

Photonic Links -- From Theory to Automated Design

Krishna Settaluri



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2019-8

<http://www2.eecs.berkeley.edu/Pubs/TechRpts/2019/EECS-2019-8.html>

April 23, 2019

Copyright © 2019, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

Acknowledgement

This PhD came at a very opportune time in my life. On the one hand, I began this journey with a more refined sense of judgment, thought, and knowledge than, say, when I entered undergrad. On the other, I still consider myself young enough to mold, adapt, and grow to the people and environment around me. Putting that together, the people I bonded with during the last five years are special in not only being unique, exceptionally talented, loving, and hilarious but also having the ability to teach and change me for the better.

Photonic Links – From Theory to Automated Design

by

Krishna Tej Settaluri

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Vladimir Stojanović, Chair

Professor Eli Yablonovitch

Professor Costas Grigoropoulos

Fall 2018

Photonic Links – From Theory to Automated Design

Copyright © 2018

by

Krishna Tej Settaluri

Abstract

Photonic Links – From Theory to Automated Design

by

Krishna Tej Settaluri

Doctor of Philosophy in Engineering - Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Vladimir Stojanović, Chair

Recent advancements in silicon photonics show great promise in meeting the high bandwidth and low energy demands of emerging applications. However, a key gating factor in ensuring this necessity is met is the utilization of a link design methodology which transcends the various levels in the hierarchy, ranging from the device and platform level up to the systems level. In this dissertation, a comprehensive methodology for link design will be introduced which takes a two-prong approach to tackling the issue of silicon photonic link efficiency. Namely, a fundamentals-based first principles approach to link optimization will be introduced and validated. In addition, physical design trade-offs connecting levels in the architectural hierarchy will also be studied and explored. This culminates in an intermediate goal of this dissertation, which is the first-ever design and verification of a full silicon photonic interconnect on a 3D integrated electronic-photonic platform. To proceed and further enable the rapid exploration of the link design architectural space, the analog macros for a majority of this dissertation were auto-generated using the Berkeley Analog Generator (BAG). With these key design tools and framework, performance bottlenecks and improvements for silicon photonic links will be analyzed and, from this analysis, the motivation for a new, single comparator-based PAM4 receiver architecture shall emerge. This architecture not only showcases the tight bond in dependency between high-level link specifications and low level device parameters, but also shows the importance of physical design constraints alongside fundamental theory in influencing end-to-end link performance.

To Mom and Dad.

Contents

Contents	ii
List of Figures	v
List of Tables	ix
1 Introduction	1
2 Background	3
2.1 Silicon Photonic Links	3
2.1.1 Photonic Building Blocks	4
2.1.2 Optical Communication Links	6
2.1.3 Circuit Design Challenges and Methodology	7
2.2 Silicon Photonic Integration Platforms	7
2.2.1 3D Integration Using Thru-Oxide Vias	8
2.2.2 3D Integration Using Flip Chip μ Bumps	8
2.3 Analog Circuit Design Challenges and Automation	9
2.3.1 BAG Architecture and Flow	9
2.4 Conclusion	13
3 End-to-End Optical Link Design Methodology	15
3.1 Introduction	15
3.2 Optical Link Modeling	17
3.2.1 High-Level Receiver Abstraction	17
3.2.2 Gain-Bandwidth product	18
3.2.3 Voltage Amplifiers	18
3.2.4 Transimpedance Amplifier	19
3.2.5 Sampler Model in Receiver	20
3.3 Macro-Parameter Derivations	25
3.3.1 Sensitivity calculation	25
3.3.2 Energy per bit	26
3.3.3 Model inputs and optimization variables	29

3.3.4	Model purpose and limitations	29
3.4	Sensitivity and energy limits	29
3.4.1	Front end noise limit	29
3.4.2	Limits at high data rates	31
3.4.3	Limits at low datarates	31
3.4.4	E/b power laws	32
3.5	Observations in Scaling and Technology	32
3.5.1	Improvements in Photonics and Interconnects	33
3.5.2	Improvements in Photonics+CMOS	33
3.6	Link Framework and First-Principles Summary	36
3.7	Conclusion	36
4	Link Framework Application for NRZ Front-End Design	38
4.1	Introduction	38
4.2	Model Application for 65nm Heterogenously Integrated Photonic Platform	39
4.2.1	Technology overview	39
4.2.2	Single slicer case ($M=1$)	39
4.2.3	Multiple slicer case ($M \geq 1$)	40
4.3	Schematic designs of model results	42
4.3.1	5Gbps optical receiver	42
4.3.2	Active-CTLE enhanced 5Gbps optical receiver	43
4.3.3	DDR 25Gbps optical receiver	44
4.3.4	QDR 25Gbps optical receiver	46
4.3.5	Switched QDR 25Gbps optical receiver	47
4.4	Conclusion	48
5	3D Integrated Silicon Photonic Interconnects	49
5.1	Introduction	49
5.2	Overview	50
5.3	3D Integration of CMOS and Photonics	51
5.4	Chip Architecture and Optical Link System	52
5.4.1	Transmitter Design	53
5.4.2	Receiver Design	55
5.4.3	Thermal Tuner Design	57
5.4.4	Link Implementation and Test setup	57
5.5	Conclusion	59
6	Single-Comparator PAM4 Architecture	63
6.1	Introduction	63
6.2	PAM4 Introduction and Link Trade-Offs	63
6.3	Single Comparator PAM4 Receiver	67

6.3.1	Motivation	68
6.3.2	Proposed Architecture and Formulation	69
6.3.3	Design Methodology and Theoretical Performance	74
6.3.4	End-to-End Photonic Co-Simulation Results	79
6.4	Conclusion	81
7	Acacia System Design	83
7.1	Introduction	83
7.2	Acacia System Design	83
7.2.1	Front-End Floorplanning and PDK Generation Constraints	83
7.2.2	Receiver AFE Design and Insights	91
7.2.3	Transmitter Design and Automation	92
7.2.4	New Transmit-Side Thermal Tuning	93
7.2.5	Tx-Rx Self Test Setup	94
7.2.6	Putting together the Acacia System	95
7.3	Results	96
7.3.1	On-Chip Clock Network	96
7.3.2	Rx AFE and Self Test Characterization	98
7.3.3	25.6Gbps Self Test Link	99
7.4	Conclusion	99
8	Conclusion	101
8.1	Thesis Contributions	101
8.2	Future Work and Final Thoughts	102
	Bibliography	103

List of Figures

2.1	The microring modulator enables OOK encoding while also enabling wavelength selectivity.	4
2.2	The photodetector is responsible for generating an output current dependent on the input optical power.	6
2.3	A WDM link composed of many transmit and receive side rings is an attractive solution which allows high bandwidth density.	7
2.4	A cross section and top view of the TOV are shown.	8
2.5	The design flow of a new block within the BAG framework is shown.	9
2.6	The design flow with feedback from the output of simulations enables rapid iteration. [4].	10
2.7	The schematic, transient, and example floorplan of the StrongArm is shown.	11
2.8	The output of the schematic and layout generators produce DRC and LVS clean instances of the StrongArm.	12
2.9	The framework allows for push-button instantiation and verification to generate quick, correct instances.	12
2.10	Pushing an instance through a full design flow is simple within this codified environment.	13
3.1	Demonstrated Optical Link Efficiencies [6-15], Against Objectives From [6]. Further information at linksurvey.eecs.berkeley.edu	16
3.2	Optical Link System Overview	16
3.3	StrongArm Sampler Schematic	20
3.4	Sampler Timing Evaluation Breakdown	23
3.5	Ideal Transfer Function of A System with Equalization	30
3.6	Technology Dependent Performance Prediction	36
4.1	Optimal Energy Per Bit Versus Data rate For Optimal Topologies With Parameters From Table 3.1. Only One Slicer is Allowed in This Case	40
4.2	Optimal Energy Per Bit Versus Data rate For Optimal Topologies With Parameters from Table 3.1 with the Possibility of Multiple Slicers . .	41
4.3	5Gbps Receiver Topology	42

4.4	5Gbps Model-Predicted Receiver Topology with Active-CTLE	44
4.5	25Gbps Optical Receiver With Variable Interleaving Stages	45
4.6	Switching Time-Interleaved 25Gbps QDR Receiver	47
5.1	A comparison of the integration technology categories shows pros and cons for both.	50
5.2	An example diagram of the full WDM link using a heterogeneous platform shows independent photonic and CMOS wafers on both the Tx and Rx side.	51
5.3	A cross section of the 3D TOV heterogeneous integration process shows the photonic and CMOS wafers and associated connections.	52
5.4	A full system of the heterogeneously integrated EPHI system.	53
5.5	The Tx Macro consists of high speed serializers and drivers shift the ring resonance.	54
5.6	The Tx microring resonance characteristics and eye diagrams are shown.	55
5.7	The receiver AFE as well as the photodiode is shown.	55
5.8	The Rx photonic ring responsivity and response versus frequency.	56
5.9	Measured receiver average photo-current sensitivity over different data rates and BER bathtub curves for both receiver slices.	57
5.10	The thermal tuner block diagram used to control the microring resonance is shown.	58
5.11	The progression of the transient eye along with the resonance location for the thermal tuner.	59
5.12	The test setup of the EPHI chip contains the Tx and Rx macros connected by a 100m fiber reel.	60
5.13	Full optical link with optical power budget and performance.	61
5.14	Electrical energy breakdown for the Tx and Rx macros in a 5Gb/s link.	62
5.15	Comparison with previous work.	62
6.1	The traditional PAM4 architecture comprises of three comparators after the AFE to slice the 4-level eye.	64
6.2	Removing all constraints on the receiver architecture shows that the PAM4 architecture is superior to NRZ under <i>particular system constraints</i>	65
6.3	A comparison of the E/b for a PAM4 versus NRZ link comprised of only a single interleaving way for slicing shows the benefits of PAM4 over NRZ.	66
6.4	A comparison of the E/b for a PAM4 versus NRZ link comprised of only a single comparator for slicing shows the benefits of PAM4 over NRZ.	67
6.5	Full, end-to-end drawing of the photonic link along with a point to the critical node – the primary concern of this chapter.	68

6.6	Preliminary link performance results show the benefit of scaling down the sampler capacitance by $3\times$	69
6.7	Behavior of a PAM4 receiver with a single comparator yield “NRZ-equivalent” behavior.	70
6.8	The StrongArm Sense Amplifier schematic along with two sets of waveforms (red and green) showing the output due to a large and small input signal, respectively.	70
6.9	A fully differential PAM4 front-end will have two complementary outputs centered about the common-mode.	72
6.10	A time-to-digital circuit takes an input the differential outputs of the StrongArm sense amplifier.	73
6.11	The sample waveforms of the TDC block show pulse width’s dependency on the output StrongArm waveforms.	73
6.12	The TDC receiver architecture composed of the AFE and a single interleaving way composed of the StrongArm, D2S, and new TDC (from Figure 6.10).	74
6.13	Non-idealities in the AFE result in a voltage ratio between the BigBit and LittleBit that is smaller than theoretical $3\times$	76
6.14	The double tail sense amplifier has benefits in the new PAM4 context that outweigh the traditional StrongArm sense amplifier.	77
6.15	The double tail sense amplifier evaluation-time “gain” may be characterized using Equations 6.9 and 6.10. The MATLAB simulation results are plotted here.	78
6.16	A comparison of the new, TDC-based PAM4 receiver and the traditional, three-comparator architecture show the potential benefits when viewing the link energy consumption. These results reflect not only the three-comparator difference, but also any secondary limitations on g_m due to the presence of the TDC.	80
6.17	The end-to-end photonic co-simulation schematic shows the Tx driver, photonic components, and the CMOS receiver.	81
6.18	The receiver input current eye is shown. This signal was produced by a CMOS transmitter driving a photonic microring. The modulated light goes into a VerilogA photodiode to produce this eye.	81
6.19	The output of the Rx AFE is shown. This signal subsequently traverses into a single slicer prior to digitizing.	82
7.1	The AFE’s signal-critical blocks are placed first to ensure optimized performance and minimize path lengths and parasitics.	85
7.2	The full Rx front-end layout is composed of the signal-critical blocks, deserializers, and DACs.	86

7.3 The bumps (shown using the artistically rendered red squares) are spaced evenly, with the purpose of interposing with the components on the photonic reticle. 87

7.4 Bump pitch variations result in very simple changes to the macro script in order to produce the two flavors of Rx above (drawn to scale relative to one another). 88

7.5 Routing channels, although constrain the maximum width of the Rx front-end, provide much needed space to allow proper routing between the front-end digital bits and the digital backend. 89

7.6 The body-connected and non-body connected transistor flavors are vastly different from the standpoint of the generator. This image shows the two flavors, with similar widths and 5 fingers. 90

7.7 The Rx AFE (simplified) schematic shows the main blocks along the critical signal path. 91

7.8 The Tx-side AC driver uses a push/pull architecture to maximize the voltage swing across the microring modulator. 92

7.9 The closed loop thermal tuner serves to ensure that the microring resonance remains in lock. This scheme has the added benefit of not requiring an additional drop-port. 93

7.10 The self test schematic shows the full data path to characterize the AFEs. 94

7.11 The full Acacia system realizes a high speed, high performance end-to-end optical link with all necessary critical peripherals within the system. 95

7.12 By modifying the capacitive DAC code, the output frequency of the global clock distribution can be modified according to the plot above. 97

7.13 The link’s performance is shown for low frequency and operating in MSB-mode. 98

7.14 The link’s performance is shown for low frequency and operating in MSB-mode. 99

7.15 The bathtub curve of the Acacia link operating at 25.6Gbps is shown. 100

List of Tables

3.1	Model inputs and optimization variables	28
3.2	Minimum power laws for E/b limits dependence	33
3.3	Performance Comparison of Model-Predicted and Schematic-Simulated Optical Receivers	35
6.1	This table summarizes the alternate interpretation of the PAM encod- ing scheme, using a single comparator's "timing information".	72

Acknowledgments

This PhD came at a very opportune time in my life. On the one hand, I began this journey with a more refined sense of judgment, thought, and knowledge than, say, when I entered undergrad. On the other, I still consider myself young enough to mold, adapt, and grow to the people and environment around me. Putting that together, the people I bonded with during the last five years are special in not only being unique, exceptionally talented, loving, and hilarious but also having the ability to teach and change me for the better.

To begin, none of this would be possible without my advisor, Professor Stojanović. He gave me a chance and an opportunity to succeed at a time when I needed it the most. His knowledge was inspiring and his optimism has been the bane of my existence for the last half decade. And, most recently, his incredible dance moves have shown me that I have much to learn. Oh, and thanks for not caring about the many coffee-related issues in Cory.

During my PhD, I had the wonderful opportunity to engage with colleagues and faculty members both inside and outside of the circuits space. Thanks to Professor Eli Yablonovitch, whose brain-scratching questions left me dazed for days and to Professor Ming Wu for valuable insight into the photonic-device space. To Chris Keraly and Nicolas Andrade, thank you for your help refining and innovating the link framework. To Professor Elad Alon, looking forward to our next adventure together. To Professor Ali Niknejad, I wish we could have played more soccer, but the times we did play were boundlessly fun. And, to Erman Timurdogan, Zhan Su, and Prof Michael Watts at MIT for support with EPHI and Acacia.

This thesis would not be possible without the kind, caring direction of my fellow BWRC lab mates. To begin, teamvlada has been an endless source of literally everything. Sen Lin and Sajjad Moazeni have been my battle buddies throughout the years. Nandish Mehta has been a loving, caring force of nature for many a tape-outs. Pavan Bhargava, you are my fiercest ally and friend. To the Acacia team composed of Sidney, Kouros, Ruocheng, Zhaokai, Eric KJ, and Nick, I really hope you learned a great deal during the process and had as much "fun" as I did. To Taehwan, Panos, and Christos thanks for your help and support over the years.

I would like to thank others at BWRC. Thank you to Greg and Luke who have never changed and I sincerely hope never will. Your support and intimidatingly strong background in circuits has been beyond helpful. To Antonio and John, thank you for your continued friendship and beards. To Bonjern, thanks for being the most excited Warriors fan I know. Nathan, thanks also for the beard and coming to my wedding. Looking forward to what's next. And thank you to Sameet, who was the best wedding officiant a guy could ask for. You are hilarious, logical, sharp, and caring. I hope there are no more pasta aglio e olios in our future.

Thank you to Chen, Mark, and Ranko for being fantastic mentors during my PhD. You have taught me many a lesson and I'm sure that that will continue.

Thank you Candy for basically running the center, for providing granola, and for plants/life/food advice. To Ajith, Brian, Fred, Melissa, Olivia, Yessica, Erin, and Leslie thanks for being the backbone of this center.

To those outside of BWRC, thank you to Vinay (and James, I guess) for still putting up with my incessant altitude humor. Thank you, Korok, for many fascinating conversations and chin (not a typo). Thank you, Jasdeep and Pingy for TiddleWinkle Pingydinkle and so much more. Thanks to the Bandits for letting me never sub out of ball games. This might be a good time to also thank Twalve for his support and ability to hold many many cars in his garage (and many hearts in his fish bowl). Thank you Kevin, Katrina, Jorge, and Dani for continued support. To Pallovy, Abhiram, Jiwon, Susan, Ankush, Riley, Anusha (blabla), Josh Hsaio, and Lakshmi, thank you all for being there for me.

I would not be an iota of the person I am now today without my family. To my dad, Dr. Raghu Kumar Settaluri, and mom, Sandhya Rani Settaluri. You two are jaw-droppingly inspiring and loving. Seriously, well done! To my dearest sister, Keertana. I'm so grateful to have you in my life and in BWRC. You will always have a special place in my heart, booboo.

Last, but certainly not least, thank you to my dearest wife, Pallavi. You have been a part of my life during all its ups and downs, yet you remain grounded, logical, and hilarious. I cannot think of a reality without you and I am endlessly in your debt. I am sure we have a lot more adventures to go on and I know we are ready for anything. And I'll be sure to take care of you anywhere we go – after all, I'm a doctor now.

-Krishna

Chapter 1

Introduction

With the recent increase in demand of high performance computing and with the emergence of the 5G mobile market, global IP traffic has steadily increased over the past 5 years and is projected to continue doing so for the foreseeable future. To meet these consumer communication bandwidth demands, modern data-centers and computers require optimization at every level in the system hierarchy. Silicon photonics integration within large-scale systems-on-chip (SoCs) has emerged as a primary contender in not only enhancing the capabilities of CMOS technologies, but also meeting the high bandwidth and low energy demands of these next-generation computing interconnects. Recent years have seen great strides in the development and commercialization of silicon photonic technologies.

However, as complexity and performance of these systems ever increases, many problems still remain and numerous new ones emerge. Specifically, to continue optimizing system efficiencies given new and ever changing process, technologies, and devices, a unified framework needs to be in place to quickly characterize and explore the design landscape. Moreover, to truly realize rapid design exploration, automation at every level in the hierarchy becomes mandatory – from high-level place and route optimization to low-level analog automation.

This thesis delves into optimizing the silicon photonic link performance by means of two parallel thrusts. First, the optimization of performance shall be approached by taking a fundamentals, theory-oriented approach. A link design framework which is versatile in process, technology, and device will be introduced in Chapter 3. This framework allows for quick optimization contingent on the CMOS as well as photonic parameters. From here, the validity of the framework will be further enforced in Chapter 4, wherein particular design points for the 65nm heterogeneously integrated photonic platform will be selected and explored. This work was published in “*First Principles Optimization of Opto-Electronic Communication Links*” in the IEEE Transactions on Circuits and Systems I: Regular Papers [1]. The authors of this work were Krishna T. Settaluri, Christopher Lalau-Keraly, Eli Yablonovitch, Vladimir Stojanovic. The author of this dissertation contributed in developing the

theory and methodology for the link design framework.

In addition, transmit-side exploration will also take place with the utilization of a new, nanoLED-based transmitter model. The nanoLED-based transmit work was published in “*Optical Antenna nanoLED Based Interconnect Design*” in the IEEE IPC 2018 [2]. The authors of this work were Nicolas M. Andrade, Krishna T. Settaluri, Seth A. Fortuna, Sean Hooten, Kevin Han, Eli Yablonovitch, Vladimir Stojanovic, and Ming C. Wu. The author’s contribution to this work were in the form of the end-to-end link model incorporation with the nanoLED transmitter model.

Next, in Chapter 5, a detailed analysis and design of an end-to-end 5Gbps NRZ link will be introduced. Here, we get a taste of the full design flow or “package” which comes about from aiming to design a full fledged optical link. This work was taped out and tested. The results were published in “*Demonstration of an Optical End-to-End Link in a 3D Integrated Electronic-Photonic Platform*” in ESSCIRC 2015, and authored by Krishna Settaluri, Sajjad Moazeni, Chen Sun, Erman Timurdogan, Michele Moresco, Zhan Su, Yu-Hsin Chen, Gerald Leake, Douglas LaTulipe, Colin McDonough, Jeremiah Hebding and Douglas Coolbaugh [3].

In the last portion of this dissertation, the performance bottleneck of optical links will be studied and it will be proven why the newly introduced single-comparator based PAM4 receiver architectures proves superior. Moreover, the rapid design flow and physical design challenges attributed to the various layers in the hierarchy will be looked into in detail.

Chapter 2

Background

In the parlance of circuit design, silicon photonics has stepped up as a clear contender in enhancing the capabilities of CMOS technologies. Indeed, photonics alongside CMOS systems can potentially improve energy efficiency and realize applications requiring higher bandwidths. Additionally, photonic links have the added benefit of having the channel loss (i.e. loss through an optical fiber) be weakly dependent on distance, allowing for better performance in long-range networks such as within the data center. Lastly, by incorporating multiple wavelengths of light within the same fiber through a technique called Wavelength Division Multiplexing (WDM), bandwidth density is greatly increased as compared with traditional, CMOS-only techniques.

In this chapter, we will introduce the foundational building blocks of silicon photonic links, namely components like photodiodes and microring resonators. Additionally, this chapter will showcase how these components fit together to enable true, end-to-end link operation. Lastly, with increasing complexity in analog design, tools that simplify the tasks of designing and laying out analog IP blocks are a treat to any mixed signal and systems engineer. This chapter will introduce the Berkeley Analog Generator (BAG), which aims to re-think the design process and redirect effort towards research and optimization of circuits rather than drawing polygons and laying out. Indeed, BAG is the foundation for a majority of this thesis and, as such, the mechanics and methodology for the tool will be “laid out” in this chapter.

2.1 Silicon Photonic Links

Silicon photonics-based links provide benefits over counterparts that use conventional, discrete optics or components. By utilizing a manufacturing platform that enables mass production and low cost, silicon photonic links and their associated designs follow an approach analogous with traditional CMOS-only systems, but with

added functionality and while retaining all the benefits. Silicon photonics also have the luxury of minimal overhead when incorporating with CMOS processes, thereby enabling very large-scale co-designs, where photonic components sit alongside CMOS circuits and transistors.

2.1.1 Photonic Building Blocks

2.1.1.1 Modulators and Microring Resonators

An optical modulator is responsible for “modulating” light on the transit side of an optical link. The mechanism for modulating the light as well as the means for encoding it are both free variables, dependent on the system architecture and purpose. However, in general, the modulator contains an electrical driver, which takes as input a digital rail-to-rail signal, which drives an optical component in a way that encodes the incoming digital sequence into variations in light intensity. The microring-based optical modulator, shown in Figure 2.1 is an attractive device in part because it enables on-off keyed (OOK) encoding easily, while also allowing for wavelength selectivity.

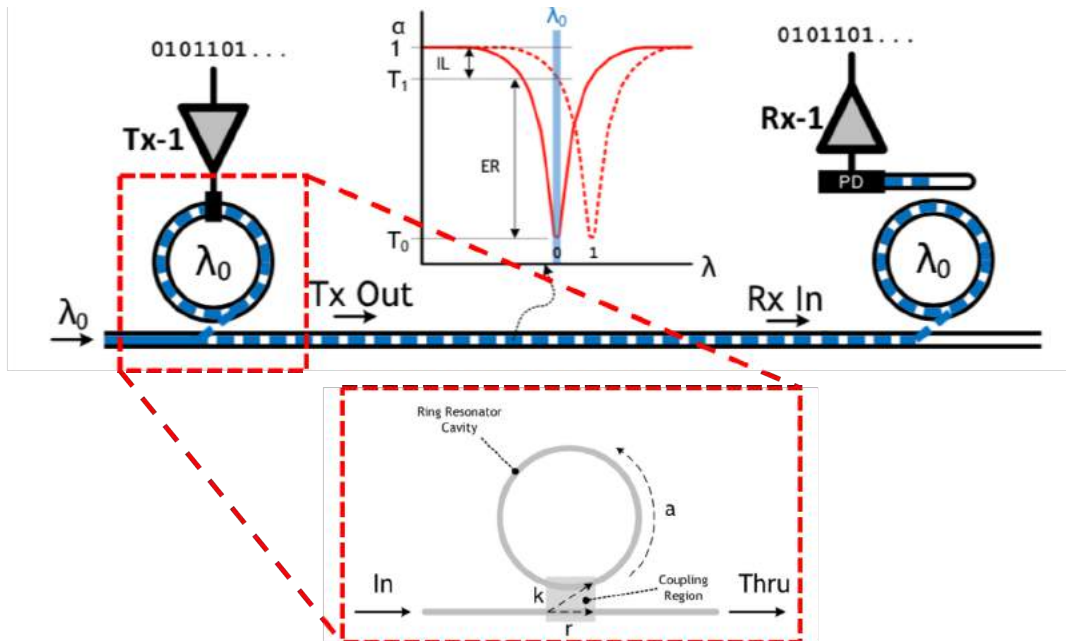


Figure 2.1: The microring modulator enables OOK encoding while also enabling wavelength selectivity.

A microring modulator, when coupled with a bus waveguide (shown as a zoom-in in Figure 2.1), acts as a notch filter that either passes or circulates light from the “In-port” to the “Thru-port”. The location of the notch in the wavelength-space

is dependent on not only the length of the ring, but also the index of refraction. The location of the notch is also very dependent on the incoming electrical stimulus. This stimulus modifies the depletion width of the p-n junctions with the microring, thereby causing a change in the effective index of refraction and thus a shift in the notch location. Additionally, the ring introduces periodic notches in the wavelength space, which is contingent on the round trip distance that light traverses before either constructively or destructively interfering with the incoming light. This is summarized with the equation below:

$$\lambda_0 = \frac{n_{eff}L}{m} \quad m = 1, 2, 3, \dots \quad (2.1)$$

Aside from the location of the resonance itself, from the context of an end-to-end optical link, a few other parameters of the modulator are also important in dictating performance and behavior. Firstly, the insertion loss (IL) of the modulator, shown in Figure 2.1, dictates the loss caused by the ring when *inserted* into the system. Secondly, the extinction ratio (ER), also shown in Figure 2.1, is the power ratio between the ring driven by a “0”-signal and “1”-signal. Together, these two factors greatly influence the performance and design of an end-to-end optical link. In general, to have the best behavior and ease the receiver-side design, smaller IL and larger ER are desirable. This manifests itself as a large optical modulator amplitude (OMA), while reducing the incoming laser power.

These specifications on the modulator become critical when optimizing end-to-end link performance. As will be evident in later chapters, the design and co-optimization of these parameters, along with the design attributes of the CMOS technology itself, drastically influence the best-case energy per bit (E/b) of the link.

2.1.1.2 Photodetectors

A photodetector (or photodiode) is a device which converts the incoming light into electrical current which may be decoded via the receiver circuitry. An example photodetector layout is shown in Figure 2.2.

The incoming waveguide provides the optical input to the photodetector. Based on the absorption properties of the detector, which are heavily dependent on the characteristics of the material composition, the detector generates a current proportional to the magnitude of the input optical signal. Indeed, because of the broad-band nature of the photodetector, a microring is generally appended before the photodiode itself to allow for wavelength selectivity.

Once again, from the context of an optical link system, particular specifications at the photodetector device level become critical in influencing the overall energy per bit. Namely, the three biggest contributors in performance are the optical bandwidth, responsivity, and added capacitance. The optical bandwidth of the photodetector,




Figure 2.2: The photodetector is responsible for generating an output current dependent on the input optical power.

or equivalently the Full Width Half Max (FWHM) of the pulse response for a linear photodiode, is a byproduct of transit time limitations of the material itself. Any degradations in generating carriers slows down rise/fall times of the incoming bit stream. The responsivity of the photodiode dictates how efficiently the device is able to generate electron/hole pairs. Lastly, the capacitance of the photodiode influences the electrical bandwidth. Contingent on the resistance seen by the photodiode, the resulting R-C time constant needs to be sufficiently high to not cause signal power roll-off. All of these specs, again combined with the CMOS properties of the receiver, influence the best-case performance of the link. Moreover, co-designing the various sub-components of the link (both CMOS and photonic) are critical in attaining the global optima.

2.1.2 Optical Communication Links

Putting together the sub-components detailed above along with an electronic backbone to encode, decode, and control the optics, a full optical link, shown in Figure 2.3 can be designed.

Figure 2.3 shows a bank of microring modulators on the transmit side and a similar bank on the receive side which drops the light onto the respective photodiodes. Notice also the presence of serializers and deserializers which aim to provide a low speed communication means to the rest of the digital backend. Additionally, the presence of thermal tuners on both sides allow the ring to maintain resonance lock

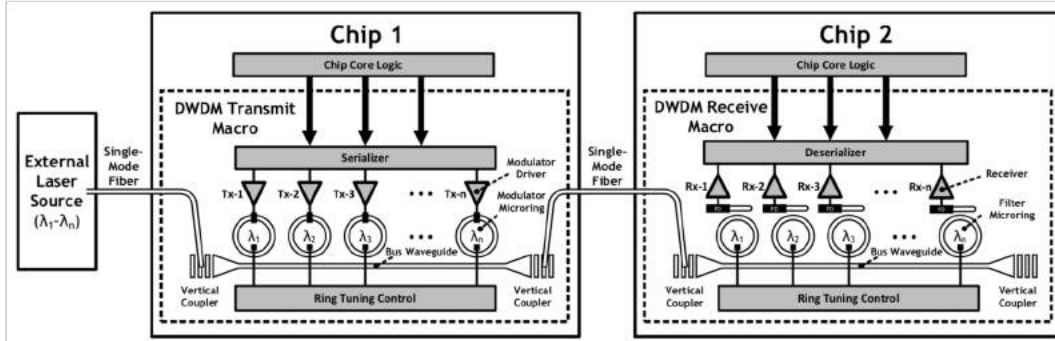


Figure 2.3: A WDM link composed of many transmit and receive side rings is an attractive solution which allows high bandwidth density.

despite thermal aggressors surrounding the system. In general, a separate drop port alongside the microring picks up fractions of the incoming light to allow for passive, background monitoring and tuning. Later in this work, during the design process for the Acacia system, a new type of thermal tuner will be introduced which does not rely on this added drop port. Rather, careful monitoring of the data sequence itself enables one to extract the appropriate resonance information to ensure locking.

2.1.3 Circuit Design Challenges and Methodology

The best case energy per bit performance of an optical link is dependent on a multitude of factors. Parameters like data rate, technology, circuit topology, parasitics, etc. all influence the behavior and design of the system. Indeed, trade-offs between the devices, the technology, and system architecture all motivate the utilization of an all-encompassing photonic plus CMOS co-design. As will be the objective of later chapters, this precise co-design approach will be enforced in producing designs with optimal performance.

However, as is painfully evident to all high speed circuit designers, layout parasitics, matching, and other layout-specific design attributes greatly influence performance and functionality. As such, a design and layout engine or framework with the objective of simplifying the design and execution of these blocks is paramount. As will be detailed in Section 2.3, this framework combined with the co-design mantra detailed above, high speed, high performance optical link systems may be realized with ease.

2.2 Silicon Photonic Integration Platforms

The integration platform with which the photonic and CMOS components intertwine is a key parameter which influencing the overall performance. Namely, the

parasitic characteristics associated with the interconnect may easily become a bottleneck in performance. This section will detail the two main platforms used in this dissertation. Both of these platforms are heterogenously integrated platforms, meaning that the CMOS and photonics are designed independently on separate wafers. Shortly thereafter, they are attached to one another either using through-oxide vias (TOVs) and C4 (μ Bumps).

2.2.1 3D Integration Using Thru-Oxide Vias

The first wafer-scale 3D integration using TOVs was developed by the SUNY College of Nanoscale Sciences (CNSE).




Figure 2.4: A cross section and top view of the TOV are shown.

The TOV, shown in Figure 2.4, connects the top metal layer of the CMOS wafer with that of the photonics. The process relies on firstly flipping the CMOS wafer, etching away the substrate to further reduce capacitance, and then punching TOVs to allow connectivity. With this approach, approximately 3fF of capacitance per TOV was observed. Indeed, this technique was utilized for demonstrating a full end-to-end link and will be detailed in a later chapter.

2.2.2 3D Integration Using Flip Chip μ Bumps

Using μ Bumps is another popular technique to allow integration of the CMOS and photonics. In this technique, small, conductive balls are first placed on the top passivation opening of the wafer. The other wafer is then flipped and aligned onto the first. Once the alignment is correct, the balls are melted and collapsed to effectively

solder the connection between the two wafers. Using this technique, past work has demonstrated 20-50fF of capacitance per bump. However, their high yield and process dependence proves this technique promising.

2.3 Analog Circuit Design Challenges and Automation

As was detailed beforehand, analog and mixed signal design motivates the necessity of an alternative means of design. The Berkeley Analog Generator (BAG) is a framework that captures the methodology of the designer within the confines of a new framework that provides many useful features. More specifically, this framework allows not only for layout parameterization (which produces push-button design rule check clean and layout-versus-schematic clean designs), but rapid iteration on designs. Function calls enable generation of new designs, calls to simulators, parsing the simulator output results, and iterating quickly based on the feedback. BAG was utilized in the Acacia system for all analog and mixed signal sub-blocks.

2.3.1 BAG Architecture and Flow

2.3.1.1 Overview

An example BAG-based design flow which an engineer may undertake is shown in Figure 2.5.

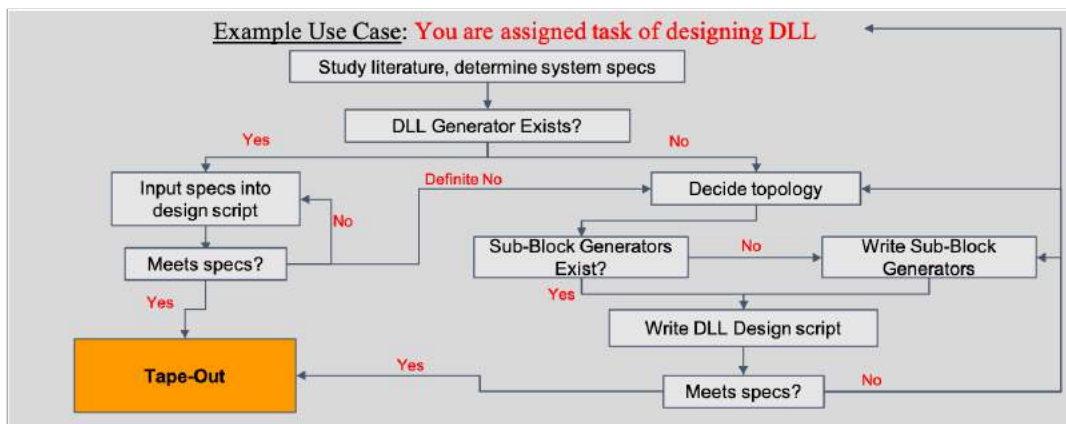


Figure 2.5: The design flow of a new block within the BAG framework is shown.

Here, the objective is to design a DLL within the BAG framework. Due to the rapid iteration abilities of the framework, along with the reusability of past designs, the designer begins by checking the existing work for a similar script to run. Additionally, due to the parameterizable attributes of BAG-generated layouts, a multitude

of DLLs may be generated with a single design script, all at the push of a button. This allows for rapid reusability, contingent on the generator writer's scope for usage. Should the original generator not meet the needs of this specific architecture, the designer may choose to write his or her own generator, specific to the methodology employed by said designer. Lastly, iteration is highly possible based on the output of characterization scripts and simulations. Verification of specifications can also be done within the framework. This feedback-based design flow information is summarized in Figure 2.6.

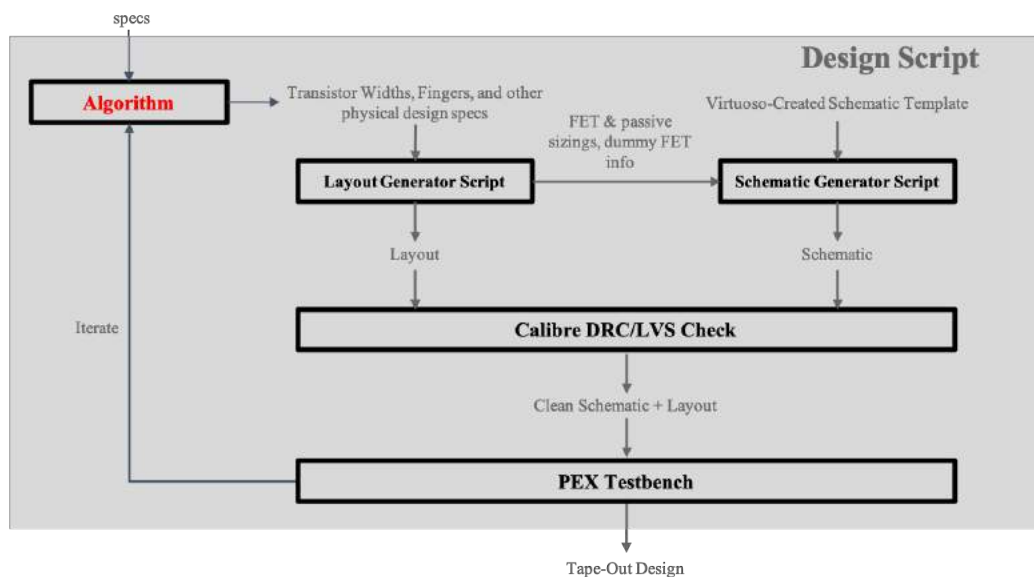


Figure 2.6: The design flow with feedback from the output of simulations enables rapid iteration. [4]

2.3.1.2 Design Example: StrongArm Sense Amplifier Generator

To further understand the design usage of the BAG framework, an example generator for a simple sub-block is outlined below. In this case, a StrongArm sense amplifier is chosen as the block of interest.

The StrongArm sense amp, shown in Figure 2.7, functions as a rudimentary analog-to-digital converter. The block takes as input two small analog voltages and generates rail-to-rail swing outputs. It does this by relying on the current discharge path of the input pair, followed by the cross-coupled action of the inverters.

Figure 2.7 not only shows the schematic, but also the example transient waveform of the outputs during the evaluation phase. Notice that the evaluation phase is composed of the initial linear (or integration) phase of the StrongArm followed by

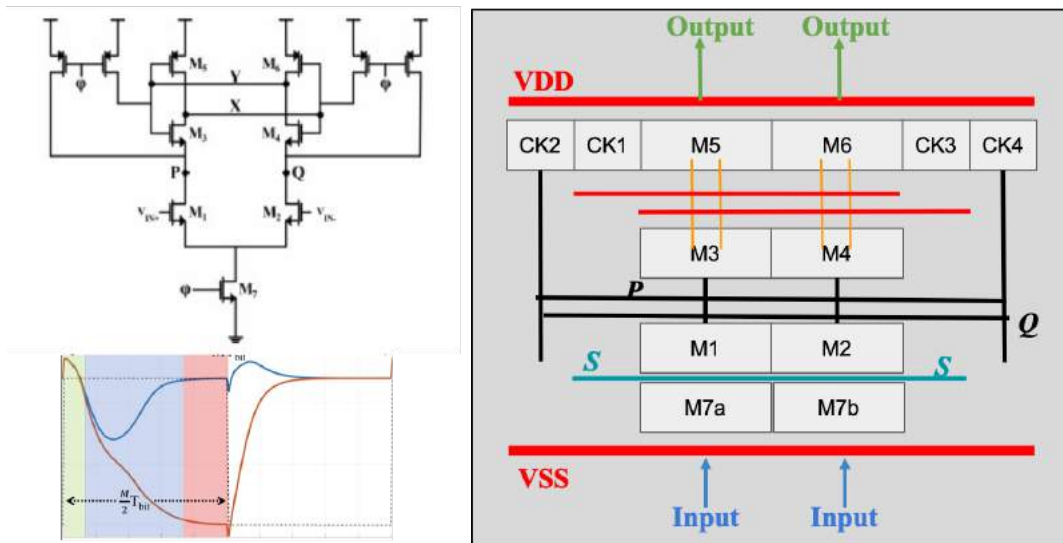


Figure 2.7: The schematic, transient, and example floorplan of the StrongArm is shown.

the evaluation phase wherein the outputs cross-couple to opposite railed values. This is then proceeded by the reset phase, wherein both outputs take a V_{DD} value.

An important first step when undertaking designing a block within the BAG framework is to draw a sample floorplan of the block (this presumes that the topology is “fixed” or needs little maneuverability). This floorplan not only shows the positions of the various transistors and important signal wires, but also shows the direction of growth for the various transistors. The drawing of the important signal wires aid in identifying critical layout-dependent characteristics, such as presuming differential matching or minimizing trace lengths. The floorplan drawing aids the designer in properly coding up routings without the mental hassle of remembering all growth conditions and layout design constraints. The floorplan is then coded within the Python-based framework as a layout generator. The code itself follows a “boilerplate” template, beginning with importing the parameter values, specifying the total width per row (within the framework, each horizontal row of transistors assumes a fixed FET width but can have variable number of fingers), drawing the basic template layout, instantiating the transistors themselves, creating the wire connections, and finally adding pins and fillers.

The schematic generator is also written by the designer, but is substantially easier to implement than the layout generator script. A simple template schematic followed by appropriately updating the parameters is sufficient to create a generator which outputs parameter-specific schematics. The layout generator script, combined with the schematic generator, produce unique schematic and layout instances dependent on an input list of parameters, such as FET widths, and number of fingers. A

sample output of the generators is shown in Figure 2.8.

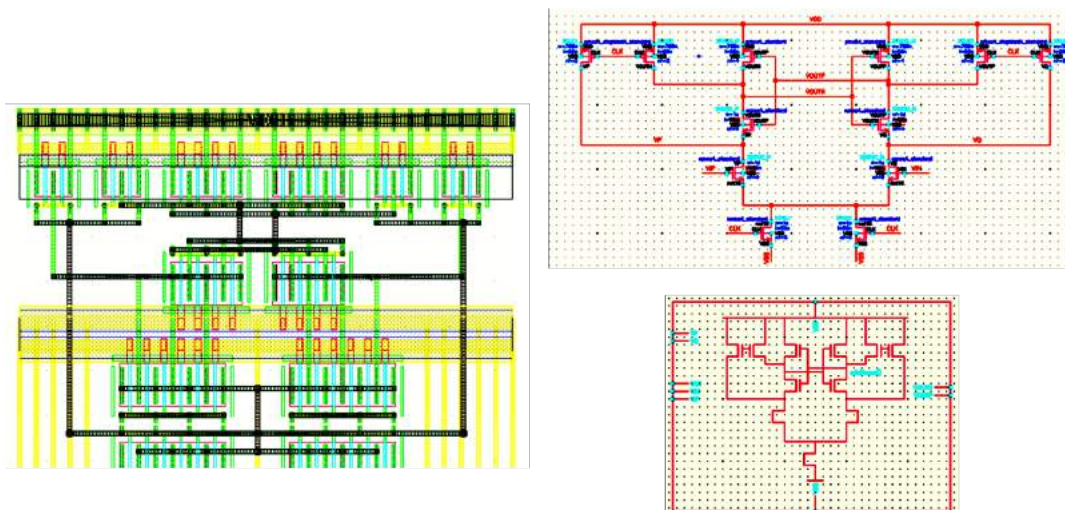


Figure 2.8: The output of the schematic and layout generators produce DRC and LVS clean instances of the StrongArm.

Indeed, with this layout and schematic generator, various instances of the StrongArm sense amplifier can be designed with the push and execution of the code. By simply modifying the dimensions of the transistors, instances such as those in Figure 2.9 can very easily be generated and verified. Notice that because the generator was written with sizing parameters incorporated within the layout, any permutation of transistor widths would easily yield DRC and LVS clean layout designs.




Figure 2.9: The framework allows for push-button instantiation and verification to generate quick, correct instances.

To further ensure that a layout and schematic are behaving properly when sim-

ulated, any particular instance can be pushed through the “entirety” of the flow, enabling rapid layout and schematic generation, verification, and post-PEX simulation to ensure the design specifications are met. This is highlighted in Figure 2.10. Input parameters, both within the layout and also within the simulation environment, can be modified and the design can proceed through the flow. The log shows the generation of the layout and schematic, run LVS, run PEX simulations, and analyze the resulting outputs. The waveforms are processed and shown in the right side of Figure 2.10.

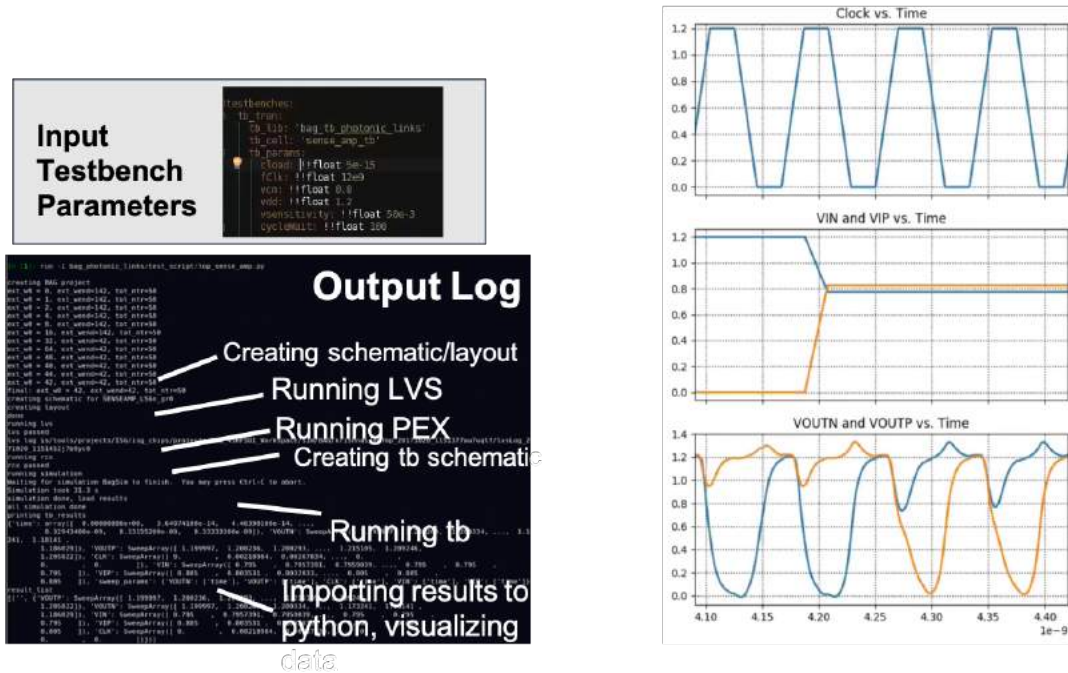


Figure 2.10: Pushing an instance through a full design flow is simple within this codified environment.

2.4 Conclusion

In this chapter, we introduced the foundational blocks of both silicon photonics and analog automation. In providing this introduction, the reader should take careful note on the power of growing vast systems which rely on the marriage of these underlying foundations. Notably, understanding the device and block level parameters that influence macro-specifications at the system level is a powerful step in designing high performance optical links. Moreover, realizing that operating at these high speed-high performance corners rely heavily on layout parameters is the key to motivating the need for analog automation. With these two in mind, the rest of this

thesis aims to operate at the intersection of these foundations – one that focuses on multi-hierarchical fundamental theory and the other which focuses on executing and optimizing beginning with the “optimized” circuits and ending with physical design.

Chapter 3

End-to-End Optical Link Design Methodology

3.1 Introduction

Optical interconnects, after having completely replaced electrical interconnects for long haul communication, are forecast to continue their expansion to shorter and shorter links, eventually bringing data directly to the processing chips, and even potentially replacing some of the longer interconnects on the chip itself [5]. This is due to several key aspects of optical interconnects: their potential for extremely high bandwidth, distance insensitivity of optical channel loss when compared with electrical, and better optical components and technology. Nevertheless, as illustrated in Fig. 3.1, in a world where energy dissipation from computing units is becoming increasingly important, optical links must still prove that they can offer a more energy efficient communication means than electrical links for shorter distances.

Commonly cited objectives for chip to chip links range in the $\sim 100\text{fJ}$ per bit, and drop to $\sim 10\text{fJ}$ per bit when considering on-chip interconnects [6]. These energy requirements, when combined with the extremely high bandwidths needed, still pose a number of challenges for optical links. The emergence of Silicon Photonics is offering new possibilities and prospects in this regard by enabling seamless integration of photonics and electronics on a single platform, thereby increasing energy efficiency. The purpose of this work is to model these links and optimize them in order to explore what limits can be reached in terms of energy efficiency and how these limits depend on the specific technology available.

Prior literature in this space has made strides in accurately modeling particular aspects of the link data path, namely the front-ends and the systems-level energy breakdowns ([7–9]). However, a proper marriage between "analog"-dominated and "digital"-dominated constraints has yet to be demonstrated. More specifically, in the context of optical receiver design, specifications on the sensitivity and power




Figure 3.1: Demonstrated Optical Link Efficiencies [6-15], Against Objectives From [6]. Further information at linksurvey.eecs.berkeley.edu.




Figure 3.2: Optical Link System Overview

of the signal are contingent on the interaction of the front-end and the follow-on samplers that ultimately convert the analog signal into a digital bit, to set the overall energy, bandwidth, gain and noise properties. Linking all of these relevant interactions together, this work shows the behavior of the full optical link under different regimes of operation from the context of energy-efficiency and noise.

In this chapter, we will introduce and analyze the optical modeling framework, beginning with a high-level link picture and slowly delving down into the various sub-components. The theory and “interface” between these blocks will also be studied. Once this foundation has been laid out, the focus will shift to the link-level, where macro-parameters will be derived and initial trade-offs will be studied. Lastly, using the framework, performance projections will be made which, in turn, gives insight into the direction of possible fabrication improvements in the future.

3.2 Optical Link Modeling

We consider a very general model for the optical link, which enables us to perform optimizations on its topology and estimate the optimal energy per bit which can be achieved at particular data rates given the technology constraints. The parameterized model topology is depicted in Fig. 3.2 and is detailed as such: the receiver element is constructed with a transimpedance front end followed by N amplification stages and terminated with a sampling unit composed of M individual samplers. The number of amplification stages N , the size of each stage, the number of sampling units M and the sizing of its transistors constitute optimization variables. There is of course a variety of other receiver topologies or variations on the one suggested. The framework we describe next will be readily extendable to these topologies.

The energy consumed in the receiver can be computed from the bias currents and circuit capacitances, and its sensitivity is determined by two constraints: a noise constraint, and a system output voltage constraint. Finally, the energy consumed by the transmitter can be calculated starting from the receiver sensitivity requirement, and back-propagating that through the data path losses to the transmitter. The total energy is the sum of the receiver and transmitter energy, which is minimized with respect to the optimization variables at hand.

3.2.1 High-Level Receiver Abstraction

The receiver is modeled as illustrated in Fig. 3.2. The front end consists of a transimpedance amplifier (TIA) that converts the input photocurrent to an output voltage signal, and is followed by N chained gain stages forming a voltage amplifier (VA) to further amplify the signal. All these amplifiers are considered to be first order stages (except for the TIA which has two poles: one from the photodiode capacitance and feedback resistor, and one from the input capacitance of the VA). The chaining of such stages causes the overall bandwidth to degrade. The bandwidth B_{chain} resultant from N first order stages of bandwidth f_S is [19]:

$$B_{chain} \sim f_S \frac{0.9}{\sqrt{N+1}} \quad (3.1)$$

We set the target end-to-end bandwidth to $0.7 \times f_{data}$, where f_{data} is the Nyquist rate of the input data stream. This implies that the bandwidth f_S of each stage must be

$$f_S > 0.7 f_{data} \frac{\sqrt{(N+2)+1}}{0.9} \quad (3.2)$$

in order to satisfy this constraint. The factor of 2 comes as a result of the two poles imposed by the TIA.

3.2.2 Gain-Bandwidth product

While the unity current gain-bandwidth of a technology is f_T , the actual gain bandwidth that is achieved in an individual gain stage that is loaded by its replica will be lower due to various parasitics and non idealities. Additionally, different gain stage topologies will yield different GBWs. For example, inductive peaking is a popular way of enhancing the bandwidth and will yield a higher GBW than simple resistively-loaded stages. Therefore we use a parameter α which describes what fraction of f_T is achieved by each individual gain stage. The GBW of a replica-loaded stage is therefore $f_a = \alpha f_T$.

3.2.3 Voltage Amplifiers

Here, we introduce the analysis of the follow-on voltage amplifiers, which helps lay the foundation for the analysis of the transimpedance amplifier stage. Every stage in the voltage amplifier is defined by input transistor gate width $W_{gate,i}$ (where "i" denotes its position in the amplifier chain), which then also defines its transconductance $g_{m,i}$, gate capacitance $C_{ox,i}$ and bias current $I_{d,i}$. To simplify the problem, we assume that g_m , C_{ox} , and I_d are simply proportional to W_{gate} , which implies that the biasing for each transistor is relatively similar– a reasonable assumption to first order. The GBW of each stage depends on the capacitance seen at the output, and in the case of simple resistively loaded stages, we have $GBW_i = g_{m,i}/(C_{out,i} + C_{in,i+1})$. We define $\beta = C_{out}/C_{in}$ as the ratio of output to input capacitance of a gain stage. Similar to α , β is dependent on the stage topology.

In the model, two factors are used to characterize the individual gain stages: $\alpha = \frac{f_a}{f_T}$, the ratio of the gain bandwidth to f_T of a replica loaded stage, and $\beta = \frac{C_{out}}{C_{in}}$, the ratio of input to output capacitance. Here we calculate α and β for simple $g_m R_L$ topology and cascode stages for the 65nm platform used.

3.2.3.1 α -factor Derivation

For a simple $g_m R_L$ topology we have

$$C_{in} = C_{ox} + AC_{gd} \quad (3.3)$$

where the second term accounts for the Miller Effect, and $C_{out} = C_{gd} + C_{ds}$. For a cascode stage, we have

$$C_{in} = C_{ox} + C_{gd} \quad (3.4)$$

Notice that the C_{GD} seen by the input does not see the Miller effect due to the intermediary FET between the input FET and the output node.

Given that $C_{ox} = 0.5fF/\mu m$, $C_{gd} = 0.2fF/\mu m$, $C_{gs} = 0.27fF/\mu m$, we have $\alpha = 0.36$ for a standard $g_m R_L$ stage and $\alpha = 0.4$ for a cascode stage.

3.2.3.2 β -factor Derivation

With the expressions given above, it is easy to show that $\beta = 0.29$ for $g_m R_l$ stages and $\beta = 0.4$ for cascode stages.

To summarize,

$$f_a = \frac{g_m}{(1 + \beta)C_{in}} \quad (3.5)$$

and we can derive the GBW of every stage as:

$$GBW_i = \frac{g_{m,i}}{C_{out,i} + C_{in,i+1}} = f_a \frac{1 + \beta}{\beta + \frac{W_{gate,i+1}}{W_{gate,i}}} \quad (3.6)$$

As mentioned earlier, each gain stage must also have a 3-dB bandwidth of f_s , so that the DC gain of stage i in the linear amplifier is:

$$G_{DC,i} = \frac{f_a}{f_s} \frac{1 + \beta}{\beta + \frac{W_{gate,i+1}}{W_{gate,i}}} \quad (3.7)$$

The maximum gain is capped by the intrinsic gain of the devices $g_m r_0$. For the last stage, the capacitance driven is the sampler's input capacitance C_{SA} . Finally the power consumed by each stage is $V_{DD} I_{bias,i}$, where $I_{bias,i} = g_{m,i} V_{ov}$, where V_{ov} is the stage overdrive voltage, which is considered to be the same for every stage. The motivation for the constant overdrive voltage stems from insight gained while executing the optimization framework. Namely, adding V_{ov} as an optimization parameter yielded little performance improvement over holding it constant, while adding a significant time overhead in terms of optimization convergence.

3.2.4 Transimpedance Amplifier

The transimpedance amplifier (TIA) is composed of a gain stage similar to those in the VA, with a feedback resistor chosen to meet the bandwidth requirement per stage (f_s). The open loop gain is calculated similar to the VA stage gain. The feedback resistor is therefore set to:

$$R_{FB} = \frac{G_{DC,TIA}}{2\pi f_s (C_{PD} + C_{in,TIA})} \quad (3.8)$$

where C_{PD} is the photo-detector parasitic capacitance including the interconnect between the photo-detector and the TIA, and $C_{in,TIA}$ is the TIA input capacitance. The two poles resulting from the TIA designed in this fashion are not real, and the

extremely low signal voltages. The sampler analyzed here, and depicted in Fig. 3.3 is known as the StrongArm, but the presented analysis and trends can be generalized to a large family of sampler topologies, such as CML-based samplers or more exotic techniques such as double-tail sampling.

3.2.5.1 StrongArm Operating Principle

Before the sampler starts evaluating, the clock is down, and the nodes P,Q,X and Y are brought up to VDD by the reset transistors driven by clock, ϕ . The evaluation starts when the clock goes up, and is composed of two periods,: the sampling period, where in the nodes P,Q, X and Y discharge through M1, M2, M3, M4 and M7, building a differential voltage on nodes X and Y. The sampling period ends when $V_{X,Y}$ reach $V_{DD} - V_{th,P}$ and the cross coupled inverters composed of M3, M4, M5 and M6 turn on. The regeneration then starts and the differential voltage on nodes X and Y is amplified to logic level by the latch.

3.2.5.2 Sampling Period

The sampling phase can itself be divided into two separate phases. The first, during which only M1 and M2 are on, discharges nodes P and Q until they reach $V_{DD} - V_{th,N}$. The common mode voltage V_{PQ} behaves as $V_{DD} - \frac{I_1 t}{C_{PQ}}$ where $I_1 = g_{m1,2} V_{CM}$ is the current drawn by the common mode and lasts $t_1 = \frac{V_{th,N} C_{PQ}}{I_1}$

The second phase starts when M3 and M5 are also on, therefore discharging nodes X and Y. It ends when $V_{XY} = V_{DD} - V_{th,P}$. The common mode behaves according to

$$V_{XY} = V_{DD} - \frac{I_1}{C_{PQ} + C_{XY}} \left[(t - t_1) + \tau \left(\exp\left(-\frac{t - t_1}{\tau}\right) - 1 \right) \right] \quad (3.10)$$

$$\text{where } \tau = \frac{C_{XY} C_{PQ}}{g_{m,3} (C_{XY} + C_{PQ})} \quad (3.11)$$

There is no closed form solution to determine when nodes XY reach $V_{DD} - V_{th,P}$, but if τ is small compared to $V_{th,P} (C_{PQ} + C_{XY}) / I_1$, which is usually the case, the end time of the second sampling phase may be approximated as

$$t_2 \sim \frac{V_{th,P} (C_{PQ} + C_{XY})}{I_1} + \tau + t_1 \quad (3.12)$$

The differential mode, during the second phase, can be shown [20] to follow the equation:

$$\frac{d\Delta V_{XY}}{dt} = \frac{g_{m3,4}}{C_{XY}} \left(1 - \frac{C_{XY}}{C_{PQ}}\right) \Delta V_{XY} - g_{m3,4} \frac{\Delta It}{C_{PQ} C_{XY}} \quad (3.13)$$

$$\Delta V_{XY}(t) = \int \left[\frac{g_{m3,4}}{C_{XY}} \left(1 - \frac{C_{XY}}{C_{PQ}}\right) \Delta V_{XY} - g_{m3,4} \frac{\Delta It}{C_{PQ} C_{XY}} \right] dt \quad (3.14)$$

$$\Delta V_{XY}(t) = \frac{g_{m,1}}{C_{XY} - C_{PQ}} (t - t_1 + \tau_{\Delta} (1 - \exp(\frac{t - t_1}{\tau_{\Delta}}))) \quad (3.15)$$

$$\tau_{\Delta} = \frac{g_{m,3}}{C_{XY}} \left(1 - \frac{C_{XY}}{C_{PQ}}\right) \quad (3.16)$$

Since C_{XY} is usually greater than C_{PQ} , τ_{Δ} is usually negative, and there is no regeneration gain during the sampling period. The sampling gain can be approximated as

$$G \sim \frac{V_{thresh}}{V_{CM} - V_{thresh}} \frac{C_{PQ} + C_{XY}}{C_{XY} - C_{PQ}} \quad (3.17)$$

3.2.5.3 Regeneration Period

Once the top PMOS transistors turn on, the regeneration period starts. The approximation is made that only the cross-coupled inverter pairs are on, providing positive feedback gain, with a time constant

$$\tau_{reg} = \frac{g_{m,3} + g_{m,5}}{C_{in,D2S} + C_{out,SA}} \quad (3.18)$$

3.2.5.4 StrongArm Model within Framework

The modeled sampling stage is made of M interleaved StrongArm samplers (also referred to as Sense Amplifiers (SA)), that evaluate the bits sequentially. This means each individual StrongArm has a cycle $M \times T_{bit}$ long. Half of this period is dedicated to the resetting of the sampler, while the other half is dedicated to the integration and regeneration of the bit (minus the setup time of the follow-on flip-flop T_{D2S}) so that the actual time the sampler is evaluating is T_{SA} , given in (3.21). The schematic of an individual sampler is depicted in Fig. 3.3 and sample transient waveforms are shown in Fig. 3.4. The blue and red lines show the complementary outputs of the StrongArm sampler (nodes X and Y in Fig. 3.3). The integration period lasts while the input pair discharges nodes P,Q,X and Y until nodes X and Y reach $V_{DD} - V_{th,P}$ which dictate when the cross coupled pair turns on and the regeneration period starts [20]

($V_{th,P}$ is the threshold voltage of the PMOS). Fig. 3.4 shows a StrongArm’s transient characteristics with the three main regimes of operation highlighted. The regeneration gain is generated by a cross coupled pair forming a latch, is exponential with time, and brings the output signal to logic levels.




Figure 3.4: Sampler Timing Evaluation Breakdown

The optimization variables available are the common mode voltages at the input, the gate widths of the input transistors, and the gate widths for the cross coupled pair transistors. These define the length of the integration period (which must stay under T_{bit} in order to avoid intersymbol interference), the integration gain, and the regeneration gain. The size of the tail transistor, M_7 , is not considered to be an optimization parameter and is sized to be at least 2x larger than the input pair, M_1 and M_2 , and therefore not current-limiting the signal path.

The sampler then drives a dynamic to static (D2S) converter stage which we simply characterize as a load capacitance to the sampler, $C_{in,D2S}$ [21]. The D2S requires a fixed amount of time $T_{D2S} \sim \frac{2}{f_T}$ to latch, which is taken out of the total evaluation time. Approximations are nevertheless given here:

$$T_{int} \sim \frac{V_{TH}(2C_{PQ} + C_{XY})}{g_{m,1}(V_{CM} - V_{TH})} \quad (3.19)$$

$$G_{int} \sim \frac{V_{TH}}{V_{CM} - V_{TH}} \frac{C_{PQ} + C_{XY}}{C_{XY} - C_{PQ}} \quad (3.20)$$

$$T_{SA} = M/2 \times T_{bit} - T_{D2S} \quad (3.21)$$

$$G_{SA} \sim G_{int} \exp\left(\frac{T_{SA} - T_{int}}{\tau_{reg}}\right) \quad (3.22)$$

$$\tau_{reg} = \frac{g_{m,3} + g_{m,5}}{C_{in,D2S} + C_{out,SA}}; \quad (3.23)$$

where V_{TH} is the absolute value of the threshold voltages and G_{SA} is the final sampler gain. Finally the input capacitance of the SA seen by the front end is given by $M \times C_{ox,SA}$. The fanout M is detrimental to the gain of the front end, and, as will be shown, can be amortized by using switches that connect only one sampler at a time to the output of the VA. In this case, the input capacitance seen by the front end is approximately $C_{ox,SA}$ neglecting wire capacitance and junction capacitance effects of the sampling switches and the RC time associated with them. This assumption holds true for reasonable number of samplers:

$$M < \frac{f_T}{f_{data}} \frac{C_{ox}}{C_{gd}} \quad (3.24)$$

Indeed the size of the transistor serving as a switch can be made substantially smaller than the input cap of the SA, by a factor $\sim \frac{f_T}{f_{data}}$ to minimize it's effect on the circuit bandwidth, and the only capacitance it presents to the circuit is it's gate-drain capacitance, justifying (3.24).

The energy consumed by the sampler comes from the charging and discharging of all it's capacitances at each cycle, as well as the dynamic power burned by the cross coupled inverter during the latching process:

$$E_{samp} = E_{Cap} + E_{latch} \quad (3.25)$$

$$E_{Cap} = C_{SA} V_{DD}^2 \quad (3.26)$$

$$E_{latch} \sim (g_{m,3} + g_{m,5}) \left(\frac{V_{DD}}{2} - V_{TH}\right) V_{DD} (T_{SA} - T_{int}) \quad (3.27)$$

where C_{SA} comprises all the capacitances that will have to be charged to V_{DD} during the reset period.

3.3 Macro-Parameter Derivations

3.3.1 Sensitivity calculation

The sensitivity calculated for the receiver can be separated into two parts: the sampler swing requirement, and the circuit noise requirement. The final sensitivity is the sum of the two.

3.3.1.1 Swing Based Sensitivity Requirement

The swing requirement represents the signal needed to ensure that the differential voltage at the output of the sampler reaches V_{DD} by the appropriate time, and is calculated from the sampler gain, the TIA gain and the VA gain:

$$I_{req,swing} = 2 \frac{V_{DD}}{R_{tot}G_{SA}} \quad (3.28)$$

The factor of 2 comes from the fact that the signal is only half the actual photon current magnitude for an optical ONE (during a ZERO, the photon current is assumed to be nil). Slight changes must be made if the modulation extinction ratio of the transmitter is finite, but the general framework is the same.

3.3.1.2 Noise Based Sensitivity Requirement

The noise requirement necessitates the calculation of the input referred noise generated by the amplification circuit. These include the feedback resistor thermal noise, the Johnson noise from the TIA's transistors, and the transistor noise from the follow-on transistors as well as the noise from the samplers. The TIA transistor and resistor noises are calculated using Personick's method, with all the Personick integrals set to unity [22], while the follow-on stages are estimated using approximations consistent with literature [23]. The photon shot noise (or PD shot noise) is neglected as it is always much lower than the circuit noise sources for incoherent detection systems (roughly one order of magnitude). Indeed for a BER of 10^{-12} , the limit that would be imposed by photon shot noise is 27 photons per bit during a ONE (also known as the quantum limit), which is a current of 44 nano-Amps at 10 Gbps. Naturally when the other noise sources impose a higher photon current, the photon shot noise's absolute value also goes up, but it will necessarily be smaller than the other noise sources.

$$I_{noise,in,Rfb}^2 = \frac{4k_b\theta}{T_{bit}R_{fb}} \quad (3.29)$$

$$I_{noise,in,TIA}^2 = \frac{16\pi^2 k_b\theta\gamma(C_{PD} + C_{in,TIA})^2}{g_{m,TIA}T_{bit}^3} \quad (3.30)$$

$$I_{noise,i}^2 = \frac{4k_b\theta\gamma}{g_{m,i}[T_{bit}R_{fb}\prod G_{DC,j}]^2} \quad (3.31)$$

$$V_{noise,SA}^2 = \frac{8k_b\theta\gamma}{t_2g_{m1}} + \frac{8k_b\theta\gamma g_{m,3}}{t_{12}g_{m1}^2} + \frac{2k_b\theta}{C_{out,SA}G_{sample}^2} \quad (3.32)$$

Here, k_b is the Boltzmann Constant and θ is 273 Kelvin. The sampler voltage noise is approximated using the methodology presented in [24].

Finally the sensitivity is calculated using a current SNR of 14 in order to achieve a bit error rate of 10^{-12} . Please note that this is for current magnitude SNR and not power SNR.

$$I_{req,noise} = SNR I_{noise,input} \quad (3.33)$$

The total photon current requirement at the input of the photodiode is the sum of the swing current requirement and the noise current requirement:

$$I_{req,input} = I_{req,noise} + I_{req,swing} \quad (3.34)$$

3.3.2 Energy per bit

The total energy per bit that is consumed by the link is the sum of the energy burned in the transmitter and the receiver

$$E_{bit} = E_{RX} + E_{TX} \quad (3.35)$$

$$E_{RX} = T_{bit}V_{DD} \sum I_{bias} + E_{samp} \quad (3.36)$$

$$E_{TX} = T_{bit}V_{TX}(I_{req,noise} + I_{req,swing}) + E_{mod} \quad (3.37)$$

E_{RX} includes the power burned in the amplification stages as well as the energy consumed by the samplers. E_{TX} includes laser energy and modulator energy E_{mod} , where V_{TX} represents the energy cost of transmitted photons that represent a bit

successfully detected at the receiver. It encompasses all the efficiencies, η , encountered from the generation of photons to their absorption into useful photocurrent in the receiver photodiode, such as the laser wall plug efficiency, coupler inefficiencies, waveguide losses, modulator loss, photodiode quantum efficiency, etc.

$$V_{TX} = \frac{h\nu}{q} \frac{1}{\eta} \quad (3.38)$$

$$\eta = \prod \eta_{system} \quad (3.39)$$

Table 3.1: Model inputs and optimization variables

Model inputs	Variable description	65nm heterogeneous integration
f_T	Technology unit current gain frequency	150 GHz
C_{PD}	Photodiode capacitance	20 fF
f_{data}	Datarate	1 Gbps to 50 Gbps
α	Fraction of f_T for self loaded stage GBW	0.29 (standard $g_m R_L$ stages) 0.4 (cascode stages)
β	Fraction of input to output cap of a gain stage	0.67
$C_{out,D2S}$	D2S input capacitance	3 fF
t_{D2S}	D2S latching time requirement	25ps
V_{DD}	Supply voltage	1.6V
V_{TX}	Voltage cost of photons	580 V
$(g_m t_0)_{max}$	Maximum voltage gain per stage	4
V_{ov}	Overdrive voltage	0.3 V
γ	noise factor	2/3
E_{mod}	Modulator energy per bit	0
Optimization variables		Bounds
N	Number of amplification stages	0 to 4
M	Number of samplers (M-DR)	1 to 64, in powers of 2
$W_{gate,TIA}$	Input transistor gate width for the TIA	> 150nm
$W_{gate,1,\dots,N}$	Input transistor gate width for each stage	> 150nm
$W_{gate,in,SA}$	Input transistor gate width for sampler	> 150nm
$W_{gate,CC,SA}$	Transistor gate width for cross coupled pair	> 150nm
V_{CM}	Common mode voltage at SA input	0.8V < ... < 1.4V = $V_{DD}/2$ if N=0

3.3.3 Model inputs and optimization variables

The model described enables us to rapidly predict the performance of a given optical receiver characterized by the number of amplification and sampling stages, the technology available, and the size of the transistors involved. These different parameters can therefore also be optimized in order to reach minimal total link energy. The optimization variables and model parameters are described in Table 3.1, and the optimized links are presented in Figs. 4.1, 4.2 and Table 3.3.

3.3.4 Model purpose and limitations

The goal of the model is to accurately encompass all the most important effects and limits that fundamentally constrain the performance of an optical link. Naturally, no model can include all practical limitations, such as systematic and random transistor mismatches, kickback, jitter, layout imperfections, etc. In particular, transistor offset and mismatch can be of significant importance and its effects have been extensively studied [25]. However, through calibration techniques, which indeed add design complexity, the effects of mismatch can be corrected while still adding a minimal power penalty. Additionally, exotic amplification schemes such as higher order stages, or multiple interleaving schemes are not included. While these considerations are important in practical circuit design, we believe that our modeling approach is readily extendable to include these considerations. The presented model will however allow us to derive some general conclusions about critical link trade-offs. It is also precise-enough to provide optimal transistor sizing and accurate sensitivity predictions leading to functional circuits as those shown in Section IV.

3.4 Sensitivity and energy limits

While the model enables us to choose optimal transistor sizings and achieve optimal system link efficiencies, it does not immediately provide us with a deep understanding of the different limits experienced by such a system. In this section we derive these limits. As shown earlier, it is possible to alleviate the swing requirement by using an appropriate amount of interleaved samplers. In a similar way, if a track and hold method is used as in Section 4.3.5 to negate the effect of interleaving fanout, we can make sure the dominant noise source comes from the very front end. We will therefore focus on the limits imposed by noise in the TIA/VA front end.

3.4.1 Front end noise limit

The noise in the front end is dominated by the first amplification stage, which is the TIA in this case. The two major sources of noise have been given in (3.29)

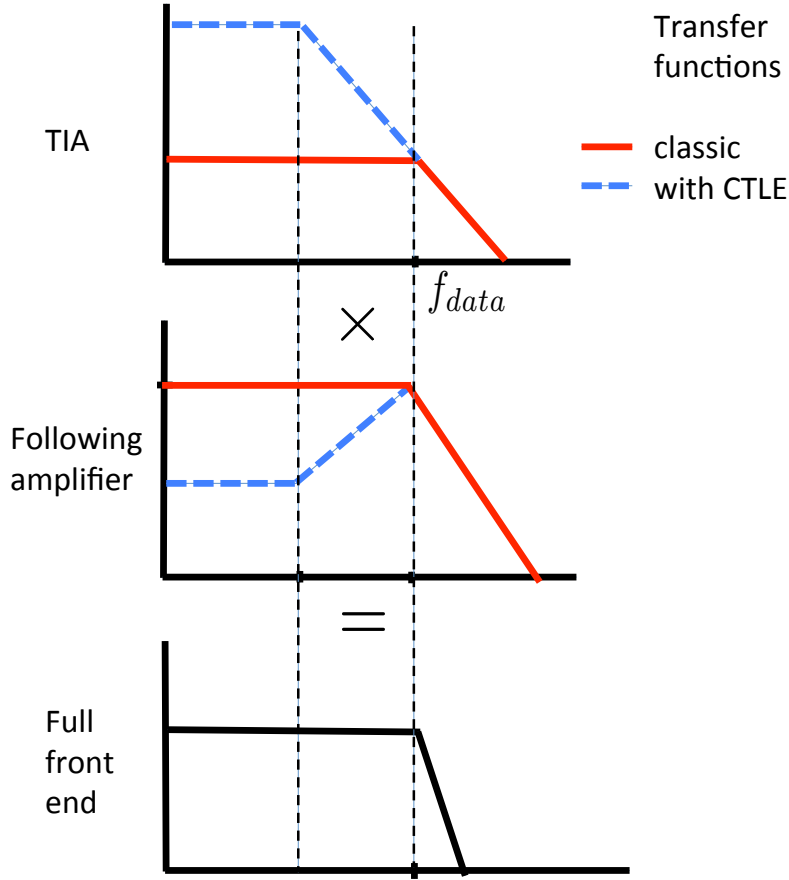


Figure 3.5: Ideal Transfer Function of A System with Equalization

and (3.30), and their input referred noise current is given in (3.40) and (3.41). They depend on the input capacitance $C_{in,TIA}$ of the transimpedance amplifier (though the feedback resistor noise does not explicitly depend on it, the bandwidth requirement sets its value, so there is an implicit dependence), but they have lower limits, which are given by

$$I_{n,R}^2 = (qf_{data})^2 8\pi \frac{C_{PD} + C_{in,TIA}}{6.4aF} \frac{f_{TIA}}{f_{data}} \frac{1}{g_m r_0} \quad (3.40)$$

$$I_{n,amp}^2 = (qf_{data})^2 8\pi \frac{(C_{PD} + C_{in,TIA})^2}{6.4aF \times C_{in,TIA}} \frac{f_{data}}{f_T} \frac{\gamma}{\alpha(1 + \beta)} \quad (3.41)$$

Where f_{TIA} is the bandwidth of the TIA, and $6.4aF = q/V_{th}$ where V_{th} is the thermal noise voltage. When both terms are present, the optimal $C_{in,TIA}$ is somewhere between 0 and C_{PD} . The energy used in the laser to overcome the resistance noise is

constant with datarate if the bandwidth of the TIA scales linearly with the datarate, whereas the energy needed to overcome the transistor noise increases with datarate. Thus at lower datarates, the resistance term is likely to dominate. For the 5Gb/s case when using the parameters of table 3.1, we can calculate that the energy burned in the laser to compensate for this noise is $E_{TX} = V_{TX} SNR I_{R,min}/f_{data} = 152\text{fJ/bit}$, which is close to the energy predicted by the model ($\sim 200\text{fJ/bit}$). Naturally, since the model also optimizes for receiver energy, it is expected that it is not entirely optimized for feedback resistor noise.

Nevertheless, the feedback resistor noise can be overcome to some extent by increasing the value of the feedback resistor, and compensating for the bandwidth degradation by including equalization such as a CTLE stage, as we show in the example circuit of Fig. 4.4. The total front end bandwidth is not enhanced in any way since the TIA and the CTLE stage compensate each other, as illustrated in Fig. 3.5, but this enables the use of a higher resistor value and therefore translates to smaller input referred noise. In (3.40), this is illustrated by the fact that f_{TIA} is reduced, therefore reducing the input referred noise. In this way, it appears that the transistor noise is somewhat more fundamental than the feedback resistor noise.

3.4.2 Limits at high data rates

At high data rates, the input referred noise contributed from the transistors is high enough that laser energy required to overcome it will be the dominant source of power consumption. In this case the optimal receiver will be optimized purely for noise and not its own power consumption, since it will be negligible. We can easily show from (3.41) that the optimal sizing for the input transistors will be $C_{in,TIA} = C_{PD}$. This yields the transistor noise limit, which, expressed in terms of photons per bit for a ONE is:

$$n_{ph,min} = SNR \sqrt{32\pi \frac{C_{PD}}{6.4aF} \frac{f_{data}}{f_T} \frac{\gamma}{\alpha(1+\beta)}} \tag{3.42}$$

Naturally if this value comes close to the the quantum limit of 27 photons per bit for a ONE, the photon shot noise will start to take over.

3.4.3 Limits at low datarates

At lower data rates, the power will not necessarily be dominated by the laser power. If we consider only the noise from the TIA transistors and the power consumption of the TIA and the laser, the energy per bit consumption of the link is:

$$E_{bit} = SNR V_{TX} I_{n,amp} T_{bit} + I_{TIA} V_{DD} T_{bit} \quad (3.43)$$

$$\begin{aligned} &= SNR V_{TX} q \sqrt{8\pi \frac{(C_{PD} + C_{in,TIA})^2 f_{data}}{6.4aF \times C_{in,TIA}} \frac{\gamma}{f_T \alpha(1+\beta)}} \\ &\quad + C_{in,TIA} (1+\beta) 2\pi f_a V_{ov} V_{DD} T_{bit} \end{aligned} \quad (3.44)$$

In this case, there is an optimal size for $C_{in,TIA}$. The lower the data rate, the smaller the input capacitance of the TIA will be in order to minimize power consumption for that stage. To obtain an analytic expression, we assume that $C_{in,TIA} \ll C_{PD}$, which leads to:

$$E_{bit,opt} = 3[\pi SNR V_{TX} C_{PD}]^{2/3} [V_{DD} V_{ov} \gamma k_B \theta]^{1/3} \quad (3.45)$$

Concluding from (3.45), the optimal energy per bit in this case does not depend on the data rate or the speed of the transistors f_T when the link energy is not dominated by the laser power.

3.4.4 E/b power laws

The limit between these two regimes is when we can no longer use the approximation $C_{in,TIA} \ll C_{PD}$ which is only valid when

$$4 \left(\frac{SNR V_{TX}}{2V_{DD} V^*} \right)^{2/3} \left(\frac{qV_{th}\gamma}{C_{PD}} \right)^{1/3} \frac{1}{\alpha(1+\beta)} \ll \frac{f_T}{f_{data}} \quad (3.46)$$

With the photonics platform which will be further elaborated in Chapter 4, this leads to $\frac{f_T}{f_{data}} \sim 15$ which clearly states why 25Gb/s is in the laser limited regime, whereas 5Gb/s is in the full link limited regime. The power laws for optimal Energy/bit of these different regimes is summarized in Table 3.2.

3.5 Observations in Scaling and Technology

With performance limitations arising from both the quality of the CMOS and photonic devices, this section aims to study the effects of an improved design platform with respect to optimized energy per bit. Following the previous analytic analysis, here we utilize the model and optimization procedure described in section 3.2, and apply it to different hypothetical technology platforms. This enables the capture of additional effects such as sampler energy not described in section 3.4. In doing so, we hope to target key bottlenecks in performance and potential for improvements in the next-generation of integration technologies.

Table 3.2: Minimum power laws for E/b limits dependence

Variable	TX dominated regime	TX and RX balanced regime
	Regimes defined by (3.46)	
C_{PD}	1/2	2/3
V_{TX}	1	2/3
f_{data}	1/2	0
f_t	-1/2	0

3.5.1 Improvements in Photonics and Interconnects

Parasitics such as coupler losses and photodiode capacitance dominate the platform described in Table 3.1 and limit the achievable energy efficiency. To study the importance of the photonic performance, we replace the existing metrics for coupler losses, modulator loss, laser efficiency and photodiode capacitance, C_{PD} from 3.5dB/coupler, 5dB/modulator, 10% laser efficiency and 20fF to 1dB/coupler, 3dB/modulator, 30% laser and 3fF, respectively, implying $V_{TX} = 15V$. In addition, modulator efficiency as low as 1 fJ/bit have been demonstrated, justifying their omission from this analysis [27].

The results of the analysis are shown in Fig. 3.6. As compared with the existing heterogeneous integration platform, using better photonics shows more than an order-of-magnitude improvement in link efficiency. Because the price to convert from the photonic to electrical domain, V_{TX} , is so cheap now, the optimized links at the various data rates are more receiver-performance limited, as expected intuitively.

3.5.2 Improvements in Photonics+CMOS

To push the boundary of integration technologies altogether, we now turn to the case where the photonics and CMOS are both pushed to their bounds. In particular, we utilize the same best photonic specifications from before, but, now, scale the technology node to reflect a theoretical f_T of 1THz. The results of the study are shown in Fig. 3.6. For lower data rates, the performance improvement from scaling f_T from 150GHz to 1THz is observable but not drastic and stems mostly from the lower energy consumption of the samplers themselves and not the front end amplifier or the laser, as expected from the limits of section 3.4. For the 25G DDR case, however, the improvement is almost an order of magnitude since the faster amplifiers can provide gain at these speeds. Notice that the last column in this bar plot shows a *100G DDR* receiver, with a theoretical best end-to-end link efficiency of 20fJ/bit.

While the previous sections show the performance for given technologies, we can

reverse the exercise to deduce the necessary technology properties for a given link efficiency. To achieve sub 1fJ/bit efficiency at 5Gbps and $f_t=1000$ THz, this would require $C_{PD}=200$ aF, $V_{DD}=0.5$ V, $V_{ov}=0.1$ V and $V_{TX}=10$ V. These small photodiode capacitances would require such a small device that some sort of absorption enhancement would be necessary, such as a cavity or a metaloptic focusing scheme [28]. At this point the link energy itself is so small that effort must be redirected to the energy overhead of peripheral blocks such as clock networks and bias generators.

The performance results for these higher data rates have another interesting trend – as the CMOS platform performance improves, the energy consumption of the receiver is mostly limited by the sampler itself. Because we have assumed a StrongArm topology for the sampler for all data rates of operation, the minimum achievable E/b of this sampler is far greater than the rest of the link put together. This yields the conclusion that within the confines of a better platform where photon efficiency is so high, using a simple gain stage such as an inverter as the sampler is more optimal than having a StrongArm or CML latch.

Table 3.3: Performance Comparison of Model-Predicted and Schematic-Simulated Optical Receivers

Specification	5Gbps Standard RX		5Gbps CTLE RX		25Gbps DDR		25Gbps QDR		25Gbps Switching QDR	
	Modeled	Designed	Modeled	Designed	Modeled	Designed	Modeled	Designed	Modeled	Designed
Total Front-end Gain (Ω)	15500	20200	15500	15500	636	770	1055	760	1750	1840
Total Front-end BW (GHz)	3.5	3.4	3.5	3.4	17.5	18.8	17.5	20.5	17.5	17.8
Total AFE E/b (fJ/bit)	79	52	79	91	340	300	260	395	134	135
Total Sampler E/b (fJ/bit)	32	35	32	37	110	97	93	153	88	91
Total RX Sensitivity (μA)	3.37	3.8	2.95	3.2	356	288.2	33.9	86.2	28	77.4
SA Input Common Mode (mV)	500	500	500	530	850	840	850	810	850	950
SA Minimum Swing (mV)	0.2	6	0.2	5	210	165	1.6	4	1.5	11
RX Swing Sensitivity (μA)	0.01	0.8	0.01	2.1	326	214	1.7	5	0.9	6
14 σ RX Noise Sensitivity (μA)	3.4	2.9	2.9	1.1	30.8	74.3	32.2	81.2	26.6	71.4
Link RX E/b (fJ/bit)	110	88	110	128	440	400	225	550	221	227
Link Laser E/b (fJ/bit)*	392	—	337	—	16000	—	840	—	655	—




Figure 3.6: Technology Dependent Performance Prediction

3.6 Link Framework and First-Principles Summary

This work introduced an fundamentals-influenced optimization approach for true end-to-end optical links incorporating a TIA, a linear amplifier chain, and follow-on StrongArm latches. In portraying this "digital IN" to "digital OUT" end-to-end link structure, distinct regimes of operation were evident. For low data rate regimes, overcoming front-end noise posed as the dominant contributor to overall link budget. However, for high data rate regimes, the StrongArm voltage evaluation requirement quickly dominated and yielded the swing-limited regime. Circuit techniques such as sampler time interleaving were used to greatly reduce this swing requirement by exponentially amplifying the sampler gain. In doing so, the swing requirement no longer became dominant at higher data rates. Rather, sampler noise instead quickly took its place in this regime of operation. Further circuit techniques such as placing interleaving switches between the AFE and the StrongArm latches can be used to further reduce this sampler noise contribution. This is done by effectively reducing the total sampler load capacitance, thereby allowing for higher front-end gain.

This work continues by using this fundamental model in order to extrapolate the performance of next-generation, best-case technologies that are optimized for photonic as well as CMOS performances. As expected, the best-case energy per bit for these optimized technologies scales to show more than an order-of-magnitude improvement in performance. Moreover, new limitations arise that are a result of the "weakest link" technology. For example, using the best-case photonics with standard CMOS platform reveals the E/b is receiver-performance limited. These trends and next-generation platform studies showcase the importance of various parameters and their ultimate relationship to end-to-end link performance.

3.7 Conclusion

In this chapter, we introduced the link design framework, which will act as one of the pinnacles of the remainder of this dissertation. The framework will be used in not only engineering physical designs (i.e. Project: EPHI), but also bring insight into means of improving link efficiency. We will see how the link framework

motivates the necessity of better transmitters, showcased by the nanoLED transmitter analysis in Chapter 4. Moreover, on the receiver side, we will see how this framework motivates the necessity of PAM4 links and various new insights that emerge from that playground.

Chapter 4

Link Framework Application for NRZ Front-End Design

4.1 Introduction

To further realize optimal topologies under unique system parameters as well as to provide validation for the link design framework introduced in the last chapter, the main objective of this chapter will be to propose and design circuits and systems under realistic PDK and technology specs. In addition, based on these unique technology specifications and sizings, the end-to-end link performance will also be calculated and analyzed. In Section , the link design framework will be applied to a 65nm, heterogeneously integrated CMOS plus photonics platform. The physical parameters and PDK landscape will be studied in Chapter 5 but, for the purposes of this analysis, these parameters will be abstracted to simple numerical quantities only. Upon introducing this platform and applying the framework, the resulting schematic results will be showcased in Section 4.3. More specifically, data points in both the noise and swing dominated regimes will be studied and their resulting performance will be derived and analyzed. Finally, in the conclusion of this chapter, we will see how the analysis presented in this chapter sets the foundation for proceeding to not only tape out physical designs, but also bring awareness to potential performance bottlenecks. An attractive solution to one of the key performance bottlenecks will be studied in Chapter 6 with the introduction of the single comparator-based PAM4 link architecture.

4.2 Model Application for 65nm Heterogenously Integrated Photonic Platform

4.2.1 Technology overview

The circuit optimization has been applied for use in a heterogeneous integration platform which combines a 65nm CMOS technology node with a custom-SOI photonic node [18]. In this process, separate 300mm photonic and CMOS wafers are face-to-face oxide bonded in the CNSE 300mm foundry.

To reduce the capacitance between the CMOS and photonic wafers, the silicon substrate is first removed from the photonic wafer and through-oxide-VIAs (TOVs) with a lumped capacitance of 3fF per TOV are used to interconnect the two wafers. The major technology parameters are listed in Table 3.1, in Chapter 3.

Using Figure 3.2 as reference, light from the laser source experiences multiple sources of loss before contacting the photodiode on the receiver side. Firstly, the laser source itself is assumed to have a wall-plug efficiency of 10%. The three vertical grating couplers, which measured 3.5dB/coupler of loss, are also in the critical path of the signal. The germanium photodiode has a measured responsivity of 0.8 A/W at 1510nm [26]. In addition, the modulator insertion loss is 5dB. In this study, we assume no waveguide loss. However, that can be easily implemented. The above path losses translate to an overall photon energy cost, V_{TX} , of 580 V. The modulator energy in this platform is 20fJ/bit and will therefore be neglected for the purposes of this analysis.

4.2.2 Single slicer case (M=1)

The results of the optimization for the optimal performance of the link are plotted in Figure 4.1. The laser energy to accommodate Noise and Swing are respectively the quantities described in (3.38).

Two clear regimes are visible: the "Noise limited regime" at low datarates, where the sensitivity of the receiver is constrained by the noise, and the "Swing limited regime" at high datarates, where the sensitivity is dominated by the output swing requirement ($V_{out} = V_{DD}$). The regeneration gain of the sampler is exponential with time, so it is natural that at higher datarates it drops significantly. While the VA can compensate for this drop in gain by increasing its number of stages (and this happens at ~ 7.5 GHz for M=1), there is a limit to the amount of aggregate gain achievable by chaining amplifiers due to the bandwidth requirement, as described in (3.2).

The justification for adding multiple slicers is now obvious. This relaxes the condition on the regeneration time being less than the bit duration, and can thus push the swing limited regime to much higher datarates.

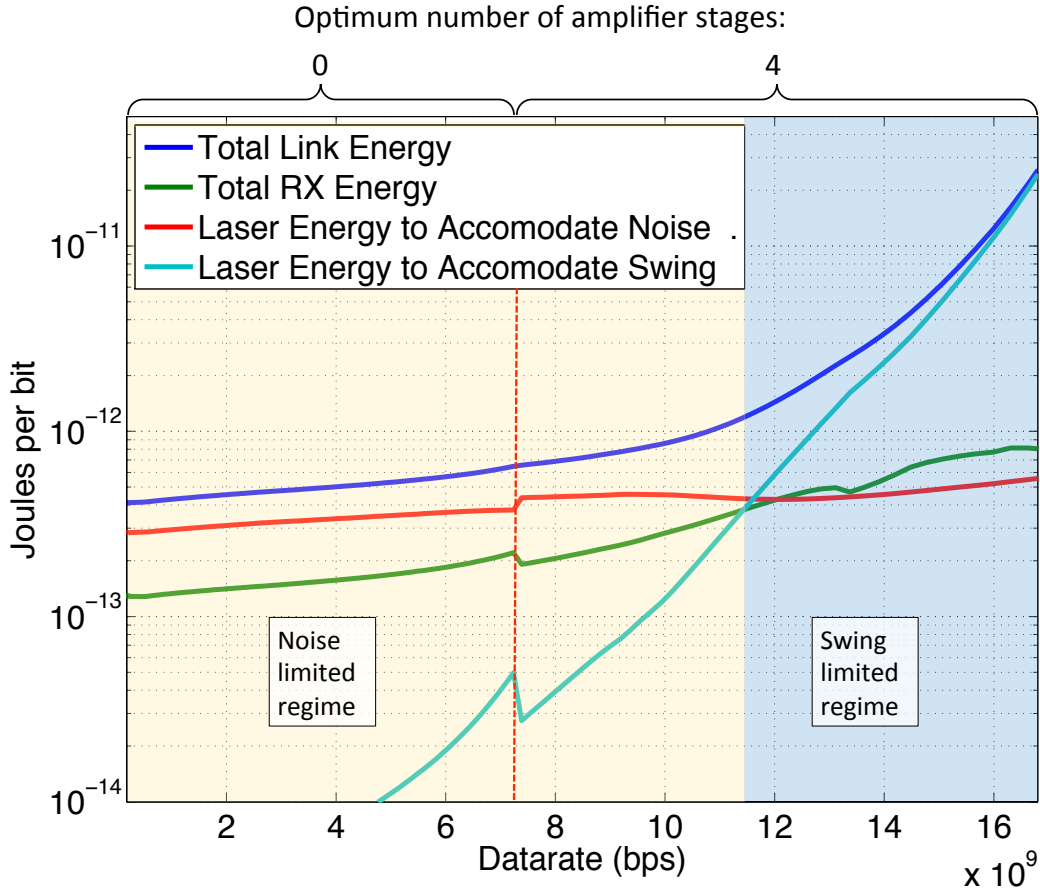


Figure 4.1: Optimal Energy Per Bit Versus Data rate For Optimal Topologies With Parameters From Table 3.1. Only One Slicer is Allowed in This Case

4.2.3 Multiple slicer case ($M \geq 1$)

The results of the optimization when the number of samplers is not constrained to 1 is plotted on Figure 4.2. We can see that there is no longer a "Swing limited regime", since the optimal topologies have several samplers in order to benefit from much higher regeneration time and gain. While the energy per bit is greatly reduced at higher data rates, eventually the sampler noise starts to dominate. This comes about because as the data rate goes up, the bandwidth requirement on the VA reduces the possible achievable gains. Additionally, adding several samplers increases the fan-out of the VA by a factor M , further reducing the gain. Eventually the gain of the amplifier stages drops below 1, so that the input referred noise coming from the samplers becomes greater than that coming from the front end. We therefore observe a front end noise limited regime and a sampler noise limited regime, which is different from the sampler swing limited regime discussed in the single-slicer case.




Figure 4.2: Optimal Energy Per Bit Versus Data rate For Optimal Topologies With Parameters from Table 3.1 with the Possibility of Multiple Slicers

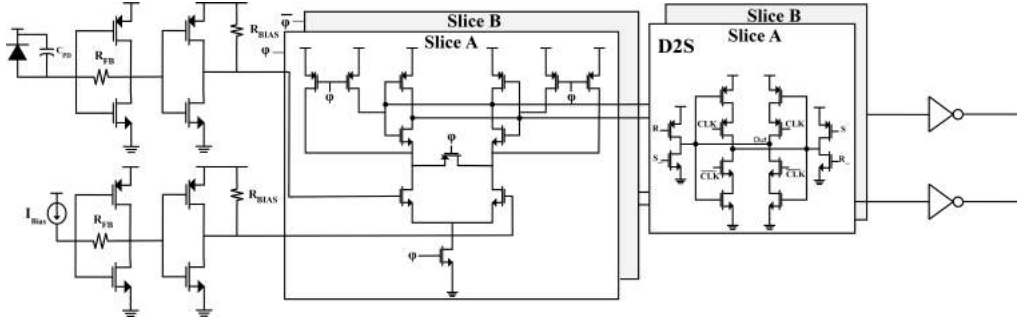


Figure 4.3: 5Gbps Receiver Topology

4.3 Schematic designs of model results

4.3.1 5Gbps optical receiver

To highlight the performance in the noise-limited regime, we introduce the schematic design of an optimized receiver topology operating at 5Gbps, with no active equalization, running off of a 1.2V supply. Figure 4.3 shows the overall topology of the front-end pre-amplifier and slicers. While the number of slicers and samplers does not match the optimal values of Figure 4.2, these values were chosen because they yielded performance within a few percent of the optimum, and were easier to implement. Nonetheless the transistors sizings were still produced by the algorithm.

4.3.1.1 Design Overview

The photodiode, with a total capacitance C_{PD} of 20fF, inputs into a TIA amplifier with a feedback resistance, R_{FB} , valued at $2.3k\Omega$. The output of this stage enters a single pre-amplifier gain stage with a gain of 2 before entering the optimized, dual-data-rate (DDR) triggered StrongArm Sense Amplifiers and follow-on dynamic-to-static converters. The sense amplifiers and dynamic-to-static converters are triggered on clock and clockB (Φ and Φ_B), which each operate at half the data rate or 2.5GHz. The sampler transistor sizes as well as the front-end sizings are optimized using the algorithm. Additionally, the biasing at each stage is also dictated by the algorithm. More specifically, the common-mode voltage at the input of the samplers was selected to be 850mV while the constrained common-mode voltages at the TIA's input and output were set at $V_{DD}/2$ or 600mV. The output of the samplers, which are effectively a 1-to-2 deserialized version of the input data sequence, was verified in simulation.

4.3.1.2 Simulation Results

The above design has been implemented at the simulation level and its performance was verified with respect to the values predicted from the model. Table 3.3

summarizes specifications for the model and simulated results. The optimized circuit had an overall front end gain of $5.1k\Omega$ and from the StrongArm sampler’s standpoint, the minimum required swing at the input (neglecting noise) to resolve successfully at 5Gbps, or 200ns of evaluation time per sampler, was measured to be 6mV. This translates to a $0.8\mu\text{A}$ receiver sensitivity due to the swing requirement of the sampler. From a noise perspective, the total input-referred noise contribution from the front-end is $0.21\mu\text{A}$ (1σ). Thus, the total simulated input sensitivity $i_{swing} + 14 \times i_{noise}$, or $3.8\mu\text{A}$. The total energy per bit for the full RX block is 280 fJ/bit, with the front-end consuming 115 fJ/bit and the samplers plus D2S consuming 165 fJ/bit total. The front-end E/b in this case takes into account the the dummy front-end as well. From an overall link perspective, the energy in the laser and TX macro is 392fJ/bit.

4.3.2 Active-CTLE enhanced 5Gbps optical receiver

We now use the algorithm to design a continuous-time linear equalizer (CTLE)-based optical receiver front-end. The purpose of this design is to show that it is possible to reduce the noise contribution of the feedback resistor in the front end, as will be discussed in section 3.4. We nevertheless present the circuit results here for consistency. The full schematic is shown in Figure 4.4.

4.3.2.1 Design Overview

The CTLE-based pre-amplifier allows for the preceding TIA stage’s R_{FB} to increase drastically, from $2.3k\Omega$ to $70k\Omega$. This enhances noise performance while keeping the overall gain-bandwidth the same. The new low-pass pole of the TIA front-end is then compensated with the peaking of the CTLE block, which adds a *zero* in the transfer function from the degenerated R_S and C_S (see Figure 3.5). The *zero* location is chosen to match the dominant pole-location of the TIA, thereby enhancing the overall bandwidth to the target specification for operation at 5Gbps. The two poles of the CTLE are set to the same frequency in order to maximize the effective gain-bandwidth of the stage [9].

4.3.2.2 Results

The results are summarized in Table 3.3. The overall receiver gain and bandwidth of the CTLE are approximately that of the standard RX topology, at 5460 k Ω and 5.3GHz, respectively. The CTLE-based front-end consumes 250 fJ/bit with the samplers consuming 141 fJ/bit. This yields an overall RX E/b of 391 fJ/bit. The main advantage in using a CTLE-based scheme comes from the input referred noise sensitivity, as will be further elaborated in Section V. Here, we observe $0.2\mu\text{A}$ input sensitivity whereas the standard RX topology had almost double that. In the

CTLE topology, the feedback resistor contributes only 15% of the total front-end noise, whereas for the standard RX topology, that contribution is almost 50%.




Figure 4.4: 5Gbps Model-Predicted Receiver Topology with Active-CTLE

4.3.3 DDR 25Gbps optical receiver

4.3.3.1 Design Overview

To better characterize the universality of the model, we now present an optimized optical receiver design operating in the sampler swing-limited regime. We choose to operate at V_{DD} of 1.6V in order to allow for enough voltage headroom to utilize cascode-amplifiers as the basis design for the VA stages, which have an α of 0.4 as opposed to the standard amplifiers which have an α of 0.29. We retain the StrongArm topology for the samplers and also retain the topology of the D2S converters. In this design, we choose to operate the system as DDR to show the importance of relaxed timing margin on the sampler's evaluation period.

Under these constraints, the model-predicted topology is shown in Figure 4.5. All front-end FETs, resistances, and sampler FETs, are all sized based on the constraints presented by the algorithm. In the DDR case, $M=2$

4.3.4 QDR 25Gbps optical receiver

4.3.4.1 Design Overview

In the subsequent analysis, we retain the same technology parameters as in the previous section. However, now, we present a quadrature-data-rate (QDR), $M=4$ from Figure 4.5, operation of the receiver, wherein four samplers are utilized to parse the amplified photodiode signal. Once again, the design of the front-end as well as samplers is fully optimized with our tool taking into account the added capacitive load factor on the final stage of the VAs. In using four phases, we alleviate the timing evaluation requirements of the samplers by doubling the allocated time for sampling and reset phases, while adding clocking overhead in the form of quadrature phase generation. In the context of links, this drastically improves efficiency and extends the crossover point of the noise-limited and sampler-limited regimes to past 25Gbps, as seen in Figure 4.2. Although we acknowledge the added overhead of generating quadrature phases versus dual phases, the purpose of this analysis is to highlight the importance of easing timing requirements for the samplers to improve the overall performance. Indeed, an order of magnitude reduction in link power efficiency was observed (mostly through the increase in SA gain), not taking into account the cost overhead of clock phase generation. According to Figure 4.2, the optimal number of samplers is actually 8, yet the energy consumption benefit of going from 4 to 8 samplers is not enough to justify the additional design complexity.

4.3.4.2 Results

The QDR receiver performed on par with the DDR in power, gain, and bandwidth metrics. However, from a swing sensitivity standpoint, the QDR receiver performed an order of magnitude better. The simulations yielded a swing sensitivity of under $5\mu\text{A}$, with a front-end gain of 760Ω and 20.5GHz net bandwidth. The four samplers and D2S's were consuming 153fJ/bit while the front-end was consuming 395fJ/bit for a total 550fJ/bit being burned on the receiver end. The input referred noise sensitivity for the receiver was $1.8\mu\text{A}$, now mostly dominated by the sampler noise.

Because of this ultra-low sensitivity, even though the RX total power stayed approximately the same for the DDR and QDR cases, the required laser power was substantially reduced, as shown in Table 3.3.

Discrepancy in the noise sensitivity values between modeled and designed may be attributed to not only the sampler noise approximation error (which may be as high as 50%) but also to the first-order noise calculation methodology being used [24]. The swing sensitivity discrepancy stems from the following: (1) the simulated swing sensitivity looks at settling at the output of the D2S, an effect not captured by the model; (2) lower input sensitivities rely heavily on capturing the effects of regeneration properly. The error on the regeneration side shows up as an exponential variation

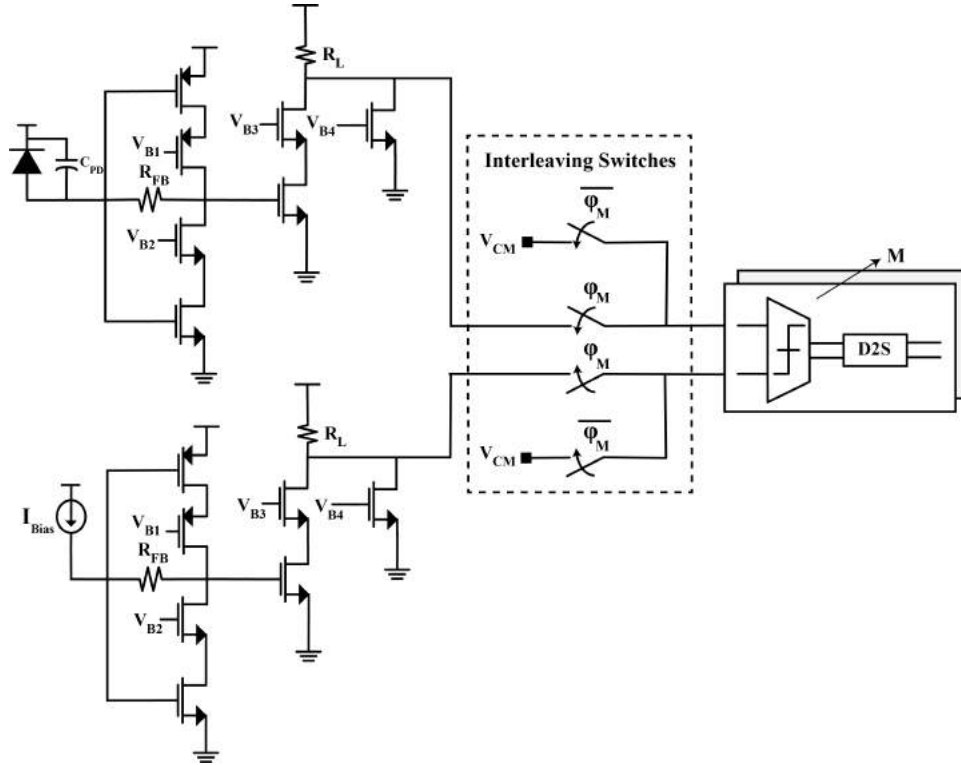


Figure 4.6: Switching Time-Interleaved 25Gbps QDR Receiver

in the sampler gain. For the purposes of this study, however, the 2x variability in sensitivity is considered acceptable.

4.3.5 Switched QDR 25Gbps optical receiver

4.3.5.1 Design Overview

To alleviate the sampling noise contribution of the StrongArm sense amplifiers, a time-interleaved switching topology was implemented, reducing the load on the VA and allowing it to provide more gain for a given bandwidth constraint. The schematic is shown in Figure 4.6. By placing a track and hold circuit prior to the sampler array, not only does the sense amplifier input load capacitance diminish, but the potential effects of kickback from other sampler clocks is also theoretically reduced. The receiver topology and design process is similar to the QDR receiver in Figure 4.5. All transistor sizings are optimized by the tool with the biggest difference being in how the $C_{ox,SA}$ capacitance scales. $C_{ox,SA}$ now goes up linearly with sampler input FET size and is completely independent of the slicing count, M , as detailed in Section 2. For the purposes of this study, the non-idealities of these sampling switches (i.e. finite junction capacitance) were not taken into account within this study. However,

the simulation results reflect performance with these non-idealities in place, and we see no significant difference between the predicted and simulated specifications. This is because the sampler count, M , is kept to a reasonable value according to (3.24). Additionally, charge injection, which was not modeled analytically, is a common mode issue and, therefore, does not affect sensitivity drastically.

4.3.5.2 Results

The results in 3.3 for the 25Gbps Switching QDR receiver show similar performance to the non-switching. However, the total noise sensitivity is reduced by 10% on account of the reduction in sampler noise contribution, while the noise from the front-end stays relatively constant. The sensitivity required to overcome the sampler swing is also relatively constant, with small adjustments made to the input sampler FETs on account of the switching.

4.4 Conclusion

In this chapter, we showcased the capabilities of the link design framework and presented real-world tape-out-ready designs that focus on the optimal energy per bit given a set of input technology parameters. More specifically, the 65nm heterogeneous integration platform was introduced and designs in both the noise and swing-limited regimes were outputted and characterized. Furthermore, transmit-side versatility of the framework was validated with the design and optimization of the nanoLED-based optical transceiver. To conclude, this chapter lays the second layer to the framework foundation. More specifically, by generating physical designs and validating them at the schematic-level, we are able to understand trade-offs extending from the technology platform up to the link layer.

Moving on, we will delve into two developing thrusts with this analysis. First, in the spirit of *true* end-to-end link link design, physical silicon must be tried and tested. As such, Chapter 5 will focus on the 5Gbps data point and present post-tape out silicon results of the end-to-end link, which went on to be published in ESS-CIRC 2013. Second, in further analyzing performance bottlenecks under particular system constraints, Chapter 6 will introduce a new, single-comparator based PAM4 architecture which not only promises better energy efficiency, but also embraces the rapid-design approach to taping out complicated silicon photonic links. In doing so, architectural trade-offs which extend beyond their hierarchy will be studied in depth.

Chapter 5

3D Integrated Silicon Photonic Interconnects

5.1 Introduction

In order to realize the *true* end-to-end link design approach highlighted in the previous chapter, this chapter will take to silicon a 5Gbps end-to-end optical transceiver in a 65nm heterogeneously integrated platform. To enable full optical links for interconnection networks, high speed and low power optical transmitters as well as high bandwidth and high sensitivity optical receivers are required. These necessitate the need for close integration in order to achieve small parasitic capacitance between electronics and photonic devices. Furthermore, a two-wafer solution is desirable to separately optimize the performance of the photonic components and the CMOS circuits. This chapter will demonstrate for the first time an optical chip-to-chip link built in a heterogeneous, 3D integration platform using thru-oxide via (TOV) technology. The TOV technology overcomes the challenges of close integration of electronic and photonic components, by simultaneously enabling separate wafer optimization of electronic and photonic components while providing a low-capacitance, high-density connection between the photonic and electronic wafers.

This chapter will detail the analysis and design of a full, end-to-end 5Gbps NRZ link in the first-ever through-oxide via (TOV)-based integration technique. The chapter will begin with a description of the integration platform. Next, we will showcase the top-level design for the Electronic-Photonic Integration (EPI) system, which was taped out and published in ESSCIRC 2015. Each sub-block will then be analyzed and characterized before concluding with final remarks.

5.2 Overview

Optical interconnects targetted towards high speed and low power networks require specialized integration technologies to optimize not only photonic and CMOS technologies independently, but allow for tight integration between the two to facilitate optimized co-performance. On the photonics side, the fabrication platform may target low polysilicon loss for best waveguide performance, or the inclusion of germanium for better photodiode performance, as examples. On the CMOS side, ensuring fast transistors (by having high f_T) and having low wire, gate, or fringing capacitances improves overall circuit performance. From the standpoint of co-integration, very low parasitic capacitance between the CMOS and photonic connections is required.

A multitude of integration platforms exist that realize this connection between CMOS and photonics. The two main categories are monolithic integration, where the photonic components are placed alongside CMOS devices on the same die, and hybrid integration which relies on specialized connectivity between the independent CMOS and photonic reticles. Figure 5.1 highlights the key pros and cons of each.

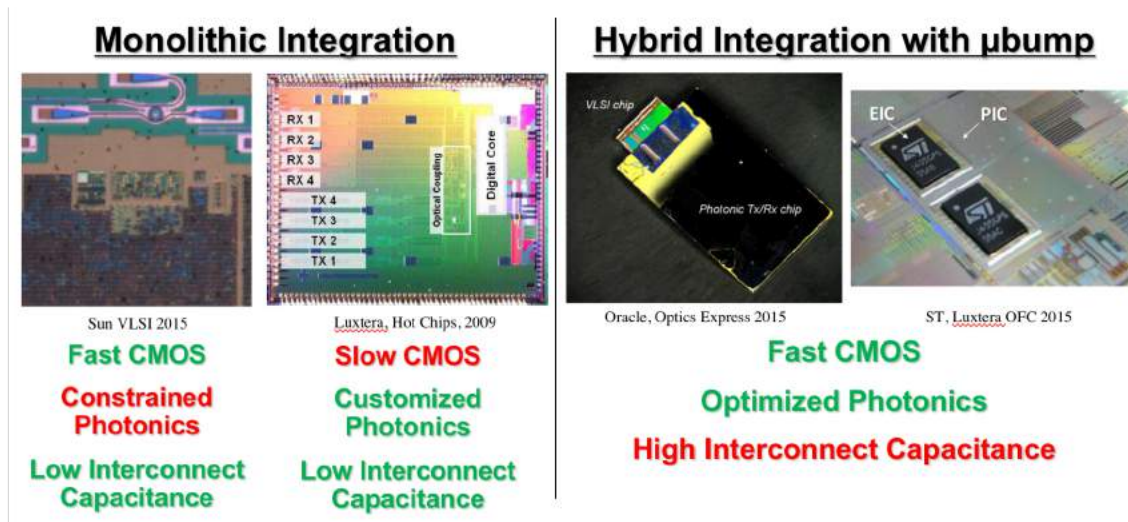


Figure 5.1: A comparison of the integration technology categories shows pros and cons for both.

Monolithic technologies offer very low interconnect capacitances as they rely on the traditional metal stack for connectivity. This allows close proximity for the photonic components and transistors, say. However, these technologies generally target either for best CMOS performance or for best photonics performance. Unfortunately, their respective attributes are “zero-sum” so maximizing performance on one generally hinders performance on the other. This sets the motivation for hybrid integration technologies, where the CMOS and photonics may be optimized independently. However, because they are on separate reticles, a tight integration platform is necessary.

This chapter demonstrates for the first time an optical chip-to-chip link built in a heterogeneous, 3D integration platform using thru-oxide via (TOV) technology. The TOV technology overcomes the challenges of close integration of electronic and photonic components, by simultaneously enabling separate wafer optimization of electronic and photonic components while providing a low-capacitance, high-density connection between the photonic and electronic wafers.

5.3 3D Integration of CMOS and Photonics

Traditional heterogeneous platforms capitalize on the ability to individually optimize the photonic and electronic macros, an element missing in other forms of integration. An example WDM system utilizing a heterogeneous platform is shown in Figure 5.2.




Figure 5.2: An example diagram of the full WDM link using a heterogeneous platform shows independent photonic and CMOS wafers on both the Tx and Rx side.

However, the large interface capacitance associated with thru-silicon via (TSV) and -bump technologies limits the overall system performance as well as energy-efficiency.

As illustrated in Figure 5.3, in this process, 300mm photonic and electronic wafers are manufactured separately in CNSE 300mm foundry and then bonded face-to-face using oxide bonding.

The silicon substrate is then removed on the photonic SOI wafer and TOVs are punched through at 4m pitch to connect the top metal layer (M2) of the photonic wafer to the top-layer metal on the 65nm bulk CMOS wafer. Back metal is patterned on the flipped photonic wafer to create bonding pads and connections to the TOVs.

For packaging, wire-bonded back metal pads are deposited on top of the selected TOVs. The connection from the CMOS wafer to the photonic device is achieved through the TOVs passivated on top with an oxide layer, which minimizes the parasitic capacitance. Our measurements estimate the TOV capacitance to be 3fF, which enables low-power and high-sensitivity electronic-photonic systems for a variety of applications. This represents an order of magnitude reduction in parasitic capacitance,




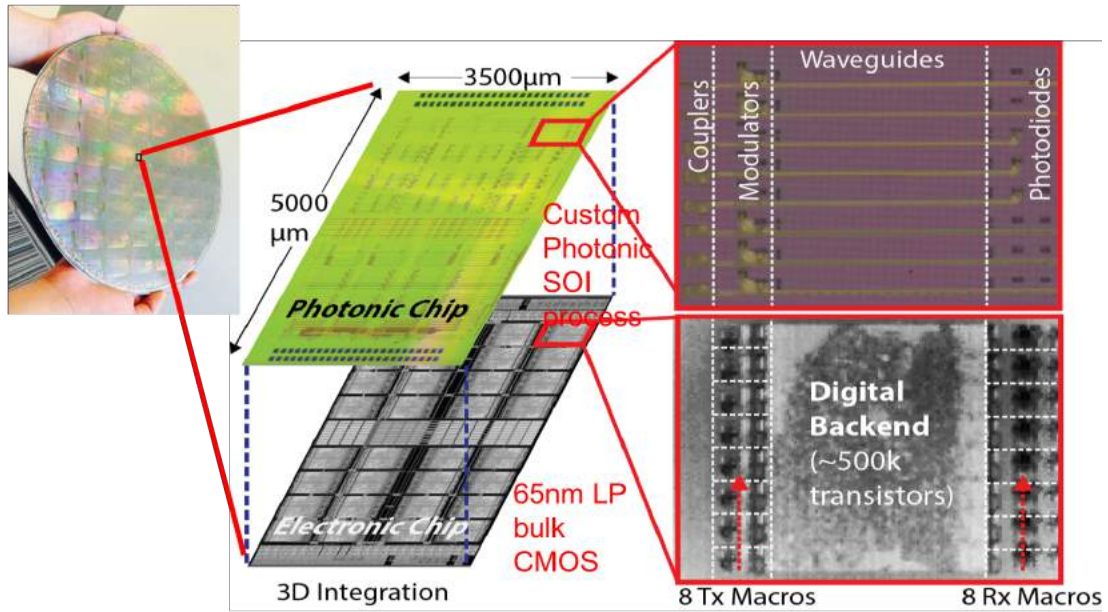
Figure 5.3: A cross section of the 3D TOV heterogeneous integration process shows the photonic and CMOS wafers and associated connections.

and two-orders of magnitude higher density compared to previously demonstrated -bump flip-chip electronic-photonic integration [3].

5.4 Chip Architecture and Optical Link System

The optical chip-to-chip link is a part of the wafer-scale heterogeneously integrated technology-development and demonstration platform with low-energy optical transmitters, receivers, and comprehensive backends for performance characterization, as shown in Figure 5.4. Apart from containing vertical junction depletion mode microdisk modulators [4] within the photonics die, hetero-epitaxially grown Germanium photodiodes and body crystalline silicon low-loss waveguides are also used to enable electro-optic transceiver functionality.

The 16M transistor electronic chip contains 32 Multicell sub-blocks that enable a full self-test of modulators and receivers within the link. Each Multicell is composed of eight RX as well as eight TX macros, enabling in-situ testing of a wide variety of photonic devices. The Multicell also contains an expansive digital backend infrastructure to enable full, self-contained characterization of each of the eight TX and RX sites. Characterization is accomplished through on-chip, self-seeding PRBS generators and counters. The 231-1 length PRBS data sequence gets fed into one of the TX macro sites, which serializes the data and drives the resonant modulator device imprinting the data sequence on the light in photonic waveguide. On the RX



4

Figure 5.4: A full system of the heterogeneously integrated EPHI system.

side, this modulated light is fed into one of the eight RX macros. The output of this RX macro is an eight-channel bus, marking the deserialized input optical data. These eight channels proceed on into the backends bit-error-rate (BER) checkers, which count the total number of errors between the received data from the RX macro and the ideal sequence provided by the seeded PRBS generator.

5.4.1 Transmitter Design

The TX macro, shown in Figure 5.5, consists of a tunable vertical junction depletion-mode ring resonator similar to [2, 4] driven by an 8 to-1 serializer and driver head with on-chip PRBS input. The applied reverse-bias voltage to the junction via the driver head depletes free carriers and perturbs the refractive index of silicon, which in turn shifts the resonance wavelength (or frequency) of the optical modulator. The cathode of the modulator diode is connected to 1.2V while the anode is modulated from 0 to 1.2V. The modulator p-n junction is reverse-biased during modulation.

Given that the leakage current is small, the energy is consumed only when the transitions charge the reverse-biased junction capacitor. With a total modulator driver capacitive load of 12.4fF (modulator diode and TOV), at 6Gb/s the whole

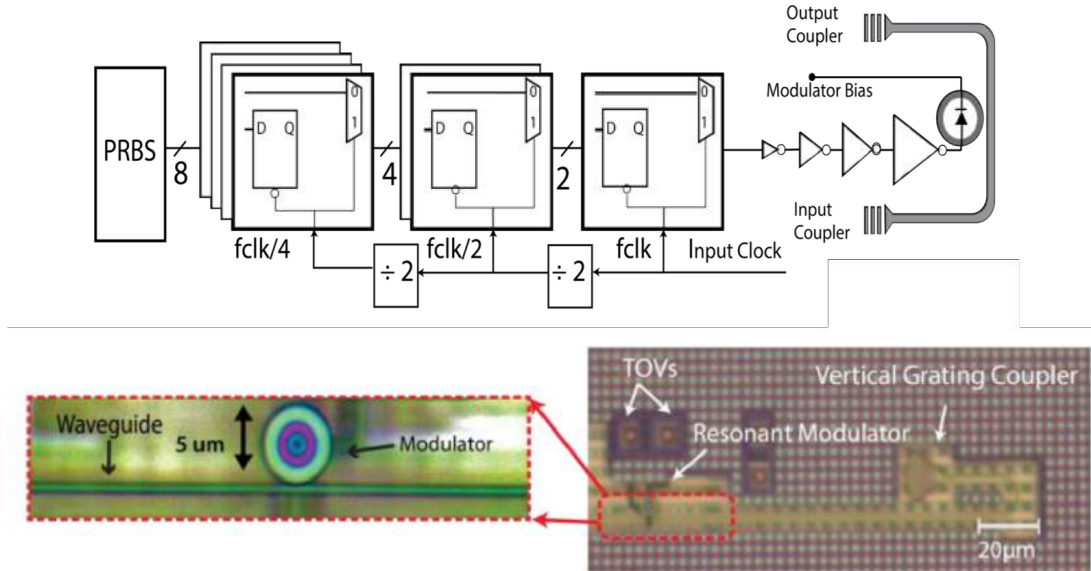


Figure 5.5: The Tx Macro consists of high speed serializers and drivers shift the ring resonance.

macro consumes 100fJ/b (5fJ/b modulator, 15fJ/b driver, and 80fJ/b serializer).

Heterogeneous integration allows us to use the state-of-the-art ring resonant modulators with a large electro-optic response of 150pm/V (20GHz/V), which enables low power modulation using small voltage swing (1.2V) while still maintaining sufficient extinction ratio (Fig. 4(a)). Measured from the modulator transmission spectra at 0V and -1.2V dc biases, the device should ideally achieve 6.2dB extinction ratio (ER) and 1.8dB insertion losses (IL). The modulator can also be modulated between a slightly forward-biased regime and depletion regime by lowering the bias voltage of the anode (i.e. -0.2 to 1.0V). This will further improve extinction ratio of the modulator.

A tunable CW laser source was coupled to an on-chip silicon waveguide through a vertical grating coupler. The laser frequency was aligned adjacent to the resonance frequency of the modulator ring (1520nm, see the left graph of Figure 5.6).

The TX circuits drive the 31-bit PRBS sequence into the modulator, achieving the non-return to zero on-off keying (NRZ-OOK) modulation eye at 6Gb/s, as shown in the right side of Figure 5.6, with 6dB extinction ratio and 2dB of insertion loss, which agrees well with the transmission spectrum. The fast rise-time indicates the potential for faster operation, but the results are currently limited by the global high-speed clock distribution network that spans the whole chiplet and supplies the clock to all the Multicell macros.

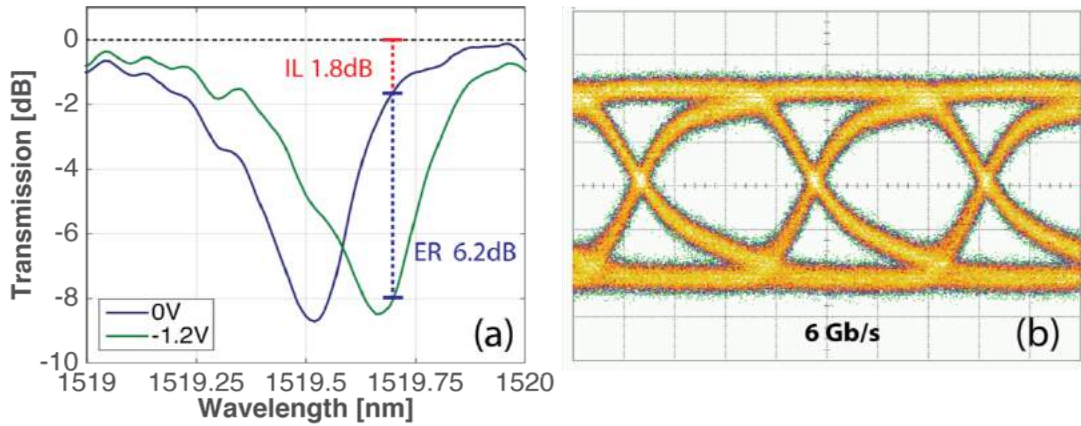


Figure 5.6: The Tx microring resonance characteristics and eye diagrams are shown.

5.4.2 Receiver Design

The receiver, shown in Figure 5.7, consists of a Ge photodiode placed on top of the electronics and connected to the receiver circuitry via TOVs with minimal parasitic capacitance.

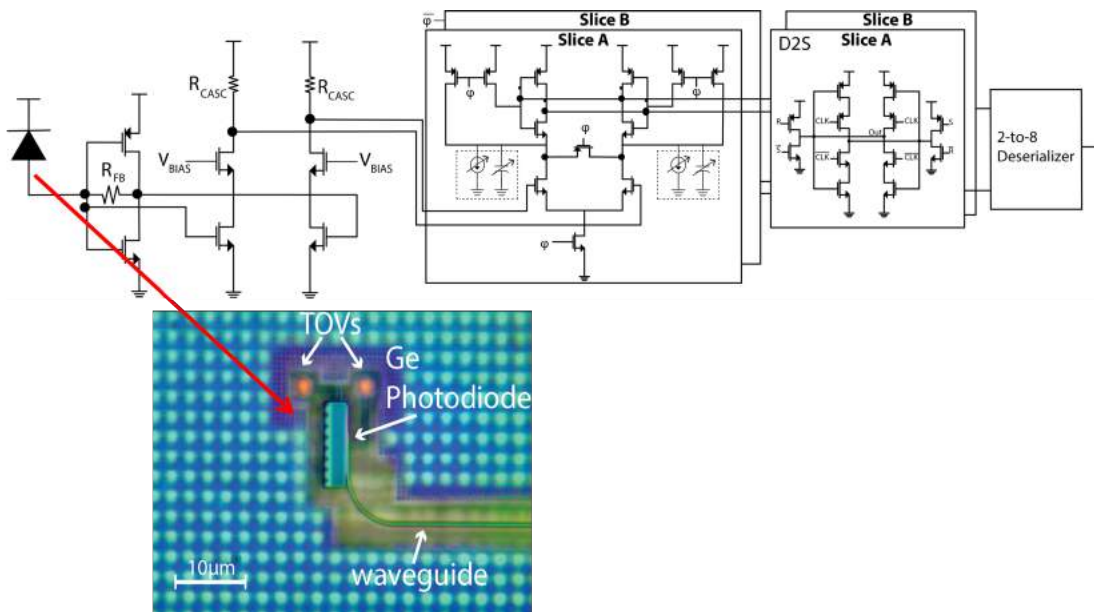


Figure 5.7: The receiver AFE as well as the photodiode is shown.

The TIA-based receiver circuit has a pseudo-differential front-end with a cascode pre-amplifier feeding into double-data rate (DDR) sense-amplifiers and dynamic-to-static converters (D2S). The TIA stage with 3kOhm feedback contains a 5-bit current bleeder at the input node, which is set to the average current of the photodiode. This

allows the TIA input and output to swing around the midpoint voltage of the inverter. The TIA input and output are directly fed into a cascode amplifier with resistive pull up.

The bias voltage of the cascode is tuned through a 5-bit DAC. Adjusting this bias voltage results in a trade-off between the output common-mode voltage and the signal gain of the cascode stage. More specifically, increasing this bias voltage results in a higher cascode gain but lower output common-mode voltage that reduces the sense-amplifier speed. For a given data rate, an optimal bias voltage is determined so as to minimize the overall evaluation time of the sense amplifier. The proceeding sense amplifiers then evaluate the cascode outputs before getting deserialized and fed into on-chip BER checkers. Each sense amplifier has a coarse, 3-bit current bleeding DAC as well as a fine, 5-bit capacitive DAC for offset correction. An external Mach-Zehnder modulator with extinction ratio of about 10dB driven by an FPGA-sourced PRBS sequence is coupled into the chip to enable stand-alone receiver characterization. During the initial seeding phase, the incoming receiver data are used to seed the on-chip PRBS generators for the BER check. The receiver and deserializer achieve 7Gb/s with a BER below 10^{10} .

The responsivity and bandwidth of this process variant of the Ge photodiode in [2], are shown in Figure 5.8.

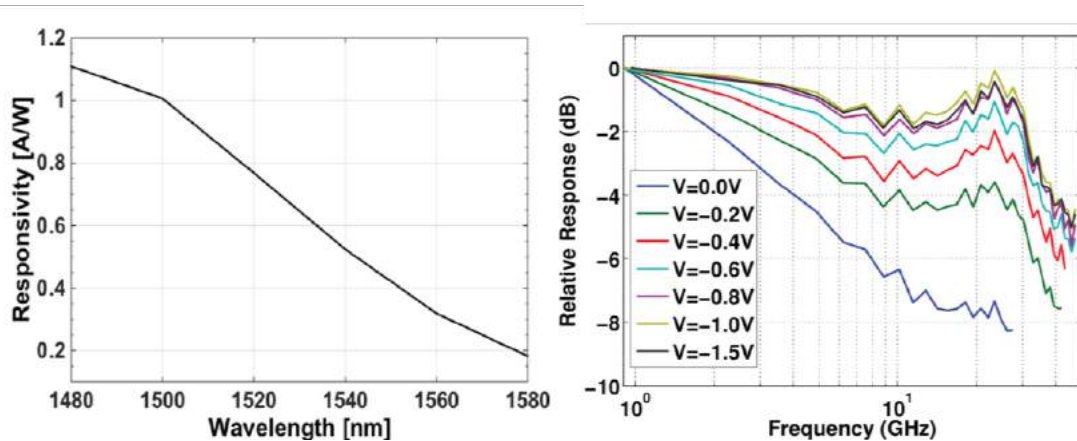


Figure 5.8: The Rx photonic ring responsivity and response versus frequency.

At 1520nm, the responsivity is 0.73A/W, resulting in optical RX sensitivity of 14.5dBm at 7Gb/s, for electrical sensitivity of 26A. The overall energy consumption is 340fJ/bit. The TIA+cascode pre-amplifier stage consumes 70fJ/bit. The sense amplifier, current plus capacitive correction DACs, and the dynamic-to-static converter together consume 120fJ/bit. Finally, the deserializer consumes 150fJ/bit. Figure 5.9 shows the sensitivity of the receiver as a function of data rate. Additionally, bathtub curves for the two slices of the DDR receiver are also shown.

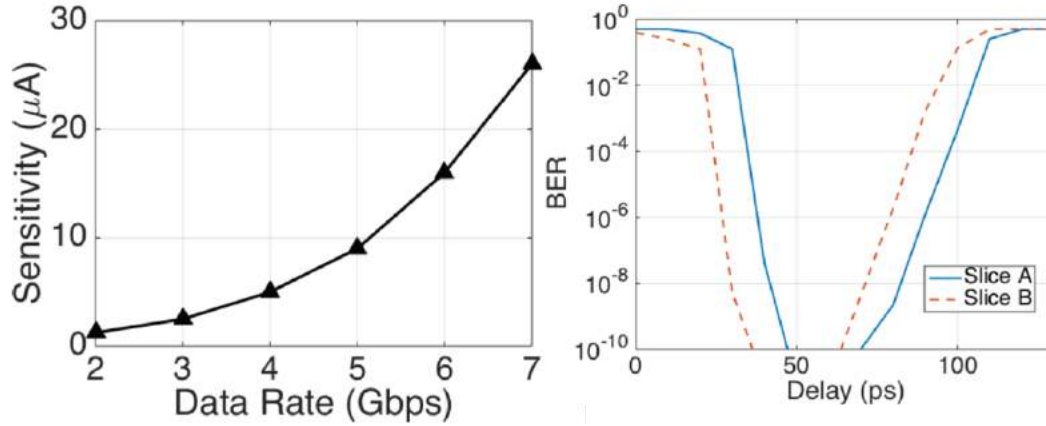


Figure 5.9: Measured receiver average photo-current sensitivity over different data rates and BER bathtub curves for both receiver slices.

5.4.3 Thermal Tuner Design

We designed thermal tuning circuits to stabilize the resonance wavelength of microring resonators in order to compensate process variations and temperature fluctuations. The thermal tuner for microring transmitters is based on a bit-statistical tuning algorithm [31]. The similar thermal tuning backend is implemented in 65nm process. The system diagram of the tuning backend is shown in Figure 5.10.

As shown in Figure 5.10, a drop port waveguide is weakly coupled to microring resonator to detect power level inside the microring. The photocurrent at the drop port is then integrated and quantized by a ring oscillator based SAR ADC. The power strengths for optical level 1 and 0 can be calculated by the tuning backend based on the knowledge of transmitted data. With the goal to maximize the optical eye opening, a thermal controller actively sets the coefficients for a sigma-delta heater driver. This heater DAC drives the embedded silicon heater inside the microring and controls the local temperature and thereby the resonance of the microrings. For initial locking, the heater strength is swept to search for the laser wavelength and optimal locking point (Figure 5.11).

The optimal heater strength for maximizing optical eye diagram is stored in this initial sweeping process. The heater strength is then reset to this optimal value while the thermal tuning loop continues to thermally lock the microring. The captured eye diagrams in a slowed down thermal locking process show that the thermal tuning loop works as expected.

5.4.4 Link Implementation and Test setup

A 100-meter optical link operating at 5Gb/s is demonstrated (Figure 5.13) illustrating the functionality of all the required optical and electrical components in

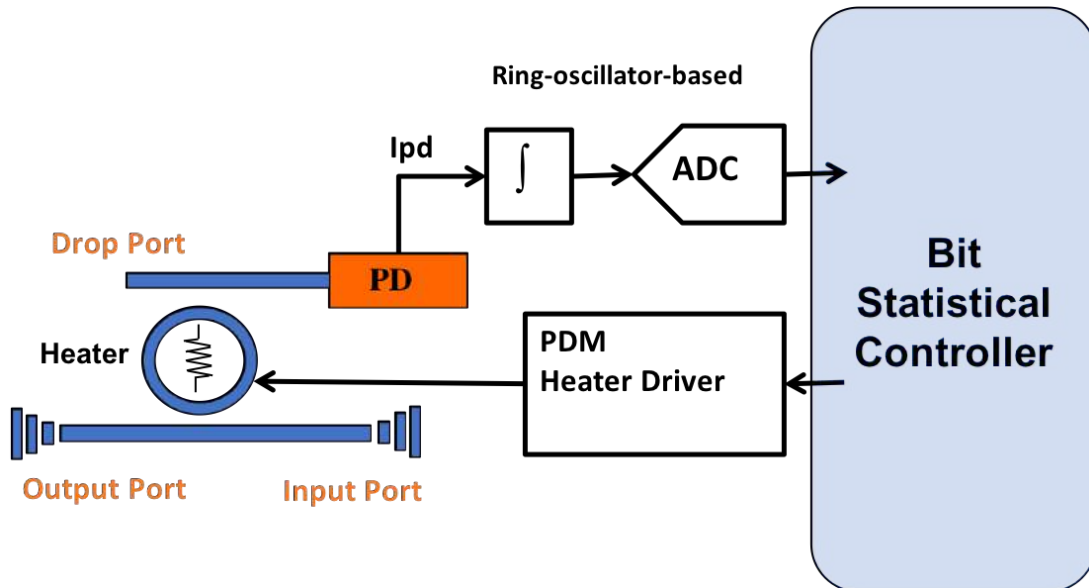


Figure 5.10: The thermal tuner block diagram used to control the microring resonance is shown.

this heterogeneous platform. Figure 5.13 also shows the optical power breakdown per stage within the full link.

A CW laser at 1517nm is coupled to the on-chip TX macro of Chip 1 using a vertical grating coupler. The coupler results in 7.5dB of loss in optical power. A PRBS generated data within this TX macro are fed into the modulator driver, which in turn modulates the ring resonator. The output of the TX macro including the coupler is the modulated light with 6dB extinction ratio. This light is fed into an optical amplifier providing 8dB of gain. The 8dB amplifier is necessary to mitigate part of the 15 dB chip-to-chip coupler loss in the optical data path (7.5dB per coupler) due to unoptimized coupler designs. The amplifier feeds into the 100 meter fiber preceded by a 90/10 power splitter. A monitoring scope, using the 10% output, is used to ensure that an optical eye is visible. The 90% output is coupled into the RX macro. The Ge photodiode is used within the RX macro to convert incoming optical data to an electrical bit stream. This photodiode sees 12.3dB and -18.3dB optical power for a bit 1 and 0, respectively. Figure 8 shows the output BER plot indicating at least 10-10 bit accuracy. This BER plot sweeps two parameters within the RX macro. First, the delay of the RX clock with respect to the TX clock is shown on the x-axis. Second, the corrective capacitor DAC within the receiver sense amplifiers is swept and shown on the y-axis. For particular delays and capacitive DAC values, a steady BER $\geq 10^{-10}$ is observed, illustrating the margins for the robust operation of the link. The transceiver electrical energy cost is 560fJ/bit and the optical energy cost is 4.2pJ/bit (taking into account the amplifier gain). With optimized couplers (< 3 dB

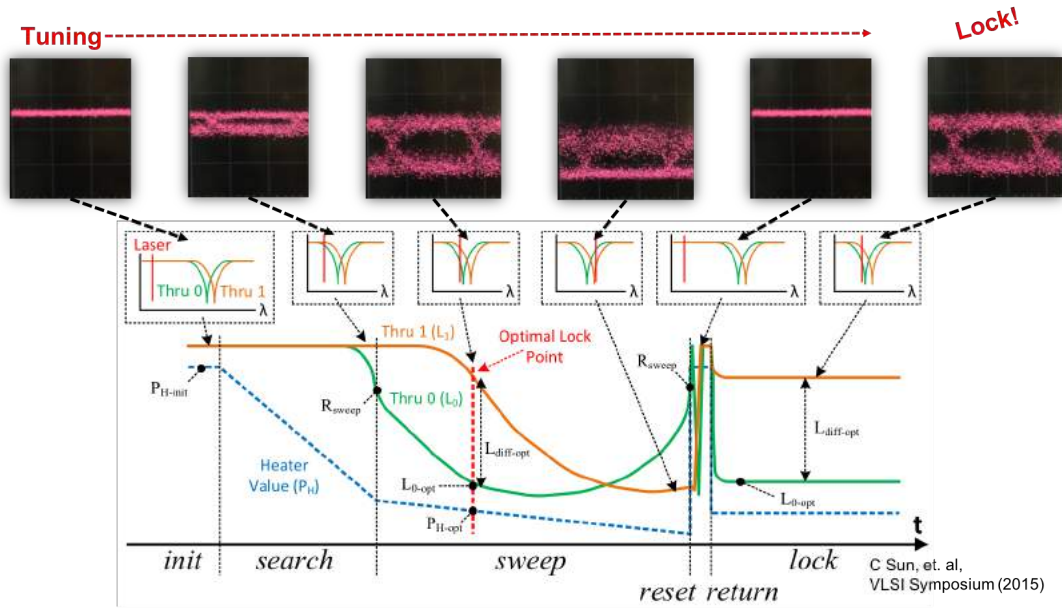


Figure 5.11: The progression of the transient eye along with the resonance location for the thermal tuner.

readily achievable in literature [6]), the required optical energy would scale down to below 0dBm (200fJ/bit) thereby eliminating the need for the optical amplifier.

Figure 5.14 shows the electrical power breakdown of TX and RX macros within the link at 5Gb/s data rate. Table I presents the comparison to previous non-monolithic electronic-photonic transceiver works.

5.5 Conclusion

This work demonstrates the first large-scale 3D integrated photonic chip-to-chip link manufactured in a 300mm CMOS foundry. The functional 3D-assembled chips with 16M transistors and 1000s of photonic devices illustrate the high yield of the CMOS, photonic fabrication and 3D integration processes.

A full optical chip-to-chip link is demonstrated for the first time in a wafer-scale heterogeneously integrated platform, where the photonics and CMOS chips are 3D integrated using wafer bonding and low-parasitic capacitance thru-oxide vias (TOVs). This development platform yields thousands of functional photonic components as well as 16M transistors per chip module. The transmitter operates at 6Gbps with an energy cost of 100fJ/b and the receiver at 7Gbps with a sensitivity of 26 μ A (-14.5dBm) and 340fJ/bit energy consumption. A full 5Gbps chip-to-chip link, with the on-chip calibration and self-test, is demonstrated over a 100m single mode optical fiber with 560fJ/bit of electrical and 4.2pJ/bit of optical energy. These results show




Figure 5.12: The test setup of the EPHI chip contains the Tx and Rx macros connected by a 100m fiber reel.

that the 3D integrated electronic-photonics platform holds great promise for future energy-efficient high-speed WDM communication links.

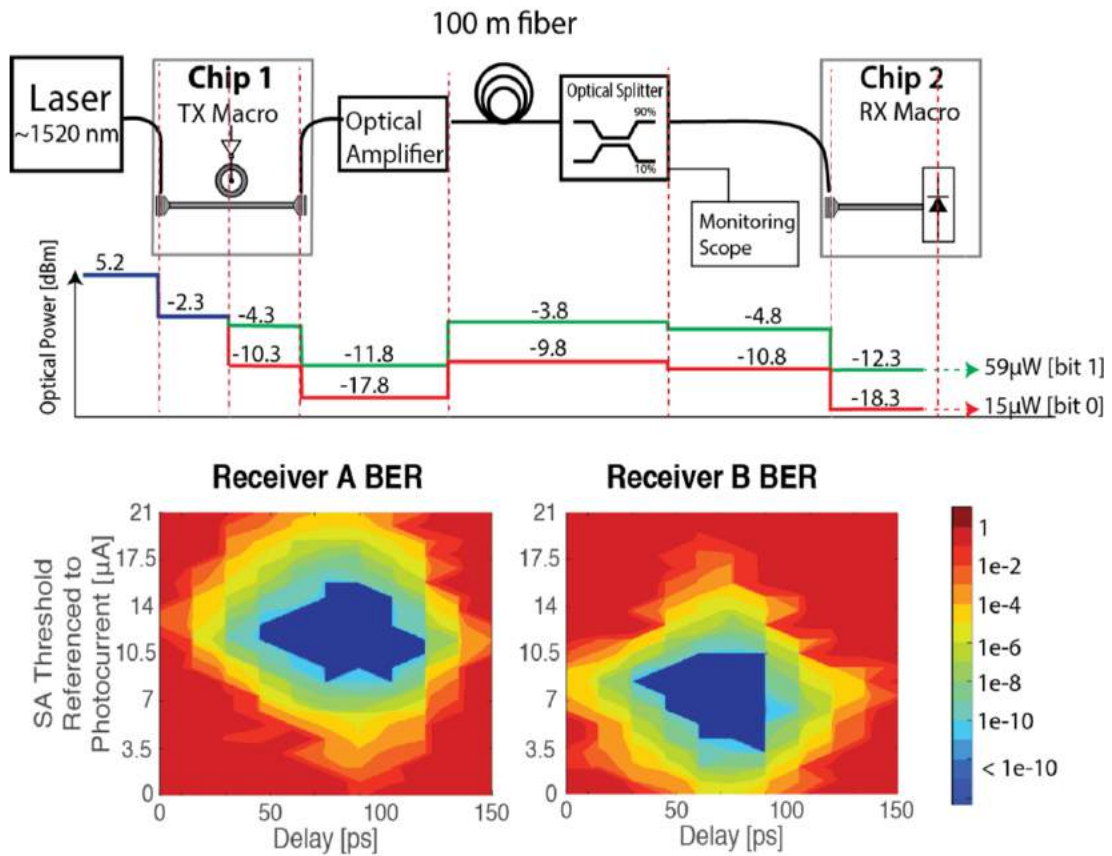


Figure 5.13: Full optical link with optical power budget and performance.




Figure 5.14: Electrical energy breakdown for the Tx and Rx macros in a 5Gb/s link.

	This work	[3]	[5]
CMOS Technology	65nm	40nm	65nm
Integration	Flip-Wafer 3D TOV	Flip-Chip μ bump	Flip-Chip
RX Data rate [Gb/s]	7	10	8
RX Energy [fJ/b]	340	395	275
RX Area [mm²]	0.0025	0.02	0.036
C_{interconnect} [fF]	3	30	200
RX Ckt. Sens. [μA]	26	20	53.7
RX Opt. Sens. [dBm]	-14.5	-15	-12.7
TX Data rate [Gb/s]	6	10	5
TX Energy [fJ/b]	100	140	808
TX Area [mm²]	0.0015	0.0012	0.04
Link DR [Gb/s]	5	10	-
Link Energy [fJ/b]	560	535	

Figure 5.15: Comparison with previous work.

Chapter 6

Single-Comparator PAM4 Architecture

6.1 Introduction

Pulse amplitude modulation (PAM) is an attractive link technique to double the number of bits per symbol while trading off front-end loading and reduced sampler swing. In particular link scenarios, as will be detailed in this chapter, the PAM4 architecture proves beneficial over its inferior NRZ-based counterpart. Moreover, a critical bottleneck associated with PAM4 receivers, namely the capacitive loading imposed on the front-end by the comparators, will be addressed and solved.

This chapter will begin by analyzing the trade-offs of the PAM4 (i.e. two bits per symbol, with a total of four unique levels) and NRZ links under the context of the link framework developed in Chapter 3. Conclusions will be drawn as to exactly when the traditional PAM4 architecture is superior to the NRZ architecture. Next, in Section 6.3, the author will detail how the new single-comparator based PAM4 receiver architecture proves superior over the traditional architecture. By targetting and eliminating two extra comparators per way, the new architecture provides a much more efficient AFE design. This claim, once again, will be verified in the context of the link framework from Chapter 3. Moving on, the architecture's performance will be further analyzed from a link-level standpoint with a photonic and CMOS co-simulation. The chapter will conclude by setting a foundation for the physical design of the single-comparator architecture.

6.2 PAM4 Introduction and Link Trade-Offs

The PAM4 architecture is an attractive alternative to the traditional NRZ flavor. At each symbol, the PAM4 link transmits and receives two bits, while compromising certain link characteristics that influence the overall energy per bit. Figure 6.1 shows

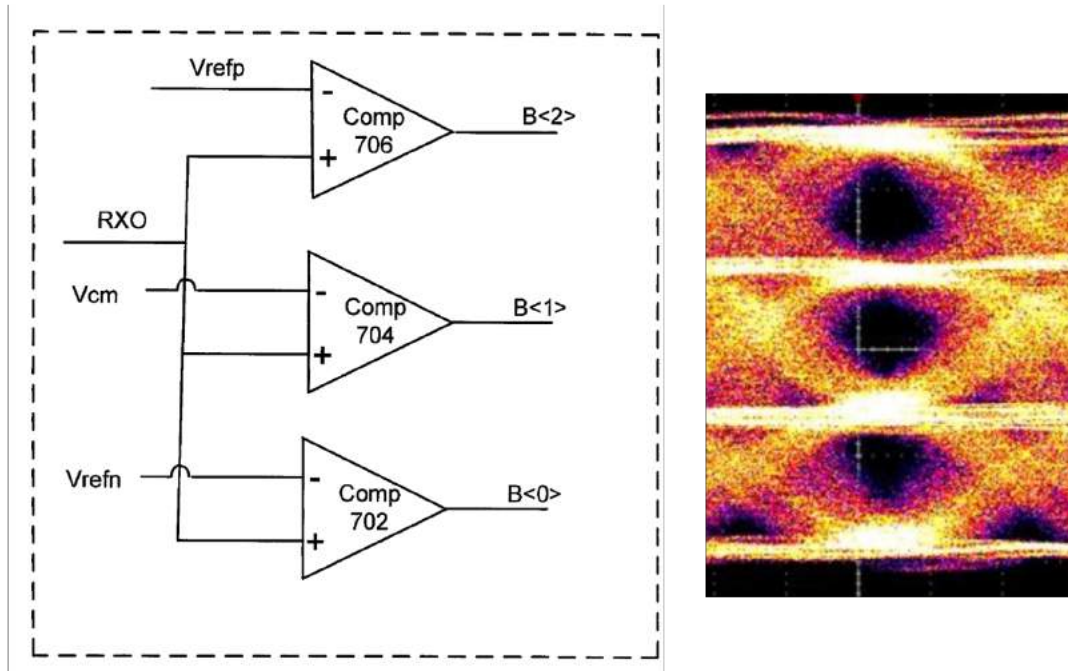


Figure 6.1: The traditional PAM4 architecture comprises of three comparators after the AFE to slice the 4-level eye.

the standard architecture of the PAM4 receiver, with three comparators at the output of the AFE (shown as the *RX0* wire). Each comparator is responsible for slicing one of the three eyes, shown in the right side of Figure 6.1.

From a high level, the comparison between PAM4 and NRZ architectures shows pros and cons on both sides. For instance, even though the PAM4 flavor offers twice as many bits per symbol, the number of comparators increases by $3\times$ per interleaving way. In addition, each comparator will see only $\frac{1}{3}$ of the effective front-end swing.

To rigorously analyze the pros and cons of each flavor, the PAM4 architecture was implemented within the context of the link framework developed in Chapter 3. For now, the technology parameters reflect the heterogeneous integration platform, similar to those in the previous chapter. Controlled restrictions were placed on the maximum allowable interleaving. This was done to truly highlight the space where PAM4 betters NRZ. Figure 6.2 shows the model results of this comparison.

The plot shows a multitude of different sweeps superposed on one another. All the dashed lines are the traditional, NRZ (i.e. PAM2) flavor. The solid lines are the PAM4 flavor. The colors represent the maximum allowable interleaving (i.e. $M=2$, $M=4$, $M=8$) for each sweep. Notice that when no restriction on topology is placed, both PAM4 and NRZ perform very similarly from an energy per bit standpoint. For instance, the *PAM4*, $M=2$ and *PAM2*, $M=4$ flavor perform virtually identically. However, for practical purposes, certain system constraints may impose a restriction




Figure 6.2: Removing all constraints on the receiver architecture shows that the PAM4 architecture is superior to NRZ under *particular system constraints*.

on the maximum allowable interleaving. Clocking overhead, for instance, is far higher for quadrature versus dual data rate operation. Not only is routing more complicated, but generating two extra clock signals that are aligned appropriately in phase is energy-heavy and requires more design. This, alone, may be enough of a burden on the designer to err towards lower interleaving. If so, this is the design space where PAM4 truly wins over NRZ. By restricting the amount of interleaving, one can see a noticeable difference in energy efficiency for the two flavors. At 30Gbps, for example, and restricting the interleaving to $M=2$, the E/b for PAM2 and PAM4 are 40pJ/b and 1pJ/b, respectively.

This difference in performance is further highlighted when restricting the number of interleaving ways to 1, even. Figures 6.3 and 6.4 show the full link energy per bit as well as the receiver side energy per bit, respectively. Once again, the heterogeneous platform parameters were utilized.

In this case, breaking down the full energies and the receive-side energies only showcases a few interesting secondary points to the overall link comparison. Mainly, even though the full link efficiency degrades significantly with data rate for the NRZ flavor, the receive-side energy for both the PAM4 and NRZ flavors are very similar,




Figure 6.3: A comparison of the E/b for a PAM4 versus NRZ link comprised of only a single interleaving way for slicing shows the benefits of PAM4 over NRZ.

with minimal overhead for shifting to PAM4. Moreover, there exist points in the data rate space where the PAM4 architecture is slightly more efficient than the NRZ flavor itself. However, the rationale for this is counterintuitive when considering the receiver side alone. To accurately understand this phenomena, the overall link picture must be taken into account. From a link standpoint, the energy trade-off between the Tx and Rx side at these high data rates shows benefit in expending more energy on Rx side for the NRZ versus the PAM4 flavor.

To conclude, even though adding comparators for slicing two more eyes imposes a slight energy penalty, the savings from a sensitivity standpoint are astronomically higher. This yields the conclusion that, once again, PAM4 is superior to NRZ under the restriction that a cap on the maximum interleaving is imposed. By giving the comparators $2 \times$ the time for each bit (assuming the same bit rate), the savings on energy are exponentially better due to the gain characteristics of the comparator, as detailed in Chapter 3.

Perhaps contrarily, when looking at this comparison from a sampler count-limited regime where the maximum number of samplers is restricted, it is apparent




Figure 6.4: A comparison of the E/b for a PAM4 versus NRZ link comprised of only a single comparator for slicing shows the benefits of PAM4 over NRZ.

from Figure 6.2 that there is no clear distinction between PAM4 and NRZ links when studying the energy per bit. Indeed, reducing the total number of comparators alone is the key to ensuring better link efficiency. Thus, techniques which target a multibit solution while reducing the comparator count promise better performance. This sets the foundation for the introduction of the single-comparator based PAM4 receiver in the next section.

6.3 Single Comparator PAM4 Receiver

Performance for traditional PAM4 receivers has been bottlenecked due to the significant loading effects imposed on the front-end by the samplers (i.e. $3\times$ loading per time-interleaved way).

Methods to reduce this loading could yield fruitful performance benefits, provided that the secondary effects of such techniques do not add burden to the front-end.

In this section, we will consider the benefits of capacitance reduction from the

samplers on overall link performance. This, in turn, will set the foundation for the new single comparator PAM4 architecture, at which point we will delve into the theory and results.

6.3.1 Motivation

To preliminarily quantify the effects of capacitance loading from the samplers, consider again the link framework developed in Chapter 3. The full link is shown again in Figure 6.5, with the critical node highlighted for emphasis. Reducing capacitance at this node is of paramount importance because all signal transistors in the front-end “see” a high speed data signal. To ensure maximum gain in the front-end chain, reducing capacitance is key.




Figure 6.5: Full, end-to-end drawing of the photonic link along with a point to the critical node – the primary concern of this chapter.

Now, to study the importance of this node, the PAM4 link optimization scripts from the previous section were run under similar conditions, except that the sampler loading imposed on the front-end was scaled down by $3\times$ (i.e. to mimic NRZ links). Said differently, this optimization was run under the condition that a *single* comparator was required per time-interleaved way to extract both the MSB and LSB. Moreover, this comparator is “similar” to that of an NRZ link (we will study in a future section the subtle differences in the comparator behavior when using the yet-to-be-proposed technique).

The results of this simulation are shown in Figure 6.6. The blue line shows the traditional, three-comparator-per-way, PAM4 link optimization results. The red line shows the theoretical performance of a PAM4 receiver that relies on a single comparator. All simulations were run using the 45nm SOI PDK platform metrics listed in Chapter 3. Notice that, at 50Gbps, an approximately 40% improvement is attained by reducing the effective capacitance seen by the front-end by a $3\times$ factor. Neglecting secondary non-idealities, this result shows the motivation for targetting capacitance reduction within and around the high speed front-end chain.




Figure 6.6: Preliminary link performance results show the benefit of scaling down the sampler capacitance by $3\times$.

6.3.2 Proposed Architecture and Formulation

With the motivation for capacitance reduction in place, the challenge now becomes how to actually realize a PAM4 receiver architecture utilizing only a single comparator (i.e. with a capacitance reduction). At present, simply tying up a single comparator to the end of a PAM4 receiver AFE will yield nothing but NRZ-equivalent behavior, as shown in Figure 6.7. More specifically, the StrongArm simply compares the relative magnitudes of V_{SA1} and V_{SA2} . Depending on which signal is larger, the outputs evaluate to either produce a dip in V_{O1} or a dip in V_{O2} . Indeed, if a differential PAM4 signal is assumed at the input of the AFE, this single-comparator flavor will extract nothing but the MSB. The LSB, at present, is not extracted by simply connecting a StrongArm to a PAM4 AFE.

To proceed with the formulation of this architecture, let's begin by considering what we already have – the StrongArm sense amplifier, shown in left of Figure 6.8. Aside from the schematic diagram, Figure 6.8 also shows two sets of waveforms in red and green (four waveforms in total). Each of the two sets corresponds with a particular input voltage stimulus into the StrongArm. The green set shows the output waveforms of the StrongArm when a relatively large input voltage stimulus




Figure 6.7: Behavior of a PAM4 receiver with a single comparator yield “NRZ-equivalent” behavior.

is applied. The red set, on the other hand, shows the output due to a small input voltage stimulus.

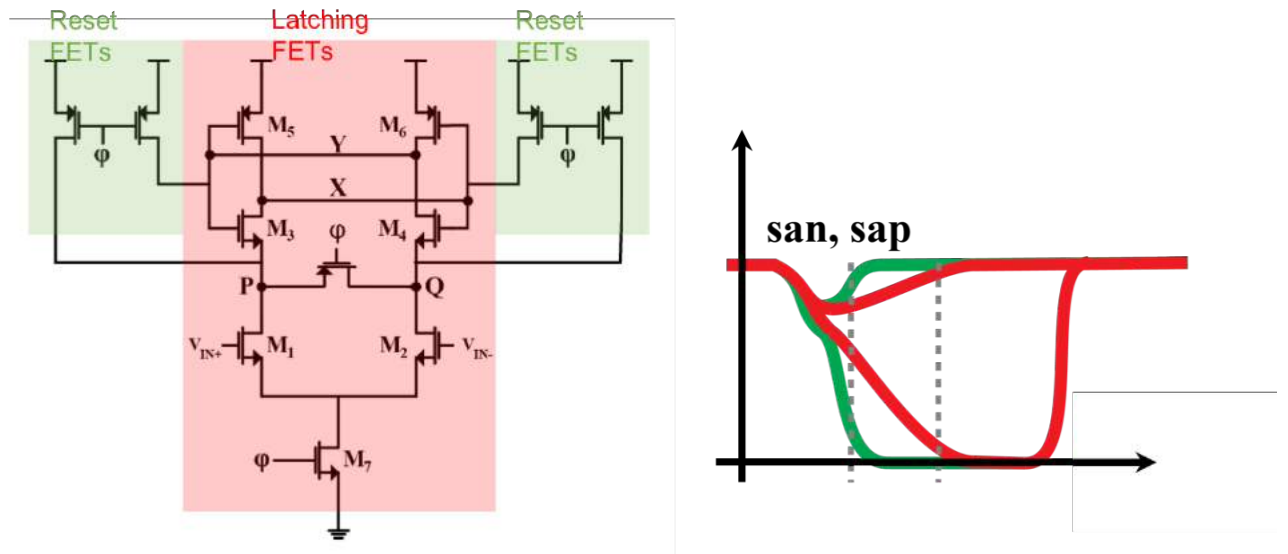


Figure 6.8: The StrongArm Sense Amplifier schematic along with two sets of waveforms (red and green) showing the output due to a large and small input signal, respectively.

The key distinction between the two sets comes down to the evaluation time of the StrongArm. Said differently, a small input voltage requires a longer time in order of the StrongArm to evaluate, whereas a large input voltage evaluates quicker. Aside from this intuition, the behavior is further reinforced through the simple StrongArm regeneration time approximation below:

$$T_{reg} = \tau_{regen} \ln \left[\frac{\Delta V_o}{\Delta V_i} \right], \quad (6.1)$$

$$\tau_{regen} = \frac{C_L}{g_m} \quad (6.2)$$

As can be seen in Equation 6.1, for a fixed output voltage, ΔV_o , the time to evaluate, T_{reg} , shortens with larger input voltages, ΔV_i . Indeed, the *true* time-dependent nature of the StrongArm is more involved (with slightly differing integration times as well), but the basis for that analysis can be studied in Chapter 3.

To summarize our findings, a single comparator appended to a traditional PAM4 receiver provides two sources of information:

1. The polarity of the incoming signal.
2. The duration of the pulse at the output of the StrongArm, which is dependent on the input differential magnitude.

With that said, the questions to address to realize a functioning single-comparator PAM4 link are as follows:

- **Item 1** How do you ensure a different magnitude input dependent on the PAM4 signal?
- **Item 2** How are the fast and slow evaluation times differentiated?

Let's analyze each question in depth.

6.3.2.1 Item 1

Assuming a fully differential PAM4 front-end (including differential input signals), sample waveforms at the input to the StrongArm may look like those in Figure 6.9.

The dashed and solid lines show complementary PAM4 signals. Notice that the signals are not only centered about a “common mode”, but their respective differential magnitude (i.e. $|V_{dashed\ line} - V_{solid\ line}|$) is dependent on the logic level of the incoming PAM4 bit stream. Furthermore, notice that the absolute differential magnitude only takes on 2 values – this is labeled as “BigBit” and “LittleBit” in the figure. Notice, in particular, that the “BigBit” has an absolute voltage difference that is $3\times$ that of the “LittleBit”. This information, which may be looked at as how we view PAM4 signal encoding, is summarized in Table 6.1.

A combination of the raw StrongArm output (i.e. X and Y signals from Table 6.1) and the time difference in evaluation can now be used to extract the MSB and LSB.

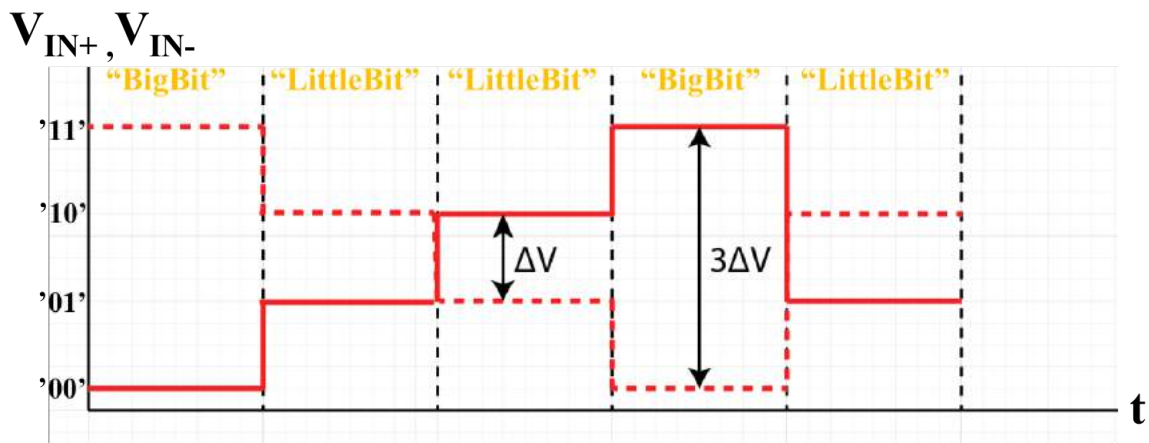


Figure 6.9: A fully differential PAM4 front-end will have two complementary outputs centered about the common-mode.

PAM Alternate Signal Interpretation Encoding Scheme						
Input Characteristics			Output Characteristics			
MSB	LSB	$V_{IN+} - V_{IN-}$	X	Y	Bit Type	Eval Time
0	0	$\frac{3}{2}V_{LSB}$	0	1	BigBit	Fast
0	1	$\frac{1}{2}V_{LSB}$	0	1	LittleBit	Slow
1	0	$-\frac{1}{2}V_{LSB}$	1	0	LittleBit	Slow
1	1	$-\frac{3}{2}V_{LSB}$	1	0	BigBit	Fast

Table 6.1: This table summarizes the alternate interpretation of the PAM encoding scheme, using a single comparator’s “timing information”.

6.3.2.2 Item 2

With the notion that bit times are dependent on the input bit stream, the question now becomes “how do we distinguish between the fast and slow evaluations?” More specifically, what circuit can be used to differentiate between the red and green lines in Figure 6.8. This circuit would be placed at the output of the StrongArm and ideally have as little capacitive overhead as possible.

Although many options exist, the most straightforward solution (and, as will be apparent later, the solution with the least capacitive overhead) comes down to the baseline distinction between the two waveforms: time. By focusing on doing a direct time to digital conversion, inefficiencies caused by having intermediary mediums such as charge or current, are avoided.




Figure 6.10: A time-to-digital circuit takes an input the differential outputs of the StrongArm sense amplifier.

The circuit in Figure 6.10 is an example of a time to digital converter. The NAND gate takes as input the differential outputs of the StrongArm sense amplifier. The output of the NAND gate produces a pulse with its width dependent on the quickness of the StrongArm during the evaluation phase – if the StrongArm evaluates quickly, the NAND triggers earlier. Keep in mind that the NAND always falls back to zero during the reset phase of the StrongArm (when signals *san* and *sap* are both set to 1). The follow-on buffer then cleans up and strengthens the signal before going into flip-flop. The flip-flop uses a reference clock that is parked in between a fast and slow evaluation (see Figure 6.11).




Figure 6.11: The sample waveforms of the TDC block show pulse width’s dependency on the output StrongArm waveforms.

If the input to the StrongArm is small, the StrongArm takes a longer time to evaluate, causing the NAND to produce a narrower pulse, which lastly causes the flip-flop’s clock edge to “miss” the input signal causing a logic 0. On the other hand, if the input to the StrongArm is large, this yields a quicker evaluation and a wider NAND pulse, which then the flip-flop clock “sees”, resulting in a logic 1. The narrowest this pulse difference can be depends on the aperture of the flip-flop itself.

When viewing the TDC-based approach from the context of the full link, a few additional insights may be extracted. These insights are grounded on the research presented in Chapter 3 and take into account bit time constraints and constraints

imposed by the nature of interleaving StrongArms. The insights, combined with the above analysis, are summarized below:

- **MSB Fidelity** – The StrongArm outputs must evaluate successfully (where success is considered based on the follow-on latch’s correct evaluation; a rigorous analysis of this success criterion takes into account the BER characteristics of the full link itself).
- **LSB Fidelity** – The TDC pulse difference between a BigBit and LittleBit input is long enough to result in correct evaluation by the TDC flip-flop.
- **Integration Time Restriction** – The integration time of the StrongArm is smaller than one bit time (from Chapter 3).

To summarize, the single comparator PAM4 receiver architecture offers benefits by reducing the sampler capacitance exposed to the front-end by $3\times$. In addition, by using both the StrongArm inputs differentially, the necessity of external voltage levels (for PAM4 decoding) become unnecessary. As such, DACs or reference generators along the signal path are avoided when using this new architecture scheme.

Next, we delve into the theory and design methodology of this single comparator architecture before detailing the design and layout process.

6.3.3 Design Methodology and Theoretical Performance

The aforementioned discussion has led us to the high-level link architecture shown in Figure 6.12.

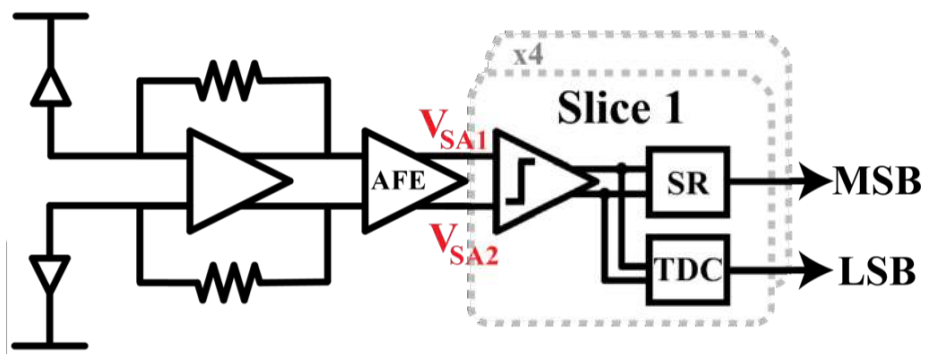


Figure 6.12: The TDC receiver architecture composed of the AFE and a single interleaving way composed of the StrongArm, D2S, and new TDC (from Figure 6.10).

To proceed with the design, we now look further into the time-dependent characteristics of the TDC as well as discuss the added constraints on the StrongArm sense amplifier.

Let's begin by analytically calculating the approximate time difference that results from a BigBit versus a LittleBit. For this calculation, assume that the follow-on flip-flop in the TDC has a constant propagation delay (independent of the input signal swing or rise/fall times). Additionally, assume that the integration time of the StrongArm is also fixed (and independent of input signal swing or rise/fall times). Now, the evaluation time of the StrongArm becomes input signal-dependent and based on the regeneration time, as shown in Equation 6.4.

$$\tau_{regen} = \frac{C_L}{g_m} \quad (6.3)$$

$$T_{reg} = \tau_{regen} \ln \left[\frac{\Delta V_o}{\Delta V_i} \right] \quad (6.4)$$

Here, V_o and V_i are the output and input voltages of the StrongArm, respectively. In the context of the PAM4 link, the BigBit and LittleBit voltage stimuli have a 3× difference between themselves. As such, the time difference between the BigBit and LittleBit stimuli can be equated as follows:

$$T_{reg,diff} = |T_{reg,BigBit} - T_{reg,LittleBit}| \quad (6.5)$$

$$T_{reg,diff} = \left| \tau_{regen} \ln \left[\frac{\Delta V_o}{3\Delta V_{LSB}} \right] - \tau_{regen} \ln \left[\frac{\Delta V_o}{\Delta V_{LSB}} \right] \right| \quad (6.6)$$

$$T_{reg,diff} = \tau_{regen} \ln(3) = \frac{C_L}{g_m} \ln(3) \quad (6.7)$$

Notice that the time difference between the two stimuli is only dependent on the regeneration time constant and the BigBit to LittleBit voltage fraction. However, when introducing non-idealities such as ISI or front-end bandwidth degradation, this voltage fraction is reduced.

As shown in Figure 6.13, the worst-case evaluation time difference occurs when comparing voltages between the smallest BigBit and the biggest LittleBit (notice the particular size of the red curly braces in Figure 6.13). Preliminary simulations show that these degradations result in a level difference of approximately 1.9. Moreover, this degradation, along with the nominal regeneration time constants for the 45RFSOI PDK, show the time difference to be at approximately 15ps.

From Equation 6.7, the initial design methodology to maximize the pulse difference may be to make the regeneration time constant, τ_{regen} , larger. This can be done by either increasing C_L , or reducing g_m .

6.3.3.1 The Problem With the StrongArm

When studying Equation 6.7, it becomes apparent that maximizing the time difference requires tuning either C_L or g_m or both. Indeed, ideally, tuning each




Figure 6.13: Non-idealities in the AFE result in a voltage ratio between the BigBit and LittleBit that is smaller than theoretical $3\times$.

parameter would minimally if not all influence the other parameters and a singular dependence between the time difference and each parameter would exist. However, in the case of a StrongArm sense amplifier, a clean relationship between the input and output parameters does not exist. For instance, altering the C_L , for example, yields secondary effects on the integration time (slowing the integration time). This, in turn, bottlenecks the “Integration Time Restriction” from before (*The integration time of the StrongArm is smaller than one bit time*). To review, the integration time expression is repeated below:

$$T_{int} \approx \frac{V_{TH}(2C_{PQ} + C_L)}{g_m(V_{CM} - V_{TH})} \quad (6.8)$$

Notice the dependence on C_L . Thankfully, a cleaner solution for the sampler exists. By using a double tail sense amplifier instead, the direct dependence between C_L and integration time is removed.

The double tail sense amplifier, shown in Figure 6.14, follows very similar mechanics as the StrongArm. The integration phase of the double tail is controlled by the first stage diff amp structure composed of M_1 thru M_5 . Moreover, the integration time is strictly controlled by the loading of the first stage and the g_{m1} of the input pair. The regeneration time is dependent on the second stage’s cross-coupled inverter pair. This is summarized in the expressions below:




Figure 6.14: The double tail sense amplifier has benefits in the new PAM4 context that outweigh the traditional StrongArm sense amplifier.

$$T_{int,DTSA} = \frac{2 * V_{thn} * C_L}{I_{CC}} \quad (6.9)$$

$$T_{regen,DTSA} = \frac{C_L}{g_{m,cc}} \ln \frac{V_{T,D2S}}{V_{o,integ}} \quad (6.10)$$

Here, I_{CC} is the current drawn from the input stage differential amplifier composed of M_1 through M_5 . $V_{T,D2S}$ is the trip point of the follow-on D2S latch. Lastly, $V_{o,integ}$ is the voltage difference at the end of the integration period, detailed below:

$$V_{o,integ} = 2 \frac{V_{thn} * g_m * \Delta V_{diff}}{I_{CC}} \quad (6.11)$$

$$V_{diff} = \frac{t_{int} * g_{m,in} * \Delta V_{in}}{C_{diff}} \quad (6.12)$$

C_{diff} is the capacitance at the output of the differential pair.

Perhaps the most important rationale to motivate the double tail over the StrongArm is the following: the primary (and even secondary) loading effects caused by the

output load capacitance now only dictate the regeneration behavior of this latch. Said differently, the presence of transistors M_6 and M_7 isolate the first and second stages altogether. This yields a much cleaner relationship between the input parameters (C_L and g_m from Equation 6.7) and the maximum output time difference.

We can plot the expression in Matlab, as shown in Figure 6.15 to reveal the effective time gain for various load capacitances. This shows that the simulated time difference between the large and small voltage differences is upwards of 8ps for a 100mV input swing.




Figure 6.15: The double tail sense amplifier evaluation-time “gain” may be characterized using Equations 6.9 and 6.10. The MATLAB simulation results are plotted here.

6.3.3.2 TDC-Induced Sampler Sensitivity

Now, the question becomes “what attributes of the sampler are modified adversely due to the presence of the TDC?”. To a first order, the added capacitance on the output node of the sampler by the TDC impacts the overall performance. However, not only is that capacitance minimal (minimum sized inverter), but the upsizing of the cross-coupled pair should yield similar timing performance (for a slight power penalty).

Moreover, from studying Equation 6.7 further, it is apparent that a direct connection between the time difference and the input swing to the sampler does not exist. Rather, a maximum constraint on g_m is imposed (assuming a fixed C_L). As an example, for values of $T_{reg,diff}$ and C_L of 15ps and 5fF, respectively, the effective g_m is about $220\mu\text{S}$ in the 45RFSOI PDK. This restriction can be placed within the overall link analysis presented in Chapter 3. As one can see, this maximum g_m influences not only the sampler gain but the input referred sampler noise as well.

6.3.3.3 Model Results with new PAM4 receiver

To quantify theoretically the performance improvement over the traditional PAM4 receiver, the link optimization framework developed in Chapter 3 may be applied in this context.

Within the context of the framework, two important changes took place to facilitate the accurate comparison of these two architectures. Firstly, the comparator model was altered to reflect the utilization of the double tail sense amplifier as opposed to the StrongArm sense amplifier. Secondly, a maximum cross-coupled g_m constraint was placed on the double tail. Recall from Equation 6.7 that the maximum time difference between the *BigBit* and *LittleBit* is inversely dependent on the cross-coupled g_m . When fixing a minimum required time difference, a maximum g_m must, thus, be enforced.

Figure 6.16 shows the model results comparing the two architectures. In this case, similar PDK and integration styles as what is presented in Chapter 3 was used. Moreover, no restriction was placed on the overall link architecture, i.e. the number of amplifiers and comparators in parallel was left unrestricted. At 50Gbps, an improvement of over 30% is observed in the overall link E/b. This difference is further amplified at 100Gbps, where the improvement is over 55%.

6.3.4 End-to-End Photonic Co-Simulation Results

With the preliminary architecture and sizing methodology introduced above, we can now proceed to analyze the schematic behavior within the context of a full, photonic plus CMOS co-simulation. The co-simulation uses the VerilogA framework developed in the group to accurately model photonic components such as microrings, waveguides, and photodiodes. More information for the framework can be found in Sen Lin's dissertation [29].

The end-to-end link schematic is shown in Figure 6.17. The schematic is composed of the transmitter (serializer plus driver), the photonics test structures (modulator, waveguide, and microring), and the receiver (single comparator PAM4 architecture). All of the analog IP blocks are designed using Bag, with parameterizable sizing options for the various sub-blocks.




Figure 6.16: A comparison of the new, TDC-based PAM4 receiver and the traditional, three-comparator architecture show the potential benefits when viewing the link energy consumption. These results reflect not only the three-comparator difference, but also any secondary limitations on g_m due to the presence of the TDC.

The transmitter contains two 16-to-1 serializers and drivers, one for the MSB and one for the LSB. The driver signal then passes through to the microring modulator. The ring’s resonance shifts according to the incoming voltage. The output modulated light passes to the photodetector, which in turn supplies the receiver with an input current proportional to the incoming light intensity. The photonic components, as mentioned, utilize the VerilogA modeling framework to accurately depict waveguide loss, ring insertion loss, photodiode responsivity, and other key link-pertinent photonic parameters.

Once instantiated within the Cadence environment, the system can then be simulated similar to traditional circuits-only simulations. Particular care has to be taken within the simulation environment to ensure no “aliasing” occurs between the simulation time step and the round-trip time of the microring. Again, these details are highlighted in [29].

The photonic components in this simulation use parameters extracted from real photonic components designed using the AIM PDK. The testing to yield these results was done by AnalogPhotonics and MIT.

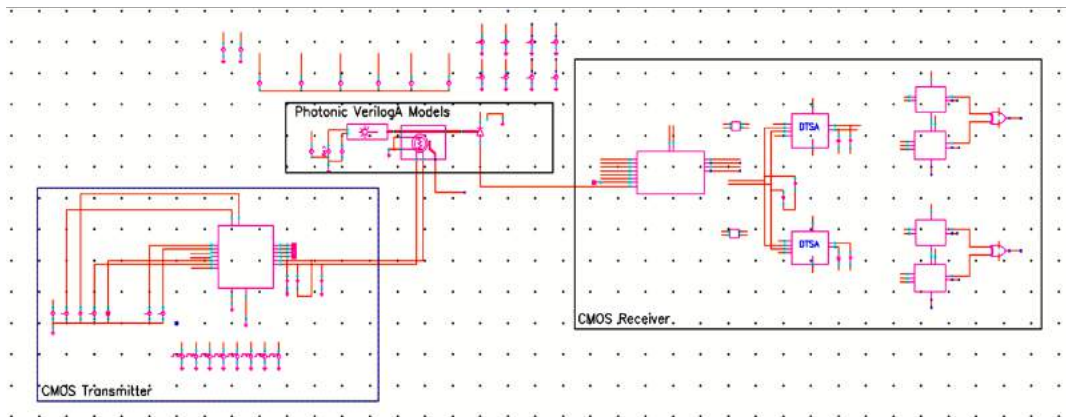


Figure 6.17: The end-to-end photonic co-simulation schematic shows the Tx driver, photonic components, and the CMOS receiver.

The respective eye diagrams of the input and output of the Rx macro are shown in Figure 6.18 and 6.19.




Figure 6.18: The receiver input current eye is shown. This signal was produced by a CMOS transmitter driving a photonic microring. The modulated light goes into a VerilogA photodiode to produce this eye.

6.4 Conclusion

In this chapter, we introduced the single comparator-based PAM4 receiver architecture. After motivating the necessity of PAM4 over NRZ encoding, the methodology for the architecture formulation was detailed. Next, the theoretical performance in the context of the link framework was studied. Schematic-level simulations were also

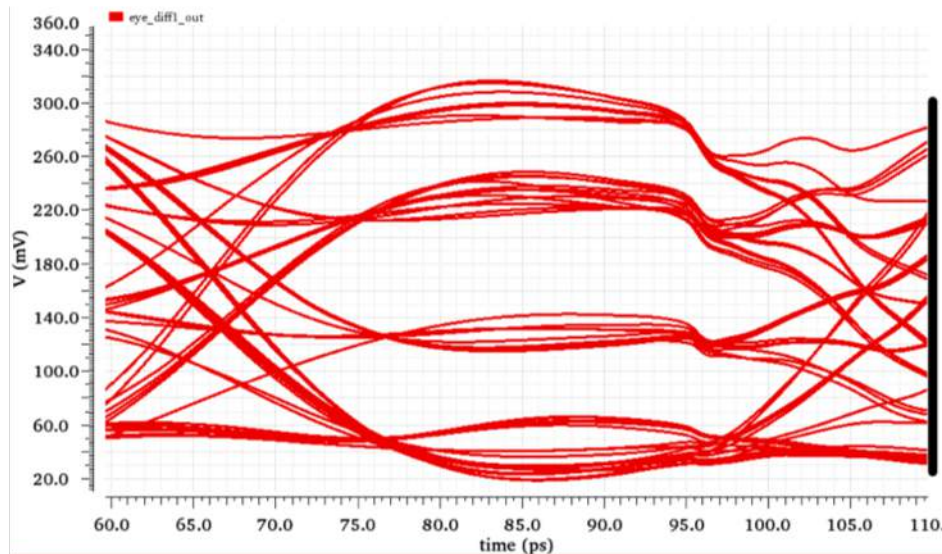


Figure 6.19: The output of the Rx AFE is shown. This signal subsequently traverses into a single slicer prior to digitizing.

run to characterize the design and showcase the signal fidelity for decoding both the MSB and LSB. Lastly, a full photonic plus CMOS cosimulation environment was set up, which placed realistic photonic VerilogA components alongside the optimized PAM4 transceiver.

This chapter shows how understanding the link framework and the trade-offs that transcend the hierarchy bring to light limitations and techniques to better optimize the link energy per bit. In the next chapter, this single comparator architecture will be further realized with the tape out of the Acacia system. The Acacia system aims to not only validate the technique, but also exercise the various physical trade-offs that exist during the design process. Moreover, we will see how these physical trade-offs marry tightly with the performance specifications we aim to optimize.

Chapter 7

Acacia System Design

7.1 Introduction

7.2 Acacia System Design

The new single-comparator PAM4 receiver was placed for testing and design in the Acacia SoC. The tape out occurred in May 2018, with collaboration from AnalogPhotonics and MIT (both partners through the AIM project). The tape out contained not only the single comparator circuitry, but a fully functional Rx backend, a Tx macro with a high-speed PAM4 transmitter, new non-drop port-based thermal tuner, and on-chip oscillator and clock distribution network. This section will detail the planning and trade-offs of the physical design of the PAM4 receiver and associated surroundings. The section will begin with insights and design aspects of the AFE macro itself, which is composed of all the core analog and mixed-signal blocks necessary for the link. In particular, the BAG generators necessary for the Tx and Rx AFEs are explored. We will then move on to detailing the full SoC itself, with particular emphasis on the components surrounding the AFE core such as the Tx and Rx macros, associated backends, and clock distribution blocks.

7.2.1 Front-End Floorplanning and PDK Generation Constraints

Perhaps the easiest way to plan and design the full receiver front-end (i.e. AFE core, slicers, DACs, and deserializers) is to analyze individually the constraints placed on the macro itself. These constraints help “close the loop” on the design process by, for example, altering the sizings of the FETs if necessary. Once again, this entire closed-loop approach is captured through the BAG environment. Any new changes in layout are instantly captured and fed into the Python code. The newly generated schematics and layouts are produced and DRC/LVS clean. The pre and post-PEX

simulations to verify performance are then run on the newly altered design before being stamped as tape-out ready.

7.2.1.1 Constraints

To ensure end-to-end link performance and functionality, the Rx front-end is surrounded by all necessary blocks to quantify link performance. Thus, even before formalizing the physical constraints, the functional design of the macro-system is itself constrained by the objectives. Said differently, these constraints impose the necessity of the digital backend with PRBS generators and BER checkers, the on-chip clock generation and distribution network, and various external bias and scan signals for further validation and functionality. Moreover, seeing as this is a photonic link, the CMOS circuitry will be bonded to a wafer consisting of the necessary photonic components (supplied by AnalogPhotonics). The CMOS and photonic wafers are flip-chip packaged, with the top metal layer of the CMOS wafer bumped to the photonics wafer metal layer. The bumping itself introduces parasitic capacitance which was taken into account within the design script for the front-end macro.

Signal-Critical Path Optimization

Perhaps the biggest byproduct of operating a link at such high speeds is the constraint placed on the routing of the signal-critical blocks. It comes as no surprise that the block sizings and inter-block trace sizings are heavily influenced by parasitics, which add capacitance, reduce net bandwidth, and thus, hinder overall link performance. To reduce these effects, the signal-critical blocks (AFE and slicers from the schematic in Figure 7.7) are placed first, with a key emphasis on reducing the wire connection length.

Figure 7.1 shows the AFE layout along with the slicers, D2S latch, and TDC circuitry. The blocks are all placed in a center-symmetric manner, with the inputs and outputs of each respective block straddling the vertical axis of symmetry. Upon hitting the quadrature slicers right after the differential amplifiers, the signals symmetrically extend left and right to reach the input of the samplers.

Indeed, should the size of any particular block vary, the width or length of the block may change. But, the block is always centered about the global vertical axis of symmetry. This ensures that the signal-critical wire lengths are as minimal as possible.

Process Variability and Offset

Another obvious constraint placed on the design of the front-end is the necessity of correction blocks to compensate for process variability and offset. Additionally, to avoid the need for a multitude of different testing blocks (i.e. current sources, voltage sources, etc.) on-chip DACs are used to generate the required biases for the various blocks. After design and placement, it is easy to see that these periphery

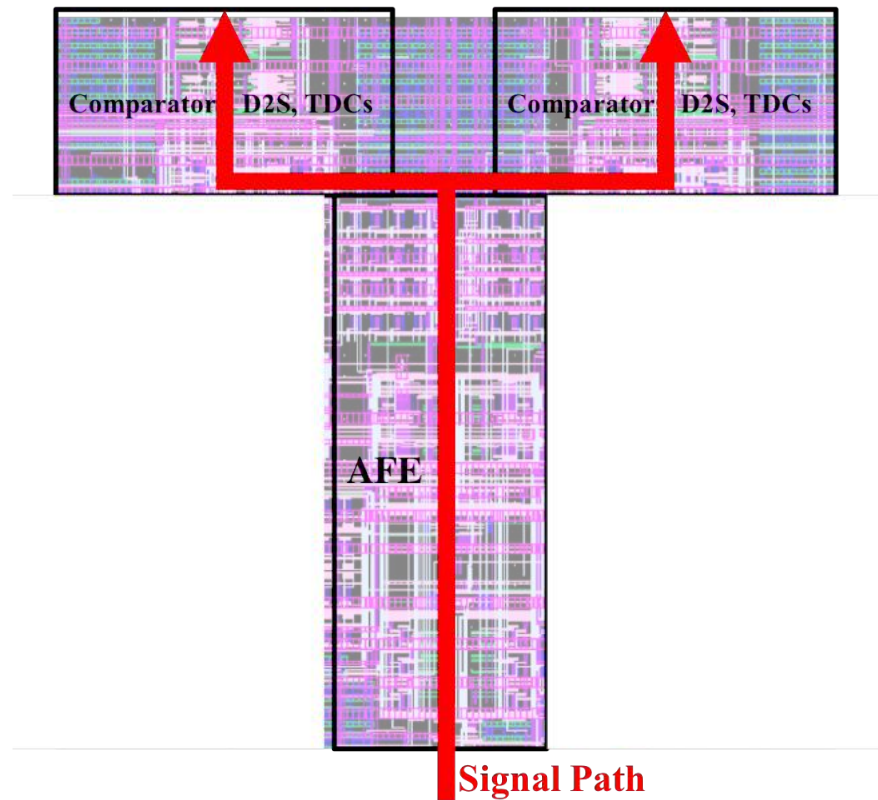


Figure 7.1: The AFE’s signal-critical blocks are placed first to ensure optimized performance and minimize path lengths and parasitics.

blocks dominate the area and impose the biggest constraints on floorplanning the macro (see Figure 7.2). Moreover, the DACs themselves add parasitics to the signal-critical blocks. Thus, the design script takes into account these added parasitics when optimizing. Indeed, all “higher-level” constraints on the area or size of the front-end macro alter greatly the design and placement of these DACs. Many times, the number of bits or the DAC architecture is sacrificed or altered for the sake of area management.

Figure 7.2 shows the Rx front-end blocks. The various blocks are generated in BAG for the 45RFSOI PDK. The deserializers and custom, standard-cell blocks (such as the TDC) are integrated as black boxes within the TemplateBase class of BAG. The overall front-end macro is $320\mu\text{m}$ by $320\mu\text{m}$. This size is restrained by the bump pitch constraint, which will be detailed next. The deserializer output signals, along with all necessary DAC control bits, are all connected to the digital back end of the RxMacro, where necessary scan chains or bit-error-rate checkers process or drive the various signals. Lastly, keep in mind that the input signals to the front-end come

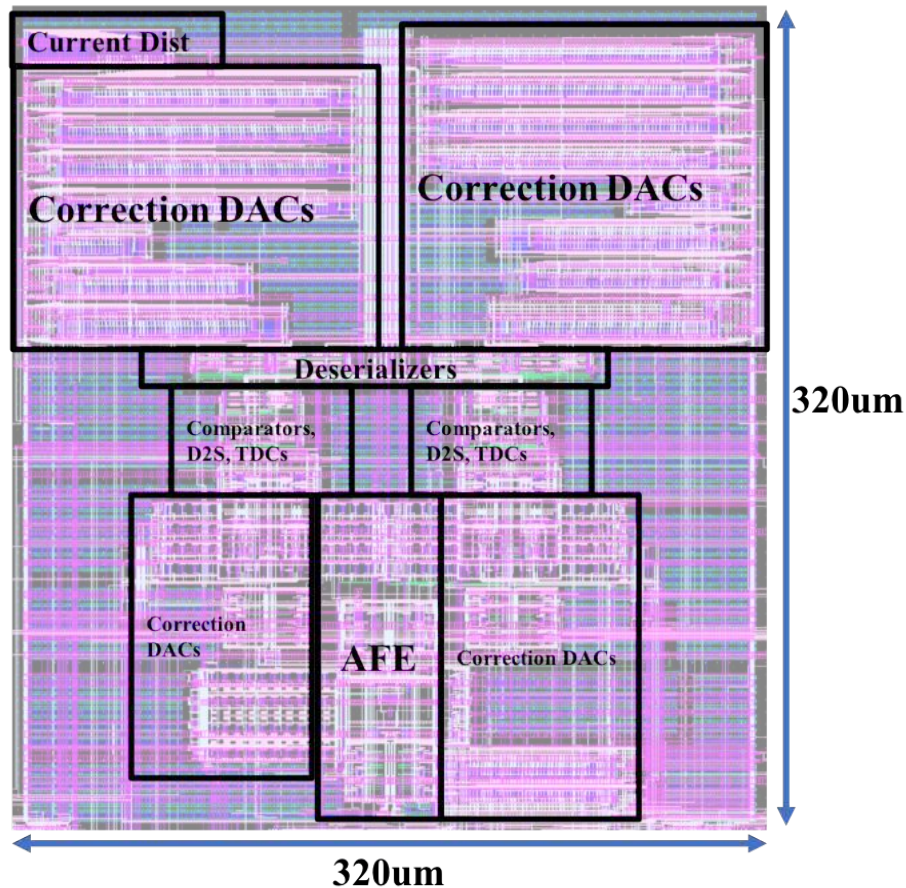


Figure 7.2: The full Rx front-end layout is composed of the signal-critical blocks, deserializers, and DACs.

in from the photodiode on the photonic reticle (which is bumped to the top metal layer of the CMOS chip). These top level pad connections happen using the digital place-and-route tool with very tight constraints on the signal length.

Bump Pitch Upper Bound

The photonics reticle contains the necessary components to realize a fully functional photonic link. As such, these components require proper electrical stimuli or recording to properly complement their behavior. The electronics needs to be pitch-matched with the photonic pads before being flip-chip bumped and connected. Figure 7.3 shows the placement of adjacent Rx front-ends (labeled as “Rx Site” in the figure). Each front-end macro is $320\mu\text{m}$ wide, with a bump pitch of $160\mu\text{m}$. The pads above the Rx Sites are the photodiode reverse bias signal and the TIA input signal. These two signals alternate when traversing left to right along the chip. Moreover, there are

two rows of these alternating bias and signal pads, which necessitates the presence of two rows of Rx front-ends. Additionally, each photonic site (i.e. ring filter) has an associated heater driver for thermal locking. Those heater drivers also have top level pads on the CMOS chip, with appropriate drive circuitry.

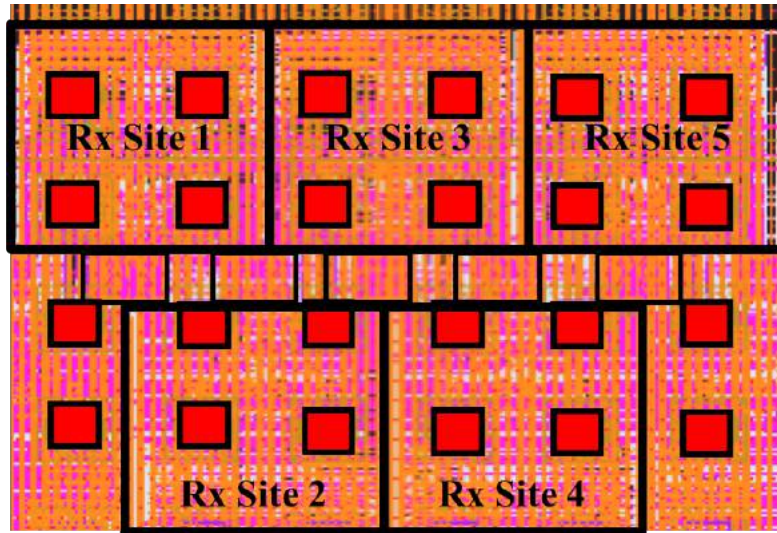


Figure 7.3: The bumps (shown using the artistically rendered red squares) are spaced evenly, with the purpose of interposing with the components on the photonic reticle.

For the purposes of minimizing the total routing on the signal-critical signals (both on the CMOS and photonic designs), the bump pitch was selected as $160\mu\text{m}$. Thus, the maximum width of each front-end macro is constrained to at most $2\times$ the bump pitch. Indeed, in order to interface between the digital backend for signals such as the digital control bits and outputs of the deserializers, gaps between the blocks are placed. This will be explored further in the next constraint.

A worthwhile discussion to have at this time is to consider the effects of varying (or reducing) the bump pitch on the overall floorplan of the Rx front-end. Of course, the signal-specific blocks and internal wires are the least favorable to change should the block size need reduction, seeing as their properties influence the link performance critically. Perhaps the biggest source of size reduction or change comes from varying the DAC resolution or architecture. Indeed, as it currently is, all of the DACs combine to occupy a majority of the usable area of the front-end macro.

Using the generator-based approach, variations in the bump pitch result in trivial iterations of the macro scripts. Figure 7.4 shows two particular flavors of the front-end, should the bump pitch either increase or decrease (left and right floorplans). Notice that aside from variations in the top current DAC widths, bias signals of the AFE along with the resolution of the capacitor DAC in the CTLE can also easily change, contingent on the allocated area of the macro.




Figure 7.4: Bump pitch variations result in very simple changes to the macro script in order to produce the two flavors of Rx above (drawn to scale relative to one another).

Routing Channels

As mentioned before, the Rx front-end contains a vast number of tunable DACs, offsets, and deserializer outputs which interface with the extensive digital backend. As such, these low speed signals need to be routed from the front-end into the digital backend. The presence of routing channels in between the Rx front-end macros alleviate the stress on the place-and-router by providing lanes for the digital bits to traverse.

Figure 7.5 shows the placement of these routing channels, laced in between the Rx front-end macros. The width of these channels is contingent on the minimum pitch of the routable metals, along with the number of digital bits that require routing. For instance, by allowing three metal layers to route in the channel, with each metal layer having a minimum pitch of $0.3\ \mu\text{m}$, and a total of 350 digital signals, the minimum required space of the channel would be $35\ \mu\text{m}$. Using such an approach and initial first pass estimate for the width of the channel avoids routing congestion and the tool freezing.




Figure 7.5: Routing channels, although constrain the maximum width of the Rx front-end, provide much needed space to allow proper routing between the front-end digital bits and the digital backend.

Metal Stack Restriction

The 45RFSOI technology was the chosen platform for the Acacia system. Although the platform has many benefits from a transistor performance standpoint, our selected metal stack is certainly the biggest drawback. This particular metal option was meant for RF applications. As such, “top” level metal layers were meant to be used for placing inductors and other RF components. Unfortunately, a byproduct of this design attraction is very, very large metal spacing and pitch for these metal layers. Moreover, the stack-up only allowed for 8 metal routing layers. Putting this together, this meant that this stack up allowed for 3 “normal” metal layers (i.e. M1, M2, M3) before exploding in dimension and size. M4 and M5 are substantially larger and require over a $6\times$ increase in pitch. M6 and M7 are substantially larger than M4 and M5, making them horrendously unattractive for a multitude of reasons. Lastly, M8 is the top layer of the chip, which interfaces with the photonics chip. M8 contains the pads and, due to DRC restrictions, no routing may extend from these pads.

Let us review: routing layers M1 to M7 are “usable”. However, lest we forget, because the technology standard cells (as well as the BAG primitive transistors) use M1 and M2 routing in their cells, this prevents the Rx front-end macro from utilizing these layers excessively. Moreover, dense power grids are necessary to ensure proper distribution of the various V_{DD} and V_{SS} signals. The top two metal layers not used for the flip chip pads (M6 and M7) are thus allocated for the global power grid. To

avoid complications with aligning local power grids in the Rx front-end, the power grid within the front-end macros are generally one or two layers below the global grid. Thus, ideally, M4 and M5 would be utilized for the local horizontal and vertical power grid. This leaves M3 only for routing! Seeing as that is vastly infeasible for any level of routing, sacrifices had to be made with either the allocated power grid layers or the avoidance of higher (and thicker) metal layers for routing. Thus, to allow inter sub-block routing and routing between the front-end and backend macros, M3 through M6 were used for routing. The local grid quality was sacrificed by just interspersing the power grid on M5 and M6. The global grid quality was also sacrificed by not having a dense grid of M6 and M7 over the macros. Instead, the grid was segmented based on the approximate signal routing density.

PDK-Specific BAG Primitives Aside from the poor metal stack options for the 45RFSOI technology, the transistors themselves require “special treatment” in the context of the BAG framework in order to ensure that DRC checks are handled properly. The 45RFSOI PDK offers two unique flavors of transistors: body-connected (BC) and non-body connected (NBC). These two flavors differ in a multitude of ways, with each flavor having its own unique set of pros and cons. Figure 7.6 shows the two flavors for visual comparison.

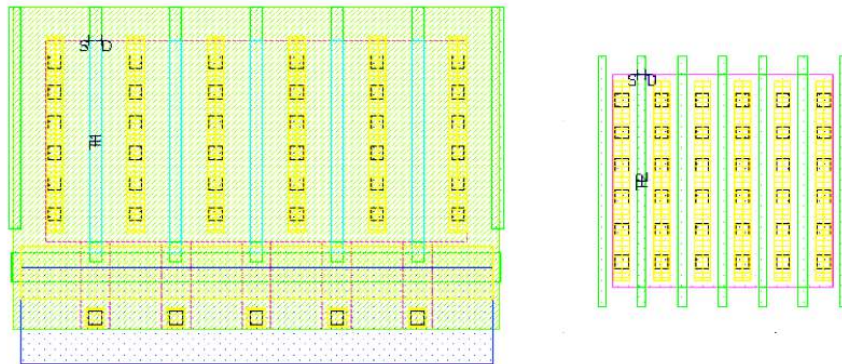


Figure 7.6: The body-connected and non-body connected transistor flavors are vastly different from the standpoint of the generator. This image shows the two flavors, with similar widths and 5 fingers.

Aside from the reduced gate pitch for the NBC flavor, the NBC flavor have a higher extrinsic f_T with reduced wire parasitics due to the smaller pitch requirement. However, the NBC flavor was intended for digital-like operation, meaning that the gates of the FETs were meant to be driven from rail to rail. This characteristic

manifested itself due to the hysteretic effects on the threshold voltage of the NBC transistor, which was measured in the lab. Due to channel charge leaking into the body area of the transistors, the threshold voltage varied. Moreover, this threshold voltage change was heavily dependent on the incoming data stream. For that reason, this flavor became unattractive for pure analog operation, or any situation where rail to rail operation was not achievable.

The biggest question for its usage versus the usage of the BC flavor came up when considering the StrongArm sense amplifier, or similar comparator. Because particular transistors were driven rail to rail within the block itself, it was thought that certain FETs may benefit from being NBC. However, experimentation showed that any combination of BC/NBC FETs within the StrongArm did not provide significant enough benefits over the all BC flavor-based architecture.

7.2.2 Receiver AFE Design and Insights

The core Rx AFE block is composed of many signal-path blocks and peripherals. Many of the PDK-specific constraints were taken into account while planning and designing this block. Aside from being able to decode the high-speed PAM signal, periphery voltage, current, resistor, and capacitor DACs exist to correct for process variability and offset.

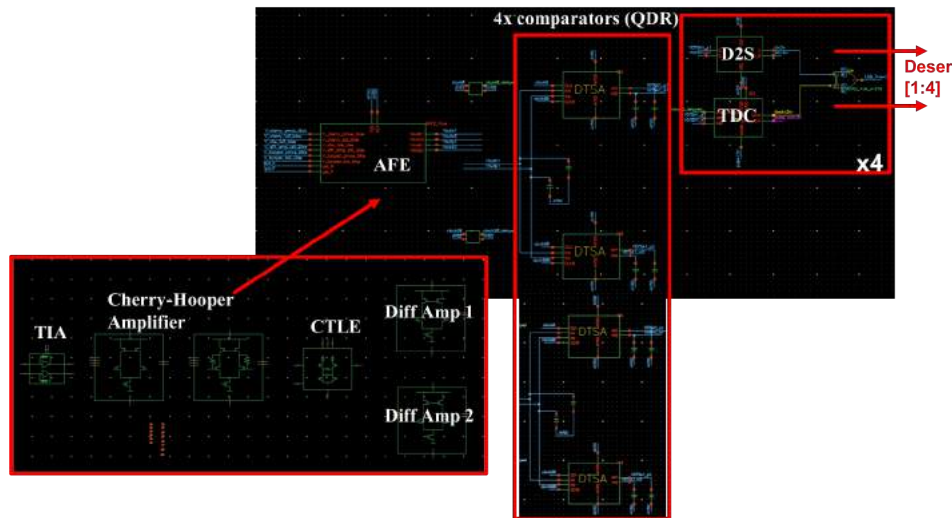


Figure 7.7: The Rx AFE (simplified) schematic shows the main blocks along the critical signal path.

Figure 7.7 shows the schematic of the single comparator PAM4 receiver. The schematic is composed of all the necessary signal-critical blocks along one slice of the receiver. The AFE core is composed of the TIA, cherry-hooper, CTLE, and differential amplifier block. Notice that, henceforth, all schematic and layout diagrams should

be assumed as BAG generated. The schematic generators, sub-block generators, and layout generators are all parameterizable and coded within the Python-based BAG environment. As an example, the pseudo-differential inverter-based TIA block within the AFE contains parameters to vary the NMOS and PMOS sizes, along with the value and size of the feedback resistor. Similarly, the follow-on amplifiers or equalizers all follow a similar generator methodology. In order to write the design script of the AFE generator, a similar methodology as to what was introduced in Chapter 3 was applied. Moreover, the added constraints from the TDC which were detailed in the last section, were also taken into account during the design and optimization of the single comparator PAM4 receiver.

7.2.3 Transmitter Design and Automation

Similar to the Rx AFE, the Tx AFE was designed as an executable generator. In this case, however, to optimize for FET speed and efficiency, the core of the designs (i.e. for the serializer) was done using Laygo. Laygo is analogous to BAG (or XBase) in that both give the designer power to optimally design and close the loop. However, Laygo has the benefit of allowing for “standard-cell-like layout” while also giving the designer finer abilities to write more compact layout generator scripts (at the expense of complexity). Additionally, Laygo allows the user access to the NBC flavored transistors which exhibit higher speeds at the expense of undesirable secondary effects.




Figure 7.8: The Tx-side AC driver uses a push/pull architecture to maximize the voltage swing across the microring modulator.

The Tx AFE driver is shown in Figure 7.8. The inputs to the driver (labeled as MSB and LSB) are the high speed, serialized signals outputted from the PRBS generator bit stream in the backend. Both the MSB and LSB serializers were designed using Laygo.

7.2.4 New Transmit-Side Thermal Tuning

Within the Acacia system, another critically important feature is the thermal tuning block. The thermal tuner ensures resonance lock of the microring's Lorentzian features. Traditional thermal tuning blocks required an external drop port on the ring, which acts as a small tap outputting a fraction of the ring's internal power. This fractional power was then low-pass filtered and decoded in order to back-calculate the resonance of the microring. However, this traditional approach not only wasted valuable internal ring resonance power, but also added to the microring footprint and made routing and floorplanning difficult for the necessary peripheral blocks.

The Acacia system utilized a different kind of thermal tuning on the transmit end. In this scheme, as shown in Figure 7.9, a drop-port is not necessary and the thermal tuning AFE may be placed and co-planned along with the transmit-side macro.

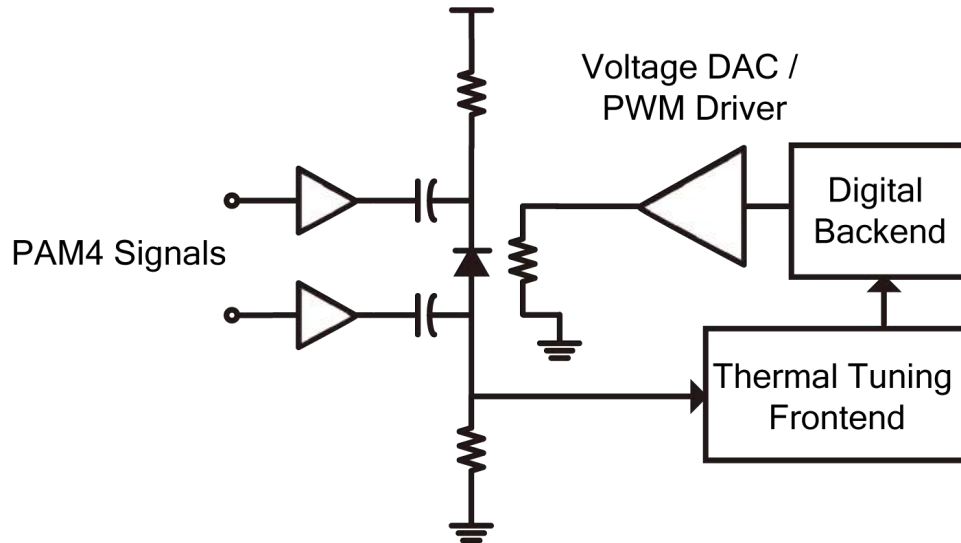


Figure 7.9: The closed loop thermal tuner serves to ensure that the microring resonance remains in lock. This scheme has the added benefit of not requiring an additional drop-port.

The scheme relies on observing the nominal current drawn by the microring's main port (which is driven by the Tx AC driver) and averaging its value in order to back calculate the resonance shift. This current draw, detected using the *Thermal Tuning Frontend* labelled block in Figure 7.9, is then fed into the *Digital Backend*. The *Digital Backend* is then responsible for calculating the effective number of 1s and 0s, as based on the Bit-Statistics tuning scheme. The output of the backend controls the PWM heater driver, which necessarily heats the ring based on the detected offset in resonance.

7.2.5 Tx-Rx Self Test Setup

In order to efficiently characterize the Rx AFE while having the abilities to control the dominant performance characteristics, a self-test-based setup was used. The full self test schematic is shown in Figure 7.10.

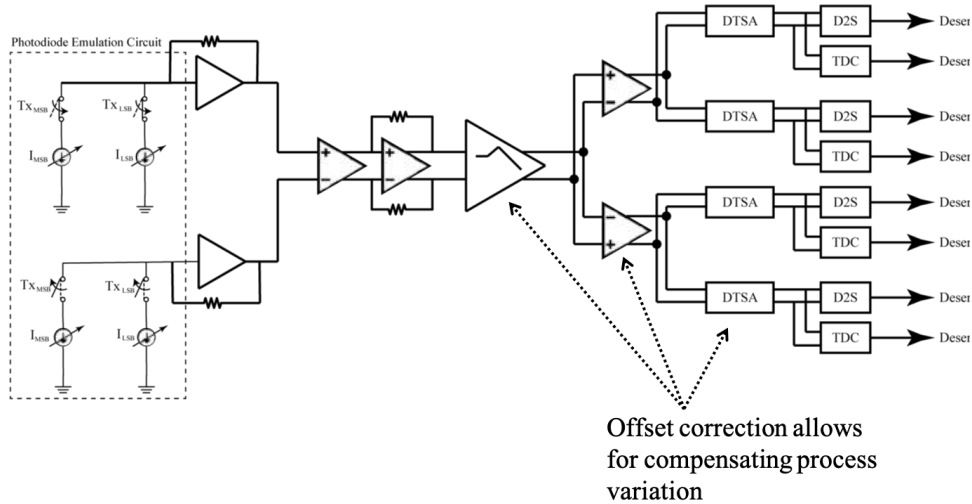


Figure 7.10: The self test schematic shows the full data path to characterize the AFEs.

The schematic shows the Rx AFE being driven by a Photodiode Emulation Circuit (PDE). The PDE enables the designer fine control over the input magnitude of the current, giving independent tuning abilities to both the LSB and MSB currents in each leg. The input to the PDE is driven by the high speed serialized Tx outputs. The output of the Tx AFE is routed with minimal distance to ensure as small a parasitic capacitance as possible. In addition, both the MSB and LSB traces are matched and symmetric to ensure minimal phase delays between the two signals.

Aside from the signal path characteristics, the Rx AFE has numerous knobs to control and compensate for offset. In the CTLE, differential amplifiers, and the double tail sense amplifiers, DAC-controlled offset calibration exists just for this sole purpose. Indeed, the current magnitude DACs in the PDE for both the LSB and MSB can be used for offset calibration as well. Within the DTSA, offset calibration in the form of fine control (using capacitive DACs) and coarse control (current DACs) exist to give a fine resolution and large dynamic correction range.

The output of the self test AFE sites proceed into the deserializers and finally into the digital backend. Within the backend, the bit stream gets verified using a bit-error-rate checker.

7.2.6 Putting together the Acacia System

Putting together the system objectives along with the listed constraints, the top level GDS in Figure 7.11 was produced.

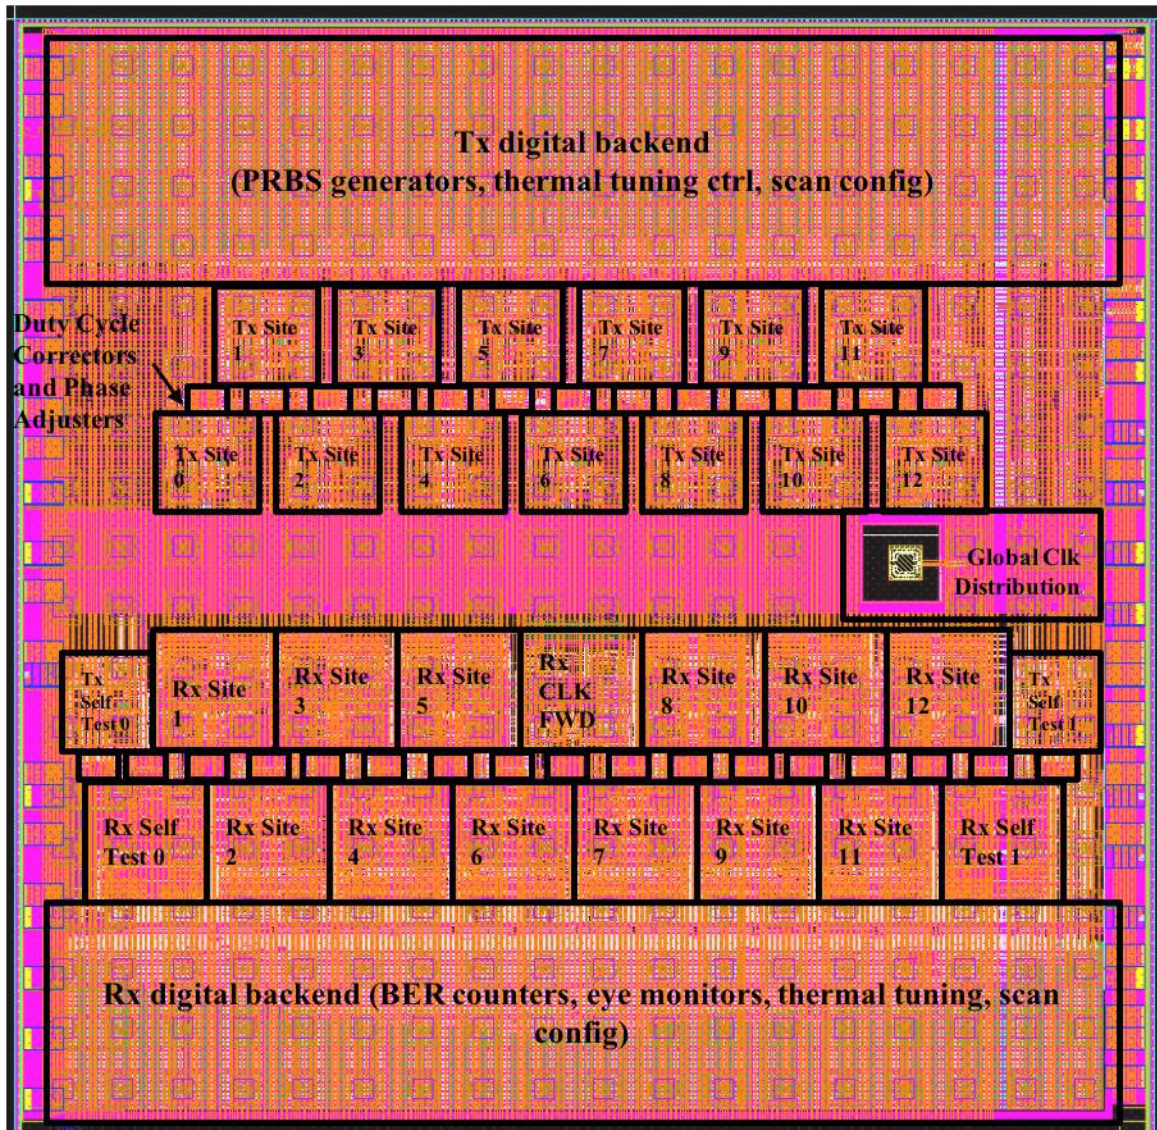


Figure 7.11: The full Acacia system realizes a high speed, high performance end-to-end optical link with all necessary critical peripherals within the system.

A list of the unique features of this CMOS system is shown below as summary:

- A new, single-comparator based PAM4 receiver that is completely auto-generated using the BAG framework

- A new, thermal tuning scheme that relies on the data port and does not require a separate drop-port filter for ring stabilization.
- Full, end-to-end characterization and supporting backends of the links (i.e. PRBS generators, BER checkers)
- On chip oscillator and clock distribution network for the high speed macros
- Duty cycle correctors and phase adjusters per front-end macro
- On chip heater drivers for ring thermal tuning

The self test sites are located in the left and right sides of the macro. The duty cycle adjusters and phase shifters are located in the smaller blocks in between the rows of Rx sites. Each AFE has an appropriately designated Rx duty cycle corrector block. Additionally, the global clock distribution site, which contains an LO as well as the ability to bypass the LO generation with an external clock, is located in between the Tx and Rx macros. The output of the clock distribution network is fed into the Tx and Rx macros independently. The clocks go through a localized clock distribution network within each macro which buffers and cleans up the clock signal as necessary.

7.3 Results

7.3.1 On-Chip Clock Network

To validate the fidelity of the on chip LO as well as the clock distribution network, a sweep of the DAC configuration codes to adjust the LO clock frequency was initially performed. In this test, the capacitive DAC of the LO output was swept. In doing so, the output frequency of the global clock distribution block was modified. The clock, in turn, propagated through the clock tree, duty cycle correctors, Rx AFE where it was divided down, and then through the flop-synchronization logic prior to entering the backend. Within the backend, a counter was used in order to back-calculate the frequency of the global clock distribution block.

Figure 7.12 shows the output clock frequency of the global clock network. The global clock distribution network contains a divider which is controlled by the global scan configuration network. This divider allows the output clock to either be undivided or divided by 2. Indeed, the nominal clock frequency is targetted to around 10 to 12GHz. However, the maximum permissible speed of the backend (without failure in either the clock distribution network or the backend itself) required the divider ratio be set to 2. Thus, the frequencies shown in Figure 7.12 are between 5.6 and 6.6GHz.

To further characterize the fidelity and quality of the incoming clock signal, a phase noise and jitter measurement was also conducted. In this case, three variants of




Figure 7.12: By modifying the capacitive DAC code, the output frequency of the global clock distribution can be modified according to the plot above.

inputs were studied and characterized. In the first test, a simple phase noise and jitter measurement was done on a direct feed-through path between the external clock source and the phase noise spectrum analyzer. Next, a test wherein the external clock is fed into the Acacia chip, then a divided clock was fed out into the spectrum analyzer was set up. Lastly, an “unlocked” LO was activated and the output clock frequency (as found by measuring the divided clock path) was measured and characterized. In this last test, the “unlocked” refers to a free-running oscillator test set up. As such, the jitter measurement takes into account the frequency drift of the clock signal.

To summarize these results, the result of the first test case shows about 130fs of accrued jitter. The external clock path case (test 2) showed about 30ps. Lastly, the LO activated clock jitter was 140ps. Notice that the jitter takes into account not only the clock source (either external or LO) jitter but also the divider and pad driver jitters. Additionally, the free running nature of the LO causes frequency shifts to be seen as jitter noise. Thus, when locked, the effective jitter will be far less for the LO case.

7.3.2 Rx AFE and Self Test Characterization

Rx AFE signal fidelity characterization was done using the photodiode emulation circuit (PDE) detailed in the last section. In this case, the optimal PDE current as well as the offset for the AFE signal path were adjusted manually. Once adjusted, a BER measurement was taken using the snapshot outputs of the Rx digital backend.

To characterize the MSB data path alone (i.e. operate the receiver in NRZ mode initially), the offset codes were adjusted in such a way as to yield the edge of a bath tub curve for low data rates. Figure 7.13 shows the low frequency sweep. In this case, the input current from the PDE was fixed at $20\mu\text{A}$ peak-to-peak.




Figure 7.13: The link's performance is shown for low frequency and operating in MSB-mode.

The initial test to validate the TDC-based receiver was conducted at 10Gbps (collectively for both MSB and LSB streams). In this test, the MSB path was programmed in such a way as to yield a 10^{-5} BER (as shown by the vertical dashed line in Figure 7.13). The offsets and minimum PDE current were both adjusted in such a way as to yield the target BER at 5Gbps (for MSB alone). Then, the phase offset of the TDC-parked clock was adjusted.

The offset code, as can be seen by Figure 7.14, sweeps the location of Φ with respect to the data path. Notice that if the clock is parked before the opening of the TDC eye, the resulting BER will show garBAGe. Similarly, if the clock is parked after the eye opening, the BER will be garage again. However, in between this area, when the clock is parked in such a way as to distinguish between the slow and fast




Figure 7.14: The link’s performance is shown for low frequency and operating in MSB-mode.

evaluations, the BER of the LSB data path shows a working link. Thus, at this point, the BER of the LSB data path easily hits the 10^{-5} BER.

7.3.3 25.6Gbps Self Test Link

A 6.4GHz clock was fed into the Acacia System’s Rx macro backend, with the goal of creating a 25.6Gbps link. In this circumstance, the on-chip LO was activated and programmed to output a 6.4GHz clock. The PDE codes were set in such a way as to produce a peak to peak current of $90\mu\text{A}$. From here, the bathtub curves were measured based on the snapshot outputs produced by the digital backend. The resulting bathtubs are shown in Figure 7.15.

Figure 7.15 shows individually the bathtub curves of both the LSB and MSB slices. In this case, the floor of the BER was limited due to the rate of snapshot updating as opposed to any fundamental performance limit imposed by the AFE itself. Moreover, the bathtub shows flat openings for both these eyes for a substantial range of the unit interval. The asymmetry in the LSB BER curve is partially due to lack of correction abilities in both the clock phase of the on chip phase adjusters as well as due to an offset correction bug which limited the maximum correction ability. In spite of these issues, the two curves still show promising openings.

7.4 Conclusion

The new single comparator PAM4 receiver was tested and designed in the Acacia SoC. This chapter summarizes the various physical design trade-offs when designing and instantiating this architecture. Aside from having just the receiver, a PAM4 transmitter as well as a new drop-port less thermal tuner was also implemented. From a receiver standpoint, the various design trade-offs in the architecture were closely




Figure 7.15: The bathtub curve of the Acacia link operating at 25.6Gbps is shown.

studied and their impact on the Rx AFE was formalized. While testing the PAM4 receiver, a self-test based architecture was used to quantify the link performance. The TDC-based operation was verified by sweeping the delay of the parked clock to the TDC flop. Additionally, the link was characterized for 25.6Gbps operation and showed promising bathtub openings for both the MSB and LSB slices.

Chapter 8

Conclusion

8.1 Thesis Contributions

Silicon photonic links are extremely complicated and rely on a tight overseeing in every level of the hierarchy, beginning at the device level, going into the circuit sub-block domain, climbing up into the macro system level composed of the backend and infrastructure, and proceeding higher into the system integration level and beyond. This dissertation focused on understanding the bridge between these various levels in hierarchy and connecting it to the optimal performance for silicon photonic links.

Chapter 3 introduced and analyzed the optical modeling framework, beginning with a high-level link picture and slowly delving into the various sub-components. The theory and “interface” between these various blocks was also studied. Once this foundation was laid out, the focus shifted to the link-level, where macro-parameters were derived and trade-offs studied. Lastly, using the framework, performance projections were made which, in turn, provided insight into the direction of possible fabrication improvements moving forward.

To further realize and validate the model methodology, Chapter 4 focused on optimizing circuit topologies given unique system parameters. Said differently, the main objective of Chapter 4 was to propose and design circuits and systems under realistic PDK and technology specifications. In addition, based on these unique technologies, the end-to-end link performance was calculated and studied further. The framework was applied to a 65nm, heterogenously integrated CMOS plus photonics platform. The resulting schematic results were also showcased. Furthermore, the swing and noise-dominated regimes of operation were introduced and specific design points in these spaces were used to create schematics. From this all, a framework was in place which could then be used to create actual silicon designs, or intuit and engineer techniques for better circuit optimization.

Chapter 5 detailed the analysis and design of a full, end-to-end 5Gbps NRZ link in the first-ever through-oxide va-based integration technique. The chapter also described the integration platform and showcased the top-level design for the

Electronic-Photonic Integration (EPHI) system, which was taped out and published in ESSCIRC 2015. Subsequent sub-blocks were then analyzed further.

In Chapter 6, the trade-offs of the PAM4 versus NRZ link within the context of the link framework were showcased. Conclusions were drawn as to exactly when the traditional PAM4 architecture was superior to the NRZ architecture. Next, the single-comparator PAM4 architecture was introduced and proved its superiority over the traditional architecture. By targetting and eliminating two extra comparators per way, the new architecture provided a much more efficient AFE design. This claim was verified in the context of the link framework from Chapter 3. The architecture's performance was further analyzed from a link-level standpoint with a photonic and CMOS co-simulation.

Finally, in Chapter 7, the physical design of the single-comparator PAM4 architecture drove the primary motivatoin for the Acacia SoC. In the Acacia SoC, not only was this technique validated, but new control schemes and high performance transmit circuitry was also introduced. Additionally, an in depth look at the physical design trade-offs which transcended hierarchy showed the value in having a rapid design flow in order to quickly iterate and explore the space.

8.2 Future Work and Final Thoughts

As has been reiterated throughout this document (and throughout my PhD), silicon photonics has indeed stepped up as a clear contender in enhancing the capabilities of CMOS technologies. For the past five years, I have been studying and playing within this landscape and to understand these hierarchical trade-offs enough to project better performance and also provide designer's insights to reduce the pain of taping out and testing these massively complex systems. However, it is clear from the conclusion of my dissertation that although many questions were answered and design approaches were made concrete, the "show still goes on". There are an uncountable number of questions which emerged from my work and even more foundations for dissertations.

To enumerate a few possible future directions, we can begin by understanding the context and bounds of the existing link framework. Indeed, in this dissertation, we primarily focused on single- λ links and optimizing end-to-end performance that takes into account both the transmit and receive sides. However, a logical next step of this framework is to study and optimize end-to-end WDM links. WDM links have the benefit of realizing very high bandwidth densities. For the context of this work, we presumed that simply cascading single- λ links impose minimal secondary effects on adjacent channels. However, formalizing these trade-offs within the link framework will truly promise the best end-to-end link performance while potentially exposing critical performance bottlenecks. Factors like resonance spacing, thermal fluctuations, and many more can be taken into account to further this research.

Another interesting avenue stems from utilizing existing machine learning frameworks to truly optimize and close the design loop. Because BAG gives designers rapid access to generating new layouts and running/parsing post-extracted results, one may find benefit in generating BAG-based training sets and applying machine learning techniques to optimize the design. Although existing work in this area uses machine learning as a glorified optimizer, it is foreseeable that with proper circuit abstraction, true “learning” may take place and new topologies or techniques may emerge.

Lastly, the definition of “end-to-end” may indeed be brought into question. Truly end-to-end might imply going from theory, to design, to post-extracted results, to even physical characterization. One may indeed aim to bring in physical experimental characterization into the loop. Imagine that, at the push of a button, a designer is able to generate not only a schematic and layout, but also FPGA-bit files and a list of test equipment necessary to characterize the physical chip. Additionally, an all encompassing end-to-end characterization framework may take into account the FPGA setup, clock source jitter, etc. in order to characterize (pre-tape out) the expected behavior from the chip. Obviously, this may be used to close the loop and modify schematic designs accordingly.

These are just a few examples of the many research paths that this work opens the door to. Thankfully, that chapter in my life has come to a close and I leave it to the reader of this work to make steadfast progress in one of the many directions made apparent throughout this dissertation.

Bibliography

- [1] Settaluri, Krishna T., et al. "First principles optimization of opto-electronic communication links." *IEEE Transactions on Circuits and Systems I: Regular Papers* 64.5 (2017): 1270-1283.
- [2] Andrade, Nicolas M., et al. "Optical Antenna NanoLED Based Interconnect Design." *2018 IEEE Photonics Conference (IPC)*. IEEE, 2018.
- [3] Settaluri, Krishna T., et al. "Demonstration of an optical chip-to-chip link in a 3D integrated electronic-photonic platform." *European Solid-State Circuits Conference (ESSCIRC), ESSCIRC 2015-41st*. IEEE, 2015.
- [4] Chang, Eric, et al. "An Automated SerDes Frontend Generator Verified with a 16NM Instance Achieving 15 GB/S at 1.96 PJ/Bit." *2018 IEEE Symposium on VLSI Circuits*. IEEE, 2018.
- [5] Krishnamoorthy, Ashok V., et al. "Computer systems based on silicon photonic interconnects." *Proceedings of the IEEE* 97.7 (2009): 1337-1361.
- [6] Miller, David Device Requirements for Optical Interconnects to Silicon Chips, *Proc. IEEE* 97, 1166 - 1185 (2009)
- [7] Georgas, Michael, et al. "Addressing link-level design tradeoffs for integrated photonic interconnects." *Custom Integrated Circuits Conference (CICC), 2011 IEEE*. IEEE, 2011.
- [8] Li, J., et al. "Scaling Trends for Picojoule-per-Bit WDM Photonic Interconnects in CMOS SOI and FinFET Processes."
- [9] Jung, Kwangmo, Yue Lu, and Elad Alon. "Power analysis and optimization for high-speed I/O transceivers." *Circuits and Systems (MWSCAS), 2011 IEEE 54th International Midwest Symposium on*. IEEE, 2011.
- [10] Zheng, Xuezhe, et al. "Ultra-low power arrayed CMOS silicon photonic transceivers for an 80 Gbps WDM optical link." *National Fiber Optic Engineers Conference*. Optical Society of America, 2011.

- [11] Proesel, Jonathan, Clint Schow, and Alexander Rylyakov. "25Gb/s 3.6 pJ/b and 15Gb/s 1.37 pJ/b VCSEL-based optical links in 90nm CMOS." 2012 IEEE International Solid-State Circuits Conference. IEEE, 2012.
- [12] Li, Cheng, et al. "A ring-resonator-based silicon photonics transceiver with bias-based wavelength stabilization and adaptive-power-sensitivity receiver." 2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers. IEEE, 2013.
- [13] Ye, Yaoyao, et al. "3D optical networks-on-chip (NoC) for multiprocessor systems-on-chip (MPSoC)." 3D System Integration, 2009. 3DIC 2009. IEEE International Conference on. IEEE, 2009.
- [14] Rylyakov, Alexander, et al. "A 40-Gb/s, 850-nm, VCSEL-based full optical link." Optical Fiber Communication Conference. Optical Society of America, 2012.
- [15] Schow, Clint L., et al. "25-Gb/s 6.5-pJ/bit 90-nm CMOS-driven multimode optical link." IEEE Photonics Technology Letters 24.10 (2012): 824-826.
- [16] Pan, Huapu, et al. "High-speed receiver based on waveguide germanium photodetector wire-bonded to 90nm SOI CMOS amplifier." Optics express 20.16 (2012): 18145-18155.
- [17] Wu, Xiaotie, et al. "A 20Gb/s NRZ/PAM-4 1V transmitter in 40nm CMOS driving a Si-photonics modulator in 0.13m CMOS." 2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers. IEEE, 2013.
- [18] Settaluri, Krishna T., et al. "Demonstration of an optical chip-to-chip link in a 3D integrated electronic-photonics platform." European Solid-State Circuits Conference (ESSCIRC), ESSCIRC 2015-41st. IEEE, 2015.
- [19] Razavi, Behzad. "Design of Integrated Circuits for Optical Communications", Mc Graw Hill, 2003
- [20] Razavi, Behzad. "The StrongARM Latch [A Circuit for All Seasons]." Solid-State Circuits Magazine, IEEE 7.2 (2015): 12-17.
- [21] Nikolic, Borivoje, et al. "Improved sense-amplifier-based flip-flop: Design and measurements." Solid-State Circuits, IEEE Journal of 35.6 (2000): 876-884.
- [22] Personick, Stewart D. "Receiver design for digital fiber optic communication systems, I." Bell system technical journal 52.6 (1973): 843-874.
- [23] J. Kim, B. S. Leibowitz, J. Ren, and C. J. Madden, Simulation and Analysis of Random Decision Errors in Clocked Comparators, vol. 56, no. 8, pp. 18441857, 2009.

- [24] Nuzzo, Pierluigi, et al. "Noise analysis of regenerative comparators for reconfigurable ADC architectures." *IEEE Transactions on Circuits and Systems I: Regular Papers* 55.6 (2008): 1441-1454.
- [25] Kinget, Peter R. "Device mismatch and tradeoffs in the design of analog circuits." *IEEE Journal of Solid-State Circuits* 40.6 (2005): 1212-1224.
- [26] Timurdogan, Erman, et al. "An ultra low power 3D integrated intra-chip silicon electronic-photonic link." *Optical Fiber Communication Conference*. Optical Society of America, 2015.
- [27] Timurdogan, Erman, et al. "An ultralow power athermal silicon modulator." *Nature communications* 5 (2014).
- [28] Going, Ryan, et al. "Metal-optic cavity for a high efficiency sub-fF Germanium photodiode on a Silicon waveguide" *Optics Express*, vol .21, Issue 19, 2013
- [29] Lin, Sen. *Electronic-Photonic Co-Design of Silicon Photonic Interconnects*. Diss. UC Berkeley, 2017.