

for microring. During this process, each microring is locked to a desired transmitting wavelength. After the final microring is tuned, the algorithm uses the stochastic gradient method online to ensure wavelength tracking with temperature drift. Note that once a microring is tuned to a wavelength, the other microrings will not compete for it as the incoherent RF power will decrease when other microrings try to tune towards the same wavelength.

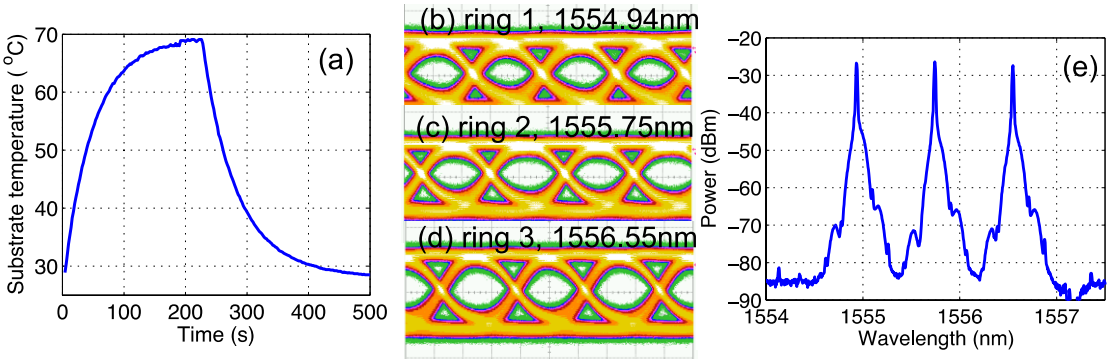

Figure 4.8 (a) Temperature cycle during active locking. (b)-(d) Eye diagrams of three channels during the temperature cycle shown in (a), while the rings are simultaneously locked to the input wavelengths. (e) Locked modulation spectrum.

Figure 4.8(b)-(d) illustrate the optical eye diagrams of the three modulating microrings, depicting that all the eye diagrams have a reasonable opening over a temperature change of more than 40 °C. **Figure 4.8(e)** shows the spectrum of the modulated WDM signals. We further characterized the BER of the three microrings with different optical launching powers to a 40 Gb/s optical receiver and a BER tester. The results are shown in **Figure 4.9**. With enough launching power (at least -8 dBm, optical receiver limiting), the three microrings can maintain a BER no greater than 10^{-9} within an approximate temperature change just over 40 °C, thus demonstrating the robust locking of our method. Further increasing of the launching power (at least

-6 dBm) can allow zero BER to be observed for the three microrings during the temperature cycling period of about 500 seconds.

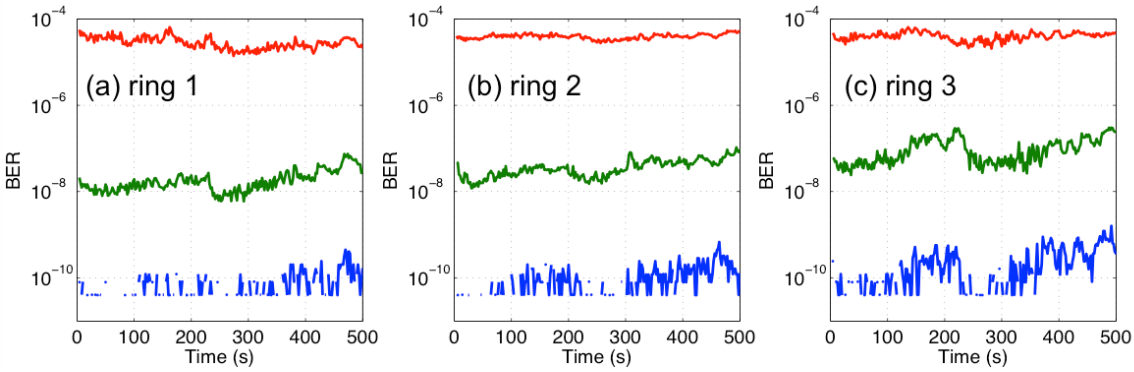

Figure 4.9 BERs during the temperature cycles for the three wavelength-locked microrings. Different curves represent the cases with different launching power to the optical receiver before BER tester with the red, green, and blue curves representing low, medium, and high optical powers respectively. (a) ring 1 with launching powers -11 dBm, -10 dBm, and -8 dBm; (b) ring 2 with launching powers -12 dBm, -10 dBm, and -9 dBm; and (c) ring 3 with launching powers -13 dBm, -11.5 dBm, and -10.5 dBm.

The same locking technique can be applied to other advanced intensity modulation formats such as discrete multi-tone modulation (DMT). DMT has been used in the past for digital subscriber lines (DSL) to boost the bandwidth of these legacy copper technologies. Likewise, it can be used to further boost the bandwidth of optical transceivers. A 10 to 20 Gb/s OOK optical device can boosted to achieve 100 Gb/s by converting the modulation to DMT. This technique utilizes many subcarriers and optimizes the modulation format for each one based on its achievable signal-to-noise ratio (SNR) in the optical link – this is typically known as bit loading. In [77], high-capacity DMT modulation using silicon microring modulators was demonstrated for the first time with a channel rate of about 90 Gb/s. In [87], the channel rate was pushed to 128 Gb/s. Here, we used the same setup as in [77] to

generate and detect DMT signals with three microrings wavelength-locked – see **Figure 4.10**.

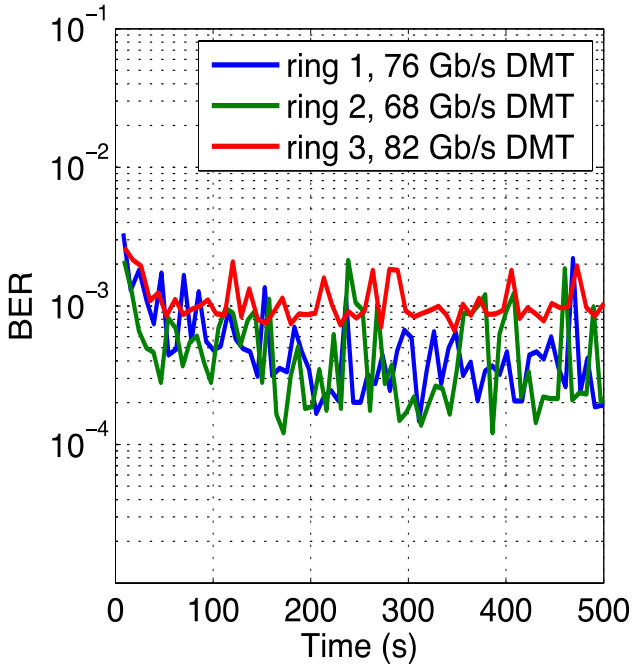

Figure 4.10 BERs for DMT modulation while three rings are wavelength-locked.

The input laser power was set to 10 dBm per wavelength, and we demonstrated wavelength locking with three microrings modulated at 76, 68, and 82 Gb/s using DMT. During temperature cycling, all the BER was measured to be below 3.8 x 10^{-3}, which is the hard-decision forward error correction (FEC) limit with 7 percent overhead. Thus, our locking technique is potentially applicable to advanced intensity modulation formats such as DMT.

4.6 **Wavelength locking of a WDM receiver**

Another critical device in a silicon PIC is the WDM receiver which typically consists of optical filters such as microrings or arrayed waveguide gratings (AWGs) to

demultiplex and/or reroute each wavelength channel. As with the transmitter case, we look at using a microring-based WDM receiver and look at avenues to close loop control them for wavelength locking and tracking. We also look at how to close loop control a WDM receiver that is polarization diverse since in most short-reach optical interconnects not on-chip, there is a glass single mode fiber that is not polarization maintaining. Thus, the polarization is changing as the signal propagates through the bending and pinching of the optical fiber. Silicon photonic WDM receivers are highly sensitive to the polarization input, which considerably limits the receivers' performance when the polarization changes. Thermal tuning of polarization insensitive WDM receivers has been done in previous works [12], [62], [63]; and automated microring filter tuning has been explored [45], [64]–[67], [69]–[71], [88]. Designs of future WDM receivers should incorporate both polarization insensitivity and mutli-wavelength channel locking while maintaining a compact, simple, low power closed-loop system design.

We take a polarization insensitive receiver similar to those in [12], [63], but develop an automated system with some key differences [89]. A schematic of the system is depicted in **Figure 4.11**. As with the WDM transmitter, we utilize a single MPD integrated at the combined through port of the WDM receiver, which would still allow locking of multiple wavelengths simultaneously. Instead of using the optical WDM signal's RF power as in the transmitter, we exploit the DC component of the signal to provide closed-loop control.

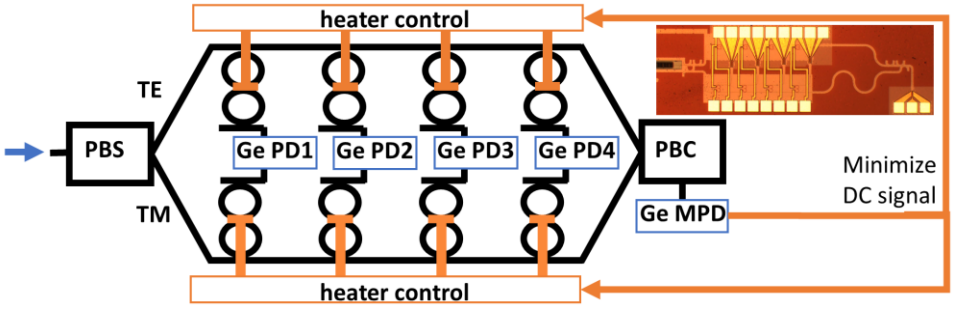

Figure 4.11 Proposed closed-loop control of a polarization insensitive microring WDM receiver with the detection of the DC signal in the combined through port. PBS/C: polarization beam splitter/combiner; PD: photodetector; MPD: monitoring photodetector. Picture of an unpackaged fabricated WDM receiver in the inset.

In our polarization insensitive WDM receiver, the input WDM signal is split into two polarization branches by a polarization beam splitter (PBS), and each branch is wavelength demultiplexed by cascaded second-order microring filters and photocurrent combined with germanium photodetectors (PDs) at the drop ports between the branches. The through ports of each branch are combined by a polarization beam combiner (PBC) and then are detected by an MPD. The MPD outputs an electrical signal that is proportional to the amount of optical power. This electrical signal can be used as feedback to the circuitry that simultaneously controls all the microring heaters. When no microrings are locked to the input wavelengths, the MPD returns the maximum signal; and when all microrings are properly locked and all the wavelength channels are dropped to the PDs, the MPD returns the minimum signal.

As with the transmitter, the main advantage of this locking mechanism is simplification of the optical circuit by exploiting only one monitoring signal. Also, by requiring only the DC components of the WDM signal, which is simply proportional to the intensity of the light incident on the photodetector, the mechanism is

independent of the modulation format. More importantly, by only using the DC components, we can convert the photocurrent outputted by the MPD to a voltage with a single resistor, requiring no additional power to amplify the MPD signal. This design, however, has an inherent pitfall due to its polarization diversity scheme. When a multi-wavelength signal is present the MPD signal will have multiple minima. **Figure 4.12** illustrates the range of MPD responsivities of a WDM receiver configured for 2 input wavelength channels.

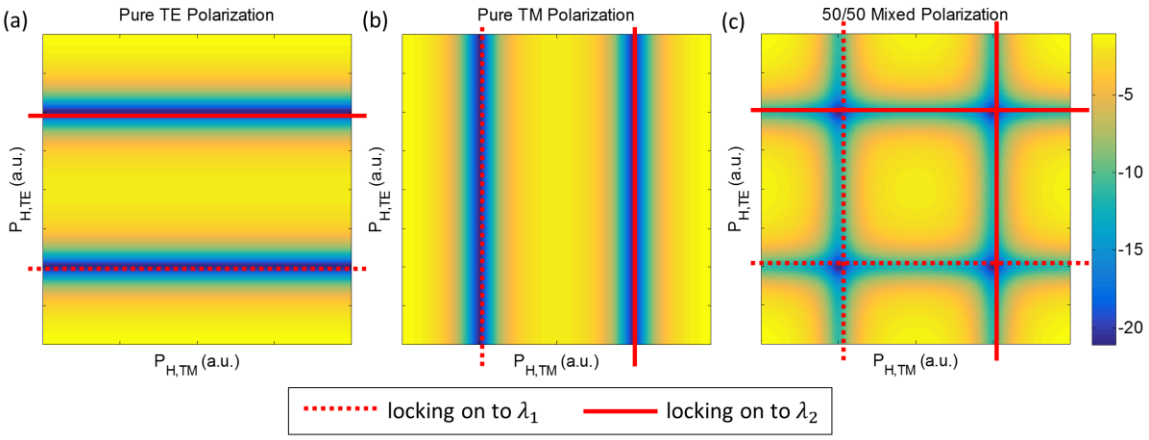

Figure 4.12 Illustrated calculations of MPD responsivity versus the heating powers of the first pair of microring heaters for (a) TE, (b) TM, and (c) mixed polarization inputs. An intersection of a solid and dashed line represents a minimum with no demultiplexing.

When the input signal is purely TE-polarized or TM-polarized, there is no problem – a minimized MPD signal will result in proper demultiplexing as shown in **Figure 4.12(a)-(b)**. However, when the input signal is of a mixed polarization – which would undoubtedly occur in real applications using no polarization maintaining fibers – potentially no demultiplexing could occur despite having a minimized MPD signal. This situation is depicted in **Figure 4.12(c)**. In other words, the first filter in the TE branch could properly lock on to the first wavelength channel,

but the first filter in the TM branch could lock on to the second wavelength channel. We prevent this pitfall by tuning each TM microring filter proportionally to the tuning of its corresponding TE microring filter and vice versa. This is possible if we can calibrate the device based on its microring filters' room temperature resonances and tunability By using proper calibration of the TE and TM microring heaters, we can demultiplex at any polarization input.

The wavelength locking and tracking algorithm is implemented in the same microcontroller used in the transmitter; however, this time, we used the adaptive gradient (ADAGRAD) method, which has proven often to have faster convergence then that of the classic stochastic gradient method. To help avoid microrings locking onto the wrong wavelength, the microrings are tuned such that the resonance of each microring filter is set at a resonant wavelength away from the wavelengths of the received WDM channels. This maximizes the optical power to the MPD and allows the largest dynamic range to lock onto to each wavelength channel. The ADAGRAD method is then used to minimize the signal, consecutively tuning and locking each microring filter in the TE branch on to a wavelength channel. The order of the locking can be configured into the algorithm. Each TM microring filter is paired (or mapped) to its corresponding TE microring filter, based on their room temperature resonances and tunabilities. Therefore, when the resonance of a TE microring filter is changed by a certain heating power, the TM microring filter is tuned proportionally to that change, locking on to the same wavelength channel and avoiding the possibilities of partially demultiplexed signals. Subsequent fine tuning of the TM microring filters is done with the ADAGRAD method. After the final microring filter is initially tuned, the

ADAGRAD method is endlessly run over each microring filter to maintain the best microring tuning with temperature drift.

We demonstrated this locking method on a real silicon photonic polarization insensitive WDM receiver as shown in the inset of **Figure 4.11**. The free spectral range (FSR) of each microring is about 6.4nm with room temperature resonances and tunability experimental data listed in **Table 4.1**. TE branch microring tunability is inherently larger than that of the TM branch due to larger optical confinement in the waveguides. Tunability variation can be attributed to different metal volume in the microrings' proximities. The average resistances of the metal heaters are 485 Ω and 588 Ω for TE and TM microrings, respectively. The heater metal width is 2 μm, and each second-order ring has a single heater which is symmetrically laid on its two rings to relax the misalignment impact from fabrication errors. Similar devices have shown that the worst-case channel crosstalk for 200 GHz and 100 GHz channel spacing is less than -21 dB and -12 dB respectively [63]. The fiber-to-PD responsivities of the MPD as the device temperature changes from 23 °C to 60 °C are shown in **Figure 4.13**. The wavelength red shift is approximately 0.092 nm/°C.

Table 4.1 Microring Filter Resonance and Tunability

	λ (nm)[a]		Tunability (pm/mW)	
Devices	TE	TM	TE	TM
PD1 microring filters	1546.54	1549.37	70.1	27.5
PD2 microring filters	1547.48	1549.76	63.1	36.6
PD3 microring filters	1548.41	1549.77	74.7	43.4
PD4 microring filters	1549.49	1550.45	70.9	47.5

[a]Device is at 23 °C

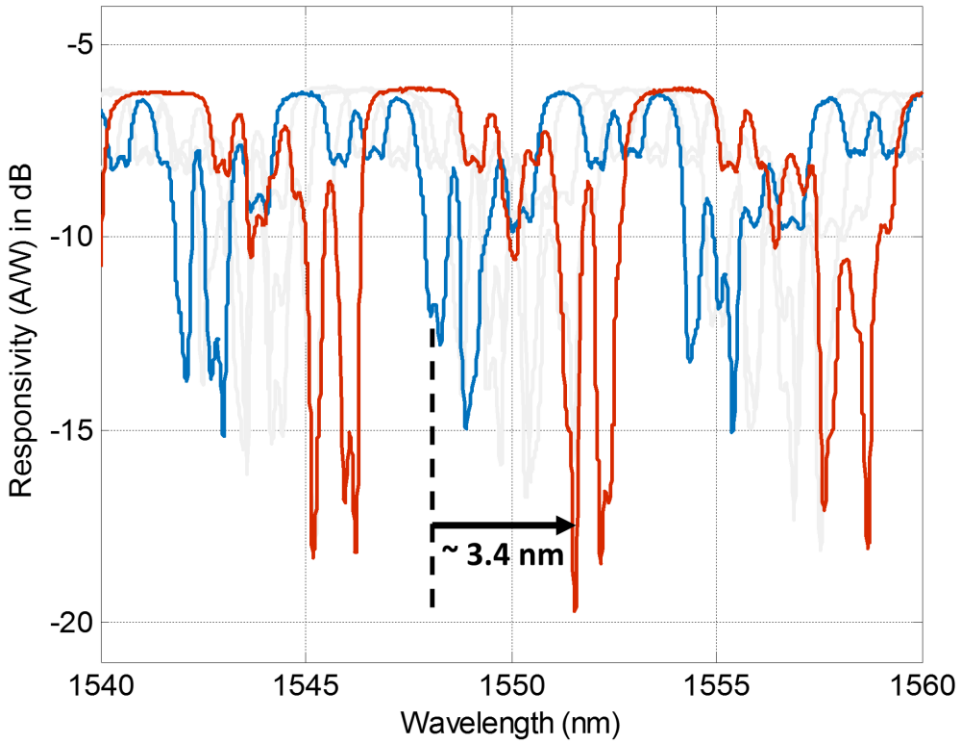

Figure 4.13 Fiber-to-PD responsivity of the MPD when the WDM receiver is at 23 °C (blue) and 60 °C (red) – intermediate responsivities in light gray. The microrings in the TE branch are more sensitive to the device temperature than those in the TM branch, resulting in a change to the general responsivity profile at higher temperatures.

In this experiment, we launch four WDM channels of 1550.52 nm, 1551.72 nm, 1552.93 nm, and 1554.13 nm (150 GHz channel spacing) through a wideband optical modulator driven by an OOK PRBS with a length of 2^{15}-1, generating a 4 x 25 Gb/s WDM signal. The modulated WDM signal passes through a 5 km long optical fiber with a dispersion of 17 ps/km/nm to decorrelate the signals. In real applications, each wavelength channel should have different data streams and, thus, will be inherently decorrelated. The WDM signal also passes through a polarization scrambler, is amplified by an EDFA, and is finally launched into a packaged receiver chip.

Each drop port PD is reverse biased at 2 V and can be used to observe the electrical eye diagram when their corresponding microrings are locked onto a

wavelength channel. Since the signals are decorrelated by the 5-km fiber, we can determine if there is any significant channel crosstalk in the eye diagrams. The through ports of each polarization branch are combined to be detected by the MPD, and a 1 kΩ resistor converts the MPD's photocurrent into a voltage signal. This simple conversion provides enough dynamic range, so a transimpedance amplifier (TIA) is unnecessary, resulting in a zero-power consumption feedback signal.

We demonstrated active wavelength locking of a 4 x 25 Gb/s WDM signal as shown in **Figure 4.14.** At initialization, we de-tuned the microring filters and sent most of the optical power to the MPD. At this point, the microring filters are set to have some resonant wavelength between 1555.33 nm and 1555.72 nm. Then we consecutively locked on to each wavelength first on the TE branch and finally on the TM branch. While locking, the heaters have ambient effects – thermal crosstalk – towards microrings for which they are not designed to tune. The algorithm systematically compensates for this and fine tunes the microrings. We configured the system such that the locked wavelength for PD1 is 1550.52 nm, PD2 is 1551.72 nm, PD3 is 1552.93 nm, and PD4 is 1554.13 nm. In this configuration, their corresponding microring filters have the maximum tuning range for when the device temperature increases from room temperature. The heaters are configured to tune each microring filter approximately over a 7.5 nm range – more than one FSR.

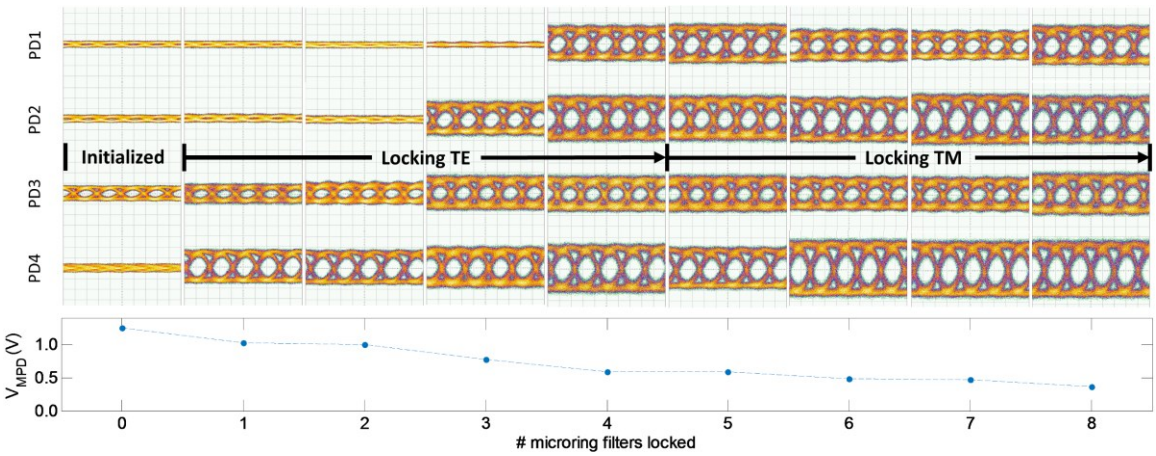

Figure 4.14 Locking and demultiplexing of a mixed polarization 4×25 Gb/s WDM signal. Initialization runs only a few times, hence the eye diagram for PD3 not being completely closed before locking. After each microring filter locks on a wavelength, the corresponding eye diagram widens. In some cases, due to thermal crosstalk, the currently controlled heater slightly tunes/detunes neighboring microrings, but the algorithm fine tunes the microrings after initial locking.

Furthermore, we demonstrated tracking while a temperature controller cycles the device temperature between 23 °C and 60 °C and a polarization controller scrambles the input polarization. The 150 GHz WDM spectrum observed at the input of the device is shown in **Figure 4.15(a)**, modulating at 25 Gb/s, and the temperature applied to the entire device over time is shown in **Figure 4.15(b)**. **Figure 4.15(c)** illustrates the electrical eye diagrams from the four high-speed PDs with actively tracked and simultaneously demultiplexed 25 Gb/s OOK channels. All eye diagrams are shown to have a clean opening during polarization scrambling and over a temperature change just over 37 °C. PD3 has the worst eye diagram due to higher polarization dependent loss from its microrings. The main limitation to the temperature control range is the used epoxy between the input/out fibers and the facet of the device, which begins to liquefy at temperatures near 70 °C. To avoid increased insertion loss due to misalignment, we set the max temperature to a safe

60 °C. We have also characterized the BER for each wavelength channel as depicted in **Figure 4.15(d)**. The 4 x 25 Gb/s WDM signal is launched with various launching powers, demonstrating reasonable demultiplexing. No TIAs are packaged with chip, so instead we use an external power amplifier with a gain of 30 dB, limiting the sensitivity. In these experiments, we have maintained tracking with temperature rates ≤0.64 °C/second with BER rippling roughly maintained within an order of the average BER. Channel power fluctuation should not drastically affect the closed-loop control unless the fluctuation is on the order of the control circuitry's electrical bandwidth, which may result in an incorrect estimation of the gradient.

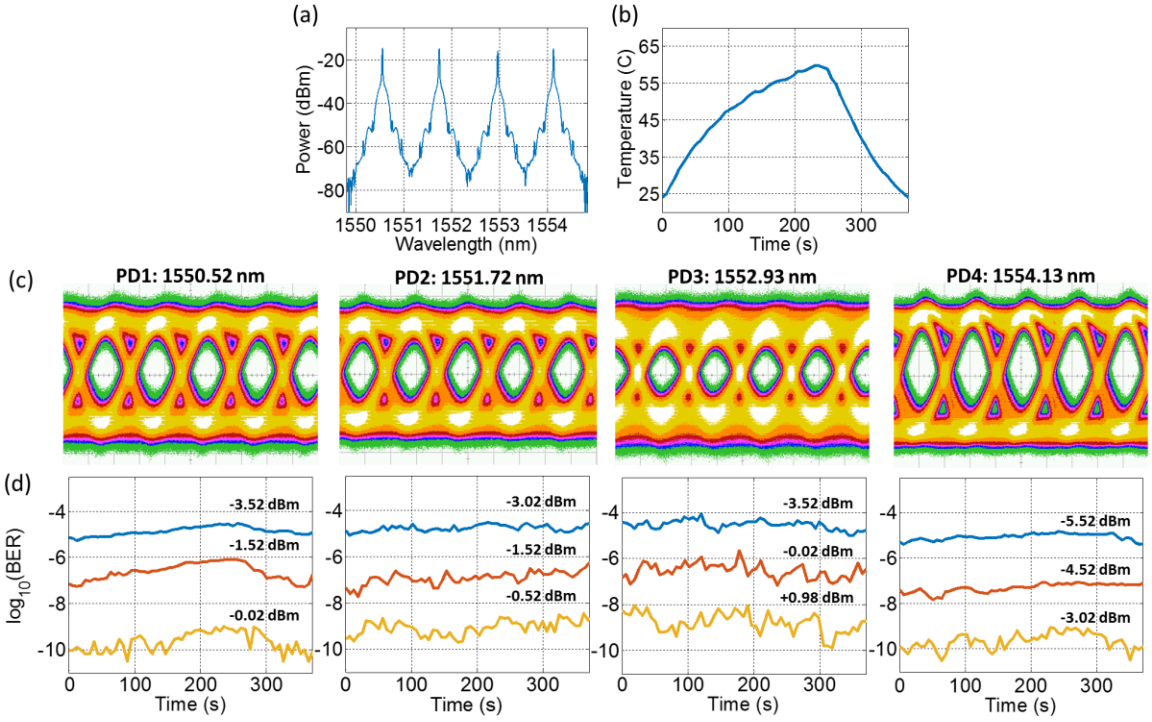

Figure 4.15 (a) Modulated optical spectrum. (b) Temperature cycle during wavelength locking. (c) 25 Gb/s received electrical eye diagrams and (d) BERs from the four high-speed PDs during both temperature cycling and polarization scrambling.

4.7 **Conclusion**

In summary, we have proposed and demonstrated a novel technique for simultaneous wavelength locking of multiple ring modulators to generate WDM signals. The monitoring signal for the closed-loop control is derived from the RF power detection of the modulated components. Maximizing this signal not only enables simultaneously tracking multiple rings to the input laser wavelengths but also allows maximizing the OMA. Furthermore, we have also demonstrated simultaneous wavelength locking and tracking with a polarization insensitive WDM receiver. By using a single MPD at the combined through port of the microring filter array, we can implement a feedback control loop for automatic control of all the microrings, which distinguishes this work from previous reports of single microring filter locking techniques. This novel approach enables low-power consumption using a single monitoring signal for multiple channels and has no additional power requirement to amplify this signal. Moreover, we synchronized the tuning between two polarization branches. Opposite microring filters from two branches may mix two wavelength channels, but synchronizing the tuning avoids this pitfall.

The number of microrings that can be simultaneously locked in our transmitter mainly depends on the bit number of the ADC that reads the output of the RF power detector. In our receiver, the number of bits in the ADC and the chosen feedback resistor can limit the detection sensitivity of the feedback signal. The number of bits in the DACs limits the resolution of the voltage driving the microring heaters. As heating power is the square of the DAC output voltage, this is even more of a concern for heaters tuning at a high voltage and therefore potential instability in

the transmitted and received bit streams. Assuming a 10-bit ADC's dynamic range is fully utilized, and each ring requires a 10-dB dynamic range for locking, the number of microrings in a WDM transmitter array that can be locked becomes about 100. For devices like our WDM receiver, the full dynamic range of a microring filter can be detected under 7 bits of resolution. If we utilize most of the ADC's 12-bit dynamic range, we estimate that we can lock as many as 36 microring filters (or 18 wavelength channels). However, these can be easily improved with higher-bit ADCs and DACs, and in our experiments, the current bit number is enough.

As the thermal tuning speed of silicon photonic devices is typically on the order of 10 μs [60], the tracking speed would be mainly limited by DACs, RF power detector, and ADCs used in the experiment. The electrical bandwidth of our implemented feedback loop – which includes our encoders, the microcontroller, and all microring heaters – is estimated to be on the order of 120 kHz and tracking is maintained with thermal fluctuations as high as 0.64 °C/second. Performance can still be improved with faster control circuitry and more efficient tracking algorithms. The additional RF power detector in our transmitter, compared with low-power CMOS control circuit in [84], can be built on CMOS as well, which costs very little power. In addition, an analog proportional–integral–derivative controller (PID) [85] or a comparator circuit [84] can be used for closed-loop control, instead of a microcontroller used here.

Chapter 5: Conclusions & Future Work

5.1 Summary of this work

Silicon photonics proves to be viable for large scale development in various practical applications. Silicon as a semiconductor material is abundant, well-studied, and compatible with many, if not all, foundry processes. It has already found a competitive edge in long-haul and short-reach telecommunications with continued research and development of silicon photonic transceivers in both academia and industry. Silicon photonics is an attractive technology for scaling up, but it has its own set of challenges to overcome.

In this dissertation, we covered various problems and impactful avenues of science and engineering to continue scaling up. First, we altered our perception of simplified waveguide loss models to consider intermodal-coupling induced loss in wider waveguides. Instead of considering just the loss from the injected mode, we claimed that it is useful to perceive the loss within a range of the losses of all guided modes. From an engineering standpoint, by considering the worst-case loss, we can rightfully engineer the width of our waveguides to reduce the overall loss regardless of the injected mode. In a similar manner, this perception can provide an avenue to optimize for space division multiplexed networks. Ultimately, waveguide loss reduction will undoubtedly reduce the power budget and performance of silicon photonic systems. Second, we proposed and demonstrated a waveguide superlattice bend – a bending dense waveguide array with low crosstalk and small footprint. For scaling up, footprint becomes ever more valuable and the method of waveguide

packing can be a chokepoint, Waveguide bends tend to be particularly large, especially when there is a need to spread out the bends to reduce crosstalk; however, through a waveguide superlattice design approach, we can maintain compactness and optical signal fidelity. For large scale subsystems requiring the flexibility of bending waveguides – such as optical interposers – this is a necessity. Third, we provided methods of automated tuning for highly-sensitive microring networks. Microring-based devices are attractive as they have a small footprint and various functions, including filtering and modulation. The tradeoff of such scalable features is in the sensitivity of microrings to fabrication and ambient conditions. Our method takes advantage of electronic subsystems that provide closed-loop control of microring-based transceivers, ensuring scalability of the method itself to more microrings, well-written control algorithms that optimize for high tuning speed and low electrical power, and endless tracking of microrings allowing quick adaptation to changing environments.

5.2 **Suggestions for future work**

5.2.1 Understanding phase error high-index contrast waveguides

Phase changes induced by rough waveguides is particularly concerning in applications involving coherent communication and OPAs. The exact phase changes caused by fabrication deviations is difficult to predict, but we can estimate the phase error [90]. The variance in the phase change of a waveguide of waveguide width w and length L is defined as

$$\sigma_{\delta\phi}^2 \cong \begin{cases} L^2 \cdot \left(\left\{\frac{\partial\beta}{\partial w}\right\}^2 \sigma^2 + \left\{\frac{\partial\beta}{\partial n_1}\right\}^2 \sigma_{\delta n_1}^2 + \left\{\frac{\partial\beta}{\partial n_2}\right\}^2 \sigma_{\delta n_2}^2\right) + \beta^2(\delta L), & L < L_C \\ L \cdot L_C \cdot \left(\left\{\frac{\partial\beta}{\partial w}\right\}^2 \sigma^2 + \left\{\frac{\partial\beta}{\partial n_1}\right\}^2 \sigma_{\delta n_1}^2 + \left\{\frac{\partial\beta}{\partial n_2}\right\}^2 \sigma_{\delta n_2}^2\right) + \beta^2(\delta L), & L \geq L_C \end{cases} \quad (5.1)$$

where σ, $\sigma_{\delta n1}$, and $\sigma_{\delta n2}$ correspond to the standard deviations of width, core refractive index, and cladding refractive index. δL describes the small deviation from the designed waveguide length. We are only particularly interested in the sidewall roughness. If the refractive indices of the core and cladding are constant, and the length of the waveguide is constant, the phase error can be simplified as

$$\sigma_{\delta\phi}^2 \cong \begin{cases} L^2 \cdot \left\{\frac{\partial\beta}{\partial w}\right\}^2 \sigma^2, & L < L_C \\ L \cdot L_C \cdot \left\{\frac{\partial\beta}{\partial w}\right\}^2 \sigma^2, & L \geq L_C \end{cases} \quad (5.2)$$

Figure 5.1 depicts a surface plot of the phase error variance from a waveguide with varying RMS roughness and correlation lengths at λ_0 = 1550 nm. Such as in the case for estimating propagation loss, the theory for phase error estimation is based on a simplified one-dimensional model. Therefore, to apply it for a two-dimensional channel waveguide, the effective index method is used.

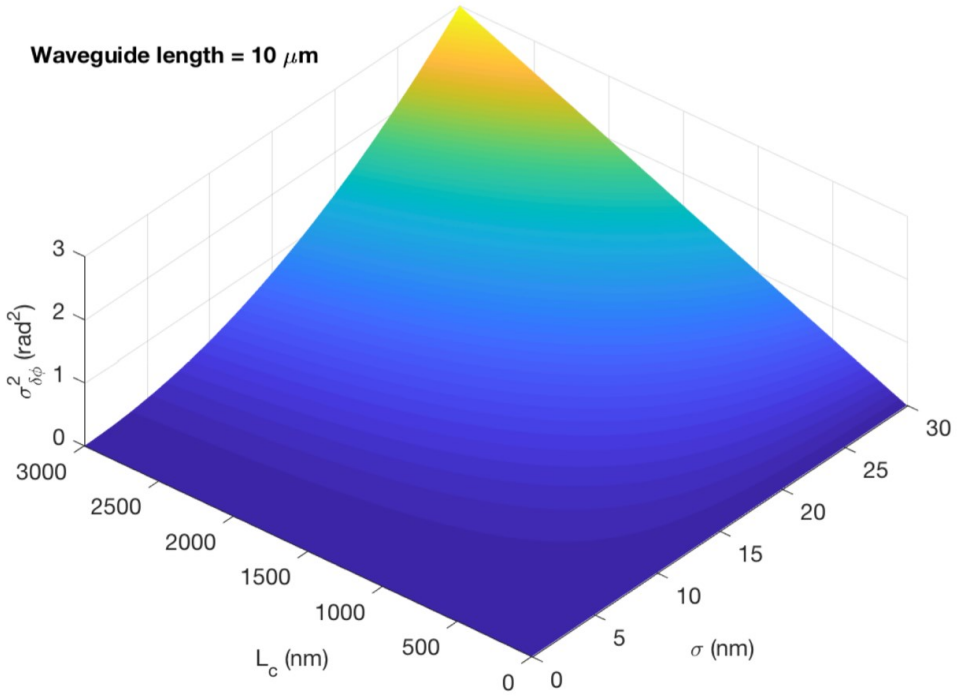

Figure 5.1 Phase error variance of a 260 nm × 420 nm (height × width) silicon waveguide at various sidewall roughness parameters. The waveguide is cladded with silica.

We have yet to verify this theory in experiment, and it should be also noted that not many previous works have extensively verified it and have indirectly shown ways to lower phase errors in waveguides. We also expect – as with the waveguide loss – intermodal-coupling between guided modes may result in larger phase errors than expected due to the assumption that only the fundamental mode of the waveguide is excited.

From an experimental perspective, phase information is lost when light is measured directly with a photodetector. However, we can acquire the phase through interferometric methods. One method in doing so is with a Mach-Zehnder interferometer. One of the arms would have the DUT while the other is a reference. The DUT would induce some phase delay, resulting in a reduced intensity at the

combined output. An on-chip ohmic heater would equalize the phase between the two arms, and its heating power would be proportional to the phase delay induced by the DUT. This method was demonstrated in [90] to determine the phase error in silica waveguides. Another method also uses Mach-Zehnder interferometry, but in this method, we use multi-port intensity measurements to calculate the phase. Specifically, we use the input, the rough waveguide arm, the reference arm, and the combined (interference) arm intensity measurements. **Figure 5.2** depicts a schematic of such a setup.

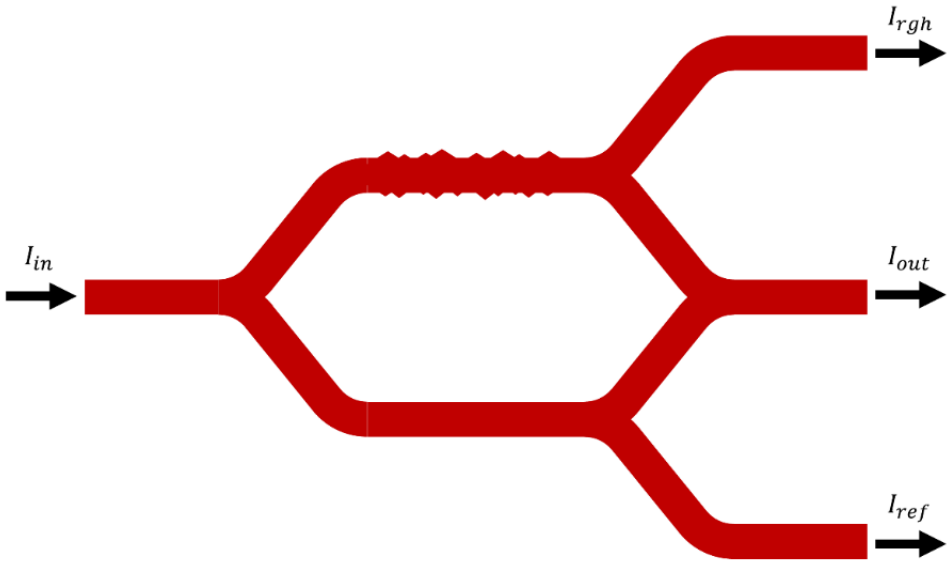

Figure 5.2 Schematic of a Mach-Zehnder interferometer to determine phase error.

Intensity measurements can be attained through the use a photodetector using a similar measurement setup from **Figure 5.2**, where the detected power is proportional to the intensity. We can calculate the phase error from the reference arm using the two-beam interference equation:

$$\delta\phi = cos^{-1}\left(\frac{I_{rgh} - I_{out} - I_{ref}}{2\sqrt{I_{in}I_{ref}}}\right). \quad (5.3)$$

An even more advanced method to acquire phase is through swept wavelength interferometry, which uses interferograms produced via wide bandwidth optical measurements. Optical vector analyzers (OVAs) such as Luna Technologies' OVA 5000 can make measurements in less than three seconds, providing comprehensive optical component characterizations by utilizing SWI.

5.2.2 Exploring other designs of waveguide superlattice bends

Waveguide bends can exist in many different forms other than the typical radial waveguide bend. Two commonly used alternative designs are bends based on Bézier curves [91] and Euler bends [92], [93]. Both types of bends are adiabatic, whose designs smooth out the transition of the waveguide mode as it propagates through the bend to prevent inter-modal crosstalk for multimode bends, reduce insertion loss caused by coupling to radiative modes, and minimize back reflections. In such waveguide bends, the radius of the bend changes depending on what point of the bend the light is in. Bézier curves can be defined by a set of control points $\boldsymbol{P_0}, \boldsymbol{P_1}$, …, $\boldsymbol{P_n}$ for an nth-order curve with the mathematical expression

$$\boldsymbol{B}(t) = \sum_{i=0}^{n} \binom{n}{i} (1-t)^{n-1} t^i \boldsymbol{P_i}, \quad (5.4)$$

where t is between 0 and 1 and $\boldsymbol{B}$ corresponds to the rectangular coordinates of the interpolated curve. For Euler bends, the curvature of the bend changes linearly with the bend length s such that it follows the relation

$$\frac{1}{R} = \frac{d\theta}{ds} \propto s \propto \sqrt{\theta}, \tag{5.5}$$

where θ is the desired bending angle and R is the bending radius. In [92], they describe a modified Euler bend formula which allows for low-loss and low-crosstalk multimode bends. In all types of modified curvature geometries, it may be worth exploring how the constant change in bending physics would hopefully allow for more compact waveguide superlattice bend designs.

5.2.3 Closed-loop control of dense silicon photonic optical phased arrays

We demonstrated automatic closed-loop tuning that could be applied for scalable silicon photonics with microring modulators and arrays, and others have demonstrated closed-loop control in other systems as well [69], [86], [89], [94]–[98]. With regards to silicon photonic AWGs, automatic tuning has been demonstrated via phase retrieval with swept wavelength interferometry (SWI) [99]. While effective, this approach requires a complex phase retrieval method typically demonstrated with fiber optics. To demonstrate this with integrated optics could prove to be overly-complex, especially since the SWI optics would also need to be tuned and calibrated. We propose two approaches as steps toward simplified AWG/OPA tuning and control: alternating projection and particle swarm optimization.

Unfortunately, phase retrieval is not so trivial. Photodetectors are essentially photon counters, which in a manner of speaking, implies that the phase information from a propagating wave is lost. Clearly, there are indeed methods to retrieve the phase – some more complex than others (e.g. SWI). An alternative set of methods for phase retrieval are known as alternating projection methods. One of the earliest, if not the earliest, implementations is the Gerchberg-Saxton (GS) algorithm [100]. Mathematically, the algorithm implements the following relations:

$$G_n(u) = \mathcal{F}\{g_n(x)\} = |G_n(u)|e^{i\phi_{G_n}(u)} \tag{5.6}$$

$$G_n'(u) = |F(u)|e^{i\phi_{G_n}(u)} \tag{5.7}$$

$$g_n'(x) = \mathcal{F}^{-1}\{G_n'(u)\} = |g_n'(x)|e^{i\phi_{g_n'}(x)} \tag{5.8}$$

$$g_{n+1}(x) = |f(x)|e^{i\phi_{g_n'}(x)} \tag{5.9}$$

Requiring only the intensity information from the image plane $|g(x)|^2$ and the far-field diffraction plane $|G(x)|^2$, we can potentially retrieve the phase information from the image plane. This is possible due to the Fourier relation in far-field diffraction optics. This is also known as the Fraunhofer diffraction regime:

$$\frac{W_s^2}{L_f\lambda} \ll 1, \tag{5.10}$$

where W_s is the width of the slit (or aperture), L_f is the focal length, and λ is the wavelength. By alternating between the near and far-field projections, the square error between actual and estimated far-field projections $M^{-2}[\Sigma_u|F(u) - G_n(u)|]$ will be minimized, and the estimated image plane information $\varphi_{g'}(x)$ will be retrieved. **Figure**

5.3 depicts a flow chart of the algorithm. The loop stops when a specified square error is reached or when the change is nearly negligent.

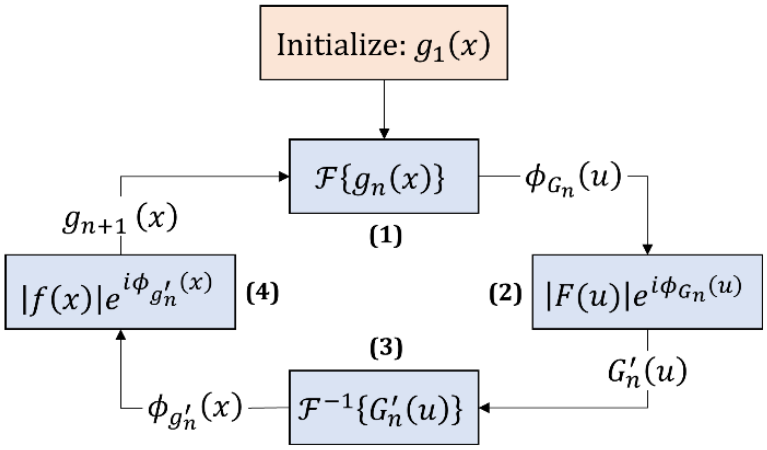

Figure 5.3 Gerchberg-Saxton algorithm for phase retrieval.

Initialization of the GS algorithm is typically the amplitude pattern of the expected input field with uniformly distributed random phases for each input aperture, i.e.,

$$g_1(x) = |f(x)|e^{i\cdot rand[\phi(x)]} \tag{5.11}$$

The GS algorithm, however, can unfortunately settle into a local minimum error, resulting in incorrect phase information. There have been various additional alternating projection methods to avoid this local minimum which modify steps in the baseline GS algorithm [101]–[107]. A popular one, known as the hybrid input-output (HIO) algorithm replaces the last step of a GS iteration with

$$g_{n+1}(x) = \begin{cases} g_n'(x), & x \notin \gamma \\ g_n(x) - \beta_f g_n'(x), & x \in \gamma \end{cases} \tag{5.12}$$

where β_f is a feedback constant (typically between 0.5 and 1.0) and γ s the set of points that are not within the image plane's constraints. Compared to the baseline GS algorithm, the HIO algorithm performs reasonably better and, in some cases, a combined GS/HIO algorithm can push the performance further [107]. Unfortunately, under certain conditions, these algorithms are still susceptible converging towards a local minimum. Another interesting approach to avoid local minima is setting the initial phase guess more deterministically as opposed to completely at random. We propose that if we are already around the neighborhood of the correct phases – the global minimum – then we potentially can decrease the chances of reaching a local minimum. Also, the phase retrieval algorithm may work quicker if the initial guess is within the neighborhood, thus saving precious computation time for the control circuitry.

If the computational power of the control circuitry is constrained, especially for low-power applications, then the Fourier transforms used in alternating projection methods may be too slow. Alternating projection thus is impractical. Instead, we look towards other promising methods such as particle swarm optimization (PSO). In summary, PSO is an evolutionary algorithm inspired by behavior of a flock of birds converging to a point. Each particle in the search space adjusts its trajectory in subsequent iterations using the information from itself and those of neighboring particles. While many variations of PSO have existed, the latest version of the Standard PSO (SPSO-2011) was formulated to provide straightforward adaptation of PSO with scalable behavior, demonstrating consistent performance up to 50-dimensional problems [108].

For a D-dimensional search space, the position and velocity of the i^{th} particle are defined as $\boldsymbol{X}_i = x_{i1}, x_{i2}, \ldots, x_{iD}$ and $\boldsymbol{V}_i = v_{i1}, v_{i2}, \ldots, v_{iD}$ respectively. Each particle's search performance is defined by an application-specific figure of merit, which is then used to update $\boldsymbol{X}_i$. The best-known position of the i^{th} particle – its personal best – is $\boldsymbol{P}_i = p_{i1}, p_{i2}, \ldots, p_{iD}$, and the best-known position within that particle's neighborhood – its personal best – is $\boldsymbol{L} = l_1, l_2, \ldots, l_D$. SPSO-2011 exploits the idea of rotational invariance to update each particle. At each iteration, a center of gravity $\boldsymbol{G}_i$ is defined by three points: current position $\boldsymbol{X}_i^t$, a point slightly beyond the previous best personal position $\boldsymbol{p}_i^t$, and a point slightly beyond the previous best neighborhood position $\boldsymbol{l}_i^t$. Mathematically, this is

$$\boldsymbol{p}_i^t = \boldsymbol{X}_i^t + c_1 \boldsymbol{U}_1^t \otimes (\boldsymbol{P}_i^t - \boldsymbol{X}_i^t) \tag{5.13}$$

$$\boldsymbol{l}_i^t = \boldsymbol{X}_i^t + c_2 \boldsymbol{U}_2^t \otimes (\boldsymbol{L}^t - \boldsymbol{X}_i^t) \tag{5.14}$$

$$\boldsymbol{G}_i^t = \frac{(\boldsymbol{X}_i^t + \boldsymbol{p}_i^t + \boldsymbol{l}_i^t)}{3} \tag{5.15}$$

where c_1 and c_2 are cognitive and social acceleration coefficients, both usually equal to 0.5 + ln(2), $\boldsymbol{U}_1^t$ and $\boldsymbol{U}_2^t$ are independent uniformly distributed random vectors within the range [0, 1]. Each i^{th} particle is updated for each t^{th} iteration as

$$\boldsymbol{V}_i^{t+1} = \omega \boldsymbol{V}_i^t + \mathcal{H}_i(\boldsymbol{G}_i^t, \|\boldsymbol{G}_i^t - X_i^t\|) - \boldsymbol{X}_i^t \tag{5.16}$$

$$\boldsymbol{X}_i^{t+1} = \boldsymbol{X}_i^t + \boldsymbol{V}_i^{t+1} \tag{5.17}$$

Where ω is the inertia weight, which is typically $\omega = 1/[2\cdot\ln(2)]$. The function H_i returns a random point in a hypersphere (uniform distribution) with its center defined at $\boldsymbol{G}_i^t$ and radius defined as $||\boldsymbol{G}_i^t - \boldsymbol{X}_i^t||$. Particle swarm size is user-defined in SPSO-2011, and the initial conditions are initialized as

$$X_i^0 = U(min_d, max_d) \tag{5.18}$$

$$V_i^0 = \frac{U(min_d, max_d) - X_i^0}{2} \tag{5.19}$$

$$p_i^0 = X_i^0 \tag{5.20}$$

$$l_i^0 = min[f(p_i^0)] \tag{5.21}$$

$U(min_d, max_d)$ is a uniformly distributed random number in the range $[min_d, max_d]$. These boundaries are also user-defined. In the algorithm, when a particle reaches these boundaries, its velocity is reset to zero. SPSO-2011 runs until some user-defined tolerance error is met or after a certain number of iterations.

PSO has recently gained much traction as a tool to optimize integrated device design. For our purposes, we expect PSO to be beneficial in closed-loop control systems with programmable control circuitry such as microcontrollers or field programmable gate arrays (FPGAs). This is particularly interesting in actively controlling the phase output of the waveguides in OPAs and AWGs. Because SPSO-2011 is scalable for high-dimensional problems, it can potentially handle the sheer number of phase-controlled waveguides in these devices. In addition to this advantage, as we proposed for improving the performance of the GS algorithm, we

can hopefully improve the performance of SPSO-2011 if we constrain the search space within the neighborhood of the desired phase control.

Bibliography

[1] P. J. Winzer, "Making spatial multiplexing a reality," *Nat. Photonics*, vol. 8, no. 5, pp. 345–348, Apr. 2014.

[2] D. J. Richardson, J. M. Fini, and L. E. Nelson, "Space-division multiplexing in optical fibres," *Nat. Photonics*, vol. 7, no. 5, pp. 354–362, Apr. 2013.

[3] C. T. DeRose *et al.*, "Electronically controlled optical beam-steering by an active phased array of metallic nanoantennas," *Opt. Express*, vol. 21, no. 4, p. 5198, Feb. 2013.

[4] R. T. Chen, "Single-Stage Large-Angle Beam Steering Optical Phased Array on Silicon Nanomembrane APPROVED BY SUPERVISING COMMITTEE : Single-Stage Large-Angle Beam Steering Optical Phased Array on Silicon Nanomembrane by David Nien Kwong , BS Master of Science in Engin."

[5] C. V Poulton *et al.*, "Large-scale silicon nitride nanophotonic phased arrays at infrared and visible wavelengths," *Opt. Lett.*, vol. 42, no. 1, p. 21, Jan. 2017.

[6] S. A. Miller *et al.*, "512-Element Actively Steered Silicon Phased Array for Low-Power LIDAR," in *Conference on Lasers and Electro-Optics*, 2018, p. JTh5C.2.

[7] Y. Shen *et al.*, "Deep learning with coherent nanophotonic circuits," *Nat. Photonics*, vol. 11, no. 7, pp. 441–446, Jun. 2017.

[8] T. Rudolph, "Why I am optimistic about the silicon-photonic route to quantum computing," *APL Photonics*, vol. 2, no. 3, p. 030901, Mar. 2017.

[9] M. Gimeno-segovia, "Towards Practical Linear Optical Quantum Computing," 2015.

[10] R. Soref, "Silicon-based optoelectronics," *Proc. IEEE*, vol. 81, no. 12, 1993.

[11] P. Dong, "Silicon Photonic Integrated Circuits for Wavelength-Division Multiplexing Applications," *IEEE J. Sel. Top. Quantum Electron.*, vol. 22, no. 6, pp. 370–378, Nov. 2016.

[12] D. Y. Lee *et al.*, "Error-free operation of a polarization-insensitive 4λ x 25 Gbps silicon photonic WDM receiver with closed-loop thermal stabilization of Si microrings," *Opt. Express*, vol. 24, no. 12, p. 13204, Jun. 2016.

[13] K. Okamoto, "Wavelength-Division-Multiplexing Devices in Thin SOI: Advances and Prospects," *IEEE J. Sel. Top. Quantum Electron.*, vol. 20, no. 4, pp. 248–257, 2014.

[14] W. Bogaerts *et al.*, "Silicon-on-Insulator Spectral Filters Fabricated With CMOS Technology," *IEEE J. Sel. Top. Quantum Electron.*, vol. 16, no. 1, pp. 33–44, 2010.

[15] L.-W. Luo *et al.*, "WDM-compatible mode-division multiplexing on a silicon chip," *Nat. Commun.*, vol. 5, p. 3069, Jan. 2014.

[16] D. Dai, "Multimode optical waveguide enabling microbends with low inter-mode crosstalk for mode-multiplexed optical interconnects.," *Opt. Express*, vol. 22, no. 22, pp. 27524–34, 2014.

[17] L. H. Gabrielli, D. Liu, S. G. Johnson, and M. Lipson, "On-chip transformation optics for multimode waveguide bends," *Nat. Commun.*, vol. 3, p. 1217, 2012.

[18] P.-K. Shen, X. Xu, A. Hosseini, Z. Pan, and R. T. Chen, "Multiple-input multiple-

output based high density on-chip optical interconnect," in *Optical Interconnects XV*, 2015, vol. 9368, p. 936812.

[19] J. Wang, S. Chen, and D. Dai, "Silicon hybrid demultiplexer with 64 channels for wavelength / mode-division multiplexed on-chip optical interconnects," *Opt. Lett.*, vol. 39, no. 24, pp. 6993–6996, 2014.

[20] L. Liu, "Densely packed waveguide array (DPWA) on a silicon chip for mode division multiplexing," *Opt. Express*, vol. 23, no. 9, p. 12135, 2015.

[21] N. K. Fontaine, "Optical MIMO Demultiplexing and Space-Division Multiplexing using a Photonic Integrated Circuit," in *Advanced Photonics 2013*, 2013, p. SPT4D.1.

[22] C. R. Doerr, N. Fontaine, M. Hirano, T. Sasaki, L. Buhl, and P. Winzer, "Silicon photonic integrated circuit for coupling to a ring-core multimode fiber for space-division multiplexing," *37th Eur. Conf. Expo. Opt. Commun.*, p. Th.13.A.3, 2011.

[23] Y. Ding, H. Ou, J. Xu, and C. Peucheret, "Silicon Photonic Integrated Circuit Mode Multiplexer," *IEEE Photonics Technol. Lett.*, vol. 25, no. 7, pp. 648–651, Apr. 2013.

[24] M. Raval, C. V Poulton, and M. R. Watts, "Unidirectional waveguide grating antennas with uniform emission for optical phased arrays," *Opt. Lett.*, vol. 42, no. 13, p. 2563, Jul. 2017.

[25] K. Van Acoleyen, K. Komorowska, W. Bogaerts, and R. Baets, "One-Dimensional Off-Chip Beam Steering and Shaping Using Optical Phased Arrays on Silicon-on-Insulator," *J. Light. Technol.*, vol. 29, no. 23, pp. 3500–3505, Dec. 2011.

[26] M. J. R. Heck, "Highly integrated optical phased arrays: photonic integrated circuits for optical beam shaping and beam steering," *Nanophotonics*, vol. 6, no. 1, pp. 93–107, Jan. 2017.

[27] A. Yaacobi, J. Sun, M. Moresco, G. Leake, D. Coolbaugh, and M. R. Watts, "Integrated phased array for wide-angle beam steering," *Opt. Lett.*, vol. 39, no. 15, p. 4575, Aug. 2014.

[28] H. Abediasl and H. Hashemi, "Monolithic optical phased-array transceiver in a standard SOI CMOS process," *Opt. Express*, vol. 23, no. 5, p. 6509, Mar. 2015.

[29] J. Notaros, C. V. Poulton, M. J. Byrd, M. Raval, and M. R. Watts, "Integrated optical phased arrays for quasi-Bessel-beam generation," *Opt. Lett.*, vol. 42, no. 17, p. 3510, Sep. 2017.

[30] J. Sun, E. Timurdogan, A. Yaacobi, E. S. Hosseini, and M. R. Watts, "Large-scale nanophotonic phased array.," *Nature*, vol. 493, no. 7431, pp. 195–9, Jan. 2013.

[31] M. Hochberg and T. Baehr-Jones, "Towards fabless silicon photonics," *Nat. Photonics*, vol. 4, no. 8, pp. 492–494, Aug. 2010.

[32] Y. Vlasov and S. McNab, "Losses in single-mode silicon-on-insulator strip waveguides and bends.," *Opt. Express*, vol. 12, no. 8, pp. 1622–1631, 2004.

[33] D. Marcuse, "Mode Conversion Caused by Surface Imperfections of a Dielectric Slab Waveguide," *Bell Syst. Tech. J.*, vol. 48, no. 10, pp. 3187–3215, Dec. 1969.

[34] J. P. R. Lacey and F. P. Payne, "Radiation loss from planar waveguides with random wall imperfections," *IEE Proc. J Optoelectron.*, vol. 137, no. 4, p. 282, 1990.

[35] F. P. Payne and J. P. R. Lacey, “A theoretical analysis of scattering loss from planar optical waveguides,” *Opt. Quantum Electron.*, vol. 26, pp. 977–986, 1994.
[36] F. Ladouceur and L. Poladian, “Surface roughness and backscattering,” *Opt. Lett.*, vol. 21, no. 22, p. 1833, Nov. 1996.
[37] T. Barwicz and H. a. Haus, “Three-dimensional analysis of scattering losses due to sidewall roughness in microphotonic waveguides,” *J. Light. Technol.*, vol. 23, no. 9, pp. 2719–2732, 2005.
[38] C. G. Poulton *et al.*, “Radiation modes and roughness loss in high index-contrast waveguides,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 12, no. 6, pp. 1306–1320, 2006.
[39] F. Grillot, L. Vivien, S. Laval, and E. Cassan, “Propagation loss in single-mode ultrasmall square silicon-on-insulator optical waveguides,” *J. Light. Technol.*, vol. 24, no. 2, pp. 891–895, 2006.
[40] D. Melati, F. Morichetti, and a Melloni, “A unified approach for radiative losses and backscattering in optical waveguides,” *J. Opt.*, vol. 16, no. 5, p. 055502, 2014.
[41] S. Zhu, G. Q. Lo, and D. L. Kwong, “Low-loss amorphous silicon wire waveguide for integrated photonics: effect of fabrication process and the thermal stability.,” *Opt. Express*, vol. 18, no. 24, pp. 25283–91, Nov. 2010.
[42] K. K. Lee, D. R. Lim, L. C. Kimerling, J. Shin, and F. Cerrina, “Fabrication of ultralow-loss Si/SiO(2) waveguides by roughness reduction.,” *Opt. Lett.*, vol. 26, no. 23, pp. 1888–1890, 2001.
[43] D. K. Sparacin, S. J. Spector, and L. C. Kimerling, “Silicon waveguide sidewall smoothing by wet chemical oxidation,” *J. Light. Technol.*, vol. 23, no. 8, pp. 2455–2461, 2005.
[44] R. Gatdula, S. Abbaslou, M. Lu, A. Stein, and W. Jiang, “Evidence of Intermodal-Coupling-Induced Loss in Rough Silicon Waveguides,” Manuscript in Preparation.
[45] R. G. G. Beausoleil *et al.*, “A Nanophotonic Interconnect for High-Performance Many-Core Computation,” in *2008 16th IEEE Symposium on High Performance Interconnects*, 2008, pp. 182–189.
[46] K. Chen, S. Wang, S. Chen, C. Wang, ShipengZhang, D. Dai, and L. Liu, “Experimental demonstration of simultaneous mode and polarization-division multiplexing based on silicon densely packed waveguide array,” *Opt. Lett.*, vol. 40, no. 20, pp. 4655–4658, 2015.
[47] K. Murray, Z. Lu, H. Jayatilleka, and L. Chrostowski, “Dense dissimilar waveguide routing for highly efficient thermo-optic switches on silicon,” *Opt. Express*, vol. 23, no. 15, p. 19575, 2015.
[48] W. Song *et al.*, “High-density waveguide superlattices with low crosstalk,” *Nat. Commun.*, vol. 6, no. May, p. 7027, May 2015.
[49] M. Mrejen *et al.*, “Adiabatic elimination-based coupling control in densely packed subwavelength waveguides.,” *Nat. Commun.*, vol. 6, no. May, p. 7565, 2015.
[50] V. Donzella, S. Talebi Fard, and L. Chrostowski, “Study of waveguide crosstalk in silicon photonics integrated circuits,” in *Photonics North 2013*, 2013, vol.

8915, pp. 89150Z-89150Z–8.
[51] W.-P. Huang, “Coupled-mode theory for optical waveguides: an overview,” *J. Opt. Soc. Am. A*, vol. 11, no. 3, p. 963, Mar. 1994.
[52] A. W. Snyder and J. Love, *Optical Waveguide Theory*. London, England: Chapman & Hall, 1983.
[53] B. E. A. Saleh and M. C. Teich, *Fundamentals of Photonics*, 2nd ed. Wiley-Interscience, 2007.
[54] Y. Murakami, “Coupling between curved dielectric waveguides.,” *Appl. Opt.*, vol. 19, no. 3, pp. 398–403, Feb. 1980.
[55] S. J. Garth, “Modes on a bent optical waveguide,” *IEE Proc. J Optoelectron.*, vol. 134, no. 4, p. 221, 1987.
[56] K. Thyagarajan, M. R. Shenoy, and A. K. Ghatak, “Accurate numerical method for the calculation of bending loss in optical waveguides using a matrix approach,” *Opt. Lett.*, vol. 12, no. 4, p. 296, Apr. 1987.
[57] R. Gatdula, S. Abbaslou, M. Lu, A. Stein, and W. Jiang, “Guiding light in bent waveguide superlattices with low crosstalk,” *Optica*, vol. 6, no. 5, p. 585, May 2019.
[58] A. Arbabi and L. L. Goddard, “Dynamics of Self-Heating in Microring Resonators,” *IEEE Photonics J.*, vol. 4, no. 5, pp. 1702–1711, Oct. 2012.
[59] W. Bogaerts *et al.*, “Silicon microring resonators,” *Laser Photon. Rev.*, vol. 6, no. 1, pp. 47–73, Jan. 2012.
[60] F. Gan *et al.*, “Maximizing the Thermo-Optic Tuning Range of Silicon Photonic Structures,” in *2007 Photonics in Switching*, 2007, pp. 67–68.
[61] P. Dong *et al.*, “Thermally tunable silicon racetrack resonators with ultralow tuning power,” *Opt. Express*, vol. 18, no. 19, p. 20298, Sep. 2010.
[62] P. De Heyn *et al.*, “Polarization-Insensitive 5x20Gb/s WDM Ge Receiver using Compact Si Ring Filters with Collective Thermal Tuning,” in *Optical Fiber Communication Conference*, 2014, p. Th4C.5.
[63] P. Dong, Y.-K. Chen, and L. L. Buhl, “Reconfigurable Four-Channel Polarization Diversity Silicon Photonic WDM Receiver,” in *Optical Fiber Communication Conference*, 2015, p. W3A.2.
[64] H. Jayatilleka *et al.*, “Wavelength tuning and stabilization of microring-based filters using silicon in-resonator photoconductive heaters,” *Opt. Express*, vol. 23, no. 19, p. 25084, Sep. 2015.
[65] Y. Li and A. W. Poon, “Active resonance wavelength stabilization for silicon microring resonators with an in-resonator defect-state-absorption-based photodetector,” *Opt. Express*, vol. 23, no. 1, p. 360, Jan. 2015.
[66] C. V. Poulton, P. Dong, and Y.-K. Chen, “Photoresistive Microring Heater with Resonance Control Loop,” in *CLEO: 2015*, 2015, p. SM2I.3.
[67] S. Grillanda *et al.*, “Non-invasive monitoring and control in silicon photonics using CMOS integrated electronics,” *Optica*, vol. 1, no. 3, p. 129, Sep. 2014.
[68] H. Yu *et al.*, “Using carrier-depletion silicon modulators for optical power monitoring,” *Opt. Lett.*, vol. 37, no. 22, p. 4681, Nov. 2012.
[69] J. C. C. Mak, W. D. Sacher, Tianyuan Xue, J. C. Mikkelsen, Zheng Yong, and J. K. S. Poon, “Automatic Resonance Alignment of High-Order Microring Filters,” *IEEE J. Quantum Electron.*, vol. 51, no. 11, pp. 1–11, Nov. 2015.

[70] J. A. Cox, A. L. Lentine, D. C. Trotter, and A. L. Starbuck, "Control of integrated micro-resonator wavelength via balanced homodyne locking," *Opt. Express*, vol. 22, no. 9, p. 11279, May 2014.
[71] K. Padmaraju, D. F. Logan, T. Shiraishi, J. J. Ackert, A. P. Knights, and K. Bergman, "Wavelength Locking and Thermally Stabilizing Microring Resonators Using Dithering Signals," *J. Light. Technol.*, vol. 32, no. 3, pp. 505–512, Feb. 2014.
[72] F. Morichetti, S. Grillanda, and A. Melloni, "Breakthroughs in Photonics 2013: Toward Feedback-Controlled Integrated Photonics," *IEEE Photonics J.*, vol. 6, no. 2, pp. 1–6, Apr. 2014.
[73] K. Padmaraju and K. Bergman, "Resolving the thermal challenges for silicon microring resonator devices," *Nanophotonics*, vol. 3, no. 4–5, pp. 269–281, Jan. 2014.
[74] Q. Xu, B. Schmidt, J. Shakya, and M. Lipson, "Cascaded silicon micro-ring modulators for WDM optical interconnection," *Opt. Express*, vol. 14, no. 20, p. 9431, 2006.
[75] C.-H. Chen *et al.*, "A comb laser-driven DWDM silicon photonic transmitter based on microring modulators," *Opt. Express*, vol. 23, no. 16, p. 21541, Aug. 2015.
[76] J. Müller *et al.*, "Silicon photonics WDM transmitter with single section semiconductor mode-locked laser," *Adv. Opt. Technol.*, vol. 4, no. 2, Jan. 2015.
[77] P. Dong, J. Lee, K. Kim, Y.-K. Chen, and C. Gui, "Ten-Channel Discrete Multi-Tone Modulation Using Silicon Microring Modulator Array," in *Optical Fiber Communication Conference*, 2016, p. W4J.4.
[78] H. Li *et al.*, "A 25 Gb/s, 4.4 V-Swing, AC-Coupled Ring Modulator-Based WDM Transmitter with Wavelength Stabilization in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3145–3159, Dec. 2015.
[79] C. Li *et al.*, "Silicon Photonic Transceiver Circuits With Microring Resonator Bias-Based Wavelength Stabilization in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 6, pp. 1419–1436, Jun. 2014.
[80] K. Padmaraju, D. F. Logan, X. Zhu, J. J. Ackert, A. P. Knights, and K. Bergman, "Integrated thermal stabilization of a microring modulator," *Opt. Express*, vol. 21, no. 12, p. 14342, Jun. 2013.
[81] S. Lin *et al.*, "Wavelength Locked High-speed Microring Modulator Using an Integrated Balanced Homodyne CMOS Control Circuit," in *Optical Fiber Communication Conference*, 2016, p. Th3J.4.
[82] K. Padmaraju, J. Chan, L. Chen, M. Lipson, and K. Bergman, "Thermal stabilization of a microring modulator using feedback control," *Opt. Express*, vol. 20, no. 27, p. 27999, Dec. 2012.
[83] W. A. Zortman, A. L. Lentine, D. C. Trotter, and M. R. Watts, "Bit-Error-Rate Monitoring for Active Wavelength Control of Resonant Modulators," *IEEE Micro*, vol. 33, no. 1, pp. 42–52, Jan. 2013.
[84] X. Zheng *et al.*, "A high-speed, tunable silicon photonic ring modulator integrated with ultra-efficient active wavelength control," *Opt. Express*, vol. 22, no. 10, p. 12628, May 2014.
[85] A. Melikyan, K. Kim, Y.-K. Chen, and P. Dong, "Tapless Locking of Silicon Ring

Modulators for WDM Applications," in *Optical Fiber Communication Conference*, 2017, p. Tu2H.6.
[86] P. P. Dong *et al.*, "Simultaneous wavelength locking of microring modulator array with a single monitoring signal," *Opt. Express*, vol. 25, no. 14, p. 16040, Jul. 2017.
[87] X. Wu, C. Huang, K. Xu, C. Shu, and H. K. Tsang, "128-Gb/s Line Rate OFDM Signal Modulation Using an Integrated Silicon Microring Modulator," *IEEE Photonics Technol. Lett.*, vol. 28, no. 19, pp. 2058–2061, Oct. 2016.
[88] D. Miller, "Device Requirements for Optical Interconnects to Silicon Chips," *Proc. IEEE*, vol. 97, no. 7, pp. 1166–1185, Jul. 2009.
[89] R. Gatdula *et al.*, "Simultaneous four-channel thermal adaptation of polarization insensitive silicon photonics WDM receiver," *Opt. Express*, vol. 25, no. 22, p. 27119, Oct. 2017.
[90] T. Goh, S. Suzuki, and A. Sugita, "Estimation of waveguide phase error in silica-based waveguides," *J. Light. Technol.*, vol. 15, no. 11, pp. 2107–2113, 1997.
[91] L. Chrostowski and M. Hochberg, *Silicon Photonics Design: From Devices to Systems*. 2015.
[92] X. Jiang, H. Wu, and D. Dai, "Low-loss and low-crosstalk multimode waveguide bend on silicon," *Opt. Express*, vol. 26, no. 13, p. 17680, Jun. 2018.
[93] M. Cherchi, S. Ylinen, M. Harjanne, M. Kapulainen, and T. Aalto, "Dramatic size reduction of waveguide bends on a micron-scale silicon photonic platform," *Opt. Express*, vol. 21, no. 15, p. 17814, Jul. 2013.
[94] X. Zhu *et al.*, "Fast Wavelength Locking of a Microring Resonator," vol. 4, pp. 11–12, 2014.
[95] M. Ma *et al.*, "Silicon Photonic Polarization Receiver with Automated Stabilization for Arbitrary Input Polarizations," in *Conference on Lasers and Electro-Optics*, 2016, p. STu4G.8.
[96] W. Bogaerts, M. Fiers, and P. Dumon, "Design Challenges in Silicon Photonics," *IEEE J. Sel. Top. Quantum Electron.*, vol. 20, no. 4, pp. 1–8, 2014.
[97] C. R. Doerr and L. Chen, "Monolithic PDM-DQPSK receiver in silicon," in *36th European Conference and Exhibition on Optical Communication*, 2010, vol. 1–2, pp. 1–3.
[98] Z. Wang and C. Xie, "Automatic optical polarization demultiplexing for polarization division multiplexed signals," *Opt. Express*, vol. 17, no. 5, p. 3183, Mar. 2009.
[99] M. Gehl, D. Trotter, A. Starbuck, A. Pomerene, A. L. Lentine, and C. DeRose, "Active phase correction of high resolution silicon photonic arrayed waveguide gratings," *Opt. Express*, vol. 25, no. 6, p. 6320, Mar. 2017.
[100] R. W. Gerchberg and W. O. Saxton, "A practical algorithm for the determination of phase from image and diffraction plane pictures," *Optik (Stuttg).*, vol. 35, no. 2, pp. 237–246, 1972.
[101] M. Guizar-Sicairos and J. R. Fienup, "Understanding the twin-image problem in phase retrieval," *J. Opt. Soc. Am. A*, vol. 29, no. 11, p. 2367, 2012.
[102] J. R. Fienup, "Phase retrieval algorithms: a personal tour [Invited].," *Appl. Opt.*, vol. 52, no. 1, pp. 45–56, 2013.
[103] Y. Shechtman, Y. C. Eldar, O. Cohen, H. N. Chapman, J. Miao, and M. Segev,

"Phase Retrieval with Application to Optical Imaging," *Signal Process. Mag.*, no. April 2015, pp. 1–25, 2015.
[104] J. a. Rodriguez, R. Xu, C. C. Chen, Y. Zou, and J. Miao, "Oversampling smoothness: An effective algorithm for phase retrieval of noisy diffraction intensities," *J. Appl. Crystallogr.*, vol. 46, no. 2, pp. 312–318, 2013.
[105] J. Fienup and C. C. Wackerman, "Phase retrieval stagnation problems and solutions," *J. Opt. Soc. Am. A*, vol. 3, no. 11, pp. 1897–1907, 1986.
[106] J. R. Fienup, "Phase retrieval algorithms: a comparison.," *Appl. Opt.*, vol. 21, no. 15, pp. 2758–2769, 1982.
[107] C. Guo, S. Liu, and J. T. Sheridan, "Iterative phase retrieval algorithms I: optimization," *Appl. Opt.*, vol. 54, no. 15, p. 4698, 2015.
[108] M. Zambrano-Bigiarini, M. Clerc, and R. Rojas, "Standard Particle Swarm Optimisation 2011 at CEC-2013: A baseline for future PSO improvements," in *2013 IEEE Congress on Evolutionary Computation*, 2013, pp. 2337–2344.
[109] R. C. Jaeger, *Introduction to Microelectronic Fabrication*, 2nd ed. Prentice Hall, 2002.
[110] R. J. Bojko, J. Li, L. He, T. Baehr-Jones, M. Hochberg, and Y. Aida, "Electron beam lithography writing strategies for low loss, high confinement silicon optical waveguides," *J. Vac. Sci. Technol. B Microelectron. Nanom. Struct.*, vol. 29, no. 6, p. 06F309, 2011.
[111] V. R. Manfrinato *et al.*, "Resolution limits of electron-beam lithography toward the atomic scale.," *Nano Lett.*, vol. 13, no. 4, pp. 1555–8, Apr. 2013.
[112] J. K. W. Yang and K. K. Berggren, "Using high-contrast salty development of hydrogen silsesquioxane for sub-10-nm half-pitch lithography," *J. Vac. Sci. Technol. B Microelectron. Nanom. Struct.*, vol. 25, no. 6, p. 2025, 2007.
[113] Y. Chen, H. Yang, and Z. Cui, "Effects of developing conditions on the contrast and sensitivity of hydrogen silsesquioxane," *Microelectron. Eng.*, vol. 83, no. 4–9, pp. 1119–1123, Apr. 2006.
[114] J. Lee, G. Yeom, J. Lee, and J. Lee, "Study of shallow silicon trench etch process using planar inductively coupled plasmas," *J. Vac. Sci. Technol. A*, vol. 15, no. 3, pp. 573–578, 1997.
[115] C. W. Holzwarth, T. Barwicz, and H. I. Smith, "Optimization of hydrogen silsesquioxane for photonic applications," *J. Vac. Sci. Technol. B Microelectron. Nanom. Struct.*, vol. 25, no. 6, p. 2658, 2007.
[116] Y. K. Siew, G. Sarkar, X. Hu, J. Hui, a. See, and C. T. Chua, "Thermal Curing of Hydrogen Silsesquioxane," *J. Electrochem. Soc.*, vol. 147, no. 1, p. 335, 2000.
[117] S. Choi, M. J. Word, V. Kumar, and I. Adesida, "Comparative study of thermally cured and electron-beam-exposed hydrogen silsesquioxane resists," *J. Vac. Sci. Technol. B Microelectron. Nanom. Struct.*, vol. 26, no. 5, p. 1654, 2008.
[118] J. Batey and E. Tierney, "Low-temperature deposition of high-quality silicon dioxide by plasma-enhanced chemical vapor deposition," *J. Appl. Phys.*, vol. 60, no. 9, p. 3136, 1986.
[119] M. F. Ceiler Jr., P. A. Kohl, and S. A. Bidstrup, "Plasma-Enhanced Chemical Vapor Deposition Silicon Dioxide Deposited at Low Temperatures," *Journal of Electrochemical Society*, vol. 142, no. 6. pp. 2067–2071, 1995.
[120] W. Kern, "The Evolution of Silicon Wafer Cleaning Technology," *J. Electrochem.*

Soc., vol. 137, no. 6, p. 1887, 1990.
[121] M. Itano, F. W. Kern, M. Miyashita, and T. Ohmi, "Particle removal from silicon wafer surface in wet cleaning process," *IEEE Trans. Semicond. Manuf.*, vol. 6, no. 3, pp. 258–267, 1993.

Appendix A: Microfabrication Methods

Silicon photonics is an attractive option for on-chip optical communications because it is easily compatible with the high-yield industrial processes already used for fabricating microelectronic chips [109]. In a similar manner, this makes it easy to prototype structures in an academic clean room – given that the clean room can do it in the first place. Since the waveguide superlattice is a passive photonics structure – meaning that there are no electrical components, including ohmic heaters or PN junctions that modulate refractive indices – fabrication of it is straightforward.

Fabrication of single layer silicon waveguides consists of four main processes: electron-beam lithography [110]–[113], reactive ion etching [114], deposition of spin-on glass [115]–[117], and plasma enhanced chemical vapor deposition [118], [119].

Electron-beam (e-beam) lithography is a serial patterning process in which a resist's solubility is modified when a focused e-beam is directed toward it. If the resist becomes more soluble, it is known as a positive resist; if it becomes less soluble, it is known as a negative resist. E-beam lithography has its advantages over optical lithography due to its ability to reach deep sub-micron resolutions on the order of a few nanometers. Moreover, it allows for fast prototyping since there is no physical lithography mask needed. However, e-beam lithography does have its own set of issues. The actual exposure time is significantly longer and, due to stage movement for large area exposures, there could be misalignments between exposure fields. Also, e-beam lithography is susceptible to exposure dose variations due to forward

scattering and backscattering of the firing electrons in the resist and substrate. Lastly, if the forward acceleration voltage of the e-beam gun is set too high, there is potential for substrate damage.

Reactive ion etching (RIE) is a type of dry etching – gas chemistry – that is generally anisotropic. Typically, a pair of parallel plates generate plasma in a chamber using a radio-frequency (RF) electromagnetic field. Ions are accelerated towards the material being etched not only to react with the material but also to knock off some with high kinetic energy. A special type of RIE uses inductively coupled plasma in which plasma is active instead using an RF powered magnetic field, allowing for high-density plasma to be achieved.

Spin-on glass deposition is not necessarily a categorized microfabrication process but more so the use of a flowable oxide in a spinning process. A spinner is typically used to spread lithography resist evenly over a substrate. After the resist is dripped on it, spinning causes spreading via centrifugal forces. A spinner is versatile enough to allow the use of various viscous materials, including flowable oxide. After the fluid is evenly spread, a curing process over a hot plate or in an oven dehydrates and solidifies it. An important property of the spinning fluid is its ability to planarize the substrate. In other words, some fluids will spread, and the substrate's topology will be transferred to the top surface after the curing process. If a fluid can planarize, the top surface will have less of an imprint and be flatter. This particularly useful when multiple layers of materials will be used. In addition, it also simplifies any optical properties among and between the different layers (e.g., no unintended diffraction effects). Spin-on glass has two attractive properties: it can planarize over

feature sizes less than 100 nm, and it has similar optical properties as thermally grown silicon dioxide after curing.

Plasma-enhanced chemical vapor deposition (PECVD) is a process in which multiple gases react in a chamber and deposit a solid thin film onto a substrate. Plasma activated by RF and/or DC power is used to increase the reaction rates of the precursor gases, allowing deposition at lower temperatures. PECVD materials easily conform to substrates and can be configured to be of high quality. PECVD is one of the faster CVD processes that also allow high film quality; however, the tradeoff is some non-uniformity. Common useful materials in integrated optics that can be made via PECVD include silicon dioxide, silicon nitride, amorphous silicon, and silicon oxynitride.

The following is a CMOS compatible structure that was fabricated in an academic cleanroom using the following recipe (a cross-sectional process flow visualization shown in **Figure A.1**.

Process Description	**Time**
(1) Begin with a Silicon-On-Insulator (SOI) wafer with a 260 nm device layer (Silicon) thickness.	[N/A]
(2) Clean the wafer using RCA process [120], [121]: - Acetone with sonication - H_2SO_4+H_2O_2 (1:1) - NH_4OH+H_2O_2+H_2O (1:1:5) in a water bath @ 80°C - HCl+H_2O_2+H_2O (1:1:5) in a water bath @ 80°C - 10% HF - Deionized water (DIW) rinse and blow dry with N_2 gas	 [10 minutes] [10 minutes] [10 minutes] [10 minutes] [5 minutes] [N/A]

H_2SO_4 Sulfuric Acid H_2O_2 Hydrogen Peroxide NH_4OH Ammonium Hydroxide HF Hydrofluoric Acid HCl Hydrochloric Acid	
(3) Spin XR-1541-006 (HSQ) e-beam lithography resist: - Dehydration bake on a hotplate @ 150°C - Spin HSQ @ 4000 RPM - Softbake on a hotplate @ 80°C	[10 minutes] [60 seconds] [4 minutes]
(4) Expose with e-beam lithography: - Base dose range: 5000 $\frac{\mu C}{cm^2}$ to 8000 $\frac{\mu C}{cm^2}$ - $t_{expose} \approx \frac{Area^2 \times Dose}{Current}$	[t_{expose}]
(5) Develop and clean: - NaOH+NaCl+DIW (1:4:95 %wt) - DIW rinse - Acetone - Isopropyl alcohol (IPA) - DIW rinse and blow dry with N_2 gas	 [4 minutes] [5 minutes] [5 minutes] [5 minutes] [N/A]
(6) Etch with ICP etcher: - HBr + Cl based recipe @ 20°C - High anisotropy and smooth sidewalls - Selectivity of Si:SiO_2 is approximately 50:1 (Note that cured HSQ has similar properties to SiO_2) - $t_{etch} = etch\ rate \times 260nm$	[t_{etch}]
(7) Spin and thermally cure HSQ for thin cladding and planarization (2 coatings): - Spin HSQ at 4000 RPM - Bake on a hotplate @ 90°C - Bake on a hotplate @ 150°C - Bake on a hotplate @ 225°C - Bake on a hotplate @ 400°C - Repeat spinning and baking for 2nd coating - Bake on a hotplate @ 400°C	 [45 seconds] [2 minutes] [2 minutes] [2 minutes] [10 minutes] [16 min. 45 sec.] [10 minutes]
(8) Final cladding with PECVD SiO_2: - Low flow deposition @ 400°C - $t_{dep} = dep.rate \times 200nm$	[t_{dep}]

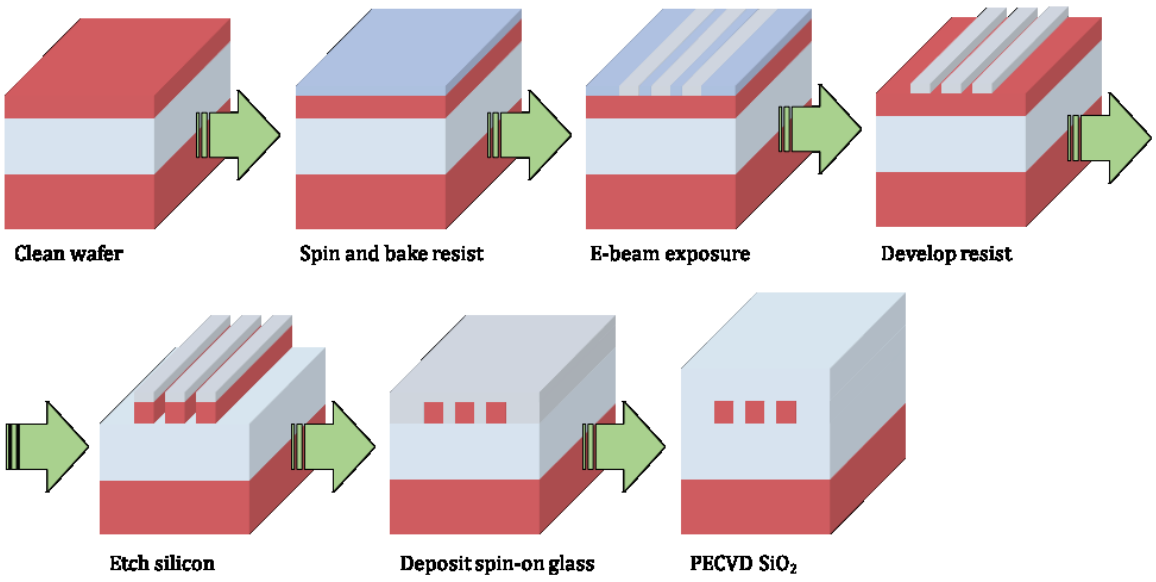

Figure A.1 Process flow for passive device fabrication.

Appendix B: Passive Device Test and Measurement Methods

To characterize the passive performance of a device, we assume the refractive index of the photonic structure is not changing during the measurement – such as when there is a thermo-optic effect from an ohmic heater or the plasma dispersion effect from PN junctions. **Figure B.1** depicts an example of our typical test and measurement setup.

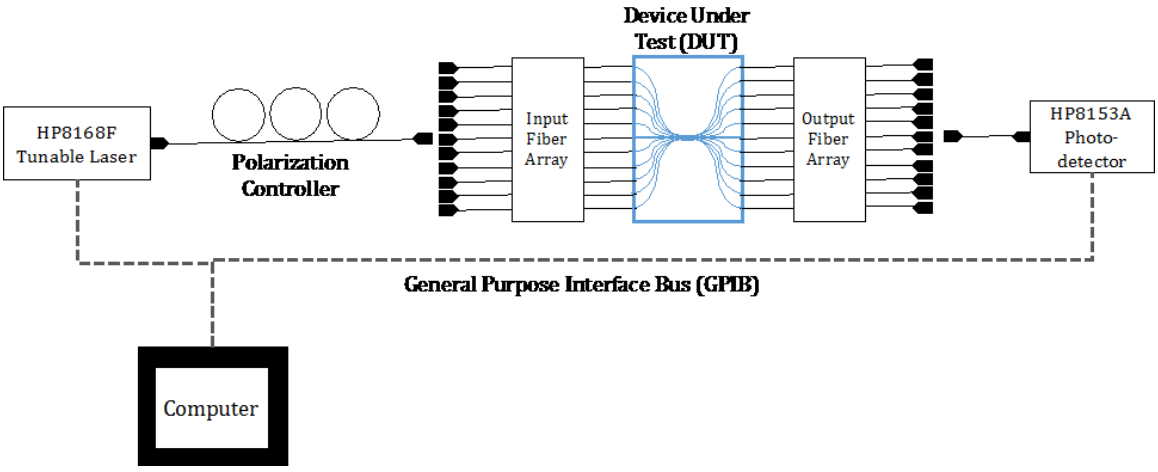

Figure B.1 Measurement diagram for a passive DUT. The DUT currently illustrated is a waveguide superlattice, but this setup can be used for any device (or set of devices) which have multiple optical I/O ports.

Measurement of a silicon photonic device under test (DUT) is done by coupling TE or TM-polarized light from an HP8168F tunable laser via single-mode fiber array into the waveguide facets (or on-chip grating couplers). The output facets/grating couplers couple light into another single-mode fiber array. The input and output fibers to be measured can either be chosen manually or via a 1 x N optical switches.

Testing of silicon photonic DUTs passively is mostly trivial, so automation can help drastically reduce the measurement time. If testing instruments support some

sort of remote communication protocol – such as those listed in **Table B.1** – we can program a test executive to automatically control each instrument to operate in sync.

Table B.1 Common Remote Control Interfaces

Protocol	Max speed (Mb/s)	Connector	Multi-connections
GPIB (IEEE-488)	64	24-pin micro ribbon	Connector stacking
Serial (RS-232)	0.1152	D-sub	Serial Hub
TCP/IP	10000	RJ45, WiFi	Ethernet switch/WiFi Router
USB	10000	USB A/B/C	Serial Hub

General Purpose Interface Bus (GPIB) – formerly known as the Hewlett-Packard Interface Bus (HPIB) – is a legacy protocol that is still widely used in modern instrumentation. Its main disadvantages are transfer speed and its bulky cables and connectors. However, for most small-scale instrumentation platforms, its capabilities are enough. Many old test instruments are often maintained well, refurbished or rebranded. As a result, they do not have more modern remote interfaces, and so GPIB is many times the only option to control it. Serial is another legacy protocol often used as well. With the adaptation of the 9-pin version over the 25-pin version, the D-sub connector became more compact. However, its bitrate is considerably lower than other protocols, so it is often used for simple instructions and small amounts of data transfer. Also, for computers without a built-in serial port, typically a USB to serial port adapter is needed. Instrument control via TCP/IP is a particularly attractive way for instrument control as it can transfer data much faster than other protocols. Sacrificing some speed, one can also make the instrumentation platform completely wireless, saving lab space considerably. Unfortunately, the adaptation of TCP/IP-

based instrument control is relatively recent, so legacy instruments would not be able to be controlled unless older protocols such as a GPIB and serial are used. Similarly, USB-based control is not readily available for most instruments despite its attractive data transfer rate. In our experiments, our instrumentation is controlled via GPIB.

Each measurement is done between 1480 nm and 1580 nm by using HP8153 photodetectors at the fiber array output. This range is within the S, C, and L optical communication bands. The transmission spectra of the DUTs are normalized to a reference waveguide (or the average of a set of reference waveguides) to analyze their respective insertion losses. This also de-embeds the insertion loss of the test and measurement setup itself. An illustration of insertion loss de-embedding is shown in **Figure B.2**.

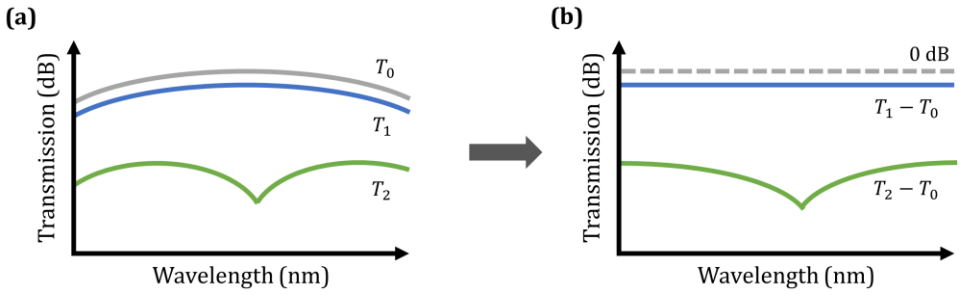

Figure B.2 Insertion loss normalization/de-embedding. (a) Illustrated spectrum T_0 of a reference waveguide with similar length to that of an evanescently coupled asymmetric waveguide pair with transmission spectra T_1 and T_2. (b) The transmission spectra of an asymmetric waveguide pair are normalized to the reference waveguide to get a more accurate analysis of their respective insertion losses.

During testing, however, the optical input power should be high enough to pass light through the DUT and be detecting by the photodetector. For all our measurements, the optical input power of the tunable laser was set to 3 dBm (or 2.0 mW), and the noise floor of our detector was approximately -70 dBm (or 1.0 x 10^{-7}

mW). In low crosstalk measurements of our waveguide superlattice and high loss measurements of our rough waveguides, the actual insertion losses may be lower than the measured performance due to the limitations of our test setup. Lastly, for some temperature-sensitive devices, such as microring modulators and filters, the DUT mount temperature is stabilized during a passive measurement. This is done by either placing the DUT on a Peltier thermo-electric cooler or a hot plate programmed to maintain the desired temperature.