Optical Receiver Techniques for Integrated Photonic Links

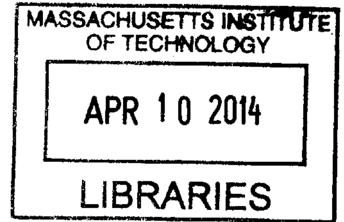
ARCHIVES
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
APR 10 2014
LIBRARIES

by

Michael S. Georgas

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

Doctor of Philosophy in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

February 2014

© Massachusetts Institute of Technology 2014. All rights reserved.

Author ....
Department of Electrical Engineering and Computer Science
November 8, 2013

Certified by ....
Vladimir Stojanović
Thesis Supervisor

Accepted by ....
Leslie A. Kolodziejski
Chair, Department Committee on Graduate Theses

Optical Receiver Techniques for Integrated Photonic Links

by

Michael S. Georgas

Submitted to the Department of Electrical Engineering and Computer Science
on November 8, 2013, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

Integrated photonics has emerged as an I/O technology set to disrupt the communication fabric of many-core computer systems. The optical technology uses wavelength-division-multiplexing and a high degree of integration in order to surpass electrical I/O in both throughput and energy by more than an order of magnitude. However, integrated photonic systems need to be properly designed in order to reach their full potential, and electronic design techniques need to be updated to take full advantage of tight integration with photonics. This thesis explores the engineering of integrated photonics systems, focusing on optical data receiver design and techniques. We develop a representation of a photonic communication system based on circuit and device models, and perform a system-level optimization to find the optimal operating point for each of the components. This operating point sets the specification for the receiver circuits developed. An equivalent model of the receiver is used to develop a host of split-photodiode topologies that enable new ways to implement double-data-rate and decision-feedback-equalization operation. The receivers are fabricated as part of an integrated photonic test platform, and the measurements are detailed. This work presents the first-ever monolithically-integrated optical receiver with μA-sensitivity in a zero-foundry-change, commercial SOI CMOS logic process, and the first-ever monolithically-integrated optical receiver in a bulk CMOS memory process as part of a 9x5Gb/s DWDM receiver bank.

Thesis Supervisor: Vladimir Stojanović

Acknowledgments

The past six years have been a long journey, and there are many people that I would like to thank for their contributions to my experience at MIT.

First and foremost, I would like to thank my advisor, Prof. Vladimir Stojanović. He is a great researcher, mentor, and leader, and from my arrival in the group has always been a model to follow. Prof. Stojanović not only taught me how to conduct disruptive research, think logically, and tackle complex systems, but he also pushed me to work hard with determination and perseverance - much-needed in the photonics project. I'm very much looking forward to applying the skillset that he helped me to develop in the next phase of my career, and I won't forget where that skillset was built.

I would also like to thank my thesis committee for providing guidance and feedback. This thesis is much stronger as a result. Prof. Rajeev Ram helped me tremendously in understanding optical phenomena and how they impact integrated systems. Prof. Anantha Chandrakasan taught me about digital systems early in my graduate studies, which helped greatly with the implementation of the low-power receivers in this work.

Across all of the schools and groups, I would like to thank the entire integrated photonics team. I spent many days and nights designing and working with you all - oftentimes under the most extremes of pressures and deadlines. On the electrical side, I want to thank my fellow ISG members: Ben Moss, Chen Sun, Yu-Hsin Chen, and Jonathan Leu. Together we put more transistors on more chips than I care to count. There's no other group I'd rather have built those chips with. On the optical side, I want to especially thank Jason Orcutt. Everything I know about the optical table, fibers and coupling I learned from him. I also want to thank the team in Colorado. Prof. Milos Popovic, Jeff Shainline, and Mark Wade have always been amazing researchers and friends. I'd also like to thank our collaborators at Micron Technology for helping us to develop the optical memory portion of this thesis.

The Integrated Systems Group was a great team to be a part of, and I want to

thank all of the members past and present that I got a chance to work with. Fred Chen and Ranko Sredojevic are not only great friends, but also great engineers that I learned a great deal from. I'm very lucky to have sat next to Yan Li for my entire graduate career - I enjoyed the many off-topic conversations. I'd also like to thank all of the other friends I've made these past few years for their support during the tough times: Patrick Mercier, Marcus Yip, George Hwang, Jackie Hu, Jason Chang (who is also my finance partner), John Sun, Sushmit Goswami, and many more. I want to give a special thanks to the MIT Men's Hockey team. Playing with them has been one of the most rewarding experiences of my life.

I would like to especially thank my family: Mom, Dad, Deana, and Peter. Your support and help has always been a source of strength for me, and I'm very lucky to have you all behind me.

Most importantly, I want to thank my fiancee Diana, who supports me continuously and unconditionally. Her belief in me is the single strongest reason that I push forward every day.

Support My work over the past four years has been funded by DARPA and NSERC. I thank each group for helping to fund my exploration of this interesting research.

Contents

1 Introduction **19**

1.1 The Emergence of Multi-core Computing 20

1.2 Integrated Photonic System 24

1.3 Photonic Integration Strategies and Receivers 25

1.3.1 Discrete Components 25

1.3.2 Hybrid Integration 28

1.3.3 Monolithic Integration 29

1.3.4 Receiver Performance Summary 31

1.4 Thesis Contributions 32

1.5 Thesis Overview 33

2 Link-level Anlaysis of a Photonic Network **37**

2.1 Photonic Link Components 38

2.1.1 Modulator 38

2.1.2 Optical Data Receiver 42

2.1.3 Optical Clock Receiver 47

2.1.4 Single Channel Link Tradeoffs 48

2.2 Towards a Full WDM Link 50

2.2.1 Optical Clock Distribution 51

2.2.2 Ring Resonance Mismatch 51

2.2.3 Ring Tuning Techniques 53

2.2.4 Ring Tuning Model 53

2.3 WDM Photonic Link Evaluation 55

2.4 Summary 58

3 Optical Data Receiver Design and Techniques **61**

3.1 Photodiode Design 62

3.2 Optical Data Receiver: Latching Sense Amplifier 65

3.2.1 Architecture 66

3.2.2 Effect of Capacitance on Receiver Settling 67

3.2.3 Sensitivity 70

3.3 Optical Data Receiver: Transimpedance Amplifier 71

3.3.1 TIA Front-end 72

3.3.2 Sense Amplifier 74

3.4 GSG Transimpedance Amplifier Test Site 75

3.5 Photodiode Parameter Impact on Receiver Suitability 77

3.6 Photodiode Splitting 78

3.6.1 Double Data Rate 80

3.6.2 Decision Feedback Equalization 80

3.7 Summary 81

4 EOS Photonic Technology Platform **83**

4.1 EOS Chip Architecture 84

4.2 Integrated Electronic-Photonic Test Setup 85

4.3 Enabling Low Loss Waveguides in 45nm SOI: Substrate Removal Electronic Verification 88

4.4 Optical Data Receiver: Logic 90

4.4.1 Low Bandwidth Detection (EOS4) 90

4.4.2 Gigabit Optical Receivers 93

4.4.3 Multi-Gigabit Optical Receivers 97

4.5 Optical Data Receiver: Memory 102

4.5.1 Stand-alone Receiver Performance 102

4.5.2 DWDM Receivers 108

4.6 Integrated Photonic Links 110

4.6.1 Single-λ Link . . . 110
4.6.2 Multi-λ Link . . . 113
4.7 Summary . . . 115

5 Conclusion **117**
5.1 Future Work . . . 118
5.1.1 Roadmap to 1Tb/s . . . 118
5.1.2 Platform Development . . . 120
5.1.3 New Application Areas . . . 121

A Die Photos **123**

List of Figures

1-1 The emergence of multi-core processors. 20

1-2 Electrical transceiver trends. 21

1-3 Off-chip communication costs. 23

1-4 An example optical link with chip-to-chip and intra-chip communication links shown. A CW laser source is coupled onto Chip A through a vertical grating. Two ring-resonant modulators imprint data onto two wavelength-channels, λ_0 and λ_1, which propagate along the waveguide. The bus is routed over an optical fiber to Chip B. The drop rings on Chip B are each tuned to either λ_0 and λ_1 to select that channel from the bus and direct it to the correct data receiver. A second set of wavelengths, λ_2 and λ_3 carry data from Chip B to Chip A. 24

1-5 Discrete-component optical receivers. 27

1-6 Hybrid-integration optical data receivers. 29

1-7 Monolithically-integrated optical data receivers. 30

1-8 Optical transceiver trends. 32

2-1 An integrated WDM photonic link. A continuous-wave (CW) multi-λ laser is coupled onto the chip through a vertical-coupling grating structure. Once on chip, frequency selective ring-resonant modulators encode digital bitstreams onto their resonant wavelengths. Each wavelength propagates along the waveguide (and possibly off-chip) until it is routed through a matching drop ring to an integrated photodiode (PD). An optical receiver forms a bit decision based upon the PD photocurrent. Clock signals are routed both optically along the waveguide and electrically through local H-trees. Ring tuning circuits are used to tune the resonance of the modulator and drop rings. 39
2-2 Device requirements to reach the target $ER_{dB} = 6$dB for various values of IL_{dB}. 41
2-3 Electrical model of modulator device and driver. 42
2-4 Energy costs to reach the target $ER_{dB} = 6$ dB for various values of IL_{dB}. 43
2-5 Resistive receiver sensitivity. 44
2-6 Transimpedance Amplifier. 46
2-7 TIA design example at C_p=25 fF, DR=5 Gb/s. 46
2-8 TIA performance. 47
2-9 Integrating receiver design and performance. 48
2-10 Data rate tradeoffs for a single photonic link for 4 integration scenarios. C_P=5fF represents monolithic integration, while C_P=25fF is expected for a TSV connection to an optical die. Channel losses of 10dB and 15dB correspond to on-chip and chip-to-chip links, respectively. . . . 49
2-11 Power vs. datarate per channel, jitter is fixed to be within 3% UI, and C_{PD}=5fF 52
2-12 Tuning power vs process variation at various channelizations for the full thermal tuning scenario. 54
2-13 Tuning power vs process variation at various channelizations with an electrical backend capable of bit reshuffling. 54

2-14 Tuning power vs process variation at various channelizations with a bit reshuffling electrical backend and electrical assisted tuning. 55

2-15 Optimized power vs. data-rate for different aggregate link throughputs for Loss=10dB, C_P=5fF. For tuning, we assume a bit-reshuffler backend and electrically-assisted tuning with local variation σ_{rL}=40GHz and systematic variation σ_{rS}=200GHz. Note that the number of WDM channels changes with data rate (Channels = Throughput / Data-Rate). 57

2-16 Bandwidth density of integrated photonics compared to electrical. . . 58

3-1 Optical absorption coefficients for Silicon and Germanium. 62

3-2 Photodiode connection options. 63

3-3 Photocurrent as a function of reverse bias, for SiGe detectors in logic process. Measurement performed by Jason Orcutt. 64

3-4 Optical data receiver architecture. The LSA (a) is followed by an output buffer stage (d) and dynamic-to-static converter (e), before being fed into the digital backend infrastructre (f). The chip has receivers connected to integrated PDs (g) or electrical diode-emulation circuits (h). The simulation model is shown in (i). A cross-section of the implemented PD and optica mode is shown in (j). 66

3-5 LSA sensitivity model. 68

3-6 LSA sensitivity model. 69

3-7 Receiver circuit noise CDF. 71

3-8 Receiver sensitivity and PD split. C_{PD}=10fF, $C_w = 5fF$. 72

3-9 TIA optical receiver (EOS18). 73

3-10 TIA front-end in 0.180μm bulk memory process. 74

3-11 EOS16 TIA optical receiver. 76

3-12 GSG TIA stage gain and phase (simulation). 76

3-13 GSG TIA performance (simulation). 77

3-14 Single-rate, split-diode, and DFE photodiode configurations. 79

4-1 EOS platform architecture (EOS18). 85

4-2 Micron platform architecture (D1L). 86
4-3 EOS platform chip testbench. 87
4-4 Fiber coupling into an EOS chip. 88
4-5 Optical Table 2 and link demonstration setup. 89
4-6 EOS platform 45nm SOI ring oscillator. 90
4-7 EOS platform 45nm SOI stage delay before and after substrate removal and transfer. 91
4-8 Die photo showing photodetector layout options. 92
4-9 The receiver's ability to distinguish between a DC *optical-1* and *optical-0*. 93
4-10 EOS4 Receiver: Measured performance. 94
4-11 EOS4 Receiver: Electrical eye diagram and energy-cost breakdown. . 94
4-12 SiGe photodiode design and measurement (EOS16, measured by Jason Orcutt). 95
4-13 Optical data reception (EOS16). 95
4-14 SiGe photodiode performance (EOS18, measured by Ashwyn Srinivasan). 96
4-15 GSG-testable gigabit TIA in 45nm SOI. 97
4-16 Gigabit receiver eye diagrams in 45nm SOI (EOS18). 98
4-17 EOS18 TIA receiver eye diagrams. 99
4-18 TIA-based receiver performance vs data-rate (EOS18). 100
4-19 Receiver schematic (D1L). 103
4-20 Memory process test setup. 104
4-21 Receiver eye diagrams at 3Gb/s on the Memory chips with fabrication errors (D1L). 105
4-22 Receiver eye diagrams at 2Gb/s on the correctly-processed memory die (D1LL). 105
4-23 Receiver eye diagrams at 5Gb/s on the correctly-processed memory die (D1LL). 106
4-24 Receiver bathtube curves at 5Gb/s. 107
4-25 Receiver energy-cost and sensitivity vs data-rate for the Memory process (D1LL). 108

4-26 DWDM result showing each receiver receiving a different wavelength. 109

4-27 Single-λ Memory-Memory link test setup. 111

4-28 RX link eye diagrams at 2Gb/s from the correctly-processed memory die. 112

4-29 Single-λ Memory-Memory link at 5Gb/s, using an SOA between the TX and RX chips. 113

4-30 Link eye diagrams for WDM configuration. 114

5-1 Roadmap to 1Tb/s. 119

A-1 EOS chip photos. 124

List of Tables

1.1 Optical transceiver performance for various integration options - prior work. 35

2.1 Link Evaluation Parameters 50

2.2 Tuning model evaluation parameters 53

3.1 TIA front-end performance in 45nm SOI (simulation). 73

3.2 TIA front-end performance in 0.180μm bulk (simulation). 74

4.1 Logic-process transceiver performance summary. Optical receiver sensitivity calculated based on PD responsivity of 0.02A/W. 101

4.2 Memory-process receiver performance summary. Receiver optical sensitivity based on photodiode responsivity of 0.2 A/W. 109

Chapter 1

Introduction

As high performance computing continues to advance toward the many-core regime, processors require ever-increasing bandwidth from the on-chip interconnect network and off-chip interfaces. These interconnect fabrics already occupy large portions of the chip area and consume a significant fraction of the total processor power [1]. However, unlike the past we cannot rely on process improvements to help achieve higher data-rates and more efficient links. Projections show poor scaling of on-chip wires and I/O bandwidth density with technology [2,3], as links are becoming channel-limited and not process-limited. This emphasizes the need for a disruptive interconnect technology in order to meet the bandwidth-density demands and energy-efficiency requirements of many-core systems.

Integrated silicon photonics is an interconect technology that demonstrates significant advantages over traditional electrical links for on-chip core-to-core [4,5] and off-chip core-to-DRAM [3,6,7] applications. The technology stands apart from electrical IO in two key ways. Firstly, integrated photonics makes use of dense wavelength-division-multiplexing (DWDM) to send multiple channels of data on a single fiber, each encoded on its own wavelength. Energy-efficiency can be achieved through tight, monolithic integration of the components with the transceiver circuits. Additionally, the energy-cost is relatively distance-insensitive due to low-loss optical cables, making links $100mm$ apart as feasible as links $100m$ apart.

While significant work has been done in the field of silicon photonics to integrate

optical devices into CMOS processes, existing circuit techniques must also be adapted in order to leverage the small device capacitance and availability of a source-forwarded clock in a DWDM monolithically-integrated link. This new integration scenario also demands a system-level study in order to drive the performance for each link component. This thesis focuses on the design of an integrated optical receiver, motivated through such a system level study.

1.1 The Emergence of Multi-core Computing

The recent trend in high-performance computing towards multi-core architectures is shown in Figure 1-1. With the end of frequency scaling, computer architects needed to find a way to keep computational performance increasing, and the answer was found in parallelism. Since the year 2000 the number of cores on chip has doubled roughly once every 18 months. In order for the cores to perform any meaningful task, however, they must be provided with data on which to compute. These bytes of data may come from other cores on the same chip, other processors, or memory.

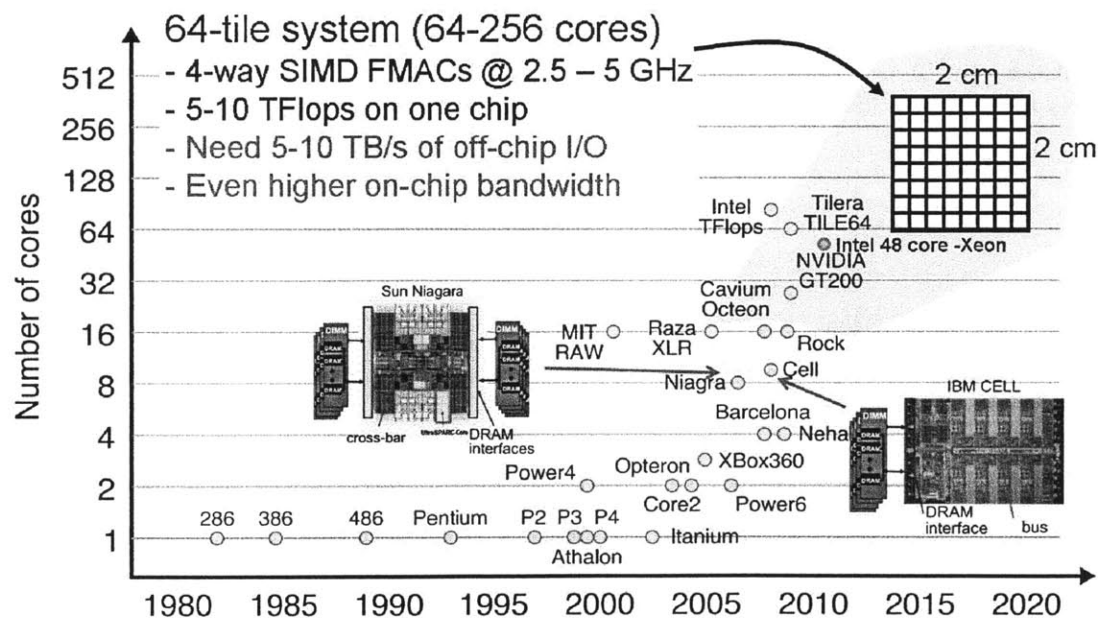

Figure 1-1: The emergence of multi-core processors.

Looking at the energy-cost associated with providing these cores with data, the IO bottleneck emerges. Figure 1-2 shows the recent state-of-the-art in chip I/O, with energy-cost plotted against the link data-rate. While there are works that have approached 1pJ/bit, they have not advanced farther below. Furthermore, we can see that many of the more efficient works are achieved at lower rates, emphasizing the point that it is difficult for electrical I/O to be both energy-efficient and high-bandwidth at the same time. The solid line drawn in Figure 1-2 passes through 10Gb/s and 1pJ/bit, and is intended to help the reader appreciate that higher data-rate designs are less energy-efficient. The line does not represent a fundamental barrier.

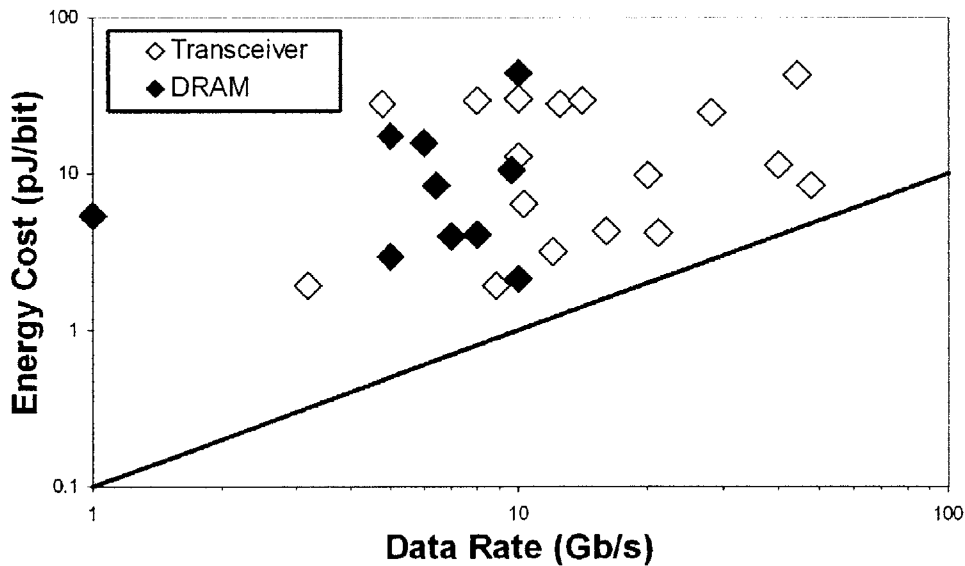

Figure 1-2: Electrical transceiver trends.

It's also important to notice that the DRAM points on the plot all fall below 10Gb/s. This is partly due to the fact that differential signaling is not a feasible solution in a pin-constrained environment [8], as single-ended solutions achieve a higher data-rate-per-pin (despite increased supply noise). In fact, packaging is a major part of the problem in the processor as well, as will be made evident in our photonic system model.

In order to appreciate the severity of the IO problem, we consider projected computation demands and link energy costs. Figure 1-3a shows the computation required per chip in both the server and mobile-client space. From the plot, we can see that a high-performance server will require upwards of 10 TFlops/s by around 2016 [9]. The data for these computations is coming from off-chip DRAM, and the associated memory-access energy-cost is shown in Figure 1-3b. While the energy-cost is decreasing, it is not doing so as quickly as the computation requirements are increasing, further emphasizing the barrier shown in Figure 1-2.

If we assume 1 byte is required per flop, we can multiply the curves of Figure 1-3a and Figure 1-3b in order to obtain the total off-chip I/O power requirements in Figure 1-3c. The figure shows that in the near future, the power required for I/O will exceed the budget for the entire chip in both the server and mobile client space. As a result, we will not be able to take full advantage of emerging many-core systems, and computation capability will cease to advance. This is compounded by the fact that packaging technology is also advancing relatively slowly. Figure 1-3d shows the package pin count projections. In order to satisfy the throughput demands in the server space, packages will require 16,000 pins in 2017, but will only have around 6,000.

In addition to energy-cost, bandwidth density is an important metric for off-chip I/O. Consider the electrical I/O at the die level. A 100 μm C4 bump pitch results in 100 $bumps/mm^2$, of which only half will be available for I/O (assuming the other half are required for power supply). 25 differential links operating at 20 Gb/s each yields a bandwidth density of 500 $Gb/s/mm^2$. At the package level, we assume a pin count around 8000, resulting in 2000 differential I/Os again at 20 Gb/s. A $40mm$ by $40mm$ socket then achieves a bandwidth density of only 25 $Gb/s/mm^2$. On the other hand, at the package-level, photonic interconnects offer bandwidth densities of around 1 Tb/s per fiber, where fibers can be positioned at a roughly 100μm pitch to the chip, achieving bandwidth densities of $\approx 50 Tb/s/mm^2$. In order to understand the constraints for the photonic case at the die-level, system-level modeling must be performed (the reader can skip ahead to Figure 2-16 at the end of Chapter 2 to find

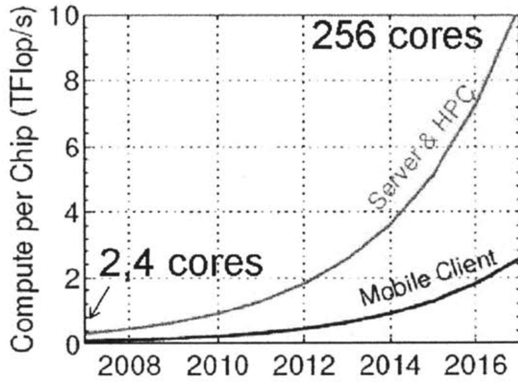

(a) Computation requirements.

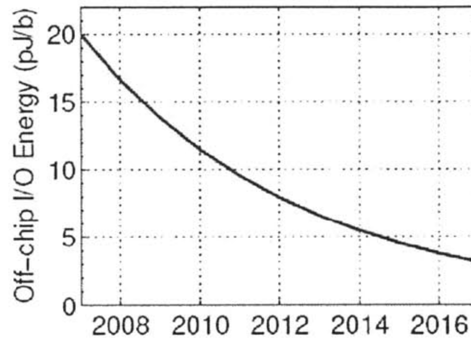

(b) Off-chip IO Energy.

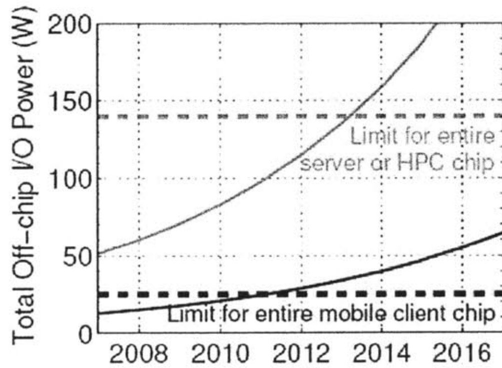

(c) Total Off-chip I/O Power.

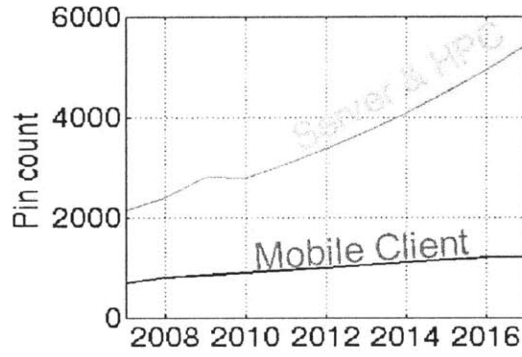

(d) Package pin count.

Figure 1-3: Off-chip communication costs.

out).

To summarize, the problem with electrical I/O is that it can be either energy-efficient enough, or fast enough, but not both. In order to keep computation scaling, we need a disruptive technology to advance the field in both directions at once. Integrated photonic interconnects are such a technology.

1.2 Integrated Photonic System

An example of a DWDM, monolithically-integrated photonic link between two chips is shown in Figure 1-4. This link could, for example, represent the connection of a processor with off-chip DRAM. A continuous wave (CW), multi-λ laser is coupled from an optical fiber onto the chip through a vertical grating coupler. The light is then routed throughout the chip along waveguides fabricated using either gate poly-silicon or the SOI body.

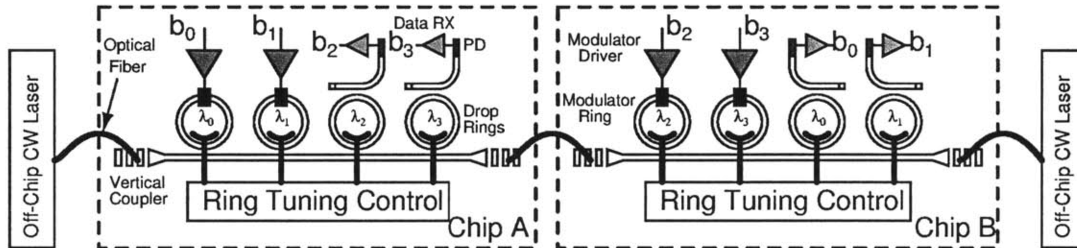

Figure 1-4: An example optical link with chip-to-chip and intra-chip communication links shown. A CW laser source is coupled onto Chip A through a vertical grating. Two ring-resonant modulators imprint data onto two wavelength-channels, λ_0 and λ_1, which propagate along the waveguide. The bus is routed over an optical fiber to Chip B. The drop rings on Chip B are each tuned to either λ_0 and λ_1 to select that channel from the bus and direct it to the correct data receiver. A second set of wavelengths, λ_2 and λ_3 carry data from Chip B to Chip A.

Resonant drop rings form notch filters that can pull a particular wavelength-channel off of the optical bus. This can redirect the wavelength-channel to a different waveguide or be used to modulate a data signal. The modulator leverages the free-carrier-dispersion effect to modulate a P-N junction located around the ring in order to change the ring's refractive index, and therefore its resonant frequency. By tuning

a particular ring's resonance to a wavelength channel, light is confined to the ring and prevented from traveling down the waveguide, yielding an optical-0. De-tuning the ring shifts the notch filter away from the wavelength-channel and light propagates down the waveguide, yielding an optical-1.

The modulated light is routed to another location on the die (e.g. core-to-core) or to another die (e.g. socket-to-socket). At the destination, the ring-tuning control block selects the channel to be removed from the optical bus by setting the resonance of a drop ring filter to the particular wavelength-channel. An optical receiver, such as the one presented in this work, then converts the data back into the electrical domain by detecting the PD photocurrent. An optical clock signal can also be forwarded along with data, which is desirable as little intereference exists in the optical domain.

One of the key characteristics of a monolithically-integrated photonic link is that all of the components are tightly integrated with each other on the same die. As such, power can be optimized at the system level, setting the specification of each component. The system-level design trade-offs are explored in Chapter 2.

1.3 Photonic Integration Strategies and Receivers

In the previous sections, we have motivated some of the potential advantages of integrated photonic interconnects. This section examines some of the work previously done in the field in the context of the different integration strategies adopted. By studying the following works, we can develop an understanding of the benefits and costs of different integration types. In particular we will be examining the implementation of the optical receiver - a key component in an integrated optical link.

1.3.1 Discrete Components

Optical receivers have traditionally been designed as discrete components for telecom applications. By separating the electrical and optical parts of the system into two separate chips, the responsivity (photocurrent-per-unit-optical-power) could be optimized through semiconductor material selection. This is due to the fact that the

designer is free to choose any material available for the optical device - it does not have to be already present in the circuit's process. The discrete integration approach does, however, result in fairly large packaging capacitances, decreasing the sensitivity of the receiver. To mitigate the bandwidth limitation at the circuit's input node, a power-hungry transimpedance amplifier (TIA) can be implemented to drive down the impedance looking into the circuit while preserving a large transimpedance gain.

The TIA presented in [10,11] is shown in Figure 1-5a and targets low-power, short-range applications. The discrete integration of the photodector results in a fairly large capacitance of 320fF, composed of 220fF depletion capacitance for the photodiode and 100fF parasitic capacitance for the input pad. A modifided regulated cascode circuit is described, with peaking inductors used to operate in the low headroom of 80nm CMOS process. We note that the peaking inductors are unsuitable for applications in a many-core processor where the receiver cicuit will be instantiated many hundreds of times and its area cost is closely monitored.

Figure 1-5b shows the receiver block diagram disclosed in [12]. In an effort not to use an amplifier that burns static current, the topology implemented is an integration-based receiver, where the photocurrent is integrated onto a capacitance to produce a voltage different. The voltage difference is then sampled and evaluated using time-interleaved voltage-sense-amplifiers. In this manner, the photocurrent is able to integrate on the capacitance for an entire bit time, and the reset phase of one sampler is hidden in the evaluate phase of the other. The discrete integration results in a large capacitance at the input node of 420fF, though this capacitance does help to mitigate the effects of comparator kick-back and charge-sharing. A current-source draws the recent-average-photocurrent from the the input, keeping the system properly biased. The receiver is able to resolve a 11-μA input current at 1.6-Gb/s with a power consumption of only 3-mW. [15] shows the next generation of this receiver as part of an optical transceiver.

Figure 1-5c shows a another discrete-integration trasceiver. The circuits are implemented in a GP 65nm CMOS process, with the photodiode wire-bonded to the circuits chip. The total input capacitance to the receiver is reported as greater than

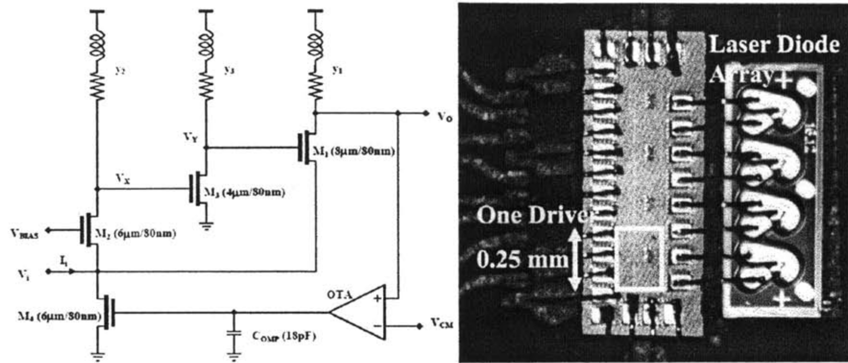

(a) [10,11]

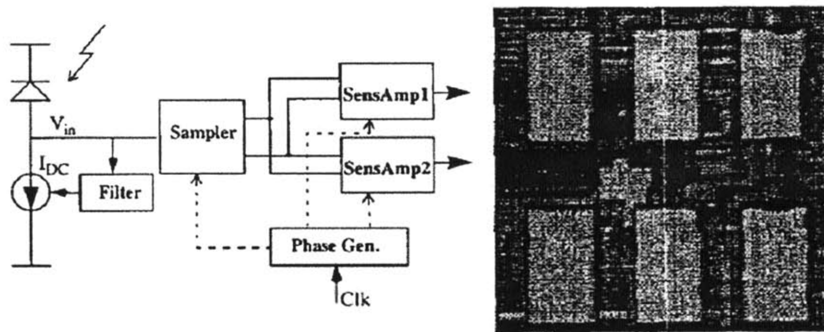

(b) [12]

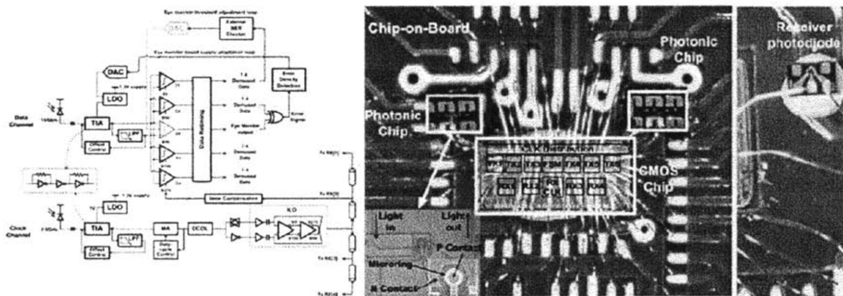

(c) [13]

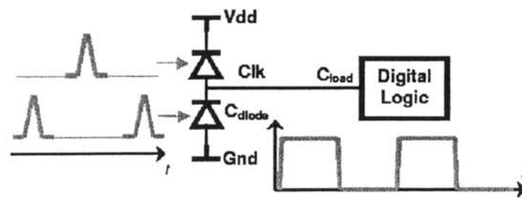

(d) [14]

Figure 1-5: Discrete-component optical receivers.

200fF. The receiver topology here uses a simple inverter-based TIA with a resistor in feedback. Similar to [12], interleaved voltage samplers are used after the front-end. The receiver's energy-cost is 275fJ/bit at 8Gb/s.

The clearest example of an integrating receiver is presented in [14], and is shown in Figure 1-5d. In this *receiver-less* design, two opposite-phase optical clock pulses are propagated to two photodiodes stacked on top of each other. As the optical clock pulses reach their target photodiodes, the photocurrent that is generated is used to charge/discharge the node looking into the digital logic block. The capacitance at this node is composed of the diode's own capacitances, as well as any parasitic capacitances. Assuming that the diodes are well matched, and that there is enough optical power and high-enough conversion efficiency, the input node to the digital logic will have enough range to drive the following logic blocks. The detector was flip-chip bonded to the chip.

1.3.2 Hybrid Integration

More recently, integrated photonics has addressed chip I/O bottlenecks through hybrid-packaged solutions [16, 17]. The decreased photodiode and parasitic capacitances have helped to improve the sensitivity and energy efficiency, but also caused designers to turn back to TIA-based designs that avoid integrator-specific problems such as kick-back. A capacitance of 90 fF is reported in [16]. In [18] a 25fF photodiode capacitance connected through a 20 fF microsolder bump leads to a receiver sensitivity of 9 μA but energy-cost of 690 fJ/bit at 5 Gb/s.

Figure 1-6c shows another example of hybrid-integrated photonics with a receiver. The circuits process used is a 40nm CMOS process, while the photonics are created on a separate die. The photodetector is a Germanium photodiode grown on top of a silicon waveguide and connected to the VLSI chip through fine-pitch solder bumps, with an estimated 60fF capacitance for both the bonds and photodiode. The interfacing circuit consists of a TIA with interleaved voltage sense-amplifiers that make use of a receiver-side clock for energy-efficient regeneration. The data-rate achieved was 10Gb/s at an energy-cost of 0.375pJ/bit and sensitivity of $22.1\mu A$.

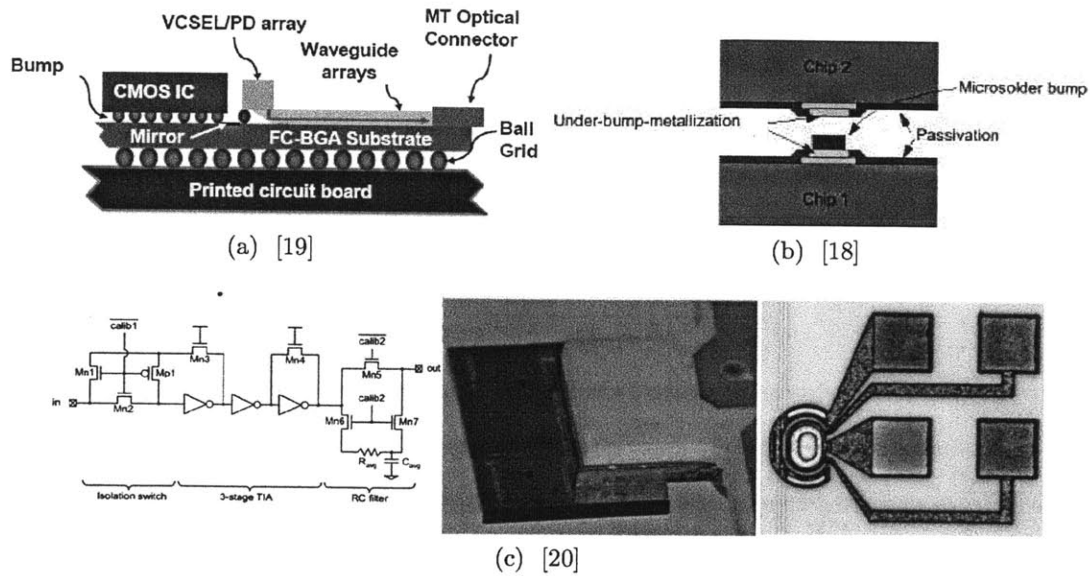

Figure 1-6: Hybrid-integration optical data receivers.

Further capacitance reduction will improve energy-cost and sensitivity, which maps directly to the system's laser power, resulting in designs more competitive with electrical solutions already at 1 pJ/bit [21].

1.3.3 Monolithic Integration

In contrast to large photodiode and parasitic capacitances that cause traditional optical receivers to utilize various power-hungry TIA topologies, monolithic integration offers low PD parasitic capacitances, and in this section we explore recent work in the field.

Figure 1-7a shows the receiver in [22], which uses a Ge-on-SOI detector stated to have a responsivity of 0.35 A/W and intrinsic capacitance of 30-fF. The TIA is a differential, common-gate topology, that also uses peaking inductors and is followed by a 5-stage limiting amplifier, and then a buffer to drive a 50-Ω load. The complete receiver operates at 12-Gb/s, and the authors note the low photodiode biases used which will help in implementation in a complex digital system.

The transceiver in [23] (shown in Figure 1-7b) is implemented in a 130nm CMOS

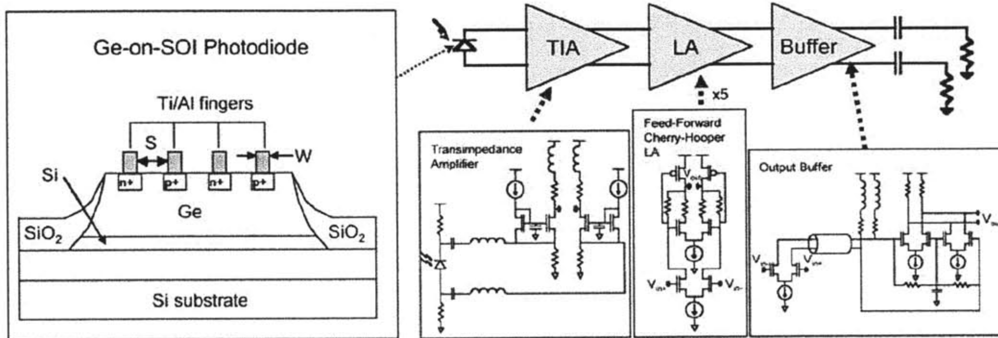

(a) [22]

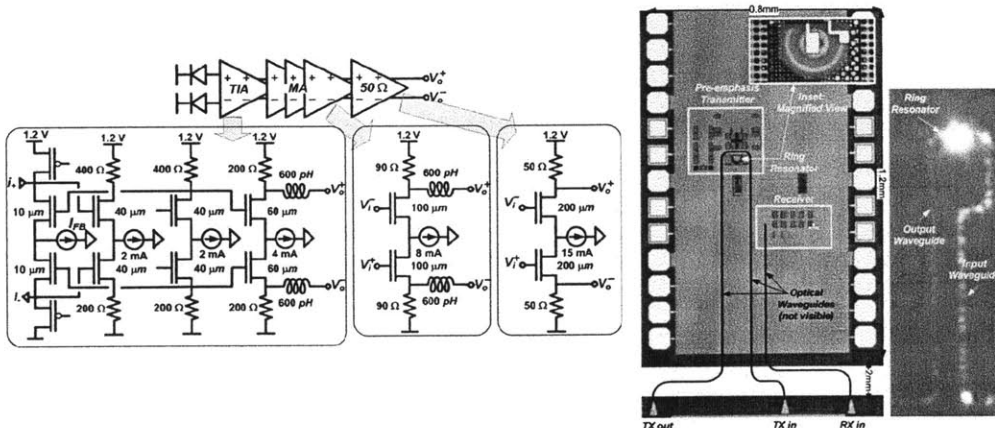

(b) [23]

Figure 1-7: Monolithically-integrated optical data receivers.

SOI process. Again we see the use of a TIA front-end in the receiver, which achieves a sensitivity of $180\mu A$ at 25Gb/s with an input-node capacitance 20fF. The energy-cost of the receiver is quite large at 1.2pJ/bit, due to the use of an expensive and complex TIA. This makes [23] more suitable for transceiver applications than the VLSI applications we are focused on in this thesis.

While the monolithic examples presented succeed in achieving low device capacitance, they are not identical to the work presented in this thesis. In Chapter 4 we demontrate our monolithically-integrated photonic platform in two processes: an SOI process and a bulk process. Unlike the thick-BOX (greater than $1\mu m$) SOI processes presented here, we make use of a thin-BOX SOI process suitable for high-performance computing. Additionally, our work requires zero process changes at the foundry. While the the bulk CMOS work presented does require foundry changes, to our knowledge it is the first photonic platform enabled in bulk CMOS to date.

1.3.4 Receiver Performance Summary

Having looked at prior optical receiver art in the previous subsections, we briefly look at the trend. Table 1.1 shows transceiver performance for several works in similar technology, and compares them to the performance achieved in this thesis.

Figure 1-8 plots transceiver energy-cost against data-rate (similar to Figure 1-2) for recent optical works. From the plot we note that several works have already broken past the barrier below which no electrical links penetrated. While we are primarily interested in monotlithic integration in this thesis, the best-performing transceivers (in terms of both data-rate and energy-efficiency) have been hybrid transceivers. An important reason for this is that hybrid transceivers pre-date their monolithic counterparts, and many of the monolithic results to date (including the one we disclose in this work), are very early-generation works. Large performance gains can be achieved during this development phase, and it will be important to monitor Figure 1-8 as the field continues to mature. Only then will a fair comparison between hybrid and monolithic approaches be possible.

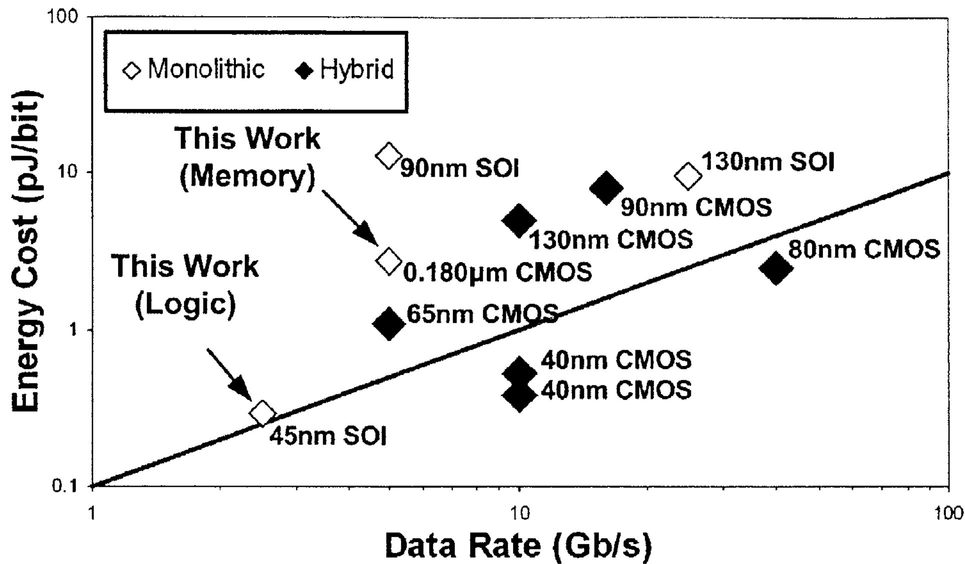

Figure 1-8: Optical transceiver trends.

1.4 Thesis Contributions

This thesis develops circuit designs and techniques for optical receivers in integrated photonic interconnect systems. The main contributions fall into two categories: integrated photonic system modeling, and optical receiver circuit design and validation.

The main contributions of this thesis are:

- A study of technology trends in order to motivate the need for integrated photonic links as a disruptive technology for next-generation IO.
- A system-level analysis of integrated photonic systems, built upon compact models of each of the circuit blocks and components.
- The development of receiver circuits tailored for implementation in a monolithically-integrated phtoonic system.
- Split-photodiode techniques that enable double-data-rate and decision-feedback-equalization topologies previously unrealizable.

- Development of a monolithically-integrated photonic technology platorm and experimental validation of the receiver designs and techniques:
  - First ever optical receiver and platform in a zero-foundry-change, commercial SOI CMOS process. Receivers are shown to operate with μA-sensitivity at multi-gigabit rates with very low energy-cost.
  - First ever optical receiver and platform in a bulk CMOS process. Receivers are implemented as part of a DWDM bank capable of 9x5Gb/s.

1.5 Thesis Overview

The technical portion of this thesis is divided into three main parts. In Chapter 2 we examine and model an integrated photonic link system. The system is broken down into each of its main components (modulators, receivers, photodiodes, drivers, thermal tuning blocks, and waveguides), and a compact model for the components is developed. The models are then simulated and optimized in order to gain intuition on system-level configuration and performance. The chapter concludes with a comparison between integrated photonics and electrical I/O.

Based on the insight achieved in Chapter 2, we develop three receiver circuits suitable for integrated photonics in Chapter 3. The first receiver is a current-sense-amp that makes use of a source-forwarded clock to implement a regenerative latch. The design is aggressive, digital, compact, and has a very low energy-cost. In order to gain insight into the effect of photodiode and wiring capacitance on the receiver, we develop a simple equivalent-circuit model. The second receiver presented is based on a more traditional transimpedance amplifier (TIA). The TIA-based receiver consists of a TIA front-end followed by a voltage-sense-amplifier that again makes use of a receiver-side clock for energy-efficient decision resolution. Finally, the third receiver presented is a simple TIA followed by limiting amplifiers. This design is presented as an easily-testable "canary in a coal mine" that is implemented on our photonic device development chips, and is used to efficiently test a wide array of

photodiodes using GSG and DC probes. Chapter 3 concludes with a discussion of several novel split-photodiode techniques that enable both double-data-rate (DDR) and decision-feedback-equalization (DFE). The split-photodiode is motivated based on the receiver's equivalent circuit model.

With both a system-level understanding and receiver circuits in place, Chapter 4 presents the EOS integrated photonic technology platform, which is the vehicle through which we implement the photonic devices and circuits. As this work is focused on processor-memory communication, two main families of chips are presented: logic chips that are implemented in a zero-foundry-change commercial 45nm SOI CMOS processor, and memory chips that are implemented in a $0.180\mu m$ bulk CMOS memory periphery process. The architecture of the two chips is detailed, showing what the potential capabilities of the chips are. The receiver circuits are then characterized, with the ultimate result showing integrated receivers in each of the two processes capable of energy-efficient, multi-gigabit operation. The chapter concludes with single-link and WDM results demonstrating the viability of the receivers in a complete link system.

	This Work (Logic)	**This Work (Memory)**	[24]	[23]	[25]	[26]
Technology Photonics	45nm SOI Monolithic	0.180μm Bulk Monolithic	90/130nm SOI Monolithic	130nm SOI Monolithic	40nm Hybrid	65nm Hybrid
RX Data Rate (Gb/s)	2.0 / 2.5	5	5	25	10	10 / 24
RX Energy (pJ/bit)	0.165 / 0.220	1.57	5.6	1.2	0.3	0.43 / 0.4
RX Area (mm^2)	0.0003	0.005	0.015	0.03	-	0.003
RX C_{in}	15	50	-	20	60	200
Sensitivity (Average Photocurrent, μA)	10 / 15	125	-	180	24	27 / 80
Sensitivity (Average Optical Power, dBmW)	-3 / -1.3	-2	-	-6	-14	-12.5 / -4.7
TX Data Rate (Gb/s)	3.5	5	5	25	10	-
TX Energy (pJ/bit)	0.07	1.22	7.2	8.3	0.08	-
TX Area (mm^2)	0.002	0.006	0.026	0.05	-	-

Table 1.1: Optical transceiver performance for various integration options - prior work.

Chapter 2

Link-level Anlaysis of a Photonic Network

To realize the full potential of integrated silicon photonics, it is essential for device, circuit and system designers to understand the relationships and trade-offs among components, as well as the impact of the different device-integration scenarios. In this chapter, we illustrate these design trade-offs on an example integrated WDM photonic link. We create the component models that connect device, process and circuit parameters to link performance and power consumption. We show that careful design-space exploration coupled with new tuning techniques can result in an optimal set of link, circuit, and device parameters. We demonstrate that the choice of channel-rate (per-wavelength) allows us to achieve minimum energy-costs by balancing laser, tuning and circuits backend power. This exploration also highlights the significance of monolithic-integration to minimize the receiver parasitics, reducing the laser power and enabling low-energy, high throughput-density interconnect fabrics.

This chapter largely consists of the photonic link analysis paper that we presented at CICC in 2011 [27]. It was wholely a team effort, and I thank the co-authors and the rest of the photonics team that contributed. There are several subsections that I did not play a large role in, but they are critical in understanding how the larger system model was put together. The equations defining receiver sensitivity in Section 2.1.2 were changed to the form presented in [28] in order to be clearer to the reader. Also,

detail on system-level ring-tuning has been omitted, as this as a new technique (rather than part of an illustrative model) and detailed in a separate thesis.

2.1 Photonic Link Components

We begin with a discussion concerning the operation of modulators and receivers, the primary data-path elements that form a WDM photonic link (Figure 2-1). As these two components depend highly on the characteristics of the devices that they use, we motivate a device-technology-driven analysis. We consider two integration scenarios - a monolithic integration of photonic components into the CMOS front-end (polysilicon photonics or thin-BOX SOI photonics), and a hybrid integration scenario with optimized SOI photonic die attached to the CMOS chip via through-silicon vias (TSVs). The monolithic integration will typically have smaller parasitic capacitances between circuits and photonic components, while potentially having higher optical lossess due to fabrication process constraints.

2.1.1 Modulator

The optical ring-resonant modulator and driver convert electrical data into the optical domain by on-off keying one CW λ from the multi-λ laser source. Light is modulated by shifting the ring-resonant filter's passband in and out of the optical channel's wavelength. The passband is shifted most efficiently using the free-carrier plasma dispersion effect [29] to change the refractive index of the ring material. To avoid the high energy cost of carrier-injection modulators, which have a high on-current due to carrier recombination [30], we focus on reverse-bias driver designs that modulate the depletion width of a vertical P-N junction fabricated in the polysilicon/silicon ring [31].

Modulator Design

A key system tradeoff exists between the extinction ratio ER of the modulator (a ratio of the on-to-off light intensity), its insertion loss IL, and its total energy cost.

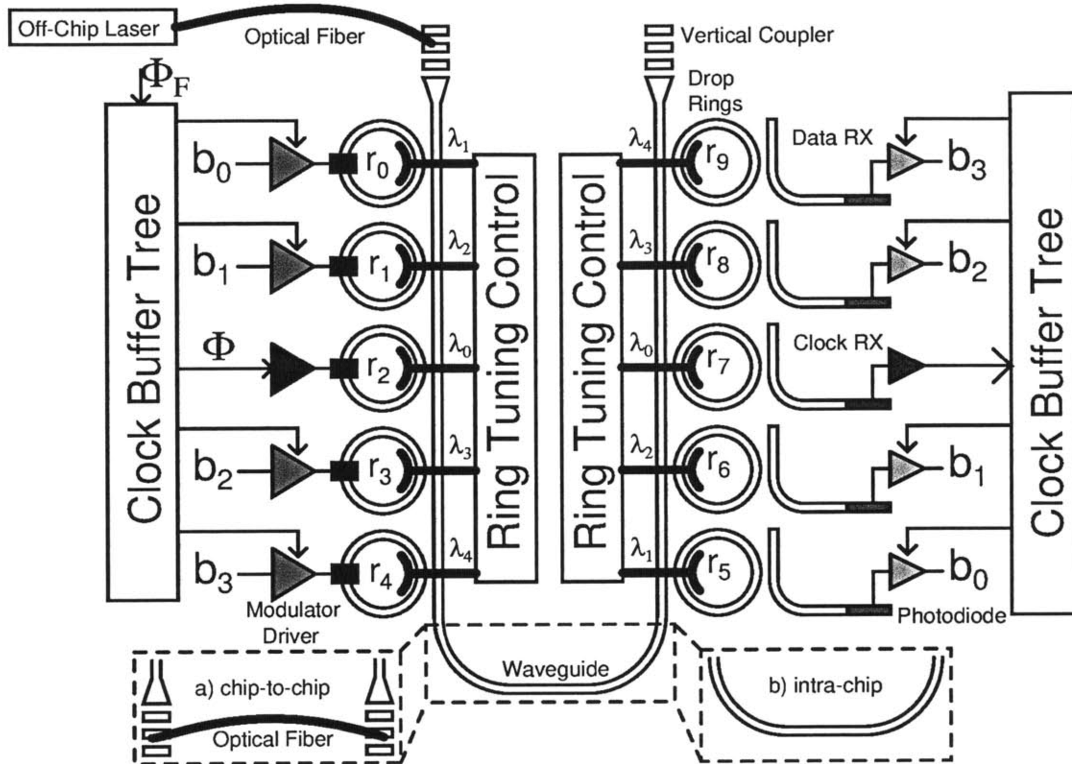

Figure 2-1: An integrated WDM photonic link. A continuous-wave (CW) multi-λ laser is coupled onto the chip through a vertical-coupling grating structure. Once on chip, frequency selective ring-resonant modulators encode digital bitstreams onto their resonant wavelengths. Each wavelength propagates along the waveguide (and possibly off-chip) until it is routed through a matching drop ring to an integrated photodiode (PD). An optical receiver forms a bit decision based upon the PD photocurrent. Clock signals are routed both optically along the waveguide and electrically through local H-trees. Ring tuning circuits are used to tune the resonance of the modulator and drop rings.

A small shift in the ring's Lorentzian frequency response requires a small energy cost, but results in a low ER. The location of the CW laser resonant wavelength with respect to the resonance of the ring is set by the desired IL, ER and data rate DR (bandwidth) requirements, which then set the modulation energy cost. The system designer must balance the modulation energy cost with receiver and laser power, which depend strongly on the extinction ratio, insertion loss and data rate.

Device Energy

In this section we calculate the modulation energy required to achieve given specifications. The necessary charge difference between the on- and off-states, ΔQ, is determined in Equation 2.1 by evaluating the Lorenzian transfer function, where $T_1 = 1/IL$ is the ring's transmissivity in its shifted on-state, $T_0 = 1/(ER \cdot IL)$ is its transmissivity in its off-state, and T_n is its transmissivity at the ring's resonant wavelength, Figure 2-3. Q_0 is the charge difference necessary to shift the ring by its half-width-half-max bandwidth [30], where q is the charge of an electron, $n_g = 4$ is the group index of the ring, and some typical ring modulator parameters are set as $V_{tot} = 2.6e - 12$ cm^3 is the total volume of the ring, $n_f = 3e - 21$ cm^3 is the carrier-induced index change per unit carrier density at $\lambda_0 = 1300$ nm, and $\Gamma = 0.4$ is the overlap of the optical mode with the ring cross-section. The quality factor $\mathbb{Q}_f$ of the ring is set by the required data rate as $\mathbb{Q}_f = \frac{8 \cdot \pi}{3\lambda_0 \cdot DR}$, with a maximum value of 1e5, limited by practically achievable optical losses.

$$\Delta Q = Q_0 \cdot \left(\sqrt{\frac{T_1 - T_n}{1 - T_1}} - \sqrt{\frac{T_0 - T_n}{1 - T_0}} \right), Q_0 = \frac{q \cdot n_g \cdot V_{tot}}{2 \cdot \mathbb{Q}_f \cdot n_f \cdot \Gamma} \tag{2.1}$$

The modulator diode operates as a varactor in the reverse-bias regime. The required charge difference is integrated onto the nonlinear junction capacitance (Equation 2.2) to determine the minimum reverse-bias drive voltage. V_a is plotted in Figure 2-2b as a function of data rate, at fixed ER and various values of IL.

$$\Delta Q = Q(V_a) - Q(0) = \int_0^{V_a} \frac{C_{j0}}{\sqrt{1 + \frac{V}{V_{bi}}}} dV \tag{2.2}$$

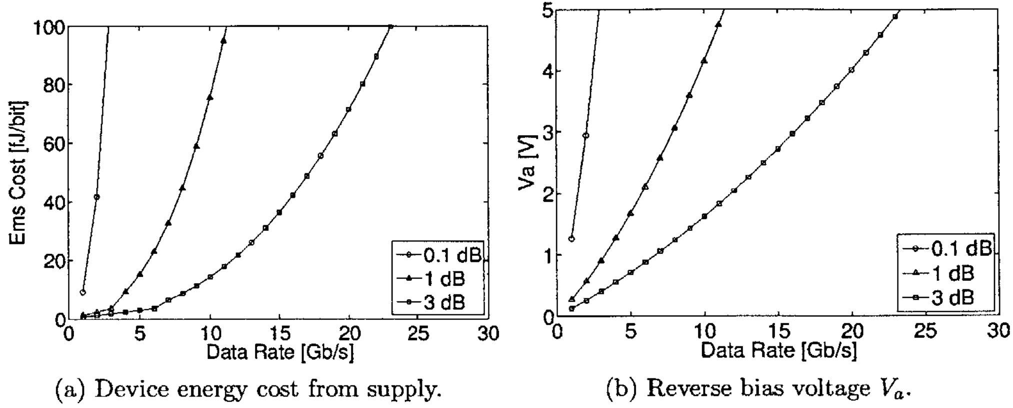

(a) Device energy cost from supply.

(b) Reverse bias voltage V_a.

Figure 2-2: Device requirements to reach the target $ER_{dB} = 6\text{dB}$ for various values of IL_{dB}.

If V_a is less than the supply V_{DD}, the energy to charge the junction comes from V_{DD} and the energy-per-bit is $E_{ms} = \Delta Q \cdot V_{DD}/4$ assuming a random data pattern. Otherwise the energy must come from a higher-voltage source, which we assume is generated from the supply with a conversion efficiency of $\eta = 0.5$. With $V_a > V_{DD}$, the energy-per-bit drawn from the supply is $E_{ms} = \Delta Q \cdot V_a \cdot (1 + \eta)/4$ (Figure 2-2a).

Circuit Energy

The driver model is shown in Figure 2-3 as an inverter chain pre-driver followed by a final driver stage. The circuit topology of the final drive stage will change based on V_a; If $V_a \leq V_{DD}$, a low-swing topology can be used (Figure 2-3a); otherwise a voltage-boosting circuit may be necessary (Figure 2-3b). The final stage is modeled as an effective resistance R_{eff} and a parasitic capacitance C_{par}, connected to wiring capacitance C_{wire}. Logical effort analysis is used to size the driver (W) and pre-driver chains (FO) to meet the data-rate requirements.

$$E_{dr} = \frac{C_{par} + C_{wire}}{4} \cdot \max\left(V_a \cdot V_{DD}, V_a^2\right) + \frac{1}{4} \cdot \frac{3 \cdot C_g \cdot W}{1 - \frac{1}{FO}} \cdot V_{DD}^2 \tag{2.3}$$

Figure 2-4 shows the driver energy-per-bit cost to reach $ER = 6$ dB for various values of IL. For current modulator device technology, to satisfy data rates up to 30 Gb/s, the intrinsic RC time constant of the device allows designers to choose R_{mod}

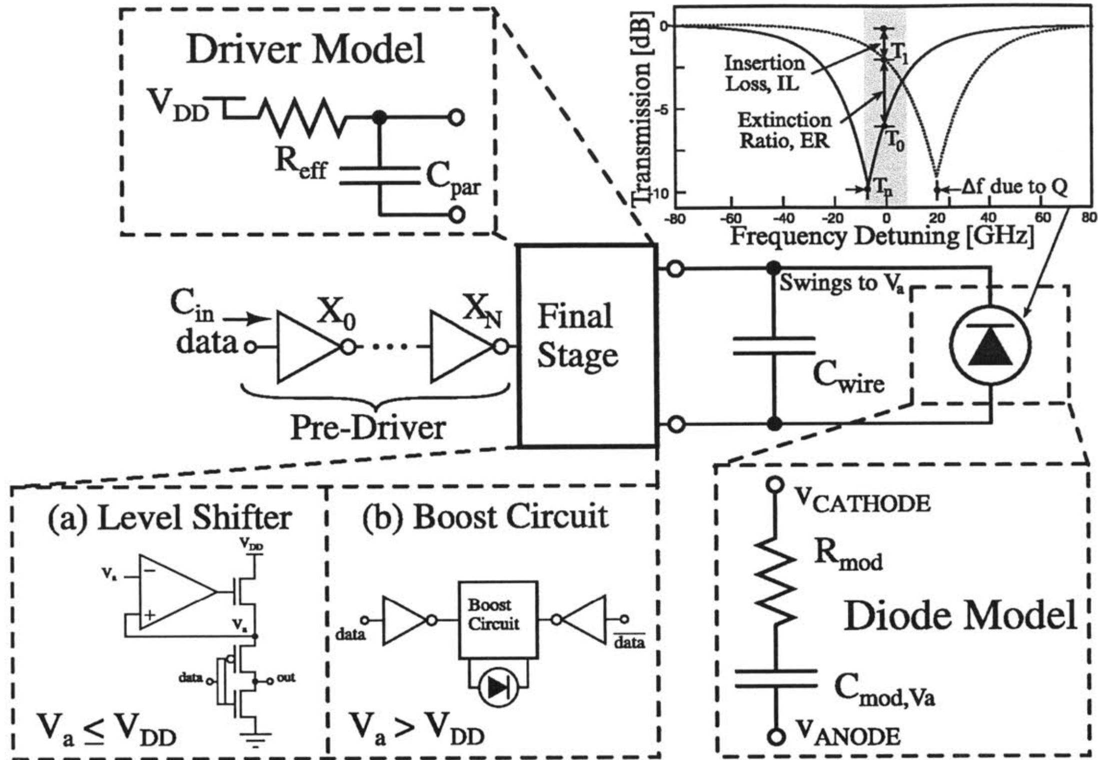

Figure 2-3: Electrical model of modulator device and driver.

up to 1 kΩ without significant energy penalty, relaxing the modulator optical losses due to contact placement.

2.1.2 Optical Data Receiver

The optical receiver converts optically-modulated data back into the electrical domain by sensing a photocurrent produced by the PD. In contrast to traditional optical receivers which utilize various power-hungry trans-impedance amplifiers (TIA) to combat the large PD parasitic capacitance, monolithic integration offers the opportunity for much-simpler, energy-efficient receiver circuits due to low PD parasitic capacitances. In this section, we illustrate the relationship between sensitivity and power consumption across ranges of data rates and parasitic capacitances, for various receiver topologies, including transimpedance amplifiers (TIA) and integrating receivers.

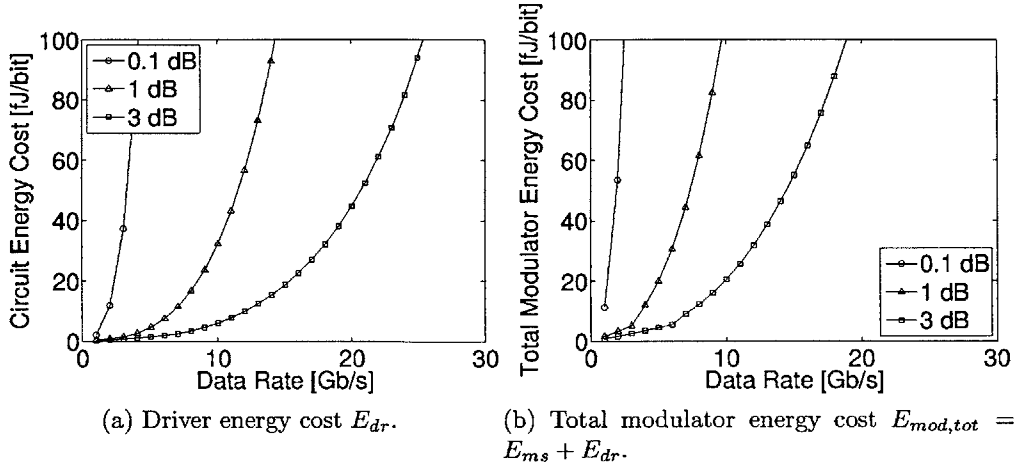

(a) Driver energy cost E_{dr}.

(b) Total modulator energy cost $E_{mod,tot} = E_{ms} + E_{dr}$.

Figure 2-4: Energy costs to reach the target $ER_{dB} = 6$ dB for various values of IL_{dB}.

Photodiode

An equivalent model of the PD is shown in Figure 2-5a, consisting of a capacitance in parallel with a photocurrent-generating source and series resistance. The PD is connected to the front-end through either a Through-Silicon Via (TSV, $C_p \approx 25$fF) or low-level metal routing ($C_p \approx 5$fF). All parasitic series resistances are assumed negligible.

Sense Amplifier

A sense amplifier (SA) is used as a comparator to regenerate the full-swing digital signal. The main factors affecting SA sensitivity are mismatch, settling time, supply noise, and circuit noise. The minimum input signal that allows the latch's decision nodes to settle to the rails is $v_{sense} = V_{DD}e^{-T_{bit}/(2\tau)}$, where τ is the time constant of the exponential regeneration. Residue offset due to mismatch is compensated by a 5-bit DAC [32] resulting in $v_{OS,res} = 3v_{OS}/2^5$, with v_{OS} of 40mV taken from [33]. Deterministic ($v_{supply,det}$) and random ($v_{supply,rand}$) supply noise estimates were taken from [34] [35]. The PD's noise is shot-noise dominant [36], and is given by $\sigma_{i,PD} = \sqrt{2qI_{PD}\Delta f}$. v_{margin} accounts for any other un-modeled noise or non-idealities, and was set to 20mV.

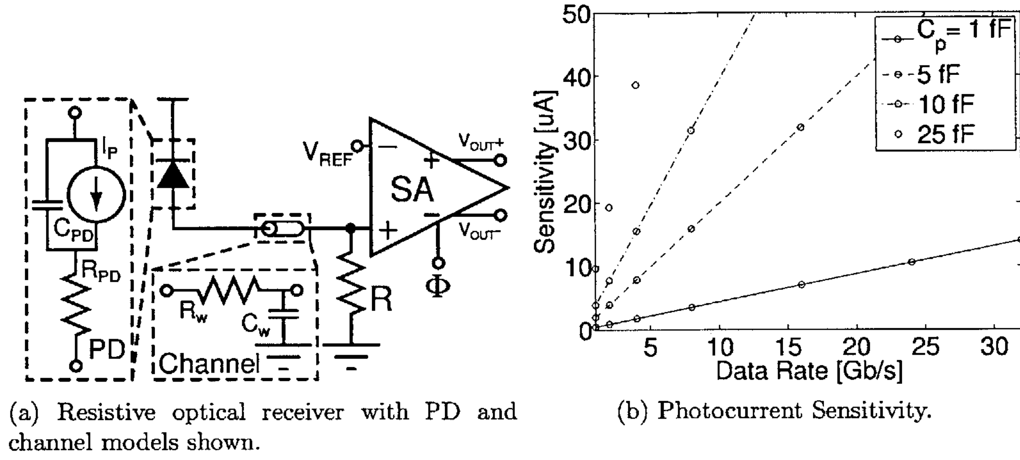

(a) Resistive optical receiver with PD and channel models shown.

(b) Photocurrent Sensitivity.

Figure 2-5: Resistive receiver sensitivity.

By input-referring the SA's input swing requirements across each of the front-ends considered through the front-end's transimpedance, R_{FE}, Equation 2.4 describes the receiver's input current-swing requirement in terms of input photocurrent [2, 37].

$$\Delta I = \underbrace{i_{sense}}_{v_{sense}/R_{FE}} + \underbrace{i_{OS,res}}_{v_{OS,res}/R_{FE}} + \frac{v_{supply,det}}{CMRR \cdot R_{FE}} + \underbrace{i_{margin}}_{v_{margin}/R_{FE}} + \sqrt{SNR}\sigma_n \quad (2.4)$$

where

$$\sigma_n = \sqrt{\sigma^2_{i,circuit} + \frac{v^2_{supply,rand}}{CMRR^2 \cdot R^2_{FE}} + \sigma^2_{i,PD}}. \quad (2.5)$$

The input sensitivity of the receiver can then be computed as $I_{ON} = \Delta I/(1 - 10^{-ER_{dB}/10})$, where ER is the extinction ratio of the modulator and $\Delta I = I_{ON} - I_{OFF}$ is the difference in photocurrents required to meet a given BER requirement.

Resistive Receiver

Figure 2-5a shows a resistive receiver with a SA. Photocurrent is driven across the resistance, R, which is the front-end's gain. The dominant pole is at the input node. Equation 2.6 (with $R_f = R_{in} = R$) shows that the resistor is penalized for its parasitic capacitance through the parameter k_R, assumed to be 0.4fF/kΩ [38]. For each data rate and $C_p = C_{PD} + C_w + C_{front-end}$, the maximum R is computed from Equation 2.6. BER requirements set the minimum ΔI.

Figure 2-5b shows that the receiver is able to sense photocurrents of approximately 10μA for low C_p and data rate. The sensitivity worsens linearly with data rate as R is traded for bandwith. The energy-efficiency remains constant due to the dominance of the SA's digital switching power.

TIA

A TIA (Figure 2-6) breaks the gain-bandwidth limitation of the resistive receiver. Equation 2.7 shows the receiver's gain and use of feedback to decrease input impedance [36].

$$BW = \frac{1}{2\pi R_{in}\left(C_{PD} + k_R R_f\right)} \tag{2.6}$$

$$R_{TIA} = \frac{g_m - g_f}{g_f\left(g_m + g_{ds}\right)} \tag{2.7}$$

$$R_{in} = \frac{g_{ds} + g_f}{g_f\left(g_m + g_{ds}\right)} \tag{2.8}$$

Figure 2-7a shows sensitivity-optimized designs for different TIA bias powers. In this relatively small C_p environment, large designs are penalized for their increased gate capacitance, requiring a reduction in R_{in} and therefore R_{TIA} and sensitivity. Figure 2-7b shows sensitivity optimum for various values of v_{margin}.

Figure 2-8 summarizes TIA performance for various values of C_p. Though the TIA achieves sensitivity superior to the resistive receiver, the power consumption is considerably worse.

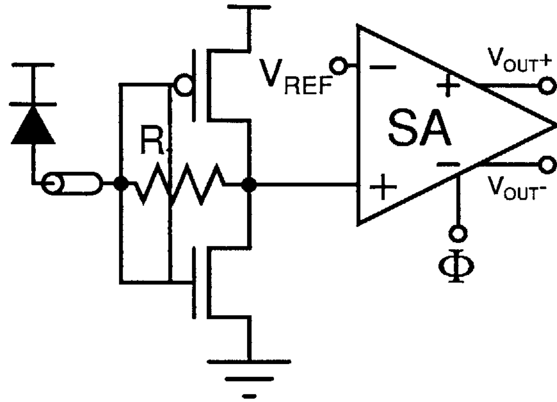

Figure 2-6: Transimpedance Amplifier.

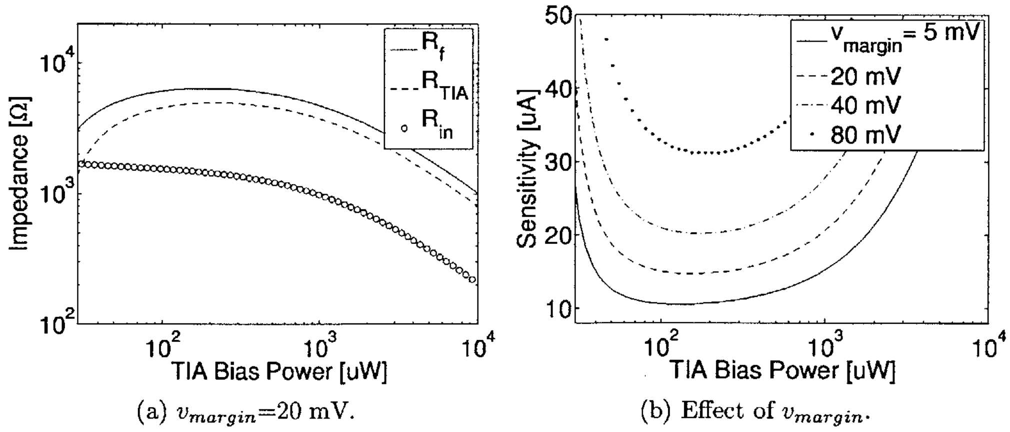

(a) v_{margin}=20 mV.

(b) Effect of v_{margin}.

Figure 2-7: TIA design example at C_p=25 fF, DR=5 Gb/s.

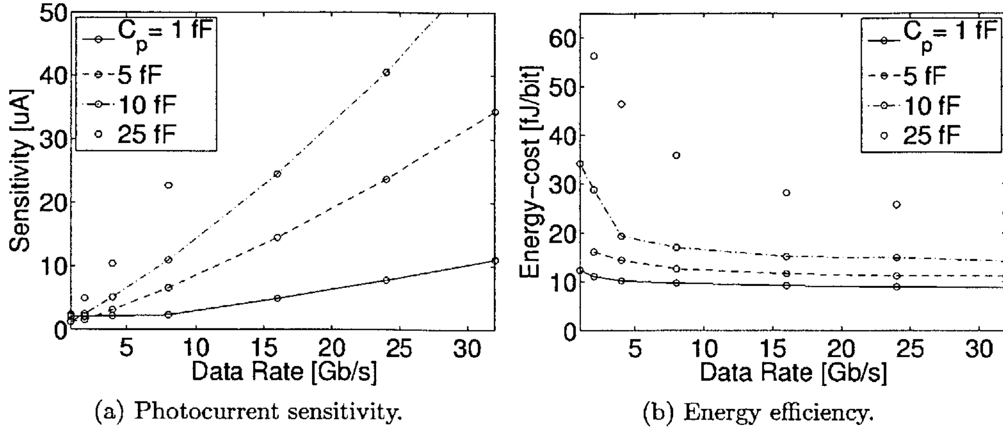

(a) Photocurrent sensitivity.

(b) Energy efficiency.

Figure 2-8: TIA performance.

Current-Integrating Receiver

The third topology considered is an integrating receiver (Figure 2-9a), where the photocurrent is converted to a voltage by integrating it onto a capacitor $C_{INT} = C_{PD} + C_w + C_{SA,in}$. The photocurrent is integrated over a fraction (k_{INT}=0.7) of a bit time yielding a front-end gain given by Equation 2.9.

$$R_{INT} = \frac{k_{INT} \cdot T_{bit}}{C_{INT}} \tag{2.9}$$

Figure 2-9 shows that the integrating receiver is the best performing of the three receivers considered. The energy-efficiency of the receiver is dominated by the SA as in the resistive receiver. It should be noted that this simple model has several hidden challenges remaining. The voltage on C_{INT} must be reset or at least charge-shared [12], which is partially accounted for through k_{INT}. A small C_{INT} will also suffer from SA kickback, while increasing C_{INT} degrades sensitivity.

2.1.3 Optical Clock Receiver

Clock distribution is critical in synchronizing communication channels and functional blocks in high-performance processors. The simplest optical clock receiver

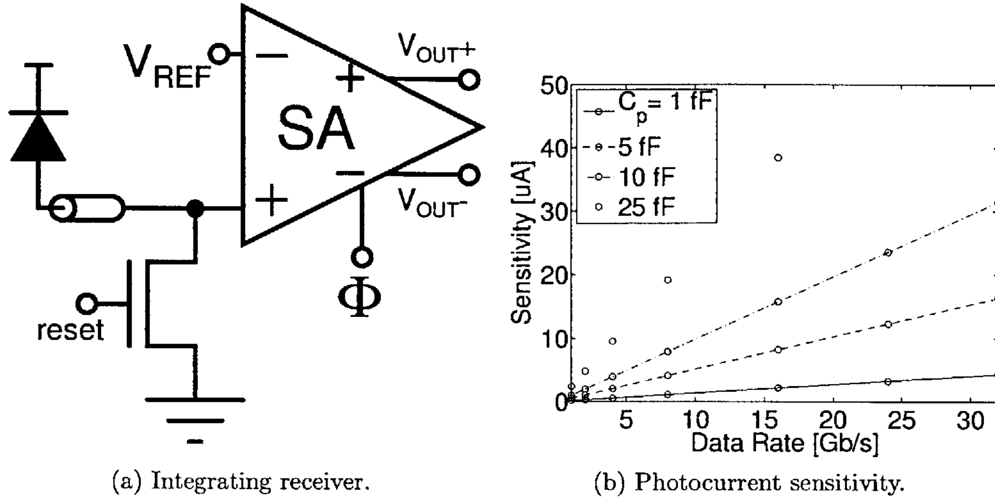

(a) Integrating receiver. (b) Photocurrent sensitivity.

Figure 2-9: Integrating receiver design and performance.

considered is a receiverless clocking scheme [39]. By alternately illuminating two PDs stacked in series, a clock signal is generated at the internal node. The only circuitry between the PDs and the clocked node are buffers, minimizing added jitter and reducing circuit power consumption. However, a large optical power is required to create the rail-to-rail voltage swing at the PD. As shown in Section 2.1.2, differential TIAs can be used to amplify the signal at the cost of added noise and circuit power [36] [40].

In addition to bandwidth and sensitivity constraints, a clock receiver output must meet a given jitter specification. While C_p is relatively small, it is still much larger than any circuit loading. The voltage transient slope is then approximated by C_{PD}/I_{ON}. Any transistor or power supply noise is input referred onto this slope, yielding a timing jitter.

2.1.4 Single Channel Link Tradeoffs

To illustrate the interactions between the modulator and receiver and the impact on wall-plug laser power, we perform a power optimization across modulator insertion loss, extinction ratio, and receiver topologies for different link data-rates. Figure 2-10 shows the energy-per-bit breakdowns for four integration scenarios.

In all plots, the laser power is the dominant energy consumer, increasing quickly

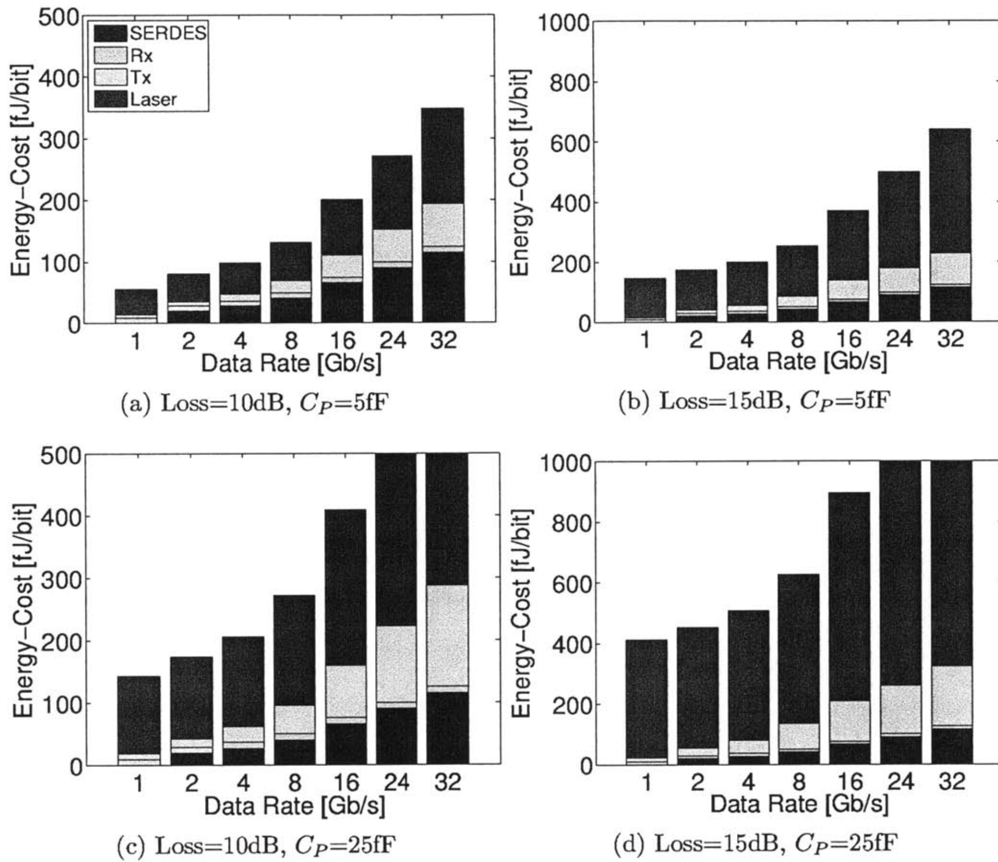

(a) Loss=10dB, C_P=5fF

(b) Loss=15dB, C_P=5fF

(c) Loss=10dB, C_P=25fF

(d) Loss=15dB, C_P=25fF

Figure 2-10: Data rate tradeoffs for a single photonic link for 4 integration scenarios. C_P=5fF represents monolithic integration, while C_P=25fF is expected for a TSV connection to an optical die. Channel losses of 10dB and 15dB correspond to on-chip and chip-to-chip links, respectively.

Table 2.1: Link Evaluation Parameters

Parameter	**Value**
Process Node	32nm Bulk CMOS
V_{DD}	1.0V
Device to Circuit Parasitic Cap C_P	5-25fF
Wavelength Band λ_0	1300nm
Photodiode Responsivity	1.1A/W
Wall-plug Laser Efficiency P_{laser}/P_{elec}	0.3
Channel Loss	10-15dB
Insertion Loss IL_{dB} (Optimized)	0.05-5.0dB
Extinction Ratio ER_{dB} (Optimized)	0.01-10dB
Bit Error Rate (BER)	10^{-15}
Core Frequency	1GHz
SERDES Topology	Mux/Demux Tree

with data rate as aggressive modulation rates force a relaxation of modulator insertion loss and extinction ratio. We can see that the laser power is highly sensitive to C_P, as the laser power for C_P=25fF is roughly 5X that of C_P=5fF. Though the modulator tries to offset the laser cost by increasing its extinction ratio and decreasing insertion loss, it inevitably reaches a limit on its capabilities. The higher loss simply amplifies the laser power component, resulting in a 3X laser power difference between the 15dB and 10dB loss cases. Matching previous analysis, the optimization chose the integrating receiver as the optimal receiver in all scenarios. Though our results present a grim outlook for the C_P=25fF (optical die with TSV) scenario, we note that lower losses may be achievable with a dedicated optical die, allowing TSV integration to remain competitive.

2.2 Towards a Full WDM Link

Expanding upon our analysis for a single-channel data link, we explore the additional backend components required in a high-speed multi-channel WDM link. We present a model of optical clock distribution and source-synchronous clocking. Then, we outline techniques for tackling resonance mismatches of optical ring resonators, a key challenge in nanophotonic integration.

2.2.1 Optical Clock Distribution

An example point-to-point optical WDM source-forwarded link is shown in Figure 2-1. Clock transmission occurs on λ_0. At the differential receiver, the signal is regenerated, buffered, and used to clock the other data receivers. Since clock-TX and data-TX share the same clock fabric, relative jitter between the sent clock and data is minimal.

On top of previously discussed benefits of optical links, optical clock signaling does not suffer power from rail injected noise or crosstalk, so no jitter is added in the channel. This obviates the need for an RX PLL/DLL, greatly reducing the power and area overhead. The low latency of optical waveguides also means that all data receivers can operate on the same clock phase.

By accounting for the total clock load at the TX or RX, including wire routing, the power consumption of the clock distribution is modeled for a given clock frequency. The expected timing jitter can be derived from [41]. In Figure 2-11a we fix C_{PD} at 5fF, and plot the total clock power across frequency, with the criteria that the timing jitter at the data receivers is less than 3% unit interval (UI). With higher clock frequency, fewer channels are needed for a given total data throughput, decreasing the distribution endpoint capacitance. As we increase the frequency, the jitter requirement (in seconds) tightens, requiring an increase in power. When I_{ON} increases, the jitter performance of the receiver improves proportionally, allowing for less electrical power. Figure 2-11b shows the power consumption of the TX clock tree, RX clock tree and clock receiver circuit for I_{ON}=10μA and jitter of 3% UI.

2.2.2 Ring Resonance Mismatch

An integrated WDM link relies heavily upon optical ring resonators to perform channel selection using the ring's resonant frequency. Dependent upon both device geometry and the index of refraction, large ring resonance mismatches can arise from limited process tolerances and temperature changes. For rings built with gate polysilicon on commercial CMOS bulk processes, process variation can result in reso-

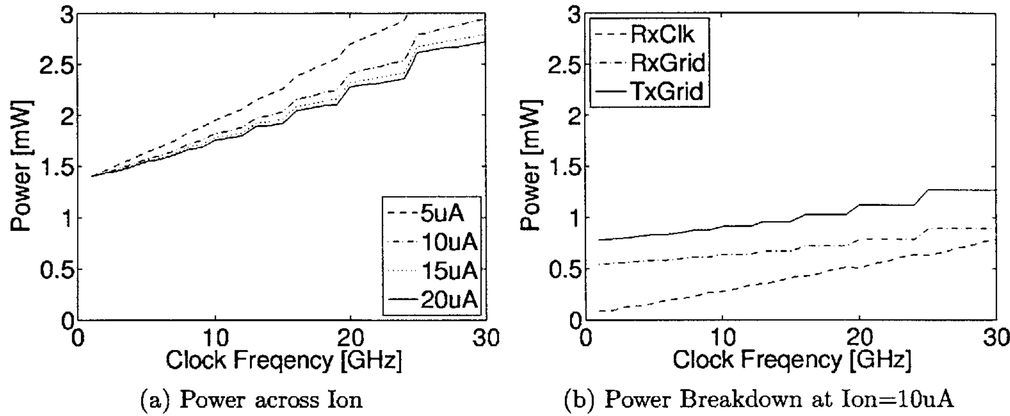

(a) Power across Ion (b) Power Breakdown at Ion=10uA

Figure 2-11: Power vs. datarate per channel, jitter is fixed to be within 3% UI, and C_{PD}=5fF

nance mismatches of up to 90GHz and absolute die-to-die mismatches of 600GHz or more [42]. Similarly, mismatches with standard deviations in the range of 20-70GHz for same-die and 150-220GHz for die-to-die have been observed for rings built with SOI [43]. As local- and systematic-level process mismatches differ greatly in magnitude and tuning implications, we model them using σ_{rL} and σ_{rS}, corresponding to the standard deviations characteristic of local ring-to-ring and global systematic mismatches, respectively.

A strong thermal dependence in the index of refraction of silicon causes ring resonances to drift with temperature. $\frac{\Delta f}{\Delta T}$ in the range of -10GHz/K have been observed [42] [44], implying that a shift of several hundred GHz can be expected in a hostile thermal environment, such as that of a high-performance processor. Unlike static process variations, however, thermal fluctuations are time-dependent, requiring active tuning to stabilize ring resonances. At the same time, strong temperature dependence allows for simple and effective thermal compensation of process mismatches. Recently, *athermal* ring resonators using polymer-based cladding [45] have also been proposed as a solution to undesired thermal-induced resonance drifts. Their inability to be thermally tuned, however, means that any process-induced mismatch must be compensated by UV trimming on a per ring basis, potentially limiting commercial

Table 2.2: Tuning model evaluation parameters

Parameter	Value
Aggregate Link Throughput	64Gb/s
Free Spectral Range (FSR) R=3um ring	4THz
Heating Efficiency	44K/mW [42]
Tuning Efficiency $\frac{\Delta f}{\Delta T}$	10GHz/K
Local Process Variation σ_{rL}	varies
Systematic Process Variation σ_{rS}	varies
Temperature Range	300-360K
Process (for electrical backend)	32nm Bulk CMOS
Tuner Controller Power	10uW/Ring
Electrical Tuning Limit	50GHz
Yield Target	99%

scalability.

2.2.3 Ring Tuning Techniques

As explained at the start of this chapter, this section has been omitted from this thesis. Please read [27] for more detail on the ring tuning techniques.

2.2.4 Ring Tuning Model

To evaluate the tuning strategies in [27], we develop a *Monte Carlo* based tuning model. For each tuning scenario, a set of rings in a ring filter bank with some desired resonances is fabricated. To simulate the effects of local process variations, we randomize the resonance of each ring using σ_{rL} and apply global systematic variations using σ_{rS}. The model then attempts to tune the set of rings across a range of temperatures. If successful, the tuning power cost is reported. The experiment is performed 1000 times for each parameter combination (fabrication bias, number of extra rings, etc.) to find the optimum tuning strategy for a given yield target.

The power needed to perform full thermal tuning is shown in Figure 2-12 across a range of process variations (σ_{rS}, σ_{rL}) and channelizations. Against process variations, the power cost of full thermal tuning is linear with both σ_{rS} and σ_{rL}, stemming from the increase in fabrication bias needed in order to maintain the same yield given higher

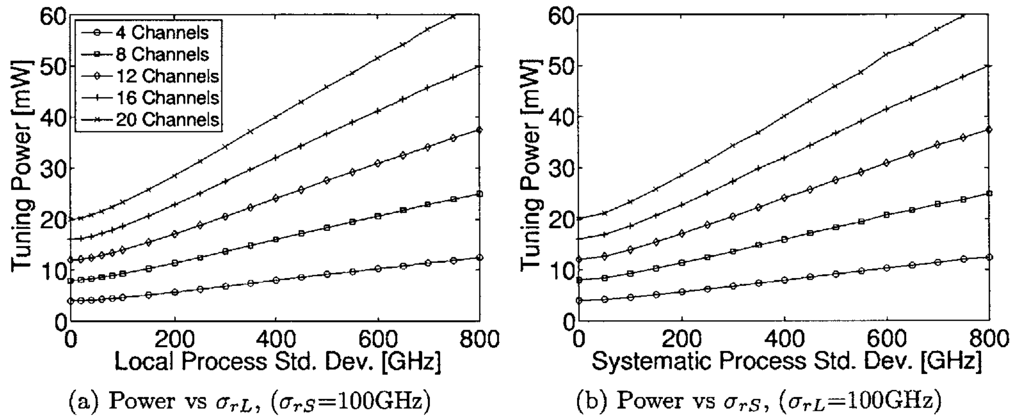

(a) Power vs σ_{rL}, (σ_{rS}=100GHz) (b) Power vs σ_{rS}, (σ_{rL}=100GHz)

Figure 2-12: Tuning power vs process variation at various channelizations for the full thermal tuning scenario.

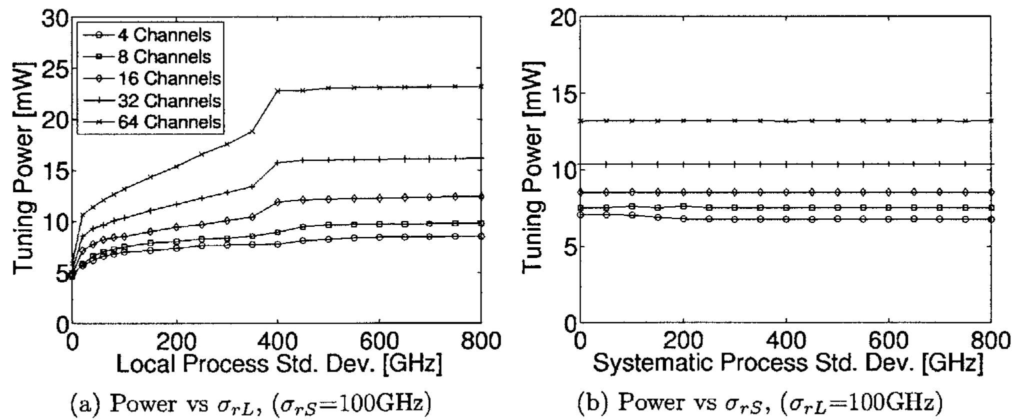

(a) Power vs σ_{rL}, (σ_{rS}=100GHz) (b) Power vs σ_{rS}, (σ_{rL}=100GHz)

Figure 2-13: Tuning power vs process variation at various channelizations with an electrical backend capable of bit reshuffling.

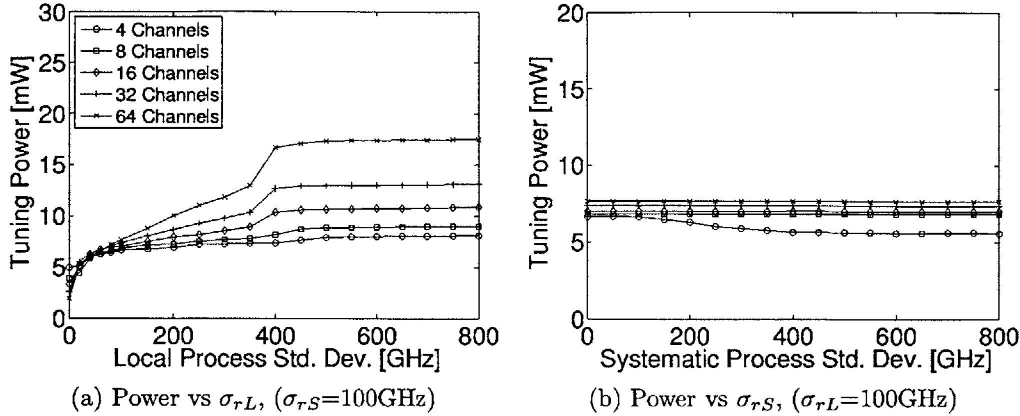

(a) Power vs σ_{rL}, (σ_{rS}=100GHz)

(b) Power vs σ_{rS}, (σ_{rL}=100GHz)

Figure 2-14: Tuning power vs process variation at various channelizations with a bit reshuffling electrical backend and electrical assisted tuning.

process variations. The increase in tuning power is also linear with the number of channels, tracking the increase in the number of rings that require tuning.

Using an electrical backend to perform bit reshuffling, we show that tuning power can be successfully decoupled from σ_{rS} (Figure 2-13). Local variations (σ_{rL}) still affect the tuning power, as a larger σ_{rL} requires a larger degree of bit-reorder multiplexing. The tuning power also scales gracefully with the number of channels, owing to the decrease in channel-to-channel separation (and tuning distance) of each ring. Electrically-assisted tuning with an electrical backend allows for even further reductions in tuning power. As shown in Figure 2-14, cases with high numbers of channels benefit most as the channel separation is small enough to be covered electrically, without using heaters. Using this backend, we demonstrate a 5-10X tuning power reduction at dense WDM channelizations while maintaining tuning robustness across a range of process variations.

2.3 WDM Photonic Link Evaluation

In this section, we perform a full link-level optimization and evaluation of a WDM link to quantify energy consumption tradeoffs. In our evaluation, we explore links with 4 different aggregate throughput design points, 64Gb/s, 256Gb/s, 512Gb/s,

1024Gb/s, corresponding to minimum, medium, high, and maximum bandwidth scenarios.

Figure 2-15 shows that tuning power dominates at lower data-rates (since there are more channels given fixed throughput) and decreases with data-rate. Modulator, laser, SERDES, and receiver energies increase with data-rate and dominate at high rate-rates. At all throughput scenarios, an optimal energy balance is achieved at around 4-8Gb/s. An overall energy-optimal point occurs at less than 200fJ/bit for a link with 256Gb/s of aggregate throughput and 4Gb/s data-rate.

At the energy optimal point, we see that the energy consumption is roughly an even 3-way split between tuning, laser, and mod/rx/SERDES. As tuning power is now mostly dominated by the backend electrical components, this energy will scale favorably with technology and can be optimized using custom design. A full electrical tuning backend is also unnecessary on both modulate- and receive-side – barrel-shifts and bit-reordering only need to be performed once – meaning backend power can be cut by another 50%. Refinement of photodetector responsivity and parasitic capacitances as well as lower-loss optical devices with improved electrical laser efficiencies can bring about further reductions in wall-plug laser power. It can be expected that energy/bit will drop to sub-100fJ with device development, process scaling and overall link component refinement.

We conclude this section by answering a question brought up early on in Section 1.1: how the cost of photonics compares with its electrical counterpart. We previously calculated that electrical I/O is limited at the package level to $\approx$ 25Gb/s/mm^2, and that at the package-level, photonics can support $50Tb/s/mm^2$. Figure 2-16 shows the bandwidth density of integrated photoncis for a throughput of 256Gb/s compared to the electrical bottlenecks discussed earlier. Again taking the lower, more constraining curve, we can see that at the optimal data-rate, integrated photonics achieve a bandwidth density of more than two orders of magnitude greater than then electrical case, solidifying its position as the next-generation I/O technolgy.

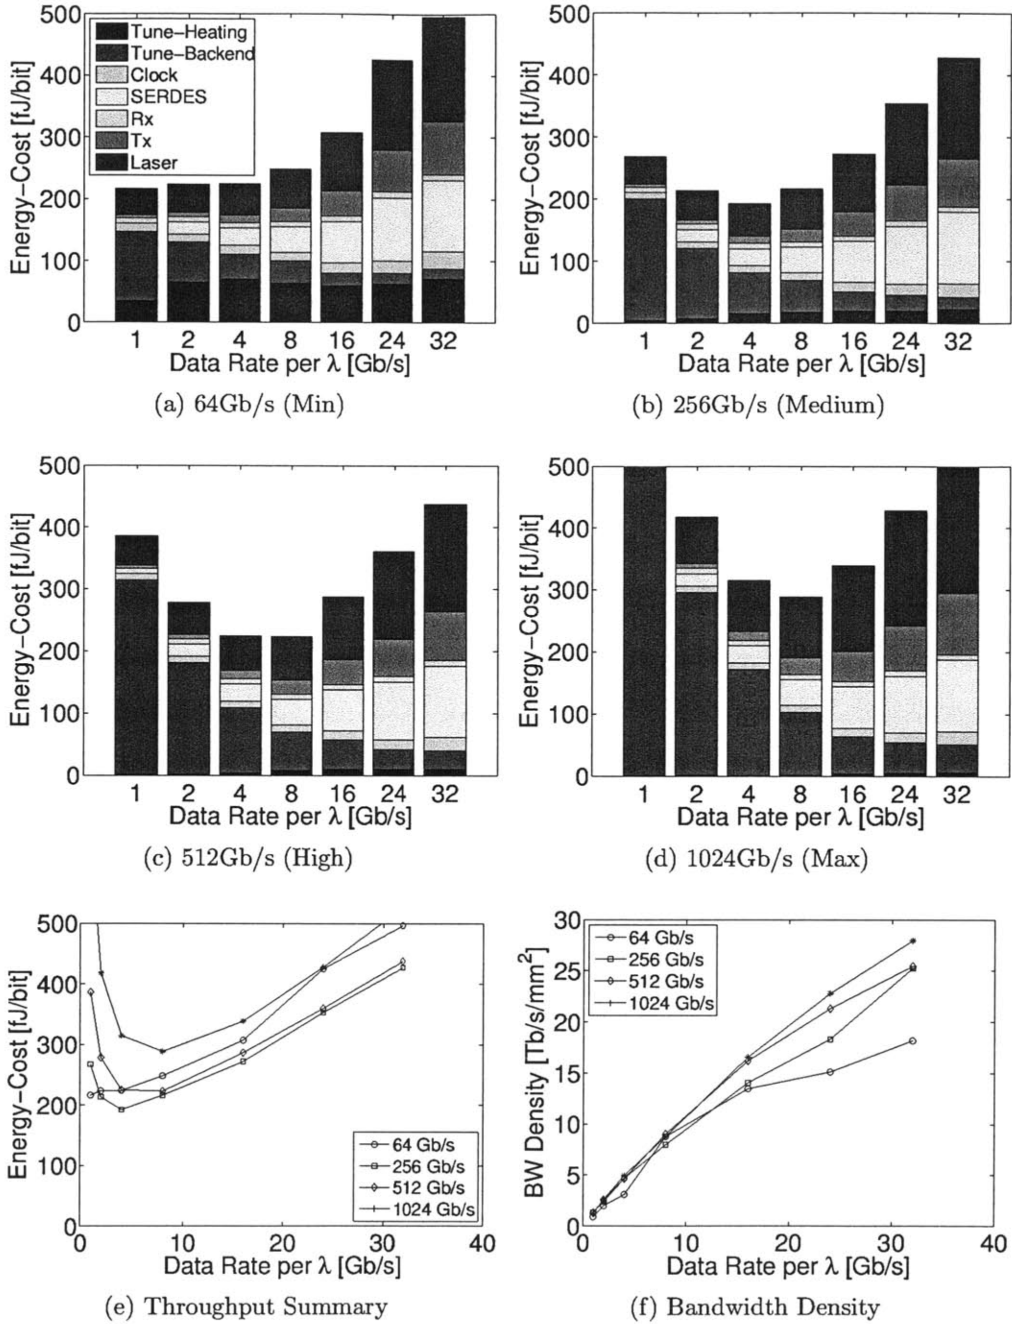

(a) 64Gb/s (Min)

(b) 256Gb/s (Medium)

(c) 512Gb/s (High)

(d) 1024Gb/s (Max)

(e) Throughput Summary

(f) Bandwidth Density

Figure 2-15: Optimized power vs. data-rate for different aggregate link throughputs for Loss=10dB, C_P=5fF. For tuning, we assume a bit-reshuffler backend and electrically-assisted tuning with local variation σ_{rL}=40GHz and systematic variation σ_{rS}=200GHz. Note that the number of WDM channels changes with data rate (Channels = Throughput / Data-Rate).

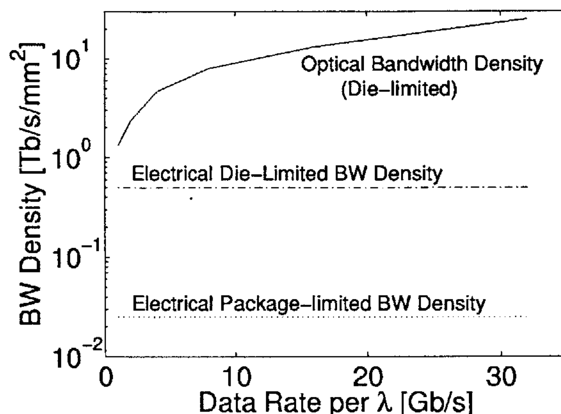

Figure 2-16: Bandwidth density of integrated photonics compared to electrical.

2.4 Summary

Integrated photonic interconnects are a promising solution to the throughput demands of future many-core processors. As an emerging technology, circuit, device and architecture designers require insights concerning the impact of device and circuit parameters on link-level figures of merit. This chapter presented a design-space exploration of a WDM integrated photonic link, facilitated through a set of circuit and device models that captured the optical-electrical tradeoffs of each link component. The modulator model showed the relationship between the modulation energy and the laser power, set through the extinction ratio and insertion loss specifications. Similarly, the optical receiver models demonstrated the degradation of sensitivity with data rate, which translated directly into increased laser power requirement. The impact of clock distribution was factored into the link-level analysis. Finally, ring tuning power was computed, allowing for dense WDM and robustness against process and thermal variations.

Using the models, we performed co-optimization across all link components for a complete WDM integrated photonic link. We found that relatively low (sub 10Gb/s) data-rates per link yielded optimal energy-efficiency accross a range of system throughputs. We showed that the photonic link is highly sensitive to parasitic capacitances present at the receiver input and that optical integration using TSVs to connect to

an optical die could result in significant overhead in the laser power. This study illustrated that monolithic integration of photonic components can offer interconnect solutions with high throughput-density and energy-efficiency.

Chapter 3

Optical Data Receiver Design and Techniques

In the previous chapters, we looked at traditional optical receivers designed for telecom applications as well as the work of other groups moving towards more integrated photonic systems. We created a system-level model based on compact representations of each of the components in an integrated photonic system. An optimization of the model showed that the energy-cost can be minimized for a fixed system throughput, and in particular for the throughputs examined in this work the energy-cost-minimum occur at relatively low data rates (below 10Gb/s), as microring tuning costs balances against other electrical costs.

Building on insight from the previous sections, this chapter develops receiver architectures and techniques that target the data-rates motivated by the system-level analysis. The designs we develop take advantage of the availability of a receiver-side clock to implement an energy-efficient receiver with a regenerative comparator. This chapter also develops several split-photodiode techniques that are enabled by monotlithic integration, where the circuit designer has unprecedented access to the design of the photodiode structure on the same chip.

Three different receiver designs will be detailed. The first is an aggressive, current-sensing latch that interfaces directly with the compact photodiode created through monolithic-integration. The design is completely digital, compact, energy-efficient,

and particularly well-suited for process scaling. The second design is based on a more traditional transimpedance amplifier (TIA). While the energy-cost increases due to the TIA stage, the use of a clocked sense amplifier to resolve the bit decision into digital values results in far lower power than the use of a limiting amplifier. The third receiver is also a TIA, but has been designed as a stand-alone circuit that can be tested with a GSG probe and two DC probes.

Chapters 3 and 4 are largely organized as design and experimental results, respectively. However, we bring measured results into this chapter where it is necessary to motivate certain design decisions that may have changed over the course of the work. One of those decisions comes up very early in this chapter.

3.1 Photodiode Design

Two types of photodetectors are used throughout this work. The first is a SiGe-based photodiode, and the second is a resont-ring-based defect detector. The focus of this thesis is the receiver design itself, but as will become clear, it is critical to understand the operation of the photodetector in order to design a suitable receiver.

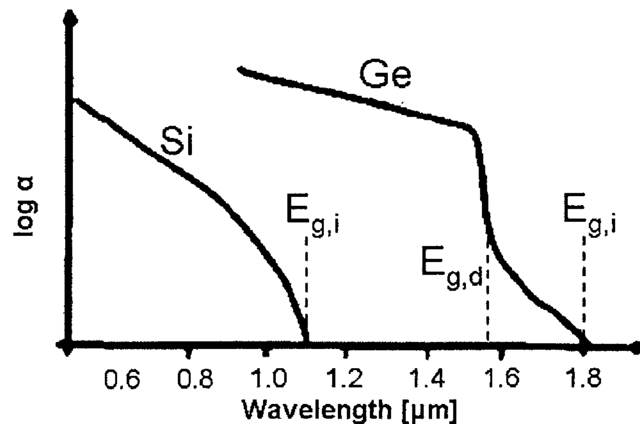

Figure 3-1: Optical absorption coefficients for Silicon and Germanium.

For the SiGe photodiode, consider the absorption curves as a function of wave-

length for Silicon and Germanium, shown in Figure 3-1. From the plot, we can see that if a wavelength between ≈1100nm and 1600nm is used for the link, then the light will propagate through Silicon, but be absorbed by Germanium. In this manner, we achieve both propagation through waveguides, and absorption in the detector. Note that the absorption coefficient curve for SiGe can be represented as a linear combination of the two curves shown depending on the fraction of Ge.

The defect detector utilizes free-carrier generation in ridge-waveguide microrings through sub-bandgap optical transitions involving defect states in the polySi waveguide core [46].

Once a suitable photodetector has been created, we look at how to connect that detector's terminals to a receiver. Figure 3-2 shows several photodiode connection options. In the first case, the photodiode is placed differentially across the inputs of the receiver. One of the key assumptions with this configuration is that the optical photocurrent will not significantly degrade with a low reverse-bias. If this is true, then with this configuration the receiver circuit does not have to tune out significant dark current. The fully-differential placement also helps to de-couple noise sources such as those from the supply.

Figure 3-2: Photodiode connection options.

In the second configuration, the photodiode is connected to one of the receiver's inputs, with the receiver's second input connected to a reference bias (which may be generated by the receiver itself). This configuration does not rely upon the assumption

that the photocurrent is strong at small reverse-biases. If, however, the dark current is unknown or has potentially large variation, then the circuit will have to be designed to accomodate a large range of dark currents.

The third configuration - and the one that has been used in the latest-generations of our chips - is one where the photodiode is connected with a large reverse-bias, and a second, non-optically-connected photodiode is connected to the reference terminal of the receiver. The advantage of this configuration is that the dummy photodiode can serve to match the dark dark current and capacitance of the optically-connected device. This greatly reduces the requirements on the circuit's dark-current tuning blocks. As the receivers are monolithically-integrated and fairly compact, the addition of a second photodiode does not significantly increase the area cost of the receiver, as it would in discrete- or hybrid-integration scenarios.

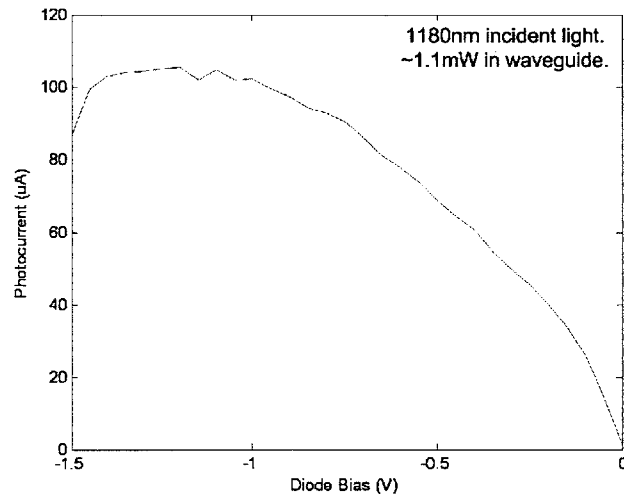

Figure 3-3: Photocurrent as a function of reverse bias, for SiGe detectors in logic process. Measurement performed by Jason Orcutt.

One of the largest challenges in this thesis was the development of novel, energy-efficient receivers for a technology platform that was still in development. Many of the early assumptions on photodiode performance were challenged and proven incorrect. Figure 3-3 shows measured data from one of the first working photodiode devices in the 45nm SOI process. The photodiode is made using SiGe and is tested at a

wavelength of 1180nm. From the plot, we can see that the optical photocurrent decreases significantly with reverse bias, rendering our initial assumption incorrect. Unless otherwise stated, the receiver designs use the third photodiode configuration in Figure 3-2. However, we note that this assumption is worth re-visiting in future chips, where the photodetector device may have changed (in fact, the most recent generation of the detector, in EOS18, has substantial responsivity at zero-bias). The reason is that with the photodetector placed differentially (and therefore hopefully with little dark current), any issues with matching the reference path go away. A low-dark-current detector reduces the motivation to move back to the first configuration in Figure 3-2.

3.2 Optical Data Receiver: Latching Sense Amplifier

With an understanding of how the photodiodes are designed, we move on with the design of the first receiver topology in this work: a current-detecting, latching sense-amplifier (LSA) [28]. The main principle behind this design is that the latch itself is metastable, and the introduction of a small photocurrent in one of its branches will provide enough imbalance to render a correct bit decision. What makes this topology possible at high data-rates in the first place is the fact that the photodiode capacitances in monolithic integration are relatively small, and therefore should not significantly alter the operation of the latch. Furthermore, by connecting the photodiodes directly to the branches of the latch, we obviate the need for a TIA stage that consumes constant current (hurting energy-cost at lower data rates). In a many-core application, the data receiver could be instantiated several hundred times, and therefore its aggregate energy-cost grows more significant.

3.2.1 Architecture

One variant of the receiver architecture is shown in Figure 3-4. The photodiode is connected differentially across the LSA, followed by a dynamic-to-static (DS) converter and an on-chip high-speed digital testing backend. The receiver operates in two clock phases, receiving one bit per clock period. In a second variation, the optically-connected photodiode is instantiated between the input node and a reverse-bias voltage. A second, non-optically-connected photodiode is connected to the receiver's negative (reference) terminal in order to match capacitance and dark current. This is the preferred topology.

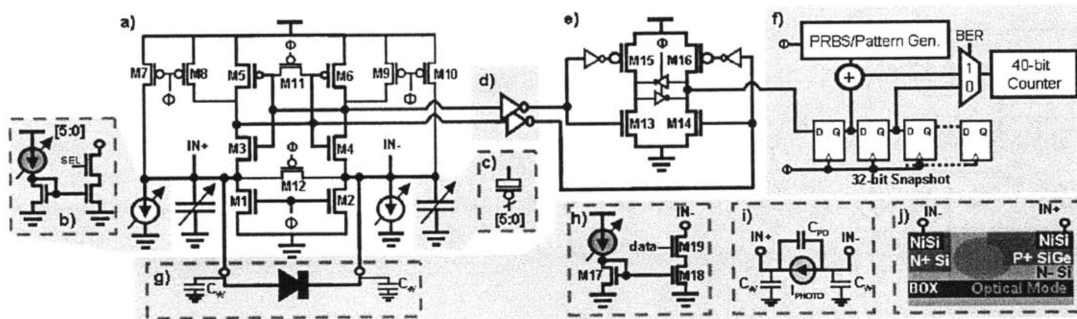

Figure 3-4: Optical data receiver architecture. The LSA (a) is followed by an output buffer stage (d) and dynamic-to-static converter (e), before being fed into the digital backend infrastructre (f). The chip has receivers connected to integrated PDs (g) or electrical diode-emulation circuits (h). The simulation model is shown in (i). A cross-section of the implemented PD and optica mode is shown in (j).

In the photodiode (Figure 3-4g,j) we make use of P+ SiGe, which is integrated in the SOI process for PMOS strain engineering and is suitable for optical absorption in the near-IR range [47]. The photodiode is extremely compact and has an estimated capacitance of 10 fF. Early in the project we assumed that since the photodiode is not transit-time limited, increasing the reverse bias does not increase the speed of the device, but will increase the dark current. For this reason we started out with the differential photodiode placement.

The LSA (Figure 3-4a) senses the differential photocurrent and makes a bit decision. During the reset phase (Φ=0), the LSA's nodes pre-charge high. During the

decision phase (Φ=1), the two branches, $M_{1,3,5}$ and $M_{2,4,6}$, discharge. If an *optical-1* is received, photocurrent flows from node IN- to IN+, slowing the discharge of branch $M_{1,3,5}$ and causing it to latch *high*. Otherwise, imbalance programmed through offset compensation causes branch $M_{1,3,5}$ to latch *low*. Without the programmed imbalance, an *optical-0* will cause the branches to discharge at the same rate, resulting in a random bit decision.

The LSA transistors are carefully sized according to [48] in order to adjust the sampling aperture (time resolution) of the receiver. In particular, transistors $M_{3,4}$ are sized large relative to $M_{5,6}$. This lowers the trip-point voltage of the cross-coupled inverters and ensures that they do not activate too early, which would increase the noise bandwidth of the LSA. Offset compensation is implemented as programmable current-steering (Figure 3-4b) and capacitive (Figure 3-4c) DACs [32], for coarse- and fine-compensation, respectively.

Figure 3-4h shows a diode-emulation circuit that is used to characterize the receiver's performance when decoupled from the optical devices. When the input *data* is *1*, the circuit pulls current from IN-, emulating the photocurrent sourced from that node. A *0*-bit sources no current. The diode-emulation circuit is driven by a pattern generator on a separate, programmable clock phase from the rest of the receiver.

The output of the LSA is buffered (Figure 3-4d) to isolate the LSA decision nodes from the data-dependent capacitance looking into the DS (Figure 3-4e). The bits stored in the DS are fed into the on-chip digital test backend for *in situ* processing (Figure 3-4). The backend, consisting of synthesized PRBS and pattern generators, snapshots, and counters, gathers bit-error-rate and receiver decision threshold data *in situ* and exports only the collected statistics off-chip.

3.2.2 Effect of Capacitance on Receiver Settling

To provide qualitative analysis of the impact of parasitic capacitances and operation frequency on the receiver decision settling time, an equivalent model of the LSA is shown in Figure 3-5. The model is developed on the receiver with the photodiode instantiated differentially between the inputs, but the concepts are easily adapted for

other photodiode configurations.

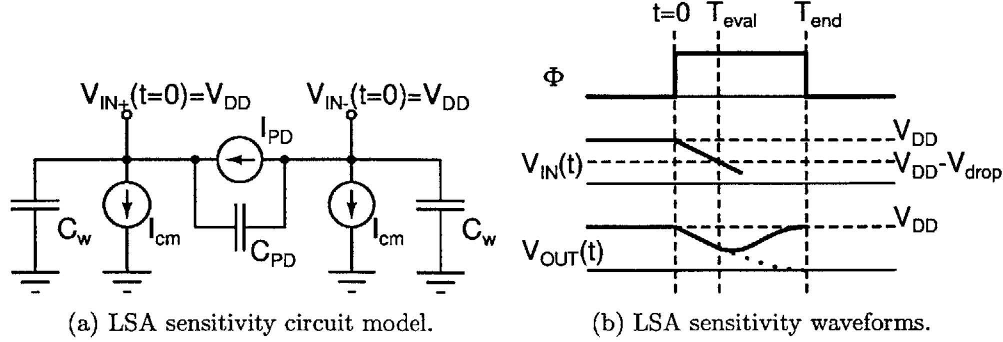

(a) LSA sensitivity circuit model.

(b) LSA sensitivity waveforms.

Figure 3-5: LSA sensitivity model.

Figure 3-5a shows the input nodes at the end of LSA reset ($t = 0$), pre-charged *high*. I_{cm} models $M_{1,2}$ pulling down on the input nodes until cross-coupled inverters M_{3-6} turn on. C_w represents the wiring capacitance from the photodiode to the receiver. The model divides the decision phase into two steps: integration, and evaluation (Figure 3-5b). During the integration phase of duration T_{eval} (Equation 3.1), the photocurrent is integrated across $C_{int} = C_{PD} + C_w/2$, resulting in a voltage difference, $V_{diff} = V_{IN+} - V_{IN-}$, at the onset of evaluation (Equation 3.3).

$$T_{eval} = \frac{C_w V_{drop}}{I_{cm}} \tag{3.1}$$

$$V_{diff} = \frac{I_{PD} T_{eval}}{C_{int}} \tag{3.2}$$

$$= \frac{I_{PD} C_w V_{drop}}{C_{int} I_{cm}} \tag{3.3}$$

$$V_{out} = A V_{diff} e^{\frac{T_{end} - T_{eval}}{\tau}} \tag{3.4}$$

$$I_{PD} = \frac{V_{out} C_{int} I_{cm}}{A C_w V_{drop}} e^{-\frac{T_{end}}{\tau}} e^{C_w \frac{V_{drop}}{I_{cm} \tau}} \tag{3.5}$$

The output voltage of the LSA is related to the input voltage difference, V_{diff}, through a proportionality constant, A. During the evaluation phase, V_{out} regenerates exponentially until T_{end} according to Equation 3.4. Rearranging the formulas, the current-sensitivity of the receiver can be expressed by Equation 3.5.

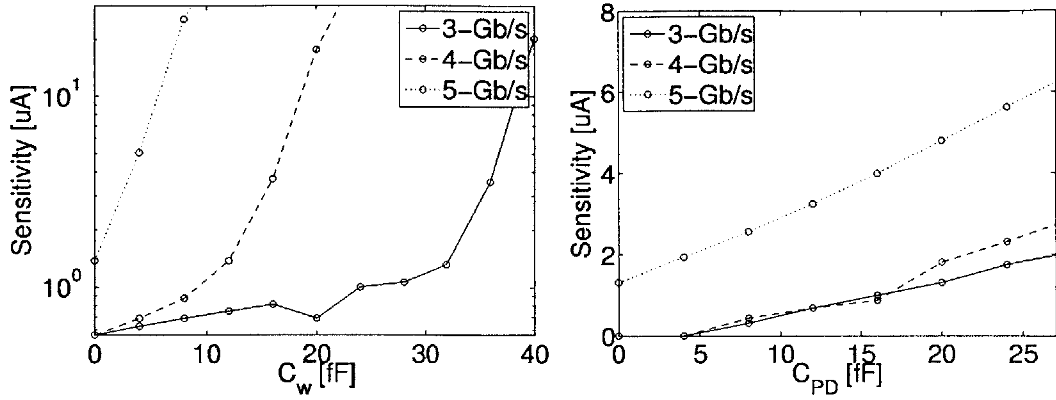

(a) Effect of C_w on sensitivity. $C_{PD} = 10fF$. (b) Effect of C_{PD} on sensitivity. $C_w = 2.5fF$.

Figure 3-6: LSA sensitivity model.

Figure 3-6a shows through noiseless extracted simulation that for high data-rates where the exponential is not completely settling, wire capacitance, C_w, delays the onset of evaluation, shortening the evaluation time and therefore demanding exponentially more input photocurrent. Sensitivity is computed based on an output voltage settling constraint. Note that this is different than the sensitivity definition used in Chapter 4, which is based more directly on the signal strength required to obtain a particular bit-error-rate. The proposed topology may suffer in scenarios where a second die provides the optical transport layer, necessitating TSV or microsolder bumps where C_w may increase above 20 fF [49]. Figure 3-6a shows that for C_w in this range and data rates above 4 Gbps, the sensitivity becomes prohibitively poor. As our photodiode was implemented on the same die as the receiver, the low-metal-layer routing between the photodiode and the receiver results in a small C_w, $\approx$ 2.5 fF based on slightly more than $10\mu m$ of wire at $0.2\text{fF}/\mu m$, exploiting the benefits of monolithic integration. Figure 3-6b shows that the photodiode capacitance, C_{PD}, reduces V_{diff} linearly, demanding only proportionally more photocurrent (Figure 3-6b).

3.2.3 Sensitivity

In addition to an output voltage settling-time constraint, it is critical to evaluate the impact of noise and mismatch on the sensitivity of the receiver. We compute the minimum input current signal from a BER requirement as we did in Section 2.1.2.

The minimum input signal required for the exponential to evaluate to the rails is given by $i_{sense} = V_{DD}Ge^{-(T_{bit}-T_{eval})/\tau}$. The time constant of the exponential term, τ, and conductance, G, are measured in simulation. As in our equivalent model, as the end of the bit time starts to approach T_{eval}, the receiver's sensitivity degrades exponentially.

Mismatch in the differential latching receiver also leads to a threshold offset. In order to avoid large latch sizing, which results in increased power, we employ offset compensation circuitry in the form of a 6-bit capacitive DAC. The threshold offset is measured through Monte Carlo simulation. The residual threshold is then given by $i_{OS,res} = i_{OS}/2^6$, found to be negligably small.

The circuit noise was computed in a transient noise simulation by sweeping the receiver's input photocurrent threshold and recording the decision statistics. The resulting input-referred noise cumulative distribution function as a function of the LSA's decision threshold is shown in Figure 3-7, depicting a standard deviation of $\approx 1\mu A$.

Photodiode shot noise was computed using the receiver's sampling bandwidth [50]. Supply noise sources were ignored and no additional margin was added.

Figure 3-8 shows the receiver's predicted sensitivity as a function of data-rate for $C_{PD} = 10fF$ and $C_w = 5fF$. We can see that as the data-rate increases, T_{end} in our model decreases, demanding exponentialy more input photocurrent starting around 5 Gb/s.

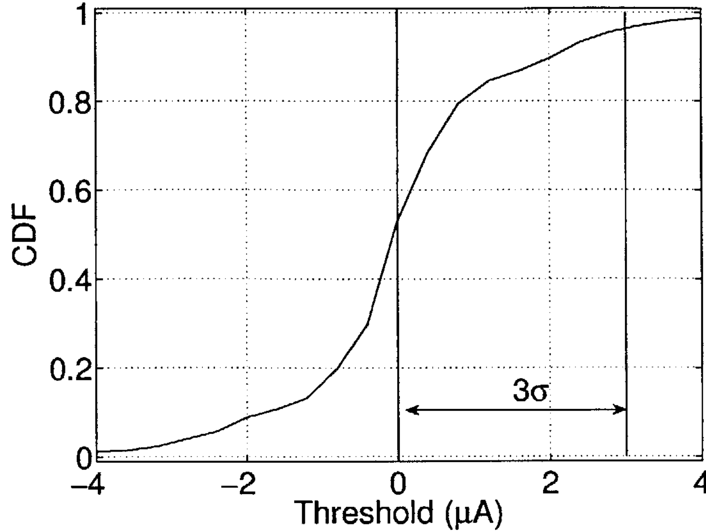

Figure 3-7: Receiver circuit noise CDF.

3.3 Optical Data Receiver: Transimpedance Amplifier

While the LSA is designed with energy-cost in mind, we also implement a more traditional TIA-based receiver to serve as a benchmark. The main disadvantage of this second receiver is that in addition to a clocked comparator, it requires a TIA front-end at burns static power. The analog front-end also may not scale well with technology. However, the TIA-based receiver does have several advantages. Firstly, it is a very standard receiver, and its operation and governing equations are well understood. In the context of this work as a whole, this is an important characteristic. We are developing a new technology - integrated photonic interconnects - with new photodiode devices still in development and not completely understood. It's important to minimize energy-cost (as we did with the LSA), but also important to hold some of the variables constant in order to characterize the new devices. The photodiode is attached to the TIA's input node, resulting in a reverse bias of roughly half of the supply voltage. In contrast to the LSA receiver, the photodiode's reverse bias does not dramatically change during the course of operation.

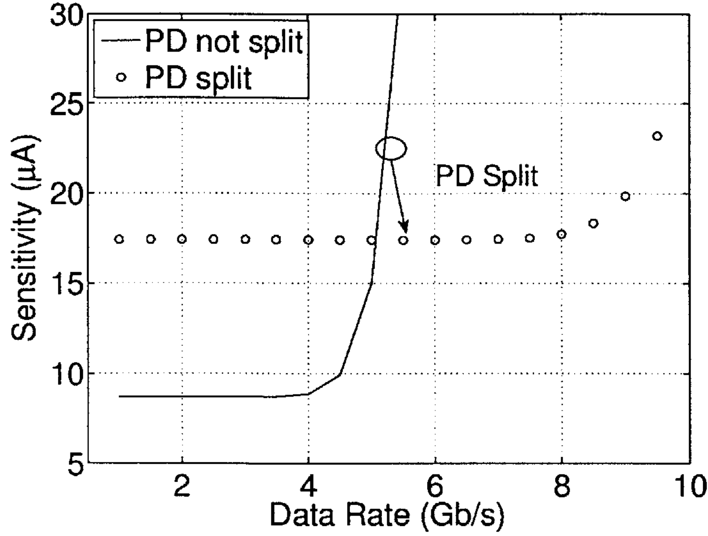

Figure 3-8: Receiver sensitivity and PD split. C_{PD}=10fF, $C_w = 5fF$.

3.3.1 TIA Front-end

Figure 3-9 shows a schematic of the TIA front-end. It consists of an inverter with an impedance placed in feedback between the input and output nodes. The impedance can be implemented as a resistor, or by using the effective resistance between the drain and source of a feedback transistor. If a transistor is used, an NMOS and PMOS should be used in parallel to improve the linearity of the amplification.

As with other amplifiers, the key points of study for the TIA front-end are its bandwidth and transimpedance gain. The transimpedance gain is given by Equation 2.7, which as with Figure 2-7a shows that the gain is roughly equal to the feedback resistance, R_f. Where the TIA is advantageous is in the fact that the impedance looking into the input node is decreased through the negative feedback. Even in monolithic integration we estimate that the photodiode has a capacitance of 10fF to 20fF. When combined with the gate capacitance from the TIA itself (≈3fF) and the wiring from the photodiode to the receiver's input, the input pole is will still dominate the circuit's bandwidth.

The performance of the TIA front-end in the 45nm SOI logic process is summarized in Table 3.1. The impedance of the feedback transistors is roughly equal to the TIA's

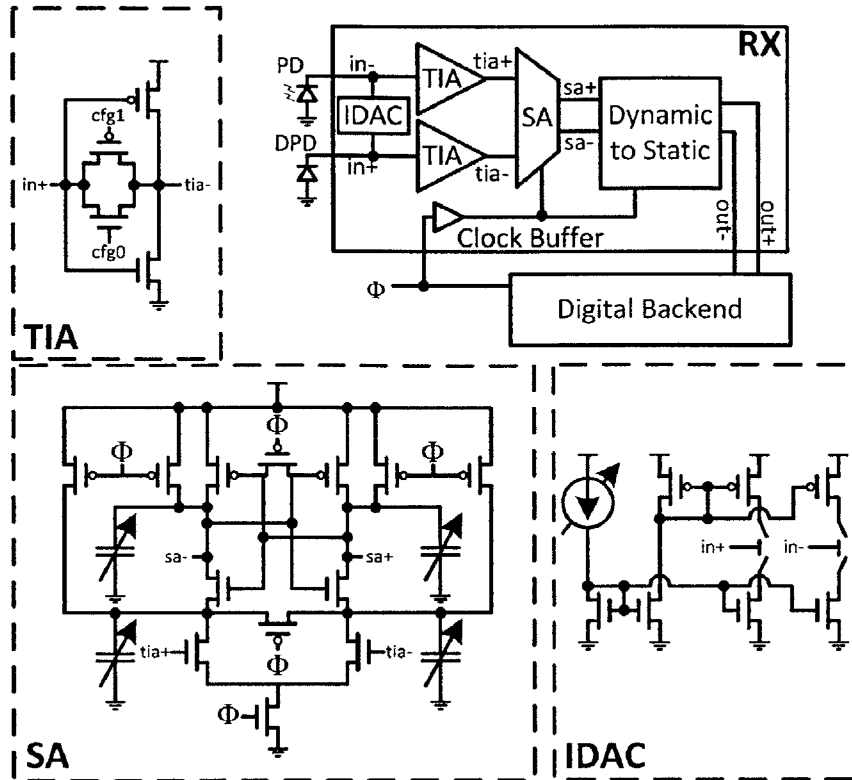

Figure 3-9: TIA optical receiver (EOS18).

transimpedance gain.

Feedback	NMOS	PMOS	NMOS and PMOS
Gain (kΩ)	8.6	3.7	2.4
Bandwidth (GHz)	5.4	7.9	8.9
Linearity (ENOB)	5.1	5.48	7.8
R_{in} (kΩ)	1.28	0.88	0.77

Table 3.1: TIA front-end performance in 45nm SOI (simulation).

For the memory-process chip, the feedback impedance was implemented with resistors (Figure 3-10). The sizing of transistors M_N and M_P again follows from Equations 2.6, 2.7, 2.8. As with the TIA in the logic platform, two gain settings were implemented to facilitate a study of gain and sensitivity, and provide more configuration options for use with a photodiode whose performance was still unknown at the time of design. The performace of that frontend is given in Table 3.2. The capaci-

tance at the input node due to the TIA is $\approx$ 15fF - a considerable load. As a point of reference, a 40fF capacitance at the input node combined with a $1k\Omega$ impedance looking into the TIA results in a 3dB bandwidth of 4GHz, able to support a data-rate of little more than 5.5Gb/s.

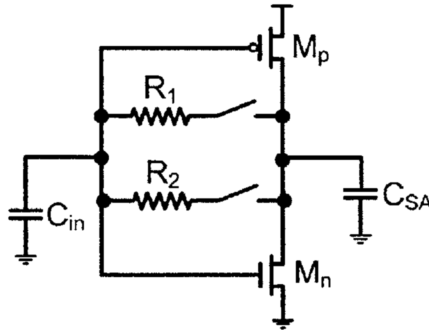

Figure 3-10: TIA front-end in $0.180\mu m$ bulk memory process.

Feedback	R_1	R_2	$R_1 \| R_2$
Resistance (kΩ)	2.8	8.4	2.1
Gain (kΩ)	2.8	8.1	1.9
Bandwidth (GHz)	4.4	2.4	4.3
R_{in} **(kΩ)**	0.6	0.9	0.5

Table 3.2: TIA front-end performance in $0.180\mu m$ bulk (simulation).

Figure 3-9 also shows a block labeled IDAC. The purpose of this block is to provide a programmable current-DAC that can tune out potentially large dark currents generated by the photodiodes (even when not optically-connected). In practice, this current-DAC has also been used to improve the bias points of the TIA and sense-amplifier.

3.3.2 Sense Amplifier

Once the TIA has converted the input photocurrent to a voltage signal (with as much gain as possible while satisfying bandwidth requirements), the voltage signal is regenerated to a fully-digital signal by a traditional sense-amplifier, shown in Figure 3-9. As with the LSA detailed earlier in this chapter, the sense-amplifier is employed

here due to its energy-efficient operation facilitated through the presence of a receiver-side clock.

The design and sizing of the sense amplifier is similar to that of the LSA in Section 3.2.1, with the sizing again following [48]. Furthermore, most of the heavy-lifting has already been done by the TIA front-end. The signal has been amplified, and just has to be converted to a rail-to-rail signal. We do, however, note that the equivalent model developed in Section 3.2.2 applies to a voltage-sense-amplifier as well: if an increasing data-rate results in the reset signal decreasing the latch's evaluation time, then an exponentially-larger input signal will be required in order to compensate. In fact, we will see this in Chapter 4 where this design is implemented in a slow, $0.180\mu m$ bulk CMOS memory process.

3.4 GSG Transimpedance Amplifier Test Site

One of the major themes of this thesis is experimentation and iteration. When we get chips back several months after submission, we often need to quickly evaluate whether the devices perform as intended (with the circuits), as there may be another tapeout where we have an opportunity to make corrections. To mitigate the tight timeline, we developed a two-chip pipeline. On each run, one chip contains an array of devices, while the other contains the larger digital circuit system. As the device chips became more sophisticated, we began implementing very simple circuits that can be tested using only GSG- and dc-probes. The advantage here is that we can quickly try different processing steps on the die and then test them immediately, rather than waiting for costly bonding and packaging.

In order to evaluate the photodiode performance, we implemented a simple TIA-based receiver (Figure 3-11) that is easily configurable and able to be tested with GSG and DC probes and a single optical fiber. The design is similar [51].

The TIA is configurable through two current mirrors that either add or subtract current from the input node, effectively raising or lowering the threshold of the receiver. Figure 3-12 shows the gain and phase margins of each of the two TIA stages.

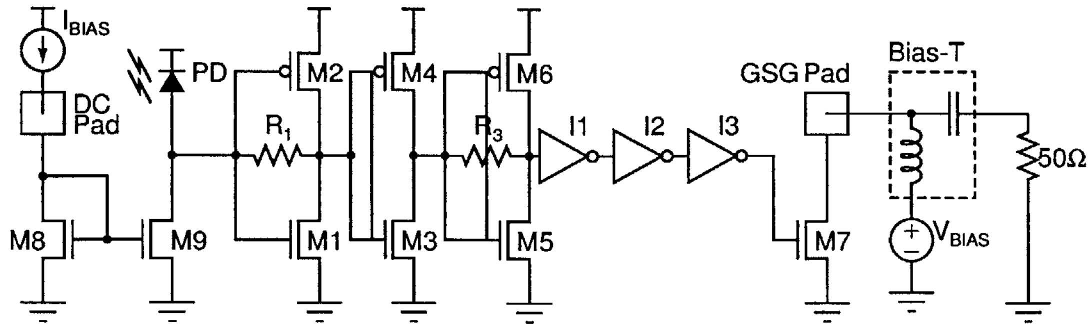

Figure 3-11: EOS16 TIA optical receiver.

The plots demonstrated the stability of the receiver stages (phase margin does not decrease by 180-degrees before unity-gain is reached).

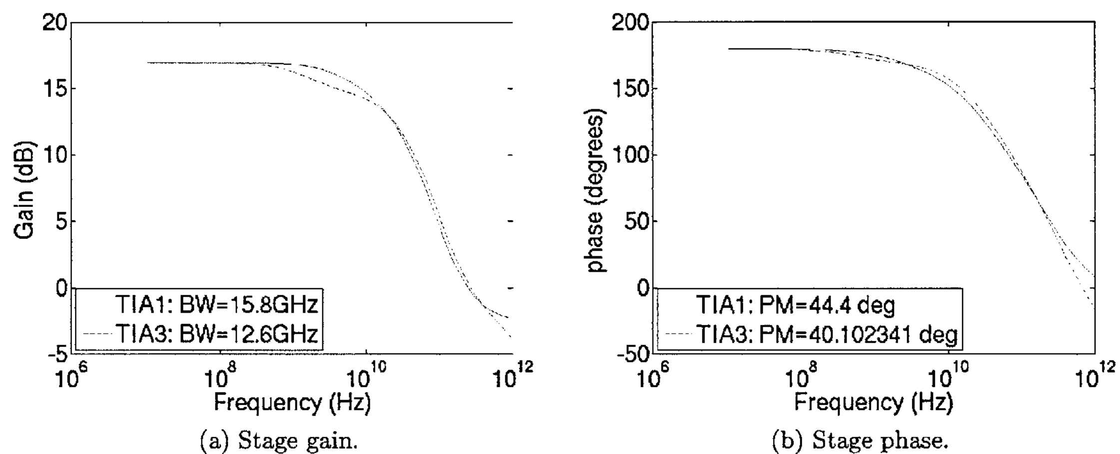

(a) Stage gain.

(b) Stage phase.

Figure 3-12: GSG TIA stage gain and phase (simulation).

Figure 3-13 shows the performance of the GSG TIA. The total gain of the TIA stages is $12k\Omega$. For the reader that may want to use the receiver in an analog application such as an ADC, the linearity of the TIA stages is measured to an effective number of bits (ENOB) of 6.8 for an input sinusoid at 5GHz with amplitude of $5\mu A$. Note that this is not the aggregate lineary of the TIA with the buffers, as the purpose of the buffers is to rail the signal in order to drive the GSG pad itself. Total simulated power consumption is 5mW, with the output dominating (Figure 3-13b).

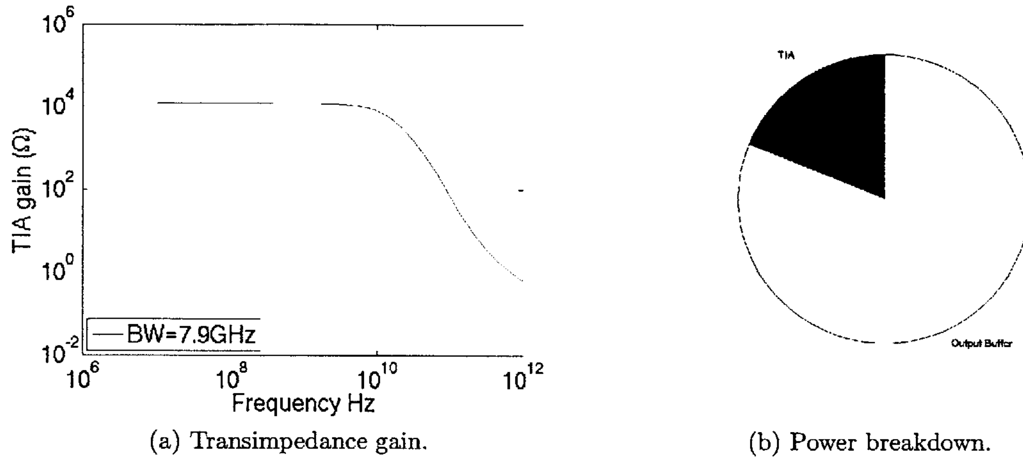

(a) Transimpedance gain.

(b) Power breakdown.

Figure 3-13: GSG TIA performance (simulation).

3.5 Photodiode Parameter Impact on Receiver Suitability

Photodetectors can be characterized by several technological parameters, such as bandwidth, responsivity, dark current, and capacitance. We briefly touch on the suitability of each of the receivers based on some of these parameters.

Both TIA and LSA receivers require sufficient bandwidth from the photodetector in order to detect the bit correctly. This is fundamental to system design and not unique to optical systems. Both designs also benefit from increased photodiode responsivity, which increases the photocurrent for a fixed optical power and therefore improves the receiver's optical sensitivity.

Dark current and photodiode capacitance can have a stronger impact on the choice of topology. Consider first the LSA, where we saw that the input capacitance can affect the large-signal operation of the circuit. A large capitance slows its evaluation and causes the evaluation phase to run into the end of the bit-time, exponentially decreasing sensitivity. In this high-capacitance scenario (such as hybrid or discrete integration), the LSA is a poor choice. The TIA on the other hand, can trade gain for bandwidth, and accomodate a larger input capacitance. Additional gain stages

may be required after the front-end (consuming more power), but the topology has been shown to accomodate larger capacitances.

Dark current also has different effects on the two receivers. With a dummy-photodiode implemented, the two branches of the LSA will simply discharge more quickly in the presence of a large dark current. The TIA, however, has a more critical problem. While a dummy-photodiode can provide a suitable reference for a TIA-based receiver, the large dark current may drastically affect the bias point of the TIA. For example, a $100\mu A$ dark current (which is in the range that we have experimentally observed), being fed into a TIA front-end with a transimpedance gain of $10k\Omega$ results in a difference of 1V across the TIA, which is well outside of its linear region of operation. While offset-compensation and dark current tuning circuitry has been applied to both the TIA and the LSA, we note that the TIA is particularly vulnerable not only to dark current mismatches, but also large dark current magnitudes.

3.6 Photodiode Splitting

Understanding the impact of the settling time and noise, it is possible to further optimize the receiver, leveraging the monolithic integration once again for a closer interaction between the photodiode and the receiver circuit.

Figure 3-14 summarizes the different ways that the monolithically-integrated-photonic designer can adapt the photodiode design to create different receiver-system functionality. Figure 3-14a shows a standard, single-ended configuration with a dummy-diode, detailed earlier in this chapter. Figure 3-14b shows how DDR functionality can be enabled, with two copies of a receiver circuit operating on two different clock phases. Figure 3-14c shows how to connect two receiver to a split-photodiode in order to implement at DFE. Figure 3-14d shows another method that effectively splits the photodiode. Here, a waveguide splits the optical power into two and directs the signal to two separate photodiodes. While this technique is likely less area-efficient, it can be used with a photodiode design that may not be easily inter-digitated.

Figure 3-14: Single-rate, split-diode, and DFE photodiode configurations.

3.6.1 Double Data Rate

The limiting sensitivity factor at higher data rates is the settling-time term (i_{sense} in Equation 2.4). By operating the receiver at half the rate, we can double the value of T_{end}, giving the exponential regeneration phase much more time to settle. In order to keep the data-rate on the channel the same, we need two receivers and a DEMUX. Since monolithic integration affords us a high degree of control over the design of the photodiode, we can simply interdigitate metal contacts to break it into two separate photodiodes, each connected to a separate receiver. While one receiver is integrating and evaluating the input signal, the other is resetting. As a result of the photodiode splitting, only half of the total photocurrent will go to each receiver, requiring 2X the laser power, but this is still better than a higher exponential factor.

Figure 3-8 plots the sensitivity of the receiver for both photodiode split and unsplit cases. By splitting the photodiode and doubling T_{end}, the exponential term begins to dominate only at much higher data rates. A factor of 2 is applied to the total sensitivity computed in Equation 2.4, reflecting that each receiver only gets half of the photocurrent. Should higher rates be desired, the photodiode can be split again, further increasing T_{end} and suppressing the exponetial i_{sense}. Careful partitioning of the photodiode fingers ensures that each receiver's photodiode gets a roughly equivalent share of the incident optical power.

This split-photodiode enables double-data-rate (DDR) receivers which are very useful in parallel source-forwarded links, where a data-pattern of 101010 on one of the transmitted wavelength-channels can be used as a receive clock and directly applied to all of the DDR receivers.

3.6.2 Decision Feedback Equalization

While the split-photodiode is presented in the context of enabling DDR, another interesting application exists in decision-feedback-equalization (DFE). DFE can be used to mitigate bandwidth limitations, such as those at the input to the receiver or even at the modulator device or driver. A DFE works by adjusting the decision

threshold of the receiver based on the previous bit. For example, if the most recently received bit is an optical-0, then the decision threshold is moved slightly above the optical-0 threshold level. If the most recently received bit is an optical-1, then the decision threshold is moved slighlty below the optical-1 threshold level. The result is that the receiver is able to easily resolve both the case where the next bit is the same as the most recently received on and the case where the next bit is different. A DFE decreases the amount that a signal needs to change in order to have the bit-decision change, which helps bandwidth-limited signals that transition more slowly.

A DFE can operate by having the value of the most recently-received bit dictate the decision threshold of the receiver. By using the split-photodiode technique, however, we can set up two receivers, each with a decision threshold set for one of the two cases described (Figure 3-14). The receivers operate on the same clock phase, and both make a bit decision each clock period. The receivers are fed into a MUX, which selects which receiver-decision to output based on the previous bit value. The advantage here is that low-bandwidths earlier in the link can be mitigated. The DFE implementation comes at the cost of reduced data-rate and sensitivity though, since a single optical signal is driving two photodiodes and receivers.

3.7 Summary

This chapter presented designs and techniques for monolithically-integrated optical receivers. The design of three main receivers were details. First, we examined a compact, all-digital current-sense-amplifier with aggressive energy-cost. An equivalent circuit model of the receiver was developed in order to provide insight on the effect of photodiode and wiring capacitance. The second receiver topology detailed is a more conservative TIA-based design. The TIA-based receiver suffers from an increased energy-cost due to the addition of an analog front-end, but provides a more stable interface with the photodiode. A simple TIA used for photodiode characterization was also shown.

After the three receiver topologies were examined, the equivalent circuit model

of the LSA was used to develop a split-photodiode technique enabled through monolithic integration. The technique was shown to enable double-data-rate operation to achieve higher data-rates, and decision-feedback-equalization to mitigate the effects of bandwidth bottlenecks earlier in the channel.

Chapter 4

EOS Photonic Technology Platform

This thesis has so far described how an integrated-photonic system should be put together from a system-level standpoint. Energy-costs and throughputs were taken into consideration to develop a methodology for determining the optimal data-rate for each wavelength-channel. The analysis also facilitated a fair comparison between electrical and optical IO, demonstrating that optical IO achieves bandwidth densities better than two orders of magnitude over its electrical counterpart.

After developing a systems-level perspective, we focused on the design of one of the critical components in the optical link: the optical data receiver. We explored two main receiver topologies - a current-sensing latching receiver and a more traditional TIA-based receiver - each of which made use of the presenece of a receive-side clock and sense-amplifier in order to efficiently regenerate the signal. A third receiver was developed for ease of testing. We developed multiple techniques surrounding a split-photodiode design. These techniques are particularly relevant here as they are only possible in a tightly-integrated environment, where a circuit designer has unprecedented control over the design of the photodiode.

In the following sections we detail the integrated photonic technology platform upon which work from this thesis is built and verified. We describe an overview of the chips, the experimental setup, and the key results throughout several generations

of testing that ultimately enabled record-setting optical receivers and link demonstrations.

4.1 EOS Chip Architecture

The EOS integrated photonic technology platform is a series of chips developed in order to test a wide range of photonic components. Through the platform, we have characterized a host of individual devices and circuits, as well as demonstrated WDM chip-to-chip links. While this thesis covers several generations of EOS chips, we'll detail the overall architecture through EOS18 - a recent test chip. Figure 4-1 shows an overview of the EOS18 system chip. The chip has several functional areas, of which we'll focus on the array of test cells. Each cell is an independent test site with its own modulator driver circuit, receiver circuit, and high-speed digital backend. The circuits are connected to a number of different photonic component variants, enabling the engineer to quickly measure a variety of devices through simply reprogramming the chip and moving the optical fiber inputs if necessary.

On the TX side, the high-speed digital backend consists of dual 31-bit PRBS generators that are muxed inside the modulator circuit block in order to create a DDR modulated signal. The PRBS generators also drive on-chip snapshots that can be used to double-check the pattern that is being fed to the modulator. On the RX side, the receiver circuit drives two snapshots with two data-patterns that have already been demuxed. The receiver output is also fed into a high-speed counter that can be programmed to either count the output bits or the bit-errors (through comparison with an RX-side self-seeding PRBS). The digital backend is very reconfigurable and enables measurements to be performed *in situ*, with only the statistics exported off-chip.

The EOS technology platform was also implemented in a memory process through collaboration with Micron Technologies. Figure 4-2 shows the chip architecture of one of the memory chips fabricated - D1L. From the figure, it is clear that the overall architecture is quite similar to that of the logic EOS chips. One of the key dis-

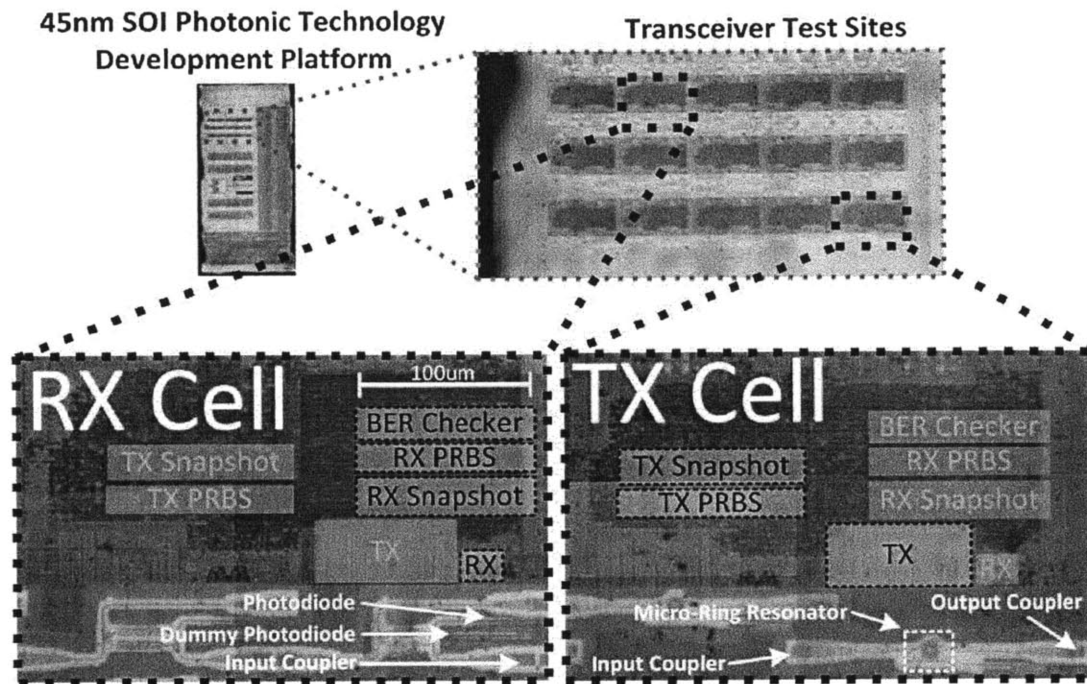

Figure 4-1: EOS platform architecture (EOS18).

tinctions is the fact that these chips were fabricated in a 0.18um process, which is considerably slower than the 45nm SOI logic process. As a result, a higher degree of (de-)serialization (8-to-1 as opposed to 2-to-1) was required in order to get the channel data-rates up to specification. Modifications were also made to the receiver and transmitter.

4.2 Integrated Electronic-Photonic Test Setup

The EOS platform integrates electronic and optical components to the highest degree possible, and as a result demands a unique hybrid test setup that is capable of both electronic and optical testing. The lab testbench must house an EOS platform chip, providing it with the necessary electronic control (low-speed) and clock (high-speed) signals. On the optical side, the testbench must position optical fibers above the chip at various angles while aligning them with micrometer precision. Additionally, due to the experimental nature of the work, the chips must be interchanged

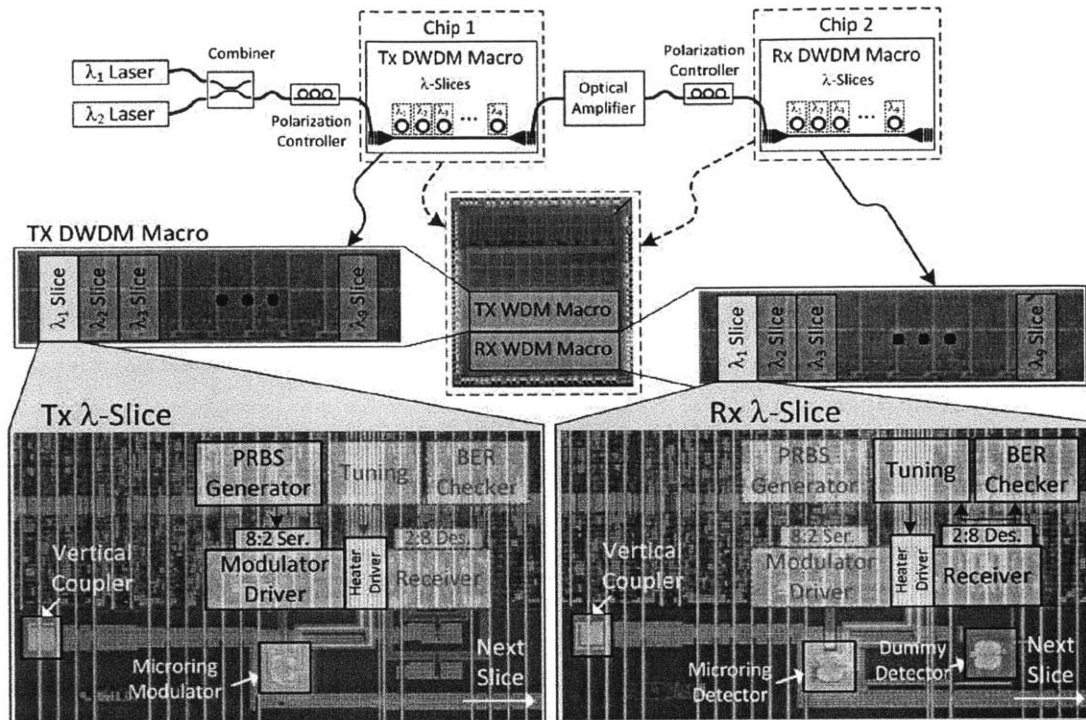

Figure 4-2: Micron platform architecture (D1L).

quickly, allowing different designs, chips, and wafers to be tested efficiently.

Figure 4-3 shows the electronic-photonic testbench. Light from a laser source (not shown) propagates down an FC/PC optical fiber. A polarization controller changes the polarization of the incident light such that maximum coupling is achieved through the on-chip gratings. The polarization-corrected light then connects to either a lensed fiber (for the 45-nm SOI process) or a cleaved fiber (for the $0.180\mu m$ bulk process), aligned to the grating through the Left Fiber Mount. The Right Fiber Mount aligns the output fiber to the output grating. The optical signal can then be fed to a power meter for waveguide-loss measurement, or to an oscilloscope to observe an optical eye-diagram. An IR camera mounted above the chip provides visualization for both fiber alignment and optical absorption throughout the chip.

The EOS platform chip itself is mounted onto a PCB through which the power, ground, electrical control, and high-speed clock signals are routed to the chip package. Much care is taken with the PCB design in order to keep board components from

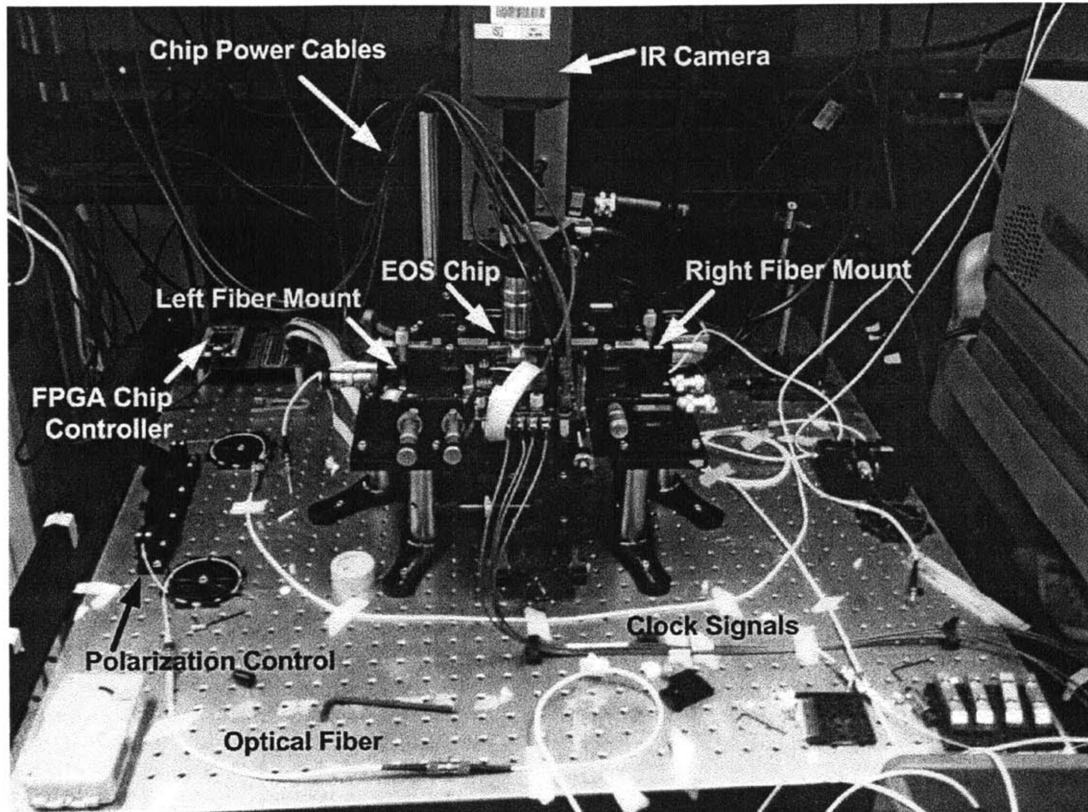

Figure 4-3: EOS platform chip testbench.

interfering with the optical fiber positioning stages. The PCB is firmly mounted, and electrical cabling tied down to avoid the board (and chip) from moving while the fibers are coupled in. Movement could result in the fibers becoming misaligned, or much worse the chip becoming damaged through fiber contact.

Figure 4-4 shows a closer view of an optical fiber coupling into an EOS chip. On the left, we can see an optical fiber mounted to a variable-angle fiber holder being positioned above an EOS chip. The chip die sits in an open-cavity package, mounted to the PCB through a quick-release Zero-Insertion-Force (ZIF) socket. No output coupler is mounted on the Right Fiber Mount in this figure. On the right-hand figure, we can see a picture captured from a side-view camera. The image shows two fibers positioned above an integrated photonic die, and provides the reader with a sense of scale. The optical fibers are roughly $100\mu m$ in diameter, and at this zoom

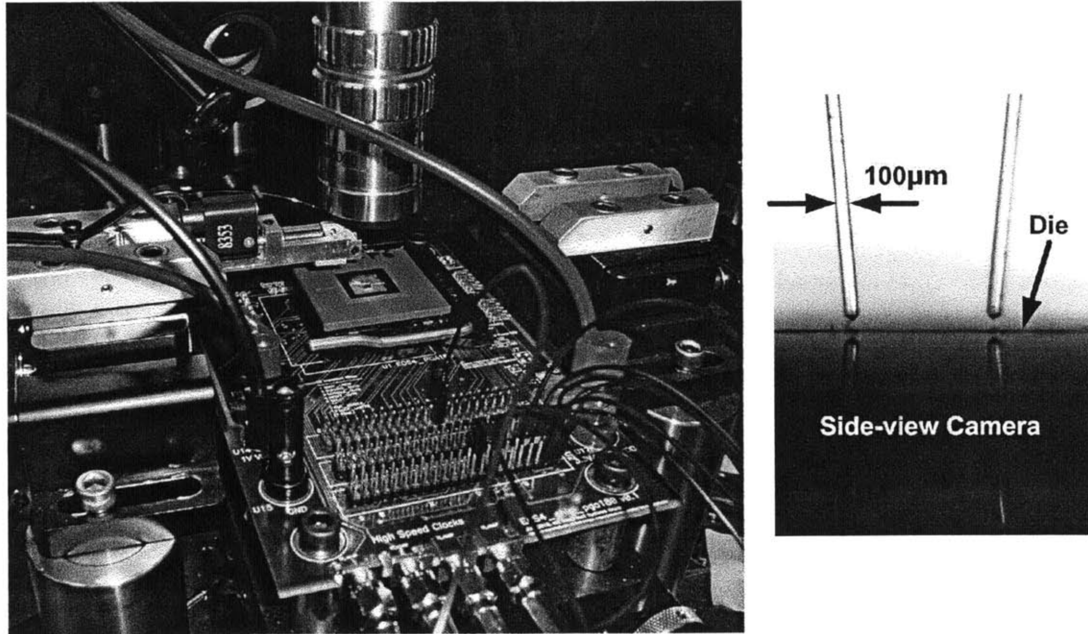

Figure 4-4: Fiber coupling into an EOS chip.

the lensing at the tip of the fiber is visible. Note that grating couplers that do not require lensed fibers can be designed and were in fact used on the memory process chips.

A second optical table was built to serve two purposes. First, it is a highly-portable test station that was used as a demo at ISSCC 2013 [52]. Second, it enabled a chip-to-chip link experiment, where two optical chips could be optically-connected through a fiber patch-cord. Figure 4-5 shows the second table and its configuration with the first table in a link setup.

4.3 Enabling Low Loss Waveguides in 45nm SOI: Substrate Removal Electronic Verification

Significant prior-work has been done in this project to drive down the loss of the waveguides [53], and one of the key challenges has been how to enable low loss waveguiedes while ensuring that electronics remain un-affected. This resulted in

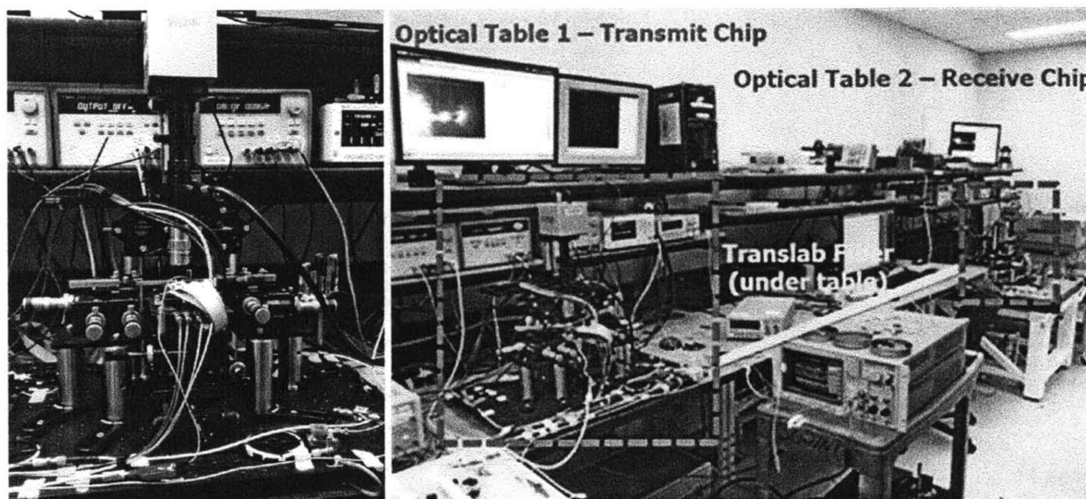

Figure 4-5: Optical Table 2 and link demonstration setup.

more-complex-than-usual packaging and processing strategies that standard circuit designers would not have to consider (for example, we use a fairly expensive ceramic package that is compatible with clean room facilities for post-processing).

Partial-substrate-removal is one of the processes that our team has developed in order to decrease waveguide loss. However, the process must be fine-tuned for each new technology, and this takes time. In order to assess the viability of a full-substrated-removal strategy, we implemented a set of ring-oscillators on the logic chips and tested their performance before and after substrate removal. Figure 4-6 shows the ring oscillator schematic and Figure 4-7 shows the stage delay results on both non-processed (NP) and substrate-transferred (ST) die. Four different oscillator lengths were used (3-, 5-, 7-, and 9-stage oscillators). A series of buffers drive the output NMOS transistor with a 50Ω external load. A spectrum analyzer was used to measure the output frequency.

From the figure, we can see that the not only do the transistors still function after a full substrate removal, but their stage delays are not noticeably affected. While this is a small circuit and may not suffer from the temperature variations that a large, digital design will, it provides enough evidence that more complex circuits will survive the full substrate removal and transfer.

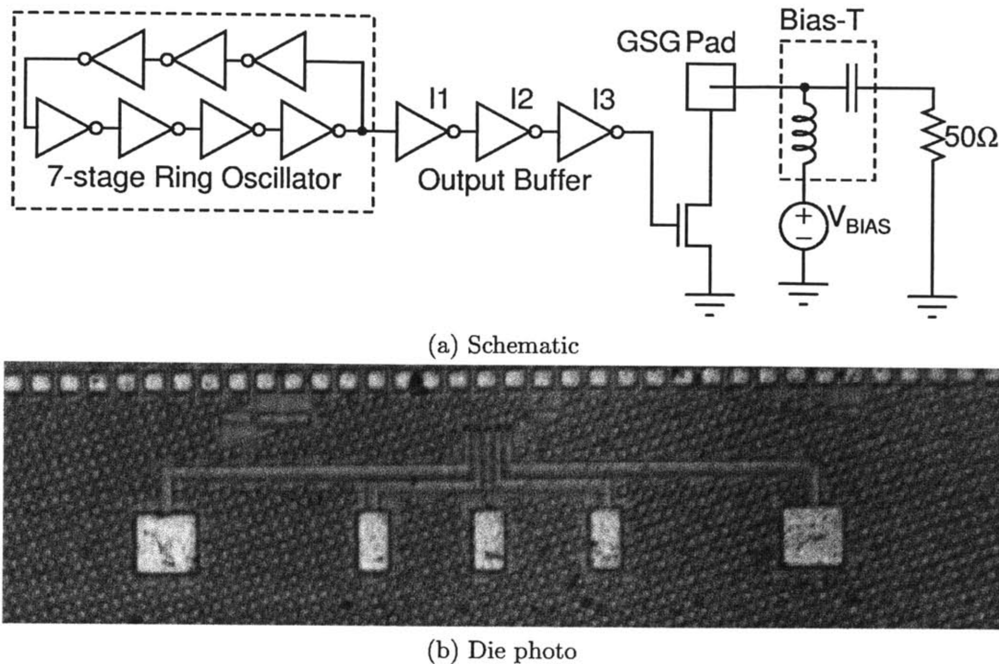

(a) Schematic

(b) Die photo

Figure 4-6: EOS platform 45nm SOI ring oscillator.

4.4 Optical Data Receiver: Logic

With a testing-friendly chip architecture, an electronic-photonic testbench, and a process to enable photonics all in place, we proceed with the testing of the receivers designed in Chapter 3. We begin with the EOS chips implemented in a logic process: commercial 45nm SOI CMOS.

4.4.1 Low Bandwidth Detection (EOS4)

The first optical receiver results obtained from the EOS platform were on the EOS4 test chip, with the LSA receiver detailed in Figure 3-4. The photodiode designs that were implemented with the receiver are shown in Figure 4-8. In Figure 4-8a the photodiode is implemented as an absorption-type receiver where the optical power is absorbed along the length of the device. In order to maximize the photocurrent, the device must be relatively long. As a result, the capacitance is increased and reduces the receiver's sensitivity. Figure 4-8b shows a second option where the detector is

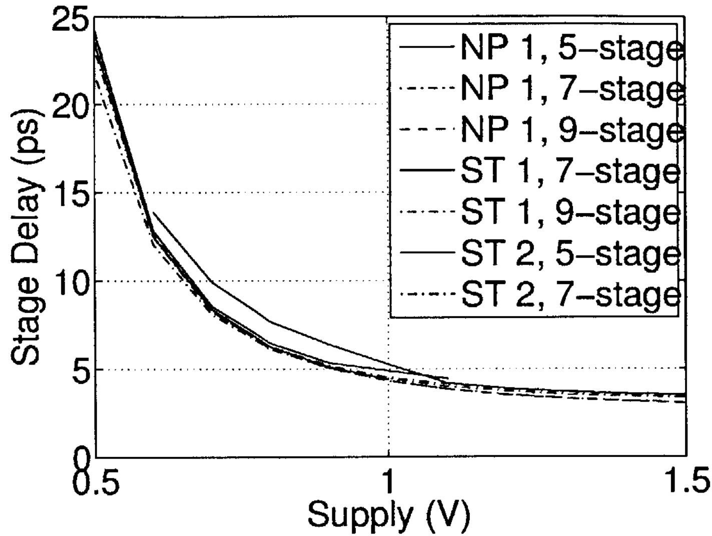

Figure 4-7: EOS platform 45nm SOI stage delay before and after substrate removal and transfer.

integrated into a resonant ring. When the ring is tuned to the particular wavelength channel of the incoming signal, the light becomes confined in the ring. During each round-trip, part of the light is absorbed, allowing for a smaller PD length and therefore less capacitance and better sensitivity. The figure also shows how two PDs can be implemented in the same wavelength-channel, enabling the PD split described.

Figure 4-9 shows two DC photocurrents generated by a 1310-nm wavelength laser, coupled into the chip through a vertical coupler and horizontal waveguide made with front-end body Si. The receiver's threshold is swept using the offset circuitry (Figure 3-4b,c) while recording the output decision statistics. Photocurrent values were de-embedded from DAC settings through simulation. Though the receiver was able to detect photocurrent from the PD, a foundry error in the SiGe mask definition limited the achievable PD bandwidth in this chip.

Figure 4-10a shows the receiver's sensitivity vs date rate for different supply voltages. Sensitivity is measured on a PD-connected receiver (Figure 3-4g,j) as the width of the transition region (Figure 4-9) of an optical-0. As clock frequency increases, sensitivity degrades exponentially as predicted by our model due to the decrease in T_{end}. The sensitivity degrades with reduced supply voltage, as I_{CM} in Equation 3.5

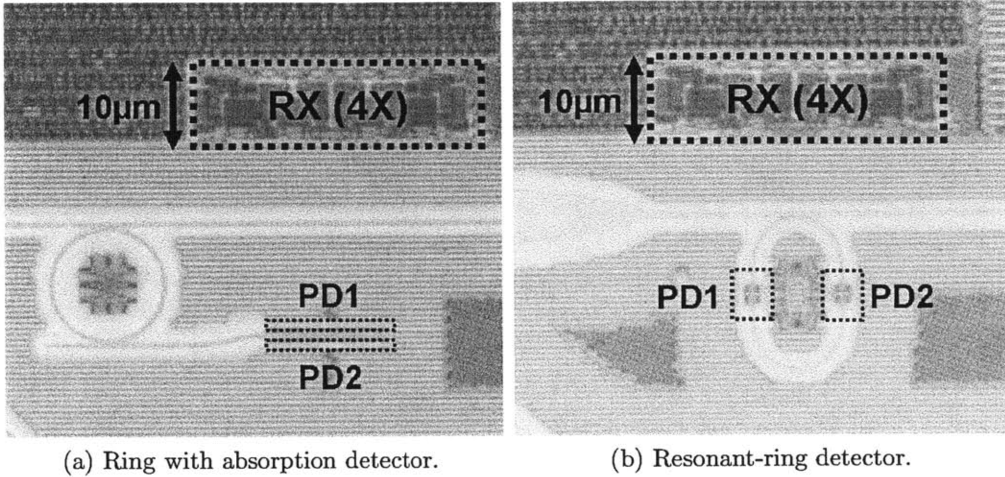

(a) Ring with absorption detector.
(b) Resonant-ring detector.

Figure 4-8: Die photo showing photodetector layout options.

will decrease with reduced V_{DD}, increasing the length of the integration phase, but decreasing the length of the exponential evaluation phase. We note that this is a purely electrical measurement.

Figure 4-10b shows the energy-cost of the receiver. The linearity emphasizes the digital design, with power following $P_{digital} = fCV_{DD}^2$, keeping the receiver energy-cost $\approx$ 50fJ/b across a range of frequencies. The power breakdown is shown in Figure 4-11b. While the latch power is dominant, experimental infrastructure such as capacitive DACs used for link analysis and eye-diagram measurement are shown to be equally expensive. This cost can be reduced in the future. The receiver's current-DAC was not needed and was shut off, consuming no power.

Figure 4-11a shows the bit-error-rate eye diagram of the receiver when configured with a PD-emulation circuit (Figure 3-4h). Clock phase and receiver threshold were swept for a 31-bit PRBS data sequence at 3.5 Gb/s and a supply of 1.1 V, and error statistics were gathered *in situ* using the digital backend. Clock rates above 3.7 GHz caused timing violations in the digital testing backend. A die photo is shown in Figure A-1b. The chip contains 72 test cells that implement combinations of optical modulators and receivers. Each receiver has a circuit area of 108 μm^2 and PD area of 416 μm^2.

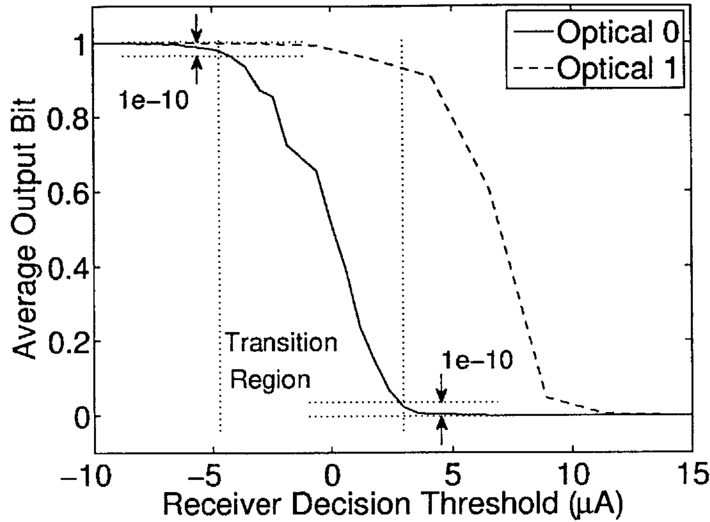

Figure 4-9: The receiver's ability to distinguish between a DC *optical-1* and *optical-0*.

4.4.2 Gigabit Optical Receivers

In order to evaluate the PD performance early on, we implemented a simple TIA-based receiver (Figure 3-11), that is easily configurable and able to be tested with GSG and DC probes and a single optical fiber. In combination with independent photodiode test sites, we search for the best-yielding photodiode designs from the run.

Figure 4-12 shows the photodiode performance from one of the better-performing designs on the EOS16 device chip (Figure A-1i). The figures show several promising features. First of all, the dark-current of the photodiode is very small (less than a μA). While the receiver circuits have been designed to offset large dark currents, the nominal operating point of the circuits is for small dark currents. We can also see that the photodiode exhibits a large photocurrent (more than 10 μA) that should be detectable by reasonably sensitive receiver circuits. We note that the photocurrent generated increases with decreasing wavelength, as the wavelength further enters the absorption range of the SiGe.

The receiver in Figure 3-11 was connected to the more promising photodetectors and an optical clock pattern (1010) coupled into the chip through a grating coupler.

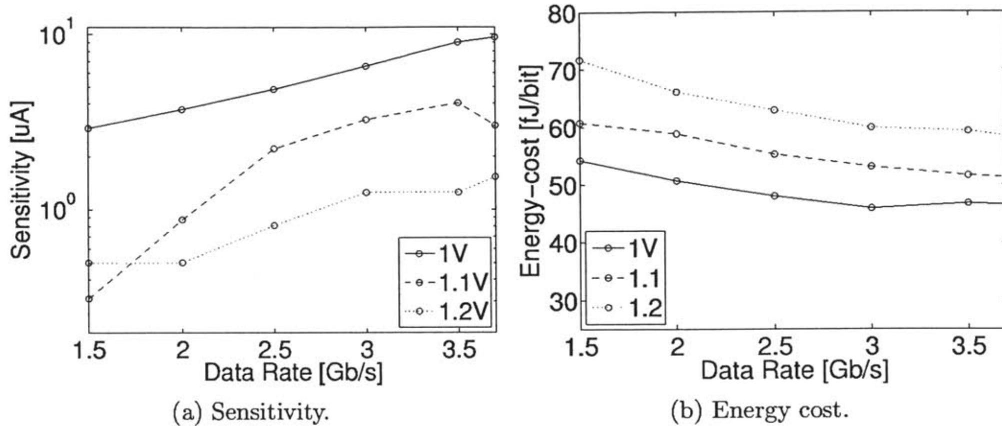

(a) Sensitivity. (b) Energy cost.

Figure 4-10: EOS4 Receiver: Measured performance.

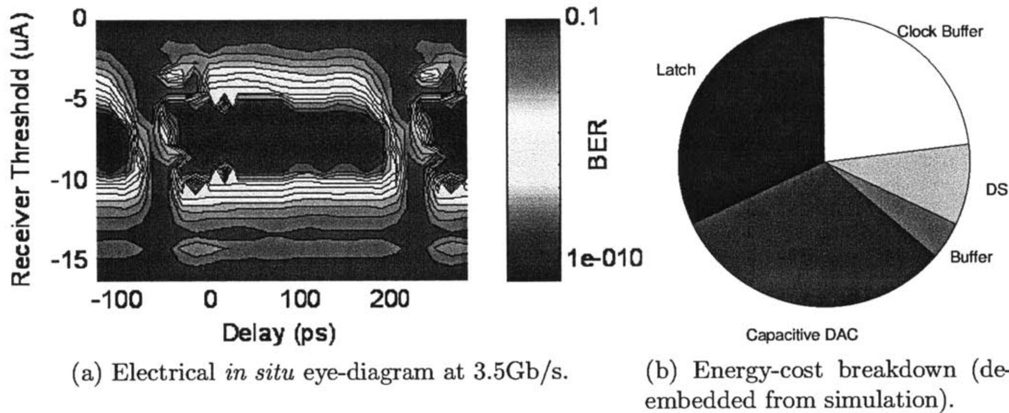

(a) Electrical *in situ* eye-diagram at 3.5Gb/s. (b) Energy-cost breakdown (de-embedded from simulation).

Figure 4-11: EOS4 Receiver: Electrical eye diagram and energy-cost breakdown.

Figure 4-13 shows the input optical signal and the TIA's output electrical signal (measured with a GSG probe), for input frequencies of 15MHz and 500MHz. From the plot, we can see that there are actually two effects contributing to the output waveform. There is the optical signal itself, which is a relatively high-bandwidth signal, but there is also a large, lower-bandwidth thermal response that dies away with higher data-rates.

After a modification to the photodiode designs, the first circuits to break the gigabit barrier were on the EOS18 device chip (Figure A-1k). The SiGe photodiode connected to the receiver exhibits sub-μA dark current, responsivity of approximately

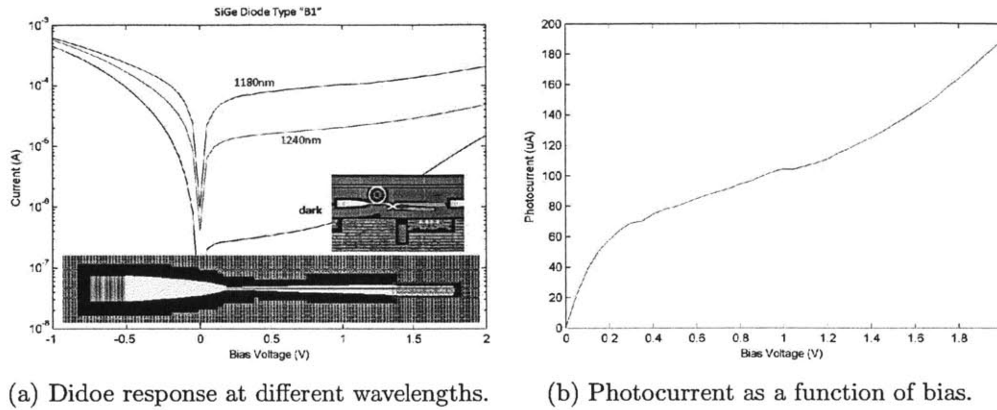

(a) Didoe response at different wavelengths. (b) Photocurrent as a function of bias.

Figure 4-12: SiGe photodiode design and measurement (EOS16, measured by Jason Orcutt).

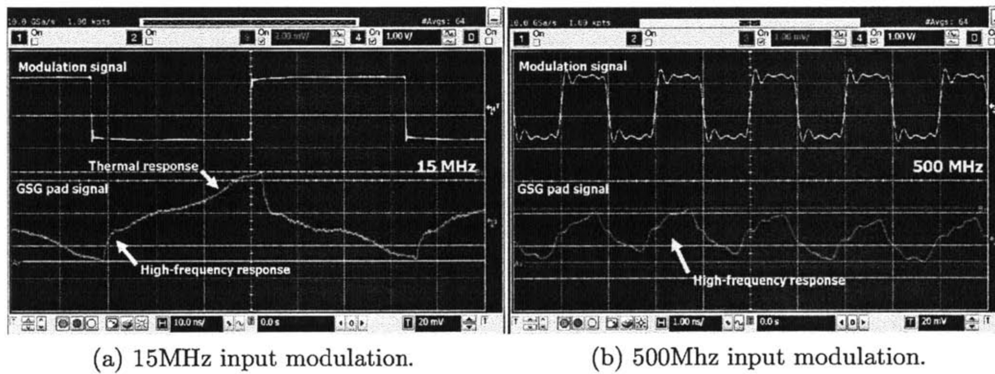

(a) 15MHz input modulation. (b) 500Mhz input modulation.

Figure 4-13: Optical data reception (EOS16).

0.02A/W, and bandwidths between 200MHz and 5GHz for reverse biases of 0.5V and 5V. Figure 4-14 shows the performance of the promising photodiode. The circuit again consists of a TIA followed by an output buffer that drives a GSG pad. A small addition in this design is a PMOS current mirror to set the bias of the TIA (bias0 and bias1 are current-mirror inputs used to source current to or from the photodiode node, providing offset compensation in both directions).

The TIA and GSG structure was repeated 15 times on the device chip in order to test 15 different photodiode devices. The photodiodes were replicated elsewhere on the chip with their connections made directly to DC pads so that their I-V char-

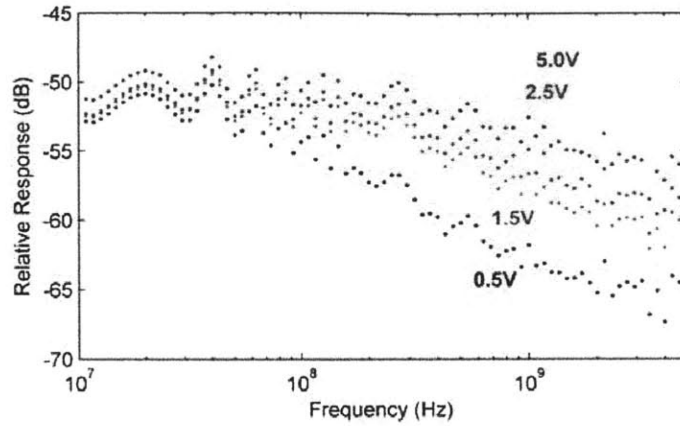

(a) Photodetector bandwidth.

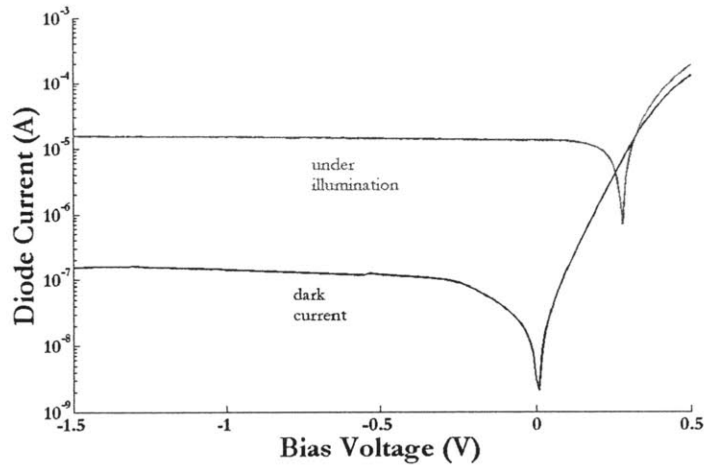

(b) Measured I-V characteristics in the dark and under 0.25mW illumination at 1180nm.

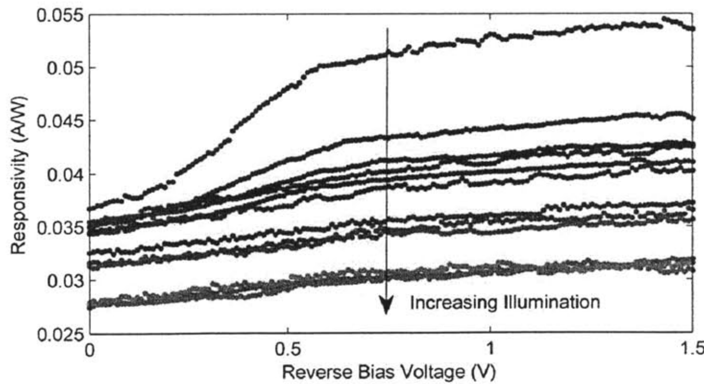

(c) Responsivity as a function of voltage and illumination intensity from 0.024mW to 0.6mW.

Figure 4-14: SiGe photodiode performance (EOS18, measured by Ashwyn Srinivasan).

acteristics can be measured. Figure 4-15 shows the layout of one of the photodiode test sites connected to a GSG-TIA receiver. A fiber and grating couple a modulated optical signal into the chip from the right. Once in the chip, the optically-modulated signal travels down a short waveguide until it reaches an absorption-type photodiode (under the area shown in the diagram). The photodiode is connected to the receiver TIA, which converts the optical signal to an electrical one, and buffers up the signal to drive an output pad.

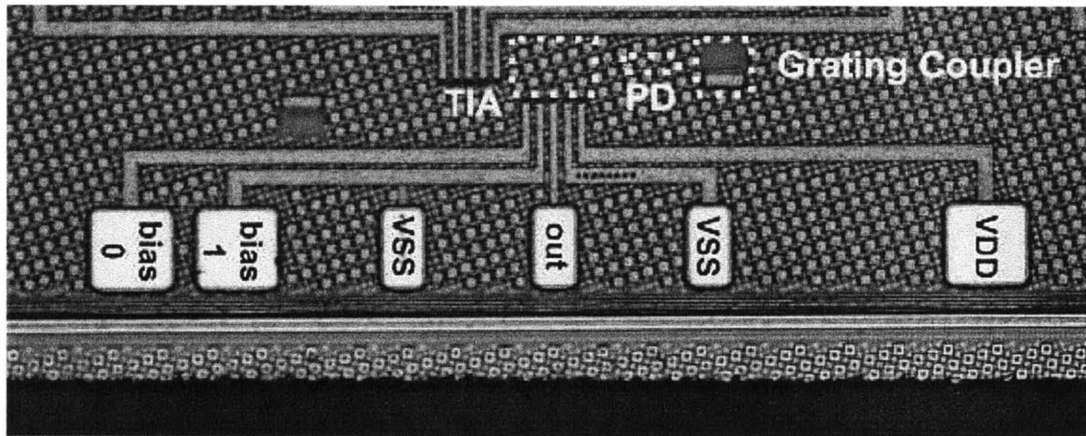

Figure 4-15: GSG-testable gigabit TIA in 45nm SOI.

The resulting eye diagram is viewable on an electrical oscillscope. Figure 4-16 shows the receive eye diagrams at 1Gb/s and 1.5Gb/s. This result represents the first gigabit optical detection for the project, and marks a major milestone.

4.4.3 Multi-Gigabit Optical Receivers

The work from the previous sections has shown that photodetectors and receivers can be integrated onto the same die and operate at gigabit rates. In this section, we detail a fully-monolithic, multi-gigabit receiver integrated into a synthesized, high-speed digital backend.

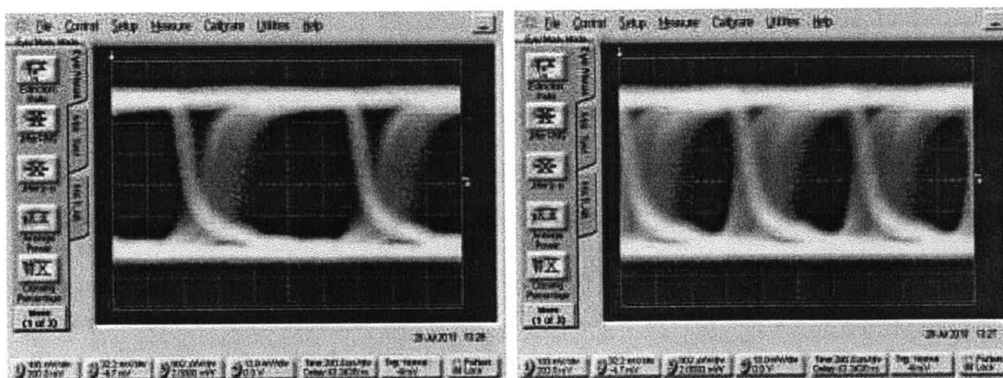

(a) Schematic. (b) Layout.

Figure 4-16: Gigabit receiver eye diagrams in 45nm SOI (EOS18).

TIA-biased receiver

As discussed earlier, TIA-based receivers that make use of a limiting amplifier are expensive and not well-suited to many-core applications where they will be replicated many times throughout a chip. We employ a TIA front-end followed by a clocked comparator known for energy-efficient operation. In a many-core link, the comparator will make use of a source-forwarded clock traveling on a parallel wavelength-channel. The photodiode was a version of the diode measured in Figure 4-14, shortened by a factor of ≈ 2.

The TIA (designed in Section 3.3)uses both NMOS and PMOS transistors for feedback resistance to improve the linearity of the amplifier. A current-steeering DAC connected to the input of the TIA is used to offset dark current, bias the TIA, and/or sweep the decision threshold. The comparator consists of a traditional Sense Amplifier (SA) thtat uses capacitive DACs for offset compensation and threshold-sweeping functionality required to measure and eye diagram. The output of the SA, which resets once every bit period, is fed into a dynamic-to-static block. The high-speed digital backend records the receiver's bit decisions and errors *in situ*, exporting only statistics and pattern snapshots off-chip.

Figure 4-17 shows the receiver's performance when operated from a 0.9V supply, with photodiode reverse bias of 0.45V. Although the bandwidth of the photodiode

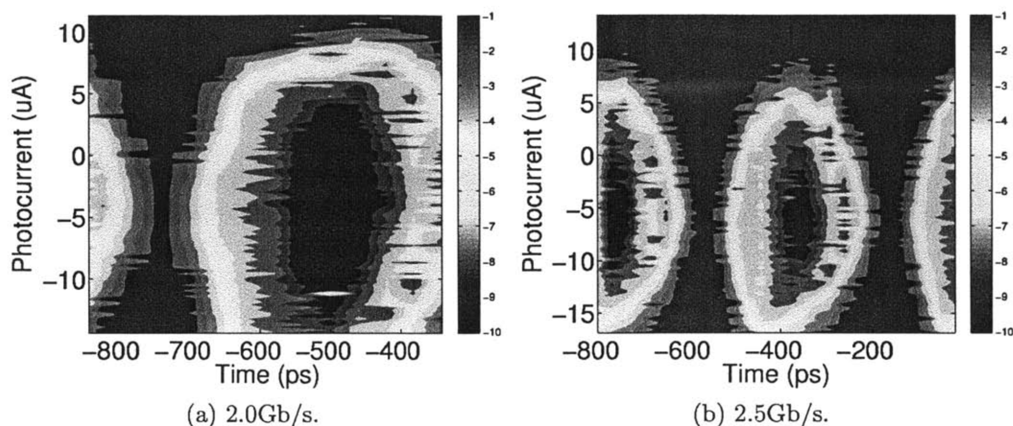

(a) 2.0Gb/s.

(b) 2.5Gb/s.

Figure 4-17: EOS18 TIA receiver eye diagrams.

is only 200MHz at this bias voltage, the relatively mild roll-off enables operation at up to 2.5Gb/s. A separate bias would enable operation at rates well above 10Gb/s in future designs. The BER eye diagram shows an open ey at 2.0Gb/s, at a sensitivity of $10\mu A$ (average photocurrent) and BER $<$ 1e-10 for a PRBS31 input. At this operating point, the receiver does not require the use of the current-DAC, resulting in an energy-cost of 165fJ/bit. At 2.5Gb/s where the current-DAC is required to bring the center of the eye within the capacitive-DAC's range, the energy-cost increases to 220fJ/bit while sensitivity decreases to $15\mu A$. The resulting optical sensitivities at 2.0Gb/s and 2.5Gb/s are -3dBmW and -1.3dBmW, respectively. We note that we are not using a split-diode and DDR in this instance, effectively wasting half of a bit time. The implication is that a 2.5Gb/s result will easily translate to a 5.0Gb/s result with sensitivity degraded by a factor of 2, provided the photodiode can support that bandwidth.

The energy-cost of the receiver as a function of data-rate is shown in Figure 4-18a. The cost is dominated by the TIA and current-DAC, though both components' costs are amortized as the data-rate increases. The receiver's digital blocks (SA, dynamic-to-static, clock buffer) have a flat energy-cost component as their power will follow fCV_{DD}^2. The static-current component can likely be reduced in future chips, now that more is known about the dark currents and threshold variations that can be

expected.

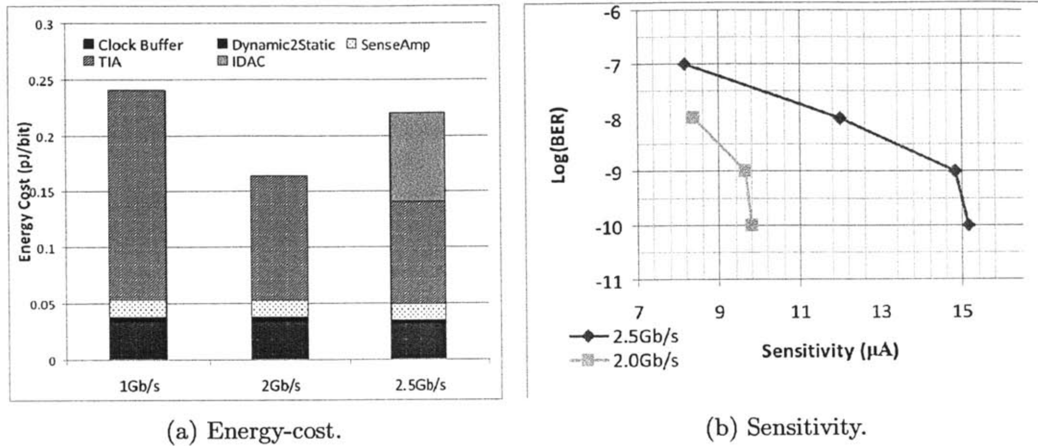

(a) Energy-cost.

(b) Sensitivity.

Figure 4-18: TIA-based receiver performance vs data-rate (EOS18).

The sensitivity of the receiver is shown in Figure 4-18b. The sensitivity was measured by measuring the optical power levels off-chip and applying a coupling loss of 10dB and responsivitiy of 0.02A/W. The eye diagram was then recorded in both statistical and BER mode. From the statistical mode we obtained the separation of the 1 and 0 levels in bits. From that difference we subtract the 1e-10 BER eye opening from the BER eye diagram (which represents the vertical margin in the measurement). The power level at the photodiode is then scaled by the margin. Using this method we obtain a photocurrent sensitivity of $10\mu A$ at 2Gb/s and $15\mu A$ at 2.5Gb/s.

Table 4.1 summarizes the receiver's performance. Modulator performance is also listed in order to give the reader perspective on overall link performance.

Latching Sense Amplifier (LSA)

In Section 3.2 we presented the current-detecting, Latching Sense Amplifier (LSA) as an energy-efficient, aggressive digital design that uses the small capacitance available in monolithic integration to connect the photodiode directly to the current branches of the latch. While the circuit has been used for several of the low-bandwidth results presented earlier in this chapter (including the first-ever optical detection by

Technology Photonics	45nm SOI CMOS Monolithic
RX Supply	0.9V
RX Data-rate (Gb/s)	2.0 / 2.5
RX Efficiency (pJ/bit)	0.165 / 0.220
Sensitivity (μA) (Average Photocurrent)	10 / 15
RX Sensitivity (dBmW) (Average Optical Power)	-3 / -1.3
RX Area (mm^2)	0.0003
RX C_{in} (fF)	15
TX Data-rate (Gb/s)	3.5
TX Energy (pJ/bit)	0.07
TX Area (mm^2)	0.002

Table 4.1: Logic-process transceiver performance summary. Optical receiver sensitivity calculated based on PD responsivity of 0.02A/W.

a circuit in this project), penetration into higher data-rates has never been achieved. We briefly discuss some possibilties for why the latching receiver has been unable to work at higher rates.

The leading theory as to why the LSA may have trouble detecting higher-rate data has to do with the biasing of the diode. We saw in Figure 3-3 that the photocurrent generated by the diode increases significantly with reverse bias. In its *reset* state, the LSA should be strongly reverse-biasing the photodiode, generating the most photocurrent for an optical-1. As the latch evaluates, if an optical-1 is detected, then the bias on the photodiode will decrease more quickly, reducing the photocurrent generated. Furthermore, the bandwidth of the photodiode will drop quickly during the evaluation phase, further hurting the possibility of gigabit optical detection.

Properties of the latch unique to an SOI process were also investigated, such as the SOI body effect and hysteresis in the latch's decision. However, LSA implementations in the bulk process in the next section were similarly unsuccessful. Furthermore, the LSA was shown to operate correctly using an electrical diode-emulation circuit in Figure 4-11a, strengthening a theory based on the photodiode's interface with the circuit. Since a larger photodiode reverse-bias needs to be applied to increase

bandwidth, the relative drop due to the LSA will decrease, and we can expect the LSA to operate correctly and with the sensitivity demonstrated with the diode-emulation circuit. If the photodiode has a large dark current, this large reverse bias comes at the cost of increased power. An additional bias for the photodiode also increases the complexity of the circuit.

4.5 Optical Data Receiver: Memory

The over-arching goal of this project is optical memory-access in future many-core processors, and so integrated photonic circuits working in a memory (DRAM) process are a critical component. This section describes the measured results from the memory chips that were fabricated. All of the system components were fabricated in a 0.180um bulk flash periphery process - unsurprisingly slower than the 45nm SOI logic process in the previous section. As a result, multiple design changes were required for both the analog front-end circuits and high-speed digital backend.

In contrast to the Logic work presented earlier, there was only a single generation of designs for the memory part of the project (though there was significant iteration in on the process and photonic devices). As a result, this section highlights the results from one batch of chips: the Micron D1L chips. The first batch of D1L chips had a fabrication error that resulted in transistor poly gate lengths that were 40nm longer (slower) than designed. As a result, the achieved data-rates from the first batch of chips were slower than expected.

4.5.1 Stand-alone Receiver Performance

The optical data receiver (Figure 4-19) consists of a configurable-gain transimpedance amplifier (TIA) followed by a clocked sense-amplifier. To mitigate the slow speed of the process (F04 $\approx$100ps), we utilize two half-rate receivers operating on opposite clock phases (Section 3.6). The microring photodiode is separated into two electrically-isolated half-photodiodes (half-PD), each connected to one half-rate receiver. A dummy half-PD and TIA serve as a reference for each sense-amplifier while

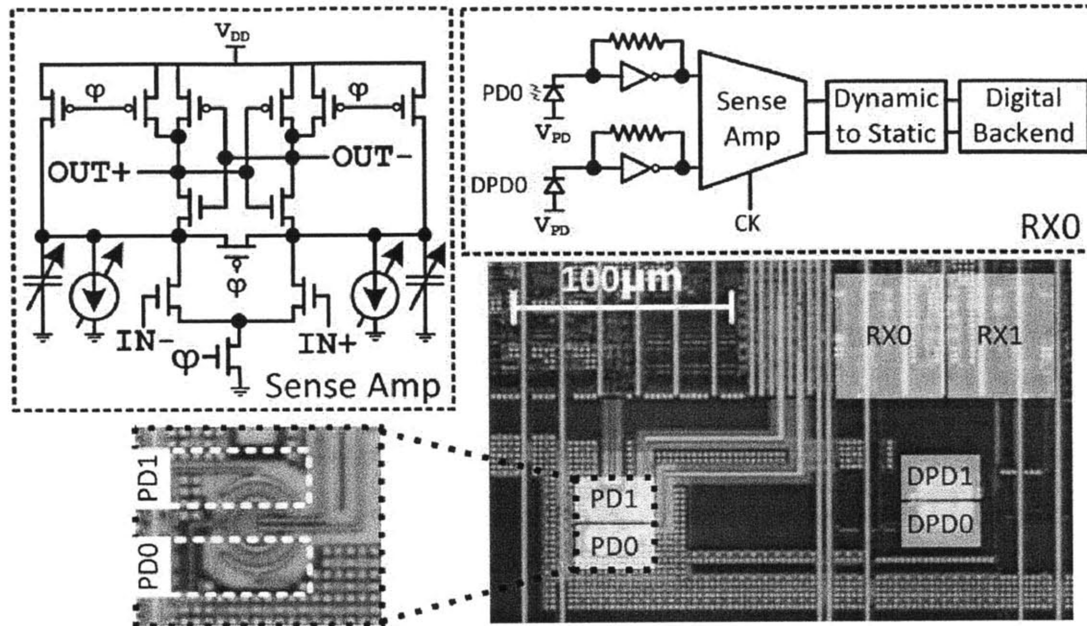

Figure 4-19: Receiver schematic (D1L).

current and capacitive DACs provide offset compensation and eye-measurement functionality. In this process, higher data-rates squeeze sense-amplifier evaluation time, requiring much larger photocurrents to compensate. The TIA's static power consumption dominates the receiver energy cost.

The PIN photodetectors utilize free-carrier generation in ridge-waveguide microrings through sub-bandgap optical transitions involving defect states in the polySi waveguide core [46]. Responsivities of ≈0.2 A/W were measured at low biases. Note that this is significantly better than the responsivity measured in the SiGe detectors in the logic process. Device bandwidth is strongly voltage-dependent, exhibiting 3dB bandwidths of 1.5GHz and 9.7GHz at -1V and -15V biases, respectively. Notably, the defect detectors work in both the low-1200nm and the 1300nm wavelength ranges. Within each range the wavelength must be tuned to the ring of the defect detector.

The inset in Figure 4-19 shows a die photo of the receiver location tested. From photo, we can see how two separate detectors were placed in one microring, with each detector routed to a separate receiver. We can also see the placement of a dummy microring, with two reference detectors. The figure shows that photodetector place-

ment relative to the receiver circuits is an area for improvement - closer placement will drastically reduce wire length (and therefore capacitance).

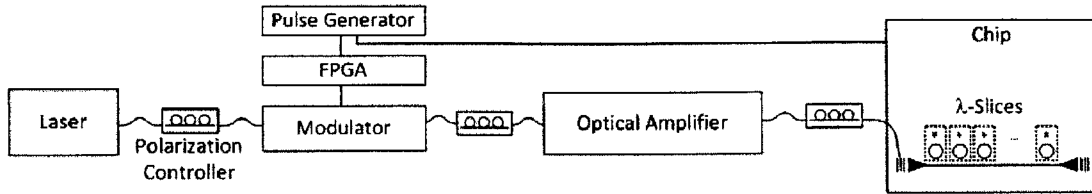

Figure 4-20: Memory process test setup.

Figure 4-20 shows the experimental test setup for the Memory chip. Light from an off-chip laser source is polarization-corrected and input into a discrete MZI modulator. The modulator is driven by eight time-interleaved PRBS generators programmed onto an ML605 FPGA development board. The optically-modulated light is amplified, again polarization corrected, and coupled into the chip through a grating coupler. The wavelength is tuned until it matches the wavelength-channel of a single slice of the WDM receiver bank. At the correct wavelength-channel, light is coupled into the microring and detectors, and converted to an electrical signal that can be detected by the integrated receiver.

As an early result, Figure 4-21 shows the receiver's eye diagram for a data-rate of 3 Gb/s. From the plot, we can see that there is a clear opening in the center of the eyes, and that the receiver achieves a BER of better than 1e-9. This batch of chips contained the mask error that resulted in slower transistors. We mark this result as proof that photonics can be enabled in the process, and continue to test chips that have correct gate lengths.

With the gate lengths corrected on a new set of wafers, the data-rates were pushed from lower rates (2Gb/s in Figure 4-22) towards their target 5Gb/s and eye diagram were recorded (Figure 4-23). From the plot it is clear that there is a clean, open eye that achieves a BER of better than 1e-10. An adaptive BER eye diagram measurement script was run to sweep out the eye. The script continues to check a particular phase and threshold value (DAC code) until a bit-error is recorded or until the total number of correct bits is greater than 1e10. This means that the algorithim sweeps

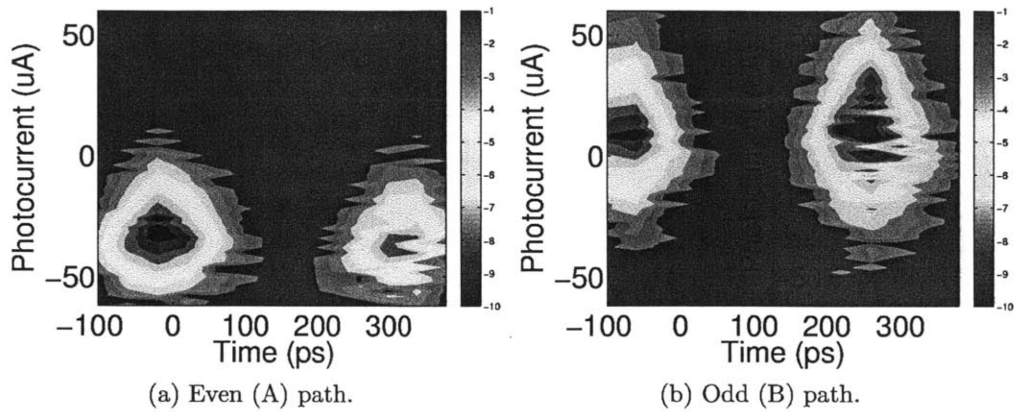

(a) Even (A) path.

(b) Odd (B) path.

Figure 4-21: Receiver eye diagrams at 3Gb/s on the Memory chips with fabrication errors (D1L).

quickly over the high bit-error parts of the eye, but spends time correctly measuring the bit-error-rate at the center.

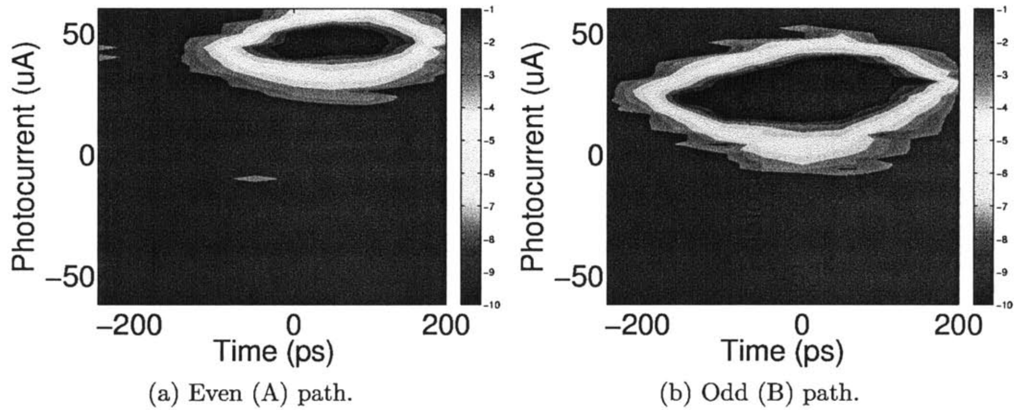

(a) Even (A) path.

(b) Odd (B) path.

Figure 4-22: Receiver eye diagrams at 2Gb/s on the correctly-processed memory die (D1LL).

One of the unique challenges of this work is that unlike conventional (electrical) high-speed link testing, many of the optical components drift in and out of alignment with time constants that range anywhere from hours to tens of seconds. Without thermal locking loops activated in the receiver, even the receiver's or the electrical backend's own activity can change the temperature of the chip enough to cause the

resonant frequency to drift. As a result, a seemingly simple experiment can often take a day or more to perform, and this is once an optical photodetector and circuit have been identified. With the development of more advanced system-level control (e.g. thermal feedback loops), and more sophisticated photonic packaging technology (which will keep optical fibers better aligned), these testing problems go away. We discuss them here to emphasize the importance of system-level controls and packaging solucations, and also to give the reader perspective on where time is spent in lab work.

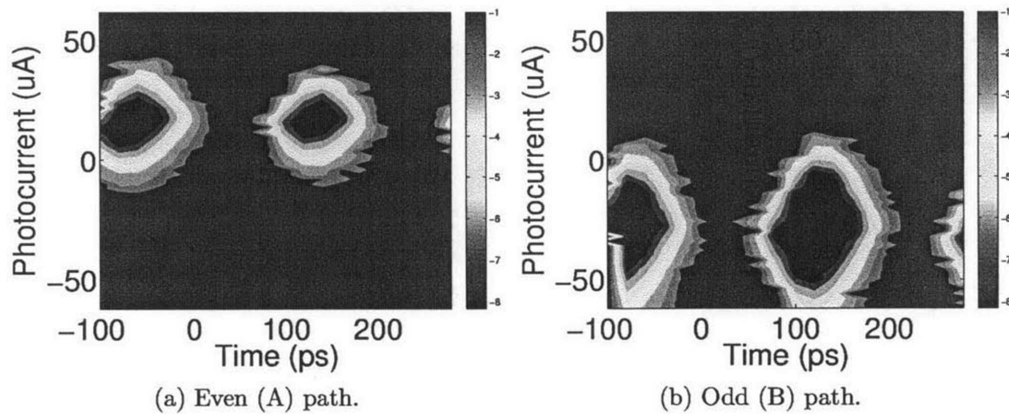

(a) Even (A) path.

(b) Odd (B) path.

Figure 4-23: Receiver eye diagrams at 5Gb/s on the correctly-processed memory die (D1LL).

That bathtub curves generated from the eye diagrams of Figure 4-23 are shown in Figure 4-24. When looking at a double-data-rate receiver, it is important to ensure that there exists a phase where both the even and odd phases of the receiver are error-free. Figure 4-24 shows the BER as a function of phase for one threshold code of each half-receiver. From the plot, we can see that there does in fact exist a range of receiver clock phases where both of the receivers have bit error rates of better than 1e-10.

In Chapter 2 we saw how critical the receiver's sensitivity was in determining the main operating point for several other link components. The energy-cost and sensitivity of the receiver as a function of data-rate are shown in Figure 4-25. From the plot, we can see that the TIA dominates the energy-cost of the receiver, though its

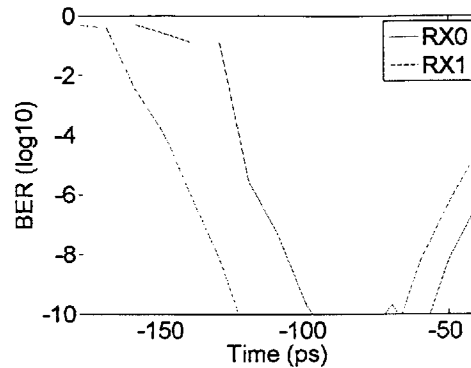

Figure 4-24: Receiver bathtube curves at 5Gb/s.

static power consumption results in a energy-cost that decreases with data-rate. The digital parts of the receiver (the sense amplifier and dynamic to static converter) have a flat energy-cost as expected. The sensitivity of the receiver is fairly respectable at low data-rates, with a photocurrent sensitivity of around $2\mu A$. At higher rates such as 5Gb/s, exponentially more photocurrent is needed to reach the required BER.

There are two contributing factors to this. The first, described earlier, is the fact that in this slow process, an increased data-rate reduces the bit time and as a result reduces the evaluation phase of the receiver, during which the cross-coupled latch regenerates exponentially. The second part of the problem pertains to a wiring mistake, visible in Figure 4-19. The photograph shows that the wire that connects the photodiode to the receiver was made unnecessarily long. As a result, there is additional capacitance at the input node that was unaccounted for in the initial design. To give the reader more perspective, examine the eye-diagrams shown in Figures 4-21, 4-22, 4-23. In all of the plots we can see that the Even (A) path of the receiver has a smaller eye opening than the Odd (B) path. This is likely attributable to the fact that the Even path has considerably longer wiring than the Odd path ($\approx 100\mu m$ longer). While a seemingly careless mistake, keep in mind that the Memory process work consists of multiple chip templates that each contain more than a hundred devices with different variations. More sophisticated electronic-photonic development tools

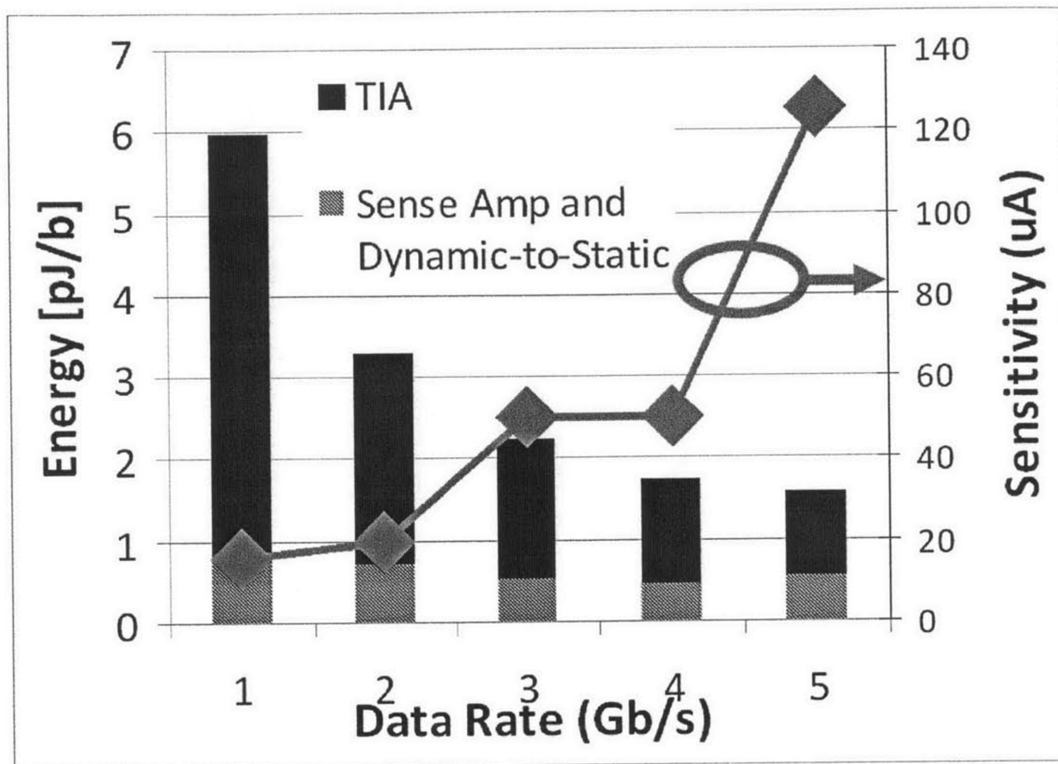

Figure 4-25: Receiver energy-cost and sensitivity vs data-rate for the Memory process (D1LL).

will help the designer avoid mistakes like this.

4.5.2 DWDM Receivers

While we have focused so far on individual receiver results, the end-goal here is to demonstrate DWDM link results. Towards that end, we performed a DWDM receiver test, where the wavelength of the optically transmitted signal was tuned to each of the 9 receivers in a DWDM bank. Figure 4-26 shows the testbench diagram, a die photo with the DWDM test row, and the receiver eye diagrams at 5Gb/s for each of the 9 channels as the input wavelength was swept. The result is an aggregate-throughput of 45Gb/s.

The receiver's performance is summarized in Table 4.2. The modulator and driver performance is listed to give the reader perspective on link performance.

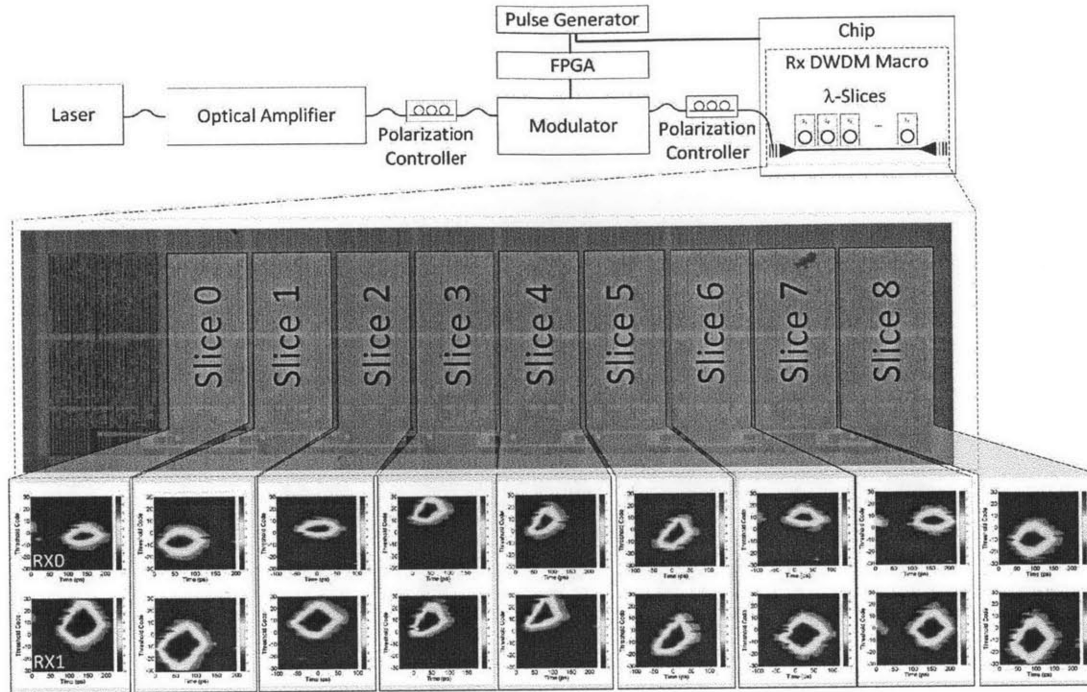

Figure 4-26: DWDM result showing each receiver receiving a different wavelength.

Technology **Photonics**	$0.180\mu m$ Bulk CMOS Monolithic
RX Supply	2.5
RX Data-rate (Gb/s)	5
RX Efficiency (pJ/bit)	1.57
Sensitivity (μA) **(Average Photocurrent)**	125
RX Sensitivity (dBmW) **(Average Optical Power)**	-2
RX Area (mm^2)	0.005
RX C_{in} (fF)	50
TX Data-rate (Gb/s)	5
TX Energy (pJ/bit)	1.22
TX Area (mm^2)	0.006

Table 4.2: Memory-process receiver performance summary. Receiver optical sensitivity based on photodiode responsivity of 0.2 A/W.

To summarize the importance of this section, by developing the photonic platform in Micron's memory process, we have shown that photonic interconnects are viable in

cost-aware mainstream bulk CMOS processes. The optical data receiver was shown to operate at 5Gb/s, and while the sensitivity needs to be improved, methods for doing so were proposed. Even in this relatively old flash-periphery process, we then demonstrated a DWDM receiver bank with a total data-rate of 45Gb/s. This is particularly significant given the pin constraints discussed in Chapter 1. For comparable packaging pitch (fiber vs C4), we achieve 45Gb/s/waveguide - more than an order of magnitude better than DDR4 at 1.5Gb/s/pin. With a working photodiode and receiver combination now in place, a more tailored receiver can be designed to further increase performance.

4.6 Integrated Photonic Links

This thesis developments an integrated photonic platform, a system-level understanding of component trade-offs, and the design of an optical receiver suitable for integrated optical links. Towards that end, we have demonstrated the performance of the optical receiver on its own, operating in a testbench with a fairly ideal, discrete, Lithium Niobate modulator. In this section, we demonstrate the receiver's operation in chip-to-chip optical links. We start off by examining a single-wavelength link, from one Memory chip to another.

4.6.1 Single-λ Link

Figure 4-27a shows a diagram of the test setup for the single-λ link test. A photograph of the setup was presented earlier in this chapter (Figure 4-5). Light from an off-chip laser is coupled onto the first Memory chip through vertical grating couplers at a cost of about 5dB per coupler. The light propagates down an integrated waveguide until it reaches the ring resonant modulator. The modulator imprints data from eight multiplexed on-chip 31-bit PRBS generators onto the optical signal. The optically-modulated signal then propagates down the waveguide and exits through a second grating coupler. The signal travels down a fiber patch-cord and couples into a second Memory die through a third grating coupler. We emphasize that no amplifier

was used to regnerate the signal between the chips.

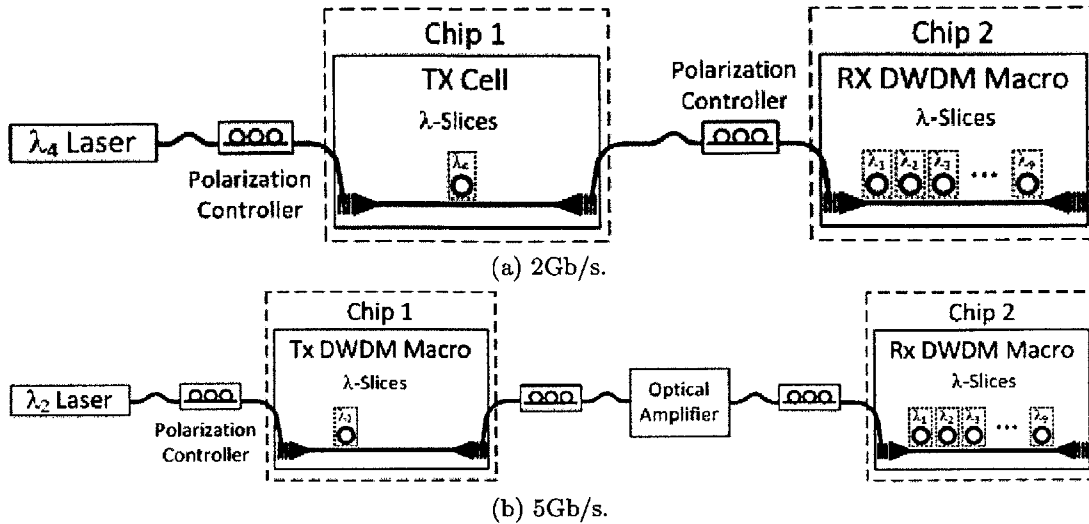

(a) 2Gb/s.

(b) 5Gb/s.

Figure 4-27: Single-λ Memory-Memory link test setup.

At the second memory chip, the optically-modulated signal propagates down an integrated waveguide until it reaches a defect detector tuned to the particular wavelength-channel. In this case, we use the 4th ring in the RX megacell row. While use of the fourth ring increases waveguide loss on the receiver-side chip, that ring was already fairly closely tuned to the transmit wavelength-channel, and so is a convenient receive location. The defect detector creates an photocurrent that is detected by a TIA-based receiver that regenerates the bit decision using a sense-amplifier. The data checked using another eight 31-bit PRBS generators on the RX chip.

Figure 4-28 shows the receiver BER eye-diagrams for the single-λ link. From the plots we can see that a BER of better than 1e-10 was achieved at a data-rate of 2Gb/s.

In order to get the link operating at its target 5Gb/s, the optical power into the second chip had to be increased. This was achieved through the addition of an optical amplifier between the two chips (Figure 4-27b). Figure 4-29 shows the transmit eye diagram coming out of Chip 1, as well as the eye diagrams for the Even and Odd phases of the receiver. We can see that the eye diagrams are open with a BER of better than 1e-10. This is a significant result with respect to today's memory I/O

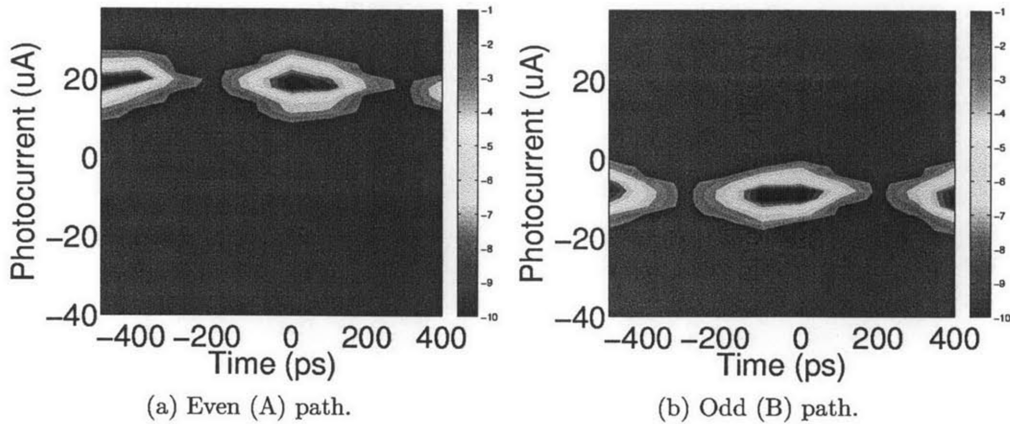

(a) Even (A) path.

(b) Odd (B) path.

Figure 4-28: RX link eye diagrams at 2Gb/s from the correctly-processed memory die.

speeds, especially considering that the fanout-of-4 in this process is 100ps making a 5Gb/s link only a fanout-of-2. As a point of reference for the reader, the 5m patch of cable between the two chips will have approximately 125 bits traveling on it at any time.

While an optical amplifier as used between the chips for this demonstration, it should not be regarded as a major point of concern. The amplifier will only add noise to the signal, degrading the SNR, so the successful detection of a clean receiver-side eye is significant. As for the overall optical (laser) power levels, there are several areas where improvement will obviate the need for an amplifier. We have already discussed how receiver sensitivity can be improved by reducing photodiode wiring length, giving the sense amplifier more time to evaluate (possibly through further photodiode splitting), and further tuning the receiver based on knowledge from the photodetector measurements (e.g. low dark currents). The optical link budget can also be improved by continued process development (lowering waveguide loss), more agressive grating coupler designs (decreasing coupler loss from $\approx$ 5dB/coupler to less than 3dB/coupler), and further photodetector responsivity improvement.

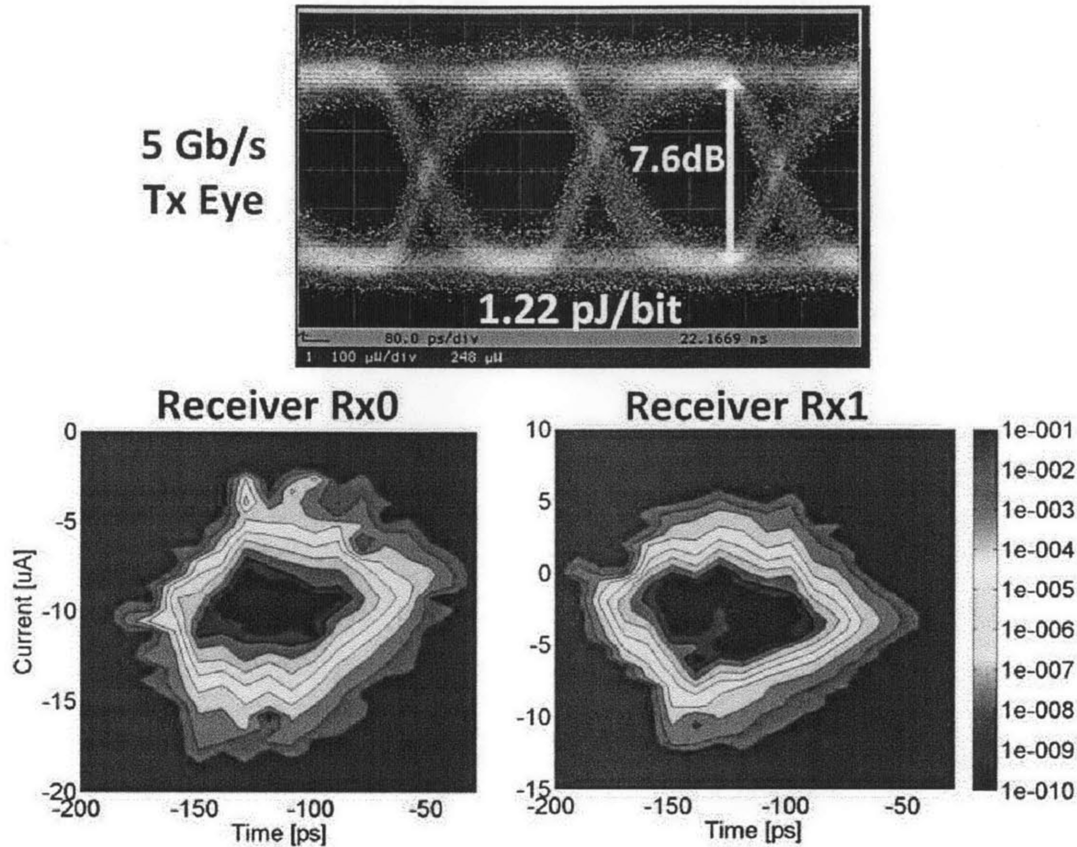

Figure 4-29: Single-λ Memory-Memory link at 5Gb/s, using an SOA between the TX and RX chips.

4.6.2 Multi-λ Link

Having demonstrated a single-λ link from one chip to another, we demonstrate WDM. In order to do this, we make use of the Megacell rows on the TX chip, where multiple modulator rings are located on each waveguide, each tuned to a separate wavelength channel (Figure 4-1). Integrated heaters fine-tune the channel alignment.

The top diagram of Figure 4-30 shows the test setup for the multi-λ link test. Two laser sources are used for the two wavelength-channels. The continuous-wave lasers are combined using an optical coupler, and coupled into the chip through the gratings. Two modulator slices on Chip 1 imprint data onto their respective wavelength-channels, and the two modulated signals continue on off-chip to be amplifier. At the receiver, Slice 0 and Slice 2 are aligned to the wavelengths channels.

The bottom portion of Figure 4-30 shows the resulting eye diagrams. We note here that non-adjacent slices were used on the receiver chip as these were the slices that were closest aligned to the wavelength channels. A miscalculation in the thermal tuning portion of the chip resulted in relatively small tuning ranges for each of the rings. A decrease in the resistance used for thermal tuning in the next generation of chips will solve this problem and enable a wider tuning range for each of the rings. With that functionality in place, adjacent channels will be easily tested.

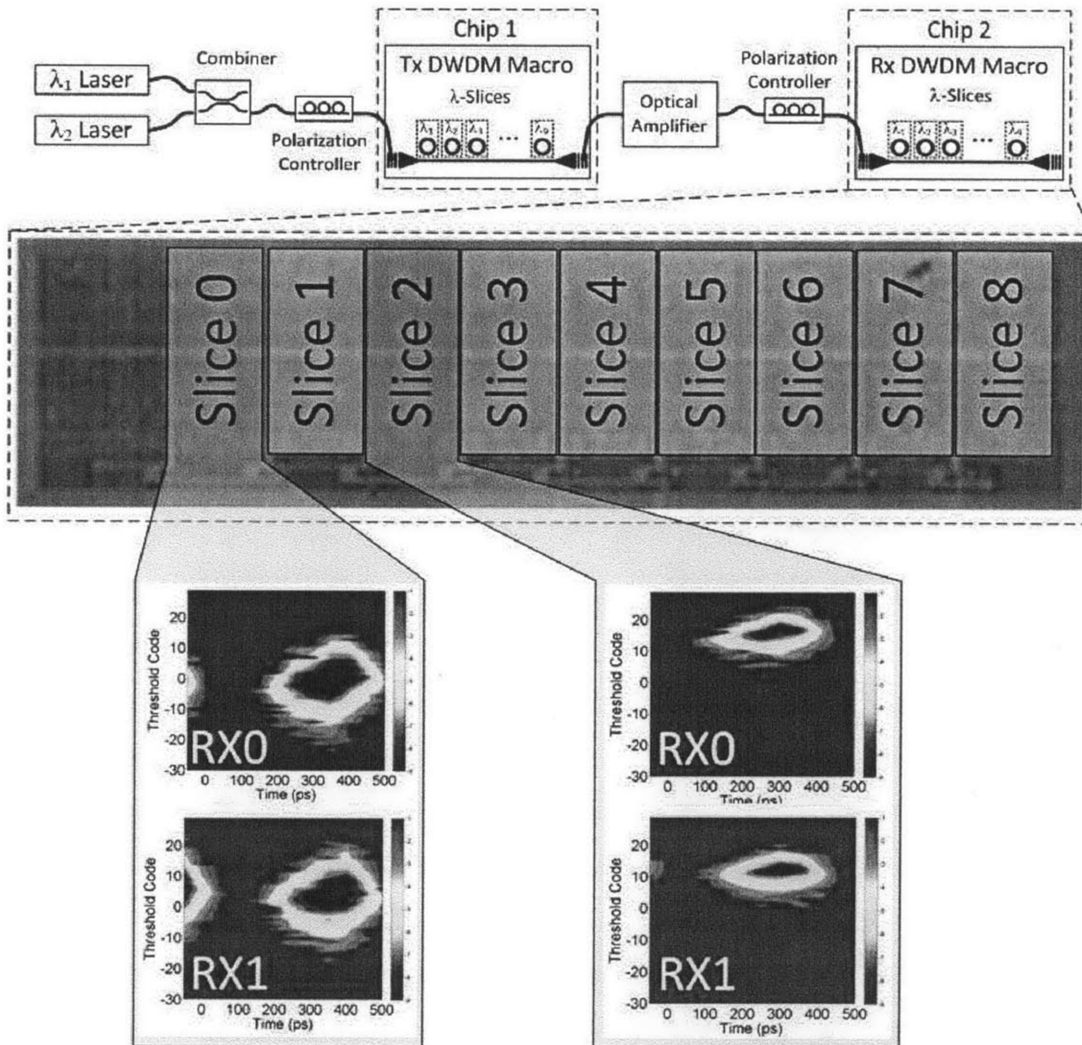

Figure 4-30: Link eye diagrams for WDM configuration.

4.7 Summary

In this chapter we developed the EOS photonic technology platform. We started by detailing the platform's architecture, which consists of numerous flexible test cells, each of which interfaces with a variety of optical devices. Platform architectures were explored in both the memory and logic chips. We then looked at several smaller test circuits, such as ring oscillators and GSG-testable transimpedance amplifiers, which were used to evaluate processing steps in the technology as well as provide the first signs of life from promising photodetectors. With a foundation built upon simple circuits and functional photonic devices, we detailed the first-ever operation and performance of multi-gigabit optical receivers in both the bulk CMOS memory and zero-foundry-change, commercial SOI CMOS logic processes. The SOI receiver was shown to be compact, energy-efficient, and achieved excellent μA sensitivity. The viability of the bulk CMOS memory receivers was demonstrated through chip-to-chip link demonstrations that show both high-speed links (5Gb/s) as well as the first wavelength-division multiplexed links for this project. Additionally, we demonstrated a 9x5Gb/s DWDM receiver bank that achieves 45Gb/s/waveguide - an order of magnitude better than pin-constrained DRAM solutions in the low-Gb/s/pin range.

Chapter 5

Conclusion

This thesis presented integrated photonics as a technology that is set to disrupt traditional communication systems. The platform was presented in the context of memory access in future multi-core processors. After first understanding previous optical receiver designs and their application space, we focused on monolithic integration where the electronics and photonics are integrated onto the same die.

We developed a system-level model of the integrated photonic link, composed of smaller models of each of the individual components (modulators and drivers, rings, waveguides, couplers, ring tuning, and optical receivers). Link-level analysis was performed for a range of total throughputs, and the overall energy-costs of the systems were compared. The work showed that for the range of throughputs explored, an energy-cost optimum can be achieved, and is typically in the range of 5Gb/s to 10Gb/s per wavelength channel. Note that the designer is free to set the channel-rate higher (perhaps in an effort to comply with existing telecom standards), but this may not be the most efficient region of operation.

Once the system-level photonic trade-offs were understood, we focused on the design of the optical receiver. We focused on two different receiver designs: a TIA-based receiver that uses a standard sense amplifier for signal regeneration, and a latching current-sense-amplifier. The TIA base receiver is a more traditional design, while the current-sensing latch is intended to reduce energy cost. An equivalent-circuit model was derived for the LSA, but as the TIA-based receiver has a sense

amplifier, the principles of the model can be applied to both receivers.

With a system-level view in place and two receivers designed, this work presented the EOS integrated photonic technology platform. The EOS platform is designed to enable the testing of a wide variety of circuits and devices, including photonic links. Chapter 4 presented a number of measured results from the EOS platform, focusing on gigabit optical receivers. The work demonstrated the first-ever μA-sensitivity gigabit optical receiver implemented in a zero-foundry-change, commercial SOI CMOS process. The SOI receiver is compact, energy-efficient, and well-suited for VLSI applications. A receiver implemented in a cost-aware memory process was also measured as a part of a chip-to-chip optical link, in both single-channel and WDM modes. This chip-to-chip link represents the first monolithically-integrated optical transceiver implemented in a bulk CMOS process ever.

The era of integrated photonic interconnects is upon us. This thesis presented numerous circuit and system techniques and integration demonstrations that pave the way to the commercialization and adoption of this high-potential technology.

5.1 Future Work

This thesis demonstrated an integrated photonic platform that combines electronics and photonics onto a single microchip. The platform was developed as part of a link system aimed at solving the bandwidth-energy-cost bottleneck of memory access in emerging many-core processors. However, the integrated photonic platform developed is very flexible and can be adapted to solve a host of other problems. We briefly look at other applications for integrated photonics. We will also outline what the next steps are in order to mature the platform itself.

5.1.1 Roadmap to 1Tb/s

Early on in this thesis we described the potential of integrated photonics as being able to provide energy-efficient links at bandwidths up to 1Tb/s. In this work, the highest data-rate achieved was 45Gb/s aggregate through the WDM receiver bank in

the memory process. While this is significantly smaller than 1Tb/s, there are several small modifications that are being developed that will pave the way.

The first step towards increasing the total bandwidth is increasing the data-rate-per-channel. While the system currently operates at data-rates of 5Gb/s, we can increase the rates to 10Gb/s or 16Gb/s while still operating near the efficiency-optimum. With the channel-rate increased, we can increase the number of channels by both increasing the available bandwidth, and decreasing the space between adjacent channels (Figure 5-1).

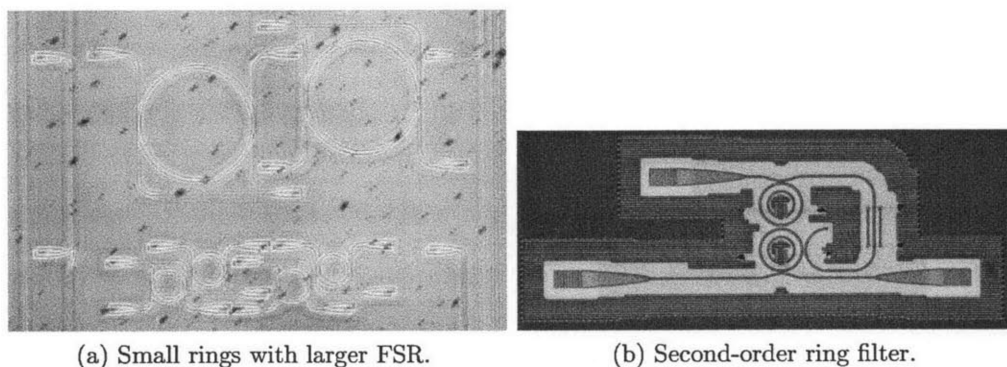

(a) Small rings with larger FSR. (b) Second-order ring filter.

Figure 5-1: Roadmap to 1Tb/s.

At the current ring radii, the FSR (usable system bandwidth) is roughly 2THz. By decreasing the radius (Figure 5-1a), we increase the FSR and the available system bandwidth. As the rings get smaller, however, they become lossier due to increased bending-loss, and the Q degrades. Reasonable advances in the ring design should enable us to double or triple the FSR, creating 2x to 3x the system bandwidth.

Once the system bandwidth has increased, we can also decrease the spacing between channels (and therefore pack more channels into a fixed system bandwidth). In order to decrease channel spacing without suffering increased channel-crosstalk, the order of the system (and therefore the Q) must be increased. Figure 5-1b shows preliminary work on designing second-order ring filters and modulators that have a sharper channel response. The second-order filters come at the cost of increased area and thermal-tuning control power.

These three advances are already in development and represent relatively incremental changes to the design that translate to a large increase in total throughput.

5.1.2 Platform Development

The integrated photonic platform has developed significantly from the project's inception. Photonic structures have been parameterized and layouts can now be generated quickly and to specification. The electrical tool flow follows a traditional synthesis-place-and-route flow, and can accomodate the range of required blocks, from custom link components to multi-core processors. Despite this there are several areas for improvement and additional infrastructure that needs to be brought up.

The first (and most critically needed) piece of tool infrastructure is more advanced integrated photonic verification tools. Currently, the photonics and electronics are designed separately, with the physical designs being instantiated together at higher levels of the tool flow. Circuits are checked with DRC and LVS, Photonics are checked with DRC, and automated fill routines create photonic-compatible fill structures. However, the capability to perform a full-system LVS (even treating photonic componets as black boxes) does not exist. This has resulted in several hand-drawn connections failing, as they could not be caught by LVS. Additionally, DRC/LVS rules that can test the layers used to create photonic devices are needed. We recently suffered an error where a doping-block layer was accidentally omitted from a small section of waveguide, creating upwards of 10dB loss in that section and rendering the connected system untestable. Once the verification tools have been upgraded, more complex capabilities such as photonic place and route can be developed. Such a system could enable the designer to instantiate functionality, similar to verilog, and have the tool determine where resonant rings, photodiodes, waveguides, and modulators be placed, and tuned to which wavelength-channels. This level of abstraction will be critical as integrated photonic systems continue to develop and scale.

Outside of the tool flow and die-level improvements, both the laser and packaging strategies need to be addressed. The integrated photonic platform chips presented in this work were tested with a variety of laser sources, from commercial tunable

multi-band lasers, to custom-built 1200nm lasers. One of the main advantages of our lower-rate DWDM model is that we assume that a centralized power source is used, amortizing the laser cost across many links. This breaks from the previous paradigm whereby a separate laser is needed for each link, driving data-rates higher in order to decrease the cost-per-bit. While laser energy-cost was taken into consideration in our system analysis, the comb laser source to be used has yet to be finalized. Its integration with the rest of the system is also not clear, and depends on the application area. In a processor-memory-communication application, the laser must be tightly integrated into the computer system, occupying a small footprint. In larger data-center applications, there is a little more flexibility, as one larger laser unit could potentially provide the source for all of the servers in the rack. As far as packaging goes, our strategy thus far has been centered on testability. We have used open-cavity ceramic packages that are compatible with the clean-room tools and enable optical fibers to be moved to various locations on the chip's surface. A commercial system will have each fiber fixed to a pre-determined grating structure, and the method for fixing and aligning that fiber is not yet known.

5.1.3 New Application Areas

In the computational application area, we focused on processor-memory access, where multiple processor cores can talk to each other as well as off-chip memory (DRAM). Extending this idea up one level, there are several applications in the data center where this technology can be applied. In particular, we can take advantage of the distance-insensitivity of optical interconnects and apply that trait to the connection of multiple servers within a rack, or between multiple compute racks. In a single rack, the degree of photonic integration presented in this work could enable multiple server processors to be optically-connected, possibly creating a single, virtual, many-core processor, where the total number of cores is the sum of the cores of all of the optically-connected processors. Again, while such architectures are *possible* through today's interconnect, integrated photonics makes such an architecture *viable.* A further extesion of this idea could enable the memory of each of the servers to be

grouped into one large pool.

Outside of the data-center there are several application areas that are a much larger departure from the research's origin, such as remote-sensing and advanced radar. There are also applications in medical imaging, and in particular Optical Coherence Tomography (OCT). OCT is a non-invasive imaging technology that uses near-infrared light to obtain cross-sectional and volumetric images of tissue. Notably, the technology can image specimens non-invasively and *in vivo*, making it an important tool in ophthalmology, gastroenterology and interventional cardiology. However, even the smallest modern OCT systems are not much smaller than a briefcase or small cart, with a hand-held component attached to a cord. The units are also fairly costly and as such are only feasible for use in hospitals or in specialized practices. OCT system consists of little more than a laser, fiber coupler, reference delay, imaging optics and a photodetector. Through discussions with experts in the field, we have learned that a current bottleneck exists between the photodetector and the DSP backend. The integrated photonic platform developed in this work could potentially be ideally suited to address issues of size and cost (through integration onto a single chip), as well as bandwidth (through the integration of a programmable digital backend near the optical path).

Appendix A

Die Photos

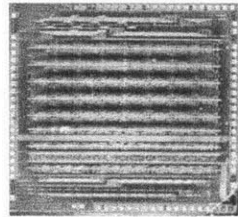
(a) EOS2 (32nm CMOS)

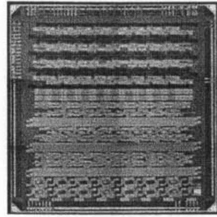
(b) EOS4 (45nm SOI)

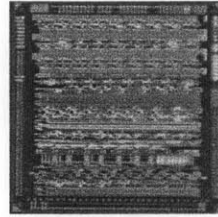
(c) EOS6 (45nm SOI)

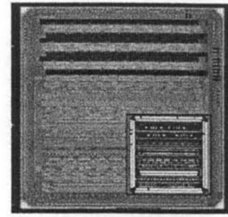
(d) EOS8 (45nm SOI)

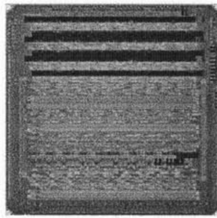
(e) EOS10 (45nm SOI)

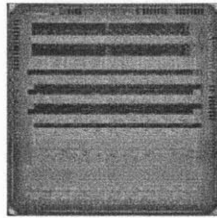
(f) EOS12 (45nm SOI)

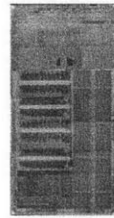
(g) EOS14 (45nm SOI)

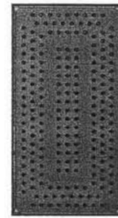
(h) EOS16 (45nm SOI)

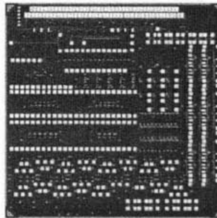
(i) EOS16 photonics (45nm SOI)

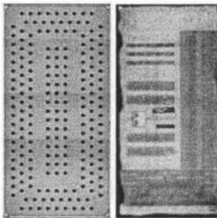
(j) EOS18 (45nm SOI)

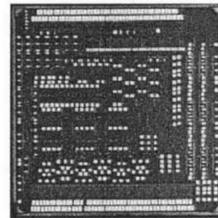
(k) EOS18 photonics (45nm SOI)

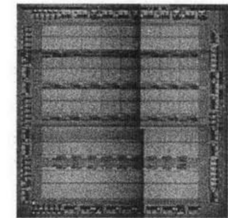
(l) Micron D1S (0.18um bulk)

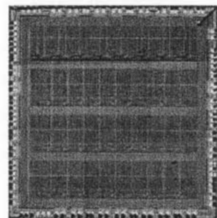
(m) Micron D1L (0.18um bulk)

Figure A-1: EOS chip photos.

Bibliography

[1] Reid J. Riedlinger, Rohit Bhatia, Larry Biro, Bill Bowhill, Eric Fetzer, Paul Gronowski, and Tom Grutkowski. A 32nm 3.1 billion transistor 12-wide-issue Itanium; processor for mission-critical servers. In *ISSCC '11*, pages 84 –86, feb. 2011.

[2] S. Palermo, A. Emami-Neyestanak, and M. Horowitz. A 90 nm cmos 16 gb/s transceiver for optical interconnects. *Solid-State Circuits, IEEE Journal of*, 43(5):1235–1246, May 2008.

[3] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Holzwarth, M. Popovic, Hanqing Li, H. Smith, J. Hoyt, F. Kartner, R. Ram, V. Stojanovic, and K. Asanovic. Building manycore processor-to-dram networks with monolithic silicon photonics. *HOTI '08*, pages 21–30, Aug. 2008.

[4] George Kurian, Jason E. Miller, James Psota, Jonathan Eastep, Jifeng Liu, Jurgen Michel, Lionel C. Kimerling, and Anant Agarwal. ATAC: a 1000-core cache-coherent processor with on-chip optical network. In *Proceedings of PACT '10*, pages 477–488, New York, NY, USA, 2010. ACM.

[5] A. Joshi et al. Silicon-photonic clos networks for global on-chip communication. In *Proceedings of NoCS '09*, page 91, may 2009.

[6] Dana Vantrease, Robert Schreiber, Matteo Monchiero, Moray McLaren, Norman P. Jouppi, Marco Fiorentino, Al Davis, Nathan Binkert, Raymond G. Beausoleil, and Jung Ho Ahn. Corona: System implications of emerging nanophotonic technology. In *Proceedings of ISCA '08*, pages 153–164, Washington, DC, USA, 2008. IEEE Computer Society.

[7] Scott Beamer, Chen Sun, Yong-Jin Kwon, Ajay Joshi, Christopher Batten, Vladimir Stojanović, and Krste Asanović. Re-architecting DRAM memory systems with monolithically integrated silicon photonics. In *Proceedings of ISCA '10*, pages 129–140, New York, NY, USA, 2010. ACM.

[8] Kyung-Soo Ha, Lee-Sup Kim, Seung-Jun Bae, Kwang-Il Park, Joo-Sun Choi, Young-Hyun Jun, and Kinam Kim. A 6gb/s/pin pseudo-differential signaling using common-mode noise rejection techniques without reference signal for dram interfaces. In *Solid-State Circuits Conference - Digest of Technical Papers, 2009. ISSCC 2009. IEEE International*, pages 138–139,139a, 2009.

[9] International Technology Roadmap for Semiconductors. System drivers. 2011.

[10] C. Kromer, G. Sialm, T. Morf, M.L. Schmatz, F. Ellinger, D. Erni, and H. Jackel. A low-power 20-ghz 52-db-ohm transimpedance amplifier in 80-nm cmos. *Solid-State Circuits, IEEE Journal of*, 39(6):885–894, June 2004.

[11] C. Kromer, G. Sialm, C. Berger, T. Morf, M. Schmatz, F. Ellinger, D. Erni, Gian-Luca Bona, and H. Jackel. A 100mw 4 times;10gb/s transceiver in 80nm cmos for high-density optical interconnects. In *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, pages 334–602 Vol. 1, 2005.

[12] A. Emami-Neyestanak, D. Liu, G. Keeler, N. Helman, and M. Horowitz. A 1.6 gb/s, 3 mw cmos receiver for optical communication. *VLSI Circuits Digest of Technical Papers, 2002. Symposium on*, pages 84–87, 2002.

[13] Cheng Li, Rui Bai, A. Shafik, E.Z. Tabasy, Geng Tang, Chao Ma, Chin-Hui Chen, Zhen Peng, M. Fiorentino, P. Chiang, and S. Palermo. A ring-resonator-based silicon photonics transceiver with bias-based wavelength stabilization and adaptive-power-sensitivity receiver. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2013 IEEE International*, pages 124–125, 2013.

[14] C. Debaes, A. Bhatnagar, D. Agarwal, R. Chen, G.A. Keeler, N.C. Helman, H. Thienpont, and D.A.B. Miller. Receiver-less optical clock injection for clock distribution networks. *Selected Topics in Quantum Electronics, IEEE Journal of*, 9(2):400–409, March-April 2003.

[15] S. Palermo, A. Emami-Neyestanak, and M. Horowitz. A 90 nm cmos 16 gb/s transceiver for optical interconnects. *Solid-State Circuits, IEEE Journal of*, 43(5):1235–1246, May 2008.

[16] I.A. Young, E. Mohammed, J.T.S. Liao, A.M. Kern, S. Palermo, B.A. Block, M.R. Reshotko, and P.L.D. Chang. Optical i/o technology for tera-scale computing. *Solid-State Circuits, IEEE Journal of*, 45(1):235 –248, jan. 2010.

[17] A. Rylyakov, C. Schow, B. Lee, W. Green, J. Van Campenhout, Min Yang, F. Doany, S. Assefa, C. Jahnes, J. Kash, and Y. Vlasov. A 3.9ns 8.9mw 4x4 silicon photonic switch hybrid integrated with cmos driver. In *ISSCC*, pages 222 –224, feb. 2011.

[18] Guoliang Li, Xuezhe Zheng, Jon Lexau, Ying Luo, Hiren Thacker, Thierry Pinguet, Po Dong, Dazeng Feng, Shirong Liao, Roshanak Shafiiha, and et al. Ultralow-power silicon photonic interconnect for high-performance computing systems. *Proceedings of SPIE*, 7607(Optoelectronic Interconnects and Component Integration IX):760703–760703–15, 2010.

[19] I. Young, E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, and P. Chang. Optical I/O technology for tera-scale computing. In *ISSCC*, pages 468 –469,469a, feb. 2009.

[20] F. Liu, D. Patil, J. Lexau, P. Amberg, M. Dayringer, J. Gainsley, H.F. Moghadam, Xuezhe Zheng, J.E. Cunningham, A.V. Krishnamoorthy, E. Alon, and R. Ho. 10 gbps, 530 fj/b optical transceiver circuits in 40 nm cmos. In *VLSI Circuits (VLSIC), 2011 Symposium on*, pages 290–291, 2011.

[21] K. Fukuda, H. Yamashita, G. Ono, R. Nemoto, E. Suzuki, N. Masuda, T. Takemoto, F. Yuki, and T. Saito. A 12.3-mw 12.5-gb/s complete transceiver in 65-nm cmos process. *JSSC*, 45(12):2838 –2849, dec. 2010.

[22] C.L. Schow, L. Schares, S.J. Koester, G. Dehlinger, R. John, and F.E. Doany. A 15-gb/s 2.4-v optical receiver using a ge-on-soi photodiode and a cmos ic. *Photonics Technology Letters, IEEE*, 18(19):1981–1983, Oct.1, 2006.

[23] J.F. Buckwalter, Xuezhe Zheng, Guoliang Li, K. Raj, and A.V. Krishnamoorthy. A monolithic 25-gb/s transceiver with photonic ring modulators and ge detectors in a 130-nm cmos soi process. *Solid-State Circuits, IEEE Journal of*, 47(6):1309–1322, 2012.

[24] S. Assefa, W.M.J. Green, A. Rylyakov, C. Schow, F. Horst, and Y.A. Vlasov. Monolithic integration of silicon nanophotonics with cmos. In *Photonics Conference (IPC), 2012 IEEE*, pages 626–627, 2012.

[25] Xuezhe Zheng, Ying Luo, J. Lexau, F. Liu, Guoliang Li, H.D. Thacker, I. Shubin, Jin Yao, R. Ho, J.E. Cunningham, and A.V. Krishnamoorthy. 2-pj/bit (on-chip) 10-gb/s digital cmos silicon photonic link. *Photonics Technology Letters, IEEE*, 24(14):1260–1262, 2012.

[26] M.H. Nazari and A. Emami-Neyestanak. A 24-gb/s double-sampling receiver for ultra-low-power optical communication. *Solid-State Circuits, IEEE Journal of*, 48(2):344–357, 2013.

[27] M. Georgas, J. Leu, B. Moss, Chen Sun, and V. Stojanovic. Addressing link-level design tradeoffs for integrated photonic interconnects. In *Custom Integrated Circuits Conference (CICC), 2011 IEEE*, pages 1–8, 2011.

[28] M. Georgas, J. Orcutt, R.J. Ram, and V. Stojanovic. A monolithically-integrated optical receiver in standard 45-nm soi. *Solid-State Circuits, IEEE Journal of*, 47(7):1693–1702, 2012.

[29] R. Soref et al. Electrooptical effects in silicon. *IEEE Journal of Quantum Electronics*, 23(1), January 1987.

[30] Q. Xu et al. 12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators. *Optical Society of America*, 15(2), 2007.

[31] M.R. Watts, D.C. Trotter, R.W. Young, and A.L. Lentine. Maximally confined silicon microphotonic modulators and switches. In *IEEE Lasers and Electro-Optics Society, 2008. LEOS 2008. 21st Annual Meeting of the*, pages 457 –458, nov. 2008.

[32] K.-L.J. Wong and C.-K.K. Yang. Offset compensation in comparators with minimum input-referred supply noise. *JSSC*, 39(5):837 – 840, may 2004.

[33] S.-H. Woo, H. Kang, K. Park, and S.-O. Jung. Offset voltage estimation model for latch-type sense amplifiers. *Circuits, Devices Systems, IET*, 4(6):503 –513, 2010.

[34] E. Alon, V. Stojanovic, and M.A. Horowitz. Circuits and techniques for high-resolution measurement of on-chip power supply noise. *Solid-State Circuits, IEEE Journal of*, 40(4):820 – 828, april 2005.

[35] S. Naffziger, B. Stackhouse, T. Grutkowski, D. Josephson, J. Desai, E. Alon, and M. Horowitz. The implementation of a 2-core, multi-threaded itanium family processor. *Solid-State Circuits, IEEE Journal of*, 41(1):197 – 209, jan. 2006.

[36] C. Kromer, G. Sialm, T. Morf, M.L. Schmatz, F. Ellinger, D. Erni, and H. Jackel. A low-power 20-GHz 52-dB-Ohm transimpedance amplifier in 80-nm CMOS. *JSSC*, 39(6):885–894, June 2004.

[37] Michael Georgas, Jonathan Leu, Benjamin Moss, Chen Sun, and Vladimir Stojanovic. Addressing link-level design tradeoffs for integrated photonic interconnects. In *Custom Integrated Circuits Conference (CICC), 2011 IEEE*, pages 1 –8, sept. 2011.

[38] S. Assefa, F. Xia, W.M.J. Green, C.L. Schow, A.V. Rylyakov, and Y.A. Vlasov. Cmos-integrated optical receivers for on-chip interconnects. *Selected Topics in Quantum Electronics, IEEE Journal of*, 16(5):1376 –1385, sept.-oct. 2010.

[39] A. Bhatnagar et al. Receiverless clocking of a cmos digital circuit using short optical pulses. In *LEOS '02*, volume 1, pages 127 – 128, 2002.

[40] F. Tavernier et al. High-speed optical receivers with integrated photodiode in 130 nm cmos. *JSSCC*, 44(10):2856 –2867, 2009.

[41] A. Strak et al. Analysis of timing jitter in inverters induced by power-supply noise. In *DTIS '06*, pages 53 –56, sept. 2006.

[42] Jason S. Orcutt et al. Nanophotonic integration in state-of-the-art CMOS foundries. *Opt. Express*, 19(3):2335–2346, Jan 2011.

[43] S. Selvaraja et al. Fabrication of uniform photonic devices using 193nm optical lithography in silicon-on-insulator. *ECIO '08*, 2008.

[44] Magdalena S. Nawrocka et al. Tunable silicon microring resonator with wide free spectral range. *Applied Physics Letters*, 89(7):071110, 2006.

[45] Vivek Raghunathan et al. Athermal silicon ring resonators. In *Integrated Photonics Research, Silicon and Nanophotonics*, page IMC5. Optical Society of America, 2010.

[46] Kyle Preston, Yoon Ho Daniel Lee, Mian Zhang, and Michal Lipson. Waveguide-integrated telecom-wavelength photodiode in deposited silicon. *Opt. Lett.*, 36(1):52–54, Jan 2011.

[47] J.L. Polleux and C. Rumelhard. Optical absorption coefficient determination and physical modelling of strained SiGe/Si photodetectors. In *High Performance Electron Devices for Microwave and Optoelectronic Applications, 8th IEEE International Symposium on*, pages 167 –172, 2000.

[48] S. Mukhopadhyay, R.V. Joshi, Keunwoo Kim, and Ching-Te Chuang. Variability analysis for sub-100nm PD/SOI sense-amplifier. *ISQED*, pages 488–491, March 2008.

[49] Jung-Sik Kim, Chi Sung Oh, Hocheol Lee, Donghyuk Lee, Hyong-Ryol Hwang, Sooman Hwang, Byongwook Na, Joungwook Moon, Jin-Guk Kim, Hanna Park, Jang-Woo Ryu, Kiwon Park, Sang-Kyu Kang, So-Young Kim, Hoyoung Kim, Jong-Min Bang, Hyunyoon Cho, Minsoo Jang, Cheolmin Han, Jung-Bae Lee, Kyehyun Kyung, Joo-Sun Choi, and Young-Hyun Jun. A 1.2V 12.8GB/s 2Gb mobile wide-I/O DRAM with 4x128 I/Os using TSV-based stacking. In *ISSCC*, pages 496 –498, feb. 2011.

[50] M. Jeeradit, J. Kim, B. Leibowitz, P. Nikaeen, V. Wang, B. Garlepp, and C. Werner. Characterizing sampling aperture of clocked comparators. In *VLSI Circuits, 2008 IEEE Symposium on*, pages 68 –69, june 2008.

[51] F. Paillet and T. Karnik. Low power and noise tolerant 20 gb/s cmos tia for short-distance optical interconnect. In *Mixed-Signal Design, 2003. Southwest Symposium on*, pages 49–53, Feb. 2003.

[52] B.R. Moss, Chen Sun, M. Georgas, J. Shainline, J.S. Orcutt, J.C. Leu, M. Wade, Yu-Hsin Chen, K. Nammari, Xiaoxi Wang, Hanqing Li, R. Ram, M.A. Popovic, and V. Stojanovic. A 1.23pj/b 2.5gb/s monolithically integrated optical carrier-injection ring modulator and all-digital driver circuit in commercial 45nm soi. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2013 IEEE International*, pages 126–127, 2013.

[53] C.W. Holzwarth, J.S. Orcutt, Hanqing Li, M.A. Popovic, V. Stojanovic, J.L. Hoyt, R.J. Ram, and H.I. Smith. Localized substrate removal technique enabling strong-confinement microphotonics in bulk si cmos processes. *Lasers and Electro-Optics, 2008 and 2008 Conference on Quantum Electronics and Laser Science. CLEO/QELS 2008. Conference on*, pages 1–2, May 2008.