

PHOTONICS Research

High capacity, low power, short reach integrated silicon photonic interconnects

ANDREW NETHERTON,¹ MARIO DUMONT,¹ ZACHARY NELSON,¹ JAHYUN KOO,¹ JINESH JHONSA,¹ ALICE MO,¹ DAVID MCCARTHY,¹ SKYLAR DECKOFF-JONES,² YUN GAO,¹ NOAH PESTANA,² JORDAN GOLDSTEIN,² REN-JYE SHIUE,² CHRISTOPHER POULTON,² M. J. KENNEDY,¹ MARK HARRINGTON,¹ BOZHANG DONG,¹ JOCK BOVINGTON,³ MICHAEL FRANKEL,⁴ LUKE THEOGARAJAN,¹ MICHAEL WATTS,² DANIEL BLUMENTHAL,¹ AND JOHN E. BOWERS^{1,*}

¹Department of Electrical and Computer Engineering, University of California, Santa Barbara, Santa Barbara, California 93106, USA

²Analog Photonics, Boston, Massachusetts 02210, USA

³Cisco Systems, San Jose, California 95134, USA

⁴Ciena Corporation, Hanover, Maryland 21076, USA

*Corresponding author: bowers@ece.ucsb.edu

Received 29 January 2024; revised 2 April 2024; accepted 21 April 2024; posted 22 April 2024 (Doc. ID 520203); published 28 October 2024

The architecture and component technology of a low power, high capacity, short reach optical interconnect are detailed. Measurements from high-performance 300 mm silicon photonics components that comprise the system are shown, along with a quantum-dot mode-locked laser 20-channel comb source with free space wall plug efficiencies up to 17%, advanced packaging techniques for 3D silicon photonic-electronic integration, and schematics for integrated electronics that control the photonic integrated circuits. Techniques for operating such a system in the presence of changing ambient temperature are addressed. Experiments on a 1 Tbps design are conducted with an optical link experiment indicating sub-picojoule/bit energy consumption at scale. © 2024

Chinese Laser Press

<https://doi.org/10.1364/PRJ.520203>

1. INTRODUCTION

Optical interconnects are the preferred means of transmitting data at distances in excess of 1 m, but electrical interconnects are currently the better-suited choice for on-board distances and very short lengths [1,2]. Changing this paradigm requires the development of optical transceivers that are highly compact, scalable, energy efficient, with low latency, and thermally robust in terms of the underlying constituent components: integrated photonics, multi-wavelength laser source, control and driving electronics, and packaging.

Developing short reach optical interconnects faces some fundamental problems, but there are established strategies that introduce some tradeoffs. For instance, energy efficiency encourages dense wavelength division multiplexing (DWDM) with lower baud rate and simpler modulation format [3]; while it is possible to transmit more than 1 Tbps of capacity onto a single carrier using coherent modulation formats, experimental demonstrations have required more than 13 pJ/bit to achieve this feat [4]. This is in contrast with DWDM approaches that have claimed energy consumption rates below 1 pJ/bit [5]. Lower baud rates are preferred to an extent because current ASIC internal logic clock speeds are limited to approximately 3–5 GHz, and using a SERDES gearbox to convert from low

rate ASICs to higher symbol rate consumes power at each gearbox stage. This ultimately results in a maximum data rate to target before the SERDES power negatively impacts the energy per bit consumed. Additionally, 3D integration of electronics to photonics reduces parasitics between chips and improves energy efficiency.

With regards to a multi-wavelength data carrier for DWDM links, quantum-dot mode-locked lasers (QD-MLLs) offer high wall plug efficiency, a flat comb, a fixed channel separation, a simple electrical pumping scheme, and a compactness that cannot be matched simultaneously by alternative options, namely distributed feedback laser (DFB) arrays and nonlinear optical frequency combs [6–8]. Some comb sources may also need to be thermally stabilized or require optical isolation; this adds to power consumption and may prevent integration of the light source onto a photonic integrated circuit (PIC) [9–12].

Compactness encourages minimizing component size, placing components as close together as possible, and minimizing the number of components used. Maximizing the per-channel data rate within energy consumption limits aids compactness by reducing the number of transmit–receive unit cells needed in the system. Silicon photonics is a clear choice of platform for compact short reach transceiver PICs due to its many

features, including tight bends, high core confinement, phase shifters, modulators, photodetectors, and low-loss passive components [5,13,14]. Key components in these silicon PICs are wavelength-selective microring resonators used as modulators and WDM demux filters. Electronic unit cells for data encoding and reception must also be able to fit into a tight area constraint, and the 3D integration of the electronics and photonics must occur at relatively small flip chip pitches. Simple drivers with distributed logic for serialization are preferred over complex unit cells. Additionally, clocking and power distribution become challenging at tight pitches and need to be carefully managed.

Another aspect to consider is the port count required to transport a given requisite capacity. When the number of comb lines available from the laser source and the data rate per channel have reached their maximum limit, a port's capacity can be increased through multiplexing parallel data streams across fibers, polarization, or modes of the fiber; polarization multiplexing comes with the benefits of using more commonly available single-mode polarization-maintaining fibers with considerably fewer coupling losses compared to few mode fiber needed for mode division multiplexing with the tradeoff of multiplexing fewer parallel data streams on one port [13,15].

This work outlines our architecture for sub-picojoule/bit links. A description of each component in the system, both electronic and photonic, is provided. Experiments demonstrating ring resonator alignment with a controller IC and a single-channel link with QD-MLL comb source are shown. The results of the link experiment will be extrapolated to a fully integrated link for demonstrating the viability of reaching less than 0.4 pJ/bit. Finally, methods for making an even smaller, more energy efficient system will be discussed.

2. SYSTEM OVERVIEW

Figure 1 illustrates the photonic architecture of the proposed transceiver. A 20-wavelength O-band QD-MLL comb source with a channel separation of 60 GHz has its outputs coupled to two separate 1 Tbps silicon PICs; the 60 GHz channel separation chosen is set by the symbol rate to prevent crosstalk. The input comb power coupled to each PIC is split across two separate paths equally and is then deinterleaved into "odd" and "even" combs with 120 GHz channel separation that are then fed into micro-ring modulator (MRM) banks for data encoding. Deinterleaving is done both to reduce off-resonance losses by having fewer rings in series on the bus and to reduce inter-modulation crosstalk by increasing the spectral separation between channels [16,17]. Each MRM operates at 26.4 Gbps NRZ with a 5.6% forward error correction (FEC) overhead. A discussion on the case for FEC in compact, energy efficient links and its impact on link latency can be found in the following subsection. This modulation signal is provided by an electronic driver chip, and the ring resonances are stabilized by a controller IC. NRZ is the preferred modulation format over alternatives such as PAM-4 because PAM formats are disadvantageous from an energy perspective; their optical power penalty accumulates much faster than capacity gain. Coherent modulation formats are also possible to encode with ring modulators, but the receiver is too complex for high energy efficiency. Additionally, phase shift keying with ring modulators typically has higher excess loss than amplitude modulation for the same ring modulator quality factor [18].

The modulated odd and even comb lines are interleaved together, and a shallow 209 MHz clock signal is encoded onto all channels for optical clock transmission via a variable optical




Fig. 1. Diagram illustrating the photonic transceiver architecture of the system.

attenuator (VOA). Optical clock transmission allows for simpler clock recovery circuitry and can reduce system power [19]. Prior proposals for optical clock transmission in similar links have allocated one of the optical channels for the data clock [20,21]. A downside of this approach is the need for a dedicated ring site on the TX and RX that must align to one another even when the TX and RX are at different temperatures and experience different fabrication errors; this conflicts with the goal of minimizing power consumption introduced by DC tuning of rings.

The net-500G lanes then pass through off-chip quantum-dot semiconductor optical amplifiers (QD-SOAs) before being polarization multiplexed onto a single 1T port for propagation over polarization-maintaining (PM) fiber; since the propagation distance is short, input power is low, and the carrier frequencies are in the O-band (but slightly shifted from the zero-dispersion wavelength), signal distortions due to chromatic dispersion and optical nonlinearity in fiber are negligible. It is optimal to place the optical amplification at this point in the link because it is energetically preferable to place it at a spot where optical power is already low; i.e., given the desired target optical gain it will minimize optical output power, which is proportional to electrical power consumption.

The receiver (RX) portion of the PIC first polarization demultiplexes the two 500G lanes; a small fraction of the power of each channel is sent to a photodiode (PD) that detects the clock signal for synchronization on the RX side of the driver IC. The rest of the power is deinterleaved, and then each individual 25G channel is demultiplexed by a ring resonator add-drop filter with a high-speed PD receiver on its drop port; the electrical output of the PD is connected to a receiver cell input on the driver IC.

As shown in the architecture, a complete optical link is composed of several photonic components that can potentially be at very different temperatures and come from different wafer fabrication lots. It is also noted that the laser, TX-side PIC, and RX-side PIC can all be at physically different locations and may experience uncorrelated temperatures. However, components on the same PIC typically experience both highly correlated temperatures and correlated manufacturing deviations. Individual channels coming from a laser have to be aligned to the transmit side deinterleavers, modulators, and interleavers, as well as to receiver side deinterleavers and channel drop filters. These have to be maintained over a specified operating temperature range, which is specified for commercial grade optics as 0°C–70°C or a more relaxed 15°C–70°C [22]. Further frequency adjustment allocations must be made available for fabrication-induced perturbations of initial component set points at a fixed temperature so that all the components in the system can be appropriately aligned to the spectral grid.

The integrated phase shifter commonly used for component alignment in silicon photonics is the heater, and it is recognized that reducing static link heater power consumption is a primary focus of increasing link efficiency [3]. There are ways to reduce static heater power consumption, such as substrate removal and optimized heater placement [23,24]. Substrate removal may make the PIC more delicate and add complexity to packaging. Furthermore, it has recently been shown that the 3D

integration of electronics onto silicon photonics with a flip chip pitch of 50 μm degrades thermal tuning efficiency by 43.4% and increases the thermal crosstalk by 44.4% [25]. These thermal tuning performance metrics are expected to further degrade with tighter flip chip pitches in the aforementioned study. Another approach is index trimming by locally amorphizing and annealing the waveguide [26]; this allows for correction of fabrication alignment errors, but not for thermal drifts. Athermalization of elements such as the ring modulator using techniques such as negative thermo-optic coefficient cladding is not worth the incurred penalty to performance from reduced core confinement [27]. Overall, the most effective way to eliminate heater power would be to eliminate the need for heaters, but phase shifters of some sort are still needed.

The lasers selected for our system also have three temperature-related effects that need to be considered. First, optical efficiency generally decreases with increasing temperature, impairing overall system efficiency. Second, optical gain shifts at a rate of roughly 80 GHz/°C for the III/V material system. Third, optical channel frequencies shift at a rate of 18 GHz/°C due to III/V cavity optical length changes.

Given the above considerations, we performed a detailed system energy efficiency analysis of various combinations for system thermal stabilization (including manufacturing offsets) and have selected the following combination as optimal for our system design. We implement athermal deinterleavers/interleavers, ensuring that the TX and RX sides are well aligned. These serve as a thermally invariant frequency reference for the optical link. The laser is designed to operate close to the system thermal range mid-point, and a thermo-electric element stabilizes temperature by either cooling or heating, as needed. This provides a good tradeoff between minimizing TEC power consumption and not penalizing laser efficiency at higher temperatures. Finally, ring modulators and demultiplexers are stabilized using the free carrier plasma dispersion effect (FCPDE, colloquially referred to as electro-optic or EO throughout this work), rather than with thermal heating. The tradeoffs and design considerations are described in more detail in the following section.

A. Forward Error Correction

FEC is a powerful technique that allows communication channels to get closer to the limit established by the theory of communication over a noisy channel. For short reach links, FEC would be considered feasible and beneficial to implement if it consumes less electrical power than an equivalent increase in optical signal power, if FEC encoding and decoding fit within system latency budget specification, and if the FEC circuit fits within the area of the driver IC.

Referring to Eq. (1), the signal to noise ratio (SNR) is proportional to power squared and inversely proportional to bandwidth. When thermal noise dominates, as is the case in nearly all practical communication links without optical amplification, this equation reduces to

$$\text{SNR} \approx \frac{R_i P_s^2 \rho^2}{2k_B T F_n B}. \quad (1)$$

In the above equation, R_i is the load resistance, P_s is the input power, ρ is the responsivity, k_B is the Boltzmann

constant, T is the temperature, F_n is the noise figure, and B is the baud rate. A $\sqrt{10} = 3.16\times$ increase in power produces $10\times$ increase in capacity. In general, short thermally limited links are best operated as fast as possible within the bandwidth limits of available electro-optic components and available laser power. Parallelizing is not helpful as, for example, $10 \times 10\text{G}$ links are $3.16\times$ less efficient than a single 100G link. Further, since these are typically single-span links, the received signal power P_s is inversely proportional to link loss, L : $\text{SNR} \approx 1/BL^2$. Net effective coding gain (NECG), which is the SNR reduction allowed by introducing FEC, is proportional to

$$\text{NECG}_{\text{dB}} \approx B_{\text{dB}} + 2L_{\text{dB}}. \quad (2)$$

Thus, NECG can be spent on data rate increase or on link budget but in different proportions. For example, an FEC code with $\text{NECG}_{\text{dB}} = 6$ dB can be used to expand data rate by 6 dB for an optical-power limited channel, i.e., a factor of $4\times$. However, the same code can only extend link loss by 3 dB:

$$\begin{aligned} \text{NECG}_{\text{dB}} = & 20 \log_{10}(\text{erfc}^{-1}(\text{BER}_{\text{ref}})) \\ & - 20 \log_{10}(\text{erfc}^{-1}(\text{BER}_{\text{in}})) + 10 \log_{10}(R_C). \end{aligned} \quad (3)$$

Practical optical links may deviate from the idealized computation presented above. Specifically, channel distortions induced by imperfect electro-optic components, finite crosstalk, etc., frequently make FEC more helpful than what this computation suggests; i.e., FEC allows us to clean up error floors or flaring that a simple increase in optical power is unable to overcome. For example, Fig. 2(a) shows a typical BER curve for a highly optimized laboratory instrument built using discrete components and operating as a single channel, single polarization, and without optical filtering—note 2.2 dB penalty from BER 10^{-5} and 10^{-12} [28]. In contrast, highly parallelized optical systems are susceptible to bandwidth narrowing from EO components and optical filtering, to linear crosstalk between neighbors, polarizations, modes, etc. Such links have exhibited optical power penalty in excess of 18 dB at 10^{-12} BER, and 10 dB at 10^{-5} BER with $16 \text{ Gbps}/\lambda$, shown in Fig. 2(d) of Ref. [5].

This effect can also be shown in link performance simulations, when various bandwidth and crosstalk effects are properly accounted for, shown in Fig. 2(b). FEC performance gain

on an idealized link is indeed fairly small (2.4 dB less input optical power needed for target BER), while a realistic optical link benefits from FEC more significantly (3.5 dB). FEC provides even more performance gain for an APD receiver (6.0 dB), and the APD requires 8.9 dB less input optical power than the p-i-n receiver at the FEC threshold. For clarity, FEC performance gain can be estimated from Fig. 2(b) by computing the ratio between a link's unamplified photocurrent "1" level at the red–yellow and green–yellow region boundaries.

We constrain ourselves to a maximum latency (excluding time of flight) of 50 ns, and we target FEC NECG ≈ 4 dB at corrected $\text{BER}_{\text{ref}} \approx 10^{-12}$, which indicates a pre-corrected $\text{BER}_{\text{in}} \approx 10^{-5}$. Such a stringent upper bound on latency opens up the possibility for using these interconnects in memory disaggregation applications [29,30]. Preliminary estimates of optical link QD SOA marginal power efficiency are 0.06 pJ/bit at 2 dB optical gain (i.e., $\text{NECG}/2$). Therefore, FEC is likely to be beneficial if its power consumption is $<0.06 \text{ pJ/bit}$.

A relatively simple BCH $N = 511$, $K = 484$ code with code rate $R_c \approx 0.947$ and NECG ≈ 4.25 dB was selected for this link. The code has been fully implemented in MATLAB and tested on various data and error sequences. This code is capable of completely correcting up to three errors per 511-bit block. If there are four or more errors per block, they are passed through without correction. Infrequent additions of three artificial errors are also possible when there are four or more real errors, but this effect is observed to have a negligible impact on performance, as shown in Fig. 2(c).

In general, robust optical link operation in practical deployments mandates the use of FEC, which not only provides reduction in total system power consumption, but also provides link immunity against dribbling errors that may be induced by occasional control loop errors, RF channel reflections, crosstalk fluctuations, etc.

FEC encoding blocks require interleaving to provide tolerance against error bursts longer than 3 bits. However, care must be taken about how interleaving is implemented as it increases the system memory depth and corresponding latency associated with FEC. FEC coding on a single-wavelength tributary is not supportable due to the low latency requirements. Figure 3 shows a single wavelength operating at 26.4 Gbps and a 511-bit block occupying 19.4 ns . Interleaving 20 blocks




Fig. 2. Single-channel, discrete component 10G link performance. (a) Simulated link performance comparison of idealized (blue) and realistic (red) optical links. (b) The yellow region is where the BCH(511, 484) FEC code can be used to correct for all errors, the green region is FEC-free, and the red region exhibits too many errors to be adequately corrected. (c) Simulated FEC performance with simple algebraic decoder compared with binomial approximation.




Fig. 3. FEC latency with 20x interleaving of 511-bit blocks streamed through a single-wavelength channel operating at 26.4 Gbps.

increases FEC delay to 387 ns, which is considerably higher than the given latency constraint.

The system's 20 wavelengths are clock-synchronous, and FEC blocks can be spread across all of these. In addition, robustness is increased further by adding 26-way time-domain interleaving. This method is shown in Fig. 4, and shows an interleaving latency of 26 ns.

Total FEC latency can be estimated assuming an implementation in Taiwan Semiconductor Manufacturing Corporation (TSMC) 7 nm CMOS process. The XOR gate count is $511 \text{ bits} \times 9 \text{ bits/syndrome} \times 3 \text{ syndromes} = 13,797$ gates. FEC decoder delay involves computing the third and fifth power, square root, and cubic root of elements in GlobalFoundries (GF) via a 512 element look-up table (LUT). Further, addition via XOR of 9-bit sequences, the product of two GF elements via LUTs, and division via 512 element LUTs are also performed. Thus, total decoder delay is due to the encoder (2 ns), interleaving (26 ns), and decoder (5 ns), 33 ns in total, which is well within a latency specification of 50 ns.

This system has another specific benefit from FEC with regards to latency. EO ring tuning is restricted in range such that occasional channel switching is required over the full temperature range ramp. At the same time, EO tuning can be quite fast, with tuning within 1 ns quite achievable; thermo-optic tuning, on the other hand, is orders of magnitude slower. This switching speed corresponds to 26.4 bits at 26.4 Gbps NRZ. With reasonable block interleaving, at most one to two errors per FEC block are expected, which are within the FEC correction range. Thus, with careful channel switching algorithm design, error-free operation is expected.

In terms of consumed CMOS die area, a receiver area associated with the PIC is determined by the areal bandwidth density, which is in turn limited by the pad pitch. The system pad pitch in the high-speed region of $36 \mu\text{m}$ leads to a footprint of $4700 \mu\text{m}^2$ for a 500 Gbps RX composite. The area is largely




Fig. 4. FEC latency with 20x interleaving across 20 wavelengths and 26-way time-domain interleaving of 511-bit blocks.

dominated by XOR gates, and the area consumed can be estimated using published results as $0.22 \mu\text{m}^2 \times 13,797 \text{ XORs} = 3000 \mu\text{m}^2$, with some additional area for LUTs and some logic [31]. Therefore, FEC should be able to fit into the allocated area with area to spare, assuming an advanced 7 nm CMOS node.

Similar estimates of FEC power consumption indicate that $\approx 13 \text{ fJ/bit}$ will be consumed by XOR operations, with some additional power for LUTs and some logic, which fits comfortably within the FEC-allocated 20 fJ/bit system power budget.

In summary, FEC should be considered mandatory for highly efficient and operationally resilient parallelized optical links. Simple FEC codes provide sufficient gain, mitigate the effect of many impairments, and can be realized in advanced CMOS nodes meeting specifications for energy efficiency, latency, and areal density in short reach applications.

3. COMPONENTS

This section provides descriptions of the design and performance of the constituent photonic, electronic, and packaging elements that comprise the system in isolation of one another.

A. Quantum-Dot Mode-Locked Laser Comb

The QD-MLLs used in this system consist of a Fabry–Perot laser structure with a saturable absorber (SA) made from the same material as the gain medium. The SA is formed by electrically isolating part of the laser ridge from the gain section by etching through the p-contact and part of the p-cladding layers in the region between the gain and SA regions. By using this laser design, it eliminates the costly grating patterning and regrowth steps involved in commercial DFB fabrication. The saturable absorber is placed precisely in the middle of the cavity to use the colliding pulse effect, which increases the comb channel spacing from 30 to 60 GHz in our case. The performance of the MLL can be tuned by changing the reverse bias voltage applied. This is considerably effective in tuning the bandwidth of the comb to match the exact requirements of the system. The details of the device fabrication and performance are provided in Ref. [32].

With extremely stringent power consumption requirements of short reach interconnects, the wall plug efficiency (WPE) of the laser is extremely important. Figure 5 shows the light-current and WPE curves for a 60 GHz QD-MLL, which reaches a peak WPE of 17% in free space. However, the bandwidth of the comb also affects the efficiency of generating data channels; having a too narrow or too broad bandwidth

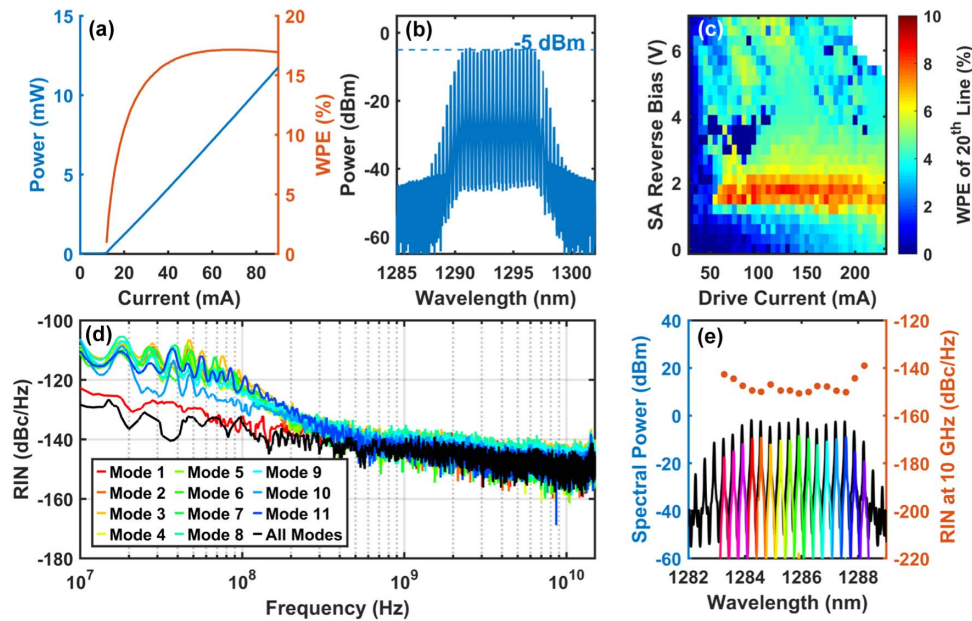


Fig. 5. (a) Light-current output and wall plug efficiency of a QD-MLL. (b) Optical spectrum and (c) wall plug efficiency of generating the 20th comb line assuming 20 comb lines are used for data transmission. (d) Comparison of the measured RIN for different comb lines versus the entire spectrum. (e) Comb state and RIN at 10 GHz for different comb lines in a QD-MLL spectrum; colored spectra correspond to bandpassed individual comb lines measured in (d).

produces channels with low power and efficiency. To track this, the 20 contiguous comb lines with the highest power are identified from a spectrum [shown in Fig. 5(b)], and the comb line with the lowest power is divided by one-twentieth of the electrical power consumption of the laser. This calculates the WPE of the worst comb line that will be used as a data channel, referred to as WPE₂₀. These devices have been verified to emit equal power from both facets so the output collected by one facet is multiplied by a factor of two in calculating WPE₂₀. Figure 5(c) shows the WPE₂₀ of a device as a function of drive current and SA reverse bias to identify the most efficient comb state.

For IM-DD links, the relative intensity noise (RIN) of the laser is important because high RIN levels degrade the overall link performance. The RIN characteristics of an FM comb were investigated and are shown in Fig. 5(e). Each comb line was filtered with a tunable OBPF to evaluate the RIN of a single line versus the entire comb spectrum. In this study, the measured RIN accounts for the shot noise and the thermal noise without average [33], which is dependent on the output power. Considering the additional loss from the optical filter, the maximum fiber-coupled power of an individual comb line is limited to -9 dBm. As such, we measured the RIN of the whole spectrum at a power level of -9 dBm for a fair comparison. The entire comb spectrum has a similar RIN value to each individual line from 0.1 to 18 GHz, which is as low as -151 dBc/Hz at the frequency offset of 10 GHz. It should be noted that the RIN measurement is limited by the thermal noise of the photodetector when the fiber-coupled power is below 3 dBm. For our device studied, the RIN at 10 GHz of the whole spectrum further decreases to -165 dBc/Hz once the measurement is no longer limited by the thermal noise. In Fig. 5(d), the RIN

of each comb line is plotted to show that there is no change across the entire spectrum.

B. Silicon Photonics

Figure 6 depicts a fully processed 300 mm silicon photonics wafer and an image of an individual PIC. A number of silicon photonic components needed to be designed and working in tandem for the system to be operational. This includes athermal interleavers, high-speed modulators, variable optical attenuators, low-loss edge couplers, RX demux rings, and finally photodetectors. The performance of the athermal deinterleaver used in this system, as well as techniques to optimize it for enhanced fabrication tolerance, is discussed in detail in Refs. [34,35]. The other components will be evaluated in the following subsections.

1. Edge Couplers

The inverse taper edge coupler in this system is optimized for coupling to standard SMF-28 fiber, which has a mode field diameter of ~ 10 μm . In order to support polarization multiplexed links, the edge coupler is designed for both TE and TM polarizations. The buried oxide (BOX) layer between the device layer and the silicon substrate that silicon foundries typically use is 2 μm thick [36], which is thin enough to increase loss of edge couplers that are designed to match SMF-28 fiber modes due to substrate coupling. Substrate etch and oxide deposition are performed near the start of the fabrication process in order to effectively increase the BOX thickness, as shown in the cartoon and micrograph in Fig. 7(a). Figure 7(b) plots measurements of excess loss of the edge coupler for each mode when coupling to SMF-28 fiber. A uniform broadband insertion loss of ~ 1 dB was measured for each mode.




Fig. 6. (a) 300 mm wafer of silicon photonic transceivers. (b) Micrograph of a silicon photonic 1 Tbps transceiver PIC.

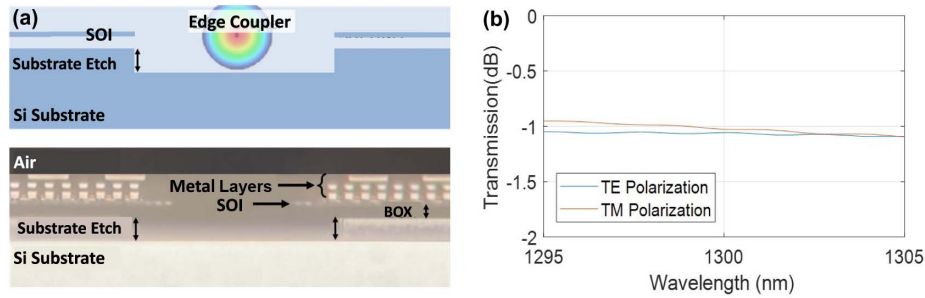


Fig. 7. (a) Low-loss edge coupler concept and micrograph with silicon substrate etch. (b) Measured transmission of edge coupler to SMF-28 fiber.

2. Polarization Beam Splitter and Rotator

An integrated polarization beam splitter and rotator is required to polarization demultiplex the two 500G lanes at the receiver. Figure 8(a) shows that the insertion losses for TE and TM input light are about 0.06 and 0.40 dB at 1300 nm, respectively, for the fabricated device. Additionally, measured polarization crosstalk results are shown in Fig. 8(b). The peak crosstalk magnitude for the TM input to TE output path is approximately -23.5 dB, which is likely sufficient to close the link. The crosstalk on the TE input to TE output, ranging between -13.0 and -17.3 dBm, may be significant enough to impair signal integrity.

3. Ring Resonators

There are two types of ring resonators: the TX-side modulator that encodes the high-speed optical waveform and the RX-side filter that wavelength demultiplexes each channel. Figure 9(a) shows an example transmission response of the ring modulators. These devices have a loaded Q of ~ 4600 and a free spectral range (FSR) of 1.2 THz. As shown in Fig. 9(b), power buildup inside the cavity introduces nonlinear losses and self-heating that limits the incident power on the ring near -5 dBm before the resonance shape distorts and red shifts [37]; this effectively limits the amount of useful power the QD-MLL can provide for each carrier without splitting the

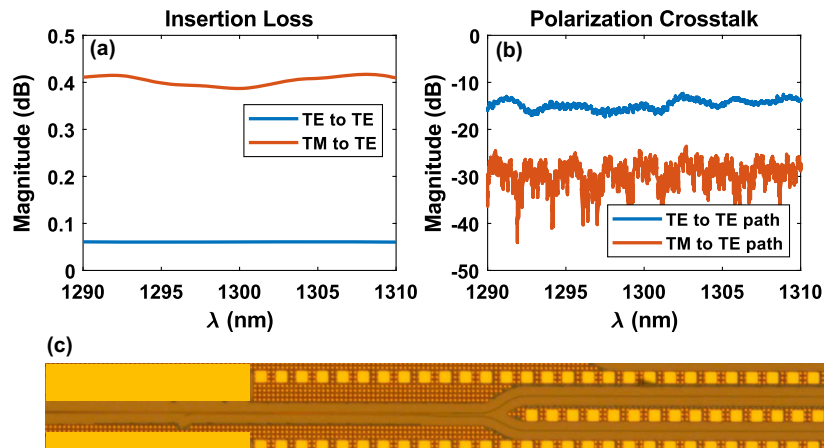


Fig. 8. (a) Insertion losses of the silicon photonic polarization beam splitter and rotator. (b) Measurement of polarization crosstalk for each path when the undesired input polarization is excited. (c) Micrograph of PSR test structure.

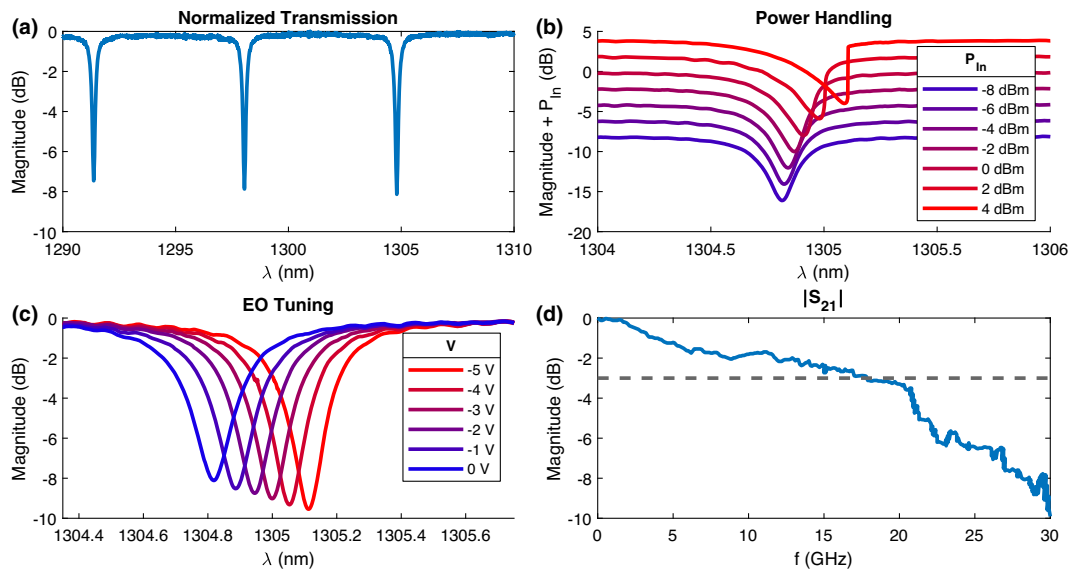


Fig. 9. (a) Transmission spectrum of a ring modulator. (b) Transmission spectrum as input power to ring is increased. (c) Transmission as integrated phase shifter is reverse biased. (d) Measured electro-optic S_{21} magnitude with the carrier +32 GHz detuned from resonance.

comb power over multiple PICs. To encode the data onto a carrier, the resonance is red shifted by reverse biasing an integrated p-n junction phase shifter within the cavity to deplete charge; Fig. 9(c) shows that the resonance can be shifted by 52 GHz over a 5 V span before exhibiting the onset of breakdown characteristics. This large DC tuning can also be used to correct for low-speed wavelength misalignment from temperature drifts and fabrication deviations. Figure 9(d) plots the

small signal electro-optic transfer function of the device. The 3 dB small signal bandwidth is about 18 GHz, which is sufficient for the target data rate.

These rings are designed as a series with varied optical path lengths to introduce a passively fabricated channel separation between neighboring resonances [38]. This is seen in Fig. 10(a) where the resonance location of different ring channels is plotted and one line is a single reticle. While the absolute

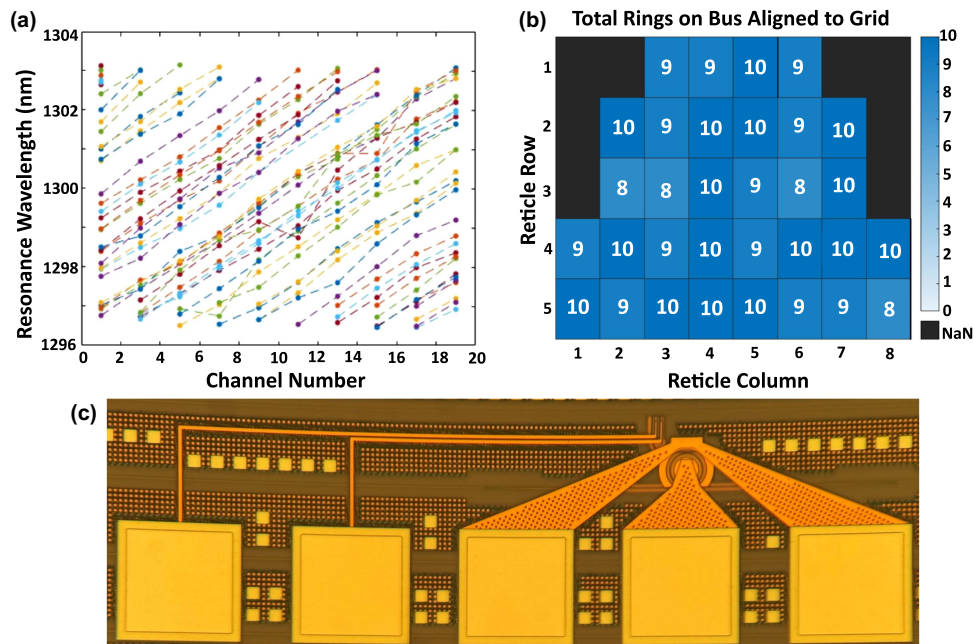


Fig. 10. (a) Resonance location of different ring odd-deinterleaved channels on a bus. One line of 10 points is a single reticle's bus of rings, showing strong correlation between different channels. The line may be broken into two segments if the Channel 1 resonance did not occur near 1297 nm. (b) Wafer map of the number of rings that can be tuned to the correct channel on each reticle. The maximum value is 10. (c) Micrograph of a ring modulator test structure.

resonance wavelength of a particular design varies between reticles as waveguide dimensions fluctuate across the wafer, the majority of the lines follow a monotonic trend with the expected spacing. Given the original resonance location of the ring, a large number of reticles support the tuning of all rings to their correct location even without the ability to tune a full FSR. This is shown in Fig. 10(b) where all reticles on half a 300 mm wafer can support the operation of at least 8 out of the 10 rings on the odd-deinterleaved bus at a fixed temperature without the need for integrated heaters that would consume more power, and half of all reticles can support operation on all 10 channels. State-of-the-art substrate undercut ring modulators are currently budgeting 336 fJ/bit for TX modulator DC tuning for both fabrication errors and thermal drifts [39]; this excludes dynamic power consumption due to modulation [40] and the static power consumption of rings on the RX side of the PIC. This EO DC tuning result was performed with a p-n junction with a capacitance of 50 fF that will be driven with a high-speed waveform with $1.2V_{pp}$, leading to fabrication errors being compensated for with no static power consumption and an estimated 18 fJ/bit of dynamic power consumption in the ideal case. However, to use EO tuning to also handle temperature drifts would require roughly tripling the FCPDE magnitude (which would in turn increase the capacitance of the ring modulator) and introducing a barrel shifter channel addressing circuit [19,41,42]. The implementation of a trimming mechanism would allow for only needing to double the tuning strength because only temperature drifts would need to be handled dynamically [26]. Factoring that an EO tuned RX ring consumes little dynamic power, it is apparent that realizing rings that do not require static thermal tuning has the potential to save hundreds of fJ/bit in a link budget even with the increased capacitance and dynamic power consumption that is introduced.

While the EO tuning electrical circuit consumes less power than the thermo-optic solution, introducing the integrated p-n junction phase shifter into the ring introduces more round-trip losses and degrades cavity Q . While commonly understood as part of the modulation bandwidth versus optical modulation amplitude design tradeoff for ring modulators, this also impacts the design of RX demux filters [43]. This Q spoiling is not an insurmountable issue for the RX ring since the Q factor required to adequately demultiplex a 26.4 Gbps channel is a similarly low value to that required by the modulator [44]. However, if the filter bandwidth is too large, more inter-channel crosstalk can be picked up by the filter; this can be corrected for by reducing the coupling coefficients of the rings to target at the expense of less demultiplexed power received at the drop port of the ring. Thus, there are tradeoffs between tuning range, filter drop bandwidth, induced optical loss, and crosstalk. The round-trip losses of an RX filter can be much less than those of a modulator since the time constants associated with thermal drifts are much slower than the target channel rate; this allows for changing the integrated phase shifter design by pulling high concentration implants away from the core at the expense of increased series resistance and a slower RC time constant. Since the same amount of charge is still depleted at the junction, these devices should have the same frequency tuning

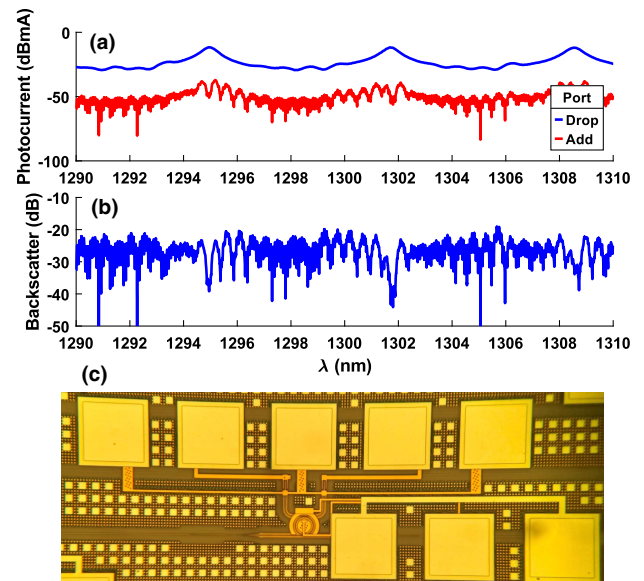


Fig. 11. (a) Measured spectral response of RX ring drop and add ports. (b) Backscatter spectrum estimate from the drop and add responses. (c) Micrograph of RX ring test structure for backscatter measurement.

performance as the TX modulators. Additionally, forming the demux filter out of multiple cascaded rings can produce a flatter passband, but doing so increases the complexity of the locking circuitry, electronic and photonic footprint, and drop port losses [45]; therefore, a design featuring a single ring to form each filter was selected.

Figure 11(a) shows the drop spectrum of a measured demux filter with 0 V applied bias to the integrated phase shifter. A full width at half maximum of approximately 99.4 GHz is observed. Additionally, the spectral response at the add port was measured, which allows for extracting a backscatter spectrum. This result is shown in Fig. 11(b) and suggests crosstalk magnitudes lower than -40 dB near resonance peaks [46]. It is possible that this low backscatter could allow for two parallel data streams with identical wavelength allocations to be demultiplexed by the same ring with minimal crosstalk in the two counterpropagating directions of the cavity; this would permit halving the number of RX rings needed in the PIC, reducing complexity, footprint, and energy consumption.

4. Variable Optical Attenuator

The VOA is formed with a forward biased p-i-n junction in a silicon ridge waveguide; using a p-i-n junction instead of a p-n junction reduces excess optical loss in the “off” state. In order to attenuate the light (i.e., modulate the clock), it is slightly forward biased at a ~ 200 MHz rate. Only a small amount of modulation is required, on the order of 1 dB, for the RX electronics to extract and divide the clock on each channel. Figure 12(a) shows the measured I - V curve and insertion loss (attenuation) of the VOA. Since it is a forward biased silicon diode, it has a turn on voltage of ~ 0.7 V. There is a ~ 0.3 dB excess loss at 0 V due to the optical mode interaction with the p-i-n junction. In order to achieve the required 1 dB modulation, the VOA needs to swing from 0 to 0.8 V

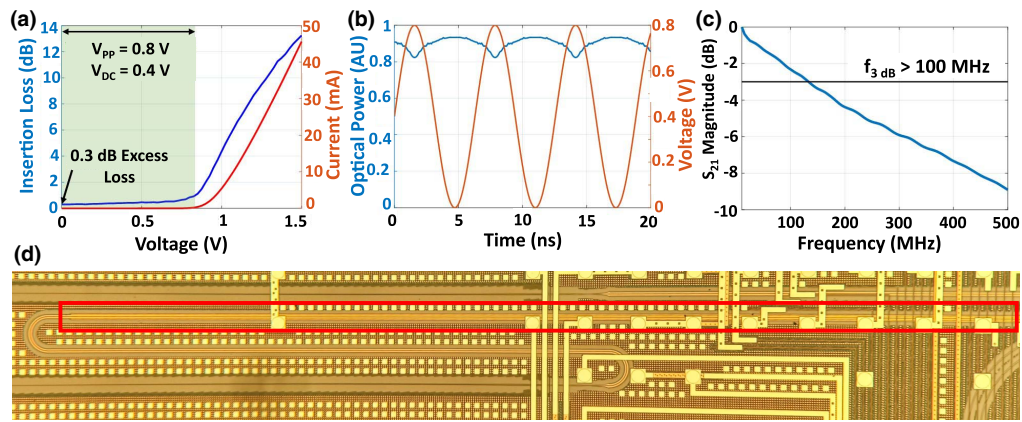


Fig. 12. (a) Measured I - V curve and optical attenuation of the variable optical attenuator. (b) Measurement of clock signal using the VOA from 0 V to 0.8 V. (c) EO bandwidth measurement of the VOA. (d) Micrograph of 1 Tbps PIC with segment of VOA shown in red box (d).

($V_{pp} = 0.8$ V and $V_{DC} = 0.4$ V). At 0.8 V, the VOA draws less than 1 mA.

Figure 12(b) shows time domain measurements of the VOA output when a voltage input with 200 MHz sinusoid shape of 0–0.8 V is applied. One can see the modulation on the optical output that enables the extraction of a clock later in the receive chain. Figure 12(c) shows the measured EO bandwidth of the VOA is ~ 125 MHz.

5. Photodetector

The waveguide coupled, linear-mode vertical p-i-n junction photodetector used in the receiver is based on epitaxially grown germanium on silicon [47]. Figure 13(a) shows the measured responsivity and bandwidth of the photodiode. The bandwidth was measured as a function of reverse bias, and a bandwidth of ~ 25 GHz is measured at 1 V reverse bias. A responsivity of 0.8 A/W is measured across the O-band. Figure 13(b) shows the dark current of each PD on a reticle across the wafer at a 1 V reverse bias. The average dark current is ~ 8 nA, which is more than sufficient for low-noise receiver performance.

C. Electronics

The electronics design consists of two electronic ICs (EICs): the driver and controller. Each chip is produced in 65 nm CMOS. The TX on the driver EIC uses on-chip PRBS sources to generate the data traffic, and the RX side of the driver converts the

input PD current into a digital signal. The drift of the ring modulator resonance wavelength is corrected using the control EIC. To drive a single PIC, two control chips and one driver chip are required. Each control chip drives the rings of the TX and RX channels, respectively, while the driver chip operates all 40 channels of data transmission in the PIC.

1. Control IC

The control IC is specifically designed to drive the ring resonator modulators and demux filters within the PIC. The resonance frequency of each channel in the PIC is determined by the DC bias voltage applied to the ring's integrated phase shifter; this, in turn, affects the magnitude of current measured by the monitor photodiode on each ring's drop port for a fixed laser power and emission wavelength. Consequently, by accurately tracking the photodiode current, the optimal bias voltage for efficient operation can be determined. The monitor photocurrent is digitized using a VCO-based ADC and compared to a set value. The error signal then passes through a digital loop filter and is converted into a $\Sigma\Delta$ bit stream. This bit stream is converted to a high voltage signal (0–6 V) via a level shifter and filtered to yield the bias voltage for the ring.

The control IC chip features on-chip regulators to reduce power supply noise, phase locked loops (PLLs), 40 channel controllers, and an SPI interface to control the operation using an external FPGA/microcontroller. The total chip size is 2 mm \times 1.8 mm.




Fig. 13. (a) Bandwidth measurement of the receiver photodetector. Inset shows the responsivity. (b) Wafer-scale measurement of receiver dark current at 1 V reverse bias in nA. (c) Micrograph of photodiode test structure.

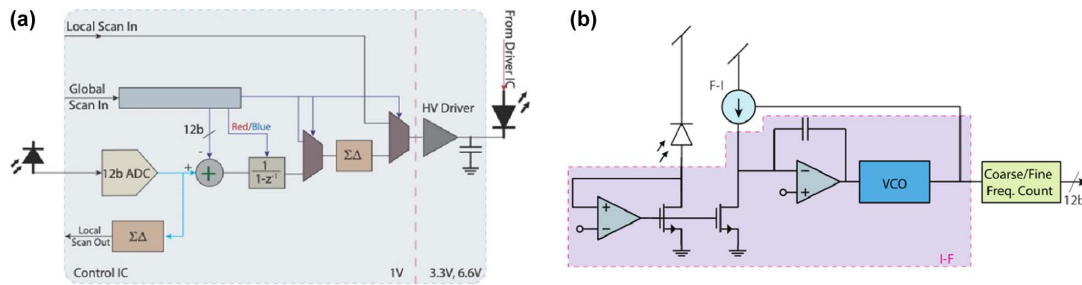


Fig. 14. (a) Schematic of control IC chip. (b) Schematic of the photocurrent sensing ADC.

Figure 14(a) shows the schematic of a unit cell with PLL and low-dropout (LDO) regulator. Each controller has a 12-bit current sensing ADC, red/blue locking digital loop filter, and high voltage EO driver.

Figure 14(b) shows the schematic of the photocurrent sensing ADC. The system comprises an analog front end designed to get current input, a loop filter, a voltage-controlled oscillator (VCO), and a switched capacitor-based resistor. The switched capacitor-based resistor's effective resistance is determined by the VCO's frequency. With the help of a feedback loop in the control IC, the VCO frequency is dynamically adjusted to align with the input current, thereby achieving a balance between the input and resistor currents. The frequency is solely influenced by the input current through feedback operation, while a course/fine frequency counter converts this frequency into a digital code with a resolution of 12 bits.

2. Driver IC

Figure 15 shows the overall architecture of the driver IC. There are two PLL blocks for TX and RX, on-chip LDOs, and 10 TX and RX blocks. The overall operation is controlled by a serial peripheral interface (SPI) scan chain.

Each TX block consists of four serializers, a pattern generation block, and an injection-locked eight-phase oscillator. It drives four channels and generates various patterns including PRBS signal ($2^{31} - 1$). The main PLL generates a 3 GHz clock signal, and it applies an injection signal to the replica oscillator for injection locking. From the injection locked replica oscillator, eight phase clocks serialize the PRBS data and generate a 26.4 Gbps data stream for the EO driver. The 209 MHz forwarded clock signal is generated by the driver via a 16× down-sampling of the main PLL.

Each RX block consists of four transimpedance amplifiers (TIAs), four deserializers, an error check block, and an injection-locked eight-phase oscillator. It receives the data and checks the number of errors. The photodiode converts the received optical waveform into electrical current. In the RX block, the TIA receives this current input and generates a 26.4 Gbps data stream. To deserialize this data stream, it is passed through an eight-phase injection-locked oscillator; the reference clock of the PLL on the TX side is sent as an amplitude modulated wave, recovered, and used as a reference for the RX PLL to enable deserialization of the data. Utilizing the output from this process, the pseudo-random bit sequence is employed to check for the number of errors within the received data. This error check block is merely an in-lab debugging tool

that would not be included in a field deployment and is hence left absent from Fig. 15. Since FEC was not implemented, it is absent from Fig. 15 as well.

D. Packaging of 3D Integrated Transceiver

Packaging unifies the separate electronic and photonic chips to operate together as a system and facilitates testing of a complicated system. Figure 16(a) shows the cross-section vision for a 1 Tbps unit. The PIC is flip chip bonded to an interposer containing an embedded driver IC, TSVs, and backside redistribution layers (RDLs) with a backside bonded controller IC, all of which are soldered onto a daughter card PCB [shown in Fig. 16(b)]. An aluminum nitride block can be adhered onto the backside of the PIC both to act as a heat spreader and to increase the bond area for a fiber array, and a heat sink can be attached to the other side of this block.

Since this packaging vision is quite intricate to realize, an intermediate package designed to validate the processing steps of the process is shown in Fig. 16(c). This lower-complexity package still requires embedding the controller within an interposer with multiple RDL layers that has a PIC bonded to it; however, electrical continuity to the PCB is formed via wire bonds. Additionally, the controller IC is attached separately to the PCB and has its contacts wirebonded. A sacrifice made in this simpler package is system complexity; rather than running a full 1 Tbps link, this package could only output a single 25 Gbps channel of the system. As this section continues, work outlining forming TSVs, ohmic contacts to TSVs, and fiber array attach was done in isolation in preparation for the full system, and the remaining processing steps were conducted to form the single-channel package.

The TSVs are formed by etching pillars in the degenerately doped silicon interposer wafer, and these silicon TSVs are projected to have resistances on the order of 100s of mΩ. The space between the TSVs is filled with BCB for both electrical isolation and structural support, as shown in Fig. 17(a). The interposer wafer would then be polished to release the TSVs from the substrate. Low resistance ohmic contacts to the silicon TSVs are formed within the temperature specifications of BCB using a Ni/Ti/Au stack; the resultant contact resistivity is approximately $3.1 \times 10^{-6} \Omega \text{ cm}^2$. An example of contacts to the TSVs is shown in Fig. 17(b).

The die slot is produced during the same etching step as the TSVs. The die slot mask feature includes pillars periodic throughout the slot to tailor the gas transport in the slot during the etch, and this helps control the side wall angle uniformity.

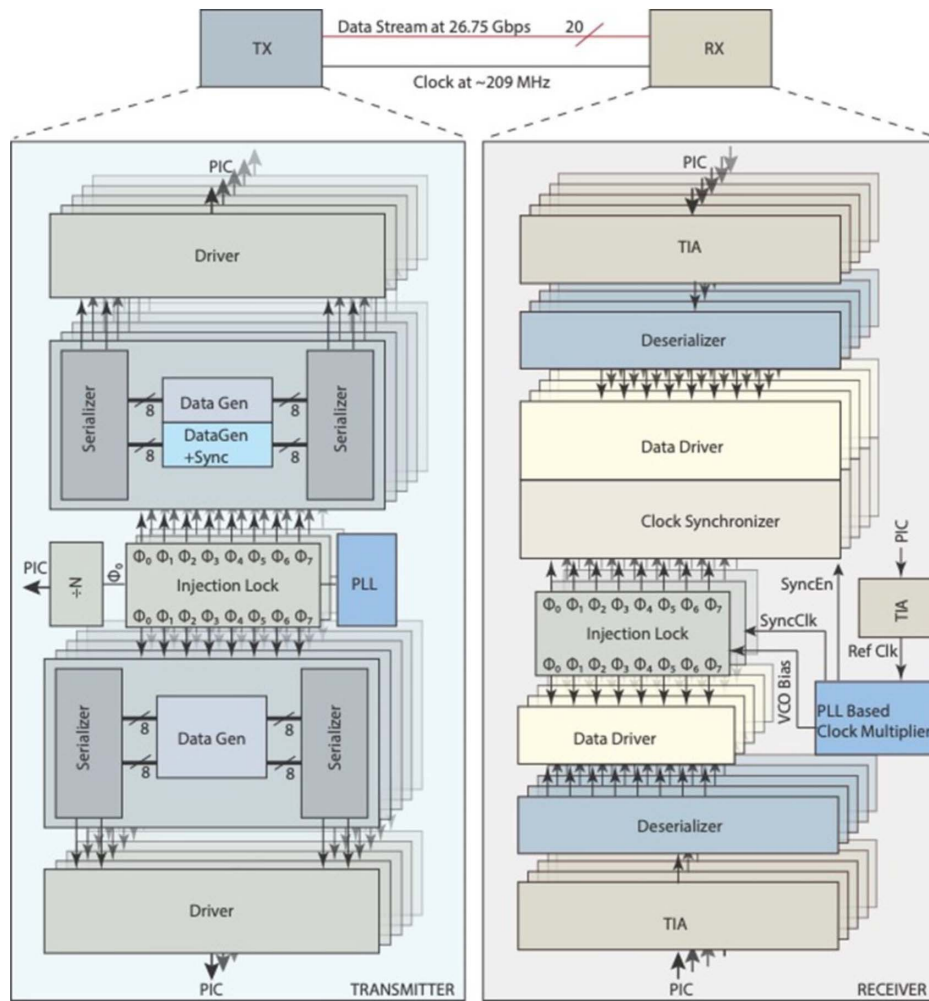


Fig. 15. Overall architecture of the driver IC.

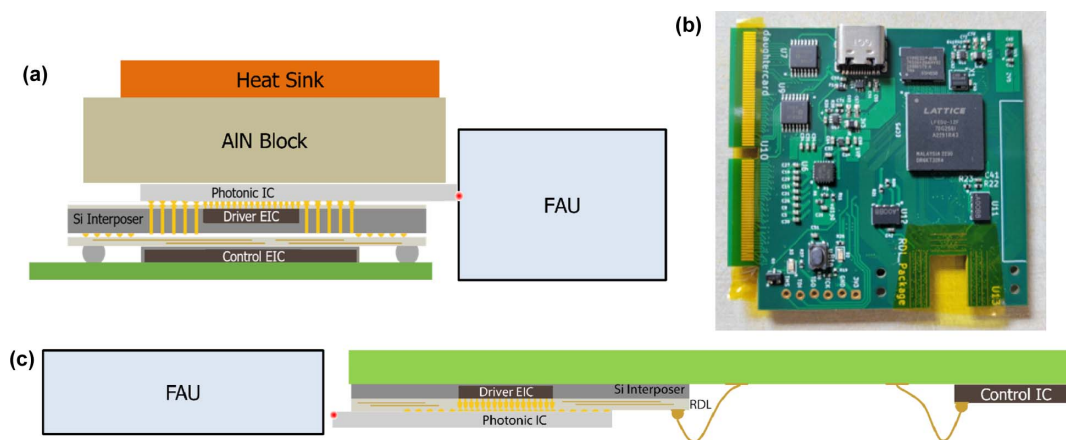


Fig. 16. (a) Cross-section diagram for a 1 Tbps system. (b) Daughter card printed circuit board for 1 Tbps system. (c) Cross-section diagram of simplified 25 Gbps system package.

This results in a high yield for driver placement into these slots across the wafer, which is performed with a flip chip bonder. A socketed driver IC is shown in Fig. 17(c). RDL metallization is

formed by gold plating on BCB and through photoresist. An image of the multiple metallization layers in BCB can be seen in Fig. 17(d).

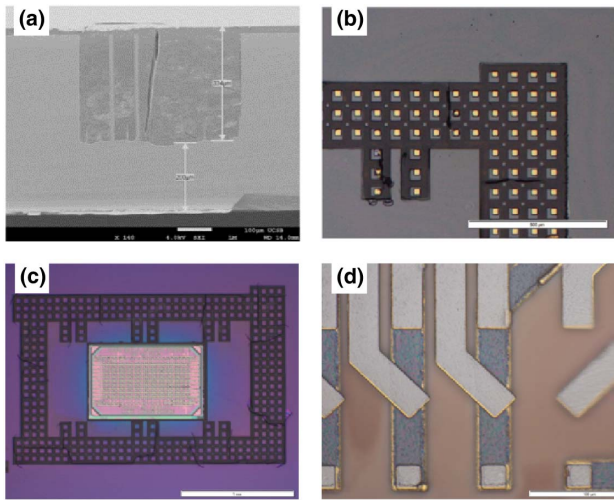


Fig. 17. (a) Doped silicon TSVs with BCB gap fill prior to backside substrate release. (b) Ohmic contacts to the TSVs. (c) Interposer with embedded driver IC. (d) Multiple redistribution layers on BCB.

Indium is deposited onto the pads of the interposer RDL and the embedded driver IC in preparation for flip chip bonding of the PIC onto the interposer; the PICs have a pad pitch of 36 μm in the RF region that bonds to the driver IC and 72 μm for connections to TSVs on the full-complexity interposer. Prior to flip chip bonding, the PIC undergoes several post-processing steps. The silicon photonics wafer leaves the foundry

unsingulated with bare aluminum pads that need to be gold plated for the flip chip bonding procedure. Wafer-scale electroplating was performed on the silicon photonics wafer after it had been cored to 150 mm. Laser stealth dicing was used to singulate the die over saw dicing to ensure a cleaner pad surface and minimize impact on chip-to-chip bonding; pad quality post-singulation is shown in Fig. 18(a). PICs are then facet polished to remove a ledge on the optical interface that would otherwise prevent a standard fiber array from being aligned and attached to the package; the overpolishing is limited to at most 5 μm , which should only perturb the edge coupler performance to a small degree. Figure 18(b) demonstrates the cleanliness of the facet after polishing. The interposer and PIC are then aligned on the stage of the flip chip and bonded, and one such completed single-channel package is shown in Fig. 18(c).

That interposer package is then attached to a printed circuit board. A fiber array can be directly epoxied to the PIC after active alignment; additional fixtures are included to preserve the integrity of the fiber attach process. An example of a successful fiber attach, along with the additional optomechanical and thermal hardware, is shown in Fig. 19(a). One such packaged device was placed into an oven and underwent wavelength scans at temperatures ranging between 20°C and 80°C, and the results are shown in Fig. 19(b). The data show that the insertion loss is stable across wavelength and temperature ranges of interest. With a loss of about 4.2 dB in the optical pathway consisting of two edge couplers with 1 dB per facet loss and 8.1 mm of integrated silicon photonics waveguides, the quality




Fig. 18. (a) PIC pads that bond to the driver IC after plating and singulation. (b) PIC optical facet for fiber coupling after facet polishing. (c) An assembled 25 Gbps package.

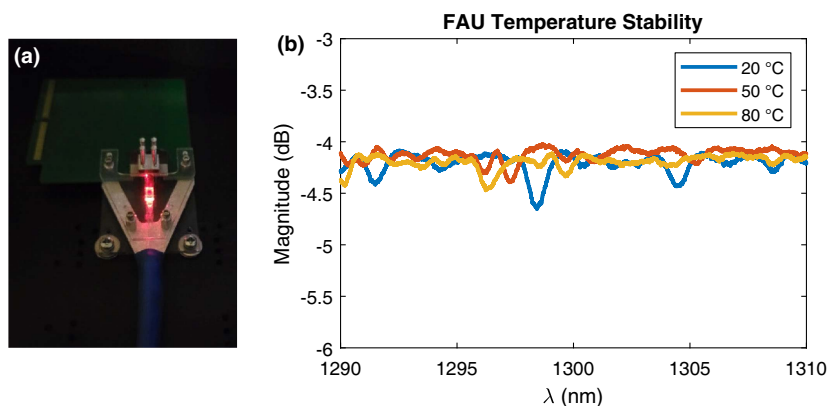


Fig. 19. (a) FAU attach hardware secured to a daughter card. (b) Insertion loss through FAU and PIC alignment loopback across wavelength and temperature.

of the facet polishing and fiber attach is confirmed. The small ripples in the transmission spectrum are assumed to be due to minor polarization fluctuations occurring during the scan, as the fibers in the fiber array attach test were SMF-28 rather than PM fibers. The optical pathway on the PIC includes bends that would filter out any TM light that may be present at any moment.

While much of the work on the package is promising, there are still minor issues that exist in the first generation that require addressing before full system tests can be conducted. Namely, improvements are needed in eliminating shorts present in the RDLs and for increasing the flip chip bond yield between the PIC and the driver EIC.

4. SUBSYSTEM TESTING

There were two main experiments conducted of electronic-photonic systems. The first is a demonstration of a controller IC evaluation board successfully interacting with a silicon photonic ring modulator chip, and the second details our efforts in copackaging the photonics and electronics for a fundamental validation of all sub-components functioning in a single-wavelength 27 Gbps optical link.

A. Ring Locking

Figure 20(a) shows a ring modulator chip that is optically and electrically probed in a manner allowing for light to be passed in and out of the PIC with the capability of applying an EO bias to the phase shifter and to read the monitor photocurrent at the drop port of the ring. These electrical inputs are connected to a PCB with a packaged control IC, shown in Fig. 20(b). Figure 20(c) shows that the resonance wavelength shifted according to the EO driver voltage generated by the control IC.

The next objective after verifying that the controller can shift the ring resonance with an applied voltage is to verify

the controller can convert the monitor photocurrent into an appropriate ADC value and see a semblance of a ring lineshape. This is successfully shown in Fig. 20(d), in which a tunable laser wavelength is set to be on the red side of the resonance at a controller drive voltage of 0 V; as the drive voltage increases, the ring resonance is red shifted, which causes an initial increase in the ADC value as the monitor photocurrent reaches its maximum when the laser is now aligned with the ring resonance. Further increases to the driver voltage past this point lead to a decay in the photocurrent as the laser is further red shifted away from the ring resonance.

Finally, locking the ring is done by instructing the controller to achieve an input photocurrent ADC value on either the red or blue side of the resonance. For a full 1 Tbps system, it is required that all ring modulators should lock to the same side of the ring resonance, preferably on the blue side of the resonance [37]. However, it is likely that some rings will have their MLL channel start on the red side of the resonance. Figure 20(e) demonstrates that the locking circuit is capable of identifying that it is on the red side of the resonance from the amplitude derivative and switches to the highest possible reverse bias output to get on the blue side of the resonance.

B. Discrete 1 – λ Link with QD-MLL Source

Figure 21 shows the experimental setup diagram and results for a single-channel modulation experiment; the integrated silicon photonic ring modulator and PD used in these experiments come from test structure sites with RF probe inputs [shown in Figs. 10(c) and 13(c)] in the same wafer as the fully integrated PICs and are otherwise the same designs used in the 1 Tbps system. One line of a packaged QD-MLL has data encoded onto it with an integrated ring modulator driven by an arbitrary waveform generator (AWG) producing a 27 Gbps NRZ signal. This driving signal has a $V_{pp} = 1$ V and is pre-emphasized to compensate for cable losses. The overall

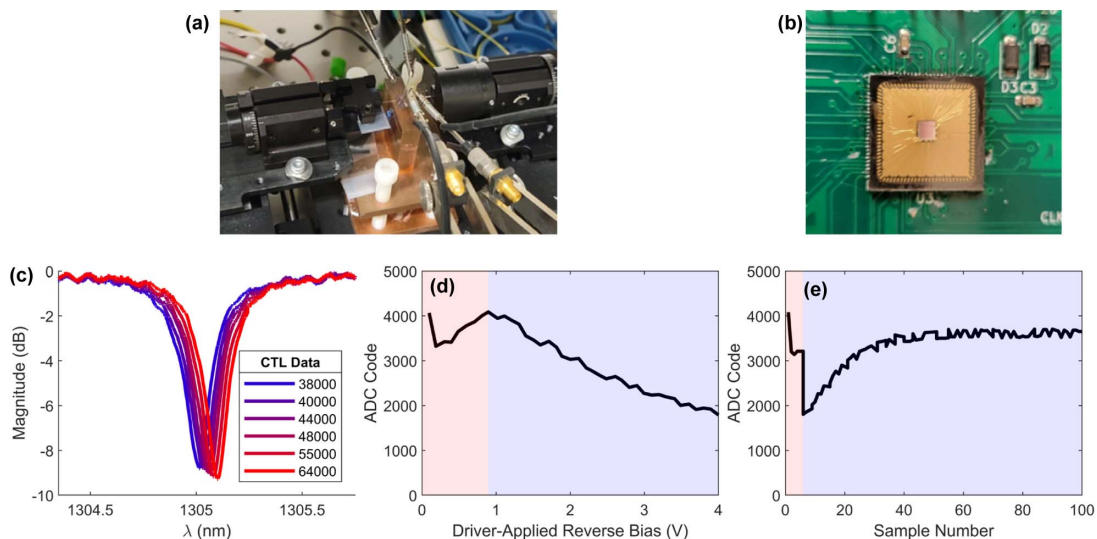


Fig. 20. (a) Ring modulator fiber coupled and DC needle probed on its p-n junction phase shifter and its monitor photodiode. (b) Packaged control IC chip. (c) Ring frequency response shifting according to the applied bias voltage change from the control IC. (d) Scanning a ring modulator resonance by stepping the EO phase shifter bias and reading its monitor photocurrent both with the control IC. (e) Locking a laser that begins on the red side of a ring resonance to a particular photocurrent on the blue side of the resonance. Blue and red regions in (d) and (e) indicate whether the laser is on the blue or red side of the ring's resonance, respectively.

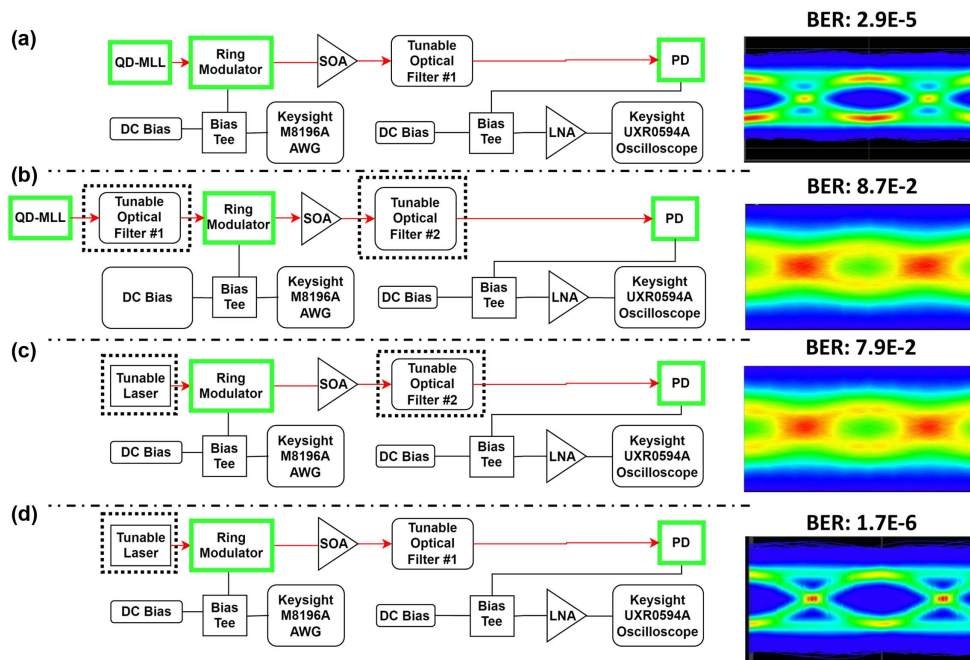


Fig. 21. Link experiments and resultant eye diagrams. Boxed elements indicate variations from the original experimental setup in (a), and elements within a box with a green border represent integrated components. (b) Modification of the original experiment where a single QD-MLL carrier is filtered out of the comb and passed through the ring modulator. (c) Recreating the single-carrier experiment with a laser capable of higher power at the same carrier conditions as prior experiment. (d) Increasing the input power of the single carrier experiment.

comb power from the QD-MLL before entering the TX chip is limited to about +11 dBm (about -2 dBm per line), and the edge couplers on these test structure chips (different than the ones on the system PICs) have an insertion loss of about 3 dB/facet. This leads to an input power of about -5 dBm per comb line at the modulator input.

The optical signal is amplified with a commercial QD-SOA (Innolume SOA-1300-30-PM-35 dB), and a tunable optical filter is then used to remove the unmodulated comb lines; this tunable filter has an excess loss of about 5 dB. The modulated line is coupled into a PD that has a low-noise amplifier (LNA) on its backend that is connected to an oscilloscope for analysis. The QD-SOA in this experiment has a noise figure of 7.5 dB and is providing about 11 dB of gain. The achieved bit error rate (BER) in the preliminary experiment of 2.9×10^{-5} is close to the system's FEC threshold of 1.3×10^{-5} ; however, the simplified TX photonic element used in the experiment lacks a deinterleaver. The absence of a deinterleaver increases the number of comb lines that couple power into the modulator, and the excess input power from the additional comb lines causes self-heating of the ring and degrades signal integrity [as shown in Fig. 9(b)].

To prevent self-heating, a second iteration of the experiment is conducted in which the QD-MLL is filtered down to a single carrier before passing into the MRM [Fig. 21(b)]. This iteration experiences a much worse BER, but this is likely due to the extra 3.5 dB insertion loss introduced by the second filter. This hypothesis is confirmed by substituting the QD-MLL with a commercial tunable laser in Fig. 21(c) with the same wavelength (1298.152 nm) and power conditions (1.7 μ A of average dropped photocurrent in the ring modulator), which

achieves a similar BER; the similarities in the resultant BER between carriers suggest that link performance is in fact being constrained by losses. Since the tunable laser carrier can be set to have a much higher individual line power than the QD-MLL, we use it as an analog to explore whether the QD-MLL could clear the FEC threshold in this link if losses were lower and the QD-MLL was deinterleaved. This tunable laser carrier is capable of a BER almost an order of magnitude smaller than the FEC threshold when the incident power on the ring modulator is increased to about +1 dBm. The single-channel link experiment conducted here is anticipated to introduce 6 dB more loss than the fully integrated WDM demo. Therefore, an appropriately deinterleaved QD-MLL comb source with -5 dBm per line incident at the ring input is expected to be sufficient for a 1 Tbps link.

5. RESULTS AND DISCUSSION

While we have not shown a fully copackaged link, we have shown that the constituent components are likely viable once capable of being brought together. Some of the figures of merit of concern for this link are power consumption and areal bandwidth density.

Figure 22 projects the energy consumption of a 1 Tbps link based upon the MLL and SOA conditions used in the single-wavelength demo as 0.76 pJ/bit. However, if the losses in this experiment matched those expected for the copackaged 1 Tbps link, the energy consumption of the link is expected to reduce to 0.39 pJ/bit due to needing less optical amplification. Beyond lower losses than what were achieved in the single-wavelength demonstration, a 1 Tbps link would likely experience a higher

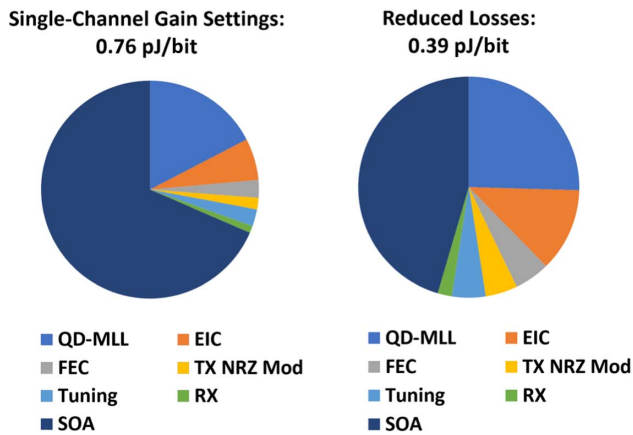


Fig. 22. Energy budget extrapolations for a 1 Tbps system utilizing a QD-SOA.

signal to noise ratio due to the benefits of a flip chip bonded driver chip instead of RF probing, lower input power into the ring modulators, and ring stabilization circuitry; this may further reduce power consumption expectations in a full system demonstration.

While there are many approaches that could be leveraged to further drive down power consumption, Fig. 22 indicates that one of the most substantive ways to do so would be to reduce the amount of optical amplification needed to close the link. One such approach to do so would be to replace the optical gain of QD-SOAs with the electrical gain of integrated avalanche photodiode (APD) receivers; APDs with a gain-bandwidth product of at least 200 GHz would enable removing the QD-SOAs for this application [48]. APDs provide a substantial link loss gain (8 dB) for only marginal increase in bias circuit power consumption. Since this method of gain amplifies each channel individually, it may lead to improved signal integrity by avoiding inter-channel crosstalk mechanisms QD-SOAs can introduce such as cross gain modulation [49]. Additionally, the use of APDs could reduce the amount of power needed to be tapped to monitor photodiodes and the optical clock; since both of these circuits operate at much slower speeds than the data rate, it is possible that they could be biased at a higher gain point since lower bandwidths are needed.

In the driver IC bond region of the PIC, the contiguous area that the 1 Tbps worth of ring modulators and RX demux filters and PDs occupy is about 0.19 mm^2 . This leads to the local area bandwidth density in the RF region being 5.3 Tbps/mm^2 ; the local bandwidth density is currently the commonly reported figure of merit [50].

Despite the high local bandwidth density, the overall PIC footprint has an area 28.6 times larger than the RF region; this disparity can be driven down by iterating on the PIC footprint allocation. The PIC discussed in this work is limited by metal routing and flip chip pitch, which could be addressed with pitch reductions and introducing additional metal routing layers [51]. However, the photonic footprint allocation could also be reduced. For example, a simple change to the architecture would be to perform the TX deinterleaving operation before power splitting. Deinterleavers are much bigger than

power splitters, and they can also require phase shift correction; this therefore also can reduce the number of pads, corresponding drive circuitry, and energy per bit. Integrating the MLL and gain sections to distribute one comb source over many more 500G unit cells would mean needing only one deinterleaver on the TX side per MLL, but may not be the best overall choice for link performance and energy efficiency [9,52,53].

6. CONCLUSIONS

We developed an architecture capable of very low power, very high capacity data transmission. The individual components worked well and operated within a power budget that allowed 1 Tbps transmission per PIC on one port over 20 m with a local PIC bandwidth density of 5.3 Tbps/mm^2 and a total energy consumption below 0.4 pJ/bit with a p-i-n receiver. Optimizations to further compactify the PIC and reduce its energy consumption have been identified.

The mode-locked laser source developed for this experiment demonstrated very flat combs that were very efficient (17% free space wallplug efficiency) and met the demanding needs for links of this nature.

A packaging that allows very high density 3D heterogeneous combination of the PIC with driver and control ICs was developed. The package was not completed due to a lack of time and resources. However, the approach has been developed and appears viable. The fiber attach to this package was successful and demonstrated stable, low loss connection across wavelength and temperature.

Funding. Defense Advanced Research Projects Agency (HR0011-19-C-0083); Semiconductor Research Corporation (2023-JU-3132).

Acknowledgment. The authors thank PHIX Photonics Assembly, Optelligent, Collier Ventures Inc., Micross, Pro Precision Process & Reliability, and Advotech for their backend process and packaging services. The authors also thank Kaiyin Feng, Eamonn Hughes, Michael Zylstra, and Jerome Jahn for their support and helpful discussions. A portion of this work was performed in the UCSB Nanofabrication Facility, an open access laboratory. The views, opinions, and/or findings expressed are those of the authors and should not be interpreted as representing the official views or policies of the Department of Defense or the U.S. Government. Released under Distribution Statement "A" (Approved for Public Release, Distribution Unlimited).

Disclosures. The authors declare no conflicts of interest.

Data Availability. Data underlying the results presented in this paper are not publicly available at this time but may be obtained from the authors upon reasonable request.

REFERENCES

1. G. A. Keeler, "Optical microsystem technologies and applications," in *Optical Fiber Communication Conference* (Optica Publishing Group, 2021), paper M1A-1.

2. D. A. Miller, "Attojoule optoelectronics for low-energy information processing and communications," *J. Lightwave Technol.* **35**, 346–396 (2017).
3. R. Polster, Y. Thonnart, G. Waltener, *et al.*, "Efficiency optimization of silicon photonic links in 65-nm CMOS and 28-nm FDSOI technology nodes," *IEEE Trans. Very Large Scale Integr. Syst.* **24**, 3450–3459 (2016).
4. E. Berikaa, M. S. Alam, S. Bernal, *et al.*, "Next-generation O-band coherent transmission for 1.6 Tbps 10 km intra-datacenter interconnects," *J. Lightwave Technol.* **42**, 1126–1135 (2023).
5. A. Rizzo, A. Novick, V. Gopal, *et al.*, "Massively scalable Kerr comb-driven silicon photonic link," *Nat. Photonics* **17**, 781–790 (2023).
6. M. Dumont, S. Liu, M. Kennedy, *et al.*, "High-efficiency quantum dot lasers as comb sources for DWDM applications," *Appl. Sci.* **12**, 1836 (2022).
7. X. Zheng, S. Lin, Y. Luo, *et al.*, "Efficient WDM laser sources towards terabyte/s silicon photonic interconnects," *J. Lightwave Technol.* **31**, 4142–4154 (2013).
8. E. Lucas, S.-P. Yu, T. C. Briles, *et al.*, "Tailoring microcombs with inverse-designed, meta-dispersion microresonators," *Nat. Photonics* **17**, 943–950 (2023).
9. B. Buscaino, E. Chen, J. W. Stewart, *et al.*, "External vs. integrated light sources for intra-data center co-packaged optical interfaces," *J. Lightwave Technol.* **39**, 1984–1996 (2021).
10. Z. Zhang, D. Jung, J. C. Norman, *et al.*, "Linewidth enhancement factor in InAs/GaAs quantum dot lasers and its implication in isolator-free and narrow linewidth applications," *IEEE J. Sel. Top. Quantum Electron.* **25**, 1900509 (2019).
11. C. Xiang, W. Jin, J. Guo, *et al.*, "Narrow-linewidth III-V/Si/Si₃N₄ laser using multilayer heterogeneous integration," *Optica* **7**, 20–21 (2020).
12. D. Huang, P. Pintus, and J. E. Bowers, "Towards heterogeneous integration of optical isolators and circulators with lasers on silicon," *Opt. Mater. Express* **8**, 2471–2483 (2018).
13. K. Y. Yang, C. Shirpurkar, A. D. White, *et al.*, "Multi-dimensional data transmission using inverse-designed silicon photonics and microcombs," *Nat. Commun.* **13**, 7862 (2022).
14. K. Hosseini, E. Kok, S. Y. Shumarayev, *et al.*, "8 Tbps co-packaged FPGA and silicon photonics optical IO," in *Optical Fiber Communications Conference and Exhibition (OFC)* (IEEE, 2021), pp. 1–3.
15. Y. Bian, T. Hirokawa, V. Karra, *et al.*, "Monolithically integrated self-aligned SiN edge coupler with <math><0.6/0.8</math> dB TE/TM insertion loss, <math><-39</math> dB back reflection and >520 mW high-power handling capability," in *Optical Fiber Communication Conference* (Optica Publishing Group, 2023), paper M3C-3.
16. K. Padmaraju, X. Zhu, L. Chen, *et al.*, "Intermodulation crosstalk characteristics of WDM silicon microring modulators," *IEEE Photonics Technol. Lett.* **26**, 1478–1481 (2014).
17. A. James, A. Novick, A. Rizzo, *et al.*, "Scaling comb-driven resonator-based DWDM silicon photonic links to multi-Tb/s in the multi-FSR regime," *Optica* **10**, 832–840 (2023).
18. K. Padmaraju, N. Ophir, Q. Xu, *et al.*, "Error-free transmission of microring-modulated BPSK," *Opt. Express* **20**, 8681–8688 (2012).
19. M. Georgas, J. Leu, B. Moss, *et al.*, "Addressing link-level design tradeoffs for integrated photonic interconnects," in *IEEE Custom Integrated Circuits Conference (CICC)* (IEEE, 2011), pp. 1–8.
20. A. Malik, S. Liu, E. Timurdogan, *et al.*, "Low power consumption silicon photonics datacenter interconnects enabled by a parallel architecture," in *Optical Fiber Communication Conference* (Optica Publishing Group, 2021), paper W6A-3.
21. C. Xiang, W. Jin, D. Huang, *et al.*, "High-performance silicon photonics using heterogeneous integration," *IEEE J. Sel. Top. Quantum Electron.* **28**, 8200515 (2021).
22. "400G-FR4 QSFP-DD OCP optical transceiver specification revision 0.1," <https://www.opencompute.org/documents/meta-400g-fr4-optical-transceiver-specification-for-ocp-rev0-1-1-pdf> (accessed on 12 December 2023).
23. M. van Niekerc, V. Deenadalan, A. Rizzo, *et al.*, "Wafer-scale-compatible substrate undercut for ultra-efficient SOI thermal phase shifters," in *2022 Conference on Lasers and Electro-Optics (CLEO)* (IEEE, 2022), pp. 1–2.
24. A. Ribeiro, S. Declercq, U. Khan, *et al.*, "Column-row addressing of thermo-optic phase shifters for controlling large silicon photonic circuits," *IEEE J. Sel. Top. Quantum Electron.* **26**, 6100708 (2020).
25. D. Coenen, M. Kim, H. Oprins, *et al.*, "Thermal modeling of hybrid three-dimensional integrated, ring-based silicon photonic–electronic transceivers," *J. Opt. Microsyst.* **4**, 011004 (2023).
26. H. Jayatilaka, H. Frish, R. Kumar, *et al.*, "Post-fabrication trimming of silicon photonic ring resonators at wafer-scale," *J. Lightwave Technol.* **39**, 5083–5088 (2021).
27. S. S. Djordjevic, K. Shang, B. Guan, *et al.*, "CMOS-compatible, athermal silicon ring modulators clad with titanium dioxide," *Opt. Express* **21**, 13958–13968 (2013).
28. "10 GB/s high-sensitivity limiting PIN-TIA optical receiver," https://www.discoverysemi.com/Product_Pages/DSCR603.php (accessed on 12 December 2023).
29. H. Al Maruf and M. Chowdhury, "Memory disaggregation: advances and open challenges," *ACM SIGOPS Oper. Syst. Rev.* **57**, 29–37 (2023).
30. Z. Wang, H. Huang, J. Zhang, *et al.*, "Shuhai: benchmarking high bandwidth memory on FPGAs," in *IEEE 28th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)* (IEEE, 2020), pp. 111–119.
31. L. H. Brendler, A. L. Zimpeck, C. Meinhardt, *et al.*, "Multi-level design influences on robustness evaluation of 7 nm FinFET technology," *IEEE Trans. Circuits Syst. I, Reg. Papers* **67**, 553–564 (2019).
32. B. Dong, M. Dumont, O. Terra, *et al.*, "Broadband quantum-dot frequency-modulated comb laser," *Light Sci. Appl.* **12**, 182 (2023).
33. J. Duan, Y. Zhou, B. Dong, *et al.*, "Effect of p-doping on the intensity noise of epitaxial quantum dot lasers on silicon," *Opt. Lett.* **45**, 4887–4890 (2020).
34. Y. Gao, N. Pestana, S. Deckoff-Jones, *et al.*, "Passive integrated athermal (de)multiplexers on 300 mm silicon photonics wafers," in *Optical Fiber Communication Conference* (Optica Publishing Group, 2023), paper M4I-3.
35. A. M. Netherton, Y. Gao, N. Pestana, *et al.*, "Athermal, fabrication-tolerant Si-SiN FIR filters for a silicon photonics foundry platform," *Opt. Express* **31**, 23952–23965 (2023).
36. N. M. Fahrenkopf, C. McDonough, G. L. Leake, *et al.*, "The aim photonics MPW: a highly accessible cutting edge technology for rapid prototyping of photonic integrated circuits," *IEEE J. Sel. Top. Quantum Electron.* **25**, 8201406 (2019).
37. M. De Cea, A. H. Atabaki, and R. J. Ram, "Power handling of silicon microring modulators," *Opt. Express* **27**, 24274–24285 (2019).
38. A. V. Krishnamoorthy, X. Zheng, G. Li, *et al.*, "Exploiting CMOS manufacturing to reduce tuning requirements for resonant optical devices," *IEEE Photonics J.* **3**, 567–579 (2011).
39. A. Rizzo, V. Deenadayalan, M. van Niekerc, *et al.*, "Ultra-efficient foundry-fabricated resonant modulators with thermal undercut," in *CLEO: Science and Innovations* (Optica Publishing Group, 2023), paper SF2K-6.
40. D. A. Miller, "Energy consumption in optical modulators for interconnects," *Opt. Express* **20**, A293–A308 (2012).
41. D. Gostimirovic and N. Y. Winnie, "Ultralow-power double vertical junction microdisk modulators," *IEEE J. Sel. Top. Quantum Electron.* **27**, 3400907 (2021).
42. H. Gevorgyan, A. Khilo, M. T. Wade, *et al.*, "Moscap ring modulator with 1.5 μm radius, 8.5 THz FSR and 30 GHz/V shift efficiency in a 45 nm SOI CMOS process," in *Optical Fiber Communications Conference and Exhibition (OFC)* (IEEE, 2021), pp. 1–3.
43. H. Yu, D. Ying, M. Pantouvaki, *et al.*, "Trade-off between optical modulation amplitude and modulation bandwidth of silicon micro-ring modulators," *Opt. Express* **22**, 15178–15189 (2014).
44. M. Bahadori, D. Nikolova, S. Rumley, *et al.*, "Optimization of microring-based filters for dense WDM silicon photonic interconnects," in *IEEE Optical Interconnects Conference (OI)* (IEEE, 2015), pp. 84–85.
45. W. Bogaerts, P. De Heyn, T. Van Vaerenbergh, *et al.*, "Silicon microring resonators," *Laser Photonics Rev.* **6**, 47–73 (2012).
46. A. Li, T. Van Vaerenbergh, P. De Heyn, *et al.*, "Backscattering in silicon microring resonators: a quantitative analysis," *Laser Photonics Rev.* **10**, 420–431 (2016).
47. J. Michel, J. Liu, and L. C. Kimerling, "High-performance Ge-on-Si photodetectors," *Nat. Photonics* **4**, 527–534 (2010).

48. J. Zhang, B. P.-P. Kuo, and S. Radic, "64 Gb/s PAM4 and 160 Gb/s 16QAM modulation reception using a low-voltage Si-Ge waveguide-integrated APD," *Opt. Express* **28**, 23266–23273 (2020).
49. J. Kim, M. Laemmlin, C. Meuer, *et al.*, "Theoretical and experimental study of high-speed small-signal cross-gain modulation of quantum-dot semiconductor optical amplifiers," *IEEE J. Quantum Electron.* **45**, 240–248 (2009).
50. S. Daudlin, A. Rizzo, S. Lee, *et al.*, "3d photonics for ultra-low energy, high bandwidth-density chip data links," *arXiv*, [arXiv:2310.01615](https://arxiv.org/abs/2310.01615) (2023).
51. S. Jangam and S. S. Iyer, "Silicon-interconnect fabric for fine-pitch ($\leq 10 \mu\text{m}$) heterogeneous integration," *IEEE Trans. Compon. Packag. Technol.* **11**, 727–738 (2021).
52. M. L. Davenport, S. Liu, and J. E. Bowers, "Integrated heterogeneous silicon/III–V mode-locked lasers," *Photonics Res.* **6**, 468–478 (2018).
53. K. Feng, C. Shang, E. Hughes, *et al.*, "Quantum dot lasers directly grown on 300 mm Si wafers: planar and in-pocket," *Photonics* **10**, 534 (2023).