

Silicon Photonics Chip I/O for Ultra High-Bandwidth and Energy-Efficient Die-to-Die Connectivity

Yuyang Wang^{1,*}, Songli Wang¹, Robert Parsons¹, Asher Novick^{1,3}, Vignesh Gopal¹, Kaylx Jang¹, Anthony Rizzo^{1,4}, Chia-Pin Chiu², Kaveh Hosseini², Tim Tri Hoang², Sergey Shumarayev², and Keren Bergman¹

¹Columbia University, New York, NY 10027, USA

²Intel Corporation, Santa Clara, CA 95054, USA

³Now with Xscape Photonics, New York, NY 10027, USA

⁴Now with Air Force Research Laboratory, Rome, NY 13441, USA

*yw3831@columbia.edu

Abstract—Embedded silicon photonics (SiPh) is promising to enable ultra-high bandwidth system-wide connectivity with vastly reduced energy consumption by integrating optics deeply within computing sockets. We present the design and characterization of a dense wavelength-division multiplexing (DWDM) SiPh transceiver chip, featuring a unique architecture in the multi-FSR regime and targeting a shoreline bandwidth density of 2 Tbps/mm with a sub-pJ/b energy consumption, providing a viable path toward die-to-die connectivity at scale.

1. Introduction

Traffic demands in data centers and high-performance computing systems have grown exponentially over the past decade, driven by the proliferation of data-intensive workloads in machine learning, big data analytics, and especially deep learning (DL)–based artificial intelligence (AI) applications. The recent demonstrations of huge potentials of large language models in natural language processing and content generation have further accelerated the technological progress toward data ubiquity with the adoption of increasingly larger DL models and datasets [1]. The ongoing trend has sparked tremendous effort improving the capabilities of the computational hardware, notably through aggressive parallelism and specialization [2, 3], far outpacing the progress of the underlying communication infrastructure [4]. As a result, moving massive data off and between chips has become the bottleneck of the computing performance and energy efficiency, gating the continuous scaling of such systems toward exascale [5].

Optical interconnects have been widely recognized as a viable solution to providing uniformly high-bandwidth data communication across a wide range of distances, leveraging the low-loss and multiplexing nature of light in fiber optics [6]. However, the gain from replacing electrical links with optical ones starts to plateau at shorter distances, predominantly bottlenecked by the energy consumption associated with the electrical-optical (EO) conversion at chip edges [7]. Existing solutions based on pluggable optics require electrical signals to travel along centimeter-long copper wires before reaching the EO interface. Such a form factor significantly limits the achievable data rate per interface without incurring substantial energy cost, prohibiting further upscaling of the system connectivity [8]. To this end, embedded silicon photonics (SiPh) optical interconnects have emerged as a promising avenue for bringing the optics closer to the computing sockets [9–12], effectively eliminating the distance dependency of the bandwidth density–energy efficiency product for data links across multiple network hierarchies [13]. In particular, SiPh microresonator-based modulators and filters [14, 15] demonstrate excellent wavelength selectivity with compact footprints that—when combined with recent advances in optical frequency comb sources [16–18]—provide an elegant solution to realizing dense wavelength-division multiplexing (DWDM) links [19, 20] without relying on bulky distributed feedback (DFB) laser arrays or dedicated (de-)multiplexers.

Silicon nitride (Si₃N₄) Kerr combs operating in the normal group velocity dispersion (GVD) regime are of growing interest to DWDM applications, thanks to their capability of generating hundreds of evenly-spaced low-noise frequency channels from a single continuous-wave (CW) laser source, with better conversion efficiency, power per line, and spectral flatness compared to alternative technologies [21]. The massive parallel frequency channels open up the possibility of achieving petabit-scale package escape bandwidth at sub-

pJ/b energy consumption [13], while also posing challenges in terms of how physical limitations of resonant devices can be circumvented by architectural innovations to support such scaling. Specifically, due to the periodic nature of the resonator resonances in the frequency domain, the traditional link architecture for microresonator-based DWDM—typically featuring a single bus of cascaded microresonators [22–25]—requires all communication channels to be placed within a single free spectral range (FSR) of the resonator. The maximum optical bandwidth supported by this architecture, hereafter referred to as the *single-FSR* regime, is thus upper bounded by the resonator’s FSR, which, in turn, is physically limited by the minimum resonator radius [26] and only a fraction of the spectral bandwidth of state-of-the-art frequency comb implementations [21]. To address this challenge, the authors have recently proposed a scalable link architecture in the *multi-FSR* regime [20, 27, 28] that allows for resonator FSRs smaller than the total optical bandwidth of the link while still being able to communicate over all channels with minimal crosstalk penalties, unlocking the design space for ultra-broadband Kerr comb–driven DWDM links.

In this study, we present our latest design and characterization of a SiPh microresonator-based DWDM transceiver (TRx) chip in the multi-FSR regime, featuring 64 channels per fiber output at 100 GHz channel spacing and targeting a chip shoreline bandwidth density of over 2 Tbps/mm with a sub-pJ/b energy consumption. In Section 2, we briefly review the link architecture and the photonic integrated circuit (PIC) design with considerations for electronic co-packaging. In Section 3, we present the design and characterization of key PIC enabling devices and components, including even-odd (de)-interleavers with automatic tuning, ultra-efficient vertical-junction microdisk modulators, and add-drop filter arrays. In Section 4, we discuss the energy implications of fabrication process variations and foundry-compatible substrate undercuts based on wafer-scale measurement data. In Section 5, we conclude the paper with a summary of the work and future outlook.

2. Link Architecture and Chip Design

Traditional microresonator-based DWDM link architectures in the single-FSR regime face fundamental challenges when the total optical bandwidth of the comb source scales beyond the resonator FSR. Specifically, since a microresonator exhibits periodic resonances at integer multiples of its FSR, there exist unwanted resonances other than the design target, a.k.a., *resonance aliases*, that could potentially overlap with non-target comb lines and result in severe crosstalk. Meanwhile, packing more channels into a single FSR risks introducing excessive inter-modulation crosstalk due to reduced channel spacing, with past experiments showing < 100 GHz channel spacing impractical [29]. While it is possible to have microresonators with larger FSRs by adopting smaller radii [30–32], it is an unsustainable solution to massive wavelength scaling due to physical limitations on the minimum microresonator size and accompanying design/manufacturing complexities regarding integrated heaters, radio frequency (RF) contacts, as well as electrical parasitic and process variation containment. Another approach, a.k.a., *band-interleaving*, tackles the problem by subdividing the incoming comb channels into multiple narrower bands, each fitting into a resonator FSR [33]. However, this approach incurs stringent requirements on the performance of the (de)-interleaving devices, such as sharp pass-band roll-off, broadband crosstalk suppression, as well as decent tunability, and thus relies on the maturity of design and fabrication of novel devices such as on-chip dichroic filters [34, 35] and contra-directional couplers [36, 37]. In this study, we base our transceiver PIC design on the multi-FSR regime enabled by an *even-odd interleaving* scheme, as previously described in [27]. This section briefly reviews the link architecture and the PIC design.

2.1. Scalable Link Architecture

As described in [27] and shown in Fig. 1a, the TRx link is designed to be driven by a Si₃N₄ dual-ring normal GVD Kerr comb that generates up to hundreds of low-noise frequency channels with 100 GHz channel spacing centered at 1505 nm (S-Band). At the transmitter (Tx) side, the incoming comb lines are subdivided by two stages of de-interleavers—the FSRs of which designed to be 200 GHz and 400 GHz, respectively—onto four buses before traversing

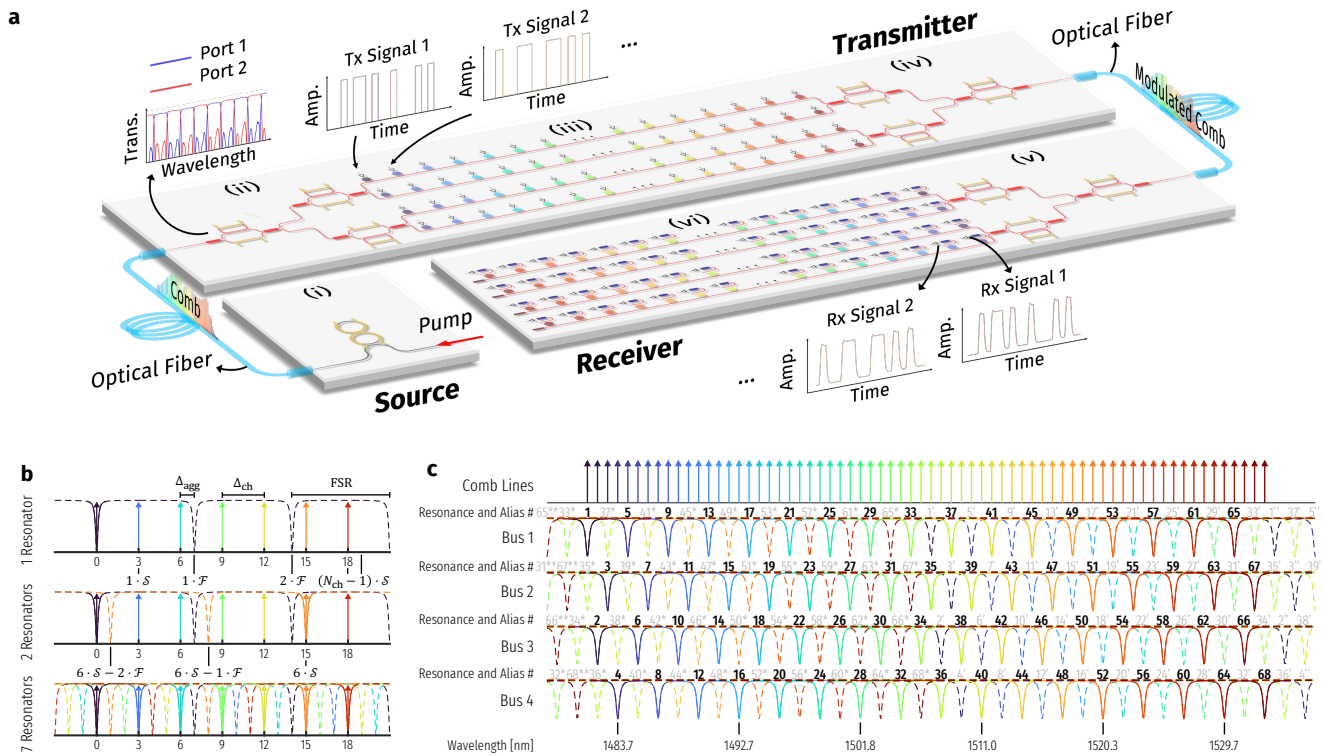


Figure 1. (a) Scalable link architecture based on even-odd interleaving. The end-to-end link consists of (i) a normal GVD Kerr comb, (ii) Tx de-interleavers, (iii) microresonator modulators, (iv) Tx interleavers, (v) Rx de-interleavers, and (vi) microresonator filters and photodetectors. (b) An exemplar multi-FSR arrangement for $(S, \mathcal{F}) = (3, 7)$ [Eqs. (1a) and (1b)] showing the placement of resonance aliases (dashed) between communication channels (solid) ensured by the co-prime rule [13]. (c) The multi-FSR arrangement adopted for the transceiver link in this study, with $(S, \mathcal{F}) = (2, 17)$, allowing up to 17 channels per bus achievable by microresonators with a moderate 25.69 nm FSR. Asterisks (primes) denote the aliases to the blue (red) end of the nominal resonance.

separate banks of cascaded microresonator modulators. Each stage of de-interleaver splits the incoming wavelengths into “even” and “odd” groups, effectively doubling the channel spacing while nearly preserving the total optical bandwidth seen on each output port. Each resonant modulator can modulate a distinct wavelength channel while nearly appearing transparent to other channels on the bus. The modulated signals from the four buses are recombined by two stages of even-odd interleavers into a single fiber output. Symmetrically, at the receiver (Rx) side, the modulated signals are de-interleaved and sent to respective banks of cascaded microresonator filters that drop each channel onto a photodetector (PD) for sensing.

Since the comb lines on each bus span an optical bandwidth of over 50 nm, it is challenging, if not infeasible, to design microresonator modulators with integrated electronic contacts for an even larger FSR due to a prohibitively small physical dimension that raises manufacturability concerns [38]. As a result, the resonator FSR must be carefully chosen to prevent resonance aliases from overlapping with non-target comb lines in the multi-FSR regime. We have mathematically formulated the multi-FSR design principles in [13], from which a valid multi-FSR channel arrangement can be elegantly derived by picking a pair of co-prime integers (S, \mathcal{F}) satisfying

$$\begin{cases} S = \frac{\Delta_{ch}}{\Delta_{agg}}, & (1a) \\ \mathcal{F} = \frac{FSR}{\Delta_{agg}}, & (1b) \end{cases}$$

where Δ_{ch} is the effective channel spacing on each bus after de-interleaving, Δ_{agg} is the reduced spacing between a channel and its nearest aggressor alias, and FSR is the resonator free spectral range. Fig. 1b visualizes how the co-prime rule ensures the resonance aliases to fall between the communication channels with an example for $(S, \mathcal{F}) = (3, 7)$, meaning that up to 7 resonators can be arranged along a bus with resonance aliases placed at every one third of the channel spacing.

For the TRx architecture in Fig. 1a, which features a total of 64 channels

arranged into four buses, we choose $(S, \mathcal{F}) = (2, 17)$, as illustrated in Fig. 1c, allowing up to 17 channels per bus achievable by microresonators with a moderate 25.69 nm FSR. The optimal selection of S , \mathcal{F} , and FSR involves various trade-offs between the aggregated data rate, crosstalk penalties, and link yield, for which a more holistic design space exploration is provided in [28]. With 16 microresonator modulators/filters placed 400 GHz apart on each bus, the transceiver can effectively perform 64-channel DWDM targeting a 1.024 Tbps/fiber bandwidth capacity with a moderate data rate of 16 Gbps/channel.

2.2. Photonic Integrated Circuit Design and Fabrication

We designed a transceiver PIC featuring the abovementioned scalable link architecture in a fabrication process co-developed with AIM Photonics, which includes custom vertical-junction implants with optimized doping conditions and wafer-scale substrate undercuts for improved thermal tuning efficiency [39, 40]. Fig. 2 shows the microscope images of the fabricated PIC. Each PIC consists of four groups of Tx/Rx arrays, each targeting a 4.096 Tbps/group data rate. Each group of Tx/Rx arrays consists of four 1.024 Tbps links with the architecture described in Sec. 2.1 and Fig. 1a. For electronic integrated circuit (EIC) co-packageability, the PIC signal pads are placed at a 55 μm pitch and routed with considerations for minimizing the parasitic resistance and capacitance of the RF traces. Fig. 2b shows the pad arrays of the Tx cells as an example. The PDs after the resonator drop-ports provide electrical feedback for locking the resonance wavelengths to their target comb lines. Notable implementations of feedback control logics have been reported leveraging the drop-port PD signal [41–45]. The PIC pads will be μ -bumped after fabrication for flip-chip bonding with the EIC driver chip designed accordingly, an approach demonstrated feasible in past works [12, 46]. Wafer-scale substrate undercuts are placed near the resonant modulators/filters as well as the (de-)interleavers (Fig. 2c) for improving the thermal tuning efficiency of and reducing the thermal crosstalk between the thermal phase shifters [39, 40, 47]. The edge coupler array is placed at a 127 μm pitch

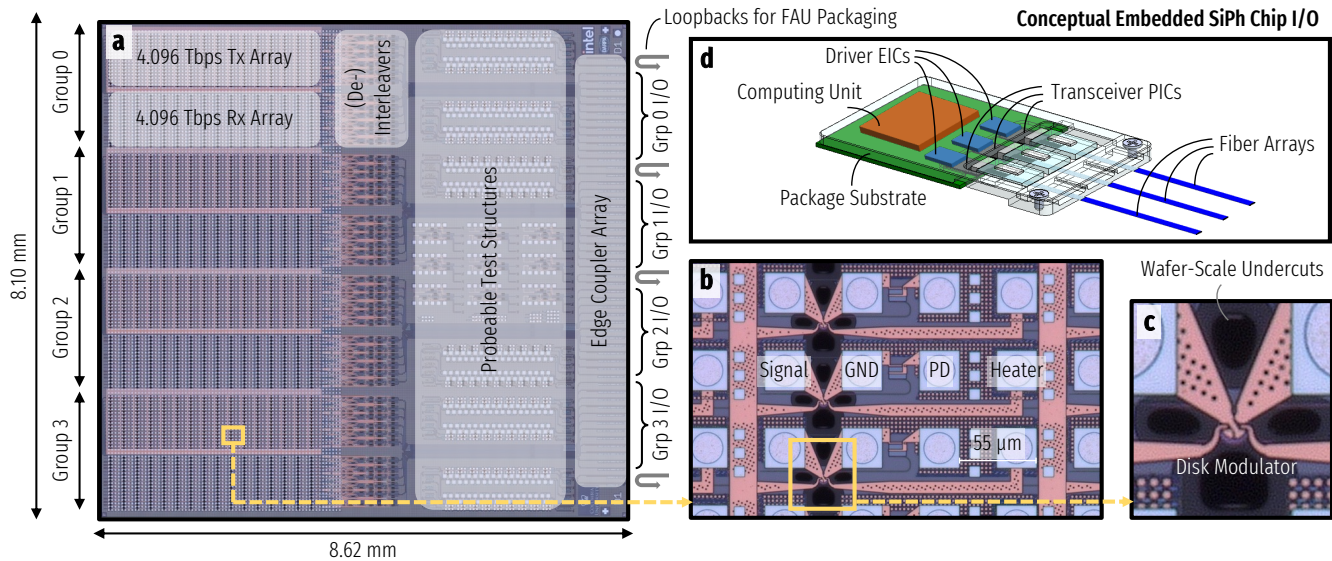


Figure 2. Microscope images of the designed and fabricated transceiver PIC, where each (a) 8.10 mm×8.62 mm chip consists of four groups of Tx/Rx arrays targeting 4.096 Tbps/group, probable test structures of link subcomponents for pre-packaging characterization of process variations and screening of KGDs, and custom suspended edge couplers [12] providing low-loss optical I/O. Each 4.096 Tbps transceiver group consists of four 1.024 Tbps links with the architecture of Fig. 1a. The Tx/Rx arrays feature (b) electrical pads at a 55 μm pitch for EIC flip-chip bonding and (c) wafer-scale substrate undercuts for improving the thermal tuning efficiency of and reducing the thermal crosstalk between the thermal phase shifters. (d) A conceptual chip I/O configuration shows multiple transceiver PICs embedded onto a single platform, next to the computing socket.

along a single side of the PIC for optical input/output (I/O), including four pairs of loopbacks (Fig. 2a) for fiber array unit (FAU) alignment. Custom suspended edge couplers are adopted, which have shown to improve the coupling loss down to -1.1 dB/facet [12], compared to the -3 dB/facet of the standard edge couplers from the process design kit. Between the edge couplers and the (de-)interleavers is a keep-out zone (KOZ) reserved for packaging. However, grating-coupled probable test structures of the link subcomponents are placed in the KOZ (Fig. 2a) and can be accessed before packaging for identifying known good dies (KGDs) and providing qualitative insights on link performance and yield. Finally, a 16.384 Tbps aggregated bandwidth is able to escape from a single 8.10 mm side of the PIC, resulting in a shoreline bandwidth density of 2.023 Tbps/mm at 1.024 Tbps/fiber. In Fig. 2d, a conceptual configuration for embedded SiPh chip I/O is shown, where multiple transceiver PICs can be co-packaged with the computing unit on a single platform, bringing high-bandwidth optical connectivity closer to the computing socket.

3. Component Design and Validation

The design of the transceiver PIC involves the tight integration of multiple constituent components, notably even-odd (de-)interleavers and microresonator-based modulators and filters. In this section, we present component-level design and validation of the key building blocks of the transceiver PIC.

3.1. Compact Even-Odd (De-)Interleavers

As described in Sec. 2.1 and illustrated in Fig. 1a, even-odd (de-)interleavers are required to expand the effective channel spacing to accommodate the resonance aliases present in the multi-FSR channel arrangement of the cascaded resonator arrays. While basic Mach-Zehnder interferometer (MZI)-based interleavers are compact and relatively straightforward to design, they are prone to fabrication/environmental perturbations and have a limited channel capacity due to the group velocity dispersion of silicon-on-insulator waveguides. In our link design, we adopt a modification of the MZI, known as a ring-assisted MZI (RAMZI), for the required even-odd (de-)interleaving operation. RAMZIs incorporate ring resonators to achieve flat-top pass-bands [48, 49], making it more resilient to both perturbations and the FSR mismatch w.r.t. the comb source, while having a more compact footprint than alternative structures, such as cascaded MZIs [50, 51], that achieve a similar flat-top response.

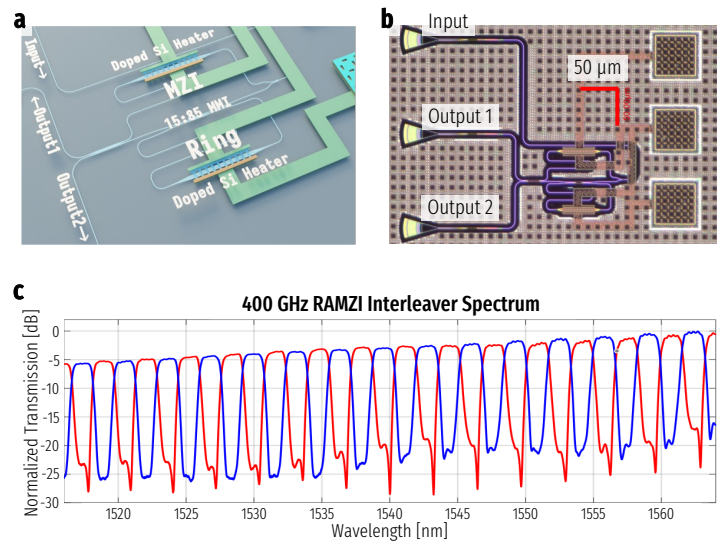


Figure 3. (a) Design rendering of a RAMZI with a 400 GHz FSR. (b) Microscope image of the fabricated device. (c) Measured transmission spectra of the interleaver. The red and blue curves correspond to the two output ports.

3.1.1. Device Design and Characterization

The RAMZIs are designed with an FSR twice of the channel spacing of the incoming signals to achieve even-odd (de-)interleaving. In the case of this study, the first stage and the second stage (de-)interleavers (Fig. 1a) are designed for 200 GHz and 400 GHz FSRs, respectively. To achieve a broadband flat-top response, a compact multi-mode interferometer-based coupler with 15:85 splitting ratio is implemented for effective coupling into the ring. The effective path length of the ring needs to be approximately twice the MZI arm length difference. Fig. 3a shows a schematic rendering of a 400 GHz RAMZI as an example, and Fig. 3b is the microscope image of the fabricated device. The measured transmission spectra, as shown in Fig. 3c, exhibit a flat-top response after applying 1 V to the phase shifter on the MZI arm, with an extinction ratio of > 20 dB over 50 nm bandwidth for both output ports. It should be noted that the unevenness in the envelope of the spectrum

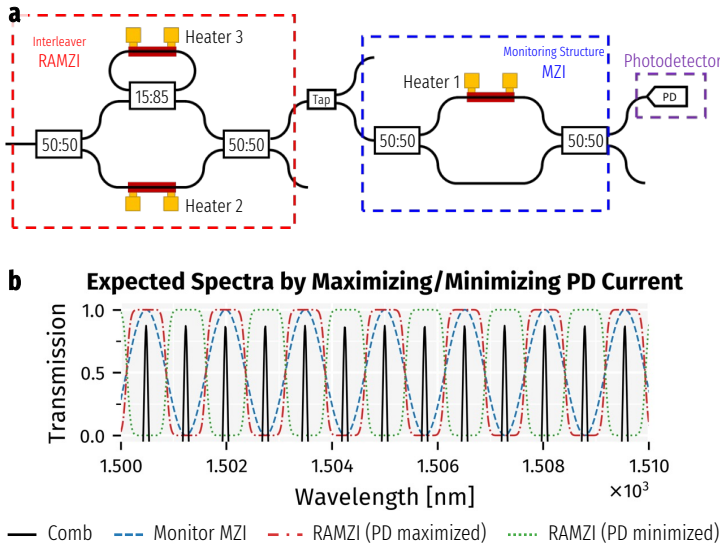


Figure 4. (a) Schematic of the RAMZI interleaver with an auxiliary monitoring structure. (b) Simulated RAMZI and monitor MZI spectra when the PD reading is maximized/minimized.

is primarily introduced by the grating coupler used in the test structure, rather than the interleaver itself. The custom edge couplers used in the transceiver PIC design are expected to have a uniformly low insertion loss across the optical band of interest [12].

3.1.2. Automatic Tuning

The (de-)interleavers in the transceiver PIC are also equipped with a monitoring PD for implementing automatic compensation of phase errors—vital to achieving the desired flat-top response—as well as automatic alignment of the pass-/stop-bands with the DWDM channels. Despite that monitored auto-tuning architectures have been demonstrated for various types of optical filtering devices [52–54], directly tapping either one of the RAMZI output ports for photocurrent monitoring would result in a constant readout, as explained in [51], regardless of the voltages applied on the thermal phase shifters. To address this challenge, we introduce an auxiliary monitoring structure composed of an MZI with an identical FSR to that of the RAMZI, followed by a PD, as illustrated in Fig. 4a. The resulting photocurrent as a function of the thermal tuning voltages is no longer a constant. We have verified through equation-based analyses as well as simulations that the PD current reading will reach its maximum when the pass-bands of both the RAMZI and the monitor MZI are aligned with the DWDM channels, and will reach its minimum when their pass-bands offset with one another by a single channel spacing, as illustrated in Fig. 4b.

We experimentally verified the automatic tuning of the interleavers in a test setup outlined in Fig. 5a. The setup consists of eight DFB lasers with a 200 GHz channel spacing, acting as the DWDM source. The DWDM channels are directed into the interleaver through an angled fiber array and grating couplers, combined with a broadband light source. The broadband source is used for visualizing the interleaver spectrum shape under the optical spectrum analyzer (OSA), and its power is controlled to be much smaller than the DWDM source so that it has negligible impact on the auto-tuning process. The output light from the interleaver is examined by a Yokogawa OSA, allowing for real-time spectrum inspection during the optimization process. It should be noted that the broadband source and the OSA serve only for visualization purposes and do not assist the auto-tuning process. For interleaver thermal control, a direct current (DC) probe is landed on the six pads, connecting the interleaver structure to the DC power supplies. A Keysight multi-channel power supply delivers voltages to the three heaters, acting as the optimization inputs, while a Keithley high-precision power supply provides bias voltage to the PD, reading the generated photocurrent as the optimization target. The adjustment to the heater power at each step is performed in a binary search manner as

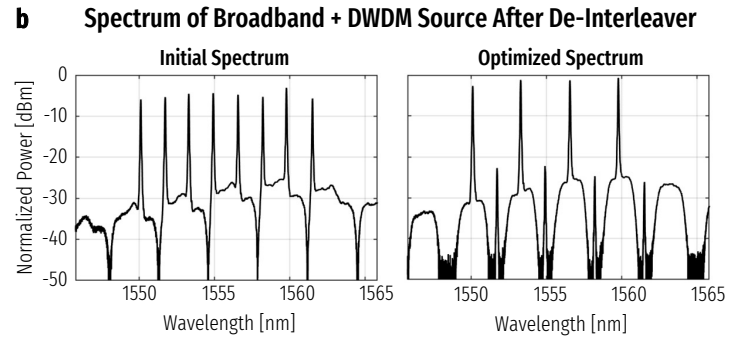
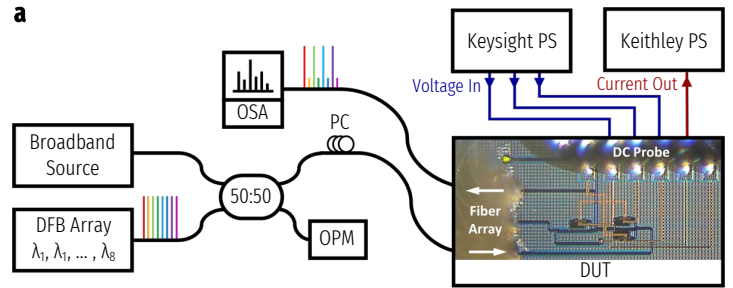


Figure 5. (a) Experimental setup for the interleaver auto-tuning process, where PC: polarization controller, OPM: optical power meter, PS: power supply, and DUT: device under test. (b) RAMZI interleaver spectra with DWDM source before and after the auto-tuning, showing an optimized extinction ratio of over 20 dB and the precise alignment of the pass-bands to the DWDM channels after tuning [55].

described in [51]. Fig. 5b shows the interleaver transmission spectra before and after the automatic tuning process, demonstrating an optimized extinction ratio of over 20 dB and the precise alignment of the pass-bands to the DWDM channels after tuning. More details on the interleaver device design, tuning algorithm, and evaluations are provided in [55].

3.2. Resonant Modulator Arrays

In order to validate the transmitter functionalities before EIC integration, we designed and fabricated a test link that has the same modulator and edge coupler designs as the PIC. As shown in Fig. 6a, the test link follows the multi-FSR configuration of Fig. 1c, consisting of 68 Tx modulators arranged into 4 buses. Each modulator is designed to match a distinct wavelength of a 100 GHz comb, which is being co-developed with the PIC by our collaborators. Each bus of cascaded modulators is edge coupled from both ends, without traversing (de-)interleaver stages which are independently verified as described above. Each modulator has two RF pads for modulation and two DC pads for thermal tuning of resonance wavelengths. The RF and DC pads are placed two rows apart, as shown in Fig. 6b, to allow for simultaneous landing of the RF and DC probes. Fig. 6c plots the transmission spectrum of a Tx bus as an example, confirming that the resonances aliases fall between the nominal resonances, satisfying multi-FSR requirements. The channel arrangement closely matches the architectural design shown in Fig. 1c. In the following, we present the design and characterization of the custom microresonator modulators.

3.2.1. Custom Vertical-Junction Microdisk Modulators

Integrated resonant modulators allow for wavelength-selective data encoding in an efficient and compact form factor. Microdisk modulators with vertical-junction depletion-mode modulation are particularly attractive for energy-efficient DWDM implementations, as they can be driven by modest CMOS-compatible voltages [26]. The microdisk's whispering gallery mode requires that all electrical contacts be placed within the external radius of the modulator, as there must be a clean full etch at this boundary to ensure that radiative

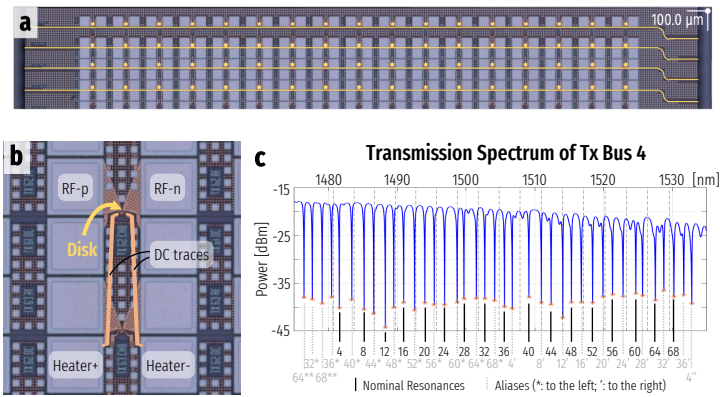


Figure 6. (a) Microscope image of the test link for microresonator modulator arrays, where (b) the RF and DC pads are placed two rows apart to allow for simultaneous landing of the RF and DC probes. (c) Exemplary transmission spectrum from the 4th Tx bus, matching and validating the multi-FSR architectural design of Fig. 1c.




Figure 7. Illustrations of modulator RF contact placement schemes: (a) the traditional left-right contacts as seen in [13]; (b) the quadrant contacts featured in this transceiver PIC design; and (c) the experimental interleaved contacts demonstrated in [56] for incorporation into future chip design. The layer thicknesses are not to scale in order to show the contacts and vias.

bend loss does not degrade the resonator's quality factor (Q) [38]. As the modulator radius decreases, the series resistance—key determinant of the modulator's RC bandwidth—increases, because the foundry design rules regarding minimum metal spacing limit the number of parallel RF contacts. This limit is exacerbated with the inclusion of an integrated heater, which requires room for two additional contacts within an already constrained area [56].

To address this challenge, we developed new contact placement schemes that distribute the RF contacts more evenly across the junctions to reduce the series resistance and improve the modulation bandwidth and efficiency. Compared to the traditional *left-right* scheme (Fig. 7a), where the RF contacts to the P and N junctions are placed onto two distinct halves of the microdisk [13, 31, 57], the transceiver PIC in this study adopts a *quadrant* scheme, as illustrated in Fig. 7b and shown in Fig. 2c, where we divide the doping region into four quadrants and arrange the P/N contacts in a criss-cross pattern. We also took one step further by introducing a new *interleaved* scheme, as illustrated in Fig. 7c and detailed in [56], where the contacts alternate between P and N, analogous to the arrangement of the interleaved lateral-junction ring modulators [58].

Quadrant Contact Modulator Characterization The quadrant contact modulators, featured in the transceiver PIC, are characterized in the test link described in Fig. 6. From the transmission spectra of the modulator buses (an example given in Fig. 6c), we extract the average FSR of the microdisks to be 25.64 nm, closely matching the design target of 25.69 nm. We extract the resonator Q, full width at half maximum (FWHM), and extinction ratio (ER) by fitting a Lorentzian function to each nominal resonance, as seen in Fig. 8a. The average Q, FWHM, and ER are found to be ~ 7000 , 0.22 nm, and 20 dB, respectively. The integrated heater exhibits a 250Ω resistance (Fig. 8b left), reasonably close to the design target of 200Ω and expected to be further

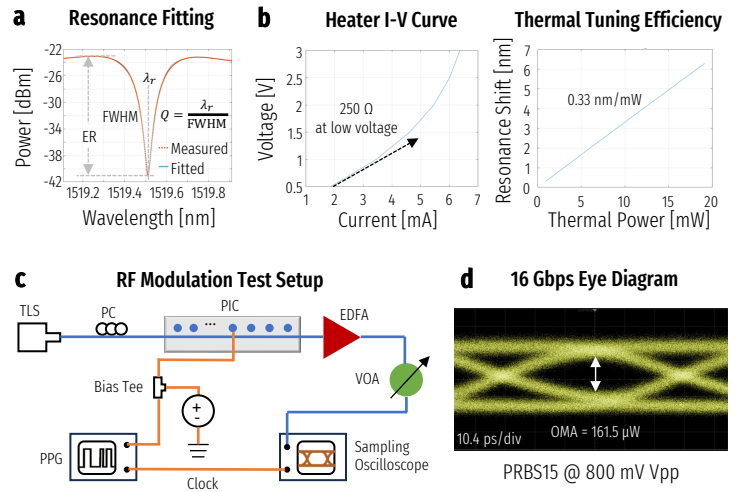


Figure 8. Characterization of the quadrant contact modulator: (a) average Q, FWHM, and ER extracted as ~ 7000 , 0.22 nm, and 20 dB, respectively; (b) thermal tuning efficiency measured as 0.33 nm/mW, achieved by a 250Ω integrated heater; (c) RF modulation test setup; and (d) open eye diagram at 16 Gbps with PRBS15 driven by 800 mV peak-to-peak voltage.




Figure 9. (a) Microscope image of a fabricated interleaved contact microdisk modulator. (b) Open eye diagram at 32 Gbps with PRBS15 driven by 800 mV peak-to-peak voltage. (c) Measured S11 frequency response after de-embedding the pads and connected metal traces, showing reduced back reflections due to better impedance matching with greater number of parallel interleaved contacts [56].

improved in future design. By applying a series of voltages to the heater and observing the shift in the resonance wavelength, a 0.33 nm/mW tuning efficiency is characterized for the modulators (Fig. 8b right).

As shown in Fig. 8c, we characterize the RF performance of the quadrant contact modulators with a non-return-to-zero (NRZ) pseudo random bit sequence (PRBS) that is $(2^{15} - 1)$ -bit long, generated by an Anritsu MP1900A pulse pattern generator (PPG). A HyperLabs HL9447 bias tee was used for combining a DC reverse bias of -1.2 V with the RF signal of an 800 mV peak-to-peak voltage. A Keysight 81606A tunable laser source (TLS) sends light at 1478 nm into a bus of 17 cascaded modulators. The light, modulated by the modulator at corresponding resonance, goes through an erbium doped fiber amplifier (EDFA) before being measured by a Keysight N1092C sampling oscilloscope, with a Thorlabs EVOA1550A variable optical attenuator (VOA) placed before the optical input. The oscilloscope receives a clock trigger directly from the PPG and is set to pattern lock on the 16 Gbps NRZ PRBS15 pattern. Fig. 8d shows a completely open eye diagram captured by the oscilloscope with $\text{OMA} \approx 161.5 \mu\text{W}$, demonstrating high promise for achieving the target 1 Tbps/fiber data rate at ultra-low energy.

Interleaved Contact Modulator Characterization The interleaved contact microdisk modulators are fabricated and characterized in the form of stand-alone test structures (Fig. 9a). As detailed in [56], the optimized design achieves $\text{ER} = 36.5$ dB, $Q \approx 6000$, and $\text{FSR} = 27.15$ nm, with higher-order modes suppressed across the full C-band. With contact resistance further reduced through greater parallelism, we verified the improved RC bandwidth with an open eye diagram obtained at 32 Gbps, driven by PRBS15 signals with as low as 800 mV Vpp. The results are cross-validated with S11

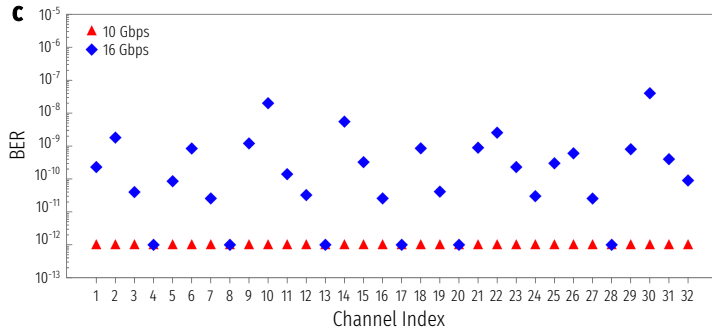



Figure 10. (a) Die photo of the 32-channel wire-bonded and optically packaged receiver chip featuring a two-stage de-interleaved multi-FSR architecture. (b) Optically packaged and DC wire-bonded receiver chip. (c) BER measurements for all 32 channels on the chip at both 10 Gbps and 16 Gbps [20, 59].




Figure 11. (a) Custom microdisk filters featuring conformal couplers and an integrated doped silicon heater have demonstrated exceptional suppression of higher order modes, while matching the modulator's resonance and dispersion behavior. (b) A representative filter spectrum showing key optical parameters for extractions. (c) Thermal tuning efficiency of the microdisk filters extracted as 0.73 nm/mW.

measurements, as shown in Fig. 9c, where the interleaved contact design results in significantly reduced back reflections compared to the quadrant design. This novel implementation of vertical-junction microdisk modulators thus provides a viable path toward achieving 100 Tbps escape bandwidth in a package similar to what conceived in Fig. 2d, by providing a double-folded data rate per channel at the same driving voltage.

3.3. Resonant Filter Arrays

Microresonator-based filter arrays perform channel-selective filtering of the modulated signal onto individual PDs for converting it back to the electrical domain. In light of the multi-FSR regime, the same de-interleaving structure can be employed to subdivide the channels onto separate banks of resonant filters for detection. In [20, 59], we have built a proof-of-concept 32-channel receiver chip, featuring two stages of de-interleavers and 4 buses of eight-channel add/drop ring filter arrays, optically packaged and DC wire-bonded. The receiver chip achieves natively error-free operation for all channels at 10 Gbps and for out of 32 channels at 16 Gbps, with anticipated improvement in receiver sensitivity by closely integrating with a transimpedance amplifier.

In addition to validating the symmetric link architecture, we also pushed along the device customization frontier by exploring custom microdisk filters. Addressing one of the design emphases of the link architecture in this study, namely the energy efficiency, microdisk filters are desirable due to




Figure 12. (a) 1σ resonance wavelength variation as function of disk radius. 448 resonator devices in total were measured across 64 reticles. Wafer-scale average thermal tuning efficiency is also shown in units of mW/nm to emphasize the inverse relationship with 1σ variation. (b) Intra-reticle resonance variation of custom microdisk modulators across wafer. Row and column indicated reticle location on wafer.

their inherent robustness to process variations [60], as well as an increased FSR upper bound that allows for a greater link design space. As illustrated in Fig. 11a, our custom microdisk filters are designed with conformal couplers and an integrated doped silicon heater that demonstrably suppress higher order modes [27]. We characterized the optical specifications of these custom microdisk filters following the same procedure as the quadrant modulators. Key optical specifications were extracted, including $Q \approx 11\,000$, $\text{FWHM} = 0.14\text{ nm}$, and $\text{ER} = 23\text{ dB}$. These measured optical specs are in line with the anticipated design values, providing the optical bandwidth required by a dedicated EIC to detect and recover the RF signal dropped by the filter. We also extract the thermal tuning efficiency of the custom filters to be 0.73 nm/mW. The custom filters have a greater thermal tuning efficiency than the modulators mainly due to the lack of an oxide barrier between the heater and the RF circuitry, which is required by the modulators for electrical isolation. In addition, the integrated heater of the filters features a radial shape that is closer to the optical mode, and less metal is present in the vicinity of the heater that would otherwise conduct heat away from the resonator device.

4. Process Variation and Energy Estimation

Critical to limiting energy consumption, our architecture benefits significantly from process-tolerant devices and integration of wafer-scale substrate undercut. In this section, we present measured wafer-scale process variations of our constituent devices, thermal tuning efficiency improvement from substrate undercut, and the associated implications on energy efficiency.

4.1. Wafer-Scale Process Variations

Custom add-drop microdisk filters over a range of radii from 1.6 μm to 4.5 μm were fabricated on a dedicated 300 mm wafer run through AIM Photonics, offering us the ability to collect comprehensive reticle-to-reticle statistics. Fig. 12a shows the effect of variation in fabrication on resonances of a particular microdisk measured across 64 reticles. Microdisks do not have an inner waveguide wall, giving them one less dimension affected by variations in fabrication than microring filters. Decreased variation in fabrication, which is inversely proportional to radius, leads to reduced need for thermal tuning to offset fabrication variation-induced resonance shift. Wafer-scale measurements indicate a linear relationship between radius and mW/nm tuning efficiency, and a reciprocal relationship between radius and 1σ resonant wavelength variation due to variations in fabrication [61]. We also demonstrate a wafer map of intra-reticle variations in resonant wavelength of our custom modulators, shown in Fig. 12b, indicating the resonant devices are robust to process variations at both the wafer-scale and locally.

4.2. Wafer-Scale Substrate Undercut

Substrate undercut is a promising technique to reduce the thermal tuning energy required for rectifying the fabrication variations [40]. We have worked with our foundry partner, AIM Photonics, to co-develop a wafer-scale substrate undercut process. Due to the experimental nature of the process, and thus the stringent design rules keeping the undercut area away from the active

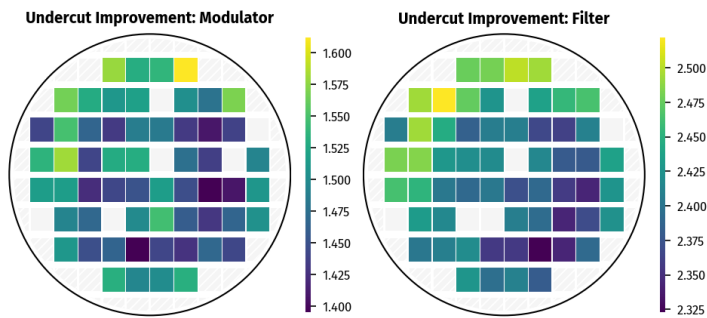


Figure 13. Measured thermal tuning efficiency improvement across wafer from wafer-scale substrate undercut. Up to $\sim 1.6\times$ and $\sim 2.5\times$ improvements are observed for disk modulators and filters, respectively.

devices by certain distances, the effectiveness of the process is expected to vary across the wafer and also depend on device layout. In Fig. 13, we show the measured improvement in thermal tuning efficiencies across the entire 300 mm wafer for both disk modulators and filters. The improvement is quantified by the ratio of the thermal tuning efficiency with and without substrate undercut. For disk modulators, the recorded improvement is up to $\sim 1.6\times$, while for disk filters, the improvement can be extended to $\sim 2.5\times$. This is likely due to the fact that the disk filters have less metal traces in the vicinity of the heaters, allowing for larger undercut area, and meanwhile, having fewer heat sinks that can potentially harm the thermal efficiency. In order to further improve the undercut effectiveness, future designs will look into optimizing metal contact and trace layouts in relation to the undercut area.

4.3. Energy Efficiency Implications

With latest knowledge of wafer-scale process variation magnitudes and effectiveness of substrate thermal undercuts, we estimate the thermal energy consumption for control of most constituent devices within the link. Integrated micro-heaters are used for either of the dual rings of the Kerr comb for initialization and stability. Previously fabricated integrated Kerr combs indicate a thermal energy consumption of 49–100 fJ/b. We note that each 1.024 Tbps link consists of 9 (de-)interleavers, where each (de-)interleaver contains 3 thermo-optic phase shifters for the RAMZI arm, the assistive ring, and the monitoring MZI arm, respectively. These integrated heaters, as shown in Fig. 4, are necessary for interleaver alignment and flat-top response. The disk modulators and filters also utilize integrated heaters, as shown in Figs. 7 and 11, respectively, to achieve high-efficiency thermal tuning. Based on the latest characterization of wafer-scale process variations and tuning efficiencies, and assuming a 3σ design offset below the target values for the resonant devices, we derive Table 1 that summarizes the anticipated thermal energy consumption of the transceiver PIC. The anticipated improvement in tuning efficiency with substrate undercut for the Kerr comb and interleavers are based on the data point from the microdisk filters, due to its geometry design being the most friendly to undercut placement. We will continue to work toward improving the effectiveness of substrate undercuts, and thus the energy efficiency of the link, through a co-design of device and undercut geometries.

5. Conclusion and Future Work

In conclusion, we have presented the design and characterization of a groundbreaking Kerr comb-driven SiPh DWDM transceiver chip, which offers a scalable architecture for ultra-high bandwidth and energy-efficient die-to-die connectivity. Our results demonstrate the feasibility of achieving a chip shore-line bandwidth density of over 2 Tbps with a sub-pJ/b energy consumption, which is significant in the context of exponentially growing traffic demands in data centers and high-performance computing systems. The multi-FSR regime adopted in this work enables the utilization of microresonators with moderate FSR for ultra-broadband DWDM with minimal crosstalk penalties. The preliminary characterization of crucial components such as the even-odd (de-)interleavers and microresonator modulators and filters has laid the

Table 1. Expected thermal energy consumption of both constituent components and total link, obtained by dividing the expected power dissipation by the link data rate. The expected energy efficiency for devices with substrate undercuts is also listed based on latest characterization results.

Component	w/o Undercuts	w/ Undercuts
Comb Source [fJ/b]	49–100	19.6–40
(De-)Interleavers [fJ/b]	70	28
Disk Modulators [fJ/b]	329	206
Disk Filters [fJ/b]	149	59.6
Total Thermal Energy [fJ/b]	597–648	313.2–333.6

groundwork for subsequent system-wide validations.

For future work, we aim to further enhance the modulator and filter designs to substantially boost the chip escape bandwidth. Link-level demonstrations will be carried out to validate the transceiver's performance systematically. With a focus on mitigating the impact of process variations and maximizing energy efficiency, future designs will also look into improving the effectiveness of wafer-scale thermal undercuts and adopting fabrication-robust device designs. The progress made in this study is a key step towards realizing the vision of seamlessly integrated photonics and electronics in future computing systems.

Acknowledgments

This work was supported in part by the U.S. Defense Advanced Research Projects Agency (DARPA) under Common Heterogeneous Integration and IP Reuse Strategies (CHIPS) program contract number HR00111830002, in part by DARPA under Photonics in the Package for Extreme Scalability (PIPES) program contract number HR00111920014, and in part by the Center for Ubiquitous Connectivity (CUbic), sponsored by the Semiconductor Research Corporation (SRC) and DARPA under the JUMP 2.0 program.

References

- [1] D. Narayanan *et al.*, "Efficient large-scale language model training on GPU clusters using megatron-LM," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, ser. SC '21. New York, NY, USA: Association for Computing Machinery, Nov. 2021, pp. 1–15.
- [2] W. J. Dally *et al.*, "Evolution of the Graphics Processing Unit (GPU)," *IEEE Micro*, vol. 41, no. 6, pp. 42–51, Nov. 2021, conference Name: IEEE Micro.
- [3] N. Jouppi *et al.*, "TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, ser. ISCA '23. New York, NY, USA: Association for Computing Machinery, Jun. 2023, pp. 1–14.
- [4] S. Pati *et al.*, "Computation vs. Communication Scaling for Future Transformers on Future Hardware," May 2023, arXiv:2302.02825 [cs].
- [5] R. Lucas *et al.*, "DOE Advanced Scientific Computing Advisory Subcommittee (ASCAC) Report: Top Ten Exascale Research Challenges," Office of Scientific and Technical Information (OSTI), U.S. Department of Energy, Tech. Rep. 1222713, Feb. 2014.
- [6] D. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proceedings of the IEEE*, vol. 88, no. 6, pp. 728–749, Jun. 2000.
- [7] B. G. Lee *et al.*, "Beyond CPO: A Motivation and Approach for Bringing Optics Onto the Silicon Interposer," *Journal of Lightwave Technology*, vol. 41, no. 4, pp. 1152–1162, Feb. 2023, conference Name: Journal of Lightwave Technology.
- [8] N. Margalit *et al.*, "Perspective on the future of silicon photonics and electronics," *Applied Physics Letters*, vol. 118, no. 22, p. 220501, May 2021.
- [9] K. Hosseini *et al.*, "8 Tbps Co-Packaged FPGA and Silicon Photonics Optical IO," in *Optical Fiber Communication Conference (OFC) 2021*. Washington, DC: Optica Publishing Group, 2021, p. Th4A.2.
- [10] K. Hosseini *et al.*, "5.12 Tbps Co-Packaged FPGA and Silicon Photonics Interconnect I/O," in *2022 IEEE Symposium on VLSI Technology and Circuits (VLSI Technology and Circuits)*. Honolulu, HI, USA: IEEE, Jun. 2022, pp. 260–261.

- [11] R. Mahajan *et al.*, "Co-Packaged Photonics For High Performance Computing: Status, Challenges And Opportunities," *Journal of Lightwave Technology*, vol. 40, no. 2, pp. 379–392, Jan. 2022.
- [12] S. Daudlin *et al.*, "3D photonics for ultra-low energy, high bandwidth-density chip data links," Oct. 2023, arXiv:2310.01615 [physics].
- [13] A. Rizzo *et al.*, "Petabit-Scale Silicon Photonic Interconnects With Integrated Kerr Frequency Combs," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 29, no. 1: Nonlinear Integrated Photonics, pp. 1–20, Jan. 2023, conference Name: IEEE Journal of Selected Topics in Quantum Electronics.
- [14] Q. Xu *et al.*, "Micrometre-scale silicon electro-optic modulator," *Nature*, vol. 435, no. 7040, pp. 325–327, May 2005.
- [15] W. Bogaerts *et al.*, "Silicon microring resonators," *Laser & Photonics Reviews*, vol. 6, no. 1, pp. 47–73, Jan. 2012.
- [16] A. L. Gaeta *et al.*, "Photonic-chip-based frequency combs," *Nature Photonics*, vol. 13, no. 3, pp. 158–169, Mar. 2019.
- [17] L. Chang *et al.*, "Integrated optical frequency comb technologies," *Nature Photonics*, vol. 16, no. 2, pp. 95–108, Feb. 2022.
- [18] Y. Okawachi *et al.*, "Chip-scale frequency combs for data communications in computing systems," *Optica*, vol. 10, no. 8, pp. 977–995, Aug. 2023, publisher: Optica Publishing Group.
- [19] H. Shu *et al.*, "Microcomb-driven silicon photonic systems," *Nature*, vol. 605, no. 7910, pp. 457–463, May 2022.
- [20] A. Rizzo *et al.*, "Massively scalable Kerr comb-driven silicon photonic link," *Nature Photonics*, vol. 17, no. 9, pp. 781–790, Sep. 2023, number: 9 Publisher: Nature Publishing Group.
- [21] B. Y. Kim *et al.*, "Turn-key, high-efficiency Kerr comb source," *Optics Letters*, vol. 44, no. 18, p. 4475, Sep. 2019.
- [22] Q. Xu *et al.*, "Cascaded silicon micro-ring modulators for WDM optical interconnection," *Optics Express*, vol. 14, no. 20, p. 9431, 2006.
- [23] C.-H. Chen *et al.*, "A comb laser-driven DWDM silicon photonic transmitter based on microring modulators," *Optics Express*, vol. 23, no. 16, p. 21541, Aug. 2015.
- [24] D. Kong *et al.*, "Intra-Datacenter Interconnects With a Serialized Silicon Optical Frequency Comb Modulator," *Journal of Lightwave Technology*, vol. 38, no. 17, pp. 4677–4682, Sep. 2020.
- [25] M. Wade *et al.*, "TeraPHY: A Chiplet Technology for Low-Power, High-Bandwidth In-Package Optical I/O," *IEEE Micro*, vol. 40, no. 2, pp. 63–71, Mar. 2020.
- [26] E. Timurdogan *et al.*, "An ultralow power athermal silicon modulator," *Nature Communications*, vol. 5, no. 1, p. 4008, Sep. 2014.
- [27] Y. Wang *et al.*, "Scalable architecture for sub-pJ/b multi-Tbps comb-driven DWDM silicon photonic transceiver," in *Next-Generation Optical Communication: Components, Sub-Systems, and Systems XII*, G. Li *et al.*, Eds. San Francisco, United States: SPIE, Mar. 2023, p. 55.
- [28] A. James *et al.*, "Scaling comb-driven resonator-based DWDM silicon photonic links to multi-Tb/s in the multi-FSR regime," *Optica*, vol. 10, no. 7, pp. 832–840, Jul. 2023, publisher: Optica Publishing Group.
- [29] K. Padmaraju *et al.*, "Intermodulation Crosstalk Characteristics of WDM Silicon Microring Modulators," *IEEE Photonics Technology Letters*, vol. 26, no. 14, pp. 1478–1481, Jul. 2014.
- [30] E. Timurdogan *et al.*, "An Interior-Ridge Silicon Microring Modulator," *Journal of Lightwave Technology*, vol. 31, no. 24, pp. 3907–3914, Dec. 2013.
- [31] M. Gehl *et al.*, "Operation of high-speed silicon photonic micro-disk modulators at cryogenic temperatures," *Optica*, vol. 4, no. 3, pp. 374–382, Mar. 2017, publisher: Optica Publishing Group.
- [32] H. Gevorgyan *et al.*, "Miniature, highly sensitive MOSCAP ring modulators in co-optimized electronic-photonics CMOS," *Photonics Research*, vol. 10, no. 1, p. A1, Jan. 2022.
- [33] S. Wang *et al.*, "Integrated, Compact, and Tunable Band-Interleaving of a Kerr Comb Source," in *CLEO 2023*. San Jose, CA: Optica Publishing Group, 2023, p. STh3J.6.
- [34] E. S. Magden *et al.*, "Transmissive silicon photonic dichroic filters with spectrally selective waveguides," *Nature Communications*, vol. 9, no. 1, p. 3009, Aug. 2018.
- [35] Asher Novick *et al.*, "Tunable and Compact SiP Quasi-Dichroic Filter with ≥ 10 dB/nm Roll-Off Across C- & L-bands," in *2022 International Conference on Numerical Simulation of Optoelectronic Devices (NUSOD)*, Sep. 2022, pp. 171–172, iSSN: 2158-3242.
- [36] W. Shi *et al.*, "Ultra-compact, flat-top demultiplexer using anti-reflection contra-directional couplers for CWDM networks on silicon," *Optics Express*, vol. 21, no. 6, p. 6733, Mar. 2013.
- [37] J. Chen *et al.*, "Flat-Top CWDM (De)Multiplexers Based on Contra-Directional Couplers With Subwavelength Gratings," *IEEE Photonics Technology Letters*, vol. 31, no. 24, pp. 2003–2006, Dec. 2019.
- [38] A. Novick *et al.*, "High-bandwidth density silicon photonic resonators for energy-efficient optical interconnects," *Applied Physics Reviews*, vol. 10, no. 4, p. 041306, Nov. 2023.
- [39] M. van Niekerk *et al.*, "Wafer-Scale-Compatible Substrate Undercut for Ultra-Efficient SOI Thermal Phase Shifters," in *Conference on Lasers and Electro-Optics*. San Jose, California: Optica Publishing Group, 2022, p. JTh3B.24.
- [40] A. Rizzo *et al.*, "Ultra-efficient foundry-fabricated resonant modulators with thermal undercut," in *CLEO 2023*, ser. CLEO 2023. Optica Publishing Group, 2023, p. SF2K.6.
- [41] K. Padmaraju *et al.*, "Thermal stabilization of a microring modulator using feedback control," *Optics Express*, vol. 20, no. 27, p. 27999, Dec. 2012.
- [42] D. Y. Lee *et al.*, "Error-free operation of a polarization-insensitive $4\lambda \times 25$ Gbps silicon photonic WDM receiver with closed-loop thermal stabilization of Si microrings," *Optics Express*, vol. 24, no. 12, p. 13204, Jun. 2016.
- [43] H. Li *et al.*, "A 25 Gb/s, 4.4 V-Swing, AC-Coupled Ring Modulator-Based WDM Transmitter with Wavelength Stabilization in 65 nm CMOS," *IEEE Journal of Solid-State Circuits*, vol. 50, no. 12, pp. 3145–3159, Dec. 2015.
- [44] K. Yu *et al.*, "A 25 Gb/s Hybrid-Integrated Silicon Photonic Source-Synchronous Receiver With Microring Wavelength Stabilization," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 9, pp. 2129–2141, Sep. 2016.
- [45] S. Agarwal *et al.*, "Wavelength Locking of a Si Ring Modulator Using an Integrated Drop-Port OMA Monitoring Circuit," *IEEE Journal of Solid-State Circuits*, vol. 51, no. 10, pp. 2328–2344, Oct. 2016.
- [46] M. Rakowski *et al.*, "Hybrid 14nm FinFET - Silicon Photonics Technology for Low-Power Tb/s/mm² Optical I/O," in *2018 IEEE Symposium on VLSI Technology*. Honolulu, HI: IEEE, Jun. 2018, pp. 221–222.
- [47] D. Coenen *et al.*, "Thermal Modelling of Silicon Photonic Ring Modulator with Substrate Undercut," *Journal of Lightwave Technology*, vol. 40, no. 13, pp. 4357–4363, Jul. 2022.
- [48] L.-W. Luo *et al.*, "High bandwidth on-chip silicon photonic interleaver," *Optics Express*, vol. 18, no. 22, p. 23079, Oct. 2010.
- [49] A. Rizzo *et al.*, "Ultra-Broadband Interleaver for Extreme Wavelength Scaling in Silicon Photonic Links," *IEEE Photonics Technology Letters*, vol. 33, no. 1, pp. 55–58, Jan. 2021.
- [50] F. Horst *et al.*, "Cascaded Mach-Zehnder wavelength filters in silicon photonics for low loss and flat pass-band WDM (de-)multiplexing," *Optics Express*, vol. 21, no. 10, p. 11652, May 2013.
- [51] T. Akiyama *et al.*, "Cascaded AMZ triplets: a class of demultiplexers having a monitor and control scheme enabling dense WDM on Si nano-waveguide PICs with ultralow crosstalk and high spectral efficiency," *Optics Express*, vol. 29, no. 6, p. 7966, Mar. 2021.
- [52] H. Matsuura *et al.*, "Fast Frequency Tuning of Silicon-Photonic Thermo-optic MZI Filters using "Turbo Pulse" Method," in *Optical Fiber Communication Conference*. San Diego, California: OSA, 2018, p. M4H.2.
- [53] M. W. AlTaha *et al.*, "Monitoring and automatic tuning and stabilization of a 2x2 MZI optical switch for large-scale WDM switch networks," *Optics Express*, vol. 27, no. 17, p. 24747, Aug. 2019.
- [54] K. Entesari *et al.*, "Automated Tuning for Silicon Photonic Filters," in *Optical Fiber Communication Conference (OFC) 2022*. San Diego, California: Optica Publishing Group, 2022, p. Th1D.6.
- [55] S. Wang *et al.*, "Automated tuning of ring-assisted MZI-based interleaver for DWDM systems," in *Optical Fiber Communication Conference (OFC)*, 2024, under review.
- [56] A. Novick *et al.*, "Ultra-efficient interleaved vertical-junction microdisk modulator with integrated heater," in *Optical Fiber Communication Conference (OFC)*, 2024, under review.
- [57] A. Biberman *et al.*, "Adiabatic microring modulators," *Optics Express*, vol. 20, no. 28, pp. 29223–29236, Dec. 2012, publisher: Optica Publishing Group.
- [58] L. Alloatti *et al.*, "High-speed modulator with interleaved junctions in zero-change CMOS photonics," *Applied Physics Letters*, vol. 108, no. 13, p. 131101, Mar. 2016.
- [59] V. Gopal *et al.*, "Simultaneous Error-Free Data Modulation with Silicon Microdisks in the Multi-FSR Regime for Scalable DWDM Links," in *2023 Optical Fiber Communications Conference and Exhibition (OFC)*, Mar. 2023, pp. 1–3.
- [60] A. Rizzo *et al.*, "Fabrication-robust silicon photonic devices in standard sub-micron silicon-on-insulator processes," *Optics Letters*, vol. 48, no. 2, p. 215, Jan. 2023.
- [61] R. Parsons *et al.*, "Efficient silicon photonic add-drop microdisk filters for dwdm systems," in *CLEO 2023*. Optica Publishing Group, 2023, p. JW2A.127.