

Research Article

Robert Parsons*, Alexander Oh, James Robinson, Songli Wang, Michael Cullen, Kaylx Jang, Aneek James, Yuyang Wang and Keren Bergman

Foundry-enabled wafer-scale characterization and modeling of silicon photonic DWDM links

<https://doi.org/10.1515/nanoph-2025-0439>

Received September 5, 2025; accepted November 25, 2025;

published online December 11, 2025

Abstract: AI/ML compute clusters are driving unprecedented bandwidth demands at the package boundary, motivating co-packaged integrated photonics closely co-located with the compute unit. We present a scalable silicon-photonics transceiver platform and a measurement-driven design methodology that together enable dense, energy-efficient DWDM links suitable for in-socket integration. Automated wafer-scale probing on 300 mm active photonic wafers extracts waveguide and resonator statistics using index fitting and comprehensive device characterization. The resulting wafer-scale measurements highlight design points such as wider robust waveguides, whispering gallery mode resonators, and thermally efficient undercut devices, that reduce required thermal tuning power and tighten insertion loss distributions. We propagate the measured distributions through a system model via large-scale Monte Carlo simulations to derive realistic link margins and source power targets. Together, the scalable architecture and wafer-scale measurement-informed design process offer a practical path to high-bandwidth, low energy consumption DWDM links with robust yield.

Keywords: silicon photonics; wafer-scale photonics; DWDM links

1 Introduction

With the extensive popularity of compute clusters in data centers for artificial intelligence and machine learning applications, the demand for higher bandwidth interconnects continues to grow. However, current interconnect solutions cannot match the performance scaling of compute resources while having significant environmental impact due to their power consumption [1]–[5]. Silicon photonics has emerged a promising platform for pushing higher data bandwidths with lower energy consumption and smaller footprint than their electrical counterparts, while being feasible at an industrial scale. By utilizing existing CMOS technology and fabrication, silicon has become the leading choice in larger-scale photonic systems [6]–[9]. Despite the maturing fabrication process, the importance of statistical insight for device performance at the wafer-scale cannot be understated, especially for silicon photonics where non-uniformities can significantly degrade device performance [10], [11], and thus impacting link performance. Devices such as microring and microdisk resonators are the foundation of a scalable dense wavelength-division multiplexing (DWDM) photonic link, which can see significant changes in resonance location, quality factor, and modulation efficiency across a single wafer [12]–[16]. To compensate for static fabrication variations, devices are commonly integrated with micro-heaters driven by DC voltages for thermo-optic tuning [17], [18]. Hence, deviations from nominal design parameters heavily influence the energy consumption, potentially leading to inefficient operation. Additionally, resonant modulator device variations can impact the insertion loss of the device, leading to variation of associated power penalties in the link budget [19], [20]. An understanding of the statistics governing device variations allows us to model the system performance of a DWDM silicon photonic link. Here, we present a comprehensive methodology for wafer-scale measurement and data extraction of individual devices to analyze trends and correlations across 300 mm active photonic wafers. We demonstrate this in the context of a scalable architecture for co-packaged integrated

*Corresponding author: Robert Parsons, Department of Electrical Engineering, Columbia University, New York, NY 10027, USA, E-mail: rp3020@columbia.edu. <https://orcid.org/0000-0001-5141-5597>

Alexander Oh, James Robinson, Songli Wang, Michael Cullen, Kaylx Jang, Aneek James and Keren Bergman, Department of Electrical Engineering, Columbia University, New York, NY 10027, USA

Yuyang Wang, Department of Electrical Engineering, Columbia University, New York, NY 10027, USA; and Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT 06269, USA, E-mail: yuyang.wang@uconn.edu

photonic transceivers. Through analysis of the extracted parameters across a wafer, we estimate the expected energy efficiency of the overall link. Furthermore, we investigate insertion loss and power penalties at the wafer-scale to inform the expected variation of the link budget.

2 Scalable architecture

Microresonator-based DWDM transceivers driven by optical frequency combs offer compelling footprint and energy advantages for short-reach datacom [21], [22] over CW-laser arrays [23] or non-resonant modulators/filters that need extra (de-)MUX stages [24]. A common design, one bus waveguide with cascaded, slightly detuned rings, assigns each ring a wavelength while remaining transparent off-resonance [13]. The resonators are designed to have as-fabricated resonant wavelengths slightly below that of their respective comb carrier wavelengths. Integrated thermo-optic phase shifters are used to tune the modulator or filter resonant wavelength to match their carrier wavelength. Scaling breaks down once comb bandwidth exceeds the rings' free spectral range (FSR): periodic resonances introduce resonance aliases that coincide with non-target comb lines, producing crosstalk, while packing more channels within one FSR by shrinking spacing incurs strong inter-modulation; spacings below ~ 100 GHz are experimentally unacceptable [25]. Increasing FSR via smaller radii helps in principle [15], [26]–[28] but scales poorly due to fabrication variability, tight heater/RF tolerances, and higher off-resonance loss. Band-interleaving partitions the comb into sub-bands that fit within an FSR [19], yet demands steep-roll-off, low-crosstalk, tunable elements – e.g., dichroic filters and contra-directional couplers – that are still maturing [29]–[33]. To overcome these limits, we propose an architecture that scales along two orthogonal axes: even–odd channel interleaving and multi-FSR channel allocation. Together, they support much denser combs across bandwidths far exceeding any single resonator's FSR while maintaining low crosstalk and high efficiency – enabling massively parallel, energy-efficient optical interconnects [19].

Even–odd channel interleaving alternates comb lines into “even” and “odd” groups that are routed through separate banks of cascaded microresonators [19]. With a (de-)interleaver whose FSR is twice the comb-line spacing – and assuming ideal alignment – each de-interleaving stage halves the channel density per bus while doubling the inter-channel spacing. In practice, silicon waveguide group velocity dispersion (GVD) and fabrication variations can misalign the interleaver pass/stop bands and the comb tones. Broadband Mach-Zehnder interferometer (MZI)-based [34]

(de-)interleavers – e.g., ring-assisted MZIs (RAMZIs) and MZI lattices – provide flat-top pass/stop responses, easing insertion-loss and crosstalk constraints and enabling tighter comb spacings (e.g., ≤ 100 GHz) while preserving the larger effective spacing seen by individual resonators [35]–[40].

Following even–odd de-interleaving, resonators with FSRs exceeding the retained comb bandwidth are still infeasible – small radii would be required, aggravating bend loss and fabrication tolerances [41]. We therefore adopt a multi-FSR allocation that places resonance aliases between active channels to maximize spectral isolation [19]. We define

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

where Δ_{ch} is the post-de-interleaved channel spacing, Δ_{agg} is the spacing to the closest alias (“aggressor”), and FSR is the resonator free spectral range. A configuration is valid when (1) S and \mathcal{F} are coprime integers and (2) $\mathcal{F} \geq N_{\text{ch}}$ (channels per bus). For example, $S = 3$, $\mathcal{F} = N_{\text{ch}} = 7$ separates each alias by one-third of the channel spacing; smaller S enlarges Δ_{agg} but demands larger FSR, with $S = 1$ reducing to the (often infeasible) single-FSR case. Increasing \mathcal{F} while remaining coprime with S accommodates more channels without degrading minimum aggressor spacing. In the 4×16 transceiver of Figure 1a, two de-interleaving stages with $S = 2$, $\mathcal{F} = 17$ permit up to 17 channels per bus at a moderate resonator FSR of 25.69 nm; with 100 GHz comb spacing, the effective aggressor spacing is 200 GHz, a regime where inter-channel crosstalk is negligible [25], [42]. Other coprime choices (e.g., $S = 2$, $\mathcal{F} = 19$) offer additional channels and can be exploited using comb-envelope knowledge or thermal tuning [43]; comprehensive design trade-offs appear in refs. [19], [44].

Our approach towards a scalable implementation uses a dual-ring Si_3N_4 Kerr comb in normal GVD for high efficiency, uniform line power, and stable coherence [45]. Transmit-side, two cascaded RAMZI de-interleavers split the spectrum into four groups, each modulated by its own bank of cascaded microresonator modulators; a symmetric two-stage RAMZI interleaver recombines the outputs to a single fiber. The receiver mirrors this with de-interleaving, microresonator filtering, and per-channel photodiodes. With 100 GHz base comb spacing and 16 devices per bank at 400 GHz spacing, the system supports 64 DWDM channels; at 16 Gbps per channel, the aggregate capacity is 1.024 Tbps per fiber. A silicon photonic transceiver using this scalable architecture can be co-packaged directly in the socket

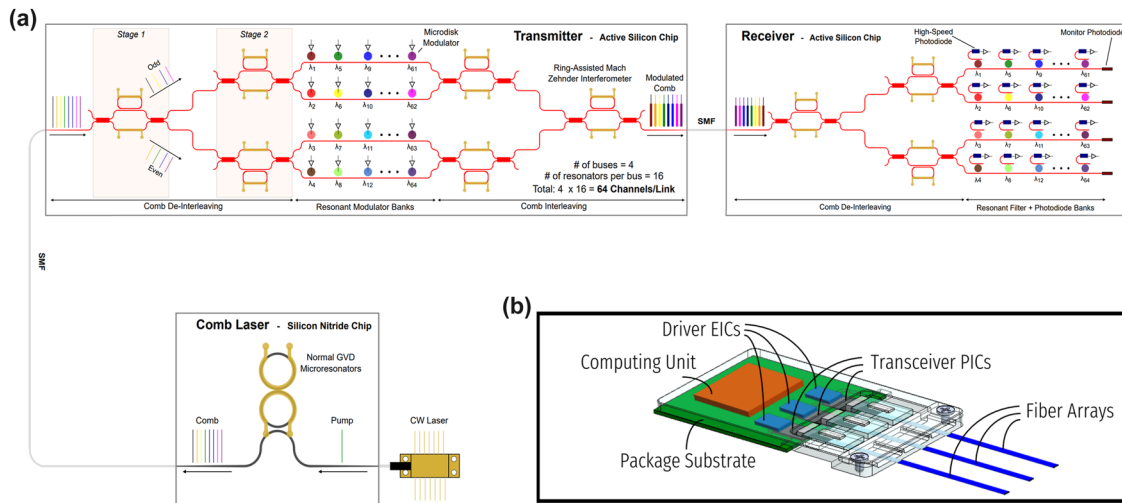


Figure 1: Co-packaged scalable link architecture. (a) Scalable link architecture. Normal-GVD comb laser is used as a multi-wavelength source, sent to the transmitter consisting of multiple stages of (de-)interleavers and microdisk modulators in multi-FSR regime. Encoded signals are recovered in the receiver with de-interleaver stages, microdisk filters in the multi-FSR regime, and photodetectors. (b) Conceptualization of optical transceiver utilizing scalable link architecture co-packaged with compute units in the socket.

with compute units to enable high bandwidth, low energy consumption links (Figure 1b) [46].

3 Results

We utilize a ficonTEC WLT-1200 wafer-level tester for automated probing of 300 mm photonic wafers [47]. Shown in

Figure 2a, the wafer prober operates within an enclosure with an air filter to minimize dust and contaminants. Inside, the wafer chuck can translate, rotate, and regulate the temperature of the wafer from 20 to 80 degrees C. An image of the wafer on the wafer chuck is shown in Figure 2b. There are four probe mounts available for use – two six-axis stages on the East and West for optical probing, and two five-axis stages on the North and South for electrical probing. Both DC

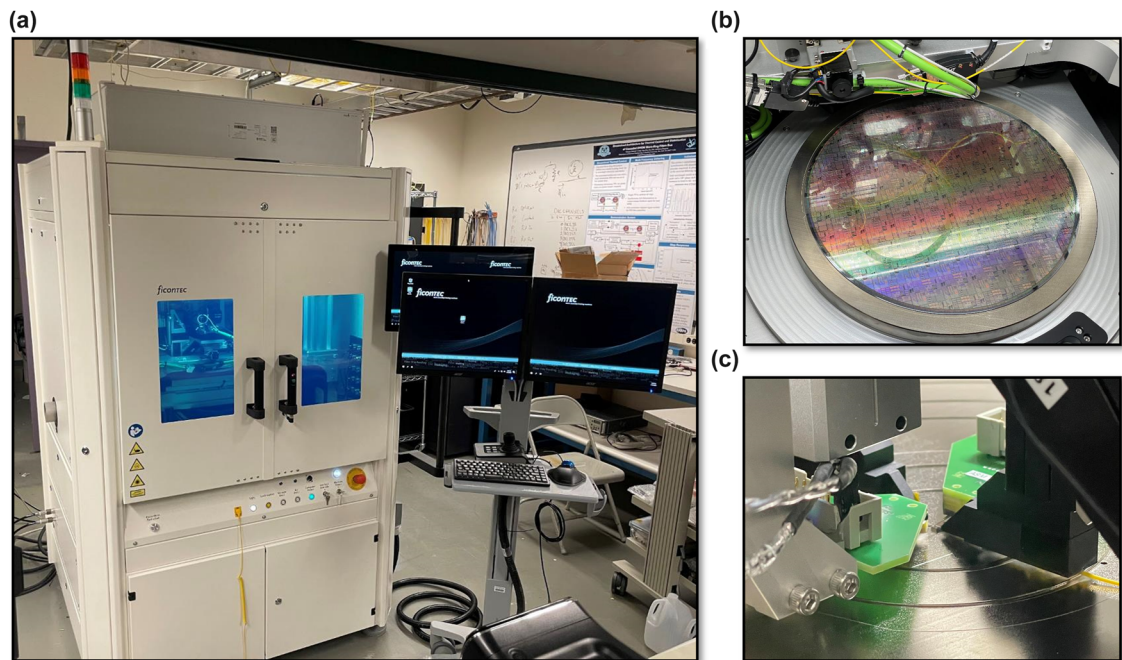


Figure 2: Wafer-level tester. (a) Image of ficonTEC wafer prober. (b) Image of 300 mm active photonic wafer on wafer table. (c) Image of four probes, electrical and optical, being used simultaneously on a single die.

and RF electrical probing is possible with the wafer prober. For wafer-scale optical probing, either angled fiber arrays can be used for grating (vertical) coupling, or periscope fiber arrays can be used for edge coupling. The periscope fiber arrays utilize a 3D-printed optic with a mirror and lens that is lowered into the dicing trench which enables edge coupling at the wafer-scale. Additionally, single dies can be measured with the wafer prober, as long as the die is placed carefully over the center vacuum hole. Figure 2c displays a single die being measured with all four probes in place. Calibration of the probes is achieved with three different cameras and visual recognition techniques. Similarly, a combination of die maps generated from the layout files and visual recognition enables the wafer prober to evaluate where each device is and where to land the necessary probes. Measurement equipment is connected externally and controlled with the wafer prober's process

control master via SCPI commands over TCP/IP. Generally, a broadband tunable laser and high-precision DC power supply are interfaced with to allow measurements ranging from passive optical sweeps to thermo-optic tuning and depletion response measurements.

3.1 Wafer-scale optical parameter extraction

First, we investigate the optical properties and geometric parameters of silicon waveguides. We use a large sweep of asymmetric, or imbalanced, Mach-Zehnder interferometers (AMZIs) with varying waveguide width and arm length imbalance (Figure 3e) across two wafers fabricated by AIM Photonics in dedicated full-reticle wafer runs. The MZIs are grating-coupled with de-embedding structures to subtract the loss envelope of the grating couplers, as shown

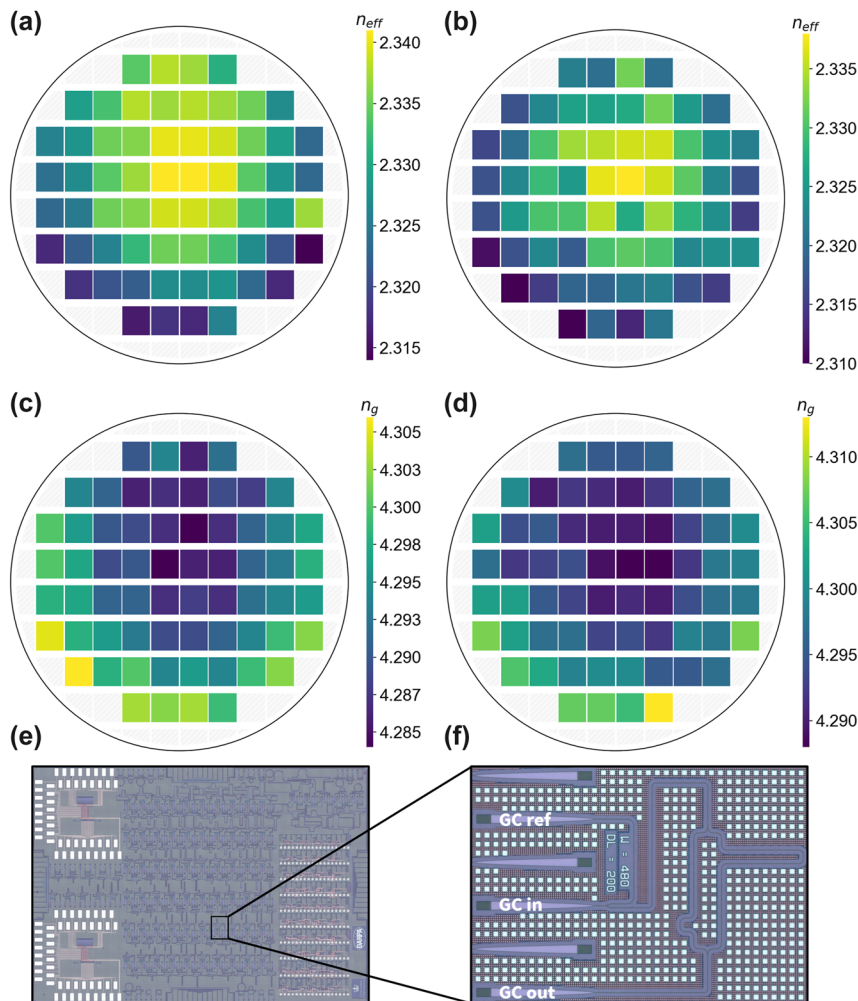


Figure 3: Wafer-scale characterization of waveguide effective and group index. (a) Wafermap of effective index at wavelength of 1,550 nm across wafer 1. (b) Wafermap of effective index at wavelength of 1,550 nm across wafer 2. (c) Wafermap of group index at wavelength of 1,550 nm across wafer 1. (d) Wafermap of group index at wavelength of 1,550 nm across wafer 2. (e) Micrograph of chip containing asymmetric Mach-Zehnder interferometers (AMZIs) with varying width and length imbalances. (f) Micrograph of representative grating-coupled AMZI test structure.

in a representative test structure in Figure 3f. Because a large sweep of waveguide widths is used, Euler bends are employed to maintain single-mode operation within multi-mode waveguides; interference fringes from higher order modes could interfere with the optical parameter extraction. The transmission spectrum of each AMZI across the wafer is measured with an optical sweep using a broadband tunable laser covering the S-, C-, and L-bands. A peak-finding algorithm is used to extract the wavelength locations of each interference fringe. From here, we utilize an ordinary least squares (OLS) regression to derive fitting parameters for the group index of the AMZI following the procedure outlined in ref. [48]. Because the fringe order of the interferometer must be known to find the effective index, we must use a method to infer it. The smaller the length imbalance of the AMZI, the larger the accuracy of the inferred fringe order. A mode solver simulation sweep across waveguide width and wavelength was performed to allow us to compare the extracted group index to that of the simulated group index. By minimizing the error between these, we can find the waveguide width that corresponds to the extracted group index. The effective index at this waveguide width can then be used to infer the fringe order of the measured device, ensuring a more accurate effective index extraction.

AMZIs were measured on two wafers. We plot the wafermaps of the effective index and group index for each

of the two wafers at the single mode waveguide width of 480 nm. The effective index is shown in Figure 3a and b for the first and second wafer, respectively, and the group index is shown in Figure 3c and d for the first and second wafer, respectively. It is clear that there is a radial relationship shown in each of the four wafermaps. To demonstrate this statistically, we mapped the locations of the reticles to polar coordinates in reference to the center of the wafer and performed an OLS regression. The correlations for radial relationships for effective index were 0.507 and 0.623 for each wafer, respectively, and for group index, 0.527 and 0.502, respectively. As suggested by the slanted bullseye appearance of the wafermaps, there also exists a planar relationship. Again, we performed an OLS regression, this time using Cartesian coordinates and projecting across the wafer at varying angles. This way, we understand not only the strength of the planar relationship, but also the angle at which the plane is situated. The correlations and angles for planar relationships for effective index were 0.351 (264.15 deg) and 0.174 (245.18 deg) for each wafer, respectively, and for group index, 0.355 (259.72 deg) and 0.291 (258.72 deg), respectively. These trends could be explained by temperature variations in the process.

The resulting extracted wafer-scale optical parameters are displayed in Figure 4. The effective and group indices averaged over the 64 reticles across each wafer are shown

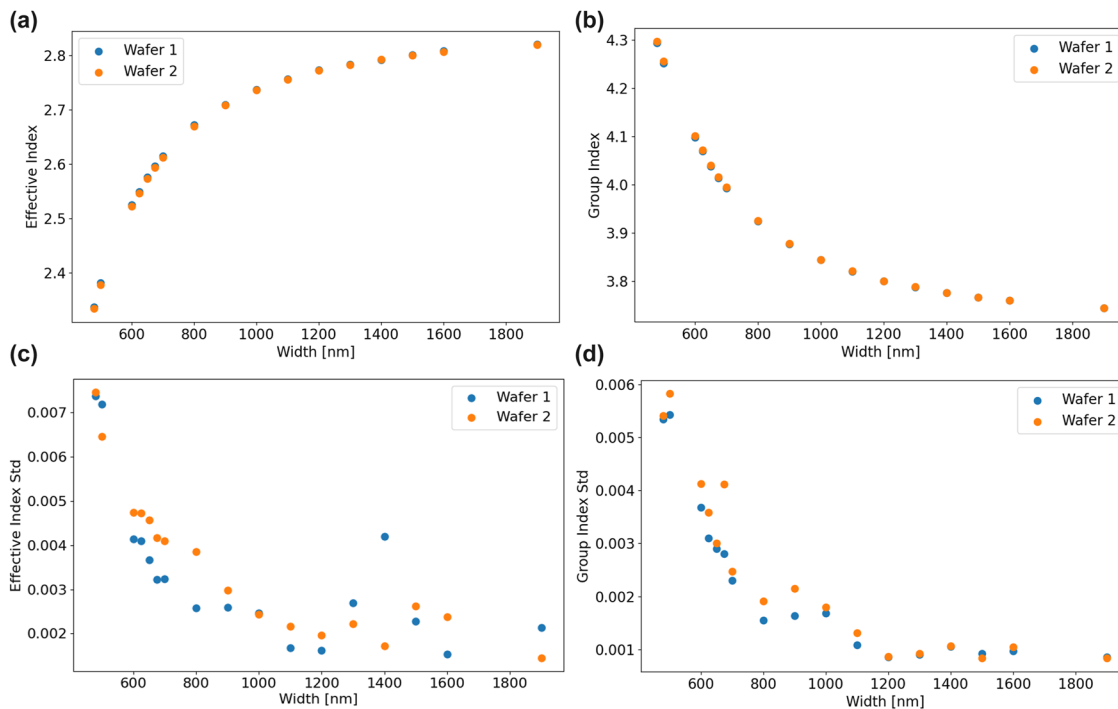


Figure 4: Effective and group index statistics across waveguide width. (a) Effective index averaged across each wafer at different waveguide widths at wavelength of 1,550 nm. (b) Group index averaged across each wafer at different waveguide widths at wavelength of 1,550 nm. (c) Standard deviation of effective index averaged across each wafer at different waveguide widths. (d) Standard deviation of group index averaged across each wafer at different waveguide widths.

in Figure 4a and b, respectively. The effective and group indices are calculated across waveguide widths supporting only a single mode up to those supporting many higher order modes. As expected, both parameters head toward an asymptote at highly multimode waveguide widths, as their behavior starts to exhibit that of a slab mode. The high consistency between wafers is noteworthy for each parameter. Additionally, the standard deviation for both the effective index and group index is displayed in Figure 4c and d, respectively. There is a clear trend that larger waveguide widths experience drastically reduced statistical variation

in both effective index and group index. This is explained by the mode being less perturbed proportionally by etching biases due to the much larger width. Further, the mode of larger waveguides is more confined and is less affected by changes in sidewall. There is relatively low variation between the two measured wafers, despite the two wafers being from different lots. We expect wafer-to-wafer variations to be lower within the same lot compared to variability between wafers from different lots.

We measured microdisk modulators and filters on our latest 300 mm wafer taped out with AIM Photonics.

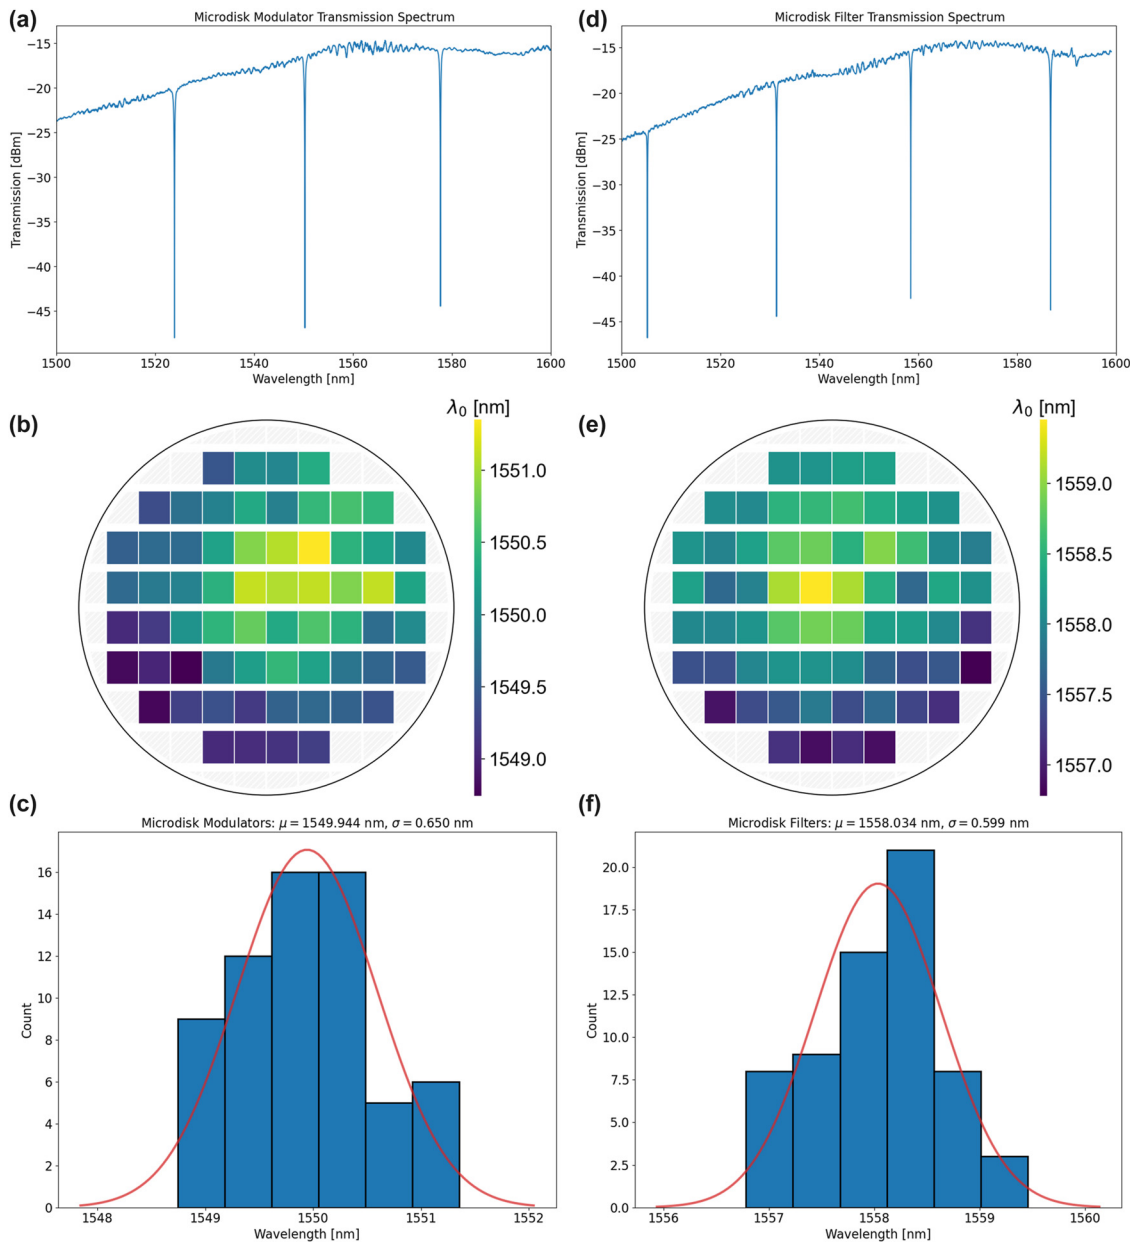


Figure 5: Wafer-scale microresonator characterization. (a) Transmission spectrum of a microdisk modulator. (b) Wafermap of microdisk modulator resonance location across the wafer. (c) Histogram and statistical data of microdisk modulator. (d) Transmission spectrum of a microdisk filter. (e) Wafermap of microdisk filter resonance location across the wafer. (f) Histogram and statistical data of microdisk filter.

The transmission spectrum of a representative microdisk modulator is shown in Figure 5a, the wafermap displaying the resonance locations across the wafer is shown in Figure 5b, and a histogram of the distribution is displayed in Figure 5c. Similarly, the transmission spectrum, wafermap, and histogram for the microdisk filter are shown in Figure 5d–f, respectively. The variation of the modulator is slightly higher than that of the filter (standard deviation of 0.65 nm vs. 0.6 nm) due to the modulator junction width being slightly smaller to reduce series resistance and factoring in variations of the doped junction in the modulator. However, both demonstrate high robustness to variations compared especially to microring resonators. Microrings of a similar radius demonstrated standard deviations on the order of multiple nanometers, further solidifying these as robust devices.

3.2 Implications for link energy consumption

With a suite of devices now measured and analyzed at the wafer-scale, we can use these results to calculate power consumption to inform component and link design (Table 1). The variation of effective index, extracted from the AMZI test structures, can be converted mathematically to waveguide phase error. This phase error due to fabrication variations must be compensated for by thermal tuning of the device. In our scalable link architecture, the devices that are sensitive to these waveguide phase errors are the interleavers. Therefore, we can calculate the required thermal tuning per thermo-optic phase shifter within the interleaver. Accounting for 3σ variation allows us to cover about 99.7% of all devices, statistically speaking. From the effective index standard deviation of 480 nm wide waveguides, we calculate a 3σ phase error of 7.887 radians. Assuming a tuning efficiency of 25 mW P_π , on average 31.38 mW would be required. However, as we noted previously, moving to a larger waveguide width reduces the variation of optical parameters significantly. Therefore, using wider waveguides in the interleavers enables higher robustness

Table 1: Energy consumption of link components due to thermal tuning requirements.

Component	3σ phase error [rad]	Power [mW] (nominal)	Power [mW] (undercut)
Interleaver (480 nm)	7.887	31.38	1.34
Interleaver (1,200 nm)	2.154	8.57	0.37
Microdisk modulator (4 μm)	0.451	5.91	1.18
Microdisk filter (4 μm)	0.416	2.46	0.49

to fabrication variations and an associated reduced energy consumption. Using the same tuning efficiency, 1,200 nm wide waveguides would exhibit a 3σ phase error of 2.154 radians, leading to an average tuning power of 8.57 mW. This is a significant reduction. Moreover, the 3σ phase error being less than π enables easier alignment to channels. Moving on to the microdisks, using the numbers reported in Figure 5 and the device FSR [49], we have 3σ phase errors of 0.451 and 0.416 radians for the microdisk modulators and filters, respectively. This leads to an associated average tuning power of 5.91 and 2.46 mW, respectively. The filter exhibits a larger tuning efficiency due to no RF contact vias and closer placement to the optical mode. Using a thermal undercut of the substrate significantly improved the thermal tuning efficiency of devices [50]. If we apply those improvement factors here to these already robust components, then the average tuning power required would now be 1.34 mW, 0.37 mW, 1.18 mW, and 0.49 mW for the 480 nm waveguide phase shifter, 1,200 nm waveguide phase shifter, microdisk modulator, and microdisk filter, respectively. The resulting power consumption that each component contributes to the link is calculated and displayed in Table 1. Clearly, utilizing robust component designs and strategies to improve tuning efficiency, such as thermal undercut, are imperative to reducing the power consumption of links due to thermal tuning requirements.

3.3 Wafer-scale link parameter extraction

First, we look at one of the primary contributors to the link loss budget – the power penalties associated with modulation. Important to understanding the power penalties associated with the resonant modulator, shown in Figure 6a, is the depletion response of the device (Figure 6b). The insertion loss (IL) and extinction ratio (ER) vary depending on the driving voltage of the modulator and carrier detuning from the resonant wavelength. As shown in Figure 6c, the IL, ER, and associated power penalties can be calculated based on the detuning of the carrier from the resonant wavelength. The total power penalty is given by the following equation:

$$PP_{\text{total}} = IL_{\text{mod}} + PP_{\text{ER}} + PP_{\text{OOK}} \quad (2)$$

where

$$PP_{\text{ER}} = 10 \log_{10} \left(\frac{r+1}{r-1} \right), \quad (3a)$$

$$PP_{\text{OOK}} = 10 \log_{10} \left(\frac{2r}{r+1} \right), \quad (3b)$$

and r is the linear extinction ratio. As the extinction ratio increases and approaches infinity, the power penalty associated with the ER approaches zero, while the power penalty

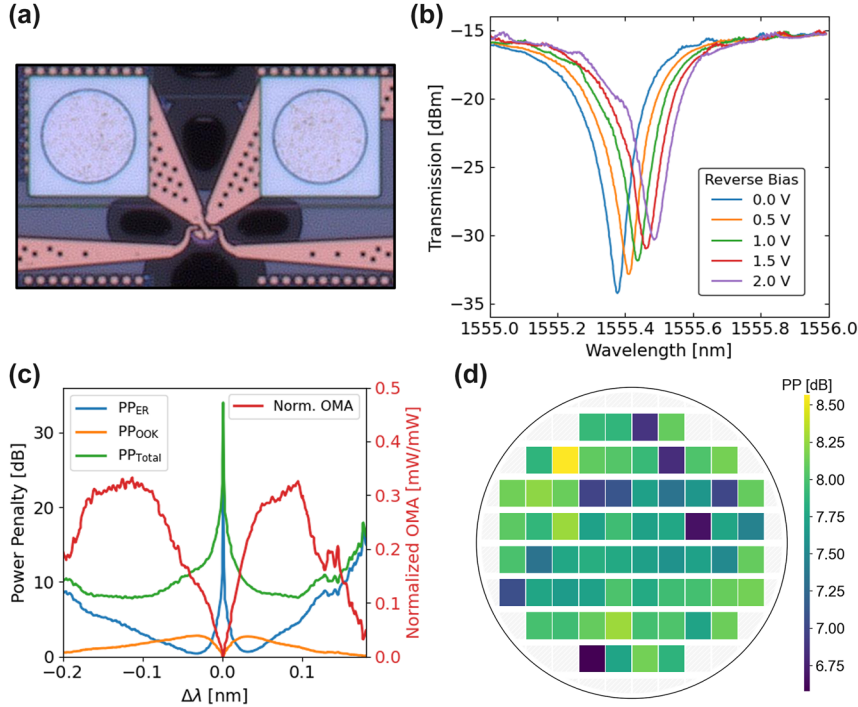


Figure 6: Microdisk modulator power penalties. (a) Micrograph of a microdisk modulator. (b) Depletion response of microdisk modulator demonstrating changing spectrum with varied applied reverse bias to the junction. (c) Power penalties, insertion loss, and optical modulation amplitude of modulator as function of detuning wavelength. (d) Wafermap of total power penalty of modulators in each reticle.

associated with on–off keying approaches 3 dB (equal distribution of ‘1’s and ‘0’s), for 3 dB total. There are diminishing returns of power penalty after achieving an ER of at least 10 dB. A large depletion response is necessary for a small insertion loss while maintaining a low driving voltage. In addition, the optical modulation amplitude (OMA) is shown to be maximal at the minimum of the total power penalty; the OMA normalized to the input power is defined as

$$\text{OMA}_{\text{norm}} = (P_1 - P_0)/P_{\text{in}}, \quad (4)$$

where P_1 is the ‘on’ state optical power, P_0 is the ‘off’ state optical power, and P_{in} is the input optical power. With the ability to calculate these power penalties from the depletion response, we measured the depletion response of disk modulators across a wafer. The calculated total power penalty, including OOK and ER penalties and IL, is shown in the wafermap in Figure 6d. With a driving voltage of only 1 V, the mean total power penalty is 7.78 dB, with a standard deviation of only 0.036 dB. These sub-dB edge couplers demonstrate high robustness to variations in both fabrication and positioning of the fiber array (similar process to that used when attaching the fiber array), critical to realizing ultra-broadband links.

Coupling between fiber and chip has traditionally also been a challenge in terms of contending with excess losses [51]–[53]. Although grating couplers are available, we use edge couplers for both higher coupling efficiency and larger optical bandwidth, important for broadband DWDM links. A representative edge coupler is shown in Figure 7a. The substrate below the silicon nitride edge coupler is partially removed and backfilled with silicon dioxide. This dramatically improves the mode matching between the cleaved SMF28 mode and that of the silicon nitride waveguide, thereby dramatically improving the possible coupling efficiency. Low-loss and broadband escalators enable coupling between silicon nitride and silicon device layers. Using a periscope fiber array, we obtained wafer-scale per-facet coupling losses by measuring a loopback structure in each of the 64 reticles. The loss per coupler is shown in Figure 7b, with a mean coupling efficiency of 0.924 dB and standard deviation of only 0.036 dB. These sub-dB edge couplers demonstrate high robustness to variations in both fabrication and positioning of the fiber array (similar process to that used when attaching the fiber array), critical to realizing ultra-broadband links.

On the receiver side, wavelength-selective resonant filters must be used to filter out and drop each encoded carrier. Assuming a large enough filter bandwidth (or alternatively,

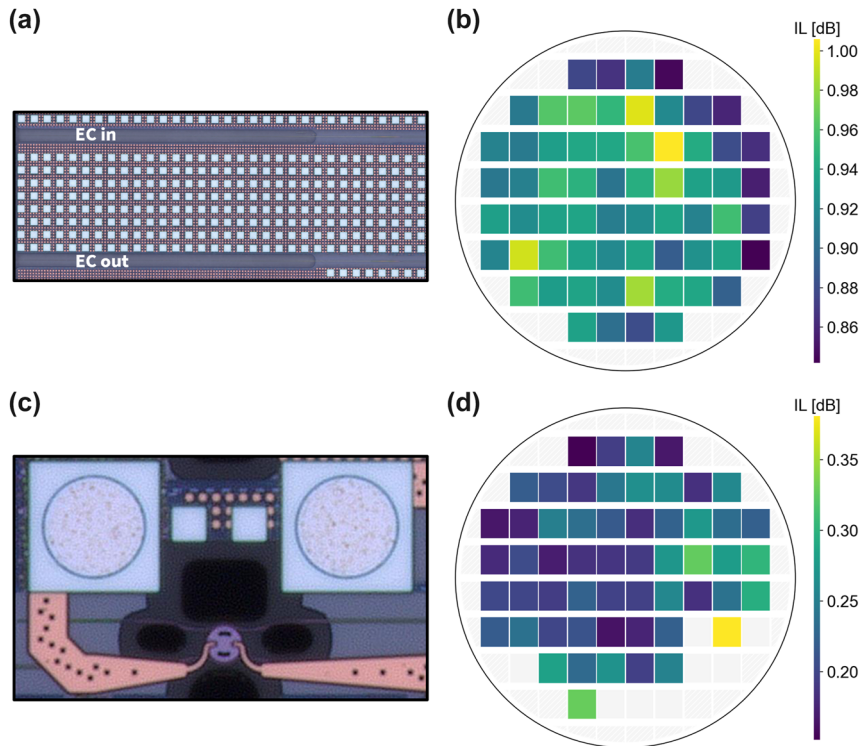


Figure 7: Wafer-scale insertion loss of edge couplers and microdisk filters. (a) Micrograph of an edge coupler with substrate removed underneath. (b) Wafermap of per-facet coupling losses between edge couplers and cleaved SMF28. (c) Micrograph of a microdisk filter. (d) Wafermap of insertion loss seen at the drop port of add-drop microdisk filters.

lower quality factor) is used to avoid attenuating parts of the signal, the main parameter of interest is the insertion loss seen at the drop port of the device. If critical coupling is not achieved, the insertion loss can become quite high, straining the link loss budget. We utilize microdisk filters for their robustness to fabrication variations (Figure 7c). Another advantage is that the cavity loss is small, owing to the supported whispering gallery mode. This allows a simple symmetric coupling scheme to produce low insertion loss add-drop filters; a wafermap of the insertion loss seen at the drop port of microdisk filters is shown in Figure 7d. The mean and standard deviation insertion loss is 0.225 dB and 0.047 dB, respectively. Introducing a slight asymmetry to the coupling can reduce this loss even further.

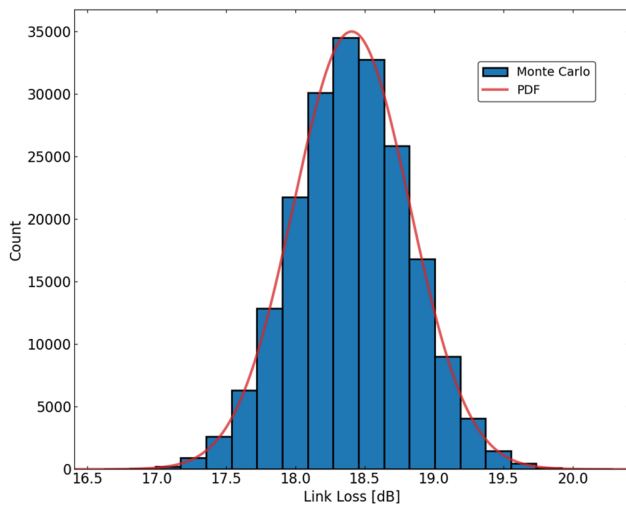
3.4 Implications for link optical power budget

The various insertion losses and power penalties of each component within the link measured from test structures are tallied in Table 2. Statistical information is included where mass-measurement was possible, however all values are based on measured data. To better understand

the system-level implications of component-level variability in integrated photonic links, we perform a Monte Carlo simulation using the available statistical data obtained from wafer-scale measurements. While one-off measurements provide estimates of link loss, they do not capture the full distributional behavior of the link budget. A Monte Carlo approach enables us to propagate measured variations, including insertion losses, coupling efficiencies, and modulation power penalties, through the full link architecture. This allows for a more accurate estimation of performance metrics such as link yield for a given power budget and worst-case design margins (e.g., 3σ link loss). Using the available wafer-scale measurements, 200,000 Monte Carlo simulations were performed on the link budget (Figure 8). The resulting mean link loss was 18.407 dB, and the corresponding standard deviation was 0.418 dB. Given a receiver sensitivity for error-free operation at 16 Gbps per channel [49], incorporating the 3σ link loss into our tabulated link budget as the margin allows us to describe the expected behavior of 99.7 % of links (Table 2). By statistically sampling from the known distributions of each component, we gain insights into opportunities for targeted yield improvement through component-level optimization.

Table 2: Link loss budget with tabulated per-component losses, including statistical information where available.

Loss component	Value	1σ	Unit	Multiplier	Source
Transmitter loss total	15.06	–	dB		
Comb to fiber	2.20	–	dB		Measured
Edge coupler in	0.924	0.036	dB		Measured
De-interleaver stages	0.35	–	dB	×2	Measured
Modulator off-resonance	0.10	–	dB	×15	Measured
Modulator power penalties	7.78	0.41	dB		Measured depletion response
Crosstalk penalty	–	–	dB		>100 GHz channel spacing
Interleaver stages	0.35	–	dB	×2	Measured
Propagation loss	0.33	–	dB		PIC layout
Edge coupler out	0.924	0.036	dB		Measured
Receiver loss total	3.35	–	dB		
Edge coupler in	0.924	0.036	dB		Measured
De-interleaver stages	0.35	–	dB	×2	Measured
Filter off-resonance	0.10	–	dB	×15	Measured
Filter on-resonance	0.225	0.047	dB		Measured
Crosstalk penalty	–	–	dB		>100 GHz channel spacing
Link loss total	18.41	–	dB		
+ Margin	1.253	–	dB		3σ link loss variation
+ Receiver sensitivity	–22.37	–	dBm		For 16 Gbps/channel
= Min. power per line	–2.71	–	dBm		

**Figure 8:** Monte Carlo simulations ($n = 200,000$) showing total link loss incorporating statistical variations of components from wafer-scale measurements. Mean and standard deviation link losses were 18.407 dB and 0.418 dB, respectively.

4 Conclusions

We have presented a wafer-scale, measurement-driven framework that links device-level statistics to system-level performance for scalable, co-packaged silicon photonic interconnects. Building on a scalable, comb-driven DWDM architecture that combines even-odd channel interleaving with multi-FSR channel allocation, we quantified the impact of fabrication variability across 300 mm wafers and

propagated these statistics through a realistic system model. The resulting methodology connects waveguide and resonator non-uniformities to thermal tuning power requirements, power penalties, and end-to-end link margins. Our wafer-scale extraction of effective and group indices revealed strong radial and planar trends, and demonstrated that wider waveguides substantially suppress statistical variation, thereby reducing required thermo-optic tuning power in phase shifters. Microdisk modulators and filters exhibited tight resonance distributions and low insertion losses at scale, while per-facet edge coupling losses remained sub-dB with minimal spread. Incorporating the measured depletion responses, we mapped modulation power penalties, and then executed a 200,000-trial Monte Carlo analysis of the full link. The resulting mean link loss (18.41 dB) and variation (0.418 dB) translate into clear source power targets when combined with receiver sensitivity and a 3σ margin. Importantly, design choices identified here—wider interleaver waveguides, whispering gallery mode resonators for both modulators and filters, and thermal undercut for improved heater efficiency—drive tuning power reductions to improve energy per bit. These results validate the proposed architectural path for dense, comb-driven DWDM. The combination of measured device distributions and system-level Monte Carlo yields a practical path forward for scaling channel count and aggregate bandwidth while maintaining low energy consumption for in-socket, co-packaged integrated photonics.

Acknowledgments: The wafer/chip fabrication and custom device processing were provided by AIM Photonics/SUNY Poly Photonics engineering team and fabricator in Albany, New York.

Research funding: 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 HR0011830002, in part by DARPA under Photonics in the Package for Extreme Scalability (PIPES) program contract number HR0011920014, 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.

Author contributions: RP, AJ, and YW conceived of the characterization methodology. Designs, mask layout, and testing were completed by RP, SW, MC, AJ, and YW. RP led wafer-scale measurements and analysis with assistance from AO, JR, and KJ. YW and KB supervised the project. All authors have accepted responsibility for the entire content of this manuscript and consented to its submission to the journal, reviewed all the results and approved the final version of the manuscript.

Conflict of interest: Authors state no conflict of interest.

Data availability: The datasets generated during and/or analyzed during the current study are available from the corresponding author on reasonable request.

References

- [1] Q. Cheng, M. Bahadori, M. Glick, S. Rumley, and K. Bergman, “Recent advances in optical technologies for data centers: A review,” *Optica*, vol. 5, no. 11, pp. 1354–1370, 2018.
- [2] D. Miller, “Rationale and challenges for optical interconnects to electronic chips,” *Proc. IEEE*, vol. 88, no. 6, pp. 728–749, 2000.
- [3] E. Strubell, A. Ganesh, and A. McCallum, “Energy and policy considerations for modern deep learning research,” *AAAI*, vol. 34, no. 9, pp. 13693–13696, 2020.
- [4] Z. Wu *et al.*, “Peta-scale embedded photonics architecture for distributed deep learning applications,” *J. Lightwave Technol.*, vol. 41, no. 12, pp. 3737–3749, 2023.
- [5] S. Xing *et al.*, “Seamless optical cloud computing across edge-metro network for generative ai,” *Nat. Commun.*, vol. 16, no. 1, p. 6097, 2025.
- [6] A. H. Atabaki *et al.*, “Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip,” *Nature*, vol. 556, no. 7701, pp. 349–354, 2018.
- [7] 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,” *Nat. Photonics*, vol. 4, no. 1, pp. 37–40, 2010.
- [8] K. Hosseini *et al.*, “8 tbps co-packaged fpga and silicon photonics optical io,” in *2021 Optical Fiber Communications Conference and Exhibition (OFC)*, IEEE, 2021, pp. 1–3.
- [9] 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)*, IEEE, 2022, pp. 260–261.
- [10] M. J. Hossain *et al.*, “Compact modeling of key passive silicon photonic components for a process design kit,” in *2023 IEEE Photonics Conference (IPC)*, 2023, pp. 1–2.
- [11] L. G. Carpenter *et al.*, “Towards low propagation losses in active photonic multi-project wafer runs,” in *Integrated Photonics Research, Silicon and Nanophotonics*, Optica Publishing Group, 2021, p. ITu3A-5.
- [12] W. Bogaerts *et al.*, “Silicon microring resonators,” *Laser Photonics Rev.*, vol. 6, no. 1, pp. 47–73, 2012.
- [13] C.-H. Chen *et al.*, “A comb laser-driven dwdm silicon photonic transmitter based on microring modulators,” *Opt. Express*, vol. 23, no. 16, pp. 21541–21548, 2015.
- [14] Y. Wang *et al.*, “Energy efficiency and yield optimization for optical interconnects via transceiver grouping,” *J. Lightwave Technol.*, vol. 39, no. 6, pp. 1567–1578, 2021.
- [15] R. Parsons, A. Novick, M. Hattink, A. James, A. Rizzo, and K. Bergman, “Efficient silicon photonic add-drop microdisk filters for dwdm systems,” in *CLEO 2023*, Optica Publishing Group, 2023, p. JW2A.127. [Online]. Available at: https://opg.optica.org/abstract.cfm?URI=CLEO_SI-2023-JW2A.127.
- [16] A. Rizzo *et al.*, “Massively scalable Kerr comb-driven silicon photonic link,” *Nat. Photonics*, vol. 17, no. 9, pp. 781–790, 2023.
- [17] A. Rizzo *et al.*, “Fabrication-robust silicon photonic devices in standard sub-micron silicon-on-insulator processes,” *Opt. Lett.*, vol. 48, no. 2, pp. 215–218, 2023.
- [18] K. Padmaraju and K. Bergman, “Resolving the thermal challenges for silicon microring resonator devices,” *Nanophotonics*, vol. 3, nos. 4–5, pp. 269–281, 2014.
- [19] A. Rizzo *et al.*, “Petabit-scale silicon photonic interconnects with integrated Kerr frequency combs,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 29, no. 1, pp. 1–20, 2022.
- [20] A. Novick *et al.*, “High-bandwidth density silicon photonic resonators for energy-efficient optical interconnects,” *Applied Physics Reviews*, vol. 10, no. 4, p. 041306, 2023.
- [21] Y. London *et al.*, “Energy efficiency analysis of comb source carrier-injection ring-based silicon photonic link,” *IEEE J. Sel. Top. Quantum Electron.*, vol. 26, no. 2, pp. 1–13, 2020.
- [22] Y. London *et al.*, “Performance requirements for terabit-class silicon photonic links based on cascaded microring resonators,” *J. Lightwave Technol.*, vol. 38, no. 13, pp. 3469–3477, 2020.
- [23] X. Zheng *et al.*, “A 33mW 100Gbps CMOS silicon photonic WDM transmitter using off-chip laser sources,” in *Optical Fiber Communication Conference/National Fiber Optic Engineers Conference 2013*, Anaheim, California, OSA, 2013, p. PDP5C.9.
- [24] C. Zhang, S. Zhang, J. D. Peters, and J. E. Bowers, “8 × 8 × 40 Gbps fully integrated silicon photonic network on chip,” *Optica*, vol. 3, no. 7, p. 785, 2016.
- [25] K. Padmaraju, X. Zhu, L. Chen, M. Lipson, and K. Bergman, “Intermodulation crosstalk characteristics of WDM silicon microring modulators,” *IEEE Photonics Technol. Lett.*, vol. 26, no. 14, pp. 1478–1481, 2014.
- [26] E. Timurdogan, C. M. Sorace-Agaskar, E. S. Hosseini, and M. R. Watts, “An interior-ridge silicon microring modulator,” *J. Lightwave Technol.*, vol. 31, no. 24, pp. 3907–3914, 2013.

- [27] M. Gehl *et al.*, “Operation of high-speed silicon photonic micro-disk modulators at cryogenic temperatures,” *Optica*, vol. 4, no. 3, p. 374, 2017.
- [28] H. Gevorgyan, A. Khilo, M. T. Wade, V. M. Stojanović, and M. A. Popović, “Miniature, highly sensitive MOSCAP ring modulators in co-optimized electronic-photonic CMOS,” *Photonics Res.*, vol. 10, no. 1, p. A1, 2022.
- [29] E. S. Magden *et al.*, “Transmissive silicon photonic dichroic filters with spectrally selective waveguides,” *Nat. Commun.*, vol. 9, no. 1, p. 3009, 2018.
- [30] A. Novick, S. Wang, A. Rizzo, R. Parsons, K. Jang, and K. Bergman, “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)*, Turin, Italy, IEEE, 2022, pp. 171–172. [Online]. Available at: <https://ieeexplore.ieee.org/document/9894746/>.
- [31] W. Shi *et al.*, “Ultra-compact, flat-top demultiplexer using anti-reflection contra-directional couplers for CWDM networks on silicon,” *Opt. Express*, vol. 21, no. 6, p. 6733, 2013.
- [32] J. Chen and Y. Shi, “Flat-top CWDM (de)multiplexers based on contra-directional couplers with subwavelength gratings,” *IEEE Photonics Technol. Lett.*, vol. 31, no. 24, pp. 2003–2006, 2019.
- [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. 3ThJ3.6.
- [34] L. Cao, A. Elshaari, A. Aboketaf, and S. Preble, “Adiabatic couplers in soi waveguides,” in *Conference on Lasers and Electro-Optics*, Optica Publishing Group, 2010, p. CThAA2.
- [35] R. M. de Ridder and C. G. H. Roeloffzen, “Interleavers,” in *Wavelength Filters in Fibre Optics*, vol. 123, W. T. Rhodes *et al.*, Eds., Berlin, Heidelberg, Springer Berlin Heidelberg, 2006, pp. 381–432.
- [36] L.-W. Luo *et al.*, “High bandwidth on-chip silicon photonic interleaver,” *Opt. Express*, vol. 18, no. 22, p. 23079, 2010.
- [37] F. Horst, W. M. Green, S. Assefa, S. M. Shank, Y. A. Vlasov, and B. J. Offrein, “Cascaded Mach-Zehnder wavelength filters in silicon photonics for low loss and flat pass-band WDM (de-)multiplexing,” *Opt. Express*, vol. 21, no. 10, p. 11652, 2013.
- [38] 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,” *Opt. Express*, vol. 29, no. 6, p. 7966, 2021.
- [39] A. Rizzo, Q. Cheng, S. Daudlin, and K. Bergman, “Ultra-broadband interleaver for extreme wavelength scaling in silicon photonic links,” *IEEE Photonics Technol. Lett.*, vol. 33, no. 1, pp. 55–58, 2021.
- [40] S. Wang *et al.*, “Automated tuning of a ring-assisted MZI-based interleaver for Kerr frequency combs,” *Opt. Lett.*, vol. 50, no. 2, p. 698, 2025.
- [41] E. Timurdogan, C. M. Sorace-Agaskar, J. Sun, E. Shah Hosseini, A. Biberman, and M. R. Watts, “An ultralow power athermal silicon modulator,” *Nat. Commun.*, vol. 5, no. 1, p. 4008, 2014.
- [42] V. Gopal *et al.*, “Simultaneous error-free data modulation with silicon microdisks in the multi-FSR regime for scalable DWDM links,” in *Optical Fiber Communication Conference (OFC) 2023*, San Diego, California, Optica Publishing Group, 2023, p. M4E.4.
- [43] 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*, vol. 12429, SPIE, 2023, pp. 271–288.
- [44] 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, 2023.
- [45] B. Y. Kim *et al.*, “Turn-key, high-efficiency Kerr comb source,” *Opt. Lett.*, vol. 44, no. 18, p. 4475, 2019.
- [46] A. Rovinski *et al.*, “Scaling co-packaged optical interconnects using hybrid 2.5D/3D integration,” in *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2025, pp. 1–5.
- [47] ficonTEC, “Photonics automated assembly & testing,” 2024. Available at: <https://www.ficontec.com/>.
- [48] A. James *et al.*, “Process variation-aware compact model of strip waveguides for photonic circuit simulation,” *J. Lightwave Technol.*, vol. 41, no. 9, pp. 1–14, 2023.
- [49] Y. Wang *et al.*, “Co-designed silicon photonics chip I/O for energy-efficient petascale connectivity,” *IEEE Trans. Compon. Packag. Manuf. Technol.*, vol. 15, no. 8, pp. 1581–1591, 2025.
- [50] R. Parsons *et al.*, “Highly uniform thermally undercut silicon photonic devices in a 300 mm cmos foundry process,” *Sci. Rep.*, vol. 15, no. 1, p. 29906, 2025.
- [51] P. Yin *et al.*, “Low connector-to-connector loss through silicon photonic chips using ultra-low loss splicing of smf-28 to high numerical aperture fibers,” *Opt. Express*, vol. 27, no. 17, pp. 24188–24193, 2019.
- [52] H. Huang, Z. Shi, G. Talli, M. Kuschnerov, R. Penty, and Q. Cheng, “Arbitrary mode size conversion with 3d-nanoprinted couplers: A generic coupling strategy,” in *2024 Optical Fiber Communications Conference and Exhibition (OFC)*, IEEE, 2024, pp. 1–3.
- [53] M. Hall *et al.*, “Innovations in fibre array coupling and integration for high-bandwidth fpga multichip packages,” in *Optical Interconnects and Packaging 2025*, vol. 13372, SPIE, 2025, pp. 21–27.