Analog Integrated Photonics

by

Cheryl M. Sorace-Agaskar

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

Doctor of Philosophy

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2015

Signature redacted

Author .....
Department of Electrical Engineering and Computer Science
January 30, 2015

Signature redacted

Certified by .....
Michael R. Watts
Associate Professor
Thesis Supervisor

Signature redacted

Certified by .....
Franz X. Kaertner
Adjunct Professor
Thesis Supervisor

Signature redacted

Accepted by .....
Leslie Kolodziejski
Chairman, Department Committee on Graduate Theses

Analog Integrated Photonics

by

Cheryl M. Sorace-Agaskar

Submitted to the Department of Electrical Engineering and Computer Science
on January 30, 2015, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

Silicon photonics promises to revolutionize the field of optics by allowing for cheap, compact, low-power and low-noise optical systems on chip. In the past decade and a half, the basic functionality and acceptable performance of many individual integrated photonic components have been demonstrated, particularly in the digital regime. However, there are several challenges remaining before these advances can truly be exploited to create large-scale, commercial, *analog* integrated photonic systems. In this thesis, we address three of these challenges: (1) managing photonic layout and design of large-scale, complex systems jointly with CMOS driving circuitry, (2) integrating analog optical components in silicon, and (3) integrating photonic light sources in silicon.

First, we present a comprehensive VerilogA modeling toolkit for the simulation of large, joint photonic plus CMOS systems as part of the creation of a full photonic process design kit (PDK) and demonstrate its use. Other smaller contributions to the PDK and process are also described. Next, we describe the development of two modulators meant for analog applications: an integrated, linearized Mach-Zehnder modulator and an integrated single-sideband modulator, both of which are measured to have impressive performance. Then, we discuss the development of an integrated mode-locked laser to serve as an on-chip light source for precision, low-noise optical applications. Finally, we describe preliminary work toward creating fully integrated analog systems, with the ultimate aim of demonstrating a compact, low-noise microwave oscillator.

Thesis Supervisor: Michael R. Watts
Title: Associate Professor

Thesis Supervisor: Franz X. Kaertner
Title: Adjunct Professor

Acknowledgments

As I finish my PhD, I am extremely grateful to the many wonderful people with whom I have worked over the years, and without whom this work would not have been possible. First and foremost I have to thank my wonderful advisors, Michael Watts and Franz Kärtner, who provided me with the opportunity and funding to do the research presented in this thesis and whose insights were always helpful. Their support was invaluable. I would also like to thank Dr. Paul Juodawlkis for serving on my thesis committee, for his many helpful thoughts and questions and for being the driving force behind the IPI (now IQI) collaboration between campus and MIT Lincoln Laboratory from which I benefited significantly. I am privileged to have had the chance to work with Prof. Vladimir Stojanovic whose dedication has been inspiring and whose comments have been invaluable. Prof. Leslie Kolodziejski has been a fantastic academic advisor; I couldn't ask for better. Prof. Terry Orlando was a great professor to TA for — I learned a lot from him. All of the professors I took classes from over the years get my thanks for the knowledge they so wonderfully imparted.

I want to acknowledge Anatoly Khilo who took me under his wing when I started and explained the work with patience and care. Steve Spector was also instrumental in getting my photonics career off the ground. His dedication to getting things done, no matter what, was wonderful. Jonathan Leu did the CMOS driving circuitry for all the integrated devices and collaborated closely with me on the VerilogA code. His thoughtfulness, care, and attention to detail made working with him a joy. Patrick Callahan was also a joy to work with; I am glad that the mode-locked laser work is being left in such capable hands. I am grateful, too, for the work that the rest of the mode-locked laser team: Anna Baldcheva, Salih Magden, and Jonathan Bradley for the Al_2O_3 depositions and Katia Shtyrkova for measuring the full laser structures.

My officemate for the last three and half years, Ami Yaacobi, has been an invaluable source of sound advice, technical and otherwise, and I have benefited a lot from his insights. Dave Cole has also been a great source of wisdom on RF design and life

as well as great fun to work with. I give my appreciation to Michele Moresco, who worked tirelessly to make all of the mask runs happen, who was always extremely supportive of me and my work, and has been great to know. Speaking of fabrication runs, designs are nothing if one cannot test them, so I extend my thanks to the fabrication team at MIT Lincoln Laboratory, as well as at CNSE Albany for their fabrication efforts on this project. I would especially like to thank Douglas D. Coolbaugh, Gerald Leake, who always found a way to make things work.

Thanks always go to Dorothy, our wonderful administrative assistant in the OQE group, who truly makes the magic happen. And also to the many wonderful administrative assistants we've had in the Photonics group during my time here: Laura von Bosau, Donna Gale, Gabrielle Brewington, Susan Davco and Stacy McGrath. (Thanks particularly to Laura for always having food available to satisfy my pregnancy-related need to eat constantly.)

There are others too numerable to name. All of my group mates, in both the Photonics and OQE groups, have been a wonderful source of support, whether it's bouncing ideas around, getting technical help in the lab, or advice on courses. I am blessed to have worked with so many wonderful, dedicated, talented people. My colleagues in the IPI/IQI initiative have also been fantastic and I am grateful to have worked with them and to be able to do so again in the future.

I couldn't have done it without the love and support of my friends and family. Kendra and Tom have been great sources of strength. My brother, Brian; parents; sister-in-law, Manishika; and mother-in-law have also all been amazing through all the ups and downs. And, of course, I am grateful to Ameya, the rock on which my world stands, and to Simon around whom it now turns.

God Bless,

Cheryl Agaskar

Contents

1 Introduction **23**

1.1 Silicon Photonics 23

1.2 Low Noise Microwave Oscillators 26

1.3 Integrated Photonic Low Noise Oscillator 28

2 Joint Modeling and a Process Design Kit (PDK) for a Photonics Platform **33**

2.1 Introduction to the Silicon Photonics Process and Platform 35

2.2 A PDK for photonics - Layout and Verification 41

2.2.1 Layout 42

2.2.2 Verification 44

2.3 A PDK for photonics - VerilogA System Modeling 48

2.3.1 Why VerilogA? 51

2.3.2 Simulation Architecture and Philosophy 51

2.3.3 Mathematical Representation of Primitive Component Blocks 57

2.3.4 Simple Passive and Active Composite Devices 61

2.3.5 Full System Simulations 72

2.3.6 Advanced features 78

2.3.7 Summary and Future work 81

3 Developing Modulators for Analog Applications **83**

3.1 Integrated Linear Mach-Zehnder Modulators 84

3.1.1 Previous Art and a Note on Units 84

3.1.2 Approach and theory 88
3.1.3 Design and DC results 95
3.1.4 AC and Linearity results 100
3.1.5 Digital Performance 101
3.1.6 Summary and Future work 103
3.2 Integrated Single Sideband Modulator 103
3.2.1 Theory and Prior art 104
3.2.2 Design, Layout, and Results 114
3.2.3 Summary and Future Work 119

4 Integrated Mode-locked Laser in the Silicon Photonics+CMOS platform 121

4.1 Introduction 121
4.1.1 Desired MLL features 122
4.2 Platform Review and Basic Structures 123
4.3 Laser Components 124
4.3.1 Waveguide 126
4.3.2 Bends 132
4.3.3 Loop Mirror 136
4.3.4 SA tapers 138
4.3.5 Output tapers 141
4.3.6 Erbium Trench Transitions, Active Laser Devices and Other Components 142
4.4 MLL Architectures and Lasing Results 147
4.4.1 Passive Designs 147
4.4.2 Active Designs 154
4.5 Summary and Future Work 158

5 Integrated Pound-Drever-Hall Loops and Other Integrated Systems 161

5.1 Other Devices and Integration with CMOS circuitry 161
5.2 Fully integrated PDH Loop 165

5.3 Balanced Detector and Frequency Divider 172
5.4 Summary and Future Work 173

6 Summary of Contributions and Parting Thoughts **175**

A Overview of Modulator Approach and Design Process **177**

List of Figures

1-1 Silicon photonics provides a way forward for drastically reducing the size, weight, and power of large optical systems, including phased array antennas (top) and LIDAR systems (bottom). Figures from [5, 116, 7, 8], figure based off figure by Ami Yaacobi. 24

1-2 A worker on a CMOS silicon fabrication line. Image from [135]. . . . 25

1-3 Silicon photonic systems on a chip. Clockwise from top left: 4-channel on chip WDM link [126], on-chip phased-array [5], integrated heterodyne interferometer [115], on-chip optical trapping for biomolecule manipulation [129], integrated photonic front-end for an analog-to-digital converter [1], on-chip LIDAR steerable antenna [116]. 27

1-4 (a.) Poseidon Scientific (now Raytheon) sapphire loaded cavity oscillators. Rack mounted sapphire loaded cavity oscillator (top) and Shoebox Oscillator ®(bottom), From [123] (b.) Schematic of NIST frequency divider oscillator. From [118]. 29

1-5 Proposed integrated low-noise microwave oscillator 30

1-6 Schematic diagram of proposed oscillator [119, 120]. 30

2-1 (a) Photonic-only chip layer stack. All layers other than the Er Doped Al_2O_3 (light pink) are fabricated at CNSE Albany. The Al_2O_3 is deposited at MIT. (b) The full 3D integrated layer stack showing photonics wafer flipped and wafer bonded to the CMOS wafer (From [47]). The only difference among the photonics layers compared to (a) is that the laser trench now comes in from the back (silicon handle side) instead of the the front (metal side) to allow for post-3D-integration deposition of the Al_2O_3 at MIT. As before, the rest of the fabrication is carried out at CNSE Albany. 36

2-2 Example final doping distribution showing density of p-type (blue) and n-type (orange,red) dopants in a MZ diode modulator (top), a disk diode modulator (middle) and an integrated heater (bottom) 39

2-3 Full Silicon MZM with CMOS driving circuitry is laid out in cadence. Simple shapes (such as basic waveguides and tapers) are combined to create more complicated structures (such as the ridge taper shown), these more complicated structures (ridge taper, coupler, phase shifter, heater) are combined to create the full photonic layout. The full circuit layout is added to the photonic layout. Thus, the photonic and electronic pcells together create the layout of the full system. Designs therefore have the ability to include all 100+ photonic and CMOS layers, allowing the full system to be viewed together. A zoom in of the heater and heater driver section of the full layout is shown at bottom. 45

2-4 (a) Schematic of labeling for connectivity for LVS verification (b) Example photonic layout including labels 47

2-5 A ring resonator (diagrammed at top) is broken into components. The VerilogA modules for these components are placed and connected in a schematic in Cadence (middle). A schematic symbol is created to represent this schematic block (bottom). This symbol can be inserted wherever a ring is needed as opposed to having to redraw the middle block. Relevant variables are passed through this symbol to the underlying schematic. 56

2-6 (a) Diagram showing the components of a ring resonant filter and accompanying layout using our toolkit in Cadence Virtuoso. (b). A full link schematic including laser and photodetector is created to simulate the performance of the device. (c). Amplitude and phase response of the ring showing good agreement between VerilogA simulations and MATLAB transfer function code. 62

2-7 (a). Diagram showing the components of a ring resonant filter with drop port and accompanying layout using our toolkit in Cadence Virtuoso. (b). Amplitude response of the ring showing output intensity in the through and drop port. One can see that the critical coupling has better extinction in the through port but less power in the drop port output, while matched coupling has worse extinction and higher drop port output power, as expected. Good agreement is seen between VerilogA simulations and MATLAB transfer function code 65

2-8 (a). Diagram showing the components of a two ring series filter with drop port and accompanying layout using our toolkit in Cadence Virtuoso. (b). Amplitude response of the ring showing output intensity in the through and drop port showing good agreement between VerilogA simulations and MATLAB transfer function code. 66

2-9 (a) . Diagram showing the components of a single ring modulator and accompanying layout using our toolkit in Cadence Virtuoso. (b). Modulator amplitude transfer function as a function of applied phase-shifter bias, showing good agreement between VerilogA simulations and MATLAB transfer function code. (c). Schematic of full like. Output from a CW laser is feed into a ring modulator driven by a PRBS source. The output of the modulator is sent to a photodiode. 68

2-10 Simulation output, showing electrical driving voltage, field amplitude inside the ring, through port field amplitude, and output of photodiode. 69

2-11 (a.) Diagram of single side band modulator component system (from reference [129]) and accompanying layout using our toolkit in Cadence Virtuoso using basic components (couplers, heaters, phase-shifters). (b) Fouier transform of SSB output as a function of modulation amplitude. Input CW signal is shown in black in upper left. Following graphs show steadily increasing modulation depth. Increase in power transferred to desired sideband and accompanying increase in power in undesired harmonics is seem. VerilogA and MATLAB code show good agreement. 70

2-12 (a). Schematic of WDM link with four PRBS channels and one clock channel. (b) real part of signal on WDM waveguide (top), input and output of modulator channels (bottom four channels). 74

2-13 Schematic of PDH loop (after citen19) and schematic layout in Cadence Virtuoso. (b) PDH error function in low and high frequency regime, showing good agreement between MATLAB and VerilogA code. (c). Phase noise simulation done using our toolkit showing noise of the free-running CW laser (red) and of laser locked to cavity (black) showing successful phase-noise suppression. 76

2-14 (a) Schematic layout of PDH loop using nonideal photonics and full driving circuitry in Cadence Virtuoso. (b) Left: Output PDH error function in low frequency regime. Distortions are clearly visible, but sign is still correct allowing loop to lock. Right: Laser frequency locks to cavity reference in 100nm (red), real part of laser electric field as locking occurs (green). 77

2-15 A Gaussian pulse undergoes dispersion in a waveguide. 81

3-1 Schematic of MZM Layout. From [98]. 89

3-2 (a) Cross-section of phase-shifter showing doping distribution, (b) Simulated effective index and effective field loss of phase-shifter as a function of applied bias. From [94]. 97

3-3 Simulation output showing predicted distance of the third harmonic below the fundamental as a function of input drive power. Predicted input IP3 was 64dBm. From [94]. 98

3-4 Schematic of Measurement Setup. From [94]. 98

3-5 (a) Output of MZM as a function of heater bias, showing the full range of the output can be achieved. Heater is 100μm in length, length of phase-shifter section (unbiased) varied from 500-2000μm (b) Output of the MZM as a function of DC phase-shifter bias, showing full device operating range can be achieved, (c) Extracted phase as a function of bias for devices of different length. Results match well to original theory. Measured VπL = 0.27Vcm. (d) Graph of reverse-bias current showing breakdown around 7V. From [94]. 99

3-6 High speed performance of MZM, showing 3dB roll-off of 25GHz in 550μm device (left) and 10GHz in 2mm device (right). From [94]. . . 100

3-7 Comparison of measured and simulation output of predicted distance of the third harmonic below the fundamental as a function of input drive power. Good agreement is seen between simulation and theory. Predicted input IP3 was 64dBm. Measured input IP3 was 68.4dBm. From [94]. 101

3-8 SFDR measurement of the device. Showing (a) the SFDR intercept with the noise floor at -127dBm in our measurement setup. The SFDR was measured to be 82.1dBHz2/3. (b) The fundamental at -44.9dBm and the third harmonic in the noise floor. From [94] 102

3-9 Open eye diagrams at 1 Gb/s, 5 Gb/s, and 12.5 Gb/s. 102

3-10 (a) Schematic of Single Sideband Modulator (b) Equivalent Layout with all heaters in inner arms, (c) Frequency domain representation of the signal at each point along the modulation path. 106

3-11 Labeled diagram of SSB modulator 107

3-12 As RF drive voltage increases power transfer to desired sideband increases, but undesired harmonic terms increase faster, making choice of operating voltage a trade off in loss vs. nonlinearity. 111

3-13 (a.) Effects of variation in RF signal phase on SSB output. (b.) Effect of incorrect bias on one heater inside one of the nexted MZMs. Note the appearance of the even-ordered harmonics and the carrier. 115

3-14 Measurement set-up. * = laser, OSA = optical spectrum analyzer, SG = signal generator, DC = DC power source. From [?]. 116

3-15 Output of single sideband modulator, showing carrier suppression greater than 18dB and suppression of other sideband and tones more than 15dB. The expected power loss is -9.3dB and the measured loss is -10.4dB. From [?]. 117

3-16 Shift in lower sideband location as a function of modulation frequency. From [?]. 118

3-17 Single-sideband operation in same device at 1570nm and 1590nm. Carrier and other sideband suppression >15dB can be seen at 1570nm, and carrier suppression >10dB and other sideband suppression >25dB can be seen at 1590nm. From [?]. 118

4-1 (a) Full Photonic + CMOS layer stack after [47] and chapter 2, (b) Short-loop layer stack, which includes only the nitride and Er:Al_2O_3 layers. 123

4-2 (a) Diagram of passive spiral laser architecture with main component parts labeled, (b) Diagram of passive compact s-bend laser with main component parts labeled. 125

4-3 (a) cross section of gain waveguide where orange = nitride, blue = oxide, and green = Er:Al_2O_3. Marbled region at edges denotes roughness (see below), (b) Fundamental TE mode in waveguide for laser center and pump wavelengths, (c) Overlap of the laser mode with the pump are of the gain medium. 127

4-4 (a) Loss vs. roughness width (width of the center region is equal to 50 microns minus two times the roughness width). Grey region represents expected waveguide gain. As can be seen loss in the Fundamental TE mode remains low while loss in the other modes is high, (b) Top view of trench width in roughened waveguide along with cross sections at two locations in the roughness. 129

4-5 Experimental transmission results for roughened waveguides: multimode beating is clearly seen for the case of no roughness (blue line), a roughness width of 5 or 7 microns (red and green lines) removes unwanted modes (as seen by lack of beating), but does not add to loss, while a roughness width of 10 microns (black line) causes extra loss. . 130

4-6 Material Dispersions 130

4-7 Total dispersion for the best (left) and worst (right) Al2O3 material dispersion. Anomalous (negative) dispersion is achieved in the best case for thick layer heights. For thinner layers or the worst case Al2O3 material dispersion total dispersion is normal but, in most cases, small. 132

4-8 Simulated bending loss vs. radius (in mm) for gain waveguide bend. Inset figures show modes at different film heights relative to bending radii (center of bend is on right). From left to right: 1μm film height and 3mm bending radius: mode is well confined to center of the guide, 1.6μm and 9.3mm radius: one can see that the mode is starting to slide toward the edge of the trench, 2.0μm film height and 15mm radius: mode is no longer confined by nitride guide, but rather by side of the Al_2O_3 trench. 134

4-9 Simulated bending loss vs. radius for nitride waveguide bend 135

4-10 Diagram of bend, where dark green is nitride waveguide and olive green is the trench. bend parameters 1-3 are described in the table below. . 136

4-11 (a). Diagram of loop mirror as drawn on the mask. (b) Comparison of coupled mode theory vs. FDTD simulations. 137

4-12 (a) Measured loop mirror reflection is in good agreement with theoretical predictions. (b) Simulated reflection for various wavelengths. Reflection percentages are similar across the 1530-1570nm Erbium gain band. Reflection at the 1480nm is lower, as desired, but could be further optimized. 139

4-13 (a) Taper type one where both nitride layers are tapered together. (b) Taper type two where each nitride layer is tapered separately. 140

4-14 Loss in dB vs. wavelength for round-trip coupling to SA. From [47] . 140

4-15 Raw data of best performing taper coupled to 3μm fiber for each taper type. 142

4-16 Measured and predicted loss vs. taper width for taper type one. As can be seen an extra 1-1.5dB of loss is seen in the measured results, though the trend agrees. Additionally, length has an effect in the measured results, with the longer device performing worse. 143

4-17 Diagram of Er trench transition structure (top) and Fimmwave™ modesolver (left) and transmission simulation of same (right) 145

4-18 (a) Diagram of passive spiral laser architecture, (b) photograph of spiral lasers on chip. 148

4-19 Green florescence from high power 1480nm pump outlines the laser structure. 149

4-20 (a) Ouput spectrum of cw laser as a function of pump power. (b) Laser output power vs. input power on chip. oscillation threshold is around 1.6W. This laser had a mirror reflectivity measured to be 30%. From [47]. 150

4-21 (a) RF spectrum of laser in Q-switched mode-locking regime, (b) Oscilloscope trace of same, (c) Zoomed in oscilloscope trace of pulse burst. From [47]. 151

4-22 (a) Diagram of compact passive s-bend laser architecture. 152

4-23 Schematic of compact s-bend laser to be pumped at (a) 980nm and (b) 1480 nm. 153

4-24 Measured transmission through passive s-bend laser architecture as a function of wavelength. 154

4-25 Fluorescence in compact s-end laser designed for 980nm pumping. A 3 way pump splitter (out of the frame) couples the pump light into each of the three gain sections. 155

4-26 Diagram of actively mode-locked s-bend laser configuration including MZ modulator for active pulse shaping and intracavity controls for comb stabilization. Phase-shifters are shown in gray, the intracavity heater in brown, the silicon in gold, and the nitride in red. 157

5-1 (a.) Mask layout of integrated heater, (b.) Phase change through heater vs. input DC power 162

5-2 Integrated Germanium Photodetector (a.) Layout of two different device variations, (b) Schematic of TIA drving circuit for PD, (c.) Output current vs, optical power out of fiber input to chip for 8x1μm (red) and 8x3μm (green) 3D integrated PDs with 5kΩ TIA resistance tested at 1530nm, (d.) Speed measurements of 8x1μm (blue) and 8x3μm (red) 3D integrated photodiodes with 5kΩ TIA resistance. 164

5-3 Plot of single sideband phase noise due to thermal fluctuations for ring variations on mask. 166

5-4 Diagram of PDH Loop. Blue denotes electrical components, and pink photonic ones. 168

5-5 PDH error signal in the low frequency (left) and high frequency (right) operating regions. Figure from [69]. 170

5-6 Simulation output of Pound-Drever-Hall Loop in the low frequency regime. Results previously shown in Figure 2-14. 172

5-7 Cadence layout of PDH photonics-only loop (top, variation with nitride rings used as high-Q reference cavity and integrated ring lasers), and photonics+CMOS loop (bottom, same variation). 173

5-8 (a.) Zoom in of Cadence layout of balanced detectors. Detectors shown are photonics only. Edge of pads for CMOS integrated balanced detectors which are interwoven are visible at right. (b.) Layout of CMOS integrated frequency divider loop consisting of SSB modulator and balanced detector. 174

A-1 Diagram of simulation flow for modeling silicon reversed biased diode plasma dispersion effect phase shifters. Figure previously appeared in [[36]. 180

List of Tables

4.1 Summary of measurements of compact Er:Al_2O_3 bend structures. . . 136
4.2 Summary of measurements of taper structures. 144
4.3 Summary component loss measurements. 155

Chapter 1

Introduction

1.1 Silicon Photonics

Silicon photonics is the field in which silicon integrated-circuit technology is used to transmit, manipulate, and process optical signals. It promises to bring the advantages of semiconductor process technology to the optical systems that have created the backbone of the telecommunications industry and that promise improved performance on a variety of electric systems from analog-to-digital converters (ADCs) to arbitrary-waveform generators (AWGs) [1, 2, 3, 4]. Now, in particular, is a time of great promise in the field of silicon photonics, as many of the basic devices needed for these systems have been demonstrated and optimized to the point where they have acceptable performance metrics.

There are a number of advantages to integrated photonics in general and integrated silicon photonics in particular. First, it leverages the huge advances made in semiconductor processing technologies over the last half century and brings them to the optical regime, thus, enabling small precision devices to be fabricated cheaply and in large number. Thus, silicon photonics promises to unite the discrete components currently used to create optical communication and other table-top systems onto a single chip and allow for affordable mass-production of these chips for use in products. Movement of optical systems on chip significantly reduces the size, weight, and power requirements of these systems, further widening the situations in which an ap-

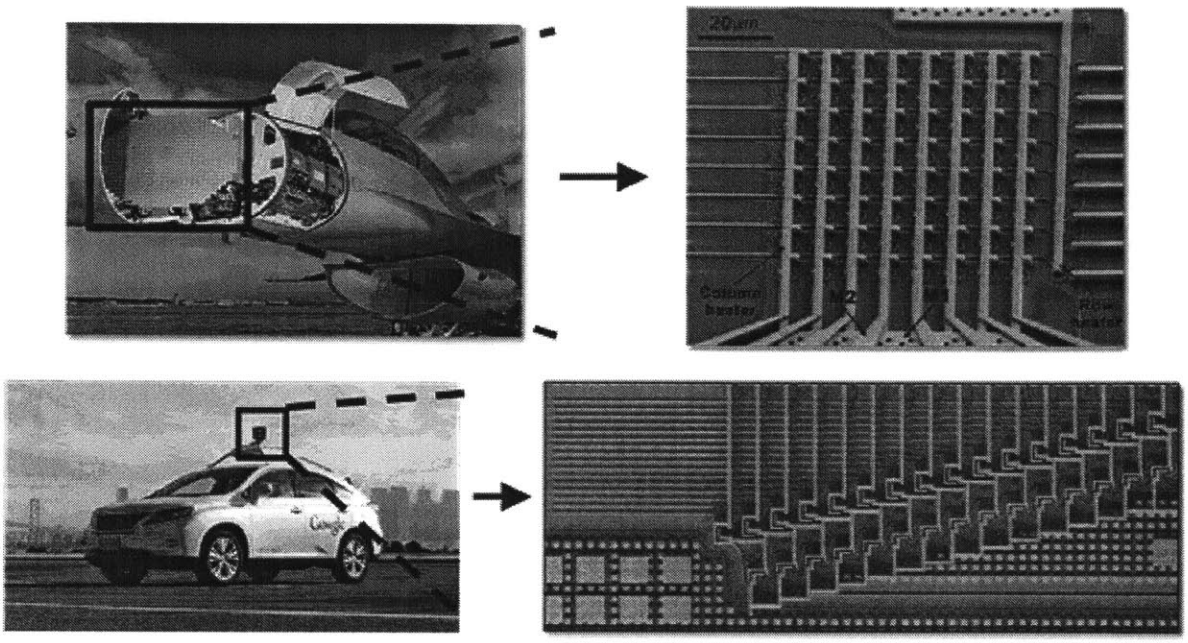

Figure 1-1: Silicon photonics provides a way forward for drastically reducing the size, weight, and power of large optical systems, including phased array antennas (top) and LIDAR systems (bottom). Figures from [5, 116, 7, 8], figure based off figure by Ami Yaacobi.

plications for which they can be used [4]. For example, silicon photonics is currently working to replace large microwave phased array beam steering and LIDAR applications, where giant antenna arrays and large systems are currently in use at long wavelengths, with small on-chip systems that work at shorter optical wavelengths as illustrated in Fig.1-1 [5, 116].

This miniaturization is made possible not just by the fact that small features can be fabricated, but also by the fact that optical devices of such small size can work. A key reason for this is the high index contrast between silicon and SiO_2, silicon's native oxide. This high index contrast means that optical mode areas can be kept small and guided through tight bends. It reduces optical power requirements, and can even facilitate nonlinear interactions, should they be desired [116]. While the high index contrast can lead to fabrication sensitivities, much recent work has been done to overcome this either through the creation of robust devices or improvement in fabrication techniques (see, for example, [9]). Moreover, the use of silicon allows for the

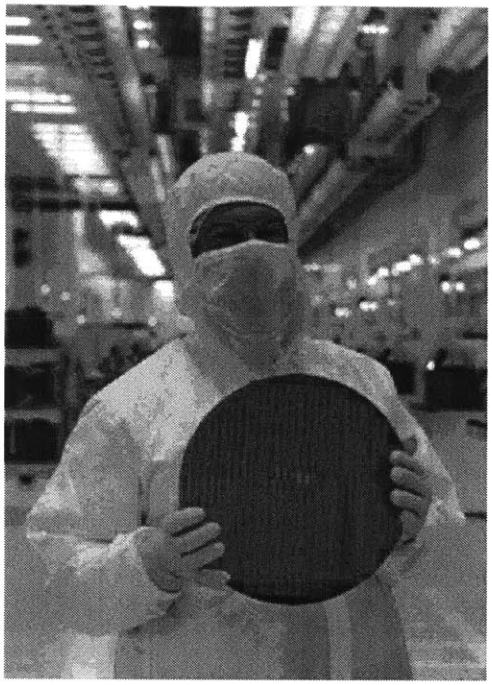

Figure 1-2: A worker on a CMOS silicon fabrication line. Image from [135].

creation of active devices, such as diodes, used for the creation of modulators (which transfer electrical signals onto an optical carrier), and the silicon process is materially compatible with the inclusion of germanium layers for optical detection (which turns optical signals into electrical ones). Importantly, Silicon is transparent in both telecommunications bands (1.3μm and 1.55μm) as well as at longer IR wavelengths, and its process is also materially compatible with the inclusion of nitride layers which are transparent in the visible as well as the IR, and are useful for nonlinear as well as quantum applications [10, 11]. Further, though its indirect bandgap prevents it from being used to create a light source, various laser sources have been successfully bonded onto or grown on silicon photonic chips [12, 13, 14]. Finally, the use of a silicon platform, the same material system used for the CMOS electronics industry, for photonics allows for easy integration with complicated electrical driving circuitry through either monolithic integration [15] or wafer bonding [126].

In the last decade and a half, much work has been done on individual silicon photonic devices, including modulators [9, 93], photodetectors [19], waveguides [125],

ring filters [21, 22], bonded lasers [23, 14], erbium doped lasers [24], vertical couplers and antennas [5], among many others. This has allowed more recent work to focus on using these devices to create systems, and to apply the technology to create state-of-the-art or near state-of-the-art integrated photonic versions of both existing systems and novel optical systems capable of surpassing existing system performance, such as WDM communication links [126], phased arrays [5], LIDARs [116], heterodyne interferometers [115], photonically-assisted ADCs [1, 128], AWGs [?, 3], sensing and optofluidics [129], and advanced signal processing applications [131], among many others. However, more work needs to be done to truly leverage silicon photonics' advantages in the systems world, as a complete silicon photonics platform needs to be established. This will be discussed more in Chapter 2.

1.2 Low Noise Microwave Oscillators

One application of particular interest is that of low noise microwave oscillators. There are many applications, from radar to coherent communications to global positioning systems, to precision noise measurement systems where low noise microwave oscillators are important [117, 118, 119]. Very low noise oscillators lower the noise floor in coherent detection and communication systems so that extremely weak signals can be detected, such as the Doppler shift from a person walking in the observation region of a Doppler radar system [119]. Oscillator noise performance is often characterized by giving the single-sideband phase noise at a given frequency offset from the carrier frequency —typically a 1 Hz offset is used in very low noise systems. For very low noise systems, the given frequency offset is usually a 1Hz offset. The single-sideband phase noise is equal to half the power spectral density of the phase error ($\mathcal{L}(f) = \frac{1}{2}\int_{-\infty}^{\infty}\langle\Delta\phi(t)\Delta\phi(t+\tau)\rangle e^{i2\pi f/\tau}d\tau$ where f is the frequency offset form the carrier and $\mathcal{L}(f)$ is the single-sideband phase noise [119]).

While very low noise oscillators have been demonstrated, these oscillators are often large and require cryogenic cooling or extensive temperature or cavity stabilizing circuitry [117, 122, 121]. They are also very expensive. The current state of the art is

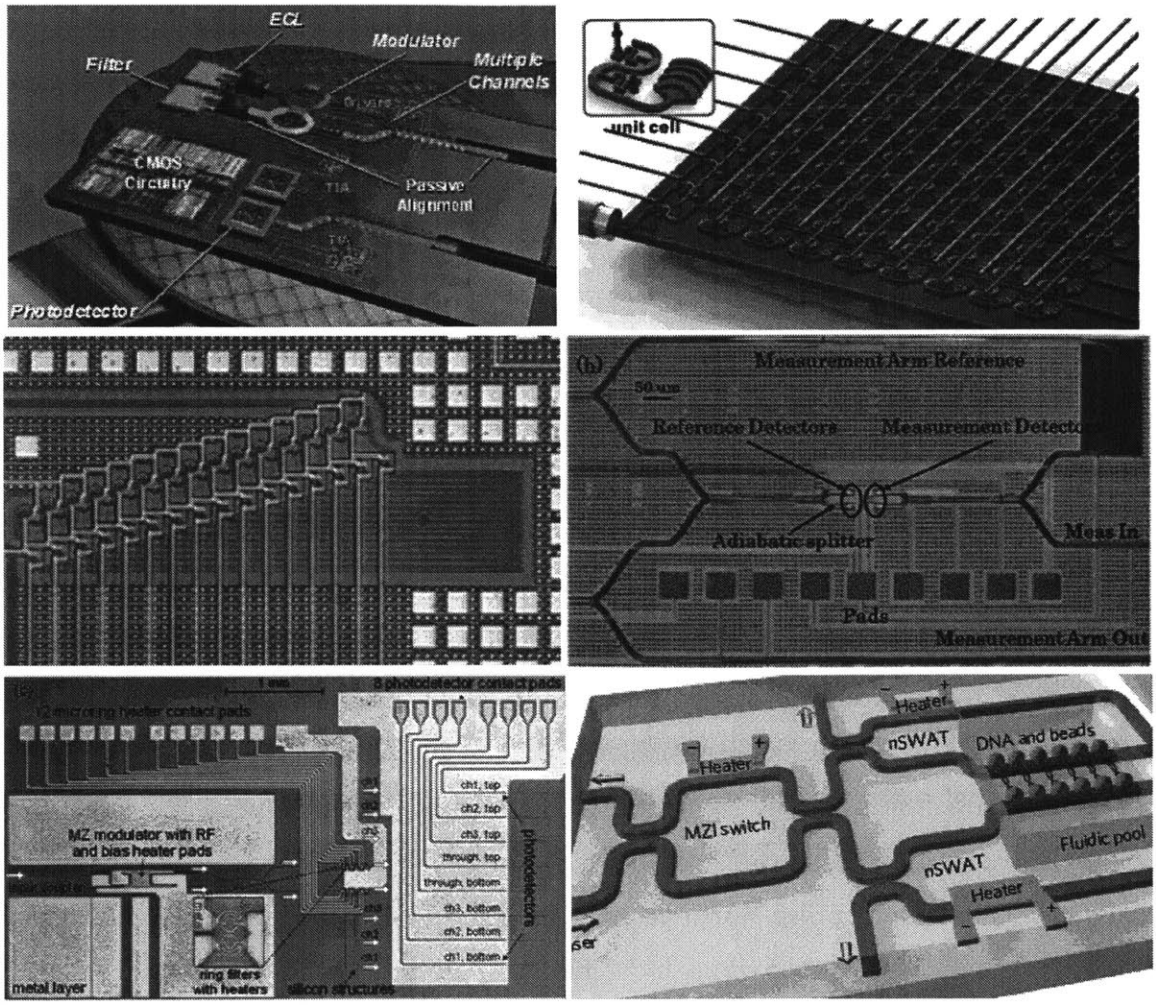

Figure 1-3: Silicon photonic systems on a chip. Clockwise from top left: 4-channel on chip WDM link [126], on-chip phased-array [5], integrated heterodyne interferometer [115], on-chip optical trapping for biomolecule manipulation [129], integrated photonic front-end for an analog-to-digital converter [1], on-chip LIDAR steerable antenna [116].

the Poseidon sapphire loaded cavity oscillator pictured in Fig.1-4.a., which achieves a phase noise of less than -100dBc/Hz at a 100Hz offset in a "compact" rack device or shoebox sized device [123, 124, 130]. This latter weights 14.5 pounds [130]. While this is a significant advancement over other systems, it is still too large to be included in, say, a hand-held radio or small system designed for field use. Additionally, it is desired to make an oscillator even less sensitive to environmental noise factors, especially vibrations, to truly allow for practical use. Cost is another consideration as a cheap, mass-producible component can be used in a much wider variety of systems. However, using sapphire cavity technology, it is not clear how oscillator size, sensitivity, and cost can be significantly more reduced.

Recent developments out of NIST on frequency dividers provide a possible path forward. In [118], a cavity stabilized continuous wave laser was used to stabilize a comb line of a mode-locked laser (MLL). Together with a lock on the mode-lock laser's carrier-envelop-offset frequency, this allowed for a highly stabilized MLL comb, which, when beat on a photodetector created a very low noise microwave signal (see Fig.1-4.b.). A single sideband phase noise of less than -100dBc/Hz at a 1Hz offset was achieved in this system at a microwave carrier frequency of 10 GHz — .which is the lowest phase noise that has been achieved to date at this carrier frequency [118]. At first, because this is a large, table top system, it may not seem to help in the quest for a small, robust, cheap, portable low noise microwave oscillator. However, because this system is based on optical components, system size, weight, and power requirements can be drastically reduced through on-chip integration — that is through movement to a silicon photonic platform.

1.3 Integrated Photonic Low Noise Oscillator

A proposed erbium silicon photonic integrated oscillator for generation of an ultra-low noise microwave signal, inspired by the NIST system, is shown in Fig.1-5 and is under development in our group [119, 120]. As the details of this system have been presented elsewhere [119] only a quick outline will be given here. The goal of this

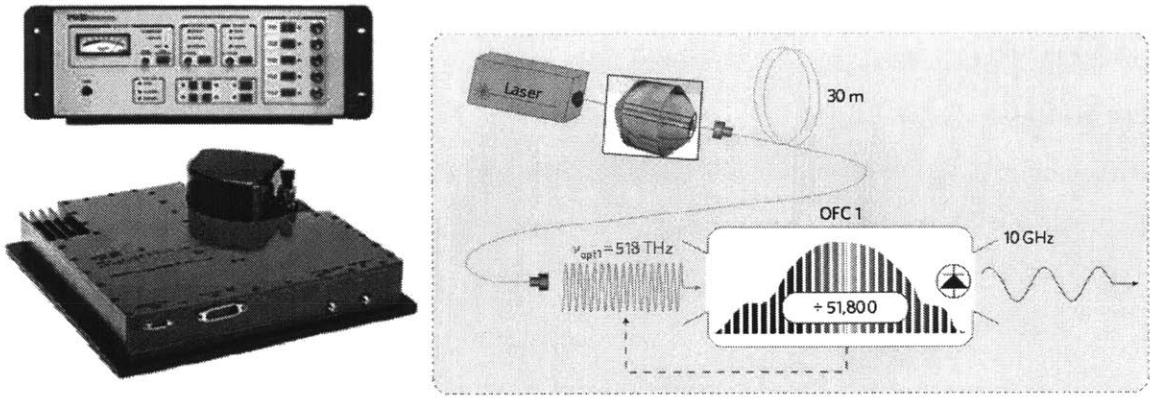

(a) (b)

Figure 1-4: (a.) Poseidon Scientific (now Raytheon) sapphire loaded cavity oscillators. Rack mounted sapphire loaded cavity oscillator (top) and Shoebox Oscillator ®(bottom), From [123] (b.) Schematic of NIST frequency divider oscillator. From [118].

system is to reach a single-sideband phase noise of -115dBc/Hz at a 1Hz frequency offset.

A schematic diagram of this oscillator is shown in Fig.1-6. The basis of the system is a continuous wave (CW) laser locked to a reference cavity, here a high-Q, passive, athermal ring resonator, through a Pound-Drever-Hall locking loop (bottom left box of Fig.1-6). This loop will be discussed further in Chapter 5. The output of the stabilized CW laser is split onto three different waveguide paths.

One of these outputs is used to transfer the lower frequency noise of the locked CW laser to a mode-locked laser (MLL) comb (upper left block in Fig.1-6). In this loop, a single line is filtered out of the output of the mode-locked laser, combined with the CW laser tap in an adiabatic 3dB coupler, and beat on a photodetector. More precisely, the two outputs of the coupler each feed one photodiode in a balanced detector. The frequency difference of the two lasers is reflected in the frequency of the signal output from the balanced detector. The use of a balanced detector not only increases the signal-to-noise ratio (SNR) by 3dB, but also serves to cancel the common-mode noise. The output is then fed back on the MLL in such a way as to adjust the repetition rate of the laser and lock the comb line to the CW laser

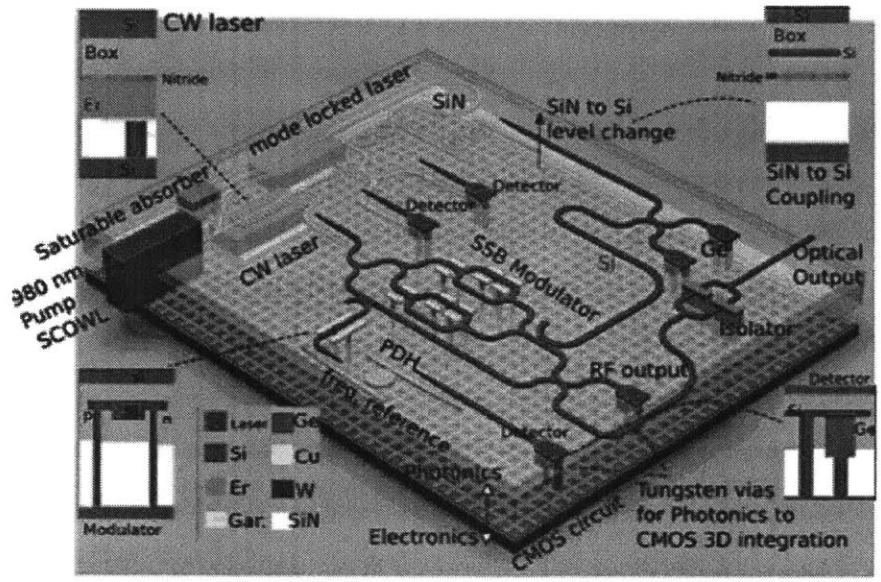

Figure 1-5: Proposed integrated low-noise microwave oscillator

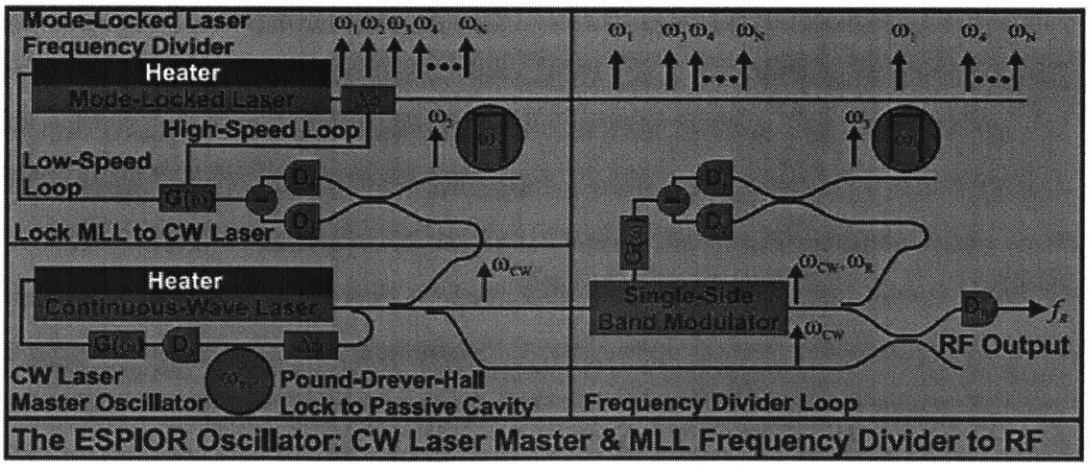

Figure 1-6: Schematic diagram of proposed oscillator [119, 120].

frequency, while preventing any change in the carrier-envelope-offset frequency. This loop, thus, locks the MLL comb in place. Because of the extremely low timing jitter and noise in the MLL, the output pulse train is very stable, and would create a very-low noise microwave signal if incident on an appropriate photodetector. However, due to the on-chip nature of our components, it is difficult to get enough power out of the MLL to create a strong microwave signal; we therefore add an extra loop (right box in Fig.1-6). This is where our system deviates from the NIST demonstration. Up to this point the system has basically mirrored the NIST system in integrated form, however, whereas, in the NIST frequency divider the output of the MLL is sent directly to a photodetector, here we have this extra step.

In this last loop, a second line is filtered out of the MLL comb, separated from the first by the desired microwave frequency. This line is then beat against one of the remaining CW outputs on a second balanced detector. However, in this case a single-sideband modulator (SSB) is used to frequency shift this CW laser output and lock *it* to the second filtered MLL line. The output signal from the balanced detector is fed back on to a voltage controlled oscillator (VCO) that serves as the SSB driver and controls this frequency shift. The frequency shifter CW laser line is then combined with the last, unmodified CW laser line and the two are beaten together on a final detector to create the desired microwave signal. This last loop serves to transfer the low noise of the MLL comb to the split CW laser output, while allowing the oscillator to take advantage of the larger power (and therefore lower shot noise) in the CW laser output. The ultimate goal is an RF microwave oscillator at 20GHz with a single-sideband phase noise less than -115dBc/Hz at a 1Hz offset, an RF oscillator output power of 100mW, a power consumption of less than 15W, and a substrate area requirement of less than $1cm^2$.

The work described in this thesis aims to support the development of the above oscillator from beginning to end: starting with the development of the integrated platform, layout, and systems codes necessary to fabricate the oscillator system, through the development of oscillator devices, and ending with some preliminary work on the control loop building blocks of system, in particular the Pound-Drever-Hall locking

loop. The layout of this thesis is as follows. In Chapter 2 we discuss the silicon photonics platform under development in our group. We summarize the photonics process and present work done to facilitate the creation of an integrated photonics + CMOS electronic process design kit. A key piece of this work is the development of a VerilogA modeling toolkit to facilitate joint photonic+CMOS modeling of large, complex systems, such as our oscillator loop. Chapters 3 and 4 discuss the creation of necessary devices in this platform. Chapter 3 focuses on the silicon layer and discusses the creation of modulators. A linear Mach-Zehnder modulator (MZM) is discussed first, and then two such MZMs are nested to form the single-sideband modulator needed in the oscillator. Chapter 4 discusses the creation of the integrated mode-locked laser. Chapter 5 builds on these devices and presents preliminary work on the sub-loops in the overall oscillator loop. Finally Chapter 6 provides a summary of the contributions of this work and a discussion of directions for future work.

Chapter 2

Joint Modeling and a Process Design Kit (PDK) for a Photonics Platform

Integrated photonics stands at a tipping point similar to that of the electronics industry a half century ago. In the past fifteen years most work has focused on the development of individual devices, aiming to demonstrate basic functionality and achieve performance metrics necessary for use in larger systems. During this time impressive results have been achieved for many individual devices including modulators [25, 9], photodetectors [19], ring filters [21, 22], integrated lasers [12, 13, 14], couplers [26, 27], on chip tapers [28], and many other components. Because they are integrated, many of these components promise equivalent or better performance than existing table-top components, while being smaller, lighter, and more power efficient. They are often less sensitive to environmental variations, such as vibrations, and have other specific use-case related advantages (such as higher confinement leading to smaller mode area and tighter power confinement yielding increased nonlinear effects) [29].

The state of the art in optical components is now advanced enough that the focus has shifted to leveraging them for use in optical systems. Stand alone integrated components have been included to improve the performance of table top systems (see, for example [30]), but the ultimate goal is fully integrated systems that both have state of the art performance and are cheaply mass producible (for example [1]). While many integrated systems have been demonstrated [5, 31, 129, 126], these systems have

either been limited to no more than a couple dozen components [1, 31, 129, 126] or, have had a simple architecture, which kept performance easily predictable and allowed for easy layout, despite having many components [5]. Design has usually been done in a pieced together manner with component simulations done separately and pieced together and together with the outputs of electrical driving circuity (when there has been any electrical driving circuity) roughly at the end. In order to be able to design and layout truly complex systems, especially ones that include electrical drivers, a design platform is needed. This chapter will discuss such a platform.

Because the development of the photonics industry really parallels that of the electronics industry and because many silicon photonics devices are even fabricated in the same foundries and clean rooms that make electronic devices (as is the case in this work among many others), it makes sense to look to the electronics industry to see how to design such a platform. One key idea that we can take from that industry is the idea of a multi-project run (aka a multi-project wafer service). This is where a foundry combines designs from many different organizations (small companies, students, universities, etc...) and fabricates them all on the same wafers using the same mask sets. Because foundry services are extremely expensive, putting many projects on the same run allows for the costs to be spread out between groups and greatly reduces the barrier to entry. In order to make the foundry services accessible to these groups, foundry companies create and distribute process design kits (PDKs). PDKs aim to give designers all the information that they need in order to design into the foundry process including fabrication information (available material layers, design rules that need to be followed) and standard device libraries. They will be discussed in more detail in the sections below, as we discuss the creation of the photonics version of such a PDK.

There have been some forays into providing silicon photonic multi-project runs, including [32], which has since ceased operation, [34], which does not offer light source integration or a complete design kit (no co-simulation with CMOS is available) and [33], which is a custom design studio which does in house fabrication. A slightly out-of-date review of foundry services is provided in [35]. But this area of the field is in its

infancy. In our group we have found the need to develop equivalent services for use in house and we coordinate internally across many projects, designers, and groups. Our group's wafers are in some ways their own mini one organization multi-project runs.

2.1 Introduction to the Silicon Photonics Process and Platform

In order to set up a multi-project run for photonics, it is necessary to create a photonics platform. A true platform consists of a fabrication process through which devices and systems can be physically realized as well as a process design kit (PDK), usually including a device library, that allows designers to truly utilize the fabrication process. That is, the PDK contains all the relevant and necessary information for designers to create and submit designs to the fabrication facility for production.

The first part of the platform, therefore is the fabrication process. The devices discussed in this thesis were fabricated at the 300mm CMOS foundry at the College of Nanoscale Science and Engineering (CNSE) at the University of Albany in Albany New York using a process they developed in collaboration with our group, and Professor Vladimir Stojanovic's group at Berkeley. Some early work on this process was also done in collaboration with MIT Lincoln Laboratory. The process will be outlined here. It consists of a photonics wafer 3D bonded to a CMOS wafer via through oxide vias (TOVs). The CMOS wafer is made using a standard CMOS process on a 65nm node and contains electronic driving and other desired circuitry. The photonics wafer, which contains the photonic part of the designs, is fabricated using the developed, custom process for silicon photonics. It consists of both active and passive photonic devices as well as a trench in which erbium doped Al2O3 is deposited to act as a gain medium for on chip lasers. Both photonics-only (without CMOS) and 3D-integrated (with CMOS) wafers can be returned from the fabrication facility as desired.

Fig. 2-1 shows a diagram of the resulting layer stack for both the photonics-

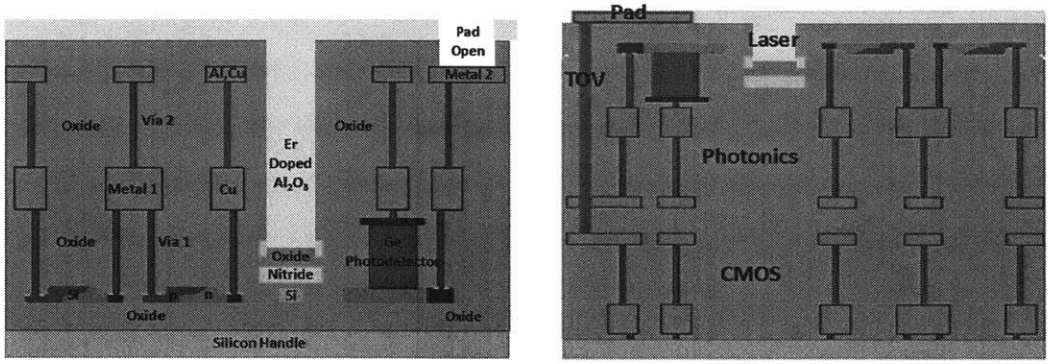

(a) (b)

Figure 2-1: (a) Photonic-only chip layer stack. All layers other than the Er Doped Al_2O_3 (light pink) are fabricated at CNSE Albany. The Al_2O_3 is deposited at MIT. (b) The full 3D integrated layer stack showing photonics wafer flipped and wafer bonded to the CMOS wafer (From [47]). The only difference among the photonics layers compared to (a) is that the laser trench now comes in from the back (silicon handle side) instead of the the front (metal side) to allow for post-3D-integration deposition of the Al_2O_3 at MIT. As before, the rest of the fabrication is carried out at CNSE Albany.

only and 3D-integrated cases. As shown, the platform centers around a silicon layer patterned to provide waveguides for routing the light. The high refractive index of silicon allows for tight bends, high confinement, and thus, small devices, while the infrastructure developed over many years for the electronics industry and the fact that silicon is transparent at telecom wavelengths allows for low loss. (The former drives down losses from roughness, etching, and patterning, and the later yields low material loss). The silicon layer has both a full height and a half-height, ridge etched section. The ridge etch serves several purposes, among other things: it allows for the creation of rib waveguides, which in turn allow for easy contacting to active devices [36]; it can be used to create polarization rotating devices which convert between the TE and TM modes of a waveguide [37]; it can be used for symmetry breaking in vertical couplers to allow for more than 50% upward coupling [?].

The silicon layer is also used to support active devices. The fabrication process contains four doping masks; one each for a high and a low doping of acceptors and donors (n+, p+, n, and p). There are also two doping impants to support the photodetectors. The low dopants (n and p) are used where the doping overlaps with

the optical mode, allowing the creation of an active device while keeping optical losses low. In early work, low doping concentrations were quite low (on the order of 10^{17}, or sometimes even high 10^{16}) [38]. However, because higher doping levels allow for faster and smaller diode devices with minimal loss or heater performance penalties, doping levels used in the literature have been continuously increasing [9] and those picked for use in this process were chosen to be higher – on the order of $1 - 5 \times 10^{18}$. The high dopants (n+ and p+) included in the process contribute a prohibitively high optical loss and are used away from the optical mode to allow for lower resistance, high speed contacts.

While they can be used however desired, the doping impants in our process are generally used to create two types of active devices: diodes and resistors. The resistors serve as heaters, shifting significantly, but at low speed, the index of refraction of the waveguide through the thermo-optic effect, and thereby creating a phase shifter. In the thermo-optic effect, the index of the waveguide is changed according to $\Delta n = n_0 + (dn/dt)\Delta T$ where $dn/dt \approx 2 \times 10^{-4}/^oC$ in silicon, and where the loss change is minimal [39]. Such resistive heaters are useful in setting the bias points of Mach-Zehnder modulators [36] and ring modulators and filters [22, 40], and for controlling phases in antenna elements of phased arrays where their small size is essential [5].

Diodes structures aim to take advantage of the faster, but comparatively weaker, plasma dispersion effect (PDE). Here, a change of carrier concentration creates a change in loss and refractive index of a material through its effect on the material band structure. The PDE can be theoretically modeled through the Drude equations [41], and has also been experimentally measured by Soref et al [42]. In the literature an empirical fit to Soref's data for index of refraction, n, and loss coefficient, α, given by:

$$\begin{aligned} \Delta n &= -8.8e - 22\Delta N_e - 8.5e - 18\Delta N_h^{0.8} \\ \Delta\alpha &= 8.5e - 18\Delta N_e + 6e - 18\Delta N_h \end{aligned} \tag{2.1}$$

is often used. However, this equation does not fit the data precisely and does not

account for wavelength. Because of this often other fits to the same data set are used when more precision is desired [93].

A p-n diode is the simplest, non-metal, device useful for changing carrier concentrations in a waveguide. It can be operated in one of two regimes: forward or reversed bias. In the forward bias regime, minority carriers are injected across the p-n junction in large numbers allowing for moderately large phase shifts ($dn/dV \approx 1 \times 10^{-4} - 1 \times 10^{-3}$. The large change in carriers also leads to an added loss, which can be helpful or harmful depending on the application. Unfortunately, forward bias is still relatively slow; operation is limited to a few GHz at most by the minority carrier diffusion time necessary to remove injected carriers [43, 36]. Some work has been done on decreasing this time by adding defects that can serve as recombination centers. This helps but both also adds optical loss and harms sensitivity at all speeds as it limits the number of carriers that can build up. In tightly confined silicon wire waveguides it is usually the recombination time at the silicon-oxide interface on the **edges** that dominates the minority carrier lifetime [43].

The reverse bias regime is the opposite: fast but relatively weak. Here the change in carriers, and thus refractive index, is caused by a change in depletion region width with applied bias. Since the carriers shifted are majority carriers, it is the majority carrier relaxation time or RC time constant, whichever is longer (usually the RC time constant), that determines the speed. Unlike their forward biased counterparts, reverse biased devices can operate in the 10s of GHz range. However, since it comes only from the change in depletion width, the magnitude of the carrier concentration change is limited to the doping concentration. This means that reverse biased diodes are weaker in terms of effect size, with $dn/dV \approx 1 \times 10^{-5} - 1 \times 10^{-4}$. The sensitivity presented in Chapter 3 of this work of $dn/dV \approx 1.3 \times 10^{-4}$ is actually very high for a diode operated in reverse bias. This is achieved in part by making sure that the peak of the optical mode is spatially lined up to the edge(s) of the depletion region so that the mode sees the maximum difference.

In all, the doping implants for these two types of device — resistors and diodes — were optimized by the author and are shown in Fig. 2-2. Optimization of vertical

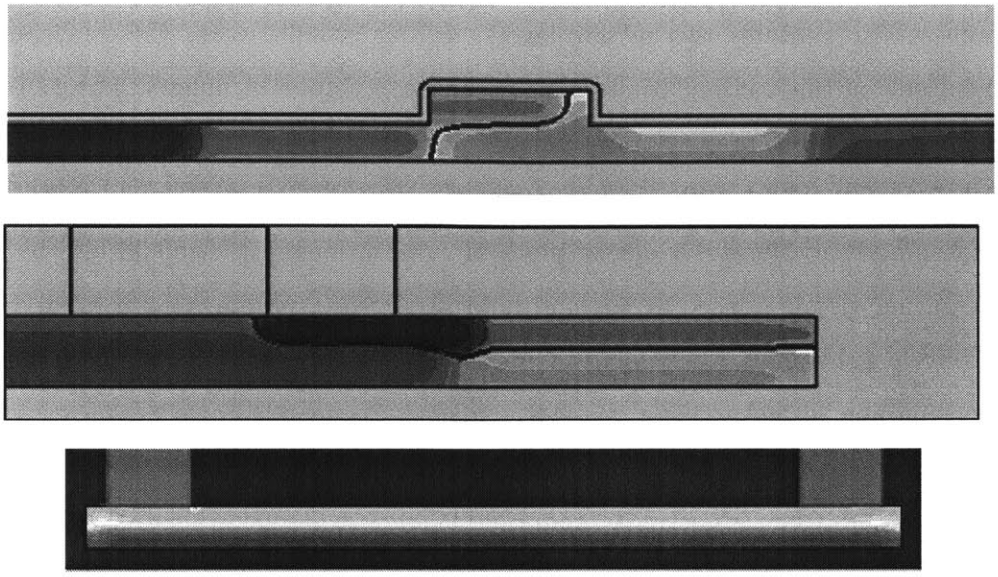

Figure 2-2: Example final doping distribution showing density of p-type (blue) and n-type (orange,red) dopants in a MZ diode modulator (top), a disk diode modulator (middle) and an integrated heater (bottom)

junction diodes, for both rib and disk phase shifters, aimed to maximize phase-shift while maintaining high intrinsic speed. In both cases the devices are not limited by thee intrinsic RC time constant of the junction, but rather by the optical resonance width in ring and disk devices or by the lumped electrode capacitance for linear phase-shifting devices. A vertical junction was chosen to allow for maximal phase shift per length in the waveguide and for easier, more symmetric contacting in disks. The dopants were also chosen to provide acceptable resistances for thermo-optic heaters in the range of 400-2000Ω.

While the silicon layer is the workhorse of the platform, other layers must be integrated for the platform to be truly versatile and capable of creating a wide variety of systems. Photodetection is accomplished in an integrated germanium layer (germanium absorbs in the 1550nm telecommunications band we are interested in). The germanium is grown in specified trenches and a vertical p-i-n diode is incorporated, thus enabling the creation of reversed-biased photodiodes. Metal via layers support electrical contact to both the silicon and germanium layers and two metal layers on the photonics side allow for easy routing of signals.

The process also contains two nitride layers. In our group's work, these layers are mainly used for the guiding of on-chip lasers (see Chapter 4 below), but can also be used for a number of other purposes. For example, unlike silicon, nitride is transparent at visible wavelengths and therefore this layer can route (etc..) visible signals, something that is useful in many quantum applications [11]. It could also be useful for creating visible phased arrays and holograms (e.g a visible version of [5]). The large wavelength range over which nitride is transparent additionally makes it good for sensing and lasing applications [44]. Nitride's higher nonlinear coefficients making it useful for many on-chip nonlinear optics applications, while its lower absortion makes it useful as a laser cavity material [45]. For example, others have used dispersion matching in nitride resonators allows for FWM generated frequency combs [46]. Moreover, nitride has a lower thermo-optic coefficient ($dn/dT \approx 4.5 \times 10^{-5}/^{o}C$ [39], making nitride structures more robust against thermal fluctuations from the environment, on chip and off. Additionally, because it is used in the laser cavity, nitride structures can be matched to laser structures to allow for automatic compensation of thermal drift.

However, as mentioned, the primary reason for the inclusion of the nitride layers is to serve as the guiding layer inside laser cavities. Our platform supports Er doped glass waveguide lasers integrated on-chip [24, 47]. One nitride layer serves as a waveguide layer for both the laser light and the pump light (either a 980nm or 1480nm pump can be used as our nitride is low loss and transparent at both). A second nitride layer serves as an etch stop for a large trench. Outside the trench, this layer serves as a second nitride waveguide layer. The trench etch denotes where Er doped Al_2O_3 will be deposited. All fabrication up to this point is done at CNSE, however, the Er:Al_2O_3 deposition is carried out at MIT as a backend step. A blanket film is deposited over individual chips. This film fills the trenches to create the lasers. The oxide encapulation over the rest of the devices prevents their operation from being affected. The Al_2O_3 is scratched off the pads and polished off the chip edges where necessary. In the 3D integrated case, once the photonic wafers are fabricated they bonded to the CMOS wafers to create a single stack. If Er deposition is wanted

on these wafers, it is carried out on the back after bonding (see Fig. 2-1).

Obviously there are other layers (nitride liners, pad openings, fiber trench etches, etc...) present in the process, but the above list highlights those most important for optical systems. Of course, having such a process is useless if there is no way for designers to access it. This is where the other part of the platform — a process design kit (PDK) becomes important.

2.2 A PDK for photonics - Layout and Verification

Of course, having such a process is useless if there is no way for designers to access it. This is where the other part of the platform — a process design kit (PDK) becomes important. The goal of a PDK is to give a designer all the information that they need to design a chip in a given fabrication process and send it out to be produced. It usually consists of written information on the process, including layer definitions and design rules (rules governing what the foundry can and cannot make), layout information, verification decks, and a library of standard devices in the process including simulation models and device layouts.

During the past decade plus of silicon photonic design, such kits have been bare bones to nonexistent. A lot of fabrication has been carried out in house, often by designers themselves. Designs are simulated by hand (if at all) in tools (e.g. FDTD) separate from those used to create mask layouts (e.g. autocad), which are also often drawn by hand. However, as system complexity and size increases such a design flow is no longer sustainable, and the need for a new infrastructure becomes acute. The integrated (electronic) circuits industry has been dealing with this problem for decades now, and has developed a variety of constantly improving CAD tools for the creation of large and complex systems. The integrated CMOS driving circuitry is often developed completely — modeled, laid out, and verified — in these tools. Because many aspects of design and layout for silicon photonics are similar, if not directly analogous, and because the electronic parts of joint CMOS+photonics circuits already exist in these tools, it makes sense to "simply" add photonic capabilities into a pre-

existing CMOS design kit. In our case the software tool used is Cadence®Virtuoso®.

To integrate photonics into the design flow, we add in the corresponding items to each electronic part: the photonic layers to the layer definition file, the design rules for the photonic layers into the design rule deck, the photonic devices to the standard cell library, etc. Keeping things as analogous as possible allows us to use the same data structure, user interface, phyiscal layout, verification tools, and simulators that are already used to design the electronic parts. We can thus design photonic systems just as if they were CMOS ones. Furthermore, the unified environment allows many designers to work collaboratively on a project in a synchronized database, and for the optics and electronics designers (if different) to easily see both sides of the system and the interface. It also allows for explicit verification of that interface.

A combined photonics+CMOS PDK (PCPDK) consists of two main parts: a physical architecture and a simulation architecture. The physical architecture allows for the chip to be laid out and for verification to be preformed. A lot of the physical architecture that we use, especially in regards to the layout, was begun as part of a previous project and grew out of the work done in [48]. It has been expanded on, particularly in regards to LVS verification, and will be summarized below. The simulation architecture allows for the system performance to be predicted and was developed as part of this work. It will be described in detail in the following section.

2.2.1 Layout

The first part of the PCPDK's physical architecture allows for the layout of the completed chip design. This includes the definition of the process layers, the parameterized layout cells for the photonics devices, and layout verification (both design rule and layout vs. schematic checking).

The most basic component of the PDPCK's physical architecture is the layer definition. In an abstract sense the layer definition consists of a list of the mask layers that need to be sent to the fabrication facility and their purposes (drawing of mask shapes, labeling, etc...). In a formal sense, the layer definition consists of two files. One is the techfile which maps the cadence layout layers onto the gds layers

used in the foundry mask files. The second is the "display.drf" file which governs how these layers appear (color, pattern, borderline) in the layout GUI of Cadence ®Virtuoso ®. Since we start from a set of layer files used in the existing CMOS PDK, we simply need to add in the layers for drawing and labeling on the photonics wafer and in the 3D integration steps. In addition, we create a set of export files to stream out the final mask(s) once the layout is complete. For the full process, the final exported file contains the layers for the CMOS wafer, those for the photonics wafer, and those for the TOV connections. Once at the foundry, these layers are broken out, the appropriate masks cut, and sent to make their respective wafers.

The next part of the physical architecture consists of the parameterized layout cells (pcells) [48]. These layout cells are the physical representations of the devices that will be placed on the mask. More simply, they are parameterized drawings of the basic system components that the designer can place and modify to their needs. In CMOS PDKs these cells correspond to vias, transistors capacitors, resistors, and etc, as well as some smaller common circuit blocks such as current mirrors and ESD protection circuits [49]. The electronic part remains unchanged in the PCPDK, and we add to it a set of photonics pcells created in house. On the photonics side, such pcells include waveguides, couplers, photodetectors, pads and phase-shifters as well as small systems like modulators and filters. The joint photonics+electronics pcell library, thus, allows for the complete, joint layout of all the mask layers used in a photonics+electronics system

Much like in the electronic case, where the electronic pcells are designed using traditional lower-level-physics tools (such as sentarurs, atlas, or experimental data), the photonic pcells are also designed using traditional the low-level-physics tools, only in this case those relevant to optics, such as FDTD, Fimmwave, modesolves, and Sentaurus (for active devices). Pcell layouts of the resulting designs are then created in Cadence ®Virtuoso ®using the SKILL coding interface.

Once pcells for the most basic cells have been created, they can be combined to create higher level cells. Those composite cells can, in turn, be combined to create yet higher level cells. In this way a system can be built up from basic cells (couplers,

phase-shifters) to full devices (modulators) to complete systems (communications links). An example of the such a process is shown in Fig 2-3. As additional photonic devices are designed, their pcells can be easily added into the device libraries for use in future systems. Similarly, because of the layouts are created using codes, existing pcells can be easily tweaked to reflect device optimization and experimental results. Moreover, as can be seen in Fig. 2-3, because the photonics and electronics pcells exist in the same set of libraries, the optical and electrical cells can be laid out jointly. Driving circuitry can even be included in lower level device blocks if desired — the layout contains all the layers in the 3D stack. One can turn layers on and off to looks at just the photonics devices, just the CMOS devices, or, importantly, just the connections between them. This makes verifying and trouble shooting the whole system considerably easier than if the two halves are designed separately and brought together later or only at the fabrication steps, and thereby minimizes errors. That joint layout can be accomplished in the same SKILL (and GUI) environment, makes it easy to do.

Note that because we have added the photonic pcells into Cadence®, the electrical pcells from the CMOS foundry do not need to be transferred into any non-standard platform. Moreover, the full design, including optical parts, can take advantage of Cadence's ®verification and simulation framework.

2.2.2 Verification

The verification framework forms the final part of the physical architecture in the PCPDK. While manual checking may be sufficient for small or repetitive devices and systems, (e.g.[1, 5, 31, 129, 126]), as systems scale up in both size and complexity, it becomes impractical and unreliable. Thus, in any large scale design, automated checking tools are necessary. Two tools that have proven themselves to be particularly useful in the electronics realm, and that translate directly to the optical one, are design rule checking (DRC) tools and layout vs. schematic (LVS) checking tools. These have been adopted to our joint PCPDK.

Design rule checking operates on the physical mask layout of a design and ensures

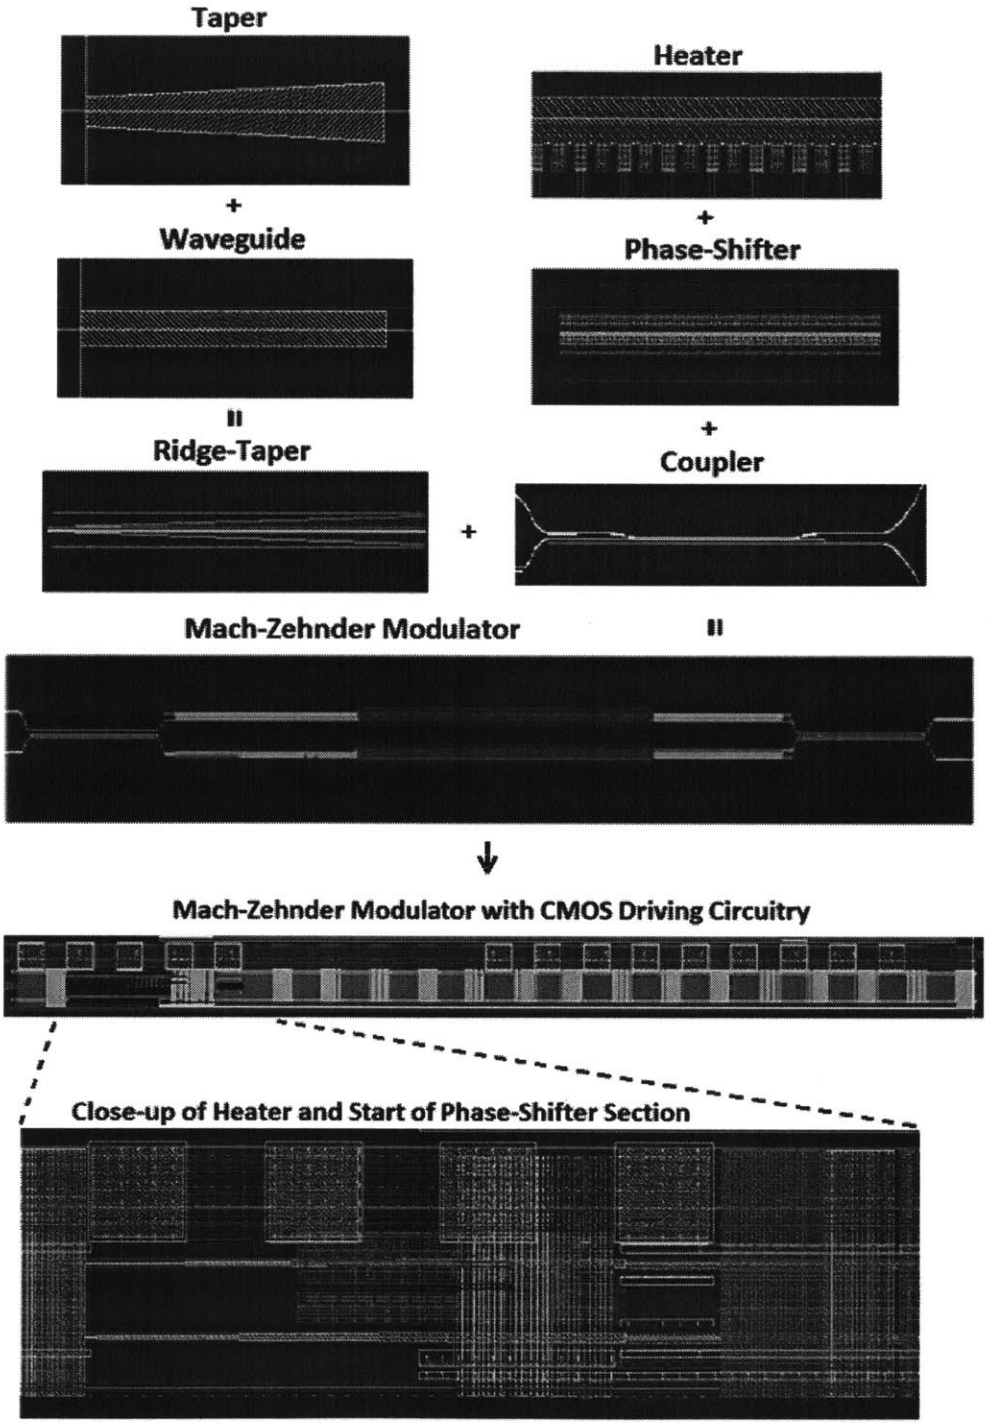

Figure 2-3: Full Silicon MZM with CMOS driving circuitry is laid out in cadence. Simple shapes (such as basic waveguides and tapers) are combined to create more complicated structures (such as the ridge taper shown), these more complicated structures (ridge taper, coupler, phase shifter, heater) are combined to create the full photonic layout. The full circuit layout is added to the photonic layout. Thus, the photonic and electronic pcells together create the layout of the full system. Designs therefore have the ability to include all 100+ photonic and CMOS layers, allowing the full system to be viewed together. A zoom in of the heater and heater driver section of the full layout is shown at bottom.

it follows the design rules provided by the foundry. These rules govern what the foundry can and cannot physically produce and include things like minimum and maximum line width and spacing, density rules, and rules governing layer overlaps. Each mask layer, or combination of mask layers, has a different set of rules to ensure that yield targets are achieved and equipment isn't damaged. Modifying this rule deck is straight forward - the rules for the photonic layers as well as the 3D integration layers are simply added into the existing CMOS DRC deck. This last set of rules — the rules governing the interaction between the electronic and photonic layers can only be done with a merged PCPDK as it requires both sets of layers simultaneously. (Note: because we are often working on just the photonic or CMOS side, we also split these rules out into two distinct desks to allow for faster subchecks before putting the whole chip through together).

Merging layout vs. schematic checking is more complicated. LVS checks verify that the device instantiations on the layout correspond correctly to a desired simulation schematic. This guarantees that if a designer draws a system schematic, simulates and optimizes it, and then lays it out; the final layout matches the final schematic with no missing or extra devices and connections and no crossed wires. In traditional electrical PDKs, this extends all the way down to recognizing the devices as drawn on the mask, that is determining if something is a transistor or capacitor or other device only from the layout. The ultimate goal is to have the LVS checker able to do this on the photonic side as well; to correctly identify a modulator, photodetector or other device. However, we have yet to enable layout extraction to recognize photonic devices. (Note: at present, younger students in the Stojanovic group are making good progress). When the designs created in this thesis underwent LVS checking, only the electronic devices were fully extracted. On the photonic side, we currently attached labels within each pcell to the contacts of the photonic devices. The labels then tell the LVS what the connection is to (for example: the positive side of a modulator or the negative side of a photodetector) and allow it to check that the connectivity between the photonics and the circuitry is complete (no floating connections) and accurate (no shorts, no modulator drivers connected to photodetectors, devices con-

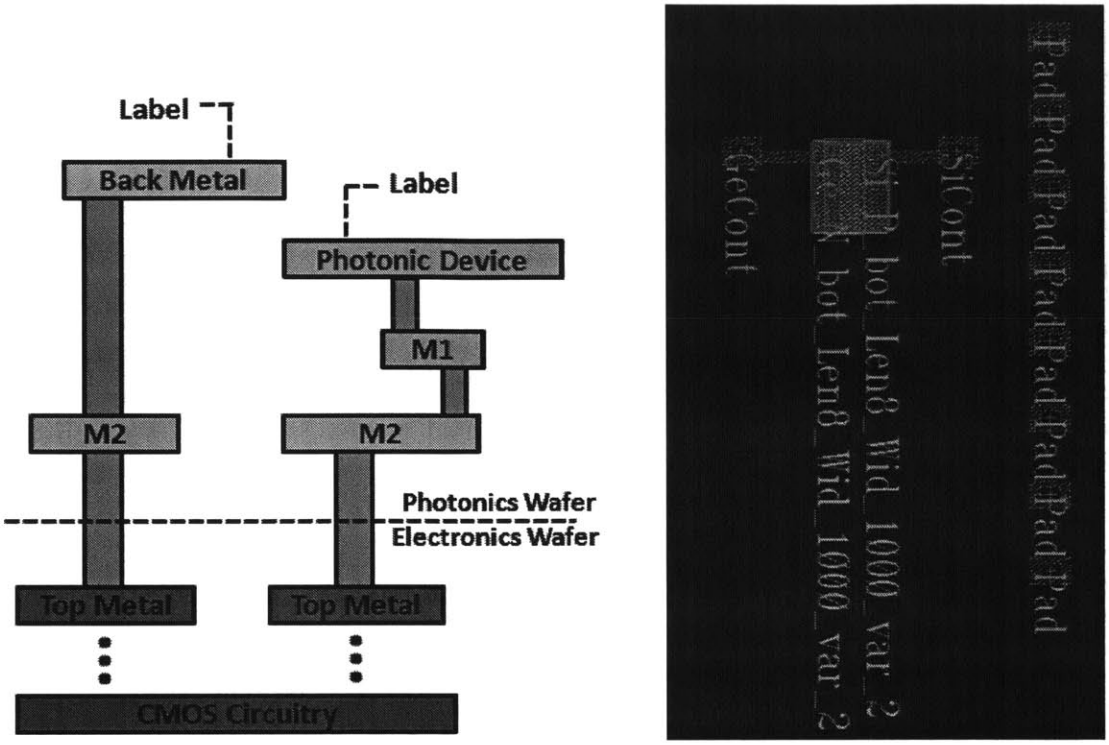

(a) (b)

Figure 2-4: (a) Schematic of labeling for connectivity for LVS verification (b) Example photonic layout including labels

nect to pads where appropriate, etc...). A diagram of this labeling process and the resulting layout it shown in Fig. 2-4.

These three parts: the joint layer definitions, pcell library, and verification infrastructure form the physical architecture part of the PCPDK and allows for the physical realization of joint photonics-electronic systems on chip. However, it is not enough to simply be able to make such systems, one must also be able to design them. Thus, the second part of the PCPDK allows for joint simulation of the electronic and photonic parts of an integrated optical circuit.

2.3 A PDK for photonics - VerilogA System Modeling

As mentioned above, as heterogeneous electronic-photonic circuits increase in scale and complexity, a scalable modeling framework becomes necessary to correctly anticipate system behavior. Such a framework is not necessary for simple systems such as a single ring resonator or simple photonic link, just like it is not necessary for simple circuits such as band pass filters and current mirrors; the behavior of these systems can be adequately captured with simple analytic equations, rough numerical simulations, and basic calculations. However, as the electronics industry made the transition from a few components on a breadboard to millions (and then billions) of components on an integrated chip, and from a situation where components could be easily swapped if a design tweak was needed to a situation where they were set in stone (or silicon) once a chip went out, such simple models were no longer sufficient. The same is true for heterogenous circuits that include photonics as they scale.

In particular, there is need for a joint system simulation code. Mature specialty software such as MEEP, Lumerical, Fimmwave ®, Sentaurus ®, and custom home-built scripts exist and are widely used to simulate photonic components at the physical, device level. These programs mirror programs like Sentaurus, Atlas, and HFSF that are used to model electronic devices, such as MOSFETs, at the physical, device level. Such programs allow for the maximum amount of detail and flexibility and the most accurate results, but are too computationally expensive to be feasible, let alone practical, for simulating the behavior of hundreds, thousands, or millions of devices, especially over "long" timescales. To solve this issue, the electronics industry has moved to translating the results from such detailed physical simulations, via careful approximations and assumptions, into simple governing equations for devices and using them in a separate system simulation in a different programming framework — often in Verilog or Spice. This allows necessary and fundamental behavior to be captured without requiring as many input parameters to be provided or the full physical complexity to be simulated, and, thus, allow for faster simulation times while still

letting the designer change key device parameters. Such models already exist in the Cadence architecture for the CMOS electronic devices, and tour goal is to supplement them with photonic device models, allowing for the high-fidelity simulation of heterogenous electronic-photonic circuits in Cadence itself.

Previous work has attempted to provide such behavioral marcomodels for photonic devices. In fact, several simulation tools or individual models have been presented [50, 51, 52, 53, 54]. Only some of these [50, 51, 52] are able to capture transient behavior of the photonics, handle frequency shifts and model both amplitude and phase effects (usually while removing the optical carrier for faster simulation time). Of this subset of tools all are written as standalone software and are therefore difficult to adapt into existing circuit design and layout infrastructure. It is also, therefore, difficult to modify or add to existing devices. For example, OptiSpice [51] is an excellent tool that correctly captures transient effects as well as amplitude and phase effects, but is written in as proprietary standalone software suite.

VerilogA is already widely used by the circuit design community to model analog devices in continuous-time and has been integrated into leading electronic circuit design platforms (including Cadence ®) [55]. There have been some attempts to write photonic models into VerilogA, but all have fallen short of designer needs. [53] provides a basic framework for simulation at fixed wavelength, which is sufficient for simple links, but cannot organically capture the full behavior including the phase phenomena, cross-talk, etc. of more complex systems and devices like single-sideband modulators, coherent optics and wavelength division multiplexed waveguide links. Here a higher fidelity representation of the underlying physics is needed. [54] presents a outline of a complete PDK similar to that created here, but the VerilogA models presented are overly simplistic for the type of models needed in many systems. For example, they pass around the output power and phase of the signals, instead of real and imaginary parts. Working in power and phase will create convergence issues (e.g. what is the phase when there is no power?) and will prevent appropriate modeling of interference based systems such as rings. Indeed, they do not show any ring or similar systems in their work. They also do no appropriately account for the interaction of

refractive and group index. The waveguide model simply outputs the input amplitude and phase delayed by a specified amount, but in reality the light output at the end of the waveguide will have an additional phase shift due to the waveguide refractive index. This effect is not captured and will lead to incorrect phases. Further they say that they are working with the optical envelope, but no discussion is made of the proper way to handle the carrier signal, making this almost a de facto single wavelength code, with the same problems mentioned above. On a more minor note, too few optical variables are parameters, making it difficult to model complicated systems with many components (for example, there is no way to change refractive or group index (which as mentioned above are not used properly anyway), modulator phase-shift vs. voltage is hard wired, etc...). There are other similar problems, and citen46 seems to be the most recent publication from this group. Finally, there has been some work on models for individual devices, including [56, 57]. Simulation blocks for single components can be very useful in certain simulations but do not provide flexible simulation frameworks in and of themselves. Moreover, it is desirable not to need to calculate the transfer functions of even moderately complicated devices by hand and program them in — it would be much nicer if such things could be organically generated during simulation.

As described in further detail below, we have aimed to create a fully functioning simulation toolkit written in VerilogA, that combines the strengths of previous work. This toolkit correctly simulates both the amplitude and phase of optical signals, as well as optical-electrical interactions, and does so without any need to specify optical carrier frequency or other "magic" parameters. It organically captures transient behavior and transfer functions of composite components without need for further input. It is also fully compatible with the Cadence®simulation environment where electronic simulation models already exist, thus preventing the need to import or rewrite provided electronic models in another simulation platform. Most of this section first appeared in [70], a paper in preparation to be submitted to Optics Express.

2.3.1 Why VerilogA?

VerilogA was chosen because it allows photonic components to be simulated in combination with the standard CMOS cell models provided by most foundries. Since it is fully compatible with the Cadence ®simulation environment, photonic components can be included easily and simulations can take advantage of the many simulation engines and options that are available already in the Cadence framework (e.g., Spectre). Indeed, once these models are dropped into the schematic, DC, ac, transient, and noise simulations can be run identically to how they would be for pure electronic designs. Noise can even be added, either using standard VerilogA noise blocks, or directly into the component model. There is no need to modify the underlying simulation program - all the standard simulations and parameters sweeps available for CMOS VLSI designs can be run on these joint systems. There are of course drawbacks to using VerilogA and Cadence, over, say MATLAB. Most notably, VerilogA is a time domain code, and simulating at the optical frequency is too time consuming. Additionally, Verilog does not support complex numbers, a functionality we need to capture the phase effects of the signal. Resolutions to these two issues will be discussed below.

2.3.2 Simulation Architecture and Philosophy

There are a number of traits that are important for a useful, flexible simulation platform. The first and most basic, is that we want to make a full simulation toolkit, not just create a few device models compatible only with a few specific or specific types of systems. We therefore need to think both about how best to write *specific models* and about how best to construct overall system *architecture*. On an architectural level we would like photonic models to (1) be easily combinable with electronic standard cells, (2) be usable in a variety of simulation types (DC, ac, transient, etc), (3) be fast, (4) capture phase effects, (5) support forward and backward propagation, (6) handle multiple frequencies and WDM channels, (7) capture appropriate electro-optic interactions, (8) be modular and flexible, (9) organically generate behavior

of composite devices, that is to not require a new block for every new component combination, and (10) require a limited amount of hacking so new types of systems can be easily modeled without changes to the underlying toolkit. This last one is basically a preference for elegance in solution.

The first two goals on this list are automatically satisfied by the choice of VerilogA as the modeling language and the creation of the models in the Cadence framework. However, this choice adds a constraint of its own: (11) component models must exist only in the time domain, as VerilogA is a time domain modeling language. We next tackle desired trait number three: speed. Modeling at the optical frequency timescale, 200THz or 5fs and below, is prohibitively time consuming. To get around this, the optical responses are shifted down and calculated relative to a reference frequency. This is made possible because the optical frequency band of interest is generally much smaller (1THz to <100GHz, timescales on the order of picoseconds). Capturing amplitude and phase effects of all the components in this smaller band is generally sufficient. To accomplish this shift in a simple manner, we choose to work on the analytic representation of the signal (that is we retain only the positive frequency components of the real signal), as it allows such a shift to be made solely by multiplying the signal by $e^{j\omega_{ref}t}$ where ω_{ref} is the arbitrarily chosen reference frequency. The justification for choosing and meaning of the analytic signal will be explained in more detail in the section below. It is important to note that the chosen reference frequency, which will correspond to a frequency of zero in the shifted simulation, is arbitrary. This frequency does not have to have any physical meaning to the system; it does not need to be the laser frequency, a ring reference frequency, the center of a channel, or anything else. However, because it defines zero frequency, placing it near the center of the desired simulation bandwidth allows for the largest time steps and yields the largest improvement in simulation speed.

Working with the analytic signal also has the added benefit of making it easy to model phase effects, and satisfy desired trait number (4) - a phase or frequency shift is just a multiplication by an exponential. However, it also means that the signal is no longer real. Because it usually works on the full real signal, and models

phase changes through time delays, VerilogA does not inherently support complex numbers. We don't want to go back to working on the full real signal, thus we need to add complex algebra into VerilogA. To do so we write a set of complex algebra modules that can be called from inside our photonic component model blocks. We also add optical and thermal disiples into the disciplines file ("disciplines.vams"), so that we can pass around optical signals as a distinct entities from electrical signals. This avoids the cludge of passing optical amplitudes around as "voltages," allows them to be their own thing, and prevents accidental cross wiring with electrical ports. Because a complex signal has two independent degrees of freedom, we represent it as a two signal bus, with one signal representing the real and one signal representing the imaginary part. The real and imaginary parts are chosen instead of the amplitude and phase because working with these avoids the numerical convergence issues present in the amplitude and phase representations due to discontinuities in the phase signal when it wraps around 2π and the lack of a clearly defined phase for a zero amplitude signal. Expanding the representation into a four signal bus, two signals for the real and imaginary parts of the forward propagating signal and two for the backward propagating signal, allows us to satisfy requirement (5) and to capture behavior and errors caused by reflections as well as the effects of having forward and backward propagation on the same waveguide (should it ever be relevant).

We further note that representing the signal(s) in the manner described above, as well as carefully representing the components as described below, will allow for multiple optical frequencies or spreads of frequencies to be modeled *without* any need to make explicit reference to these frequencies (desired trait 6). Thus the optical signal can propagate through the system, including through frequency selective components such as ring filters or frequency changing elements such as single sideband modulators, without any explicit reference made to the frequency of the light or the frequency transfer function of these elements. All the appropriate interactions will be handled automatically. This avoids a common issue in previous work, where the center frequency of the channel must be passed as a variable through the system and is used to look up what behavior is appropriate – something that can severely

limit accuracy in the case of wide-band signals (for example, if a modulated signal is put through a narrow ring filter) [53]. Avoiding this means that we have a true time domain code that is compatible with VerilogA, as opposed to some hybrid time-frequency domain code. Finally, we note that since VerilogA was developed for use in the electronics industry, adding electrical structures to optical components to model electro-optic effects (and satisfy desired trait 7) is straightforward. The details of these implementations are described below.

The next thing on our list of desired traits is for the code to be modular and flexible, that is for it to be able to accurately model a large variety of system types for a large variety of applications with a minimum need for changes or special set ups. While the architecture above helps with this, it is also critical to choose and design optical component models for maximum flexibility, applicability, and adaptability. To this end optical blocks are defined at the lowest possible level; they consist of primitive building blocks that cannot be broken down any further: waveguides, couplers, phase-shifters, photodetectors, heaters, light sources, etc. Physical equations are used within these blocks to capture optical behavior and to allow the insides to be easily modified to add complexity and higher order effects if desired. Where helpful an ideal and a non-ideal version with these effects included are created separately. The designer can then choose the one most applicable to their simulation. The blocks include both the optical response and, if applicable, the electrical response of the component.

More complex structures, such as modulators and filters, are then built out of these components. The transfer functions of these composite devices are created automatically during the simulation; they do not need to be recalculated by hand. This differs from approaches in the literature where transfer-functions for these more complicated devices, especially modulators, are calculated explicitly and programmed into VerilogA as the most basic building blocks (e.g [56, 57]). The phase information in these blocks is also often dropped. Under these approaches, small changes to devices requires the recalculation and separate modification of the inside of each composite block, making it hard to maintain consistency. It also means that the transfer function of each new configuration needs to be calculated and programmed by hand. Under

our approach where these composite blocks are built out of primitive blocks, changes to the bottom blocks are automatically propagated throughout the system during simulations. New configurations can similarly be built out of existing components and their behavior modeled without the need to create any new code. This will be demonstrated in more detail below.

Some complex devices, such as modulators and ring resonators, may be used repeatedly within the same system. It is desirable not to have to recreate these small systems from scratch every time they appear. For combinations of components that are commonly used, schematic symbols can be created. Inside these symbols the component is still built out of our primitive blocks - thus adaptability and flexibility is maintained. The schematic simply creates a block so the device or sub-system can be dropped into larger simulation layouts without needing to be redrawn each time. Parameters can be passed into the building blocks as desired. Such symbols can also be useful for holding more complicated components, even ones that might not be repeated often, so that the top layout is cleaner and easier to understand. Figure 2-5 illustrates the creation and inside of such a schematic symbol for a ring resonator

In summary, we have identified a set of fundamental photonic components from which larger systems are built. We have described the optical and electrical (and in some cases thermal) behavior of each at the physical level, such that the characteristics of a composite device will be created organically. We have created a guiding philosophy and architecture in which these components exist that allows for joint photonics + CMOS circuitry simulations that satisfy our list of desired traits, and have written the necessary complex algebra sub-functions and added the optical and thermal disciplines needed to support this architecture. We next seek to explain the inner workings of these component blocks.

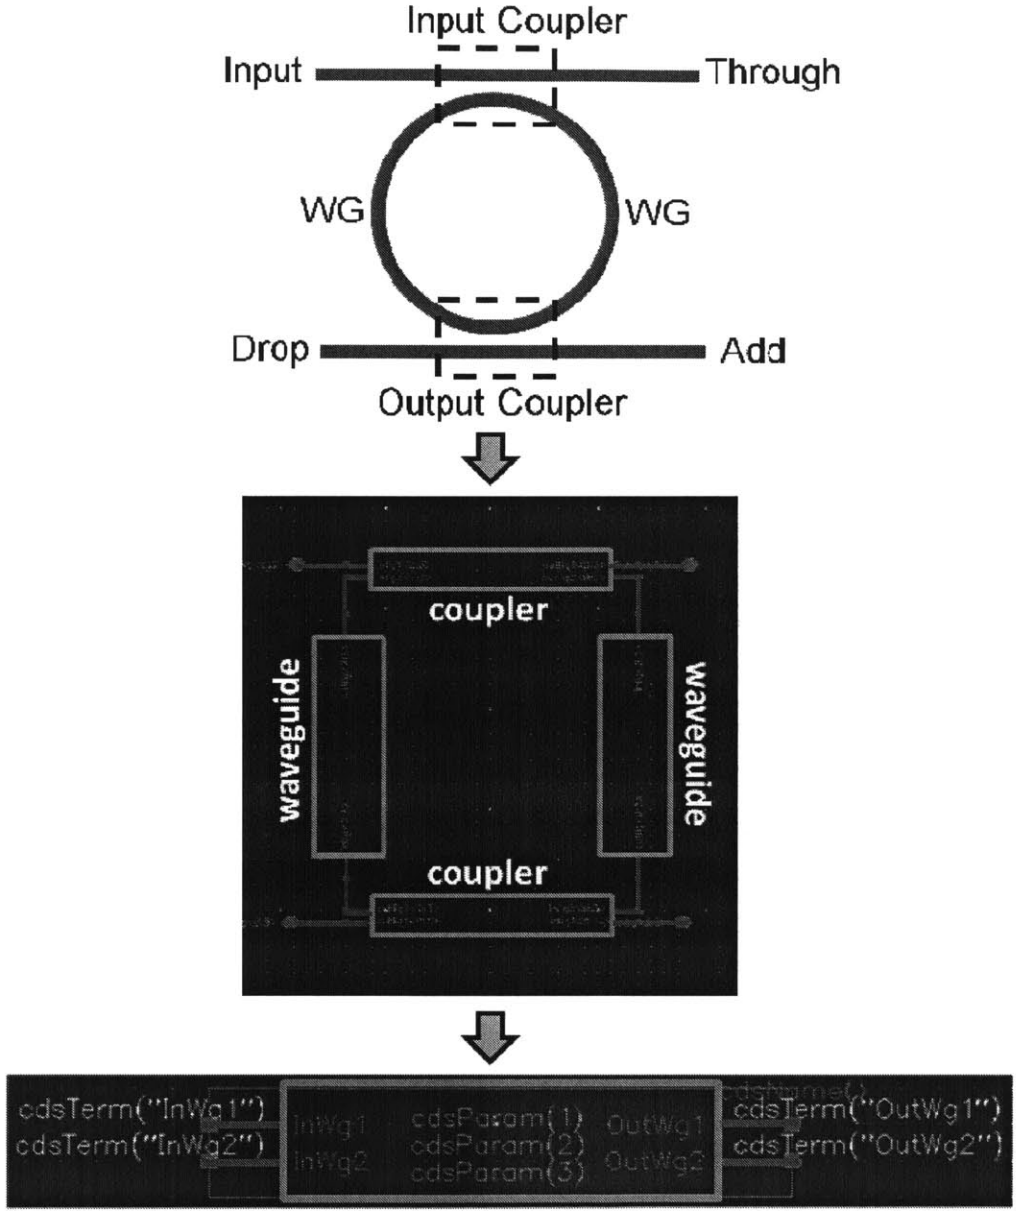

Figure 2-5: A ring resonator (diagrammed at top) is broken into components. The VerilogA modules for these components are placed and connected in a schematic in Cadence (middle). A schematic symbol is created to represent this schematic block (bottom). This symbol can be inserted wherever a ring is needed as opposed to having to redraw the middle block. Relevant variables are passed through this symbol to the underlying schematic.

2.3.3 Mathematical Representation of Primitive Component Blocks

Before we go on to show the results and versatility of the simulation toolkit, we will describe the inside of the optical blocks and the mathematical underpinning of the simulations in more detail. The most fundamental building block of optical systems is the optical waveguide. To represent this component in VerilogA, it is necessary to come up with a time domain representation of the waveguide transfer function. To do so, we start by considering the electric field. The physical electric field is real, but many operations are difficult to do on the full real signal. To get around this, we can take advantage of the fact that real signals are symmetric around zero in the frequency domain. We can therefore ignore all the negative frequencies without loss of information. An inverse Fourier transform can then be done to get back to the time domain, where the result will be a complex signal. In the signal processing literature, this is called the analytic representation of the signal or analytic signal (see [58]), though it is often assumed implicitly in the electromagnetic literature (see, for example, [59, 60]). Many operations are simplified by working with this representation; for example, a phase shift corresponds to multiplying by an exponential. At the end, to recover the real, observable signal we can take the real part. Because our component code is very modular, we don't know where the end is until we hit some sort of detector. We therefore pass the complex, analytic signal though all the components. We note that, because our light source models output the time-domain analytic signal representation of the laser output to start with, there are no Fourier transforms in the code, nor is there any explicit tracking or specification of light frequencies, all the components are operated in the time domain as is required by the VerilogA simulator. Below we will go through the component math explicitly.

We begin with an optical waveguide. The electric field at the start of the wavguide ($z = 0$) is $\tilde{E}(0, t)$. The wave equation, derived from Maxwell's equations, states that

at a location, z, in the waveguide the field is equal to:

$$\tilde{E}(z,t) = \int_{-\infty}^{\infty} \tilde{E}(0,\omega)e^{(-j\beta z + j\omega t)}d\omega = \int_{0}^{\infty} \tilde{E}(0,\omega)e^{(-j\beta z + j\omega t)}d\omega \tag{2.2}$$

Where $\beta = \beta(\omega)$ is the wavevector in the waveguide and the second equality follows because there are no negative frequency components. As explained above, in order to avoid simulation at the optical time step we choose some reference frequency, ω_R, such that $\omega = \omega_R + \Delta\omega$ and shift our signal down by this frequency. To do this, we first Taylor expand $\beta(\omega)$ around this frequency yielding:

$$\beta(\omega) = \beta(\omega_R) + \Delta\omega \frac{\partial\beta}{\partial\omega}\Big|_{\omega_R} \tag{2.3}$$

We note that $\frac{\partial\beta}{\partial\omega}\big|_{\omega_R} = \frac{1}{v_g}$, where v_g is the group velocity of light in the waveguide at the reference frequency and is equal to c/n_g, where n_g is the group index. This brings us to:

$$\tilde{E}(z,t) = \int_{0}^{\infty} \tilde{E}(0,\omega_R + \Delta\omega)e^{(-j\beta(\omega_R)z - j\frac{\Delta\omega}{v_g}z + j(\omega_R+\Delta\omega)t)}d\omega \tag{2.4}$$

We note that we can define a function $\tilde{E}_{shift}(z,\Delta\omega) = \tilde{E}(z,\omega_R + \Delta\omega)$ which re-labels the frequency to shift the amplitude coefficients by the reference frequency. We note that the reference frequency can be thought of as a carrier frequency and $\tilde{E}_{shift}(z,\Delta\omega)$ as an envelope function, but want to emphasis that the reference frequency is not the center (or necessarily any) frequency of the propagating light and that $\tilde{E}_{shift}(z,\Delta\omega)$ is complex (it is not amplitude modulation) and holds for an arbitrarily (or infinitely) large frequency range. It is the accuracy of the Taylor Expansion of β and simulation time constraints that will limit the workable frequency range. Using this function, we can change the variable of integration to $\Delta\omega$ and pull the constants out of the

integral function:

$$\begin{aligned}\tilde{E}(z,t) &= e^{-j\beta(\omega_R)z+j\omega_R t}\int_{-\omega_R}^{\infty}\tilde{E}_{shift}(0,\Delta\omega)e^{-j\frac{\Delta\omega}{v_g}z+j\Delta\omega t}d\Delta\omega \\ &= e^{-j\beta(\omega_R)z+j\omega_R t}\int_{-\infty}^{\infty}\tilde{E}_{shift}(0,\Delta\omega)e^{-j\frac{\Delta\omega}{v_g}z+j\Delta\omega t}d\Delta\omega \end{aligned} \tag{2.5}$$

Where, again, the second equality follows from the fact that there are no frequency components between $-\omega_R$ and $-\infty$. We note that the integral is just the Fourier transform of the time shifted $\tilde{E}_{shift}(z,t)$. We, thus have:

$$\tilde{E}(z,t) = e^{-j\beta(\omega_R)z+j\omega_R t}\tilde{E}_{shift}(0,t-\frac{z}{v_g}) \tag{2.6}$$

A similar line of reasoning yields: $\tilde{E}(z,t) = e^{j\omega_R t}\tilde{E}_{shift}(z,t)$ Combining these two equations and recalling that $v_g = c/n_g$ and that $z = L$ at the end of the waveguide finally yields:

$$\tilde{E}_{shift}(L,t) = e^{-j\beta(\omega_R)L}\tilde{E}_{shift}(0,t-\frac{Ln_g}{c}) \tag{2.7}$$

We note that $\tilde{E}_{shift}$ is the signal we pass around in our code, as it has had the reference frequency removed. We further note that the measureable E-field would be the real part of $\tilde{E}(z,t)$. To implement this waveguide equation in continuous time in our VerilogA code, we output the phase shifted $\tilde{E}_{shift}$ after a delay of $n_g L/c$. We also add a propagation loss term, α_{field}. The VerilogA component block for the waveguide therefore looks like:

```
`include "../../../../cad/constants.vams"
`include "../../../../cad/disciplines.vams"

module Optical_Waveguide(inlig, outlig);
        // DEFINE PORTS
        inout   [0:3]           inlig;  //[0:1]left input [2:3]left output
        inout   [0:3]           outlig; //[0:1]right output [2:3]right input
        optical [0:3]           outlig, inlig;
        // OPTICAL PATAMETERS
        parameter real L = 0.0005;
        parameter real ng  = 3.97;
        parameter real np = 2.1;
        parameter real alphaA = 287.6;
        parameter real G_freq = 1.93e14;
        // INTERNAL VARIABLES
        optical [0:1] outNodly; // cart
        optical [0:1] outNodlyB;
        optical [0:1] ringres;
```

```
        optical [0:1] ringConv;
        // INTERNAL POLAR/CARTESIAN CONVERSION
        pol2cart convs1(ringres, ringConv);
        cartmul mulout1(ringConv, inlig[0:1], outNodly);
        cartmul mulout1B(ringConv, outlig[2:3], outNodlyB);

        analog begin
                //calculate phase shift and attenuation
                E(ringres[0]) <+ (-L*np*2*`M_PI*G_freq/`P_C)%(2*`M_PI);
                E(ringres[1]) <+ exp(-alphaA*L);
                //forward propagating wave:
                E(outlig[0]) <+ absdelay(E(outNodly[0]), L*ng/`P_C);
                E(outlig[1]) <+ absdelay(E(outNodly[1]), L*ng/`P_C);
                //backward propagating wave:
                E(inlig[2]) <+ absdelay(E(outNodlyB[0]), L*ng/`P_C);
                E(inlig[3]) <+ absdelay(E(outNodlyB[1]), L*ng/`P_C);
        end
endmodule
```

The equations governing other components can be similarly derived. Phase-shifters can be modeled as waveguides where the effective group and refractive index, as well as the loss, vary, possibly nonlinearly, with applied voltage. In a realistic phase shifter, the dependence on voltage can be nonlinear and a voltage-dependent loss can also be included. In practice, reverse bias silicon diode phase shifters that rely on the plasma dispersion effect to create a phase shift are commonly used. These phase shifters can be modeled like the non-ideal phase shifter with the addition of the underlying electrical diode structure. This allows for the capture of electrical effects, such as speed limitations due to diode resistance and capacitance, as well as optical effects. We have implemented all three types of phase shifters in our toolbox. Similarly a simple low-pass filter can be used to limit the rate of change in a thermal, heater based phase-shifter.

Waveguide couplers are modeled in our code as point couplers where part of the light is transferred across, with appropriate phase shift, and the rest remains. Both adiabatic 3dB couplers and directional or ring couplers, where the field coupling fraction is specified, have been implemented. Similar principles to those above are used to write models for photodetectors. Photodiodes form a common class of integrated photodetectors. In the ideal photodetector, the instantaneous optical intensity is simply measured and transferred to the electrical domain. Such an ideal diode can be modeled by specifying a responsivity. In the non-ideal case, the underlying diode

structure is included to allow for electrical responses such as frequency roll-off and dark current to be captured. We note that due to the short length of the device, which can be as little as 4μm [5], it does not make sense to include the optical delays in the device model. Other useful components, such as continuous-wave (CW) lasers, as well as some debugging components, such as phase detectors, are also included in our toolkit. We note that this frameworks makes it easy to both add new components and modify old ones, for example, to add higher order effects or noise. Once this simulation framework has been established and these simple component models have been implemented, complex optical systems can be built and combined with CMOS devices to model complicated joint systems. As the next sections will demonstrate, using these basic building blocks and this framework, effects like optical resonances and phase-effects are automatically captured, thus allowing for convenience and ease of use, especially for electronic designers less familiar with optical simulation techniques. Most importantly, larger scale electronic-photonic system simulations can be carried-out with high-fidelity.

2.3.4 Simple Passive and Active Composite Devices

Because the optical components are implemented as basic physical blocks, they must be combined together to create full devices. Our simulation technique can be used to model very complicated devices and systems that are beyond the capabilities of more traditional optical simulation techniques. To validate our technique, however, we also modeled several devices simple enough for existing techniques to handle and found that our results matched theirs. Since the traditional techniques have been validated against experiments, this is equivalent to experimentally validating our technique.

The first of these devices is a ring resonator, which is created out of a waveguide connected in feedback between two ports of a coupler. A diagram showing this system is given in Fig. 2-6.a. The frequency response of the ring resonator is shown in Fig. 2-6.c. Because this is a time domain code, to get the frequency response it was necessary to sweep the input cw laser frequency. This is done by having a voltage controlled phase ramp inside the laser (as a phase ramp in the time domain corresponds to a

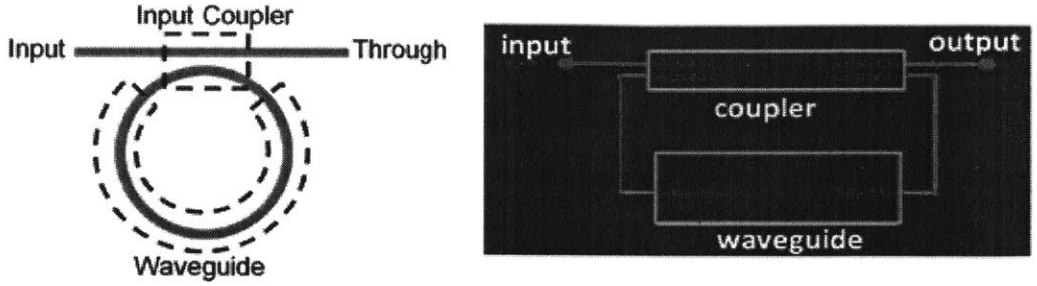

(a)

(b)

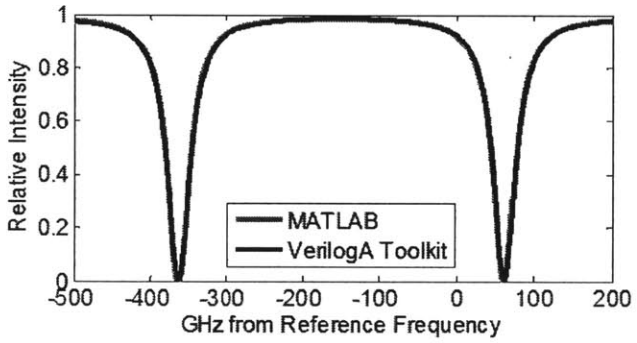

(c)

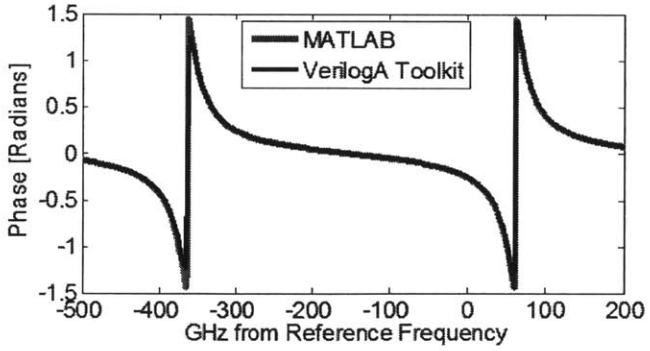

(d)

Figure 2-6: (a) Diagram showing the components of a ring resonant filter and accompanying layout using our toolkit in Cadence Virtuoso. (b). A full link schematic including laser and photodetector is created to simulate the performance of the device. (c). Amplitude and phase response of the ring showing good agreement between VerilogA simulations and MATLAB transfer function code.

frequency shift in the frequency domain) and running multiple DC simulations with different voltage inputs to this phase ramp. Similar results can also be obtained by linearly sweeping the input voltage to the phase ramp, though in this case one needs to be sure to sweep at a slow enough rate to avoid the resonance pulling effects that have been observed in experiment and in our simulations. A diagram of the full simulated system is shown in Fig. 2-6.b. The obtained spectral response of the ring is compared and shown to match well to the results from an existing transfer matrix code implemented in MATLAB.

We can further compare the results to theoretical expectations. The free spectral range (FSR) of a ring resonator can be analytically found to be $FSR = c/n_g L$, where c is the speed of light in vacuum, n_g is the group index, and L is the effective waveguide path length. The results in Fig 2-6.c. are for a SiN waveguide 200nm tall and 1μm wide that has a group index of 1.72 and is formed into a ring with a 66μm radius. The expected FSR is therefore 423GHz, as it is in the simulation. The effective refractive index of this waveguide is 1.516. We therefore expect the closest reference frequency to 1550nm (=193THz), the reference frequency used in this simulation, to be at $\lambda = n_{eff}L/m = 1549.5nm$ or 60.9GHz above the reference frequency. This value is also the same as that seen in Fig. 2-6.c. Finally, the half width at half maximum (HWHM) of the resonance was found to be 17.4GHz, also exactly what is expected given the input waveguide field loss, α, or $312.7m^{-1}$ and that $HWHM = c\alpha/n_g\pi$. Moreover, not only does the amplitude response of the ring match that expected, but the phase response (measured using a debug block that outputs the instantaneous phase), shown in Fig. 2-6.d. also matches, going to $\pm\pi/2$ as the laser frequency approaches the resonant frequency. Thus, it can be seen that our VerilogA toolkit also captures phase effects correctly.

To further expand on this example we can consider the case where the ring has a drop port and the case of a two ring filter. Both modifications are very simple to make using our toolkit, making it easy for designers to change systems on the fly. We compare both cases below to the results of transfer matrix codes implemented in MATLAB. No new calculations or equations were required when using our VerilogA

toolkit; however, the new transfer matrices needed to be calculated and input into our MATLAB codes.

To add a drop port, we simply add an output coupler as shown in Fig.2-7.a. One can then use this model in a system or to explore how to best optimize the component for a desired application. For example, in a single ring resonator with drop port, it has been shown [61] that, given a target HWHM, the drop port output intensity is maximized when the input and output coupler coefficients are equal. On the other hand, the through port intensity is minimized at critical coupling where the input coupler coefficient equals the sum of the waveguide and output coupler loss. VerilogA simulations of both cases are shown in Fig.2-7.b. and compared to transfer matrix codes, where we observe very good agreement. Note that the waveguide loss was exaggerated in these cases to highlight the difference.

The two-ring series filter case is shown in Fig. 2-8.a. As can be seen, creating this case with our toolkit is similarly straight forward. One simply adds a second ring to the bottom of the first. Further, basic schematic symbol of this and other ring filters exist in our tool kit making it even easier for a designer to switch filters to tune the roll-off, ripple, passband width, or other figures of merit, and re-simulate. Two-ring filters are often helpful in systems such as switching matrices where having a wider and flatter passband allows for a wider bandwidth signal to propagate through. However, because two resonances are combined to make the larger resonance, ripples are often introduced into the passband. All these effects are captured in the simulation output which is shown in Fig 2-8.b and also matches exactly the expected output calculated via transfer matrices in MATLAB.

The above demonstrations all involve passive devices, but simulating active devices in our toolkit is equally straight forward. They can be built in our tool kit in a way exactly analogous to the passive devices described above. In this section we will demonstrate the design and simulation of a ring modulator of the type often proposed for digital communication links. We will also demonstrate the simulation of a single sideband modulator. Single sideband modulators are often used to create frequency shifts, and are important for both analog applications, such as arbitrary waveform

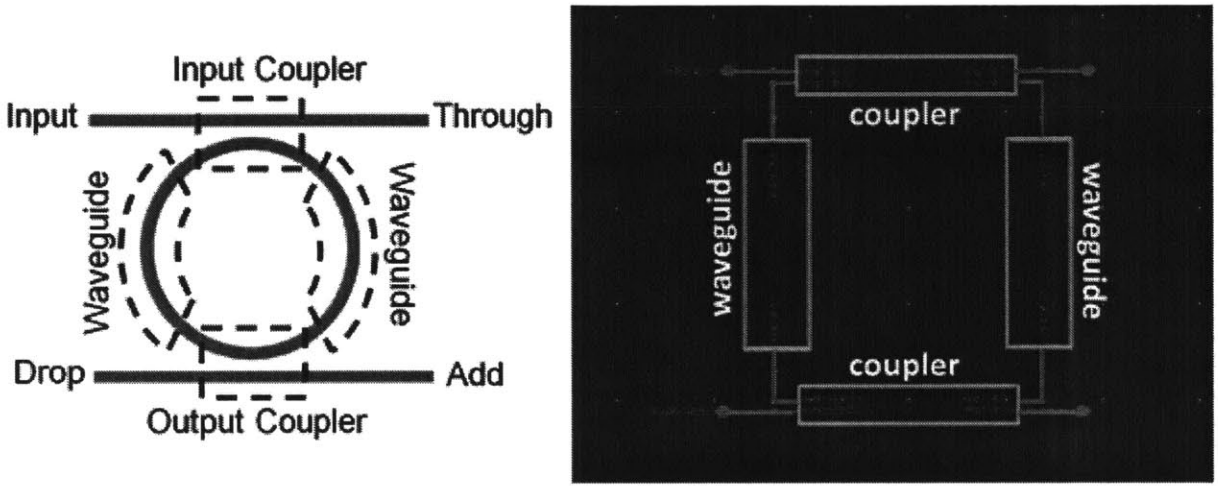

(a)

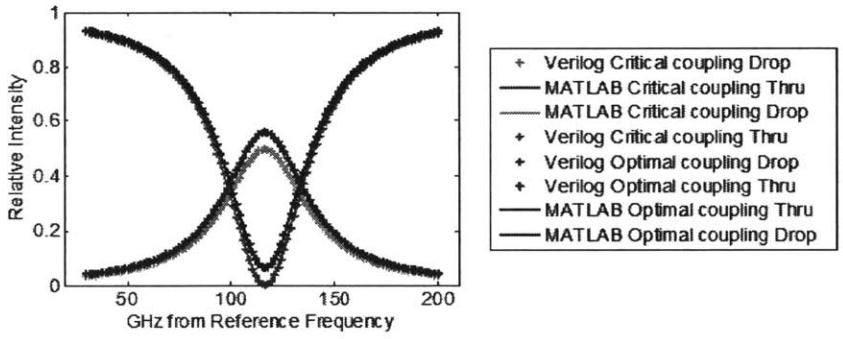

(b)

Figure 2-7: (a). Diagram showing the components of a ring resonant filter with drop port and accompanying layout using our toolkit in Cadence Virtuoso. (b). Amplitude response of the ring showing output intensity in the through and drop port. One can see that the critical coupling has better extinction in the through port but less power in the drop port output, while matched coupling has worse extinction and higher drop port output power, as expected. Good agreement is seen between VerilogA simulations and MATLAB transfer function code

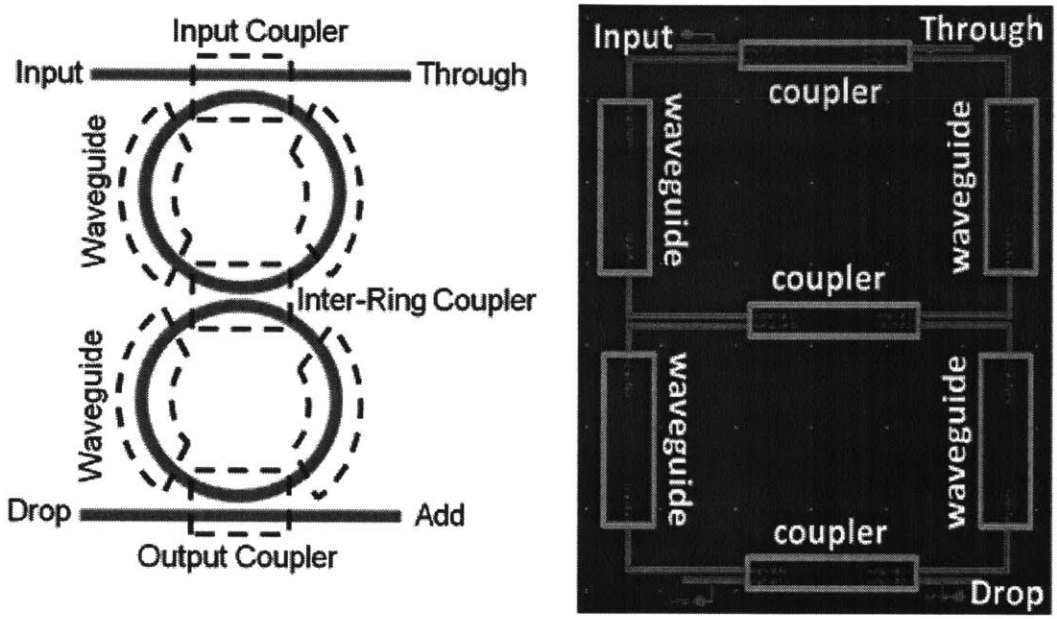

(a)

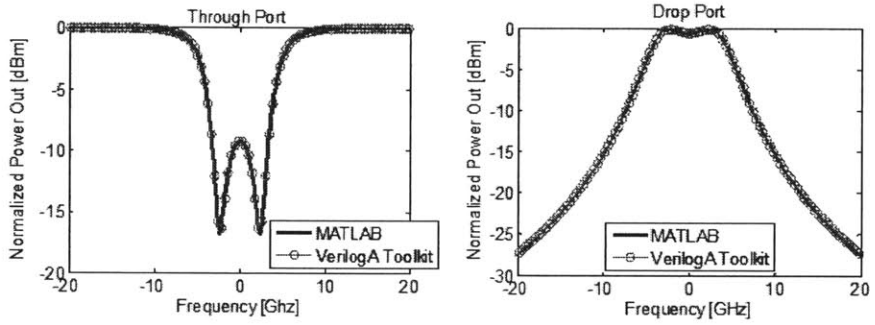

(b)

Figure 2-8: (a). Diagram showing the components of a two ring series filter with drop port and accompanying layout using our toolkit in Cadence Virtuoso. (b). Amplitude response of the ring showing output intensity in the through and drop port showing good agreement between VerilogA simulations and MATLAB transfer function code.

generators, and digital applications, like dense wavelength division multiplexing systems [62].

Ring modulators are important for low power optical links and wavelength division multiplexed systems [63]. A ring modulator can be made in our toolkit by replacing the waveguide component with a phase shifter in the resonant ring (see Fig.2-9.a). As described above, the phase shifter allows an input voltage to change the refractive index, group index, and loss of the waveguide, and thus the location of the resonance frequencies in a ring modulator. The shift in frequency of the resonance location of the modulator with voltage is shown in Fig.2-9.b. The simulation is done for a five micron radius disk modulator made in a 220nm tall silicon layer. The output is compared to calculations done in MATLAB. Again, it matches the expected output.

However, the transient output, and not the DC output, is often what is of real interest in modulators. It is easy to simulate this using our toolkit: the input laser is simply set at one wavelength and the input voltage to the phase shifter changed from a DC to a pseudo-random-bit-sequence (PRBS) source (using a standard electrical PRBS block from the "ahdlLib" library). The layout of this link is shown in Fig.2-9.c, and the output in Fig.2-10 As can be seen, when the input voltage is low, the optical amplitude inside the ring modulator is high and the through port output is low. Conversely, when the input voltage is high, and the resonance is shifted away from the laser frequency, the amplitude inside the ring is low, and the through port output is high, as expected. Further, we can clearly see the usefulness of a transient simulation by noting the roll-off in the output through the modulator. In this simulation it is the optical time constant (due to high Q) that ultimately limits the frequencies at which the device works. If the Q were lower or the electrical circuit in the phase-shifter or photodetector slower, the frequency roll-off would be caused by that instead. It should be noted that because the code models the optical interference inside the ring with each time step, these phenomenon are neatly captured without any explicit modeling. No changes to the inside of the blocks or the overall system code are needed; we just have to draw the schematic for the desired systems and run the preferred simulation.

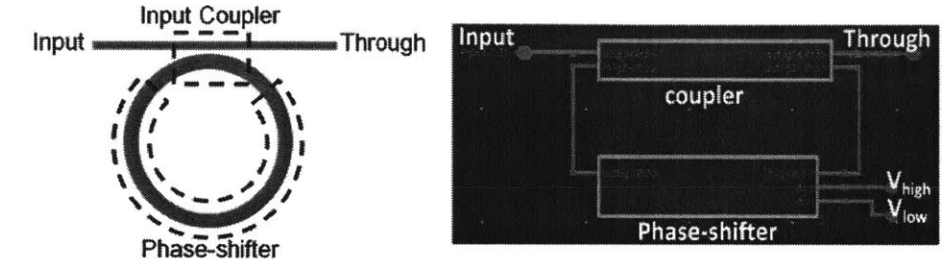

(a)

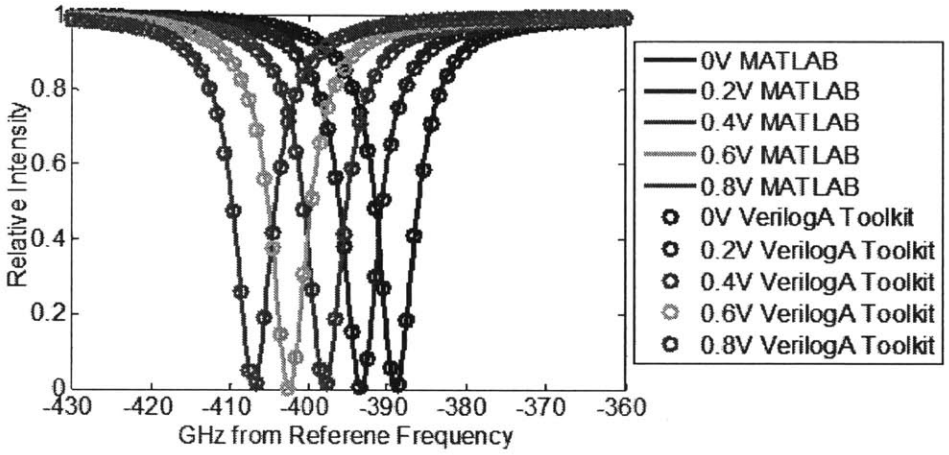

(b)

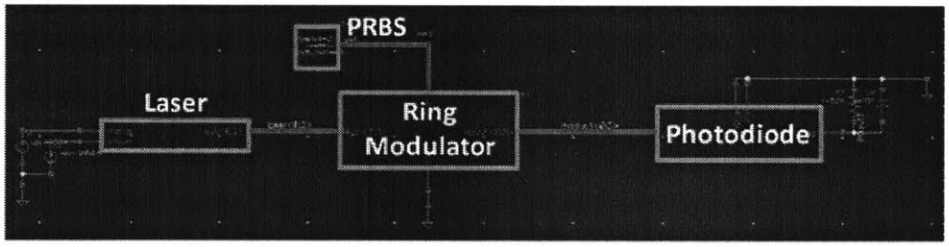

(c)

Figure 2-9: (a) . Diagram showing the components of a single ring modulator and accompanying layout using our toolkit in Cadence Virtuoso. (b). Modulator amplitude transfer function as a function of applied phase-shifter bias, showing good agreement between VerilogA simulations and MATLAB transfer function code. (c). Schematic of full like. Output from a CW laser is feed into a ring modulator driven by a PRBS source. The output of the modulator is sent to a photodiode.

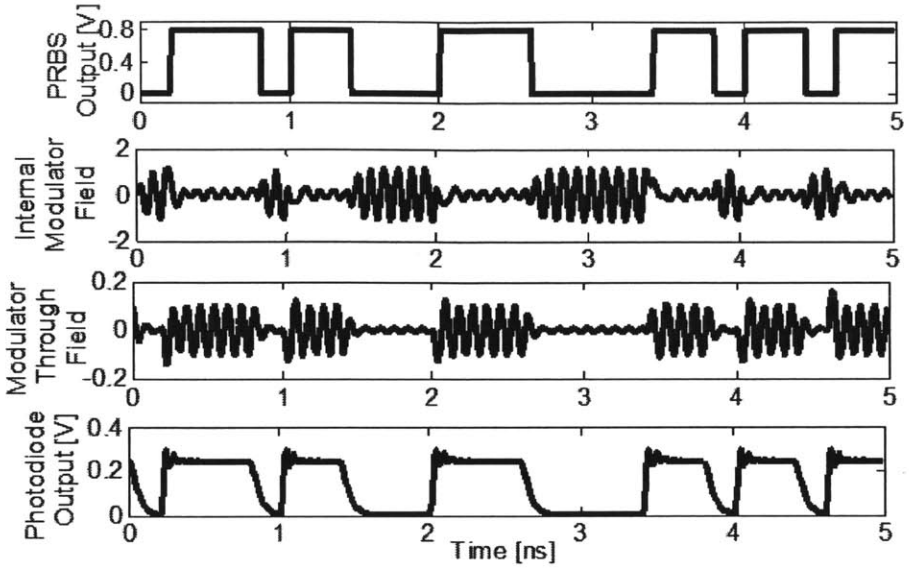

Figure 2-10: Simulation output, showing electrical driving voltage, field amplitude inside the ring, through port field amplitude, and output of photodiode.

In addition to the signal amplitude, for some applications a designer may be interested in the phase or frequency effects of a device or system. Fortunately, our toolkit is also well equipped to handle and accurately model such effects. An example of such a device where phase and frequency effects are of overriding importance is a single sideband modulator (SSB). As mentioned, SSBs are important for many analog applications as well as for high density WDM communication systems as they allow for closer channel spacings and lower powers [64]. Additionally, the same topology can be used to transmit data using more complicated modulation formats (BPSK, QAM, etc). Correctly modeling the phase effects and delays of components is vital to simulating single sideband operation, and as we will demonstrate below, our VerilogA Toolkit is capable of doing this. We note again that the transfer function of the SSB is never calculated explicitly, instead the correct output is generated automatically from the combination of basic optical blocks.

The SSB topology is shown in Fig.2-11.a. along with its schematic layout using our VerilogA toolkit in Fig 2-11.b. The SSB is made from two nested Mach-Zehnder modulators (MZMs) and generates a single sideband carrier suppressed signal. To do

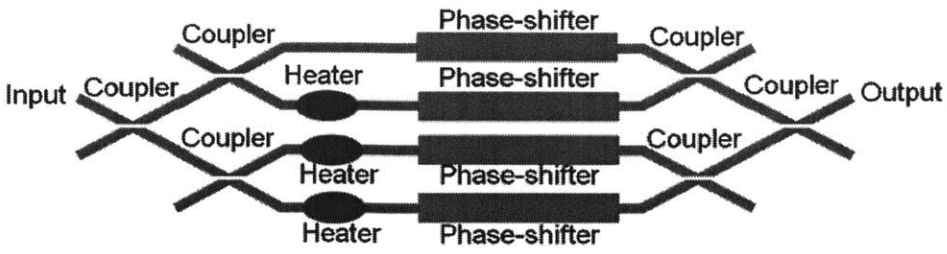

(a)

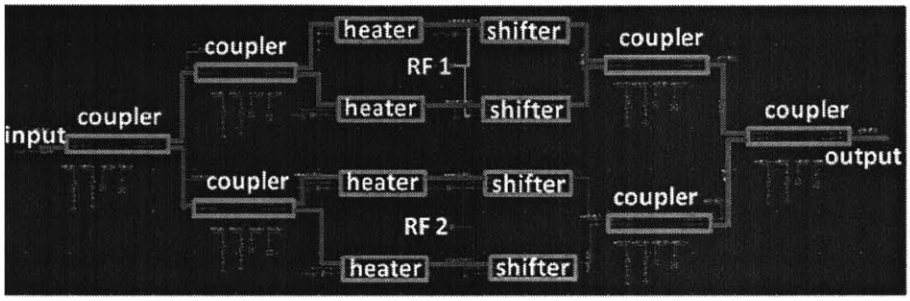

(b)

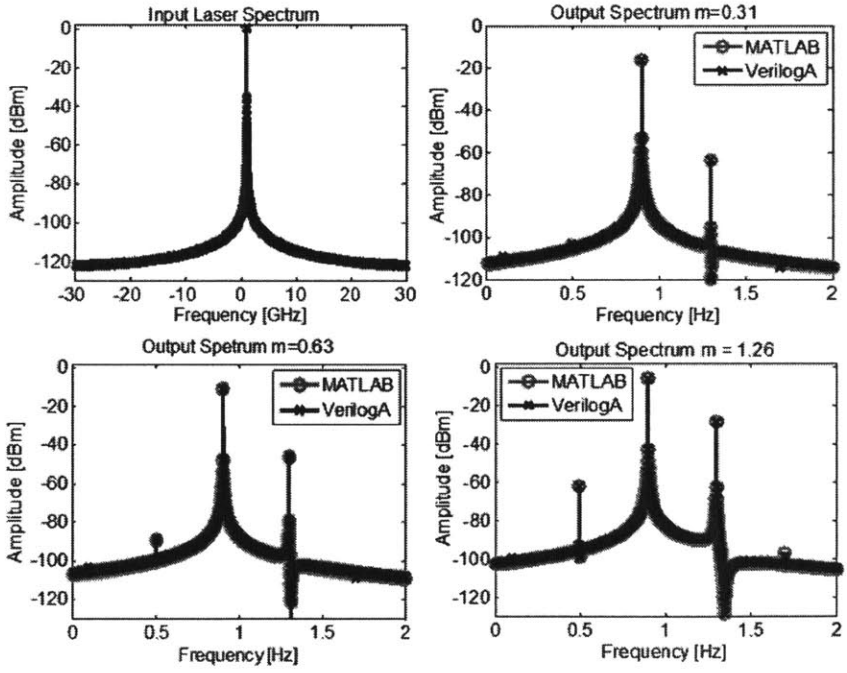

(c)

Figure 2-11: (a.) Diagram of single side band modulator component system (from reference [129]) and accompanying layout using our toolkit in Cadence Virtuoso using basic components (couplers, heaters, phase-shifters). (b) Fouier transform of SSB output as a function of modulation amplitude. Input CW signal is shown in black in upper left. Following graphs show steadily increasing modulation depth. Increase in power transferred to desired sideband and accompanying increase in power in undesired harmonics is seem. VerilogA and MATLAB code show good agreement.

so each individual MZM is biased using a heater at the zero output point (thereby suppressing the carrier), and driven in push-pull. The electrical RF input signals to each modulator are out of phase by $\pi/2$, such that one arm receives a cosine modulation and the other a sine modulation. An extra $\pi/2$ phase added by heaters to one of the MZM signals therefore allows one sideband to cancel and the other to add. Which sideband is kept depends on which arm receives the extra $\pi/2$ phase relative to which receives the sine (vs. the cosine) RF input signal. More detail about SSB operation, along with measurements of an integrated version of the device will be presented in Chapter 3 below. This surviving sideband is the output of the device, and is given by $f_{out} = f_{opt} f_{rf}$ or $f_{out} = f_{opt} + f_{rf}$, where f_{opt} is the original optical frequency and f_{rf} is the modulation frequency. We choose to keep the lower sideband.

The output of the simulation using our VerilogA toolkit is shown in Fig.2-11.c. Here a continuous wave (CW) laser block inputs a CW signal to the SSB. Since our code operates entirely in the time domain, the Fourier transform was computed in a postprocessing step. In a real system model, the time domain output samples would simply be passed on to the next block with no transform necessary. To verify that the VerilogA model is in fact working correctly, we compare it to a MATLAB system code for the case of a 0dBm laser output that is 1GHz from the reference frequency and is modulated at 100MHz to yield a sideband frequency of 900MHz. The presence of a single, frequency shifted tone at the output can clearly be seen — the carrier and unwanted sideband are well suppressed in all cases. As the modulation depth increases the amount of power transferred to the desired sideband increases as does the relative height of the higher harmonic distortions. These last are caused by the increasing interaction with the nonlinear portion of the MZM transfer function and are given by the Bessel function expansion for the sine within the sine [64, 65]. All these expected effects can be clearly seen in the output spectra shown in Fig.2-11.c. Further, there is very good agreement between the MATLAB and the VerilogA models in all cases, indicating that our VerilogA framework is working correctly. It is worth noting that writing and running the MATLAB code to simulate the output of the SSB

modulator took more time and was more complicated than setting up and running the SSB modulator schematic in VerilogA, thus demonstrating the convenience and one of the benefits of having a single, flexible simulation platform.

2.3.5 Full System Simulations

In the preceding section, we demonstrated the ability of our system to correctly simulate simple, composite components, and basic devices. But the ultimate goal of our work is to be able to simulate complex systems, particularly ones that include the electrical driving circuitry and noise models. Now that we have shown that we can build accurate, flexible device systems out of our optical toolkit blocks and demonstrated their use in simulations, in this section we will take those blocks and use them to build full heterogeneous electrical and optical systems. Note how similar the process is to building circuits out of transistors, by simply instantiating components, setting parameters, and connecting the waveguides and electrical wires, further allowing for easy combination with existing electrical simulation infrastructure and ease of use by designers familiar with such. While the toolkit can of course be used to build a multitude of different systems and simulate them at numerous operating points under a myriad of conditions, we will focus on two examples here: a wavelength division multiplexed (WDM) link and a Pound-Drever-Hall laser stabilization loop.

We begin with the WDM link. One of the main driving forces for integrated silicon photonics is to use waveguide division multiplexing on a single waveguides to replace copper wires for data transmission on chip and, thus, solve the bandwidth density and energy consumption problems in current electrical interconnects [66]. Currently more than half the power on an electrical chip goes to these communication needs [67]. In a WDM link a laser source provides a comb of laser frequencies each of which is modulated with a differnt data stream and transmitted through the same waveguide. On the reciever side filters pick off each individual line and drop it into a photodiode where it is converted back into the electrical domain. As there are many optical and electrical components involved, there are many tradeoffs involved in optimizing link performance for any given application. While architecture-level analysis [68] can

give useful insight to identify major bottlenecks, for deeper understanding a physical simulation tool, such as that developed here, is necessary.

The schematic layout of a WDM link is shown in Fig.2-12.a. The link uses a WDM laser block, which outputs a five continuous wave frequencies (but adding their phase ramps together). The output is passed through five modulators and five ring filters connected in series. The drop port of each ring goes to a photodiode. Four of the modulators are driven by four different PRBS data streams. The fifth is driven by a clock stream. The simulated input and output electrical waveforms are shown in Fig.2-12.b. Note that the optical rise and fall times and signal delays are neatly captured. Since all five channels are represented as one complex field signal in the waveguide, very subtle crosstalk can also be seen from the modulators modulating the neighboring channels slightly. Depending on other tradeoffs, the optimal channel spacing might be widened to reduce this crosstalk or pressed even closer if the cross talk is still acceptable and bandwidth is tight. Either way, our toolkit will allow the designer to inspect the resulting bit-error-rates (BER) in the presence of channel crosstalk. Similarly, other non-idealities, such as thermal effects, thermal noise, transistor noise, can be added to the simulation to provide a more realistic prediction of a systems behavior.

As a counterpoint to the above digital link, we next explore the operation of a purely analog, and extremely phase sensitive system, a Pound-Drever-Hall (PDH) locking loop, using our code. This set of simulations is also more complex as it additionally makes use of noise blocks and of standard CMOS cells to model more complicated electrical driving circuits. The theory of PDH operation is described in [69] and in Chapter 5 of this thesis. The basic concept is that the output of a laser is sent through a phase modulator modulated by a sine wave from an RF oscillator and then through a resonant cavity, in our case a ring resonator. The optical carrier picks up a different phase-shift in relation to the modulation sidebands depending on its frequency relative to the ring reference frequency. The carrier and the sidebands are beat together on a photodiode and the resulting RF signal is mixed down to baseband using the same, equally delayed, signal from the RF oscillator. The DC

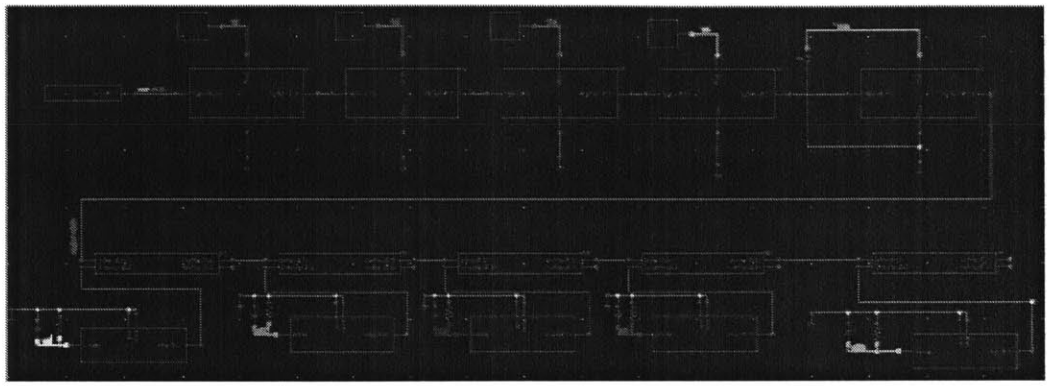

(a)

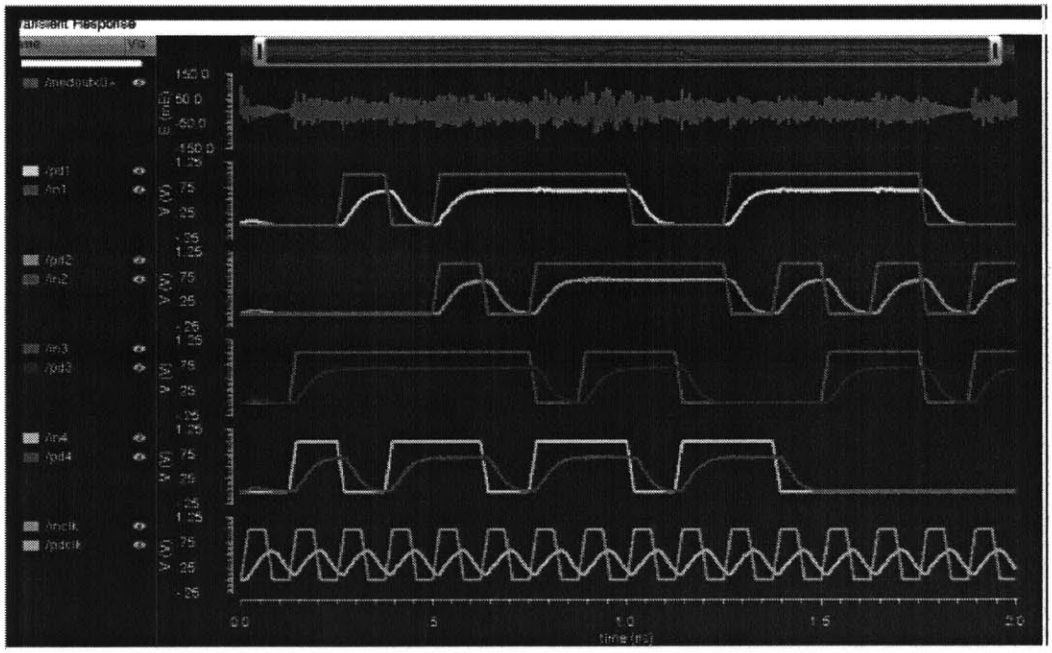

(b)

Figure 2-12: (a). Schematic of WDM link with four PRBS channels and one clock channel. (b) real part of signal on WDM waveguide (top), input and output of modulator channels (bottom four channels).

component of the output is the PDH error function and its sign indicates whether the laser frequency needs to be adjusted up or down to match that of the cavity. The error function is fed back on the laser and decreases laser phase noise by "locking" the laser frequency to that of the resonant cavity. The PDH scheme has two operation regimes: a low modulation frequency one where both the laser and the modulation sidebands fall within the cavity resonance, and a high modulation frequency one where the sidebands remain outside of the cavity.

A diagram of this system is shown in Fig.2-13.a, along with a schematic layout of ideal photonic and electrical components. The error function output in both the low modulation and high modulation regimes is shown in Fig.2-13.b., where they are compared to results generated using hand calculated transfer-functions and joint time and frequency domain simulations in MATLAB. As in the SSB case writing and running the MATLAB code to simulate the output of the PDH took more time and was more complicated than setting up the schematic and running the VerilogA code. Note that the VerilogA output closely matches the MATLAB output.

To further demonstrate the code, we add phase noise to our laser using standard VerilogA noise blocks, and modify the schematic to close the loop and suppress the laser phase noise (we allow the resonator in the simulation to remain ideal and noiseless). A noise analysis is performed in Cadence Virtuoso, as would be done for purely electrical circuits. The calculated laser phase noise in both the open loop (unsuppressed) and closed loop (suppressed) case are shown in Fig.2-13.c. For this particular set of loop parameters, mid-frequency phase noise is suppressed by >10dB.

To predict the performance of a real system with non-ideal components, a new schematic is created using accurate, non-idealized blocks for the photonics components and the actual CMOS driving circuitry. The parameter values in blocks for the photonics components are based on measured results from devices our group has designed and tested in the lab. The driving circuitry blocks shown in the schematic are designed in-house using the CMOS component models provided by the CMOS foundry, to satisfy the needs and constraints of the photonic components and of the full system. Fig.2-14.a. shows a schematic of this full system. As can be seen, this

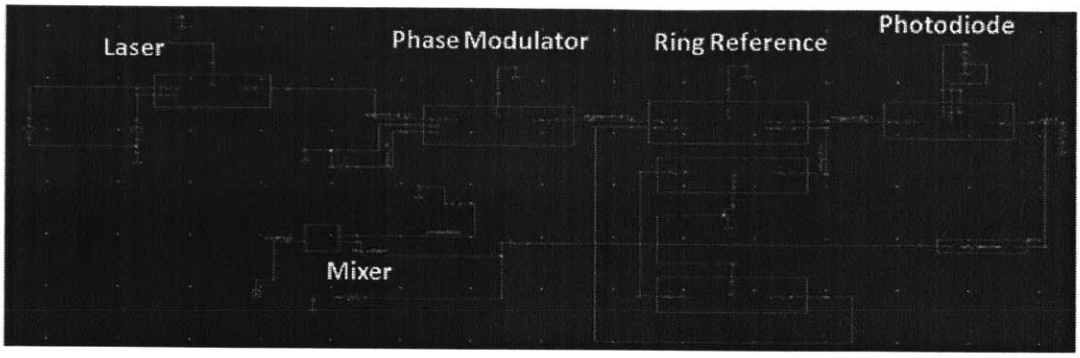

(a)

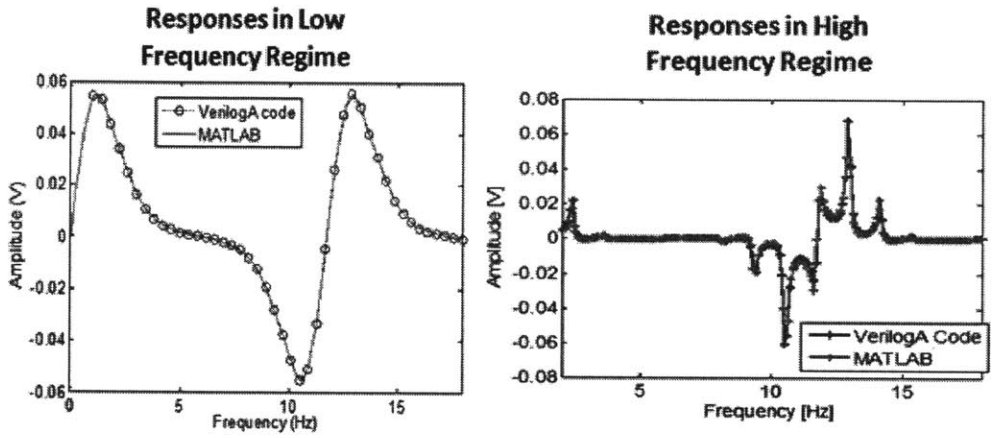

(b)

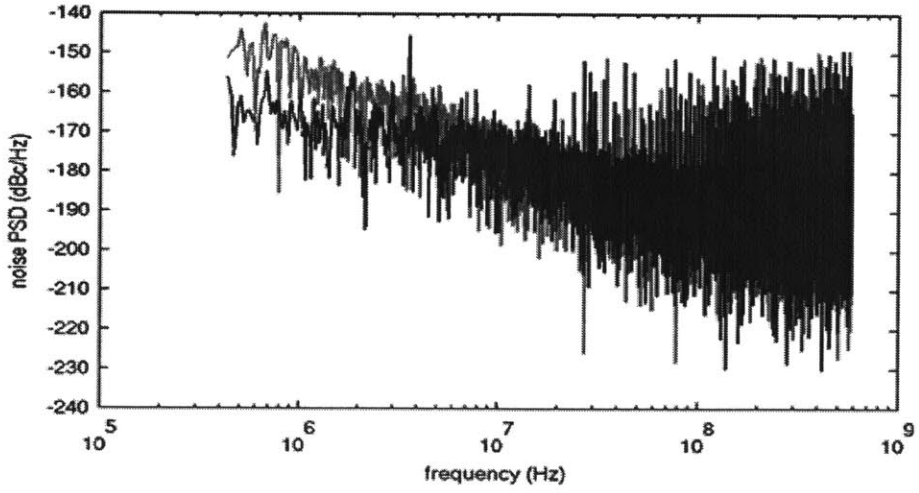

(c)

Figure 2-13: Schematic of PDH loop (after citen19) and schematic layout in Cadence Virtuoso. (b) PDH error function in low and high frequency regime, showing good agreement between MATLAB and VerilogA code. (c). Phase noise simulation done using our toolkit showing noise of the free-running CW laser (red) and of laser locked to cavity (black) showing successful phase-noise suppression.

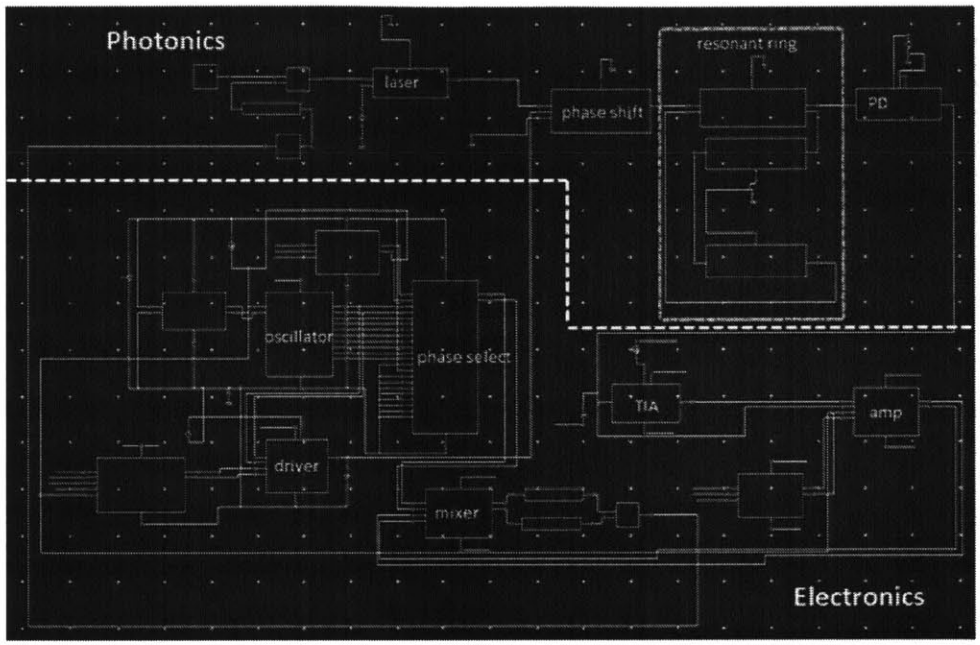

(a)

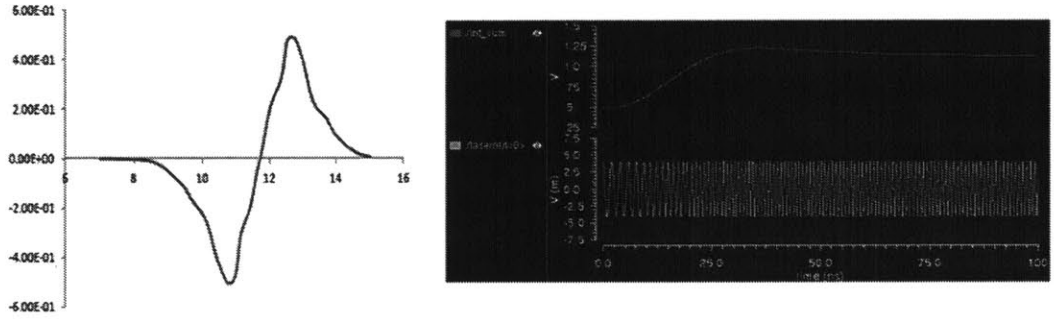

(b)

Figure 2-14: (a) Schematic layout of PDH loop using nonideal photonics and full driving circuitry in Cadence Virtuoso. (b) Left: Output PDH error function in low frequency regime. Distortions are clearly visible, but sign is still correct allowing loop to lock. Right: Laser frequency locks to cavity reference in 100nm (red), real part of laser electric field as locking occurs (green).

schematic is significantly more complex that the idealized schematic in Fig.2-13.a.

Simulations of the whole system can then be run, and performance achievable in the lab can be predicted. Furthermore, sensitivity to operating point, operating conditions, fabrication variations, various optical and electronic noise sources, and changes in component type or design can be explored. This allows for a much better understanding of system performance and robustness than previously possible. Predicted output of the current system at a chosen operating point is shown in Fig.2-14.b. The effects of modulator and oscillator nonlinearity are clearly visible in the ripples present in the DC error signal. However, it can also be seen that the sign of the error signal is preserved, and therefore that the PDH loop will still lock. A closed loop simulation further shows this, as the signal setting the laser frequency changes and then stabilizes to the cavity frequency. Thus, our VerilogA toolkit allows us to successfully and usefully model the performance of complex hybrid systems.

2.3.6 Advanced features

In addition to the above there are some advanced features that we would like to be able to add to the optical primitive blocks that are not straight forward. In particular, because it is a time domain code, adding waveguide dispersion is difficult. (We note that adding only waveguide nonlinearity is trivial — since this is a time domain code fields can simply be squared, cubed, or etc... inside the waveguide blocks). To do so it is necessary to write a block that does convolution in the time domain (as multiplication in the frequency domain is equal to convolution in the time domain). Before we get to this however, we derive the time domain transfer function for a waveguide including dispersion.

In order to add dispersion into the waveguide we follow the same general strategy as for the dispersion-less waveguide, and start with:

$$\tilde{E}(z,t) = \int_{-\infty}^{\infty} \tilde{E}(0,\omega) e^{(-j\beta z + j\omega t)} d\omega = \int_{0}^{\infty} \tilde{E}(0,\omega) e^{(-j\beta z + j\omega t)} d\omega \qquad (2.8)$$

Where $\beta = \beta(\omega)$ is the wavevector in the waveguide and the second equality follows

because there are no negative frequency components. Again, in order to avoid simulation at the optical time step, we choose some reference frequency, ω_R, such that $\omega = \omega_R + \Delta\omega$ and shift our signal down by this frequency. Again, we first Taylor expand $\beta(\omega)$ around this frequency, however here we keep the second order term, to account for dispersion (note that, right now, we are not accounting for higher order dispersion):

$$\beta(\omega) = \beta(\omega_R) + \Delta\omega \frac{\partial\beta}{\partial\omega}\Big|_{\omega_R} + (\Delta\omega)^2 \frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} \tag{2.9}$$

Where again we note that $\frac{\partial\beta}{\partial\omega}\Big|_{\omega_R} = \frac{1}{v_g}$, where v_g is the group velocity of light in the waveguide at the reference frequency and is equal to c/n_g, where n_g is the group index. This brings us to:

$$\tilde{E}(z,t) = \int_0^\infty \tilde{E}(0,\omega_R+\Delta\omega) e^{(-j\beta(\omega_R)z - j\frac{\Delta\omega}{v_g}z - j(\Delta\omega)^2\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z + j(\omega_R+\Delta\omega)t)} d\omega \tag{2.10}$$

As before we can define a function $\tilde{E}_{shift}(z,\Delta\omega) = \tilde{E}(z,\omega_R+\Delta\omega)$ which re-labels the frequency to shift the amplitude coefficients by the reference frequency.Using this function, we can change the variable of integration to $\Delta\omega$:

$$\begin{aligned}\tilde{E}(z,t) &= e^{-j\beta(\omega_R)z + j\omega_R t} \int_{-\omega_R}^\infty \tilde{E}_{shift}(0,\Delta\omega) e^{-j\frac{\Delta\omega}{v_g}z - j(\Delta\omega)^2\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z + j\Delta\omega t} d\Delta\omega \\ &= e^{-j\beta(\omega_R)z + j\omega_R t} \int_{-\infty}^\infty \tilde{E}_{shift}(0,\Delta\omega) e^{-j\frac{\Delta\omega}{v_g}z + j\Delta\omega t} e^{-j(\Delta\omega)^2\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z} d\Delta\omega \end{aligned} \tag{2.11}$$

Where, again, the second equality follows from the fact that there are no frequency components between $-\omega_R$ and $-\infty$. We note that the integral is just a multiplication of two functions: the Fourier transform of $\tilde{E}_{shift}(z,t)$ that we had before in the ideal waveguide case and an exponential in $(\Delta\omega)^2$. A multiplication in the Fourier domain is a convolution in the time domain (of the inverse Fourier transform of the two quantities). The first is the same as we had before in the simple waveguide case:

$$\tilde{E}(z,t) = e^{-j\beta(\omega_R)z + j\omega_R t} \tilde{E}_{shift}(0, t - \frac{z}{v_g}) \tag{2.12}$$

The second quantity is:

$$g(z,t) = \int_{-\infty}^{\infty} e^{-j(\Delta\omega)^2 \frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z + j\omega_R t} d\Delta\omega = \sqrt{\frac{2\pi}{\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z}} e^{-j\left(\frac{t^2}{2\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z} - \frac{\pi}{4}\right)} \tag{2.13}$$

Putting these two together yields:

$$\tilde{E}(z,t) = e^{-j\beta(\omega_R)z + j\omega_R t} \tilde{E}_{shift}(0, t - \frac{z}{v_g}) \otimes \sqrt{\frac{2\pi}{\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z}} e^{-j\left(\frac{t^2}{2\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z} - \frac{\pi}{4}\right)} \tag{2.14}$$

We recall a similar line of reasoning yields: $\tilde{E}(z,t) = e^{j\omega_R t}\tilde{E}_{shift}(z,t)$ Finally, combining these two equations and recalling that $v_g = c/n_g$ gives:

$$\tilde{E}_{shift}(z,t) = e^{-j\beta(\omega_R)z} \tilde{E}_{shift}(0, t - \frac{z n_g}{c}) \otimes \sqrt{\frac{\pi}{\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z}} e^{j\left(\frac{t^2}{4\frac{\partial^2\beta}{\partial\omega^2}\Big|_{\omega_R} z} - \frac{\pi}{4}\right)} \tag{2.15}$$

Where at the end of the waveguide $z = L$, and where $\tilde{E}_{shift}(z,t)$ is the signal we pass around in our code, and where again we will add a propagation loss term, α_{field}. Modeling this in VerilogA, however, becomes a little tricky, because VerilogA does not support convolutions by arbitrary functions, we therefore write a block that keeps an array of the last n time steps and performs a convolution. In order for this to work, the simulation time step needs to be set to be constant and input to the block as a parameter. It should be noted that while this block works, it is a little finicky and difficult to use. In particular it is necessary to carefully choose the time step and the time window over which the convolution is performed in order to prevent clipping of the pulse or other numerical problems. It is also necessary that the waveguide be long enough as similar numerical issues occur in very short waveguides.

2.3.7 Summary and Future work

As can be seen our toolkit provides a versitile platform written in VerilogA for simulating optical devices in conjunction with electronic circuits. This toolkit has been presented and verified. Both the real and the imaginary parts of the optical electric field are simulated allowing for the capture of phase-phenomena and interference effects, and optical electrical interactions are also captured where relevant. A set of fundamental photonic components are identified and each is described at the physical level, such that the characteristics of a composite device will be created organically. We show that resonant effects in composite devices such as ring filters and modulators and phase-effects in composite devices like single sideband modulators are correctly, organically captured by these models. We further show the ability to simulate complete systems including a WDM link, and a Pound-Drever-Hall laser stabilization loop complete with standard cell models of all the CMOS driving circuitry and noise. Our toolkit opens the possibility for system designers to easily build and simulate complex heterogeneous electronic+photonic circuits.

To make this tool kit work we chose to work with the analytic signal,and keep the complex representation of the E field. We further chose to work in the real and the imaginary parts to facilitate numerical convergence. We wrote the complex algebra modules necessary to make this work in VerilogA and also added optical and thermal disciplines to the program. We calculated time domain transfer functions for a set of primitive optical components and verified that these components worked by checking them against MATLAB results where possible. We then went on to succssfully simulate larger, complex systems using these primitive blocks. However, there is more work to be done, specifically it would be good to have a way to model both dispersion and nonlinearity at the same time. This may necessitate droping down into some sort of spilt-step code written in another language. Adding more components as needed, is also on the list, as is trying to find a way to handle antenna design, and adding in more parasitic effects.

Finally, we aim to link these VerilogA models to layout representations of the

components. This brings us back to the physical architecture part of the PDK. Here our group inherited the basic physical part of the PDK and have expanded on it. In the future, we aim to link each layout to a VerilogA model and to add LVS extraction to photonic devices to enable true LVS checking. Further, we aim to imporve the DRC rule deck. Because photonic shapes are often not on a Manhattan grid, the current DRC deck gives many spurious error. Eliminating these errors would be beneficial to designers.

In conclusion we outline a full platform that can be used to create photonic+electronic chips.

Chapter 3

Developing Modulators for Analog Applications

Now that we have a photonic platform in which to work, we can look at the ability to make devices in that platform. Key components in many analog optical systems include waveguides, ring resonators, photo-detectors, modulators, and lasers. While work has been done during this thesis on all of these components, the focus has been on modulators (and phase-shifters) and mode-locked lasers. This chapter will focus on the work on modulators, while the next will focus on the work on mode-locked lasers.

There are two main sections of the modulator work. The first focuses on the development of an integrated Mach-Zehnder Modulator with the aim of creating a very linear device. The design of the modulator phase-shifter is also optimized to achieve a cutting-edge value for $V_{\pi}L$ and to allow for good performance. The second section focuses on the demonstration of an integrated single-sideband modulator. To the best of the author's knowledge this is the first demonstration and characterization of an integrated silicon photonic version of this device.

The work presented in this chapter is based off work first presented in [94] and [65], two papers in preparation.

3.1 Integrated Linear Mach-Zehnder Modulators

A high degree of linearity is important in many analog systems such as radio-over-fiber, high-definition cable television, subcarrier transmission, wireless networks (including wireless local area networks), analog-to-digital converters, other types of signal processing, and analog optical links. Antenna and sensing remoting applications, where signals picked up in one location are transferred, before digitization, to another location for processing (often because of size, weight, or power limitations at the receiver) are another widespread, important use case for analog optical links, and are common in drilling and field operations [98, 99]. In these systems, the modulator often provides the limit for system linearity. Low nonlinearity is important because nonlinearits create distortions in the signal that can limit system performance (for example, limiting the achievable effective number of bits in an analog-to-digital converter). Because the harmonics grow faster than the fundamental with input power, one solution to decrease unwanted contributions from nonlinearities (spurs) is to decrease modulation depth. However, this limits the signal-to-noise ratio (SNR) of the system as it necessarily decreases the amount of power in the fundamental signal tones. Noise may then take over from nonlinearity in limiting system performance. It is, thus, desired to find a way to make the optical modulator transfer function more linear and to limit distortions at high SNR. In this chapter, a scheme to achieve good linearity over a wide rf and optical bandwidth by using silicon phase-shifter nonlinearity to cancel out Mach-Zehnder transfer function nonlinearity and thereby linearize Mach-Zehnder modulators is presented.

3.1.1 Previous Art and a Note on Units

Many schemes have been proposed to create modulators that are more linear [95, 96, 97, 100], but very few have been demonstrated in integrated systems. Furthermore, simple schemes that achieve high linearity, enough to be used in links with linearity greater than $120 dB Hz^{2/3}$, integrated or not, are hard to come by. Standard linearization techniques can be roughly split into two different categories: electrical or optical

schemes. In Mach-Zehnder (MZ) modulators, electrical schemes focus on modifying the driving circuitry or electrical responses to somehow compensate for modulator distortions. Examples include pre-emphasis (pre-distortion of the input signal to counteract nonlinearities), feedback schemes, and reversing the transfer function in post-processing [96, 98]. Optical schemes aim to modify the optical system or components such that the new transfer function is more linear that that of the modulator on its own. Examples include cascading modulators, splitting signals, adding elements (for example, rings), and modifying elements (for example, using non-3 dB couplers). While the resulting systems, electrical or optical, do work, they are usually more complicated, larger, and more power-hungry than a single MZ modulator [96]. They often also have other, scheme specific, drawbacks. For example, pre-emphasizing the signal, while simple to implement in practice, significantly reduces the usable RF bandwidth of the modulator, as the modulator must operate at harmonics of the desired signal in order to cancel them out [96]. Post processing relies on an understanding of the other elements in the channel in order to back out the original signal and can become impractical if there are many nonlinear elements after the modulator or if the transfer function of the channel or elements between the modulator and the post-processing circuits is unknown. For optical schemes, complexity and the necessity for careful balancing, plus larger power consumption, are usually the main drawbacks.

While all the above schemes have been proposed and shown to work for discrete component systems, a smaller, though significant, amount of work has been done on linearity in integrated systems. Before we discuss this work, however, we wish to discuss an important point about units: there are many different metrics that people use to measure linearity. These include both input and output IP3, dB, dBc, CDR3, and SFDR (which is expressed in $dBHz^{2/3}$ for the third harmonic – often the limiting harmonic in these systems). To understand the literature on this topic, and to be able to compare performances (whether measured or predicted), it is, thus necessary to be able to convert between these metrics. The first step to doing this is to be able to distinguish between metrics that reflect just the linearity of the modulator (IP3, dB, dBc, CDR3) and those that reflect the performance of the entire link (SFDR).

Because measurements of linearity always have to go into some sort of measurement equipment and can also be affected by nonlinearity in the input signal, they are always to some extent affected by systems. The photodetector, used detect the modulated signal so that it can be input into an ESA, in particular can often contribute to measured nonliearity for very linear modulators. However, beyond these parasitics, there are more fundamental ways that measurement systems can effect these results. Take for example, Spur Free Dynamic Range (SFDR). The SFDR of a system is defined as the height of the fundamental signal above the tallest unwanted signal, and because nonlinear signals grow more with power, is usually defined as the height of the fundamental above the noise floor when the highest spur is no longer visible above the noise. This means that the noise figure (NF) of a system - often set by the sensitivity of the photodetector and available laser power, not the modulator, among other things is critical. This will be discussed in more detail below.

In contrast, the third order intercept point (IP3) is, except for parasitic, more of a measure of just the modulator performance. However, even here there are important caveats. First, the input(output) IP3 is defined as the input(output) power at which the amount of power in the fundamental, desired signal is equal to the amount of power present in the third harmonic or the third order intermodulation tone. The number achieved depends on whether the harmonic or intermodulation tone is used. In particular the IP3 based on an intermodulation tone will be 4.8dB *lower* than that based on the third harmonic. Thus, it is important that this choice be specified [96]. Further, it is important to know whether the input or output IP3 is used. The input IP3 is more robust to measurement, though it can be affected by drive power requirements and loss in the modulator. The output IP3 can be affected by how much power is put into the system from the laser. One thing to note is that IP3 is really only useful for weakly nonlinear devices and is generally calculated by taking several input power vs. output power points and extrapolating to the intercept point. The intercept point itself is often above the safe operation threshold of the device. It should also be noted that using this metric assumes that the third harmonic is the distortion that limits performance, an assumption that works for integrated modulators. In other

systems IP2, IP5, etc might be more appropriate [96].

Compared to IP3, dB, dBc, and CDR3 are even more slippery. Often they are simply various ways of saying the difference, in dB, between the power in the fundamental and that in the third harmonic (or occasionally the third order intercept point) at the output of the system for some arbitrary, but specified, set of input conditions. Thus, getting an apples to apples comparison when using these numbers is very tricky — the exact way the system is operated will often have an effect on these metrics and it is important to read the fine print when comparing numbers. While imperfect, it seems that the literature has standardized on input IP3 (IIP3) and SFDR as the most useful of these metrics, despite the fact that the latter measures performance in a link. Often, too, achievable modulation depth over which the intermodulation distortion can be suppressed by a given amount will be specified, so as to indicated what performance may be possible in a real system. In the discussion of results in the literature below, the linearity specified in the paper will be given, but, as it is the most stable way to measure linearity in a modulator, the IIP3 will also be provided if it was possible to determine it from the data presented.

First, [99] provides an experimental and analytic analysis of a silicon MZ modulator, similar to the one used here. A carrier-to-third-order distortion ratio (CDR3) — that is the difference between the power in the carrier and that in the highest third order intermodulation term — of -50dB is measured for a 5GHz signal. Similarly, in [101], which cites the theoretical work described here, measuments of a MZ modulator link resulted in an SFDR of 97dBHz$^{2/3}$ and an input IP3 of 42dBm (this second figure being read off the graphs shown in the paper). These papers simply measured the linearity of existing devices, but did not attempt optimization of the device to enhance performance. A long the same lines is [102], which also cites the theoretical work described here as well as that in [99], and which creates an analytic model and compares it to experiment in what is basically a combination of the work in [98] and [99]. It appears to have best SFDR measurements around the mid-80dBHz$^{2/3}$ and best input IP3 measurements around 35dBm. Again they do not appear to optimize their devices.

There have also been attempts to create more linear integrated modulators. [103] proposes an integrated linear modulator made from a polymer on a silicon substrate and using a modulated directional coupler as the underlying structure. A measured link linearity of 110dBHz$^{2/3}$ is achieved in a frequency range of 2-8GHz, which here corresponded to an input IP3 of 40dBm. [104] demonstrates a table top system that uses injection locking of a resonant cavity containing a gain element with a linear response to achieve modulation with a linearity of 120 dBHz$^{2/3}$, but this system isn't integrated. In [105], which cites our work described below, the authors propose a more compact variation made with SiGe heterojunction bipolar transistor phase shifters that is predicted to achieve a linearity of "42dB." In, [106] a scheme for linearizing the device using a ring resonator cavity was demonstrated and an SFDR of 106dBHz$^{2/3}$ was achieved, the input IP3 was 25dBm (again as read off the graph in the paper), but the results only hold over a narrow optical bandwidth.

3.1.2 Approach and theory

We desire a highly linear, high speed, low noise, optically wideband, versatile, integrated linear modulator. We choose to work with a Mach-Zehnder configuration as it is optically broadband and less sensitive than ring configurations to heating and other environmental noise. We use silicon reversed-biased diode phase-shifters that rely on the plasma dispersion effect to achieve phase shifts. The plasma-dispersion effect in silicon diode phase-shifters is chosen because of its high speed and sensitivity and comparative easy of fabrication in relation to other options. More detail on this choice can be found in Appendix A.

In previous work, we proposed a scheme for linearizing these devices [98, 107, 36]. In this thesis, we demonstrate this scheme. As mentioned, we use silicon reversed biased diode phase-shifters (RBDPS) in a Mach-Zehnder configuration. Because the MZ transfer function is sinusoidal, even perfectly linear phase-shifters do not yield linear operation. Further, silicon RBDPSs are themselves nonlinear and could further add to the nonlinearity of the device. However, because these are two different sources of nonlinearity it is theoretically possible to design one to counter act the other. Because

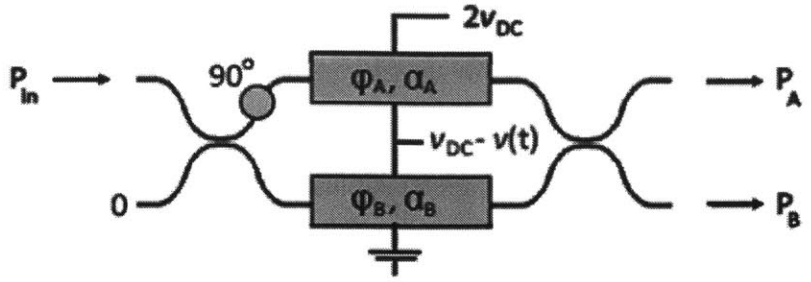

Figure 3-1: Schematic of MZM Layout. From [98].

such as scheme balances the intrinsic components of the device, our approach is simple to implement — it does not require extra devices or complex splitting, feedback, or post-processing, merely the careful choice of an operating region. Furthermore, it does not adversely impact device performance, since the modulator need not function at the harmonics of the signal; it is the harmonics that are automatically generated which will cancel. Nor does the modulator size or power consumption increase. The theoretical work and early simulation results were presented in [98, 107, 36], but will be summarized here. Some additional results on fabrication sensitivities will also be discussed in this section. Design of the linearized modulator along with detailed simulation and measurement results will be discussed in the following sections.

To see how the proposed scheme works, we calculate the Mach-Zehnder transfer function allowing the phase-shifting sections in the arms to have an arbitrary, voltage dependent phase and loss. A schematic of the device is shown in figure 3-1. The MZ modulator is assumed to be biased in push-pull, to have a heater section that adds an additional phase, ϕ_h, in one arm, and to have a phase-shifting section in each arm. The heater allows the modulator to be biased in quadrature. Let $\phi_1 = \phi(v_{DC} + v)$ represent the phase added by the rf phase-shifter in the top arm as a function of voltage, where v_{DC} is the dc bias voltage and v is the rf signal. Let $\phi_2 = \phi(v_{DC} + v)$ be the equivalent for the bottom arm, and $\alpha_1 = \alpha(v_{DC} + v)$ and $\alpha_2 = \alpha(v_{DC} + v)$ be the voltage dependent losses through the respective phase-shifters. It is further assumed that the input and the output couplers are ideal 50-50 splitters, that light is input only on the top arm, and that the input light has power P_0. Calculating the

transfer function for the modulator then yields [98]:

$$P_{top} = P_0 \exp\left[-\left(\frac{\alpha_1}{2}+\frac{\alpha_2}{2}\right)\right]\sin^2\left(\frac{\phi_h}{2}+\frac{\phi_1-\phi_2}{2}\right)+\frac{P_0}{4}\left[\exp\left(-\frac{\alpha_1}{2}\right)-\exp\left(-\frac{\alpha_2}{2}\right)\right]^2 \tag{3.1}$$

$$P_{bottom} = P_0 \exp\left[-\left(\frac{\alpha_1}{2}+\frac{\alpha_2}{2}\right)\right]\cos^2\left(\frac{\phi_h}{2}+\frac{\phi_1-\phi_2}{2}\right)+\frac{P_0}{4}\left[\exp\left(-\frac{\alpha_1}{2}\right)-\exp\left(-\frac{\alpha_2}{2}\right)\right]^2 \tag{3.2}$$

A Taylor expansion of the phase and losses gives:

$$\phi_m(v) \approx \left(\phi_{mDC} + a_m v + \frac{b_m v^2}{2} + \frac{c_m v^3}{6}\right)L \tag{3.3}$$

$$\alpha_m(v) \approx \left(\alpha_{mDC} + x_m v + \frac{y_m v^2}{2} + \frac{z_m v^3}{6}\right)L \tag{3.4}$$

where a_m, b_m, c_m, x_m, y_m and z_m are the Taylor coefficients, and ϕ_{mDC} and α_{mDC} are the DC phase and loss vaules. Combining Equs. 3.1 & 3.2 with Equs. 3.3 & 3.4, Taylor expanding the result, and simplifying yields:

$$P_{top} \approx P_0 e^{-\alpha_{DC}L}\left[\frac{1}{2}+(aL)v+\frac{1}{4}(x^2L^2-yL)v^2+\left(-\frac{2}{3}a^3L^3+\frac{c}{6}L-\frac{1}{2}ayL^2\right)v^3\right] \tag{3.5}$$

$$P_{bottom} \approx P_0 e^{-\alpha_{DC}L}\left[\frac{1}{2}-(aL)v+\frac{1}{4}(x^2L^2-yL)v^2-\left(-\frac{2}{3}a^3L^3+\frac{c}{6}L-\frac{1}{2}ayL^2\right)v^3\right] \tag{3.6}$$

where we have assumed that the modulator is biased in quadrature and the phase shifters in the arms are exactly identical (a discussion of the unbalanced case occurs below) [98].

To understand these equations, it is useful to understand the origin of each term. The first term in Equ. 3.5 is the output at the DC bias point. Here, the bias in quadrature simply splits the input light evenly between the two output channels. The second term is the desired, linear response of the modulator: it is simply the input RF voltage scaled by a constant determined by the linear response at the bias point and the length of the phase-shifting section. These first two terms are the same as what would be found in the case of a perfectly linear phase-shifter. However, the third term (proportional to v^2), would not be present in the ideal case.

It corresponds to the second order nonlinearity, and, assuming all higher harmonics are negligible, gives the second harmonic response of the device. Evidently, the second order nonlinearity comes only from the voltage-dependent loss. It would not be present if the phase-shift sections were lossless, or if the loss were independent of voltage. This seems to imply that the silicon modulator case is worse than the ideal linear case (in which the phase-shifters are, say, lithium niobate and relatively linear and lossless). However, this term has the same sign in both the top and the bottom arms, while the linear terms have opposite signs. Therefore, taking the difference between the output terminals will leave the linear term intact but cancel the second-order nonlinearity term. In fact, subtracting the two arms will cancel all even order nonlinearities (this can easily be understood by considering the alternating signs in the Taylor expansion of a sine function). Such differential detection also has the benefit of increasing the fundamental power by 3dB [98].

However, it is not always possible to use differencial detection. This is especially true in systems where there are many elements between the modulator and detector and where it is impractical, or impossible, to create two identical, but separate, paths for the two complimentary channels. In that case, it is possible to cancel the second order nonlinearity by moving the modulator operation point away from the quadrature point — that is by tuning the DC (heater phase) bias point of the device. This introduces an additional quadratic term into the modulator transfer function which can cancel the loss-induced quadratic nonlinearity in the desired modulator output (top or bottom) only. Here, the output of the top channel is given by:

$$\begin{aligned} P_{top} = P_0 e^{-\alpha_{DC} L} \Big[& \sin^2\left(\frac{\phi_h}{2}\right) + 2\sin\left(\frac{\phi_h}{2}\right)\cos\left(\frac{\phi_h}{2}\right) aLv \\ & + \Big[\Big(\frac{x^2L^2}{4} - \frac{1}{2}\sin^2\left(\frac{\phi_h}{2}\right)Ly\Big) + \Big(\cos^2\left(\frac{\phi_h}{2}\right) - \sin^2\left(\frac{\phi_h}{2}\right)\Big)aL\Big]v^2 \\ & + \sin\left(\frac{\phi_h}{2}\right)\cos\left(\frac{\phi_h}{2}\right)\Big[-\frac{4}{3}a^3L^3 + \frac{1}{3}cL - ayL^2\Big]v^3\Big] \end{aligned} \tag{3.7}$$

where ϕ_h is choosen to take the second-order nonlinearity to zero. Looking at the third-order nonlinearity, one can see that one retains the same equation for the fundamental and the third-order nonlinearity as before, just multiplied by a different

constant (but one that is the same for both). Thus, aside from a change in SNR, the movement in bias point does not affect the performance of the modulator.

It is also worth noting that another way to deal with the presence of the second-order nonlinearity is just to leave it. In many applications the signal bandwidth is relatively small and the second harmonic is far enough away from the band of interest that it can be filtered out. The third harmonic would also be filtered in this case, but the third order intermodulation terms present at $2\omega_{rf1} - \omega_{rf2}$ and $2\omega_{rf2} - \omega_{rf1}$ would generally still fall in the bandwidth of interest and, thus, pose a problem.

Finally, the last set of terms in equation 3.5 gives the third order nonlinearity, and, assuming higher order nonlinearities can be ignored, the third harmonic. (In the case where higher order nonlinearites cannot be ignored, all higher order odd nonlinearities will contribute to the third harmonic.) This term has the same sign as the linear term in both cases and so cannot be canceled with subtraction. The third harmonic is usually the limiting factor for the spur free dynamic range (SFDR) for MZ modulators [96].

To understand this third-order nonlinearity, we look at the sub terms. We note that the first of these terms, $-2a^3L^3v^3/3$, comes from the interaction of the linear phase change with the sinusoidal transfer function of the MZ structure, it therefore represents the third-order nonlinearity from the MZ transfer function itself. It is the only contribution to the third harmonic that would be present in a MZ modulator with ideal, lossless, linear phase-shifters, and it would (usually) place the limit modulator linearity and possibly, if modulator nonlinearity dominated, system performance. The next contribution to the third order nonlinearity, $cLv^3/6$, comes only from the third-order phase-shifter nonlinearity and is a direct result of the phase shifters themselves being nonlinear. The final contribution , $ayL^2v^3/2$, arises from the interplay between the MZ transfer function and the voltage-dependent phase-shifter loss. [36]. To achieve more linear performance than in the case of a MZM with ideal phase shifters, the combination of these second two terms simply has to be such that they counteract part of the first term. To fully cancel the third order nonlinearity, these two terms have to balance the first so that these three terms sum to zero. Looking at the equation for

the third order nonlinearity, one can see that it gives a quadratic equation in terms of length, and that the resulting value is zero when:

$$L = \frac{1}{4a^2}\left(\frac{-3y}{2} \pm \sqrt{\frac{9y^2}{4} + 4ca}\right). \tag{3.8}$$

. It should be noted that this cancels the third order nonlinearity, not the third harmonic, which will have contributions from the fifth, seventh, ninth, etc order nonlinearities. In cases where these higher order nonlinearities are sufficiently small, canceling just the third-order nonlinearity might be sufficient. In other cases, the equation governing the height of the third harmonic (or the third order intermodulation terms) will simply contain terms from these higher order terms in the tayor expansion. While this will give a slightly different equation for the cancellation L, the general idea remains unchanged and should still work.

We note that in the case of a linear-lossless phase shifter, $y = 0 = c$, and the expression for the length at which the third order nonlinearity vanishes yields zero. Thus, for a MZ modulator with linear-lossless phase-shifters, no choice of bias or length can lead to a useful cancelation of the third-order nonlinearity. In order for this linearization scheme to be useful for real applications, we need to show that the length calculated is amenable to high modulation depths in reasonable lengths, and that the design is tolerant to fabrication variations. These will both be shown in the next sections. Because showing the first requires simulation of the phase shifters themselves it will be covered in the next section. The question of fabrication variations, will be addressed first.

It is useful to consider the effects of fabrication variations and imperfections in experimental setup, so we know what to expect in the lab. In the above theory, we are seeking to balance two effects exactly. This is equivalent to trying to get a very small number by subtracting two large ones. Even a tiny fractional error in the large signals could destroy the cancellation. We wish to assure ourselves that we will be able to recover good cancellation even if the device and setup are imperfect. We first consider the case of mis-matched phase-shifting sections. Fabrication variations guarantee that

the two phase-shifting sections in the MZ modulator will not have exactly identical responses. It is therefore necessary to consider the impact of slight differences between the phase-shifters on linearity. Perhaps the most important qualitative difference is that, given different phase responses, use of the push-pull scheme will no longer allow for cancellation of the large second-order nonlinearity in the response of the phase. Performing the same set of Taylor expansions as above yields a set of coupled, transcendental equations for the second and third order nonlinearity [36]. Whether an acceptable length and bias phase can be found to solve these again depends on the exact coefficients. Because these equations are somewhat complicated, direct simulation, like that described in the nest section, using simulated or measured phase-shifters performance, is a necessary and more informative way to determine the details of the achievable linearity. However, it does turn out that, while the exact location varies, cancellation is still possible in many cases. This can be seen on an intuitive level, by considering the simplified case where the loss is assumed to be constant with voltage. Here:

$$\begin{aligned} P_{top} =& P_0 e^{-\alpha_{DC} L}\Big[\frac{1}{2} + aL\cos(2\Delta\phi)v + \Big(\frac{L\Delta b}{2} - 4L^2a^2\sin(2\Delta\phi)\Big)v^2 \\ &+ \Big(-\frac{2}{3}a^3L^3\cos(2\Delta\phi) + \frac{cL}{6}\cos(2\Delta\phi) - \Delta b a L^2\sin(2\Delta\phi)\Big)v^3\Big] \end{aligned} \tag{3.9}$$

where a and c are the sums of the coefficients in each arm: $a = a1 + a2$,$c = c1 + c2$ (as before, except, of course the values may be different), where $\Delta b = b1 - b2$, and $\Delta\phi = \phi_{mDC1} - \phi_{mDC2}$ are the differences in the coefficients between both arms, and again it is assumed that the heater is used to bias the modulator in quadrature [36]. We note that in the case where $\Delta b = 0$, $\Delta\phi = 0$, we recover the result in the symetric case where $y = 0$ as expected. The result is that the third harmonic again vanishes at a length very similar to the original one; since $\Delta\phi$ and Δb are presumably small, we can see that we only have a small tweak on the original equation and presumably only a slight adjustment will be needed. At first this might seem like a problem, as the length of the device cannot be adjusted in post processing. However, because of the nature of the nonlinearity in the phase-shifters (which have something approaching a

square root dependence on voltage), the values of the Taylor coeffcients are dependent on the choice of bias point and a slight movement in bias point should be sufficient to reestablish third order nonlinearity cancelation. However, in this situation the second order nonlinearity does not vanish. That is, if we cannot do differential detection and if the bias point is being used to correct for fabrication error and cancel the third order intermodulation term in a device of a desired length, then the bias cannot also be tuned to cancel the second harmonic. Thus, in a wideband signal the difference between the quadratic coefficients of the phase on the two arms and the shape of the voltage dependent loss will limit performance. If the contribution from these terms is small enough, the impact on linearity will be minimal; if they are large, achieving linear operation over a broad band, where the second harmonic cannot simply be filtered away, may not be possible.

From the above paragraphs it should be clear what changing the heater bias does - it changes the values of the Taylor expansion coefficients. It is therefore critical to bias the device at the correct point. Fortunately, as this is an active control, the heater bias can be adjusted in the lab until the best result is achieved.

The above theory is all well and good, but to truly understand more it is necessary to look at simulation and experimental results. The next section will deal with the design of a reversed biased diode silicon phase-shifter and its experimental performance.

3.1.3 Design and DC results

The above analysis is general and can be applied to any MZ modulator with non-ideal phase shifters. However, we wish to apply it specifically to the case of a silicon reversed-biased diode phase-shifter embedded in a rib waveguide and show that it can work in practice. Extensive simulations verifying that the lengths required to cancel the third harmonics in these modulators are real and reasonable (on the order of a few hundred microns — long enough to allow for good modulation depth and short enough to remain compact on chip) have been presented in [98] and [36] and will not be repeated here. Instead, we focus on the design and demonstration of an optimal

device.

To design the phase-shifter, process and electrical simulations were carried out in Synopisis's Sentaurus® software suite. Simulations in Sprocess® mimic all the fabrication steps, including doping and diffusion, and allow for an accurate simulation of the details of the junction profile. Simulation in Sdevice® than yield the electrical reponse of the device. The output of these simulations, 2D cross-sections of carrier concentrations, is sent to an optical modesolver and the change in phase-shifter refractive index and loss as a function of bias voltage is then calculated. This response is then fed into a system code to which models the full MZ modulator and outputs the predicted linearity. For more detail on this process see Appendix A or [36]. Based on the outputs, parameters such as doping locations, modulator geometry, and doping doses and energies are changed and the design iterated until the desired linearity is achieved. We aim to avoid changing doping doses and energies as these effect the designs of others as well.

Additionally, when designing the modulator we want to work with a vertical doping junction (where the p-doping is on top of, instead of just next to the n-doping) in order to maximize the overlap between the optical mode and regions of index change and thereby maximize phase-shift per length and achievable modulation depth and minimize device length. Further, we would also like our device to work at high speeds (>20GHz), and have reasonably low loss. Fortunately, it turns out that all these targets can be met in one device.

Ideally, we would also like to work low drive voltage, if possible something compatible with CMOS driving circuits, 1V, we find that the same device that promises to have good linearity can work well as a digital modulator at these voltages, but the best linearity point is at slightly higher bias voltages (2V).

A cross section of the optimized device is shown in Fig. 3-2.a. Fig. 3-2.b shows the predicted change in effective index and loss as a function of bias voltage. Before predicting the linearity, from these simulations we estimate the $V_{\pi}L$ of the modulator, when driven in push-pull, to be a very low 0.3 Vcm in push-pull. Next the linearity response of the device is simulated, and is plotted in Fig. 3-3. The predicted input IP3

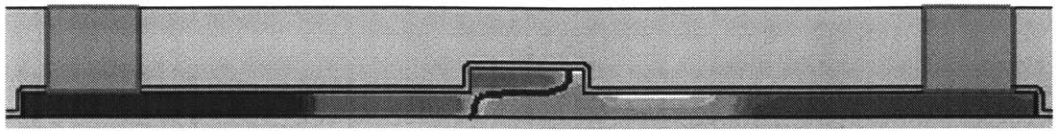

(a)

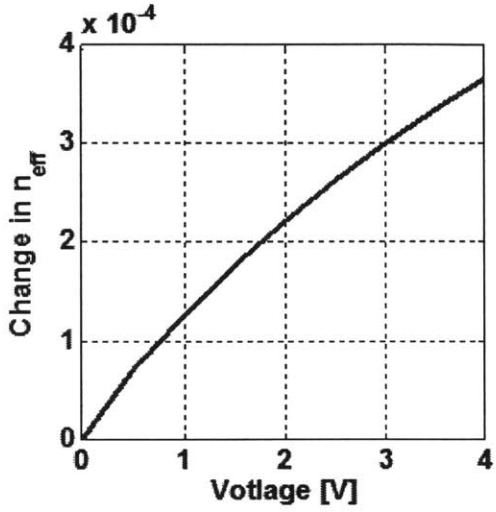

(b)

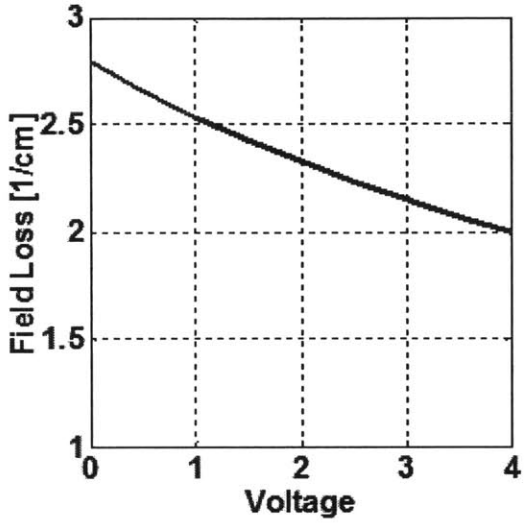

(c)

Figure 3-2: (a) Cross-section of phase-shifter showing doping distribution, (b) Simulated effective index and effective field loss of phase-shifter as a function of applied bias. From [94].

for the third harmonic of this device was 64dBm. It should be noted that the outcome of these simulations is somewhat sensitive to numerical parameters as detailed in [98], and there is therefore about $\pm 5 - 10dBm$ of uncertainty in this prediction.

A MZM using these phase shifters was designed and laid out using the Cadence framework described in Chapter 2. Adiabatic 3dB couplers were used at the input and the output of the modulator and an integrated heater was included in one arm in order to correct for fabrication errors and bias the device in quadrature. The device was fabricated at CNSE Albany also using the process described in Chapter 2. The device was tested on the setup diagrammed in Fig.3-4, and the DC output is shown in Fig.3-5. Fig.3-5.a. shows the the modulator output as a function of heater bias; clearly our integrated heater allows us to tune the MZM bias-point through the full 2π range. The power consumption of the heater was 38mW for a full swing (see

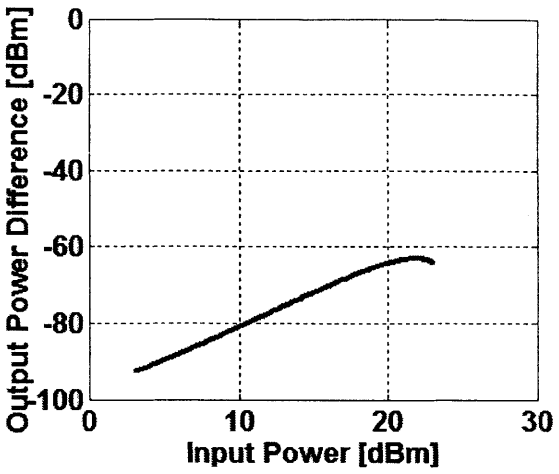

Figure 3-3: Simulation output showing predicted distance of the third harmonic below the fundamental as a function of input drive power. Predicted input IP3 was 64dBm. From [94].

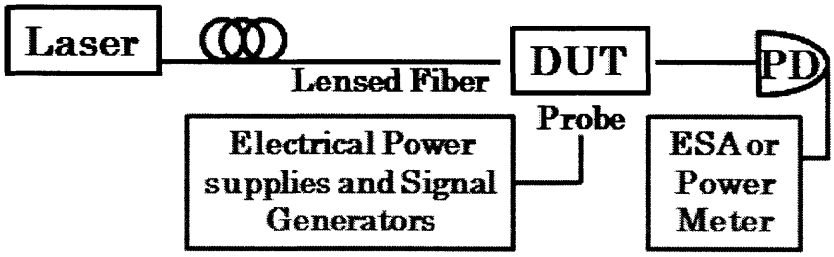

Figure 3-4: Schematic of Measurement Setup. From [94].

Chapter 5 for more details). Fig.3-5. b. shows the modulator output as a function phase-shifter bias voltage in push-pull and shows that a full π phase-shift can be swept through with only 1.35V in a 2mm device, yielding a $V\pi L$ of 0.27Vcm, which is in good agreement with our predicted value. This number can be further verifies by inverting the MZM transfer function to extract the resulting phase change as a function of voltage. These curves are shown in Fig.3-5.c. for MZMs with phase-shifters of different length, and also yield a $V\pi L$ value of 0.27Vcm. To ensure that the highly doped junction is functioning as desired and is not in danger of breaking down, a current sweep of this device was also measured and is shown in Fig.3-5. d. The predicted breakdown voltage was 7V reverse bias, as is the measured breakdown voltage.

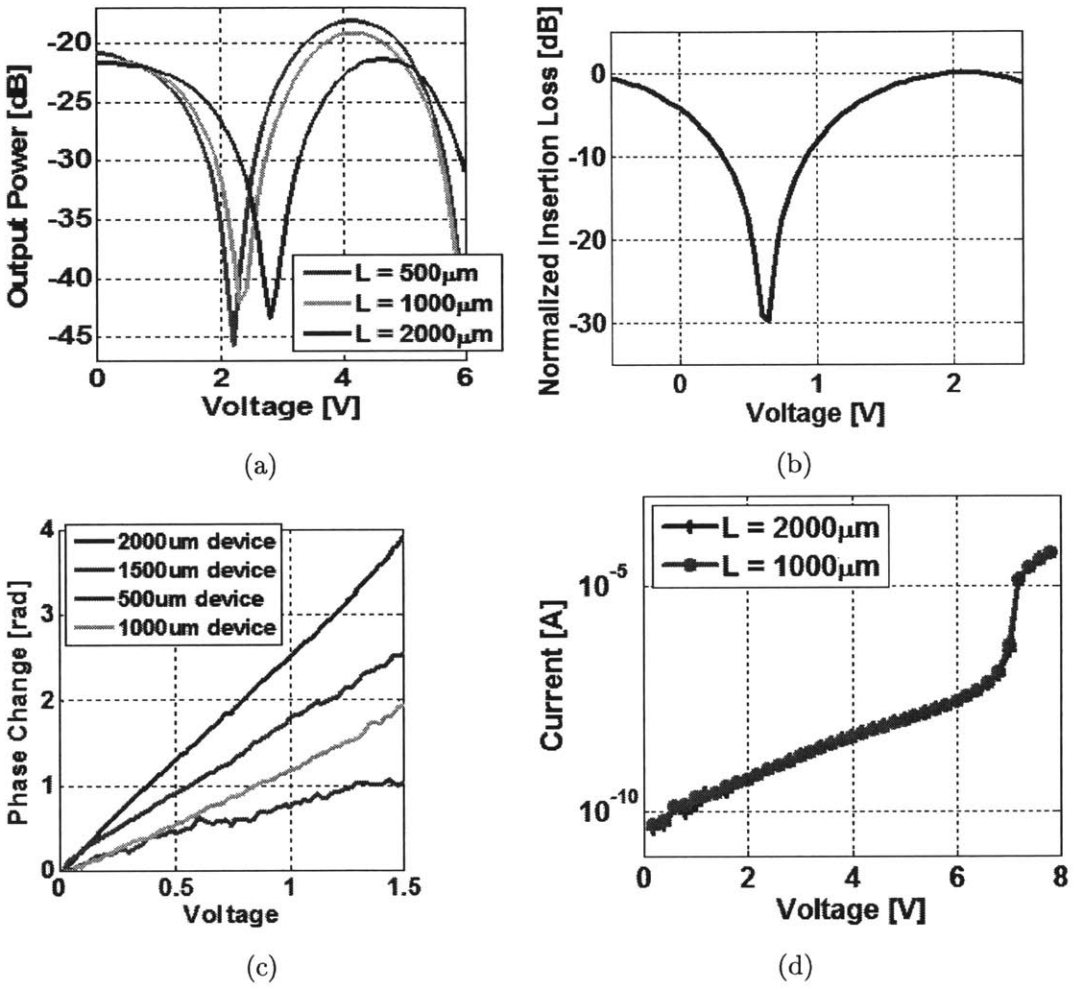

Figure 3-5: (a) Output of MZM as a function of heater bias, showing the full range of the output can be achieved. Heater is 100μm in length, length of phase-shifter section (unbiased) varied from 500-2000μm (b) Output of the MZM as a function of DC phase-shifter bias, showing full device operating range can be achieved, (c) Extracted phase as a function of bias for devices of different length. Results match well to original theory. Measured VπL = 0.27Vcm. (d) Graph of reverse-bias current showing breakdown around 7V. From [94].

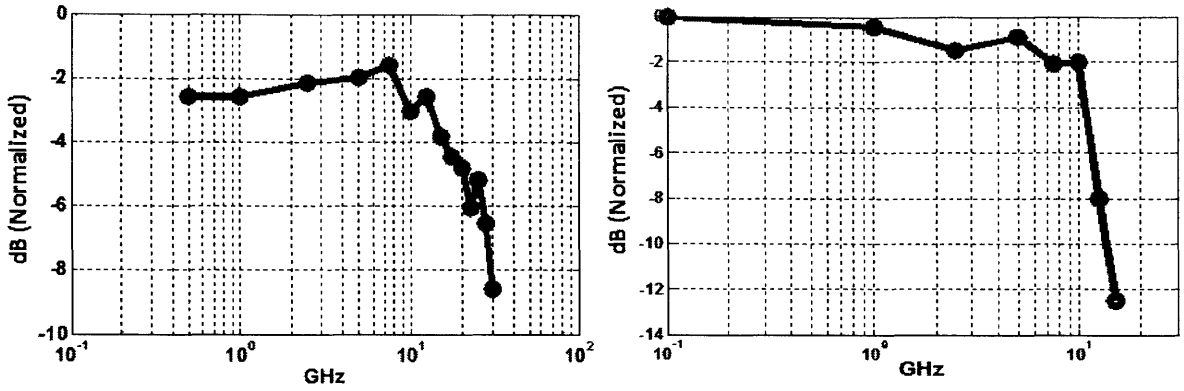

Figure 3-6: High speed performance of MZM, showing 3dB roll-off of 25GHz in 550μm device (left) and 10GHz in 2mm device (right). From [94].

3.1.4 AC and Linearity results

A signal generator, photodetector, and electrical spectrum analyzer were used to test the device. The ac response of the 2000μm device as a function of frequency and the 550μm device (from a different wafer) are shown in Fig. 3-6. The 2000μm device was found to have a roll-off of 10GHz, and the 550μm device was found to have a 3dB roll-off of 25GHz. Next the linearity was measured on the 550μm device — which was expected to perform best. The DC bias of the phase shifters in the device was set to $2V_DD$ on the top contact and 1V on the middle contact, and the fundamental and third harmonic powers as a function of the input driving power of a 5GHz ac signal were measured. The results of this measurement, plotted against the theoretical prediction are shown in Fig.3-7. Good agreement is observed between experiment and theory. The input IP3 of the device was measured to be 68.4dBm, matching well with the prior, previously explained simulation results. It should be noted that this value is >20dB better than the best in the literature for integrated modulators, and 28.7dBm better than the theoretical value, 39.7 dBm, for a MZM with ideal, linear lossless phase shifters and equivalent modulation depth [96].

For a practical operating point with an 55% modulation depth we measured a 3rd harmonic suppression of more than 60dB below the fundamental, more than enough for many analog application. For example this would allow for 10 ENOB to be achieved in a photonic-assisted ADCs [96, 135]

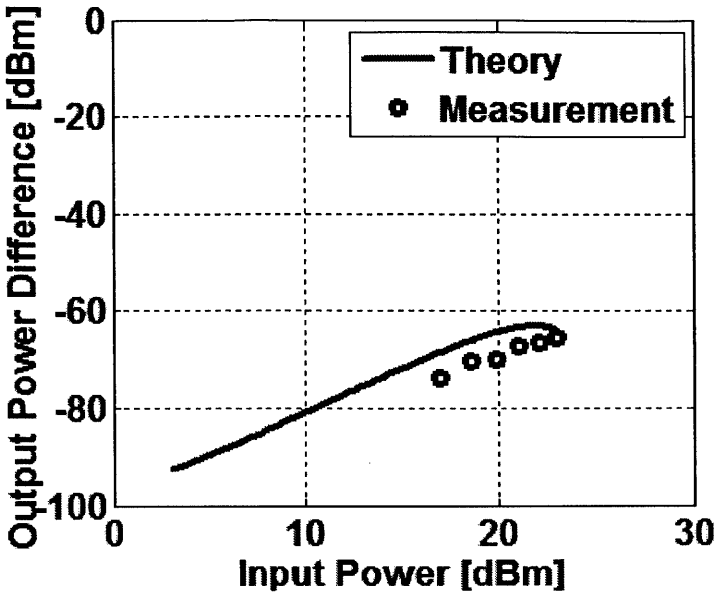

Figure 3-7: Comparison of measured and simulation output of predicted distance of the third harmonic below the fundamental as a function of input drive power. Good agreement is seen between simulation and theory. Predicted input IP3 was 64dBm. Measured input IP3 was 68.4dBm. From [94].

Finally, SDFR of the system was measured, and the results are shown in Fig.3-8. The SFDR was measured to be 82.1dBHz$^{2/3}$. The SFDR was limited in part by photodiode nonlinearity and in part by the noise floor of the system which was at -127dBm (most measurements in the literature are taken with a noise floor of -160dBm. If our noise floor were at this level we would expect to measure an SFDR of 104dBHz$^{2/3}$). Future work consists of optimizing the measurement set-up to get rid of parasitic nonlinearities and to lower the noise floor. Measurements at other DC bias points, which may perform better, and at different RF and optical frequencies, to demonstrate the bandwidth of the device are also planned.

3.1.5 Digital Performance

For completeness the digital performance of the modulator was also measured on a 2000μm device. Open eyes were observed out until 12.5 Gb/s, the maximum output speed of the pseudo-random-bit-source (PRBS).

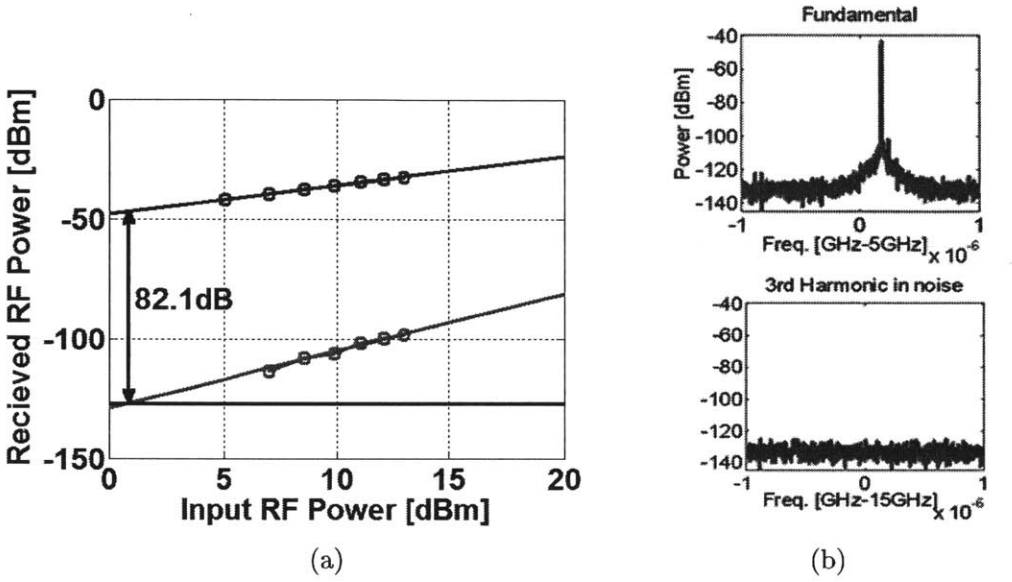

(a) (b)

Figure 3-8: SFDR measurement of the device. Showing (a) the SFDR intercept with the noise floor at -127dBm in our measurement setup. The SFDR was measured to be 82.1dBHz2/3. (b) The fundamental at -44.9dBm and the third harmonic in the noise floor. From [94]

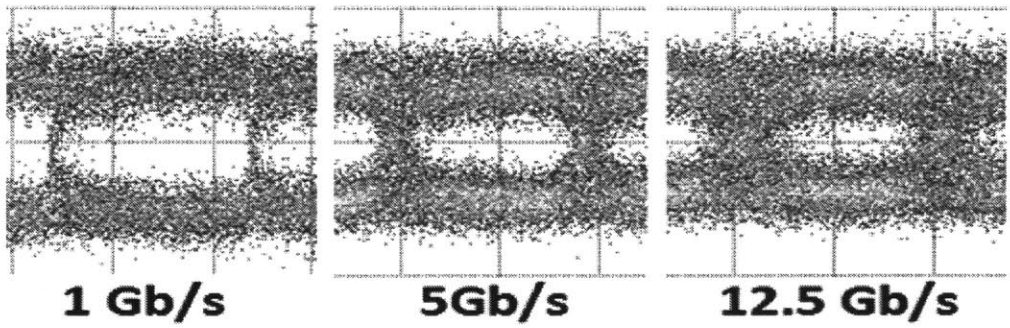

Figure 3-9: Open eye diagrams at 1 Gb/s, 5 Gb/s, and 12.5 Gb/s.

3.1.6 Summary and Future work

A broadband, linear silicon MZM is demonstrated based on our theory proposed in [98]. Our modulated device had an extremely good $V\pi L$ of 0.27Vcm, and a measured input IP3 of 68.4dBm, 20dBm better than the best in the literature. A link SFDR measurement was also performed and an SFDR of 82.1dBHz$^{2/3}$ was measured. Photodiode nonlinearity and a high noise floor limited this measurement. Finally, we carried out digital measurements where we observed open eyes out to 12.5 Gb/s, the limit of our measurement system.

Short term future work consists of getting a cleaner, more complete measurement set. Longer term work focuses on exploring options to improve on chip linearity and also digital modulator performace. One way to do the former is to try to integrate some of the more successful tabletop schemes. One way to do the later is to look at other underlying electrical devices such as MOScaps, MOSFETs, JFETs, ad BJTs. These devices may or may not help device linearity. Looking at alternatives to the plasma-diserpsion effect is also an option. In a slightly different vein, this work could also be expanded to different material systems and different optical wavelengths.

3.2 Integrated Single Sideband Modulator

Single Sideband Modulators (SSBs) are important for many analog applications such as laser noise suppression loops, low noise microwave oscillators, and arbitrary waveform generators. This is especially true of suppressed carrier SSBs. Furthermore, SSBs are useful in digital applications such as dense wavelength division multiplexing (DWDM) in long-haul fiber systems since the lower amount of necessary optical power allows for decreased nonlinear effects. In such digital systems, the modulated signal can be decoded through mixing with a local oscillator. [108]. SSBs are particularly interesting for inclusion on chip because they provide a way to create a frequency shift, and acousto-optic modulators (AOMs), the more traditional, table top way to do so, are difficult to integrate. Several different style of SSBs have been used in years past. Including filtering out one sideband from amplitude modulation, AOMs, and

nested Mach-Zehnder (MZs) interferometers [108, 109, 110].

For our oscillator system, described in the introduction, an integrated SSB is necessary. We desired this modulator to be relatively optically broadband and to tolerate small thermal fluctuations well. Because amplitude modulation followed by sideband filtering requires that a ring resonant filter be lined up to the sideband, something that requires retuning for each modulation frequency and optical wavelength, and may also need to be re-tuned with temperature fluctuations and because the filter needs to be very narrow for small frequency shifts to separate out the sidebands, this method is not practical here. Likewise, AOMs are impractical due to the need to integrate exotic materials onto the chip. Thus, we focused our attention on nested MZMs as the obvious integration choice. This thesis presents the first, to the author's knowledge, integrated silicon photonic nested MZM suppressed-carrier single-sideband modulator made in a CMOS-compatible silicon photonics platform.

3.2.1 Theory and Prior art

Before the design, fabrication, and experimental characterization of this modulator is presented, it is useful to take the time to understand how this modulation scheme works. A picture of the modulator topology is shown in Fig. 3-10(a) and was first proposed using LiNbO3 modulators in [111], and was elaborated on in [108, 112, 113]. As can be seen the SSB consists of two nested Mach-Zehnder modulators (MZMs) within a larger MZ super-structure. Each nested MZM is biased at the zero output point (that is there is a π phase shift in DC between the two arms). This choice of bias point prevents the optical carrier from being transmitted through the device — thus the suppressed carrier. The two arms of each nested MZ modulators are operated in push-pull; allowing for the largest achievable phase offset for a given voltage swing, but it is important to note that the use of push-pull, while it gives better results, is not a necessary part of the device architecture: a single arm phase-shifter operated with twice the voltage in the voltage swing would work equivalently well from a topological standpoint. The inner modulation yields a pure amplitude modulation on each arm. A $\pi/2$ phase shift between the input RF tones going to each nested MZM causes one

to be modulated with a sine and the other a cosine wave. In the frequency regime, the two sidebands from the cosine modulation point in the same direction, while those of the sine modulation are imaginary and point in opposite directions. A further DC $\pi/2$ shift in the optical phase of the light between the two outer arms puts them in the same plane and, thus, allows one sideband to cancel out and the other to add; thereby yielding the single-sideband operation. We note that an extra phase shift on each of the inner arms of one of the modulators is equivalent here, as illustrated in Fig. 3-10(b). For illustration, the frequency domain output at each step is shown in figure 3-10(c). Which RF signal is delayed in relation to which arm has the extra $\pi/2$ from the heater (and whether it is $\pi/2$ or $3\pi/2$ determines whether it is the upper sideband or the lower sideband that is passed through to the output.

To better understand the operation of this device (and it's benefits, drawbacks, and failure modes), we work through the math. As discussed in the previous chapter, we will work with the analytic representation of the signal. Suppose that a pure frequency of laser light is input into the modulator at point A in Fig. 3-11. This signal is given by:

$$A = E_0 e^{j\omega_0 t} \tag{3.10}$$

Where E_0 is the complex amplitude of the input field, and ω_0 is the frequency of the light. Since A is the large input of an adiabatic 3dB coupler, after the first adiabatic 3dB coupler the signal on the upper and lower arms is given by:

$$B = C = \frac{E_0 e^{j\omega_0 t}}{\sqrt{2}} \tag{3.11}$$

Where we have assumed that the adiabatic coupler is ideal with a 3dB splitting ratio and no loss. We note that since adiabatic couplers, instead of directional couplers, are used there is no $\pi/2$ phase shift between the arms of the couplers. We further note that the presence or absence of this term is irrelevant as it can be corrected out with the heaters. After going through the large input in the second set of adiabatic

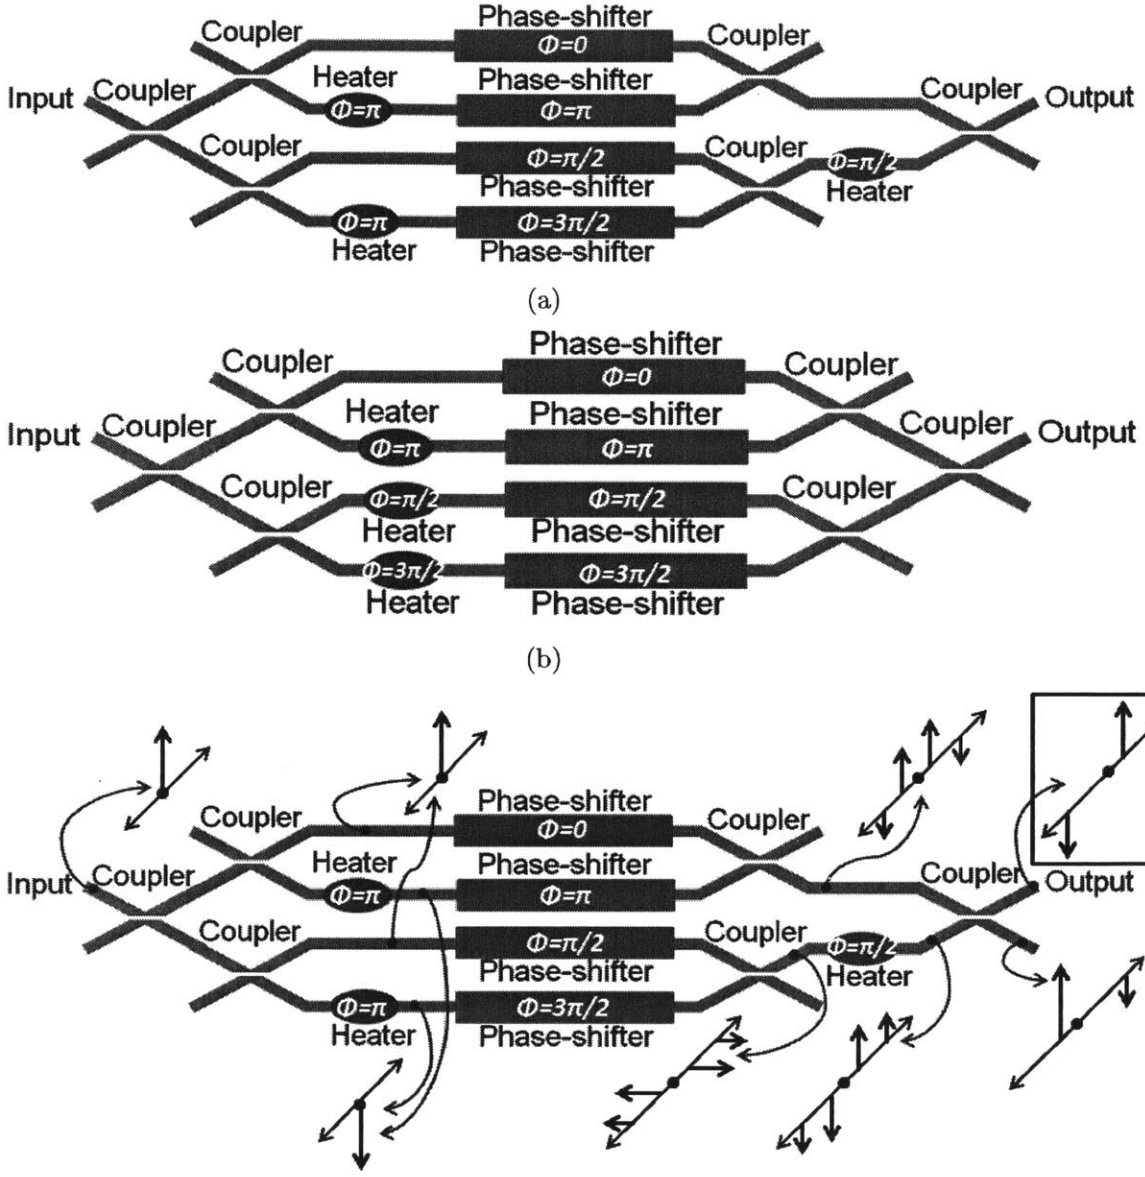

Figure 3-10: (a) Schematic of Single Sideband Modulator (b) Equivalent Layout with all heaters in inner arms, (c) Frequency domain representation of the signal at each point along the modulation path.

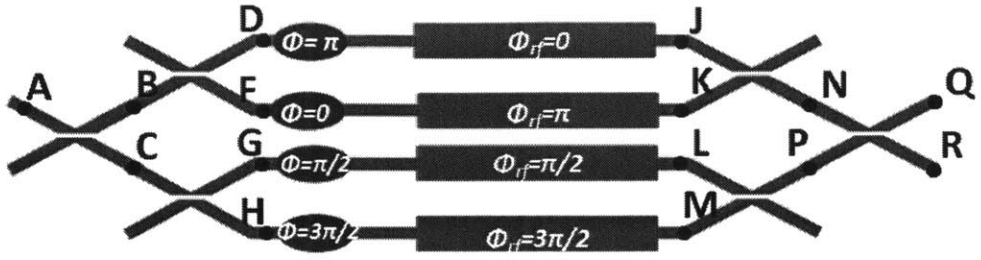

Figure 3-11: Labeled diagram of SSB modulator

couplers the fields are given by:

$$D = F = G = H = \frac{E_0 e^{j\omega_0 t}}{2} \tag{3.12}$$

The light then hits the heaters and the phase-shifters. For the moment we will assume that the fabrication is perfect so that the phase-shifters and heaters act exactly the same. After passing through the heaters and phase-shifters we have:

$$J = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+n(v))L - j\phi_{h1}}}{2} = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t))L - j\phi_{h1}}}{2}$$

$$K = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+n(v))L - j\phi_{h0}}}{2} = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t+\theta_\pi))L - j\phi_{h0}}}{2}$$

$$L = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+n(v))L - j\phi_{h2}}}{2} = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t+\theta_{\pi/2}))L - j\phi_{h2}}}{2}$$

$$M = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+n(v))L - j\phi_{h3}}}{2} = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t+\theta_{3\pi/2}))L - j\phi_{h3}}}{2} \tag{3.13}$$

At this point the signals are recombined at the output couplers to determine the final signal. After the first set of output adabatic couplers we have:

$$N = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t))L - j\phi_{h1}}}{2\sqrt{2}} + \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t+\theta_\pi))L - j\phi_{h0}}}{2\sqrt{2}}$$

$$P = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t+\theta_{\pi/2}))L - j\phi_{h2}}}{2\sqrt{2}} + \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n \sin(\omega_{RF} t+\theta_{3\pi/2}))L - j\phi_{h3}}}{2\sqrt{2}} \tag{3.14}$$

And after the final adabatic coupler the top output is given by:

$$Q = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t))L - j\phi_{h1}}}{4} + \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t+\theta_\pi))L - j\phi_{h0}}}{4} + \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t+\theta_{\pi/2}))L - j\phi_{h2}}}{4} + \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t+\theta_{3\pi/2}))L - j\phi_{h3}}}{4} \tag{3.15}$$

and the bottom output is given by:

$$R = \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t))L - j\phi_{h1}}}{4} + \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t+\theta_\pi))L - j\phi_{h0}}}{4} - \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t+\theta_{\pi/2}))L - j\phi_{h2}}}{4} - \frac{E_0 e^{j\omega_0 t - jk_0(n_0+\Delta n\sin(\omega_{RF}t+\theta_{3\pi/2}))L - j\phi_{h3}}}{4} \tag{3.16}$$

. In order to be able to intuitively understand the equations above, it is useful to pursue two different paths. One is to create a device model in MATLAB® or VerilogA and run simulations to explore what happens as different parameters change. This will be explored more later. The second, more straight forward one it to simplify the above into a set of analytic equations that are easier to understand. Here we follow this latter path. To begin, it is necessary to expand out the sine term. The traditional way to do this for just a sine in the exponent is the Jacobi-Anger expansion (a Bessel function expansion). Because of the addition of the phase terms in the RF and heater offsets, we use Chowning's version of the expansion instead, which is given by [114]:

$$Ae^{j\omega t + j\phi + j\beta\sin(\Omega t+\theta)} = Ae^{(j\omega t + j\phi)} \sum_{m=-\infty}^{\infty} J_m(\beta) e^{jm\Omega t + jm\theta} \tag{3.17}$$

Applying this yields:

$$Q = \frac{E_0 e^{j\omega_0 t - jk_0 n_0 L}}{4}\Big[e^{-j\phi_{h1}} \sum_{m=-\infty}^{\infty} J_m(k_0\Delta nL) e^{-jm\omega_{RF}t} + e^{-j\phi_{h0}} \sum_{m=-\infty}^{\infty} J_m(k_0\Delta nL) e^{-jm\omega_{RF}t - jm\theta_\pi} + e^{-j\phi_{h2}} \sum_{m=-\infty}^{\infty} J_m(k_0\Delta nL) e^{-jm\omega_{RF}t - jm\theta_{\pi/2}} + e^{-j\phi_{h3}} \sum_{m=-\infty}^{\infty} J_m(k_0\Delta nL) e^{-jm\omega_{RF}t - jm\theta_{3\pi/2}}\Big] \tag{3.18}$$

Expanding out the sums and collecting like terms yields:

$$
\begin{aligned}
Q = \frac{E_0 e^{j\omega_0 t - jk_0 n_0 L}}{4} \Big[& J_0(k_0 \Delta n L)\big[e^{-j\phi_{h1}} + e^{-j\phi_{h0}} + e^{-j\phi_{h2}} + e^{-j\phi_{h3}}\big] \\
& + J_1(k_0 \Delta n L)\big[- e^{-j\phi_{h1} + j\omega_{RF}t} + e^{-j\phi_{h1} - j\omega_{RF}t} - e^{-j\phi_{h0} + j\omega_{RF}t + j\theta_\pi} \\
& + e^{-j\phi_{h0} - j\omega_{RF}t - j\theta_\pi} - e^{-j\phi_{h2} + j\omega_{RF}t + j\theta_{\pi/2}} + e^{-j\phi_{h2} - j\omega_{RF}t - j\theta_{\pi/2}} \\
& - e^{-j\phi_{h3} + j\omega_{RF}t + j\theta_{3\pi/2}} + e^{-j\phi_{h3} - j\omega_{RF}t - j\theta_{3\pi/2}}\big] \\
& + J_2(k_0 \Delta n L)\big[e^{-j\phi_{h1} + j2\omega_{RF}t} + e^{-j\phi_{h1} - j2\omega_{RF}t} + e^{-j\phi_{h0} + j2\omega_{RF}t + j2\theta_\pi} \\
& + e^{-j\phi_{h0} - j2\omega_{RF}t - j2\theta_\pi} + e^{-j\phi_{h2} + j2\omega_{RF}t + j2\theta_{\pi/2}} + e^{-j\phi_{h2} - j2\omega_{RF}t - j2\theta_{\pi/2}} \\
& + e^{-j\phi_{h3} + j2\omega_{RF}t + j2\theta_{3\pi/2}} + e^{-j\phi_{h3} - j2\omega_{RF}t - j2\theta_{3\pi/2}}\big] \\
& + J_3(k_0 \Delta n L)\big[- e^{-j\phi_{h1} + j3\omega_{RF}t} + e^{-j\phi_{h1} - j3\omega_{RF}t} - e^{-j\phi_{h0} + j3\omega_{RF}t + j3\theta_\pi} \\
& + e^{-j\phi_{h0} - j3\omega_{RF}t - j3\theta_\pi} - e^{-j\phi_{h2} + j3\omega_{RF}t + j3\theta_{\pi/2}} + e^{-j\phi_{h2} - j3\omega_{RF}t - j3\theta_{\pi/2}} \\
& - e^{-j\phi_{h3} + j3\omega_{RF}t + j3\theta_{3\pi/2}} + e^{-j\phi_{h3} - j3\omega_{RF}t - j3\theta_{3\pi/2}}\big] + ...\Big]
\end{aligned}
\tag{3.19}
$$

where we have taken advantage of the fact that $J_{-n}(B) = (-1)^n J_n(B)$. If we make no more simplifying assumptions this where we get stuck, and we have calculate the Bessel coefficient and phase of each frequency term to figure out what it is. However, to gain intuition for the problem, we can consider the ideal case. In this case we assume that the RF and heater biases are perfect and unchanging (that is that: $\phi_{h0} = 0$, $\phi_{h1} = \pi$, $\phi_{h2} = \pi/2$, $\phi_{h3} = 3\pi/2$, $\theta_\pi = \pi$, $\theta_{\pi/2} = \pi/2$ and $\theta_{3\pi/2} = 3\pi/2$. Substituting these values in and combining like terms gives:

$$
\begin{aligned}
Q = E_0 e^{j\omega_0 t - jk_0 n_0 L}\big[& - J_1(k_0 \Delta n L) e^{-j\omega_{RF}t} + J_3(k_0 \Delta n L) e^{j3\omega_{RF}t} \\
& - J_5(k_0 \Delta n L) e^{-j5\omega_{RF}t} + J_7(k_0 \Delta n L) e^{j7\omega_{RF}t} + ...\big]
\end{aligned}
\tag{3.20}
$$

A similar analysis shows that the complimentary output gets:

$$
\begin{aligned}
R = E_0 e^{j\omega_0 t - jk_0 n_0 L}\big[& J_1(k_0 \Delta n L) e^{j\omega_{RF}t} - J_3(k_0 \Delta n L) e^{-j3\omega_{RF}t} \\
& + J_5(k_0 \Delta n L) e^{j5\omega_{RF}t} - J_7(k_0 \Delta n L) e^{-j7\omega_{RF}t} + ...\big]
\end{aligned}
\tag{3.21}
$$

The above is the equation commonly cited in the literature and first worked out in [108]. Looking at these equations a few things are quickly obvious. One of the more important is that not all the power in the carrier can be transferred into the desired sideband. At full modulation (m=1) 7.1dB of the input power is lost. Even at maximum power transfer (m=1.84) 4.7dB of power is lost. That there is some loss makes intuitive sense: the two output ports each get one (set of) sideband(s), so at least half the power must be lost to the other port. Additionally, some of the power is lost out of the rejected port in each internal MZM. While at first it might seem that there might be another 3dB here, this is not the case. One way to see this is to consider the average power transmitted. When the modulation is very weak, around the bottom of the sinusoidal transfer curve, most of the light is lost. This increases until the full 0 to 1 of the transfer curve is traversed. Here, symmetry dictates that exactly 1/2 the light goes though yielding a 3dB loss (plus another 1.1dB loss from light sent into higher harmonics - higher order Bessel terms - at this depth for a total of 7.1dB lost). However, as the modulator is driven harder the parts of the curve past the "1" that get swept out still transfer more than half the input light, continuing to pull more light through than the half averaged when the swing was 0->1 and therefore pulling the average up. This is true even more so as the sinosoidal swing spends more time on the parts of the curve near the extremes of its range. This continues until the parts of the curve swept out on the other side of the peak dip below the new average by enough to start pulling the amount of light passed back down. Thus "overdriving" allows for more power transfer and at the maximum transfer point (m=1.84) only an extra m = 1.7dB is lost out of this rejected port. As m continues to increase loss again increases and undesired high harmonics continue to climb. This progression is illustrated using both our VerilogA code from Chapter 2 and a MATLAB system code in Fig. 3-12.

Another thing that becomes clear from the above equations is that the two ports give the complimentary outputs – it's not so much that the light is all frequency shifted as that the modulation bands are picked off and separated spatially. In this way it is similar to an amplitude modulator with a filter, except with the benefit that

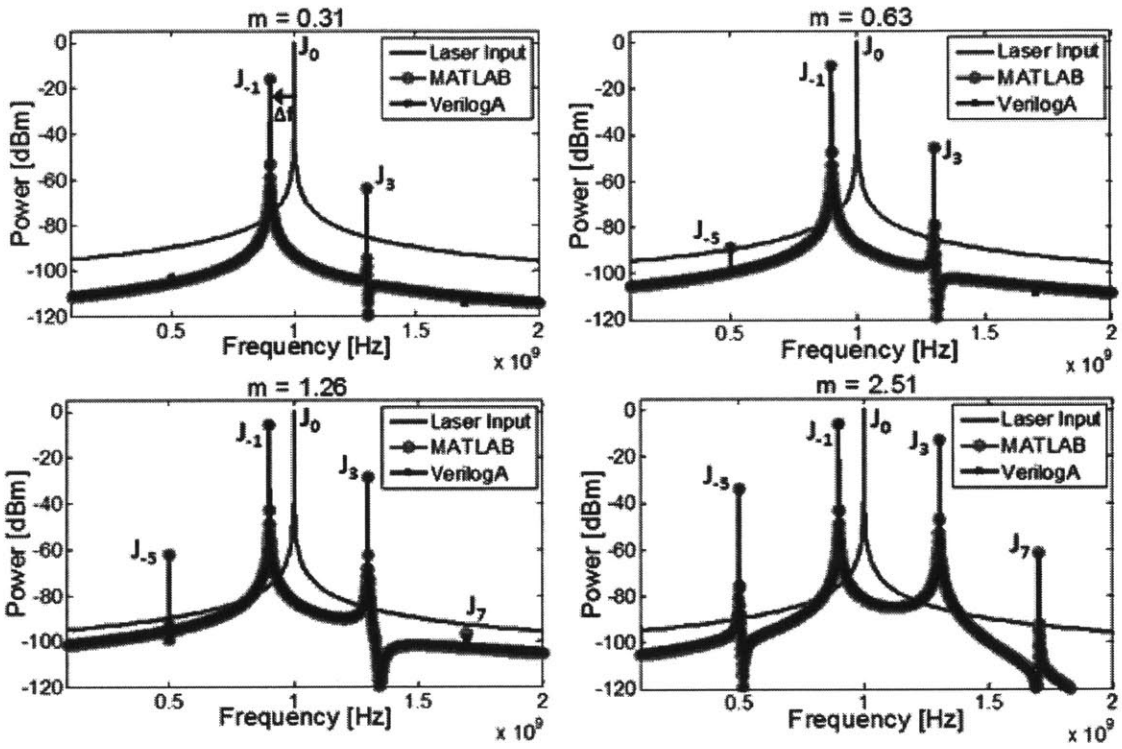

Figure 3-12: As RF drive voltage increases power transfer to desired sideband increases, but undesired harmonic terms increase faster, making choice of operating voltage a trade off in loss vs. nonlinearity.

the filter does not need to be re-tuned as the RF modulation frequency or optical input wavelength changes. Moreover, as can be seen from the Bessel function expansion, the modulation scheme is not perfectly linear and, as the power transferred to the desired output goes up, so to do does the power in these unwanted harmonics. This is fairly standard, as harmonics are generally caused by nonlinear processes and therefore rise at a faster rate than fundamentals with input power. [96].

One final thing that can be seen in these equations is that there are a lot of active knobs to turn. Since tuning a lot of knobs at the same time is difficult, it is useful to consider what happens if these various knobs are tuned imperfectly (and how one might go about figuring out which needs to be changed in the lab). Uneven splitting in the 3dB coupler can also create problems. First, looking at the J_0 term in Eq. 3.19 we can see that the carrier suppression comes entirely from setting the heater biases in the MZ modulators to ensure destructive interference at the output ports. This carrier suppression is unaffected by errors in the RF signals, which don't appear in any form in this term of the equation. In order to fully suppress the carrier, it is only a requirement that all four heater terms cancel out, however in order to ensure proper single-sideband operation (as will be seen looking at later terms), it is actually important that each nested MZ be biased at the zero output point. Thus, if an outer heater is used instead of three inner heaters (Fig. 3-10.a. instead of Fig. 3-10.b.), the bias on that heater doesn't matter for carrier suppression, either. It should be noted that in a practical lab setting, the silicon diode phase shifters in the SSB will also create a DC phase shift based on their average voltage bias point, and that fabrication differences may create path differences between the arms, meaning that appropriate heater bias to achieve a π difference in output phase between each set of inner arms will need to be determined experimentally. In addition to the heater biases, carrier suppression also relies on splitting and loss difference in the couplers and phase-shifters respectively. This is not explicitly shown in 3.19, but can be seen intuitively by the fact that different amplitudes can't cancel each other completely. Therefore, if the coupler split isn't exact, there will be a limit on how much the carrier can be suppressed. However, since the appropriate bias point corresponds

to exact destructive interference within each nested MZM, it is only the splitting ratio imperfections within the two nested MZMs that matter for carrier suppression. Differences in splitting into the two super-structure arms (the outer set) will result in imperfect unwanted sideband suppression but will not affect carrier suppression or power transfer to the desired sideband. Since the splitting affects carrier suppression, it is clear that unequal losses in the arms will also affect carrier suppression. For a standard DC difference in loss (due to say a defect in one of the waveguides) one can consider it equivalent to a difference in splitting ratio and the same results apply. For a different DC phase shift in the two arms due to different average bias points, this can dynamically affect carrier supression and is discussed more below. In all, the DC phase and amplitude differences between the two arms inside the two nested MZMs are drivers behind carrier suppression, with the differences and imperfections in the RF signal and between the two arms in the super structure playing basically no role.

We next consider effects on unwanted sideband suppression and harmonic suppression. The latter is easier to treat. Simply put, if the modulation and coupling was ideal and linear no higher harmonics would be generated. Here they come from the sinusoidal nature of the transfer function of the MZMs, and so are affected by anything that causes the less linear parts of the sinusoidal curve to come into play more than necessary. That is they get worse when the modulator is driven harder and the and the more nonlinear parts of the transefer function are seen [96]. The even order harmonics are canceled by being in push-pull exactly. Errors in DC bias point, differences in loss between the two arms, splitting ratio differences, and lack of exact matching between the modulators can cause these terms to appear (see the previous section on linearity for more detail). In general though, for an SSB modulator, it is the height of the largest undesired spur that is of interest, and as these harmonics are, for the integrated devices presented here, generally less tall than the spur from the unsupressed carrier or unwanted first order sideband, they are less important.

It is more important to consider detrimental effects on the unwanted other first order sideband. We can see from 3.19 that bias point will have a detrimental effect on this, as will differences in RF phase offset (away from $\pi/2$) between the RF signals,

incorrect phase difference in output between the two MZMs as set by the heaters (that is, different than $\pi/2$), and splitting ratios that differ from 50:50. Note that differences in RF signal amplitude will not affect the amount of power in the unwanted sideband (though RF loss will of course mean that there is less power available to be transferred to the desired sideband than would otherwise be present). Thus, except for the difference in m caused by the lower RF power, differences in RF amplitude will not matter – they don't affect any necessary cancelation to create a *higher* number. Fig.3-13.a. and b. provides a brief overview of some of the failure modes of the SSB that might be observed in the lab. One interesting thing to note is that the outside heater bias can be used to compensate for errors in RF phase difference and vice-versa. This is also illustrated in Fig.3-13.a.

Finally, we take a moment to discuss linearity. Much attention was paid in the last section to the creation of a linear MZM where the third harmonic was canceled. In that section the second harmonic was ignored because it could be filtered if it was out of band, canceled with differential detection, or removed by changing the DC bias point of the device. None of these three options are possible here: the band is usually wide enough to include the second harmonic, differential detection is obviously not compatible with the application, and the heater biases are being used to suppress the carrier. Therefore the second harmonic, which if one would recall is caused by un-equal losses in the arms, will create an unwanted tone that can limit performance. However, another source of unequal losses in the arms — unequal coupling through the adiabatic couplers — often dominates over second harmonics caused by the phase-shifters or ideal MZ transfer function, to limit carrier suppression and create unwanted even order tones in the output signal. We suspect that this is what creates the J_2 signal we see in the outputs below.

3.2.2 Design, Layout, and Results

To create an on-chip CMOS compatible version of the SSB modulator, the same reversed-biased diode silicon phase-shifters, resistive heating elements, and adiabatic couplers were used as were used to create the linear modulator in the section above.

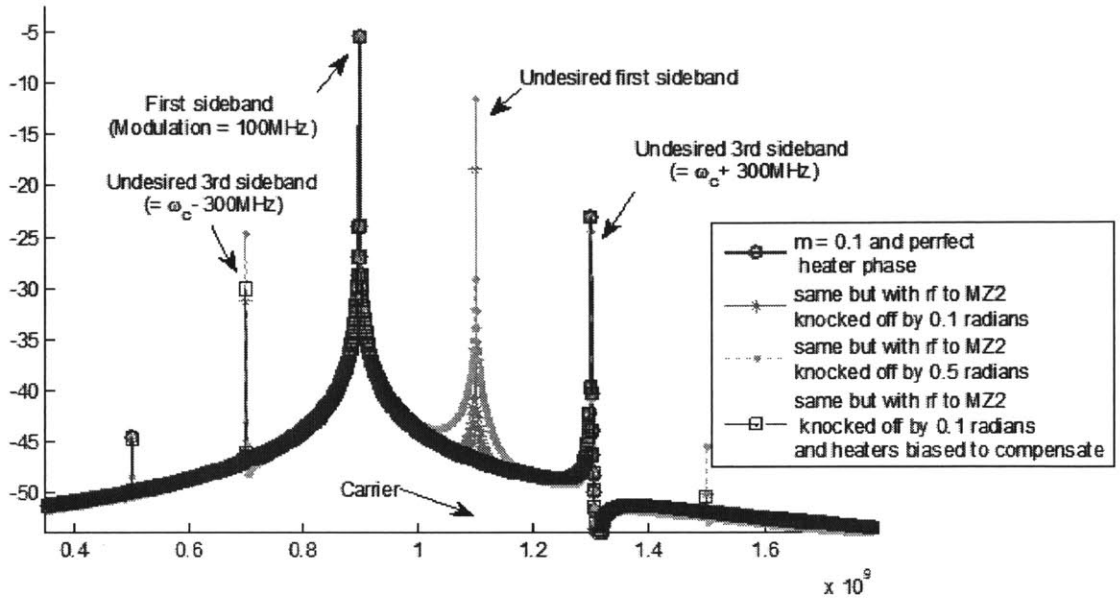

(a)

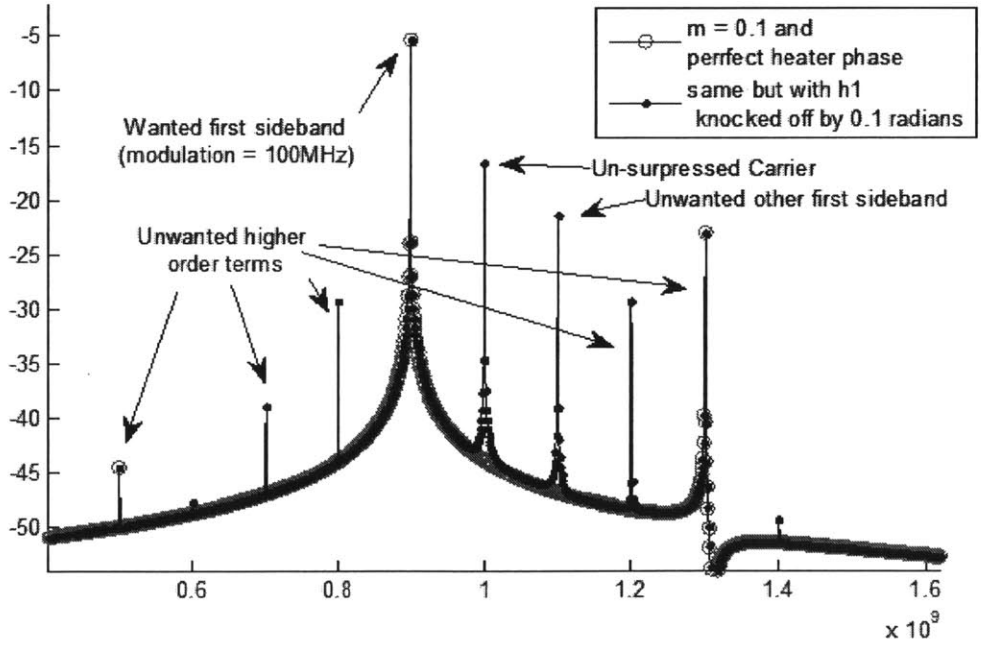

(b)

Figure 3-13: (a.) Effects of variation in RF signal phase on SSB output. (b.) Effect of incorrect bias on one heater inside one of the nexted MZMs. Note the appearance of the even-ordered harmonics and the carrier.

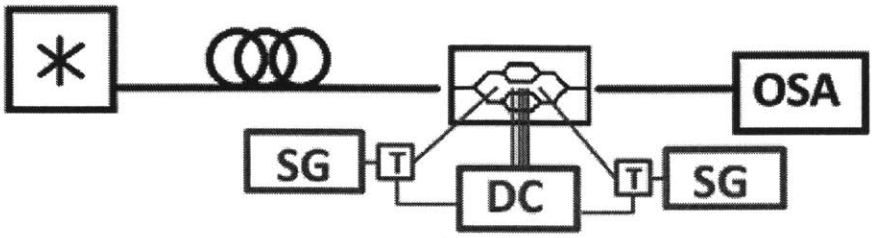

Figure 3-14: Measurement set-up. * = laser, OSA = optical spectrum analyzer, SG = signal generator, DC = DC power source. From [?].

The device was fabricated at CNSE Albany in the silicon photonics process described in Chapter 2. The experimental test set-up is shown in Fig 3-14. The output from a tunable laser is put through a polarization controller and a lensed fiber and coupled onto the chip. Inverse tapers are used on chip to capture the light into and out of the waveguides that connect to the device. Another lensed fiber is used at the output of the chip to couple the output light and input it to an optical spectrum analyzer (OSA). The electrical signals are brought onto the chip through three GSGSG probes. Three of the six signals are the heater bias signals and are provided by a Keithley power source. A forth is the VDD for the top of the modulators. The other two signals are the RF+DC signals that go to the nested MZMz in each arm. The RF signals are provided by two signal generators that are locked together. Each signal generator output goes through a bias-T where it is combined with a DC signal to set the center bias between the diodes in each arm.

A device with 2mm long phase shifters was tested. The 3dB roll-off of the device was measured to be 10GHz as shown in Fig.3-6 in the previous section. A shorter device would have a higher frequency roll-off and also a more linear performance, but would require higher drive voltages to achieve the same modulation depth and power transfer to the desired sideband. We therefore choose to work with the longer device. Single sideband operation at 1550nm was achieved and the results are shown in Fig.3-15. More than 18dB suppression of the carrier and more than 15dB suppression of the spurious sideband and tones was achieved. The presence of the unwanted sideband is likely due to a difference in RF signal amplitude due to different paths from the two

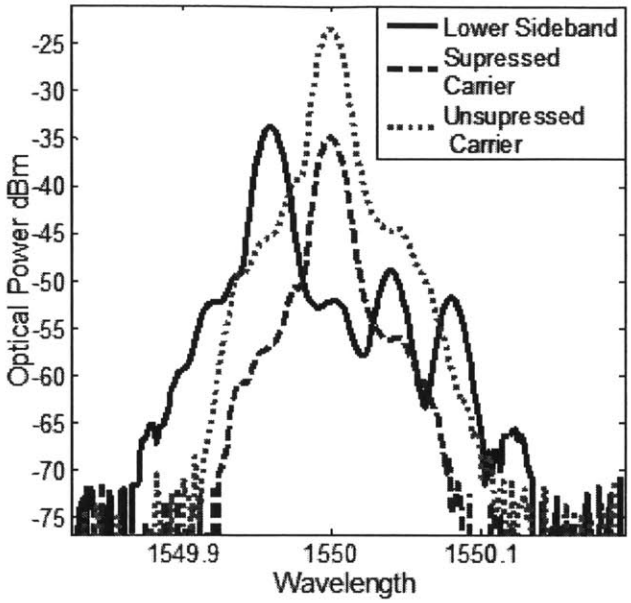

Figure 3-15: Output of single sideband modulator, showing carrier suppression greater than 18dB and suppression of other sideband and tones more than 15dB. The expected power loss is -9.3dB and the measured loss is -10.4dB. From [?].

signal generators as well as to error in bias point. This latter is also likely responsible for the observed J2 term. At the measured value of m = 0.74 we would expect about -9.3dB of power to be lost between the sideband and the unsuppressed carrier. This matches well with the measured value of -10.4dB. The slight discrepancy is again likely due to errors in bias which degrade power transfer in addition to shifting power into unwanted extra tones. Finally, at this point we would expect that the J3 term is suppressed below the sideband by 32.6dB. This also matches well with the measured value of 31.9dB.

The device was also tested at different optical wavelengths and RF frequencies. Fig. 3-16. shows the shift in the lower sideband location with modulation frequency. Fig.3-17. shows device performance at 1570nm and 1590nm. As can be seen the device works out to 1590+. The carrier is suppressed by >10dB in both the 1570nm and 1590n case and the unwanted sideband and tones are suppressed by >15dB at 1570nm and >25dB at 1590nm.

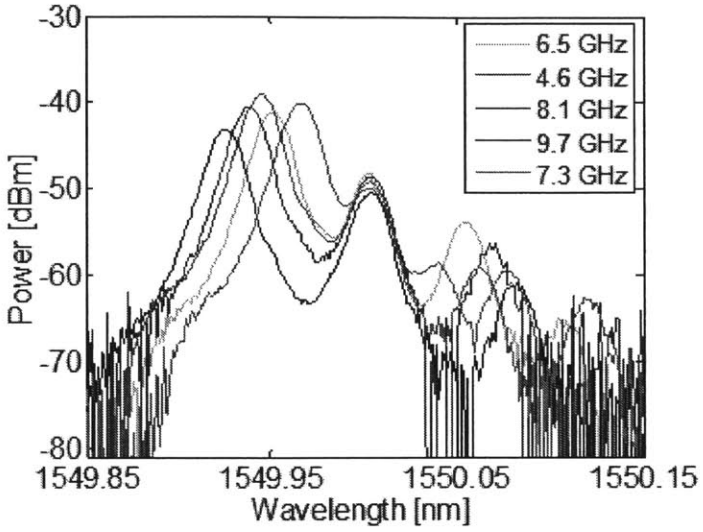

Figure 3-16: Shift in lower sideband location as a function of modulation frequency. From [?].

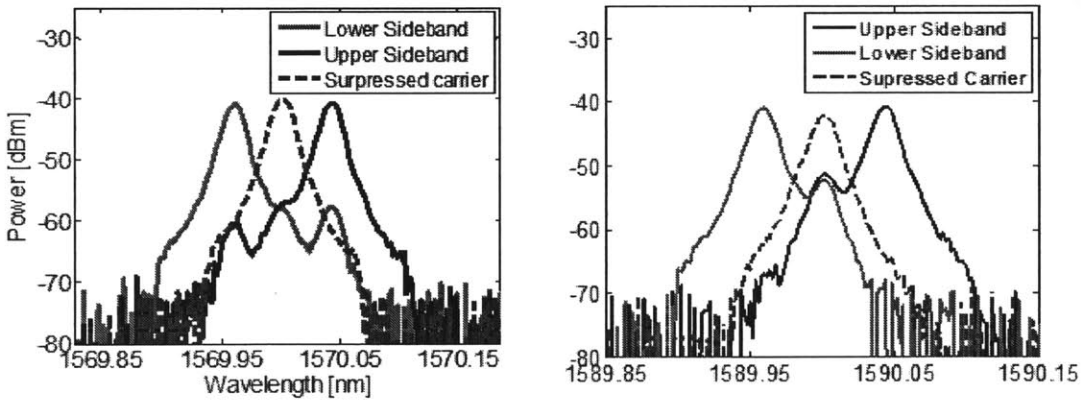

(a) (b)

Figure 3-17: Single-sideband operation in same device at 1570nm and 1590nm. Carrier and other sideband suppression >15dB can be seen at 1570nm, and carrier suppression >10dB and other sideband suppression >25dB can be seen at 1590nm. From [?].

Total insertion loss (including on chip couplers) though the device is 20dB. insertion loss through device only is estimated to be 7dB (for the unsuppressed carrier), through comparison with a straight waveguide containing no device.

3.2.3 Summary and Future Work

Above we have demonstrated, to the best of the authors' knowledge, the first CMOS compatible integrated silicon photonic single sideband modulator. The device works to 10GHz and from 1550-1590nm+ and has more than 10dB suppression of the carrier and unwanted tones. Future work on this device includes designing to counteract failure modes by, for example, including integrated optical attenuators in the arms to correct for loss or coupling differences and, thus, further suppress the carrier and unwanted sideband. Applying the linearization techniques from the previous session may also help suppress higher harmonics. Modifying the device to work over a different wavelength range is also a possibility as is using the, now functional device in a variety of systems. Such systems may include analog systems such as herterodyne interferometers [115] and on-chip lidar systems [116], as well as our oscillator loop. This topology can also be used to transmit a variety of digital modulation formats [113].

Chapter 4

Integrated Mode-locked Laser in the Silicon Photonics+CMOS platform

4.1 Introduction

Mode-locked lasers (MLLs) are important for many applications. In particular, the low timing jitter between pulses makes them attractive for low jitter analog systems such as low-noise microwave generation, optical sampling and for frequency combs. Creating an integrated version of an MLL is attractive as it would allow for integrated versions of these systems that have small size and weight, but still have the precision made possible by low jitter MLLs. In addition, integrating an MLL promises to significantly reduce the sensitivity of these lasers to environmental variations, such as vibrations, as well as to make them more compact and easier to package. Further, an integrated laser design would allow for mass-produciblity, and, via economies of scale, cheaper laser systems, and, since all the components are fixed in place, would not require expert maintenance or realignment. Finally, an integrated MLL could have a shorter cavity than is practically feasible in free space or fiber laser systems, and thus, have a higher repetition rate.

Previous work on integrated MLLs has been done in the group by Michelle Sander and Hyunil Byun [81, 82, 83, 84] on an integrated MLL using Er doped aluminosilicate glass on a silicon chip. This work achieved 440 fs pulses at a 394MHz repetition

rate with 30 pJ of intracavity pulse energy, and 285 fs pulses at 500MHz with 10 pJ pulse energy [81, 82, 83, 84]. However, this material system cannot be directly integrated with active silicon photonics or CMOS for further data processing, a major downside given our goals. Integrated MLLs have also been shown in III-V materials on silicon [85]. While this allows for bonding to active photonic and CMOS circuitry, the noise properties of III-Vs ultimately limit the performance of the MLL to levels below that achievable in solid-state MLLs, and to levels that are unacceptable for our oscillator application. There have been a number of other demonstrations of integrated waveguide lasers in the literature, but these have generally relied on free space components to facilitate mode-locking (see for example [86, 87, 88, 89, 90]). Regardless, we wish to demonstrate a MLL that is truly compatible with integrated silicon photonics and CMOS systems.

4.1.1 Desired MLL features

While there are many different applications for which an integrated MLL would be useful, we particularly would like to make a MLL for use in the low noise microwave oscillator system described in the introduction. In order to ensure that we can meet the specificed oscillator performance, we would like a mode-locked laser that has timing jitter less than 100fs and ultimately less than 1fs, takes up less than $1cm^2$, and RMS amplitude variation $<0.2\%$ and ultimately less than 0.08%. We would also like a high peak power of at least 50mW and preferably 200mW. We would also like a repetition rate high enough that we can filter out individual comb lines using ring resonant filters. Because of the required very low timing jitter, we choose to pursue passive, as opposed to active, mode-locking mechanisms as they have demonstrated the shortest pulses and due to the nature of photon number statistics, shorter pulses generally have lower timing jitter. We further note that for the case of passive soliton mode-locking the pulse width and the pulse energy are related via the parameters of the saturable absorber and of the cavity dispersion and nonlinearities, and that these parameters along with repetition rate will determine the peak and average power in the laser. The repetition rate of the laser is of course determined by cavity length

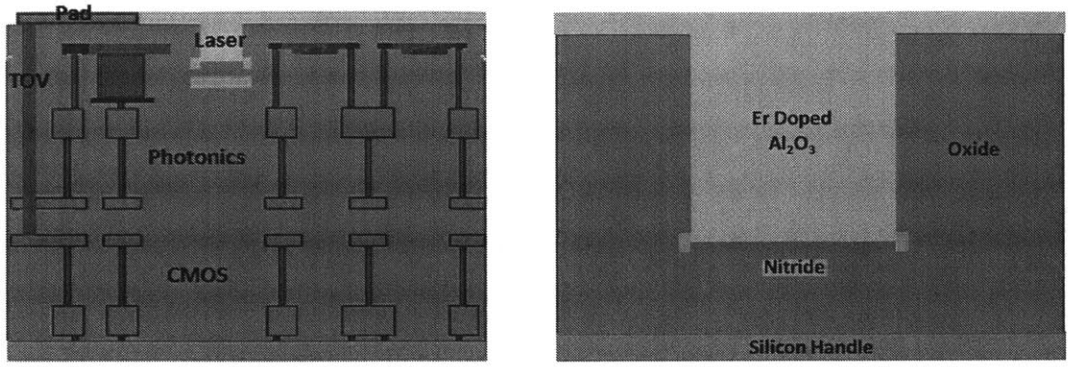

(a) (b)

Figure 4-1: (a) Full Photonic + CMOS layer stack after [47] and chapter 2, (b) Short-loop layer stack, which includes only the nitride and Er:Al_2O_3 layers.

which will affect device area. Thus, as can be seen, our target specifications are highly inter-related [80].

4.2 Platform Review and Basic Structures

The full photonic platform in which our lasers are fabricated was described in detail in Chapter 2, but a review will be provided here and an overview will be given of the basic laser structures. In particular, to save on costs during development, the MLL designs and test structures were fabricated in a series of laser short-loop runs where only the nitride, trench, and Erbium doped Al_2O_3 layers were fabricated. As a reminder, Fig.4-1.a. shows the diagram of the full layer stack from Chapter 2. Fig.4-1.b. shows a diagram of the cheaper, short-loop layer stack. As can be seen, the laser waveguide is comprised of a patterned nitride layer to guide the light and an etched trench in which an erbium-doped Al_2O_3 gain medium is deposited. A second nitride layer is used as an etch stop for the trench. Outside the trench etch, this layer is present to help guide the light. The nitride deposition and trench etching is done at CNSE Albany, while the Er:Al_2O_3 deposition is carried out at MIT.

The passive and active laser designs and lasing results will be discussed in more detail below, after we discuss the laser components. However, in order to facilitate component discussions, we will present a brief outline of two main passive architec-

tures here. A diagram of each of these architectures is shown in Fig. 4-2

Like all lasers, these consist of a gain medium inside a resonant cavity. Here the gain medium is the Er:Al_2O_3 present in the gain waveguide. The cavity is formed by a loop mirror, which serves as the output coupler, on one side and a saturable absorber that is butt-coupled to the chip and serves as a mirror on the other side. In order to achieve mode-locking, some kind of pulse shaping device is required in the cavity. The purpose of this device is to force all the longitudinal modes in the cavity to enter into a defined phase relationship such that when superposed they form a single pulse. In this case, the pulse shaping device is the saturable absorber. A saturable absorber has a nonlinear response to intensity — it absorbs the light hitting it up to a certain point, after which it transmits, or more commonly reflects, the remaining light. This serves to facilitate pulse narrowing as only the peak of the pulse survives. Thus, pulse shaping and mode-locking can be achieved. The gain waveguide, loop mirror and other cavity components are described in detail below. The saturable absorber is described in detail in [71].

4.3 Laser Components

Each laser cavity is made out of a series of integrated components. Because the platform is integrated, we cannot buy these components off the shelf to make the laser system, and we therefore need to design them ourselves. These designs will be summarized below. As mentioned above, the gain waveguide, loop mirror, and taper to the butt-coupled saturable absorber are important. Other components are also key. Bends are needed to allow for a long enough gain region to be realized in the compact chip area. Transitions are needed between components, particularly into and out of the gain region. A taper is needed to get the light off the chip once it is outside of the laser cavity. Structures are needed to get the pump light onto the chip and into the gain cavity. We designed lasers for both 980nm and 1480nm pump light. In the 1480nm case the pump can be manipulated along with the laser light and many components can be used for both. In the 980nm case a separate set of components

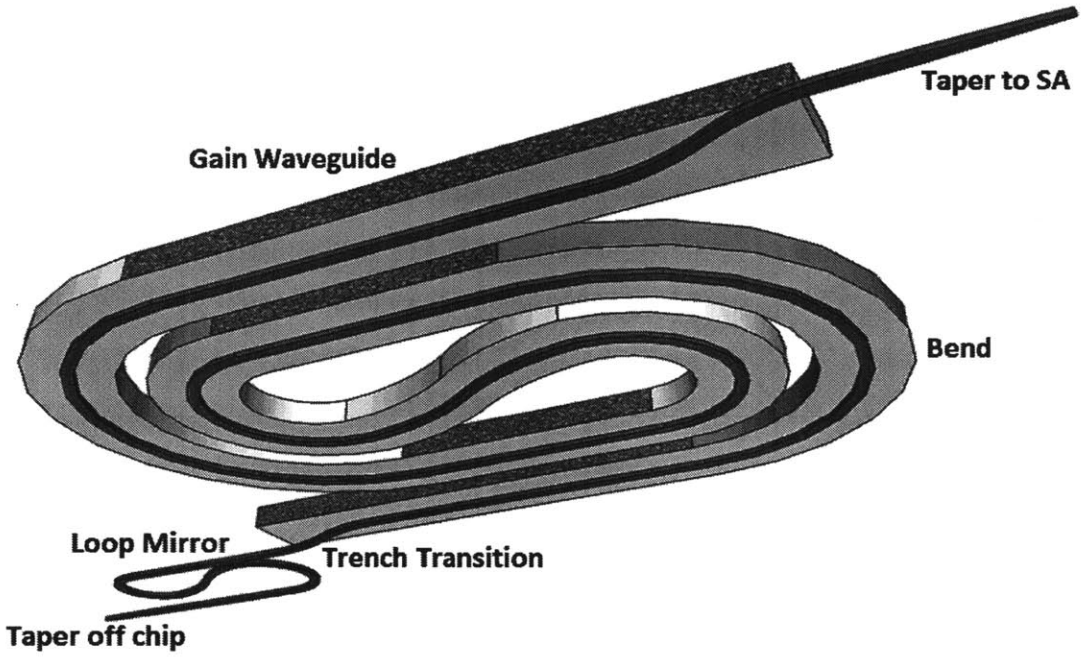

(a)

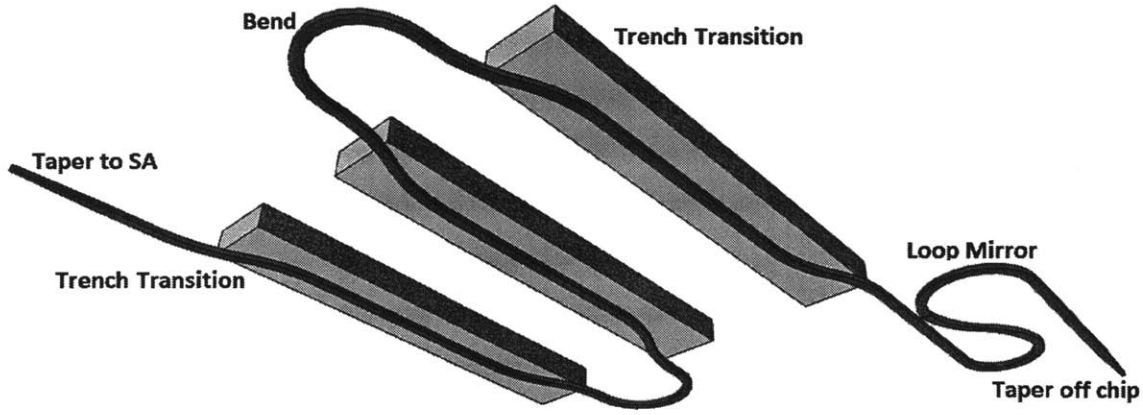

(b)

Figure 4-2: (a) Diagram of passive spiral laser architecture with main component parts labeled, (b) Diagram of passive compact s-bend laser with main component parts labeled.

that manipulates the pump light while leaving the laser light unchanged is necessary.

In all, the components used in the passive mode-locked laser structures are: an input taper on-chip, transitions into and out of the trench at both the pump and laser wavelength, a loop mirror, pump couplers and splitters, the gain waveguide, bends, and tapers to the off-chip saturable absorber. In addition the actively mode-locked structures need other components to do the active mode-locking. These include a silicon MZM, heaters, and phase-shifters, and a SiN-to-Silicon taper that moves the laser light from the initial nitride layers to the silicon layer in which the active devices are formed. We will go through these devices in more detail below.

4.3.1 Waveguide

While all components are ultimately important for laser operation, perhaps the most critical is the design of the gain waveguide itself. In our platform, the gain waveguide is a pseudo-rib type waveguide formed from a nitride layer which serves to guide the laser light, and an Er:Al_2O_3 filled trench, separated from the nitride by 200nm, which serves to provide the gain. Because the nitride has a higher refractive index (1.9-2) than the Al_2O_3 (1.55-1.65), both the laser and the pump light prefer to travel in the nitride guiding layer as opposed to the gain waveguide. To overcome this issue, the nitride waveguide is cut into a number of small pieces, in our case five, thus lowering the effective refractive index of the nitride block and forcing both the laser and the pump light up into the gain region. A cross-section of the gain waveguide and the resulting laser and pump modes is shown in Fig. 4-3.a-b. However, it is not enough just for the light to be in the gain material; there must also be good overlap between the laser mode and the pump mode in this region. Fig.4-3.c. shows the calculated overlap of the laser mode with the 980 and 1480nm pump modes as a function of Er:Al_2O_3 layer height. As can be seen, there is more than 70% overlap in all but one case and a >90% overlap can be achieved with a layer heights greater than 1.6μm with a 1480nm pump.

The width of the nitride slab is chosen to ensure that this pseudo-rib waveguide supports only a single transverse TE mode. This is important to ensure that different

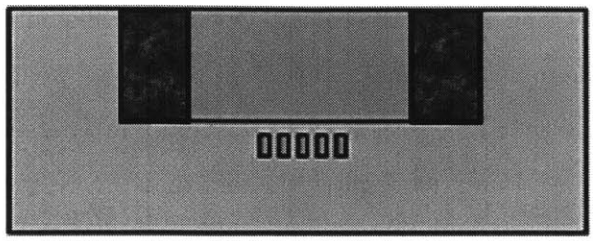

(a)

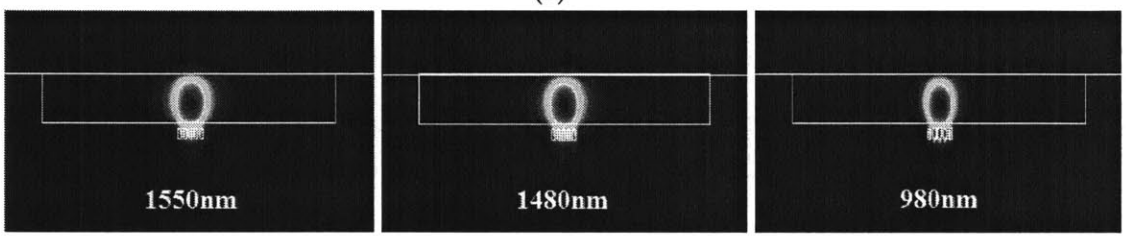

(b)

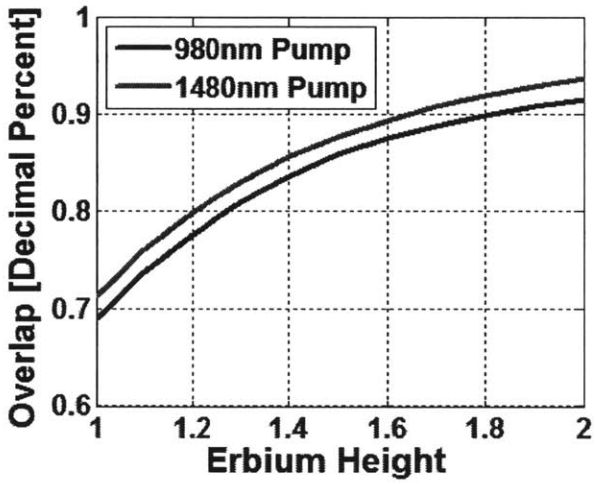

(c)

Figure 4-3: (a) cross section of gain waveguide where orange = nitride, blue = oxide, and green = Er:Al_2O_3. Marbled region at edges denotes roughness (see below), (b) Fundamental TE mode in waveguide for laser center and pump wavelengths, (c) Overlap of the laser mode with the pump are of the gain medium.

transverse modes do not compete for available cavity gain and, thus, destroy the defined phase relationship between the longitudinal cavity modes necessary for mode-locking. A single TM mode is also be supported by this waveguide, but this mode will be attenuated by other elements in the cavity. However, while only a single TE mode is guided by the nitride multiple modes are still supported by the trench (analogous to how the slab will support extra modes in a rib waveguide). These modes will prohibit mode-locking. To remove them, we exploit the fact that the desired mode is confined to the center of the trench by the nitride rib, while the undesired modes, which are guided by the trench itself, are exposed to the trench sidewall. We roughen the sidewall by creating a large, random change in trench width, increasing the scattering loss for modes that have significant field distribution at the edges while leaving the desired TE mode, confined to the trench center, untouched. A top view and cross sections of such a waveguide along with the predicted modal loss from the roughness are shown in Fig.4-4. We choose a roughness of 7μm, which as can be seen from the graph contributes basically no loss to the TE fundamental mode (>0.001dB/cm in theory) while contributing loss greater than the waveguide gain to the higher order modes.

To experimentally verify that this technique works we measured light output through a series of waveguides with different roughness widths. The results can be seen in Fig.4-5. The Al_2O_3 deposited in the trench in these samples was not doped with erbium this allowing for passive, un-pumped measurements. The blue trace, which corresponds to a waveguide with no roughness clearly shows the periodic nulls in transmission (fringes) caused by the destructive interference from multimode beating. The $5\mu m$ and 7 μm roughness widths show that his beating has been removed – the green and red trances have a flat transmission profile – without adding significant propagation loss (there is no decrease in amplitude as the roughness width is increased). As the roughness width is widened further to $10\mu m$, the black line, the induced scattering loss starts to affect the fundamental, desired mode as well, resulting in higher loss as can be seen in the figure. Thus, this technique of roughening the trench walls works as expected. As a further check that the roughness will not affect

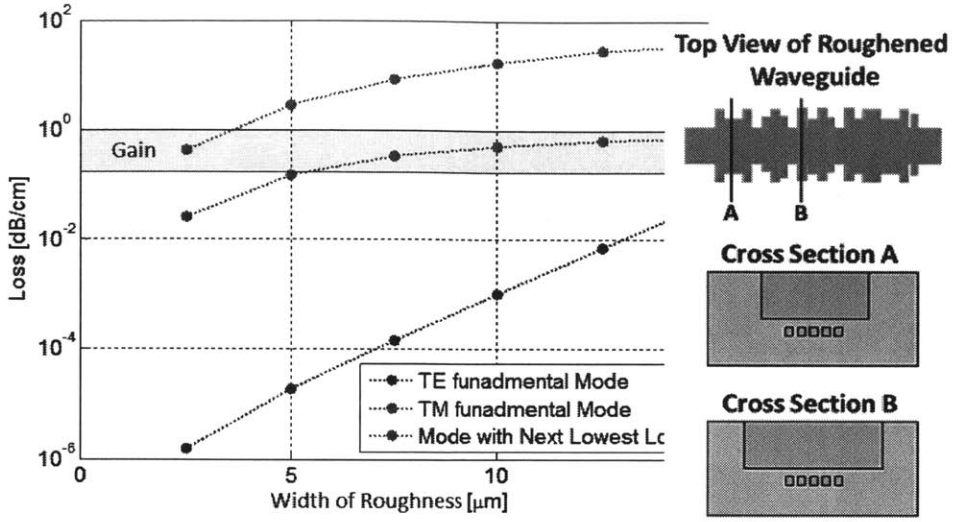

(a) (b)

Figure 4-4: (a) Loss vs. roughness width (width of the center region is equal to 50 microns minus two times the roughness width). Grey region represents expected waveguide gain. As can be seen loss in the Fundamental TE mode remains low while loss in the other modes is high, (b) Top view of trench width in roughened waveguide along with cross sections at two locations in the roughness.

the waveguide gain, transmission as a function of pump power was measured for a roughened and un-roughened waveguide, and was seen to look identical.

We next consider the gain in this waveguide. The gain was measured by comparing the amount of light through a 10cm long and 20cm spiral of gain waveguide as a function of input pump power. The net waveguide gain at 1550nm was measured to be around 0.5 dB/cm. This was lower than the original prediction (of 1dB/cm) and appears to be a function of material growth. However, it is still high enough to achieve laser oscillation in the longer structures as will be discussed below.

Because we wish to use this waveguide as part of a mode-locked laser, the waveguide dispersion is also important. Mode-locking requires that the pulse shape remain unchanged after one round trip of the laser cavity, this is often accomplished by balancing the effects of anomalous dispersion against those of self-phase modulation. In order for this to work, the overall laser cavity has to have anomalous dispersion. This turns out to be difficult to achieve.

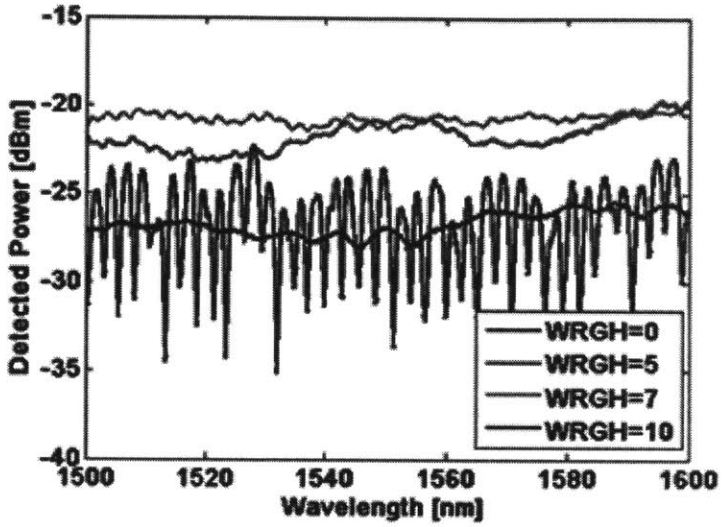

Figure 4-5: Experimental transmission results for roughened waveguides: multimode beating is clearly seen for the case of no roughness (blue line), a roughness width of 5 or 7 microns (red and green lines) removes unwanted modes (as seen by lack of beating), but does not add to loss, while a roughness width of 10 microns (black line) causes extra loss.

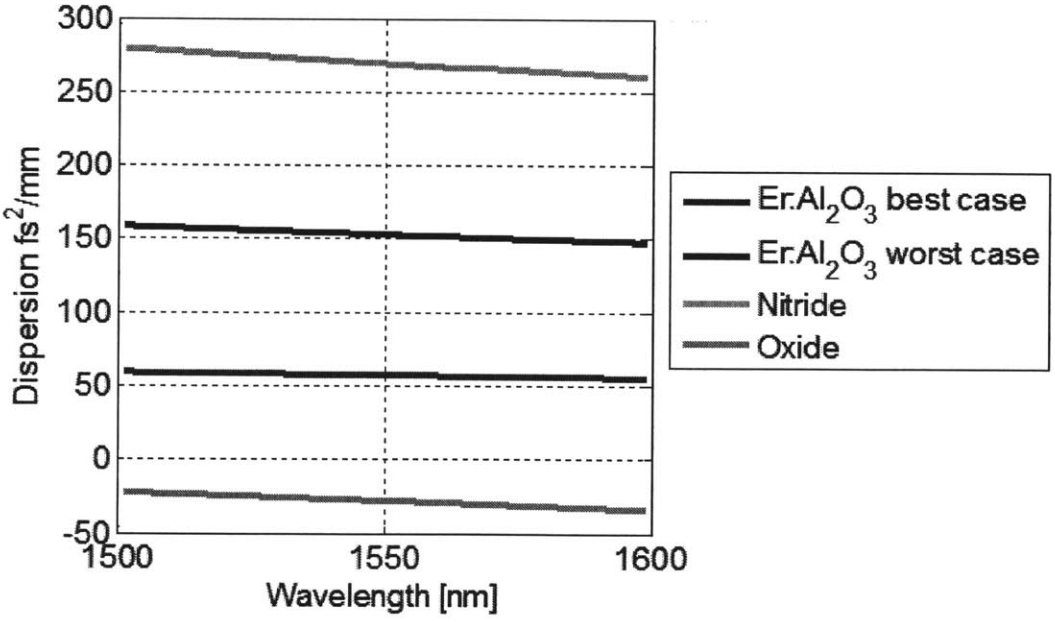

Figure 4-6: Material Dispersions

The material dispersion contributed by the nitride, oxide, and Al_2O_3 are shown in Fig.4-6. The two different lines for Al_2O_3 represent dispersion measurements from different deposition runs. The "best case" line represents the lowest normal dispersion and was the only data we had when the original designs were carried out. One thing that is immediately apparent from Fig.4-6 is that the contribution from material dispersion to the total dispersion will be normal; the majority of the mode exists in either the nitride or the Al_2O_3 both of which have normal dispersion. While the oxide does have anomalous dispersion, it is only weakly anomalous and only a small fraction of the mode sees the oxide. Increasing the fraction of the mode in the oxide would only help slightly and would come at the expense of the gain. Moreover, increasing the fraction of the mode in the oxide is difficult in our process as fabrication considerations dictate that the Al_2O_3 filled trench must remain wide. Fortunately, the material dispersion is not the only contribution to the overall dispersion, the waveguide dispersion also contributes.

In the literature, there are two main ways to achieve anomalous waveguide dispersion: slots and very loose ribs [72, 73, 74, 75]. Adding slots to our nitride does not help with the dispersion as not enough of the mode is there, and adding slots to the alumina is not feasible in our fabrication platform (at least not adding *vertical* slots — the direction needed for dispersion engineering in the TE mode), we therefore go for the thick rib option. After all, we do basically have a thick rib structure to start with. We aim to keep the dispersion either anomalous or normal but close enough to zero that we can compensate for it with dispersion added by the saturable absorber. (An external piece of fiber could also be used, if there is enough gain, but this stops the laser from being fully on-chip). Total dispersion as a function of Al_2O_3 layer height for the best and worst case is shown in Fig 4-7. As can be seen for the best material case and the thicker heights, anomalous dispersion can be achieved. For the worst material case and the thinner heights overall dispersion is normal, but small. Dispersion considerations will be discussed further below.

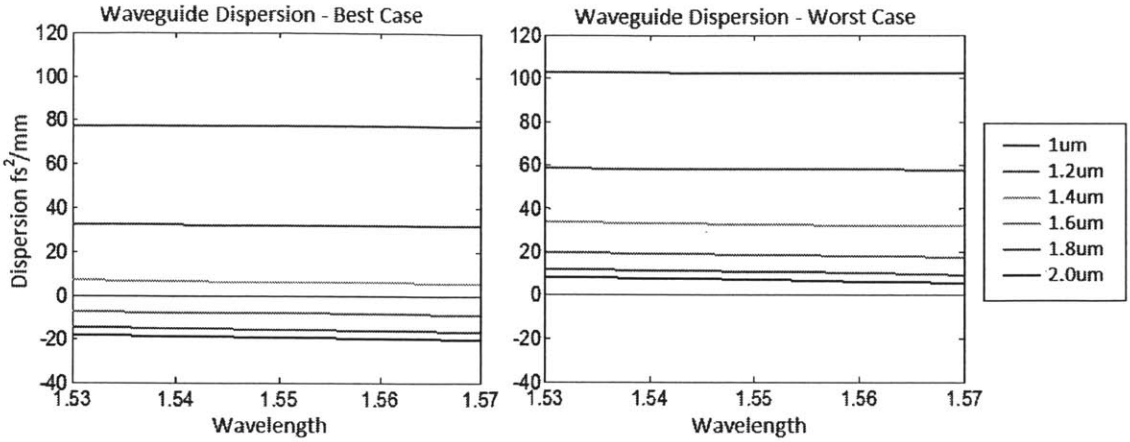

Figure 4-7: Total dispersion for the best (left) and worst (right) Al2O3 material dispersion. Anomalous (negative) dispersion is achieved in the best case for thick layer heights. For thinner layers or the worst case Al2O3 material dispersion total dispersion is normal but, in most cases, small.

4.3.2 Bends

Because the reticle is only 2.6cm wide by 3.2cm long, if we want a longer gain region, we will need to bend the gain waveguide. In our case, a longer gain region is necessary to accumulate enough total gain to overcome total cavity loss. The extra gain also allows for more pulse energy, necessary to hit our pulse energy targets, and easier mode-locking. Higher pulse energy also leads to lower timing jitter, also desirable for our oscillator application. Additionally, a longer laser cavity yields a lower repetition rate, which also makes mode-locking simpler, particularly at low powers. However, when a waveguide is bent, the eigenmodes of the waveguide change. This causes two issues. The first is that, at abrupt changes in curvature, the incident field will excite both fundamental and, if present, higher order local modes as well as radiation modes in the connected waveguide [76]. This means that when transitioning from a straight to bent waveguide section, or vice versa, we have to make sure that we have good overlap between the fundamental modes in which we desire to propagate, otherwise there will be significant loss at the transition. The second issue is that the bend itself can be lossy. In the bend, the planes of constant phase exist along radial lines and, thus, the phase velocity increases linearity with distance from the center of the arc.

At some point the phase velocity will exceed the speed of light in the oxide and at that point that part of the field will radiate away [76]. We need to make sure our radii of curvature are large enough that very little of the field is radiated through this effect. However, in order to conserve mask space and keep the laser as compact as possible, we also want to keep the bend as tight as possible.

We adopt two different strategies in our two different laser architectures shown in Fig.4-2. In the larger, spiral lasers, we aim to keep the cavity as simple as possible. We therefore simply bend the gain waveguide we have at the minimum radii that allows for negligible radiation loss and for >99% overlap between the straight and bent waveguide modes for both the bandwidth around 1550nm at which the laser travels and at the 1480nm pump bandwidth. While there is also >99% overlap at the 980nm pump wavelength, the bend is multimode at this pump wavelength so we do not count on it working and instead plan to pump at 1480nm. However, we are also constrained in bend radius by the size of the chip. It turns out that this minimum bending radius is around 3mm —large but easily accommodated — for Er:Al_2O_3 film heights in the 1-1.3mm range, but expands to be huge, >2.5cm, the width of the reticle, for the 2mm film height case. Fig.4-8 shows minimum bending radius as a function of Er:Al_2O_3 film height for our gain waveguide as well as diagrams of the bent modes. We therefore plan to work at thinner film heights. This has consequences for dispersion as discussed above and below.

Obviously we would like a way to create a more compact bend, especially one that will work with a thicker Er:Al_2O_3 film height. One option is to couple the light out of the trench, bend the light in the nitride layer, and then coupled back into the gain region using another erbium trench transition. This allows for a tight bend as the minimum bending radius in the nitride is on the order of 20 μm, as seen in Fig. 4-9. However, transitioning into and out of the trench add loss to the system both because distance traveled in the nitride during the transition and in the bend is distance not traveled in the gain medium and loss from the mode manipulation in the transition itself. This later loss, discussed in more detail below, was measured to be about 1dB per bend, and, given that only about 1.5-2.5dB of gain are added with

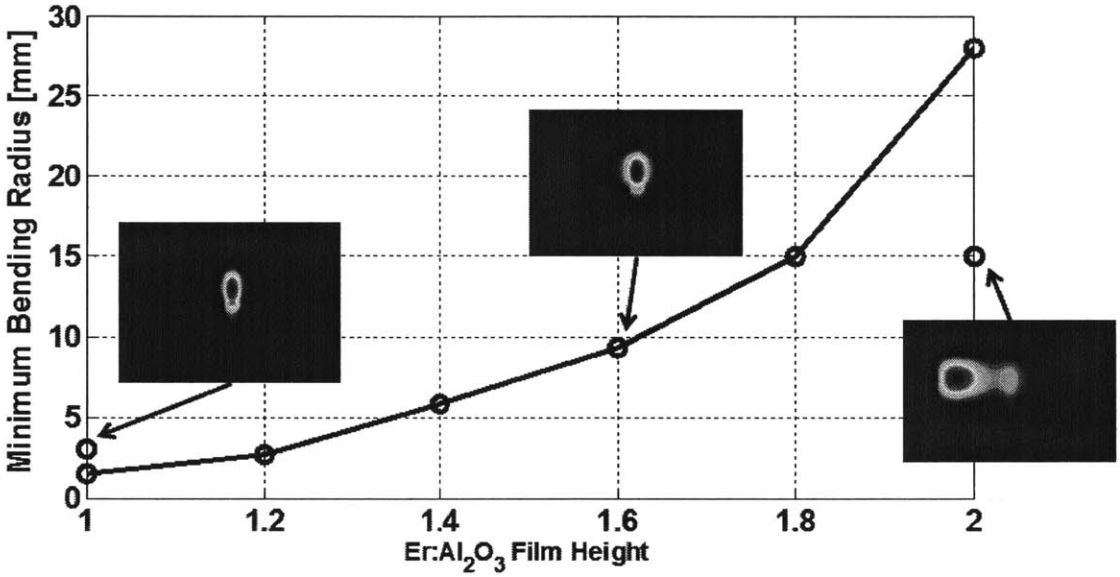

Figure 4-8: Simulated bending loss vs. radius (in mm) for gain waveguide bend. Inset figures show modes at different film heights relative to bending radii (center of bend is on right). From left to right: 1μm film height and 3mm bending radius: mode is well confined to center of the guide, 1.6μm and 9.3mm radius: one can see that the mode is starting to slide toward the edge of the trench, 2.0μm film height and 15mm radius: mode is no longer confined by nitride guide, but rather by side of the Al_2O_3 trench.

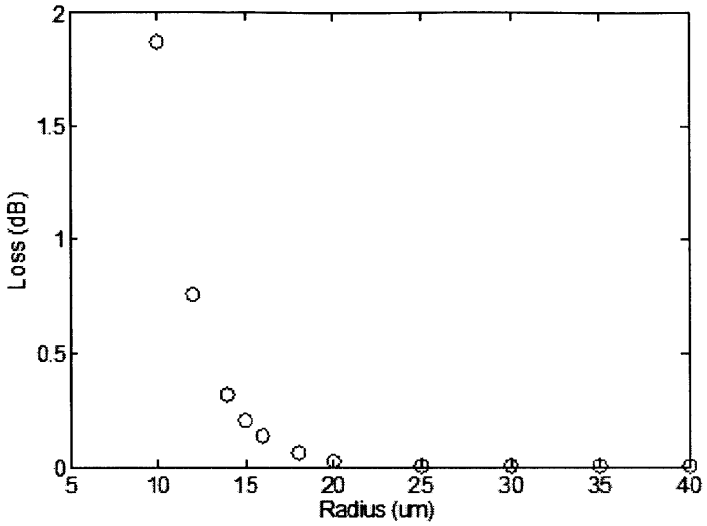

Figure 4-9: Simulated bending loss vs. radius for nitride waveguide bend

each gain section greatly diminishes the gain returns for bending the waveguide. The trench transition design itself is also very sensitive to fabrication issues so this loss may drift higher or lower in somewhat unpredictably. We, therefore, seek a solution that allows us to stay within the laser trench.

To design these bends, we first tackle the problem of radiation loss. We find that pulling the mode into the nitride by transitioning from a five piece to a one piece block and bringing the trench wall next to the nitride waveguide to create a hard outerwall that greatly reduces radiation loss and allows for bends as tight as 110μm in radius even for thick Er:Al_2O_3. In fact, simulations show that the Er:Al_2O_3 height has little effect on the performance of these bends, thus allowing us to work at thicker layer heights were anomalous dispersion is achievable. However, bringing the trench wall in repels the light toward the center of the waveguide in the straight section, while the angular momentum of the light in the bent section still presses it against the outer waveguide wall, thus destroying mode overlap in the transition. To fix this problem an offset was inserted into the waveguide to align the fundamental modes at 1550nm and 1480nm in the straight and the bent section. A diagram of this bend, taken from the mask is shown in Fig.4-10, and Table 4.1 shows the measurement results. The bends were found to work and to contribute 1dB of loss per complete bend. While

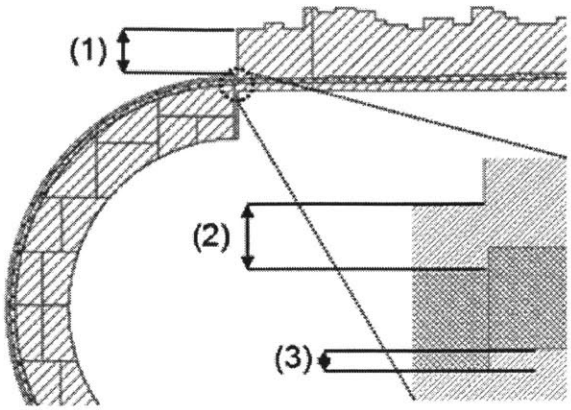

Figure 4-10: Diagram of bend, where dark green is nitride waveguide and olive green is the trench. bend parameters 1-3 are described in the table below.

Bend Type	Width of SiN [μm]	Bend Radius [μm]	Offset in Er [μm] (1)	Edge space in bend [μm] (2)	Offset in SiN [μm] (3)	Predicted loss per transition [dB]	Measured loss per transition [dB]
1	3	110	24.5	1.5	0.65	0.252	0.496
2	3.5	110	24	1.5	0.7	0.176	0.519
3	2.5	110	24	1.5	0.55	0.315	inconcl. (meas 0)

Table 4.1: Summary of measurements of compact Er:Al_2O_3 bend structures.

this is similar to the nitride bends above, the extra ability to pump around the bend offsets this, making the net loss smaller. Further optimization should also allow for future improvement in performance.

4.3.3 Loop Mirror

The output coupler is based on an integrated loop mirror like that presented in [77]. A diagram of this component is shown in Fig.4-11.a. The loop mirror is formed out of a directional coupler "looped" back on itself. Light is input into one port on one side, the two ports on the other side are connected together, and the second port on the first side serves as the output for the transmitted light. The reflected light is sent back out the input port. The amount of reflection is determined by the coupling

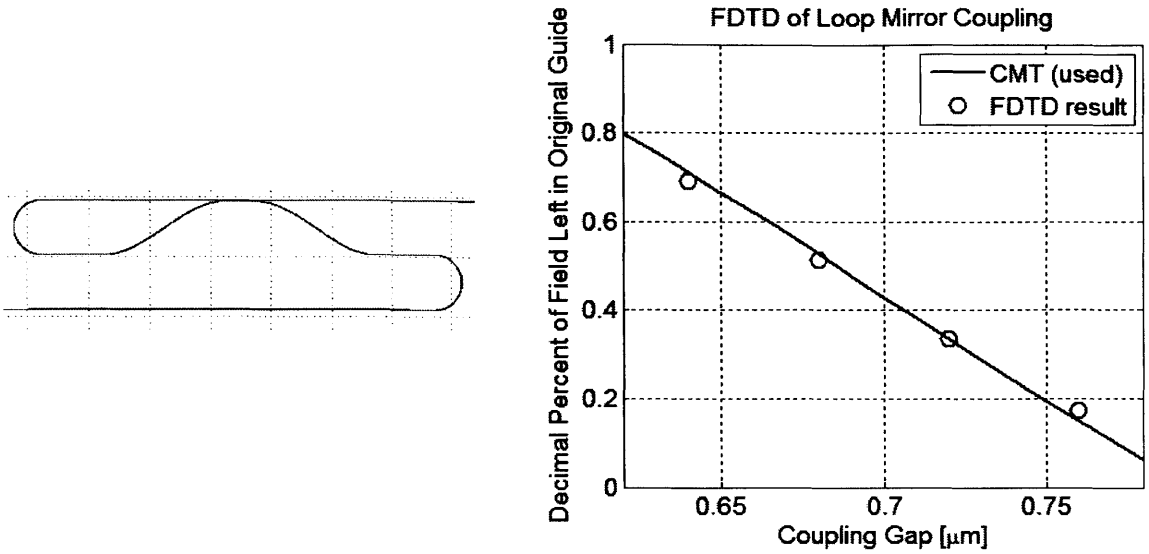

(a) (b)

Figure 4-11: (a). Diagram of loop mirror as drawn on the mask. (b) Comparison of coupled mode theory vs. FDTD simulations.

coefficient in the directional coupler, and is given by:

$$P_{reflected} = 4(\kappa - \kappa^2) \tag{4.1}$$

Where κ is the power coupling coefficient. The coupling coefficients for the loop mirror were calculated in two ways. First, a coupled mode theory script was written. This script used the 2D electromagnetic modes output from a modesolver to determine the coupling coefficient. This script is fast, but coupled mode theory is not always as accurate as necessary in tightly confined structures. Therefore a finite difference time domain (FDTD) code, a more accurate but far slower (especially given the 50-100μm lengths involved) simulation technique, was also run to determine the coupling coefficients and the results compared to the coupled mode theory (CMT) based code output. As can be seen in Fig.4-11.b., the results of the two codes match. Once it had been verified the CMT code was used to change and optimize the design of the loop mirror.

To verify the loop mirror design, a set of test structures were included on the mask.

These consisted of two loop mirrors connected to form a resonant cavity. The reflection of the loop mirrors could then be determined from the depth of the interference fringes. The resulting measurements are plotted against our theoretical predictions in Fig.4-12.a. As can be seen there is very good agreement between the predicted and measured results. Such good agreement indicated not only that our simulation code is working well, but also that the fabrication variations in the process are small enough not to have a detrimental effect on component performance. Because a mode-locked laser has a wide bandwidth, and this is a bandwidth sensitive component, we also want to make sure that our loop mirror structure will have a similar reflection across the whole bandwidth range of interest. To this end simulations were carried out for the 1530-1570nm range and the reflection was found to change by 10%. This is acceptable because higher losses in the rest of the cavity make this a small overall variation in wavelength vs. loss. Finally, in the 1480nm case, we would like to pump the laser waveguides through the loop mirror. In this case we would like the 1480nm to be transmitted while the 1550nm light is reflected. In this design reflection of the 1480nm light tended to be about 20% lower than that of the 1550nm light — workable but not great. Reflections vs. gaps for all these wavelengths are shown in Fig.4-12.b. In a more recent design of this component the gap and length were further optimized to create better separation between the 1550nm and 1480nm light and to further flatten the reflection curve in the 1530-1570nm band.

4.3.4 SA tapers

The loop mirror forms one end of the laser cavity, and the saturable absorber mirror (SA), in our case a saturable bragg reflector, forms the other end. Because the SA is butt-coupled to the chip, that is pressed against the side, eliminating losses coupling off and back on chip to and from it is important. Due to the high intensity of the confined waveguide mode, we need to expand the mode size at the chip edge to avoid burning the SA. Further, a larger mode size is helpful in limiting beam divergence in the SA or in any air gap that may exist between the chip and the SA, as larger beams diverge more slowly. This allows for a greater percentage of the light to be successful

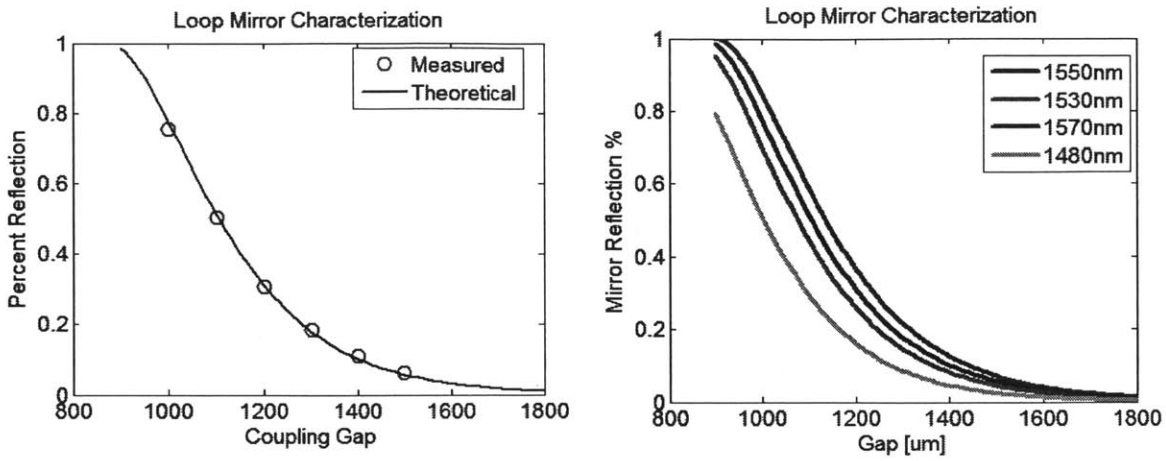

(a) (b)

Figure 4-12: (a) Measured loop mirror reflection is in good agreement with theoretical predictions. (b) Simulated reflection for various wavelengths. Reflection percentages are similar across the 1530-1570nm Erbium gain band. Reflection at the 1480nm is lower, as desired, but could be further optimized.

coupled back into the laser waveguide on-chip. To successfully expand the mode in both directions we first couple out of the laser trench. Coupling out is necessary as the Er:Al_2O_3-air boundary limits the extent to which the mode is able to expand upward and we do not have a good way to taper the film height or width to overcome this. Once in the nitride layers an inverse taper can be employed to expand the mode in both the vertical and horizontal directions. In this work we looked into two different taper designs: one where both nitride layers were tapered together and one where the two nitride layers were tapered separately (see Fig. 4-13). The first is simpler but does not allow for the mode to be expanded as much as the second. This is bad in terms of beam divergence, but good in terms of limiting losses dues to coupling to the substrate. However, waeguide roughness loss matters as well, especially as the waveguide gets narrow and the center of the mode is more exposed to the sidewall. It is this last effect that turned out to dominate in our case, making the first type of taper (shown in Fig.4-13.a.) the lower loss option.

The resulting loss in dB as a function of wavelength for coupling to the SA and back is shown in Fig.4-14. The loss was measured by pressing a silvered mirror against

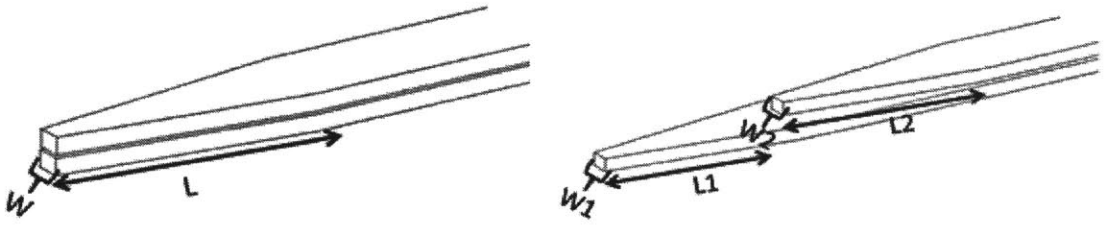

(a) (b)

Figure 4-13: (a) Taper type one where both nitride layers are tapered together. (b) Taper type two where each nitride layer is tapered separately.

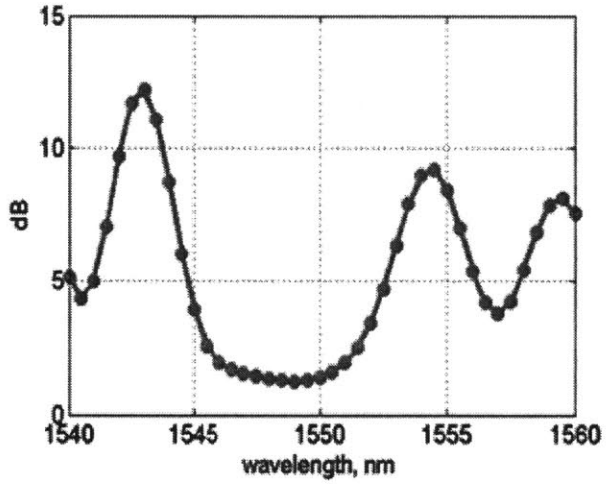

Figure 4-14: Loss in dB vs. wavelength for round-trip coupling to SA. From [47]

a tapered test structure; a straight nitride waveguide with the taper of interest on the end; the effects of the straight nitride section and input taper were calibrated out using a reference guide. As can be seen the loss remains low (>2dB, 1.42dB at 1550nm) for a 5nm bandwidth aroud 1550nm. This is acceptable for a relative long pulsed MLL though not for short pulses that will have wider spectral bandwidth. However, we suspect that the oscillations seen on either side are due to reflections in the measurement setup and test structures and do not represent a limit cause by the taper itself. Loss in these tapers will be further discussed in the next section.

4.3.5 Output tapers

Coupling to the off-chip saturable absorber is not the only place where we need to taper off chip: we also need a good way to get the laser pulses off chip once they have exited the loop mirror. Because this taper occurs outside of the laser cavity performance constraints are not as strict. We consider the same two basic structures shown in Fig.4-13. for this off chip taper. In this case the aim is to match the optical mode at the end of the taper to that of a 3μm spot-size fiber or that of a cleaved fiber (which will ultimately be able to handle more power). To this end, we aim to match the fiber mode diameter and shape at the end of the taper. Based on these numbers, a set of test structures consisting of a taper on each side and a straight nitride waveguide in the middle was placed on the mask and measured for each taper type. Fig 4-15 shows the raw throughput data from the best performing taper of each type. As elsewhere, taper type one consistently performed better than taper type two. However, because longer tapers also consistently performed better than shorter ones (as is also seen in Fig 4-16 and Table 4.2), it is suspected that this difference in performance is due to the longer lengths inherent in this design. Because theoretically longer tapers should perform better as they are more adiabatic, because material losses as measured on ring test structures are not large enough to account for this difference, and because changing the thickness of the oxide layer between the nitride and the silicon box did not account for the difference in either theory or measurement, it is suspected that this difference has something to do with sidewall roughness. As the taper progresses and the nitride edges are closer to the center of the mode sidewall roughness would contribute more loss, and longer tapers where the sidewalls are near the center of the mode for a greater distance would suffer more from this effect. Given its better performance, we choose to use taper type one in our designs.

Fig. 4-15 shows the loss coupling to a 3μm diameter lensed fiber as a function of taper width for taper type one. Material losses from the nitride have been removed and the data divide to give the loss per facet. As can be seen, off-chip coupling losses

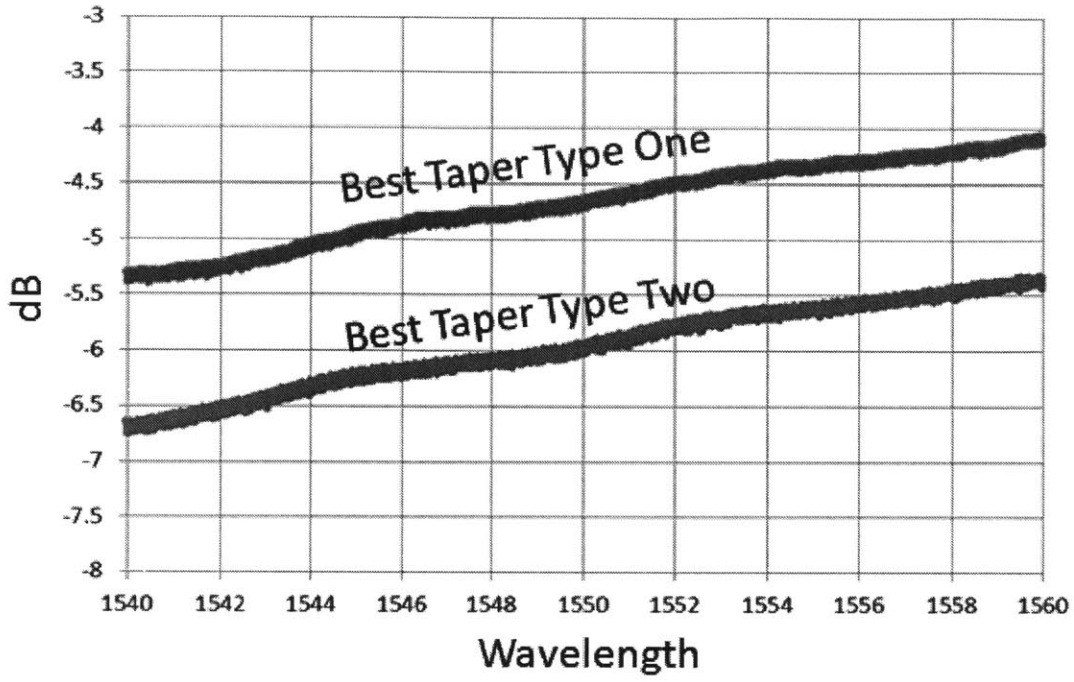

Figure 4-15: Raw data of best performing taper coupled to 3μm fiber for each taper type.

of less than 1.5dB can be achieved. The shape of the losses also matches reasonably well to the theoretical prediction, but the overall loss is about 1-1.5dB more than expected, again we assume this is due to roughness. For completeness, the data is summarized in Table 4.2.

In designs which use a 1480nm pumping wavelength, we wish to use the same taper as a way to bring pump light on chip. Fortunately, these tapers are broadband and losses at 1480nm are similarly small to those at 1550nm. (In fact, even measured losses at 980nm are in the 1.3dB-4 dB range). Thus, these tapers allow for efficient pump coupling on chip.

4.3.6 Erbium Trench Transitions, Active Laser Devices and Other Components

There are a number of other components necessary to make the lasers work. Key among these is the transition into and out of the Erbium trench. This transition was

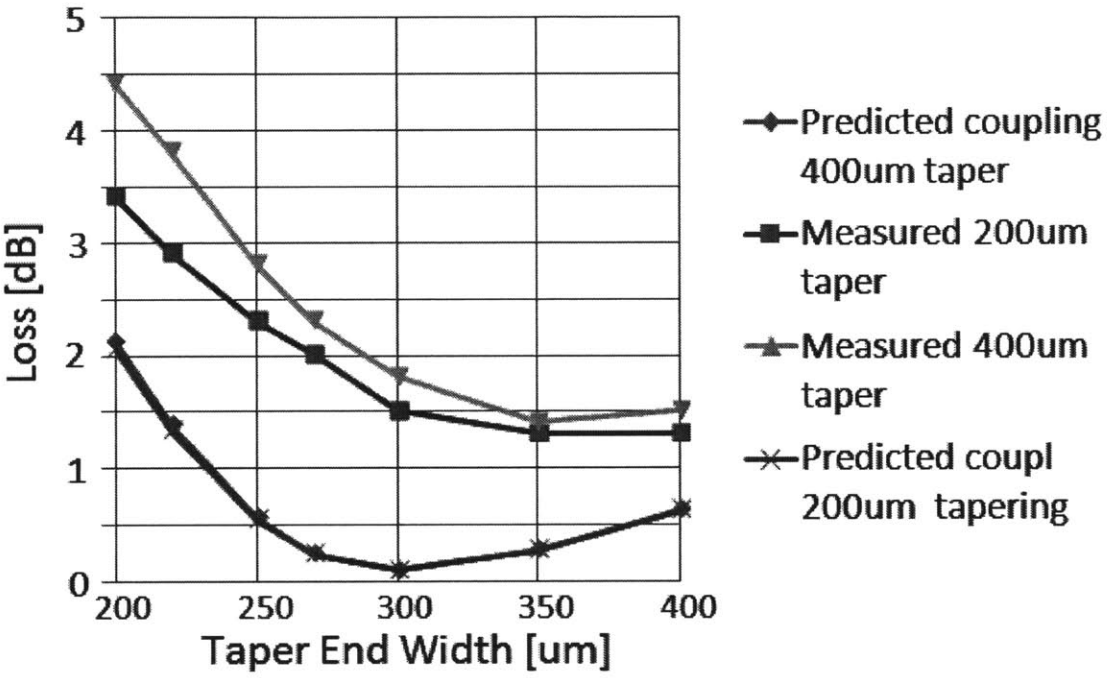

Figure 4-16: Measured and predicted loss vs. taper width for taper type one. As can be seen an extra 1-1.5dB of loss is seen in the measured results, though the trend agrees. Additionally, length has an effect in the measured results, with the longer device performing worse.

Number	Taper width [nm]	Taper length [μm]	X-width of mode at 1550nm [μm]	Y-width of mode at 1550nm [μm]	Predict. Coupling loss cleaved fiber [dB]	Predict. Coupling loss 3μm spot size lensed fiber [dB]	Raw measured coupling loss 3μm spot size lensed fiber
1	270	300	2.35	2.75	-1.3983	-0.2335	-2.5
2	270	150	2.35	2.75	-1.3983	-0.2335	-2.2
3	270	100	2.35	2.75	-1.3983	-0.2335	-1.95
4	400	400	1.6	1.45	-5.3474	-0.6303	-1.8
5	350	400	1.75	1.75	-3.9927	-0.2727	-1.7
6	300	400	2	2.25	-2.3465	-0.0970	-2.1
7	270	400	2.35	2.75	-1.3983	-0.2335	-2.6
8	250	400	2.55	3.05	-0.8001	-0.5268	-3.1
9	220	400	3.3	4.05	-0.2539	-1.2842	-4.1
10	200	400	3.8	4.75	-0.1335	-1.9656	-4.7
11	400	200	1.6	1.45	-5.3474	-0.6303	-1.6
12	350	200	1.75	1.75	-3.9927	-0.2727	-1.6
13	300	200	2	2.25	-2.3465	-0.0970	-1.8
14	270	200	2.35	2.75	-1.3983	-0.2335	-2.3
15	250	200	2.55	3.05	-0.8001	-0.5268	-2.6
16	220	200	3.3	4.05	-0.2539	-1.2842	-3.2
17	200	200	3.8	4.75	-0.1335	-1.9656	-3.7

Table 4.2: Summary of measurements of taper structures.

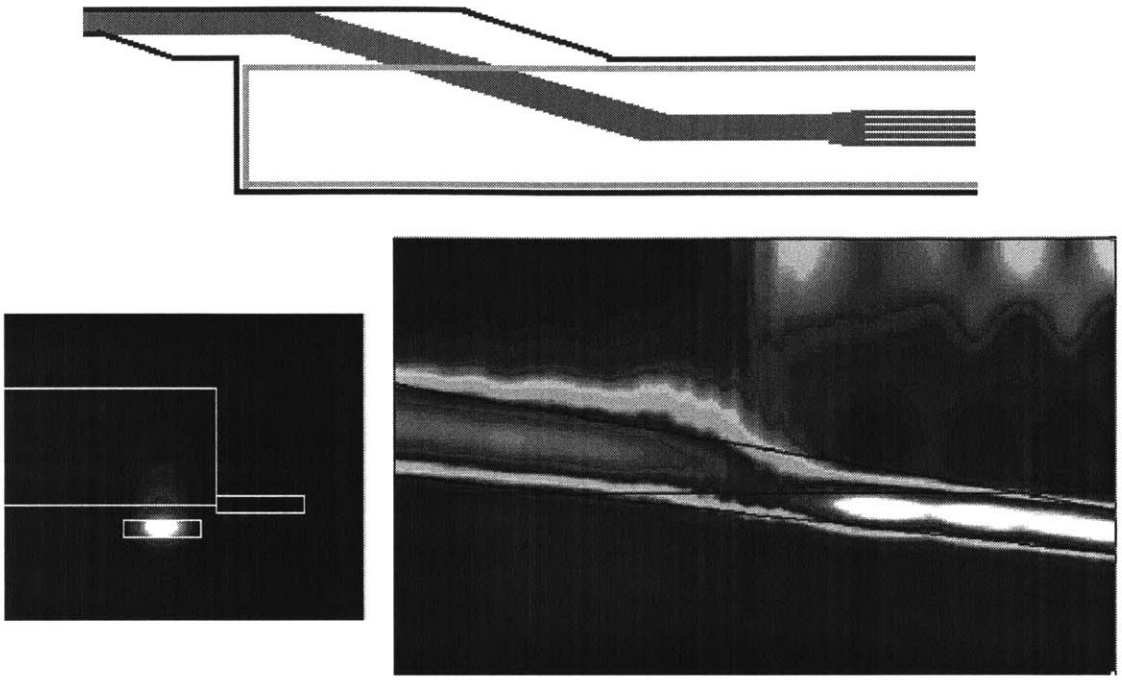

Figure 4-17: Diagram of Er trench transition structure (top) and Fimmwave™ modesolver (left) and transmission simulation of same (right)

originally designed by Purnawirman and the cw laser team for use in cw lasers, and is pictured in Fig. 4-17. We included test structures of this structure on our chips and measured the loss per transition to be between 0.25dB and 0.5dB. While we use this structure in our lasers, future work should aim to find a more robust transition structure, as this current structure is sensitive to fabrication variations, especially to changes in Er:Al_2O_3 height and index (it has a tendency to fail for lower indexes and thicker layers as well as for lower nitride indexes).

While the 1480nm pump can transverse many of the same elements as the 1550nm laser band, when a 980nm pump is desired other structures are needed. A separate taper was used to couple the 980nm light on chip. This taper was designed via a similar process to that described in the "Output Tapers" section above and was measured to have loss <1.4dB, thus allowing for good on-chip coupling of pump light. Because bend performance at 980nm is questionable, it is desirable to pump each straight gain section separately thereby avoiding relying on the bends. To do

that, it is necessary to split the light and send a fraction to each straight gain section. A three-way 980nm pump splitter was designed. These components were measured and found to successfully split the 980nm pump light into three different gain sections allowing those sections to be pumped all the way down the length of the chip. Once the pump light is split it is necessary to get it into the laser trench. Confinement of the 980nm light to wide waveguide in the nitride is high enough to allow for the 980nm pump to simply be routed under the trench without any special structures — more than 90% of the light is transmitted when the trench wall is hit head on. Getting into the trench, however, is not the same as getting into the laser waveguide. To accomplish this an adiabatic pump coupler structure, which doubled as a pump stripper structure where necessary, was again borrowed from Purnawirman and the cw laser group. Test structures were included in our chips and the devices were found to separate and combine the 1550nm and 980nm light properly, while adding less than 0.5dB of loss to the cavity at 1550nm.

In addition to the above, specialty components are needed in designs for actively mode-locked lasers to couple light down to the silicon layer and modulate and manipulate it there. Such devices may also be important for feedback loops aimed at controlling MLL repition rate and carrier envelope offset frequency. The first device needed is a taper to couple between the silicon nitride and silicon layers. One challenge in designing this taper is that the width of the side layer between the silicon and the nitride layers can vary. A taper were both nitride layers taper down together and force the mode into the silicon layer (which tapers up) was designed in Fimmwave and simulated to have more than 99% transmission for oxide separation distances across the range 70-130nm, therefore demonstrating its robustness to fabrication variations. Test structures of these tapers were measured and were found to have $<$0.25dB loss or 94% transmission.

Once the light is in the silicon layer, it needs to be modulated and otherwise manipulated. A Mach-Zehnder modulator is chosen to modulate the light due to its ability to simultaneously modulate large optical bandwidths. The modulator design is based on that described in Chapter 3. Two different types of modulators are

used. A reversed-biased modulator exactly like that in Chapter 3 is used to ensure high enough speed variations. However, because loss is also a consideration and the repetition rate is low enough in our long cavities, a p-i-n diode silicon modulator designed for operation in forward bias is also included. This device allows for greater extinction with lower loss. In this second case the added loss is expected to be >1dB as compared with 2.5dB in the reversed biased case, and higher modulation depths are achievable in a shorter modulator length. Finally, to allow for active control of the devices, a separate heater and phase-shifter are included to allow for manipulation of the phase and amplitude of the light in the cavity in order to stabilize the MLL comb.

4.4 MLL Architectures and Lasing Results

Now that we have designed all the necessary components, they can be combined to create complete MLL systems. As mentioned above there are two main passive architectures: a simple, safe architecture based on a large spiral gain section, and a more complicated, but more compact s-bend laser structure. In addition to this there are actively mode-locked laser structures. These are based on the s-bend design to conserve mask space in the more expensive full-build run, and contain an integrated MZ modulator to act as a pulse-shaping element. Each of these structures will be described below.

4.4.1 Passive Designs

Spiral Lasers

The spiral laser architecture was shown in Fig. 4-2.a., and is repeated here in Fig.4-18.a. to save the reader the trouble of flipping back. As outlined in the diagram, this structure consists of a loop mirror on one end and a taper to a satruable absorber (SA) on the other end. In between the two, the light is coupled into the laser trench and a long spiral section of the gain waveguide is present. The advantage of this

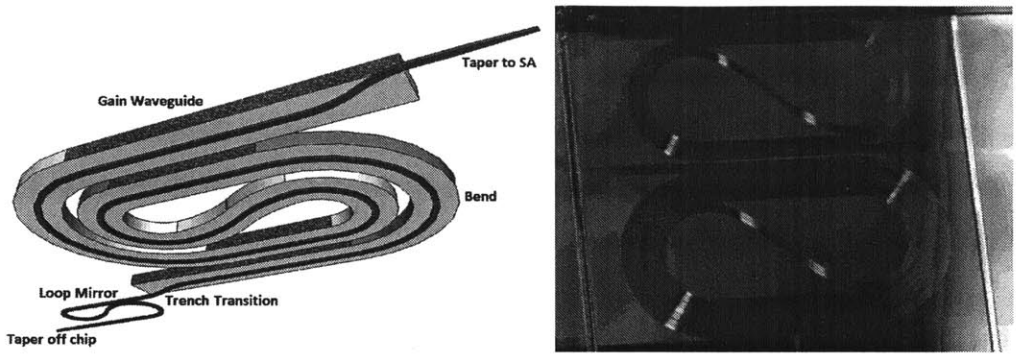

(a) (b)

Figure 4-18: (a) Diagram of passive spiral laser architecture, (b) photograph of spiral lasers on chip.

architecture is its simplicity. There are few components in the cavity and therefore there are few component-contributed losses at the laser wavelengths or the pump wavelength. Further, long structures can be laid down. We place two different sets of these lasers on the mask, a 20cm version and a 10cm version. These lengths correspond to repetition rates of approximately 0.97 GHz and 0.47GHz respectively and a single laser takes up approximately 1.5cm^2 of area on the chip ($2.5 \times 0.6cm$ for the 10cm spiral laser and $2.5 \times 0.65cm$ for the 20cm spiral laser). In each version we vary the key laser parameters (loop mirror gap, taper to SA, gain waveguide) to allow for the best chance of working. We interleave these variations on the mask to conserve space. A picture of the spiral laser section of the fabricated chip is shown in Fig.4-18.b.

A high power 1480nm pump was used to pump the laser. A high power pump was needed in order to pump through the entire structure due to both high absorption of the pump in the waveguide and the large pump (and laser) mode size. (A large laser mode is desired so that high output power can eventually be achieved). An image of the green fluorescence in the pumped laser spiral is shown in Fig. 4-19. Spiral laser structures with both a high mirror reflectivity (85%) and a low mirror reflectivity (20-30%) were fabricated. In the high reflectivity case, much of the pump laser was also reflecting, and, consequently, only the 10cm spirals could be fully pumped and made to exhibit laser oscillation. In the low reflectivity case, the 5-7dB loss at the

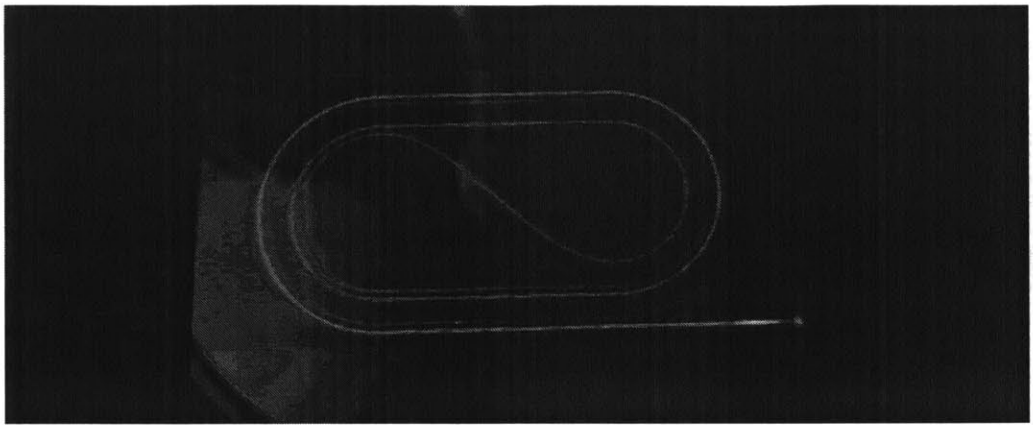

Figure 4-19: Green florescence from high power 1480nm pump outlines the laser structure.

output coupler was too high to allow laser oscillation to occur in the 10cm structures, as there was not enough gain. However, the 20cm spirals could support the higher losses and, due to the lower mirror reflectivity, be fully pumped. These structures were observed to oscillate and to exhibit Q-switched mode-locking.

An Er:Al_2O_3 film with a thickness of 1100nm and an erbium concentration of $1.6 \times 10^{20} cm^{-3}$ was deposited on these laser chips and the output facets were polished to allow for good alignment of and coupling to the SA. The 20cm spiral lasers were first tested in continuous-wave (CW) operation by butt-coupling a mirror, instead of a SA, to the chip facet. The resulting laser spectrum is shown in Fig.4-20.a., and the output power as a function of input pump power in Fig.4-20.b. One can see the narrow spectral line-width and exponential increase in output power indicative of laser oscillation above threshold. The laser wavelength was around 1560nm, which happens to be around the peak of the net gain at maximum pump power in our gain waveguides (0.67dB/cm gain at 1559nm). The oscillation threshold is about 1.6W of on-chip pump power.

Once CW operation was observed, the mirror was exchanged for a SA and the laser was retested. An RF spectrum and oscilloscope trace of the resulting pulses is shown in Fig.4-21. The expected 470MHz repetition rate is observed, however the amplitudes of the comb lines are highly variable due to the additional presence of

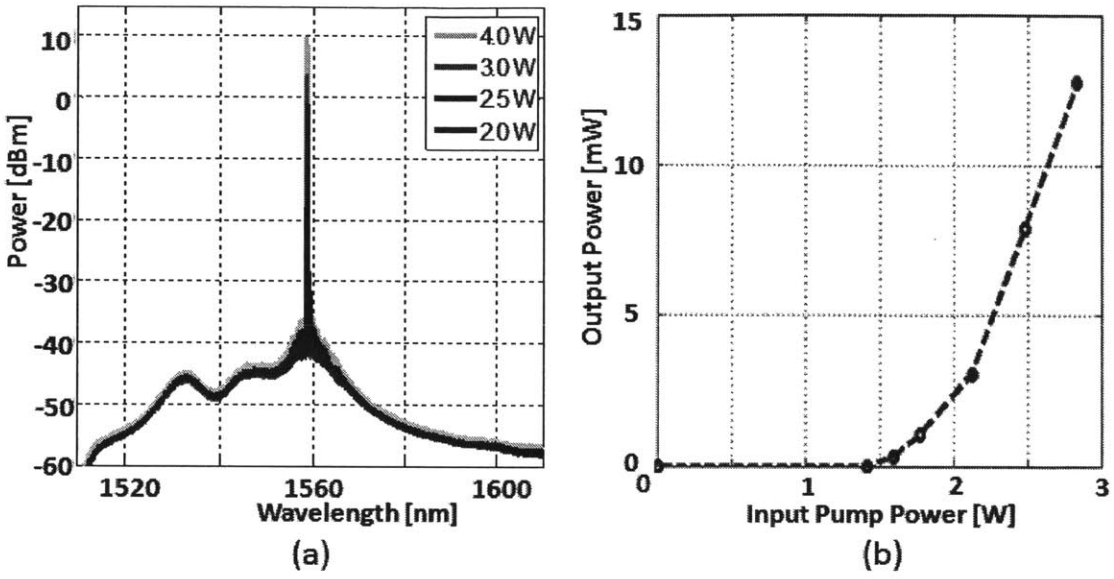

Figure 4-20: (a) Ouput spectrum of cw laser as a function of pump power. (b) Laser output power vs. input power on chip. oscillation threshold is around 1.6W. This laser had a mirror reflectivity measured to be 30%. From [47].

Q-switching in the cavity. The Q-switched pulse bursts can clearly be seen on the oscilloscope trace. Each Q-switched pulse burst is comprised of a series of pulses spaces by the MLL repetition time (see Fig.4-21), thus confirming that we have *both* Q-switching and mode-locking present in the cavity [78, 79]. Because an excess of normal dispersion is a possible cause of Q-switched mode-locking [79], and because we know, due the thinness of our Er:Al_2O_3 film that our cavity dispersion will be normal (see the section on the gain waveguide above), we suspect that this excess of dispersion is, in fact, what is causing the observed Q-switching. Because a thicker film cannot be used as it expands the bending radius by too much, a redesign of the gain and bend waveguide is necessary to achieve anomalous dispersion. This will such a redesign will be discussed in the future work section below.

Compact S-bend Lasers

We next turn our attention to the compact s-bend laser structure layouts. This architecture was also shown in Fig 4-2.b., and is repeated here in Fig.4-22. to save

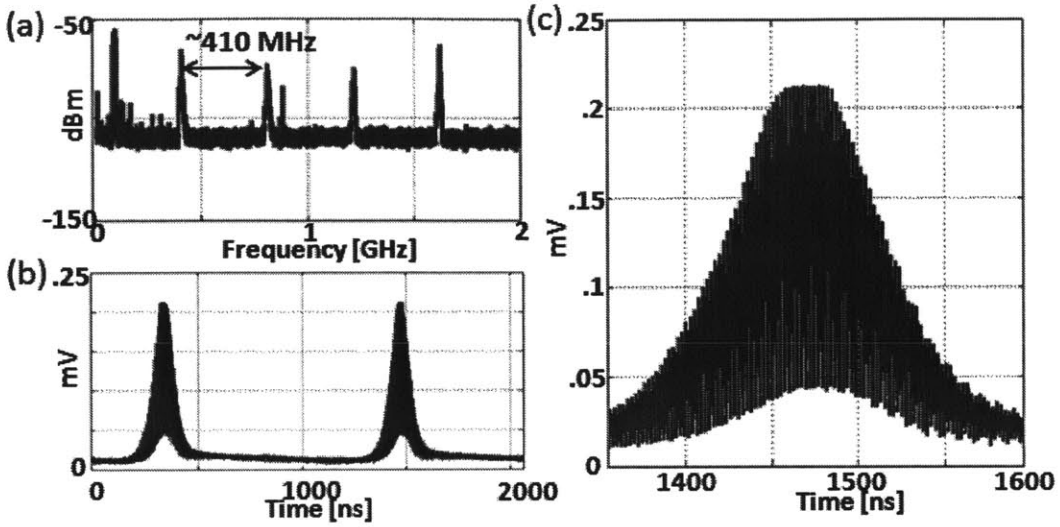

Figure 4-21: (a) RF spectrum of laser in Q-switched mode-locking regime, (b) Oscilloscope trace of same, (c) Zoomed in oscilloscope trace of pulse burst. From [47].

the reader the trouble of flipping back. As outlined in the diagram, this structure consists of multiple straight gain waveguides connected in a serpentine layout by tight waveguide bends. The advantage of this architecture is that it is very compact, requiring $< 0.175cm^2$ of chip area (2.5x0.025cm to 2.5x0.06cm). In the future, this compact size would facilitate integration of this laser with larger systems, especially those present on the full silicon or silicon+CMOS runs where space is at a premium. Another advantage of this structure is that both the straight waveguide sections and then bends function with thick Er:AL_2O_3 layers, thus allowing for anomalous dispersion to be achieved in the gain waveguide. Further , the lack of wrapping means that it can be more easily scaled to higher (or to any particularly desired) repetition rates. The disadvantage of this layout is its complexity; many components in the cavity means many chances for a component to fail and increased component-contributed losses.

While Fig.4-22 shows one version of this laser, it is worth knowing that there are actually several variations. First, some of these lasers have three gain sections, for a total of 6cm in length and a repetition rate of 1.5GHz, and some have five, for a total of 10cm in length and a repetition rate of 0.92GHz.. Of the lasers with

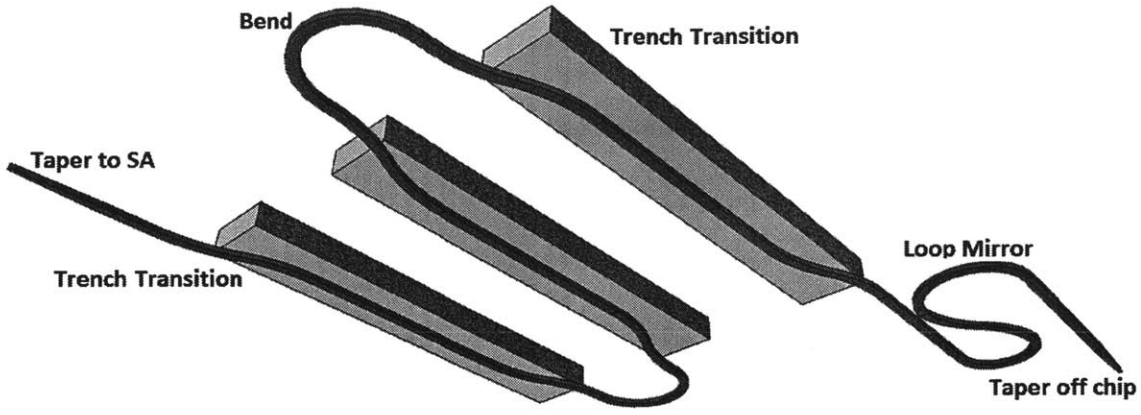

Figure 4-22: (a) Diagram of compact passive s-bend laser architecture.

three gain sections some are made to work at the 1480nm pumping wavelength we have been using in the other structures and some have the extra pump input couplers and components necessary to work with a 980nm pump instead. Fig.4-23. shows the difference between the 980nm and 1480nm pump designs. Additionally, we include structures with both nitride bends (where the light is coupled out of and into the laser trench before and after the bend) and the compact trench bends described in the section on bends above. Just like in the case of the spiral architecture, for each version here we vary the key laser parameters (loop mirror gap, taper to SA, gain waveguide) to allow for the best chance of achieving oscillation and mode-locking.

Because there are so many components, we begin by testing the passive loss of the lasers. To do this, we measure the transmission at 1550nm though one of our laser structures with an undoped film. We expect to see 12-15dB of loss in this transmission measurement. The contributions to this loss are summarized in table 4.3. The measured result is shown in Fig. 4-24. As can be seen, about 15-16dB of loss is observed across the wavelength of interest, a number which agrees with the predicted value. From this measurement, we can be fairly certain that our understanding of the component losses is correct. The total cavity loss in this case would then be 11.5-16.5dB depending on the exact construction of the cavity. We note that this translates into needing a net waveguide gain in the range of 0.7-1dB/cm in order to reach threshold and achieve oscillation. When these designs were originally done with an estimate

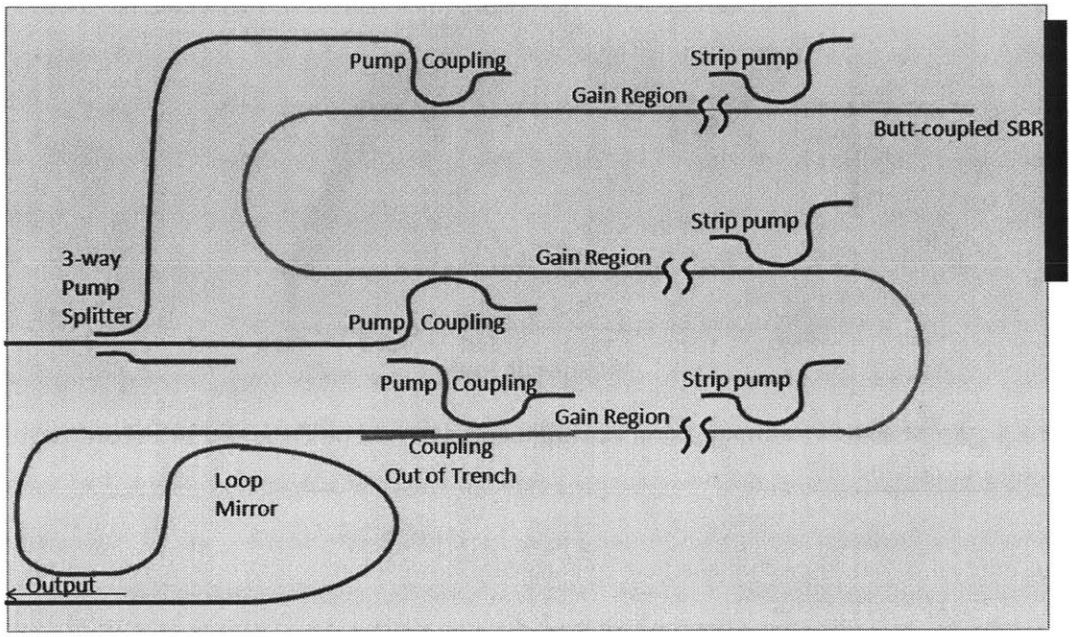

(a)

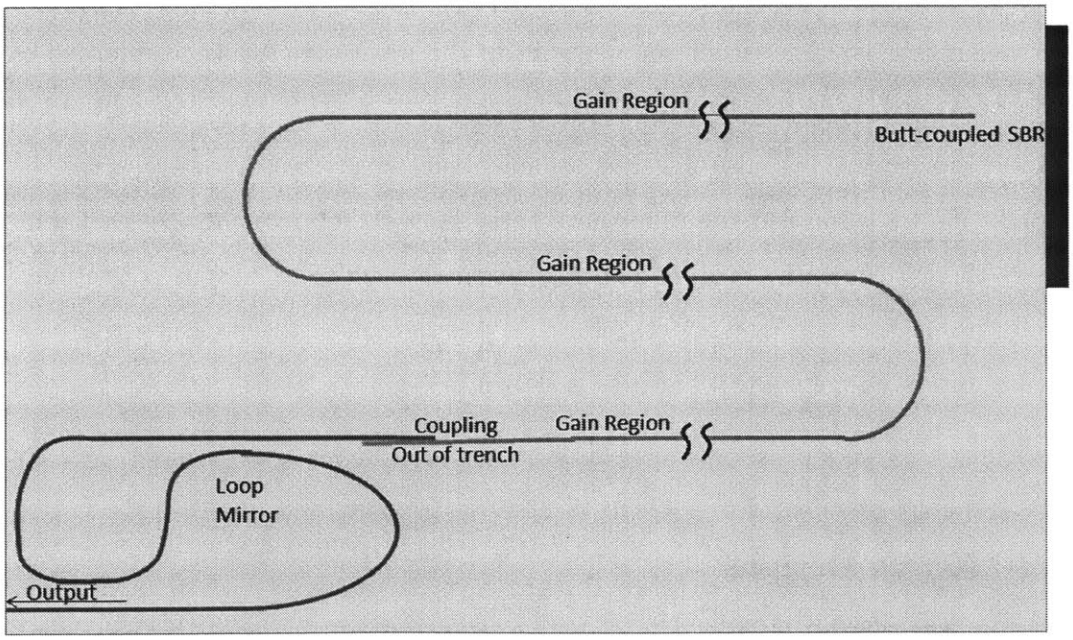

(b)

Figure 4-23: Schematic of compact s-bend laser to be pumped at (a) 980nm and (b) 1480 nm.

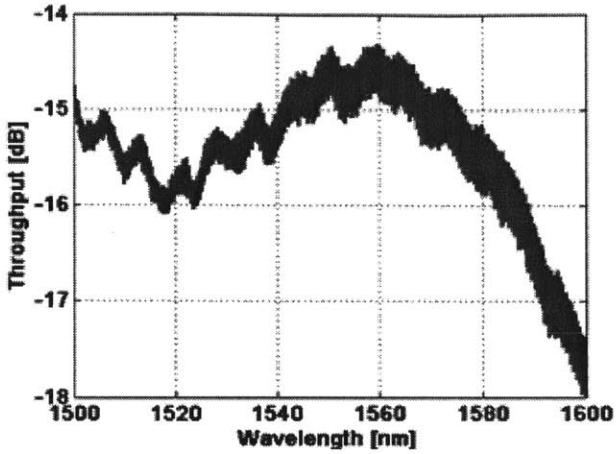

Figure 4-24: Measured transmission through passive s-bend laser architecture as a function of wavelength.

of the gain of 1-2dB/cm this seemed reasonable, however now having measured the gain to be in the 0.5-0.7dB/cm range it is clear that these losses cannot be supported and these structures will not oscillate. In order to achieve laser oscillation in these structures it is clear that either the net gain in the waveguide needs to improve or components need to be redesigned to have lower loss.

For completeness, we also checked to make sure that these structures can be pumped. We have coupled the pump into these laser structures and observed fluorescence though the entire erbium gain regions, indicating that the pump light successfully propagates through the structure. A photograph of the florescence from a 980nm pump in one of these structures is shown in Fig.4-25. along with a picture of the chip.

4.4.2 Active Designs

Finally, we have also designed actively mode-locked laser structures. These structures are based on the s-bend footprint from above in order to save space on the expensive full silicon photonics masks. A diagram of this architecture is shown in Fig.4-26. As can be seen, it mimics the s-bend laser footprint, but also contains a Mach-Zehnder modulator to facilitate pulse-shaping. The modulator is switches between the "on"

Component	Measured Loss	Comments
Taper on-chip	-1.5 dB	Coupling loss to 33μm spot-size lensed fiber.
Loop Mirror	-2 dB	Mirror transmission (30% reflection)
Erbium trench transition	-0.25 to -0.5 dB	Six transitions total
Nitride Bend	negligible	known from ring measurements
Gain Waveguide	-1 dB	Loss in passive Al_2O_3 layer, three gain sections
Taper off-chip	-4.5 dB	Taper to SA coupling loss to 3μm spot-size lensed fiber.
TOTAL	-12 to -15 dB	Expected single-pass transmission loss

Table 4.3: Summary component loss measurements.

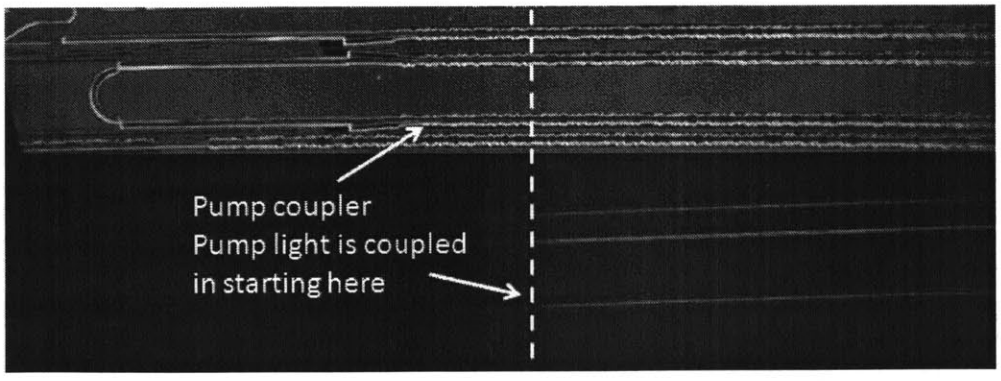

Figure 4-25: Fluorescence in compact s-end laser designed for 980nm pumping. A 3 way pump splitter (out of the frame) couples the pump light into each of the three gain sections.

and "off" states such that the peak of the pulse travels through the modulator when it is "on" and is transmitted with little loss while the edge of the pulse and the spontaneous emission in the cavity pass through while the modulator is "off" and are extinguished. This suppresses other oscillations in the cavity and also shortens the pulse in time, creating the desired mode-locked-laser pulse output. Note that the choice of modulator location, right before the taper to the SA is not random. Because we have a linear, not a circular, cavity, the pulse will have to pass through the modulator twice in any given round trip. Placing the modulator near one cavity end insures that these two passes are as close as possible and allows the modulator to still be "on" for both of them, preventing the pulse from being extinguished on one pass, while ensuring that there are not periods of time in which the modulator is "on" but the pulse is not there that would allow another pulse to develop in the cavity. (While harmonic mode-locking, where more than multiple pulses are circulating in the cavity with a constant time spacing, can be useful in some cases, due to added difficulties in stabilizing the pulse energies and reducing the timing jitter of such a laser it is not our goal here.) Placing the modulator at the end of the cavity also allows us to easily pump the rest of the laser without having to worry about pumping the length after the modulator [80].

In active mode-locking Gaussian pulses of the form $e^{-\frac{t^2}{2\tau^2}}$ are common. Here, $\tau = (\frac{2g}{\Omega_g^2 M \omega_M^2})^{(1/4)}$, where g is the gain, Ω_g is the gain bandwidth, M is the modulation depth of the modulator, and $\omega_M = 2\pi f_{ref}$ where f_{rep} is the repetition rate. Plugging in the numbers for our system, we find that we can expect pulses in the 10-30ps range. As this is rather larger than we would like, so the taper to the SA and SA are also still included, this allows the SA to further shape (clip) the pulse to create shorter pulses and allow for the better timing jitter achievable from passively mode-locked systems [80, 81]. In this joint system, essentially the modulator is used to kick start the mode-locking and the SA to really create a nice pulse.

We note that, in addition to the modulator in the cavity, there is also a phase-shifter and heater. These are needed for the simultaneous control amplitude and phase for feedback to stabilize f_{rep} and f_{ceo} that simulations and table-top experiments

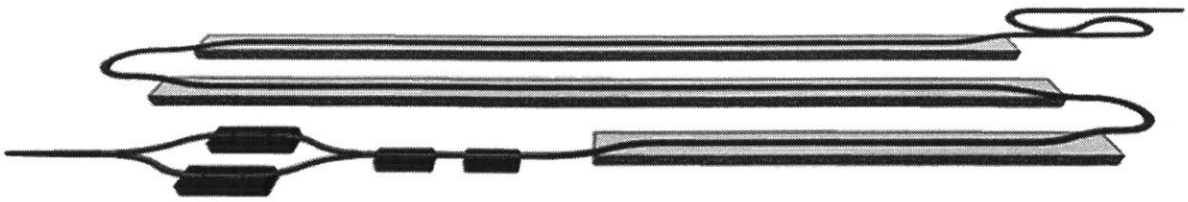

Figure 4-26: Diagram of actively mode-locked s-bend laser configuration including MZ modulator for active pulse shaping and intracavity controls for comb stabilization. Phase-shifters are shown in gray, the intracavity heater in brown, the silicon in gold, and the nitride in red.

have indicated will be necessary to achieve the ESPIOR oscillator specifications. In particular, we would like orthogonal controls, that is to be able to change f_{rep} without changing f_{ceo} and vice-versa. The three active devices shown in the cavity in Fig.4-26. should allow this. The MZ modulator, in addition to allowing for active mode-locking can also serve as an amplitude modulator. The silicon phase-shifter can be forward biased to control cavity loss (and thus pulse power and nonlinear phase shift) or can be biased to control round-trip phase. The heater can also change round-trip cavity phase or remove any unwanted phase added by the other controls. Together these three elements should allow for appropriate stabilization of the MLL comb.

Like the s-bend laser case, variations exist to allow for 980nm pumping,1480nm pumping, and erbium trench vs. nitride bends. However, due to space constraints only five-gain-section lasers were placed on the mask for 1480nm pumping wavelengths. Variations in the MZ modulator phase-shifter design to allow for wither forward bias or reverse bias operation were also included as described in the component section above. Unfortunately, because they are based on the high loss s-bend configuration, it is unlikely that these lasers will oscillate. (They have not yet been received back from the fabrication facility). To solve this problem, either component (especially bend) losses need to be reduced, or gain needs to be improved (perhaps by lowering background material loss or but inclusion of ytterbium doping along with the erbium doing in the Al_2O_3 film).

4.5 Summary and Future Work

The above sections outline the architecture for an integrated mode-locked laser. Two different passive and one active architectures are described. The necessary integrated components for these lasers (gain waveguides, bends, loop mirror, tapers, etc...) are designed, simulated in detail, and experimentally verified. One passive architecture consists of a large spiral gain waveguide between a loop mirror output coupler and a butt-coupled satrurable absorber. This structure has the advantage of being simple and low loss, but does not allow for anomalous dispersion to be achieved. These structures were shown achieve laser oscillation and Q-switched mode-locking. In CW operation they had a threshold power of 1.6W and a lasing wavelength of 1560nm. In the Q-switched mode-locking regime they exhibited a repetition rate of 410MHz, about that predicted. The other passive architecture was based on a compact s-bend footprint. This architecture has the advantage of small size, greater flexibility, and anomalous dispersion in the gain waveguide. Unfortunately, it also has more cavity loss, and the net gain, which was lower than that predicted, was not enough to overcome this loss and allow for laser oscillation. Finally, a number of actively mode-locked structures based on the compact s-bend footprint where also designed and laid out. They include both a silicon MZ modulator in the cavity and an SBR butt-coupled to the cavity end. The active pulse-shaping provided by the MZ modulator should make mode-locking simple to achieve and the SA should then act to further shorten pulses. An additional phase-shifter and heater in the cavity also allow for comb stabilization. However, while these chips have not returned from the fabrication facility and been tested, their reliance on the s-bend footprint makes it unlikely that they will oscillate as they, too, will likely have too much cavity loss.

To improve the s-bend designs it is necessary to lower component losses or increase waveguide gain. This later, highly dependent on the material is unlikely, though experiments of co-doping the Al_2O_3 with both erbium and ytterbium are one option (that is being pursued). Lowering component losses is a more promising path forward. In particular, optimizing the bend structure (which is repeated with every gain section

and currently has 1dB loss per 180^o bend) would help significantly. Creating an adiabatic tight bend that avoids the current offset in the laser trench and nitride guide (and the change in nitride guide geometry), while remaining in the laser gain medium is on potential option. This would allow for more gain structures to be strung together. Changing the trench transitions element to have less loss and be less sensitive to fabrication variations would also help. Careful managing of the pump to make sure that the whole structure can be pumped, even at longer lengths, will also be necessary. Fixes to this laser structure should also allow for the active structures to work.

Improving the spiral laser structure mainly consists of fixing the gain and bend waveguide dispersion. In this case movement to a w-waveguide structure (where the one cut up nitride block in the center is joined by two nitride blocks on the side to create anomalous dispersion in a way analogous to that in [91] and where oscillation would then occur in the central, third-order mode) has shown great promise in simulations. A slight variation on this structure can also be implemented in the bends to allow for anomalous dispersion there. A combination of these changes to the gain waveguide and shortening the parts of the laser cavity outside of the waveguide allow for the total cavity dispersion to remain anomalous and should allow for stable mode-locking.

Another important piece of future work has to with determining system parameters. We now have a measurement of the net gain and the material dispersion. It would be good to have a rigorous measurement of the dispersion and nonlinearity in the cavity. Dispersion tests structures like those in [74] have been included with the new spirals and the hope is to be able to measure or develop a good model for the nonlinearities as well. Knowing these parameters would greatly help in designing laser systems able to meet the required specifications.

Finally, perhaps the most interesting avenue for future work rests with the integrated active structures. First, in order to achieve laser oscillation, component losses need to be lowered as in the case of the s-bend lasers. However, if and when mode-locking is achieved, these structures would provide a fascinating platform from which

to study the different contributions to mode-locked laser noise and how those contributions and their effects and the effects of the various feedback methods available for removing them differ between free-space, fiber, and on-chip MLL systems.

Chapter 5

Integrated Pound-Drever-Hall Loops and Other Integrated Systems

The work in this thesis has aimed to support the development of an integrated ultra-low-noise microwave oscillator. In previous chapters we have discussed the silicon photonics platform necessary to design and create the system, as well as two devices, an integrated single-sideband modulator and an integrated mode-locked laser, that are key to the oscillator. In this chapter we discuss some early work on creating the actual subsystems of the oscillator, and mainly with creating the integrated Pound-Drever-Hall locking loop.

5.1 Other Devices and Integration with CMOS circuitry

Before we continue, we note that many more devices than those discussed in the previous two chapters are needed to make these systems. Many of these devices, such as tapers onto the chip and the waveguide used for propagation, do not merit discussion on their own, however others deserve a brief mention.

Adiabatic 3dB couplers where used in these designs. These couplers were based off of a previous design created by Michael Watts and translated by the author into

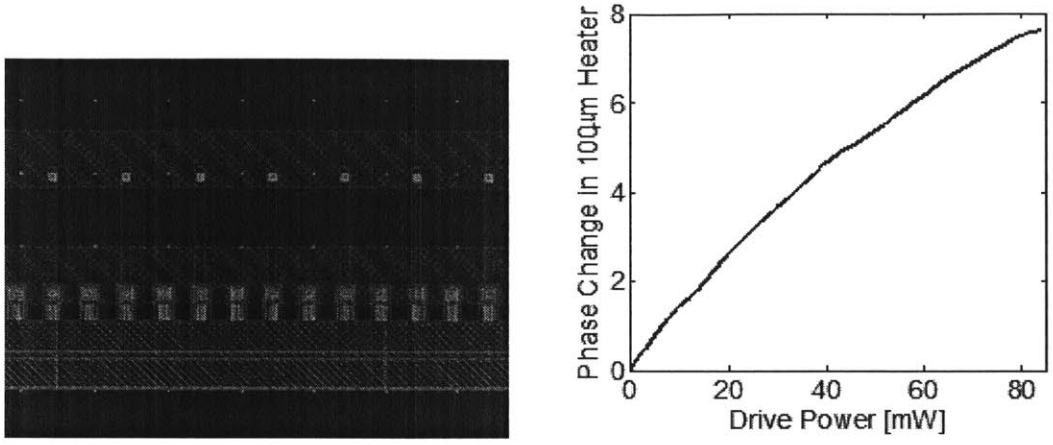

(a) (b)

Figure 5-1: (a.) Mask layout of integrated heater, (b.) Phase change through heater vs. input DC power

this process. Simulations of these couplers were performed in Fimmwave, and test devices showed the splitting to be consistently correct to within $\approx \pm 3\%$.

An integrated heater was created to bias the modulators and serve as an external phase shifter where needed. Simulations were run in COMSOL Multiphysics ®Modeling Software in order to model device heating and ensure that a π phase-shift could be achieved. Performance of this heater was also discussed briefly in Chapter 3. A schematic of the heater, along with the phase shift versus DC input power of the device, is shown in Fig. 5-1. The heater was 100μm long and made in the rib waveguide structure used for the modulator. This particular architecture was chosen not for efficiency, but because it required lower drive voltages and thus could conceivably be driven by CMOS circuitry. However, in our case we always connected the heaters directly to external pads, just in case more voltage was needed, as did, indeed, turn out to be the case.

A common device to all the oscillator subsystems is an integrated photodetector. A Germanium photodetector was designed based on the work presented in [19]. The layout of this detector is shown in Fig.5-2.a. The responsivities of the devices at 1550nm were measured to be in the 0.6-0.85 A/W range depending on the exact variation of the device tested. Dark currents initially varied from the nA to the 100s

of μA, but, due to process refinements at CNSE, have stabilized in the nA region. In our control loops these diodes need to be driven by CMOS circuitry. The CMOS TIA driver for this device, designed by Jonathan Leu, is shown in Fig.5-2.b. 3D integrated Silicon + CMOS devices have recently been returned from the foundry at CNSE Albany. Preliminary tests of these devices show dark currents in the nA range, with responsivities roughly estimated to be in the 0.5-1.0 A/W range for 1530-1550nm light (the range tested). A plot of measured output current vs. input power on these 3D integrated devices is given in Fig.5-2.c.

The speed of these integrated photodiode devices is more problematic. While the non-3D integrated device works out to 10s of GHz, the speed of the 3D integrated PDs with driving circuitry appears to be limited to 2-3GHz and does not quantitatively follow expected trends with TIA resistance, diode capacitance or circuit bias location, though qualitatively the trends are there. So, for example, the speed does improve as the TIA resistance is decreased, but not by the factor expected. The reason for this is unclear and current work is ongoing. Many of these analog control loops only need to work at low speed, and, therefore, performance may be alright despite the current speed limitation. Obviously for the last PD in the oscillator, which needs to work at the RF oscillator speed of 20GHz, this issue has to be overcome.

Another key part of our oscillator system is the high-Q, athermal ring resonator to which the CW laser is locked. It is the frequency stability of this resonator that ultimately determines the frequency stability of the whole system. An analysis of the noise sources in the ring quickly reveals that it is thermal fluctuations and, thus, the thermal stability of the resonance that will dominate resonator frequency noise. The high-Q athermal resonator is currently under development by another student in the group. The idea is to create a nitride ring under an open trench (in the same way the laser architecture in Chapter 4, above, uses a nitride waveguide under a trench) and to deposit a material in the trench that has a thermal change in refractive index that counteracts that of the nitride guide and oxide cladding similar to what was done in [132]. For the purposes of the PDH loop describe below, three different rings were used. The first was a less stable, silicon ring, this allows for preliminary tests and

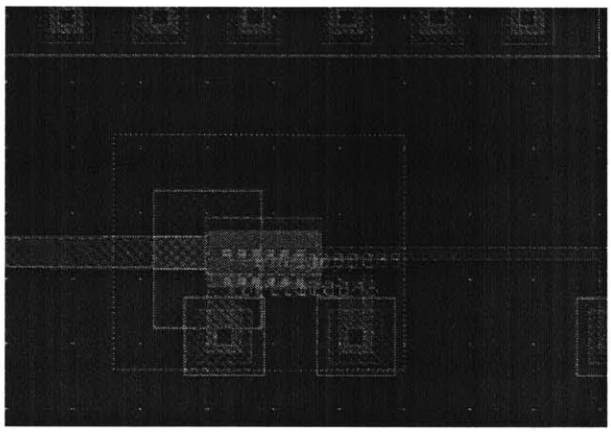

(a)

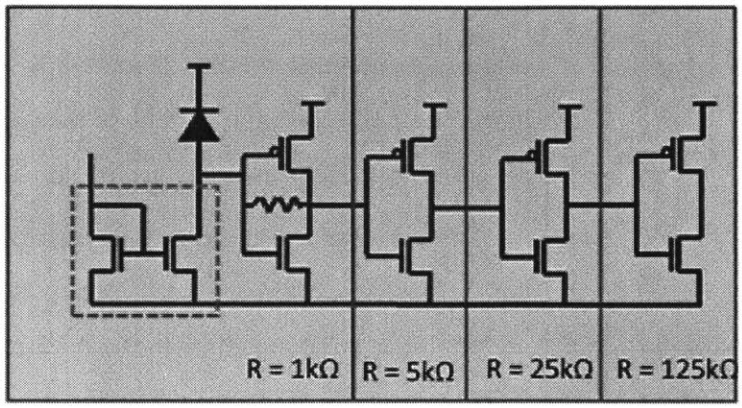

(b)

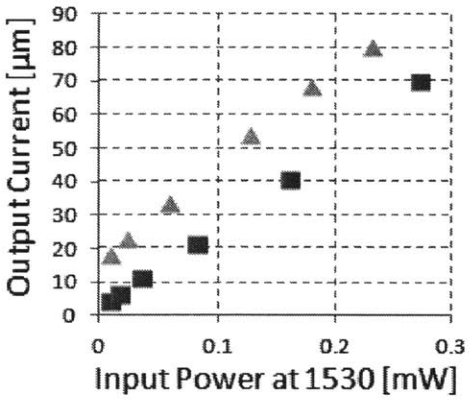

(c)

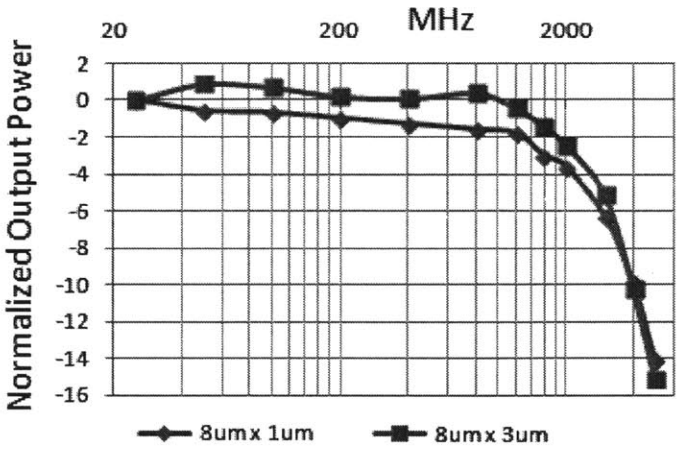

(d)

Figure 5-2: Integrated Germanium Photodetector (a.) Layout of two different device variations, (b) Schematic of TIA drving circuit for PD, (c.) Output current vs, optical power out of fiber input to chip for 8x1μm (red) and 8x3μm (green) 3D integrated PDs with 5kΩ TIA resistance tested at 1530nm, (d.) Speed measurements of 8x1μm (blue) and 8x3μm (red) 3D integrated photodiodes with 5kΩ TIA resistance.

the quickest turnaround time from the foundry. The second was a nitride ring, this allows for more thermal stability, though not enough to hit the oscillator target, but requires a slightly longer fabrication time. It is probably the safest option. Finally, a nitride ring under an open trench was included as well. Because the material and geometry needed for athermal operation was not known at the time the mask was cut, this option allowed for the possibility of a true athermal cavity, but might not meet the geometry requirements needed. Fig.5-3 shows the expected thermal noise in the resonance frequency for the various ring choices. The rings were designed using a combination of FDTD and transfer matrix codes (note that the VerilogA code described in Chapter 2 could be used instead of the transfer matrix code, but wasn't completed at the time). The bending radius was chosen to be large enough that bending loss was trivial compared to material loss, and, thus, allow for the highest Q. The circumference of the ring was chosen to be very large to maximize the thermal stability while still allowing for critical coupling (not too much loss in the ring) and for a large enough free spectral range that modulator sidebands could be located in a flat area between resonances. The silicon and nitride rings were bent into a race-track or a "C" shape to conserve mask space and chosen to have circumferences of 1mm or 6.26mm. In preliminary measurements the silicon rings were found to have a FWHM of 2GHz.

Finally, for some of the Pound-Drever-Hall loops and integrated laser was included on the chip. This tunable laser was designed by Dr. Anna Baldcheva based on the heater design above and the laser design in [133].

5.2 Fully integrated PDH Loop

The Pound-Drever-Hall (PDH) technique is familiar from table top optics as a way to stabilize the frequency output of a laser, by locking it to a frequency reference cavity. For a table top system this is usually a Fabry-Perot cavity. In our integrated system a ring resonator will serve the purpose of the reference. In addition to its necessity for our oscillator, the motivation for integrating a Pound-Drever-Hall locking loop

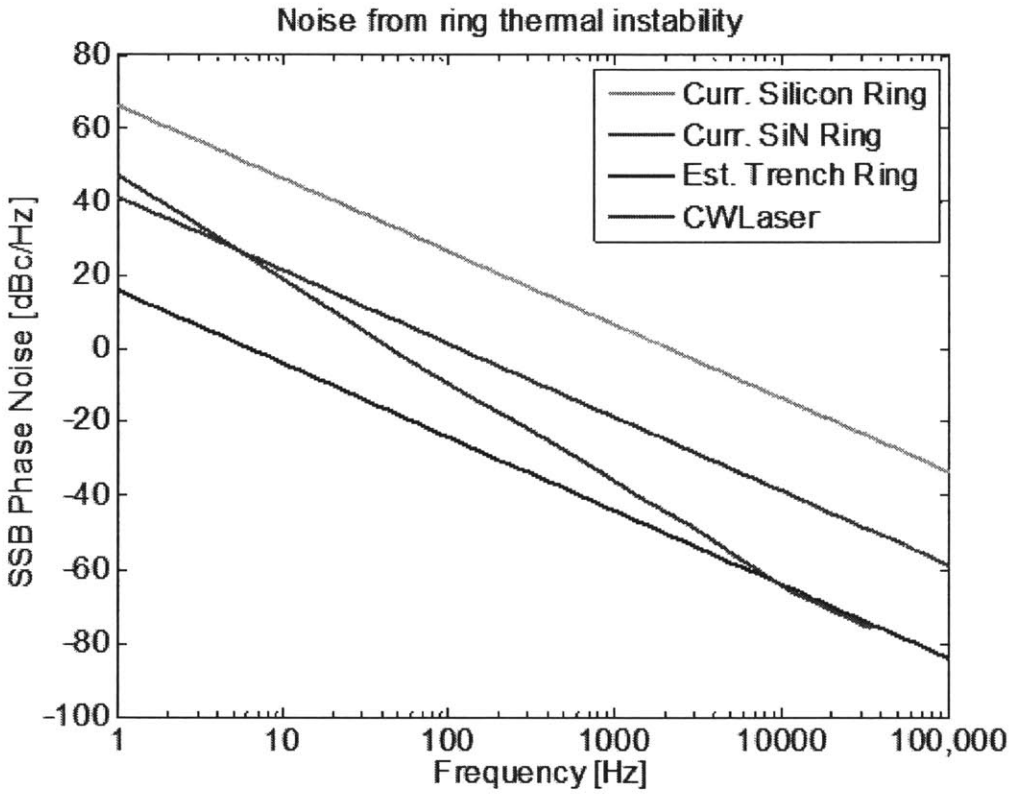

Figure 5-3: Plot of single sideband phase noise due to thermal fluctuations for ring variations on mask.

is two-fold. First an integrated version of the loop will be more compact, and may also be more stable to environmental perturbations such as vibrations. Secondly, one of the ultimate goals of integrated photonic systems is to integrate the light source on chip. However, precision analog systems often require sources with extremely low phase noise, and the locking loops to stabilize such integrated lasers must also be included on chip. One of the most versatile and widely used of these loops is the Pound-Drever-Hall loop.

There is very little literature on integrated PDH loops, perhaps because most integrated systems demonstrated to date are much simpler. In generally, there is very little literature on complicated, integrated, analog systems with both photonic and electronic components. However, in [30], a PDH locking loop is created that uses an integrated ring resonator, instead of a Fabry-Perot cavity as its frequency reference. The rest of the loop is off-chip. This paper is also motivated by the increased stability of an on-chip cavity, as well as by the ease of changing the ring design for various environments.

PDH Loop Theory

A diagram of the proposed PDH loop is shown in Fig.5-4. A brief description of how the locking loop works, following the treatment given in [69] will be given here. PDH loops were first proposed in 1983 [134]. The basic idea behind the system is to stabilize the output of a laser by locking it to a frequency reference cavity, in our case an integrated ring resonator. A feedback loop is established where the interaction between the modulation sidebands and laser carrier frequency with the ring is measured to determine which side of the ring resonance the laser is on, and the laser is then corrected. Ultimately, the width and stability of the ring resonance (the narrower and more stable the better) and the noise in the locking loop will be the limiting factors of the frequency stability of the laser output.

A more detailed description follows. The laser outputs a CW beam at some frequency, f_C. This beam is the fed into a phase modulator (a) which is driven by a local electronic oscillator (b) at some frequency, f_M. The phase modulator generates

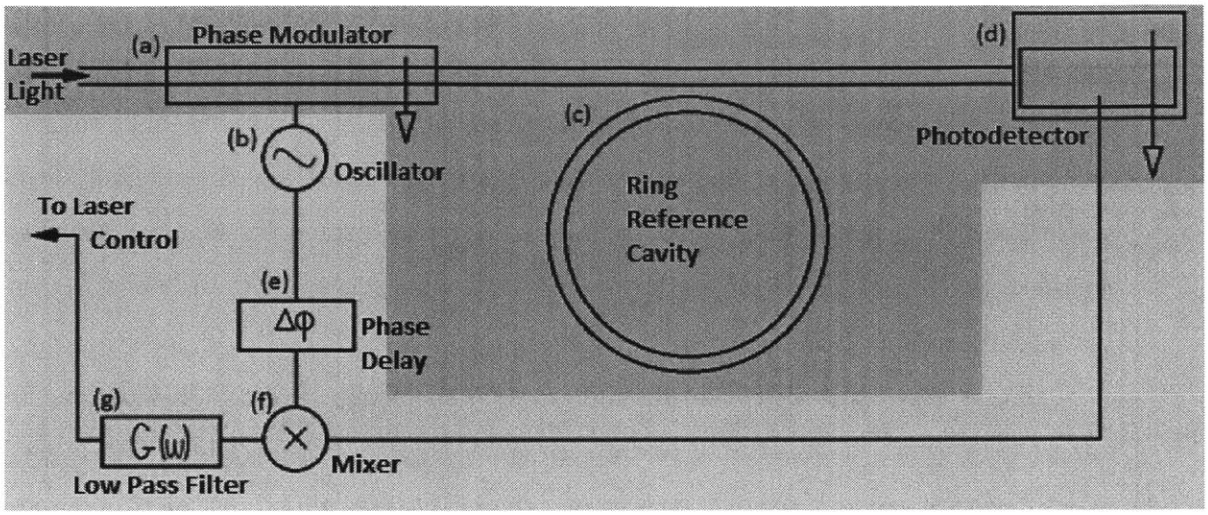

Figure 5-4: Diagram of PDH Loop. Blue denotes electrical components, and pink photonic ones.

two new frequency sidebands of opposite sign at $f_C + f_M$ and $f_C - f_M$. In many table top systems, the phase modulator is made with a Pockel's cell modulator [69]. In our integrated system, the phase modulator will be a reverse-biased silicon diode phase-shifter of the type described for the MZ phase-shifters in chapter 3.

After the phase modulator, the new, modulated signal sees the ring resonator (c) used in our system as a frequency reference. The modulator and the ring cavity are connected on chip via a standard silicon waveguide. Part of the signal couples into the ring and is either dissipated in the cavity or couples back out of the cavity to interfere with the signal that did not enter the ring. The result of the interference is present in the so-called "through port". Looking at light in this port is equivalent to looking at light reflected from a Fabry-Perot cavity. The effect of the ring is to adjust the amplitude and phase of the carrier and sidebands according to how far they are in frequency from the resonance of the cavity, with the sign of the phase change depending on whether the band is higher or lower than the ring's resonant frequency [69].

The modified signal is then detected on a photodetector (d). In our system the photodetector is an integrated germanium photodetector composed of n-type doped germanium grown on a p-type silicon block to form a vertical p-n junction. A poly-

silicon contact is used on the top to contact the germanium. Contact to the silicon is made off to the side. The photodetector, in essence, serves as an optical mixer, mixing the side bands with the carrier and each other. However, the photodiode has limited bandwidth (it cannot respond at the carrier frequency), and therefore only passes a DC component and the components at $f_C - (f_C \pm f_M) = \pm f_M$, which represent the cross-terms between the carrier and each sideband. If the photodiode is AC coupled, the DC component is also lost. (It should be noted that the photodetector may also pass a component equal to twice the local oscillator frequency, but we do not care about this component as it will be lost after the mixer and filter as will be seen below). The photodiode is the last optical component in the system [69].

The resulting signal from the photodiode is then electrically processed. First it passes through a mixer (f), where it is mixed with the signal from the local oscillator (b). This is the same oscillation that was used to drive the modulator and therefore it also has frequency, f_M. Because of this, only the components in the photodetector output that are at f_M will be mixed down to base band. To counter any drift in the frequency of the local oscillator (b), it is beneficial to add a phase-delay (e) between the oscillator and the mixer that equalizes the effective path lengths of the optical and RF paths. Perhaps more importantly, the phase delay serves to line up the signal with the local oscillator at no phase offset, ensuring that we measure the correct amplitude and the correct sign on that amplitude and canceling the common mode noise from the oscillator in the two signals. A low pass filter (g) is then placed after the mixer to filter out the baseband component. As noted above, this component comes from the beating of the carrier with the sidebands [69]. The sign and magnitude of this DC component is determined by the difference in frequency between the ring and the laser — the magnitude being larger the further from resonance (to a point), and the sign determining the sign of the resonance. The PDH loop has two different ranges of operation: a low modulation frequency range where both the modulation sidebands and the laser line fall within the resonance, and a high frequency range where only the laser frequency falls within the resonance and the modulation sidebands remain outside it. The high frequency operation regime has better noise performance. The

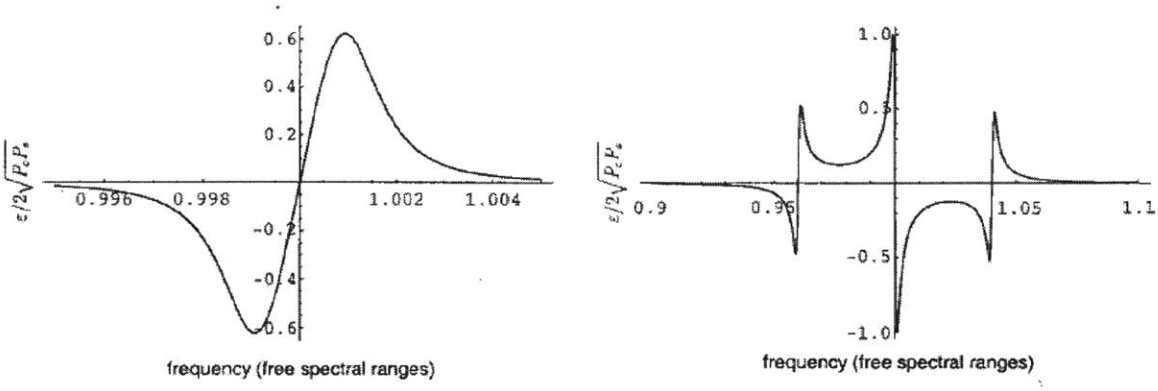

Figure 5-5: PDH error signal in the low frequency (left) and high frequency (right) operating regions. Figure from [69].

error functions of the PDH response are shown in Fig.5-5 [69]. The signal can then be fed back to change the laser operating wavelength. In a free-space cavity this might be accomplished by changing the cavity length; for an integrated laser the frequency could be changed by use of a heating element to change the refractive index of the laser cavity waveguide, and thus the effective cavity length.

PDH Loop Modeling and Design

In order to successfully design an integrated version of the PDH loop, models for the components, including circuitry, as well as models for the system also a whole need to be developed. Having a complete, accurate system model is particularly important for understanding the ultimate noise suppression possible in the system. Details of the components were touched on above. The phase modulator used was based on the phase-shifters created for the MZ modulators and detailed in Chapter 3. The photodetector and ring resonator were discussed in the previous section. We note that a 2GHz FWHM, as measured in the silicon resonator, means that achievable modulation rates in the few GHz range will serve to put us in the lower noise, high frequency operating regime. The intrinsic photodiode bandwidth should eventually allow this as well, though it is questionable at the moment. Because the nitride has lower loss than the silicon (0.7-2.0dB/cm vs. 3-4dB/cm), we expect that FWHM of those resonances to be even narrower and the high frequency range to be even easier

to reach. However, because it is simpler implement in practice (it's easier to align phases, etc...), we begin by considering the low frequency regime. All the necessary components, with the exception of the athermal and nitride rings, were included on an earlier mask and experimental results used for system design. (Test structures for the nitride rings were included on the most recent mask but were not available for system design).

The system performance was simulated using the VerilogA toolkit presented in Chapter 2. Real, imperfect photonic components with the performance outlined in the sections above (photodiode responsivity = 0.8A/W, dark current = 1nA-1μA, modulator $V_\pi L$ = 0.27V/cm, linearity given by that of device in Chapter 3, ring FWHM of 2GHz, etc...) were used in simulations. The CMOS circuitry was designed and modeled by Jonathan Leu using the standard cells provided by the foundry. The output of this simulation is shown in Fig 5-6. One system concern, is the effect of nonlinearities in the phase-shifter response and the driving oscillator output signal. Ripples caused by these nonlinearities are clearly present in error signal, but, as evident, don't affect error signal sign, and have relatively little effect near the error signal zero crossing. Thus, they will not prevent oscillator locking or significantly degrade performance. Noise is another issue; as stated above, limits to system performance will ultimately come from thermal noise in the ring resonator. These can be modeled here. However, noise is also present from the photodiode caused by shot noise and by the noise in the dark current. With low powers or high dark currents, this noise comes to dominate instead (for example, if the dark current is above 1μm and the power on the detector is in the 10μA range, depending on the other system parameters, the dark current may prevent locking).

PDH Loop Layout

Based on the above modeling several versions of the full PDH system were laid out the most recent set of chips. The versions included variations in ring resonators, circuits, and modulator length. They also included versions with and without an integrated laser (as opposed to an off chip laser) to stabilize. Ring variations included coupling

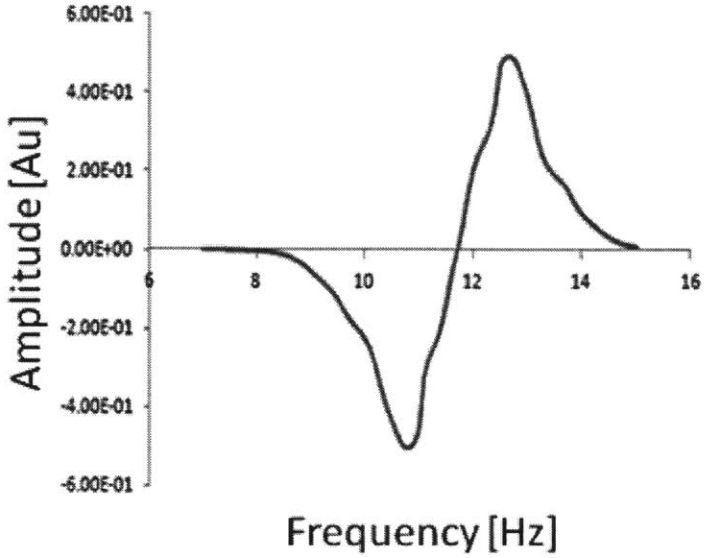

Figure 5-6: Simulation output of Pound-Drever-Hall Loop in the low frequency regime. Results previously shown in Figure 2-14.

gap variations, as well as the three types of rings described above. Both photonics-only versions, designed to work with off-chip circuits, and photonics+CMOS version with 3D integrated circuitry were included. Fig.5-7. shows both the photonic-only and photonic+CMOS layouts of the variations with nitride rings and integrated ring lasers.

5.3 Balanced Detector and Frequency Divider

Integrated versions of the balanced detector loop and the frequency divider loop, where the SSB output is sent to a balanced detector, were also designed and included on the mask. Again both photonics-only and photonics+CMOS versions were included. Variations in the balanced detectors were mainly limited to variations in device size and in driving circuitry to try to truly find the optimal point in the sensitivity vs. speed trade-off. Photodiodes were placed as physically close to each other as possible to minimize fabrication variation and achieve the most balanced response. Variations in the frequency divider loop were severely limited due to mask space,

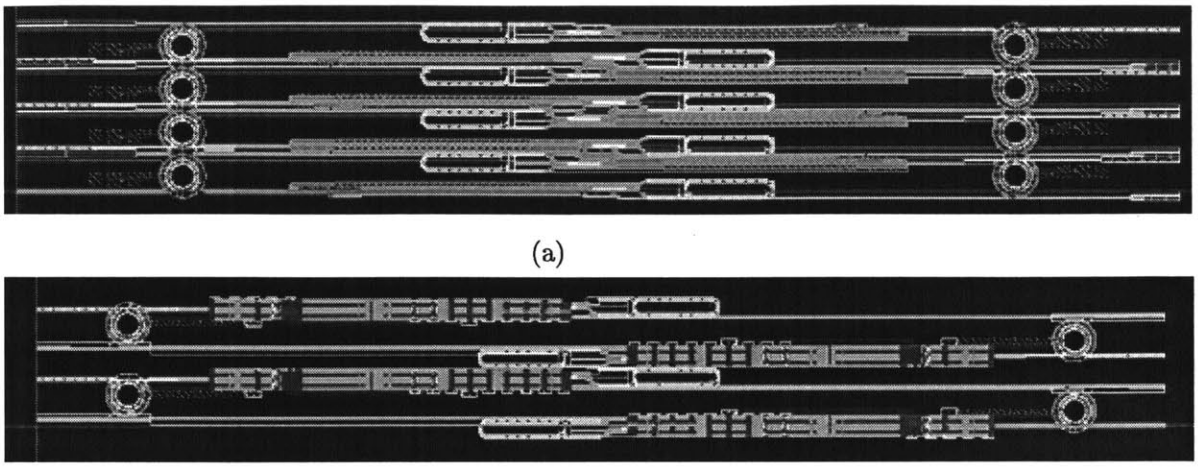

(a)

(b)

Figure 5-7: Cadence layout of PDH photonics-only loop (top, variation with nitride rings used as high-Q reference cavity and integrated ring lasers), and photonics+CMOS loop (bottom, same variation).

and mainly consisted balanced detector photodiode variations on the photonic-only structures, and circuit and photodiode variations on the photonic+CMOS devices. As above, layouts of some of these variations are shown in Fig.5-8. below.

5.4 Summary and Future Work

A first pass of the oscillator subsystems has been designed and layout. Early chips have also been fabricated and returned back from the foundry. Obvious future work consists of testing the performance of these chips as well as of 3D integrated test structures. Improvements in the speed performance of the 3D integrated photodetectors are also necessary; as is careful characterization of the noise of these loops and components, with particular attention paid to the photodetectors and the athermal ring. Further simulations in VerilogA using the outputs of these measurements can then be used to refine the subsystems and create the full oscillator system.

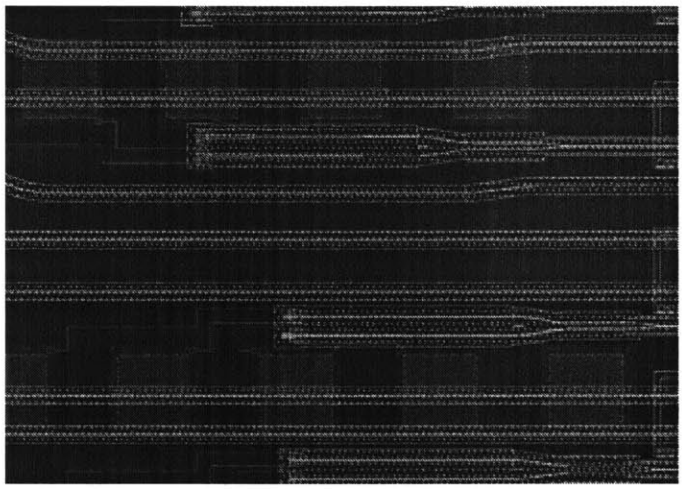

(a)

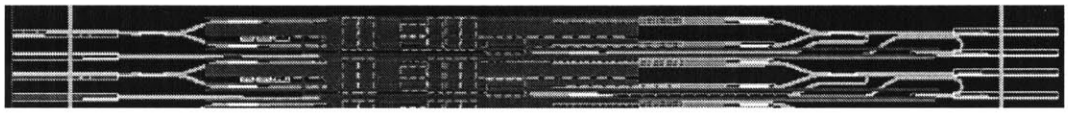

(b)

Figure 5-8: (a.) Zoom in of Cadence layout of balanced detectors. Detectors shown are photonics only. Edge of pads for CMOS integrated balanced detectors which are interwoven are visible at right. (b.) Layout of CMOS integrated frequency divider loop consisting of SSB modulator and balanced detector.

Chapter 6

Summary of Contributions and Parting Thoughts

Silicon photonics truly stands at a tipping point similar to that of the electronics industry fifty years ago. It is on the cusp of moving from moderately complex, custom systems to large, integrated, extremely complex, mass-producible ones. Much work has been done in recent years to bring the industry to this point, particularly in regards to individual component design and digital systems, such as WDM optical links. However, the development of analog devices and systems has lagged behind, as has the integration of light sources. Further, the need for a complete infrastructure to allow for design, layout, high-yield fabrication, and testing of large, complex integrated systems has become acute. Through the lens of an ultra-low noise oscillator, work in this thesis has attempted to address these needs.

In Chapter 2, and overview of a complete silicon photonics platform under development in our group was given. One of the key, new pieces of this work, the development of a VerilogA toolkit to allow for joint system modeling, was presented and verified. The toolkit was used to correctly simulate the responses of simple passive and active devices and systems as well as to model more complex joint photonics+CMOS systems. Once the platform was established, Chapter 3 focused on the development of analog modulators. A linear Mach-Zehnder modulator was presented and shown to have a $V\pi L$ of 0.27Vcm, and a measured input IP3 of 68.4dBm, which is 20dBm

better than the best in the literature. Next, a single sideband modulator was presented and shown to have >15dB suppression of the carrier and unwanted tones. Te device was shown to work with >10dB suppression across many wavelengths from 1550-1590nm, and at a variety of RF drive frequencies.

Chapter 4 tackled the problem of an integrated light source for low-noise analog applications, such as our low-noise microwave oscillator. The necessary components for such a mode-locked laser were designed and tested, and Q-switched mode-locking was observed in a fully integrated laser system. It is suspected that excess normal dispersion lead to the simultaneous presence of Q-switching. Future work aims to lower dispersion. The ability to integrate with active devices also opens up a fascinating possibility for future research directions. Finally, Chapter 5 focused on putting the above together to create fully integrated analog systems including driving circuitry and, in some cases, lasers. These systems were designed and sent to the foundry for fabrication.

There is much work still to be done in the field before commercial products can be developed and sent to market with the same ease as in the CMOS electronics industry — from further development on current projects like those above, to efforts needed in packaging design and automated testing setups, to the availability of reliable multiproect runs — but it's clear that both the will and necessary firm foundations exist. The sheer variety of systems now capable of being designed and built is exciting and astonishing: radar and beam steering applications; sensing and biomolecular manipulations; quantum circuits, logic gates, and photon entanglement devices; nonlinear comb generators; thermal control structures; ADCs; AWGs; advanced signal processing devices; high-speed optical links; and the list could go on. It is truly an exciting time to be working in the field.

Appendix A

Overview of Modulator Approach and Design Process

Here we outline our choices regarding choice of fundamental modulation effect and modulator structure and present a Sentaurus® and MATLAB® model for modeling active integrated devices, particularly a silicon reversed biased diode plasma dispersion effect phase shifter. A longer version of this discussion can be found in [36].

Making a modulator out of silicon involves picking both the physical effect to be used to create modulation and the structure in which it will be implemented. There are many different physical effects that can give rise to modulation. Among the most common are the Pockels effect, Kerr effect, Franz-Keldysh effect, quantum confined Stark effect (QCSE), thermo-optic effect and plasma dispersion effect. The Pockels, Kerr, thermo-optic and plasma dispersion effect are used primarily as mechanisms to change the index of refraction of an optical waveguide, thereby creating a phase modulation. The Franz-Keldysh and Quantum Confined Stark effect are used primarily for absorption based modulators. In general, modulation schemes that rely on changes in the index of refraction, while trying to minimize absorption, are considered to be more useful for analog modulation.

The modulators discussed in this work use the plasma dispersion effect for modulation and the thermo-optic effect for additional bias tuning. The Pockels effect is not present in silicon and the Kerr and Franz-Keldysh effects are very weak, yielding

index changes of <10-4 for fields close to break down. While calculations indicate that these could be useful at high voltages, at the low voltages we desire for CMOS compatibilty they do not yield enough of a change to be helpful. The quantum confined Stark effect is stronger, but requires special fabrication steps. The thermo-optic effect is quite strong, but thermal tuning is also quite slow and thermal modulators are limited to speeds of <1MHz [39]. The plasma dispersion effect is chosen because it has the advantage of being relatively strong (index changes of 110^{-3} are achieved for hole injections of 110^{18}) and relatively high speed (in the GHz range), as well as practically realizable [36].

The plasma-dispersion effect is caused by free carrier injection or extraction. The addition of these carriers changes the index of refraction and absorption of a material. While a theoretical derivation of this effect is possible, starting from the Drude model, the resulting equations are not as accurate as is necessary for good device design. Because of this, the empirical equations derived by Soref [42], are used in the design process instead. These equations are:

$$\Delta n = -8.8e - 22\Delta N_e - 8.5e - 18\Delta N_h^{0.8} \tag{A.1}$$

$$\Delta \alpha = 8.5e - 18\Delta N_e + 6.0e - 18\Delta N_h \tag{A.2}$$

though several other fits to the data have been proposed and used. In this work we a fit to the data set given by a spline extrapolation to of data curves pulled from the original paper is usually used.

A diode was chosen as the underlying electrical device to facilitate the movement of carriers. Reverse biased operation was decided upon because of the high speed at which the carrier distribution can be changed during reverse biased operation as a well as the fact that a revise biased modulator draws very little current, and, thus, dissipates relatively little power. A Mach-Zehnder structure was chosen over a more compact resonant structure due to its significantly larger optical bandwidth which can be on the order of 10s of nm. A rib waveguide structure was chosen to confine the optical mode to the area of index change and away from the lossy, metal contacts

necessary to access to the diode structure. A heater was placed on top of one arm of the Mach-Zehnder structure. Index changes caused by the thermo-optic effect from the heater were used to provide an additional, static offset to ensure that the modulator was biased in quadrature. In addition, both arms of the Mach-Zehnder contained a silicon plasma-dispersion effect modulator and these modulators where biased in push-pull, effectively doubling the sensitivity of the device [92].

It is important to note that the use of a reverse biased diode modulator has some drawbacks. Because the only carrier movement in a reversed biased diode comes from increasing and decreasing the size of the depletion region, overall sensitivity is low meaning that device size is large. Sensitivity is limited by the doping concentration to substantially lower than that in the forward biased case. Sensitivity can be improved by optimizing the overlap of the mode and the area experiencing carrier change. Also, close examination of Soref's equations show that holes contribute a larger index change and a lower optical loss than electrons, so making sure that the center of the optical mode (where most of the light is) sees a change due to holes and not electrons helps as well [93]. We also create a vertical junction in order to help increase the sensitivity per length then process through which this design is done is explained in the paragraph below.

Accurate modeling of the phase shifting sections involves modeling both the electronic carrier responses of the device to an applied bias and the optical responses of the modes to those carrier changes. The modeling in this work consisted of two main parts. First, an electronic structure was designed and the carrier distributions were calculated using Synopsys's Sentaurus™ TCAD software suite. Then, those carrier distributions were imported into MATLAB™ and used to calculate the optical response of the device. Optical simulations for straight and rib waveguides were performed using a mode solver previously developed by Milos Popovic to calculate effective indices and losses. Because results were originally output from Sentaurus™ in a unique data format, a parser was written during this work to convert the output into âĂIJ.matâĂİ format. When Senturus™ updated its visualization software, a new set of codes were written; one in SVisual™ to output a .csv file with the

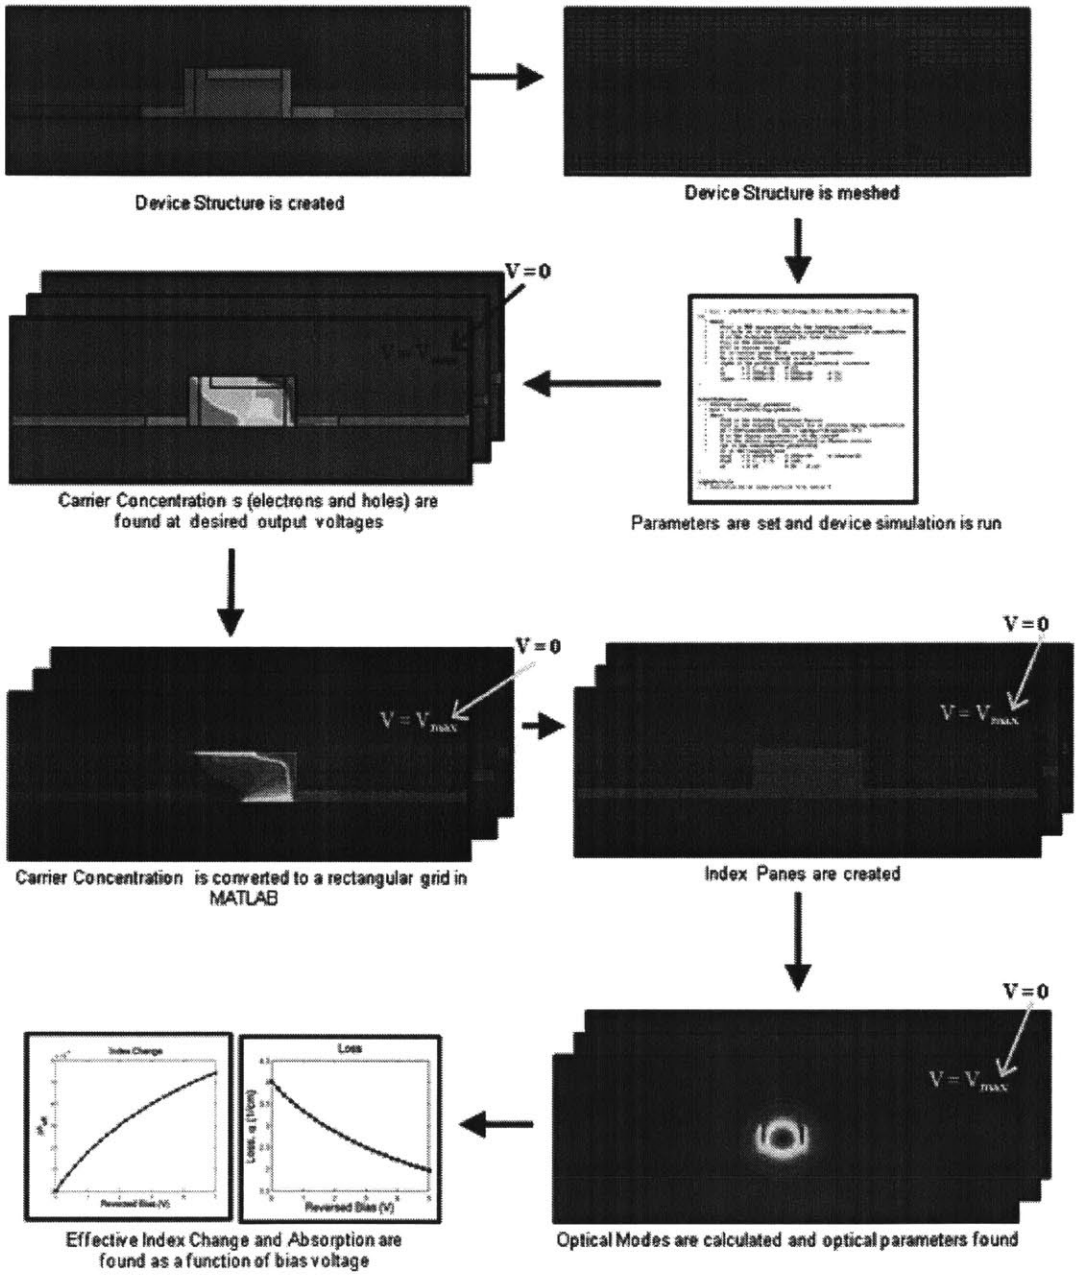

Figure A-1: Diagram of simulation flow for modeling silicon reversed biased diode plasma dispersion effect phase shifters. Figure previously appeared in [[36].

data and another in MATLAB^{TM} to reconstruct a 2D distribution from this .csv file. Additionally, the carrier distributions or other desired parameters were re-gridded from a finite element mesh onto a rectangular finite difference mesh to allow use of Prof. Popovic's modesolver and for ease and speed in later processing. The overall simulation flow for these files is outlined in figure A-1.

Bibliography

[1] A. Khilo, S. Spector, M. Grein, A. Nejadmalayeri, C. Holzwarth, M. Sander, M. Dahlem, M. Peng, M. Geis, N. DiLello, J. Yoon, A. Motamedi, J. Orcutt, J. Wang, C. Sorace-Agaskar, M. Popovi?, J. Sun, G. Zhou, H. Byun, J. Chen, J. Hoyt, H. Smith, R. Ram, M. Perrott, T. Lyszczarz, E. Ippen, and F. KÃd'rtner, "Photonic ADC: overcoming the bottleneck of electronic jitter," Opt. Express 20, 4454-4469 (2012).

[2] Jalali, B., P. Kelkar, and V. Saxena. "Photonic arbitrary waveform generator." Lasers and Electro-Optics Society, 2001. LEOS 2001. The 14th Annual Meeting of the IEEE. Vol. 1. IEEE, 2001.

[3] W. Jiang, F. Soares, S. Seo, J. Baek, N. Fontaine, R. Broeke, J. Cao, J. Yan, K. Okamoto, F. Olsson, S. Lourdudoss, A. Pham, and S. Yoo, "A Monolithic InP-Based Photonic Integrated Circuit for Optical Arbitrary Waveform Generation," in Optical Fiber Communication Conference/National Fiber Optic Engineers Conference, OSA Technical Digest (CD) (Optical Society of America, 2008), paper JThA39.

[4] Pavesi, Lorenzo, and David J. Lockwood, eds. Silicon photonics. Vol. 1. Springer Science & Business Media, 2004.

[5] Sun, Jie, et al. "Large-scale nanophotonic phased array." Nature 493.7431 (2013): 195-199.

[6] Yaacobi, Ami, et al. "Integrated phased array for wide-angle beam steering." Optics letters 39.15 (2014): 4575-4578.

[7] John Keller: http://www.militaryaerospace.com/articles/2012/10/passive-radar-surveillance.html

[8] Google in BGR http://bgr.com/2012/08/07/google-self-driving-car-accidents-300000-miles/

[9] E. Timurdogan, C. Sorace-Agaskar, J. Sun, E.S. Hosseini, A. Biberman, M.R. Watts ., "An ultralow power athermal silicon modulator," Nat. Commun., vol. 5, p. 11,(2014).

[10] Miller, S., Luke, K., Okawachi, Y., Cardenas, J., Gaeta, A. L., and Lipson, M., On-chip frequency comb generation at visible wavelengths via simultaneous second-and third-order optical nonlinearities, Opt. Express, 22, 26517, 17 Nov. 2014.

[11] âĂIJManipulation of multiphoton entanglement in waveguide quantum circuitsâĂİ Jonathan C. F. Matthews, Alberto Politi, AndrÃľ Stefanov & Jeremy L. O'Brien Nature Photonics 3, 346 - 350 (2009) Published online: 24 May 2009

[12] J. Liu, X. Sun, R. Camacho-Aguilera, L. Kimerling, and J. Michel, "Ge-on-Si laser operating at room temperature," Opt. Lett. 35, 679-681 (2010).

[13] J. Van Campenhout, P. Rojo Romeo, P. Regreny, C. Seassal, D. Van Thourhout, S. Verstuyft, L. Di Cioccio, J. Fedeli, C. Lagahe, and R. Baets, "Electrically pumped InP-based microdisk lasers integrated with a nanophotonic silicon-on-insulator waveguide circuit," Opt. Express 15, 6744-6749 (2007).

[14] . Liang, Di, and John E. Bowers. "Recent progress in lasers on silicon." Nature Photonics 4.8 (2010): 511-517.

[15] Orcutt, Jason S., et al. "Open foundry platform for high-performance electronic-photonic integration." Optics express 20.11 (2012): 12222-12232.

[16] M.R. Watts, et. al. âĂIJAn Electronic-Photonic 3D Integration Platform,âĂİ GORMAC, Las Vegas, NV, USA, (2013)

[17] http://www.scienceforbrazil.com/brazil-powers-up-semiconductor-research/

[18] Spector, S.J.; Sorace, C.M.; Geis, M.W.; Grein, M.E.; Yoon, J.U.; Lyszczarz, T.M.; Ippen, E.P.; Kartner, F.X.; , "Operation and Optimization of Silicon-Diode-Based Optical Modulators," Selected Topics in Quantum Electronics, IEEE Journal of , vol.16, no.1, pp.165-172, Jan.-feb. 2010

[19] DeRose, Christopher T., et al. "Ultra compact 45 GHz CMOS compatible Germanium waveguide photodiode with low dark current." Optics express 19.25 (2011): 24897-24904.

[20] K. K. Lee, D. R. Lim, L. C. Kimerling, J. Shin, and F. Cerrina , "Fabrication of ultralow-loss SSiO2 waveguides by roughness reduction," Opt. Lett. 26, 1888-1890 (2001).

[21] Z. Su, M. Cherchi, E. Timurdogan, J. Sun, M. Moresco, G. Leake, D. Coolbaugh, and M. Watts, "Silicon wavelength-selective partial-drop broadcast filter bank," Opt. Lett. 39, 5459-5462 (2014).

[22] E. Timurdogan, E. Hosseini, G. Leake, D. Coolbaugh, and M. Watts, "L-Shaped Resonant Microring (LRM) Filter with Integrated Thermal Tuner," in CLEO: 2013, OSA Technical Digest (online) (Optical Society of America, 2013), paper CTh4F.2.

[23] A. W. Fang, H. Park, 0. Cohen, R. Jones, M. J. Paniccia, and J. E. Bowers, "Electrically pumped hybrid AlGaInAs-silicon evanescent laser," Opt. Express 14, 9203-9210 (2006).

[24] Hosseini, Ehsan Shah, et al. "CMOS-compatible 75 mW erbium-doped distributed feedback laser." Optics letters 39.11 (2014): 3106-3109.

[25] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, âĂIJSilicon optical modulators,âĂİ Nature Photonics 4(8), 518âĂŞ526 (2010).

[26] T. Ramadan, R. Scarmozzino, and R. Osgood, "Adiabatic Couplers: Design Rules and Optimization," J. Lightwave Technol. 16, 277- (1998).

[27] L. Cao, A. Elshaari, A. Aboketaf, and S. Preble, "Adiabatic Couplers in SOI Waveguides," in Conference on Lasers and Electro-Optics 2010, OSA Technical Digest (CD) (Optical Society of America, 2010), paper CThAA2.

[28] Mitomi, O.; Kasaya, K.; Miyazawa, H., "Design of a single-mode tapered waveguide for low-loss chip-to-fiber coupling," Quantum Electronics, IEEE Journal of , vol.30, no.8, pp.1787,1793, Aug 1994

[29] Reed, Graham T., and Andrew P. Knights. Silicon photonics: an introduction. John Wiley & Sons, 2004.

[30] J. P. Chambers; Christi K. Madsen; Haiqiao Lin, âĂIJHigh frequency Pound-Drever-Hall sensing of ring resonator cavitiesâĂİ Fiber Optic Sensors and Applications V. Edited by Udd, Eric. Proceedings of the SPIE, Volume 6770, pp. 67700W (2007).

[31] Doylend, Jonathan K., et al. "Two-dimensional free-space beam steering with an optical phased array on silicon-on-insulator." Optics express 19.22 (2011): 21595-21604.

[32] Opsys Foundry: http://opsisfoundry.org/

[33] PhotonIC Corp: http://www.photonic-corp.com/photonic-services.htm

[34] Europractice-ic: http://www.europractice-ic.com/SiPhotonics_general.php

[35] Lim, A.E.-J.; JunFeng Song; Qing Fang; Chao Li; Xiaoguang Tu; Ning Duan; Kok Kiong Chen; Tern, R.P.-C.; Tsung-Yang Liow, "Review of Silicon Photonics Foundry Efforts," Selected Topics in Quantum Electronics, IEEE Journal of , vol.20, no.4, pp.405,416, July-Aug. 2014

[36] C. Sorace, âĂIJAdvanced Silicon Photonic ModulatorsâĂİ, S.M thesis, Massachusetts Institute of Technology, September 2010.

[37] Watts, M. R., and H. A. Haus. "Integrated mode-evolution-based polarization rotators." Optics letters 30.2 (2005): 138-140.

[38] Gardes, F. Y., et al. "High-speed modulation of a compact silicon ring resonator based on a reverse-biased pn diode." Optics express 17.24 (2009): 21986-21991.

[39] R. Amatya, âĂIJOptimization of tunable silicon compatible microring filters,âĂİ S.M thesis, Massachusetts Institute of Technology, February, 2008.

[40] DeRose, Christopher T., et al. "Silicon microring modulator with integrated heater and temperature sensor for thermal control." Conference on Lasers and Electro-Optics. Optical Society of America, 2010.

[41] F. Gan, High-Speed Silicon Electro-Optic Modulator for Electronic Photonic Intergrated Circuits, Ph.D. thesis, Massachusetts Institute of Technology, June 2007.

[42] Soref, Richard A., and Brian R. Bennett. "Electrooptical effects in silicon." Quantum Electronics, IEEE Journal of 23.1 (1987): 123-129.

[43] Zhou, Gui-rong, et al. "Effect of carrier lifetime on forward-biased silicon Mach-Zehnder modulators." Optics express 16.8 (2008): 5218-5226.

[44] Belt, Michael, and Daniel J. Blumenthal. "Erbium-doped waveguide DBR and DFB laser arrays integrated within an ultra-low-loss Si 3 N 4 platform." Optics express 22.9 (2014): 10655-10660.

[45] Moss, David J., et al. "New CMOS-compatible platforms based on silicon nitride and Hydex for nonlinear optics." Nature Photonics 7.8 (2013): 597-607.

[46] Okawachi, Yoshitomo, et al. "Octave-spanning frequency comb generation in a silicon nitride chip." Optics letters 36.17 (2011): 3398-3400.

[47] Sorace-Agaskar, C., et al. âĂIJIntegrated Mode-Locked Lasers in a CMOS-Compatible Silicon Photonic PlatformâĂİ submitted to CLEO 2015.

[48] Orcutt, Jason S., and Rajeev J. Ram. "Photonic device layout within the foundry CMOS design environment." Institute of Electrical and Electronics Engineers, 2010.

[49] http://en.wikipedia.org/wiki/Process_ design_ kit

[50] J. Morikuni and S.-M. Kang, âĂIJOptoelectronic simulation at the device and circuit level,âĂİ in Integrated Optoelectronics, 1994., Proceedings of IEE/LEOS Summer Topical Meetings, Jul 1994.

[51] T. Smy, M. Freitas, and V. Ambalavanar, âĂIJSelf-consistent opto-thermalelectronic simulation of micro-rings for photonic macrochip integration,âĂİ in Optical Interconnects Conference, 2012 IEEE, pp. 68âĂŞ69, May 2012

[52] A. Lowery, O. Lenzmann, I. Koltchanov, R. Moosburger, R. Freund, A. Richter, S. Georgi, D. Breuer, and H. Hamster, âĂIJMultiple signal representation simulation of photonic devices, systems, and networks,âĂİ Selected Topics in Quantum Electronics, IEEE Journal of, vol. 6, pp. 282âĂŞ296, March 2000.

[53] B. Wang, I. OâĂŹConnor, E. Drouard, and L. Labrak, âĂIJBottom-up verification methodology for cmos photonic linear heterogeneous system,âĂİ in Specification Design Languages (FDL 2010), 2010 Forum on, pp. 1âĂŞ6, Sept 2010.

[54] S. Fathpour and B. Jalali, Silicon Photonics for Telecommunications and Biomedicine (CRC Press, 2012). Chap. 12.

[55] K. Kundert, H. Chang, D. Jefferies, G. Lamant, E. Malavasi, and F. Sendig, âĂIJDesign of mixed-signal systems-on-a-chip,âĂİ Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on, vol. 19, pp. 1561âĂŞ1571, Dec 2000.

[56] K. Zhu, V. Saxena, and W. Kuang, âĂIJCompact Verilog-A modeling of silicon raveling-wave modulator for hybrid CMOS photonic circuit designâĂİ in Proceedings of IEEE 57th International Midwest Symposium on Circuits and Systems (IEEE, 2014), pp. 615-618.

[57] Moss, Ben. âĂIJHigh-Speed Modulation of Resonant CMOS Photonic Modulators in Deep-Submicron Bulk-CMOSâĂİ thesis, Massachusetts Institute of Technology, September 2010.

[58] A.V. Oppenheim, R.W. Schafer, J.R. Buck, Discrete-time Signal Processing (Prentice Hall, 1999). Chap. 2.

[59] J.D. Jackson, Classical Electrodynamics (Wiley, 1998), Chap. 7.

[60] D.J. Griffiths, Introduction to Electrodynamics (Prentice Hall, 1999), Chap. 9.

[61] M. Dasic and M. Popovic, âĂIJMinimum drop-loss design of microphotonic microring-resonator channel add-drop filters,âĂİ in TelecommunicationsForum (TELFOR), 2012 20th, pp. 927âĂŞ930, Nov 2012.

[62] C.Sorace-Agaskar, Massachusetts Institute of Technology, Research Laboratory of Electronics, 50 Vassar St. Cambridge, MA, USA and M.R.Watts and collaborators are preparing a manuscript to be called âĂIJIntegrated Silicon Photonic Single-Sideband Modulator.âĂİ

[63] E. Timurdogan, C. Sorace-Agaskar,J. Sun, E.S. Hosseini, A. Biberman, M.R. Watts, âĂIJAn ultralow power athermal silicon modulatorâĂİ, Nature Communications, vol. 5, 4008 (2014).

[64] S. Shimotsu, S. Oikawa, T. Saitou, N. Mitsugi, K. Kubodera, T. Kawanishi, and M. Izutsu, âĂIJSingle side-band modulation performance of a linbo3 integrated modulator consisting of four-phase modulator waveguides,âĂİ Photonics Technology Letters, IEEE, vol. 13, pp. 364âĂŞ366, April 2001.

[65] C.Sorace-Agaskar, Massachusetts Institute of Technology, Research Laboratory of Electronics, 50 Vassar St. Cambridge, MA, USA and M.R.Watts and collaborators are preparing a manuscript to be called âĂIJIntegrated Silicon Photonic Single-Sideband Modulator.âĂİ

[66] C. Batten, A. Joshi, J. Orcutt, C. Holzwarth, M. Popovic, J. Hoyt, F. Kartner, R. Ram, V. Stojanovic, and K. Asanovic, âĂIJBuilding manycore processor-to-dram networks with monolithic cmos silicon photonics,âĂİMicro, IEEE, vol. PP, no. 99, p. 1, 2009.

[67] D. Miller, "Device Requirements for Optical Interconnects to Silicon Chips," Proc. IEEE 97, 1166 (2009).

[68] M. Georgas, J. Leu, B. Moss, C. Sun, and V. Stojanovic, âĂIJAddressing link-level design tradeoffs for integrated photonic interconnects,âĂİ in Custom Integrated Circuits Conference (CICC), 2011 IEEE, pp. 1âĂŞ8,Sept 2011.

[69] Black, Eric D. "An introduction to PoundâĂŞDreverâĂŞHall laser frequency stabilization." American Journal of Physics 69.1 (2001): 79-87.

[70] C.Sorace-Agaskar, Massachusetts Institute of Technology, Research Laboratory of Electronics, 50 Vassar St. Cambridge, MA, USA and J. Leu and collaborators are preparing a manuscript to be called âĂIJElectro-Optical Co-Simulation for Integrated CMOS Photonic Circuits with Verilog-A.âĂİ

[71] M.Y. Sander; âĂIJHigh Repetition Rate Fiber and Integrated Waveguide Femtosecond LasersâĂİ, Ph.D. thesis, Massachusetts Institute of Technology, June 2012

[72] Mas, Sara, et al. "Tailoring the dispersion behavior of silicon nanophotonic slot waveguides." Optics express 18.20 (2010): 20839-20844.

[73] YÄśldÄśrÄśm, Hasan. "Dispersion parameters and nonlinear optical properties of silicon nitride rib waveguides." Optics Communications 284.7 (2011): 2031-2035.

[74] Tan, D. T. H., et al. "Group velocity dispersion and self phase modulation in silicon nitride waveguides." Applied Physics Letters 96.6 (2010): 061101.

[75] Boggio, J. M., et al. "Group velocity dispersion manipulation in integrated waveguides." Lasers and Electro-Optics Europe (CLEO EUROPE/IQEC), 2013 Conference on and International Quantum Electronics Conference. IEEE, 2013.

[76] Snyder, Allan W., and J. Love. Optical waveguide theory. Vol. 190. Springer Science & Business Media, 1983.

[77] Doran, N. J., and David Wood. "Nonlinear-optical loop mirror." Optics Letters 13.1 (1988): 56-58.

[78] http://www.rp-photonics.com/q_ switched_ mode_l ocking.html

[79] Byun, Hyunil. Integrated high-repetition-rate femtosecond lasers at 1.55 [mu] m. Diss. Massachusetts Institute of Technology, 2010.

[80] Franz X. Kartner, unpublished course notes.

[81] M.Y. Sander; âĂIJHigh Repetition Rate Fiber and Integrated Waveguide Femtosecond LasersâĂİ, Ph.D. thesis, Massachusetts Institute of Technology, June 2012

[82] H. Byun, D. Pudo, S. Frolov, A. Hanjani, J. Shmulovich, E. P. Ippen, and F. X. Kartner, âĂIJIntegrated Low-Jitter 400-MHz femtosecond waveguide laser," IEEE Photonics Technology Letters, vol. 21, no. 12, pp. 763-765, 2009.

[83] H. Byun, D. Pudo, S. Frolov, A. Hanjani, J. Shmulovich, E. P. Ippen, and F. X. Kartner, âĂIJIntegrated 2 GHz femtosecond laser based on a planar Er-doped lightwave circuit," in Lasers and Electro-Optics (CLEO) and Quantum Electronics and Laser Science Conference (QELS), Conference on, 2010.

[84] H. Byun, âĂIJIntegrated high-repetition-rate femtosecond lasers at 1.55Âţm," Ph.D. dissertation, Massachusetts Institute of Technology, 2010.

[85] B.R. Koch, et al. âĂIJMode-locked silicon evanescent lasers,âĂİ Optics Express, vol. 15, no. 18. pp.11225-11233 (2007)

[86] B. Callicoatt, J. Schlager, R. Hickernell, R. Mirin, and N. Sanford, "Compact solidstate waveguide lasers," Circuits and Devices Magazine, IEEE, vol. 19, no. 5, pp. 18 - 27, Sep. 2003.

[87] E. Thoen, E. Koontz, D. Jones, F. Ksrtner, E. Ippen, and L. Kolodziejski, "Erbiumytterbium waveguide laser mode-locked with a semiconductor saturable absorber mirror," Photonics Technology Letters, IEEE, vol. 12, no. 2, pp. 149 -151, Feb. 2000.

[88] G. Della Valle, R. Osellame, G. Galzerano, N. Chiodo, G. Cerullo, P. Laporta, 0. Svelto, U. Morgner, A. G. Rozhin, V. Scardaci, and A. C. Ferrari, "Passive mode locking by carbon nanotubes in a femtosecond laser written waveguide laser," Applied Physics Letters, vol. 89, no. 23, pp. 231115-231115-3, Dec. 2006.

[89] S. J. Beecher, R. R. Thomson, N. D. Psaila, Z. Sun, T. Hasan, A. G. Rozhin, A. C. Ferrari, and A. K. Kar, "320 fs pulse generation from an ultrafast laser inscribed waveguide laser mode-locked by a nanotube saturable absorber," Applied Physics Letters, vol. 97, no. 11, pp. 111114-111 114-3, Sep. 2010.

[90] D. Jones, S. Namiki, D. Barbier, E. Ippen, and H. Haus, "116-fs soliton source based on an Er-Yb codoped waveguide amplifier," Photonics Technology Letters, IEEE, vol. 10, no. 5, pp. 666 -668, May 1998.

[91] Poletti, Francesco, et al. "All-solid highly nonlinear singlemode fibers with a tailored dispersion profile." Optics express 19.1 (2011): 66-80.

[92] A. Khilo, Integrated Optical Analog-to-Digital Converter, SM Thesis, Massachusetts Institute of Technology, February 2008.

[93] Spector, S.J.; Sorace, C.M.; Geis, M.W.; Grein, M.E.; Yoon, J.U.; Lyszczarz, T.M.; Ippen, E.P.; Kartner, F.X.; , "Operation and Optimization of Silicon-Diode-Based Optical Modulators," Selected Topics in Quantum Electronics, IEEE Journal of , vol.16, no.1, pp.165-172, Jan.-feb. 2010

[94] C.Sorace-Agaskar, Massachusetts Institute of Technology, Research Laboratory of Electronics, 50 Vassar St. Cambridge, MA, USA and M.R.Watts and collaborators are preparing a manuscript to be called âĂIJIntegrated Silicon Linear ModulatorâĂİ

[95] B. G. Lee, A. Biberman, J. Chan, and K. Bergman, âĂIJHigh-performance modulators and switches for silicon photonic networks-on-chip,âĂİ IEEE J. Sel. Top. Quantum Electron. 16(1), 6âĂŞ22 (2010).

[96] C.H. Cox III, Analog Optical Links: Theory and Practice. Cambridge University Press, New York, NY. 2006.

[97] Bridges, W.B.; Schaffner, J.H.; , "Distortion in linearized electrooptic modulators," Microwave Theory and Techniques, IEEE Transactions on , vol.43, no.9, pp.2184-2197, Sep 1995

[98] Khilo, Anatol, Cheryl M. Sorace, and Franz X. KÃd'rtner. "Broadband linearized silicon modulator." Optics express 19.5 (2011): 4485-4500.

[99] Vacondio, F.; Mirshafiei, M.; Basak, J.; Ansheng Liu; Ling Liao; Paniccia, M.; Rusch, L.A.; , "A Silicon Modulator Enabling RF Over Fiber for 802.11 OFDM Signals," Selected Topics in Quantum Electronics, IEEE Journal of , vol.16, no.1, pp.141-148, Jan.-feb. 2010

[100] C. Chang, ed., RF photonics technology fiber links (Cambridge University Press, 2002).

[101] Streshinsky, Matthew, et al. "Highly linear silicon traveling wave Mach-Zehnder carrier depletion modulator based on differential drive." Optics express 21.3 (2013): 3818-3825.

[102] Gutierrez, Ana M., et al. "Analytical Model for Calculating the Nonlinear Distortion in Silicon-Based Electro-Optic MachâĂŞZehnder Modulators." Lightwave Technology, Journal of 31.23 (2013): 3603-3613.

[103] X. Zhang, B. Lee, C. Lin, A. Wang, A. Hosseini, and R. T. Chen, "Highly Linear Electro-optic Polymer Based Traveling Wave MMI-fed Directional Coupler Modulator," in CLEO: Science and Innovations, OSA Technical Digest (online) (Optical Society of America, 2012), paper CF1A.6.

[104] N. Hoghooghi, J. Davila-Rodriguez, S. Bhooplapur, and P. Delfyett, "120 dB.Hz2/3 Spur Free Dynamic Range from a Resonant Cavity Interferometric Linear Intensity Modulator," in CLEO: Science and Innovations, OSA Technical Digest (online) (Optical Society of America, 2012), paper CTu2A.2.

[105] P. Wu, S. Deng, and Z. Rena Huang âĂIJCompact SiGe HBT EO Modulator for Analog ApplicationsâĂİ in CLEO: Science and Innovations, OSA Technical Digest (online) (Optical Society of America, 2012), paper JW4A.77.

[106] Cardenas, Jaime, et al. "Linearized silicon modulator based on a ring assisted Mach Zehnder inteferometer." Optics express 21.19 (2013): 22549-22557.

[107] C. C. Sorace, A. Khilo, and F. X. KÃďrtner, âĂIJBroadband linear silicon Mach-Zehnder modulators,âĂİ in Integrated Photonics Research (IPR), Silicon and Nanophotonics, 2010, paper IWA4.

[108] S. Shimotsu et al. âĂIJSingle Side-Band Modulation Performance of a LiNbO3 Integrated Modulator Consisting of Four-Phase Modulator WaveguidesâĂİ, IEEE Photonics Technology Letters, V.13, No. 4, 2001

[109] Fandino, J.S.; Rius, M.; Mora, J.; Mufioz, P.; Capmany, J., "Optical single sideband transmitter using phase modulation and a photonic integrated filter," Microwave Photonics (MWP), 2013 International Topical Meeting on , vol., no., pp.154,157, 28-31 Oct. 2013

[110] Kodama, Y., M. ap Romagnoli, and S. Wabnitz. "Stabilisation of optical solitons by an acousto-optic modulator and filter." Electronics Letters 30.3 (1994): 261-262.

[111] S. Shimotsu et al. âĂIJLiNbO3 Optical Single-Sideband ModulatorâĂİ, OFC, 2000

[112] K. Higuma et al. X-cut lithium noibate optical single-sideband modulatorâĂİ, IEEE Electronics Letters, 2001.

[113] T. Kawanishi et al. âĂIJHigh-speed Control of Lightwave Amplitude, Phase, and Frequency by Use of Electrooptic Effect,âĂİ JSTQE V. 13. No. 1, 2007

[114] Schottstaedt, Bill. "An introduction to FM." Center for Computer Research (2003).

[115] David Cole Massachusetts Institute of Technology, Research Laboratory of Electronics, 50 Vassar St. Cambridge, MA, USA and C. Sorace-Agaskar and collaborators are preparing a manuscript to be called âĂIJIntegrated heterodyne interferometer with on-chip modulators and detectors.âĂİ

[116] Yaacobi, Ami, et al. "Integrated phased array for wide-angle beam steering." Optics letters 39.15 (2014): 4575-4578.

[117] Nelson, C. W., D. A. Howe, and A. Sen Gupta. Ultra-low-noise cavity-stabilized microwave reference oscillator using an air-dielectric resonator. NATIONAL INST OF STANDARDS AND TECHNOLOGY BOULDER CO, 2004.

[118] T. M. Fortier, M. S. Kirchner, F. Quinlan, J. Taylor, J. C. Bergquist, T. Rosenband, N. Lemke, A. Ludlow, Y. Jiang, C. W. Oates & S. A. Diddams Generation of ultrastable microwaves via optical frequency division, Nature Photonics 5, 425âĂŞ429 (2011) doi:10.1038/nphoton.2011.121

[119] Li, Duo. Attosecond timing jitter modelocked lasers and ultralow phase noise photonic microwave oscillators. Diss. Massachusetts Institute of Technology, 2014.

[120] Li, Duo, et al. "Fiber-Optic Demonstration of Optical Frequency Division for Erbium Silicon Photonics Integrated Oscillator." CLEO: Science and Innovations. Optical Society of America, 2014.

[121] Tsarapkin, D.P.; Shtin, N.A., "Sapphire loaded cavity microwave oscillator with improved temperature stability," Frequency Control Symposium and PDA Exhibition, 2002. IEEE International , vol., no., pp.572,579, 2002 doi: 10.1109/FREQ.2002.1075947

[122] New-generation of cryogenic sapphire microwave oscillators for space, metrology, and scientific applications Giordano, Vincent and Grop, Serge and Dubois, BenoÃČÂőt and Bourgeois, Pierre-Yves and KersalÃČÂľ, Yann and Haye, Gregory and Dolgovskiy, Vladimir and Bucalovic, Nikola and Di Domenico, Gianni and Schilt, StÃČÂľphane and Chauvin, Jacques and Valat, David and Rubiola, Enrico, Review of Scientific Instruments, 83, 085113 (2012), DOI:http://dx.doi.org/10.1063/1.4747456

[123] http://www.raytheon.com/news/technology_ today/2014_ i1/eyeontech_ poseidon.html

[124] Green, D.; McNeilage, C.; Searls, J.H., "A Low Phase Noise Microwave Sapphire Loop Oscillator," International Frequency Control Symposium and Exposition, 2006 IEEE , vol., no., pp.852,860, June 2006 doi: 10.1109/FREQ.2006.275501

[125] K. K. Lee, D. R. Lim, L. C. Kimerling, J. Shin, and F. Cerrina , "Fabrication of ultralow-loss SSiO2 waveguides by roughness reduction," Opt. Lett. 26, 1888-1890 (2001).

[126] Alduino, Andrew, et al. "Demonstration of a high speed 4-channel integrated silicon photonics WDM link with hybrid silicon lasers." Integrated Photonics Research, Silicon and Nanophotonics. Optical Society of America, 2010.

[127] New-generation of cryogenic sapphire microwave oscillators for space, metrology, and scientific applications Giordano, Vincent and Grop, Serge and Dubois, BenoÃČÂőt and Bourgeois, Pierre-Yves and KersalÃČÂľ, Yann and Haye, Gregory and Dolgovskiy, Vladimir and Bucalovic, Nikola and Di Domenico, Gianni and Schilt, StÃČÂľphane and Chauvin, Jacques and Valat, David

and Rubiola, Enrico, Review of Scientific Instruments, 83, 085113 (2012), DOI:http://dx.doi.org/10.1063/1.4747456

[128] Valley, George C. "Photonic analog-to-digital converters." Optics Express 15.5 (2007): 1955-1982.

[129] Soltani, M., Lin, J., Forties, R. A., Inman, J. T., Saraf, S. N., Fullbright, R. M., Lipson, M., and Wang, M. D., Nanophotonic trapping for precise manipulation of biomolecular arrays, Nature Nanotech., doi:10.1038/nnano.2014.79, 28 April 2014.

[130] Poseidon Scientific Instruments Pty, Ltd. âĂIJA Mobile Ultra-low Phase Noise Sapphire OscillatorâĂİ Microwave Journal, 2002 [Online] (2002)http://www.microwavejournal.com/articles/3367-a-mobile-ultra-low-phase-noise-sapphire-oscillator

[131] K. Goda, and B. Jalali, âĂIJDispersive Fourier transformation for fast continuous single-shot measurementsâĂİ Nature Photonics, v. 7, 102âĂŞ112 (2013)

[132] Guha, B., Cardenas, J., Lipson, M., Athermal silicon microring resonators with titanium oxide cladding, Opt. Express, 21, 26557, 04 Nov. 2013.

[133] J. Bradley, E. Hosseini, . Purnawirman, Z. Su, T. Adam, G. Leake, D. Coolbaugh, and M. Watts, "Monolithic erbium- and ytterbium-doped microring lasers on silicon chips," Opt. Express 22, 12226-12237 (2014).

[134] Drever, R. W. P., Hall, J. L., Kowalski, F. V., Hough, J., Ford, G. M., Munley, A. J. H. Ward (1983). "Laser phase and frequency stabilization using an optical resonator". Appl Phys B 31 (2): 97.

[135] Khilo, A.; Sorace, C.M.; Birge, J.R.; Kartner, F.X., "Accurate photonic analog-to-digital conversion," General Assembly and Scientific Symposium, 2011 XXXth URSI , vol., no., pp.1,4, 13-20 Aug. 2011