Silicon-Photonics for VLSI Systems

by

Chen Sun

B.S. Electrical Engineering and Computer Sciences
University of California, Berkeley (2009)
S.M. Electrical Engineering and Computer Science
Massachusetts Institute of Technology (2011)

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

June 2015

Author .....
Department of Electrical Engineering and Computer Science
May 20, 2015

Certified by .....
Vladimir Stojanović
Associate Professor of Electrical Engineering and Computer Science
Thesis Supervisor

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

Silicon-Photonics for VLSI Systems

by

Chen Sun

Submitted to the Department of Electrical Engineering and Computer Science
on May 20, 2015, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Electrical Engineering and Computer Science

Abstract

As raw compute power of a single chip continues to scale into the multi-teraflop regime, the processor I/O communication fabric must scale proportionally in order to prevent a performance bottleneck. As electrical wires suffer from high channel losses, pin-count constraints, and crosstalk, they are projected to fall short of the demands required by future memory systems. Silicon-photonic optical links overcome the fundamental tradeoffs of electrical wires; dense wavelength division multiplexing (DWDM) – where multiple data channels share a single waveguide or fiber to greatly extend bandwidth density – and the potential to combine at chip-scale with a very large scale integrated (VLSI) CMOS electrical chip make them a promising alternative for next-generation processor I/O. The key device for VLSI photonics is the optical microring resonator, a compact micrometer-scale device enabling energy-efficient modulation, DWDM channel selection, and sometimes even photo-detection. While these advantages have generated considerable interest in silicon-photonics, present-day integration efforts have been limited in scale owing to the difficulty of integration with advanced electronics and the sensitivity of microring resonators to both process and thermal variations.

This thesis develops and demonstrates the pieces of a photonically-interconnected processor-to-memory system. We demonstrate a complete optical transceiver platform in a commercial 45 nm SOI process, showing that optical devices can be integrated into an advanced, commercial CMOS SOI process even without any changes to the manufacturing steps of the native process. To show that photonic interconnects are viable even for commoditized and cost-sensitive memory, we develop the first monolithic electronic-photonic links in bulk CMOS. As the stabilization of ring resonators is critical for use in VLSI systems, we contribute to the understanding of process and thermal variations on microring resonators, leading to the demonstration of a complete auto-locking microring tuning system that is agnostic to the transmitted data sequence and suitable for unencoded low-latency processor-to-memory traffic. Finally, the technology and methods developed in this work culminate in the demonstration of the world's first processor chip with integrated photonic intercon-

nects, which uses monolithically integrated photonic devices to optically communicate to main memory.

Thesis Supervisor: Vladimir Stojanović
Title: Associate Professor of Electrical Engineering and Computer Science

Acknowledgments

The work in this thesis has been an almost six-year journey that has spanned both east coast and west coast, at both MIT and UC Berkeley. During this time, I was surrounded by a number of top-notch researchers and a caring group of friends and family. Each of them contributed to this thesis in his/her own way and, as such, I owe thanks to these many individuals.

First, I would like to thank Professor Vladimir Stojanović, my thesis supervisor, for his mentorship and guidance throughout this endeavor. I am lucky to have been part of his group, which exposed me to a wealth of different research areas, well beyond just circuit design. Though his specific contributions to this thesis are too numerous to list, one particularly commendable quality of Vladimir is his boundless optimism. During the many disasters over the course of this thesis when things did not work out as planned, Vladimir always remained the everlasting beacon of hope who could motivate us to carry forward. Despite how often I found myself testing his resolve, Vladimir has shown me time and time again that the challenges we faced, no matter how daunting, were not insurmountable. In the end, the magnitude of success achieved in this thesis is, in no small part, due to Vladimir's absolute faith and commitment in making the project work.

The photonics project is a multi-team, multi-university effort and I had the pleasure to work alongside many brilliant collaborators. The work in this thesis would not have been possible without Jason Orcutt, the graduate student and research scientist who jump-started our integration efforts on commercial CMOS platforms. Jeffrey Shainline, who contributed the "one ring to rule them all" on both the bulk CMOS and SOI photonic platforms, was another critical person who helped us hit our early project milestones. Mark Wade, Luca Alloatti, Karan Mehta, Amir Atabaki, and Fabio Pavanello refined all our optical devices and contributed the high-performance devices that went into the full platform demonstrations and the the photonically interconnected processor chip. Professor Milos Popović and Professor Rajeev Ram provided technical knowledge, leadership, and guidance from the optics side that were critical to the demonstrations of the integrated photonic platforms. I also thank Erman Timurdogan, Jie Sun, and Professor Michael Watts for helpful discussions and for starting another potential photonics platform for demonstrating my ideas. On the foundry side, I could not have asked for better industry collaborators than Roy Meade, Ofer Tehar-Zahav, Daniel Miller, Zvika Sternberg, Reha Bafrali, and Gurtej Sandhu from Micron Technology. All were instrumental in getting the bulk photonics platform up and running. I would also like to thank our contacts at the Trusted Access Program Office and IBM Foundry Services, who had to work overtime to push the polygons from our designs though the foundry.

I consider myself lucky to have been a part of a crack team of circuit and systems designers. Michael Georgas, Benjamin Moss, and Jonathan Leu lead the circuits side of the project early on and I benefited greatly from their experience and guidance when I joined the project. Yu-Hsin Chen assisted me with composing digital link testing infrastructure on several chips and improved my understanding of microring

resonator thermal behavior through a thermal simulator tool we codeveloped. I thank Yunsup Lee, Andrew Waterman, Rimas Avizienis, and the rest of the RISC-V team for the design of the processor (the "programmable data generator") and for warping my sleep schedule around when we were getting our first RISC-V chips taped out and tested. I do not know of other research teams in the world who could hit the type of tapeout schedules we have had and it has been an amazing experience working with everyone on the team.

Before the photonics platforms started working in silicon, I had the opportunity to work on DSENT and Graphite with Chia-Hsin Owen Chen, George Kurian, Jason Miller, Professor Lionel Kimerling, Professor Anant Agarwal, and Professor Li-Shiuan Peh. These projects provided a high-level perspective on how photonics could fit into a larger system and identified the core unresolved issues to focus on in this thesis. I was also able to collaborate with Scott Beamer, Yong-Jin Kwon, Professor Ajay Joshi, Professor Chris Batten, and Professor Krste Asanovic to propose photonically interconnected memory, which defined the vision for the final system we built in this thesis. I also want to thank these same individuals for introducing me to the photonics project in my final year as an undergraduate at UC Berkeley and for enabling a quick start on the project when I had just arrived at MIT.

My fellow group members in the Integrated Systems Group, at both MIT and UC Berkeley, provided friendship and interesting philosophical discussions on a number of topics. It was especially interesting to sit next to Ranko Sredojevic, who complains about how terrible the tools are for designing hardware even more than I do. In addition I had the pleasure of sharing much of my time at BWRC with Krishna Settaluri, Sen Lin, Sajjad Moazeni, Nandish Mehta, Pavan Bhargava, and Taehwan Kim, all of whom helped with my transition to UC Berkeley and with getting our new lab up and running there.

Finally, I would like to thank all my friends and family for their unwavering support and for the occasional shoulder to cry on during the most stressful moments. I was also lucky enough to meet (or rather, re-meet) Michelle (Xiaoxi) during my time here and it was wonderful to have her be a part of this experience. Overall, graduate school has been exciting and constantly full of surprises, and I owe this experience to all who have helped to define it.

Support The work in this thesis is funded in part by DARPA, NSF, and the Berkeley Wireless Research Center

Contents

List of Figures

List of Tables

Chapter 1

Introduction

As CMOS technology continues to scale deep into the sub-100 nm regime, power density constraints have put a halt on clock frequency scaling. Performance scaling has since continued by leveraging limited instruction and data-level parallelism, manifested in ever-higher core counts (Figure 1-1) or a large number of specialized accelerators in order to take advantage of continued process improvements that enhance both transistor density and performance. With continued compute throughput scaling, the amount of interconnect bandwidth needed to supply the ever-increasing thirst for data increases as well, dominated by the chip-to-chip interconnect linking main memory to the processing chip. Today, the proportion of compute power to memory bandwidth offered by commercial high-performance processors is between 5–10 DP-FLOP/Byte [26, 60] (double precision floating-point operations per byte of memory bandwidth). As scaling continues, future >10 DP-TFLOPS systems will require in excess 1 TB/s of memory bandwidth.

Scaling to these bandwidths will be problematic due to a doubly constrained problem. Unlike logic gates and electrical transistors, copper wires improve marginally or not at all with each process generation. At the chip-to-chip level, a limited number of chip package pins and wire routability in the chip package or the mounted printed circuit board have forced high-performance chips to push for increasingly higher data-rates on each available signal pin (Figure 1-2.) However, electrical wires are fundamentally limited by the tradeoff between signal data-rate and channel loss.

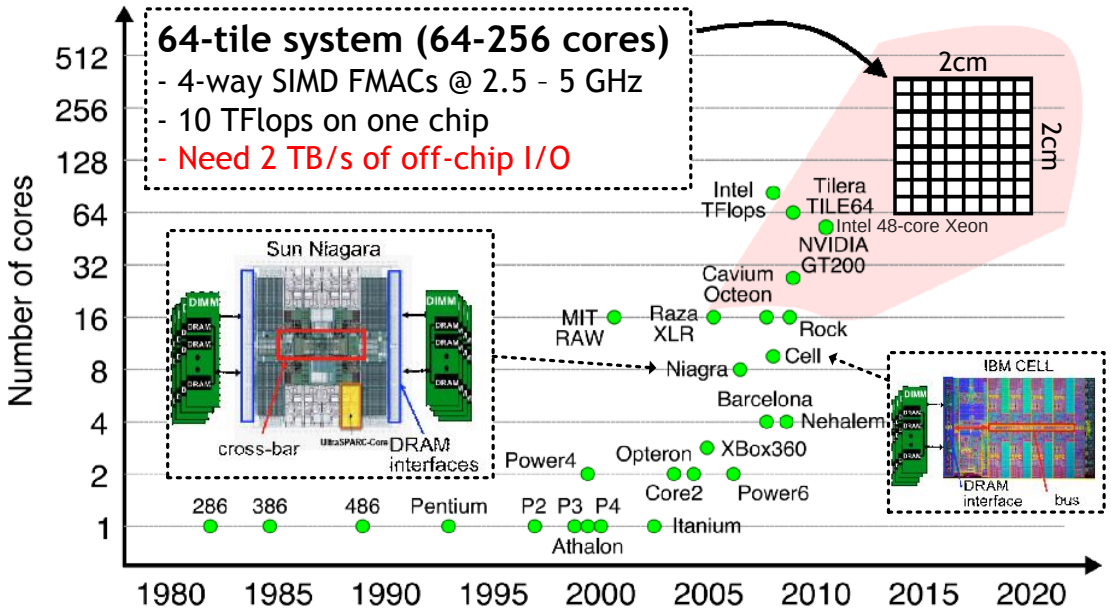

Figure 1-1: Processor core count scaling into the many-core era

The channel loss further grows with both distance – due to non-zero resistive and dielectric losses of the wire – and from the socket-based packaging architecture of chips – due to contact points creating discontinuities in the transmission lines. While techniques such as equalization are effective in extending the achievable data rate per pin, interconnect power is already consuming a significant fraction of the chip's power budget in modern processors [1, 23]. The power cost to push to even higher data rates and regimes becomes unsustainable (Figure 1-3). As a result, the combination of pin constraints and a chip's total power limit losses limit the achievable chip-to-chip interconnect bandwidth; without a sufficiently scalable approach for bandwidth density and energy-efficiency, high-performance chips, such as processors, are doomed to bandwidth starvation.

Silicon photonics is a highly disruptive technology for chip-to-chip communications. Nanoscale integration of optical devices with CMOS electronics enables high channel capacity through dense wavelength-division multiplexing (DWDM) and low-loss bit transport over long distances. As such, there is significant interest in lever-

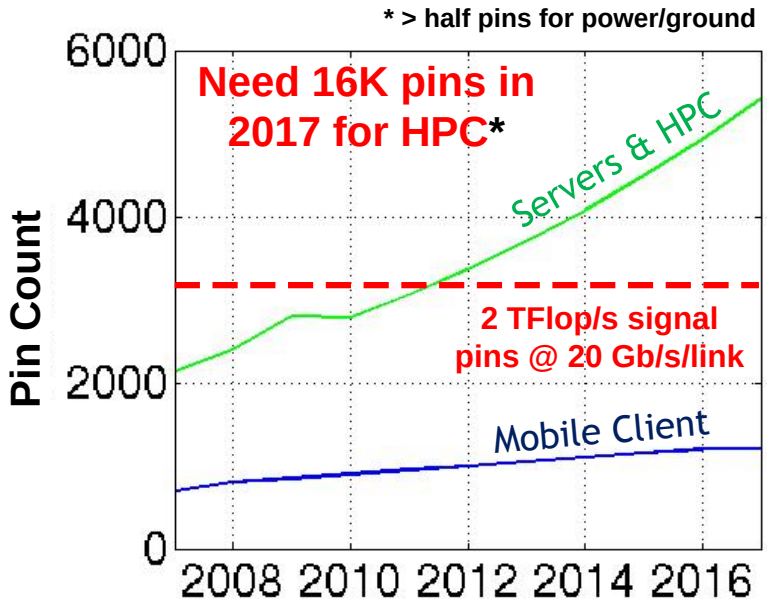

Figure 1-2: Package pin count scaling [27]. Package pins are needed for both power delivery and electrical signaling.

aging this technology to satisfy the interconnect demands for future many-core systems [6, 29, 35, 85]. While the advantages of optical signaling over electrical signaling are pronounced, two major technological hurdles prevent the adoption of optics for at the chip-to-chip level inside a VLSI system.

The first hurdle is the issue of electronic-photonic integration. While independent and discrete optical devices are suitable for long haul applications, where only a few devices are needed and where the size of the endpoints is a secondary factor to performance, optical devices for chip-to-chip communications must be integrated tightly with nanoscale electronics. Integration approaches can be grouped into two categories: monolithic integration or heterogeneous integration. The monolithic approach, where optical devices and electronics sit on the same chip, has been traditionally difficult due to chip manufacturing conflicts between optical devices and electronics. The heterogeneous approach, where optical devices and electrical devices sit on separate chips, suffers from 3D-integration and chip-stacked packaged complexity. As such, a VLSI electronic-photonic platform has yet to be developed. The second major hurdle stems from thermal and process variations; DWDM links rely upon compact,

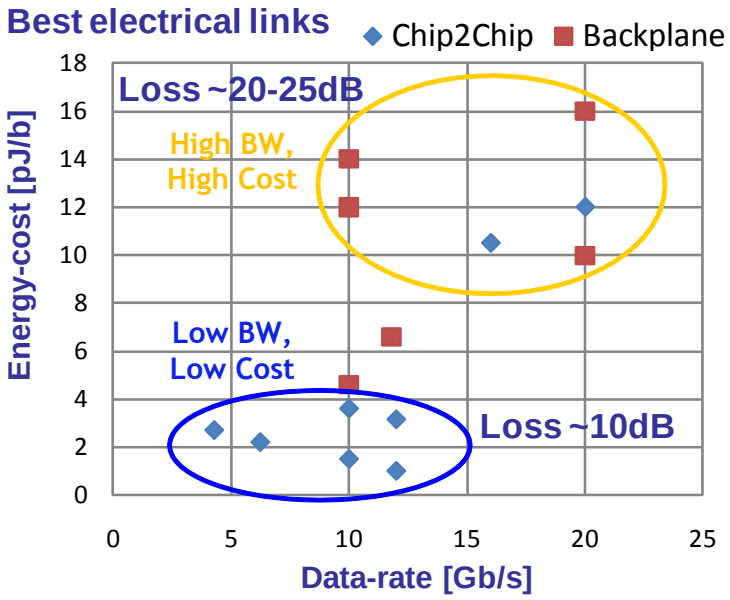

Figure 1-3: Chip-to-chip interconnect efficiency vs. data-rate. Source: ISSCC

high quality-factor microring resonator devices whose resonances drift around with process variations and temperature changes. For this reason, current optical systems rely upon millimeter-scale Mach-Zehnder interferometer devices for modulation and bulky arrayed waveguide gratings for coarse wavelength-division-multiplexing. In order for VLSI-compatible DWDM systems to profilerate, methods for resonance stabilization must be found.

This thesis will develop methodologies and platforms for integrating photonics in a VLSI system. In the first half of this thesis, I will develop an understanding of microring tuning. In particular, I will focus on the effects of laser-induced self-heating and their implications on methods for microring control. Using these insights, I will propose a new methodology for stabilizing ring resonators to enable their use inside a VLSI electronic-photonic system. In the second half of the thesis, I present the architecture and system-level demonstrations of electronic-photonic integration platforms, which include the implementation of the proposed tuning scheme in silicon. The integration platforms span both commercial SOI and bulk CMOS processes. These represent commodity and high-performance platforms, respectively, suitable for the

processor and memory sides of an optically-connected main memory system. Finally, I will show a complete demonstration of a photonically interconnected processor-to-memory system, which marks the first demonstration of a processor chip with photonic interconnects.

1.1 Thesis Organization and Contributions

Chapter 2 of the thesis provides an overview of integrated photonics technology. We introduce, in detail, the fundamentals of an optical microring resonator, a critical building block of integrated photonic links. We also discuss the tradeoffs and challenges of electronic-photonic integration.

Chapter 3 develops the models for understanding the dynamics of microring resonator wavelength stabilization, leading to the proposal of a bit-statistical tracking and tuning approach in Chapter 4. This approach was used and demonstrated in silicon in:

- [78] C. Sun, M. Wade, M. Georgas, S. Lin, L. Alloatti, B. Moss, R. Kumar, A. Atabaki, F. Pavanello, R. Ram, M. Popović and V. Stojanović. A 45nm SOI Monolithic Photonics Chip-to-Chip Link with Bit-Statistics-Based Resonant Microring Thermal Tuning. *2015 Symposium on VLSI Circuits [to appear]*. June 2015.

To address both the low-cost high-volume and the high-cost high-performance needs of VLSI photonics integration, we demonstrate its viability on two independent electro-optic platforms. Chapter 5 describes our monolithically-integrated bulk CMOS photonics platform, representing low-cost applications such as memory or other commodity bulk CMOS parts, and the first chip-to-chip optical link demonstrated in bulk CMOS. Parts of this chapter appear in:

- [76] C. Sun, M. Georgas, J. Orcutt, B. Moss, Y.-H. Chen, J. Shainline, M. Wade, K. Mehta, K. Nammari, E. Timurdogan, D. Miller, O. Tehar-Zahav, Z. Sternberg, J. Leu, J. Chong, R. Bafrali, G. Sandhu, M. Watts, R. Meade,

M. Popović, R. Ram, V. Stojanović. A Monolithically-Integrated Chip-to-Chip Optical Link in Bulk CMOS. *IEEE Journal of Solid-State Circuits*. April 2015.

- [77] C. Sun, M. Georgas, J. Orcutt, B. Moss, Y-H. Chen, J. Shainline, M. Wade, K. Mehta, K. Nammari, E. Timurdogan, D. Miller, O. Tehar-Zahav, Z. Sternberg, J. Leu, J. Chong, R. Bafrali, G. Sandhu, M. Watts, R. Meade, M. Popović, R. Ram, V. Stojanović. A Monolithically-Integrated Chip-to-Chip Optical Link in Bulk CMOS. *2014 Symposium on VLSI Circuits*. June 2014.

- [49] R. Meade, J. Orcutt, K. Mehta, O. Tehar-Zahav, D. Miller, M. Georgas, B. Moss, C. Sun, Y.-H. Chen, J. Shainline, M. Wade, R. Bafrali, Z. Sternberg, G. Machavariani, G. Sandhu, M. Popović, R. Ram, V. Stojanović *2014 Symposium on VLSI Technology*. June 2014.

- [80] C. Sun, E. Timurdogan, M. Watts, V. Stojanović. Integrated Microring Tuning in Bulk CMOS. *IEEE Optical Interconnects Conference*. May 2013.

Chapter 6 discusses a thin-BOX SOI photonics platform fabricated in a 45nm commercial CMOS foundry, representing a high-end part, and is the first such platform in a deep sub-100nm production CMOS. Parts of this chapter appear in:

- [78] C. Sun, M. Wade, M. Georgas, S. Lin, L. Alloatti, B. Moss, R. Kumar, A. Atabaki, F. Pavanello, R. Ram, M. Popović and V. Stojanović. A 45nm SOI Monolithic Photonics Chip-to-Chip Link with Bit-Statistics-Based Resonant Microring Thermal Tuning. *2015 Symposium on VLSI Circuits [to appear]*. June 2015.

- [19] M. Georgas, B. Moss, C. Sun, J. Shainline, J. Orcutt, M. Wade, Y.-H. Chen, K. Nammari, J. Leu, A. Srinivasan, R. Ram, M. Popović, V. Stojanović. A monolithically-integrated optical transmitter and receiver in a zero-change 45nm SOI process *2014 Symposium on VLSI Circuits*. June 2014.

- [62] J. Orcutt, B. Moss, C. Sun, J. Leu, M. Georgas, J. Shainline, E. Zgraggen, H. Li, J. Sun, M. Weaver, S. UroŽevic, M. Popović, R. Ram, V. Stojanović *Optics Express*. May 2012.

The knowledge and building blocks developed over the course of this thesis all culminate in Chapter 7, where we demonstrate the world's first processor with integrated photonic interconnects.

Chapter 2

Background

2.1 Silicon Photonic Interconnects

For the past several decades, optical communication has been steadily replacing their electrical counterparts. The benefits of optical signaling over electrical signaling are clear: low channel losses for high data-rates over long distances, immunity to electromagnetic interference, and high bandwidth capacity per cable. As communication bandwidth requirements continue to scale, the advantages of optics over electronics become more pronounced for gradually shorter and shorter distances. This has lead to the replacement of electrical cables with optical fibers in modern long- and medium-range communication systems, beginning with the first trans-atlantic cable [69] and expanding to server rooms and racks [70] in datacenters. Today, the fundamental tradeoff between signaling data-rates and channel loss in copper wires bottleneck the performance and energy-efficiency of even very short-reach chip-to-chip links that span only a few centimeters, such as the interface between processor and main memory. Despite the numerous challenges faced by high-speed short-reach electrical interconnects, this space has traditionally been untouched by optics due to the sheer number of optical components needed per chip and the bulkiness of discrete optical devices.

In recent years, the development of silicon-based photonics has enabled the possibility of direct optical chip-to-chip communication. The key enabler of silicon pho-

tonics lies in its inherent compatibility with modern CMOS integrated circuit manufacturing; like electrical transistors, thousands of micrometer-scale silicon photonic devices can be fabricated on a microchip using the same lithographic processes used to make electronics. DWDM, in particular, provides a method for overcoming pin-limited bandwidth density limitations of modern microchips.

2.1.1 Waveguides

A waveguide routes near-IR light on a chip and is the fundamental building block of all on-chip optical devices. Waveguides used in silicon-photonic microchips consist of a waveguide core, which has a high index of refraction, surrounded by cladding material of a lower index of refraction (Figure 2-1). A high index contrast between the core and the cladding material allows for bend radii on the order of a few microns to be achieved with minimal radiative losses, enabling optical routing within tight geometries, such as that of a chip. An evanescent field extends out from the waveguide core into the cladding material, allowing waveguides to optically interact with each other when brought together in close enough proximity. The field decays exponentially the farther it is from the waveguide core and the amount of coupling between interacting waveguides can be adjusted using the spacing. They extent of the evanescent fields also defines the minimum spacing between two non-interacting waveguides and the minimum cladding thickness. Waveguide losses stem from the bulk material absorption loss, line-edge etch induced sidewall roughness scattering loss, surface scattering loss, and bending loss. The presence of free electrons or holes (from a doped silicon waveguide or from a piece of metal in close proximity to the waveguide) can introduce an additional loss from free carrier absorption. Waveguides used for silicon-photonics are usually designed to be single-mode, though multi-mode structures can be used in straight waveguides to create low-loss "highways" or waveguide crossings [47] or create beat-patterns that avoid certain lossy obstacles such as contacts or high-doping regions in active devices [46, 72].

The choice of materials for the waveguide core and the cladding varies greatly depending on the silicon-photonic platform. The waveguide core is silicon, which

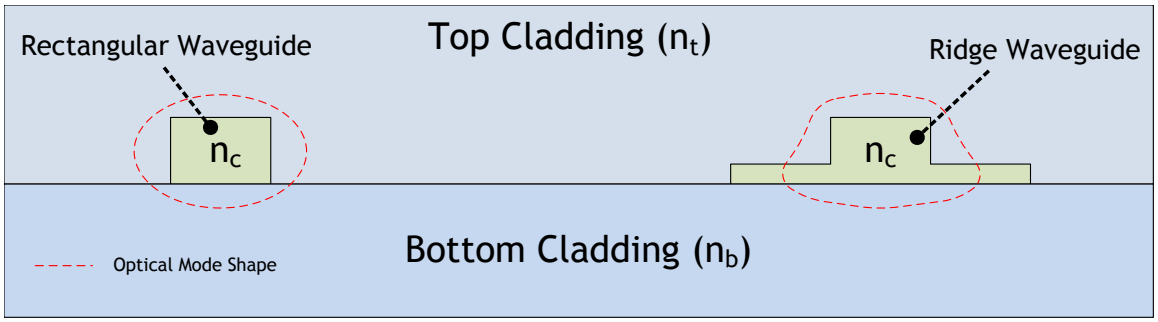

Figure 2-1: Rectangular and ridge waveguides, two examples of waveguide geometries for silicon photonics, and their respective optical mode profiles. Both waveguide geometries can support an optical mode if $n_c > n_t$ and $n_c > n_b$.

has a high refractive index, and can be crystalline [4, 8, 62], epitaxially-grown crystalline [9], polysilicon [61], or amorphized polysilicon [50, 76]. The top and bottom cladding cladding material can be a combination of silicon dioxide, nitrides, deposited polymers [67], and even air [19, 61], all of which have a lower refractive than silicon.

2.1.2 Chip-to-Fiber Couplers

Vertical grating couplers [81, 87] allow light to be directed both into and out-of the plane of the chip. Light traveling in waveguides on the chip can be coupled (at an angle) into optical fibers to be brought off-chip and off-chip light can be coupled into on-chip waveguides. While many different geometries of grating couplers exist, the goal of a grating coupler is to create a structure that shapes the optical on-chip waveguide mode into the shape that an optical fiber expects.

2.1.3 Laser Sources

Both on-chip and off-chip lasers are possible options as sources of light for on-chip optics. Leveraging the existing fiber-optic communication infrastructure, uncooled off-chip continuous wave (CW) lasers are commercially available at reasonably high efficiencies [66]. Integrated on-chip laser sources can also be found in the form of Ge-based [11, 45] or through a myriad of hybrid silicon/indium phosphide [25, 44] laser variants. While on-chip lasers have markedly lower efficiencies due to process

constraints and device immaturity, we note that close proximity to electronics enables lower laser distribution losses (no coupling loss to go on-chip, in the case of off-chip lasers). In addition, the proximity of on-chip lasers to electronics creates the possibility of direct laser modulation or aggressive laser power gating techniques [36] to reduce static laser power consumption.

Lasers are not limited to a single wavelength output per laser; comb-based lasers can output a full wavelength comb [92], producing a set of wavelengths needed to power up all the wavelengths of a DWDM link on its own. In general, because the total output power of the laser is limited, the maximum power per wavelength produced by a comb laser is less than what can be achieved by individual, single-wavelength lasers. In addition, comb lasers cannot provide the same output power on all of its output wavelengths; the output power spectrum tracks the broadband gain characteristics of the laser's gain medium and the output power per wavelength falls off quickly for wavelengths on the edge of the gain spectra.

We note that a laser's output power and its output wavelength(s) are strongly dependent on its temperature, which affects nearly all aspects of the laser, from the gain material to the physical dimensions of the laser cavity due to thermal expansion. As such, a laser diode is often paired with a thermoelectric cooler connected in a temperature-stabilized control loop. The power cost of the cooler is significant, often times dominating the overall wall-plug efficiency of the laser. If the temperature control is omitted – typical for on-chip lasers or for power savings in off-chip lasers – the output wavelengths will drift with temperature. In a comb laser, a change in temperature moves all the comb wavelengths together, which maintains the channel spacing to the first order. Wavelengths from a bank of individual lasers will move independently of one another depending on the temperature of each laser. The spacing between wavelengths changes as a function of the temperature difference between lasers. As a result, the bank of lasers needs to share the same thermally conductive substrate, otherwise wavelengths will collide with each other if the temperature difference is too great.

In addition to the potential efficiency loss due to cooling, the power-loss due to the finite turn-on threshold of the laser diode favors the use of higher-power lasers, which is in contrast to the low power-per-wavelength requirements of individual photonic links. Silicon-photonics and tight-integration with transistors offers the opportunity to utilize the higher-power lasers and locally on-chip manage and administer the laser power to the parts of the system where this power is needed [5, 6].

We note that despite the many laser designs currently available, the design of an optimal laser system for use in high-density integrated silicon-photonic systems is still an active research topic. Such lasers should be tightly codesigned to coordinate with the capabilities of the photonic resonance tracking loops and laser power management capabilites found in complex silicon-photonic systems.

2.1.4 Optical Microring Resonator

The key building block of a DWDM link is the optical microring resonator, which is a waveguide looped around in a ring to form a resonant cavity. As opposed to traditional WDM multiplexers implemented using cumbersome arrayed waveguide gratings [17], which are on the scale of hundreds of micrometers on a side, ring radii in DWDM links are on the order 5 µm, which allows for thousands of ring resonators to fit on a single die.

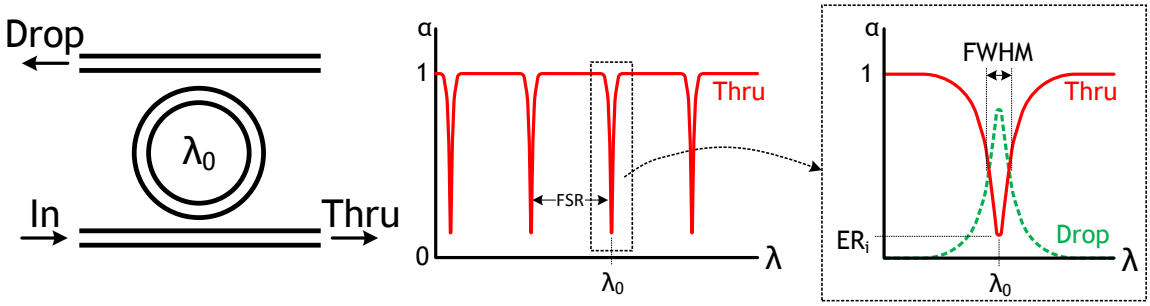

Figure 2-2: The optical transfer characteristics of a microring resonator (a waveguide looped around in a ring to form a resonating cavity) that is resonant at λ_0.

At a high level, a microring coupled to a bus waveguide acts as a notch filter from in-port to through-port, shown in Fig. 2-2; wavelengths not on resonance are allowed

to pass by while wavelengths close to the resonant wavelengths are captured in the resonator. Resonances are periodic with a spacing defined as the free spectral range (FSR), which grows with a smaller microring circumference. Due to this periodicity, all wavelengths used in a DWDM link must fit within one FSR, where each ring has single-wavelength selectivity.

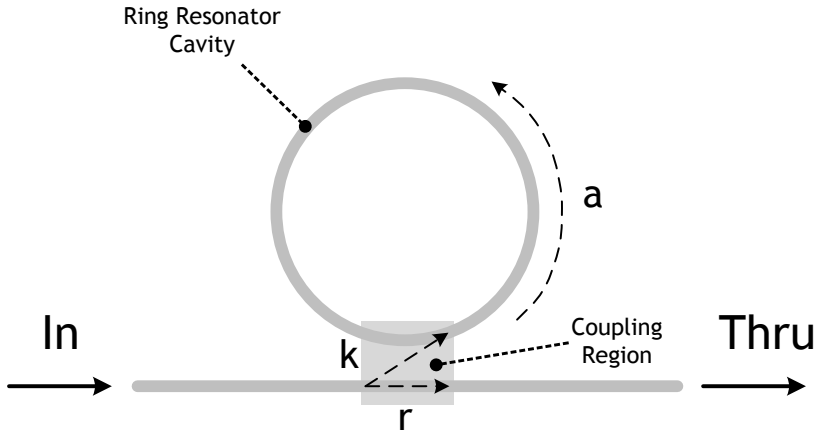

Figure 2-3: Microring resonator coupled to a bus waveguide through a coupling region.

We can derive quantitative relationships of a microring resonator system. In Figure 2-3, we define quantities a, r, and k, which correspond to the transmission of the microring for one roundtrip, the self-coupling of the bus waveguide, and the cross-coupling between the bus waveguide and the microring waveguide, respectively. Note that $r^2 + k^2 = 1$ as r^2, k^2 are the power splitting ratios of the coupler and the coupling region is assumed to be lossless. Using these quantities, we can derive the microring's input port to through port wavelength transmission characteristics, α, following the analysis given in [7]:

$$\alpha = \frac{I_{thru}}{I_{input}} = \frac{a^2 - 2ar\cos(\phi) + r^2}{1 - 2ar\cos(\phi) + a^2r^2} \tag{2.1}$$

I_{thru} and I_{input} are the intensities at the through port and input port, respectively. The quantity ϕ is the single-pass phase shift provided by the microring waveguide

and is given by $\phi = BL$. B is the propagation constant, equivalent to $2\pi n_{eff} L/\lambda$, and L is the round-trip length of the microring. n_{eff} is the effective index.

The ring is on resonance when the phase shift ϕ is a multiple of 2π. This condition is satisfied for a set of resonant wavelengths, λ_0, when:

$$\lambda_0 = \frac{n_{eff} L}{m}, \quad m = 1, 2, 3, ... \tag{2.2}$$

These multiple resonances give rise to the free spectral range, which is the spacing between resonances, given by:

$$FSR = \frac{\lambda^2}{n_g L} \tag{2.3}$$

where n_g is the group index, which takes into account the dispersion of the silicon waveguide. When on resonance, Equation 2.1 reduces to:

$$\alpha(\lambda_0) = 1 - A = \frac{a^2 - 2ar + r^2}{1 - 2ar + a^2 r^2} = \frac{(a - r)^2}{1 - 2ar + a^2 r^2} \tag{2.4}$$

When the ring is critically coupled, defined as when the power coupled from the bus waveguide matches the energy decay rate due to losses in the microring, $r = a$ and $\alpha(\lambda_0) = 0$. If the ring is overcoupled ($r > a$) or undercoupled ($r < a$), $\alpha\lambda_0 > 0$. Consequently, in order for a ring to have the highest extinction in its notch response, the ring needs to be critically coupled. a will vary depending on the loss of the microring cavity, which is a function of many factors. To critically couple a ring, we can tune r and k by changing the spacing between the waveguide and microring to adjust the amount of evanescent coupling. Note that in Equation 2.4, we have also defined a new quantity, A, corresponding to the decrease in through-port transmisivity (from unity) when the ring is on resonance. We can also define a microring's intrinsic extinction ratio, ER_i, such that $ER_i = -10 \cdot \log_{10}(1 - A)$. When the ring is critically coupled, $A = 1$ and ER_i is infinite.

A microring's full-width half-maximum (FWHM) bandwidth, $\Delta\lambda$, around a resonance λ_0 is indicative of the sharpness of the resonance notch and the microring's

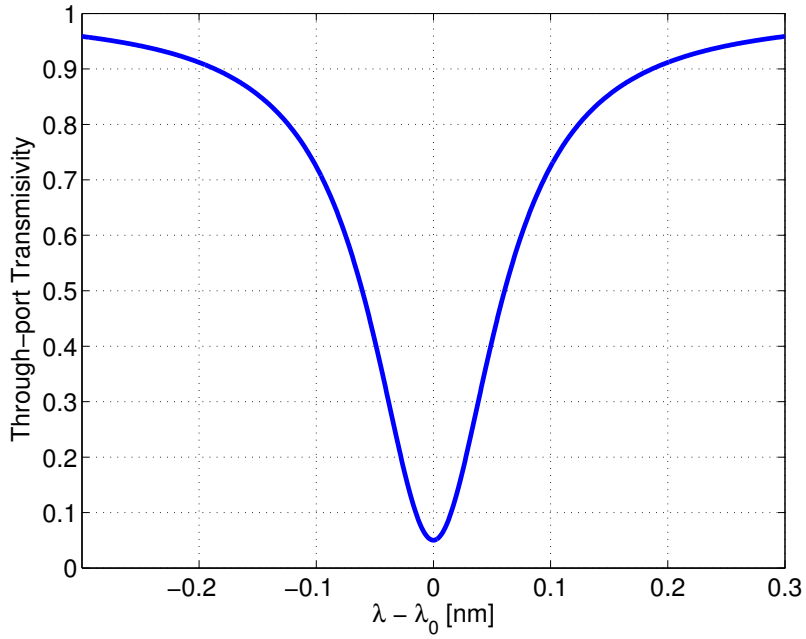

Figure 2-4: Plot of a ring resonator's through-port transfer characteristic, α, using Equation 2.7 versus $\lambda - \lambda_0$. In this plot, we use $A = 0.95$, corresponding to an $ER_i = 20\,\text{dB}$. The linewidth ($\Delta\lambda$) is 128 pm.

selectivity. From [7], this is given by:

$$\Delta\lambda = \frac{(1 - ra)\lambda_0^2}{\pi n_g L\sqrt{ra}} \tag{2.5}$$

. The quality factor (Q-factor) of the microring is related to the $\Delta\lambda$ FWHM through:

$$Q = \frac{\lambda_0}{\Delta\lambda} \tag{2.6}$$

The quality factor of the ring is inversely proportional to the loss of the ring cavity; as a decreases, which indicates lower round-trip transmission and higher loss in the microring, $\Delta\lambda$ increases and Q drops. Doping the ring in active structures (modulators or a resonated photodetector), for example, will lower the quality factor by introducing additional round-trip loss due to free carrier absorption. Likewise, electrical contacts placement must avoid overlap with the optical mode without sacrificing contact resistance. This mechanism also sets a lower-bound on the radius of the ring;

a radius that is too small will suffer from very high bend losses and a low quality factor.

While Equation 2.1 captures the full dynamics of a microring, it is cumbersome to work with due to the use of low-level physical quantities as opposed to the characteristic quantities that can be derived from a wavelength spectrum measurement (A, λ_0, Q, and $\Delta\lambda$). For the rest of the quantitative analysis in this thesis, we will instead adopt a simplified model which fits a microring's through-port transfer characteristic, α_λ to a Lorentzian line shape [8] using the characteristic quantities:

$$\alpha(\lambda) = 1 - \frac{A}{1 + 4\left(\frac{\lambda - \lambda_0}{\Delta\lambda}\right)^2} \tag{2.7}$$

We note that this model is fitted around a single resonant wavelength, λ_0 and does not take into account the other resonant wavelengths from the same microring. Nevertheless, this model provides an accurate approximation since we avoid placing DWDM channels outside of a microring's FSR in order to maintain single wavelength selectivity.

2.1.5 Microring Modulators

We can realize an electro-optic microring modulator by modulating the resonance (λ_0) of the ring to perform on-off-keying (OOK) of input light aligned close to λ_0 (Fig. 2-5, left). Changes in free carrier concentration are used to shift λ_0 by changing the material's index of refraction [74]. Carrier-injection modulators are p-i-n junctions that inject carriers into the intrinsic region during forward-bias, blue-shifting λ_0. Carrier-depletion modulators are p-n junctions that deplete the carriers from the junction during reverse bias, red-shifting λ_0. Carrier-injection modulators are limited in speed by minority carrier lifetimes, necessitating pre-emphasis schemes to reach higher data-rates [40, 54, 94]. Forward-biased operation of the junction also results in static power dissipation and poor energy-efficiency. Carrier-depletion designs avoid these issues, but require mid-level doping control (often difficult in a CMOS process) to balance λ_0 shift with Q-factor degradation from free carrier absorption. Figure 2-

6 shows a simulated through-port transfer characteristic that is representative of a depletion modulator, showing a shift in the ring's resonance between the *0* and *1* states. Here, the *1* state corresponds to the depleted state as it is red-shifted relative to the *0* state. We note that a higher Q-factor yields a greater modulation depth given the same change in λ_0 but also translates to a higher lifetime for photons resonating in the ring. Photon lifetimes comparable to or greater than the bit time result in optical ISI due to residual light left in the ring from bit-to-bit. A ring's Q-factor sets an optical bound on its maximum data-rate.

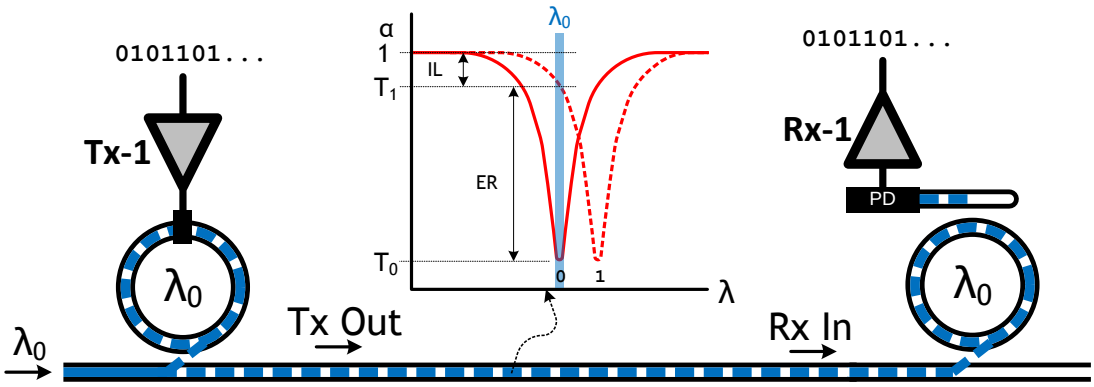

Figure 2-5: Ring resonators used in a modulator and a receiver. The input light is at the same wavelength as λ_0. The modulation insertion loss (IL) and the extinction ratio (ER) are defined as $10 \cdot \log_{10} T_1$ and $10 \cdot \log_{10} \frac{T_1}{T_0}$, respectively. T_1 and T_0 represent modulated output powers for optical ones and zeros, normalized to the input power.

Compared to traditional Mach-Zehnder interferometer-based modulator structures [15, 56], which even the smallest are hundreds of micrometers long, the resonant structure of rings increases the optical length, allowing them to perform modulation in a much smaller form factor.

2.1.6 Photodetectors and Receivers

A photodetector is responsible for converting optical power into electrical current, which can then be sensed by a receiver circuit [20, 40, 59] and resolved to electrical ones and zeros. Both pure germanium and SiGe (which are already used in modern processes for strain engineering) are potential photodetector materials [33, 62,

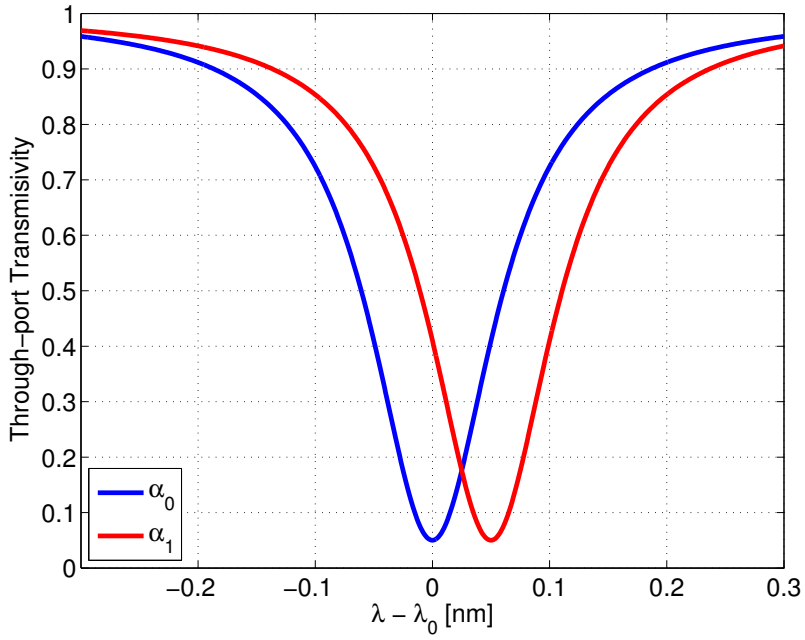

Figure 2-6: Plot of a ring modulator's through-port transfer characteristic versus $\lambda - \lambda_0$. Note that the modulator has two states, α_0 and α_1, corresponding to the modulator's *0* and *1* states, respectively. In this plot, we use $A = 0.95$, corresponding to an $ER_i = 20\,\text{dB}$. The linewidth ($\Delta\lambda$) is 128 pm.

84]. Defect-trap-based poly-Si [51] and photonic crystal photodetectors [52] are also promising silicon-based photodetector of interest.

As the absorption spectrum of the detector depends on the material itself, standalone photodetectors are wideband. Instead, to obtain wavelength selectivity, photodetectors are always paired with microring resonators, which performs the channel selection (Fig. 2-5, right). The photodetector is placed on a waveguide that is coupled to the drop-port of a microring, such that light resonating in the ring will get dropped onto the the drop-port waveguide and illuminate the photodetector. This creates an arrangement where only light at the resonant wavelength of the ring illuminates the photodetector, whereas all other wavelengths pass by. Alternatively, the PD may be directly embedded into the structure of the ring itself [20, 51]. In addition to providing channel selectivity, the ring also enhances the absorption of the photodetector, as resonating light makes multiple passes through it. As a result, we can reduce the overall size of the photodetector, minimizing its footprint and its parasitics, to improve

overall receiver performance [18, 20]. We note that integration of the photodetector material (typically Germanium) into the microring structure can be challenging. However, this method has been very attractive method for polysilicon-based defect detectors [51]. The trade-off here is between the drop-loss of the passive ring and absorption length (capacitance of the slab PD), versus a broader spectral response (lowered Q) of the resonant PD, in potential DWDM applications.

2.1.7 Microring Tuning Requirements

In order to perform its roles, a microring resonator's λ_0 must be aligned to λ, the laser wavelength it is modulating or filtering. Unfortunately, a practical system contributes several factors that invalidates this assumption. First, as a resonant device, λ_0 is very sensitive to process (geometric) variations. Second, as the laser itself may not be temperature stabilized, the output wavelength of the laser could drift time with time. Third, an increase or decrease in temperature changes the index of refraction through the thermo-optic effect, creating large temperature-dependent offsets in λ_0. The strong thermal dependence provides a way for us to tune λ_0 to combat the first two sources of λ_0 misalignment but also necessitates active resonance control when microrings are integrated into an electrical system, where temperatures fluctuate. We note to the reader that ring tuning requirements will be covered in much greater detail in Chapter 3.

2.1.8 A Silicon Photonic Link

The architecture of an n-wavelength DWDM chip-to-chip silicon-photonic link using an off-chip laser source is shown in Fig. 2-7. The laser source—either a comb laser or a bank of single-wavelength lasers shared across all links in the system—produces CW light of n wavelengths (λs). The laser light travels from the off-chip laser source to the chip through a single-mode fiber. The light then couples into an on-chip single-mode waveguide through a vertical grating coupler (VGC). An n-λ DWDM transmit macro, built using a bank of n resonant microrings, each tuned to

one of the wavelengths, modulates the wavelengths, imprinting an independent digital bitstream upon each of the wavelengths. The modulated wavelengths of light exit the transmit chip through a second VGC into a single-mode fiber bound for the receive chip. Once at the receive chip, they couple into an n-λ DWDM receive macro. Here, resonant filter microrings tuned to each wavelength drop the light onto photodetectors to produce photocurrent, which the receivers resolve back into data. To simplify clock recovery, the clock can also be source-forwarded as one of the datastreams using one of the wavelengths [18, 39].

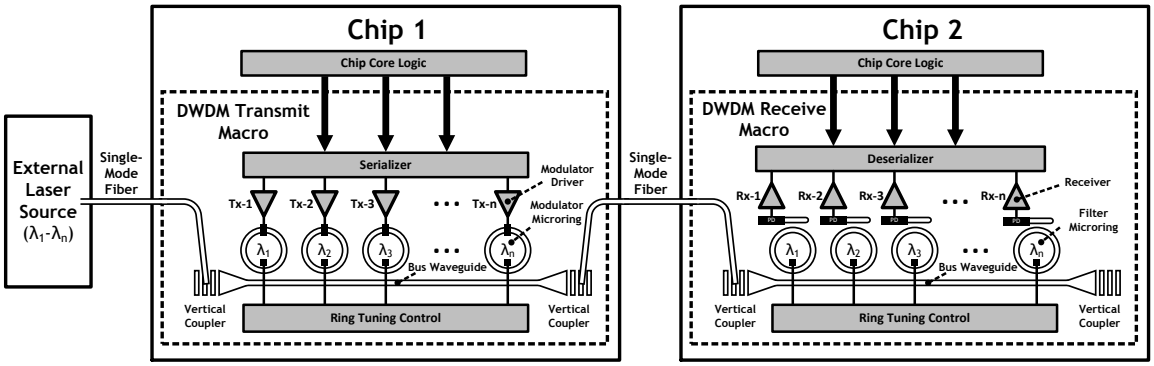

Figure 2-7: A chip-to-chip DWDM optical link using silicon-photonics.

2.2 Silicon Photonic Integration

Despite the numerous efforts underway [4, 9, 19, 43, 84], the integration of optical devices with advanced electronics in a VLSI system remains a key challenge. We can classify integration strategies as either monolithic, with optical devices and transistors both on the same chip [4, 19, 43], or heterogeneous, with separate photonic and electronic dies bonded in a multi-chip solution [40, 59, 98], each with their own advantages. In addition to the integration method, another degree of freedom is the choice of the substrate for building optical devices on. Here, both silicon-on-insulator (SOI) and bulk substrates are promising and address different needs. This section provides a brief overview of the various integration strategies.

2.2.1 Heterogeneous vs. Monolithic Integration

In a heterogeneous platform, the electronic wafer and the photonic wafer are fabricated separately in different processes or foundries and attached together using through-silicon-vias (TSVs) [32, 93], through-oxide-vias (TOVs) [84], or face-to-face microbumps [98]. The split nature of fabrication enables both sides to be individually optimized, allowing process optimizations for photonics to be completely decoupled from those of the CMOS electronics it attaches to. As optical devices are not bound by CMOS processing constraints, heterogeneous platforms can demonstrate better standalone device performance than on a monolithic platforms. The main downside of heterogeneous integration is the overhead of 3D integration or packaging. Aside from the added cost and yield penalty, the chip-to-chip interconnect adds parasitic capacitance [98] both from the 3D interconnect – 20–50 fF for microbumps but can be as low as 3 fF for wafer-level stacked TOVs [84] – and from the wiring through the entire metal stack of the CMOS chip (often more than 10 metal layers) to reach where the microbumps and the TSVs contact the chip. State-of-the-art 3D packaging solutions are pitch limited – 4 µm for wafer-level stacked TOVs and 40 µm for microbumps – which inherently limit the density of contacted optical devices. These factors ultimately degrade the performance and energy efficiency [18] of the final part. As such, the benefits of heterogeneous photonics must be weighed against the complexity of the packaging.

A monolithic platform simplifies packaging and minimizes cost through a single-chip solution containing both CMOS electronics and optics. Monolithically integrated electro-optic chips feature tight device-to-circuit proximity; connections to optical devices leverage on-chip wires to minimize device-to-circuit interconnect parasitics and enable a very high density of interconnections to the optical devices. The approach offers near-perfect yield, in line with that of CMOS transistors. As both optics and electronics are now on the same chip, however, processing optimizations for optics and electronics cannot be performed independently of each other. As such, the tran-

sistors in monolithic platforms tend to derive from older CMOS processes [4, 8], where transistor properties are not so sensitive to processing changes for optics.

2.2.2 SOI and Bulk Platforms

To date, the vast majority of both monolithic and heterogeneous photonic platforms use thick buried-oxide (BOX) SOI processes as the substrate of choice. Here, the crystalline silicon layer on top of the BOX serves as a suitable material for the waveguide core and the thick BOX serves as the bottom cladding, providing optical mode confinement [4, 8, 43, 98] and preventing light in the waveguide from coupling into the substrate. Together, thick BOX platforms have demonstrated waveguide losses that are sub-1 dB/cm. As a comparison point, losses that are < 20 dB/cm are sufficient in enabling resonant devices of a high enough quality factor for use in DWDM applications.

We note, however, that thick BOX SOI is generally avoided in commercial high-performance CMOS SOI processes due to the higher thermal impedance. These processes use a thin BOX layer, which is not sufficiently thick to prevent the propagating optical mode in the waveguide from leaking into the silicon substrate. To enable optics, a post-processing etch using XeF_2 can be applied to remove the high-index silicon substrate [19, 54, 62, 88] from the chip, utilizing the BOX layer as a high selectivity stopper for the etchant. Afterwards, a lower-index material (such as glass or silicon carbide) can be attached to form the new substrate [54, 62] or the chip can be left exposed left to air [19, 88] (which has a low index) to enable confinement. With these steps, we have demonstrated a complete electro-optic platform in a state-of-the-art commercial CMOS SOI process, with no foundry changes needed to enable optics.

The majority of photonics integration efforts presented so far have really only focused on addressing a niche market; both heterogeneous approaches and SOI are costly and avoided by mainstream foundry customers. To gain traction among high-volume cost-driven applications, such as memory, electronic-photonic integration must be demonstrated monolithically in bulk silicon. Compared to SOI, a bulk platform faces two additional challenges. The first is the lack of a thin crystalline

silicon layer present natively on the wafer, necessitating polysilicon-based waveguides [50, 61], which can be much more lossy due to crystal grain imperfections. Alternatively, crystalline silicon may be expitaxially grown as the waveguide material [9], but the high temperature processing has thus far proved to be a risk to process-native transistors. The second challenge is the lack of a thick BOX layer for waveguide isolation from the substrate. The undercut technique [61] is suitable for this, but requires post-processing, constrains circuit-photonics placement, and weakens the chip mechanically. The extension of shallow-trench isolation to make deep oxide-filled trenches underneath optical waveguides [9, 50] is an alternative that enables tighter integration without post-processing, but requires additional process integration.

Chapter 3

Microring Resonance Tuning

The optical microring resonator is the key enabling device for integrated DWDM optical links. However, microrings are difficult devices to work with due to the sensitivity to process and temperature variations. Our past work explored high-level tuning strategies based on microring-to-wavelength assignments in DWDM links [18, 79]. In this chapter, we explore the lower-level challenges associated with keeping a microring tuned to an assigned wavelength channel. We first contribute models that describe microring behavior when both variations (thermal or process) and laser-induced self-heating come into play. Using these features, we perform a review of representative tuning techniques to classify the effectiveness of each solution against a set of tuning challenges. What we find is that no existing solution can address the full list of issues, motivating the need for a more sophisticated approach. We use the experience, insights, and considerations gained from these discussions to support the design of the generalized bit-statistical tuner in Chapter 4, where we address the shortcomings of all previous attempts.

3.1 Process Variations and Thermo-Optic Effect

As a modulator, the ring imprints digital *1*s and *0*s onto a wavelength positioned close to its resonance, performing both the channel selection and *on-off* keying of light. In the receiver, the ring acts as a channel-selective filter, picking off a specific

wavelength to drop to a broadband photodetector while letting other wavelengths pass by unaffected. In both cases, the resonance of the ring (λ_0) must be precisely aligned to the wavelength of interest; any deviation in resonance will immediately translate to a reduction in the transmit or receive margin (a closing of the link's eye), leading to bit-errors. Unfortunately, the resonant wavelength of a ring is highly sensitive to both ring geometry and temperature.

Variations in silicon layer thicknesses on SOI layers and lithographic precision (cross-chip, cross-wafer and lot-to-lot) impact the post-fabrication resonances of ring resonators. Processing variations have both a random and systematic component; ring resonators in close proximity experience less overall variations relative to each other than ring resonators far away. The authors of [71] characterized this effect in an SOI process for 193 nm lithography using identically designed ring resonators separated at different distances away from each other. Ring resonator devices only 25 µm away had a λ_0 mean standard deviation of 0.15 nm (18 GHz at 1550 nm). When this distance increased to 1700 µm, representative of an intra-chip distance, the mean standard deviation increased to 0.55 nm (69 GHz). At 10 000 µm and 20 000 µm, which are representative of inter-chip distances, the mean standard deviation increased further to 1.3 nm (162 GHz) and 1.8 nm (225 GHz), respectively. The variations of the microring filter banks we fabricate in [62] are in agreement with the reported scale of variations. In [61], identical microrings fabricated adjacent to each other in a state-of-the-art bulk CMOS process employing 193 nm immersion lithography have been measured to differ in resonance by up 90 GHz. These variations are similar in magnitude to the variations for an SOI process [71]. If uncompensated for, the magnitude of process variations will easily bring a microring's post-fabrication resonance out of its useful wavelength range.

The strong thermo-optic effect of silicon can produce large refraction index changes; an increase or decrease in temperature of 1 K produces a resonance red-shift (increase in wavelength) or blue-shift (decrease in wavelength), respectively. The resultant wavelength shift due to temperature can be obtained through the following relation-

ship [58]:

$$\frac{\partial \lambda_0}{\partial T} = \frac{\lambda_0}{n_g}\frac{\partial n_{eff}}{\partial T} \tag{3.1}$$

where T is temperature, λ_0 is the ring's resonant wavelength, n_g is the group refractive index, and n_{eff} is the effective refractive index. It can also be expressed in frequency as:

$$\frac{\partial f_0}{\partial T} = -\frac{f_0}{n_g}\frac{\partial n_{eff}}{\partial T} \tag{3.2}$$

where f_0 is the ring's resonant frequency ($f_0 = c/\lambda_0$). Note that these equations ignore the change in physical dimension from the coefficient of thermal expansion, as they are negligible compared to the shift caused by the change in index. The effective index, n_{eff}, is also dependent upon the optical mode confinement in both the silicon waveguide core and the cladding material. For typical silicon-based microrings, $\partial f/\partial T$ is on the order of $-10\,\mathrm{GHz/K}$ (or $0.05\,\mathrm{nm/K}$) [58, 61] around the wavelengths of interest.

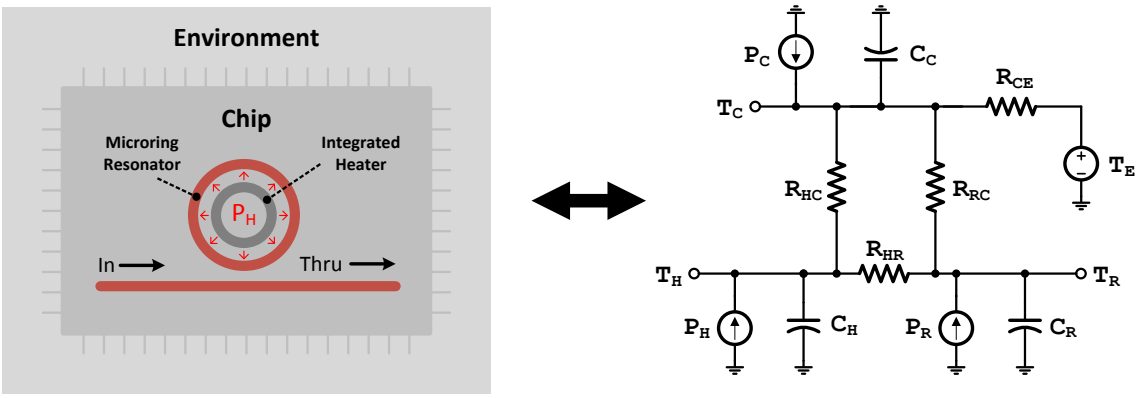

Figure 3-1: Mapping the thermal behavior of a ring resonator to an equivalent electrical circuit. Thermal capacitance maps directly to electrical capacitance and thermal resistances map to electrical resistors. Elements that generate heat are modeled as a current source and voltages represent temperature. The circuit shown contains four nodes: the embedded microring heater, the microring waveguide, the rest of the chip (and package), and the environment. Elements corresponding to those nodes are represented by subscripts H, R, C, and E, respectively.

The strong thermo-optic effects enable the most severe process variations to be compensated thermally through a temperature change of a few tens of degrees. We

can create temperatures changes in the ring by placing a resistive heating element alongside the ring. When voltage is applied across the heater, the heat created by the heater flows into the ring and changes its temperature. The temperature change then causes a shift in λ_0. This is a widely employed technique [16, 41, 58, 62, 80, 88, 99] and can tune a ring across an entire FSR, provided there is enough power delivered to the heater. We can exploit the thermal-electrical circuit duality [13, 14, 24] to build an equivalent thermal circuit model of a heater-embedded microring resonator system, shown in Figure 3-1. This model captures temperature activity at four temperature nodes: the heater, the ring waveguide, the chip, and the environment. For ease of understanding, the equivalent circuit in Figure 3-1 uses lumped circuit elements as opposed to distributed elements. The use of distributed elements can be performed using the thermal simulation framework we describe in a prior work [13, 14].

For the analyses we perform in the rest of this thesis, we can make several simplifications to this model. First, as microring heaters are typically embedded within the ring waveguide itself [80, 88] or placed in close proximity, the heater-to-ring thermal impedance, R_{HR}, is very small compared to the heater-to-chip and ring-to-chip impedances, R_{HC} and R_{RC}. As such, we can group the heater and ring waveguide nodes into one combined node. Second, due to the large thermal capacitance of a chip (from the sheer amount of matter that is present), heating an individual ring has a negligible impact on the temperature of the chip node. In addition, from prior analysis performed in [13], the temperature of the chip changes very slowly with power changes on the chip; chip temperature fluctuations happen on the order of several milliseconds to seconds [13], which is more than an order of magnitude slower than the temperature changes to the ring from excitations occurring from the embedded ring heater or from the ring waveguide, which are on the order of tens of microseconds. The faster transients dominate the response and the ring sees the chip node as if it was tied to a temperature (voltage) source that changes very slowly with time. Finally, if we consider only changes in temperature changes as opposed an absolute temperature, any fixed temperature (voltage) sources in our circuit model become shorts to a temperature "ground".

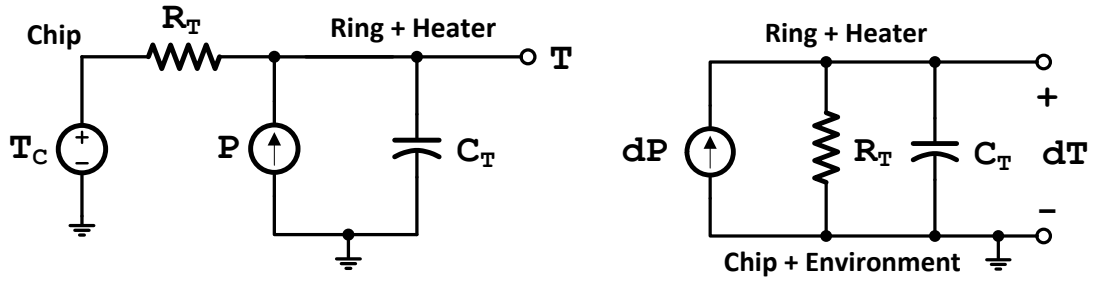

(a) Simplified circuit (b) Temperature change due to heaters

Figure 3-2: Simplified analytical model of ring resonator temperature variations (a) and dT, when subjected to a change in heater power, dP, when chip temperature is constant (b). R_T is the thermal impedance between the ring resonator to the chip, and C_T is the thermal capacitance (heat capacity times density times volume). $R_T = R_{HC}||R_{RC}$ and $C_T = C_H + C_R$.

With these simplifications, the circuit reduces down to the circuit in Figure 3-2. By superposition, we can relate dP, a change in heater output power, and to a corresponding change in steady-state temperature, dT, through the circuit in Figure 3-2b. The relationship is given by:

$$dT = R_T \cdot dP \tag{3.3}$$

where R_T is the thermal impedance from the microring resonator to the rest of the chip. We define the *tuning efficiency*, H_T, of a ring to be:

$$H_T = \frac{\partial \lambda_0}{\partial P} = \frac{\partial \lambda_0}{\partial T} \cdot \frac{dT}{dP} = \frac{\partial \lambda_0}{\partial T} \cdot R_T \tag{3.4}$$

which has units of nm/mW. We note that the choice of the chip substrate has a huge impact on the thermal resistance between the microring and the rest of the chip (R_T) which, consequently, impacts the tuning efficiency (H_T) as well. As crystalline silicon is a very thermally conductive material, much of the heat transfer from a ring resonator to the rest of the chip is through a chip's bulk substrate. As such, an SOI or bulk process chip with the substrate removed and/or replaced with a thermally insulating material will have order-of-magnitude higher R_T and H_T [13].

This is known as an undercut technique and has been exploited widely to enhance the efficiency of thermal tuners [16, 61, 62, 88].

If we take into account the fact that it takes time for the ring to heat up and cool down, the self-heating response will be frequency dependent. We can derive the following expression for $H'_T(\omega)$, which is the change in λ_0 with respect to input power, taking into account a first-order temperature response caused by the heat capacitance and thermal impedance of the microring to the substrate:

$$H'_T(\omega) = H_T \cdot \frac{1}{1 + j\omega\tau_T} \tag{3.5}$$

where τ_T is the thermal time constant. In our simplified circuit model, $\tau_T = R_T \cdot C_T$. By using the superposition property of the power (current) and temperature (voltage) sources in Figure 3-2a, τ_T is the same between the case where the heater is responsible for the change in ring temperature and the case where a change in chip temperature is responsible. However, as we stated earlier, changes in the chip temperature occur at far slower timescales than τ_T.

The strong temperature dependence is a mixed blessing, however. In a hostile thermal environment, such as that of a modern microprocessor chip, the ring could be subject to very large environmental temperature swings. As shown in Figure 3-3, even a 1 K temperature change results in a significant degradation in modulation depth for a modulator ring or, similarly, a significant loss of power in a receive filter ring. From these observations, it is clear that a mechanism for stabilizing ring temperatures is essential to microring operation.

3.2 Ring Self-Heating

In our discussions so far, we have only considered ring resonator temperature changes from ambient sources. We now discuss the implications of microring self-heating, which includes the heat generated from the ring waveguide itself. These

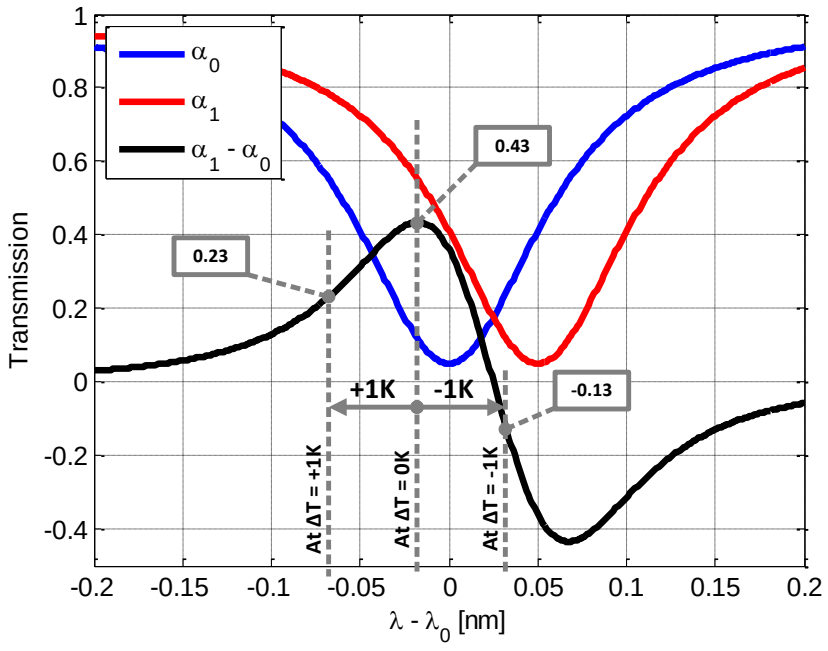

Figure 3-3: The effect of a 1 K change in temperature on a microring resonator with a quality factor of 10000 at a wavelength of 1280 nm. λ is the laser wavelength and λ_0 is the ring's resonance on optical *0*s. α_0 and α_1 denote the thru-port transmission of the *0* and *1* states of the modulator ring. $\alpha_1 - \alpha_0$ represents the achieved modulation depth. The modulator ring is nominally biased at the point of maximum modulation depth. When temperature increases, the resonance red-shifts, such that $\lambda - \lambda_0$ is smaller. When temperature decreases, the resonance blue-shifts, such that $\lambda - \lambda_0$ is larger.

effects result in phenomena, such as optical bistability of ring resonators [65, 95, 99], that are essential to consider in a thermal tuning solution.

3.2.1 Change in Resonance From Self-Heating

There are two primary sources of self-heating power: laser power (primary from free carrier absorption) and $I \cdot V$ power dissipated by electrical junctions in the ring (such as in carrier-injection modulators or from photocurrents in a resonated photodetector). We define a value, P_S, for the self-heating power that includes these terms:

$$P_S = P_J + P'_L \tag{3.6}$$

where P_J is the power dissipation from the electrical junctions and P'_L is the laser power that self-heats the ring. P'_L is related to the input port laser power, P_L, through the following relation:

$$P'_L = \kappa \cdot P_L \cdot \gamma(\lambda) \tag{3.7}$$

where $\gamma(\lambda)$ is the fraction of input power "lost" by the ring as a function of the laser wavelength and κ is the fraction of the "lost" power that self-heats the ring. From [99], typical values of κ can vary from 0.3–0.9 , dependent upon the proportion of laser power absorbed through free carrier absorption in the cavity or radiated out from the cavity. A ring with a strongly coupled drop-port, for example, would have a low value of κ as most of the power lost by the ring goes to the drop port. The amount of laser power lost by the resonator depends on how close the laser wavelength is to the resonance, λ_0, and is captured by γ. We note that in steady-state, γ is related to α, the through-port transmission, through the following relationship:

$$\gamma(\lambda) + \alpha(\lambda) = 1 \tag{3.8}$$

which states that whatever proportion of laser power not transmitted to the through port is "lost" by the ring. Using Equation 2.7, we then rewrite P_S as:

$$P_S = P_J + \kappa \cdot P_L \cdot \gamma(\lambda) \tag{3.9a}$$

$$= P_J + \kappa \cdot P_L \cdot (1 - \alpha(\lambda)) \tag{3.9b}$$

$$= P_J + \kappa \cdot P_L \cdot \frac{A}{1 + 4\left(\frac{\lambda - \lambda_0}{\Delta\lambda}\right)^2} \tag{3.9c}$$

P_S causes a change in λ_0 through the tuning efficiency, H_T. From this, we can finally derive an expression for λ_0, taking into account the self-heating effects:

$$\lambda_0 = \lambda'_0 + H_T \cdot P_S \tag{3.10a}$$

$$= \lambda'_0 + H_T \cdot \left[P_J + \kappa \cdot P_L \cdot \frac{A}{1 + 4\left(\frac{\lambda - \lambda_0}{\Delta\lambda}\right)^2} \right] \tag{3.10b}$$

where λ_0' is the resonance of the ring without self-heating considered. This equation can be directly solved for λ_0 and we plots the solutions for λ_0 as a function of the laser wavelength, λ, in Figure 3-4. A higher laser power causes the shape of the curves to "stretch" upward and right, extending the maximum change in λ_0 as the laser is able to heat the ring up more. Microrings with lower quality factors will see the self-heating effect across a broader range of wavelengths, as it is able to capture laser light when the laser is farther away from λ_0. P_J, when independent of the laser power, provides an additional shift in λ_0. As the change in λ_0 also affects P_L', a fixed P_J does not equate a fixed change in λ_0.

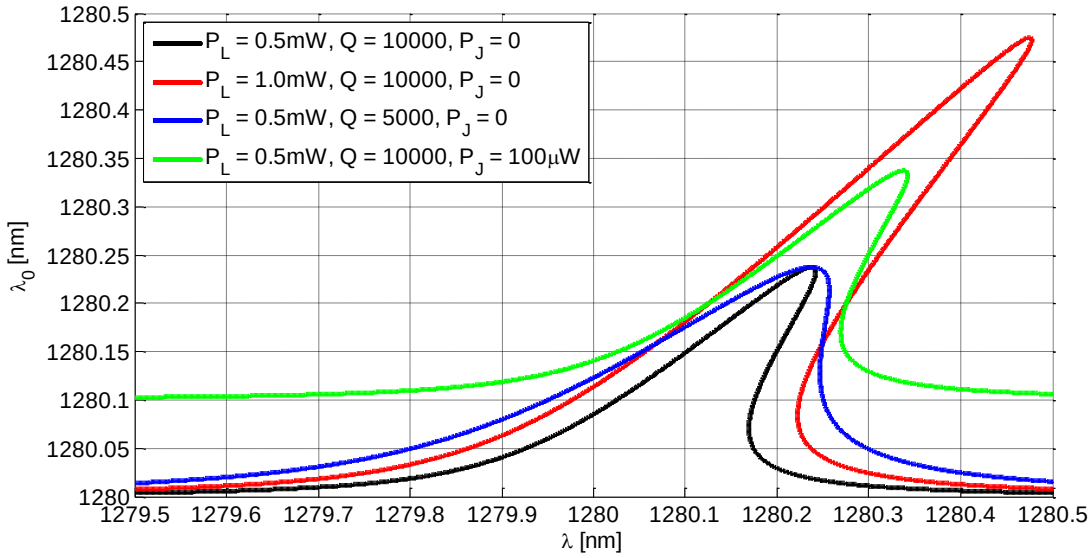

Figure 3-4: Calculated λ_0 vs. the laser wavelength (λ). The plotted lines shows all solutions of Equation 3.10, each with different ring parameters. The ring self-heats from the resonating laser power to produce a changing λ_0. The ring has $\lambda_0' = 1280\,\text{nm}$, $H_T = 1\,\text{nm/ mW}$, $P_J = 0$, and $\kappa = 0.5$.

3.2.2 The Feedback Relation of Self-Heating

The solutions of Equation 3.10 take on many interesting forms and carry significant implications. In particular, the regions where there are three different solutions for a given λ imply multiple stability points. To begin discussion, we consider the case of an optical ring resonator in thermal steady-state that is perturbed by a small

change in power, dP_E, generated from an arbitrary source. dP_E has an impact on the resonator's resonant wavelength through the tuning efficiency H_T:

$$d\lambda_0 = H_T \cdot dP_E \tag{3.11}$$

As λ_0 moves, $\alpha(\lambda)$ and hence $\gamma(\lambda)$ changes as well, producing a change in the self-heating power of the ring. If we assume, for now, that P_J is independent of λ_0, we obtain:

$$dP_S = \frac{\partial P_S}{\partial \lambda_0} \cdot d\lambda_0 \tag{3.12a}$$

$$= \kappa \cdot P_L \cdot \frac{\partial \gamma(\lambda)}{\partial \lambda_0} \cdot d\lambda_0 \tag{3.12b}$$

However, the change in P_S will also affect λ_0 through the tuning efficiency:

$$d\lambda_0 = H_T \cdot dP_S \tag{3.13}$$

which will again affect P_S. Hence, an intrinsic feedback relationship is present in ring resonator self-heating.

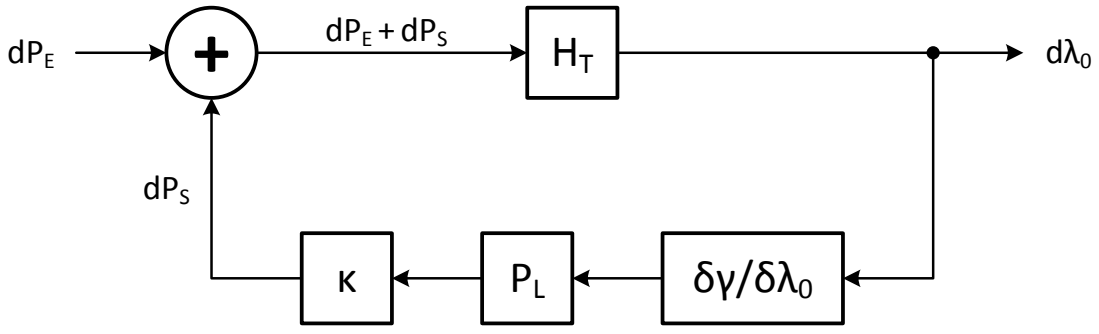

Figure 3-5: The feedback relationship of self-heating. For any power perturbation to the ring, dP_E, such as from the chip heating up or from the ring's heaters, the feedback path will either amplify or attenuate the change in λ_0, depending on the sign of $\partial\gamma/\partial\lambda_0$.

We can model and analyze the effects of self-heating as the feedback system shown in Figure 3-5. The loop gain of the system is given by:

$$G_L = \kappa \cdot P_L \cdot H_T \cdot \frac{\partial \gamma}{\partial \lambda_0} \tag{3.14}$$

Using the loop gain, we can derive an overall transfer function between a change in power and the corresponding change in resonant wavelength:

$$\frac{d\lambda_0}{dP_E} = \frac{H_T}{1 - G_L} \tag{3.15a}$$

$$= \frac{H_T}{1 - \left(\kappa \cdot P_L \cdot H_T \cdot \frac{\partial \gamma}{\partial \lambda_0}\right)} \tag{3.15b}$$

To understand the feedback relationship, we will examine the loop gain, G_L in Equation 3.14. We observe that the system has positive feedback when $G_L > 0$ and negative feedback when $G_L < 0$. In the expression for G_L, κ, P_L, and H_T are positive physical terms that cannot be negative. As such the $\partial\gamma/\partial\lambda_0$ term alone determines whether self-heating contributes positive or negative feedback.

We plot $\partial\gamma/\partial\lambda_0$ for a ring resonator as a function of λ in Figure 3-6. Observe that when the ring is *red-biased* ($\lambda < \lambda_0$), $\partial\gamma/\partial\lambda_0$ is negative, implying negative feedback. When the ring is *blue-biased* ($\lambda > \lambda_0$), $\partial\gamma/\partial\lambda_0$ is positive, implying positive feedback. The crossover wavelength at which $\partial\gamma/\partial\lambda_0$ switches signs is λ_0. The magnitude of each of the terms of G_L determine the strength of the feedback. Intuitively, a higher laser power, a higher the fraction of laser power that is converted to heat, or a higher tuning efficiency all translate to a stronger self-heating feedback effect. For $\partial\gamma/\partial\lambda_0$, the magnitude is greater for rings with a higher quality factor, as that sharpens the wavelength notch. As such, rings with a higher quality factor experience stronger positive and negative feedback.

3.2.3 Self-Heating Optical Bistability

From the loop stability criteria of $G_L < 1$, we note that the transfer function can become thermally unstable when *blue-biased*. When unstable, any perturbations to

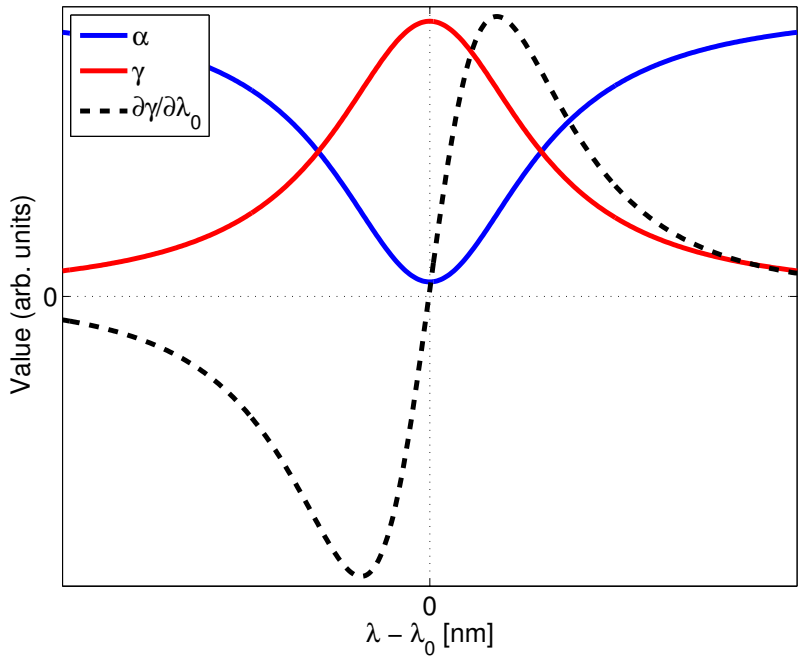

Figure 3-6: A plot of $\partial\gamma/\partial\lambda_0$ across λ. $\partial\gamma/\partial\lambda_0$ is negative for $\lambda < \lambda_0$ (*red-biased*) and positive for $\lambda > \lambda_0$ (*blue-biased*).

the ring are continuously amplified and the ring resonance immediately drifts away from the original point. As λ_0 moves, $\partial\gamma/\partial\lambda_0$ changes as well and ring resonance continues to drift until it reaches back to a point where a lower value of $\partial\gamma/\partial\lambda_0$ again enables the stability criteria to be satisfied. This leads to thermo-optical bistability; given constant P_L, a ring biased in the unstable regime will drift away from the metastable point and settle to one of two thermally-stable steady states.

We show simulated self-heating stability curves for a microring resonator in Figure 3-7. In the case with 1 mW of laser power, the laser is bistable over a large range of wavelengths. When the ring is *red-biased* (region above or to the left of the $\lambda = \lambda_0$ line), the ring is always in a stable regime. When the ring is *blue-biased*, the ring can enter a metastable regime, where any deviation from the metastable point grows, amplified by the positive feedback of the self-heating behavior, before settling to one of two stable states where $A < 1$. The range of wavelengths with bistability behavior shrinks with smaller laser power. Bistability is not present below 0.2 mW of laser power. However, the self-heating still causes noticeable asymmetry in the response.

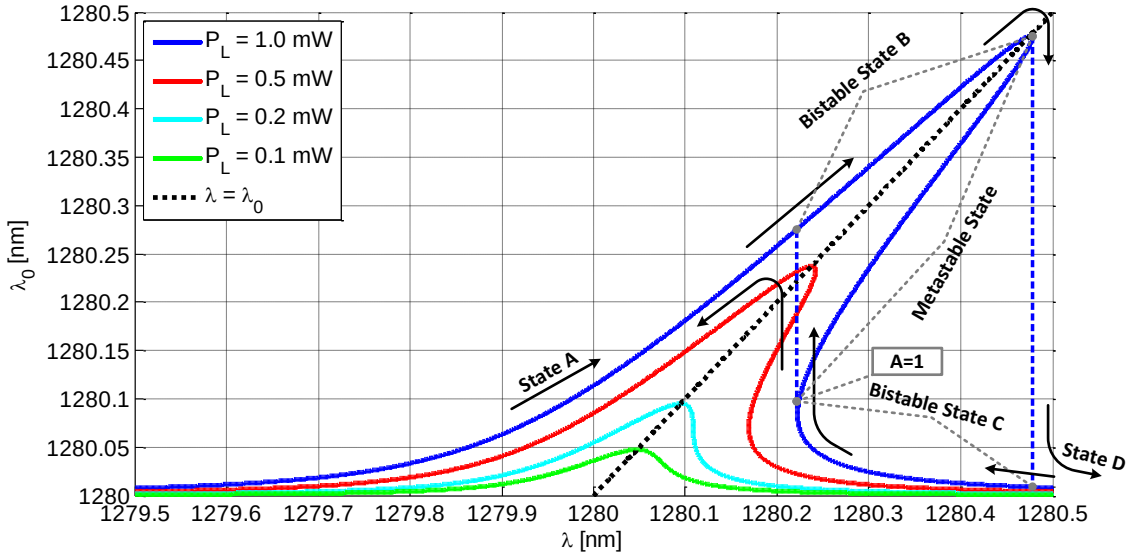

Figure 3-7: Microring self-heating stability curves for a ring in steady state. A plot of λ_0 vs. the laser wavelength (λ). Each plotted line shows all the solutions of Equation 3.10 at a different laser power (P_L). The line where $\lambda = \lambda$ is also shown. The ring self-heats from the resonating laser power to produce a changing λ_0. The ring has a nominal $\lambda_0 = 1280\,\text{nm}$, $Q = 10000$, $H_T = 1\,\text{nm/mW}$, and $\kappa = 0.5$. We annotate the region of bistability for the case of 1 mW of laser power, showing the range of wavelengths with two stable states and the metastable state. The directions of the arrows show the trajectory λ_0 takes if we sweep the laser from low λ to high λ and from high λ to low λ.

When the laser wavelength, λ, is in the region where the system is bistable, the system can settle into either state B or state C. To reach the bistable state B, the system must first be *red-biased* into state A. Then, either λ increases or λ_0' decreases to move the microring into state B. Similarly, to reach bistable state C, the ring must first be *blue-biased* into state D followed by λ decreasing or λ_0' increasing to move into state C.

3.2.4 Non-Steady-State

So far, we have limited our analysis to the steady-state case. $H'_T(\omega)$ can be directly substituted into the loop gain expression of Equation 3.14, giving:

$$G_L(\omega) = \kappa \cdot P_L \cdot H'_T(\omega) \cdot \frac{\partial \gamma}{\partial \lambda_0} \tag{3.16a}$$

$$= \kappa \cdot P_L \cdot H_T \cdot \frac{\partial \gamma}{\partial \lambda_0} \cdot \frac{1}{1 + j\omega\tau_T} \tag{3.16b}$$

$$= G_L(0) \cdot \frac{1}{1 + j\omega\tau_T} \tag{3.16c}$$

Intuitively, the loop-gain rolls off at higher frequencies as the temperature change can no longer keep up with the perturbation. Note that the temporal behavior of other factors of the equation are much faster than the timescales of interest. Hence, they are assumed to be instantaneous.

3.3 Thermal Effects on Active Elements

The feedback system of Figure 3-5 and Equation 3.15 can be generalized for different sources of perturbations, which we can use to use to model the impact of active elements. In all cases, the loop gain, G_L, of the self-heating effect works to scale the effect of the perturbation on the resonance. We define $S(\omega) = 1/(1 - G_L(\omega))$. Figure 3-8 shows the steady-state behavior of $S(\omega)$ when laser wavelength changes relative to λ_0. With the given parameters in steady state, 1 mW of laser power will attenuate the perturbation by up to 8 dB when *red-biased*, while amplifying the perturbation to the point of instability when *blue-biased*.

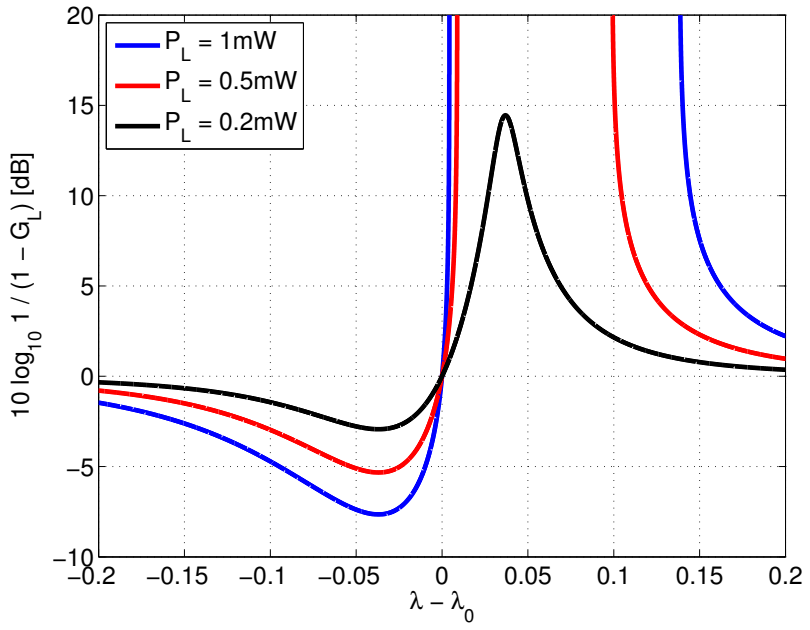

Figure 3-8: Self-heating's scaling effect on perturbations, plotted at different laser powers for a ring with $\Delta\lambda = 128\,\text{pm}$. $H_T = 1\,\text{nm/ mW}$ and $\kappa = 0.5$. We plot $1/(1 - G_L)$ on a dB scale.

The frequency domain response of $S(\omega)$ can be written as:

$$S(\omega) = \frac{1}{1 - G_L(\omega)} \tag{3.17a}$$

$$= \frac{1}{1 - G_L(0) \cdot \left(\frac{1}{1+j\omega\tau_T}\right)} \tag{3.17b}$$

$$= \frac{1 + j\omega\tau_T}{1 + j\omega\tau_T - G_L(0)} \tag{3.17c}$$

$$= \frac{1}{1 - G_L(0)} \cdot \frac{1 + j\omega\tau_T}{1 + j\omega \cdot \frac{\tau_T}{1 - G_L(0)}} \tag{3.17d}$$

The system has a zero at $1/\tau_T$ and a pole at $(1 - G_L(0))/\tau_T$. When G_L is negative (*red-biased*), the zero frequency is lower than the pole frequency. If G_L is positive, the zero frequency is higher than the pole frequency. We show a plot of $|S(\omega)|$ for the case of both positive and negative feedback $G_L(0)$ in Figure 3-9. The rest of this section will develop models for analyzing the impact of $S(\omega)$ on active microring structures. To keep the analysis intuitive, we will model active effects as perturbations

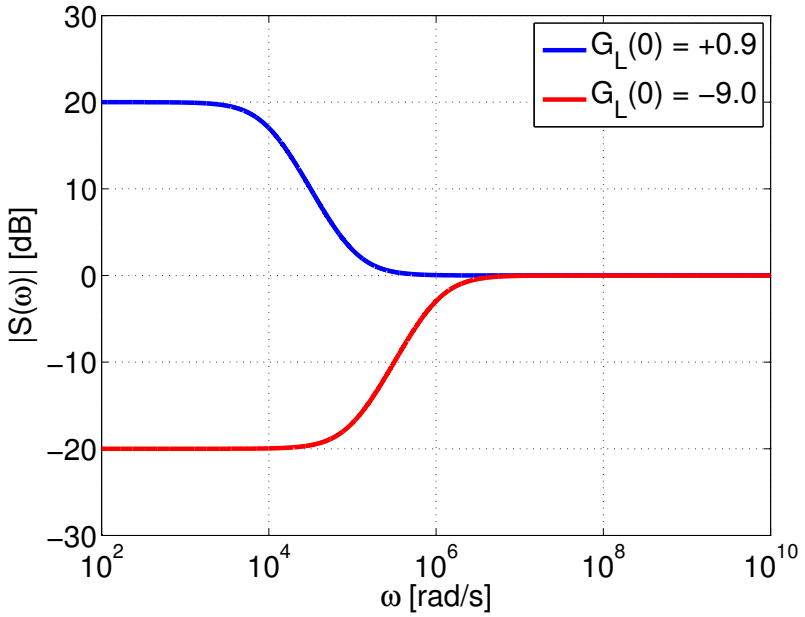

Figure 3-9: Plot of $|S(\omega)|$ for different values of $G_L(0)$, showing the effects of positive and negative feedback relationships. We use a $\tau_T = 10\,\mu s$.

to a microring system. As active effects are typically large, we essentially make the approximation to linearize non-linear quantities, such as $G_L(0)$, around a region of interest.

3.3.1 Impact on Active Thermal Tuning

Consider the case of an embedded ring heater, which we use to control the λ_0 of the ring (to compensate for process variations, for example). The heater outputs a small change in power, dP_E in order to achieve this. If self-heating is not considered, the steady-state change in the resonant wavelength, $d\lambda_0$, is just $d\lambda_0 = H_T \cdot dP_E$. When self-heating is considered, the same dP_E results in a smaller $d\lambda_0$ shift for a red-biased ring and a larger $d\lambda_0$ shift for a blue-biased ring. This effect stems from the the attenuation or amplification caused by the factor of $1/(1 - G_L)$ in steady-state. As such, a *red-biased* ring will experience finer-grained λ_0 control from the heater and a *blue-biased* ring will experience coarser-grained λ_0 control from the heater.

With time constants considered, the transfer function for heating takes the form of:

$$\frac{d\lambda_0}{dP} = H_T'(\omega) \cdot S(\omega) \tag{3.18a}$$

$$= \frac{H_T}{1 - G_L(0)} \cdot \frac{1}{1 + j\omega \frac{\tau_T}{1 - G_L(0)}} \tag{3.18b}$$

The scaling factor from self heating also moves the pole in the system to a frequency given by $\tau_T/(1 - G_L(0))$. Figure 3-10 shows a typical heater tuning curve. The bistability creates hysteresis in λ_0 when responding to the heater power.

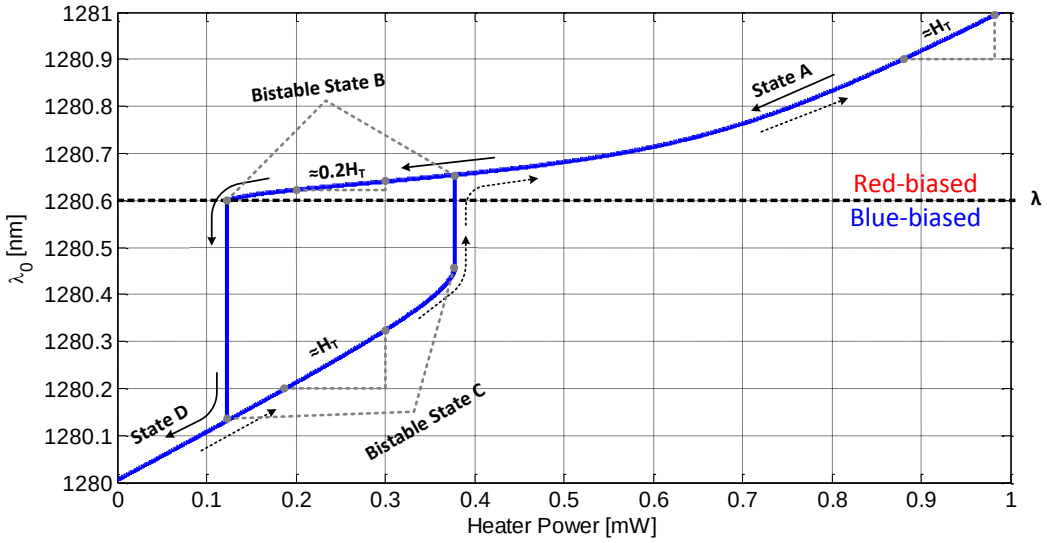

Figure 3-10: Simulated heater tuning curve showing hysteresis behavior, solved from Equation 3.10 for each different heater power (which changes λ_0'). The dashed arrow lines indicate the path taken if the heater is swept from low power to high power, the solid arrow lines indicate the path if swept from high power to low power. When the heater is off, the ring's $\lambda_0 = 1280\,\text{nm}$ and the laser is parked at $\lambda = 1280.6\,\text{nm}$ (the ring is *blue-biased* when heater power is 0). Depending on the output heater power and the bistable state, the ring may be either *red-biased* ($\lambda_0 > 1280.6\,\text{nm}$) or *blue-biased* ($\lambda_0 < 1280.6\,\text{nm}$). $P_L = 1\,\text{mW}$, $H_T = 1\,\text{nm/mW}$ and $\kappa = 0.5$. The slope of the curve is labeled at a few representative locations.

If the heater power is determined by the output of a digital-to-analog converter (DAC), we obtain either an enhancement of the DAC resolution for a *red-biased* ring or a degradation of the resolution of the DAC for a *blue-biased* ring. Even at low

amounts of laser, where the ring is stable, the DAC driving the embedded heater needs to be extremely precise if the ring is to be operated *blue-biased*. Whereas in the case of the *red-biased* ring, we can use a very coarse DAC and gain a few extra bits of precision for free (conveniently located around where the resonance is, where we would want the extra precision) from the self-heating. The frequency response also indicates that λ_0 takes a shorter time to settle when responding to the heater when *red-biased* and a longer time to settle when *blue-biased* due to the frequency change in the pole.

3.3.2 Impact on Modulators

We can model the modulation as a small-signal perturbation in λ_0 with an output of $d\alpha$, the modulation depth (the change in the thru-port transmission), shown in Figure 3-11. The overall transfer function is:

$$\frac{d\alpha}{d\lambda_0'} = \frac{d\alpha}{d\lambda_0} \cdot S(\omega) \tag{3.19}$$

with $S(\omega)$ given by Equation 3.17. To gain an understanding of the behavior, let us first consider the simple case of single-tone modulation. The $S(\omega)$ term attenuates the modulation depth for *red-biased* rings ($G_L(0) < 0$) and enhances the modulation depth for *blue-biased* rings ($G_L(0) > 0$). The effect comes into play only at lower frequencies, while leaving higher frequencies unaffected. The corner frequency for which a lower frequency produces an attenuation or enhancement depends on the relative position of the pole and zero in $S(\omega)$. This occurs at $(1 - G_L)(0)/(\tau_T)$ and $1/(\tau_T)$ for *red-biased* and *blue-biased* rings, respectively.

When modulating bitstreams, the self-heating effect manifests itself as a noise in the ring resonance similar to baseline wander noise in AC-coupled electrical systems. To continue analysis, we make an approximation to linearize $d\alpha/d\lambda_0$ around our region of interest (or alternatively, assume a small modulation signal). We can quantify the noise variance as:

$$\sigma^2 = \frac{1}{2\pi} \int_{-\infty}^{\infty} PSD_{TX}(\omega) \cdot |H_M(\omega)|^2 \, d\omega \tag{3.20}$$

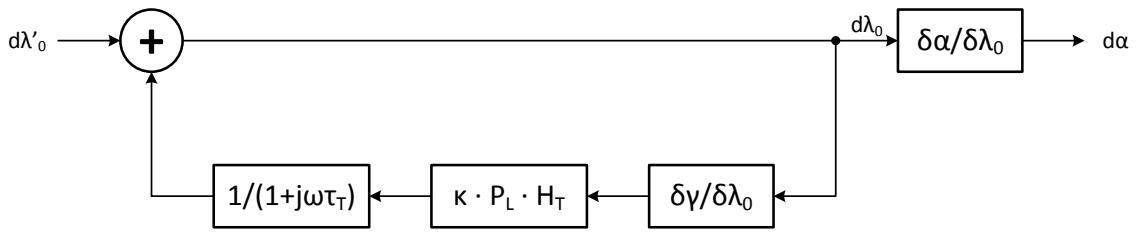

Figure 3-11: Modulation system, with the feedback due to self-heating shown.

where H_M is the portion of the transfer response of the power that contributes to noise caused by the drift in resonance. H_M consists of the signal power that is lost due to the self-heating response, given by:

$$H_M = \frac{d\alpha}{d\lambda_0} - \frac{d\alpha}{d\lambda_0'} \tag{3.21a}$$

$$= \frac{d\alpha}{d\lambda_0} \cdot (1 - S(\omega)) \tag{3.21b}$$

$$= -\frac{d\alpha}{d\lambda_0} \cdot \frac{G_L(0)}{1 - G_L(0)} \cdot \frac{1}{1 + i\omega\tau_T'} \tag{3.21c}$$

The $|H_M(\omega)|^2$ quantity is evaluated to be:

$$|H_M|^2 = \left(\frac{d\alpha}{d\lambda_0}\right)^2 \cdot \left(\frac{G_L(0)}{1 - G_L(0)}\right)^2 \cdot \frac{1}{1 + \omega^2\tau_T'^2} \tag{3.22}$$

where $\tau_T' = \tau_T/(1-G_L(0))$. If we assume a random modulated bit sequence where the change in λ_0' between *1*s and *0*s is of a magnitude $\Delta\lambda_0'$, we can obtain an expression for PSD_{TX}:

$$PSD_{TX}(\omega) = \Delta\lambda_0'^2 \cdot t_{bit} \cdot \mathrm{sinc}^2\left(\frac{\omega \cdot t_{bit}}{2}\right) \tag{3.23}$$

The noise variance can then be written as:

$$\sigma^2 = \left[\frac{\Delta\alpha \cdot G_L(0)}{1 - G_L(0)}\right]^2 \frac{t_{bit}}{2\pi} \int_{-\infty}^{\infty} \mathrm{sinc}^2\left(\frac{\omega \cdot t_{bit}}{2}\right) \cdot \frac{1}{1 + \omega^2\tau_T'^2} \cdot d\omega \tag{3.24}$$

with $\Delta\alpha$ given by:

$$\Delta\alpha = \Delta\lambda_0' \cdot \left(\frac{d\alpha}{d\lambda_0}\right) \tag{3.25}$$

Note that $\Delta\alpha$ is equivalent to the change in modulator transmission if there were no self-heating effects at all ($S(\omega) = 1$). Since the noise variance, σ, scales proportionally with $\Delta\alpha$, we can express the noise as a fraction of the modulation depth with no self-heating effects through the quanity $\sigma/\Delta\alpha$:

$$\frac{\sigma}{\Delta\alpha} = \left|\frac{G_L(0)}{1-G_L(0)}\right| \sqrt{\frac{t_{bit}}{2\pi}\int_{-\infty}^{\infty} \mathrm{sinc}^2\left(\frac{\omega\cdot t_{bit}}{2}\right)\cdot\frac{1}{1+\omega^2\tau_T'^2}\cdot d\omega} \tag{3.26}$$

which corresponds to the fractional eye-closure metric. Note that typically, $t_{bit} \ll \tau_T$ and $t_{bit} \ll \tau_T'$. As such, $\mathrm{sinc}^2(\omega \cdot t_{bit})$ term is ≈ 1 for frequencies smaller than the τ_T' custoff. We can therefore simplify Equation 3.26 to:

$$\frac{\sigma}{\Delta\alpha} = \left|\frac{G_L(0)}{1-G_L(0)}\right| \sqrt{\frac{t_{bit}}{2\pi}\int_{-\infty}^{\infty} \frac{1}{1+\omega^2\tau_T'^2}\cdot d\omega} \tag{3.27a}$$

$$= \left|\frac{G_L(0)}{1-G_L(0)}\right| \sqrt{\frac{t_{bit}}{2\tau_T'}} \tag{3.27b}$$

$$= \frac{|G_L(0)|}{\sqrt{1-G_L(0)}}\sqrt{\frac{t_{bit}}{2\tau_T}} \tag{3.27c}$$

Figure 3-12 shows plots of $\sigma/\Delta\alpha$ versus $G_L(0)$ for different ratios of t_{bit}/τ_T. A larger $|G_L(0)|$ increases the eye closure; when red-biased, a more negative $G_L(0)$ moves τ_T' to affect higher frequencies. When blue-biased, a more positive $G_L(0)$ creates larger positive feedback and drastically amplifies the wander noise. Using representative parameters corresponding to a red-biased ring – $t_{bit} = 100\,\mathrm{ps}$, $G_L = -5$, and $\tau_T = 10\,\mu\mathrm{s}$ – the calculated $\sigma/\Delta\alpha = 0.0046$, which means a $0.46\,\%$ eye-closure. Consequently, a 10σ eye-closure (which corresponds to a rate of 10^{-15}) would be $4.6\,\%$. While this can be acceptable, any non-random sequence with unbalanced run lengths will suffer much larger eye-closure. As a result, modulator rings require techniques for self-heating mitigation, unless the transmitted data sequences are specifically encoded to avoid unbalanced run lengths.

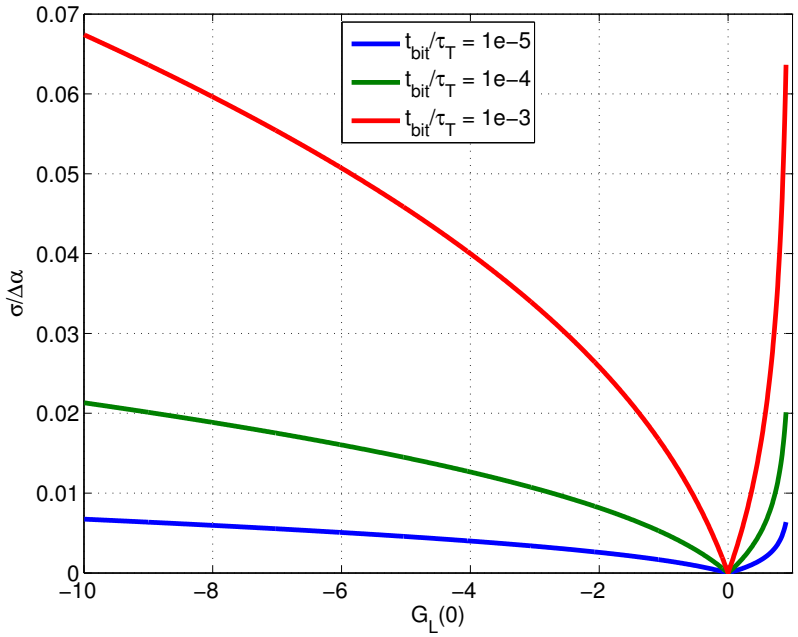

Figure 3-12: Modulator eye-closure, $\sigma/\Delta\alpha$, due to self-heating for different self-heating loop gains, $G_L(0)$, and ratios of the bit-time to the thermal time constant, t_{bit}/τ_T.

3.3.3 Impact on Filter and Receiver Rings

We can model the input on-off keyed data that a receiving ring sees as a perturbation in laser power, dP'_L. We define the output of a receiver ring, dP_L as the change in the amount of power that is "lost" by the ring, whether due to cavity absorption (such as from a resonant photodetector) or dropped by the ring. The model of the receiver is shown in Figure 3-13 and the linearized transfer function of the system can be evaluated as:

$$\frac{dP_L}{dP'_L} = \alpha S(\omega) \tag{3.28}$$

Note that this expression takes up a similar form as that of the modulator and the rest of the analysis flows identically. For a receiver sending random NRZ data, we arrive at the following equation for eye-closure metric due to self-heating

$$\frac{\sigma}{dP'_L} = \frac{|G_L(0)|}{\sqrt{1-G_L(0)}}\sqrt{\frac{t_{bit}}{2\tau_T}} \tag{3.29a}$$

which is identical to 3.27

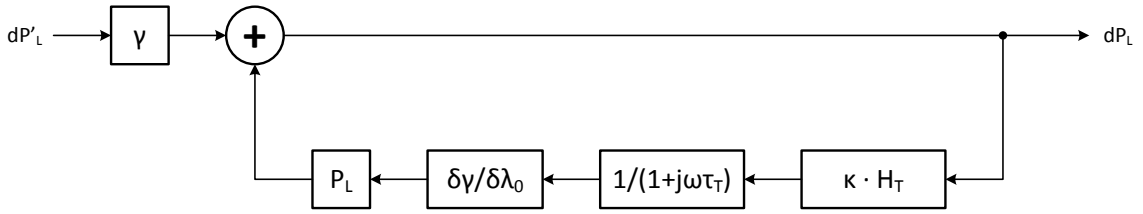

Figure 3-13: Modulation system, with the feedback due to self-heating shown.

We make one final observation that typically receiver rings will sit 10–20 dB downstream on an optical path compared to a modulator. As a result, the amount of P_L seen by a receiver ring is smaller and self-heating effects are far less pronounced.

3.4 Active Tuning

To combat thermal and process variations, a number of different solutions have been proposed. In this section, we review the different thermal tuning techniques.

3.4.1 Open-Loop vs. Closed Loop Approaches

In order to compensate for process or thermal variations, both open-loop and closed-loop approaches have been proposed. Open loop solutions typically rely upon microring athermalization, where we eliminate or reduce the sensitivity of the ring's resonance to temperature. Microring athermalization is achieved primarily through one of two ways. The first is the introduction of cladding materials with a negative thermo-optic coefficient to cancel the positive thermo-optic coefficient of silicon [53, 67]. A careful balance can bring $\partial n_{eff}/\partial T$ to zero which effectively renders the ring insensitive to temperature. However, the athermal behavior is difficult to maintain over a wide wavelength or temperature range due to the higher order temperature and wavelength dependencies in n_{eff}. The second method of athermalization is to couple the resonator to an interferometer such that the optical path length change from temperature introduced by the ring and the arm of the interferometer

cancel out [21]. Compared to the first method, the introduction of an interferometer consumes much more area but does not require any new materials. In both cases, however, athermalized solutions require perfect matching; any material or dimensional variations result in temperature-dependent behavior. Moreover, we lose our ability to thermally compensate for random process variations, as temperature is no longer a knob we can use to tune resonances. Though we can always trim athermal rings at manufacturing time to compensate for process variations, this has to be performed on a ring-by-ring basis. As a result, this is not a scalable solution for mass-produced systems with thousands of rings, limiting the commercial viability of athermal rings. Alternatively, open-loop solutions that do not athermalize rings enable process variations to be tuned away but will suffer from temperature variations.

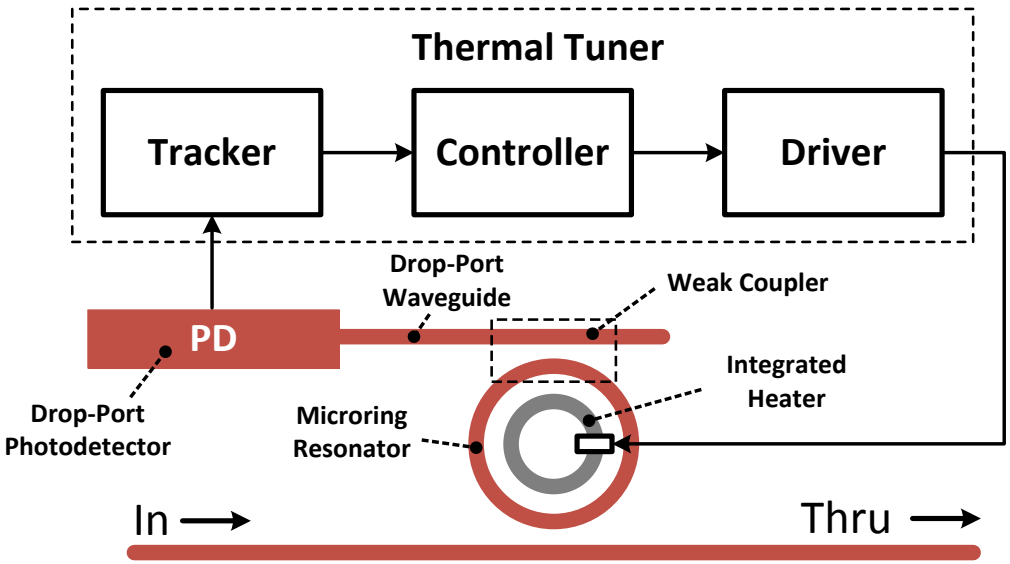

Figure 3-14: Architecture of a wavelength locking thermal tuner, consisting of a drop-port photodetector to monitor photocurrent, a photocurrent tracker circuit, a controller that makes heating decisions, and a driver for the integrated ring heater.

Closed-loop solutions build a feedback control loop around the ring to *wavelength-lock* the resonance of the ring, λ_0, to that of the laser, λ. Closed-loop solutions treat both process and temperature variations the same way and can compensate for both. Figure 3-14 shows an architecture of closed-loop wavelength locking thermal

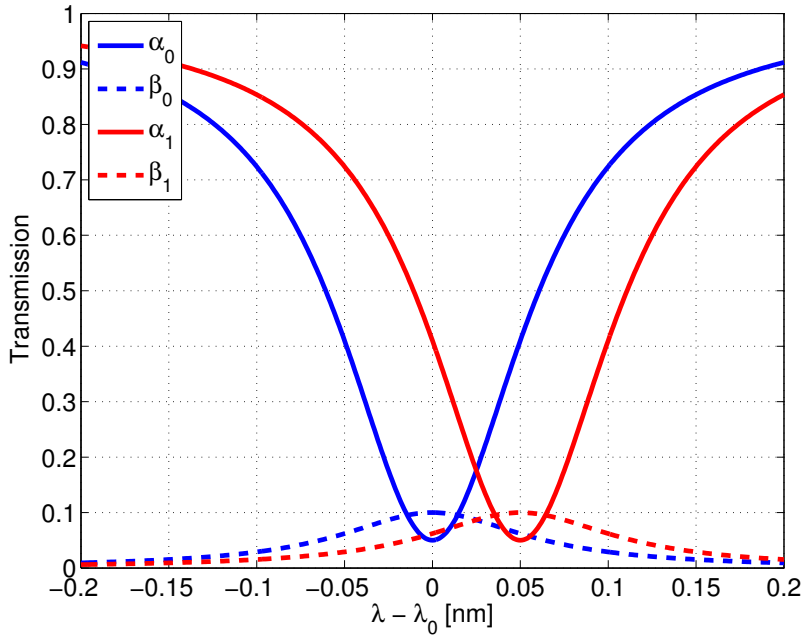

Figure 3-15: Example drop-port and through-port optical transmission characteristics of a microring modulator plotted versus $\lambda - \lambda_0$, the difference in wavelength between the microring's resonance in the *0* state (λ_0) and the laser (λ). α_0 and α_1 denote the thru-port transmission for the *0* and *1* states of the modulator ring, respectively. β_0 and β_1 denote the drop-port transmission for the *0* and *1* states of the modulator ring, respectively. The drop port waveguide is weakly coupled to the resonator and receives 1/10 of the input power when the ring is on resonance.

tuner that is representative of closed-loop tuning solutions to date. The system uses a resistive heater embedded within the ring [88] or alongside the ring to change the ring's temperature, providing the means to control λ_0 through the thermo-optic effect. λ_0 tuning through carrier-injection [40] or carrier-depletion [83] have also been proposed. However, these are limited in tuning range and cannot tune out the magnitude of process variations. To detect the relative position of λ_0 to λ, a waveguide terminated by a photodetector is weakly coupled to the drop-port of the ring. A fraction of the power resonating in the ring couples to the drop-port waveguide and illuminates the photodetector to produce photocurrent. We show an example of a drop-port transmission response of a microring modulator as λ changes relative to λ_0 in Figure 3-15. Here, β_0 and β_1 denote the drop-port transmission for the *0* and *1* states of the modulator ring, respectively. When λ is far from λ_0 (ring is

off-resonance), very little amount of light resonates in the ring and very little power couples into the drop port to generate only a small amount of photocurrent. When $\lambda = \lambda_0$ (the ring is on resonance), a significant amount of power resonates in the ring and couples to the drop-port to generate a lot of photocurrent. As such, the drop-port photodetector is able to convert the quantity $|\lambda - \lambda_0|$ into a photocurrent amplitude response. Recall that the drop-port power, β_0 and β_1, constitute a fraction of the power that is "lost" by the microring cavity. As a result they are proportional to γ_0 and γ_1 (the power lost by the cavity), respectively. Because the data is amplitude modulated, the amount of drop-port power changes is different for whether we are in the *1* state or *0* state (reflected in the difference between β_0 and β_1 at the same λ.) The tracker circuit reads the photocurrent and performs the calculations needed to create an appropriate error signal to output to the controller. The controller applies a control algorithm to determine the amount of heating necessary and the driver circuit drives the integrated heater to apply the requested amount of heat.

3.5 Overview of Tuning Techniques

There are two functionally-important criteria for a tuning scheme. First, the tuning scheme's main goal is to tune out both static process variations and dynamic temperature fluctuations. Second, a tuner should be able to support the transmission of arbitrary data. In this section, we review the different thermal tuning schemes that have been proposed using these functional criteria. In doing so, we will decouple discussion of the controller and the tracker because, in most cases, they fulfill orthagonal roles and can be mixed and matched with each other to complete the list of features for the tuner.

3.5.1 Photocurrent Trackers

The first and arguably the most commonly used type of tracker is the averaging tracker, found in [2, 40, 42, 63, 82]. An averaging tracker simply reports the drop-port photocurrent to the controller for the controller to make its decisions. An averaging

tracker can be very sensitive, as it can average over an interval of time comparable to the thermal time constant of the ring (on the order of tens of microseconds), which is far slower than the bit-time (200 ps for a 5 Gb/s data rate). The chief drawback to an averaging tracker is that it requires a fixed ratio of *1*s to *0*s over the averaged interval to function correctly; since the data is amplitude modulated, a change in the ratio of *1*s to *0*s will change the average amount of light resonating in the cavity. As such, the average drop-port photocurrent will change and the corresponding controller will misinterpret this as a change in the ring's resonance and erroneously attempt to compensate.

The bit-error-rate (BER) tracking approach of [100] can address the issue of changing *1/0* ratios on the transmit side. In this approach, the tracker includes a monitoring receiver connected directly to the drop-port photodiode of the microring, resolving the photocurrent into bits. The tracker outputs BERs by comparing the output bitstream of the monitoring receiver to the bitstream driving the modulator. BERs for *1*s and *0*s are kept independently of each other, which allows it to operate independently of *1/0* ratios. Each BER is an indicator of degrading signal levels for the *1* and *0* levels and the controller can discern which way the ring resonance is drifting by just looking at the two BERs. However, the dynamic range of this tracker is very limited. Too low photocurrent results in a tracker that is inaccurate and too much photocurrent could result in too few bit-errors being generated, resulting in a blind tracker (or very slowly reacting) until the ring has drifted far enough off the desired lock point to cause a higher *1* or *0* BER. As such, this tracker will only work for photocurrents in the narrow window for which the receiver produces bit-errors but not too many. Finally, a long run length data sequence would prevent updates to one of the two BERs, blinding the tracker to one direction of resonance drift. Efficiency-wise, this approach requires a power-hungry receiver that resolve bits at data rate (shortest time interval in the system). At the same time, the controller cannot leverage the full-rate of information available from each individual resolved bit due to the wait for statistics collection.

	Averaging	BER	Level-Track	Peak Detect
Changing 1/0 Ratios	✗	✓	✓	✓
Long Run-Length	✗	✗	✓when $\times 2$	✓when $\times 2$
Low Bandwidth Frontend	✓	✗	✗	✗

Table 3.1: Comparison between the different types of trackers. A component labeled "when $\times 2$" applies when at least two trackers are needed.

A more refined approach utilizes conditioned level-tracking circuitry to decouple and monitor the *1* and *0* photocurrent levels independently [2, 76]. Here, a comparator performs a comparison between a value, L, and the drop-port photocurrent, I_{PD}. This forms a bang-bang control loop, whereby L increases if the $L > I_{PD}$ and decreases if $L < I_{PD}$. Updates to L are conditioned by the transmitted or received bitstream. If we only allow L to update on bit *1*s, L tracks the *1*-level. If we only let L update on bit *0*s, L tracks the *0*-level. Allowing updates to L on both *0*s and *1*s will let L track the average photocurrent and work as an averaging tracker. To conserve power, the comparator may also sub-sample. The peak detector of [96] functions similarly to the level-tracker, though the conditioning is performed at the analog frontend through a source-follower circuit as opposed to in the digital domain. A single level-tracker or peak detector, however, can still be blinded by long run lengths (or antagonistic data sequences, if subsampled) of the datastream with only the "no-update" bit, which will cause L to never update. This can be resolved by putting down two level-trackers or peak detectors, one for tracking *0*s and one for *1*s, such one tracker will always be able to update. In this case, however, the controller must be given knowledge of which of the two trackers to "trust" (such as by counting number of updates each of the trackers received) prior to making its decisions. Both level-tracking and peak detector architectures require a high-speed TIA frontend or a high-speed sampler, consuming a large amount of power at higher data-rates. The higher bandwidth of the system also leaves it more noisy than an averaging tracker. While this component may be shared with the high-speed datapath in the receiver, this is not an option in the case of a transmitter. As a result, using these tuning options in the transmitter will add a significant power overhead.

Table 3.1 summarizes and compares the capabilities of the discussed trackers.

3.5.2 Wavelength-Locking Controllers

The controller controls the heating of the ring to keep the ring wavelength-locked. A controller's input is the data from the tracker, L. The controller makes a decision on and outputs a heating value, P_H, which is carried out by the heater driver. The simplest type of controller is a *lock-to-reference* controller (LTR) [40, 42, 63, 100], which attempts to lock L (an average photocurrent in [40, 42, 63] or a BER in [100]) to a reference value, L_{REF}. Every time an LTR controller makes a decision, it simply increases or decreases P_H based on whether L is smaller or larger than L_{REF}. Though simple, an LTR controller can only lock onto a region for which L is monotonic. For the modulators in [63], this meant that it could not lock onto the region for which modulation depth was maximized. A second drawback of an LTR controller is the non-automatic nature in picking L_{REF}. Though L_{REF} can be calibrated for manually, an automatic optimization of L_{REF} is necessary for cases when the optical insertion loss can change dynamically, such as the case in switched optical paths. To accomplish this, we can build a separate search loop [40] to find an appropriate value for L_{REF} during initialization. If L_{REF} is picked appropriately, however, an LTR controller can guarantee that the ring is never *blue-biased*, preventing the system from falling into an unstable state from which it cannot recover from.

An alternative approach to LTR is *lock-to-maximum* (LTM) [82]. Here, the controller will always attempt to maximize the value of L from the tracker. One such implementation of an LTM controller will make the same decision to P_H if it corresponds to an increase in L from the last time it checked. Otherwise, if the last decision resulted in L decreasing, then make the opposite of the last decision. An LTM controller does not require a separate L_{REF} to be provided and will always find the local maximum (or minima). In the case where L is tracking average power, an LTM controller will bring the ring to the point of maximum average power. In a receiver or filter ring, this is the point where the laser $\lambda = \lambda_0$ and the optimal lock point. However, in a modulator, this is not necessarily the optimal lock point. A drawback of LTM is that it may unknowingly bring the ring into a *blue-biased* state.

Under high laser power, positive feedback from the laser-induced self-heating can make the system unstable and the LTM controller can no longer recover. Another drawback of LTM is the lack of directional awareness of where the ring is relative to the desired lock point; when the ring drifts off-resonance, the LTM controller has only a 50/50 chance that the first decision it takes is in the correct direction. Though the controller will eventually move the ring back to the optimal point, this increases the coarseness of the lock.

To mitigate the lack of directional awareness, [64] uses a dithering technique to create an error signal that is monotonic around the optimal lock point. In this case, when the ring is perturbed, an LTR controller knows the exact direction to take to move the ring back to the lock point. When tracking average power, the optimal lock point corresponds to the point where the error signal is zero, obviating the need for a secondary loop to find L_{REF} of the desired lock point (though some calibration may still be needed to remove any DC offsets from the error signal). With this knowledge, the controller can also avoid the *blue-biased* regime.

	Open-Loop		Closed-Loop			
	Thermal	Athermal	LTR	LTM	Dither+LTR	Eye-Max
Process Variations	✓	✗	✓	✓	✓	✓
Ambient Temperature	✗	✓	✓	✓	✓	✓
Optimal Lock Point			✗	✗/✓	✗/✓	✓
Direction Awareness			✓	✗	✓	✗
Blue-Bias-Safe			✓	✗	✓	✗
Self-Heat Eye Close			✗	✗	✗	✗

Table 3.2: Comparison of the supported features of different types of tuning controllers. We include the open-loop thermal and athermal approaches for comparison. LTR is a lock-to-reference controller, LTM is a lock-to-maximum controller, Max-Eye locks to the point of the maximum eye opening.

If the tracker can track both *0*-levels and *1*-levels, L_0 and L_1, we can enable an *eye-max* controller, a variant of LTM. An *eye-max* controller maximizes the value of $L_1 - L_0$, corresponding to the point where the difference between L_1 and L_0 seen at the drop port. For a receiver, this value corresponds directly to the receive eye-height. In a transmitter, the point where $L_1 - L_0$ is maximized at the drop port is also coincident to the point of maximum transmit eye-height (which is from the through

port). As such, the ring will always be locked to the lock point giving greatest eye-height, an optimal case for both transmit microrings and receiver microrings. Like an LTM controller, however, Eye-Max also lacks directional awareness and can bring the ring into the unstable *blue-biased* regime.

We note that all the presented controllers have a decision rate that is slower than the time constant of wavelength actuation ($\tau_T/(1 - G_L(0))$ using heaters), otherwise the controller response will overshoot or oscillate. Though this is fast enough to account for ambient temperature changes (which affect the rings slower), controllers cannot address the eye-closure effects of laser-induced self-heating which happens right in the ring waveguide. Thus, whenever the ratio of *1*s to *0*s changes, the change in power creates a change in λ_0 before the controller is able to react to it, resulting in a change in vertical eye-height and a shift in the vertical eye location. The shift in the eye location due to self-heating can be mitigated at the receiver through dynamic decision threshold adjustment [2, 12]. However, the decrease in the eye height cannot be mitigated at the receiver end. Table 3.2 summarizes the supported features of the discussed tuning controllers.

3.6 Summary

Microring resonators are essential to the energy-efficiency and bandwidth density of DWDM optical links. As a high-Q device, however, a microring is highly sensitive and requires precise alignment of its resonant wavelength to the laser wavelength in order to perform modulation or filtering functions, necessitating active feedback mechanisms to stabilize the resonance. In this chapter, we discussed the effects of both static process variations and dynamic temperature fluctuations on microring resonators. In particular, we focused on the dynamics of laser-induced self-heating of the microring, which produces a feedback relation that affects how the ring reacts to different perturbations. In particular, the tuning efficiency factor, H_T, can be thought of as a double-edged sword; a higher H_T improves the efficiency of the thermal tuning but also increases the effects of self-heating, creating stability challenges and

degradation of active devices. With these insights, we classified and compared the landscape of proposed controller-based tuning solutions. We found that many existing tuning solutions lack the ability to handle non-encoded data channels, which can have long run-lengths and arbitrarily changing ratios of 1s to 0s. These effects cause transient eye-closure due to self-heating and create the need for more sophisticated photocurrent tracking circuits to obtain the appropriate error signal for the tuning control logic. In the next chapter, we describe the tracking and control solution that addresses the shortcomings of the attempts presented here.

Chapter 4

Bit-Statistical Thermal Tuning

In the previous chapter, we discussed the effects of thermal variations as well as the advantages and disadvantages of the control schemes proposed to maintain microring functionality in a hostile thermal environment. We noted the necessity for independent level-tracking of 1-levels or 0-levels to enable the controller to track the ring resonance correctly if the data sequences changes the ratios of *1*s to *0*s. Level-tracking comes at a cost, however; it requires a high-speed sampler (which requires a bandwidth identical to that of a receiver) and sub-bit-time alignment of the sampling position in order to get an accurate read on the power level. These factors make the level-tracker approach more costly in energy and less sensitive than an averaging photocurrent tracker. Additionally, we noted that the level-tracker method is blind for specific data sequences; if we track only 0-levels, the tracker cannot update for a data sequence containing only 1s. Conversely, if we track 1-levels, the tracker cannot update for a data sequence containing only 0s. As such, there is a minimum required density of 1s or 0s in the data stream in order to accurately track the optical power. If we choose to duplicate the circuitry to track both 1-levels and 0-levels simultaneously (which increases the energy cost of the tracker), it is still unclear to the controller which of the two values it should use to make tuning decisions, as either one could be unreliable due to insufficient updates. Despite these shortcomings, level-trackers do introduce one important concept: the tracker can use the transmitted or received bitstream as a conditioning signal to independently track 1-levels and 0-levels. Could

we leverage statistics of this bitstream to develop a more generalized tracker that is just as sensitive and efficient as an averaging tracker but is also truly pattern independent? Furthermore, can we leverage additional information that the tracker provides to resolve the drawbacks of previous controller designs?

4.1 A Bit-Statistical Tracking Method

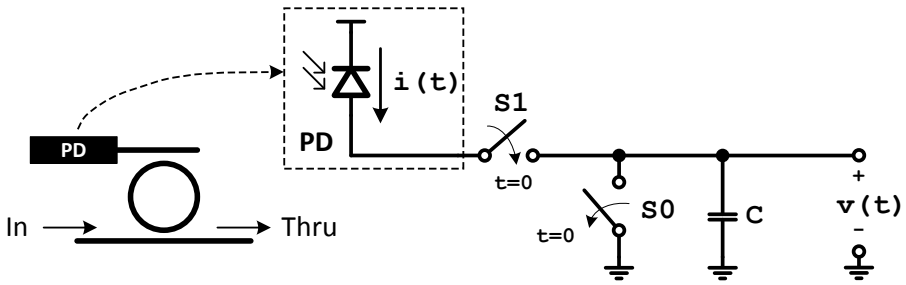

Figure 4-1: A drop-port photodetector connected to a photocurrent integrator. Prior to $t = 0$, switch *S0* is closed and discharges the capacitor to ground, hence $v(0) = 0$. For $t > 0$, switch *S0* opens and switch *S1* closes, charging the capacitor with photocurrent, $i(t)$, from the photodetector.

In order to design a tracker that combines the best the both worlds – both independent level-tracking as well as high sensitivity – we first consider the case of a drop-port photodetector that is connected to a photocurrent-to-voltage integrator, shown in Figure 4-1. We can write an equation for the voltage output of the integrator:

$$v(t) = \frac{1}{C}\int_0^t i(t)dt \tag{4.1}$$

where $v(t)$ is the voltage across the capacitor, C is the capacitance and $i(t)$ is the photocurrent. Now, consider the case when light imprinted with a digital bitstream, with a bit-time of t_{bit} $(1/f_{data})$, illuminates the drop-port photodetector. The bit-time for a gigabit link is on the order of hundreds of picoseconds and is many orders of magnitude shorter than the microsecond-scale thermal time constant of the ring $(t_{bit} \ll \tau_T)$. As such, the temperature of the ring changes negligibly during a bit time and $i(t)$ is a function of only the bit pattern. As such, we can rewrite Equation 4.1

as a photocurrent sum for an interval N discrete bits:

$$v_N = \frac{t_{bit}}{C} \sum_{n=1}^{N} i(n) \tag{4.2}$$

Next, we define i_1 and i_0 as the photocurrents of an optical *1* or *0*, respectively. If we pick N to be small enough such that $N \cdot t_{bit} \ll \tau_T$, then i_0 and i_1 are constant over the N-bit interval. As it is a binary sequence, we can decompose N into N_1 and N_0, which are the numbers of *1s* and *0s*, respectively, during this interval ($N_0 + N_1 = N$). Substituting into Equation 4.2, we can remove the summation:

$$v_N = \frac{t_{bit}}{C} \left(N_0 \cdot i_0 + N_1 \cdot i_1\right) \tag{4.3}$$

Equation 4.3 defines what the voltage will be if we integrate the photocurrent over an interval of N bits. Note that the values of i_0 and i_1 are precisely what the tracker wants to track and what we should ultimately solve for. N_0 and N_1 are values that can be calculated using the transmitted or received bitstream and v_N can be measured. Unfortunately, this is still an equation with two unknowns (i_0 and i_0) and not directly solvable. However, consider the case where we integrate two voltages for N bits starting at different times t_a and t_b, obtaining two expressions for the voltage at the end of the two intervals:

$$v_N(t_a) = \frac{t_{bit}}{C} \left[N_0(t_a) \cdot i_0(t_a) + N_1(t_a) \cdot i_1(t_a)\right] \tag{4.4a}$$

$$v_N(t_b) = \frac{t_{bit}}{C} \left[N_0(t_b) \cdot i_0(t_b) + N_1(t_b) \cdot i_1(t_b)\right] \tag{4.4b}$$

As an example of the notation, $N_0(t_a)$ and $N_0(t_b)$ denote the number of *0*s transmitted or received during the interval starting at t_a and t_b, respectively. $N_0(t_a) + N_1(t_a) = N_0(t_b) + N_1(t_b) = N$ If t_a and t_b are close enough in time, such that $t_a - t_b \ll \tau_T$ then

$i_0(t_a) \approx i_0(t_b) \approx i_0$ and $i_1(t_a) \approx i_1(t_b) \approx i_1$, and we arrive at:

$$v_N(t_a) = \frac{t_{bit}}{C} \left[N_0(t_a) \cdot i_0 + N_1(t_a) \cdot i_1 \right] \tag{4.5a}$$

$$v_N(t_b) = \frac{t_{bit}}{C} \left[N_0(t_b) \cdot i_0 + N_1(t_b) \cdot i_1 \right] \tag{4.5b}$$

If the numbers of ones and zeros are different in the two intervals, where $N_0(t_a) \neq N_0(t_b)$ and $N_1(t_a) \neq N_1(t_b)$, the two equations are linearly independent and i_0, i_1 can be solved for directly:

$$i_0 = \frac{C}{t_{bit} \cdot N} \left[\frac{v_N(t_b) \cdot N_1(t_a) - v_N(t_a) \cdot N_1(t_b)}{N_0(t_b) - N_0(t_a)} \right] \tag{4.6a}$$

$$i_1 = \frac{C}{t_{bit} \cdot N} \left[\frac{v_N(t_a) \cdot N_0(t_b) - v_N(t_b) \cdot N_0(t_a)}{N_0(t_b) - N_0(t_a)} \right] \tag{4.6b}$$

With Equation 4.6, we have derived a basic method to obtain both i_0 and i_1, the photocurrent levels we wish to track, using statistics of a transmitted or received bitsteam. We will explain in Section 4.2 how to deal with the case where $N_0(t_a) = N_0(t_b)$, which causes a division-by-zero in the calculations of i_0 and i_1.

4.1.1 Towards a Digital Implementation

From the nature of the calculations (several multiplications and a divide), the implementation of the computation inside the tracker will be digital. As such, all inputs to the computation must be digital. From Equation 4.6, the only non-constant inputs are N_0, N_1, and v_N. The N_0, N_1 counts can be obtained directly by using a counter to count the stream of bits going towards the transmitter or coming from the receiver. The integrated voltage, v_N, necessitates an analog to digital converter (ADC) in the frontend to convert v_N to a digital value L_N:

$$L_N = v_N \cdot G_A \tag{4.7}$$

G_A is the analog voltage to digital conversion ratio of the ADC, and has units of LSBs/V. The overall structure of the tracker is shown in Figure 4-2.

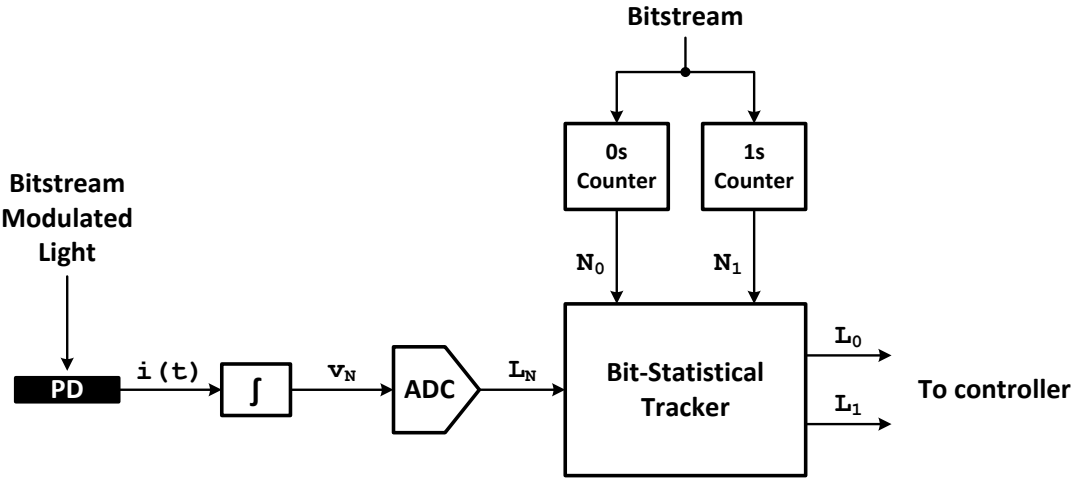

Figure 4-2: Structure of the bit-statistical tracker, showing the inputs and outputs. The ADC and photocurrent integrator form an analog frontend.

We define the variables L_0 and L_1 to be what the *0-level* and *1-level* are, respecively, represented digitally. L_0 and L_1 are related to i_0 and i_1 through a composite of C, t_{bit}, N, and G_A constants:

$$L_0 = \frac{G_A \cdot N \cdot t_{bit}}{C} \cdot i_0 \tag{4.8a}$$

$$L_1 = \frac{G_A \cdot N \cdot t_{bit}}{C} \cdot i_1 \tag{4.8b}$$

Taken physically, L_0 and L_1 are equivalent to the ADC output, L_N, if the integration interval consisted of N bits of all *0*s or all *1*s, respectively. We can then substitute i_0, i_1, and v_n of Equations 4.5 and 4.6 for L_0, L_1, and L_N:

$$L_N(t_a) = \frac{1}{N}\left[N_0(t_a) \cdot L_0 + N_1(t_a) \cdot L_1\right] \tag{4.9a}$$

$$L_N(t_b) = \frac{1}{N}\left[N_0(t_b) \cdot L_0 + N_1(t_b) \cdot L_1\right] \tag{4.9b}$$

$$L_0 = \frac{L_N(t_b) \cdot N_1(t_a) - L_N(t_a) \cdot N_1(t_b)}{N_0(t_b) - N_0(t_a)} \tag{4.9c}$$

$$L_1 = \frac{L_N(t_a) \cdot N_0(t_b) - L_N(t_b) \cdot N_0(t_a)}{N_0(t_b) - N_0(t_a)} \tag{4.9d}$$

Equation 4.9 forms the basics of a digital implementation of the tracker.

4.2 Achieving Pattern Independence

As explained previously, the current tracker implementation is not fully bit pattern independent; if $N_0(t_a) = N_0(t_b)$, we cannot recover L_0 and L_1. This condition occurs for many data sequences; a constant stream of *0*s, a constant stream of *1*s, or any bit pattern in which N_0 and N_1 are always constant over an N-bit integration interval (such as a clock pattern) will trigger the condition and prevent new values for L_0 and L_1 from being calculated. Moreover, the introduction of an ADC with finite resolution, linearity, and noise performance causes an accuracy issue even in the case where $N_0(t_a)$ and $N_0(t_b)$ are different, but are close in value. Here, the denominator terms of Equation 4.9 will be small, amplifying errors in L_N from the ADC and any loss of precision from subsequent algebraic operations. This distorts the calculated L_0 and L_1 significantly.

To mitigate this error, we can choose to disallow updates to L_0 and L_1 unless $N_0(t_b) - N_0(t_a)$ is greater than some threshold, N_{thres}:

$$L_0, L_1 = \begin{cases} \text{new values} & \text{if } |N_0(t_b) - N_0(t_a)| > N_{thres} \\ \text{keep old values} & \text{otherwise} \end{cases} \tag{4.10}$$

Though the errors to L_0 and L_1 decrease with increasing N_{thres}, the probability that the tracker can make an update to L_0 and L_1 also decreases. Therefore, as N_{thres} grows, the number of adversarial data patterns for which the tracker is blind (never able to update L_0 and L_1) also grows.

To resolve this issue, we introduce a *level difference* term, L_d, such that $L_d = L_1 - L_0$. We can formulate an expression for L_d that consists of only the ADC

output, L_N, and bit-counts N_0, N_1 by making a few substitutions into Equation 4.9:

$$L_N(t_a) = \frac{1}{N}\left[N_0(t_a)\cdot(L_1 - L_d) + N_1(t_a)\cdot L_1\right] \tag{4.11a}$$
$$= \frac{1}{N}\left[(N_0(t_a) + N_1(t_a))\cdot L_1 - N_0(t_a)\cdot L_d\right] \tag{4.11b}$$
$$= L_1 - \frac{N_0(t_a)}{N}\cdot L_d \tag{4.11c}$$
$$L_N(t_b) = \frac{1}{N}\left[N_0(t_b)\cdot(L_1 - L_d) + N_1(t_b)\cdot L_1\right] \tag{4.11d}$$
$$= \frac{1}{N}\left[(N_0(t_b) + N_1(t_b))\cdot L_1 - N_0(t_b)\cdot L_d\right] \tag{4.11e}$$
$$= L_1 - \frac{N_0(t_b)}{N}\cdot L_d \tag{4.11f}$$
$$L_N(t_a) - L_N(t_b) = L_1 - \frac{N_0(t_a)}{N}\cdot L_d - L_1 + \frac{N_0(t_b)}{N}\cdot L_d \tag{4.11g}$$
$$= \frac{1}{N}\left[N_0(t_b)\cdot L_d - N_0(t_a)\cdot L_d\right] \tag{4.11h}$$
$$L_d = \frac{N\left[L_N(t_a) - L_N(t_b)\right]}{N_0(t_b) - N_0(t_a)} \tag{4.11i}$$

We can then rewrite L_0, L_1 in terms of L_d:

$$L_1 = L_N + \frac{N_0}{N}\cdot L_d \tag{4.12a}$$
$$L_0 = L_1 - L_d \tag{4.12b}$$

Plots of L_0, L_1, and L_d for typical transmit and receive rings are shown in Figure 4-3.

The calculation of L_1 and L_0 through an intermediate term, L_d, decouples updates to L_d from updates to L_1 and L_0; though we cannot compute a new value for L_d in the case when $N_0(t_a) = N_0(t_b)$, new values for L_1 and L_0 may always be calculated using the current value of L_d. Any new value of L_1 and L_0 calculated without a corresponding update to L_d, however, will be erroneous, as L_d will certainly change as the ring drifts around. Despite this, the new values of L_1 and L_0 still allow a ring to be tuned correctly through the following observation.

Consider the case of a ring that is wavelength-locked by a *lock-to-reference* controller, which keeps the ring at a point where $L_1 = L_{ref}$. For thermal stability, the controller red-biases the ring ($\lambda < \lambda_0$), red-shifting λ_0 if $L_1 > L_{ref}$ and blue-shifting

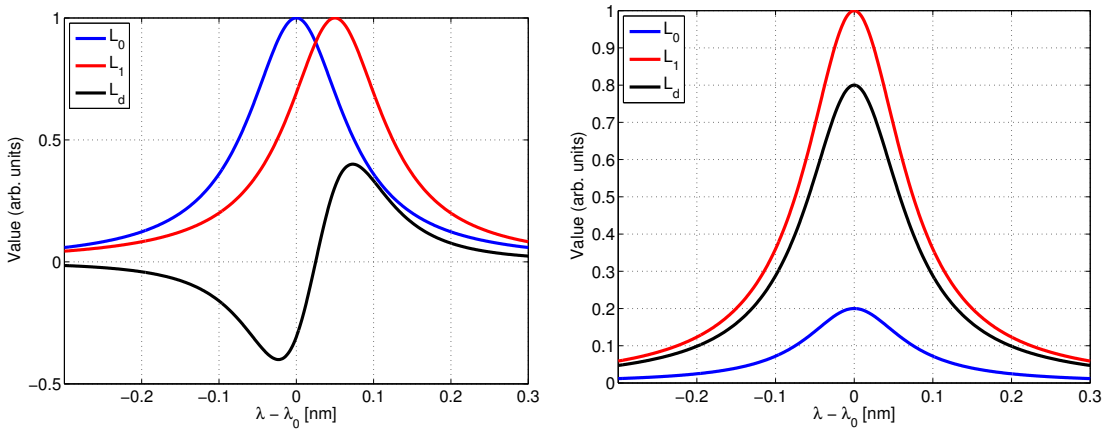

(a) Transmit Ring (b) Receiver Ring

Figure 4-3: L_0, L_1, and L_d plots for a transmit ring and a receiver ring, taken using a drop-port photodetector. For a receiver ring, the photodetector could also be embedded in the ring itself. λ is the laser wavelength, and λ_0 is the ring's nominal resonance. The X-axis indicates the frequency offset between λ and λ_0. In both cases, we assume rings with a full-width half-maximum bandwidth of 0.15 nm. The transmitter ring has a modulation shift (between a zero and a one) of 50 pm and the receiver ring is receiving on-off-keyed optical data with a 5-to-1 (7 dB) on-to-off ratio.

if $L_1 < L_{ref}$. At the lock point, $L_d = L_{d-ref}$. Now, suppose the transmitted or received data sequence changes to that of a clock waveform, such that L_d may never be able to update again and is stuck forever at L_{d-ref}, shown in Figure 4-4. For both the receive ring and the transmit ring cases, observe that though the calculated L_1 begins to deviate from the real L_1 the farther the ring is from the lock point, the calculated L_1 still follows the same trend as the real L_1 (they will both increase or decrease) close to the lock point. Consequently, if the ring drifts off the lock-point, both the real and calculated L_1 change in the same direction to be greater than or equal to L_{ref}. The *lock-to-reference* controller will thus make the same decision with either the real L_1 or the calculated L_1, returning the ring correctly back to lock-point, where both the real and calculated $L_1 = L_{ref}$.

Though the example assumes a 50 % ratio of *0*s in the data sequence with a controller using L_1, we note this property is true even for arbitrary ratios of *0*s or with a controller that locks using L_0. These results have major implications. First, we have achieved complete pattern independence; by decoupling L_d calculations, the

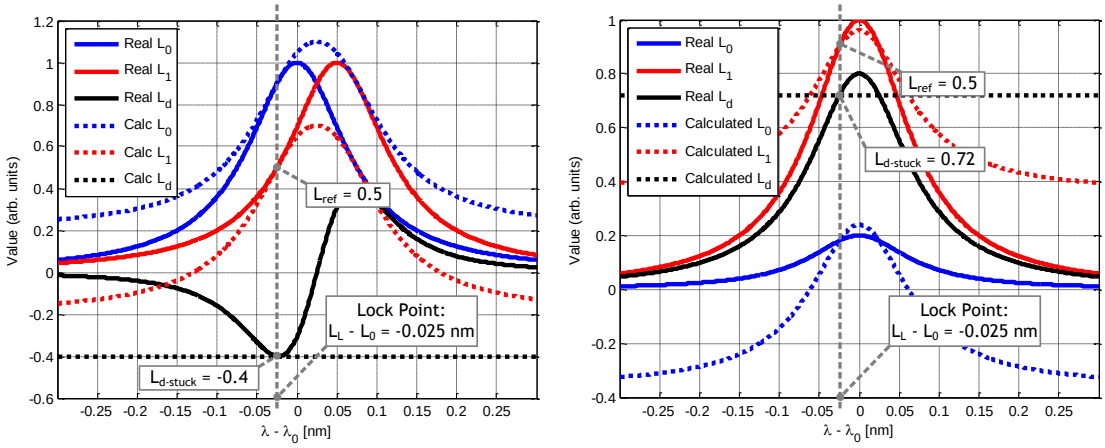

(a) Transmit Ring

(b) Receive Ring

Figure 4-4: Real L_0, L_1, and L_d versus calculated L_0, L_1, L_d, plotted at various laser wavelengths (λ) relative to the nominal resonance λ_0. The calculated values assume that L_d cannot update and is stuck at the value of the original lock point, L_{d-ref}. The desired lock point is indicated by the vertical line, corresponding to a red-biased ring with $\lambda - \lambda_0 = 0.025\,\text{nm}$. Plots are for a 50 % ratio of *0*s in the data sequence. Note that the receive ring is *red-biased* slightly off resonance such that the *lock-to-reference* controller has a monotonic error signal around the lock point.

tracker is never blind and provides sufficient information for the controller to lock the ring, regardless of the transmitted or received bit pattern. Second, we note that updates to L_d are not even necessary in locked operation; after the tracker *learns* the L_d at the lock point (which it can do during initialization), we no longer need to update it further to keep the ring locked. As such, circuitry performing the L_d calculation may be powered down completely in normal locked operation after the initial lock has been achieved. Only the circuitry responsible for calculating L_0 and L_1 needs to continue working to provide sufficient information to the controller to maintain the lock.

4.3 Controller Design for a Statistical Tracker

The controller of a tuning system must perform two roles. The first role is initialization, for which the controller must first find the lock point to lock to and then

move the ring λ_0 to that point. The second role is to maintain the lock to the chosen lock point. A controller compatible with the bit-statistical tracker will take as inputs L_0, L_1, and L_d, and output P_H, the desired heater value output. In this section, we will design this controller.

4.3.1 Controller Observations

Prior to initialization, it is reasonable to assume that λ_0 is initially some distance away from the laser wavelength, λ, such that the ring is completely off-resonance and that the drop-port photodetector is picking up only minute amounts of laser power ($L_0 = L_1 = L_d = 0$). As such, the first goal of initialization is to to use the heater to sweep λ_0 to search for λ and generate sufficient drop-port laser power to perform more complex decision-making. The "search" task can be performed in one of two ways. We can perform either a *sweep up* search, where the heater output power P_H starts at 0 and is stepped up, or a *sweep down* search, where the heater output power starts at the maximum value and is stepped down.

The optimal lock point for both transmit and receive microrings is one that provides the greatest optical eye height. Using the value of L_d computed by the tracker, which corresponds to the eye height for a receiver and is indicative of the eye height for a transmitter, the controller can also find the optimal lock point to lock onto. Once the "search" task brings λ_0 close to λ, the controller can continue the heater sweep to map out L_d over all potential $\lambda - \lambda_0$ lock points to find the lock point that maximizes the eye height. In a modulator, there are two points that achieve this: the first is a *red-biased* point where L_d is negative and minimized, the second is a *blue-biased* point where L_d is positive and maximized. In a receiver, the lock-point of maximum eye-height occurs at $\lambda_0 = \lambda$. Depending on the laser power, the optimal lock points may fall into a bistable state, the unstable state, or a stable state. As a result, the choice of *sweep up* or *sweep down* depends on the state the desired lock point is in. Using the notation of Figure 3-10, if the desired lock point falls within bistable state B, we can only reach it via a *sweep down* sweep. On the other hand, if the final lock position is in bistate C, we would require a *sweep up* sweep. Both

sweep up and *sweep down* can reach lock points that fall into states A and D. In the case of the transmitter microring, the *red-biased* optimal lock point will generally fall into bistable state B while the *blue-biased* optimal lock point is often unstable at moderate laser powers. In a receiver, the optimal lock point is at the edge of bistable state B, where it sits dangerously close to the potentially unstable region. Note that during initialization, we are not transmitting real data and may use a data sequence favorable to the operation of the bit-statistical tracker so that L_d can be updated frequently.

For generality, we define two additional variables, L_{opt} and L_{track}, which correspond to the variable to optimize (maximize or minimize) when picking the lock point and the variable the controller tracks to maintain the lock, respectively. L_{opt} and L_{track} can each be picked to be L_0, L_1, or L_d. As an example, A controller which picks an "eye-maxed" lock point but tracks L_0 to maintain the lock would set $L_{opt} = L_d$ and $L_{track} = L_0$.

The decoupling of the tracked variable from the optimized variable allows us to pick a suitable pair of L_{track} and L_{opt}. To understand the advantages, consider the case where these two variables are not decoupled and we track the same variable we optimize for ($L_{opt} = L_{track}$). At the point where L_{opt} is maximized, $dL_{opt}/d\lambda_0 = dL_{track}/d\lambda_0 = 0$. If the controller has limited resolution (such as from limited ADC resolution) on L_{track}, a large change in λ_0 needs to occur before the controller registers a 1 LSB change to L_{track}, lowering the lock resolution. In addition, if L_{opt} drops from the optimal value, the controller will lack the appropriate directional awareness to steer λ_0 back and may need to make the wrong choice first. These are precisely the issues faced by the *lock-to-max* and *eye-max* controllers we compared in Chapter 3, which can be thought of as controllers where $L_{opt} = L_{track} = L_{0,1}$ and $L_{opt} = L_{track} = L_d$, respectively. An approach leveraging the decoupled L_{track} and L_{opt} would instead pick an L_{track} that is monotonic around the region where L_{opt} is optimal. Using Figure 4-3 for the transmitter, we can see that with an *eye-max* optimization, where $L_{opt} = L_d$, picking $L_{track} = L_0$ or $L_{track} = L_1$ would provide us precisely this behavior. This enables us to use a *lock-to-reference* controller to maintain the lock while still

maintaining the optimal lock point advantages that an *eye-max* controller brings. Note that in the case of a receiver ring, the point of optimal L_d coincides with the point where $dL_0/d\lambda_0 = dL_1/d\lambda_0 = 0$ and there is no signal we can pick for L_{track} that is monotonic around an optimal L_d. In this situation, we can still leverage the search to find the optimal L_{opt} but simply back away from the optimal by a few LSBs on L_{track}. This way, we can still use a *lock-to-reference* controller to maintain the lock while remaining negligibly close to the optimal point.

4.3.2 Controller Design

We can use these observations to build a suitable controller design and we describe the functionality of the controller design as follows. At the start, *init* switches the transmitted data to the training sequence and sets P_H to an initial value P_{H-init}. The controller then enters the *search* state which steps P_H in large strides ($P_{H-stride}$)until L_{track} is above a set threshold R_{sweep}, indicating that λ_0 is close to λ. Next, *sweep* steps P_H by small steps (P_{H-step}) as the controller maps the shape of the resonance, finding P_{H-opt}, $L_{track-opt}$, and $L_{opt-opt}$ at the optimal lock point, where L_{opt} is maximized or minimized. When L_{track} is again smaller than R_{sweep}, indicating that the sweep has again moved λ_0 far from λ, the controller exits the *sweep* state. To return back to the optimal lock point, the *reset* and *return* states set $P_H = P_{H-init}$ before setting $P_H = P_{H-opt}$. These two states ensure that the heater sweep direction (*sweep up* or *sweep down*) is the same when returning to the optimal lock point as when it was reached during the *sweep* state. This eliminates the possibility of returning to the wrong bistable state given optical bistability. The *lock* state maintains the lock under thermal perturbations. Once in the *lock* state, the controller indicates to the tracker that it is locked, and the tracker no longer needs to update L_d. At this point, normal data can be sent again, as the tuner becomes agnostic to the data sequence and is able to keep the ring locked under any condition.

Figure 4-5 illustrates the simulated state of the tuning controller as it transitions through all the states. It shows the levels of L_0, L_1, L_d, P_H, and the through-port

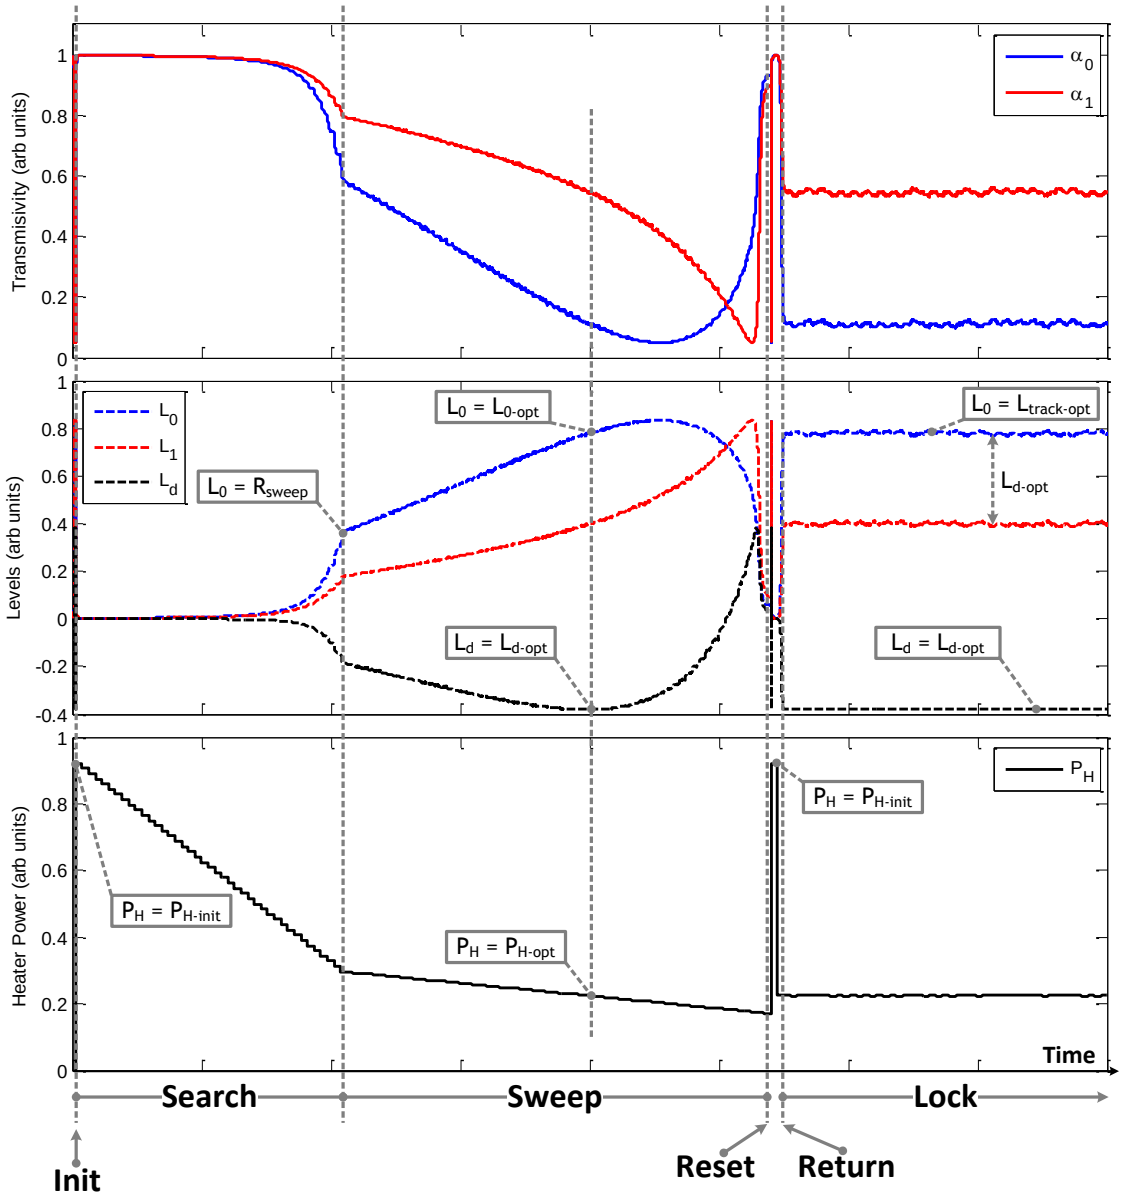

Figure 4-5: Controller state during the lock-on process. As the controller uses $L_{opt} = L_d$ and $L_{track} = L_0$, $L_{opt-opt} = L_{d-opt}$ and $L_{track-opt} = L_{0-opt}$. The value of the important levels are labeled at each significant point in the state evolution of the system.

transmisivity α_0 and α_1 at all times as the controller transitions from *init* to *lock*. Here, the controller locks a modulator ring using $L_{opt} = L_d$ and $L_{track} = L_0$.

4.4 Transient Self-Heating Cancellation

Self-heating from the laser produces very large heating power perturbations in the ring waveguide and causes transient eye closures whenever the ratios of *0*s and *1*s changes due to amplitude modulated data. This an issue that the main controller cannot tackle due to the far faster response compared to ambient temperature variations, necessitating a different approach.

We consider the case of self-heating in a modulator. For a modulator, if r_0 and r_1 are the ratios of *0*s and *1*s in the transmitted data, then the average amount of self-heating power from the laser is:

$$\overline{P_L'} = \kappa \cdot P_L \cdot [r_0 \cdot \gamma_0(\lambda) + r_1 \cdot \gamma_1(\lambda)] \tag{4.13a}$$

$$= \kappa \cdot P_L \cdot [r_0 \cdot \gamma_0(\lambda) + (1 - r_0) \cdot \gamma_1(\lambda)] \tag{4.13b}$$

$$= \kappa \cdot P_L \cdot [\gamma_1(\lambda) + r_0 \cdot (\gamma_0(\lambda) - \gamma_1(\lambda))] \tag{4.13c}$$

Since r_0 could change depending on the transmitted data sequence, this is the component that causes a change in self-heating power and the resultant fast transient eye closure. We will write this term as $\Delta\overline{P_L'}$:

$$\Delta\overline{P_L'} = \kappa \cdot P_L \cdot r_0 \cdot (\gamma_0(\lambda) - \gamma_1(\lambda)) \tag{4.14a}$$

In order to prevent the transient eye closure, we must find a way to negate the data-dependent $\Delta\overline{P_L'}$ term. A simple approach to this is by just using the heaters to deliver an extra amount of power, P_C, that is exactly opposite in sign to $\Delta\overline{P_L'}$. This power would be added to the power output specified by the controller and would quickly

change based on r_0:

$$P_C = -\Delta\overline{P_L'} \tag{4.15a}$$
$$= -\kappa \cdot P_L \cdot r_0 \cdot (\gamma_0(\lambda) - \gamma_1(\lambda)) \tag{4.15b}$$
$$= \kappa \cdot P_L \cdot r_0 \cdot (\gamma_1(\lambda) - \gamma_0(\lambda)) \tag{4.15c}$$

In order to measure r_0, we can again leverage bit-statistics from the datapath, so that $r_0 = N_0/N$ for an appropriate choice of N (which sets the bandwidth of the cancellation loop):

$$P_C = \kappa \cdot \frac{P_L}{N} \cdot (\gamma_1(\lambda) - \gamma_0(\lambda)) \cdot N_0 \tag{4.16a}$$
$$= K_C' \cdot N_0 \tag{4.16b}$$

Note that we defined a quantity, $K_C' = \kappa \cdot \frac{P_L}{N} \cdot (\gamma_1(\lambda) - \gamma_0(\lambda))$. As such, one implementation of P_C would just involve a single multiplier which multiplies a configurable value of K_c' to N_0. In the expression for K_C', however, only κ and N can be assumed constant. $\gamma_1(\lambda)$, $\gamma_0(\lambda)$, and P_L (if insertion loss to the transmitter changes) can all change with time or with the lock point of the system. As such, this would require K_C' to be constantly reconfigured and is not practical.

As an alternative approach, recall that the drop-port transmisivity (β) constitutes a fraction of the light lost by the cavity and is proportional to γ through a factor, which we will call K_d. As a result, the drop-port photocurrent is proportional to $\gamma(\lambda)$:

$$i_0 = P_L \cdot \beta_0(\lambda) \tag{4.17a}$$
$$= P_L \cdot K_d \cdot \gamma_0(\lambda) \tag{4.17b}$$
$$i_1 = P_L \cdot \beta_1(\lambda) \tag{4.17c}$$
$$= P_L \cdot K_d \cdot \gamma_1(\lambda) \tag{4.17d}$$

Since L_0 and L_1 from the tracker are proportional to i_0 and i_1 through Equation 4.8, we can rewrite γ_0 and γ_1 in terms of L_0 and L_1:

$$\frac{C}{G_A \cdot N \cdot t_{bit}} \cdot L_0 = P_L \cdot K_d \cdot \gamma_0(\lambda) \tag{4.18a}$$

$$\frac{C}{G_A \cdot N \cdot t_{bit}} \cdot L_1 = P_L \cdot K_d \cdot \gamma_1(\lambda) \tag{4.18b}$$

$$\gamma_0(\lambda) = \frac{C}{G_A \cdot N \cdot t_{bit} \cdot P_L \cdot K_d} \cdot L_0 \tag{4.18c}$$

$$\gamma_1(\lambda) = \frac{C}{G_A \cdot N \cdot t_{bit} \cdot P_L \cdot K_d} \cdot L_1 \tag{4.18d}$$

In light of this, the quantity $\gamma_1 - \gamma_0$ can be directly written as a function of L_d.

$$\gamma_1(\lambda) - \gamma_0(\lambda) = \frac{C}{G_A \cdot N \cdot t_{bit} \cdot P_L \cdot K_d} \cdot L_d \tag{4.19}$$

L_d is a measure of the difference in the amount of power resonating in the ring between the *0* and *1* states. Intuitively, this would mean it is proportional to $\gamma_1 - \gamma_0$, the difference in the power lost by the ring. Substituting this into the expression for P_C in Equation 4.16, we obtain:

$$P_C = \kappa \cdot \frac{P_L}{N} \cdot \frac{C}{G_A \cdot N \cdot t_{bit} \cdot P_L \cdot K_d} \cdot L_d \cdot N_0 \tag{4.20a}$$

$$= \frac{C \cdot \kappa}{G_A \cdot N^2 \cdot t_{bit} \cdot K_d} \cdot L_d \cdot N_0 \tag{4.20b}$$

$$= K_C \cdot L_d \cdot N_0 \tag{4.20c}$$

Here, $K_C = \frac{C \cdot \kappa}{G_A \cdot N^2 \cdot t_{bit} \cdot K_d}$ and is dependent only on static constants that do not change with time, lock position, or laser power. By leveraging L_d in our calculation for P_C, we have allowed K_C to be a configurable constant that only needs to be determined once. Putting it all together, the self-heating cancellation circuitry can be implemented by two multipliers.

So far, we have only considered the case of a modulator microring in the transmitter. For completeness, we can perform the same analysis for the cancellation of the receiver, to find $\Delta\overline{P_L'}$. Here, instead of different γ_0 and γ_1, we get a change in the

laser power for *0*s and *1*s, corresponding to P_{L0} and P_{L1}, respectively:

$$\overline{P'_L} = \kappa \cdot \gamma(\lambda) \cdot [r_0 \cdot P_{L0} + r_1 \cdot P_{L1}] \tag{4.21a}$$

$$= \kappa \cdot \gamma(\lambda) \cdot [P_{L1} + r_0 \cdot (P_{L0} - P_{L1})] \tag{4.21b}$$

$$\Delta\overline{P'_L} = \kappa \cdot \gamma(\lambda) \cdot r_0 \cdot (P_{L0} - P_{L1}) \tag{4.21c}$$

In the receiver, the tracked photocurrent depends on the change in laser power:

$$i_0 = P_{L0} \cdot K_d \cdot \gamma(\lambda) \tag{4.22a}$$

$$i_1 = P_{L1} \cdot K_d \cdot \gamma(\lambda) \tag{4.22b}$$

and that $P_{L1} - P_{L0}$ can be written in terms of L_d:

$$P_{L1} - P_{L0} = \frac{C}{G_A \cdot N \cdot t_{bit} \cdot \gamma(\lambda) \cdot K_d} \cdot L_d \tag{4.23}$$

As such, the expression for P_C can be similarly derived to be:

$$P_C = \frac{C \cdot \kappa}{G_A \cdot N^2 \cdot t_{bit} \cdot K_d} \cdot L_d \cdot N_0 \tag{4.24a}$$

$$= K_C \cdot L_d \cdot N_0 \tag{4.24b}$$

where the expression for K_C is identical to that of the transmitter.

In both cases, we have only cancelled out the data-dependent portion ($\Delta\overline{P'_L}$) of self-heating power as opposed to the entirety of P'_L. We note that we can go a step further and cancel out laser self-heating altogether. If appropriately tuned, this could in theory nullify the full effects of self-heating and potentially stabilize *blue-biased* rings.

4.5 Other Considerations

4.5.1 Non-zero Rise and Fall Times on Integrated Voltage

Equation 4.2 and subsequent equations all implicitly assume that the photocurrent transitions are instantaneous. This approximation has a negligible impact for the longer integration windows, such as $N = 64$ or $N = 128$, but is significant in the case of small N and rise and fall transition times (t_R, t_F) comparable to t_{bit}. If t_R and t_F are smaller than t_{bit}, we can account for this effect by counting the bit previous to the start of the integration interval and the last bit of our integration interval through a weighting factor to add to N_0 or N_1. This approach can be generalized to ISI spanning across multiple t_{bit} by adding the contributions of even more bits.

4.5.2 Cavity Losses in Depleted and Non-Depleted States

In a modulator ring, the free-carrier absorption loss (from the modulator junctions) dominates the loss of the ring cavity. Switching the modulator into a reverse-biased state on optical *1*s depletes the junction of carriers and lowers the loss of the cavity. The fraction of optical power that couples into the drop port depends on both the drop-port coupling coefficient (fixed by the spacing of the drop-port waveguide to the ring) and the loss of the cavity; both the drop port and free carriers are ways for the cavity to lose light, when the loss due to free carriers shrinks, the proportion lost to the drop-port increases. Hence, even for the same amount of input power coupled into the ring, the drop-port photocurrent will vary based on the bias across the ring. This has the implication that L_d calculated using drop-port photocurrent is not an entirely accurate measure of the eye-height. For a red-biased ring, this will lead to an underestimation of L_d compared to the true eye-height as the fraction of power coupling to the drop-port in the depleted state (optical *1* state) is higher than in the unbiased (optical *0*) state. This effect shows up as a constant offset to L_0 and L_1 and does not affect the controller's ability to lock on. Additionally, the point of optimal L_d is still coincident with the point of maximum modulator eye-height so

long as the modulation depth is high enough to overcome the change in drop-port coupling coefficient. For our demonstrated modulator [88], the difference in drop-port photocurrent for the optical *0* (0.6 V) and *1* (−0.6 V) states for the same input power level coupled on resonance is less than 10 %. Here, the change in cavity loss impacts the tuner negligibly.

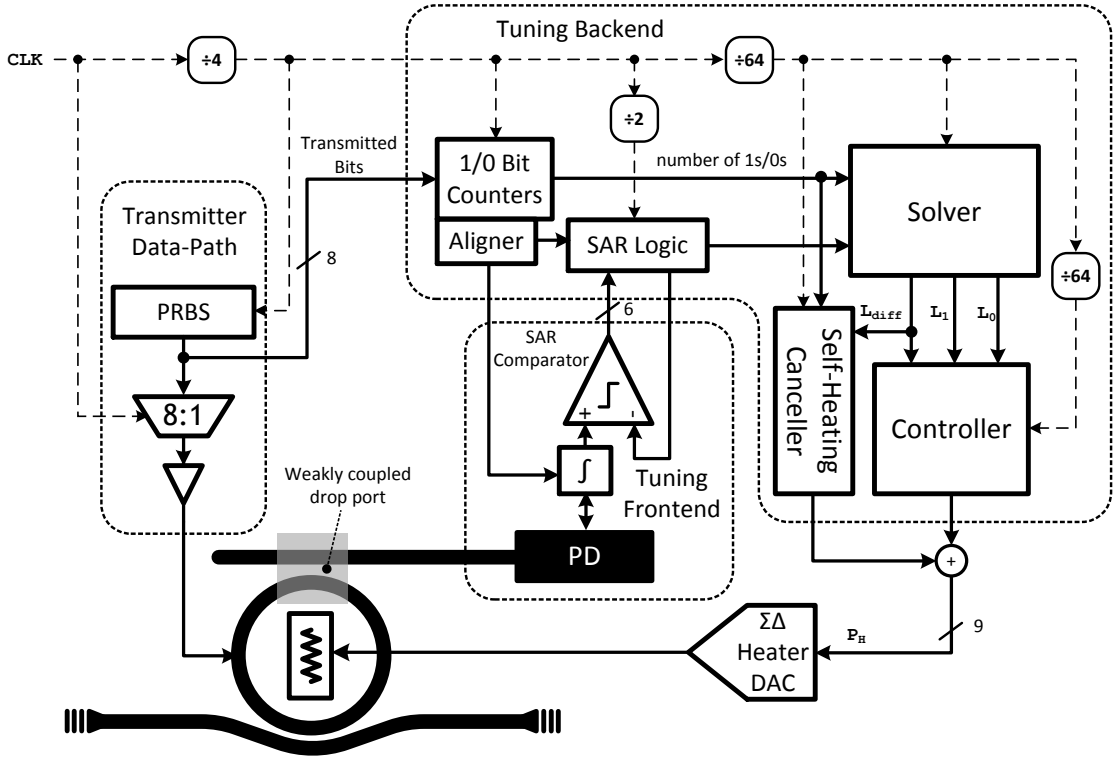

Figure 4-6: Complete bit-statistical thermal tuning circuit with tracker, solver, controller, self-heating power cancellation, shown for a transmitter-side implementation. We note that the tuning implementations operates off of various divided clocks for energy efficiency and the divided clock domains are also shown.

4.6 Summary

In this Chapter, we proposed a bit-statistical tracking method for tuning resonant microring devices, culminating in a complete thermal tuning system, shown in Figure 4-6. The bit-statistical tracker utilizes statistics of the number of *1*s and *0*s in the datastream in order to efficiently track optical signal levels. The solver uses

the technique of updating the signal levels (L_0, L_1) through an intermediate variable corresponding to the eye-height difference (L_d). This allowed the proposed tracker to become oblivious to changing 1/0 ratios, maintain correctness for long run-lengths, and require only a low bandwidth integrating frontend. No other tracker from Table 3.1 can fulfill these same requirements. We next complemented the abilities of the tracker with a new controller design, which takes advantage of the additional information that the tracker provides. In doing so, we enabled a controller that is able to automatically determine the optical lock point and to enable the selection of a monotonic error signal around this lock point. In order to eliminate transient eye closures due to the data-dependent self-heating power, we proposed a self-heating cancellation technique. This technique uses the eye-height information from the tracker (L_d) and a count of the number of zeros in the data stream (N_0) to create an appropriately large heater power response to maintain eye opening, all without any manual intervention from the user. The controller and self-heating cancellation circuit meet all the requirements outlined in Table 3.2, which no other type of controller has been able to accomplish. Combined, the proposed tracker, controller, and self-heating canceller fulfill all the requirements for an efficient, data-oblivious tracker suitable for latency-critical non-encoded links.

Chapter 5

A Monolithically-Integrated Photonics Platform in Bulk CMOS

For reasons of cost, in order for photonics to be adopted into mainstream commodity parts, such as memory, we must demonstrate its viability in a bulk process. Compared to an SOI process, a bulk process has to overcome additional challenges, such as the need for a low-loss waveguide material and bottom-side cladding for the waveguides. Because of these challenges, electro-optic transceivers and links in monolithic bulk processes have yet to be demonstrated. This chapter introduces a monolithically-integrated bulk photonics platform and presents the devices and circuits that form the components of a DWDM link, culminating in a chip-to-chip link that demonstrates the feasibility of this platform. We take an existing bulk process and enable photonics in the most CMOS-friendly way possible—all in polysilicon—while identifying the best DWDM-suitable devices and circuits that can be built effectively under these constraints.

We organize this chapter as follows, In Section 5.1, we introduce our bulk silicon-photonic platform. Sections 5.2 and 5.3 describe the transmitter, receiver, and DWDM transceiver macros. In Section 5.4, we present a data-conditioned tuning circuit that makes microring resonators robust in a hostile thermal environment. Finally, we demonstrate a chip-to-chip electro-optic link in Section 5.5.

5.1 Monolithic Photonics Platform in Bulk

The monolithic bulk platform is demonstrated in a 0.18 µm 3-metal-layer bulk CMOS process. We construct all optical devices in polysilicon, including the photodetector. Electronics are built from the power-optimized NOR flash periphery transistors native to the process, with FO4 delays of approximately 80 ps and 65 ps with 2 V and 2.5 V supplies, respectively. In contrast to standard logic processes, the platform represents a low-cost high-volume application, such as memory, where transistors are slower.

5.1.1 Process Integration

To enable photonics, we make three key modifications to the original CMOS process [49], shown in Fig. 5-1. The first is the addition of a silicon implant amorphization step for the polysilicon used in optical devices. This lowers the loss of waveguides built using process-native gate polysilicon from 40 dB/cm to 18 dB/cm [50]. Further nitride spacer optimizations bring the loss down to 10.5 dB/cm at process end-of-line [49]. The losses compare favorably to the 2 dB/mm (20 dB/cm) losses reported in [9] for waveguides built using expitaxially grown crystalline silicon in an electronic-photonic process flow.

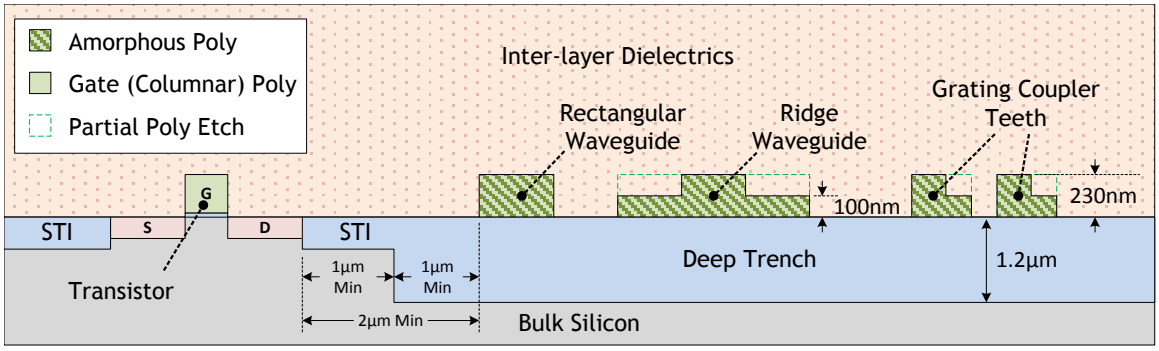

Figure 5-1: Photonics platform cross-section. The step-like shape of the polysilicon ridge waveguide and vertical coupler grating teeth are enabled via the partial polysilicon etch.

The second is the inclusion of a deep-trench isolation step (DTI), which allows us to achieve greater yield and robustness for bulk integration compared to a post-process undercut [61]. DTI creates 1.2 µm trenches filled with SiO_2 prior to polysilicon deposition. The DTI serves as bottom cladding for the polysilicon waveguides, providing optical waveguide mode confinement and isolation from the bulk silicon substrate. The challenge of a DTI process is in the polishing step to ensure uniformity of thicknesses of both DTI and shallow-trench isolation (STI, used for transistor isolation) across the wafer. However, its numerous additional benefits include enabling wafer-level optical testing, improved process development efficiency, as well as enhanced thermal isolation, since SiO_2 is much less thermally conductive than bulk silicon. Optical waveguides and devices can be placed as close as 2 µm from circuits.

The third is a partial polysilicon etch (PPE) step to break the vertical symmetry of grating couplers. The PPE enables gratings to direct light either up or down, avoiding the 3 dB ideal efficiency limit for vertically-symmetric gratings without having to embed reflectors [28] in the trenches. The two different polysilicon heights additionally enable ridge waveguide structures, convenient for forming electrically-contacted optical structures such as microring modulators.

After these process changes, transistor performance remains within process corners, allowing use of existing process-native standard cells and simulation models in the design. We note that we purposefully avoid the use of expitaxial crystallization of silicon or the introduction of germanium or silicon-germanium in our platform. Though these methods are adopted by prior art [4, 9, 43], they also introduce additional high-temperature processing and frontend process interactions that impact transistor properties. By constraining ourselves to only amorphized polysilicon and minimizing high-temperature process steps, we incorporate photonics with working transistors and improve the photonics platform's compatibility with advanced processes.

5.1.2 Microring Modulators and Polysilicon Photodetectors

The optical modulator device is a carrier-depletion microring modulator constructed using a polysilicon ridge waveguide (enabled by PPE) and doped with mid-level implants as a pn junction (Fig. 5.1.2). The ridge structure confines the optical mode to the center of the ridge, allowing electrical contacts on the sides to avoid overlap with the optical mode. The modulator ring has a radius of 7.25 µm with an FSR of 1.6 THz (9 nm). We pick this radius conservatively to make radiative tight-bend losses negligible; rings with 3 µm radii and an FSR of 3.7 THz have been demonstrated previously in the same platform [80]. Measured optical transfer characteristics of this device under different DC voltages (Fig. 5-3) shows a resonance shift of 2.7 GHz/V (15 pm/V) at 1280 nm and that weak forward-biases can be applied to increase the total shift. The ring has a Q-factor of 8000 (28.8 GHz FWHM) and an ER_i of 15 dB. Note the higher ER_i under reverse bias and the lower ER_i in forward-bias. This is indicative of slight undercoupling, as the depletion of carriers (which lowers the free-carrier loss) moves the ring closer to critical coupling.

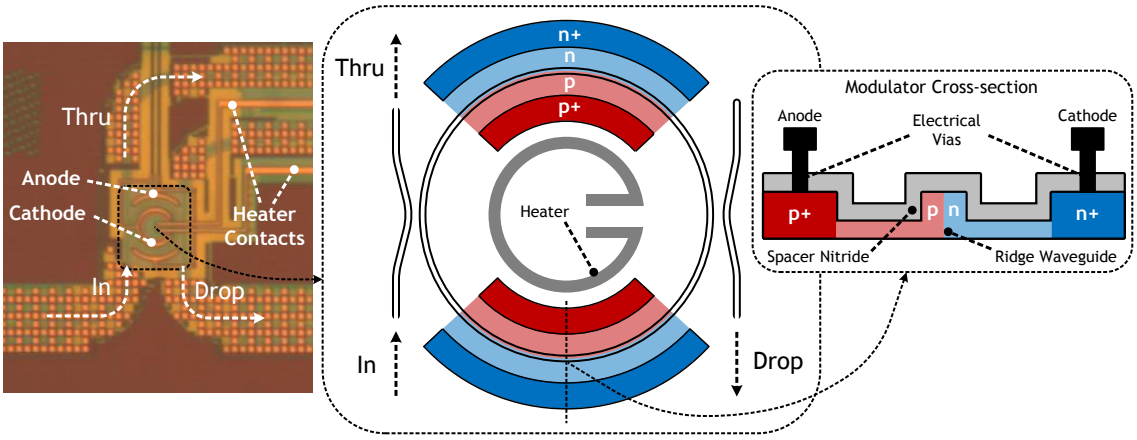

Figure 5-2: Structure of the carrier-depletion ridge waveguide modulator used in the transmitter. We estimate a modulator junction capacitance of 20 fF and a series resistance of 500 Ω, corresponding to a 15.9 GHz device RC bandwidth.

To avoid the introduction of germanium—found in all custom photonics platforms to date [3, 8, 9, 98]—we employ a completely polysilicon-based photodetector utilizing

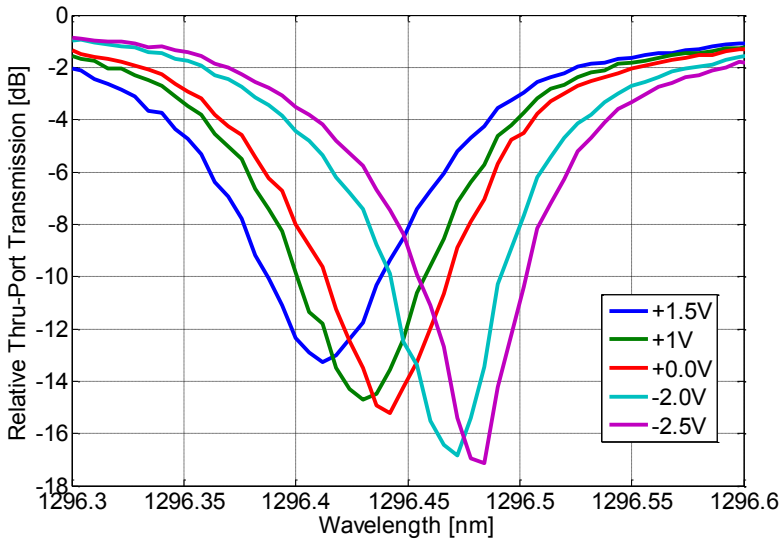

Figure 5-3: Ridge modulator thru-port transfer characteristics under different applied DC voltages.

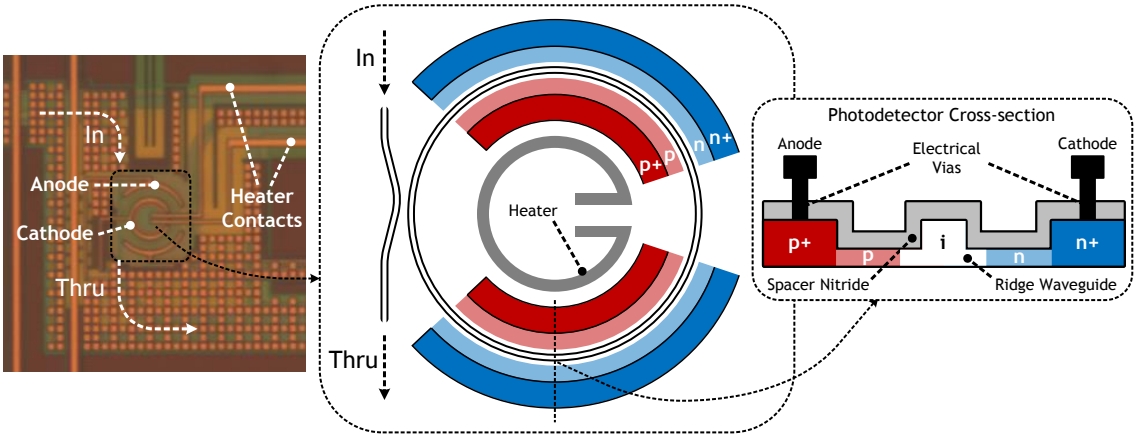

Figure 5-4: Structure of the resonant polysilicon defect-based photodetector used in the receiver. The photodetector capacitance is estimated to be 15 fF.

absorption from defect states [51]. This photodetector is a ridge-waveguide microring (with matching radius and FSR as the modulator device) doped with a p-i-n junction (Fig. 5-4). When light resonates in the ring, sub-bandgap photoabsorption stemming from defect states in the polysilicon generates free-carriers. Though this absorption is nominally weak, the resonant structure significantly enhances the effective absorption length, achieving a PD responsivity of 0.2 A/W in both the 1280 nm and 1550 nm

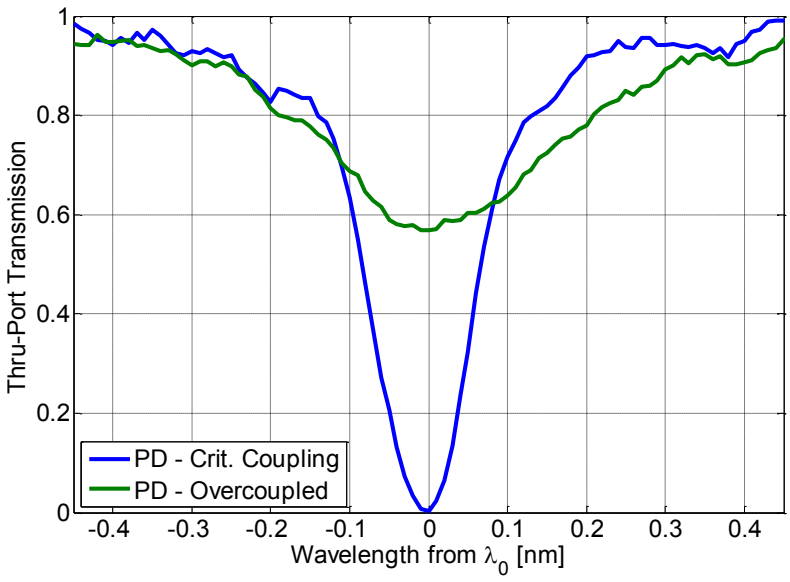

Figure 5-5: Measured thru-port transfer characteristics of the resonant polysilicon defect-based photodetector. We show the case for both a critically-coupled ring and the severely overcoupled ring (which is the version connected to the receiver circuits). Note the linear scale in the transfer characteristic.

wavelength bands. The device exhibits 3 dB photoresponse bandwidths of 1.5 GHz and 7.9 GHz at -1 V and -10 V biases, respectively. Due to much lower end-of-line waveguide losses than expected during design time, all receive macros with the best-performing receiver circuit topologies connect to PD-microrings that are severely overcoupled (Fig. 5-5). These exhibit a Q-factor of only 4000 (60 GHz FWHM) and an ER_i of 2.4 dB. This can be fixed in the layout by increasing the ring-to-waveguide gap to weaken the coupling and the gap can be set appropriately at design time once the waveguide loss is known. Critically-coupled, but otherwise identical, PD rings appearing elsewhere on the chip achieve a quality factor of more than 9000 (26 GHz FWHM) and an ER_i of 25 dB.

5.1.3 Technology Development Platform

We build the optical link components as part of a 24 mm $\times$ 24 mm technology development reticle for testing a variety of optical devices and circuits. The reticle is divided into standalone optical device regions and an array of 10 5 mm $\times$ 5 mm

transceiver chiplets, which are individually wirebonded and packaged electrically. Each chiplet hosts an array of 8 single-λ electro-optic transceiver macros (Fig. 5-6) and 3 9-λ DWDM transceiver macros for a total of 5.5 million transistors and 100 optical devices per chiplet.

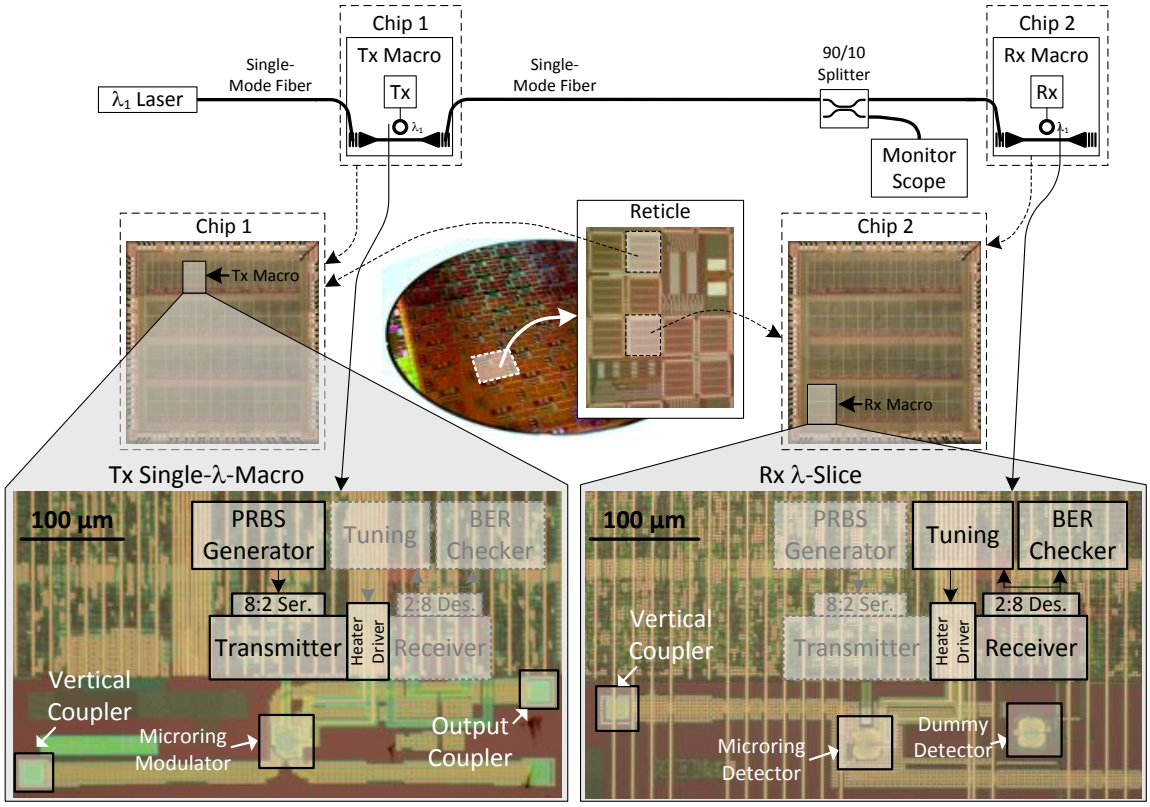

Figure 5-6: Monolithically-integrated photonics platform in bulk. Shown here is an example of a single-λ chip-to-chip optical link using a transmit macro from chip 1 and a receive macro from chip 2.

Each single-λ macro consists of a synthesized digital backend and custom high-speed transmit and receive heads that connect to the optical devices. The backend runs at one-fourth the data clock and interfaces with the custom heads through 8-to-2/2-to-8 CMOS mux/demux tree SerDes. PRBS31 generators and bit-error-rate (BER) checkers in the backend perform *in situ* characterization of the transceivers. Heater driver and tuning components drive integrated ring heaters and contain programmable logic for closed-loop wavelength-locking. Transmit and receive macros contain identical circuits and are configured into their respective roles via scan chain.

Digital Backend Test												
		Clock Frequency										
		1.5	1.6	1.7	1.8	1.9	2	2.1	2.2	2.3	2.4	2.5
External Supply Voltage	1.8	263.8	276.2	287.7	302.2	314.3	319.3	336.5	346.6	345.2	336.4	
	1.9	298.4	313.2	325.5	342.7	353.5	369.6	382.4	397.1	404.7	411.0	
	2	336.5	351.6	368.9	385.3	398.7	415.1	429.4	446.4	459.0	474.0	463.4
	2.1	377.0	394.6	412.7	431.5	447.7	465.2	479.6	499.2	513.5	529.6	536.1
	2.2	420.4	438.0	459.2	480.1	497.7	517.1	536.5	555.5	573.0	587.0	603.9
	2.3	420.4	438.0	459.2	480.1	497.7	517.1	536.5	555.5	573.0	587.0	603.9
	2.4	465.9	485.2	509.2	531.8	552.2	569.2	594.5	615.2	634.4	650.8	669.2
	2.5	566.7	589.9	618.7	645.5	670.0	693.8	721.4	746.5	769.6	787.8	814.7

Receiver Self-Test Loopback												
		Clock Frequency										
		1.5	1.6	1.7	1.8	1.9	2	2.1	2.2	2.3	2.4	2.5
External Supply Voltage	1.8	4.90	5.05	5.18	5.66	5.71						
	1.9	5.81	6.03	6.18	6.34	6.31	6.73					
	2	6.59	7.04	7.19	7.40	7.56	7.76	7.90	8.17			
	2.1	7.73	8.21	8.36	8.58	8.78	9.02	9.21	9.38	9.71		
	2.2	9.20	9.45	9.86	10.34	10.42	10.86	10.94	11.10	11.10	11.04	
	2.3	10.39	10.62	11.24	11.29	11.69	11.75	12.24	12.39	12.77	13.26	
	2.4	12.23	12.21	12.64	13.38	13.53	14.08	14.42	14.77	14.77	14.94	
	2.5	13.52	14.06	14.43	14.52	14.74	15.02	15.74	15.95	16.28	16.54	

(a) Chip on a wafer without process changes

Digital Backend Test												
		Clock Frequency										
		1.5	1.6	1.7	1.8	1.9	2	2.1	2.2	2.3	2.4	2.5
External Supply Voltage	1.8	240.1	251.7	257.9	269.9	281.2	293.1	305.0	316.3	317.3		
	1.9	269.2	284.7	297.2	312.6	323.9	331.3	346.1	358.8	368.8	380.1	
	2	304.0	319.9	335.9	352.0	365.4	381.2	395.5	412.1	416.0	428.4	414.2
	2.1	340.7	358.1	375.9	393.8	408.4	426.6	441.3	460.0	472.6	488.9	485.1
	2.2	379.6	398.2	418.0	438.1	454.0	474.3	492.0	511.2	526.6	544.6	547.1
	2.3	421.1	441.3	463.1	485.1	504.4	525.1	545.8	565.5	584.3	601.1	616.1
	2.4	464.3	485.4	510.8	534.9	557.0	577.8	601.5	623.7	644.5	662.5	683.8
	2.5	510.7	533.6	561.6	588.4	612.0	634.8	660.9	685.7	708.7	726.6	753.2

Receiver Self-Test Loopback												
		Clock Frequency										
		1.5	1.6	1.7	1.8	1.9	2	2.1	2.2	2.3	2.4	2.5
External Supply Voltage	1.8	5.74	5.90									
	1.9	6.37	6.50	6.69	7.15	7.07						
	2	7.56	7.76	7.96	8.15	8.26	8.73	8.73	8.89	8.90		
	2.1	8.63	8.84	9.17	9.35	9.51	9.77	9.89	10.34	10.34		
	2.2	10.04	10.45	10.61	10.67	10.83	11.05	11.37	11.66	11.92		
	2.3	11.34	11.77	11.64	12.06	12.43	12.49	12.84	13.10	13.49	13.76	
	2.4	12.57	12.88	13.39	13.74	14.10	14.49	14.39	14.66	14.90	15.49	
	2.5	14.52	14.81	15.26	15.31	15.79	15.85	16.20	16.85	17.17	17.53	

(b) Chip on a wafer with process changes

Figure 5-7: Measured voltage and frequency shmoo plot of chips with (a) and without (b) the introduced process changes. A red box indicates a test failure. Power consumption of the digital backend at each voltage/frequency point is shown in each box of the backend test. The combined transmitter and receiver power consumption are shown in the boxes for the receiver loopback test.

Optical waveguides, couplers, and active devices are instantiated in rows alongside the circuits. Single-mode fibers (with a cleaved tip) are positioned over the VGCs (using probe positioners or mounted to the package) to couple light in or out of on-chip waveguides. The 9-λ DWDM transceiver macros are similar to 9 adjacent single-λ macros, but string together all 9 modulator- or receive-microrings on a single bus waveguide to provide DWDM functionality.

To verify that the additional process steps did not significantly alter the electrical performance of the original flash periphery process, we perform an electrical self-test shmoo characterization, shown in Figure 5-7. We compare between chips on wafers that go through only the electrical processing steps and chips on a wafer that go through both the electrical processing steps and the additional changes to enable photonic devices. The digital test measures the functionality of counters in the digital backend, ensuring that they count properly for a given clock frequency. The receiver loopback test tests the full path from the transmit PRBS generators, through the serializer, to the receiver, through the deserializer and into the bit-error-rate measurement modules. The receiver is placed in self-test mode, with the receiver input driven by output of the serializer through a diode emulation circuit in the receiver. The receiver passes the test if there are no bit-errors for 4×10^9 bits. The differences in performance between the two cases are minimal, well within the standard process corners. We note that for up to a 2.3 V external supply, the internal chip voltage remains below 2 V. The voltage droop stems from insufficient power grid density due to the availability of only 3 metal layers the process.

5.2 Microring-based Optical Transmitter

The transmitter consists of a depletion-ridge modulator device driven by a transmitter frontend circuit, shown in Fig. 5-8. The frontend circuit consists of a 2-to-1 DDR serializer followed by an inverter-based push-pull driver. The NMOS pull-up transistor in the final driver stage on the anode terminal is used to limit the applied forward-bias voltage. On logic 1s, the transmitter circuit applies a voltage of

-VDD to the modulator junction, depleting the junction of carriers and red-shifting the resonance. On logic 0s, the circuit weakly forward-biases the device to a voltage of $V_{REF} - V_{Tn}$ to inject carriers into the junction and blue-shift the resonance.

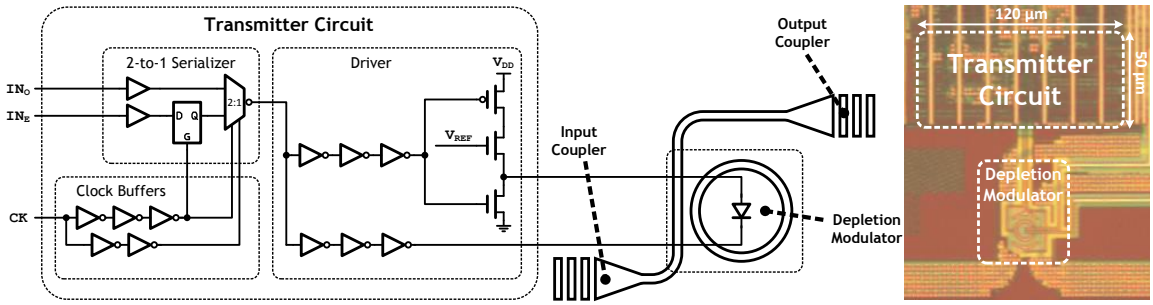

Figure 5-8: Carrier-depletion-based transmitter. The output impedance of the driver is matched for both the pull-up and pull-down at approximately 700 Ω.

The choice of where to bias the laser wavelength relative to the resonance is a degree of freedom for a ring modulator. We define an eye-height metric, ΔT, as the difference in optical powers for modulated logic 1 and logic 0 levels, normalized by modulator input optical power:

$$\Delta T = |T_1 - T_0| = \left|10^{-IL/10} - 10^{-(IL+ER)/10}\right| \tag{5.1}$$

where IL is the modulator insertion loss and ER is the modulator extinction ratio $(10 \cdot log_{10}\frac{T_1}{T_0})$, both given in dB. Using this metric, we perform an optimization to find the maximum eye-height that this device can support, shown in Fig. 5-9. Note that the optimal eye-height is not located at the same wavelength as the one that gives the greatest ER; though T_0 is supressed far below T_1 at this point, the noticeably higher IL degrades the level of T_1. Conversely, bias points far away from the resonance have low IL but provides too little ER to create sufficient modulation depth. The optimum occurs at a point where IL and ER are balanced against each other, e.g. $IL = 2.0$ dB and $ER = 6.6$ dB for the shown modulator, corresponding to $\Delta T = 0.49$.

We measure the modulator under two sets of modulator operating voltages, *case 1* and *case 2*, achieving open-eye data-rates of 4 Gb/s and 5 Gb/s, respectively, with a PRBS31 sequence (Fig. 5-10). The maximum data-rate of the transmitter is limited

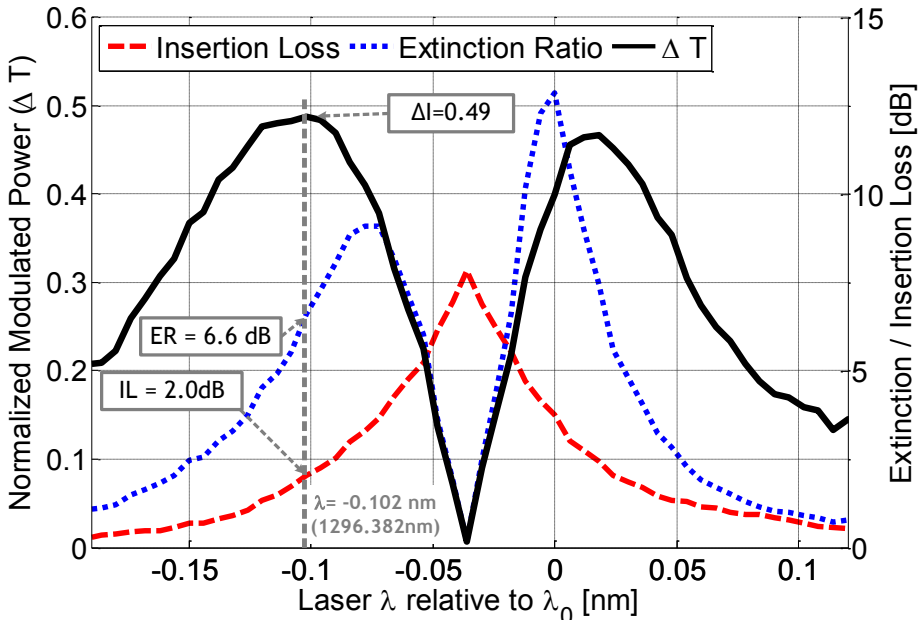

Figure 5-9: Trade-off analysis between extinction ratio, insertion loss, and the eye-height metric ΔT versus the laser wavelength. We use a a drive swing of $-2.5\,\mathrm{V}$ to $1.5\,\mathrm{V}$ and the device in Fig. 5-3. Note that the polarity of the modulator flips when T_1 is smaller than T_0 to keep ΔT positive.

by transistor performance; At $5\,\mathrm{Gb/s}$, the higher voltage in *case 2* is necessary to maintain correct digital functionality in the 2-to-1 serializer and sufficient overdrive on the pull-up NMOS transistor for a fast edge. Photon lifetime effects from the microring linewidth itself (Q = 8000, $28.8\,\mathrm{GHz}$ optical bandwidth) are negligible at these data-rates. Note that the experimentally measured ER and IL are better than what is expected from the DC wavelength scans. This is because under sufficiently large forward bias (blue-shift), self-heating from diode on-current acts to red-shift the ring slightly back, making the shift captured by the DC measurement appear smaller. As data-transmission is at a much higher rate than the self-heating time constant, the modulated eye captures the true extinction of the device. The measured energy-per-bit of the modulator across data-rates is shown in Fig. 5-11. The clock buffers driving the serializer dominate total power due to aggressive sizing for 3x FO4 target bit-times and phase-matching of the inverted phase for the 2-to-1 CMOS DDR multiplexer. At the V_{REF} used in *case 2*, the driver applies a sufficiently large

forward-bias voltage to weakly carrier-inject the modulator diode, drawing additional static current on logical 1s. The static current is amortized at higher data-rates, hence the small efficiency improvement with data-rate. *Case 1* pushes the device into weak carrier-injection to a smaller extent due to lower V_{REF}. The driver energy efficiency is 200 fJ/bit at 4 Gb/s and 350 fJ/bit at 5 Gb/s for *case 1* and *case 2*, respectively, with an overall energy-efficiency of 0.7 pJ/b and 1.16 pJ/b for the full transmit circuit, which includes the clock buffers and 2-to-1 serializer. Photon lifetime effects from the microring itself (28.8 GHz optical bandwidth) are negligible at these data-rates. For the macro floorplan to be compatible across all types of optical devices on the platform, we had to place the modulator device 60 µm away from the frontend circuit, adding wiring capacitances of 13 fF and 14 fF (from layout extraction) on the anode and cathode nodes, respectively. We estimate that an optimized device-to-circuit placement using the minimum 2 µm spacing will reduce wiring capacitances down to sub-2 fF and bring the total load on the modulator driver circuit down to 22 fF. Avoiding speed-limitations due to other components in the driver chain, this should increase the achievable data rate to 8 Gb/s while lowering the energy cost even further.

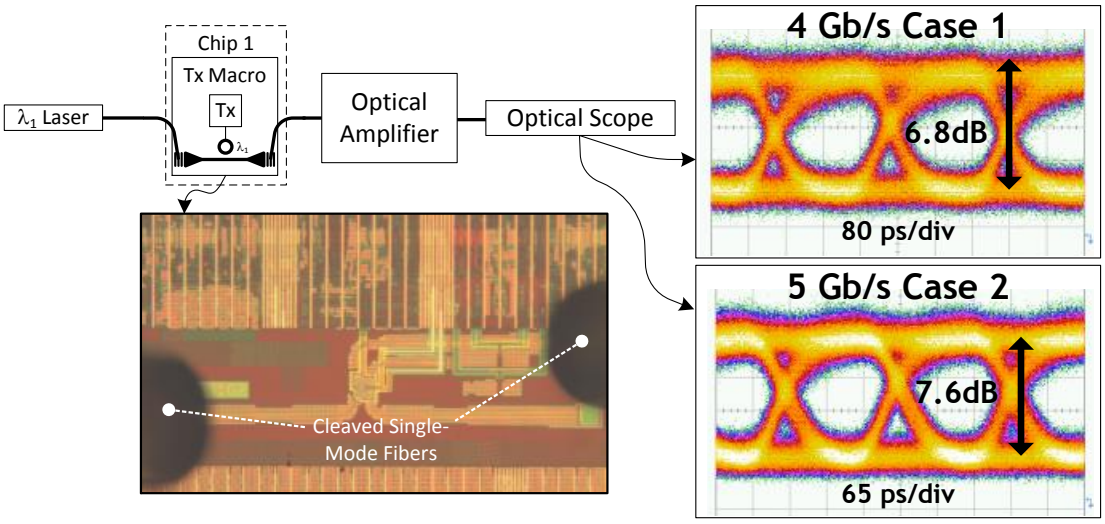

Figure 5-10: Transmitter characterization setup. The eye-diagrams are shown for *case 1* at 4 Gb/s and *case 2* at 5 Gb/s.

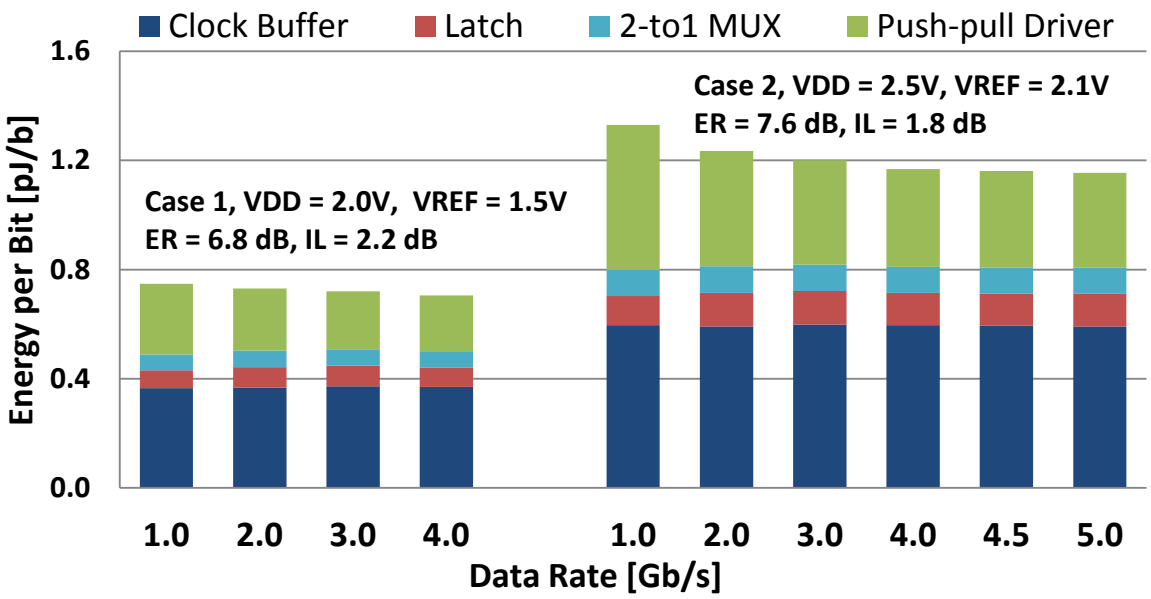

Figure 5-11: Measured transmitter energy-per-bit. The component-level power breakdowns are deembedded through simulation using extracted netlists that include the wiring capacitances on the modulator anode and cathode (extracted to be 13 fF and 14 fF, respectively) and a 20 fF modulator junction capacitance.

We build the transmitter circuit as part of a 9-λ DWDM transmit macro, shown in Fig. 5-13. The rings are spread across the 9 nm FSR and we step the radii to achieve a nominal 1 nm channel-to-channel spacing (Fig. 5-12), achieving more than 20 dB of cross-talk isolation between adjacent channels. Across the 3.5 mm span of the rings, no ring experiences more than 0.5 nm deviation from its nominal resonance and local variations are not large enough to flip the ordering of adjacent channels on any measured chip or wafer. We use the tuning circuits in the backend to drive integrated microring heaters to move rings back to the grid. We verify that the macro is capable of an aggregate 45 Gb/s of data transmission through a single waveguide or fiber by individually capturing an open transmit eye on each of the 9 λ-slices at 5 Gb/s (Fig. 5-13). The spectral efficiency of the macro, given a 1.6 THz FSR, is 0.028 bits/s/Hz. The I/O density of the DWDM macro is approximately 110 Gb/s/mm^2, including all transmitter circuits (120 µm $\times$ 50 µm per slice) and a conservative (3500 µm $\times$ 100 µm) trench for the photonic devices. Note that the ring-to-ring placement pitch of 384 µm and the overall size of DWDM macro (4000 µm $\times$ 900 µm) is limited by the area of the

digital testing backend; the pitch can be as tight as 40 µm without violating design rules.

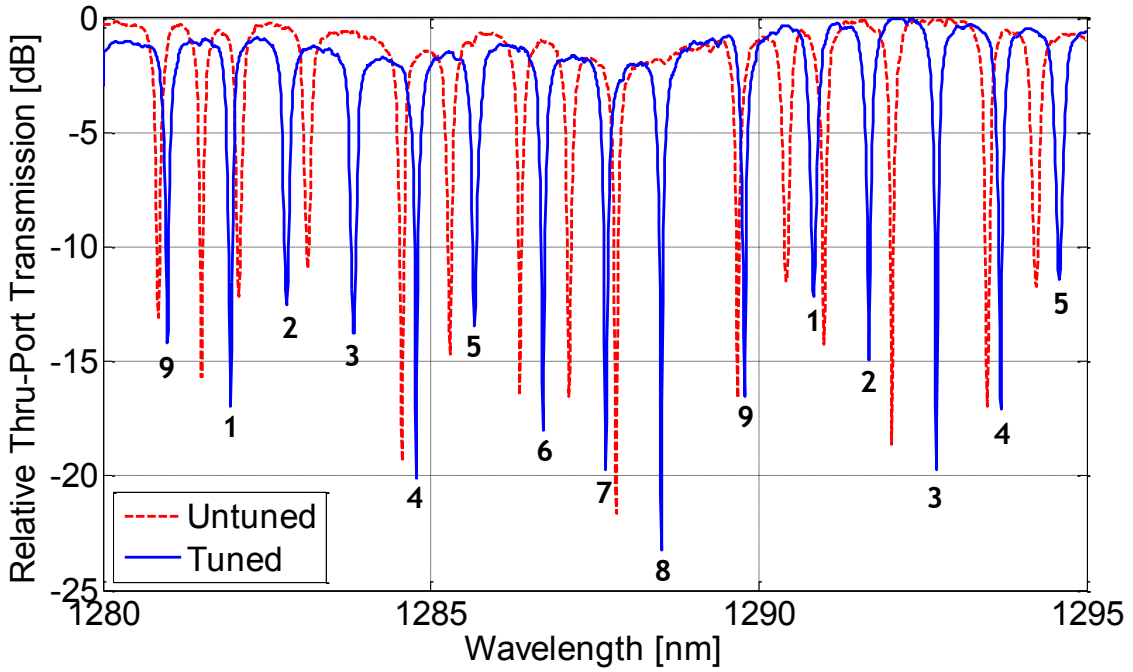

Figure 5-12: Optical thru-port transfer characteristic of the 9-λ DWDM transmit bank. The tuned spectra is after coarse tuning to a 1 nm grid. Each resonance is numbered with its corresponding ring.

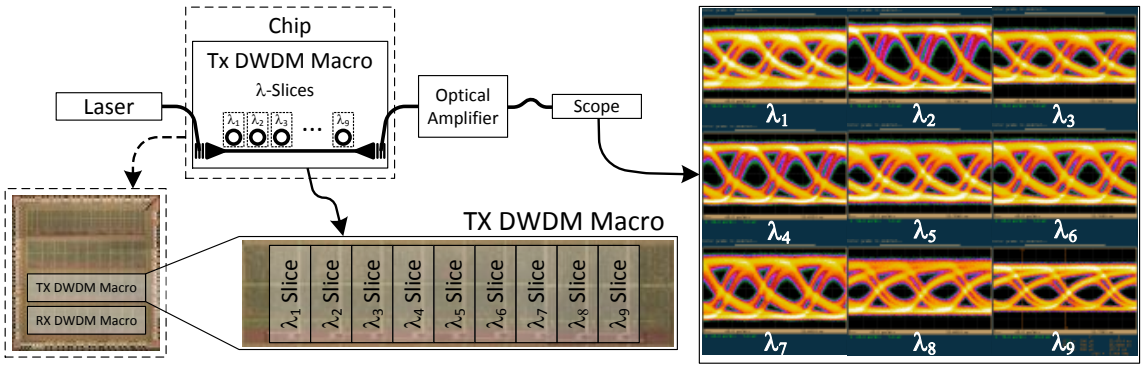

Figure 5-13: Demonstration of the 9-λ DWDM transmit macro. The double-edge transition in the eyes is due to the previous edge not fully settling after one bit-time.

5.3 Polysilicon-based Resonant Optical Receiver

The receiver block is composed of a ring resonant defect-detector connected to a receiver circuit, shown in Fig. 5-14. To mitigate the slow speed of the process, we adopt a split-diode technique [20]; each PD microring is separated into two electrically-isolated PD-halves (PD-0, PD-1), each connected to a receiver-half (RX-0/1) running at half-rate on opposite clock phases. In effect, each receiver-half gets 1/2 of the total photocurrent (I_{PD}), but is given twice the evaluation time. Each receiver-half circuit consists of an inverter-based TIA followed by a clocked sense-amplifier and RS latch. The TIA transimpedance gain can be adjusted by configuring the feedback resistor to be 12 kΩ, 4 kΩ, or 12 kΩ || 4 kΩ. Current and capacitive DACs attached to the sense amplifier provide offset compensation and eye-measurement capability for the receiver. We design the circuit to accommodate a potentially large dark current range across all PD variants on the platform reticle; the dummy PDs and TIAs serve as dark current references, keeping the sense-amps balanced for large dark currents. Under extremely high dark currents, the current DAC at the input of each TIA cancels dark currents to keep the TIA biased in a linear regime. We note, however, that there is less than 50 pA of dark current at a −10 V bias for the defect microring PD [51]. Hence, the dark currents are low enough for the dark current cancellation DAC and dummies to be removed in a receiver tailored specifically for this device.

A sense-amp undergoes two phases during evaluation time: a linear integration phase in which v_{in} is integrated onto the cap through the g_m of the input transistor and an exponential regeneration phase where the positive feedback of the cross-coupled inverters drives v_{out} to the supply rails. The behavior of the sense-amp can be modeled, to first-order, as:

$$v_{out} = v_{in}\frac{g_m}{C_{int}} \cdot t_{int} \cdot \cdot e^{\frac{t_{regen}}{\tau_{inv}}} \tag{5.2}$$

where $t_{int}+t_{regen} = t_{eval}$ The onset of regeneration (and the end t_{int}) is triggered when the voltage at the drain of the input transistor drops sufficiently low to trigger a loop-gain > 1 in the cross-coupled inverters. In traditional sense-amps, this is set by the sizing of the footer and the common mode of the input transistor and does not change

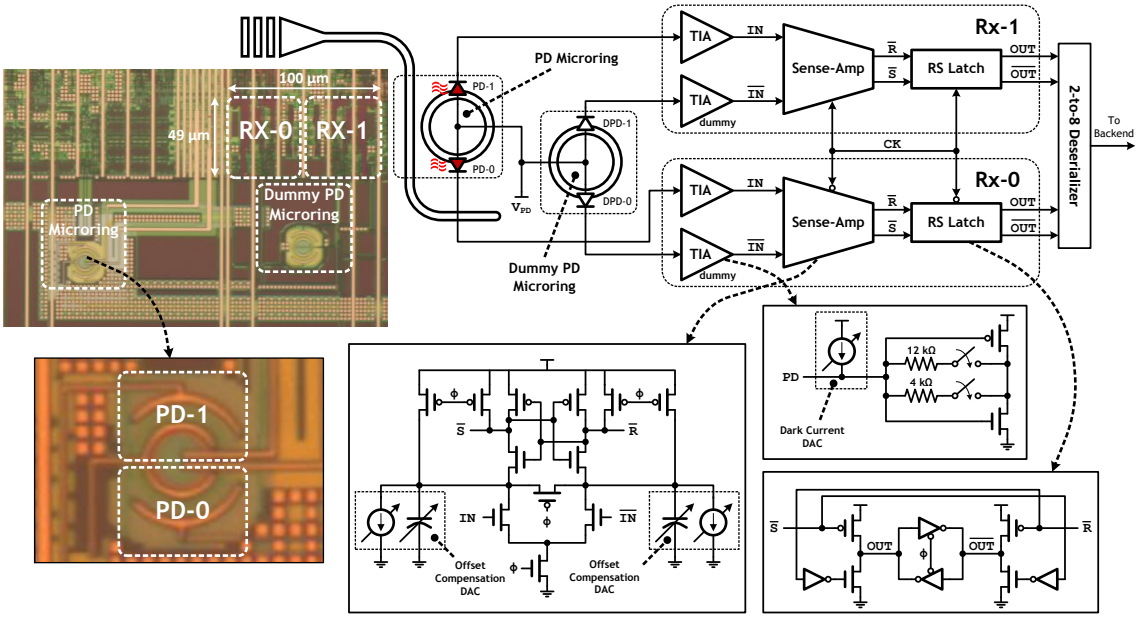

Figure 5-14: Optical micrograph and schematic of the receiver architecture

automatically when t_{eval} shrinks with data-rate. Accordingly, a shorter evaluation time will directly squeeze t_{regen}, degrading the sense-amp sensitivity exponentially once t_{regen} is no longer sufficient to regenerate rail-to-rail. Like-wise, a t_{int} that is too short results in a t_{regen} that is more than sufficient to fully regenerate rail-to-rail, wasting time that could have been used to integrate a larger input. To minimize sense-amp sensitivity v_{in}, the t_{int} should be traded off with t_{regen} to balance linear growth with exponential growth. We perform this optimization in the receiver by switching on both sides of the current compensation DAC simultaneously. We size the receiver for relatively long t_{int} for better sensitivity at low data-rates and switch on the current DACs to reduce t_{int} and boost t_{regen} for higher rates as t_{eval} shrinks. This comes at a small cost in static power, but allows for an optimal split between t_{int} and t_{regen} across a wide data-rate range.

We measure the receiver in a high-performance mode ($V_{DD} = 2.5\,\text{V}$) and a low-power mode ($V_{DD} = 2.0\,\text{V}$), biasing the PD at $V_{PD} = -10\,\text{V}$ for all experiments (Fig. 5-15). The receiver runs with no bit-errors up to 5 Gb/s and 3 Gb/s for 2×10^{12} bits (10^{12} bits on each receiver-half) in the high-performance and low-

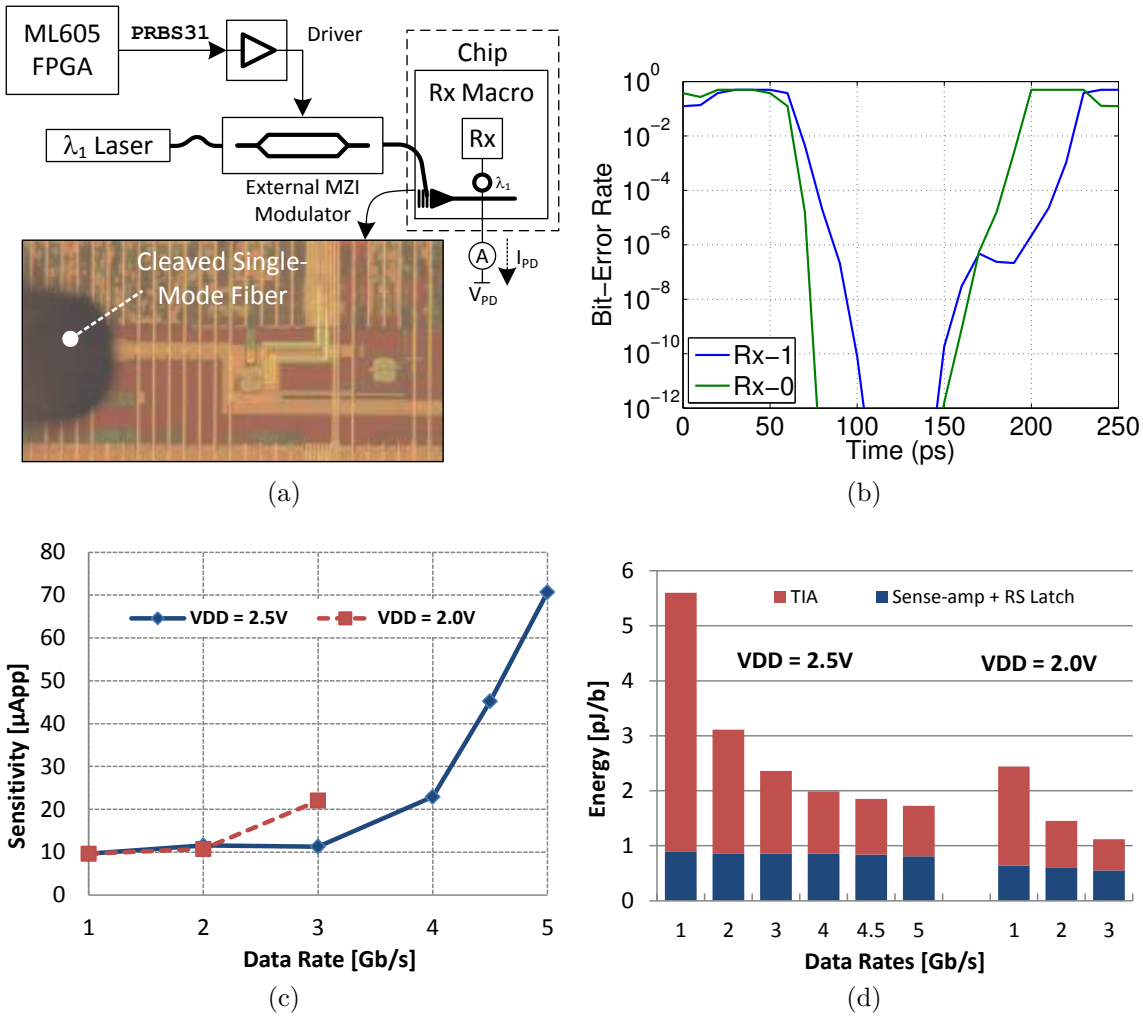

Figure 5-15: Receiver measurement test setup (a), optical BER measurement at 5 Gb/s (b), measured optical sensitivity (c), and measured receiver energy-per-bit (d). Sensitivity is defined using the total peak-to-peak photocurrent (both halves) necessary for error-free operation after 2×10^{12} total bits. The 5 Gb/s receiver optical BER measurement is performed using a laser power that results in $I_{PD} = 71\,\mu A_{pp}$. The bottom of the curve indicates no bit-errors for 2×10^{12} total bits. All measurements are performed *in situ* using the on-chip digital backend.

power modes, respectively. Beyond 5 Gb/s, the sense-amp is too slow to maintain correct functionality. The high-performance mode achieves a 10^{-12} BER sensitivity of $12\,\mu A_{pp}$ (-15.2 dBm with a 0.2 A/W PD) for 1–3 Gb/s. At 5 Gb/s, sensitivity degrades to $126\,\mu A_{pp}$ and $71\,\mu A_{pp}$ (-7.5 dBm) before and after the sense-amp t_{int} optimization, respectively. The low-power mode achieves comparable sensitivity at half the power up to 3 Gb/s, when the t_{eval}-limited sense-amp begins degrading sensitivity.

A layout error resulted in the placement location of the dummy PD microring and the active PD microring to be swapped, causing the active PD to be placed 150 µm farther from the receiver than the dummy PD. The wiring capacitances at the active PD input nodes are 32 fF for Rx-1 and 21 fF for Rx-0, compared to 11 fF for Rx-1 and 13 fF for Rx-0 at the dummy PDs, estimated using layout extraction. This error lowers the TIA bandwidth by 33 % from the original, causing the 12 kΩ feedback resistor to have insufficient bandwidth to support data-rates beyond 3 Gb/s. The reduction of the TIA feedback resistance to 4 kΩ, coupled with the smaller sense-amp t_{eval}, results in the degradation of sensitivity past 3 Gb/s. Power consumption of the full receiver is dominated by the static TIA power, which is amortized at higher rates. Elimination of the dummy dark-current matching TIAs in future designs will halve the TIA power component. The receive macro follows the same floorplan as that of the transmit macro and can likewise be optimized to move devices closer to circuits to lower wiring capacitances.

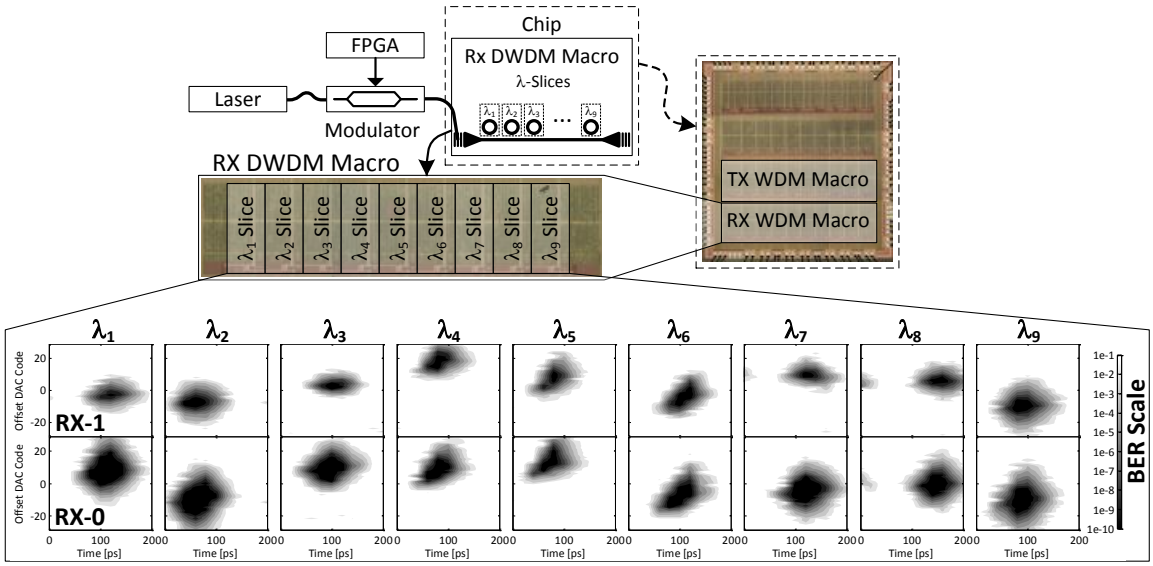

Figure 5-16: Demonstration of the 9-λ DWDM receiver macro. The receive eye diagrams are generated by sweeping the receiver threshold using the offset compensation DACs.

We verify receivers integrated into a 9-λ DWDM receiver macro by individually aligning the laser to each λ-slice and measuring the BER (Fig. 5-16). We record error-

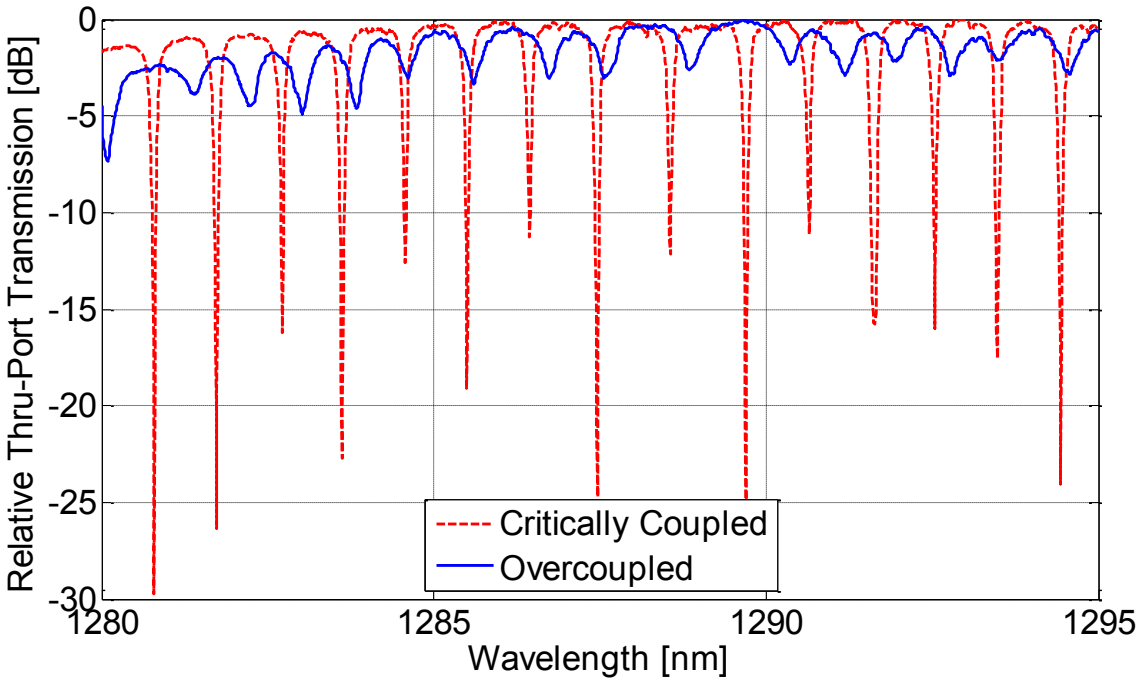

Figure 5-17: Optical transfer characteristics of the measured DWDM receiver macro, which has overcoupled PD rings, and a DWDM receiver macro with critically-coupled PD rings. The rings are coarsely tuned to a 1 nm grid.

free receiver eye openings for 2×10^{10} bits on each slice at 5 Gb/s, demonstrating that the macro is capable of 45 Gb/s aggregate receive bandwidth per waveguide or fiber (114 Gb/s/mm^2). Figure 5-17 shows the macro's optical spectrum. The macro is FSR-matched to the DWDM transmit macro with the same channel spacing (1 nm). PDs in the tested DWDM macro are also severely overcoupled with $ER_i < 3$ dB (Q = 4000, 15 dB crosstalk isolation at 1 nm spacing). A separate DWDM receiver macro with critically-coupled PD rings exists elsewhere on the same chip (Q = 9000, 23 dB crosstalk isolation at 1 nm spacing), though it is connected to a different kind of receiver.

5.4 Ring Resonator Wavelength Locking

In this section, we demonstrate an on-chip wavelength-locking circuit that maintains the receiver eye opening under changing temperatures. The synthesized receive-

side tuning sub-system (Fig. 5-18) contains an optical power meter circuit, a configurable data path, a programmable lookup-table (LUT), and a $\Delta\Sigma$-DAC circuit [80] that drives an integrated microring heater.

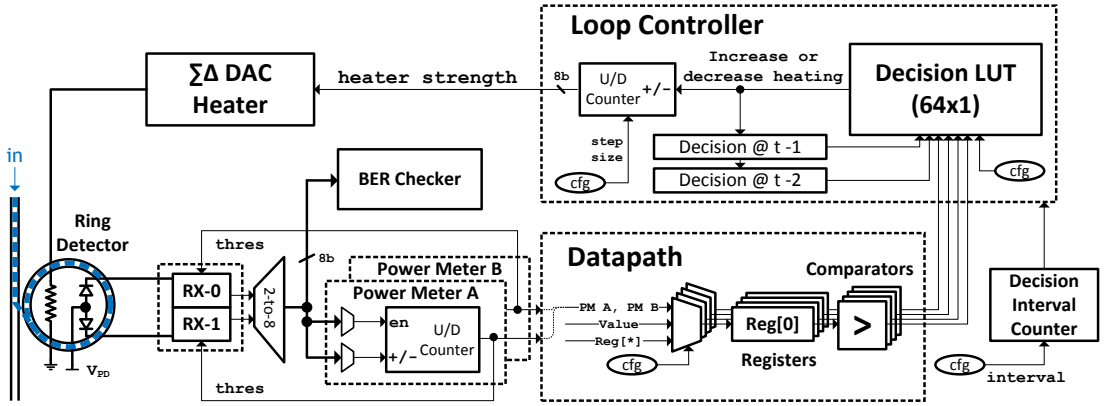

Figure 5-18: Architecture of the programmable tuning sub-system.

5.4.1 Heater Driver System

The open-loop heater driver system consists of a heater driver circuit and a a microring resonator with an integrated heater. When current flows through the heater resistor, the power dissipated by the heater heats up the ring and moves the ring resonance. The driver circuit follows an all-digital $\Delta\Sigma$ architecture and consists of a digital accumulator and a NMOS driver head. Given a binary-encoded digital input, D_H, corresponding to the desired heater output level, the accumulator produces a pulse-density modulated (PDM) waveform. This waveform is a digital pulse train with a duty-cycle corresponding to D_H. This waveform drives the gate of the NMOS transistor in the driver head. When the waveform is *1*, the transistor turns on, allowing current to flow through the heater from the supply, dissipating heat. When *0*, the transistor shuts off and no current flows through the heater.

We rely upon the the slow thermal response to act as a low-pass filter, smoothing the ripples in temperature (and ring transmission) introduced by the binary heater drive waveform. This holds true provided that the driver's clock frequency (and hence

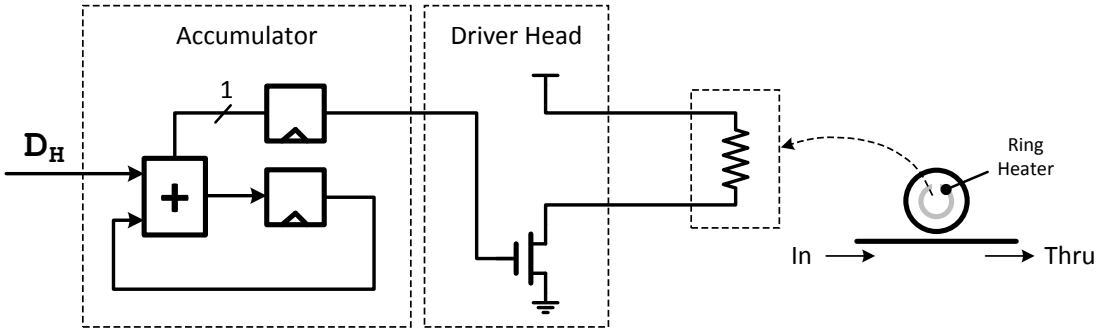

Figure 5-19: Heater driver architecture, which drives an integrated ring heater. The accumulator keeps track of the error from the input, producing a *1* in the highest carry-out bit and wrapping around back to zero when the count overflow.

the PDM's oversampling ratio) is set significantly higher than the ring's thermal time constant. Since power consumption of the accumulator scales with the clock frequency, driver power can be optimized by fine-tuning the oversampling ratio of the PDM. Working together, the system forms a digital to analog converter (DAC), converting a digital code, D_H to a change in temperature, ΔT. We note that the duty-cycled nature of the driver guarantees a monotonic and linear relationship between heater power and input setting.

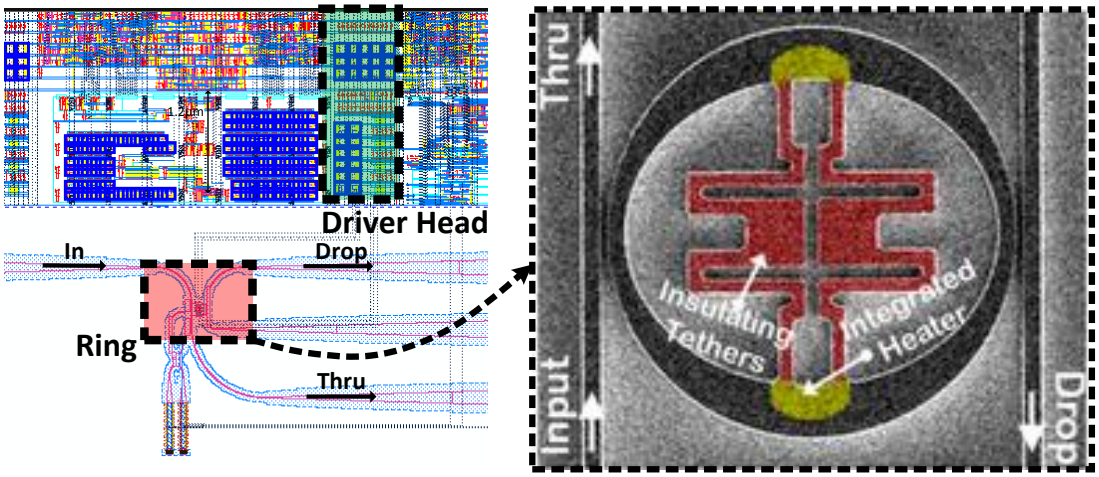

Figure 5-20: Layout view of the heater driver head and the connected adiabatic microring resonator filter implemented on the bulk platform.

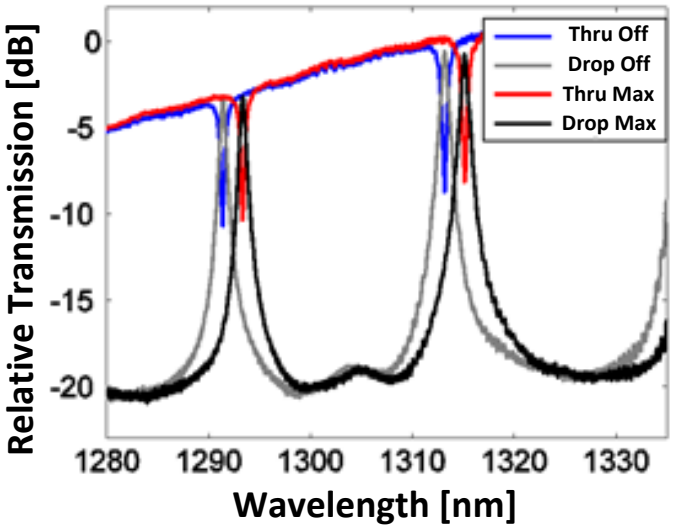

Figure 5-21: Optical transmission vs. wavelength when driver is off and at max power and total power consumed at various driver power settings. For reference, *off* corresponds to a driver strength setting of 0, *max* corresponds to a strength setting of 255.

To evaluate the tuning efficiency of integrated tuners on the deep-trench bulk platform, we implement an 8-bit design of the heater driver system driving two different types of rings. On this platform, the 8-bit accumulator, and driver head occupy 28 µm^2, 1500 µm^2, and 2500 µm^2 of area, respectively, including area used for decoupling capacitances and fill cells. The implementation shown here is aggressively pipelined to hit frequency targets in excess of 400 MHz and sized to drive more than 5 mA of current. Design constraints can be relaxed to lower driver power and area.

The first type of ring we drive is a 6 µm diameter adiabatic microring resonator (ARM), similar to the one in [91], fabricated to work at $\approx$ 1290 nm, shown in Figure 5-20. The integrated heater is formed by n type doping in the adiabatic region with $n+$ and silicided silicon tethers to minimize contact resistance and to insulate the tethers. The fabricated resistance is $\approx 1.7\,\text{k}\Omega$. Single radial mode propagation is preserved through the adiabatic region, achieving an uncorrupted free spectral range (FSR) of 3.7 THz (Figure 5-21). Here, the heater output power is 3.5 mW at the max power

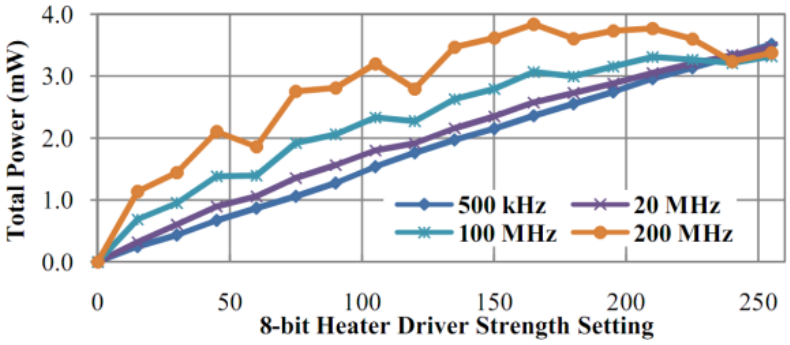

Figure 5-22: Total tuning power, which includes the power dissipated by the heater and the power overhead of the driver, vs. heater power setting at various driver operating frequencies.

setting, limited by supply voltage (2.5 V) and heater resistance (1.7 kΩ). We achieve a maximum resonance shift of 1.9 nm (350 GHz). At laser power levels low enough for self-heating to be negligible, this corresponds to an LSB step size of 7.4 pm (1.37 GHz). Compared to the power delivered to the ARM's heater, the power overhead for the driver circuit is negligible at a clock frequency of 500 kHz but substantial at 200 MHz (Figure 5-22). The tuning efficiency of the ring is 0.55 nm/mW (10 µW/GHz).

The second type of the ring we drive are the polysilicon-based microrings used in the receiver, for which we will use to demonstrate receiver wavelength locking. Here, the ring uses a 1 kΩ integrated heater. With a 2.5 V supply, the heater driver outputs 6 mW of maximum power and achieves 0.5 nm (90 GHz) of tuning range. Each LSB step size corresponds to 2 pm (0.35 GHz). The tuning range will increase by using a lower heater resistance or by improving the tuning efficiency, which we estimate to be 0.082 nm/mW (67 µW/GHz).

The ΣΔ converter in the heater driver produces a quantization noise power spectral density of:

$$|H_{\Sigma\Delta}(\omega)|^2 = \frac{\Delta^2}{12} \cdot \frac{1}{2\pi f_S} \cdot 4 \left| \sin\left(\frac{\omega}{2f_S}\right) \right|^2 \tag{5.3}$$

where Δ is the quantization level and f_S is the sampling frequency of the heater driver. We can calculate a variance on λ_0 of the ring that the heater driver causes due to the quantization noise and the oversampling ratios, with thermal effects factored in:

$$\sigma^2 = \frac{1}{2\pi}\int_{-\infty}^{\infty} |H_{\Sigma\Delta}(\omega)|^2 \cdot |H'_T(\omega)|^2 \, d\omega \tag{5.4a}$$

$$= \frac{1}{2\pi}\int_{-\infty}^{\infty} |H_{\Sigma\Delta}(\omega)|^2 \cdot \left(\frac{1}{1-G_L(0)}\right)^2 \cdot \frac{1}{1+\omega^2\left(\frac{\tau_T}{1-G_L(0)}\right)^2} d\omega \tag{5.4b}$$

In Table 5.1, we use Equation 5.4 to calculate the variance in wavelength at each of the heater driver frequencies for both types of rings. From here, we see that operating the heater driver at a 20 MHz rate provides a minimal power overhead while only adding a 3σ noise that is about 1 LSB. Note that these calculations are calculated with $G_L = 0$ (no self-heating). However, the evaluated values do not change much for different G_L as G_L changes both the pole location and the magnitude of $S(\omega)$ and the effects on σ^2 mostly cancels out.

	ARM $\Delta = 1.9\,\text{nm}$, $\tau_T = 10\,\mu\text{s}$		Polysilicon Detector Ring $\Delta = 0.5\,\text{nm}$, $\tau_T = 10\,\mu\text{s}$	
	σ	$10\sigma/LSB$	σ	$10\sigma/LSB$
500 KHz	104 pm	140.5	27.5 pm	140.5
20 MHz	2.7 pm	3.65	0.72 pm	3.65
50 MHz	1.1 pm	1.49	0.29 pm	1.49
100 MHz	0.55 pm	0.74	0.14 pm	0.74
200 MHz	0.27 pm	0.36	0.07 pm	0.36
312.5 MHz	0.18 pm	0.24	0.05 pm	0.24

Table 5.1: Variance in wavelength produced by the $\Delta\Sigma$ heater driver for both the ARM and the polysilicon detector ring. We also show the ratio between 10σ value and the LSB of the 8-bit heater driver DAC for comparison.

5.4.2 Tracker and Controller

As opposed to tracking an averaged photocurrent [40, 63, 82, 97], which can mistake changes in the 1/0-balance of the data (which is on-off encoded) for a drift in resonance, the power meters are conditioned on the received data to track the one or zero levels directly. Level-trackers were previously proposed in [2, 12, 75] to actively

adjust receiver threshold levels given variances in signal power. Here, we use the tracked photocurrent-level as an indication for how the ring's resonance has drifted in order to move the resonance back to its desired position. Due to area constraints in the platform, we reuse the receiver-half circuit in a bang-bang loop to act as a 5-bit sense-amp based ADC (with the capacitative offset compensation DACs serving as the voltage DAC) in the power meter. During wavelength-locked receiver operation, the power meter takes control of one receiver-half and uses the data stream from the other half, which receives data normally, to use as the conditioning signal. The controller consists of several registers and LUTs that provide flexibility in programming a variety of control schemes, based on current and previous power meter outputs and arbitrary threshold values. The output of the controller is an 8-bit binary value for the the heater strength. A $\Delta\Sigma$-DAC (described in [80]) drives the integrated heater. The tracker and controller are clocked by the data-clock through a divide-by-64 divider.

5.4.3 Wavelength-Locking Demonstration

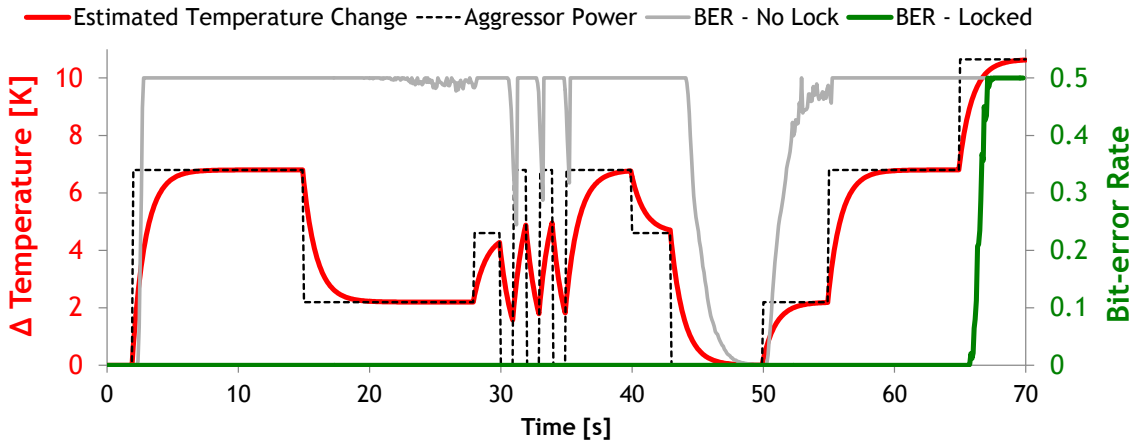

Figure 5-23: Transient lock experiment with the tuning controller locked and without lock. The estimated temperature change is calculated from the aggressor power profile using 20 K/W and a time constant of 1 s.

We perform a 2.5 Gb/s single-rate transient wavelength-lock experiment, shown in Fig. 5-23. We clock-gate on/off the various components in the digital backend of adjacent transceiver macros to create a temperature aggressor. These induce a

microring temperature change of approximately 20 K for every Watt of power they dissipate. In this experiment, we program the controller with a simple tuning scheme where we simply heat more if the power meter reading exceeds a set threshold and heat less if the readings are less.

During initialization (completed prior to $t = 0$), we begin at the maximum heating value and step-down, which keeps the ring in a stable red-biased regime. This is in contrast to step-up initialization, which blue-biases the ring and enters a region of self-heat instability. During this phase, the power meter ignores the conditioning bitstream and updates continuously. Initialization continues until a sufficient power meter reading threshold is crossed, at which point the power meter switches to normal data conditioning mode The wavelength-locked receiver is completely error-free until 65 seconds into the test, when we apply a deliberately large temperature change (approximately 11 K) to exceed the lock range to force failure. By contrast, we see that an unlocked receiver fails immediately with any temperature perturbations caused by the aggressor due to a drop in photocurrent caused by the drift in the ring's resonance. The tuning backend consumes 0.43 mW at 2.5 Gb/s (171 fJ/bit) and 0.024 mm^2, excluding the heater driver and the receiver. The area is largely dominated by the LUT and data-path and the tuning algorithm can also be synthesized directly into gates at design time to conserve power and area.

The lock range of the tuner is 0.5 nm (90 GHz) [77], corresponding to a 9 K change in temperature. In the context of the 9-λ transmit and receive DWDM macros, this range is marginally sufficient to tune out local process mismatch but is still less than the total desired (1.5 nm); applying the ring-to-λ assignment shifting schemes from [18], the tuning range needs to be around that of the channel spacing (1 nm) plus the worst-case local process variation (0.5 nm) to guarantee that each ring can be assigned and tuned to a laser channel λ across all temperatures.

5.5 Monolithic Chip-to-Chip Link in Bulk

Using the transceiver macros, we build a chip-to-chip optical link through 5 m of single-mode fiber interconnecting the two chips (Figure 5-24). We note that this demonstration uses a single λ and clocks are forwarded to both chips electrically. Normally, in a DWDM configuration with multiple λs, the clock would be available as a transmit-forwarded signal on one of the other wavelengths. We demonstrate a full-rate 2 Gb/s chip-to-chip link that is error-free for 2×10^{12} bits (Fig. 5-25). The maximum data-rate of the link is currently limited by the degradation of receiver sensitivity at higher rates and the 10 mW maximum output power of the off-chip laser. We currently hit this limit due to an extra 9 dB of optical loss caused by a sub-optimal permutation of optical devices in the transceiver macros; VGCs in the transceiver macros contribute 5 dB of loss per coupler, compared to the 3 dB loss VGCs present elsewhere on the platform. Additionally, the severely overcoupled photodetector rings (2.4 dB ER_i) effectively results in another loss of 3.7 dB. To overcome the loss with the current combination of devices, we insert an optical amplifier between the transmit chip and the receive chip, which adds approximately 8 dB of optical gain and enables 5 Gb/s operation for the link. The optical power at each point in the link is shown in Figure 5-26 for both cases. The link consumes 4 pJ/b electrical and 5 pJ/b optical energy at 2 Gb/s. At 5 Gb/s, the link consumes 3 pJ/b electrical and 12 pJ/b optical energy.

We perform a link power analysis across a range of data-rates using the metrics from the measured circuits (Fig. 5-27). We calculate the wall-plug laser power of the link as

$$P_L = \frac{1}{\eta_e} \cdot \frac{S_i}{R_{pd}} \cdot \frac{1}{1 - 10^{-ER/10}} \cdot 10^{\sum L/10} \tag{5.5}$$

where η_e is the laser wall-plug efficiency, S_i is the receiver sensivity (in A_{pp}), R_{pd} is the photodetector responsivity, ER is the modulator extinction ratio and $\sum L$ is the total optical loss in the path from the laser to receiver. For this analysis, we assume an $\eta_e = 25\%$. We show a case using devices currently in the transceiver macros and a case using best-known devices on the current platform. In both cases, the link is most

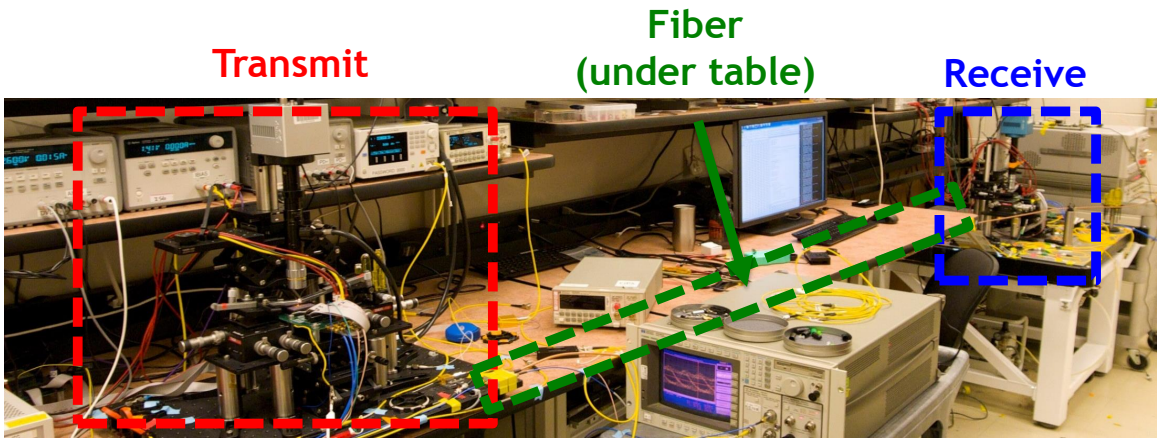

Figure 5-24: Test setup for the chip-to-chip optical link across the room

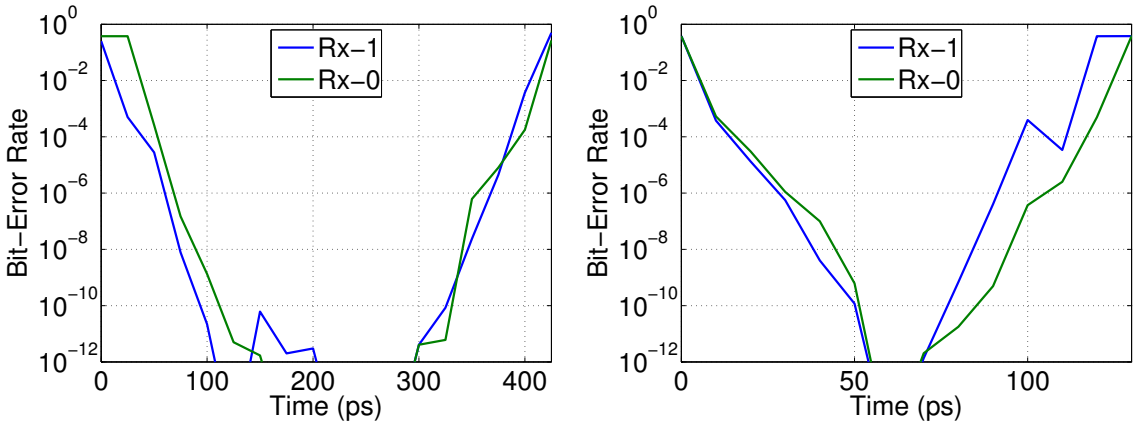

(a) 2 Gb/s BER (b) 5 Gb/s BER

Figure 5-25: 2 Gb/s (a) and 5 Gb/s (b) link BER measurements for 2×10^{12} bits.

energy-efficient at a data-rate of 3–4 Gb/s. Static power (microring heating/tuning, receiver TIA, the receiver) is amortized at higher rates. Energy-per-bit from the laser also decreases initially from 1–3 Gb/s, where S_i (and hence P_L) remains flat. At the higher rates, S_i increases drastically and offsets the electrical energy-per-bit improvement from the higher data-rate. For the case using the current set of devices in the transceiver macros, the laser wallplug power dominates the energy-per-bit, reaching an optimal wall-plug link energy of 18.3 pJ/b at 3 Gb/s. In the second case, the better devices reduce optical loss by more than 9 dB. As such, the laser is no longer dominant for 1–4 Gb/s and an optimal energy-efficiency of 6.5 pJ/b is reached

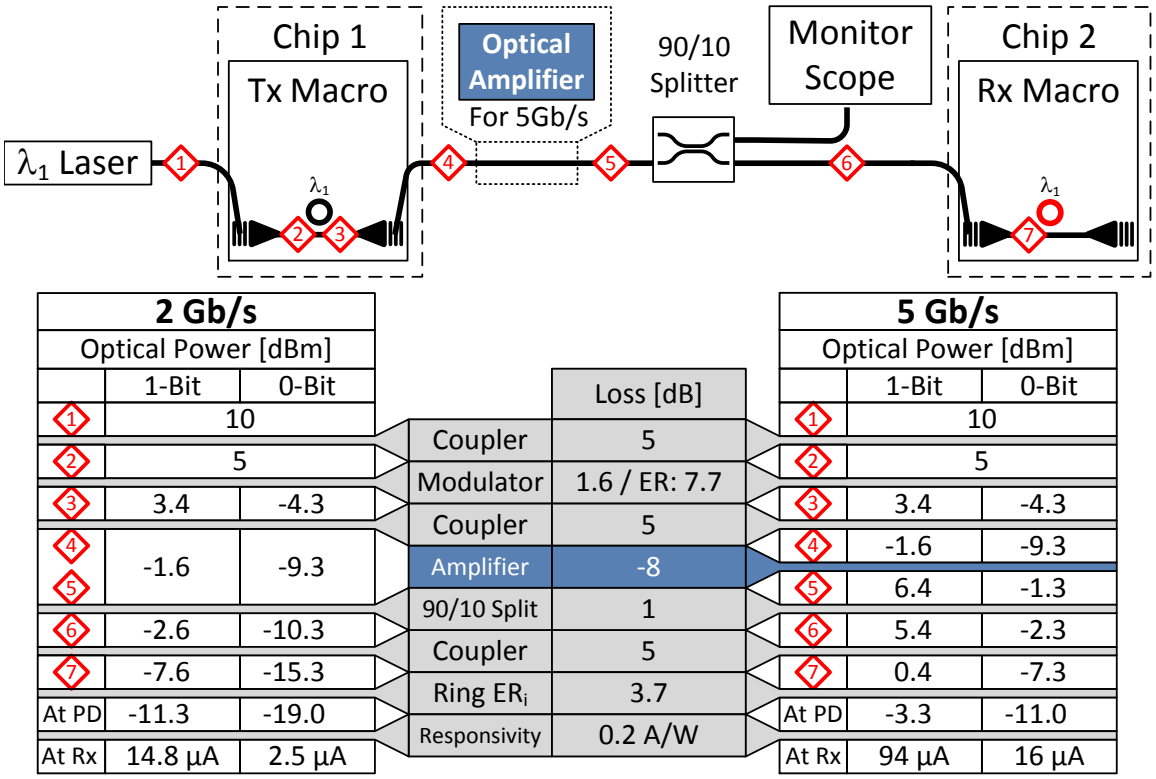

Figure 5-26: Estimated optical power breakdowns of the 2 Gb/s and 5 Gb/s chip-to-link links.

at 4 Gb/s. We note that the demonstrated bandwidth density of 110 Gb/s/mm^2 and energy cost of 6–18 pJ/bit far exceeds these same metrics for traditional memory modules.

5.6 Summary

To facilitate the adoption of photonics for mainstream CMOS applications, we must first remove its intimidation factor: prove that the technology is viable for electronics integration without exotic processing steps, customized SOI wafers, or complicated packaging. To this end, we demonstrate a polysilicon-only monolithic photonics platform in bulk CMOS, to show that low-loss waveguides and active optical structures can be integrated through a minimal number of process changes. By avoiding epitaxially grown crystalline silicon for waveguides and finding an alternative

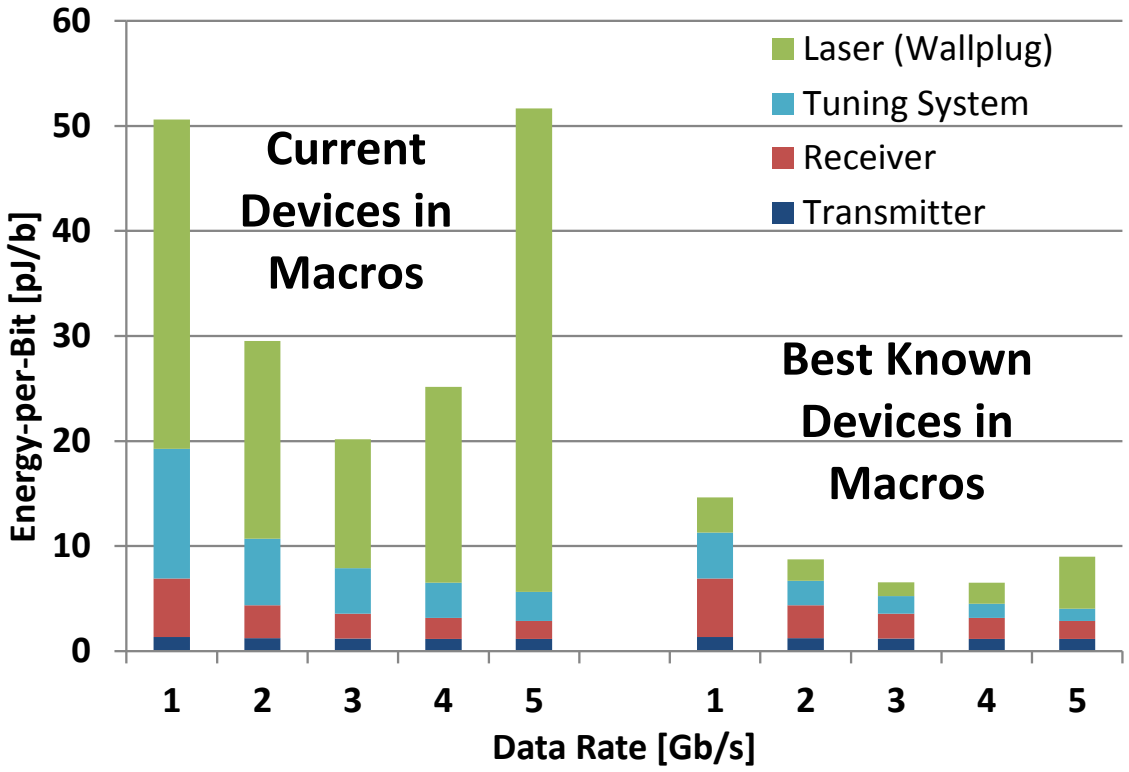

Figure 5-27: Link power analysis across data-rates using the current set of devices placed in the transmit/receive macros and if the best-known measured devices are placed into the transmit/receive macros.

to germanium integration for photodetectors, the polysilicon-only photonics module minimizes risks to process-native transistors and improves technology portability.

We build a portfolio of DWDM link components–a DWDM transmitter and a DWDM receiver–each demonstrating competitive performance and efficiency despite the early stage of platform development. We further improve the robustness of the wavelength-locking techniques by resolving the data-dependency challenge, demonstrating that robust DWDM ring tuning is feasible in a hostile temperature environment. We compose an optical chip-to-link, which, to the knowledge of the authors, is the first demonstration of a monolithically-integrated optical link in a bulk CMOS process. The unique polysilicon aspects of this photonic platform make it independent of the front-end integration processes, enabling deployment in more advanced bulk CMOS process nodes to further scale the link performance and energy-efficiency.

Chapter 6

A Monolithically-Integrated VLSI Photonics Platform in SOI CMOS

Monolithic integration of photonic devices is a promising approach to unlocking the full energy-efficiency and bandwidth density advantages of silicon-photonic interconnects. So far, however, monolithic efforts have been primarily restricted to older customized processes; the quest for the most optimal optical devices possible have lead to process customizations that conflict conflict with advanced transistor design. In order to enable photonics in advanced process nodes, we can alternatively adopt a zero-change approach, whereby photonics devices are designed to fit within an existing electronics flow. Though the optical device designs are more constrained, access to superior transistors allows this drawback to be compensated for through circuit techniques. In addition, the ability to leverage existing IP available for a mature electronics process process, combined with a high-volume volume manufacturing, yields a platform immediately capable of supporting a VLSI electronic-photonic system. As an embodiment of the zero-change approach, this chapter presents a monolithically-integrated photonics platform in a high-performance commercial 45 nm CMOS SOI process. We describe the platform in detail in Section 6.1 and present the transceiver components in Sections 6.2 and 6.3, including the demonstration of an integrated bit-statistical thermal tuner. Using the transceivers, we demonstrate a chip-to-chip

link in Section 6.4. Finally, in Section 6.5, we evaluate the potential of this platform for DWDM.

6.1 Zero-Change Monolithic Photonics Platform

The monolithic SOI platform is a 45 nm CMOS SOI process with 11 metal layers, fabricated in a commercial IBM foundry. Compared to the bulk platform of Chapter 5, the SOI process represents a higher performance (and potentially higher cost) platform designed to support flagship performance-oriented products. Notably, this 45 nm SOI process was also used to fabricate the Playstation 3 Cell processor, the WiiU Expresso processor, and the IBM Power 7 [30]. Photonic devices are integrated into the process through a "zero-change" approach where the designs to conform to an existing (purely-electrical) design flow. We submit our designs as part of a multi-project shuttle run and the design is integrated alongside other purely electronic chips without affecting the yield of the process. From the point of view of the foundry, the integration method is seamless; to them, we are just submitting another electronics design, as no extra processing steps, mask sets, or special handling is required to enable our devices.

6.1.1 Platform Overview

The process frontend consists of a crystalline silicon layer, patterned to form the body of transistors. The crystalline silicon layer is sandwiched between the buried oxide layer and the interlayer dielectric layer. Layers of nitride used for spacers and strain engineering are also present on top of the silicon layer. The high-quality crystalline silicon body layer enables a high-index waveguide core. To form vertically asymmetric structures, such as vertical grating couplers or ridge-shaped waveguides, we can deposit the polysilicon layer used for transistor gates on top of the silicon layer. The dimensions of each layer are representative of a scaled microelectronics process in SOI and a number of process specifications can be found in published articles from IBM [30, 37, 57].

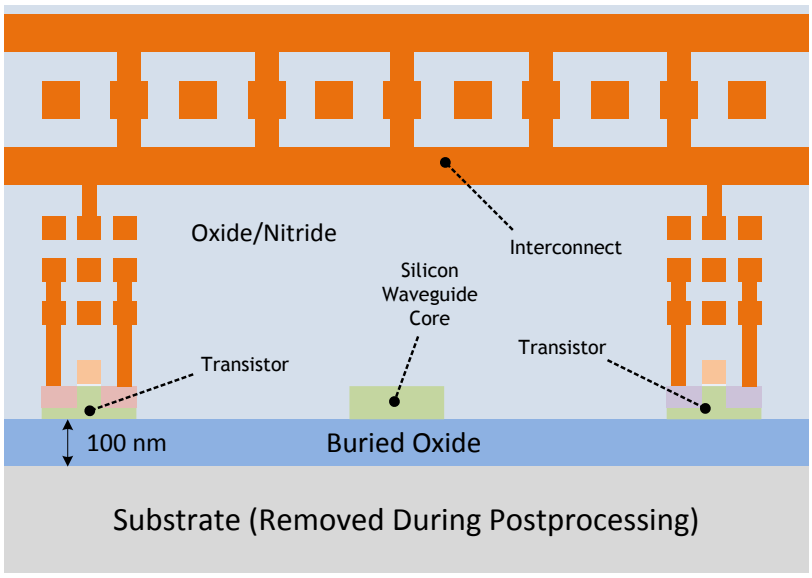

Figure 6-1: Cross-section of the SOI platform. Transistors and waveguides are formed by patterning the crystalline silicon layer. A layer of thin-buried oxide separates the crystalline silicon layer from the chip substrate. We keep the region around optical waveguides clear of any metal interconnect by adding metal fill blockages during design time.

On the top of the high-index silicon waveguide core, the lower-index interlayer dielectric and the nitride layer provide optical mode confinement, provided the lower metal layers are clear. Below the waveguide core, the buried oxide layer (BOX) in a scaled SOI process is not sufficiently thick to provide confinement; the thin-BOX layer cannot completely isolate the silicon waveguide core from the silicon substrate. To enable a suitable waveguide undercladding, we perform a post-process substrate removal step to remove the substrate [22, 62]. After substrate removal, waveguides built using the crystalline silicon layer demonstrate a waveguide loss of approximately 3 dB/cm [62]. Waveguides built using the polysilicon transistor gate layer demonstrate a loss in excess of 50 dB/cm, which is consistent with the unoptimized polysilicon loss from [61].

As a substrate-less chip is both difficult to handle and fragile to work with, the post-processing step must take into account the method for electronics packaging. For chips undergoing wirebonded packaging, the substrate removal step is performed prior wirebonding. First, the chips are flipped upside down and mounted (pad-side

down) to a silicon wafer, leaving the backside of the chip exposed. Next, the chips are placed inside an XeF_2 chamber so that the gas etches away the silicon substrate. As XeF_2 has a 1000-to-1 etching selectivity between silicon and SiO_2, the BOX forms a natural stopper for the etching process. The etch leaves behind a thin film containing just the buried oxide, the silicon layer, and the metal stackup. To maintain structural integrity of the die during the wirebond process, we attach a silicon-carbide or glass handle to the film using an optical adhesive prior to detaching from the mounted silicon wafer. Finally, we wirebond the part to a package.

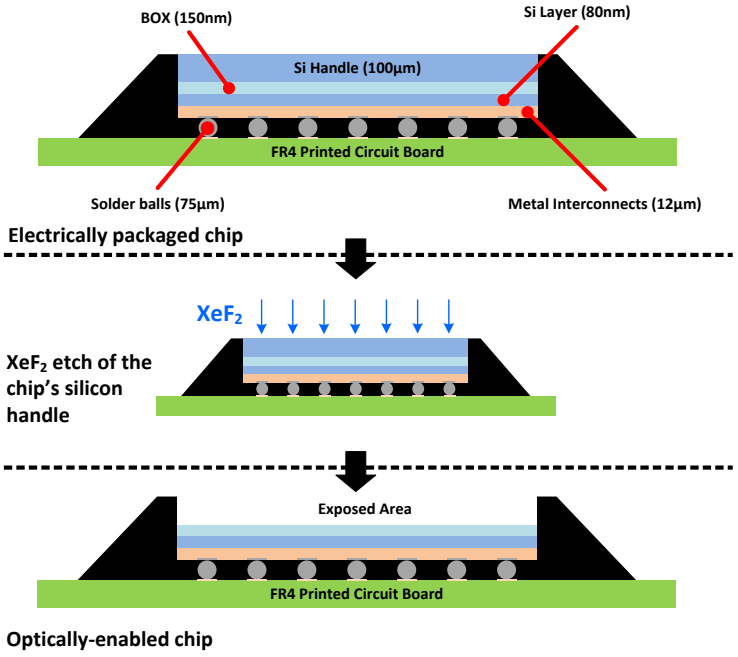

(a) Die cross-section and substrate removal

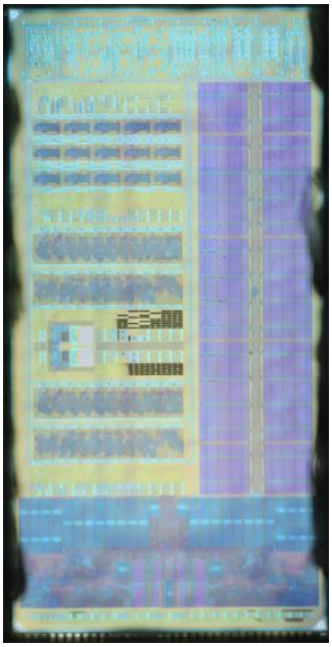

(b) Substrate Removed Die

Figure 6-2: The C4 flip-chip package structure, where the chip is flip-chip-attached to a printed circuit board, with post-processing steps to remove the substrate to enable on-chip photonic structures (a). A backside die-photo of a substrate-removed chip (b). Normally, the substrate would block the view of the chip from the back and we would not be able to see the structures on the chip.

The flip-chip packaging approach is more attractive due for high-performance parts due to better power delivery, pin counts, and signal integrity of the I/O pins. In a flip-chip package, the final packaged part is mounted pad-side down to a ceramic package or circuit board using controlled collapse chip connection (C4) balls, which

form the electrical pin connections. Epoxy encapsulation is added to the mounted part for additional mechanical support. We incorporate substrate-removal post-processing for a flip-chipped part, shown in 6-2, where the pad-side down orientation of the chip post-packaging enables a package-first substrate removal strategy. Here, we package the chip first and place the packaged chip into the XeF_2 chamber. The epoxy encapsulation prevents the etchant from attacking the sides of the chip while the backside is etched away. The board or package provides the mechanical support and there is no need to attach a new handle. The package-first approach avoids putting mechanical stress (from the packaging steps) on post-processed die, when it is far more more fragile. In addition, we can electrically test the dies both before and after the substrate removal process. As a result, the post-processing yield is above 80 %, which is considerably higher than the package-last approach of wirebonded packages of $< 50\,\%$. We expect the post-processing yield to improve even further once we more precisely define the conditions for when to stop the etch (currently through just visual inspection) and once the pump cycling procedures for the etch are optimized. Substrate removal is a commercially viable technique and can be performed at wafer-level, as shown previously to make suspended-layer microphones in standard CMOS [68]. In addition to providing the necessary waveguide isolation, substrate removal also improves the tuning efficiency of the platform's resonator structures [88].

To verify that the electrical performance of the platform is not degraded by the substrate removal process, we measure the oscillation frequency of a ring oscillator, shown in Figure 6-3 and extract the delay value of the logic stages. The measurements show that the removal and transfer steps had a negligible impact on transistor performance. Notably, a stage delay of 5 ps at the nominal process voltage of 1 V makes this the fastest electrical performance of any monolithic photonics platform to date. As transistor performance is unaffected, all foundry IP are compatible with the platform, a fact we leveraged highly to accelerate the development of sophisticated digital electronics on the platform.

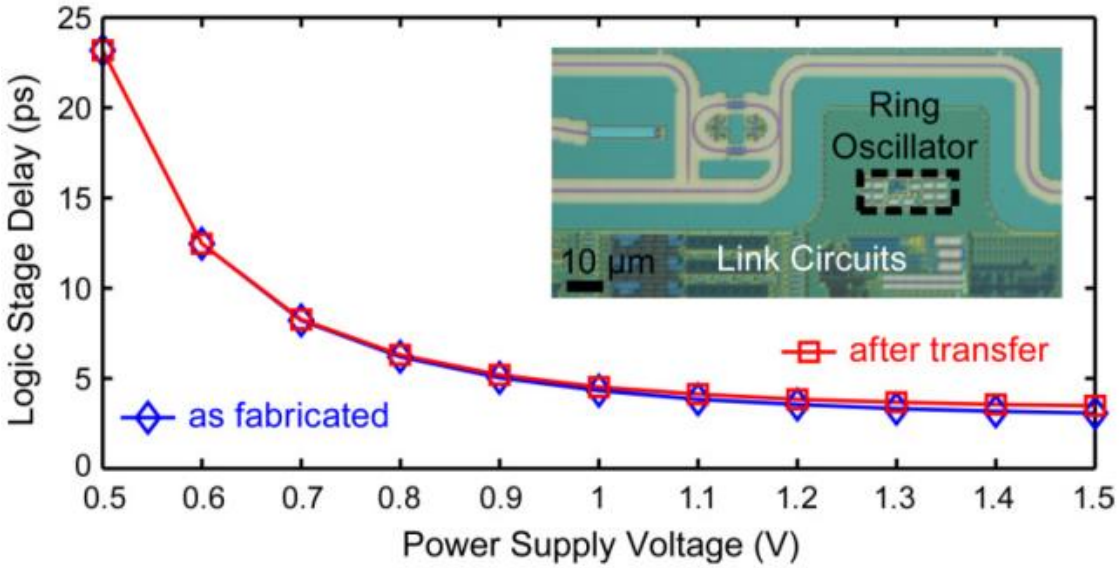

Figure 6-3: Stage delay of a ring oscillator versus supply voltage before and after substrate removal, calculated from the measured oscillation frequency. The nominal voltage of the process is 1 V. We observed no change in ring oscillator stage delays from the substrate removal step.

6.1.2 Electro-Optic Modulators and Photodetectors

The modulator is a multimode microring structure supporting the "whispering-gallery" modes [73, 88], shown in Figure 6-4. The fundamental mode is constrained to the outer radius of the modulator. As such, electrical contacts to the device along the inner circle edge will not introduce free carrier absorption to the fundamental mode. To avoid exciting the higher-order modes, a propagation length-matched coupler with a long interaction length is used to couple the ring to the bus waveguide. Modulation of carriers in the device is performed through interdigitated lateral p-n junctions constructed using mid-level p and n implants. Advanced 45 nm process lithography enables fine definition of the junctions, allowing 84 junctions to fit across the circumference of a 5 µm radius ring. The "spoked ring" modulator operates as a carrier-depletion modulator. Shown in Figure 6-5, the free spectral range of the is 3.02 THz with a measured quality factor of $\approx 10,000$. The ring embeds a resistive heater with an approximate resistance of 400 Ω, enabling thermal adjustment of the

ring resonances. A monitoring photodetector can be connected to the weakly coupled drop-port for a closed-loop tuning solution.

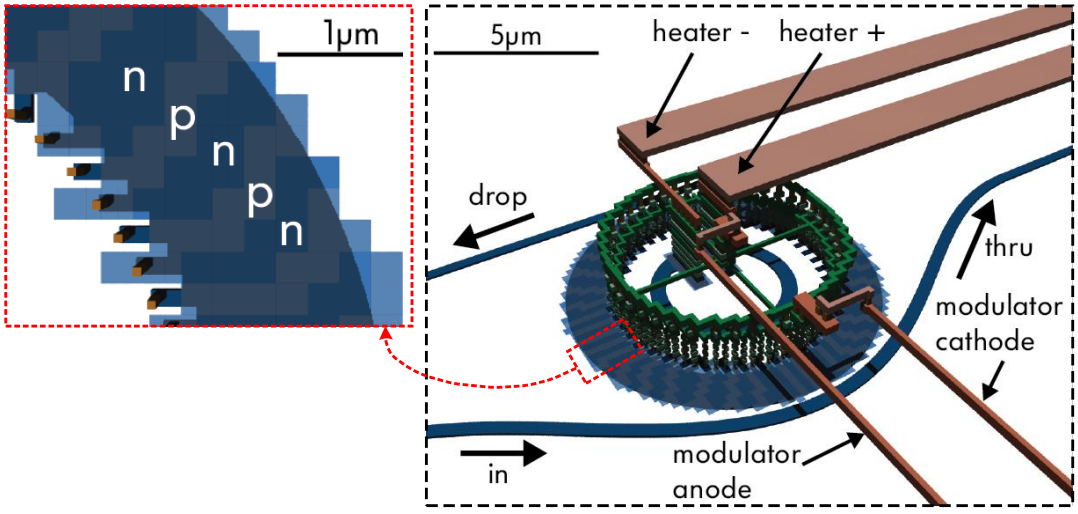

(a) Modulator structure

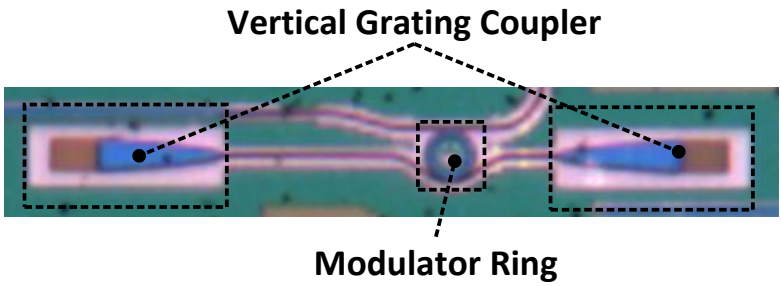

(b) Modulator micrograph

Figure 6-4: Structure (a) and test-site micrograph (b) of the "spoked ring" modulator device.

Compared to the carrier injection modulator designs earlier in the platform's history [54, 62], the depletion-based design of the spoked ring modulator does not consume forward-biased static current, enabling greater energy efficiency. The design also avoids the previous carrier-lifetime-limited modulation speeds. The multimode design also enables the device to be contacted effectively without resorting to the use of ridge-waveguides of the rings in [54, 62]; ridge-waveguide rings in this platform require the use of the lossy polysilicon layer and severely degrades ring resonator Q-factor.

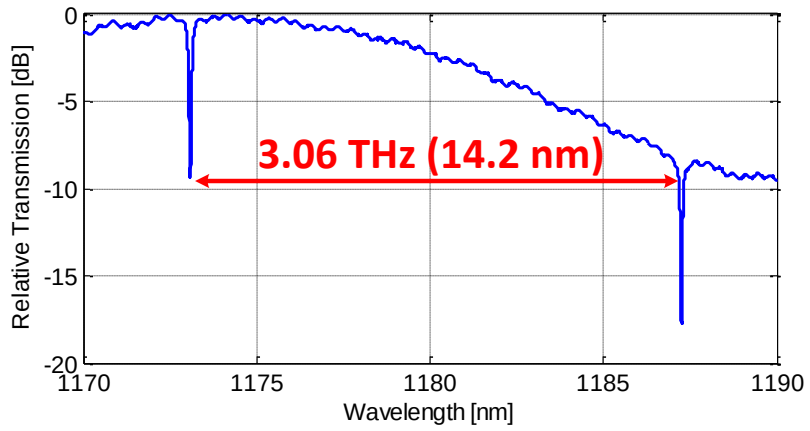

(a)

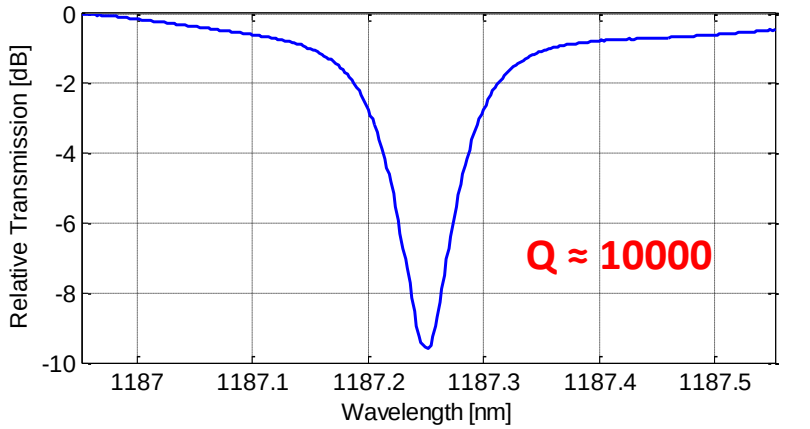

(b)

Figure 6-5: Optical spectrum plots of the "spoked ring" modulator device, showing both the device free spectral range and the quality factor.

Nominally, advanced processes introduce SiGe as a material to strain engineer electrical transistors for improved performance. In this platform, we can use the same SiGe to form photodetectors, shown in Figure 6-6. Currently, the native process places two constraints on the efficiency of the detector. First, unlike pure germanium detectors, which can absorb light at wavelengths in the 1550 nm range, the low mole fraction of germanium in the SiGe limits the absorption band to a far lower wavelength. Though the compressive stress created by the SiGe reduces the bandgap and extends the absorption band, the photodetector is still limited to about 1200 nm and is weakly absorbing, requiring longer photodetectors to completely absorb the

light. Second, as the SiGe material is used only in PMOS source drain regions to apply compressive stress to the PMOS transistor channel, the SiGe material is inherently heavily p-doped. As such, the maximum efficiency of the detector is limited by the balance between light lost to free-carrier absorption and light absorbed for photogeneration of carriers.

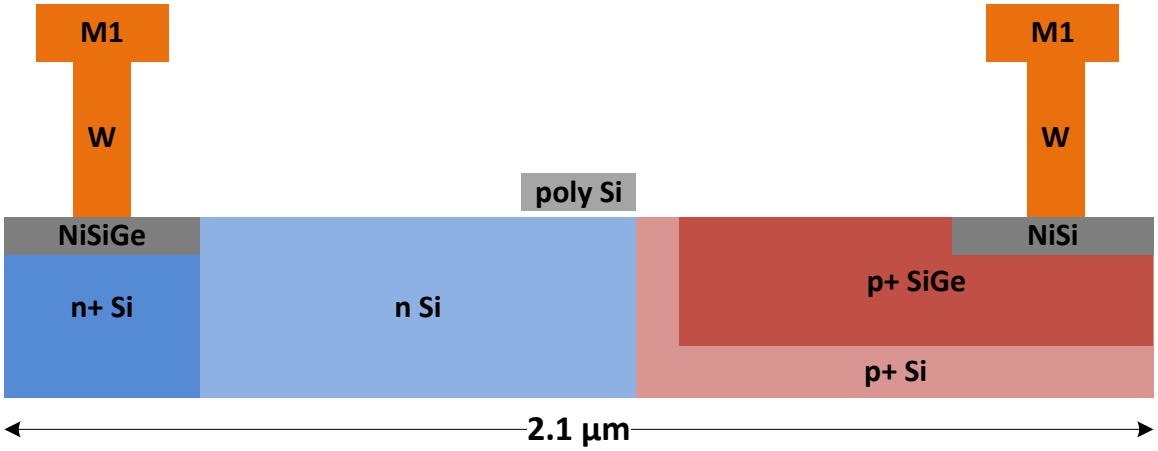

(a) Photodetector structure

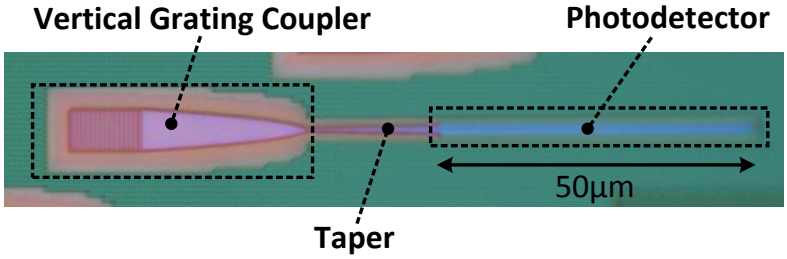

(b) Photodetector micrograph

Figure 6-6: Structure (a) and test-site micrograph (b) of the photodetector constructed using SiGe available in the process for compressive stress engineering of PMOS transistor channels.

The long straight detector shown in Figure 6-6 has a responsivity of 0.02 A/W at 1180 nm. As shown in Figure 6-7, the device exhibits a bias-dependent bandwidth with a gentle roll-off. When reverse-biased at 5 V, the device has > 4 GHz of bandwidth. We note that at a length of only 50 µm, the responsivity of this detector is limited by the length of this device.

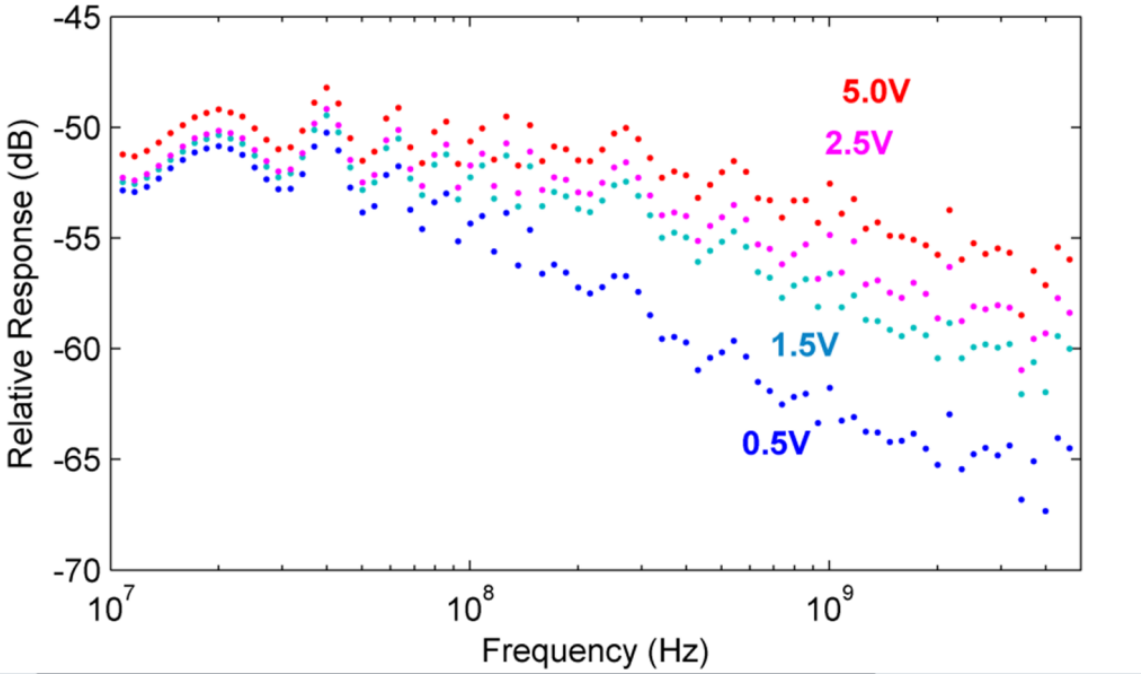

Figure 6-7: Bandwidth measurements of the SiGe photodetector, applied under different bias voltages.

Interestingly, a spoked-ring modulator, when used as a resonant photo-detector, exhibits an optical responsivity of 0.1 A/W, shown in Figure 6-8. Though the resonant structure certainly helps with absorption, the device does so without any use of SiGe. There are two possible explanations to this behavior. First, as the entire spoked-ring modulator is doped, the dopant atoms may create a sufficient number of defect states in the crystalline silicon lattice, which cause photoabsorption in a similar manner as the detectors in Chapter 5. However, this is unlikely; defect states lower minority carrier lifetimes and, as a result, are kept to a minimum in the crystalline silicon (the transistor body) through various annealing processes in advanced processes. The second (and more likely) cause of the behavior may be bandgap narrowing effects; the absorption coefficient of silicon at 300 K drops off rapidly beyond ≈ 1100 nm, where the photon energy is equivalent to the 1.12 eV bandgap of silicon. An reduction in bandgap energy of less than ≈ 100 meV will move the drop-off to 1200 nm. The bandgap energy reduction could be caused by the presence of silicon-nitride stress

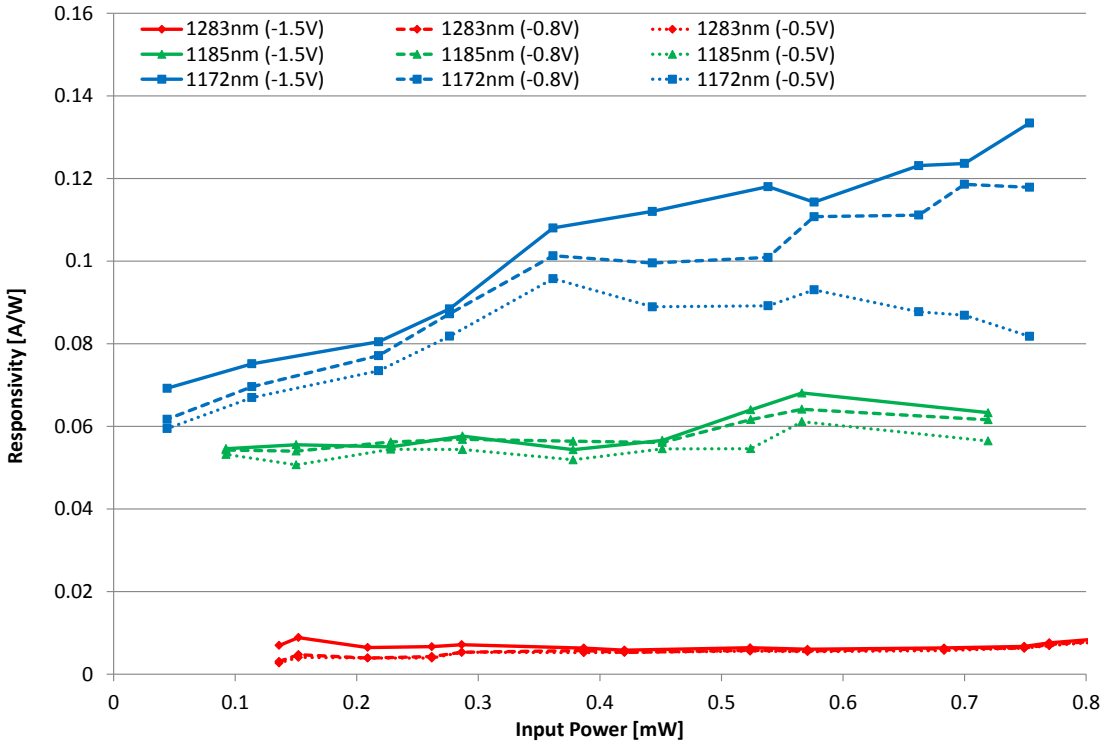

Figure 6-8: Optical responsivity plot at different bias voltages at different wavelengths of the spoked ring modulator, used as a photodetector. Each measured wavelength corresponds to a different resonance of the measured device.

liners [57], which are used for aggressive strain engineering of transistors and can cause a reduction in bandgap energy [10, 31, 55]. In addition, the dopants used to form the spoked ring junctions will also cause dopant-based bandgap narrowing [34, 89]. Both effects combined can result in a sufficiently large reduction in bandgap energy. If these effects can be taken advantage of, we can potentially build novel silicon-based photodetector designs. Such work, however, is outside the scope of this thesis.

6.2 Bit-Statistical Wavelength-Locked Transmitter

We integrate the modulator device with a modulator driver circuit, consisting of a 2-to-1 serializer, a voltage level-shifter, and a push-pull driver head. The serializer and driver head use V_{DD} and V_{DDH}, respectively, as their supply voltages, while the voltage level-shifter interfaces between the two voltage domains. Both the level-shifter

and driver head use thick-oxide transistors (used in this process for I/O transistors) to enable voltage drives above the process nominal logic voltage of 1 V. We use V_{REF} and V_{DDH} to control the modulator drive voltages in the forward and reverse regimes, respectively. On data *zeros*, NMOS transistor N1 drives the device anode terminal to 0 V while the thick-oxide inverter I2 brings the cathode cathode to V_{DDH}, applying a voltage of $-V_{DDH}$ across the device. On data *ones*, the NMOS pullup N2 pulls the anode voltage up to $V_{REF} - V_{Tn}$ (where V_{Tn} is the threshold voltage of the NMOS N2), while I2 drives the cathode to 0 V, providing $+(V_{REF} - V_{Tn})$ across the device.

The level-shifter gate is ratioed and draws significant amount of current when it switches as the pull-down NMOS transmit fights the pull-up PMOS. In addition, the device achieves reasonable extinction ratio even under voltages compatible with standard digital logic. As such, the thick-oxide transistors are unnecessary. These insights allow us to simplify the design of the driver to that of a simple chain of standard logic inverters, shown in Figure 6-10. The driver now drives the device single-ended from the cathode terminal, providing a peak-to-peak voltage swing equivalent to the modulator supply voltage, V_{DD}. The anode terminal of the device is connected to a bias voltage, V_{BIAS}. On data *zeros*, the voltage across the modulator device is V_{BIAS}. On data *ones*, the voltage across the device is $V_{BIAS} - V_{DD}$.

We integrate the inverter-based driver with the wavelength-locking thermal tuner we proposed in Chapter 4 and comes complete with the bit-statistical tracker, a optimizing controller design, and the self-heating cancellation logic. We separate the implementation of the tuner into a frontend and a backend. The frontend implements the integrator which integrates drop-port photocurrent for a configurable interval of 16-128 transmitted bits. A 6-bit successive approximation register (SAR) ADC converts the integrator output to digital values. An aligner FSM in the backend coordinates the start and reset times of the integrator and counter to match the integration window with the counted bits. Using the ADC output and the counts of *1*s and *0*s counts, the solver implements the bit-statistical tracker calculations, outputting the power levels of optical *1*s and *0*s (L_1 and L_0) as well as the modulation depth ($L_d = L_1 - L_0$). The tuning controller auto-lock the ring using the eye-maximizing

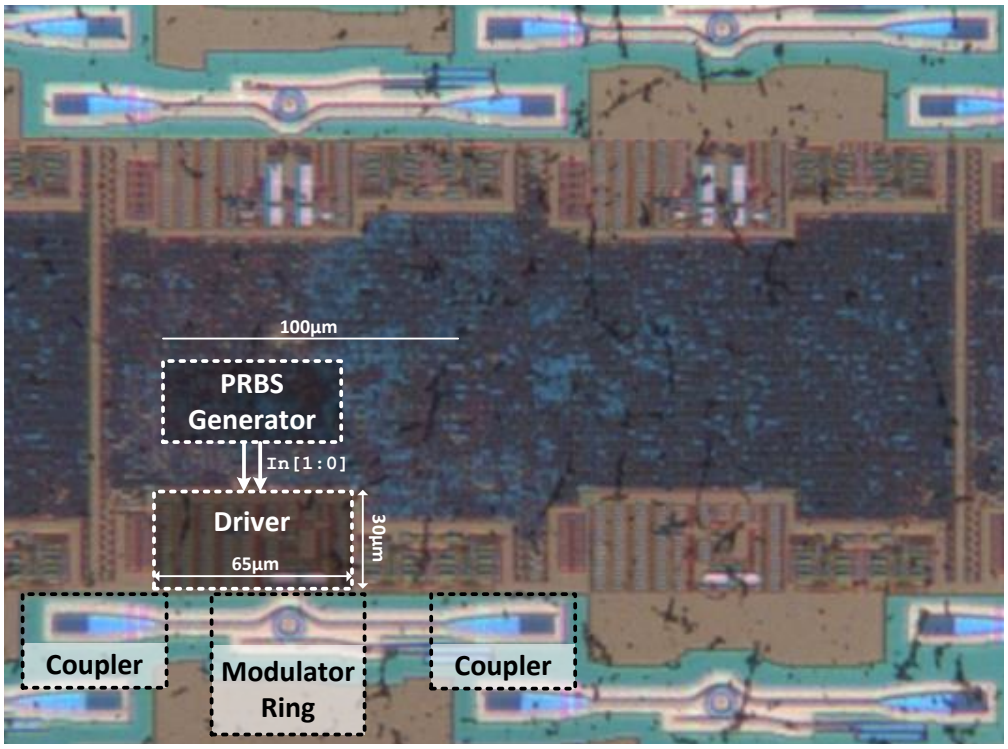

(a) Test site for transmitter with push-pull driver

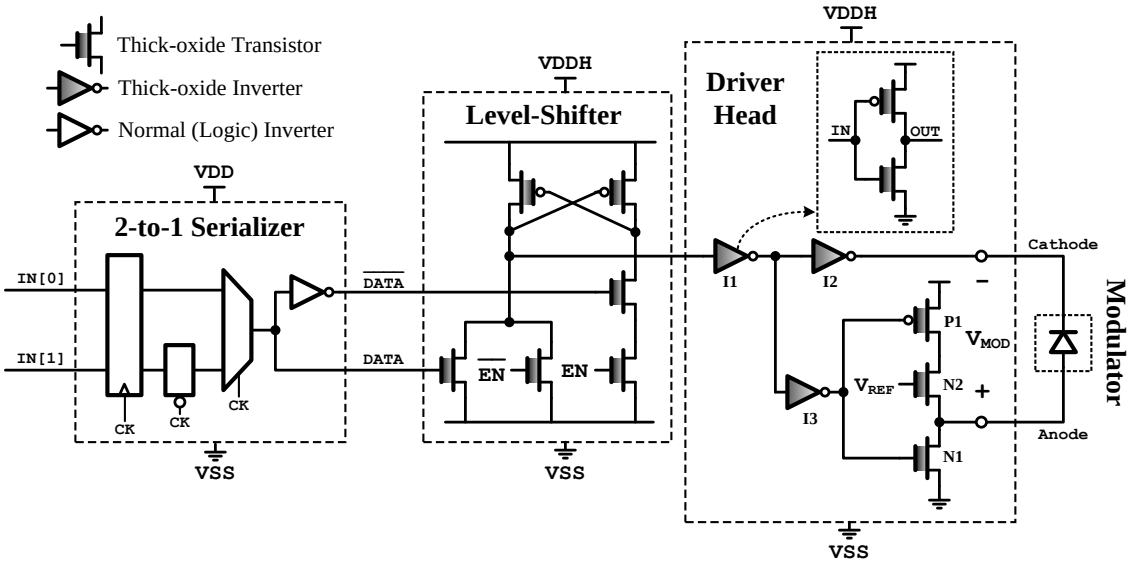

(b) Transmitter with push-pull modulator

Figure 6-9: Implementation of the push-pull modulator driver. Thick-oxide devices were used to enable higher voltage operation. The driver circuit consumes an area of 65 µm × 30 µm.

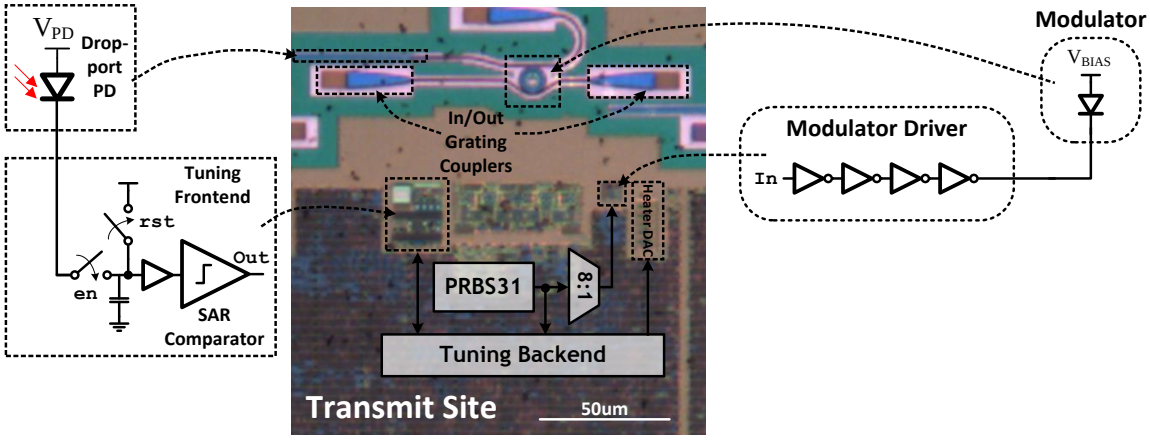

Figure 6-10: Implementation of the inverter-based modulator driver. The driver occupies only a $7\,\mu m \times 7.5\,\mu m$ area footprint. The driver is integrated with a wavelength-locking thermal tuner. Implementation of the photocurrent integrating frontend of the thermal tuner is shown.

method of Chapter 4. The self-heating canceller nullifies the data-dependent self-heating power changes to keep the resonance stable. The backend of the system operates on aggressively-divided clock domains to minimize power consumption.

In Figure 6-11, we demonstrate the capabilities of the wavelength-locking transmitter. In this test, the tuning circuitry auto-locks the modulator ring (with a non-heated resonance at $\approx 1187.2\,nm$) to a non-tunable 1189 nm laser outputting 2 mW of power. During the lock-on process, the controller finds and locks to the point that maximizes L_d, the maximum eye height. A PRBS31 sequence is used as both the lock-on training sequence and as the signal for the eye-measurement. Under a swing of $-0.7\,V$ to $0.5\,V$, ($V_{DD} = 1.2\,V$, $V_{BIAS} = 0.5\,V$) at $5\,Gb/s$, we achieve a $6.5\,dB$ extinction ratio (ER) at an insertion loss (IL) of $4\,dB$. With forward-biased voltages, electrical junction characteristics limit the data-rate to $5\,Gb/s$. An open modulator eye can be maintained up to $8\,Gb/s$ using fully reverse-biased swings ($-1.2\,V$ to $0\,V$), though ER degrades to $2.3\,dB$. The driver energy cost is $30\,fJ/bit$ with $V_{DD} = 1.2\,V$ for all data rates. We note that the limited achievable data-rate is likely due to an implant mask error; the contacts to the p-type junctions were doped with a mid-level p-implant as opposed to a p+ implant. As a result, the resistance to the modulator junction is very high and degrades modulator transition times.

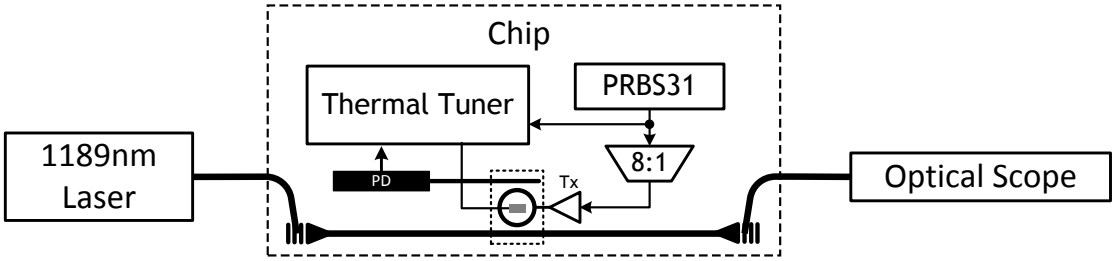

(a) Test setup for the transmitter wavelength-lock experiment

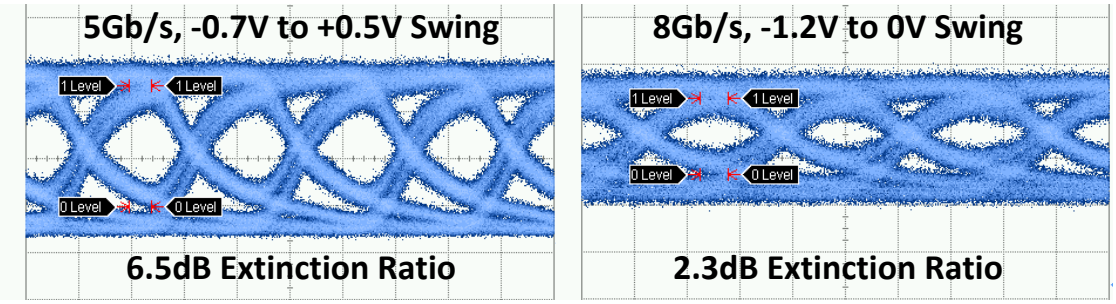

(b) Transmitter eyes at 5 Gb/s and 8 Gb/s

Figure 6-11: Test-setup (a) and optical eye diagrams (b) of the demonstration of the wavelength-locking transmitter.

We use a controller decision rate of 76 kHz. The upper-bound on lock time is 7 ms, assuming a 1 LSB step size for *search* and a worst-case P_{H-init} from the final lock position. The lock time can be decreased by using a larger step size during the *search* state. The tuner frontend sensitivity is < 500 nA, though the drop-port PD provides ≈ 3.5 µA of current for the laser power used. The $\Sigma\Delta$ heater driver DAC provides 600 GHz (2.7 nm) of tuning range (≈ 60 K of temperature change) and delivers 3.8 µW/GHz overall tuning efficiency. The DAC has 9 bits of resolution, corresponding to 1.18 GHz per LSB of resolution. At the laser power used to test the modulator, however, self-heating enhances the DAC resolution to approximately 12 bits when the laser wavelength is close to λ_0. The frontend (integrator, SAR ADC) and backend consume 39 fJ/bit and 152 fJ/bit, respectively. De-embedding from simulation, we note that approximately 80 % of the backend power comes from the synchronous *1/0* bit counters (to produce N_0), the aligner FSM, and the associated clock buffers. These components are all on the least divided clock and were synthesized for a much

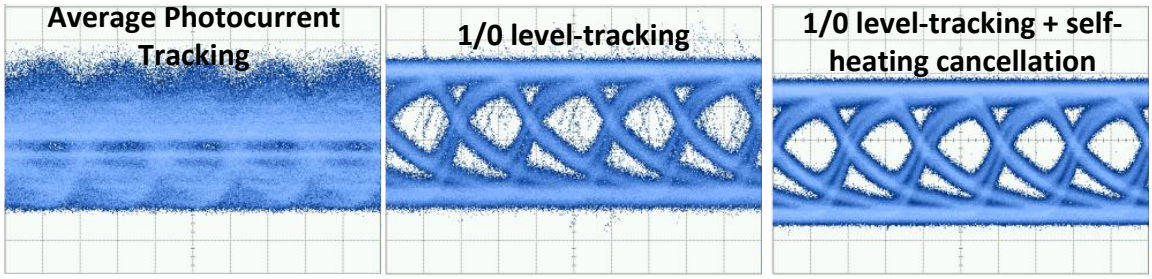

Figure 6-12: 5 Gb/s eye diagrams for the auto-locking transmitter when transmitting data with arbitrarily changing ratios of *1*s to *0*s. Each diagram shows the controller used in a different mode. To use the controller in the average photocurrent tracking mode, we force $L_{diff} = 0$, such that the controller cannot distinguish the different levels for *1*s and *0*s.

higher frequency (3.125 GHz on a slow-slow corner) than intended, leading to aggressive gate size-ups and logic for timing closure. To support a 5 Gb/s, these components only need to run at 625 MHz. In addition to using the correct clock frequency during synthesis, a higher deserialization ratio or an asynchronous counter can both reduce the energy overhead of the bit counters.

To demonstrate the full capabilities of the level-tracking circuitry and the self-heating canceller, we randomly change the 1/0 balance (uniform distribution) of the transmitted data every 200ms, including the all 0s and all 1s cases. The resultant eye diagrams are shown in Figure 6-12 With the tuning circuit operating in average power lock mode (tracking average photocurrent only), the modulator eye is completely closed as the controller is unable to discern a change in 1/0 balance from a drift in resonance. Independent 1/0 level tracking enabled by the bit-statistical tracker allows the controller to lock correctly. However, the sudden change in heating from the laser when the 1/0 balance changes causes transient eye closure until the controller can react. When enabled, the self-heating canceller compensates for this effect, allowing the eye to remain open.

6.3 Data Receiver

The receiver consists of a SiGe photodetector connected to a TIA, followed by two data slicers operating on opposite clock phases. The low parasitic capacitance afforded by monolithic integration enables high-bandwidth and sensitivity from a single-stage inverter TIA design with a high feedback resistance of 5 kΩ. The achieved TIA bandwidth is approximately 7 GHz for a combined PD and wiring cap of 10–20 fF. Capacitive and current DACs in the two data slicers and the current DAC at the TIA input allow for adjustment of the receiver decision threshold, providing a method for offset compensation and eye-monitoring sweeps.

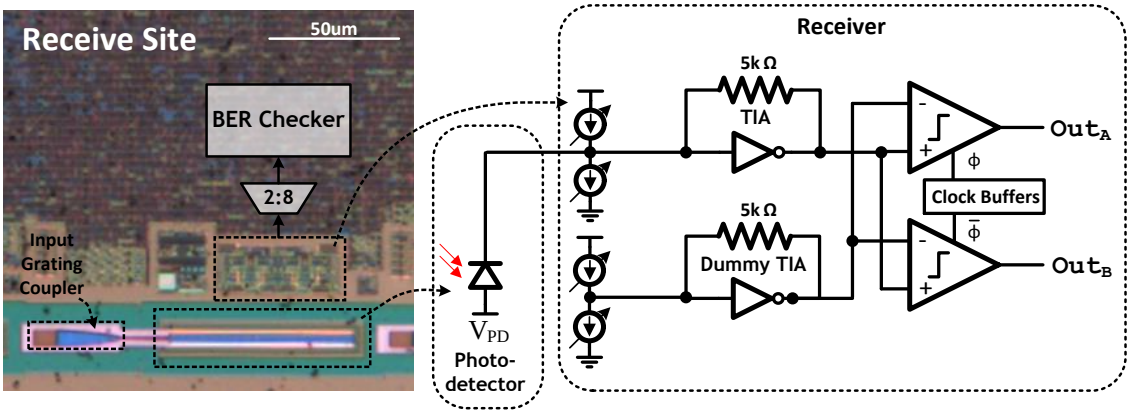

Figure 6-13: Implementation of the data receiver circuit.

In Figure 6-14, we demonstrate the optical receiver using the 50 µm long linear photodetector (the resonant detector is not connected to any receiver circuits). We externally modulate a 1185 nm laser with a PRBS31 sequence using an external modulator and couple the modulated data sequence into the chip directly into the photodetector. For higher data rates, we use the current DACs in the data slicers to apply the sense-amp integration time adjustment technique from Section 5.3 to scale the receiver sensitivity appropriately with data-rate. At the nominal 1 V supply, the receiver achieves a maximum data rate of 10 Gb/s for a $< 10^{-12}$ bit-error-rate at a peak-to-peak photocurrent sensitivity of 15.5 μA_{pp} (−21.1 dBm normalized sensitivity). At 8 Gb/s and 6 Gb/s, the BER $< 10^{-12}$ sensitivity improves to 7.8 μA_{pp} and

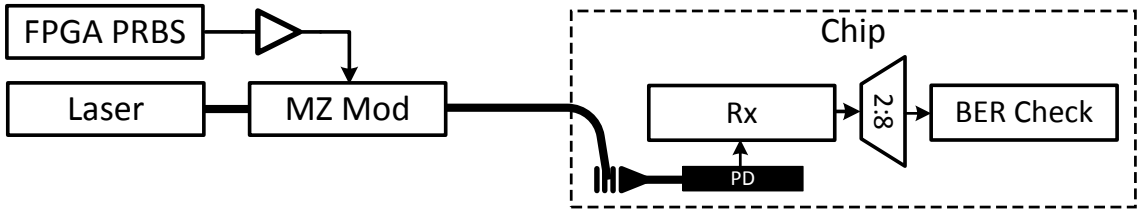

(a) Test setup for the receiver

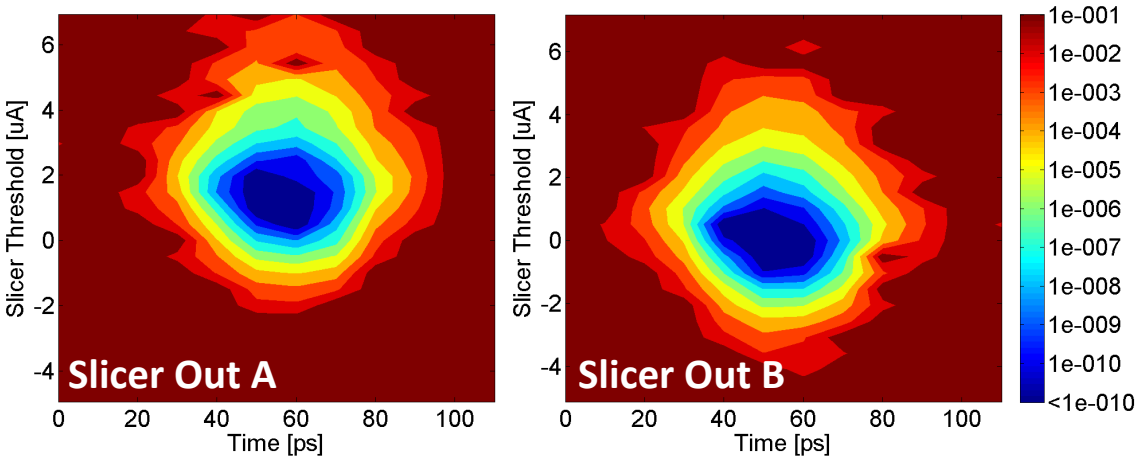

(b) 10 Gb/s bit-error-rate eye

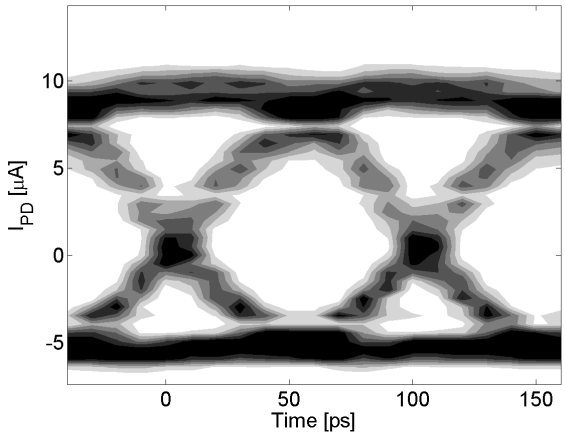

(c) 10 Gb/s eye seen by the receiver.

Figure 6-14: Test-setup (a), 10 Gb/s bit-error-rate eye diagrams of each of the slicer outputs of the receiver (b), and statistical eye shape measured using the output of the data receiver (c). Each point in the BER eye diagram is measured for 10^{10} bits for each data slicer output and the centers of the eyes are error-free during this interval.

$5.9\,\mu A_{pp}$ (-24.1 dBm and -25.3 dBm normalized sensitivity), respectively. Using the 0.02 A/W detector, the achieved optical sensitivity is -4.1 dBm at 10 Gb/s. If paired with the 0.1 A/W spoked ring detector, the optical sensitivity becomes -11.1 dBm.

The power consumption of the receiver consists of both a static power component from the TIA as well as a dynamic power component from the slicer and clock buffers. At 10 Gb/s, the receiver consumes 2.97 mW of power, split between 0.59 mW for the TIA and 2.37 mW for the slicers and clock buffers. The energy efficiencies are 297 fJ/bit, 313 fJ/bit, and 336 fJ/bit at 10 Gb/s, 8 Gb/s, and 6 Gb/s, respectively. Energy efficiency is worse for lower data-rates as the static TIA power consumption is amortized over fewer bits.

6.4 Chip-to-Chip Link

We demonstrate the full functionality of the transmitter, receiver, and bit-statistical thermal tuner in a chip-to-chip optical link. The chip-to-chip link runs at 5 Gb/s with a $< 10^{-10}$ BER, shown in Figure 6-15. The link uses the 50 µm SiGe linear photodetector with a 0.02 A/W responsivity and vertical grating couplers with a loss of approximately 4 dB. The modulator is biased at a point which provides 7.3 dB of on-off extinction and a 3 dB insertion loss. In this link test site, we use an input laser power of 3.8 dBm. In order to overcome the very limited responsivity from the connected photodetector, we insert an optical amplifier between the two chips to bring the incident laser power at the photodetector to a level sufficient that is sufficient for the receiver to function. The amplifier provides approximately 8 dB of optical gain, though it also degrades the extinction ratio by 0.6 dB. Figure 6-16 shows the power breakdown of the circuit components in the link.

Note that we could, in theory, omit the amplifier and simply increase the input laser power by 8 dB to 11.8 dBm to obtain sufficient power to complete the link. This would, however, couple approximately 7.8 dBm of laser power into the transmit ring, causing the self-heating power from the laser to completely dominate over the power output of the embedded heaters in the microring's thermal tuning circuitry.

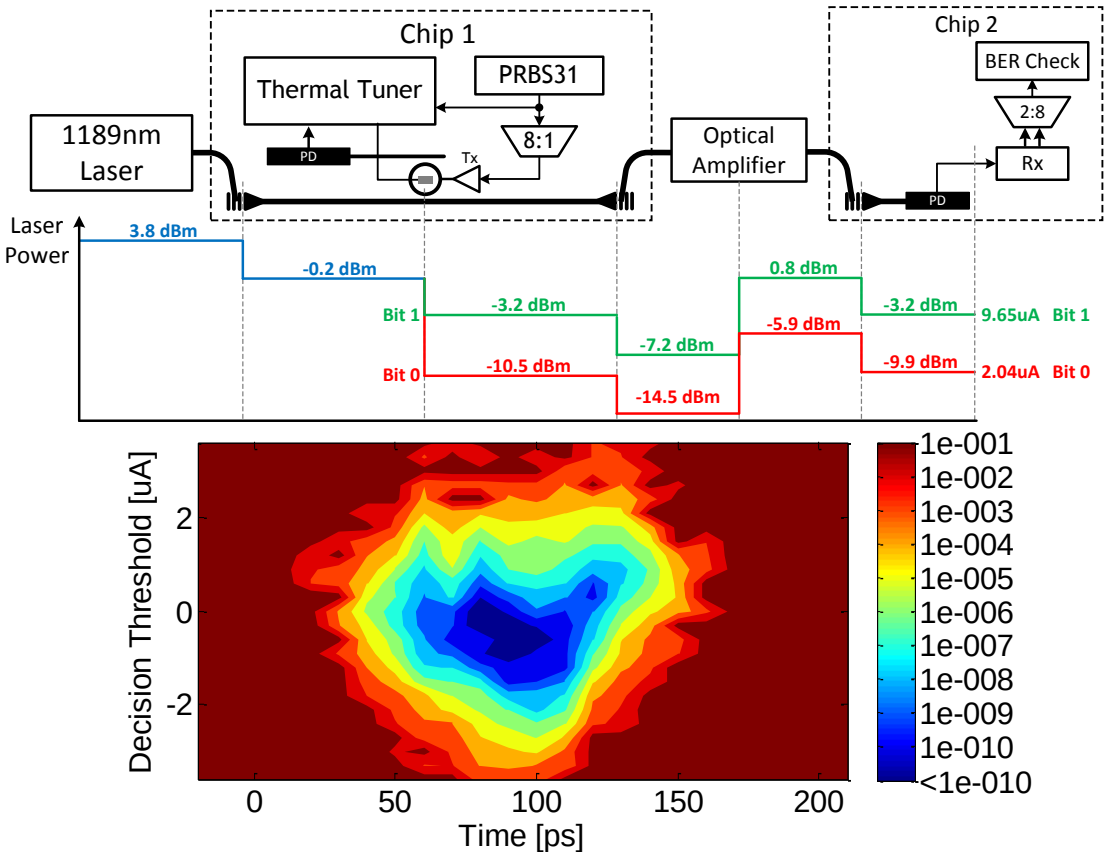

Figure 6-15: Demonstration of the optical chip-to-chip link with the amount of optical power annotated at each point of the link. We show the combined receiver BER eye diagram (from both data slicers). The center of the eye is bit-error-free for $> 2 \times 10^{10}$ received bits.

Experimentally, the strong feedback behavior from self-heating (very high G_L) would leave the heaters with very little control over the ring's resonance (λ_0). Furthermore, the strong self-heating would result in significant eye closure from increasing baseline wander noise and the power output from the heaters would be insufficient to cancel the self-heating transients. As a result of these practical limitations, an amplifier was necessary to complete the link using the current set of devices.

We note, however, that the link test site did not contain the optimal combination of circuits and devices experimentally demonstrated on the platform due to the limited number of device and circuit combinations we could place on the chip, In particular, the photodetector used in the link has a responsivity that is five times worse than the spoked ring (0.1 A/W) and the 4 dB loss couplers are far worse than the 1 dB couplers

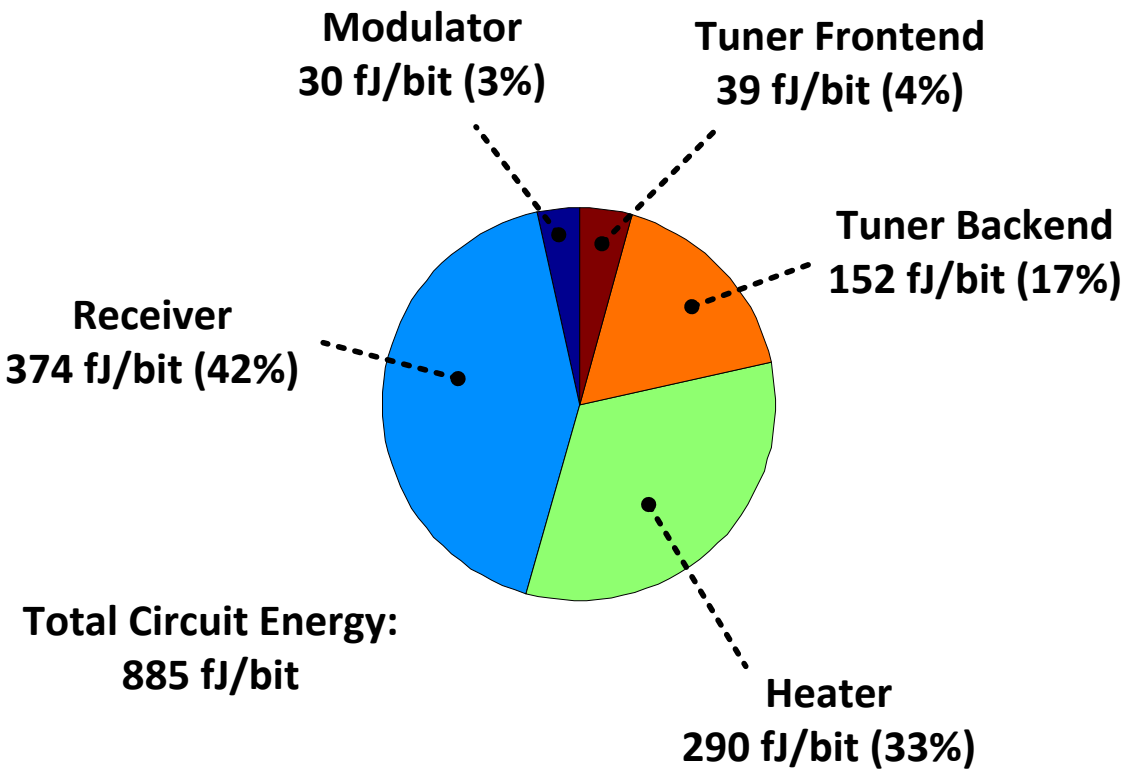

Figure 6-16: Circuit energy per bit breakdown of the 5 Gb/s link demonstration. The heater power corresponds to a tuning range of 1.8 nm. The aggregate circuit cost of the link is 882 fJ/bit.

we showed in [86]. With the optimal set of devices used in the link, we expect the laser power requirement to drop by 16 dB. This would eliminate the amplifier and require only −4.2 dBm of input laser power for the full link.

We can next calculate the link's total optical energy per bit. Using the current set of devices and the optical amplifier, we obtain a total optical power of 2.4 mW from the laser and 0.73 mW from the amplifier. The total optical energy per bit in this configuration is 626 fJ/bit. Hypothetically, if we omit the amplifier and use a single 11.8 dBm laser, the optical energy per bit is 3.02 pJ/bit. If we use the optimal set of devices, the total laser power is 0.38 mW, corresponding to a total optical energy per bit of 76 fJ/bit.

6.5 Towards DWDM Transceivers

6.5.1 DWDM Filter Bank Alignment

We fabricated a filter bank alignment test structure, shown in Figure 6-17. The inputs to the 8-bit DAC heater driver cell (Figure 6-17a) are externally set through scan-based serial I/O and drives the heaters of an 8-channel WDM filter bank (Figure 6-17). The heater driver cell successfully tuned all 8 ring filters to a 250 GHz grid on a 2 THz free spectral range (Figure 6-18) while consuming a total ring tuning power of 10.3 mW and achieving an overall filter tuning efficiency of 2.6 mW/nm (14.2 µW/GHz).

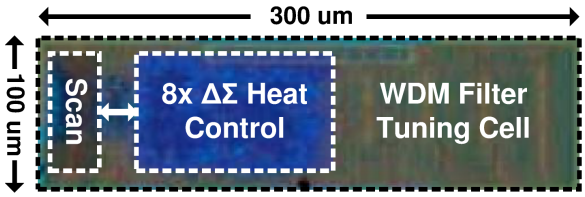

(a) Heater Driver Cell

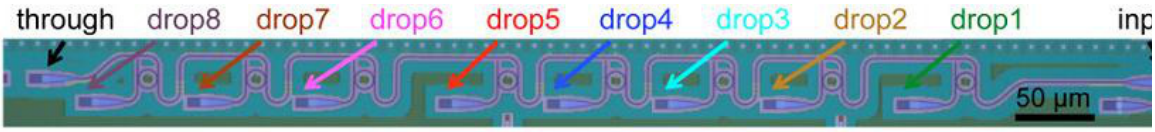

(b) Fabricated 8-channel WDM filter bank

Figure 6-17: Picture of the fabricated all-digital ΣΔ heater driver cell and wavelength transfer characteristics of a fabricated filter bank before and after tuning. Note that the size of the heater driver cell is oversized for ease-of-floorplan reasons. More than 60 heater drivers can actually fit in the 300 um-by-100 um region.

6.5.2 DWDM Transmitter Macro Demonstration

We use the wavelength-locked transmitter components to construct an $11 - \lambda$ DWDM transmitter macro in this platform, shown in Figure 6-19. Each modulator microring couples to the bus waveguide which spans across the entire macro. The radii of the microrings are stepped, such that each ring is resonant at a different wavelength, shown in Figure 6-20. Using the auto-locking circuitry, we demonstrate the

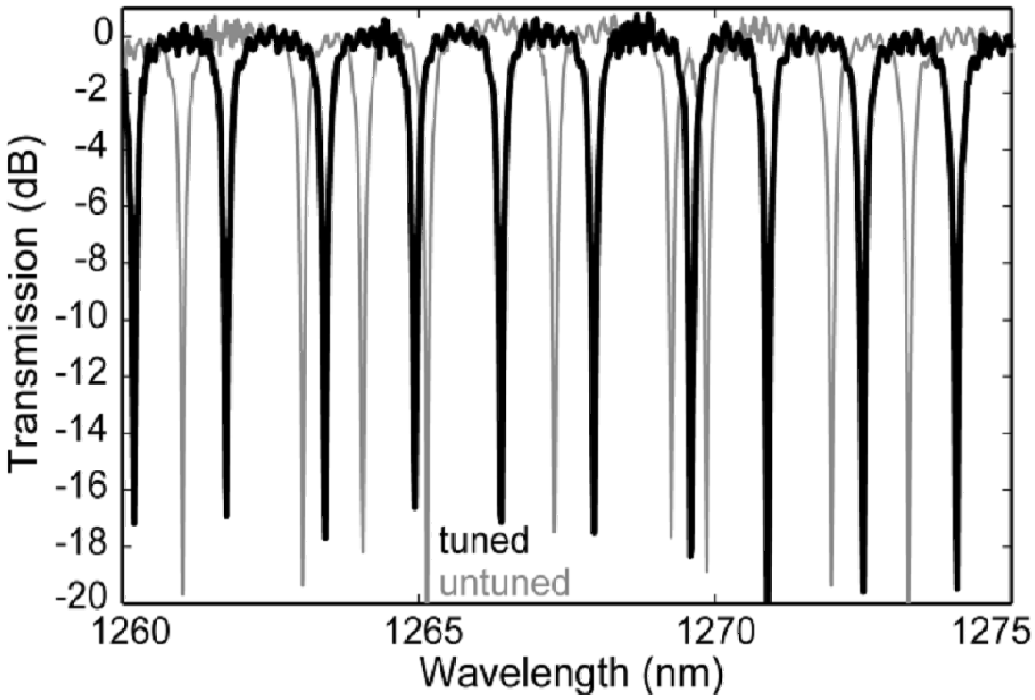

Figure 6-18: Wavelength transfer characteristics of the fabricated filter bank before and after tuning

functionality of each of the transmitter sites, one at a time individually, in Figure 6-21, demonstrating an open eye at 8 Gb/s on each of the 11 channels. The DWDM transmitter bank can achieve up to 11×8 Gb/s aggregate transmit data rate, potentially enabling 88 Gb/s of throughput on a single on-chip waveguide or optical fiber. Given a 3.06 THz FSR, the macro achieves a spectral efficiency of 0.029 bits/s/Hz. The full macro, which includes the digital backend, thermal tuning logic, and transceiver frontends, is 1450 µm × 200 µm, corresponding to a bandwidth density of 303 Gb/s/mm^2. Microrings are placed at a pitch of 128 µm but can be made as tight as approximately 20 µm before becoming fill density limited.

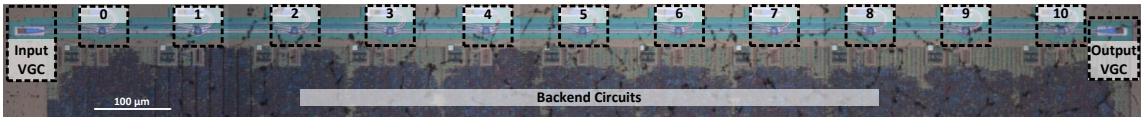

Figure 6-19: An 11λ DWDM transmitter.

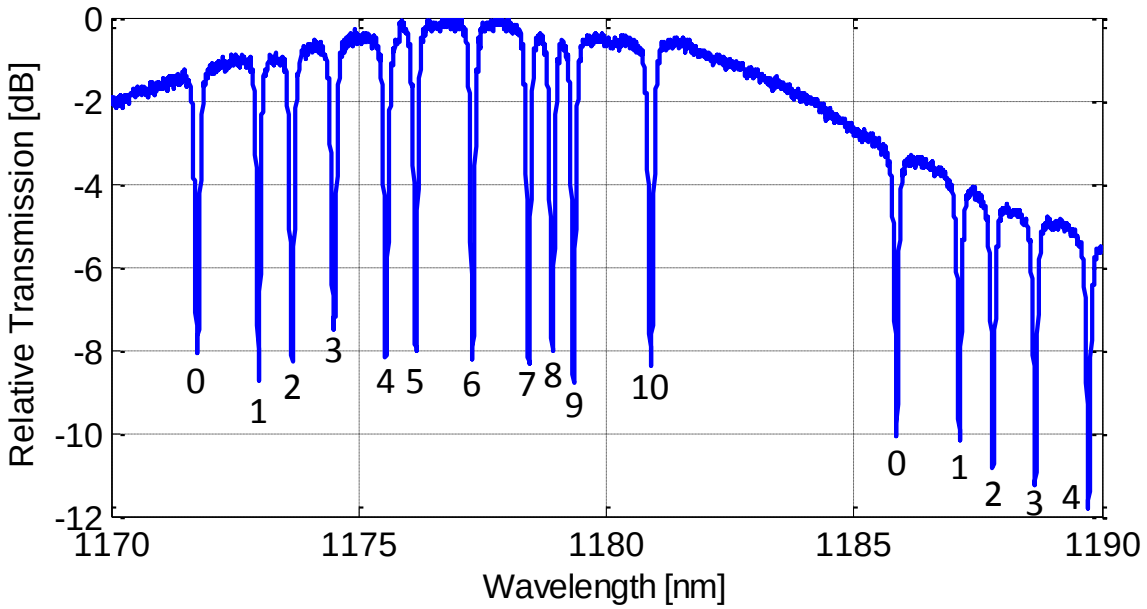

Figure 6-20: Optical transmisivity measurements of the DWDM transmitter, with the resonant wavelengths of each ring labeled.

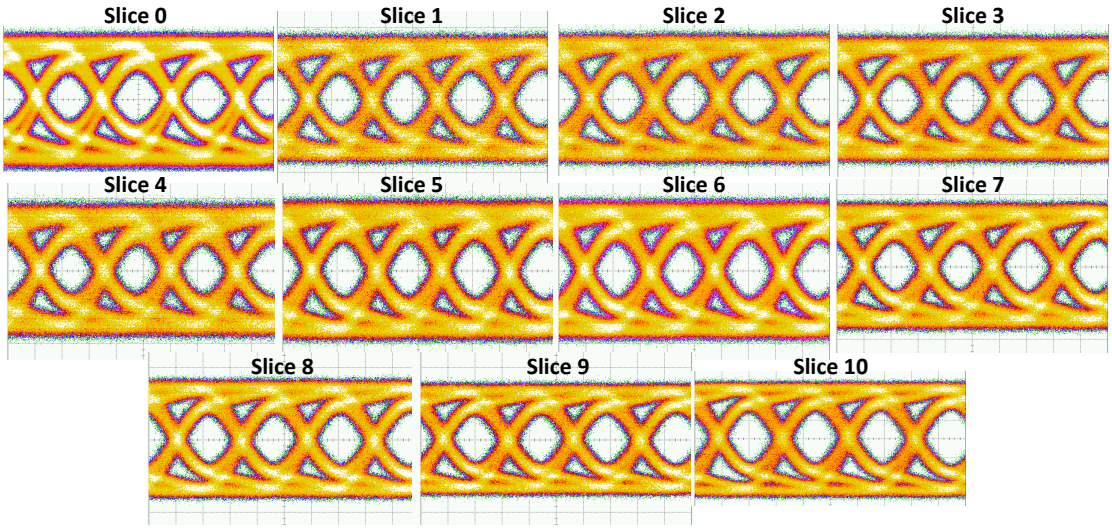

Figure 6-21: 8 Gb/s eye diagrams on each of the 11 channel slices that make up the DWDM transmit macro. The eyes are taken one at a time using a tunable laser.

6.6 Summary

If photonics are to adopted as a next-generation interconnect technology, we must demonstrate a seamless method for integrating it within a larger electrical system, ideally without introducing any additional cost. There are two major hurdles to this. The first is the challenge of photonic-electronic process compatibility and the second is the difficulty of tuning sensitive microring resonantor devices. In this Chapter, we have overcome these two hurdles through the demonstration of a zero-change optical transceiver platform and a bit-statistical thermal tuner suitable for general-purpose high-speed optical communications. Despite the zero-change constraint, the transceivers demonstrate a level of performance and energy-efficiency, along with the potential for DWDM, that makes the platform a compelling milestone for VLSI-compatible photonics. The possibilities for such a platform are quite limitless, and we shall explore one such application in the next chapter.

Chapter 7

Monolithically-Integrated Processor with Photonic Interconnects

At the beginning of this thesis, we set out to remove the barriers towards adoption of VLSI electro-optic systems. Up to this point, however, the work may still be considered incomplete; Chapters 5 and 6 demonstrate only the potential of the zero-change SOI and bulk platforms for supporting electro-optic transceivers, falling short of showing that these same platforms can yield a reliable electro-optic system of a reasonable scale. Our demonstrations of microring stabilization techniques still use artificially-generated data patterns and temperature stimuli, and have not yet been subject to real temperature fluctuations or data pattern workloads. To show that these techniques are applicable in a real system, this chapter combines the proposed designs and methodologies to demonstrate a true VLSI electro-optic system: a monolithically-integrated processor with optical interconnects. We focus on the components of the processor's optical memory system, which enables communication to main memory located an arbitrary distance away by optical fiber. We conclude this chapter with the demonstration of the entire system, which to the best of our knowledge, is the world's first processor chip with integrated optical interconnects.

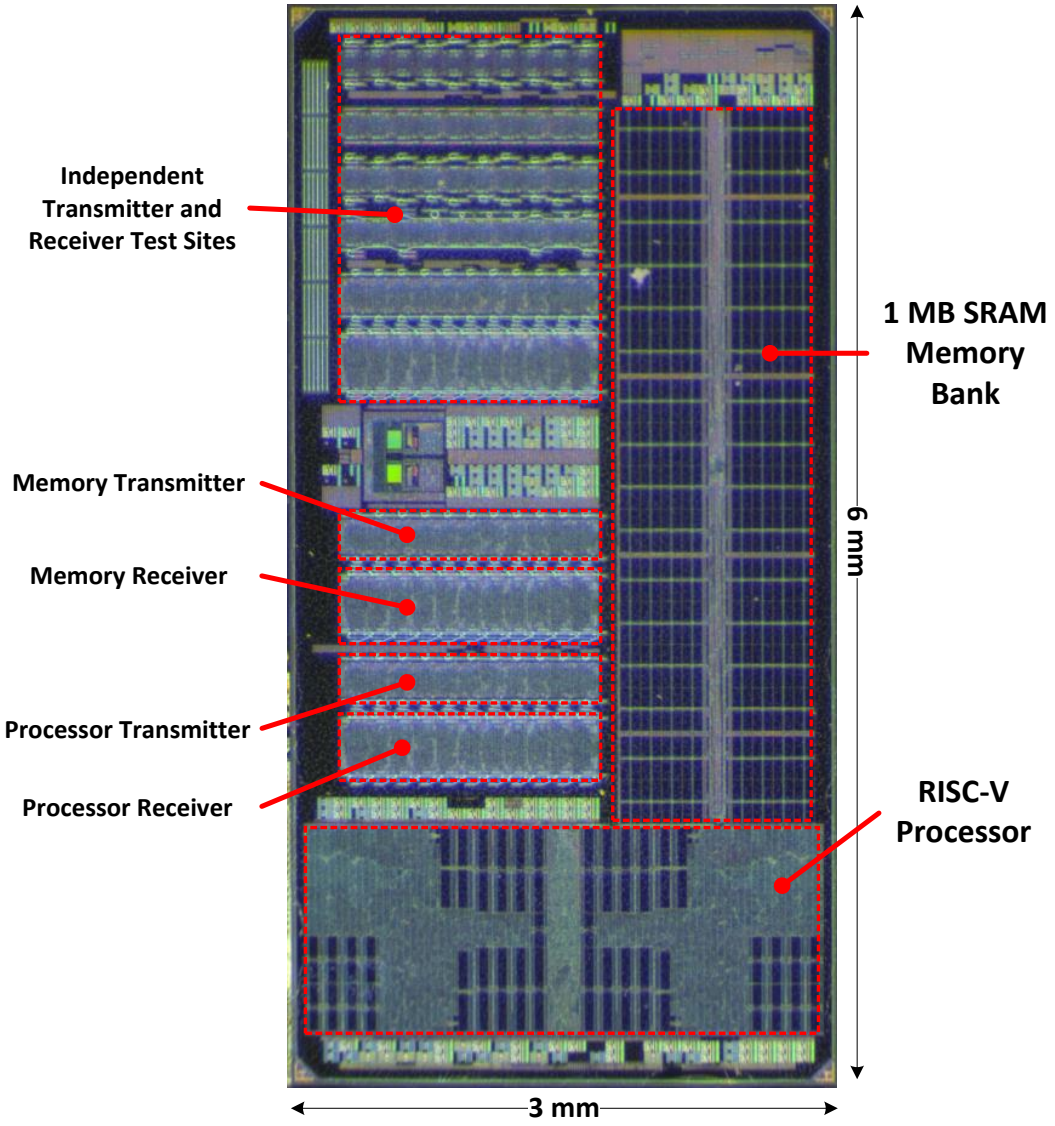

Figure 7-1: Die photo taken of the back-side of the 3 mm × 6 mm chip, with important features labeled. The optical transmitters and receivers belong to the processor and memory blocks enable optical chip-to-chip communication.

7.1 Processor Chip Overview

We show the processor chip in Figure 7-1. The dual-core processor executes the RISC-V [38, 90] instruction set architecture (ISA). The individual cores incorporate a 64-bit scalar core, floating-point unit, vector accelerator, and private L1 instruction and data caches [38]. The microprocessor runs at a maximum speed of 1.65 GHz,

and executes arbitrary compiled programs compiled to the RISC-V ISA, including modern operating systems such as Linux. A 1 MB bank of static random access memory (SRAM) serves as the processor's main memory. The processor accesses the main memory memory bank through an integrated memory controller. The the SRAM memory bank implements an emulated DRAM interface, modeling timing, bank conflicts, and states of a DDR3-like interface. Together, the memory controller and the memory bank implement a DRAM memory access interface.

We use two sets of low-speed electrical I/O interfaces to control and configure the chip. We communicate to the processor through the host target interface (HTIF), a full duplex parallel bus (16 wires each way) with valid/ready flow control and a forwarded clock with a roughly 100 MT/s maximum transaction rate. We configure the transceiver blocks through a 1-bit serial I/O interface. Both the memory controller and memory bank connect to integrated optical transceivers on the chip, enabling optical communication between them. In parallel, an additional set of on-chip wires directly link the ports of the memory controller to the ports of the memory bank. We may also configure the processor to not use the SRAM memory bank entirely; the processor can time-multiplex memory traffic onto the HTIF interface used for the processor control interface. This way, the processor can utilize an external bank of memory to execute programs that require more than 1 MB of storage, such as a full Linux kernel. To summarize, the processor chip has several modes of operation depending on the means to access memory:

- *Extended memory mode* – The processor uses an external bank of memory accessed through the HTIF control interface. This bypasses both the memory controller and the SRAM memory bank entirely.

- *Electrical bypass mode* – The memory controller communicates directly to the SRAM memory bank on the same chip using the on-chip electrical wires that directly link the two together.

- *Optical loopback mode* – The memory controller communicates to the SRAM memory bank optically using the integrated optical transceivers. The optical

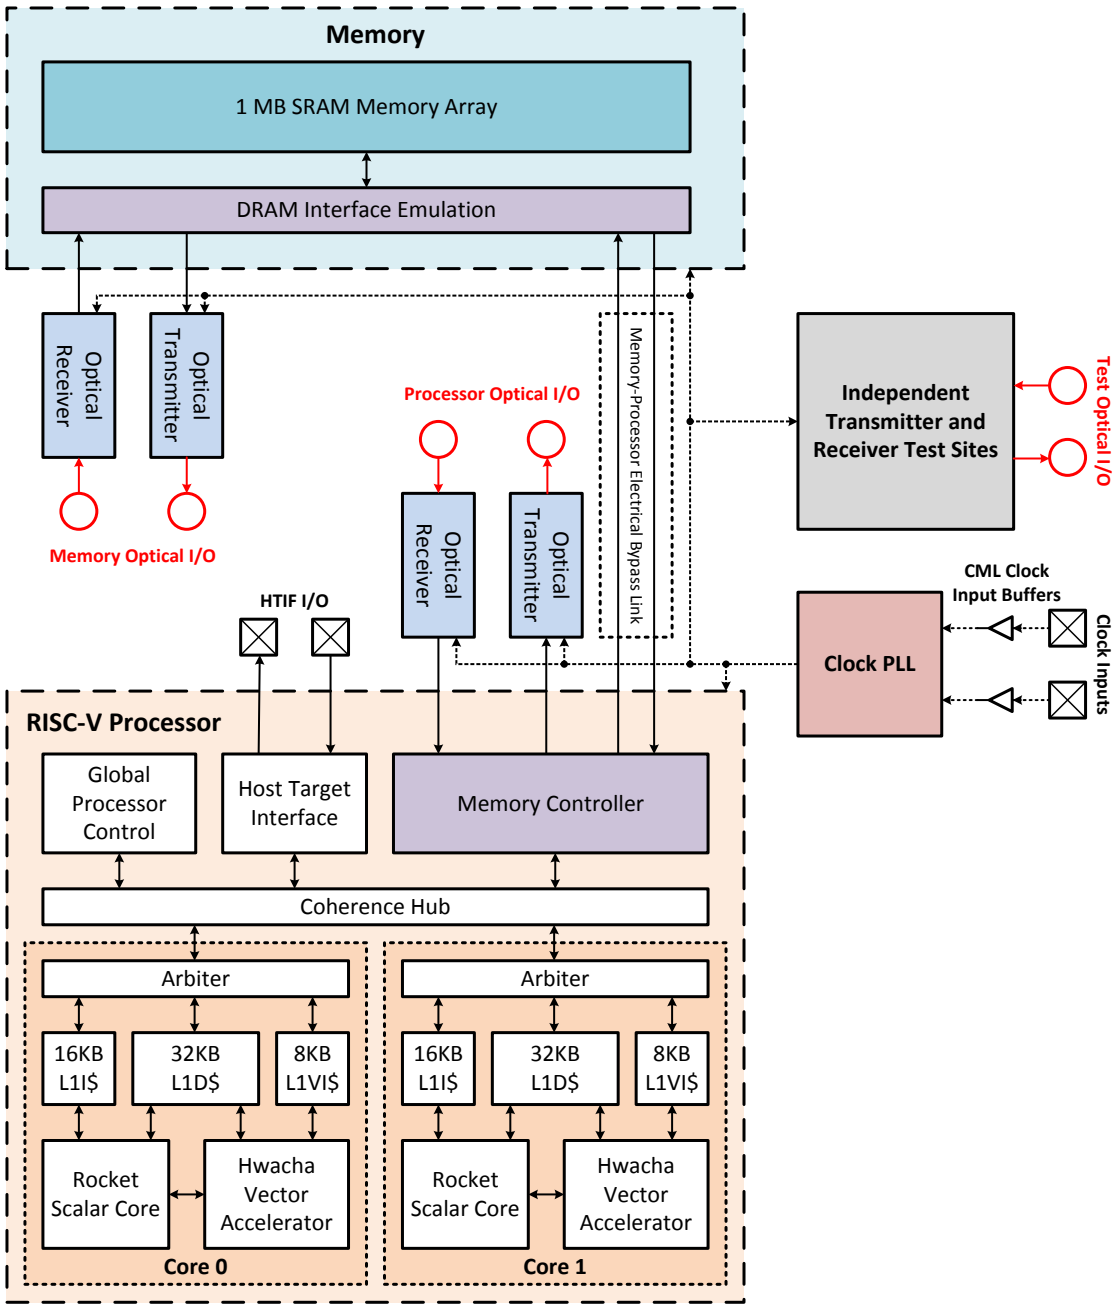

Figure 7-2: Functional diagram of the processor chip, showing electrical and optical I/O ports and the connectivity between parts.

input and output ports of the processor's optical transceivers are looped back to the optical output and input ports of the SRAM memory bank *on the same chip* through a set of optical fibers.

- *Optical processor/memory mode* – The memory controller communicates to the SRAM memory bank *on a different chip* using the integrated optical transceivers. This mode involves two separate chips: chip A is configured to be in *memory mode*, where we use it only like a memory chip, chip B is in *processor mode*, where it functions only as a processor chip. The optical I/O ports of the processor on chip B connect to the optical I/O ports of the SRAM memory bank on chip A through optical fibers.

The chip is fabricated in the thin-BOX SOI platform introduced in Chapter 6. As opposed to the full substrate removal process, we perform selective substrate removal on the chips after electrical packaging to only etch away the silicon substrate under regions with optical devices, shown in Figure 7-3. We retain the silicon handle over the microprocessor and memory, which dissipate the most power, to allow a heat sink to be contacted to them. As electronics are not sensitive to the substrate removal process, the definition of the removed and non-removed regions can be done very coarsely. In our process, we just used klapton tape applied manually by hand.

7.2 Optical Memory Scheduler

Here, we describe the design and implementation of the processor's optical memory scheduler, which enables the processor to communicate to memory located a distance away through an optical fibre. A memory scheduler takes read/write requests from the cache refill state-machine of the last-level processor cache. Using these requests, the memory controller schedules the read/write accesses to the connected DRAM memory, taking into account bank conflicts and data bus scheduling.

An electrical DRAM system typically shares the same set of wires for both DRAM read data and write data, implemented as a bidirectional DQ bus. This is architected around a pin-constrained design and we pay additional bus-reversal latencies whenever we transition from a read request to a write request or the other way around. A DRAM system with DWDM optical I/O will have dedicated read and write buses. This is not only because there is now sufficient bandwidth density for dedicated buses,

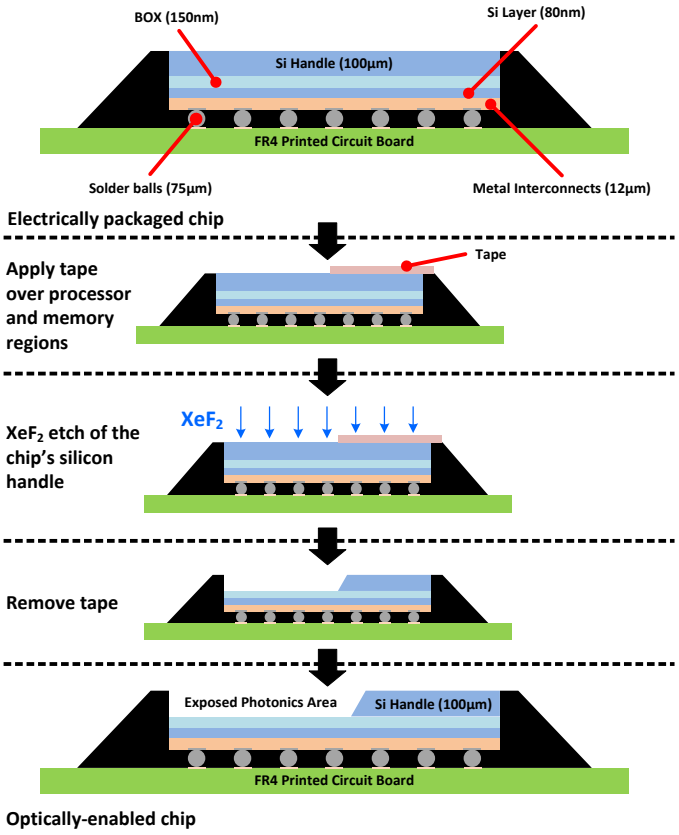

(a) Selective substrate removal processing steps

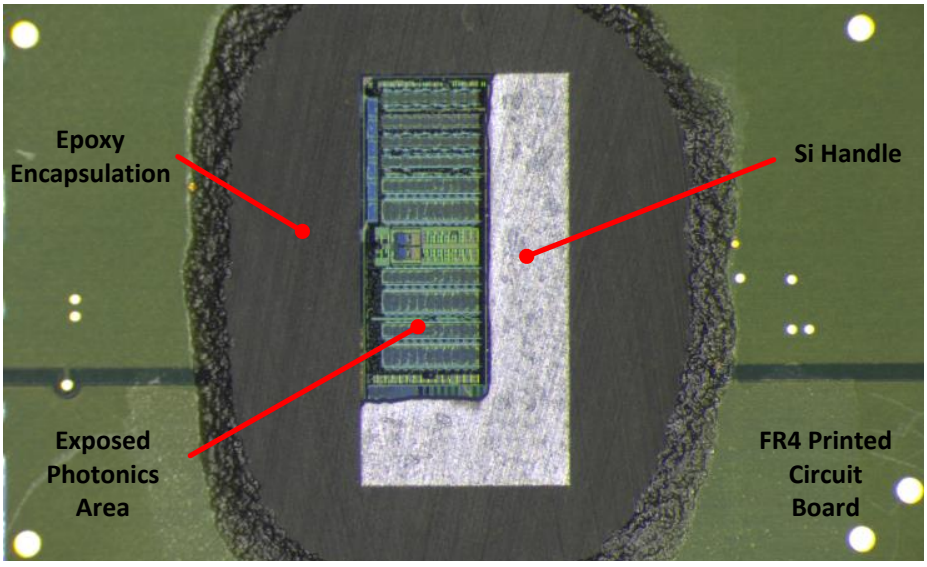

(b) 10 Gb/s bit-error-rate eye

Figure 7-3: The selective substrate removal process (a) and backside view of a packaged processor chip after selective substrate removal (b).

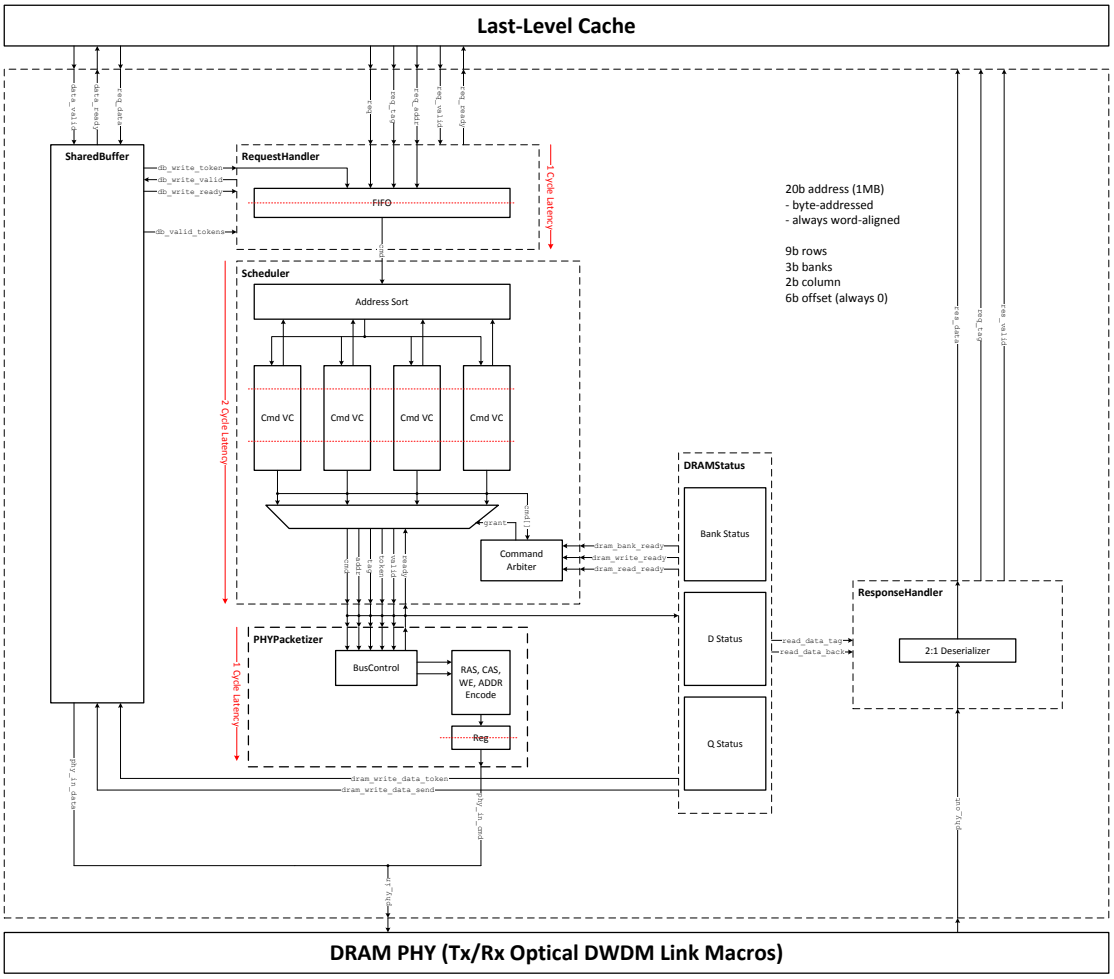

Figure 7-4: Optical memory controller architecture supporting simultaneous reads and writes and re-ordering of requests to different banks.

but also because it is difficult to reverse the directionality of an optical bus without an optical switch network.

We design a memory controller implementing dedicated read/write bus functionality, shown in Figure 7-4. The request and response handler modules present a valid/ready interface to the last level cache. The data buffer holds write data prior to being written to DRAM. The scheduler will re-order requests to different banks to maximize throughput, while guaranteeing read-after-write and write-after-read consistency. The DRAM status maintains bus and bank states and the PHY packetizer issues the DRAM commands to the optical PHY (physical layer optical links). To

minimize engineering time, we adopted a minimalistic approach to the design of the memory controller and omit additional features available in a commercial memory controller, such as low-power modes and DRAM refreshes.

Due to a lack of a real photonic DRAM part, we design a 1 MB bank of emulated DRAM (Figure 7-5 to allow testing of the full system. The DRAM emulator mimics behavior of a real DRAM with configurable DRAM timings. The storage cells are implemented using SRAM arrays available from the foundry.

7.3 DWDM-Enabled Optical PHY

A DRAM PHY abstracts away the physical realities of DRAM links and presents a clean synchronous digital interface to the memory controller. Here, we design a DRAM PHY that targets an optical physical medium, shown in Figure 7-6.

The DWDM transmit PHY is reponsible for transmitting optical data using 10 data wavelengths and 1 clock forwarding wavelength. The PHY is divided into 11 transmit wavelength slices. Each slice consists of a ring modulator, its driver, thermal tuning circuitry, all from the components described in Chapter 6. In addition to the data-path components, The PHY includes PRBS31 data generators and can be switched into a link test mode. The PHY presents an 80-bit interface to the memory scheduler or the DRAM emulator and has a bandwidth of 80-bits per input clock cycle (memory scheduler or DRAM emulator clock cycle). A variable ratio serializer serializes the 80-bits down to the 10 data wavelengths for a nominal 8-to-1 serialization ratio. The choice of the number of wavelengths to use can be configured by changing the serialization ratio. For example, if we use only one wavelength, the serializer performs an 80-to-1 serialization. Naturally, the input clock rate must also be slowed appropriately to rate-match the lower total bandwidth.

The receive PHY is responsible for outputting 80 bits per output cycle of optically received data using 10 data wavelengths and 1 clock wavelength. Each wavelength slice in the receive PHY consists of a receive filter microring and two photodetectors coupled to the drop-port of the filter. The photodetector that is strongly coupled (and

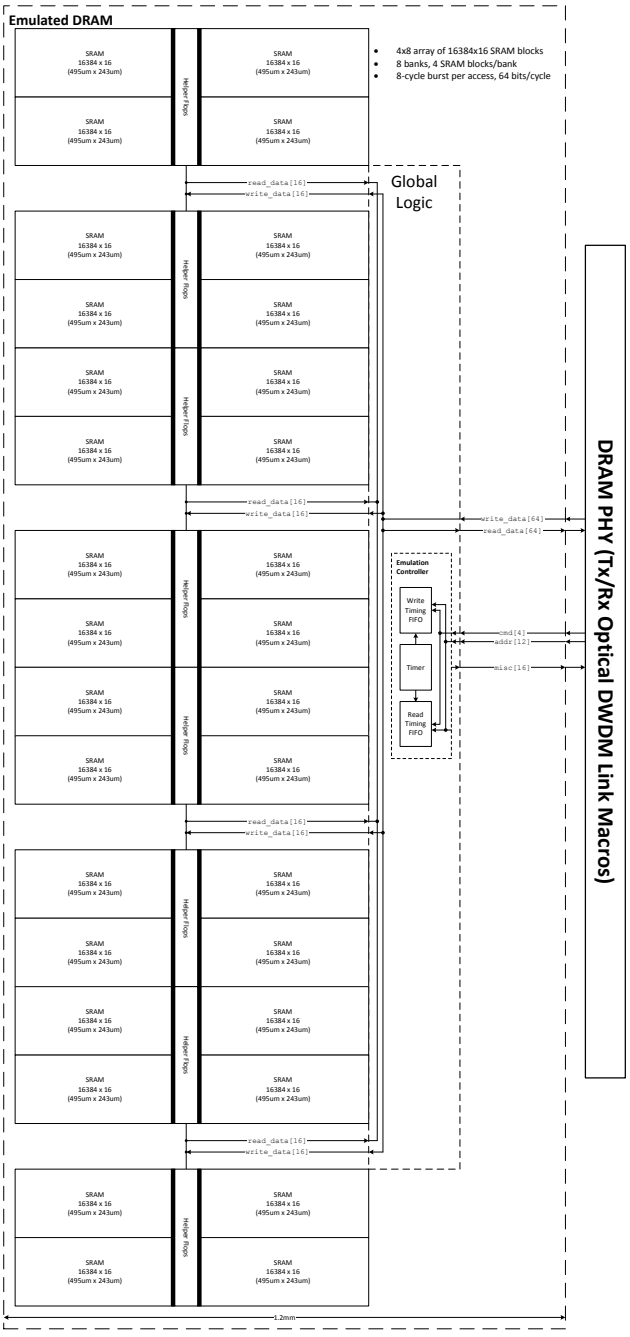

Figure 7-5: Architecture of the DRAM emulator.

receives more drop-port power) connects to the data receiver. The photodetector that is weakly coupled is used by the thermal tuner for stabilizing the filter ring's resonance. In addition, the ring filter may be bypassed completely through the bypass coupler

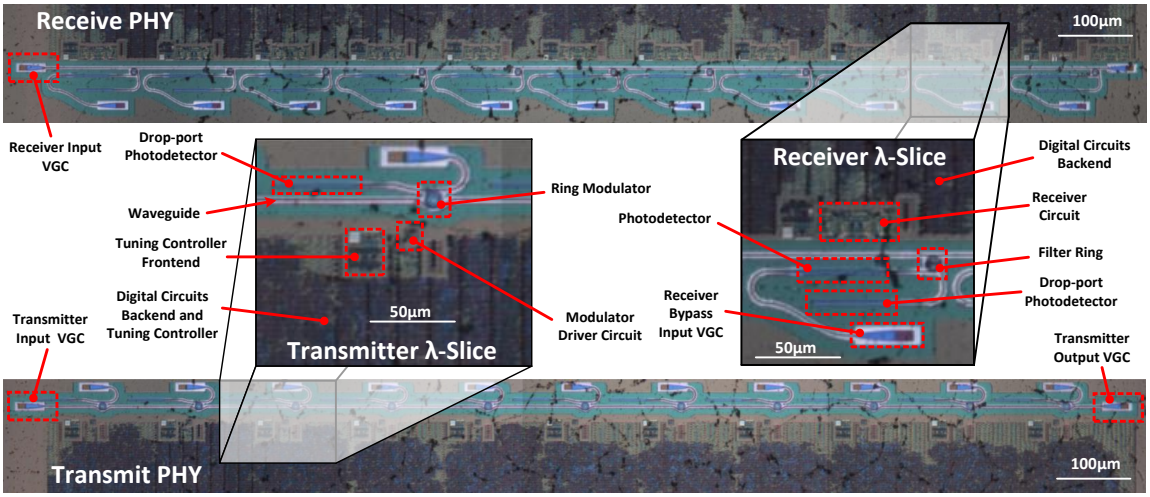

Figure 7-6: DWDM Optical PHY for both transmitters and receivers

port, which enables light to be directly coupled into the data receiver photodetector. The PHY has a dedicated clocking subsystem designed to receive and distribute the clock forwarded by the transmitter. Like its transmit counterpart, the receive PHY can also be configured to use an arbitrary number of wavelengths and supports in self-test or real data modes.

7.4 Optically-Connected Memory Demonstration

We setup the processor's chip-to-chip optical memory system demonstration in Figure 7-8 and the block diagram is shown in Figure 7-7. Here, we configure one chip to be in *processor mode*, allowing it to communicate optically to the 1 MB memory array on a second chip in *memory mode*, located an arbitrary distance away. The *processor mode* chip sends requests (a *read* or *write*), the memory address (location in memory to *read* or *write*), and write data (for *write* requests) via the microprocessor to memory (P2M) link. The memory to microprocessor (M2P) link returns read data for *read* requests. For this demonstration, we use only a single laser wavelength (as we currently do not have equipment to multiplex different wavelengths together).

For both links, the laser light first couples into a transmit PHY; laser light arriving in a single-mode (SM) fiber couples into the bus waveguide through a vertical grating

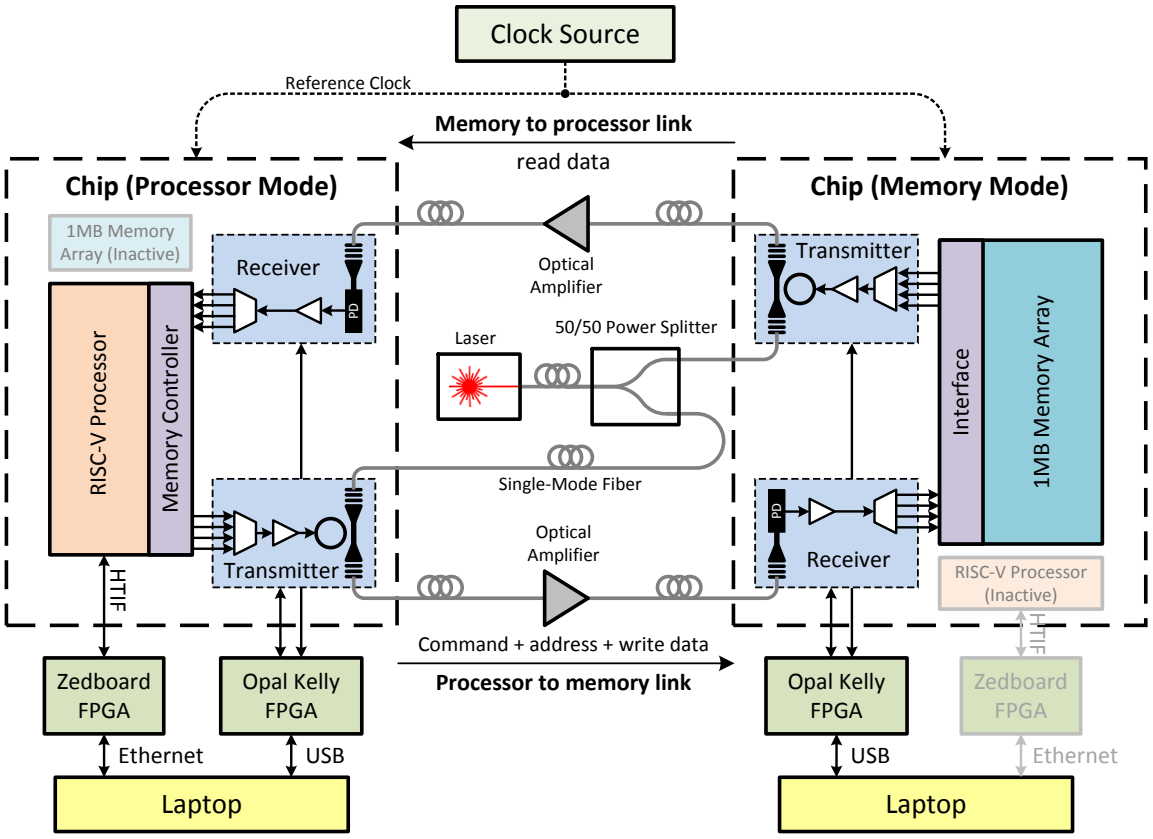

Figure 7-7: Block diagram of the optical memory system demonstration. The system uses one chip acting as the processor and the other acting as *memory*, connected together by a full duplex optical link with a round-trip distance of 40 m by fiber.

coupler (VGC). The optical modulator, driven by circuits, modulates light in the waveguide and imprints it with on-off keyed binary data from the source. The light then exits the chip through a second vertical grating into an SM fiber bound for the other chip. Once there, the light couples into the receive PHY through a VGC, illuminating a photodetector connected to a data receiver. We use a single 1183 nm continuous wave (CW) off-chip solid-state laser as the laser source and split its output power 50/50 to share it across both the P2M and M2P links. Since we use only one wavelength, we bypass the filter ring of the receive PHY and couple directly into photodetector. To overcome the 4–6 dB coupling losses on each VGC and the low photodetector sensitivity, we insert an optical amplifier, which provides about 8 dB of gain, to get sufficient optical power at the receiver to resolve the signal. With a

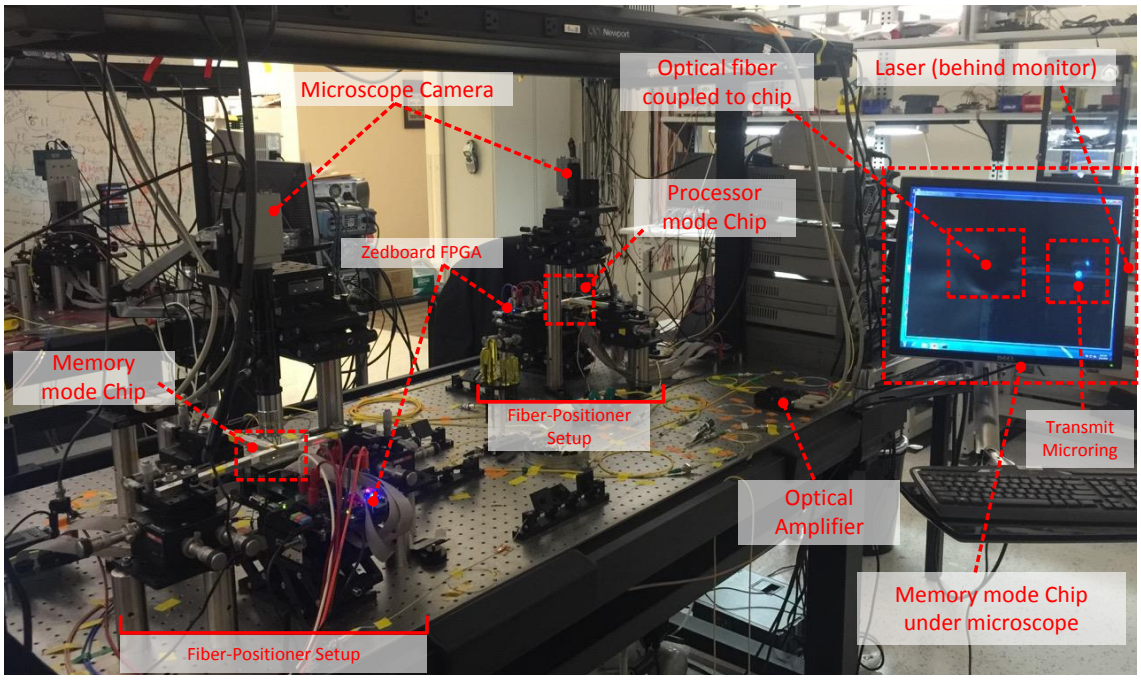

Figure 7-8: Test setup of the demonstration. One chip serves as memory and is configured in *memory mode* and the other chip configured in *processor mode*. Each chip sits on a microscope and 3-axis fiber positioner setup, which are used to position optical fibers over the optical ports (vertical couplers) of both chips. In this demo, we require 3 fibers coupled into each chip.

single wavelength, both the transmit and receive DWDM PHYs operate in a 80-to-1 serialization/deserialization ratio mode. As such, the input clock, and hence the processor and DRAM emulator clocks, are divided down by a factor of 80 from the link data rate.

We demonstrate a fully working optical processor-to-memory system in Figure 7-9, which shows the processor terminal output for two test programs we run. For each test program, the FPGA connected to the processor chip uses the HTIF interface to directly access the processor's memory controller. In doing so, the FPGA writes the program's contents (which flows to the memory through the P2M optical link) into the memory bank on the memory chip. Once program loading is complete, the FPGA resets the processor. When reset, processor begins to read instructions from the memory starting at an initial program counter value and begins executing these

```
root@zynq:~/eos22-1mb# ./fesvr-zedboard-head.1MB +divisor=1
CPU reset complete
uncore slowio divisor=1, hold=1
host_clk frequency = 15.63 MHz
cpu_clk frequency  = 31.27 MHz
hello world with photonics!
```

(a) The optically-connected processor executing "Hello World"

```
root@zynq:~/eos22-1mb# ./fesvr-zedboard-head.1MB +divisor=1 +hold=1 +memt
CPU reset complete
uncore slowio divisor=1, hold=1
host_clk frequency = 15.63 MHz
cpu_clk frequency  = 31.27 MHz
running memory test for 1 MB...
wrote 5ccdb01c74e3f122 7d4b595a6b0593c4 7dc10af410c30bbc 00b900d56a24fc57
read  5ccdb01c74e3f122 7d4b595a6b0593c4 7dc10af410c30bbc 00b900d56a24fc57
done memory test BER=0/8388608=0.000000000...
-------------------------------------------------------------
STREAM version $Revision: 5.10 $
-------------------------------------------------------------
This system uses 8 bytes per array element.
-------------------------------------------------------------
Array size = 32768 (elements), Offset = 0 (elements)
Memory per array = 0.2 MiB (= 0.0 GiB).
Total memory required = 0.8 MiB (= 0.0 GiB).
Each kernel will be executed 10 times.
 The *best* time for each kernel (excluding the first iteration)
 will be used to compute the reported bandwidth.
-------------------------------------------------------------
Your clock granularity/precision appears to be 1 microseconds.
Each test below will take on the order of 815 microseconds.
   (= 815 clock ticks)
Increase the size of the arrays if this shows that
you are not getting at least 20 clock ticks per test.
-------------------------------------------------------------
WARNING -- The above is only a rough guideline.
For best results, please be sure you know the
precision of your system timer.
-------------------------------------------------------------
Function    Best Rate MB/s  Avg time     Min time     Max time
Copy:            624.8     0.000843     0.000839     0.000852
Scale:           541.3     0.000972     0.000969     0.000982
Add:             571.7     0.001381     0.001376     0.001396
Triad:           572.8     0.001382     0.001373     0.001396
-------------------------------------------------------------
Solution Validates: avg error less than 1.000000e-13 on all three arrays
-------------------------------------------------------------
```

(b) The optically-connected processor executing a memory test and the STREAM memory benchmark

Figure 7-9: The processor executing the "Hello World" (a) and STREAM memory benchmark (b) programs. In both demonstrations, the processor uses the optical links to communicate to main memory to both fetch the instruction data of the program and for memory requests made by the program.

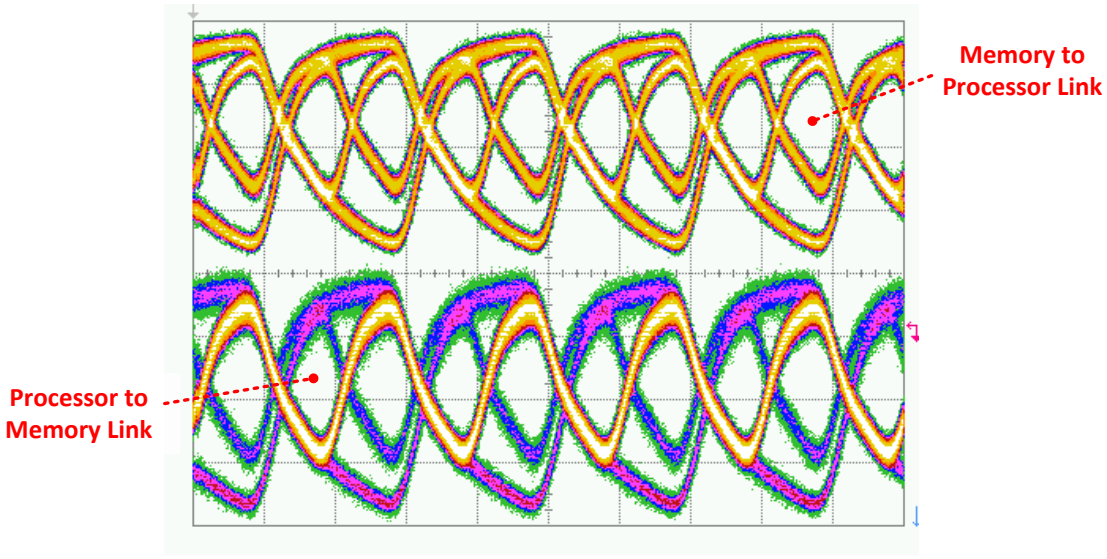

Figure 7-10: Optical eye diagrams from the P2M and M2P optical links during the execution of the STREAM benchmark.

instructions. During the execution of the program, the processor will perform accesses to memory to both read additional instructions and to read and write memory data necessary for program execution. In all demonstrations, we run both P2M and M2P optical links at 2.5 Gb/s, corresponding to a processor and DRAM emulator clock frequency of 31.25 MHz.

Figure 7-9a shows the successful execution of a "Hello World" program, which tests the basic functionality of the processor through the optical link. Figure 7-9b shows the successful execution of a memory test and a STREAM memory benchmark application [48]. The memory test simply writes to and reads from all addresses of the optically-connected 1 MB memory array, verifying that there are no bit-errors. The STREAM memory benchmark application, compiled to the RISC-V ISA, tests the memory bandwidth of the system and performs a check to make sure all memory locations are correct. Figure 7-10 shows the transmitter eye diagrams for both the P2M and M2P links simultaneously as STREAM was being executed. We note that the system is quite stable and can execute an arbitary number of programs. The time to failure is between 30 minutes to an hour, limited by the stability of the 3-axis mechanical fiber positioner stages, which drift away from the desired spot over

time, causing the fibers to become misaligned from the grating couplers and adding insertion loss. A properly packaged part with fixed fiber locations will remove this limitation.

Chapter 8

Final Thoughts and Conclusions

Almost five years ago, when I just started work on my PhD, we proposed a photonically interconnected DRAM memory system (PIDRAM). This system leveraged monolithically-integrated silicon photonic transceivers to overcome the bandwidth density and energy-efficiency bottlenecks of chip-to-chip memory I/O, creating a scalable memory interconnect to support an era of aggressive many-core computing [6]. At that point in time, however, integrated silicon photonics was still in its infancy; while architectural proposals leveraging the technology proliferated, actual practical demonstrations of the technology in silicon were rare or nonexistent. With a goal to eventually build a PIDRAM memory system, I embarked on this five-year long journey (Figure 8-1) to overcome the hurdles of electronic-photonic integration.

The work in thesis resolved the two gating issues that currently prevent the realization of VLSI electronic-photonic systems: microring resonator stability and photonic integration. To address the former issue, this thesis developed a model of the thermal variations experienced by microring resonators – the key building blocks of integrated DWDM systems. Using the insights gained from this model, I proposed bit-statistical tuning and self-heating cancellation as techniques necessary to properly tune microring resonances, enabling them to be used in generic, unencoded data links. To address the latter issue, I presented a low-cost polysilicon-only bulk photonics platform and a zero-change SOI platform. In both platforms, we adopted a minimal-change approach to integrate photonics that aimed to create the least amount of disruption

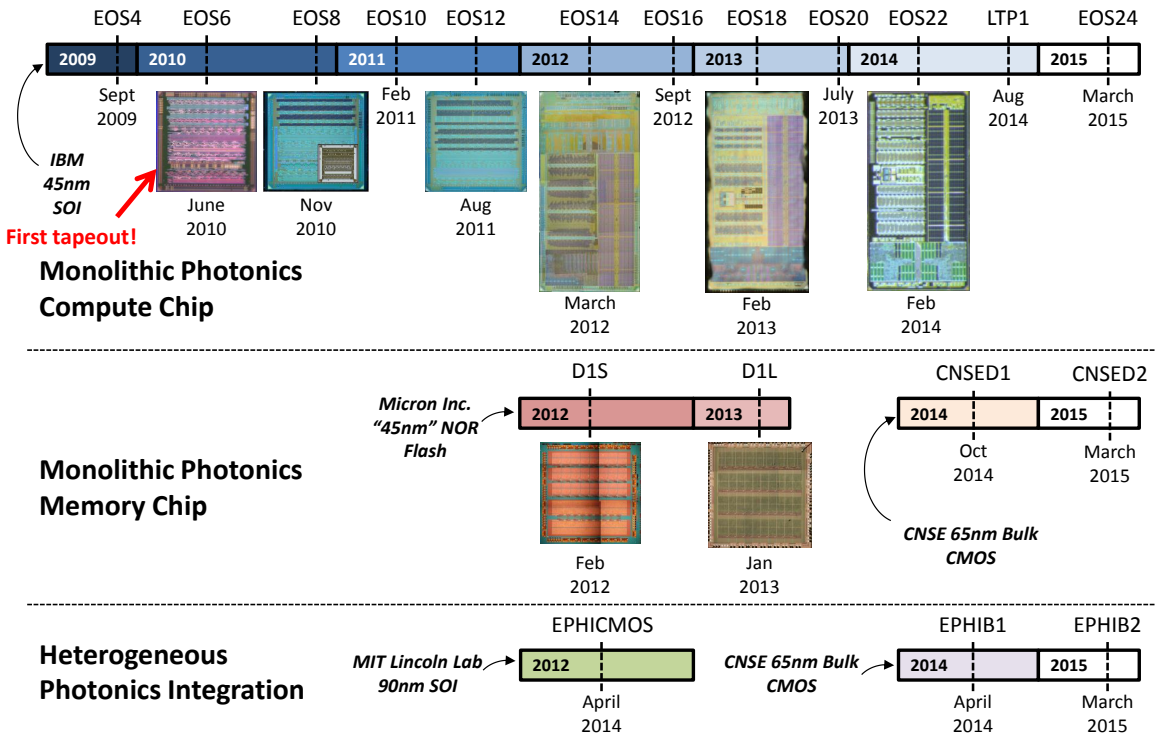

Figure 8-1: A five year long journey of electronic-photonic integration across multiple different photonics integration platforms. Chip die photos are shown for the chips that marked important project milestones. A total of 15 tapeouts were done over the course of this PhD.

to existing electronics, overcoming the challenges associated with electronic-photonic process compatibility and potentially accelerating photonics adoption by the mainstream electronics industry.

I concluded the work in this thesis with the demonstration of a processor chip with photonically interconnected memory, which is the first demonstration of a photonically interconnected processor chip and perhaps the first demonstration of a silicon photonic VLSI system. In doing so, this thesis answers the question of whether electronics and photonics can coexist alongside each other as a system on a chip, and the answer is a resounding *yes*. When compared to the technology specifications we projected in our initial architectural study almost five years ago (Figure 8-2), the platforms demonstrated in this thesis are not just functional, but also competitive in performance and efficiency, despite using only first or second generation circuits and

Metric	[Beamer ISCA 2010] Conservative Estimates	45nm SOI Platform	Bulk Photonics Platform*
Waveguide Loss	4 dB/cm	3.7 dB/cm	10.5 dB/cm
Vertical Coupler Loss	1 dB	1 dB	3 dB
Tx Data Rate	10 Gb/s	8 Gb/s	5 Gb/s
Tx Energy Per Bit	120 fJ/b	30 fJ/b	350 fJ/b
Rx Data Rate	10 Gb/s	10 Gb/s	5 Gb/s
Rx Energy Per Bit	80 fJ/b	297 fJ/b	1700 fJ/b
Rx Sensitivity	10 µA	8.3 µA	36 µA
PD Responsivity	0.9 A/W	0.1 A/W	0.2 A/W
Thermal Tuning Efficiency	1.6 µW/GHz	3.8 µW/GHz	10 µW/GHz

Figure 8-2: Comparison between the technology projections used in [6] and the demonstrated platforms of Chapters 5 and Chapter 6. Note that the bulk platform is on a significantly older transistors node than what was assumed in [6], hence the big difference in energy efficiency between them.

devices. Interestingly, this means that even the current platforms are not too far off from the point where a photonic memory system could gain an order of magnitude advantage over conventional electrical memory systems. Nevertheless, such a memory system is but one of a virtually limitless number of emerging photonic applications. Therefore, it is my hope that this thesis serves as a catalyst for an era of integrated electronic-photonic systems.

Bibliography

[1] A 32nm 3.1 billion transistor 12-wide-issue Itanium processor for mission-critical servers. February 2011.

[2] P. Amberg, E. Chang, F. Liu, J. Lexau, Xuezhe Zheng, Guoliang Li, I Shubin, J.E. Cunningham, AV. Krishnamoorthy, and Ron Ho. A sub-400 fj/bit thermal tuner for optical resonant ring modulators in 40 nm CMOS. In *Solid State Circuits Conference (A-SSCC), 2012 IEEE Asian*, pages 29–32, Nov 2012.

[3] 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, Sept 2012.

[4] S. Assefa, S. Shank, W. Green, M. Khater, E. Kiewra, C. Reinholm, S. Kamlapurkar, A Rylyakov, C. Schow, F. Horst, Huapu Pan, T. Topuria, P. Rice, D.M. Gill, J. Rosenberg, T. Barwicz, Min Yang, J. Proesel, J. Hofrichter, B. Offrein, Xiaoxiong Gu, W. Haensch, J. Ellis-Monaghan, and Y. Vlasov. A 90nm CMOS integrated nano-photonics technology for 25Gbps WDM optical communications applications. In *Electron Devices Meeting (IEDM), 2012 IEEE International*, pages 33.8.1–33.8.3, Dec 2012.

[5] C Batten, A Joshi, J Orcutt, C Holzwarth, M Popovic, J Hoyt, F Kartner, R Ram, V Stojanovic, and K Asanovic. Building manycore processor-to-DRAM networks with monolithic CMOS silicon photonics. *Micro, IEEE*, PP(99):1, 2009.

[6] 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 *International Symposium on Computer Architecture*, pages 129–140, New York, NY, USA, 2010. ACM.

[7] Wim Bogaerts, Peter De Heyn, Thomas Van Vaerenbergh, Katrien De Vos, Shankar Kumar Selvaraja, Tom Claes, Pieter Dumon, Peter Bienstman, Dries Van Thourhout, and Roel Baets. Silicon microring resonators. *Laser & Photonics Reviews*, 6(1):47–73, 2012.

[8] J.F. Buckwalter, Xuezhe Zheng, Guoliang Li, K. Raj, and AV. 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, June 2012.

[9] Hyunil Byun, Jinkwon Bok, Kwansik Cho, Keunyeong Cho, Hanmei Choi, Jinyong Choi, Sanghun Choi, Sangdeuk Han, Seokyong Hong, Seokhun Hyun, T. J. Jeong, Ho-Chul Ji, In-Sung Joe, Beomseok Kim, Donghyun Kim, Junghye Kim, Jeong-Kyoum Kim, Kiho Kim, Seong-Gu Kim, Duanhua Kong, Bongjin Kuh, Hyuckjoon Kwon, Beomsuk Lee, Hocheol Lee, Kwanghyun Lee, Shinyoung Lee, Kyoungwon Na, Jeongsik Nam, Amir Nejadmalayeri, Yongsang Park, Sunil Parmar, Junghyung Pyo, Dongjae Shin, Joonghan Shin, Yong hwack Shin, Sung-Dong Suh, Honggoo Yoon, Yoondong Park, Junghwan Choi, Kyoung-Ho Ha, and Gitae Jeong. Bulk-si photonics technology for dram interface. *Photon. Res.*, 2(3):A25–A33, Jun 2014.

[10] Jingnan Cai, Yasuhiko Ishikawa, and Kazumi Wada. Strain induced bandgap and refractive index variation of silicon. *Optics express*, 21(6):7162–7170, 2013.

[11] Rodolfo E. Camacho-Aguilera, Yan Cai, Neil Patel, Jonathan T. Bessette, Marco Romagnoli, Lionel C. Kimerling, and Jurgen Michel. An electrically pumped germanium laser. *Optics Express*, 20(10):11316–11320, May 2012.

[12] E. Chang, F. Liu, P. Amberg, J. Lexau, and Ho. Efficient techniques for canceling transceiver noise. In *VLSI Design, Automation, and Test (VLSI-DAT), 2013 International Symposium on*, pages 1–4, April 2013.

[13] Yu-Hsin Chen. *Investigating thermal dependence on monolithically-integrated photonic interconnects*. PhD thesis, Massachusetts Institute of Technology, 2013.

[14] Yu-Hsin Chen, Chen Sun, and Vladimir Stojanovic. Scalable electrical-optical thermal simulator for multicores with optical interconnects. In *Optical Interconnects Conference*, May 2013.

[15] P. Dainesi et al. CMOS compatible fully integrated machŰzehnder interferometer in SOI technology. *IEEE Photonics Technology Letters*, 12:660–662, June 2000.

[16] Po Dong, Wei Qian, Hong Liang, Roshanak Shafiiha, Dazeng Feng, Guoliang Li, John E. Cunningham, Ashok V. Krishnamoorthy, and Mehdi Asghari. Thermally tunable silicon racetrack resonators with ultralow tuning power. *Opt. Express*, 18(19):20298–20304, September 2010.

[17] P. Dumon, W. Bogaerts, D. Van Thourhout, D. Taillaert, R. Baets, J. Wouters, S. Beckx, and P. Jaenen. Compact wavelength router based on a silicon-on-insulator arrayed waveguide grating pigtailed to a fiber array. *Opt. Express*, 14(2):664–669, Jan 2006.

[18] 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, Sept 2011.

[19] M. Georgas, B.R. Moss, C. Sun, J. Shainline, J.S. Orcutt, M. Wade, Y.-H. Chen, K. Nammari, J.C. Leu, A Srinivasan, R.J. Ram, M.A Popovic, and V. Stojanovic. A monolithically-integrated optical transmitter and receiver in a zero-change 45nm SOI process. In *VLSI Circuits Digest of Technical Papers, 2014 Symposium on*, pages 58–59, June 2014.

[20] 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, July 2012.

[21] Biswajeet Guha, Bernardo BC Kyotoku, and Michal Lipson. Cmos-compatible athermal silicon microring resonators. *Optics Express*, 18(4):3487–3493, 2010.

[22] C. Holzwarth et al. Localized substrate removal technique enabling strong-confinement microphotonics in bulk Si CMOS processes. May 2008.

[23] J. Howard *et al.* A 48-core IA-32 message-passing processor with DVFS in 45nm CMOS. In *ISSCC*, pages 108 –109, Feb 2010.

[24] Wei Huang, Shougata Ghosh, Sivakumar Velusamy, Karthik Sankaranarayanan, Kevin Skadron, and Mircea R Stan. Hotspot: A compact thermal modeling methodology for early-stage vlsi design. *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, 14(5):501–513, 2006.

[25] Intel hybrid silicon laser. Website, `http://techresearch.intel.com/ProjectDetails.aspx?Id=149`.

[26] Intel Xeon Phi product comparison. Website, `http://ark.intel.com/compare/71992,75801`.

[27] ITRS 2010 update, assembly and packaging chapter. Website, `http://www.itrs.net/Links/2010ITRS/Home2010.htm`.

[28] Ho-Chul Ji, K.H. Ha, I S. Joe, S.G. Kim, K. W. Na, D.J. Shin, S.D. Suh, Y.D. Park, and C.H. Chung. Optical interface platform for DRAM integration. In *Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011 and the National Fiber Optic Engineers Conference*, pages 1–3, March 2011.

[29] Ajay Joshi, Christopher Batten, Yong-Jin Kwon, Scott Beamer, Imran Shamim, Krste Asanovic, and Vladimir Stojanovic. Silicon-photonic clos networks for global on-chip communication. In *Proceedings of the 2009 3rd ACM/IEEE International Symposium on Networks-on-Chip*, NOCS '09, pages 124–133. IEEE Computer Society, 2009.

[30] Ron Kalla, Balaram Sinharoy, William J Starke, and Michael Floyd. Power7: Ibm's next-generation server processor. *IEEE micro*, 30(2):7–15, 2010.

[31] Ting-Kuo Kang. Evidence for silicon bandgap narrowing in uniaxially strained mosfets subjected to tensile and compressive stress. *Electron Device Letters, IEEE*, 33(6):770–772, 2012.

[32] 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 *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2011 IEEE International*, pages 496–498, Feb 2011.

[33] Lionel Kimerling. Silicon microphotonics. *Applied Surface Science*, 159:8–13, June 2000.

[34] DBM Klaassen, JW Slotboom, and HC De Graaff. Unified apparent bandgap narrowing in n-and p-type silicon. *Solid-State Electronics*, 35(2):125–129, 1992.

[35] George Kurian et al. ATAC: A 1000-core cache-coherent processor with on-chip optical network. *Parallel Architectures and Compilation Techniques*, September 2010.

[36] George Kurian, Chen Sun, C-HO Chen, Jason E Miller, Jürgen Michel, Lan Wei, Dimitri A Antoniadis, Li-Shiuan Peh, Lionel Kimerling, Vladimir Stojanovic, et al. Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads. In *Parallel & Distributed Processing Symposium (IPDPS), 2012 IEEE 26th International*, pages 1117–1130. IEEE, 2012.

[37] Sungjae Lee, Basanth Jagannathan, Shreesh Narasimha, Anthony Chou, Noah Zamdmer, Jim Johnson, Richard Williams, Lawrence Wagner, Jonghae Kim, J-O Plouchart, et al. Record rf performance of 45-nm soi cmos technology. In *Electron Devices Meeting, 2007. IEDM 2007. IEEE International*, pages 255–258. IEEE, 2007.

[38] Yunsup Lee, Andrew Waterman, Rimas Avizienis, Henry Cook, Chen Sun, Vladimir Stojanovic, and Krste Asanovic. A 45nm 1.3 ghz 16.7 double-precision gflops/w risc-v processor with vector accelerators. In *European Solid State Circuits Conference (ESSCIRC), ESSCIRC 2014-40th*, pages 199–202, 2014.

[39] J. Leu and V. Stojanovic. Injection-locked clock receiver for monolithic optical link in 45nm SOI. In *Solid State Circuits Conference (A-SSCC), 2011 IEEE Asian*, pages 149–152, Nov 2011.

[40] Cheng Li, Rui Bai, A Shafik, E.Z. Tabasy, Binhao Wang, Geng Tang, Chao Ma, Chin-Hui Chen, Zhen Peng, M. Fiorentino, R.G. Beausoleil, P. Chiang, and S. Palermo. Silicon photonic transceiver circuits with microring resonator bias-based wavelength stabilization in 65 nm CMOS. *Solid-State Circuits, IEEE Journal of*, 49(6):1419–1436, June 2014.

[41] Guoliang Li, Xuezhe Zheng, Jin Yao, Hiren Thacker, Ivan Shubin, Ying Luo, Kannan Raj, John E. Cunningham, and Ashok V. Krishnamoorthy. 25Gb/s 1V-driving CMOS ring modulator with integrated thermal tuning. *Optics Express*, 19(21):20435–20443, October 2011.

[42] Hao Li, Zhe Xuan, Alex Titriku, Cheng Li, Kunzhi Yu, Binhao Wang, Ayman Shafik, Nan Qi, Yang Liu, Ran Ding, et al. A 25gb/s 4.4 v-swing ac-coupled si-photonic microring transmitter with 2-tap asymmetric ffe and dynamic thermal tuning in 65nm cmos. In *Solid-State Circuits Conference-(ISSCC), 2015 IEEE International*, pages 1–3. IEEE, 2015.

[43] Jun Li, Guoliang Li, Xuezhe Zheng, K. Raj, AV. Krishnamoorthy, and J.F. Buckwalter. A 25-gb/s monolithic optical transmitter with micro-ring modulator in 130-nm SoI CMOS. *Photonics Technology Letters, IEEE*, 25(19):1901–1903, Oct 2013.

[44] Di Liang and John E Bowers. Recent progress in lasers on silicon. *Nature Photonics*, 4(8):511–517, 2010.

[45] Jifeng Liu et al. Ge-on-si laser operating at room temperature. *Opt. Lett.*, 35(5):679–681, March 2010.

[46] Yangyang Liu and Milos A Popovic. Air-suspended high-q ring microcavities with scatterer-avoiding Şwiggler Ť supermode fields. In *Optical Fiber Communications Conference and Exhibition (OFC), 2014*, pages 1–3. IEEE, 2014.

[47] Yangyang Liu, Jeffrey M Shainline, Xiaoge Zeng, and Miloš A Popović. Ultra-low-loss cmos-compatible waveguide crossing arrays based on multimode bloch waves and imaginary coupling. *Optics letters*, 39(2):335–338, 2014.

[48] John D McCalpin. Stream: Sustainable memory bandwidth in high performance computers, 1995.

[49] R. Meade, J. Orcutt, K. Mehta, O. Tehar-Zahav, D. Miller, M. Georgas, B. Moss, C. Sun, Y.-H. Chen, J. Shainline, M. Wade, R. Bafrali, Z. Sternberg, G. Machavariani, G. Sandhu, M. Popovic, R. Ram, and V. Stojanovic. Integration of silicon photonics in bulk CMOS. In *VLSI Technology, 2014. Digest of Technical Papers. 2014 Symposium on*, pages 228–229, June 2014.

[50] R. Meade, O. Tehar-Zahav, Z. Sternberg, E. Megged, G. Sandhu, J.S. Orcutt, R. Ram, V. Stojanovic, M.R. Watts, E. Timurdogan, J. Shainline, and M. Popovic. Integration of silicon photonics in a bulk CMOS memory flow. In *Optical Interconnects Conference, 2013 IEEE*, pages 114–115, May 2013.

[51] Karan K. Mehta, Jason S. Orcutt, Jeffrey M. Shainline, Ofer Tehar-Zahav, Zvi Sternberg, Roy Meade, Miloš A. Popović, and Rajeev J. Ram. Polycrystalline silicon ring resonator photodiodes in a bulk complementary metal-oxide-semiconductor process. *Opt. Lett.*, 39(4):1061–1064, Feb 2014.

[52] Karan K Mehta, Jason S Orcutt, Ofer Tehar-Zahav, Zvi Sternberg, Reha Bafrali, Roy Meade, and Rajeev J Ram. High-q cmos-integrated photonic crystal microcavity devices. *Scientific reports*, 4, 2014.

[53] Milan M Milosevic, Frederic Y Gardes, David J Thomson, and Goran Z Mashanovich. Temperature insensitive racetrack resonators for near infrared applications. In *Optical Fiber Communication Conference*, pages OTh4D–5, 2012.

[54] 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, Feb 2013.

[55] J Munguía, G Bremond, JM Bluet, JM Hartmann, and M Mermoux. Strain dependence of indirect band gap for strained silicon on insulator wafers. *Applied Physics Letters*, 93(10):102101, 2008.

[56] A. Narasimha, B. Analui, Yi Liang, T.J. Sleboda, and C. Gunn. A fully integrated 4x10gb/s dwdm optoelectronic transceiver in a standard 0.13 µmCMOS SOI. In *Solid-State Circuits Conference, 2007. ISSCC 2007. Digest of Technical Papers. IEEE International*, pages 42 –586, feb. 2007.

[57] S Narasimha, K Onishi, HM Nayfeh, A Waite, M Weybright, J Johnson, C Fonseca, D Corliss, C Robinson, M Crouse, et al. High performance 45-nm soi technology with enhanced strain, porous low-k beol, and immersion lithography. In *Electron Devices Meeting, 2006. IEDM'06. International*, pages 1–4. IEEE, 2006.

[58] Magdalena S. Nawrocka, Tao Liu, Xuan Wang, and Roberto R. Panepucci. Tunable silicon microring resonator with wide free spectral range. *Applied Physics Letters*, 89(7):071110, 2006.

[59] 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, Feb 2013.

[60] Nvidia Tesla K40 and K80 GPU accelerators for servers. Website, `http://www.nvidia.com/object/tesla-servers.html`.

[61] Jason S. Orcutt, Anatol Khilo, Charles W. Holzwarth, Milos A. Popović, Hanqing Li, Jie Sun, Thomas Bonifield, Randy Hollingsworth, Franz X. Kärtner, Henry I. Smith, Vladimir Stojanović, and Rajeev J. Ram. Nanophotonic integration in state-of-the-art CMOS foundries. *Opt. Express*, 19(3):2335–2346, Jan 2011.

[62] Jason S. Orcutt, Benjamin Moss, Chen Sun, Jonathan Leu, Michael Georgas, Jeffrey Shainline, Eugen Zgraggen, Hanqing Li, Jie Sun, Matthew Weaver, Stevan Urošević, Miloš Popović, Rajeev J. Ram, and Vladimir Stojanović. Open foundry platform for high-performance electronic-photonic integration. *Opt. Express*, 20(11):12222–12232, May 2012.

[63] Kishore Padmaraju, Johnnie Chan, Long Chen, Michal Lipson, and Keren Bergman. Thermal stabilization of a microring modulator using feedback control. *Opt. Express*, 20(27):27999–28008, Dec 2012.

[64] Kishore Padmaraju, Dylan F Logan, Takashi Shiraishi, Jason J Ackert, Andrew P Knights, and Keren Bergman. Wavelength locking and thermally stabilizing microring resonators using dithering signals. *Lightwave Technology, Journal of*, 32(3):505–512, 2014.

[65] Gino Priem, Pieter Dumon, Walter Bogaerts, Dries Van Thourhout, Geert Morthier, and Roel Baets. Optical bistability and pulsating behaviour in silicon-on-insulator ring resonator structures. *Optics express*, 13(23):9623–9628, 2005.

[66] QD Laser. Website, `http://www.qdlaser.com`.

[67] Vivek Raghunathan, Juejun Hu, Winnie N. Ye, Jurgen Michel, and Lionel C. Kimerling. Athermal silicon ring resonators. In *Integrated Photonics Research, Silicon and Nanophotonics*, page IMC5. Optical Society of America, 2010.

[68] A Roger. Breaking a new sound barrier; it's a mic-on-a-chip. *Electronic Design*, 54, 2006.

[69] P. K. Runge. Undersea lightwave systems. *Optics and Photonics News*, 1(11):9–12, 1990.

[70] J. R. Sauer et al. Photonic interconnects for gigabit multicomputer communications. *LTS, IEEE*, 3(3):12–19, August 1992.

[71] Shankar Kumar Selvaraja, Wim Bogaerts, Pieter Dumon, Dries Van Thourhout, and Roel Baets. Subnanometer linewidth uniformity in silicon nanophotonic waveguide devices using cmos fabrication technology. *Selected Topics in Quantum Electronics, IEEE Journal of*, 16(1):316–324, 2010.

[72] Jeffrey M Shainline, Jason Orcutt, Mark T Wade, Roy Meade, Ofer Tehar-Zehav, Zvi Sternberg, Vladimir Stojanovic, and Milos Popovic. Multi-modal

optical microcavities for loss avoidance. In *CLEO: Science and Innovations*, pages CM1F–2. Optical Society of America, 2013.

[73] Jeffrey M Shainline, Jason S Orcutt, Mark T Wade, Kareem Nammari, Benjamin Moss, Michael Georgas, Chen Sun, Rajeev J Ram, Vladimir Stojanović, and Miloš A Popović. Depletion-mode carrier-plasma optical modulator in zero-change advanced cmos. *Optics letters*, 38(15):2657–2659, 2013.

[74] R. Soref and B. Bennett. Electrooptical effects in silicon. *IEEE Journal of Quantum Electronics*, 23(1), January 1987.

[75] V. Stojanovic, A. Ho, B.W. Garlepp, F. Chen, J. Wei, G. Tsang, E. Alon, R.T. Kollipara, C.W. Werner, J.L. Zerbe, and M.A. Horowitz. Autonomous dual-mode (pam2/4) serial link transceiver with adaptive equalization and data recovery. *Solid-State Circuits, IEEE Journal of*, 40(4):1012–1026, April 2005.

[76] C. Sun, M. Georgas, J. Orcutt, B. Moss, Y.-H. Chen, J. Shainline, M. Wade, K. Mehta, K. Nammari, E. Timurdogan, D. Miller, O. Tehar-Zahav, Z. Sternberg, J. Leu, J. Chong, R. Bafrali, G. Sandhu, M. Watts, R. Meade, M. Popovic, R. Ram, and V. Stojanovic. A monolithically-integrated chip-to-chip optical link in bulk CMOS. *Solid-State Circuits, IEEE Journal of*, 50(4):828–844, April 2015.

[77] C. Sun, M. Georgas, J.S. Orcutt, B.R. Moss, Y.-H. Chen, J. Shainline, M. Wade, K. Mehta, K. Nammari, E. Timurdogan, D. Miller, O. Tehar-Zahav, Z. Sternberg, J.C. Leu, J. Chong, R. Bafrali, G. Sandhu, M. Watts, R. Meade, M.A Popovic, R.J. Ram, and V. Stojanovic. A monolithically-integrated chip-to-chip optical link in bulk CMOS. In *VLSI Circuits Digest of Technical Papers, 2014 Symposium on*, pages 56–57, June 2014.

[78] C. Sun, M. Wade, M. Georgas, S. Lin, L. Alloatti, B. Moss, R. Kumar, A. Atabaki, F. Pavanello, R. Ram, M. Popović, and V. Stojanović. A 45nm soi monolithic photonics chip-to-chip link with bit-statistics-based resonant microring thermal tuning. In *VLSI Circuits Digest of Technical Papers, 2015 Symposium on*, June 2015.

[79] Chen Sun, C-HO Chen, George Kurian, Lan Wei, Jason Miller, Anant Agarwal, Li-Shiuan Peh, and Vladimir Stojanovic. Dsent-a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling. In *Networks on Chip (NoCS), 2012 Sixth IEEE/ACM International Symposium on*, pages 201–210. IEEE, 2012.

[80] Chen Sun, Erman Timurdogan, Michael R. Watts, and Vladimir Stojanovic. Integrated microring tuning in deep-trench bulk CMOS. In *Optical Interconnects Conference*, May 2013.

[81] Dirk Taillaert, Peter Bienstman, and Roel Baets. Compact efficient broadband grating coupler for silicon-on-insulator waveguides. *Optics Letters*, 29(23):2749–2751, December 2004.

[82] Erman Timurdogan, Aleksandr Biberman, Douglas C. Trotter, Chen Sun, Michele Moresco, Vladimir Stojanovic, and Michael R. Watts. Automated wavelength recovery for microring resonators. In *Conference on Lasers and Electro-Optics 2012*, page CM2M.1. Optical Society of America, 2012.

[83] Erman Timurdogan, Cheryl M Sorace-Agaskar, Jie Sun, Ehsan Shah Hosseini, Aleksandr Biberman, and Michael R Watts. An ultralow power athermal silicon modulator. *Nature communications*, 5, 2014.

[84] Erman Timurdogan, Zhan Su, Krishna Settaluri, Sen Lin, Sajjad Moazeni, Chen Sun, Gerry Leake, Douglas D Coolbaugh, Ben Moss, Michele Moresco, et al. An ultra low power 3d integrated intra-chip silicon electronic-photonic link. In *Optical Fiber Communication Conference*, pages Th5B–8. Optical Society of America, 2015.

[85] 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 the 35th Annual International Symposium on Computer Architecture*, ISCA '08, pages 153–164. IEEE Computer Society, 2008.

[86] M. Wade, F. Pavanello, R. Kumar, C. Gentry, A. Atabaki, R. Ram, V. Stojanović, and M. Popović. 75% efficient wide bandwidth grating couplers in a 45 nm microelectronics CMOS process. In *Optical Interconnects Conference*, April 2015.

[87] Mark T Wade, Rajesh Kumar, Kareem Nammari, Cale M Gentry, Jeffrey Shainline, Jason S Orcutt, Ananth Tamma, Rajeev Ram, Vladimir Stojanovic, and Miloš A Popović. Unidirectional chip-to-fiber grating couplers in unmodified 45nm cmos technology. In *CLEO: Science and Innovations*, pages STh3M–5. Optical Society of America, 2014.

[88] Mark T. Wade, Jeffrey M. Shainline, Jason S. Orcutt, Chen Sun, Rajesh Kumar, Ben Moss, Michael Georgas, Rajeev J. Ram, Vladimir Stojanovic, and Milos Popovic. Energy-efficient active photonics in a zero-change, state-of-the-art CMOS process. In *Optical Fiber Communication Conference*, page Tu2E.7. Optical Society of America, 2014.

[89] Joachim Wagner and Jesus A del Alamo. Band-gap narrowing in heavily doped silicon: A comparison of optical and electrical data. *Journal of applied physics*, 63(2):425–429, 1988.

[90] Andrew Waterman, Yunsup Lee, David A Patterson, and Krste Asanovi. The risc-v instruction set manual. volume 1: User-level isa, version 2.0. Technical report, University of California, Berkeley, Department of Electrical Engineering and Computer Sciences, 2014.

[91] Michael R. Watts. Adiabatic microring resonators. *Opt. Lett.*, 35(19):3231–3233, October 2010.

[92] Gregory L Wojcik, Dongliang Yin, Alexey R Kovsh, Alexey E Gubenko, Igor L Krestnikov, Sergey S Mikhrin, Daniil A Livshits, David A Fattal, Marco Fiorentino, and Raymond G Beausoleil. A single comb laser source for short reach wdm interconnects. In *SPIE OPTO: Integrated Optoelectronic Devices*, pages 72300M–72300M, 2009.

[93] M. Wordeman, J. Silberman, G. Maier, and M. Scheuermann. A 3D system prototype of an eDRAM cache stacked over processor-like logic using through-silicon vias. In *Solid-State Circuits Conference Digest of Technical Papers (ISSCC), 2012 IEEE International*, pages 186–187, Feb 2012.

[94] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson. 12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators. *Optical Society of America*, 15(2), January 2007.

[95] Qianfan Xu and Michal Lipson. Carrier-induced optical bistability in silicon ring resonators. *Optics letters*, 31(3):341–343, 2006.

[96] Kunzhi Yu, Hao Li, Cheng Li, Alex Titriku, Ayman Shafik, Binhao Wang, Zhongkai Wang, Rui Bai, Chin-Hui Chen, Marco Fiorentino, et al. A 24gb/s 0.71 pj/b si-photonic source-synchronous receiver with adaptive equalization and microring wavelength stabilization. In *Solid-State Circuits Conference-(ISSCC), 2015 IEEE International*, pages 1–3. IEEE, 2015.

[97] Xuezhe Zheng, Eric Chang, Philip Amberg, Ivan Shubin, Jon Lexau, Frankie Liu, Hiren Thacker, Stevan S. Djordjevic, Shiyun Lin, Ying Luo, Jin Yao, Jin-Hyoung Lee, Kannan Raj, Ron Ho, John E. Cunningham, and Ashok V. Krishnamoorthy. A high-speed, tunable silicon photonic ring modulator integrated with ultra-efficient active wavelength control. *Opt. Express*, 22(10):12628–12633, May 2014.

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

[99] Xuezhe Zheng, Ying Luo, Guoliang Li, Ivan Shubin, Hiren Thacker, Jin Yao, Kannan Raj, John E Cunningham, and Ashok V Krishnamoorthy. Enhanced optical bistability from self-heating due to free carrier absorption in substrate removed silicon ring modulators. *Optics express*, 20(10):11478–11486, 2012.

[100] William A Zortman, Anthony L Lentine, Douglas C Trotter, and Michael R Watts. Bit-error-rate monitoring for active wavelength control of resonant modulators. *IEEE Micro*, (1):42–52, 2013.