

# UC Santa Barbara

## UC Santa Barbara Electronic Theses and Dissertations

### Title

Wavelength-Selective Photonic Switches for Energy Efficient Reconfigurable Data Center Networks

### Permalink

<https://escholarship.org/uc/item/2bx8b9bw>

### Author

Hirokawa, Takako

### Publication Date

2020

Peer reviewed|Thesis/dissertation

University of California  
Santa Barbara

# Wavelength-Selective Photonic Switches for Energy Efficient Reconfigurable Data Center Networks

A dissertation submitted in partial satisfaction  
of the requirements for the degree

Doctor of Philosophy  
in  
Electrical and Computer Engineering

by

Takako Hirokawa

Committee in charge:

Professor Clint L. Schow, Chair  
Professor John E. Bowers  
Professor James F. Buckwalter  
Professor Adel A. M. Saleh

December 2020

The Dissertation of Takako Hirokawa is approved.

---

Professor John E. Bowers

---

Professor James F. Buckwalter

---

Professor Adel A. M. Saleh

---

Professor Clint L. Schow, Committee Chair

December 2020

Wavelength-Selective Photonic Switches for Energy Efficient Reconfigurable Data  
Center Networks

Copyright © 2020

by

Takako Hirokawa



## Acknowledgements

There is a very large number of people that I would like to thank for helping me get through graduate school. First and foremost, I need to thank my advisor, Clint. I've learned a great many things both technical and non-technical alike from you over the years. Thank you for not only supporting my research endeavors but also encouraging my other non-technical grad school interests and activities. You created an environment in which I was unafraid to be ambitious and set me up for success.

Secondly, I'd like to thank my committee members. I'd like to thank John for guidance on the AIM projects that I worked early on in my graduate school days. Thank you to Jim for patiently explaining circuit concepts to someone who is nominally getting a degree in electrical engineering but doesn't know anything about circuits past Ohm's law. Thank you to Adel, who always made time to explain networking concepts in great detail and with many colors.

I would like to thank Dr. Roger Helkey for providing a lot of input and practical, helpful advice on the AIM projects on which I worked.

Next, I would like to thank my group who include Yujie, Hector, Steven, Aaron, Stephen, JQ, and Xinhong as well as my mentees, Garey, Amalu, Adi, Nano, Shayan, Uriel, and Sean. There's nothing like getting asked a "stupid" question to really make you think about what it is you're actually doing.

I'd like to thank those who weren't in my group but with whom I still worked both on AIM and ARPA-E projects. These people are Akhilesh, Andy, Navid, Mitra, Luis, Sergio, Robert, Fabrizio, Thomas, and Mario. Especially Akhilesh, with whom I was thrown in the deep end with the first AIM tape-out!

As someone who didn't have more senior group members to go talk to, I'd like to thank the following people for their advice when I had trouble navigating the grad school

hoops and experiment woes: Alex, Tin, Minh, Tony, Eric, Paolo, Nicolas, Bowen, and Hongwei. In addition, I'd like to thank the ECE and IEE admin and staff who helped me with all the interstitial things that help move research along especially Val and Amanda.

As I write this, we are in the midst of the coronavirus pandemic, and I am wishing that I was back interacting with members of my office on a regular basis. Thank you to Brandon, Victoria (the OG officemates!), Joe, Michael, Fengqiao, Yujie, Hector, Steven, and Simone for providing much everyday color.

Of course, I have to thank those who have been involved in the Photonics Society over the years. It started as an escape from the lack of progress in my research as well as providing what felt like a tangibly productive outlet. These people include Victoria (again!), Philip, Andy (again!), Demis, Shereen, Caroline, Warren, Eric (again!), Tanya, and many more. Through the Photonics Society, I got to know the amazing staff at the Center for Science and Engineering Partnerships (CSEP), particularly Wendy, who have provided many opportunities in mentoring, outreach, and teaching that complemented my research.

I'd like to thank friends and roommates who haven't been mentioned already: Marissa, Daniel, Bess, Anisa, Kyle, David H., David V., Haw-Tyng, Danilo, and many more as well as my skating coach and friends. And of course, a thank you to my friends from pre-graduate school days for cheering me on from afar.

I'd like to thank my family for providing support and encouragement. I'd like to thank my mom for sending recipes, food, and updates of the dogs. I'd like to thank my dad who encouraged and helped foster my interest in math and science from a young age. I'd like to thank my brother who became my grad school 'twin' (though at a different university) after a childhood of being mistaken for being my biological one.

Last and certainly not least, I'd like to thank Josh. Thank you for being there for me every day no matter how close or how far.

# Curriculum Vitæ

## Takako Hirokawa

### Education

- 2020 Ph.D. in Electrical and Computer Engineering (Expected), University of California, Santa Barbara.
- 2016 M.S. in Electrical and Computer Engineering, University of California, Santa Barbara.
- 2012 B.S. in Engineering Physics, University of Colorado, Boulder.
- 2012 B.S. in Applied Mathematics, University of Colorado, Boulder.

### Publications

1. “Analog Coherent Detection for Energy Efficient Intra-Data Center Links at 200 Gbps per Wavelength,” **T. Hirokawa**, S. Pinna, N. Hosseinzadeh, A. Maharry, H. Andrade, J. Liu, T. Meissner, S. Misak, G. Movaghar, L. Valenzuela, Y. Xia, S. Bhat, F. Gambini, J. Klamkin, A. A. M. Saleh, L. A. Coldren, J. F. Buckwalter, and C. L. Schow, *Journal of Lightwave Technology*, doi: 10.1109/JLT.2020.3029788.
2. “A Wavelength-Selective Multiwavelength Ring-Assisted Mach-Zehnder Switch,” **T. Hirokawa**, M. Saiedi, S. Pillai, A. Nguyen-Le, L. Theogarajan, A. A. M. Saleh, and C. L. Schow, *Journal of Lightwave Technology*, 32, 6292-6298, doi: 10.1109/JLT.2020.3011944.
3. “Analysis and Monolithic Implementation of Differential Transimpedance Amplifiers,” H. Andrade, A. Maharry, **T. Hirokawa**, L. Valenzuela, S. Pinna, C. L. Schow, and J. F. Buckwalter, *Journal of Lightwave Technology*. vol. 38, no. 16, pp. 4409-4418, 15 Aug.15, 2020, doi: 10.1109/JLT.2020.2990107.
4. “Ring-Assisted Mach-Zehnder Interferometer Switch with Multiple Rings Per Switch Element”, **T. Hirokawa**, M. Saiedi, L. Theogarajan, A. A. M. Saleh, C. L. Schow, *Proc. SPIE 11286, Optical Interconnects XX*, 1128612 (28 February 2020); doi: 10.1117/12.2546865
5. “High-Speed Silicon Photonic Optical Interconnects for Cryogenic Readout,” S. B. Estrella, **T. Hirokawa**, A. Maharry, D. S. Renner, C. L. Schow, *Proc. SPIE 11286, Optical Interconnects XX*, 112860B (March 2020)
6. “Comparison of three monolithically integrated TIA topologies for 50 Gb/s OOK and PAM4,” H. Andrade, A. Maharry, **T. Hirokawa**, L. Valenzuela, S. Simon, C. L. Schow, and J. F. Buckwalter, *Proc. SPIE 11286, Optical Interconnects XX*, 112860W (28 February 2020); doi: 10.1117/12.2548762
7. “A Novel Architecture for a Two-Tap Feed-Forward Optical or Electrical Domain Equalizer Using a Differential Element,” A. Maharry, H. Andrade, **T. Hirokawa**, J.F. Buckwalter, and C. L. Schow, *IEEE Photonics Conference (2019)*, San Antonio, TX.

8. "A  $4 \times 4$  Electrooptic Silicon Photonic Switch Fabric with Net Neutral Insertion Loss," N. Dupuis, F. Doany, R. A. Budd, L. Schares, C. W. Baks, D. M. Kuchta, **T. Hirokawa**, and B. G. Lee, *Journal of Lightwave Technology*. doi: 10.1109/JLT.2019.2945678
9. "A Spectrally-Partitioned Crossbar Switch with Three Drops per Cross-point Controlled with a Driver," **T. Hirokawa**, M. Saeidi, A. Maharry, R. Helkey, J. E. Bowers, L. Theogarajan, A. A. M. Saleh, and C. L. Schow, *IEEE Photonics Conference (IPC) 2019*, San Antonio, TX.
10. "A Novel Architecture for a Two-Tap Feed-Forward Optical or Electrical Domain Equalizer using a Differential Element," A. Maharry, H. Andrade, **T. Hirokawa**, J. F. Buckwalter and C. L. Schow, *2019 IEEE Photonics Conference (IPC)*, San Antonio, TX, USA, 2019, pp. 1-2.
11. "Demonstration of a Spectrally-Partitioned  $4 \times 4$  Crossbar Switch with 3 Drops per Cross-point," **T. Hirokawa**, A. Maharry, R. Helkey, J. E. Bowers, A. A. M. Saleh, and Clint L. Schow, *2019 24th OptoElectronics and Communications Conference (OECC) and 2019 International Conference on Photonics in Switching and Computing (PSC)*, Fukuoka, Japan, 2019, pp. 1-3.
12. "Energy Efficiency Analysis of Coherent Links for Datacenters," **T. Hirokawa**, S. Pinna, J. Klamkin, J. F. Buckwalter, and C. L. Schow, *2019 IEEE Optical Interconnects Conference (OI)*, Santa Fe, NM, USA, 2019, pp. 1-2.
13. "High-Speed Optical Interconnect for Cryogenically Cooled Focal Plane Arrays," S. Estrella, D. Renner, **T. Hirokawa**, A. Maharry, M. Dumont, and C. L. Schow, (2019). *GOMACTech 2019*, (18-1, p. 311-314).
14. "Light-based educational outreach activities for pre-university students," K. W. Hamdy, **T. Hirokawa**, P. Chan, W. Jin, V. Rosborough, E. Stanton, A. M. Nether-ton, M. Garza, W. Ibsen, D. D. John, and J. E. Bowers, *Fifteenth Conference on Education and Training in Optics and Photonics: ETOP 2019*, ETOP 2019 Papers (Optical Society of America, 2019), paper 11143071.
15. "A Nonblocking  $4 \times 4$  Mach-Zehnder Switch with Integrated Gain and Nanosecond-Scale Reconfiguration Time," N. Dupuis, F. Doany, R. A. Budd, L. Schares, C. W. Baks, D. M. Kuchta, **T. Hirokawa**, and B. G. Lee, in *Optical Figure Conference (OFC) 2019*, OSA Technical Digest (Optical Society of America, 2019), paper W1E.2.
16. "Monolithically-Integrated 50 Gbps 2pJ/bit Photoreceiver with Cherry-Hooper TIA in 250nm BiCMOS Technology," H. Andrade, **T. Hirokawa**, A. Maharry, A. Rylyakov, C. L. Schow, and J. F. Buckwalter, in *Optical Fiber Communication Conference (OFC)*, 2019, OSA Technical Digest (Optical Society of America, 2019), paper M3A.5.
17. "On-chip Wavelength Locking for Photonics Switches," A. S. P. Khope, **T. Hirokawa**, A. M. Nether-ton, M. Saeidi, Y. Xia, N. Volet, C. L. Schow, R. Helkey, L.

Theogarajan, A. A. M. Saleh, J. E. Bowers, and R. C. Alferness, in *Opt. Lett.*, 42, 4934-4937 (2017).

18. “Elastic WDM Optoelectronic Crossbar switch with On-Chip Wavelength Control,” A. S. P. Khope, A. M. Netherton, **T. Hirokawa**, N. Volet, E. Stanton, C. Schow, R. Helkey, A. A. M. Saleh, J. E. Bowers, and R. C. Alferness, in *Advanced Photonics 2017 (IPR, NOMA, Sensors, Networks, SPPCm, PS)*, OSA Technical Digest (online) (Optical Society of America, 2017), paper PTh1D.3.
19. “Forward bias operation of silicon photonic Mach-Zehnder modulators for RF applications,” R. L. Chao, J. W. Shi, A. Jain, **T. Hirokawa**, A. S. P. Khope, C. L. Schow, J. E. Bowers, R. Helkey, and J. F. Buckwalter, *Opt. Express*, 25, 23181-23190 (2017).

## Abstract

Wavelength-Selective Photonic Switches for Energy Efficient Reconfigurable Data  
Center Networks

by

Takako Hirokawa

Wavelength-selective switches have been proposed for datacenter use to enhance data-center scalability and to aid in meeting ever-increasing traffic demands and the resulting energy consumption. In silicon photonics, photonic integrated circuit (PIC) designers can take advantage of the high contrast between silicon and silicon dioxide—the latter of which acts as the cladding for the silicon nanowire waveguides—to design compact microring resonators with large free spectral ranges (FSRs). Furthermore, as commercial silicon photonics foundry offerings become more widely available, the ability to produce larger and more complicated PICs has become easier, as well as providing a clearer path towards large-scale manufacturability and adoption of such technologies. Thus, ring-based wavelength-selective switches are a particularly well-suited application for silicon photonics. Another major design consideration for PICs is low energy consumption.

A discussion of simulation results from a model for next generation energy efficient photonic links for data centers motivates and the two wavelength-selective switch designs that are presented in this thesis. The first design is an  $N \times N$  crossbar switch with  $L$  ring pairs to route up to  $L$  wavelengths at each cross-point. The second design is an  $N \times N$  ring-assisted Mach-Zehnder interferometer (RAMZI) switch with  $L$  ring pairs per switch element. In both designs, multiple ring pairs of differently sized rings were utilized to partition the FSR such that any one ring pair does not have to move far in the spectrum to complete the switching, thus saving in the power consumption.

# Contents

<b>Curriculum Vitae</b>	<b>vi</b>
<b>Abstract</b>	<b>ix</b>
<b>List of Figures</b>	<b>xii</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Data Center Networks . . . . .	3
1.2 Switches in the Data Center . . . . .	4
1.3 Silicon Photonics . . . . .	6
1.4 Overview . . . . .	8
<b>2 Energy Efficiency Analysis of Analog Coherent Links</b>	<b>9</b>
2.1 Introduction . . . . .	9
2.2 Coherent Link Energy Efficiency Model . . . . .	13
2.3 Results and Discussion . . . . .	21
2.4 Discussion . . . . .	27
2.5 Conclusion . . . . .	32
<b>3 Crossbar Switches</b>	<b>34</b>
3.1 Ring-Based Switch Fundamentals . . . . .	34
3.2 C-Band 4×4 Crossbar Switch . . . . .	43
3.3 O-Band 4×4 Crossbar Switch . . . . .	48
<b>4 RAMZI Switches</b>	<b>56</b>
4.1 Introduction . . . . .	56
4.2 Experimental Methods . . . . .	61
4.3 Scaling Up the Switch . . . . .	68
4.4 Conclusion and Future Work . . . . .	72
<b>5 Summary and Outlook</b>	<b>74</b>
5.1 Future Directions . . . . .	75

<b>A</b>	<b>Test Structures</b>	<b>78</b>
A.1	General Test Structures . . . . .	79
A.2	Device-Specific Test Structures . . . . .	81
A.3	Some Practical Considerations . . . . .	86
<b>B</b>	<b>Crossbar Switch Practical Details</b>	<b>88</b>
B.1	Wirebonding . . . . .	88
B.2	Some Testing Considerations . . . . .	91
B.3	Some Design Considerations . . . . .	91
<b>C</b>	<b>RAMZI Switch Tuning Operation</b>	<b>95</b>
C.1	Wirebonding . . . . .	95
C.2	Switching Operation . . . . .	97
C.3	Some Design Considerations . . . . .	104
<b>D</b>	<b>Code listing</b>	<b>107</b>
D.1	ACD Link Simulation Code . . . . .	107
D.2	RAMZI 4×4 with 2 Ring Pairs Simulation Code . . . . .	126
D.3	RAMZI 2×2 for an Add-Drop or All-Pass Ring Pair . . . . .	133
	<b>Bibliography</b>	<b>137</b>



# List of Figures

1.1	A schematic of a fat-tree data center network architecture, as implemented by Facebook, from Ref. [1] . . . . .	4
2.1	This figure shows the link implementation for a QPSK link, where (a) shows the schematic for the QPSK transmitter (TX) considered in the model, while (b) shows the schematic for the receiver (RX), including the OPLL. Note that the MUX/DEMUX is included in the design for the TX and RX, respectively. . . . .	14
2.2	(a) Electric field and associated output optical power of an MZM and the bias point of the MZM for QPSK modulation. (b) shows the modulation factor vs. the ratio of $V_{sig}/V_{\pi}$ in dB. . . . .	15
2.3	Three-section polarization controller utilizing phase shifters (PS) after the polarization splitter rotator (PSR) with a polarization controller circuit. This scheme would be implemented for each wavelength. . . . .	18
2.4	Simulation results for (a) a 3-mm-long Si Tx/Si Rx TW-MZM and (b) a 1-mm-long InP Tx/Si Rx TW-MZM. The EPB curve (black) and LO power curve (red) correspond to BER = $1 \times 10^{-5}$ , below the KR4-FEC threshold of $2.1 \times 10^{-5}$ . . . . .	23
2.5	The proportion of power taken up by each component in the link at the operating point indicated in Figure 2.4. . . . .	24
2.6	Simulation results when we compare the minimum EPB (black) and drive voltage (red) for TW-MZMs in InP and Si. SEG-MZMs yield similar minimum EPB, and lower drive voltage. . . . .	25
2.7	Starting from the conditions used to generate the results for 3-mm-long Si TW-MZMs shown in Figure 2.4(a), LO power and EPB vs. TX power for BER $10^{-5}$ with a 4 dB link margin and BER $10^{-12}$ with a 13 dB link margin cases are shown. The drive voltage is set to $3 V_{pp-d}$ . LB = link budget. . . . .	27

2.8	Starting from the conditions used to generate the results for 3-mm-long Si TWMMZMs shown in Figure 2.4(a), and taking the operating point where the TX and LO laser powers are equal, this plot shows how the EPB changes for a given drive voltage and the required TX/LO laser powers to close the link. LB = link budget. . . . .	28
2.9	Comparison of unallocated link margin in coherent and IMDD links, assuming MZM drive, and a target BER of $10^{-5}$ . The QPSK curve assumes analog coherent link performance as described in this work, while the other curves assume representative link performance projections for next-generation IMDD and digital coherent links [2]. . . . .	30
2.10	Preliminary hardware for (a) a Si TX modulator and driver, (b) Si RX PIC packaged with an OPLL. . . . .	32
3.1	(a) and (b) shows a perspective view of a strip and ridge waveguide, respectively. (c) and (d) show the intensity distribution of the fundamental mode for the strip and ridge waveguides, respectively. . . . .	35
3.2	This plot shows the effective index corresponding varying widths of waveguide for a given thickness. When the waveguide becomes wide enough, it can support more than one mode. . . . .	37
3.3	(a) shows the schematic of an all-pass ring, while (b) shows a schematic of an add-drop ring. (c) is a schematic of a two serially coupled add-drop ring configuration. . . . .	38
3.4	The responses at the thru and drop ports for an add-drop ring. . . . .	39
3.5	A block diagram of a switch cell in which there are two inputs and two outputs. . . . .	42
3.6	A schematic diagram of a spectrally partitioned $N \times N$ crossbar switch with $M$ wavelengths per port and up to $L$ wavelength drops per cross-point. (a) A high-level block diagram of the switch. (b) Possible realization of an $L$ -MRR cross-point. (c) Details of the spectral partitioning. . . . .	44
3.7	(a) shows the experimental setup with the single-mode optical fiber while (b) shows a close-up of the wirebonded chip. . . . .	46
3.8	(a) shows the tuning spectrum of all three MRRs within each sub-band using SMUs. Each ring is tuned with 0 V, 4 V, 8 V. Figure (b) shows independent tuning of one ring with respect to the other rings. . . . .	47
3.9	(a) shows the tuning spectrum of all three MRRs within each sub-band. Each ring is tuned with 0 V, 4 V, 8 V using a custom driver. Figure (b) shows independent tuning of one ring with respect to the other rings. . . . .	47
3.10	(a) shows the switch wirebonded assembly, while (b) shows a close-up of the switch chip. (c) shows the resistance of the heaters across the switch chip. . . . .	49
3.11	(a) shows switch operation across one full FSR using the heaters, while (b) shows the second MRR tuning independently of the first and third MRRs. . . . .	51

3.12	The expected spectrum of the redesigned 1310 ring switch . . . . .	52
3.13	Description of our wavelength assignment algorithm. (a) A wavelength demand. (b) Resulting wavelength assignment using our algorithm, represented by simple 2-factor graphs. (c) Single-wavelength representation of (b). (d) A wavelength assignment using a standard algorithm, which would have required some of the MRRs to tune over the entire FSR. . . . .	54
4.1	(a) shows single RAMZI cell with two ring pairs, (b) shows a screenshot of the design for one of the rings, and finally, (c) shows an example of the paths taken through the switch. . . . .	57
4.2	The phase change across the resonance of a ring is shown for the over-coupled ( $r > a$ ), critically coupled ( $r = a$ ), and undercoupled ( $r < a$ ) conditions. To generate these plots, $r$ was set to 0.85. . . . .	59
4.3	This figure shows the test bed schematic to test the switch. PC=polarization controller, DUT = device under test, SMU = source measurement units. The red arrows represent the optical path, while the black arrows represent all the electrical connections made in the experimental setup. . . . .	61
4.4	(a) shows the wirebonded switch on a custom PCB, while (b) shows a closeup of the chip on which the switch is located. In (b), the RAMZI switch area is outlined in red, while the area right below the box are the edge couplers. . . . .	62
4.5	The bar state of the both the ring pairs is shown in (a) while the cross state for both ring pairs is shown in (b). The wavelengths being switch are highlighted in blue. . . . .	64
4.6	The outputs of each of the switch states for both ring pairs is shown. The wavelengths are routed from input 1 to (a) output 1, (b) output 2, (c) output 3, and (d) output 4. An inset for each switch configuration is also shown. The data was obtained utilizing SMUs to tune the rings. . . . .	66
4.7	The outputs of each of the switch states for both ring pairs is shown. The wavelengths are routed from input 1 to (a) output 1, (b) output 2, (c) output 3, and (d) output 4. The dark lines represent data taken with the switch controlled with a driver, while the more transparent lines represent data taken for the same switch state with the SMUs. . . . .	67
4.8	The resulting spectra of the driver-controlled RAMZI when controlling the rings from input 1 to two different outputs, ports 2 and 3. . . . .	68
4.9	A $2 \times 2$ RAMZI cell with taps by converting the all-pass rings to add-drop rings. The taps can be routed to a grating coupler to monitor during the initial calibration. . . . .	73
4.10	(a) shows the $2 \times 2$ RAMZI response utilizing an all-pass ring, while (b) shows the RAMZI response utilizing an add-drop ring where the self-coupling coefficient of the second bus waveguide is large. (c) and (d) show the ring response for the respective cases. . . . .	73

A.1	An example of a TLM design. The dark grey squares represent the pads, while the lighter grey represent the doped silicon region. It is often easiest to have one large piece of silicon under all of the pads in the TLM, with highly doped regions right under the contacts. . . . .	79
A.2	An example of a crossing design, including measuring the crosstalk of the crossing. . . . .	82
A.3	A heater test structure design is shown here. There are eight different heater designs, each routed to a pair of pads. . . . .	84
A.4	A diode test structure design is shown here. There is a diode on each arm of the MZI. . . . .	85
B.1	All of the wirebonds made for the crossbar switch mounted on a custom PCB. The large blue-green pad on which the chip is located is a ground plane, while the smaller blue-green squares represent pads on the PCB. .	89
B.2	The first layers of the wirebonds made for the crossbar switch, as denoted by the different colors. In this case, the black lines represent connections from ground pads on the switch to the ground plane on the PCB, while the red lines represent a layer of connections for signal pads. . . . .	89
B.3	The next layers of wirebonds made for the crossbar switch, as denoted by the yellow and orange lines. These layers are above those depicted in figure B.2. . . . .	90
B.4	The next layers of wirebonds made for the crossbar switch, as denoted by the green and pink lines. These layers are above those depicted in figure B.3. . . . .	90
B.5	The final layers of wirebonds made for the crossbar switch, as denoted by the light blue and black lines. These layers are above those depicted in figure B.4. In this case, another ground connection could be made between the chip and PCB. . . . .	91
B.6	A single cell design of the crossbar switch. . . . .	93
B.7	A single cell design of the crossbar switch, but this time in a larger switch. Note that the waveguides are joined with neighboring cells. . . . .	93
C.1	All of the wirebonds made for the RAMZI switch mounted on a custom PCB. The large blue-green pad on which the chip is located is a ground plane, while the smaller blue-green squares represent pads on the PCB. .	96
C.2	The first layers of the wirebonds made for the RAMZI switch, as denoted by the red, orange, and black. In this case, the black lines represent connections from ground pads on the switch to the ground plane on the PCB, while the red lines represent a layer of connections for signal pads. .	96
C.3	The final layers of the wirebonds made for the RAMZI switch, as denoted by the grey and pink lines. These layers are above those depicted in figure C.2. . . . .	97
C.4	Procedure for testing the RAMZI. . . . .	98

C.5	The RAMZI switch with a random phase put on each ring. . . . .	99
C.6	The RAMZI switch with the first stage tuned such that the resonances for the two ring pairs are matched. . . . .	100
C.7	The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are matched. . . . .	100
C.8	The RAMZI switch with the third stage tuned such that the resonances for the two ring pairs are matched and outputs to port 3. . . . .	101
C.9	The RAMZI switch with the third stage tuned such that the resonances for the two ring pairs are slightly detuned from one another and outputs to port 4. . . . .	102
C.10	The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are slightly detuned from one another. . . . .	102
C.11	The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are slightly detuned from one another. . . . .	103
C.12	The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are matched. The resonance of the first ring pair in the second stage is tuned to match the resonance of the first ring pair in second cell in the third stage. The second ring pair in the third stage is redshifted to match the resonance in the second stage. . . . .	104
C.13	The RAMZI switch with the third stage now tuned such that the resonances on each arm are matched and the outputs are directed to port 1. . . . .	105
C.14	The RAMZI switch with the third stage now tuned such that the resonances on each arm are slightly detuned from one another and the outputs are directed to port 2. . . . .	106
C.15	A single cell design of the RAMZI switch. . . . .	106

# Chapter 1

## Introduction

The internet is one of the greatest technological inventions of the modern era, and have given rise to some of the largest corporations today, such as Google, Amazon Web Services, and Facebook. Most of data traffic today is handled within data centers [1, 3]. Many of these large corporations have built and use so-called hyper-scale data centers to deliver their services to customers. Data centers are physical facilities dedicated to housing computer systems and associated infrastructure such as power supplies, cooling systems, and communications equipment, and are located around the world, with many built in the U.S. Data centers are immensely power-hungry facilities due to the large number of servers they must power in addition to the large number of interconnects. Like personal computers, servers generate their own heat which can compromise operation of some components. Thus, data centers must also find ways to cool the servers they house, further adding to the energy consumption. 50% of energy consumed by data centers goes to cooling, while 10% goes to network hardware [4]. Most mid-sized or larger companies have a data center [5], though many companies may choose to co-locate their data center with another company or buy space from another data center company. Data centers remain in operation almost all the time. Planned outages can range from almost

30 hours a year to four hours during a 5-year period depending on how sophisticated the facility is [5].

The energy consumption of data centers is unlikely to decrease in the coming years as demand for more bandwidth from the consumer continues to increase with every passing year [6]. In 2009, data centers consumed 1.5% of the total energy consumed in the U.S. [5]. In 2018, data centers contributed to 0.3% of the overall carbon emissions [7]. A group at Huawei projected that the energy consumption of data centers worldwide could increase from 1% of the global energy consumption in 2010 to between 3% (best case) and 13% (worst case) in 2030 [3]. Furthermore, they project that the energy consumption due to communication technology could contribute as much as 23% of greenhouse gases emissions worldwide by 2030. Within the U.S. as of 2016, the projected growth in electricity consumption between 2015 to 2020 decreased to only 4% from 24% largely due to significantly more energy efficient servers [8]. In addition, some data center companies are working towards utilizing more renewable energies specifically to power their data centers [9, 10]. It is thus of great interest to find more energy-efficient solutions for next generation data centers while being able to deliver on basic tenets of large-scale computing components—namely low cost and power, high reliability and yield, and small size—while being able to continue to increase capacity. While much work is being done to make more reliable computing chips, the energy density within a data center is continuing to increase.

Hyper-scale data centers refer to the ability to continue to scale the data centers to much larger number of servers. They are typically over 400,000 ft<sup>2</sup> and contain advanced cooling systems and redundant power [8]. These data centers arose when large companies such as Google and Amazon needed to deploy hundreds of thousands of servers. Due to the potential energy consumption associated with the size of hyper-scale data centers, data centers have been built to include bare-bones servers with no lights or video. Compared to typical data centers, hyper-scale are able to decrease the ratio of total energy

for everything to computing energy by 40%. This ratio is known as the power usage efficiency and is typically about 2 for a traditional data center and 1.2 for hyper-scale data centers. Hyper-scale data centers made up about 20% of total data center electricity usage in 2018. [7]

## 1.1 Data Center Networks

Hyper-scale data centers must necessarily be connected with a well-defined network architecture that also accommodate the growing number of servers. A typical structure for hyper-scale data centers is a so-called fat-tree or Clos architecture [1, 11] as shown in figure 1.1. There are typically two or three layers of switches above the server racks. Within a server rack, there is the top of rack (ToR) or edge switch, which is connected to each server within the rack and to one or more second layer switch. Second layer or aggregation switches are connected to one or more so-called core or layer 3 switches above [1, 12]. The links below the ToR switches to the servers are typically electrical links, while links above the ToR between the switches are typically optical links and are operated at higher speeds than the links from ToR to servers. Since the second and third layer switches are not necessarily located close to other switches, the links connecting them can be as long as 2 km. Currently, optical links use some sort of intensity-modulated direct detection (IMDD) modulation scheme to transport data between switches, but as the data rate continues to grow, continuing the increase the bandwidth of the IMDD signals will no longer be a viable option from a cost perspective. As such, energy efficient links for higher data rates must be proposed and studied. This is the subject of chapter 2.






Figure 1.1: A schematic of a fat-tree data center network architecture, as implemented by Facebook, from Ref. [1]

## 1.2 Switches in the Data Center

Current switches in the data center convert the optical signal to an electrical signal to perform the switching and converted back to an optical signal, known as OEO conversion. While this possesses important advantages, such as the ability to monitor the health of the links, it is also a very power-hungry implementation. Datacenter switches such as the Mellanox EDR 100 Gb/s Infiniband Smart Switches consume 100s of pJ/bit [13]. Table 1.2 compares the energy efficiency of current commercial data center switches. A common metric to measure the energy efficiency of network components such as transceivers is the

Vendor	Switch Model	Power (W)	Max data rate	pJ/bit	# ports	Ref.
Mellanox	EDR Infiniband Switch S7800	136	100 Gbps	1360	36	[13]
Dell	EMC PowerSwitch Z9332F-ON	900	25.6 Tbps	35	32	[16]
HPE	FlexFabric 5700 48G 4XG 2QSFP+	175	336 Gbps	520	48	[17]
HPE	FlexFabric 5700 40XG 2QSFP+	162	960 Gbps	168	40	[17]
HPE	FlexFabric 5700 32XGT 8WG 2QSFP+	350	960 Gbps	365	32	[17]
Extreme Networks	ExtremeSwitching X465-24W	2485	316 Gbps	7864	24	[18]
Extreme Networks	ExtremeSwitching X465-48T	176	364 Gbps	484	48	[18]
Extreme Networks	ExtremeSwitching X465-48P	1746	364 Gbps	4797	48	[18]
Extreme Networks	ExtremeSwitching X465-48W	4024	364 Gbps	2813	48	[18]
Extreme Networks	ExtremeSwitching X465-24MU	1722	648 Gbps	2657	24	[18]
Extreme Networks	ExtremeSwitching X465-24MU-24W	3941	696 Gbps	5662	24	[18]
Extreme Networks	ExtremeSwitching X465-24S	173	316 Gbps	547	24	[18]
Extreme Networks	ExtremeSwitching X465-24XE	207	888 Gbps	120	24	[18]
Cisco	Nexus 9316D-GX	420	6.4 Tbps	66	16	[19]
Cisco	Nexus 93600CD-GX	590	6.0 Tbps	98	16	[19]
Juniper Networks	QFX5100-48S	150	1.44 Tbps	66	16	[20]
Juniper Networks	QFX5100-48T	335	1.44 Tbps	233	6	[20]
Juniper Networks	QFX5100-24Q	161	2.56 Tbps	63	24	[20]
Juniper Networks	QFX5100-24Q-AA	175	2.56 Tbps	68	24	[20]
Juniper Networks	QFX5100-96S	263	2.56 Tbps	102	8	[20]
Arista	7060CS2-32S	187	6.4 Tbps	29	32	[21]
Arista	7060SX2-48YC6	240	3.6 Tbps	67	6	[21]
Arista	7060CS-32S	187	6.4 Tbps	29	32	[21]
Arista	7260QX-64	315	5.12 Tbps	62	64	[21]
Arista	7260CS-64	1672	12.8 Tbps	130	64	[21]

Table 1.1: Commercial data center switch power, maximum data rate, and energy efficiency.

energy per bit, usually reported in pJ/bit. As can be seen from the table, many current data center switches are not necessarily optimized for energy efficient operation, with energy per bit metrics ranging from almost 30 pJ/bit to over 5 nJ/bit. As bandwidth demands increase, so, too does the power consumption of the switch. Furthermore, the power consumption of the switch application specific ICs (ASICs) are starting to reach the thermal limits for cooling ICs [14, 15].

Replacing some of the electronic switching layers with optical circuit switches can improve the energy efficiency, reduce the amount of OEO conversion (i.e., eliminating transceivers) and enhance the data center performance (by reducing the packet latency throughout the system. Data center network topology also affects the performance and energy efficiency of the whole system; implementing optical switches that allow for flexible

switching could assist in flattening the network topology to reduce energy efficiency overall and enhance performance [22].

## 1.3 Silicon Photonics

Silicon photonics provides an attractive solution to implement next-generation all optical switches. Complementary metal-oxide-semiconductor (CMOS) fabrication technology that enables computer chip manufacturing can be exploited to manufacture a large number of silicon photonic chips for a relatively low cost. In addition to the mass manufacturing capabilities of CMOS, there is high index contrast between silicon and silicon dioxide. Therefore photonic components and photonic integrated circuits (PICs) on chip can be made compact due to the tight bends possible from the high confinement of the light in the silicon waveguides. Despite the advantages that silicon holds over many other materials, there are notable downsides. The primary disadvantage is that silicon possesses an indirect bandgap; thus, it cannot be used to generate light without significant phononic interaction—unlike other semiconductor materials that possess a direct bandgap like indium phosphide (InP) materials—making silicon a very poor material with which to fabricate a laser. Much work is being done to integrate III-V semiconductor materials, most notably InP, onto silicon in a standard silicon photonics process. Furthermore, silicon photonics has a comparably weaker electro-optic effect, making it a less efficient material than other III-V materials to be used in active devices such as modulators. In addition, significant challenges remain in packaging silicon PICs into packages. While the electronic packaging is not necessarily an issue due to the large electronic IC ecosystem, the optical packaging still presents a large hurdle towards mass adoption of photonic technologies. Typically, the light on chip is confined in a waveguide that is several hundred nanometers wide and no more than a couple hundred nanometers

tall. Meanwhile, the mode of light at telecom wavelengths in single mode fiber is around  $10\ \mu\text{m}$  in diameter [23]. Simply butt-coupling a waveguide to a single mode fiber results in significant loss due to the large mode mismatch. It is possible to make the waveguides larger. In fact, this has been proposed and demonstrated by Rockley Photonics [24] at the expense of the ability to make tight bends and therefore very compact photonic devices. Another notable effort towards packaging silicon photonic parts comes from IBM with the use of polymer waveguides [25, 26]. When silicon became the material of choice for mass manufacturing of electronics, this spurred research interest in silicon as a material for photonic devices in the mid to late 1980s. There had previously been research efforts in other materials such as the aforementioned InP, as well as gallium arsenide (GaAs) and lithium niobate ( $\text{LiNbO}_3$ ) [27]. The interest in silicon photonics was driven primarily by the desire to integrate the photonics with the electronics [27, 28]. It was discovered that silicon is transparent to wavelengths above  $1.2\ \mu\text{m}$ , including around  $1.3\ \mu\text{m}$  and  $1.5\ \mu\text{m}$ —wavelengths typically used for optical communications [28], motivating optical communications as an application for silicon photonics. Despite research activity beginning in the 1980s, it wasn't until 2001 that Luxtera was able to demonstrate the first commercial silicon photonic product [29, 30] manufactured in a commercial CMOS process. Today, there are currently several foundries that provide silicon photonic processes, often through a multi-project wafer (MPW) run. Table 1.3 provides a list of some of the current silicon photonics foundries. MPW runs are crucial for smaller organizations that cannot afford a full fabrication run, let alone maintain a manufacturing facility to be able to design silicon photonic devices, paving the way towards more so-called fab-less operation [29, 30]. In this arrangement, the foundry provides access to the process to the customers who can their desired devices as long as they adhere to the design rules of the process. All of the silicon PICs presented in this thesis were designed in MPW runs.

Foundry	Location	Silicon Thickness	Wafer Size	Ref.
GlobalFoundries	New York	–	300 mm	[31]
AIM Photonics	New York	220 nm	300 mm	[32]
TowerJazz	California	–	200 mm	[33]
IHP Microelectronics	Germany	220 nm	200 mm	[33, 34]
AMF	Singapore	–	200 mm	[35]
imec	Belgium	220 nm	200 mm, 300 mm	[36]

Table 1.2: Select listing of commercial silicon photonic foundries that offer MPW services.

## 1.4 Overview

The rest of the thesis is organized as thus. The second chapter presents an energy efficiency model for next generation data center links. These links assume a link budget of 13 dB between transmitter and receiver, which can be used by a photonic switch. In Chapter 3, silicon photonic crossbar ring-based wavelength-selective switches are presented. These ring-based crossbar switches are relatively easy implementation of wavelength-selective switches but do not provide a good path towards large-scale wavelength-selective photonic switches. Thus in Chapter 4, an alternative approach to a wavelength-selective switch in silicon using ring-assisted Mach-Zehnder interferometers (RAMZIs) are proposed and demonstrated. I conclude and propose further research directions in Chapter 5. The appendices are organized as thus: Appendix A contains brief descriptions of test structures, Appendix B and C includes practical details when testing crossbar switches and RAMZI switches, including the wirebonding diagrams and tuning procedures, and Appendix D contains code listing for simulations presented throughout the thesis.

# Chapter 2

## Energy Efficiency Analysis of Analog Coherent Links

Work presented in this section has appeared in [37] and [38].

### 2.1 Introduction

With ever-increasing demand for cloud services, evaluating interconnect technology benefits and tradeoffs anticipates future deployments of the data center through scaling baud rates, higher order modulation formats with more bits/symbol, polarization multiplexing, and adding additional wavelength division multiplexed (WDM) channels. Current data center links rely on intensity-modulated direct detection (IMDD) schemes due to their relative simplicity and correspondingly relatively low cost and power consumption. However, scaling IMDD links to 200 Gbps/lane will require a large jump in complexity and power consumption. A recent study showed the potential of a 100 GBd PAM-4 link to operate over a 400 m link distance [39]. However, heavy equalization was required, with 71 feedforward equalizer (FFE) taps and 15 decision feedback equalizer

(DFE) taps, just to achieve a pre-FEC (Forward Error Correction) bit error ratio (BER) slightly below the soft decision (SD-FEC) limit of  $2 \times 10^{-2}$ . With such power-hungry equalization, the required received optical power was  $> +7$  dBm, likely demanding an unfeasible output power from the transmitter (TX) source laser [39]. The limited prospects for scaling IMDD links to 200 Gbps/lane and beyond have driven substantial interest in developing a new generation of energy-efficient coherent links designed specifically for intra-datacenter applications [2, 37, 40, 41]. A recent paper by authors from the Alibaba Group presents a detailed comparison of several variants of IMDD (PAM4, CAP16, DMT) against digital coherent (PDM-16QAM) for 400G links, backed up with experimental results, using metrics of minimizing laser and ASIC power consumption [41]. The authors conclude that coherent links have lower laser power requirements and comparable ASIC power dissipation and digital signal processing (DSP) complexity compared to the IMDD approaches. Recent work from Google provides a comparison up to 1.6 Tb/s, analyzing in detail multiple digital coherent (16, 32, 64QAM) and IMDD (PAM4, 6, 8) architectures [2]. The coherent links are projected to consume somewhat more power—on the order of 10-20%—but offer substantial advantages: greater tolerance to fiber impairments, higher spectral efficiency, and a large advantage in receiver (RX) sensitivity. For modulator drive swings less than 1V?, the gains in RX sensitivity are found to be mostly offset by large modulator losses and the PAM links are projected to achieve larger link budgets. The coherent links operate at 2X higher total bit rates, and with higher modulator drive voltages achieve 5-9 dB more link budget than the IMDD variants [2].

Digital coherent architectures commonly used in telecom interconnects are implemented with a free-running local oscillator (LO) which requires an RX chain consisting of a linear receiver front end followed by an analog-to-digital converter (ADC) to digitize incoming data. Doing so enables the DSP to perform functions such as carrier recovery,

polarization demultiplexing and channel equalization to remove fiber propagation impairments such as chromatic dispersion (CD) and polarization mode dispersion (PMD). An alternative approach to coherent detection is analog coherent detection (ACD) which utilizes a highly integrated optical phase-locked loop (OPLL) to directly lock the frequency and phase of the LO laser to an incoming wavelength channel. Chip-scale integration enables low feedback loop delay and therefore high loop bandwidth, enabling the use of more easily integrated tunable LO lasers with MHz-scale linewidth [42–46]. Furthermore, the OPLL approach provides for the direct demodulation of complex signals at low uncorrected bit error rates (BER), with previous proof-of-concept demonstrations achieving  $\text{BER} < 10^{-12}$  for BPSK modulation up to 35 Gb/s [45].

Although latency may not be especially critical for our primary target application of intra-datacenter links where the use of FEC is ubiquitous, the potential to construct FEC-free coherent links offers a substantial advantage for highly latency-sensitive applications such as high-performance computing (HPC). Another key benefit of OPLL-based coherent detection in general, and offered by our OPLL ACD architecture is inherent wavelength selectivity. When the LO is locked to an incoming wavelength channel, other channels are rejected by the RX. For example, if the system channel spacing is 200 GHz, when the LO is locked to one of the wavelength channels in the incoming optical signal, the locked signal is down-converted to the baseband while the other wavelength channels are converted to 200 GHz or higher—far above the operating bandwidth of the receiver electronics. This wavelength selectivity can be exploited to reduce crosstalk requirements for future networks that incorporate photonic routing/switching and eases channel crosstalk requirements of on-chip wavelength multiplexing/demultiplexing components.

It is widely accepted that much of the complexity of traditional coherent DSP can be removed for datacenter applications [2, 41] where O-band operation of links up to 2 km present negligible fiber impairments. Consequently, the biggest power savings offered by



ACD arises not through the elimination of DSP, but through the removal of linear RX frontends and ADCs. QPSK as a modulation format uniquely takes full advantage of the direct demodulation capability enabled by ACD. At the output of the  $90^\circ$  hybrids in an ACD receiver, the I and Q channels have been separated and low-power electronics using limiting amplifiers can be used to make a binary decision, just like in the most power efficient non-return to zero (NRZ) on-off keying (OOK) links [47]. State-of-the-art ADCs have been developed with sufficient sampling rate and effective number of bits (ENOB) for 224 Gbps DP-16QAM coherent receivers with power consumption ranging between 235 mW [48] to 702 mW [49]. A dual-polarization I-Q receiver would require four such ADCs, resulting in a total ADC power consumption between 940-2808 mW or 4.2-12.5 pJ/bit based upon the efficiencies reported in [48, 49]. Our QPSK link architecture does not require these power-hungry components and full-link energy efficiencies of less than 5 pJ/bit are feasible. The substantial power savings advantage for QPSK does not straightforwardly scale to higher order QAM formats which require multiple decision thresholds for both I and Q channels, driving the need for A/D conversion.

In this paper we present a multi-wavelength analog coherent detection (ACD) architecture utilizing a chip-scale OPLL and based on 50 GBd polarization-multiplexed QPSK (PM-QPSK) for an aggregate data rate of 200 Gbps/ $\lambda$ . In addition to the link-level advantages in optical budget and power efficiency offered by QPSK-based ACD, we believe it will be advantageous to scale to bit rates of 800 Gb/s and beyond by using four or more WDM lanes, each carrying 200 Gb/s, as opposed to fewer lanes at higher per- $\lambda$  bit rates. The large optical loss budget enabled by ACD further opens a wider space for network architecture designs offering greater flexibility and scalability through the insertion of optical wavelength-level routing and/or circuit switching devices in the data center network. Keeping the per- $\lambda$  bandwidth granularity lower expands opportunities for network architectures with substantial power savings and enhanced operational

flexibility as discussed in Section 2.4.2.

## 2.2 Coherent Link Energy Efficiency Model

In this section, we present an ACD link model that supports a quantitative exploration of the design space of modulator length, drive voltage, and TX source and LO laser powers. For ease of reference, we refer to the operating baud rates as 50 GBd, but all simulations are conducted at 56 GBd to allow for coding and FEC overhead. Furthermore, although our link architecture is capable of operating at uncorrected BER  $< 10^{-12}$ , we assume a target BER of  $1 \times 10^{-5}$ , compatible with KR4-FEC (BER threshold =  $2.1 \times 10^{-5}$ ), and the KP4-FEC (BER threshold =  $2.2 \times 10^{-4}$ ) that is widely implemented in data center network switches [50]. The ACD link model consists of a quadrature phase-shift keying (QPSK) transmitter, a low-loss optical link ( $< 2$  km), and a homodyne coherent receiver. Figure 2.1 illustrates a schematic of the dual-polarization QPSK (DP-QPSK) ACD link for the Si-based architectures. In Figure 2.1, the transmitter (TX) laser light is split into two single-mode waveguides and modulated with IQ modulators. We consider two modulator architectures and two photonic integrated circuit (PIC) platforms. The first modulator is based on a traveling wave modulator (TW-MZM) design [51], while the second utilizes a segmented modulator (SEG-MZM) [52]. Both modulators have been demonstrated in Si and InP platforms [53–62]. We find that the choice of TX architecture and PIC technology both have a significant impact on the overall link performance and power budget. Previous comparisons of SEG-MZMs and TW-MZMs have been made for high-bandwidth radio over fiber (RoF) photonic systems [60–62]. RoF links based on SEG-MZM generally showed improvements in gain and noise figure over TW-MZM implementations at high frequency, but at the expense of higher power consumption.

Considering a differential driving signal, a phase modulated signal is realized when

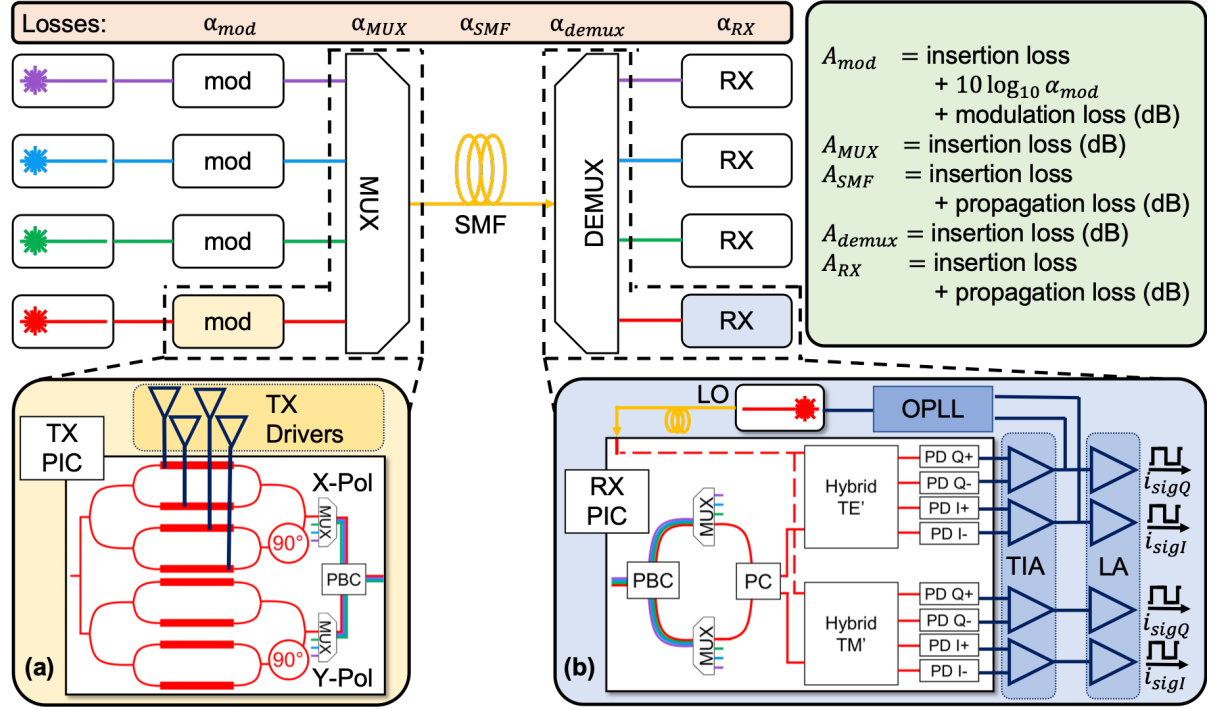


Figure 2.1: This figure shows the link implementation for a QPSK link, where (a) shows the schematic for the QPSK transmitter (TX) considered in the model, while (b) shows the schematic for the receiver (RX), including the OPLL. Note that the MUX/DEMUX is included in the design for the TX and RX, respectively.

an MZM is biased at its null point, where the electric field transmission is 0, as shown in Figure 2.2(a). At this bias point, the optical carrier undergoes a  $180^\circ$  phase shift when the input signal transitions from the logic 0 to 1 and vice versa even when the voltage swing is less than twice the full modulator half-wave voltage ( $V_\pi$ ) of the MZM. However, driving the modulator with a signal amplitude smaller than  $2V_\pi$  leads to increased loss. Such loss can be estimated using the modulation factor ( $F_M$ ), defined in linear units by:

$$F_M = \frac{1}{2} \left( 1 - \cos \left( \pi \frac{V_{sig}}{2V_\pi} \right) \right), \quad (2.1)$$

where  $V_{sig}$  is the peak-to-peak drive voltage. The optical loss due to the modulation factor with respect to the drive voltage is shown in Figure 2.2(b). A  $F_M$  of 1, which corresponds to a  $2V_\pi$  drive voltage swing, leads to no modulation-induced loss, while

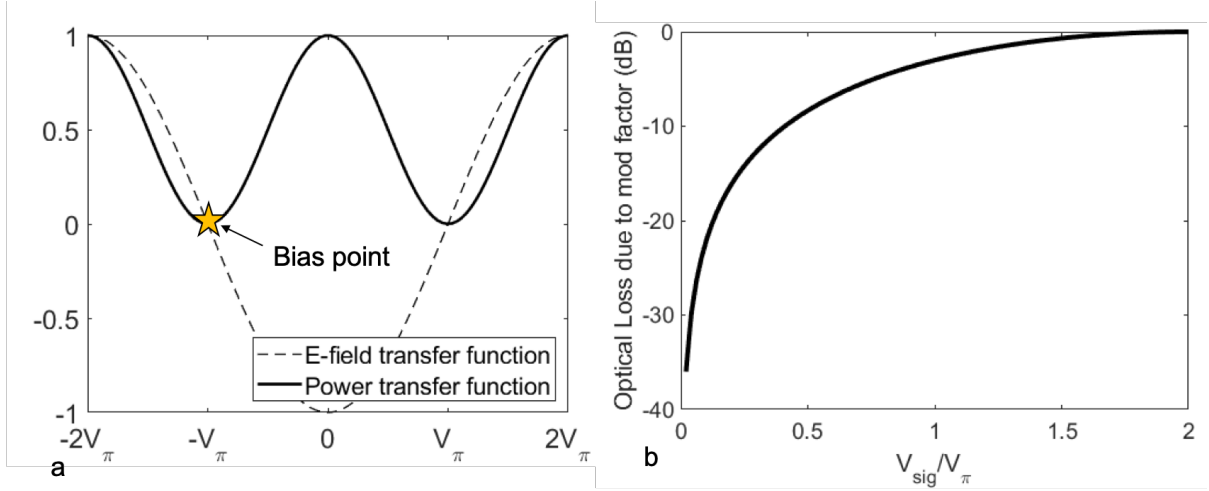


Figure 2.2: (a) Electric field and associated output optical power of an MZM and the bias point of the MZM for QPSK modulation. (b) shows the modulation factor vs. the ratio of  $V_{sig}/V_\pi$  in dB.

$F_M$  of 0.5 corresponds to a  $V_\pi$  voltage swing and 3 dB of induced optical loss. The modulation factor therefore presents a fundamental power consumption tradeoff: larger drive voltages reduce modulation loss at the expense of higher power dissipation for the electrical modulator driver circuits. Conversely, lower drive voltages reduce driver power but increase optical losses that need to be compensated by higher source and/or LO laser power levels. The modulation loss is therefore controlled by the drive voltage amplitude and is independent of MZM insertion loss. The MZM length is also a key parameter that trades off optical propagation losses against electrical power dissipation in the driver circuits. The relationship between MZM length and optical propagation loss is given by:

$$P = P_{in} e^{-\alpha_{opt} L_{MZM}}, \quad (2.2)$$

where  $P_{in}$  is the input power,  $L_{MZM}$  is the active length of the modulator, and  $\alpha_{opt}$  is the loss per length of the active region. In InP, we measured this to be 4.34 dB/mm, while in Si we used a value of 1.5 dB/mm. Modulator  $V_\pi$  inversely depends on modulator length:

longer modulators exhibit lower  $V_\pi$  and require relatively lower drive voltages at the expense of higher optical propagation losses; shorter modulators have lower optical losses but higher  $V_\pi$  with correspondingly higher drive voltage requirements and accompanying driver power consumption.

We assume differential drive for both SEG-MZM and TW-MZM transmitters and incorporate polarization multiplexing to increase the link capacity by a factor of two.  $V_{sig}$  was swept from 0 to  $V_\pi$ . In the TW-MZMs, we assume that the microwave and optical velocities are matched such that the phase shift induced by the traveling wave electrodes is integrated along the length of the MZM. If the optical and microwave velocities are not perfectly matched, there is a well-known degradation in bandwidth that would introduce an additional inter-symbol interference (ISI) power penalty. In the TW-MZMs, we account for electrode loss along the length of the modulator, which we estimated from simulations of traveling wave electrodes designed for operation above 50 GHz. The electrode loss was estimated to be 0.2 Np/mm and 0.4 Np/mm from simulations in Si and InP, respectively. For the SEG-MZMs, we assume that the driver accounts for the time delay between phase shifter sections and that the voltage delivered to each segment is the same for all segments.

For both driver power calculations, we consider only the power dissipation in the output stage. In both calculations, we assume a 45 nm CMOS technology and that  $\eta_{dr}$  is the efficiency of the driver. For the TW-MZM driver power consumption, we assume differential drive, current-mode logic operation, and that  $V_{sig}$  is the single-ended peak output voltage of the driver. The single-ended voltage swing delivered to the MZM is also dependent on the characteristic impedance of the transmission line ( $Z_0$ ) which was set to 40  $\Omega$  for Si and 30 $\Omega$  for InP. 40  $\Omega$  was chosen for Si based on typical values found in the literature [63–66], while 30  $\Omega$  for InP was measured from initial designs. Changing the TW-MZM impedance to say, 50  $\Omega$ , would not change the bandwidth or the phase

efficiency in the calculations, as this relationship is not captured by the model. The power consumption for the SEG-MZM is dependent on the length ( $l_{seg}$ ) and capacitance ( $C_{seg}$ ) of each segment, the number of segments being driven, and the baud rate ( $R_b$ ). Like the TW-MZM driver calculation,  $V_{sig}$  is the single-ended peak output voltage of the driver. To calculate the number of segments, we first assumed that the length of each segment was  $200 \mu\text{m}$  to ensure they would behave as lumped circuit elements under 50 GBd operation. The active length of the modulator was then divided by the segment length and rounded up to give an integer number of segments.  $C_{seg}$  was measured in initial test structures to be  $0.27 \text{ fF/mm}$  and  $0.94 \text{ fF/mm}$  for Si and InP, respectively. We define the driver efficiency as a ratio of the capacitance of each segment to the sum of the segment capacitance and the output capacitance of the driver. It is given by

$$\eta_{dr,SEG} = 1 - G_C \frac{V_{sig}}{t_r}, \quad (2.3)$$

where  $G_C$  is the ratio of output capacitance to drain current for a given transistor, and  $t_r$  is the signal rise time. In other words, for a given drive voltage and rise time, the efficiency is set by the physical transistor parameters. We found that decreasing the process node to 22 nm, did not yield a significantly higher efficiency due to only slight changes between the processes in  $G_C$ . A 45 nm CMOS process, consistent with our calculations and that will feature full monolithic integration with high-performance Si photonic devices is currently under development [67].

After propagating through up to 2 km of SMF, the WDM signal is coupled into the RX as shown in Figure 2.1. While propagating through the fiber, the light undergoes random polarization rotation, necessitating polarization recovery in the RX PIC. In all of our analysis presented here, we assume a Si photonic implementation of the coherent RX PIC, as low-loss, on-chip polarization de-multiplexing can be much more readily realized

compared to monolithic InP platforms. When light enters the RX, it first goes through a polarization splitter rotator (PSR) that separates incoming TE and TM components and rotates the TM component to TE for propagation through the on-chip waveguides that natively only support low-loss propagation of TE polarized light. After the PSR, we refer to the two propagating polarizations as TE' and TM', the latter of which has been rotated to TE. Both TE' and TM' contain a mix of the original transmitted polarizations and must be further processed to recover and separate the original X- and Y-polarizations modulated at the TX.

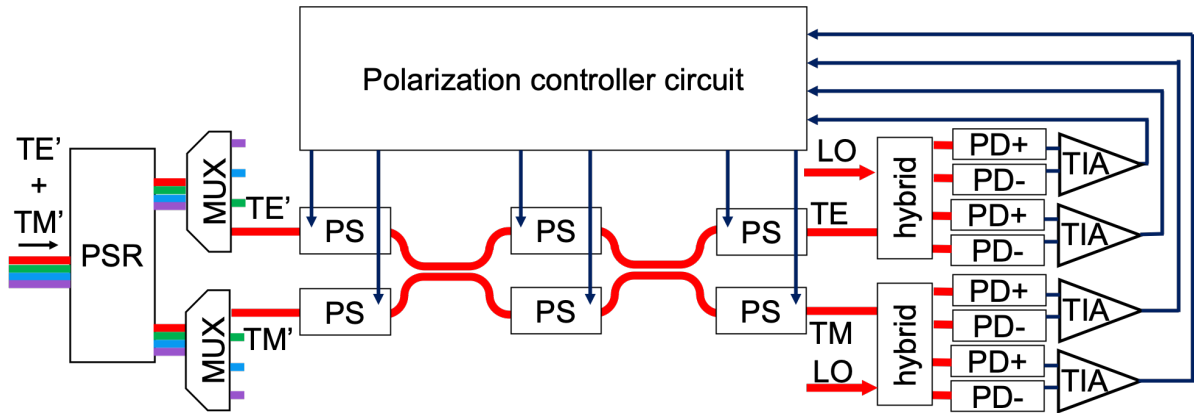


Figure 2.3: Three-section polarization controller utilizing phase shifters (PS) after the polarization splitter rotator (PSR) with a polarization controller circuit. This scheme would be implemented for each wavelength.

The TE' and TM' signals next go through separate wavelength demultiplexers that separate individual wavelengths. For each wavelength, the TE' and TM' signals are processed by a polarization controller. Since polarization recovery is conducted for each wavelength, future networks incorporating wavelength routing or switching are readily supported—it doesn't matter if each wavelength entering the receiver has traversed a different path through the data center network. We have implemented a polarization controller using a three-stage device, described in greater detail in [40], and shown schematically in Figure 2.3. Six thermo-optic phase shifters can be configured to fully

separate the original X- and Y-polarizations from the received TE' and TM' signals. After polarization recovery, the X and Y signals are sent to separate 90° hybrids. The polarization recovery scheme exploits a low frequency (few MHz) pilot tone impressed on the I-component of the X-polarization at the TX. In the RX, the low frequency pilot tone is separated from the information-bearing signal by means of a low-pass filter and fed to a low-speed microcontroller. A feedback loop tunes the six phase shifters in the polarization controller, minimizing the pilot tone for all the 90° hybrid output channels except the I-component of the X-polarization. The thermo-optic phase shifters have a response time on the order of tens of microseconds, sufficient to track polarization variations under nominal operating conditions, and can be controlled by a low-cost and ultra-low power microcontroller. We have previously demonstrated a multi-channel thermal phase shifter driver in [68]. The printed circuit board (PCB) implementation consumed a total of 400 mW in the unloaded configuration with an additional average 20 mW per thermal phase shifter connected to the driver. However, to be adapted for the polarization controller described here, the total channels would be reduced from 96 channels to six, with a total of three used at any time, thus significantly reducing the total power consumption of the circuit to be less than 100 mW.

The receiver includes an integrated LO laser for each wavelength channel that is split and then mixed with the incoming signal in separate 90° optical hybrids for the X and Y polarizations. Each optical hybrid produces four outputs, the I+/I- and Q+/Q- signal components, which are subsequently detected by high-speed photodiodes. The detected photocurrents are converted to voltage signals by transimpedance amplifiers (TIAs) and then fully converted into digital signals via limiting amplifiers (LAs) that make hard binary decisions. A key advantage of our approach is being able to utilize TIA and LA circuits similar to those in proven NRZ designs that typically achieve the best energy efficiencies with lowest BER [47]. The TIA and LA outputs in both I and Q paths are



tapped as inputs for the OPLL circuitry that keeps the LO frequency and phase-locked to the incoming signal. The OPLL is implemented as a Costas loop, providing phase and frequency detection, and is specifically designed for QPSK modulation [44]. The Costas loop architecture has been demonstrated for robust 40 Gbps BPSK operation ( $10^{-12}$  at 35 Gbps) across temperature variations of  $2.6^\circ\text{C}$  [44, 45]. Furthermore, wide frequency pull-in range of  $\pm 30\text{-}40$  GHz and phase-lock in less than 10 ns have been achieved [44, 45]. OPLLs also provide a high-level of wavelength selectivity through the rejection of all other incoming wavelength channels, reducing the sensitivity to optical crosstalk when scaling to higher numbers of WDM channels. [44]

The BER is determined by the Q-factor at the receiver, which is directly related to SNR. For QPSK,  $\text{SNR} = Q^2$  [69]. All analysis here assumes a  $\text{BER} = 10^{-5}$ , or  $Q \approx 4.26$ , unless stated otherwise. This is sufficient to reach the KR4-FEC BER threshold of  $2.1 \times 10^{-5}$ . For homodyne detection, the SNR is

$$\text{SNR} = \frac{\langle I_{ac}^2 \rangle}{\sigma^2} \quad (2.4)$$

$$= \frac{4R^2\alpha_{coh}P_{laser}P_{LO}}{2qR\delta f(P_{LO} + \alpha_{coh}P_{laser} + I_d/R) + \sigma_T^2} \quad (2.5)$$

where  $R$  is the responsivity of the photodiode,  $\alpha_{coh}P_{laser}$  and  $P_{LO}$  are the transmitter and LO laser powers at the at the photodiode, respectively,  $\delta f$  is the bandwidth of the signal,  $I_d$  is the dark current in the photodiode, and  $\sigma_T^2$  is the thermal noise power. In our model,  $F_M$ -induced losses are included in the total link attenuation,  $\alpha_{coh}$ . [69] The energy per bit calculation is given by

$$EPB = (P_{RXIC} + P_{OPLL} + P_{TXdr} + \eta P_{TXlaser} + \eta P_{LO})/R_b \quad (2.6)$$

where  $\eta P_{TX,laser}$  and  $\eta P_{LO}$  represent the wall plug efficiency of the TX and LO lasers,

TW-MZM Driver	SEG-MZM Driver	TIA	OPLL
$\frac{V_{sig}^2}{8Z_0\eta_{dr,TW}}$	$\frac{L_{MZM}}{I_{seg}}$	$\frac{C_{seg}V_{sig}^2f_b}{\eta_{dr,SEG}}$	198 mW 53 mW

Table 2.1: Circuit power consumption

respectively;  $P_{TXdr}$  represents the MZM driver, as described in the first two columns of Table 2.1;  $P_{RXIC}$  and  $P_{OPLL}$  represent the receiver chain—including TIA, LA, and output buffer (OB)—and OPLL including the polarization control loop, respectively; and  $R_b$  represents the total bit rate. The circuit power dissipation was extracted from transistor-level Spectre simulations Cadence Virtuoso in the GlobalFoundries SiGe 8XP BiCMOS process for all circuits except the driver for the segmented modulator driver that were calculated as indicated in Table 2.1.

Other technology-dependent losses, such as waveguide passive attenuation, are included in the model and are listed in Table 2.2. The laser efficiencies of the TX and LO lasers are both set to 20%. The TX laser power was swept from 0 to 30 dBm, while the LO laser power was swept from -10 to 20 dBm, though we do not expect to feasibly operate the lasers at powers over 15 dBm. In all simulations that follow, unless otherwise noted, the unallocated link budget and symbol rate are set to 13 dB and 56 GBd, respectively.

## 2.3 Results and Discussion

Figure 2.4 plots show the LO power as a function of the TX laser power required to achieve a BER of  $10^{-5}$ . Figure 2.4(a) and (b) refer to Si photonic and InP TW-MZM, respectively. The plots also present the energy per bit (EPB) required to achieve  $BER = 10^{-5}$  ( $Q=4.26$ ) as a function of the transmitted power. Figure 2.4 therefore informs the available design and operating space for ACD links, indicating that sub-

Parameter	InP (dB)	Si (dB)
FC	1.5	1.5
MUX	1.5	1.5
PSR	1	1
Splitter	3	3
Excess (TX)	2	2
PC	2	2
Hybrid	6	6
Excess (RX/LO)	1	1
TX Loss	9	10.5
RX Loss	13	13
LO Loss	10	11.5

Table 2.2: Invariant loss parameters through the link are given in the table. The TX loss is equal to the sum of twice the FC loss, the loss of the MUX, PSR, and splitter, and excess loss for the Si TX. In the InP TX, the fiber coupling loss occurs once in the sum since it is assumed that there is an on-chip laser. The RX loss is equal to the sum of losses due to FC, MUX, PSR, polarization controller, hybrid, and excess losses. The loss in the LO path is the sum of FC, splitter, hybrid, and excess losses for the Si RX. For the InP implementation, the FC is neglected since it is assumed that there is an on-chip laser. FC = fiber coupling.

5 pJ/bit energy efficiency is possible for MZMs of practical active lengths in both Si and InP technologies. Figure 2.5 from the various components included in the link at the operating point indicated in Figure 2.4. We found the expected trade-off between drive voltage and TX and LO laser powers; that is, one can reduce the drive voltage but must increase the laser powers to overcome the higher incurred losses. However, the reduction of drive voltage does not necessarily achieve the minimum EPB. In general, the EPB decreases with increasing MZM length because an increase in the modulation length leads to a reduction in the required drive voltage. This is true until the increasing electrical and optical losses due to the growing total length of active sections overcome any potential increase in modulation efficiency, thereby increasing the overall EPB, as shown in Figure 2.6. Based on our modeling, this inflection does not occur until the length of the MZM grows to about 3-4 mm in Si and 1 mm in InP. For Si MZMs, we chose a 3-mm active length as a reasonable tradeoff from a device-density and packaging

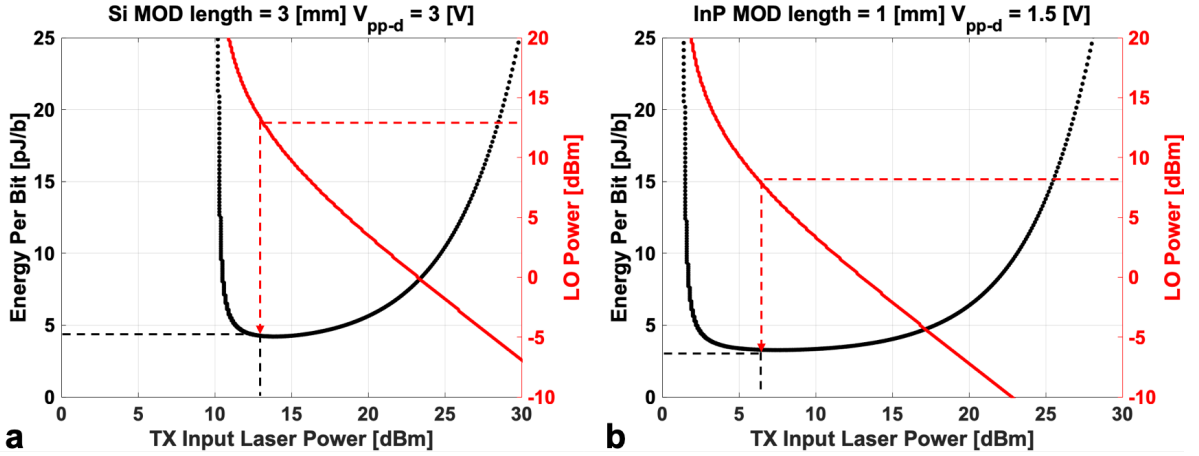


Figure 2.4: Simulation results for (a) a 3-mm-long Si Tx/Si Rx TW-MZM and (b) a 1-mm-long InP Tx/Si Rx TW-MZM. The EPB curve (black) and LO power curve (red) correspond to  $\text{BER} = 1 \times 10^{-5}$ , below the KR4-FEC threshold of  $2.1 \times 10^{-5}$ .

perspective. From the literature, we assume  $V_{\pi}L_{\pi} = 19$  V-mm [32, 52, 70–72], which we have also confirmed through device testing. In estimating the TX loss, we included the optical waveguide loss due to undoped and doped sections as well as the splitters and couplers. Finding a point along the LO power curve in Figure 2.4(a) that balances the TX and LO laser powers indicates around 13 dBm is required for both the LO and TX for the Si TW MZM based link. For Si SEG-MZM links, balanced TX and LO laser powers were found to be  $\sim 12.5$  dBm, and a similar EPB value of around 4 pJ/bit was projected for a voltage of  $1.2 V_{pp-d}$ .

Transmitters incorporating InP MZMs offer improvements in both link efficiency and laser power requirements compared to Si MZMs. Since the InP platform shows higher modulation efficiency (lower  $V_{\pi}L_{\pi}$ ) and higher passive waveguide loss compared to the Si platform, we constrained the design to a 1-mm long device. We also found that a 1-mm long InP device is the optimal design point for achieving the minimum EPB, which can be seen in Figure 2.6. For the InP simulations, we model  $V_{\pi}L_{\pi}$  as 2 V-mm—an order of magnitude lower than Si. The InP TW-MZM design point is very realizable, requiring

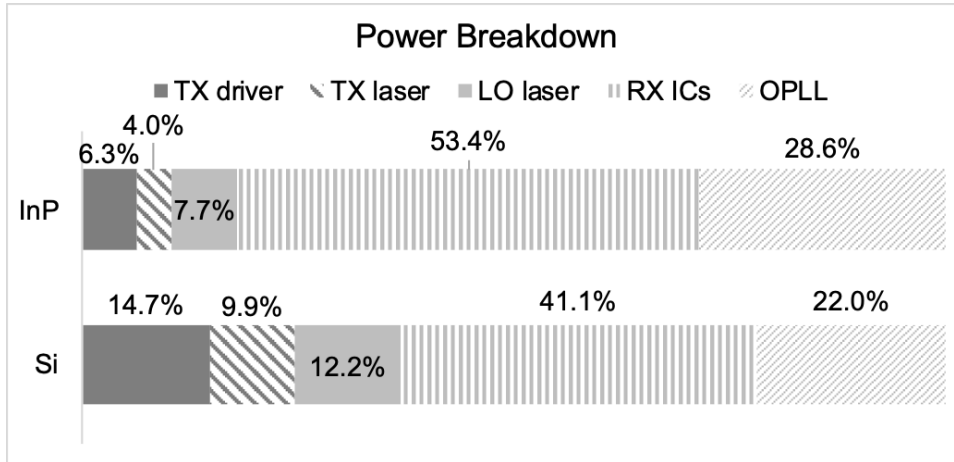


Figure 2.5: The proportion of power taken up by each component in the link at the operating point indicated in Figure 2.4.

only  $1.5 V_{pp-d}$  and roughly +7 dBm from both the input and LO lasers as can be seen in Figure 2.4(b), while the InP SEG-MZM design point suggests that TX and LO laser powers of +5 dBm and drive voltage of  $1 V_{pp-d}$  to be the ideal operating point for a similar EPB. InP device parameters are based on measured data of PICs fabricated in the UCSB Nanofabrication Facility.

Note that although InP TXs offer a more efficient solution than Si, a Si RX implementation is more favorable as on-chip polarization de-multiplexing is much more readily implemented. In addition, an all Si implementation likely offers advantages in electronic and photonic integration, packaging, and cost. While it is possible to achieve low EPB, Si TXs have higher loss due to less efficient MZMs that degrades link efficiency and drives laser power requirements to challenging levels. Furthermore, both InP and Si SEG-MZMs will face significant challenges in integrating a large number of drivers alongside or flip-chipped on the modulator due to the large number of segments that make up the total active length required to achieve sufficient modulation. Thus, we focus on TW-MZMs for the rest of the analysis presented here. However, SEG-MZMs with integrated drivers could be a very compelling solution in monolithic processes that offer high-performance




Figure 2.6: Simulation results when we compare the minimum EPB (black) and drive voltage (red) for TW-MZMs in InP and Si. SEG-MZMs yield similar minimum EPB, and lower drive voltage.

CMOS capable of 50 GBd operation [67].

In Figure 2.6, we compare InP and Si TW-MZM EPB and drive voltage with respect to modulator length. The much lower  $V_{\pi}L_{\pi}$  of InP contributes to a much lower drive voltage compared to the Si structure. However, a steep rise in EPB is projected for longer InP modulators. In this regime, propagation losses in the active sections dominate, driving up laser power requirements that overwhelm any power savings due to reduced drive voltage. This optimal point in EPB occurs at about 1 mm for InP due to much higher propagation losses in its active sections but is not a significant factor for modulators less than 6 mm long in Si. The 3-mm Si TW-MZM link can achieve energy efficiency approaching the InP MZMs, although still requiring higher power operating points for both TX and LO lasers. Combining this observation with the fact that Si offers a cost-effective platform for

building large-scale, highly integrated PICs in 300 mm wafer manufacturing processes, we focus on Si TW-MZMs for the rest of our analysis in this paper. However, the observations and conclusions that follow are applicable to links incorporating InP MZMs.

To further explore the design space available for ACD link architectures, simulations were conducted for lower BER targets, reduced link budgets, and single polarization operation. Selected results are presented here for links with Si TW-MZM transmitters. Moving to a more aggressive BER target,  $10^{-12}$ —often referred to as “error-free”—has a minor impact on the achievable EPB but drives both source and LO lasers powers considerably higher, to around +16 dBm, as seen in Figure 2.7. Increasing the output drive voltage can also potentially result in decreased bandwidth of the driver circuits. Reducing the link budget from 13 to 4 dB has the opposite effect on link operation. Effects on the minimum achievable EPB are also minimal, but the required LO and input laser powers are reduced significantly to about +6 dBm. The EPB curve is also substantially flattened, indicating a wider range of choices for LO and TX laser power that achieve the optimal EPB. Figure 2.8 shows how EPB changes with voltage, assuming an operating point where the TX and LO laser powers are equal based upon the same link configurations analyzed in Figure 2.7. Here it is evident that while the minimum EPB may occur at lower drive voltages, the TX and LO output powers may be unfeasible. Finally, we investigated the case of a single polarization link, starting from the baseline Si TW-MZM case presented in Figure 2.4(a). One may expect that the EPB will decrease with lower complexity of the PIC due to the lack of polarization-specific components, but in fact, the change is not significant. This is because ICs take up a significant portion of the total link energy consumption at these optical powers (100s of mW compared to 10s of mW). Therefore, if we halve the number of ICs in the receiver, we reduce the IC energy consumption of the link almost by half for half the number of bits. On the other hand, by eliminating the need for sharing TX and LO lasers between two polarizations,




Figure 2.7: Starting from the conditions used to generate the results for 3-mm-long Si TWMMZMs shown in Figure 2.4(a), LO power and EPB vs. TX power for BER  $10^{-5}$  with a 4 dB link margin and BER  $10^{-12}$  with a 13 dB link margin cases are shown. The drive voltage is set to  $3 V_{pp-d}$ . LB = link budget.

the operating points of both lasers are reduced by roughly 3 dB to 8 dBm.

## 2.4 Discussion

In the previous section we have shown how we can achieve sub-10 pJ/b, 200 Gbps per wavelength links. Increasing the data rate in a single lane by implementing a higher order modulation format is another path to higher aggregate link bandwidth but will decrease the unallocated link budget [2]. Utilizing the additional loss budget by inserting optical switches can decrease cost, latency, and power consumption of data centers. We will show in this section that increasing to higher order modulation rates decreases the unused link margin and may restrict connectivity between servers.






Figure 2.8: Starting from the conditions used to generate the results for 3-mm-long Si TWMZMs shown in Figure 2.4(a), and taking the operating point where the TX and LO laser powers are equal, this plot shows how the EPB changes for a given drive voltage and the required TX/LO laser powers to close the link. LB = link budget.

### 2.4.1 Unallocated Link Budget

Optical switching is the subject of worldwide research, motivated by the promise of adding reconfigurability to data center networks and potentially improving overall data center energy efficiency [73–75]. The principle of adding a layer of arrayed waveguide grating routers (AWGRs) or optical switches layer to a data center to enhance scalability while reducing cost, power and latency, is described in [73] and [74], respectively. However, for optical switching to be practical, the links traversing the switches must either have enough budget to accommodate the losses of the switches, or the switches must be made transparent by incorporating optical gain. The latter approach, usually relying on semiconductor optical amplifiers, presents integration challenges in Si photonic platforms and also raises operational issues including added noise, gain uniformity

across wavelengths, and crosstalk [76]. We believe the best approach to enable photonic switching is through expanding available link budgets to accommodate the insertion loss of switching or passive wavelength routing components. In order to assess the achievable link budgets offered by candidate link architectures, we follow the analysis approach in [2] to compare QPSK to both IMDD alternatives as well as 16QAM. The results are presented in Figure 2.9 for an analysis conducted under a consistent set of assumptions for each link: the same laser powers, MZM modulators, and target BER of  $10^{-5}$ . For drive swings above  $V_\pi$ , 16QAM can offer some improvement in budget compared to IMDD, but the advantages of QPSK are much more substantial. At full  $2V_\pi$  drive levels, QPSK expands link budgets by 8 dB compared to PAM4 and 12 dB compared to PAM8. At a more practically realizable drive voltage of  $0.6 V_\pi$ , QPSK offers increases of 2 dB and 6 dB compared to PAM4 and PAM8, respectively. In addition to enabling photonic switching in data center networks, the expanded link budgets offered by QPSK-based ACD can also potentially be used to improve transceiver yields owing to the reduced sensitivity to optical loss as well as reduce transceiver power consumption by lowering the operating points of the lasers and drivers.

### 2.4.2 Optical switch-based networks

Here we illustrate the design flexibility offered by scaling bandwidth by adding additional 200 Gbps/ $\lambda$  WDM channels as opposed to increasing per- $\lambda$  bit rates. A relevant design example is a data center network connecting 131,072, 50-Gb/s servers, using next-generation 51.2 Tb/s switches. A conventional design, using 800-Gb/s inter-switch links, would employ three levels of electronic switches, with 51.2 Tb/s switches used in both the spine and aggregation layers, and a smaller 6.4 Tb/s ToR switch (supporting 64 servers per rack). If the 800 Gbps transceivers were realized using four independent, separable

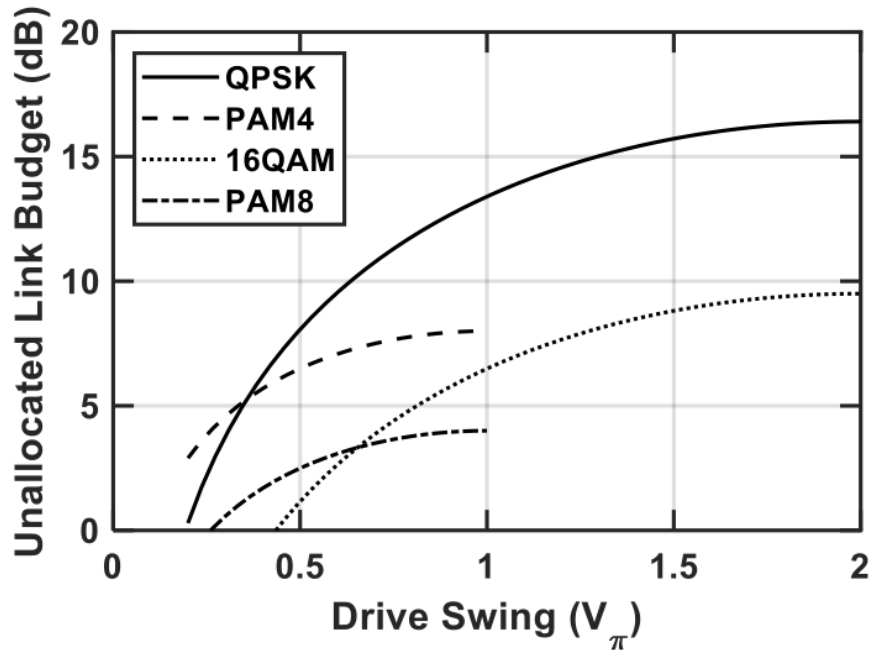


Figure 2.9: Comparison of unallocated link margin in coherent and IMDD links, assuming MZM drive, and a target BER of  $10^{-5}$ . The QPSK curve assumes analog coherent link performance as described in this work, while the other curves assume representative link performance projections for next-generation IMDD and digital coherent links [2].

lanes operating at 200-Gb/s, the same number of servers could be supported using only two levels of 51.2 Tb/s switches interconnected by 200-Gb/s links. The drawback in this scenario is that four times as many fibers would be needed. However, the number of fibers can be reduced back to the original number of fibers as in the classic design by using four 200-Gb/s wavelengths per fiber, and inserting an optical wavelength-routing layer between the two switching levels, described in detail in [73]. With only two levels of switches instead of three, this optimized design, taking full advantage of WDM parallelism, results in lower cost, latency, and power consumption. Thus, as described in the Section 2.1 and as shown in Figure 2.1, to realize an 800G link, we include four lanes multiplexed into a single fiber. Likewise, to scale to higher data rates, such as 1.6 Tb/s, eight lanes would be multiplexed into a single fiber, without changing the data format to

take advantage of the inherent link budget advantage of QPSK as well as the efficiency advantages presented by being able to implement low-power electronics similar to the circuits used in NRZ link implementations that we have implemented in our ACD link. The wavelength routing layer used to flatten the data center network in the above example could be based on all passive elements such as AWGRs, but there is also the possibility to add reconfigurability to the wavelength routing layer in the form of small (e.g.  $4 \times 4$ ,  $8 \times 8$ ,  $16 \times 16$ ) WDM photonic switches. Furthermore, the power consumed by optical switches are independent of the data rate of the signals they route. Optical switches do not need to perform power-hungry optical-electrical-optical conversions and instead perform the switching in the optical domain while being transparent to data format and data rate [77]. Consequently, as links move to higher data rates, their impact on the total energy per bit decreases. Focusing on switches implemented in planar Si platforms, which offer a realizable path towards mass manufacturing, there have been several noteworthy recent demonstrations of Si photonic switches, including: a  $4 \times 4$  switch with integrated gain [76], a  $32 \times 32$  port polarization diverse switch [78], and a  $240 \times 240$  port MEMS-based switch [79], among many others. We have demonstrated a wavelength-selective crossbar switch with multiple wavelength-selective elements at each cross-point [80]. Each switch offers a promising feature: the large port count from stitching of multiple die in [79], the path-independent loss in addition to the polarization diversity in [78], and the gain and custom ASIC integration in [76]. While these demonstrations differ in port count and switching time, they are non-blocking and exhibit losses less than the unallocated link budget assumed in the link model presented above. Having a large unallocated link budget eases the requirement of having ultra-low-loss photonic components in the switch design. This ensures that the insertion loss of the switch remains low so that the port count can be scaled to 32-64, offering flexibility in data center network architectures.

## 2.5 Conclusion

A comprehensive QPSK coherent link model has been presented and indicates that EPB under from 5-10 pJ/bit is possible with substantial improvements in optical loss budget. The simulation tool allows exploration of optical and electrical parameters that impact PIC design. Measured hardware will verify and refine the parameters used in the link analysis. Figure 2.10 shows functional hardware that we have built and characterized to provide hardware-derived inputs to our modeling. Design and characterization results for the transmitter in Figure 2.10(a) have been reported in [81], while receiver subsystems shown in Figure 2.10(b) and (c) as well as other transmitters will be reported in forthcoming publications. Finally, we showed that QPSK links increase unallocated link budget. Analog coherent detection based on QPSK modulation has the potential to enable novel network designs incorporating wavelength routing and switching while simultaneously maximizing energy efficiency, facilitation future lower power data center network architectures that maximize overall data center efficiency.

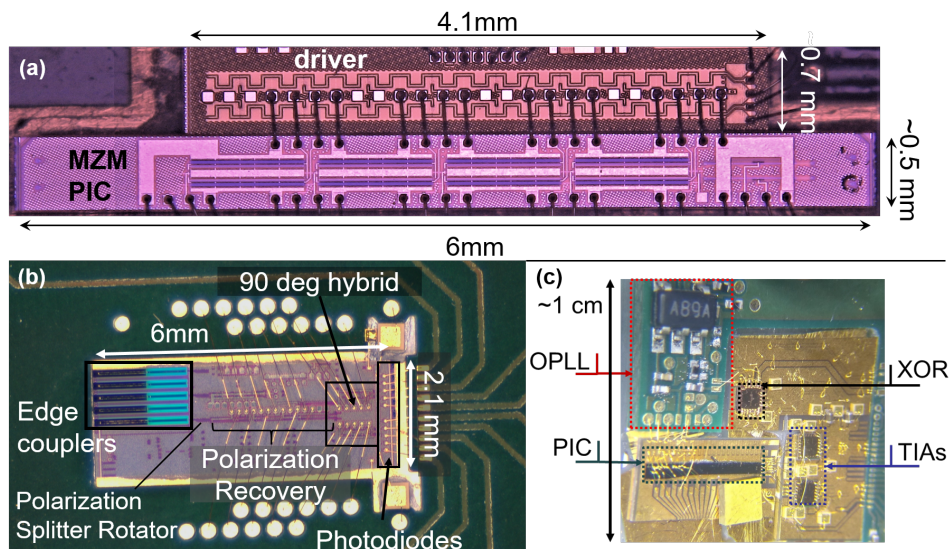


Figure 2.10: Preliminary hardware for (a) a Si TX modulator and driver, (b) Si RX PIC packaged with an OPLL.

Finally, while there was greater focus on implementing energy efficient links in Si, the analysis indicated that InP provided a path towards lower pJ/b operation as well as lower TX and LO laser operating points. Current research efforts on III-V integration on Si remains very active, and could provide an even greater reduction in energy over an all-Si link solution. If a III-V modulator and laser can be integrated on Si in a high-yield, reproducible manner, one can take advantage of the inherently more energy efficient properties that III-V materials to modulate the light and therefore reduce the driving voltage of the MZM. Furthermore, a III-V-on-Si modulator and laser solution can reduce the loss in the optical path of the link and in the RX, allowing the lasers to be driven at a lower power than in an all-Si solution. In short, one can take advantage of the increased energy efficiency that III-V materials have to offer all while taking advantage of the compactness, polarization diversity, and ease of manufacturability of Si photonics.

# Chapter 3

## Crossbar Switches

In the preceding chapter, an energy efficiency analysis of analog coherent links for data centers was presented. A link budget of 13 dB was included such that a reconfiguration of the network, as described in chapter 1 is possible. One component that can consume part of this link budget is optical switches. It has been proposed that having optical switches can help flatten or restructure data center network topologies for further energy savings [73, 75].

### 3.1 Ring-Based Switch Fundamentals

In this section I briefly describe the underlying physics that are necessary to understanding the performance of the switches presented in the rest of this thesis.

#### 3.1.1 Waveguiding

The most fundamental building block of any PIC is the waveguide. As the name suggests, it guides waves and acts as the optical equivalent of an electrical wire. It is comprised of two materials of differing indices of refraction, where the higher index

material is used as the waveguiding material and the lower index material is known as the cladding. Two common cross-sectional geometries of a waveguide are shown in Figure 3.1.




Figure 3.1: (a) and (b) shows a perspective view of a strip and ridge waveguide, respectively. (c) and (d) show the intensity distribution of the fundamental mode for the strip and ridge waveguides, respectively.

The light forms certain lateral spatial distributions based on the lateral dimensions known as modes as it travels along the waveguide. Mathematically-speaking, these modes are eigensolutions to the wave equation. The lateral dimensions determine the number



of modes that are allowed to propagate through the waveguide. It is desirable to have so-called single-mode waveguides on PICs because the lowest order mode (also known as the fundamental mode) typically carries the information and it is typically difficult to separate multiple modes from each other in a waveguide. For the photonic devices presented in this thesis, it is assumed that the waveguides are single-mode unless stated otherwise. Figure 3.2 shows the effective index vs waveguide width for particular heights of waveguide. The effective index is a weighted average of the index of refraction over the spatial distribution of the light, typically in the lateral dimension. The effective index determines how fast the light propagates through the medium. Due to the boundary conditions that result from Maxwell's equations, the light is not fully contained in the waveguide; while the majority of the power is contained in the waveguide, some evanescently 'leaks' out.

The different waveguide geometries shown in Figure 3.1 are typically used for different purposes. The strip waveguide shown in Figure 3.1(a) is typically used for passive purposes, such as routing. The confinement of the optical mode is tighter thus allowing for tighter bends. The ridge waveguide geometry shown in Figure 3.1(b) is typically used for active elements which will be described in more detail in a later section.

The changing the cross-sectional geometry affects the propagation of the light through the waveguide. An abrupt change in width results in so-called scattering and can result in significant loss. This abrupt change in waveguide geometry either by changing the lateral dimensions of the waveguide or bringing another waveguide close to the first should be avoided in order to reduce scattering losses. In many of the devices designed and presented in this thesis, adiabatic tapers to and from different waveguide geometries and adiabatic couplers between waveguides were used to minimize loss.

For a more rigorous treatment, the reader is directed to the following resources: [69, 82, 83].




Figure 3.2: This plot shows the effective index corresponding varying widths of waveguide for a given thickness. When the waveguide becomes wide enough, it can support more than one mode.

### 3.1.2 Microrings

In this section, I will give a brief introduction of rings, which gives rise to the wavelength-selectivity of the switches presented in this thesis.

A resonator can be formed on-chip by connecting the input and output ports of a waveguide together such that it forms a ring. A resonance occurs when the wavelength of light is an integer multiple of the roundtrip length. That is,

$$m\lambda_r = n_{eff}L, \quad (3.1)$$

where  $m$  is a natural number,  $L$  is the roundtrip length of the ring,  $n_{eff}$  is the effective index, and  $\lambda_r$  is the resonant wavelength. The spectral distance between peaks is known




Figure 3.3: (a) shows the schematic of an all-pass ring, while (b) shows a schematic of an add-drop ring. (c) is a schematic of a two serially coupled add-drop ring configuration.

as the free spectral range (FSR), and is given by

$$\text{FSR} = \frac{\lambda_r^2}{n_g L}, \quad (3.2)$$

where  $n_g$  is the group index. Note that a small roundtrip length corresponds to a large FSR, which is desired for ring-based switches such that there is only one resonance within the wavelength band of interest. If there are multiple resonances within a wavelength band in a switch, two wavelength channels would not be switched independently. The roundtrip amplitude is given as  $a$ . There are two ring configuration types, add-drop and all-pass ring, which are shown in Figure 3.3. The latter is discussed in more detail in chapter 4. Due to the light ‘leakage’ from the waveguide, light can ‘leak’ into a nearby waveguide in a phenomenon called coupling. How much light couples into the second waveguide is dependent primarily on the distance between the two waveguides and the length of the so-called coupling region. Rather than having a point coupling region, as would be the case with a circular ring resonator, it is common to utilize a so-called race-track ring resonator in which straight sections are designed to change to fine-tune the length of the coupling region and therefore the coupling to the ring. We denote the

cross-coupling and self-coupling coefficients to be  $k$  and  $r$ , where  $k^2$  and  $r^2$  denote the power transferred and remaining in each waveguide, respectively. For an add-drop ring, there are two pairs of  $r$  and  $k$ . In an add-drop ring, light at the resonance wavelength couples into the ring from the first waveguide and then couples to the second waveguide, while light that does not meet the resonance condition continues to pass through. This property forms the basis of the wavelength-selectivity of the switches presented later in this chapter. The power transmission of the through and drop ports are given by

$$\frac{I_{thru}}{I_0} = \frac{r_2^2 a^2 - 2r_1 r_2 a \cos \phi + r_1^2}{1 - 2r_1 r_2 a \cos \phi + (r_1 r_2 a)^2} \quad (3.3)$$

$$\frac{I_{drop}}{I_0} = \frac{(1 - r_1^2)(1 - r_2^2)a}{1 - 2r_1 r_2 a \cos \phi + (r_1 r_2 a)^2}, \quad (3.4)$$

where  $I_0$  is the input power and  $\phi$  is the phase change incurred by the ring. The transmission are plotted in Figure 3.4. [84]




Figure 3.4: The responses at the thru and drop ports for an add-drop ring.

As can be seen from Figure 3.4 the width of the peaks is narrow. The width required for wavelength-selective switches is dependent on the width of the wavelength channel. To obtain a flatter, wider response, it is common to use a second serially coupled ring [85], as shown in Figure 3.3(c). For a maximally flat response, the coupling coefficients between the rings and bus waveguides are assumed to be the same. The change in transmission is shown in Figure 3.4. [86]

### 3.1.3 Ring Tuning Mechanisms

Tuning the resonant wavelength of a ring is often desired either by design or to correct some deviation in the designed wavelength due to manufacturing variability. By observing the resonance condition of the ring given in equation 3.1, it is clear that to change the resonant wavelength of a ring one can change one or both of the following: round-trip length or effective index. Once the ring is produced in a commercial silicon photonics foundry, its dimensions cannot be changed; thus, the only method of change the resonant wavelength of the ring is to change the effective index. This is typically done by utilizing either a thermo-optic or electro-optic effect.

Thermal tuning is used exploit that there is temperature dependence in the index of refraction. This is typically achieved with a heater either with a strip of metal, which is typically located a few microns above the waveguiding layer in commercial foundry services, or with silicon. While metal heaters are attractive because they typically incur low loss, they are relatively inefficient, especially if the metal is highly conductive. Embedding a heater in silicon by comparison can be more efficient, especially if it is possible to connect the heater to the waveguide with a shallow etch silicon for more direct heat transfer. In doing so, care must be taken not to disturb the mode traveling through the waveguide. Thermo-optic switches typically have switching times on the

order of microseconds

Electrooptic tuning can also be used to change the refractive index by exploiting the effect that charge carriers have when interacting with light. This is typically achieved by doping the waveguide with p- and n-type dopants to create a pn-junction. While this provides switching times on the order of nanoseconds [87], electrooptic tuning is typically associated with higher loss since the charge carriers are more likely to absorb the photons. Electrooptic tuning has been investigated in great detail by Soref et al. in [88].

Alternative tuning methods to thermo-optic and electro-optic methods that require utilizing novel materials integrated into the rings. One approach includes using a piezoelectric effect to deform the ring and thus changing its round-trip length [89,90]. Tuning as low as 0.5 mW has been demonstrated using this technique. Another alternative approach is to integrate InP on top of a Si ring, where the ring acts as the cavity of a laser, as demonstrated in [91–93].

### 3.1.4 Photonic Switch Topologies

An optical switch cell is the smallest unit that consists of one or more input ports with a mechanism between the inputs and outputs that redirects the light to different outputs. A typical optical switch cell has two inputs and two outputs, as shown in Figure 3.5. For an add-drop ring, the thru and drop ports act as the two possible output ports.

Typical metrics by which to characterize switches are included in Table 3.1 [87].

How switching elements are connected together determines the routing that is possible through the switch. The network that is built from the switching elements must not be blocking. That is, any input can connect to any output either regardless of pre-existing connections through the network (strictly non-blocking) or by rearranging pre-existing



Figure 3.5: A block diagram of a switch cell in which there are two inputs and two outputs.

Metric	Description
switching time	Minimum time to configure the switch
path loss	Loss through a path of the switch, usually reported in dB
crosstalk	Output power ratio between the intended output and an unintended output
bandwidth	FWHM range of output
area	Area on the chip taken up by the switch. Sometimes the optical I/O are included in the area
power consumption	Power consumed for a particular switch configuration

Table 3.1: Typical switch metrics [87]

connections through the network (rearrangeably non-blocking) [87]. There are many different switch topologies that have been proposed [87]. In this thesis, only crossbar and Beneš topologies are used.

### 3.1.5 Notable Silicon Photonic Switch Demonstrations

There have been a few notable demonstrations of switches in silicon photonics. There have been demonstrations of a  $240 \times 240$  MEMS-based switch [79], a  $32 \times 32$  polarization diverse switch [78], as well as a  $4 \times 4$  switch with integrated gain [76]. While the MEMS-based switch has very high port count and low loss per port, the voltage required to actuate the MEMS is high, on the order of a few tens of volts. Thus, it has not been

demonstrated with a custom driver. The 32-port polarization diverse switch reports roughly 10 dB of loss. Finally, the 4-port switch with integrated gain and controlled with a custom driver IC reports net neutral insertion loss though crosstalk values range from -10 dB to -15 dB.

These switch demonstrations have wide switching bandwidth to redirect a large number of wavelength channels together. In this thesis, I will be describing wavelength-selective switches that can reroute wavelength channels from within a WDM signal. The idea of silicon photonic ring-based crossbar switches was first proposed by S. Emelett and R. Soref in [94–96] in 2005. These crossbar switches included multiple rings per crosspoint and demonstrated that serially cascading rings at the crosspoint made the edges of the passband sharper. The Bergman group at Columbia University has demonstrated a number implementations. Most recently, they have demonstrated an  $8 \times 8$  Banyan-type ring-based switch [97] and a polarization-diverse ring switch with a switch-and-select architecture [98]. For the former, the on-chip loss was almost 10 dB for the worst-case, and an average crosstalk of -16 dB. For the latter, only a  $2 \times 2$  switch cell is demonstrated, with polarization-dependent losses on the order of 1.6 dB, on-chip loss to be about -4 dB, and crosstalk measured to be over 45 dB. For both these demonstration, only one ring pair is used per crosspoint, and thus can switch only one wavelength for each WDM input.

## 3.2 C-Band $4 \times 4$ Crossbar Switch

The work in this section has been published in [99, 100].



### 3.2.1 Introduction

To manage the growth of data centers and high-performance computing facilities, and to mitigate their corresponding unsustainable growth in energy consumption, it has been advocated that wavelength-selective optical switches be deployed in their interconnect networks [74,101–103]. An integrated silicon photonic,  $N \times N$  WDM crossbar switch based on thermally tuned, microring resonators (MRRs) was demonstrated in [104] and [105]. At a high level, that switch has the same architecture as that of the switch we study here, which is shown schematically in Figs. 3.6(a) and 9(b). The switch is flexible in that, with  $M$  wavelengths incident at each of the input ports, one can drop up to  $L$  wavelengths ( $1 \leq L \leq M$ ) at each of the outputs ports. This is made possible by having  $L$  MRRs inserted at each switching cross-point, as shown in Fig. 3.6(b). Each MRR consists of a pair of coupled microrings to reduce the inter-channel crosstalk and provide wider spectral response [106].

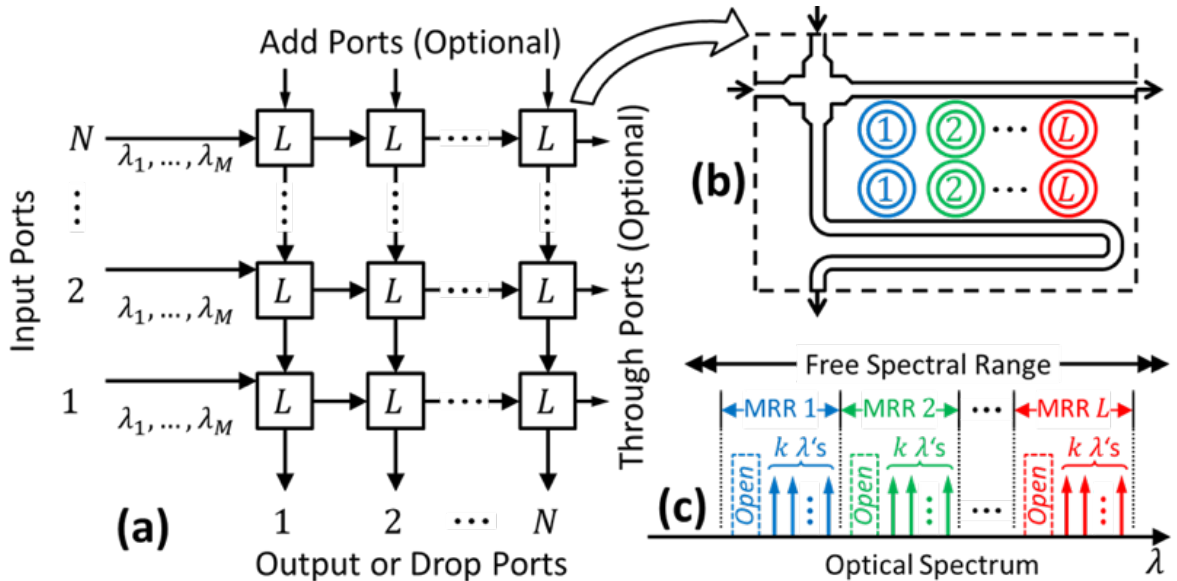


Figure 3.6: A schematic diagram of a spectrally partitioned  $N \times N$  crossbar switch with  $M$  wavelengths per port and up to  $L$  wavelength drops per cross-point. (a) A high-level block diagram of the switch. (b) Possible realization of an  $L$ -MRR cross-point. (c) Details of the spectral partitioning.

One of the operational properties of the previous version of the switch [104], [105] is that each MRR is required to be tunable over the entire free spectral range (FSR) of the rings. However, this stresses the ring heaters and increases the tuning power consumption. We mitigate these problems by restricting the tuning range of each of the  $L$  MRRs to  $1/L$  of the FSR, as shown in Figure 3.6(c). Each MRR is tunable over  $k = M/L$  wavelengths, falling within its partitioned part of the spectrum. Note that each tuning partition has an open wavelength channel to which one can tune the corresponding MRR to allow a wavelength to bypass the cross-point. The use of standard wavelength assignment algorithms with our spectral partitioning leads to wavelength blocking. We have devised a novel wavelength assignment algorithm to allow arbitrary input-output connectivity despite our spectral partitioning constraint.

### 3.2.2 Experimental Results

The  $4 \times 4$  switch was fabricated in the AIM Photonics foundry. Each crossing in the switch is comprised of three add-drop serially-coupled MRR pairs. While the switch was designed for flip-chip integration into a custom AIM Photonics silicon photonics interposer, the chip was packaged on a custom PCB as shown in Fig. 3.7. The design of the switch fits in a  $1.1 \times 2.2 \text{ mm}^2$  area, set by the large number of pads and pad pitch.

Figure 3.8(a) shows the fiber-to-fiber spectral response through one path of the switch, which was chosen to be from input 1 to output 2. The shortest path (input 1 to output 1) was inaccessible due to the tight pitch dictated by the design constraints imposed by the interposer design rules. In the measured path, the light must pass through one crossing before getting dropped at the second output. Source-measurement units (SMUs) were used to make the measurement. The average loss due to passing by an MRR is roughly 0.4 dB per MRR for a total of 1.2 dB for passing through the first crossing. Thus, the loss




Figure 3.7: (a) shows the experimental setup with the single-mode optical fiber while (b) shows a close-up of the wirebonded chip.

along the longest path (i.e. input 4 to output 4) due the thru loss of the MRR crosspoint is predicted to be 9.4 dB. The loss of the testing system, including an external polarization controller (PC), and the on-chip edge couplers are not included in the response shown in the figure. The PC was used to optimize for TE polarization. The additional loss is likely due to imperfect coupling due to an optical I/O pitch that is too tight for two side-by-side single-mode fibers, dictated by the design rules for interposer integration. The crosstalk for channels at a 200 GHz spacing is roughly 20 dB. In Fig. 3.8(a), all three MRRs were given the same bias so all three ring pairs were tuned together across their respective sub-bands. As can be seen in Fig. 3.8(b), each MRR can tune fully independently across its own sub-band, or a quarter of the total FSR. Here, the second and third MRRs were held fixed at 3 V and 6 V, respectively. The average power dissipation to tune an MRR across the 6.5 nm sub-band is 13.7mW. The 100  $\mu\text{m}$  distance between the MRRs—determined by the 100  $\mu\text{m}$  pad pitch—is large enough that no thermal crosstalk was observed.

We tuned the MRRs using a custom PCB-based 64-channel 10 V driver. As a proof of concept, twelve channels were used to control the rings establishing the desired path through the switch network. Smooth wavelength tuning is achieved via a 12-bit digital to

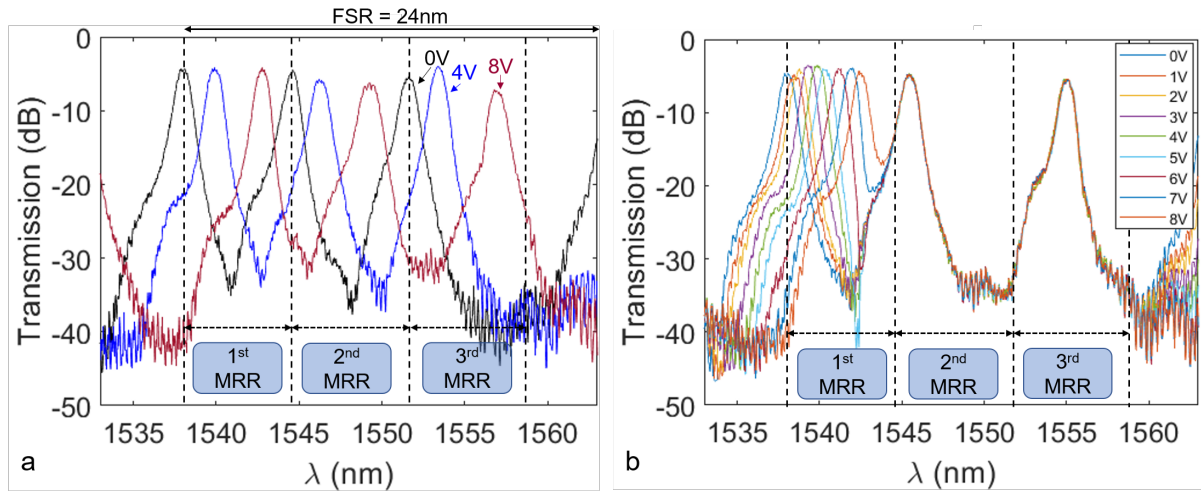


Figure 3.8: (a) shows the tuning spectrum of all three MRRs within each sub-band using SMUs. Each ring is tuned with 0 V, 4 V, 8 V. Figure (b) shows independent tuning of one ring with respect to the other rings.

analog converter (DAC), resulting in a 2.44 mV or  $\sim 1.63$  pm step size. The experiments were repeated to verify driver operation, the results of which are shown in Figure 3.9. Here, the 1st and 3rd MRRs were held fixed at 0 V and 4 V, respectively.

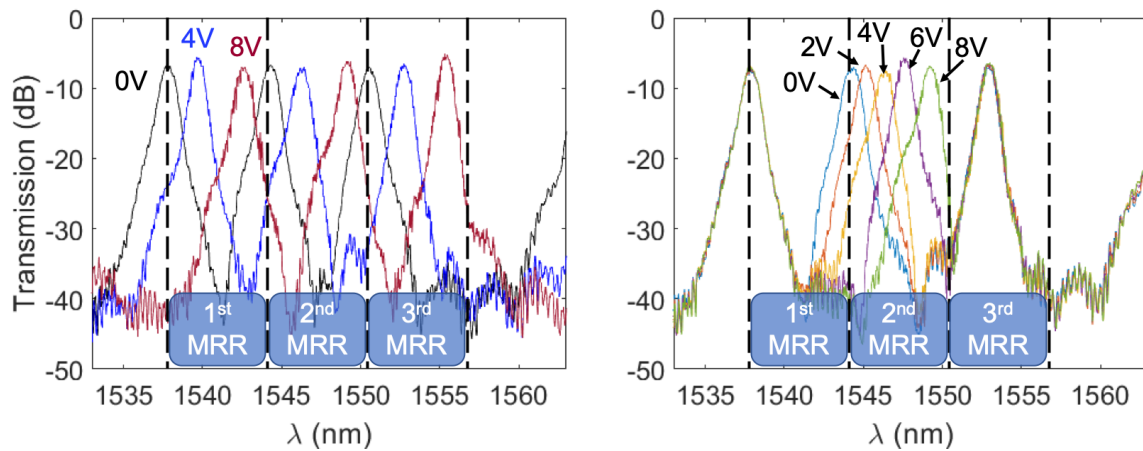


Figure 3.9: (a) shows the tuning spectrum of all three MRRs within each sub-band. Each ring is tuned with 0 V, 4 V, 8 V using a custom driver. Figure (b) shows independent tuning of one ring with respect to the other rings.

## 3.3 O-Band $4 \times 4$ Crossbar Switch

The work in this section appears in part in [80].

### 3.3.1 Introduction

Many chip-scale switches that have been demonstrated are in the C-band [104, 105, 107], however data centers use O-band wavelength channels. The authors in [108, 109] have demonstrated switches in the O-band, but these switches were broadband. In this section, we present a  $4 \times 4$  chip-scale wavelength-selective switch in the O-band, as demonstrated in the previous section, but modified for operation for the O-band, which required a redesign of the MRRs.

### 3.3.2 O-band Switch Design

The switch was fabricated in the TowerJazz silicon photonics process using their multi-project wafer service. We designed custom MRRs of a racetrack geometry as described in [85] to provide a maximally flat filter response. Three pairs of MRRs were included at each crossing, with resonances at 1300 nm, 1306 nm, and 1311.5 nm, corresponding to straight coupling sections of lengths  $4.75 \mu\text{m}$ ,  $3.25 \mu\text{m}$ , and  $2 \mu\text{m}$  and FSRs of 17 nm, 19 nm, and 21 nm, respectively. All the racetrack rings had an effective radius of  $5 \mu\text{m}$ . A heater was embedded in the center of each MRR by doping a smaller silicon MRR with the standard n-type dopant provided by the foundry process. A shallow etch region of un-doped silicon connects the heater to the MRR waveguide in order to provide an efficient method of heat transfer from the heater to the waveguide. The heaters were designed to be on the order of a few  $\text{k}\Omega$  such that a custom driver can provide the necessary voltage and currents required to control the switch. Uniformity of the heater resistances is desired such that each channel of a switch driver need not be customized

for each ring, which would add unnecessary complexity to the driver design. One side of each heater was routed directly to pads that were placed almost directly above the rings to simplify the design. The other side of each heater was routed to an on-chip ground plane which could be connected to an external ground such as that of a driver circuit via several dedicated ground pads.

### 3.3.3 Results and Discussion

To demonstrate functionality of the switch, the shortest path through the switch—that is, input 1 to output 1—was tested. It was packaged on and wirebonded to a custom printed circuit board (PCB) to fan out the electrical signals from the on-chip switch, as shown in Fig. 3.10a. A closeup of the switch chip is shown in Fig. 3.10b. The size of the switch—which was measured to be roughly  $900\ \mu\text{m}$  by  $2\ \text{mm}$ —was largely dominated by the  $70\ \mu\text{m}$  square pad size and  $100\ \mu$  pitch.

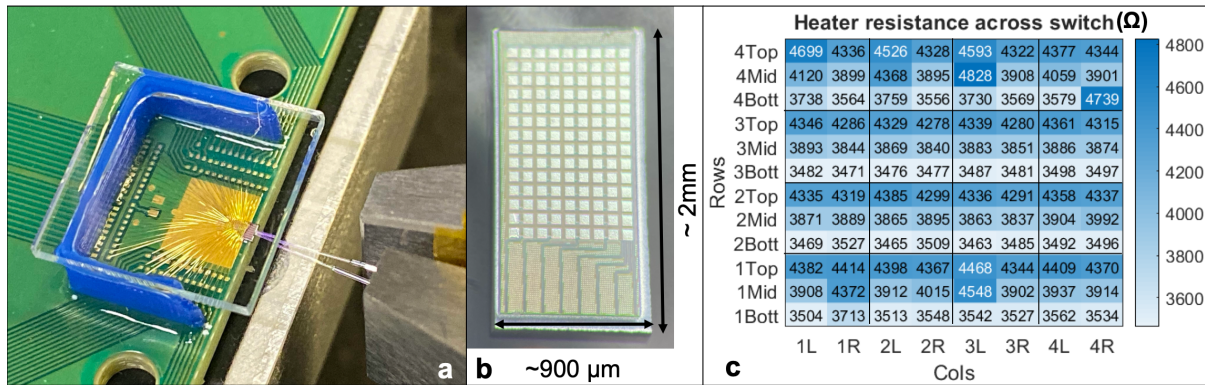


Figure 3.10: (a) shows the switch wirebonded assembly, while (b) shows a close-up of the switch chip. (c) shows the resistance of the heaters across the switch chip.

The custom edge couplers were designed to have a mode spot size of  $3.5\ \mu\text{m}$  and were characterized to be 3.1 dB of loss. Figure 3.10c shows the uniformity of the resistance of the heaters across the switch. The average resistances of the heaters were measured to be  $4370 \pm 90\ \Omega$ ,  $3985 \pm 227\ \Omega$ , and  $3586 \pm 230\ \Omega$  for the large, medium, and small MRRs,

respectively. Due to the process variations, the rings were not equally spaced in the original wavelength region of interest of 1298-1316 nm; instead a different region within the O-band where the rings are more equally spaced was chosen within a single FSR. As a proof of concept, we tested the shortest path through the switch, which corresponds to input 1 to output 1 in Figure 3.6. Figure 3.11a shows the response of the switch across the full 11.6 nm FSR, not including the loss due to the edge couplers. Each MRR was tuned across their own sub-band using source-measurement units (SMUs) to control the voltages on the heaters and measure the current. In the configuration in which all the rings are tuned to the end of their respective sub-partitions, the total power consumption of the cross-point was found to be 32.2 mW. The voltage required to tune any single ring across the FSR sub-band did not exceed 4.75 V. The full-width half maximum (FWHM) was measured to be 2 GHz, which is too narrow for data center applications, but will be corrected in future iterations of the switch design. Figure 3.11b shows the first MRR tuning independence of the second MRR across its own sub-band with respect to the first and third MRRs, which were held at roughly 2 V and 3.5 V, respectively. Independent tuning behavior is repeatable with the other MRRs. Due to the 100  $\mu\text{m}$  pitch between MRR pairs—which was set by the signal pad pitch—we did not observe any thermal crosstalk.

The additional loss of the rings is likely due to unoptimized coupling regions between the MRR pairs as well as between the MRR and the bus waveguides, as this was a first-pass design using the foundry process. Unoptimized coupling regions between bus waveguides and ring and between rings can result in loss waveguide coupling and therefore loss in the transmission of an add-drop ring pair. A mistake in the design led to sub-optimal field coupling coefficients ( $\kappa$ ) between the bus waveguide and the ring and between the rings in the ring pairs. which led to a narrower response than expected. This mistake has been fixed in later designs and the FWHM is expected to be 35 GHz,



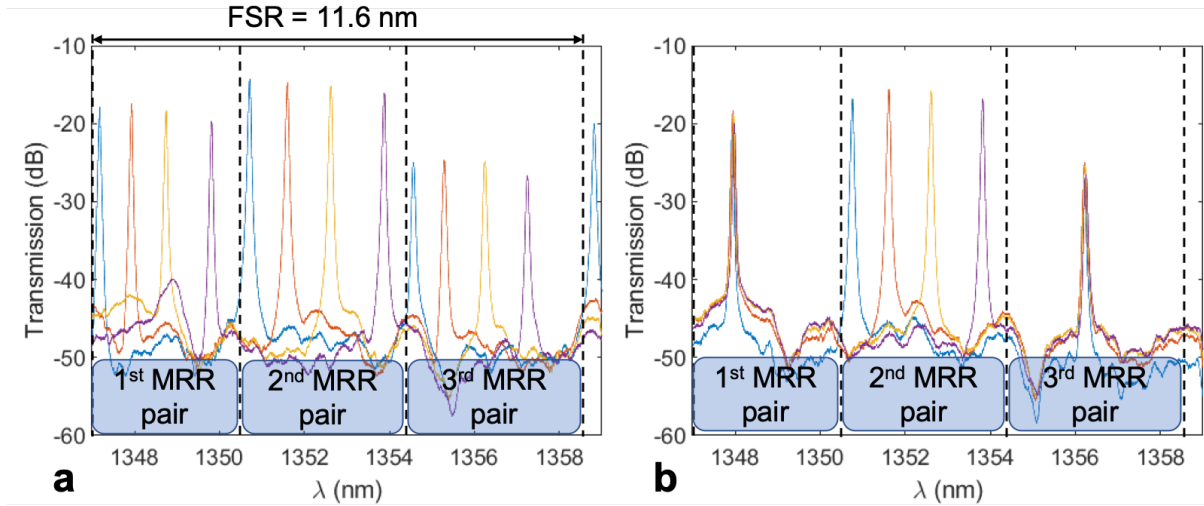


Figure 3.11: (a) shows switch operation across one full FSR using the heaters, while (b) shows the second MRR tuning independently of the first and third MRRs.

as shown in Figure 3.12. The original coupling coefficients  $k_1$  and  $k_2$  were 12.6% and 0.80%. These were revised to be 0.44% and 11.2%.

### 3.3.4 Wavelength Assignment

Here we provide a brief description of our algorithm at a high level by way of an example. Consider a  $4 \times 4$  WDM switch of the type discussed above, capable of dropping up to  $L = 3$  wavelengths per cross-point. Let the number of wavelengths per port be  $M = 6$ . Thus, the switch has three MRRs per cross-point, each tunable to  $k = ML = 2$  wavelengths within its own spectral partition, each covering one third of the FSR. Let the desired input-output wavelength connectivity demand be represented by the graph of Fig. 3.13(a), which shows connectivity varying from zero to three. To find the optimum wavelength assignment, our algorithm applies a form of the Ford-Fulkerson maximum-network-flow algorithm [110], modified with the flow restricted through simple (i.e., having no multiple edges)  $k$ -regular bipartite graphs (i.e., bipartite graphs with all vertices of degree- $k$ ), where  $k = 2$  in our example. In effect, our algorithm is used for






Figure 3.12: The expected spectrum of the redesigned 1310 ring switch

factorization of the bipartite demand graph, Fig. 3.13(a), into  $L$   $k$ -regular graphs, known as  $k$ -factors [111]. In our case, the factorization leads to the solution of Fig. 3.13(b), represented by three degree-2 bipartite graphs, where, with  $l = 1, 2, 3$ , the solid and dashed lines, respectively, represent connectivity at wavelength  $\lambda_l$  and  $\lambda'_l$ , both falling within the spectrally partitioned corresponding to MRR- $l$ , as depicted in Fig. 3.6(c). Thus, Fig. 3.13(b) represent the desired, blocking-free, spectrally partitioned, wavelength assignment solution for the wavelength demands of Fig. 3.13(a). Conventional wavelength assignment algorithms (e.g., those used in [101], [74]) expand the wavelength demand graph into single-wavelength components (called graph matchings). In fact, our solution of Fig. 3.13(b) can be further expanded into such single-wavelength components (Fig.

3.13(c)), which in turn can be pairwise combined to get back Fig. 3.13(b). But, there is no guarantee that a standard algorithm will lead to such a solution. For example, Fig.3.13(d) shows a typical single-wavelength expansion using a standard algorithm. It can be verified that there are no pairwise combinations of the components of Fig. 3.13(d) that yields a simple 2-factor solution similar to Fig. 3.6(b). All of the possible pairwise combinations will yield 2-factors with at least one multiple edge, which violates the spectral partitioning assumption. That is, in the standard wavelength assignment algorithm depicted in Figure 3.13(d), there are three connections from input 3 to output 4, one using the blue band and two using the green band, with none using the red band. This would require one of the red rings to be tuned to the green band. On the other hand, in our algorithm depicted in (b) and (c), each of the three connections uses a different band, thus each of the three rings stays in its own band

### 3.3.5 Conclusion and Future Work

We have successfully realized and demonstrated a  $4 \times 4$ ,  $L = 3$  spectrally-partitioned switch controlled by a driver, and a compatible blocking-free, wavelength assignment algorithm. In addition, we have demonstrated operation of the first chip-scale wavelength-selective all-optical O-band switch. Future work includes testing larger number of paths through the switch with a custom 64-channel 10 V driver prototyped on a PCB. A 12-bit digital to analog converter (DAC) will be used to achieve smooth wavelength tuning.

Additional future work on the O-band crossbar includes correcting the coupling constants in the next designs as well as evaluation scaling up the work to a larger port count. For example, if the rings were scaled up to  $16 \times 16$ , the required space for the I/O would increase to larger than 4 mm of space on a side given  $127 \mu\text{m}$  pitch. More importantly, the number of signal pads required would increase from 96 signals to 1536 signals. These

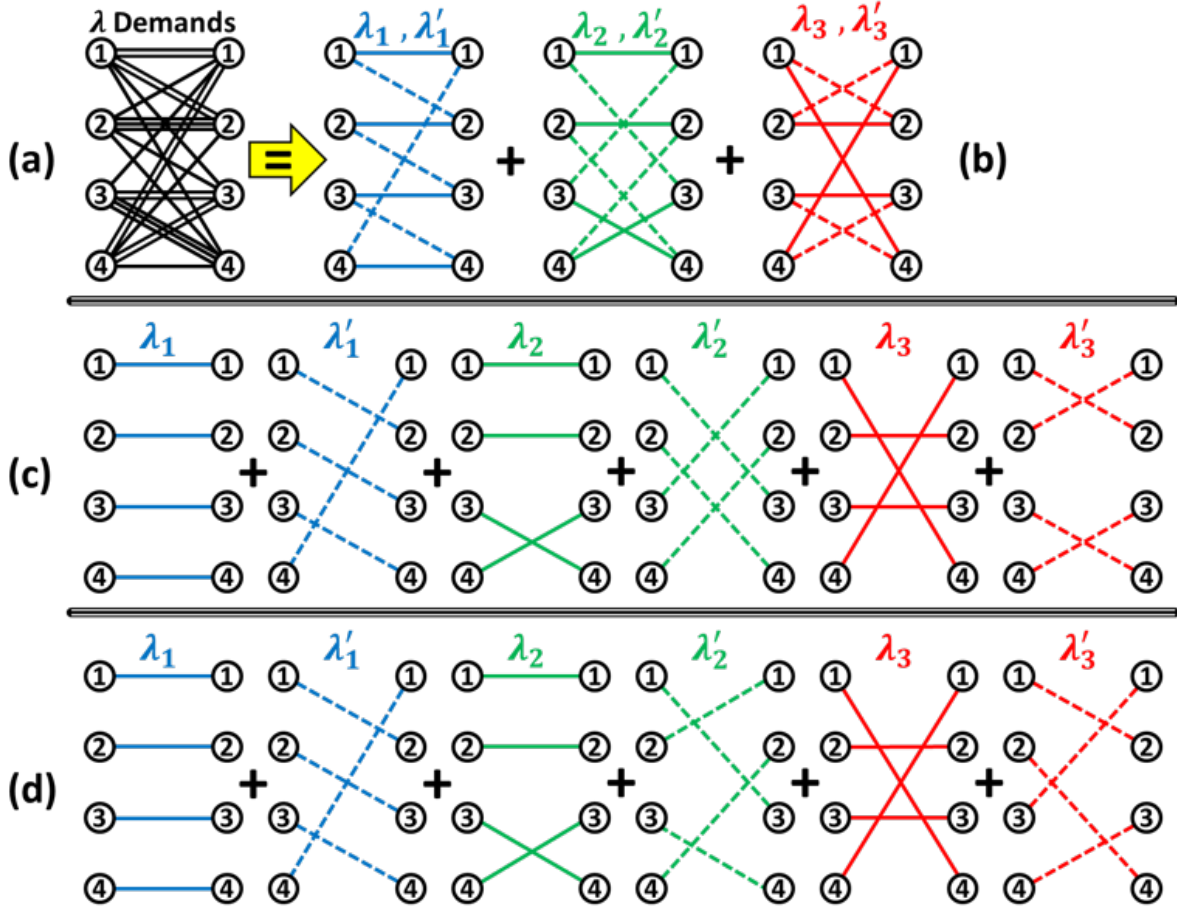


Figure 3.13: Description of our wavelength assignment algorithm. (a) A wavelength demand. (b) Resulting wavelength assignment using our algorithm, represented by simple 2-factor graphs. (c) Single-wavelength representation of (b). (d) A wavelength assignment using a standard algorithm, which would have required some of the MRRs to tune over the entire FSR.

pads would take up  $4.8 \text{ mm} \times 7.2 \text{ mm}$  based on a  $150 \mu\text{m}$  pitch. If a driver were to be flip-chipped onto a switch of this size, this pitch is compatible with IC nodes capable of delivering the maximum voltage ( $\approx 5 \text{ V}$ ) required to control the switch. Additional space is required for the numerous ground pads required to provide a good external ground connection given the number of signals. If the average power dissipation of each crossing is taken to be  $16.1 \text{ mW}$ —half the power required to tune across the whole FSR—a  $16 \times 16$  switch would dissipate an average of  $4.1 \text{ W}$ . This would require temperature monitoring

on the PIC and a corresponding circuit in the driver.

# Chapter 4

## RAMZI Switches

Work in this section has appeared in [112] and [68].

### 4.1 Introduction

We have described demonstrated wavelength-selective switches in a silicon photonics process in the previous chapter. These switches utilized serially coupled double ring pairs in a crossbar configuration. Some concerns inherent in this switch topology include path-dependent loss and complexity in controlling a higher radix switch. For the former, the difference between the shortest and longest path is  $2(N-1) \times (\text{loss per crossing})$ , while for the latter, the pad count scales as  $2LN^2$  for an  $N \times N$  switch with  $L$  ring pairs per crossing when the  $2 \times 2$  switch cells are arranged in the so-called Beneš configuration. Ring-assisted Mach-Zehnder interferometers (RAMZI) modulators were first proposed in [113] in order to overcome the non-linear response of Mach-Zehnder interferometers (MZIs) for analog applications. Since then, RAMZI switches have previously been proposed and demonstrated [43,114], but with only one ring pair per switching element. Since switching more than one wavelength would be required, in this chapter, we demonstrate a RAMZI

switch with  $L$  wavelength-selective elements per switch cell with path-independent loss.




Figure 4.1: (a) shows single RAMZI cell with two ring pairs, (b) shows a screenshot of the design for one of the rings, and finally, (c) shows an example of the paths taken through the switch.

The switch utilizes an element called a RAMZI, which is comprised of a Mach-Zehnder interferometer (MZI) with two all-pass rings on each arm, as shown in Figure 4.1(a). A Mach-Zehnder interferometer switch cell typically consists of two input waveguides that are joined either with a  $2 \times 2$  multi-mode interferometer (MMI) or directional coupler, split into two arms, and rejoined again, with either a  $2 \times 2$  MMI or directional coupler with two output waveguides [82]. For the switches presented in this thesis, directional couplers are used. The response of a 50/50 split directional coupler can be written in transfer matrix notation as thus:

$$T_{matrix} = \frac{1}{\sqrt{2}} \begin{bmatrix} 1 & j \\ j & 1 \end{bmatrix} \quad (4.1)$$

[83] A phase shifting element is used in one or both arms to control the phase at the second directional coupler where the light can recombine destructively or constructively, thus switching the light from one port to another. This forms the basis of all

Mach-Zehnder-based switches, though the phase shifting element differs between implementations. Mach-Zehnder-based switches have been investigated in a wide variety of switches [76, 109, 115–118]. If light is injected into the top left port, as shown in Figure 4.1(a), the top right port is known as the bar port, while the bottom right port is known as the cross port. [87]

In the case of a RAMZI, a ring is the phase-shifting element on both arms. Instead of utilizing an add-drop ring, as described in Chapter 3, an all-pass ring, as seen in Figure 4.1(b) is used. The resonance condition remains the same as that described in equation 3.1. The effective phase shift due to the ring resonator is given by:

$$\Phi = \pi + \phi + \arctan \frac{r \sin \phi}{a - r \cos \phi} + \arctan \frac{ra \sin \phi}{1 - ra \cos \phi}, \quad (4.2)$$

where  $\phi = \beta L$ , and  $r$  and  $a$  are the self-coupling coefficient and single-pass amplitude.  $\beta = \frac{2\pi n_{eff}}{\lambda}$ , known as the wavenumber, and  $L$  is the round-trip length. It is assumed that the coupling region of the rings is lossless (i.e.  $r^2 + k^2 = 1$ ). The phase change across the resonance must be smooth and is dependent on the relationship between  $r$  and  $a$ . If  $r > a$ , this condition is known as overcoupling, while if  $r < a$ , this is known as undercoupling. At  $r = a$ , a condition known as critical coupling, the phase shift is an abrupt  $\pi$  phase shift. Figure 4.2 shows the phase change across the resonance for the three conditions. The rings in the RAMZI switch were designed to be in the over-coupled region. The racetrack rings were designed such that they are both compact had have large FSR [85].

The power transmission response of the output of an all-pass filter is

$$E_{ring} = \frac{a^2 - 2ra \cos \phi + r^2}{1 - 2ra \cos \phi + (ra)^2}. \quad (4.3)$$




Figure 4.2: The phase change across the resonance of a ring is shown for the over-coupled ( $r > a$ ), critically coupled ( $r = a$ ), and undercoupled ( $r < a$ ) conditions. To generate these plots,  $r$  was set to 0.85.

[84] If there are two rings in series, the output after the second ring is  $E_{arm1} = E_{ring1}E_{ring2}$ .

The bar and cross states can be calculated by cascading the responses in the transfer matrix:

$$\begin{bmatrix} E_{bar} \\ E_{cross} \end{bmatrix} = \frac{1}{2} \begin{bmatrix} 1 & j \\ j & 1 \end{bmatrix} \begin{bmatrix} E_{arm1} & 0 \\ 0 & E_{arm2} \end{bmatrix} \begin{bmatrix} 1 & j \\ j & 1 \end{bmatrix} \begin{bmatrix} E_{in} \\ 0 \end{bmatrix} \quad (4.4)$$



The normalized spectrum for the cross and bar states can be found to be

$$\left| \frac{E_{bar}}{E_{in}} \right|^2 = \frac{1}{4} [ |E_{arm1}|^2 + |E_{arm2}|^2 - 2E_{arm1}E_{arm2} \cos(\theta_{arm1} - \theta_{arm2}) ] \quad (4.5)$$

$$\left| \frac{E_{cross}}{E_{in}} \right|^2 = \frac{1}{4} [ |E_{arm1}|^2 + |E_{arm2}|^2 + 2E_{arm1}E_{arm2} \cos(\theta_{arm1} - \theta_{arm2}) ] \quad (4.6)$$

where  $E_{arm2}$  is the E-field right before the directional coupler of the other arm, and  $\theta_{arm1,2}$  is the cumulative effective phase shift due to both rings in each arm. In other words,  $\theta_{arm1} = \phi_{ring1} + \phi_{ring2}$ .

The resonance of the ring can be altered by changing the index of refraction within the waveguide. This can be done by increasing the temperature of the ring. Thus, a heater comprising of doped silicon was added to the center of the ring. Figure 4.1b is a screenshot of the design of one of the rings, which include a heater integrated in the center of the ring. Heat is transferred through the  $1\mu\text{m}$  of silicon dioxide that separates the waveguide from the heater. The heater was separated from the waveguide by  $1\mu\text{m}$  to minimize interference with the waveguide mode. In addition, a thermal phase tuner was placed on each arm of the MZI in order to symmetrize the path length. Directional couplers, crossings, and edge coupler designs provided by the foundry were used in the design. [114]

The switch was designed in a Beneš configuration, with each switching element comprised of two ring pairs—with one all-pass ring from each pair on each arm of a symmetric MZI. The Beneš configuration was chosen due to its rearrangeably non-blocking property [112]. Figure 4.1c shows the path of the optical signal taken through the  $4 \times 4$  switch from the first input to each of the four outputs. Since the MZI has symmetric arms, if the corresponding ring pairs are tuned to the same wavelength  $\lambda_s$ ,  $\lambda_s$  is sent to the bar port. This occurs because the phase changes by  $\pi$  across the resonance [119]. Therefore, when  $\lambda_s$  reaches the second directional coupler, it destructively interferes such that it

leaves the switch cell through the bar port.

## 4.2 Experimental Methods

The  $4 \times 4$  RAMZI switch was fabricated the AIM Photonics process and packaged on a custom printed circuit board (PCB). A total of 45 wirebonds were made, including ground connections, for each of the rings and thermal MZI phase tuners. The size of the switch is 0.5 mm by 2.1 mm with an additional  $900 \mu\text{m}$  by  $650 \mu\text{m}$  required for the optical I/O. To test the switch, a C-band Yenista Tunics T100S tunable laser was used in conjunction with the Yenista CT-440—a passive component analyzer to take the spectral data of the switch. Figure 4.3 shows the experimental setup, while Figure 4.4 shows a closeup of the switch. Since the edge couplers were spaced by  $127 \mu\text{m}$ , the switch was tested using a pair of cleaved single mode fibers. The polarization of the incoming light was controlled using an external polarization controller.

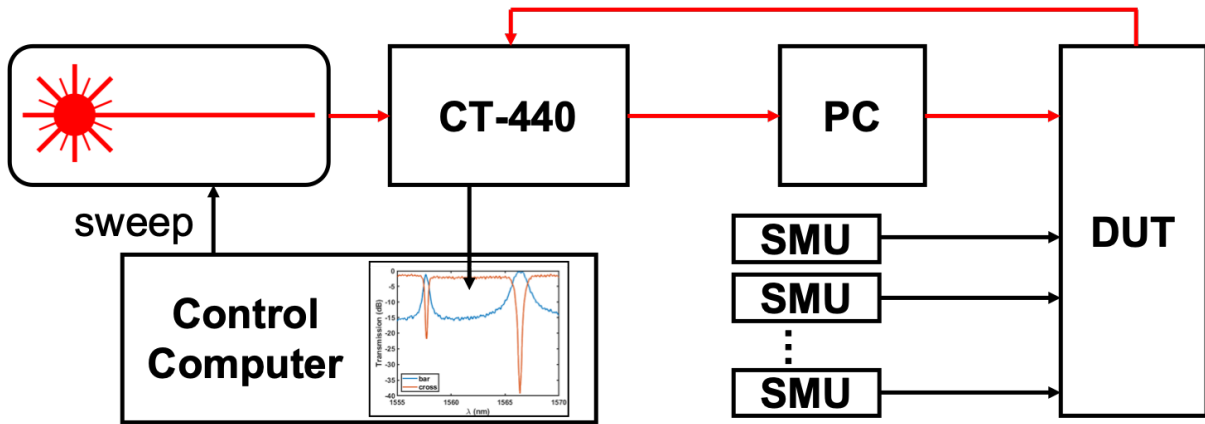


Figure 4.3: This figure shows the test bed schematic to test the switch. PC=polarization controller, DUT = device under test, SMU = source measurement units. The red arrows represent the optical path, while the black arrows represent all the electrical connections made in the experimental setup.

A 96-channel custom-designed driver was used for switch tuning, the 96-channels were realized using six banks of commercial off-the-shelf (COTS) 16-channel 12-bit Digital to




Figure 4.4: (a) shows the wirebonded switch on a custom PCB, while (b) shows a closeup of the chip on which the switch is located. In (b), the RAMZI switch area is outlined in red, while the area right below the box are the edge couplers.

Analog Converters (DACs) assembled on a printed circuit board (PCB). A Graphical-User-interface (GUI) designed in MATLAB controlled both the DACs and individual channels through a serial to parallel (SPI) bus. DAC outputs and the photonic IC were connected via a cable. The DAC resolution was dictated by the ability to enable both fine tuning of the wavelength and large dynamic range. The PCB was carefully designed to

avoid noise coupling due to electrical switching. This was achieved by separating analog and digital signals and ground planes.

We chose LTC2668 as the DAC with a 10 V full-scale range to allow for large dynamic range, while the 12-bit resolution (2.44 mV step size) enabled fine wavelength tuning. An Arduino board was used to control the DACs through the GUI.

The  $2 \times 2$  RAMZI switch cell with two ring pairs was measured first using source measurement units (SMUs), the results of which are shown in Figure 4.5. The response of the grating couplers has been calibrated out. Two cleaved single mode fibers (SMFs) were used to vertically couple to the chip, while probe needles were used to inject current into the heaters. The test switch was not wirebonded. The layout of the test switch is the same as a single cell in the larger  $4 \times 4$  switch; thus the test cell is representative of a switch cell within the larger switch. The small ring pair resonance is centered at 1558 nm, while the large ring pair resonance is centered at 1566 nm. In the bar state, the crosstalk for the small ring pair is -20 dB, while the crosstalk for the large ring pair is -40 dB. In the cross state, the crosstalk is -20 dB and -19 dB for the small and large ring pairs, respectively. The bar state consumed 5.0 mW, while the cross state consumed 7.9 mW. Due to process variations, the ring pairs were tuned about the most redshifted ring within each pair since thermal tuning redshifts ring responses. This represents the minimum power needed to achieve the correct switching behavior. The narrow response and relatively low crosstalk between the bar and cross outputs of the small ring pair centered at 1558 nm is due to an unoptimized coupling region. The tuning efficiency of the rings was found to be 2.06 mW/nm and 2.62 mW/nm for the large and small rings, respectively. The effective Q factors were found to be 1300 and 3900 for the large and small ring pairs. The pass-by loss was found to be -1.1 dB and -0.19 dB for the small and large ring pairs.

To measure the  $4 \times 4$  switch, the optical signal was injected into the first input, and

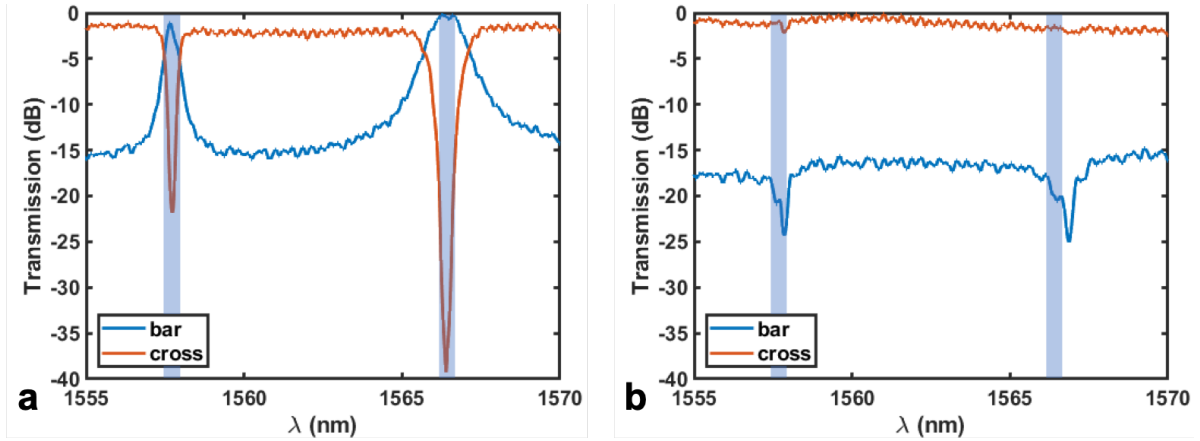


Figure 4.5: The bar state of the both the ring pairs is shown in (a) while the cross state for both ring pairs is shown in (b). The wavelengths being switched are highlighted in blue.

Output	Small Ring Pair	Large Ring Pair
1	-15 dB	-13 dB
2	-14 dB	-14 dB
3	-17 dB	-21 dB
4	-17 dB	-22 dB

Table 4.1: Measured crosstalk of  $4 \times 4$  RAMZI switch

the response was measured at all four outputs. The responses were overlaid to show wavelength-routing of two different wavelengths using the different ring pairs. Paths through the  $4 \times 4$  switch from the first input to each of the outputs were first controlled using SMUs. Figure 4.6 shows the responses when tuning both ring pairs to each of the four outputs, with the loss of the edge couplers removed. The switch configuration consumed 25.2 mW, 27.9 mW, 30.26 mW, and 31.1 mW to go from input 1 to output 1, output 2, output 3, and output 4, respectively. Similar to the  $2 \times 2$  RAMZI test cell, the rings were tuned to the most redshifted ring within the two ring sizes across the switch. That is, if there was a small ring that was significantly redshifted from the rest of the small rings, all the small rings in each switch cell in the path were tuned about the response of the outlying small ring. Due to the large number of rings in the switch, not all the rings in the switch were controlled. In fact, only those rings in the path were

tuned. Two rings were controlled in the first stage in order to symmetrize the response and direct all the light in the spectrum to the cross port. All four rings in bottom switch cell in the second stage, as seen in Figure 4.1c, were controlled in order to switch the light either to the top cell or bottom cell in the third stage. Only the four rings on the cell corresponding to the desired output port were controlled with the remaining SMUs. In other words, if the wavelengths of interest were being output at port 3, only the rings in the bottom cell of the third stage were controlled. Likewise, if the wavelengths were being routed to port 2, only the rings in the top cell of the third stage were controlled. If more SMUs were available, the other cell in the third stage could be controlled and the crosstalk performance of the corresponding outputs improved with respect to the desired output port. In Figure 4.6, a different region of the spectrum was chosen that allowed the responses of the ring pairs to be more clearly separated from each other compared to the single-stage results presented in Figure 4.5. The small ring pair response is centered at roughly 1548 nm, while the large ring pair response is centered at 1558 nm. Again, the narrow response of the small ring pair is due to the unoptimized coupling section of the ring. Note that compared to Figure 4.5, the small and large ring pair responses have been flipped. Due to process variations, the central resonances of the ring pairs were shifted from the  $2 \times 2$  RAMZI test cell, thus we chose the spectrum that was shifted half an FSR from the original spectrum. The crosstalk is measured as the minimum extinction between the output port of interest and the output port with the next greatest amount of output power. Table 4.1 includes the crosstalk in dB for each ring pair and the output port. The smaller ring pair has an FSR of 23 nm, 0.4 nm bandwidth, and an average resistance of  $1.6 \pm 0.4$  k $\Omega$ , while the larger ring pair has a 19 nm FSR, 1.2 nm bandwidth, and an average resistance of  $3.4 \pm 0.9$  k $\Omega$ . The increased insertion loss seen in Figures 4.6(a), (b), and (d), compared to 4.6(c) is due to suspected increase in the loss in the bar port. This phenomenon was seen in several other chips that were tested.

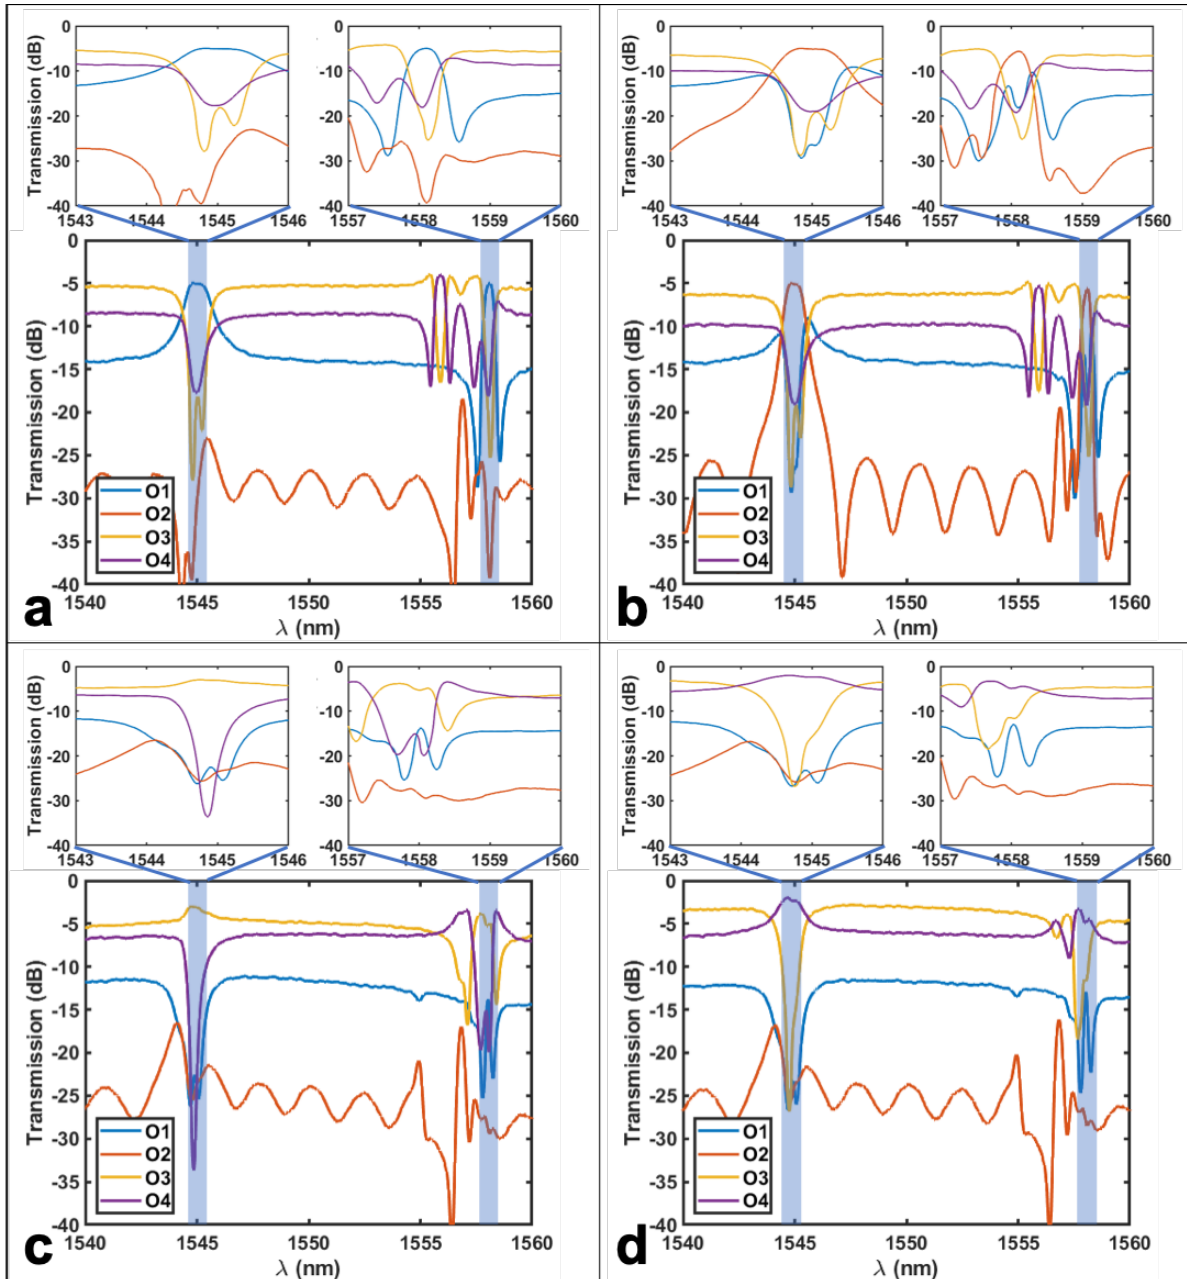


Figure 4.6: The outputs of each of the switch states for both ring pairs is shown. The wavelengths are routed from input 1 to (a) output 1, (b) output 2, (c) output 3, and (d) output 4. An inset for each switch configuration is also shown. The data was obtained utilizing SMUs to tune the rings.

The experiment was repeated with the 96-channel driver. The resulting spectra are shown in Figure 4.7, in which the driver can control the switch accurately. The spectra



were normalized with respect to the output peaks of the respective control mechanisms. The control PCB (including Arduino board) consumes 400 mW under unloaded conditions. It consumes 700 mW while driver the 14 heaters. The DACs contribute the most to the power consumption (60%) with each DAC approximately consuming 10% of the power.

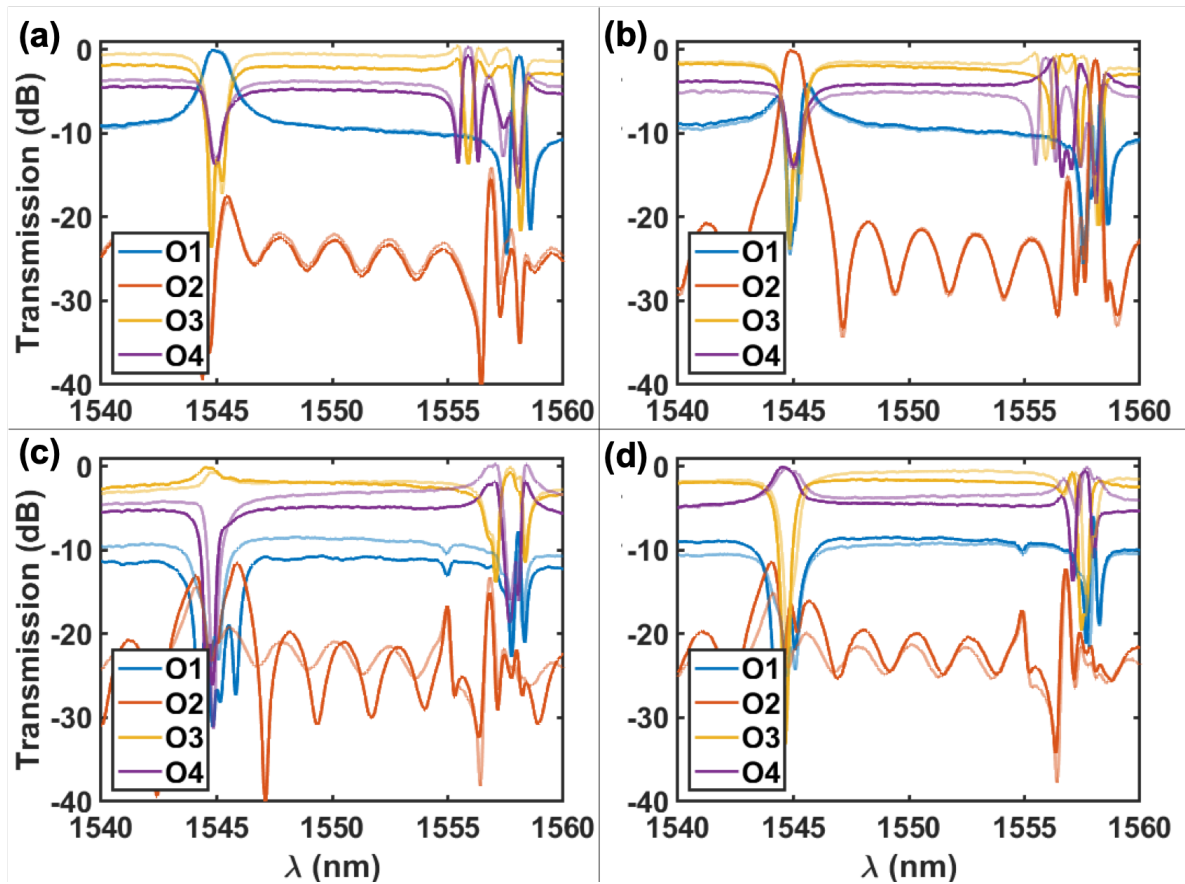


Figure 4.7: The outputs of each of the switch states for both ring pairs is shown. The wavelengths are routed from input 1 to (a) output 1, (b) output 2, (c) output 3, and (d) output 4. The dark lines represent data taken with the switch controlled with a driver, while the more transparent lines represent data taken for the same switch state with the SMUs.

To demonstrate independent switching between the ring pairs, Figure 4.8 shows wavelength-selective switching to two different outputs using the driver. 1545 nm light is routed to output 3, while 1548 nm light is routed to output 2. In this configuration,



the driver consumed 726.6 mW. Output 3 exhibits -15 dB of crosstalk, while output 2 has -13 dB of crosstalk.




Figure 4.8: The resulting spectra of the driver-controlled RAMZI when controlling the rings from input 1 to two different outputs, ports 2 and 3.

### 4.3 Scaling Up the Switch

In Chapter 3, we demonstrated a  $4 \times 4$  wavelength-selective crossbar switch. However, a major drawback to these switches is that having multiple drops per cross-point does not lend itself well to scaling up the number of ports and particularly, the number of electrical connections needed to control the switch. We assume that there is one signal pad for every ring and that the other side of the ring heater is tied to an on-chip ground plane. The signal pad count scales by  $2LN^2$  for an  $N \times N$  crossbar switch with  $L$  ring pairs per cross-point. On the other hand, for an  $N \times N$  RAMZI switch with  $L$  ring pairs

per switching element, the pad count scales by  $2L(N \log_2 N - \frac{N}{2})$ . It is important to note that the ground plane will also require pads for external connections, but the exact number of ground pads is dependent on the PIC and external IC co-design requirements, and thus difficult to predict. However, the number of ground pads is smaller than the number of rings, and scales much slower than the number of signal pads.

To illustrate, we compare a  $4 \times 4$  switch with 3 ring pairs per switching element in the crossbar and RAMZI configurations. It is important to note the distinction in configuration between the two ring-based switches. In a crossbar switch, we use  $L = 3$  ring resonator pairs per crosspoint for a total of  $2L$  per cell. Such a crossbar switch would require 96 signal pads, while such a RAMZI switch requires only 36 signal pads. The longest path and shortest paths through the crossbar switch requires passing by 20 and 2 rings, respectively, while only 9 rings are passed when taking any path through the RAMZI. If we consider next a  $16 \times 16$  switch, again with 3 ring pairs per switching element, the crossbar switch would require 1536 signal pads, while the equivalent RAMZI would require 336 signal pads. The longest path through the crossbar switch would require passing by 92 rings, while the number of rings passed in the shortest path remains the same. In the case of the RAMZI switch, the optical signal would pass only 21 rings regardless of the path taken through the switch.

An additional drawback of such crossbar switches is that the loss through the switch is dependent on the path taken through the switch. Indeed, this is due to the unequal lengths of the paths through the switch and therefore a different number of photonic elements that the signal must pass which all contribute to the loss. In the crossbar switch, the difference between the shortest and longest paths is  $2(N - 1)$  cross-points. On the other hand, a RAMZI switch with the switching elements in a Beneš configuration will incur the loss of passing through  $2 \log_2 N - 1$  stages. Assuming that the pass-by loss of a single ring is  $A$ , the pass-by loss per cell is  $AL$ , where  $L$  is the number of ring pairs.

(This assumes that the pass-by loss of each ring is the same on both arms.) Then, the total pass-by loss of the rings in the full switch PIC is  $AL(2 \log_2 N - 1)$ . The average of the pass-by loss of the rings is 0.6 dB. If we use this and let  $L = 3$  the loss for an  $L = 3$  cell is 1.8 dB. For a  $4 \times 4$  RAMZI, the loss is 5.4 dB. In the crossbar switch, the loss to pass by a ring is lower than in the RAMZI because the light that passes by does not meet the resonance condition of the ring it is passing by. Furthermore, to complete the switching, there is loss associated with the cross state, which we shall denote as  $B$ . Therefore, the loss in the longest path of an  $N \times N$  crossbar switch is calculated as  $AL(2N - 1) - A + B$ . In  $4 \times 4$  crossbar switch demonstrated in [99] and [100],  $A$  is 0.3 dB,  $B$  is 5.1 dB. The loss in a  $4 \times 4$   $L = 3$  switch is 11.7 dB. If scaling up the port count to  $16 \times 16$ , the pass-by loss in the longest path in the crossbar switch is 35.5 dB, while for the RAMZI the pass-by loss is 12.6 dB. Since the coupling regions of the rings pairs have not been optimized, we expect that the loss number due to passing by the rings would reduce with improved ring design.

In terms of spectral bandwidth, assuming that all the rings are designed with the current response of the large ring pair, up to 8 wavelengths, and thus 8 ring pairs can be supported, with an optical bandwidth of 150 GHz. The resulting crosstalk, assuming the ring pairs are spaced at 1.2 nm (150 GHz) would be around -25 dB to the neighboring rings. These specifications are sufficient to support 100G QPSK and 100G PAM-4.

Due to the change in architecture, the power consumption of a RAMZI will not increase as rapidly as a crossbar with an increase in port count. If we let  $X$  be the average tuning distance in nm, and  $H$  be the heater tuning efficiency in mW/nm, the power consumption is  $2 \times H \times L \times X \times N(\log_2 N - \frac{1}{2})$ . If we let the average tuning distance be  $\frac{\text{FSR}}{2L}$  (half of an FSR partition), the above expression reduces to  $H \times \text{FSR} \times N(\log_2 N - \frac{1}{2})$ . We have previously demonstrated that spectrally partitioned switches do not need to tune the full FSR, while preserving the wavelength-routing flexibility [99, 100]. By contrast,

the power consumption of a square spectrally partitioned crossbar switch is  $H \times \text{FSR} \times N^2$ . A large FSR is desirable to maximize the number of wavelength channels that can be routed through the switch. Therefore in order to reduce the total power consumption of the switch, the key parameter that should be improved is the heater tuning efficiency, reducing the value of  $H$ . The average tuning efficiency was found to be 2.34 mW/nm, thus the total power consumption of a  $4 \times 4$  RAMZI is 267 mW as compared to over 700 mW for a  $4 \times 4$  crossbar switch. For a  $16 \times 16$  switch, these numbers become 2.7 W and 11.3 W for the RAMZI and crossbar, respectively. These power consumption numbers can be reduced; in [80], we have demonstrated rings with tuning efficiencies of 1.76 mW/nm—a 25% reduction—by connecting the heater to the waveguide with undoped shallowly etched silicon for higher thermal conductivity. Furthermore, it has been shown that undercutting photonic devices in a silicon process reduces the power consumption of heaters by over an order of magnitude [120]. This conclusion suggests that a 64-port RAMZI switch, with improved heater designs with undercuts, could achieve power consumption of about 1 W.

Discussions of power consumption lead to questions of energy efficiency. The energy efficiency of the switch can be found by dividing the power consumption by the aggregate data rate passing through the switch. Assuming that there are as many ring sizes for each wavelength channel of a WDM input signal, for the RAMZI, the energy efficiency is  $\frac{H\text{FSR}(\log_2 N - \frac{1}{2})}{LB_r}$ , where  $B_r$  is the bit rate per wavelength. For the crossbar switch, the energy efficiency is  $\frac{H\text{FSR}N}{LB_r}$ . For a heater efficiency of 1.76 mW/nm, FSR of 20 nm, and WDM input signals with 4 200 Gb/s channels, a  $16 \times 16$  RAMZI would have an energy efficiency of 0.156 pJ/b, while a crossbar switch of the same size would have an energy efficiency of 0.68 pJ/b. While this will add to the link efficiency analysis findings demonstrated in chapter 2, the links will still not exceed 10 pJ/b.

## 4.4 Conclusion and Future Work

We have successfully demonstrated switching operation of a  $4 \times 4$  RAMZI switch with two ring pairs per cell to enable switching two wavelengths independently. The switching time is estimated to be  $\sim 16\mu\text{s}$  based on [105], which features rings of a very similar design. Future work includes moving the switch wavelengths within the sub-partition to give a more direct comparison of switch power efficiency with respect to crossbar switches with partitioned FSRs [99, 100]. Further design improvements include optimizing the coupling region in the smaller ring pair to improve crosstalk and width of the passband, integrating the driver more directly through monolithic integration or other dense integration means, and more efficient heaters—the latter easily implemented by connecting the silicon heater to the waveguide with a slab waveguide region for efficient heat transfer rather than keeping the waveguide and heater separated by  $1\mu\text{m}$  of silicon dioxide.

Immediate research directions is to measure data through the switch and show any degradation in the SNR as a result of traversing the switch. Additionally, one can work to improve the design of the RAMZI switch. One action that can be taken is scaling up the RAMZI switch to a higher port count with more ring pairs per switch cell. In addition, to ease the burden of the initial calibration, optical taps can be implemented after each cell. This tap can be in the form of another bus waveguide to each ring pair to a grating coupler, as shown in figure 4.9, transforming the all-pass rings to add-drop rings. It is important to note that the coupling coefficient to the second bus waveguide must be very small, such that the response of the RAMZI switch cell is undisturbed, as seen in Figure 4.10.

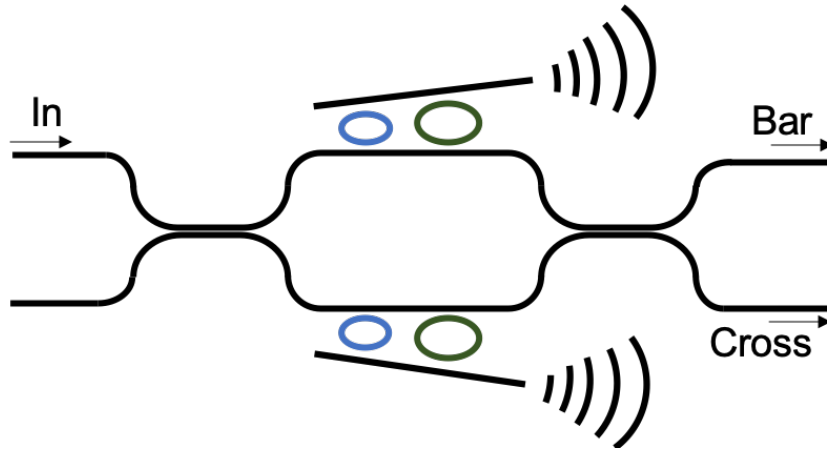


Figure 4.9: A  $2 \times 2$  RAMZI cell with taps by converting the all-pass rings to add-drop rings. The taps can be routed to a grating coupler to monitor during the initial calibration.

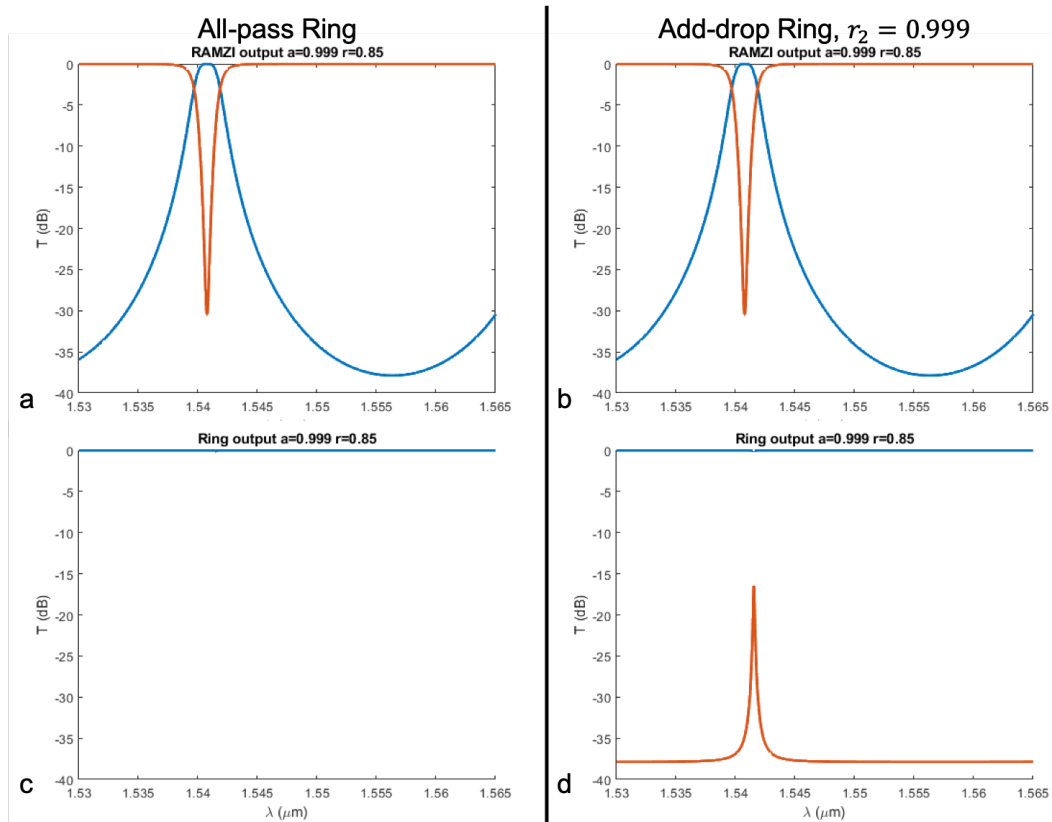


Figure 4.10: (a) shows the  $2 \times 2$  RAMZI response utilizing an all-pass ring, while (b) shows the RAMZI response utilizing an add-drop ring where the self-coupling coefficient of the second bus waveguide is large. (c) and (d) show the ring response for the respective cases.

# Chapter 5

## Summary and Outlook

In this thesis, I have described and presented an energy efficiency analysis of next generation analog coherent links for data center links at 200 Gbps/wavelength projecting link energy efficiencies of sub-10 pJ/b, which encourages a movement from direct detection to analog coherent detection. The energy efficiency analysis included a 13 dB link budget between the transmitter and receiver to allow for other optical components to be implemented to flatten the network architecture, such as an optical switch.

Two optical switch types were presented: a crossbar switch and a ring-assisted Mach-Zehnder interferometer (RAMZI) switch, both with the ability to route multiple wavelengths at each crosspoint or switch cell, respectively. A C-band and an O-band ring-based switch was demonstrated with multiple ring pairs per crosspoint to partition the FSR, reducing the overall power needed to tune over a full FSR and thereby the switch overall. The C-band switch was measured to have an average power dissipation of roughly 70 mW per crosspoint, and projected to have an overall maximum power dissipation of 1.4 W including the PCB driver power. The O-band switch was measured to have 16 mW per crosspoint, due to the improved ring heater design, and projected to consume 658mW including the PCB driver power. The energy efficiency of  $16 \times 16$  crossbar switches is

projected to be 0.68 pJ/b, assuming 4 wavelengths at 200 Gb/s per port.

While the crossbar switch is easier to design and implement, a major concern is how to package and drive a large number of rings as the port count increases. A RAMZI in a Beneš configuration on the other hand, significantly reduces the number of electrical connections for the same port count compared to the crossbar switch, as well as providing lower loss, due primarily to the change in switch architecture. The  $4 \times 4$  RAMZI switch consumed 667 mW of power, and is projected to consume 0.16 pJ/b, assuming 4 wavelengths at 200 Gb/s per port.

Compared to current commercial data center switches, all-optical switches show promise in further energy efficiency reduction in next generation data center links. The energy efficiency of the switches presented in this dissertation are one to two orders of magnitude lower than typical commercial data center switches. However, this is at the cost of reduced SNR, increased error rate due to the inability to regenerate the optical signal, and inability to multicast the optical signal.

## 5.1 Future Directions

The switches presented in this thesis were prototypes. However, to be deployed in a system, they must satisfy the following conditions:

- support higher aggregate data rates per port compared to current data center switches;
- have as many or more ports as current data center switches;
- have more uniform rings across the switch chip;
- have lower loss per element; and



- have greater robustness, including polarization and temperature insensitivity.

Optical switches must be able to support higher aggregate data rates compared to current data center switches. Since OEO conversion does not need to occur in optical switches, they are comparatively agnostic to data rate and data types. As long as the crosstalk between neighboring ports is low, the effect of optical switches on data passing through the switch should be minimal. In this aspect, optical switches have a clear advantage over electrical switches. It should be noted however, that this behavior must be confirmed for the switches presented in this thesis.

Both the RAMZI and crossbar switches presented in Chapters 3 and 4 were both  $4 \times 4$ . As shown in Table 1.2 in Chapter 1, typical port counts for data center switches range from 6 to 64. There will be challenges associated in increasing the radix, including the need for a different method of controlling the switch, such as moving from a PCB driver as described in Chapter 3 and 4 to a more compact package with a custom driver IC as a proof of concept that photonic switches can be physically integrated into a larger system. A more compact package can be in the form of say, a silicon photonic interposer—such as one offered by AIM Photonics—on which the PIC and the custom driver ICs (EICs) would be flip-chipped into a silicon chip.

As the port count of the switch increases, so too must the uniformity of rings across a chip. Currently, due to process variations, rings of nominally the same size can have slightly differing resonant wavelengths and resistances on a single chip. This variability in rings will increase the power consumption of the switch to compensate. Thus, rings that are more process-tolerant must be designed. Just as pressing is the need for lower-loss switch elements, especially when the radix increases. Though some of the requirement for lower-loss switch elements in a wavelength-selective switch was mitigated by arranging RAMZI switch cells in a Beneš configuration to reduce the number of switch cells passed

through for the same port count, low-loss switch elements is still of utmost importance in designing future iterations of a wavelength-selective switch.

Polarization remains a relatively unexplored facet of the switches were presented in this thesis. While there has been work done on polarization independent switches [121], these switches were not wavelength-selective. Adding polarization diversity to the switch would further enable implementation in a data center, especially those in the vein of the ACD links presented in Chapter 2, in which polarization was used to increase the bitrate of the links. A major challenge that would arise in co-packaging PICs and EICs in a single package would be the temperature of the PIC and EICs. For the former, the operating temperature is especially important due to the temperature sensitive nature of the microrings. Wavelength-locking in switches studied in [122] has shown that the resistance of the heater in the ring can be used a temperature monitor. Thus having a low-power monitoring circuit on the EIC with minimal self-heating could allow for minimally-cooled or uncooled operation of the switch.

Despite several challenges that must be overcome before photonic switches can be integrated into physical systems, optical switches provide many advantages over electrical switches. As Ballani et al. from Microsoft outlined in [123], the two main challenges facing switches to be integrated into a larger system include cost and reliability. If the challenges I have outlined here can be overcome, Si photonic switches will be deployed due to the low cost and high level of maturity of Si manufacturing compared to other material platforms capable of waveguiding.

# Appendix A

## Test Structures

When designing devices, it is important to have a few test structures in order to characterize each part of the desired device. Unlike circuits with discrete components, one cannot test individual components of an integrated circuit. Thus, including test structures can help debug and manage expectations of the device performance. Furthermore, when designing in a new process at a foundry, or designing in a foundry process for the first time, the final fabrication run parameters may not be as promised at the beginning of the design phase. These parameters can include the doping levels of the waveguiding material or the etch quality of the narrowest gap between adjacent waveguides. Finally, foundries may be unwilling or unable to share information that is required to design a successful device. In this section, I will focus on lessons I have learned while designing in new silicon photonics multi-project wafer (MPW) fabrication runs at AIM Photonics and TowerJazz as well as in somewhat more established processes such as the 90WG process at GlobalFoundries and the SG25-PIC process at IHP Microelectronics. Note that I will be assuming that I am assuming a silicon-on-insulator (SOI) process with full silicon height of around a couple hundred nanometers.




Figure A.1: An example of a TLM design. The dark grey squares represent the pads, while the lighter grey represent the doped silicon region. It is often easiest to have one large piece of silicon under all of the pads in the TLM, with highly doped regions right under the contacts.

## A.1 General Test Structures

General test structures are those that can be used to characterize the process regardless of device.

### A.1.1 Electrical

The main electrical test structures are known as Transmission Line Measurements (TLMs). This helps characterize the sheet resistance of the doping level as well as the resistance of the pads. If the latter must be known—say to drive high-speed photonic devices—the pad structure from the top metal down to the connection to the silicon should be copied or very similar to how the connection is made in the device. To design this electrical test structure, one varies the distance between adjacent pads such that one measures different silicon sheet resistances, as shown in Fig. A.1. To extract sheet resistance, one plots the resistance versus length and finds the slope of the line. The y-intercept represents twice the pad resistance, as the current must pass through two pads in making the measurement. [124] Varying the distances from 10–100  $\mu\text{m}$  in 10  $\mu\text{m}$  increments is a safe design, however if space-constrained, using 5 increments to go from 20–100  $\mu\text{m}$  or 10–50  $\mu\text{m}$  should be sufficient to obtain the required information.

### A.1.2 Optical

The most basic optical structure in integrated photonics is the waveguide. Thus, most of the optical test structures described in this subsection are passive—that is, they do not require any electrical current to change the properties of the waveguide. The first basic test structure that must be included is the optical I/O test structure. This can take the form of a grating coupler to grating coupler test structure or an edge coupler looped back to another edge coupler on the same side. For test structures (and single ring structures), it is especially important to make sure that there is enough space between the optical I/O for two fibers. It may be tempting to put two grating couplers back to back or edge couplers next to each other with a tight loopback, but a single single-mode fiber will shadow the other optical port. Thus, for edge couplers on the same chip edge and grating couplers with the same direction of orientation, they should be spaced by at least  $127\ \mu\text{m}$ . Grating couplers with opposite orientation for I/O are trickier in that the distance between them will vary based on the test bench, as the vertical fiber holders tend to determine the how close two fibers can come. However, typically at least  $200\ \mu\text{m}$  is sufficient. If an edge coupler test structure must be on opposite sides of the chip, make sure to include an offset of about  $50\ \mu\text{m}$ . Edge couplers placed exactly opposite each other on the chip are extremely difficult to couple into (especially without the help of an IR camera to image the chip). This is because when one is aligning the fiber to edge couplers opposite each other across the chip, the power loss by pointing two fibers at each other above the chip (or through the Si substrate, or through oxide, etc.) can be less than 10 dB than coupling into the waveguide. A very crude back-of-the-envelope calculation reveals that when two fibers are pointed towards each other in free space at a distance of 1 mm, one could have as little as 6 dB of loss. Thus, optimizing fiber alignment when the noise floor is -10 dBm and the peak is roughly 1-3 dBm is very

difficult. On the other hand, when one is trying to align two single-mode fibers to the same side of the chip, one often starts at the noise floor of whatever power meter is being used (such as -65 dBm) because very little light ends up in the out-going fiber. It is thus significantly easier to start optimizing the positions of the fibers with respect to the waveguide because the signal is so much greater than the noise floor. Another basic optical test structures includes characterizing waveguide loss. As waveguide loss should be low (typically 0.5-1.5 dB/cm in a standard silicon photonics foundry [31, 125]), this will require a large amount of area to be able to see appreciable loss. A standard way to do this is to use a spiral that coils in on itself and has a total length of several centimeters. Such a structure only has one optical input must be measured using optical backscatter reflectometry (OBR). If there is not enough space for such a structure, one may have to follow up with the foundry for waveguide loss numbers.

## A.2 Device-Specific Test Structures

In this section, test structures for typical sub-components of the devices described in the thesis are discussed.

### A.2.1 Ring-based Crossbar switch

In designing a ring-based crossbar switch, it often quickens the design process if a sub-component consisting of a single crossing including the switch rings and associated signal pads is created and then copied to create the full switch. Designing in this way makes it easier to account for all of the required test structures. The main components therefore of a single switch crosspoint are the crossing and the switch rings. To test the crossings, the cutback method can be used; similar to the TLMs, include different numbers of serially cascaded crossings between between two optical I/O. When this collection of test




Figure A.2: An example of a crossing design, including measuring the crosstalk of the crossing.

structures is measured, plot the optical loss versus the number of crossings at a single wavelength. The slope will give an indication of loss per crossing while the y-intercept represents twice the optical loss of a single optical I/O. This number can be used to confirm a dedicated optical I/O loss test structure. Additionally, the crosstalk of the crossing can be measured if a waveguide is routed out from one of the perpendicularly oriented waveguides to the optical I/O. A sample crossing test structure is shown in figure A.2. In order to test the switch rings, it is useful to have a test structure in which all four ports of each variant of the add-drop switch rings are connected to some sort of optical I/O, with all of the electrical test structures connected to pads. Again, care should be taken such that the optical I/O, especially if grating coupled, are far enough away from one another. This is especially the case if there are electrical pads that will be connected via electrical probes, as these tend to take up a significant portion of the real estate on the test bench and can determine how close the tip can get to the fiber and other probes. With such a test structure, one can isolate the behavior of a single ring from the others. This is crucial because in, for example, a  $4 \times 4$  switch ring with three ring pairs per crosspoint, one may want to characterize a single ring but will not be able to isolate its performance from all the other rings in the same path. For example, one can use this to characterize such parameters as the  $Q$ , the resistance of the heaters, and the tuning efficiency with ease. To determine the pass-by loss of the rings, again, the cutback method can be used.

### A.2.2 Mach-Zehnder Modulator

Mach-Zehnder modulators (MZMs) by comparison may seem less complex when looking at the design, but comprises of many different components. Typical test structures for MZMs include characterizing the  $1 \times 2$  splitter—whether that comes in the form of a multi-mode interferometer (MMI) or a directional coupler—tapers, pn junctions, heater, and in the case of traveling-wave MZMs, the electrodes. A basic test structure for a  $1 \times 2$  splitter involves several cascaded splitters along say, the bottom output if oriented from left to right. The other outputs can be routed to a grating coupler. One can either apply the cutback method to find the excess loss (since each output port will have at least 3 dB loss compared to the power at the input port, regardless of direction) or simply measure the output power at the last  $1 \times 2$  splitter and divide by the number of  $1 \times 2$  splitters passed through to get a very rough number. If there is some sort of persistent bias towards one output towards the other, this may not be shown. A more rigorous method to testing the response of  $1 \times 2$  splitters is given in [126], using  $1 \times 2$  MMIs as the splitting method. The methodology can also be applied to directional couplers. Generally, taper test structures again use the cutback method. It is a good idea to get an idea of how much loss a single taper element will incur with simulations that employ 3D finite difference time domain (FDTD) type solvers. This will help determine a rough minimum number of elements and what increment to use in designing the cutback test structures. For example, if a taper element is expected to incur 1 dB of additional loss, a good minimum number to use would be two pairs to go from rib to ridge to rib waveguide and use increments of 2 pairs because 4 dB increments are very easily measured and would stand out from the noise due small fabrication errors or coupling differences. Heaters are often used to correct any errors in fabrication and to hold the bias of the MZM at a particular point. An unbalanced Mach-Zehnder interferometer (UMZI) or asymmetric Mach-Zehnder in-






Figure A.3: A heater test structure design is shown here. There are eight different heater designs, each routed to a pair of pads.

terferometer (AMZI) with a path length difference of 150–200  $\mu\text{m}$  between the two arms with multiple heater designs can be used to test efficiency. When designing a heater, it is best to aim for roughly 1  $\text{k}\Omega$  such that to achieve powers in the mW range, neither the required voltage nor current are too high to be driven by an ASIC. A sample design is shown in figure A.3. Typically the most efficient design is a heater on full-height silicon with undoped partial etch silicon several microns wide (or at least many times the  $1/e$  decay length in power of the mode of interest) connecting the heater to the waveguide. The undoped partial-etch silicon is a significantly better thermal conductor than the silicon dioxide. Using the various heaters in the larger UMZI, heater efficiency of each variant can be determined by measuring the shift in spectral fringe pattern versus the electrical power dissipated by the heater. If the MZM in question has traveling wave electrodes, it may be desirable to measure their electrical only response. If the S-parameters are to be measured, then it is necessary to include calibration structures. [127] has a thorough explanation of what test structures should be used. It should be said that when designing an MZM, make sure to include enough separation between the optical I/O and the RF pads. For a first time design, include at least 1-2 mm of space between the different types of I/O, especially if RF probes will be used. For mechanical stability, optical and RF

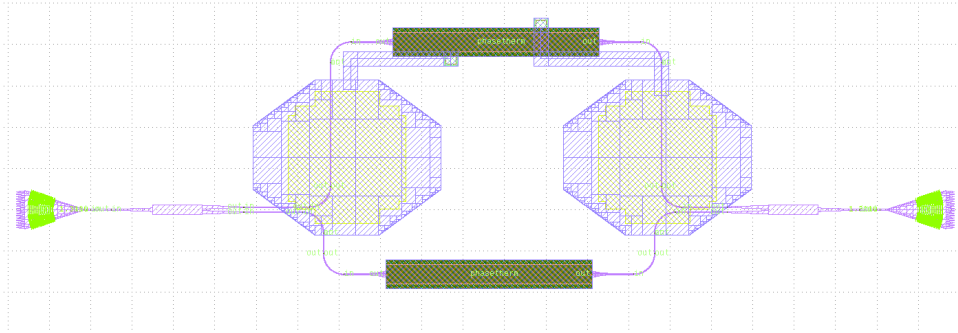


Figure A.4: A diode test structure design is shown here. There is a diode on each arm of the MZI.

probe holders tend to be very large. If optical and RF probe holders have to be close together, one may have to let a long segment (1 in.) of optical fiber hang outside of the fiber holder, which will introduce mechanical instability in the fiber and noise in the measurement. It is recommended that for ease of testing and/or packaging, the optical and RF I/O should be oriented perpendicular to each other.

### Diode test structures

Finally we come to the crucial electrooptical test structure for MZMs. MZMs rely on diodes implanted into the waveguide to incur a change in the index of refraction. A versatile test structure to find the effect of the diodes on the test structure is a diode implanted in a waveguide in a UMZI with path length difference of  $200\ \mu\text{m}$ . From the diodes, one can confirm the IV behavior from which the series resistance can be extracted as well as the capacitance of the diode in reverse bias—which is the region that the diode will likely be operated for high-speed operation. The series resistance can be extracted from the forward bias region of the IV. With both the resistance and the capacitance, the RC-limited bandwidth of the segment can be extracted. Since the diodes are embedded in waveguides, the wavelength shift of the notch of the UMZI and therefore  $V_\pi$  of the segment due to the diode can be measured. If designing MZMs for the first time in a

new process, it is worth trying different diode configurations varying the length, doping profile within the waveguide, and lateral versus interdigitated pn junctions. This will lead to insights into how to improve the design in the future.

### A.3 Some Practical Considerations

While edge couplers typically are lower loss than grating couplers [31,125], optical test structures should utilize grating couplers as much as possible. This is because grating coupled structures do not need to be placed on an edge in order to be tested. Furthermore, edge space may be at a premium since foundries may restrict the number of edge couplers per chip. Using grating couplers on test structures allows for placement flexibility when putting together the mask. Completely passive optical test structures can be placed under pads of other devices or test structures, for example, thereby saving space. Test structures can also be placed between adjacent devices such that if the latter were to be diced out, the test structures can be easily sacrificed. If there is dedicated space for test structures, align the grating couplers to the same orientation, and preferably with the same grating coupler and pad layout. If the chip is mounted on a moveable stage during testing, one can move the chip to go from test structure to test structure with very minimal adjustment to the fiber position; optimizing fiber position often takes up a very large chunk of time in testing time as this requires an extreme amount of precision.

It may seem counterintuitive, but it may also be worth designing a test structure designed to fail in order to compare the performance of other test structures. This is especially important for passive optical elements where the loss of a single component could be very low. Take for example, a taper from rib to ridge waveguide types. One may vary the length of the taper from 5  $\mu\text{m}$  to 50  $\mu\text{m}$  and try to flare out linearly or quadratically, etc., but having a test structure where one abruptly changes from rib to

ridge ultimately sets the baseline for improvements on a design. In designing a known-bad test structure may also yield surprisingly good results and allow you to focus more energy on other parts of device design improvement.

While it is tempting to include every single test structure possible to characterize each individual part of the device, there are often severe space constraints barring one from doing so. Therefore, one must make an educated guess and prioritize the parts of the device that are most crucial to the device performance. The desired performance of the device in question will dictate the priority list of test structures. For example, when designing a ring-based crossbar switch for the first time in a new technology, the desired performance could be as simple as a single yielded device. Because putting several variations of the same switch will likely consume a vast amount of design area, one should design a “best” switch based on simulated performance. The main test structures one might try to characterize the optical I/O, crossings, and rings present in the design. However, since the process variations cannot usually be characterized in the simulations, designing test structures that are variants of the main test structures can help lead to easily actionable improvements on the design.

Finally, any empty area can be filled with test structures if they cannot be filled by devices. Any empty area on a mask is a missed opportunity to understand a device or process more deeply. Even if it may feel like all of the test structures designed have covered the gamut of subcomponents to the device, take the opportunity to design a novel device. Do not leave any empty space on the mask!

# Appendix B

## Crossbar Switch Practical Details

In this section I detail practical details of the crossbar switch, including wirebonding diagrams and design considerations.

### B.1 Wirebonding

While it is undesirable to have to wirebond over 100 pads, it is one of the simpler ways to connect the switch electrically. Wirebonding the full  $4 \times 4$  switch which consists of 96 signal pads and a few ground pads, includes concerns such as pad pitch on both the switch and the PCB, maximum and minimum lengths of wirebonds, and not crossing wirebonds. The first two concerns can be addressed in the PCB and chip designs, but in this section, I will address the final issue. With a high density grid of pads, it is much more difficult to avoid wirebonds crossing. To ensure that the wires didn't short by touching, the wirebonds in the  $4 \times 4$  crossbar switch were made in layers. To illustrate this, figures B.1 through B.5 show all of the wirebonds on a single build, and then broken up into every layer.




Figure B.1: All of the wirebonds made for the crossbar switch mounted on a custom PCB. The large blue-green pad on which the chip is located is a ground plane, while the smaller blue-green squares represent pads on the PCB.




Figure B.2: The first layers of the wirebonds made for the crossbar switch, as denoted by the different colors. In this case, the black lines represent connections from ground pads on the switch to the ground plane on the PCB, while the red lines represent a layer of connections for signal pads.




Figure B.3: The next layers of wirebonds made for the crossbar switch, as denoted by the yellow and orange lines. These layers are above those depicted in figure B.2.




Figure B.4: The next layers of wirebonds made for the crossbar switch, as denoted by the green and pink lines. These layers are above those depicted in figure B.3.




Figure B.5: The final layers of wirebonds made for the crossbar switch, as denoted by the light blue and black lines. These layers are above those depicted in figure B.4. In this case, another ground connection could be made between the chip and PCB.

## B.2 Some Testing Considerations

Before beginning to test a crossbar switch, it is helpful to have a naming convention for each ring in the switch. First, have a global naming scheme for each crosspoint; the most helpful is two numbers where the first number corresponds to which numbered input, while the second number to corresponds to the output. Then, within each crosspoint, have a naming convention to denote something about each ring. For example, perhaps the rings are all different sizes, so one could use ‘big’, ‘medium’, and ‘small’.

## B.3 Some Design Considerations

It is perhaps counterintuitive to put a section on design considerations last, but it is really only after understanding some of the potential problems that arise from testing that one begins to appreciate some design concerns. It is important to include a loop-back



structure on the same side of the chip as the inputs and outputs. (Refer to Appendix A for more information on loop-back structures). This will not only help with calibrating the optical I/O out of the chip with minimal disturbance to the test setup, but when coupling to the chip for the first time, it serve as a known-good optical measurement with low risk of measurement failure by which one can set some expectation for "good" coupling and "bad" coupling. Another design consideration is the pitch of the optical and electrical I/O. For the electrical I/O, it is important to be cognizant of the minimum pitch and pad opening of whatever wirebonder or wirebonding service can provide for the type of pad metal. A tiled pattern of pads is the most efficient use of space. Depending on the space constraints, it may be helpful to stagger the pads between rows and/or columns. To facilitate the tiling of pads, one should consider making a repeatable cell consisting of each crosspoint, complete with waveguides, rings, a ground plane, and connections from active elements in the rings to pads and the ground plane. The pad pitch will determine the size of the cell. If one places the rings between the pads (and therefore at the same pitch as the pads), one can rest easy that the thermal crosstalk is minimized. The layout of a single cell is shown in figures B.6 and B.7.

When designing like this, it is helpful, especially when designing in Cadence, to take advantage of the hierarchical design flow. That is, having a separate cell for the ring and bus waveguide only, another cell that uses the first cell to connect to a higher metal layer, a cell that includes all the optical and electrical connections within the crosspoint, and finally a cell of the full-sized switch and connections to the optical I/O. When designing the optical I/O, the standard pitch is  $127\mu\text{m}$  or  $250\mu\text{m}$  on many optical fiber arrays, though one's optical packaging requirements may differ. When considering the placement of the optical I/O one is most likely to be constrained by other designers' space requirements. If this isn't a constraint, it may be tempting to put the optical I/O on opposite sides of the chip. If this is a desired design feature and one is using edge

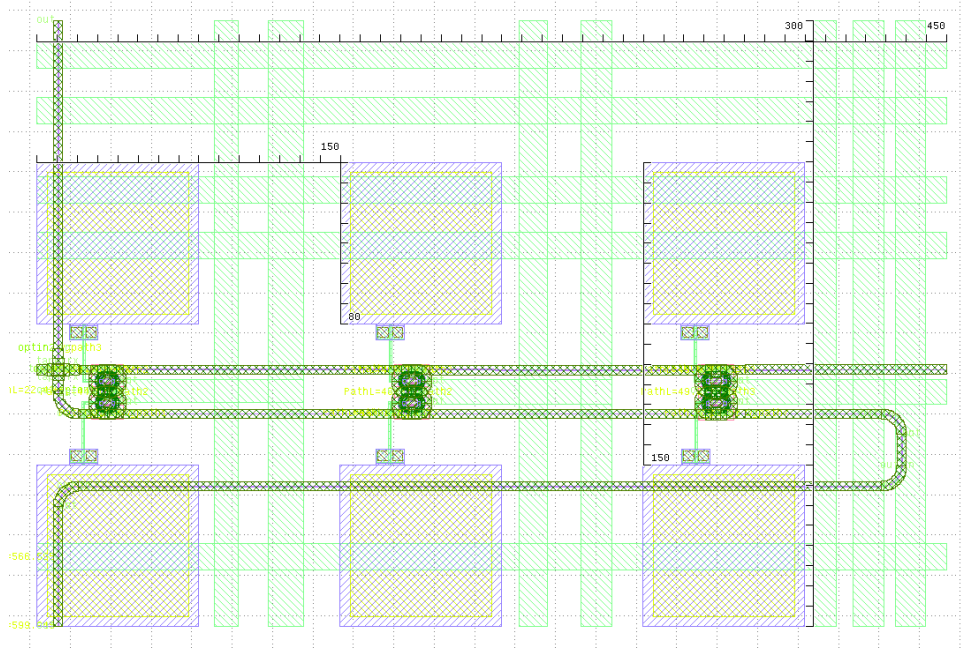


Figure B.6: A single cell design of the crossbar switch.

couplers, consider staggering the output edge couplers from the input edge couplers. When coupling to the chip for the first time, it can be difficult to discern whether the

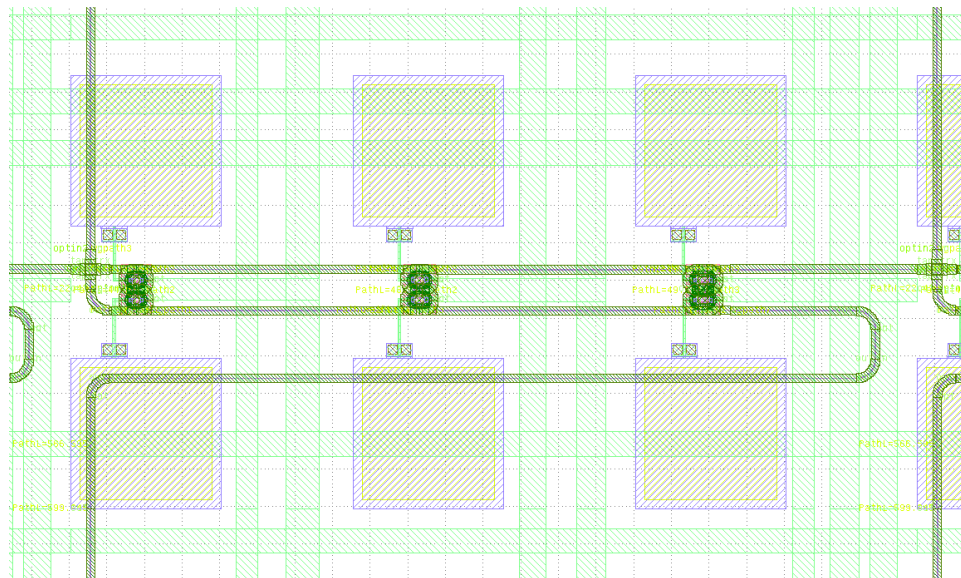


Figure B.7: A single cell design of the crossbar switch, but this time in a larger switch. Note that the waveguides are joined with neighboring cells.

---

optical mode is coupling into the substrate or into the waveguide if the input and outputs are directly opposite from one another, especially in the absence of an IR camera to image the chip. Anecdotally, one can still observe about 10dB of loss across a 1mm-wide chip when coupling into the substrate. While this is lower than the ideal performance of most edge couplers, this is not an unreasonable power loss number when optimizing the positions of the fibers. If one will be packaging the switch with wirebonds, putting optical I/O on opposite sides of the chip may mean that there is less area on the PCB that the chip is mounted on, and longer wirebonds. As can be seen in the Wirebonding section of this Appendix, being able to use all three sides to wirebond out.

# Appendix C

## RAMZI Switch Tuning Operation

In this section I detail RAMZI Switch Tuning operation. Since the RAMZI contains many cascaded elements, which often need some sort of tuning due to fabrication errors, mapping the correct voltages for each ring for each switch cell is required. This process is not as straightforward as it is for the crossbar switch, so I have added some simulated results to illustrate the steps in the process. In addition, I provide some practical details such as wirebonding diagrams and design considerations.

### C.1 Wirebonding

Because there are so few wirebonds for a  $4\times 4$  RAMZI switch compared to a crossbar switch of the equivalent size, layering the wirebonds such that they do not cross is less critical. Nevertheless, the figures C.1 thru C.3 are cartoons of the wirebonding connections that were made. Note that the same PCB that was used for the  $4\times 4$  crossbar switch was reused for the  $4\times 4$  RAMZI switch.




Figure C.1: All of the wirebonds made for the RAMZI switch mounted on a custom PCB. The large blue-green pad on which the chip is located is a ground plane, while the smaller blue-green squares represent pads on the PCB.




Figure C.2: The first layers of the wirebonds made for the RAMZI switch, as denoted by the red, orange, and black. In this case, the black lines represent connections from ground pads on the switch to the ground plane on the PCB, while the red lines represent a layer of connections for signal pads.




Figure C.3: The final layers of the wirebonds made for the RAMZI switch, as denoted by the grey and pink lines. These layers are above those depicted in figure C.2.

## C.2 Switching Operation

Figure C.4 shows the testing procedure one should follow for testing the RAMZI switch in order to map the operating voltages. This procedure should be followed to obtain the voltages for both the bar and cross states of each switch cell. Note that for the cells in the first stage, only the voltage for the cross state must be maintained. For simplicity, all ring pairs should be tuned during this procedure. To confirm that the tuning of the cells are correct, it is important to either monitor all outputs at once or switch between outputs.

### C.2.1 Switch Calibration Process

To help develop the procedure, a MATLAB simulation was developed to model the output from input 1 to each of the outputs. The simulated switch response was first with a random phase put on each ring in the path to simulate how a RAMZI switch fabricated

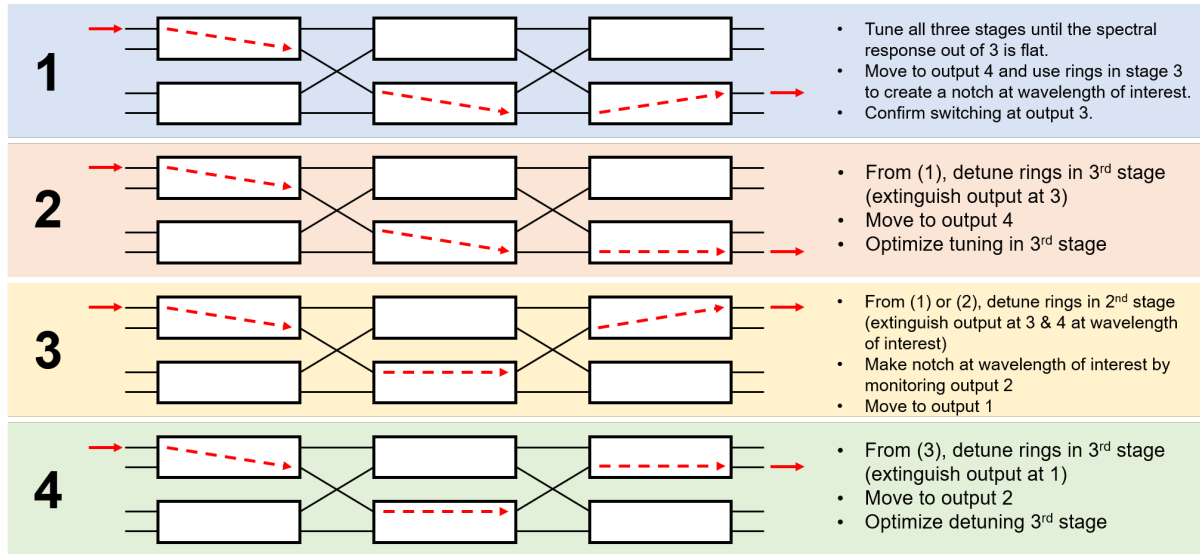


Figure C.4: Procedure for testing the RAMZI.

on a chip might behave on a first-pass optical sweep, though for the purposes of this demonstration, the phase difference between ring pairs is exaggerated. When calibrating the switch, it is recommended that all ring pairs are calibrated at each step.

The first stage is tuned such that both ring resonances are the same on both arms. We notice in Figure C.6 that there are fewer valleys in each of the outputs. In the lab, we would likely focus on the third output given that if we tune all of the rings in the path that this would result in a flat spectrum assuming that the edge coupler response is broadband and relatively flat in the region we are looking. Furthermore, when in the lab, simply tune one ring at a time. When one ring gets closer to its pair in the spectrum, the valley caused by the difference in resonance between the pair will narrow and become shallower the closer together and vice versa.

Next we tune the rings in the second stage, again looking to flatten the response in the output of port 3, as shown in figure C.7. We notice that the outputs of first and second ports are suppressed and do not appear. In the lab, the responses for these two ports would still appear in the spectrum, but would be minimized; the simulation assumes that




Figure C.5: The RAMZI switch with a random phase put on each ring.

the directional couplers are ideal (i.e. lossless). Record the voltage for the first and third stage.

Next, we tune the final stage to flatten output 3, which is shown in Figure C.8. It is worth going back to the previous stages to adjust the voltage on the rings to get a flatter response, especially for the rings in the first stage. It does not matter where the resonance falls for the rings in the first stage given that the wavelength-selectivity does not occur until the second stage for the routing algorithm chosen. Again, note that even though the response from outputs 1, 2, and 4 do not appear, this is due to the assumption that the directional couplers are lossless. In the lab, these outputs will be visible but will be suppressed compared to the output from port 3.

The rings in the third stage can now be tuned to output light to the fourth port. This can be seen in Figure C.9. At this point, record the voltage required to tune to the






Figure C.6: The RAMZI switch with the first stage tuned such that the resonances for the two ring pairs are matched.




Figure C.7: The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are matched.




Figure C.8: The RAMZI switch with the third stage tuned such that the resonances for the two ring pairs are matched and outputs to port 3.

fourth voltage.

Now that the calibration has been determined for the 3rd and 4th ports in the third state, the calibration can be determined for ports 1 and 2. First the rings in the second stage need to be detuned such that the wavelengths of interest are routed to the top  $2 \times 2$  cell in the third stage. The response can be seen in Figure C.10.

As can be seen from the response, either the rings in the second stage are somewhat detuned from the rings in the 3rd stage connected to ports 3 and 4, or the rings in the 3rd stage connected to ports 1 and 2 do not have the same resonances. Since applying heat to the rings red-shifts the resonances, when tuning the rings in the second and third stages, it is important to tune them to the longest resonance wavelength for each ring pair of all the cells in these cells. Thus, we tune the rings in the second stage such that they match the resonances shown in Figure C.9, namely red-shifting the first ring pair in the second stage to match the first ring pair in the third stage. In addition, we re-tune

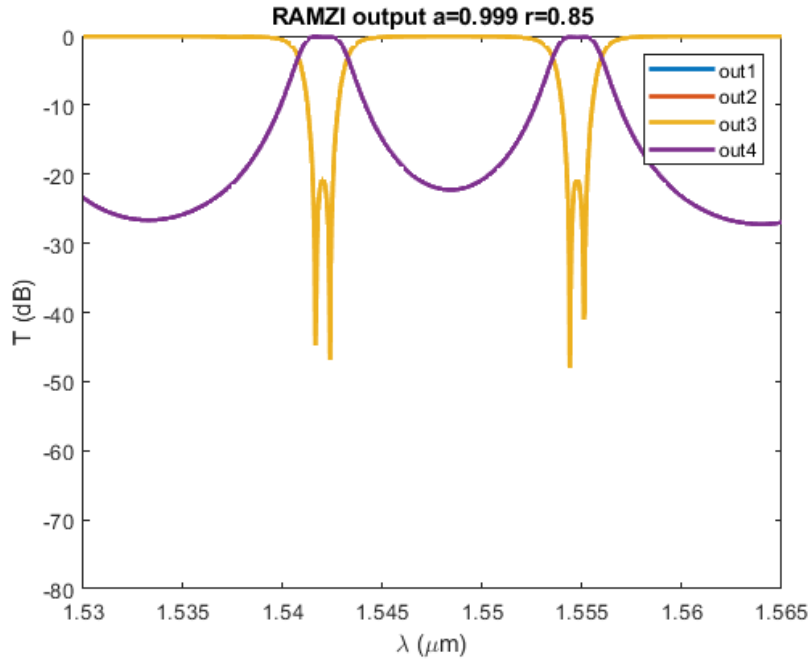


Figure C.9: The RAMZI switch with the third stage tuned such that the resonances for the two ring pairs are slightly detuned from one another and outputs to port 4.




Figure C.10: The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are slightly detuned from one another.

the “voltage” to output 4 to redshift the second ring pair centered at roughly  $1.554 \mu\text{m}$  to  $1.556 \mu\text{m}$ . The result is shown in figure C.11. Record the voltage in the third stage to switch to output 4, as well as the voltage in the second stage.




Figure C.11: The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are slightly detuned from one another.

Now, we make a second attempt to switch to ports 1 and 2 by detuning the rings slightly from one another in the second stage, the result of which can be seen in figure C.12.

Next, we can focus on finding the voltages to switch to ports 1 and 2. From the spectrum shown in Figure C.12, we tune the ring resonances such that they match on each arm in the third stage to output to port 1. This spectrum is shown in Figure C.13. Record the voltage in the third stage.

Finally, we detune the rings in the third stage to output to port 2, as shown in Figure C.14. Record the voltage in the third stage. At this point, it may be evident that the resonances recorded for ports 1 and 2 may be different than the resonances for ports 3

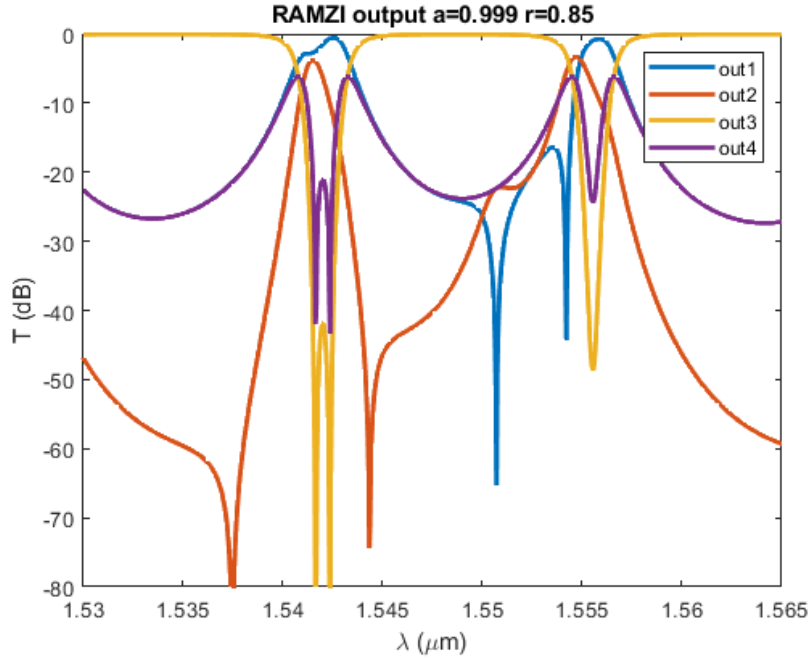


Figure C.12: The RAMZI switch with the second stage tuned such that the resonances for the two ring pairs are matched. The resonance of the first ring pair in the second stage is tuned to match the resonance of the first ring pair in second cell in the third stage. The second ring pair in the third stage is redshifted to match the resonance in the second stage.

and 4. Thus, it is important to go back to the previous appropriate step to continue to red-shift rings in either the second or third stages such that the resonances all match. Code for the simulation in this section can be found in appendix D.

### C.3 Some Design Considerations

Like the crossbar switch, designing a switch cell building block that can be reused facilitates not only building the full switch design, but also when using the design rule checker (DRC). Since the arms of the MZI for each switch cell must be balanced, making the arms of the MZI as short as possible while maintaining a large distance between the rings sharing the same bus such that thermal crosstalk isn't observed is important. For




Figure C.13: The RAMZI switch with the third stage now tuned such that the resonances on each arm are matched and the outputs are directed to port 1.

this reason, the switch cell did not have any bends between the directional coupler. The space between the arms can be filled with a ground plane to which the active elements of the rings can be connected. Make sure to connect the ground planes of each switch cell together with large shapes of metal. Tenuous connections from rings to ground planes, especially will introduce extra resistance in the paths that cannot be overcome easily. Due to the long and thin aspect ratio of the switch cell, the same approach of putting the pads on top of the switch cell that was employed in the crossbar switch is not necessarily appropriate here, but may be dependent on the available space. The layout of a single RAMZI cell is shown in C.15




Figure C.14: The RAMZI switch with the third stage now tuned such that the resonances on each arm are slightly detuned from one another and the outputs are directed to port 2.




Figure C.15: A single cell design of the RAMZI switch.

# Appendix D

## Code listing

### D.1 ACD Link Simulation Code

```
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%  
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%  
% Last update 03/12/2020  
% Last modification by: Takako  
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%  
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%  
  
clear all; clc  
  
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%  
% simulation choices  
Si = 0; % 1 = silicon; 0 = InP
```



```

cool = 1; % 1 = the cool way; 0 = Jim's
        way
diff_opt = 1; % 0-single-ended, 1-differential
seg_opt = 1; % 0-traveling wave, 1-segmented
dual_pol = 1; % 1 - dual pol, 0 - single pol
savedata = 0; % 1 - save data, 0 - don't save
        data
gen = 2; % 1 - 1st gen, 2 - 2nd gen

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Sweep parameters
if(cool)
    close all
    if (Si)
        lstart = 3; %set to 3
    else
        lstart = 1;
    end
    lend = lstart;
    NVsig = 30; % Number Vsig points
    Nlengths = 1; % Number of length points
else
    lstart = 0.5;
    lend = 6;
    NVsig = 1000;

```

```

    Nlengths = 100; %default = 20, go up to 100 for publication
end
l = linspace(lstart ,lend ,Nlengths);           % Modulator
    length [mm]
l2 = [];

Plaser = (-30:0.1:40);           %sweep TX laser power [dBm]
    normally -10 to 40
Npoints = length(Plaser);
Plo = (-30:0.1:40);           %sweep LO laser power [dBm]
    normally -10 to 40
%Plo = linspace(-10, 50, Npoints);

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Device parameters

q=1.6e-19;
if (Si)
    material = 'Si';
else
    material = 'InP';
end

[VpLp, Cperl, Zo, alpha_opt_log, alpha_elec, eta_laser, R,
    TXLoss0, RXLoss, LOLoss] = load_variables(material,gen);

```

```

alpha_opt = log(10^(-alpha_opt_log/10)); %Attenuation nepers per
      mm of modulator; Jim's val = 0.032 Np/mm

Rseg = 15; %Segmented driver collector
      resistance
eta_drv = 0.2; %Efficiency of Driver

%Microwave parameters
% alpha_elec = 0.1; %Attenuation nepers per mm of
      modulator original=0.1Np/mm
Skin_corner = 10e9; %approximate skin-effect factor
Skin_coef = 0.1; %strength relative to normal
      losses

%Circuits
% P_TIA_lin = 2^(dual_pol)*2*0.15; % first gen
      150mW, 2* for single pol, 4* for dual pol
% P_OPLL_lin = 0.104*2; % for each polarization,
      104 for OPLL, factor 2 for margin of other circuitry
P_TIA_lin = 2^(dual_pol)*0.198; % 2nd gen 198
      mW, 2* for dual pol
P_OPLL_lin = 0.053*2; % for each polarization,
      53mW for OPLL, factor 2 for margin of other circuitry

```

```

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% System requirements
Baud = 56e9; %Symbol rate rate
Rb = 2^(dual_pol)*2*Baud; %Bit rate (2 for QPSK, 2
    for dual pol)
ROF = 0.7; %Roll off factor
BW = Baud*ROF; %Bandwidth [Hz]
Qdes = 4.26; %desired Q 7 for
    10^-12, 3 for 10^-3, 4.09 for 2.1e-5 (KR4), 4.26 for 1e-5
Ipp = 50e-6; %Desired input sensitivity
    . Default 50uA
LinkMargin = 13; %Desired Link Margin.
    Normally 13, but can be set to 4dB (typical in data centers)

Vpi = VpLp./l; %Vpi of the modulator [V
    ]
Ctot = Cperl*l; %Total capacitance for a
    given modulator length;
VarTherm = (Ipp/2/Qdes)^2; %Thermal noise
Plaserlin = 10.^((Plaser-30)/10); %laser optical power at
    output of laser [W]
Ploin = 10.^((Plo-30)/10); %LO optical power at
    output of LO [W]

```

```

% Elec_atten = exp(-alpha_elec/2*(1+Skin_coef*sqrt(BW/
    Skin_corner)));

for ll = 1:length(l)
    TXLoss = TXLoss0 + alpha_opt_log*l(ll) + 4*(dual_pol-1); %
        Transmitter loss [dB] (GratCoup loss (2*3dB) + MMI loss
        (3dB) + phase shifter loss , -5 if single pol
    TotalLoss = (TXLoss+RXLoss+LinkMargin);
    TotalLosslin = 10^(-TotalLoss/10);

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
% Calculations
% Elec_atten = exp(-alpha_elec*l(ll)/2*(1+Skin_coef*sqrt(BW/
Skin_corner)));
    Elec_atten = exp(-alpha_elec*l(ll));

    PlaserPD = Plaser - TotalLoss;           %Laser power at the
        PD
    PlaserPDlin = 10.^((PlaserPD-30)/10);    %Laser power at the
        PD [W]

    PloPD = Plo - LOLoss;                   %LO power at the PD
    PloPDlin = 10.^((PloPD-30)/10);         %LO power at the PD
        [W]

```

```

Vsig = linspace(Vpi(11)/100, Vpi(11), NVsig); %Peak
    Electrical signal amplitude in the driver
minEPB_temp_vec = [];
TXlaser_minEPB_vec = [];
minV_temp_vec = [];

if(seg_opt)
    SegLength = 200e-3; %200um, in mm
    NumSeg = ceil(1(11)/SegLength);
    active_length = NumSeg*SegLength; %therefore what's the
        total active length
end

for vv = 1:NVsig
    %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
    %%%% DRIVER POWER CONSUMPTION %%%%
    %%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
    if (seg_opt == 1)
        %Segmented calculation for driver, assuming
            monolithic
        %integration into 90nm CMOS. eta_drv is dependent on
            frequency
        %in driving a capacitive load.
%           g = 1e-15; %1 fF/um
%           f = 1e-3; %1 mA/um

```

```

%           pptn = f/g;
%           pptn = 135/1e-9; % 135 V/ns, 90nm CMOS | f/g,
where g = C_out_drvr/w_transistor, f = i_out_drvr/
w_transistor, vdd= 1V
           pptn = 334/1e-9; % 334 V/ns, 45nm CMOS | f/g, where
           g = C_out_drvr/w_transistor, f = i_out_drvr/
           w_transistor, vdd = 1.1V
%           pptn = 370/1e-9; % 370 V/ns, 22nm CMOS | f/g,
where g = C_out_drvr/w_transistor, f = i_out_drvr/
w_transistor, vdd = 0.8V
           tr = 5e-12; %5 ps rise time of signal.
%           eta_drv = 1 - (g/f)*(Vsig(vv)/tr);
eta_drv = 1 - (1/pptn)*(Vsig(vv)/tr); % NOTE: THIS
           CAN BE NEGATIVE WHICH IS UNPHYSICAL
if eta_drv <= 0 || eta_drv > 1
           continue
end
CperSeg = Cperl*SegLength;
Pdc = (1/2)*(1/2)*CperSeg.*(Vsig(vv).^2)*Baud*(1/
           eta_drv)*NumSeg*4; % (1/2)*C*V^2*(datarate)*(
           probability that state is one that consumes
           energy=1/2)/(eta_drv)*(number of segments/arm)*(
           number of arms)
elseif (seg_opt == 0)
           %Traveling wave driver

```

```

Pdc = ((Vsig(vv)/(2^diff_opt)).^2).*2^(1+diff_opt)/(
    eta_drv*Zo*8);    %DC power consumption assuming
                    differential amplif. both for differential and
                    single ended mod, factor of 2^(1) is for IQ
end

Pdc = Pdc*ones(1,length(Plo));

%%% calculate ModFactor
if(seg_opt)
    VLequiv = 2^(diff_opt)*Vsig(vv)*active_length;
    ModFactor(ll,vv) = 0.5*(1+cos(pi*(1+VLequiv./(2*VpLp
        ))));
else
    VLequiv = (Vsig(vv)./2)./alpha_elec.*(1-exp(-
        alpha_elec*l(ll)))*2^diff_opt; %Vsig(vv)*sqrt(2)
    % to convert to Vpeak, divide by 2 to get voltage
    % at beginning of mzm
    ModFactor(ll,vv) = 0.5*(1+cos(pi*(1+VLequiv./(2*VpLp
        ))));
end

%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%
%%%%%%%% SHOT NOISE %%%%%%%%%
%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%%

```



```

VarShot=2*q*R*BW*(PloPDlin'+Plaserlin*ModFactor(11,vv)*
    TotalLosslin);

SNR=(4*R^2*(PloPDlin'.*Plaserlin*ModFactor(11,vv)*
    TotalLosslin))./(VarTherm+VarShot); %SNR for QPSK
Q = sqrt(SNR);
V =[0,7];

PlaserQ_lin = [];
Plaserout = [];
Pdc_laser = [];
Pdc_lo_lin = [];
EPB = [];
all_Ps = [];

for k = 1:length(Plo)

    Qbound(1,k) = min(Q(k,:));
    Qbound(2,k) = max(Q(k,:));

    if (Qbound(2,k)<Qdes)
        continue
    end

    ind = find(Q(k,:) >= Qdes,1,'first');

```

```

PlaserQ_val = Plaserlin(ind);
PlaserQ_lin = [PlaserQ_lin PlaserQ_val];

%           Plaserout_val = PlaserQ_val/exp(-alpha_opt*l(ll))
%           /((1/2*(1+cos((-1+ModFactor(ll ,vv)/2)*pi))); %optical loss
%%           Plaserout_val = exp(-alpha_opt*l(ll)); %optical
%           loss
%           Plaserout = [Plaserout Plaserout_val];
%           Pdc_laser_val = Plaserout_val/eta_laser;
%DC power of laser
Pdc_laser_val = PlaserQ_val/eta_laser;
Pdc_laser = [Pdc_laser Pdc_laser_val];
%           Pdc_lo(k) = Plolin(k)/eta_laser;           %DC
%           power of LO
Pdc_lo_val = Plolin(k)/eta_laser;
Pdc_lo_lin = [Pdc_lo_lin Pdc_lo_val];

all_P = Pdc(k) + Pdc_laser_val + Pdc_lo_val +
        P_TIA_lin + P_OPLL_lin;
EPB_val = (Pdc(k) + Pdc_laser_val + Pdc_lo_val +
        P_TIA_lin + P_OPLL_lin) ./Rb;           %Energy
%           per bit traveling wave
EPB = [EPB EPB_val];
power_entries = (10^2/all_P)*[Pdc(k) Pdc_laser_val
        Pdc_lo_val P_TIA_lin P_OPLL_lin];

```

```

    all_Ps = [all_Ps; power_entries];           %output
        percentage of power from which source

end

PlaserQ = 10*log10(PlaserQ_lin)+30;
Pdc_lo = 10*log10(Pdc_lo_lin*eta_laser)+30; % multiply
    by eta_laser to undo the eta_laser division above to
    plot what the actual value of P_dc,LO is
% [minEPB_temp, min_ind] = min(EPB);
if max(Qbound(2,:))>Qdes
    [TXLOlaser_diff, min_ind] = min(abs(PlaserQ-Pdc_lo))
        ; %find value where PlaserQ ~ Pdc_lo
    minEPB_temp = EPB(min_ind);
    minEPB_temp_vec = [minEPB_temp_vec minEPB_temp];
    minV_temp_vec = [minV_temp_vec Vsig(vv)];
    TXlaser_minEPB = PlaserQ(min_ind);
    TXlaser_minEPB_vec = [TXlaser_minEPB_vec
        TXlaser_minEPB];
end

if (cool)
    fig=figure;

% subplot(2,1,1)

```

```
%          plot(10*log10(PlaserQ)+30, exp(-alpha_opt*l(l))/1
           e-12,'b')
%
%          subplot(2,1,2)

%          PlaserQ = 10*log10(PlaserQ_lin)+30;
yyaxis right
colormap(jet)
%          contour(PlaserQ,Pdc_lo,Q,V);
plot(PlaserQ,Pdc_lo,'r','linewidth',3)
xlabel('TX Input Laser Power [dBm]')
ylabel('LO Power [dBm]')
axis([0 30 -10 20])

grid on
hold on
AX = gca;
set(AX,'fontsize',14);
set(AX,'fontweight','bold');
set(AX,'fontname','Arial');
set(AX,'ycolor','r');

yyaxis left
plot(PlaserQ,EPB/1e-12,'.k','MarkerSize',10)
```

```

xlabel('TX Input Laser Power [dBm]')
ylabel('Energy Per Bit [pJ/b]')
grid on
axis([0 30 0 25])
hold on
AX = gca;
set(AX, 'fontsize', 14);
set(AX, 'fontweight', 'bold');
set(AX, 'fontname', 'Arial');
set(AX, 'ycolor', 'k');

%adds break down of power source by percentage if
    you hover
%over the curve.
dcm_obj = datacursormode(fig);
set(dcm_obj, 'UpdateFcn', {@myupdatefcn, all_Ps})

if(Si)
    title(['Si MOD length = ' num2str(l(11)) ' [mm]
        V-pp-d = ' num2str(round(Vsig(vv)/(2^
            diff_opt),1)) ' [V]'])
else
    title(['InP MOD length = ' num2str(l(11)) ' [mm]
        V-pp-d = ' num2str(round(Vsig(vv)/(2^
            diff_opt),1)) ' [V]'])

```

```
        end

    end

end

if(cool)
    figure
    yyaxis right
    plot(minV_temp_vec/(2^diff_opt), minEPB_temp_vec*1e12)
    xlabel('V_{sig}')
    ylabel('EPB_{min}')
    hold on
    ylim([0 15])

    yyaxis left
    plot(minV_temp_vec/(2^diff_opt), TXlaser_minEPB_vec)
    ylabel('TX, LO laser powers')
    ylim([0 15])
end

if length(min(minEPB_temp_vec))>0
    [minEPB(11) ind] = min(minEPB_temp_vec);
%     minEPB = [minEPB tempminEPB];
end
```

```

    minVsig(l1) = minV_temp_vec(ind);
    TXlaser_minEPB_min(l1) = TXlaser_minEPB_vec(ind);
    if(seg_opt == 1)
        l2(l1) = active_length;
    else
        l2 = [l2 l(l1)];
    end
end

% figure;
% plot(Vsig,EPBmin(l1,:) *1e12)
% xlabel('Vsig [V]')
% title(['EPBmin length = ' num2str(l(l1)) ' [mm]'])

end

if(~cool)
    %figure
    yyaxis right
    plot(l2, minEPB*1e12, 'k', 'linewidth',3)
    xlabel('Mod Length [mm]')
    ylabel('Energy Per Bit [pJ/b]')
    grid on
    axis([lstart lend 0 15]) %50
    hold on
    AX = gca;

```

```
set(AX, 'fontsize', 14);
set(AX, 'fontweight', 'bold');
set(AX, 'fontname', 'Arial');
set(AX, 'ycolor', 'k');

yyaxis left
% plot(1, Vpi, '--')
% hold on
minVsig = minVsig/(2^diff_opt);
if Si ==1
    windowSize = 5;
    b = (1/windowSize)*ones(1, windowSize);
    a = 1;
    filtered = filter(b,a, minVsig);
    minVsig=filtered;
    plot(l2(5:end), minVsig(5:end), 'r', 'linewidth', 3)
else
    plot(l2, minVsig, 'r', 'linewidth', 3)
end
xlabel('Mod Length [mm]')
ylabel('V-pp-d [V]')
grid on
axis([lstart lend 0 3]) %10
hold on
AX = gca;
```



```

set (AX, 'fontsize', 14);
set (AX, 'fontweight', 'bold');
set (AX, 'fontname', 'Arial');
set (AX, 'ycolor', 'r');

%   ylabel = {'EPB', 'V_sig', 'TXlaser'};
%   plotyyy (l2, minVsig, l2, minEPB*1e12, l2, TXlaser_minEPB_min,
%           ylabel)
%   xlabel ('Mod Length [mm]')

end

if (savedata) && (cool == 0)
    minEPB = minEPB*1e12;
    if (seg_opt)
        if (Si)
            save ('C:\Users\Takako Hirokawa\Documents\ARPA-E\
                Power Budget\IMDD\Coherentvars_Si_SEG', 'l2', '
                minEPB', 'minVsig', 'TXlaser_minEPB_min');
        else
            save ('C:\Users\Takako Hirokawa\Documents\ARPA-E\
                Power Budget\IMDD\Coherentvars_InP_SEG', 'l2', '
                minEPB', 'minVsig', 'TXlaser_minEPB_min');
        end
    elseif ~(seg_opt)

```

```

    if(Si)
        save('C:\Users\takako Hirokawa\Documents\ARPA-E\
            Power Budget\IMDD\Coherentvars_Si_TW', '12', '
            minEPB', 'minVsig', 'TXlaser_minEPB_min');
    else
        save('C:\Users\takako Hirokawa\Documents\ARPA-E\
            Power Budget\IMDD\Coherentvars_InP_TW', '12', '
            minEPB', 'minVsig', 'TXlaser_minEPB_min');
    end
end
end
end

```

```

function txt = myupdatefcn(~, event_obj, all_Ps)
% Customizes text of data tips
pos = get(event_obj, 'Position');
I = get(event_obj, 'DataIndex');
txt = {[ 'X: ', num2str(pos(1)) ], ...
       [ 'Y: ', num2str(pos(2)) ], ...
       [ 'P_{TXdriver}: ', num2str(all_Ps(I,1)), '%'] , ...
       [ 'P_{TXlaser}: ', num2str(all_Ps(I,2)), '%'] , ...
       [ 'P_{LOlaser}: ', num2str(all_Ps(I,3)), '%'] , ...
       [ 'P_{TIA}: ', num2str(all_Ps(I,4)), '%'] , ...
       [ 'P_{OPLL}: ', num2str(all_Ps(I,5)), '%'] ];
end

```

## D.2 RAMZI 4×4 with 2 Ring Pairs Simulation Code

```

% Simulates a 4x4 RAMZI switch response from port 1 to all four
    ports , with
% 2 ring pairs per switch cell.

clear all

% close all

Lres1 = 35e-6; %25um diameter of ring
Lres2 = 36e-6; %2nd ring
a = 0.999; %single pass amplitude transmission
r = 0.85; %self-coupling coefficient , 0.65
lambda = linspace(1.530e-6,1.565e-6,1000); %1550 nm
neff = 2.2;
randomphase = 0;
if randomphase %add a phase difference from the resonance
    condition.
    % phase difference in the (1,1) and (2,2) switching elements
        for first
    % ring pair:
    phasediff_1 = pi/4*randi([-1000,1000],[4,1])/1000;
    % phase difference in the rings in the 3rd stage switching
        elements for
    % first ring pair:
    phasediff2_1 = pi/4*randi([-1000,1000],[4,1])/1000;
    % phase difference in the (1,1) and (2,2) switching elements

```

```
        for second
% ring pair:
phasediff_2 = pi/4*randi([-1000,1000],[4,1])/1000;
% phase difference in the rings in the 3rd stage switching
    elements for
% second ring pair:
phasediff2_2 = pi/4*randi([-1000,1000],[4,1])/1000;

%display the phase difference to record later for the second
    part of
%the if loop.
disp(phasediff_1)
disp(phasediff2_1)
disp(phasediff_2)
disp(phasediff2_2)
else
%Use this space to manipulate the phases of the rings with known
    phases.

phasediff_1 = [ 0.5050; -0.7618; -0.7179; -0.5199];
phasediff2_1 = [ 0.2340; 0.3644; 0.2325; -0.0770];
phasediff_2 = [ 0.0738; -0.3204; 0.3848; -0.4885];
phasediff2_2 = [ 0.2937; -0.4972; -0.2066; 0.1971];
%
%   phasediff_1 = [ 0.5050; 0.5050; 0.2325; 0.5825];
```

```
%    phasediff2_1 = [ 0.2325; 0.5825; 0.2325; 0.5825];
%    phasediff_2  = [ 0.0738; 0.0738; 0.3848; 0.6848];
%    phasediff2_2 = [ 0.3848; 0.6848; 0.3848; 0.6848];

end

phase = pi/10;
n1 = neff;

%calculate the phase difference for the first two stages for
    both ring
%pairs
theta_base_1 = 2*pi*n1*Lres1./lambda;
theta_base_2 = 2*pi*n1*Lres2./lambda;
theta0_1 = [theta_base_1;theta_base_1;theta_base_1;theta_base_1;
    theta_base_1;theta_base_1];
theta0_2= [theta_base_2; theta_base_2; theta_base_2;
    theta_base_2; theta_base_2; theta_base_2];
theta_1 = theta0_1;
theta_2 = theta0_2;
theta_1(1:4,:) = theta0_1(1:4,)+phasediff_1;
theta_2(1:4,:) = theta0_2(1:4,)+phasediff_2;

%calculate the responses at each of the ports
```

```

output1 = cross(a,r,theta_1(1,:),theta_1(2,:),theta_2(1,:),
    theta_2(2,:)).* ...
    bar(a,r,theta_1(3,:),theta_1(4,:),theta_2(3,:),
    theta_2(4,:)).* ...
    cross(a,r,(theta_1(5,)+phasediff2_1(1)),(theta_1
    (6,)+phasediff2_1(2)), ...
    (theta_2(5,)+phasediff2_2(1)),(theta_2
    (6,)+phasediff2_2(2)));
output2 = cross(a,r,theta_1(1,:),theta_1(2,:),theta_2(1,:),
    theta_2(2,:)).* ...
    bar(a,r,theta_1(3,:),theta_1(4,:),theta_2(3,:),
    theta_2(4,:)).* ...
    bar(a,r,(theta_1(5,)+phasediff2_1(1)),(theta_1(6,)+
    phasediff2_1(2)), ...
    (theta_2(5,)+phasediff2_2(1)),(theta_2(6,)+
    phasediff2_2(2)));
output3 = cross(a,r,theta_1(1,:),theta_1(2,:),theta_2(1,:),
    theta_2(2,:)).* ...
    cross(a,r,theta_1(3,:),theta_1(4,:),theta_2(3,:),
    theta_2(4,:)).* ...
    cross(a,r,(theta_1(5,)+phasediff2_1(3)),(theta_1(6,)+
    phasediff2_1(4)), ...
    (theta_2(5,)+phasediff2_2(3)),(theta_2(6,)+
    phasediff2_2(4)));
output4 = cross(a,r,theta_1(1,:),theta_1(2,:),theta_2(1,:),

```

```

theta_2(2,:)).* ...
    cross(a,r,theta_1(3,:),theta_1(4,:),theta_2(3,:),
        theta_2(4,:)).* ...
    bar(a,r,(theta_1(5,)+phasediff2_1(3)),(theta_1(6,)+
        phasediff2_1(4)), ...
        (theta_2(5,)+phasediff2_2(3)),(theta_2(6,)+
            phasediff2_2(4)));

%convert the responses into log10
logoutput1 = 10*log10(output1);
logoutput2 = 10*log10(output2);
logoutput3 = 10*log10(output3);
logoutput4 = 10*log10(output4);

% figure(1)
figure
plot(1e6*lambda,logoutput1,'LineWidth',2)
hold on
plot(1e6*lambda,logoutput2,'LineWidth',2)
plot(1e6*lambda,logoutput3,'LineWidth',2)
plot(1e6*lambda,logoutput4,'LineWidth',2)
title(strcat('RAMZI output a=',num2str(a),' r=',num2str(r)))
% legend('bar', '\Delta \theta = \pi/10','cross', '\Delta \theta = \
    pi/10')
xlim(1e6*[lambda(1) lambda(end)])

```

```

ylim([-80 0])
legend('out1','out2','out3','out4')
xlabel('\lambda (\mu m)')
ylabel('T (dB)')

%% function to calculate the all-pass ring response
function [ringEthru,ringthru, phi]= ring_response(a,r,theta)
phi = theta + atan((r.*sin(theta))./(a-r.*cos(theta)))+ ...
            atan((r*a.*sin(theta))./(1-r*a.*cos(theta)));
ringthru = (a^2-2*r*a*cos(theta)+r^2)./(1-2*a*r*cos(theta)+(r*a)
            ^2);
ringEthru = sqrt(ringthru);
end

% calculates the bar state of a 2x2 RAMZI cell with two ring
  pairs
function barstate = bar(a, r, theta1_1, theta2_1, theta1_2,
    theta2_2)
[ringEthru1_1,ringthru1_1, phi1_1]= ring_response(a,r,theta1_1);
[ringEthru2_1,ringthru2_1, phi2_1]= ring_response(a,r,theta2_1);
[ringEthru1_2,ringthru1_2, phi1_2]= ring_response(a,r,theta1_2);
[ringEthru2_2,ringthru2_2, phi2_2]= ring_response(a,r,theta2_2);
ringthru1 = ringthru1_1.*ringthru1_2;
ringthru2 = ringthru2_1.*ringthru2_2;
ringEthru1 = ringEthru1_1.*ringEthru1_2;

```



```

ringEthru2 = ringEthru2_1.*ringEthru2_2;
barstate = (1/4)*(ringthru1+ringthru2-2*ringEthru1.*ringEthru2.*
    cos(phi1_1+phi1_2-phi2_1-phi2_2));
end

% calculates the cross state of a 2x2 RAMZI cell with two ring
    pairs
function crosstate = cross(a, r, theta1_1, theta2_1, theta1_2,
    theta2_2)
[ringEthru1_1, ringthru1_1, phi1_1]= ring_response(a, r, theta1_1);
[ringEthru2_1, ringthru2_1, phi2_1]= ring_response(a, r, theta2_1);
[ringEthru1_2, ringthru1_2, phi1_2]= ring_response(a, r, theta1_2);
[ringEthru2_2, ringthru2_2, phi2_2]= ring_response(a, r, theta2_2);
ringthru1 = ringthru1_1.*ringthru1_2;
ringthru2 = ringthru2_1.*ringthru2_2;
ringEthru1 = ringEthru1_1.*ringEthru1_2;
ringEthru2 = ringEthru2_1.*ringEthru2_2;
crosstate = (1/4)*(ringthru1+ringthru2+2*ringEthru1.*ringEthru2
    .* cos(phi1_1+phi1_2-phi2_1-phi2_2));
end

```

## D.3 RAMZI $2 \times 2$ for an Add-Drop or All-Pass Ring Pair

```

%2x2 RAMZI simulation , can switch between add-drop ring and all-
    pass ring
%response

clear all
% close all
Lres = 35e-6; %25um diameter of ring
a = 0.999; %single pass amplitude transmission
r = 0.85; %self-coupling coefficient , 0.65
r2 = 0.999; %self-coupling coefficient to second buswaveguide
lambda = linspace(1.530e-6,1.565e-6,1000); %1550 nm
neff = 2.2;
% deltan = 0.05;
phase = pi/10; % phase difference between rings on each arm
adddrop = 1; %add-drop or all-pass ring response switch

n1 = neff;

theta1 = 2*pi*n1*Lres./lambda;
theta2 = theta1+phase;

phil = theta1 + atan((r.*sin(theta1))./(a-r.*cos(theta1)))+ ...

```

```

        atan((r*a.*sin(theta1))./(1-r*a.*cos(theta1)));
phi2 = theta2 + atan((r.*sin(theta2))./(a-r.*cos(theta2)))+ ...
        atan((r*a.*sin(theta2))./(1-r*a.*cos(theta2)));

%calculate ring resopnse
if adddrop
    ringthru1 = (r2^2*a^2-2*r*r2*a*cos(phi1)+r^2)./...
        (1-2*r*r2*a*cos(phi1)+(r*r2*a).^2);
    ringthru2 = (r2^2*a^2-2*r*r2*a*cos(phi2)+r^2)./...
        (1-2*r*r2*a*cos(phi2)+(r*r2*a).^2);
else
    ringthru1 = (a^2-2*r*a*cos(phi1)+r^2)./(1-2*a*r*cos(phi1)+(r
        *a)^2);
    ringthru2 = (a^2-2*r*a*cos(phi2)+r^2)./(1-2*a*r*cos(phi2)+(r
        *a)^2);
end

ringEthru1 = sqrt(ringthru1);
ringEthru2 = sqrt(ringthru2);

%% cross terms
% cross1 = exp(1j.*(theta1-theta2)).*...
%
%         (a^2+2*a*r.*cos((theta1+theta2)./2).*exp(1j.*(
%         theta2-theta1)/2)+r^2.*exp(1j.*(theta2-theta1)./2))./ ...
%
%         (1-2*a*r.*cos((theta1+theta2)./2).*exp(-1j.*(
%         theta1-theta2)./2)+a^2*r^2.*exp(1j.*(theta1-theta2)));

```

```
% cross2 = conj(cross1);

%calculate bar and cross states
barstate = (1/4)*(ringthru1+ringthru2-2*ringEthru1.*ringEthru2.*
    cos(phi1-phi2));
crosstate = (1/4)*(ringthru1+ringthru2+2*ringEthru1.*ringEthru2
    .*cos(phi1-phi2));

%convert to dB
logbarstate = 10*log10(barstate);
logcrosstate = 10*log10(crosstate);

%plot response of RAMZI
figure(1)
plot(1e6*lambda, logbarstate, 'LineWidth', 2)
hold on
plot(1e6*lambda, logcrosstate, 'LineWidth', 2)
title(strcat('RAMZI output a=', num2str(a), ' r=', num2str(r)))
xlim(1e6*[lambda(1) lambda(end)])
ylim([-40 0])
xlabel('\lambda (\mum)')
ylabel('T (dB)')

%plot ring response
```

```

figure(2)
hold on
plot(1e6*lambda, 10*log10(ringthru2), 'LineWidth',2)
title(streac('Ring output a=',num2str(a), ' r=',num2str(r)))
xlim(1e6*[lambda(1) lambda(end)])
ylim([-40 0])
xlabel('\lambda (\mum)')
ylabel('T (dB)')
if adddrop
    hold on
    ringdrop = ((1-r^2).*(1-r2^2).*a) ./ ...
                (1-2*r*r2*a*cos(phi2)+(r*r2*a).^2);
    plot(1e6*lambda, 10*log10(ringdrop), 'LineWidth',2)
end

```

# Bibliography

- [1] A. Andreyev, “Introducing data center fabric, the next-generation facebook data center network,” <https://engineering.fb.com/production-engineering/introducing-data-center-fabric-the-next-generation-facebook-data-center-network/> (2014). [online; accessed June 16, 2020].
- [2] X. Zhou, R. Urata, and H. Liu, “Beyond 1 tb/s intra-data center interconnect technology: Im-dd or coherent?” *Journal of Lightwave Technology* **38**, 475–484 (2020).
- [3] A. S. G. Andrae and T. Edler, “On global electricity usage of communication technology: Trends to 2030,” *Challenges* **6**, 117–157 (2015).
- [4] M. Dayarathna, Y. Wen, and R. Fan, “Data center energy consumption modeling: A survey,” *IEEE Communications Surveys Tutorials* **18**, 732–794 (2016).
- [5] K. Darrow and B. Hedman, “Opportunities for combined heat and power in data centers,” Tech. rep., Oak Ridge National Laboratory (2009).
- [6] “Cisco annual internet report (2018-2023,” Tech. rep., Cisco (2020).
- [7] N. Jones, “How to stop data centres from gobbling up the world’s electricity,” *Nature* **561**, 163–166 (2018).
- [8] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, “United states data center energy usage report,” Tech. rep., Lawrence Berkeley National Laboratory (2016).
- [9] “Google data center renewable energy,” <https://www.google.com/about/datacenters/renewable/>. [online; accessed June 16, 2020].
- [10] “On our way to lower emissions and 100% renewable energy,” <https://about.fb.com/news/2018/08/renewable-energy/>. [online; accessed June 16, 2020].
- [11] M. Al-Fares, A. Loukissas, and A. Vahdat, “A scalable, commodity data center network architecture,” in “Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication,” (Association for Computing Machinery, New York, NY, USA, 2008), SIGCOMM ‘08, pp. 63–74.

- [12] Cisco Systems, Inc., *Cisco Data Center Infrastructure 2.5 Design Guide* (2011).
- [13] [https://www.mellanox.com/sites/default/files/related-docs/prod\\_lib\\_switch\\_systems/pb\\_sb7800.pdf](https://www.mellanox.com/sites/default/files/related-docs/prod_lib_switch_systems/pb_sb7800.pdf). [online; accessed June 17, 2020].
- [14] B. G. Lee, “Photonic switch fabrics in computer communications systems,” in “2018 Optical Fiber Communications Conference and Exposition (OFC),” (2018), pp. 1–22.
- [15] K. E. Grutter and T. Horton, “Findings of the 2nd photonics and electronics technology for extreme scale computing workgroup (repete): Design challenges for socket level photonic io,” Tech. rep., Laboratory for Physical Sciences.
- [16] <https://www.dellemc.com/resources/en-us/asset/data-sheets/products/networking/dell-emc-networking-z9332f-spec-sheet.pdf>. [online; accessed June 17, 2020].
- [17] [https://buy.hpe.com/emea/europe/en/networking/switches/fixed-port-l3-managed-ethernet-switches/flexfabric-5700-switch-products/hpe-flexfabric-5700-switch-series/p/7268889#:~:text=The%20HPE%20FlexFabric%205700%20Switch%20Series%20expands%20your%20enterprise%20network, CB\)%20switch%20improving%20server%20connectivity.](https://buy.hpe.com/emea/europe/en/networking/switches/fixed-port-l3-managed-ethernet-switches/flexfabric-5700-switch-products/hpe-flexfabric-5700-switch-series/p/7268889#:~:text=The%20HPE%20FlexFabric%205700%20Switch%20Series%20expands%20your%20enterprise%20network, CB)%20switch%20improving%20server%20connectivity.) [online; accessed June 17, 2020].
- [18] <https://cloud.kapostcontent.net/pub/ea2398f1-8d18-46bc-916f-8f890ce567a4/x465-series-data-sheet>. [online; accessed June 17, 2020].
- [19] <https://www.cisco.com/c/dam/en/us/products/switches/nexus-9000-series-switches/Nexus-9300-400-GE-Switches.html>. [online; accessed June 17, 2020].
- [20] <https://www.juniper.net/us/en/products-services/switching/qfx-series/datasheets/1000480.page>. [online; accessed June 17, 2020].
- [21] <https://www.arista.com/en/products/7060x-series>. [online; accessed June 17, 2020].
- [22] S. Aleksic and M. Fiorani, “The future of switching in data centers,” in “Optical Switching in Next Generation Data Centers,” , F. Testa and L. Pavesi, eds. (Springer International Publishing, 2019), chap. 15, pp. 301–328.
- [23] [https://www.thorlabs.com/newgrouppage9.cfm?objectgroup\\_id=334](https://www.thorlabs.com/newgrouppage9.cfm?objectgroup_id=334). [online; accessed June 17, 2020].

- [24] A. J. Zilkie, P. Srinivasan, A. Trita, T. Schrans, G. Yu, J. Byrd, D. A. Nelson, K. Muth, D. Lerose, M. Alalusi, K. Masuda, M. Ziebell, H. Abediasl, J. Drake, G. Miller, H. Nykanen, E. Kho, Y. Liu, H. Liang, H. Yang, F. H. Peters, A. S. Nagra, and A. G. Rickman, “Multi-micron silicon photonics platform for highly manufacturable and versatile photonic integrated circuits,” *IEEE Journal of Selected Topics in Quantum Electronics* **25**, 1–13 (2019).
- [25] A. La Porta, J. Weiss, R. Dangel, D. Jubin, N. Meier, J. Hofrichter, C. Caer, F. Horst, and B. J. Offrein, “Silicon photonics packaging for highly scalable optical interconnects,” in “2015 IEEE 65th Electronic Components and Technology Conference (ECTC),” (2015), pp. 1299–1304.
- [26] A. L. Porta, R. Dangel, D. Jubin, F. Horst, N. Meier, D. Chelladurai, B. W. Swatowski, A. C. Tomasik, K. Su, W. K. Weidner, and B. J. Offrein, “Optical coupling between polymer waveguides and a silicon photonics chip in the o-band,” in “2016 Optical Fiber Communications Conference and Exhibition (OFC),” (2016), pp. 1–3.
- [27] G. T. Reed, W. R. Headley, and C. E. J. Png, “Silicon photonics: the early years,” in “Optoelectronic Integration on Silicon II,” , vol. 5730, J. A. Kubby and G. E. Jabbour, eds., International Society for Optics and Photonics (SPIE, 2005), vol. 5730, pp. 1 – 18.
- [28] R. Soref, “The past, present, and future of silicon photonics,” *IEEE Journal of Selected Topics in Quantum Electronics* **12**, 1678–1687 (2006).
- [29] M. Hochberg and T. Baehr-Jones, “Towards fabless silicon photonics,” *Nature Photonics* **4**, 492 – 494 (2010).
- [30] A. E. Lim, J. Song, Q. Fang, C. Li, X. Tu, N. Duan, K. K. Chen, R. P. Tern, and T. Liow, “Review of silicon photonics foundry efforts,” *IEEE Journal of Selected Topics in Quantum Electronics* **20**, 405–416 (2014).
- [31] K. Giewont, K. Nummy, F. A. Anderson, J. Ayala, T. Barwicz, Y. Bian, K. K. Dezfulian, D. M. Gill, T. Houghton, S. Hu, B. Peng, M. Rakowski, S. Rauch III, J. C. Rosenberg, A. Sahin, I. Stobert, and A. Stricker, “300-mm monolithic silicon photonics foundry technology,” *IEEE Journal of Selected Topics in Quantum Electronics* **25**, 1–11 (2019).
- [32] N. M. Fahrenkopf, C. McDonough, G. L. Leake, Z. Su, E. Timurdogan, and D. D. Coolbaugh, “The aim photonics mpw: A highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits,” *IEEE Journal of Selected Topics in Quantum Electronics* **25**, 1–6 (2019).



- [33] A. Rahim, T. Spuesens, R. Baets, and W. Bogaerts, “Open-access silicon photonics: Current status and emerging initiatives,” *Proceedings of the IEEE* **106**, 2313–2330 (2018).
- [34] D. Knoll, S. Lischke, A. Awany, M. Kroh, E. Krune, C. Mai, A. Peczek, D. Petousi, S. Simon, K. Voigt, G. Winzer, R. Barth, and L. Zimmermann, “High-performance bimos si photonics platform,” in “2015 IEEE Bipolar/BiCMOS Circuits and Technology Meeting - BCTM,” (2015), pp. 88–96.
- [35] <http://www.advmf.com/wp-content/uploads/2019/08/AMF-Brochure-Web.pdf> (2020). [online; accessed June 18, 2020].
- [36] <https://www.imec-int.com/drupal/sites/default/files/2020-03/SILICON-PHOTONICS-V06.pdf> (2020). [online; accessed June 18, 2020].
- [37] T. Hirokawa, S. Pinna, J. Klamkin, J. F. Buckwalter, and C. L. Schow, “Energy efficiency analysis of coherent links for datacenters,” in “2019 IEEE Optical Interconnects Conference (OI),” (2019), pp. 1–2. [©2019 IEEE. Reprinted, with permission, from T. Hirokawa et al, Energy Efficiency Analysis of Coherent Links for Datacenters, 2019 IEEE Optical Interconnects Conference (OI), April 2019].
- [38] T. Hirokawa, S. Pinna, N. Hosseinzadeh, A. Maharry, H. Andrade, J. Liu, T. Meissner, S. Misak, G. Movaghar, L. Valenzuela, Y. Xia, S. Bhat, F. Gambini, J. Klamkin, A. A. M. Saleh, L. A. Coldren, J. F. Buckwalter, and C. L. Schow, “Analog coherent detection for energy efficient intra-data center links at 200 gbps per wavelength,” *Journal of Lightwave Technology* pp. 1–12 (2020). [©2020 IEEE. Reprinted, with permission, from T. Hirokawa et al., Analog Coherent Detection for Energy Efficient Intra-Data Center Links at 200 Gbps per Wavelength, *Journal of Lightwave Technology*, November 2020].
- [39] X. Pang, O. Ozolins, R. Lin, L. Zhang, A. Udalcovs, L. Xue, R. Schatz, U. Westergren, S. Xiao, W. Hu, G. Jacobsen, S. Popov, and J. Chen, “200 gbps/lane im/dd technologies for short reach optical interconnects,” *Journal of Lightwave Technology* **38**, 492–503 (2020).
- [40] J. K. Perin, A. Shastri, and J. M. Kahn, “Design of low-power dsp-free coherent receivers for data center links,” *Journal of Lightwave Technology* **35**, 4650–4662 (2017).
- [41] J. Cheng, C. Xie, Y. Chen, X. Chen, M. Tang, and S. Fu, “Comparison of coherent and imdd transceivers for intra datacenter optical interconnects,” in “2019 Optical Fiber Communications Conference and Exhibition (OFC),” (2019), pp. 1–3.
- [42] S. Ristic, A. Bhardwaj, M. J. Rodwell, L. A. Coldren, and L. A. Johansson, “An optical phase-locked loop photonic integrated circuit,” *Journal of Lightwave Technology* **28**, 526–538 (2010).

- [43] M. Lu, H. Park, E. Bloch, L. A. Johansson, M. J. Rodwell, and L. A. Coldren, “An integrated heterodyne optical phase-locked loop with record offset locking frequency,” in “OFC 2014,” (2014), pp. 1–3.
- [44] M. Lu, H. Park, J. S. Parker, E. Bloch, A. Sivananthan, Z. Griffith, L. A. Johansson, M. J. Rodwell, and L. A. Coldren, “A heterodyne optical phase-locked loop for multiple applications,” in “2013 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC),” (2013), pp. 1–3.
- [45] M. J. W. Rodwell, H. C. Park, M. Piels, M. Lu, A. Sivananthan, E. Bloch, Z. Griffith, M. Uteaga, L. Johansson, J. E. Bowers, and L. A. Coldren, “Optical phase-locking and wavelength synthesis,” in “2014 IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS),” (2014), pp. 1–4.
- [46] M. Lu, H. Park, E. Bloch, A. Sivananthan, J. S. Parker, Z. Griffith, L. A. Johansson, M. J. W. Rodwell, and L. A. Coldren, “An integrated 40 gbit/s optical costas receiver,” *Journal of Lightwave Technology* **31**, 2244–2253 (2013).
- [47] L. Szilagy, J. Pliva, R. Henker, D. Schoeniger, J. P. Turkiewicz, and F. Ellinger, “A 53-gbit/s optical receiver frontend with 0.65 pj/bit in 28-nm bulk-cmos,” *IEEE Journal of Solid-State Circuits* **54**, 845–855 (2019).
- [48] L. Kull, D. Luu, C. Menolfi, M. Brandli, P. A. Francese, T. Morf, M. Kossel, A. Cevrero, I. Ozkaya, and T. Toifl, “A 24.72-gs/s 8-b time-interleaved sar adc with 2.03.3-pj/conversion and  $\geq 30$  db snr at nyquist in 14-nm cmos finfet,” *IEEE Journal of Solid-State Circuits* **53**, 3508–3516 (2018).
- [49] K. Sun, G. Wang, Q. Zhang, S. Elahmadi, and P. Gui, “A 56-gs/s 8-bit time-interleaved adc with enob and bw enhancement techniques in 28-nm cmos,” *IEEE Journal of Solid-State Circuits* **54**, 821–833 (2019).
- [50] Y. Yue, Q. Wang, J. Yao, J. O’Neill, D. Pudvay, and J. Anderson, “400gbe technology demonstration using cfp8 pluggable modules,” *J. of App. Science* **8**, 2055 (2018).
- [51] M. Streshinsky, R. Ding, Y. Liu, A. Novack, Y. Yang, Y. Ma, X. Tu, E. K. S. Chee, A. E.-J. Lim, P. G.-Q. Lo, T. Baehr-Jones, and M. Hochberg, “Low power 50 gb/s silicon traveling wave mach-zehnder modulator near 1300 nm,” *Opt. Express* **21**, 30350–30357 (2013).
- [52] B. G. Lee, N. Dupuis, R. Rimolo-Donadio, T. N. Huynh, C. W. Baks, D. M. Gill, and W. M. J. Green, “Driver-integrated 56-gb/s segmented electrode silicon mach zehnder modulator using optical-domain equalization,” in “2017 Optical Fiber Communications Conference and Exhibition (OFC),” (2017), pp. 1–3.

- [53] K. Goi, H. Kusaka, A. Oka, K. Ogawa, T. Liow, X. Tu, G. Lo, and D. Kwong, “128-gb/s dp-qpsk using low-loss monolithic silicon iq modulator integrated with partial-rib polarization rotator,” in “OFC 2014,” (2014), pp. 1–3.
- [54] P. Dong, L. Chen, C. Xie, L. L. Buhl, and Y.-K. Chen, “50-gb/s silicon quadrature phase-shift keying modulator,” *Opt. Express* **20**, 21181–21186 (2012).
- [55] K. Goi, H. Kusaka, A. Oka, Y. Terada, K. Ogawa, T. Liow, X. Tu, G. Lo, and D. Kwong, “Dqpsk/qpsk modulation at 40?60 gb/s using low-loss nested silicon mach-zehnder modulator,” in “2013 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC),” (2013), pp. 1–3.
- [56] N. Kono, T. Kitamura, H. Yagi, N. Itabashi, T. Tatsumi, Y. Yamauchi, K. Fujii, K. Horino, S. Yamanaka, K. Tanaka, K. Yamaji, C. Fukuda, and H. Shoji, “Compact and low power dp-qpsk modulator module with inp-based modulator and driver ics,” in “2013 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC),” (2013), pp. 1–3.
- [57] E. Yamada, S. Kanazawa, A. Ohki, K. Watanabe, Y. Nasu, N. Kikuchi, Y. Shibata, R. Iga, and H. Ishii, “112-gb/s inp dp-qpsk modulator integrated with a silica-plc polarization multiplexing circuit,” in “OFC/NFOEC,” (2012), pp. 1–3.
- [58] N. Kikuchi, E. Yamada, Y. Shibata, and H. Ishii, “High-speed inp-based mach-zehnder modulator for advanced modulation formats,” in “2012 IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS),” (2012), pp. 1–4.
- [59] P. Evans, M. Fisher, R. Malendevich, A. James, P. Studenkov, G. Goldfarb, T. Valaitis, M. Kato, P. Samra, S. Corzine, E. Strzelecka, R. Salvatore, F. Sedgwick, M. Kuntz, V. Lal, D. Lambert, A. Dentai, D. Pavinski, J. Zhang, B. Behnia, J. Bostak, V. Dominic, A. Nilsson, B. Taylor, J. Rahn, S. Sanders, H. Sun, K. . Wu, J. Pleumeekers, R. Muthiah, M. Missey, R. Schneider, J. Stewart, M. Reffle, T. Butrie, R. Nagarajan, C. Joyner, M. Ziari, F. Kish, and D. Welch, “Multi-channel coherent pm-qpsk inp transmitter photonic integrated circuit (pic) operating at 112 gb/s per wavelength,” in “2011 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference,” (2011), pp. 1–3.
- [60] N. Hosseinzadeh, A. Jain, K. Ning, R. Helkey, and J. F. Buckwalter, “A 0.5-20 ghz rf silicon photonic receiver with 120 db?hz<sup>2</sup>/3 sfdR using broadband distributed im<sub>3</sub> injection linearization,” in “2019 IEEE Radio Frequency Integrated Circuits Symposium (RFIC),” (2019), pp. 99–102.

- [61] N. Hosseinzadeh, A. Jain, K. Ning, R. Helkey, and J. F. Buckwalter, “A linear microwave electro-optic front end with sige distributed amplifiers and segmented silicon photonic mach?zehnder modulator,” *IEEE Transactions on Microwave Theory and Techniques* **67**, 5446–5458 (2019).
- [62] N. Hosseinzadeh, A. Jain, K. Ning, R. Helkey, and J. F. Buckwalter, “A 1 to 20 ghz silicon-germanium low-noise distributed driver for rf silicon photonic mach-zehnder modulators,” in “2019 IEEE MTT-S International Microwave Symposium (IMS),” (2019), pp. 774–777.
- [63] D. Patel, S. Ghosh, M. Chagnon, A. Samani, V. Veerasubramanian, M. Osman, and D. V. Plant, “Design, analysis, and transmission system performance of a 41 ghz silicon photonic modulator,” *Opt. Express* **23**, 14263–14287 (2015).
- [64] A. Samani, E. El-Fiky, M. Morsy-Osman, R. Li, D. Patel, T. Hoang, M. Jacques, M. Chagnon, N. Abadía, and D. Plant, “Silicon photonic mach?zehnder modulator architectures for on chip pam-4 signal generation,” *Journal of Lightwave Technology* **37**, 2989–2999 (2019).
- [65] R. Ding, Y. Liu, Y. Ma, Y. Yang, Q. Li, A. E. Lim, G. Lo, K. Bergman, T. Baehr-Jones, and M. Hochberg, “High-speed silicon modulator with slow-wave electrodes and fully independent differential drive,” *Journal of Lightwave Technology* **32**, 2240–2247 (2014).
- [66] M. Li, L. Wang, X. Li, X. Xiao, and S. Yu, “Silicon intensity mach-zehnder modulator for single lane 100-gb/s applications,” *Photon. Res.* **6**, 109–116 (2018).
- [67] M. Rakowski, C. Meagher, K. Nummy, A. Aboketaf, J. Ayala, Y. Bian, B. Harris, K. McLean, K. McStay, A. Sahin, L. Medina, B. Peng, Z. Sowinski, A. Stricker, T. Houghton, C. Hedges, K. Giewont, A. Jacob, T. Letavic, D. Riggs, A. Yu, and J. Pellerin, “45nm cmos—silicon photonics monolithic technology (45clo) for next-generation, low power and high speed optical interconnects,” in “2020 Optical Fiber Communications Conference and Exhibition (OFC),” (2020), pp. 1–3.
- [68] T. Hirokawa, M. Saeidi, S. Pillai, A. Nguyen-Le, L. Theogarajan, A. Saleh, and C. Schow, “A wavelength-selective multiwavelength ring-assisted mach-zehnder interferometer switch,” *Journal of Lightwave Technology* **38**, 6292–6298 (2020). [©2020 IEEE. Reprinted, with permission, from T. Hirokawa et al., A Wavelength-Selective Multiwavelength Ring-Assisted Mach-Zehnder Interferometer Switch, *Journal of Lightwave Technology*, November 2020].
- [69] G. P. Agrawal, *Fiber-optic communication systems* (Wiley-Blackwell, 2011).
- [70] H. Yu, M. Pantouvaki, J. V. Campenhout, D. Korn, K. Komorowska, P. Dumon, Y. Li, P. Verheyen, P. Absil, L. Alloatti, D. Hillerkuss, J. Leuthold, R. Baets,

- and W. Bogaerts, “Performance tradeoff between lateral and interdigitated doping patterns for high speed carrier-depletion based silicon modulators,” *Opt. Express* **20**, 12926–12938 (2012).
- [71] A. Rahim, J. Goyvaerts, B. Szelag, J. Fedeli, P. Absil, T. Aalto, M. Harjanne, C. Littlejohns, G. Reed, G. Winzer, S. Lischke, L. Zimmermann, D. Knoll, D. Geuzebroek, A. Leinse, M. Geiselmann, M. Zervas, H. Jans, A. Stassen, C. Domínguez, P. Muñoz, D. Domenech, A. L. Giesecke, M. C. Lemme, and R. Baets, “Open-access silicon photonics platforms in europe,” *IEEE Journal of Selected Topics in Quantum Electronics* **25**, 1–18 (2019).
- [72] A. L. Giesecke, A. Prinzen, H. Fueser, C. Porschatis, H. Lerch, J. Bolten, S. Suckow, B. Chmielak, and T. Wahlbrink, “Ultra-efficient interleaved depletion modulators by using advanced fabrication technology,” in “ECOC 2016; 42nd European Conference on Optical Communication,” (2016), pp. 1–3.
- [73] A. A. M. Saleh, “Scaling-out data centers using photonics technologies,” in “Advanced Photonics for Communications,” (Optical Society of America, 2014), p. JM4B.5.
- [74] A. A. M. Saleh, A. S. P. Khope, J. E. Bowers, and R. C. Alferness, “Elastic wdm switching for scalable data center and hpc interconnect networks,” in “2016 21st OptoElectronics and Communications Conference (OECC) held jointly with 2016 International Conference on Photonics in Switching (PS),” (2016), pp. 1–3.
- [75] G. Michelogiannakis, Y. Shen, M. Y. Teh, X. Meng, B. Aivazi, T. Groves, J. Shalf, M. Glick, M. Ghobadi, L. Dennison, and K. Bergman, “Bandwidth steering in hpc using silicon nanophotonics,” in “Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis,” (Association for Computing Machinery, New York, NY, USA, 2019), SC ?19.
- [76] N. Dupuis, F. Doany, R. A. Budd, L. Schares, C. W. Baks, D. M. Kuchta, T. Hirokawa, and B. G. Lee, “A  $4 \times 4$  electrooptic silicon photonic switch fabric with net neutral insertion loss,” *Journal of Lightwave Technology* **38**, 178–184 (2020).
- [77] L. P. Barry, J. Wang, C. McArdle, and D. Kilper, “Optical switching in datacenters: Architectures based on optical circuit switching,” in “Optical Switching in Next Generation Data Centers,” , F. Testa and L. Pavesi, eds. (Springer International Publishing, 2019), chap. 2, pp. 23–44.
- [78] K. Suzuki, R. Konoike, N. Yokoyama, M. Seki, M. Ohtsuka, S. Saitoh, S. Suda, H. Matsuura, K. Yamada, S. Namiki, H. Kawashima, and K. Ikeda, “Nonduplicate polarization-diversity  $32 \times 32$  silicon photonics switch based on a sin/si double-layer platform,” *Journal of Lightwave Technology* **38**, 226–232 (2020).

- [79] K. Kwon, T. J. Seok, J. Henriksson, J. Luo, and M. C. Wu, “Large-scale silicon photonic switches,” in “2019 PhotonIcs Electromagnetics Research Symposium - Spring (PIERS-Spring),” (2019), pp. 268–273.
- [80] T. Hirokawa, A. Netherton, M. Saeidi, H. Andrade, L. Theogarajan, J. E. Bowers, A. A. Saleh, and C. L. Schow, “An all-optical wavelength-selective o-band chip-scale silicon photonic switch,” in “2020 Conference on Lasers and Electro-Optics (CLEO),” (Optical Society of America, 2020).
- [81] N. Hosseinzadeh, K. Fang, L. Valenzuela, C. Schow, and J. Buckwalter, “A 50-gb/s optical transmitter based on co-design of a 45-nm cmos soi distributed driver and 90-nm silicon photonic mach-zehnder modulator,” in “IEEE MTT-S Int. Microw. Symp.”, (2020).
- [82] L. Chrostowski and M. Hochberg, *Silicon Photonics Design* (Cambridge University Press, 2015).
- [83] L. A. Coldren, S. W. Corzine, and M. L. Mashanovitch, *Diode Lasers and Photonic Integrated Circuits* (John Wiley & Sons, 2012).
- [84] W. Bogaerts, P. D. Heyn, T. V. Vaerenbergh, K. D. Vos, S. K. Selvaraja, T. Claes, P. Dumon, P. Bienstman, D. V. Thourhout, and R. Baets, “Silicon microring resonators,” *Laser & Photonics Review* **6**, 47–73 (2012).
- [85] C. L. Manganelli, P. Pintus, F. Gambini, D. Fowler, M. Fournier, S. Faralli, C. Kopp, and C. J. Oton, “Large-fsr thermally tunable double-ring filters for wdm applications in silicon photonics,” *IEEE Photonics Journal* **9**, 1–10 (2017).
- [86] D. G. Rabus, *Integrated Ring Resonators* (Springer, 2007).
- [87] B. G. Lee and N. Dupuis, “Silicon photonic switch fabrics: Technology and architecture,” *Journal of Lightwave Technology* **37**, 6–20 (2019).
- [88] R. Soref and B. Bennett, “Electrooptical effects in silicon,” *IEEE Journal of Quantum Electronics* **23**, 123–129 (1987).
- [89] W. Jin, E. J. Stanton, N. Volet, R. G. Polcawich, D. Baney, P. Morton, and J. E. Bowers, “Piezoelectric tuning of a suspended silicon nitride ring resonator,” in “2017 IEEE Photonics Conference (IPC),” (2017), pp. 117–118.
- [90] W. Jin, R. G. Polcawich, P. A. Morton, and J. E. Bowers, “Piezoelectrically tuned silicon nitride ring resonator,” *Opt. Express* **26**, 3174–3187 (2018).
- [91] C. Zhang, D. Liang, G. Kurczveil, J. E. Bowers, and R. G. Beausoleil, “Thermal management of hybrid silicon ring lasers for high temperature operation,” *IEEE Journal of Selected Topics in Quantum Electronics* **21**, 385–391 (2015).

- [92] D. Liang, M. Fiorentino, T. Okumura, H.-H. Chang, D. T. Spencer, Y.-H. Kuo, A. W. Fang, D. Dai, R. G. Beausoleil, and J. E. Bowers, “Electrically-pumped compact hybrid silicon microring lasers for optical interconnects,” *Opt. Express* **17**, 20355–20364 (2009).
- [93] D. Liang, M. Fiorentino, S. Srinivasan, J. E. Bowers, and R. G. Beausoleil, “Low threshold electrically-pumped hybrid silicon microring lasers,” *IEEE Journal of Selected Topics in Quantum Electronics* **17**, 1528–1533 (2011).
- [94] S. J. Emelett and R. A. Soref, “Synthesis of dual-microring-resonator cross-connect filters,” *Opt. Express* **13**, 4439–4456 (2005).
- [95] S. J. Emelett and R. A. Soref, “Electro-optical and optical-optical switching of dual microring resonator waveguide systems,” in “Advanced Optical and Quantum Memories and Computing II,” , vol. 5735, H. J. Coufal, Z. U. Hasan, and A. E. Craig, eds., International Society for Optics and Photonics (SPIE, 2005), vol. 5735, pp. 14 – 24.
- [96] S. J. Emelett and R. Soref, “Design and simulation of silicon microring optical routing switches,” *Journal of Lightwave Technology* **23**, 1800–1807 (2005).
- [97] Y. Huang, Q. Cheng, Y. Hung, H. Guan, X. Meng, A. Novack, M. Streshinsky, M. Hochberg, and K. Bergman, “Multi-stage  $8 \times 8$  silicon photonic switch based on dual-microring switching elements,” *Journal of Lightwave Technology* **38**, 194–201 (2020).
- [98] H. Yang, Q. Cheng, R. Chen, and K. Bergman, “Polarization-diversity microring-based optical switch fabric in a switch-and-select architecture,” in “2020 Optical Fiber Communications Conference and Exhibition (OFC),” (2020), pp. 1–3.
- [99] T. Hirokawa, A. Maharry, R. Helkey, J. E. Bowers, A. A. M. Saleh, and C. L. Schow, “Demonstration of a spectrally-partitioned 4x4 crossbar switch with 3 drops per cross-point,” in “2019 24th OptoElectronics and Communications Conference (OECC) and 2019 International Conference on Photonics in Switching and Computing (PSC),” (2019), pp. 1–3. [©2019 IEEE. Reprinted, with permission, from T. Hirokawa et al., A Spectrally-Partitioned Crossbar Switch with Three Drops Per Cross-Point Controlled with a Driver, 2019 24th OptoElectronics and Communications Conference (OECC) and 2019 International Conference on Photonics in Switching and Computing (PSC), July 2019].
- [100] T. Hirokawa, M. Saeidi, A. Maharry, R. Helkey, J. E. Bowers, L. Theogarajan, A. A. M. Saleh, and C. L. Schow, “A spectrally-partitioned crossbar switch with three drops per cross-point controlled with a driver,” in “2019 IEEE Photonics Conference (IPC),” (2019), pp. 1–2. [©2019 IEEE. Reprinted, with permission, from T. Hirokawa et al., Demonstration of a Spectrally-Partitioned 4x4 Crossbar

- Switch with 3 Drops per Cross-point, 2019 IEEE Photonics Conference (IPC), September 2019].
- [101] G. Porter, R. Strong, N. Farrington, A. Forencich, P. Chen-Sun, T. Rosing, Y. Fainman, G. Papen, and A. Vahdat, “Integrating microsecond circuit switching into the data center,” *SIGCOMM Comput. Commun. Rev.* **43**, 447–458 (2013).
  - [102] R. Yu, S. Cheung, Y. Li, K. Okamoto, R. Proietti, Y. Yin, and S. J. B. Yoo, “A scalable silicon photonic chip-scale optical switch for high performance computing systems,” *Opt. Express* **21**, 32655–32667 (2013).
  - [103] Q. Cheng, S. Rumley, M. Bahadori, and K. Bergman, “Photonic switching in high performance datacenters,” *Opt. Express* **26**, 16022–16043 (2018).
  - [104] A. S. P. Khope, A. A. M. Saleh, J. E. Bowers, and R. C. Alferness, “Elastic wdm crossbar switch for data centers,” in “2016 IEEE Optical Interconnects Conference (OI),” (2016), pp. 48–49.
  - [105] A. S. P. Khope, M. Saeidi, R. Yu, X. Wu, A. M. Netherton, Y. Liu, Z. Zhang, Y. Xia, G. Fleeman, A. Spott, S. Pinna, C. Schow, R. Helkey, L. Theogarajan, R. C. Alferness, A. A. M. Saleh, and J. E. Bowers, “Multi-wavelength selective crossbar switch,” *Opt. Express* **27**, 5203–5216 (2019).
  - [106] K. Oda, N. Takato, and H. Toba, “A wide-fsr waveguide double-ring resonator for optical fdm transmission systems,” *Journal of Lightwave Technology* **9**, 728–736 (1991).
  - [107] Y. Huang, Q. Cheng, Y. Hung, H. Guan, A. Novack, M. Streshinsky, M. Hochberg, and K. Bergman, “Dual-microring resonator based 8 x 8 silicon photonic switch,” in “2019 Optical Fiber Communications Conference and Exhibition (OFC),” (2019), pp. 1–3.
  - [108] K. Kwon, T. J. Seok, J. Henriksson, J. Luo, L. Ochikubo, J. Jacobs, R. S. Muller, and M. C. Wu, “128 × 128 silicon photonic mems switch with scalable row/column addressing,” in “2018 Conference on Lasers and Electro-Optics (CLEO),” (2018), pp. 1–2.
  - [109] N. Dupuis, B. G. Lee, A. V. Rylyakov, D. M. Kuchta, C. W. Baks, J. S. Orcutt, D. M. Gill, W. M. J. Green, and C. L. Schow, “Design and fabrication of low-insertion-loss and low-crosstalk broadband 2 × 2 mach–zehnder silicon photonic switches,” *J. Lightwave Technol.* **33**, 3597–3606 (2015).
  - [110] J. L.R. Ford and D. R. Fulkerson, “Maximal flow through a network,” *Can. J. of Mathematics* **8**, 399–404 (1956).



- [111] J. Akiyama and M. Kano, *Factors and Factorizations of Graphs* (Springer, Berlin, Heidelberg, 2011).
- [112] T. Hirokawa, M. Saeidi, L. Theogarajan, A. A. M. Saleh, and C. L. Schow, “Ring-assisted mach-zehnder interferometer switch with multiple rings per switch element,” in “Proc. SPIE 11286, Optical Interconnects XX, 1128612,” (2020).
- [113] Xiaobo Xie, J. Khurgin, Jin Kang, and F. . Chow, “Linearized mach-zehnder intensity modulator,” *IEEE Photonics Technology Letters* **15**, 531–533 (2003).
- [114] Z. Guo, L. Lu, L. Zhou, L. Shen, and J. Chen, “16 x 16 silicon optical switch based on dual-ring-assisted mach-zehnder interferometers,” *J. of Lightwave Technology* **36**, 225–232 (2018).
- [115] S. Zhao, L. Lu, L. Zhou, D. Li, Z. Guo, and J. Chen, “16×16 silicon mach-zehnder interferometer switch actuated with waveguide microheaters,” *Photon. Res.* **4**, 202–207 (2016).
- [116] K. Suzuki, G. Cong, K. Tanizawa, S.-H. Kim, S. Namiki, and H. Kawashima, “50-db extinction-ratio in 2×2 silicon optical switch with variable splitter,” in “CLEO: 2014,” (Optical Society of America, 2014), p. SM4G.2.
- [117] D. Celo, D. J. Goodwill, Jia Jiang, P. Dumais, Chunshu Zhang, Fei Zhao, Xin Tu, Chunhui Zhang, Shengyong Yan, Jifang He, Ming Li, Wanyuan Liu, Yuming Wei, Dongyu Geng, H. Mehrvar, and E. Bernier, “32×32 silicon photonic switch,” in “2016 21st OptoElectronics and Communications Conference (OECC) held jointly with 2016 International Conference on Photonics in Switching (PS),” (2016), pp. 1–3.
- [118] N. Dupuis, A. V. Rylyakov, C. L. Schow, D. M. Kuchta, C. W. Baks, J. S. Orcutt, D. M. Gill, W. M. J. Green, and B. G. Lee, “Ultralow crosstalk nanosecond-scale nested 2 × 2 mach-zehnder silicon photonic switch,” *Opt. Lett.* **41**, 3002–3005 (2016).
- [119] P. Kolman, “On nonblocking properties of the beneš network,” in “Algorithms — ESA’ 98,” , G. Bilardi, G. F. Italiano, A. Pietracaprina, and G. Pucci, eds. (Springer Berlin Heidelberg, Berlin, Heidelberg, 1998), pp. 259–270.
- [120] J. S. Orcutt, A. Khilo, C. W. Holzwarth, M. A. Popović, H. Li, J. Sun, T. Bonifield, R. Hollingsworth, F. X. Kärtner, H. I. Smith, V. Stojanović, and R. J. Ram, “Nanophotonic integration in state-of-the-art cmos foundries,” *Opt. Express* **19**, 2335–2346 (2011).

- [121] K. Suzuki, R. Konoike, N. Yokoyama, M. Seki, M. Ohtsuka, S. Saitoh, S. Suda, H. Matsuura, K. Yamada, S. Namiki, H. Kawashima, and K. Ikeda, “Polarization-diversity 32 x 32 si photonics switch with non-duplicate diversity circuit in double-layer platform,” in “Optical Fiber Communication Conference (OFC) 2019,” (Optical Society of America, 2019), p. Th1E.2.
- [122] A. S. P. Khope, T. Hirokawa, A. M. Netherton, M. Saeidi, Y. Xia, N. Volet, C. Schow, R. Helkey, L. Theogarajan, A. A. M. Saleh, J. E. Bowers, and R. C. Alferness, “On-chip wavelength locking for photonic switches,” *Opt. Lett.* **42**, 4934–4937 (2017).
- [123] H. Ballani, P. Costa, I. Haller, K. Jozwik, K. Shi, B. Thomsen, and H. Williams, “Bridging the last mile for optical switching in data centers,” in “2018 Optical Fiber Communications Conference and Exposition (OFC),” (2018), pp. 1–3.
- [124] D. K. Schroder, *Semiconductor Material and Device Characterization* (John Wiley & Sons, Inc., 2006), 3rd ed.
- [125] N. M. Fahrenkopf, C. McDonough, G. L. Leake, Z. Su, E. Timurdogan, and D. D. Coolbaugh, “The aim photonics mpw: A highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits,” *IEEE Journal of Selected Topics in Quantum Electronics* **25**, 1–6 (2019).
- [126] M. A. Tran, T. Komljenovic, J. C. Hulme, M. L. Davenport, and J. E. Bowers, “A robust method for characterization of optical waveguides and couplers,” *IEEE Photonics Technology Letters* **28**, 1517–1520 (2016).
- [127] S. Sakamoto, “Reduced voltage substrate-removed gallium arsenide/aluminum gallium arsenide electro-optic modulators,” Ph.D. thesis, University of California Santa Barbara (2006).