CORA Cork Open Research Archive
Cartlann Taighde Oscailte Chorcaí

Title	Silicon photonic 2.5D multi-chip module transceiver for high-performance data centers
Authors	Abrams, Nathan C.;Cheng, Qixiang;Glick, Madeleine;Jezzini, Moises;Morrissey, Padraic;O'Brien, Peter;Bergman, Keren
Publication date	2020-01-16
Original Citation	Abrams, N. C., Cheng, Q., Glick, M., Jezzini, M., Morrissey, P., O'Brien, P. and Bergman, K. (2020) 'Silicon photonic 2.5D multi-chip module transceiver for high-performance data centers', Journal of Lightwave Technology, 38(13), pp. 3346-3357. doi: 10.1109/JLT.2020.2967235
Type of publication	Article (peer-reviewed)
Link to publisher's version	10.1109/JLT.2020.2967235
Rights	© 2020, the Authors. This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see http://creativecommons.org/licenses/by/4.0/ - https://creativecommons.org/licenses/by/4.0/
Download date	2026-10-30 03:04:00
Item downloaded from	https://hdl.handle.net/10468/10659



UCC
University College Cork, Ireland
Coláiste na hOllscoile Corcaigh

Silicon Photonic 2.5D Multi-Chip Module Transceiver for High-Performance Data Centers

Nathan C. Abrams, Qixiang Cheng, *Member, IEEE*, Madeleine Glick, *Senior Member, IEEE, Senior Member, OSA*, Moises Jezzini, Padraic Morrissey, Peter O'Brien, and Keren Bergman, *Fellow, IEEE, Fellow, OSA*

(Invited Paper)

***Abstract*—Widespread adoption of silicon photonics into data-centers requires that the integration of the driving electronics with the photonics be an essential component of transceiver development. In this article, we describe our silicon photonic transceiver design: a 2.5D integrated multi-chip module (MCM) for 4-channel wavelength division multiplexed (WDM) microdisk modulation targeting 10 Gbps per channel. A silicon interposer is used to provide connectivity between the photonic integrated circuit (PIC) and the commercial transimpedance amplifiers (TIAs). Error free modulation is demonstrated at 10 Gbps with −16 dBm received power for the photonic bare die and at 6 Gbps with −15 dBm received power for the first iteration of the MCM transceiver. In this context, we outline the different integration approaches currently being employed to interface between electronics and photonics—monolithic, 2D, 3D, and 2.5D—and discuss their tradeoffs. Notable demonstrations of the various integration architectures are highlighted. Finally, we address the scalability of the architecture and highlight a subsequent prototype employing custom electronic integrated circuits (EICs).**

***Index Terms*—Optical interconnections multichip module, silicon interposer, silicon photonics, wavelength division multiplexing.**

I. Introduction

With global internet traffic expected to reach nearly 400 exabytes per month by 2022, the demand on data center interconnect bandwidths continues to increase at an exponential rate [1]. Similarly, data center energy consumption trends predict a worldwide 3 PWh electricity usage by data centers by 2030, with worst case estimates as high 8 PWh [2]. Keeping up with the internet traffic requires data center node bandwidths approaching 10 Tb/s [3], all while reducing energy consumption from tens of pJ/bit to sub-pJ/bit to combat the data center energy consumption trends [4]. The number of I/O pins per package tends to double every six years, which is outpaced by total I/O bandwidth doubling every three to four years. Resolving these rate discrepancies requires that every three to four years the bandwidth per I/O is doubled [5].

Increasing the I/O data rate can be attempted by simply increasing the electrical data rate; however, attenuation and intersymbol interference from dispersion become significant issues at higher frequencies. Intersymbol interference can be combatted with additional equalization circuitry, with the tradeoff of higher energy consumption. The data rate at which equalization is necessary depends on the electrical cable type—for an eight-inch high-performance cable, the cutoff is 20 Gbps [6]. Equalization can increase the achievable data rate, but a second cutoff rate exists for which equalization becomes prohibitively energy expensive due to timing path limitations of the CMOS technology node. For the 45 nm and 16 nm nodes the cutoff rates are 20 Gbps and 40 Gbps, respectively [6]. As a result, electrical interconnects are facing a bottleneck in keeping pace with the data center bandwidth and energy consumption requirements.

Integrating silicon photonics into the data center provides a path to keep up with the above requirements due to photonics' minimal signal attenuation, low energy consumption, high bandwidth, and the ability to leverage the mature CMOS ecosystem [7]. Commercial silicon photonic transceivers have already been introduced to the market by Luxtera [8], Elenion [9], MACOM [10], and Intel [11]. Silicon photonic modulators have been demonstrated with 64 Gbaud signals in both Mach Zehnder modulators (MZMs) [12] and microring resonator (MRR) modulators [13]. Additionally, Germanium on Silicon photodiodes have been demonstrated with bandwidths above 100 GHz [14]. While device performance and yield can be improved, one of the paramount challenges for widespread adoption of silicon photonics in data centers is the integration with both the driving electronic integrated circuits (EICs) and the compute nodes: application-specific integrated circuits (ASICs)—such as CPUs, GPUs, and memory—or field programmable gate arrays (FPGAs).

Manuscript received October 15, 2019; revised January 9, 2020; accepted January 10, 2020. Date of publication January 16, 2020; date of current version July 14, 2020. This work was supported in part by the U.S. Department of Energy Advanced Research Projects Agency—Energy under ENLITENED Grant DE-AR000843. *(Corresponding author: Nathan C. Abrams.)*

N. C. Abrams, M. Glick, and K. Bergman are with the Lightwave Research Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027, USA (e-mail: nca2123@columbia.edu; msg144@columbia.edu; bergman@ee.columbia.edu).

Q. Cheng was with the Lightwave Research Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA. He is now with the Electrical Engineering Division, Department of Engineering, University of Cambridge, Cambridge CB3 0FA, U.K. (e-mail: qc223@cam.ac.uk).

M. Jezzini, P. Morrissey, and P. O'Brien are with the Tyndall National Institute, T12 R5CP Cork, Ireland (e-mail: moises.jezzini@tyndall.ie; padraic.morrissey@tyndall.ie; peter.obrien@tyndall.ie).

Color versions of one or more of the figures in this article are available online at https://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JLT.2020.2967235

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see http://creativecommons.org/licenses/by/4.0/

In this paper, we provide a background on silicon photonic integration and report on our efforts in developing integrated multi-chip modules (MCMs). In Section II, we provide an overview of different integration approaches, their tradeoffs, and highlight notable demonstrations. In Section III, we report on the performance of the photonic components used in our MCMs. In Section IV, we present the results of our first version of MCMs. In Section V, we discuss the scalability of our MCM and outline a subsequent prototype currently in fabrication. We finish with conclusions in Section VI.

II. Integration Background

The purpose of introducing silicon photonics into data centers is to increase I/O bandwidth and minimize energy consumption. How the photonic integrated circuit (PIC) is integrated with the driving EICs and the compute ICs can have a major impact on the areal bandwidth density, edge bandwidth density, and packaging parasitics. In turn, these factors directly influence the transceiver's I/O bandwidth and energy consumption—thus, improper integration of photonics with electronics can negate all the potential benefits of silicon photonics. Parasitics between the electronic modulator driver and photonic modulator introduce parasitic poles that impact the modulator's electro-optic frequency response. In the extreme case, the parasitics can be the limiting factor in the modulator's bandwidth. Even if the modulator's bandwidth is acceptable, the effective driving voltage may be reduced, resulting in a smaller optical extinction ratio. Compensating the smaller extinction ratio requires either higher driving voltages or power-hungry multi-tap equalization circuits within the TX or RX ICs [6]. On the receive side, the major concern is the parasitic capacitance. The total photodiode capacitance (photodiode's junction capacitance plus the parasitic capacitance) dictates the transimpedance amplifier's (TIA) dominant pole, which ultimately limits the TIA's bandwidth [15]. Receiver signal-to-noise (SNR) is inversely proportional to the squared receiver bandwidth (BW) multiplied by the total photodiode capacitance (C_{PD}): SNR $\sim (BW^2 \times C_{PD})^{-1}$ [16]. Increased parasitic capacitance results in a lower SNR for a given bandwidth, or requires reducing bandwidth to maintain the SNR. Finally, the components that make hybrid integration possible—pads, wirebonds, solder bumps, copper pillars, and through silicon vias (TSVs)—are difficult to control from an impedance perspective. Reflections from impedance mismatch can reduce the effective driving voltage and introduce additional noise to the signal [5].

MCMs are used to integrate PICs and driving EICs into transceivers, and to integrate these transceivers with compute nodes. In this paper, we will focus on the integration between the PIC and the EIC, and how that integration enables connectivity to compute nodes. The main PIC to EIC integration architectures are monolithic, 2D, 3D, and 2.5D, as shown in Fig. 1. Examples of the integration architectures are outlined in Table I.

A. Monolithic Integration

Monolithic integration is when the photonic components are developed into an existing electronic process node with minimal to no process alterations, as shown in Fig. 1(a). In this approach, parasitics are kept to an absolute minimum as the active photonics and their driving electronics are co-located within the same die. Removing the need for pads and bumps to interface between the PIC and EIC minimizes the potential for impedance mismatch due to packaging. By combining two potential dies into one, the required packaging is simplified—a monolithic transceiver's I/O to a compute node can be through wirebonds or flip-chipped to an interposer.

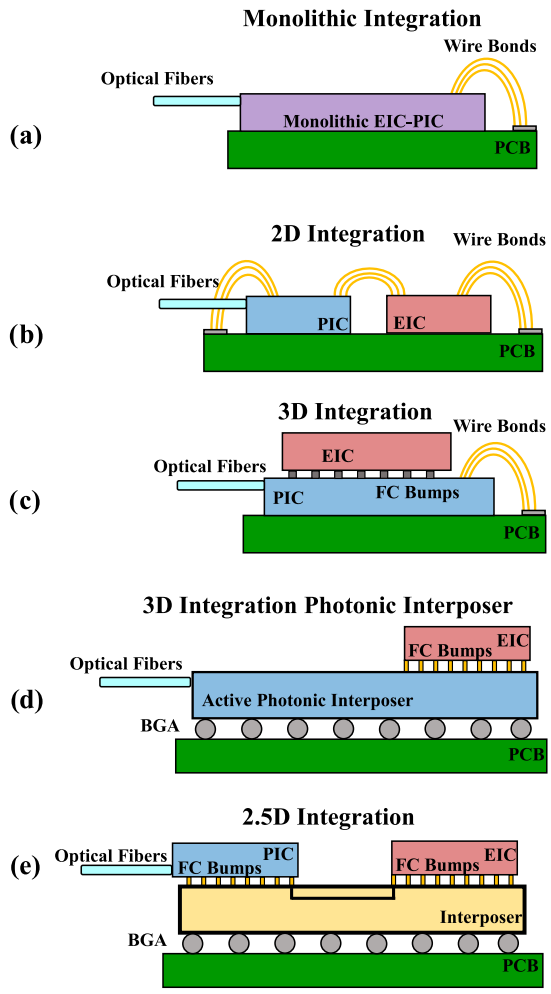


Fig. 1. The various integration approaches to integrate photonic integrated circuits (PICs) with driving electronic integrated circuits (EIC). Flip chip (FC) bumps are used to interface between flipped dies and ball grid arrays (BGAs) are used to interface between an interposer and a PCB. (a) Monolithic integration. (b) 2D integration. (c) 3D integration. (d) 3D integration with a photonic interposer. (e) 2.5D integration.

While theoretically ideal, in practice monolithic integration uses older CMOS nodes which are not optimized for photonic performance. The most advanced nodes developed for monolithic integration are the 45 nm [17] and 32 nm [18] processes, which still lag in performance compared to the cutting edge FinFET 10 nm and below nodes [19]. Additionally, these monolithic processes suffer from high waveguide loss, low photodiode responsivity, and low photodiode bandwidth—in the 45 nm node these values at 1310 nm are 3.7 dB/cm, 0.5 A/W, and 5 GHz, respectively [18]. Other monolithic processes have

TABLE I
NOTABLE DEMONSTRATIONS OF MCM INTEGRATED TRANSCEIVERS

Integration Approach	EIC-PIC Bump Type	Interposer Type	MCM Interface	Channel Data Rate	Number of Channels	Energy/Bit[a]	EIC Node	Modulator Type	Refs.
Monolithic	N/A	N/A	-	25 Gbps NRZ (TX, RX)	4	14.6 pJ/b	90 nm	Mach-Zehnder	[22, 43, 44]
Monolithic	N/A	N/A	Wirebonds	10 Gbps NRZ (TX, RX)	4	65.5 pJ/b	130 nm	Mach-Zehnder	[21]
Monolithic	N/A	N/A	Wirebonds	26 Gbps NRZ (TX, RX)	4	16.25 pJ/b	130 nm	Mach-Zehnder	[24]
Monolithic	N/A	N/A	Wirebonds	10, 7 Gbps NRZ (TX, RX)	10	600 fJ/b	65 nm	Microring	[25]
Monolithic	N/A	N/A	-	8 Gbps NRZ (TX, RX)	4	400 fJ/b	45 nm	Microring	[18]
2D	Wirebonds	N/A	Wirebonds	25 Gbps NRZ (TX)	5	5.34 pJ/b	65 nm	Microring	[27]
2D	Wirebonds	N/A	Wirebonds	112 Gbps PAM-4 (TX)	1	6 pJ/b[b]	28 nm	Microring	[30]
3D	Solder Bumps	N/A	Wirebonds	10 Gbps NRZ (TX, RX)	64, 42 (TX, RX)	530 fJ/b	40 nm	Microring	[45]
3D	Cu-Pillars	N/A	Wirebonds and Probe Pads	56 Gbps NRZ (TX, RX)	2	11.2 pJ/b	130 nm	Mach-Zehnder	[37, 38]
3D	Not Specified	N/A	Wirebonds	20 Gbps NRZ (TX, RX)	4	1.9 pJ/b	40 nm	Microring	[46]
3D	Solder Bumps	N/A	Wirebonds	10 Gbps NRZ (TX, RX)	8	1.3 pJ/b	40 nm	Microring	[47]
3D	Solder Bumps	N/A	Wirebonds	10 Gbps NRZ (TX, RX)	8	1.37 pJ/b	40 nm	Electro-Absorption	[48]
3D	Solder Bumps	Build-Up Substrate	C4 Bumps	25 Gbps NRZ (TX, RX)	1	4.9 pJ/b	28 nm	Microring	[33]
3D	Cu-Pillars	N/A	Wirebonds	25 Gbps NRZ (RX)	1	170 fJ/b[c]	28 nm	N/A	[15]
3D	Cu-Pillars	N/A	Wirebonds	50 Gbps PAM-4 (TX, RX)	2	-	-	Mach-Zehnder	[49]
3D	Solder Bumps	Glass Ceramic	C4 Bumps and Wirebonds	25 Gbps NRZ (TX, RX)	16	5.91 pJ/b	28 nm	Mach-Zehnder	[50]
3D	Cu Bumps	N/A	Wirebonds and Probe Pads	40 Gbps NRZ (TX, RX)	4, 1 (TX, RX)	229 fJ/b	14 nm	Microring	[31]
3D Photonic Interposer	Solder Bumps	Silicon	Cu-Pillars	40 Gbps NRZ (TX, RX)	12	4.4 pJ/b[b]	-	VCSELs	[51]
2.5D	Not Specified	LTCC Substrate	BGA	200 Gbps 16-QAM (TX, RX)	1	-	-	Mach-Zehnder	[42]
2.5D	Cu-Pillars	Silicon	Wirebonds	50 Gbps NRZ (TX, RX)	4	-	-	Mach-Zehnder	[41]
2.5D	Cu-Pillars	Ceramic	BGA	272 Gbps DP 16-QAM (TX, RX)	1	-	-	Mach-Zehnder	[52]
2.5D	Solder Bumps	Silicon	BGA	[Targeted] 10 Gbps NRZ (TX, RX)	4	-	-	Microdisk	This work

[a]Laser, thermal tuning and stabilization, clocking and serialization not included.
[b]Only TX energy/bit.
[c]Only RX energy/bit.

demonstrated improved photonic performance at the expense of decreased transistor performance, as they exist in technology nodes between 90 nm and 250 nm [20]–[23]. As energy per bit often scales with technology node size, monolithic integration in older CMOS nodes results in greater than 14.5 pJ/b EIC transceiver consumption [21], [22], [24]. The monolithic integrations that use 65 nm nodes and below have sub pJ/b consumption for the EIC drivers [18], [25]. From a practical perspective, the process development cost for monolithic is very expensive, and the resulting technology is less flexible than heterogenous process development. Luxtera initially developed a monolithic transceiver [24], but ultimately switched to hybrid 2.5D integration for the above reasons [26].

B. 2D Integration

On the other end of the integration spectrum is 2D integration. In this approach, the PIC and EIC are placed side by side, typically on a PCB, as shown in Fig. 1(b). Wirebonds are used to connect the two dies and to interface to the PCB. An example of this can be seen in [27], where a five channel microring transmitter operating at 25 Gbps per channel PIC was wirebonded to an EIC that provided PRBS generation, differential drivers, and thermal stabilization. The benefit of 2D integration is the ease of packaging; however, the reliance on wirebonds has its drawbacks. While wirebonds can reach pitches of 25 μm [28], the connections between the PIC and EIC are

limited to a single edge, severely limiting the aggregate I/O. Additionally, wirebonds introduce a significant parasitic inductance in the typical range of 0.5–1.0 nH/mm [29]. Techniques such as ribbon bonding can reduce the parasitic inductance, but do so by using ribbons that are over 100 μm wide, placing further limitations on the wirebond pitch. Limiting the I/O to the edge of the EIC also impacts the connection to the compute node, as the remaining edge space must accommodate I/Os to the compute node, grounds, supplies, and bias signals. The two demonstrations outlined in Table I consume more than 5 pJ/b for the EIC drivers [27], [30], which is an order of magnitude higher than the best 3D integration demonstrations.

C. 3D Integration

Minimizing packaging parasitics between the EIC and PIC makes 3D integration an attractive solution. In 3D integration, the EIC is flip chipped on top of the PIC, as shown in Fig. 1(c). Typically, the I/Os to the compute node and DC signals for both the EIC and PIC are routed out through to the edge of the PIC and wirebonded to a PCB. The most common types of flip chip bumps (FC bumps) are copper pillars and microsolder bumps. Copper pillars between PICs and EICs have been demonstrated with parasitic capacitances below 30 fF, parasitic resistances below 1 Ω, and neglible parasitic inductance [31]. Microsolder bumps between PICs and EICs have been demonstrated with parasitic capacitances below 25 fF and parasitic resistances below 1 Ω [32]. Microsolder bumps and copper pillars in transceiver MCMs have achieved similarly dense pitches—in the range of 40 μm to 50 μm [31], [32], [33]. Microsolder bumps and copper pillars are expected to be able to reach pitches of 20 μm and 10 μm, respectively [34]. For denser pitches, direct bonding Cu-Cu has been demonstrated at 15 μm pitch [34], and is expected to be able to support pitches below 4 μm [35]. When fabricating bumps at these ultra-dense pitches, the most common approach is bump growth at wafer scale. While this is acceptable for transceiver production at large scale, procuring a full wafer can be unfeasible for smaller prototypes utilizing multi-project wafer (MPW) runs. However, there are some vendors offering bare die bumping at dense pitches for solder bumps and copper pillars—35 μm and 25 μm, respectively [36]. The 3D integrated demonstrations exhibit a range of energy consumptions per bit. The energy per bit roughly scales with the EIC technology node. On one end, the 14 nm node EIC transceiver consumes 229 fJ/b [31]; on the other end, the 130 nm node EIC transceiver consumes 11.2 pJ/b [37], [38] While 3D integration allows for dense pitches and minimal packaging parasitics, it does not provide the best approach for thermal dissipation from the EIC and offers minimal thermal isolation between the EIC and PIC. Heat generated by the flip chip assembled EIC can be transferred to the PIC. Temperature increases for the PIC of 20 °C above ambient temperature have been observed for IC input powers of approximately 0.5 W [39]. This can be especially problematic for thermally sensitive photonic components, such as microrings. Moving beyond lab prototypes can require thermo-electric coolers (TECs) with unique configurations to allow for thermal management [40]. A second potential constraint of 3D integration is the I/O to the compute node, as the most common method of interfacing to the 3D integrated PIC-EIC is through wirebonds to the PIC. This places limitations on the transceiver bandwidth and introduces non-trivial parasitic inductance.

Interposers can be used to combat bandwidth limitations and parasitic inductance. An interposer is a substrate that commonly serves as passive electrical redistribution for active chips that sit on top of the interposer. The interposer's redistribution is commonly for fanning signals out to larger pitches or for enabling connectivity between multiple active chips. The material used to construct the interposer varies, but common types are silicon, glass, and organic substrates. With silicon interposers, through silicon vias (TSVs) can be used to enable connectivity between the front and back of the interposer. Passive optical redistribution can also be included in silicon interposers by adding a silicon nitride waveguide layer to route optical signals between photonic chips. Finally, active silicon photonics can be incorporated into silicon interposers to form active photonic interposers to combine the interposer's redistribution functionality with the PIC's functionalities. One specific 3D integration solution utilizing an interposer is to hang the EIC partially off the edge of the PIC, and then use a high-density glass ceramic interposer (CGIP) to interface to the package [41]. A second potential solution is to replace the PIC with an active photonic interposer to enable dense, low-parasitic I/O on both the front and back side of the active photonic interposer [5], [42], as shown in Fig. 1(d).

D. 2.5D Integration

Active photonic interposers are an ideal combination of heterogenous process optimization, low EIC-PIC packaging parasitics, and high I/O bandwidth to a compute node. However, they are still a relatively recent technology, and as a result come with high process development costs and long fabrication timelines. A compromise is 2.5D integration, where both the EIC and PIC are flip chipped onto a passive interposer with TSVs to interface to the back side, as shown in Fig. 1(e). The interposer can be electrical only, or also include a waveguide layer for optical redistribution. The flip-chipping can be achieved with micro solder bumps or copper pillars with the same pitch and parasitics as with 3D integration. If the compute node die is flipped onto the same interposer, the I/O connections can be made with micro solder bumps or copper pillars. A drawback of 2.5D integration is that the parasitics between the EIC and PIC are higher than 3D integration as critical signals must pass through two bump interfaces and a trace on an interposer. However, this additional source of parasitics can still be manageable, as dies on an interposer can be placed several hundred micrometers apart. For longer connections to the compute node, high quality transmission lines can be developed within silicon interposers, with 8 mm long coplanar waveguides transmission lines demonstrating better than 50 GHz bandwidth with less than 2 dB S21 loss [41]. If the compute node is housed in a larger package, connections from the back of the transceiver's interposer can be made to a PCB with a ball grid array (BGA) [42], [52], as shown in Fig. 2. Standard BGA packages can be

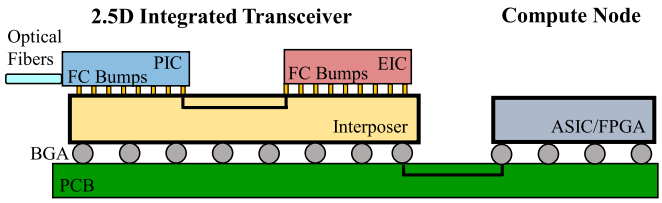


Fig. 2. The integration approach for how the 2.5D integrated transceiver would integrate with the compute node, such as an ASIC or FPGA. While the 2.5D approach is highlighted here, other integrations would interface to a compute node in a similar manner.

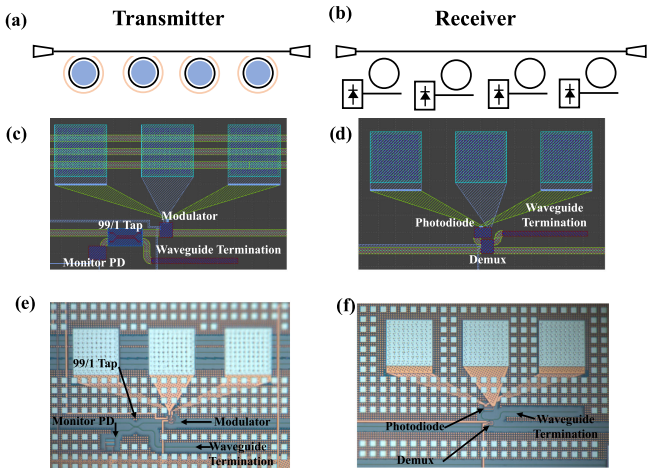


Fig. 3. The architecture of the transceiver. (a) Schematic of the transmitter. (b) Schematic of the receiver. (c) Layout of one channel of the transmitter. (d) Layout of one channel of the receiver. (e) Image of one channel of the fabricated transmitter. (f) Image of one channel of the fabricated receiver.

developed with pitches as low as 0.5 mm [53], with limitations a result of PCB spacing and routing. PIC integration to a PCB through a BGA package has demonstrated signal paths with better than 3 dB insertion loss with a bandwidth of 40 GHz [54]. The 2.5D demonstrations did not report energy per bit metrics. For our first iteration of prototypes, we elected to use a 2.5D integration approach to provide a balance of performance with cost and fabrication turnaround.

III. MCM 2.5D Integrated Prototype Architecture

A. Photonic Architecture

The silicon photonic transceiver architecture, for both transmitter and receiver, is built on bus waveguides with microdisk elements coupled to the waveguide, as shown in Fig. 3. The photonic integrated circuit was fabricated at AIM Photonics on an MPW using the AIM process design kit (PDK) components. On the transmit side, four microdisk modulators are coupled to the bus waveguide. The modulators were measured to be evenly spaced across a 25 nm free spectral range (FSR), resulting in channel spacings of approximately 6.25 nm. A 99/1 tap is positioned after each modulator, with the 99% output continuing on the bus waveguide. The 1% output leads to a monitor photodiode, which allows for feedback signals for thermal stabilization of the microdisk modulators. The difference between the monitor photodiode before and after a modulator corresponds to the power coupled into that modulator. Thermal stabilization architectures have been demonstrated with a single monitor photodiode at the end of a bus of microring modulators, but they typically require measuring RF power [55] or dithering [27]. Other thermal stabilization architectures require a drop port on the modulator [56], [57], reducing the Q factor of the ring. Using a more direct measurement of the power coupled into the modulator, without requiring a drop port, allows for a simpler implementation of the feedback. The bus waveguide is routed to edge couplers at a 127 μm pitch to allow for coupling to a standard zero polish fiber array. On the receive side, four microdisk demux filters are coupled to the bus waveguide. The resonances are fabricated to match the four transmit resonances. At the drop port of each demux filter is a high-speed photodetector. The end of the bus waveguide is split with a 50/50 tap: one end is routed to a monitor photodiode and the other end is routed to the edge coupler array. The thermal stabilization feedback signal is simpler on the receive side for two reasons. One, the output of the high-speed photodetectors can be used as the feedback signal to thermally stabilize the data dropped by the demux microdisks. Two, the optimal operating condition is directly on resonance, unlike the transmit side.

The goal for this iteration of prototypes was to provide a highly scalable, ultralow energy consumption, and very dense bandwidth platform that future prototypes could build on. For this reason, the target data rate for each channel is 10 Gbps. The microdisk modulator occupies approximately 250 μm^2 and the combination of the microdisk demux and high-speed photodiode occupies approximately 750 μm^2. The small microdisk footprints allows transceivers to designed very densely and to still achieve system data rates in the Tbps regime [31]. Additionally, microdisk modulators have lower device capacitances and typically require smaller swing voltages compared to Mach Zehnder modulators [58], allowing the microdisk modulators to consume less modulation power. To fully compare the Mach Zehnder and microdisk modulators, the thermal power consumption for tuning the microdisks must be considered. In [18], the thermal tuning control logic contributed 18 fJ/b of the total 700 fJ/b for the EIC. The energy consumption for microdisk heaters is very dependent on link design and the fabrication variability of microdisks. Thermal tuning energy per bit is dependent on data rate because the thermal tuning power is independent of rate, so higher data rates will reduce the energy per bit. If the laser wavelengths are fixed, then the burden of accounting for resonance fabrication variability is placed on microdisk's thermal heaters. The standard deviation for the microdisk's resonance is 2.14 nm [59]. If the thermal heaters are to account for this variation, the operating wavelength must be set above the expected resonance because the heaters can only tune the microdisk in one direction. Ensuring that 95% of the fabricated resonances fall below the operating wavelength requires an operating wavelength 3.5 nm above intended fabricated resonance. Using the same thermal tuning efficiency and channel rate as [18] results in an average expected heater energy per bit of approximately 186 fJ/b. Paths towards reducing the heater energy per bit are operating at higher

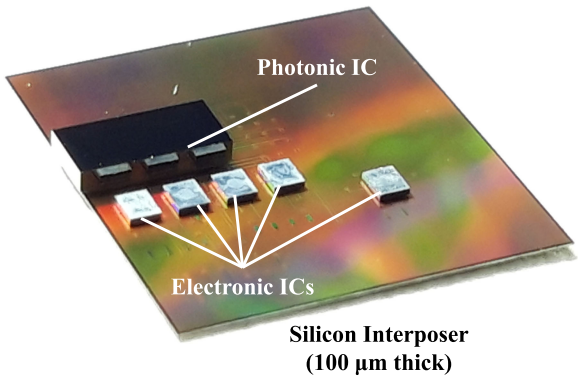


Fig. 4. The assembled MCM showing the silicon interposer, which provides connectivity between the four TIA EICs and the PIC. The isolated EIC is not connected to the PIC and used for TIA-only tests.

data rates, developing more efficient heaters, and reducing the fabrication variability.

B. Integration Architecture

The integration approach for the prototype was a 2.5D integrated MCM. The 2.5D integration is achieved with a silicon interposer. The silicon interposer has TSVs to provide a connection between the metal layers on the front side and the back side of the interposer. Both the PIC and EICs are flipped on top of the interposer, as shown in Fig. 4. The backside of the interposer is used to connect to a PCB to fan out the signals and to house decoupling capacitors and resistors for Received Signal Strength Indicators (RSSIs). Assembly of the prototype was done at the Tyndall National Institute. The flip chip connections for both the PIC and EIC to the interposer is done using stud bumps at a minimum of 100 μm pitch. Stud bumps are gold wire bond segments that are bonded to aluminum pads on the PIC, EIC, and interposer. The initial wirebond can be attached to the aluminum pads and serves as a pillar that the solder can adhere to. BGA-type connections on 300 μm pads are used to connect the backside of the interposer to the PCB with ENEPIG finish at 500 μm pitch. The PIC and EICs are connected to the interposer via stud-bumps prior to attaching the interposer to the PCB, so the reflow temperature of the BGA-type solder must be below the reflow temperature of the stud bump solder. The PIC is positioned such that it overhangs off the interposer by 200 μm to allow for a clear interface to the edge couplers and to provide a visual sight of the edge couplers to aid alignment.

C. Transceiver Architecture

For the first iteration of prototypes, the transceiver design used a commercial Texas Instruments ONET8551 TIA. The driver features a TIA, voltage amplifier, and output buffer. The TIA is designed for wirebonds, but the pad pitch is 115 μm and the aluminum terminated pads were ideal candidates for stud bumps, making flip chipping feasible. While the TIA is single channel, its small size (870 μm by 1036 μm) make it possible to place four dies to interface to the four received channels on the PIC. The 9 GHz bandwidth is designed for a maximum data rate of 11.3 Gbps, which matches with the target photonic rate of 10 Gbps. The TIA has a sensitivity of -20 dBm and differential outputs. A single VCC of $+3.3$ V can be used to power both the input stage for the TIA and the output stage for the voltage amplifier and buffer. The EIC consumes 9.2 pJ/b when receiving a 10 Gbps signal. The EIC also features RSSIs that are connected to ground via surface mount resistors on the PCB to translate to a voltage indicator signal. An additional test TIA is placed on the same interposer, designed to receive external input signals from SMAs on the PCB for debugging purposes. For this prototype, the microdisk modulators on the PIC are externally driven as we were unable to procure an appropriate commercial bare die modulator driver that had a pad configuration conducive to flip chipping. Future prototypes, outlined in Section V, will feature separate custom TX and RX EIC drivers flip chipped onto the interposer.

The fabricated PIC is 4.15 mm by 1.87 mm, and the photodiodes are evenly spaced out at a 0.75 mm pitch along the edge of a 4.15 mm. The spacing between the TIAs and between the TIA and PIC is 0.5 mm. As a result, the TIAs are not able to be placed directly in line with the photodiodes on the PIC. The top-level aluminum trace between the photodiode and the TIA input is 30 μm wide and ranges from 0.75 mm to 2.6 mm long. Each photodiode to TIA input connection is a cathode-anode-cathode configuration, and the TIA provides a 2.35 V reverse bias to the photodiode. The trace configuration and interposer stack-up yields an approximate effective relative permittivity of 2. As a result, a 10 GHz signal yields a $\lambda/4$ length of 3.75 mm. This means that even the longest trace length can be treated as a lumped element, meaning that reflections from impedance mismatch should not impact performance.

All electrical signals—DC and RF signals for the PIC and EIC—are routed to the back of the interposer through the TSVs to BGA-type connections. The RF signals use coplanar waveguide and microstrip transmission lines on the PCB to fan out to SMAs. The DC signals connect to a dense DC mezzanine to interface to a second DC fanout board. All optical connections on the PIC route to edge couplers at a 127 μm pitch. The overhanging PIC is coupled with a standard flat polish fiber array, which is connected to the prototype's mechanical substrate.

IV. MCM 2.5 Integrated Prototype Performance

A. Photonic Die Performance

In this section, we present the performance of the bare photonic die. The DC response for the modulator's heater can be seen in Fig. 6(a). The heater was tested up to 7.5 volts, which produced over a 3.5 nm shift in the microdisk's resonance. The heater efficiency for these modulators was extracted to be 0.54 ± 0.05 nm/mW. These measurements allowed additional parameters to be extracted: the FSR was 24.3 nm, the full width at half maximum (FWHM) was 0.18 nm, and the extinction ratio (ER) was 12 dB. The DC response of the depletion modulator can be seen in Fig. 6(b). A voltage range of 0 to 2 volts was applied to the depletion modulator. From the response, the modulator efficiency was extracted to be 62 ± 2 pm/V. As the reverse bias voltage is increased, the ER of the ring decreased, demonstrating

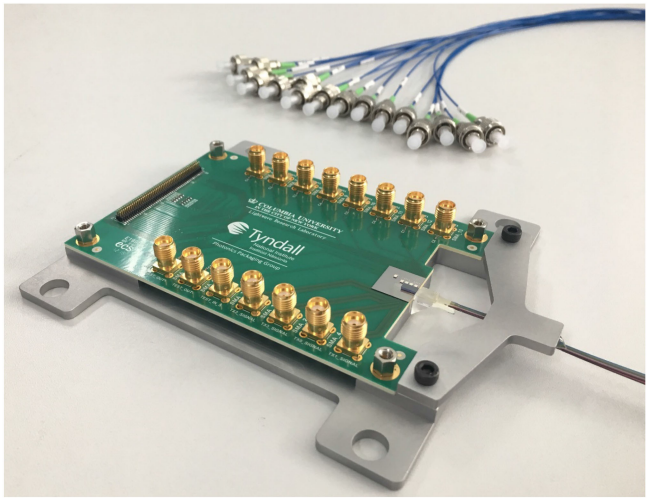


Fig. 5. The fully packaged MCM 2.5D integrated transceiver. The interposer sits on a fan out PCB, which sits on a mechanical substrate that mounts to an optical table for testing. The optical fiber array couples to the PIC and connects to the mechanical support in several locations.

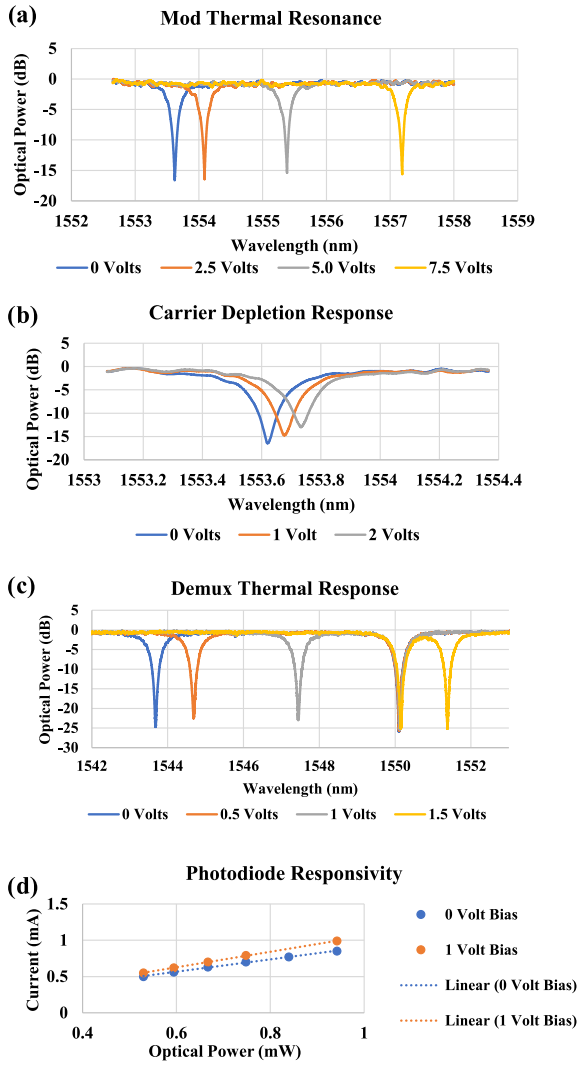


Fig. 6. The DC performance of the bare die PIC. (a) The resonance response of the microdisk modulator's heater. (b) The resonance response of the depletion modulator. (c) The resonance response of the demultiplexer microdisk's heater. (d) The responsivity of the photodiode.

that the microdisk was overcoupled. The eye diagram of the microdisk modulator can be seen in Fig. 7(a). The eye diagram at 10 Gbps was recorded with a drive voltage of 1.5 Vpp biased at −0.5 volts. The measurements for the modulator's bandwidth and the bit error rate curve can be seen in Fig. 7(b) and 7(c), respectively. From these measurements, it can be seen that the modulator produces an error free rate (1E-12) at a received power of −16.0 dBm using a commercial receiver. The bandwidth was measured to be 20 GHz. The demux filter's through response can be seen in Fig. 5(c). From the response, the heater's efficiency was extracted to be 0.66 ± 0.03 nm/mW. From this resonance response, static parameters can be extracted: the filter's FSR was 25.1 nm, the filter's FWHM was 0.54 nm, the filter's ER was 24 dB. The photodiodes at the drop port of the demux filter are designed to be operated with 1 volt of reverse bias. The measured DC photodiode response can be seen in Fig. 6(d). The input power was varied and the output current from the photodiode was measured. This was done with 0 volts and 1 volt reverse bias. Higher voltages were tested, but there was minimal difference from the 1 volt reverse bias. The photodiode responsivity was 0.94 A/W when the photodiode was biased at 0 volts. The photodiode responsivity was 1.11 A/W when the photodiode was reverse biased at 1 volt.

B. MCM Prototype Performance

In this section, the performance of the assembled prototype is presented. Initial coupling losses were measured to be approximately 5 dB per facet. The DC responses of the heaters for the TX and RX were characterized, as well as the DC response of the PN depletion modulator—these can be seen in Fig. 8. The heater efficiency was extracted to be 0.49 ± 0.01 nm/mW. The RX heater response was used to extract the demux heater efficiency, which was determined to be 0.71 ± 0.10 nm/mW. The energy per bit for thermally tuning the microdisks will be dependent on the resonance shift required. If the required resonance shift is 2.14 nm—the standard deviation in the fabricated microdisk resonance reported in [59]—then at 10 Gbps the heaters will consume 437 fJ/b and 301 fJ/b for the modulator and demux, respectively. The modulation efficiency was extracted to be 60 ± 3 pm/V. The bandwidth response of the four channels on the TX of the prototype can be seen in Fig. 9(a). While the response of the four channels are in good agreement, a significant drop in the bandwidth occurs at 8.5 GHz. We are currently investigating to determine the exact cause of the bandwidth limitation. The two most likely culprits are impedance mismatch (which would be resolved when the modulator driver is integrated near the microdisk modulator) and electrical resonance from a combination of interposer trace parasitics and packaging parasitics. The interposer design for the connection to the microdisk modulators was not impedance controlled, as the main concern was keeping the trace length below λ/4 at 10 GHz. Subsequent

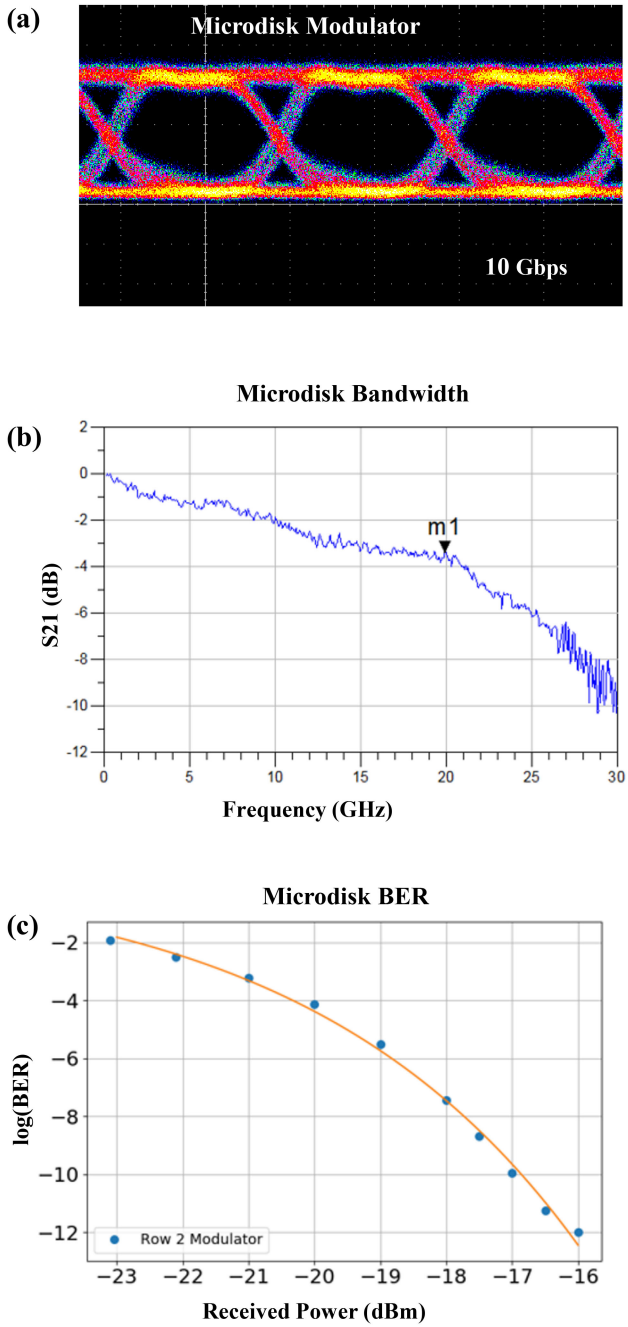


Fig. 7. The microdisk modulator's AC performance. (a) The eye diagram at 10 Gbps driven with a 1.5 V_{PP} signal. (b) The bandwidth of the modulator. (c) The BER at 10 Gbps driven with a 1.5 V_{PP} signal.

prototypes were designed with 50-ohm transmission lines connecting to the RF I/Os of the EIC. Additionally, subsequent designs have data generation occurring within the TX EIC, and deserialization occurring within the RX EIC to bring a 10 Gbps data rate down to 2.5 Gbps, which is supported by this prototype's bandwidth. The BERs for various data rates can be seen in Fig. 8(c). For data rates of 6 Gbps and below, error free performance is achievable at −15 dBm received power or better. Data rates at or above the electrical bandwidth degradation produced significantly worse BERs.

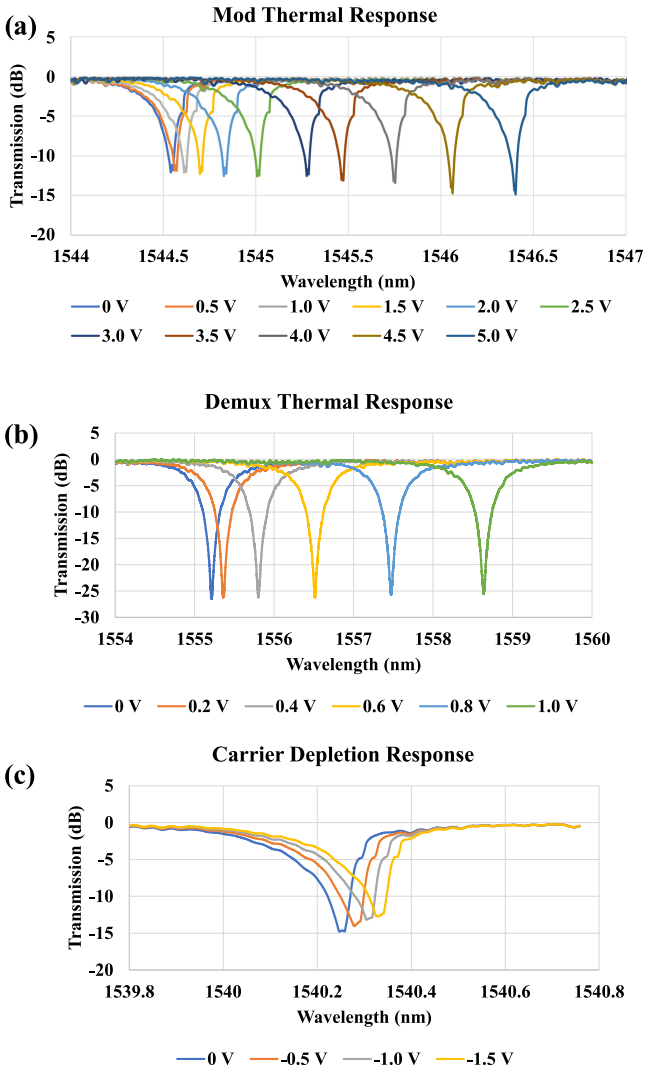


Fig. 8. The DC performance of the PIC components of the MCM integrated transceiver. (a) The resonance response of the microdisk modulator's heater. (b) The resonance response of the demultiplexer microdisk's heater. (c) The resonance response of the depletion modulator.

To verify functionality of the test TIA, a small current supply was approximated by placing a voltage signal in series with a 22 kΩ resistor to provide a Thevenin equivalent of a μA current supply. For this test, alligator clips were used to clip to the resistor, which degraded the bandwidth and added noise to the test. A 400 MHz square wave signal was varied from 0.05 V_{PP} to 1.1 V_{PP} to replicate a current source of 7 μA_{PP} to 50 μA_{PP}. We see that with a current source of 14 $\mu A_{PP,}$ the output was a clean 0.16 V_{PP} square wave. At current swings below 14 μA_{PP}, the signal was severely degraded, as seen in Fig. 10. Increasing the current source's swing beyond this threshold did not impact the output voltage, as the EIC has a limiting amplifier in the output portion of the driver. The TIA was expected to exhibit a sensitivity of −20 dBm. With a 1 A/W photodiode, this would correspond to an input current swing of 10 μA_{PP}, which is in relative agreement with the results observed from the test TIA. Testing the photodiodes with the TIAs will occur when a second

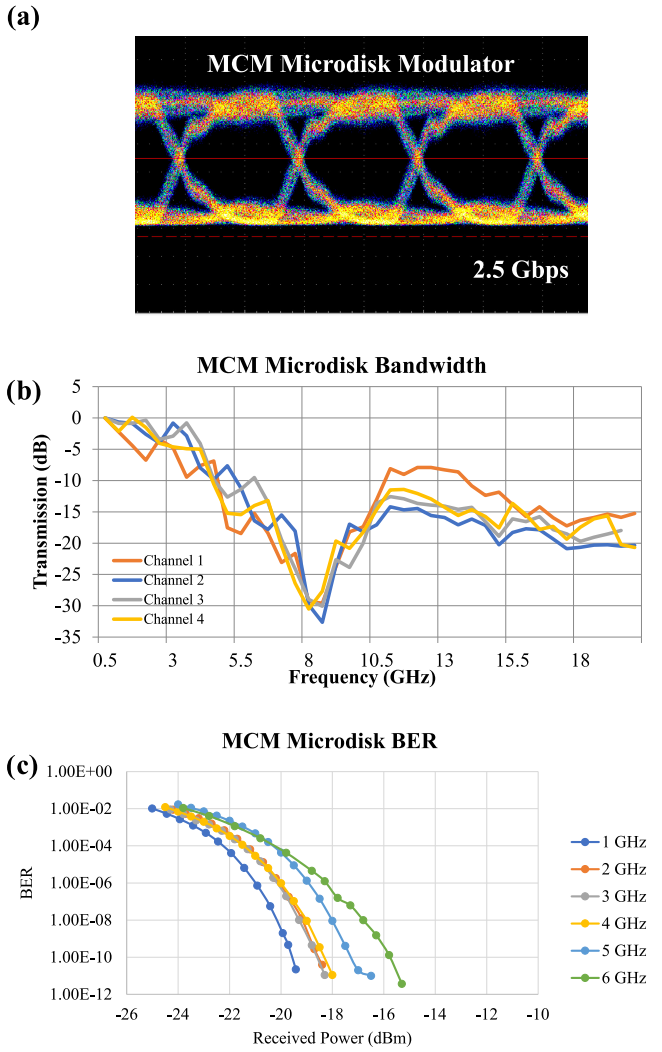


Fig. 9. The microdisk modulator's AC performance of the PIC components of the MCM integrated transceiver. (a) The eye diagram at 2.5 Gbps driven with a 1.5 V_{PP} signal. (b) The bandwidth of the microdisk modulators of the four channels. (c) The BER driven with a 1.5 V_{PP} signal at different data rates.

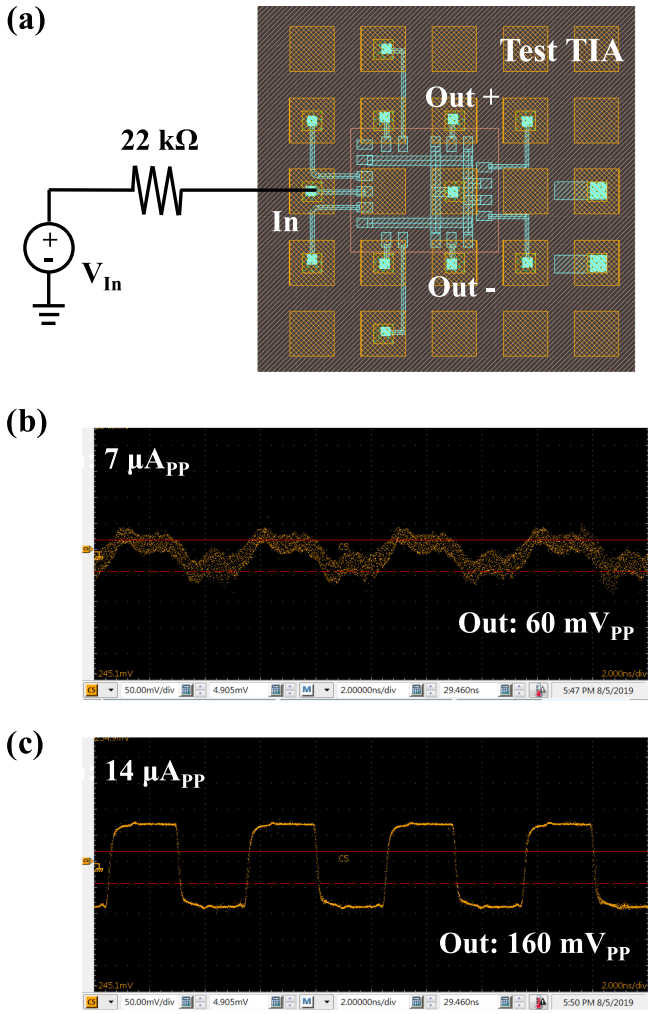


Fig. 10. The results of the test TIA of the MCM prototype. (a) the test setup to mimic a small swing input current. (b) The signal from Out+ when driven with the equivalent of a 7 μA_{PP} input signal. (c) The signal from Out+ when driven with the equivalent of a 14 μA_{PP} input signal.

prototype is repackaged to correct for an anode and cathode mismatch.

V. MCM Scalability and Subsequent Prototypes

Two main approaches for scaling up the total throughput of a WDM MCM are increasing the data rate per wavelength and increasing the number of wavelengths. Scaling up the data rate per wavelength can be achieved by multiplexing channels of data together on the same wavelength or by increasing the data rate on the wavelength. Multiplexing multiple channels to a single wavelength can be achieved via polarization division multiplexing (PDM) [60] and mode division multiplexing (MDM) [61]. PDM is a mature enough approach that silicon photonic foundries are beginning to offer polarization splitters and rotators on MPW runs [62], [63]. One of the drawbacks for PDM is polarization drift that will occur in fiber-based transmission. While work has been explored to expand PDM beyond dual polarization systems [64], practical integrated systems are limited to the dual polarization of TE and TM, limiting the scalability of PDM. MDM presents another avenue for multiplexing on a single wavelength. While challenges related to MDM exist—such as mode crosstalk and coupling MDM data to a multimode fiber—MDM systems with 11 modes have been demonstrated on a single polarization [65]. The direct approach to scale up the data rate per wavelength is to increase the data rate. Scaling up to 50 Gbps per wavelength is possible with NRZ modulation. Microdisk modulators have been demonstrated with 64 Gbaud signals [13] and photodiodes have been demonstrated with bandwidths above 100 GHz [14]. Interposer transmission lines with bandwidths up to 50 GHz have been demonstrated [66], [67]. Driving EICs to interface to silicon photonic MZMs have been shown to operate up to 56 Gbaud [37], [38]. Combining all these elements together presents a path for NRZ 50

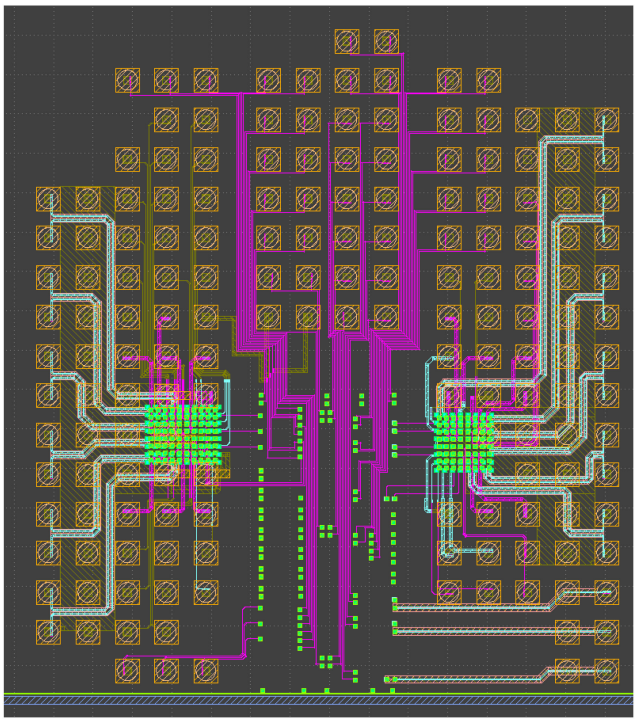


Fig. 11. The interposer layout for the future prototype utilizing custom EICs. On the left side is the EIC for the receiver portion of the PIC. On the right side is the EIC for the transmitter portion of the PIC.

Gbps. To achieve data rates of 100 Gbps and beyond, it is necessary to move to higher order modulation formats. Data rates of 200 Gbps on a single wavelength have been demonstrated with MZM-based 16-QAM in an MCM transceiver [42]. While coherent modulation schemes present the path towards the highest data rates, they often require power hungry signal processing to recover the signal. A compromise is to use PAM-4, which has been demonstrated at 128 Gbps in microring resonators [68].

Our preferred approach for scaling up the total throughput is to increase the number of wavelengths in the transceiver. The cascaded bus-microdisk modulator and demux architecture is scalable to tens of channels. Ultimately, the number of channels that can be placed on a single bus waveguide is bounded by the FSR of the resonant devices and minimum channel spacing that can be supported by the link. To further increase the scalability of the architecture, an interleaving structure can be used to route different groups of channels to different cascaded buses of modulators or demuxes [69]. Such an approach can be scaled to hundreds of wavelengths to provide a clear path towards links with multiple Tbps of data. This system provides a natural integration with CMOS-compatible comb lasers that produce hundreds of evenly spaced wavelengths [70]. Constructing massively parallel links with 10 Gbps channel rates provides an energy efficient path towards Tbps links. With over 100 channels, it is reasonable to allocate one or more wavelengths for clock forwarding, removing the need to perform clock recovery at the receiver. The 10 Gbps channel data rate avoids a heavy reliance on serialization and deserialization (SERDES), which typically dominate power consumption of optical transceivers [5]. SERDES with clock and data recovery (CDR) can consume up to 60% of a transceivers total power [71]. Using 10 Gbps channel rates also avoids the reliance on digital signal processing (DSP). In coherent commercial optical transceivers, the DSP ASIC consumes as much power as the rest of the transceiver combined [72]. Additionally, channel rates of 10 Gbps are close to the most efficient rate for resonant modulators/demuxes in silicon photonics requiring thermal stabilization [73], [74]. This architecture provides the foundation for a highly parallel system that combines energy efficiency, high throughput, and high areal bandwidth density.

In addition to repackaging the TI TIA prototype, a 2.5D integrated prototype with custom EICs fabricated in a 28 nm process is currently in the final stages of fabrication for the interposer, which can be seen in Fig. 11. The custom EICs are separated into a TX and a RX chip. Both the TX and the RX EIC are designed for two data channels and one clock channel. The TX chip contains PRBS generators at 2.5 Gbps, serialization to 10 Gbps, and an output driver to differentially drive the modulator. The RX chip contains a TIA, voltage amplifier, deserializer, and an output buffer.

VI. Conclusion

In this paper, we outline the different integration approaches pursued to interface between EICs and PICs in MCM transceivers. We provide an overview of the initial prototype of our 2.5D integrated, 4-channel WDM transceiver targeting a channel operation of 10 Gbps. The prototype utilizes silicon photonic microdisk modulators, microdisk demux filters, high speed photodiodes, and commercial single-channel bare die TIAs. A silicon interposer facilitates the connection between the PIC and the EICs, and provides connections to the PCB through TSVs. The results of our characterization of the PIC and integrated MCM transceiver are summarized—error free modulation is demonstrated at -16 dBm at 10 Gbps for the PIC and -15 dBm at 6 Gbps for the MCM. The initial prototype is currently being repackaged to correct the photodiode connections. Subsequent prototypes include a 2.5D integrated transceiver with custom EICs to provide full TX and RX driving functionality, and a 3D integrated network-on-chip utilizing an active photonic interposer. Continuing to adopt silicon photonics into datacenters requires careful co-integration of silicon photonics with the driving electronics to ensure that the optical advantages can be realized.

Acknowledgment

The authors would like to thank Yudong Zhang and Peter Kinget for development of the custom EICs that will be used in the future prototype. Fabrication of the PIC was through a MPW at AIM Photonics. The authors further thank SUNY CNSE for fabrication of the interposer and Matt Traverso at Cisco for helpful integration discussions.

REFERENCES

[1] "Cisco visual networking index Forecast and trends, 2017-2022," White Paper, 2019. [Online]. Available: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html

[2] A. S. G. Andrae and T. Edler, "On global electricity usage of communication technology: Trends to 2030," *Challenges*, vol. 6, pp. 117–157, Feb. 2015.

[3] E. J. Fluhr *et al.*, "The 12-core POWER8TM processor with 7.6 Tb/s IO bandwidth, integrated voltage regulation, and resonant clocking," *IEEE J. Solid-State Circuit*, vol. 50, no. 1, pp. 10–23, Jan. 2015.

[4] S. Rumley *et al.*, "Optical interconnects for extreme scale computing systems," *Parallel Comput.*, vol. 64, pp. 65–80, May 2017.

[5] A. Krishnamoorthy *et al.*, "From chip to cloud: Optical interconnects in engineered systems," *J. Lightw. Technol.*, vol. 35, no. 15, pp. 3103–3115, Aug. 2017.

[6] I. A. 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.

[7] Q. Cheng, M. Bahadori, M. Glick, S. Rumley, and K. Bergman, "Recent advances in optical technologies for data centers: A review," *OSA Optica*, vol. 5, no. 11, pp. 1354–1370, 2018.

[8] P. De Dobbelaere, "Silicon photonics transceivers for hyper-scale datacenters: Deployment and roadmap," 2016. [Online]. Available: http://www.phoxtrot.eu/wp-content/uploads/2017/01/ECOC-2016-PeterDe-Dobbelaere.pdf.

[9] "Elenion technologies announces availability of silicon photonic integrated modulator/receiver assembly for coherent applications," Mar. 17, 2017. [Online]. Available: https://elenion.com/news/2017/3/17/elenion-technologies-announces-availability-of-silicon-photonic-integrated-modulator-receiver-assembly-for-coherent-applications

[10] "Optoelectronics & photonics: Product selection guide," 2019. [Online]. Available: https://www.macom.com/files/live/sites/ma/files/contributed/aboutMacom/pdf/2019%20OptoPhotonics%20Brochure.pdf

[11] "100G parallel single mode data center connectivity," May 2017. [Online]. Available: https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/optical-transceiver-100g-psm4-qsfp28-brief.pdf

[12] A. Samani *et al.*, "Experimental parametric study of 128 Gb/s PAM-4 transmission system using a multi-electrode silicon photonic Mach–Zehnder modulator," *Opt. Express*, vol 25, no. 12, pp. 13252–13262, 2017.

[13] J. Sun, R. Kumar, M. Sakib, J. B. Driscoll, H. Jayatilleka, and H. Rong, "A 128 Gb/s PAM4 silicon microring modulator with integrated thermo-optic resonance tuning," *J. Lightw. Technol.*, vol. 37, no. 1, pp. 110–115, Jan. 2019.

[14] L. Vivien *et al.*, "Zero-bias 40 Gbit/s germanium waveguide photodetector on silicon," *Opt. Express*, vol. 20, no. 2, pp. 1096–1101, 2012.

[15] S. Saeedi, S. Menezo, G. Pares, and A. Emami, "A 25 Gb/s 3D-integrated CMOS/silicon-photonic receiver for low-power high-sensitivity optical communication," *J. Lightw. Technol.*, vol. 34, no. 12, pp. 2924–2933, Jun. 2016.

[16] A. V. Krishnamoorthy *et al.*, "Progress in low-power switched optical interconnects," *IEEE J. Sel. Topics Quantum Electron.*, vol. 17, no. 2, pp. 357–376, Mar. 2011.

[17] C. Sun *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, vol. 528, pp. 534–538. Dec. 2015.

[18] V. Stojanović *et al.*, "Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes," *Opt. Express*, vol. 26, no. 10, 2018, Art. no. 13106.

[19] B. Yu *et al.*, "FinFET scaling to 10 nm gate length," in *Proc. Digest. Int. Electron Devices Meeting*, 2002, pp. 251–254.

[20] J. F. 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.

[21] A. Narasimha *et al.*, "A 40-Gb/s QSFP optoelectronic transceiver in a 0.13 um CMOS silicon-on-insulator technology," in *Proc. Opt. Fiber Commun. Conf.*, 2008, pp. 1–3, Paper OMK7.

[22] N. B. Feilchenfeld *et al.*, "An integrated silicon photonics technology for o-band datacom," in *Proc. IEEE Int. Electron Devices Meeting*, 2015, pp. 25.7.1–25.7.4.

[23] L. Zimmermann *et al.*, "BiCMOS silicon photonics platform," in *Proc. Opt. Fiber Commun. Conf.*, 2015, pp. 1–3, Paper Th4E.5.

[24] A. Narasimha *et al.*, "An ultra low power CMOS photonics technology platform for H/S optoelectronic transceivers at less than $1 per Gbps," in *Proc. Opt. Fiber Commun. Conf.*, 2010, pp. 1–3, Paper OMV4.

[25] A. H. Atabaki *et al.*, "Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip," *Nature*, vol. 556, pp. 349–354, Apr. 2018.

[26] P. De Dobbelaere *et al.*, "Advanced silicon photonic technology platform leveraging a semiconductor supply chain," in *Proc. IEEE Int. Electron Devices Meeting*, 2017, pp. 34.1.1–34.1.4.

[27] H. Li *et al.*, "A 25 Gb/s 4.4 V-swing, AC-coupled ring modulator-based WDM transmitter with wavelength stabilzation in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 12, pp. 3145–3159, Dec. 2015.

[28] A. Inmann and D. Hodgins, *Implantable Sensor Systems for Medical Applications*, Cambridge, U.K.: Woodhead, 2013, p. 119.

[29] I. Ndip, A. Öz, S. Guttowski, H. Reichl, K. Lang, and H. Henke, "Modeling and minimizing the inductance of bond wire interconnects," in *Proc. IEEE Workshop Signal Power Integrity*, 2013, pp. 1–4.

[30] H. Li *et al.*, "A 112 Gb/s PAM4 silicon photonics transmitter with microring modulator and CMOS driver," in *Proc. Opt. Fiber Commun. Conf.*, 2019, pp. 1–3, Paper Th4A.4.

[31] M. Rakowski *et al.*, "Hybrid 14 nm FinFET silicon photonics technology for low-power Tb/s/mm^2 Optical I/O," in *Proc. Symp. VLSI Technol. Digest Tech. Papers*, 2018, pp. 221–222.

[32] H. D. Thacker *et al.*, "Flip-chip integrated silicon photonic bridge chips for sub-picojoule per bit optical links," in *Proc. Electron. Components Technol. Conf.*, 2010, pp. 240–246.

[33] Y. Chen, *et al.*, "A 25 Gb/s hybrid integrated silicon photonic transceiver in 28 nm CMOS and SOI," in *Proc. IEEE Intz. Solid-State Circuits Conf.*, 2015, pp. 1–3.

[34] R. Yarema, "Advances in bonding technologies," 2014. [Online]. Available: https://indico.cern.ch/event/309449/contributions/1679992/attachments/591500/814219/Bonn_talk.pdf

[35] M. Garcia-Sciveres and N. Wermes, "A review of advances in pixel detectors for experiments with high rate and radiation," *Rep. Profess Phys.*, vol. 81, no. 6, May 2018.

[36] Collier Ventures Inc., "Die Bumping," 2019. [Online]. Available: http://covinc.com/our-services/die-bumping/

[37] G. Denoyer, A. Chen, B. Park, Y. Zhou, A. Santipo, and R. Russo, "Hybrid silicon photonic circuits and transceiver for 56 Gb/s NRZ 2.2 km transmission over single mode fiber," in *Proc. Eur. Conf. Opt. Commun.*, 2014, pp. 1–3, Paper PD.2.4.

[38] G. Denoyer *et al.*, "Hybrid silicon photonic circuits and transceiver for 50 Gb/s NRZ transmission over single-mode fiber," *J. Lightw. Technol.*, vol. 33, no. 6, pp. 1247–1254, Mar. 2015.

[39] S. Straullu *et al.*, "Demonstration of a partially integrated silicon photonics ONU in a self-coherent reflective FDMA PON," *J. Lightw. Technol.*, vol. 35, no. 7, pp. 1307–1312, Apr. 2017.

[40] L. Carroll *et al.*, "Photonic packaging: Transforming silicon photonic integrated circuits into photonic devices," *Appl. Sci.*, vol. 6, no. 12, p. 426, Dec. 2016.

[41] D. Kim *et al.*, "2.5D silicon optical interposer for 400 Gbps electronic-photonic integrated circuit platform packaging," in *Proc. IEEE Electron. Packag. Technol. Conf.*, 2017, pp. 1–4.

[42] C. Doerr *et al.*, "Silicon photonics coherent transceiver in a ball-grid array package," in *Proc. Opt. Fiber Commun. Conf.*, 2017, pp. 1–3, Paper Th5D.5.

[43] J. S. Orcutt *et al.*, "Monolithic silicon photonics at 25 Gb/s," in *Proc. Opt. Fiber Commun. Conf.*, 2016, pp. 1–3, Paper Th4H.1.

[44] J. S. Orcutt, "Design of monolithic silicon photonics at 25 Gb/s," in *Proc. IEEE Compound Semicond. Integr. Circuit Symp.*, 2017, pp. 1–4.

[45] F. Y. Liu *et al.*, "10-Gbps, 5.3-mW optical transmitter and receiver circuits in 40-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 47, no. 9, pp. 2049–2067, Sep. 2012.

[46] M. Rakowski *et al.*, "A 4 × 20 Gb/s WDM ring-based hybrid CMOS silicon photonics transceiver," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2015, vol. 58, pp. 408–409.

[47] I. Shubin *et al.*, "All solid-state multi-chip multi-channel WDM photonic module," in *Proc. IEEE Electron. Components Technol. Conf.*, 2015, pp. 1293–1298.

[48] H. D. Thacker *et al.*, "An all-solid-state, WDM silicon photonic digital link for chip-to-chip communications," *Opt. Express*, vol. 23, no. 10, pp. 12808–12822. May 2015.

[49] Y. De Koninck *et al.*, "Advanced silicon photonics transceivers," in *Proc. Eur. Conf. Opt. Commun.*, 2017, pp. 1–3.

[50] T. Aoki *et al.*, "Low-crosstalk simultaneous 16-channel × 25 Gb/s operation of high-density silicon photonics optical transceiver," *J. Lightw. Technol.*, vol. 36, no. 5, pp. 1262–1267, Mar. 2018.

[51] B. Sirbu *et al.*, "3D silicon photonics interposer for Tb/s optical interconnects in data centers with double-side assembled active components and integrated optical and electrical through silicon via on SOI," in *Proc. IEEE Electron. Components Technol. Conf.*, 2019, pp. 1052–1059.
[52] A. Novack *et al.*, "A silicon photonic transceiver and hybrid tunable laser for 64 gbaud coherent communication," in *Proc. Opt. Fiber Commun. Conf.*, 2018, pp. 1–3, Paper Th4D.4.
[53] Xilinx, "Recommended design rules and strategies for BGA devices," 2016. [Online] Available: https://www.xilinx.com/support/documentation/user_guides/ug1099-bga-device-design-rules.pdf
[54] S. Bernabé *et al.*, "On-board silicon photonics-based transceivers with 1-Tb/s capacity," *IEEE Components, Packag. Manuf. Technol.*, vol. 6, no. 7, pp. 1018–1025, Jul. 2016.
[55] P. Dong *et al.*, "Simultaneous wavelength locking of microring modulator array with a single monitoring signal," *Opt. Express*, vol. 25, no. 14, pp. 16040–16046, Jul. 2017.
[56] K. Padmaraju, D. R. Logan, T. Shiraishi, J. J. Ackert, A. P. Knights, and K. Bergman, "Wavelength locking and thermally stabilizing microring resonators using dithering signals," *J. Lightw. Technol.*, vol. 32, no. 3, pp. 505–512. Feb. 2014.
[57] W. A. Zortman, A. L. Lentine, D. C. Trotter, and M. R. Watts, "Bit-error-rate monitoring for active wavelength control of resonant modulators," *IEEE Micro*, vol. 33, no. 1, pp. 42–52, Feb. 2013.
[58] E. Timurdogan, C. M. Sorace-Agaskar, J. Sun, E. S. Hosseini, A. Biberman, and M. R. Watts, "An ultralow power athermal silicon modulator," *Nature Commun.*, vol. 5, pp. 1–11, Jun. 2014.
[59] E. Timurdogan *et al.*, "APSUNY process design kit (PDKv3.0): O, C and L band silicon photonics component libraries on 300 mm wafers," in *Proc. Opt. Fiber Commun. Conf.*, 2019, pp. 1–3, Paper Tu2A.1.
[60] Z. Zhou, B. Bai, and L. Liu, "Silicon on-chip PDM and WDM technologies via plasmonics and subwavelength grating," *IEEE J. Sel. Topics Quantum Electron.*, vol. 25, no. 3, pp. 1–13, May/Jun. 2019.
[61] L. Luo *et al.*, "WDM-compatible mode-division multiplexing on a silicon chip," *Nature Commun.*, vol. 5, pp. 1–7, Jan. 2014.
[62] "GLOBALFOUNDRIES silicon photonics platform," 2018. [Online]. Available: http://soiconsortium.eu/wp-content/uploads/2018/10/GLOBALFOUNDRIES_SiPh_platform_v2.pdf
[63] N. Fahrenkopf, C. McDonough, G. Leake, Z. Su, E. Timurdogan, and D. Coolbaugh, "The AIM photonics MPW: A highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits," *IEEE J. Sel. Topics Quantum Electron.*, vol. 25, no. 5, pp. 1–6, Aug. 2019.
[64] Z. Chen *et al.*, "Use of polarization freedom beyond polarization division multiplexing to support high-speed and spectral-efficient data transmission," *Light: Sci. Appl.*, vol. 6, pp. 1–7, Aug. 2016.
[65] J. He *et al.*, "Silicon high-order mode (De)multiplexer on single polarization," *J. Lightw. Technol.*, vol. 36, no. 24, pp. 5746–5753, Dec. 2018.
[66] B. Snyder *et al.*, "Packaging and assembly challenges for 50G silicon photonics interposers," in *Proc. Opt. Fiber Commun. Conf.*, 2018, pp. 1–3, Paper Tu2A.3.
[67] K. Chang, R. Li, L. Ding, and S. Zhang, "Study of transmission line performance on through silicon interposer," in *Proc. IEEE Electron. Packag. Technol. Conf.*, 2014, pp. 284–287.
[68] J. Sun, R. Kumar, M. Sakib, J. Driscoll, H. Jayatilleka, and H. Rong, "A 128 Gb/s PAM4 silicon microring modulator with integrated thermo-optic resonance turning," *J. Lightw. Technol.*, vol. 37, no. 1, pp. 110–115, Jan. 2019.
[69] L. Luo *et al.*, "High bandwidth on-chip silicon photonic interleaver," *Opt. Express*, vol. 18, no. 22, pp. 23080–23087, Oct. 2010.
[70] J. S. Levy A. Gondarenko, M. A. Foster, A. C. Turner-Foster, A. L. Gaeta, and M. Lipson, "CMOS-compatible multiple-wavelength oscillator for on-chip optical interconnects," *Nature Photon.*, vol. 4, no. 1, pp. 37–40, Dec. 2009.
[71] Y. Audzevich, P. M. Watts, A. West, A. Mujumdar, S. W. Moore, and A. W. Moore, "Power optimized transceivers for future switched networks," *IEEE Trans. VLSI Syst.*, vol. 22, no. 10, pp. 2081–2092, Oct. 2014.
[72] F. Frey, R. Elschner, and J. Fischer, "Estimation of Trends for Coherent DSP ASIC Power Dissipation for different bitrates and transmission reaches," in *Proc. Photon. Netw. Conf.*, May 2017, pp. 1–8.
[73] R. Polster, Y. Thonnart, G. Waltener, J. Gonzalez, and E. Cassan, "Efficiency optimization of silicon photonics links in 65-nm CMOS and 28-nm FDSOI technology nodes," *IEEE Trans. VLSI Syst.*, vol. 24, no. 12, pp. 3450–3459, Dec. 2016.
[74] M. Bahadori *et al.*, "Energy-performance optimized design of silicon photonic interconnection networks for high-performance computing," in *Proc. Des., Autom., and Test Europe Conf.*, 2017, pp. 326–331.

Nathan C. Abrams received the B.S. and M.S. degrees in electrical engineering from Columbia University, New York, NY, USA, in 2014 and 2016, respectively. He is currently working toward the Ph.D. degree with Lightwave Research Laboratory, Columbia University, New York, NY, USA. His research interests relate to silicon photonic integration for optical interconnects.

Qixiang Cheng (Member, IEEE) received the B.S. degree from the Huazhong University of Science and Technology, Hubei, China, in 2010 and the Ph.D. degree from the University of Cambridge, Cambridge, U.K., in 2014, in the field of III/V integrated optical switches. He was a Research Assistant with the University of Cambridge. In March 2015, he joined Shannon Laboratory, Huawei, Beijing, China, researching future optical computing systems. He is currently a Research Scientist with the Lightwave Research Laboratory, Columbia University, New York, NY, USA. He has authored or coauthored more than 60 papers. His current research interests include design, simulation, and characterization of large-scale optical integrated devices for data center and optical computing applications.

Madeleine Glick (Senior Member, IEEE) received the Ph.D. degree in physics from Columbia University, New York, NY, USA, for research on electro-optic effects of GaAs/AlGaAs quantum wells. After receiving the degree, she joined the Department of Physics, Ecole Polytechnique Federale de Lausanne (EPFL), Lausanne, Switzerland, where she continued her research in electro-optic effects in GaAs and InP-based materials. From 1992 to 1996, she was a Research Associate with CERN, Geneva, Switzerland, as a part of the Lightwave Links for Analogue Signal Transfer Project for the Large Hadron Collider. From 2002 to 2011, she was a Principal Engineer with Intel (Intel Research Cambridge U.K., Intel Research Pittsburgh) leading research on optical interconnects for computer systems. Her research interests are in applying photonic devices and interconnects to computing systems. She is a fellow of the Institute of Physics and a Senior Member of OSA.

Moises Jezzini received the B.Sc. degree in physics and the M.Sc. degree in electronic systems from Tec de Monterrey, Monterrey, Mexico, on the propagation of special beams in conical structures, and the Ph.D. degree from Tyndall National Institute, Cork, Ireland, in 2018, in microwave design of high-speed multilayer interposers for the packaging of PICs. He worked for two years as a Test Engineer and five years as a Manufacturing Product Engineer in the optoelectronics industry at CyOptics/Avago working with 10G/40G/100G receivers and transmitters.

Padraic Morrissey received the B.Sc. degree in physics from University College Cork, Cork, Ireland, in 2009, and the Ph.D. degree in physics from the Tyndall National Institute, Cork, Ireland, in 2014, with a thesis entitled "Photonic Integrated Circuits for the Generation of Coherent Optical Signals." He is a Researcher with the Photonic-Packaging Group, Tyndall National Institute, focusing on the design of new optical and electrical packaging schemes for Si and InP-based photonic systems. Before joining the Photonic Packaging Group, he worked on PIC research and development with the Integrated Photonics Group, Tyndall.

Peter O'Brien received the master's degree in electronic engineering and the Ph.D. degree in physics from University College Cork, Cork, Ireland, and the primary degree in physics from Trinity College Dublin, Dublin, Ireland. He was a Postdoctoral Researcher with the California Institute of Technology and NASA's Jet Propulsion Laboratory in Pasadena, where he was involved in the development of submillimeter wave devices for remote sensing applications. He is currently the Head of the Photonic Packaging Group, Tyndall National Institute, Cork, and the Deputy Director of the Irish Photonic Integration Centre, Ireland. His group has grown steadily to become an internationally recognized research team, focused primarily on the development of packaging solutions for high-density integrated photonic systems. Prior to founding the Photonics Packaging Group, he founded a photonics company developing products for bioimaging and specialized machine vision applications, which he sold in 2008.

Keren Bergman (Fellow, IEEE) received the B.S. degree from Bucknell University, Lewisburg, PA, USA, in 1988, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology, Cambridge, MA, USA, in 1991 and 1994, respectively, all in electrical engineering. She is currently a Charles Bachelor Professor with Columbia University, New York, NY, USA, where she also directs the Lightwave Research Laboratory. She leads multiple research programs on optical interconnection networks for advanced computing systems, data centers, optical packet switched routers, and chip multiprocessor nanophotonic networks-on-chip. She is a Fellow of OSA.