INTEGRATING SILICON PHOTONIC INTERCONNECTS WITH CMOS: FABRICATION TO ARCHITECTURE

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Nicholas Ramsey Sherwood

August 2011

INTEGRATING SILICON PHOTONIC INTERCONNECTS WITH CMOS: FABRICATION TO ARCHITECTURE

Nicholas Ramsey Sherwood, Ph.D.

Cornell University 2011

While it was for many years the goal of microelectronics to speed up our daily tasks, the focus of today's technological developments is heavily centered on electronic media. Anyone can share their thoughts as text, sound, images or full videos, they can even make phone calls and download full movies on their computers, tablets and phones. The impact of this upsurge in bandwidth is directly on the infrastructure that carries this data. Long distance telecom lines were long ago replaced by optical fibers; now shorter and shorter distance connections have moved to optical transmission to keep up with the bandwidth requirements. Yet microprocessors that make up the switching nodes as well as the endpoints are not only stagnant in terms of processing speed, but also unlikely to continue Moore's transistor-doubling trend for much longer.

Silicon photonics stands to make a technical leap in microprocessor technology by allowing monolithic communication speeds between arbitrarily spaced processing elements. The improvement in on-chip communication could reduce power and enable new improvements in this field. This work explores a few aspects involved in making such a leap practical in real life. The first part of the thesis develops process techniques and materials to make silicon photonics truly compatible with CMOS electronics, for two different stack layouts, including a glimpse into multilayerd photonics. Following this is an evaluation of the limitations of integrated devices and a post-fabrication/stabilizing solution using thermal index shifting. In the last parts we explore higher level device design and architecture on the SOI platform.

BIOGRAPHICAL SKETCH

Nick Sherwood was born and raised in San Juan, Puerto Rico, where he spent the first 18 years of his life not going to the beach, but instead playing with computers. He completed a Bachelor of Science degree in Electrical Engineering in 2002; attained from Worcester Polytechnic Institute with High Distinction in Worcester, MA. After a brief stint driving back and forth across the U.S. he came back to WPI where he completed a Master of Science degree in 2005 in the same field. His research focused on satellite communication systems and protocols under the guidance of Prof. David Cyganski.

Nick joined Prof. Michal Lipson's group at Cornell University in 2006, instead of going to Boston University, after a convincing 5-minute meeting on an extended layover between France and Puerto Rico. He explored several research topics as a Ph.D. student including dynamic optical devices, photonic networking systems and integration practices, but always focused on moving integreated optics from a research topic to a practical application. His motivation echoed Mr. Zapp Brannigan's motto: "If we can hit that bullseye, the rest of the dominoes will fall like a house of cards. Checkmate."

To my parents, who would support me in any and every endeavor, no matter how wild or crazy. And to my sister who told me it was ok to be a nerd.

ACKNOWLEDGMENTS

I have to thank first and foremost my advisor and friend, Prof. Michal Lipson. All the awards and accolades she has received (and there are many) do not capture the extent of her character and wonderful personality. The trust and respect she bears her students makes her a pleasure to work with, and attracts a wondefully diverse and smart group of people to her team. I could not have finished a Ph.D. with any other advisor. I would also like to thank my comittee members Prof. Alexander Gaeta and Prof. Alyssa Apsel for their support with this work.

I am also indebted to the Cornell Nanophotonics Group. They made this work possible, and my Ph.D. time not only bearable, but enjoyable. Most of all I would like to thank Dr. Jacob Robinson and Dr. Alexander Gondarenko, who took me in and with whom I spent many nights and days in the office and lab truly applying the motto "work hard, play hard". I would also like to thank Dr. Carl Poitras, Dr. Po Dong, Dr. Bradley Schmidt, Dr. Amy Turner, Dr. Long Chen, Dr. Sasikanth Manipatruni, Dr. Arthur Nitkowski and Dr. Kyle Preston who welcomed me when I arrived in 2006 and guided me until the day they left. I thank also those who came after me, but who helped me no less: Jacob Levy, Lian-Wee Luo, Dr. Jaime Cardenas, Biswajeet Guha, Dr. Mohammad Soltani, Dr. Michael Menard and especially the Brazilian contingency of Hugo Lira, Lucas Gabrielli, Dr. Gustavo Wiederhecker, Bernardo Kyotoku and Dr. Danilo Spadoti with whom I always felt the warmth of the tropics nearby.

Outside of the lab my friends made Ithaca a wonderful place to live, even in the winter. There are too many to mention but I wouldn't have made it without the constant support of Irma Servatius, Amanda Schlenker, and Uziel Colón who never calls me back. I specifically would never have come here if it wasn't for my cousin Juan Lube. He is one of the main reason I'm still here. I also miss our flatmate, Ryan Gorman, with whom we put together a sweet place to live.

Finally I thank my family: Edna, Tim and Maia. Highly educated, but also incredibly smart, they taught me to explore the world and enjoy my life. Every lego set, book, computer, project, and family trip, made me the person I am; I owe them them everything.

TABLE OF CONTENTS

LIST OF TABLES

List of Figures

Chapter 1

Introduction

In the last few decades, we have had an explosion of technology the likes of which has no precedent. Everything from mobile phones to TVs to car radios to home appliances now all look and act much like personal computers. This invasion of technology into our day to day life can be explained, to a large degree, as an effect of the expansion of the Internet. While many gadgets have practical functions, most technological improvements today seem to revolve around interconnectivity. With over 2 billion users on the Internet right now [5], anything that simplifies exchanging text, pictures, music and videos can be the new "killer app". Remixing media on the Internet has become an art form, open to all.

The Internet has become a super-massive repository of data with only more growth in sight. In 2005 Eric Schmidt, the CEO of Google, estimated the size at roughly 5 million terabytes of data [5]. Since then, more and more social networking sites have popped up, video rental has essentially moved online, and voice and video communication are both available "over IP", all in crisp digital high-definition.

Figure 1.1 highlights how this year 5 out of the top 6 applications in North America, using the most aggregate Internet traffic, are sharing media and most of the remaining top 10 are for communication. Needless to say all of this data bandwidth is supported by light and fiber optics. Today's glass fibers can carry terabits of data over many wavelengths for thousands of miles, connecting most of the world's countries together [6]. Such a technology was only necessary to connect large sections of the population together, such as states or at most cities. Today we are starting to see fiber optics reaching directly into our homes [7]. It might take some time, but eventually optics will also have to reach into our own computers.

Table 1 - North America, Fixed Access, Peak Period, Top Applications by Bytes

	Upstream		Downstream		Aggregate	
Rank	Application	Share	Application	Share	Application	Share
1	BitTorrent	52.01%	Netflix	29.70%	Netflix	24.71%
2	HTTP	8.31%	HTTP	18.36%	BitTorrent	17.23%
3	Skype	3.81%	YouTube	11.04%	HTTP	17.18%
4	Netflix	3.59%	BitTorrent	10.37%	YouTube	9.85%
5	PPStream	2.92%	Flash Video	4.88%	Flash Video	3.62%
6	MGCP	2.89%	iTunes	3.25%	iTunes	3.01%
7	RTP	2.85%	RTMP	2.92%	RTMP	2.46%
8	SSL	2.75%	Facebook	1.91%	Facebook	1.86%
9	Gnutella	2.12%	SSL	1.43%	SSL	1.68%
10	Facebook	2.00%	Hulu	1.09%	Skype	1.29%
	Top 10	83.25%	**Top 10**	84.95%	**Top 10**	82.89%

SOURCE: SANDVINE NETWORK DEMOGRAPHICS

sandvine

Figure 1.1: Internet top-usage statistics, emphasizing media applications. ©2011 Sandvine, inc.

1.1 Shrinking Electronics

The economy of integrated circuits has enabled the continued improvement of technology for a long time. The trend known as Moore's Law has continued for more than five decades and tells us that on-chip transistor density doubles approximately every 2 years [8]. Of course, since the prediction was made by one of the co-founders of Intel, one of the largest companies involved in the production of microprocessors, the prediction became the standard by which all other microprocessors were made. Unfortunately the performance increase that came with this trend depended heavily on shrinking the transistors to make them faster. The gate size of the transistor became the standard by which new "technology nodes" were described. At the time of this writing, microprocessors are being fabricated at a 32 nm technology node while transition to 22 nm will occur over the course of this year. The economy of integrated circuits has enabled the continued improvement of technology for a long time. The trend known as

Moore's Law has continued for more than five decades and tells us that on-chip transistor density doubles approximately every 2 years [8]. Of course, since the prediction was made by one of the co-founders of Intel, one of the largest companies involved in the production of microprocessors, the prediction became the standard by which all other microprocessors were made. Unfortunately the performance increase that came with this trend depended heavily on shrinking the transistors to make them faster. The gate size of the transistor became the standard by which new "technology nodes" were described. At the time of this writing, microprocessors are being fabricated at a 32 nm technology node while transition to 22 nm will occur over the course of this year.

Many people have weighed in on the issue and so far there is no consensus as to when exactly CMOS (Complementary Metal Oxide Semiconductor) electronics will physically reach their limit [9, 10]; however it is clear that we will reach atomic scales in the next few years (e.g. silicon atomic spacing is $\tilde{0}.5$ nm). The International Technology Roadmap for Semiconductors reported in 2009 that new technologies "Beyond CMOS" need to be developed now to update semiconductor practices in the near future [11]. Unfortunately the static power losses that face smaller and smaller transistors are a problem that cannot be easily solved [12, 13]. They have led to the development of multi-core processors (CMPs), which operate at the same speeds as previous processors but on multiple times the load. This has displaced the problem of scaling for the time being, but at the same time has focused the need on improving on-chip communication.

Among the biggest limiting factors in processor scaling now is that of interconnects. Typical electrical lines all have some amount of resistance, resulting in loss of power as well as switching speed. While many clever tricks, mostly

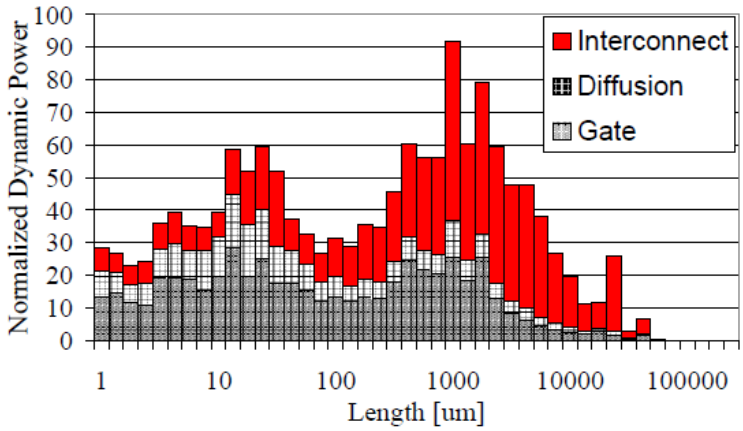

Figure 1.2: Dynamic switching power breakdown versus interconnect length. ©2004 ACM, Inc. [1]

material-related, have been played over the years to improve on this fundamental limitation of electric transmission lines, the limit itself turns out to be "scale-invariant" [14]. Which is to say the limit cannot be solved either by making the wires thicker or thinner, but power loss can still become worse with increased speed, bandwidth, and particularly, length. Figure 1.2 shows the percentage of lost power that goes to interconnects on a typical processor versus channel length. In fact [1] tells us that over 50% of the dynamic power is being lost to interconnects. Which is why we find ourselves at a similar cross-roads to that of the 70's, when fibers finally showed a significant improvement over electrical transmission lines [15]. Particularly now, as systems-on-chip (SoC) emerge as the replacement for serial processing, on-chip communication will play a much bigger role in the performance of computers in general.

1.2 SILICON PHOTONICS

While issues arise in scaling electronics, advances in photonics have steadily continued, including intense research in the area of optical interconnects. Photonic Networks-on-Chip (NoC) have become popular as a viable option for increasing the bandwidth, lowering the latency and reducing the power in chip multiprocessors (CMPs) [16, 17, 18]. By utilizing an optical network to link multiple processors, the full capability of large on-chip parallel systems can be achieved [19] and potentially with considerable power benefits [20]. However, such an optical network is only practical if it can be fabricated monolithically as part of the electronics, reducing the complexity of manufacturing to a few additional CMOS steps. For this reason, silicon photonics has become the ideal candidate for future NoCs. Research groups as well as companies like Kotura, Luxtera, Lightwire, Intel and Sun are all working on public projects relating to silicon photonics. Whether for integration or for independent devices, the economics of using existing CMOS fabrication facilities seems to make sense [17, 21].

Specifically, work in silicon photonics and its related compounds has shown that although the most used element in electronics is not optically active, it can be used in a variety of ways to manipulate and direct light in ways necessary to make integrated communication links possible [22, 23]. Silicon photonics has gained enough interest that a large group of researches are still working on making a compatible laser with many promising options. These include using the raman effect [24], erbium-doped silicon [25] and ge-si [26] as well as many bonded hybrid-material lasers [27, 28].

Using high-contrast material systems, sub-wavelength optical wires can be made along with a host of miniature optical devices to lead the way to feasi-

ble and economic integration of optics and electronics. This, in turn, opens the door to high-bandwidth, low-power interconnects that can be used for applications such as chip multi-processor interconnects or low-jitter chip-wide clock distribution networks [16, 29, 30]. Researchers have now shown full electro-optic links [30, 27] consisting of integrated modulators [31], detectors [32]. More complex devices such as hitless wavelength switches [33] and amplifiers [34] are constantly appearing and could be likewise integrated.

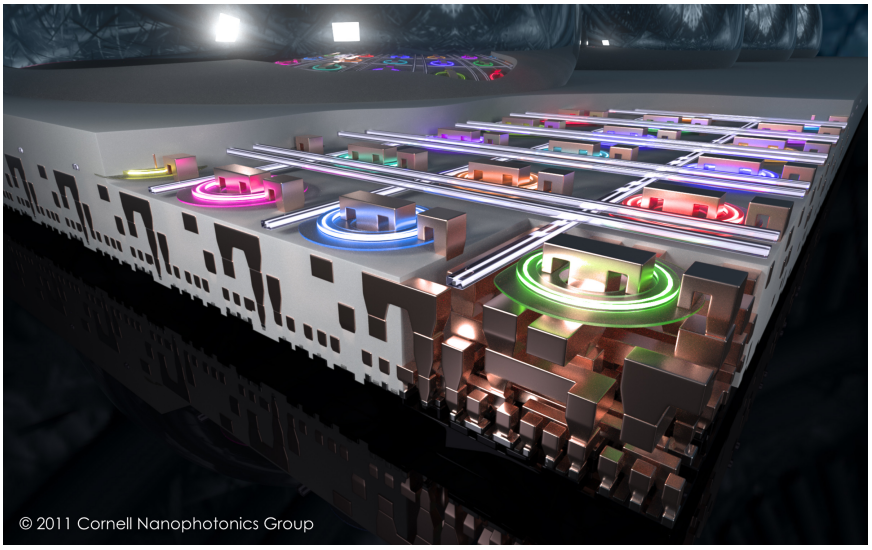

Figure 1.3: Vision of system: deposited optics over CMOS electronics. Multiple bus waveguides sit at different vertical layers above a set of an array of micro-ring resonators.

Unfortunately, even with the extensive exposure it has been getting in the last few years, interest in integrated electro-optic silicon photonics has not really been embraced and developed as a standard. While there may be clear advantages over electronic communication, the integration of photonic interconnects with on-chip electronics requires a costly change to well-established CMOS processes. Even for the companies that have their own fabrication facilities this can be hard, but for smaller companies who wish to enter this market

it becomes all but impossible. It is clear that further integration research is necessary to achieve CMOS photonics; we need more than similar processing, but completely compatible photonic processes that can be added as-is to any CMOS fabrication run.

We envision in our research, an integrated chip with electrical and optical devices working in unison to improve performance. Due to economic, scaling, and physical constraints such a design may not have the best electrical configuration. Likewise it may not be the fastest or lowest-power photonics available, but the combination might just be greater than the sum of its parts. Figure 1.3 shows our vision of a 3D optical network comprising of several photonic layers above an electronics die, assisting but not replacing all communication links within the chip.

1.3 Thesis Focus and Layout

This thesis covers several subjects, all of which focus on the same goal of optical CMOS integration. The document is organized from the bottom up. It begins with low-level materials and process integration and moving towards higher-level device and system integration.

Chapter 2 explores the limitation of the silicon-on-insulator (SOI) platform, which although common and easy to use, still limits us to making discreet optical components. The goal of this work was to develop a fabrication process which could be used on ordinary bulk silicon wafers to make optical devices, and by the same token then replicated on the front-end of a CMOS stack. We limited ourselves only to fully compatible materials and techniques, and developed a method for adding optics to electronics processes in the near future.

Chapter 3 looks at a fabrication alternative to add optics to the back-end of

the CMOS stack. Given the low temperature limitations of this area, we examined the possibilities available to us in terms of materials and processes, and the best way to accomplish our vision of integrated, and still useful, optics. This method allows for the least amount of effort on the part of the fabrication facility in order to combine electronics and photonics without any change to the original CMOS layers.

In Chapter 4 we discuss at a device level some of the limitations of optics due to material and process issues. We take it as a given that these limitation cannot be changed in the near future, and we investigate how to manage the resulting problems. Specifically, we look at thermal control to solve these problems, and beyond that to create opportunities for dynamic systems.

Chapter 5 moves to the architectural level where we design a higher level implementation of a networking system. Using thermal control discussed in the previous chapter, we implement a 4x4 dynamic non-blocking optical router. The configuration is built on an SOI platform, then tested and measured electrically and optically using a variety of methods.

Finally, in Chapter 6 we take a look at the possibility of using our devices in the field. While typically designed for on-chip integration, silicon photonics is wavelength-compatible with telecom equipment making it potentially useful in this area. A thermally tuned wavelength filter is taken to a geographically-distributed fiber network to process high-speed data operation under real-world conditions.

CHAPTER 2

TRANSISTOR-LAYER INTEGRATION OF PHOTONICS USING A LOCALIZED BOX

This chapter uses material extracted with permission from the following references:

N. Sherwood-Droz, A. Gondarenko, and M. Lipson, "Oxidized Silicon-On-Insulator (OxSOI) from bulk silicon: a new photonic platform," Opt. Express 18, 5785-5790 (2010).

2.1 INTRODUCTION

Integration of photonics with electronics would ideally be done in what is called the front-end of the CMOS stack. This is where the transistors reside, which also means where high temperatures and quality materials can be found. In general pure materials are preferred for photonics, much like with fibers, to avoid optical losses. Unfortunately integrating photonics is not as easy simply displacing transistors.

One of the main drawbacks of the integration of optics on the transistor layer is the need for extensive cladding material, typically a few micrometers of silica surrounding the silicon core waveguide. Although high-contrast optics requires less real estate, the optical wire - a typical Si channel waveguide (e.g. 450 nm wide x 250 nm thick), still needs around 1 µm of SiO_2 cladding on all sides. This optical buffer is particularly important underneath the waveguide in order to eliminate crosstalk to the nearby substrate. In electronics however, only a small layer of SiO_2 is needed. In fact, having more than a few hundred nanometers of SiO_2 underneath the Si waveguide creates incompatibility with the CMOS transistor layer, which necessitates thermal conduction through the silicon substrate [35].

Modern processes which use SOI (Silicon-On-Insulator) technology rely on a "thick" silica layer around 70 nm and better transistor performance can be achieved with a 40 nm BOX (Buried OXide) layer [36]. Yet even with older SOI CMOS stacks which rely on up to a 400 nm BOX [35], optical losses from radiation would be in the thousands of dB/cm. Comparatively from our simulations, 1 µm of BOX would lead to 0.35 dB/cm and 3 µm would lead to negligible loss assuming chips with several centimeters of waveguides between nodes.

One option for the integration of optics with electronics is the fabrication

of the optics layer in the back-end of the CMOS stack, alongside the metalization layers [37]; however, single-crystalline silicon would not be available as a guiding material, since it cannot be deposited. In addition, this approach limits fabrication to low-temperature processes when done monolithically. Alternatively, one could envision using flip-chip bonding; however, this would reduce the already limited area reserved for off-chip contacts. Instead it would be preferable if the optics could be built alongside the transistor layer where high-temperature processes are allowed and close contact with electronics is maintained, allowing the architect full flexibility in design.

Here we propose a fabrication process for enabling the integration of optics on the transistor layer. The process is based on the traditional electronic medium, namely a bulk silicon wafer, and produces a small or non-existing BOX layer for the electronics while providing an optical buffer for the optics sharing the layer. We refer to this as an Oxidized Silicon-on-Insulator (OxSOI) stack.

2.2 Process Definition and Simulation

The overall process relies on creating high quality oxide out of bulk silicon beneath our waveguides using thermal oxidation, having protected them with a material with very high oxidation temperature and CMOS compatibility: Si_3N_4 . A similar technique is used in some CMOS fabrication in a process called LOCOS (LOCal Oxidation Of Silicon), in which oxide spacers are created between transistors to avoid parasitic effects. LOCOS has been used to make waveguides, though typically they are large ridge waveguides with little confinement [38, 39]. Smith et al. have previously shown the formation of waveguides by excavating air pockets underneath waveguides made of polycrystalline silicon

[40]. This process requires depositing the guiding layer since it needs a buffer layer to protect the guiding material from the silicon-etching chemical. By maintaining the original crystalline silicon above our oxide layer we ensure a high quality optical guiding medium is used.

We induce oxidation from beneath the waveguide in order to optically isolate the waveguide from the substrate. To do this we use a full three-sided nitride cap as a protective cover over an etched pillar that will be formed into the final waveguide (see Figure 2.1). In electronics, a similar process called SWAMI (SideWAll Masked Isolation) is used in some transistor fabrication [41]. Using the Silvaco modeling program Athena, we simulate the oxidation growth for our waveguide structure, shown in Figure 2.1. We generally find good agreement in the resulting shape of the structure in these simulations, though the simulation underestimated the oxidation time.

In Figure 2.1(a) the waveguide-like structure has been clad with Si_3N_4 and over-etched by just under 1 μm before being subjected to oxidation. Under oxidation the structure narrows from either side (Figure 2.1(b)) until it separates from the substrate, at which point we can begin to see the formation of a "buried oxide layer", shown in Figure 2.1(c). At this point we continue to oxidize for two reasons: 1) to increase the distance from the waveguide to the substrate and 2) to reduce the triangular "tail" of the waveguide to a flat surface as seen in Figure 2.1(c). Further flattening can be achieved with increased oxidation time until achieving a shape similar to a typical channel waveguide (Figure 2.1(d)).

2.3 FABRICATION

The fabrication process was developed with compatibility and simplicity in mind. We constantly remind ourselves that the CMOS industry is extremely

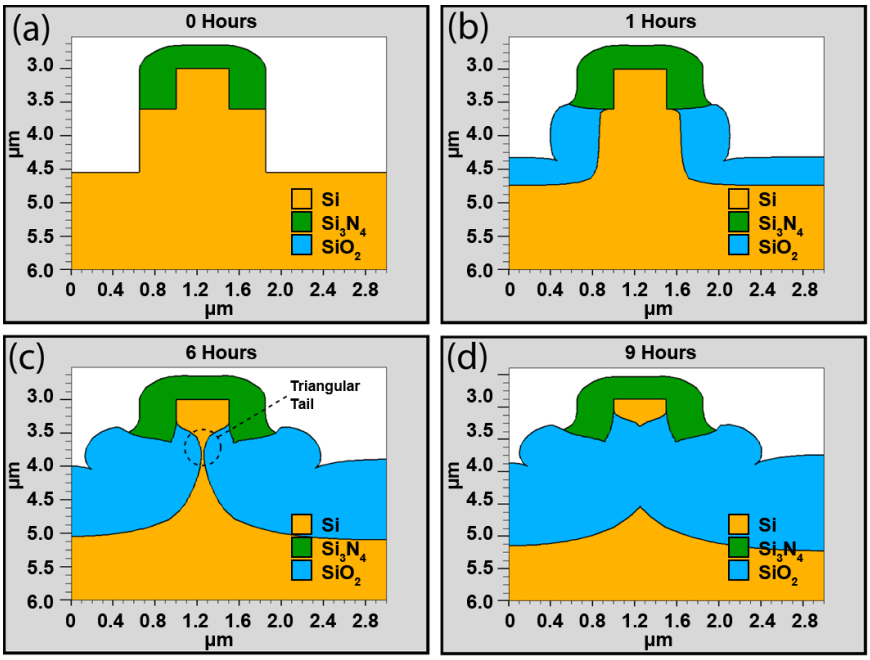

Figure 2.1: Simulated effects of wet oxidation on waveguide structure: (a) post over-etching and prior to oxidation, (b) 1 hour of oxidation, (c) 6 hours of oxidation, (d) 9 hours of oxidation.

inflexible and costly. Therefore, the fabrication process described here relies on a single step of lithography, deposition of Si_3N_4, etching and oxidation. The details are discussed here in detail.

We begin with a blank silicon wafer (Figure 2.2(a)). The assumption lies that this process could be recreated at the transistor layer of the CMOS stack, or just below it, where bulk silicon is accessible. On this layer we deposit a 300 nm layer of Si_3N_4 using PECVD (Plasma-Enhanced Chemical Vapor Deposition) or LPCVD (Low-Pressure Chemical Vapor Deposition). This layer serves as the top cover of the protective mask shown in Figure 2.2(b). Stoichiometric nitride, as produced by the LPCVD process, tends to be a stronger material and could be deposited thinner. We will see below that thickness plays a role in the final shape of the waveguides due to limitations in the etching process.

We then define the waveguides patterns over the nitride layer. These patterns are usually exposed on negative resist using an electron-beam lithography tool. For the simple reasons of cost and time (e-beam sweeps a single beam, rather than exposing a full pattern), the industry does not use this tool to make wafer patterns. At this point, however, current lithography has reached a point where it can pattern 0.5 µm lines with consistently low roughness using DUV (248 nm) [42]. While we did not have this technology at the time at our research facility we did develop an i-line (365 nm) photolithographic process which would give us 0.5 µm lines as well to maintain compatibility with CMOS. Photolithography typically uses a positive resist, in our case we use SPR 955CM 0.9. It is important that we invert our patterns in order to use this method. This step results in a hardened resist material over the areas to be protected (Figure 2.2(c)). The waveguides are defined to be 450 nm wide, our standard geometry to maintain single-mode operation and tight bending radii [43].

The waveguide layer is then etched vertically down using an ICP (Inductively-Coupled Plasma) RIE (Reactive-Ion Etching) tool. The same anisotropic silicon etch composed of Cl_2, BCl_2 and H_2 which etches silicon is also used to etch the silicon nitride at a slower rate. We continue to etch using the same recipe into the bulk silicon, in this case for 600 nm. The depth of this first etch determines how much room we will have later during the oxidation for a buried oxide layer as well as for smoothing the underside of the waveguides. We made the pillars purposefully taller than the designed waveguide thickness of 250 nm specifically to try to achieve the flat bottom surface (Figure 2.2(d)). Removing the photoresist with an oxygen asher results in a waveguide-like pillar with an Si_3N_4 layer on top.

Oxidation at this point would result in a step which could be used as a

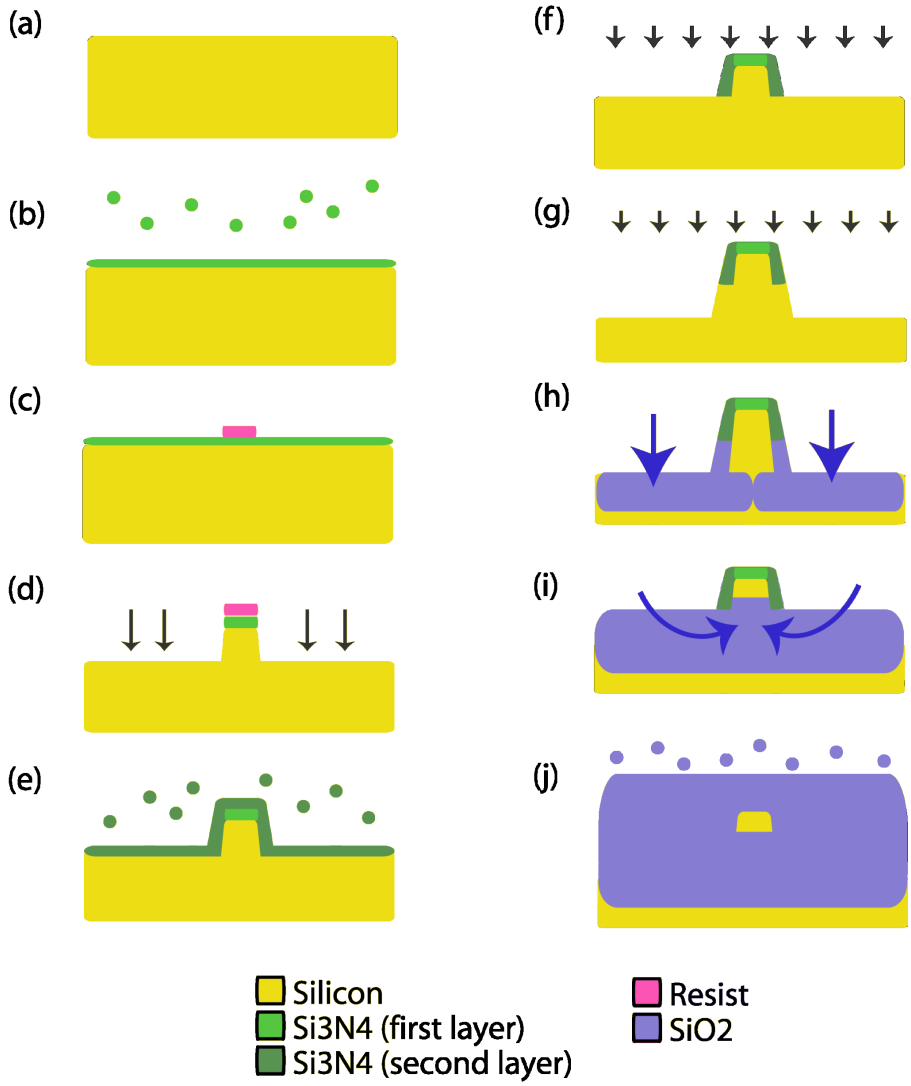

Figure 2.2: Process flow: (a) bulk silicon wafer, (b) Si_3N_4 deposition, (c) lithographic waveguide definition, (d) ICP etch, (e) Si_3N_4 deposition, (f) ICP cap etch, (g) extended etch for quicker oxidation, (h) wet oxidation for buffer layer growth, (i) extended wet oxidation for waveguide under-layer flattening, (j) upper layer oxide deposition to complete optical buffer.

waveguide only if the beginning stack was SOI and had a buried oxide layer (BOX). The following step however separates us from previous work in the oxidized waveguide category by allowing us to make channel waveguides as well as creating a full BOX. The addition of sidewall masking allows the waveguides to be first, finite in width, and second, completely isolated. We deposit another 350 nm of Si_3N_4 , shown in Figure 2.2(e), over the full wafer, which covers the sidewalls of our structures in silicon nitride.

In order to leave the three-sided mask over the waveguide regions while exposing the bulk silicon to oxidation, we need to etch all of the nitride that does not reside over our pattern. This would typically require an extra step of lithography, however we we use a clever trick to avoid this. When the Si_3N_4 layer is etched (by the same amount as was deposited in the last step), the verticality of the etch leaves Si_3N_4 on the sidewalls of the pillars. Given the initial layer of Si_3N_4 deposited above the silicon, we have as a result a three sided mask over all waveguiding structures (Figure 2.2(f)). We refer to this step as being self-aligned, since the mask is left only over the waveguides without the need for the additional lithography. As an optional step, we can continue to etch into the silicon substrate to reduce the amount of oxidation time and SiO_2 generated on the sides (Figure 2.2(g)). If etched as trenches, this step would reduce the spacing necessary between adjacent waveguides or directionally coupled devices.

The wafer is then placed in a wet oxidation chamber at 1100°C . This step grows SiO_2 in the regions of exposed silicon. The SiO_2 reaches underneath the structure creating the BOX layer shown in Figure 2.2(h). Figure 2.4 shows the rate at which SiO_2 grows from the silicon. At first a triangular bottom is formed at the bottom of the waveguide due to the isotropic nature of the oxidation on both sides. Such a waveguide can still work and is discussed in the results sec-

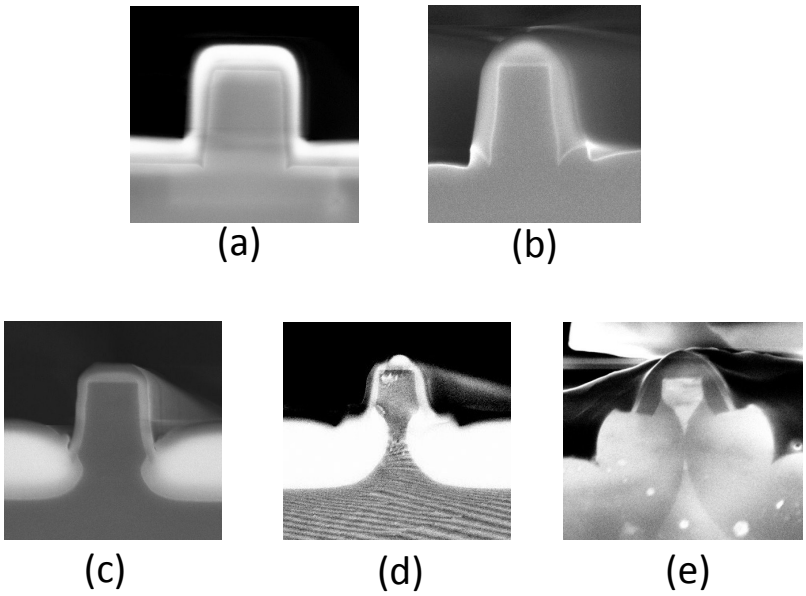

Figure 2.3: Cross-sectional SEM views of various waveguides during different process steps: (a) after second deposition of silicon nitride, (b) after second nitride etching (ICP birdbeak features notable), (c) after 2 hours of oxidation (no over-etch), (d) after 6 hours of oxidation (no over-etch), (e) after 6 hours of oxidation (800 nm of over-etch).

tion, but extended oxidation into the underside of the Si pillars can reduce the pinched point into a flat underside (Figure 2.2(i)). Figure 2.3 shows qualitatively what the process looks like closeup. Figures 2.3(a) and (b) show a prepared 3-side mask. Figures 2.3(c)-(e) show how the oxidation looks over the course of a few hours and can be compared to the simulation considerably well. At 1100°C we can leave the oxidation for a total of 16 hours to create the final waveguide we tested. This is clearly a limitation but could be reduced considerably with a higher temperature process and by designing a narrower waveguide.

The Si_3N_4 layers can be removed at this point, though from our simulation it does not disturb the mode profile and so we leave it on the waveguide to avoid additional processing steps. The structures are clad with 3 µm PECVD SiO_2 to

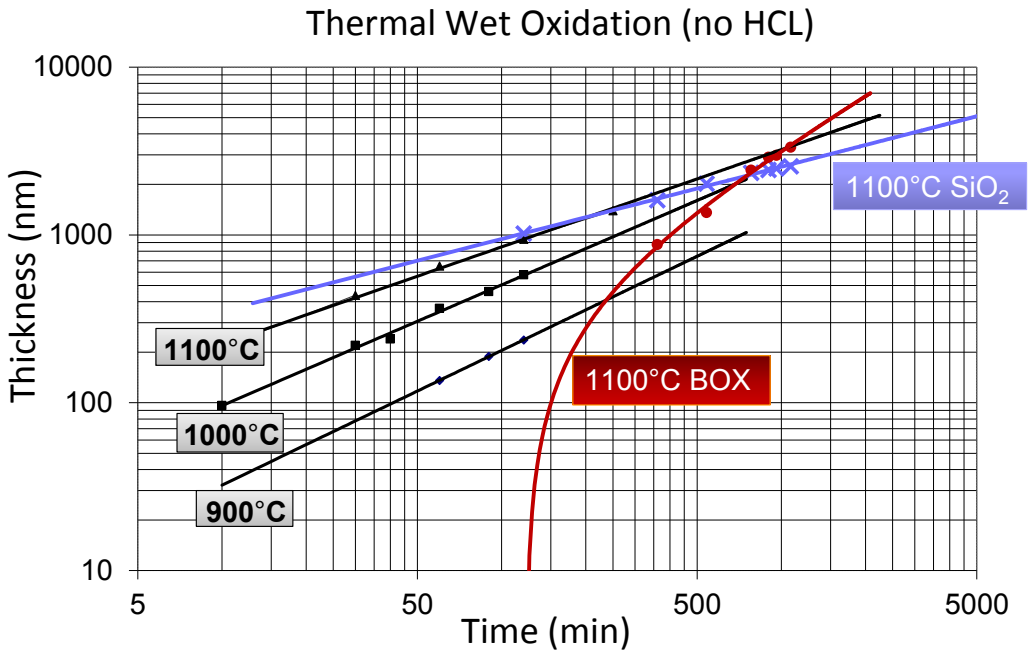

Figure 2.4: Graph of wet oxidation times without HCl for silicon at various temperatures, along with measured rate for our structures, along with the plot of the buried oxide layer thickness, measured from the bottom of the "waveguide" to the beginning of the substrate.

complete the surrounding optical buffer (Figure 2.2(j)).

2.4 RESULTS

We tested our waveguides at different points in the process to quantify our results, and we divide them into the "first set" and the "second set". They are both waveguiding structures; however, the second set more closely approaches what we envisioned as a channel waveguide, while the first shows a triangular peak which makes it into a 5-sided waveguide.

2.4.1 First Set

After six hours of oxidation on a set of waveguides with an over-etch of 800 nm, we achieve separated waveguides with a 0.5 µm BOX layer. Theoretically,0 this might not be enough to prevent leaky modes from existing, but the uncommon geometry pushes the mode upward and away from the triangular peak making it able to guide light well. Removal of the silicon nitride can be done, but it results in some surface roughness of the silicon. Alternatively, as we see in the optical mode in Figure 2.5(b), the small amount of nitride does not affect our mode adversely; therefore, we can safely leave it on.

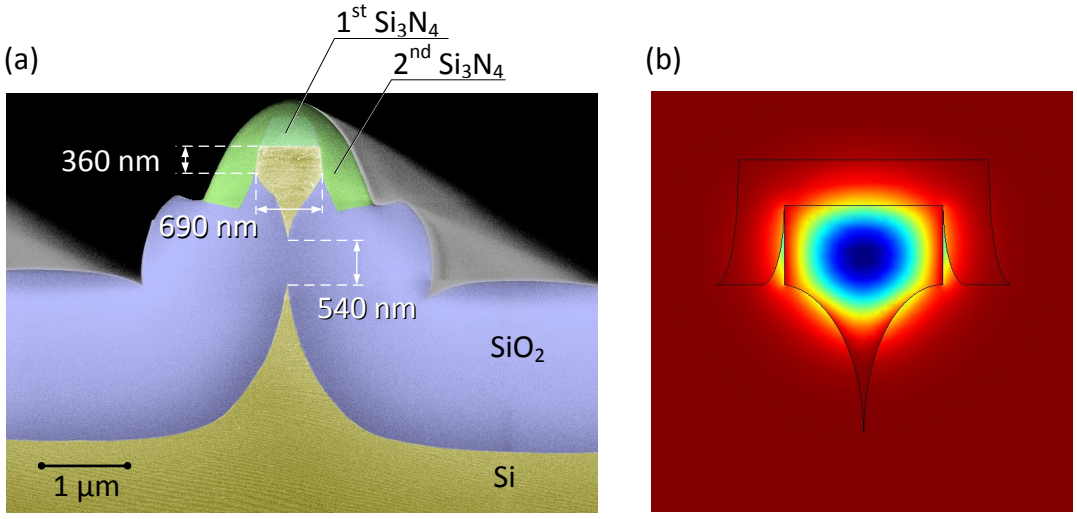

Figure 2.5: First waveguiding structure: (a) false-color SEM of waveguiding structure; and (b) numerically calculated optical TE mode for given geometry.

We measured the losses of the waveguides using the cutback method. This involves making waveguides of different lengths and measuring the absolute loss of each, which can then be compared to the set to extract the loss component from the length variation. Fiber coupling was used to guide light at the standard telecom wavelength of 1550 nm, incident from the polished edge of

the chip. Light was collected at the output edge with a collimating objective and power detector. Figure 2.6 shows the measurement results for two different chips. The propagation loss fits for both are on the order of 10 dB/cm, which is not great, but encouraging. This is our first demonstration of crystalline silicon waveguides made directly on a bulk silicon wafer.

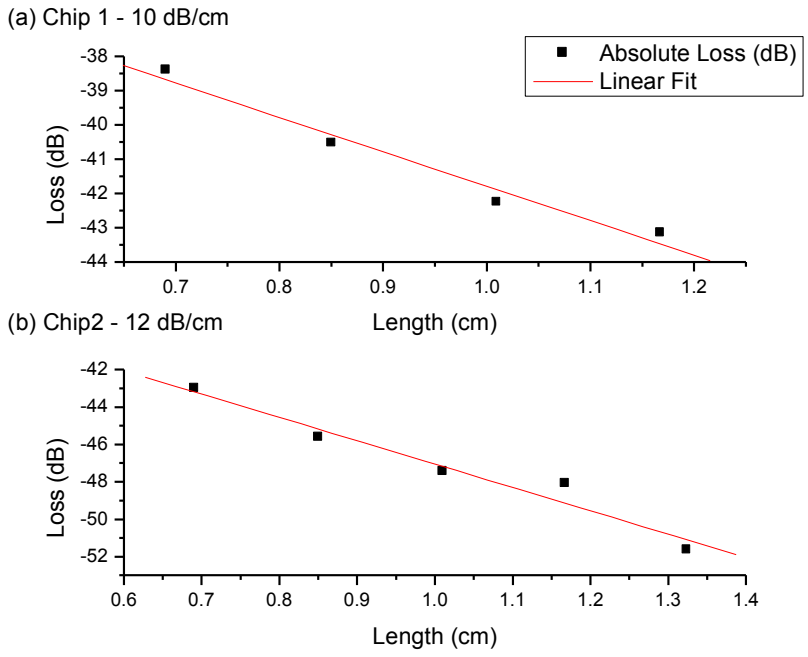

Figure 2.6: Propagation loss measurements and fits for two chips, shown in (a) and (b).

While our propagation losses were relatively low the absolute loss was very high; yet they were similar between multiple chips. We found a likely issue after some additional simulation. While leaving the silicon nitride above the full size waveguides is not a problem, we realized that the high index material above our reverse-tapered region might be constraining the mode. Figure 2.7 shows a comparison of two nanotapers. Reference [43] explains the need for a nanotaper

for proper facet coupling; the small waveguide delocalizes the mode enough to match that size of fiber mode in order to properly couple. However, when we have a nitride cladding, the mode is still contained in a high index region. This reduces the efficiency of light both coming in and going out of the chip. Such a limitation makes a case for either removing the nitride reducing the size of the taper further.

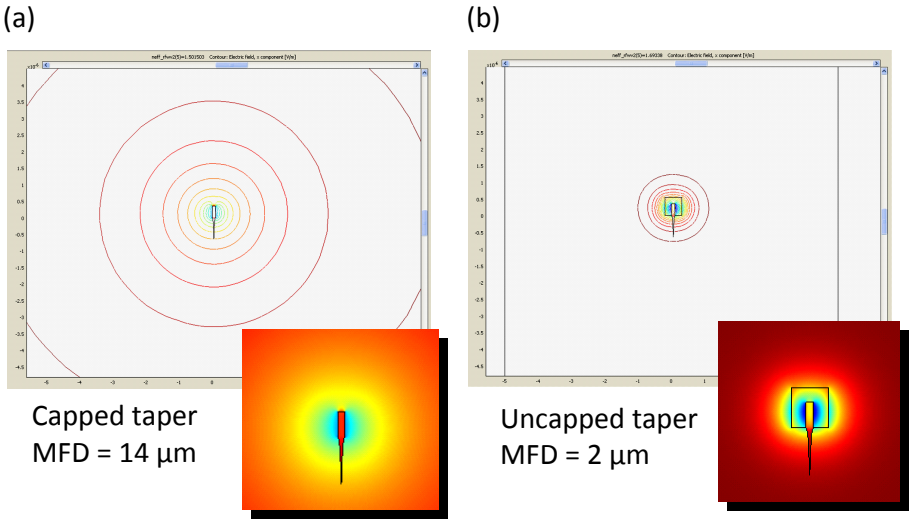

Figure 2.7: Mode field diameter simulation for a 100 nm taper (a) capped with Si_3N_4 and (b) clad only with SiO_2 .

2.4.2 Second Set

In this second set of waveguides, the overall shape of the fabricated waveguide itself is rectangular (seen in Figure 2.8(a)) and provides a confined optical mode shown in Figure 2.8(b), still unaffected by the leftover Si_3N_4 . Using this process we can indeed create a sizable BOX layer beneath our crystalline silicon waveguide. Figure 2.8(c) shows the 3.3 µm BOX layer beneath three adjacent

waveguides.

The thicknesses of the two nitride layers here were determined through various trials in order to avoid over-etching (visible in the thin top corners of the Si_3N_4 layer in Figure 3(a)) that would lead to unwanted oxidation of the waveguides. The thicker cover however, coupled with the 12 degree sidewall angle produced by our ICP etcher and an additional patterning layer (150 nm SiO_2 hard-mask) resulted in an increase of the waveguide width. The outcome is a waveguide 740 nm wide. A smaller waveguide can be made in two ways: either writing 160 nm patterns which would result in 450 nm waveguides, or ideally using an etcher with higher verticality that would also reduce the necessary thicknesses of the two Si_3N_4 layers. Given a perfectly vertical etcher, the thickness of each of the Si_3N_4 layers would need to be slightly over 100 nm each to provide the necessary protection for the given oxidation time [44].

We tested the fabricated waveguides for propagation losses again using the cut-back method. We observed similar losses for TE and TM modes and report here on unpolarized light. We measured 2.92 dB/cm loss shown in Figure 2.9. From our simulations, a BOX greater than 3 µm should result in radiation losses into the substrate of less than 0.001 dB/cm, and likely were the major loss source in the previous set of waveguides. We believe the propagation losses could still be significantly reduced with improved etching techniques and should be limited to single crystal silicon waveguide losses. It should be noted that a low quality PECVD Si_3N_4 was used in these samples which could result in some material absorption. This can easily be resolved through the use of LPCVD Si_3N_4 .

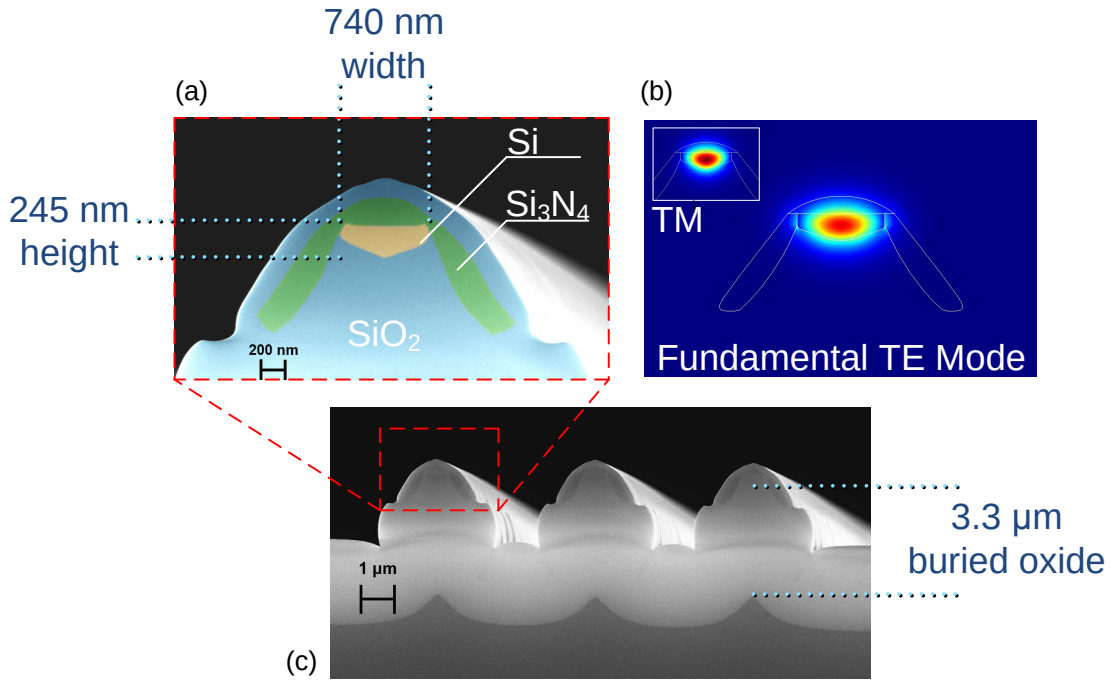

Figure 2.8: Final experimental waveguide shape shown: (a) closeup false color SEM of waveguide cross-section; (b) fundamental TE mode with effective index of 2.63, TM shown for comparison; (c) SEM showing buried oxide layer.

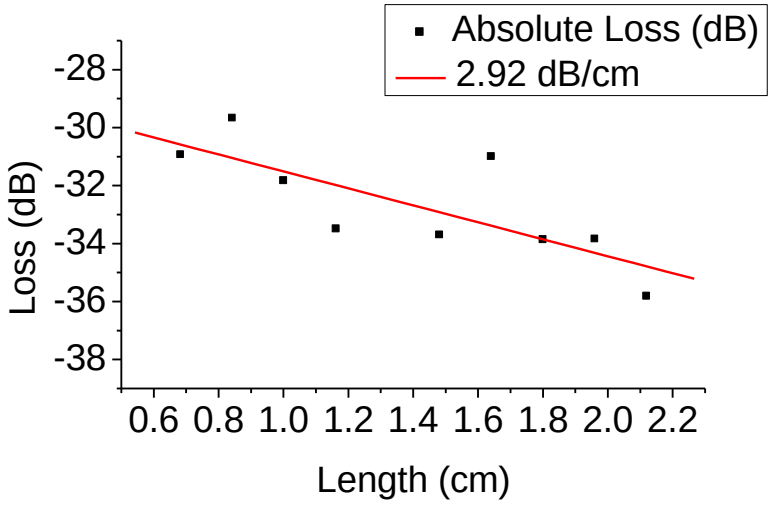

Figure 2.9: Experimental loss for varying lengths of waveguide measured at 1550 nm, and loss per length fit of 2.92 dB/cm.

2.5 CONCLUSIONS

We have developed here a process which produces small (sub 1 μm width) waveguides with channel geometry, on a common bulk silicon wafer with losses of 2.92 dB/cm. The final propagation losses achieved are indeed comparable to those of typical SOI waveguides, and likely come from sidewall etching roughness to be improved at the foundry level. The waveguides are made of single-crystalline silicon, which is important for maintaining high contrast guiding in small footprint devices with low losses. The optical layer is made directly on a real CMOS compatible wafer with real CMOS compatible processes. This process can be integrated at the transistor layer of a currently fabricated electronics chip. Finally the process requires only one lithography step, thereby reducing the cost and the variability in the fabrication of both the photonic material stack as well as the devices.

CHAPTER 3

SCALABLE 3D DENSE INTEGRATION OF SILICON-BASED PHOTONICS

This chapter uses material extracted with permission from the following references:

N. Sherwood-Droz and M. Lipson, "Scalable 3D Dense Integration of Photonics on Bulk Silicon," to be published.

A. Biberman, N. Sherwood-Droz, X. Zhu, M. Lipson, and K. Bergman, "High-Speed Data Transmission in Multi-Layer Deposited Silicon Photonics for Advanced Photonic Networks-on-Chip," in CLEO:2011 - Laser Applications to Photonic Applications, OSA Technical Digest (CD) (Optical Society of America, 2011), paper CThA1.

3.1 INTRODUCTION

The reality of integrating photonics and electronics is that beyond compatibility, a solution needs to provide enough enhacement to the system to be an economically worthwhile improvement [8]. Specifically, optical interconnects might need to provide not just one optical Wavelength Division Multiplexing (WDM) link, but many in order to surpass the multi-Tbps electric interconnects now in use in recent multi-core processors [45]. Given the limits of chip real estate, the best way to dramatically improve the bandwidth density is to break out of the single optical-plane limitations of crystalline-silicon photonics.

Multilayer photonics can go beyond simple integration, providing interesting advantages to the system as a whole. Single layer interconnect photonics is typically criticized for not providing enough performance gain in terms of bandwidth and power to justify the integration challenges. Losses from waveguide crossings for example, a necessary limitation of single layer optical networks, are often cited as one of the biggest obstacles in these systems [21, 46]. A multi-layered system could reduce or eliminate these restrictions altogether by avoiding physical crossings. We would also give architects a new dimension to explore, perhaps leading to denser as well as more complex networks with radically higher cross-sectional bandwidth and reduced communication power consumption. Many researchers have already published work which shows optical networks winning against electronic ones, but only when provided with enough WDM channels [47, 48, 49]. 3D stacking now opens the door to a new area of systems research where photonics can be implemented in ways that have only been theorized.

3.2 VERTICAL INTEGRATION

A major challenge in physically integrating optics with electronics still lies in the general incompatibility of optical materials and structures with CMOS processes. Research groups often rely on a Silicon-on-Insulator (SOI) stack with a guiding medium of single-crystal silicon (c-Si) above several micrometers of optically insulating silica (SiO_2). While the materials are directly compatible (unlike III-V materials), the c-Si must still be grown independently, implanted, annealed, cut, oxidized and bonded to another structural Si wafer before it can be used for device fabrication. More important than the additional cost however, is the limitation of having only a single layer of optics, one which still typically relies on a packaged or bonded electronics die for operation. Being able to deposit multiple layers monolithically made CMOS scalable and successful; a limited separate optical solution will likely not do as well.

The main limitation in integration is still optical isolation: large amounts of dielectric which directly translate to reduced real estate and increased thermal isolation for electronics. The size of the necessary optical buffer depends on the geometry of the guiding structure as well as the refractive index contrast between the two layers. Unfortunatly, even for high-contrast waveguides such as silicon, this is typically larger than 1 µm in single mode operation. Even electronic fabrication on SOI uses too thin of an oxide layer. At its largest an electronic SOI BOX is 400 nm [35], while newer processes use as little as 40 nm [36], making them unsuitable for this work. Figure 3.1 shows such a diagram of such an SOI stack. The transistor layer is packed in the front-end of the wafer, barely separated from the Si substrate. And in fact research groups including ours have shown that this layer can be re-purposed for optics, as was described in the previous chapter, and in [50, 40]. However, this form of integration comes

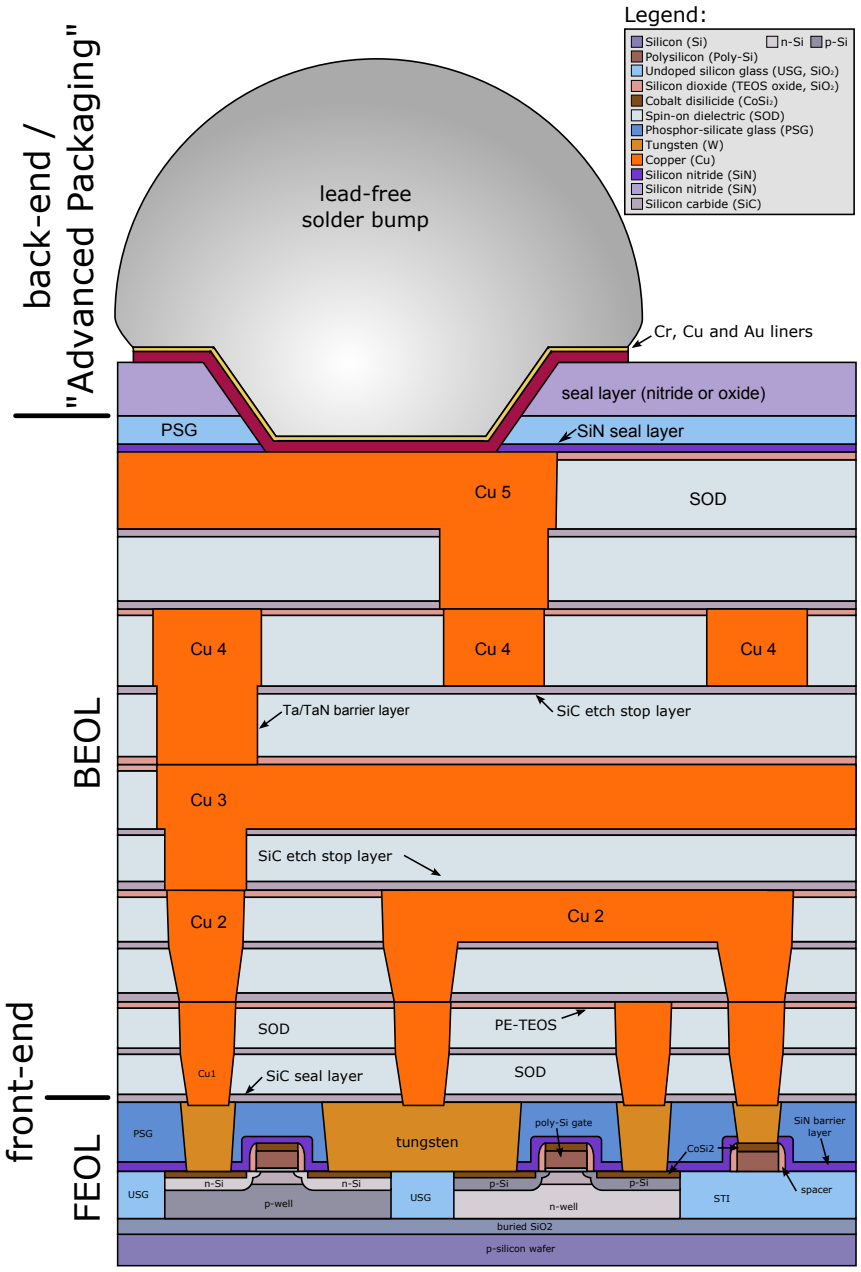

Figure 3.1: Schematic structure of a CMOS chip, common in the early 2000s. The graphic shows field effect transistors on an SOI silicon substrate with five metalization layers and solder bump for flip-chip bonding. It also shows the section for FEOL (front-end of line), BEOL (back-end of line) and first parts of back-end process. ©Henrik Schumacher, from Wikimedia Commons

at the cost of more real estate along with some changes to the standard fabrication. In this work we propose that a practical and useful integration of optics and electronics in the near-term depends on the vertical back-end stacking of deposited photonics with no change to the front-end electronics.

We report here on our first-generation multi-plane wavelength-switched optical link made from fully CMOS-compatible materials and processes. We believe a low-temperature many-layer process is critical in bringing optics and electronics together in a practical way, as there is so far no collective decision in industry regarding the best process or material stack for integrated optoelectronics. Some research has proposed exploring this route, such as with air-clad devices [51], but with no visible adoption by industry. Providing a method to expand further vertically will directly increase the bandwidth/area available to the point where it can improve over electronics. The International Technology Roadmap for Semiconductors (ITRS) 2009 specification notes that "for [cost-efficient] optical solutions to be viable, the development of CMOS-compatible optical components is of paramount importance". [11]

3.3 PECVD Materials

Along with using group IV compatible materials, the most important detail in back-end photonics is using processes around or below 400 °C to avoid distressing the metalization and changing the dopant diffusion. Our fabrication needs to be governed by plasma enhanced chemical vapor deposition (PECVD) of all layers, planarization to allow for vertical stacking, and proper patterning and spacing of structures to allow for coupling and insulation. Regarding materials, we need a set of two with an appreciable refractive index ratio, which will keep bends and inter-waveguide spacing small. Silica is commonly used as an

Table 3.1: CNF GSI PECVD Material Recipes.

Parameter	Si_3N_4	SiO_2
Gas 1 (sccm)	28 (SiH_2)	18 (SiH_2)
Gas 2 (sccm)	1900 (NH_3)	1800 (N_2O)
RF Power at 13.56 MHz (W)	350	100
Temperature (°C)	400	400
Pressure (Torr)	3.5	1.4

outer cladding layer as it has a fairly low refractive index of around 1.45 at telecom wavelengths (1.55 µm), and it is found in abundance in CMOS. Some common waveguide materials include silicon nitride (Si_3N_4) [52, 2] and amorphous silicon (a-Si) [53, 37] as low optical loss waveguides. Polycrystalline silicon (p-Si) and poly-germanium (poly-Ge) have been explore for active devices like switches and modulators [54] and for detectors [55] respectively.

We chose to work with Si_3N_4 as a guiding medium for its good optical characteristics. Si_3N_4 is readily found in CMOS for passivation, masking, and dielectric layers [56], and provides lower optical losses than c-Si in the NIR [57, 58] when using SiO_2 cladding which is also a CMOS dielectric. Although low pressure chemical vapor deposition (LPCVD) silicon nitride has been shown to have losses as small as 0.1 dB/cm [59], most transmission distances at the chip-scale will be less than a few centimeters. This gives us leeway to consider other alternatives.

Given the expected absorption of Si-H and N-H bonds centered on $\lambda = 1.5$µm and shown in Figure 3.2, we explored a variety of PECVD recipes to reduce potential losses. One group suggests adding helium into the plasma to mitigate the hydrogen bonding environment [60]. Others utilize the lower frequency options of the PECVD tool to reduce bombardment of lower weight hydrogen atoms [52, 61]. While preliminary testing did see some overall im-

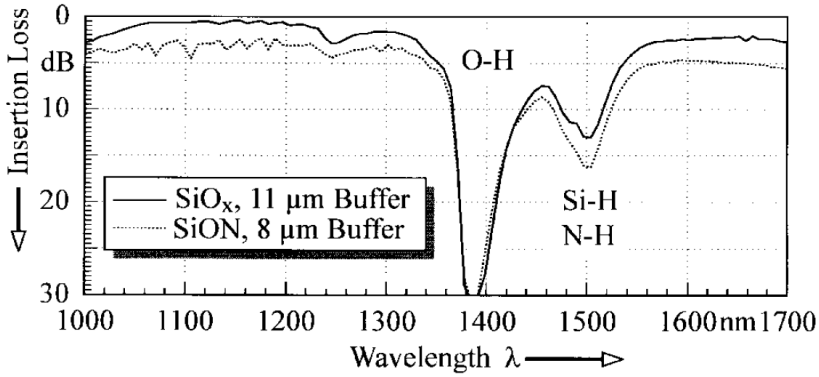

Figure 3.2: Characteristic absorption peaks in the NIR for silicon nitride compounds. (reproduced with permission from [2] ©1997 IEEE)

provements in some of these cases of 1-2 dB/cm, the trade-off was often extensive non-uniformity in thickness over a full 100 mm wafer, on the order of 100 nm center to edge. These numbers are specific to the PECVD tool at the Cornell NanoScale Facility, and could be improved in commercial fabrication. We include in Table 3.1 the details of the two final material recipes used in this work.

3.4 FABRICATION

The full process starting with a blank silicon wafer can be described in 9 independent steps. We begin with a 100 mm silicon wafer which is meant to replicate the back-end of an electronics CMOS stack, one in which we are not allowed access to post-fabrication. On this wafer we deposit 3 µm of PECVD SiO_2 , according to the details in Table 3.1, which will provide the bottom optical insulation for the waveguiding layer. Typically in an SOI stack this layer is grown thermally; however, we have neither the silicon real estate nor the tem-

perature freedom to do this. This is shown in Figure 3.3(a).

We deposit PECVD Si_3N_4 on this layer (Figure 3.3(b)), once again using the recipe given in Table 3.1. Using positive photoresist SPR 955CM 0.9 on an i-line stepper lithography tool, we pattern bus waveguides and laterally-coupled micro-ring resonators as shown in Figure 3.3(c). We term this layer "L1" as the first photonic layer on the stack. We clad this first optical layer with the same PECVD SiO_2 used for the BOX layer.

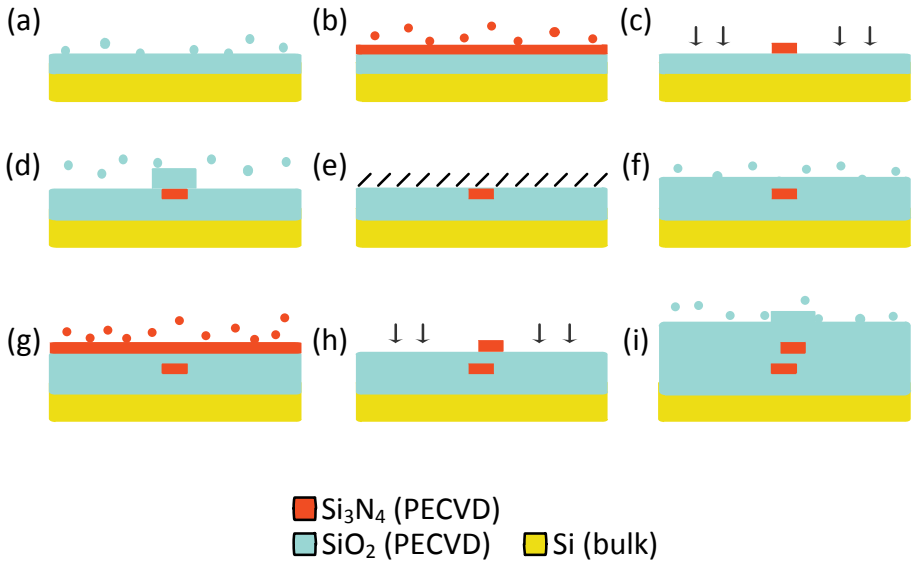

Figure 3.3: Fabrication of a 2-layer optical link: (a) deposit SiO_2 as BOX; (b) deposit Si_3N_4 for first guiding layer; (c) pattern/etch L1 waveguides and rings; (d) deposit SiO_2 for lateral buffer; (e) polish and planarize SiO_2 layer with CMP; (f) deposit SiO_2 as spacer; (g) deposit Si_3N_4 for second guiding layer; (h) pattern/etch L2 waveguides; (i) deposit SiO_2 as final cladding layer.

The conformal PECVD deposition process creates a ridge above all nitride features as seen in Figure 3.3(d). These "bumps" limit the ability to stack vertical layers, so the wafer is planarized to the waveguide level, or just above it using a chemical mechanical polishing (CMP) tool (Figure 3.3(e)). Two meth-

ods exist to do this. The first involves using a cesium oxide (Celexis CX94S) slurry, which selectively removes SiO_2 over Si_3N_4 in a process known as damascene. This method polishes one material and ideally stops when it reaches the second. However, to properly achieve this, the wafer requires high feature-uniformity and density. For example, large alignment marks found across the wafer can affect the performance of the polish. Because we expect the CMP to stop at the waveguide layer, the penalty for error is partial or full removal of the waveguiding layer. Instead, a KOH and silica slurry (Cabot SS12), which has a repeatable polishing rate for SiO_2 along with a spectral reflectance tool is used to polish and measure continuously. While slower, this process has a higher yield and is more controllable.

A second deposition of PECVD SiO_2 creates the vertical gap between the waveguiding layers, which is shown in Figure 3.3(f). Given the precision with which we can deposit new material, this step can be used to precisely set the gap between the two layers with very high precision. Adjusting the gap, and therefore the coupling by deposition gives us considerably more control than the best lateral coupling possible. Lateral coupling is dependent on resist, lithography, and etching limitations which are generally harder to control.

The second Si_3N_4 layer is then deposited above this last layer and patterned the same way as the first (Figure 3.3(g)-(h)). We term this layer "L2", the second photonic layer. Finally we clad again in PECVD SiO_2 . This last layer is again 3 µm thick and is used as a top isolation layer (Figure 3.3(i)).

3.5 Device and Stucture Design

We fabricated a multi-layer optical link, traversing two layers through one passive wavelength switch. The first layer (L1) was patterned with waveguides

and rings of 30 µm radius. L1 spans the full die starting and ending in inverse-tapered couplers on both ends of the chip to provide an input and through port. L2, the second photonic layer begins above the micro-ring resonator and spans to the output facet where light can be coupled out to read the drop port (see Figure 3.4(a)). The coupling of L2 waveguides is set by the vertical and horizontal offsets provided by the 800 nm mid-layer SiO_2 02 and lithographic positioning shown in Figure 3.4(b). Figures 3.4(c) and (d) show false-color SEM images of the cross section of the chip with emphasis on the L1 waveguide.

3.6 RESULTS

We tested our devices using a standard set metrics. We were interested to see how well our waveguides guide light given the low-temperature materials we used. This directly affects our wavelength resonators since it changes the quality factor as well as the coupling conditions if the losses are too high. Furthermore, since we are advocating 3D integration, we want to see what losses we can expect from vertically overlapping waveguides and if this is an improvement from intersecting waveguides. We also test high-speed data transmission through these devices to confirm error-free operation.

3.6.1 CHARACTERIZATION

Losses on the low-temperature silicon nitride waveguides were tested using the cut-back method whereby a set of equal geometry waveguides were patterned with varying lengths so as to isolate the optical losses generated by propagation. Using a tapered fiber input to the chip facet with a tunable NIR laser source, we measure at output with a collimating objective and power detector. We show

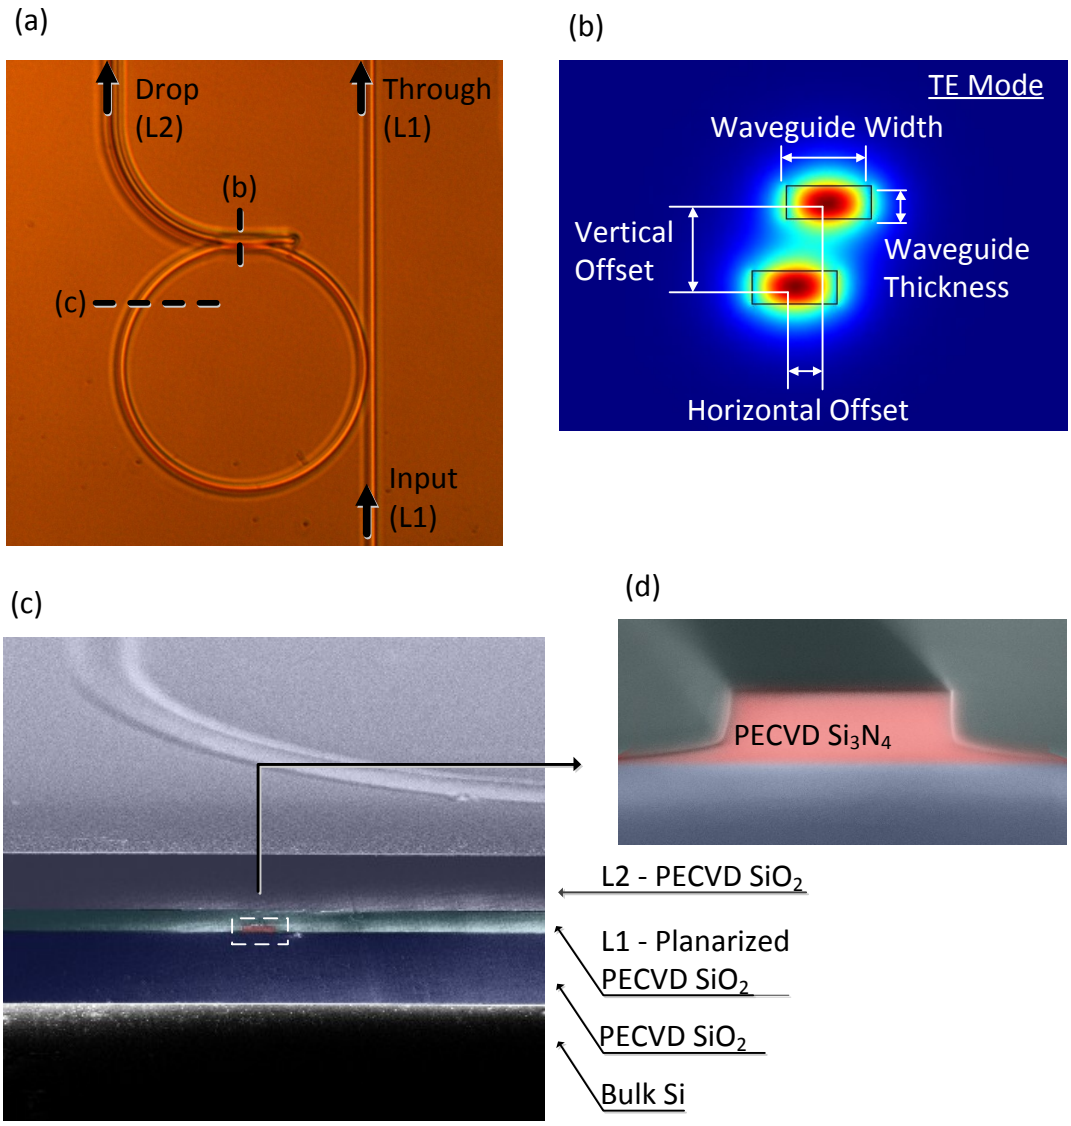

Figure 3.4: Microscope image of bus waveguides and microring resonator; (b) vertically coupled structure with TE mode and geometry parameters; (c) false color SEM image of ring cross section; and (d) close-up SEM of silicon nitride waveguide cross section.

the waveguide loss in Figure 3.5(a) to be just over 1 dB/cm for most of the L-band in our 400 nm x 1 µm geometry. The loss increases steadily into the C-band due to Si-H and N-H bonds in the film [2]. These losses are comparable to those shown in the literature for silicon nitride and are still better than most comparable silicon single-mode waveguides [58].

We then measure losses for the waveguide crossings using a similar approach. We fabrticate many waveguides with increasing number of over-layer crossings to isolate the loss generated by this perturbation. We measure the number to be -0.04 ± 0.002 dB/cross for waveguides separated vertically by 800 nm for the full range. The incrementing overlapping waveguides in L2 corresponded to a maximum of 25 crossings over L1. Figure 3.5(b) shows this loss as compared to a 3D finite-difference time-domain (FDTD) simulation using perfect materials, estimating the ideal value around -0.015 dB/cross for this gap. Additional loss is to be expected from impurities in the PECVD. This number is a major improvement over single-layer silicon crossings, which in theory can only approach this number [62] and could still be improved with increased vertical layers by separating the waveguides further. One design example could be two waveguide bus layers with a switching layer, using rings or directional couplers in between, separating the bus waveguides by more than 2x the coupling gap.

Lastly, we test the optical path through the vertically-coupled filter and measure a channel bandwidth of 25 GHz and an extinction ratio of -24 dB. Figure 3.5(c) shows both the drop port and through port normalized by the input power to the device. We also measure an insertion loss of -0.6 dB from the drop port response, and we calculate the loaded quality factor of the resonator from the linewidth to be $Q_l = 7.5 \times 10^3$. Using temporal coupled-mode theory [63],

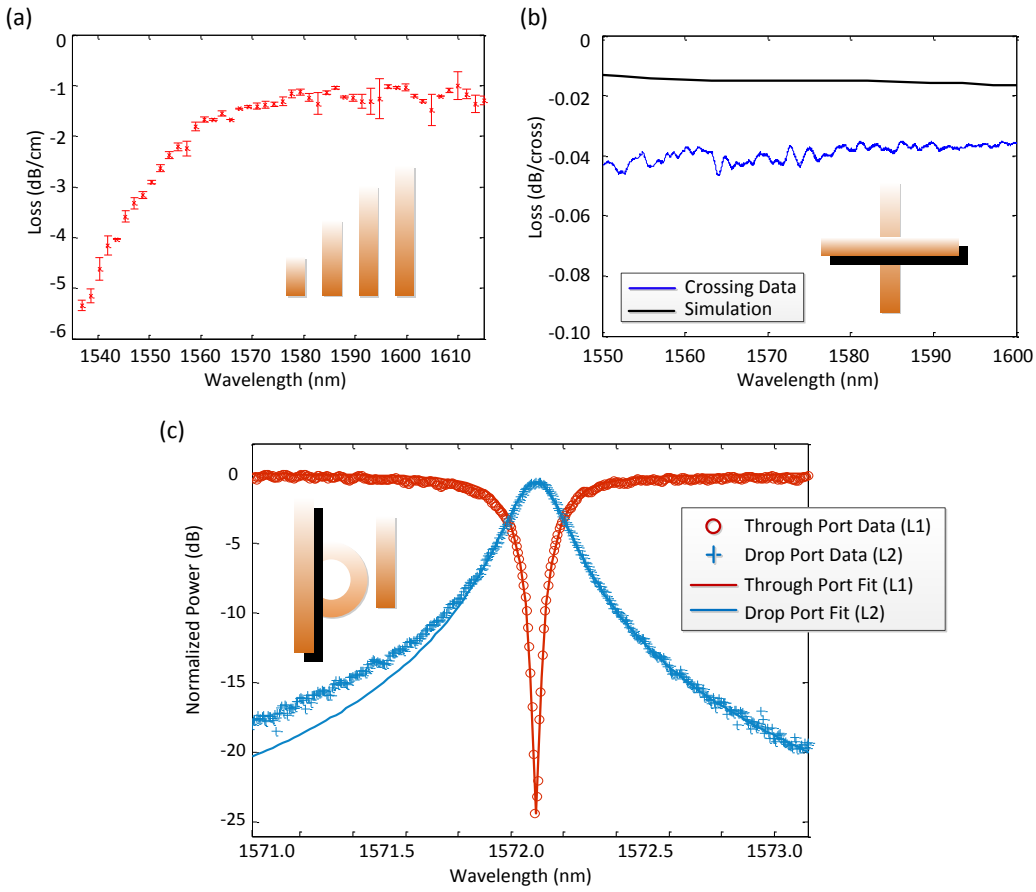

Figure 3.5: Results of fabrication: (a) propagation losses of PECVD Si_3N_4 waveguides (400 nm x 1000 nm) over wavelength of interest; (b) averaged loss per vertical crossing (800 nm separation); (c) microring resonator (30 µm radius) through and drop port responses from both layers.

we fit the experimental data and calculate an intrinsic quality factor greater than 105 for this resonator. We further validate our numbers through simulation using finite-element method (FEM) including material and bending losses from which we estimate an ideal intrinsic value of $Q_{i-sim} = 1.7 \times 10^5$. The spectral width is engineered for on-chip network applications in the L-Band while maintaining the critical coupling condition for high filter efficiency. These rings are

considered mostly radiation limited and can be tailored to specific bandwidth values through proper coupling design.

3.6.2 HIGH-SPEED DATA TEST

Following our characterization, the devices were tested at Columbia University for high-speed data transfer. Using again a tunable laser source, but this time amplitude-modulated at several speeds, we test the power penalty at low BER rates as well as explore the eye diagrams from these data links.

Figure 3.6 shows a diagram of the experimental setup used to test data transmission through the multi-layered device. A NIR source is modulated using a non-return-to-zero (NRZ) signal, encoded with a $2^{31} - 1$ pseudo-random bit sequence (PRBS) by a pulse pattern generator (PPG). The signal is set to the quasi-TE polarization and coupled using a tapered fiber. The output is collected through another fiber and converted back to the electrical domain with a PIN photodiode, a transimpedance amplifier (PIN-TIA), and a limiting amplifier (LA). The signal is evaluated using a bit error rate tester (BERT) and digital communications analyser (DCA).

As we are using a passive-wavelength switch, light is directed between outputs using the tunable laser source. In a real implementation there would be different signals encoded into each of the wavelengths and would be routed accordingly. Setting the source to $\lambda = 1568.5$nm, the data encoding speed is varied to explore its effect on the resulting signal quality. The ring does not resonate at this wavelength and so ignores it, thereby letting the light pass through to the so-called "through port" on the first level. Setting $\lambda = 1565.7$nm we are now in the resonance condition of the ring and light switches direction; it is now coupled through the ring into the top layer waveguide, to the "drop port".

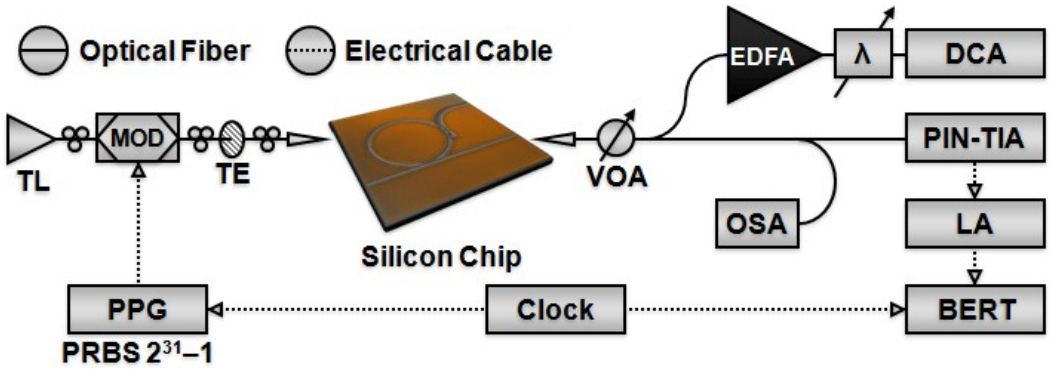

Figure 3.6: High-speed data experimental setup including modulated optical signal and bit error rate tester.

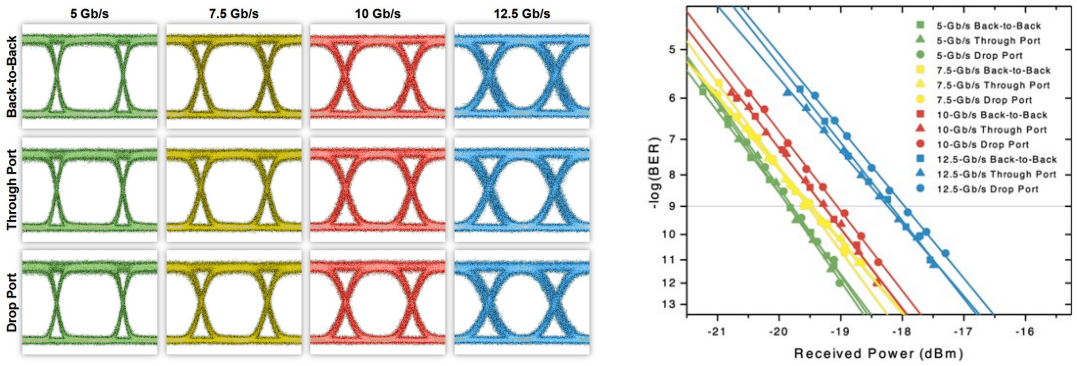

(a) (b)

Figure 3.7: Measured results through 3D switch at both output ports and back-to-back comparisons: (a) eye diagrams from optical signals with 5, 7.5, 10, and 12.5 Gb/s data rates; and (b) BER curves and fits from each of the specified data rates.

Figure 3.7(a) shows the eye diagrams for each of the speeds tested, at 5, 7.5, 10 and 12.5 Gb/s, all quite open and seemingly unaffected by the device. Figure 3.7(b) quantifies the penalty as compared to measurements through the system without the chip but using a variable optical attenuator (VOA) to mimic the fiber-to-fiber loss of the chip (back-to-back). The through port shows error free operation, defined as having a BER less than 10^{-12}. The drop shows negligible power penalties for speeds under and up to 7.5 Gb/s, and 0.22 and 0.23 dB for 10 and 12.5 Gb/s rates, respectively. This rate dependence points to a limitation in channel size, and could be likely solved with a spectrally larger ring resonator design.

3.7 Conclusion

We have demonstrated here a process for integrating multiple photonic layers on the backend of CMOS electronics. Using low-temperature PECVD Si_3N_4 as a guiding layer and PECVD SiO_2 cladding, we show multiple stacked optical layers, coupled by micro-ring resonators, all with i-line photolithography. Furthermore, we show that a major deterrent of optical networks, single-layer waveguide crossing losses, can be avoided using multiple-layer photonics. Data streams up to 12.5 Gb/s were switched through the 3D optical device with error-free operation being shown up at 7.5 Gb/s and minimal loss above that. This method for fabricating on-chip optics opens the doors to a new area of optical network design previously limited by real estate and WDM channel limitations. All processes and materials used here are compatible and common to the standard electronics CMOS process.

CHAPTER 4

STABLE TUNING AND FABRICATION CORRECTION OF RESONATING STRUCTURES

4.1 INTRODUCTION

Resonating structures have the clear benefit of producing sharp spectral features without electrically active elements. The more pronounced these features become, the more material and physical effects we see acting on them. These effects can be positive or negative depending on the application we are targeting. The first productive use that comes to mind when we think of spectral tuning, for example, is that of light modulation. Simply turning on and off a light can be universally understood as a medium for communicating information. Modulation is achieved in resonating structures by shifting a spectral passband back and forth over the wavelength of light used for transmission. The most common ways to do this include using the All-Optic (AO) effect [64], the Thermo-Optic (TO) effect [65], and the Electro-Optic (EO) effect [31] which in silicon is the name given to carrier injection. It can be understood that whatever the effect, the process acts by changing the refractive index of the waveguide as a whole (effective index). In a resonating structure this has the effect of changing the effective path length which in turn changes the resonance condition. For waveguiding rings the resonance condition is defined by Equation (4.1).

$$m\lambda_{resonance} = L_{effective} = 2\pi R_{ring} n_{effective} \tag{4.1}$$

Many demonstrations have shown fast modulation speeds using AO and EO in the range of 10's of Gb/s, including a recent showing of an EO 40 Gb/s modulator [66]. However there are many possible problems with making real devices out of any of these technologies, not least of which are the competing effects. While light-induced temperature change can be easily prevented, conduction-based external temperature is a common variable found in electronics and especially microprocessor designs. We will see that while cooling options exist the

easiest solution will be to fight heat with heat. Additionally there is the pervasive issue of fabrication. The physical limitations of fabricating these devices will also cause spectral changes which need to be corrected. We will discuss these issues here.

4.2 Fabrication Limits

The practical use of ring resonators in mass-production environments is not a trivial problem. While it is clear that the resonating properties of a microring resonator depend on the dimensions of the structure, it is not commonly acknowledged how sensitive the dimensions are to the proper operation of these devices. In fact allowing for one of the most advanced fabrication tools available now, an electron beam (e-beam) lithography tool, the dimensions we commonly use are still not consistently reproducible. Figure 4.1(a) shows a common silicon ring resonator, using typical single-mode dimensions, and made using this technology. While there are ways to alleviate the visible sidewall roughness the precision of the geometry width at this scale is already lost.

Figure 4.1(b) illustrates how the spectral features for a microring would change with a few nanometers of variation by Equation (4.1). In comparison, we show the limitations of an e-beam tool, specifically the JEOL 9300FS available at the Cornell NanoScale Facility at the time of this writing. While still very precise at only a few nanometers, the spectral feature of a modest $Q = 10,000$ ring resonator at Full Width Half-Maximum (FWHM) is around 0.15 nm. Which is to say that if the beam is displaced by only a quarter of its size the spectral feature might shift more than 6 times its linewidth. Likewise Figure 4.1(c) shows what the displacement might mean in a coupling scenario. Unfortunately foundries do not actually use e-beam due to cost, but rather less accurate photolithogra-

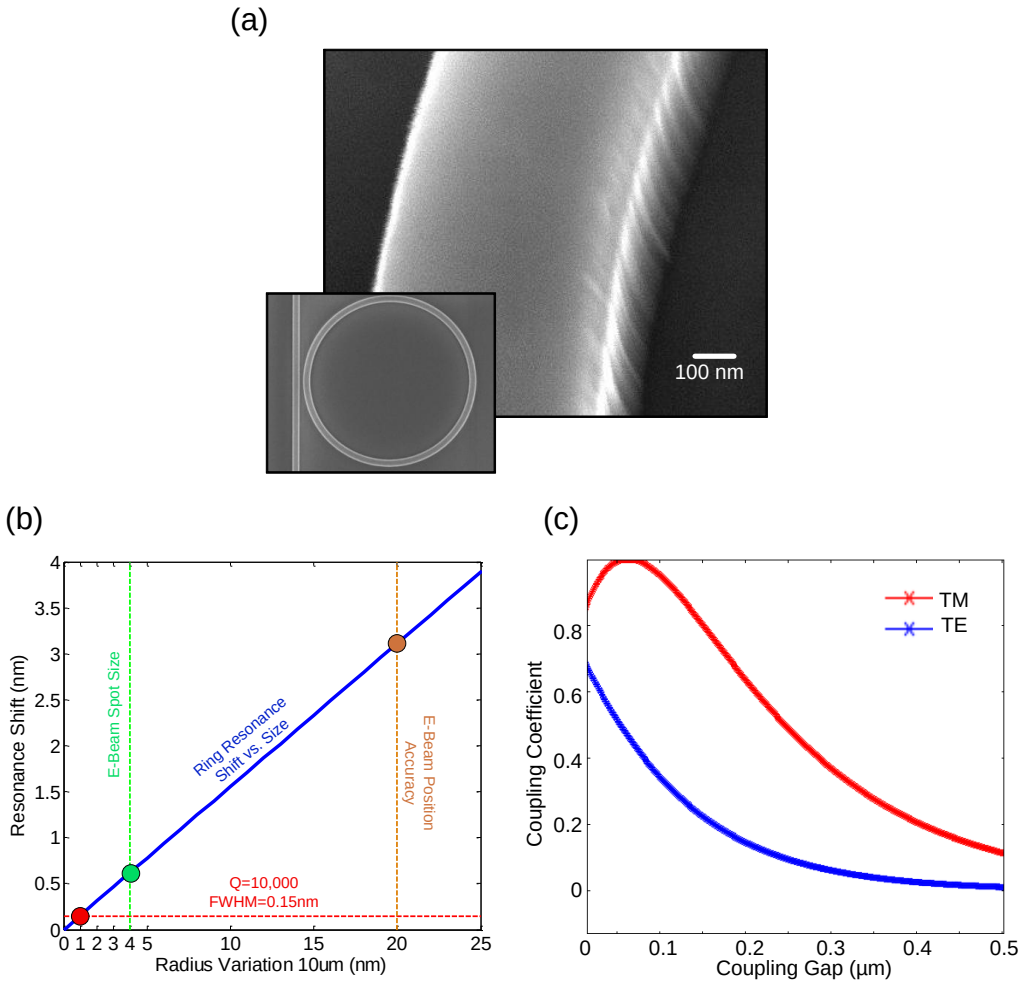

Figure 4.1: Fabrication limitations for ring resonators: (a) closeup of waveguide sidewall fabricated using electron-beam lithography; (b) ring resonator spectral shift versus change in resonator radius; and (c) coupling variation given a change in gap spacing.

phy tools where the tolerances are higher.

Very few methods exist for post-fabrication correction. They include laser correction, either with photosensitive cladding [67] or direct manipulation of the material [68], as well as direct oxidation of the material [69]. All of these processes of course must be done during the fabrication process, though the benefit is that the trimming is done only once. Instead we can do our tuning

post-fabrication at the cost of active power on the device.

4.3 On-Chip Temperature Variations

A commonly cited problem with silicon photonics is that of temperature sensitivity [70, 71]. Even within regular CMOS electronics the power dissipation has become a critical problem, limiting the achievable performance of these chips. The issue is that reducing gate lengths to increase performance results in increased source-drain leakage, which in turn drives power consumption up. The increased power then dissipates increased heat [72, 71]. Of course the power used and heat produced are variables, both in space as well as in time [73]. Spatially the different elements of a processor consume more or less power, e.g. an arithmetic logic unit versus memory. Temporally the processor will use more or less power depending on the kind of processing being done, for example FEM simulation versus web surfing. Figure 4.2 shows an example of a microprocessor heat map, which can have hotspots in different locations and can vary over time.

The unpredictability of the heat distribution over the surface of a microprocessor is a major concern for the on-chip optics. The variations can be between $\Delta T = 30$°C to 50°C on a typical processor [72, 73]. The TO coefficient of silicon however induces a change in refractive index of silicon per degree Kelvin given by [74]:

$$\left(\frac{\mathrm{d}n}{\mathrm{d}T}\right)_{Si@1.5\mu\mathrm{m}} = 1.86 \times 10^{-4}\mathrm{K}^{-1} \tag{4.2}$$

Such a change in index can vary the resonance feature of a resonator considerably, roughly about 1 nm per 10 °C . Using the TO coefficient we can roughly see

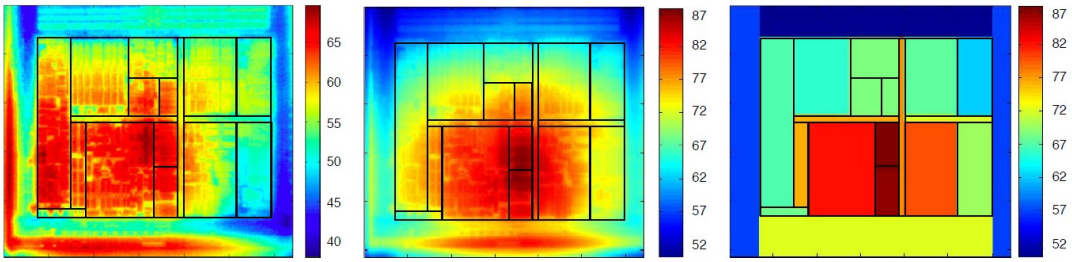

Figure 4.2: Thermal hotspots appear on a microprocessor with spatial and temporal variation: (a) full-thermal image with floorplan; (b) hottest captured image; (c) and its average temperature per block (reproduced with permission from [3] ©2008 IEEE).

how the temperature relates to the effective index in Silicon and in turn shifts the resonant wavelength of a ring. The graph is shown in Figure 4.3.

While architects work on stabilizing hotspots though, device engineers have also found a few ways to control variation. Mostly the techniques rely on changing the effective TO coefficient of a device either by changing the cladding of a waveguide [75] or altering the mode profile to move the field to the cladding and out of the core [76].

There is an alternative here, which is to use heaters to control temperature. This may seem counter-intuitive at first but if a designer knows the amount of ΔT variation that exists in a processor she could in principle bias the heaters by that amount, and cool or heat as necessary. Metal heaters have been shown to tune at least over 20 nm of spectrum [77]. With proper design they could likely shift more to handle a larger variation, at the cost of more power.

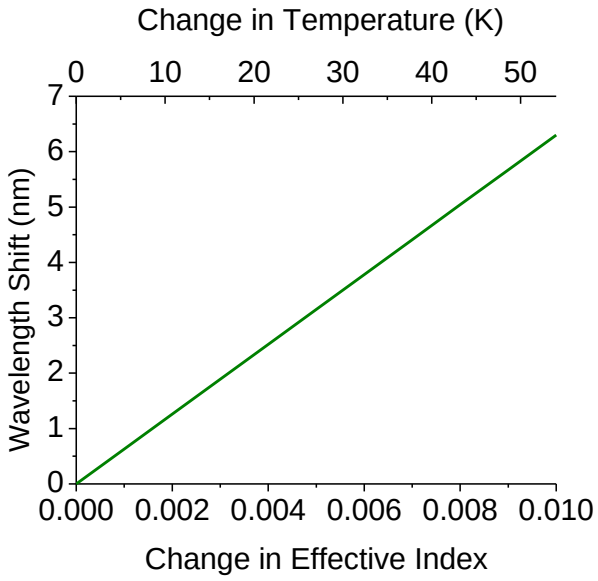

Figure 4.3: The resonance wavelength shift in the microrings that is caused by the shift in effective index and equivalent temperature shift.

4.4 Generating Localized Heat

Heating of the waveguide can be done in a number of ways including radiation, convection and conduction, though of these the last is the most convenient and efficient at these scales. Furthermore of all the methods to generate heat we would like to use an electric source, given that we plan to integrate our devices with electronics. In fact all metals and even doped semiconductors can be used as ohmic or resistive heaters. By passing a current through a wire, majority carriers are accelerated and will often collide with ions in their path converting some of their energy through vibrations into some amount of heat, proportional to the square of the current multiplied by the resistance [78] as shown in Equation (4.3).

Table 4.1: Resistivity of Materials

Material	$\rho(\Omega\,\mathrm{m})$	$\alpha[(^{\circ}\mathrm{C})^{-1}]$
Aluminum	2.75×10^{-8}	0.0039
Copper	1.72×10^{-8}	0.00393
NiChrome	100×10^{-8}	0.0004
Tungsten	5.25×10^{-8}	0.0045

$$Q \propto I^2 \cdot R \tag{4.3}$$

This simple equation however has a double edge to it, at least with regards to metals. Increased heat has a direct effect on certain materials, specifically with regards to its resistivity. The more vibrations are generated through collisions the more collisions can occur, further increasing the resistance of the material which in turn generates more heat. Catastrophic failure can easily ensue from such a scenario. This means that we need to be careful to find a proper material for ohmic heating. The resistivity change in metallic conductors can be approximated using Equation (4.4):

$$\rho(T) = \rho_0[1 + \alpha(T - T_0)] \tag{4.4}$$

where ρ_0 is the resistivity at a reference temperature T_0, and $\rho(T)$ is the temperature-dependent resistivity. α is the temperature coefficient of resistivity. Table 4.1 lists a few common metals along with their respective values. We list among them our preferred alloy NiChrome (80/20 wt.%) which has an order of magnitude smaller α, making this material less responsive to temperature and more convenient as a heater [79].

4.4.1 LOCALIZING HEAT

The configuration of the rings and heaters is determined so that the optical mode is well separated from the heater to avoid absorption while still close enough in proximity to couple heat into the waveguide. We find the best arrangement of heater and waveguide using 3D Finite Element Modelling (FEM) as shown in Figure 4.4. Because our heater is created from the deposition of a thin film of metal, most of the heat will be conducted from the largest surfaces above and below the heater. To conduct the majority of heat efficiently into the ring we trace the heater wire above the ring resonator. The shape that is traced is similar to the Ω symbol, since we need to contract a positive and negative lead at point in the ring [80].

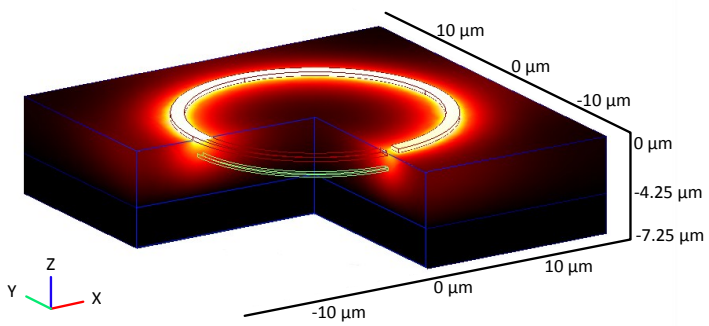

Figure 4.4: Finite Element Modelling of microring resonator with vertical localized thermal source.

We take into consideration the size of the optical mode for our waveguide dimensions as well as the heat conduction of our heaters through the cladding material. We make our waveguides 450 nm wide by 250 nm thick. From our Finite-Difference (FD) mode solver we determine that for the TE (polarized parallel to the substrate) mode, 1 µm of distance is sufficient to limit the metal absorption to $a \approx -2.17 \times 10^{-4}\,\mathrm{cm}^{-1}$ (shown in Figure 4.5(b)).

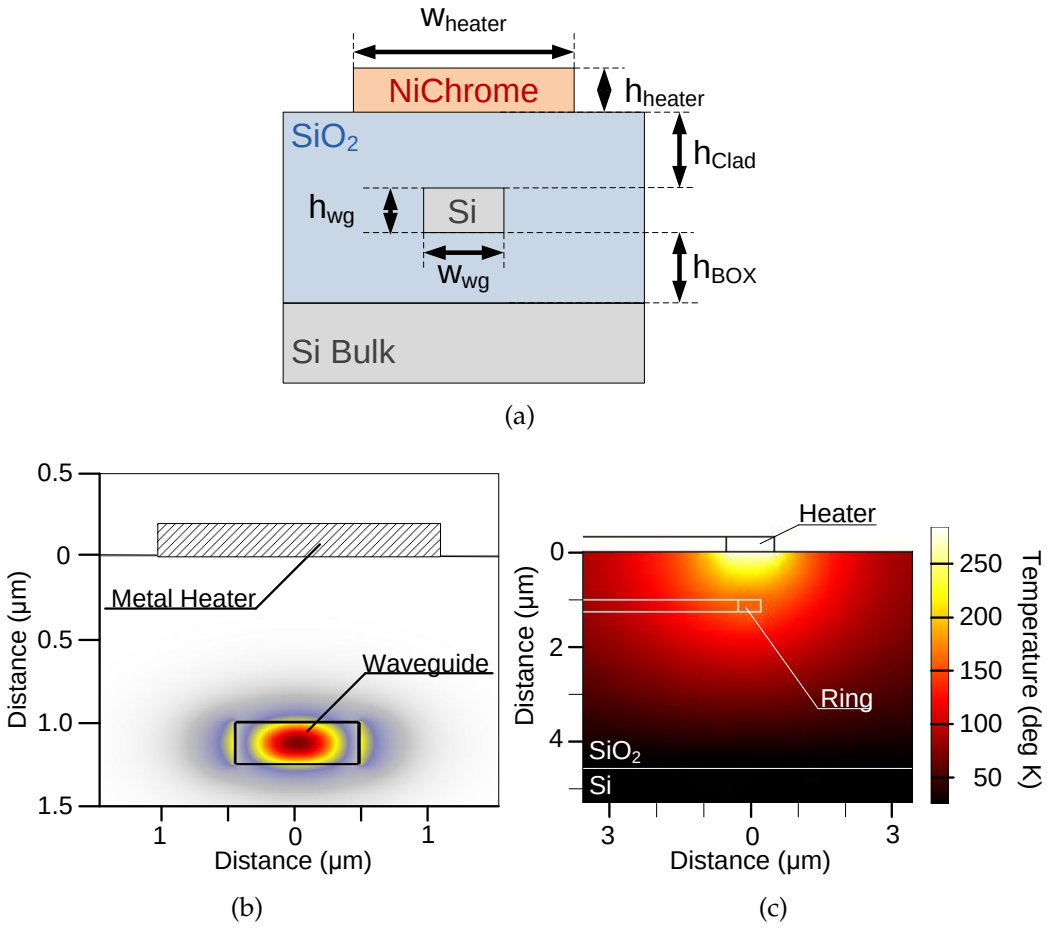

Figure 4.5: Heater and waveguide design for thermal and optical efficiency: (a) variable positioning relative to waveguide and dimensions; (b) the optical mode is separated from the heater to avoid absorption while (c) maintaining enough proximity to couple heat into the waveguide.

Confinement of heat is important not only to save power but also because in our designs we might have adjacent rings with potential crosstalk issues. Thermally tuned microrings which are physically close would experience thermal crosstalk if the heaters were designed too wide or too close to each other. Using once again our 3D FEM tool we explored the most efficient way of conducting heat to our rings, while keeping the heaters small. We keep in mind that thermal coefficients of silicon and silica are two orders of magnitude apart, which

means that we will conduct poorly from the heater to the ring, though we will confine well locally:

$$\kappa_{Si} = 156 \frac{\mathrm{W}}{\mathrm{m\,K}}$$
$$\kappa_{SiO_2} = 1.4 \frac{\mathrm{W}}{\mathrm{m\,K}}$$

Our final design settles on 1 µm wide wires directly above the waveguide ring in order to get a uniform, but laterally contained, heat distribution. We see in Figure 4.5(c) that at 1 µm distance we can still conduct more than 56% of the temperature generated by the heater into the waveguide. Locally we still confine the heat in the lateral directions to about 10 µm from the heat source with minimal effect. Figure 4.6(a) shows two side by side ring structures (clad so that the rings are not visible), one of which has a 100 °C heater operating above it. Figure 4.6(b) shows the cross-sectional decay of the heat with the overlaid ring structures. At 5 µm the heat decays to about 13% and at 10 µm it is at 1%.

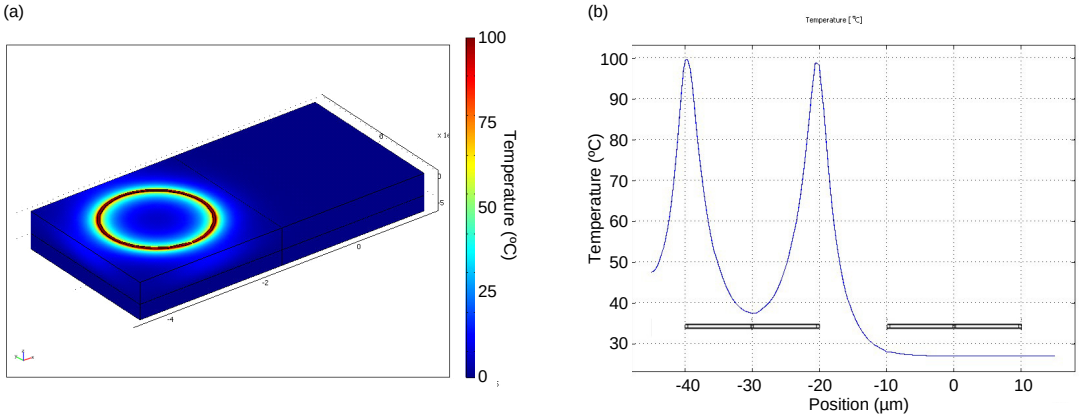

Figure 4.6: Simulated thermal crosstalk: (a) 3D FEM structure encapsulating two 10 µm radius ring resonators with one active heater; and (b) center cross-sectional view of the thermal decay showing temperature as a function of distance.

4.5 THERMAL SPEED

While we know that speed is not the purpose of this device, rather tuning and stabilizing, we make a small note here regarding the limitations of a thermal switch. Most experimental demonstration have shown thermal switching on the order of several µs [81, 82]. One group was able to show just under 1 µs using a mach-zhender and a clever push pull configuration [83]. However the problem is in fact fundamental due to the diffusion limitations of the materials: even if one achieves fast thermal injection the time it takes to release the heat is still in the µs range. Fischer et al. provide a good equation for finding the frequency cutoff of such devices [84]:

$$f_{cutoff} = \frac{1}{\pi\lambda\rho\epsilon_{th}} \left(\frac{P_\pi}{A}\right) \left(\frac{\mathrm{d}n}{\mathrm{d}T}\right)_{Si} \tag{4.5}$$

where ρ in this case is the density and ϵ_{th} is the specific heat of the heater material, and A is the cross-sectional area of that material. P_π is the switching power and λ of course is the wavelength.

We verity the speeds reported with a simulation using again our 3D FEM tool. We first pulse our electrical heater for 1 µs and measure the temperature response both at the heater and the silicon ring. The result shown in Figure 4.7(a) is characteristic and shows the µs response time. We do notice at least that the turn-on response time seems to be dependent on the amount of thermal power injected so in Figure 4.7(b) we show a simulation where we pulse at higher temperature for 0.25 µs and quickly drop the power to where we want to maintain the heater. The result is in fact a few hundreds of nanoseconds response, however with once again the diffusion-limited turn-off time of a few microseconds. In short without a cooling solution thermal speeds are limited to this range.

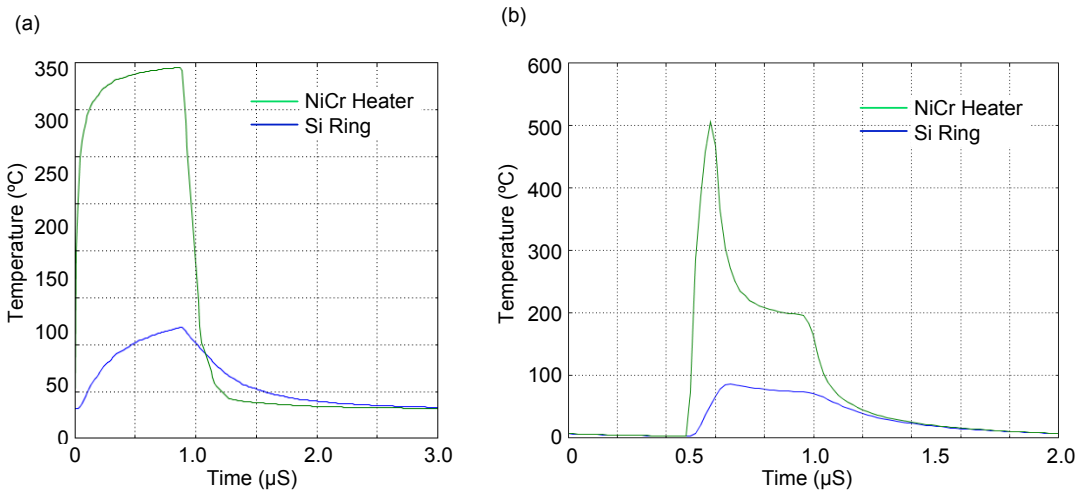

Figure 4.7: Heating speed of SOI waveguides is limited by diffusion: (a) FEM simulation showing on/off limits; and (b) Pre-pulising effects on speed.

4.6 MICROHEATER FABRICATION

Fabricating heaters requires a straight-forward process that is compatible with regular SOI photonic processing. The key to the design is that we make these heaters above the photonic regions after they have been fully processed, which creates minimal disturbance to the sensitive photonic regions. The process we describe is shown graphically in Figure 4.8.

We begin with a regular SOI wafer with a 250 nm silicon layer. We cover the wafer with negative e-beam resist, in this case XR-1541, a hydrogen silsesquioxane (HSQ). This resist becomes very similar to SiO_2 when exposed and developed, such that we can leave it above our waveguides where it will optically act like the rest of our cladding. The resist is exposed with our waveguide pattern, 450 nm lines which define our single mode geometry, using our electron beam patterning tool. The pattern is developed using 300 MIF.

The pattern is etched intothe silicon using an anisotropic chlorine etch. High

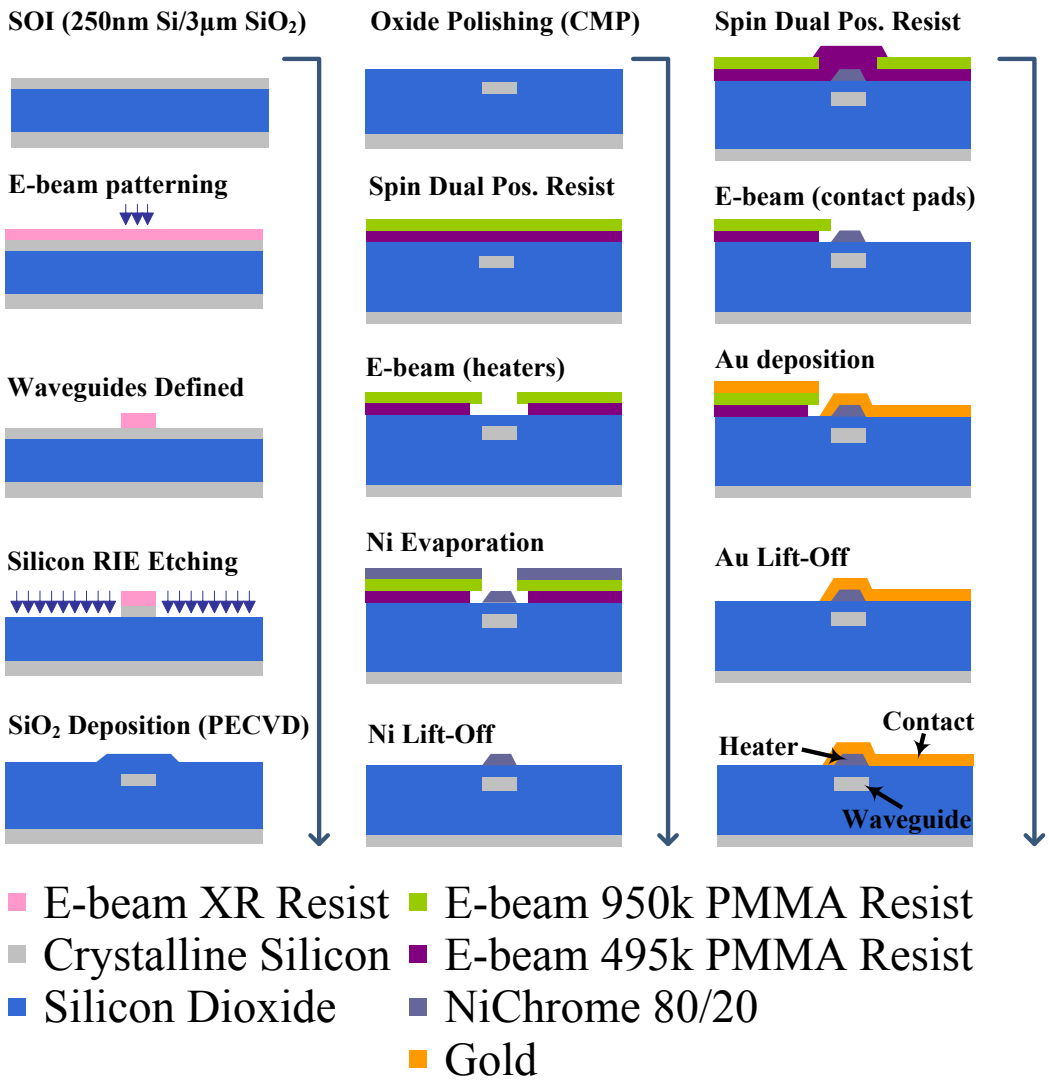

Figure 4.8: Fabrication steps for making ohmic heaters.

verticality of the sidewalls is important to maintain our designed geometry, which affects not only our optical mode but more importantly our coupling between waveguides. We use an inductively coupled plasma (ICP) reactive ion etching (RIE) tool to realize this. There are some limitations to this process including an apparent non-linear etch rate due to gas-lighting steps, but these can be calibrated in with some testing. The resist is removed using some form of oxygen plasma or appropriate chemical remover.

SiO_2 is deposited above the waveguide layer using plasma-enhanced chemical vapor deposition (PECVD). Once again we see that this process produces a conformal film which is "bumpy" on the waveguide regions. As was discussed in section 3.4 we have the option to planarize the film using chemical mechanical polishing (CMP). In this case the benefit lies in avoiding heater failure points that might arise during metal deposition, though we will discuss a method for avoiding these. Since the heaters are localized strictly above the waveguide layer we only need to concern ourselves with the layer being the appropriate thickness of 1 µm above this region for proper heat conduction and optical mode isolation.

To define our heaters we can use e-beam or photolithography given the larger dimensions of the heater features (defined in Table 4.2). In order to fabricate metal regions it is common to use a lift-off method, rather than etching, since many of the chemical necessary have undesirable effects on our substrate layers. To do this we define an inverse sacrificial layer using positive resist. This step exposes the areas to be lifted off the wafer [56]. A single layer lift-off has several issues that we need to be concerned with however. Because the resist will be removed by solvent it needs to be accesible at points around the wafer, typically at the edges of the defined regions. If the edges are covered in

Table 4.2: Device Parameters for Ohmic Heaters

Parameter	Symbol	Value
Heater Width	w_{heater}	1µm
Heater Thickness	h_{heater}	200nm
Heater Clad Gap	h_{clad}	1µm
Waveguide Width	w_{wg}	450nm
Waveguide Thickness	h_{wg}	250nm
Buried Oxide Layer Thickness	h_{BOX}	3µm

deposited metal, then the liftoff process might fail. Also edges may collect additional metal during deposition creating what are known as "ears" which then might tip over and create undesired features. Shearing of the metal may also occur during lift-off from having continuous metal regions between resist and defined area.

All of these problems are resolved using a two layer lift-off process. In essence we have two resist layers which are exposed in the same way but develop at different rates (or vice versa). If the bottom layer is removed faster we create a cantilever structure (as seen in Figure 4.8) through which metal can be deposited without the problems of having resist directly adjacent to the defined area. In our case we use polymethyl methacrylate (PMMA) for our resists, a positive e-beam polymer resist with differing molecular weights [85]. 950 K PMMA is used above for its higher resolution: this layer will define the heater areas. Underneath (but spun first) a layer of 495 K PMMA is used which will create an undercut when developed. We recommend the thickness of this layer to be 2X the layer of metal being deposited. After exposure the developing of this resist is done using a ratio of methyl isobutyl ketone to isopropanol alcohol (MIBK:IPA) of 1:3 to achieve best resolution.

The metal, in this case NiChrome, is deposited using an electron beam evap-

oration tool. For proper adhesion to the dielectric stack 5 nm of titanium is evaporated first. 200 nm of NiCr is then deposited. In the case where CMP is not done, a vertical evaporation of metal would result in thinner metal films and even disconnected regions where large steps are encountered, for example at the edge of the waveguide regions. These spots change the cross-sectional area of the heater which in turn changed the resistance at that point. These are the most likely failure points for these types of heaters [80]. We can avoid this problem by using an angled spinner inside the evaporation tool. This will ensure a conformal and uniform film deposition even on the sidewalls of the raised oxide features. An example of such a film is shown in Figure 4.9. The wafer is then left in acetone to remove the resist and lift-off the metal above it. The bilayer resist process is then repeated again for the contact pads using gold as the deposited metal. This secondary metal has considerably lower resistance and results in better wires and contact pads.

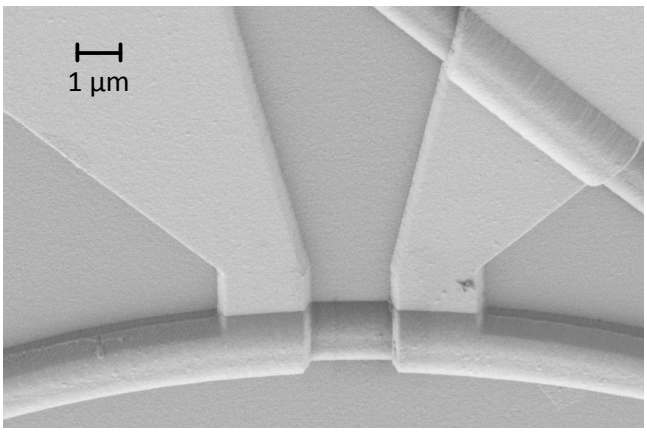

Figure 4.9: Closeup SEM of evaporated NiCr film using angled rotation platform, and bilayer lift-off.

Figure 4.10(a) shows an SEM of a finished heater device over a 10 µm radius ring resonator. The NiCr heater curves over the surface of the ring and

flares out to a large region where gold overlaps and extents to wirebond points. The sharp corners of the flaring region are filleted to avoid hotspots. Figure 4.10(b) shows through-port measurements of the heated microring along with consumed power. Using a DC source as input we can already see that retuning a fabricated ring can correct for fabrication limitations by adjusting the loss-coupling ratio, which leads to a high extinction ratio.

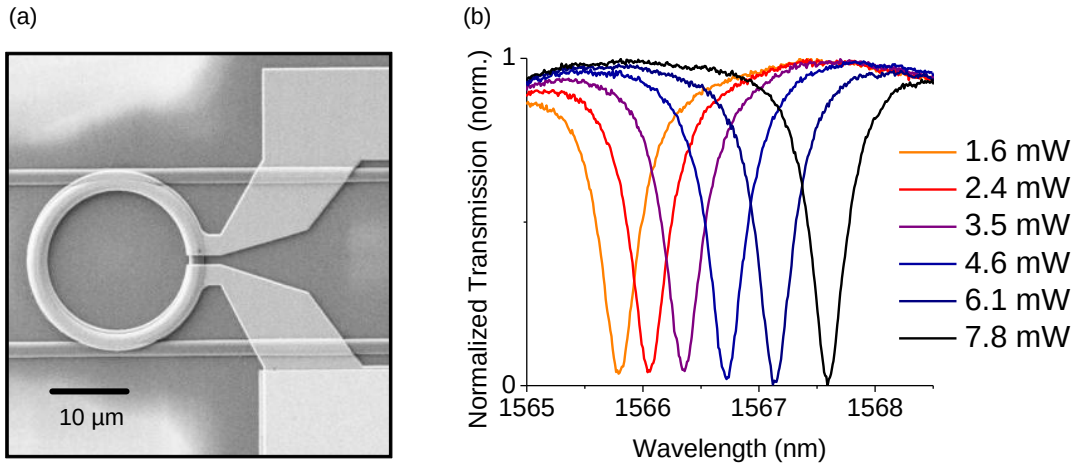

Figure 4.10: A simple Ω heater above a 10 µm radius ring resonator: (a) SEM of the device and (b) its electrical performance.

4.7 Conclusion

As we have seen the thermo-optic effect of silicon can be incorporated very easily on processed photonics to alleviate a variety of problems. Most notably the issue of fabrication limitations, which sees no likely end in the near future, can be corrected using individually biased heaters above sensitive resonator devices. Actively correcting for temperature can also be done in the cases of thermally variant electronic chips. Beyond this the microheater can be used as a

tuning device, capable of dynamically switching optical signals when used with waveguide-coupled wavelength filters. This is all achieved while maintaining a stable state over time, unlike other optical effects in silicon.

CHAPTER 5

SILICON PHOTONIC NETWORKS AND ROUTING DEVICES

This chapter uses material extracted with permission from the following references:

N. Sherwood-Droz, H. Wang, L. Chen, B. Lee, A. Biberman, K. Bergman, and M. Lipson, "Optical 4x4 hitless slicon router for optical networks-on-chip (NoC)," Opt. Express 16, 15915-15922 (2008).

B. Lee, A. Biberman, K. Bergman, N. Sherwood-Droz, and M. Lipson, "Multi-Wavelength Message Routing in a Non-Blocking Four-Port Bidirectional Switch Fabric for Silicon Photonic Networks-on-Chip," in Optical Fiber Communication Conference, OSA Technical Digest (CD) (Optical Society of America, 2009), paper OMJ4.

5.1 INTRODUCTION

An interconnection network does not necessarily have the same requirements as other types of networks. It will however, be described ultimately by a graph similar to that of Figure 5.1. The graph shows the average latency of a packet traveling through the network, versus the amount of data being put into the network [4]. At its best the packet will take some minimum amount of time, referred to as the "zero-load latency" T_0 to traverse the distance from start to end. When more and more packets show up, contention at various points in the network will slow down the average time of all packets. At some point if too much data is injected the latency goes to infinity at some saturation throughput. Where that happens can vary. For example if there is a bottleneck somewhere, a hotspot can develop which limits all traffic, like an accident on a main exit bridge in a city. The ideal throughput though, Θ_I, is defined by the physical topology of the network, and that topology is limited by its routers.

One of the most important elements in any networked system is the router. It provides the ability to dynamically select a destination for any input source of information in an interconnected group of nodes. Similar to having exits on a highway, or street intersections, data on a network have the ability to change direction to reach a specific location. As anyone who has ever been in a city traffic jam knows, moving large amounts of people (or data) is no easy task. Software algorithms that manage and distribute data could do wonders for traffic in New York City, like they do with the data of billions of people on the Internet, yet it is the physical foundation that ultimately limits routing capabilities.

In optical networks there are two ways to move light: passively and actively. Passive routing is performed by routing signals based on their wavelengths [86, 87, 88, 89, 90]. The traveling analogy to this design is that of the train system.

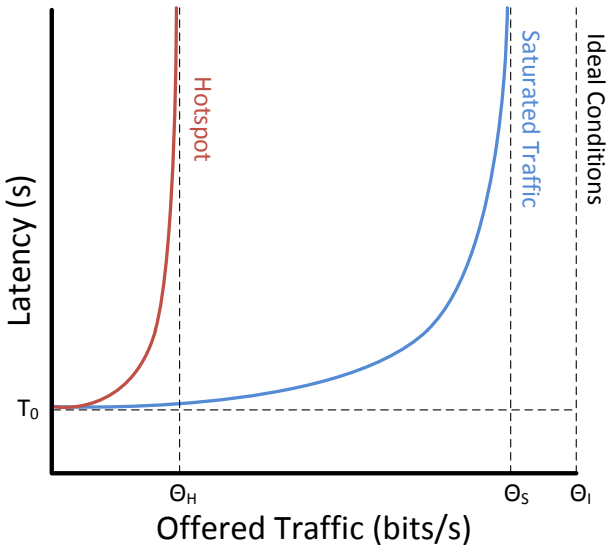

Figure 5.1: Latency versus offered traffic curves for a typical interconnection network. Latency of packets in a network is fundamentally limited by the amount of traffic being injected; however the network capacity depends more on the structure of the network itself [4].

You choose a train at the beginning of the journey which will take you directly to the destination. It is efficient in certain environments but it requires as many tracks as you have routes. In our terms, it means that you need as many light sources as the network has unique paths. This leads to problems in scaling due to the necessary physical space and power consumption required by each light source. This is also how old circuit-switched telephone systems used to work.

An actively routed network, though it requires constant managing, can use a single wavelength to get anywhere, much like you can with a car. It can also change routes mid-flight to circumvent problems in the network. The downside is that this method suffers from contention problems, since multiple signals cannot overlap on the same optical wire. The design can be optimized in order to avoid such contention; the ideal case is known as a hitless solution. For our case,

we define hitless as being spatially non-blocking, such that any of the inputs can be routed to any of the available outputs without using the same physical path for a given wavelength. We will see that this can be also scaled to multiple wavelengths, or WDM operation. We demonstrate here a 4x4 hitless router for use in an electronically-integrated optical network.

5.2 Router Architecture

Redirecting or routing light from one path, a waveguide in this case, to another is a matter of coupling two paths together in a way that moves the light completely from one to the other. Simply putting two waveguides together however, will result in broadband coupling [91]. Since we want the switches to be dynamic we need to make them wavelength specific. This is why we turn again to microring resonators which will couple a certain amount of light from some specified set of frequencies, essentially filtering for select wavelength channels.

Switching can, in principle, be accomplished by adding microring resonators alongside waveguide intersections [92]. A microring resonator coupled to two waveguides can allow light to pass through, or if activated, switch the light to the intersecting waveguide. We can quickly improve on this basic design, since such a switch does not take full advantage of all the available paths. If instead we have two rings opposite each other on a single intersection, they can utilize both waveguides fully by creating what we call a "deflection switch", shown in Figure 5.2(a). The inset in the figure shows light passing through since the rings are not activated to the given wavelength. Arranging single or double ring switches in a simple mesh configuration to create a 4x4 router has been shown before [4], but presents the obvious disadvantage of being spatially blocking (seen in Figure 5.2(b)); a limiting factor in modern parallel processing devices

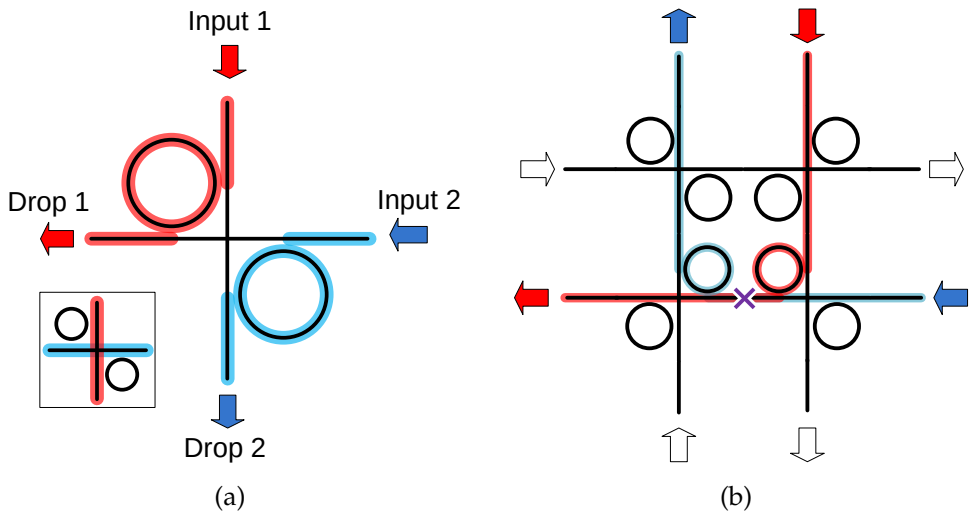

Figure 5.2: Microring resonators are used as individual routing devices; (a) an active switching element switches light at an intersection (insert shows switches in off state), (b) a simple 4x4 routing device highlighting a contention state where a red and a blue signal have used the same physical path.

such as CMPs.

Working in collaboration with Keren Bergman's group in Columbia University, we built a practical 4x4 router as described in their work [93, 94], which allows for parallel hitless routing using the same microring resonators. Hitless routing is accomplished by having one dedicated waveguide for each input-output combination. In this topology, known as a torus mesh, signals are never routed to their direction of origin or to the same direction as another signal, such that there are 12 possible physical paths; 3 for each input direction as shown in Table 5.1. When these paths are operating in parallel there can be no overlap so that the number of router states (where each state specifies four paths) is reduced to 9 as seen in Table 5.2. Each of the paths uses at most one ring resonator, so that the router uses zero, two or four ring resonators to establish each of the router states.

In Figure 5.3 the router is shown with a sample configuration of four optical

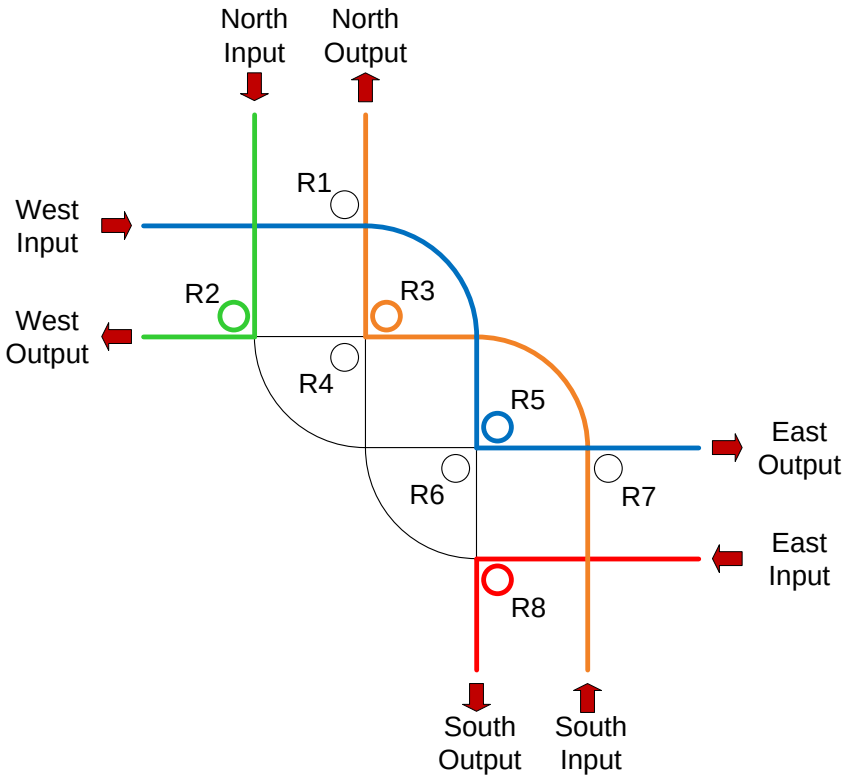

Figure 5.3: The new 4x4 routing device is characterized by its multiple internal paths and ring resonator switching elements, allowing for dynamic hitless networking. Shown in color are four paths highlighting an arbitrary configuration using the maximum number of microrings.

Table 5.1: Input to Output Path Active Resonators.

		Input			
		N	**S**	**E**	**W**
Output	**N**		R3	none	R1
	S	R6		R8	none
	E	none	R7		R5
	W	R2	none	R	

Table 5.2: Full Available Router States with Active Resonators.

		I/O Combination				
		N	**S**	**E**	**W**	**Rings Used**
State Number	**1**	W	N	S	E	R2,R3,R8,R5
	2	W	E	N	S	R2,R7
	3	W	E	S	N	R2,R7,R8,R1
	4	S	N	W	E	R6,R3,R4,R5
	5	S	W	N	E	R6,R5
	6	S	E	W	N	R6,R7,R4,R1
	7	E	W	S	N	R8,R1
	8	E	W	N	S	NONE
	9	E	N	W	S	R1,R4

non-interacting paths which are established by switching four of the rings. The multiple colors depict different signal paths, which all share the same wavelength. In the state depicted, number 1 from Table 5.2, the router uses the maximum number of simultaneous rings as all channels are being switched. Figure 5.4 shows the router within an eight-node implementation of the torus mesh with additional injection and ejection switches as well as the overlapping electrical network. Such a configuration requires 16 routers, and is accessed optically by each of the eight nodes through one of the gateways which would house both modulators and detectors.

5.3 Switching Design

The microring resonators used in the proposed router are particularly useful since they can selectively filter a specific frequency with a device-designed amount of bandwidth. They can pass along the given frequency band to a different waveguide, and in so doing switch the direction of light. Specifically, an input signal is "switched" when the following resonance condition is met [91]:

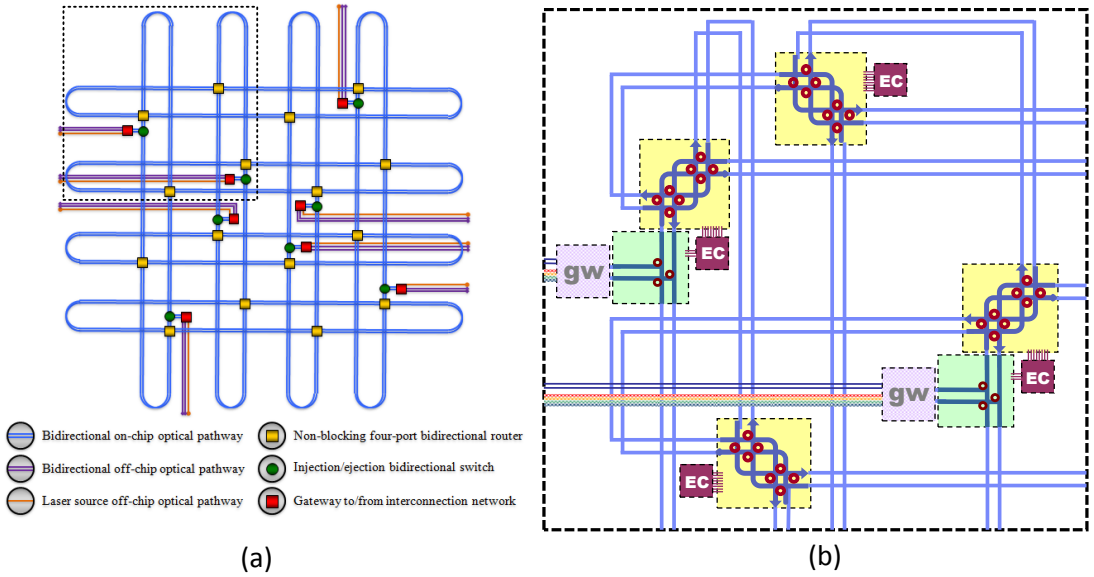

Figure 5.4: Topology of an eight-port non-blocking torus photonic NoC: (a) diagram of optical pathways and photonic network elements; and (b) closeup deliniating waveguides and interfaces to photonic routers.

$$m\lambda_0 = n_{eff}L \tag{5.1}$$

where λ_0 is the resonant wavelength, m is an integer, n_{eff} is the effective index of the optical mode, and L is the length of the resonating cavity.

When the resonance condition is satisfied, light is passed to the drop port instead of the through port as seen back in Figure 5.2(a). By tuning the refractive index of the silicon and hence the effective index of the mode, we can switch a single wavelength between the through port and the drop port. A shift of the effective index Δn_{eff} causes a shift of the resonant wavelength $\Delta\lambda = \lambda_0 \Delta \frac{n_{eff}}{n_{eff}}$ [95].

We have chosen to use the Thermo-Optic effect (TO) [96] because of its strong refractive index tuning in silicon ($\Delta n \approx 1 \times 10^{-2}$). Fabrication effects vary from ring to ring and aligning many of these to one wavelength requires a large

amount of tuning flexibility. The TO effect has been shown to provide a large wavelength shift on the order of 20 nm [97] while operating on a time scale of a few µs using direct heating [98], suitable for circuit-switched networks. we note again that the Electro-Optic effect (EO), known also as the free carrier plasma dispersion effect in silicon, has a faster switching time in this material below 100 ps and could also be used for switching in this device. However, the amount of wavelength shift given by the EO effect is limited to about $\Delta\lambda \approx 2\,\text{nm}$. In practice we need to shift a resonance as far as one Free Spectral Range (FSR) to adjust for fabrication misalignment of an individual ring's resonant wavelength.

5.4 Device Structure and Fabrication

We start with commercial Silicon-on-Insulator (SOI) wafers with a 3 µm buried oxide layer. Using standard e-beam lithography we pattern the waveguides and etch them using reactive ion etching. The waveguide dimensions are 450 nm wide by 250 nm thick, at which we maintain single mode operation. The waveguides are then clad with 1 µm of plasma-enhanced chemical vapor deposition SiO_2 to protect the optical mode. We then planarize the top of the oxide to minimize unevenness in the following metal deposition. This step will reduce the chances of heater failure due to weak points in the metal wire.

We evaporate 300 nm of NiChrome above the cladding and using a double-resist liftoff method creates the 1 µm heaters. Finally another SiO_2 cladding step protects the heaters from outside elements. The microrings use the same waveguide dimensions and have a 10 µm radius with spacing from the waveguides of 200 nm. The final system is seen in Figure 5.5(a).

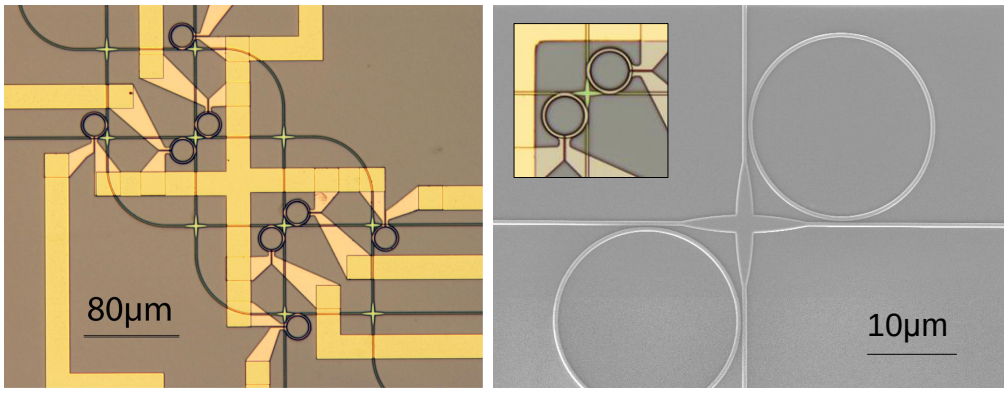

(a) (b)

Figure 5.5: Images of the fabricated router: (a) microscope image of full device shows gold contacts to nichrome heaters above the microrings; (b) an SEM image shows the details of the fabricated waveguide crossing and coupled rings (insert shows close-up of Ω heaters).

5.4.1 Waveguide Crossings

We also reduce the loss typically seen in waveguide intersections by transitioning adiabatically to a larger waveguide cross-section for the intersecting region. By increasing the amount of light confined within the waveguide we limit the interaction of light with the variations of the sidewalls, specifically the large scattering points created by the walls of the crossing waveguide.

We ran finite-difference time-domain (FDTD) simulations of various sized crosses, shown as power drop versus cross size in Figure 5.6 to find an optimum design. Unsurprisingly, all crosses did better with longer transition regions. However, since we wanted to couple a fixed-size ring to both intersecting waveguides, the tapering distance was limited to about the radius size of the ring. Our chosen crossing, shown in Figure 5.5(b), has an arm width of 2 μm and a transition region of 6 μm . It is simulated to have an insertion loss of -0.18 dB and better than -20 dB crosstalk to the intersecting waveguide. We have made

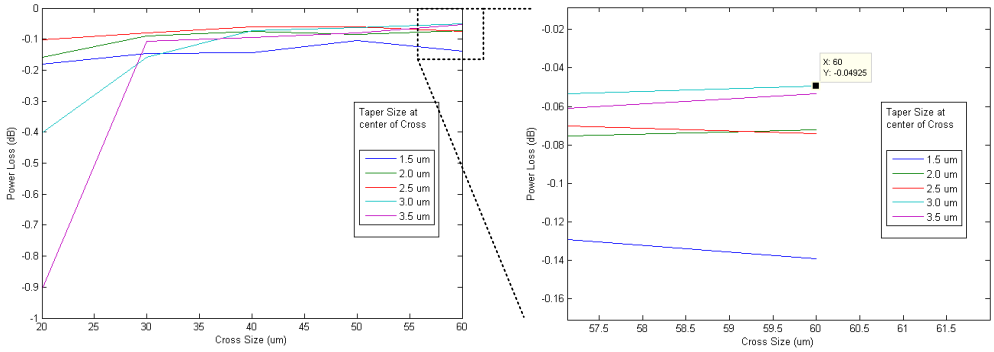

Figure 5.6: FDTD simulation of adiabatically tapered waveguide intersections showing waveguides tapered to between 1.5 and 3.5 µm over a range of lengths.

and tested similar crossings with -0.51 dB insertion loss. MMI-based crossing designs were also considered, and could be used in future design iterations to lower the loss further [99].

5.5 Results

We tested the router for dynamic routing functionality in two separate cases. The first was done at Cornell University using a single tunable wavelength laser source, measuring the ability and effectiveness of the router to actively switch CW light using the eight microheaters in unison. The next test was performed at Columbia University and used several wavelengths of modulated light to test the BER of WDM signals passing through the router.

5.5.1 Single Wavelength CW Routing

We began testing the router by independent ring characterization. Using an Optical Spectrum Analyzer (OSA) and a broadband source, we first found all

microring resonances by tuning each independently from the others. In this way we also tested the switching capabilities of the TO effect. We reproducibly tuned the ring resonances by 10 nm, which is above the FSR of these rings (8.8 nm), without any destructive effects. Switching is achieved at a rate of around 0.25 nm/mW as seen in Figure 5.7, comparable to recent results for SOI microring heaters [77, 100]. An extinction ratio of 17 dB can be achieved with less than 2 mW, as much as 21 dB can be achieved at 6.5 mW. We can also see that our resonances have a FWHM of 0.31 nm which corresponds to a bandwidth of around 38.5 GHz.

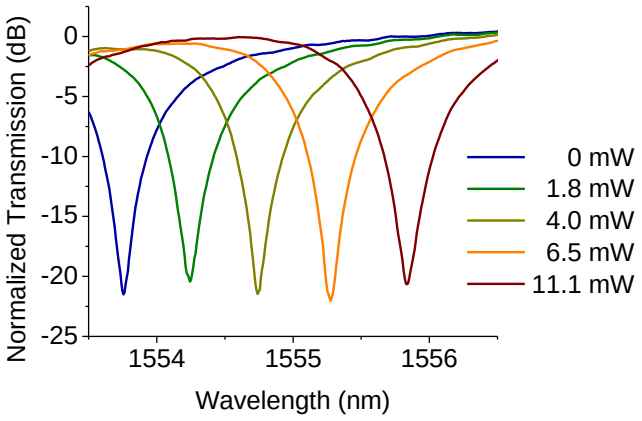

Figure 5.7: DC power applied to the heaters shifts the microring resonances, causing a switch in the transmitted optical power of a given wavelength.

We tested the router by coupling a movable tapered-lens fiber to each of the input waveguides and coupling the output from the chip to another fiber using a collimating lens. A tunable CW laser was used at 1553.67 nm in all of the tests as the input. For each input port the light was routed to all available outputs where power measurements were taken. Each of the input ports has three routing possibilities for a total of 12 unique routing paths. Because the chip was not packaged, we were physically limited from testing the third output of each

input directly such that only eight of the paths were tested directly. Four paths which routed in the direction of the input facet of the chip were measured as drops of power at their respective through ports and were seen to be consistent with the results for the other eight paths.

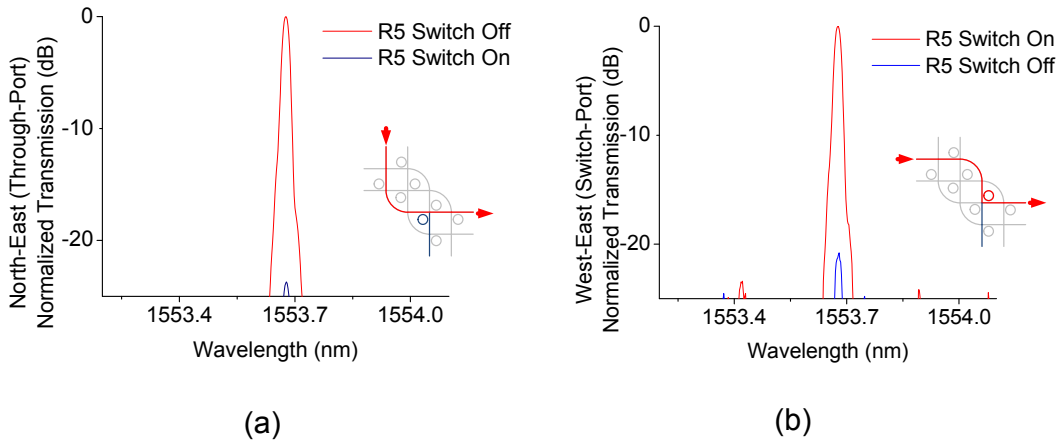

Figure 5.8: Extinction ratio comparison for two switched paths: (a) the power measured at the through port before and after switching, and (b) the power measured at the drop port before and after switching.

For each of these tested paths we looked at the extinction ratio between the signal being routed to the measurement output and the remaining signal when routed away. Figure 5.8 shows two of the inputs routing to one output. In Figure 5.8(a), North to East is routed directly through the waveguide and bypasses four rings along its path, which are detuned from the laser frequency. It is then routed by the thermally tuned R6 microring to the south direction, which is seen at the output as a reduction of power (shown in the same graph). The two measurements shown are taken using the same input and output to compare the extinction ratio. Conversely, in Figure 5.8(b) we switch inputs to measure the signal being routed West to East through R5 and compare it to the same output when R5 is detuned. All other available optical paths were similarly

tested showing a maximum extinction ratio of 20.79 dB for a routed signal.

5.5.2 MULTIPLE WAVELENGTH DATA ROUTING

Subsequently, we tested the router at Columbia University for signal integrity using pseudo-random data and BER measurement tools. In this case three wavelength sources were used in combination to highlight the flexibility of the router to handle WDM signals. Figure 5.9(a) highlights the key elements of the testbed. The three sources are combined equally into a fiber and amplitude-modulated with a 10 Gb/s non-return-to-zero (NRZ) signal. The signals are decorrelated, amplified using an erbium-doped fiber amplifier and then focused into the input waveguide using free-space optics.

At each of the router outputs we couple back into a tapered-lens fiber. The signal is amplified again through an EDFA and filtered for the given channel, as is typical in an optical receiver. The signal is piped into a communications signal analyzer (CSA) and a bit-error-rate tester (BERT) for evaluation. Figure 5.9(b) highlights the three optical paths tested and the signals measured at the outputs are shown in 5.9(c).

Eye diagrams are recorded to show visually the proper transmission of the signal through the router in each of the path combinations (shown in Figure 5.10(a)). The BER tests shown in Figure 5.10(b) quantify the signal degradation and show the exponent of the BER as a function of received power. A BER of 10^{-9} is a typical measure of good performance, at which point we measure power penalties below 1 dB. The exception is that of the *North-to-West* case which loses 1.3 dB, though we attribute this to the physical limitations of this path, which requires collection through the free space optics method instead of fiber coupling method.

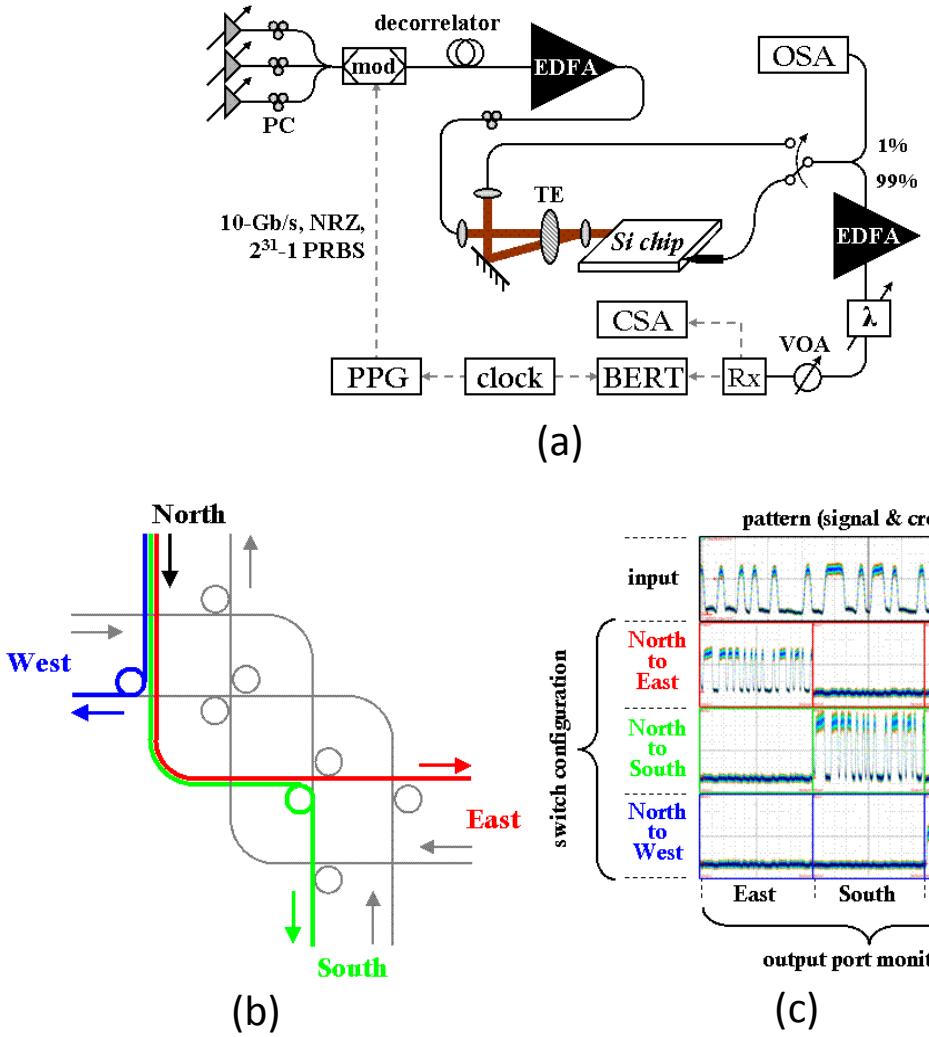

Figure 5.9: WDM router data experiment: (a) experimental setup delineating electrical paths (dashed line) and optical paths (solid lines); (b) optical paths tested within the router; and (c) pattern traces of each of the outputs, each window shows 5 ns in time.

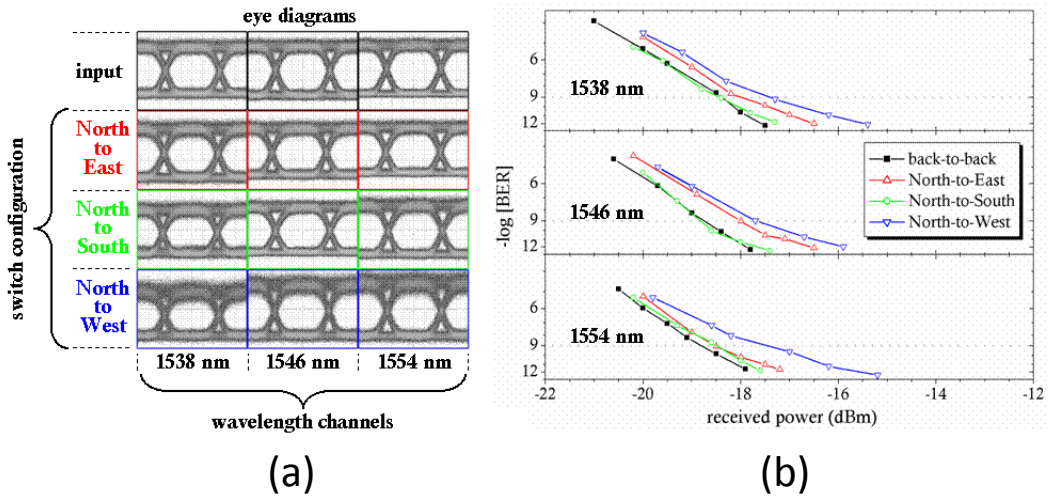

Figure 5.10: Signal analysis results: (a) eye diagrams for each of the wavelength channels routed to each of the destinations; and (b) BER measurements for each of the switch combinations.

5.6 Conclusion

In summary, we have fabricated and shown a working 0.07 mm2 4x4 dynamic hitless router on SOI technology for use in future CMP optical networks. The device is defined by eight 10 µm radius microring resonators in an arrangement that provides efficient and spatially unblocking internal routing. The resonators are individually tuned by means of micro-heaters fabricated out of thin films of nichrome, which are powered by DC current to create a resonance shift. The heater/ring design has a switching power of 0.25 nm/mW. Each of the available paths is individually tested using a tunable laser and OSA for routing extinction ratio which provides a maximum of 20.79 dB and a bandwidth of 38.5 GHz. The router is further tested with PRBS data at multiple wavelengths, and shows power penalties of less the 1.3 dB for three wavelengths traversing the three available paths.

CHAPTER 6

PHOTONIC DEVICE FIELD TESTING

This chapter uses material extracted with permission from the following references:

J. Robinson, J. Marconi, N. Sherwood-Droz, A. Cerqueira, Jr., H. Figueroa, H. Fragnito, and M. Lipson, "Implementation of Silicon Microphotonic Devices in a Geographically-Distributed Fiber Optic Network," in Integrated Photonics and Nanophotonics Research and Applications, (Optical Society of America, 2008), paper IME4.

The benefits of high speed silicon electro-optical modulation, switching, wavelength conversion and gain are useful not only on an on-chip scheme but also at a telecom scale. The technologies we have developed were done with telecom light sources in mind, such that we use the same wavelength ranges [101]. While silicon photonic devices are tested regularly, independently from networks, they are not often tested as part of larger systems [102]. That is because testing with large fiber networks is often difficult and costly, requiring extensive equipment and large spools of fiber. But even these tests might not tell us everything about a real geographically-distributed fiber network. For a real system experiment we need to bring our devices to the field.

6.1 Fiber Network Test Design

We investigated the feasibility of using silicon photonic devices in large scale networks through the use of the active KyaTera network in Campinas, Brazil [103]. The network spans the state of São Paulo and connects multiple universities through buried and aerial optical fibers. These fibers exposed to the real world can experience chromatic dispersion and polarization rotation over the length of the links, creating drastic changes in the conditions of a passing signal. We wanted to know if one of our laboratory chips could perform alongside regular telecom equipment.

The device under test is shown in Figure 6.1. An input waveguide from one side of a chip couples to a microring resonator and exits on the other side of the chip in what we call the through-port. In parallel, a second waveguide also couples to the same ring and exits the chip, acting as a drop-port for the device. The chip is made using e-beam lithography out of an SOI stack. The radius of the ring used in our experiment is approximately 10 µm and the waveguide

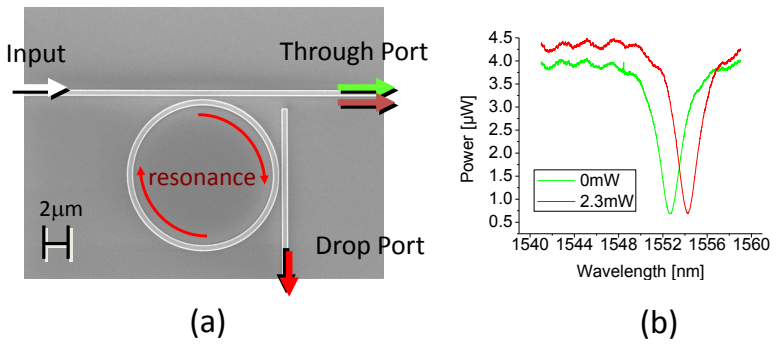

Figure 6.1: Silicon microring device under test: (a) SEM image of device and (b) device optical response with electro-thermal response.

dimensions are 250 nm thick by 450 nm wide.

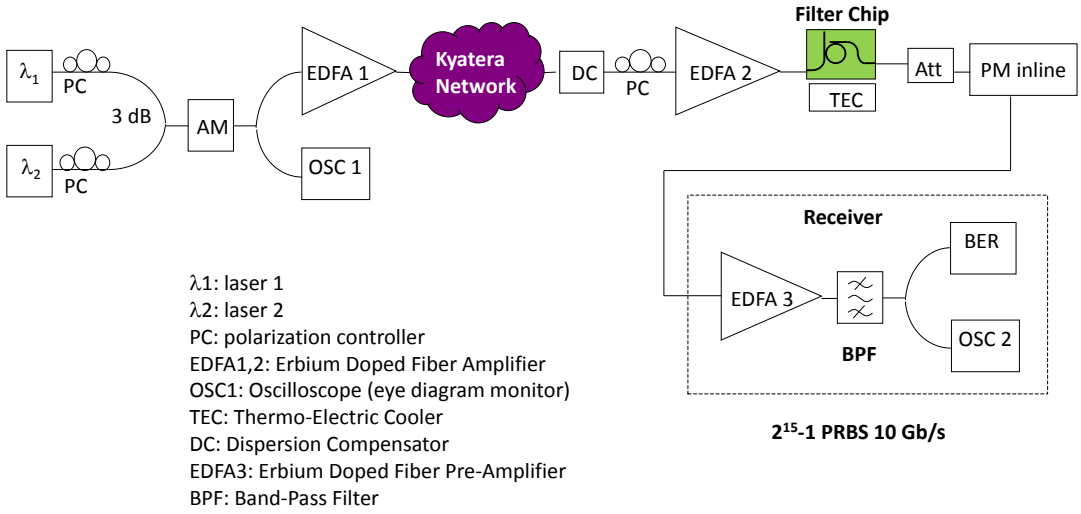

Figure 6.2: KyaTera Network silicon microring filter test setup.

6.2 Experimental Setup

Figure 6.2 shows a schematic of the experimental setup used to measure the microring add-drop filter. Two external cavity lasers (ECL) were externally modulated with a pseudo-random bit sequence (PRBS) at 10 Gb/s with a lithium-niobate-based electro-optic amplitude modulator. These signals were then amplified with an erbium doped fiber amplifier (EDFA) and sent through a 40 km path of laid fiber in the KyaTera Network. To compensate for the added dispersion in the KyaTera Network we included a section of dispersion compensating (DC) fiber. We have also included a dynamic polarization controller (DPC), which fixes the polarization of a randomly polarized input, in order to minimize the polarization effects. Using an inline static polarization controller (PC) we can then select the TM mode for input to the device. Light is then collected from the drop port of the microring filter. With a variable optical attenuator (VOA) we control the amount of signal reaching the receiver for bit error rate (BER) characterization. The receiver consists of an EDFA, a band pass filter ($\tilde{1}$ nm linewidth) and a BER analyzer or and an oscilloscope to record the eye diagram.

6.3 Results

The optical filter was tested for increased Bit Error Rate (BER) within a live fiber network. A comparison was first made between the full network configuration with, and without the device as an in-line filter between the fiber network and the receiver. Figure 6.3 shows the BER comparison of the two paths as a function of the input power of the signal with the pseudo-random bit sequence (PRBS). The results show no added loss to the system.

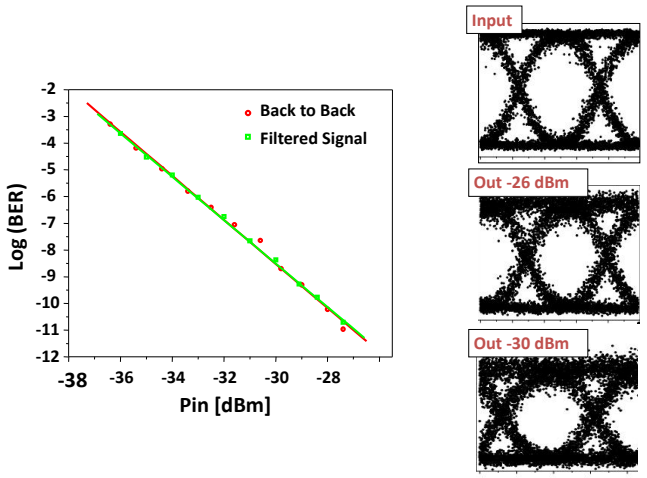

Figure 6.3: BER comparison of 10 Gb/s PRBS signal back-to-back in the KyaTera Network and through our silicon ring filter.

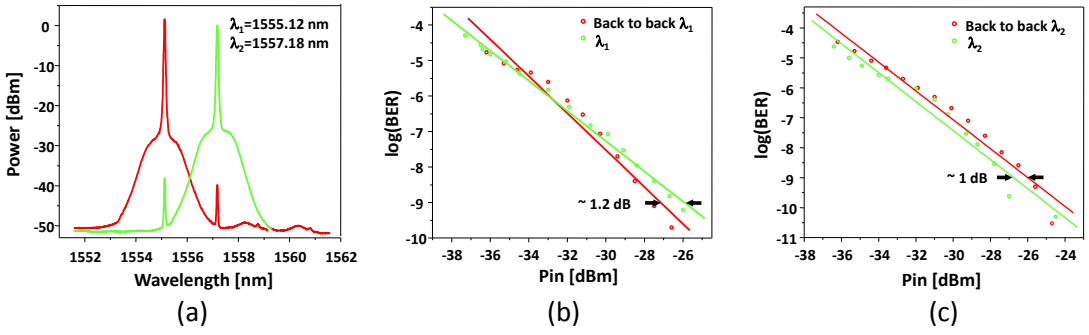

Figure 6.4: Results of channel switching in a silicon filter through the KyaTera Network: (a) OSA measurement of the two channels; (b) back-to-back BER comparison to first channel at 1555.12nm; (b) back-to-back BER comparison to second channel at 1557.18nm

As a second test, we added an additional channel through the use of a secondary laser. The resonance wavelength of the microring is tuned using the thermo-optic effect with a thermo-electric cooler (TEC) below the chip. We switch from one signal at 1555.12 nm to another at 1557.18 nm and test the BER of each to look for the effects of crosstalk. Here again there is no additional loss of signal integrity due to the addition of the filter chip. Figure 6.4(a) shows the output spectra of the channels measured with an optical spectrum analyzer (OSA) at the output of the BPF of the receiver. At a BER of 10^{-9} we see in Figures 6.4(a) and (b) the power penalty for both channels is between 1 and 1.2 dB.

6.4 CONCLUSION

Testing in a real geographically-distributed network, we can show that our devices can effectively operate in complex systems outside of the lab with minimal effect on the overall system. We have shown our silicon microring wavelength filters switch 10 Gb/s signals over 40 Km of laid fiber in the state of São Paulo, Brazil. Minimal power penalty is seen across the device though we note that polarization effects caused by physically exposed fiber can be problematic since our devices are polarization dependent. This issue can be avoided with a DPC.

Chapter 7

Conclusions and Future Directions

Within this thesis we have explored various aspects of a future integration of silicon photonics with CMOS microelectronics. Such an integration, while difficult and costly to an established high-volume based market, might be one of the best ways to continue improving integrated microprocessors in the future. Optics can provide high-bandwidth through single "wires" and low-power and latency to create integrated circuits properly built for the media-driven world we live in today.

First we looked at the standard fabrication of CMOS and compared it to standard silicon photonics practices in research labs. While similar in nature, we found that a very specific limiting factor keeps us from simply attaching photonic elements into an electronic stack, namely the optical isolation layer. Specifically for silicon photonics, we focus on the "buried oxide layer" or BOX, a low refractive-index dielectric that is necessary for the operation of every waveguide and optical device. The problem is the BOX in microelectronics creates a large thermal hinderance to high-performance transistors.

We solve the BOX layer limitation in two separate ways. In one method we integrate our optics at the transistor layer of the CMOS stack in a method we call OxSOI. This approach allows for high temperature processes and high-purity materials, a general assumption in photonics research. We do this by creating a localized BOX layer using a typical high-temperature oxidation of silicon with a special 3-sided silicon nitride mask. The mask is fabricated around each waveguide using a self-aligned method during the waveguide patterning step, which does not require additional lithography. This method produces comparable optics to SOI using only a bulk silicon wafer. However, for its intended application we have to note that it requires space to be taken from the transistors. Such a disadvantage can be considerable since the real estate is already limited due to

dense integrated circuits.

The second method avoids taking space from the transistor layer, or even altering the CMOS process, by using backend photonics. This technique leapfrogs silicon photonics by making future optical inteconnects scalable. It does so at the cost of not using crystalline silicon, an electrically and optically-versatile material. Using the backend of the CMOS stack to integrate photonics with silicon-compatible deposited materials, the opics layer can avoid affecting existing electronics and further scale vertically. This is posible in contrast to c-Si which cannot be deposited, but instead must be grown. Deposited materials have imperfections different to crystalline silicon, but mutiple materials specialized to different tasks can potentially be utilized to make a better system than that available with SOI.

Regardless of which fabrication method is used, lithographic techniques still do not posses the resolution necessary to eliminate fabrication errors which alter the optical properties of critical resonant devices. Beyond this limitation, thermal effects produced by ohmic heating in the electronics can also affect these devices. The solution to both of these problems is active thermal management through localized heaters. We explored this option using waveguide-scale metalized heaters and extended it to an implementation of a 4x4 optical router. Such a router utilizing 8 microring resonators would be difficult to repeat consistently without post-fabrication correction due to the sensitivity of the microrings to the dimensions of its geometry.

From the various work explored, it is this author's view that the nearest-term chip-scale photonics implementation will take the route of a deposited-materials multi-layered implementation. This is similar to the one explained in Chapter 3 along with thermally managed resonant devices as are described in

Chapter 4. The reasons are many-fold but mostly revolve around the economies of volume that CMOS can provide. If we had we been able to develop this level of experience in silicon photonics 50 years ago as transistors were being developed, the story might have been different. Instead III-V photonics dominate the optical devices landscape, but are still made independently from microelectronics, making them low-volume and expensive. In this age of silicon transistors, to deviate from the standard would not result in a commercial product.

An evaluation of such a system was conducted by the researchers at Columbia University and us in [104]. To make a compelling argument we assumed that we could use the best available deposited materials in order to highlight the benefits of each material and 3D photonics scaling. Using high temperature means that this system would be independent from the microprocessor, in contrast to the work discussed in Chapter 3. This device is a stepping stone to full integration, but highlights some of the benefits of moving in a vertically-scaled direction. In fact, one application could still be as an intermediate chip between processor cores where high bandwidth and low latency and power channels could be used much like neighboring integrated circuits. Furthermore, it still uses a fully CMOS-compatible process and materials without the need for SOI.

It was the finding of this work that deposited-photonics would have a great impact on various photonic network architectures on-chip. Using several well-studied on-chip network topologies on our material system network, performance was compared to the original single-layer versions. The result was at least 20% improvement, either in the amount of wavelength channels possible or the required injected power. One such network saw a 457% improvement in bandwidth through increased channels. Furthermore, several new topologies

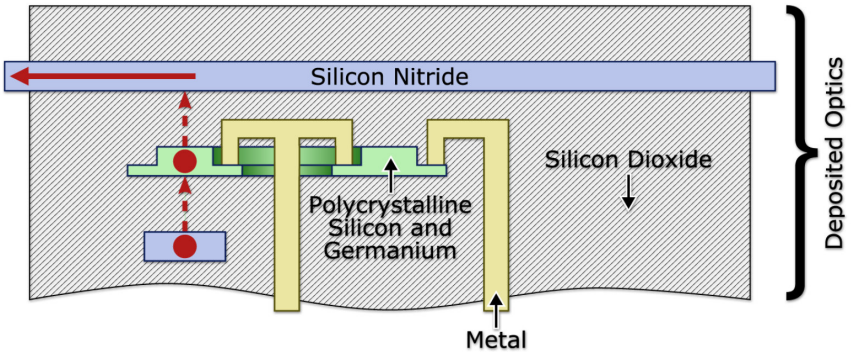

Figure 7.1: Diagram of a multi-layered poly-material photonic interconnect stack, highlighting an active microring resonator coupling between two low-loss silicon nitride waveguides at different layers.

were created that take advantage of the lower waveguide propagation loss and higher non-linear threshold, reduced waveguides crossing losses, and increased channels.

It is the recommendation, and the hope of the author, that future research continues in this direction. The optical devices that finally find their way onto microelectronic processors will have to be completely conformal and compatible with microelectronics, no matter what form they take. Likely they will be a hybrid of traditional silicon photonics and innovative fabrication processes, like much of the work discussed here. Eventually it will be a different story; as the field of microelectronics starts to depend on optics for crucial performance scaling, photonic technologies will be able to focus on efficiency rather than compatibility.

BIBLIOGRAPHY

[1] Nir Magen, Avinoam Kolodny, Uri Weiser, and Nachum Shamir. Interconnect-power dissipation in a microprocessor. In *Proceedings of the 2004 international workshop on System level interconnect prediction*, SLIP '04, pages 7–13, New York, NY, USA, 2004. ACM.

[2] M. Hoffmann, P. Kopka, and E. Voges. Low-loss fiber-matched low-temperature pecvd waveguides with small-core dimensions for optical communication systems. *Photonics Technology Letters, IEEE*, 9(9):1238–1240, 1997.

[3] F.J. Mesa-Martinez, M. Brown, J. Nayfach-Battilana, and J. Renau. Measuring power and temperature from real processors. In *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*, pages 1 –5, april 2008.

[4] W. J. Dally. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, 2004.

[5] Elinor Mills. Google eta? 300 years to index the world's info. `http://news.cnet.com/Google-ETA-300-years-to-index-the-worlds-info/2100-1024_3-5891779.html`, Oct 2005.

[6] A.K. Dutta, N.K. Dutta, and M. Fujiwara. *WDM Technologies: Optical Networks*. Academic Press, 2004.

[7] J. Prat. *Fiber-to-the-Home Technologies*. Kluwer Academic Publishers, 2002.

[8] S. E. Thompson and S. Parthasarathy. Moore's law: the future of si microelectronics. *Materials Today*, 9(6):20–25, 2006.

[9] J. D. Meindl, Q. Chen, and J. A. Davis. Limits on silicon nanoelectronics for terascale integration. *Science*, 293(5537):2044–2049, 2001.

[10] M. Lundstrom. Moore's law forever? *Science*, 299(5604):210–211, 2003.

[11] International technology roadmap for semiconductors. Technical report, 2009.

[12] N.S. Kim, T. Austin, D. Baauw, T. Mudge, K. Flautner, J.S. Hu, M.J. Irwin, M. Kandemir, and V. Narayanan. Leakage current: Moore's law meets static power. *Computer*, 36(12):68 – 75, dec. 2003.

[13] D. J. Frank. Power-constrained cmos scaling limits. *IBM Journal of Research and Development*, 46(2.3):235 –244, march 2002.

[14] D. A. B. Miller and H. M. Ozaktas. Limit to the bit-rate capacity of electrical interconnects from the aspect ratio of the system architecture. *J. Parallel Distrib. Comput.*, 41:42–52, February 1997.

[15] J.C. Palais. *Fiber optic communications*. Prentice Hall, 1988.

[16] A. Shacham, K. Bergman, and L. P. Carloni. On the design of a photonic network-on-chip. In *Networks-on-Chip*, pages 53–64, 2007.

[17] R. G. Beausoleil, P. J. Kuekes, G. S. Snider, Wang Shih-Yuan, and R. S. Williams. Nanoelectronic and nanophotonic interconnect. *Proceedings of the IEEE*, 96(2):230–247, 2008.

[18] J. Lexau, X. Zheng, J. Bergey, A. V. Krishnamoorthy, R. Ho, R. Drost, and J. Cunningham. Cmos integration of capacitive, optical, and electrical interconnects. *International Interconnect Technology Conference, IEEE 2007*, pages 78–80, 2007.

[19] M. Brière, B. Girodias, Y. Bouchebaba, G. Nicolescu, F. Mieyeville, F. Gaffiot, and I. O'Connor. System level assessment of an optical noc in an mpsoc platform. In *Proceedings of the conference on Design, automation and test in Europe*, DATE '07, pages 1084–1089, San Jose, CA, USA, 2007. EDA Consortium.

[20] Hoyeol Cho, Pawan Kapur, and Krishna C. Saraswat. Power comparison between high-speed electrical and optical interconnects for interchip communication. *J. Lightwave Technology*, 22(9):2021, Sep 2004.

[21] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C.W. Holzwarth, M.A. Popovic, H. Li, H.I. Smith, J.L. Hoyt, et al. Building many-core processor-

to-dram networks with monolithic cmos silicon photonics. *Micro, IEEE*, 29(4):8–21, 2009.

[22] L. Pavesi and D.J. Lockwood. *Silicon Photonics*. Springer Verlag, 2004.

[23] R. Soref. The past, present, and future of silicon photonics. *Selected Topics in Quantum Electronics, IEEE Journal of*, 12(6):1678 –1687, nov.-dec. 2006.

[24] Hanjo Rhee, Oliver Lux, Stefan Meister, Ulrike Woggon, Alexander A. Kaminskii, and Hans Joachim Eichler. Operation of a raman laser in bulk silicon. *Opt. Lett.*, 36(9):1644–1646, May 2011.

[25] C. A. Barrios and M. Lipson. Electrically driven silicon resonant light emitting device based on slot-waveguide. *Optics Express*, 13(25):10092–10101, Dec 2005.

[26] X. Sun, J. Liu, L. C. Kimerling, and J. Michel. Room-temperature direct bandgap electroluminesence from ge-on-si light-emitting diodes. *Optics Letters*, 34(8):1198–1200, Apr 2009.

[27] A. Alduino, L. Liao, R. Jones, M. Morse, B. Kim, W. Lo, J. Basak, B. Koch, H. Liu, H. Rong, M. Sysak, C. Krause, R. Saba, D. Lazar, L. Horwitz, R. Bar, S. Litski, A. Liu, K. Sullivan, O. Dosunmu, N. Na, T. Yin, F. Haubensack, I. Hsieh, J. Heck, R. Beatty, H. Park, J. Bovington, S. Lee, H. Nguyen, H. Au, K. Nguyen, P. Merani, M. Hakami, and M. Paniccia. Demonstration of a high speed 4-channel integrated silicon photonics wdm link with hybrid silicon lasers. In *Integrated Photonics Research, Silicon and Nanophotonics*, page PDIWI5. Optical Society of America, 2010.

[28] Matthew N. Sysak, Joel O. Anthes, John E. Bowers, Omri Raday, and Richard Jones. Integration of hybrid silicon lasers and electroabsorption modulators. *Optics Express*, 16(17):12478–12486, Aug 2008.

[29] A. Shacham, K. Bergman, and L. P. Carloni. Photonic networks-on-chip for future generations of chip multiprocessors. *IEEE Transactions on Computers*, 57(9):1246–1260, 2008.

[30] L. Chen, K. Preston, S. Manipatruni, and M. Lipson. Integrated ghz silicon photonic interconnect with micrometer-scale modulators and detectors. *Optics Express*, 17(17):15248–15256, 2009.

[31] S. Manipatruni, Qianfan Xu, B. Schmidt, J. Shakya, and M. Lipson. High

speed carrier injection 18 gb/s silicon micro-ring electro-optic modulator. In *Lasers and Electro-Optics Society, 2007. LEOS 2007. The 20th Annual Meeting of the IEEE*, pages 537 –538, oct. 2007.

[32] L. Chen and M. Lipson. Ultra-low capacitance and high speed germaniumphotodetectors on silicon. *Optics Express*, 17(10):7901–7906, May 2009.

[33] H. L. R. Lira, S. Manipatruni, and M. Lipson. Broadband hitless silicon electro-optic switch for on-chip optical networks. *Optics Express*, 17(25):22271–22280, Dec 2009.

[34] M. A. Foster, A. C. Turner, J. E. Sharping, B. S. Schmidt, M. Lipson, and A. L. Gaeta. Broad-band optical parametric gain on a silicon photonic chip. *Nature*, 441(7096):960–963, 2006. 10.1038/nature04932.

[35] R. Koh. Buried layer engineering to reduce the drain-induced barrier lowering of sub-0.05 µm soi-mosfet. *Japanese Journal of Applied Physics*, 38(Part 1, No. 4B):2294–2299, 1999.

[36] C. "Fenouillet-Beranger, T. Skotnicki, S. Monfray, N. Carriere, and F." Boeuf. "requirements for ultra-thin-film devices and new materials for the cmos roadmap". *"Solid-State Electronics"*, "48"("6"):"961 – 967", "2004". "Silicon On Insulator Technology and Devices".

[37] J.M. Fedeli, M. Migette, L. Cioccio, L. El Melhaoui, R. Orobtchouk, C. Seassal, P. Rojo-Romeo, F. Mandorlo, D. Marris-Morini, and L. Vivien. Incorporation of a photonic layer at the metallizations levels of a cmos circuit. In *Group IV Photonics, 2006. 3rd IEEE International Conference on*, pages 200 –202, 2006.

[38] F. Y. Gardes, G. T. Reed, A. P. Knights, G. Mashanovich, P. E. Jessop, L. Rowe, S. McFaul, D. Bruce, and N. G. Tarr. Sub-micron optical waveguides for silicon photonics formed via the local oxidation of silicon (locos). volume 6898, page 68980R. SPIE, 2008.

[39] R. Pafchek, R. Tummidi, J. Li, M. A. Webster, E. Chen, and T. L. Koch. Low-loss silicon-on-insulator shallow-ridge te and tm waveguides formed using thermal oxidation. *Applied Optics*, 48(5):958–963, Feb 2009.

[40] C. W. Holzwarth, J. S. Orcutt, H. Li, A. Popovic, M, V. Stojanovic, J. L. Hoyt, R. J. Ram, and H. I. Smith. Localized substrate removal technique enabling strong-confinement microphotonics in bulk si cmos processes. In *Conference on Lasers and Electro-Optics/Quantum Electronics and*

Laser Science Conference and Photonic Applications Systems Technologies, page CThKK5. Optical Society of America, 2008.

[41] D. Andriukaitis and R. Anilionis. Thermal oxidation in locos, pbl and swami micro and nano structures. *tC*, 10:6.

[42] P. Dumon, W. Bogaerts, V. Wiaux, J. Wouters, S. Beckx, J. Van Campenhout, D. Taillaert, B. Luyssaert, P. Bienstman, D. Van Thourhout, and R. Baets. Low-loss soi photonic wires and ring resonators fabricated with deep uv lithography. *Photonics Technology Letters, IEEE*, 16(5):1328 –1330, may 2004.

[43] V. R. Almeida, R. R. Panepucci, and M. Lipson. Nanotaper for compact mode conversion. *Optics Letters*, 28(15):1302–1304, Aug 2003.

[44] T. Enomoto, R. Ando, H. Morita, and H. Nakayama. Thermal oxidation rate of a si3n4 film and its masking effect against oxidation of silicon. *Japanese Journal of Applied Physics*, 17:1049, 1978.

[45] D. Wentzlaff, P. Griffin, H. Hoffmann, Bao Liewei, B. Edwards, C. Ramey, M. Mattina, Miao Chyi-Chang, J. F. Brown, and A. Agarwal. On-chip interconnection architecture of the tile processor. *Micro, IEEE*, 27(5):15–31, 2007.

[46] J. Chan, G. Hendry, A. Biberman, and K. Bergman. Architectural exploration of chip-scale photonic interconnection network designs using physical-layer analysis. *Journal of Lightwave Technology*, 28(9):1305–1315, 2010.

[47] A. Joshi, C. Batten, Y. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. Silicon-photonic clos networks for global on-chip communication. In *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*, NOCS '09, pages 124–133, Washington, DC, USA, 2009. IEEE Computer Society.

[48] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn. Corona: System implications of emerging nanophotonic technology. *SIGARCH Comput. Archit. News*, 36:153–164, June 2008.

[49] Yan Pan, J. Kim, and G. Memik. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *High Performance Computer Ar-*

chitecture (HPCA), 2010 IEEE 16th International Symposium on, pages 1 –12, jan. 2010.

[50] N. Sherwood-Droz, A. Gondarenko, and M. Lipson. Oxidized silicon-on-insulator (oxsoi) from bulk silicon: a new photonic platform. *Optics Express*, 18(6):5785–5790, 2010.

[51] M. Ghulinyan, R. Guider, G. Pucker, and L. Pavesi. Monolithic whispering-gallery mode resonators with vertically coupled integrated bus waveguides. *Arxiv preprint arXiv:1102.2370*, 2011.

[52] E. Cianci, A. Schina, A. Minotti, S. Quaresima, and V. Foglietti. Dual frequency pecvd silicon nitride for fabrication of cmuts' membranes. *Sensors and Actuators A: Physical*, 127(1):80–87, 2006. 0924-4247 doi: DOI: 10.1016/j.sna.2005.11.053.

[53] D.K. Sparacin, R. Sun, A.M. Agarwal, M.A. Beals, J. Michel, L.C. Kimerling, T.J. Conway, A.T. Pomerene, D.N. Carothers, M.J. Grove, D.M. Gill, M.S. Rasras, S.S. Patel, and A.E. White. Low-loss amorphous silicon channel waveguides for integrated photonics. In *Group IV Photonics, 2006. 3rd IEEE International Conference on*, pages 255 –257, 0-0 2006.

[54] K. Preston, S. Manipatruni, A. Gondarenko, C. B. Poitras, and M. Lipson. Deposited silicon high-speed integrated electro-optic modulator. *Optics Express*, 17(7):5118–5124, 2009.

[55] L. Colace, G.M.A. Altieri, and G. Assanto. Waveguide photodetectors for the near-infrared in polycrystalline germanium on silicon. *Photonics Technology Letters, IEEE*, 18(9):1094 –1096, 1 2006.

[56] M.J. Madou. *Fundamentals Of Microfabrication: The Science Of Miniaturization*, volume 1 of *CRC Press LLC.* Taylor & Francis, Inc., 1997.

[57] E. Herth, B. Legrand, L. Buchaillot, N. Rolland, and T. Lasri. Optimization of sinx:h films deposited by pecvd for reliability of electronic, microsystems and optical applications. *Microelectronics Reliability*, 50(8):1103–1106, 2010. 0026-2714 doi: DOI: 10.1016/j.microrel.2010.04.011.

[58] M. Melchiorri, N. Daldosso, F. Sbrana, L. Pavesi, G. Pucker, C. Kompocholis, P. Bellutti, and A. Lui. Propagation losses of silicon nitride waveguides in the near-infrared range. *Applied Physics Letters*, 86(12):121111–3, 2005.

[59] A. Gondarenko, J. S. Levy, and M. Lipson. High confinement micron-scale silicon nitride high q ring resonator. *Optics Express*, 17(14):11366–11370, 2009.

[60] G. N. Parsons, J. H. Souk, and J. Batey. Low hydrogen content stoichiometric silicon nitride films deposited by plasma-enhanced chemical vapor deposition. *Journal of Applied Physics*, 70(3):1553 –1560, aug 1991.

[61] S. C. Mao, S. H. Tao, Y. L. Xu, X. W. Sun, M. B. Yu, G. Q. Lo, and D. L. Kwong. Low propagation loss sin optical waveguideprepared by optimal low-hydrogen module. *Optics Express*, 16(25):20809–20816, 2008.

[62] W. Bogaerts, P. Dumon, D. V. Thourhout, and R. Baets. Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides. *Optics Letters*, 32(19):2801–2803, 2007.

[63] C. Manolatou, M. J. Khan, S. Fan, P. R. Villeneuve, H. A. Haus, and J. D. Joannopoulos. Coupling of modes analysis of resonant channel add-drop filters. *Quantum Electronics, IEEE Journal of*, 35(9):1322–1331, 1999.

[64] V. R. Almeida, C. A. Barrios, R. R. Panepucci, and M. Lipson. All-optical control of light on a silicon chip. *Nature*, 431(7012):1081–1084, 2004. 10.1038/nature02921.

[65] F. Gan, T. Barwicz, M.A. Popovic, M.S. Dahlem, C.W. Holzwarth, P.T. Rakich, H.I. Smith, E.P. Ippen, and F.X. Kartner. Maximizing the thermo-optic tuning range of silicon photonic structures. In *Photonics in Switching, 2007*, pages 67 –68, aug. 2007.

[66] F. Y. Gardes, D. J. Thomson, N. G. Emerson, and G. T. Reed. 40 gb/s silicon photonics modulator for te and tm polarisations. *Opt. Express*, 19(12):11804–11814, Jun 2011.

[67] S.T. Chu, W. Pan, S. Sato, T. Kaneko, B.E. Little, and Y. Kokubun. Wavelength trimming of a microring resonator filter by means of a uv sensitive polymer overlay. *Photonics Technology Letters, IEEE*, 11(6):688 –690, jun 1999.

[68] H. Haeiwa, T. Naganawa, and Y. Kokubun. Wide range center wavelength trimming of vertically coupled microring resonator filter by direct uv irradiation to sin ring core. *Photonics Technology Letters, IEEE*, 16(1):135 –137, jan. 2004.

[69] Yiran Shen, Ivan B. Divliansky, Dimitri N. Basov, and Shayan Mookherjea. Perfect set-and-forget alignment of silicon photonic resonators and interferometers. In *Optical Fiber Communication Conference*, page PDPC3. Optical Society of America, 2011.

[70] D.A.B. Miller. Rationale and challenges for optical interconnects to electronic chips. *Proceedings of the IEEE*, 88(6):728 –749, jun 2000.

[71] S.J.B. Yoo. Future prospects of silicon photonics in next generation communication and computing systems. *Electronics Letters*, 45(12):584 –588, 4 2009.

[72] H. F. Hamann, A. Weger, J. A. Lacey, Z. Hu, P. Bose, E. Cohen, and J. Wakil. Hotspot-limited microprocessors: Direct temperature and power distribution measurements. *Solid-State Circuits, IEEE Journal of*, 42(1):56 –65, jan. 2007.

[73] J.H.J. Janssen and H.J.M. Veendrick. Spatial and temporal temperature variations in cmos designs. In *Thermal Investigations of ICs and Systems, 2009. THERMINIC 2009. 15th International Workshop on*, pages 31 –35, oct. 2009.

[74] G. Cocorullo, F.G. Della Corte, I. Rendina, and P.M. Sarro. Thermo-optic effect exploitation in silicon microstructures. *Sensors and Actuators A: Physical*, 71(1-2):19–26, 1998.

[75] Payam Alipour, Ehsan Shah Hosseini, Ali Asghar Eftekhar, Babak Momeni, and Ali Adibi. Temperature-insensitive silicon microdisk resonators using polymeric cladding layers. In *Conference on Lasers and Electro-Optics/International Quantum Electronics Conference*, page CMAA4. Optical Society of America, 2009.

[76] B. Guha, B. B. C. Kyotoku, and M. Lipson. Cmos-compatible athermal silicon microring resonators. *Optics Express*, 18(4):3487–3493, Feb 2010.

[77] F. Gan, T. Barwicz, M. A. Popovic, M. S. Dahlem, C. W. Holzwarth, P. T. Rakich, H. I. Smith, E. P. Ippen, and F. X. Kartner. Maximizing the thermo-optic tuning range of silicon photonic structures. In *Photonics in Switching*, pages 67–68, 2007.

[78] H.D. Young, R.A. Freedman, and L. Ford. *University physics*. Addison Wesley Longman, 1996.

[79] Imam H. Kazi, P. M. Wild, T. N. Moore, and M. Sayer. The electromechanical behavior of nichrome (80/20 wt. *Thin Solid Films*, 433(1-2):337–343, 2003.

[80] D. Geuzebroek, E. J. Klein, H. Kelderman, and A. Driessen. Wavelength tuning and switching of a thermooptic microring resonator. *ECIO*, pages pp. 395–398, 2003. wavelength resonance shift citation.

[81] A. H. Atabaki, E. Shah Hosseini, A. A. Eftekhar, S. Yegnanarayanan, and A. Adibi. Optimization of metallic microheaters for high-speed reconfigurable silicon photonics. *Optics Express*, 18(17):18312–18323, Aug 2010.

[82] Timo T. Aalto, Markku Kapulainen, Sanna Yliniemi, Pivi Heimala, and Matti J. Leppihalme. Fast thermo-optical switch based on soi waveguides. In Yakov S. Sidorin and Ari Tervonen, editors, *Integrated Optics: Devices, Materials, and Technologies VII*, volume 4987, pages 149–159. SPIE, 2003.

[83] M. Harjanne, M. Kapulainen, T. Aalto, and P. Heimala. Sub- mu;s switching time in silicon-on-insulator mach-zehnder thermooptic switch. *Photonics Technology Letters, IEEE*, 16(9):2039 –2041, sept. 2004.

[84] U. Fischer, T. Zinke, B. Schuppert, and K. Petermann. Singlemode optical switches based on soi waveguides with large cross-section. *Electronics Letters*, 30(5):406 –408, mar 1994.

[85] Z. Cui. *Micro-nanofabrication: Technologies and Applications*. Springer Verlag, 2005.

[86] B. A. Small, B. G. Lee, K. Bergman, Q. Xu, and M. Lipson. Multiple-wavelength integrated photonic networks based on microring resonator devices. *Journal of Optical Networking*, 6(2):112–120, 2007.

[87] F. Xia, M. Rooks, L. Sekaric, and Y. Vlasov. Ultra-compact high order ring resonator filters using submicron silicon photonic wires for on-chip optical interconnects. *Optics Express*, 15(19):11934–11941, 2007. high order filters.

[88] S. T. Chu, B. E. Little, W. Pan, T. A. Kaneko, S. A. Sato, and Y. A. Kokubun. An eight-channel add-drop filter using vertically coupled microring resonators over a cross grid. *Photonics Technology Letters, IEEE*, 11(6):691–693, 1999.

[89] B. E. Little, S. T. Chu, W. Pan, and Y. A. Kokubun Y. Kokubun. Microring resonator arrays for vlsi photonics. *Photonics Technology Letters, IEEE*, 12(3):323–325, 2000.

[90] L. Zhang, M. Yang, Y. Jiang, E. Regentova, and E. Lu. Generalized wavelength routed optical micronetwork in network-on-chip. *Proceedings of the 18th IASTED International Conference on Parallel and Distributed Computing and Systems*, 2006. network wron aonoc.

[91] C. R. Pollock and M. Lipson. *Integrated Photonics*. Kluwer Academic Publishers, 2003.

[92] C. K. Madsen and J. H. Zhao. *Optical Filter Design and Analysis*. John Wiley & Sons, Inc., 2001.

[93] A. Shacham, B. G. Lee, A. Biberman, K. Bergman, and L. P. Carloni. Photonic noc for dma communications in chip multiprocessors. In *Proceedings of the 15th Annual IEEE Symposium on High-Performance Interconnects*, HOTI '07, pages 29–38, Washington, DC, USA, 2007. IEEE Computer Society.

[94] H. Wang, M. Petracca, A. Biberman, B. G. Lee, L. P. Carloni, and K. Bergman. Nanophotonic optical interconnection network architecture for on-chip and off-chip communications. In *Optical Fiber communication/National Fiber Optic Engineers Conference*, pages 1–3, 2008.

[95] M. S. Nawrocka, T. Liu, X. Wang, and R. R. Panepucci. Tunable silicon microring resonator with wide free spectral range. *Applied Physics Letters*, 89(7):071110–3, 2006.

[96] W. M. J. Green, H. F. Hamann, L. Sekaric, M. J. Rooks, and Y. A. Vlasov. Ultra-compact reconfigurable silicon optical devices using micron-scale localized thermal heating. In H. F. Hamann, editor, *Optical Fiber Communication and the National Fiber Optic Engineers Conference*, pages 1–3, 2007.

[97] M. Lipson. Compact electro-optic modulators on a silicon chip. *Selected Topics in Quantum Electronics, IEEE Journal of*, 12(6):1520–1526, 2006.

[98] M. W. Geis, S. J. Spector, R. C. Williamson, and T. M. Lyszczarz. Submicrosecond submilliwatt silicon-on-insulator thermooptic switch. *Photonics Technology Letters, IEEE*, 16(11):2514–2516, 2004.

[99] F. Xu and A. W. Poon. Silicon cross-connect filters using microring resonator coupled multimode-interference-based waveguide crossings. *Optics Express*, 16(12):8649–8657, 2008.

[100] H.Y. Ng, M.R. Wang, D. Li, X. Wang, J. Martinez, R.R. Panepucci, and K. Pathak. 4x4 wavelength-reconfigurable photonic switch based on thermally tuned silicon microring resonators. *Optical Engineering*, 47(4):044601, 2008.

[101] M. Lipson. Guiding, modulating, and emitting light on silicon-challenges and opportunities. *J. Lightwave Technol.*, 23(12):4222, Dec 2005.

[102] A. Melloni, M. Martinelli, G. Cusmai, and R. Siano. Experimental evaluation of ring resonator filters impact on the bit error rate in non return to zero transmission systems. *Optics Communications*, 234(1-6):211 – 216, 2004.

[103] São Paulo Research Foundation (Fapesp). Kyatera project. `http://www.kyatera.fapesp.br`.

[104]