

A Monolithically-Integrated Chip-to-Chip Optical Link in Bulk CMOS

Chen Sun, *Student Member, IEEE*, Michael Georgas, Jason Orcutt, Benjamin Moss, *Member, IEEE*, Yu-Hsin Chen, *Student Member, IEEE*, Jeffrey Shainline, Mark Wade, Karan Mehta, Kareem Nammari, Erman Timurdogan, *Student Member, IEEE*, Daniel Miller, Ofer Tehar-Zahav, Zvi Sternberg, Jonathan Leu, Johanna Chong, Reha Bafrafi, Gurtej Sandhu, *Fellow, IEEE*, Michael Watts, *Member, IEEE*, Roy Meade, *Senior Member, IEEE*, Miloš Popović, *Member, IEEE*, Rajeev Ram, *Senior Member, IEEE*, and Vladimir Stojanović, *Member, IEEE*

Abstract—Silicon-photonics is an emerging technology that can overcome the tradeoffs faced by traditional electrical I/O. Due to ballooning development costs for advanced CMOS nodes, however, widespread adoption necessitates seamless photonics integration into mainstream processes, with as few process changes as possible. In this work, we demonstrate a silicon-photonics link with optical devices and electronics integrated on the same chip in a 0.18 μm bulk CMOS memory periphery process. To enable waveguides and optics in process-native polysilicon, we introduce deep-trench isolation, placed underneath to prevent optical mode leakage into the bulk silicon substrate, and implant-amorphization to reduce polysilicon loss. A resonant defect-trap photodetector using polysilicon eliminates need for germanium integration and completes the fully polysilicon-based photonics platform. Transceiver circuits take advantage of photonic device integration, achieving 350 fJ/b transmit and 71 μA_{pp} BER = 10^{-12} receiver sensitivity at 5 Gb/s. We show high fabrication uniformity and high-Q resonators, enabling dense wavelength-division multiplexing with 9-wavelength 45 Gb/s transmit/receive data-rates per waveguide/fiber. To combat perturbations to variation- and thermally-sensitive resonant devices, we demonstrate an on-chip thermal tuning feedback loop that locks the resonance to the laser wavelength. A 5 m optical chip-to-chip link achieves 5 Gb/s while consuming 3 pJ/b and 12 pJ/bit of circuit and optical energy, respectively.

Index Terms—Memory, optical interconnects, optoelectronics, process integration, silicon-photonics, wireline transceivers.

Manuscript received August 31, 2014; revised November 04, 2014; accepted November 18, 2014. Date of publication January 19, 2015; date of current version March 24, 2015. This paper was approved by Guest Editor Jeffrey Gealow. This work was supported by DARPA POEM award HR0011-11-C-0100 and contract HR0011-11-9-0009, led by Dr. Jagdeep Shah. Support is also acknowledged from NSF, FCRP IFC, Trusted Foundry, Intel, MIT CICS, Berkeley Wireless Research Center, Santec, and NSERC. The views expressed are those of the authors and do not reflect the official policy or position of the DoD or the U.S. Government.

C. Sun, M. Georgas, J. Orcutt, B. Moss, Y.-H. Chen, K. Mehta, E. Timurdogan, J. Leu, J. Chong, M. Watts, and R. Ram are with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA.

C. Sun and V. Stojanović are with the Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, CA 94709 USA.

J. Shainline, M. Wade, K. Nammari, and M. Popović are with the Department of Electrical, Computer, and Energy Engineering, University of Colorado, Boulder, CO 80309 USA.

D. Miller, O. Tehar-Zahav, and Z. Sternberg are with Micron Semiconductor Israel, Kiryat Gat 82109, Israel.

R. Bafrafi, G. Sandhu, and R. Meade are with Micron Technology, Inc., Boise, ID 83707 USA.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSSC.2014.2382101

I. INTRODUCTION

WHILE electrical links remain mainstream for short-reach chip-to-chip applications, high channel losses, pin-count constraints, and crosstalk limit the energy-efficiency and bandwidth. Optical links demonstrate low loss at high symbol rates, have high distance insensitivity, and are immune to electromagnetic interference, making them a promising alternative for bandwidth-bound systems [1] such as interconnection networks or memory systems [2], [3]. Integrated silicon-photonics, in particular, holds promise of mass-produced photonics in low-cost silicon ICs and enables dense wavelength division multiplexing (DWDM), where multiple data channels share a single waveguide or fiber to greatly extend bandwidth density through an optical I/O port.

Despite the numerous efforts underway [4]–[7], the integration of optical devices with advanced electronics in a VLSI system remains a key challenge. Current integration strategies are either monolithic, with optical devices and transistors both on the same chip [4]–[6], or heterogeneous, with separate photonic and electronic dies bonded in a multi-chip solution [8]–[10]. While heterogeneous integration decouples process optimization for photonics and electronics, the necessary multi-chip TSV/microbump packaging adds significant parasitic capacitance [11], [12] and limits density, degrading performance and energy efficiency. Monolithic integration, on the other hand, simplifies packaging and enables tighter device-to-circuit proximity to lower parasitics, but faces integration challenges with electronics. To date, the vast majority of both approaches use custom silicon-on-insulator (SOI) processes with thick buried-oxide (BOX), utilizing the crystalline silicon layer on top of the BOX to form low-loss waveguides. The thick BOX provides optical mode confinement [4], [5], [8], [13], preventing light in the waveguide from coupling into the substrate. Monolithic thin-BOX photonics has also been demonstrated in a commercial CMOS SOI process [6], [14]–[16]. Here, the BOX is not sufficiently thick to completely confine the optical mode, but serves as a stopper for post-process substrate removal to use air (or an attached low-K material) as isolation. While SOI is attractive for photonics, it is a niche process compared to bulk; the high per-unit cost of SOI creates a hurdle for mainstream foundry customers. Heterogeneous approaches adds additional 3D integration

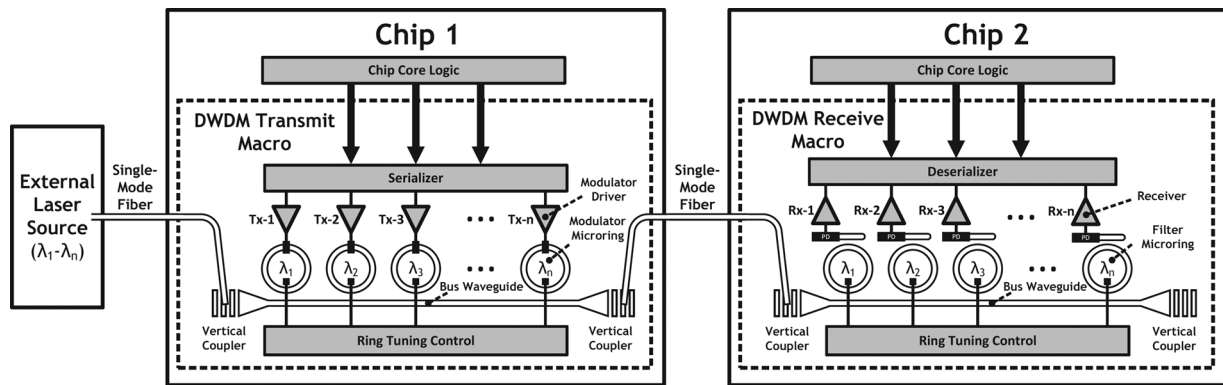


Fig. 1. A chip-to-chip DWDM optical link using silicon-photonics.

complexity and packaging yield concerns, ultimately raising even more barriers to widespread photonics adoption.

To gain traction among high-volume 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 [17], [18], which can be much more lossy due to crystal grain imperfections. Alternatively, crystalline silicon may be epitaxially grown as the waveguide material [7], 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 [17] requires post-processing and constrains circuit-photonics placement. The extension of shallow-trench isolation to make deep oxide-filled trenches underneath optical waveguides [7], [18] is an alternative that enables tighter integration without post-processing, but requires additional process integration. Because of these challenges, electro-optic transceivers and links in monolithic bulk processes have yet to be demonstrated.

This paper 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.

The rest of the paper is as follows. Section II provides an overview of optical microring resonator devices—a fundamental building block for DWDM—to motivate the design choices of the platform. In Section III, we introduce our bulk silicon-photonics platform. Section IV and Section V describe the transmitter, receiver, and DWDM transceiver macros. In Section VI, 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 VII.

II. THE OPTICAL RING RESONATOR

The architecture of an n - λ DWDM chip-to-chip silicon-photonics link is shown in Fig. 1. An off-chip laser source—either

a comb laser or a bank of lasers shared across all links in the system—produces continuous-wave light of n wavelengths (λ s) which couple 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 tuned to each λ , modulates each λ with an independent bit-stream. The modulated λ s exit the transmit chip through a second VGC into a single-mode fiber bound for the receive chip, where they couple into an n - λ DWDM receive macro. Here, filter microrings tuned to each λ drop the light onto photodetectors to produce photocurrent, which the receivers resolve into data. To simplify clock recovery, the clock can also be source-forwarded on one of the wavelengths [19], [20].

The key building block of a DWDM link is the optical microring resonator (Fig. 2). When coupled to a waveguide, the ring captures only light at its resonant λ s, forming a notch filter. A λ is resonant when the ring circumference is an integer multiple of λ . Resonances are periodic with a spacing defined as the free spectral range (FSR), which grows with a smaller circumference. Due to this periodicity, all λ s used in a DWDM link must fit within one FSR, where each ring has single- λ selectivity. The quality factor of the microring resonator is dictated by the full-width half-maximum (FWHM) bandwidth of the notch and the intrinsic extinction ratio (ER_i) is defined as the depth of the notch. Doping and contacting the ring in active structures (modulators or photodetectors) lowers the quality factor by introducing loss from free carrier absorption. ER_i is primarily a factor of the coupling between the ring and the bus waveguide, tuned by the gap between them. A ring that is perfectly critically coupled (defined as when the power coupled from the bus waveguide matches the energy decay rate due to losses in the microring) provides infinite ER_i . 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.

An electro-optic microring modulator can be realized by modulating the resonance (λ_0) to perform on-off-keying (OOK) of input light aligned close to λ_0 (Fig. 3, left). Changes in free carrier concentration are used to shift λ_0 by changing the material's index of refraction [21]. Carrier-injection modulators are p-i-n junctions that inject carriers into the intrinsic region during forward-bias, blue-shifting λ_0 . Carrier-depletion

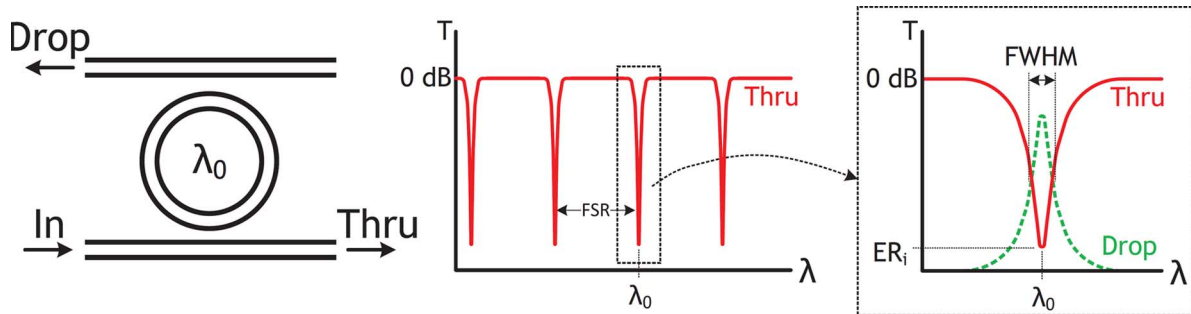


Fig. 2. The optical transfer characteristics of a microring resonator (a waveguide looped around in a ring to form a resonating cavity) that is resonant at λ_0 .

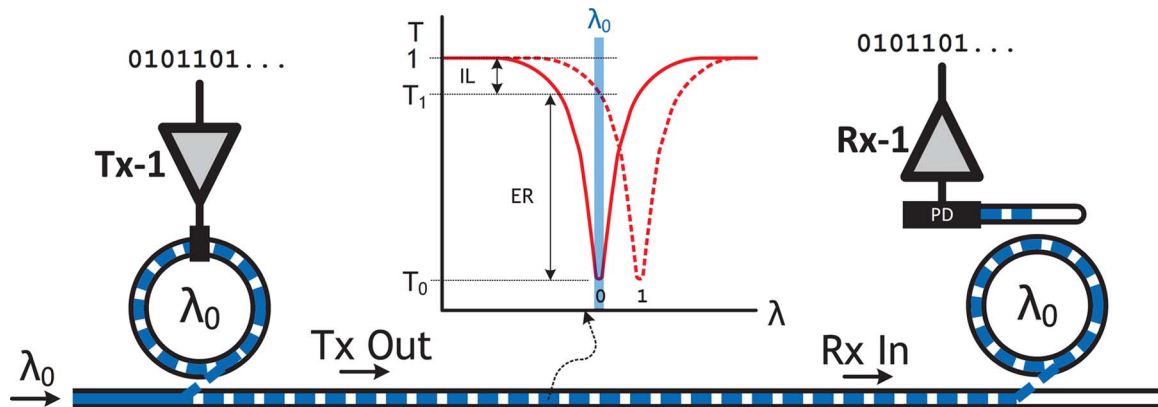


Fig. 3. Ring resonators used in a modulator and a receiver. The input light is at the same wavelength as λ_0 . The modulation insertion loss (IL) and the extinction ratio (ER) are defined as $10 \cdot \log_{10} T_1$ and $10 \cdot \log_{10} T_1/T_0$, respectively. T_1 and T_0 represent modulated output powers for optical ones and zeros, normalized to the input power.

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 [9], [15], [22]. Forward-biased operation of the junction also results in static power dissipation and poor energy-efficiency. Carrier-depletion designs avoid these issues, but require better doping control to balance λ_0 shift with Q-factor degradation from free carrier absorption. Rings also perform receiver channel wavelength selection, dropping only light at λ_0 onto a photodetector (PD), which is broadband (Fig. 3, right). Embedding the PD in the ring itself [23] can enhance optical absorption for a given PD size. However, integration of the PD material (typically Germanium) into the ring can be challenging.

An increase or decrease in temperature causes a resonance red-shift or blue-shift (through a change in index of refraction), respectively, of roughly 10 GHz/K (55 pm/K at 1280 nm). The strong thermal dependence enables integrated microring heaters to combat resonance variations from limited process tolerances [19], [24], but also necessitates active resonance control when microrings are integrated in an electrical system, where temperatures fluctuate. Because thermal variations are slow, however, wavelength-locking can be achieved with low overheads [9], [25]–[28].

III. MONOLITHIC PHOTONICS PLATFORM IN BULK

The monolithic bulk platform is demonstrated in a $0.18 \mu\text{m}$ 3-metal-layer bulk CMOS process. We construct all optical de-

vices 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.

To enable photonics, we make three key modifications to the original CMOS process [29], shown in Fig. 4. 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 [18]. Further nitride spacer optimizations bring the loss down to 10.5 dB/cm at process end-of-line [29]. The losses compare favorably to the 2 dB/mm (20 dB/cm) losses reported in [7] for waveguides built using epitaxially grown crystalline silicon in an electronic-photonic process flow. The second is the inclusion of a deep-trench isolation step (DTI) to avoid post-process undercut. DTI creates a $1.2 \mu\text{m}$ oxide-filled trench underneath polysilicon waveguides to provide optical waveguide mode confinement and isolation from the bulk silicon substrate. Optical waveguides and devices can be placed as close as $2 \mu\text{m}$ 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 [30] in the trenches. After these process changes, transistor

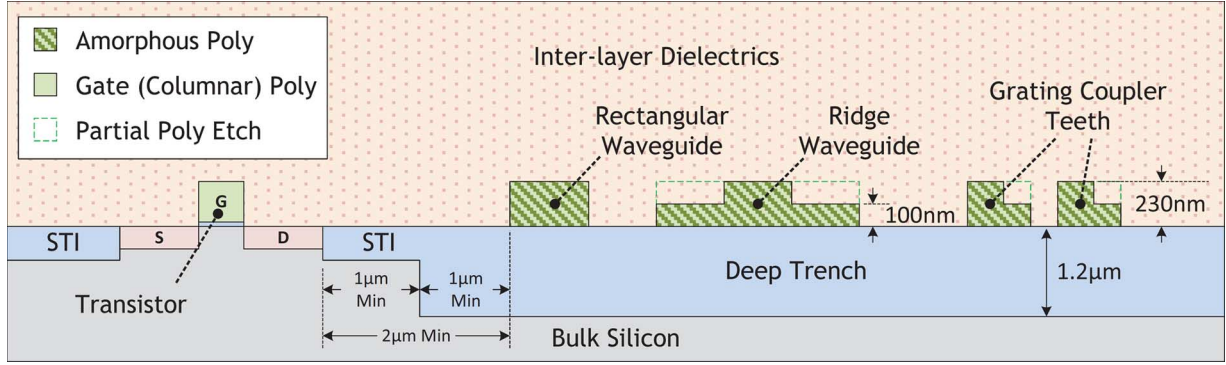


Fig. 4. Photonics platform cross-section. The step-like shape of the polysilicon ridge waveguide and vertical coupler grating teeth are enabled via the partial polysilicon etch.

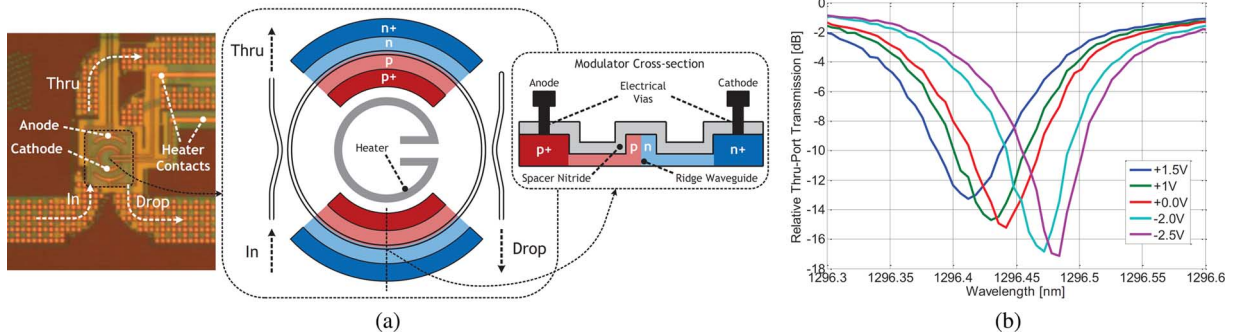


Fig. 5. Structure of the carrier-depletion ridge waveguide modulator used in the transmitter (a) and its measured thru-port transfer characteristics under different applied DC voltages (b). We estimate a modulator junction capacitance of 20 fF and a series resistance of 500 Ω , corresponding to a 15.9 GHz device RC bandwidth.

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 epitaxial crystallization of silicon or the introduction of germanium or silicon-germanium in our platform. Though these methods are adopted by prior art [4], [5], [7], they also introduce additional high-temperature processing and front-end 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.

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 p-n junction (Fig. 5). 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 [31]. Measured optical transfer characteristics of this device under different DC voltages (Fig. 5(b)) shows a resonance shift of 2.7 GHz/V (15 pm/V) 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.

To avoid the introduction of germanium—found in all custom photonics platforms to date [7], [8], [13], [32]—we employ a completely polysilicon-based photodetector utilizing absorption from defect states [33]. This photodetector is a ridge-waveguide microring (with matching radius and FSR as the modulator device) doped with a p-i-n junction (Fig. 6(a)). 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 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 connect to PD-microrings that are severely overcoupled (Fig. 6(b)). 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.

We build the optical link components as part of a 24 mm \times 24 mm technology development reticle for testing a

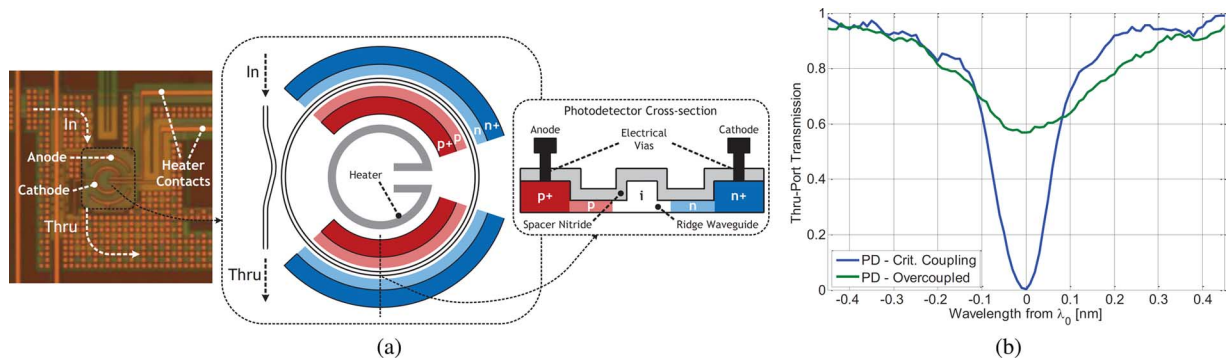


Fig. 6. Structure of the resonant polysilicon defect photodetector used in the receiver (a) and its measured thru-port transfer characteristics (b). We show the case for both a critically-coupled ring and the severely overcoupled ring (which is the version connected to the receiver circuits). Note the linear scale in the transfer characteristic. The photodetector capacitance is estimated to be 15 fF.

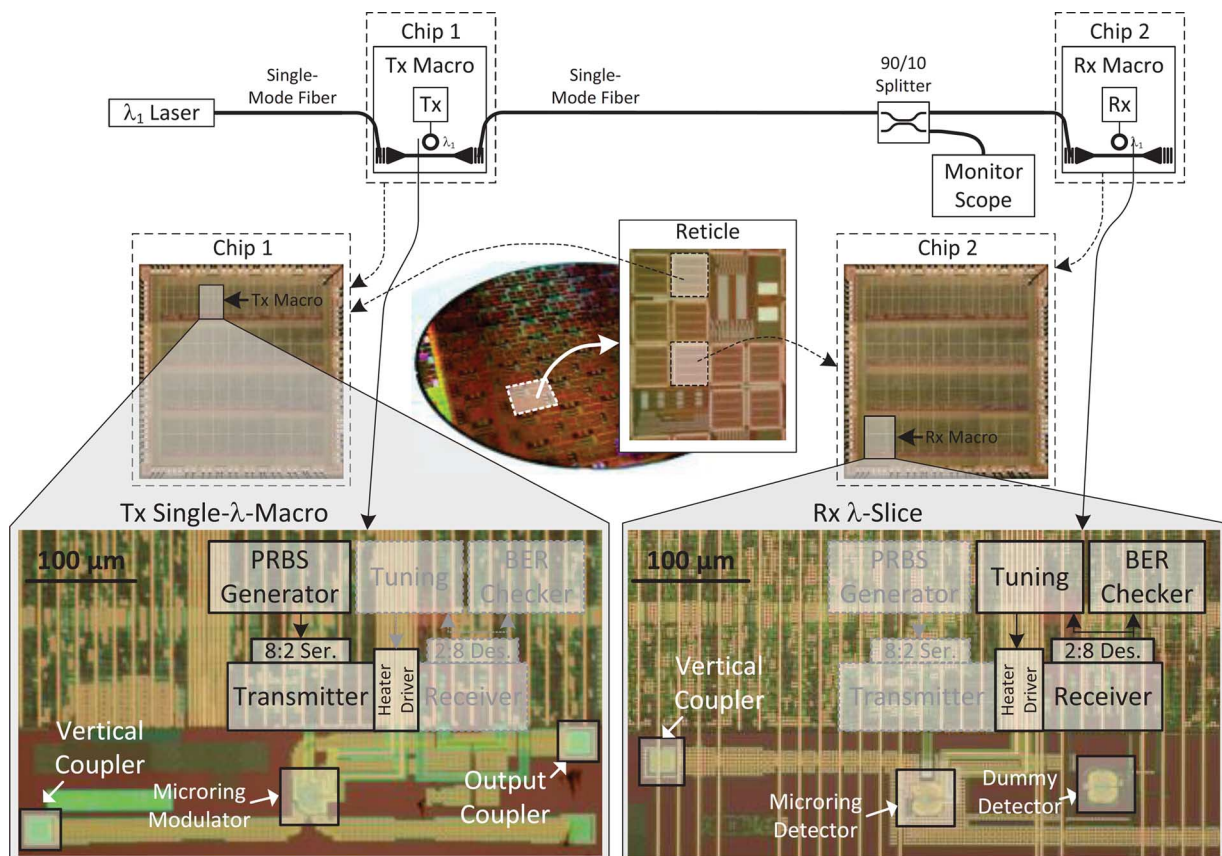


Fig. 7. Monolithically-integrated photonics platform in bulk. Shown here is an example of a single- λ chip-to-chip optical link using a transmit macro from chip 1 and a receive macro from chip 2.

variety of optical devices and circuits. The reticle is divided into standalone optical device regions and an array of ten $5\text{ mm} \times 5\text{ mm}$ transceiver chiplets, which are individually wirebonded and packaged electrically. Each chiplet hosts an array of 8 single- λ electro-optic transceiver macros (Fig. 7) and three 9- λ DWDM transceiver macros for a total of 5.5 million transistors and 100 optical devices per chiplet.

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 gener-

ators 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. 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 nine adjacent single- λ macros, but string together all nine modulator- or receive-microrings on a single bus waveguide to provide DWDM functionality.

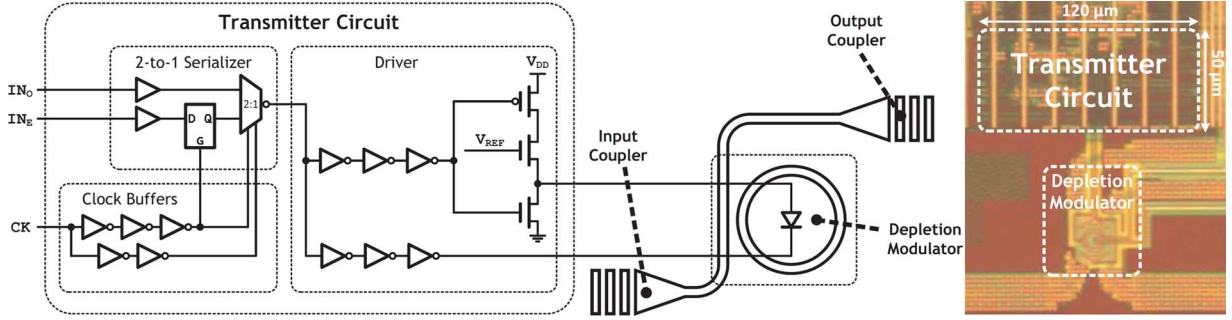


Fig. 8. Carrier-depletion-based transmitter. The output impedance of the driver is approximately 700Ω .

IV. DEPLETION-MODE MICRORING OPTICAL TRANSMITTER

The transmitter consists of a depletion-ridge modulator device driven by a transmitter frontend circuit, shown in Fig. 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 1 s, the transmitter circuit applies a voltage of $-V_{DD}$ to the modulator junction, depleting the junction of carriers and red-shifting the resonance. On logic 0 s, 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.

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:

$$\Delta T = |T_1 - T_0| = |10^{-IL/10} - 10^{-(IL+ER)/10}| \quad (1)$$

where IL is the modulator insertion loss and ER is the modulator extinction ratio, 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. 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 suppressed 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. 10). The maximum data-rate of the transmitter is limited by transistor performance; At 5 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 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

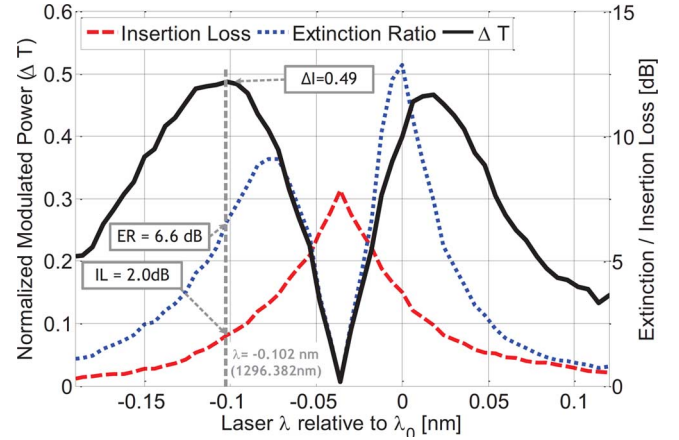


Fig. 9. Trade-off analysis between extinction ratio, insertion loss, and the eye-height metric ΔT versus the laser wavelength. We use a drive swing of -2.5 V to 1.5 V and the device in Fig. 5(b). Note that the polarity of the modulator flips when T_1 is smaller than T_0 to keep ΔT positive.

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 thermal 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. 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. For the macro floorplan to be compatible across all types of optical devices on the platform, we had to place the modulator device $60 \mu\text{m}$ 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 $2 \mu\text{m}$ spacing will reduce wiring capacitances

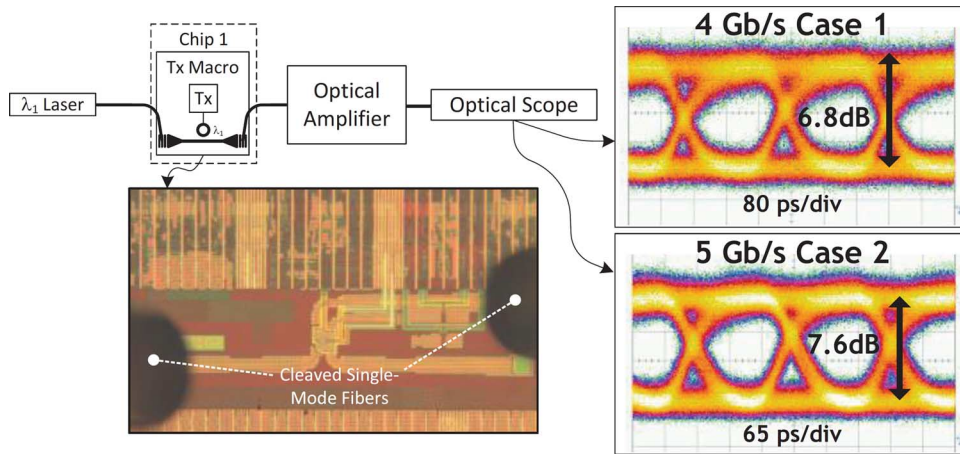


Fig. 10. Transmitter characterization setup. The eye-diagrams are shown for *case 1* at 4 Gb/s and *case 2* at 5 Gb/s.

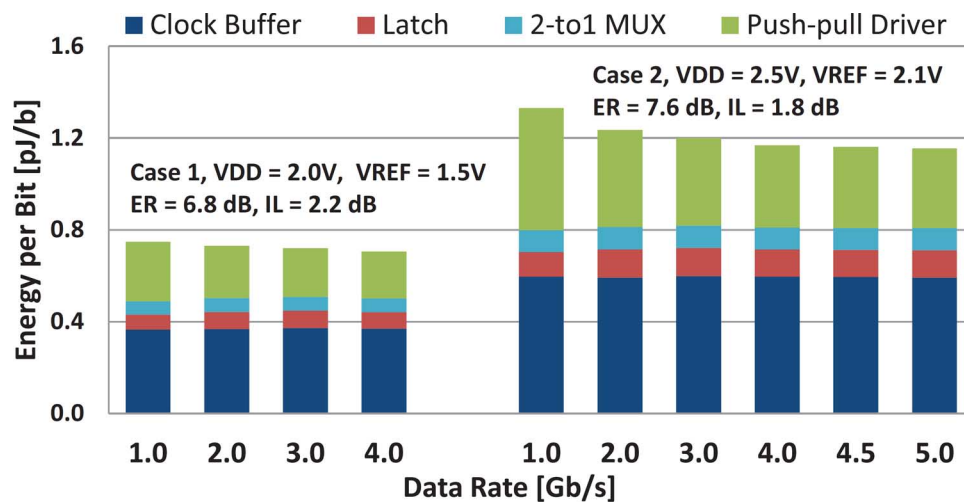


Fig. 11. Measured transmitter energy-per-bit. The component-level power breakdowns are deembedded through simulation using extracted netlists that include the wiring capacitances on the modulator anode and cathode (extracted to be 13 fF and 14 fF, respectively) and a 20 fF modulator junction capacitance.

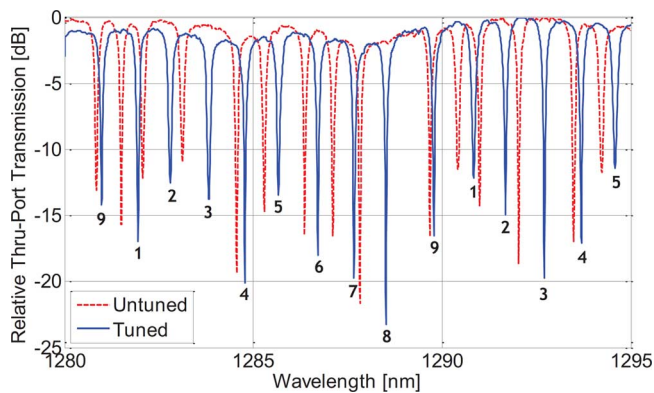


Fig. 12. Optical thru-port transfer characteristic of the 9- λ DWDM transmit bank. The tuned spectra are after coarse tuning to a 1 nm grid. Each resonance is numbered with its corresponding ring.

down to sub-2 fF, extending the achievable data-rate and lowering the energy cost further.

We build the transmitter circuit as part of a 9- λ DWDM transmit macro, shown in Fig. 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. 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. 13). The I/O density of the DWDM macro is approximately 110 Gb/s/mm², including all transmitter circuits (120 μ m \times 50 μ m per slice) and a conservative (3500 μ m \times 100 μ m) trench for the photonic devices. Note that the ring-to-ring placement pitch of 384 μ m and the overall size of DWDM macro (4000 μ m \times 900 μ m) is limited by the area of the digital testing backend; the pitch can be as tight as 40 μ m without violating design rules.

V. POLYSILICON-BASED RESONANT OPTICAL RECEIVER

The receiver block is composed of a ring resonant defect-detector connected to a receiver circuit, shown in Fig. 14. To

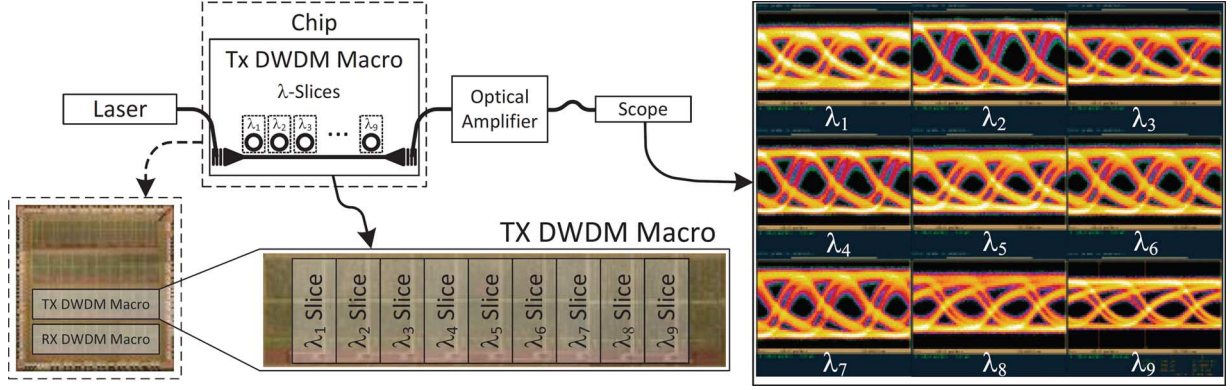


Fig. 13. Demonstration of the 9- λ DWDM transmit macro. The double-edge transition in the eyes is due to the previous edge not fully settling after one bit-time.

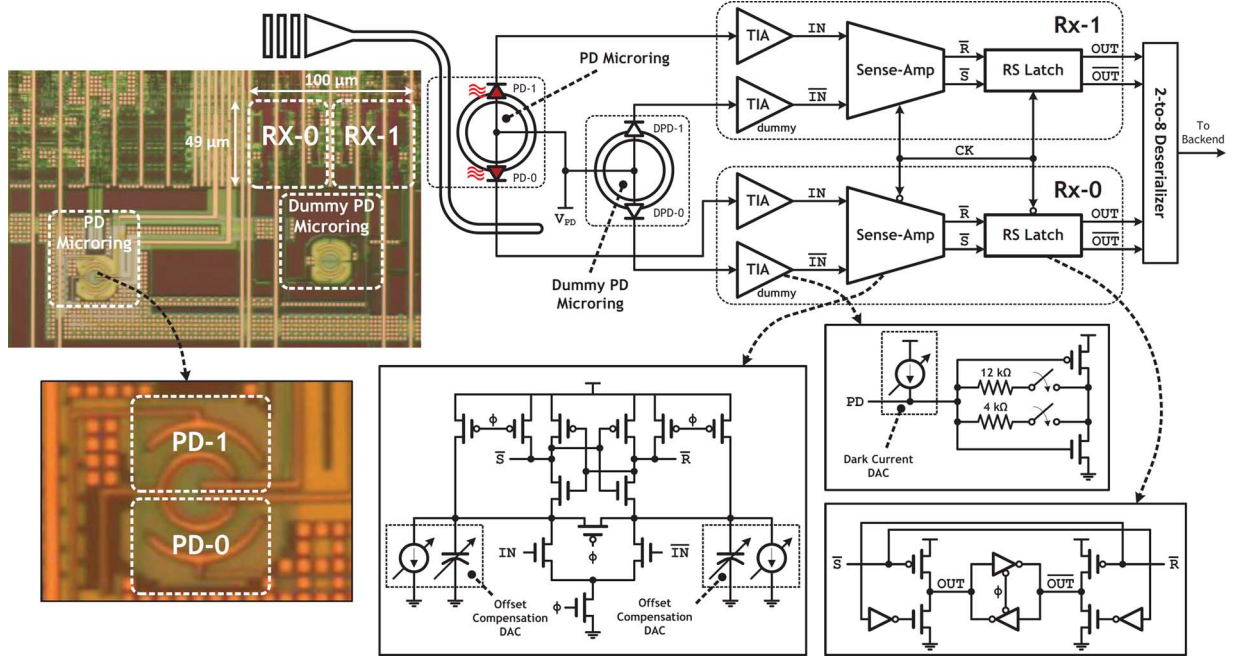


Fig. 14. Optical micrograph and schematic of the receiver architecture.

mitigate the slow speed of the process, we adopt a split-diode technique [23]; 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 k Ω , 4 k Ω , or 12 k Ω ||4 k Ω . 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 [33]. 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:

$$v_{out} = v_{in} \frac{g_m}{C_{int}} \cdot t_{int} \cdot e^{t_{regen}/\tau_{inv}} \quad (2)$$

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 automatically when t_{eval}

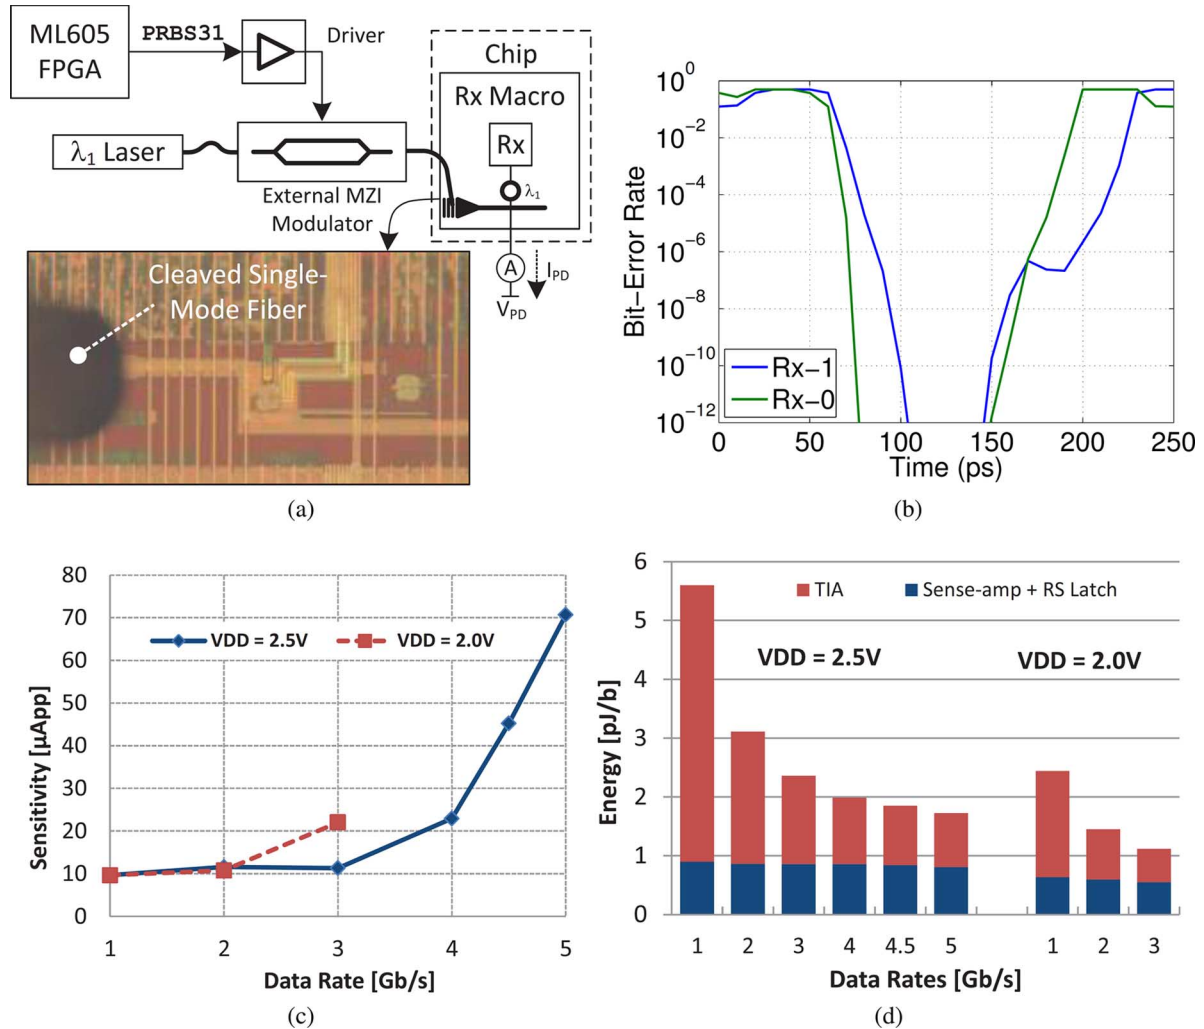


Fig. 15. Receiver measurement test setup (a), optical BER measurement at 5 Gb/s (b), measured optical sensitivity (c), and measured receiver energy-per-bit (d). Sensitivity is defined using the total peak-to-peak photocurrent (both halves) necessary for error-free operation after 2×10^{12} total bits. The 5 Gb/s receiver optical BER measurement is performed using a laser power that results in $I_{PD} = 71 \mu\text{A}_{pp}$. 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.

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. 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-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 $12 \mu\text{A}_{pp}$ (-15.2 dBm with a 0.2 A/W PD) for 1 Gb/s–3 Gb/s. At 5 Gb/s, sensitivity degrades to $126 \mu\text{A}_{pp}$ and $71 \mu\text{A}_{pp}$ (-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 $150 \mu\text{m}$ farther from the receiver than the dummy PD. From layout extraction, this error lowers the TIA bandwidth by approximately 1/3, causing the $12 \text{ k}\Omega$ feedback resistor to have insufficient bandwidth to support data-rates beyond 3 Gb/s. The reduction of the TIA feedback resistance to $4 \text{ k}\Omega$, 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

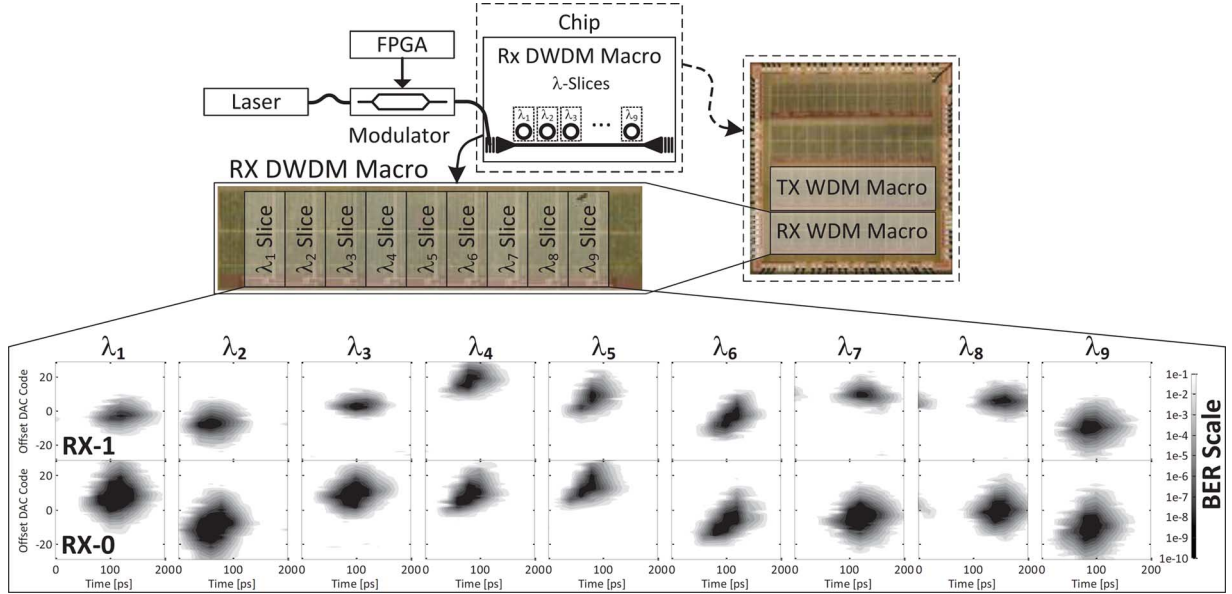


Fig. 16. Demonstration of the 9- λ DWDM receiver macro. The receive eye diagrams are generated by sweeping the receiver threshold using the offset compensation DACs.

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.

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. 16). We record error-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/mm²). Fig. 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.

VI. 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. 18) contains an optical power meter circuit, a configurable data path, a programmable lookup-table (LUT), and a $\Delta\Sigma$ -DAC circuit [31] that drives an integrated microring heater. As opposed to tracking an averaged photocurrent [9], [25], [26], [28], 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 [27], [34], and [35] to actively adjust receiver threshold levels given variances in signal power. Here,

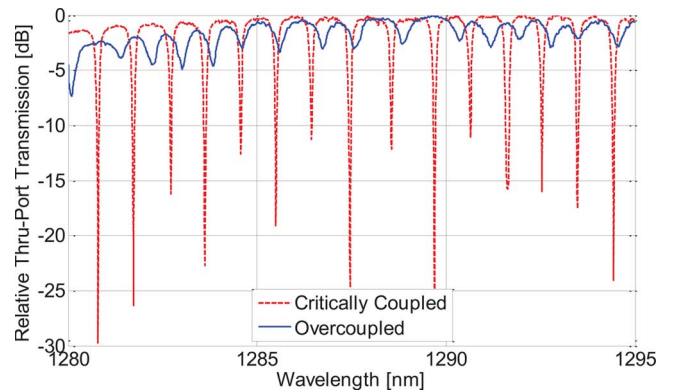


Fig. 17. Optical transfer characteristics of the measured DWDM receiver macro, which has overcoupled PD rings, and a DWDM receiver macro with critically-coupled PD rings. The rings are coarsely tuned to a 1 nm grid.

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 capacitive 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 heater strength. A $\Delta\Sigma$ -DAC (described in [31]) drives the integrated heater. All tuning backend circuits are clocked by the data-clock through a divide-by-64 divider.

The lock range of the tuner is 0.5 nm (90 GHz) [36], corresponding to a 9 K change in temperature. In the context of

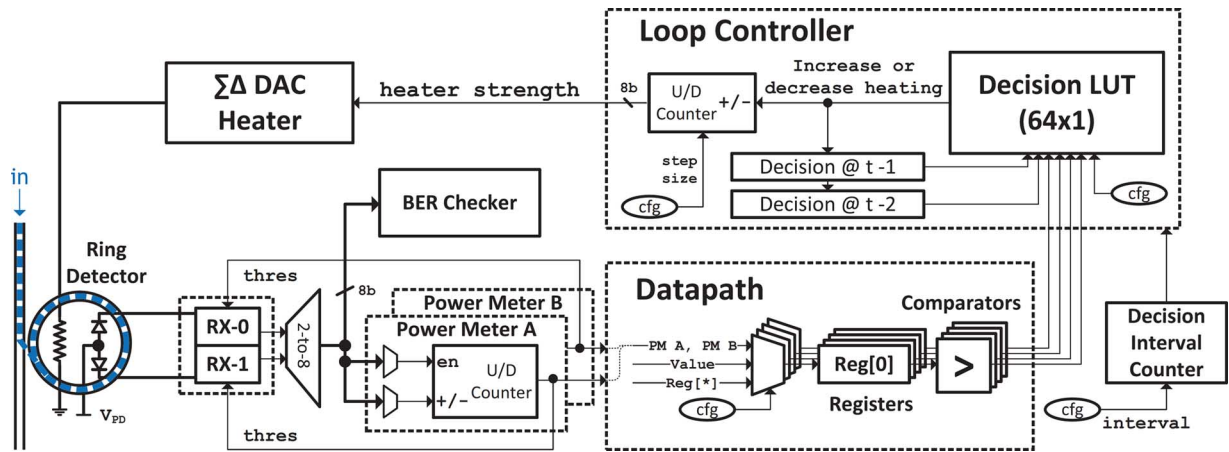


Fig. 18. Architecture of the programmable tuning sub-system.

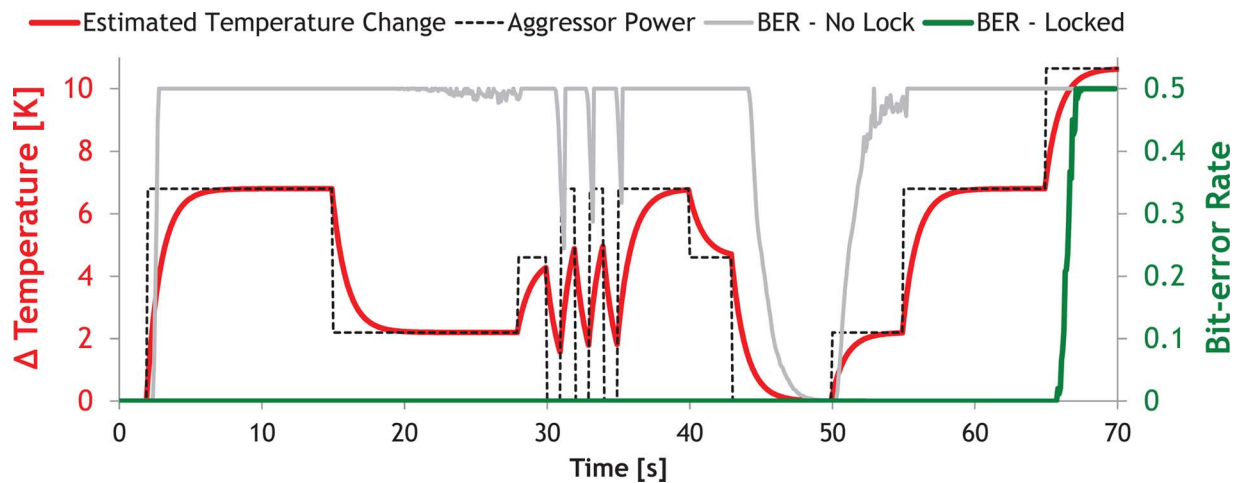


Fig. 19. Transient lock experiment with the tuning controller locked and without lock. The estimated temperature change is calculated from the aggressor power profile using 20 K/W and a time constant of 1s.

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 [19], 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. Currently, the tuning range is limited by the maximum heater output power of 6 mW (limited by the 1 k Ω resistance of the integrated heater and the 2.5 V supply). The tuning range will increase by using a lower heater resistance or by improving the tuning efficiency, which we estimate to be 67 μ W/GHz (108 mW/FSR). As a comparison point, a 3 μ m radius ring on the same platform with optimized heater placement demonstrated a tuning efficiency of 10 μ m/GHz [31].

We perform a 2.5 Gb/s single-rate transient wavelength-lock experiment, shown in Fig. 19. We clock-gate on/off the various components in the digital backend of adjacent transceiver macros to create a temperature aggressor. These induce a mirroring 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. 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, 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 mm², 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.

VII. 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 (Fig. 20). We note that this demonstration uses a single λ and clocks are forwarded to both chips electrically. Normally, in a DWDM configuration with multiple λ s, 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

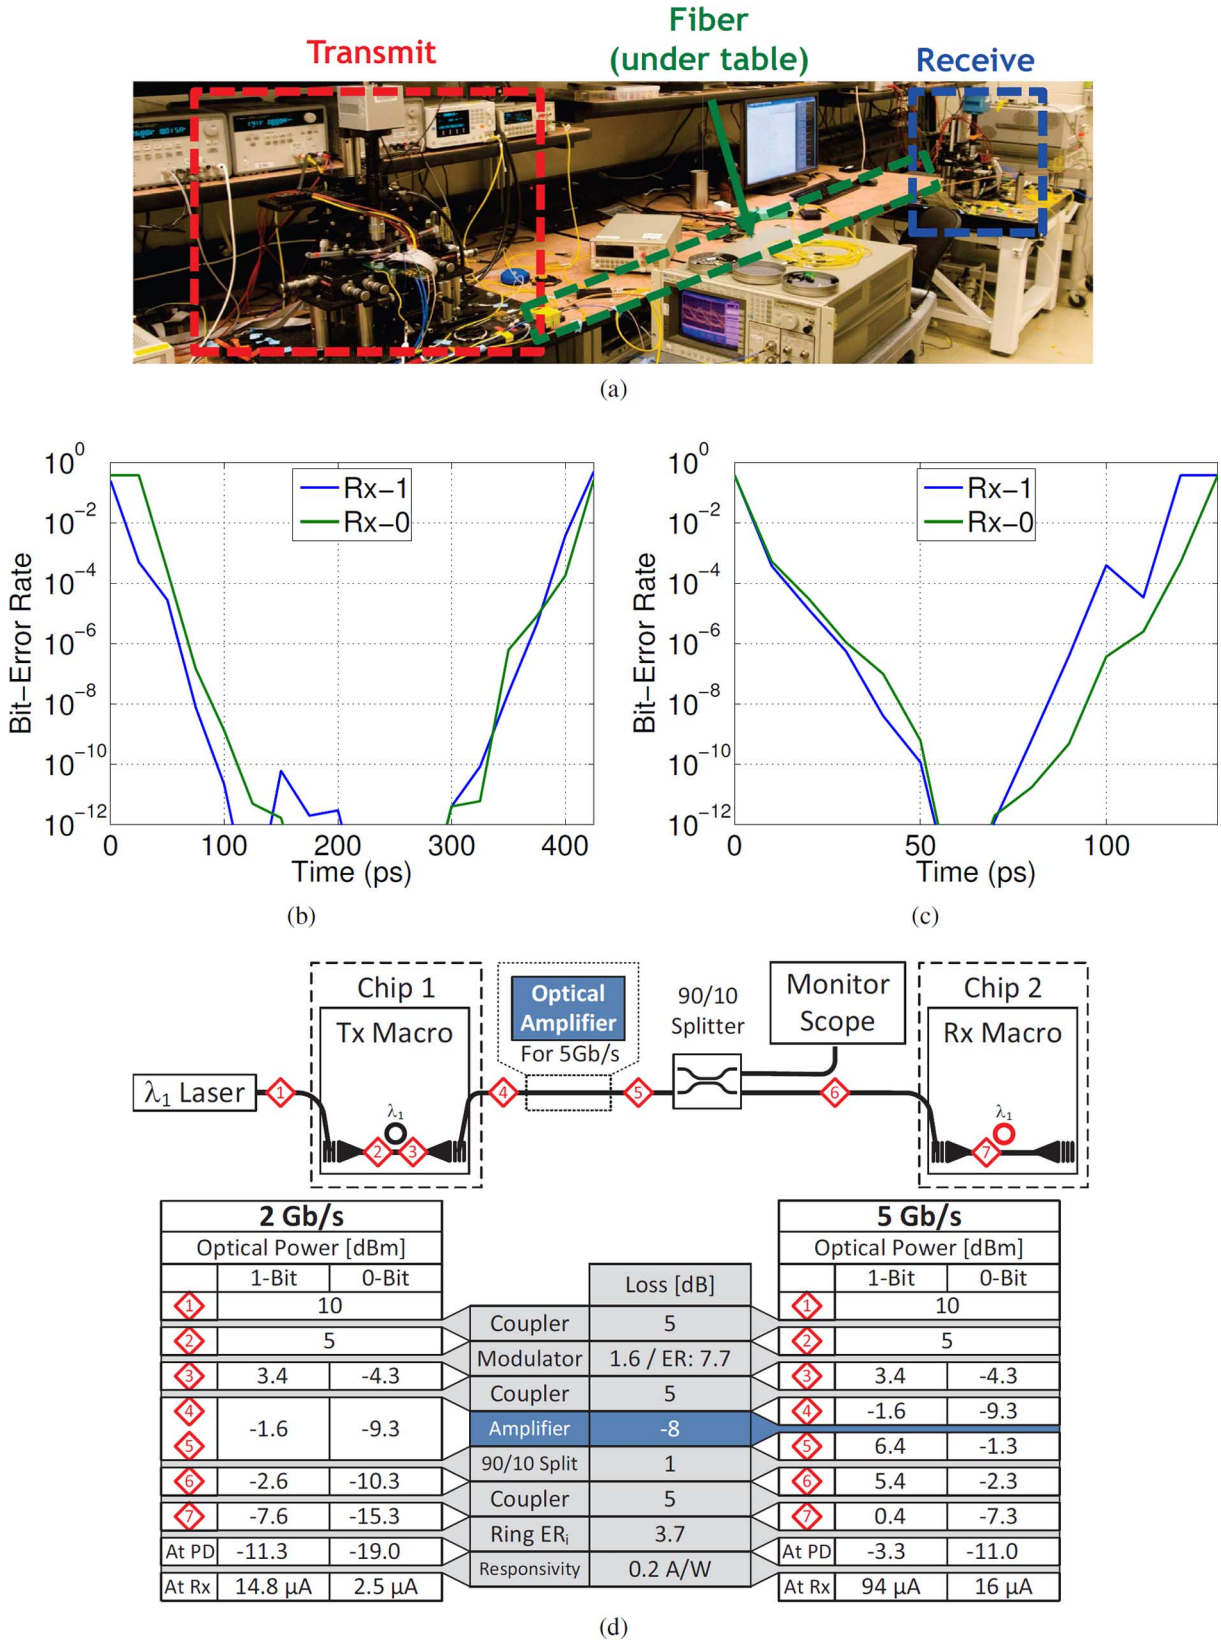


Fig. 20. Chip-to-chip optical link across the room (a), 2 Gb/s (b) and 5 Gb/s (c) link BER measurements for 2×10^{12} bits, and estimated optical power breakdowns of the 2 Gb/s and 5 Gb/s chip-to-link links (d).

link that is error-free for 2×10^{12} bits. 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

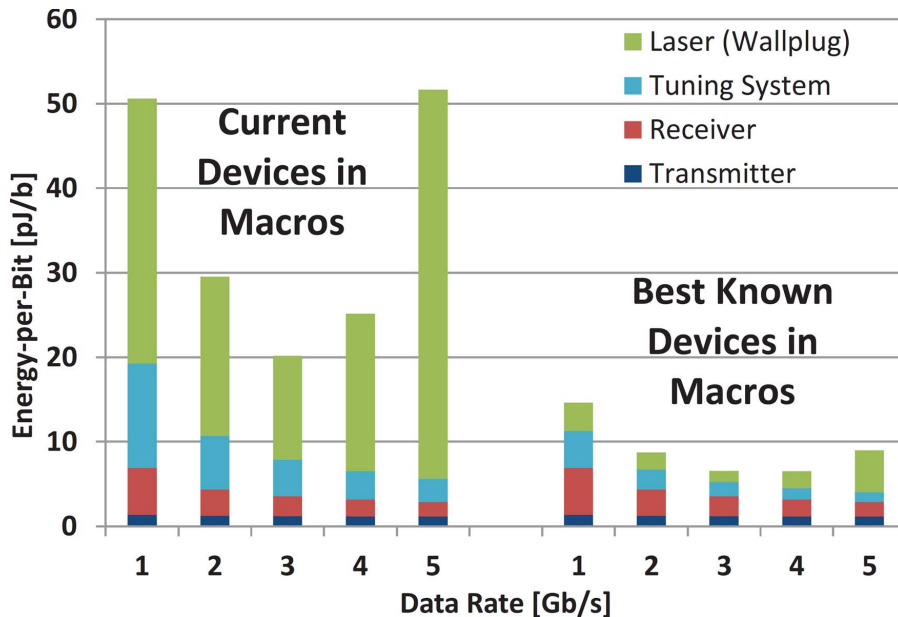


Fig. 21. Link power analysis across data-rates using the current set of devices placed in the transmit/receive macros and if the best-known measured devices are placed into the transmit/receive macros.

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 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. 21). We calculate the wall-plug laser power of the link as

$$P_L = \frac{1}{\eta_e} \cdot \frac{S_i}{R_{pd}} \cdot \frac{1}{1 - 10^{-ER/10}} \cdot 10 \sum L/10 \quad (3)$$

where η_e is the laser wall-plug efficiency, S_i is the receiver sensitivity (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 energy-efficient at a data-rate of 3 Gb/s–4 Gb/s. Static power (mirroring heating/tuning, receiver TIA, the receiver) is amortized at higher rates. Energy-per-bit from the laser also decreases initially from 1 Gb/s–3 Gb/s, where S_i (and hence P_L) remains flat. At the higher rates, S_i increases drastically and offsets the 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 Gb/s–4 Gb/s and an optimal energy efficiency of 6.5 pJ/b is reached at 4 Gb/s.

VIII. CONCLUSION

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 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.

ACKNOWLEDGMENT

The authors gratefully acknowledge the help and support from all POEM team members at Micron Technology Inc., MIT, CU Boulder, and UC Berkeley.

REFERENCES

- [1] I. Young *et al.*, "Optical I/O technology for tera-scale computing," *IEEE J. Solid-State Circuits*, vol. 45, no. 1, pp. 235–248, Jan. 2010.
- [2] S. Beamer *et al.*, "Re-architecting DRAM memory systems with monolithically integrated silicon photonics," in *Proc. 37th Annu. ACM Int. Symp. Comput. Architecture, ISCA'10*, 2010, pp. 129–140.
- [3] N. Binkert *et al.*, "The role of optics in future high radix switch design," in *Proc. 38th Annu. ACM Int. Symp. Comput. Architecture, ISCA'11*, 2011, pp. 437–448.
- [4] S. Assefa *et al.*, "A 90 nm CMOS integrated nano-photonics technology for 25 Gbps WDM optical communications applications," in *2012 IEEE Int. Electron Devices Meeting (IEDM)*, 2012, pp. 33.8.1–33.8.3.
- [5] J. Li, G. Li, X. Zheng, K. Raj, A. Krishnamoorthy, and J. Buckwalter, "A 25-Gb/s monolithic optical transmitter with micro-ring modulator in 130-nm SOI CMOS," *IEEE Photonics Technol. Lett.*, vol. 25, no. 19, pp. 1901–1903, Oct. 2013.
- [6] M. Georgas *et al.*, "A monolithically-integrated optical transmitter and receiver in a zero-change 45 nm SOI process," in *2014 Symp. VLSI Circuits Dig. Tech. Papers*, 2014, pp. 58–59.
- [7] H. Byun *et al.*, "Bulk-Si photonics technology for DRAM interface," *Photon. Res.*, vol. 2, no. 3, pp. A25–A33, Jun. 2014.
- [8] X. Zheng *et al.*, "2-pJ/bit (on-chip) 10-Gb/s digital CMOS silicon photonic link," *IEEE Photon. Technol. Lett.*, vol. 24, no. 14, pp. 1260–1262, Jul. 2012.
- [9] C. Li *et al.*, "Silicon photonic transceiver circuits with microring resonator bias-based wavelength stabilization in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 6, pp. 1419–1436, Jun. 2014.
- [10] M. Nazari and A. Emami-Neyestanak, "A 24-Gb/s double-sampling receiver for ultra-low-power optical communication," *IEEE J. Solid-State Circuits*, vol. 48, no. 2, pp. 344–357, Feb. 2013.
- [11] J.-S. Kim *et al.*, "A 1.2 V 12.8 GB/s 2Gb mobile wide-I/O DRAM with 4x128 I/os using TSV-based stacking," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2011, pp. 496–498.
- [12] 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 *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2012, pp. 186–187.
- [13] J. Buckwalter, X. Zheng, G. Li, K. Raj, and A. Krishnamoorthy, "A monolithic 25-Gb/s transceiver with photonic ring modulators and ge detectors in a 130-nm CMOS SOI process," *IEEE J. Solid-State Circuits*, vol. 47, no. 6, pp. 1309–1322, Jun. 2012.
- [14] J. S. Orcutt *et al.*, "Open foundry platform for high-performance electronic-photonic integration," *Opt. Express*, vol. 20, no. 11, pp. 12,222–12,232, May 2012.
- [15] B. Moss *et al.*, "A 1.23 pJ/b 2.5 Gb/s monolithically integrated optical carrier-injection ring modulator and all-digital driver circuit in commercial 45 nm SOI," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2013, pp. 126–127.
- [16] M. T. Wade *et al.*, "Energy-efficient active photonics in a zero-change, state-of-the-art CMOS process," in *Opt. Fiber Commun. Conf. Opt. Soc. Am.*, 2014, p. Tu2E.7.
- [17] J. S. Orcutt *et al.*, "Nanophotonic integration in state-of-the-art CMOS foundries," *Opt. Express*, vol. 19, no. 3, pp. 2335–2346, Jan. 2011.
- [18] R. Meade *et al.*, "Integration of silicon photonics in a bulk CMOS memory flow," in *2013 IEEE Opt. Interconnects Conf.*, 2013, pp. 114–115.
- [19] M. Georgas, J. Leu, B. Moss, C. Sun, and V. Stojanovic, "Addressing link-level design tradeoffs for integrated photonic interconnects," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, 2011, pp. 1–8.
- [20] J. Leu and V. Stojanovic, "Injection-locked clock receiver for monolithic optical link in 45 nm SOI," in *Proc. IEEE Asian Solid State Circuits Conf. (A-SSCC)*, 2011, pp. 149–152.
- [21] R. Soref and B. Bennett, "Electrooptical effects in silicon," *IEEE J. Quantum Electron.*, vol. 23, no. 1, Jan. 1987.
- [22] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, "12.5 Gbit/s carrier-injection-based silicon micro-ring silicon modulators," *Opt. Soc. Am.*, vol. 15, no. 2, Jan. 2007.
- [23] M. Georgas, J. Orcutt, R. Ram, and V. Stojanovic, "A monolithically-integrated optical receiver in standard 45-nm SOI," *IEEE J. Solid-State Circuits*, vol. 47, no. 7, pp. 1693–1702, Jul. 2012.
- [24] S. Selvaraja, W. Bogaerts, D. V. Thourhout, and R. Baets, "Fabrication of uniform photonic devices using 193 nm optical lithography in silicon-on-insulator," in *Proc. 14th Eur. Conf. Integr. Opt. (ECIO)*, 2008.
- [25] K. Padmaraju, J. Chan, L. Chen, M. Lipson, and K. Bergman, "Thermal stabilization of a microring modulator using feedback control," *Opt. Express*, vol. 20, no. 27, pp. 27,999–28,008, Dec. 2012.
- [26] E. Timurdogan *et al.*, "Automated wavelength recovery for microring resonators," in *Proc. Conf. Lasers Electro-Optics Opt. Soc. Am.*, 2012, p. CM2M.1.
- [27] P. Amberg *et al.*, "A sub-400 fJ/bit thermal tuner for optical resonant ring modulators in 40 nm CMOS," in *Proc. IEEE Asian Solid State Circuits Conf. (A-SSCC)*, 2012, pp. 29–32.
- [28] X. Zheng *et al.*, "A high-speed, tunable silicon photonic ring modulator integrated with ultra-efficient active wavelength control," *Opt. Express*, vol. 22, no. 10, pp. 12,628–12,633, May 2014.
- [29] R. Meade *et al.*, "Integration of silicon photonics in bulk CMOS," in *Symp. VLSI Technol. Dig. Tech. Papers*, 2014, pp. 228–229.
- [30] H.-C. Ji *et al.*, "Optical interface platform for DRAM integration," in *Opt. Fiber Commun. Conf. and Exposition (OFC/NFOEC), 2011 and the Nat. Fiber Optic Engineers Conf.*, 2011, pp. 1–3.
- [31] C. Sun, E. Timurdogan, M. Watts, and V. Stojanovic, "Integrated microring tuning in deep-trench bulk CMOS," in *Proc. IEEE Opt. Interconnects Conf.*, 2013, pp. 54–55.
- [32] S. Assefa, W. Green, A. Rylyakov, C. Schow, F. Horst, and Y. Vlasov, "Monolithic integration of silicon nanophotonics with CMOS," in *Proc. IEEE Photonics Conf. (IPC)*, 2012, pp. 626–627.
- [33] K. K. Mehta *et al.*, "Polycrystalline silicon ring resonator photodiodes in a bulk complementary metal-oxide-semiconductor process," *Opt. Lett.*, vol. 39, no. 4, pp. 1061–1064, Feb. 2014.
- [34] E. Chang, F. Liu, P. Amberg, J. Lexau, and Ho, "Efficient techniques for canceling transceiver noise," in *Proc. Int. Symp. VLSI Design, Automation, and Test (VLSI-DAT)*, 2013, pp. 1–4.
- [35] V. Stojanovic *et al.*, "Autonomous dual-mode (PAM2/4) serial link transceiver with adaptive equalization and data recovery," *IEEE J. Solid-State Circuits*, vol. 40, no. 4, pp. 1012–1026, Apr. 2005.
- [36] C. Sun *et al.*, "A monolithically-integrated chip-to-chip optical link in bulk CMOS," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2014, pp. 56–57.



Chen Sun (S'09) received the B.S. degree in electrical engineering and computer science from the University of California, Berkeley, CA, USA, in 2009, and the S.M. degree in electrical engineering in 2011 from the Massachusetts Institute of Technology, Cambridge, MA, USA, where he is currently a Ph.D. candidate.

He held internships at Nvidia in 2008 and 2009 and at Rambus in 2010. Since 2013, he has been a visiting student at the Berkeley Wireless Research Center, focusing on system-level integration of silicon-photonics. His research interests include design of energy efficient interconnects, modeling and implementation of electro-optic devices and systems, and techniques for dense wavelength-division multiplexed (DWDM) optical transceivers.



Michael Georgas (S'07–M'13) received the B.A.Sc. in engineering science from the University of Toronto, Toronto, ON, Canada, in 2007, the Master degree from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2009, and the Ph.D. in electrical engineering and computer science from MIT in 2013.

He completed a Postdoctoral Associate position in the Research Lab of Electronics Translational Fellows Program at MIT. His research interests lie in the field of integrated photonics and high-speed links, and more recently the application and commercialization of integrated photonics in data-centers.



Jason S. Orcutt received the B.S. degree in electrical engineering from Columbia University, New York, NY, USA, in 2005, and the M.S. and Ph.D. degrees in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2008 and 2012, respectively.

From 2012 to 2013, he worked as a research scientist at MIT. Since 2013, he has been a research staff member in the Physical Sciences department of IBM's T. J. Watson Research Center. His research interests range from ultrafast and integrated optics to

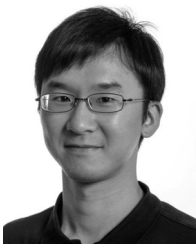
bioelectronics and stochastic systems.

Dr. Orcutt was the recipient of a National Sciences Foundation Graduate Research Fellowship from 2005 to 2008.



Benjamin Moss (S'07–M'08) received the B.S. degrees from the Missouri University of Science and Technology, Rolla, MO, USA, in 2006, the M.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2009, and the Ph.D. degree in electrical engineering and computer science from MIT in 2014.

He is interested in the applications of high-speed photonic interconnects in industry applications



Yu-Hsin Chen (S'11) received the B.S. degree in electrical engineering from National Taiwan University, Taipei, Taiwan, in 2009, and the S.M. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2013, where he is currently working toward the Ph.D. degree.

His research focuses on energy-efficient algorithm, architecture, and VLSI design for computer architecture and computer vision systems.



Jeffrey M. Shainline received the B.S. degree in physics from the University of Colorado, Boulder, CO, USA, in 2005, and the Ph.D. degree from Brown University, Providence, RI, USA, in 2010.

From 2010 to 2013, he was a postdoctoral researcher in the laboratory of Dr. Miloš Popović at the University of Colorado. In 2013, he joined the Division of Optoelectronics and Photonics at the National Institute of Standards and Technology in Boulder, Colorado as a National Research Council Postdoctoral Fellow. His work concerns the integra-

tion of chipscale quantum photonic devices with superconducting single-photon detectors.



Mark Wade received the B.S. degrees in electrical engineering and physics from Louisiana Tech University, Ruston, LA, USA, in 2010, and the M.S. degree in electrical engineering from the University of Colorado, Boulder, CO, USA, in 2013. He is a Ph.D. candidate in Dr. Miloš Popović's group at the University of Colorado Boulder and a Research Affiliate at MIT in Prof. Rajeev Ram's group.

He was a NIST Measurement, Science and Engineering Fellow and is currently a National Science Foundation Graduate Research Fellow. In his graduate work, he has focused on realizing high performance photonic devices in advanced CMOS processes which include efficient photonic filters, optical modulators, grating couplers, and detectors. His interests are in using photonic/electronic integration to create efficient communication links.



Karan Mehta received the B.S. degrees in electrical engineering and physics from the University of California at Los Angeles (UCLA), CA, USA, in 2010, and the S.M. in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 2012. He is currently pursuing research in integrated optics for trapped ion quantum systems.



Kareem Nammari received the M.S. degree in 2014 and the B.S. degree in 2012 in electrical engineering from the University of Colorado, Boulder, CO, USA. He was an M.S. student in Dr. Miloš Popović's group at the University of Colorado Boulder. His graduate work focused on energy-efficient grating couplers in advanced CMOS processes.

His interests include optical-electrical interconnects and swept wavelength sources. He is currently employed with Insight Photonic Solutions.



Erman Timurdogan (S'13) received the B.S. (Hons.) degree in electrical and electronics engineering from Koc University, Istanbul, Turkey, in 2010. He then joined the Photonics Microsystems Group, the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA. He received the S.M. degree in electrical engineering from MIT in 2013 in "automated wavelength recovery for silicon photonics." He is currently a member of the Photonics Microsystems Group at MIT and working toward the Ph.D. degree.

He was with Mercedes-Benz Turk, Istanbul, Turkey, in 2008 and Bruker Optics, Karlsruhe, Germany, in 2009. His current research interests include modeling and testing integrated electro-optic and thermo-optic photonic devices and systems for low-power, high speed, wavelength division multiplexed (WDM) communications. He has more than 35 authored and coauthored conference and journal publications and an issued patent.

Mr. Timurdogan is a member of the Optical Society of America (OSA).



Daniel L. Miller holds a Ph.D. in Theoretical Physics from the Hebrew University in Jerusalem. He worked in the Weizmann Institute of Science from 1996 to 1999 in the field of quantum chaos. In 1999, he joined Intel and worked for 10 years as a device engineer on sustaining and developing numerous technologies. He joined Micron Israel in 2010, working in R&D on 45 nm NOR devices and later on silicon photonics. In 2014, he returned to Intel and is currently working on 14 nm product development.



Ofer Tehar-Zehav worked for Intel and Numonyx on 45 nm NOR process integration, before joining Micron Technology. While at Micron, he integrated silicon photonics process modules covered in this work. In 2013, he returned to Intel in Haifa, Israel.



Zvi Sternberg worked for Intel and Numonyx on 45 nm NOR process integration, before joining Micron Technology. While at Micron, he was an engineering manager for process integration, contributing to the development of silicon photonics process modules.



Jonathan Leu was born in Columbia, SC, USA, and grew up in Michigan, USA, and Hsinchu, Taiwan. He received the Bachelor degree in 2007 from National Chiao Tung University, Hsinchu, Taiwan, and the Master degree in 2010, both in electrical engineering and computer science.

His interests are focused on novel integrated photonic timing circuits, and developing simulation tools that are CMOS compatible.

Johanna Chong received the B.S. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2012. She received the M.Eng. degree in electrical engineering in February 2014 from MIT, having completed her thesis on hybrid-CMOS based multiwavelength lasers.

She is an R&D Electronics Engineer at Sandia National Laboratory, Albuquerque, NM, USA, where she works on embedded systems.



Reha Bafrali received the B.S. degree in chemical engineering from Carnegie-Mellon University, Pittsburgh, PA, USA, in 1984, and the Ph.D. degree in chemical engineering from the University of California, Berkeley, CA, USA, in 1995.

He worked for Intel Corporation from 1995 to 2008, initially as a process engineer on wafer cleans and photolithography before making the transition to developing models and recipes for optical proximity correction (OPC) in 1999. He joined Numonyx BV at its formation in 2008 and Micron in 2010 as part

of the acquisition of Numonyx.



Gurtej S. Sandhu (M'91–SM'06–F'11) received the degree in electrical engineering from the Indian Institute of Technology, New Delhi, India, and the Ph.D. degree in physics from the University of North Carolina, Chapel Hill, NC, USA, in 1990.

He is Senior Fellow and Director of Advanced Technologies Research at Micron Technology Inc. In his current role, he manages the Advanced Memory Materials and Technologies research and developments at Micron. At Micron he has served in a variety of engineering and management roles

responsible for process technology development, pilot manufacturing and transfers to manufacturing. He has been actively involved with a broad range of process technologies for making computer chips. He has been associated with microelectronics technology for over 24 years and has pioneered several technologies which are currently employed in mainstream chip manufacturing. More recently, he was involved with introduction of a number of Atomic Layer Deposition (ALD) based processes and innovative patterning techniques for memory chip technology used by all major chip manufacturing companies around the world.

Dr. Sandhu has served as a guest speaker at numerous international events, authored many technical papers and holds over 1100 issued U.S. patents. He has received several honors and recognition for these contributions from international technical societies and organizations. Dr. Sandhu is a Fellow of IEEE.



Michael R. Watts (M'10) is a principal investigator in the Research Laboratory of Electronics (RLE) and a member of the Electrical Engineering and Computer Science Department (EECS) at the Massachusetts Institute of Technology. He received his Bachelor of Science in Electrical Engineering from Tufts University in 1996. He then joined Draper Laboratory as a Member of Technical Staff in their Fiber Optics Group. In 1999 he became a Draper Fellow and received his SM and PhD degrees from MIT in 2001 and 2005, respectively. In 2005 he joined Sandia National Labs where he led their silicon photonics effort as a Principal Member of Technical Staff. In 2010 he returned to MIT as an Associate Professor.

Prof. Watts' research focuses on photonic microsystems for low-power communications, sensing, and microwave-photonics applications. His current interests include the modeling, fabrication, and testing of large-scale implementations of microphotonic circuits, systems, and networks that are being integrated, directly or through hybrid techniques, with CMOS electronics for high-speed transmitting, switching, and routing applications of digital signals. Additional interests include large-scale microphotonic sensing and imaging arrays, along with optical phased arrays, nanophotonic antennas, nonlinear optics, and manipulations of optical-electromagnetic fields on-chip.

Prof. Watts' research focuses on photonic microsystems for low-power communications, sensing, and microwave-photonics applications. His current interests include the modeling, fabrication, and testing of large-scale implementations of microphotonic circuits, systems, and networks that are being integrated, directly or through hybrid techniques, with CMOS electronics for high-speed transmitting, switching, and routing applications of digital signals. Additional interests include large-scale microphotonic sensing and imaging arrays, along with optical phased arrays, nanophotonic antennas, nonlinear optics, and manipulations of optical-electromagnetic fields on-chip.



Roy Meade (M'98–SM'09) holds the BME and MSME from the Georgia Institute of Technology, Atlanta, GA, USA, and the MBA from Duke University, Durham, NC, USA, graduating as a Fuqua Scholar.

He is a Senior Member of Technical Staff and a Senior Program Manager for High Bandwidth Memory at Micron Technology, Inc. He is also the program manager for Micron's POEM Program, a DARPA funded silicon photonics collaboration with MIT, UC Berkeley, and UC Boulder. He has over 15

years in CMOS manufacturing and R&D spanning ASIC, DRAM, resistive RAM, and photonics. His research interests include the material science and device physics of beyond CMOS devices as well as silicon photonic device fabrication. He is a Senior Member of the IEEE, and has 25 US patents.



Miloš Popović (S'98–M'05) was born in Zajecar, Serbia (formerly Yugoslavia), in 1977. He received the B.Sc.E. degree in electrical engineering from Queen's University, Kingston, ON, Canada, in 1999, and the M.S. and Ph.D. degrees in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2002 and 2008, respectively.

Since 2010, he has been an Assistant Professor in the Department of Electrical, Computer and Energy Engineering, University of Colorado, Boulder, CO,

USA. In 2012, he was named a Fellow of the David & Lucile Packard Foundation. His research interests include silicon photonic circuits, CMOS photonics integration and optical interconnects, nonlinear and quantum integrated photonics, and light-forces based, nanomechanical photonic devices. More generally, his interests are in the application of electromagnetic theory and solid-state physics toward the invention of devices that function based on novel aspects of the physics of light-matter interaction. He is also interested in electromagnetic theory and numerical methods related to photonics.

Dr. Popović has served on the technical program committees for the Optical Fiber Communication conference (OFC), the Conference on Lasers and Electro-Optics (CLEO), Integrated Photonics Research (IPR) and Group IV Photonics (GFP). He is an inventor or co-inventor on over 16 patents, and author or co-author of over 130 research papers and two book chapters.



Rajeev J. Ram (S'94–M'96–SM'07) received the B.S. degree in applied physics from the California Institute of Technology, Pasadena, CA, USA, in 1991, and the Ph.D. degree in electrical engineering from the University of California, Santa Barbara, CA, USA, in 1997.

He is currently a Professor at the Massachusetts Institute of Technology, Cambridge, MA, USA. His research focuses on physical optics and electronics, including the development of novel components and systems for communications and sensing and studies of fundamental interactions between electronic materials and light.



Vladimir Stojanović (S'96–M'04) received the Dipl. Ing. degree from the University of Belgrade, Serbia, in 1998, and the Ph.D. in electrical engineering from Stanford University, Stanford, CA, USA, in 2005.

He is an Associate Professor of electrical engineering and computer science at the University of California, Berkeley, CA, USA. His research interests include design, modeling and optimization of integrated systems, from CMOS-based VLSI blocks and interfaces to system design with emerging devices like NEM relays and silicon-photonics. He is also interested in design and implementation of energy-efficient electrical and optical networks, and digital communication techniques in high-speed interfaces and high-speed mixed-signal IC design. He was also with Rambus, Inc., Los Altos, CA, USA, as a Principal Engineer from 2001 through 2004, and with MIT as an Associate Professor from 2005 to 2013.

Dr. Stojanović received the 2006 IBM Faculty Partnership Award, and the 2009 NSF CAREER Award, as well as the 2008 ICCAD William J. McCalla, 2008 IEEE TRANSACTIONS ON ADVANCED PACKAGING, and 2010 ISSCC Jack Raper Best Paper Awards. He was an IEEE Solid-State Circuits Society Distinguished Lecturer for the 2012–2013 term.