

DESIGN OF OPTICAL INTERCONNECT TRANSCEIVER CIRCUITS AND  
NETWORK-ON-CHIP ARCHITECTURES FOR INTER- AND INTRA-CHIP  
COMMUNICATION

A Dissertation

by

CHENG LI

Submitted to the Office of Graduate and Professional Studies of  
Texas A&M University  
in partial fulfillment of the requirements for the degree of

DOCTOR OF PHILOSOPHY

Chair of Committee,	Samuel Palermo
Committee Members,	Paul V. Gratz
	Christi K. Madsen
	Duncan M. Walker
Head of Department,	Head of Department

December 2013

Major Subject: Electrical Engineering

Copyright 2013 Cheng Li

## ABSTRACT

The rapid expansion in data communication due to the increased multimedia applications and cloud computing services necessitates improvements in optical transceiver circuitry power efficiency as these systems scale well past 10 Gb/s. In order to meet these requirements, a 26 GHz transimpedance amplifier (TIA) is presented in a 0.25- $\mu\text{m}$  SiGe BiCMOS technology. It employs a transformer-based regulated cascode (RGC) input stage which provides passive negative-feedback gain that enhances the effective transconductance of the TIA's input common-base transistor; reducing the input resistance and providing considerable bandwidth extension without significant noise degradation or power consumption. The TIA achieves a 53 dB $\Omega$  single-ended transimpedance gain with a 26 GHz bandwidth and 21.3 pA $\sqrt{\text{Hz}}$  average input-referred noise current spectral density. Total chip power including output buffering is 28.2 mW from a 2.5 V supply, with the core TIA consuming 8.2 mW, and the chip area including pads is 960  $\mu\text{m} \times 780 \mu\text{m}$ .

With the advance of photonic devices, optical interconnects becomes a promising technology to replace the conventional electrical channels for the high-bandwidth and power efficient inter/intra-chip interconnect. Second, a silicon photonic transceiver is presented for a silicon ring resonator-based optical interconnect architecture in a 1V standard 65nm CMOS technology. The transmitter circuits incorporate high-swing drivers with non-linear pre-emphasis and automatic bias-based tuning for resonance wavelength stabilization. An optical forwarded-clock adaptive inverter-based transimpedance amplifier (TIA) receiver trades-off power for varying link budgets by employing an on-die eye monitor and scaling the TIA supply for the required sensitivity. At 5Gb/s operation, the ring modulator under 4V<sub>pp</sub> driver achieves 12.7dB extinction ratio with 4.04mW power consumption, while

a 0.28nm tuning range is obtained at  $6.8\mu\text{W}/\text{GHz}$  efficiency with the bias-based tuning scheme implemented with the  $2V_{pp}$  transmitter. When tested with a wire-bonded 150f-F p-i-n photodetector, the receiver achieves -12.7dBm sensitivity at a BER= $10^{-15}$  and consumes 2.2mW at 8Gb/s.

Third, a novel Nano-Photonic Network-on-Chip (NoC) architecture, called LumiNoC, is proposed for high performance and power-efficient interconnects for the chip-multi-processors (CMPs). A 64-node LumiNoC under synthetic traffic enjoys 50% less latency at low loads versus other reported photonic NoCs, and  $\sim 25\%$  less latency versus the electrical 2D mesh NoCs on realistic workloads. Under the same ideal throughput, LumiNoC achieves laser power reduction of 78%, and overall power reduction of 44% versus competing designs.

## ACKNOWLEDGEMENTS

It has been my great fortune to work with many wonderful people during my PhD study at Texas A&M University. First and foremost, my advisor, Prof. Samuel Palermo, has inspired me to this interesting and exciting field with his enthusiasm and guided me through my research work with his brightness thinking. I sincerely thank Prof. Paul V. Gratz, Prof. Christi K. Madsen and Prof. Duncan M. Walker for serving as my thesis committee. Your valuable suggestions and discussion are very important for my research. I would especially like to thank Prof. Paul V. Gratz for his guidance on my research on photonic network-on-chip project with his kindness, openness, brightness and patience.

My sincere gratitude also goes to my colleagues (Ehsan Zhian Tabasy, Geng Tang and Alex Titriku) for their collaboration. I will never forget those sleepless nights we spent together trying to catch the deadline of the chip tapeout. I am also thankful to my collaborators outside Texas A&M University. I performed the photonic transceiver test with Rui Bai from Oregon State University at HP Labs, and we exchanged the opinions and experience of circuit design. It is my great pleasure to work with Chin-Hui Chen and Marco Fiorentino from HP Laboratories. They provided tremendous support for the photonic device design and optical testbench set up.

Above all, my research would not be possible without the support from my family. This thesis is dedicated to my wife, who became a great mother of two kids during my PhD study, and my parents who have supported me unconditionally during their life. I feel sorry to my son Jonathan, since most of the time he needed me to be with him, I was staying in the lab for the research work. But Dad promises you we will enjoy more happy family time together, no matter how busy dad will be. My love is always with you.



## TABLE OF CONTENTS

	Page
ABSTRACT . . . . .	ii
ACKNOWLEDGEMENTS . . . . .	iv
TABLE OF CONTENTS . . . . .	v
LIST OF FIGURES . . . . .	vii
LIST OF TABLES . . . . .	xii
1. INTRODUCTION . . . . .	1
2. BACKGROUND . . . . .	6
2.1 Integrated Silicon Photonic Devices for Optical Interconnects . . . . .	6
2.1.1 Laser Sources . . . . .	6
2.1.2 Optical Modulators . . . . .	8
2.1.3 Optical Drop Filter and Waveguide Photodetector . . . . .	12
2.2 Silicon Ring Resonator Based Photonic Interconnects . . . . .	14
3. DESIGN OF OPTICAL RECEIVER FRONT-END CIRCUITS FOR HIGH-SPEED OPTICAL TRANSMISSION* . . . . .	15
3.1 High-Speed Transimpedance Amplifier Design Challenges and Potential Solutions . . . . .	17
3.2 Overview of Bandwidth Extension Techniques . . . . .	18
3.2.1 Series Inductive Peaking . . . . .	18
3.2.2 Conventional RGC Topology . . . . .	20
3.3 Transformer-based RGC Input Stage . . . . .	21
3.4 TIA Design . . . . .	22
3.4.1 TIA Topology . . . . .	22
3.4.2 Bandwidth Extension Analysis . . . . .	26
3.4.3 Transformer Design . . . . .	33
3.5 Experimental Results . . . . .	36
3.6 Summary . . . . .	44
4. SILICON RING RESONATOR TRANSCEIVER DESIGN* . . . . .	45
4.1 Silicon Ring Resonator Based Photonic Interconnect Design Considerations	47
4.2 Silicon Ring Resonator Modeling . . . . .	50
4.3 Transceiver Architecture . . . . .	50

4.4	Non-Linear Pre-emphasis Modulator Driver Transmitter . . . . .	51
4.5	Automatic Bias-based Wavelength Stabilization . . . . .	55
4.6	Optical Forwarded-Clock Receiver . . . . .	62
4.7	Experimental Results . . . . .	66
4.8	Summary . . . . .	74
5.	EXPLORATION OF PHOTONIC NETWORK-ON-CHIP ARCHITECTURES*	75
5.1	Photonic Network-on-chip Technical Background . . . . .	77
5.2	Related Work . . . . .	79
5.3	Power Efficiency in Photonic Interconnect . . . . .	82
5.4	LumiNOC Architecture . . . . .	86
5.4.1	LumiNOC Subnet Design . . . . .	88
5.4.2	Router Microarchitecture . . . . .	93
5.5	Implementation . . . . .	94
5.5.1	Photonic Power Model . . . . .	97
5.5.2	Power Comparison . . . . .	99
5.6	Evaluation . . . . .	101
5.6.1	Methodology . . . . .	101
5.6.2	Synthetic Workload Results . . . . .	102
5.6.3	Realistic Workload Results . . . . .	103
5.6.4	Power Efficiency . . . . .	104
5.6.5	Discussion . . . . .	104
5.7	Summary . . . . .	105
6.	PROJECTION OF SILICON PHOTONICS INTEGRATION . . . . .	106
6.1	Chip Area Estimation for 128-node PNoC . . . . .	106
6.2	Silicon Ring Based Transceiver Energy Efficiency Projection . . . . .	107
7.	CONCLUSION . . . . .	112
	REFERENCES . . . . .	115

## LIST OF FIGURES

FIGURE	Page
1.1 A typical structure of optical transmission system. . . . .	1
1.2 Future chip multiprocessor (CMP) with 256 compute tiles utilizing a global interconnect network-on-chip (NoC). . . . .	3
2.1 Compare of (a) quantum well laser; (b) quantum dot laser. . . . .	6
2.2 Normalized quantum dot comb laser spectrum with channel spacing of 43 GHz (left). A relative intensity noise plot from 100kHz to 10GHz for one channel (right). (Figure reproduced from [1]). . . . .	7
2.3 Ring modulator configuration. . . . .	8
2.4 Simulated transmission spectrum at ring resonator throughput port. . . . .	9
2.5 Cross section view of silicon ring resonators: (a) carrier-injection mode; (b) depletion mode. . . . .	11
2.6 Ring filter configuration. . . . .	12
2.7 Simulated optical spectrum at ring filter drop port. . . . .	13
2.8 Silicon ring resonator-based wavelength-division-multiplexing (WDM) link. . . . .	14
3.1 Optical receiver system block diagram. . . . .	15
3.2 Bandwidth enhancement by inserting a series inductor between the photodiode and the TIA. . . . .	18
3.3 Frequency response of inductive series peaking $\pi$ -network for various $m$ values ( $k=0.3$ ). . . . .	19
3.4 Regulated cascode input stage: (a) conventional topology, (b) proposed transformer-based topology. . . . .	20
3.5 Transformer-based RGC TIA schematic. . . . .	22
3.6 Simulated 2.5-turn transformer coupling coefficient vs frequency with different turn ratios. . . . .	24

3.7	(a) TIA schematic without the output buffer, (b) Equivalent small-signal model, (c) Equivalent analysis model. . . . .	26
3.8	Simulated TIA frequency response with various series inductance values: (a) normalized transimpedance gain, (b) group-delay of input $\pi$ -network. The frequency axis in both curves is normalized to the 3-dB bandwidth without series inductive peaking. . . . .	29
3.9	Simulated 40 Gb/s deterministic jitter performance of the proposed TIA with a $2^{31}$ -1 pattern. . . . .	30
3.10	Simulated transimpedance frequency response with different transformer turn number, transformer turn ratio is fixed at $n = 2$ . . . . .	31
3.11	Simulated transimpedance frequency response with different transformer turn ratio, transformer turn number is fixed at 2.5. . . . .	32
3.12	Simulated TIA performance versus series inductance $L_1$ for different transformer turn ratios: (a) bandwidth, (b) group delay variation. Here the series inductance is normalized to the optimum value of 830 $pH$ . . . . .	32
3.13	Monolithic transformer used for input stage $g_m$ -boosting. . . . .	33
3.14	Simulated transformer coupling coefficient at 20 GHz vs turn number with different turn ratios. . . . .	35
3.15	Die photograph. . . . .	36
3.16	Measured TIA single-ended S-parameters. . . . .	37
3.17	Single-ended simulated/measured transimpedance gain and measured group delay. . . . .	38
3.18	Measurement setup for eye diagram and BER test. . . . .	39
3.19	Measured 27 Gb/s single-ended eye-diagram with a $125 \mu A_{pp}$ $2^{15}$ -1 PRBS input signal. . . . .	40
3.20	Post-layout simulated single-ended 40 Gb/s eye-diagram of the proposed TIA with $100 \mu A_{pp}$ input current. . . . .	40
3.21	BER jitter bathtub plot with 25 Gbps $150 \mu A_{pp}$ PRBS $2^{15}$ -1 input. . . . .	41
3.22	Measured BER versus input current. . . . .	41
3.23	Measured TIA single-ended integrated output noise. . . . .	42

3.24	Simulated and calculated input-referred current noise density for the proposed transformer-based RGC input-stage TIA and a simple common-base input-stage TIA. . . . .	43
4.1	(a) Top and cross section views of carrier-injection silicon ring resonator modulator, (b) optical spectrum at through port. . . . .	47
4.2	Measured quality factor and resonance wavelength of nine $2.5\mu\text{m}$ radius silicon ring modulators fabricated on an 8" 130nm CMOS SOI wafer. . .	48
4.3	Photonic transceiver circuits prototype block diagram. . . . .	50
4.4	Simulated carrier-injection ring resonator modulator response to 200ps data pulses with: (a) $2V_{pp}$ simple modulation, (b) $4V_{pp}$ simple modulation, , (c) $4V_{pp}$ modulation with pre-emphasis. . . . .	53
4.5	Non-linear pre-emphasis modulator driver transmitters: (a) transmitter block diagrams, (b) per-terminal 2V pre-emphasis driver, (c) tunable delay cell, (d) pulsed-cascode output stage. . . . .	55
4.6	Ring resonator modulator transmission curves with high and low modulation voltage levels when: (a) resonance wavelength is not aligned with input laser wavelength; (b) resonance wavelength is aligned with input laser wavelength. . . . .	56
4.7	Measured carrier-injection ring resonator modulator performance: (a) optical transmission spectrum at different bias levels; (b) resonance wavelength shift versus bias voltage. . . . .	57
4.8	Bias-based ring resonator modulator semi-digital wavelength stabilization loop. . . . .	58
4.9	9-bit non-linear bias tuning DAC. . . . .	59
4.10	Ring resonator bias-based tuning algorithm. . . . .	59
4.11	Simulated tuning waveforms and final optical transmission curves for (a) static tuning mode, and (b) dynamic tuning mode. . . . .	61
4.12	Extinction ratio versus modulated wavelength shift for static and dynamic tuning modes. . . . .	61
4.13	Adaptive sensitivity-power data receiver. . . . .	62
4.14	Inverter-based TIA front-end: (a) schematic, (b) simulated TIA common-mode output response to a 5mV power supply step. . . . .	64

4.15	Optical receiver sensitivity-power adaption algorithm. . . . .	65
4.16	Optical clock receiver. . . . .	65
4.17	Optical transceiver circuits prototype bonded for electrical characterization and optical testing. (a) Optical transmitter configuration with silicon ring resonator modulators. (b) Optical receiver configuration with commercial photodetectors. . . . .	66
4.18	Modulator drivers' electrical eye diagrams. 9Gb/s operation with $2V_{pp}$ driver with (a) minimum pre-emphasis, (b) maximum pre-emphasis. 8Gb/s operation with $4V_{pp}$ driver with (c) minimum pre-emphasis (d) maximum pre-emphasis. . . . .	67
4.19	5 Gb/s optical eye diagrams with silicon carrier-injection ring resonator modulators driven by the $4V_{pp}$ transmitter: (a) minimum pre-emphasis settings; (b) optimized pre-emphasis settings. . . . .	68
4.20	Ring resonator bias-based wavelength stabilization measurements: (a) ring 1's 500Mb/s eye diagrams demonstrating the automatic bias tuning stabilizing to 1286.93nm, (b) ring 2's 800Mb/s eye diagrams with input laser wavelengths of 1311.86 and 1311.96nm. . . . .	69
4.21	8Gb/s receiver supply scaling measurements: (a) sensitivity (BER= $10^{-15}$ ) and power versus TIA supply voltage, (b) BER bathtub plot for a power supply of 0.96V. . . . .	70
4.22	Integrated photodetector emulator circuit. . . . .	71
4.23	10Gb/s receiver supply scaling measurements with integrated photodetector emulator. . . . .	71
4.24	Optically forwarded-clock receiver measurements: (a) 2GHz recovered clock waveform, (b) jitter versus input optical power. . . . .	72
5.1	Basics of photonic on-chip interconnect. . . . .	77
5.2	Optical link budgets for the photonic data channels of various photonic NoCs. . . . .	84
5.3	Optical power overhead of arbitration channels in various photonic NoCs. . . . .	85
5.4	LumiNOC interconnection of CMP with 16 tiles - (a) One-row interconnection, (b) Two-rows interconnection, (c) Four-rows interconnection. . . . .	86
5.5	Bold circles (TX and RX) represent groups of rings, and each pair in the oval are for a single node. . . . .	89

5.6	Arbitration on 4 a node subnet. . . . .	92
5.7	Router microarchitecture. . . . .	93
5.8	One-row LumiNOC with 64 tiles. . . . .	94
5.9	Electrical Laser Power (W) contour plots for networks with the same aggregate throughput (assuming 30% efficient electrical to optical power conversion) . . . . .	98
5.10	Synthetic workloads showing LumiNOC vs. Clos LTBw and electrical network. . . . .	103
5.11	Message Latency in PARSEC benchmarks for LumiNOC compared to electrical network. . . . .	104
6.1	Schematic of a single-stage inverter-based resistive shunt feedback CMOS TIA with a photodetector. . . . .	109
6.2	Transceiver circuitry power efficiency vs. data rate under different CMOS technologies. . . . .	110

## LIST OF TABLES

TABLE	Page
3.1 Key parameters of the input-stage transformer . . . . .	23
3.2 TIA performance comparisons . . . . .	44
4.1 Performance summary and comparisons. . . . .	73
5.1 Components of optical loss. . . . .	96
5.2 Configuration comparison of various photonic NoC architectures - $N_{core}$ = number of cores in the CMP, $N_{node}$ = number of nodes in the NoC, $N_{rt}$ = total number of routers, $N_{wg}$ = total number of waveguides, $N_{wv}$ = total number of wavelengths, $N_{ring}$ = total number of rings, ITP = Ideal Throughput. . . . .	99
5.3 Power efficiency comparison of different photonic NoC architectures - ELP = Electrical Laser Power, TTP = Thermal Tuning Power, ERP = Elec- trical Router Power, EO/OE = Electrical to optical/Optical to electrical conversion power, ITP = Ideal Throughput, TP = Total Power. . . . .	100
6.1 Area of building blocks in silicon ring-based photonic transceiver. . . . .	107
6.2 Technology roadmap for CMOS transistors. . . . .	108
6.3 128-node PNoC area estimation, via/pad area vs circuits area. . . . .	108



## 1. INTRODUCTION

Bandwidth demands continue to scale up rapidly to meet the development of bandwidth-hungry multimedia applications and cloud computing services. There is an urgent need to maximize the capacity that can be transported by optical backbone networks in order to meet both business and residential customers' requirements, while as always at a cost-effective way (lower energy cost per bit). As a result, energy-efficient optical transceiver circuits are paramount for the high-speed long-haul and mid-range optical transmission system.




Figure 1.1: A typical structure of optical transmission system.

Fig. 1.1 illustrates a typical structure of optical transmission system. At the transmitter side, the laser source is modulated by the optical modulator driven by the electrical driver, and then the modulated optical signal is transmitted via single-mode fiber. Depending on the transmission distance, optical amplifiers might be inserted in certain span to compensate the optical signal loss in the transmission. At receiver side, a photodiode first converts the optical signal to electrical current. Then an optical receiver front-end circuits, transimpedance amplifier (TIA), converts and amplifies the current to a voltage

signal, which is further amplified by the limiting amplifier to achieve a signal sufficient for the reliable operation of the subsequent clock and data recovery (CDR) circuits. In this dissertation, chapter 3 presents an optical receiver front-end TIA serving as a very important building block in the optical receiver system. As the optical pre-amplifier, TIA dominates the entire receiver's power, bandwidth, sensitivity and noise performance.

In addition, energy-efficient interconnects are paramount for next-generation high-performance networking and computing applications. However, conventional inter/intra-chip electrical interconnects will not be able to increase their pin-bandwidths significantly due to channel-loss limitations. Signal attenuation, dispersion and cross-talk severely limit the reach of copper-based links beyond 10 Gb/s. While many have proposed techniques to overcome these limitations and extend the reach of copper, such techniques are usually complicated or have high power consumption requirements, and will not scale to higher data rates. Therefore, optical short reach interconnects are emerging as a replacement for the conventional electrical link as the inter-chip or even future intra-chip communication method. Chapter 4 describes a silicon photonic transceiver circuits for a ring resonator-based optical interconnect architecture, providing the potential for silicon photonic links that can deliver distance-independent connectivity whose pin-bandwidth scales with the degree of wavelength-division multiplexing.

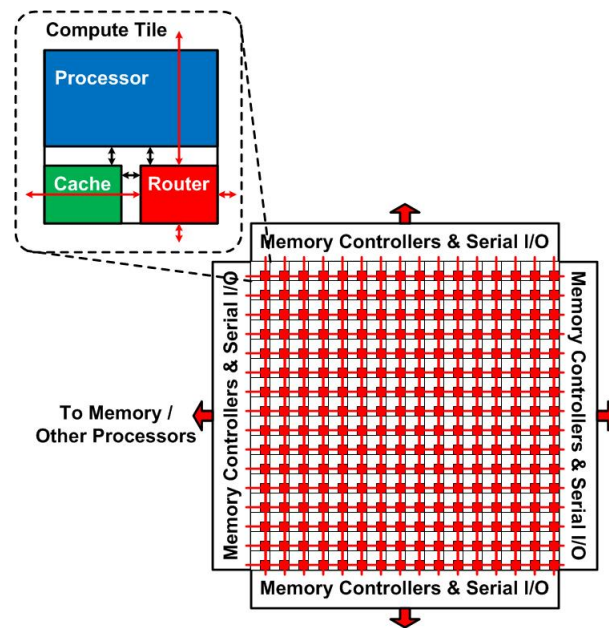


Figure 1.2: Future chip multiprocessor (CMP) with 256 compute tiles utilizing a global interconnect network-on-chip (NoC).

The amount of data communicated between cores and off-chip to memory or other processors also scales as the number of cores increases in future many-core systems shown in Fig. 1.2. Projections based on data from the International Technology Roadmap for Semiconductors (ITRS) [2] shows that greater than tera-byte aggregate bandwidth is required for the future many-core chip-multiprocessors (CMPs). This explosion in both intra- and inter-chip bandwidth requires interconnect systems to achieve very high energy efficiency in order to comply with power budgets that have plateaued near 100W due to thermal constraints. However, the electrical channel limitations is even more obvious for the on-chip wires. Serious challenges exist in achieving projected communication bandwidth over electrical channels while still satisfying I/O power and density constraints due to high-frequency loss of electrical traces, reflections caused from impedance discontinuities, and crosstalk from adjacent signals. While advanced signaling and circuit techniques, such

as equalization [3], capacitive drivers [4], and RF-interconnects [5], can be leveraged to extend on-chip wire bandwidth, the significant power costs incurred with this additional complexity is prohibitive for future compute systems. The efficiency of current state-of-the-art Network-on-chip (NoCs) with simple CMOS inverter-based repeaters is near 2pJ/bit [6], allowing for only near 1TB/s throughput with a typical 20% allowance from the total 100W processor power budget. Monolithic silicon photonics, which offers high-speed photonic devices, THz-bandwidth waveguides and immense bandwidth-density via wavelength-division-multiplexing (WDM) [7–12], provides architectures suitable to efficiently scale to meet future many-core systems' bandwidth demands. Typical optical channels, including glass fibers and on-chip waveguides display signal loss characteristics which varies only fractions of dBs over wide wavelength ranges (tens of nanometers), allowing for data transmission of several Tb/s without the requirement of channel equalization. This simplifies design of optical links in a manner similar to non-channel limited electrical links. Another important feature of optical interconnects is the ability to combine multiple data channels on a single waveguide via wavelength-division-multiplexing (WDM) and greatly improve bandwidth density.

Due to the benefits of silicon photonics, recently photonic NoCs (PNoCs) have emerged as a potential replacement for electrical NoCs. Much of the current research in PNoCs focuses on leveraging the high bandwidth of photonic interconnect. Some designs propose using electrical interconnect to coordinate and arbitrate a shared photonic medium, effectively trading increased latency for higher bandwidth. While increased bandwidth without regard for latency is useful for some applications, it eschews the primary benefit of PNoCs over electrical NoCs for CMPs, low latency. Other recent photonic NoC proposals attempt to address the latency of arbitration. In particular, several groups have proposed crossbar or Clos topologies to improve the latency of multi-core photonic interconnect arbitration. While these designs do provide low latency and high bandwidth, it

comes at a high cost in terms of bandwidth per Watt of static power due to the need to significantly over-provision the network to achieve low latency. There is a clear need for a PNoC architecture that is energy-efficient and scalable while maintaining the goals of low latency and high bandwidth. In chapter 5, a novel photonic NoC architecture, called LumiNOC, is proposed to address the issues of power and resource overhead due to channel over-provisioning, while reducing latency and maintaining high bandwidth.

This dissertation is organized as below. Chapter 2 describes the background of integrated silicon photonics. A TIA design which employing a novel transconductance boosting and series inductive peaking to efficiently obtain significant bandwidth extension and low-noise performance is discussed in Chapter 3. Chapter 4 presents a silicon photonic transceiver circuits for a ring resonator-based optical interconnect architecture that address limited modulator bandwidth, variations in ring resonator resonance wavelength and link budget, and efficient receiver clocking. Moving to the architecture level, Chapter 5 discusses a novel photonic NoC (LumiNOC) architecture to address the issues of power and resource overhead due to channel over-provisioning, while reducing latency and maintaining high bandwidth. Chapter 6 projects the 128-node PNoC area cost and link power efficiency under different CMOS technologies. Finally, Chapter 7 concludes the thesis.

## 2. BACKGROUND

This chapter gives an overview of the silicon photonic components for energy-efficient and compact-size intra/inter-chip interconnects. The theory of ring resonator is discussed and two types of silicon ring resonator, carrier-injection ring and depletion ring, are also compared. This section ends with an introduction of a silicon ring resonator based wavelength-division-multiplexing (WDM) link as a potential alternative of the electrical channel for future high-speed, power-efficient interconnects.

### 2.1 Integrated Silicon Photonic Devices for Optical Interconnects

#### 2.1.1 Laser Sources

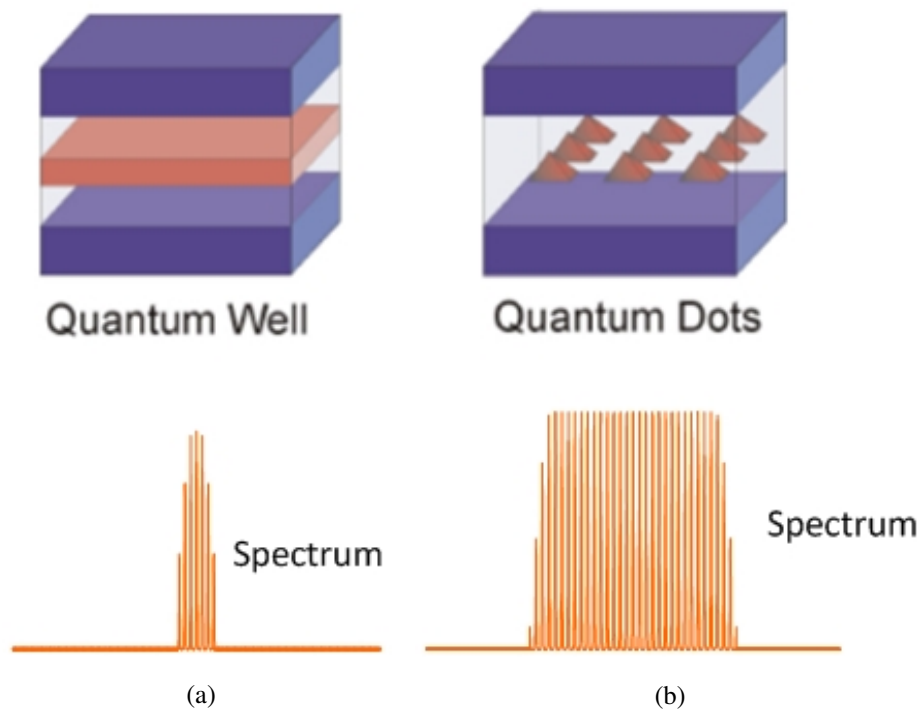


Figure 2.1: Compare of (a) quantum well laser; (b) quantum dot laser.

An impediment to adopt high density WDM for short reach optical interconnects is that each WDM wavelength currently requires its own expensive DFB laser [13]. The alternative is to use a power and area efficient broad-spectrum light emitter replacing the gang of DFB lasers. Comb laser [1] is developed to meet the short range optical interconnect requirement. The comb laser injects multiple wavelengths at 1310nm wavelength range to the silicon waveguide via the optical grating coupler. Unlike the conventional laser sources (Fig. 2.1a), which are composed of multiple quantum wells, the comb laser uses quantum dots (Fig. 2.1b) and generates multiple wavelengths simultaneously. Instead of using fewer wavelength with each modulated at very high speed (e.g. using multiple DFB laser sources), which dramatically increases the power at electrical side, comb laser uses more wavelengths with each modulated at lower data rate to achieve the same aggregate bandwidth, however, at much better overall power efficiency. Typically, a comb laser can generate 16-64 effective wavelengths with optical power of 0.2-1mW on each channel, as shown in Fig. 2.2. Channel spacing is in the range of 50GHz to 100GHz.




Figure 2.2: Normalized quantum dot comb laser spectrum with channel spacing of 43 GHz (left). A relative intensity noise plot from 100kHz to 10GHz for one channel (right). (Figure reproduced from [1]).

### 2.1.2 Optical Modulators

**Silicon Ring Resonator:** Silicon ring resonator is a potential candidate to enable the platform for large-scale monolithic integration of optics and microelectronics. It can be configured either as an optical modulator or a WDM drop filter. Silicon ring resonator modulators/filters offer advantages of small size, relative to Mach-Zehnder modulators [14, 15], and increased filter functionality, relative to electro-absorption modulators [16]. Silicon photonic links based on ring resonator devices provide a unique opportunity to deliver distance-independent connectivity whose pin-bandwidth scales with the degree of wavelength-division multiplexing.




Figure 2.3: Ring modulator configuration.

A basic silicon ring resonator consists of a straight waveguide unidirectionally coupled with a circular waveguide, as shown in Fig. 2.3. When an optical input signal with power  $P_{t1}$  is launched into the input port (left side) of the waveguide, its intensity is split into an output signal at through port (right side) and a feedback signal which is either coupled to the through port of the waveguide or trapped inside the ring. With the input power normalized to unity, the transmission power  $P_{t2}$  at throughput port can be obtained by the following equation 2.1:



$$P_{t2} = \frac{a^2 + |t|^2 - 2a|t|\cos(\theta + \phi_t)}{1 + a^2|t|^2 - 2a|t|\cos(\theta + \phi_t)}, \quad (2.1)$$

where  $|t|$  represents the coupling losses,  $\phi_t$  is the phase of the coupler and  $\alpha$  is the loss coefficient of the ring. In the equation,  $\theta=4\pi^2 n_{eff} \frac{r}{\lambda}$ , where  $n_{eff}$  is the effective refractive index and  $\lambda$  is the optical wavelength. At the resonance, most energy will be trapped in the ring resonator due to the destructive interference within the ring. Fig. 2.4 shows the normalized output power at ring modulator throughput port as a function of the wavelength. The spectrum displays a notch-shaped characteristic at periodic resonance wavelengths, repeating over a free spectral range (FSR) defined as equation 2.2.

$$FSR = \frac{\lambda^2}{n_{eff}L}, \quad (2.2)$$

where  $n_{eff}$  is the effective refractive index of the ring waveguide, and  $L$  is the circumference of the ring device. Because FSR is inversely proportional to the size of the ring resonator, the ring must be small in order to achieve a high FSR.

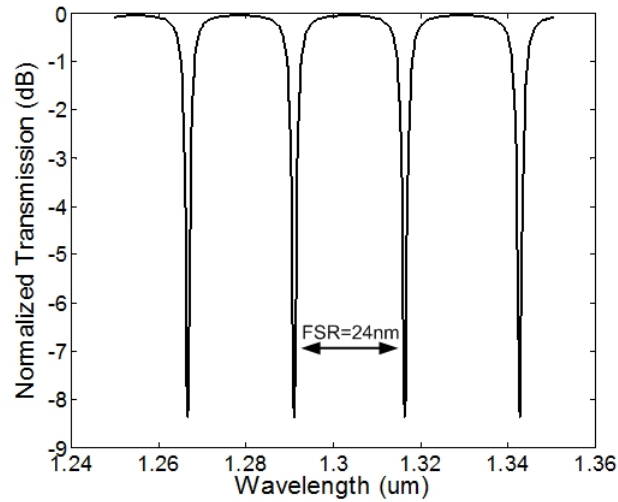


Figure 2.4: Simulated transmission spectrum at ring resonator throughput port.

The full width and half maximum (FWHM) of the ring device can be written as equation 2.3, with  $k$  represents the normalized coupling coefficient of the coupler between the straight waveguide and circular waveguide of the ring device.

$$FWHM = \frac{k^2 \lambda^2}{\pi * L * n_{eff}}, \quad (2.3)$$

The quality factor ( $Q$ ) is another key specification of the ring resonator, which is a measure of the sharpness of the resonance. It is defined as the ratio of the operation wavelength and the resonance width, shown in equation 2.4. A typical silicon ring resonator can achieve a relatively large  $Q$  of 8000 [17].

$$Q = \pi \frac{n_{eff} L}{\lambda} \frac{t}{1 - t^2}, \quad (2.4)$$

Fig. 2.5 shows the cross section view of two types of silicon ring resonator modulators. The p-i-n junction-based carrier-injection devices [11, 17], shown in Fig. 2.5a, operate primarily in forward-bias. The waveguide region defining the optical mode is confined within the intrinsic region to avoid optical absorption losses in the heavily doped p-type and n-type regions. When the junction is forward-biased, carriers can be injected into the intrinsic region, where the refractive index is also changed. Modulation based on carrier-injection ring generally can achieve large extinction ratio. However, the modulation speed is generally limited due to the long minority carrier lifetime ( $\sim 1$ ns) of the p-i-n junction. This limitation can be partially mitigated by using modulation equalization technic [11]. The carrier-depletion devices [18], shown in Fig. 2.5b, operate primarily in reverse-bias. The waveguide in depletion device is lightly doped, resulting a p-n diode can be operated in reverse-bias to deplete carriers from a central region [19]. Although a depletion ring generally achieves higher modulation speeds relative to a carrier-injection ring due

to the ability to rapidly change the depletion width, its modulation depth is limited due to the relatively low doping concentration in the waveguide to avoid excessive loss. In contrast, carrier-injection ring modulators can provide large refractive index changes and high modulation depths, but are limited by long minority carrier lifetimes.




Figure 2.5: Cross section view of silicon ring resonators: (a) carrier-injection mode; (b) depletion mode.

However, a major barriers to widespread adoption of ring-based silicon photonics is the non-uniformity in the fabrication at both the die and wafer scales [20]. For example, the resonance wavelengths of silicon ring resonators depend on the device dimensions, effective refractive index and etch depths across a wafer. Due to the fabrication variation, identical rings at different locations in the wafer can cause significant variation in their passband wavelengths. it is very difficult to fabricate ring resonated at the exactly required wavelength under current silicon-on-insulator (SOI) process. In addition, silicon ring resonator performance is also sensitive to the temperature variation, which causes the resonance wavelength drifts and degrades the modulation extinction ratio. Therefore tunability is essential for the piratical application of ring modulator. Two methods are commonly used for tuning the resonance wavelength when the ring modulator is in operation, which are the thermo-optic tuning and electro-optic tuning. Thermo-optic tuning is implemented by implanting a heater nearby the ring waveguide to heat the entire device.

The heat changes the refractive index of the material which in turn shifts the resonances towards the larger wavelength. It should be noted that the thermo-optic tuning process is fairly slow ( $\sim$ ms) and significant power is needed to maintain the tuned status. However it is suitable for the case where large refractive index change is required. Another tuning method is electro-optic tuning. An electric field is applied over the ring to change carrier density in the waveguide, leading to the refractive index change. The tuning range is relatively small compared to the thermal tuning method [8], but it has the advantage of fast tuning speed and better tuning efficiency.

### 2.1.3 Optical Drop Filter and Waveguide Photodetector

**Ring Drop Filter:** Silicon ring resonator can also be configured as the optical filter by adding an additional straight waveguide functioning as add-drop port, as shown in Fig. 2.6. The four ports of the ring resonator are referred as input port, throughput port, drop port and add port. The simulated optical spectrum at ring filter drop port is shown in Fig. 2.7. The resonance peaks at periodic resonance wavelengths also repeat over a free spectral range (FSR). The modulated signal can be filtered out by aligning the resonance peak with the carrier wavelength.




Figure 2.6: Ring filter configuration.




Figure 2.7: Simulated optical spectrum at ring filter drop port.

**Waveguide Photodetector:** Photodetector absorbs incident light, and then creates accumulated charge carriers that can be measured by electronic circuits. Conventional lateral p-i-n photodetector is not suitable for the integrated photonics applications due to the large area required to improve the photo-response. Recently, significant research in silicon photonics has been focused on realizing individual components of photonic integrated circuits. However, silicon is transparent to the standard telecommunication wavelengths used for short-range photonic interconnects, and therefore cannot be used as an active element of a photodetector. In recent years the epitaxial integration of germanium with silicon waveguides has led to several device structures that are promising for high-bandwidth interconnects. For example, a germanium waveguide photodetector has been demonstrated using selective growth on a siliconon-insulator platform [21, 22], and a SiGe waveguide photodetector has been developed to reduce the lattice mismatch experienced by Ge photodetectors [23]. However, their dark current densities are typically higher than conventional

III-V photodetectors primarily due to dislocations from the growth on a silicon substrate. In addition, their absorption is typically lower at wavelengths beyond 1550nm, leading to lower responsivity at longer wavelengths.

## 2.2 Silicon Ring Resonator Based Photonic Interconnects




Figure 2.8: Silicon ring resonator-based wavelength-division-multiplexing (WDM) link.

Silicon photonic links based on ring resonator devices provide a unique opportunity to deliver distance-independent connectivity whose pin-bandwidth scales with the degree of wavelength-division multiplexing. As shown in Fig. 2.8, multiple wavelengths ( $\lambda_1 - \lambda_4$ ) generated by an off-chip comb laser are coupled into a silicon waveguide via an optical coupler. At the transmit side, ring modulators insert data onto a specific wavelength through electro-optical modulation. These modulated optical signals propagate through the waveguide and arrive at the receiver side where ring filters drop the modulated optical signals of a specific wavelength at a receiver channel with photodetectors (PD) that convert the signals back to the electrical domain.

### 3. DESIGN OF OPTICAL RECEIVER FRONT-END CIRCUITS FOR HIGH-SPEED OPTICAL TRANSMISSION\*

The continuous growth of data volume due to increased multimedia applications and cloud computing services requires that the data rates of optical communication systems scale to supply this demand. This rapid expansion in data communication also necessitates improvements in optical transceiver circuitry power efficiency as these systems scale well past 10Gb/s.




Figure 3.1: Optical receiver system block diagram.

A typical optical receiver architecture is shown in Fig. 3.1. The photodetector detects the optical signal and converts it to electrical current. A transimpedance amplifier (TIA) then converts this current signal into a voltage which is passed through a limiting amplifier (LA) to achieve a signal sufficient for reliable operation of the subsequent clock and data recovery (CDR) circuits. After the CDR, a demultiplexer is used to generate multiple low-speed data streams for further processing. Transimpedance amplifiers as the optical receiver front-end circuits typically determine the overall optical link performance, as their

---

\*Reprinted with permission from "A Low-Power 26-GHz Transformer-Based Regulated Cascode SiGe BiCMOS Transimpedance Amplifier" by Cheng Li, 2013, IEEE Journal of Solid-State Circuits, Volume: 48, Issue: 5, Page(s): 1264 - 1275, Copyright 2013 by IEEE

speed and sensitivity set the maximum data rate and tolerable channel loss.

This chapter describes a 26 GHz transimpedance amplifier (TIA) that employs a transformer-based regulated cascode (RGC) input stage which provides passive negative-feedback gain that enhances the effective transconductance of the TIA's input common-base transistor; reducing the input resistance and isolating the parasitic photodiode capacitance. This allows for considerable bandwidth extension without significant noise degradation or power consumption. Further bandwidth extension is achieved through series inductive peaking to isolate the photodetector capacitance from the TIA input. The optimum choice of series inductive peaking value and key transformer parameters for bandwidth extension and jitter minimization is analyzed. Fabricated in a 0.25- $\mu\text{m}$  SiGe BiCMOS technology and tested with an on-chip 150  $fF$  capacitor to emulate a photodiode, the TIA achieves a 53  $\text{dB}\Omega$  single-ended transimpedance gain with a 26 GHz bandwidth and 21.3  $\text{pA}\sqrt{\text{Hz}}$  average input-referred noise current spectral density. Total chip power including output buffering is 28.2 mW from a 2.5 V supply, with the core TIA consuming 8.2 mW, and the chip area including pads is 960  $\mu\text{m} \times 780 \mu\text{m}$ .

This chapter is organized as following. Design challenges of high-speed TIA and potential solutions are discussed in Section 3.1. Common bandwidth extension techniques, including series inductive peaking and the active regulated cascode topology, are reviewed in Section 3.2. Section 3.3 discusses the transformer-based RGC input stage, where the mutual magnetic coupling of the on-chip transformer provides a negative feedback between the emitter and base terminals of the common-base input stage; improving the input transistor effective transconductance and allowing for extended bandwidth. The complete TIA topology is detailed in Section 3.4, along with analysis on the optimization of the series inductance and key transformer parameters to extend bandwidth, while limiting frequency peaking and group delay variation. Experimental results of the TIA, fabricated in a 0.25  $\mu\text{m}$  SiGe BiCMOS technology, are presented in Section 3.5. Finally, Section 3.6



concludes the chapter.

### 3.1 High-Speed Transimpedance Amplifier Design Challenges and Potential Solutions

One TIA design challenge stems from the potentially large photodiode parasitic capacitance, which deteriorates both the bandwidth and noise performance of the system. Various input stages have been proposed [24–27] to relax this bandwidth limitation. A popular technique to obtain a very small input resistance involves modifying a conventional common-gate/common-base (CG/CB) input stage to a regulated cascode (RGC) architecture which employs active negative feedback gain to boost the input transconductance [24, 25]. This reduced input resistance pushes the input pole to a higher frequency, relaxing trade-offs between TIA gain and bandwidth. However, conventional RGC topologies require additional voltage headroom due to the cascode topology. Moreover, extra power is required in the feedback stage in order to avoid excessive TIA frequency peaking and obtain sufficient noise performance [24].

An efficient way to boost transistor transconductance involves passive transformer-based negative feedback. In this method, magnetic coupling between the transformer primary and secondary windings is utilized to realize negative feedback gain without introducing additional power and noise. While this approach has been employed in narrow-band LNA design [28], applying this in broad-band TIA design requires tight control on frequency peaking and group delay variation, particularly when combined with other bandwidth extension techniques [29].

Series inductive peaking is another technique to extending TIA bandwidth. Placing inductors in series between amplifier stages forms an equivalent  $\pi$ -network which isolates the capacitance of the stages [30, 31]. In TIA design, this is often used to isolate the photodetector capacitance from the TIA input capacitance. While this approach is effective, the inductance should be optimized to limit frequency peaking and group delay variation.

### 3.2 Overview of Bandwidth Extension Techniques

This section reviews the two key bandwidth extension techniques used in the presented TIA design, series inductive peaking and input transistor transconductance-boosting via the regulated cascode topology.

#### 3.2.1 Series Inductive Peaking




Figure 3.2: Bandwidth enhancement by inserting a series inductor between the photodiode and the TIA.

Series inductive peaking [30–34] is an effective method to extend bandwidth in multi-stage amplifiers by isolating a stage’s output capacitance from the subsequent stage’s input capacitance. This technique is often leveraged in TIA design by interposing an inductor between the photodiode and the circuit input, as shown in Fig. 3.2. From the equivalent small-signal model, the series inductor  $L_1$  isolates the two parasitic capacitors ( $C_{pd}$  and  $C_{in}$ ), forming a  $\pi$ -network which extends the bandwidth relative to a lumped  $RC$  system. Following a similar approach as in [33], the current-mode transfer function of this  $\pi$ -network can be expressed as the following third-order expression.




Figure 3.3: Frequency response of inductive series peaking  $\pi$ -network for various  $m$  values ( $k=0.3$ ).

$$\begin{aligned}
 \frac{I_{in}}{I_{pd}} &= \frac{1}{s^3 R_{in} L_1 C_{pd} C_{in} + s^2 L_1 C_{pd} + s R_{in} (C_{pd} + C_{in}) + 1} \\
 &= \frac{1}{\left(\frac{s}{\omega_0}\right)^3 \frac{k}{m} (1-k) + \left(\frac{s}{\omega_0}\right)^2 \frac{1-k}{m} + \frac{s}{\omega_0} + 1}
 \end{aligned} \tag{3.1}$$

where  $k = \frac{C_{in}}{C_{in} + C_{pd}}$  and  $m = \frac{R_{in}^2 (C_{pd} + C_{in})}{L_1}$ . Significant bandwidth extension ratios (BWER) can be achieved by choosing different  $k$  and  $m$  values, as shown in Fig. 3.3, where the frequency is normalized to the 3-dB frequency ( $\omega_0 = \frac{1}{(C_{pd} + C_{in}) R_{in}}$ ) of the uncompensated case with  $L_1=0$ . However, it is important to avoid values which cause large gain ripple in the frequency response, as this introduces large group delay variation and results in significant signal distortion [31]. A more detailed analysis of the relationship between the bandwidth extension and group-delay variation in the proposed TIA can be found in Section 3.4.

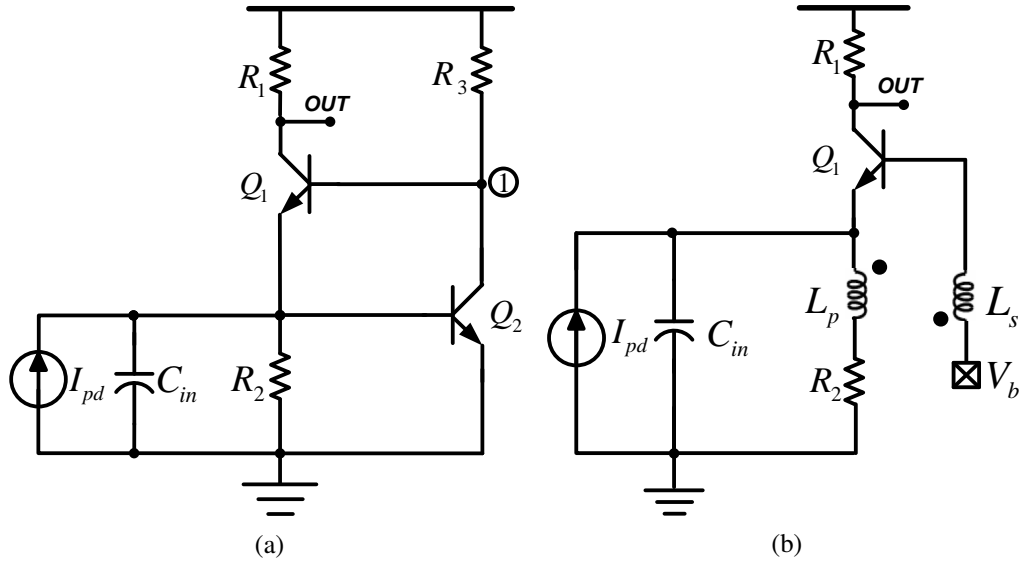


Figure 3.4: Regulated cascode input stage: (a) conventional topology, (b) proposed transformer-based topology.

### 3.2.2 Conventional RGC Topology

TIA bandwidth extension can also be achieved by reducing the input resistance. The regulated cascode topology [24], shown in Fig. 3.4a, achieves this by using a common-base input stage ( $Q_1$ ) with local active feedback ( $Q_2$ ) to boost the transconductance of  $Q_1$  and provide a small signal input resistance of

$$R_{in} \simeq \frac{1}{g_{m1}(1 + g_{m2}R_3)}. \quad (3.2)$$

An important feature of this gain-boosted common-base input stage is that it isolates the photodiode capacitance from subsequent amplifier stages. Used in combination with a subsequent feedback TIA, a high transimpedance can be achieved while maintaining stability over a wide input capacitance range.

However, a conventional RGC topology has a power overhead due to the headroom

necessary to support the two base-emitter voltages and maintain a suitable frequency response. In addition, the local feedback stage introduces a zero ( $z_1$ ) in the transimpedance transfer function. This zero can be estimated by  $z_1 = (r_{\pi 1} || r_{o2} || R_3)C_{p1}$ , where,  $r_{\pi 1}$  is the base-emitter resistance of transistor  $Q_1$ ,  $r_{o2}$  is the collector resistance of transistor  $Q_2$ , and  $C_{p1}$  is the total parasitic capacitance at node 1 of Fig. 3.4a. In order to avoid the frequency peaking of transimpedance gain, a smaller resistor  $R_3$  is normally used to set this zero in the roll-off region of the gain curve [35]. For a given negative-feedback gain and  $g_m$ -boosting factor, this results in an increased  $Q_2$  bias current. In addition, the local feedback transistor can contribute substantial thermal noise at high frequency, thus degrading the system noise performance.

### 3.3 Transformer-based RGC Input Stage

The TIA proposed in this work employs passive transformer-based negative feedback, shown in Fig. 3.4b, in order to provide input  $g_m$ -boosting. Relative to a conventional RGC input stage, this approach trades off increased area from the large transformer to avoid the power and noise of an added active amplifier stage. Here the transformer consists of primary ( $L_p$ ) and secondary inductors ( $L_s$ ), with the bias voltage  $V_b$  provided externally. Feedback via mutual magnetic coupling in the transformer provides anti-phase operation between the emitter and base terminals, thus boosting the transconductance of the common-base transistor to

$$g'_m = (1 + nk)g_m. \quad (3.3)$$

Here the turn ratio  $n$  is

$$n = \sqrt{L_s/L_p} \quad (3.4)$$

and the coupling coefficient  $k$  is

$$k = M/\sqrt{L_s L_p}, \quad (3.5)$$

where  $M$  is the mutual inductance between the primary and secondary windings [36]. The coupling coefficient indicates the magnetic coupling strength in the transformer and is intrinsically less than unity due to magnetic flux leakage. In order to effectively boost the transconductance of the input transistor, thus reducing the effective resistance at input node and extending the TIA bandwidth, the monolithic transformer should be designed to achieve a relatively high magnetic coupling coefficient over a wide frequency range. This implies careful design of the wires that comprise the transformer windings, the turn number, and the turn ratio. The details of transformer design for achieving considerable bandwidth extension and low deterministic jitter are described in Section IV.

### 3.4 TIA Design

#### 3.4.1 TIA Topology

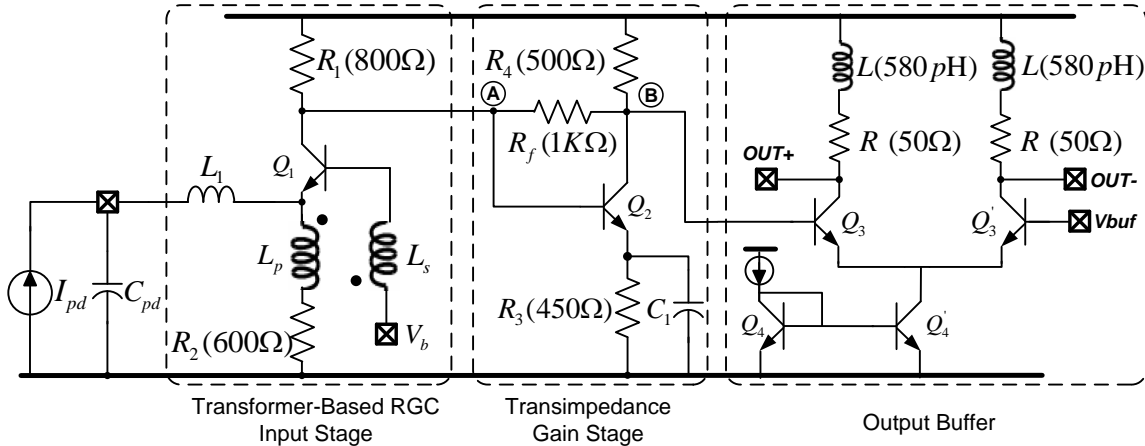


Figure 3.5: Transformer-based RGC TIA schematic.

The complete schematic of the transformer-based regulated cascode TIA is shown in Fig. 3.5. Both series inductive peaking ( $L_1$ ) and input transistor  $g_m$ -boosting via transformer-based negative feedback are leveraged in order to extend TIA bandwidth. Table 3.1 gives the key parameters for the input-stage transformer. The  $g_m$ -boosted common-base input stage isolates the photodiode capacitance from the second-stage feedback TIA. This common-emitter gain stage, consisting of ( $Q_2$  and  $R_4$ ) with local shunt feedback resistor ( $R_f$ ) connected between the base and collector terminals, provides the majority of the transimpedance gain.  $R_3$  and  $C_1$  are inserted in the second stage to provide an appropriate level shift for the DC voltage at emitter terminal. The final stage is a differential output buffer which converts the TIA's single-ended output to differential outputs and drives the  $50\Omega$  load of the measurement equipment. Here shunt inductive peaking [37] is also used to achieve broadband operation. As this simple output buffer is not a major point of emphasis in this design, emitter degeneration is not included in the output buffer current mirror.

Table 3.1: Key parameters of the input-stage transformer

$L_p$	$L_s$	Turn Number	Turn Ratio ( $n$ )	Coupling Coefficient ( $k$ ) at 20 GHz
0.28 nH	1.09 nH	2.5	2	0.68

Assuming sufficiently large boosted  $g_{m1}$  and  $g_{m2}R_f$  values, the low frequency transimpedance from the input to node  $B$  of Fig. 3.5 is approximately

$$Z_T(0) \simeq g_{m2}R_4 \left( R_1 \parallel \frac{R_f}{1 + g_{m2}R_4} \right) \quad (3.6)$$

At the input, since the ( $Q_1$ ) transconductance is boosted by the negative feedback from

the on-chip transformer, the resistance at the input node can be expressed as

$$R_{in}(s) \simeq \frac{1}{g_{m1}(1 + nk(s))} \quad (3.7)$$

where  $k(s)$  displays a high-pass response, as shown in Fig. 3.6. With turn ratio  $n = 2$  and coupling coefficient  $k$  near 0.7 at 20GHz, the input resistance can be reduced to a relatively low value. Note, this high-pass coupling coefficient can also be leveraged to compensate for bandwidth degradation caused by the circuits poles.




Figure 3.6: Simulated 2.5-turn transformer coupling coefficient vs frequency with different turn ratios.

The trade-offs between input resistance, noise, and voltage headroom pose challenges in the design of a high-speed TIA based on a single common-base topology. In order to obtain a low input resistance the bias current needs to be large, which, for a given voltage headroom, limits the load resistor  $R_1$  which sets the transimpedance gain and the TIA noise performance [38]. Moreover, a large load resistor reduces the TIA output bandwidth. The proposed TIA architecture alleviates these trade-offs by utilizing a transformer-based input



stage that enhances the transconductance without increasing the bias current, thus reducing the input resistance. At node  $A$  of Fig. 3.5, the local shunt feedback lowers the node resistance down by the factor of the open loop gain of the second amplifier stage. Although the relatively low impedance caused by the local resistive shunt feedback sacrifices the transimpedance gain of first stage, it overcomes the bandwidth limit due to the large output impedance of the simple common-base topology. Finally, the effective resistance at node  $B$  is  $\frac{1}{g_{m2}} \parallel R_4$ , which is inherently small. Due to this TIA topology, all major signal path poles reside at relatively high frequencies, making this architecture suitable for wideband, high speed applications.

Since passive transformer-based negative feedback is utilized to boost the input transconductance, the noise penalty of the active feedback amplifier in the conventional RGC input stage is avoided. Considering transistor collector current shot noise and base resistance thermal noise [39], the input-referred noise current can be derived as

$$\begin{aligned} \overline{i_{n,in}^2} \simeq & \frac{4kT}{R_1} + \frac{4kT}{R_2} + \frac{4kT}{R_f} \\ & + 2kT \left( \frac{1}{g_{m2}} + 2r_{b2} + \frac{2}{g_{m2}^2 R_4} \right) \omega^2 C_A^2 \\ & + 4kT \left( \frac{g_{m1}}{2} + r_{b1} g_{m1}^2 + \frac{1}{R_f} + \frac{1}{R_1} \right) \frac{\omega^2 C_{in}^2}{g_{m1}^2} \end{aligned} \quad (3.8)$$

where  $r_{b1}$  and  $r_{b2}$  are the base resistance of  $Q_1$  and  $Q_2$ , respectively. Here,  $g'_{m1}$  is the boosted transconductance and  $C_{in}$  and  $C_A$  are the total parasitic capacitance at the input node and  $Q_1$ 's collector, respectively. From (3.8), the boosted  $g'_{m1}$  value, which is a function of the feedback from the transformer comprised of  $L_p$  and  $L_s$ , should also provide a reduction in the last noise term at no power overhead.

### 3.4.2 Bandwidth Extension Analysis

While on-chip inductors and transformers can be used to enhance broadband amplifier bandwidth and overcome a given process's transimpedance limit, improperly designed inductor values can cause frequency peaking and lead to relatively large group delay variation and signal distortion. This has been well studied in the work of [31] and [40]. In this subsection, we model the frequency response of the proposed TIA, neglecting the output buffer (Fig. 4.4a), in order to select the series inductance value and transformer design parameters. We extend the approach of [40] for TIA modeling with series inductive peaking to include both the frequency-dependent response of the transformer-based  $g_m$ -boosting and a more accurate 2-pole and 1-zero feedback TIA model.

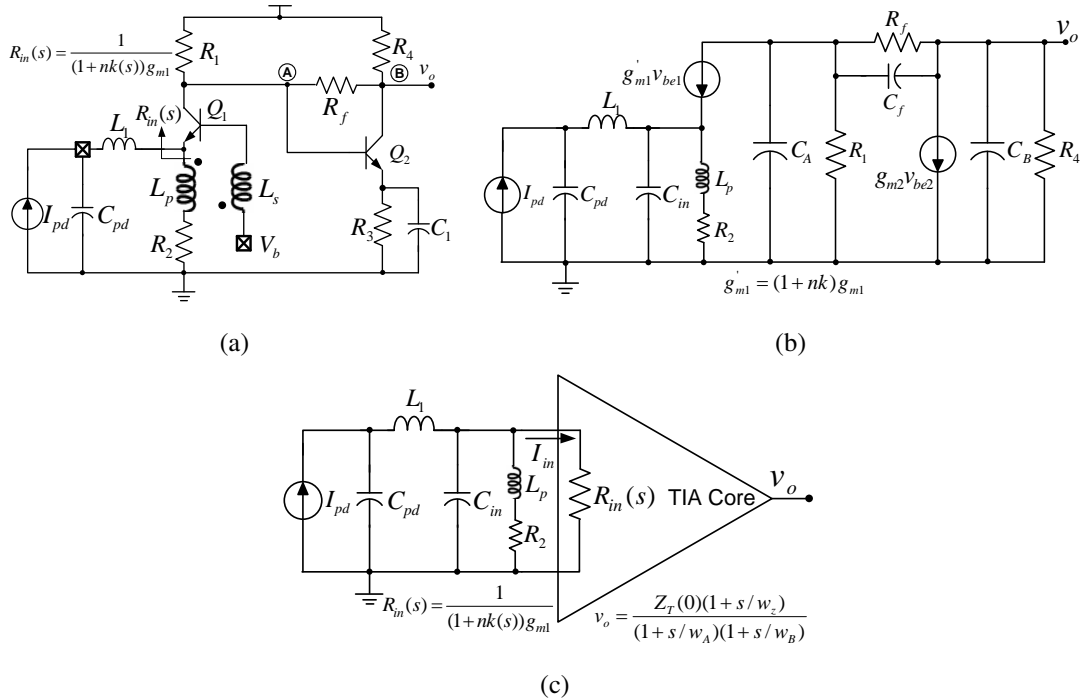


Figure 3.7: (a) TIA schematic without the output buffer, (b) Equivalent small-signal model, (c) Equivalent analysis model.

Fig. 3.7b shows the equivalent TIA small-signal model. The equivalent resistance seen into the emitter of transistor  $Q_1$  is taken from (3.3). As shown in Fig. 3.7c, the proposed TIA is then simplified to a passive  $\pi$ -network followed by the feedback TIA model. Including the transformer's primary inductor, the transfer function of this  $\pi$ -network can be written as a fourth-order expression. The feedback TIA is modeled with two poles located at nodes  $A$  and  $B$  in Fig. 3.7a and a zero from the parasitic capacitance in parallel with the local feedback resistor. Overall, the complete TIA transfer function is approximated as (3.9),

$$Z_T(s) \simeq \frac{R_1 R_4}{R_1(1 + g_{m2} R_4) + R_f} \times \frac{(s R_f C_f + 1 - g_{m2} R_f)}{(s R_A C_A + 1)(s R_B C_B + 1)} \times \frac{L_p s + R_2}{a_4 s^4 + a_3 s^3 + a_2 s^2 + a_1 s + a_0} \quad (3.9)$$

where

$$a_0 = R_2 + R_{in}(s),$$

$$a_1 = R_2 R_{in}(s)(C_{pd} + C_{in}) + L_p,$$

$$a_2 = L_1 C_{pd}(R_2 + R_{in}(s)) + R_{in}(s) L_p (C_{pd} + C_{in}),$$

$$a_3 = L_1 R_2 R_{in}(s) C_{pd} C_{in} + L_1 L_p C_{pd},$$

$$a_4 = L_1 L_p C_{pd} C_{in} R_{in}(s),$$

$$R_A = R_1 \parallel \frac{R_f}{1 + g_{m2} R_4},$$

$$R_B = \frac{1}{g_{m2}} \parallel R_4,$$

$$C_A = C_{cs1} + C_{be2} + C_f(1 + g_{m2} R_4) + C_{bc1},$$

$$\text{and } C_B = C_f + C_{cs2}.$$

Here  $C_{pd}$  denotes the parasitic photodiode capacitance and the bond pad capacitance,  $C_{cs1}$  and  $C_{bc1}$  are the  $Q_1$  collector-substrate and base-collector capacitances, respectively,  $C_{be2}$  and  $C_{cs2}$  are the  $Q_2$  base-emitter and collector-substrate capacitances, respectively, and  $C_f$  is the depletion capacitance of the collector-base junction of  $Q_2$ .

Using this model, the series peaking inductance and transformer design parameters are chosen for a flat frequency response, low group delay variation, and low deterministic jitter. The series inductance is selected to achieve a Butterworth response with maximally flat gain magnitude and the total TIA's response is optimized by varying the transformer turn number and turn ratio. As  $L_1$ ,  $L_p$ , and  $L_s$  need to be jointly optimized, this iterative process is outlined in the following steps.

*Step 1* : Using initial reasonable transformer parameters (e.g. 2 turns and turn ratio of 2:1), optimize the series inductance  $L_1$  for reasonable bandwidth extension, low jitter and group delay variation;

*Step 2* : Using the  $L_1$  value found in Step 1 and the initial turn ratio, optimize the transformer turn number;

*Step 3* : Using the  $L_1$  value found in Step 1 and turn number found in Step 2, optimize the transformer turn ratio;

*Step 4* : Using the transformer parameters found in Steps 2 and 3, re-optimize the series inductance;

*Step 5* : If necessary, re-optimize the transformer parameters and finalize the design.

The following sub-sections provide key design insights on how the series inductance and transformer parameters impact the TIA performance, with the assumption for each of the parameters that the other design parameters are already optimized.

#### 3.4.2.1 Series Peaking Inductance

In order to achieve a flat frequency response, low group-delay variation, and low deterministic jitter, the inductance of  $L_1$  in Fig. 3.5 needs to be carefully selected. The simulated TIA transimpedance frequency response is shown in Fig. 3.8a for various  $L_1$  inductance values, with a finite- $Q$  inductor model employed, and a 220fF capacitance to model the photodetector and input bondpad. Also, an initial transformer design with 2

turns and a 2:1 turn ratio is assumed. Here both the transimpedance gain is normalized to one and the frequency axis is normalized to the 3-dB bandwidth without series inductive peaking ( $L_1=0$ ). A Butterworth response with maximally flat magnitude and  $1.8\times$  bandwidth extension is achieved when employing a proper series inductor value of  $L_1=820$   $pH$ . Note that higher inductance values also cause peaking in the frequency response, thus leading to relatively large group-delay variations, as shown in Fig. 3.8b. The chosen series inductor value of  $L_1=820$   $pH$  and  $Q$  of  $\sim 8$  achieves a low group delay variation of  $\pm 10\%$  and, as shown in Fig. 3.9, minimal deterministic jitter with a 40Gb/s  $2^{31}-1$  PRBS pattern. Post-layout simulations indicate that this series inductance value is suitable for photodetector capacitance variations near  $\pm 20\%$ , while still maintaining  $< 1$  dB gain peaking and 5% bandwidth degradation. While an octagonal-shaped inductor is employed in the final design, this geometry choice is not essential, as post-layout simulations indicate that a  $Q$  of approximately three can be used without degrading the bandwidth more than 10%.




Figure 3.8: Simulated TIA frequency response with various series inductance values: (a) normalized transimpedance gain, (b) group-delay of input  $\pi$ -network. The frequency axis in both curves is normalized to the 3-dB bandwidth without series inductive peaking.

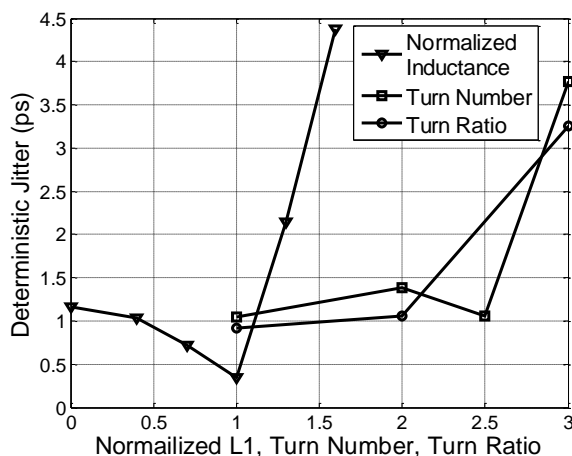


Figure 3.9: Simulated 40 Gb/s deterministic jitter performance of the proposed TIA with a  $2^{31}-1$  pattern.

Realizing  $L_1$  fully with an on-chip series peaking inductor value is directly applicable for an optical receiver with monolithically integrated photodetectors [41]. For optical receivers which have off-chip photodetectors, a portion of the series peaking inductor  $L_1$  could be realized with the bondwire inductance between the photodetector and the TIA input pad. In this case, a smaller on-chip peaking inductor could still be included to isolate the bond pad capacitance from the TIA input capacitance for further bandwidth extension [31].

### 3.4.2.2 Transformer Turn Number

The total TIA response is optimized by setting the transformer turn number and ratio. Using the  $L_1 = 820 \text{ pH}$  value to optimize the input  $\pi$ -network and assuming an initial 2:1 transformer turn ratio, the turn number is varied to observe how the change in coupling coefficient affects the TIA's frequency response. Here the transformer area is increased in order to increase the turn number, which results in increased parasitic resistance and capacitance. For example, the 3-turn transformer was designed by adding an extra turn to

the 2-turn design, incurring a 72% area increase. As shown in Fig. 3.10, increasing turn number allows for bandwidth extension up to a point. However, when the transformer becomes large, as in the 3-turn case, the incurred parasitics cause a steep roll-off in the frequency response. A 2.5 turn number allows for a maximally flat bandwidth response, and, as shown in Fig. 3.9, minimal deterministic jitter with a 40Gb/s  $2^{31}$ -1 PRBS pattern.

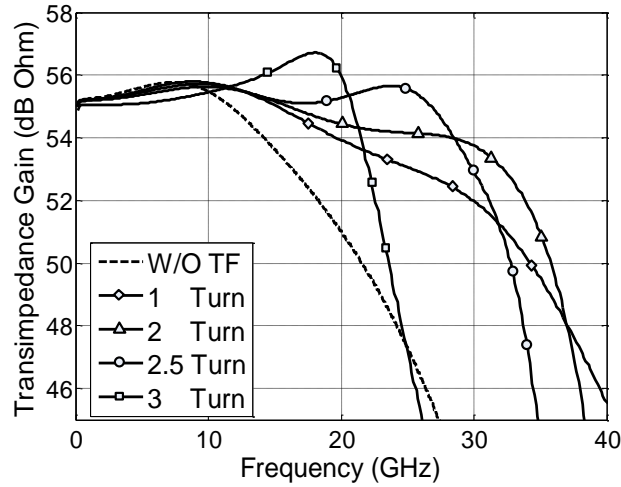


Figure 3.10: Simulated transimpedance frequency response with different transformer turn number, transformer turn ratio is fixed at  $n = 2$ .

### 3.4.2.3 Transformer Turn Ratio

Transformer turn ratio is another important parameter which sets the amount of input transistor  $g_m$ -boosting. Using a 2.5-turn transformer value, the turn ratio is optimized for maximum bandwidth enhancement and minimum magnitude variation. As shown in Fig. 3.11, increasing turn ratio allows for bandwidth extension due to increased input transistor transconductance. However, again due to transformer size issues, the incurred parasitics cause excessive frequency peaking and a steep roll-off with  $n = 3$ . Also, as shown in Fig. 3.9, a large increase in deterministic jitter is observed for a turn ratio larger than two.

The final transformer design uses  $n = 2$  and 2.5 turns, which allows for a simulated TIA -3dB-bandwidth of 32 GHz.




Figure 3.11: Simulated transimpedance frequency response with different transformer turn ratio, transformer turn number is fixed at 2.5.




Figure 3.12: Simulated TIA performance versus series inductance  $L_1$  for different transformer turn ratios: (a) bandwidth, (b) group delay variation. Here the series inductance is normalized to the optimum value of 830  $pH$ .



As mentioned previously, the overall design procedure is an iterative process to optimize the series peaking inductance and key transformer parameters. Fig. 3.12 shows how the TIA bandwidth and group delay vary over a more complete design space of various series inductor values, normalized to the optimum  $820 \text{ pH}$  value, and transformer turn ratios. Overall, a turn ratio of two yields the maximum bandwidth and minimum group delay variation. Note that while a smaller value relative to the chosen  $820 \text{ pH}$  series inductor yields a potentially wider bandwidth, this would result in sub-optimum group delay variation.

### 3.4.3 Transformer Design

Trade-offs exist in the design of the wires that comprise the transformer windings. For a transformer with a given number of turns, magnetic flux coupled between windings increases as the metal width decreases. However, if the metal width is too narrow, this may lead to excessive ohmic losses. On the contrary, a larger width comes at the cost of relatively higher parasitic capacitance.




Figure 3.13: Monolithic transformer used for input stage  $g_m$ -boosting.

The monolithic transformer used in this work is shown in Fig. 3.13, where the two sections of primary winding are connected in parallel to form a 2:1 transformer and an inverting configuration is implemented in order to form the negative feedback for the transconductance boosting in the input stage. A simple square shape is utilized in order to more accurately control the turn ratio. As critical electromagnetic effects, such as substrate eddy currents and frequency-dependent metal loss need to be considered in the transformer design, an electromagnetic simulator, SONNET, is used to model the transformer. Each process layer is accurately modeled in SONNET according to the specifications in the 0.25- $\mu\text{m}$  SiGe BiCMOS process design kit. In order to reduce the parasitic effects which cause high-frequency loss, the transformer is realized with the top metal layer (M6) which is the thickest and farthest from the substrate. The other metal layers (M3, M5) are used to facilitate convenient connections to other circuitry. Considering the wideband application of this design, the smallest linewidth ( $5\mu\text{m}$ ) is used to reduce the parasitic capacitance and the minimum line spacing ( $3\mu\text{m}$ ) allowed in the technology is used to achieve the best magnetic coupling. A grounded polysilicon bar shield is also introduced to provide increased isolation to the transformer from the silicon substrate, at the cost of increasing the capacitance to the shield to  $10.5\text{ fF}$  from a potential value of  $7.8\text{ fF}$  to the substrate. Note, as this design was implemented in a process with a high-resistivity substrate, the impact of the ground shield has a negligible impact on the quality factors of the inductors used to implement the transformer.




Figure 3.14: Simulated transformer coupling coefficient at 20 GHz vs turn number with different turn ratios.

In addition to the physical size and spacing of metal wires, the number of turns is another important factor that affects the coupling coefficient. Fig. 3.14 shows the simulated transformer coupling coefficient  $k$  at 20 GHz versus turn number and for different turn ratios. While a large coupling coefficient improvement is observed from one to three turns, it tends to saturate around 0.8 with further increase in turn number. This is due to increased magnetic coupling between adjacent lines causing a large improvement in  $k$ -factor as the turn number is increased from one to two. However, as the turn number is increased further, the separation between the multiple parallel conductors increases and causes the  $k$ -factor to saturate near turn values of three to four [42]. Note that for multi-turn designs the turn ratio has little impact on the coupling coefficient in the 26 GHz frequency range of interest, with a lower turn ratio displaying only marginally better performance. However, at higher frequencies the  $k$ -factor of the 1:1 transformer falls off faster due to its lower self-resonant frequency.

It is seen from (3.3), that increasing the turn ratio between the secondary and primary windings can also boost the TIA input transconductance. For a monolithic transformer

implemented by conductors interwound in the same plane, the mutual inductance of the transformer increases with the length of each winding. The simulated coupling coefficient of a 2.5-turn transformer versus frequency and for different transformer turn ratios is shown in Fig. 3.6, where the frequency response displays a high-pass shape that is mostly independent of turn ratio over the frequency range of interest. Therefore, a transformer with a turn ratio larger than one can be implemented with a single top-layer metal by sectioning the continuous primary winding into multiple individual turns [42]. While a larger turn ratio can further boost input transconductance, a transformer with turn ratio 2:1 is chosen in this work to balance TIA bandwidth extension and group-delay variation.

### 3.5 Experimental Results

The TIA was fabricated in a 0.25- $\mu\text{m}$  SiGe BiCMOS technology with bipolar npn transistor peak  $f_t$  of 137 GHz. As shown in the die photograph of Fig. 3.15, the total chip area including pads is  $960\ \mu\text{m} \times 780\ \mu\text{m}$ . The chip is encapsulated in a QFN package only for connection of external power supply and DC voltage biases.

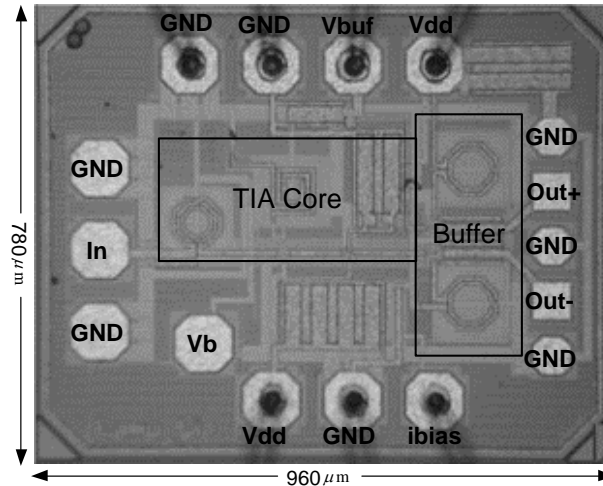


Figure 3.15: Die photograph.

High frequency testing is performed with the package cover removed and the high-speed input/output signals and transistor base bias voltage applied via probing. Thus, there are no wire bonds on these critical signals. The S-parameters shown in Fig. 3.16 are measured using an Agilent N5230A network analyzer. In order to emulate the parasitic capacitance of a potential photodiode, a  $150\text{ fF}$  on-chip metal capacitor is included in the design. This capacitance, along with the  $70\text{ fF}$  bondpad, yields a total effective  $C_{PD}$  of  $220\text{ fF}$ . Due to equipment availability, testing is performed in single-ended mode with the unused ports terminated at  $50\ \Omega$ . The differential transimpedance gain is calculated from the measured S-parameters based on

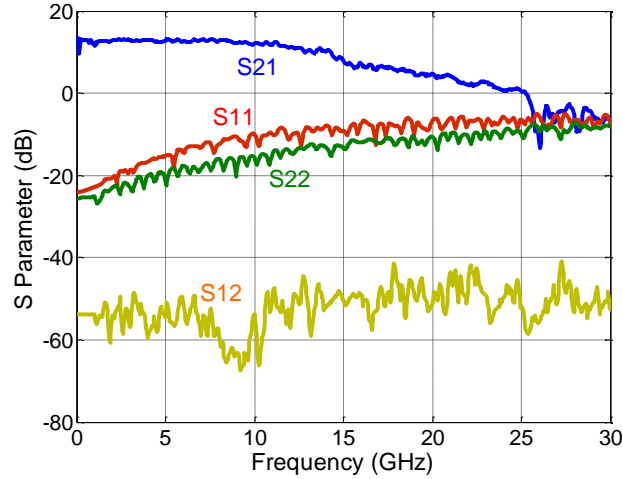


Figure 3.16: Measured TIA single-ended S-parameters.

$$Z_t = \frac{2S_{21}}{(1 - S_{11})(1 - S_{22}) - S_{21}S_{12}} \times Z_0, \quad (3.10)$$

where  $Z_0$  is  $50\ \Omega$  [43]. A single-ended transimpedance gain of  $53\text{ dB}\Omega$  with a -3dB frequency of  $26\text{ GHz}$  is achieved, as shown in Fig. 3.17. The TIA group delay is also extracted from the measured S-parameters, with group delay variation below  $19\text{ ps}$  from

near DC to 26 GHz. With data encoding schemes, the low-frequency cut-off frequency of interest is increased. If group delay values below 5 GHz are neglected [31], then the group delay variation is only 10 ps.




Figure 3.17: Single-ended simulated/measured transimpedance gain and measured group delay.

Fig. 3.18 illustrates the measurement setup for the high-speed eye diagram test. Two uncorrelated 13.5 Gb/s  $2^{15}-1$  bit sequences from a pattern generator (Agilent N4872A) are multiplexed by a 2:1 MUX (SHF 23210A) to form a 27 Gb/s data sequence. A programmable step attenuator is then used to attenuate the input signal to an appropriate range for the TIA. The 27 Gb/s input/output signals pass through external DC blocks and are applied to the TIA via probing. Note that while this test setup is not optimal for TIA characterization, as the  $50 \Omega$  system has the potential to artificially increase the TIA bandwidth, this does not have a dramatic increase in the experimental results due to the transformer-based negative feedback at the input node reducing the TIA input impedance to near  $15 \Omega$ . Post-layout simulations indicate only a 2.7% bandwidth differential with the  $50 \Omega$  measurement system due to the low TIA input impedance and the contribution

of other TIA poles setting the overall bandwidth. An improved test setup which better emulates a photodiode would include a large series resistor at the TIA input [44].

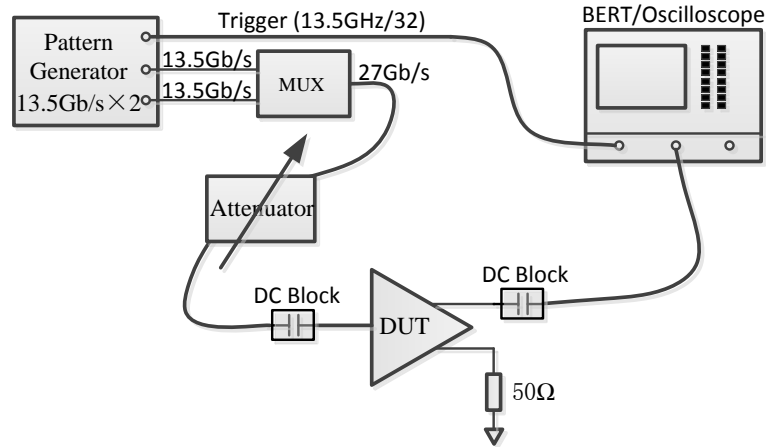


Figure 3.18: Measurement setup for eye diagram and BER test.

Fig. 3.19 shows a single-ended  $2^{15}-1$  PRBS eye diagram obtained at 27 Gb/s with an estimated input current of  $125 \mu A_{pp}$ . At this data rate the TIA output displays healthy margins, with the output rise/fall time mainly limited by the edge rate of the input signal, rather than the TIA circuit. Additionally, some bandwidth limitation is introduced by the microwave attenuators, connectors and cables. The TIA only introduces a small amount of jitter relative to the original input signal, as the measured output peak-to-peak jitter is 17 ps with an input signal that has 14 ps peak-to-peak jitter. While the available equipment limited testing at higher data rates, a post-layout simulated single-ended eye diagram at 40 Gb/s with a  $100 \mu A_{pp}$   $2^{31}-1$  PRBS pattern, shown in Fig. 3.20, also achieves minimal vertical and horizontal eye closure.




Figure 3.19: Measured 27 Gb/s single-ended eye-diagram with a  $125 \mu A_{pp}$   $2^{15}-1$  PRBS input signal.




Figure 3.20: Post-layout simulated single-ended 40 Gb/s eye-diagram of the proposed TIA with  $100 \mu A_{pp}$  input current.



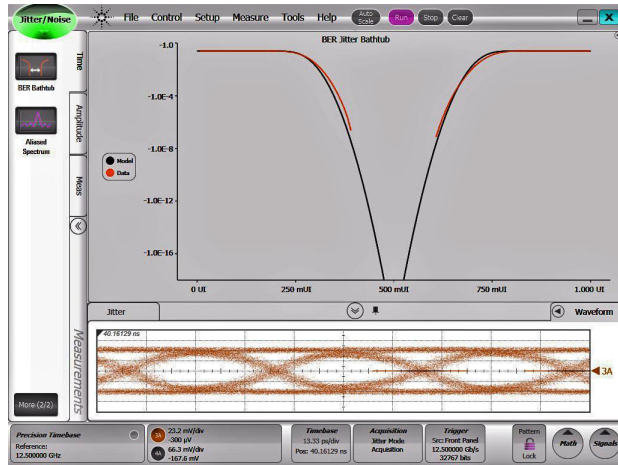


Figure 3.21: BER jitter bathtub plot with 25 Gbps 150  $\mu A_{pp}$  PRBS  $2^{15}-1$  input.

The experimental setup of Fig. 3.18 was also used to perform BER testing. Fig. 3.21 shows a 25 Gb/s BER bathtub curve, with approximately 20% UI timing margin for a 150  $\mu A_{pp}$  input signal. A sensitivity of 93  $\mu A_{pp}$  is achieved at a BER of  $10^{-12}$ , as shown in Fig. 3.22.




Figure 3.22: Measured BER versus input current.




Figure 3.23: Measured TIA single-ended integrated output noise.

The integrated single-ended output noise, measured via the oscilloscope histogram function with the absence of any input signal source, is shown in Fig. 3.23. After deconvolving the inherent oscilloscope noise of  $0.476 \text{ mV}_{rms}$ , the single-ended integrated output noise of the TIA is estimated at  $1.54 \text{ mV}$ . The integrated input-referred noise of the TIA's differential output can be calculated as

$$I_{n,in} = \frac{2\sqrt{(1.6125\text{mV})^2 - (0.476\text{mV})^2}}{59\text{dB}\Omega} = 3.45\mu\text{A}_{rms}, \quad (3.11)$$

and the average input-referred noise current density is

$$I_{n,in,avg} = \frac{I_{n,in}}{\sqrt{BW}} = 21.3\text{pA}/\sqrt{\text{Hz}}. \quad (3.12)$$

This measured result matches well with the simulated input-referred noise current density, shown as the solid curve in Fig. 3.24, which is below  $20 \text{ pA}/\sqrt{\text{Hz}}$  up to 30 GHz. Fig. 3.24 also compares the simulated input-referred current noise with the expressions for the proposed transformer-based RGC TIA (3.8) and a simple common-base input stage. The noise for a simple common-base input stage without any gain boosting increases at a rel-

atively low frequency due to a reduced input pole frequency which effectively amplifies the input-referred current noise. Whereas the increased input bandwidth provided by the transformer-based RGC topology reduces the frequency-dependent slope of the last term in (3.8), and allows for a slower increase in the noise current with frequency. Note that the calculated noise level of the proposed transformer-based RGC TIA is lower than the simulated results at low frequencies due to neglecting the noise of the output buffer.




Figure 3.24: Simulated and calculated input-referred current noise density for the proposed transformer-based RGC input-stage TIA and a simple common-base input-stage TIA.

Table 3.2 compares this work with recent TIA designs. The use of passive transformer-based negative feedback to reduce the input resistance and extend TIA bandwidth allows the design to achieve a comparable transimpedance gain with a power consumption of only  $28.2\text{ mW}$  from a  $2.5\text{ V}$  power supply, of which the TIA core is only  $8.2\text{ mW}$ . Also, since the passive transformer contributes little noise, superior noise performance is achieved.

Table 3.2: TIA performance comparisons

	[25]	[45]	[31]	This Work
BW (GHz)	28	42	29	26
Gain (dB $\Omega$ )	53.6	65	50	59 (diff)
Noise (pA $\sqrt{Hz}$ )	36.5	34.2	51.8	21.3
GDV (ps)	NA	10	16	13
Power (mW)	110.0	600.0	45.7	28.2
Area (mm <sup>2</sup> )	0.56	1.0	0.4	0.75
Architecture	RGC	CE	CS	Transformer based RGC
Technology	0.13 $\mu$ m- BiCMOS	InP- InGaAs	0.13 $\mu$ m- CMOS	0.25 $\mu$ m- BiCMOS
$f_t$ (GHz)	160	150	85	137

### 3.6 Summary

This chapter describes a TIA design which employs two key bandwidth extension techniques; input  $g_m$ -boosting via transformer-based negative feedback and series inductive peaking. Utilizing mutual magnetic coupling of a passive on-chip transformer to provide negative feedback between the emitter and base terminals of a common-base input stage provides  $g_m$ -boosting in the input common-base stage, without the power and noise penalties associated with a conventional regulated cascode topology. Further bandwidth extension is achieved with series inductive peaking at the TIA input. This series inductive peaking value and key transformer design parameters were selected to obtain a broadband flat frequency response, low group delay variation, and low deterministic jitter. These techniques allow for relaxed voltage headroom, low power, improved system noise performance, and high bandwidth operation, making the proposed topology suitable for high-speed, low power, and low noise applications.

#### 4. SILICON RING RESONATOR TRANSCEIVER DESIGN\*

Realizing the required intra-chip bandwidth is difficult due to on-chip wires being limited by resistance-capacitance (RC) time constants that increase with each new technology node, resulting in shorter repeater distances with CMOS technology scaling, which can severely degrade the efficiency of cm-length global interconnects. Recently, photonic interconnect has emerged as a potential replacement for the conventional electrical inter/intra-chip interconnect. Optical channels provide the potential to overcome key interconnect bottlenecks and greatly improve data transfer efficiency due to their flat channel loss over a wide frequency range and also relatively small crosstalk and electromagnetic noise [46]. Another important feature of optical interconnects is the ability to combine multiple data channels on a single waveguide via wavelength-division-multiplexing (WDM) and greatly improve bandwidth density. In order to take advantage of these attractive properties, silicon photonic platforms are being developed to enable tightly integrated optical interconnects and future photonic interconnect network architectures [9, 14–20, 47–55]. One promising photonic device is the silicon ring resonator [9, 11, 17–19], which can be configured either as an optical modulator or WDM drop filter. Silicon ring resonator modulators/filters offer advantages of small size, relative to Mach-Zehnder modulators [14, 15], and increased filter functionality, relative to electro-absorption modulators [16]. Silicon photonic links based on ring resonator devices provide a unique opportunity to deliver distance-independent connectivity whose pin-bandwidth scales with the degree of wavelength-division multiplexing. As shown in Fig. 2.8, multiple wavelengths ( $\lambda_1$ - $\lambda_4$ ) generated by an off-chip laser are coupled into a silicon waveguide via an optical cou-

---

\*Reprinted with permission from "A ring-resonator-based silicon photonics transceiver with bias-based wavelength stabilization and adaptive-power-sensitivity receiver" by Cheng Li, IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2013, Page(s): 124 - 125, Copyright 2013 by IEEE

pler. At the transmit side, ring modulators insert data onto a specific wavelength through electro-optical modulation. These modulated optical signals propagate through the waveguide and arrive at the receiver side where ring filters drop the modulated optical signals of a specific wavelength at a receiver channel with photodetectors (PD) that convert the signals back to the electrical domain.

This chapter presents silicon photonic transceiver circuits for a ring resonator-based optical interconnect architecture that address limited modulator bandwidth, variations in ring resonator resonance wavelength and link budget, and efficient receiver clocking. Section 4.1 introduces the basics of silicon ring resonator and the design considerations for the silicon ring based photonic interconnects. The architecture of the transceiver circuits prototype is outlined in Section 4.3. Section 4.4 describes transmitters with independent dual-edge pre-emphasis to compensate for the bandwidth limitations of the carrier-injection microring resonators used in this work. A novel bias-based resonance wavelength stabilization scheme for the modulators that offers advantages in tuning speed at comparable efficiency levels is presented in Section 4.5. Section 4.6 discusses an optical forwarded-clock adaptive sensitivity-power receiver that accommodates variations in input capacitance, modulator/photodetector performance, and link budget. Experimental results of the transceiver circuits prototype, fabricated in a 65nm CMOS technology and integrated via wire-bonding to photonic devices, are presented in Section 4.7. Finally, Section 4.8 summarizes the chapter.

#### 4.1 Silicon Ring Resonator Based Photonic Interconnect Design Considerations

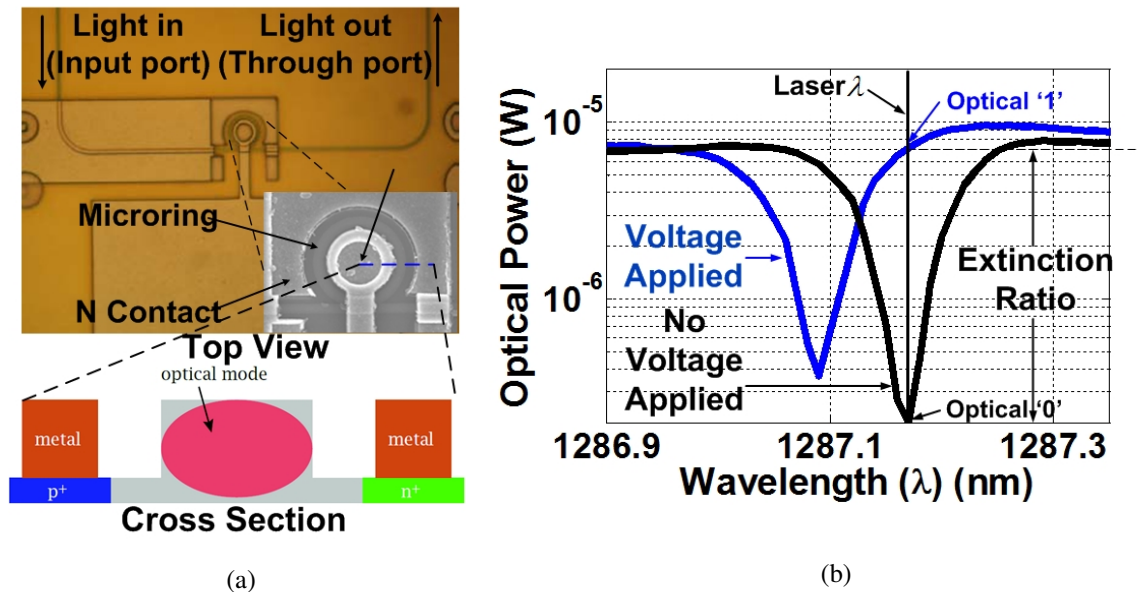


Figure 4.1: (a) Top and cross section views of carrier-injection silicon ring resonator modulator, (b) optical spectrum at through port.

A basic silicon ring resonator consists of a straight waveguide coupled with a circular waveguide, as shown in Fig. 4.1a. Input light at the resonance wavelength mostly circulates in the circular waveguide, with only a small amount of optical power observed at the through port, resulting in the ring's spectrum at the through port displaying a notch-shaped characteristic (Fig. 4.1b). This resonance wavelength of the ring device is periodic, repeating over a free spectral range (FSR), and can be shifted by changing the effective refractive index of the waveguide through the free-carrier plasma dispersion effect [19]. Two common implementations of silicon ring resonator modulators include p-i-n junction-based carrier-injection devices [11, 17], operating primarily in forward-bias, and

carrier-depletion devices [18], operating primarily in reverse-bias. Although a depletion ring generally achieves higher modulation speeds relative to a carrier-injection ring due to the ability to rapidly change the depletion width, its modulation depth is limited due to the relatively low doping concentration in the waveguide to avoid excessive loss. In contrast, carrier-injection ring modulators can provide large refractive index changes and high modulation depths, but are limited by long minority carrier lifetimes.




Figure 4.2: Measured quality factor and resonance wavelength of nine  $2.5\mu\text{m}$  radius silicon ring modulators fabricated on an  $8''$   $130\text{nm}$  CMOS SOI wafer.

While ring-resonator-based photonic interconnects have the potential to offer both improved power efficiency and bandwidth density, reliability and robustness are major barriers to widespread adoption of ring-based silicon photonics [20]. A key challenge is the variation in resonance wavelength with temperature changes and fabrication tolerances. For example, Fig. 4.2 shows that while a high quality factor is maintained for nine  $2.5\mu\text{m}$  radius ring resonators spread across an  $8''$   $130\text{nm}$  CMOS-compatible silicon-on-insulator (SOI) wafer, the  $4.96\text{nm}$  resonance wavelength variation implies the need for a potentially



wide resonance tuning range for robust operation. In order to relax this, system-level WDM channel-shuffling techniques are proposed that reduce the tuning to the order of  $FSR/N$ , where  $N$  is the WDM channel number [20, 52]. A commonly proposed resonance wavelength tuning technique is to adjust the device's temperature with a resistor implanted close to the photonic device to heat the waveguide, thus changing the refractive index [53, 54]. One potential issue with this approach is that the tuning speed, which is limited by the device thermal time constant (ms), may necessitate long calibration times. Also, tuning power overhead can degrade overall link power efficiency [18, 54].

Achieving reliable and efficient operation in silicon photonic interconnect systems with large variations in link budget components, such as photonic device properties and interface parasitics, is another important consideration. The link budget determines the receiver sensitivity, with various front-end circuits proposed for optical interconnects, such as regulated-cascode transimpedance amplifiers (TIA) [24, 55, 56], feedback TIAs [53, 55, 57], and integrating topologies [58, 59]. Achieving the necessary receiver sensitivity for a given bit-error rate (BER) at the maximum data rate and under a worst-case link budget scenario generally sets these optical receiver circuits' power consumption, and can lead to sub-optimal power efficiency at lower data rates and in the common presence of link budget margin. One efficient approach to optimize receiver power efficiency versus data rate is to utilize supply-scaling with CMOS inverter-based feedback TIAs [57]. However, in order to leverage this approach for large channel-count systems, efficient control loops with per-receiver voltage regulators are required that allow for self-adaptation to the desired data rate and link budget conditions. While efficient clocking architectures for receiver-side data retiming and de-serialization are often neglected in optical interconnect designs [55, 57], they are necessary to form a complete link. One approach is to utilize a continuously-running clock-and-data recovery (CDR) system [58] which allows the potential for plesiochronous operation between the transmitter and receiver. Howev-

er, this generally consumes more power and area relative to mesochronous architectures which only require periodic training to optimize the receiver sampling position [53]. For mesochronous architectures, key considerations include achieving efficient receiver-side clock generation and sufficient jitter tracking of the incoming data to achieve the desired BER.

## 4.2 Silicon Ring Resonator Modeling

## 4.3 Transceiver Architecture




Figure 4.3: Photonic transceiver circuits prototype block diagram.

Fig. 4.3 shows a block diagram of the CMOS photonic transceiver circuits prototype, with six transmitter and five receiver modules integrated in a  $2\text{mm}^2$  65nm CMOS die. At the transmitter side, a half-rate CML clock is distributed to the 6 transmitter modules where 8-bit parallel data is multiplexed to the full output data rate before being buffered

by the modulator drivers. Two versions of the drivers are implemented. A differential driver, with approximately 0V average bias level, provides a 4Vpp output swing to allow for high-speed operation, while a single-ended driver provides a 2Vpp output swing on the modulator cathode and utilizes a bias-tuning DAC on the anode for an adjustable DC-bias level. These drivers are wire-bonded to carrier-injection silicon ring resonator modulators (Fig. 4.1), where continuous wavelength light near 1300nm from a tunable laser is vertically coupled into the photonic device's input port. The modulated light is then coupled from the modulator's through port into a single-mode fiber for routing to the bias-based tuning photodetector used to stabilize the resonant wavelength and to the optical receiver modules for high-speed data recovery. At the receiver side, data is recovered by adaptive inverter-based TIA front-ends that trade-off power for varying link budgets by employing on-die eye monitors and scaling the TIA supplies for the required sensitivity. The receive-side sampling clocks are produced from an optically-forwarded quarter-rate clock which is amplified by a fixed-supply TIA before being passed to an injection-locked oscillator which produces four quadrature clocks that are routed to the four receiver data channels.

#### 4.4 Non-Linear Pre-emphasis Modulator Driver Transmitter

While carrier-injection silicon ring modulators are capable of high extinction ratio operation in a low area footprint, the operating speed is limited by relatively slow carrier dynamics in forward-bias and parasitic contact resistance in reverse-bias. Fig. 4.4 shows device model simulation results of the carrier-injection silicon microring modulators used in this work, with positive and negative 200ps pulse responses overlaid. Observe that a simple 2Vpp waveform yields a high extinction ratio, but the optical rise time is excessively long due to the carrier recombination lifetime [47]. Increasing the modulation voltage to 4Vpp dramatically improves the optical rise time at the expense of high-level ringing and a high steady-state charge value. Unfortunately, this large amount of charge results in

a slow optical fall time due to the modulator's series resistance limiting the drift current to extract the excess carriers in the junction. The conflicting requirements for fast rising and falling transitions are addressed through the use of a pre-emphasis waveform [11]. During a rising-edge transition the positive voltage overshoots (2V) for a fraction of a bit period to allow for a high initial charge before settling to a lower voltage (1V) corresponding to a reduced steady-state charge. A similar shaped waveform is used for the falling-edge transition to increase the drift current to extract the carriers. As the rising and falling-edge time constants are different, a non-linear modulation waveform is applied. This work adjusts the amount of over/under-shoot time of the pre-emphasis waveform for a specific modulator, with the rising-edge pre-emphasis pulse typically wider than the falling-edge. Relative to approaches which change the pre-emphasis settings with different voltage levels [48], adjusting the pre-emphasis time allows the optimization of the transient response to be decoupled from the steady-state extinction ratio value.

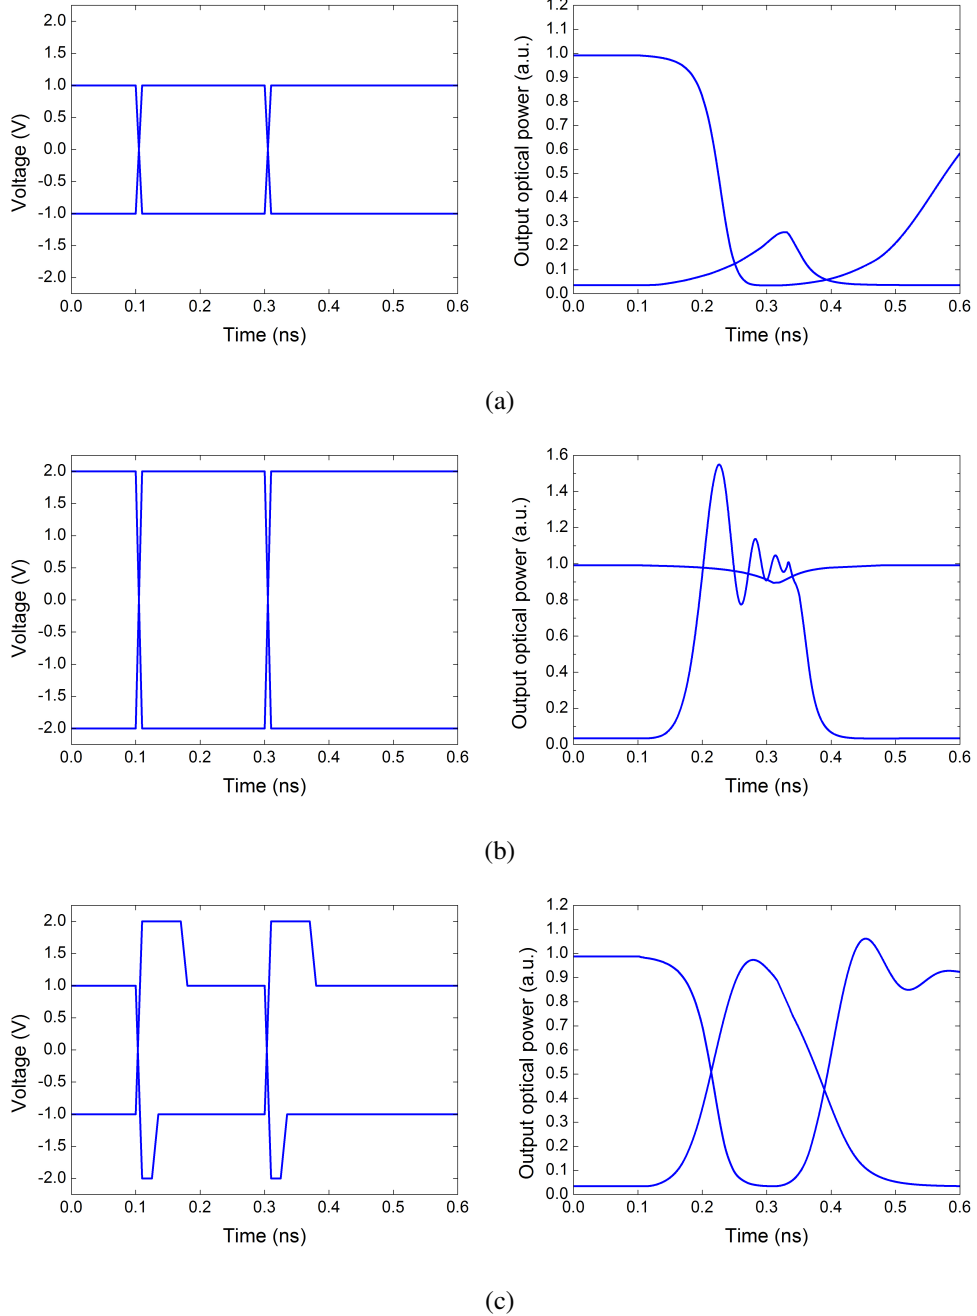


Figure 4.4: Simulated carrier-injection ring resonator modulator response to 200ps data pulses with: (a)  $2V_{pp}$  simple modulation, (b)  $4V_{pp}$  simple modulation, (c)  $4V_{pp}$  modulation with pre-emphasis.

As shown in the block diagram of Fig. 4.5a, two optical transmitter versions are developed to demonstrate high-speed operation with a differential 4V<sub>pp</sub> output driver and explore bias-based modulator tuning capabilities with a single-ended 2V<sub>pp</sub> output driver. Serialization of eight bits of parallel input data is performed in both transmitter versions with three 2:1 multiplexing stages, with the serialization clocks generated from a half-rate CML clock which is distributed to six transmitter modules, converted to CMOS levels, and subsequently divided to switch the mux stages. The serialized data is then transmitted by the modulator drivers, with both output stage versions utilizing a main driver, positive-edge and negative-edge pre-emphasis pulse drivers in parallel (Fig. 4.5b) to generate the pre-emphasis output waveform. Tunable delay cells, implemented with digitally-adjustable current-starved inverters (Fig. 4.5c), allow for independent control of the rising and falling-edge pre-emphasis pulse duration. Finally, pulsed-cascode output stages (Fig. 4.5d) with only thin-oxide core devices [60] reliably provide a final per-terminal output swing of twice the nominal 1V supply. A capacitive level shifter and parallel logic chain generate the signals  $IN_{low}$ , swinging between GND and the nominal VDD, and  $IN_{high}$ , level-shifted between VDD and 2\*VDD, that drive the final pulsed-cascode output stages.




Figure 4.5: Non-linear pre-emphasis modulator driver transmitters: (a) transmitter block diagrams, (b) per-terminal 2V pre-emphasis driver, (c) tunable delay cell, (d) pulsed-cascade output stage.

### 4.5 Automatic Bias-based Wavelength Stabilization

One problem with silicon ring resonators is that their resonance wavelength is sensitive to temperature variation and fabrication tolerances. As shown in Fig. 4.6, a poor extinction ratio results when the modulator's resonance is not aligned with the input continuous-wave

laser wavelength. Hence, a closed-loop adaptation scheme is therefore necessary to stabilize the ring's resonance to match the input laser. Thermal tuning schemes with closely integrated heating resistors [49, 54, 61–63], which red-shift the resonance wavelength as the device is heated up, are commonly proposed for this tuning. However, thermal time constants in the ms-range limit the speed of this tuning approach. Another important consideration is the tuning power efficiency, which varies for thermal tuning depending on the fabrication complexity. Doping a section of the ring waveguide differently to realize a thermal resistor is relatively simple, but has been shown to have a relatively poor  $42\mu\text{W}/\text{GHz}$  tuning efficiency [49]. Improved efficiencies near  $10\text{--}15\mu\text{W}/\text{GHz}$  has been demonstrated using approaches such as substrate removal and transfer for an SOI process [61] and deep-trench isolation for a bulk CMOS process [54]. Finally, superior efficiencies in the  $1.7\text{--}2.9\mu\text{W}/\text{GHz}$  has been achieved with localized substrate removal or undercutting [62, 63], but this comes at the cost of complex processing steps.




Figure 4.6: Ring resonator modulator transmission curves with high and low modulation voltage levels when: (a) resonance wavelength is not aligned with input laser wavelength; (b) resonance wavelength is aligned with input laser wavelength.



Compared with the conventional heater-based tuning approaches, the proposed bias-based tuning method of this work has advantages of fast tuning speed and flexibility in the tuning direction, while displaying comparable tuning efficiency. As shown in Fig. 4.7, increasing the resonator p-i-n diode anode voltage causes the resonance wavelength to blue-shift to shorter wavelengths due to the accumulation of free carriers in the ring waveguide. This provides the potential for a very fast tuning mechanism. While some optical loss and quality factor reduction is observed with increased forward-bias due to the additional carriers, an extinction ratio in excess of 10dB is achieved for a 190mV tuning range.




Figure 4.7: Measured carrier-injection ring resonator modulator performance: (a) optical transmission spectrum at different bias levels; (b) resonance wavelength shift versus bias voltage.

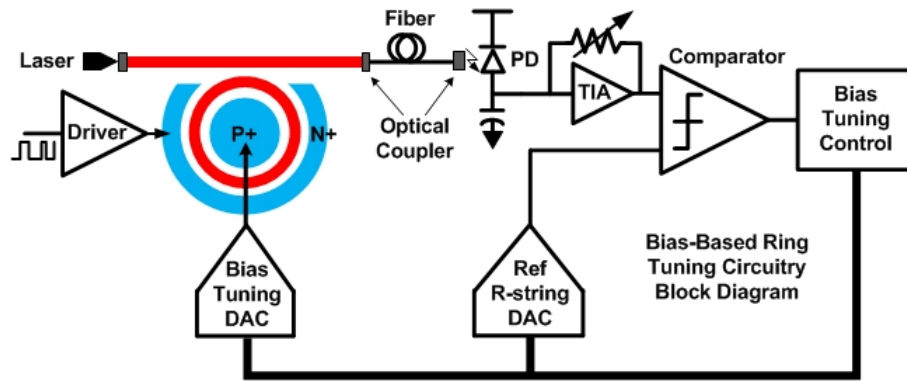


Figure 4.8: Bias-based ring resonator modulator semi-digital wavelength stabilization loop.

Fig. 4.8 shows the semi-digital control loop for the bias-based resonator tuning approach which is utilized with the 2Vpp transmitters. A monitor PD and low-bandwidth TIA is used to sense the average resonator power levels for comparison with a DAC-programmable reference level. This comparator output signal is digitally filtered by a bias tuning control finite-state machine that adjusts the setting of the bias tuning DAC that drives the resonator anode terminal, while the 2Vpp high-speed modulation signal is applied at the resonator cathode terminal. Fig. 4.9 shows the schematic of the 9-bit segmented bias DAC, which utilizes a coarse 3-bit non-linear R-string DAC to match the p-i-n I-V characteristics and a fine 6-bit linear R-2R DAC.

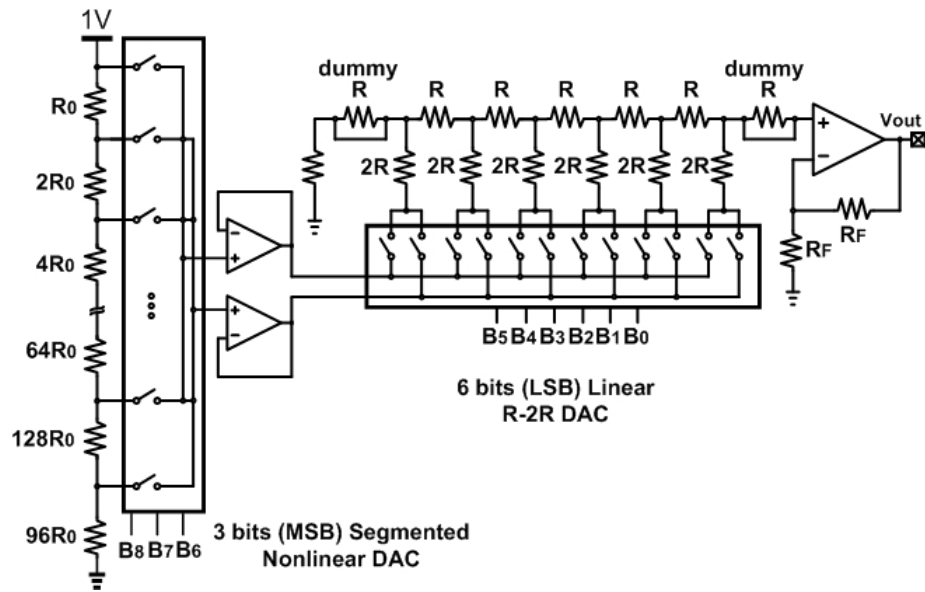


Figure 4.9: 9-bit non-linear bias tuning DAC.

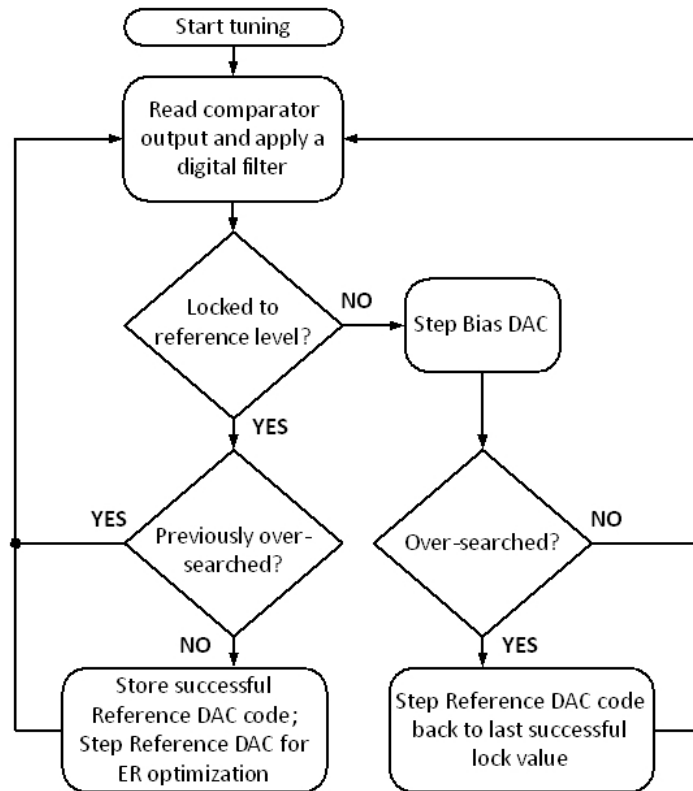


Figure 4.10: Ring resonator bias-based tuning algorithm.

The flowchart of Fig. 4.10 and simulation results of Fig. 4.11 summarize the bias tuning system operation, which can operate both in a static tuning mode with a constant maximum forward-bias across the modulator and a dynamic tuning mode with the modulator driven with a random data signal. For simplicity, first consider the static tuning of Fig. 4.11a. The tuning system works by initially locking the monitor PD and low-bandwidth TIA output to a conservative reference DAC voltage that maps to a reliable point on the resonator curve. After an initial lock is achieved, the reference DAC code is saved as a successful lock point and adjusted in order to maximize the extinction ratio. Since the current system monitors the modulator through-port power, the objective is to minimize the received power to obtain the maximum extinction ratio. As illustrated by the stair-step curves in the simulation results, the tuning procedure then undergoes several cycles of locking and reference level adjustment until the loop "over-searches" and can no longer lock to a minimum power point. The system then steps back to the last successful reference level to obtain the final lock point near the target resonance wavelength and achieve an extinction ratio near 10dB. Tuning with randomly modulated data (Fig. 4.11b), which obviates bringing the link down, is achieved by utilizing the same procedure. Note that the convergence time increases due to higher digital filtering for sufficient optical power averaging. Also, in order to achieve a similar extinction ratio in dynamic-tuning mode, the modulated resonance shift, shift, between data '0' and data '1' should be larger than the resonator's full width half maximum,  $\text{FWHM} = \lambda_0/Q$ , as shown in Fig. 4.12.




Figure 4.11: Simulated tuning waveforms and final optical transmission curves for (a) static tuning mode, and (b) dynamic tuning mode.




Figure 4.12: Extinction ratio versus modulated wavelength shift for static and dynamic tuning modes.

Note that while these simulation results and the experimental tuning results of Section VI are obtained by monitoring the modulator through-port, perhaps a more efficient approach for silicon photonic systems is to use an additional drop-port waveguide coupled to the ring modulator that has a waveguide PD for local power monitoring. This is accommodated in the current tuning system state machine via digital control to switch to drop-port monitoring mode, where the optical power is maximized to lock to the resonance point.

#### 4.6 Optical Forwarded-Clock Receiver

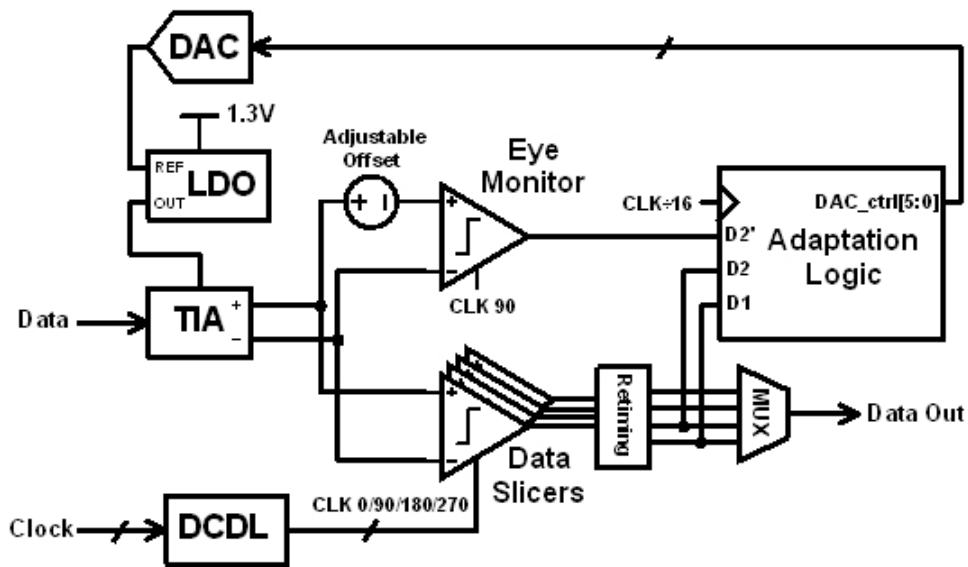


Figure 4.13: Adaptive sensitivity-power data receiver.

As shown in Fig. 4.13, the data-channel receivers consist of an inverter-based TIA front-end followed by a bank of four quadrature-clocked comparators whose offsets are digitally calibrated to optimize receiver sensitivity. The quadrature sampling clocks, generated from an optical forwarded-clock receiver, are passed through a local digitally-controlled delay line for timing margin optimization and phase-spacing calibration. An additional parallel comparator with a 6-bit programmable threshold is introduced that

serves as an eye monitor, setting the minimum voltage margin needed to correctly slice the input signal for a required bit-error rate. By comparing its output with the normal data comparator on the same clock phase, eye-closure can be detected before a bit-error actually occurs. This information is used to control a 6-bit R-2R voltage DAC that sets the LDO-generated TIA supply voltage to the minimum level required to achieve the sensitivity and bandwidth for a given bit-error rate. Fig. 4.14 shows the TIA front-end [57], which consists of three inverter stages with resistive feedback in the first and third stages. These inverter stages are biased around the trip-point for maximum gain with an offset control loop that subtracts the average photocurrent from the input node. The front-end's power supply level has a significant impact on gain, bandwidth, and noise performance [57], allowing for an efficient mechanism to trade-off receiver sensitivity with power consumption. However, excessive fluctuations can result in the front-end output common-mode variation if a simple single-ended low-pass filter is used in the offset control loop, which can impact overall receiver sensitivity. In order to reduce this common-mode variation, the feedback RC filter capacitor is split into equal decoupling to ground and the adaptive supply. A differential transconductance stage then amplifies the difference between this filtered node and half the adaptive supply to produce the offset correction current. This reduces the output common-mode disturbance with a 5mV power supply step from 92mV with a simple single-ended low-pass filter to 1.5mV with the adaptive-supply referenced implementation.

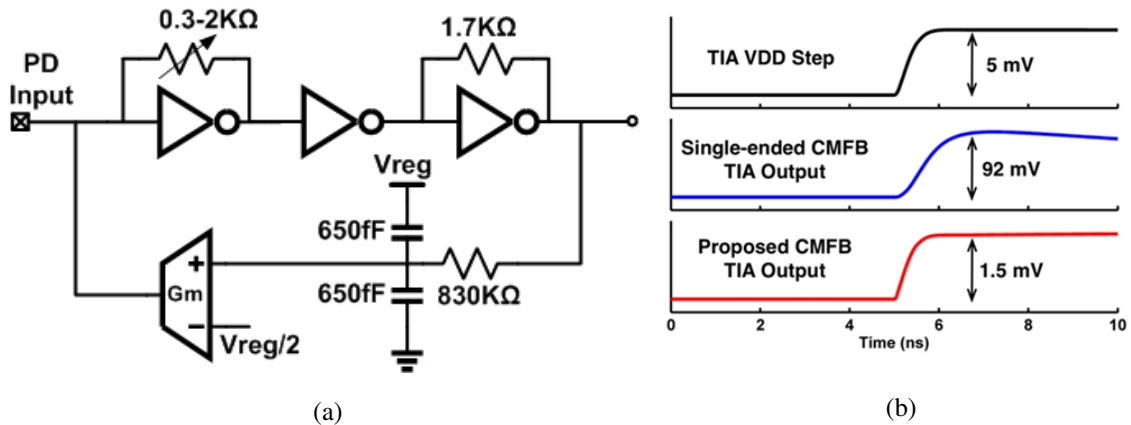


Figure 4.14: Inverter-based TIA front-end: (a) schematic, (b) simulated TIA common-mode output response to a 5mV power supply step.

The optical receiver sensitivity-power adaptation is done partially with a software-controlled outer loop that monitors the bit-error rate and adjusts the voltage margin with the eye-monitor comparator threshold through a serial test interface, and an on-chip state machine that scales the front-end power supply level. Fig. 4.15 summarizes the eye monitor and supply scaling state machine. The adaptation algorithm captures two consecutive bits D1 and D2, and proceeds only with a '01' pattern for the worst case ISI condition. Next, the data comparator output (D2) is compared with eye monitor output (D2') on the same clock phase, and an error is recorded if there is a difference. After a certain amount of total bits, a decision is made to reduce the power supply if no error is observed, or increase the power supply if the error rate exceeds a preset threshold. In order to minimize dithering without the overhead of a large averaging counter, the power supply doesn't change if the error rate is below a certain threshold.



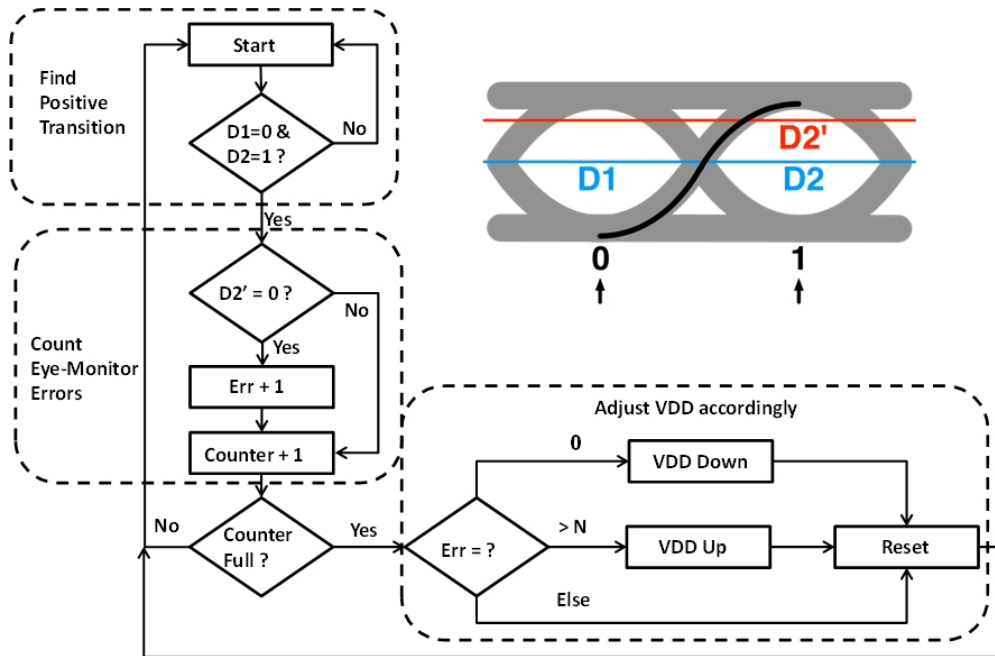


Figure 4.15: Optical receiver sensitivity-power adaption algorithm.

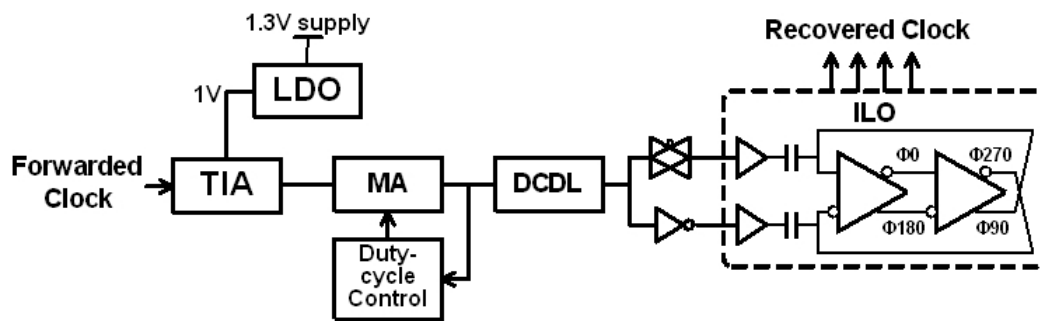


Figure 4.16: Optical clock receiver.

Fig. 4.16 shows a block diagram of the clock receiver, which utilizes the same inverter-based TIA front-end, but with a constant 1V supply for minimal jitter. The TIA output is amplified to full CMOS levels by a multi-inverter stage main amplifier (MA) that also contains a duty-cycle control loop. Global skew adjustment between the clock and data

channels is achieved by a subsequent digitally-controlled delay line, which provides approximately 130ps de-skew range. This clock is then converted from singled-ended to differential before being injected by AC-coupling into a two-stage differential oscillator that generates the quadrature clocks that are distributed to the four data receiver channels.

#### 4.7 Experimental Results

The optical transceiver circuits prototype was fabricated in a 65nm CMOS general purpose process. As shown in the photographs of Fig. 4.17,




Figure 4.17: Optical transceiver circuits prototype bonded for electrical characterization and optical testing. (a) Optical transmitter configuration with silicon ring resonator modulators. (b) Optical receiver configuration with commercial photodetectors.

A chip-on-board test setup is utilized, with the CMOS die wirebonded both to PCB traces for electrical characterization and to silicon ring resonator chips and commercial photodetectors for optical testing. Different bonding configurations are used for transmitter and receiver characterization. In order to verify the functionality of the pre-emphasis transmitters, electrical characterization is performed with the  $2V_{pp}$  and  $4V_{pp}$  transmitters driving a single-ended  $50\Omega$  and differential  $100\Omega$  termination, respectively. Fig. 4.18

shows  $2^7-1$  PRBS electrical eye diagrams with minimum and maximum pre-emphasis settings for the  $2V_{pp}$  transmitter module, which operates error-free up to 9Gb/s, and the  $4V_{pp}$ , which achieves 8Gb/s operation. A clear over/undershoot is observed for both drivers with the maximum pre-emphasis settings enabled. In both cases the maximum electrical data rate is limited by attenuation in the on-chip global clock distribution path. For high-speed optical testing, a continuous-wavelength laser is vertically coupled to a waveguide connected to a silicon ring resonator which is driven by a  $4V_{pp}$  transmitter. After vertically coupling the modulated light out into a single-mode fiber, the light is observed with an optical oscilloscope to produce the 5Gb/s eye diagrams of Fig. 4.19.

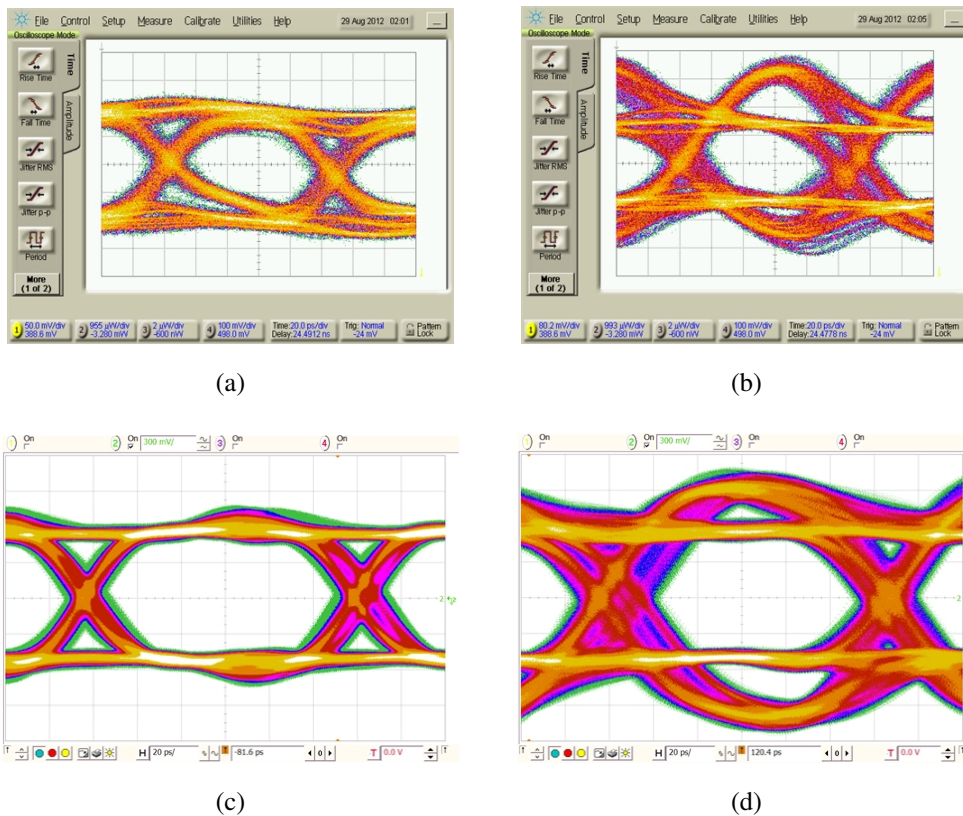


Figure 4.18: Modulator drivers' electrical eye diagrams. 9Gb/s operation with  $2V_{pp}$  driver with (a) minimum pre-emphasis, (b) maximum pre-emphasis. 8Gb/s operation with  $4V_{pp}$  driver with (c) minimum pre-emphasis (d) maximum pre-emphasis.

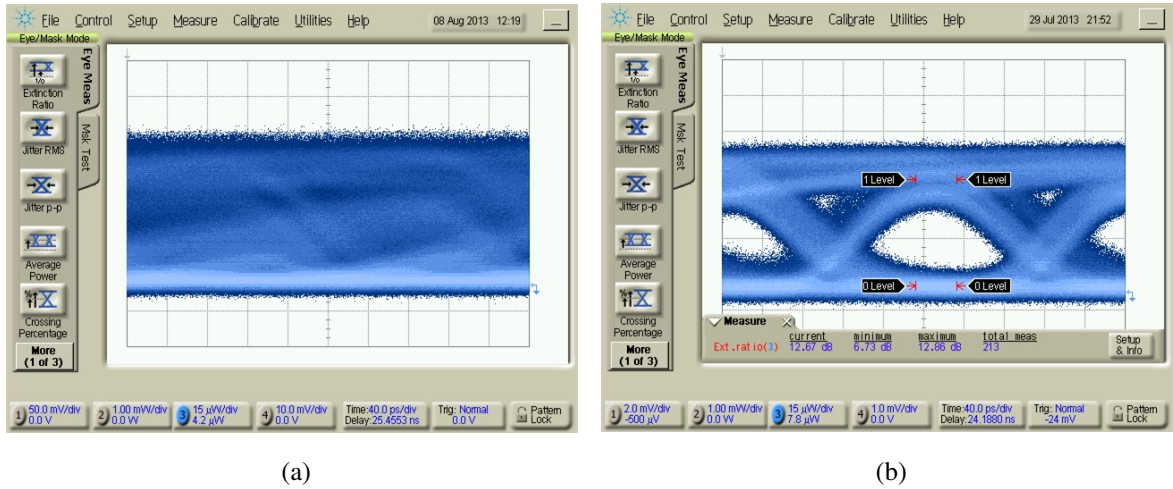
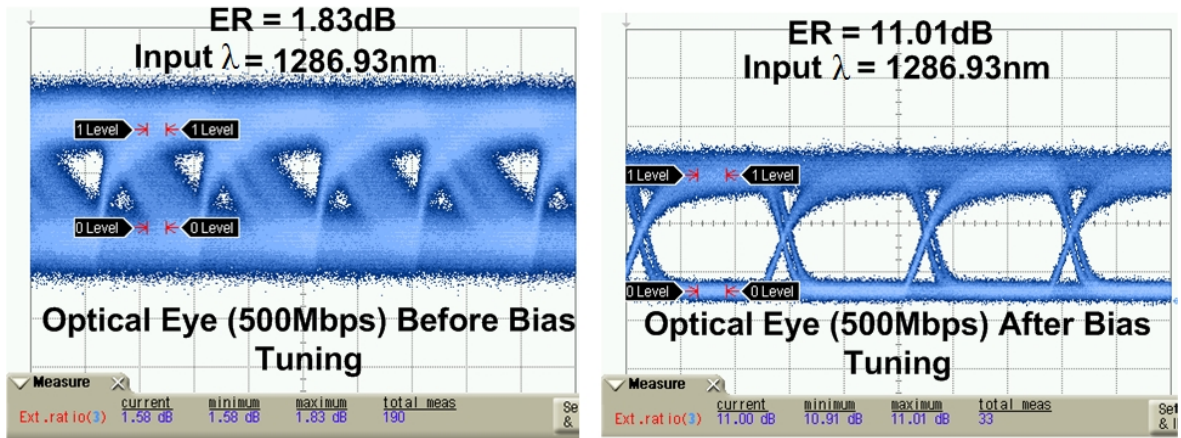


Figure 4.19: 5 Gb/s optical eye diagrams with silicon carrier-injection ring resonator modulators driven by the 4V<sub>pp</sub> transmitter: (a) minimum pre-emphasis settings; (b) optimized pre-emphasis settings.

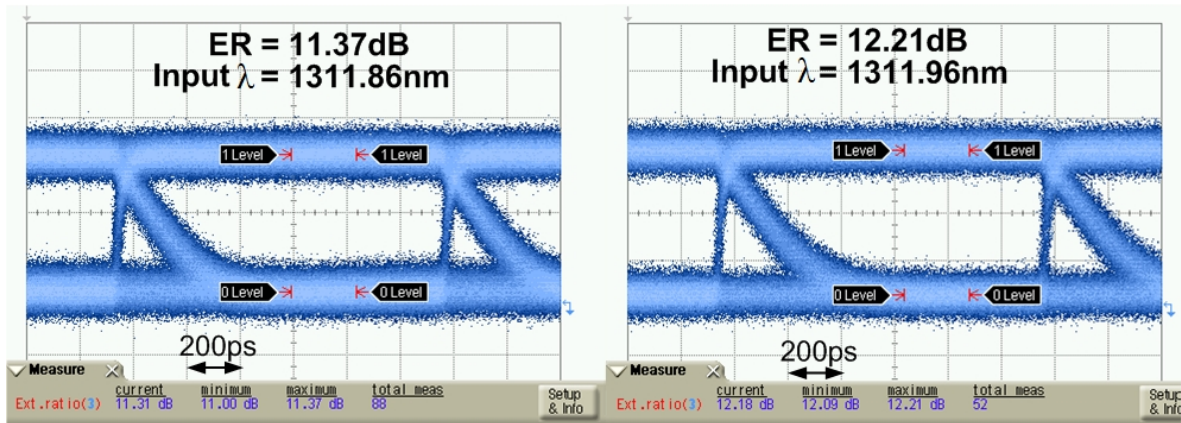
The eye is completely closed with the minimum pre-emphasis settings, while optimizing the pre-emphasis settings allows for an open eye with a 12.7dB extinction ratio. Here the maximum optical data rate is limited to 5Gb/s due to the unanticipated excessive contact resistance ( $2k\Omega$ ) of the ring resonator modulator. The bias-based resonance wavelength tuning effectiveness is demonstrated with two different ring resonator modulators with 0V-bias resonance wavelengths of 1287.01 and 1312.06nm. Fig. 4.20a shows the effectiveness of this bias-based control, with the extinction ratio improving from 1.8dB to 11.0dB after activating the tuning loop to lock to a 1286.93nm laser wavelength. A high extinction ratio is maintained as a given ring is tuned over different wavelengths, as shown by Fig. 4.20b where the bias-based tuning loop is able to lock to input wavelengths spaced by 0.1nm and obtain extinction ratios in excess of 11dB. The overall tuning range is 0.28nm for a tuning power of  $340\mu\text{W}$ , which results in a tuning efficiency of  $6.8\mu\text{W}/\text{GHz}$ . Note, this tuning power includes the entire tuning feedback loop of Fig. 4.8. Here the speed is



limited to 500-800Mb/s due to the bias-based tuning being implemented with the 2Vpp driver. Improving the excessive 2k  $\Omega$  contact resistance of the current ring resonator modulator to a more reasonable sub-100  $\Omega$  value should allow for simultaneous high-speed modulation and bias-based tuning capabilities.



(a)



(b)

Figure 4.20: Ring resonator bias-based wavelength stabilization measurements: (a) ring 1's 500Mb/s eye diagrams demonstrating the automatic bias tuning stabilizing to 1286.93nm, (b) ring 2's 800Mb/s eye diagrams with input laser wavelengths of 1311.86 and 1311.96nm.

In order to characterize the optical performance of the data receiver, an externally-modulated laser source is vertically coupled to a 150fF commercial photodiode which is wirebonded to the receiver input.

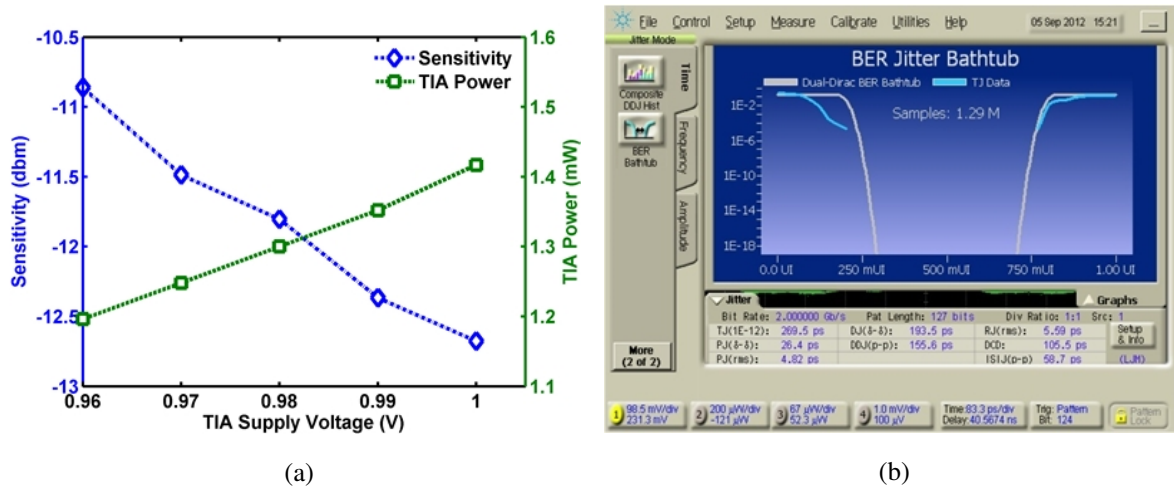


Figure 4.21: 8Gb/s receiver supply scaling measurements: (a) sensitivity (BER=10<sup>-15</sup>) and power versus TIA supply voltage, (b) BER bathtub plot for a power supply of 0.96V.

When the nominal 1V front-end power supply is utilized, Fig. 4.21 shows that a sensitivity of -12.7dBm is achieved at 8Gb/s for a BER=10<sup>-15</sup>. Relaxing the input sensitivity by 2 dB with increased optical input power enables the adaptive TIA supply to decrease by 5%, resulting in a 14% reduction in TIA power while still maintaining a healthy timing margin. As the data rate of the current optical characterization is limited by 1.5mm bond-wires and 200fF total capacitance, an on-chip current source (Fig. 4.22) is used to emulate a high-speed waveguide photodetector capable of being tightly integrated with the optical receiver [46, 52, 55]. Fig. 4.23 shows that this enables operation at a higher data rate of 10Gb/s with an improved sensitivity of -18dBm, assuming a unity responsivity. This on-chip test setup also enables a wider range of supply scaling, with the automated control loop reducing the TIA power 40% as the input current is scaled from 16 to 60 $\mu$ A with

a 50-100mV eye monitor margin. Refining the control state machine and using a more aggressive margin level could potentially achieve even more power savings, as overriding the automated control loop yields 60% power reduction.




Figure 4.22: Integrated photodetector emulator circuit.





Figure 4.23: 10Gb/s receiver supply scaling measurements with integrated photodetector emulator.

A similar optical test set-up is used to characterize the optical clock receiver. An

optical clock signal is amplified by the clock receiver and quadrature clocks are generated by the ILO, with one of the 2GHz quadrature clocks used for the 8Gb/s data receiver clocking shown in Fig. 25. The recovered clock jitter performance is a function of the input clock jitter and power, with the clock path introducing an additional  $0.25ps_{rms}$  jitter for  $-12\text{dBm}$  input power and able to generate sub- $2ps_{rms}$  total jitter down to  $-16\text{dBm}$ . Table 4.1 shows the transceiver circuits performance summary and compares this design with recent ring resonator optical interconnect work utilizing hybrid integration via face-to-face micro solder bonding [53] and monolithic integration [55]. An extinction ratio of 12.7dB is achieved with the injection-mode ring resonator modulators used in this work, which exceeds the 7dB extinction ratios achieved with the depletion-mode devices of [53, 55]. The  $4V_{pp}$  transmitter achieves 808fJ/bit at 5 Gb/s, while the  $2V_{pp}$  transmitter demonstrates bias-based resonator tuning with a 10% power overhead. While the optical receiver test configuration contributed to a dramatically higher input capacitance, a superior energy efficiency of 275 fJ/bit is achieved with the adaptive power-sensitivity receiver.



(a)




Figure 4.24: Optically forwarded-clock receiver measurements: (a) 2GHz recovered clock waveform, (b) jitter versus input optical power.



	This Work	[16]	[18]
<b>Ring Radius</b>	2.5 $\mu$ m	12 $\mu$ m	7.5 $\mu$ m
<b>Ring Mode</b>	Carrier-injection	Depletion	Depletion
<b>Ring Quality Factor (Q)</b>	~8,000	~15,000	~13,000
<b>Technology</b>	65nm CMOS	40nm CMOS	130nm CMOS SOI
<b>Supply Voltage</b>	1V, 2V, 2.5V	1V, 1.3V, 2V	1.2V/-1.2V
<b>TX Optical Data Rate</b>	5Gb/s (4V <sub>pp</sub> TX)	10Gb/s	25Gb/s
<b>TX Electrical Data Rate</b>	9Gb/s	10Gb/s	25Gb/s
<b>ER</b>	12.7dB	7dB	6.9dB
<b>Tuning Method</b>	Voltage Bias	Thermal	N/A
<b>Wavelength Tuning Range</b>	0.28nm	1.6nm	N/A
<b>TX Area</b> 4V <sub>pp</sub> TX 2V <sub>pp</sub> TX (w/ Bias Tuning)	0.04mm <sup>2</sup> 0.08mm <sup>2</sup>	0.00012mm <sup>2</sup>	N/A
<b>TX Power</b> 4V <sub>pp</sub> TX 2V <sub>pp</sub> TX (w/ Bias Tuning) Tuning	4.04mW 3.48mW 0.34mW* (6.8 $\mu$ W/GHz)	1.35mW  1.25mW (6.3 $\mu$ W/GHz)	207.6mW
<b>RX Input Capacitance</b>	>200fF	40~60fF	20fF
<b>RX Sensitivity</b> Optical input data Optical input clock Electrical input data Electrical Input clock	-12.7dBm @ 8Gb/s -18dBm @ 2GHz <16 $\mu$ A @ 10Gb/s <8 $\mu$ A @ 2.5GHz	-15dBm @ 10Gb/s	-6dBm @ 25Gb/s
<b>RX Power</b> TIA Comparators/other	1.42mW 0.78mW	3.95mW	48mW
<b>RX Area</b> Clock RX Data RX	0.032mm <sup>2</sup> 0.036mm <sup>2</sup>	0.008mm <sup>2</sup>	N/A

\*Includes the total power of the tuning loop shown in Fig. 9.

Table 4.1: Performance summary and comparisons.

## 4.8 Summary

This chapter presented silicon photonic transceiver circuits for a ring resonator-based optical interconnect architecture that addresses limited modulator bandwidth, variations in ring resonator resonance wavelength and system link budget, and efficient receiver clocking. The photonic transmitters incorporate high-swing non-linear pre-emphasis drivers to overcome the limited bandwidth of carrier-injection ring resonator modulators and an automatic bias-based tuning loop for resonance wavelength stabilization. An adaptive receiver trades-off sensitivity versus power to accommodate variations in input capacitance, modulator/photodetector performance, and link budget. The receive-side data sampling clocks are produced from an optically-forwarded quarter-rate clock which is amplified before being passed to an injection-locked oscillator for efficient quadrature clock generation. Overall, these circuits provide the potential for silicon photonic links that can deliver distance-independent connectivity whose pin-bandwidth scales with the degree of wavelength-division multiplexing.

## 5. EXPLORATION OF PHOTONIC NETWORK-ON-CHIP ARCHITECTURES\*

Over the past decade, single-chip multiprocessors (CMPs) have emerged to address power consumption and performance scaling issues in current and future VLSI process technology. On-chip interconnection networks, networks-on-chip (NoCs), have concurrently emerged to serve as a scalable alternative to traditional, bus-based interconnection between processor cores [64]. Conventional NoCs in CMPs use wide, point-to-point electrical links to relay cache-lines between private mid-level and shared last-level processor caches [65, 66], however, electrical on-chip interconnect is severely limited by bandwidth, power and latency constraints [67, 68]. These constraints are placing practical limits on the viability of future CMP scaling. For example, communication latency in a typical NoC connected multiprocessor system increases rapidly as the number of nodes increases [69]. Furthermore, power in electrical interconnects has been reported as high 12.1W for a 48-core, 2D-mesh CMP at 2GHz [66], a significant fraction of the system's power budget. Monolithic silicon photonics have been proposed as a scalable alternative to meet future many-core systems bandwidth demands, however current photonic NoC architectures suffer from high static power demands, high latency and low efficiency, making them less attractive than their electrical counterparts. In this chapter, a novel photonic NoC architecture is presented, which significantly reduces latencies and power consumption versus competing electrical and photonic NoC designs.

Monolithic silicon photonics can be efficiently scaled to meet future many-core systems bandwidth demands by leveraging high-speed photonic devices [9, 10, 70], THz-bandwidth waveguides [71, 72], and immense bandwidth-density via wavelength-division-

---

\*Reprinted with permission from "LumiNOC: a power-efficient, high-performance, photonic network-on-chip for future parallel architectures" by Cheng Li, Proceedings of the 21st international conference on Parallel architectures and compilation techniques (PACT'12), Page(s): 421 - 422, Copyright 2012 by ACM

multiplexing (WDM) [73, 74]. Recently, several NoC architectures leveraging the high bandwidth of silicon photonics have been proposed. These works can be categorized into two general types: 1). Hybrid optical/electrical interconnect architecture [75–78], in which a photonic packet-switched network and an electronic circuit-switched control network are combined to respectively deliver large size data messages and short control messages; 2). Crossbar or Clos architectures, in which the interconnect is fully photonic [79–87]. Although these designs provide high and scalable bandwidth, they either suffer from relatively high latency in the optical/electrical designs due to the electrical control circuits for photonic path setup, or significant power/hardware overhead in the crossbar/clos designs due to significant over-provisioned photonic channels. In future latency and power constrained CMPs, these characteristics promise to hobble the utility of photonic interconnect.

In this chapter, a novel photonic NoC (LumiNOC) architecture is proposed to address the issues of power and resource overhead due to channel over-provisioning, while reducing latency and maintaining high bandwidth. LumiNOC is divided into many small sub-networks, called subnets. Utilizing smaller subnets rather than a single, chip-encompassing, photonic crossbar, greatly increases efficiency by reducing on-chip photonic losses. Within a given subnet, a novel arbitration scheme and waveguide layout are employed to achieve extremely low-latency and low static power due to using fewer photonic devices and wavelengths compared to the other photonic architectures with the same aggregate bandwidth. At the same ideal throughput, LumiNOC consumes less than half the power of competing photonic network designs. LumiNOC is evaluated using both synthetic traffic and traces from PARSEC benchmark suite. The synthetic results show that low-load latency is reduced by  $\sim 50\%$  versus prior photonic NoCs and an electrical 2-D mesh network. Under realistic workloads, LumiNOC decreases packet latencies by 25% versus a 2-D mesh electrical NoC.

The remainder of this chapter is organized as follows. Section 5.1 presents some back-

ground on silicon nanophotonics in NoCs. Section 5.2 briefly reviews previous work on nanophotonics NoCs and discusses the pros and cons of each architecture. Section 5.3 addresses the issue of inefficient static power utilization, which motivated us to design the LumiNOC. The design details of LumiNOC are presented in Section 5.4, along with the novel SCDA arbitration scheme, router microarchitecture and flow control. A particular LumiNOC design implementation with 64 nodes is described and analyzed in Section 5.5. Section 5.6 evaluates the system performance and power efficiency, and compares the results with alternative PNoCs. Section 5.7 summarize the chapter.

### 5.1 Photonic Network-on-chip Technical Background

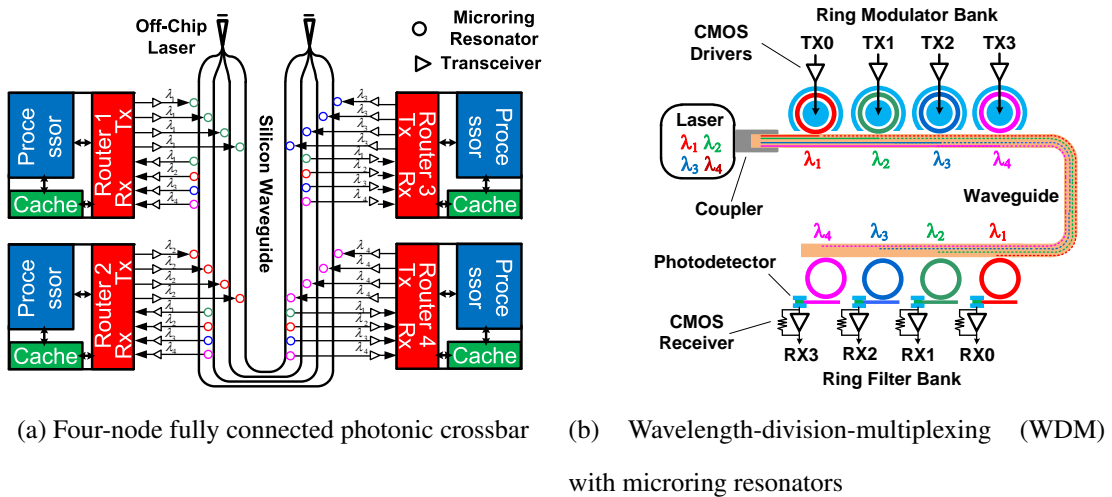


Figure 5.1: Basics of photonic on-chip interconnect.

Photonic NoCs have emerged as a potential replacement for electrical NoCs due to the high bandwidth, low latency and low power of nanophotonic channels. Figure 5.1a shows a small CMP with 4 compute tiles interconnected by a photonic NoC. Each tile consists of a processor core, private caches, a fraction of the shared last-level cache, and a router

connecting it to the photonic network. Figure 5.1a also shows the details of an example photonic NoC, organized as a simple, fully connected crossbar interconnecting the four processors. The photonic channel connecting the nodes is shown as being composed of microring resonators (MRR) [11, 74], integrated photodetectors [10](small circles) and silicon waveguides [71, 72] (black lines connecting the circles). Transceivers (small triangles) mark the boundary between the electrical and photonic domain. While the network shown is non-optimal in terms of scalability, it is sufficient for introducing the components of a simple photonic NoC.

**Silicon ring Resonators:** Silicon ring resonator can serve as either optical modulators for sending data or as filters for dropping and receiving data from on-chip photonic network. The basic configuration of a ring resonator consists of a silicon ring coupled with a straight waveguide. When the ring circumference equals an integer number of an optical wavelength, called resonance condition, most of the light from the straight waveguide circulates inside the ring and the light transmitted by the waveguide is suppressed. The resonance condition can be changed by applying electrical field over the ring, thus achieving electrical to optical modulation. ring resonance is sensitive to temperature variation, therefore, thermal/bias trimming is required to tune the ring to resonate at the working wavelength.

**Silicon Waveguides:** In photonic on-chip networks, silicon waveguides are used to carry the optical signals. In order to achieve higher aggregated bandwidth, multiple wavelengths are placed into a single waveguide in a wavelength-division-multiplexing (WDM) fashion. As shown in Figure 5.1b, multiple wavelengths generated by an off-chip laser ( $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ ) are coupled into a silicon waveguide via an optical coupler. At the sender side, microring modulators insert data onto a specific wavelength through electro-optical modulation. The modulated wavelengths propagate through integrated silicon waveguide and arrive at the receiver side, where microring filters drop the corresponding wavelength

and integrated photodetectors (PD) convert the signals back to the electrical domain.

In this work, silicon nitride waveguides are assumed to be the primary transport layer. Similar to electrical wires, silicon nitride waveguides can be deployed into multiple layers to eliminate in-plane waveguide crossing, thus reducing the optical power loss [88].

**Three-dimensional Integration:** In order to optimize system performance and efficiently utilize the chip area, three-dimensional integration (3DI) is emerging for the integration of silicon nanophotonic devices with conventional CMOS electronics. In 3DI, the silicon photonic on-chip networks are fabricated into a separate silicon-on-insulator (SOI) die or layer with a thick layer of buried oxide (BOX) that acts as bottom cladding to prevent light leakage into the substrate. This photonic layer stacks above the electrical layers containing the compute tiles.

In Figure 5.1a, the simple crossbar architecture is implemented by provisioning four send channels, each utilizing the same wavelength in four waveguides, and four receiving channels by monitoring four wavelengths in a single waveguide. Although this straightforward structure provides strictly non-blocking connectivity, it requires a large number of transceivers  $O(r^2)$  and long waveguides crossing the chip, where  $r$  is the crossbar radix, thus this style of crossbar is not scalable to a significant number of nodes. Researchers have instead proposed a number of more scalable photonic NoC architectures than fully connected crossbars, as described in the following section.

## 5.2 Related Work

Many photonic NoC architectures have been recently proposed which may be broadly categorized into four basic architectures: 1) Electrical-photonic 2) Crossbar 3) Multi-stage and 4) Free-space designs.

**Electrical-Photonic Designs:** Shacham et al. propose a hybrid electrical-photonic NoC using electrical interconnect to coordinate and arbitrate a shared photonic medi-

um [75–78]. These designs achieve very high photonic link utilization by effectively trading increased latency for higher bandwidth. While increased bandwidth without regard for latency is useful for some applications, it eschews a primary benefit of photonic NoCs over electrical NoCs, low latency. Recently, Hendry et al. addressed this issue by introducing an all optical mesh network with photonic time division multiplexing (TDM) arbitration to set up communication path. However, the simulation results show that system still suffers from relatively high average latency [89].

**Crossbar Designs:** Other recent photonic NoC work attempts to address the latency issue by providing non-blocking point-to-point links between nodes. In particular, several works have proposed crossbar topologies to improve the latency of multi-core photonic interconnect. While fully connected crossbars [81] do not allow practical scaling, some researchers have examined channel sharing crossbar architectures, called Single-Write-Multiple-Read (SWMR) or Multiple-Write-Single-Read (MWSR), with various arbitration mechanisms for coordinating shared sending and/or receiving channels. Vantrease et al. proposed Corona, a MWSR crossbar, in which each node listens on the dedicated channel, but with the other nodes competing to send data on this channel [84, 85]. To implement arbitration at sender side, the author implemented a token channel [85] or token slot [84] approach similar to token rings used in early LAN network implementations. Alternately, Pan et al. proposed Firefly, a SWMR crossbar design, with a dedicated sending channel for each node, but all the nodes in a crossbar listen on all the sending channels [83]. Pan et al. proposed broadcasting the flit-headers to specify a particular receiver.

In both SWMR and MWSR crossbar designs, over-provisioning of dedicated channels, either at the receiver (SWMR) or sender (MWSR), is required, leading to under utilization of link bandwidth and poor power efficiency. Pan et al. also proposed a channel sharing architecture, FlexiShare [82], to improve the channel utilization and reduce channel over-



provisioning. The reduced number of channels, however, limit the system throughput. In addition, FlexiShare requires separated dedicated arbitration channels for sender and receiver sides, incurring additional power and hardware overhead.

**Multi-stage Designs:** Recently, Joshi et al. proposed a photonic multi-stage Clos network with the motivation of reducing the photonic ring count, thus reducing the power for thermal ring trimming [79]. Their design explores the use of a photonic network as a replacement for the middle stage of a three-stage Clos network. While this design achieves an efficient utilization of the photonic channels, it incurs substantial latency due to the multi-stage design.

Koka et al. present an architecture consisting of a grid of nodes where all nodes in each row or column are fully connected by a crossbar [86]. To maintain full-connectivity of the network, electrical routers are used to switch packets between rows and columns. In this design, photonic “grids” are very limited in size to maintain power efficiency, since fully connected crossbars grow at  $O(n^2)$  for the number of nodes connected. Morris et al. [90] proposed a hybrid multi-stage design, in which grid rows (x-dir) are subnets fully connected with a photonic crossbar, but different rows (y-dir) are connected by a token-ring arbitrated shared photonic link.

**Free-Space Designs:** Xue et al. present a novel free-space optical interconnect for CMPs, in which optical free-space signals are bounced off of mirrors encapsulated in the chip’s packaging [91]. To avoid conflicts and contention, this design uses in-band arbitration combined with an acknowledgment based collision detection protocol. Packets not acknowledged as correctly received are resent after a timeout period. This acknowledgment period significantly reduces channel utilization and increases latency, however, relative to the other designs discussed. Abousamra et al. proposed to limit the point-to-point link in each row and column to address the scalability issue of conventional one-hop all-to-all free-space architecture [92]. Electrical routers are used to switch packets be-

tween rows and columns. However, the number of optical transceivers grows at  $O(n^3)$  for the number of nodes in the NoC, increasing photonic resource and power overhead.

Our proposed architecture, LumiNOC, attempts to address the issues found in competing designs. Similar to FlexiShare [82] and Clos [79], LumiNOC focuses on improving the channel utilization to achieve better efficiency and performance. Unlike these designs, however, LumiNOC leverages the same channels for arbitration, parallel data transmission and flow control, efficiently utilizing the photonic resources. Similar to Clos [79], LumiNOC is also a multi-stage design, however unlike Clos, the primary stage (our subnets) is photonic and the intermediate is electrical, leading to much lower photonic energy losses due to waveguide length and lesser latency due to simplified intermediate node electronic routers. Similar to Xue et al.’s design [91], in-band arbitration with collision detection is used to coordinate channel usage; however, in LumiNOC, the sender itself detects the collision and may start the retransmit process immediately without waiting for an acknowledgment, which may increase latency due to timeouts and reduce channel bandwidth utilization. These traits give LumiNOC better performance in terms of latency, energy efficiency and scalability.

### 5.3 Power Efficiency in Photonic Interconnect

Power efficiency is an important motivation for photonic on-chip interconnect. In photonic interconnect, however, the static power consumption (due to off-chip laser, ring thermal tuning, etc) dominates the overall power consumption, potentially leading to energy-inefficient photonic interconnects. In this section, we examine prior photonic NoCs in terms of static power efficiency. We use bandwidth per watt as the metric to evaluate power efficiency of photonic interconnect architectures, showing that it can be improved by optimizing the interconnect topology, arbitration scheme and photonic device layout.

**Channel Allocation:** We first examine channel allocation in prior photonic intercon-

nect designs. Several previous photonic NoC designs, from fully connected crossbars [81] to the blocking crossbar designs [80, 82–85], provision extra channels to facilitate safe arbitration between sender and receiver. Although conventional photonic crossbars achieve nearly uniform latency and high bandwidth, channels are dedicated to each node and cannot be flexibly shared by the others. Due to the unbalanced traffic distribution in realistic workloads [93], channel bandwidth cannot be fully utilized. This leads to inefficient energy usage, since the static power is constant regardless of traffic load. Over-provisioned channels also implies higher ring resonator counts, which must be maintained at the appropriate trimming temperature, consuming on-chip power. Additionally, as the network size increases, the number of channels required may increase quadratically, complicating the waveguide layout and leading to extra optical loss. An efficient photonic interconnect must solve the problem of efficient channel allocation. Our approach leverages this observation to achieve lower power consumption than previous designs.

**Topology and Layout:** Topology and photonic device layout can also cause unnecessary optical loss in the photonic link, which in turn leads to greater laser power consumption. Many photonic NoCs globally route waveguides in a bundle, connecting all the tiles in the CMP [80, 83–85]. In these designs, due to the unidirectional propagation property of optical transmission, the waveguide must double back to reach each node twice, such that the signal being modulated by senders on the outbound path may be received by all possible receivers. The length of these double-back waveguides leads to significant laser power losses over the long distance.




Figure 5.2: Optical link budgets for the photonic data channels of various photonic NoCs.

Figure 5.2 shows the optical link budgets for the photonic data channel of Corona [85], Firefly [83], Clos [79] and LumiNOC under same radix and chip area, based on our power model (described in Section 5.5.1). Flexishare [82] is not compared, since not enough information was provided in the paper to estimate the optical power budget at each wavelength. The figure shows that waveguide losses dominate power loss in all three designs. This is due to the long waveguides required to globally route all the tiles on a chip. For example, the waveguide length in Firefly and Clos network in a  $400 \text{ mm}^2$  chip are estimated to be 9.5cm and 5.5cm, respectively. This corresponds to 9.5dB and 5.5dB loss in optical power, assuming the waveguide loss is 1dB/cm [79]. Moreover, globally connected tiles imply a relatively higher number of rings on each waveguide, leading to higher ring through loss. Despite a single-run, bi-directional architecture, even the Clos design shows waveguide loss as the largest single component.

In contrast to other losses (e.g. coupler and splitter loss, filter drop loss and photodetector loss) which are relatively independent of interconnect architecture, waveguide and ring through loss can be reduced through layout and topology optimization. We propose a network architecture which reduces optical loss by decreasing individual waveguide length

as well as the number of rings along the waveguide.

**Arbitration Mechanism:** The power and overhead introduced by the separated arbitration channels or networks in previous photonic NoCs can lead to further power efficiency losses. Corona, a MWSR crossbar design, requires a token channel or token slot arbitration at sender side [84, 85]. Alternatively, Firefly [83], a SWMR crossbar design, requires the broadcasting arbitration at receiver side. FlexiShare [82] requires both token stream arbitration and head-flit broadcasting arbitration. These arbitration mechanisms require significant overhead in form of dedicated channels and photonic resources, consuming extra optical laser power. For example, the radix-32 Flexishare [82] with 16 channels requires 416 extra wavelengths for arbitration, which accounts for 16% of the total wavelengths. Firefly [83] and FlexiShare [82] also incur higher optical power to facilitate a multi-receiver broadcast of header-flits. Thus, in these designs optical power is wasted since all but the header-flit is unicast. Arbitration mechanisms are a major overhead for these architectures, particularly as network radix scales.




Figure 5.3: Optical power overhead of arbitration channels in various photonic NoCs.

We analyzed the optical power overhead of arbitration channels in different architec-




Figure 5.4: LumiNOC interconnection of CMP with 16 tiles - (a) One-row interconnection, (b) Two-rows interconnection, (c) Four-rows interconnection.

tures, and show the results in Figure 5.3. Corona [85], a token ring arbitration based MWSR crossbar, achieves the least power overhead of arbitration channels at the cost of high transmission latency due to the delay of token passing through the arbitration waveguide. The power overhead in Firefly [83] and FlexiShare is relatively higher, since a large amount of optical power is required to broadcast the head-flit for receiver selection. Note that photonic clos [79] has no optical power overhead due to arbitration channels, since it uses electrical buffering and arbitration to control the middle photonic routers.

There is a clear need for a PNoC architecture that is energy-efficient and scalable while maintaining low latency and high bandwidth. In the following sections, we propose the LumiNOC architecture which reduces the optical loss by partitioning the global network into multiple smaller sub-networks. Further, a novel arbitration scheme is proposed which leverages the same wavelengths for channel arbitration and parallel data transmission to efficiently utilize the channel bandwidth and photonic resources, without dedicated arbitration channels or networks which lower efficiency or add power overhead to the system.

#### 5.4 LumiNOC Architecture

In our analysis of prior PNoC designs, we found a significant amount of laser power consumption was due to the waveguide length required for propagation of the photonic

signal across the entire network. Based on this, the LumiNOC design breaks the network into several smaller networks (subnets), with shorter waveguides. Figure 5.4 shows three example variants of the LumiNOC architecture with different subnet sizes, in an example 16-node CMP system: the one-row, two-rows and four-rows designs. In the one-row design, a subnet of four tiles is interconnected by a photonic waveguide in the horizontal orientation. Thus four non-overlapping subnets are needed for the horizontal interconnection. Similarly four subnets are required to vertically interconnect the 16 tiles. In the two-row design, a single subnet connects 8 tiles while in the four-row design a single subnet touches all 16 tiles. In general, all tiles are interconnected by two different subnets, one horizontal and one vertical. If a sender and receiver do not reside in the same subnet, transmission requires a hop through an intermediate node's electrical router. In this case, transmission experiences longer delay due to the extra O/E-E/O conversions and router latency. To remove the overheads of photonic waveguide crossings required by the orthogonal set of horizontal and vertical subnets, the waveguides can be deposited into two layers with orthogonal routing [88].

Another observation from prior photonic NoC designs is that channel sharing and arbitration have a large impact on design power efficiency. Efficient utilization of the photonic resources, such as wavelengths and ring resonators, is required to yield the best overall power efficiency. To this end, we leverage the same wavelengths in the waveguide for channel arbitration and parallel data transmission, avoiding the power and hardware overhead due to the separated arbitration channels or networks. Unlike the over-provisioned channels in conventional crossbar architectures, channel utilization in LumiNOC is improved by multiple tiles sharing a photonic channel.

A final observation from our analysis of prior photonic NoC design is that placing many wavelengths within each waveguide through deep wavelength-division multiplexing (WDM) leads to high waveguide losses. This is because the number of rings that each in-

dividual wavelength encounters as it traverses the waveguide is proportional to the number of total wavelengths in the waveguide times the number of waveguide connected nodes, and each ring induces some photonic power losses. We propose to limit LumiNOC's waveguides to a few frequencies per waveguide and increase the count of waveguides per subnet, to improve power efficiency with no cost to latency or bandwidth, a technique we call "ring-splitting". Ring-splitting is ultimately limited by the tile size, with a reasonable waveguide pitch of  $15\mu\text{m}$  required for layout of microrings as we will discuss in implementation.

#### 5.4.1 *LumiNOC Subnet Design*

Figure 5.5 details the shared channel for a LumiNOC one-row subnet design. Each tile contains  $W$  modulating "Tx rings" and  $W$  receiving "Rx Rings", where  $W$  is the number of wavelengths multiplexed in the waveguide. Since the optical signal unidirectionally propagates in the waveguide from its source at off-chip laser, each node's Tx rings are connected in series on the "Data Send Path", shown in a solid line from the laser, prior to connecting each node's Rx rings on the "Data Receive Path", shown in a dashed line. In this "double-back" waveguide layout, modulation by any node can be received by any other node; furthermore, the node which modulates the signal may also receive its own modulated signal, a feature that is leveraged in our collision detection scheme in the arbitration phase. The same wavelengths are leveraged for arbitration and parallel data transmission.



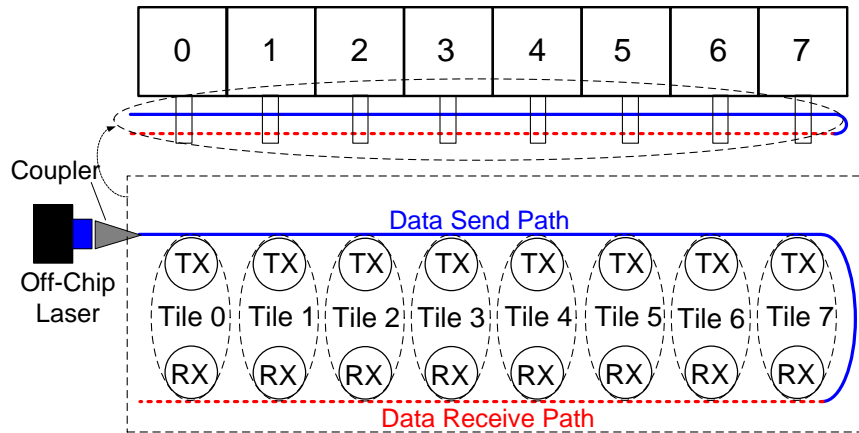


Figure 5.5: Bold circles (TX and RX) represent groups of rings, and each pair in the oval are for a single node.

During data transmission, only a single sender is modulating on all wavelengths and only a single receiver is tuned to all wavelengths. However, during arbitration (i.e. any time data transfer is not actively occurring) the Rx rings in each node are tuned to a specific, non-overlapping set of wavelengths. Up to half of the wavelengths available in the channel are allocated to this arbitration procedure. with the other half available for credit packets as part of credit-based flow control. This particular channel division is designed to prevent optical broadcasting, the state when any single wavelength must drive more than one receiver, which if allowed would severely increase laser power [94]. Thus, at any given time a multi-wavelength channel with  $N$  nodes may be in one of three states: **Idle** - All wavelengths are un-modulated and the network is quiescent. **Arbitration** - One more sender nodes are modulating  $N$  copies of the arbitration flags; one copy to each node in the subnet (including itself) with the aim to gain control of the channel. **Data Transmission** - Once a particular sender has established ownership of the channel, it modulates all channel wavelengths in parallel with the data to be transmitted.

In the remainder of this section, we detail the following: **receiver addressing and**

**sender-side arbitration** - the mechanism by which the photonic channel is granted to one sender, avoiding data corruption when multiple senders wish to transmit, and by which the receivers are addressed and **data transmission** - the mechanism by which data is transmitted from sender to receiver.

**Receiver Addressing and Sender-side Arbitration:** We propose an optical statistical collision detection and avoidance (SCDA) technique to coordinate access of the shared photonic channel at sender-side. This approach can achieve efficient channel utilization without the latency of electrical arbitration schemes [75–78], or the overhead of wavelengths and waveguides dedicated to standalone arbitration [82, 83, 85]. In this scheme, a sender works together with its own receiver to ensure message delivery in the presence of conflicts.

*Receiver:* Once any receiver detects an arbitration flag, it will take one of three actions: if the arbitration flag is uncorrupted (single-sender) and the forthcoming message is destined for this receiver, it will enable all its Rx rings for the indicated duration of the message, capturing it. If the arbitration flags are uncorrupted, but the receiver is not the intended destination, it will detune all of its Rx rings for the indicated duration of the message to allow the recipient sole access. Finally, if a collision is detected (described below), the receiver circuit will take no action.

*Sender:* To send a packet, a node will first wait for any on-going messages to complete. Then, it will modulate one copy of the arbitration flags to the appropriate arbitration wavelengths for each of the  $N$  nodes. The arbitration flags for an example 4 node subnet are depicted in Figure 5.6. The arbitration flags are a  $t_{arb}$  cycle long header (2 in this example) made up of the destination node address (D0-D1, encoded to always have at least 1 bit set), a bimodal packet size indicator (Ln), and a “1-hot” source address (S0-S3) which serves as a guard band or collision detection mechanism: since the subnet is operated synchronously, any time multiple nodes send overlapping arbitration flags, the “1-hot” precondition

will be violated and all nodes will be aware of the collision. We make use of the fact that a node will receive its own signal: right after sending, the node will monitor the incoming arbitration flags. If they are uncorrupted, then the sender was successful at arbitrating the channel and the two nodes proceed to the Data Transmission phase. If the arbitration flags are corrupted, then a conflict has occurred. Any data already sent is ignored and the conflicting senders enter the SCDA regime. Inspired by traditional CSMA/CD protocols, all conflicting senders will attempt resending with exponentially increasing periods between resend attempts. So, if two nodes conflict, each will wait either 0 or 1 cycles before attempting a resend. If either of those attempts fail, the conflicting nodes will wait between 0 and 2 cycles, then 0 and 4, etc. In this way, the conflict rate self regulates proportionally to the subnet activity.

If the subnet size increases without proportionally increasing the available wavelengths per subnet, then the arbitration flags will take longer to serialize as more bits will be required to encode the source and destination address. However, if additional wavelengths are provisioned to maintain the bandwidth/node, then the additional arbitration bits can be sent in parallel.

The physical length of the waveguide incurs a propagation delay,  $t_{pd}$  (cycles), on the arbitration flags traversing the subnet. The “1-hot” collision detection mechanism will only function if the signals from all senders are temporally aligned, so if nodes are physically further apart than the light will travel in 1 cycle, they will be in different clocking domains to keep the packet aligned as it passes the final sending node. Furthermore, the arbitration flags only start on cycles that are an integer multiple of the  $t_{pd} + 1$  to assure that no nodes started arbitration during the previous  $t_{slot}$  and that all possibly conflicting arbitration flags are aligned. This also means that conflicts only occur on arbitration flags, not with data.

Note that a node will not know if it has successfully arbitrated the channel until after  $t_{pd} + t_{arb}$  cycles, but will begin data transmission after  $t_{arb}$ . In the case of an uncontested

link, the data is captured by the receiver immediately. When it detects conflict, the senders will cease sending unusable data.

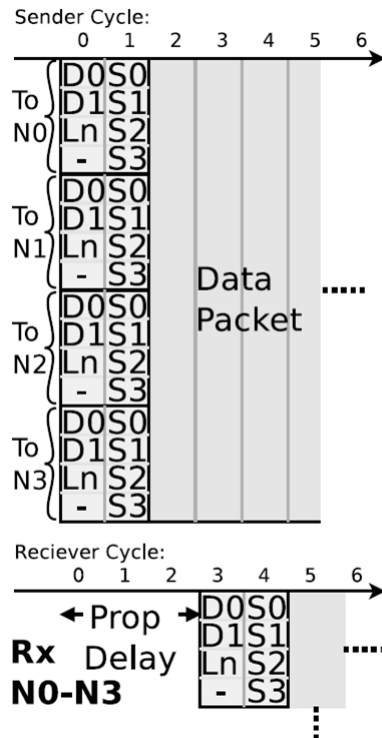


Figure 5.6: Arbitration on 4 a node subnet.

As an example, say that the packet in Figure 5.6 is destined for node 2 with no conflicts. At cycle 5, Nodes 1, 3, and 4 would disable their arbitration receivers, but node 2 would enable them all and begin data transfer.

**Data Transmission:** In this phase the sender transmits the data over the photonic channel to the receiving node. All wavelengths in the waveguide are used for bit-wise parallel data transmission, so higher throughput is expected when more wavelengths are multiplexed into the waveguide.

### 5.4.2 Router Microarchitecture

The electrical router architecture for LumiNOC is shown in Figure 5.7. Each router serves both as an entry point to the network for a particular core, as well as an intermediate node interconnecting horizontal and vertical subnets. If a processor must send data to another node on the same vertical or horizontal subnet, packets are switched from the electrical input port to the vertical photonic output port with one E/O conversion. Packets which are destined for a different subnet must be first routed to an intermediate node via the horizontal subnet before being routed on the vertical subnet. Each input port is assigned with a set of flit buffers as virtual-channels (VCs) to hold the incoming flits. The local control unit performs routing computation, virtual-channel allocation and switching allocation in crossbar. The LumiNOC router's complexity is similar to that of a electrical, bi-directional, 1-D ring network router, with the addition of the E/O-O/E logic.

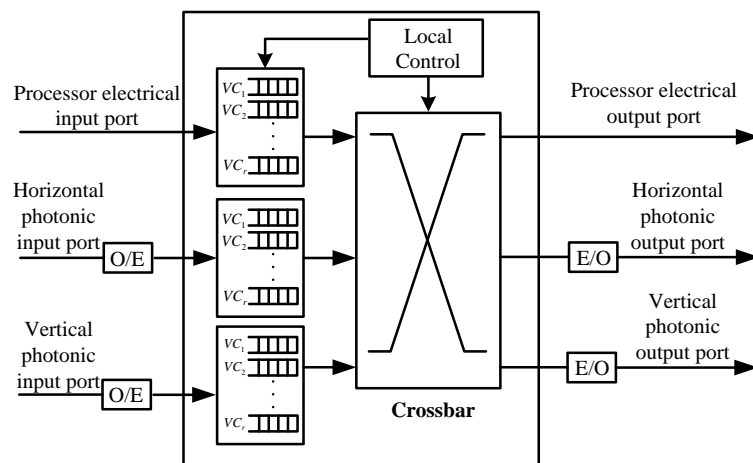


Figure 5.7: Router microarchitecture.

## 5.5 Implementation

In this section we discuss one particular baseline physical implementation of the general LumiNOC architecture specified in Section 5.4. We assume a  $400 \text{ mm}^2$  chip implemented in a 22 nm CMOS process and containing 64 square tiles that operate at 5GHz, as shown in Figure 5.8. A 64-node LumiNOC design point is chosen here as a reasonable network size which could be implemented in a 22nm process technology. The tiles each contain a processor core, private caches, a fraction of the shared last-level cache, and a router connecting it to one horizontal and one vertical photonic subnet. Each router input port contains seven virtual channels (VCs), each five flits deep. Credit based flow control implemented via a single-bit, pipelined wire is used between nodes within a subnet.




Figure 5.8: One-row LumiNOC with 64 tiles.

A 64-node LumiNOC may be organized into three different architectures: the one-row, two-row and four-row designs (shown in Figure 5.4), which represent a trade-off between

interconnect power, system throughput and transmission latency. For example, power decreases as row number increases from one-row to two-row, since the single waveguide is roughly with the same length, but fewer waveguides are required. The low-load latency is also reduced due to more nodes residing in the same subnet, reducing the need for intermediate hops via an electrical router. The two-row subnet design, however, significantly reduces throughput due to the reduced number of transmission channels. As a result, we choose the “one-row” subnet architecture of Figure 5.4a, arranged as shown in Figure 5.8 for the remainder of this section. In both the horizontal and vertical axes there are 8 subnets which are formed by 8 tiles that share a photonic channel, resulting in all tiles being redundantly interconnected by two subnets. As discussed in Section 5.1, 3DI is assumed, placing orthogonal waveguides into different photonic layers, eliminating in-plane waveguide crossings [88].

We assume a 10GHz network modulation rate, while the routers and cores are clocked at 5GHz. Muxes are placed on input and output registers such that on even network cycles, the photonic ports will interface with the lower half of a given flit and on odd, the upper half. With a  $400 \text{ mm}^2$  chip with photonic ports placed closest to the center-line of the chip, an effective datapath length of 3.0 cm can be realized, yielding a propagation delay of  $t_{pd} = 3.5$  network cycles.

When sender and receiver reside in the same subnet, data transmission is accomplished with a single hop, i.e. without a stop in an intermediate electrical router. Two hops are required if sender and receiver reside in different subnets, resulting in a longer delay due to the extra O/E-E/O conversion and router latency. The “one-row” subnet based network implies that for any given node 15 of the 63 possible destinations reside within one hop, the remaining 48 destinations require two hops.

Loss Component	Value	Loss Component	Value
Coupler	1 dB	Waveguide	1 dB/cm
Splitter	0.2 dB	Waveguide Crossing	0.05 dB
Non-linearity	1 dB	Ring Through	0.001 dB
Modulator Insertion	0.001 dB	Filter Drop	1.5 dB
Photodetector	0.1 dB		

Table 5.1: Components of optical loss.

Considering the link width, or the number of wavelengths per logical subnet, if the number of wavelengths and thus channel width is increased, it should raise ideal throughput and theoretically reduce latency due to serialization delay. We are constrained, however, by the 3.5 network cycle propagation delay of the link, and the small packet size of single cache line transfers in typical CMPs. There is no advantage to sending the arbitration flags all at once in parallel when additional photonic channels are available; the existing bits would need to be replaced with more guard bits to provide collision detection. Thus, the arbitration flags would represent an increasing overhead. Alternately, if the link were narrower, the 3.5 cycle window would be too short to send all the arbitration bits and a node would waste time broadcasting arbitration bits to all nodes after it effectively “owns” the channel. Thus, the optimal link width is 64 wavelengths under our assumptions for clock frequency and waveguide length. If additional spectrum or waveguides are available, then we propose to implement multiple parallel, independent network layers. Instead of one network with a 128-bit data path, there will be two parallel 64-bit networks. This allows us to exploit the optimal link width while still providing higher bandwidth. When a terminal is injecting into the network, it round-robins through the available input ports, dividing the traffic amongst the layers evenly.



We limit the wavelengths per waveguide to 32, to reduce the ring through loss as described in Section 5.4. This implies a trade off of waveguide area for lower power.

### 5.5.1 Photonic Power Model

In this section, we describe our total power model and compare the baseline LumiNOC design against prior work PNoC architectures. In order for a fair comparison versus other reported PNoC architectures, we refer to the photonic loss of various photonic devices reported by Joshi et al. [79] and Pan et al. [82], shown in Table 5.1. Equation 5.1 shows the major components of our total power model.

$$TP = ELP + TTP + ERP + EO/OE \quad (5.1)$$

TP = Total Power, ELP = Electrical Laser Power, TTP = Thermal Tuning Power, ERP = Electrical Router Power and EO/OE = Electrical to Optical/Optical to Electrical conversion power. Each components is described below.

**ELP:** Electrical laser power is converted from the calculated optical power. Assuming a  $10\mu\text{W}$  receiver sensitivity, the minimum static optical power required at each wavelength to activate the farthest detector in the PNoC system is estimated based on Equation 5.2. This optical power is then converted to electrical laser power using 30% efficiency.

$$P_{optical} = N_{wg} \cdot N_{wv} \cdot P_{th} \cdot K \cdot 10^{(\frac{1}{10} \cdot l_{channel} \cdot P_{WG.Loss})} \cdot 10^{(\frac{1}{10} \cdot N_{ring} \cdot P_{t.Loss})} \quad (5.2)$$

In Equation 5.2,  $N_{wg}$  is the number of waveguide in the PNoC system,  $N_{wv}$  is the number of wavelength per waveguide,  $P_{th}$  is receiver sensitivity power,  $l_{channel}$  is waveguide length,  $P_{wg.loss}$  is optical signal propagation loss in waveguide (dB/cm),  $N_{ring}$  is the number of rings attached on each waveguide,  $P_{t.loss}$  is modulator insertion and filter ring through loss (dB/ring) (assume they are equal),  $K$  accounts for the other loss components in the optical path including  $P_c$ , coupling loss between the laser source and optical wave-




Figure 5.9: Electrical Laser Power (W) contour plots for networks with the same aggregate throughput (assuming 30% efficient electrical to optical power conversion)

uide,  $P_b$ , waveguide bending loss, and  $P_{splitter}$ , optical splitter loss. Figure 5.9 shows electrical laser power contour plot, derived from Equation 5.2, showing the photonic device power requirements at a given electrical laser power, for a SWMR photonic crossbar(Corona) [85], Clos [79] and LumiNOC with equivalent throughput (20Tbps), network radix and chip area. In the figure, x and y-axis represent two major optical loss components, waveguide propagation loss and ring through loss, respectively. A larger x- and y-intercept implies relaxed requirements for the photonic devices. As shown, given a relatively low 1W laser power budget, the two-layer LumiNOC can operate with a maximum 0.012dB ring through loss and waveguide loss of 1.5dB/cm.

**TTP:** Thermal tuning is required to maintain microring resonant at the work wavelength. In the calculation, a ring thermal tuning power of  $20\mu\text{W}$  is assumed for a 20K temperature tuning range [79,82]. In a photonic NoC, total thermal tuning power (TTP) is proportional to ring quantity.

**ERP:** The baseline electrical router power is estimated by the power model reported by Kim et al. [95]. We synthesized the router using TSMC 45nm library. Power is measured via Synopsis Power Compiler, using simulated traffic from a PARSEC [67] workload to

Literature		$N_{core}$	$N_{node}$	$N_{rt}$	$N_{wg}$	$N_{wv}$	$N_{ring}$	ITP (Tbps)
EMesh [66]		128	64	64	NA	NA	NA	10
Corona [85]		256	64	64	388	24832	1056K	160
FlexiShare [82]		64	32	32	NA	2464	550K	20
Clos [79]		64	8	24	56	3584	14K	18
LumiNOC	1-Layer	64	64	64	32	1024	16K	10
	2-Layers	64	64	64	64	2048	32K	20
	4-Layers	64	64	64	128	4096	65K	40

Table 5.2: Configuration comparison of various photonic NoC architectures -  $N_{core}$  = number of cores in the CMP,  $N_{node}$  = number of nodes in the NoC,  $N_{rt}$  = total number of routers,  $N_{wg}$  = total number of waveguides,  $N_{wv}$  = total number of wavelengths,  $N_{ring}$  = total number of rings, ITP = Ideal Throughput.

estimate its dynamic component. Results are analytically scaled to 22nm.

**EO/OE:** The power for conversion between the electrical and optical domains (EO/OE) is based on the model reported by Joshi et al. [79], which assumes a total transceiver energy of 40 fJ/bit data-traffic dependent energy and 10 fJ/bit static energy. Since previous photonic NoCs consider different traffic loads, it is unfair to compare the EO/OE power by directly using their reported figures. Therefore, we compare the worst-case power consumption when each node was arbitrated to get a full access on each individual channel. For example, Corona is a MWSR  $64 \times 64$  crossbar architecture. At worst-case, 64 nodes are simultaneously writing on 64 different channels. This is combined with a per-bit activity factor of 0.5 to represent random data in the channel.

### 5.5.2 Power Comparison

Table 5.2 lists the photonic resource configurations for various photonic NoC architectures, including one-layer, two-layer and four-layer configurations of the LumiNOC. While the crossbar architecture of Corona has a high ideal throughput, the excessive number of rings and waveguides results in degraded power efficiency. In order to support equal 20Tbps aggregate throughput, LumiNOC requires less than  $\frac{1}{10}$  the number of rings

Literature		ELP (W)	TTP (W)	ERP (W)	EO/OE (W)	RTP/ITP (Tbps)	TP (W)	ITP/Watt (Tbps/W)
EMesh [66]		NA	NA	NA	NA	10	26.7	0.37
Corona [85]		26.00	21.00	0.52	4.92	160	52.4	3.1
FlexiShare [82]		5.80	11.00	0.13	0.60	9/20	17.5	1.1
Clos [79]		3.30	0.14	0.10	0.54	8/18	4.1	4.4
LumiNOC	1-Layer	0.35	0.33	0.13	0.30	3.5/10	1.1	9.1
	2-Layers	0.73	0.65	0.26	0.61	6.25/20	2.3	8.9
	4-Layers	1.54	1.31	0.52	1.22	13/40	4.6	8.7

Table 5.3: Power efficiency comparison of different photonic NoC architectures - ELP = Electrical Laser Power, TTP = Thermal Tuning Power, ERP = Electrical Router Power, EO/OE = Electrical to optical/Optical to electrical conversion power, ITP = Ideal Throughput, TP = Total Power.

of FlexiShare and almost the same number of wavelengths. Relative to the Clos architecture, LumiNOC requires around  $\frac{4}{7}$  wavelengths, though approximately double number of rings.

The power and efficiency of the network designs is compared in Table 5.3. The individual components of optical power are calculated as described in Section 5.5.1. A  $6 \times 4$  2GHz electrical 2D-mesh [66] was scaled to  $8 \times 8$  nodes operating at 5GHz, in a 22nm CMOS process, to compare against the photonic networks.

The table shows that LumiNOC has the highest power efficiency of all designs compared in ITP/Watt, doubling the efficiency of the nearest competitor, Clos [79]. By reducing wavelength multiplexing density, utilizing shorter waveguides, and leveraging the data channels for arbitration, LumiNOC consumes the least ELP among all the compared architectures. A 4-layer LumiNOC consumes half the ELP of a competitive Clos architecture, with double ideal throughput. Corona [85] contains 256 cores with 4 cores sharing an electrical router, leading to a 64-node photonic crossbar architecture; however, in order to achieve throughput of 160Gbps, each channel in Corona consists of 256 wavelengths, 4X the wavelengths in a 1-layer LumiNOC. In order to support the highest ideal throughput,

Corona consumes the highest electrical router power in the compared photonic NoCs.

## 5.6 Evaluation

### 5.6.1 Methodology

To evaluate the performance of our design, we use a cycle-accurate, microarchitectural network simulator, *ocin\_tsim* [96]. The network was simulated under both synthetic and realistic workloads. As described in Section 5.5, LumiNOC designs with multiple layers are simulated to show results for different bandwidth design points. These layers are independent; the injection terminals will choose a layer in round-robin fashion when first injecting a packet. Results are presented for 1, 2, and 4 layers.

**Photonic Networks:** The baseline, 64-node LumiNOC system, as described in Section 5.5, was simulated for all evaluation results. Synthetic benchmark results for the Clos LTBw network are presented for comparison against the LumiNOC design. We chose the Clos LTBw design as the most competitive in terms of efficiency and bandwidth as discussed in Section 5.5. Clos LTBw data points were extracted from the paper by Joshi et al [79].

**Baseline Electrical Network:** In the results that follow, our design is compared to a electrical 2-D mesh network. Traversing the dimension order network consumes 3cycles/hop; 1 cycle for link delay and 2 within each router. The routers have 2 virtual channels per port, each 10 flits deep, and implement wormhole flow control.

**Workloads:** Both synthetic and realistic workloads were simulated. The traditional synthetic traffic patterns, *uniform random* and *bit-complement* represent nominal and worst-case traffic for this design. These patterns were augmented with the *P8D* pattern, proposed by Joshi et al. [79], designed as a best-case for staged or hierarchical networks where traffic is localized to individual regions. In P8D, nodes are assigned to one of 8 groups, made up of topologically adjacent nodes and nodes only send random traffic

within the group.

In these synthetic workloads, all packets contain data payloads of 512-bits, representing four flits of data in the baseline electrical NoC.

Realistic workload traces were captured for a 64-core CMP running PARSEC benchmarks [67]. The Netrace trace dependency tracking infrastructure was used to ensure realistic packet interdependencies are expressed as in a true, full-system CMP system [97]. The traces were captured from a CMP composed of 64 in-order cores with 32-KB, private L1I and L1D caches and a shared 16MB LLC. Coherence among the L1 caches was maintained via a MESI protocol. A 150 million cycle segment of the PARSEC benchmark “region of interest” was simulated. Packet sizes for realistic workloads vary bimodally between 1 and 5 data flits for miss request/coherence traffic and cache line transfers.

### 5.6.2 Synthetic Workload Results

In Figure 5.10, the LumiNOC design is compared against the electrical and Clos networks under uniform random, bit complement, and *P8D*. The figure shows the low-load latencies of the LumiNOC design are much lower than the competing designs. This is due primarily to the lower diameter of the LumiNOC topology, destinations within one subnet are one “hop” away while those in a second subnet are two. At higher bandwidths, however, the overheads of arbitration and lower link utilization due to collisions cause the saturation bandwidth to be lower than other designs. We note that, as the figure shows, higher bandwidth is achievable through adding LumiNOC layers; there is no such simple tweak to lower its latency at low loads. Adding layers to achieve higher bandwidth does come with the cost of increasing power and area, however, as shown in Table 5.3 even the 4-layer LumiNOC design consumes roughly equal power as the Clos LTBw network while providing significantly lower latency.

The different synthetic traffic patterns bring out interesting relationships. On the *P8D*




Figure 5.10: Synthetic workloads showing LumiNOC vs. Clos LTBw and electrical network.

pattern, which is engineered to have lower hop counts, all designs have universally lower latency than on other patterns. However, while both the electrical and LumiNOC network have around 25% lower low-load latency than uniform random, Clos only benefits by a few percent from this optimal traffic pattern. At the other extreme, the electrical network experiences a 50% increase in no-load latency under the bit-complement pattern compared to uniform random while both Clos and the LumiNOC network are only marginally affected. This is due to the LumiNOC having a worst-case hop count of 2 and not all routes go through the central nodes as in the electrical network. Instead, the intermediate nodes are well distributed through the network under this traffic pattern. However, as the best-case hop count is also 2 with this pattern, the LumiNOC network experiences more contention and the saturation bandwidth is decreased as a result.

### 5.6.3 Realistic Workload Results

Figure 5.11 shows the performance of the LumiNOC network in 1-, 2- and 4-layers, normalized against the performance of the baseline electrical NoC. This shows that per-

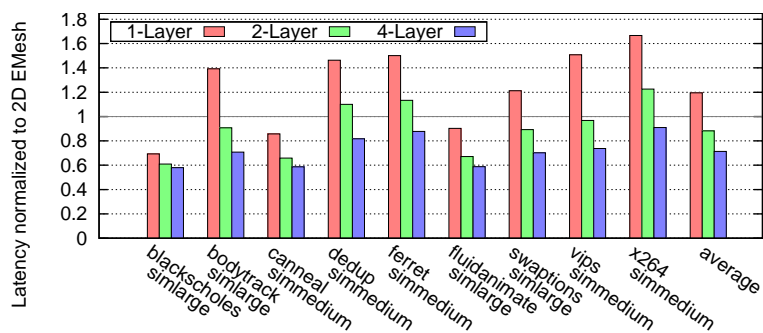


Figure 5.11: Message Latency in PARSEC benchmarks for LumiNOC compared to electrical network.

formance improves substantially with more than 1 layer: by 10% for 2-layers and 25% for 4-layers. These results can be explained by examining the bandwidth-latency curves in Figure 5.10. The overall average injection rates for the PARSEC benchmarks are of the order of 1%, so they benefit from our design’s low latency while being less affected by the lower saturation bandwidth.

#### 5.6.4 Power Efficiency

On realistic workloads, 2 layers are sufficient to achieve better performance than an electrical mesh, and still have lower overall power usage than the best competing photonic network, Clos. By going with 4 layers, power is roughly on the same order as Clos, but we can see performance improvements of 25% over the electrical mesh. Although Joshi et al. [79] do not report results for realistic workloads, we see similar performance to our baseline electrical network under uniform random in the low bandwidth regions, so there is some potential for comparison.

#### 5.6.5 Discussion

While implementing this design and exploring various means of increasing performance, it became clear that a challenge our design faces is the overhead associated with



arbitration packets and dealing with conflicts due to our statistical arbitration design. Of note is that collision detection in Ethernet networks has been phased out in favor of a switched fabric. However, on-chip networks have some advantages that may allow us to succeed where Ethernet struggled. On chip, subnets are small enough to prevent severe contention and long exponential back off periods. Furthermore, the availability of a global clock source allows other synchronizing mechanisms to assist in preventing collisions. We feel there is opportunity to further improve the performance of the LumiNOC design.

Despite arbitration contention, realistic workloads perform very well on both the 2 and 4 layer designs and power is better than or on par with other leading photonic network designs.

## 5.7 Summary

Photonic NoCs are a promising replacement for electrical NoCs in future many-core processors. In this chapter, prior photonic NoCs are analyzed, with an eye towards efficient system power utilization and low-latency. The analysis of prior photonic NoCs reveals that power inefficiencies are mainly caused by channel over-provisioning, unnecessary optical loss due to topology and photonic device layout and power overhead from the separated arbitration channels and networks. LumiNOC addresses these issues by adopting a shared-channel, photonic on-chip network to efficiently utilize power, achieving a high performance and scalable interconnect with extremely low latency. Optical loss in LumiNOC is reduced by partitioning the global network into multiple smaller sub-networks. Unlike the other photonic NoC designs, LumiNOC leverages the same channels for arbitration, and data transmission to further efficiently utilize the photonic resources. The simulations show that LumiNOC reduces the average latency by 50% at low-loads under synthetic traffic, while achieving power savings compared with the best alternative design.

## 6. PROJECTION OF SILICON PHOTONICS INTEGRATION

In this chapter, the area and power of silicon ring-based photonic transceiver on CMOS 65nm process described on chapter 4 are first summarized. Based on these design results and transistor technology roadmap from ITRS, the area for a 128-node photonic network-on-chip with both flip-chip bonding and 3D TSV integration are estimated for different CMOS process technologies. It utilizes the constant current density technique as the optimization methodology to find the transceiver circuitry power efficiency [98].

### 6.1 Chip Area Estimation for 128-node PNoC

According the transceiver circuits area (Chapter 4) in CMOS 65nm and the CMOS transistor technology roadmap (Table 7.2) from ITRS, the transceiver circuits area are scaled down for different CMOS process technologies and summarized in Table 7.1. The task is to survey the feasibility of integrating 128-node photonic network-on-chip Chip-Multi-Processors (CMPs) on a silicon die through 3D Through-Silicon-Via (TSV) integration or on different dies using flip-chip bonding technology. The PNoC are interconnected by a global silicon waveguide routing through all the 128 nodes. 65 wavelengths are multiplexed in the waveguide, with 64 wavelengths for the data transmission and 1 wavelength for clock distribution. Therefore, each node needs a dedicate set of 130 silicon rings attached on the waveguide to fulfil communication with the other nodes (65 transceiver ring pairs). Assuming each ring at least requires 5 metal pads for flip-chip bonding, or 5 TSVs to enable 3D integration of integrated circuits, therefore, total 288 equally distributed pads or vias are required to route signal between the CMOS transistors and silicon photonic devices. Table 7.3 estimates the necessitate via/pad area versus the CMOS circuits area. For the flip-chip bonding integration with metal pad size of  $45\mu\text{m}$  and pitch of  $60\mu\text{m}$ , the pads take  $298\text{mm}^2$  which matches the circuit area when process scales down to 32nm.

<b>TX Area (<math>\mu\text{m}^2</math>)</b>	<b>65nm</b>	<b>45nm</b>	<b>32nm</b>	<b>22nm</b>	<b>16nm</b>
<b>Ring Driver</b>	6000	4200	2700	1450	880
<b>Ring Tuning</b>	22800	14600	11000	7450	6130
<b>SerDes/Clocking/PRBS</b>	6000	3680	1860	930	580
<b>Total</b>	34800	22480	15560	9830	7590
<b>RX Area (<math>\mu\text{m}^2</math>)</b>	<b>65nm</b>	<b>45nm</b>	<b>32nm</b>	<b>22nm</b>	<b>16nm</b>
<b>TIA</b>	3600	3400	1800	1050	760
<b>Sense Amplifier/Retiming</b>	3600	1800	900	420	225
<b>Ring Tuning</b>	22800	14600	11000	7450	6130
<b>Eye Monitor/PRBS Checker/PRBS Counter</b>	8400	4120	2100	1030	525
<b>Total</b>	38400	23920	15800	9950	7640

Table 6.1: Area of building blocks in silicon ring-based photonic transceiver.

Moving to the TSV technology with via pitch of  $60\mu\text{m}$ , the CMOS technology needs to scale down to 16nm, otherwise, it leads to an area constraint at CMOS circuits side. TSV with via pitch of  $30\mu\text{m}$  requires a more aggressive CMOS size scaling down, which is not surveyed in this chapter.

## 6.2 Silicon Ring Based Transceiver Energy Efficiency Projection

In this section, a constant current density technique for the optimization methodology is utilized to find the photonic transceiver circuits power efficiency under different CMOS technologies. In order to achieve accurate modeling results for the transceiver circuitry, the model utilize the normalized transistor transconductance ( $g_m/W$ ), capacitances ( $C_{gg}/W$ ) and output conductances ( $g_{ds}/W$ ).

Transmitter is designed such that the serialization and pre-driver circuits have transition time of one third of the bit period to avoid excessive inter-symbol interference. Lower fan-outs is required at higher circuitry data rates due to a higher percentage of self-loading capacitance, leading to large transistor sizes and more pre-amplifier stages with increased power to satisfy the transition time specification. However, power of final output driver

Gate Length	65nm	45nm	32nm	22nm	16nm
EOT (nm) <sup>1</sup>	1.5	1.3	1.2	0.9	0.72
Nominal Saturation Current <sup>2</sup> ( $\mu\text{A}/\mu\text{m}$ )	900	980	1020	1333	1654
Nominal Power Supply (V)	1	1	1	0.9	0.77
NMOS $\tau$ <sup>3</sup> (ps)	1.6	1.2	0.87	0.66	0.3
Gate Capacitance (F/ $\mu\text{m}$ )	NA	7.40E-16	5.73E-16	6.37E-16	4.71E-16
Parasitic Overlap Cap (F/ $\mu\text{m}$ )	NA	2.40E-16	NA	2.40E-16	1.70E-16
C <sub>g</sub> Total	12E-16	9.8E-16	8.13E-16	7.94E-16	6.70E-16
C <sub>g</sub> (F/ $\mu\text{m}^2$ )	1.85E-14	2.18E-14	2.54E-14	3.61E-14	4.19E-14
f <sub>t</sub> (GHz)	190	180	180	200	228
MIM Cap (fF/ $\mu\text{m}^2$ )	2	2	4	7	10
Poly Res ( $\Omega/\square$ )	200-300	200-300	200-300	200-300	200-300

1. Equivalent Oxide Thickness
2. Nominal saturation current drive is defined as the MOSFET drain current at room temperature with the gate bias and the drain bias set equal to the nominal power-supply voltage; all MOSFET device dimensions are assumed to be at their nominal/target values. Nominal PMOS saturation current drive value is assumed to be 40-50% of the nominal NMOS saturation current-drive value.
3. Intrinsic delay metric for NMOS device; PMOS CV/I metric assumed to scale proportionally.

Table 6.2: Technology roadmap for CMOS transistors.

	65nm	45nm	32nm	22nm	16nm
Single Transceiver Area (mm <sup>2</sup> )	0.0732	0.0464	0.0314	0.0198	0.0152
Total TRX area for PNoC* (128 node and 65 wavelength WDM) (mm <sup>2</sup> )	609	386	261	165	127
FC Bonding Pad Area (mm <sup>2</sup> ) (Pad Pitch: 60 $\mu\text{m}$ )	298	298	298	298	298
TSV Area (mm <sup>2</sup> ) (Via Pitch: 40 $\mu\text{m}$ )	132	132	132	132	132
TSV Area (mm <sup>2</sup> ) (Via Pitch: 30 $\mu\text{m}$ )	75	75	75	75	75

\*Assume a Photonic Network-on-Chip (PNoC) architecture with 128 tiles connected with waveguide bus containing 64 data wavelength and 1 clock wavelength.

Table 6.3: 128-node PNoC area estimation, via/pad area vs circuits area.

stage and tuning circuitry remains constant due to the limited scalability of the photonic ring device.

TIA as the optical receiver front-end dominates the receiver circuits power, bandwidth and noise performance. Fig. 6.1 shows the schematic of an inverter-based single-stage resistive shunt feedback TIA in common-source configuration used in Chapter 4. The dc transimpedance gain and the input impedance  $Z_i$  can be calculated as equation 6.1 and 6.2 [26].




Figure 6.1: Schematic of a single-stage inverter-based resistive shunt feedback CMOS TIA with a photodetector.

$$Z_{TIA} = \frac{g_m - g_f}{g_f(g_m + g_{ds})} \quad (6.1)$$

$$Z_i = \frac{g_{ds} + g_f}{g_f(g_m + g_{ds})} \quad (6.2)$$

where  $g_f=1/R_f$ ,  $g_m=\alpha/I_d$ , and  $g_{ds}=g_m/A_0$ . The parameter  $A_0$  is the intrinsic gain

of the transistors. The parameters  $R_f$  denotes the feedback resistance and  $I_d$  the drain current. Finally,  $g_m$  is the total nMOS and pMOS small-signal transconductance, and  $g_{ds}$  is the total nMOS and pMOS transistor output conductance at the bias point.

Due to the limited gain of the single-stage inverter, the input resistance of the shunt-feedback TIA cannot be lowered down to a value very small. As the result, a dominant pole at the input node limits the TIA bandwidth due to the large parasitic capacitance of the bonding pad or silicon through via [99]. The data rate scaling needs to meet receiver's SNR standard. For the high data rate, TIA needs to use relatively smaller  $R_f$  but maintain the constant transconductance in the amplifier stage to achieve relatively smaller input resistance, extending bandwidth. Extra stages of inverter-based TIA is needed to achieve the same transimpedance gain for meeting the SNR requirement but at the cost of degraded energy-efficiency.

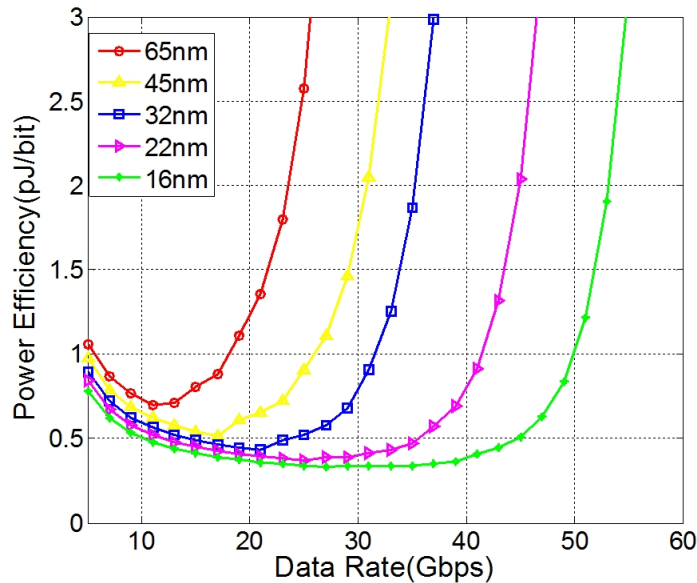


Figure 6.2: Transceiver circuitry power efficiency vs. data rate under different CMOS technologies.

Transceiver's power scaling with different CMOS technologies is shown in Fig. 6.2. Scaling from 65nm to 16nm has allowed the data rate to increase from 18 to 50 Gb/s at power efficiency of sub-1pJ/bit. Comparing the link at 20Gb/s, transistor scaling down achieves 50% improvement in power efficiency relative to the 65nm CMOS process.

## 7. CONCLUSION

The rapid expansion in data communication due to the increased multimedia applications and cloud computing services requires the energy-efficient, high-speed interconnects between two communication nodes. Although the long-haul WDM optical transmission already achieved very large aggregate bandwidth ( 100Gb/s), the energy consumption due to the transceiver circuits becomes a major concern. On the other hand, with the advance of the photonic devices, optical link offers a promising alternative solution relative to the conventional electrical link for the inter/intra-chip interconnects due to the optical channel's negligible frequency dependent loss. There is the potential to fully leverage photonic technology advances with low-power transceiver circuits for high-efficiency electrical-optical transduction. In this thesis, Chapter 3 presented an energy-efficient optical receiver front-end circuit architecture which achieves high data rates and low group-delay variation by leveraging a novel transformer-based transconductance boosting technique to meet the low-power optical long-haul transmission requirements. The boosted transconductance effectively decreased the input impedance of the transimpedance amplifier, thus isolating the photodiode's large parasitic capacitance. Considerable bandwidth extension has been achieved with but at no significant noise degradation or power consumption overhead. Further bandwidth extension is achieved through series inductive peaking to isolate the photodetector capacitance from the TIA input. The optimum choice of series inductive peaking value and key transformer parameters for bandwidth extension and jitter minimization is well analyzed. The TIA achieves a  $53 \text{ dB}\Omega$  single-ended transimpedance gain and  $21.3 \text{ pA}/\sqrt{\text{Hz}}$  average input-referred noise current spectral density. Total chip power including output buffering is 28.2 mW from a 2.5 V supply, with the core TIA consuming 8.2 mW, and the chip area including pads is  $960 \mu\text{m} \times 780 \mu\text{m}$ .



Recently, optical interconnects are also introduced into the inter/intra-chip communication with the advance of silicon compatible photonic devices. To meet the bandwidth demands of next-generation high-performance computing systems, silicon photonic transceiver circuits are presented in Chapter 4 for a silicon ring resonator-based optical interconnect architecture in a 1V standard 65nm CMOS technology. A high-swing ( $2V_{pp}$  and  $4V_{pp}$ ) drivers with non-linear pre-emphasis is incorporated to address the ring modulator issues which include a relatively low modulation bandwidth due to the slow carrier injection speed of the p-i-n junction serving as the electrical-optical transduction interface. The temperature and process variability issue is also addressed by employing an automatic bias-based tuning for resonance wavelength stabilization. At receiver side, an optical forwarded-clock adaptive inverter-based transimpedance amplifier (TIA) trades-off power for varying link budgets by employing an on-die eye monitor and scaling the TIA supply for the required sensitivity. At 5Gb/s operation, the  $4V_{pp}$  transmitter achieves 12.7dB extinction ratio with 4.04mW power consumption, excluding laser power, when driving wire-bonded modulators designed in a 130 nm SOI process, while a 0.28nm tuning range is obtained at  $6.8\mu\text{W}/\text{GHz}$  efficiency with the bias-based tuning scheme implemented with the  $2V_{pp}$  transmitter. When tested with a wire-bonded 150fF p-i-n photodetector, the receiver achieves -12.7dBm sensitivity at a  $\text{BER}=10^{-15}$  and consumes 2.2mW at 8Gb/s. Testing with an on-die test structure emulating a low-capacitance waveguide photodetector yields  $16\mu\text{App}$  sensitivity at 10Gb/s and more than 40% power reduction with higher input current levels.

To meet energy-efficient performance demands, the computing industry has moved to chip-multi-processors (CMPs) internally interconnected via networks-on-chip (NoC) to meet growing intra-chip communication needs. Achieving scaling performance as core counts increase to the hundreds in future CMPs will require high performance, yet energy-efficient interconnects. Silicon nanophotonics is a promising replacement for electronic

on-chip interconnect due to its high bandwidth and low latency, however, the static power needed for the laser and ring thermal tuning can be high. Finally, in Chapter 5, a novel nano-phonic NoC architecture, called LumiNOC is proposed with optimization for high performance and power-efficient interconnects. This work makes three primary contributions. First, instead of the conventional globally distributed photonic channels, which requires relatively high off-chip laser power, a novel channel sharing arrangement is proposed to interconnect sub-sets of cores into photonic subnets. Second, a novel, purely photonic, distributed arbitration mechanism, statistical collision detection and avoidance (SCDA) is proposed to achieve extremely low-latency without degrading the interconnection performance. In a 64-node NoC under synthetic traffic, LumiNOC enjoys 50% less latency at low loads versus other reported photonic NoCs, and  $\sim 25\%$  less latency versus the electrical 2D mesh NoCs on realistic workloads. Third, the proposed photonic network architecture leverages the same wavelengths for channel arbitration and parallel data transmission, allowing efficient utilization of the photonic resources, without dedicated arbitration channels or networks which lower efficiency or add latency to the system. Under the same ideal throughput, LumiNOC achieves laser power reduction of 78%, and overall power reduction of 44% versus competing designs.

## REFERENCES

- [1] G. L. Wojcik, D. Yin, A. R. Kovsh, A. E. Gubenko, I. L. Krestnikov, S. S. Mikhlin, D. A. Livshits, D. A. Fattal, M. Fiorentino, and R. G. Beausoleil, "A single comb laser source for short reach WDM interconnects," in *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, vol. 7230 of *Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series*, Feb. 2009.
- [2] *Semiconductor Industry Association (SIA), International Technology Roadmap for Semiconductors 2009 Edition, 2009.*
- [3] B. Kim and V. Stojanovic, "An energy-efficient equalized transceiver for rc-dominant channels," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 6, pp. 1186–1197, 2010.
- [4] E. Mensink, D. Schinkel, E. Klumperink, E. Van Tuijl, and B. Nauta, "Power efficient gigabit communication over capacitively driven rc-limited on-chip interconnects," *Solid-State Circuits, IEEE Journal of*, vol. 45, no. 2, pp. 447–457, 2010.
- [5] S.-W. Tam, E. Socher, A. Wong, and M. C. F. Chang, "A simultaneous tri-band on-chip rf-interconnect for future network-on-chip," in *VLSI Circuits, 2009 Symposium on*, pp. 90–91, 2009.
- [6] M. Anders, H. Kaul, S. Hsu, A. Agarwal, S. Mathew, F. Sheikh, R. Krishnamurthy, and S. Borkar, "A 4.1tb/s bisection-bandwidth 560gb/s/w streaming circuit-switched 88 mesh network-on-chip in 45nm cmos," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2010 IEEE International*, pp. 110–111, 2010.
- [7] I. Young, E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, and P. Chang, "Optical i/o technology for tera-scale computing," in *Solid-State Circuits*

- Conference - Digest of Technical Papers, 2009. ISSCC 2009. IEEE International*, pp. 468–469,469a, 2009.
- [8] X. Wang, J. A. Martinez, M. Nawrocka, and R. Panepucci, “Compact thermally tunable silicon wavelength switch: Modeling and characterization,” *Photonics Technology Letters, IEEE*, vol. 20, no. 11, pp. 936–938, 2008.
- [9] M. Lipson, “Compact electro-optic modulators on a silicon chip,” *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 12, no. 6, pp. 1520–1526, 2006.
- [10] M. Reshotko, B. Block, B. Jin, and P. Chang, “Waveguide Coupled Ge-on-Oxide Photodetectors for Integrated Optical Links,” in *The 2008 5th IEEE International Conference on Group IV Photonics*, pp. 182–184, 2008.
- [11] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, “12.5 Gbit/s carrier-injection-based silicon microring silicon modulators,” in *Conference on Lasers and Electro-Optics (CLEO) 2007*, pp. 1–2, 2007.
- [12] A. Narasimha, B. Analui, Y. Liang, T. Sleboda, and C. Gunn, “A fully integrated 4x10gb/s dwdm optoelectronic transceiver in a standard 0.13 um cmos soi,” in *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pp. 42–586, 2007.
- [13] A. Liu, L. Liao, D. Rubin, J. Basak, H. Nguyen, Y. Chetrit, R. Cohen, N. Izhaky, and M. Paniccia, “High-speed silicon modulator for future vlsi interconnect,” in *Integrated Photonics and Nanophotonics Research and Applications / Slow and Fast Light*, p. IMD3, Optical Society of America, 2007.
- [14] L. Liao, A. Liu, J. Basak, H. Nguyen, and M. Paniccia, “Silicon photonic modulator and integration for high-speed applications,” in *IEEE Intentional Electron Devices Meeting (IEDM)*, 2008.

- [15] B. G. Lee, A. V. Rylyakov, W. M. J. Green, S. Assefa, C. W. Baks, R. Rimolo-Donadio, D. M. Kuchta, M. H. Khater, T. Barwicz, C. Reinholm, E. Kiewra, S. M. Shank, C. L. Schow, and Y. A. Vlasov, "Four- and eight-port photonic switches monolithically integrated with digital CMOS logic and driver circuits," *IEEE-OSA Optical Fiber Communications Conference*.
- [16] J. E. Roth, S. Palermo, N. C. Helman, D. P. Bour, D. A. B. Miller, and M. Horowitz, "An optical interconnect transceiver at 1550nm using low-voltage electroabsorption modulators directly integrated to CMOS," *IEEE-OSA Journal of Lightwave Technology*, vol. 25, pp. 3739–3747, Dec 2007.
- [17] Z. Peng, D. Fattal, M. Fiorentino, and R. Beausoleil, "CMOS-compatible microring modulators for nanophotonic interconnect," in *Integrated Photonics Research, Silicon and Nanophotonics (IPRSN)*, July 2010.
- [18] G. Li, X. Zheng, J. Yao, H. Thacker, I. Shubin, Y. Luo, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "High-efficiency 25Gb/s CMOS ring modulator with integrated thermal tuning," *8th IEEE Intentional Conference on Group IV Photonics (GFP)*, vol. 4.
- [19] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, "Silicon optical modulators," *Nature Photonics*, vol. 4.
- [20] A. V. Krishnamoorthy, X. Zheng, G. Li, J. Yao, T. Pinguet, A. Mekis, H. Thacker, I. Shubin, Y. Luo, K. Raj, and J. E. Cunningham, "Exploiting CMOS manufacturing to reduce tuning requirements for resonant optical devices," *IEEE Photonics Journal*, vol. 3, pp. 567–579, Jun 2011.
- [21] J. Liu, D. Pan, S. Jongthammanurak, D. Ahn, C. Hong, M. Beals, L. Kimerling, J. Michel, A. T. Pomerene, C. Hill, M. Jaso, K.-Y. Tu, Y. K. Chen, S. Patel, M. Rasras, A. White, and D. Gill, "Waveguide-integrated ge p-i-n photodetectors on soi

- platform,” in *Group IV Photonics, 2006. 3rd IEEE International Conference on*, pp. 173–175, 2006.
- [22] C. Gunn, G. Masini, J. Witzens, and G. Capellini, “Cmos photonics using germanium photodetectors,” *ECS Transactions*, vol. 3, no. 7, pp. 17–24, 2006.
- [23] A. O. Splett, T. Zinke, B. Schueppert, K. Petermann, H. Kibbel, H.-J. Herzog, and H. Presting, “Integrated optoelectronic waveguide detectors in sige for optical communications,” 1995.
- [24] S. M. Park and H.-J. Yoo, “1.25-gb/s regulated cascode cmos transimpedance amplifier for gigabit ethernet applications,” *Solid-State Circuits, IEEE Journal of*, vol. 39, no. 1, pp. 112–121, 2004.
- [25] S. B. Amid, C. Plett, and P. Schvan, “Fully differential, 40 Gb/s regulated cascode transimpedance amplifier in 0.13 $\mu$ m SiGe BiCMOS technology,” in *IEEE BCTM*, pp. 33–36, 2010.
- [26] C. Kromer, G. Sialm, T. Morf, M. L. Schmatz, F. Ellinger, D. Erni, and H. Jackel, “A low-power 20-GHz 52-dB $\Omega$  transimpedance amplifier in 80-nm CMOS,” *IEEE Journal of Solid-State Circuits*, vol. 39, pp. 885–894, Jun 2004.
- [27] Z. Lu, K. S. Yeo, W. M. Lim, M. A. Do, and C. C. Boon, “Design of a CMOS broadband transimpedance amplifier with active feedback,” *IEEE Trans. On VLSI System*, vol. 18, pp. 461–472, Mar 2010.
- [28] X. Li, S. Shekhar, and D. J. Allstot, “ $G_m$ -boosted common-gate LNA and differential colpitts VCO/QVCO in 0.18- $\mu$ m CMOS,” *IEEE J. Solid-State Circuits*, vol. 40, pp. 2609–2619, Dec 2005.
- [29] C. Li and S. Palermo, “A low-power, 26-GHz transformer-based regulated cascode transimpedance amplifier in 0.25- $\mu$ m SiGe BiCMOS,” in *IEEE BCTM*, pp. 83–86,

2011.

- [30] J. Jin and S. S. Hsu, "Bandwidth enhancement with low group-delay variation for 40-Gb/s transimpedance amplifier," *IEEE J. Solid-State Circuits*, vol. 43, pp. 1449–1457, Jun 2008.
- [31] J. Kim and J. F. Buckwalter, "A 40-Gb/s transimpedance amplifier in 0.18- $\mu\text{m}$  CMOS technology," *IEEE Trans. On Circuits and Syst. I, Reg. Papers*, vol. 57, pp. 1964–1972, Aug 2010.
- [32] B. Analui and A. Hajimiri, "Bandwidth enhancement for transimpedance amplifiers," *IEEE J. Solid-State Circuits*, vol. 39, pp. 1263–1270, Aug 2004.
- [33] S. Shekhar, J. S. Walling, and D. J. Allstot, "Bandwidth extension techniques for CMOS amplifiers," *IEEE J. Solid-State Circuits*, vol. 41, pp. 2424–2439, Nov 2006.
- [34] S. Galal and B. Razavi, "40-Gb/s amplifier and ESD protection circuit in 0.18- $\mu\text{m}$  CMOS technology," *IEEE J. Solid-State Circuits*, vol. 39, pp. 2389–2396, Dec 2004.
- [35] W. Z. Chen, Y. L. Cheng, and D. S. Lin, "A 1.8-V 10-Gb/s fully integrated cmos optical receiver analog front-end," *IEEE J. Solid-State Circuits*, vol. 40, pp. 1388–1396, Jun 2005.
- [36] X. Li, S. Shekhar, and D. Allstot, "Gm-boosted common-gate lna and differential colpitts vco/qvco in 0.18-  $\mu\text{m}$  cmos," *Solid-State Circuits, IEEE Journal of*, vol. 40, no. 12, pp. 2609–2619, 2005.
- [37] S. S. Mohan, M. D. M. Hershenson, S. P. Boyd, and T. H. Lee, "Bandwidth extension in CMOS with optimized on-chip inductors," *IEEE J. Solid-State Circuits*, vol. 35, pp. 346–355, Mar 2000.

- [38] C.-F. Liao and S.-I. Liu, "A 40gb/s transimpedance-agc amplifier with 19db dr in 90nm cmos," in *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pp. 54–586, 2007.
- [39] B. Razavi, "Design of Integrated Circuits for Optical Communications," in *New York: McGraw-Hill*, 2002.
- [40] Z. Lu, K. S. Yeo, J. Ma, M. A. Do, W. M. Lim, and X. Chen, "Broad-band design techniques for transimpedance amplifiers," *IEEE Trans. On Circuits and Syst. I, Reg. Papers*, vol. 54, pp. 590–600, Mar 2007.
- [41] R. Swoboda and H. Zimmermann, "11gb/s monolithically integrated silicon optical receiver for 850nm wavelength," in *Solid-State Circuits Conference, 2006. ISSCC 2006. Digest of Technical Papers. IEEE International*, pp. 904–911, 2006.
- [42] J. R. Long, "Monolithic transformers for silicon RF IC design," *IEEE J. Solid-State Circuits*, vol. 35, pp. 105–108, Sep 2000.
- [43] C.-Y. Wang, C.-S. Wang, and C.-K. Wang, "An 18-mw two-stage cmos transimpedance amplifier for 10 gb/s optical application," in *Solid-State Circuits Conference, 2007. ASSCC '07. IEEE Asian*, pp. 412–415, 2007.
- [44] J. Mullrich, H. Thurner, E. Mullner, J. Jensen, W. Stanchina, M. Kardos, and H.-M. Rein, "High-gain transimpedance amplifier in inp-based hbt technology for the receiver in 40-gb/s optical-fiber tdm links," *Solid-State Circuits, IEEE Journal of*, vol. 35, no. 9, pp. 1260–1265, 2000.
- [45] C. Q. Wu, E. A. Sovero, and B. Massey, "40-GHz transimpedance amplifier with differential outputs using InP-InGaAs heterojunction bipolar transistors," *IEEE J. Solid-State Circuits*, vol. 38, pp. 1518–1523, Sep 2003.



- [46] I. Young, E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, and P. Chang, "Optical I/O technology for tera-scale computing," *IEEE Journal of Solid-State Circuits*, vol. 45, pp. 235–248, Jan 2010.
- [47] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Express*, vol. 15, pp. 430–436, Jan 2007.
- [48] B. R. Moss, M. G. C. Sun, J. Shainline, J. S. Orcutt, J. C. Leu, M. Wade, Y. Chen, K. Nammari, X. Wang, H. Li, R. Ram, M. A. Popovic, and V. Stojanovic, "A 1.23pJ/b 2.5Gb/s monolithically integrated optical carrier-injection ring modulator and all-digital driver circuit in commercial 45nm SOI," in *IEEE ISSCC Dig. Tech. Papers*, pp. 126–127, Feb 2013.
- [49] G. Li, X. Zheng, J. Yao, H. Thacker, I. Shubin, Y. Luo, K. Raj, J. E. Cunningham, and A. V. Krishnamoorthy, "25gb/s 1v-driving cmos ring modulator with integrated thermal tuning," *Opt. Express*, vol. 19, pp. 20435–20443, Oct 2011.
- [50] C. Li, R. Bai, A. Shafik, E. Tabasy, G. Tang, C. Ma, C.-H. Chen, Z. Peng, M. Fiorentino, P. Chiang, and S. Palermo, "A ring-resonator-based silicon photonics transceiver with bias-based wavelength stabilization and adaptive-power-sensitivity receiver," in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2013 IEEE International*, pp. 124–125, 2013.
- [51] C.-H. Chen, C. Li, R. Bai, A. Shafik, M. Fiorentino, Z. Peng, P. Chiang, S. Palermo, and R. Beausoleil, "Hybrid integrated dwdm silicon photonic transceiver with self-adaptive cmos circuits," in *Optical Interconnects Conference, 2013 IEEE*, pp. 122–123, 2013.
- [52] M. Georgas, J. Leu, B. Moss, C. Sun, and V. Stojanovic, "Addressing link-level design tradeoffs for integrated photonic interconnects," in *Custom Integrated Circuits*

- Conference (CICC), 2011 IEEE*, pp. 1–8, 2011.
- [53] F. Liu, D. Patil, J. Lexau, P. Amberg, M. Dayringer, J. Gainsley, H. Moghadam, X. Zheng, J. Cunningham, A. Krishnamoorthy, E. Alon, and R. Ho, “10-gbps, 5.3-mw optical transmitter and receiver circuits in 40-nm cmos,” *Solid-State Circuits, IEEE Journal of*, vol. 47, no. 9, pp. 2049–2067, 2012.
- [54] C. Sun, E. Timurdogan, M. Watts, and V. Stojanovic, “Integrated microring tuning in deep-trench bulk cmos,” in *Optical Interconnects Conference, 2013 IEEE*, pp. 54–55, 2013.
- [55] J. Buckwalter, X. Zheng, G. Li, K. Raj, and A. Krishnamoorthy, “A monolithic 25-gb/s transceiver with photonic ring modulators and ge detectors in a 130-nm cmos soi process,” *Solid-State Circuits, IEEE Journal of*, vol. 47, no. 6, pp. 1309–1322, 2012.
- [56] C. Li and S. Palermo, “A low-power 26-ghz transformer-based regulated cascode sigebicmos transimpedance amplifier,” *Solid-State Circuits, IEEE Journal of*, vol. 48, no. 5, pp. 1264–1275, 2013.
- [57] J. Proesel, C. Schow, and A. Rylyakov, “25gb/s 3.6pj/b and 15gb/s 1.37pj/b vcsel-based optical links in 90nm cmos,” in *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pp. 418–420, 2012.
- [58] S. Palermo, A. Emami-Neyestanak, and M. Horowitz, “A 90 nm cmos 16 gb/s transceiver for optical interconnects,” *Solid-State Circuits, IEEE Journal of*, vol. 43, no. 5, pp. 1235–1246, 2008.
- [59] M. Georgas, J. Orcutt, R. Ram, and V. Stojanovic, “A monolithically-integrated optical receiver in standard 45-nm soi,” in *ESSCIRC (ESSCIRC), 2011 Proceedings of the*, pp. 407–410, 2011.

- [60] S. Palermo and M. Horowitz, “High-speed transmitters in 90nm cmos for high-density optical interconnects,” in *Solid-State Circuits Conference, 2006. ESSCIRC 2006. Proceedings of the 32nd European*, pp. 508–511, 2006.
- [61] J. S. Orcutt, B. Moss, C. Sun, J. Leu, M. Georgas, J. Shainline, E. Zraggen, H. Li, J. Sun, M. Weaver, S. Urošević, M. Popović, R. J. Ram, and V. Stojanović, “Open foundry platform for high-performance electronic-photonics integration,” *Opt. Express*, vol. 20, pp. 12222–12232, May 2012.
- [62] P. Dong, W. Qian, H. Liang, R. Shafiq, D. Feng, G. Li, J. E. Cunningham, A. V. Krishnamoorthy, and M. Asghari, “Thermally tunable silicon racetrack resonators with ultralow tuning power,” *Opt. Express*, vol. 18, pp. 20298–20304, Sep 2010.
- [63] J. S. Orcutt, A. Khilo, C. W. Holzwarth, M. A. Popović, H. Li, J. Sun, T. Bonifield, R. Hollingsworth, F. X. Kärtner, H. I. Smith, V. Stojanović, and R. J. Ram, “Nanophotonic integration in state-of-the-art cmos foundries,” *Opt. Express*, vol. 19, pp. 2335–2346, Jan 2011.
- [64] W. J. Dally and B. Towles, “Route Packets, Not Wires: On-Chip Interconnection Networks,” in *The 38th International Design Automation Conference (DAC)*, pp. 684–689, 2001.
- [65] P. Gratz, C. Kim, R. McDonald, S. W. Keckler, and D. C. Burger, “Implementation and Evaluation of On-Chip Network Architectures,” in *2006 IEEE International Conference on Computer Design (ICCD)*, pp. 477–484, Oct. 2006.
- [66] J. Howard, S. Dighe, S. R. Vangal, G. Ruhl, N. Borkar, S. Jain, V. Erraguntla, M. Konow, M. Riepen, M. Gries, G. Droege, T. Lund-Larsen, S. Steibl, S. Borkar, V. K. De, and R. V. D. Wijngaart, “A 48-Core IA-32 Processor in 45 nm CMOS Using On-Die Message-Passing and DVFS for Performance and Power Scaling,” *IEEE Journal of Solid-State Circuits*, vol. 46, pp. 173–183, Oct 2011.

- [67] C. Bienia, S. Kumar, J. P. Singh, and K. Li, “The PARSEC Benchmark Suite: Characterization and Architectural Implications,” in *The 17th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, October 2008.
- [68] P. Gratz, K. Sankaralingam, H. Hanson, P. Shivakumar, R. McDonald, S. W. Keckler, and D. C. Burge, “Implementation and Evaluation of a Dynamically Routed Processor Operand Network,” in *1st ACM/IEEE International Symposium on Networks-on-Chip (NoCS)*, pp. 7–17, May 2007.
- [69] J. Kim, D. Park, T. Theocharides, N. Vijaykrishnan, and C. R. Das, “A Low Latency Router Supporting Adaptivity for On-Chip Interconnects,” in *2005 Design Automation Conference*, pp. 559–564, June 2005.
- [70] A. Liu, L. Liao, D. Rubin, H. Nguyen, B. Ciftcioglu, Y. Chetrit, N. Izhaky, and M. Paniccia, “High-speed optical modulation based on carrier depletion in a silicon waveguide,” *Optics Express*, vol. 15, no. 2, pp. 660–668, 2007.
- [71] C. Holzwarth, J. Orcutt, H. Li, M. Popovic, V. Stojanovic, J. Hoyt, R. Ram, and H. Smith, “Localized Substrate Removal Technique Enabling Strong-Confinement Microphotonics in Bulk Si CMOS Processes,” in *Conference on Lasers and Electro-Optics*, pp. 1–2, 2008.
- [72] L. C. Kimerling, D. Ahn, A. Apsel, M. Beals, D. Carothers, Y.-K. Chen, T. Conway, D. M. Gill, M. Grove, C.-Y. Hong, M. Lipson, J. Michel, D. Pan, S. S. Patel, A. T. Pomerene, M. Rasras, D. K. Sparacin, K.-Y. Tu, A. E. White, and C. W. Wong, “Electronic-photonic integrated circuits on the CMOS platform,” in *Silicon Photonics*, pp. 6–15, 2006.
- [73] A. Narasimha, B. Analui, Y. Liang, T. Sleboda, and C. Gunn, “A Fully Integrated 4 10-Gb/s DWDM Optoelectronic Transceiver Implemented in a Standard 0.13 m

- CMOS SOI Technology,” in *The IEEE International Solid-State Circuits Conference*, pp. 42–586, 2007.
- [74] I. Young, E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, and P. Chang, “Optical I/O Technology for Tera-scale Computing,” in *The IEEE International Solid-State Circuits Conference*, pp. 468–469, 2009.
- [75] G. Hendry, S. Kamil, A. Biberman, J. Chan, B. Lee, M. Mohiyuddin, A. Jain, K. Bergman, L. Carloni, J. Kubiawicz, L. Oliner, and J. Shalf, “Analysis of Photonic Networks for a Chip Multiprocessor using Scientific Applications,” in *The 3rd ACM/IEEE International Symposium on Networks-on-Chip (NOCS)*, pp. 104–113, 2009.
- [76] A. Shacham, K. Bergman, and L. P. Carloni, “On The Design of a Photonic Network-On-Chip,” in *The First International Symposium on Networks-on-Chip (NOCS)*, p. 53–64, 2007.
- [77] A. Shacham, K. Bergman, and L. P. Carloni, “Photonic NoC for DMA Communications in Chip Multiprocessors,” in *The 15th Annual IEEE Symposium on High-Performance Interconnects*, pp. 29–38, 2007.
- [78] A. Shacham, K. Bergman, and L. P. Carloni, “Photonic Networks-On-Chip for Future Generations of Chip Multiprocessors,” *IEEE Transactions on Computers*, vol. 57, no. 9, pp. 1246–1260, 2008.
- [79] A. Joshi, C. Batten, Y.-J. Kwon, S. Beamer, I. Shamim, K. Asanovic, and V. Stojanovic, “Silicon-Photonic Clos Networks for Global On-Chip Communication,” in *The 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip (NOCS)*, pp. 124–133, 2009.

- [80] N. Kirman, M. Kirman, R. Dokania, J. Martinez, A. Apsel, M. Watkins, and D. Albonese, “Leveraging Optical Technology in Future Bus-Based Chip Multiprocessors,” in *The 39th Annual IEEE/ACM International Symposium on Microarchitecture (Micro)*, pp. 492–503, 2006.
- [81] A. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J. Cunningham, “Computer Systems Based on Silicon Photonic Interconnects,” *Proceedings of the IEEE*, vol. 97, no. 7, pp. 1337–1361, 2009.
- [82] Y. Pan, J. Kim, and G. Memik, “FlexiShare: Channel Sharing for an Energy-Efficient Nanophotonic Crossbar,” in *The 16th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pp. 1–12, 2010.
- [83] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, “Firefly: Illuminating future network-on-chip with nanophotonics,” in *36th International Symposium on Computer Architecture (ISCA)*, 2009.
- [84] D. Vantrease, N. Binkert, R. Schreiber, and M. H. Lipasti, “Light Speed Arbitration and Flow Control for Nanophotonic Interconnects,” in *42nd Annual IEEE/ACM International Symposium on microarchitecture*, pp. 304–315, 2009.
- [85] 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,” in *35th International Symposium on Computer Architecture (ISCA)*, pp. 153–164, 2008.
- [86] P. Koka, M. O. McCracken, H. Schwetman, X. Zheng, R. Ho, and A. V. Krishnamoorthy, “Silicon-Photonic Network Architectures for Scalable, Power-Efficient Multi-Chip Systems,” in *37th International Symposium on Computer Architecture (ISCA)*, pp. 117–128, 2010.

- [87] Y. H. Kao and H. J. Chao, “BLOCON: A Bufferless Photonic Clos Network-on-Chip Architecture,” in *5th ACM/IEEE International Symposium on Networks-on-Chip (NoCS)*, pp. 81–88, May 2011.
- [88] A. Biberman, K. Preston, G. Hendry, N. Sherwood-droz, J. Chan, J. S. Levy, M. Lipson, and K. Bergman, “Photonic Network-on-Chip Architectures Using Multilayer Deposited Silicon Materials for High-Performance Chip Multiprocessors,” *ACM Journal on Emerging Technologies in Computing Systems*, vol. 7, no. 2, pp. 1305–1315, 2011.
- [89] G. Hendry, E. Robinson, V. Gleyzer, J. Chan, L. P. Carloni, N. Bliss, and K. Bergman, “Time-Division-Multiplexed Arbitration in Silicon Nanophotonic Networks-On-Chip for High-Performance Chip Multiprocessors,” *Journal of Parallel and Distributed Computing*, vol. 71, pp. 641–650, May 2011.
- [90] R. W. Morris and A. K. Kodi, “Power-Efficient and High-Performance Multi-Level Hybrid Nanophotonic Interconnect for Multicores,” in *4th ACM/IEEE International Symposium on Networks-on-Chip (NoCS)*, pp. 207–214, May 2010.
- [91] J. Xue, A. Garg, B. Ciftcioglu, J. Hu, S. Wang, I. Savidis, M. Jain, R. Berman, P. Liu, M. Huang, H. Wu, E. Friedman, G. Wicks, and D. Moore, “An intra-chip free-space optical interconnect,” in *Proceedings of the 37th annual international symposium on Computer architecture*, ISCA ’10, (New York, NY, USA), pp. 94–105, ACM, 2010.
- [92] A. Abousamra, R. Melhem, and A. Jones, “Two-Hop Free-Space Based Optical Interconnects for Chip Multiprocessors,” in *5th ACM/IEEE International Symposium on Networks-on-Chip (NoCS)*, pp. 89–96, May 2011.
- [93] P. Gratz and S. W. Keckler, “Realistic Workload Characterization and Analysis for Networks-on-Chip Design,” in *The 4th Workshop on Chip Multiprocessor Memory Systems and Interconnects (CMP-MSI)*, 2010.

- [94] M. R. T. Tan, P. Rosenberg, S. Mathai, J. Straznicky, L. Kiyama, J. S. Yeo, M. McLaren, W. Mack, P. Mendoza, and H. P. Kuo, "Photonic Interconnects for Computer Applications," in *Communications and Photonics Conference and Exhibition (ACP), 2009 Asia*, pp. 1–2, 2009.
- [95] H. Kim, P. Ghoshal, B. Grot, P. V. Gratz, and D. A. Jimenez, "Reducing network-on-chip energy consumption through spatial locality speculation," in *5th ACM/IEEE International Symposium on Networks-on-Chip (NoCS)*, pp. 233–240, 2011.
- [96] S. Prabhu, B. Grot, P. Gratz, and J. Hu, "Ocin tsim-DVFS Aware Simulator for NoCs," *Proc. SAW*, vol. 1, 2010.
- [97] J. Hestness and S. Keckler, "Netrace: Dependency-Tracking Traces for Efficient Network-on-Chip Experimentation," tech. rep., Technical Report TR-10-11, The University of Texas at Austin, Department of Computer Science, <http://www.cs.utexas.edu/~netrace>, 2010.
- [98] A. Palaniappan and S. Palermo, "Power efficiency comparisons of interchip optical interconnect architectures," *Circuits and Systems II: Express Briefs, IEEE Transactions on*, vol. 57, no. 5, pp. 343–347, 2010.
- [99] G. Katti, M. Stucchi, D. Velenis, B. Soree, K. De Meyer, and W. Dehaene, "Temperature-dependent modeling and characterization of through-silicon via capacitance," *Electron Device Letters, IEEE*, vol. 32, no. 4, pp. 563–565, 2011.