

University of Kentucky

UKnowledge

Theses and Dissertations--Electrical and
Computer Engineering

Electrical and Computer Engineering

2024

Reinventing Integrated Photonic Devices and Circuits for High Performance Communication and Computing Applications

Venkata Sai Praneeth Karempudi

University of Kentucky, kvspraneeth@gmail.com

Author ORCID Identifier:

 <https://orcid.org/0000-0003-0609-6894>

Digital Object Identifier: <https://doi.org/10.13023/etd.2024.05>

[Right click to open a feedback form in a new tab to let us know how this document benefits you.](#)

Recommended Citation

Karempudi, Venkata Sai Praneeth, "Reinventing Integrated Photonic Devices and Circuits for High Performance Communication and Computing Applications" (2024). *Theses and Dissertations--Electrical and Computer Engineering*. 197.

https://uknowledge.uky.edu/ece_etds/197

This Doctoral Dissertation is brought to you for free and open access by the Electrical and Computer Engineering at UKnowledge. It has been accepted for inclusion in Theses and Dissertations--Electrical and Computer Engineering by an authorized administrator of UKnowledge. For more information, please contact UKnowledge@lsv.uky.edu.

STUDENT AGREEMENT:

I represent that my thesis or dissertation and abstract are my original work. Proper attribution has been given to all outside sources. I understand that I am solely responsible for obtaining any needed copyright permissions. I have obtained needed written permission statement(s) from the owner(s) of each third-party copyrighted matter to be included in my work, allowing electronic distribution (if such use is not permitted by the fair use doctrine) which will be submitted to UKnowledge as Additional File.

I hereby grant to The University of Kentucky and its agents the irrevocable, non-exclusive, and royalty-free license to archive and make accessible my work in whole or in part in all forms of media, now or hereafter known. I agree that the document mentioned above may be made available immediately for worldwide access unless an embargo applies.

I retain all other ownership rights to the copyright of my work. I also retain the right to use in future works (such as articles or books) all or part of my work. I understand that I am free to register the copyright to my work.

REVIEW, APPROVAL AND ACCEPTANCE

The document mentioned above has been reviewed and accepted by the student's advisor, on behalf of the advisory committee, and by the Director of Graduate Studies (DGS), on behalf of the program; we verify that this is the final, approved version of the student's thesis including all changes required by the advisory committee. The undersigned agree to abide by the statements above.

Venkata Sai Praneeth Karempudi, Student

Dr. Ishan Thakkar, Major Professor

Dr. Daniel Lau, Director of Graduate Studies

Reinventing Integrated Photonic Devices and Circuits for High Performance
Communication and Computing Applications

DISSERTATION

A dissertation submitted in partial
fulfillment of the requirements for
the degree of Doctor of Philosophy
in the College of Engineering at the
University of Kentucky

By
Venkata Sai Praneeth Karempudi
Lexington, Kentucky

Director: Dr. Ishan G Thakkar,
Assistant Professor of Electrical and Computer Engineering
Lexington, Kentucky 2023

Copyright© Venkata Sai Praneeth Karempudi 2023

ABSTRACT OF DISSERTATION

Reinventing Integrated Photonic Devices and Circuits for High Performance Communication and Computing Applications

The long-standing technological pillars for computing systems evolution, namely Moore's law and Von Neumann architecture, are breaking down under the pressure of meeting the capacity and energy efficiency demands of computing and communication architectures that are designed to process modern data-centric applications related to Artificial Intelligence (AI), Big Data, and Internet-of-Things (IoT). In response, both industry and academia have turned to 'more-than-Moore' technologies for realizing hardware architectures for communication and computing. Fortunately, Silicon Photonics (SiPh) has emerged as one highly promising 'more-than-Moore' technology. Recent progress has enabled SiPh-based interconnects to outperform traditional electrical interconnects, offering advantages like high bandwidth density, near-light speed data transfer, distance-independent bitrate, and low energy consumption. Furthermore, SiPh-based electro-optic (E-O) computing circuits have exhibited up to two orders of magnitude improvements in performance and energy efficiency compared to their electronic counterparts. Thus, SiPh stands out as a compelling solution for creating high-performance and energy-efficient hardware for communication and computing applications.

Despite their advantages, SiPh-based interconnects face various design challenges that hamper their reliability, scalability, performance, and energy efficiency. These include limited optical power budget (OPB), high static power dissipation, crosstalk noise, fabrication and on-chip temperature variations, and limited spectral bandwidth for multiplexing. Similarly, SiPh-based E-O computing circuits also face several challenges. Firstly, the E-O circuits for simple logic functions lack the all-electrical input handling, raising hardware area and complexity. Secondly, the E-O arithmetic circuits occupy vast areas (at least $100\times$) while hardly achieving more than 60% hardware utilization, versus CMOS implementations, leading to high idle times, and non-amortizable area and static power overheads. Thirdly, the high area overhead of E-O circuits hinders them from achieving high spatial parallelism on-chip. This is because the high area overhead limits the count of E-O circuits that can be implemented on a reticle-size limited chip.

My research offers significant contributions to address the aforementioned challenges. For SiPh-based interconnects, my contributions focus on enhancing OPB by mitigating crosstalk noise, addressing the optical non-linearity-related issues through the development of Silicon-on-Sapphire-based photonic interconnects, exploring multi-level signaling, and evaluating various device-level design pathways. This enables the design of high throughput ($>1\text{Tbps}$) and energy-efficient ($<1\text{pJ/bit}$) SiPh interconnects. In the context of SiPh-based E-O circuits, my contributions include the design of a microring-based polymorphic E-O logic gate, a hybrid time-amplitude analog optical modulator, and an indium tin oxide-based silicon nitride microring modulator and a weight bank for neural network computations. These designs significantly reduce the area overhead of current E-O computing circuits while enhancing the energy-efficiency, and hardware utilization.

KEYWORDS: Photonic Devices, Photonic Links, On-Chip Communication, Photonic Computing, Optical Power Budget, Aggregate Data rate

Author's signature: Venkata Sai Praneeth Karempudi

Date: December 19, 2023

Reinventing Integrated Photonic Devices and Circuits for High Performance
Communication and Computing Applications

By
Venkata Sai Praneeth Karempudi

Director of Dissertation: Dr. Ishan G Thakkar

Director of Graduate Studies: Dr. Daniel Lau

Date: December 19, 2023

Dedicated to my Parents and Teachers

*Om Asato Maa Sad-Gamaya
Tamaso Maa Jyotir-Gamaya
Mrtyor-Maa Amrtam Gamaya
Om Shaantih Shaantih Shaantih*

Derived From: Brihadaranyaka Upanishad (Discourse 8)

ACKNOWLEDGMENTS

I express my deepest gratitude to my advisor, Dr. Ishan Thakkar, whose unwavering encouragement, steadfast support, and invaluable guidance have been pivotal throughout the entirety of my PhD journey. His insightful advice has significantly enriched my research experience, and I am beyond grateful for the mentorship that has shaped this academic endeavor.

I extend my gratitude to my esteemed committee members, Dr. Todd Hastings, Dr. Janet Lumpp, and Dr. Douglas Strachan, for their invaluable support and guidance.

I also express my sincere appreciation to my colleagues in the Unconventional Computing Architectures and Technologies (UCAT) lab—Chao-Hsuan Huang, Sairam Sri Vatsavai, Oluwaseun Alo, Samrat Patel, Bobby Bose, and David Phippen for their help and support.

I also extend my heartfelt gratitude to my parents for their unwavering encouragement and boundless love that have played a profound role in shaping me into the person I am today.

Lastly, I extend heartfelt gratitude to my exceptional friends Sai Praneeth Reddy Navari, Amit Degada, Bhamiti Sharma, Shravani Prakhya, Deepak Kumar, Srinivasa Rao, Prakash Dhungana, Ankan Bhattacharya, Rajdeep Nath, and Ashutosh Timilsina. Their unwavering help and support have been instrumental in navigating the challenges and joys of my PhD journey.

TABLE OF CONTENTS

Acknowledgments	iii
List of Tables	vii
List of Figures	ix
Chapter 1 Introduction	1
1.1 Overview	1
1.2 On-Chip Communication With Silicon Photonics	2
1.3 Computing With Silicon Photonics	14
1.4 Contributions	18
Chapter 2 Mitigating Inter-Channel Crosstalk Non-Uniformity in Microring Filter Arrays of Wavelength-Multiplexed Photonic NoCs	20
2.1 Introduction and Motivation	20
2.2 Proposed Method	23
2.3 Summary	24
Chapter 3 Redesigning Photonic Interconnects with Silicon-on-Sapphire De- vice Platform for Ultra-Low-Energy On-Chip Communication	25
3.1 Introduction	25
3.2 Motivation	26
3.3 Modelling of SOS-based Devices	27
3.4 Link-Level Modelling and Analysis	32
3.5 System-Level Evaluation	36
3.6 Related Work	38
3.7 Overheads and Challenges	39
3.8 Summary	39
Chapter 4 Photonic Networks-on-Chip Employing Multilevel Signaling: A Cross- Layer Comparative Study	40
4.1 Introduction	40
4.2 Background: Various Designs of OOK and 4-PAM Modulators from Prior Work	42
4.3 Systematic Analysis of Photonic Links with Various Modulator Imple- mentations	47
4.4 Design Tradeoffs For Photonic Links	54
4.5 System-Level Evaluation	72
4.6 Summary	82

Chapter 5	An Analysis of Various Design Pathways Towards Multi-Terabit Photonic On-Interposer Interconnects	83
5.1	Introduction	83
5.2	Preliminaries	85
5.3	Identifying the Key Design Pathways Towards Multi-Terabit On-SiPh-Interposer Links	90
5.4	Link-level evaluation	97
5.5	System-Level Evaluation	104
5.6	Key Results	114
5.7	Summary	116
Chapter 6	A Polymorphic Electro-Optic Logic Gate for High-Speed Reconfigurable Computing Circuits	118
6.1	Introduction	118
6.2	MRR-Based Polymorphic Electro-Optic Logic Gate (MRR-PEOLG)	120
6.3	Transient Analysis	125
6.4	Performance Analysis	127
6.5	Comparison with E-O Circuits from Prior Work	129
6.6	Summary	130
Chapter 7	A Hybrid Time-Amplitude Analog Photonic GEMM Accelerator	131
7.1	Introduction	131
7.2	Background	134
7.3	A Hybrid Time-Amplitude Analog Optical Modulator (TAOM)	138
7.4	Analysis of TAOM-Enabled Parallel Multiplier Circuit	148
7.5	A Hybrid Time-Amplitude Analog Optical Accelerator	154
7.6	Summary	161
Chapter 8	Indium Tin Oxide Based Silicon Nitride Microring Modulators for High Performance Photonic Integrated Circuits	163
8.1	Introduction	163
8.2	Related Work and Motivation	164
8.3	Design of Our SiN-on-SiO ₂ Modulators	165
8.4	Summary	174
Chapter 9	A Low-Dissipation and Scalable GEMM Accelerator with Silicon Nitride Photonics	175
9.1	Introduction	175
9.2	Preliminaries	176
9.3	SiNPhAR Architecture	178
9.4	Evaluation and Discussion	182
9.5	Summary	188
Chapter 10	Conclusions and Future Work	189
10.1	Conclusions	189

10.2 Future Work	192
Appendix	195
Bibliography	196
Vita	223
Education	223
Invention Disclosure	223
Publications	223

LIST OF TABLES

1.1	Typical values of various losses in a photonic link [100, 154]	13
3.1	Various types of losses and optical parameters for SOS and SOI devices.	29
3.2	Considered Q-factor, FSR, and MAOP values for our analyzed SOSPh and SOI Ph links.	33
3.3	N_λ and bitrate for different variants of CLOS PNoC.	35
4.1	Number of instances, dynamic energy-per-bit (EPB), and static power values for various hardware components required for implementing a DWDM photonic link with N_λ channels using various signaling methods and modulator types. PS is packet size in bits.	50
4.2	Sample values of per-instance EPB for modulator driver (E^{Mod}), serialization + deserialization (E^{SerDes}), trans-impedance op-amps ($E^{TI-OPAMP}$) and per-MR static power for MR tuning control circuit (P^{TC}) and microheater (P^{\muheater}).	52
4.3	Definitions of various link design parameters and notations from Eqs. 4.1 - 4.15.	57
4.3	Definitions of various link design parameters and notations from Eqs. 4.1 - 4.15. (Continued)	58
4.4	Typical values (if any) of various link design parameters and notations from Eqs. 4.1-4.15.	59
4.5	Optimal N_λ , bitrate (BR), aggregated datarate ($N_\lambda \times BR$), power budget (P_{dB}^B), $PP_{dB} + 10\log(N_\lambda)$, detector sensitivity (S), and optical laser power ($= PP_{dB} + S$) values for different datarate-BER balanced variants of CLOS and SWIFT links. S varies across different links because S depends on BR.	65
4.6	Optimal N_λ , bitrate (BR), aggregated datarate ($N_\lambda \times BR$), power budget (P_{dB}^B), $PP_{dB} + 10\log(N_\lambda)$, detector sensitivity (S), and optical laser power ($= PP_{dB} + S$) values for different BER-optimal variants of CLOS and SWIFT links. S varies across different links because S depends on BaR.	71
4.7	Parameters for modeling the 3D organizations of our evaluated PNoCs.	75
5.1	Device-layer, circuit-layer, and system-layer features that influence the performance and energy efficiency of inter-chiplet on-SiPhI interconnects.	87
5.2	Various design pathways, their target device parameters with corresponding projected values, and the likelihood of achieving the projected parameter values either in the short term (within 5 years) or long term (>5 years)	90
5.3	Design Parameters for our considered SiPh fabrication processes	97
5.4	Inter-chiplet variants derived from 45nm SOI CMOS, 32nm SOI CMOS and Deposited poly-Si platforms.	104

5.5	Insertion loss evaluated for various design pathways implemented on the 45nm SOI CMOS platform. This evaluation encompasses intra-chiplet and inter-chiplet networks within NUPLet that employ SWMR and MWMR crossbar topologies, respectively	107
6.1	Power consumed in the microheaters, programmed detuning, and required resonance shifting, used to program our MRR-PEOLG for implementing different logic functions.	125
6.2	Performance comparison of E-O circuits. A=Area, E=Energy, L=Latency	130
7.1	Definitions and values of various parameters used in Eqs. 1,2 and 3 (from [8, 263, 262]) to perform the scalability analysis.	156
7.2	Hardware description of various optical and electrical components utilized in AMW and MAW TPCs, and TAOM-TC.	158
7.3	Power consumption of various hardware components utilized in AMW and MAW TPCs, and TAOM-TC.	158
7.4	Convolutional parameters corresponding to various benchmarks utilized in the performance analysis of our TAOM-TC.	159
7.5	Definition of various convolutional parameters corresponding to different benchmarks utilized in the performance analysis of our TAOM-TC. . . .	160
8.1	Free-carrier concentration (N), real index ($\text{Re}(\eta_{ITO})$), and imaginary index ($\text{Im}(\eta_{ITO})$) for the ITO accumulation layer in our modulator. The real and imaginary effective index ($\text{Re}(\eta_{eff})$, $\text{Im}(\eta_{eff})$), operating voltage (V), and induced resonance shift ($\Delta\lambda_r$) for our modulator (ITO-SiN-ITO stack). .	169
8.2	Free-carrier concentration (N), real index ($\text{Re}(\eta_{ITO})$), and imaginary index ($\text{Im}(\eta_{ITO})$) for the ITO accumulation layer in our modulator. The real and imaginary effective index ($\text{Re}(\eta_{eff})$, $\text{Im}(\eta_{eff})$), operating voltage (V), and induced resonance shift ($\Delta\lambda_r$) for our modulator (ITO-SiO ₂ -ITO stack). .	169
8.3	Modulation bandwidth (MB) (optical (O) and Electrical (E)), modulation efficiency (ME), FSR and energy efficiency (EE) corresponding to various SiN based MRR modulators (modulator type (MT)) from prior works obtained from simulations (*) and experiments, compared with our simulated SiN-on-SiO ₂ (ITO-SiO ₂ -ITO stack) MRR modulator.	173
9.1	Definition and values of various parameters used in Eq. 1, Eq. 2, and Eq. 3 (from [8]) for the scalability analysis.	184
9.2	TPC size (N) and TPC Count ($\#$) at 4-bit precision across various data rates for various accelerator architectures.	186
9.3	Accelerator Peripherals and TPC Parameters [263].	186

LIST OF FIGURES

1.1	(a) Intel Xeon Phi – 72 core processor [165], (b) ARM ThunderX2 – 54 core processor [160].	2
1.2	Physical-layer layout of (a) CLOS PNoC [116], (b) SwiftNoC [56] and (c) LumiNoC [141]	3
1.3	An on-chip photonic link.	5
1.4	Illustration of the total internal reflection of light in the longitudinal cross-section of a silicon-photonic waveguide.	5
1.5	(a) A tunable MRM, (b) MRR with an integrated micro heater and (c) Shift in resonance of the MRR due to heating (right side of the MRR passband) and carrier tuning (left side of the MRR passband)	6
1.6	A tunable MRR as (a) an OOK modulator, (b) Switch and (c) Wavelength filter	7
1.7	MRR filter at the receiver side to detect its resonance wavelength	7
1.8	Illustration of (a) Inter-channel crosstalk and (b) Intra-channel crosstalk (reproduced from [18])	8
1.9	(a) Inter-modulation crosstalk and (b) Spectral view of spectral distortion and inter-channel crosstalk at the detector side	9
1.10	Illustration of drift in resonance passband of an MRR due to PV.	10
1.11	Wafer-level variation patterns (left) and corresponding histograms (right) of (a) Resonance wavelength, (b) Quality factor and (c) Extinction ratio of MRRs.	11
1.12	Illustration of FSR of an MRR.	12
1.13	Various types of signal losses in a silicon photonic link	13
1.14	(a) Intel Stratix FPGA [108], (b) Xilinx Vertex FPGA [278], and (c) Cerebra’s Wafer Scale Engine (WSE) [46]	15
1.15	(a) An add-drop MR-based E-O AND logic gate [225], (b) a PCM-enabled MR-based E-O XNOR gate [297], (c) an MR-based E-O XOR/XNOR logic gate [291], and (d) an MR-based polymorphic E-O logic gate [286] demonstrated in the literature.	16
2.1	An array of MRR filters at the receiver end of a silicon photonic DWDM link. The heights of the crosstalk arrows are proportional to the corresponding power penalty values.	20
2.2	Crosstalk penalty distribution across the MRR filter array for two different cases. The values are obtained for 50GHz spacing and MRR quality factor of 8000.	21
2.3	Optical laser power per channel, for the baseline and reshuffled cases, evaluated for 50GHz spacing, MR quality factor of 8000, and total link power budget of 100mW.	21

2.4	(a) Uniform distribution of crosstalk penalty and non-uniform distribution of quality factors, across the MR filters, and (b) total link-level optical laser power for three different cases.	22
3.1	(a) Distribution of optical power budget (OPB), and (b) Best Achievable aggregate data rate ($\#DWD\text{M channels } (N_\lambda) \times \text{channel bitrate}$) vs energy-per-bit (EPB), for our analyzed photonic links based on SOIPh platforms from prior works and our target (preferred) photonic platform	28
3.2	(a),(b) Depiction of the cross-sectional dimensions of and the coupling gap size (g) between a waveguide and an MR; and MR coupling coefficient (k) as a function of gap size (g) and MR radius (R) for (c) SOSPh platform and (d) SOIPh platform.	29
3.3	Quality factor (Q-factor) based on coupling gap size (g) and MR radius (R) for (a) SOSPh platform, and (b) SOIPh platform.	30
3.4	MR tuning energy versus magnitude of wavelength tuning ($\Delta\lambda_m$) for SOSPh and SOIPh MRs.	32
3.5	Aggregate datarate and total energy-per-bit (EPB) values for (a) SOSPh links, and (b) SOIPh links, for different Q-factor, FSR and MAOP value combinations from Table 3.2. The optical losses, laser efficiency, and other device parameters for this analysis are taken from [12] and [16].	34
3.6	Distribution of optical power budget (OPB) for different SOIPh and SOSPh link designs from Table 3.3.	36
3.7	(a) Average energy-per-bit (EPB), and (b) packet latency comparisons for different variants of CLOS PNoC across PARSEC benchmarks. All results are normalized to the baseline CLOS-SOI PNoC results.	37
4.1	Illustration of (a) an MR-based on-off keying (OOK) modulator, and (b) the modulator's resonance passbands and optical transmission levels. . .	43
4.2	4-PAM-SS modulator designs. (a) Design from [248] with two parallel OOK MR modulators and multi-mode interference (MMI) based asymmetric power splitter-combiner from [221]. (b) Design from [111] with two cascaded OOK MR modulators.	45
4.3	Illustration of an electrical DAC (EDAC) enabled MR-based 4-PAM modulator from [208]. Inset: Illustration of resonance passbands and optical transmission levels for an EDAC-enabled MR-based 4-PAM modulator. .	46
4.4	Optical DAC (ODAC) enabled MR-based 4-PAM signaling modulator from [167].	47
4.5	Schematics of (a) a receiver module for an OOK modulation-based link [248, 252], (b) a receiver module for a 4-PAM modulation-based link [248, 252], (c) a serialization module [174], and (d) a deserialization module [174]. N_λ is the number of DWDM signals in the link. B is number of bits per symbol; B=1 for OOK signaling, and B=2 for 4-PAM signaling . . .	49

4.6	Schematic illustration of (a) an OOK modulation-based optical link, and (b) a 4-PAM modulation-based optical link, with total four wavelengths (λ_1 to λ_4). Note that having equal number of optical signals results in $2\times$ datarate for the 4-PAM link. In other words, equal datarate can be achieved for 4-PAM links by using $2\times$ less optical signals.	51
4.7	Schematics of (a) 8-ary 3-stage CLOS PNoC architecture [116] and (b) SWIFT PNoC architecture [56].	73
4.8	Packet latency plotted across PARSEC benchmark applications for (a) datarate-BER balanced variants, and (b) BER-optimal variants of CLOS PNoC. All results are normalized to the baseline CLOS_OOK.	76
4.9	Packet latency plotted across PARSEC benchmark applications for (a) datarate-BER balanced variants, and (b) BER-optimal variants of SWIFT PNoC. All results are normalized to the baseline SWIFT_OOK.	77
4.10	Average total power dissipation for different (a) datarate-BER balanced, and (b) BER-optimal variants of CLOS PNoC. The error bars represent the minimum and maximum values of power dissipation across 12 PARSEC benchmarks.	78
4.11	Average total power dissipation for different (a) datarate-BER balanced, and (b) BER-optimal variants of SWIFT PNoC. The error bars represent the minimum and maximum values of power dissipation across 12 PARSEC benchmarks.	79
4.12	Energy-per-bit (EPB) analysis for (a) the datarate-BER balanced variants, and (b) the BER-optimal variants of the CLOS PNoC. Column heights represent EPB averaged across 100 PV maps and normalized to the CLOS_OOK variant.	80
4.13	Energy-per-bit (EPB) analysis for (a) the datarate-BER balanced variants, and (b) the BER-optimal variants of the SWIFT PNoC. Column heights represent EPB averaged across 100 PV maps and normalized to the SWIFT_OOK variant.	81
5.1	Inter-Chiplet Silicon Photonic MRR-based DWDM Link.	86
5.2	(1) Schematic of an on-SiPhI inter-chiplet link, (2), a summary of the optical power budget, and (3) evolution of the optical power budget. The laser source is assumed to be a multi-wavelength comb source, and MAOP per waveguide and MAOP per wavelength are typically forced once the light from the laser source is coupled in the link through the coupler. . .	87
5.3	Illustration of FSR of an MRR.	94
5.4	Aggregate bandwidth and energy per bit (EPB) values for different waveguide lengths ranging from 1 cm to 10 cm obtained from the analysis performed on (a) Vanilla, (b) Minimized Loss, (c) Wide FSR, (d) Increased MAOP, (e) Minimized loss + Wide FSR, (f) Minimized loss + Increased MAOP, (g) Wide FSR + Increased MAOP, and (h) Minimized Loss + Increased MAOP + Wide FSR on-SiPh variants derived from 45nm SOI CMOS [236], 32nm SOI CMOS [236] and deposited poly-Si [12] platforms.	99

5.5	(a) Breakdown of wall-plug laser power and thermal power, and (b) the number of wavelength channels supported by different single-waveguide links of 2 cm length corresponding to various design pathways and fabrication platforms	102
5.6	Chiplet based Design of NUPLet.	105
5.7	GPU based multi-chiplet module (MCM).	105
5.8	Performance comparison of on-SiPhI link variants as implemented on NUPLet architecture. These variants are based on the 45nm SOI CMOS, 32nm SOI CMOS, and deposited poly-si photonic platforms.	109
5.9	Energy comparison of on-SiPhI link variants as implemented on NUPLet architecture. These variants are based on the 45nm SOI CMOS, 32nm SOI CMOS, and deposited poly-si photonic platforms.	110
5.10	Energy-delay product comparison of on-SiPhI link variants as implemented on NUPLet architecture. These variants are based on the 45nm SOI CMOS, 32nm SOI CMOS, and deposited poly-si photonic platforms.	111
5.11	Impact of aggregate bandwidth on training time	112
6.1	Structure and cross-section of our MRR based polymorphic E-O logic gate (MRR-PEOLG).	120
6.2	(a) Step-1 to Step-4, and (b) Step-5 of the procedure used for modeling our MRR-PEOLG. The schematic simulation setup in ANSYS/Lumerical's INTERCONNECT tool for (c) frequency-domain and (d) time-domain transient analysis of our MRR-PEOLG.	123
6.3	The transmission spectra obtained at the drop port of our MRR-PEOLG for logic-gate functions (a) AND, (b) OR, and (c) XOR, and at the through port of our MRR-PEOLG for complementary logic-gate functions (d) NAND, (e) NOR, and (f) XNOR.	124
6.4	(a),(b) The electrical pulse signals of 10 Gb/s bit-rate provided as input to the PN junctions of our MRR-PEOLG. The corresponding output pulse patterns obtained at the drop port of our MRR-PEOLG for logic-gate functions (c) AND, (e) OR, and (g) XOR , and at the through port of our MRR-PEOLG for complementary logic-gate functions (d) NAND, (f) NOR, and (h) XNOR. The optical input power is 5 dBm in all cases.	126
6.5	Colormap plots for logic functions (a) AND, (c) OR, (e) XOR (obtained at the drop port of our MRR-PEOLG), and complementary logic functions (b) NAND, (d) NOR, (f) XNOR (obtained at the through port of our MRR-PEOLG) that depict the maximum achievable bit-rate for given input optical power and SOMA. These color maps are evaluated for drop-port FWHM of 1.2 nm. We also report the maximum achievable bit-rate corresponding to (g) AND, OR, and XOR functions, and (h) NAND, NOR, and XNOR functions, evaluated for different values of FWHM, 0 dBm input optical power, and -5 dBm SOMA.	128
7.1	Illustration of common MRR based analog optical TPC organizations. (a) AMW TPC and (b) MAW TPC.	136

7.2	Mapping of the input and weight matrices onto the AMW and MAW TPCs.	137
7.3	(a) Device-level schematic of our microring modulator (MRM) based hybrid time-amplitude analog optical modulator (TAOM) integrated with a balanced photo charge accumulator (BPCA) and (b) analog representation of signals (optical and electrical) at different stages of our integrated TAOM+BPCA unit.	138
7.4	Circuit-level schematic of our integrated TAOM-BPCA unit which consists of two cascaded TAOMs, connected to our BPCA circuit. The inset showcases analog representations of signals (both optical and electrical) at various stages of our circuit.	142
7.5	Schematic of the MRM of the TAOM, assembled in the INTERCONNECT solver of ANSYS/Lumerical suite [155].	145
7.6	(a) Frequency domain and (b) time-domain simulation setup of our TAOM in the INTERCONNECT solver of the ANSYS/Lumerical suite [155]. . .	147
7.7	Transmission spectra obtained at the drop port of a 3-bit TAOM for different amplitudes.	147
7.8	Colormap plots that depict the (a) accuracy and (b) precision of our TAOM for different values of input optical pulse amplitude, bit resolution, and pulse widths for a unity input value.	149
7.9	(a) WDM-based MRMs and (b) aggregate transmission spectra of the WDM-based MRMs when using a coarse channel spacing (left) and a dense channel spacing (right).	150
7.10	(a) Cascaded TAOMs that enable WDM, and (b) its time-domain (transient) simulation setup in the INTERCONNECT solver of ANSYS/Lumerical suite [155].	151
7.11	IM crosstalk induced deviation in the multiplication results of each TAOM in a 2-channel parallel multiplier circuit, depicted for various channel spacings. (1.6nm, 1.2nm and 0.7nm)	153
7.12	Architecture of the hybrid Time-Amplitude Analog Optical Modulator based Tensor Core (TAOM-TC).	155
7.13	Supported TPC size $N(=M)$ for bit-precision = 1,2,3,4-bits at data rates (DRs)=1,5,10 GS/s, for TAOM-TC, MAW TPC and AMW TPC.	157
7.14	Power consumption of the AMW, MAW and TAOM-TC architectures for a bit precision of 4-bits and a data rate of 5 GS/S.	159
7.15	Runtimes of state-of-the-art GPUs, AMW, MAW and TAOM-TC accelerator architectures estimated for four different benchmarks.	161
8.1	(a) Top view, (b) Cross-sectional view (along AA') of our SiN-on-SiO ₂ MRR modulator with ITO-SiN-ITO stack as active upper cladding. . . .	166
8.2	(a) Top view, (b) Cross-sectional view (along AA') of our SiN-on-SiO ₂ MRR modulator with ITO-SiO ₂ -ITO stack as active upper cladding. . .	166
8.3	Transmission spectra of our modulator with ITO-SiN-ITO stack.	168
8.4	Transmission spectra of our modulator with ITO-SiO ₂ -ITO stack.	168
8.5	Optical eye diagrams for (a) 30 Gb/s and (b) 55 Gb/s OOK inputs to our SiN-on-SiO ₂ (ITO-SiN-ITO stack) modulator.	168

8.6	Cross-sectional electric-field profiles of the fundamental TE mode evaluated at the coupling section (along BB' in Fig. 8.2(a)) ((a)-(c)), across the rim (along AA' in Fig. 8.2) ((d)-(f)), and at the through port of our SiN-on-SiO ₂ (ITO-SiO ₂ -ITO stack) MRR modulator ((g)-(i)), for three different free-carrier concentrations of ITO (Table 8.2) namely 1×10^{19} cm ⁻³ (for (a),(d),(g)), 9×10^{19} cm ⁻³ (for (b),(e),(f)), and 17×10^{19} cm ⁻³ (for (c),(f),(i)), using the variational FDTD (varFDTD) solver [155].	170
8.7	Optical eye diagrams for (a) 30 Gb/s and (b) 55 Gb/s OOK inputs to our SiN-on-SiO ₂ (ITO-SiO ₂ -ITO stack) modulator.	170
8.8	Modulation bandwidth, modulation efficiency and FSR (shown as the size of the bubbles and red data labels) of various Si, LN (LiNbO ₃) and SiN MRR modulators from prior work, compared with our SiN-on-SiO ₂ (ITO-SiN-ITO Stack) MRR modulator.	172
9.1	Illustration of the MWA organization of an SOI-based TPC.	177
9.2	Schematic of a TPC of our SiNPhAR GEMM Accelerator.	179
9.3	Transmission spectra measured at the through port of our SiN-on-SiO ₂ MRM weighting element for various transmission amplitudes. These different transmission amplitudes at λ_T signify the weighting of the input optical amplitude symbol.	181
9.4	Supported TPC size N(=M) for bit precision = {1,2,3,4}-bits at data rates (DRs)={1,5,10} GS/s, for SOI-MWA TPC and SiNPhAR TPC.	183
9.5	System-level implementation of SiNPhAR accelerator. DPU=TPC.	185
9.6	(a) Normalized FPS (log scale) (b) Normalized FPS/W (log scale) for SiNPhAR versus SOIPhAR accelerators with input batch size=1. Results of FPS and FPS/W are normalized w.r.t. SOIPhAR ResNet50 at 10 GS/s.	187
10.1	Schematics of how the cascaded arrays of our MRR-PEOLG can be reconfigured to implement (a) a SIMD or (b) an MIMD E-O processing unit. The reconfiguration between SIMD/MIMD can be achieved by programming the individual MRR-PEOLGs for specific logic/arithmetic functions.	194

Chapter 1 Introduction

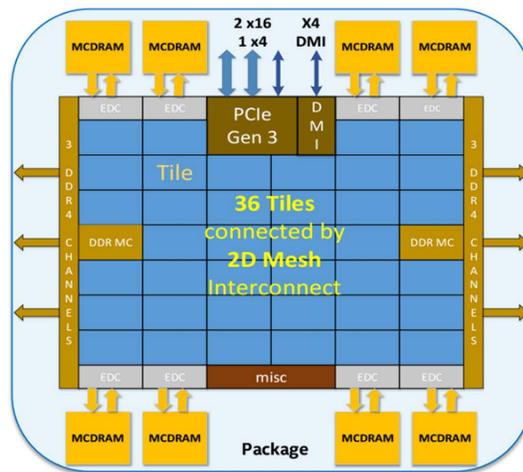
1.1 Overview

In recent decades, the computing landscape has been shaped by Moore’s law and Von Neumann architectures, with a focus on addressing computational demands through increased core integration. Von Neumann architectures necessitate the exchange of data between memory and processor cores, facilitated by electrical interconnects. However, the escalating performance requirements of contemporary data-centric applications (e.g., artificial intelligence, machine learning, big data, and internet-of-things (IoT)), including demands for higher throughput, lower latency, and enhanced energy efficiency, have rendered conventional electrical interconnects inadequate. These issues stem from challenges such as limited bandwidth, increased power consumption, and growing complexity in ensuring reliable communication. Moreover, as the pursuit of higher performance continues, the need for increased computational power also intensifies. Historically, the evolution of Moore’s law enabled us to address escalating computing demands by integrating more cores onto a single chip. To efficiently connect numerous processing cores in a scalable manner and facilitate global on-chip communication to meet performance expectations, electrical networks-on-chip (ENoCs) were introduced. ENoCs are structured and scalable communication architectures that apply methods of computer networking to on-chip communication and bring notable improvements over traditional bus and crossbar communication architectures. However, in recent years, Moore’s law has been facing fatal challenges as nanofabrication technology is experiencing physical limitations due to the exceedingly small size of transistors. Additionally, with the technology scaling, the spacing between adjacent electrical interconnects keeps shrinking which leads to an increase in the amount of coupling capacitance between the interconnects resulting in crosstalk noise that will further deteriorate the reliability of ENoCs. This challenge is poised to intensify as the demand for larger die sizes, more cores, and additional subsystems continues to surge. This has prompted both industry and academia to explore new ‘more-than-Moore’ technologies that can address the mounting demands of on-chip communication and computing. Therefore, this chapter delves into how silicon photonic interconnects and silicon photonic-based E-O computing circuits can replace their electronic counterparts, highlighting their unique advantages. Additionally, this chapter discusses various design challenges associated with silicon photonic interconnects and silicon photonic-based E-O circuits, and outlines variety of solutions proposed to overcome these challenges. The ultimate goal is to pave the way for scalable solutions that enhance both communication and computing capabilities in the era of advanced data-centric applications.

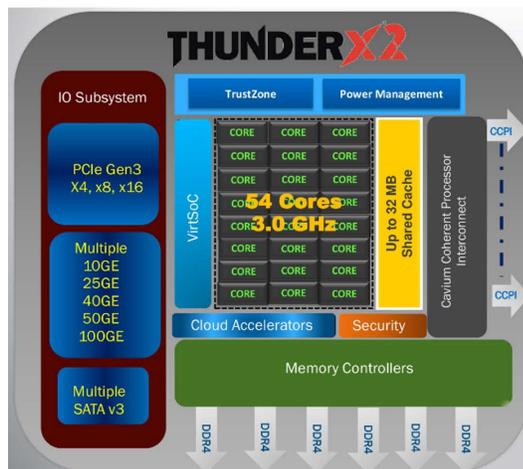
1.2 On-Chip Communication With Silicon Photonics

Background and Motivation

The evolution of mainstream computing systems has moved from the multicore to the manycore era. Manycore processors are specialist multi-core processors designed for a high degree of parallel processing, containing a large number of independent processor cores ranging from few tens to thousands of cores provided on a single chip, packaged with up to hundreds of gigabytes of memory at high bandwidth [102, 64]. Manycore processors are used extensively in embedded and high-performance computing applications. Examples include Intel Xeon Phi (up to 72 cores) [165], ARM ThunderX2 (54 cores) [160], Qualcomm Centriq 2400 (48 cores) and GPUs (100s of cores) [51].



(a)



(b)

Figure 1.1: (a) Intel Xeon Phi – 72 core processor [165], (b) ARM ThunderX2 – 54 core processor [160].

Fig. 1.1 represent Intel Xeon Phi and ARM ThunderX2 manycore processors that are used for high-performance computing applications. These manycore chips employ many processing elements that are grouped in multiple compute clusters. These compute clusters communicate with each other and with one or more on-chip memory controllers via an electrical network-on-chip (ENoC). The memory controllers connect with the employed memory modules via electrical core-to-memory interfaces. Efficient designs of interconnection fabrics (e.g., NoCs and core-to-memory interfaces) are essential to satisfy the bandwidth and latency constraints of advanced computing systems that utilize these manycore processors. It is therefore becoming evident that focus on interconnection architecture (e.g., NoCs and core-to-memory interfaces) design, customization, and exploration can provide huge performance gains in manycore processors and in advanced computing systems that utilize them [248].

To meet the growing demands of modern data-centric applications related to machine learning, artificial intelligence, big data and internet-of-things (IOT), number of cores in manycore processors keep increasing. This increase in core count in manycore systems demands for higher bandwidth and energy-efficient interconnection networks. Contrarily, traditional electrical connects ([165, 51, 160]) experience high power dissipation and reduction in performance with increase in number of cores. Electrical interconnects are also prone to crosstalk and electromagnetic interference with technology scaling, which will further dwindle the performance and reliability of electrical interconnects. This motivates the need for a new interconnect technology that can be leveraged to realize high bandwidth and energy-efficient interconnects for future manycore systems.

Silicon Photonic Interconnects

Recent developments in CMOS-photonics integration [239] have enabled an exciting solution for on-chip interconnects in the form of photonic networks-on-chip (PNoCs). Fig. 1.2 shows the physical-layer layout of several PNoC architectures namely CLOS PNoC [116], SwiftNoC [56] and LumiNoC [141].

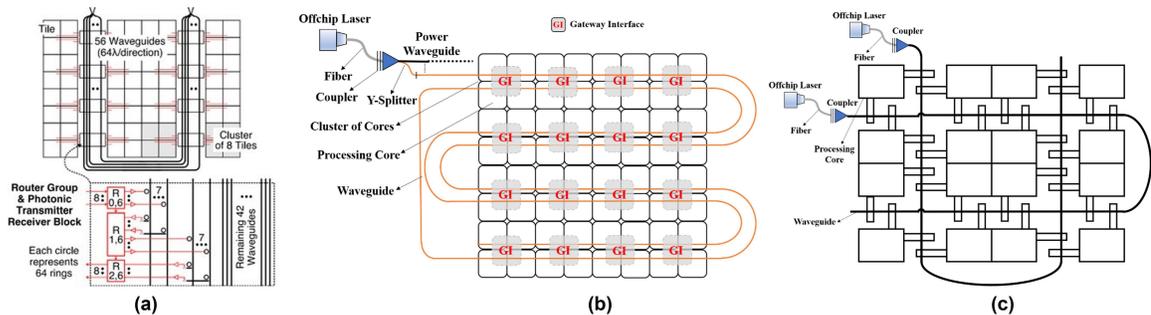


Figure 1.2: Physical-layer layout of (a) CLOS PNoC [116], (b) SwiftNoC [56] and (c) LumiNoC [141]

Typically, a PNoC comprises of multiple photonic links. Each photonic link in PNoC consists of one or more photonic waveguides spanning the PNoC chip depending on the variant of physical-layer architecture. For example, if we consider

SwiftNoC in Fig. 1.2(b), every single waveguide photonic link in this PNoC connects multiple gateway interfaces (GIs) with one another. Each GI connects to multiple parallelly laid-out photonic links, interfaces a group of four processing cores which is known as cluster with the links. Out of all the GIs that are connected to a single link, some GIs can write photonic data into the link and other can read photonic data from the link, to enable multiple-write-multiple-read (MWMR) type of crossbar configuration in the PNoC [188]. An off-chip laser source generates multiple wavelengths of light which is coupled into the PNoC via a power waveguide and a power splitter. This input multiple-wavelengths of optical power traverses individual links to all individual GIs on the chip. Each GI receives multiple wavelengths of input light as DWDM carriers for data signals, to enable communication with one or more other GIs. GIs enable inter-cluster communication by converting the received data packet from the source processing core into parallel electrical data signals, which are then modulated onto DWDM carriers converting them into parallel photonic data signals. These DWDM data signals traverse a single waveguide photonic link to a receiver GI. The receiver GI converts the incoming photonic data signals to electrical data signals and consequently, the electrical data packet which is then passed onto the destination core. Detailed discussion on structure of a photonic link, several building blocks of a photonic link and how electrical to optical (E/O) conversion at the transmitter side and optical to electrical conversion (O/E) at the receiver side is done in the next subsection.

Photonic Links

A photonic link (Fig. 1.3) comprises of one or more photonic silicon-on-insulator (SOI) waveguides with dense wavelength division multiplexing (DWDM) of multiple wavelengths into each waveguide. In a DWDM-enabled SOI waveguide, SOI microring resonator modulators (MRMs), which are arrayed along the waveguide at the source end, modulate input electric signals onto parallel photonic channels. The photonic channels travel through the waveguide and reach the destination end, where an array of SOI microring resonators (MRRs) drop the parallel photonic signals onto the adjacent photodetectors to recover the electric data signals. At the transmitter side of the photonic link, each MRM employs a serialization module and a driver circuit that can produce a sequence of signal bias voltages corresponding to the input sequence of electrical bits. The converted optical data packets are transmitted over different wavelength channels at a higher bit rate compared to the electrical data packets. Therefore, Serialization modules are used to enable the conversion between the data rates and they are implemented using parallel-in serial-out electronic buffers. Similarly, at the receiver side of the photonic link, each detector MRR employs a deserialization module (implemented using serial-in parallel out electronic buffers) and a trans-impedance amplifier (TIA) which amplifies the output signals from the photodetector to digital voltage levels.

Photonic interconnects are less susceptible to crosstalk [161] and have lower dynamic power dissipation compared to electrical interconnects. These advantages of silicon photonic interconnects make them a very promising alternative to overcome

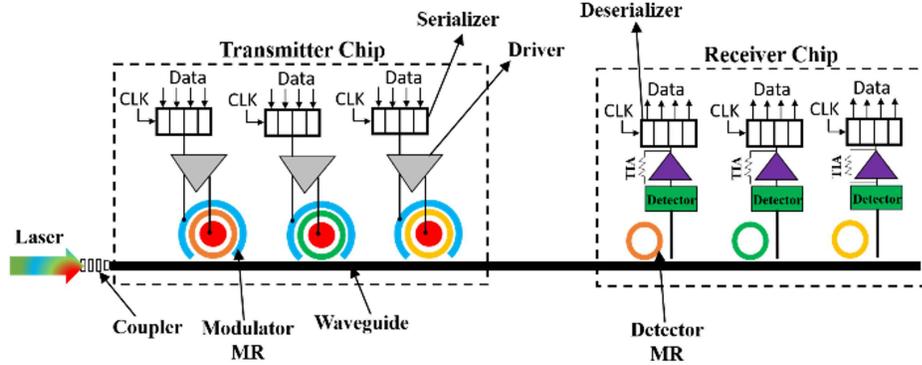


Figure 1.3: An on-chip photonic link.

the bottlenecks of electrical interconnects. Several photonic devices such as waveguides, MRRs and photodetectors are discussed in the upcoming subsections. All of these devices have been successfully fabricated and demonstrated at chip level ([279, 109, 52]) and several PNoC architectures ([189, 188, 260, 116, 56, 141]) have been designed using these devices.

Photonic Waveguide: There are several transmission mediums available which can be used to carry photons between transmitter and receiver in an on-chip photonic interconnect. The most predominantly used transmission medium in an on-chip photonic interconnect is a silicon-on-insulator (SOI) waveguide with high refractive index silicon (Si) core ($n_{si} = 3.5$) and low refractive index silicon-di-oxide (SiO_2) cladding ($n_{\text{SiO}_2} = 1.5$) [177]. There are four different configurations of SOI waveguides namely channel, ridge, slot and photonic-crystal waveguides. Channel and ridge waveguides are the most common and rely on total internal reflection (Fig. 1.4) which concentrates the light in the high index of refraction material [177]. These waveguide configurations provide single-mode propagation. Waveguides fabricated on SOI platform have advantages such as lower losses and compact footprint which requires lower drive voltage for high frequency operations.

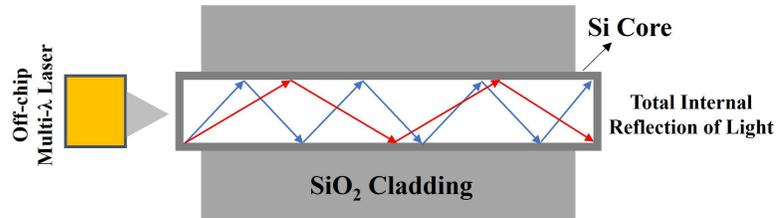


Figure 1.4: Illustration of the total internal reflection of light in the longitudinal cross-section of a silicon-photonic waveguide.

Microring Resonators: Microring resonators (MRRs) are wavelength division multiplexing (WDM) compatible devices that are compact and energy efficient and are designed to resonate when presented with specific individual wavelengths and remain quiescent at all other times [177]. Every individual MRR is capable of mod-

ulating a single wavelength. Therefore, a transmitter with a multi-bit, parallel data path can be constructed using multiple MRRs and a WDM-capable light source. The particular wavelength that a MRR resonates to is dependent upon MRR radius (R) and effective refractive index (n_{eff}). The dependence of MRR's resonance wavelength on R and n_{eff} is given by the following equation:

$$\lambda_{res} = \frac{L \times n_{eff}}{m} \quad (1.1)$$

Where λ_{res} is the resonance wavelength of the MRR, L is the round trip length of the MRR given by $2 \times \pi \times R$, n_{eff} is the effective refractive index of the MRR and m is an integer. By changing R and n_{eff} , the resonance wavelength of MRR can be altered. n_{eff} can be changed in two ways. Injection or removal of carriers from Si core of an MRR changes the n_{eff} of an MRR due to electro-optic effect [7]. Carrier tuning of an MRM is shown in Fig. 1.5(a). Heating of MRRs also alters its n_{eff} due to thermo-optic effect [176]. An MRR with an integrated heater is shown in Fig. 1.5(b). Spectral shift in the passband of an MRR due to localized heating and carrier tuning is shown in Fig. 1.5(c).

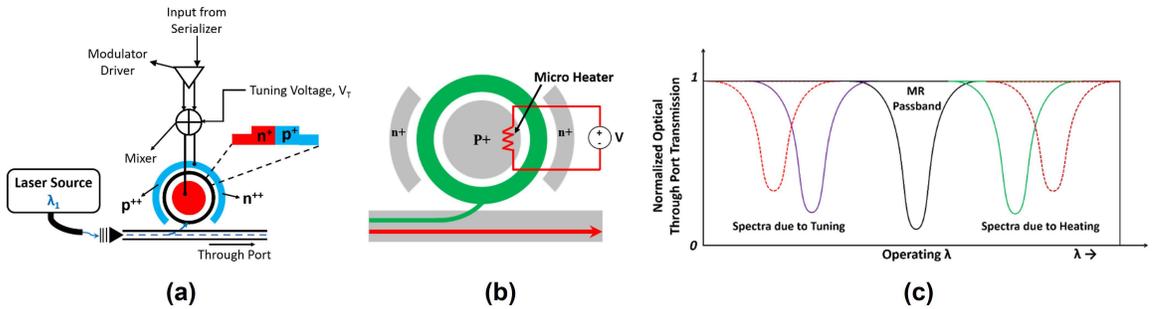


Figure 1.5: (a) A tunable MRR, (b) MRR with an integrated micro heater and (c) Shift in resonance of the MRR due to heating (right side of the MRR passband) and carrier tuning (left side of the MRR passband)

Carrier injection/removal is the most predominantly used to switch MRs between active (MRR's wavelength in resonance with wavelength from data waveguide) and passive modes (MRR's wavelength not in resonance with wavelength from data waveguide) since it is faster and consumes less power compared to heating method. In order to implement this method, MRs require a series of driver circuits which regulate carrier injection/removal rates into MRs by altering voltage to control their resonance wavelength shifts. A tunable MRR can be used for various applications as shown in Fig. 1.6. A tunable MRR can be used as a modulator for data communication (e.g., OOK modulator). In addition, a tunable MRR can also be used as a switch to route wavelengths from one waveguide to the other. It can also be used as a wavelength filter on the receiver side of a photonic link to filter a wavelength and route it towards a photodetector.

Photodetectors: MRRs at the transmitter side extract the individual wavelengths from the waveguide and direct the extracted wavelengths to photodetectors,

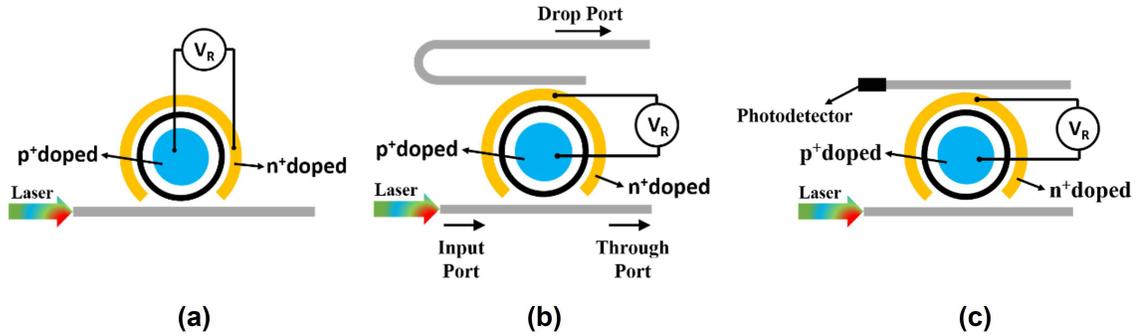


Figure 1.6: A tunable MRR as (a) an OOK modulator, (b) Switch and (c) Wavelength filter

which convert photonic power into an electrical signal (Fig. 1.7). The output voltages from photodetectors are amplified to digital voltage levels using transimpedance amplifiers (TIA) [248]. As the signals amplified by TIAs are ultimately stored and processed on the chip, their amplitudes should match the supply voltage of logic circuits. To enable amplification of signals to supply voltage, the TIAs are typically operated at 20% higher supply voltage than supply voltage of logic circuits [248].

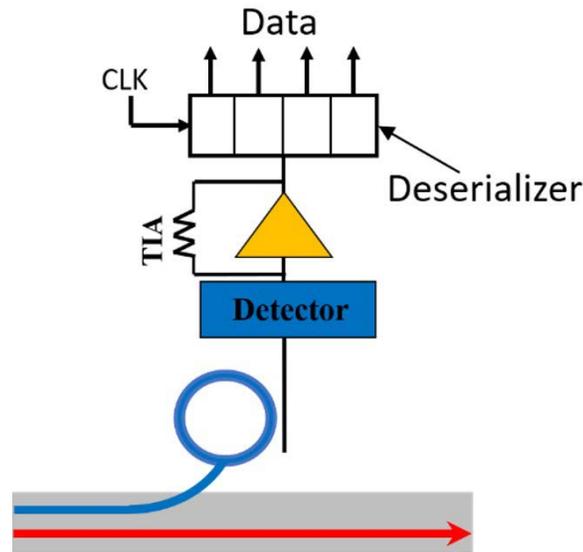


Figure 1.7: MRR filter at the receiver side to detect its resonance wavelength

Design Challenges of Silicon Photonic Interconnects

Design challenges of silicon photonic interconnects are organized into the following three categories: reliability challenges, challenges due to limited optical power budget, and challenges due to high static power dissipation. Each of these categories of challenges is described below.

Reliability Challenges

Reliability challenges in photonic interconnects include challenges due to the adverse effects of crosstalk noise, fabrication-process and on-chip temperature variations and limited free spectral range (FSR) of MRRs. Each of these challenges is elaborated below.

Crosstalk Noise: Crosstalk noise in photonic interconnects can be classified into two categories namely inter-channel crosstalk and intra-channel crosstalk. For inter-channel crosstalk, the signal power of same wavelength channel gets affected by the noise power from one or more neighboring wavelength channels (Fig. 1.8(a)) whereas for intra-channel crosstalk, the signal power of a particular wavelength channel is affected by the noise power of same wavelength channel (Fig. 1.8(b)) [250].

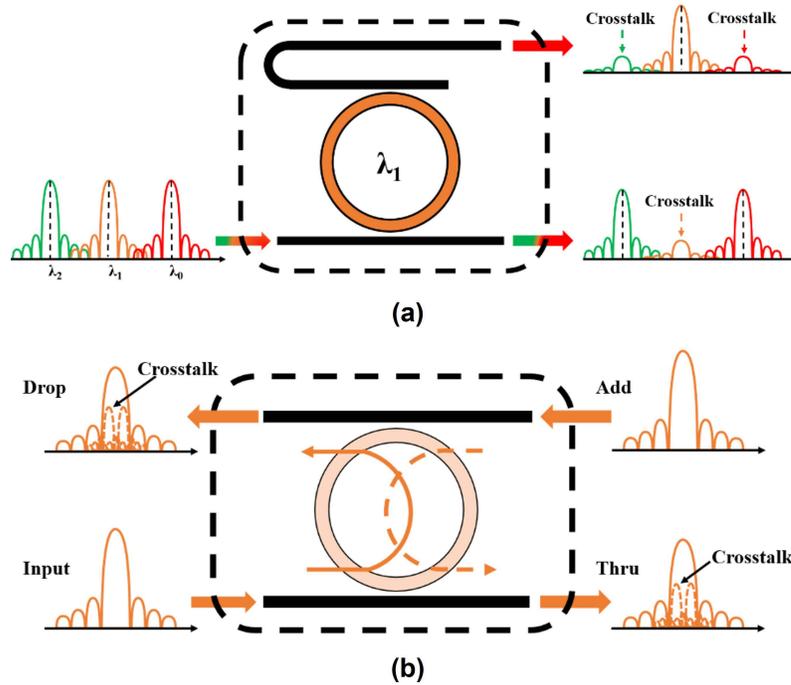


Figure 1.8: Illustration of (a) Inter-channel crosstalk and (b) Intra-channel crosstalk (reproduced from [18])

The strength of inter-channel crosstalk depends on several aspects namely quality factor of the MRRs, data rate and channel gap between the resonant wavelength of an MRR and its adjacent wavelengths. High crosstalk in photonic links degrade the optical signal-to-noise-ratio (OSNR) and the target bit-error rate (BER). In order to compensate for the effects of inter-channel crosstalk and to ensure that target BER remains unharmed, extra optical power is added to each wavelength channel at the transmitter side and the receiver side which is known as power penalty [124]. At the transmitter side, due to carrier depletion or injection of the p-n junction of MRMs, the ring spectra switches between two resonance frequencies f_1 and f_0 as shown in Fig. 1.9(a). The difference between the resonances f_1 and f_0 is given by Δf . If Δf

is too large, then the shifted spectrum captures some of power of the neighboring channel that passes by the MRM resulting in inter-modulation crosstalk.

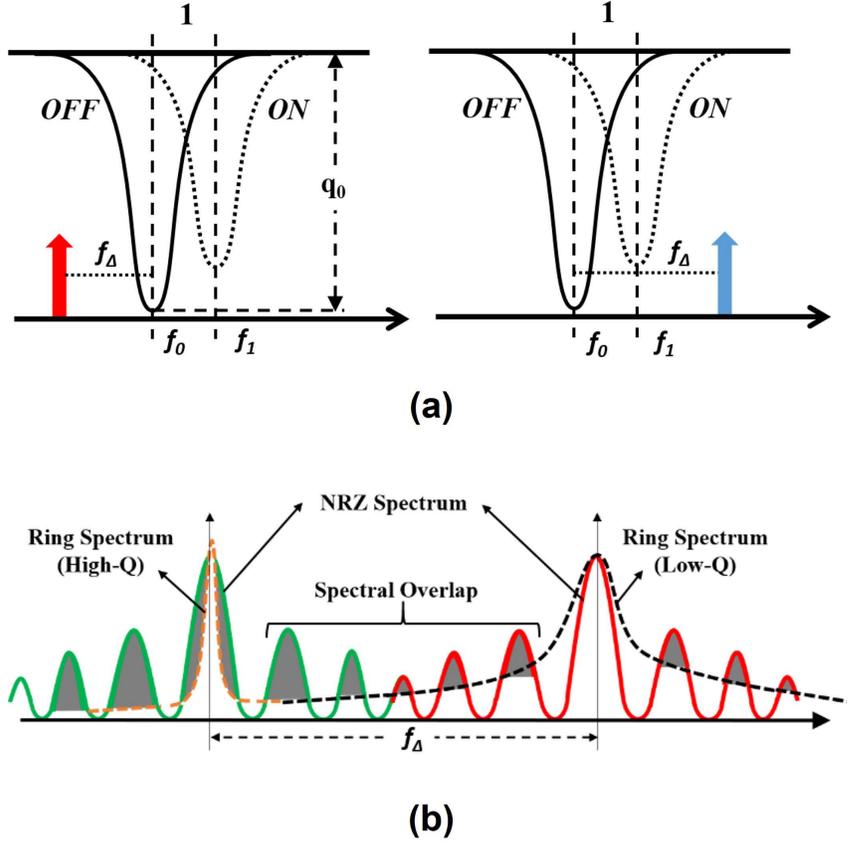


Figure 1.9: (a) Inter-modulation crosstalk and (b) Spectral view of spectral distortion and inter-channel crosstalk at the detector side

At the receive side of photonic link, power penalty is the sum of spectral distortion penalty, inter-channel crosstalk penalty and ring drop intrinsic loss (IL) [19]. For low values of FSR, the channel spacing between the adjacent wavelengths is low. Due to low channel spacing, each MRR at the receiver side drops its corresponding resonance wavelength channel and also collects some residual power from neighboring channels as shown in Fig. 1.9(b). This residual power is referred to as inter-channel crosstalk. In addition, MRRs with low quality factor have a wide resonant pass band which overlaps with signal spectra of neighboring channels resulting in crosstalk effect as shown in Fig. 1.9(b). MRR filters with high quality factor have narrow pass band spectrum which results in spectral distortion effect (Fig. 1.9(b)).

Process and Thermal Variations : Fabrication process variations (PV) are randomly occurring variations in the critical dimensions of photonic devices, such as width and thickness, when they are fabricated [217]. The PV-induced variations in width and thickness of MRRs cause drift in resonance wavelength of the MRRs . Fig. 1.10 illustrates the drift in resonance wavelength of an MRR due to PV. If the

resonance passband of an MRR drifts towards left end of the spectrum (i.e., decrease in wavelength), it is known as blue shift whereas if the resonance passband drifts towards the right end of the spectrum (i.e., increase in wavelength), it is known as red shift. This shift in resonance wavelengths of MRRs increase the crosstalk noise power and decrease the signal power, deteriorating the OSNR and BER in a waveguide. In order to counter the PV-induced resonance shifts, localized trimming [176] or thermal tuning [7] has been introduced. To counteract the PV-induced shift in MRRs resonance wavelength, localized trimming mechanism introduces free carriers to reduce the refractive index of the MRR. However, the extra free carriers increase the absorption loss in the MRR due to free carrier absorption (FCA) [32]. This increase in absorption loss reduces the quality factor of MRRs and increases the insertion loss and crosstalk noise. We model the PV induced variations in resonance wavelength, quality factor and extinction ratio of MRRs at wafer level using the spatial variation models from [268]. Fig. 1.11 illustrates the wafer level variation pattern in resonance wavelength, quality factor and extinction ratio of MRRs and their corresponding histograms. The original quality factor, extinction ratio and resonance wavelength of MRs is 6500, 10dB and 1550 nm respectively. But due to PV, quality factor varies between 2000 and 12000 which can be illustrated from the colormap scale shown in Fig. 1.11(b). Similarly, extinction ratio of MRRs vary between 4 dB and 18 dB (fig. 1.11(c)) whereas resonance wavelength of MRRs vary between 1540 nm and 1560 nm (Fig. 1.11(a)).

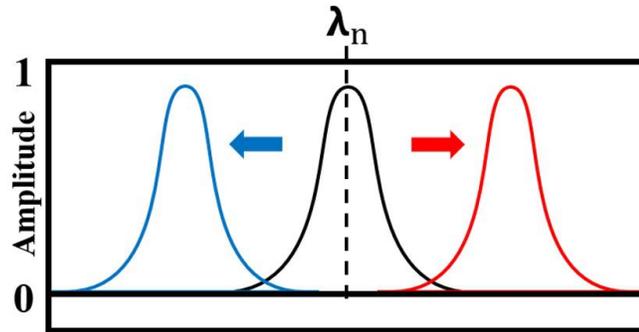


Figure 1.10: Illustration of drift in resonance passband of an MRR due to PV.

MRRs are highly susceptible to thermal variations (TV). With change in increase or decrease in temperature, the effective index of MRRs change. This change in effective index results in change in resonant wavelengths of MRRs [7]. Therefore, TV affect the reliability of the photonic link and also leads to the squandering of available bandwidth.

Limited Free Spectral Range (FSR) : MRR, which is considered as the workhorse of a photonic interconnect, is a looped waveguide in which the resonance occurs when the optical path length of the MRR is exactly a whole number of wavelengths. Therefore, MRRs support multiple resonances and the spacing between these resonances is free spectral range (FSR) (Fig. 1.12). Low values of FSR means for a given number of wavelength channels that are multiplexed in a waveguide, the

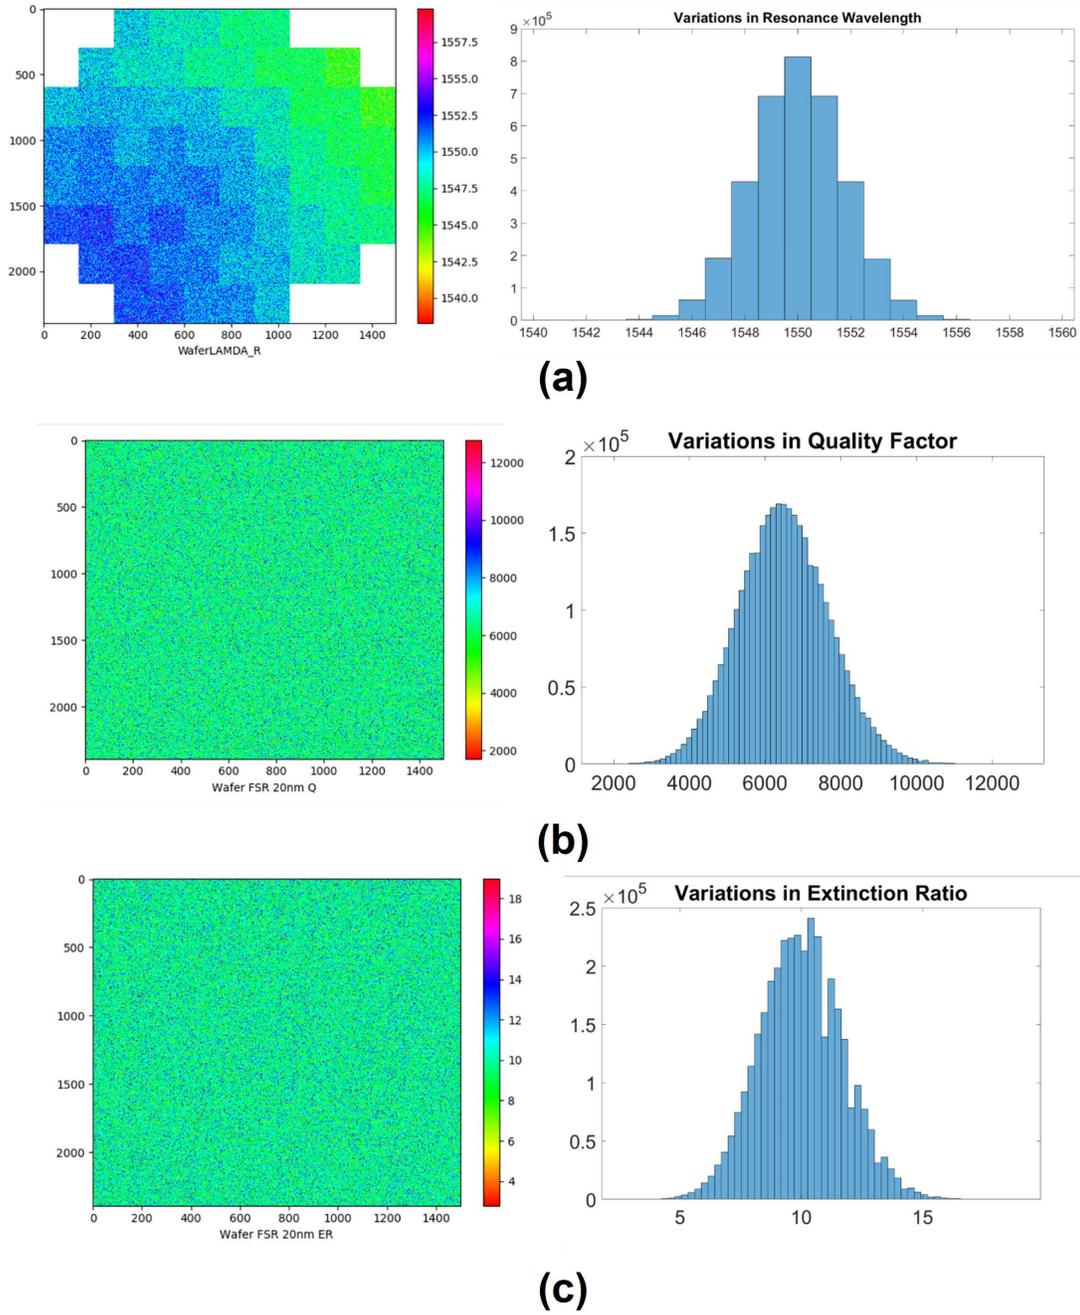


Figure 1.11: Wafer-level variation patterns (left) and corresponding histograms (right) of (a) Resonance wavelength, (b) Quality factor and (c) Extinction ratio of MRRs.

spacing between the adjacent channels is low, which in turn results in inter-channel crosstalk noise [18] worsening the optical signal-to-noise-ratio (OSNR) and bit-error rate (BER) in a waveguide.

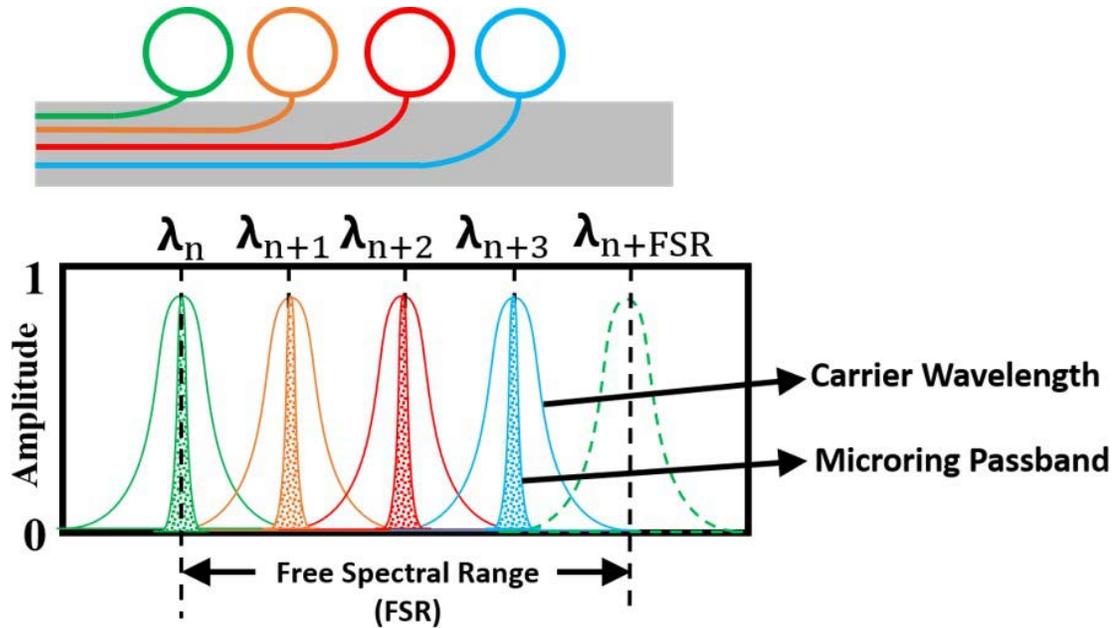


Figure 1.12: Illustration of FSR of an MRR.

Challenges Due To High Static Power Dissipation

Power dissipation challenges in photonic interconnects can be classified into two categories namely losses challenge and tuning power dissipation challenge. Each of these challenges are discussed below.

High Optical Power Due to High Optical Signal Losses : Photonic signals in photonic interconnects experience different types of losses (Fig. 1.13) namely propagation loss, bending loss, splitter and coupling loss and through loss. Typical values of various signal losses in a photonic link are provided in Table 1.1. Photonic signals propagating inside the waveguide experience propagation and bending losses. Propagation loss is the sum of absorption loss and scattering loss. Non-linear effects in Si such as TPA induce strong free carrier absorption (FCA) effect in silicon [147], significantly increasing the absorption losses in waveguides. Si waveguides are also prone to high scattering losses due to sidewall roughness of the waveguides since the refractive index contrast between Si core and SiO₂ cladding is high. In addition to these losses, splitters and couplers in photonic interconnects incur splitter and coupling loss whereas modulator and detector incur through losses. In order to ensure that detectors on the receiver end of photonic interconnect receive sufficient signal power, photonic signals demand high laser power. Therefore, high losses result in high laser power dissipation which indemnifies the energy benefits of photonic interconnects.

Tuning Power Dissipation Challenge : Trimming and tuning techniques are implemented in order to counteract PV and TV in MRRs. But if the number of MRRs or the degree of DWDM is high to support higher bandwidths, tuning power dissipation increases which in turn increases the overall power dissipation.

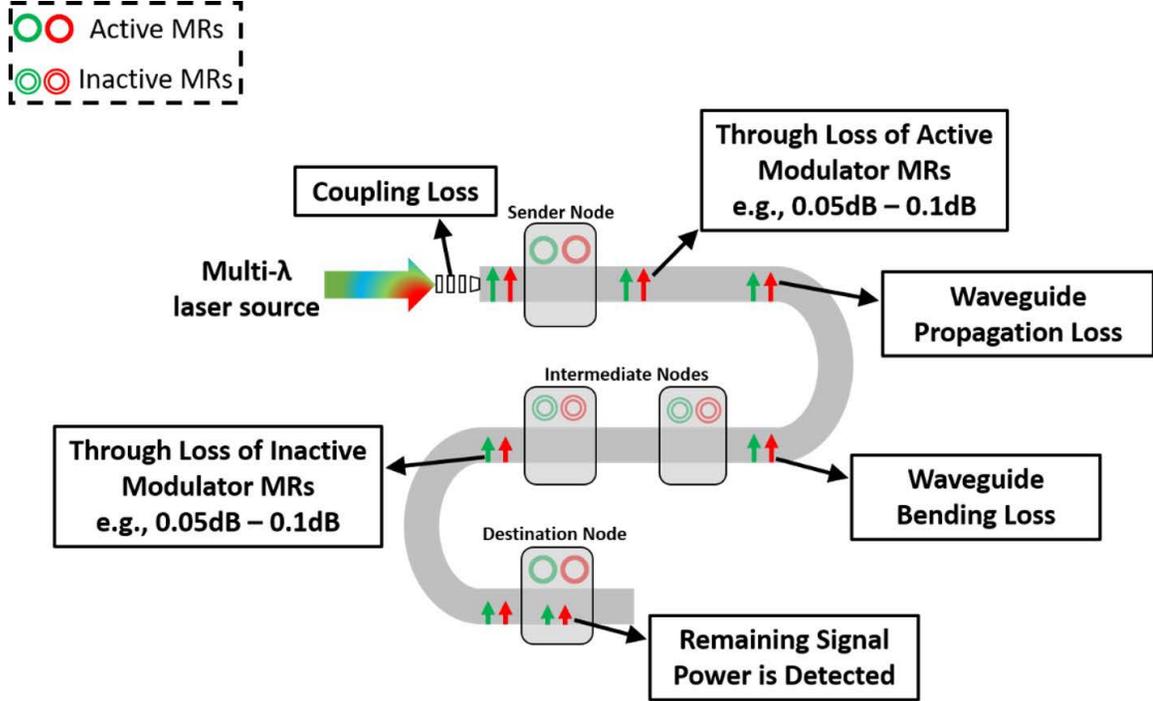


Figure 1.13: Various types of signal losses in a silicon photonic link

Table 1.1: Typical values of various losses in a photonic link [100, 154]

Parameter	Value
Waveguide Propagation Loss	1 dB/cm [154]
Waveguide Bending Loss	0.005 dB/900 [100]
Splitter Loss	0.5 dB [100]
Coupling Loss	2 dB [100]

Challenges Due To Limited Optical Power Budget

In order to optimize the design of a photonic link, optical power budget (OPB) per wavelength (λ) and OPB per waveguide (WG) are the most critical design constraints and number of DWDM wavelengths (N_λ) is the most important design parameter. OPB per λ determines the apex of signal losses and power penalties that can be allowed per λ channel propagating in the link and is calculated in dB as the difference between maximum allowable optical power (MAOP) per λ and detector sensitivity (S) (Eq. 1.2). MAOP per λ determines the ceiling of OPB per λ whereas S is the minimum amount of signal power that can be detected at the receiver side. Similarly, OPB per waveguide determines the maximum amount of allowable signal losses and power penalties in a photonic link. It is expressed as difference between MAOP per WG and S in which MAOP per waveguide determines the maximum signal power that can be fed into a waveguide (Eq. 1.3). In order to design a photonic link for a given value of N_λ , OPB per λ and OPB per WG should satisfy conditions given in Eq. 1.4 and Eq. 1.5. P_{dB}^{Loss} provided in Eqs. 1.4 and 1.5 accounts for total amount of

losses and power penalties in a photonic link.

$$OPB(\text{Per } \lambda)(dB) = MAOP(\text{Per } \lambda)(dBm) - S(dBm) \quad (1.2)$$

$$OPB(\text{Per } WG)(dB) = MAOP(\text{Per } WG)(dBm) - S(dBm) \quad (1.3)$$

$$OPB(\text{Per } \lambda)(dB) \geq P_{dB}^{\text{Loss}} \quad (1.4)$$

$$OPB(\text{Per } WG)(dB) \geq P_{dB}^{\text{Loss}} + 10 \log_{10}(N_\lambda) \quad (1.5)$$

As we can illustrate from Eq. 1.2-1.5, in order to design photonic links that can accommodate higher number of wavelengths, the amount of losses and power penalties per wavelength/waveguide should be low and MAOP per wavelength/waveguide should be high. However, Due to TPA induced effects in Si, MAOP in SOI photonic links is restricted to no more than 20dBm (MAOP per wavelength channel less than 6 dBm [147]). Also, TPA induces strong free carrier absorption (FCA) in silicon, increasing the absorption losses in Si waveguides. In addition, as discussed in previous subsections, SOI photonic links also experience high propagation losses and crosstalk penalties. Due to restrictions on MAOP per wavelength/waveguide and high losses and penalties in the link, OPB per waveguide does not accommodate higher number of wavelengths restricting the scalability of photonic links.

1.3 Computing With Silicon Photonics

Background and Motivation

Moore’s Law has been a guiding force in propelling the evolution of computing hardware since its inception. Over the past few decades, transistors have undergone a remarkable reduction in size, leading to the integration of billions of them on a single chip. This unprecedented level of transistor integration has paved the way for the design of sophisticated computational architectures, including Field Programmable Gate Arrays (FPGAs) and Graphical Processing Units (GPUs). For instance, Intel’s FPGA, illustrated in Fig. 1.14(a), boasts an impressive ~ 10.2 million logic cells [108], while Xilinx Vertex, featured in Fig. 1.14(b), incorporates approximately ~ 9 million logic cells [278]. Moreover, the trajectory of Moore’s Law has facilitated the creation of specialized hardware architectures for Artificial Intelligence (AI) acceleration. These architectures are designed to tackle the surging computational demands and inference times of Deep Neural Networks (DNNs). A prime example is the Cerebra’s WSE architecture [46], showcased in Fig. 1.14(c), which stands as the world’s largest computer chip. With a staggering ~ 2.6 trillion transistors, it not only holds this distinction but also reigns as the world’s fastest AI accelerator chip.

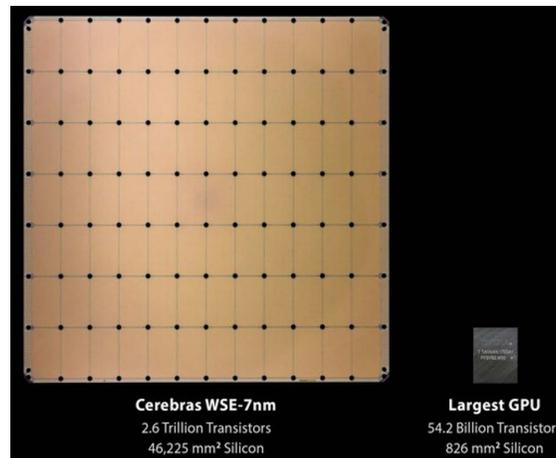
But unfortunately, in recent years, Moore’s law has faced fatal challenges as the nanofabrication technology is experiencing physical limitations due to the exceedingly small size of transistors. This has forced researchers in industry and academia



(a)



(b)



(c)

Figure 1.14: (a) Intel Stratix FPGA [108], (b) Xilinx Vertex FPGA [278], and (c) Cerebra’s Wafer Scale Engine (WSE) [46]

to develop new more-than-Moore technologies that can replace Moore’s Law and continue to provide persistently faster and efficient computing hardware for the future generations. Fortunately, silicon photonics (SiP) enabled electro-optic (E-O) circuit integration has been identified as one such promising technology. The E-O circuits built using the SiP technology are generally CMOS compatible and provide several advantages over their purely electrical counterparts. These advantages include sub-picosecond speeds, low power consumption and distance-independent bit-rate. Several prototypes of SiP-based E-O circuits for computing have been demonstrated in prior works [206, 222, 282, 81, 112, 285, 288]. A more in-depth discussion of the inception of silicon photonic-based E-O computing circuits and an elucidation of the associated design challenges are presented in the forthcoming subsections.

Silicon Photonic-Based Electro-Optic Computing Circuits

In recent years, there has been a remarkable surge in interest in Electro-Optic (E-O) computing systems, seamlessly integrating the advantages of both photonics and electronics. The allure of these hybrid systems lies in their ability to harness the precision of electronics alongside the speed of light [95]. E-O logic gates and circuits, in particular, offer additional benefits, operating with minimal latency due to their light-speed operation and achieving distance-independent, high bit rates [123][286], surpassing their electrical counterparts. Numerous prototypes of E-O logic gates/circuits have been reported in the literature, showcasing the versatility of this approach. The SiP-based E-O circuits for computing, which have been demonstrated in prior works (e.g., [297, 226, 225, 200, 287, 123, 286]) are typically used to implement the following four types of logical and arithmetic functions:

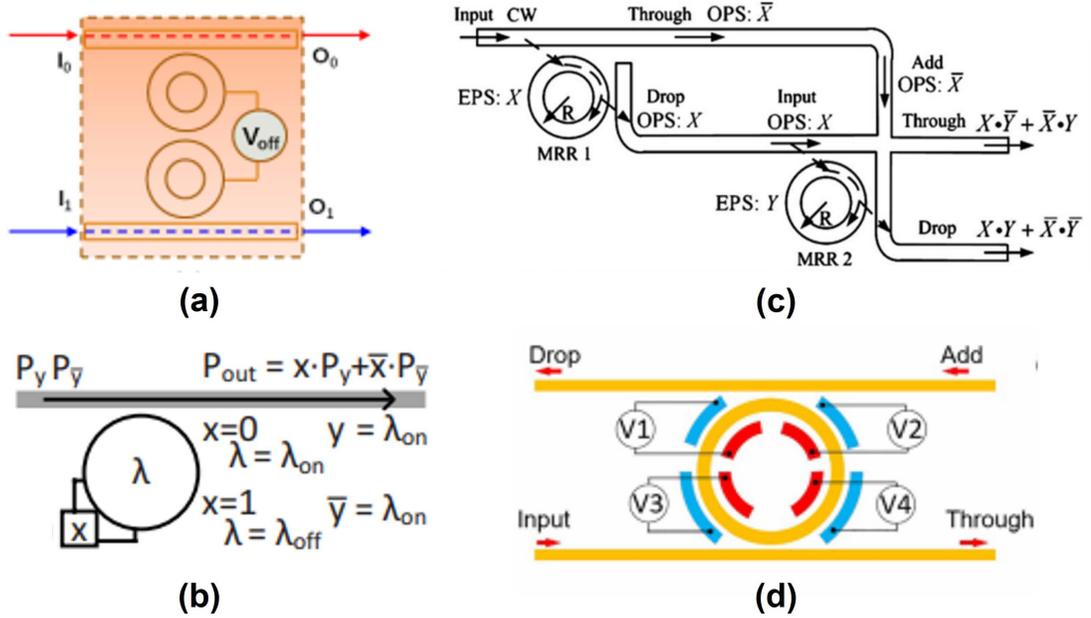


Figure 1.15: (a) An add-drop MR-based E-O AND logic gate [225], (b) a PCM-enabled MR-based E-O XNOR gate [297], (c) an MR-based E-O XOR/XNOR logic gate [291], and (d) an MR-based polymorphic E-O logic gate [286] demonstrated in the literature.

Basic Logic Gate Functions

A microring resonator (MRR) integrated with a phase change memory (PCM) device forms the basis of an XNOR gate, as depicted in Fig. 1.15(b), and is employed in [297] to accelerate binary neural networks. Similarly, in both [226] and [225], an add-drop MR-based AND gate (Fig. 1.15(a)) is utilized to enable partial multiplications of two binary operands, contributing to the acceleration of deep neural networks.

Combinational Logic Functions

Directed logic-based MRR-enabled reconfigurable E-O circuits are showcased in [200] and [287]. These circuits serve as a direct optical alternative to FPGAs.

Two-Operand Arithmetic Functions

High-speed Electro-Optic (E-O) circuits designed for partial sum accumulation and two-operand addition have been successfully demonstrated in prior work [123, 286]. These designs feature diverse configurations to support custom precision [123] and full-precision polymorphic operation [286] (Fig. 1.15(d)).

Multi-Operand Linear Arithmetic Functions

Various analog and digital Electro-Optic (E-O) circuits and architectures, utilizing MRRs and/or Mach-Zehnder Interferometers (MZIs), have been successfully demonstrated for performing operations such as Multiply-Accumulate (MAC) and General Matrix-Matrix Multiplication (GEMM) in the context of deep learning workloads. Prior works [297, 226, 152, 26] showcase the effectiveness of these E-O circuits and accelerator architectures in executing logical and arithmetic functions, meeting the demands of ultra-fast, highly parallel general-purpose computing, and accelerated deep learning applications. While MZIs introduce a significant area overhead, rendering MZI-enabled silicon photonic-GEMM accelerator architectures impractical for accelerating large-scale neural networks [244, 53], MRR-enabled accelerators exhibit disruptive performance and energy efficiencies. The compact footprint, low dynamic power consumption, and the ability of MRRs to support a large fan-in of optical signals through dense-wavelength-division multiplexing (DWDM) contribute to the exceptional performance of state-of-the-art MRR-enabled silicon photonic-GEMM accelerators, as showcased in literature [26, 263, 244].

Design Challenges of Silicon Photonic-Based Electro-Optic Computing Circuits

Despite their superiority over electronic counterparts, silicon photonic-based Electro-Optic (E-O) computing systems face three significant shortcomings, each of which is discussed below:

Lack of All-Electrical Application of Input Operands

The E-O circuits designed for simple logic-gate functions demonstrated in the literature [297, 226] often handle the two input operands differently. Typically, one operand is applied optically, while the other is applied electrically. To achieve this, one of the operands needs to be modulated onto the incoming optical wavelengths, requiring an additional optical modulator device per gate function, particularly when utilizing laser sources that provide unmodulated optical power. This necessity of providing one operand optically through an additional modulator device introduces

an increase in hardware area overhead and complicates operand handling within the E-O circuits.

High Idle Time

The E-O circuits for arithmetic functions demonstrated in the literature [226, 297] occupy up to $100\times$ more area compared to CMOS implementations. Furthermore, these E-O circuits for arithmetic functions struggle to achieve hardware utilization exceeding 60% [226]. This limitation arises from their typical integration within larger processing units, where they occupy only a fraction of the entire end-to-end datapath [26, 297, 152]. Such low hardware utilization often results in extended idle times, leading to elevated, non-amortizable area, and static power overheads.

Unsuitable to Implement Highly-Parallel Processing Architectures

The substantial area overhead of Electro-Optic (E-O) circuits renders them less suitable for the implementation of highly-parallel architectures, including Single-Instruction-Multiple-Data (SIMD), Multiple-Instruction-Multiple-Data (MIMD), and Systolic Array (SA) based processing architectures. This limitation arises because architectures such as SIMD, MIMD, and SA typically incorporate thousands of streaming processing units, each requiring multiple instances of basic logical and arithmetic functions. Implementing these functions with E-O circuits, which have an area footprint up to 100 times larger, significantly restricts the number of processing units that can be integrated onto a single chip. This constraint becomes particularly pertinent as the chip's area is typically limited by the reticle size ($\leq 900 \text{ mm}^2$ [172]).

1.4 Contributions

Sections 1.2 and 1.3 have delineated various design challenges associated with silicon photonic interconnects and silicon photonic-based E-O computing circuits. In this report, we put forth several solutions to address these challenges and make strides toward designing high-throughput, energy-efficient, and reliable photonic interconnects for future manycore computing systems, as well as scalable, reconfigurable, high-throughput, and energy-efficient E-O circuits for computing. The structure/outline of this report, highlighting our contributions, is organized as follows:

In Chapter 2, we introduce a novel design for an MRR filter array with a non-uniform quality factor distribution across individual MRR filters. This design aims to minimize crosstalk non-uniformity and achieve a uniform distribution of crosstalk penalty across channels. Uniformizing the crosstalk performance reduces overall laser power consumption in the photonic link.

In Chapter 3, we introduce a novel silicon-on-sapphire (SOS) based photonic interconnects, offering a potential solution to eliminate optical non-linearity-induced power constraints seen in conventional Silicon-On-Insulator (SOI) platforms. This innovation aims to overcome the scalability barriers and realize high-bandwidth, energy-efficient photonic interconnects for the future. We provide new compact models for

SOS devices and outline design principles for SOS-based photonic interconnects. Our chapter includes a link-level analysis assessing aggregate data rate and energy-per-bit for SOS-based photonic interconnects. Additionally, we conduct a system-level analysis on the CLOS PNoC [116], evaluating overall latency and energy-per-bit in the architecture.

In Chapter 4, we present a comparative study and a search heuristic-based method for designing DWDM-based on-chip photonic interconnects using various types of MR-based 4-PAM modulators. We conduct a comparison between different types of 4-PAM modulators and conventional OOK modulators at both link-level and system-level, considering aspects such as hardware overheads, performance, energy efficiency, and reliability. Employing a search heuristic-based method, we optimize the designs of DWDM-based photonic links using OOK and 4-PAM modulation methods. Additionally, we analyze how these optimized photonic interconnects impact the performance and energy efficiency of CLOS PNoC [116] and SWIFT PNoC architectures [56].

In Chapter 5, we explore various design pathways aimed at advancing on-silicon photonic interposer inter-chiplet interconnects to achieve multi-Terabits per second (Tb/s) aggregate bandwidth. Through an extensive link-level and system-level analysis, we investigate these design pathways both in isolation and in various combinations to assess their potential impact and effectiveness.

In Chapter 6, we present a novel design of a MRR-based polymorphic E-O logic gate (MRR-PEOLG) that can be dynamically programmed to implement different logic functions at different times. The objective of this design is to enhance the compactness and polymorphism of E-O circuits, ultimately improving operand handling and facilitating the amortization of area and static power overheads.

In Chapter 7, we present novel design of a hybrid Time-Amplitude Analog Optical Modulator (TAOM) and a balanced photo-charge accumulator (BPCA). A TAOM employs a single microring to perform a multiplication whereas a BPCA can perform a large number of spatio-temporal accumulations in situ. We arrange Multiple TAOMs and BPCAs in 2D to forge a tensor core and perform an extensive device-level, circuit-level, and system-level analyses to assess its advantages in comparison to prior works.

In Chapter 8, we present novel design of Indium Tin Oxide (ITO)-Based SiN-on-SiO₂ MRMs that can be utilized to design high-performance photonic integrated circuits for the future.

In Chapter 9, we present a novel Silicon Nitride (SiN)-Based Photonic GEMM Accelerator called SiNPhAR that employs SiN-based active and passive devices to implement analog GEMM functions. Through a cross-layer evaluation, we investigate its advantages over traditional Silicon-on-Insulator (SOI)-Based photonic GEMM accelerators, focusing on achievable spatial parallelism, throughput, and energy efficiency.

Chapter 10 concludes this report. We recap all our contributions and provide directions for future research.

Chapter 2 Mitigating Inter-Channel Crosstalk Non-Uniformity in Microring Filter Arrays of Wavelength-Multiplexed Photonic NoCs

Photonic networks-on-chip (PNoCs) employ photonic links with dense-wavelength-division-multiplexing (DWDM) of channels for parallel signal traversal, along with arrayed microring resonator (MR) filters for parallel signal reception, to enable high-bandwidth on-chip data transfers. Unfortunately, DWDM induces non-uniform inter-channel crosstalk in an MRR filter array, which degrades the communication reliability in the link. Overcoming this reliability degradation requires non-uniformly distributed signal power across the utilized data-channels in the link. This increases the total laser power consumption of the link, compared to the ideal case where the crosstalk distribution in the MRR filter array is uniform. This chapter presents a novel design of MRR filter array with minimized crosstalk non-uniformity, which can achieve total optical laser power savings of up to 34% of the link power budget.

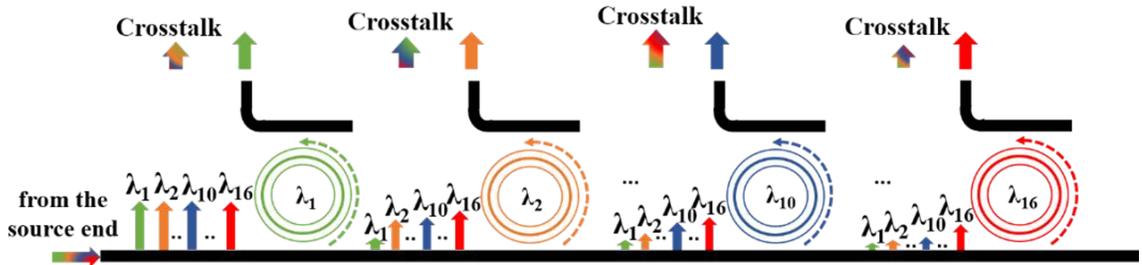


Figure 2.1: An array of MRR filters at the receiver end of a silicon photonic DWDM link. The heights of the crosstalk arrows are proportional to the corresponding power penalty values.

2.1 Introduction and Motivation

To overcome the performance bottlenecks of on-chip communication with ENoCs, recent advances in CMOS-photonics integration [239] have enabled an exciting solution in the form of photonic NoCs (PNoCs). Several PNoC architectures have been proposed to date (e.g., [188, 189, 141]). These architectures employ on-chip photonic links, each of which connects two or more clusters of processing cores. Each photonic link comprises one or more photonic waveguides with dense wavelength division multiplexing (DWDM) of multiple wavelength channels into each waveguide. In a DWDM-enabled waveguide, microring resonator (MR) modulators, which are typically arrayed along the waveguide at the source end, modulate input electrical data signals onto parallel photonic channels. The resultant photonic signals travel through the waveguide and reach the destination end, where an array of MR filters drop the parallel photonic signals onto the adjacent photodetectors to recover the electrical data signals. Thus, DWDM enables high bandwidth parallel data transfers in PNoCs.

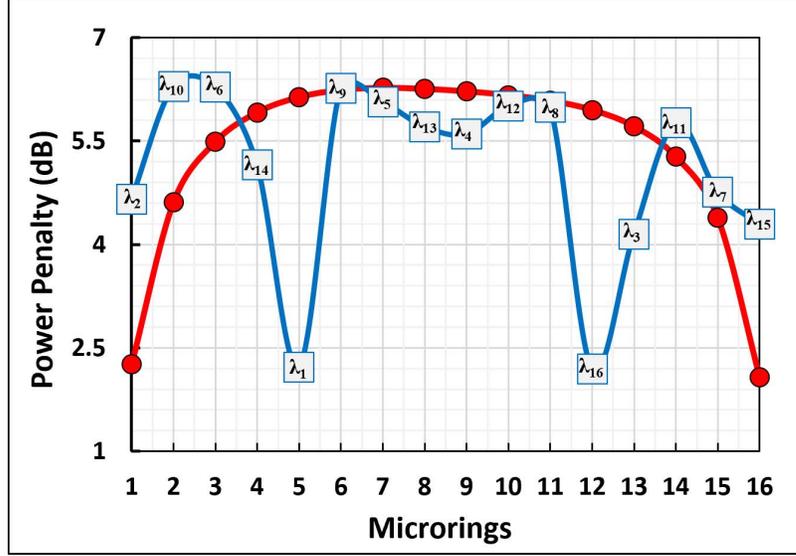


Figure 2.2: Crosstalk penalty distribution across the MRR filter array for two different cases. The values are obtained for 50GHz spacing and MRR quality factor of 8000.

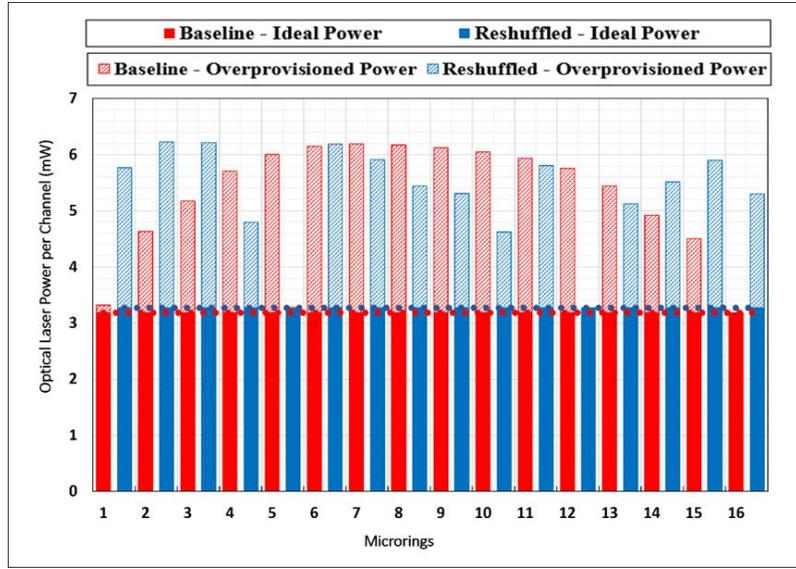
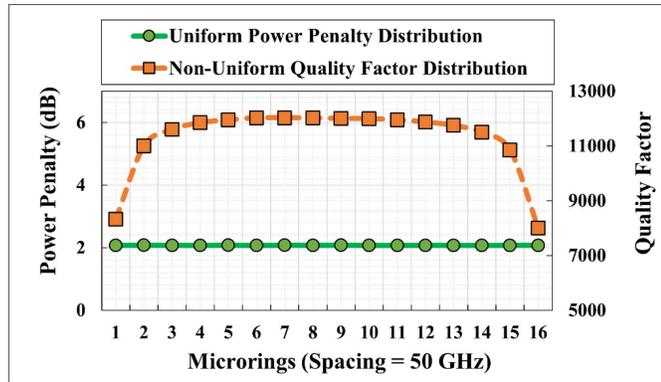


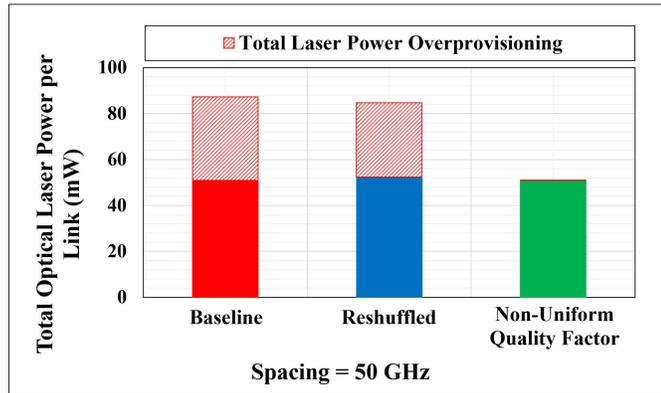
Figure 2.3: Optical laser power per channel, for the baseline and reshuffled cases, evaluated for 50GHz spacing, MR quality factor of 8000, and total link power budget of 100mW.

Unfortunately, DWDM links of PNoCs may suffer from spectral degradation of photonic channels and inter-channel crosstalk [58], which is treated as an optical power penalty in our model. The power penalty is the extra optical power required to compensate for the effects of signal degradation on bit-error-ratio (BER) [18]. As discussed in [18], due to the non-ideal transmission characteristics of DWDM photonic links and MR filter arrays, the photonic channels at the receiver end of a photonic link face non-uniform magnitudes of crosstalk and related power penalties. For example, Fig. 2.1 illustrates the MR filter array of an example single waveguide DWDM link

with 16 photonic channels (λ_1 - λ_{16}). From the figure, every MR filter in the array drops varying amount of power from the neighboring channels on its drop port as crosstalk. The first and last MR filters (λ_1 - λ_{16}) have crosstalk channels on only one side of the DWDM spectrum. Therefore, they see the least crosstalk power at their drop ports. Moreover, as the photonic signals travel along the waveguide, they are progressively dropped by the MR filters, contributing progressively varying amount of crosstalk at the drop ports of MR filters. As a result, the MR filter array sees a non-uniform distribution of crosstalk power penalties across the photonic channels (Fig. 2.2, Baseline – red curve). For example, from Fig. 2.2 (red curve), MR #16 (λ_{16}) faces the minimum crosstalk penalty of 2.1dB, whereas MR #7 (λ_7) faces the maximum crosstalk penalty of 6.3dB, yielding the variance in penalty across the array to be 4.2dB.



(a)



(b)

Figure 2.4: (a) Uniform distribution of crosstalk penalty and non-uniform distribution of quality factors, across the MR filters, and (b) total link-level optical laser power for three different cases.

Overcoming these non-uniformly distributed crosstalk penalties, which is imperative to achieve uniform BER across all the channels, requires non-uniformly distributed laser power across the channels (Baseline results in Fig. 2.3 – red bars). This in turn results in total 34.36mW of laser power overprovisioning (patterned red

bars in Fig. 2.3) for all channels in the link, compared to the ideal case (solid red bars in Fig. 2.3) where the distribution of crosstalk penalty and laser power across all the channels is uniform. Therefore, to minimize the total laser power overprovisioning in the link, the crosstalk penalty distribution in the MR filter array should be uniformized.

As examined in [18], reshuffling the assignments of the individual MR filters to the utilized photonic channels (so that MR #1 is not assigned to λ_1 channel, and so forth) can flatten the crosstalk penalty distribution and reduce the total laser power overprovisioning in the link. The blue curve in Fig. 2.2 shows crosstalk penalty distribution across the MR filter array for the best pattern out of all possible reshuffled filter-channel assignments. This pattern yields minimum crosstalk penalty of 2.2dB for MR #12 (channel λ_{16}) and maximum crosstalk penalty of 6.3dB for MR #2 (channel λ_{10}), resulting in total laser power overprovisioning of 32.3mW for the link (Fig. 2.3 – blue bars) that is 4mW less than the baseline case. This improvement in crosstalk penalty distribution and resultant reduction in the total laser power overprovisioning is negligible. As a result, the non-uniformity in the crosstalk penalty distribution across the MR filter array still exists. To overcome this problem, we propose a novel design of MR filter array with a non-uniform quality-factor distribution across the individual MR filters, as discussed next.

2.2 Proposed Method

In our proposed method, to achieve a uniform distribution of crosstalk penalty across the channels, each individual MR filter in the array is designed with a different quality factor (Fig. 2.4(a); yellow curve – right vertical axis). As a result, our designed MR filter array achieves a flat/uniform distribution of power penalty across the channels, as shown in Fig. 2.4(a) (green curve – left vertical axis). With the uniformized crosstalk penalty distribution across the channels, the total laser power over provisioning reduces to 76 μ W, which in turn reduces the total link power to 51 mW (Fig. 2.4(b) – green bar). Compared to the reshuffled design of MR filter array from [18], the total laser power over provisioning in the link for our design reduces by 32 mW. Moreover, the total link power for our design also reduces by 34 mW, which is 34% of the total link power budget of 100mW. This is because a higher quality factor for an MR filter in our designed array reduces the crosstalk power at its drop port, as the crosstalk power is inversely related to the MR quality-factor [18]. Therefore, carefully choosing the quality-factor of each MR filter in our designed array based on an exhaustive search plays a vital role in uniformizing the crosstalk penalty across the channels.

We propose to define the quality-factor of each MR filter in our filter array at the design time. For that we adopt the MR design from [250], where every MR has an embedded PN-junction at its drop-port. From [250], the carrier concentration in the drop-port PN-junction can be dynamically altered to modulate the drop-port coupling coefficient, and consequently, the loaded quality factor of an MR.

2.3 Summary

This chapter presents a novel idea of using a non-uniform quality factor distribution across an array of MR filters in a photonic link, to uniformize their crosstalk performance, and hence, decrease the total laser power consumption in the link. Our analysis shows that DWDM photonic links that utilize our designed MR filter array can achieve total optical laser power savings of up to 34 mW.

Chapter 3 Redesigning Photonic Interconnects with Silicon-on-Sapphire Device Platform for Ultra-Low-Energy On-Chip Communication

3.1 Introduction

With rapidly increasing demand for data-centric high-performance computing, future manycore processors will require exceedingly high communication bandwidth from the on-chip interconnection networks. However, traditional electrical networks-on-chip (ENoCs) already consume extravagantly large amount of chip area and total system power, which makes the energy-efficient scaling of their bandwidth improbable. This motivates the need for a new interconnect technology that can be leveraged to realize extremely high-throughput (>1 terabits/s) and energy-efficient (< 1 pJ/bit) interconnects for future manycore computing architectures.

Recent advancements in CMOS-photonics integration [240] have enabled an exciting solution in the form of photonic NoCs (PNoCs). Several PNoC architectures have been proposed thus far (e.g., [189, 188, 141]). PNoC architectures typically employ on-chip photonic links, each of which connects two or more clusters of processing cores. Each photonic link comprises of one or more photonic silicon-on-insulator (SOI) waveguides with dense wavelength division multiplexing (DWDM) of multiple wavelengths into each waveguide. In a DWDM-enabled SOI waveguide, SOI microring resonator (MR) modulators, which are arrayed along the waveguide at the source end, modulate input electric signals onto parallel photonic channels. The photonic channels travel through the waveguide and reach the destination end, where an array of SOI MRs drop the parallel photonic signals onto the adjacent photodetectors to recover the electric data signals. Thus, DWDM that utilizes SOI photonic devices enables high-bandwidth parallel data transfer in PNoCs.

A critical parameter for designing a high-throughput SOI based photonic (SOI Ph) link is its optical power budget (OPB), which determines the upper limit of the allowable signal losses and power penalties in the link for the given aggregated data rate ($\#$ DWDM channels (N_λ) \times channel bitrate) of the link. The OPB of a SOI Ph link is the difference between the photodetector noise floor (i.e., receiver sensitivity which has a dependency on bit-rate [20]) and the maximum allowable optical power (MAOP) in the link. The MAOP in a SOI Ph link is determined by the optical non-linear effects of silicon in the constituent SOI Ph waveguides and MR modulators. The primary non-linear effect in silicon at the operating wavelengths of the SOI Ph platform (i.e., $1.3\mu\text{m}$ - $1.6\mu\text{m}$) is two-photon absorption (TPA) [147], which has been shown to induce strong free carrier absorption (FCA) and free-carrier dispersion (FCD) effects in silicon [149], significantly increasing the absorption losses in SOI Ph waveguides [100] and causing self-heating and irreparable resonance shifts in SOI Ph MR modulators [147]. As recently demonstrated in [20], these TPA induced effects limit the MAOP in SOI Ph links below 20dBm, which in turn restricts the achievable link data rate below 900 Gb/s and energy-efficiency above ~ 2 pJ/bit, even for the most optimistic SOI Ph device parameters from [157]. Therefore, to achieve > 1 terabits/s aggregated

data rate and sub-pJ/bit energy-efficiency for SOIPh links, which is a very important step towards realizing the exascale computing systems of the future [224], the TPA effect in silicon must be alleviated to increase the MAOP in SOIPh links.

In this chapter, we present silicon-on-sapphire (SOS) based photonic platform as a potential solution that can mitigate the TPA related shortcomings of the SOIPh platform. The fact that underpins our rationale is that SOS-based photonic (SOSPh) waveguides and MRs have been shown to exhibit low absorption losses and no TPA for the operating wavelengths in the mid-infrared region near $4\mu\text{m}$ [157][105]. The SOSPh platform has these advantages near $4\mu\text{m}$ wavelength region, compared to the SOIPh platform, because near $4\mu\text{m}$ wavelength sapphire has lower material losses than SiO_2 [233] and silicon bandgap is smaller than the total energy of two absorbed photons [190]. Although several prior works have demonstrated the usefulness of SOSPh devices for optical signal processing (e.g., [149, 234, 144, 54, 216]), no prior work has yet explored SOSPh devices for realizing on-chip interconnects. Therefore, in this chapter, using our detailed modeling at the device- and link-level as well as extensive system-level analysis, we show for the first time that SOSPh interconnects can pave the way for realizing extreme-throughput (> 1 terabits/s) and ultra-low-energy (< 1 pJ/bit) on-chip data communication. The key contributions of this chapter are summarized below:

1. We characterize different types of losses and optical properties of SOSPh waveguides and MRs to derive compact design models;
2. We use our developed compact models to derive a new set of guidelines for designing SOSPh links and PNoC architectures;
3. We utilize our developed guidelines to optimize the designs of SOSPh links, and then compare their aggregated data rate and energy-per-bit values to the optimized designs of SOIPh links;
4. We evaluate the impact of optimized designs of SOSPh and SOIPh links on the performance and energy-efficiency of a well-known Clos PNoC architecture [116];

3.2 Motivation

To demonstrate the limitations of the SOIPh device platforms in general, we used different SOIPh platforms from [236], [190, 290, 44] to perform a design analysis for on-chip links following the more realistic design guidelines given in [20]. Results of our analysis are given in Fig. 3.1. Fig. 3.1(a) depicts how the OPB in various SOIPh links (corresponding to the SOIPh platforms from [236], [190, 290, 44]) is utilized depending on the losses present in the links, whereas Fig. 3.1(b) shows the best achievable aggregate data rate (i.e., #DWDM channels $N_\lambda \times$ channel bitrate) and energy-per-bit (EPB) values for the links. We also show our projected results for our target (preferred) photonic platform. In Fig. 3.1(a), the MAOP for the OPB values of all SOIPh links is considered to be 20dBm. Moreover, the EPB values in Fig.

3.1(b) present total EPB values that include contributions from the link laser power, thermal tuning power, modulator driver power, and receiver power, as outlined in the guidelines from [27]. From Fig. 3.1(a), different SOIPh links experience different amounts of total optical power loss (including crosstalk and signal degradation related power penalties [19]). This total power loss whittles down the OPB of all SOIPh links, leaving only a smaller portion of the OPB available to support aggregated data rate. For example, among all considered SOIPh platforms, the SOIPh platform named “zero-change” from [236] has the largest OPB of 51.5dB, which corresponds to -31.5dBm detector sensitivity and the TPA-limited MAOP of 20dBm [100]. From this 51.5dB OPB, 21.15dB portion is whittled down due to optical losses, which leaves 30.35dB of the OPB available for supporting the highest data rate in Fig. 3.1(b) of 636 Gb/s. This larger value of aggregate data rate better prorates the EPB contributions from laser power, thermal tuning, modulator power, and receiver power to yield the lowest total EPB value in Fig. 3.1(b) of 2.1pJ//bit for the SOIPh platform “zero-change”.

Clearly, higher aggregate data rate and lower EPB can be achieved for the “zero-change” platform, if the MAOP for it can be increased from 20dBm and/or total optical loss can be decreased, so that a larger portion of its OPB can be rendered available to support larger aggregate data rate (i.e., larger N_λ and/or higher channel bitrate). Therefore, we envision a target platform (Fig. 3.1) that can increase the MOAP to 22dBm and reduce the total optical loss to 11.9dB, to yield a higher OPB that can support aggregate data rate of up to 1600 Gb/s and EPB of 1 terabits/s) and ultra-low-energy (<1pJ/bit) photonic interconnects can be realized using our proposed SOSPh device links.

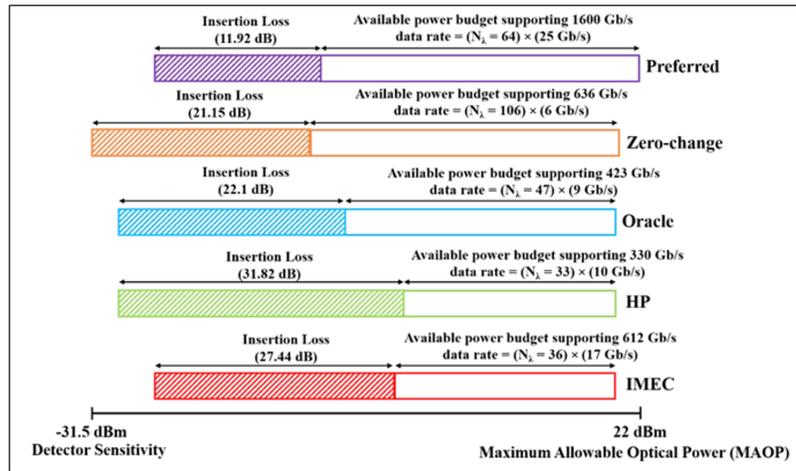
3.3 Modelling of SOS-based Devices

It is established from prior works (e.g., [20][249]) that the performance and energy-efficiency of photonic interconnects depend on the optical characteristics of the constituent waveguides and MR devices. Crucial optical characteristics for photonic interconnects include optical losses in waveguides and spectral footprints (e.g., Q-factor, free-spectral range (FSR)) of MR devices. In this section, we derive compact models for the optical characteristics of SOSPh waveguides and active/passive MR devices, and compare these models with the compact models for SOIPh devices from prior work. As SOSPh devices have been shown to exhibit low absorption losses and no TPA for wavelengths near $4\mu\text{m}$ region [157][105], we model the SOSPh devices to be operating at wavelengths near $4\mu\text{m}$.

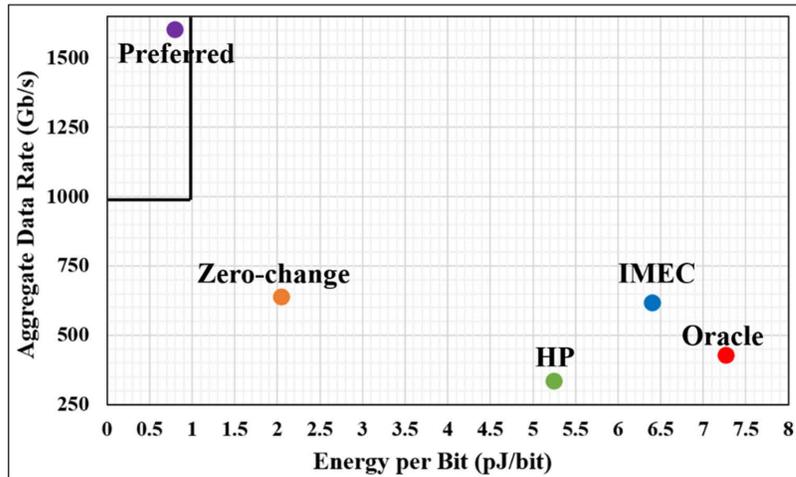
Modelling of SOS-based Passive Devices

Modelling of SOS-based Passive Waveguides

We use Fourier and finite difference time domain (FDTD) analysis methods using a commercial grade tool from Lumerical [155], to model the dimensions and losses in SOS passive waveguides. From our analysis, the cross-sectional dimensions of an SOS



(a)



(b)

Figure 3.1: (a) Distribution of optical power budget (OPB), and (b) Best Achievable aggregate data rate ($\#$ DWDM channels (N_λ) \times channel bitrate) vs energy-per-bit (EPB), for our analyzed photonic links based on SOIPh platforms from prior works and our target (preferred) photonic platform

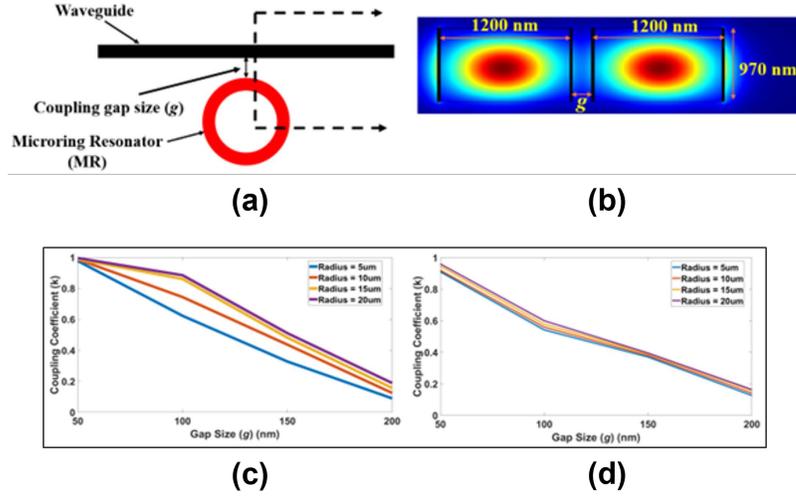


Figure 3.2: (a),(b) Depiction of the cross-sectional dimensions of and the coupling gap size (g) between a waveguide and an MR; and MR coupling coefficient (k) as a function of gap size (g) and MR radius (R) for (c) SOSPh platform and (d) SOIPh platform.

Table 3.1: Various types of losses and optical parameters for SOS and SOI devices.

Type of Loss	SOS	SOI
Waveguide Scattering Loss (dB/cm)	0.9374	1.4
Waveguide Absorption Loss (dB/cm)	10-8	0.1
Waveguide Sidewall Roughness (σ) (nm) [7]	4	6
Core-cladding Refractive-Index Contrast (Δn)	1.67	2.06
MRR Bending Loss (dB/rad)	0.004	0.0073

channel waveguide (Fig. 3.2(a)) that can support the single-mode operation near $4\mu\text{m}$ wavelength with at least 80% optical confinement were found to be $1200\text{nm} \times 970\text{nm}$, which are significantly larger than the dimensions ($450\text{nm} \times 220\text{nm}$) of a typical SOI channel waveguide operating near $1.5\mu\text{m}$. We evaluate the scattering loss and absorption loss of SOS and SOI channel waveguides using the models and methods from [111] and [157]. Our evaluated loss values are given in Table 3.1. Both silicon and sapphire exhibit lower material loss near $4\mu\text{m}$ region [233], which results in lower absorption loss for SOS waveguides. On the other hand, from [150], the scattering loss in a waveguide depends on the core-cladding refractive-index contrast (Δn) and the ratio (σ/λ) of waveguide sidewall roughness (σ) to the operating wavelength (λ). With negligible differences in Δn and σ between SOI and SOS waveguides (Table 3.1), longer operating wavelengths results in lower scattering losses for SOS waveguides.

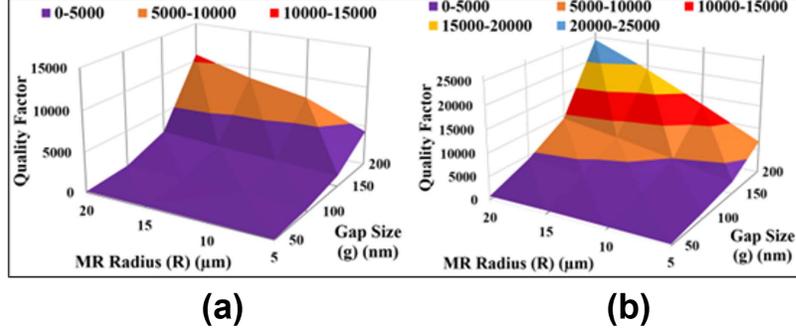


Figure 3.3: Quality factor (Q-factor) based on coupling gap size (g) and MR radius (R) for (a) SOSPh platform, and (b) SOIPh platform.

Modelling of SOS-based Passive Microring Resonators

In this subsection, we present our compact models that relate an MR's Q-factor with its radius (R) and coupling gap size (g) (i.e., gap size between the rectilinear waveguide and MR) (Fig. 3.2(a)). From [27], the Q-factor of an MR depends on the total round-trip loss in the MR's waveguide, which is the sum of scattering loss, absorption loss and bending loss (Table 3.1). To derive the bending loss values (Table 3.1), we used the Eigenmode-solver based methods described in [213].

As a first step towards deriving our intended compact models, we analyzed coupling coefficient (κ) of an MR as a function of R and g . As g increases, the power coupled into the MR from the rectilinear waveguide decreases, which in turn decreases κ . For SOS and SOI MRs, κ can be calculated using Eq. 3.1 [209]:

$$k = \sin \left(2\pi \frac{L}{\lambda_{res}} \frac{n_{eff,even} - n_{eff,odd}}{2} \right) \quad (3.1)$$

Where L is the MR circumference given as $L = 2 \times \pi \times (\text{MR radius } (R))$, λ_{res} is the MR's resonance wavelength, $n_{eff,even}$ is even-mode effective index and $n_{eff,odd}$ is the odd-mode effective index. We used FDTD simulations to extract $n_{eff,even}$ and $n_{eff,odd}$ values.

Fig. 3.2 gives κ values for SOSPh and SOIPh MRs as a function of g and R . From the figure, for $R = 10\mu\text{m}$ and $g = 50\text{nm}$, $\kappa = 0.987$ for the SOSPh MR, whereas $\kappa = 0.92$ for the SOIPh MR. Thus, SOSPh MRs achieve larger values of κ at lower gap sizes. Also, for $R = 15\mu\text{m}$, as g increases from 50nm to 150nm , κ for SOSPh MRs decreases from 0.988 to 0.4825 , whereas for SOIPh MRs κ decreases from 0.92 to 0.39 . Thus, for SOSPh MRs κ decreases less rapidly with increase in g compared to SOIPh MRs.

This type of intricate behavior of κ results into an elaborate relation of MR Q-factor with R and g . To characterize this relation, we plugged our obtained κ values from Table 3.1 in Eq. 3.2 [41]:

$$Q = \frac{\pi n_g L \sqrt{ra}}{\lambda_{res}(1 - ra)} \quad (3.2)$$

Where n_g is group index of silicon, r is cross coupling coefficient ($r = \sqrt{1 - \kappa^2}$), a is round-trip loss coefficient, with other symbols defined with Eq. 3.2. To obtain total loss for a round trip length of an MR along its circumference L is calculated based on the loss values from Table 3.1. Our obtained Q-factor values for SOSPh and SOI Ph MRs are shown in Fig. 3.3(a) and 3.3(b).

From Fig. 3.3, for given R and g values, Q-factor values for SOS MRs are lower compared to SOI MRs. This is because, r is lower for SOS MRs compared to SOI MRs (e.g., for $R = 10\mu\text{m}$ and $g = 50\text{nm}$, $r = 0.67$ for SOS MRs, where it is 0.87 for SOI MRs), which together with longer operating wavelengths for SOS MRs (i.e., $\lambda_{res} = 4\mu\text{m}$) results in lower Q-factor values for SOS MRs.

Modelling of SOS-based Active Microring Resonators

Active tuning of MRs' resonance wavelengths is required not only for realizing active devices such as modulators and switches [58], but also for counteracting the fabrication process and thermal variations induced unwanted resonant shifts [275]. A common method of achieving active resonance tuning in MRs is to change the free-carrier concentration in MR cores [275], which in turn changes the MR core's (which is made of silicon in both SOSPh and SOI Ph platforms) refractive index (Δn) and absorption loss coefficient ($\Delta\alpha$) due to the free-carrier dispersion (FCD) and free-carrier absorption (FCA) effects in silicon [232]. We model the relation of Δn and $\Delta\alpha$ with the change in free-carrier concentration using the following equations [174]:

FCD-FCA Equations for SOS (operating wavelength of $\sim 4\mu\text{m}$):

$$\begin{aligned}\Delta\alpha &= (7.45 \times 10^{-22} \Delta N_e^{1.245} + 5.43 \times 10^{-20} \Delta N_h^{1.153}) \\ \Delta n &= - (7.25 \times 10^{-21} \Delta N_e^{0.991} + 9.99 \times 10^{-18} \Delta N_h^{0.839})\end{aligned}\quad (3.3)$$

FCD-FCA Equations for SOI (operating wavelength of $\sim 1.55\mu\text{m}$):

$$\begin{aligned}\Delta\alpha &= (3.0 \times 10^{-18} \Delta N_e + 2.0 \times 10^{-18} \Delta N_h) \\ \Delta n &= - (6.2 \times 10^{-22} \Delta N_e + 6.0 \times 10^{-18} \Delta N_h^{0.8})\end{aligned}\quad (3.4)$$

Where ΔN_e is free-electron concentration and ΔN_h is free-hole concentration. For given $\Delta N_e = 10^{17}\text{cm}^{-3}$ and $\Delta N_h = 10^{18}\text{cm}^{-3}$, $\Delta\alpha$ and absolute Δn values are higher for SOS MRs (i.e., $\Delta\alpha = 4.21$, $-\Delta n = 13.1 \times 10^{-3}$ compared to SOI MRs (i.e., $\Delta\alpha = 2.3$, $-\Delta n = 1.56 \times 10^{-3}$), which means that active tuning of SOS MRs can be achieved with greater energy-efficiency. To evaluate the energy-efficiency of active tuning, we model the dynamic energy-per-bit for tuning (E_{tuning}) of SOS/SOI MRs with the following equation [275]:

$$E_{\text{tuning}} = \frac{V n_g q J}{4 \lambda_r n_f \Gamma} \Delta\lambda_m \quad (3.5)$$

Where V is the tuning voltage across the MR core required to effect the desired change in free-carrier concentration inside the MR core, n_g is group index of silicon, q is charge of an electron, J is the bulk volume of the MR core in which the change in free-carrier concentration occur, λ_{res} is MR resonance wavelength, Γ is the mode

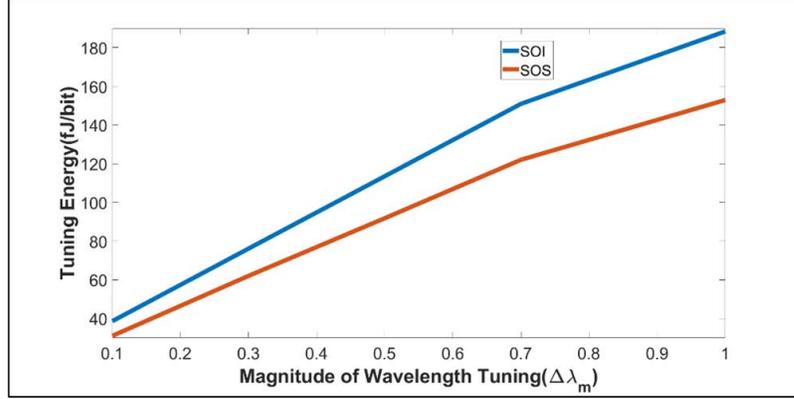


Figure 3.4: MR tuning energy versus magnitude of wavelength tuning ($\Delta\lambda_m$) for SOSPh and SOIPh MRs.

confinement factor (typically $\Gamma = 0.8$), n_f is the ratio of Δn for silicon to the electron-hole pair density that can be evaluated using the formula give in [275] (e.g., $n_f = 2.3 \times 10^{-20} \text{cm}^{-3}$ for SOS and $n_f = 2.13 \times 10^{-21} \text{cm}^{-3}$ for SOI [275]) and $\Delta\lambda_m$ is the magnitude of wavelength tuning.

Fig. 3.4 shows E_{tuning} as a function of $\Delta\lambda_m$. From the figure, E_{tuning} for SOS MRs is lower than that for SOI MRs for the entire range of $\Delta\lambda_m$, which corroborates our earlier observation that the active tuning of SOS MRs can achieve greater energy-efficiency.

Using the device-level compact models derived in this section, we develop new physical-layer design guidelines for SOIPh and SOSPh on-chip links, as described in the next section. Using these guidelines, we evaluate the achievable aggregated datarate and energy-per-bit values for SOIPh and SOSPh on-chip links.

3.4 Link-Level Modelling and Analysis

From [20], the achievable aggregated data rate and energy-per-bit (EPB) values for photonic links not only depend on the OPB of the links and optical characteristics of the constituent devices, but also on several physical-layer design parameters such as the number of DWDM wavelengths (N_λ), free-spectral range (FSR), and OPB. For designing a photonic link, N_λ is the most important design parameter and OPB is the most critical design constraint. For a link, to find the best value of N_λ that can optimally utilize its OPB, the condition given in Eq. 3.6 should be satisfied.

$$OPB(dB) \geq P_{loss}^{dB} + 10 \log_{10} (N_\lambda) \quad (3.6)$$

$$OPB(dB) = MAOP - \text{detector sensitivity} \quad (3.7)$$

P_{loss}^{dB} in Eq. 3.6 accounts for total losses in the link including the signal truncation penalty and modulator/detector crosstalk penalty [18]. From [252], the crosstalk and signal truncation penalties depend on MR Q-factor, channel bit-rate, and inter-channel spacing (which relates to FSR and N_λ [252]). Moreover, the detector sensi-

Table 3.2: Considered Q-factor, FSR, and MAOP values for our analyzed SOSPh and SOIPh links.

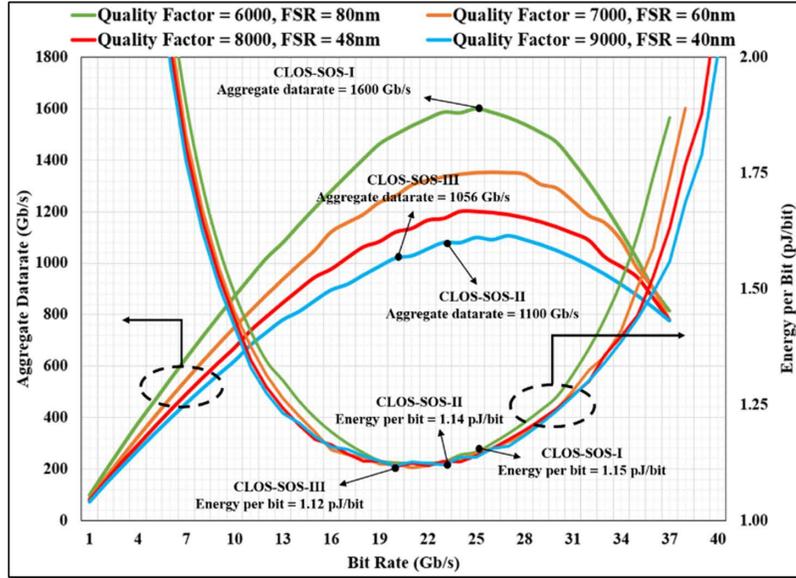
	Considered Q-factor, FSR, and MAOP Values
SOSPh Links	Q-factor=6000, FSR=80nm, MAOP=22dBm
	Q-factor=7000, FSR=60nm, MAOP=22dBm
	Q-factor=8000, FSR=48nm, MAOP=22dBm
	Q-factor=9000, FSR=40nm, MAOP=22dBm
SOIPh Links	Q-factor=6000, FSR=20nm, MAOP=20dBm
	Q-factor=7000, FSR=15nm, MAOP=20dBm
	Q-factor=8000, FSR=13nm, MAOP=20dBm
	Q-factor=9000, FSR=11nm, MAOP=20dBm

tivity in Eq. 3.7 also depends on channel bit-rate [20]. Therefore, for given values of MR Q-factor, FSR, and MAOP (Eq. 3.7), only a unique combination of N_λ and bit-rate can optimally utilize the available OPB while satisfying the condition in Eq. 3.6. This unique optimal combination of N_λ and bit-rate determines the best achievable aggregate data rate (i.e., $N_\lambda \times$ bit-rate) and energy-per-bit (EPB) for the link [20].

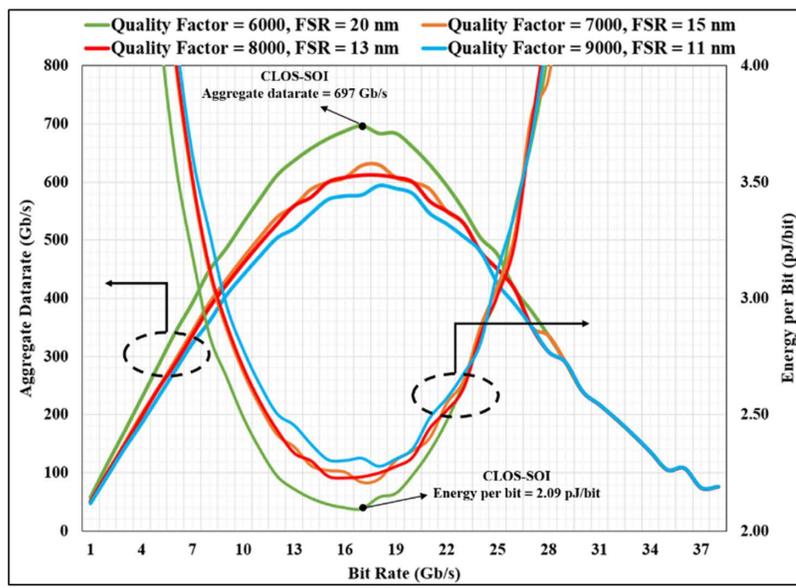
To evaluate the impacts of SOS and SOI devices on the data rate and EPB of links, we use the guidelines given in [20] to design SOIPh and SOSPh on-chip links for four different combinations of MR Q-factor, FSR, and MAOP shown in Table 3.2. For SOIPh links, we choose the TPA limited MAOP value of 20dBm. In contrast, due to the absence of TPA in SOSPh links, it is intuitive to consider a very high value of MAOP. However, we consider a conservative MAOP value of 22dBm for SOSPh links. Our rationale for being conservative is that a not-too-high value of MAOP is more likely to require a reasonable amount of per-wavelength optical power. In contrast, a very high value of MAOP (e.g., >25dBm) can require per-wavelength optical power of greater than 5dBm, which might be very difficult to extract from the state-of-the-art comb laser sources. Moreover, in Table 3.2, we choose the MR Q-factor values in the range from 6000-9000, as it is shown in [18] that this range of Q-factor values can yield minimal values of signal truncation and crosstalk penalties. For these Q-factor values in Table 3.2, we use the device-level compact models from Section 3.3 to reckon the corresponding values of MR radius R, which we use in Eq. 3.8 to reckon the corresponding FSR values.

$$FSR = \frac{\lambda_{res}^2}{2\pi R n_g} \quad (3.8)$$

We use the values from Table 3.2 to design SOSPh and SOIPh links for a well-known PNoC architecture: a 256-core 8-ary 3-stage CLOS PNoC [116]. We consider the worst-case link of CLOS PNoC that has the length of 4.5cm for 22nm technology node. Then, for each value combination in Table 3.2, we sweep the bitrate in the range from 1Gb/s to 40 Gb/s, and use the exhaustive search based heuristic from [100] to find the optimal N_λ for each considered bit-rate value. Then, for each considered bit-rate value, we evaluate aggregate data rate ($N_\lambda \times$ bit-rate) and total EPB (laser



(a)



(b)

Figure 3.5: Aggregate datarate and total energy-per-bit (EPB) values for (a) SOSPh links, and (b) SOIPh links, for different Q-factor, FSR and MAOP value combinations from Table 3.2. The optical losses, laser efficiency, and other device parameters for this analysis are taken from [12] and [16].

Table 3.3: N_λ and bitrate for different variants of CLOS PNoC.

Extracted Link Designs	N_λ	Bit-Rate (Gb/s)	Q-Factor	FSR (nm)	Power Per- λ (dBm)
CLOS-SOI	41	17	6000	18	1.31
CLOS-SOS-I	64	25	6000	80	-8.08
CLOS-SOS-II	44	25	9000	40	-4.83
CLOS-SOS-III	48	22	9000	40	-3.16

+ thermal tuning + modulator driver + receiver) values using EPB models from [19]. These evaluated data rate and EPB values are plotted in Fig. 3.5.

Fig. 3.5(a) (Fig. 3.5(b)) shows the aggregate data rate and EPB values for four different SOSPh (SOIPh) links that correspond to the four combinations of Q-factor, FSR, and MAOP values from Table 3.2. From the figures, the peak aggregate data rate values for four SOSPh links are 1600 Gb/s, 1350 Gb/s, 1200 Gb/s and 1100 Gb/s, and their corresponding EPB values are 1.15 pJ/bit, 1.14 pJ/bit, 1.13 pJ/bit and 1.12 pJ/bit, respectively. On the other hand, the peak aggregate data rate values for SOIPh links are 697 Gb/s, 630 Gb/s, 612 Gb/s and 590 Gb/s, and their corresponding EPB values are 2.09 pJ/bit, 2.22 pJ/bit, 2.23 pJ/bit and 2.28 pJ/bit, respectively. Clearly, SOSPh links achieve higher aggregate data rate and lower EPB values compared to SOIPh links.

To understand the reason behind this outcome, we extract total four link designs from Fig. 3.5(a) and 3.5(b), and list the relevant parameter values for these link designs in Table 3.2. We also present, in Fig. 3.6, how the OPB is utilized for the specific SOSPh and SOIPh link designs from Table 3.2. From Fig. 3.6, it is evident that lower losses and higher MAOP for CLOS-SOS-I, CLOS-SOS-II, and CLOS-SOS-III link designs yield greater aggregate data rate and lower EPB values for them, compared to the CLOS-SOI link design. However, note that CLOS-SOS-I, CLOS-SOS-II, and CLOS-SOS-III link designs still do not achieve sub-pJ EPB values as desired. Nevertheless, as the per-wavelength (per- λ) power requirements for the SOSPh link designs from Table 3.2 are far lower than their saturation point (i.e., 5dBm), these SOSPh link designs still have potential to achieve better (<1 pJ/bit) EPB values by simply allowing greater than 22dBm MAOP per link. Thus, from these results, we can conclude that our proposed SOSPh device platform can pave the way for realizing ultra-low-energy on-chip interconnects of the future.

Excellent link-level results for SOSPh platform cannot guarantee good performance at the system-level, especially for the real world traffic scenarios of on-chip communication. Therefore, to establish a clear winner between the SOIPh and SOSPh platforms, we present benchmark-driven system-level analysis in the next section.

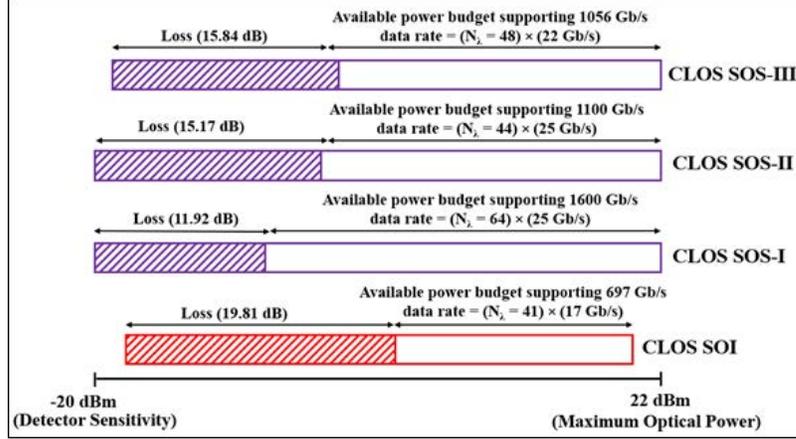


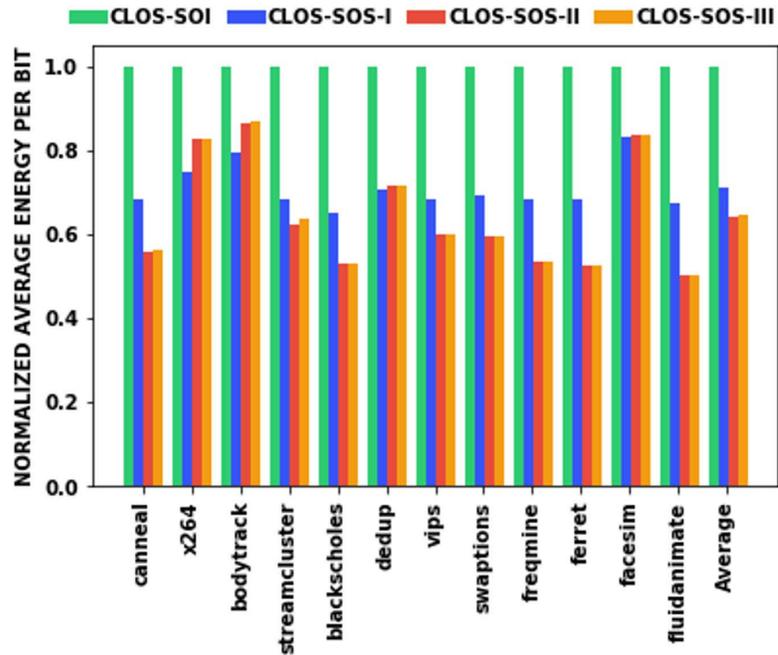
Figure 3.6: Distribution of optical power budget (OPB) for different SOIPh and SOSPh link designs from Table 3.3.

3.5 System-Level Evaluation

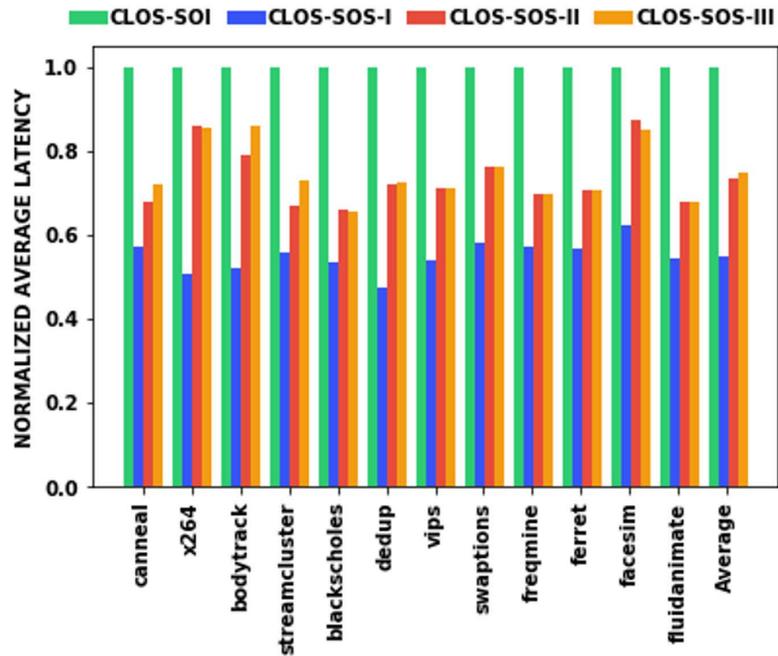
Evaluation Setup

We have done our evaluation on a 256-core system implementing 8-ary 3-stage CLOS topology PNoC [116]. The system has 8 clusters (C1-C8) with 32 cores in each cluster, a group of four cores are connected to a concentrator inside a cluster. There are 8 concentrators in each cluster, and an electrical router connected to them to realize inter-concentrator communication. Point-to-point photonic links are used for inter-cluster communication; a total of 56 single-waveguide links are used to connect all 8 clusters of the CLOS PNoC. Depending on the physical location of source and destination, the point-to-point photonic links use forward or backward propagating wavelengths. Two laser sources are used to enable forward and backward communication in PNoC. The CLOS PNoC uses 1×2 , 1×7 , and 1×4 splitters to power the 56 waveguides.

We performed benchmark-driven simulation-based analysis to evaluate the impact of SOSPh and SOIPh links from Table 3.3 on the performance and energy-efficiency of CLOS PNoC architecture. We used N_λ and bit-rate values from Table 3.3 to model four variants of CLOS PNoC using a cycle-accurate NoC simulator. We evaluated performance for a 256-core single-chip architecture at a 22nm CMOS node. We kept the number of WGs and basic floor plan of the architectures constant across all the variants. We used real-world traffic from applications in the PARSEC benchmark suite [36]. GEM5 full system simulation [38] of parallelized PARSEC applications was used to generate traces that were fed into our cycle-accurate NoC simulator. In GEM5 simulations, we set a “warmup” period of 100 million instructions and then captured traces for the subsequent 1 billion instructions. In our benchmark driven simulations, we evaluated average packet latency, and energy-per-bit (EPB) values for different variants of CLOS PNoC.



(a)



(b)

Figure 3.7: (a) Average energy-per-bit (EPB), and (b) packet latency comparisons for different variants of CLOS PNoC across PARSEC benchmarks. All results are normalized to the baseline CLOS-SOI PNoC results.

Evaluation Results

Fig. 3.7(a) represents a comparison of average packet latency values for the CLOS-SOI, CLOS-SOS-I, CLOS-SOS-II and CLOS-SOS-III PNoCs. As evident, compared to CLOS-SOI PNoC, SOS based PNoCs CLOS-SOS-I, CLOS-SOS-II and CLOS-SOS-III, respectively, have 45%, 26% and 26% lower average packet latency on average. From Table 3.3, CLOS-SOS variants have higher N_λ than CLOS-SOI PNoC, which increases the number of concurrent bits transferred over the network for the CLOS-SOS variants, reducing their average packet latency values. As CLOS-SOS-I has the highest N_λ , it has the least average packet latency. In addition to higher N_λ , SOS variants also have better bit-rate, which increases the rate at which the bits are transferred, eventually contributing to the reduced latency. We can observe that CLOS-SOS-II and CLOS-SOS-III achieve same average latency, this is because CLOS-SOS-II has higher bit-rate which is compensated by CLOS-SOS-III's higher N_λ .

As evident from Fig. 3.7(b), CLOS-SOS-I, CLOS-SOS-II and CLOS-SOS-III have 29%, 37% and 36% lower EPB compared CLOS-SOI on average. As the average latency for the SOS variants is less than CLOS-SOI, energy dissipated is also less. The EPB of CLOS-SOS-I is greater than CLOS-SOS-II and CLOS-SOS-III, as greater N_λ leads to increase in the number of MR modulators and MR detectors in CLOS-SOS-I, which in turn increases the total energy consumption.

3.6 Related Work

Significant research work (e.g., [19, 147, 100, 149]) is available in the literature that focuses on characterizing the two-photon absorption (TPA) and other types of optical non-linear effects in silicon waveguides and resonators. For example, [18] and [147] describe how TPA induced FCD and FCA effects in silicon limit the MAOP in SOIPh links, restricting the scalability of their aggregate data rate and energy-efficiency. However, no prior work has yet explored a solution to the TPA-induced scalability shortcomings of SOIPh interconnects. We for the first time presented SOS-based device platform as a potential solution to the TPA-related scalability issues in on-chip photonic interconnects.

Several SOS-based photonic devices have already been prototyped to be operated near $4\mu\text{m}$ wavelength. These prototypes include on-chip quantum cascade laser sources (e.g., [54]), photonic waveguides and MRs (e.g., [234, 149, 105]), grating couplers (e.g., [144]). Information obtained from all these prototype works, when combined with the knowledge base from this chapter, can catalyze cross-layer research in the area of SOSPh interconnects design, which can enable the widespread adoption of SOSPh platform for realizing extreme-scale on-chip and off-chip communication architectures.

3.7 Overheads and Challenges

To compare the footprint area of SOS and SOI variants of CLOS PNoC architecture from Table 3.3, SOIPh MR has footprint area of $78\mu\text{m}^2$, whereas the footprint areas for SOS-I, SOS-II and SOS-III MRs are $177\mu\text{m}^2$, $707\mu\text{m}^2$ and $708\mu\text{m}^2$ respectively. The footprint area of a 1cm long rectilinear SOIPh waveguide is $4500\mu\text{m}^2$, whereas the footprint area of 1cm long rectilinear SOSPh waveguide is $9700\mu\text{m}^2$. In terms of CLOS PNoC architecture, the total footprint area for SOI-based CLOS PNoC architecture is 0.4mm^2 , whereas the footprint area for SOS-I, SOS-II, and SOS-III based CLOS PNoC architectures are 3.1mm^2 , 2.3mm^2 and 2.4mm^2 , respectively. This comparison clearly shows that SOS links/PNoCs have higher footprint area compared to SOI links/PNoCs.

Note that the traditional fiber optics systems for inter-cluster, inter-data center, and long-haul networks still running on O, L and C optical bands. In contrast, SOSPh platform operates with wavelengths between $2.5\mu\text{m}$ - $4\mu\text{m}$. Therefore, additional specialized equipment and support are needed to introduce SOSPh interconnects in this established hierarchy, which is likely to incur extra cost. Nevertheless, it is worth bearing this extra cost, especially considering the energy and performance benefits of SOSPh platform shown here.

3.8 Summary

Conventional SOI-based photonic interconnects have limited bandwidth-energy scalability due to the optical non-linear effects in silicon, especially the two-photon-absorption (TPA) effect. In this chapter, we presented silicon-on-sapphire (SOS) device platform as a solution to the scalability limitations of SOI-based interconnects. We developed new compact models for SOS devices, utilizing which we formulated new guidelines for designing SOS links and PNoCs. Our link-level analysis showed that SOS links can achieve aggregate data rate of $>1\text{Tb/s}$, which is significantly better than SOI links. Our system-level analysis with CLOS PNoC architecture showed that PNoCs that are designed using SOS devices and links can achieve up to 45% lower latency and 37% lower EPB compared to the PNoCs implemented using the conventional SOI devices and links. These promising results prove that SOS based PNoCs can achieve high-bandwidth data transfers with low latency and greater energy-efficiency, compared to the traditional SOI-based PNoCs.

Chapter 4 Photonic Networks-on-Chip Employing Multilevel Signaling: A Cross-Layer Comparative Study

4.1 Introduction

As the core count in contemporary manycore processing chips increases, the conventional on-chip communication fabrics, i.e., electrical networks-on-chip (ENoCs), experience higher power dissipation and degraded performance. As a potential solution to these shortcomings, ENoCs have been projected to be replaced by emerging photonic network-on-chip (PNoC) fabrics. This is because the recent advancements in silicon photonics have enabled PNoCs to offer several advantages over ENoCs, such as higher bandwidth density, distance-independent data rate, and smaller bandwidth-dependent energy.

Typical PNoC architectures (e.g., [121, 21, 296, 192, 56, 30]) and processor-to-DRAM photonic interconnects (e.g., [253, 254]) utilize several photonic devices such as multi-wavelength lasers, waveguides, splitters and couplers, along with microring resonators (MRs) as modulators, detectors and switches. A broadband laser source generates light of multiple wavelengths (λ s), with each wavelength (λ) serving as a data signal carrier. Simultaneous traversal of multiple optical signals across a single photonic waveguide is possible using dense wavelength-division multiplexing (DWDM), which enables parallel data transfers across the photonic waveguide. For instance, a DWDM of 16λ s in the photonic waveguide can transfer 16 data bits in parallel. At the source node, multiple MRs modulate multiple electronic data signals on the utilized multiplexed λ s (data-modulation phase). In almost all PNoC architectures in literature, modulator MRs utilize on-off keying (OOK) modulation, wherein the high and low intensities of λ s in the waveguide are used to represent, respectively, logic ‘1’ and ‘0’. Similarly, at the destination node, multiple MRs equipped with photodetectors are used to filter and detect λ -modulated data signals from the waveguide (data-detection phase) and convert them back to proportional electrical signals. In general, using a large number of multiplexed λ s enables high-throughput parallel data transfers in PNoCs, hence boosting the bandwidth in such networks.

Leveraging a large number of multiplexed λ s, and thus the resultant high throughput, has been pivotal in PNoC architectures for efficiently amortizing their high non-data-dependent power consumption that includes the laser power and MR tuning power. However, a number of challenges related to area [56], cost [78], reliability [250], and energy-efficiency [251][193] still need to be overcome for efficient implementation of PNoCs that utilize a large number of multiplexed λ s (typically 32 or more multiplexed λ s per waveguide [38],[252]). First, generating a large number of multiplexed λ s requires a comb laser source, the ineffectiveness, complexity, and cost of which increase with the number of generated λ s [48]. Second, utilizing a larger number of multiplexed λ s to achieve higher-throughput data transfers in a PNoC results in higher area and power overheads. A large number of multiplexed λ s require larger network flit size as well as more electrical and photonic hardware such as mod-

ulator and detector MRs and their drivers. A larger network flit size can also result in larger sized electronic buffers in the network gateway interfaces, which can result in significantly higher area and power overheads. Similarly, larger number of MRs and drivers also incur greater photonic area and MR tuning power overheads. Last, the use of a larger number of multiplexed λ s can decrease the viable gap between two successive optical signals, which in turn will increase the inter-channel crosstalk noise in PNoCs, increasing the bit-error rate (BER) of communication [58][57]. As a result of the combined impact of these factors, the use of larger number of multiplexed λ s in PNoCs leads to trade-offs among the achievable throughput, BER, and energy-efficiency.

To mitigate the adverse impacts of these tradeoffs, multi-level optical signaling has been introduced in prior works. For example, in [120] and [121], Kao et al. proposed a multilevel optical signaling format four-pulse amplitude modulation (4-PAM) to achieve higher-throughput and energy-efficient data communication in PNoCs. The 4-PAM optical signaling format doubles the data rate by compressing two bits in one symbol carried out by four levels of optical intensity. In the literature, three different MR-based designs of optical 4-PAM modulators have been proposed. In [77] and [218], two cascaded on-off keying (OOK) modulators are utilized to superimpose two OOK optical signals of the same λ with 2:1 power ratio to create a 4-PAM λ -signal. But this signal superposition based 4-PAM method (referred to as 4-PAM-SS henceforth) incurs substantially high power, photonic area, and reliability overheads at the link-level. Roshan-Zamir et al. in [208] demonstrated a single-MR 4-PAM modulator that takes an electrical 4-PAM signal, generated using a segmented pulsed-cascode amplifier based electrical DAC (EDAC), as input and then converts it into an optical 4-PAM signal. But this EDAC-based conversion method (referred to as 4-PAM-EDAC henceforth) can incur significant power consumption and area overheads due to the required EDACs. In contrast, Moazeni et al. [167] utilized an optical DAC (ODAC) modulator (referred to as 4-PAM-ODAC henceforth) that directly converts two input electrical OOK signals into a 4-PAM optical signal, thereby eliminating the use of EDAC and its overheads. Thus, it is well established how various MR-based modulators can be utilized to generate 4-PAM optical signals. But what is still unknown is how different 4-PAM modulators can be utilized to design DWDM-based photonic links and PNoC architectures. Moreover, the impacts of various 4-PAM modulators on the overall energy, reliability, and performance behavior of the designed links and PNoC architectures also remain unexplored.

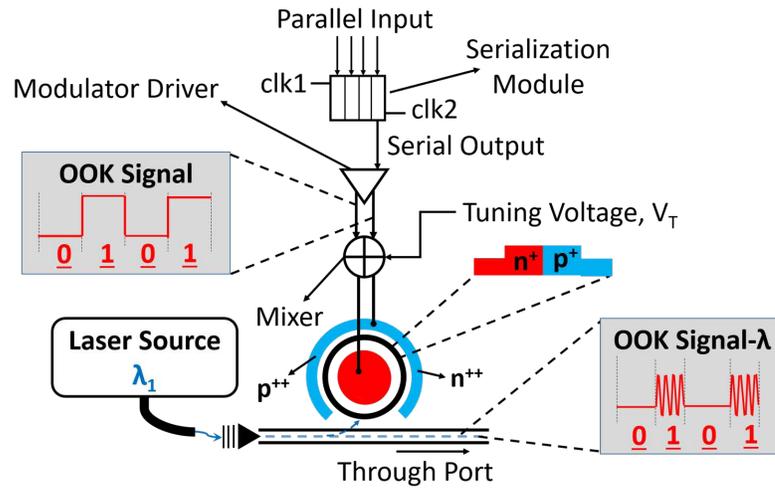
In this chapter, we present a comparative study and a heuristic-based search method for designing DWDM-based on-chip photonic links using different types of MR-based 4-PAM modulators, such as 4-PAM-SS, 4-PAM-EDAC, and 4-PAM-ODAC. We analyze how different types of MR-based 4-PAM modulators compare with the traditional OOK modulators at the photonic link-level and PNoC architecture-level while considering hardware overhead, performance, energy-efficiency, and reliability, and especially in the presence of inter-channel crosstalk. Our analysis shows that designing the constituent photonic links of PNoCs is subject to inherent trade-offs among the achievable performance (aggregated data rate), energy consumption, and reliability, irrespective of the utilized modulation method and modulator type.

Optimizing these design tradeoffs often involves finding the right balance between the photonic link’s aggregated datarate and energy-reliability behavior. We find that different modulation methods and modulator types are differently positioned to achieve this balance: i.e., which modulation method and modulator type achieves better balance really depends on the underlying PNoC architecture. Our novel contributions in this chapter are as follows:

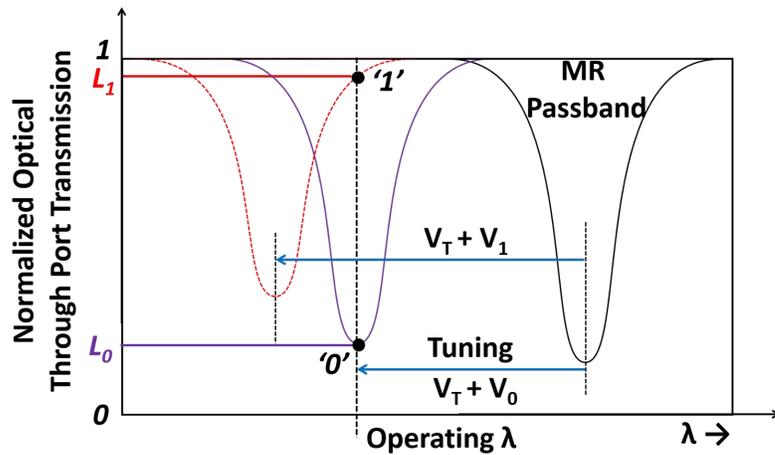
1. We present an overview of how different MR-based 4-PAM modulators generate 4-PAM optical signals, and then compare their operation with a conventional MR-based OOK modulator;
2. We present how the hardware implementation overheads for different 4-PAM modulation methods compare with one another, and with the conventional OOK modulation method;
3. We provide a systematic analysis of various design factors that affect the photonic link-level design tradeoffs for both OOK- and 4-PAM-based links;
4. We utilize a heuristic-based search method to optimize the designs of DWDM-based photonic links with OOK, 4-PAM-SS, 4-PAM-EDAC, and 4-PAM-ODAC modulation methods, to achieve the desired balance between the aggregated datarate and energy-efficiency while achieving the BER of 10^{-9} or lower;
5. We analyze how the optimized OOK and various 4-PAM photonic links affect the performance and energy-efficiency of two well-known PNoC architectures: CLOS PNoC [116] and SWIFT PNoC [56];

4.2 Background: Various Designs of OOK and 4-PAM Modulators from Prior Work

In this section, we present an overview of different MR-based OOK and 4-PAM modulator designs from prior work. In general, an MR-based modulator employs some mechanism to modulate the optical signal transmission at its through port (see Fig. 4.1(a)). In OOK modulators, the through-port optical transmission is modulated between two distinct levels, whereas for 4-PAM modulators it is modulated between four distinct levels. An MR-based modulator is fundamentally a wavelength-selective resonator whose employed modulation mechanism generally alters its resonant wavelength (λ_r) with respect to a utilized carrier (i.e., input) λ . This in turn alters the modulator’s through-port optical transmission at the carrier λ . Most of the MR-based OOK and 4-PAM modulators (shown in Fig. 4.1 to Fig. 4.4) from prior work, e.g., [120], [208], [167], [78], [74], [33], utilize voltage biasing induced free-carrier injection/depletion, and the resultant free-carrier dispersion (FCD) mechanism [257], to modulate their through-port optical transmissions. However, different modular designs differ in their physical implementations, as a result, their area-energy-reliability footprints also differ. The following subsections present the operational details of different MR-based OOK and 4-PAM modulator designs and their physical implementations.



(a)



(b)

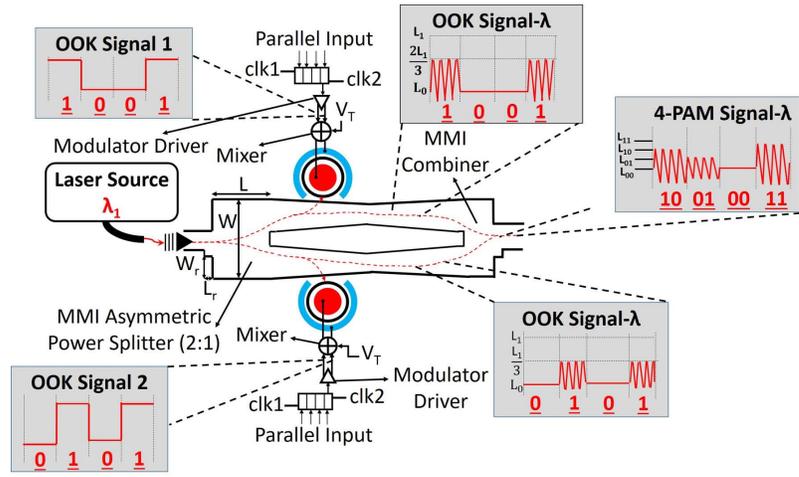
Figure 4.1: Illustration of (a) an MR-based on-off keying (OOK) modulator, and (b) the modulator's resonance passbands and optical transmission levels.

MR-Based On-Off keying (OOK) Modulator

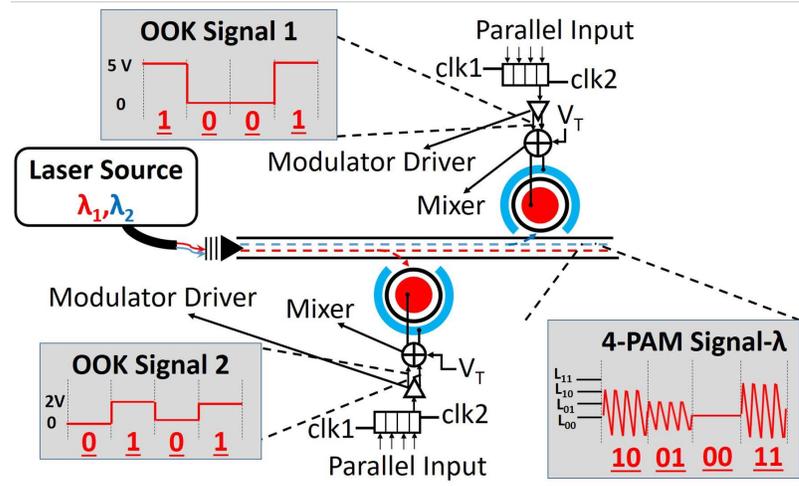
Fig. 4.1(a) illustrates a typical MR-based OOK modulator [33], which employs a serialization module and a driver circuit that can produce a sequence of signal bias voltages corresponding to the input sequence of electrical bits (i.e., ‘1’s and ‘0’s). The modulator MR’s resonance is switched in and out of alignment with signal- λ_1 by applying the sequence of signal-bias voltages to the MR. Before the MR modulator is driven by the signal bias voltages, each signal bias voltage in the sequence might be offset with a corresponding non-zero tuning-bias voltage to compensate for the resonant shift in the MR [183] that can occur because of the variations in the width and thickness of the MR during a conventional non-ideal fabrication process [45]. Such fabrication process related variations are referred to as process variations (PV) henceforth. Such resonant shift in the MR can also occur due to thermal variations (TV). For example, Fig. 4.1(b) illustrates how shifts in the resonance passband of an example OOK MR modulator can modulate its through-port optical transmission. In Fig. 4.1(b), V_T is the tuning bias voltage that depends on the magnitude of the PV-induced MR resonance misalignment, whereas V_1 and V_0 are input signal-bias voltages corresponding to logic ‘1’ and logic ‘0’ bits, respectively. Thus, from the figure, for the OOK MR modulator, the net-bias voltages of $V_T + V_1$ and $V_T + V_0$ yield, respectively, ‘on’ (L_1) and ‘off’ (L_0) levels of through-port optical transmission. As a result, an OOK MR modulator takes a sequence of bias voltages corresponding to data bits as input and generates an on-off keying (OOK) modulated optical signal as output.

MR-Enabled Signal Superposition Based 4-PAM Modulator (4-PAM-SS Modulator)

Fig. 4.2(a) illustrates a signal superposition based 4-PAM modulator design (referred to as 4-PAM-SS) for use in PNoCs, which was first proposed in [120]. From the figure, in a 4-PAM-SS modulator, two OOK MR modulators that are connected in parallel to two different waveguides generate two OOK-modulated optical signals of same λ but of different intensities in the ratio 2:1. These two OOK-modulated optical signals are superposed using a combiner to generate a 4-PAM modulated λ signal. As evident from Fig. 4.2(a), the need of an asymmetric power splitter and combiner can complicate the implementation of this design. This issue can be mitigated by using a different 4-PAM-SS design from [78] as shown in Fig. 4.2(b), which employs two cascaded OOK MR modulators coupled to a single waveguide to eliminate the need for a power splitter and combiner. Both these 4-PAM-SS modulator designs (Figs. 4.2(a) and 4.2(b)) in general require the two OOK-modulated optical signals to be in phase, which may not be possible to achieve under PV and TV, potentially causing some reliability issues that will be discussed in Section 4.3. We utilize the 4-PAM-SS modulator design from [77] (Fig. 4.2(b)) for our analysis presented henceforth.



(a)



(b)

Figure 4.2: 4-PAM-SS modulator designs. (a) Design from [248] with two parallel OOK MR modulators and multi-mode interference (MMI) based asymmetric power splitter-combiner from [221]. (b) Design from [111] with two cascaded OOK MR modulators.

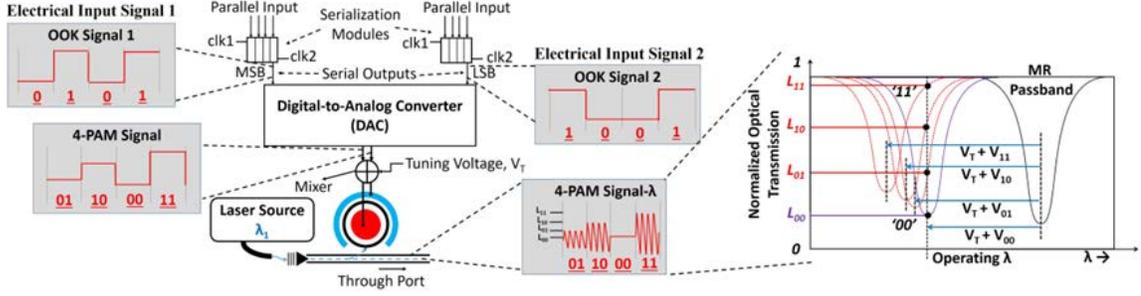


Figure 4.3: Illustration of an electrical DAC (EDAC) enabled MR-based 4-PAM modulator from [208]. Inset: Illustration of resonance passbands and optical transmission levels for an EDAC-enabled MR-based 4-PAM modulator.

4.2.3 Electrical DAC (EDAC) Enabled MR-Based 4-PAM Modulator (4-PAM-EDAC Modulator)

In [208], an MR-based 4-PAM modulator is presented that utilizes an electrical DAC (EDAC) to convert two electrical OOK signals into an electrical 4-PAM signal, as shown in Fig. 4.3. This electrical 4-PAM signal is used by the driver circuit that drives an MR modulator to generate a proportional optical 4-PAM signal. The driver circuit generates four different bias voltages corresponding to the four distinct two-bit patterns (i.e., ‘00’, ‘01’, ‘10’, ‘11’) in the input electrical 4-PAM signal. These four voltages can induce four different optical transmission levels at the through port of the MR modulator, corresponding to four different magnitudes of resonance passband shift in the MR, as shown in Fig. 4.3 (see the inset). To achieve these transmission levels L_{11} , L_{10} , L_{01} , L_{00} (shown in the inset of Fig. 4.3), the signal bias voltages V_{00} , V_{01} , V_{10} , V_{11} of the modulator have to be decided upon appropriately. This can be done efficiently using the pulsed-cascode, a digital-to-analog converter (DAC), based output driver circuit reported in [208]. This circuit from [208] has a provision for sweeping the modulator bias voltages (V_{10} , V_{01}) to determine the target transmission levels (L_{10} , L_{01}) such that they are equidistant from L_{11} and L_{00} , which allows for the in-situ corrections of any degree of aberrations in the transmission levels that can arise due to the fabrication process variation induced changes in the Q-factor and extinction ratio of the modulators. But this EDAC based 4-PAM signaling method incurs substantial area and energy overheads related to the required EDAC circuits [208], which can offset the general benefits of 4-PAM signaling.

Optical DAC (ODAC) Enabled MR-Based 4-PAM Modulator (4-PAM-ODAC Modulator)

To reduce the area and energy overheads of EDAC enabled MR modulators, an optical DAC (ODAC) enabled MR-based 4-PAM modulator was proposed in [167]. This modulator design consists of a spoked MR that functions like an ODAC to directly convert two input electrical OOK signals into a 4-PAM optical signal. A spoked MR is realized by segmenting its embedded P-N junction into multiple anodes and cathodes (e.g., 32 anodes and 32 cathodes in [167], and 15 anodes and 15 cathodes in the MR

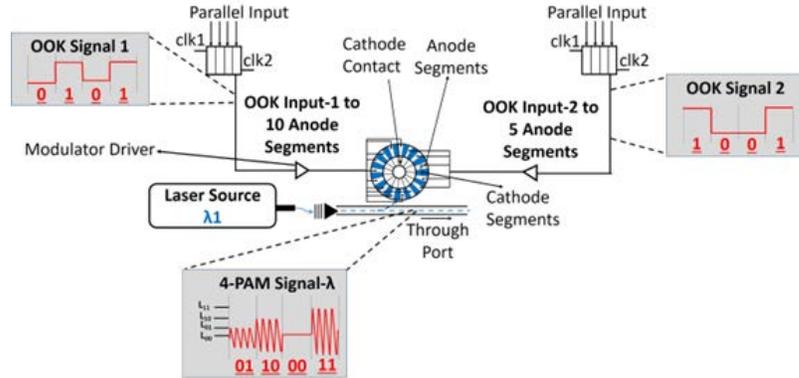


Figure 4.4: Optical DAC (ODAC) enabled MR-based 4-PAM signaling modulator from [167].

modulator shown in Fig. 4.4). All cathode segments are connected together via a spoked-ring shape metal contact in the center of the MR, while each anode segment has its own contact pin using which each anode segment can be driven independently or in some combination of other anode segments. For instance, in Fig. 4.4, a total of 10 out of 15 anode segments are connected and driven by electrical OOK signal 1, and the remaining 5 anode segments are driven by electrical OOK signal 2. This arrangement of the MR modulator’s anode connections corresponds to four distinct spectral positions of the MR’s resonance passband, which in turn corresponds to four distinct levels of optical transmission at the MR’s through port. Thus, this spoked-MR-based modulator design functions like an ODAC to reduce the typical two-stage electro-optic OOK-to-4PAM conversion process to a single-stage process. Compared to the other MR-based 4-PAM modulators, this ODAC-enabled spoked-MR based 4-PAM modulator exhibits low area overhead and dynamic energy consumption [167].

In summary, different MR-based OOK and 4-PAM modulator designs function differently at the device level. Due to these functional differences, it can be intuitive inferred that different modulator designs would have different energy-performance behavior and implementation overheads at the link- and system-level. In the next section, we systematically analyze the physical-layer design overheads and static power consumption for various photonic link implementations that are based on different modulator designs and signaling methods discussed here.

4.3 Systematic Analysis of Photonic Links with Various Modulator Implementations

Recent advancements in CMOS-photonics integration (e.g., as demonstrated in ([146, 145, 256, 198]) have enabled an exciting solution in the form of photonic network-on-chip (PNoC) architectures. Several PNoC architectures have been proposed till date ([30, 28, 173]) that employ either fully optical interconnects or hybrid optical-electrical interconnects. In this section, we identify the physical-layer hardware components of PNoCs and their building blocks (i.e., photonic links), whose implementation overheads are highly affected by the choice of signaling method and modulator design.

Typically, a PNoC comprises of multiple photonic links. A photonic link comprises of one or more photonic waveguides, which move data packets between sender and receiver nodes in the optical domain over multiple DWDM wavelength channels. However, all data packet transfers outside of the PNoC in a manycore processor chip, e.g., between the processing cores and the PNoC, still occur in the electrical domain. Therefore, in a photonic link, it is important to enable electrical-to-optical (E/O) conversion of incoming data packets, which is typically achieved using a bank of MR-based modulators at the sender node. Similarly, to enable optical-to-electrical (O/E) conversion of outgoing data packets from a link, a bank of MR-based filters and photodetectors are employed at the receiver node. Both OOK and 4-PAM optical signaling based links require E/O conversion at the sender nodes and O/E conversion at the receiver nodes. The O/E converted signals at the output of photodetectors generally follow the format of the input optical signals, i.e., an OOK (4-PAM) modulated optical signal is converted into an OOK (4-PAM) modulated electrical signal by the photodetector. The same photodetector can be used to convert both OOK and 4PAM modulated optical signals to electrical signals. These photodetector output signals are generally reshaped by trans-impedance receiver modules to make them digitally processable.

Figs. 4.5(a) and 4.5(b) show the schematics of example trans-impedance receiver modules for OOK and 4-PAM signals, respectively. From the figures, the example 4-PAM receiver module employs three trans-impedance op-amps to generate two bit-streams, compared to the example OOK receiver that employs one trans-impedance op-amp to generate one bit-stream. The E/O and O/E conversion of signals in DWDM photonic links also utilize serialization and deserialization modules. At the E/O conversion unit of a DWDM photonic link, the converted optical data packets are transferred over different channels (i.e., each wavelength is an optical channel) at a higher bitrate than the bitrate of the incoming electrical data packets. To enable this conversion between bitrates, a serialization module is utilized before each MR-based modulator at the source node, and a deserialization module is used after each MR-based detector at the receiver node. Serialization modules can be implemented using parallel-in serial-out electronic buffers, whereas deserialization modules can be implemented using serial-in parallel-out electronic buffers, as shown in Figs. 4.5(c) and 4.5(d), respectively. From Fig. 4.5 and compared to OOK signaling, for a link with N_λ wavelengths, using 4-PAM signaling (i.e., $B = 2$ in Fig. 4.5) requires $2\times$ narrower electronic buffers in each (de)serialization module of the link. This is because using 4-PAM signaling in the link requires $2\times$ number of (de)serialization modules compared to OOK signaling. As a result, for 4-PAM signaling, each incoming/outgoing data packet is striped across $2\times$ number of electronic buffers (corresponding to $2\times$ number of (de)serialization modules), allowing each buffer to be $2\times$ narrower.

In summary, for a DWDM link, the overhead (e.g., area, power consumption) of implementing (de)serialization, and E/O and O/E conversion ultimately depends on the choice of signaling method and modulator design. This is because such a choice directly controls the required number of (de)serialization modules, number of MR-based modulators and filters, and the required type and count of photodetectors and receiver modules. Table 4.1 gives the number of required modules/instances of

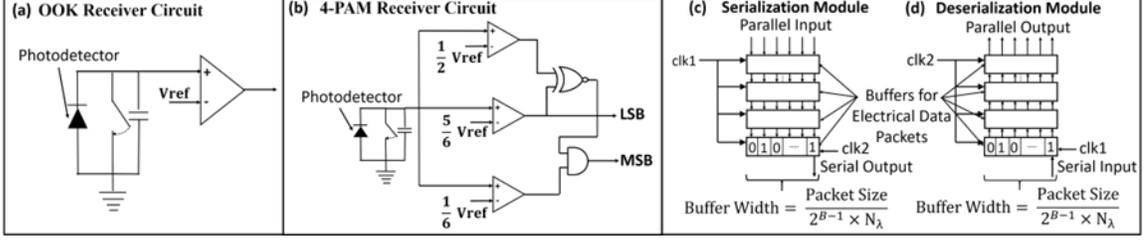


Figure 4.5: Schematics of (a) a receiver module for an OOK modulation-based link [248, 252], (b) a receiver module for a 4-PAM modulation-based link [248, 252], (c) a serialization module [174], and (d) a deserialization module [174]. N_λ is the number of DWDM signals in the link. B is number of bits per symbol; $B=1$ for OOK signaling, and $B=2$ for 4-PAM signaling

several hardware components (e.g., MR modulators, MR filters, photodetectors) for implementing DWDM photonic links with various signaling methods, such as OOK, 4-PAM-SS, 4-PAM-EDAC, and 4-PAM-ODAC. Table 2 gives example values (extracted from prior work) for dynamic energy consumption of several hardware components. The energy consumption value for a TIA based receiver front-end (ETI-OPAMP) can change with the change in the technology node. But we do not expect $E^{TI-OPAMP}$ to affect the results provided in Tables 4.1 and 4.2, because $E^{TI-OPAMP}$ does not affect any of the link configuration parameters such as N_λ , PP_{dB} , S , or BR . Nevertheless, we point the reader to [199] for more detailed analysis of how the dynamic energy consumption of a TIA circuit changes for different design parameters of the circuit. Nevertheless, note that the study presented in this chapter is independent of the parameter values provided in Table 4.1 and can be applied considering other parameter values. Moreover, we adopt the common research approach from prior works ([100, 28, 141, 173]) and select the energy consumption values for various devices from different references (Tables 4.2 and 4.4) to undertake the link and system-level evaluations presented in this chapter. We discuss the information provided in Table 4.1 in the upcoming subsections.

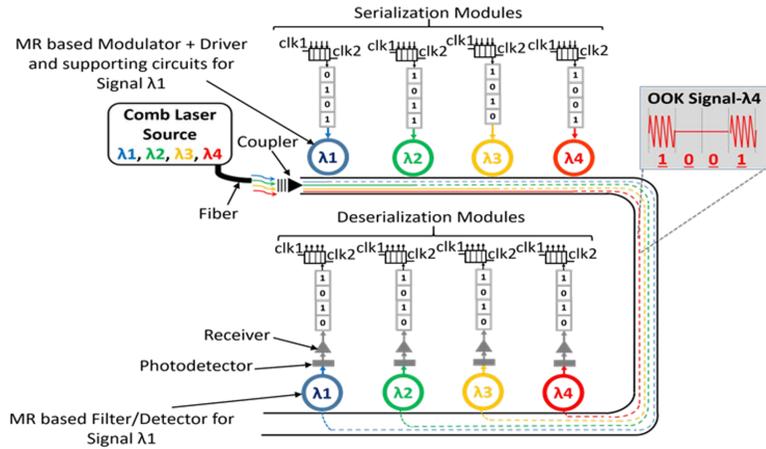
Photonic Links based on OOK Modulation

Fig. 4.6(a) shows a schematic of an OOK signaling based DWDM link with four optical channels ($N_\lambda = 4$). From the figure, the link utilizes four instances of MR modulators, modulator drivers, MR filters, photodetectors, receiver modules, serialization modules, and deserialization modules each. Therefore, one can generalize that for a DWDM OOK link with N_λ channels, it would require N_λ instances of each of the various hardware components as mentioned in Table 4.1. Moreover, the link uses one trans-impedance op-amp per receiver module (Fig. 4.5(a)), requiring N_λ trans-impedance op-amps corresponding to N_λ receiver modules (Table 4.1). Lastly, having a total of N_λ (de)serialization modules per link leads to each buffer being of size $(\text{Packet Size}/N_\lambda)$ bits wide (Figs. 4.5(c) and 4.4(d)), as $B=1$ for OOK links in Figs. 4.5(c) and 4.5(d). Moreover, the link has total energy-per-bit (EPB) and static power consumption values associated with various hardware components (see Table

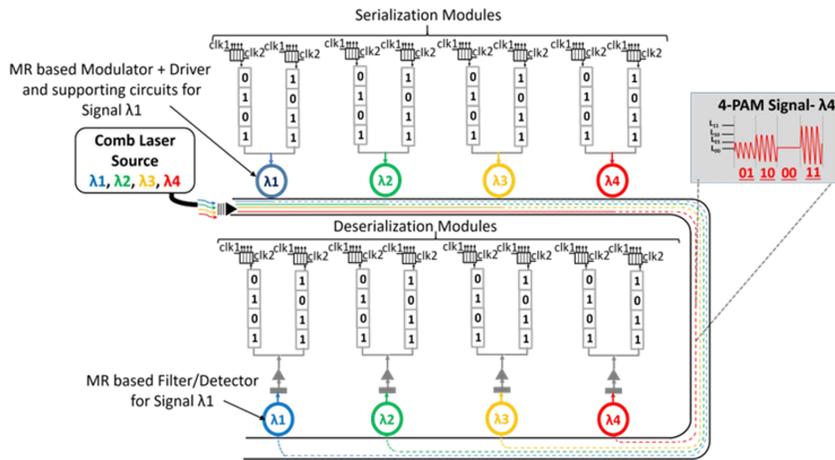
Table 4.1: Number of instances, dynamic energy-per-bit (EPB), and static power values for various hardware components required for implementing a DWDM photonic link with N_λ channels using various signaling methods and modulator types. PS is packet size in bits.

Parameter	Signaling Method/Modulator Type			
	OOK	4-PAM		
		SS	EDAC	ODAC
#instances of various hardware components				
# MR modulators	N_λ	$2 \times N_\lambda$	N_λ	N_λ
# MR filters	N_λ		N_λ	
# Photodetectors	N_λ		N_λ	
# Receiver modules	N_λ		N_λ	
# Serialization modules	N_λ		$2 \times N_\lambda$	
# Deserialization modules	N_λ		$2 \times N_\lambda$	
# Buffer width in (de)-serialization modules (Fig. 4.5)	PS/N_λ		$PS/(2 \times N_\lambda)$	
# Modulator drivers	N_λ	$2 \times N_\lambda$	N_λ	$2 \times N_\lambda$
# Total trans-impedance op-amps	N_λ		N_λ	
# Total comparator op-amps	N_λ		$3 \times N_\lambda$	
Energy-per-bit (EPB) and static power values (45nm SOI-CMOS)				
Total modulator driver EPB (pJ/bit)	$E^{Mod,OOK} \times N_\lambda$	$E^{Mod,OOK} \times 2N_\lambda$	$E^{Mod,EDAC} \times N_\lambda$	$E^{Mod,ODAC} \times 2N_\lambda$
Total serialization + deserialization EPB (pJ/bit)	$E^{SerDes} \times N_\lambda$		$E^{SerDes} \times 2 \times N_\lambda$	
Total comparator op-amps EPB (pJ/bit)	$E^{CO-OPAMP} \times N_\lambda$		$E^{CO-OPAMP} \times 3N_\lambda$	
Total trans-impedance op-amps EPB	$E^{TI-OPAMP} \times N_\lambda$		$E^{TI-OPAMP} \times N_\lambda$	
Power of MR tuning control circuit (μW)	$P^{TC} \times 2N_\lambda$	$P^{TC} \times 3N_\lambda$	$P^{TC} \times 2N_\lambda$	$P^{TC} \times 2N_\lambda$
Microheater power ($\mu W/nm$)	$P_{\mu heater} \times 2N_\lambda$	$P_{\mu heater} \times 3N_\lambda$	$P_{\mu heater} \times 2N_\lambda$	$P_{\mu heater} \times 2N_\lambda$

4.2). From Table 4.2, a typical OOK modulator driver consumes EPB of $E^{Mod,OOK} = 0.13$ pJ/bit [290], a typical serialization and deserialization module consumes EPB of $E^{SerDes} = 0.5$ pJ/bit [167], and a typical trans-impedance op-amp consumes EPB of $E^{TI-OPAMP} = 0.21$ pJ/bit [290]. As a result, an OOK link with N_λ channels consumes modulator driver EPB of $(E^{Mod,OOK} \times N_\lambda)$ pJ/bit, serialization + deserialization EPB of $(E^{SerDes} \times N_\lambda)$ pJ/bit, and trans-impedance op-amps EPB of $(E^{TI-OPAMP} \times N_\lambda)$ pJ/bit, as the link has N_λ counts of modulator drivers, serialization modules, deserialization modules, and trans-impedance op-amps each. Further, from Table 4.2, the tuning control circuit and the integrated microheater of an MR consume $P^{TC} = 385\mu W$ [218] and $P_{\mu heater} = 800\mu W/nm$ power, respectively. Therefore, the OOK link consumes $(P^{TC} \times 2 \times N_\lambda)$ μW power for the MR tuning control circuits and $(P_{\mu heater} \times 2 \times N_\lambda)$ $\mu W/nm$ power in the MR-integrated microheaters, as the link has $2 \times N_\lambda$ MRs (N_λ modulators + N_λ filters).



(a)



(b)

Figure 4.6: Schematic illustration of (a) an OOK modulation-based optical link, and (b) a 4-PAM modulation-based optical link, with total four wavelengths (λ_1 to λ_4). Note that having equal number of optical signals results in $2\times$ datarate for the 4-PAM link. In other words, equal datarate can be achieved for 4-PAM links by using $2\times$ less optical signals.

Table 4.2: Sample values of per-instance EPB for modulator driver (E^{Mod}), serialization + deserialization (E^{SerDes}), trans-impedance op-amps ($E^{TI-OPAMP}$) and per-MR static power for MR tuning control circuit (P^{TC}) and microheater ($P^{\mu heater}$).

Parameter	Signaling Method		
	OOK (pJ/bit)	4-PAM	
		EDAC (pJ/bit)	ODAC (pJ/bit)
E^{Mod}	0.13 [275]	3.04 [18]	0.04 [22]
E^{SerDes}	0.5 pJ/bit [248]		
$E^{CO-OPAMP}$	0.21 pJ/bit [275]		
$E^{TI-OPAMP}$	0.24 pJ/bit [183]		
P^{TC}	385 μ W [20]		
$P^{\mu heater}$	800 μ W/nm [147]		

DWDM Links using 4-PAM Signaling and Various 4-PAM Modulators

Fig. 4.6(b) shows a schematic of a 4-PAM signaling based DWDM link with four optical channels ($N_\lambda = 4$). It is evident from the figure that, compared to an OOK link, a 4-PAM link with $N_\lambda = 4$ requires $2\times$ more serialization and deserialization modules. Moreover, a 4-PAM receiver requires $3\times$ more trans-impedance op-amps (Fig. 4.5). Therefore, it can be generalized that contrary to an OOK link, a 4-PAM link with N_λ channels requires $2\times N_\lambda$ serialization modules, $2\times N_\lambda$ deserialization modules, $3\times N_\lambda$ trans-impedance op-amps based receiver modules (Table 4.1). Moreover, a 4-PAM link requires (Packet Size/($2\times N_\lambda$)) bits of buffer width at their E/O and O/E interfaces (Table 4.1), as the number of bits per symbol $B=2$ for a 4-PAM link in Figs. 4.5(c) and 4.5(d). On the other hand, like an OOK link, the 4-PAM link also requires N_λ counts of MR filters, photodetectors, and receiver modules each. Corresponding to these hardware component counts, a 4-PAM link with N_λ channels consumes a total serialization and deserialization EPB of ($E^{SerDes} \times 2 \times N_\lambda$) pJ/bit and total trans-impedance op-amps EPB of ($E^{TI-OPAMP} \times 3 \times N_\lambda$) pJ/bit (Table 4.1). In addition, the counts and overheads of other hardware components for a 4-PAM link, such as MR modulators and modulator drivers, depend on the specific 4-PAM modulator type, as discussed next.

4-PAM EDAC Modulator Based Links

A 4-PAM-EDAC modulator-based link with N_λ channels requires a total of N_λ MR modulators, which makes the total MRs per link to be $2\times N_\lambda$ (N_λ filters and N_λ modulators). Moreover, the link requires N_λ electrical DAC (EDAC) based modulator drivers (one EDAC per modulator as shown in Fig. 4.3), each of which consumes $E_{Mod,EDAC} = 3.04$ pJ/bit EPB [208]. Therefore, a 4-PAM-EDAC link with N_λ channels consumes modulator driver EPB of ($E^{Mod,EDAC} = 3.04 \times N_\lambda$) pJ/bit, power for MR tuning control circuits of ($P^{TC} \times 2 \times N_\lambda$) μ W, and MR microheater power

of $(P^{\mu heater} \times 2 \times N_\lambda) \mu\text{W}/\text{nm}$ (Table 4.1).

4-PAM ODAC Modulator Based Links

A 4-PAM-ODAC modulator-based link with N_λ channels requires a total of N_λ spoked MR modulators (Fig. 4.4), which makes the total number of MRs per link to be $2 \times N_\lambda$ (N_λ filters and N_λ modulators). However, unlike a 4-PAM-EDAC link, a 4-PAM-ODAC link requires $2 \times N_\lambda$ modulator drivers (2 drivers per modulator; Fig. 4.4), each of which consumes $E^{Mod,ODAC} = 0.04$ pJ/bit EPB [167]. Therefore, a 4-PAM-ODAC link with N_λ channels consumes modulator driver EPB of $(E^{Mod,ODAC} = 0.04 \times 2 \times N_\lambda)$ pJ/bit, power for MR tuning control circuits of $(P^{TC} \times 2 \times N_\lambda) \mu\text{W}$, and MR microheater power of $(P^{\mu heater} \times 2 \times N_\lambda) \mu\text{W}/\text{nm}$ (Table 4.1).

4-PAM-SS Modulator Based Links

A 4-PAM-SS modulator-based link with N_λ channels requires $2 \times N_\lambda$ MR modulators (2 modulators per channel; Fig. 4.2), which makes the total number of MRs per link to be $3 \times N_\lambda$ (N_λ filters + $2 \times N_\lambda$ modulators). Moreover, a 4-PAM-SS link requires $2 \times N_\lambda$ modulator drivers (1 driver per modulator; Fig. 4.2), each of which utilizes OOK signaling and consumes $E^{Mod,OOK} = 0.13$ pJ/bit EPB [290]. Therefore, a 4-PAM-SS link with N_λ channels consumes in total modulator driver EPB of $(E^{Mod,OOK} \times 2 \times N_\lambda)$ pJ/bit, power for MR tuning control circuits of $(P^{TC} \times 3 \times N_\lambda) \mu\text{W}$, and MR microheater power of $(P^{\mu heater} \times 3 \times N_\lambda) \mu\text{W}/\text{nm}$ (Table 4.1).

In addition, a 4-PAM-SS link also suffers from a high signal power loss in 4-PAM modulators due to the possible inter-channel crosstalk. Ideally, in a 4-PAM-SS modulator, when two OOK-modulated signals are superposed (Fig. 4.2), a 4-PAM modulated signal is generated owing to the constructive interference between the two OOK signals. However, the constructive interference happens only if both OOK signals are in phase. Unfortunately, in the presence of non-idealities such as fabrication process and on-chip temperature variations, a significant phase difference may exist between the two superposed OOK signals, which can lead to destructive interference between them. Owing to the random nature of fabrication-process variations, this incurred phase difference may fall anywhere in the range from 0 to 2π . This implies that the degree of destructive interference incurred between the OOK signals due to the phase difference (and hence the amplitude levels of the symbols of the resultant 4-PAM signal) may fall anywhere in a very large range of values. This in turn makes it very hard to ensure reliability of communication with a 4-PAM-SS modulator. We evaluate the adverse impact of random fabrication-process variations on the reliability of 4-PAM-SS links in terms of the worst-case destructive interference, as explained next.

The worst-case destructive interference in a 4-PAM-SS modulator occurs when the two superposed OOK signals are completely out of phase, i.e., when the phase difference between them is an odd multiple of π . The amount of signal loss due to the superposition of two out-of-phase OOK signals depends on their individual signal intensities. Typically, in a 4-PAM-SS modulator (Fig. 4.2), to equidistantly

space the four amplitude levels of the output 4-PAM symbol in the available range of optical transmission, the intensities of the individual OOK signals are kept at two-third and one-third of the intensity of the conventional OOK signal. Hence, for the best-case constructive interference between the superposed OOK signals, the intensity of the resultant 4-PAM signal becomes $2/3 + 1/3 = 1$. In contrast, for the worst-case destructive interference, the intensity of the resultant 4-PAM signal becomes $2/3 - 1/3 = 1/3$, causing the worst-case interference-related signal loss to be $-10 \times \log(1/3) = 4.8\text{dB}$. This interference-related signal loss in 4-PAM-SS modulators reduces the signal-to-noise ratio (SNR) and increases the bit-error rate (BER). We have considered this worst-case interference-related signal loss of 4.8dB in our trade-off analysis for 4-PAM-SS links. We have also considered the best-case scenario for which this interference-related loss is omitted for our analysis of 4-PAM-SS links. This in turn reduces the overall communication reliability, adversely affecting the trade-offs among the energy, reliability, and performance of 4-PAM-SS links.

In summary, the hardware overhead of implementing an OOK or 4-PAM signaling based photonic link depends not only on the choice of modulator design and signaling method but also on the number of parallel wavelength channels N_λ in the link (see Table 4.1). The maximum supportable N_λ for an OOK or 4-PAM signaling based photonic link is determined based on the inherent tradeoffs among the energy consumption, reliability (BER), and performance of the designed OOK or 4-PAM link, as discussed in the next section.

4.4 Design Tradeoffs For Photonic Links

Designing a photonic link is subject to inherent tradeoffs among the achievable performance (aggregated datarate), energy consumption, and reliability (BER). Optimizing these design tradeoffs often involves finding the balance between the link's aggregated datarate and energy-reliability behavior. From [58], [249], [251], and [20], for optimizing the design of a photonic link, optical power budget (P_{dB}^B) of the link is the most critical design constraint. It is calculated in dB as the difference between the maximum allowable optical power (P_{Max}) and detector sensitivity (S), as shown in Eq. 4.1. For a photonic link, P_{Max} identifies the ceiling of P_{dB}^B and ensures that the total power of all the DWDM signals (i.e., total N_λ signals) propagating through the link remains below the maximum allowable level which is limited by various non-linear effects of silicon in constituent devices [20]. On the other hand, S is the noise-limited floor of the link's P_{dB}^B and ensures that the individual signals propagating through the link reach the receiver without dropping below the minimum power level defined by S . For a photonic link, the total optical power allocated within its P_{dB}^B supports two causes, as shown in Eq. 4.2. First, it compensates for the total optical power penalty (PP_{dB}) of the link. Second, it supports total N_λ DWDM wavelength signals/channels in the link. Therefore, from Eq. 4.2, improving the performance (aggregated datarate) of a photonic link that has a fixed P_{dB}^B by increasing the supported N_λ in the link requires a corresponding decrease in the link's PP_{dB} .

$$P_{dB}^B = P_{Max} - S \tag{4.1}$$

$$P_{dB}^B \geq PP_{dB} + 10 \log_{10} (N_\lambda) \quad (4.2)$$

For optimal link design, this tradeoff between N_λ and PP_{dB} is affected by the following four factors: (i) The photodetector sensitivity S , which is the minimum detectable optical power in dBm, and depends on the baud-rate (i.e., number of amplitude/level transitions in unit time) of the individual photonic signals [20]. (ii) The cyclic dependency between PP_{dB} and achievable aggregated data rate, which is given as $N_\lambda \times \text{bitrate (BR)}$ of photonic channels. This cyclic dependency means that PP_{dB} depends on $N_\lambda \times \text{BR}$ through the available PB_{dB} , whereas $N_\lambda \times \text{BR}$ in turn depends on PP_{dB} . (iii) Several spectral parameters of MRs such as the free-spectral-range (FSR) and full-width-at-half-maximum (FWHM) bandwidths, which impact the effective value of PP_{dB} . (iv) The ultimate design goals of photonic links and PNoCs, including the goals of maximizing aggregated data rate, energy-efficiency, and/or achieving the desired BER, which also impact the effective value of PP_{dB} . In the next subsection, we systematically analyze and provide detailed models for all four aforementioned factors that affect photonic link design trade-offs, with respect to the utilized modulator types and signaling methods.

Factors that Affect Link Design Tradeoffs

Baud-Rate Dependent Detector Sensitivity (S)

From [20], detector sensitivity (S) in dBm increases with increase in signal baud-rate. Signal baud-rate is defined as the number of amplitude/level transitions in the photonic signal occurring in unit time. We consider the baseline value of $S = -22\text{dBm}$ at 10Gb/s [20] for both 4-PAM and OOK links, and adopt the model from [250] to capture how S would increase for baud-rates greater than 10Gb/s . We extract S for 4-PAM signals based on the experimentally demonstrated and validated models from [20]. From [245], it is evident that a 4-PAM signal requires 3.3dB more received power compared to an OOK signal of the same bitrate (BR), to achieve the same bit-error rate (BER) as achieved by the OOK signal. Therefore, to derive S for 4-PAM signals, we simply add 3.3 dB to the S that is obtained for OOK signals of the same BR using the BR-dependent model of S from [245]. The same value of baud-rate translates into $2 \times \text{bitrate}$ for a 4-PAM link compared to an OOK link, and for evaluating link performance, bitrate (i.e., aggregated data rate) is a more useful metric than baud-rate. Therefore, we use the following Eq. 4.3 as the relation between baud-rate and bitrate, henceforth.

$$BaR = BR / (M/2) \quad (4.3)$$

Here, BaR is baud-rate, BR is bitrate, and M is number of amplitude levels used in the signal to represent a symbol ($M=2$ for an OOK signal, and $M=4$ for a 4-PAM signal).

Cyclic Dependency Between PP_{dB} and Aggregate Datarate ($N_\lambda \times BR$)

To understand the cyclic dependency between PP_{dB} and aggregated datarate ($N_\lambda \times BR$), it is important to understand what constitutes PP_{dB} (i.e., the total optical power penalty of the link) and how it changes between OOK and 4-PAM links. For a link, PP_{dB} is comprised of the total penalty of the MR filter array (PP_{Fil}), MR modulator array crosstalk penalty (PP_{Mod}), PAM signaling penalty (PP_{PAM}), and various optical signal power losses such as waveguide propagation loss (P_L^{WGP}), waveguide bending loss (P_L^{WGB}), through loss of active MRs (P_L^{MR-Act}), through loss of inactive MRs (P_L^{In-Act}), worst-case signal interference penalty (PP_{INTRF}), and splitter/coupler loss (P_L^{SpC}). In [250], [19], and [18], PP_{Fil} and PP_{Mod} are analytically modeled, considering the general case of a bank of N_λ modulator MRs employed at the sender node and a bank of N_λ filter MRs employed at the receiver node of a link with N_λ DWDM signals. Accordingly, PP_{Mod} for an MR modulator in the bank of N_λ modulator MRs can be evaluated using Eq. 4.4 and Eq. 4.5 [17]. Similarly, PP_{Fil} for the i th MR in the bank of N_λ MR filters can be given as Eq. 4.6 [18], where formulas for some important terms in Eq. 4.6 are given in Eqs. 4.7-4.9 [18]. Here, Eq. 4.7 considers the total crosstalk contribution from all N_λ wavelength signals combined. The definitions and typical values (if any) of various terms used in Eqs. 4.1-4.9 are given in Table 4.3 and Table 4.4, respectively.

$$PP^{Mod} = -5 * \log_{10} \left(\frac{\left(\frac{2K}{FWHM}\right)^2 + q_0}{\left(\frac{2K}{FWHM}\right)^2 + 1} \right) \quad (4.4)$$

$$K = \begin{cases} f_\Delta - \Delta f, f_\Delta > 0 \\ f_\Delta, f_\Delta < 0 \end{cases} \quad (4.5)$$

$$(PP^{Fil})_{ithMR} = \left(-10 * \log_{10} \left(1 - 0.5 * Q_{BER} * \frac{P_{Xtalk}}{P_{NRZ}^{av}} * \frac{r+1}{r-1} \right) \right) \quad (4.6)$$

$$\left(\frac{P_{Xtalk}}{P_{NRZ}^{av}} \right)_{ithMR} = \sum_{j=1, j \neq i}^{N_\lambda} \Gamma_{i,j} \quad (4.7)$$

$$\Gamma_{i,j} = \left(\int_{-\infty}^{+\infty} \frac{\text{sinc}^2(F) dF}{1 + \left(\frac{F+(j-i)F_\Delta}{\xi_i}\right)^2} * \left[\prod_{k=1}^{i-1} \frac{\left(\frac{F+(j-k)F_\Delta}{\xi_k}\right)^2}{1 + \left(\frac{F+(j-k)F_\Delta}{\xi_k}\right)^2} \right] \right) \quad (4.8)$$

$$(j-i)F_\Delta = \left(\frac{v_{Si}}{\lambda_i \times \text{BaR}} - \frac{v_{Si}}{\lambda_i + \left((j-i) \times \frac{\text{FSR}(\text{nm})}{N_\lambda + 1} \right) \times \text{BaR}} \right) \quad (4.9)$$

Table 4.3: Definitions of various link design parameters and notations from Eqs. 4.1 - 4.15.

Parameter	Definition
P_{Max}	Max. allowable optical power in waveguide (dBm) [141, 250]
S	Detector sensitivity at 10Gb/s [250] (dBm)
P_L^{WGB}	Waveguide bending loss (dB per 90°) [250]
v_{Si}	Speed of light at 1550 nm in silicon (in m/s)
Q_{BER}	Signal Q-parameter for BER = 10^{-9} [239]
q_0	Extinction ratio of MR modulator in the OFF state [298]
P_L^C	Coupler loss (dB) [100]
P_L^{Sp}	Splitter Loss (dB)
P_L^{WGP}	Propagation loss (dB), @ 1dB/cm [18]
PP^{INTRF}	Worst-case signal interference penalty (dB)
PP^{PAM}	4-PAM signaling penalty (dB)
r	Extinction ratio of modulation
$FWHM$	3dB Bandwidth of an MR modulator (GHz)
PP^{ER}	Penalty (dB) due to the finite r (see above) [298]
P_{dB}^B	Photonic link power budget in dB
PP_{dB}	Total power penalty for the photonic link in dB
BR	Bitrate of a photonic signal
BaR	Baud-rate (# of level transitions per unit time) of a signal
M	# amplitude levels per symbol in a photonic signal
PP^{Mod}	Crosstalk power penalty in a modulator MR (in dB)
f_Δ	Frequency spacing between two adjacent photonic signals
Δf	Frequency spacing between an MR modulator's OFF-state and ON-state resonances
P_{NRZ}^{av}	Average power per incoming photonic signal at the i^{th} MR filter in an MR filter bank
P_{xtalk}^{av}	Cumulative crosstalk power from all N_λ signals combined at the i^{th} MR filter of an MR filter bank
PP^{Fil}	Crosstalk power penalty (in dB) at the i^{th} MR filter
$\Gamma_{i,j}$	Fraction of crosstalk power from j^{th} signal dropped at the i^{th} MR filter in an MR filter bank
F	Photonic signal frequency normalized to baud-rate (BaR)
$(j-i)F_\Delta$	Frequency spacing between the i^{th} MR filter resonance and j^{th} signal normalized to baud-rate (BaR)
	FWHM of i^{th} MR filter normalized to BaR
λ_i	Resonance wavelength of i^{th} MR filter
FSR	Free-spectral range
N_λ	Number of photonic signals DWDM in a waveguide
P_L^{MR-Act}	Through loss of an active MR
$P_L^{MR-InAct}$	Through loss of an inactive MR

Table 4.3: Definitions of various link design parameters and notations from Eqs. 4.1 - 4.15. (Continued)

P_{dB}^{BERO}	Power penalty (dB) for reliability optimal design of a link
$P_{dB}^{DR-BERBal}$	Power penalty (dB) for datarate-reliability balanced link design

Although Eqs. 4.4 to 4.9 were originally developed in [239] and [298] for OOK links, this same set of equations can be used to determine PP_{Mod} and PP_{Fil} for 4-PAM links as well. This is because OOK signals and 4-PAM signals have similar frequency spectra (i.e., in the shape of the sinc function). As a result, the utilized equations can be transformed to be based on signal baud-rate (BaR) instead of bitrate (BR), because the crosstalk at the modulators and filters can be assumed to have a Gaussian distribution, as demonstrated in [239]. Further, the parameters FWHM, QBER, and r in Eqs. 4.4 and 4.6 assume different values for OOK and 4-PAM signaling types (as shown in Table 4.4). Moreover, as a 4-PAM signal has $3\times$ less separation between its amplitude levels compared to an OOK signal, a 4-PAM signal requires ~ 3.3 dB more power at the receiver [248], compared to an OOK signal of the same BaR, to achieve the same bit-error rate (BER) of 10^{-9} . This extra required power is accounted for as PP^{PAM} (Table 4.4) in the total PP_{dB} value. Moreover, note that to evaluate PP_{Mod} for 4-PAM-SS links, we treat the 2 MR modulators required per wavelength signal as a single modulator unit (constituting a bank of N_λ MR modulator units at the sender node), and PP_{Mod} is evaluated for each MR modulator unit instead of each individual MR modulator. Also, using 8-PAM/16-PAM signals in the links can certainly reduce the hardware requirement for the links compared to 4-PAM signals, if the target aggregate datarate remains unchanged. This in turn can result in higher dynamic energy efficiency for 8-PAM/16-PAM links. However, PP_{PAM} for 8-PAM and 16-PAM photonic links would increase to 6.1 dB and 8.75 dB respectively [245], compared to PP_{PAM} of 3.3 dB for 4-PAM links (Table 4.4), due to the larger values of M for 8-PAM/16-PAM links ($M = 8, 16$ for 8-PAM, 16-PAM links respectively [245]). Larger PP_{PAM} would increase the overall penalty PP_{dB} for 8-PAM/16-PAM links, which in turn can render lower N_λ , lower BR, and hence, lower aggregate datarate to 8-PAM/16-PAM links, compared to 4-PAM links. This reduced aggregate datarate can offset the benefits obtained from achieving higher dynamic energy efficiency.

The values of PP_{Mod} and PP_{Fil} (as evaluated from Eqs. 4.4 to 4.9), along with PP_{PAM} , contribute to PP_{dB} , as shown in Eq. 4.10. Eq. 4.10 also has some other terms related to the optical signal power loss. The definitions and typical values (if any) of all these loss terms from Eq. 4.10, except P_L^{MR-Act} and $P_L^{MR-InAct}$ (which are discussed in the next paragraph), are also given in Table 4.3 and Table 4.4 respectively. Note that the values from Table 4.4 for some of the terms in Eq. 4.10 depend on the underlying signaling/modulator type and/or PNoC architecture. For example, from Eq. 4.10 and Table 4.4, PP_{INTRF} is zero for 4-PAM-ODAC, 4-PAM-EDAC, and OOK links, whereas it is 4.8dB for 4-PAM-SS links. This is because only 4-PAM-SS modulators incur signal superposition induced interference loss. Similarly, the values of r , FWHM, and PP_{ER} also change between different modulator/signaling types (Table 4.4). Moreover, the values of P_L^{SpC} and P_L^{WGP} depend on the underly-

Table 4.4: Typical values (if any) of various link design parameters and notations from Eqs. 4.1-4.15.

Parameter	Value			
P_{Max}	20 dBm [141, 250]			
S	-22.5 dBm [250] @ 10 Gb/s BaR			
P_L^{WGB}	0.005 dB per 90° [248]			
v_{Si}	8.6×10^7 m/s			
q_0	0.04 [298]			
P_L^C	0.9 dB [100]			
PNoC Architectures				
	CLOS	SWIFT		
P_L^{Sp}	5.6 dB [41]	1.2 dB [161]		
P_L^{WGP} at 1 dB/cm [19]	4.5 dB (4.5 cm long link) [188]	12 dB (12 cm long link) [161]		
Signaling Methods				
	OOK	SS	EDAC	ODAC
PP^{INTRF}	0 dB	4.8 dB	0 dB	0 dB
PP^{PAM}	0 dB	1.76 dB [248]		
r	5 dB [248]	5 dB [248]	5 dB [248]	2 dB [248]
$FWHM$	30 GHz	45 GHz [209]	18 GHz [18]	36 GHz [248]
PP^{ER}	4.2 dB [298]	4.2 dB [298]	4.2 dB [298]	7.7 dB [298]
Q_{BER}	6 dB [239]	12.5 dB [239], [188]		

ing PNoC architecture—we use CLOS [116] and SWIFT PNoC [56] architectures in this chapter—as the required count of splitters and waveguide lengths differ between CLOS and SWIFT PNoCs.

$$\begin{aligned}
 PP_{dB} = & P_L^{MR-Act} + P_L^{MR-InAct} + P_L^{WGP} + P_L^{Sp} + P_L^{INTRF} + PP^{Mod} + PP^{Fil} + \\
 & PP^{PAM} + P_L^{WGB} + P_L^C + PP^{ER}
 \end{aligned} \tag{4.10}$$

In addition, in Eq. 4.10, P_L^{MR-Act} refers to the total through loss per λ signal of an active MR bank, whereas $P_L^{MR-InAct}$ refers to the total through loss per λ signal of an inactive MR bank. We define an MR bank as an active MR bank if it actively operates on its assigned optical signals. Therefore, the resonances of the MRs of an active MR bank are typically locked to their respective optical signals with exact spectral matching. As a result, in an active MR bank, the spacing between the i th MR resonance and j th λ signal remains $\text{mod } j - i \times \text{one channel spacing}$ (i.e., spacing

between two adjacent wavelength signals). In contrast, an inactive MR bank is an MR bank that is temporarily turned off to not operate on its assigned optical signals. Typically, to turn off the constituent MRs of an inactive MR bank, their resonances are locked at spectral locations that are about half the channel spacing away from their assigned optical signals. Therefore, in an inactive MR bank, the spacing between the i th MR resonance and j th λ signal remains $\text{mod } j - i + 0.5 \times \text{one channel spacing}$. We exploit this difference in the spectral lock positioning between active MR banks and inactive MR banks to model both P_L^{MR-Act} and $P_L^{MR-InAct}$ using a common model given by Eq. 4.11. In Eq. 4.11, $\Gamma_{i,j}$ can be evaluated using Eq. 4.8, wherein the model for F_Δ changes between F_Δ^{MR-Act} (Eq. 4.12) and $F_\Delta^{MR-InAct}$ (Eq. 4.13) depending on whether P_L^{MR-Act} or $P_L^{MR-InAct}$ is being evaluated. Note that in Eqs. 4.9, 4.12, 4.13, $(FSR/(N_\lambda+1))$ equals one channel spacing between two adjacent wavelength signals with FSR being the free-spectral range.

$$P_{L,jth\lambda}^{MR} = -10 \log_{10} \left(\sum_{i=1, j \neq i}^{N_\lambda} \Gamma_{i,j} \right) \quad (4.11)$$

$$(j - i)F_\Delta^{MR-Act} = \left(\frac{v_{Si}}{\lambda_i \times BaR} - \frac{v_{Si}}{\lambda_i + \left((j - i) \times \frac{FSR(nm)}{N_\lambda+1} \right) \times BaR} \right) \quad (4.12)$$

$$(j - i)F_\Delta^{MR-InAct} = \left(\frac{v_{Si}}{\lambda_i \times BaR} - \frac{v_{Si}}{\lambda_i + \left((j - i + 0.5) \times \frac{FSR(nm)}{N_\lambda+1} \right) \times BaR} \right) \quad (4.13)$$

It is evident from Eqs. 4.1-4.13 that there is a cyclic dependency between PP_{dB} and N_λ as the achievable N_λ for a link depends on PP_{dB} from Eq. 4.1, whereas PP_{dB} in turn is determined based on the combination of N_λ and bit-rate (BR) from BaR in F_Δ in Eq. 4.4. This cyclic dependency makes it difficult to find out the optimal combination of N_λ and BR that can be supported by a link. To mitigate this problem, we employ a heuristic-based search approach that finds out the optimal combination of N_λ and BR.

Dependence of PP_{dB} on MRs Spectral Parameters

In Eq. 4.10, PP_{dB} depends on the MRs' spectral parameters such as FWHM and FSR. Parameters FWHM and FSR are defined in Table 4.3. These spectral parameters depend on the device dimensions that are utilized for implementing the MRs of the photonic links. We select different FWHM values for MR modulators and filters based on the utilized modulator/signaling type, as shown in Table 4.4. Moreover, for FSR considerations, we select a viable FSR value of 20 nm from prior work [202] for our analysis in this chapter. Our design methodology, analysis, and related link-level and system-level evaluation results are discussed in the upcoming subsections.

Dependence of PP_{dB} on Design Goals

Whether or not to consider PP_{Fil} , PP_{Mod} , and PP_{INTRF} in Eq. 4.10 to evaluate the effective value of PP_{dB} depends on whether the goal is to design photonic links and PNoCs with maximum aggregated data rate or desired bit-error rate (BER). The emanation of crosstalk noise in modulator and filter MR banks reduces the signal-to-noise ratio (SNR) in photonic links (e.g., [57], [250]), which in turn increases the BER, degrading the reliability of photonic communication. To compensate for this degradation in BER, one way is to increase the input signal power by an appropriate amount. The required increase in the input signal power to achieve the unchanged BER in the presence of crosstalk noise is termed as power penalty. In Eq. 4.10, PP_{Fil} and PP_{Mod} correspond to the crosstalk noise induced power penalties for the filter MR bank and modulator MR bank, respectively. Similarly, PP_{INTRF} corresponds to the required increase in the input signal power (i.e., caused power penalty) to compensate for the worst-case destructive signal interference in 4-PAM-SS links. From Table 4.4, our considered models and resultant values of PP_{Fil} and PP_{Mod} correspond to a BER of 10^{-9} . We select BER of 10^{-9} for our analysis, because it is often considered acceptable for optical communication links [18], [245]. From this value of BER, we calculate Q_{BER} (defined in Table 4.4) using the models presented in [18]. From Table 4.4, the evaluated QBER differs between OOK and PAM4 signaling/modulation techniques. The presence of the PP_{Fil} , PP_{Mod} and PP_{INTRF} terms in the PP_{dB} model (Eq. 4.10) increases the value of PP_{dB} , which whittles down a large portion of the power budget P_{dB}^B (Eq. 4.2), leaving only a small portion of P_{dB}^B available for supporting N_λ (Eq. 4.2). This results in a small value of aggregated data rate ($N_\lambda \times BR$) for a given bitrate (BR). Nevertheless, this ensures that the BER remains unharmed at 10^{-9} . Therefore, if achieving the desired unharmed BER is the design goal, the terms PP_{Fil} , PP_{Mod} and PP_{INTRF} should be included in the model for PP_{dB} . For easy reference in the following sections of this chapter, we identify such BER-optimal PP_{dB} as PP_{dB}^{BERO} and provide its model in Eq. 4.14, which includes the PP_{Fil} , PP_{Mod} and PP_{INTRF} terms.

$$PP_{dB}^{BERO} = P_L^{MR-Act} + P_L^{MR-InAct} + P_L^{WGP} + P_L^{Sp} + P_L^{INTRF} + PP^{Mod} + PP^{Fil} + PP^{PAM} + P_L^{WGB} + P_L^C + PP^{ER} \quad (4.14)$$

Another way of compensating for the crosstalk-induced degradation in reliability (BER) is to use forward error correction (FEC) codes (e.g., [175], [148]). FEC codes add extra redundancy bits in every data packet to enable error detection and correction. The use of FEC codes in a photonic link can improve the BER of the link to be lower than 10^{-9} , especially if the crosstalk inflicted BER of the link is above the typical FEC limit (e.g., 1.2×10^{-3} for BCH code [175]). The use of redundant bits in FEC codes (we use the popular SECDED (72, 64) FEC [148] code in this chapter) increases the packet size, and hence, the packet transfer delay and energy. Nevertheless, it does not require an increased input signal power to compensate for crosstalk-induced bit-errors. Therefore, the use of FEC codes does not whittle down the link power budget, allowing for an opportunity to support greater N_λ and aggre-

gated datarate in addition to achieving the desired reliability (BER). In other words, the use of FEC codes enables datarate-balanced BER in photonic links. Hence, if achieving the datarate-balanced desired BER using FEC codes is the design goal, the terms PP_{Fil} , PP_{Mod} and PP_{INTRF} need not be included in the formula for PP_{dB} . We identify such datarate-BER balanced PP_{dB} as $PP_{dB}^{DR-BER-Bal}$ and provide its model in Eq. 4.15, which excludes the PP_{Fil} , PP_{Mod} and PP_{INTRF} terms.

$$PP_{dB}^{DR-BER-Bal} = P_L^{MR-Act} + P_L^{MR-InAct} + P_L^{WGP} + P_L^{Sp} + PP^{PAM} + P_L^{WGB} + P_L^C + PP^{ER} \quad (4.15)$$

These design goals (i.e., BER-optimal design versus datarate-BER balanced design) are considered, along with the baud-rate dependence of the detector sensitivity and the dependence of PP_{dB} on MRs' spectral parameters, in our search-heuristic based optimization approach for photonic link designs, as discussed next.

Heuristic-Based Search For the Efficient Design of Photonic Links

Irrespective of whether the designed photonic link is BER-optimal or datarate-BER balanced, the achievable aggregated datarate (i.e., $N_\lambda \times \text{BR}$) has a cyclic dependency on the P_{dB}^B and PP_{dB} parameters of the link, which makes it difficult to obtain an optimal value of $N_\lambda \times \text{BR}$ for the link directly using Eqs. 4.1-4.15. To break this cyclic dependency and determine the optimal combination of N_λ and BR for the designed link, we employ a heuristic-based search optimization framework. The basic idea of our framework is to perform exhaustive search for the optimal combination of N_λ and BR for which the available power budget of the link (P_{dB}^B in Eq. 4.1) is fully utilized, while considering the factors that affect the photonic link design trade-offs.

We provide a set of baud-rate (BaR) and N_λ duplets as one of the inputs to our search heuristic. We use BaR instead of BR as input because the modeling equations directly depend on BaR, which can be easily converted into BR after our search using Eq. 4.3. Moreover, to limit the cost and complexity of the comb-generating laser source [48], and to be consistent with the prior works on 4-PAM optical signaling [120] and [217], we limit the maximum allowable value of N_λ to 128. Moreover, as the flit-size of a PNoC is directly proportional to the value of N_λ , and as the flit-size is usually a power-of-two value, the allowable values of N_λ should also be power-of-two values. Because of these reasons, we choose a set Λ of all allowable values of N_λ , where $\Lambda = N_\lambda \text{ --- } N_\lambda \in 128, 64, 32, 16, 8, 4, 2, 1$. Moreover, we define the set of all possible baud-rate values $R = \text{BaR} \text{ --- } \text{BaR} \in \mathbb{Q}^+$; BaR is in Gb/s; $10 \text{ Gb/s} \leq \text{BaR} \leq 30 \text{ Gb/s}$; $(\text{BaR}/0.5) \in \mathbb{N}$, which has 41 elements. The individual values for Λ and BaR combine to make a duplet in $41 \times 8 = 328$ different ways. We create a set Y of these duplets, $Y = (N_{\lambda 1}, \text{BaR}1), (N_{\lambda 2}, \text{BaR}2), \dots, (N_{\lambda 8}, \text{BaR}41)$, and give it as an input to our search heuristic. Based on the constraint in Eq. 4.2, we utilize an error function $\text{ef}(N_\lambda, \text{BaR})$ given in Eq. 4.16 to find the optimal duplet from set Y . For that, for each element (N_λ, BaR) of the set Y , we evaluate an error value $\epsilon = \text{ef}(N_\lambda, \text{BaR})$ and create a set E of all ϵ values. All (N_λ, BaR) duplets corresponding to the

positive ϵ values in set E satisfy the constraint given in Eq. 4.2. But we choose the (N_λ, BaR) duplet corresponding to the minimum positive value ϵ_{min} from set E as the optimal value, because such a duplet fully utilizes the link P_{dB}^B .

$$ef(N_\lambda, \text{BaR}) = \{P_{dB}^B - PP_{dB} - 10 \log_{10}(N_\lambda)\} \quad (4.16)$$

In Eq. 4.16, we evaluate PP_{dB} as a function of the (N_λ, BaR) duplet. We use the search heuristic to find one (N_λ, BaR) duplet for every type (i.e., OOK, 4-PAM-EDAC, 4-PAM-SS, and 4-PAM-ODAC) of photonic link. We use Eqs. 4.14 and 4.15, respectively, as the models for PP_{dB} in the error function for our search of (N_λ, BaR) duplets. To evaluate P_L^{WGP} term from Eqs. 4.14 and 4.15, we consider the maximum link length in our considered PNoCs, which is 4.5 cm for CLOS PNoC [116] and 12 cm for SWIFT PNoC [56], as provided in Table 4.4. The term P_L^{SP} from Eqs. 4.14 and 4.15 is evaluated based on the number of splitters employed by the PNoC to power its waveguides, which differs between CLOS and SWIFT PNoCs. For example, CLOS PNoC has 56 point-to-point waveguides, and to power these waveguides, the PNoC employs 1×2 , 1×7 , 1×4 splitters in series [252] [71]. Therefore, the input optical power is split in 56 parts in the CLOS PNoC. Because we consider per-split loss to be 0.1 dB, the total splitter loss P_L^{SP} in the CLOS PNoC is 5.6 dB, as shown in Table 4.4. Similarly, total splitter loss (P_L^{SP}) in the SWIFT PNoC is 1.2 dB which is also provided in Table 4.4. Note that the error-function (Eq. 4.16) evaluation differs between datarate-BER balanced and BER-optimal photonic link designs, as discussed next.

Design of Datarate-BER Balanced Photonic Links

In order to design photonic links to achieve datarate-balanced BER, we do not add the modulator penalty (PP_{Mod}), filter penalty (PP_{Fil}) and signal interference penalty (PP_{INTRF}) terms to the total PP_{dB} in Eq. 4.15. Because PP_{Mod} and PP_{Fil} model crosstalk penalty when crosstalk-induced increase in BER is mitigated by increasing input optical signal power. Instead, we use SECDED (72, 64) FEC [148] code to counter the crosstalk-induced degradation in BER. Using the FEC code enables the photonic links to achieve higher aggregate data rate while maintaining the BER at 10^{-9} , thereby enabling a datarate-balanced BER value for the links. To find the optimal datarate-BER balanced (N_λ, BaR) duplet for a given signaling/modulation method based photonic link, we use Eq. 4.15 as the model for PP_{dB} in the error function given in Eq. 4.16.

Design of BER-Optimal Photonic Links

To design BER-optimal photonic links, the modulator (PP_{Mod}), filter (PP_{Fil}) penalty terms and interference-related signal loss (PP_{INTRF}) are included in total PP_{dB} (Eq. 4.14). Including these terms in PP_{dB} in Eq. 4.14 results in a low aggregate datarate but the BER remains unscathed. To find the BER-optimal (N_λ, BaR) duplet for a given signaling/modulation method based photonic link, we use Eq. 4.14 as the model for PP_{dB} in the error function given in Eq. 4.16. We repeat this exercise of

finding the BER-optimal and datarate-BER balanced (N_λ , BaR) duplets for every CLOS and SWIFT link type (corresponding to the signaling/modulation type) for 20 nm FSR [202]. Note that our search heuristic is equitably applicable to the OOK, 4-PAM-SS, 4-PAM-EDAC and 4-PAM-ODAC links. However, the optimal (N_λ , BaR) duplets would differ for different link types, as the values of P_{dB}^B and other design parameters differ for different link types.

Results of Optimal Designs of Photonic Links using Heuristic-Based Search

In this section, we present our obtained BER-optimal and datarate-BER balanced (N_λ , BaR) duplets for different variants of the CLOS and SWIFT links (i.e., 4.5 cm long links for CLOS PNoC [116] and 12 cm long links for SWIFT PNoC [56], with various modulation methods) for an FSR value of 20 nm. We also report our evaluated aggregated data rate ($N_\lambda \times \text{BR}$) and PBdB values for different variants of CLOS and SWIFT links. To evaluate aggregated datarate, we use Eq. 4.3 to convert the BaR values found through our search heuristic in the corresponding BR values.

Results for Datarate-BER Balanced Links

Table 4.5 gives optimal N_λ , bitrate (BR) (evaluated from BaR), aggregated datarate ($N_\lambda \times \text{BR}$), and P_{dB}^B values for different datarate-BER balanced variants of CLOS and SWIFT links. It also gives $PP_{dB} + 10\log(N_\lambda)$ values for our considered link variants. For brevity, we do not provide PP_{dB} values, but these values can be easily derived from $PP_{dB} + 10\log(N_\lambda)$ values as N_λ values are already provided in Table 4.5.

From Table 4.5, the aggregated datarate values for various SWIFT links are in general lower than the aggregated datarate values for various CLOS links. This is because links in SWIFT have greater (P_L^{SP}), P_L^C , and P_L^{WGP} values in Eq. 4.15 than the CLOS links (Table 4.4), which results in larger PP_{dB} values for the SWIFT links. As a result, a relatively smaller portion P_{dB}^B is available in Eq. 4.16 for the SWIFT links to support the aggregated datarate, resulting in smaller N_λ and $N_\lambda \times \text{BR}$ (aggregated data rate) values for the SWIFT links. Further, it is interesting to note that the 4PAM-EDAC, 4PAM-ODAC, and OOK based CLOS links can achieve $>1,000\text{Gb/s}$ ($>1\text{Tb/s}$) aggregated datarate. This outcome is in strong agreement with the performance analysis done for photonic links in prior works [19] and [202]. However, none of the SWIFT links can achieve $>1\text{Tb/s}$ aggregated datarate, which corroborates the observation that to achieve terascale aggregate data rates in photonic links the losses and power penalties in the links must be minimized. As per our analysis, the CLOS links have significantly low losses and penalties compared to the SWIFT links, and as a result, the CLOS links can achieve $>1\text{Tb/s}$ datarate, whereas the SWIFT links cannot.

In addition, Table 4.5 also lists BER values for various CLOS and SWIFT links evaluated when SECDED coding was not used. These values give insights into how the crosstalk noise present in the links impacts BER. To evaluate the BER for a link, we evaluated the worst-case P_{Xtalk} from Eq. 4.7 across all the filters in the receiving

Table 4.5: Optimal N_λ , bitrate (BR), aggregated datarate ($N_\lambda \times BR$), power budget (P_{dB}^B), $PP_{dB} + 10\log(N_\lambda)$, detector sensitivity (S), and optical laser power (= $PP_{dB} + S$) values for different datarate-BER balanced variants of CLOS and SWIFT links. S varies across different links because S depends on BR.

Variants	ER (dB)	P_{dB}^B	(S)	N_λ	BR	$N_\lambda \times BR$	$PP_{dB} + 10\log(N_\lambda)$	Laser Power	BER Without SECEDED Coding
Various CLOS Links for FSR = 20nm [27]									
OOK	5	38.60	-18.6 dBm	64	17	1,088	38.29	19.69 dBm	3.39×10^{-5}
OOK	9	37.80	-17.80 dBm	64	18	1,152	37.31	19.51 dBm	2.9×10^{-5}
OOK	12	37.1	-17.1 dBm	64	19	1,216	36.31	19.21 dBm	2.6×10^{-5}
4-PAM-SS	5	42.50	-22.5 dBm	32	20	640	41.80	19.3 dBm	12×10^{-4}
4-PAM-SS	9	41.00	-21 dBm	32	27	864	40.8	19.8 dBm	7.9×10^{-4}
4-PAM-SS	12	40.35	-20.35 dBm	32	30	960	39.8	19.45 dBm	7.5×10^{-4}
4-PAM-EDAC	5	40.35	-20.35 dBm	64	30	1,920	38.00	17.65 dBm	8.83×10^{-5}
4-PAM-EDAC	9	37.9	-17.9 dBm	64	35	2,240	37.00	19.1 dBm	8×10^{-5}
4-PAM-EDAC	12	36.1	-16.1 dBm	64	40	2,560	36.00	19.9 dBm	6.2×10^{-5}
4-PAM-ODAC	2	42.50	-22.5 dBm	64	20	1,280	42.00	19.5 dBm	9.7×10^{-4}
4-PAM-ODAC	6	38.5	-18.5 dBm	64	33	2,112	38.31	19.81 dBm	8.3×10^{-5}
4-PAM-ODAC	9	37.9	-17.9 dBm	64	35	2,240	37.41	19.51 dBm	8×10^{-5}
Various SWIFT links for FSR = 20nm [48]									
OOK	5	38.60	-18.6 dBm	32	17	544	38.06	19.46 dBm	8.02×10^{-5}
OOK	9	37.80	-17.80 dBm	32	18	576	37.1	19.30 dBm	7.8×10^{-5}
OOK	12	37.1	-17.1 dBm	32	19	608	36.1	19 dBm	7.1×10^{-5}
4PAM-SS	5	42.10	-22.1 dBm	16	22	352	40.85	18.75 dBm	1.4×10^{-3}
4PAM-SS	9	40.35	-20.35 dBm	16	30	480	39.9	19.55 dBm	1×10^{-3}
4PAM-SS	12	39.1	-19.1 dBm	16	32	512	38.9	19.8 dBm	9×10^{-4}
4PAM-EDAC	5	41.00	-21 dBm	32	27	864	38.16	17.16 dBm	8.17×10^{-4}
4PAM-EDAC	9	37.9	-17.9 dBm	32	35	1,120	37.2	19.3 dBm	7×10^{-4}
4PAM-EDAC	12	41	-21 dBm	64	27	1,728	40.7	19.7 dBm	3.5×10^{-4}
4PAM-ODAC	2	42.30	-22.3 dBm	32	21	672	42.14	19.84 dBm	8.49×10^{-4}
4PAM-ODAC	6	38.5	-18.5 dBm	32	33	1,056	38.41	19.91 dBm	7.3×10^{-4}
4PAM-ODAC	9	42.5	-22.5 dBm	64	20	1,280	42.31	19.81 dBm	7.6×10^{-4}

filter MR bank when considering the P_{NRZ} in Eq. 4.7 to be the signal power reaching the worst-case filter MR after accounting for all the losses and penalties encountered in the link based on Eq. 4.15. From there, we evaluated the signal-to-noise ratio (SNR) to be P_{NRZ}/P_{Xtalk} . Based on the mathematical models and equations of BER provided for guided propagation in [245], [138], we formulate a relation between BER and SNR in Eq. 4.17. M in Eq. 4.17 is defined in Table 4.3. From Eq. 4.17, we formulated the BER equations for OOK ($M = 2$) and 4-PAM ($M = 4$) signals, which are provided in Eqs. 4.18 and 4.19, respectively. We have used the evaluated SNR value in Eq. 4.18 (from [188]) for OOK links and Eq. 4.19 (from [188]) for 4-PAM links to determine the corresponding BER values, which are reported in Table 4.5.

$$\text{BER} = \frac{2(M-1) - \log_2 M}{M \times \log_2 M} \text{erfc} \left(\frac{\sqrt{\text{SNR}}}{(M-1)\sqrt{2}} \right) \quad (4.17)$$

$$\text{BER}_{\text{OOK}} = \frac{1}{2} \text{erfc} \left(\frac{\sqrt{\text{SNR}}}{\sqrt{2}} \right) \quad (4.18)$$

$$\text{BER}_{4\text{-PAM}} = \frac{1}{2} \text{erfc} \left(\frac{\sqrt{\text{SNR}}}{3\sqrt{2}} \right) \quad (4.19)$$

From these BER values (shown in Table 4.5), it is evident that all datarate-BER balanced CLOS and SWIFT links achieve the BER values that are lower than 1.74×10^{-3} , which is a threshold BER value for the SECDED (64, 72) coding to achieve error-free transmission of data packets of size 512 bits (we consider 512-bits long packets in our system-level evaluation in the next section). We obtain this threshold value through the following reasoning: the threshold BER value in a SECDED (64, 72) coded data packet should not incur more than 1-bit of error, as only 1-bit can be corrected for, to achieve error-free transmission of the SECDED (64, 72) coded data packet. Since a 512-bit original data packet gets converted into a 576-bit data packet after it is encoded with the SECDED (64, 72) code, up to only 1-bit in the 576-bit data packet is allowed to be erroneous. Therefore, the threshold BER for this case becomes, $1/576 = 1.74 \times 10^{-3}$. Thus, all the CLOS and SWIFT links in Table 4.5 are capable of achieving error-free data transmission using SECDED (64,72) coding, as all links in Table 4.5 achieve BER of lower than 1.74×10^{-3} , which ensures that all possible bit errors in the 512-bit data packets transmitted over these links can be corrected using the SECDED (64,72) coding.

Moreover, comparing the aggregated datarate values for the links with different modulation methods, it is evident that 4PAM-EDAC and 4PAM-ODAC links in general achieve greater aggregated datarate compared to OOK links. This is because, compared to the OOK links that can have only one bit transferred per signal symbol, the 4PAM-EDAC and 4PAM-ODAC links can achieve greater BR due to their ability to transfer 2-bits per signal symbol. As a result, the 4PAM-EDAC and 4PAM-ODAC links achieve greater values of aggregated data rate ($N_\lambda \times \text{BR}$), despite achieving the same N_λ values as achieved by the OOK links. On the other hand, comparing the different types of 4PAM links with one another, it is evident that: (i) the 4PAM-SS

links achieve lower aggregated datarate values than the 4PAM-EDAC and 4PAM-ODAC links; and (ii) the 4PAM-ODAC links achieve aggregated datarate values that are higher than the 4PAM-SS links but lower than the 4PAM-EDAC links. This is because: (i) the 4PAM-SS links have the largest PPdB values due to their higher MR through losses caused due to $2\times$ more MR modulators required for them (Table 4.1, Fig. 4.2), which results in the lowest N_λ values for them, compared to the 4PAM-EDAC and 4PAM-ODAC links; and (ii) the 4PAM-ODAC links have greater PPER value compared to the 4PAM-EDAC links (Table 4.4), which results in greater PP_{dB} values for the 4PAM-ODAC links, yielding lower values of available P_{dB}^B that support lower BR values (and hence, lower $N_\lambda \times BR$) for the 4PAM-ODAC links.

Table 4.5 also provides optimal N_λ , total power penalty ($PP_{dB} + 10\log(N_\lambda)$), aggregate datarate and optical laser power consumption of OOK and PAM4 based datarate-BER balanced variants of CLOS and SWIFT links, corresponding to different extinction ratios. An increase in the extinction ratio reduces the extinction ratio penalty (PP_{ER}) and filter penalty (Eq. 4.6), which in turn reduces the total penalty in the link (PP_{dB}). This reduction in total power penalty creates more room in the link power budget (P_{dB}^B) to be leveraged to achieve larger N_λ and/or increased bit-rate (BR), and hence, achieve larger aggregate datarate (i.e., $N_\lambda \times BR$) for the link. However, the laser power consumption does not improve noticeably with the increase in extinction ratio, as evident from Tables 4.5 and 4.6. This is because since the laser power consumption is given as $PP_{dB} + S$, the laser power reduces only if the combined effect of the reduction in PP_{dB} and/or increase in S reduces $PP_{dB} + S$. More specifically, from Table 4.5, with the increase in the extinction ratio from 5 dB to 12 dB, the aggregate datarates of the OOK based CLOS and SWIFT links increase from 1.08 Tb/s to 1.22 Tb/s and from 544 Tb/s to 608 Tb/s respectively. This is because for the OOK based CLOS links the BR increases from 17 Gb/s to 19 Gb/s at the unchanged N_λ of 64, and for the OOK based SWIFT links the BR increases from 17 Gb/s to 19 Gb/s at the unchanged N_λ of 32. In terms of laser power consumption, OOK based CLOS and SWIFT links experience reduction in laser power consumption from 19.7 dBm to 19.21 dBm and 19.5 dBm to 19 dBm respectively with increase in extinction ratio from 5 dB to 12 dB. This is because OOK based CLOS and SWIFT links achieve reduced ($PP_{dB} + S$) because of combined effects of reduction in PP_{dB} due to increase in extinction ratio and increase in S due to increase in BR. Similarly, with the increase in extinction ratio from 5 dB to 12 dB, the aggregate datarates of the 4-PAM-EDAC based CLOS and SWIFT links increase from 1.92 Tb/s to 2.6 Tb/s and 864 Gb/s to 1.73 Tb/s respectively. This is because for 4-PAM-EDAC based CLOS links, the BR increases from 30 Gb/s to 40 Gb/s at the unchanged N_λ of 64, and for 4-PAM-EDAC based SWIFT links, the N_λ increases from 32 to 64. In terms of laser power consumption, 4-PAM-EDAC based CLOS and SWIFT links experience an increase in laser power consumption from 17.65 dBm to 19.9 dBm and 17.16 dBm to 19.7 dBm respectively, with increase in extinction ratio from 5 dB to 12 dB. This is because for 4-PAM EDAC based CLOS links, the decrease in PPdB due to the increase in extinction ratio is offset by the larger increase in S due to the increase in BR from 30 Gb/s to 40 Gb/s. In contrast, for the 4-PAM-EDAC based SWIFT links, the increase in N_λ from 32 to 64 increases PPdB, which in turn increases the laser

power consumption of the link. Similarly, with the increase in extinction ratio from 2 dB to 9 dB, the aggregate datarates of the 4-PAM-ODAC based CLOS and SWIFT links increase from 1.3 Tb/s to 2.24 Tb/s and 672 Gb/s to 1.3 Tb/s respectively. This is because for 4-PAM-ODAC based CLOS links, the BR increases from 20 Gb/s to 35 Gb/s with the unchanged N_λ of 64 and for 4-PAM ODAC based SWIFT links, the N_λ increases from 32 to 64 with decrease in BR from 21 Gb/s to 20 Gb/s. In terms of laser power consumption, 4-PAM-ODAC based CLOS links experience an increase in laser power consumption from 19.5 dBm to 19.51 dBm similar to 4-PAM-EDAC based CLOS links. On the other hand, 4-PAM ODAC based SWIFT links experience slight reduction in laser power consumption from 19.84 dBm to 19.81 dBm with increase in extinction ratio. Also, with increase in extinction ratio from 5 dB to 12 dB, the aggregate datarates of 4-PAM-SS based CLOS and SWIFT links increase from 640 Gb/s to 960 Gb/s and 352 Gb/s to 512 Gb/s respectively. In terms of laser power consumption, Similar to 4-PAM-EDAC based links, 4-PAMSS based CLOS and SWIFT links experience increase in laser power consumption from 19.3 dBm to 19.45 dBm and 18.75 dBm to 19.8 dBm respectively with increase in extinction ratio.

In summary, the datarate-BER balanced 4PAM-EDAC links achieve the highest datarate across the CLOS and SWIFT link types. However, it is not clear from these datarate results if the 4PAM-EDAC links can be more energy-efficient than the other types of links. To determine whether the higher overhead of the modulator driver energy for the 4PAM-EDAC links (Table 4.1 and Table 4.2) can offset their highest datarate related benefits to yield lower energy-efficiency for them, compared to the OOK and other 4PAM links, we performed a system-level analysis with real-world benchmark applications, the details of which are discussed in upcoming section.

Results for BER-Optimal Links

Table 4.6 shows the optimal N_λ , bitrate (BR) (evaluated from BaR), aggregated data rate ($N_\lambda \times \text{BR}$), P_{dB}^B values, and $\text{PP}_{dB} + 10\log(N_\lambda)$ values for different BER-optimal variants of CLOS and SWIFT links. Similar to the results for the datarate-BER balanced links presented in Table 4.5, the results presented in Table 4.6 also lead to the following observations: (i) The CLOS links achieve higher datarate than the SWIFT links across all evaluated modulation types, due to the lower P_L^{Sp} , P_L^C , and P_L^{WGP} values for the CLOS links than the SWIFT links. (ii) The 4PAM-EDAC links achieve the highest datarate values across the CLOS and SWIFT types, because 4PAM-EDAC links have the lowest PP_{dB} values, compared to the OOK and other 4PAM links. (iii) The 4PAM-SS links achieve the lowest datarate values because they have the highest PP_{dB} values due to the non-zero PP_{INTRF} for them (Table 4.4) and their higher MR through losses caused due to $2\times$ more MR modulators required for them (Table 4.1, Fig. 4.2).

Table 4.6 also provides optimal N_λ , total loss, aggregate datarate and optical laser power consumption of OOK and PAM4 based BER-optimal variants of CLOS and SWIFT links corresponding to different extinction ratios. As we can infer from Table 4.6, with increase in extinction ratio from 5 dB to 12 dB, the aggregate datarates of OOK based CLOS and SWIFT links increase from 864 Gb/s to 1.02 Tb/s and

512 Gb/s to 672 Gb/s respectively. This is because for OOK based CLOS links, the BR increases from 27 Gb/s to 32 Gb/s at the unchanged N_λ of 32 and for OOK based SWIFT links, the BR increases from 16 Gb/s to 21 Gb/s at the unchanged N_λ of 32. In terms of laser power consumption, with increase in extinction ratio from 5 dB to 12 dB, laser power consumption of OOK based CLOS links reduces from 19.9 dBm to 19.7dBm whereas laser power consumption of OOK based SWIFT links increases from 19.66 dBm to 19.8 dBm. This is because OOK based CLOS links achieve reduced ($PP_{dB} + S$) because of combined effects of reduction in PP_{dB} due to increase in extinction ratio and increase in S due to increase in BR which in turn results in reduced laser power consumption. On the other hand, for OOK based SWIFT links, the decrease in PP_{dB} due to increase in extinction ratio is offset by larger values of S with increase in BR from 16 Gb/s to 21 Gb/s. Similarly, with increase in extinction ratio from 2 dB to 9 dB, the aggregate datarates of 4-PAM-ODAC based CLOS and SWIFT links increase from 768 Gb/s to 1.6 Tb/s and 352 Gb/s to 960 Gb/s respectively. This is because for 4-PAM-ODAC based CLOS links, the BR increases from 24 Gb/s to 50 Gb/s at the unchanged N_λ of 32 and for 4-PAM-ODAC based SWIFT links, the BR increases from 22 Gb/s to 30 Gb/s with increase in N_λ from 16 to 32. In terms of laser power consumption, with increase in extinction ratio from 2 dB to 9 dB, laser power consumption of 4-PAM-ODAC based CLOS links reduces from 19.6 dBm to 19.5 dBm. On the other hand, laser power consumption of 4-PAM-ODAC based SWIFT links increases from 19.65 dBm to 19.8 dBm. This is because for 4-PAM-ODAC based CLOS links, combined effect of reduction in PP_{dB} due increase in extinction ratio and increase in S due to increase in BR from 24 Gb/s to 50 Gb/s reduces laser power consumption of the link. For 4-PAM-ODAC based SWIFT links, increase in PP_{dB} due to increase in N_λ from 16 to 32 increases laser power consumption of the link. Similarly, For 4-PAM EDAC links, with increase in extinction ratio from 5 dB to 9 dB, the aggregate datarates increase from 1.02 Tb/s to 1.5 Tb/s for CLOS links and 512 Gb/s to 736 Gb/s for SWIFT links. This is because for 4-PAM-EDAC based CLOS links, the BR increases from 32 Gb/s to 48 Gb/s at the unchanged N_λ of 32 and for 4-PAM-EDAC based SWIFT links, the BR increases from 32 Gb/s to 46 Gb/s at the unchanged N_λ of 16. In terms of laser power consumption, with increase in extinction ratio from 5 dB to 12 dB, laser power consumption of 4-PAM-EDAC based CLOS links increase from 18.13 dBm to 19.13 dBm and laser power consumption of 4-PAM-EDAC based SWIFT links reduces from 19.7 dBm to 19.6 dBm. This is because for 4-PAM-EDAC based CLOS links, the reduction in PP_{dB} due to increase in extinction ratio is nullified by larger increase in S due to increase in BR from 32 Gb/s to 48 Gb/s which in turn increases the laser power consumption of the link. In contrast, for 4-PAM-EDAC based SWIFT links, the combined effect of reduction in PP_{dB} and increase in S results in reduced laser power consumption of the link. Also, For 4-PAM-SS based links, with increase in extinction ratio from 5 dB to 12 dB, the aggregate datarates increase from 352 Gb/s to 608 Gb/s for CLOS links and 160 Gb/s to 384 Gb/s for SWIFT links. This is because for 4-PAM-SS based CLOS links, the BR increases from 22 Gb/s to 38 Gb/s at unchanged N_λ of 16 and for 4-PAM-SS based SWIFT links, the BR increases from 20 Gb/s to 24 Gb/s with increase in N_λ from 8 to 16. In terms of laser power

consumption, similar to 4-PAM-ODAC links, with increase in extinction ratio from 5 dB to 12 dB, laser power consumption of 4-PAM-SS based CLOS links decreases from 19.93 dBm to 19.1 dBm because of the combined effect of reduction in PPdB due to increase in extinction ratio and increase in S due to increase in BR from 22 Gb/s to 38 Gb/s. On the other hand, laser power consumption of 4-PAM-SS based SWIFT links increases from 17.86 dBm to 19.3 dBm due to increase in PPdB since N_λ increases from 8 to 16.

In addition, it can be observed that the OOK links achieve higher datarate values than the 4PAM-ODAC and 4PAM-SS links. This is because the inclusion of PP_{Fil} , PP_{Mod} and PP_{INTRF} terms in Eq. 4.14 increases the PP_{dB} values for the 4PAM-ODAC and 4PAM-SS links to be greater than the PP_{dB} values for the OOK links, which results in higher values of available PBdB for the OOK links, leading to higher aggregated datarate ($N_\lambda \times BR$) for the OOK links. Due to the inclusion of PP^{Fil} , PP^{Mod} and PP^{INTRF} terms in Eq. 4.14, only the 4PAM-EDAC links among all the three different 4PAM link types achieve greater datarate than the OOK links. However, it is not clear if these datarate benefits can allow 4PAM-EDAC to achieve better energy-efficiency than the OOK links. This is because the greater number of hardware components required for realizing the 4PAM-EDAC links (see the #serialization units, #deserialization units, and #transimpedance op-amps in Table 4.1) can offset their datarate benefits to render them with lower energy-efficiency, compared to the OOK links. To investigate this possibility, we performed a system-level (PNoC-level) analysis with real-world benchmark applications, the details of which are discussed in the upcoming section.

Datarate-BER Balanced vs BER-Optimal Links

From Table 4.5 and Table 4.6, it is evident that the datarate-BER balanced links in general achieve higher aggregated datarate than the BER-optimal links. This is because for the BER-optimal links, due to the inclusion of the terms PP_{Mod} , PP_{Fil} and PP_{INTRF} for the evaluation of PP_{dB} , more of the provisioned optical power is utilized for ensuring the target reliability in terms of the target BER (10^{-9} in this chapter). As a result, a relatively small amount of the total provisioned optical power remains available to support the aggregated datarate for the BER-optimal links, leading to relatively lower aggregated datarate. In contrast, for the datarate-BER balanced links, the exclusion of the terms PP_{Mod} , PP_{Fil} and PP_{INTRF} from the PP_{dB} formula keeps a relatively large amount of the total provisioned optical power available in the links, which supports relatively large values of aggregated datarate for the datarate-BER balanced links. Despite achieving large datarate values, the datarate-BER balanced links may still achieve lower average performance and energy-efficiency compared to the BER-optimal links, especially when they are utilized in a PNoC. This is because, the datarate-BER balanced links in general utilize redundant bits of the SECDED (64, 72) coding in every data packet that traverses the PNoC, which can result in a relatively higher average packet latency and per-packet dynamic energy consumption, ultimately leading to a lower value of the average throughput and energy-efficiency in the PNoC. To investigate this hypothesis, we performed a system-

Table 4.6: Optimal N_λ , bitrate (BR), aggregated datarate ($N_\lambda \times BR$), power budget (P_{dB}^B), $PP_{dB} + 10\log(N_\lambda)$, detector sensitivity (S), and optical laser power ($= PP_{dB} + S$) values for different BER-optimal variants of CLOS and SWIFT links. S varies across different links because S depends on Bar .

Variants	ER (dB)	P_{dB}^B	(S)	N_λ	BR	$N_\lambda \times BR$	$PP_{dB} + 10\log(N_\lambda)$	Laser Power
Various CLOS links for FSR = 20nm [27]								
OOK	5	30.10	-10.1 dBm	32	27	864	30.00	19.9 dBm
OOK	9	28.2	-8.2 dBm	32	30	960	27.63	19.43 dBm
OOK	12	26.6	-6.6 dBm	32	32	1,024	26.3	19.7 dBm
4PAM-SS	5	42.10	-22.1 dBm	16	22	352	42.03	19.93 dBm
4PAM-SS	9	37.9	-17.9 dBm	16	35	560	37.7	19.8 dBm
4PAM-SS	12	37.1	-17.1 dBm	16	38	608	36.2	19.1 dBm
4PAM-EDAC	5	39.10	-19.1 dBm	32	32	1,024	37.23	18.13 dBm
4PAM-EDAC	9	33.4	-13.4 dBm	32	46	1,472	32.93	19.53 dBm
4PAM-EDAC	12	32.3	-12.3 dBm	32	48	1,536	31.43	19.13 dBm
4PAM-ODAC	2	41.70	-21.7 dBm	32	24	768	41.30	19.6 dBm
4PAM-ODAC	6	32.3	-12.3 dBm	32	48	1,536	32.1	19.8 dBm
4PAM-ODAC	9	31.5	-11.5 dBm	32	50	1,600	31	19.5 dBm
Various SWIFT links for FSR = 20nm [27]								
OOK	5	39.10	-19.1 dBm	32	16	512	38.76	19.66 dBm
OOK	9	37.1	-17.1 dBm	32	19	608	36.4	19.3 dBm
OOK	12	35.3	-15.3 dBm	32	21	672	35.1	19.8 dBm
4PAM-SS	5	42.50	-22.5 dBm	8	20	160	40.36	17.86 dBm
4PAM-SS	9	37.1	-17.1 dBm	8	38	304	36.1	19 dBm
4PAM-SS	12	41.7	-21.7 dBm	16	24	384	41	19.3 dBm
4PAM-EDAC	5	39.10	-19.1 dBm	16	32	512	38.80	19.7 dBm
4PAM-EDAC	9	35.3	-15.3 dBm	16	42	672	34.5	19.2 dBm
4PAM-EDAC	12	33.4	-13.4 dBm	16	46	736	33	19.6 dBm
4PAM-ODAC	2	42.10	-22.1 dBm	16	22	352	41.75	19.65 dBm
4PAM-ODAC	6	41.7	-21.7 dBm	32	24	768	41.31	19.61 dBm
4PAM-ODAC	9	40.35	-20.35 dBm	32	30	960	40.15	19.8 dBm

level (PNoC-level) analysis with real-world benchmark applications, the details of which are discussed next.

4.5 System-Level Evaluation

Evaluation Setup and Methodology

For evaluating optimized link-level variants based on OOK and several 4-PAM modulation schemes at system-level, we have considered two separate PNoC architectures: CLOS PNoC [116] and SWIFT PNoC [56]. We particularly selected the photonic crossbar based, high-radix SWIFT PNoC architecture [56] for this system-level analysis, because SWIFT PNoC has been shown in [56] to provide significantly better throughput and energy-efficiency compared to the other classic high-radix PNoC architectures, such as [260]. In addition, to evaluate another high-radix PNoC architecture that is distinct from the photonic crossbar based SWIFT PNoC, we also selected the 8-ary 3-stage CLOS PNoC architecture from [116] that employs WDM based point-to-point photonic links. We preferred high-radix PNoC architectures to more classic, low-radix architectures such as [91] and [277], as prior works [260], and [55] have shown that high-radix PNoC architectures are extremely promising architectures to meet future on-chip bandwidth demands. These PNoC architectures were evaluated for the following modulation schemes: OOK, 4-PAM-SS, 4-PAM-EDAC, and 4-PAM-ODAC.

For CLOS-PNoC shown in Fig. 4.7(a), we have considered an 8-ary 3-stage topology for a 256 x86-core system with 8 clusters, 8 tiles in each cluster, and 4 cores in each tile. The 4 cores of each tile connect with one another via a concentrator. The 8 concentrators corresponding to the 8 tiles in a cluster communicate with one another via an electrical router. The electrical router is a simple 8×8 router, with each concentrator connected to the router using one of its ports. The concentrators and electrical routers are not shown in Fig. 4.7(a). Each router is associated with a photonic transmitter and receiver block (Fig. 4.7(a)), and the electrical-optical-electrical conversion happens at the photonic transmitter-receiver block. For inter-cluster communication, point to point photonic waveguides are supported by the photonic transmitter-receiver blocks, with forward or backward propagating wavelengths depending upon the physical location of the source and destination clusters. All the clusters are connected together using total 56 waveguides (WGs). The PNoC uses two laser sources to enable bi-directional communication.

For SWIFT PNoC shown in Fig. 4.7(b), we have again considered a 256 x86-core system. Every 4-core cluster is considered a node here and communication within a node occurs through a 5×5 electrical router. Four ports of the router connect the processing cores to the router and the fifth port of the router is connected to a gateway interface (GI) which facilitates transfers between the electrical and photonic layers. The routers use round-robin arbitration to facilitate communication between cores and the GIs. Each GI connects four nodes. The architecture utilizes a photonic crossbar topology with eight waveguide groups and four Multiple Writer Multiple

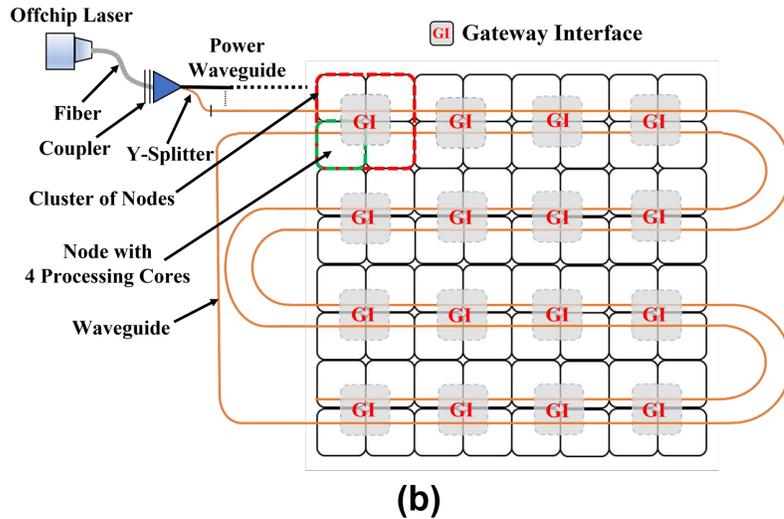
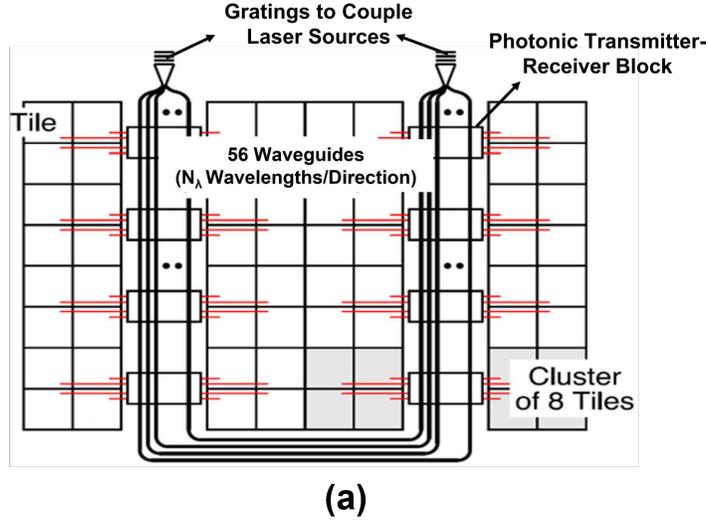


Figure 4.7: Schematics of (a) 8-ary 3-stage CLOS PNoC architecture [116] and (b) SWIFT PNoC architecture [56].

Reader (MWMR) WGs per group. A broadband off-chip laser with a laser power controller is used to power the WGs.

We consider the two-layer, 3D chip organization from [161] and [169] for each of these PNoC architectures. The bottom CMOS layer contains processing cores, caches, and electrical interconnects. The silicon-photonic top layer contains photonic transmitter-receiver (Tx-Rx) blocks (for O/E and E/O conversions), as well as photonic devices and circuits that constitute a PNoC. Through-silicon vias (TSVs) are used to vertically connect the bottom layer with the top layer at every photonic Tx-Rx block. The parameters used for modeling the 3D organizations of our considered PNoCs are given in Table 4.7. The N_λ and BR values in Table 4.7 can be taken from Table 4.5 and 4.6, depending on the utilized modulation scheme and target design goal (datarate-BER balanced or BER-optimal design).

For evaluating the impact of various signaling methods on these architectures' performance and efficiency, we performed a benchmark-driven simulation-based analysis using Gem5 full-system simulation [38] and an enhanced cycle-accurate PNoC simulator that extends the Noxim simulator [238]. For Gem5 simulations, we assumed 32KB direct mapped L1 and 128 KB direct mapped L2 caches (MOESI coherency for L2) per core and a main memory of 32 GB DDR4 RAM. Tables 4.5 and 4.6 show the number of wavelengths (N_λ) and the maximum datarate for which the simulations were run for CLOS PNoC and SWIFT PNoC. These N_λ and datarate values were utilized to model the links in different variants of the CLOS and SWIFT PNoCs that correspond to various signaling schemes and datarate-BER balanced and BER-optimal designs. The energy and power value considerations from Tables 4.1 and 4.2 were also incorporated into our simulations. The performance was evaluated at a 22nm CMOS node. The floorplan and the number of WGs were kept constant across all variants of a particular PNoC architecture, with only the link configuration parameters (e.g., N_λ , datarate, number of hardware instances from Tables 4.1 and 4.2 that depend on N_λ) changing across the variants. PARSEC benchmark applications were used to generate real-world traffic traces. The traces were generated using GEM5 full-system simulation and these traces were fed into our cycle accurate PNoC-simulator. In GEM5 simulations, the warm-up period was set as 100M cycles and the traces were captured for the subsequent 1B instructions. The simulations were used to evaluate average latency, energy-per-bit (EPB) and a breakdown of total power dissipation. Electrical energy consumption by routers and GIs was determined using the DSENT tool [238]. To obtain the laser power consumption, the total required optical power in the CLOS and SWIFT PNoC architectures were evaluated based on the P_{dB}^B and PP^{dB} values from Tables 4.5 and 4.6 for different variants, and then 15% wall-plug efficiency was assumed to convert these optical power values into the corresponding electrical laser power values. The energy and power value considerations from Tables 4.1 and 4.2 were also incorporated into our simulations. The performance was evaluated at a 22nm CMOS node. The floorplan and the number of WGs were kept constant across all variants of a particular PNoC architecture, with only the link configuration parameters (e.g., N_λ , datarate, number of hardware instances from Table 4.1 and Table 4.2 that depend on N_λ) changing across the variants.

To implement SECDED encoding for the datarate-BER balanced variants of the CLOS and SWIFT PNoC architectures, we employed a lookup table-based approach where each input 512-bit packet is encoded with the SECDED scheme via byte-level lookup tables. In other words, every Byte of the input packet gets encoded through a separate and parallelly operating lookup table. Because of this parallelism in encoding, this encoding incurs only a one cycle delay. SECDED decoding is also handled using byte-wise parallelly operating lookup tables. However, the one cycle delay of only the decoding phase comes in the critical latency path in the PNoCs, as the encoding delay can be hidden by overlapping the encoding operation with the arbitration and receiver selection phases in the PNoC. We considered the area overhead of the SRAM-implemented lookup tables (at the 22 nm node) in the encoding and decoding units, and estimated it to be 1142 μm^2 each. Each GI in the PNoC should have one encoding unit and one decoding unit, therefore, each GI would have 2284 μm^2

Table 4.7: Parameters for modeling the 3D organizations of our evaluated PNoCs.

Parameters	CLOS PNoC [116]	SWIFT PNoC [56]
Network Size	256 cores	256 cores
Network Radix	8	16
Network Diameter	1	1
Bisection Bandwidth (Gb/s)	$56 \times N_\lambda \times BR$	$32 \times N_\lambda \times BR$
Traffic Model	Multi-Threaded PARSEC Workloads	
Photonic Layer Frequency	5 GHz [169]	
Processing Core x86 Frequency	2.5 GHz [169]	
TSV Channel Configuration Per Photonic Tx-Rx Block	8 TSV Bundles	
TSV Bundle Size and Layout	2×2 TSVs per Bundle [10]	
TSV Speed	21 Gb/s [10]	
Energy of a TSV Bundle	6.7 pJ [10]	

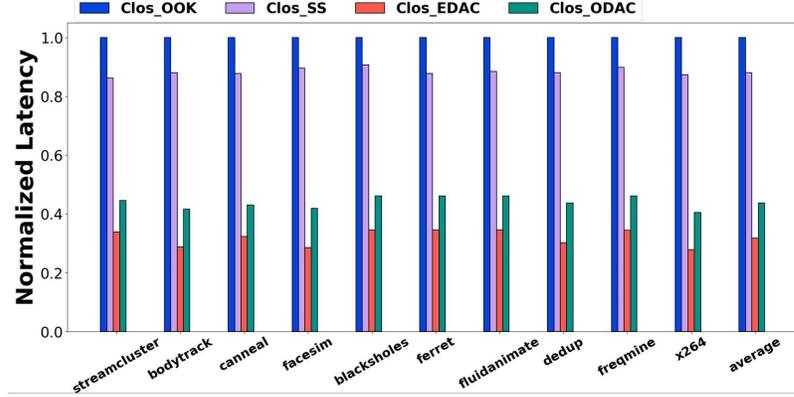
of area overhead. In addition, each encoding or decoding event for a 512-bit packet is estimated to consume 0.1 pJ energy. The area estimates were obtained using logic synthesis analysis. The energy and delay values were evaluated using CACTI-P [148] and are accounted for in our system-level analysis.

The following subsection discusses the simulation results and how the modulation schemes compare against one another, for the two considered PNoC architectures.

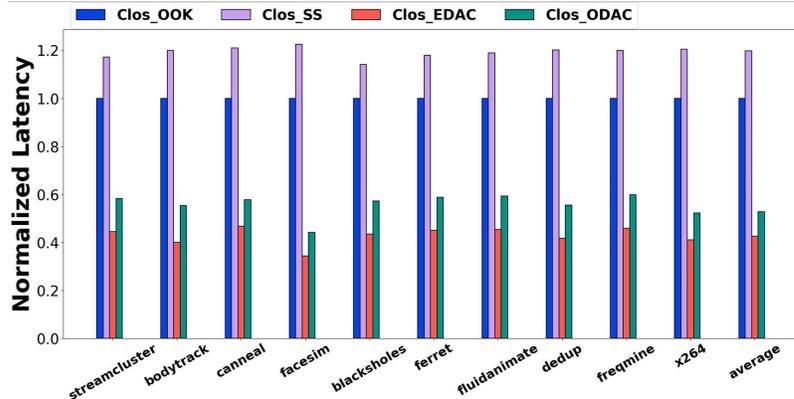
Results and Discussion

Packet Latency

Figs. 4.8(a) and 4.8(b) show the average latency for different variants of CLOS PNoC, for the different applications from the PARSEC benchmark suite, with all results normalized to CLOS_OOK for both the datarate-BER balanced and BER-optimal cases. It can be observed that the EDAC and ODAC variants of 4-PAM modulation on CLOS outperform the rest of the variants, and when compared to the baseline CLOS_OOK, they achieve 68% and 55% better latency for the balanced datarate-BER case, and 62% and 34% better latency for the BER-optimal case, on average. The 4-PAM-SS variant of CLOS displays $1.2\times$ higher latency than the baseline for BER-optimal designs. The packet latency we observe is an indicator of the combined effect of the N_λ and the datarates achieved for the links of these CLOS variants. Having a higher N_λ increases the number of concurrent bits transferred over the network, which in turn reduces the packet transfer latency. Similarly, having a higher bit-rate increases the number of bits transferred in the given time frame, which also



(a)



(b)

Figure 4.8: Packet latency plotted across PARSEC benchmark applications for (a) datarate-BER balanced variants, and (b) BER-optimal variants of CLOS PNoC. All results are normalized to the baseline CLOS_OOK.

results in a lower packet latency. In addition, Utilizing 4-PAM allows us to transmit $2\times$ bits per cycle for the same N_λ , allowing the 4-PAM signaling based variants of CLOS to have better latency values compared to their OOK based variants.

Figs. 4.9(a) and 4.9(b) show the latency results for SWIFT PNoC, with results normalized to SWIFT-OOK, which acts as the baseline for our analyses. Here, again, the EDAC and ODAC 4-PAM variants obtain better latency values when compared to the baseline, due to the higher bandwidth they can achieve, as shown in Table 4.5. For these results, we can see that the EDAC variant performs better than other variants: 65% better latency than the baseline for the datarate-BER balanced case on average, as in Fig. 4.9(a), and 53% better latency on average for the BER-optimal case, as in Fig. 4.9(b). It can be noted that, similar to CLOS variants, for SWIFT PNoC as well, the SS variant has higher latency than the baseline for BER-optimal design, with $1.35\times$ higher latency than the baseline on average. We can observe a similar trend in latency across the SWIFT PNoC variants, owing to the same reasons as discussed above, for the CLOS PNoC.

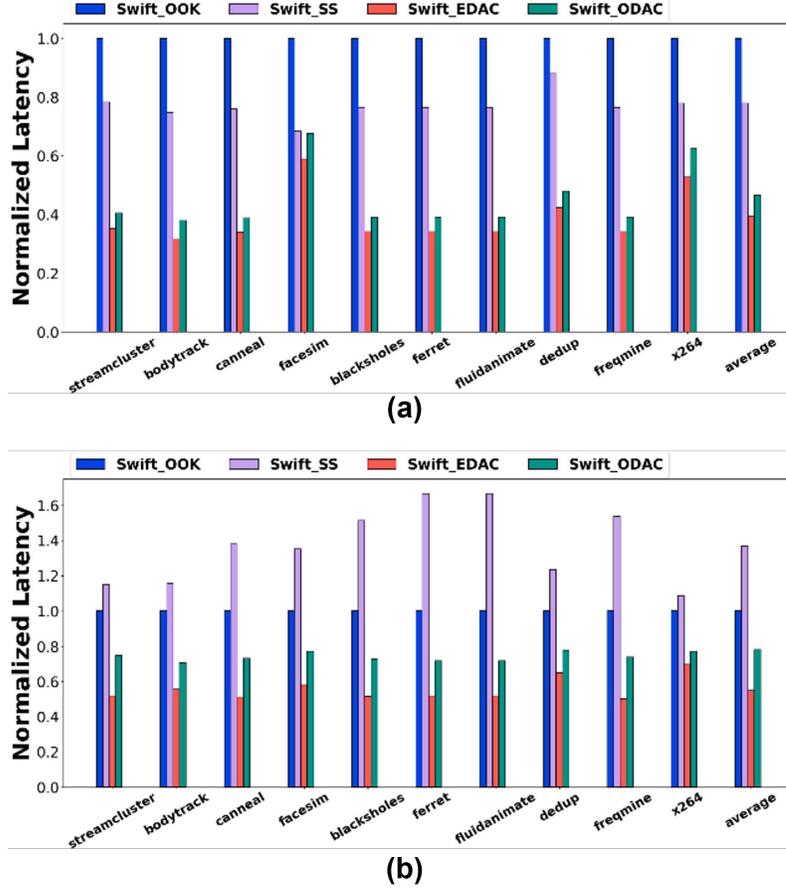
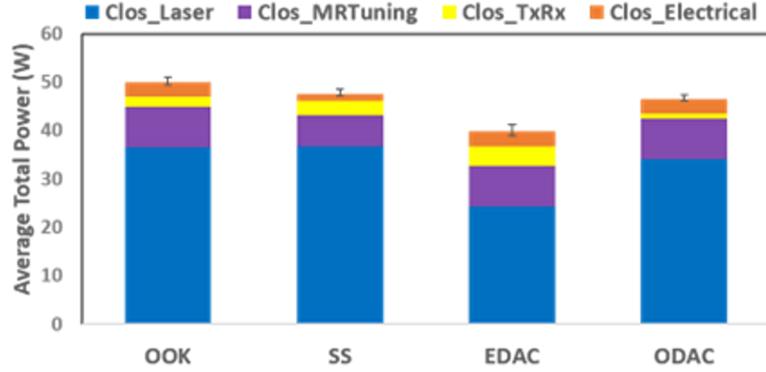


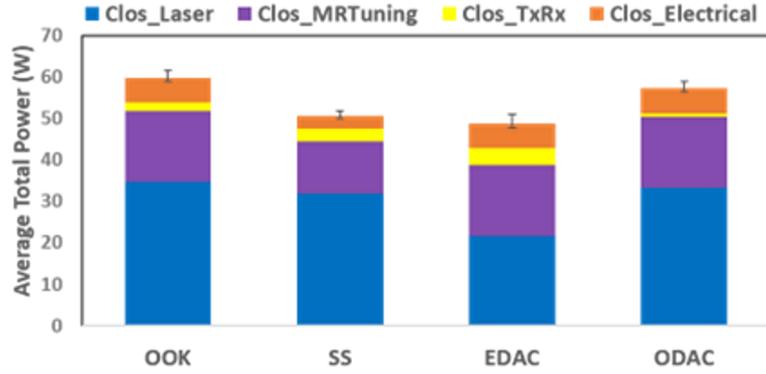
Figure 4.9: Packet latency plotted across PARSEC benchmark applications for (a) datarate-BER balanced variants, and (b) BER-optimal variants of SWIFT PNoC. All results are normalized to the baseline SWIFT_OOK.

Power Dissipation

Next, we examined the power dissipation in the considered PNoCs, with the results for CLOS PNoC shown in Fig. 4.10 and the results for SWIFT PNoC shown in Fig. 4.11. We report total power that is averaged across the considered PARSEC benchmark applications, and the corresponding error bars (with minimum and maximum values) are also shown in Figs. 4.10 and 4.11. The column heights in these figures show average values of total power, which is the sum of laser power (wall-plug power of laser sources), electrical power (power consumption of intra-cluster/intra-node communication in the electrical domain), RxTx power (dynamic power consumption of operating receiver/ transmitted modulators, other devices, and the E/O and O/E conversion modules in PNoCs; per-MR values from Table 4.2), and total MR tuning power (sum of MR tuning power + microheater power; per-MR values from Table 4.2). The link-level results obtained for various CLOS and SWIFT links that are shown in Table 4.5 and Table 4.6 are directly reflected in the power results shown in Figs. 4.10 and 4.11. The wall-plug laser power values in Figs. 4.10 and 4.11 are directly dependent on the optical laser power values given in Tables 4.5 and 4.6. The



(a)



(b)

Figure 4.10: Average total power dissipation for different (a) datarate-BER balanced, and (b) BER-optimal variants of CLOS PNoC. The error bars represent the minimum and maximum values of power dissipation across 12 PARSEC benchmarks.

higher optical laser values in Tables 4.5 and 4.6 translate into higher wall-plug laser power values in Figs. 4.10 and 4.11. Along the same lines, the TxRx and MR tuning power values in Figs. 4.10 and 4.11 depend on the N_λ values from Tables 4.5 and 4.6, the higher N_λ values translate into higher TxRx and MR tuning power values, as the number of MRs employed in a PNoC architecture depends on N_λ and TxRx and MR tuning power depend on the number of MRs. Similarly, the higher values of N_λ in Tables 4.5 and 4.6 also result in higher values of intra-cluster/intra-node electrical communication power in Figs. 4.10 and 4.11, as the sizes of the required electronic buffers in the intra-cluster electrical routers in PNoCs depend on the N_λ values, and these buffer sizes in turn control the electrical power consumption in these routers.

From Fig. 4.10, it can be observed that among the CLOS PNoC variants, the EDAC variants for both the datarate-BER balanced and BER-optimal cases dissipate the least power compared to other variants. This is because the laser power dissipation is the major contributor to the total power dissipation in CLOS PNoC, and the constituent links of the EDAC variants of the CLOS PNoC dissipate the lowest optical laser power (Tables 4.5 and 4.6) compared to the SS, ODAC, and OOK variants. To further analyze the power results, after the laser power, the second major contributor

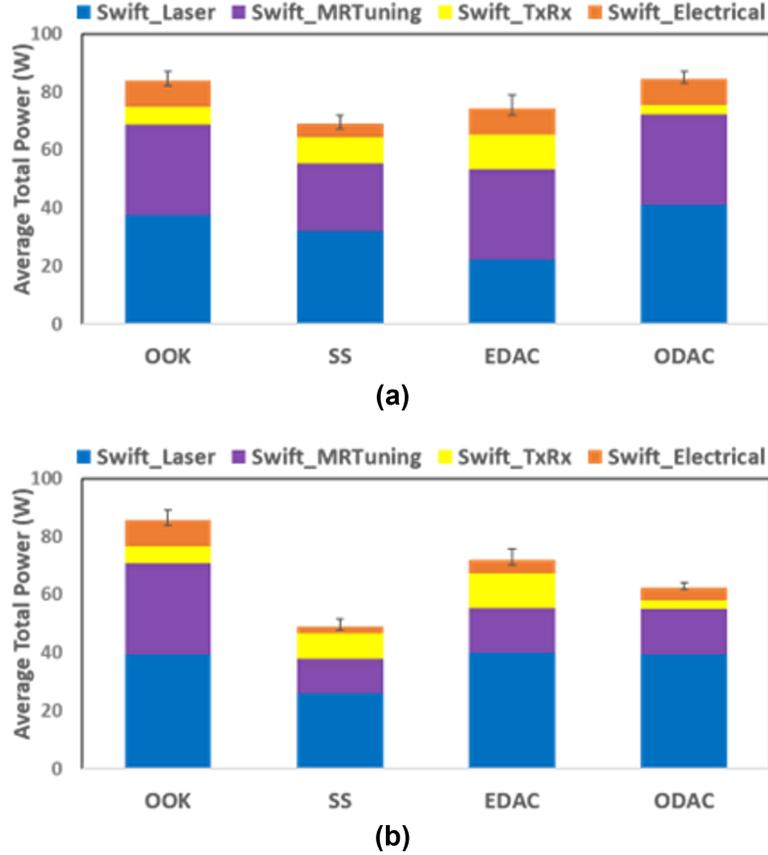


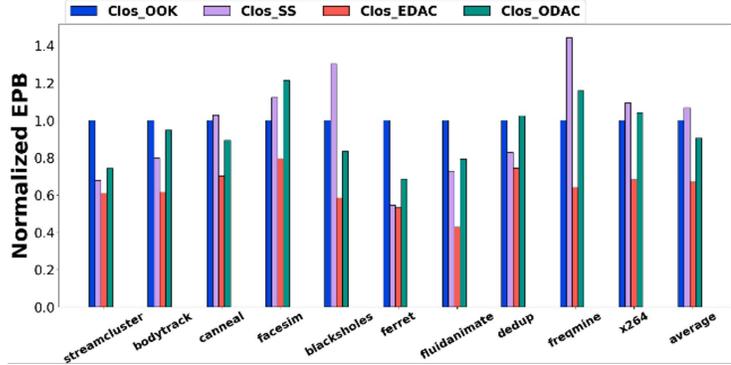
Figure 4.11: Average total power dissipation for different (a) datarate-BER balanced, and (b) BER-optimal variants of SWIFT PNoC. The error bars represent the minimum and maximum values of power dissipation across 12 PARSEC benchmarks.

to the total power dissipation in CLOS PNoC is the MR tuning power, followed by the electrical power and TxRx power. The MR tuning power varies across different variants, because different variants require different number of MRs due to different N_λ values. In contrast, the SS variants dissipate less electrical power compared to the other variants, because the SS variants achieve smaller N_λ values, which in turn reduces the complexity of the routers and GIs in the SS variants. Also, the ODAC variants dissipate less TxRx power compared to the other variants, because the ODAC variants consume less dynamic energy in the modulator drivers (as can be inferred from the EMod values in Table 4.2).

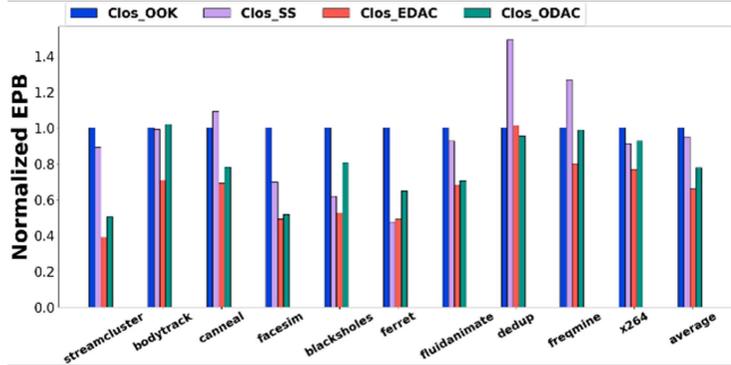
Energy-per-Bit

The energy-per-bit (EPB) results for CLOS variants are shown in Figs. 4.12(a) and 4.12(b), and for SWIFT PNoC in Figs. 4.13(a) and 4.13(b). Fig. 4.12(a) shows the EPB results for datarate-BER balanced variants of the CLOS PNoC, where EDAC and ODAC variants of the CLOS PNoC have better EPB values, on average across PARSEC benchmarks [37], in comparison with the other CLOS variants, with

15% and 11% lower EPB, respectively, than the OOK variant. This is because of the higher aggregate data rate and lower packet latencies of the EDAC and ODAC variants resulting in lower energy consumption. For the BER-optimal case, as shown in Fig. 4.12(b), the OOK variant for CLOS PNoC can be observed to have much better performance, than the ones utilizing 4-PAM techniques. Among the 4-PAM techniques, using SS has substantially higher energy utilization, with $4.9\times$ more EPB than the baseline OOK on average. Both ODAC and EDAC variants of CLOS PNoC exhibit $\sim 1.8\times$ EPB of the OOK baseline for the BER-optimal case.



(a)



(b)

Figure 4.12: Energy-per-bit (EPB) analysis for (a) the datarate-BER balanced variants, and (b) the BER-optimal variants of the CLOS PNoC. Column heights represent EPB averaged across 100 PV maps and normalized to the CLOS-OOK variant.

Among different SWIFT PNoC variants for the datarate-BER balanced case (Fig. 4.13(a)), we can see that the EDAC variant performs better across the benchmark applications, retaining $2.5\times$ less EPB than the baseline OOK, on average across the PARSEC benchmarks [37]. The ODAC variant has comparable EPB consumption to the EDAC variant, with consuming $\sim 2.4\times$ less EPB than the baseline on average across the benchmark applications. The SS variant also performs better than the baseline, consuming $1.2\times$ less EPB on average across the benchmarks.

Among the BER-optimal variants of the SWIFT PNoC (Fig. 4.13(b)), the SS variant has $\sim 1.7\times$ more EPB value than the baseline OOK variant. On the other hand, the EDAC variant consumes $2.22\times$ less EPB than the baseline OOK on av-

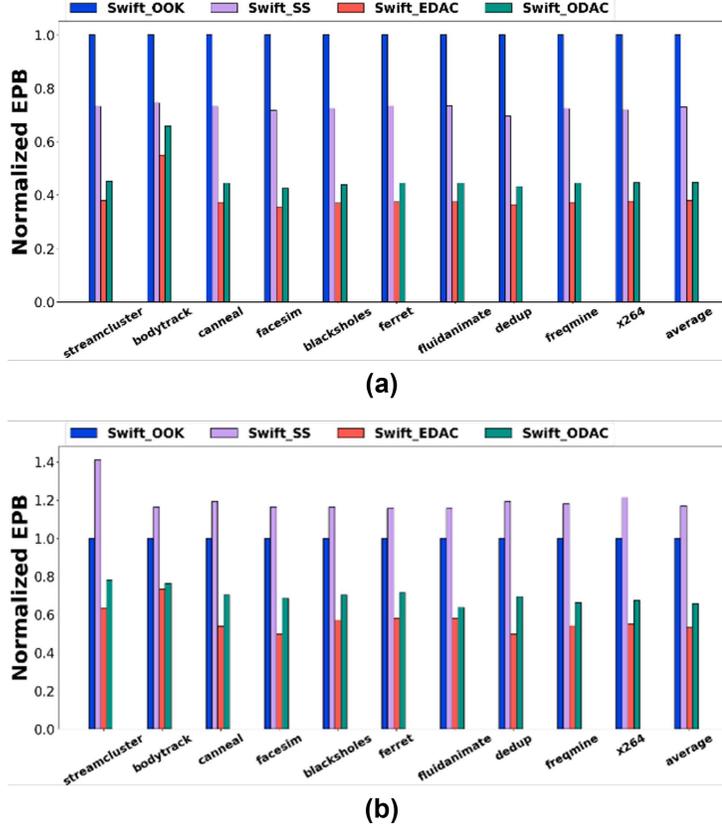


Figure 4.13: Energy-per-bit (EPB) analysis for (a) the datarate-BER balanced variants, and (b) the BER-optimal variants of the SWIFT PNoC. Column heights represent EPB averaged across 100 PV maps and normalized to the SWIFT-OOK variant.

erage across the benchmarks. The ODAC variant consumes $1.67\times$ less EPB than the baseline on average. For the SWIFT PNoC variants, as well as for the CLOS PNoC variants, the energy utilized by the laser sources and TxRx modules is the main factor controlling the EPB values, as seen in the power breakdowns in Figs. 4.10 and 4.11. The reduced power consumption of EDAC variants as shown in Figs. 4.10 and 4.11 along with their lower latency of operation, leads to better throughput, and results in better EPB values for these variants for both the PNoCs considered in our evaluations.

In summary, across all the OOK and 4-PAM variants of the CLOS and SWIFT PNoCs, the 4-PAM-EDAC variants exhibit the lowest latency and energy on average across the considered PARSEC benchmark applications. For the balanced datarate-BER case, compared to the baseline OOK variants, the 4-PAM-EDAC variants of the CLOS and SWIFT PNoCs achieve 68% and 65% of the latency, as well as 66% and 64% of the EPB. Similarly, for the optimal BER case, compared to the baseline OOK variants, the 4-PAM-EDAC variants of the CLOS and SWIFT PNoCs achieve 62% and 53% of the latency, as well as 38% and 57% of the EPB. These outcomes motivate the use of 4-PAM-EDAC signaling over OOK and other 4-PAM signaling methods to achieve significantly better energy-efficiency for on-chip communication

with PNoCs.

4.6 Summary

Conventional OOK based signaling enables high-bandwidth parallel data transfer in PNoCs, but as the number of DWDM wavelengths increases, the power, area consumption, and bit-error rate (BER) in PNoCs increase as well. To address this problem, 4-PAM signaling has been introduced which can double the aggregated datarate without incurring significant area, power, and BER overheads. In this chapter, for the first time, we performed a detailed analysis of various designs of 4-PAM modulators, including 4-PAM-SS, 4-PAM-EDAC, and 4-PAM-ODAC. We utilized these modulators to design 4-PAM photonic links and PNoC architectures with two different design goals of achieving the BER-balanced datarate (achieving maximum datarate with a desired BER of 10^{-9} using FEC codes) and optimal BER (achieving desired BER of 10^{-9} using increased input optical power). We then compared these BER-optimal and datarate-BER balanced 4-PAM links and PNoC architectures with the conventional OOK modulator based photonic links and architectures, in terms of performance (datarate and latency), BER, and energy-efficiency. Our analysis with CLOS PNoC and SWIFT PNoC architectures that are designed using the OOK, 4-PAM-SS, 4-PAM-EDAC and 4-PAM-ODAC based modulators and links showed that the 4-PAM-EDAC variants of the CLOS and SWIFT PNoCs yield the least latency and consume the least energy on average across the considered PARSEC benchmark applications. For the balanced datarate-BER case, compared to the baseline OOK variants, the 4-PAM-EDAC variants of the CLOS and SWIFT PNoCs respectively achieve 68% and 65% of the baseline latency, as well as 66% and 64% of the baseline EPB. Similarly, for the optimal BER case, compared to the baseline OOK variants, the 4-PAM-EDAC variants of the CLOS and SWIFT PNoCs respectively achieve 62% and 53% of the baseline latency, as well as 38% and 57% of the baseline EPB. These outcomes push for the PNoC architectures of the future to employ the 4-PAM-EDAC signaling over the OOK and other 4-PAM signaling methods to achieve significantly better energy-efficiency for on-chip communication.

Chapter 5 An Analysis of Various Design Pathways Towards Multi-Terabit Photonic On-Interposer Interconnects

5.1 Introduction

With the recent deluge of data-centric computing applications, such as deep learning, and graph analytics, the world's appetite for analyzing massive amounts of structured and unstructured data has grown dramatically. For instance, since 2012, the amount of compute used in the largest AI training jobs has been increasing exponentially with a 3.4-month doubling time [180], which is $50\times$ faster than the pace of Moore's Law. Fulfilling this appetite demands for increasingly high computational capacity (in terms of compute and memory bandwidths) and energy efficiency. However, consistently meeting this sustaining demand using the currently utilized large monolithic manycore chips and homogeneous multi-chip board designs (e.g., [50][61][83][94][117][264]) is becoming increasingly difficult. This challenge is primarily attributed to three fundamental reasons. First, this demand is quickly outpacing the progress realized by dwindling Moore's law, due to the fundamental physical limitations slowing the rate and increasing the complexity and cost of transition from one technology node to the next [63]. Second, the attempts to scale the size of large monolithic chips gives rise to extravagant manufacturing cost due to the limited reticle size and poor yield of stitching multiple reticles together [172]. Third, scaling the multi-chip board designs can push the package to die ratio in such designs to be greater than 10:1 [231], which in turn can dramatically increase the area overhead of computing systems that employ such multi-chip board designs.

To overcome these challenges, the industry has focused on system dis-aggregation as a solution, wherein a large monolithic system-on-chip is partitioned into multiple smaller, modular chiplets of heterogeneous types. These chiplets are then assembled into a large system-on-package using organic substrate (e.g., [110][11]), silicon interposer (e.g., [237][119][114][104][132]), or silicon wafer (e.g., [187][185][186][24][113]) as the substrate for chiplet assembly and packaging. The size of the silicon interposer based chiplet assemblies is typically limited to $<1,000\text{ mm}^2$ due to the limited reticle size [172]. Nevertheless, the silicon interposer based chiplet assemblies have several advantages over the organic substrate and silicon wafer based assemblies. Unlike the organic substrate based assemblies, the silicon interposer based assemblies have lower package-to-die ratio [231], which decreases their system area overheads. Along the same lines, unlike the silicon wafer based assemblies, the silicon interposer based assemblies are relatively less susceptible to challenges related to power delivery and thermal stability. Furthermore, silicon interposer-based assemblies offer the potential to integrate active front-end-of-line logic components directly onto the interposer. This integration opens up opportunities to enhance inter-chiplet interconnect bandwidth density and efficiency. In addition, it enables the implementation of advanced network topologies and routing logic directly within the interposer, leading to improved communication capabilities. On the other hand, the waferscale chiplet

assemblies are very large in size compared to the interposer based assemblies. But unlike the silicon interposer, to achieve >90% yield, the silicon wafer substrate, upon which various chiplets are assembled, has to remain passive. Because of these advantages, the silicon interposer based chiplet assemblies are rapidly materializing in both the industry and academia.

As such, silicon interposer based chiplet assemblies require efficient implementation of inter-chiplet communication with low end-to-end latency, high bandwidth density, and high scalability, all achieved within a strict power budget. In general, the silicon interposer substrate can be active or passive [237], and its use for assembling chiplets can be based on through silicon vias (TSVs) (e.g., as in TSMC’s CoWoS technology family [60]) or completely free of TSVs (e.g., as in Intel’s EMIB technology family [159]). Regardless of the type of the utilized interposer, the interposer based chiplet systems can support inter-chiplet interconnects with tangibly very high (potentially multi-Tb/s) bandwidth densities [62]. But such extreme-scale interconnect bandwidths are supported only for the inter-chiplet distance of less than 200-300 μm . In addition, prior work [34] has shown that as the number of chiplets on the interposer increases, the average latency of the inter-chiplet interconnects in the state-of-the-art interposer assemblies scales very poorly, regardless of the utilized interconnects topology. This is mainly because the data rates and latency of on-interposer electrical wires scale poorly due to their high impedance dependence. To overcome these shortcomings, prior works have proposed active silicon-phonic interposer (SiPhI) based chiplet systems (e.g., [29][255]). These systems consider the bandwidth density of the inter-chiplet on-SiPhI interconnects to be $\sim 1 \text{ Tb/s/mm}^2$, because there is a push for the next generation interconnects to have $\geq 1 \text{ Tb/s/mm}^2$ bandwidth density [65]. In fact, SiPh based chiplet assemblies from prior work have shown multi-Tb/s/mm² bandwidth densities for optical fiber based off-package I/Os [66][203]. As the natural progression from these excellent outcomes from prior works and driven by the increasing bandwidth needs of emerging workloads, there is impetus to achieve multi-Tb/s bandwidth across the SiPh interposer with an end-to-end latency of no more than $\sim 10\text{ns}$. However, to meet this goal, there are some daunting challenges to overcome. The major challenge is that as a SiPhI system scales to reach the reticle limit, the length of the end-to-end on-SiPhI links tends to become greater than 10 cm. For such long links, optical signal losses can become notably high, which in turn can make it very difficult to achieve even multi-Tb/s interconnect bandwidth, let alone achieving multi-Tb/s/mm² bandwidth density. Unfortunately, this challenge has not been addressed by any prior works so far.

To address this challenge, in this chapter, *for the first time*, we identify the key pathways for the design of multi-Tb/s on-SiPhI links, by taking clues from the existing literature on the design and optimization of SiPh interconnects, both in the on-chip and off-chip design domains. Our identified design pathways include: (1) increasing the available optical power budget per on-SiPhI link by minimizing the insertion losses and power penalties in the link, (2) increasing the spectral bandwidth available per on-SiPhI link (normally referred to as free-spectral range (FSR)) for higher degree of wavelength multiplexing, and (3) increasing the available optical power budget per on-SiPhI link by increasing the maximum allowable optical power (MAOP) limits

of the link. We explore these SiPhI link-level design pathways in isolation and in various combinations of one another, to investigate which of these design pathways can help achieve multi-Tb/s on-SiPhI links. Based on our link-level analysis, We also enable the following two chiplet-based systems with our designed on-SiPhI multi-Tb/s links and provide their system-level performance analysis: *(i)* a CPU based manycore multi-chiplet architecture named NUPLet [29], and *(ii)* a GPU based deep learning training system from [127] that employs a total of 512 multi-chiplet GPU modules.

The key contributions of this chapter are summarized below:

- We consider three state-of-the-art SiPh fabrication platforms from prior works [236] and [12], and then derive different variants of on-SiPhI links based on different combinations of these considered platforms and our identified design pathways mentioned earlier;
- We perform link-level analysis for all the derived on-SiPhI link variants, from which we calculate the achievable aggregate bandwidth and energy-per-bit (EPB) values for each on-SiPhI link variant for link lengths of up to 10 cm;
- We identify all viable on-SiPhI link variants that can support multi-Tb/s aggregate bandwidth;
- We use our identified viable link variants to enable and evaluate different variants of two SiPhI based multi-chiplet systems from prior work: (1) a CPU system from [29], and (2) a GPU based deep learning training system based on [127];
- We perform benchmark-driven analysis of our considered CPU based system variants to evaluate their performance (in terms of execution time), energy and energy-delay product, for PARSEC benchmark applications. Similarly, we also analyze our considered GPU based system variants to evaluate the training time-to-accuracy for deep learning applications.

5.2 Preliminaries

On-Silicon Photonic Interposer (On-SiPhI) Inter-Chiplet Links

Prior work [4] provides a survey of design methods for multi-chiplet packages that integrate silicon photonics and electronics together. The use of a SiPhI for such integration is one of the approaches advocated in this work. Based on this approach, Fig. 5.1 illustrates our envisioned schematic of an on-SiPhI link. The basic component of an on-SiPhI link is a silicon waveguide (shown in gray in Fig. 5.1) that is implemented on the SiPhI. The other SiPh components of the link that are implemented on the SiPhI include: a grating coupler; a transmitter microring resonator group (Tx MRRG); and a receiver microring resonator group (Rx MRRG). In addition, the on-SiPhI link also has other electronic and electro-optic components that are implemented on chiplets. These components are: a laser chiplet that has a comb

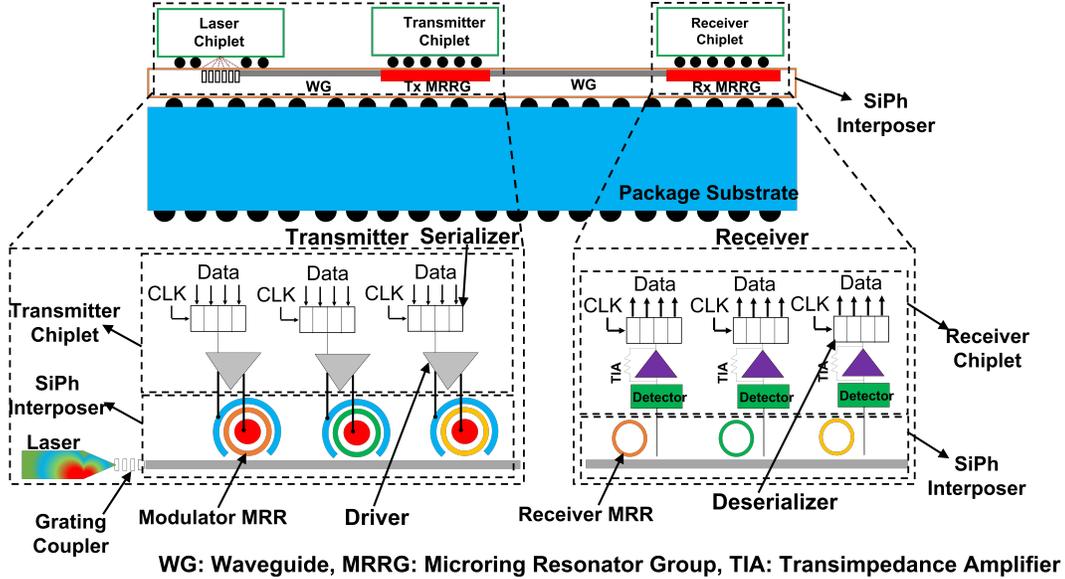


Figure 5.1: Inter-Chiplet Silicon Photonic MRR-based DWDM Link.

laser source implemented on it [235][85]; a transmitter chiplet that has Tx MRRG peripheral circuits such as modulator drivers and serializers; and a receiver chiplet that has Rx MRRG peripheral circuits such as transimpedance amplifiers, and deserializers. The comb laser source on the laser chiplet emits a comb of optical wavelengths that are coupled into the on-SiPhI silicon waveguide via the grating coupler using the dense wavelength division multiplexing (DWDM) technique. These wavelengths work as different data-carrying channels. When these wavelength channels reach the Tx MRRG, the individual MRR modulators of the Tx MRRG modulate input data signals onto these wavelength channels. These modulated wavelength channels are transmitted to the Rx MRRG at the receiver side through the on-SiPhI silicon waveguide. The Rx MRRG consists of an array filter MRRs whose resonances are tuned to the incoming wavelength channels. These MRRs drop the incoming modulated channels onto their respective photodetectors to recover the input data signals. If the number of multiplexed wavelength channels into the on-SiPhI waveguide is N_λ , and if each wavelength channel operates at bitrate of BR Gb/s, then the on-SiPhI waveguide can support $N_\lambda \times BR$ Gb/s bandwidth. Hence, to achieve ≥ 1 Tb/s bandwidth, the on-SiPhI waveguide must support sufficiently high values of N_λ and BR . Factors that impact the achievable values of N_λ and BR per on-SiPhI waveguide are discussed next.

Performance of On-SiPhI Links

The peak and average performance values of inter-chiplet on-SiPhI networks, in terms of throughput, latency, and energy efficiency, are attributed to the speeds and energy costs of the individual sender-to-receiver data transfers across the networks. Since all SiPh networks, including on-SiPhI networks, are typically circuit-switched due to the lack of appropriate means for storing optical packets on the fly, each individual

Table 5.1: Device-layer, circuit-layer, and system-layer features that influence the performance and energy efficiency of inter-chiplet on-SiPhI interconnects.

Device-Layer	Circuit-Layer	Architecture-Layer
Propagation Loss	Inter-Channel Spacing	Network Topology
Coupling Loss	Through Loss	Arbitration and Routing
Free-Spectral Range (FSR)	Tx Crosstalk Penalty	Traffic Patterns
Per-Wavelength MAOP	Rx Signal Degradation Penalty	Bandwidth Allocation
Per-Waveguide MAOP	Multiplexing Strategy	Optical Power Provisioning

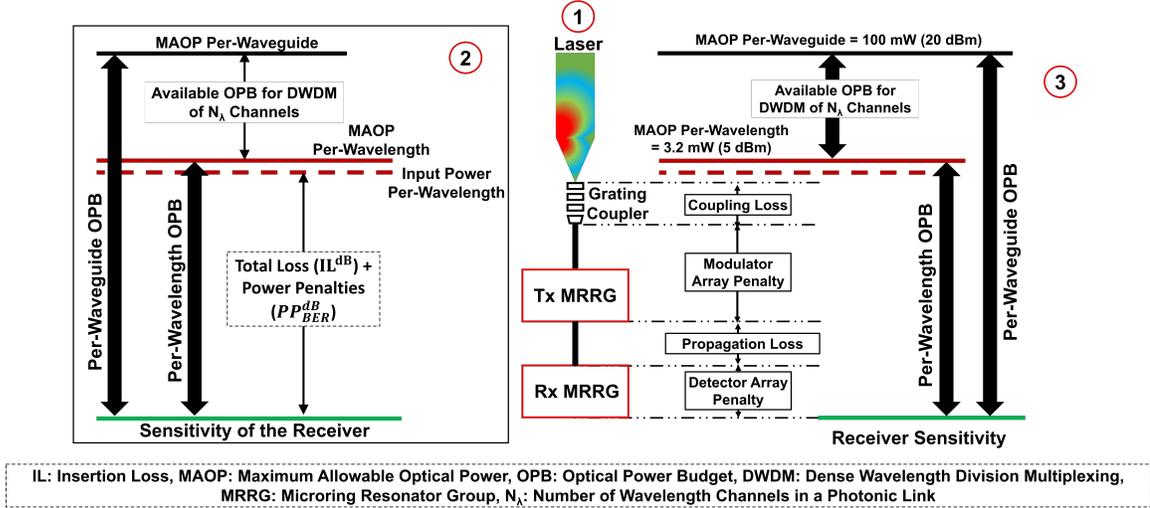


Figure 5.2: (1) Schematic of an on-SiPhI inter-chiplet link, (2), a summary of the optical power budget, and (3) evolution of the optical power budget. The laser source is assumed to be a multi-wavelength comb source, and MAOP per waveguide and MAOP per wavelength are typically forced once the light from the laser source is coupled in the link through the coupler.

sender-to-receiver data transfer takes place on a pre-established point-to-point on-SiPhI link [116, 219]. Therefore, the total energy cost of a sender-to-receiver data transfer is attributed to the data-dependent energy consumption of the link related to the dynamic switching of active photonic components of the link and the non-data-dependent consumption related to the optical and thermal stability power of the link [16]. Due to the compact size and CMOS driving voltages of MRRs, the dynamic energy portion of the total energy cost can typically be as small as a few fJ per bit [76]. This makes the optical and thermal stability power-related portion of the total energy cost the most substantial factor. To mitigate this substantial portion, a common approach has been to increase the aggregated bandwidth of data transfers on the link so as to reduce the amortized energy-per-bit cost (i.e., power per bandwidth) related to the optical and thermal power of the link [16]. Increasing the aggregated bandwidth of the link has been a common strategy even to increase the speeds of the individual sender-to-receiver data transfers on the link [20, 19]. Thus, in a nutshell, forging new techniques to increase the aggregated bandwidth of the individual links of on-SiPhI networks remains crucial to achieving high peak and average performance

values from the on-SiPhI networks.

The aggregated bandwidth of an on-SiPhI link is typically impacted by various design features across the device, circuit, and architecture layers of the hardware design stack. These design features are listed in Table 5.1. We reason using Fig. 2 and Eqs. (1) to (5) that the device-layer features such as the MAOPs, propagation + coupler losses, and FSR, among all of the features listed in Table 5.1, determine the achievable upper bound of aggregated bandwidth of an on-SiPhI link. No matter how the other circuit-layer and architecture-layer features from Table 5.1 play out for an on-SiPhI network, the achieved aggregated bandwidths of the individual sender-to-receiver data transfers across the network cannot be higher than the aggregated bandwidths of the individual links of the network that are set forth by the underlying device-layer feature values. Therefore, in this chapter, we focus on exploring the role of the MAOPs, propagation + coupler losses, and FSR in realizing on-SiPhI links with multi-terabits/second bandwidth.

It is well established in prior works that the performance (i.e., achievable N_λ and BR) of a SiPh link, whether an optical fiber based off-chip link (e.g., [19][20][15]) or a silicon waveguide based on-chip link (e.g., [125][261][252][28][193]), depends on the strict optical power budget (OPB) of the link. This holds true for the on-SiPhI waveguide based links too. In this section, we refer to Fig. 5.2 to illustrate how the OPB of an on-SiPhI link impacts its performance, i.e., its achievable N_λ and BR . As illustrated in (3) of Fig. 5.2, the OPB of a link determines the apex of the allowable optical losses and power penalties in the link. The OPB of a link has two mutually related components (see (2) in Fig. 5.2): (i) per-wavelength OPB, and (ii) per-waveguide OPB. Per-wavelength OPB determines the amount of allowable losses and power penalties for a single wavelength channel in the link, and can be defined as the difference between the per-wavelength maximum allowable optical power (MAOP) and the sensitivity of the receiver (Eq. 5.1). Similarly, per-waveguide OPB determines the amount of optical losses and power penalties allowed for all the wavelength channels in the link, and it is provided as the difference between the maximum allowable optical power (MAOP) per waveguide and the sensitivity of the receiver (Eq. 5.2). As illustrated in (2) of Fig. 5.2, the per-wavelength MAOP is restricted to 3.2 mW (5 dBm) (i.e., no more than 5 dBm optical power per wavelength is allowed). This limit has been decided upon theoretically [20][15][270] as well as empirically [147][156] to avoid the MRR modulators of the Tx MRRG from being inoperative due to the adverse impacts of optical non-linear effects such as multi-stability, self-heating, self-pulsation [67][42]. On the other hand, the per-waveguide MAOP is restricted to 100 mW (20 dBm), to avoid dramatically high optical propagation losses in on-SiPhI waveguides caused due to the increased two-photon absorption (TPA) and free-carrier absorption (FCA) [100][20][249][125].

$$OPB \text{ per Wavelength (dB)} = MAOP \text{ per Wavelength} - Receiver \text{ Sensitivity} \quad (5.1)$$

$$OPB \text{ per Waveguide (dB)} = MAOP \text{ per Waveguide} - Receiver \text{ Sensitivity} \quad (5.2)$$

Looking at the evolution of OPB provided in (2) of Fig. 5.2, a wavelength channel generated from a laser source experiences insertion loss and other power penalties as it propagates through the on-SiPhI waveguide of the link. Furthermore, insertion loss is a critical parameter that significantly impacts the aggregate bandwidth and energy efficiency of on-SiPhI links. The total insertion loss encountered by a given wavelength channel encompasses three main components: (i) the total coupling loss of the grating coupler; (ii) the waveguide propagation loss, which is the sum of the scattering loss (due to the sidewall roughness of the on-SiPh waveguide) and absorption loss (due to the material and free-carrier absorption mechanisms in the on-SiPh waveguide); and (iii) insertion loss of Tx+Rx MRRGs. It is important to note that each of these components of the insertion loss is contingent upon the specific SiPh fabrication process utilized. Section 5.3.1 delves into the specifics of each of these components, offering insights into their differences across various SiPh fabrication processes. For instance, the coupling loss pertaining to a coupler in an on-SiPhI link is ~ 1.5 dB for 45nm SOI CMOS, whereas it is significantly high at 4.9 dB for 32nm SOI CMOS. Similarly, the propagation loss in silicon waveguides is ~ 3.7 dB/cm for 45nm SOI CMOS, whereas it is 10 dB/cm for 32nm SOI CMOS. On the other hand, the power penalties experienced by a wavelength channel across the link include the modulator array penalty (i.e., the power penalty incurred due to the array of modulator MRRs of the Tx MRRG) and detector array penalty (i.e., the power penalty incurred due to the array of filter MRRs of the Rx MRRG), as shown in Fig. 5.2. The modulator array penalty consists of modulator inter-channel crosstalk [18]. Similarly, the filter array penalty consists of the total power penalty manifesting at the photodetectors due to the inter-channel crosstalk at the MRR filters and truncation of the modulated signal spectra [18]. All of these optical insertion losses (IL^{dB} in Eq. 5.3) and power penalties (PP_{BER}^{dB} in Eq. 5.3) in the link as a whole (P_{loss}^{dB} in Eq. 5.3) should amount to be less than the per-wavelength OPB (Eq. 5.4), for the link to be viable. This P_{loss}^{dB} value also whittles down a significant portion of the per-waveguide OPB to render the remaining OPB to be available for DWDM (Fig. 5.2). This outcome presents the in-equality in Eq. 5.5 as the necessary condition to accommodate N_λ wavelength channels in the link.

Therefore, for a given N_λ wavelength channels in a photonic link (N_λ), total losses and power penalties experienced by these wavelength channels should be within the optical power budget as depicted in Eq. 5.1. It is intuitive from Fig. 2 and Eq. 5.1 that to design a high bandwidth photonic link, OPB should be high and, link losses and power penalties should be low. OPB can be increased by increasing the MAOP whereas power penalties in a photonic link can be reduced by increasing the FSR. Detailed discussion regarding the impact of several link design parameters on OPB and the aggregate bandwidth of a photonic link is provided in the upcoming section.

$$P_{loss}^{dB} = PP_{BER}^{dB} + IL^{dB} \quad (5.3)$$

$$OPB \text{ per Wavelength (dB)} \geq P_{loss}^{dB} \quad (5.4)$$

$$OPB \text{ per Waveguide (dB)} \geq P_{loss}^{dB} + 10 \times \log_{10}(N_\lambda) \quad (5.5)$$

Intuitively, the bandwidth of an on-SiPhI link can be increased by increasing the $(N_\lambda \times BR)$ for the link. However, from Fig. 5.2, there should be sufficient OPB available for DWDM in the link to support such an increase in $(N_\lambda \times BR)$. But unfortunately, it is well established in prior works [20][15] that $(N_\lambda \times BR)$ in the state-of-the-art on-chip and off-chip SiPh links cannot be sufficiently increased to realize >1 Tb/s link bandwidth, due to the low values of OPB available for DWDM that is inflicted by the current, nascent state of the SiPh technology. This unpleasant shortcoming motivated us to undertake a critical thinking exercise to identify the key design pathways towards the realization of >1 Tb/s on-SiPhI links. The outcomes of this exercise are presented in the next section.

Table 5.2: Various design pathways, their target device parameters with corresponding projected values, and the likelihood of achieving the projected parameter values either in the short term (within 5 years) or long term (>5 years)

Design Pathway	Target Parameter	Projected Value	Timeline	Motivating Prior Works
Wide FSR	FSR of Filter MRR	80nm	Long Term	[140, 168]
	FSR of Modulator MRR	80nm	Short Term	[47, 35]
Minimized Loss	Propagation Loss	1 dB/cm	Long Term	[289, 73]
	Coupling Loss	1 dB	Short Term	[214, 97, 170]
Increasing MAOP	Per- λ MAOP	15 dBm (31.62 mW)	Long Term	[147]
	Per-Waveguide MAOP	Eliminate	Long Term	[273, 126, 124]

5.3 Identifying the Key Design Pathways Towards Multi-Terabit On-SiPh-Interposer Links

As per the discussion in Section 5.2.2, the device-layer features serve as key design pathways in realizing on-SiPhI links with multi-terabits/second bandwidth. The chosen design pathways, their projected values, and expected timelines along with references to the relevant prior works are provided in Table 5.2. From the discussion in Section 5.2.2, increasing the bandwidth of an on-SiPhI link requires increasing the $(N_\lambda \times BR)$ for the link, which in turn requires a sufficient increase in the 'available OPB for DWDM' in the link. From Fig. 5.2, increasing the 'available OPB for DWDM' in the link can be achieved in the following ways: (i) by decreasing the total insertion loss (IL^{dB}) in the link; (ii) by increasing the per-waveguide MAOP for a given per-wavelength input power (Fig. 5.2); (iii) by decreasing the total power penalty (PP_{BER}^{dB}) in the link.

The insertion loss (IL^{dB}) of an on-SiPhI link can be decreased by decreasing the propagation loss and coupling loss in the link (Section 5.2.2). Several optimization methods and fabrication processes pertaining to reducing the coupling loss in on-SiPhI links have been introduced in prior works [214, 97, 170, 103]. The total propagation

loss in an on-SiPhI link is the product of the waveguide length (cm) and the propagation loss constant (dB/cm). Therefore, to reduce the propagation loss in an on-SiPhI link, it is intuitive that the propagation loss constant (dB/cm) should be reduced. Another way of reducing the influence of insertion loss on the bandwidth of an on-SiPhI link is to increase the per-wavelength MAOP (Fig. 5.2). Doing so can increase the tolerance for higher total insertion loss. Increasing the per-wavelength MAOP would in turn increase the per-waveguide MAOP. All of these factors collectively can increase the 'available OPB for DWDM'.

On the other hand, the total power penalty PP_{BER}^{dB} for a link is the function of the MRR Q-factor, channel BR, Free Spectral Range (FSR), and N_λ . Prior works [125], [124], and [261] have shown that PP_{BER}^{dB} for a SiPh on-chip link can be minimized by designing the link using the optimum combination of the triplet {MRR Q-factor, channel BR, N_λ } for given FSR. This means that it is possible to minimize the increase in PP_{BER}^{dB} caused due to the intended increase in $(N_\lambda \times BR)$ by simply employing an optimal MRR Q-factor that corresponds to the increased $(N_\lambda \times BR)$. However, precisely defining the MRR Q-factor at the design time has been proven to be very difficult due to the uncertainties emanating from the unavoidable fabrication-process non-uniformity [242][269]. Moreover, the achievable operational bandwidth (i.e., the operating BR) for the MRR modulator and filter devices highly depend on the utilized device fabrication process [236][12]. Therefore, in the wake of such dependence on the fabrication process, the more practical solution for viably increasing the bandwidth of an on-SiPhI link becomes to accept the MRR Q-factor and channel BR that the utilized fabrication process provides, and then look to increase the N_λ of the link. To this end, a possible, good option for lessening the impact of increasing N_λ on PP_{BER}^{dB} would be to push for as large FSR as possible, because a large FSR renders a high spectral bandwidth available for DWDM.

Based on this discussion, we identify the following three key design pathways towards realizing >1 Tb/s on-SiPhI links, which are provided in Table 5.2.

- **Pathway 1:** increase the available OPB for DWDM in the on-SiPhI link by minimizing the insertion losses in the link;
- **Pathway 2:** increase the available OPB for DWDM in the on-SiPhI link by increasing the per-wavelength and per-waveguide MAOP limits of the link;
- **Pathway 3:** increase the spectral bandwidth available for DWDM and minimize the power penalties in the on-SiPhI link by pushing for as large FSR as possible.

The detailed discussion on each of these pathways and the considerations made for our link-level and system-level analysis are provided in the upcoming subsections.

Pathway 1: Minimize Insertion Losses

Insertion losses in an SiPh link include waveguide propagation losses and coupling losses. The amount of coupling losses incurred in on-SiPhI links depend on the utilized

fabrication process for realizing waveguides and couplers. For instance, coupling loss associated with a coupler fabricated using the 45nm SOI CMOS process amounts to 1.5dB. Likewise, a list of coupling losses pertaining to couplers created through different SiPh fabrication processes is provided in Table 5.3. Various optimization methods and fabrication processes pertaining to reducing the coupling losses in on-SiPhI links have been introduced in prior works [214, 97, 170, 103]. By utilizing these, the coupling losses per on-SiPhI link can be reduced to as low as ~ 1 dB with a relatively shorter timeline, as shown in Table 5.2.

Propagation losses in a silicon waveguide comprises of absorption losses and scattering losses. **Absorption losses:** Silicon waveguides operating at wavelengths ranging from 1500-1600 nm are prone to high absorption losses due to strong two-photon absorption (TPA), despite their moderate-to-low material absorption losses in this wavelength range. This is because, for DWDM applications, when multiple wavelengths are coupled into a silicon waveguide, the total optical power in the waveguide increases which in turn induces TPA effect in the silicon waveguide [100]. Due to TPA, free carrier concentration in the silicon waveguide increases that induces free-carrier absorption (FCA) [100][249] effect, which consequently increases the absorption losses in the silicon waveguide. **Scattering losses:** Silicon waveguides are also prone to high scattering losses mainly due to the following reasons. First, sidewall roughness of the waveguides arising from fabrication imperfections. Second, high index contrast between the core (silicon) and cladding (silicon dioxide) of the waveguides. Due to the high index contrast between the core (silicon) and cladding (silicon dioxide) of a silicon waveguide, the interaction of the guided optical mode with the rough sidewalls of the waveguide increases. This enhanced mode-roughness interaction increases the scattering losses in silicon waveguides.

Therefore, high absorption and scattering losses give rise to high propagation loss in silicon waveguides. Moreover, the propagation loss observed in silicon waveguides is contingent upon the particular SiPh fabrication process utilized. For example, silicon waveguides fabricated using the 45nm SOI CMOS process exhibit a propagation loss of 3.7 dB/cm, as provided in Table 5.3. Similarly, a list of propagation losses in silicon waveguides across various SiPh fabrication processes is also provided in Table 5.3. Furthermore, increase in propagation loss increases the amount of insertion loss present in the link. This increase in insertion loss whittles down the OPB restricting the aggregate bandwidth of photonic links. Prior works have demonstrated new photonic platforms for which TPA is absent [124][273][201], and such platforms can render decreased waveguide propagation losses. On the other hand, some prior works have reported silicon waveguide propagation losses below 1 dB/cm [106, 75]. The type of waveguide demonstrated in these prior works [106, 75] is a ridge waveguide, in which the interaction of guided mode with the sidewalls of the waveguide is low, thereby reducing the scattering losses. However, ridge waveguides are not compatible to couple with MRRs for cascaded DWDM. In contrast, channel waveguides are compatible for cascaded DWDM, but the lowest reported propagation loss for channel waveguides is greater than 2 dB/cm [289][73]. Therefore, achieving a propagation loss of 1 dB/cm for ridge waveguides represents a short-term solution, while achieving the same level of loss for channel waveguides is considered a long-term goal (Table 5.2).

From the above discussion, it is clear that *reducing the propagation loss to 1 dB/cm and coupling loss to 1 dB* is the most optimistic goal for the near future. Therefore, we have chosen these loss values for our analysis in this chapter.

Pathway 2: Increase Per-Wavelength and Per-Waveguide MAOP Limits

Per-waveguide MAOP: As discussed in Section 5.2.2, the per-waveguide MAOP limit manifests in a rectilinear on-SiPhI waveguide due to the presence of very high absorption losses at relatively high optical power density and large number of multiplexed wavelength channels in the waveguide. Such high absorption losses are caused in a DWDM based silicon waveguide due to the strong two-photon absorption (TPA) and four-wave mixing nonlinearities of the silicon material in the optical C-band of operation [33][136]. Due to the TPA effect, the free-carrier concentration in a silicon rectilinear waveguide can dramatically increase for the input optical power densities of greater than $1 \text{ W}/\mu\text{m}^2$ (corresponds to 100 mW (20 dBm) optical power in the waveguide with the cross-sectional waveguide dimensions of $520 \text{ nm} \times 220 \text{ nm}$ [136]), which consequently triggers free-carrier absorption (FCA) related very high propagation losses that can amount to up to 1 dB/cm additional loss per added multiplexed channel in the waveguide [136]. To avoid such high, power-dependent propagation losses in the waveguide, prior works limit the MAOP per waveguide to be 100 mW (20 dBm) [100][270]. Clearly, the introduction of the per-waveguide MAOP limit caps the available OPB for DWDM (Fig. 5.2), which in turn limits the achievable increase in N_λ and link bandwidth. Therefore, we can intuitively argue that the opportunities for increasing the available OPB for DWDM can be improved by increasing, or even virtually eliminating the per-waveguide MAOP limit. Prior work [182] has shown that such optical power-dependent losses are not present in silicon nitride waveguides, but due to the lack of active devices in silicon nitride material platform [273], silicon nitride waveguides are not yet commonly used in the mainstream SiPh designs. Alternatively, another prior work [124] has shown that the per-waveguide MAOP limit can be increased, or even be virtually eliminated, by designing SiPh links that can operate at relatively long wavelengths around $4\mu\text{m}$. At such long wavelengths, silicon's band gap energy is more than the energy of 2 photons, and hence, the TPA effect is absent, eliminating the optical power-dependent dramatic increase in waveguide propagation losses. Leveraging these benefits however requires adopting a new SiPh fabrication material system, referred to as silicon-on-sapphire (SOS) [124]. Although it is not clear yet if, how, and by when the SOS-based SiPh designs will replace the SOI-based SiPh designs, it is worth asking this question nevertheless: Can eliminating the per-waveguide MAOP limit in on-SiPhI links boost their bandwidth beyond 1 Tb/s? To find the answer to this question, we aim to eliminate the per-waveguide MAOP, and hence, per-waveguide OPB as part of this design pathway, for the link-level and system-level bandwidth and performance analysis. Eliminating per-waveguide MAOP is a relatively long-term solution (Table 5.2).

Per-wavelength MAOP: On the other hand, the cause for the per-wavelength MAOP limit is the interplay of the mutually conflicting free-carrier dispersion and thermal dispersion phenomena in MRR modulators that renders the modulators in-

operable for per-wavelength input optical power of greater than the MAOP limit [67][147]. Evidently, this interplay is exacerbated due to the strong TPA effect and high intra-cavity power buildup present in the silicon MRR modulators [67][147]. Nevertheless, the MRR modulators can be intelligently designed to balance the interplay of these conflicting phenomena [156], to consequently increase the per-wavelength MAOP limit to 5 mW (7 dBm) [147] (which is greater than 3,2 mW (5 dBm), as commonly assumed in several link- and system-level prior works [20][270][125]). This outcome encourages the efforts focused on eliminating the TPA effect from MRR modulators, in hopes of further increasing the per-wavelength MAOP limit to consequently increase the per-wavelength OPB (Section 5.2.2; Fig. 5.2). However, since it is not yet clear how much eliminating the TPA effect would impact the per-wavelength MAOP limit, we assume a relatively optimistic value of 31.5 mW (15 dBm) for the per-wavelength MAOP limit as part of this design pathway, which is considered a long-term solution (Table 5.2). This assumption guides us in our quest to answer the following question: Can eliminating the per-wavelength MAOP limit in long on-SiPhI links (about 10 cm long) boost their bandwidth beyond 1 Tb/s?

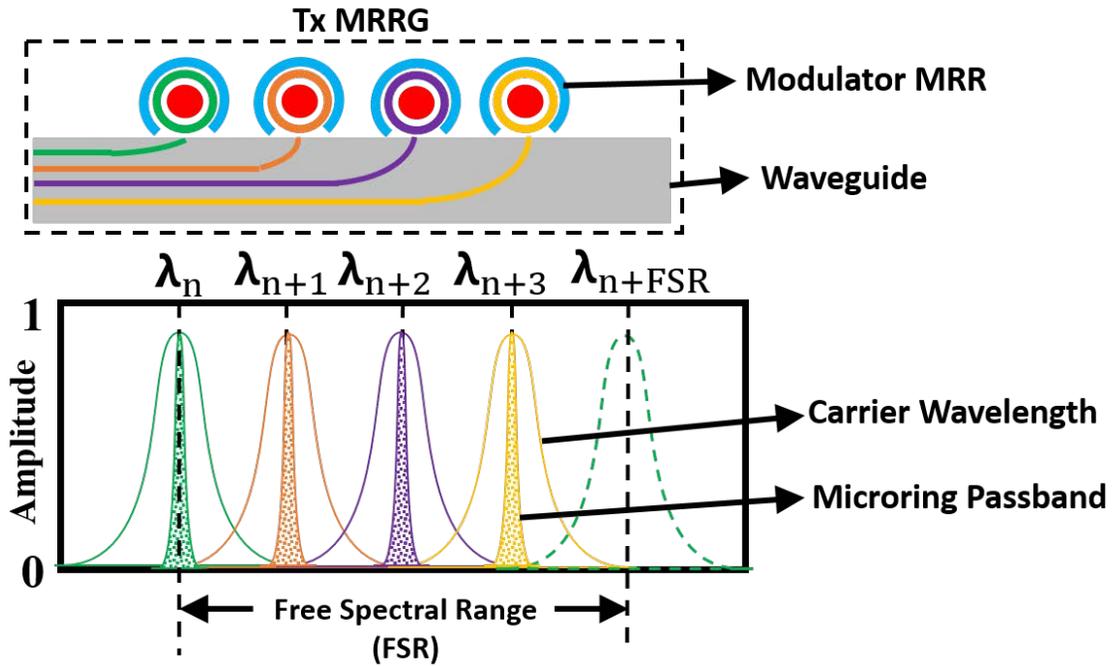


Figure 5.3: Illustration of FSR of an MRR.

Pathway 3: Push for as Wide FSR as Possible

MRR, which is considered as the workhorse of a photonic link, is a looped waveguide in which the resonance occurs when the optical path length of the MRR is exactly a whole number of wavelengths. Therefore, MRRs support multiple resonances and the spacing between these resonances is FSR as shown in Fig. 5.3. Low values of FSR means for a given number of wavelength channels in a photonic link (N_λ), spacing

between the adjacent channels is low resulting in inter-channel crosstalk [18] which in turn increases the PP_{BER}^{dB} (Section 5.2.2). Prior works have demonstrated that low FSR of MRR devices in SOI photonic links [20] restricts the aggregate bandwidth to $< 1\text{Tb/s}$ because of this increase in PP_{BER}^{dB} . Hence, it is important to enhance the FSR of constituent MRR devices to achieve the aggregate bandwidth of $\geq 1\text{Tb/s}$.

FSR of an MRR is inversely proportional to its round-trip optical length. Therefore, to widen the FSR, one way is to reduce the round-trip optical length of the MRR which would result in a compact size of the MRR. But this length cannot be infinitely reduced due to various reasons. Firstly, reducing the round-trip optical length of an MRR increases the complexity of implementing the MRR tuning mechanism. Secondly, due to the shorter coupling length, the efficient coupling between the bus waveguide and the MRR becomes difficult to realize. Finally, reducing the round-trip optical length often results in sharper bend radius, which causes extra radiation losses and scattering losses in the MRR due to the guided optical mode that overlaps with and extends beyond the rough outer wall of the MRR bend.

Alternatively, prior works have demonstrated various designs of MRR filters that can support larger FSR. Most recently, FSR-free MRR filter architectures were also demonstrated. Among the designs of MRR filters that support large FSR, Li Ang et al. in [140] demonstrated a novel method that widens the FSR by means of internal reflections inside the MRR. No extra optical loss is introduced and a wide FSR up to 150 nm can be achieved using this method. Similar design has also been demonstrated in [258] that supports FSR up to 175 nm. On the other hand, FSR-free MRR filter architectures demonstrated so far in the literature are based on either integrating the contra-directional couplers (CDCs) with the MRR or by cascading MRRs with different FSRs (popularly known as vernier scheme [90]). Eid. N et al. in [80] demonstrated FSR free MRR filters based on partially wrapping the contra-directional couplers (CDC) around the MRR. This design significantly suppresses the side-modes of the MRR resulting in FSR free response. Another similar type of FSR free MRR filter design has also been demonstrated in [164] which is based on integrating the bent CDCs into the through port coupling region of the MRR cavity which suppresses all the modes except the resonance mode of the cavity. An FSR-free MRR filter architecture based on vernier scheme is demonstrated in [168], which is polarization diverse and can be tuned beyond the range of C-band. This design of FSR-free MRR filter based on vernier scheme is CMOS compatible, making it easier to fabricate compared to other FSR-free MRR filter designs demonstrated so far. Another FSR-free MRR filter based on vernier scheme is demonstrated in [196]. An FSR free MRR filter using photonic crystal cavities was also demonstrated in [295].

Although, prior works have demonstrated MRR filters that virtually eliminate the FSR, the off-chip comb laser sources employed with on-SiPhI links, demonstrated so far [85, 281, 205, 235, 129], cannot provide consistently high optical power at every wavelength for a wide range of wavelengths. Based on what is known from these prior works, comb laser sources can consistently provide ≥ 15 dBm of optical power per wavelength for up to 80 nm range only around the C and L bands. This limitation of comb laser sources curtails the available spectral bandwidth for DWDM, which in turn has the effect of having a limited FSR because a limited FSR also curtails the

available spectral bandwidth for DWDM. Therefore, we consider the achievable FSR to be 80nm, which is a relatively long-term solution (Table 5.2).

Pathfinding analysis

Table 5.2 lists the identified design pathways and their corresponding target parameters which were discussed in previous subsections. From previous subsections, it is clear that the feasible solution for viably increasing the bandwidth of an on-SiPhI link is to accept the MRR Q-factor and channel BR that the utilized fabrication process provides, and then look to increase the N_λ of the link. However, the Q-factor and channel BR varies across different fabrication platforms. Therefore, we consider three established SiPh fabrication platforms from prior work namely 45nm SOI CMOS [236], 32nm SOI CMOS [236] and Deposited poly-Si [12] for our pathfinding analysis. Table 5.3 lists the design parameters corresponding to these fabrication platforms. The parameters listed in Table 5.3, corresponding to each platform, do not corroborate with our intended design pathway targets (Table 5.2). Hence, we have derived eight different variants of on-SiPhI inter-chiplet links, in which seven variants are derived based on our identified design pathways (Table 5.2) and one variant is derived based on the parameters innate to fabrication platforms (Table 5.3). Each of these variants are listed below:

1. **Fabrication_Platform_Name** + *Vanilla* - This variant utilizes innate design parameters corresponding to each of the considered fabrication platforms (Table 5.3)
2. **Fabrication_Platform_Name** + *Minimized Loss* - This variant employs innate design parameters corresponding to each platform except the insertion loss parameters, which are replaced with target parameters of our *Minimized Loss* design pathway (Table 5.2)
3. **Fabrication_Platform_Name** + *Wide FSR* - This variant avails innate design parameters corresponding to each platform except the FSR parameter, which is replaced with target parameter of our *Wide FSR* design pathway (Table 5.2)
4. **Fabrication_Platform_Name** + *Increased MAOP* - This variant utilizes innate design parameters corresponding to each platform except the MAOP parameters, which are replaced with target parameters of our *Increased MAOP* design pathway (Table 5.2)
5. **Fabrication_Platform_Name** + (*Minimized Loss* + *Wide FSR*) - This variant employs innate design parameters corresponding to each platform except the insertion loss and FSR parameters, which are replaced with target parameters of our *Minimized Loss* and *Wide FSR* design pathways (Table 5.2)
6. **Fabrication_Platform_Name** + (*Minimized Loss* + *Increased MAOP*) - This variant avails innate design parameters corresponding to each platform

but replaces the insertion loss and MAOP parameters with the target parameters of our *Minimized Loss* and *Increased MAOP* design pathways (Table 5.2)

7. **Fabrication_Platform_Name** + (**Wide FSR** + **Increased MAOP**) - This variant employs innate design parameters corresponding to each platform but replaces the wide FSR and MAOP parameters with the target parameters of our *Wide FSR* and *Increased MAOP* design pathways (Table 5.2)
8. **Fabrication_Platform_Name** + (**Minimized Loss** + **Wide FSR** + **Increased MAOP**) - This variant employs innate design parameters corresponding to each platform but replaces the insertion loss, wide FSR and MAOP parameters with the target parameters of our *Minimized Loss*, *Wide FSR* and *Increased MAOP* design pathways (Table 5.2)

Replacing the **Fabrication_Platform_Name** with the 45nm SOI CMOS [236], 32nm SOI CMOS [236] and Deposited poly-Si [12] platforms in the above list of variants, makes it a total of twenty four variants (eight variants corresponding to each platform). Detailed link-level and system-level analysis of each of these variants is provided in upcoming sections.

Table 5.3: Design Parameters for our considered SiPh fabrication processes

Design Parameters	45nm SOI CMOS [236]	32nm SOI CMOS [236]	Deposited Poly-Si [12]
Modulator MRRs Q	10000	6000	5000
Filter MRRs Q	8500	6500	5000
MRR Radius	5 μm	5 μm	7.5 μm
Operating-wavelength	1290 nm	1310 nm	1300 nm
FSR	12.6 nm	13 nm	8.54 nm
Modulator Bandwidth	13 GHz	13.5 GHz	16.8 GHz
Detector Bandwidth	5 GHz	12.5 GHz	11 GHz
Sensitivity (dBm)	-17.645	-11.79	-20.414
Propagation Loss	3.7 dB/cm	10 dB/cm	20 dB/cm
MAOP (per-wavelength)	1.7 mW (2.3 dBm)	2.5 mW (4 dBm)	2.8 mW (4.5 dBm)
MAOP (per-waveguide)	100 mW(20 dBm)	100 mW (20 dBm)	100 mW (20 dBm)
Per-coupler Loss	1.5 dB	4.9 dB	5.2 dB
Bit-rate	12 Gb/s	12.5 Gb/s	11 Gb/s
Per-wavelength Budget	19.945 dB	15.794 dB	24.914 dB
Per-waveguide Budget	37.645 dB	31.79 dB	40.414 dB
Waveguide Length	1-10 cm	1-10 cm	1-10 cm
Modulator IL	4.7 dB	2.8 dB	3.8 dB
Filter IL	0.18 dB	0.14 dB	0.11 dB
Coupling Loss	1.5 dB	4.9 dB	5.2 dB

5.4 Link-level evaluation

Evaluation Setup

To perform the pathfinding link-level analysis for each of the 24 derived variants (Section 5.3.4), we utilize a search heuristic based optimization framework provided in

[252]. This search heuristic consists of an error function that takes different values of N_λ and channel BR as input and evaluates an error value for each duplet of (N_λ, BR) , for a given waveguide link length. From that, the duplet of (N_λ, BR) corresponding to minimum positive value of error function is chosen as the optimal duplet since minimum positive value of error-function means the available OPB has been utilized to its maximum while satisfying the condition given in Eq. 5.5. With the obtained (N_λ, BR) duplet for each derived variant, we have calculated corresponding aggregate bandwidth which is the product of N_λ and channel BR, and energy per bit (EPB) which is sum of link laser power, thermal tuning power, modulator driver power and receiver power [27]. The results of this analysis and a detailed discussion is provided in the next subsection.

Results and Comparison

Fig. 5.4 illustrates the evaluated aggregate bandwidth (primary Y-axis) and EPB (secondary Y-axis) for different on-SiPhI inter-chiplet variants corresponding to three different SiPh fabrication platforms namely 45nm SOI CMOS [236], 32nm SOI CMOS [236] and deposited poly-Si [12], for different waveguide lengths ranging from 1 cm to 10 cm (X-axis). Based on the results obtained from this analysis, we have categorized the derived variants in to two types namely non-viable and viable variants. Non-viable variants are the variants that are unfeasible to implement as on-SiPhI inter-chiplet links due to the high insertion losses at longer waveguide lengths that exceed the amount of available OPB in the link, thereby not supporting any wavelength channels in the link and yielding no aggregate bandwidth. On the other hand, viable variants are the variants that are feasible to implement as on-SiPhI inter-chiplet links since they support some tangible aggregate bandwidth for waveguide link lengths of up to 10 cm. Detailed discussion on each category of variants is provided in next subsections.

Non-Viable Variants

Among the derived variants, *Vanilla* (Fig. 5.4(a)), *Wide FSR* (Fig. 5.4(c)), *Increased MAOP* (Fig. 5.4(d)) and *Wide FSR + Increased MAOP* (Fig. 5.4(g)) variants are considered as non-viable variants because they support no wavelength channels and therefore do not support aggregate bandwidth for longer waveguide lengths.

Among the non-viable variants, *Vanilla* variants corresponding to 32nm SOI CMOS and deposited poly-Si platforms do not support any wavelength channels due to high waveguide propagation loss of ~ 10 dB/cm and ~ 20 dB/cm respectively (Table 5.3) resulting in excess amount of insertion loss in the link. But the *Vanilla* variant corresponding to 45nm SOI CMOS platform can support wavelengths up to a waveguide length of 4 cm due to low insertion loss (3.7 dB/cm (Table 5.3)) compared to the other *Vanilla* variants. However, the aggregate bandwidth and EPB of this variant is limited to 744 Gb/s and 1.34 pJ/bit respectively. Therefore, it is intuitive that reducing the insertion loss is vital in realizing longer on-SiPhI inter-chiplet links.

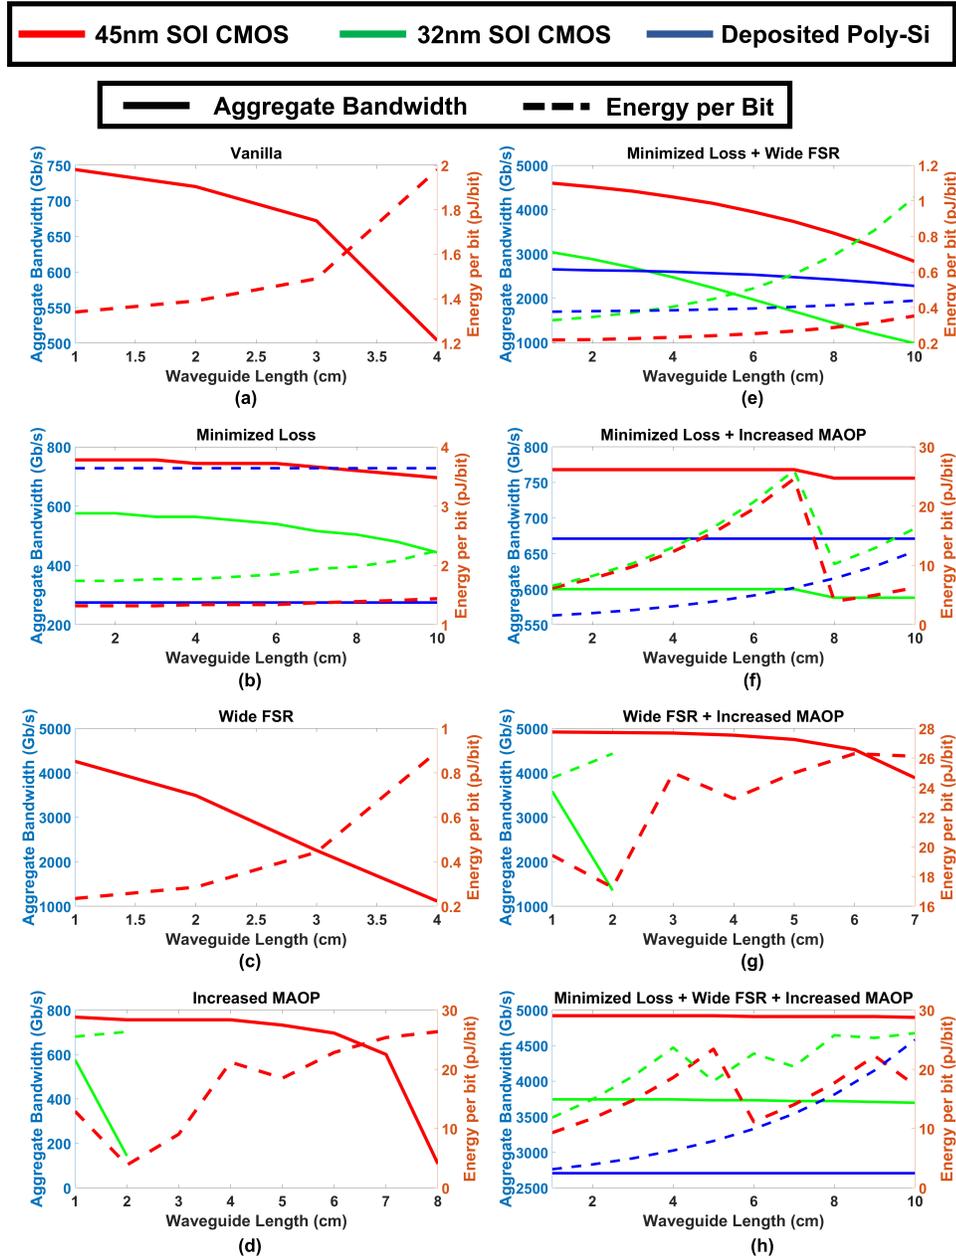


Figure 5.4: Aggregate bandwidth and energy per bit (EPB) values for different waveguide lengths ranging from 1 cm to 10 cm obtained from the analysis performed on (a) Vanilla, (b) Minimized Loss, (c) Wide FSR, (d) Increased MAOP, (e) Minimized loss + Wide FSR, (f) Minimized loss + Increased MAOP, (g) Wide FSR + Increased MAOP, and (h) Minimized Loss + Increased MAOP + Wide FSR on-SiPh variants derived from 45nm SOI CMOS [236], 32nm SOI CMOS [236] and deposited poly-Si [12] platforms.

Similarly, *Wide FSR* variants corresponding to 32nm SOI CMOS and deposited poly-Si platforms do not support any wavelength channels whereas *Wide FSR* variant corresponding to 45nm SOI CMOS platform can support wavelength channels up to a waveguide length of 4 cm with peak aggregate bandwidth of 4.3 Tb/s and corresponding EPB of 0.235 pJ/bit, and a minimum aggregate bandwidth of 1.12 Tb/s with corresponding EPB of 0.896 pJ/bit. Therefore, it is intuitive that Widening the FSR will increase the spacing between the wavelength channels in the link resulting in low power penalty in the link and thereby increasing the available OPB for DWDM. However, the presence of high insertion losses in the link is still an impediment in realizing longer on-SiPhI inter-chiplet links. Therefore, it is lucid that implementing the *Wide FSR* design pathway in combination with *Minimized Loss* design pathway will aid in realizing longer on-SiPh links with superior aggregate bandwidth and energy efficiency.

Also, *Increased MAOP* variant corresponding to deposited poly-Si platform does not support any aggregate bandwidth whereas the same variant corresponding to 45nm SOI CMOS and 32nm SOI CMOS platforms can realize on-SiPhI links up to a waveguide length of 8 cm and 2 cm respectively. In terms of aggregate bandwidth and energy efficiency, *Increased MAOP* variant corresponding to 45nm SOI CMOS platform achieve peak aggregate bandwidth of 768 Gb/s with corresponding EPB of 12.9 pJ/bit and a minimum aggregate bandwidth of 108 Gb/s with corresponding EPB of 26.34 pJ/bit whereas the same variant corresponding to 32nm SOI CMOS platform achieves peak aggregate bandwidth of 576 Gb/s with corresponding EPB of 25.51 pJ/bit and least aggregate bandwidth of 144 Gb/s with corresponding EPB of 26.32 pJ/bit. Therefore, it is intuitive that implementing the *Increased MAOP* design pathway in combination with any other design pathways, especially the *Minimized Loss* design pathway, will enable these variants to realize longer on-SiPhI links with higher aggregate bandwidth and energy efficiency.

Wide FSR + Increased MAOP variant corresponding to 45nm SOI CMOS and 32nm SOI CMOS platforms can realize on-SiPhI links up to waveguide length of 7 cm and 2 cm respectively whereas the same variant corresponding to deposited poly-Si platform does not support any wavelength channels. In terms of performance, *Wide FSR + Increased MAOP* variant corresponding to 45nm SOI CMOS platform achieves peak aggregate bandwidth of 4.92 Tb/s with corresponding EPB of 19.44 pJ/bit and a minimum aggregate bandwidth of 3.88 Tb/s with corresponding EPB of 26.13 pJ/bit whereas the same variant corresponding to 32nm SOI CMOS platform achieves peak aggregate bandwidth of 3.6 Tb/s with corresponding EPB of 24.7 pJ/bit and a minimum aggregate bandwidth of 1.4 Tb/s with corresponding EPB of 26.3 pJ/bit. Clearly, multi Tb/s aggregate bandwidth can be achieved by widening the FSR in combination with increasing the MAOP but the presence of high insertion loss in the link makes it unfeasible to realize longer on-SiPhI inter-chiplet links.

Therefore, implementing the *Minimized Loss* design pathway in combination with other design pathways is the key to realizing longer on-SiPhI inter-chiplet links with >1 Tb/s aggregate bandwidth and <1 pJ/bit energy efficiency. In addition, using repeaters can also make *Vanilla*, *Wide FSR* and *Wide FSR + Increased MAOP* variants corresponding to 45nm SOI CMOS platform, and *Wide FSR + Increased MAOP* vari-

ant corresponding to 32nm SOI CMOS platform viable for longer waveguide lengths.

Viable Variants

Among the derived variants, *Minimized Loss* (Fig. 5.4(b)), *Minimized Loss + Wide FSR* (Fig. 5.4(e)), *Minimized Loss + Increased MAOP* (Fig. 5.4(f)) and *Minimized Loss + Wide FSR + Increased MAOP* (Fig. 5.4(h)) variants are considered as viable variants to implement on-SiPhI inter-chiplet links since they support wavelength channels upto waveguide length as long as 10 cm.

Among these viable variants, *Minimized Loss* variant corresponding to 45nm SOI CMOS platform achieves peak aggregate bandwidth of 756 Gb/s with corresponding EPB of 1.32 pJ/bit and a minimum aggregate bandwidth of 696 Gb/s with corresponding EPB of 1.4 pJ/bit whereas the same variant corresponding to 32nm SOI CMOS platform achieves peak aggregate bandwidth of 576 Gb/s with corresponding EPB of 1.74 pJ/bit and a minimum aggregate bandwidth of 444 Gb/s with corresponding EPB of 2.25 pJ/bit. Similarly, *Minimized Loss* variant corresponding to deposited poly-Si platform supports peak aggregate bandwidth of 275 Gb/s with corresponding EPB of 3.64 pJ/bit. Hence, it is evident that minimizing the insertion loss in the link will enable the variants to realize on-SiPhI links up to a waveguide length as long as 10 cm. However, these variants do not achieve aggregate bandwidth of more than 1Tb/s and EPB of less than 1 pJ/bit due to high power penalty in the link resulting from the low FSR of the considered SiPh fabrication platforms and also due to low MAOP in the link resulting in less available OPB. Therefore, minimizing the insertion loss in combination with other design pathways is vital to yield extremely high aggregate bandwidth and energy-efficient on-SiPhI inter-chiplet links which is the most important step towards enabling the chiplet based systems for the future.

As illustrated in Fig. 5.4, *Minimized Loss + Increased MAOP* variant corresponding to 45nm SOI CMOS platform achieves peak aggregate bandwidth of 768 Gb/s with corresponding EPB of 6.18 pJ/bit and a minimum aggregate bandwidth of 756 Gb/s with corresponding EPB of 6.26 pJ/bit whereas the same variant corresponding to 32nm SOI CMOS platform achieves peak aggregate bandwidth of 600 Gb/s with corresponding EPB of 6.54 pJ/bit and a minimum aggregate bandwidth of 588 Gb/s with corresponding EPB of 16.25 pJ/bit. Similarly, *Minimized Loss + Increased MAOP* variant corresponding to deposited poly-Si platform yields peak aggregate bandwidth of 671 Gb/s with corresponding EPB of 1.56 pJ/bit. Here, minimizing the loss in combination with increasing the MAOP will enable the variants to support higher aggregate bandwidth compared to the *Minimized Loss* variants discussed previously. But the amount of power penalty in the link is still high and increase in MAOP along with reducing the insertion loss does not offset that penalty. Therefore, it is lucid from the observations that minimizing the loss in combination with enhancing the FSR or implementing all three design pathways in combination will enable the on-SiPhI inter-chiplet links to achieve higher aggregate bandwidth and energy-efficiency.

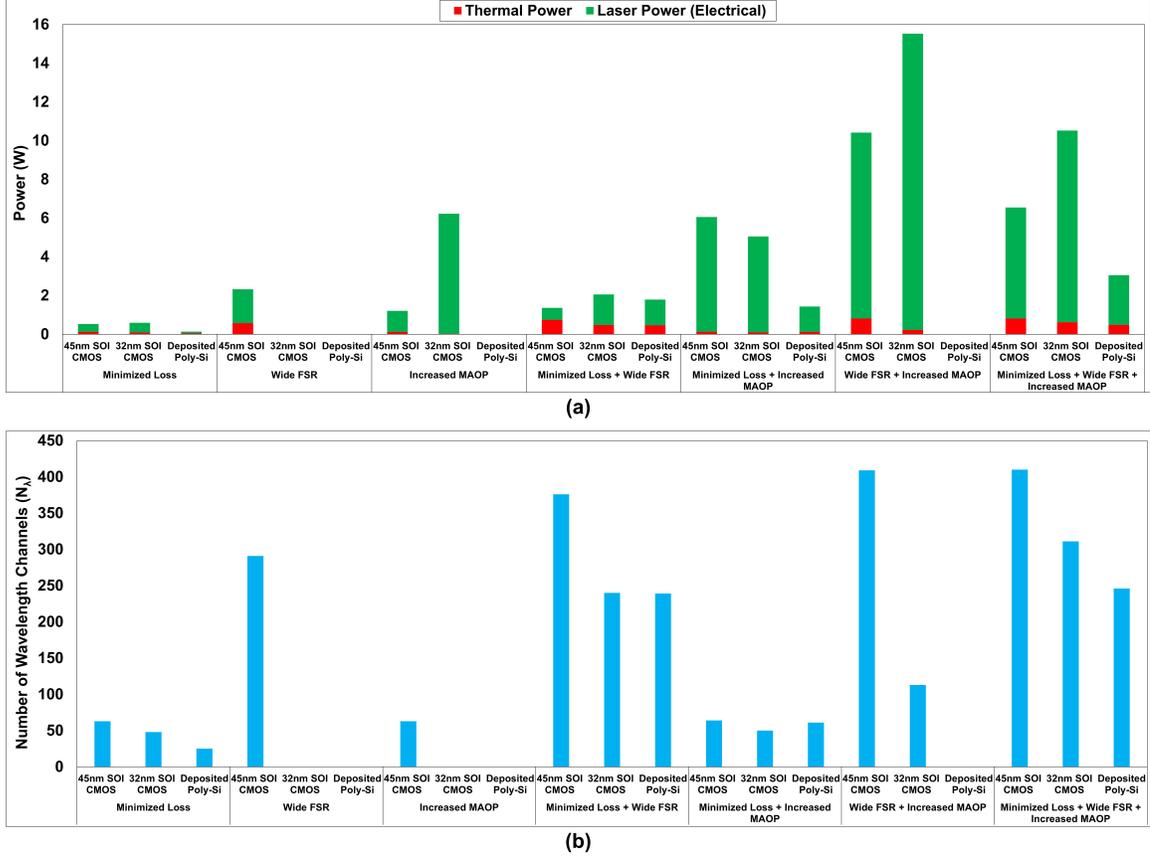


Figure 5.5: (a) Breakdown of wall-plug laser power and thermal power, and (b) the number of wavelength channels supported by different single-waveguide links of 2 cm length corresponding to various design pathways and fabrication platforms

As we can depict from Fig. 5.4(e) and Fig. 5.4(h), *Minimized Loss + Wide FSR* and *Minimized Loss + Wide FSR + Increased MAOP* variants achieve more than 1Tb/s aggregate bandwidth up to waveguide length of 10cm, for all the considered SiPh fabrication platforms, in which former variant corresponding to 45nm SOI CMOS platform achieves peak aggregate bandwidth of 4.6 Tb/s with corresponding EPB of 0.218 pJ/bit whereas the latter variant corresponding to the same fabrication platform achieves peak aggregate bandwidth of 4.92 Tb/s with corresponding EPB of 9.2 pJ/bit.

Comparing the evaluated aggregate bandwidth and EPB obtained from the analysis performed on the derived on-SiPhI inter-chiplet variants, we can deduce that for different waveguide lengths ranging from 1 cm to 10cm, *Minimized Loss + Wide FSR + Increased MAOP* variant corresponding to 45nm SOI CMOS platform achieves highest aggregate bandwidth whereas *Minimized Loss + Wide FSR* variant corresponding to the same platform achieves lowest EPB. The *Minimized Loss + Wide FSR* variant corresponding to 45nm SOI CMOS has low crosstalk and signal truncation penalty due to high FSR, and optimum values of modulator and detector Q. Therefore, due to low insertion loss and low power penalty, *Minimized Loss + Wide*

FSR variant corresponding to 45nm SOI CMOS platform achieves lowest EPB among all the variants for waveguide lengths up to 10 cm. But this variant falls short of achieving highest aggregate bandwidth due to low OPB in the link resulting from low MAOP. On the other hand, *Minimized Loss + Wide FSR + Increased MAOP* variant corresponding to 45nm SOI CMOS platform also has low insertion loss and increasing the MAOP per-wavelength enables this variant to accommodate higher number of wavelength channels which in turns enables it to achieve highest aggregate bandwidth among all the variants. However, higher number of wavelength channels in the link i.e., the high degree of DWDM leads to less channel spacing which in turn will increase the crosstalk penalty in the link resulting in higher EPB consumption. Therefore, designing on-SiPhI inter-chiplet links by implementing all three design pathways using 45nm SOI CMOS platform can achieve peak aggregate bandwidth of 4.92 Tb/s whereas EPB <1 pJ/bit with corresponding aggregate bandwidth of 4.6 Tb/s can be achieved by implementing the *Minimized Loss* design pathway in combination with *Wide FSR* for the same fabrication platform.

Fig. 5.5(a) provides a breakdown of the wall-plug laser power and thermal power for various on-SiPhI inter-chiplet variants across different design pathways and fabrication platforms (45 nm SOI CMOS, 32 nm SOI CMOS, and Deposited Poly-Si platforms) for a single-waveguide link of 2 cm length. Fig. 5.5(b) displays the number of wavelength channels (N_λ) supported by these variants. Notably, the on-SiPhI inter-chiplet link variants corresponding to 32nm SOI CMOS and Deposited Poly-Si platforms exhibit higher laser power consumption than 45 nm SOI CMOS platform due to the high propagation and coupling losses for 32 nm SOI CMOS and Deposited Poly-Si platforms (Table 5.3). This higher laser power consumption poses a real limitation because reducing it by decreasing the propagation and coupling losses may be extremely difficult if not impossible. This is because of the very low optical isolation available in the 32 nm SOI platform due to the very thin silicon layer on 32 nm SOI wafers, and the very high absorption and scattering losses present in Deposited Poly-Si substrates due to the imperfect poly-Si crystals with high number of granular boundaries. This limitation persists even when laser power is normalized with bandwidth, resulting in higher energy-per-bit (EPB) for variants corresponding to 32nm SOI CMOS and Deposited Poly-Si platforms. For instance, minimized loss + increased MAOP variants corresponding to 32nm SOI CMOS and Deposited Poly-Si platforms consume EPB of 0.35 pJ/bit and 0.38 pJ/bit respectively. In contrast, the same variant on the 45nm SOI CMOS platform consumes only 0.2 pJ/bit. In contrast, the 45 nm SOI CMOS variants have higher thermal power consumption compared to the other platforms. This is because the 45 nm SOI CMOS variants support larger N_λ values (Fig. 5.5(b)) due to the larger OPB remaining available for wavelength multiplexing owing to the lower optical losses. A larger N_λ means a higher number of MRRs per link, requiring higher total thermal stability power consumption per link. For instance, from Fig. 5.5(b), the wide FSR + increased MAOP variant corresponding to 45nm SOI CMOS platform has larger N_λ , and therefore, it consumes thermal EPB of 0.21 pJ/bit (Fig. 5.5(a)) that is higher than the thermal EPB of ~ 0.14 pJ/bit consumed by the same variant corresponding to 32nm SOI CMOS platform.

From the above observations, we can notice that in order to design longer on-SiPhI

Table 5.4: Inter-chiplet variants derived from 45nm SOI CMOS, 32nm SOI CMOS and Deposited poly-Si platforms.

VR: Viable with repeaters, V: Viable, NV: Non-viable ¹

Variants	45nm SOI CMOS			32nm SOI CMOS			Deposited Poly-Si		
		(N _λ ,BR)	ADR (Gb/s)		(N _λ ,BR)	ADR (Gb/s)		(N _λ ,BR)	ADR (Gb/s)
Vanilla	VR	(42, 12)	504	NV			NV		
Minimized Loss	V	(60, 12)	720	V	(42, 12)	504	V	(25, 11)	275
Wide FSR	VR	(93, 12)	1116	NV			NV		
Increased MAOP	V	(9, 12)	108	VR	(12, 12)	144	NV		
Minimized Loss + Wide FSR	V	(289, 12)	3468	V	(120, 12)	1440	V	(220, 11)	2420
Minimized Loss + Increased MAOP	V	(63, 12)	756	V	(49, 12)	588	V	(61, 11)	671
Wide FSR + Increased MAOP	VR	(404, 12)	4848	VR	(113, 12)	1356	NV		
Minimized Loss + Wide FSR + Increased MAOP	V	(409, 12)	4908	V	(310, 12)	3720	V	(246, 11)	2706

inter-chiplet links for the future, it is vital to keep the insertion loss to minimum. Similarly, we can also infer that combining the other design pathways such as *Wide FSR* and *Increased MAOP* with *Minimized Loss* can scale the aggregate bandwidth to more than 1Tb/s which is the most important step towards meeting the bandwidth requirements of future chiplet-based computing systems. However, it is important to see how these variants perform at system-level. Therefore, we perform a system-level analysis by implementing the derived variants on a CPU based multi-core multi-chiplet architecture and a GPU based multi-chiplet module (MCM) considered from prior work ([29][127]). Details of this analysis are provided in the next section.

5.5 System-Level Evaluation

CPU based multi-core multi-chiplet architecture

We have performed system-level analysis on a CPU based multi-core multi-chiplet architecture named NUPLet [29] and on a GPU based MCM from [127]. The architecture, inter-chiplet network of the NUPLet and the design of GPU based MCM are described in following subsections.

Architecture of NUPLet

The architecture of NUPLet employs photonic links to facilitate both inter-chiplet and intra-chiplet communication. This is because traditional electrical interconnects are constrained by inherent limitations such as signal degradation, crosstalk, and susceptibility to electromagnetic interference. These limitations become especially exacerbated in densely packed chiplet architectures such as NUPLet because of the

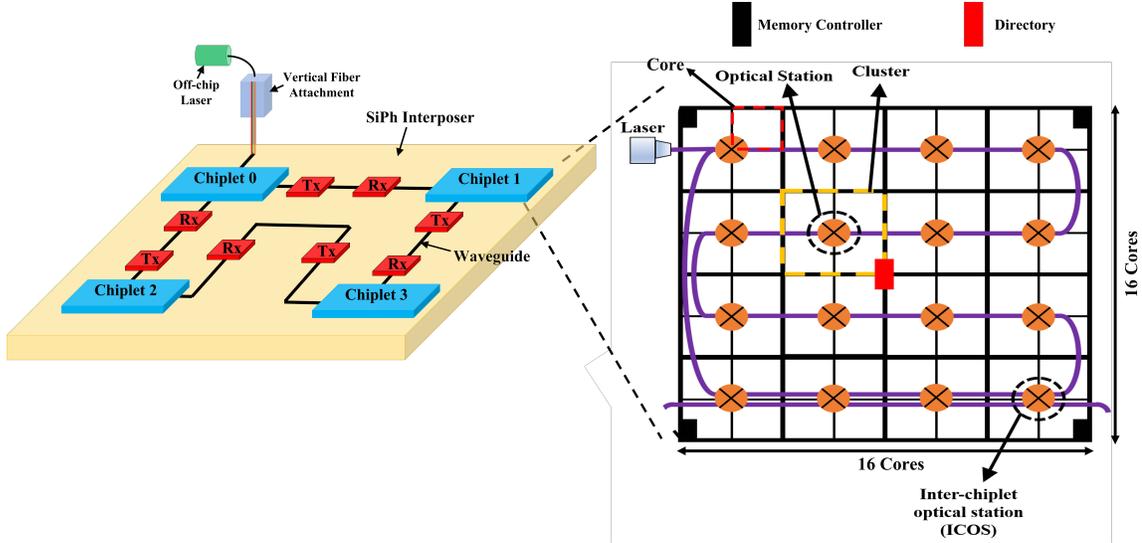


Figure 5.6: Chiplet based Design of NUPLet.

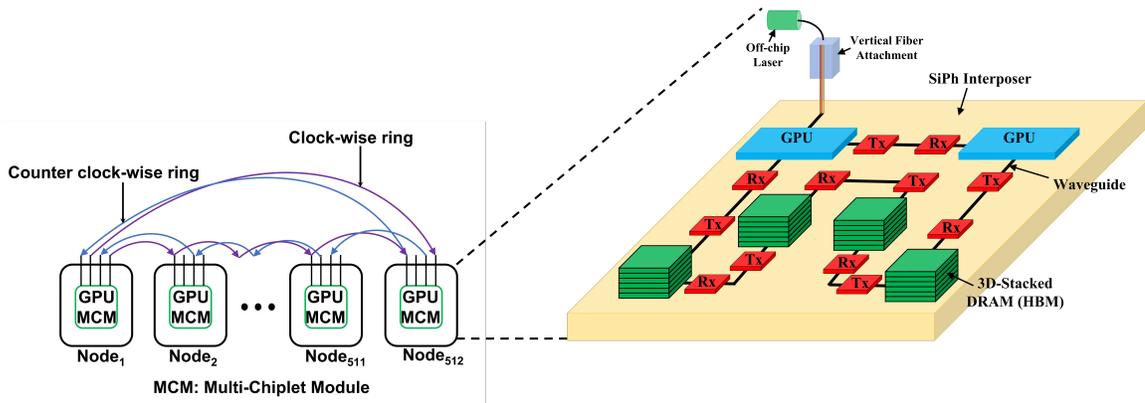


Figure 5.7: GPU based multi-chiplet module (MCM).

dense proximity of interconnects in such architectures. NUPLet architecture (Fig. 5.6) consists of four chips and each chip is called a chiplet. Each chiplet is composed of 32 cores divided into 16 clusters with 2 cores per cluster. Each chiplet in NUPLet also has an 8MB last level cache (LLC) divided into 32 cache banks in 16 clusters with 2 cache banks per cluster. Furthermore, NUPLet employs a directory-based cache coherent scheme, in which the coherence traffic is confined to a single chiplet, thereby optimizing the efficiency, and reducing the latency associated with cache coherence operations. At the interface of each cluster in a chiplet, an optical station is present which consists of a transmitter (modulator MRs array) and a receiver (filter MRs array) that enable inter-chiplet and intra-chiplet data communication. SOI based waveguides connect optical stations in a chiplet with one another in a crossbar configuration. Each optical station in NUPLet receives some amount of multi-wavelength optical power through waveguides via an off-chip laser that can generate up to 180mW of optical power. Whenever an optical station wants to send data, it redirects some portion of the light from the waveguide. This light is split into

multiple wavelengths using a comb splitter. The electrical data packet from the core is converted to parallel electrical data signals and modulated onto these wavelengths using modulator MRs. These modulated wavelengths travel along the waveguide to the destination station where a bank of MRR filters drop these wavelengths onto the adjacent photodetectors to regenerate the electrical data signals and consequently, the electrical data packet which is passed onto the destination core. Intra-chiplet network in NUPLet is based on SWMR (single writer multiple reader) crossbar topology [189] where each optical station is connected to the other optical stations in chiplet using a dedicated waveguide. Similarly, Inter-chiplet network in NUPLet is based on MWMMR (multiple writer multiple reader) crossbar topology [188]. Table 5.5 illustrates how insertion loss varies between SWMR and MWMMR crossbar topologies across various design pathways implemented on the 45nm SOI CMOS platform. For this analysis, we considered two different N_λ values of 16 and 32. In the SWMR topology-based intra-chiplet NUPLet network, there are a total of 16 nodes comprising 1 sender node and 15 receiver nodes. Conversely, in the MWMMR-based inter-chiplet NUPLet network, there are a total of 4 nodes, consisting of 2 sender nodes and 2 receiver nodes. We report the insertion loss for the longest data communication lengths, which were 4.5 cm for the SWMR topology-based intra-chiplet network and 6 cm for the MWMMR topology-based inter-chiplet network. From Table 5.5, even though the SWMR topology entails a wavelength signal traversing higher intermediate nodes (14 receiver nodes) compared to the MWMMR configuration (1 sender node and 1 receiver node), the insertion loss in the MWMMR configuration is notably higher. This counterintuitive result arises because the filter insertion loss for the 45nm SOI CMOS platform is exceptionally low at 0.18dB (Table 5.3), whereas the modulator insertion loss is comparatively high at 4.7dB (Table 5.3). Consequently, the cumulative insertion loss experienced by a wavelength channel is high in the MWMMR configuration compared to the SWMR configuration. Furthermore, the insertion loss is low for the *Minimized loss* design pathway, compared to the *Wide FSR* and *Increased MAOP* design pathways. This is because of the low propagation loss and coupling loss of 1dB/cm and 1 dB/coupler respectively for the *Minimized loss* design pathway. Detailed discussion on inter-chiplet network of the NUPLet is provided in the upcoming subsection.

Inter-Chiplet Network of NUPLet

Optical stations at the bottom of each chiplet are used for both intra-chiplet and inter-chiplet communication and are called as inter-chiplet optical stations (ICOS) as shown in Fig. 5.6. There are a total of 16 ICOSs with 4 ICOSs per chiplet. These ICOSs utilize MWMMR crossbar topology where multiple optical stations can send and receive data using their corresponding modulator and filter MRR banks respectively which enables the stations to share the available optical bandwidth. Each ICOS also consists of queues that hold intra-chiplet and inter-chiplet messages. The inter-chiplet network of NUPLet has 8 data waveguides and 8 power waveguides. If an ICOS wishes to send data, firstly it should get access to a data-power waveguide pair, then redirect some portion of light signal from the power waveguide, use comb

Table 5.5: Insertion loss evaluated for various design pathways implemented on the 45nm SOI CMOS platform. This evaluation encompasses intra-chiplet and inter-chiplet networks within NUPLet that employ SWMR and MWMR crossbar topologies, respectively

SWMR Configuration		
Design Pathway	Insertion Loss (dB) ($N_\lambda = 16$)	Insertion Loss (dB) ($N_\lambda = 32$)
Minimized Loss	15	18.8
Wide FSR	26.25	34.16
Increased MAOP	26.7	34.4
MWMR Configuration		
Design Pathway	Insertion Loss (dB) ($N_\lambda = 16$)	Insertion Loss (dB) ($N_\lambda = 32$)
Minimized Loss	16.2	20.4
Wide FSR	27.5	35.5
Increased MAOP	28.9	36.8

splitter to split the light into multiple wavelength signals, modulate the electrical data onto these wavelength channels and send it to the destination station through the data waveguide.

The power required to transmit data or an inter-chiplet message from one chiplet to other is high compared to power required for intra-chiplet communication. This is because of longer lengths and high propagation losses of inter-chiplet waveguides. In addition, there are other insertion losses such as coupler loss, splitter loss and through loss of MRs. All of these losses increase the laser power consumption and degrade the performance. In order to minimize the laser power consumption in inter-chiplet communication, NUPLet utilizes NUCA (non-uniform cache access schemes) and a unique prediction scheme.

A miss in L1 level cache prompts a request to one of the cache banks in LLC. Cache bank that contains the block of data may lie in same chiplet from which the request was prompted or in any other chiplets. If the cache bank lies in same chiplet, then it is called home bank. Otherwise, it is called non-home bank. Analysis provided in [29] shows that 57% of these prompted requests are sent to non-home banks and only 7% of these result in a hit. For a lower hit rate, large number of inter-chiplet messages are sent resulting in high laser power consumption. Restricting the access requests to local cache banks will reduce the number of inter-chiplet messages that can alleviate this drawback. For that, NUPLet utilizes NUCA schemes which enables the migration of requested cache block to cache banks that are on the same chiplet as the requesting cores. This will increase the hit rate and reduces the amount of inter-chiplet messages.

Execution time of an application is divided into several fixed size durations called epochs. Several prior works have demonstrated power reduction by predicting the traffic for the next epoch by analyzing the behavior of application in previous epochs. NUPLet utilizes a similar type of prediction scheme that predicts the number of inter-chiplet and intra-chiplet messages that will be sent in the next epoch and the

consequent laser power required. Accurate prediction of inter-chiplet and intra-chiplet messages will reduce the wastage of laser power and enhances the performance.

NUCA and prediction schemes of the NUPLet reduce the laser power consumption but insertion loss and power penalties in photonic links of NUPLet are still present that will result in significant amount of laser power consumption. Therefore, we implement our derived inter-chiplet variants on NUPLet architecture and perform a system-level analysis from which we evaluate performance and energy consumption. Details of this evaluation are provided in the upcoming subsections.

GPU based Multi-Chiplet Module

The computation requirements of modern data centric applications such as machine learning has been partially met by swift development of hardware accelerators. Although hardware accelerators have provided a notable amount of speedup but training conventional ML models can still take a significant amount of time. Several solutions have been introduced that enable distributed training on a small number of GPUs connected with a high speed electrical switch with a Tb/s bandwidth. But future ML training workloads require several Tb/s of bandwidth per device at large scales in order to reduce the training time. This raises the need for >1 Tb/s interconnects for distributed ML systems which is implausible to achieve from conventional electrical interconnects. Therefore, in [127], khani et al. proposed an end-to-end optical solution called SiP-ML for scaling of ML workloads by leveraging silicon photonic chiplets. As a part of this work, khani et al. explored two all optical architectures for scaling of ML workloads and one among them is SiP-ring shown in Fig. 5.7. This SiP-ring architecture consists of disaggregated GPU MCMs and the inter-chiplet communication in each of these modules occurs in photonic domain. This approach harnesses the unique advantages of photonics, such as high data rates and low latency, to overcome the bandwidth constraints posed by traditional electrical interconnects. In the SiP-ring architecture, each of the GPU MCMs are connected to each other in a ring topology which enables communication in both directions and is easily reconfigurable. Inside each of the GPU MCMs, there are two GPUs connected to four 3D stacked DRAMs as shown in inset of Fig. 5.7. As a part of our system-level analysis, we implemented our derived on-SiPhI inter-chiplet variants on GPU MCMs and evaluated the impact of aggregate bandwidth of the inter-chiplet variants on the training time of conventional deep neural network (DNN) models, which are widely used in computer vision and natural language processing applications. More details of this evaluation are provided in further subsections.

Evaluation setup

CPU based multi-core multi-chiplet architecture

As a part of our system-level analysis, we have implemented the derived on-SiPhI inter-chiplet variants (Table 5.3) on a CPU based multi-core multi-chiplet architecture named NUPLet [29] and performed a benchmark-driven simulation based analysis from which we have evaluated the performance (1/execution time), energy

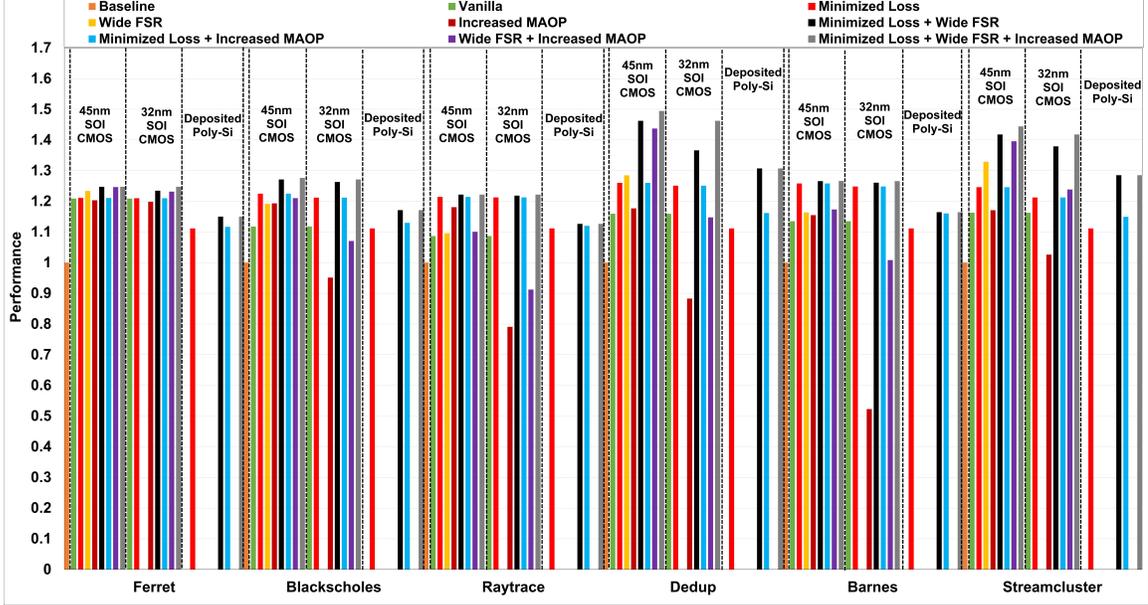


Figure 5.8: Performance comparison of on-SiPhi link variants as implemented on NUPLet architecture. These variants are based on the 45nm SOI CMOS, 32nm SOI CMOS, and deposited poly-si photonic platforms.

consumption and energy-delay product of the NUPLet architecture. We have used four 32-core chiplets in all our designs. We have evaluated our designs on a cycle architectural simulator named *Tejas* [215] for real world traffic applications in the PARSEC benchmark suite [36]. For all our experiments, we have used an epoch size of 100 cycles.

GPU based multi-chiplet module

For the system-level analysis on GPU based MCM [127], we have utilized a simulator named *Rostam* from [127] which is available online at <https://github.com/MLNetwork/rostan.git>. We implement our derived on-SiPhi inter-chiplet variants on *SiP-ring* architecture and evaluate the impact of aggregate bandwidth of inter-chiplet variants on the *time-to-accuracy* of the conventional DNN models. For this analysis, we have considered three representative DNN models namely *ResNet50* [98], *Transformer* and *Megatron* [227]. Among these models, *ResNet* is an image classification model with 25 million parameters. Similarly, *Transformer* is a model with 350 million parameters whereas *Megatron* is a model with 18 billion parameters. We evaluate *time-to-accuracy* metric corresponding to the inter-chiplet variants implemented on the *SiP-Ring* architecture for each DNN model by multiplying the time for a single iteration (obtained from the simulator) by the number of training iterations (considered from prior work [220]) required to reach the target accuracy.

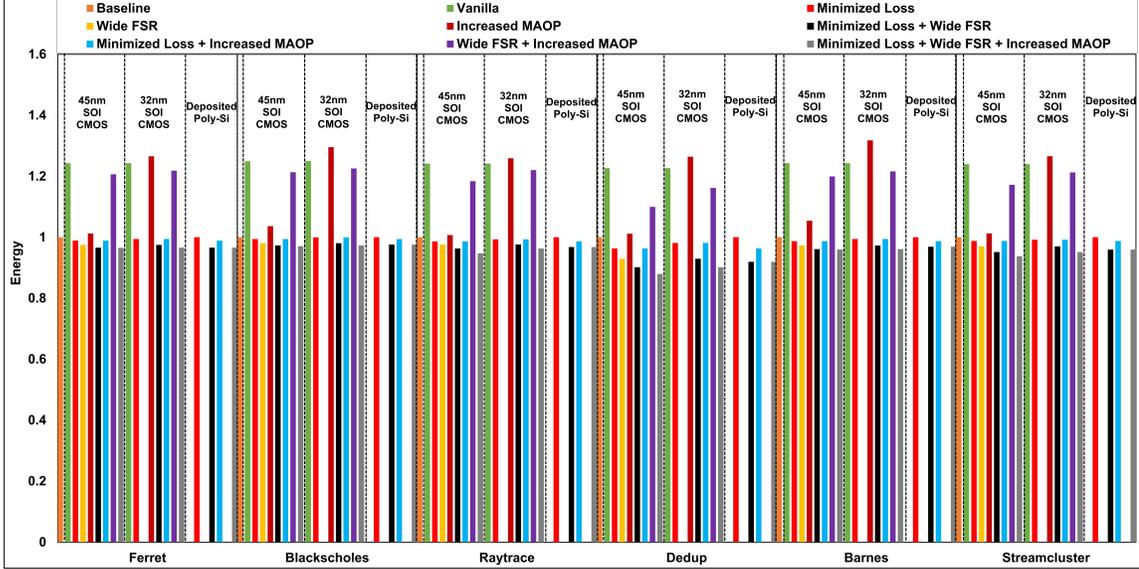


Figure 5.9: Energy comparison of on-SiPhI link variants as implemented on NUPLet architecture. These variants are based on the 45nm SOI CMOS, 32nm SOI CMOS, and deposited poly-si photonic platforms.

Evaluation Results

For the system-level analysis, we have implemented the derived on-SiPhI inter-chiplet variants (Table 5.3) on a CPU based multi-core multi-chiplet architecture named NUPLet [29] and on a GPU based MCM from [127] which is used for distributed ML training. On NUPLet, we have performed a benchmark-driven simulation based analysis from which we have evaluated performance (1/execution time) and energy consumption of the NUPLet architecture. On the GPU based MCM considered from [127], we have evaluated the impact of link-level aggregate bandwidth of our derived on-SiPhI inter-chiplet variants on training time of conventional ML models. The results of this analysis are discussed in the next subsection.

System-level Analysis on CPU based multi-core multi-chiplet module

From the system-level analysis on NUPLet, we have evaluated performance, energy consumption and energy-delay product of NUPLet architecture employed with the derived on-SiPhI inter-chiplet variants. The longest inter-chiplet waveguide length we have considered for this analysis is 8 cm. For this waveguide length, *Wide FSR* variant derived from 32nm SOI CMOS platform, and *Vanilla*, *Wide FSR*, *Increased MAOP* and *Wide FSR + Increased MAOP* variants corresponding to deposited poly-Si platform become non-viable due to high insertion loss (Fig. 5.4). Performance, energy consumption and energy-delay product of the viable architecture variants are discussed below.

Fig. 5.8, Fig. 5.9 and Fig. 5.10 illustrate the relative performance (inverse of simulated execution time), energy consumption and energy-delay product of different variants of NUPLet architecture respectively corresponding to three different fabrica-

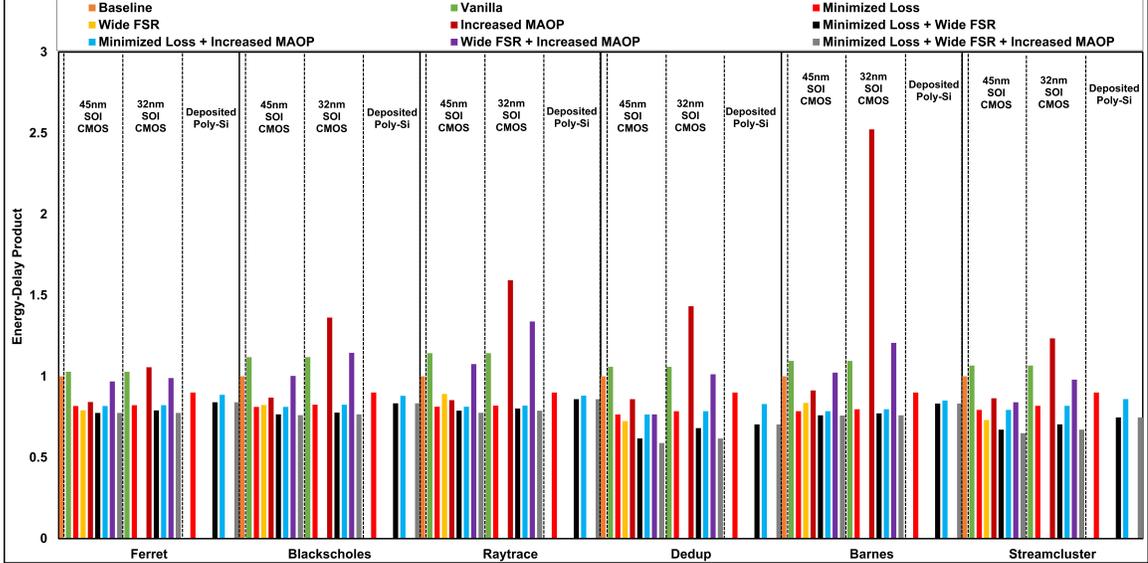


Figure 5.10: Energy-delay product comparison of on-SiPhI link variants as implemented on NUPLet architecture. These variants are based on the 45nm SOI CMOS, 32nm SOI CMOS, and deposited poly-si photonic platforms.

tion platforms for various PARSEC benchmark applications [36]. The metric energy refers to the energy consumed by cores and lasers during the execution of an application. All the results are normalized to a baseline variant of NUPLet which has an N_λ of 32 and bitrate of 10 Gb/s. As we can infer from Fig. 5.8, among variants corresponding to 45nm SOI CMOS platform, the NUPLet architecture that employs *Minimized Loss + Wide FSR + Increased MAOP*, *Minimized Loss + Wide FSR*, *Minimized Loss + Increased MAOP*, *Minimized Loss* and *Wide FSR* inter-chiplet variants achieve 33%, 31.5%, 23.6%, 23.5% and 22% better performance on average respectively across all benchmark applications compared to the baseline variant.

In terms of energy (Fig. 5.9), among variants corresponding to 45nm SOI CMOS platform, the NUPLet architecture that employs *Minimized Loss + Wide FSR + Increased MAOP*, *Minimized Loss + Wide FSR* variants consume 5.7% and 5% less energy on average respectively, followed by the NUPLet variants that employ *Wide FSR*, *Minimized Loss + Increased MAOP* and *Minimized Loss* inter-chiplet variants, across all benchmark applications compared to the baseline variant. In terms of energy-delay product (Fig. 5.10), among inter-chiplet variants corresponding to 45nm SOI CMOS platform, the NUPLet architecture that employs *Minimized Loss + Wide FSR + Increased MAOP*, *Minimized Loss + Wide FSR* and *Minimized Loss + Increased MAOP* inter-chiplet variants achieve 29%, 27% and 21% less energy-delay product on average respectively followed by *Wide FSR*, *Minimized Loss* and *Wide FSR* inter-chiplet variants across all benchmark applications compared to the baseline variant.

Therefore, the NUPLet architecture that employs *Minimized Loss + Wide FSR* and *Minimized Loss + Wide FSR + Increased MAOP* variants corresponding to 45nm SOI CMOS achieve better performance and incur less energy on average across

all benchmark applications compared to the baseline variant. This is because of low insertion loss of inter-chiplet waveguides and high bandwidth of on-SiPhI inter-chiplet links, combined with NUCA and prediction schemes of NUPLet. This is leveraged by the ICOSs of the NUPLet to send more number of inter-chiplet messages/data packets at a time without any wastage of laser power, resulting in execution of application in less number of epochs with enhanced performance and less energy consumption.

Similarl, among the inter-chiplet variants corresponding to 32nm SOI CMOS platform, the NUPLet architecture that employs *Minimized Loss + Wide FSR + Increased MAOP* and *Minimized Loss + Wide FSR* inter-chiplet variants achieve 31.2% and 28.6% better performance respectively on average across all benchmark applications compared to baseline variant. This is followed by *Minimized Loss + Increased MAOP* and *Minimized Loss* inter-chiplet variants that achieve $\sim 22\%$ better performance on average compared to the baseline variant. In terms of energy (Fig. 5.9), the NUPLet architecture that employs *Minimized Loss + Wide FSR + Increased MAOP* and *Minimized Loss + Wide FSR* inter-chiplet variants corresponding to 32nm SOI CMOS platform incur 5% and 3.3% less energy on average respectively across all benchmark applications compared to the baseline variant. In terms of energy-delay product (Fig. 5.10), *Minimized Loss + Wide FSR + Increased MAOP* and *Minimized Loss + Wide FSR* inter-chiplet variants corresponding to 32nm SOI CMOS platform achieve 27% and 25% less energy-delay product on average respectively across all benchmark applications compared to the baseline variant. Therefore, the NUPLet architecture that employs *Minimized Loss + Wide FSR + Increased MAOP* and *Minimized Loss + Wide FSR* variants corresponding to 32nm SOI CMOS platform achieve better performance and incur less energy on average across all benchmark applications compared to the baseline variant.

Similarly, among the inter-chiplet variants corresponding to deposited poly-Si platform, the NUPLet architecture that employs *Minimized Loss + Wide FSR + Increased MAOP* and *Minimized Loss + Wide FSR* inter-chiplet variants achieve 20% better performance, consume 4% less energy and achieve 20% less energy-delay product on average respectively across all benchmark applications compared to the baseline variant.

Therefore, from the system-level evaluation on NUPLet [29], we have observed that chiplet based PNoC architectures that employ *Minimized Loss + Wide FSR + Increased MAOP*, *Minimized Loss + Wide FSR* and *Minimized Loss* on-SiPhI inter-chiplet variants corresponding to 45nm SOI CMOS, 32nm SOI CMOS and deposited poly-si platforms achieve superior performance and consume less energy compared to other inter-chiplet variants.

System-level analysis on GPU based multi-chiplet module

For the system-level analysis on GPU based MCM, we have utilized the simulator provided in [127] and evaluated time-to-accuracy i.e., the training time of three conventional DNN models namely ResNet50, Transformer and Megatron based on the aggregate bandwidth of our derived inter-chiplet variants enacted in GPU based MCMs.

As we can infer from Fig. 5.11, GPU based MCMs that employ *Minimized Loss + Wide FSR* and *Minimized Loss + Wide FSR + Increased MAOP* inter-chiplet variants corresponding to 45nm SOI CMOS, 32nm SOI CMOS and deposited poly-Si platforms enable at least 1-1.75 \times , 2-8 \times , 4-30 \times faster training time for ResNet50, Transformer and Megatron respectively. This is because, both of these inter-chiplet variants achieve multi-Tb/s aggregate bandwidth at link-level.

5.6 Key Results

In this section, we summarize the results obtained from the link-level and system-level analysis of on-SiPhI inter-chiplet variants derived based on our identified design pathways (Table 5.2), corresponding to three different SiPh fabrication platforms (Table 5.3).

Key Link-Level Results

From the link-level analysis, we have evaluated the aggregate bandwidth (primary Y-axis in Fig. 5.4) and EPB (secondary Y-axis in Fig. 5.4) for different on-SiPhI inter-chiplet variants corresponding to three different SiPh fabrication platforms, for different waveguide lengths (X-axis in Fig. 5.4). Based on the results obtained from link-level analysis, we have classified the derived inter-chiplet variants into two categories namely non-viable variants and viable variants.

Non-Viable Variants

Non-viable variants are the inter-chiplet variants that do not support any wavelength channels in the link and therefore support no aggregate bandwidth for longer waveguide lengths. The non-viable variants determined from this analysis are as follows:

1. *Vanilla* and *Wide FSR* variants corresponding to 32nm SOI CMOS and deposited poly-Si SiPh platforms do not support any wavelength channels due to high insertion loss in the link whereas the same variants corresponding to 45nm SOI CMOS platform support wavelength channels up to a link length of 4cm and they can be made viable for longer waveguide lengths by employing repeaters
2. *Increased MAOP* and *Wide FSR + Increased MAOP* variants corresponding to deposited poly-Si platform does not support any wavelength channels due to high insertion loss in the link whereas the same variants corresponding to 45nm SOI CMOS and 32nm SOI CMOS platforms supports wavelength channels up to a link length of 8 cm and 2 cm respectively and they can be made viable by utilizing repeaters

Viable Variants

Viable variants are the inter-chiplet variants that support wavelength channels in the link up to link lengths as long as 10cm. The viable variants determined from this analysis are as follows:

1. *Minimized Loss*, *Minimized Loss + Wide FSR*, *Minimized Loss + Increased MAOP* and *Minimized Loss + Wide FSR + Increased MAOP* variants corresponding to three different SiPh fabrication platforms support wavelength channels up to a link length of 10cm
2. Among the viable variants, *Minimized Loss + Wide FSR + Increased MAOP* variant corresponding to 45nm SOI CMOS platform achieves highest aggregate bandwidth of 4.92 Tb/s with corresponding EPB of 9.2pJ/bit whereas *Minimized Loss + Wide FSR* variant corresponding to the same fabrication platform achieves lowest EPB of 0.218 pJ/bit with corresponding aggregate bandwidth of 4.6 Tb/s

Key System-Level Results

We have implemented the on-SiPh inter-chiplet variants on a CPU based multi-core multi-chiplet architecture named NUPLet [29] and a GPU based multi-chiplet module (MCM) [127] and performed a system-level analysis. Results of this analysis are summarized as follows.

System-Level Evaluation on CPU Based Multi-Core Multi-Chiplet Architecture

We have implemented the derived inter-chiplet variants on NUPLet architecture[29] and performed a benchmark-driven simulation based analysis from which we have evaluated the performance (Fig.; 5.8), energy consumption (Fig. 5.9) and energy-delay product (Fig. 5.10) of the NUPLet architecture. The results obtained from this evaluation are summarized as follows:

1. NUPLet architecture that employs *Minimized Loss*, *Minimized Loss + Wide FSR* and *Minimized Loss + Wide FSR + Increased MAOP* inter-chiplet variants corresponding to three considered SiPh fabrication platforms (Table 5.3) achieve 28% better performance on average compared to the baseline variant and 7.5% better performance on average compared to the other inter-chiplet variants. Similarly, NUPLet architecture that employs *Minimized Loss*, *Minimized Loss + Wide FSR* and *Minimized Loss + Wide FSR + Increased MAOP* inter-chiplet variants corresponding to three considered SiPh fabrication platforms (Table 5.3) consume 5% less energy on average compared to the baseline variant and 2% less energy on average compared to the other inter-chiplet variants.

System-Level Evaluation on GPU Based Multi-Chiplet Module

We have implemented the derived inter-chiplet variants on a GPU based MCM [127] and performed a system-level analysis utilizing the simulator provided in [127], from which we have evaluated the time-to-accuracy of three conventional DNN models namely ResNet50 (Fig. 5.11(a)), Transformer (Fig. 5.11(b)) and Megatron (Fig. 5.11(c)). The results of this evaluation are summarized as follows:

1. GPU based MCMs that employ *Minimized Loss + Wide FSR* and *Minimized Loss + Wide FSR + Increased MAOP* inter-chiplet variants corresponding to the three considered SiPh fabrication platforms (Table 5.3) accelerate the training time for ResNet50, Transformer and Megatron DNN models by at least $1-1.75\times$, $2-8\times$ and $4-30\times$ respectively.

5.7 Summary

The dwindling of Moore’s law has drastically increased the complexity and the cost of fabricating large-scale, monolithic Systems-on-Chip (SoCs). Therefore, the industry has adopted fragmentation of monolithic SoCs into several smaller chiplets, which are then assembled using silicon interposer. However, to meet the growing demands of modern data-centric workloads, it is vital to realize on-interposer inter-chiplet communication bandwidth of multi-Tb/s and end-to-end communication latency of $<10\text{ns}$. To meet these bandwidth and latency goals, prior works have focused on a potential solution of using the silicon photonic interposer (SiPhI) for integrating and interconnecting a large number of chiplets into a system-in-package (SiP). However, the designs of on-SiPhI interconnects, demonstrated so far, have to still evolve swiftly in order to meet the goal of multi-Tb/s bandwidth. But the possible design pathways that can aid in such evolution, have not been explored yet. Therefore, in this chapter, we identified several design pathways that can aid on-SiPhI inter-chiplet interconnects to meet the goal of achieving multi-Tb/s bandwidth.

Based on the identified design pathways and three different photonic fabrication platforms, namely 45nm SOI CMOS, 32nm SOI CMOS and deposited poly-Si, we derived twenty four design variants of on-SiPhI inter-chiplet interconnects. Then, we performed an extensive link-level and system-level analysis for each of these variants. From the link-level analysis, we inferred that the design pathways that simultaneously enhance the spectral range and optical power budget available for wavelength division multiplexing (WDM) provide enough impetus to the corresponding on-SiPhI inter-chiplet links to achieve aggregate bandwidth of $>4\text{Tb/s}$ and support link lengths of up to 10cm. Subsequently, leveraging this link-level analysis, we conducted a system-level analysis on state-of-the-art CPU and GPU-based SiPs, by incorporating multi-Tb/s on-SiPhI inter-chiplet links. From the system-level analysis on CPU based SiP, we observed that design pathways that simultaneously enhance the spectral range and optical power budget available for WDM achieve at least 25% better performance while consuming at least 5% less energy on average compared to other design pathways. Similarly, from the system-level analysis on GPU-based SiP, we inferred that

these same design pathways accelerate the training time of large-scale DNN models by at least $15\times$ on average compared to other design pathways. These results suggest that simultaneously enhancing the spectral multiplexing range and optical power budget of on-SiPhI interconnects would pave the way for achieving multi-terabits/second performance in the future.

Chapter 6 A Polymorphic Electro-Optic Logic Gate for High-Speed Reconfigurable Computing Circuits

6.1 Introduction

Moore's law has been steering the advancement of computing hardware since its inception. But unfortunately, in recent years, it has faced fatal challenges as the nanofabrication technology is experiencing physical limitations due to the exceedingly small size of transistors [86]. This has forced researchers in industry and academia to develop new more-than-Moore technologies that can continue to provide persistently faster and more efficient computing hardware [86]. Fortunately, silicon photonics (SiP) enabled electro-optic (E-O) circuit integration has been identified as one such promising technology [286]. The SiP-based E-O circuits are generally CMOS compatible and provide several advantages over their purely electrical counterparts. These advantages include sub-picosecond speeds, low dynamic power consumption and distance-independent bit-rate [286]. Due to these advantages compared to the CMOS-based electrical circuits, the early prototypes of SiP-based E-O circuits for computing (e.g., [297, 226, 225, 200, 287, 123, 286]) have been shown to provide up to two orders of magnitude improvements in performance and energy efficiency [26][286].

The SiP-based E-O circuits for computing, which have been demonstrated in prior works (e.g., [297, 226, 225, 200, 287, 123, 286]) are typically used to implement the following four types of logical and arithmetic functions: **(I) Basic logic-gate functions.** For instance, a microring resonator (MRR) integrated phase change memory (PCM) device based XNOR gate is employed in [297] to enable acceleration of binary neural networks. Similarly, in [226] and [225], an add-drop MRR based AND gate is employed to enable partial multiplications of two binary operands, to aid the acceleration of deep neural networks. **(II) Arbitrary combinational logic functions.** For example, the directed logic based MRR-enabled reconfigurable E-O circuits are demonstrated in [200] and [287]. These can work as the direct optical replacement of field programmable gate arrays (FPGAs). **(III) Two-operand arithmetic functions.** High-speed E-O circuits for partial sum accumulation and two-operand addition have been demonstrated (e.g., [123], [286]) with various designs supporting custom precision [123] and full-precision polymorphic operation [286]. **(IV) Multi-operand linear arithmetic functions.** Several analog and digital E-O circuits based on MRRs and/or Mach-Zehnder Interferometers (MZIs) have been demonstrated to implement Multiply-Accumulate (MAC) and Vector Dot Product (VDP) operations (e.g., [297, 226, 152, 26]) for deep learning workloads. These logical and arithmetic functions implemented using E-O circuits typically fulfill the requirements of ultra-fast, highly-parallel general purpose computing or deep learning acceleration.

However, we observe that these SiP-based E-O circuits from prior works face three major shortcomings. *First*, the E-O circuits for simple logic-gate functions intake the two input operands differently; one operand is typically applied optically and the

other operand is applied electrically. For instance, in the E-O XNOR gate from [297] and the E-O AND gate from [226], one of the two operands has to be modulated onto the incoming optical wavelengths, for which an additional optical modulator device per gate function is required, assuming that the utilized laser sources provide unmodulated optical power. Having to provide one of the operands optically through an additional modulator device increases the hardware area overhead and the operand handling complexity in the E-O circuits. Instead, there is a need to design a simpler hardware, which can be achieved by promoting all electrical provisioning of both the operands. *Second*, the E-O circuits for arithmetic functions occupy very large areas compared to CMOS implementations. For instance, the E-O MAC circuit used in [225] occupies up to $100\times$ more area compared to the all-electric MAC circuit [225]. Moreover, such E-O circuits for arithmetic functions can hardly achieve more than 60% hardware utilization [226]. This is because these circuits typically belong to larger processing units where they occupy only part of the entire end-to-end datapath [26, 297, 152]. Such low hardware utilization often leads to high idle time and consequently very high, non-amortizable area and static power overheads. This in turn motivates the need for more flexible E-O circuits that can adapt to different arithmetic/logic functions at different times, to increase the amortization of their high area and static power overheads by reducing their total idle time. *Third*, the high area overhead of E-O circuits makes them less suitable for highly-parallel Single-Instruction-Multiple-Data (SIMD), Multiple-Instruction-Multiple-Data (MIMD), and Systolic Array (SA) based processing architectures. This is because SIMD, MIMD and SA architectures typically employ thousands of streaming processing units, with each processing unit requiring multiple copies of basic logical and arithmetic functions. Implementing these functions using bulky E-O circuits with $100\times$ more area can drastically reduce the number of processing units that can be integrated on a single chip whose area is typically limited by the reticle size ($\leq 900\text{ mm}^2$ [172]). Since SIMD, MIMD, and SA based processing units have become extremely popular for executing modern Euclidean as well as non-Euclidean data workloads (i.e., workloads with grid and graph structured data), it becomes crucial to alleviate the unsuitability of E-O circuits for SIMD, MIMD and SA based designs by forging new E-O circuits with relatively low area overheads.

To address these shortcomings, in this chapter, we present a single **MRR** based **Polymorphic E-O Logic Gate (MRR-PEOLG)**. Our MRR-PEOLG can accept both input operands electrically, and its drop-port (through-port) optical response can be thermo-optically programmed to make it dynamically follow the truth table of different logic functions, such as AND, OR and XOR (NAND, NOR, and XNOR), at different times. Consequently, the E-O circuits built using our MRR-PEOLG can address the above-described shortcomings by providing (1) the ability of all-electrical application of the input operands, (2) compactness through a single-MRR structure of the E-O gate, and (3) high flexibility through the introduced polymorphism, and consequently, low idle time and improved suitability for use with SIMD/MIMD/SA based architectures.

The key contributions of this chapter are summarized below:

- We model our MRR-PEOLG using the photonics foundry-validated tools from Ansys/Lumerical [155], and then, perform the frequency, time-domain transient, and thermal analysis for different logic-gate functions;
- Based on our analysis, we evaluate the performance of our MRR-PEOLG, from which we determine the maximum achievable bit-rate and thermal tuning power for each logic-gate function supported by our MRR-PEOLG;
- We show that the use of our MRR-PEOLG in two E-O circuits from prior works can provide improvement in area-energy-delay product of up to $82.6\times$;
- We also discuss how MRR-PEOLG can be used to realize E-O reconfigurable SIMD/MIMD architectures.

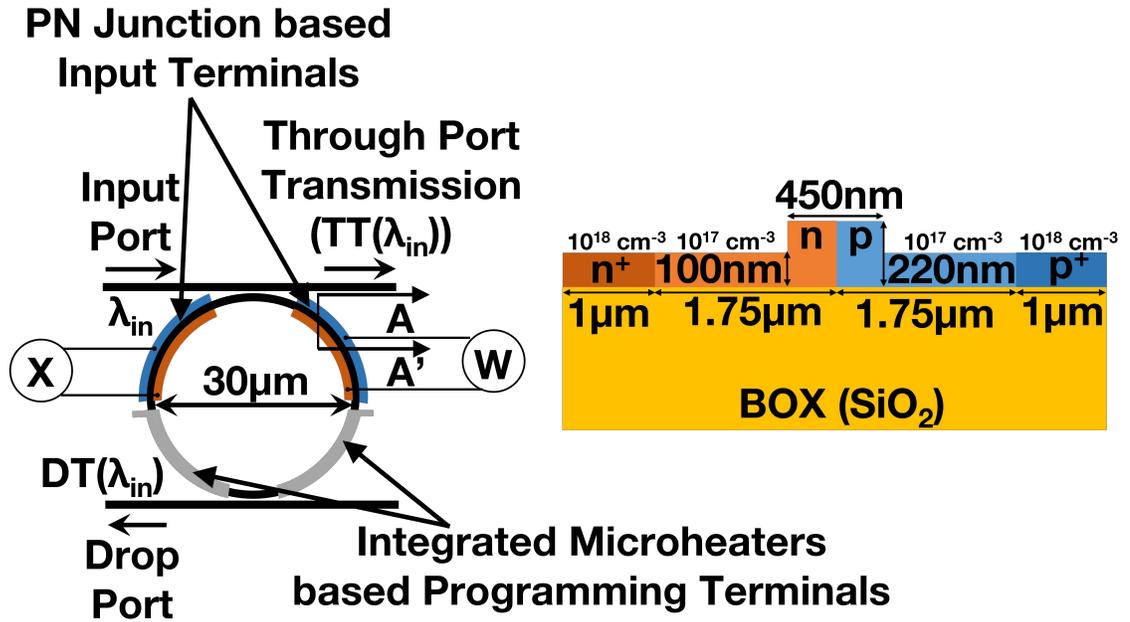


Figure 6.1: Structure and cross-section of our MRR based polymorphic E-O logic gate (MRR-PEOLG).

6.2 MRR-Based Polymorphic Electro-Optic Logic Gate (MRR-PEOLG)

Structure

Our MRR-PEOLG is basically an add-drop MRR [41] with four quarter-sized phase-shifting sections embedded in it, as shown in Fig. 6.1(a). Two quarter-sized sections of the MRR are two PN junctions which are operated in the forward bias condition, whereas the remaining two quarter-sized sections integrate micro-heaters. The cross-section of a PN-junction based section of our MRR-PEOLG is shown in the right hand side of Fig. 6.1(a), which consists of a ridge waveguide with an embedded

lateral PN junction, fabricated on the top of a buried oxide layer. The dimensions of the P-type and N-type regions, and their corresponding carrier concentrations are also provided in Fig. 6.1(a). The PN junction based sections of our MRR-PEOLG work as the input terminals where the input logic signals/operand bits are applied. On the other hand, the microheaters integrated sections of our MRR-PEOLG work as the programming terminals that are used to program the MRR-PEOLG to perform specific logic-gate functions.

Applying a voltage to the microheaters based programming terminals can increase the temperature of the MRR, which in turn can shift (red shift) the resonance of the MRR towards the longer wavelength. This is because of the thermo-optic effect in silicon ([16]). To program MRR-PEOLG to implement a specific logic-gate function, the operand-independent MRR resonance (i.e., the programmed MRR resonance) is adjusted to a specific spectral position with respect to the input optical wavelength, by applying a voltage to the programming terminals. Then, the electrical input logic signals or input operand bits (x and w) are applied to the PN junctions based input terminals of the MRR. Upon doing so, the resonance of the MRR shifts (blue shifts) towards the shorter wavelength depending on the combination of the applied input operand bits. This is because of the free-carrier plasma dispersion effect in silicon ([171]). Applying the input operand bits to the input terminals makes the through-port and drop-port optical responses of our MRR-PEOLG follow the truth-table of the logic-gate functions for which the MRR-PEOLG is programmed. In this manner, our MRR-PEOLG can perform different logic-gate functions at different times. At any given time, the through-port optical response of the MRR-PEOLG follows logical complement of the drop-port optical response. Therefore, AND, OR and XOR functions can be realized (one function at a time) at the drop port of the MRR-PEOLG. Concurrently, the through port of the MRR-PEOLG can provide complementary logic-gate functions such as NAND, NOR and XNOR as discussed below.

Modeling

We model our MRR-PEOLG using the photonics foundry-validated simulation tools from Ansys/Lumerical [155]. We break down our MRR-PEOLG design into a set of primitive elements. Fig. 6.2(a) shows a schematic of our MRR-PEOLG, whose breakdown into the primitive elements is shown in Fig. 6.2(b). We use different solvers in the Ansys/Lumerical tools [155] to model each primitive element. From these models, we extract various parameters for each primitive element. Later, we combine all of the extracted parameters in Ansys/Lumerical's INTERCONNECT tool [155] (tool for the modeling and simulations of photonic integrated circuits) to create our MRR-PEOLG in Fig. 6.2(b). Finally, we perform the frequency-domain and time-domain transient simulations of our MRR-PEOLG. Different steps for the modeling and simulation of our MRR-PEOLG using the ANSYS/Lumerical tools/solvers are summarized below.

Step-1 - modeling MRR-waveguide coupling sections: First, create coupling sections in the finite difference time domain (FDTD) solver and extract the

power coupling coefficients as a function of wavelength for the fundamental TE mode. Import these coefficients in the coupling elements C_1 and C_2 (Fig. 6.2(b)).

Step-2 - modeling straight waveguide sections: First, characterize the passive, straight, channel waveguides of the MRR-PEOLG using the finite difference eigenmode (FDE) solver. Extract the effective index, group index, and dispersion for the waveguides as functions of wavelength. Load this information into the primitive elements WGD_1, WGD_2, WGD_7, and WGD_8 (Fig. 6.2(b)).

Step-3 - modeling PN-junction based input terminals: First, create a quarter ring with an embedded lateral PN-junction in the CHARGE tool. Perform the simulation to extract the spatial distribution of the charge carriers as a function of the bias voltage. Then, export this data into the FDE solver and calculate the perturbations in the refractive index of the waveguides connected to the input terminals. Then, calculate the change in the effective index and resonance of the entire MRR-PEOLG as a function of the bias voltage. Import this information into the primitive elements WGD_6 (connected to OM_1) and WGD_5 (connected to OM_2) (Fig. 6.2(b)).

Step-4 - modeling microheaters based programming terminals: Extract the temperature profile of the MRR-PEOLG as a function of the applied microheater voltage. Then, import this data into the FDE solver to calculate the change in the effective index of the MRR-PEOLG as a function of its temperature. Import this information into the primitive elements WGD_3 (connected to OM_4) and WGD_4 (connected to OM_5) (Fig. 6.2(b)).

Step-5 - preparing for simulations: Connect the primitive-elements based model of the MRR-PEOLG (Fig. 6.2(b)) with other testing and characterization apparatus in the INTERCONNECT tool, as shown in Fig. 6.2(c) and Fig. 6.2(d).

Operation

To explain the operation of our MRR-PEOLG, we performed frequency-domain simulations using the INTERCONNECT tool [155]. Our simulation setup for this frequency-domain analysis is shown in Fig. 6.2(c). Accordingly, we connected an optical network analyzer (ONA) to our MRR-PEOLG to extract the transmission spectra at its drop and through ports. We extracted the transmission spectra for different values of the detuning of the operand-independent MRR resonance position κ with respect to the input wavelength λ_{in} . As mentioned earlier, these detuning values correspond to different logic-gate functions that the MRR-PEOLG can perform. In addition, we also extracted transmission spectra for different combinations of the input operand bits. All of these transmission spectra for different logic-gate and complementary logic-gate functions are shown in Figs. 6.3(a) to 6.3(f). Transmission spectra corresponding to logic-gate functions AND, OR and XOR are shown in Figs. 6.3(a), 6.3(b), and 6.3(c) respectively. These transmission spectra are drop-port transmission spectra (Lorentzian lineshape passbands). Similarly, the transmission spectra corresponding to complementary logic-gate functions NAND, NOR and XNOR are shown in Figs. 6.3(d), 6.3(e), and 6.3(f) respectively. These transmission spectra are through-port transmission spectra (inverse Lorentzian lineshape

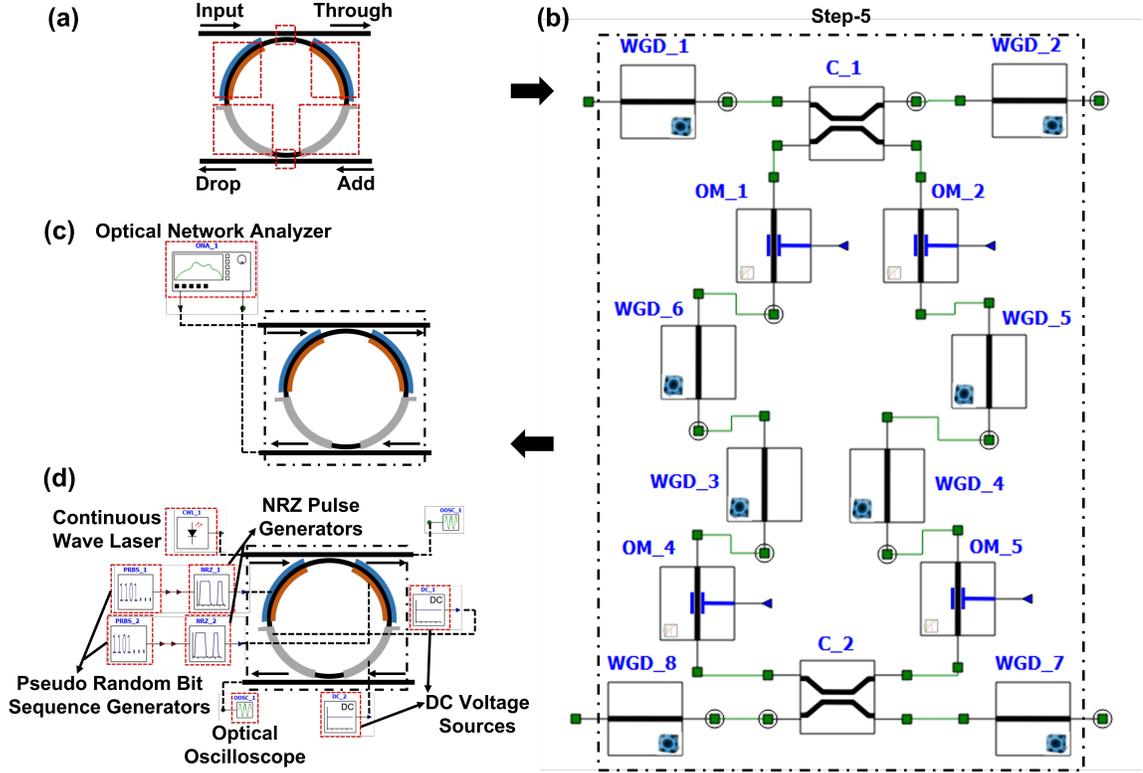


Figure 6.2: (a) Step-1 to Step-4, and (b) Step-5 of the procedure used for modeling our MRR-PEOLG. The schematic simulation setup in ANSYS/Lumerical’s INTERCONNECT tool for (c) frequency-domain and (d) time-domain transient analysis of our MRR-PEOLG.

passbands). As we can see in Fig. 6.3, the drop port and through port transmission exhibits two clearly distinguishable levels. The full transmission range at the drop port and through port of our MRR-PEOLG is divided into two areas, in which the lower part of the full transmission range is indicated with shaded gray whereas the upper part is indicated with shaded blue. If the drop port ($DT(\lambda_{in})$) and through port ($TT(\lambda_{in})$) transmission at λ_{in} falls in the lower part of the full transmission range (i.e., in the gray-shaded area), then it is referred to as logic ‘0’ transmission. On the other hand, if the drop port and through port transmission at λ_{in} falls in the upper part of the full transmission range (i.e., in the blue-shaded area), then it is referred to as logic ‘1’ transmission. However, the vertical spans of the two distinguishable transmission levels differ between the drop port and through port. This is because, similar to the transmission spectra (Fig. 6.3), the spans of transmission levels at the drop port also complement the spans of transmission levels at the through port. The difference between the minimum supported logic ‘1’ transmission and the maximum supported logic ‘0’ transmission is the sensitivity of optical modulation amplitude (SOMA). SOMA is a property of the photodetector based receiver circuit, and it affects the performance of the MRR-PEOLG (as will be discussed in Sections III and IV).

To clearly understand the operation of our MRR-PEOLG, let us consider the

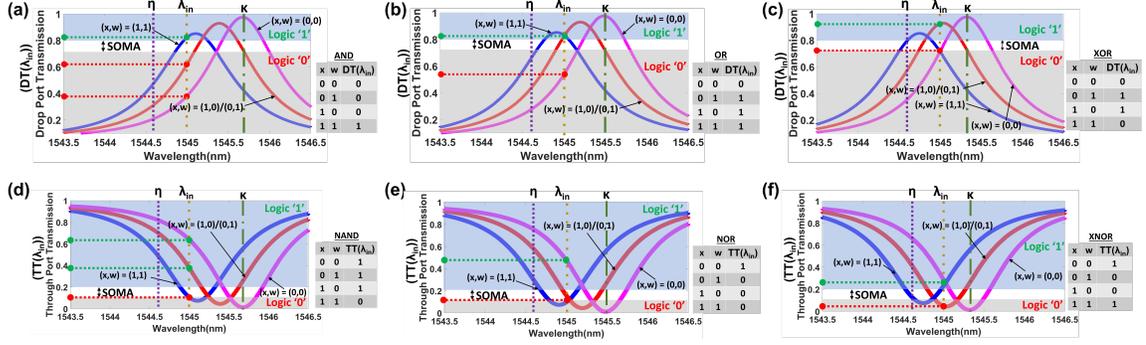


Figure 6.3: The transmission spectra obtained at the drop port of our MRR-PEOLG for logic-gate functions (a) AND, (b) OR, and (c) XOR, and at the through port of our MRR-PEOLG for complementary logic-gate functions (d) NAND, (e) NOR, and (f) XNOR.

example of AND function, as shown in Fig. 6.3(a). To program our MRR-PEOLG to implement AND function, a 0.9 V voltage (3.52 mW power) is applied to the programming terminals of the MRR-PEOLG. This shifts the resonance from the initial position, η , to the programmed position, κ , where κ has the programmed detuning of 0.7 nm with respect to λ_{in} . Then, the input operand bits x and w are applied to the input terminals of the device. Doing so induces a blueshift in the MRR resonance, the magnitude of which depends on the specific combination of the applied input operand bits (x and w , as shown in Fig. 6.3(a)). If the applied bit-combination (x,w) is $(0,0)$, the resonance position of the MRR stays at κ (magenta colored passband in Fig. 6.3(a)) and the drop port transmission at λ_{in} provides logic '0' level (the bottom red dot on the Y-axis). If the applied bit-combination (x,w) is $(0,1)$ or $(1,0)$, the position of the MRR resonance changes (red/orange colored passband in Fig. 6.3(a)), but the blueshift is the same for both $(0,1)$ and $(1,0)$ bit combinations, and the drop port transmission at λ_{in} still remains at logic '0' level (the top red dot on the Y-axis). On the other hand, if the applied bit-combination (x,w) is $(1,1)$, the MRR resonance undergoes a larger blueshift (blue colored passband in Fig. 6.3(a)), and the position of the passband with respect to λ_{in} changes. As a result, the drop port transmission at λ_{in} changes to logic '1' level (the green dot on the Y-axis). Hence, the drop port transmission at λ_{in} for our MRR-PEOLG changes with the applied input operand bits, and follows the truth table of the AND logic function (see the truth table in Fig. 6.3(a)). As discussed earlier, since the through-port response provides a logical complement to the drop-port response, this AND function at the drop port of the MRR-PEOLG corresponds to NAND function at the through port of the MRR-PEOLG as illustrated in Fig. 6.3(d).

Similarly, our MRR-PEOLG can be reconfigured to implement OR (NOR) and XOR (XNOR) gate functions as well, by applying a suitable voltage to the programming terminals of our MRR-PEOLG to set the relative position of κ with respect to λ_{in} as shown in Figs. 6.3(b) and 6.3(c) (transmission spectra corresponding to NOR and XNOR are shown in figs. 6.3(e) and 6.3(f) respectively). Table 6.1 provides the total power consumed in the microheaters, the programmed detuning ($\kappa - \lambda_{in}$), and the required resonance shifting ($\eta - \kappa$), to program our MRR-PEOLG for implement-

Table 6.1: Power consumed in the microheaters, programmed detuning, and required resonance shifting, used to program our MRR-PEOLG for implementing different logic functions.

Logic-Gate Functions	Microheater Power (mW)	Programmed Detuning ($\kappa-\lambda_{in}$) (nm)	Required Shifting ($\eta-\kappa$) (nm)
AND / NAND	3.52	0.7	-1.1
OR / NOR	2.93	0.5	-0.9
XOR / XNOR	2.3	0.3	-0.7

ing various logic functions. From Table 6.1, the power consumed in the microheaters is proportional to the required resonance shifting ($\eta-\kappa$).

6.3 Transient Analysis

Method

As illustrated in Fig. 6.2(d), to perform transient analysis of our MRR-PEOLG in the INTERCONNECT tool, we connected a pseudo random bit sequence (PRBS) generator and a non-return-to-zero (NRZ) pulse generator to each of the input terminals of the MRR-PEOLG. Each PRBS generator generates a random bit sequence of 10 Gb/s, which is given as input to the NRZ pulse generator. Each NRZ pulse generator then generates a sequence of electrical NRZ pulses of 1.5 V amplitude at 10 Gb/s. The blue and red pulses shown in Figs. 6.4(a) and 6.4(b) respectively are the electrical NRZ pulses that we have provided as inputs to the two input terminals of our MRR-PEOLG for the transient analysis. We have also connected a continuous wave (CW) laser to the input port of the MRR-PEOLG, which generates an optical signal of wavelength 1545 nm ($\lambda_{in} = 1545$ nm in Fig. 6.3) with an optical power of 5 dBm. We connected optical oscilloscopes to the drop and through ports to record the output pulse patterns corresponding to different logic functions for the given input electrical pulse signals. The results obtained from this transient analysis are discussed next.

Results and Discussion

Fig. 6.4(c), Fig. 6.4(e), and Fig. 6.4(g) illustrate the output pulse signals obtained at the drop-port of the MRR-PEOLG for different logic functions. Similarly, Fig. 6.4(d), Fig. 6.4(f), and Fig. 6.4(h) illustrate the output pulse signals simultaneously obtained at the through-port of the MRR-PEOLG for different complementary logic functions. To obtain these pulse patterns, we first reconfigured the MRR-PEOLG to implement various logic functions by changing the temperature using the integrated microheaters. We then followed the method described in Section III.A. As evident from Fig. 6.4, the output pulse signals follow the pulse-wise truth-tables of the respective logic functions, which confirms the capability of our MRR-PEOLG to correctly realize different logic functions.

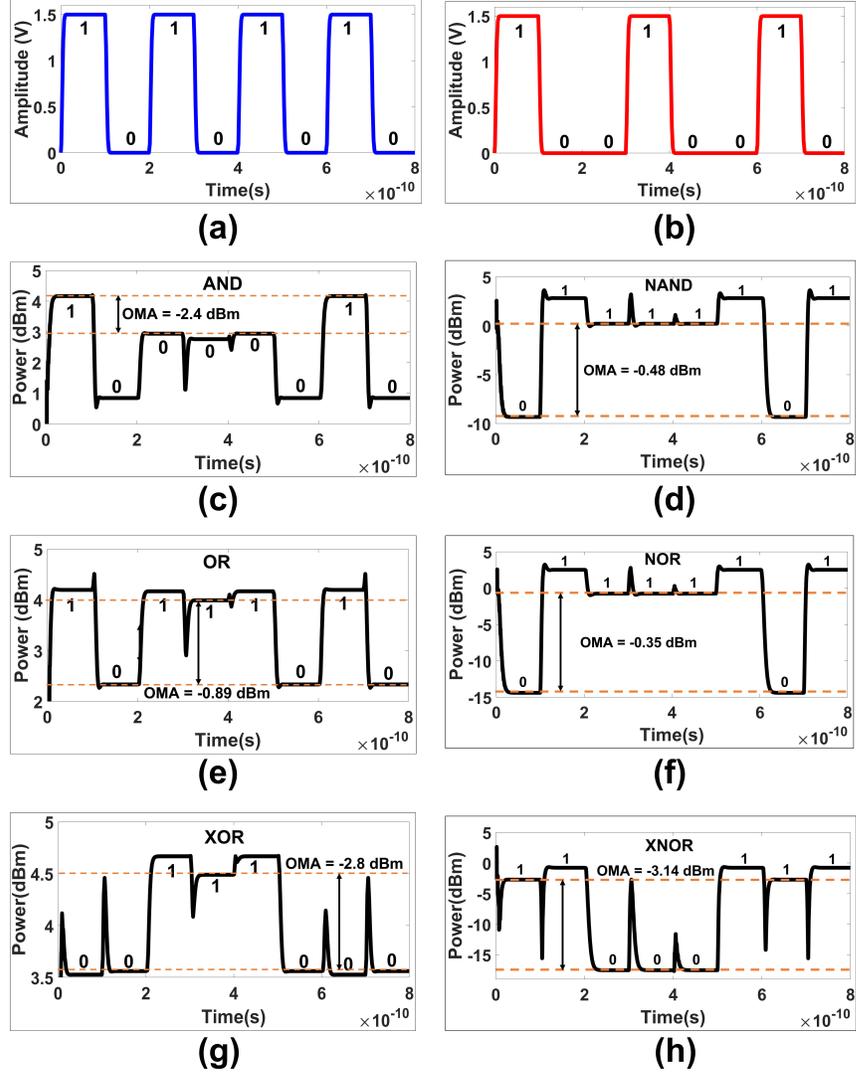


Figure 6.4: (a),(b) The electrical pulse signals of 10 Gb/s bit-rate provided as input to the PN junctions of our MRR-PEOLG. The corresponding output pulse patterns obtained at the drop port of our MRR-PEOLG for logic-gate functions (c) AND, (e) OR, and (g) XOR, and at the through port of our MRR-PEOLG for complementary logic-gate functions (d) NAND, (f) NOR, and (h) XNOR. The optical input power is 5 dBm in all cases.

From Figs. 6.4(c) - 6.4(h), the optical modulation amplitude (OMA), which is the difference between the minimum logic '1' power level and the maximum logic '0' power level in an output pulse pattern, differs for different logic functions. To clearly understand this, let us consider AND, XOR, and OR functions. For the AND function shown in Fig. 6.4(c), the OMA is ~ -2.4 dBm. This is because, as can be observed from Fig. 6.4(c), the drop port transmission at λ_{in} corresponding to the logic '1' output level (i.e., $(x,w) = (1,1)$) is ~ 0.82 (the green dot on the Y-axis), whereas the maximum drop port transmission at λ_{in} , corresponding to the logic '0' output level (i.e., $(x,w) = (1,0)$ or $(0,1)$), is ~ 0.62 (the top red dot on the Y-axis). Hence, the

OMA, i.e., the difference between the logic '1' optical power level ($0.82 \times 5 \text{ dBm} = 2.57 \text{ mW}$) and the logic '0' optical power level ($0.62 \times 5 \text{ dBm} = 1.995 \text{ mW}$) is $\sim -2.4 \text{ dBm}$ ($\sim 0.575 \text{ mW}$). Similarly, for OR (NOR) and XOR (XNOR) functions shown in Fig. 6.4, the green and red dots on Y-axis occur at different positions compared to AND (NAND) function. Therefore, our MRR-PEOLG exhibits different OMA for different logic functions.

Since the OMA for the output pulse pattern basically defines how well the logic '1's are distinguishable from logic '0's, having different OMA values renders different reliability bounds for different logic functions implemented by our MRR-PEOLG. In general, to achieve higher reliability without trying to quantify its value, it is desirable to increase the OMA of an output pulse pattern, which can be done in two ways. First, OMA can be increased by increasing the input optical power at λ_{in} . (We considered an input optical power of 5 dBm for our results discussed in previous paragraph). Second, OMA can be increased by decreasing the full width at half maximum (FWHM) or 3-dB bandwidth of the MRR-PEOLG. A lower FWHM would make the roll-off edges of the MRR passbands corresponding to (0,0), (0,1)/(1,0) and (1,1) steeper (Fig. 6.3), which in turn would increase the distance between the green dot and top red dot on the Y-axis, thereby increasing the OMA. Note that increasing OMA is not always necessary, as a low OMA would cause reliability issues only if it is lower than the OMA sensitivity (SOMA) of the receiver circuit that is employed to make sense of the output pulse pattern. Therefore, decreasing SOMA of the receiver circuit can also increase the reliability of our MRR-PEOLG. Thus, the FWHM (3-dB bandwidth) of the MRR-PEOLG, the SOMA of the receiver circuit, and the input power are the three factors that influence the impact of OMA on the reliability of our MRR-PEOLG.

Moreover, these three factors impact the maximum speed (bit-rate) at which the input pulse patterns can be driven. Increasing the bit-rate will reduce OMA because either the free-carrier concentration in the PN junctions or the optical energy inside the MRR does not change as fast as the applied input electrical pulse signals. For a given FWHM (3-dB bandwidth), it is possible to keep increasing the bit-rate until the OMA becomes less than the SOMA limit of the receiver circuit. Once the OMA for a given input power crosses the SOMA limit, the OMA can be increased to support a higher bit-rate by increasing the input power. Therefore, the maximum achievable bit-rate for our MRR-PEOLG depends on SOMA, FWHM, and input optical power. We have evaluated the maximum achievable bit-rate for our MRR-PEOLG, corresponding to various logic functions, which is discussed in next section.

6.4 Performance Analysis

For this analysis, we have used the scripting capabilities available in the ANSYS/Lumerical tools to run a performance evaluation of our MRR-PEOLG. We swept the input optical power in the range from -5 dBm to 5 dBm. Similarly, we swept SOMA in the range from -5 dBm to -20 dBm. Then, for each combination of the input optical power and SOMA, we evaluated the maximum achievable bit-rate for each logic-gate

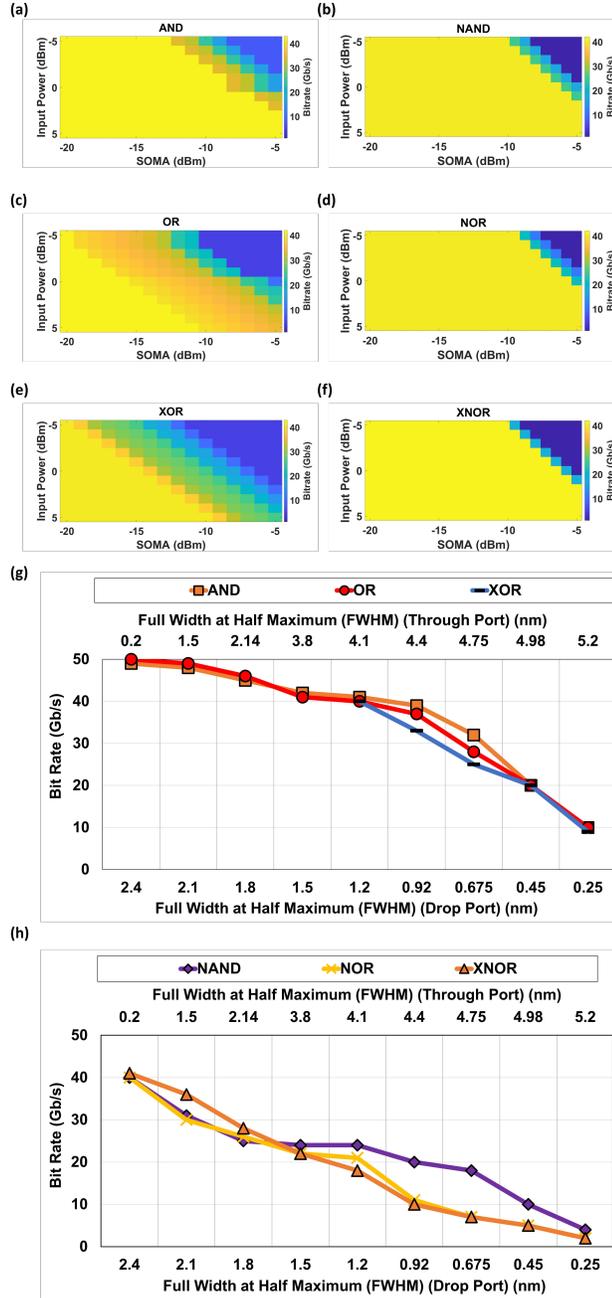


Figure 6.5: Colormap plots for logic functions (a) AND, (c) OR, (e) XOR (obtained at the drop port of our MRR-PEOLG), and complementary logic functions (b) NAND, (d) NOR, (f) XNOR (obtained at the through port of our MRR-PEOLG) that depict the maximum achievable bit-rate for given input optical power and SOMA. These color maps are evaluated for drop-port FWHM of 1.2 nm. We also report the maximum achievable bit-rate corresponding to (g) AND, OR, and XOR functions, and (h) NAND, NOR, and XNOR functions, evaluated for different values of FWHM, 0 dBm input optical power, and -5 dBm SOMA.

function supported by our MRR-PEOLG. The results of this analysis are shown in Fig. 6.5 in the form of colormap plots.

Results and Discussion

The colormap plots in Figs. 6.5(a) to 6.5(f) depict the maximum achievable bit-rate corresponding to each logic-gate function for an FWHM of 1.2 nm and different combinations of input optical power and SOMA. From the colormap plots, the AND function achieves a maximum bit-rate of 42 Gb/s across all SOMA values if the input optical power is >2 dBm, as well as across all input power values if the SOMA value is <-13 dBm. Similarly, OR and XOR functions achieve a maximum bit-rate of 41 Gb/s and 40 Gb/s respectively across all input optical power values if SOMA is <-19 dBm. Meanwhile, the NAND function achieves a maximum bit-rate of 40 Gb/s across all SOMA values if the input optical power is >2 dBm, as well as across all input power values if the SOMA value is <-11 dBm. Similarly, NOR and XNOR functions achieve a maximum bit-rate of 40 Gb/s and 41 Gb/s respectively across all input optical power values if SOMA is <-11 dBm. Moreover, we also show in Figs. 6.5(g) and 6.5(h) that increasing the drop-port FWHM (which can be achieved by increasing the cross-coupling co-efficient of the MRR-PEOLG) can increase the maximum achievable bit-rate for each logic function. These results imply that our MRR-PEOLG can be operated at up to 40 Gb/s for each of its supported logic-gate functions.

6.5 Comparison with E-O Circuits from Prior Work

We evaluated how the use of our MRR-PEOLG impacts the area, latency, and energy consumption of two E-O circuits from prior works [297] and [225]. We replaced the E-O XNOR gates with our MRR-PEOLG in the E-O XNOR-POP circuits of the binary neural network accelerator LightBulb from [297]. Similarly, we replaced the AND gates with our MRR-PEOLG in the optical bit-serial multiplier circuits of the digital CNN accelerator from [225]. As a result, the performance of these E-O circuits substantially improved as shown in Table 6.2. The energy values are the energy per bit values and include the MRR static power as well as laser power. The area and energy benefits in Table 6.2 are due to the compactness and better operand handling of our MRR-PEOLG and also our MRR-PEOLG’s ability to realize different logic functions with only a single MRR. The latency benefits are due to the fact that our MRR-PEOLG can operate at up to 40 Gb/s, whereas the original bit-serial multiplier circuit from [225] can only operate at up to 10 Gb/s. The E-O XNOR-POPCOUNT units from [297] can operate at a higher bit-rate of 50 Gb/s, but our MRR-PEOLG based variants provide better area-energy-delay product. These results corroborate the excellent capabilities and efficiency benefits of our MRR-PEOLG.

Table 6.2: Performance comparison of E-O circuits.
A=Area, E=Energy, L=Latency

Metrics	XNOR-POPCOUNT		Bit-serial Multiplier	
	[297]	MRR-PEOLG	[225]	MRR-PEOLG
A (mm ²)	0.013	0.011 (1.16×)	0.023	0.011 (2.08×)
E (nJ)	0.05	0.032 (1.53×)	0.327	0.033 (9.89×)
L (ns)	0.02	0.025 (0.8×)	0.1	0.025 (4×)
A*E*L	1.3e-5	0.9e-5 (1.44×)	75.2e-5	0.91e-5 (82.6×)

6.6 Summary

In this chapter, we demonstrated a microring resonator based polymorphic electro-optic logic gate (MRR-PEOLG) that can be dynamically reconfigured to implement different logic functions at different times. We modeled our MRR-PEOLG using the photonics foundry-validated simulation tools from ANSYS/Lumerical. Using these tools, we also performed frequency-domain, time-domain transient, and performance analysis of our MRR-PEOLG. From our analysis, we validated that our MRR-PEOLG design can implement various logic functions while operating at speeds of up to 40 Gb/s. Our evaluation shows that the use of our MRR-PEOLG in two E-O circuits from prior works can reduce their area-energy-delay product by up to 82.6×. We also show how our MRR-PEOLG can realize reconfigurable E-O SIMD/MIMD processing units.

Chapter 7 A Hybrid Time-Amplitude Analog Photonic GEMM Accelerator

7.1 Introduction

In recent years, the application of deep neural networks (DNNs) in a wide range of artificial intelligence tasks, such as image and speech recognition [98, 99], medical imaging [207], and conversational AI [72], has seen a substantial increase. This is primarily due to their superior inference accuracy and their ability to learn complex features from large amounts of data. However, DNNs are computationally intensive, due to inherently abundant linear computations such as general matrix-matrix multiplications (GEMM) and general matrix-vector multiplications (GEMV), which are at the core of DNN operations [69]. This computational intensity of DNNs is on a rapid rise owing to the ongoing rapid evolution of DNN models. As a result, the inference time of DNNs is also increasing. Although general-purpose compute engines such as graphics processing units (GPUs) have been the mainstay for DNN processing, the need to tackle this ever-increasing complexity and inference time of DNNs has led to the need for specialized hardware accelerators that are capable of efficiently performing GEMM and GEMV operations [246].

To implement these specialized hardware accelerators, the natural choice has been the use of CMOS-based electronic application-specific integrated circuits (ASICs). However, with the continuing slowdown of Moore’s law and the exponential increase of the complexity of DNN models (DNN complexity doubles every 3.4 months [181]), electronic ASIC accelerators are failing to keep up with the processing speed and energy efficiency requirements for large-scale deployment of complex DNN models [13, 181]. This has motivated researchers in industry and academia to develop new more-than-Moore technologies that can continue to provide persistently faster and energy-efficient hardware for the acceleration of DNNs for the foreseeable future.

Fortunately, Silicon Photonics (SiPh) has demonstrated remarkable potential and scalability for accelerating the GEMM and GEMV operations of DNNs. SiPh GEMM accelerators demonstrated in the literature employ linear photonic phenomena such as optical transmission and optical signal superposition to map GEMM operations onto the operating physics of photonic devices and circuits. Such physics-matched computing capability of SiPh GEMM accelerators renders them ultra-fast processing speed and sub-nanosecond input-to-output latency with $O(1)$ scaling law. Typically, a SiPh GEMM accelerator comprises multiple dot product units (DPUs) that operate concurrently, enabling parallel execution of multiple dot product operations created by unrolling the input GEMM operations. Several SiPh GEMM accelerators have been demonstrated in prior works based on various SiPh devices, such as the Mach Zehnder Interferometer (MZI) [244] and the Microring Resonator (MRR) [244]. However, the use of Mach-Zehnder Interferometers (MZIs) introduces a significant area overhead, making MZI-enabled SiPh DNN accelerators impractical for scaling in large-scale neural networks [244, 53].

Among prior SiPh GEMM accelerators, MRR-enabled accelerators have shown disruptive performance and energy efficiencies, due to the compact footprint of MRRs, low dynamic power consumption of MRRs, and the ability of MRRs to support a large fan-in of optical signals through dense-wavelength-division multiplexing (DWDM). Typically, an MRR-enabled SiPh GEMM accelerator consists of an array of photonic waveguides. Each waveguide engages with a bank of input microring modulators (MRMs) and a bank of weighting MRRs. Each MRM in the bank electro-optically modulates a sequence of amplitude-encoded electrical inputs (typically created by digital-to-analog converters (DACs)) onto an optical wavelength carrier to create a high-speed analog optical signal. A large fan-in of such high-speed optical signals is created using DWDM by aggregating parallel wavelength carriers in a single photonic waveguide. These signals then pass through a bank of weighting MRRs, where each MRR thermo-optically weights the amplitude envelope of its corresponding optical signal so that the amplitude of each symbol (period) of the signal experiences scalar weighting of the same amount. Thus, the amplitude of each symbol (period) of a weighted optical signal represents the analog product of an input and a scalar weight. These weighted optical signals are then propagated to a balanced photodetector at the end of the photonic waveguide. The balanced photodetector, during each period of its operation, performs incoherent superposition of all amplitude symbols that arrive on multi-wavelength weighted optical signals. Through incoherent superposition, the balanced photodetector essentially generates an amplitude of the electrical output current that is proportional to the signed sum of all the optical amplitude symbols incident during each period. Because each incident optical symbol represents an analog product, the balanced photodetector thus generates a signed sum of products, i.e., a dot product, during each period. Multiple waveguides per GEMM accelerator thus enable multiple parallel dot products to be generated by the accelerator in each signal period. This ability to massively parallelize processing has been shown to have rendered up to $1000\times$ more processing throughput and up to $100\times$ better energy efficiency to MRR-enabled SiPh GEMM accelerators than their CMOS-based electronic counterparts [26, 68].

However, the state-of-the-art MRR-enabled SiPh GEMM accelerators face the following shortcomings that hinder the realization of their full potential. (1) Generation of each high-speed weighted optical signal requires two dedicated devices, i.e., one MRM and one MRR. As a result, the aggregation of a total of N optical signals per waveguide may require each optical signal to actively or passively engage with up to $2N$ devices (MRMs+MRRs). This may increase the total losses faced by each optical signal because each signal-device interaction incurs a certain level of insertion loss. The increase in total losses is likely to demand more optical power per signal, thereby diminishing the energy efficiency advantages. (2) The increase in the required optical power per signal also whittles down a larger part of the limited optical power budget of the accelerator, which in turn reduces the affordable spatial parallelism in the accelerator, both in terms of the number of waveguides and number of carrier wavelengths per waveguide. The decreased spatial parallelism is likely to reduce the processing throughput of the accelerator. (3) It is established in prior works [82, 195] that to maintain the high-speed and low-latency benefits of SiPh accelerators, it is

necessary to allow the weighting of each amplitude symbol of every input high-speed optical signal by a different amount instead of allowing only static weighting. Allowing this would require fast electro-optic actuation of weighting MRMs, necessitating an additional feedback control unit per weighting MRM for enabling electro-optic actuation [195]. This would increase the number of required feedback control units per weighted optical signal to four because both the input MRM and weighting MRM would require one feedback control unit each for thermal stabilization [82] and another unit each for electro-optic actuation [195]. Requiring four feedback control units per signal would increase the cost of hardware implementation and power consumption, further diminishing the energy efficiency advantages.

To overcome these shortcomings, we present the following three innovations in this chapter. *First*, we enable the generation of a weighted optical signal using a single MRM instead of using one input MRM and one weighting MRM. For that, we introduce a novel hybrid Time-Amplitude analog Optical Modulator (TAOM) that employs a single MRM. A TAOM generates a weighted high-speed optical signal as a temporal sequence of pulse-width-amplitude-modulated (PWAM) symbols. In each PWAM symbol generated by a TAOM, an input value is encoded as the pulse width of the symbol, and a weight value is encoded as the analog amplitude of the symbol. The total optical energy contained in the PWAM symbol represents the analog product of the input and weight values. *Second*, we introduce a novel balanced photo-charge accumulator (BPCA) circuit that leverages the in-situ charge accumulation and incoherent superposition abilities of photodetectors [43][230] to generate a signed summation of a large number of temporally and spatially arriving PWAM symbols. *Third*, we organize our invented TAOMs and BPCAs in 2D arrays to design a SiPh GEMM accelerator architecture that achieves higher spatial processing parallelism to realize significantly higher processing throughput at better power efficiency compared to the existing SiPh GEMM accelerators.

Our key contributions in this chapter are summarized below.

- We present the structure and operation of our invented hybrid time-amplitude analog optical multiplier (TAOM) and balanced photo-charge accumulator (BPCA);
- We integrate TAOM and BPCA to create a dot-product processing circuit and a GEMM processing architecture, and then perform an extensive device-level, circuit-level, and system-level analysis of our invented circuit and architecture;
- At the device level, we perform detailed modeling and characterization of our invented TAOM using photonics foundry-validated, commercial-grade tools from ANSYS/Lumerical [155]. We also evaluate the achievable accuracy and bit precision of our TAOM by performing transient simulations in the INTERCONNECT solver of ANSYS/Lumerical [155];
- At the circuit level, we analyze the effect of inter-modulation crosstalk in TAOM+BPCA-enabled DWDM-based dot-product circuit on its corresponding dot-product result by performing a transient circuit simulation in the INTERCONNECT solver of ANSYS/Lumerical [155];

- At the system level, we design an accelerator architecture called hybrid time-amplitude analog optical multiplier-based tensor core (TAOM-TC), and evaluate its achievable spatial parallelism, power consumption, and performance. Furthermore, we compare the evaluation results of our TAOM-TC with two well-known SiPh GEMM accelerators from prior works.

7.2 Background

The SiPh GEMM accelerators showcased in the literature predominantly rely on two approaches: either utilizing MZIs [204, 228, 162, 223, 179] that leverage optical-interference or employing MRRs [8, 26, 243, 226] that harness optical resonance to perform matrix computations. A concise summary of accelerators built upon these devices is provided below.

Mach-Zehnder Interferometer (MZI) Based Accelerators

The MZI-based accelerators that have been demonstrated thus far in the literature are primarily coherent architectures, meaning that they rely on the manipulation of electric field amplitude and phase of an optical wavelength signal. MZI-based coherent architectures employ universal linear meshes of MZIs to efficiently implement the dot product operations. In these architectures, the weights are controlled by controlling the phase and amplitude of optical signals via the attenuators and phase shifters integrated into the arms of the MZIs, as demonstrated in [204]. Similarly, Shokraneh et al. in [228] demonstrated an MZI-based 4×4 optical matrix multiplier which was used for constructing a single-layered neural network. Furthermore, Shen et al. and Miller et al. in [223] and [162], respectively, demonstrated a singular value decomposition (SVD) technique with MZI meshes to perform matrix multiplications. SVD technique is the process utilized to perform decomposition of matrices into unitary matrices, that are encoded into the intensity and phase of light and then fed into each layer of the MZI mesh network. As mentioned earlier, MZI-based architectures are primarily coherent and utilize only a single wavelength. However, MZIs have been used to implement WDM-based, incoherent architectures as well. For instance, On et al. in [179] demonstrated a photonic matrix multiplication accelerator for recurrent neural networks (RNNs) using MZIs and Arrayed Waveguide Grating (AWG)-Multimode Interference couplers (MMIs). In this architecture, MZIs functioned as intensity modulators, while AWG-MMI coupler units, along with a coherent detection scheme, performed matrix multiplication.

Although MZIs-based coherent GEMM accelerators leverage SVD to reduce the complexity and dimensionality of GEMM operations of DNNs [162], they require precise manipulation of the electric field phase and amplitude of an optical signal to ensure reliable, accurate, and efficient matrix multiplications. Additionally, MZIs have a larger footprint (tens to hundreds of micrometers) compared to other photonic devices (e.g., MRMs and MRRs), potentially limiting the scalability of MZIs-based coherent GEMM accelerators for processing large-scale matrices. In addition, MZIs are also susceptible to phase noise errors [244]. Furthermore, the fabrication

process variations can introduce minor deviations in the dimensions of MZI arms, impacting the power splitting ratios in the arms of MZIs. These effects negatively impact the accuracy and performance of MZI-based coherent GEMM accelerators [163]. In contrast, the MRR-based incoherent GEMM accelerators offer better scalability and lower footprint because they employ photonic integrated circuits that are based on compact MRRs. Several MRR-based incoherent GEMM accelerators have been demonstrated in the prior works, which are highlighted in the next subsection.

Microring Resonator (MRR) Based Accelerators

The photonic MRR-based incoherent CNN accelerators, demonstrated thus far in the literature mainly employ multiple analog tensor processing cores (TPCs) that operate in parallel, in which each TPC is utilized to perform a dot product operation. Typically, each TPC is made up of five essential blocks (Fig. 7.1): (i) a laser block that employs N laser diodes (LDs) to generate N optical wavelength channels; (ii) an aggregation block that aggregates the optical wavelength channels generated by LDs into a single photonic waveguide through DWDM technique by employing a $N \times 1$ multiplexer, and then splits the optical power of each of these wavelength channels equally into M separate waveguides by using a $1 \times M$ splitter; (iii) a modulation block that consists of M arrays of MRMs spread across M waveguides, with each waveguide employing one MRM array; each MRM array electro-optically modulate the incoming optical wavelength channels with input values; (iv) a weighting block that consists of another M arrays of MRRs spread across M waveguides, with each waveguide employing one MRR array; each MRR array performs modulation (weighting) of the input-modulated optical wavelength channels with weight values, thereby performing an element-wise product of the input and weight values; and (v) a summation block that comprises of a total of M summation elements (SEs), with each SE employing two photo diodes in a balanced configuration, commonly referred to as balanced photo diode (BPD) configuration, connected to a transimpedance amplifier (TIA) and an analog-to-digital converter (ADC). Typically, the laser block and SE block are placed at the two ends of the TPC, whereas the aggregation, modulation, and weighting blocks are placed in between them. Moreover, based on the positioning of these intermediate blocks, the MRR-based TPC organizations demonstrated in the prior works can be classified into two categories namely *Aggregate, Modulate, Weight (AMW)* TPC and *Modulate, Aggregate, Weight (MAW)* TPC. Fig. 7.1(a) depicts the AMW TPC organization in which the aggregation block is positioned first, followed by the modulation and the weighting blocks. On the other hand, Fig. 7.1(b) illustrates the MAW TPC organization, in which the modulation block is positioned first, followed by the aggregation and the weighting blocks. Notably, the modulation block of the MAW TPC organization employs only one MRM per waveguide, enabling the imprinting of one input value per wavelength channel (N input values onto the N wavelength channels), which is then equally shared among the M waveguides. In contrast, the AMW TPC organization, also comprising of M waveguides, incorporates a dedicated MRM input array and an MRR weight bank cascaded to each waveguide.

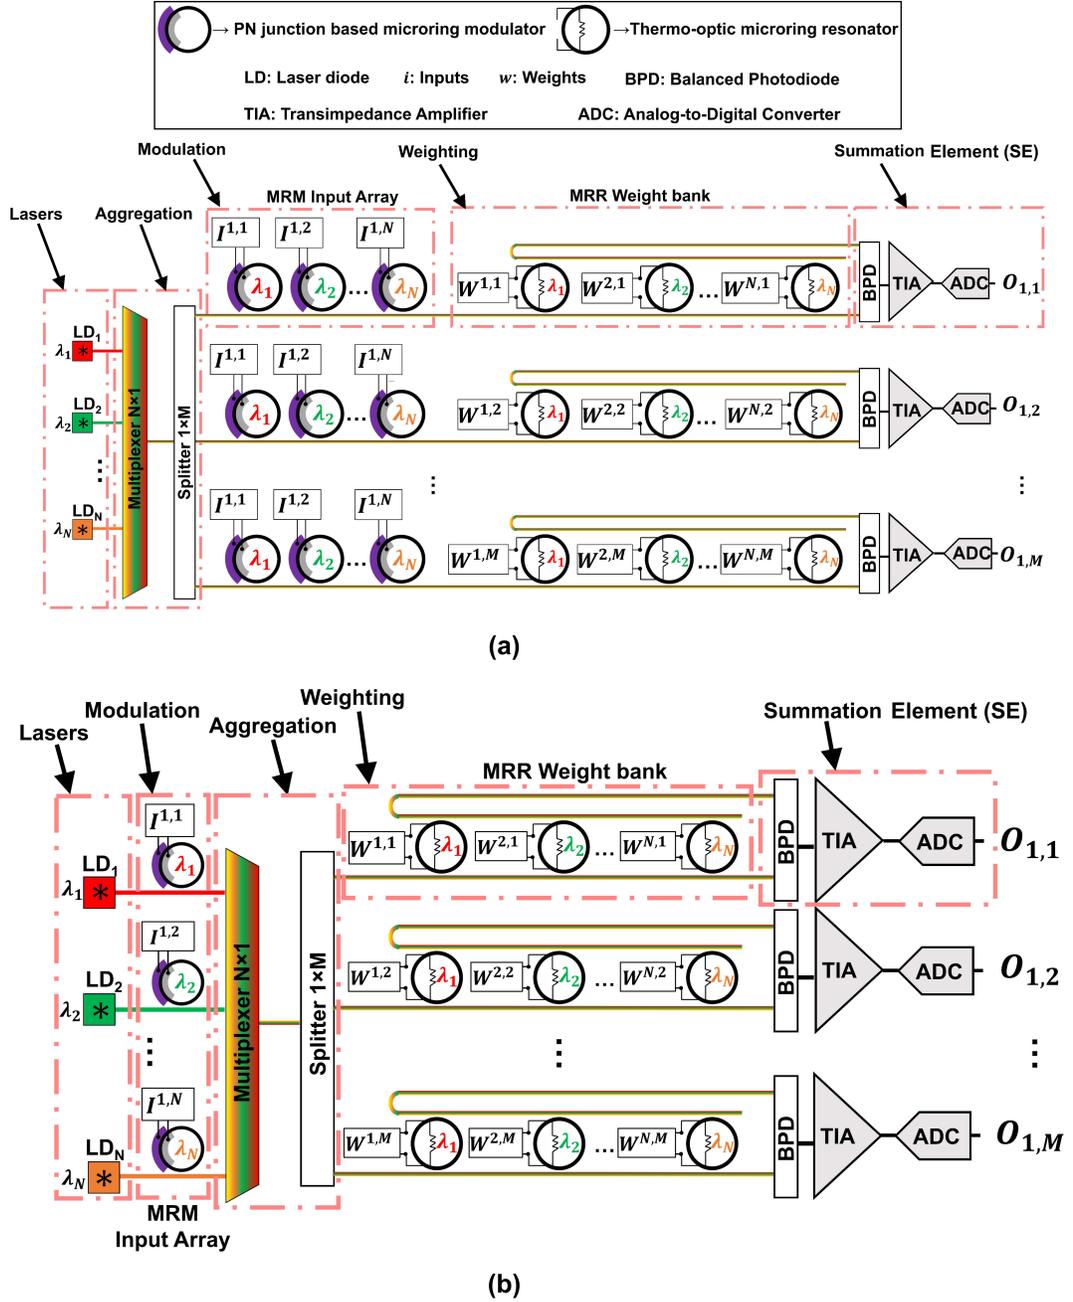


Figure 7.1: Illustration of common MRR based analog optical TPC organizations. (a) AMW TPC and (b) MAW TPC.

Fig. 7.2, in relation to Fig. 7.1, collectively provide a visual representation of how the input and weight matrices are mapped onto the AMW and MAW TPCs. As illustrated in Fig. 7.2, the rows within the input matrix are mapped onto the optical wavelength channels ($\lambda_1, \lambda_2, \dots, \lambda_N$) and each column within the input matrix represents a temporal vector (highlighted in Red in Fig. 7.2). On the other hand, the different elements across the rows within the weight matrix are distributed across the

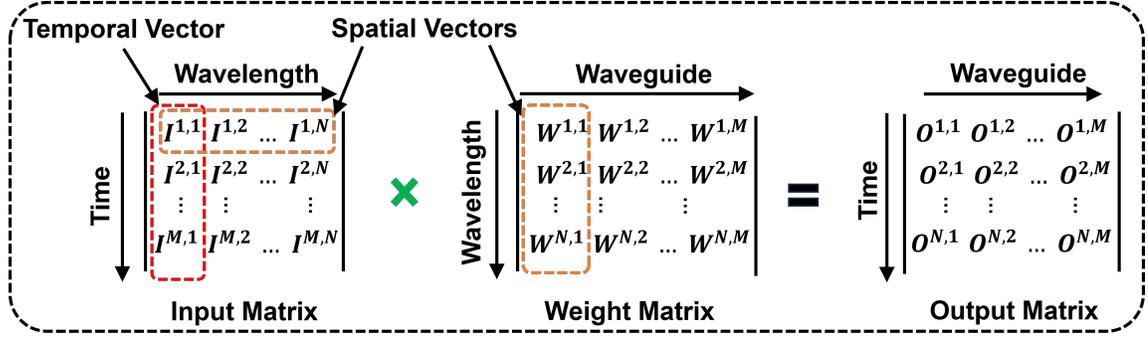


Figure 7.2: Mapping of the input and weight matrices onto the AMW and MAW TPCs.

waveguides, while the columns within the weight matrix are mapped onto the wavelength channels ($\lambda_1, \lambda_2, \dots, \lambda_N$), as depicted in Figs. 7.1 and 7.2. Within these TPCs, the MRMs electro-optically modulate the inputs onto the optical wavelength signals such that the intensity of the wavelength signals represents the input values. Here, each high-speed optical signal generated by MRM represents a temporal analog vector (highlighted in Red in Fig. 7.2). Furthermore, the MRMs within the MRR weight bank feature tunable MRR filters that can be thermo-optically adjusted to perform the static weighting of the input-encoded signals. As a result, each statically-weighted input-encoded optical wavelength signal represents a temporal product vector, which is essentially the product of the input temporal analog vector and the scalar weight. Thus, the creation of a temporal product vector enables a vector-scalar multiplication in the time domain. When multiple such temporal product vectors are created across multiple wavelength channels using DWDM, it creates a spatial fan-in of temporal product vectors in a single photonic waveguide. Subsequently, these multi-wavelength temporal product vectors within each waveguide propagate towards their respective SE. The BPD present in each SE generates an electrical current amplitude which represents a signed sum of analog product values of the corresponding spatial product vector i.e., a dot product. This amplitude is further processed by a TIA and an ADC within the SE converting it into digital domain. The output at the end of each SE in Fig. 7.1 collectively form the row vector in the output matrix illustrated in Fig. 7.2.

Furthermore, Sunny et al. in [243] demonstrated an MRR-based incoherent GEMM accelerator designed to be resilient to on-chip fabrication process variations (FPV) and thermal variations. It utilizes FPV-resilient MRR designs coupled with a thermal Eigen decomposition-based tuning approach, and an intelligent MRR placement to combat variations. Furthermore, Shiflett et al. in [226] demonstrated an accelerator that utilizes both MZIs and MRMs to perform matrix multiplications. In summary, MRR-based incoherent GEMM accelerators offer scalability, occupy less area, and achieve superior performance compared to the MZI-based implementations due to the compact footprint of MRMs and their compatibility with cascaded DWDM.

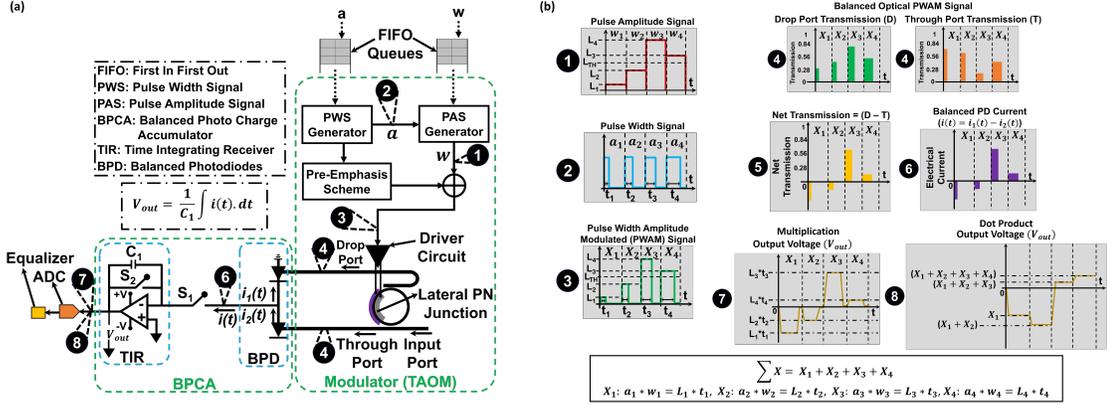


Figure 7.3: (a) Device-level schematic of our microring modulator (MRM) based hybrid time-amplitude analog optical modulator (TAOM) integrated with a balanced photo charge accumulator (BPCA) and (b) analog representation of signals (optical and electrical) at different stages of our integrated TAOM+BPCA unit.

7.3 A Hybrid Time-Amplitude Analog Optical Modulator (TAOM)

Device-Level Schematic and Operation

Schematic

Fig. 7.3(a) depicts the schematic of our invented TAOM when it is connected to a balanced photo-charge accumulator (BPCA) unit. As illustrated, our TAOM is basically an add-drop Microring Modulator (MRM) with an embedded lateral PN junction that operates in the forward bias condition. The MRM's peripheral circuitry consists of two queues of FIFO buffers, in which one of them stores the input values (from the input matrix shown in Fig. 7.2) and the other one stores the weight values (from the weight matrix shown in Fig. 7.2), both in the digital binary-radix number format. The FIFO queue for inputs connects to a pulse width signal (PWS) generator and the FIFO queue for weights connects to a pulse amplitude signal (PAS) generator. The output of the PWS is split into two parts: one part is directed to the pre-emphasis scheme, whereas the other part is provided as a reference to the PAS generator. Subsequently, the output of the PAS generator and the output of the pre-emphasis scheme are combined through a current-mode mixer, and the resulting output is a pulse-width-amplitude-modulated (PWAM) signal. For a complete understanding of the generation of PWAM signals and the underlying circuitry, we direct the readers to [130, 283]. This PWAS signal is routed to a driver circuit. The output of the driver circuit is provided as an electrical bias to the PN junction of the MRM. The output of the MRM (TAOM) is connected to a balanced photo charge accumulator (BPCA) circuit.

Our BPCA circuit is collectively inspired by the time integrating receiver (TIR) design from [230, 1] and the photodetector-based optical pulse accumulator design from [43]. As illustrated in Fig. 7.3(a), a BPCA circuit employs two photodiodes, each connected to the drop and through ports of the MRM. These photodiodes are

interlinked in a balanced configuration, commonly referred to as a balanced photodiode (BPD) configuration. The BPD is connected to a TIR via a switch (S_0). The TIR comprises an amplifier and a feedback capacitor/switch (S_1) pair (Fig. 7.3(a)). It functions as a current-to-voltage converter circuit by integrating the incoming electrical current over a period. This ensemble of the BPD and TIR makes the BPCA capable of performing temporal and spatial accumulations (this will be explained in upcoming subsections). Subsequently, the output of the TIR is connected to an analog-to-digital converter (ADC) and an equalizer.

Operation

(i) Electrical PWAM Signal Generation: Fig. 7.3(b) illustrates the sequential processing of electrical and optical signals at various stages within our integrated TAOM-BPCA circuit, demonstrating the effective execution of the multiplications and temporal accumulations. As illustrated, the FIFO queue for weight values feeds into the PAS generator, which converts the incoming sequence of digital weight values into a sequence of analog pulse amplitude symbols. This sequence is also called a pulse amplitude signal ((see ❶ in Figs. 7.3(a) and 7.3(b))). Similarly, the FIFO queue for input values feeds the PWS generator, which converts the incoming sequence of digital input values into a sequence of analog pulse-width-modulated symbols. This sequence is called a pulse width signal (see ❷ in Figs. 7.3(a) and 7.3(b)). The PWS output is divided, with one part directed to the pre-emphasis scheme, while the other part serves as a reference for the PAS generator. The output of the PAS generator (current-mode DAC), when mixed with the output of the pre-emphasis scheme, produces a sequence of pulse width amplitude modulated (PWAM) symbols ((see ❸ in Fig. 7.3(b))). This sequence is called PWAM signal. For a complete understanding of the generation of PWAM signals and the underlying circuitry, we direct the readers to [130, 283]. The PWAM signal is fed to a driver circuit, as shown in Fig. 7.3(a). From the driver circuit, the PWAM signal is provided as an electrical input to the PN junction of the MRM.

(ii) Electrical-to-Optical Signal Conversion/Balanced Optical PWAM Signal Generation: The input electrical PWAM signal induces free-carrier plasma dispersion in the MRM, which enables the MRM to dynamically adjust the transmission characteristics of the incoming wavelength channel from an external laser source. This action converts the input electrical PWAM signal into a balanced optical PWAM signal. Here, a balanced optical PWAM signal implies that the original electrical PWAM signal is encoded into optical transmissions simultaneously at both the drop and through ports of the MRM ((See ❹ in Fig. 7.3(b))).

For each symbol of this balanced optical PWAM signal, the amount of transmission at the through and drop ports of the MRM depends on the amplitude level of each symbol relative to the threshold level in the electrical PWAM signal (L_{TH} in ❸ of Fig. 7.3(b)). For instance, the amplitudes of symbols \mathbf{X}_1 and \mathbf{X}_2 in ❸ of Fig. 7.3(b) are below the defined threshold level (L_{TH}). Therefore, for symbols \mathbf{X}_1 and \mathbf{X}_2 in ❹ of Fig. 7.3(b), the transmission at the drop port of the MRM is lower than the transmission at the through port. As a result, the net transmission, represented

as the difference in transmission between the drop and through ports of the MRM (\mathbf{X}_1 and \mathbf{X}_2 in ⑤ of Fig. 7.3(b)), is negative for these symbols. On the other hand, the amplitudes of symbols \mathbf{X}_3 and \mathbf{X}_4 in ③ of Fig. 7.3(b) are above the defined threshold level (L_{TH}). Therefore, for the symbols \mathbf{X}_3 and \mathbf{X}_4 in ④ of Fig. 7.3(b), the transmission at the drop port of the MRM is higher than the transmission at the through port. As a result, the net transmission (\mathbf{X}_3 and \mathbf{X}_4 in ⑤ of Fig. 7.3(b)), is positive for these symbols. Each such symbol of a balanced optical PWAM signal (such as \mathbf{X}_1 in ④ of Fig. 7.3(b)) packetizes certain optical energy that is proportional to the analog product of the corresponding input (a) and weight values (w). For example, the energy packetized in symbol \mathbf{X}_1 represents $\mathbf{a}_1 * \mathbf{w}_1$ (or) $\mathbf{L}_1 * \mathbf{t}_1$ in Fig. 7.3(b). This balanced optical PWAM signal at the output of the TAOM is fed into the BPCA circuit.

(iii) Extraction of Multiplication/Dot Product Result: Within the BPCA circuit, the BPD transduces the incoming optical pulse sequence ($\mathbf{X}_1, \dots, \mathbf{X}_4$ in ⑤ of Fig. 7.3(b)) from the MRM into a series of differential electrical current amplitudes (⑥ in Fig. 7.3(b)). Here, the differential electrical current amplitude corresponding to each symbol is proportional to the net transmission of the respective optical PWAM symbol (⑤ in Fig. 7.3(b)). Moreover, the multiplication magnitude corresponding to each symbol is encoded as the area under the curve of the differential electrical current symbol. The direction of the electrical current symbol (incoming to the BPD and outgoing from the BPD) represents the sign of the multiplication. This series of differential electrical current symbols is directed towards the TIR via the switch S_0 . The integration of TIR with the BPD introduces a distinctive versatility that enables us to operate our integrated TAOM-BPCA circuit in one of the two distinct modes: (i) multiplier mode or (ii) multiplier and temporal accumulator mode. These modes are explained next.

(a) Multiplier Mode: For this mode, the TIR's sampling speed is matched to the arrival rate of the incoming differential electrical current symbols. At the beginning of each symbol period, opening the switch S_1 allows the electrical current symbol corresponding to that period to linearly charge the feedback capacitor C_1 of the TIR circuit. This linear charging continues for a duration equal to the pulse width of the electrical current symbol. Consequently, the accumulated charge, and therefore, the analog voltage accrued across the feedback capacitor C_1 of the TIR for that symbol period represents the multiplication result related to the corresponding optical PWAM symbol (see ⑦ in Fig. 7.3(b)). Notably, the polarity of the incoming differential electrical current pulse indicates the sign of the multiplication result; therefore, the polarity of the accrued analog voltage across the TIR's feedback capacitor becomes negative if the incoming electrical current has negative polarity. Once the accrued analog voltage is stable, it is sampled and sent to an analog-to-digital converter (ADC). Then, at the end of the symbol period, closing the switch S_1 allows resetting the charge and voltage on the feedback capacitor to be zero, to prepare the TIR for the next symbol period.

For example, consider the symbol \mathbf{X}_1 in ⑦ of Fig. 7.3(b). The feedback capacitor of the TIR circuit linearly charges until it reaches an analog voltage level of ($\mathbf{L}_1 * \mathbf{t}_1$),

which represents the signed multiplication result of the symbol \mathbf{X}_1 . After the analog voltage level on the capacitor reaches $(L_1 * t_1)$ for the symbol \mathbf{X}_1 and the capacitor reaches a steady state, the accrued analog voltage is sampled and then sent to the ADC and equalizer for further processing, before the capacitor is made to discharge (reset) by closing the switch S_1 . Here, since the polarity of the differential electrical current corresponding to symbol \mathbf{X}_1 is negative, the accrued analog voltage on the TIR's feedback capacitor is also negative, as shown for symbol \mathbf{X}_1 in ⑦ of Fig. 7.3(b). On the other hand, if the polarity of the incoming differential electrical current pulse is positive, the accrued analog voltage on the TIR's feedback capacitor is also positive, which is illustrated for symbol \mathbf{X}_3 in ⑦ of Fig. 7.3(b). Thus, in this mode, the TAOM-BPCA circuit acts as a multiplier that can produce multiplication results at a fast speed.

(b) *Multiplier and Temporal Accumulator mode:* For this case, the TIR's sampling speed is set to be very low compared to the arrival rate of the incoming differential electrical current symbols. Therefore, the series of differential electrical current symbols arriving at the TIR can sequentially charge the TIR's capacitor so that the net accumulated charge and, consequently, the analog voltage accrued on the capacitor over multiple symbol periods provides the signed sum of the individual multiplication results corresponding to different symbols. Thus, this operation essentially performs a temporal accumulation of multiplication results (products). This operation is depicted in ⑧ of Fig. 7.3(b), where the charge accumulates over time based on the incoming electrical current symbols ($\mathbf{X}_1, \dots, \mathbf{X}_4$ in ⑥ of Fig. 7.3(b)), and consequently, the resulting analog voltage accrued on the TIR's capacitor signifies the temporal accumulation result (i.e., $\mathbf{X}_1 + \dots + \mathbf{X}_4$). Moreover, if the incoming differential electrical current pulses to the TIR circuit include both positive and negative polarities, the resultant analog voltage accumulated on the capacitor over time, representing the temporal accumulation operation, is a summation of positive and negative voltages.

For instance, the temporal accumulation operation illustrated in ⑧ of Fig. 7.3(b) is a summation of both negative and positive voltages. The first two symbols of the incoming current signal (i.e., \mathbf{X}_1 and \mathbf{X}_2 in ⑥ of Fig. 7.3(b)) have negative polarity. Therefore, in ⑧ of Fig. 7.3(b), the net accrued voltage on the capacitor at the end of the second symbol period has negative polarity. The magnitude of this voltage represents $\mathbf{X}_1 + \mathbf{X}_2$. This is because, unlike multiplier mode, the switch S_1 is not closed every symbol period (rather it is kept open) during the operation of this multiplier+temporal accumulator mode. As a result, the accrued voltage does not return to zero at the end of every period; rather, it builds on top of the voltage level accrued in the previous period. In ⑥ of Fig. 7.3(b), the collective magnitude of the differential electrical current corresponding to the symbols \mathbf{X}_3 and \mathbf{X}_4 is high compared to that of the symbols \mathbf{X}_1 and \mathbf{X}_2 . Therefore, the resultant voltage accrued on the capacitor has positive polarity after all four symbol periods have elapsed. This voltage corresponds to the result of the temporal accumulation of incoming PWAM symbols. Since each PWAM symbol encodes a multiplication result, this temporal accumulation result is basically a temporal dot-product result. This dot-product result can be collected by sampling the accrued voltage and then directing it to the ADC and equalizer for further processing. After the desired number of periods for

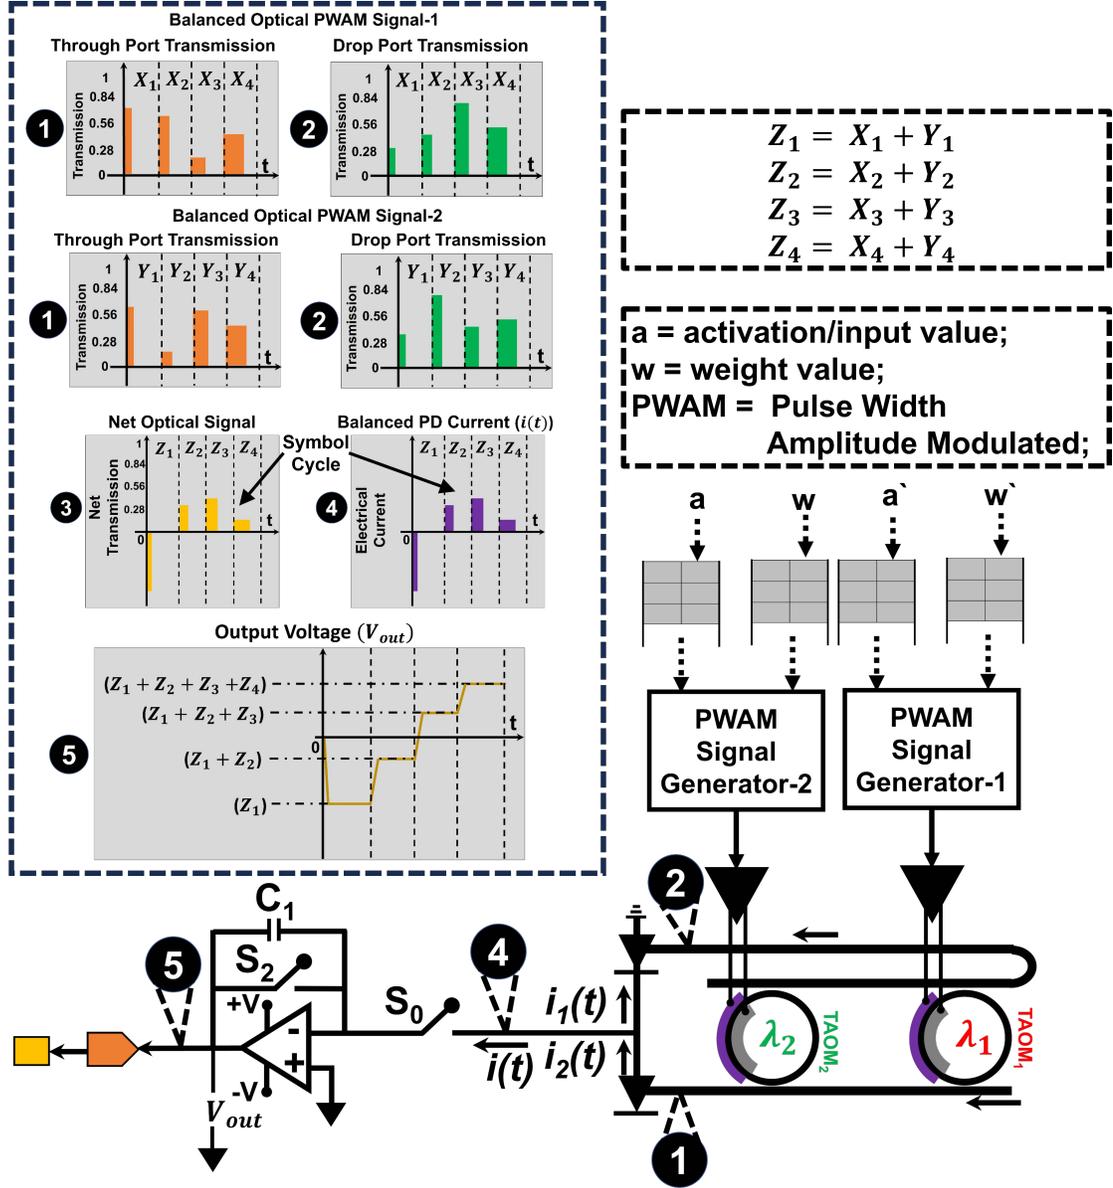


Figure 7.4: Circuit-level schematic of our integrated TAOM-BPCA unit which consists of two cascaded TAOMs, connected to our BPCA circuit. The inset showcases analog representations of signals (both optical and electrical) at various stages of our circuit.

the temporal accumulations have elapsed, the switch S_1 is closed to reset/discharge the capacitor, to prepare the circuit for the next temporal accumulation.

We model our TAOM unit using the photonics foundry-validated simulation tools from ANSYS/Lumerical [155]. Here, we perform a time-domain (transient) analysis to evaluate the accuracy and precision of our TAOM unit for different values of optical power and sample rates. Detailed discussion regarding the modelling, simulation and analysis is provided in the upcoming subsections.

Spatio-Temporal Multiply-Accumulate Operations using Cascaded TAOMs

Previously, we demonstrated that the ensemble of one TAOM and one BPCA can be used to perform temporal multiply-accumulate operations. In this subsection, we extend that idea and demonstrate that an ensemble of multiple TAOMs and a shared BPCA, by employing wavelength-division-multiplexing (WDM), can be used to perform spatio-temporal multiply-accumulate operations. To understand this, consider Fig. 7.4 that illustrates the functioning of 2-channel TAOM circuit comprised of an ensemble of two TAOMs (TAOM₁ and TAOM₂) cascaded in a WDM manner and a shared BPCA. From what we know about TAOM-BPCA ensemble discussed earlier in subsections with respect to Fig. 7.3, TAOM₁ and TAOM₂ in the 2-channel TAOM circuit of Fig. 7.4 generate balanced optical PWAM signals that are carried onto the dedicated optical wavelengths λ_1 and λ_2 respectively. Both of these balanced optical PWAM signals (shown as optical transmissions in ❶ and ❷ of Fig. 7.4) are multiplexed into the respective through and drop ports of the TAOMs, which are guided towards the shared BPCA circuit. During each symbol cycle, the BPD of the BPCA performs a balanced incoherent superposition (signed summation) of all the balanced optical PWAM symbols that arrive during the symbol cycle. Consequently, the balanced incoherent superposition first enables the creation of a net optical signal (see ❸ in Fig. 7.4), and then, allows this net optical signal to be transduced to generate a balanced photocurrent signal (see ❹ in Fig. 7.4). The area under the curve of every balanced photocurrent symbol in ❹ of Fig. 7.4 gives a spatial accumulation result (a spatial sum) of the PWAM symbols incident during the corresponding symbol cycle. The polarity of each balanced photocurrent symbol gives the sign of the corresponding spatial sum. Thus, the BPD of the BPCA enables spatial accumulation of incident PWAM symbols. Since all PWAM symbols are multiplication results, their spatial accumulation at the BPCA can also be referred to as spatial multiply-accumulate (MAC) operation or spatial dot-product operation.

This balanced photocurrent signal (which can also be called photocurrent-based spatial MAC signal) produced by the BPD of the BPCA is sent to the TIR of the BPCA, where it can be further processed in two different ways. *First*, if the sampling rate of the TIR is kept equal to the symbol rate of the incoming balanced photocurrent signal, the TIR simply converts the photocurrent-based spatial MAC signal into a voltage-based spatial MAC signal. In this signal, each symbol simply is a voltage level representing a spatial dot-product result. *Second*, if the sampling rate of the TIR is kept to be an integer multiple of the symbol rate of the incoming photocurrent signal, the TIR enables the gradual integration (temporal accumulation) of the individual symbols of the photocurrent-based spatial MAC signal to provide a single voltage level as the final output that represents the temporal sum of all the individual symbols (spatial dot-product result) of the spatial MAC signal. This occurs due to the same operational characteristics of the TIR as discussed previously. Thus, the BPCA (BPD + slowly sampled TIR) enables spatio-temporal MAC or dot-product (i.e., temporal accumulations of spatial dot-product results) in a WDM-cascaded TAOM circuit.

This spatio-temporal accumulation capability of WDM-cascaded TAOM circuit is clearly illustrated in ❺ of Fig. 7.4. During the first symbol cycle, X_1 and Y_1

are spatially accumulated resulting in an accrued V_{out} that is proportional to $(\mathbf{X}_1 + \mathbf{Y}_1)$. Here, the cumulative polarity of $(\mathbf{X}_1 + \mathbf{Y}_1)$ is negative, thereby resulting in negative V_{out} . In the next symbol cycle, the balanced photocurrent generated during the cycle corresponds to spatially accumulated symbols $(\mathbf{X}_2 + \mathbf{Y}_2)$. This balanced photocurrent accrues voltage on top of the V_{out} accrued in the previous cycle, resulting in updated V_{out} that is proportional to $(\mathbf{X}_1 + \mathbf{Y}_1 + \mathbf{X}_2 + \mathbf{Y}_2)$. Thus, the temporal accrual of V_{out} over all four symbol cycles enables the final result to be $(\mathbf{X}_1 + \mathbf{Y}_1 + \mathbf{X}_2 + \mathbf{Y}_2 + \mathbf{X}_3 + \mathbf{Y}_3 + \mathbf{X}_4 + \mathbf{Y}_4)$. The polarity of V_{out} depends on the net polarity of the final result. The final V_{out} from the BPCA circuit is given as input to an ADC to produce the final output in the digital format.

Modelling and Simulation

Modelling

We modeled our TAOM unit using the photonics foundry-validated commercial grade simulation tools from ANSYS/Lumerical [155]. As previously mentioned, our TAOM unit comprises an add-drop MRM with an embedded lateral PN junction operating in forward bias. To model this lateral PN junction-based add-drop MRM, we used the steps provided in [3], which involves modeling various primitive elements using different solvers in the ANSYS/Lumerical suite. The modeled primitive elements include MRM-waveguide coupling sections, straight and bent passive waveguides, lateral PN-junction section, and an active bent waveguide. These primitive elements are then assembled together, as shown in Fig. 7.5, to form the whole MRM model. A brief description of each modeling step is provided below:

Step-1 - modeling MRM-waveguide coupling sections: Create coupling sections of the add-drop MRM in the finite difference time domain (FDTD) solver and extract the power coupling coefficients as a function of wavelength for the fundamental TE mode.

Step-2 - modeling straight and bent passive waveguide sections: Characterize the passive, straight, and bent channel waveguides of the MRM using the finite difference eigenmode (FDE) solver; extract the effective index, group index, and dispersion for the waveguides as functions of wavelength.

Step-3 - modeling lateral PN-junction section: Create a lateral PN junction in the CHARGE tool. Then, perform a 2D CHARGE simulation to obtain the change in density and spatial distribution of charge carriers as a function of bias voltage.

Step-4 - modeling active bent waveguide: Load into the FDE solver the change in density and spatial distribution of charge carriers vs the bias voltage, obtained from Step 3. Perform two simulations to characterize the active waveguides: (i) Set the bias voltage to zero and calculate the effective index, group index, and dispersion as a function of frequency, similar to passive waveguides in step 2; and (ii) Use the scripting capabilities in ANSYS/Lumerical suite, then calculate the change in the effective index as a function of bias voltage, at the center wavelength, by sweeping the bias voltage.

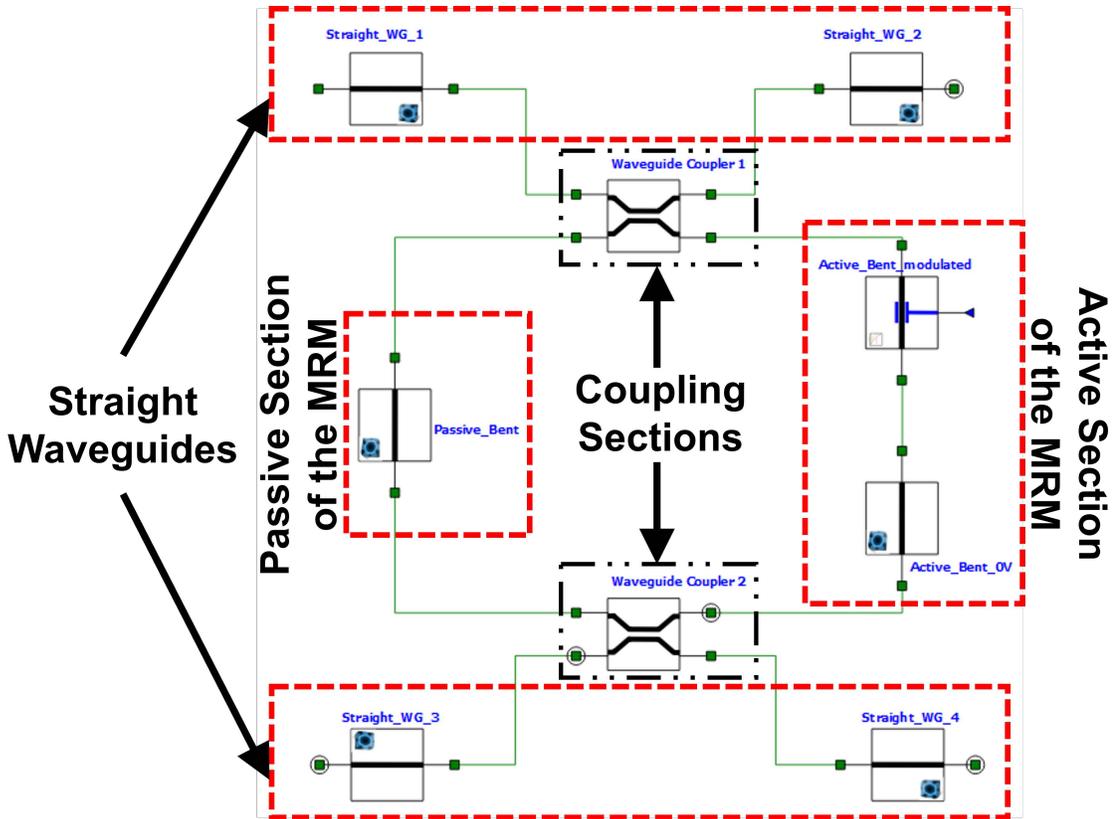


Figure 7.5: Schematic of the MRM of the TAOM, assembled in the INTERCONNECT solver of ANSYS/Lumerical suite [155].

Step-5 - assembling the sub-components: The data extracted from each of the aforementioned steps is imported into the INTERCONNECT solver. Here, primitive elements are created by utilizing the extracted data corresponding to each of the sub-components. Then, the MRM is assembled by connecting each of these primitive elements together as shown in Fig. 7.5.

Simulation

To simulate the MRM of a TAOM, assembled in the INTERCONNECT solver, we leveraged the scripting capabilities offered by the ANSYS/Lumerical suite. For better comprehension of the simulations, let us delve into a scenario involving the simulation of a 3-bit TAOM (a TAOM that operates on 3-bit operands). Different steps involved in the simulation of a 3-bit TAOM unit are discussed below:

Step-1 - Frequency-domain simulations: To perform frequency-domain simulations [2], an optical network analyzer is integrated with the MRM in the INTERCONNECT solver, as shown in Fig. 7.6(a). As previously discussed, the weights (w 's) are encoded as PAM symbols. Therefore, from the frequency domain simulations of a 3-bit TAOM, we extract the transmission spectra for 8 different amplitude levels (which represent a total of 2^3 possible magnitudes of 3-bit weight values) at the through and drop ports of the MRM.

From the extracted transmission spectra, we choose a unique target transmission value corresponding to each amplitude level from the points where the spectra intersect with the carrier wavelength λ_T , as illustrated in Fig. 7.7. This meticulous selection ensures that each amplitude level (weight value) corresponds to a unique target transmission value specific to a transmission spectrum and carrier wavelength intersection. Subsequently, the chosen carrier wavelength identified in the frequency-domain simulations is utilized as an input wavelength for the TAOM during the time-domain simulations.

Step-2 - Time-domain simulations: To perform time-domain (transient) simulations [2], we connect a continuous wave (CW) laser to the input port of the MRM in the INTERCONNECT solver (Fig. 7.6(b)), specifically targeting the carrier wavelength at which the target transmission values were chosen from Step-1 (frequency analysis). Additionally, we integrate a pseudo-random bit sequence (PRBS) generator and a non-return-to-zero (NRZ) pulse generator with the PN junction of the MRM. Utilizing the scripting capabilities of the ANSYS/Lumerical suite, we program this integrated PRBS generator-NRZ pulse generator setup to provide PWAM symbols, which we provide as bias to the PN junction of the MRM at run time.

As discussed earlier, the PWAM symbols consist of inputs/activation values (a) encoded as pulse widths and weights (w) encoded as amplitude levels. In a 3-bit TAOM, there are 8 distinct amplitude levels representing the input/activation values and 8 different pulse widths representing the weights, yielding a total of 64 possible multiplication operations. Each combination of input and weight gives rise to a unique PWAM symbol. For every multiplication operation, which corresponds to a unique PWAM symbol applied as a bias to the PN junction, the MRM performs electro-optic tuning of the transmission of the incoming carrier wavelength.

From the time-domain simulations, we extract the balanced photocurrent at the output of the BPD within the BPCA circuit corresponding to different multiplication operations. Then, we calculate the voltage across the capacitor in the TIR for each multiplication operation. These voltage values are used to evaluate the performance (achievable accuracy and precision) of our TAOM. Detailed discussion related to performance analysis is provided in the upcoming subsection.

Performance Analysis

We evaluate the performance of our TAOM in terms of accuracy and precision. To measure the accuracy of our TAOM, we calculated the logarithmic transformation of the inverse of the normalized mean absolute error (MAE) ($(\log_2(1/(\text{MAE})))$) between the actual voltages (i.e., the voltages across the capacitor (C_1) in the BPCA for different multiplication operations that are extracted from simulations) and the commanded/targeted voltages for different multiplication operations. We represent accuracy in terms of bits by considering various input optical pulse amplitudes and bit resolution values. For this analysis, we considered four values of bit resolution (2,4,6,8-bits) and three unit sizes of pulse widths (16ps,32ps,48ps). If the unit size of pulse width is 16 ps, the input values of unity would be represented as a 16 ps wide pulse. The colormap plot in Fig. 7.8(a) illustrates the accuracy evaluated for

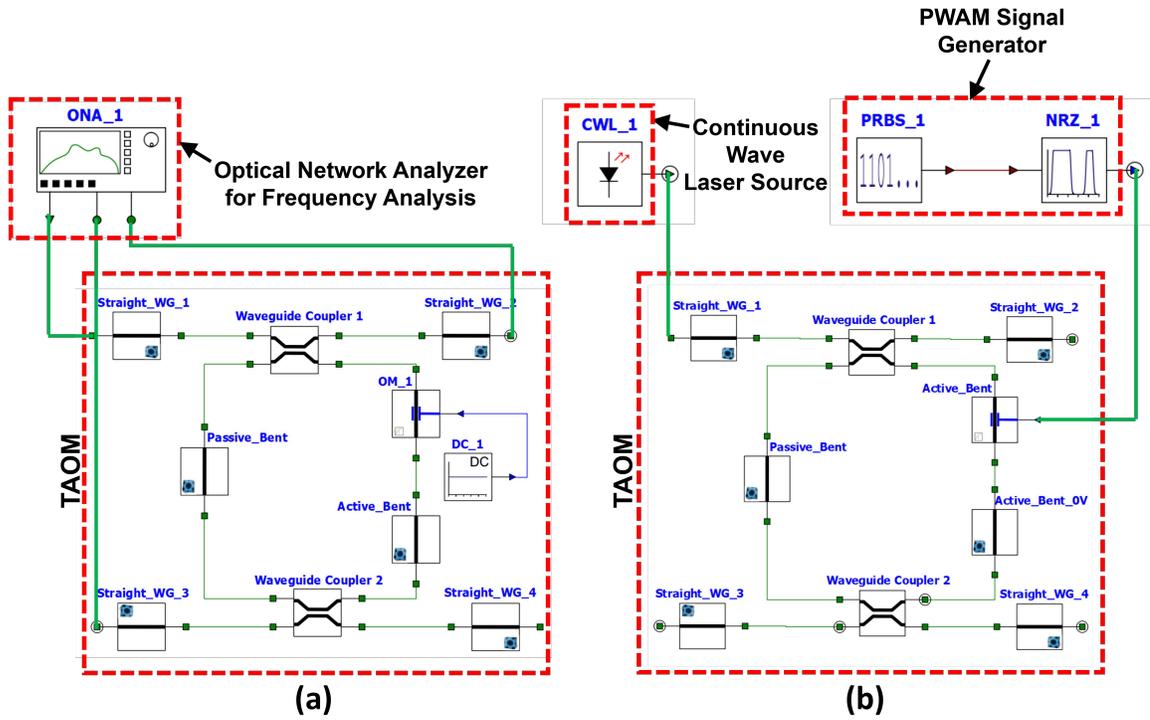


Figure 7.6: (a) Frequency domain and (b) time-domain simulation setup of our TAOM in the INTERCONNECT solver of the ANSYS/Lumerical suite [155].

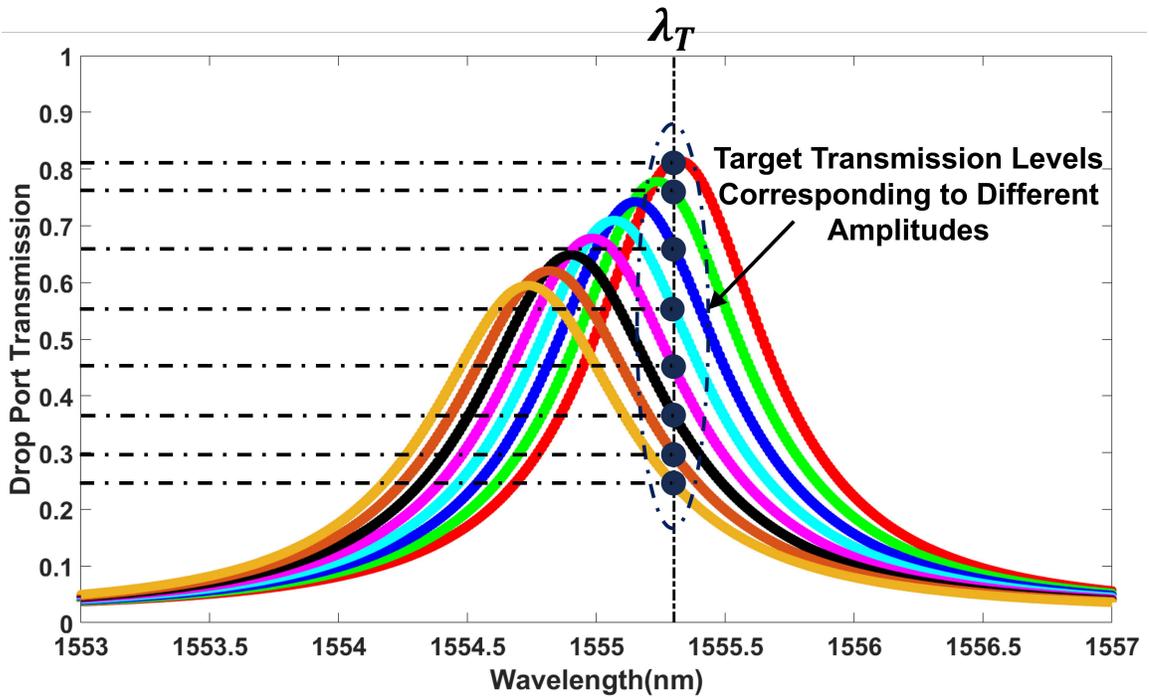


Figure 7.7: Transmission spectra obtained at the drop port of a 3-bit TAOM for different amplitudes.

different combinations of input optical pulse amplitude, bit resolution and unit pulse width. Additionally, we evaluated precision for these combinations using equations provided in [8]. The corresponding colormap plot for precision can be found in Fig. 7.8(b).

As depicted in the accuracy and precision colormap plots (Fig. 7.8), for a given bit resolution and pulse width for a unity input value, the accuracy and precision of our TAOM increase when the input optical pulse amplitude increases from 0 dBm to 10 dBm. This is because the optical pulse amplitude is basically representative of the optical power signal, and the increase in the input optical pulse amplitude means an increase in the optical power signal. This improves the signal-to-noise ratio, leading to better accuracy and precision for our TAOM. Similarly, for a given input optical pulse amplitude, the precision of our TAOM increases when the pulse width of a unity input value increases. Furthermore, for a given input optical pulse amplitude and a unit pulse width, the accuracy and precision of our TAOM increase with the increase in bit resolution. These results imply that it is possible to achieve accuracy of as high as 16-bit and precision of as high as 10-bit with our TAOM. These accuracy and precision values for our TAOM are highly competitive compared to the values achievable by analog-only photonic incoherent multipliers (or weight banks) from prior works [82][292].

7.4 Analysis of TAOM-Enabled Parallel Multiplier Circuit

In the circuit-level analysis of TAOM, we look at the impact of intermodulation (IM) crosstalk in a WDM-based TAOM-enabled parallel multiplier circuit. This impact manifests as the errors incurred in the multiplication results obtained by the individual TAOMs of the circuit. First, we examine the mechanism of IM crosstalk in WDM-based MRMs. Then, we focus on the modeling and simulation of WDM-based TAOM-enabled circuits using the INTERCONNECT solver in the ANSYS/Lumerical suite. Through these modeling and simulations, we aim to characterize the multiplication error in the individual TAOMs of the parallel multiplier circuit. Lastly, we discuss the outcomes obtained from these circuit-level simulations. The details are discussed in the following subsections.

Inter-Modulation Crosstalk in MRMs

Fig. 7.9(a) depicts the configuration of cascaded MRMs that enables wavelength division multiplexing. Each MRM in this configuration is tuned to modulate its designated wavelength channel independently of the other MRMs. As a result, increasing the WDM density (defined as the number of wavelength channels multiplexed in a waveguide) becomes a natural choice to enhance the bandwidth density of the photonic circuits. By cascading more MRMs and utilizing their individual modulation capabilities, we can accommodate a greater number of wavelength channels, leading to a higher data transmission capacity and improved performance for the photonic system.

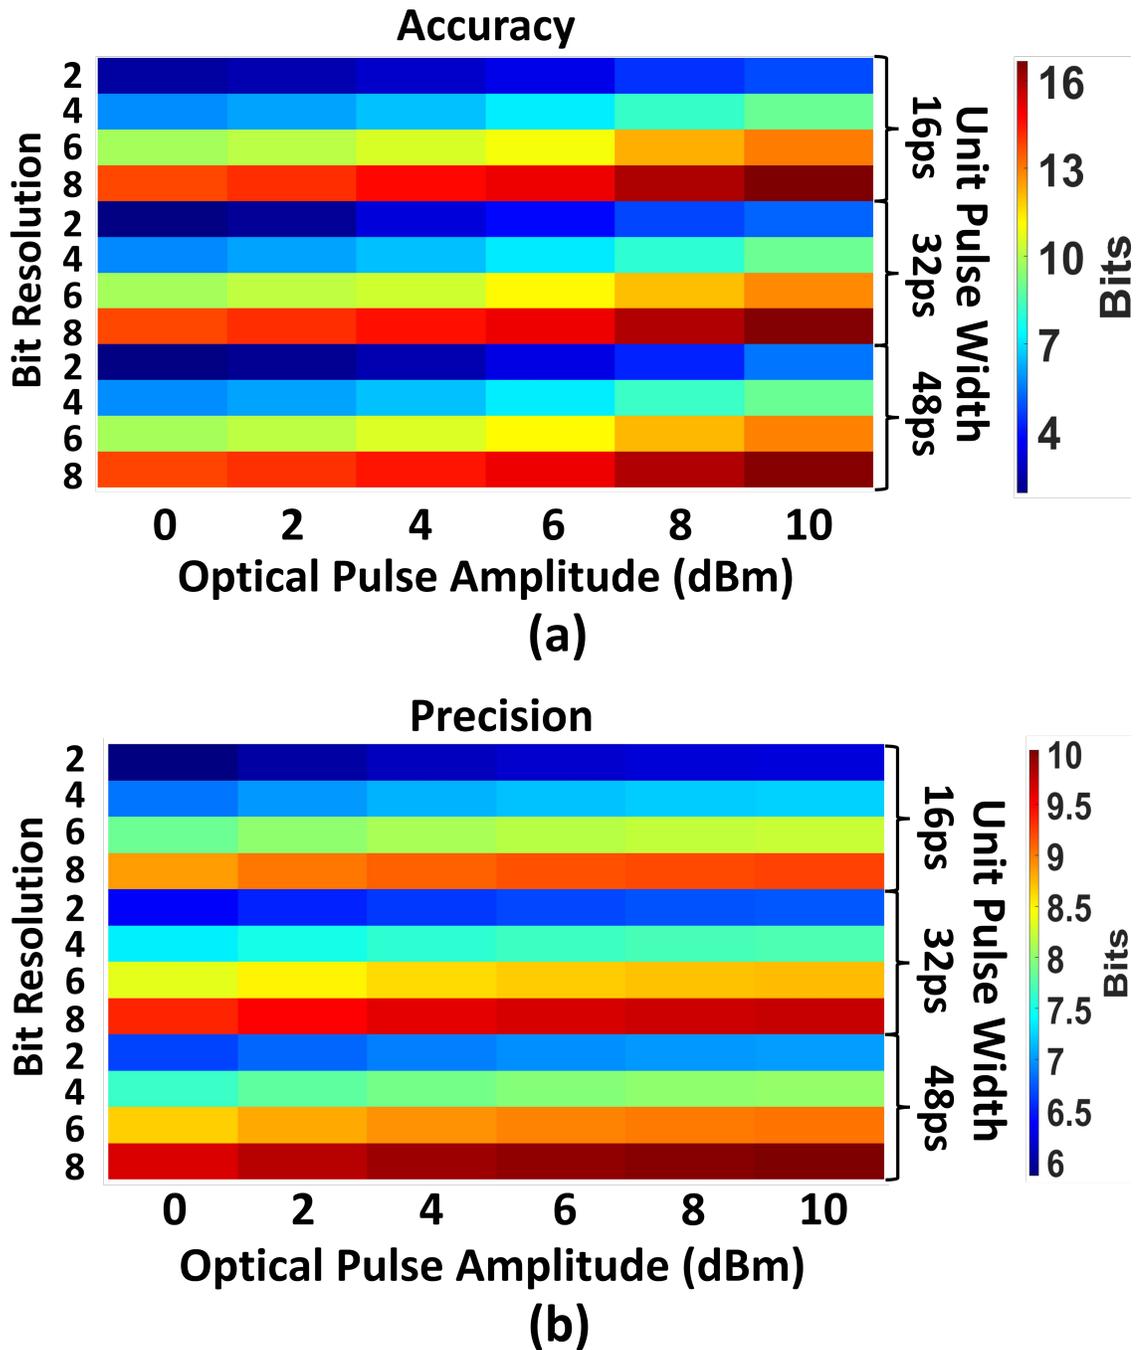


Figure 7.8: Colormap plots that depict the (a) accuracy and (b) precision of our TAOM for different values of input optical pulse amplitude, bit resolution, and pulse widths for a unity input value.

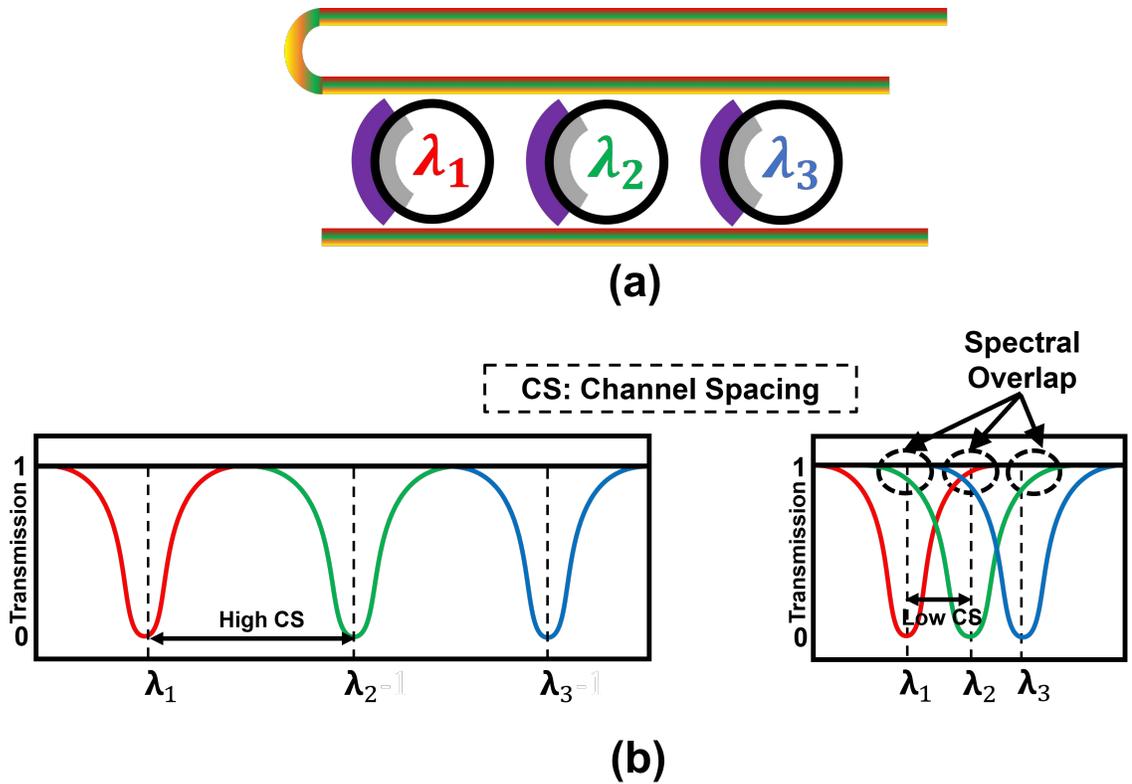


Figure 7.9: (a) WDM-based MRMs and (b) aggregate transmission spectra of the WDM-based MRMs when using a coarse channel spacing (left) and a dense channel spacing (right).

However, it is crucial to ensure that channel spacing between the individual resonances of the MRMs is sufficiently coarse. As shown in the transmission spectra on the left side of Fig. 7.9(b), the resonance wavelengths of the MRMs are deliberately kept apart to avoid any overlap. This spacing ensures that the resonances of individual MRMs do not overlap with one another. Nevertheless, as the WDM density increases and more MRMs are cascaded, the channel spacing between the resonances of the individual MRMs decreases. As a result, the resonances will begin to overlap, as illustrated in Fig. 7.9 (b) (on the right). Consequently, each MRM not only modulates its corresponding wavelength channel but also modulates the neighboring wavelength channel. This phenomenon introduces IM crosstalk [184][125] in cascaded MRMs enabling WDM, which can affect the performance and reliability of WDM-based photonic circuits and systems.

Since TAOMs employ MRMs, a cascade of TAOMs that enables WDM to realize a parallel multiplier is also naturally susceptible to IM crosstalk. To investigate the impacts of IM crosstalk in cascaded TAOMs-based parallel multiplier, we model a 2-channel cascaded TAOM circuit in the INTERCONNECT solver of ANSYS/Lumerical suite. Detailed discussion regarding the modeling and simulation is provided in the upcoming subsections.

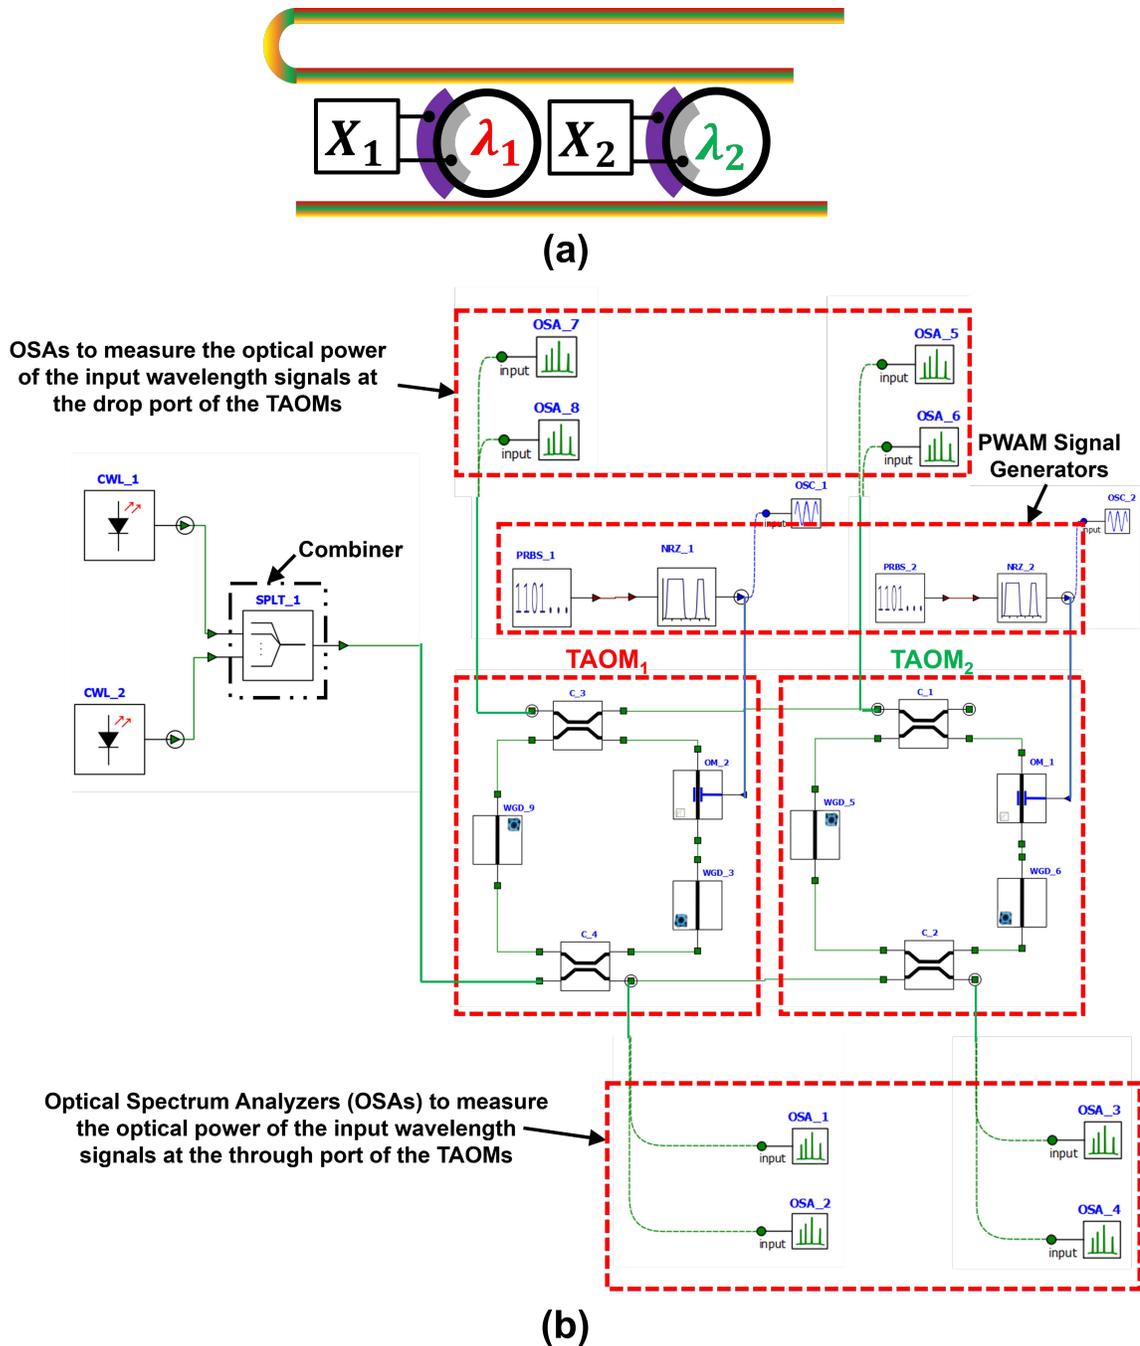


Figure 7.10: (a) Cascaded TAOMs that enable WDM, and (b) its time-domain (transient) simulation setup in the INTERCONNECT solver of ANSYS/Lumerical suite [155].

Modeling and Simulation of TAOM-Enabled Parallel Multiplier

Modeling

We modeled a 2-channel cascaded TAOM circuit in the INTERCONNECT solver of ANSYS/Lumerical suite. Following the steps described in the previous section,

we modeled two individual hybrid TAOMs of different radii and cascaded them by parallel coupling them with two bus waveguides (one at the top and one at the bottom), as shown in Fig. 7.10(a). Employing two TAOMs enables this circuit to realize two multiplication operations in parallel. This idea can be extended to a total of N cascaded TAOMs, to realize a circuit for implementing N parallel multipliers. For simulation and analysis, we performed time-domain (transient) simulations on the two-TAOM circuit from Fig. 7.10(a) for various channel spacings, and observed the change in the IM crosstalk and its impact on the multiplication result of each of the TAOMs. Comprehensive analysis and discussion regarding the simulation results are presented in the upcoming subsections.

Simulation

To simulate the cascaded TAOM circuit, we leveraged the scripting capabilities of the ANSYS/Lumerical suite. First, we connected two CWL sources, along with a combiner, to the input of our WDM cascaded TAOM circuit, as shown in Fig. 7.10(b). This CWL setup enables the provision of two distinct wavelength signals, which are then multiplexed by the combiner to the input port of our TAOM circuit. Following the steps provided for the time-domain simulations in the previous section, we integrated each of the PN junctions of the TAOMs with a PRBS generator and an NRZ pulse generator. This integration allows us to provide PWAM signals to the PN junctions of the TAOMs during the run time.

To analyze IM crosstalk phenomena in each of the TAOMs, we connected optical spectrum analyzers (OSAs) at the through and drop ports of each TAOM, as shown in Fig. 7.10 (b). As discussed earlier, when the spacing between the resonances of MRMs decreases, overlapping of resonances occurs. Consequently, each MRM not only modulates its corresponding wavelength channel but also modulates a fraction of the power in the neighboring wavelength channel. By integrating OSAs at the through and drop ports of each TAOM, we effectively monitored the fraction of power (IM crosstalk noise power) coupling into each TAOM from neighboring wavelength channels across various channel spacings. Following the methodology outlined in the previous section, we conducted time-domain simulations for various channel spacings (1.6nm, 1.2nm and 0.7nm). These simulations enabled us to calculate the capacitor voltages (in the TIR of the BPCA) associated with the multiplication results of each TAOM, for different channel spacings. Subsequently, we normalized these observed capacitor voltages with respect to the target/ideal voltages corresponding to the desired multiplication results. The observed capacitor voltages represented observed multiplication results, which deviated from the target/ideal multiplication results. The errors in the observed multiplication results with respect to the ideal results were visually illustrated by plotting the normalized multiplication results from both TAOMs on a grid plot, as depicted in Fig. 7.11. A comprehensive analysis of the outcomes of this study is presented in the next subsection.

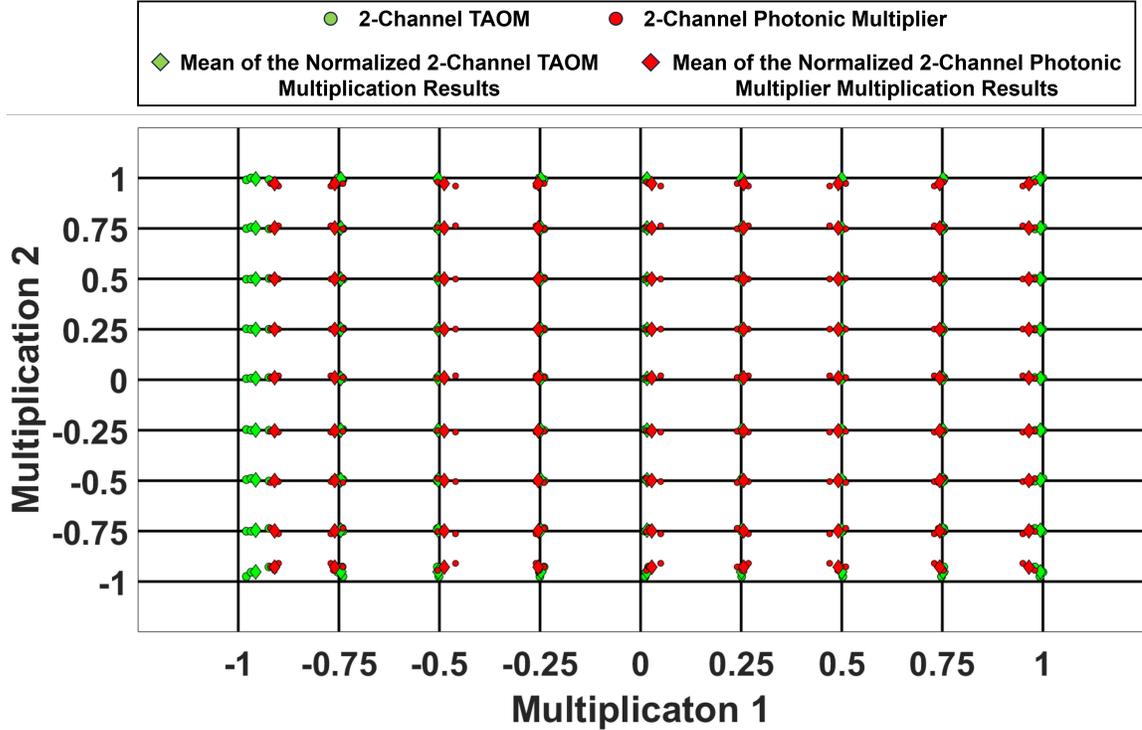


Figure 7.11: IM crosstalk induced deviation in the multiplication results of each TAOM in a 2-channel parallel multiplier circuit, depicted for various channel spacings. (1.6nm, 1.2nm and 0.7nm)

Results and Discussion

Figure 7.11 demonstrates the impact of IM crosstalk on the multiplication results of each TAOM within the 2-channel cascaded TAOM circuit, shown in Figure 7.10(a), for three different channel spacings. The grid plot, represented as a 9×9 grid, allows us to visualize the extent of deviation in the multiplication results due to IM crosstalk between the two TAOMs. The intersection points of black gridlines (horizontal and vertical gridlines) indicate the target multiplication results, where *Multiplication 1* on the horizontal axis represents the multiplication results corresponding to $TAOM_1$, and *Multiplication 2* on the vertical axis represents the multiplication results corresponding to $TAOM_2$. The green scattered circles around each target result (gridline intersection point) illustrate the observed multiplication results obtained for different channel spacings between the resonance wavelengths of the two TAOMs. In addition, the green diamond points represent the mean of the observed multiplication results obtained for different channel spacings. For comparative analysis, we evaluated multiplication results observed for a conventional 2-channel photonic multiplier, that employs 2 input MRMs and 2 weighting MRRs (for example, the AMW-styled parallel multiplier; Fig. 7.1(a)). These results are shown as red scattered circles around each target result (gridline intersection point), with the red diamond points representing the corresponding mean values.

In Fig. 7.11, comparing the observed results for a target grid line intersection

point, the green point farthest from the target multiplication result corresponds to the TAOM’s multiplication outcome with the lowest channel spacing, while the closest green point represents the TAOM’s outcome with the highest channel spacing. Consequently, as evident from the figure, as the channel spacing between the resonance wavelengths of the TAOMs decreases, the IM crosstalk noise within each TAOM increases. This escalation in IM crosstalk noise results in a deviation of the actual multiplication result from the target value, as illustrated by the observed deviations of the green scattered circles (and the green diamond mean points) from the target values (grid line intersection points) in Fig. 7.11. This IM crosstalk-induced deviation in multiplication results adversely impacts the accuracy of each TAOM. In contrast, when we compare these results with those obtained from a traditional 2-channel photonic multiplier that employs two input MRMs and two weighting MRRs, an interesting observation arises. In that setup, the use of additional MRMs and MRRs amplifies the IM crosstalk and optical losses experienced by each wavelength channel. This leads to a more substantial deviation in the multiplication results, evident from the deviation of the red scattered circles (and the red diamond mean point) from the target multiplication results, as illustrated in Fig. 7.11.

In summary, our analysis underscores that decreasing the channel spacing in TAOMs intensifies IM-crosstalk, affecting the accuracy of multiplication results. However, compared to the conventional photonic multiplier circuit, our cascaded TAOM circuit experiences reduced IM-crosstalk and, consequently, exhibits a lesser deviation/error in multiplication results. This advantage renders our TAOM-based tensor core with better output accuracy with GEMM workloads and neural network applications, as discussed later in this chapter.

7.5 A Hybrid Time-Amplitude Analog Optical Accelerator

In this section, we will discuss our invented TAOM-based tensor processing core (TAOM-TC). We will also present our analysis to evaluate the scalability, power, performance, and inference accuracy of TAOM-TC.

Design of TAOM-TC

Utilizing our invented TAOM and BPCA units, we devised a SiPh GEMM accelerator architecture named TAOM-based Tensor Core (TAOM-TC). In this architecture, we have organized our TAOMs and BPCAs in 2D arrays, as illustrated in Fig. 7.12. Our TAOM-TC design comprises a total of ‘M’ waveguides, each consisting of ‘N’ distinct TAOMs cascaded to the waveguide in a DWDM configuration. To supply input wavelength signals to each TAOM, we incorporated ‘N’ laser diodes that generate ‘N’ unique input wavelength signals, as depicted in Fig. 7.12. These ‘N’ distinct wavelength signals are multiplexed using an $N \times 1$ multiplexer. The multiplexed multi-wavelength signal is then split into ‘M’ waveguides via a $1 \times M$ splitter. At the termination of each waveguide, a BPCA circuit is deployed that performs a signed summation of a large number of temporally and spatially arriving PWAM symbols, thereby executing a dot product operation for each waveguide (as explained earlier).

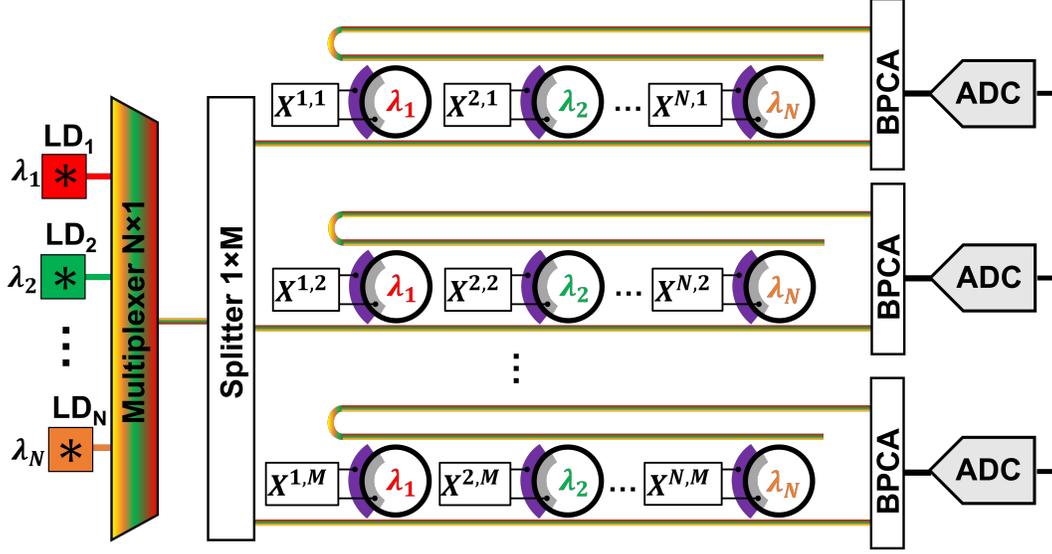


Figure 7.12: Architecture of the hybrid Time-Amplitude Analog Optical Modulator based Tensor Core (TAOM-TC).

$$B = \frac{1}{6.02} \left[20 \log_{10} \left(\frac{RP_{PD-opt}}{\left(\sqrt{2q(RP_{PD-opt} + I_d) + \frac{4KT}{R_L}} + (RP_{PD-opt})^2 RIN + \sqrt{2qI_d + \frac{4KT}{R_L}} \right) \sqrt{\frac{DR}{\sqrt{2}}}} \right) - 1.76 \right] \quad (1)$$

$$P_{output} \text{ (dBm)} = P_{laser} - P_{SMF} - P_{EC-IL} - P_{Si-att} - P_{MRM-IL} - (N-1)P_{MRM-OBL} - P_{splitter} - P_{MRR-IL} - (N-1)P_{MRR-OBL} - P_{penalty} \quad (2)$$

Furthermore, this strategic 2D arrangement of our TAOM and BPCA units within the TAOM-TC architecture allows for superior spatial parallelism. This design choice substantially enhances the processing throughput of our architecture at better power efficiency in contrast to the existing MRR-enabled SiPh GEMM accelerators.

We conducted a comprehensive evaluation of the achievable spatial parallelism (scalability), power, performance, and inference accuracy of our TAOM-TC architecture. Detailed discussion pertaining to the outcomes of each of these evaluations is presented in the following subsections.

$$ef(B, DR, N) = P_{output}(N) - P_{PD-opt}(B, DR) \quad (3)$$

Table 7.1: Definitions and values of various parameters used in Eqs. 1,2 and 3 (from [8, 263, 262]) to perform the scalability analysis.

Parameter	Definition	Value
P_{laser}	Laser Power Intensity	10 dBm
P_{SMF}	Attenuation by the Single-Mode Fiber	0 dB
P_{EC-IL}	Fiber-to-Chip Coupling Loss	1.6 dB
P_{Si-att}	Silicon Waveguide Attenuation	0.3 dB/mm
P_{MRM-IL}	Transmission Insertion Loss of the MRM (Input Vector)	4 dB
$P_{MRM-OBL}$	Out-of-Band Insertion Loss (OBL) of the MRM (Input Vector)	0.01 dB
$P_{splitter}$	Splitter Insertion Loss	0.01 dB
P_{MRR-IL}	Transmission Insertion Loss of the MRR (Weight Vector)	0.01 dB
$P_{MRR-OBL}$	Out-of-Band Insertion Loss (OBL) of the MRR (Weight Vector)	0.01 dB
$P_{penalty}$	Network Penalty for MAW	4.8 dB
	Network Penalty for AMW	5.8 dB
	Network Penalty for TAOM-TC	1.8 dB
R	PD Responsivity	1.2
q	Charge of an Electron	1.6×10^{-19} C
I_d	PD Dark Current	35 nA
K	Boltzmann Constant	1.38×10^{-23} J/K
T	Absolute Temperature	300 K
R_L	Load Resistance	50 Ohms
RIN	Relative Intensity Noise	-140 dB/Hz
B	Bit-Precision	-
P_{PD-opt}	PD Sensitivity	-

Scalability Analysis for TAOM-TC

To perform scalability analysis, we utilized the equations provided in [8], represented by Eqs. 1, 2 and 3. The parameters and their corresponding values [8, 263, 262] required to solve these equations are listed in Table 7.1. We devised a two-step procedure to determine the optimal value of N (N refers to the number of TAOMs per waveguide, and number of wavelength channels multiplexed per waveguide) for a given bit precision and data rate, as outlined below:

Step 1 - calculate PD sensitivity: Firstly, we calculate the photodiode (PD) sensitivity by solving Eq. 1 for the specified bit precision and data rate.

Step 2 - exhaustive search for optimal 'N': Next, we perform an exhaustive search to find the optimal value of N for the specified bit precision and data rate, using Eqs. 2 and 3. In this step, we consider $M=N$ and solve Eq. 3, which represents the error function (ef). The ef is the difference between the optical power reaching the photodiode (P_{output}), calculated from Eq. 2, and the PD sensitivity obtained in Step 1/Eq. 1. We iterate through different values of N , and the optimal value of N for the specified bit precision and data rate is the one for which the ef yields the minimum positive value.

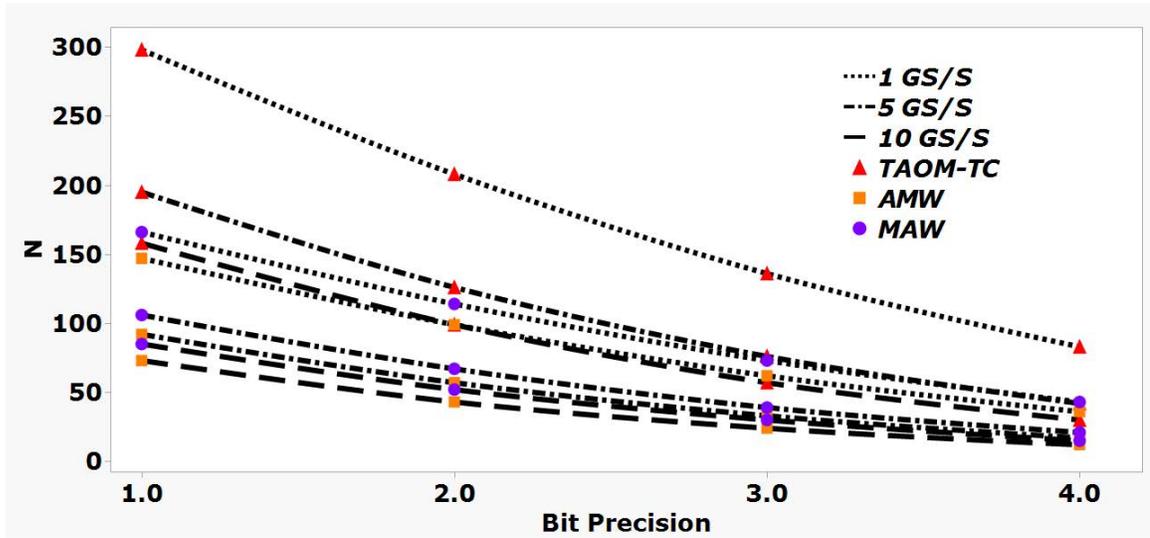


Figure 7.13: Supported TPC size $N(=M)$ for bit-precision = 1,2,3,4-bits at data rates (DRs)=1,5,10 GS/s, for TAOM-TC, MAW TPC and AMW TPC.

For our analysis, we considered bit-precision values ranging from 1-bit to 4-bits and a set of DRs namely 1GS/S, 5GS/S, and 10GS/S. The results of our analysis are illustrated in Fig. 7.13. In addition to our TAOM-TC, we conducted the scalability analysis for AMW and MAW TPCs. As we can infer from Fig. 7.13, our TAOM-TC can support a larger value of N compared to AMW and MAW TPCs. For instance, our TAOM-TC can support $N=83$ for a bit-precision of 4-bits, DR = 1GS/S and an input laser power of 10dBm, which is larger compared to AMW and MAW TPCs, which support $N=36$ and $N=43$ respectively. This advantage arises mainly due to the absence of weighting MRRs in our TAOM-TC, leading to reduced insertion losses

Table 7.2: Hardware description of various optical and electrical components utilized in AMW and MAW TPCs, and TAOM-TC.

Hardware Component	Hardware Count		
	AMW	MAW	TAOM-TC
Lasers	N	N	N
DACs	$2*M*N$	$(N+(M*N))$	$M*N$
DPCs	0	0	$M*N$
TIAAs	M	M	M
ADCs	M	M	M
MRMs	$M*N$	N	$M*N$
MRRs	$M*N$	$M*N$	0
Feedback Control Units	$4*(M*N)$	$(2*N)+(2*M*N)$	$(2*M*N)$

Table 7.3: Power consumption of various hardware components utilized in AMW and MAW TPCs, and TAOM-TC.

Hardware Component	Power
Optical Output Power per Laser	10 mW [8]
Electrical Input Power per Laser (10% Wall-Plug Efficiency)	100 mW [8]
Digital-to-Analog Converters (DACs)	26 mW [118]
Digital-to-Pulse Converters (DPCs)	2.5 mW [212]
Transimpedance Amplifiers (TIAAs)	25.1 mW [133]
Analog-to-Digital Converters (ADCs)	0.02 mW [178]
Electro-Optic Microring Modulators (MRMs) (Input Vector)	~ 0.9 mW [272]
Thermo-Optic Microring Resonators (MRRs) (Weight Vector)	~ 180 mW [88]
Power per Feedback Control Unit	~ 5 mW [292, 183]

and crosstalk-induced power penalty ($P_{penalty}$). Consequently, this creates room in the optical power budget, allowing for the accommodation of a larger N (i.e., a larger number of TAOMs/multipliers per waveguide).

Power Estimation

To evaluate the power consumption of our TAOM-TC, we selected an 'N' value of 42, corresponding to a 4-bit precision and a data rate of 5 GS/S, as determined from the scalability analysis shown in Fig. 7.13. For comparative analysis, we also evaluated the power consumption for MAW TPC with N=21 and AMW TPC with N=17,

for the same bit-precision and data rate. We then compared these evaluations with our TAOM-TC, as illustrated in Fig. 7.14. These power values include the power consumption of all components of the considered TPC architectures. For detailed hardware specifications of each TPC and the corresponding power values of their components, please refer to Tables 7.2 and 7.3 respectively.

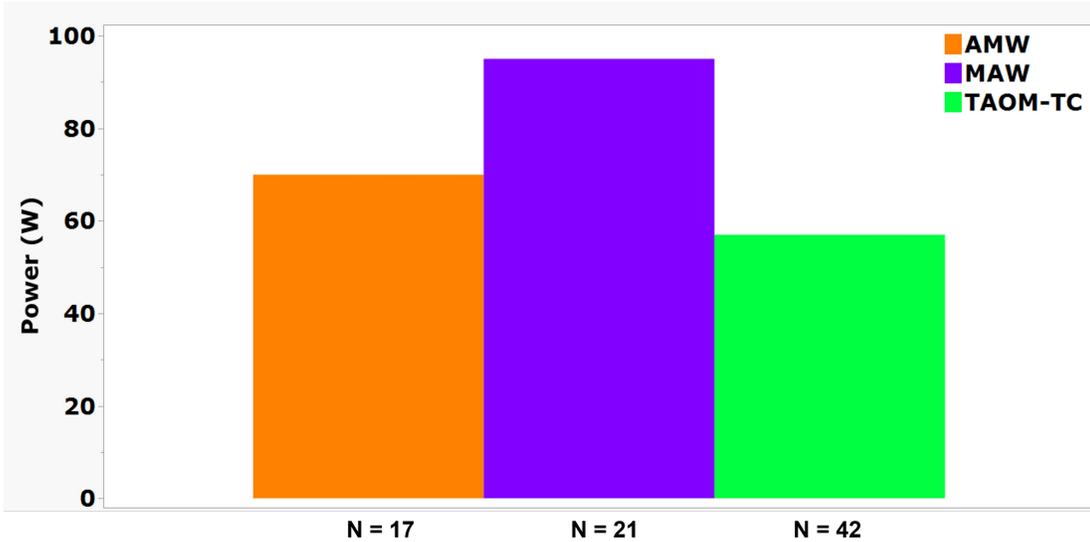


Figure 7.14: Power consumption of the AMW, MAW and TAOM-TC architectures for a bit precision of 4-bits and a data rate of 5 GS/S.

From Figure 7.14, it is evident that our TAOM-TC consumes $\approx 1.3\times$ less power than the AMW TPC and $\approx 1.7\times$ less power than the MAW TPC. This significant decrease in power consumption can be primarily attributed to the absence of weighting MRRs and the reduced number of required feedback control units in our TAOM-TC (as detailed in Table 7.2). These components contribute significantly to power consumption in both AMW and MAW TPCs.

Table 7.4: Convolutional parameters corresponding to various benchmarks utilized in the performance analysis of our TAOM-TC.

Benchmarks	W	H	D	NI	K	R_w	R_h	S
Benchmark-1	7	7	832	16	256	1	1	1
Benchmark-2	28	28	192	1	32	5	5	1
Benchmark-3	84	20	256	1	512	5	5	2
Benchmark-4	42	10	512	1	512	3	3	1

Performance Analysis

To assess the performance of our TAOM-TC architecture, we evaluated the run time required to perform a convolution on four different benchmarks from DeepBench [22]. Each benchmark encompasses a distinct set of convolutional parameters, which are

Table 7.5: Definition of various convolutional parameters corresponding to different benchmarks utilized in the performance analysis of our TAOM-TC.

Convolutional Parameter	Definition
W	Width of the Input Image Tensor
H	Height of the Input Image Tensor
D	Depth
NI	Number of Input Images
K	Tensor Count
R_W	Width of the Filter
R_H	Height of the Filter
S	Stride

detailed in Tables 7.4 and 7.5. For this evaluation, we considered a DR of 5 GS/S and bit-precision of 4-bits. The supported size $N'(=M)$ corresponding to these parameters for TAOM-TC is 42, which we determined from the scalability analysis presented in Fig. 7.13. Taking into account the convolutional parameters corresponding to each of the considered benchmarks and the hardware size of the TPC, we employed the Toeplitz matrix decomposition technique [246] to convert the convolution operations into GEMM operations. We leveraged this technique to efficiently map the convolutions onto the available hardware for acceleration. To estimate the run time of a single convolution, we considered the propagation time of light signals, i.e., the time taken for light to propagate from the multiplexer to just before the BPDs, similar to what is done in prior work [26]. Additionally, we accounted for the response time of peripheral circuitry, including DACs [118], ADCs [84], and TIAs [133]. Furthermore, we have also evaluated the run time of a convolution for MAW ($N=21$) and AMW ($N=17$) TPCs, and compared it with our TAOM-TC. Additionally, we explored the convolutional run time of various GPUs from [22] and compared their performance with that of our TAOM-TC. The run times corresponding to each of these architectures for different benchmarks are presented in Fig. 7.15.

As we can infer from Fig. 7.15, in Benchmark 1, our TAOM-TC can execute convolutions $\approx 7.3\times$ faster than GPUs and about $\approx 3\times$ faster than MAW and AMW TPCs, on average. Similarly, for Benchmark 2, our TAOM-TC can perform convolutions approximately $\approx 25\times$ faster than GPUs and roughly $\approx 2.9\times$ faster than MAW and AMW TPCs, on average. Furthermore, in both Benchmark 3 and Benchmark 4, our TAOM-TC outperforms GPUs by around $\approx 4\times$ times and MAW/AMW TPCs by approximately $\approx 3\times$ times, on average, in terms of convolution run time.

Inference Accuracy Analysis

In Section IV.C, we discussed the presence of errors in multiplication results in TAOM-TC due to IM-crosstalk, a phenomenon that also introduces errors in the multiplication outcomes of AMW TPC. To assess the impact of these inaccuracies on CNN inference, we conducted evaluations using an 8-bit integer quantized LeNet model [135]. The evaluations were performed on both TAOM-TC and AMW TPC

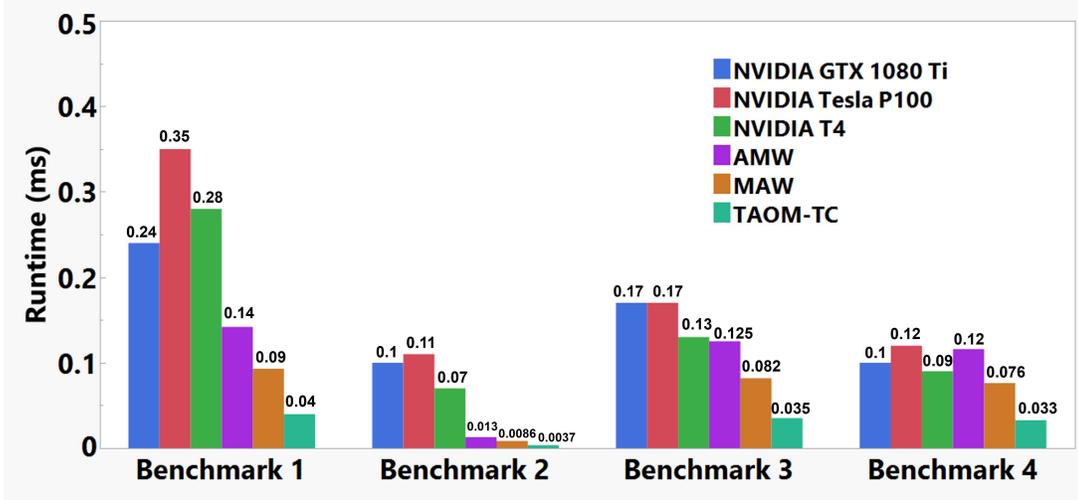


Figure 7.15: Runtimes of state-of-the-art GPUs, AMW, MAW and TAOM-TC accelerator architectures estimated for four different benchmarks.

using the PyTorch ML-framework [194], with inference executed on the MNIST validation dataset [70] comprising 10,000 images across 10 classes.

For the validation dataset, TAOM-TC achieved an overall accuracy of 98.76%, while AMW TPC attained an accuracy of 98.36%. Notably, TAOM-TC outperformed AMW TPC with a 0.4% higher accuracy. This is due to lower IM-crosstalk induced error in TAOM-TC as compared to AMW TPC. Therefore, it is evident that our TAOM-TC offers enhanced performance and accuracy compared to prior optical TPCs.

7.6 Summary

In this chapter, we presented a novel hybrid Time-Amplitude Analog Optical modulator (TAOM) that leverages only a single microring modulator (MRM) to generate a high-speed optical signal as a sequence of Pulse-Width-Amplitude-Modulated (PWAM) symbols, in which each symbol represents the analog product of an input and a weight value. This TAOM is integrated with a novel balanced photo charge accumulator (BPCA) that leverages the in-situ charge accumulation and incoherent superposition abilities of photodetectors to generate a signed summation of a large number of temporally and spatially arriving PWAM symbols. By employing a single MRM for dot product operations, we significantly reduced insertion losses and power consumption while enhancing the spatial parallelism of MRR-based photonic GEMM accelerators. Furthermore, we organized our invented TAOMs and BPCAs in 2D arrays to design a SiPh GEMM accelerator called TAOM-Tensor Core (TAOM-TC). We performed an extensive device-level, circuit-level, and system-level analysis of our invented TAOM-TC.

At the device level, we performed extensive modeling and characterization of our invented TAOM using photonics foundry-validated tools from the ANSYS/Lumerical suite. Through these device-level simulations, we evaluated the performance of our

TAOM unit in terms of accuracy and precision for varying values of input optical pulse amplitude, bit resolution, and pulse widths. From the device-level evaluation, we observed that the accuracy of our TAOM improves with higher input optical power and increased bit resolution. In addition, the precision of our TAOM increases with the increase in input optical pulse amplitude, bit resolution, and pulse width. At the circuit level, we performed transient circuit simulations on a TAOM-enabled 2-channel parallel multiplier circuit. Through these simulations, we observed the changes in inter-modulation (IM) crosstalk and its impact on the multiplication result of each TAOM, for various channel spacings. The circuit-level analysis revealed that an increase in IM-crosstalk results in a greater deviation from the target/intended multiplication results for each TAOM. However, in contrast with the conventional photonic multiplier circuit, our TAOM-enabled parallel multiplier circuit experiences reduced IM-crosstalk and consequently, less error in the multiplication results from the intended results. At the system level, we evaluated the scalability, power, performance, and inference accuracy of our TAOM-TC, and compared it with two well-known MRR-based GEMM accelerators from prior works. From the scalability and power analysis, we observed that our TAOM-TC has the capability to support at least $1.5\times$ more TAOMs/multipliers per waveguide while consuming $\approx 1.5\times$ less power compared to the considered MRR-enabled SiPh DNN accelerators from prior works. Furthermore, from the performance analysis, we inferred that our TAOM-TC performs convolutions $\approx 10\times$ faster than the GPUs and up to $\approx 3\times$ faster than the considered MRR-based GEMM accelerators from prior works. In addition, our TAOM-TC outperforms the prior MRR-based GEMM accelerators in terms of accuracy by 0.4%.

Thus, our invented TAOMs enable the use of PWAM signals for multiplications to consequently reduce the number of MRMs required. The reduced MRMs bring significant power, energy, and optical loss benefits. These benefits, when combined with the capabilities of our invented BPCA to perform a large number of spatio-temporal accumulations, provide our TAOM-TC GEMM accelerator substantial advantage in terms of achievable hardware scalability, processing parallelism, and throughput. All of these advantages collectively position our TAOM-TC architecture as a preferred solution for accelerating GEMM operations.

Chapter 8 Indium Tin Oxide Based Silicon Nitride Microring Modulators for High Performance Photonic Integrated Circuits

8.1 Introduction

Driven by the rise of CMOS-compatible processes for fabricating photonic devices, photonic integrated circuits (PICs) are inexorably moving from the domain of long-distance communications to chip-to-chip and even on-chip applications. It is common for the PICs to incorporate optical modulators, to enable efficient manipulation of optical signals, which is a necessity for the operation of active PICs. Recent advances in the CMOS-compatible silicon-on-insulator (SOI) photonic platform has fundamentally improved the applicability of SOI PICs [59], [40], [96]. But in the last few years, the silicon-nitride-on-silicon-dioxide (SiN-on-SiO₂) platform has gained tremendous attention for realizing PICs. This is because the SiN-on-SiO₂ platform has several advantages over the SOI platform. Compared to silicon (Si), the SiN material has a much broader wavelength transparency range (500nm-3700nm), lower refractive index and smaller thermo-optic coefficient [273]. The lower refractive index of SiN means that SiN offers smaller index contrast with SiO₂ compared to Si. This in turn makes the SiN-on-SiO₂ based monomode passive devices (e.g., waveguides, microring resonators (MRRs)) less susceptible to (i) propagation losses due to the decreased sensitivity to edge roughness [31], and (ii) aberrations in the realized device dimensions caused due to fabrication-process variations [273]. In addition, the smaller thermo-optic coefficient of SiN makes it possible to design nearly athermal photonic devices using SiN [87]. Moreover, SiN devices and circuits exhibit increased efficiency of nonlinear parametric processes compared to Si [139].

Despite these favorable properties of the SiN-on-SiO₂ platform, SiN-on-SiO₂ based active devices such as modulators (e.g., [197, 9, 115, 5, 101]) are scarce and lack in modulation bandwidth, modulation efficiency and free spectral range (FSR) [9]. This is because of the lack of the free-carriers based activity in the SiN material and the general difficulty of incorporating other active materials with the SiN-on-SiO₂ platform. This in turn limits the use of the SiN-on-SiO₂ platform for realizing only passive PICs. To overcome this shortcoming, there is impetus to heterogeneously integrate active photonic materials and devices with SiN-on-SiO₂ passive devices [89]. When such efforts of integrating electro-optically active materials with the SiN-on-SiO₂ platform come to fruition, it will be possible to design extremely high-performance and energy-efficient SiN-on-SiO₂ based active and passive PICs.

Different from such prior efforts, in this chapter, we demonstrate *for the first time* the use of the high-amplitude electro-refractive activity of Indium-Tin-Oxide (ITO) thin films to realize two SiN-on-SiO₂ based optical on-off-keying (OOK) modulators. Through electrostatic, transient, and finite difference time domain (FDTD) simulations using photonics foundry-validated tools from Lumerical/Ansys, we demonstrate that both of our modulators achieve exceptional performance metrics: 280 pm/V and 450 pm/V resonance modulation efficiency, 67.8 GHz and 53 GHz 3-dB modula-

tion bandwidth, approximately 19 nm and 18 nm free-spectral range (FSR), around 0.23 dB insertion loss, and 10.31 dB and 8.2 dB extinction ratio for optical OOK modulation at 30 Gb/s, respectively. *Based on the obtained simulation results, we advocate that our modulators can achieve better performance compared to the existing SiN modulators and several state-of-the-art Si and Lithium Niobate (LN) modulators from prior work.*

8.2 Related Work and Motivation

A plethora of Si, Lithium Niobate (LN) and SiN based integrated optical modulator designs have been formulated in prior work. But among these modulator designs, MRR based modulators have gained widespread attention due to their high wavelength selectivity, compact size, and compatibility for cascaded dense wavelength division multiplexing (DWDM). Here, we briefly review some relevant Si, LN and SiN MRR modulators from prior work.

Silicon (Si) Based Modulators

Over the last two decades, Si has emerged as the prominent material for fabricating PICs, mainly because the cost effectiveness of reusing the established CMOS manufacturing infrastructure promotes the use of Si for building complex PICs. Si material also exhibits high thermo-optic and electro-optic (free-carriers-induced) sensitivity, which enables the realization of optical modulators directly in Si substrate without requiring any auxiliary active materials. Si optical modulators based on free-carriers-induced plasma dispersion and absorption effects have become particularly more popular because of their low-power and high-speed operation. Although several designs of Si-based modulators have been reported, the most commonly adopted designs employ microring resonators (MRRs) (e.g., [280, 276, 166, 236, 142, 259, 191, 241, 25, 146]). The recent work [142] has also demonstrated the use of an electrically active stack of Si-SiO₂-ITO layers to substantially increase the electro-optic modulation efficiency of a Si MRR based modulator. In this design, the light-guiding core layer (i.e., Si) critically contributes to the electro-optic activity in the modulator. *In contrast*, our SiN-on-SiO₂ modulator presented in this chapter employs *for the first time* an ITO-SiN-ITO thin-film stack in which the ITO thin films act as the active upper and lower claddings of the SiN MRR based core of the modulator.

Lithium Niobate (LN) Based Modulators

Lithium Niobate (LN) has recently emerged as the promising material for designing high-performance electro-optic modulators because of its wide bandgap and large second-order electro-optic coefficient. Several LN modulators have been demonstrated so far in the literature (e.g., [50, 49, 270, 271, 137]). For instance, in [270] and [271], thin-film LN-based electro-optic modulators have been demonstrated. Similarly, in [50],[49] and [137], a hybrid Si-LN platform based MRR modulators have

been presented. These LN modulators demonstrated in prior works, however, lack in modulation efficiency compared to the Si and SiN modulators from prior works.

Silicon Nitride (SiN) Based Modulators

Recently, silicon nitride (SiN) based PICs have gained tremendous attention due to their favorable properties compared to the traditional Si based PICs. As a result, several SiN-on-SiO₂ modulators have been demonstrated (e.g., [197, 9, 115, 5, 101]). In [197], a graphene integrated electro-optic SiN MRR modulator has been reported. In [5], a hybrid SiN-LN platform based racetrack resonator modulator has been presented. Similarly, SiN modulators based on lead zirconate titanate and zinc oxide/zinc sulphide as active materials are demonstrated in [9] and [101]. In [115], a SiN modulator that achieves tuning via photo-elastic effect has been demonstrated. Compared to these modulator designs from prior work, we present a different, ITO-based electro-refractive SiN-on-SiO₂ modulators that achieves relatively better modulation bandwidth, modulation efficiency, and FSR.

8.3 Design of Our SiN-on-SiO₂ Modulators

In this section, firstly we describe the structure and operating principle of our modulator designs. Then, we discuss the characterization results for our modulators that we have obtained through photonics foundry-validated simulations. We also compare our modulators with several Si, LN and SiN based MRR modulators from prior work, in terms of modulation bandwidth, modulation efficiency, and FSR.

Structure and Operating Principle

ITO-Based SiN-on-SiO₂ Modulator with ITO-SiN-ITO Active Cladding

Fig. 8.1(a) and Fig. 8.1(b), respectively, show the top-view and cross-sectional schematics of our SiN-on-SiO₂ MRR modulator. The active region in the upper and lower claddings of the modulator consists of indium tin oxide (ITO) thin films with silicon nitride material (SiN) in between (creating an ITO-SiN-ITO thin-film stack). From Fig. 8.1(b), we have a 300 nm thick SiN-based MRR waveguide between two 10 nm thick ITO films. Upon applying voltage across the ITO-SiN-ITO stack (through the Au pads shown in Fig. 8.1(a)), free carriers accumulate in the ITO films at the ITO-SiN interfaces for up to 5 nm depth in the ITO films [59], making these accumulation regions in the ITO films high-carrier-density active regions. This is due to the free-carriers-assisted, large-amplitude modulation in the permittivity and refractive index of the ITO material previously reported in [59]. We evaluate this free-carriers based index modulation in the ITO films using the Drude-Lorentz model from [158]. It can be inferred from the Drude-Lorentz model that as the carrier concentration in the ITO accumulation regions increases, the refractive index of the ITO films decreases. Our modulator design from Fig. 8.1 leverages this electro-refractive phenomenon in ITO. The free-carriers-induced decrease in the refractive index of the ITO thin films decreases the effective refractive index of the SiN-on-SiO₂

modulator, causing a blue shift in its resonance wavelength that in turn causes a transmission modulation in the MRR modulator. The electro-refractive activity of our SiN-on-SiO₂ MRR modulator is confined only in the ITO-based claddings. This is different from the Si-SiO₂-ITO capacitor based MRR modulator from [142], which has the electro-refractive activity in both its Si-based MRR core and SiO₂-ITO based cladding.

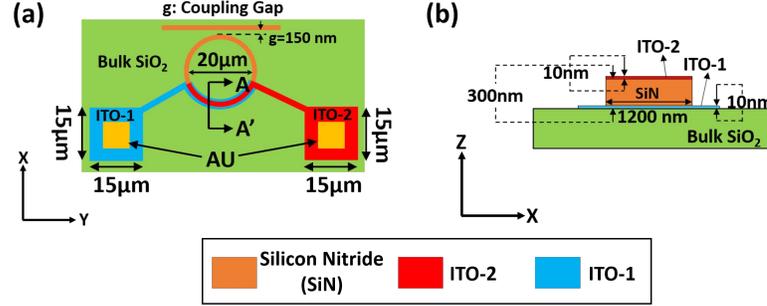


Figure 8.1: (a) Top view, (b) Cross-sectional view (along AA') of our SiN-on-SiO₂ MRR modulator with ITO-SiN-ITO stack as active upper cladding.

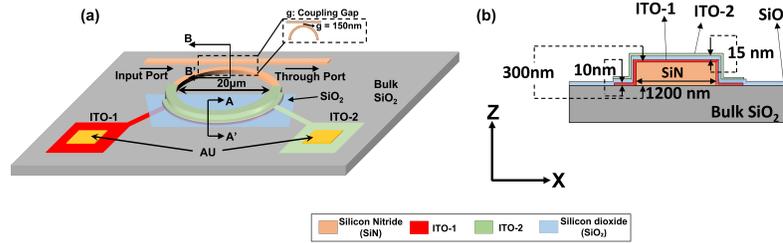


Figure 8.2: (a) Top view, (b) Cross-sectional view (along AA') of our SiN-on-SiO₂ MRR modulator with ITO-SiO₂-ITO stack as active upper cladding.

ITO-Based SiN-on-SiO₂ Modulator with ITO-SiO₂-ITO Active Cladding

Fig. 8.2(a) and Fig. 8.2(b), respectively, show the top-view and cross-sectional schematic of our SiN-on-SiO₂ MRR modulator. The active region in the upper cladding of the modulator consists of a stack of two indium tin oxide (ITO) thin films with a silicon dioxide (SiO₂) thin film in between (an ITO-SiO₂-ITO stack). From Fig. 8.2(b), we have a 300 nm thick SiN-based MRR waveguide, two 10 nm thick ITO films, and 15 nm thick SiO₂ layer. Upon applying voltage across the ITO-SiO₂-ITO stack (through the Au pads shown in Fig. 8.2(a)), free carriers accumulate in the ITO films at the ITO-SiO₂ interfaces for up to 5 nm depth in the ITO films [59], making these accumulation regions in the ITO films high-carrier-density active regions. In these regions, a free-carriers-assisted, large-amplitude modulation in the permittivity and refractive index of the ITO material has been previously reported [59]. We evaluate this free-carriers based index modulation in the ITO films using the Drude-Lorentz model from [158]. Accordingly, as the carrier concentration in the

ITO accumulation regions increases, the refractive index of the ITO films decreases. As a result, the effective refractive index of our SiN-on-SiO₂ modulator design from Fig. 8.2 also decreases, causing a blue shift in its resonance wavelength that in turn causes a transmission modulation at the through port of the modulator. The electro-refractive activity of our SiN-on-SiO₂ MRR modulator design is confined only in the ITO-SiO₂-ITO cladding. This is different from the Si-SiO₂-ITO capacitor based MRR modulator from [142], which has the electro-refractive as well as electro-absorptive activities in both its Si-based MRR core and SiO₂-ITO based cladding.

Simulations Based Characterization

ITO-Based SiN-on-SiO₂ Modulator with ITO-SiN-ITO Active Cladding

We performed electrostatic simulations of our ITO-SiN-ITO thin-film stack based SiN-on-SiO₂ modulator in the CHARGE tool of DEVICE suite from Lumerical [155], to evaluate the required voltage levels across the Au pads (Fig. 8.3(a)) for achieving various free-carrier concentrations in the ITO films. Then, based on the Drude-Lorentz dispersion model from [158], we extracted the corresponding ITO index change values for various free-carrier concentrations. These results are listed in Table 8.1. Using these index values from Table 8.1, we modeled our MRR modulator in the MODE tool from Lumerical [155] for finite-difference-time-domain (FDTD) and finite-difference eigenmode (FDE) analysis. For this analysis, we used the Kischkat model [131] of stoichiometric silicon nitride to model the MRR device. From this analysis, we extracted the effective index change and transmission spectra of our modulator (shown in Table 8.1 and Fig. 8.3 respectively) at various applied voltages for the operation around 1.615 μm wavelength (L-band). From Fig. ??, our modulator achieves ~ 4.5 nm resonance shift upon applying 17 V across the thin-film stack, which renders the resonance tuning (modulation) efficiency of ~ 280 pm/V. This is crucially significant as our MRR modulator has relatively very low overlap between the optical mode and free-carrier perturbation (only $\sim 10\%$ of the guided optical mode overlaps with the ITO-based claddings) compared to the state-of-the-art ITO-based modulators (e.g., [142]). Further, from the simulated spectra in Fig. 8.3, we evaluate the FSR of our modulator to be ~ 19 nm. Further, based on our device simulations using the Lumerical MODE tool, we evaluated the insertion loss and loaded Q-factor of our modulator to be ~ 0.23 dB and ~ 2300 respectively. We also evaluated the capacitance density of the ITO thin-films covering the MRR rim (using the Lumerical CHARGE tool) to be ~ 0.18 fF/ μm^2 for the 300 nm thick SiN layer, yielding the modulation bandwidth (3-dB RC bandwidth) of ~ 67.8 GHz for the modulator. We also modeled our modulator in Lumerical INTERCONNECT, to simulate optical eye diagrams for the modulator at 30 Gb/s and 55 Gb/s operating bitrates (Fig. 8.5). As evident (Fig. 8.5(b)), our modulator can achieve 10.31 dB extinction ratio for OOK modulation at 30 Gb/s bitrate.

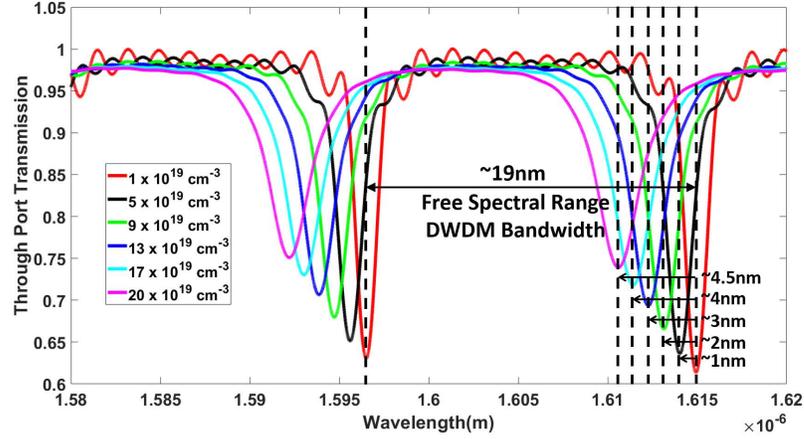


Figure 8.3: Transmission spectra of our modulator with ITO-SiN-ITO stack.

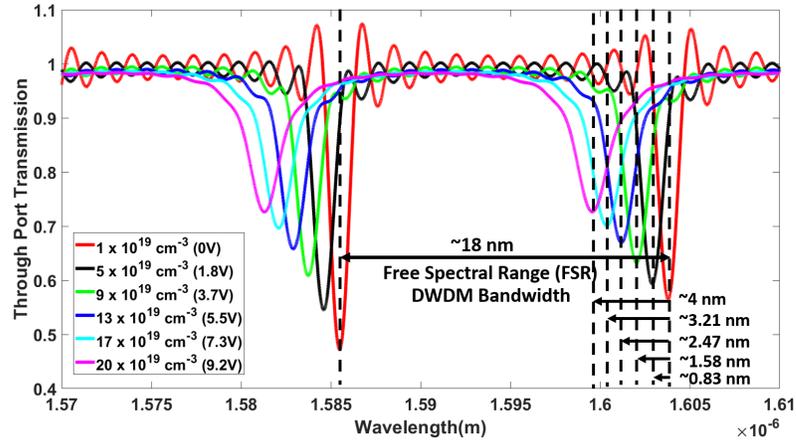


Figure 8.4: Transmission spectra of our modulator with ITO-SiO₂-ITO stack.

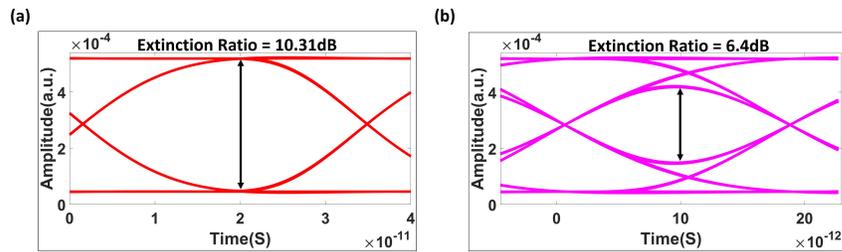


Figure 8.5: Optical eye diagrams for (a) 30 Gb/s and (b) 55 Gb/s OOK inputs to our SiN-on-SiO₂ (ITO-SiN-ITO stack) modulator.

ITO-Based SiN-on-SiO₂ Modulator with ITO-SiO₂-ITO Active Cladding

We performed electrostatic simulations of our ITO-SiO₂-ITO thin-film stack based SiN-on-SiO₂ modulator in the CHARGE tool of DEVICE suite from Lumerical [155], to evaluate the required voltage levels across the Au pads (Fig. 8.4(a)) for achieving various free-carrier concentrations in the ITO films. Then, based on the Drude-Lorentz dispersion model from [158], we extracted the corresponding ITO index

Table 8.1: Free-carrier concentration (N), real index ($\text{Re}(\eta_{ITO})$), and imaginary index ($\text{Im}(\eta_{ITO})$) for the ITO accumulation layer in our modulator. The real and imaginary effective index ($\text{Re}(\eta_{eff})$, $\text{Im}(\eta_{eff})$), operating voltage (V), and induced resonance shift ($\Delta\lambda_r$) for our modulator (ITO-SiN-ITO stack).

N (cm^{-3})	Re (η_{ITO})	Im (η_{ITO})	Re (η_{eff})	Im (η_{eff})	V	$\Delta\lambda_r$ (pm)
1×10^{19}	1.9556	0.0100	1.9973	2.651e-5	0	0
5×10^{19}	1.9111	0.0403	1.99434	2.6581e-5	3.4	991
9×10^{19}	1.8667	0.0896	1.99138	2.6587e-5	6.8	1890
13×10^{19}	1.8222	0.1289	1.98842	2.6593e-5	10.2	2970
17×10^{19}	1.7778	0.1582	1.98546	2.6598e-5	13.6	3910
20×10^{19}	1.7333	0.1874	1.9825	2.6604e-5	17.0	4470

Table 8.2: Free-carrier concentration (N), real index ($\text{Re}(\eta_{ITO})$), and imaginary index ($\text{Im}(\eta_{ITO})$) for the ITO accumulation layer in our modulator. The real and imaginary effective index ($\text{Re}(\eta_{eff})$, $\text{Im}(\eta_{eff})$), operating voltage (V), and induced resonance shift ($\Delta\lambda_r$) for our modulator (ITO-SiO₂-ITO stack).

N (cm^{-3})	Re (η_{ITO})	Im (η_{ITO})	Re (η_{eff})	Im (η_{eff})	V	$\Delta\lambda_r$ (pm)
1×10^{19}	1.9556	0.0100	1.9735	0.0001	0	0
5×10^{19}	1.9111	0.0403	1.9724	0.0003	1.8	830
9×10^{19}	1.8667	0.0896	1.9712	0.0006	3.7	1580
13×10^{19}	1.8222	0.1289	1.9701	0.0011	5.5	2470
17×10^{19}	1.7778	0.1582	1.9692	0.0017	7.3	3210
20×10^{19}	1.7333	0.1874	1.9680	0.0022	9.2	4000

change values for various free-carrier concentrations. These results are listed in Table 8.2. Using these index values from Table 8.2, we modeled our MRR modulator in the MODE tool from Lumerical [155] for finite-difference-time-domain (FDTD) and finite-difference eigenmode (FDE) analysis. For this analysis, we used the Kischkat model [131] of stoichiometric silicon nitride to model the MRR device. From this analysis, we extracted the effective index change and transmission spectra of our modulator (shown in Table 8.2 and Fig. 8.4 respectively) at various applied voltages for the operation around 1.6 μm wavelength (L-band). From Fig. 8.4, our modulator achieves up to 4 nm resonance shift upon applying 9.2 V across the thin-film stack, which renders the resonance tuning (modulation) efficiency of ~ 450 pm/V. This is a crucial outcome as our MRR modulator has relatively very low overlap between the optical mode and free-carrier perturbation (only 10% of the guided optical mode overlaps with the ITO-based upper cladding) compared to the silicon ITO-based modulators (e.g., [142]).

Fig. 8.6 illustrates the cross-sectional electric-field profiles of the fundamental TE mode evaluated for three different free-carrier concentrations of ITO, namely 1×10^{19} cm^{-3} , 9×10^{19} cm^{-3} , and 17×10^{19} cm^{-3} , at three different cross-sectional regions of

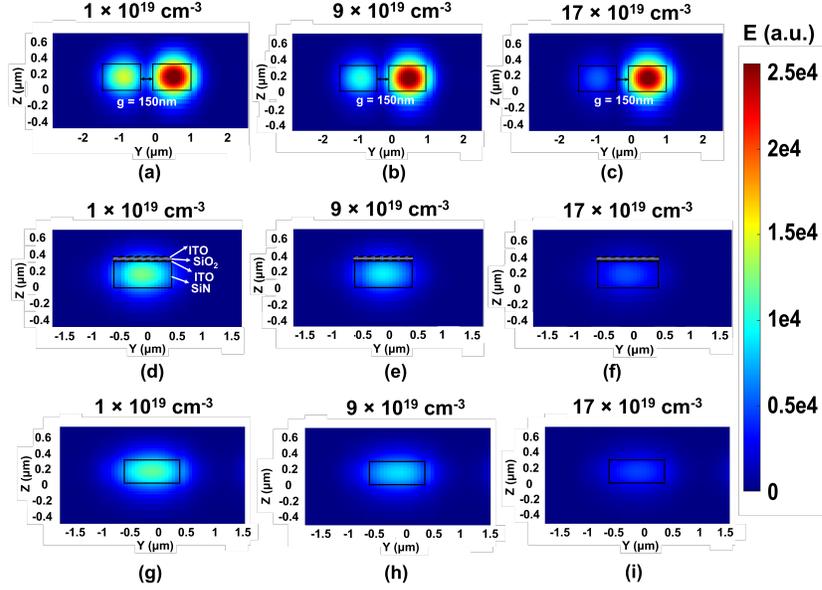


Figure 8.6: Cross-sectional electric-field profiles of the fundamental TE mode evaluated at the coupling section (along BB' in Fig. 8.2(a)) ((a)-(c)), across the rim (along AA' in Fig. 8.2) ((d)-(f)), and at the through port of our SiN-on-SiO₂ (ITO-SiO₂-ITO stack) MRR modulator ((g)-(i)), for three different free-carrier concentrations of ITO (Table 8.2) namely $1 \times 10^{19} \text{ cm}^{-3}$ (for (a),(d),(g)), $9 \times 10^{19} \text{ cm}^{-3}$ (for (b),(e),(f)), and $17 \times 10^{19} \text{ cm}^{-3}$ (for (c),(f),(i)), using the variational FDTD (varFDTD) solver [155].

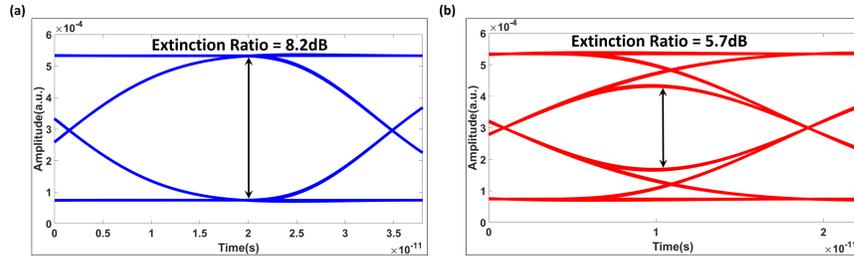


Figure 8.7: Optical eye diagrams for (a) 30 Gb/s and (b) 55 Gb/s OOK inputs to our SiN-on-SiO₂ (ITO-SiO₂-ITO stack) modulator.

our MRR modulator. To evaluate these profiles, we used the variational FDTD (varFDTD) solver of the MODE tool of the DEVICE suite from ANSYS/Lumerical. In fact, Fig. 8.6 shows a 3×3 grid of field profiles. Each row in this grid corresponds to field profiles collected for a particular cross-sectional region of our MRR modulator across different free-carrier concentrations of ITO. Similarly, each column in the grid corresponds to field profiles collected for a particular free-carrier concentration of ITO across three different regions of the modulator.

As per the discussion in the previous subsection, the increase in the free-carrier concentration in the ITO layers caused due to the increase in the applied bias across the ITO-SiO₂-ITO stack, decreases the effective index of our modulator. This in turn induces a blue shift in the resonance wavelength of our modulator. Due to

this blue shift in the resonance wavelength, the amount of optical power coupled from the input port into the MRR cavity at the coupling region decreases. This can be clearly observed from the field profiles collected at the coupling region along BB'; as the free-carrier concentration increases from Fig. 8.6(a) to Fig. 8.6(c), the intensity of the coupled field in the MRR at the cross-section BB' also decreases. The decrease in the coupled field intensity at BB' naturally results in the decrease of the steady-state field intensity inside the MRR waveguide. As a result, at the cross-section AA', the field intensity can be observed to decrease with the increase in the free-carrier concentration in the ITO layers, as we move from Fig. 8.6(d) to Fig. 8.6(f) in the middle row of Fig. 8.6. Atop the steady-state field intensity inside the MRR cavity, the field intensity at the through port (hence, the output optical power at the through port) of the MRR also decreases naturally with the increase in the free-carrier concentration. This can be observed in the bottom row of Fig. 8.6. The modulation of the optical output power at the through port with the change in the free-carrier concentration in the ITO layers corroborates the electro-refractive activity of our modulator.

In addition, as we move from the top row (coupling region field profiles) to the bottom row (through port field profile) within each column in Fig. 8.6, the field intensity slightly decreases. This provides evidence that, for each column (i.e., for each specific free-carrier concentration), the optical field intensity undergoes optical-loss-induced attenuation as the light waves travel along the propagation path from the coupling region (top row) to the through port (bottom row).

Further, from the spectra in Fig. 8.4, we evaluate the FSR of our modulator to be ~ 18 nm. We evaluated (using the Lumerical MODE tool) the insertion loss and loaded Q-factor of our modulator to be ~ 0.235 dB and ~ 2000 respectively. We also evaluated the capacitance density of the ITO thin-film stack covering the MRR rim (using the Lumerical CHARGE tool) to be ~ 2.3 fF/ μm^2 for the 15 nm thick SiO₂ layer. Moreover, we modeled our modulator in Lumerical INTERCONNECT, to simulate optical eye diagrams for the modulator at 30 Gb/s and 55 Gb/s operating bitrates (Fig. 8.7). As evident (Fig. 8.7), our modulator can achieve 8.2 dB extinction ratio for OOK modulation at 30 Gb/s bitrate.

Comparison and Discussion

ITO-Based SiN-on-SiO₂ Modulator with ITO-SiN-ITO Active Cladding

Fig. 8.8 shows a comparison of our SiN-on-SiO₂ modulator with the best performing Si (ten; [280]-[146]), LN (five; [50]-[137]) and SiN (five; [197]-[101]) MRR modulators from prior work, in terms of three key attributes, namely modulation efficiency, FSR, and modulation bandwidth. As evident from Fig. 8.8, our modulator achieves better performance compared to the existing SiN modulators and the state-of-the-art Si and LN modulators from prior works, which in turn promotes its use in DWDM-based high-performance PICs. Since our SiN-on-SiO₂ modulator achieves modulation bandwidth of ~ 67.8 GHz, it can be easily operated at the bitrate of >15 Gb/s to enable ultra-high-speed (potentially beyond Tb/s) DWDM-based PICs while ensuring

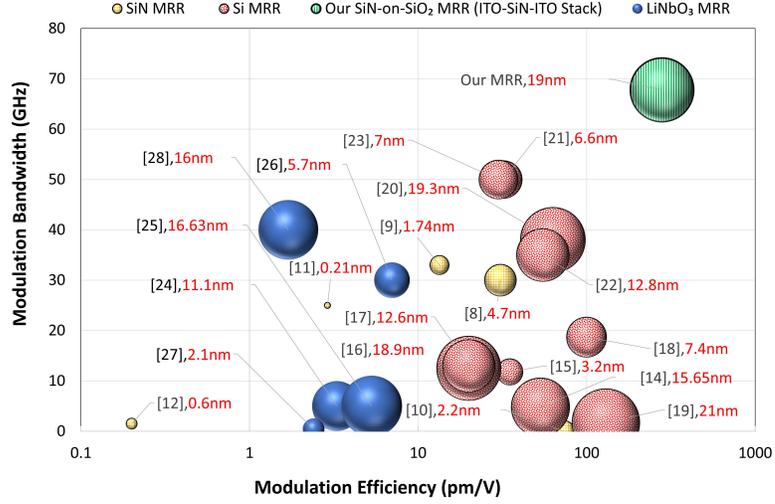


Figure 8.8: Modulation bandwidth, modulation efficiency and FSR (shown as the size of the bubbles and red data labels) of various Si, LN (LiNbO₃) and SiN MRR modulators from prior work, compared with our SiN-on-SiO₂ (ITO-SiN-ITO Stack) MRR modulator.

minimal power-penalty from crosstalk [18] and self-heating [218]. In addition, our SiN-on-SiO₂ modulator also achieves a modulation efficiency of ~ 280 pm/V. This in turn can enable dynamic operation of our modulator with energy-efficiency of < 100 fJ/bit [143]. Unfortunately, our modulator achieves relatively low loaded Q-factor of 2300. Nevertheless, we anticipate that the Q-factor can be increased to 5000-8000 by marginally trading the modulation bandwidth for better loss characteristics of the MRR cavity. Thus, future work should include an exhaustive search of design parameters, including the coupling gap, MRR waveguide width, MRR waveguide height, MRR radius, and the thicknesses of the ITO films, to minimize the coupling, bending and absorption losses in the MRR cavity without notably compromising the modulation efficiency. Having the loaded Q-factor of our modulator in the range of 5000-8000, while already having a greater than 12 nm FSR (Fig. 8.8), will enable balancing of the crosstalk penalty and modulation speed in our modulator, for high-performance DWDM based PICs [18]. Moreover, although ITO is not available in the CMOS process flow, it can be deposited at relatively low temperatures (less than 300°C) on top of the back-end-of-line (BEOL) metal layers of CMOS chips, independent of the CMOS FEOL process. This makes our SiN-on-SiO₂ modulator an excellent choice for implementing optical interconnect PICs on silicon interposers, to enable ultra-high-bandwidth inter-chiplet communication in emerging multi-chiplet systems [255].

ITO-Based SiN-on-SiO₂ Modulator with ITO-SiO₂-ITO Active Cladding

Table 8.3 shows a comparison of our SiN-on-SiO₂ MRR modulator with the simulation (marked as *) and fabrication based best-performing nine SiN MRR modulators from prior works ([197]-[101],[266]-[267],[294]), in terms of five key attributes namely

Table 8.3: Modulation bandwidth (MB) (optical (O) and Electrical (E)), modulation efficiency (ME), FSR and energy efficiency (EE) corresponding to various SiN based MRR modulators (modulator type (MT)) from prior works obtained from simulations (*) and experiments, compared with our simulated SiN-on-SiO₂ (ITO-SiO₂-ITO stack) MRR modulator.

MT	MB		ME (pm/V)	FSR (nm)	EE (pJ/bit)
	O-MB (GHz)	E-MB (GHz)			
[266]	0.06	0.02	1.6	4×10^{-3}	1×10^{-3}
[6]	1.7	N/A	0.01*	0.58	N/A
[151]	0.01	7.9	1	113	3.7×10^4
[267]	0.03	0.03	1.6	0.3	1×10^{-3}
[197]	161	30	100*	4.7	0.8
[9]	87	35.6	67*	1.74	N/A
[115]	0.19	1.3×10^{-3}	137.5*	2.2*	0.11
[5]	25	N/A	2.9	0.21*	N/A
[101]	1.55	5.9	0.2	0.6	53
[294]	1.77	N/A	5.8	0.4*	N/A
Ours	93.62*	53.1*	450*	18*	1.4*

optical modulation bandwidth (O-MB), electrical modulation bandwidth (E-MB), modulation efficiency (ME), FSR, and energy-efficiency (EE). The SiN MRR modulator in [197] achieves higher O-MB compared to the other SiN MRR modulators (Table 8.3) and our modulator. In contrast, our modulator achieves higher E-MB compared to the other SiN MRR modulators (Table 8.3). We have also evaluated that our modulator achieves the best effective MB of ~ 46.2 GHz compared to all other SiN MRR modulators, based on the formula of effective MB from [79]. Due to its superior effective MB of ~ 46.2 GHz, our modulator can be easily operated at >15 Gb/s bitrate to enable ultra-high-speed (potentially beyond Tb/s) DWDM-based PICs while ensuring minimal power-penalty from crosstalk [18]. Moreover, our modulator achieves higher ME compared to other SiN MRR modulators (Table 8.3). However, in terms of FSR, SiN MRR modulator demonstrated in [151] achieves higher FSR compared to the other SiN MRR modulators in Table 8.3 including our modulator. Nevertheless, our modulator consumes the energy of 1.4 pJ/bit which is significantly better than the energy consumption of the modulator from [151]. Its high energy efficiency and competent FSR of 18 nm make our modulator a favorable candidate for designing high-bandwidth and energy-efficient DWDM-based photonic interconnects for datacenter-scale as well as chip-scale computing and communication architectures.

Further, although ITO is not available in the CMOS process flow, it can be deposited at relatively low temperatures (less than 300°C) on top of the back-end-of-line (BEOL) metal layers of CMOS chips, in an independent manner without interfering with or contaminating the CMOS front-end-of-line (FEOL) and BEOL processes. This makes our SiN-on-SiO₂ modulator an excellent choice for implementing optical

interconnect PICs on silicon interposers, to enable ultra-high-bandwidth inter-chiplet communication in emerging multi-chiplet systems [255].

In summary, we advocate that our SiN-on-SiO₂ modulators can achieve better performance compared to the existing SiN based MRR modulators from prior work. The obtained results corroborate our modulators potential to consequently enable DWDM-based SiN-on-SiO₂ PICs that will offer highly scalable and energy-efficient solutions to a wide range of mature and emerging applications, including datacenter transceivers [44], high-performance computing [210], signal processing [128], optical computing [123], and artificial intelligence [226].

8.4 Summary

In recent years, the SiN-on-SiO₂ platform has attained tremendous attention for realizing PICs because it has several advantageous properties over the conventional SOI platform. Despite these advantages, the SiN-on-SiO₂ platform lacks high-performance active devices such as modulators. To address this drawback, we have demonstrated ITO based SiN-on-SiO₂ MRR modulators, which consists of a stack of ITO-SiN-ITO and ITO-SiO₂-ITO thin films as the active upper cladding of the SiN MRR core, respectively. This active upper cladding of our modulators leverage the free-carrier assisted, high-amplitude refractive index change in the ITO films to effect a large electro-refractive optical modulation in the device. To evaluate the performance of our SiN-on-SiO₂ MRR modulators, we performed electrostatic, transient and finite difference time domain (FDTD) simulations using the foundry-validated Ansys/Lumerical tools. Based on these simulations, our modulators achieve superior performance that demonstrates their potential to enhance the performance and energy-efficiency of SiN-on-SiO₂ based PICs of the future.

Chapter 9 A Low-Dissipation and Scalable GEMM Accelerator with Silicon Nitride Photonics

9.1 Introduction

Deep Neural Networks (DNNs) have revolutionized the implementation of various artificial intelligence tasks, such as image recognition, language translation, autonomous driving [134, 153], due to their high inference accuracy. However, DNNs are computationally intensive, due to inherently abundant linear computations such as general matrix-matrix multiplications (GEMM), which are at the core of DNN operations [69]. This computational intensity of processing the GEMM operations of DNNs is on a rapid rise owing to the ongoing rapid evolution of DNN models. This has pushed for highly customized hardware GEMM accelerators [23]. Among GEMM accelerators demonstrated in the literature, silicon-photonics accelerators have shown great promise to provide unparalleled parallelism, ultra-low latency, and high energy efficiency [152, 92, 26, 53, 284, 243]. In particular, Microring Resonator (MRR)-enabled silicon-photonics GEMM accelerators have shown disruptive performance and energy efficiencies, due to the compact footprint of MRRs, low dynamic power consumption of MRRs, and the ability of MRRs to support a massive fan-in of optical signals through dense-wavelength-division multiplexing (DWDM). These advantages have rendered up to $1000\times$ more processing throughput and up to $100\times$ better energy efficiency to MRR-enabled silicon-photonics GEMM accelerators than their electronic counterparts [26, 68].

However, the state-of-the-art MRR-enabled GEMM accelerators that are realized using the traditional silicon-on-insulator (SOI) material platform face two shortcomings. First, the high refractive index contrast between the silicon core (Si) and cladding (SiO_2) of an SOI waveguide leads to an enhanced interaction of the guided optical mode with the rough sidewalls of the waveguide. This introduces high scattering losses in the SOI waveguides [14]. Second, the presence of Two-Photon Absorption (TPA) in silicon has detrimental effects on SOI devices, particularly waveguides. These effects lead to substantial absorption losses in SOI waveguides, particularly when a moderate-to-high number of multiplexed optical signals are propagating inside an SOI channel waveguide. To counter these losses, a higher input optical power becomes necessary. However, this increased input optical power whittles down a significant part of the optical power budget, significantly hampering the achievable spatial parallelism, throughput, and energy efficiency of SOI-based photonic GEMM accelerators.

To address these shortcomings, we present a novel Silicon Nitride (SiN)-on- SiO_2 -based photonic GEMM accelerator named SiNPhAR. Our SiNPhAR accelerator integrates Indium Tin Oxide (ITO)-based SiN-on- SiO_2 MRR modulators (MRMs) within its input and weight banks, that are coupled to SiN-on- SiO_2 waveguides. These MRMs perform high-speed electro-optical encoding of electrical inputs and weights onto optical signals. These input and weight banks seamlessly integrate with our

invented balanced photo-charge accumulator (BPCA) to perform dot product operations of a large size. Unlike the SOI platform, the SiN-on-SiO₂ platform has a lower refractive index contrast between the core (SiN) and cladding (SiO₂) materials. This enables the design of ultra-low loss (<0.5 dB/cm [182, 14]) photonic waveguides. Additionally, the absence of free-carriers in the SiN material eliminates the possibility of TPA [42, 14]. This characteristic enables SiN-on-SiO₂ photonic waveguides to support a higher count of multiplexed optical signals (higher fan-in) without incurring excess absorption or scattering losses. Reduced optical losses empower our SiNPhAR accelerator to achieve superior spatial parallelism, enhanced throughput, and energy efficiency compared to prior SOI-based photonic GEMM accelerators.

Our key contributions in this chapter are summarized below:

- We illustrate the use of our ITO-based SiN-on-SiO₂ MRMs as input and weighting elements, enabling massively parallel multiplication operations;
- We design an accelerator architecture called SiNPhAR, which is based on the SiN-on-SiO₂ platform, and evaluate its achievable spatial parallelism;
- We compare the throughput and energy efficiency results of our SiNPhAR architecture with an SOI-based MRR-enabled GEMM accelerator from prior works.

9.2 Preliminaries

Background on SOI-Based Photonic GEMM Accelerators

Among the SOI-based photonic GEMM accelerators showcased in the literature, we focus on the MRR-enabled SOI-based incoherent GEMM accelerators [8, 26, 243, 226, 263, 262]. These accelerators mainly employ multiple analog tensor processing cores (TPCs) that operate in parallel, in which each TPC is utilized to perform a dot product operation. Typically, each TPC is made up of five essential blocks [263] (see Fig. 9.1 for an example TPC organization with five blocks): (i) a laser block that employs N laser diodes (LDs) to generate N optical wavelength channels; (ii) an aggregation block that aggregates the optical wavelength channels generated by LDs into a single photonic waveguide through DWDM technique by employing a $N \times 1$ multiplexer, and then splits the optical power of each of these wavelength channels equally into M separate waveguides by using a $1 \times M$ splitter; (iii) a modulation block that consists of M banks of MRMs spread across M dot product elements (DPEs), with each DPE employing one MRM bank; (iv) a weighting block that consists of another M banks of MRMs spread across the M DPEs, with each DPE employing one MRM bank; and (v) a summation block that comprises of a total of M summation elements (SEs), with each SE corresponding to a DPE and employing two photodiodes in a balanced configuration, commonly referred to as balanced photodiode (BPD) configuration, connected to a transimpedance amplifier (TIA) and an analog-to-digital converter (ADC). Typically, the laser block and SE block are placed at the two ends of the TPC, whereas the aggregation, modulation, and weighting blocks are placed in between them. Furthermore, based on the positioning of these intermediate blocks,

the MRR-based TPC organizations demonstrated in the prior works can be classified into three categories namely *Aggregate, Modulate, Weight (AMW) TPC*, *Modulate, Aggregate, Weight (MAW) TPC*, and *Modulate, Weight, Aggregate (MWA) TPC*. In the AMW TPC, the aggregation block is positioned first, followed by the modulation and the weighting blocks. On the other hand, in the MAW TPC organization, the modulation block is positioned first, followed by the aggregation and the weighting blocks. For additional details on the AMW and MAW TPCs, we direct the reader to [263].

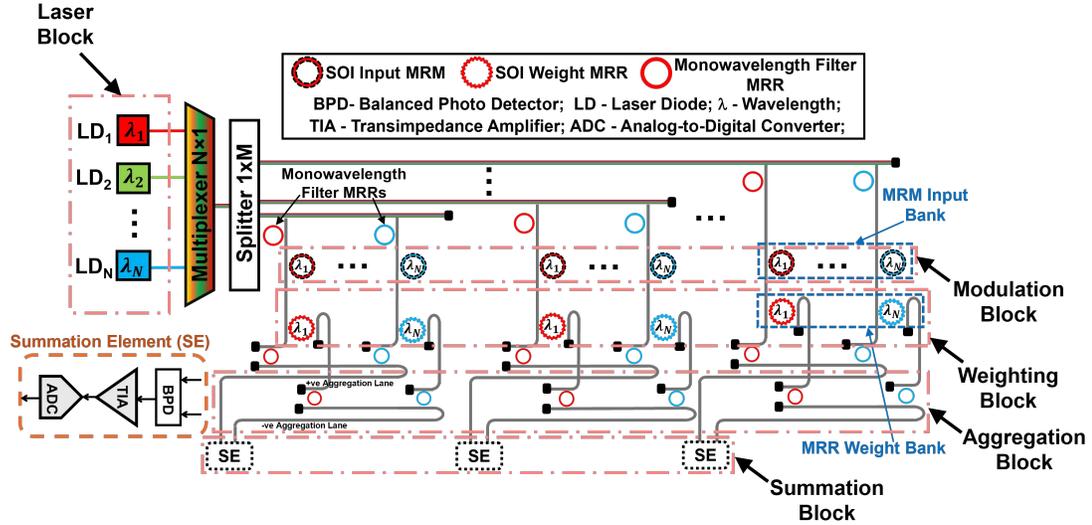


Figure 9.1: Illustration of the MWA organization of an SOI-based TPC.

For detailed elucidation, Fig. 9.1 illustrates the organization of an MWA TPC. As illustrated, the modulation and weighting blocks are placed before the aggregation block. In particular, the arrangement of each input-weight MRM pair is spectrally hitless [53] ensuring that each input-weight MRM pair produces a multiplication result and modulates this result onto a single-wavelength optical signal. This design eliminates inter-wavelength interference known as inter-modulation crosstalk [184], at the MRMs, providing a notable advantage over the AMW and MAW organizations [263]. There are a total of M DPEs in the TPC. And, in each DPE, there are a total of N input-weight MRM pairs, with each MRM pair acting as a multiplier. The modulation and weighting blocks are connected to the aggregation block via a set of mono-wavelength filter MRMs. The aggregation block consists of positive and negative aggregation lanes that guide the signals to the SE block.

Shortcomings of SOI Photonic GEMM Accelerators

The SOI-based photonic GEMM accelerators face two major shortcomings that hinder their scalability, throughput, and energy efficiency.

High Scattering Losses Due to High Index Contrast

The high refractive index contrast between the silicon core (Si, 3.5) and the cladding (SiO₂, 1.5) in SOI-based waveguides serves as a double-edged sword. While it allows for the design of compact photonic waveguides by tightly confining the guided optical modes to the core, it also makes these waveguides highly susceptible to significant scattering losses. This is because it leads to an enhanced interaction of the confined optical modes with the rough sidewalls of the waveguide. This enhanced mode-roughness interaction increases the scattering losses in the SOI waveguides. Studies have shown that even a slight RMS roughness of a few nanometers on the sidewalls, which is unavoidable due to fabrication imperfections, can result in substantial waveguide losses, often exceeding 3 dB/cm in SOI channel waveguides [122, 14].

High Absorption Losses Due to Two-Photon Absorption

The presence of free carriers in Si induces Two-Photon Absorption (TPA) in SOI-based photonic devices at telecom wavelengths. TPA increases the free-carrier density in the core (Si) material, leading to the Free-Carrier Absorption (FCA) effect in SOI waveguides [14, 122]. FCA results in higher absorption losses in SOI waveguides, particularly at elevated optical power levels. In DWDM applications, where multiple wavelengths are coupled into each SOI waveguide, the total optical power within the waveguide rises, triggering TPA-induced FCA effects and subsequently causing high absorption losses. Previous studies have shown that if the total number of wavelengths multiplexed into an SOI waveguide exceeds 20, the optical losses experienced by each additional wavelength channel propagating inside the waveguide increase by 0.1 dB/cm/wavelength [136, 182].

Motivation

To compensate for the high scattering and absorption losses in SOI waveguides, one option is to increase the input optical power. However, this increased input power whittles down a large portion of the optical power budget, leaving a small portion of the power budget available to support the scalability of size and spatial parallelism in SOI-based photonic GEMM accelerators. Additionally, the need for higher input optical power undermines the energy efficiency benefits associated with photonic GEMM accelerators. Therefore, there is a need for an alternative that can alleviate the optical signal losses and their detrimental impacts in photonic GEMM accelerators.

9.3 SiNPhAR Architecture

To alleviate the high scattering and absorption losses and related issues present in photonic GEMM accelerators, we propose to redesign photonic GEMM accelerators with the silicon nitride (SiN)-on-SiO₂ material system. Our idea is to address the root causes of high scattering and absorption losses in SOI-based designs, namely the high index contrast and TPA effect. The proposed SiN-on-SiO₂ material platform has been

shown to exhibit ultra-low waveguide propagation losses (absorption + scattering losses) (<0.5 dB/cm [182, 14]) due to its low refractive index contrast compared to the SOI platform. In addition, the absence of free carriers in the SiN material eliminates the possibility of TPA-induced increase in absorption losses [42, 14]. Our forged GEMM accelerator architecture based on the SiN-on-SiO₂ platform, which we call SiNPhAR architecture, is described in the following subsections.

Overview of SiNPhAR Tensor Processing Core (TPC)

The main processing unit of our SiNPhAR architecture is a tensor processing core (TPC) (illustrated in Fig. 9.2), which follows the MWA TPC organization described in Section II.A, with several critical modifications in the constituent blocks. Across the modulation, aggregation, and weighting blocks, all the utilized photonic devices, including the waveguides, MRMs, and filter MRRs, are based on the SiN-on-SiO₂ material platform. We take the designs of the SiN-on-SiO₂ waveguides from [39] and filter MRRs from [107], whereas we invent a new Indium Tin Oxide (ITO)-based all-pass design for SiN-on-SiO₂ MRMs (discussed in chapter 8). In addition, as the summation (SE) block, we utilize our newly invented balanced photo-charge accumulator (BPCA) (discussed in Section 9.3). Atop these modifications, a SiNPhAR TPC employs all-pass MRMs in its weighting blocks, which is different from the add-drop MRRs used in the weighting blocks of SOI-based MWA TPC. Because of this difference, a SiNPhAR TPC utilizes a filter MRR after each input-weight MRM pair. This filter MRR allows routing of the optical signal incoming from the input-weight MRM pair onto the positive or negative aggregation lanes, depending on the sign of the multiplication result produced by the input-weight MRM pair. The structure and functionality of various blocks of a SiNPhAR TPC are discussed in the upcoming subsections.

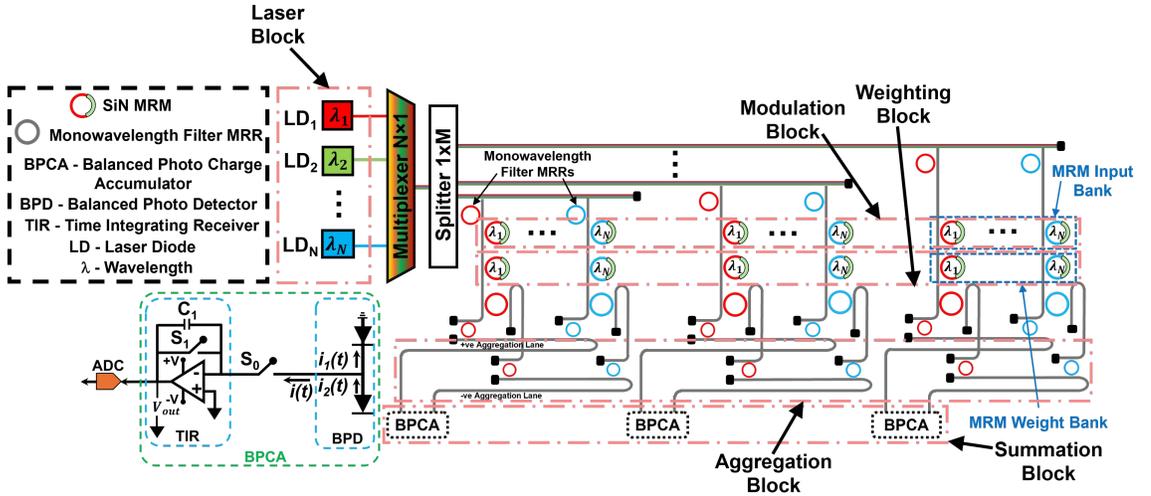


Figure 9.2: Schematic of a TPC of our SiNPhAR GEMM Accelerator.

ITO-Based SiN MRM for Input Encoding

our ITO-based SiN-on-SiO₂ MRM, illustrated in chapter 8 can produce a high-speed optical signal. This signal can be generated as a temporal train of optical amplitude symbols, similar to how an SOI MRM is used in SOI-based photonic GEMM accelerators to generate an optical signal as a temporal train of optical symbols [26]. In this signal, the amplitude of each symbol represents an analog input value. Thus, our MRM, when used for input encoding, can produce an optical signal as a temporal train of analog input values.

ITO-Based SiN MRM for Weighting

Our unique ITO-based SiN-on-SiO₂ MRM (chapter 8) serves a dual purpose in our TPC, functioning not only as a high-speed electro-optic input encoding element but also as a high-speed electro-optic weighting element. In this chapter, we explored its application in performing precise weighting of input-modulated optical signals. To assess the effectiveness of our MRM in this context, we conducted a comprehensive study with weighting values of 3-bit resolution.

Intuitively, a weighting MRM of 3-bit (4-bit) resolution should be able to alter the transmission of an input optical amplitude/symbol to one of the $2^3=8$ ($2^4=16$) distinct output amplitude levels. These 8 or 16 distinct output amplitude levels are achieved in our MRM at its through port by enabling electro-optic shifting of its resonance passband to 8 or 16 distinct spectral locations. Consequently, to imprint a certain 3-bit or 4-bit weighting on an input optical symbol, the input optical symbol is applied at the input port of the MRM, and then, the analog-converted (via a digital-to-analog converter) 3-bit or 4-bit weight value is applied to the electrical I/O pads of the MRM (see chapter 8) to effect an electro-optic shifting of the MRM's resonance passband. The shifted passband programs the through-port transmission of the input optical symbol to a corresponding output amplitude value from the 8 or 16 possible output amplitude levels. Fig. 9.3 illustrates how the shifting of our MRM's resonance passband enables weighting with 3-bit resolution. In the figure, λ_T shows the optical wavelength carrying the input optical symbol, and the 8 resonance passbands show the electro-optically shifted spectral locations corresponding to 8 output transmission amplitudes. When the spectral position of the passband is shifted, the intersection point of the passband with λ_T changes, which in turn alters the transmission amplitude. Thus, our MRM can be used to implement a weighting of an input optical symbol.

This operation of MRM weighting element can also be used to weight a high-speed optical signal output from an input-encoding MRM. For that, the MRM resonance passband is shifted to achieve different weighting amplitudes at a speed that is matched to the symbol rate of the input high-speed optical symbol. As a result, each symbol of the input optical signal is weighted with a unique weighted value to generate a weighted optical signal. Each symbol of the weighted optical signal, thus, represents a multiplication result (product) between the input and weight values. Therefore, the weighted optical signal generated by a weighting MRM is also referred

to as an optical product signal. Each symbol of this optical product signal, depending on its sign, is routed to the positive or negative aggregation lane, in the aggregation block of the SiNPhAR TPC.

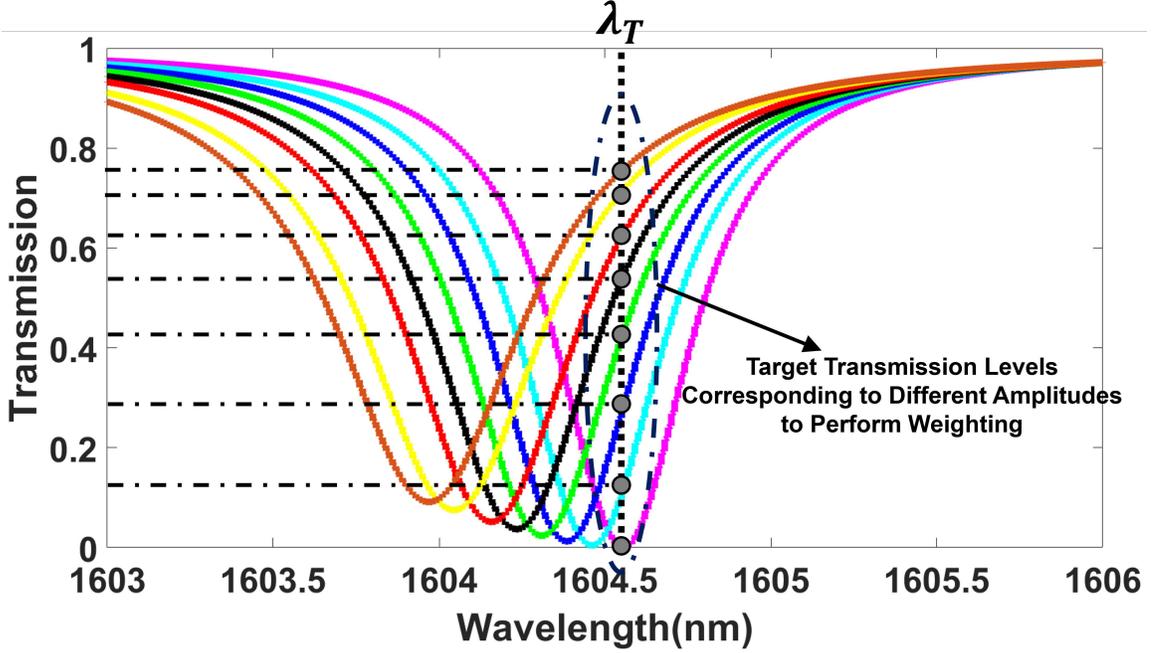


Figure 9.3: Transmission spectra measured at the through port of our SiN-on-SiO₂ MRM weighting element for various transmission amplitudes. These different transmission amplitudes at λ_T signify the weighting of the input optical amplitude symbol.

Summation with Balanced Photo-Charge Accumulator (BPCA)

In a DPE of a SiNPhAR TPC, each weighting MRM outputs one optical product signal, and a total of N such optical product signals are aggregated into the positive and negative aggregation lanes. The aggregation lanes deliver these optical product signals to the BPCA circuit for summation. Our BPCA circuit is collectively inspired by the time integrating receiver (TIR) design from [230] and the photodetector-based optical pulse/symbol accumulator design from [43]. As illustrated in Fig. 9.2, a BPCA circuit employs two photodiodes, each connected to the positive and negative aggregation lanes. These photodiodes are interlinked in a balanced configuration, commonly referred to as a balanced photodiode (BPD) configuration. The BPD is connected to a TIR. The TIR comprises an amplifier and a feedback capacitor/switch pair (Fig. 9.2).

A total of N optical product symbols arrive at the BPCA during a symbol cycle. The constituent BPD of the BPCA performs an incoherent superposition (signed summation) of all these N optical product symbols received within that cycle. Consequently, the incoherent superposition first enables the creation of a net optical symbol. The total optical energy packetized within a new optical symbol is proportional to the signed summation of the N optical product symbols. The BPD transduces this

net optical symbol into a balanced photocurrent symbol, which is further transduced by the TIR of the BPCA into an analog voltage level accrued on the capacitor of the TIR. This accrued analog voltage level, thus, represents a summation of N products, i.e., an N -sized dot product.

This N -sized dot product result, in the form of the accrued analog voltage level, can be held by the TIR. As new net optical symbols keep arriving in subsequent symbol cycles, the TIR enables a gradual integration (temporal accumulation) of the individual dot product results over multiple symbol cycles to generate a larger ($>N$ -sized) dot product result. This is possible because the N -sized dot product results arriving at the TIR can sequentially charge the TIR's capacitor so that the net accumulated charge and, consequently, the net analog voltage accrued on the capacitor over multiple symbol cycles provides the signed sum of the individual dot product results. This final sum value in the analog voltage format can be sampled and sent to the analog-to-digital converter (ADC) for conversion in the binary format. Thus, the BPCA of a SiNPhAR DPE can essentially enable the processing of very large-sized ($>N$ -sized) dot products.

9.4 Evaluation and Discussion

Scalability Analysis

To perform scalability analysis, we utilized the equations provided in [8], reproduced as Eqs. 1, 2 and 3. The parameters and their corresponding values [8, 263, 262] required to solve these equations are listed in Table 9.1. We devised a two-step procedure to determine the optimal value of N and M (N refers to the count of input-weight MRM pairs per DPE, whereas M refers to the count of DPEs per TPC) for a given bit precision and data rate (DR), as outlined below.

Step 1. We calculate the photodiode (PD) sensitivity by solving Eq. 1 for the specified bit precision and DR.

Step 2. Next, we perform an exhaustive search to find the optimal value of N (assuming $N = M$) for the specified bit precision and DR, using Eqs. 2 and 3. In this step, we solve Eq. 3, which represents the error function (ef). The ef is the difference between the optical power reaching the photodiode (P_{output}), calculated from Eq. 2, and the PD sensitivity obtained in Step 1/Eq. 1. We sweep for different values of N , and the optimal value of N for the specified bit precision, and DR is the one for which the ef yields the minimum positive value. Notably, when solving Eqs. 2 and 3, we consider P_{inc} (see Table 9.1 for definition) to be zero for ' N ' values less than 20 wavelengths/waveguide. However, beyond 20 wavelengths/waveguide, we account for a changing P_{inc} between the SOI and the SiN waveguides, as reported in Table 9.1. This is because the TPA-induced absorption losses in the SOI waveguide substantially increase if the total number of multiplexed wavelengths in an SOI waveguide exceeds 20 (as discussed in Section 9.2). In contrast, this phenomenon is not observed in SiN waveguides, as detailed in Section III.

For our analysis, we considered bit-precision values ranging from 1-bit to 4-bits and a set of DRs namely 1GS/S, 5GS/S, and 10GS/S. The results of our analysis are

$$B = \frac{1}{6.02} \left[20 \log_{10} \left(\frac{RP_{PD-opt}}{\left(\sqrt{2q(RP_{PD-opt} + I_d) + \frac{4KT}{R_L}} + (RP_{PD-opt})^2 RIN + \sqrt{2qI_d + \frac{4KT}{R_L}} \right) \sqrt{\frac{DR}{\sqrt{2}}}} \right) - 1.76 \right] \quad (1)$$

$$P_{output}(\text{dBm}) = P_L - P_{SMF} - P_C - (P_{WG-IL} \times d_{MRR} \times N) - (P_{Inc} \times d_{MRR} \times (N-20)) - (P_{sp} \times \log_2(N)) - P_{MRM} - P_{MRR} \\ - ((N-1) \times P_{MRM-OBL}) - ((N-1) \times P_{MRR-OBL}) - P_{penalty} \quad (2)$$

$$ef(B, DR, N) = P_{output}(N) - P_{PD-opt}(B, DR) \quad (3)$$

illustrated in Fig. 9.4. In addition to our SiNPhAR, we conducted the scalability analysis for an SOI-based MWA TPC (Fig. 9.1) named as *SOI-MWA*. From Fig. 9.4, our SiNPhAR can support a larger value of N compared to SOI-MWA. For instance, our SiNPhAR can support $N=52$ for a bit-precision of 3-bits, $DR = 1\text{GS/S}$, and an input laser power of 10dBm, which is larger compared to SOI-MWA that supports $N=35$. This advantage primarily stems from the reduced propagation losses in the SiN-on-SiO₂ waveguides and the lower insertion loss of the ITO-based SiN-on-SiO₂ MRMs in SiNPhAR, compared to the SOI waveguides and MRMs in SOI-MWA. Consequently, this creates a larger room in the optical power budget, allowing for the accommodation of a larger N in our SiNPhAR TPC.

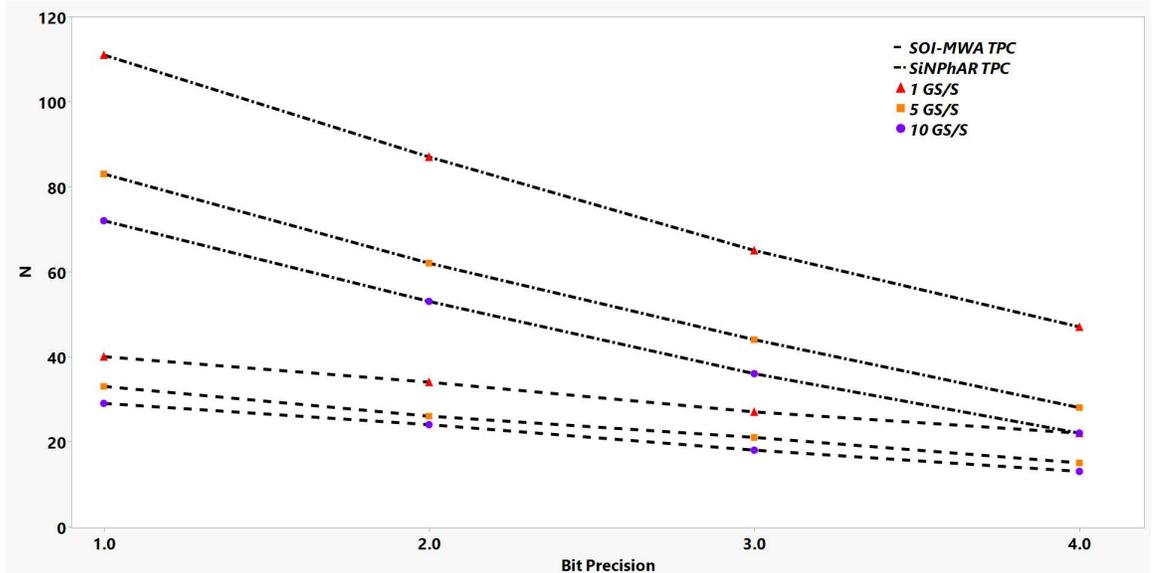


Figure 9.4: Supported TPC size $N(=M)$ for bit precision = $\{1,2,3,4\}$ -bits at data rates (DRs)= $\{1,5,10\}$ GS/s, for SOI-MWA TPC and SiNPhAR TPC.

Table 9.1: Definition and values of various parameters used in Eq. 1, Eq. 2, and Eq. 3 (from [8]) for the scalability analysis.

Parameter	Definition	Value
P_L	Laser Power Intensity	10 dBm
P_{SMF}	Attenuation by the Single Mode Fiber	0 dB
P_C	Fiber-to-Chip Coupling Insertion Loss	1.6 dB
P_{WG-IL}	Propagation Loss of SOI Waveguide	1.5 dB/cm
	Propagation loss of SiN Waveguide	0.5 dB/cm
P_{inc}	Increase in Propagation Loss of the SOI waveguide atop 20 λ s/waveguide	0.1 dB/cm/ λ
	Increase in Propagation Loss of the SiN waveguide atop 20 λ s/waveguide	0.01 dB/cm/ λ
P_{SP}	Splitter Insertion Loss	0.01 dB
P_{MRM}	Transmission Insertion Loss of the SOI MRM	4 dB
	Transmission Insertion Loss of the SiN MRM	0.235 dB
P_{MRR}	Transmission Insertion Loss of the SOI MRR	0.01 dB
$P_{MRM-OBL}$	Out-of-Band Insertion Loss (OBL) of the MRM	0.01 dB
$P_{MRR-OBL}$	Out-of-Band Insertion Loss (OBL) of the MRR	0.01 dB
$P_{Penalty}$	Network Penalty for SOI-MAW	1.8 dB
	Network Penalty for SiNPhAR	
R	PD Responsivity	1.2
q	Charge of an Electron (C)	1.6×10^{-19}
I_d	PD Dark Current	35 nA
K	Boltzmann Constant (J/K)	1.38×10^{-23}
T	Absolute Temperature (K)	300
R_L	Load Resistance (Ohms)	50
RIN	Relative Intensity Noise (dB/Hz)	-140
B	Bit-Precision	
P_{PD-OPT}	PD Sensitivity	-

System-Level Evaluation Method

System-Level Implementation

Fig. 9.5 illustrates the general system-level implementation of a photonic GEMM accelerator. It consists of global memory that stores convolutional neural network (CNN) parameters and a pre-processing and mapping unit. It has a mesh network of tiles. Each tile contains 4 dot-product units (DPUs) (a DPU is synonymous/analogous to a TPC) interconnected (via H-tree) with a unified buffer as well as pooling and activation units. Each TPC/DPU consists of multiple DPEs and each DPE is equipped with a dedicated input and output FIFO buffer [265] to store intermittent weights, inputs, and partial sum values. The generic DPUs/TPCs in the system are replaced with SiNPhAR TPCs and SOI-MWA TPCs, respectively, to derive SiNPhAR and SOIPhAR accelerator system architectures.

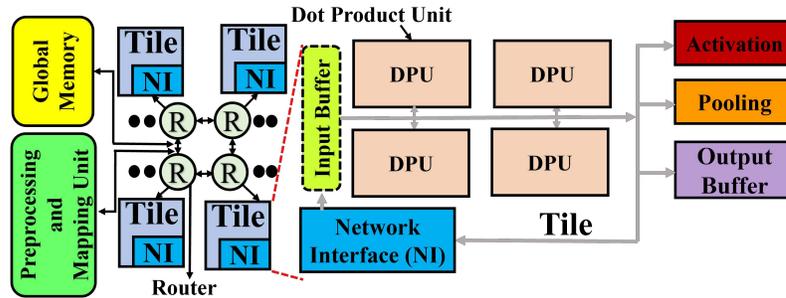


Figure 9.5: System-level implementation of SiNPhAR accelerator. DPU=TPC.

Simulation Setup

In our study, we employed a custom Python-based simulator to emulate the system-level deployment of SiNPhAR and SOIPhAR accelerator architectures. The simulation involved the inference of four distinct CNN models (with a batch size of 1): ShuffleNet V2 [293], GoogleNet [247], and ResNet50 [98]. We converted the convolutional layers and fully connected layers of these CNNs into GEMM operations using the Toeplitz matrix transformations or im2col functions [263, 262], and then accelerated these GEMM operations on our considered accelerators. We conducted a comparative analysis of SiNPhAR and SOIPhAR accelerator architectures in the context of inferring 8-bit integer quantized CNN models. Key metrics such as Frames per second (FPS) and FPS/W (energy efficiency) were evaluated. All accelerators were operated across data rates of 1GS/s, 5GS/s, and 10GS/s. Each TPC was operated at 4-bit precision; therefore, two TPCs were used with back-end shift-and-add circuits to achieve 8-bit computational precision. For these specific data rates, SOIPhAR and SiNPhAR achieve N (TPC size) as shown in Table 9.2. Our evaluation is based on output stationary dataflow. To ensure a fair comparison, we carried out an area proportionate analysis, wherein we adjusted the TPC count for each SiNPhAR and SOIPhAR variants listed in Table 9.2 so that the total area consumption of all TPCs per variant remained constant across all variants.

Table 9.1 outlines the parameters used for our evaluation, while Table 9.3 provides the parameters used for assessing the overhead of the peripherals in our evaluated accelerators. We set each laser diode to emit an input optical power of 10 mW (10 dBm) (Table 9.1). The parameters for the multiplexer and splitter were sourced from [152].

Table 9.2: TPC size (N) and TPC Count ($\#$) at 4-bit precision across various data rates for various accelerator architectures.

	Datarate					
	1 GS/s		5 GS/s		10 GS/s	
TPC	N	#	N	#	N	#
SOIPhAR	22	132	15	155	13	162
SiNPhAR	47	50	28	95	22	116

System-Level Evaluation Results

In Fig. 9.6(a), the Normalized FPS results for various accelerators with a batch size of 1, operating at different datarates are presented. These results are normalized to SOIPhAR for ResNet50 [98] at a datarate of 10 GS/s. SiNPhAR accelerators outperforms SOIPhAR accelerators in terms of gmean across four CNN models at all datarates. Specifically, at 1 GS/s, SiNPhAR achieves up to $1.7\times$ better FPS than SOIPhAR. As the datarate increases to 5 GS/s, SiNPhAR exhibits further improvements in FPS over SOIPhAR, achieving up to $1.8\times$ better FPS than SOIPhAR. These remarkable throughput improvements in SiNPhAR are attributed to two main factors. Firstly, the SiNPhAR architecture utilizes SiN-based active and passive devices to implement analog GeMM functions. The low optical signal losses in the SiNPhAR architecture, owing to the low-index contrast and absence of two-photon absorption

Table 9.3: Accelerator Peripherals and TPC Parameters [263].

	Power(mW)	Latency	Area(mm^2)
Reduction Network	0.050	3.125ns	3.00E-5
Activation Unit	0.52	0.78ns	6.00E-5
IO Interface	140.18	0.78ns	2.44E-2
Pooling Unit	0.4	3.125ns	2.40E-4
eDRAM	41.1	1.56ns	1.66E-1
Bus	7	5 cycles	9.00E-3
Router	42	2 cycles	1.50E-2
DAC [118]	12.5	0.78ns	2.50E-3
ADC(1 GS/s) [178]	2.55	0.78ns	2E-3
ADC(5 GS/s) [229]	11	0.78ns	21E-3
ADC(10 GS/s) [93]	30	0.78ns	103E-3
EO MRM Operation	1.4 pJ/bit	-	0.95E-4

(TPA) in SiN material, enable the support of a larger TPC size ($N=47$) compared to that of SOIPhAR ($N=22$). This larger TPC size, as shown in Table 9.2, increases the size of the dot product operation N and the number of parallel dot product operations M , thereby enhancing overall throughput via improved parallelism. Secondly, a larger N results in fewer buffer accesses of weight and input values, reducing the buffer access latency. This reduction in access latency improves FPS. Furthermore, as the datarate increases, the FPS of each accelerator decreases. At 5 GS/s and 10 GS/s, the N value decreases for all accelerators, as indicated in Table 9.2, leading to low parallelism and increased buffer accesses. This increase in access latency with higher datarates results in lower FPS for the accelerators.

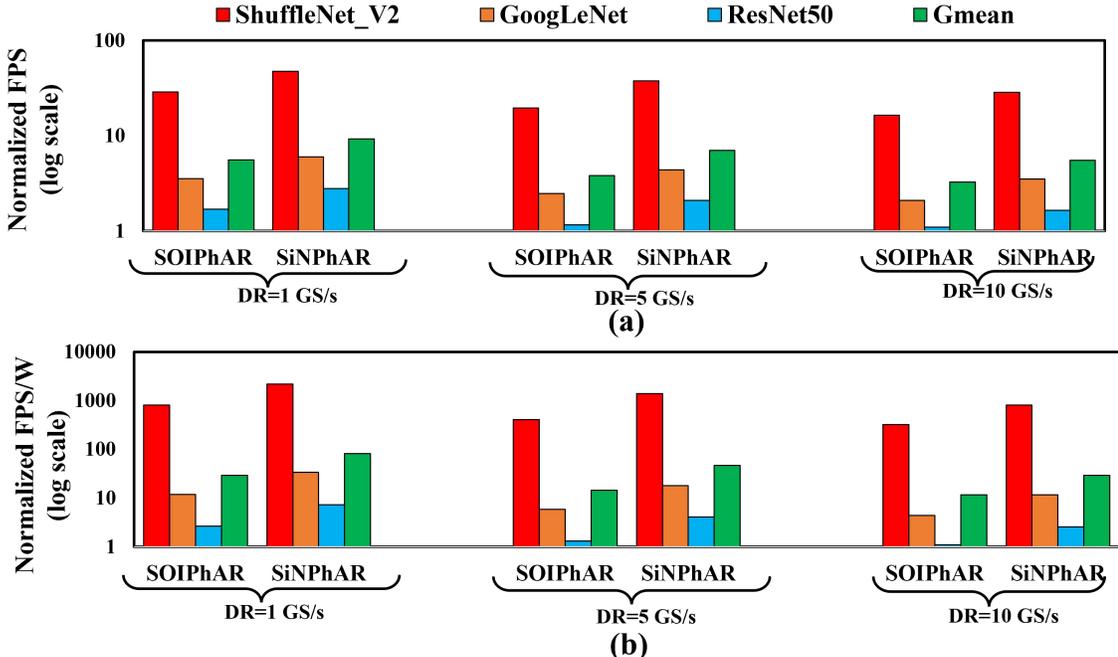


Figure 9.6: (a) Normalized FPS (log scale) (b) Normalized FPS/W (log scale) for SiN-PhAR versus SOIPhAR accelerators with input batch size=1. Results of FPS and FPS/W are normalized w.r.t. SOIPhAR ResNet50 at 10 GS/s.

In Fig. 9.6(b), the energy efficiency (FPS/W) results are presented on a log scale for both SOIPhAR and SiNPhAR accelerators, using a batch size of 1 at various datarates. These results are normalized to SOIPhAR for ResNet50 at the datarate of 10 GS/s. Notably, the SiNPhAR accelerators demonstrate superior energy efficiency compared to the SOIPhAR accelerators. Specifically, at 1 GS/s, SiNPhAR achieves $2.8\times$ better FPS/W compared to SOIPhAR, based on the Gmean across the CNNs. As the datarate increases to 5 GS/s, SiNPhAR achieves even better improvement over SOIPhAR, with $3.19\times$ better FPS/W when compared to SOIPhAR.

These energy efficiency advantages of SiNPhAR stem from several factors. First, the improved throughput and reduced energy consumption of buffer accesses contribute to enhanced energy efficiency. As discussed earlier, the higher N value supported by SiNPhAR results in improved parallelism, which, in turn, reduces dynamic

energy consumption while maintaining higher throughput. Additionally, SiNPhAR requires overall fewer buffer accesses of input and weight values, leading to energy savings by reducing the energy consumption corresponding to buffer accesses. As the datarate increases, the peripheral components of the accelerator, such as ADCs and DACs, consume more power (as indicated in Table 9.3). This additional power consumption decreases the achieved FPS/W for both SOIPhAR and SiNPhAR accelerators.

9.5 Summary

In this chapter, we presented a novel SiN-based photonic GEMM Accelerator called SiNPhAR. Our SiNPhAR accelerator employs SiN-on-SiO₂ based waveguides ITO-enabled SiN-on-SiO₂-based microring modulators (MRMs) as input and weight elements, to implement analog GEMM functions. The key advantages of SiNPhAR over traditional SOI-based photonic GEMM accelerators lie in the absence of Two-Photon Absorption (TPA) nonlinearity and the low index contrast of the SiN-on-SiO₂ devices. These features enable SiNPhAR to experience significantly low optical signal losses compared to traditional SOI-based photonic GEMM accelerators, substantially enhancing its parallelism, throughput, and energy efficiency. To validate these benefits of our SiNPhAR accelerator, we evaluated its achievable parallelism and performance and compared it with a traditional SOI-based GEMM accelerator from prior work. Our analysis reveals that SiNPhAR supports at least 1.5× more multipliers than the prior SOI-based photonic GEMM accelerator. Furthermore, from the system-level performance analysis, SiNPhAR demonstrates at least 1.7× better throughput (FPS) while consuming at least 2.8× better energy efficiency (FPS/W) compared to the prior SOI-based GEMM accelerator.

Chapter 10 Conclusions and Future Work

10.1 Conclusions

In this report, we presented several solutions to address various design challenges encountered by silicon photonic interconnects and silicon photonic-based E-O computing circuits. A recap of each of our contributions is discussed in the upcoming subsections.

Silicon Photonic Interconnects

In our first contribution, we presented a novel design of MR quality factor array that reduces overall laser power consumption in the link. At the detector side of MR filter array, each MR experiences different crosstalk noise. As a result, MR filter array experiences non-uniformity in crosstalk penalty distribution resulting in laser power overprovisioning for each channel. By designing each MR filter in the array with different quality factor, we uniformize the crosstalk penalty distribution which in turn will reduce laser power overprovisioning per channel and overall laser power consumption in the link. From our analysis, we have observed that DWDM photonic interconnects that utilized our designed MR filter array can achieve laser power savings of up to 34 mW.

In our second contribution, we presented silicon-on-sapphire (SOS) photonic platform as a solution to break the scalability barrier of silicon-on-insulator (SOI) based interconnects. At operating wavelengths of SOI platform, the predominant non-linear effect in silicon is two-photon absorption (TPA). TPA induces free-carrier absorption (FCA) and free-carrier dispersion (FCD) effects in silicon (Si) that restricts the maximum allowable optical power in the link to be no more than 20 dBm which in turn restricts the scalability of SOI photonic links. SOS platform operates in the mid-infrared region where SOS constituent devices have shown to exhibit no TPA. This advantage of SOS platform paves the way for realizing high throughput and energy-efficient photonic interconnects. We have devised new compact models for SOS devices and have formulated new guidelines for designing SOS photonic interconnects. We performed a link-level analysis from which we have evaluated that SOS photonic interconnects can achieve aggregate data rate of 1.6 Tb/s with an energy-efficiency close to ~ 1 pJ/bit. We have also performed a system-level analysis from which we have evaluated that PNoCs that employ SOS based photonic interconnects can lower the latency and energy-per-bit by 45% and 37% respectively compared to SOI based PNoCs.

In our third contribution, we conducted a comprehensive analysis of various designs of 4-PAM modulators. Initially, we presented a comparative study of various OOK and 4-PAM modulators, evaluating their performance, hardware requirements, and energy efficiency. Subsequently, we employed a heuristic-based search to optimize 4-PAM photonic link designs, targeting different goals, including BER-balanced

bitrate and optimal BER. We then compared these optimized designs of 4-PAM photonic links and architectures to conventional OOK-based photonic links and architectures in terms of aggregate data rate, latency, BER, and energy efficiency. For both design goals, 4-PAM EDAC modulator-based photonic links and architectures demonstrated lower latency and improved energy efficiency compared to conventional OOK and other 4-PAM modulator-based interconnects and architectures.

In our fourth contribution, we identified several design pathways that can aid on-silicon photonic interposer (On-SiPhI) inter-chiplet interconnects to meet the goal of achieving multi-Tb/s bandwidth. Based on the identified design pathways and three different photonic fabrication platforms, namely 45nm SOI CMOS, 32nm SOI CMOS and deposited poly-Si, we derived various design variants of on-SiPhI inter-chiplet interconnects. Subsequently, we conducted an extensive link-level and system-level analysis for each variant. From the link-level analysis, we concluded that design pathways simultaneously enhancing the spectral range and optical power budget available for wavelength division multiplexing provide significant impetus to on-SiPhI inter-chiplet links. These enhancements enable the achievement of an aggregate bandwidth exceeding 4Tb/s and support link lengths of up to 10cm. Leveraging this link-level analysis, we performed a system-level analysis on state-of-the-art CPU and GPU-based System-in-Packages (SiPs), incorporating multi-Tb/s on-SiPhI inter-chiplet links. For CPU-based SiPs, design pathways enhancing spectral range and optical power budget demonstrated at least 25% better performance while consuming at least 5% less energy on average compared to other design pathways. Similarly, for GPU-based SiPs, these design pathways accelerated the training time of large-scale Deep Neural Network models by at least 15 \times on average compared to alternative design pathways. These results suggest that enhancing the spectral multiplexing range and optical power budget of on-SiPhI interconnects concurrently could pave the way for achieving multi-Terabits/second performance in the future.

Silicon Photonic-based Electro-Optic (E-O) Computing Circuits

In our fifth contribution, we presented a novel Microring Resonator based Polymorphic Electro-Optic Logic Gate (MRR-PEOLG) that can be dynamically reconfigured to implement different logic functions at different times. We modeled the MRR-PEOLG using photonics foundry-validated simulation tools from ANSYS/Lumerical. Employing these tools, we conducted frequency-domain, time-domain transient, and performance analyses of the MRR-PEOLG. Our analysis confirmed that the MRR-PEOLG design can effectively implement various logic functions while operating at speeds of up to 40 Gb/s. Evaluation results indicate that incorporating our MRR-PEOLG into two E-O circuits from previous works can reduce their area-energy-delay product by up to 82.6 \times .

In our sixth contribution, we presented a novel hybrid Time-Amplitude Analog Optical Modulator (TAOM) utilizing a single microring modulator (MRM) to generate a high-speed optical signal represented as a sequence of Pulse-Width-Amplitude-Modulated (PWAM) symbols. Each symbol denotes the analog product of an input and a weight value. Integrated with a balanced photo charge accumulator (BPCA),

the TAOM leverages in-situ charge accumulation and incoherent superposition abilities of photodetectors to generate a signed summation of a large number of temporally and spatially arriving PWAM symbols. We organized these TAOMs and BP-CAs in 2D arrays to design an SiPh GEMM accelerator, termed TAOM-Tensor Core (TAOM-TC). Our extensive analysis covered device-level, circuit-level, and system-level evaluations of the TAOM-TC. At the device level, we observed improved TAOM accuracy with higher input optical power and increased bit resolution. Additionally, TAOM precision increased with higher input optical pulse amplitude, bit resolution, and pulse width. Circuit-level simulations revealed changes in inter-modulation (IM) crosstalk and its impact on the multiplication results of each TAOM for various channel spacings. The analysis showed that an increase in IM-crosstalk resulted in a greater deviation from the target multiplication results for each TAOM. However, compared to conventional photonic multiplier circuits, our TAOM-enabled parallel multiplier circuit experienced reduced IM-crosstalk and consequently, less error in the multiplication results. At the system level, we evaluated the scalability, power consumption, performance, and inference accuracy of the TAOM-TC. We compared it with two well-known MRR-based GEMM accelerators from prior works. From the scalability and power analysis, we found that the TAOM-TC supports at least $1.5\times$ more TAOMs per waveguide while consuming approximately $1.5\times$ less power compared to the considered MRR-enabled SiPh GEMM accelerators from prior works. Furthermore, from the performance analysis, the TAOM-TC demonstrated convolution speeds approximately $10\times$ faster than GPUs and up to approximately $3\times$ faster than the considered MRR-based GEMM accelerators from previous works. Additionally, the TAOM-TC outperformed the prior MRR-enabled SiPh GEMM accelerators in terms of accuracy by 0.4%.

In our seventh contribution, we demonstrated ITO based SiN-on-SiO₂ MRR modulators, which consists of a stack of ITO-SiN-ITO and ITO-SiO₂-ITO thin films as the active upper cladding of the SiN MRR core, respectively. This active upper cladding of our modulators leverage the free-carrier assisted, high-amplitude refractive index change in the ITO films to effect a large electro-refractive optical modulation in the device. To evaluate the performance of our SiN-on-SiO₂ MRR modulators, we performed electrostatic, transient and finite difference time domain (FDTD) simulations using the foundry-validated Ansys/Lumerical tools. Based on these simulations, our modulators achieve superior performance that demonstrates their potential to enhance the performance and energy-efficiency of SiN-on-SiO₂ based PICs of the future.

In our eighth contribution, we introduced a novel Silicon Nitride (SiN)-based photonic General Matrix-Matrix Multiplication (GEMM) Accelerator named SiNPhAR. This accelerator employs Indium Tin Oxide (ITO)-enabled SiN-on-SiO₂-based microring modulators (MRMs), demonstrated in chapter 8 as input and weight elements to implement analog GEMM functions. The SiNPhAR’s design takes advantage of the absence of Two-Photon Absorption (TPA) nonlinearity and the low index contrast of SiN-on-SiO₂ devices, resulting in significantly lower optical signal losses compared to traditional Silicon-on-Insulator (SOI)-based photonic GEMM accelerators. This characteristic substantially enhances SiNPhAR’s spatial parallelism, throughput, and energy efficiency. To validate these advantages, we conducted an evaluation of SiN-

PhAR, assessing its achievable spatial parallelism and performance. We compared SiNPhAR with a traditional SOI-based GEMM accelerator from prior work. Our analysis demonstrated that SiNPhAR supports at least $1.5\times$ more multipliers than the prior SOI-based photonic GEMM accelerator. Furthermore, from the system-level performance analysis, SiNPhAR exhibited at least $1.7\times$ better throughput (Frames Per Second, FPS) while consuming at least $2.8\times$ better energy efficiency (FPS/W) compared to the prior SOI-based GEMM accelerator.

10.2 Future Work

As the technology keeps scaling, silicon photonic interconnects and silicon photonic-based E-O computing circuits will continue to face new design challenges. Taking this into consideration, we provide the following directions for future research.

Non-Uniformity in Crosstalk Distribution

In Chapter 2, we performed an analysis of two techniques: reshuffling the resonance wavelengths (as considered in prior work [18]) and our novel non-uniform quality factor-based MR array technique. The link-level analysis parameters for both techniques included a channel spacing of 50 GHz and a bit rate of 25 Gb/s, with a quality factor of 8000 considered for the reshuffled case. For future work, we plan to enhance our link-level analysis by exploring different values of channel spacing (ranging from 50 GHz to 120 GHz) and bit rates (ranging from 5 Gb/s to 40 Gb/s). Additionally, for the reshuffled case, we intend to consider quality factors ranging from 5000 to 12000.

Our future work involves devising and analyzing two new techniques: non-uniform channel spacing between resonance wavelengths of the ring resonators in an array and the addition of dummy ring resonators at both ends of the array. For the non-uniform channel spacing technique, we plan to perform an exhaustive search to find an appropriate channel spacing that uniformizes the crosstalk penalty distribution in the array, considering a range of channel spacings from the calculated upper limit based on the Free Spectral Range (FSR) and N_λ values to a lower limit of 0.5 nm. The dummy ring resonators technique aims to uniformize crosstalk penalty distribution by adding two dummy ring resonators at both ends of the array.

Similar to the techniques analyzed in Chapter 2, our plan includes an extensive link-level analysis of these new techniques, considering various bit rates, quality factors, and channel spacings. We also intend to implement combinations of these methods to maximize energy benefits at the link level. Finally, based on the insights gained from the link-level analysis, we aim to implement these techniques at the architecture or system level on well-known Photonic Network-on-Chip (PNoC) architectures [116, 56] to evaluate their performance and energy benefits at a higher level of abstraction.

Comprehensive Investigation of Maximum Allowable Optical Power in Silicon-on-Insulator (SOI)-Based Photonic Interconnects

Prior works have experimentally demonstrated that the maximum allowable optical power per wavelength channel and per waveguide in traditional Silicon-on-Insulator (SOI)-based photonic interconnects is primarily limited by the optical non-linear effects in silicon such as two-photon absorption, especially at elevated optical power levels [147]. However, these experimentally determined MAOP limits are typically derived under conditions involving varying optical power with fixed parameters such as operating wavelength, modulation bias, quality factor, and data rate [147]. Therefore, for a more comprehensive understanding, it becomes essential to conduct a thorough exploration to ascertain the MAOP in SOI-based devices, considering a range of device-level parameters. These parameters include data rate per wavelength channel, quality factor of the SOI MRMs, modulation bias, the number of multiplexed wavelength channels per waveguide (N_λ), and wavelength detuning.

Reconfigurable Electro-Optic (E-O) SIMD/MIMD Processing Units

In chapter 6, we presented a novel design of Microring Resonator-Based Polymorphic Electro-Optic Logic Gate (MRR-PEOLG) that can be dynamically reconfigured to perform different logic functions at different times. We reason that it is possible to use the dense wavelength division multiplexing (DWDM) technique with our MRR-PEOLG design, where cascaded arrays of MRR-PEOLGs can couple with DWDM-enabled rectilinear waveguides. In these cascaded arrays, each MRR-PEOLG can be individually programmed to perform a specific logic-gate function. Moreover, from [274], it can be inferred that OR, XOR and AND logic-gate functions supported by our MRR-PEOLGs can be useful for realizing stochastic (unary) arithmetic functions such as addition, subtraction and multiplication respectively. This enables the application of the cascaded arrays of MRR-PEOLGs for realizing reconfigurable SIMD/MIMD E-O processing units (see Fig. 10.1). Such E-O SIMD/MIMD units can outperform the traditional GPUs [11] and Tensor Processing Units (TPUs) [271] due to their two-fold benefits. First, they can be operated at higher speeds (up to 40Gb/s) compared to GPUs/TPUs. Second, they can provide significantly better area \times latency product compared to their electronic counterparts.

Indium Tin Oxide (ITO)-Based Silicon Nitride (SiN)-on-Insulator Add-Drop Microring Weighting Element

In Chapter 8, we presented a novel SiN-based Photonic GEMM Accelerator (SiN-PhAR), leveraging Indium Tin Oxide (ITO)-enabled SiN-on-SiO₂ all-pass microring modulators (MRMs) as input and weighting elements to perform analog GEMM functions. For the future, we propose a refinement to the architecture by considering an add-drop configuration for the same MRM, that can be deployed as a weighting element. However, a pivotal focus will be to investigate the potential impact of inter-modulation (IM) crosstalk on the weighting in a WDM-based ITO-enabled SiN-on-SiO₂ add-drop weighting MRM circuit. To delve into this, a comprehensive

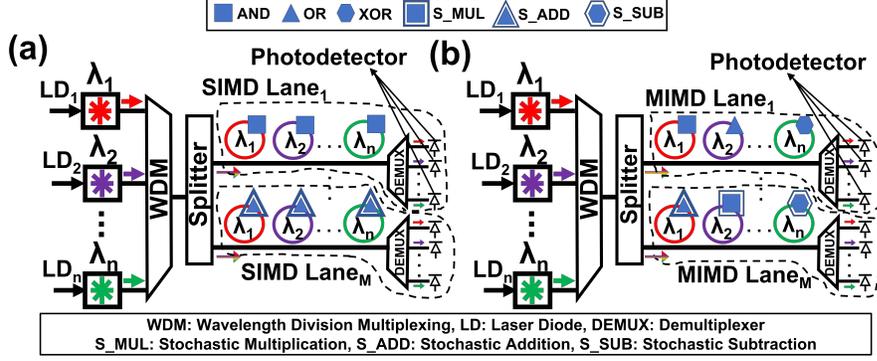


Figure 10.1: Schematics of how the cascaded arrays of our MRR-PEOLG can be reconfigured to implement (a) a SIMD or (b) an MIMD E-O processing unit. The reconfiguration between SIMD/MIMD can be achieved by programming the individual MRR-PEOLGs for specific logic/arithmetic functions.

circuit-level simulation should be performed using ANSYS/Lumerical’s INTERCONNECT tool [2], akin to the methodology outlined in Section 7.4. Moreover, for a visual representation of the IM-crosstalk influence in a cascaded ITO-enabled SiN-on-SiO₂ weighting MRM circuit, a grid plot, reminiscent of the one demonstrated in Section 7.4, would aid in understanding the IM-crosstalk induced deviation in the weighting performed by each of the MRMs in the circuit.

Zinc Oxide (ZnO)-Based Silicon Nitride (SiN)-on-Insulator Active Photonic Devices

In Chapter 8, we presented a novel Indium Tin Oxide (ITO)-based Silicon Nitride (SiN) on Silicon Dioxide (SiO₂) Microring Modulator (MRM). This device achieves intensity modulation through a free-carriers-assisted change in the permittivity and refractive index of the ITO material under the influence of an external bias [59]. Additionally, Zinc Oxide (ZnO) is identified as another promising transparent conductive oxide (TCO) that can be heterogeneously integrated with the SiN-on-SiO₂ material platform to design high-performance photonic integrated circuits. Experimental evidence from prior works indicates that ZnO exhibits free-carriers-assisted permittivity and index modulation under the influence of optical pumping power [211]. This characteristic of Zinc Oxide (ZnO) opens avenues for the design of all-optical actuation-based Microring Modulator (MRM) devices and circuits which will enable the reduction of optical-to-electrical (O/E) and electrical-to-optical (E/O) conversions, contributing to more efficient and streamlined photonic integrated circuits. By leveraging ZnO’s optical pumping responsiveness, these circuits can achieve enhanced performance and energy efficiency in all-optical signal processing applications.

Appendix

- <https://github.com/praneeth248/Inter-Channel-Xtalk> (MATLAB code implementing a heuristic-based optimization to determine the maximum number of wavelength channels in a photonic link, considering insertion losses, as well as modulator and detector crosstalk penalties.)
- <https://github.com/praneeth248/MRR-PEOLG> (Tutorial on Designing a Microring Resonator-Based Polymorphic Electro-Optic Logic Gate. (Chapter 6))
- <https://github.com/praneeth248/A-Hybrid-TAOM> (ANSYS Lumerical Simulation Files for Hybrid Time-Amplitude Analog Optical Modulator Design. (Chapter 7))

Bibliography

- [1] IVC102 data sheet, product information and support — TI.com — ti.com. <https://www.ti.com/product/IVC102>. [Accessed 17-10-2023].
- [2] Ansys lumerical. <https://www.lumerical.com/learn/whitepapers/interconnect-enabling-time-and-frequency-domain-simulation-of-photonics-integrated-circuits-with-microring-modulators/>, 2023.
- [3] Lumerical inc. <https://optics.ansys.com/hc/en-us/articles/360042322794-Ring-Modulator>, 2023.
- [4] N. C. Abrams, Q. Cheng, M. Glick, M. Jezzini, P. Morrissey, P. O'Brien, and K. Bergman. Silicon photonic 2.5 d multi-chip module transceiver for high-performance data centers. *Journal of Lightwave Technology*, 38(13):3346–3357, 2020.
- [5] A. N. R. Ahmed, S. Shi, A. J. Mercante, and D. W. Prather. High-performance racetrack resonator in silicon nitride-thin film lithium niobate hybrid platform. *Optics express*, 27(21):30741–30751, 2019.
- [6] A. N. R. Ahmed, S. Shi, M. Zablocki, P. Yao, and D. W. Prather. Tunable hybrid silicon nitride and thin-film lithium niobate electro-optic microresonator. *Optics letters*, 44(3):618–621, 2019.
- [7] J. Ahn, M. Fiorentino, R. G. Beausoleil, N. Binkert, A. Davis, D. Fattal, N. P. Jouppi, M. McLaren, C. M. Santori, R. S. Schreiber, et al. Devices and architectures for photonic chip-scale integration. *Applied Physics A*, 95:989–997, 2009.
- [8] M. Al-Qadasi, L. Chrostowski, B. Shastri, and S. Shekhar. Scaling up silicon photonic-based accelerators: Challenges and opportunities. *APL Photonics*, 7(2), 2022.
- [9] K. Alexander, J. P. George, J. Verbist, K. Neyts, B. Kuyken, D. Van Thourhout, and J. Beeckman. Nanophotonic pockels modulators on a silicon nitride platform. *Nature communications*, 9(1):1–6, 2018.
- [10] A. I. Arka, S. Gopal, J. R. Doppa, D. Heo, and P. P. Pande. Making a case for partially connected 3d noc: Nfic versus tsv. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 16(4):1–17, 2020.
- [11] A. Arunkumar, E. Bolotin, B. Cho, U. Milic, E. Ebrahimi, O. Villa, A. Jaleel, C.-J. Wu, and D. Nellans. Mcm-gpu: Multi-chip-module gpus for continued performance scalability. *ACM SIGARCH Computer Architecture News*, 45(2):320–332, 2017.

- [12] A. H. Atabaki, S. Moazeni, F. Pavanello, H. Gevorgyan, J. Notaros, L. Alloatti, M. T. Wade, C. Sun, S. A. Kruger, H. Meng, et al. Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip. *Nature*, 556(7701):349–354, 2018.
- [13] M. R. Azghadi, C. Lammie, J. K. Eshraghian, M. Payvand, E. Donati, B. Linares-Barranco, and G. Indiveri. Hardware implementation of deep network accelerators towards healthcare and biomedical applications. *IEEE Transactions on Biomedical Circuits and Systems*, 14(6):1138–1159, 2020.
- [14] R. Baets, A. Z. Subramanian, S. Clemmen, B. Kuyken, P. Bienstman, N. Le Thomas, G. Roelkens, D. Van Thourhout, P. Helin, and S. Severi. Silicon photonics: Silicon nitride versus silicon-on-insulator. In *Optical Fiber Communication Conference*, pages Th3J–1. Optica Publishing Group, 2016.
- [15] M. Bahadori and K. Bergman. Low-power optical interconnects based on resonant silicon photonic devices: Recent advances and challenges. In *Proceedings of the 2018 on Great Lakes Symposium on VLSI*, pages 305–310, 2018.
- [16] M. Bahadori, A. Gazman, N. Janosik, S. Rumley, Z. Zhu, R. Polster, Q. Cheng, and K. Bergman. Thermal rectification of integrated microheaters for microring resonators in silicon photonics platform. *Journal of Lightwave Technology*, 36(3):773–788, 2017.
- [17] M. Bahadori, D. Nikolova, S. Rumley, C. P. Chen, and K. Bergman. Optimization of microring-based filters for dense wdm silicon photonic interconnects. In *2015 IEEE Optical Interconnects Conference (OI)*, pages 84–85. IEEE, 2015.
- [18] M. Bahadori, S. Rumley, H. Jayatilleka, K. Murray, N. A. Jaeger, L. Chrostowski, S. Shekhar, and K. Bergman. Crosstalk penalty in microring-based silicon photonic interconnect systems. *Journal of Lightwave Technology*, 34(17):4043–4052, 2016.
- [19] M. Bahadori, S. Rumley, D. Nikolova, and K. Bergman. Comprehensive design space exploration of silicon photonic interconnects. *Journal of Lightwave Technology*, 34(12):2975–2987, 2016.
- [20] M. Bahadori, S. Rumley, R. Polster, A. Gazman, M. Traverso, M. Webster, K. Patel, and K. Bergman. Energy-performance optimized design of silicon photonic interconnection networks for high-performance computing. In *Design, Automation & Test in Europe Conference & Exhibition (DATE), 2017*, pages 326–331. IEEE, 2017.
- [21] S. Bahirat and S. Pasricha. Meteor: Hybrid photonic ring-mesh network-on-chip for multicore architectures. *ACM Transactions on Embedded Computing Systems (TECS)*, 13(3s):1–33, 2014.

- [22] Baidu-Research. Baidu-research/deepbench: Benchmarking deep learning operations on different hardware. <https://github.com/baidu-research/DeepBench>.
- [23] L. Baischer et al. Learning on hardware: A tutorial on neural network accelerators and co-processors, 2021.
- [24] A. A. Bajwa, S. Jangam, S. Pal, N. Marathe, T. Bai, T. Fukushima, M. Goorsky, and S. S. Iyer. Heterogeneous integration at fine pitch ($\leq 10 \mu\text{m}$) using thermal compression bonding. In *2017 IEEE 67th electronic components and technology conference (ECTC)*, pages 1276–1284. IEEE, 2017.
- [25] Y. Ban, J. Verbist, M. Vanhoecke, J. Bauwelinck, P. Verheyen, S. Lardenois, M. Pantouvaki, and J. Van Campenhout. Low-voltage 60gb/s nrz and 100gb/s pam4 o-band silicon ring modulator. In *2019 IEEE Optical Interconnects Conference (OI)*, pages 1–2. IEEE, 2019.
- [26] V. Bangari, B. A. Marquez, H. Miller, A. N. Tait, M. A. Nahmias, T. F. De Lima, H.-T. Peng, P. R. Prucnal, and B. J. Shastri. Digital electronics and analog photonics for convolutional neural networks (deap-cnns). *IEEE Journal of Selected Topics in Quantum Electronics*, 26(1):1–13, 2019.
- [27] J. Basak, L. Liao, A. Liu, H. Nguyen, M. Paniccia, Y. Chetrit, and D. Rubin. High speed photonics on an soi platform. In *2008 IEEE International SOI Conference*, pages 85–86. IEEE, 2008.
- [28] J. Bashir, E. Peter, and S. R. Sarangi. A survey of on-chip optical interconnects. *ACM Computing Surveys (CSUR)*, 51(6):1–34, 2019.
- [29] J. Bashir and S. R. Sarangi. Nuplet: A photonic based multi-chip nuca architecture. In *2017 IEEE International Conference on Computer Design (ICCD)*, pages 617–624. IEEE, 2017.
- [30] J. Bashir and S. R. Sarangi. Gpuopt: Power-efficient photonic network-on-chip for a scalable gpu. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 17(1):1–26, 2020.
- [31] J. F. Bauters, M. J. Heck, D. John, D. Dai, M.-C. Tien, J. S. Barton, A. Leinse, R. G. Heideman, D. J. Blumenthal, and J. E. Bowers. Ultra-low-loss high-aspect-ratio si 3 n 4 waveguides. *Optics express*, 19(4):3163–3174, 2011.
- [32] R. G. Beausoleil. Large-scale integrated photonics for high-performance interconnects. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 7(2):1–54, 2011.
- [33] K. Bergman, L. P. Carloni, A. Biberman, J. Chan, and G. Hendry. *Photonic network-on-chip design*. Springer, 2014.

- [34] S. Bharadwaj, J. Yin, B. Beckmann, and T. Krishna. Kite: A family of heterogeneous interposer topologies enabled via accurate interconnect modeling. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2020.
- [35] A. Biberman, E. Timurdogan, W. A. Zortman, D. C. Trotter, and M. R. Watts. Adiabatic microring modulators. *Optics express*, 20(28):29223–29236, 2012.
- [36] C. Bienia, S. Kumar, J. P. Singh, and K. Li. The parsec benchmark suite: Characterization and architectural implications. In *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*, pages 72–81, 2008.
- [37] C. Bienia, S. Kumar, J. P. Singh, and K. Li. The parsec benchmark suite: Characterization and architectural implications. In *Proceedings of the 17th international conference on Parallel architectures and compilation techniques*, pages 72–81, 2008.
- [38] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi, A. Basu, J. Hestness, D. R. Hower, T. Krishna, S. Sardashti, et al. The gem5 simulator. *ACM SIGARCH computer architecture news*, 39(2):1–7, 2011.
- [39] D. J. Blumenthal, R. Heideman, D. Geuzebroek, A. Leinse, and C. Roeloffzen. Silicon nitride in silicon photonics. *Proceedings of the IEEE*, 106(12):2209–2231, 2018.
- [40] W. Bogaerts, X. Chen, M. Wang, I. Zand, H. Deng, L. Van Iseghem, A. Ribeiro, A. D. Tormo, and U. Khan. Programmable silicon photonic integrated circuits. In *2020 IEEE Photonics Conference (IPC)*, pages 1–2. IEEE, 2020.
- [41] W. Bogaerts, P. De Heyn, T. Van Vaerenbergh, K. De Vos, S. Kumar Selvaraja, T. Claes, P. Dumon, P. Bienstman, D. Van Thourhout, and R. Baets. Silicon microring resonators. *Laser & Photonics Reviews*, 6(1):47–73, 2012.
- [42] M. Borghi, D. Bazzanella, M. Mancinelli, and L. Pavesi. On the modeling of thermal and free carrier nonlinearities in silicon-on-insulator microring resonators. *Optics Express*, 29(3):4363–4377, 2021.
- [43] F. Brücknerhoff-Plückelmann et al. A large scale photonic matrix processor enabled by charge accumulation. *Nanophotonics*, 2022.
- [44] J. F. Buckwalter, X. Zheng, G. Li, K. Raj, and A. V. Krishnamoorthy. A monolithic 25-gb/s transceiver with photonic ring modulators and ge detectors in a 130-nm cmos soi process. *IEEE Journal of Solid-State Circuits*, 47(6):1309–1322, 2012.
- [45] V. Catania, A. Mineo, S. Monteleone, M. Palesi, and D. Patti. Energy efficient transceiver in wireless network on chip architectures. In *2016 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1321–1326. IEEE, 2016.

- [46] Cerebras. <https://www.cerebras.net/product-chip/>.
- [47] D. W. Chan, X. Wu, Z. Zhang, C. Lu, A. P. T. Lau, and H. K. Tsang. Ultra-wide free-spectral-range silicon microring modulator for high capacity wdm. *Journal of Lightwave Technology*, 40(24):7848–7855, 2022.
- [48] C.-H. Chen, M. A. Seyedi, M. Fiorentino, D. Livshits, A. Gubenko, S. Mikhlin, V. Mikhlin, and R. G. Beausoleil. A comb laser-driven dwdm silicon photonic transmitter based on microring modulators. *Optics express*, 23(16):21541–21548, 2015.
- [49] L. Chen, J. Chen, J. Nagy, and R. M. Reano. Highly linear ring modulator from hybrid silicon and lithium niobate. *Optics Express*, 23(10):13255–13264, 2015.
- [50] Y. Chen, T. Luo, S. Liu, S. Zhang, L. He, J. Wang, L. Li, T. Chen, Z. Xu, N. Sun, et al. Dadiannao: A machine-learning supercomputer. In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 609–622. IEEE, 2014.
- [51] J. Cheng. Cuda by example: an introduction to general-purpose gpu programming. *Scalable Computing: Practice and Experience*, 11(4):401–401, 2010.
- [52] L. Cheng, S. Mao, Z. Li, Y. Han, and H. Fu. Grating couplers on silicon photonics: Design principles, emerging trends and practical issues. *Micromachines*, 11(7):666, 2020.
- [53] Q. Cheng, J. Kwon, M. Glick, M. Bahadori, L. P. Carloni, and K. Bergman. Silicon photonics codesign for deep learning. *Proceedings of the IEEE*, 108(8):1261–1282, 2020.
- [54] Z. Cheng, X. Chen, C. Wong, K. Xu, C. K. Fung, Y. Chen, and H. K. Tsang. Mid-infrared grating couplers for silicon-on-sapphire waveguides. *IEEE Photonics Journal*, 4(1):104–113, 2011.
- [55] S. V. R. Chittamuru, S. Desai, and S. Pasricha. Reconfigurable silicon-photonic network with improved channel sharing for multicore architectures. In *Proceedings of the 25th edition on Great Lakes Symposium on VLSI*, pages 63–68, 2015.
- [56] S. V. R. Chittamuru, S. Desai, and S. Pasricha. Swiftnoc: a reconfigurable silicon-photonic network with multicast-enabled channel sharing for multicore architectures. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 13(4):1–27, 2017.
- [57] S. V. R. Chittamuru and S. Pasricha. Crosstalk mitigation for high-radix and low-diameter photonic noc architectures. *IEEE Design & Test*, 32(3):29–39, 2015.

- [58] S. V. R. Chittamuru, I. G. Thakkar, and S. Pasricha. Hydra: Heterodyne crosstalk mitigation with double microring resonators and data encoding for photonic nocs. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 26(1):168–181, 2017.
- [59] L. Chrostowski, J. Flueckiger, C. Lin, M. Hochberg, J. Pond, J. Klein, J. Ferguson, and C. Cone. Design methodologies for silicon photonic integrated circuits. In *Smart Photonic and Optoelectronic Integrated Circuits XVI*, volume 8989, pages 83–97. SPIE, 2014.
- [60] Y.-L. Chuang, C.-S. Yuan, J.-J. Chen, C.-F. Chen, C.-S. Yang, W.-P. Changchien, C. C. Liu, and F. Lee. Unified methodology for heterogeneous integration with cowos technology. In *2013 IEEE 63rd Electronic Components and Technology Conference*, pages 852–859. IEEE, 2013.
- [61] E. Chung, J. Fowers, K. Ovtcharov, M. Papamichael, A. Caulfield, T. Massengil, M. Liu, D. Lo, S. Alkalay, M. Haselman, et al. Accelerating persistent neural networks at datacenter scale. In *Hot Chips*, volume 29, 2017.
- [62] I. Corporation. Architecture day 2020. <https://newsroom.intel.com/wp-content/uploads/sites/11/2020/08/Intel-Architecture-Day-2020-Presentation-Slides.pdf>, 2020.
- [63] S. R. Corporation. The decadal plan for semiconductors. <https://www.src.org/about/decadal-plan/>, 2020.
- [64] W. J. Dally and B. Towles. Route packets, not wires: on-chip interconnection networks. In *Proceedings of the 38th annual design automation conference*, pages 684–689, 2001.
- [65] DARPA. Pipes. <https://s3-us-west-2.amazonaws.com/instrument1/grantsgov/310031.pdf>, 2018.
- [66] S. Daudlin, A. Rizzo, N. C. Abrams, S. Lee, D. Khilwani, V. Murthy, J. Robinson, T. Collier, A. Molnar, and K. Bergman. 3d-integrated multichip module transceiver for terabit-scale dwdm interconnects. In *Optical Fiber Communication Conference*, pages Th4A–4. Optical Society of America, 2021.
- [67] M. De Cea, A. H. Atabaki, and R. J. Ram. Power handling of silicon microring modulators. *Optics express*, 27(17):24274–24285, 2019.
- [68] L. De Marinis, M. Cococcioni, P. Castoldi, and N. Andriolli. Photonic neural networks: A survey. *IEEE Access*, 7:175827–175841, 2019.
- [69] C. Demirkiran, F. Eris, G. Wang, J. Elmhurst, N. Moore, N. C. Harris, A. Basumallik, V. J. Reddi, A. Joshi, and D. Bunandar. An electro-photonic system for accelerating deep neural networks. *arXiv preprint arXiv:2109.01126*, 2021.

- [70] L. Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012.
- [71] Q. Deng, L. Liu, X. Li, and Z. Zhou. Arbitrary-ratio 1×2 power splitter based on asymmetric multimode interference. *Optics letters*, 39(19):5590–5593, 2014.
- [72] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [73] B. Dong, X. Guo, C. P. Ho, B. Li, H. Wang, C. Lee, X. Luo, and G.-Q. Lo. Silicon-on-insulator waveguide devices for broadband mid-infrared photonics. *IEEE Photonics Journal*, 9(3):1–10, 2017.
- [74] P. Dong, S. Liao, D. Feng, H. Liang, D. Zheng, R. Shafiiha, C.-C. Kung, W. Qian, G. Li, X. Zheng, et al. Low v pp, ultralow-energy, compact, high-speed silicon electro-optic modulator. *Optics express*, 17(25):22484–22490, 2009.
- [75] P. Dong, W. Qian, S. Liao, H. Liang, C.-C. Kung, N.-N. Feng, R. Shafiiha, J. Fong, D. Feng, A. V. Krishnamoorthy, et al. Low loss shallow-ridge silicon waveguides. *Optics express*, 18(14):14474–14479, 2010.
- [76] P. Dong, R. Shafiiha, S. Liao, H. Liang, N.-N. Feng, D. Feng, G. Li, X. Zheng, A. V. Krishnamoorthy, and M. Asghari. Wavelength-tunable silicon microring modulator. *Optics express*, 18(11):10941–10946, 2010.
- [77] R. Dubé-Demers, S. LaRochelle, and W. Shi. Low-power dac-less pam-4 transmitter using a cascaded microring modulator. *Opt. Lett.*, 41(22):5369–5372, Nov 2016.
- [78] R. Dubé-Demers, S. LaRochelle, and W. Shi. On-chip multi-level signal generation using cascaded microring modulator. In *2016 IEEE Optical Interconnects Conference (OI)*, pages 28–29. IEEE, 2016.
- [79] O. Dubray, A. Abraham, K. Hassan, S. Olivier, D. Marris-Morini, L. Vivien, I. O’Connor, and S. Menezo. Electro-optical ring modulator: An ultracompact model for the comparison and optimization of pn, pin, and capacitive junction. *IEEE Journal of Selected Topics in Quantum Electronics*, 22(6):89–98, 2016.
- [80] N. Eid, R. Boeck, H. Jayatileka, L. Chrostowski, W. Shi, and N. A. Jaeger. Fsr-free silicon-on-insulator microring resonator based filter with bent contra-directional couplers. *Optics express*, 24(25):29009–29021, 2016.
- [81] C. Feng, Z. Ying, Z. Zhao, J. Gu, D. Z. Pan, and R. T. Chen. Toward high-speed and energy-efficient computing: A wdm-based scalable on-chip silicon integrated optical comparator. *Laser & Photonics Reviews*, 15(8):2000275, 2021.
- [82] T. Ferreira de Lima, E. A. Doris, S. Bilodeau, W. Zhang, A. Jha, H.-T. Peng, E. C. Blow, C. Huang, A. N. Tait, B. J. Shastri, et al. Design automation of photonic resonator weights. *Nanophotonics*, 11(17):3805–3822, 2022.

- [83] J. Fowers, K. Ovtcharov, M. Papamichael, T. Massengill, M. Liu, D. Lo, S. Alkalay, M. Haselman, L. Adams, M. Ghandi, et al. A configurable cloud-scale dnn processor for real-time ai. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, pages 1–14. IEEE, 2018.
- [84] Y. Frans, J. Shin, L. Zhou, P. Upadhyaya, J. Im, V. Kireev, M. Elzeftawi, H. Hedayati, T. Pham, S. Asuncion, et al. A 56-gb/s pam4 wireline transceiver using a 32-way time-interleaved sar adc in 16-nm finfet. *IEEE journal of solid-state circuits*, 52(4):1101–1110, 2017.
- [85] A. L. Gaeta, M. Lipson, and T. J. Kippenberg. Photonic-chip-based frequency combs. *nature photonics*, 13(3):158–169, 2019.
- [86] A. Ganguly, S. Abadal, I. Thakkar, N. E. Jerger, M. Riedel, M. Babaie, R. Balasubramonian, A. Sebastian, S. Pasricha, and B. Taskin. Interconnects for dna, quantum, in-memory, and optical computing: Insights from a panel discussion. *IEEE micro*, 42(3):40–49, 2022.
- [87] G. Gao, D. Chen, S. Tao, Y. Zhang, S. Zhu, X. Xiao, and J. Xia. Silicon nitride o-band (de) multiplexers with low thermal sensitivity. *Optics Express*, 25(11):12260–12267, 2017.
- [88] D. Geuzebroek, E. Klein, H. Kelderman, and A. Driessen. Wavelength tuning and switching of a thermo-optic microring resonator. In *proc. ECIO*, volume 395, 2003.
- [89] J. Goyvaerts, A. Grabowski, J. Gustavsson, S. Kumari, A. Stassen, R. Baets, A. Larsson, and G. Roelkens. Enabling vcsel-on-silicon nitride photonic integrated circuits with micro-transfer-printing. *Optica*, 8(12):1573–1580, 2021.
- [90] G. Griffel. Vernier effect in asymmetrical ring resonator arrays. *IEEE Photonics Technology Letters*, 12(12):1642–1644, 2000.
- [91] H. Gu, Z. Wang, B. Zhang, Y. Yang, and K. Wang. Time-division-multiplexing–wavelength-division-multiplexing-based architecture for onoc. *Journal of Optical Communications and Networking*, 9(5):351–363, 2017.
- [92] J. Gu et al. Squeezelight: Towards scalable optical neural networks with multi-operand ring resonators. In *DATE*, 2021.
- [93] M. Guo et al. A 29mw 5gs/s time-interleaved sar adc achieving 48.5db snr with fully-digital timing-skew calibration based on digital-mixing. In *VLSIC*, 2019.
- [94] L. Gwennap. Graphcore makes big ai splash. *Microprocessor Rep., The Linley Group, Mountain View, CA, USA*, 2018.
- [95] J. Hardy and J. Shamir. Optics inspired logic architecture. *Optics Express*, 15(1):150–165, 2007.

- [96] N. C. Harris, J. Carolan, D. Bunandar, M. Prabhu, M. Hochberg, T. Baehr-Jones, M. L. Fanto, A. M. Smith, C. C. Tison, P. M. Alsing, et al. Linear programmable nanophotonic processors. *Optica*, 5(12):1623–1631, 2018.
- [97] A. He, X. Guo, T. Wang, and Y. Su. Ultracompact fiber-to-chip metamaterial edge coupler. *ACS Photonics*, 8(11):3226–3233, 2021.
- [98] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [99] Y. He, T. N. Sainath, R. Prabhavalkar, I. McGraw, R. Alvarez, D. Zhao, D. Rybach, A. Kannan, Y. Wu, R. Pang, et al. Streaming end-to-end speech recognition for mobile devices. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6381–6385. IEEE, 2019.
- [100] R. Hendry, D. Nikolova, S. Rumley, N. Ophir, and K. Bergman. Physical layer analysis and modeling of silicon photonic wdm bus architectures. In *Proc. HiPEAC Workshop*, pages 20–22, 2014.
- [101] A. Hermans, M. Van Daele, J. Dendooven, S. Clemmen, C. Detavernier, and R. Baets. Integrated silicon nitride electro-optic modulators with atomic layer deposited overlays. *Optics letters*, 44(5):1112–1115, 2019.
- [102] R. Ho, K. W. Mai, and M. A. Horowitz. The future of wires. *Proceedings of the IEEE*, 89(4):490–504, 2001.
- [103] C.-Y. Hsu, G.-Z. Yiu, and Y.-C. Chang. Free-space applications of silicon photonics: A review. *Micromachines*, 13(7):990, 2022.
- [104] J. Hu. System level co-optimizations of 2.5 d/3d hybrid integration for high performance computing system. In *Semicon West*, volume 2016, 2016.
- [105] T. Hu, B. Dong, X. Luo, T.-Y. Liow, J. Song, C. Lee, and G.-Q. Lo. Silicon photonic platforms for mid-infrared applications. *Photonics Research*, 5(5):417–430, 2017.
- [106] Y. Hu, Z. Yang, N. Chen, H. Hu, B. Zhang, H. Yang, X. Lu, X. Zhang, and J. Xu. 3×40 gbit/s all-optical logic operation based on low-loss triple-mode silicon waveguide. *Micromachines*, 13(1):90, 2022.
- [107] S. T. Ilie, J. Faneca, I. Zeimpekis, T. D. Bucio, K. Grabska, D. W. Hewak, H. M. Chong, and F. Y. Gardes. Thermo-optic tuning of silicon nitride microring resonators with low loss non-volatile sb 2 s 3 phase change material. *Scientific Reports*, 12(1):17815, 2022.
- [108] Intel. <https://www.digikey.com/htmldatasheets/production/2421711/0/0/1/stratix-10-gx-sx-device-overview.html>.

- [109] Y. Ishikawa, J. Osaka, and K. Wada. Germanium photodetectors in silicon photonics. In *2009 IEEE LEOS Annual Meeting Conference Proceedings*, pages 367–368. IEEE, 2009.
- [110] S. S. Iyer. Heterogeneous integration for performance and scaling. *IEEE Transactions on Components, Packaging and Manufacturing Technology*, 6(7):973–982, 2016.
- [111] E. Jaberansary, T. M. B. Masaud, M. Milosevic, M. Nedeljkovic, G. Z. Mashanovich, and H. M. Chong. Scattering loss estimation using 2-d fourier analysis and modeling of sidewall roughness on optical waveguides. *IEEE Photonics Journal*, 5(3):6601010–6601010, 2013.
- [112] N. B. Jadhav, R. Bhagat, S. Paranjpe, S. Dahitule, S. Madke, and S. Jadhav. Micro-ring resonator based all-optical arithmetic and logical unit. *Optik*, 244:167622, 2021.
- [113] S. Jangam, S. Pal, A. Bajwa, S. Pamarti, P. Gupta, and S. S. Iyer. Latency, bandwidth and power benefits of the superchips integration scheme. In *2017 IEEE 67th Electronic Components and Technology Conference (ECTC)*, pages 86–94. IEEE, 2017.
- [114] N. E. Jerger, A. Kannan, Z. Li, and G. H. Loh. Noc architectures for silicon interposer systems: Why pay for more wires when you can get them (from your interposer) for free? In *2014 47th Annual IEEE/ACM International Symposium on Microarchitecture*, pages 458–470. IEEE, 2014.
- [115] W. Jin, R. G. Polcawich, P. A. Morton, and J. E. Bowers. Piezoelectrically tuned silicon nitride ring resonator. *Optics Express*, 26(3):3174–3187, 2018.
- [116] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic. Silicon-photonic cros networks for global on-chip communication. In *2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*, pages 124–133. IEEE, 2009.
- [117] N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, et al. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th annual international symposium on computer architecture*, pages 1–12, 2017.
- [118] F. Juanda, W. Shu, and J. S. Chang. A 10-gs/s 4-bit single-core digital-to-analog converter for cognitive ultrawidebands. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 64(1):16–20, 2016.
- [119] A. Kannan, N. E. Jerger, and G. H. Loh. Enabling interposer-based disintegration of multi-core processors. In *2015 48th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pages 546–558. IEEE, 2015.

- [120] T.-J. Kao and A. Louri. Optical multilevel signaling for high bandwidth and power-efficient on-chip interconnects. *IEEE Photonics Technology Letters*, 27(19):2051–2054, 2015.
- [121] T.-J. Kao and A. Louri. Design of high bandwidth photonic noc architectures using optical multilevel signaling. In *2016 Tenth IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pages 1–4. IEEE, 2016.
- [122] V. S. P. Karempudi, J. Bashir, and I. G. Thakkar. An analysis of various design pathways towards multi-terabit photonic on-interposer interconnects. *arXiv preprint arXiv:2306.07241*, 2023.
- [123] V. S. P. Karempudi, S. Datta, and I. G. Thakkar. Design exploration and scalability analysis of a cmos-integrated, polymorphic, nanophotonic arithmetic-logic unit. In *Proceedings of the 19th ACM Conference on Embedded Networked Sensor Systems*, pages 628–634, 2021.
- [124] V. S. P. Karempudi, S. Sri Vatsavayi, and I. Thakkar. Redesigning photonic interconnects with silicon-on-sapphire device platform for ultra-low-energy on-chip communication. In *Proceedings of the 2020 on Great Lakes Symposium on VLSI*, pages 247–252, 2020.
- [125] V. S. P. Karempudi, F. Sunny, I. G. Thakkar, S. V. R. Chittamuru, M. Nikdast, and S. Pasricha. Photonic networks-on-chip employing multilevel signaling: A cross-layer comparative study. *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 18(3):1–36, 2022.
- [126] V. S. P. Karempudi, I. G. Thakkar, and J. T. Hastings. A silicon nitride microring based high-speed, tuning-efficient, electro-refractive modulator. In *2022 IEEE International Symposium on Smart Electronic Systems (iSES)*, pages 307–311. IEEE, 2022.
- [127] M. Khani, M. Ghobadi, M. Alizadeh, Z. Zhu, M. Glick, K. Bergman, A. Vahdat, B. Klenk, and E. Ebrahimi. Sip-ml: high-bandwidth optical network interconnects for machine learning training. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, pages 657–675, 2021.
- [128] A. Khilo, S. J. Spector, M. E. Grein, A. H. Nejadmalayeri, C. W. Holzwarth, M. Y. Sander, M. S. Dahlem, M. Y. Peng, M. W. Geis, N. A. DiLello, et al. Photonic adc: overcoming the bottleneck of electronic jitter. *Optics Express*, 20(4):4454–4469, 2012.
- [129] B. Y. Kim, Y. Okawachi, J. K. Jang, M. Yu, X. Ji, Y. Zhao, C. Joshi, M. Lipson, and A. L. Gaeta. Turn-key, high-efficiency kerr comb source. *Optics letters*, 44(18):4475–4478, 2019.
- [130] H.-U. Kim and J.-K. Kang. High-speed serial interface using pwam signaling scheme. In *2022 19th International SoC Design Conference (ISOCC)*, pages 255–256. IEEE, 2022.

- [131] J. Kischkat, S. Peters, B. Gruska, M. Semtsiv, M. Chashnikova, M. Klinkmüller, O. Fedosenko, S. Machulik, A. Aleksandrova, G. Monastyrskiy, et al. Mid-infrared optical properties of thin films of aluminum oxide, titanium dioxide, silicon dioxide, aluminum nitride, and silicon nitride. *Applied optics*, 51(28):6789–6798, 2012.
- [132] C.-L. Lai, H.-Y. Li, A. Chen, and T. Lu. Silicon interposer warpage study for 2.5 d ic without tsv utilizing glass carrier cte and passivation thickness tuning. In *2016 IEEE 66th Electronic Components and Technology Conference (ECTC)*, pages 310–315. IEEE, 2016.
- [133] K. R. Lakshmikumar, A. Kurylak, M. Nagaraju, R. Booth, R. K. Nandwana, J. Pampanin, and V. Bocuzzi. A process and temperature insensitive cmos linear tia for 100gbps/ λ pam-4 optical links. *IEEE Journal of Solid-State Circuits*, 54(11):3180–3190, 2019.
- [134] Y. LeCun, Y. Bengio, and G. Hinton. Deep learning. *nature*, 521(7553):436–444, 2015.
- [135] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [136] B. G. Lee, X. Chen, A. Biberman, X. Liu, I.-W. Hsieh, C.-Y. Chou, J. I. Dadap, F. Xia, W. M. Green, L. Sekaric, et al. Ultrahigh-bandwidth silicon photonic nanowire waveguides for on-chip networks. *IEEE Photonics Technology Letters*, 20(6):398–400, 2008.
- [137] Y. S. Lee, G.-D. Kim, W.-J. Kim, S.-S. Lee, W.-G. Lee, and W. H. Steier. Hybrid si-linbo 3 microring electro-optically tunable resonators for active photonic devices. *Optics letters*, 36(7):1119–1121, 2011.
- [138] T. Lengyel. *Short-range optical communications using 4-PAM*. Chalmers Tekniska Hogskola (Sweden), 2017.
- [139] J. Levy. Integrated nonlinear optics in silicon nitride waveguides and resonators. *Cornell Theses and Dissertations*, 2011.
- [140] A. Li and W. Bogaerts. A simple and novel method to obtain an fsr free silicon ring resonator. In *Silicon Photonics and Photonic Integrated Circuits V*, volume 9891, page 989115. International Society for Optics and Photonics, 2016.
- [141] C. Li, M. Browning, P. V. Gratz, and S. Palermo. Luminoc: A power-efficient, high-performance, photonic network-on-chip. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 33(6):826–838, 2014.
- [142] E. Li, B. A. Nia, B. Zhou, and A. X. Wang. Silicon microring modulator with transparent conductive oxide gate. In *2019 IEEE Optical Interconnects Conference (OI)*, pages 1–2. IEEE, 2019.

- [143] E. Li and A. X. Wang. Theoretical analysis of energy efficiency and bandwidth limit of silicon photonic modulators. *Journal of Lightwave technology*, 37(23):5801–5813, 2019.
- [144] F. Li, S. D. Jackson, C. Grillet, E. Magi, D. Hudson, S. J. Madden, Y. Moghe, C. O’Brien, A. Read, S. G. Duvall, et al. Low propagation loss silicon-on-sapphire waveguides for the mid-infrared. *Optics express*, 19(16):15212–15220, 2011.
- [145] H. Li, G. Balamurugan, M. Sakib, R. Kumar, H. Jayatilleka, H. Rong, J. Jaussi, and B. Casper. 12.1 a 3d-integrated microring-based 112gb/s pam-4 silicon-photonic transmitter with integrated nonlinear equalization and thermal control. In *2020 IEEE International Solid-State Circuits Conference-(ISSCC)*, pages 208–210. IEEE, 2020.
- [146] H. Li, G. Balamurugan, M. Sakib, J. Sun, J. Driscoll, R. Kumar, H. Jayatilleka, H. Rong, J. Jaussi, and B. Casper. A 112 gb/s pam4 silicon photonics transmitter with microring modulator and cmos driver. *Journal of Lightwave Technology*, 38(1):131–138, 2020.
- [147] Q. Li, N. Ophir, L. Xu, K. Padmaraju, L. Chen, M. Lipson, and K. Bergman. Experimental characterization of the optical-power upper bound in a silicon microring modulator. In *2012 Optical Interconnects Conference*, pages 38–39. IEEE, 2012.
- [148] S. Li, K. Chen, J. H. Ahn, J. B. Brockman, and N. P. Jouppi. Cactip: Architecture-level modeling for sram-based structures with advanced leakage reduction techniques. In *2011 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 694–701. IEEE, 2011.
- [149] T.-K. Liang and H. K. Tsang. Role of free carriers from two-photon absorption in raman amplification in silicon-on-insulator waveguides. *Applied physics letters*, 84(15):2745–2747, 2004.
- [150] M. Lipson. Guiding, modulating, and emitting light on silicon-challenges and opportunities. *Journal of Lightwave Technology*, 23(12):4222–4238, 2005.
- [151] J. Liu, H. Tian, E. Lucas, A. S. Raja, G. Lihachev, R. N. Wang, J. He, T. Liu, M. H. Anderson, W. Weng, et al. Monolithic piezoelectric control of soliton microcombs. *Nature*, 583(7816):385–390, 2020.
- [152] W. Liu, W. Liu, Y. Ye, Q. Lou, Y. Xie, and L. Jiang. Holylight: A nanophotonic accelerator for deep learning in data centers. In *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1483–1488. IEEE, 2019.
- [153] W. Liu, Z. Wang, X. Liu, N. Zeng, Y. Liu, and F. E. Alsaadi. A survey of deep neural network architectures and their applications. *Neurocomputing*, 234:11–26, 2017.

- [154] Y. London, T. Van Vaerenbergh, L. Ramini, A. J. Rizzo, P. Sun, G. Kurczveil, A. Seyedi, J. Rhim, M. Fiorentino, and K. Bergman. Performance requirements for terabit-class silicon photonic links based on cascaded microring resonators. *Journal of Lightwave Technology*, 38(13):3469–3477, 2020.
- [155] A. Lumerical. <http://www.lumerical.com/products>.
- [156] L.-W. Luo, G. S. Wiederhecker, K. Preston, and M. Lipson. Power insensitive silicon microring resonators. *Optics letters*, 37(4):590–592, 2012.
- [157] B. Luther-Davies, B. Kuyken, I. Yu, P. Ma, X. Gai, J. Van Campenhout, P. Verheyen, S. Madden, G. Roelkens, and R. Baets. Nonlinear absorption in silicon at mid-infrared wavelengths. In *Nonlinear Optics*, pages NF1A–3. Optica Publishing Group, 2013.
- [158] Z. Ma, Z. Li, K. Liu, C. Ye, and V. J. Sorger. Indium-tin-oxide for high-performance electro-optic modulation. *Nanophotonics*, 4(2):198–213, 2015.
- [159] R. Mahajan, R. Sankman, N. Patel, D.-W. Kim, K. Aygun, Z. Qian, Y. Mekkonen, I. Salama, S. Sharan, D. Iyengar, et al. Embedded multi-die interconnect bridge (emib)—a high density, high bandwidth packaging interconnect. In *2016 IEEE 66th Electronic Components and Technology Conference (ECTC)*, pages 557–565. IEEE, 2016.
- [160] F. Mantovani, M. Garcia-Gasulla, J. Gracia, E. Stafford, F. Banchelli, M. Josep-Fabrego, J. Criado-Ledesma, and M. Nachtmann. Performance and energy consumption of hpc workloads on a cluster based on arm thunderx2 cpu. *Future generation computer systems*, 112:800–818, 2020.
- [161] D. A. Miller. Device requirements for optical interconnects to silicon chips. *Proceedings of the IEEE*, 97(7):1166–1185, 2009.
- [162] D. A. Miller. Self-configuring universal linear optical component. *Photonics Research*, 1(1):1–15, 2013.
- [163] A. Mirza, A. Shafiee, S. Banerjee, K. Chakrabarty, S. Pasricha, and M. Nikdast. Characterization and optimization of coherent mzi-based nanophotonic neural networks under fabrication non-uniformity. *IEEE Transactions on Nanotechnology*, 21:763–771, 2022.
- [164] A. Mistry, M. Hammood, H. Shoman, S. Lin, L. Chrostowski, and N. A. Jaeger. Free-spectral-range-free microring-based coupling modulator with integrated contra-directional-couplers. In *Optical Components and Materials XVII*, volume 11276, page 1127607. International Society for Optics and Photonics, 2020.
- [165] S. Mittal. A survey on evaluating and optimizing performance of intel xeon phi. *Concurrency and Computation: Practice and Experience*, 32(19):e5742, 2020.

- [166] S. Moazeni, A. Atabaki, D. Cheian, S. Lin, R. Ram, and V. Stojanovic. Monolithic integration of o-band photonic transceivers in a “zero-change” 32nm soi cmos. In *2017 IEEE International Electron Devices Meeting (IEDM)*, pages 24–3. IEEE, 2017.
- [167] S. Moazeni, S. Lin, M. Wade, L. Alloatti, R. J. Ram, M. Popović, and V. Stojanović. A 40-gb/s pam-4 transmitter based on a ring-resonator optical dac in 45-nm soi cmos. *IEEE Journal of Solid-State Circuits*, 52(12):3503–3516, 2017.
- [168] F. Morichetti, M. Milanizadeh, M. Petrini, F. Zanetto, G. Ferrari, D. O. de Aguiar, E. Guglielmi, M. Sampietro, and A. Melloni. Polarization-transparent silicon photonic add-drop multiplexer with wideband hitless tunability. *Nature Communications*, 12(1):1–7, 2021.
- [169] R. W. Morris, A. K. Kodi, A. Louri, and R. D. Whaley. Three-dimensional stacked nanophotonic network-on-chip architecture with minimal reconfiguration. *IEEE Transactions on Computers*, 63(1):243–255, 2012.
- [170] X. Mu, S. Wu, L. Cheng, and H. Fu. Edge couplers in silicon photonic integrated circuits: A review. *Applied Sciences*, 10(4):1538, 2020.
- [171] J. Mulcahy, F. H. Peters, and X. Dai. Modulators in silicon photonics—heterogenous integration & and beyond. In *Photonics*, volume 9, page 40. MDPI, 2022.
- [172] S. Naffziger, N. Beck, T. Burd, K. Lepak, G. H. Loh, M. Subramony, and S. White. Pioneering chiplet technology and design for the amd epyc™ and ryzen™ processor families: Industrial product. In *2021 ACM/IEEE 48th Annual International Symposium on Computer Architecture (ISCA)*, pages 57–70. IEEE, 2021.
- [173] A. Narayan, Y. Thonnart, P. Vivet, A. Joshi, and A. K. Coskun. System-level evaluation of chip-scale silicon photonic networks for emerging data-intensive applications. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1444–1449. IEEE, 2020.
- [174] M. Nedeljkovic, R. Soref, and G. Z. Mashanovich. Free-carrier electrorefraction and electroabsorption modulation predictions for silicon over the 1-14 μ m infrared wavelength range. *IEEE Photonics Journal*, 3(6):1171–1180, 2011.
- [175] M. Nikdast, G. Nicolescu, J. Trajkovic, and O. Liboiron-Ladouceur. Chip-scale silicon photonic interconnects: A formal study on fabrication non-uniformity. *Journal of Lightwave Technology*, 34(16):3682–3695, 2016.
- [176] C. Nitta, M. Farrens, and V. Akella. Addressing system-level trimming issues in on-chip nanophotonic networks. In *2011 IEEE 17th International Symposium on High Performance Computer Architecture*, pages 122–131. IEEE, 2011.

- [177] C. J. Nitta, M. Farrens, and V. Akella. *On-chip photonic interconnects: A computer architect's perspective*. Springer Nature, 2022.
- [178] D.-R. Oh et al. An 8b 1gs/s 2.55mw sar-flash adc with complementary dynamic amplifiers. In *IVLSIC*, 2020.
- [179] M. B. On, H. Lu, H. Chen, R. Proietti, and S. B. Yoo. Wavelength-space domain high-throughput artificial neural networks by parallel photoelectric matrix multiplier. In *2020 Optical Fiber Communications Conference and Exhibition (OFC)*, pages 1–3. IEEE, 2020.
- [180] OpenAI. Ai and compute. <https://openai.com/blog/ai-and-compute/>, 2018.
- [181] OpenAI. Ai and compute. <https://openai.com/research/ai-and-compute>, 2023.
- [182] N. Ophir, A. Biberman, J. S. Levy, K. Padmaraju, K. J. Luke, M. Lipson, and K. Bergman. Demonstration of 1.28-tb/s transmission in next-generation nanowires for photonic networks-on-chip. In *2010 23rd Annual Meeting of the IEEE Photonics Society*, pages 560–561. IEEE, 2010.
- [183] K. Padmaraju, D. F. Logan, T. Shiraishi, J. J. Ackert, A. P. Knights, and K. Bergman. Wavelength locking and thermally stabilizing microring resonators using dithering signals. *Journal of Lightwave Technology*, 32(3):505–512, 2013.
- [184] K. Padmaraju, X. Zhu, L. Chen, M. Lipson, and K. Bergman. Intermodulation crosstalk characteristics of wdm silicon microring modulators. *IEEE Photonics Technology Letters*, 26(14):1478–1481, 2014.
- [185] S. Pal, J. Liu, I. Alam, N. Cebry, H. Suhail, S. Bu, S. S. Iyer, S. Pamarti, R. Kumar, and P. Gupta. Designing a 2048-chiplet, 14336-core waferscale processor. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pages 1183–1188. IEEE, 2021.
- [186] S. Pal, D. Petrisko, A. A. Bajwa, P. Gupta, S. S. Iyer, and R. Kumar. A case for packageless processors. In *2018 IEEE international symposium on high performance computer architecture (HPCA)*, pages 466–479. IEEE, 2018.
- [187] S. Pal, D. Petrisko, M. Tomei, P. Gupta, S. S. Iyer, and R. Kumar. Architecting waferscale processors-a gpu case study. In *2019 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 250–263. IEEE, 2019.
- [188] Y. Pan, J. Kim, and G. Memik. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *HPCA-16 2010 The Sixteenth International Symposium on High-Performance Computer Architecture*, pages 1–12. IEEE, 2010.

- [189] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary. Firefly: Illuminating future network-on-chip with nanophotonics. In *Proceedings of the 36th annual international symposium on Computer architecture*, pages 429–440, 2009.
- [190] M. Pantouvaki, P. De Heyn, M. Rakowski, P. Verheyen, B. Snyder, S. Srinivasan, H. Chen, J. De Coster, G. Lepage, P. Absil, et al. 50gb/s silicon photonics platform for short-reach optical interconnects. In *2016 Optical Fiber Communications Conference and Exhibition (OFC)*, pages 1–3. IEEE, 2016.
- [191] M. Pantouvaki, P. Verheyen, J. De Coster, G. Lepage, P. Absil, and J. Van Campenhout. 56gb/s ring modulator on a 300mm silicon photonics platform. In *2015 European Conference on Optical Communication (ECOC)*, pages 1–3. IEEE, 2015.
- [192] S. Pasricha and S. Bahirat. Opal: A multi-layer hybrid photonic noc for 3d ics. In *16th Asia and South Pacific Design Automation Conference (ASP-DAC 2011)*, pages 345–350. IEEE, 2011.
- [193] S. Pasricha and M. Nikdast. A survey of silicon photonics for energy-efficient manycore computing. *IEEE Design & Test*, 37(4):60–81, 2020.
- [194] A. Paszke et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*. 2019.
- [195] N. Peserico, B. J. Shastri, and V. J. Sorger. Integrated photonic tensor processing unit for a matrix multiply: a review. *Journal of Lightwave Technology*, 2023.
- [196] M. Petrini, M. Milanizadeh, F. Zanetto, G. Ferrari, M. Sampietro, F. Morichetti, and A. Melloni. Reconfigurable fsr-free microring resonator filter with wide hitless tunability. In *2021 IEEE Photonics Society Summer Topicals Meeting Series (SUM)*, pages 1–2. IEEE, 2021.
- [197] C. T. Phare, Y.-H. Daniel Lee, J. Cardenas, and M. Lipson. Graphene electro-optic modulator with 30 ghz bandwidth. *Nature photonics*, 9(8):511–514, 2015.
- [198] S. Pitris, M. Moralis-Pegios, T. Alexoudi, K. Fotiadis, Y. Ban, P. De Heyn, J. Van Campenhout, and N. Pleros. A 400 gb/s o-band wdm (8×50 gb/s) silicon photonic ring modulator-based transceiver. In *Optical Fiber Communication Conference*, pages M4H–3. Optica Publishing Group, 2020.
- [199] R. Polster, J. L. G. Jimenez, E. Cassan, and P. Vincent. Optimization of tia topologies in a 65nm cmos process. In *2014 Optical Interconnects Conference*, pages 117–118. IEEE, 2014.
- [200] C. Qiu, X. Ye, R. Soref, L. Yang, and Q. Xu. Demonstration of reconfigurable electro-optical logic with silicon photonic integrated circuits. *Optics letters*, 37(19):3942–3944, 2012.

- [201] A. Rahim, E. Ryckeboer, A. Z. Subramanian, S. Clemmen, B. Kuyken, A. Dhakal, A. Raza, A. Hermans, M. Muneeb, S. Dhoore, et al. Expanding the silicon photonics portfolio with silicon nitride photonic integrated circuits. *Journal of lightwave technology*, 35(4):639–649, 2017.
- [202] R. Rahim. Bit error detection and correction with hamming code algorithm. 2017.
- [203] M. Rakowski, Y. Ban, P. De Heyn, N. Pantano, B. Snyder, S. Balakrishnan, S. Van Huylenbroeck, L. Bogaerts, C. Demeurisse, F. Inoue, et al. Hybrid 14nm finfet-silicon photonics technology for low-power tb/s/mm² optical i/o. In *2018 IEEE Symposium on VLSI Technology*, pages 221–222. IEEE, 2018.
- [204] M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani. Experimental realization of any discrete unitary operator. *Physical review letters*, 73(1):58, 1994.
- [205] A. Rizzo, Y. London, G. Kurczveil, T. Van Vaerenbergh, M. Fiorentino, A. Seyed, D. Livshits, R. G. Beausoleil, and K. Bergman. Energy efficiency analysis of frequency comb sources for silicon photonic interconnects. In *2019 IEEE Optical Interconnects Conference (OI)*, pages 1–2. IEEE, 2019.
- [206] B. Romeira, J. Nieder, B. Jacob, R. Adão, F. Camarneiro, J. A. Alanis, M. Hejda, A. Hurtado, J. Lourenço, D. C. Alves, et al. Subwavelength neuromorphic nanophotonic integrated circuits for spike-based computing: challenges and prospects. *Emerging Topics in Artificial Intelligence (ETAI) 2021*, 11804:118040D, 2021.
- [207] O. Ronneberger, P. Fischer, and T. Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [208] A. Roshan-Zamir, B. Wang, S. Telaprolu, K. Yu, C. Li, M. A. Seyed, M. Fiorentino, R. Beausoleil, and S. Palermo. A 40 gb/s pam4 silicon microring resonator modulator transmitter in 65nm cmos. In *2016 IEEE Optical Interconnects Conference (OI)*, pages 8–9. IEEE, 2016.
- [209] S. Rumley, M. Bahadori, D. Nikolova, and K. Bergman. Physical layer compact models for ring resonators based dense wdm optical interconnects. In *ECOC 2016; 42nd European Conference on Optical Communication*, pages 1–3. VDE, 2016.
- [210] S. Rumley, D. Nikolova, R. Hendry, Q. Li, D. Calhoun, and K. Bergman. Silicon photonics for exascale systems. *Journal of Lightwave Technology*, 33(3):547–562, 2015.

- [211] S. Saha, A. Dutta, C. DeVault, B. T. Diroll, R. D. Schaller, Z. Kudyshev, X. Xu, A. Kildishev, V. M. Shalaev, and A. Boltasseva. Extraordinarily large permittivity modulation in zinc oxide for dynamic nanophotonics. *Materials Today*, 43:27–36, 2021.
- [212] B. Saif and T. Hatem. Low-complexity passive mixer-based uwb pulse generator with leakage compensation and spectrum tunability. In *2020 IEEE International Conference on Design & Test of Integrated Micro & Nano-Systems (DTS)*, pages 1–5. IEEE, 2020.
- [213] A. Sakai, H. Go, and T. Baba. Sharply bent optical waveguide silicon-on-insulator substrate. In *Physics and Simulation of Optoelectronic Devices IX*, volume 4283, pages 610–618. SPIE, 2001.
- [214] A. Sánchez-Postigo, R. Halir, J. G. Wangüemert-Pérez, A. Ortega-Moñux, S. Wang, M. Vachon, J. H. Schmid, D.-X. Xu, P. Cheben, and Í. Molina-Fernández. Breaking the coupling efficiency–bandwidth trade-off in surface grating couplers using zero-order radiation. *Laser & Photonics Reviews*, 15(6):2000542, 2021.
- [215] S. R. Sarangi, R. Kalayappan, P. Kallurkar, S. Goel, and E. Peter. Tejas: A java based versatile micro-architectural simulator. In *2015 25th International Workshop on Power and Timing Modeling, Optimization and Simulation (PATMOS)*, pages 47–54. IEEE, 2015.
- [216] G. Scalari, J. Faist, and N. Picqué. On-chip mid-infrared and thz frequency combs for spectroscopy. *Applied Physics Letters*, 114(15), 2019.
- [217] S. Selvaraja. *Wafer-scale fabrication technology for silicon photonic integrated circuits*. PhD thesis, 2011.
- [218] M. A. Seyedi, R. Wu, C.-H. Chen, M. Fiorentino, and R. G. Beausoleil. 15 gb/s transmission with wide-fsr carrier injection ring modulator for tb/s optical links. In *CLEO: Science and Innovations*, pages SF2F–7. Optica Publishing Group, 2016.
- [219] A. Shacham, K. Bergman, and L. P. Carloni. Photonic networks-on-chip for future generations of chip multiprocessors. *IEEE Transactions on Computers*, 57(9):1246–1260, 2008.
- [220] C. J. Shallue, J. Lee, J. Antognini, J. Sohl-Dickstein, R. Frostig, and G. E. Dahl. Measuring the effects of data parallelism on neural network training. *arXiv preprint arXiv:1811.03600*, 2018.
- [221] R. Shankar, I. Bulu, and M. Lončar. Integrated high-quality factor silicon-on-sapphire ring resonators for the mid-infrared. *Applied Physics Letters*, 102(5), 2013.

- [222] B. J. Shastri, A. N. Tait, T. Ferreira de Lima, W. H. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal. Photonics for artificial intelligence and neuromorphic computing. *Nature Photonics*, 15(2):102–114, 2021.
- [223] Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, et al. Deep learning with coherent nanophotonic circuits. *Nature photonics*, 11(7):441–446, 2017.
- [224] Y. Shen, X. Meng, Q. Cheng, S. Rumley, N. Abrams, A. Gazman, E. Manzhosov, M. S. Glick, and K. Bergman. Silicon photonics for extreme scale systems. *Journal of Lightwave Technology*, 37(2):245–259, 2019.
- [225] K. Shiflett, A. Karanth, A. Louri, and R. Bunescu. Bitwise neural network acceleration using silicon photonics. In *Proceedings of the 2021 on Great Lakes Symposium on VLSI*, pages 9–14, 2021.
- [226] K. Shiflett, D. Wright, A. Karanth, and A. Louri. Pixel: Photonic neural network accelerator. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 474–487. IEEE, 2020.
- [227] M. Shoeybi, M. Patwary, R. Puri, P. LeGresley, J. Casper, and B. Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- [228] F. Shokraneh, S. Geoffroy-Gagnon, M. S. Nezami, and O. Liboiron-Ladouceur. A single layer neural network implemented by a 4×4 mzi-based optical processor. *IEEE Photonics Journal*, 11(6):1–12, 2019.
- [229] Y.-S. Shu. A 6b 3gs/s 11mw fully dynamic flash adc in 40nm cmos with reduced number of comparators. In *VLSIC*, 2012.
- [230] A. Sludds, S. Bandyopadhyay, Z. Chen, Z. Zhong, J. Cochrane, L. Bernstein, D. Bunandar, P. B. Dixon, S. A. Hamilton, M. Streshinsky, et al. Delocalized photonic deep learning on the internet’s edge. *Science*, 378(6617):270–276, 2022.
- [231] A. Sodani, R. Gramunt, J. Corbal, H.-S. Kim, K. Vinod, S. Chinthamani, S. Hutsell, R. Agarwal, and Y.-C. Liu. Knights landing: Second-generation intel xeon phi product. *Ieee micro*, 36(2):34–46, 2016.
- [232] R. Soref and B. Bennett. Electrooptical effects in silicon. *IEEE journal of quantum electronics*, 23(1):123–129, 1987.
- [233] R. A. Soref, S. J. Emelett, and W. R. Buchwald. Silicon waveguided components for the long-wave infrared region. *Journal of Optics A: Pure and Applied Optics*, 8(10):840, 2006.
- [234] A. Spott, Y. Liu, T. Baehr-Jones, R. Ilic, and M. Hochberg. Silicon waveguides and ring resonators at $5.5 \mu\text{m}$. *Applied Physics Letters*, 97(21), 2010.

- [235] B. Stern, X. Ji, Y. Okawachi, A. L. Gaeta, and M. Lipson. Fully integrated chip platform for electrically pumped frequency comb generation. In *CLEO: Science and Innovations*, pages SM1D–6. Optical Society of America, 2018.
- [236] V. Stojanović, R. J. Ram, M. Popović, S. Lin, S. Moazeni, M. Wade, C. Sun, L. Alloatti, A. Atabaki, F. Pavanello, et al. Monolithic silicon-photonics platforms in state-of-the-art cmos soi processes. *Optics express*, 26(10):13106–13121, 2018.
- [237] D. Stow, Y. Xie, T. Siddiqua, and G. H. Loh. Cost-effective design of scalable high-performance systems using active and passive interposers. In *2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, pages 728–735. IEEE, 2017.
- [238] C. Sun, C.-H. O. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic. Dsent-a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *2012 IEEE/ACM Sixth International Symposium on Networks-on-Chip*, pages 201–210. IEEE, 2012.
- [239] C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin, et al. Single-chip microprocessor that communicates directly using light. *Nature*, 528(7583):534–538, 2015.
- [240] C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin, et al. Single-chip microprocessor that communicates directly using light. *Nature*, 528(7583):534–538, 2015.
- [241] J. Sun, R. Kumar, M. Sakib, J. B. Driscoll, H. Jayatilleka, and H. Rong. A 128 gb/s pam4 silicon microring modulator with integrated thermo-optic resonance tuning. *Journal of Lightwave Technology*, 37(1):110–115, 2018.
- [242] P. Sun, J. Hulme, T. Van Vaerenbergh, J. Rhim, C. Baudot, F. Boeuf, N. Vulliet, A. Seyedi, M. Fiorentino, and R. G. Beausoleil. Statistical behavioral models of silicon ring resonators at a commercial cmos foundry. *IEEE Journal of Selected Topics in Quantum Electronics*, 26(2):1–10, 2019.
- [243] F. Sunny, A. Mirza, M. Nikdast, and S. Pasricha. Crosslight: A cross-layer optimized silicon photonic neural network accelerator. In *2021 58th ACM/IEEE Design Automation Conference (DAC)*, pages 1069–1074. IEEE, 2021.
- [244] F. P. Sunny, E. Taheri, M. Nikdast, and S. Pasricha. A survey on silicon photonics for deep learning. *ACM Journal of Emerging Technologies in Computing System*, 17(4):1–57, 2021.

- [245] K. Szczerba, P. Westbergh, J. Karout, J. S. Gustavsson, Å. Haglund, M. Karlsson, P. A. Andrekson, E. Agrell, and A. Larsson. 4-pam for high-speed short-range optical communications. *Journal of optical communications and networking*, 4(11):885–894, 2012.
- [246] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.
- [247] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [248] I. G. Thakkar. *Design and Optimization of Emerging Interconnection and Memory Subsystems for Future Manycore Architectures*. PhD thesis, Colorado State University, 2018.
- [249] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha. A comparative analysis of front-end and back-end compatible silicon photonic on-chip interconnects. In *2016 ACM/IEEE International Workshop on System Level Interconnect Prediction (SLIP)*, pages 1–8. IEEE, 2016.
- [250] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha. Mitigation of homodyne crosstalk noise in silicon photonic noc architectures with tunable decoupling. In *Proceedings of the Eleventh IEEE/ACM/IFIP International Conference on Hardware/Software Codesign and System Synthesis*, pages 1–10, 2016.
- [251] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha. Run-time laser power management in photonic nocs with on-chip semiconductor optical amplifiers. In *2016 Tenth IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pages 1–4. IEEE, 2016.
- [252] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha. Improving the reliability and energy-efficiency of high-bandwidth photonic noc architectures with multilevel signaling. In *Proceedings of the Eleventh IEEE/ACM International Symposium on Networks-on-Chip*, pages 1–8, 2017.
- [253] I. G. Thakkar and S. Pasricha. 3d-wiz: A novel high bandwidth, optically interfaced 3d dram architecture with reduced random access time. In *2014 IEEE 32nd International Conference on Computer Design (ICCD)*, pages 1–7. IEEE, 2014.
- [254] I. G. Thakkar and S. Pasricha. 3d-prowiz: An energy-efficient and optically-interfaced 3d dram architecture with reduced data access overhead. *IEEE Transactions on Multi-Scale Computing Systems*, 1(3):168–184, 2015.

- [255] Y. Thonnart, S. Bernabé, J. Charbonnier, C. Bernard, D. Coriat, C. Fuguet, P. Tissier, B. Charbonnier, S. Malhouitre, D. Saint-Patrice, et al. Popstar: A robust modular optical noc architecture for chiplet-based 3d integrated systems. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1456–1461. IEEE, 2020.
- [256] Y. Thonnart, M. Zid, J. L. Gonzalez-Jimenez, G. Waltener, R. Polster, O. Dubray, F. Lepin, S. Bernabé, S. Menezo, G. Parès, et al. A 10gb/s siphotonic transceiver with $150\mu\text{w}$ $120\mu\text{s}$ -lock-time digitally supervised analog microring wavelength stabilization for 1tb/s/mm² die-to-die optical networks. In *2018 IEEE International Solid-State Circuits Conference-(ISSCC)*, pages 350–352. IEEE, 2018.
- [257] C. A. Thraskias, E. N. Lallas, N. Neumann, L. Schares, B. J. Offrein, R. Henker, D. Plettmeier, F. Ellinger, J. Leuthold, and I. Tomkos. Survey of photonic and plasmonic interconnect technologies for intra-datacenter and high-performance computing communications. *IEEE Communications Surveys & Tutorials*, 20(4):2758–2783, 2018.
- [258] D. Urbonas, A. Balčytis, M. Gabalis, K. Vaškevičius, G. Naujokaitė, S. Juodkazis, and R. Petruškevičius. Ultra-wide free spectral range, enhanced sensitivity, and removed mode splitting soi optical ring resonator with dispersive metal nanodisks. *Optics letters*, 40(13):2977–2980, 2015.
- [259] J. Van Campenhout, M. Pantouvaki, P. Verheyen, S. Selvaraja, G. Lepage, H. Yu, W. Lee, J. Wouters, D. Goossens, M. Moelants, et al. Low-voltage, low-loss, multi-gb/s silicon micro-ring modulator based on a mos capacitor. In *Optical Fiber Communication Conference*, pages OM2E–4. Optica Publishing Group, 2012.
- [260] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn. Corona: System implications of emerging nanophotonic technology. *ACM SIGARCH Computer Architecture News*, 36(3):153–164, 2008.
- [261] S. S. Vatsavai, V. S. P. Karempudi, and I. Thakkar. Proteus: Rule-based self-adaptation in photonic nocs for loss-aware co-management of laser power and performance. In *2020 14th IEEE/ACM International Symposium on Networks-on-Chip (NOCS)*, pages 1–8. IEEE, 2020.
- [262] S. S. Vatsavai, V. S. P. Karempudi, I. Thakkar, A. Salehi, and T. Hastings. Sconna: A stochastic computing based optical accelerator for ultra-fast, energy-efficient inference of integer-quantized cnns. *arXiv preprint arXiv:2302.07036*, 2023.
- [263] S. S. Vatsavai and I. G. Thakkar. Photonic reconfigurable accelerators for efficient inference of cnns with mixed-sized tensors. *IEEE Transactions on*

- Computer-Aided Design of Integrated Circuits and Systems*, 41(11):4337–4348, 2022.
- [264] S. Venkataramani, A. Ranjan, S. Banerjee, D. Das, S. Avancha, A. Jagannathan, A. Durg, D. Nagaraj, B. Kaul, P. Dubey, et al. Scaleddeep: A scalable compute architecture for learning and evaluating deep networks. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*, pages 13–26, 2017.
- [265] C.-C. Wang et al. 67.5-fj per access 1-kb sram using 40-nm logic cmos process. In *ISCAS*, 2021.
- [266] J. Wang, K. Liu, M. W. Harrington, R. Q. Rudy, and D. J. Blumenthal. Silicon nitride stress-optic microresonator modulator for optical control applications. *Optics Express*, 30(18):31816–31827, 2022.
- [267] J. Wang, K. Liu, M. W. Harrington, R. Q. Rudy, and D. J. Blumenthal. Ultra-low loss silicon nitride ring modulator with low power pzt actuation for photonic control. In *Optical Fiber Communication Conference*, pages W3D–5. Optica Publishing Group, 2022.
- [268] Y. Wang, J. Hulme, P. Sun, M. Jain, M. A. Seyedi, M. Fiorentino, R. G. Beausoleil, and K.-T. Cheng. Characterization and applications of spatial variation models for silicon microring-based optical transceivers. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2020.
- [269] Y. Wang, J. Hulme, P. Sun, M. Jain, M. A. Seyedi, M. Fiorentino, R. G. Beausoleil, and K.-T. Cheng. Characterization and applications of spatial variation models for silicon microring-based optical transceivers. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6. IEEE, 2020.
- [270] Y. Wang, M. A. Seyedi, R. Wu, J. Hulme, M. Fiorentino, R. G. Beausoleil, and K.-T. Cheng. Energy-efficient channel alignment of dwdm silicon photonic transceivers. In *2018 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 601–604. IEEE, 2018.
- [271] Y. E. Wang, G.-Y. Wei, and D. Brooks. Benchmarking tpu, gpu, and cpu platforms for deep learning. *arXiv preprint arXiv:1907.10701*, 2019.
- [272] Z. Wang and A. P. Knights. Using the intrinsic properties of silicon microring modulators for characterization of rf termination. In *Smart Photonic and Optoelectronic Integrated Circuits XIX*, volume 10107, pages 55–59. SPIE, 2017.
- [273] Q. Wilmart, H. El Dirani, N. Tyler, D. Fowler, S. Malhouitre, S. Garcia, M. Casale, S. Kerdiles, K. Hassan, C. Monat, et al. A versatile silicon-silicon nitride photonics platform for enhanced functionalities and applications. *Applied Sciences*, 9(2):255, 2019.

- [274] D. Wu, J. Li, R. Yin, H. Hsiao, Y. Kim, and J. San Miguel. Ugemm: Unary computing architecture for gemm applications. In *2020 ACM/IEEE 47th Annual International Symposium on Computer Architecture (ISCA)*, pages 377–390. IEEE, 2020.
- [275] R. Wu, C.-H. Chen, J.-M. Fedeli, M. Fournier, R. G. Beausoleil, and K.-T. Cheng. Compact modeling and system implications of microring modulators in nanophotonic interconnects. In *2015 ACM/IEEE International Workshop on System Level Interconnect Prediction (SLIP)*, pages 1–6. IEEE, 2015.
- [276] X. Xiao, H. Xu, X. Li, Y. Hu, K. Xiong, Z. Li, T. Chu, Y. Yu, and J. Yu. 25 gbit/s silicon microring modulator based on misalignment-tolerant interleaved pn junctions. *Optics express*, 20(3):2507–2515, 2012.
- [277] Y. Xie, W. Xu, W. Zhao, Y. Huang, T. Song, and M. Guo. Performance optimization and evaluation for torus-based optical networks-on-chip. *Journal of Lightwave Technology*, 33(18):3858–3865, 2015.
- [278] Xilinx. <https://www.xilinx.com/products/silicon-devices/fpga/virtex-ultrascale-plus.html>.
- [279] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson. 12.5 gbit/s carrier-injection-based silicon micro-ring silicon modulators. *Optics express*, 15(2):430–436, 2007.
- [280] Q. Xu, B. Schmidt, S. Pradhan, and M. Lipson. Micrometre-scale silicon electro-optic modulator. *nature*, 435(7040):325–327, 2005.
- [281] X. Xue, P.-H. Wang, Y. Xuan, M. Qi, and A. M. Weiner. Microresonator kerr frequency combs with high conversion efficiency. *Laser & Photonics Reviews*, 11(1):1600276, 2017.
- [282] X. Yan, S. Gitt, B. Lin, D. Witt, M. Abdolahi, A. Afifi, A. Azem, A. Darcie, J. Wu, K. Awan, et al. Silicon photonic quantum computing with spin qubits. *APL Photonics*, 6(7), 2021.
- [283] C.-Y. Yang and Y. Lee. A pwm and pam signaling hybrid technology for serial-link transceivers. *IEEE Transactions on Instrumentation and Measurement*, 57(5):1058–1070, 2008.
- [284] L. Yang et al. On-chip optical matrix-vector multiplier. In *Optics and Photonics for Information Processing*. SPIE, 2013.
- [285] Z. Ying, S. Dhar, Z. Zhao, C. Feng, R. Mital, C.-J. Chung, D. Z. Pan, R. A. Soref, and R. T. Chen. Electro-optic ripple-carry adder in integrated silicon photonics for optical computing. *IEEE journal of selected topics in quantum electronics*, 24(6):1–10, 2018.

- [286] Z. Ying, C. Feng, Z. Zhao, S. Dhar, H. Dalir, J. Gu, Y. Cheng, R. Soref, D. Z. Pan, and R. T. Chen. Electronic-photonic arithmetic logic unit for high-speed computing. *Nature communications*, 11(1):2154, 2020.
- [287] Z. Ying, C. Feng, Z. Zhao, R. Soref, D. Pan, and R. T. Chen. Integrated multi-operand electro-optic logic gates for optical computing. *Applied Physics Letters*, 115(17), 2019.
- [288] Z. Ying, Z. Wang, Z. Zhao, S. Dhar, D. Z. Pan, R. Soref, and R. T. Chen. Silicon microdisk-based full adders for optical computing. *Optics letters*, 43(5):983–986, 2018.
- [289] U. Younis, X. Luo, B. Dong, L. Huang, S. K. Vanga, A. E.-J. Lim, P. G.-Q. Lo, C. Lee, A. A. Bettiol, and K.-W. Ang. Towards low-loss waveguides in soi and ge-on-soi for mid-ir sensing. *Journal of Physics Communications*, 2(4):045029, 2018.
- [290] K. Yu, C. Li, H. Li, A. Titriku, A. Shafik, B. Wang, Z. Wang, R. Bai, C.-H. Chen, M. Fiorentino, et al. A 25 gb/s hybrid-integrated silicon photonic source-synchronous receiver with microring wavelength stabilization. *IEEE Journal of Solid-State Circuits*, 51(9):2129–2141, 2016.
- [291] L. Zhang, R. Ji, L. Jia, L. Yang, P. Zhou, Y. Tian, P. Chen, Y. Lu, Z. Jiang, Y. Liu, et al. Demonstration of directed xor/xnor logic gates using two cascaded microring resonators. *Optics letters*, 35(10):1620–1622, 2010.
- [292] W. Zhang, C. Huang, H.-T. Peng, S. Bilodeau, A. Jha, E. Blow, T. F. de Lima, B. J. Shastri, and P. Prucnal. Silicon microring synapses enable photonic deep learning beyond 9-bit precision. *Optica*, 9(5):579–584, May 2022.
- [293] X. Zhang, X. Zhou, M. Lin, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018.
- [294] Z. Zhao, Z. Zhang, J. Li, Z. Shang, G. Wang, J. Yin, H. Chen, K. Guo, and P. Yan. Mos 2 hybrid integrated micro-ring resonator phase shifter based on a silicon nitride platform. *Optics Letters*, 47(4):949–952, 2022.
- [295] H. Zhou, C. Qiu, X. Jiang, Q. Zhu, Y. He, Y. Zhang, Y. Su, and R. Soref. Compact, submilliwatt, 2×2 silicon thermo-optic switch based on photonic crystal nanobeam cavities. *Photonics Research*, 5(2):108–112, 2017.
- [296] L. Zhou and A. K. Kodi. Probe: Prediction-based optical bandwidth scaling for energy-efficient nocs. In *2013 Seventh IEEE/ACM International Symposium on Networks-on-Chip (NoCS)*, pages 1–8. IEEE, 2013.
- [297] F. Zokaee, Q. Lou, N. Youngblood, W. Liu, Y. Xie, and L. Jiang. Lightbulb: A photonic-nonvolatile-memory-based accelerator for binarized convolutional

neural networks. In *2020 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, pages 1438–1443. IEEE, 2020.

- [298] Y. Zou, S. Chakravarty, C.-J. Chung, X. Xu, and R. T. Chen. Mid-infrared silicon photonic waveguides and devices. *Photonics Research*, 6(4):254–276, 2018.

Vita

Venkata Sai Praneeth Karempudi

Education

- The University of Akron, Akron, Ohio, USA
Master of Science in Electrical Engineering, December 2019.
- Jawaharlal Nehru Technological University, Hyderabad, India
Bachelor of Technology, Electronics and Communication, May 2016.

Invention Disclosure

- *Stochastic Computing Enabled Optical Hardware Architectures for Energy-Efficient and Scalable Acceleration of Deep Neural Networks*, Ishan Thakkar, Sairam Sri Vatsavai and Venkata Sai Praneeth Karempudi, Invention Disclosure, University of Kentucky (UKRF-2831), November, 2023.

Publications

Journal Publications

- Venkata Sai Praneeth Karempudi, Sairam Sri Vatsavai and Ishan G Thakkar, *A Hybrid Time-Amplitude Analog Photonic GEMM Accelerator*, in AIP Journal of Applied Physics, pp. 1-17, Submitted on: November, 2023. (Under Review)
- Venkata Sai Praneeth Karempudi, Janibul Bashir and Ishan G Thakkar, *An Analysis of Various Design Pathways Towards Multi-Terabit Photonic On-Interposer Interconnects*, in ACM Journal on Emerging Technologies in Computing Systems (JETC), pp. 1-33, Accepted on: November, 2023.
- Venkata Sai Praneeth Karempudi, Febin Sunny, Ishan G Thakkar, Sai Vineel Reddy Chittamuru, Mahdi Nikdast, and Sudeep Pasricha, *Photonic Networks-on-Chip Employing Multilevel Signaling: A Cross-Layer Comparative Study*, ACM Journal on Emerging Technologies in Computing Systems (JETC) 18.3 (2022), pp. 1-36.

Conference Publications

- Venkata Sai Praneeth Karempudi, Sairam Sri Vatsavai, Ishan G Thakkar, Oluwaseun Alo, Todd Hastings and Justin Woods, *A Low-Dissipation*

and Scalable GEMM Accelerator with Silicon Nitride Photonics, in IEEE International Symposium on Quality Electronic Design (ISQED), pp. 1-8, Submitted on: November, 2023. (Under Review)

- Venkata Sai Praneeth Karempudi, Sairam Sri Vatsavai, Ishan G Thakkar, and Todd Hastings, *A Polymorphic Electro-Optic Logic Gate for High-Speed Reconfigurable Computing Circuits*, in IEEE International Symposium on Quality Electronic Design (ISQED), San Francisco, CA, USA, 2023, pp. 1-8.
- Venkata Sai Praneeth Karempudi, Ishan G Thakkar, and Todd Hastings, *A Silicon Nitride Microring Based High-Speed, Tuning-Efficient, Electro-Refractive Modulator*, in IEEE International Symposium on Smart Electronic Systems (iSES), Warangal, India, 2022, pp. 307-311.
- Venkata Sai Praneeth Karempudi, Shreyan Datta, and Ishan G Thakkar, *Design exploration and scalability analysis of a CMOS-integrated, polymorphic, nanophotonic arithmetic-logic unit*, in 19th ACM Conference on Embedded Networked Sensor Systems, Coimbra, Portugal, 2021, pp. 628-634.
- Venkata Sai Praneeth Karempudi, Sairam Sri Vatsavai and Ishan G Thakkar, *Redesigning Photonic Interconnects with Silicon-on-Sapphire Device Platform for Ultra-Low-Energy On-Chip Communication*, in ACM Great Lakes Symposium on VLSI (GLSVLSI), 2020, pp. 247-252.
- Venkata Sai Praneeth Karempudi and Ishan G Thakkar, *Mitigating inter-channel crosstalk non-uniformity in microring filter arrays of photonic NoCs: work-in-progress*, in IEEE/ACM International Conference on Hardware/Software Codesign and System Synthesis Companion (CODES/ISSS), 2019, pp. 1-2.

Archived Publications

- Venkata Sai Praneeth Karempudi, Ishan G Thakkar, and Todd Hastings, *A Silicon Nitride Microring Modulator for High-Performance Photonic Integrated Circuits*, 2023, arXiv preprint arXiv:2306.07238.