

# Scaling up silicon photonic-based accelerators: Challenges and opportunities

F

Cite as: APL Photonics 7, 020902 (2022); <https://doi.org/10.1063/5.0070992>

Submitted: 10 September 2021 • Accepted: 23 January 2022 • Published Online: 08 February 2022


 M. A. Al-Qadasi,  L. Chrostowski,  B. J. Shastri, et al.

## COLLECTIONS

Note: This paper is part of the APL Photonics Special Topic on Photonics and AI in Information Technologies.

F

This paper was selected as Featured

[View Online](#)[Export Citation](#)[CrossMark](#)

# Scaling up silicon photonic-based accelerators: Challenges and opportunities

Cite as: APL Photon. 7, 020902 (2022); doi: 10.1063/5.0070992  
 Submitted: 10 September 2021 • Accepted: 23 January 2022 •  
 Published Online: 8 February 2022



M. A. Al-Qadasi,<sup>1,a)</sup>  L. Chrostowski,<sup>1,2</sup>  B. J. Shastri,<sup>3,4</sup>  and S. Shekhar<sup>1</sup>

## AFFILIATIONS

<sup>1</sup>The University of British Columbia, Vancouver, British Columbia V6T 1Z4, Canada

<sup>2</sup>Stewart Blusson Quantum Matter Institute, University of British Columbia, Vancouver, British Columbia V6T 1Z4, Canada

<sup>3</sup>Queen's University, Kingston, Ontario K7L 3N6, Canada

<sup>4</sup>Vector Institute, Toronto, Ontario M5G 1M1, Canada

**Note:** This paper is part of the APL Photonics Special Topic on Photonics and AI in Information Technologies.

**a)** Author to whom correspondence should be addressed: [alqadasi@ece.ubc.ca](mailto:alqadasi@ece.ubc.ca)

## ABSTRACT

Digital accelerators in the latest generation of complementary metal–oxide–semiconductor processes support, multiply, and accumulate (MAC) operations at energy efficiencies spanning 10–100 fJ/Op. However, the operating speed for such MAC operations is often limited to a few hundreds of MHz. Optical or optoelectronic MAC operations on today's SOI-based silicon photonic integrated circuit platforms can be realized at a speed of tens of GHz, leading to much lower latency and higher throughput. In this Perspective, we study the energy efficiency of integrated silicon photonic MAC circuits based on Mach–Zehnder modulators and microring resonators. We describe the bounds on energy efficiency and scaling limits for  $N \times N$  optical networks with today's technology based on the optical and electrical link budget. We also describe research directions that can overcome the current limitations.

© 2022 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0070992>

## I. INTRODUCTION

Vector matrix multiplication operations represent the core of artificial neural networks (ANNs) and other computing applications of hardware accelerators. ANNs are realized in digital complementary metal–oxide–semiconductor (CMOS) circuits with multiple processing elements implementing multiply and accumulate (MAC) operations, which calculate the product of two numbers and add the result to an accumulator.<sup>1</sup> The processing elements can be arranged in a systolic architecture, where data are passed through connected processing elements in a rhythmic sequence, to perform MAC operations either spatially or temporally over several clock cycles.<sup>2</sup>

Integrated silicon photonics (SiP) circuits have been popularly employed in high-speed links to move data at a rate of tens of Gb/s, where optical modulation is more efficient than electronic switching for transmitting data over significant distances.<sup>3</sup> Optical modulation can be realized using Mach–Zehnder modulators (MZMs) or microring modulators (MRMs).<sup>4</sup> MZMs are broadband and easily support complex modulation schemes.<sup>3</sup> MRMs have significantly smaller footprint and driver power consumption.<sup>5</sup> As a

technology, the current generation of SiP has now matured with high volume shipments for datacenter transceivers from companies such as Intel and Cisco. SiP circuits comprising Mach–Zehnder interferometers (MZIs) or microring resonators (MRRs) have also been used for other applications, such as high-speed optical switches and filters.<sup>6–10</sup>

SiP is also being used for computing applications,<sup>11–16</sup> where devices such as MZIs, MZMs, MRRs, and MRMs are used for computation in the optical analog domain. These encompass inference and training accelerators used for machine learning and neuromorphic computing applications where convolution takes 80% of the total processing time.<sup>17–19</sup> Other integrated optical configurations implemented using field-programmable photonic arrays were shown to carry out linear transformations for signal processing and control.<sup>20,21</sup> Linear transformation circuits are also employed in Ising machines<sup>22</sup> and photonic quantum computing processors.<sup>23</sup>

In this Perspective, we describe the advantages and challenges of implementing MAC operations using SiP and comment on how to address them. This Perspective is organized as follows: Secs. II and III describe the link budget, energy efficiency, and scaling

opportunities for SiP MAC operations implemented with MZM and MRM, respectively. Section IV explores the possible approaches to further improve SiP MAC systems. It introduces the ongoing research in the field of SiP that when fully realized will lead to significant changes in the field of optical computing and communication. Section V concludes this Perspective.

## II. MZM-BASED SI-PHOTONIC IMPLEMENTATION

### A. System architecture

Figure 1 illustrates an MZM-based SiP implementation of an optical accelerator. Using off-chip lasers, light is guided by a polarization-maintaining (PM) single mode fiber (SMF), gets coupled to the SiP chip via an edge coupler, and is then split to  $N$  parts. These parts are modulated by using an array of  $N$  MZI modulators and fed into an  $N \times N$  weight transformation (multiplication) matrix,  $W_{N \times N}$ . The optical intensities at the MAC outputs,  $Y_{N \times 1}$ , are thus described by the multiplication product of the input vector,  $V_{N \times 1}$ , and the weight matrix,  $W_{N \times N}$ , as

$$Y_{N \times 1} = W_{N \times N} V_{N \times 1}. \quad (1)$$

In other words, the weight matrix performs linear transformation for the input vector,  $V_{N \times 1}$ , and delivers  $N$  outputs,  $Y_{N \times 1}$ , that are routed to an array of photodetectors (PDs). These PDs are then connected to the electrical components, including transimpedance amplifiers (TIAs), main amplifiers, and sense amplifier-based comparators, which are either built in a separate CMOS/BiCMOS chip or monolithically integrated with the SiP devices in the same process.

Singular value decomposition (SVD) is an effective approach to represent a given matrix as a factorization of multiple matrices.<sup>24,25</sup> SVD decomposes a real matrix into a product of unitary matrices and a diagonal matrix. This is useful in the experimental realization of an  $N \times N$  matrix topology in which the sequential product of rotation matrices represents the sequential arrangement of linear transformation units in the overall matrix grid.<sup>26</sup> A  $2 \times 2$  linear transformation unit in the whole grid arrangement is practically implemented using a tunable beam splitter (TBS), as seen in

Fig. 1.<sup>6,27</sup> A TBS comprises of an MZI with a phase shifter ( $\theta$ ) in at least one of the arms, along with either an outer phase shifter ( $\phi$ ) or a tunable directional coupler.<sup>28,29</sup> The transfer function for a single TBS can be described using the matrix representations for ideal 50:50 beam splitters and lossless phase shifters as

$$T_{BS} = \frac{1}{2} \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \begin{bmatrix} e^{i\theta} & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & i \\ i & 1 \end{bmatrix} \begin{bmatrix} e^{i\phi} & 0 \\ 0 & 1 \end{bmatrix} \\ = ie^{i\frac{\theta}{2}} \begin{bmatrix} e^{i\phi} \sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) \\ e^{i\phi} \cos\left(\frac{\theta}{2}\right) & -\sin\left(\frac{\theta}{2}\right) \end{bmatrix}. \quad (2)$$


Thus, any arbitrary light redistribution can be obtained by changing  $\theta$  and  $\phi$ .

The SVD of the weight matrix can be described as

$$W = D \prod_{m,n} T_{m,n}, \quad (3)$$

where  $D$  is a diagonal matrix and  $T_{m,n}$  represents the transformation matrix for a  $2 \times 2$  node between two input terminals,  $m$  and  $n$ , within the  $N \times N$  multiplication matrix,<sup>27</sup> given as

$$T_{m,n} = ie^{i\frac{\theta}{2}} \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & \cdot & 0 \\ \cdot & \cdot & \begin{bmatrix} e^{i\phi} \sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) \\ e^{i\phi} \cos\left(\frac{\theta}{2}\right) & -\sin\left(\frac{\theta}{2}\right) \end{bmatrix}_{m,n} & \cdot & \cdot \\ \cdot & \cdot & \cdot & \cdot & \cdot \\ 0 & \cdot & \cdot & \cdot & 1 & 0 \\ 0 & 0 & \cdots & \cdot & 0 & 1 \end{bmatrix}. \quad (4)$$



**FIG. 1.** SiP circuit diagram of an  $N \times N$  MZM-based accelerator. High-speed PN phase shifters and low-speed thermo-optic phase shifters are colored in blue and red, respectively.

## B. Optical network link budget

The photonic components shown in Fig. 1 are simulated in Cadence Spectre for  $8 \times 8$  and  $32 \times 32$  transformation matrix sizes to verify the optical link budget analysis. The optical components are modeled in Verilog-A to enable electronics-photonics co-simulation.<sup>30,31</sup> The models of some of the components used in this Perspective are derived from Ref. 31, with some modifications to account for the laser electrical power consumption, the wall plug efficiency,  $\eta_{WPE}$ , link losses, etc. The optical power of laser is set to 0 dBm to estimate the optical power at the output terminals of the network, incident on the PDs. Coupling light to the SiP chip introduces loss in the range of 0.6–3 dB based on the coupling scheme used.<sup>32,33</sup> The overall coupling loss from the laser to the SiP chip is collectively estimated as 1.6 dB, considering possible optimizations in the coupling efficiency. 1.6 dB is also a realistic estimate for photonic wire bonds (PWBs), an emerging technology which involves writing three-dimensional waveguides in a photosensitive polymer. PWBs have demonstrated efficient interfacing between the external sources to the silicon waveguide with coupling losses as low as 0.4 up to 1.7 dB.<sup>34–38</sup>

For an input vector size of  $N$ , light passes through  $\log_2 N$  splitters before modulation, resulting in a total insertion loss of  $10 \log_{10} N + EL_{splitter} \cdot \log_2 N$  dB, where  $EL_{splitter}$  is the estimated excess loss for a single splitter and ranges between 0.01 and 0.5 dB.<sup>39–43</sup> The estimated excess loss for beam splitters and combiners in this study is 0.01 dB. The overall attenuation in the silicon waveguide is a function of the depth of the network. The MZM-based implementation is based on Clement's arrangement, which is composed of beam splitters and phase shifters that can be programmed to implement linear transformation.<sup>27</sup> With an MZM representing one node in Clement's arrangement,<sup>27</sup> the waveguide attenuation can be approximated as  $N\eta_{wg}L_{MZI}$ , where  $\eta_{wg}$  represents the optical intensity attenuation in the Si waveguide and  $L_{MZI}$  is the length of an MZM arm, with the chosen values of 3 dB/cm and 0.5 mm, respectively. The insertion loss introduced by an MZM's PN phase shifter is approximated as 1 dB/mm.<sup>41</sup>

The insertion loss through a node is dependent on the excess loss for cross and through states,  $EL_{cross}$  and  $EL_{thru}$ , respectively, both of which are typically  $<1$  dB.<sup>44,45</sup> For simplicity, two phase shifters connected by using 3 dB adiabatic directional couplers ( $EL_{DC} \sim 0.1$  dB)<sup>39</sup> are assumed in this work for analysis.

To study the optical attenuation of the  $8 \times 8$  MZM implementation shown in Fig. 1 and verify its functionality, all inputs, except for the uppermost terminal ( $m = 1$ ), are driven by  $V_\pi$  voltage that creates a  $\pi$  phase shift difference between their MZM's arms and null their outputs. As light is set to propagate through the uppermost input terminal, the optical depth is defined by the route passing through the diagonal TBS nodes with  $i = j$ .

For a rectangular mesh arrangement, the matrix optical depth is equal to  $N$  with a total number of  $\frac{N(N-1)}{2}$  optical crossings.<sup>26,27</sup> Hence, the  $8 \times 8$  matrix implementation shown in Fig. 1 has an optical depth of 8 with 28 crossings.

The total optical link budgets are calculated based on (5), where  $P_{SMF-att}$ ,  $P_{EC-IL}$ ,  $P_{Si-att}$ ,  $P_{splitter-IL,EL}$ ,  $P_{PS-IL}$ ,  $P_{DC-IL}$ , and  $P_{penalty}$  represent the attenuation introduced by the SMF fiber, fiber to chip coupling loss, silicon waveguide attenuation, splitter insertion

and excess loss, phase shifters' insertion loss, total adiabatic coupling insertion loss, and network penalty, respectively. The network penalty takes into account further impairments due to extinction ratio, crosstalk, intersymbol interference (ISI), and laser relative intensity noise (RIN), which is caused by the random spontaneous emission over time,<sup>4,46,47</sup>

$$P_{O/p}(\text{dBm}) = P_{laser} - P_{SMF-att} - P_{EC-IL} - P_{Si-att} - P_{splitter-IL,EL} - NP_{PS-IL} - NP_{DC-IL} - P_{penalty}. \quad (5)$$

Figure 2 shows the calculated optical power throughout  $N \times N$  networks with different input vector sizes. The optical power of the laser is set to 0 dBm for ease of illustration. Besides the attenuation due to the splitting, it can be noticed that the losses introduced by the optical components in the multiplication matrix (i.e., directional couplers and phase shifters) pose a limitation for scaling the network due to the highly attenuated optical intensities reaching the outputs. Figure 3 shows the optical intensities required at the analog front-end (AFE) to detect a signal with a resolution of  $n_{ilp}$  bit. This is obtained by representing the desired output signal and current noises in terms of the received optical intensity, as given in Eq. (8). The blue dotted lines represent the optical power at the matrix outputs and the corresponding bit resolution for a laser intensity of 10 dBm. It can be shown that the maximum achievable matrix size is  $\sim 35 \times 35$  for binary networks operating at  $DR = 10$  GS/s. We revisit this calculation again in Sec. II D.

## C. Energy efficiency

The total electrical power dissipation of the whole network comprises of the power consumed by the laser, input modulators, thermo-optic tuning of the matrix phase shifters, and the AFE, including PDs. Accordingly, for a configuration of size  $N \times N$  operating at a data rate of  $DR$ , the energy efficiency ( $J/\text{Op}$ ) can be





FIG. 2. Optical link budget analysis for MZM-based matrix sizes of  $N = 8, 16, 32, 64$ , and  $128$  for  $R = 1.2$  A/W and  $DR = 10$  GS/s.





**FIG. 3.** Targeted AFE sensitivity for  $n_{i/p} = \{1, 2, 3, 4, 5, 6\}$  and the output power for MZM-based SiP matrices with sizes of  $N = 8, 16, 32, 64$ , and  $128$ .  $P_{laser} = 10$  dBm.

calculated as

$$E(J/Op) = \frac{P_{laser}}{\gamma \cdot 2N^2 \cdot DR} + \frac{NP_{i/p-drivers} + 2P_{mem-interface}}{2N^2 \cdot DR} + \frac{2(N-1)P_{mat-tuning} + P_{SOA} + P_{o/p-AFE}}{2N \cdot DR}, \quad (6)$$

where  $P_{laser}$ ,  $P_{i/p-drivers}$ ,  $P_{mem-interface}$ ,  $P_{mat-tuning}$ , and  $P_{o/p-AFE}$  represent the electrical power dissipated due to the laser, input modulator drivers, data fetch interfacing circuits, matrix tuning, and the output

AFE circuits, respectively.  $P_{SOA}$  represent the electrical power dissipated if a semiconductor optical amplifier (SOA) is used to recover the loss.

The factor  $\gamma$  refers to the energy efficiency enhancement. It can be represented as  $\gamma = \rho_{opt}^2 \rho_{SOA}$ , where  $\rho_{opt}$  represents the energy scaling due to the loss of precision factor, and will be described later in Sec. II D.  $\rho_{SOA}$  represents the efficiency enhancement due to an SOA. Assuming an SOA introducing a gain of  $\eta_{SOA}$  (in dB), the corresponding enhancement is  $\rho_{SOA} = 10^{\frac{\eta_{SOA}}{10}}$ . The use of an SOA is discussed later in Sec. IV, but it can be inferred from Eq. (6) that the use of an SOA always degrades the overall energy efficiency.

The amount of power dissipated by the laser is represented in terms of the laser's wall plug efficiency,  $\eta_{WPE}$ , the optical insertion losses introduced by the SMF fiber,  $IL_{SMF}$ , the fiber to chip coupling  $IL_{EC}$ , the silicon waveguide loss,  $IL_{WG}$ , the input MZM loss,  $IL_{i/p-MZM}$ , the weight phase shifter loss,  $IL_{weight-PS}$ , the directional coupler loss,  $IL_{DC}$ , and the receiver's PD sensitivity,  $P_{PD-opt}$ , as given in the following equation:

$$P_{laser} = \frac{10^{\frac{IL_{WG}(dB)N(L_{MZM})}{10}} N}{IL_{SMF} IL_{EC} (EL_{splitter})^{\log_2 N} IL_{i/p-MZM} (IL_{PS})^{2N}} \times \frac{P_{PD-opt}}{\eta_{WPE} (IL_{DC})^{2N} IL_{penalty}}. \quad (7)$$

The total length of the waveguide was roughly approximated as the length spanning the optical depth of the matrix only. The output sensitivity is solved based on the targeted bit resolution,  $n_{i/p}$ , and the total noise at the output front-end due to the photodetector shot noise, dark current  $I_d$ , thermal noise, and laser relative intensity noise (RIN), as given in Eq. (8), with values reported in Table I. The parameters  $R$ ,  $R_L$ ,  $k$ , and  $T$  represent the PD responsivity, output load resistance, Boltzmann constant, and the absolute temperature,

$$n_{i/p} = \frac{1}{6.02} \left[ 20 \log_{10} \left( \frac{R \times P_{PD-opt}}{\left[ \sqrt{2q(RP_{PD-opt} + I_d) + \frac{4kT}{R_L}} + R^2 P_{PD-opt}^2 RIN + \sqrt{2qI_d + \frac{4kT}{R_L}} \right] \sqrt{DR/\sqrt{2}}} \right) - 1.76 \right]. \quad (8)$$

The MZM drivers consume power that scales linearly with  $N$  as  $P_{mod-driver} = N \cdot DR \cdot E_{MZM-driver}$ , where  $E_{MZM-driver}$  represents the energy efficiency of the MZM driver. For binary resolution, the power consumed by the drivers and AFEs is extracted from recent work on PAM2 and is typically in the range of  $\sim 2$  pJ/b.<sup>49</sup>

The matrix weights are tuned using thermo-optic phase shifters (TO-PSs). Doped Si heaters on the SOI platform typically dissipate about  $\sim 20$  mW for a  $\pi$ -shift.<sup>6,50,51</sup> The efficiency of TO-PS can be improved using other heater materials such as TiN, substrate undercut to improve insulation, and deep trenches to reduce thermal crosstalk.<sup>50,52</sup> This can be shown to significantly improve the overall energy efficiency of the network, as illustrated in Fig. 5. Assuming uniformly distributed weights, the expected energy consumption of the thermo-optic phase shifter is  $E_{TO-PS} = \frac{1}{P_\pi} \int_0^{P_\pi} P_{heater} dP_{heater} = \frac{P_\pi}{2}$ , where  $P_\pi$  denotes the amount of electrical power required to create a

phase shift of  $\pi$ . Calculating for all the nodes in Clement's topology, the total average tuning power is  $\frac{N(N-1)}{4} P_\pi$ .

After optical processing, the optical data need to be converted to the electrical domain to be processed, stored, or reused in other networks. Efficient opto-electronic receivers, comprising a PD, a TIA, and main amplifiers, have been shown to have energy efficiencies of  $\sim 0.4$  to  $2.4$  pJ/b.<sup>53-58</sup> For an AFE operating at 10 Gb/s and realized in 40 nm CMOS technology, an energy efficiency of  $0.4$  pJ/b<sup>53</sup> is assumed for the calculation of binary resolution AFE, which scales with a factor of  $N$  for the whole output array. Higher AFE resolutions entail the use of linear TIAs along with analog to digital converter (ADC) circuits to recover the digital data. High-speed linear TIAs have shown efficiencies as low as  $0.6$  pJ/b.<sup>59</sup> The energy consumption for ADCs is extracted from the energy per conversion figure of merit (FOM) such that  $E_{ADC} (J/b) = 2^N \times FOM$ .

TABLE I. SNR calculation parameters.

Parameter	Description	Value
$P_{laser}$	Laser power intensity	10 dBm
$R$	PD responsivity	1 A/W <sup>48</sup>
$R_L$	Load resistance	50 $\Omega$
$I_d$	Dark current	35 nA <sup>48</sup>
$T$	Absolute temperature	300 K
$DR$	Data rate	10 GS/s
$B_o$	Optical bandwidth	25 GHz
$B_e$	Electrical bandwidth	$DR/\sqrt{2}$ GHz
$\lambda$	Wavelength	1550 nm
$RIN$	Relative intensity noise	-140 dB/Hz <sup>46,47</sup>
$WPE$	Wall plug efficiency	10%

The energy consumption values used in this work for 2b, 3b, and 4b are 1.7, 3.1, and 5.7 pJ/b based on a FOM of 0.335 pJ/conversion for an ADC designed to operate at a sampling rate of 28 GS/s.<sup>60</sup>

Providing high-speed serial inputs to the SiP accelerator requires FIFOs and multiplexers to interface the data transfer with DRAM, as shown in Fig. 4. For a fair comparison to digital CMOS implementations, the power dissipation for both input and output interfacing circuits, represented by  $P_{mem-interface}$ , is taken into consideration in the energy efficiency calculation of SiP implementations. The power dissipated by the first in first out (FIFO), multiplexers, clock dividers, and retimers is estimated as 5.77 mW in 28 nm CMOS based on the data reported in Ref. 61 in 180 nm CMOS technology.

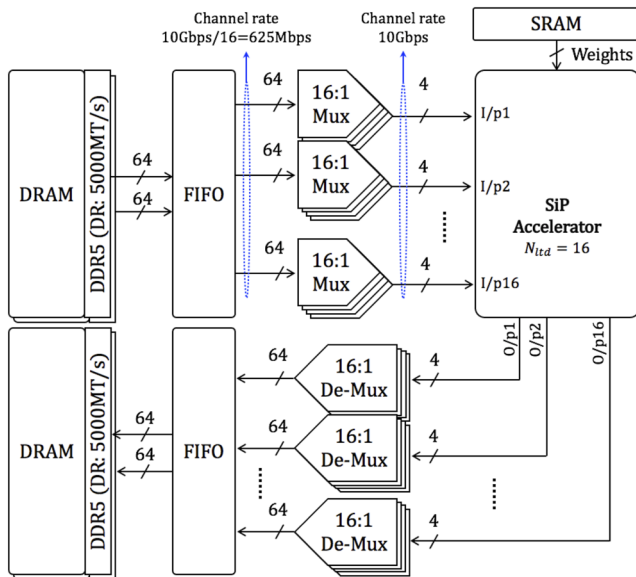


FIG. 4. Input data fetch for a SiP implementation with  $n_{ilp} = 4b$ . FIFO and serializers are used to address the high-speed throughput requirement of SiP accelerators.

## D. Efficiency trade-off factors

It can be inferred from Eq. (7) that the required input laser power increases as a function of  $N$  as a result of the exponentially increasing optical losses in the MZM-based accelerator. However, the total energy efficiency [Eq. (6)] starts improving as  $N$  scales up due to the quadratic increase in the number of accelerated operations performed by the optical matrix, as shown in Fig. 5. Taking all optical losses into account shows that there is a scaling limit beyond which optical losses grow significantly and the overall efficiency drops and an optimal network size exists for minimum energy efficiency. Unfortunately, the maximum network scale,  $N_{ltd}$ , is limited by the rated output optical power of the laser and the signal-to-noise ratio (SNR) required for any given signal resolution,  $n_{ilp} \geq 1b$ , as given in Eq. (8) and illustrated in Fig. 3.

Figure 6 shows the total energy efficiency and scaling limit for various input resolutions considering thermo-optic phase shifters with and without insulation. The energy efficiencies in Fig. 6 are calculated for accelerators to be operated at binary and higher resolution,  $n_{ilp} = \{1, 2, 3, 4\}b$ . Although the probability of transition reduces for multilevel signaling,<sup>5</sup> the requirement on driver's linearity or segmentation also increases. Furthermore, the energy consumed in the serializing and clocking remains the same.<sup>5</sup> Thus, we assume similar energy efficiency for multilevel signaling as PAM2,  $\sim 2$  pJ/b, for MZM modulators.<sup>62</sup> For binary resolution, the power consumed by the drivers and AFEs is extracted from recent work on PAM2 transceivers.<sup>5,59,62,63</sup> Therefore, the energy efficiency for 2b, 3b, and 4b input MZM drivers is estimated in our calculation as  $\sim 4$ , 6, and 8 pJ per symbol, respectively.

It can be concluded from Fig. 6 that opting for PDs with higher responsivities improves the energy efficiency of the network. This





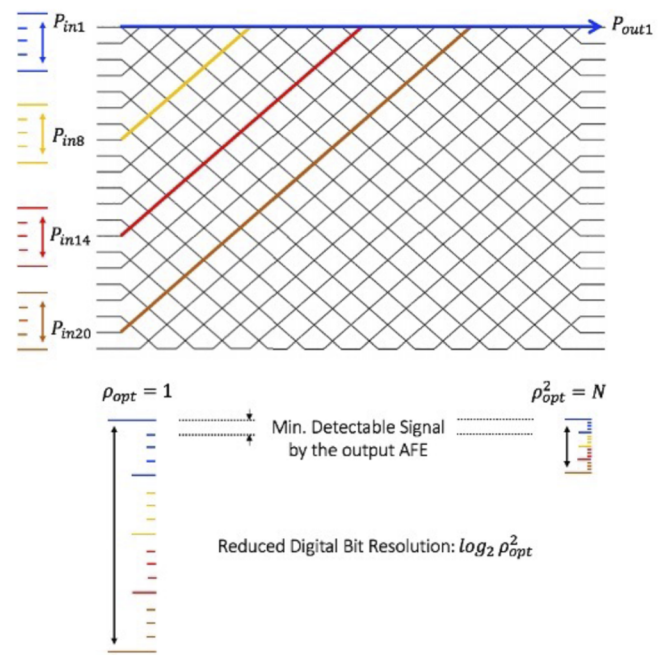
FIG. 5. Total energy efficiency (pJ/Op) for an  $N \times N$  MZM implementation with PD responsivity  $R = 1.2$  A/W and binary resolution,  $n_{ilp} = 1b$ . Energy efficiency improves as the network scales up due to the increased number of operations. Improving the tuning efficiency with insulation has significant enhancement on the overall energy efficiency.



**FIG. 6.** Total energy efficiency and scaling limit of an MZM-based network with an input resolution of  $n_{ip} = \{1, 2, 3, 4\}b$  considering thermo-optic phase shifters with and without insulation at the scaling limit,  $N_{ltd}$ , where the laser rated power output (10 dBm) is reached. Implementing networks with higher resolution requires scaling down the network to improve the SNR at the AFE. This leads to a degradation of the energy efficiency due to the reduced number of operations.

compensates for the optical system loss, relaxes the need to inject high optical power at the network input, and improves the overall energy efficiency. Utilizing avalanche PDs (APDs) is a possible way to significantly improve the optical sensitivity.<sup>64</sup> Figure 6 also suggests that taking advantage of the loss of precision, when possible, shows minor improvement in the energy efficiency and the network scaling (Table II).

Considering a matrix that scales with  $N$ , the laser optical intensity should be typically scaled by a factor of  $N$  to account for the splitting loss in a lossless network. For mesh-like configurations similar to Fig. 1, scaling the input vector size increases the dynamic range of the output intensities. In other words, for a given matrix



**FIG. 7.** The loss of precision factor for a case when four input signals ( $2b$  digital resolution) combine to form the output  $P_{out1}$ . Scaling down the input optical power by  $N = 4$  can be achieved at the expense of reducing the digital precision at the output by  $\log_2 4 = 2b$ .

output, the intensity can be as low as that of a single input or as high as  $N$  times that amount. With the input's digital resolution being  $n_{ip}$ , the effective overall output resolution due to the network scaling is  $n_{ip} + \log_2 N$ .

The conservative estimate of scaling the input power by  $N$  may not be necessary in some computational context, such as convolutional neural network (CNN) layers with adaptable hidden layer resolutions,<sup>65</sup> the increased output resolution might be higher than that needed by the AFE to detect. Therefore, an energy scaling vs loss of precision trade-off factor,  $\rho_{opt}$ , can be introduced to take advantage of the network scaling,<sup>66,67</sup> as illustrated in Fig. 7. Full accuracy is described by  $\rho_{opt} = 1$  corresponding to reduced output precision of  $\log_2(\rho_{opt}) = 0$ , at which the input optical intensity is scaled by  $N$ . Generally, for  $\log_2(\rho_{opt})$  bit reduction, the input is scaled by  $N/\rho_{opt}$ . Therefore, the maximum amount of energy saving is achieved when the  $\log_2 N$  bit reduction is tolerable at the optical output (AFE input).

To get a meaningful sense of the trade-off between the energy scaling and the loss of precision,  $\rho_{opt}$  is quantified in the following equation in terms of the probability of bit errors for binary networks

**TABLE II.** MZM-based implementation characteristics.

$\eta_{WPE}$	$IL_{SMF}$ (dB)	$IL_{EC}$ (dB)	$IL_{WG}$ (dB/mm)	$EL_{Splitter}$ (dB)	$IL_{MZI}$ (dB/mm)	$L_{MZI}$ (mm)	$IL_{DC}$ (dB)
0.1	0	1.6	0.3	0.01	1	0.5	0.01

(networks with binary weights) at the output such that

$$Prob_{error} = Q\left(\frac{P_{opt-o/p}}{\rho_{opt}} \frac{R}{2i_{irn}}\right), \quad (9)$$

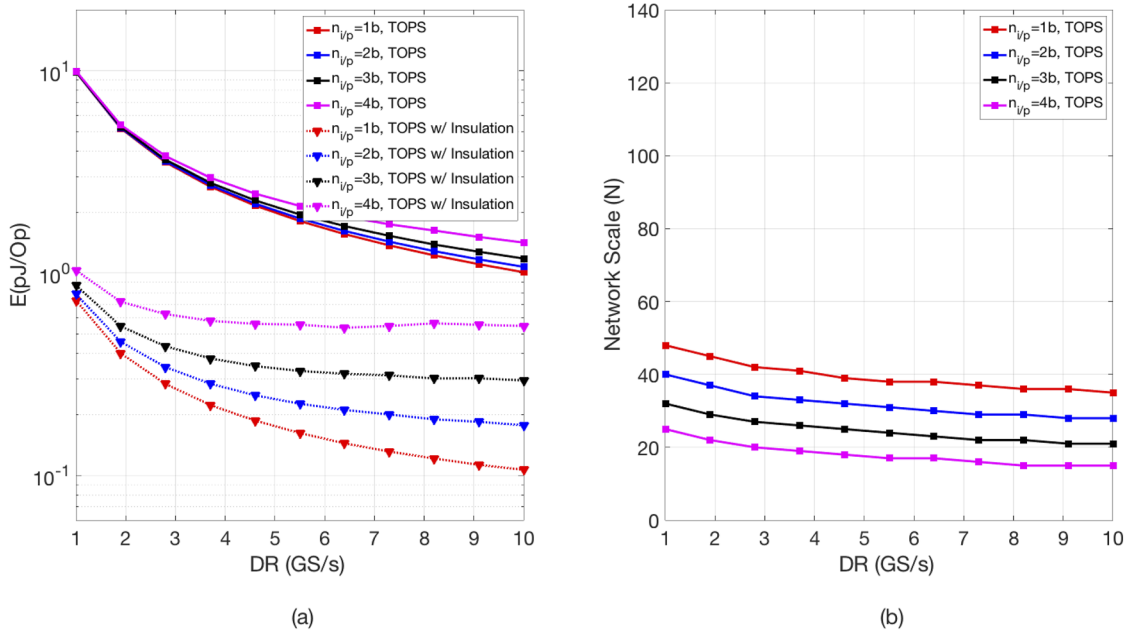
where the  $Q$  function is defined as  $Q(x) = \int_x^\infty \frac{1}{\sqrt{2\pi}} e^{-u^2/2} du$  and  $i_{irn}$  represents the total input referred noise at the AFE input with contributions from the PD, TIA, main amplifiers, and comparators (if applicable). Therefore, the laser power,  $P_{laser} = NP_{opt-o/p}$ , can be traded off for loss of output resolution.

The IL difference between the bar and cross states of a tunable beam splitter impacts the interference between the nodes in the mesh. For an MZM with intensity loss of  $\alpha_1$  in one arm and  $\alpha_2$  in the other, it can be shown that the output intensity at the cross state is given by  $I_{cross} = I_{in1}[\alpha_1 + \alpha_2 + 2\sqrt{\alpha_1\alpha_2}\cos(\theta_2)]$  when  $I_{in2}=0$ , where  $I_{in1}$  and  $I_{in2}$  represent the MZM input intensities at its input ports 1 and 2, respectively. To get the transmission response of an MZM with equal losses,  $\alpha_1$  on both arms,  $\theta$  should be modified such that  $\cos(\theta) = \frac{\Delta\alpha + 2\alpha_1\cos(\theta_1)}{2\sqrt{\alpha_1\alpha_2}}$  in order to account for the loss difference between the two MZM arms, where  $\Delta\alpha = \alpha_1 - \alpha_2$  and  $\theta_1$  describes the phase shift when both arms have attenuation of  $\alpha_1$ . This difference in insertion losses can be observed in single-arm beam splitters in which phase shifters are controlled by a single arm only. Using dual-arm tunable beam splitters, where  $\theta$  is implemented differentially using phase shifters on both arms, introduces equal insertion losses for the bar and cross transmissions of each node. A dummy phase shifter can also be used in single-arm topologies to obtain equal losses.

For sake of comparison, the energy consumption for a digital MAC is estimated based on the energy consumed by multiplication and accumulation operations and register file access in a 28 nm CMOS implementation.<sup>68</sup> With an estimated energy consumption of 0.046 pJ for 8b MAC and 0.0117 pJ for register file access, the calculated energy consumption is  $\sim(0.046 + 0.0117 \text{ pJ})/2 = 28.85 \text{ fJ}$  for a single operation. Conversely, it can be observed from Fig. 6 that SiP networks based on MZMs need to be scaled down to achieve higher resolutions, which further degrades their energy efficiency. Compared to their 8b digital CMOS counterpart, the energy efficiencies for SiP MZM MAC operations (using low-power thermo-optic phase shifters with insulation and 1.2 A/W<sup>69</sup> PD responsivity) at  $n_{ilp} = 1b\text{--}4b$  resolutions are  $3.5\times$  to  $17.5\times$  worse. Despite the lower energy efficiency, MZM-based MAC operations are performed at  $2N\times$  higher operating speed and lower latency (for a weight-stationary systolic array with an input vector size of  $N \times 1$ ) than the corresponding digital CMOS implementation.

In addition, multiple clock cycles are needed for digital multipliers and adders to provide the MAC output in systolic arrays. Operating at  $\sim 10\times$  lower clock speeds further decreases their throughput in comparison to optical implementations. Assuming a number  $\alpha$  of clock cycles needed for digital MAC, the latency of a digital systolic array with a  $1 \times N$  input vector size and a  $N \times N$  matrix size is  $2N\alpha/f_{CLK-CMOS}$ , where  $f_{CLK-CMOS}$  represents the clock frequency of the digital CMOS implementation. Hence, the throughput ratio of the optical to CMOS implementations is  $2N\alpha f_{CLK-OPT}/f_{CLK-CMOS}$ , where  $f_{CLK-OPT}$  represents the clock speed at which an optical implementation operates at.

We also investigate the energy efficiency and network size at lower data rates in Fig. 8. Intuitively, the energy efficiency degrades



**FIG. 8.** (a) Total energy efficiency and (b) MZM network sizes for  $R = 1.2 \text{ A/W}$  and  $DR = 1\text{--}10 \text{ GS/s}$ . Operating at higher data rates helps reduce the contribution of the static power consumption of thermo-optic phase shifters to the energy efficiency. Network scales up at lower data rates due to the SNR improvement.

at lower data rates because less number of operations is conducted with respect to the dissipated static power. On the other hand, the network size, shown in Fig. 8, can be increased by making use of the SNR improvement at lower data rates, as inferred from Eq. (8). If maximizing the optical throughput is not an overarching goal, opting for lower data rates (relative to 10 GS/s) leads to larger networks while not sacrificing much on the energy efficiency in implementations incorporating phase shifters with insulation [which do not consume much static power, as shown by the dotted curves in Fig. 8(a)].

### III. MRR-BASED SI-PHOTONIC IMPLEMENTATION

#### A. System architecture

The SiP MRR-based implementation of an optical accelerator is illustrated in Fig. 9. A comb CW laser source is used to provide wavelengths  $\lambda_1$  through  $\lambda_n$ , which are coupled into the chip and then modulated by an array of  $N$  MRMs. Unlike mesh-like topologies, implementing vector matrix multiplication in the form of dot products has the advantage of maintaining equal path loss for all the outputs. The modulated input vector,  $X_{i/p}(\lambda)$ , is then split (broadcast) into  $N$  branches to be modulated by the weight bank arrays;<sup>70</sup> each output represents the dot product of the input vector and one of the row arrays of the weight matrix. In order to achieve weights with positive and negative polarities, the thru and drop transmissions of the weight arrays are routed to balanced PDs in a push-pull configuration at the receiver. The current difference at the output,  $Y_{O/p}$ , can be represented as<sup>14</sup>

$$Y_{O/p} = \int_{-\infty}^{\infty} |E_0(\lambda)|^2 X_{i/p}(\lambda) W_{dt}(\lambda) R(\lambda) d\lambda, \quad (10)$$

where  $E_0(\lambda)$ ,  $W_{dt}(\lambda)$ , and  $R(\lambda)$  represent the amplitude of the input optical field, the difference between the weight's drop and thru intensity transmissions, and the PD responsivity at a wavelength  $\lambda$ , respectively.

#### B. Optical network link budget

Similar to Eq. (5), the optical link budget is calculated based on the following equation:

$$\begin{aligned} P_{O/p}(\text{dBm}) = & P_{\text{laser}} - P_{\text{SMF-att}} - P_{\text{EC-IL}} - P_{\text{Si-att}} \\ & - P_{\text{MRM-I/p-IL}} - (N-1)P_{\text{MRM-I/p-OBL}} \\ & - P_{\text{splitter-IL,EL}} - P_{\text{MRR-W-IL}} \\ & - (N-1)P_{\text{MRR-W-OBL}} - P_{\text{penalty}}, \end{aligned} \quad (11)$$

where  $P_{\text{MRM-I/p-IL}}$  represents the transmission insertion loss of the MRM for the input vector,  $P_{\text{MRM-I/p-OBL}}$  represents the out-of-band insertion loss (OBL) of the MRM for the input vector when the MRM resonance wavelength does not match the input vector wavelength,  $P_{\text{MRR-W-IL}}$  represents transmission insertion loss of the MRR for the weight vector, and  $P_{\text{MRR-W-OBL}}$  represents the out-of-band insertion loss of the MRR for the weight vector. Other terms have been defined already when describing Eq. (5).

Figure 10 shows the calculated optical power throughout an MRR-based implementation with different input vector sizes,  $N$ , using the values shown in Table III. The optical power of the laser is set to 0 dBm for the ease of illustration. Similar to MZM-based implementations, the attenuation introduced due to the splitting and cascading of microrings significantly degrades the optical power and poses a limitation on the energy efficiency, as will be discussed in Sec. III C. Figure 11 shows the optical intensities required at the AFE to detect a signal with a resolution of  $n_{i/p}$  bit. This is obtained by representing the desired output signal and current noises in terms of the received optical intensity, as given in Eq. (8). It can be shown that the maximum achievable matrix size is  $\sim 85 \times 85$  for binary networks. We revisit this calculation again in Sec. III D.

#### C. Energy efficiency

The energy efficiency (J/Op) of the MRM-based implementation with size  $N \times N$  operating at a data rate of  $DR$  can be calculated

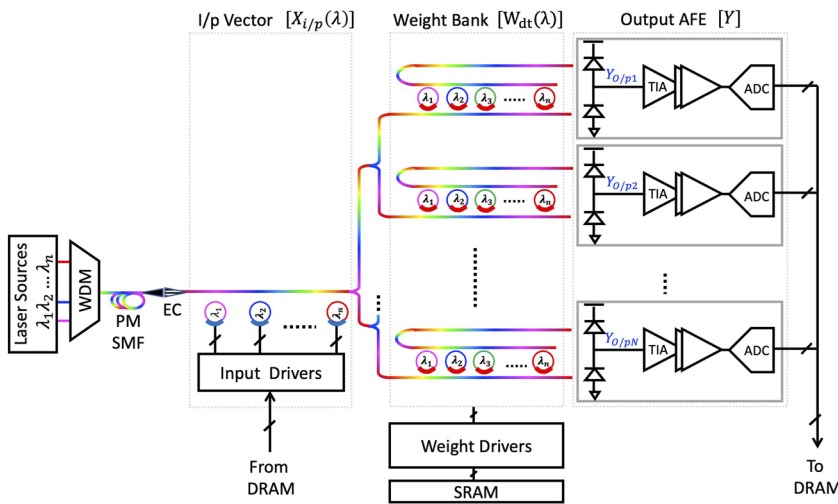


FIG. 9. SiP circuit diagram of an  $N \times N$  MRR-based accelerator.



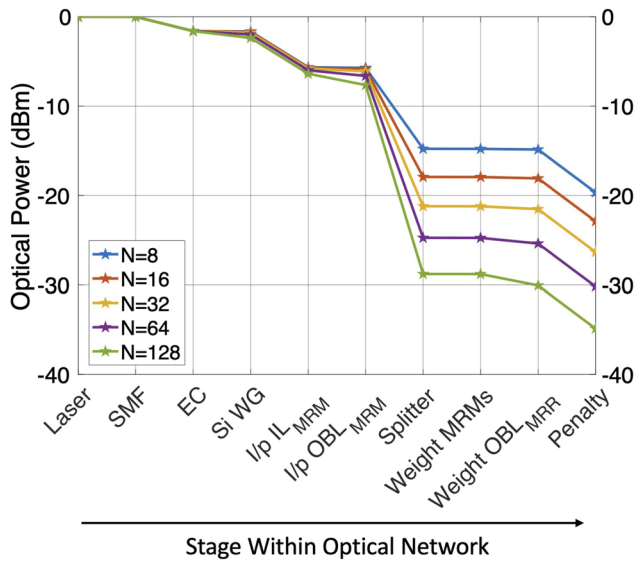


FIG. 10. Optical link budget analysis for MRR-based matrix sizes,  $N = 8, 16, 32, 64$ , and  $128$  for  $R = 1.2 \text{ A/W}$  and  $DR = 10 \text{ GS/s}$ .

as

$$E(\text{J/Op}) = \frac{P_{\text{laser}}}{\rho_{\text{SOA}} \cdot 2N^2 \cdot DR} + \frac{NP_{\text{i/p-drivers}} + 2P_{\text{mem-interface}}}{2N^2 \cdot DR} + \frac{NP_{\text{mat-tuning}} + P_{\text{SOA}} + P_{\text{o/p-AFE}}}{2N \cdot DR}. \quad (12)$$

Here,  $P_{\text{laser}}$  represents the total electrical power consumed by the optical source (either a single comb laser source or multiple sources generating all the desired input wavelengths). To better study the energy efficiency based on a targeted signal resolution,  $n_{\text{ip}}$ , the power consumption of the input laser source is formulated as a

TABLE III. MRR-based implementation characteristics.

Parameter	Value
$\eta_{\text{WPE}}$	0.1
$IL_{\text{SMF}}$ (dB)	0
$IL_{\text{EC}}$ (dB)	1.6
$IL_{\text{WG}}$ (dB/mm)	0.3
$EL_{\text{Splitter}}$ (dB)	0.01
$IL_{\text{MRM}}$ (dB) <sup>71</sup>	4
$OBL_{\text{MRM}}$ (dB)	0.01
$IL_{\text{MRR}}$ (dB)	0.01
$d_{\text{MRR}}$ ( $\mu\text{m}$ )	20
$IL_{\text{penalty}}$ (dB)	4.8




FIG. 11. Targeted AFE sensitivity for  $n_{\text{ip}} = \{1, 2, 3, 4, 5, 6\}b$  and the output power for MRR-based SiP matrices with sizes of  $N = 8, 16, 32, 64$ , and  $128$ .  $P_{\text{laser}} = 10 \text{ dBm}$ .


function of the optical power reaching the PDs at the output AFE, as given in the following equation:

$$P_{\text{laser}} = \frac{10^{\frac{\eta_{\text{WG}}(\text{dB}) N(d_{\text{MRR}})}{10}} N}{\eta_{\text{SMF}} \eta_{\text{EC}} IL_{\text{i/p-MRM}} (OBL_{\text{MRM}})^{N-1} (EL_{\text{splitter}})^{\log_2 N}} \times \frac{P_{\text{PD-opt}}}{\eta_{\text{WPE}} IL_{\text{weight-MRR}} (OBL_{\text{weight-MRR}})^{N-1} IL_{\text{penalty}}}, \quad (13)$$

where  $d_{\text{MRR}}$  represents the gap between the centers of two adjacent microrings and is dictated by the thermal crosstalk, which should be taken for design considerations. A  $d_{\text{MRR}}$  of  $15 \mu\text{m}$  has been shown to be sufficient to avoid thermal crosstalk in a photonic switch implementation.<sup>72</sup> We assume a  $d_{\text{MRR}}$  of  $20 \mu\text{m}$  in this work for an optimistic realization of the system with  $R_{\text{MRR}} = 6 \mu\text{m}$ .

Since the number of operations scales quadratically with the weight vector size,  $N$ , while the energy consumption of the modulators and AFE increases linearly as can be seen in (12), the overall energy efficiency improves with scaling, as shown in Fig. 12. We assume similar MRM energy efficiency for multilevel signaling as PAM2,  $\sim 0.3 \text{ pJ/b}$ ,<sup>5</sup> which takes into account both the contribution of the modulator driver and the serializers. Therefore, the energy efficiency for  $2b$ ,  $3b$ , and  $4b$  input MRM drivers is estimated in our calculation as  $\sim 0.6, 0.9$ , and  $1.2 \text{ pJ}$  per symbol, respectively.

MRMs offer smaller footprint and lower input capacitance, which leads to a significant reduction in their driving power. However, they are also more sensitive to fabrication mismatch and thermal drift, which entails the need to use heaters for calibration across a wide spectral range. Excluding heater power, the power consumed by a closed loop controller implemented for a low-power wavelength division multiplexing (WDM) topology is  $\sim 0.2 \text{ mW}$ .<sup>73</sup> The average energy efficiency for the state-of-the-art MRR heaters on an SOI platform is  $\sim 20 \text{ mW}/\pi$ .<sup>7,74,75</sup> The scaling of the power consumed by the heaters with  $O(N^2)$  degrades the overall energy efficiency of the network. Figure 13 shows the total energy efficiency



**FIG. 12.** Total energy efficiency (pJ/Op) for an  $N \times N$  MRM implementation with PD responsivity  $R = 1.2$  A/W and binary resolution,  $n_{ip} = 1b$ . Energy efficiency improves as the network scales up due to the increased number of operations. Improving the tuning efficiency with insulation has significant enhancement on the overall energy efficiency.


and scaling limit for various input resolutions considering thermo-optic phase shifters with and without insulation. We assume power consumption values,  $P_{heater}$ , of 2.8<sup>75</sup> and 40 mW<sup>7</sup> for phase shifters with and without insulation, respectively, to provide a phase shift of one free spectral range (FSR).

Driving inputs at high speeds entails the need to multiplex data fetched from the memory, as shown in Fig. 4. As per the calculations in Sec. II C,  $P_{mem-interface}$  is taken as 5.77 mW for the energy efficiency calculation of input or output memory interfacing circuits.

#### D. Scaling limitations

It can be inferred from Fig. 13 that it is feasible to implement vector matrix multiplication using MRR-based networks with sizes scaling up to  $N = 85$ . For  $n_{ip} = 2b$  or above, the network size is within the maximum number of microrings permitted for WDM implementations due to FSR limitations and crosstalk.<sup>76</sup> Attempting to engineer the MRR's dimensions and coupling ratio compromises the quality factor, which degrades the channel spacings. This translates to a limit in the MRR vector size of  $N < FSR/\Delta\lambda$ . Channel spacings are typically set according to the amount of acceptable crosstalk between channels (interchannel interference). As an example, for a 50 nm transmission window with channel spacings of 0.8 nm, the maximum number of channels is 62.<sup>76,77</sup> Thus, for  $n_{ip} = 1b$ , FSR may set a limitation to the overall network size.

Using series coupling to increase the filter order has been experimentally shown to reduce both interchannel and intrachannel crosstalks, thus maximizing the filter finesse and the channel count.<sup>77</sup> However, this comes at the expense of higher footprint,

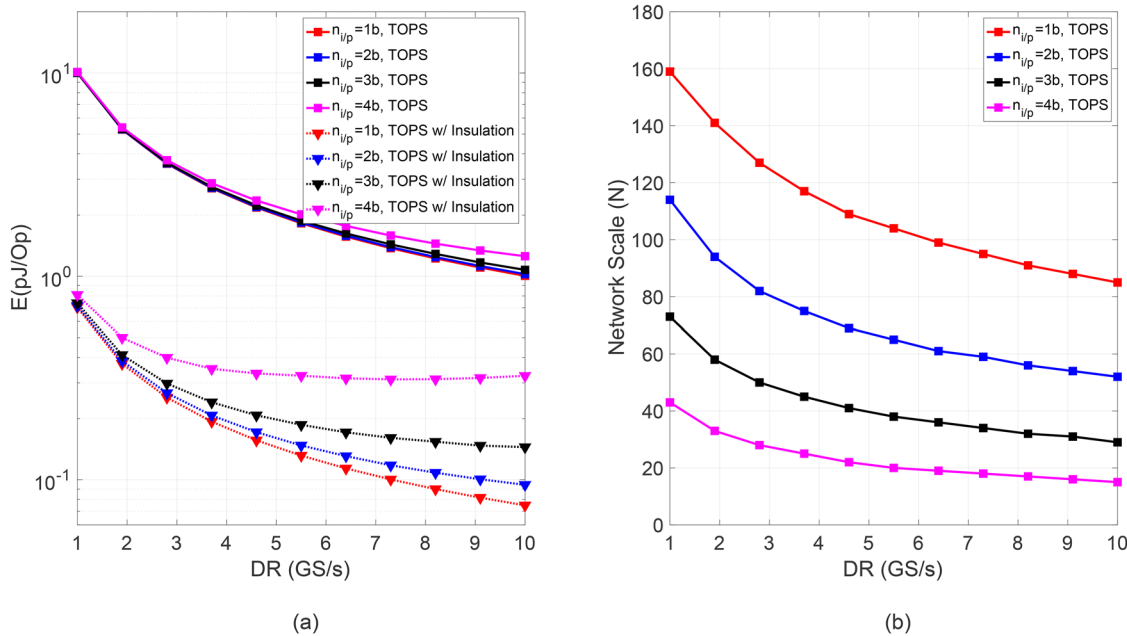


**FIG. 13.** Total energy efficiency and scaling limit of an MRM-based network with an input resolution of  $n_{ip} = \{1b, 2b, 3b, 4b\}$ , considering thermo-optic phase shifters with and without insulation at the scaling limit,  $N_{tld}$ , where the laser rated power output (10 dBm) is reached. Implementing networks with higher resolution requires scaling down the network to improve the SNR at the AFE. This degrades the energy efficiency due to the reduced number of operations.

lower drop port transmission, and extra tuning power. To cascade several MRRs for MAC operations, it is necessary to maintain channel spacings to avoid the adjacent weight-dependent crosstalk. The number of channels that can be supported by optimized MRRs with finesse of 368 and 540 are calculated to be 108 and 148, respectively.<sup>76,78,79</sup>

The two-point coupling scheme has been proposed to address the post-fabrication correction of MRM spectral features for large-scale MRM implementations.<sup>80</sup> Although it mitigates the secondary resonances of an MRM and doubles the FSR, an extra micro-heater is introduced to correct for the coupling, which increases the power consumption. Another attempt to achieve an FSR-free filter has been demonstrated using tunable couplers along with modified vernier filters that use higher-order coupled MRRs.<sup>81</sup> However, this topology is associated with penalty in terms of design complexity, increased footprint, and tuning power.

Introducing contra-directional coupling (CDC) in a microring combines the wavelength selectivity of the CDC with the compact feature size of the MRR, thus reaping the advantages of both and providing an FSR-free response.<sup>82</sup> Implementing this design technique allows for the potential use of several channels in MRR-based



**FIG. 14.** (a) Total energy efficiency and (b) MRM network sizes for  $R = 1.2$  A/W and  $DR = 1$ –10 GS/s. Operating at higher data rates helps reduce the contribution of the static power consumption of thermo-optic phase shifters to the energy efficiency. Network scales up at lower data rates due to the SNR improvement.

accelerators. This comes with a trade-off of using extra heaters in the CDC and in the region of the MRR that does not include corrugated structures.


As shown in Fig. 13, the optimum energy per operation of binary SiP networks based on MRRs ( $\sim 75$  fJ) is obtained at  $N = 85$  for a PD responsivity of  $R = 1.2$  A/W. It can also be shown that reducing the power consumption of weight tuning circuits by one order of magnitude improves the energy efficiency by roughly one order of magnitude as well. Compared to their 8b digital CMOS counterpart, the energy efficiencies for SiP MRM MAC operations (using low-power thermo-optic phase shifters with insulation and 1.2 A/W PD responsivity) at  $n_{ip} = 1b$ –4b resolutions are  $2.6\times$  to  $13\times$  worse. In comparison to MZM-based implementations, MRR-based implementations can have  $1.8\times$  bigger network scale and achieve  $1.3\times$  lower energy consumption per operation. Similar to MZM-based implementations, MRR MAC operations are performed at a  $2Naf_{CLK-OPT}/f_{CLK-CMOS}$  higher throughput than its digital CMOS counterparts.

Figure 14 shows the energy efficiency and scaling at lower data rates. Similar to MZM implementations, reducing the data rate degrades the energy efficiency while scaling up the network size due to the reduced noise levels at the AFE.

#### IV. RESEARCH OPPORTUNITIES

As summarized in Secs. II and III, SiP accelerators operate at much higher speed and lower latency than their CMOS counterparts. Nevertheless, it is further desired to improve the size of the MAC networks in SiP, especially for neural network

applications, and the energy efficiency. There have been several promising research studies in the field of SiP. Classifying the existing commercial SiP technology as the first generation, we describe several emerging technologies that will make up the next generation of SiP. Figure 15 summarizes the advancements in SiP that can be leveraged by SiP-based accelerators to reduce optical loss, improve the energy efficiency, and incorporate heterogeneous integration techniques for performance improvement.



**FIG. 15.** Evolution of SiP technology from current (1.0) to the next generation (2.0). Evolution of the existing SiP technologies—which comprise of grating couplers (GCs), Edge Couplers (ECs), V-grooves, silicon waveguides, thermal heaters, sub-wavelength gratings (SWGs), PN junction modulators, germanium photodetectors (Ge PDs), in-resonator photoconductive heaters (IRPHs), and silicon nitride (SiN) escalators—to include emerging technologies, such as photonic wire bonds (PWBs), low-loss waveguides, liquid crystal on silicon (LCOS), nano-opto-electromechanical systems (NOEMS), semiconductor optical amplifiers (SOAs), polymers, phase change materials (PCMs), indium tin oxide (ITO), and avalanche photodetectors (APDs).

### A. Optical loss reduction

As described in Secs. II and III, optical losses limit the scalability of the SiP 1.0 networks. Losses must be minimized at the coupling interfaces and in the components. PWB is one way to ensure efficient coupling between the chip and the optical fiber with insertion loss  $\sim 1$  dB with negligible variation.<sup>35</sup> Passive alignment to SMF optical fibers can be accomplished using V-grooves arrays. Such fiber to chip self-alignment has been shown to have a coupling efficiency of  $\sim -1.3$  dB.<sup>83</sup> In another demonstration, coupling losses as low as  $\sim 0.5$  dB and  $\sim 0.35$  dB have also been reported for passive and active alignments, respectively.<sup>84</sup>

For scaling up the networks, the optical signal attenuation can be compensated by using SOAs. On-chip SOAs can be utilized to pre-amplify the input signal and also exploited as weight matrix elements to provide weights magnitudes  $>1$ .<sup>85</sup> However, the non-linear gain-current curve entails a need for calibration.

Improving the responsivity of the AFE is yet another way to tolerate the optical losses. It relaxes the need to increase the laser power to compensate for the losses. The high multiplication gain and responsivity of APDs have been shown to improve the sensitivities of optoelectronic receivers front-end.<sup>64</sup> Improving dark current and quantum efficiency by a careful design of the APD geometry has been projected to improve the sensitivity of Si-Ge APD receivers of up to  $-29$  dBm at 12.5 Gb/s<sup>86</sup> as compared to  $-18.5$  dBm for Ge PIN detectors.<sup>87</sup> Limiting the bandwidth of the AFE and using equalization techniques,<sup>88</sup> such as continuous time linear equalization (CTLE) and decision feedback equalization (DFE),<sup>89</sup> can reduce the input referred noise of the AFE and further improve the sensitivity.

### B. Improving energy efficiency

Commercial CW lasers suffer from low WPE in the range of  $\sim 1\% - 10\%$ , which impacts the energy efficiency on the system.<sup>90,91</sup> Hybrid-integrated silicon photonic lasers have been shown to provide  $\sim 12.2\%$  WPE.<sup>92</sup>

Although introducing on-chip CW lasers mitigates the coupling losses, the feasibility of using them, especially for networks using WDM, requires wavelength stabilization and reflection cancellation.<sup>93,94</sup>

Reducing the power consumption of the phase shifters in the weight matrix is a critical requirement, given that their overall energy consumption scales quadratically with network size [Eqs. (6) and (12)]. Thermo-optic phase shifters dissipate high-power consumption, given their resistive nature. Introducing trenches, undercuts, and back-side substrate removal has been shown to improve the tuning efficiency of the rings by an order of magnitude with a measured reported power consumption of  $\sim 4$  mW per FSR.<sup>75,95,96</sup> However, thermal isolation and substrate removal exacerbate self-heating and must be taken into consideration while designing a CMOS controller.<sup>97</sup>

Several post-fabrication schemes have been investigated to correct for the fabrication-induced variations. Reducing the process variations was investigated by patterning SiN on top of the Si waveguide to introduce field perturbations, which effectively adjusts the optical path length.<sup>98</sup> Another demonstrated technique relies on trimming using Ge ion implantation followed by laser annealing to tune the MRR resonant wavelength across the whole FSR without

introducing any excess loss. Its accuracy, CMOS-compatibility, and feasibility for wafer-scale correction render it a potential technique to be utilized in optical neuromorphic implementations to reduce the tuning power.<sup>99</sup>

Alternatives such as nano-opto-electro-mechanical systems (NOEMS)<sup>16,100</sup> and liquid crystal on silicon (LCOS)<sup>101</sup> have the potential to reduce the tuning power overhead significantly. The dynamic energy consumption of NOEMS was reported in the range of 0.13 and 0.32 fJ for digital pulse signals<sup>100</sup> and is assumed to be  $\sim 1$  fJ for our study. On the other hand, LCOS has been shown to dissipate power as low as 2 nW.<sup>101</sup>

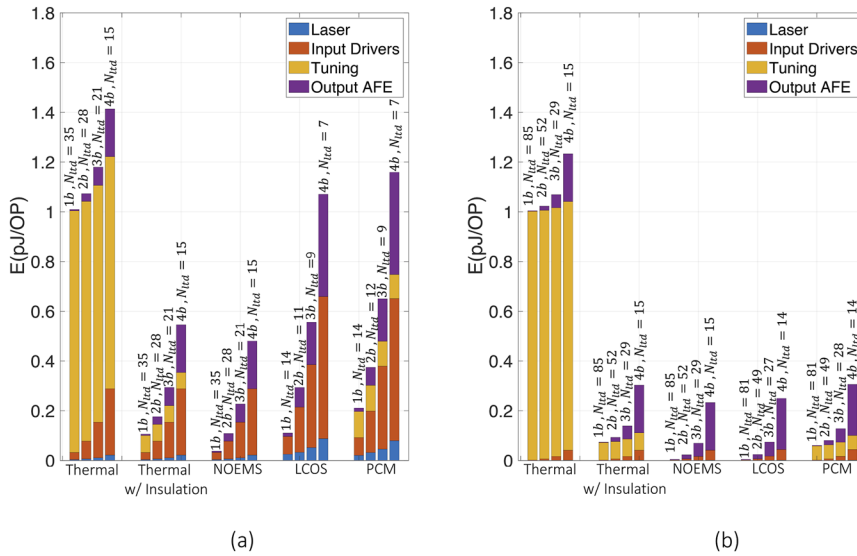
Phase change materials (PCMs) such as  $\text{Ge}_2\text{Sb}_2\text{Se}_4\text{Te}$  (GSST) have been demonstrated as compact phase shifters in which the optical phase shift is obtained by tuning the state of the material from amorphous and crystalline.<sup>102,103</sup> Being able to sustain their crystallization state with the absence of power renders them as good candidates for tuning low-speed weights in SiP implementations with no static power consumption. Given their compact sizes and non-volatile nature, the efficiency of implementing them in large-scale SiP networks is investigated, as shown in Fig. 16. Although not as lossy as PN phase shifters, PCMs have  $IL = 0.32$  dB, which is relatively high for cascaded phase shifters in a large-scale implementation.<sup>102</sup> This limits the network sizes for computation with several bits of resolutions. The resolution of the weights can be set by adjusting the level of crystallization of a PCM cell.<sup>104</sup>

The pulse energy consumption for writing and erasing levels 1–7 was reported in the range of 372–601 and 562–373 pJ,<sup>104</sup> respectively. Assuming the weights to be uniformly distributed, the average energy consumption,  $E_{PCM}$ , for setting the PCM to various weights can be calculated, as in Eq. (14), where  $E_A$  and  $E_C$  represent the pulse energy required to write (amorphization) and erase (crystallization) the first level ( $L_1$ ), respectively. For levels,  $L_i$ , where  $i > 1$ ,  $\Delta E_A$  and  $\Delta E_C$  represent the average amount of energy required to transition to one level higher or a lower, respectively, with all levels assumed to be equally spaced for the sake of simplicity,

$$E_{PCM} = \frac{2^n - 1}{2^{2n}} (E_A + E_C) + \frac{(1/3)(2^{2n} - 1)2^{n-1} - (2^n - 1)}{2^{2n}} (\Delta E_A + \Delta E_C). \quad (14)$$

Assuming that  $E_A = 372$  pJ,  $E_C = 373$  pJ,  $\Delta E_A = (601 - 372)/(2^n - 2)$  pJ, and  $\Delta E_C = (562 - 373)/(2^n - 2)$  pJ, the estimated average energy consumption for a PCM phase shifter with  $n = \{1, 2, 3, 4\}$   $b$  equals  $\{186, 231, 165, 121\}$  pJ. For phase shifters with zero static power dissipation, the dynamic energy consumption is divided by the number of times weights have been reused for vector matrix multiplication. This is done by introducing a weight reuse factor,  $\alpha_w$ , such that  $P_{mat-tuning} = P_{NOEMS,PCM}/\alpha_w$ , where  $\alpha_w$  typically ranges between  $2^6$  and  $2^{18}$  in general matrix multiplications (GEMMs).<sup>105</sup> A value of  $\alpha_w = 4096$  is chosen for the calculation of energy efficiencies in this work. For networks where the weight reuse is low, the contribution of the dynamic energy per operation can be considerably higher for PCM than all the other weight tuning alternatives, degrading the energy efficiency by orders of magnitude.

Figure 16 shows the energy efficiency breakdown for both the MZM-based and MRR-based architectures for weight tuning that



**FIG. 16.** Energy efficiency breakdown considering various weight tuning options for bit resolutions:  $n_{ip} = \{1, 2, 3, 4\}b$  for (a) an MZM-based implementation and (b) an MRM-based implementation. Phase shifters with low insertion loss and static power consumption (e.g., TOPS with insulation and NOEMS) are good candidates to enhance the energy efficiency of SiP implementations. The high insertion loss of LCOS and PCM poses limitations on the network sizes for a target resolution. In addition, PCMs, with their large dynamic power consumption, are promising for MRM implementations, provided that a high weight reuse is possible.

rely on thermo-optic phase shifters without and with insulation,<sup>75</sup> NOEMS, LCOS, and PCM. For each implementation, energy calculations are based on the network scales that can satisfy the SNR requirements to compute with bit resolutions,  $n_{ip} = \{1, 2, 3, 4\}b$ . The insertion loss for a 35  $\mu\text{m}$  long LCOS used to realize a  $\pi$  phase shift is taken as 0.35 dB.<sup>101</sup> For MZM implementations, thermo-optic phase shifters with insulation seem currently attractive for energy efficiency and network size. For a large weight reuse factor, NOEMS-based phase shifters promise further energy reduction.

For MRM implementations, similar conclusions can be drawn except that LCOS-based phase shifters also seem promising.

For both MZM-based and MRR-based architectures, it is evident that opting for matrix weight tuning alternatives with almost zero power consumption significantly improves the total energy efficiency with values approaching <100 fJ/Op for both architectures. Further research is still needed to demonstrate the feasibility of these approaches in high volume production to realize such an energy efficiency regime,

$$\text{SNR} = \frac{(RGP_{o/p})^2}{\left[ \sqrt{\left( 2q(RGP_{o/p} + I_d) + \frac{4kT}{R_L} + 2\rho_{ASE}R^2GP_{o/p} + \rho_{ASE}^2R^2(2B_o - B_e) + R^2P_{o/p}^2RIN \right)} + \sqrt{\left( 2qI_d + \frac{4kT}{R_L} + \rho_{ASE}^2R^2(2B_o - B_e) \right)} \right]^2 B_e}. \quad (15)$$

### C. SOA cascadability

SOAs are used to amplify optical signals over a given spectrum and can be implemented off-chip or using hybrid integration. Cascading SOAs has been conventionally used to restore signal levels in interconnect links and has been recently explored in deep neural network implementations.<sup>85</sup>

Although SOAs help compensate for the optical losses to increase the scale of the network, they contribute significantly to the energy consumption. In addition, they suffer from several downsides, including the noises produced due to the optical amplification, the ripples in their gain spectrum, and amplification nonlinearity.<sup>106–108</sup> The major noise component is attributed to the amplified spontaneous emission (ASE) of photons toward the input and output of an SOA.<sup>108,109</sup> Therefore, the build-up of ASE noise due to the cascade degrades the SNR of the optical signal reaching the output.<sup>106,108</sup>

To investigate the effect of incorporating SOAs on the scalability of the network and the received signal resolution, the overall SNR is quantified in Eq. (15), with an SOA gain value as  $G = 17$  dB.<sup>46,110</sup>  $B_o$  and  $B_e$  stand for the optical bandwidth of the amplifier and electrical bandwidth of the AFE, respectively.  $\rho_{ASE}$  represents the ASE noise and is calculated as

$$\rho_{ASE} = 2N_{SOA}n_{sp} \frac{hc}{\lambda} (G - 1), \quad (16)$$

where  $N_{SOA}$  and  $n_{sp}$  represent the number of SOAs used in the network and the spontaneous emission factor of the optical amplifier, respectively.

The parameters given in Table I are used for calculating the SNR and  $N_{SOA}$ . For a target network resolution,  $n_{ip}$ , the number of SOAs is calculated based on the resultant SNR whose signal and noise power values vary with the network scale. The SNR is




calculated as follows:


$$n_{i/p} = \frac{\text{SNR (dB)} - 1.76}{6.02}. \quad (17)$$

As can be inferred from Figs. 2 and 10 for an SOA-less network, the SNR degrades since the signal optical intensity at the AFE  $P_{o/p}$  is attenuated with the scale of  $N$ . Incorporating an SOA helps replenish the signal intensity. SOAs can be added in the network before the SNR degrades below the threshold for a given resolution due to the insertion losses. However, the signal dependent noises are also amplified, which do not align in favor of the network resolution. For calculating the resolution, the AFE is assumed to tolerate signals with maximum  $P_{opt}$  intensity as high as 10 dBm, beyond which the number of SOAs is limited. Regions with no SOAs at the right-hand side of Figs. 17 and 18, shown in the Appendix, represent regions where the desired SNR cannot be achieved, indicating an infeasible resolution for a given scale.

SOAs require introducing III-V or II-VI compound semiconductor materials to the SiP platform, which is non-compliant with the standard SiP CMOS foundry runs. Back propagating ASE noise from the SOA emphasizes the need to employ an optical isolator for the input laser source.<sup>106</sup> Use of narrowband filters is also possible to reduce the out-of-band noise,<sup>85,106</sup> but these further reduce the maximum channel count, thus limiting the scaling of the network. To be used for linear analog computations, SOAs should deliver gains that are independent of the input intensities. Cross-gain modulation (XGM) is one type of non-linearity observed in SOA amplifiers in



**FIG. 17.** Network scale vs resolution for various numbers of SOAs for an MZM-based implementation. The scale of the network can be traded off with its resolution. Incorporating SOAs helps in extending the network scale for a fixed resolution or increasing the resolution for a given network scale. The number of SOAs that can be added to a network is limited to 1 or 2.



**FIG. 18.** Network scale vs resolution for various numbers of SOAs for an MRM-based implementation. Compared to their MZM counterpart, SOAs have bigger impact in scaling up the network for a given resolution.

which the combination of all the input intensities impacts the gain of a single channel.<sup>111,112</sup>

The aforementioned SOA limitations should be addressed in order to maintain the linearity of the weight matrix and allow for further scaling of networks. Designing highly efficient SOAs has the potential to increase the size of the networks. However, improving the network energy efficiency requires that the SOAs have low injection current and high optical signal-to-noise ratio. Recent attempts to use SOAs for optical networks have reported a power consumption of 42 mW per SOA, which leads to an energy consumption of  $\sim 4.2$  pJ/Op at a  $DR = 10$  GS/s.<sup>85</sup> To carry out four weighted additions, 16 SOAs were used for the weight tuning, along with extra SOAs for optical pre-amplification and input selection. A crosstalk of 0.6 dB was reported for the SOAs even for a small-scale circuit implementation of the arrayed waveguide grating (AWG) filter, which necessitated the use of feedback loops for gain calibration. For accelerators with SOA integration, the contribution to the overall network energy efficiency scales with  $O(N^2)$ , setting the efficiency to the pJ/Op regime.

## V. CONCLUSION

We describe the behavior of MZM- and MRM-based SiP implementations for MAC accelerators based on today's SiP technology. Both MZM and MRM implementations share similar optical and electrical challenges. In comparison to digital CMOS accelerators, SiP implementations have relatively higher energy consumption and operate at lower bit resolutions. In addition, they cannot be scaled to

large network sizes because of the optical losses. Implementing MAC operations using SiP has two distinct advantages:<sup>113</sup>

1. Optical MAC operations can be scaled to frequencies at tens of GHz, whereas MAC operations in digital CMOS are limited to a few hundreds of MHz or at most GHz operating speeds. For tasks where memory access is not the bottleneck, such as inference with fixed weights, an optical implementation can reduce latency and improve the energy efficiency at such high speeds. Digital CMOS counterparts, on the other hand, are limited by the clock frequency.
2. Multiplication operations can be intrinsically implemented in parallel in which the analog nature of the computation allows all matrix operations to take place at the same time for each input fetch.<sup>114</sup> Therefore, optical MAC implementations can increase the throughput and improve their energy efficiency at such high speeds. MAC operations implemented with digital circuits in CMOS are limited by the wiring interconnect density.

There are also some challenges of implementing MAC operations using SiP:

1. The losses in optical circuits severely limit the size of the MAC networks that can be physically realized, in comparison to a digital CMOS implementation where signal gain and regeneration are easily available. This, in turn, limits the applications of SiP MAC operations.
2. Although the power consumed in the optical MAC is less, when accounting for the losses and the power consumed by the laser and CMOS electronic circuits that drive and control the optical circuits, the energy efficiency is degraded.
3. Unlike digital CMOS implementations that can support 16b/32b resolution, analog photonic MAC operations have a maximum demonstrated resolution of 8.5b.<sup>115</sup> Nevertheless, such a resolution has been shown to be adequate for many inference tasks.<sup>116–118</sup>
4. Achieving a high throughput optical network entails accessing data at high speed. High-speed input/output (I/O) data can be streamed in/out from/to an off-chip DRAM; the corresponding energy consumption in moving the data must be considered for the overall implementation of an accelerator.<sup>119</sup> The energy consumption in data fetch from off-chip DRAM is significantly large. However, for a given dataset, the DRAM associated penalty is similar for both optical and digital CMOS implementations. We exclude that penalty in our work. For weight-stationary implementations where weights do not need to change frequently, an on-chip SRAM can be used, which can be adequately large due to the limited network size of the photonic accelerators.
5. Most of the low-loss phase shifters have a reconfiguration speed in the range of  $\mu$ s-to-ms, making the weight reconfiguration in photonic MAC operations significantly slower than their electronic counterparts. This limits the use of photonic MAC operations to weight-stationary systolic array implementations, where the incoming data are high-speed, but the weight does not get updated quickly. For other scenarios, a fine weight retuning can be done with high-speed plasma-dispersion phase shifters, where the loss

is controlled due to the need for a fine weight tuning range only.

6. To carry out optoelectronic computing with several bits of resolutions at high-speed, CMOS or biCMOS drivers and transimpedance amplifiers (TIAs) are needed that must operate with multi-level pulse-amplitude modulation (PAM) signaling. Although many PAM2 (1b) and PAM4 (2b) transceivers have been demonstrated,<sup>62,120</sup> higher levels of modulation require linear drivers and TIAs, which are challenging to design at high speed and good energy efficiency.<sup>3</sup> However, if data rates in optical computing are limited to a few tens of GBaud, this challenge is surmountable.
7. Packaging considerations in optics (e.g., laser, fiber, and SOA attach) are far more challenging than the packaging considerations for electronic dies due to alignment accuracy and thermal management requirements.<sup>121</sup>

The mesh-like interconnections of MZM-based implementations that extend the optical dynamic range at the AFE place a trade-off between the output resolution and the power requirement of the laser. MZM drivers also consume higher electrical driver power because MZMs cannot be made very long due to the losses associated with their larger footprints. Therefore, the power consumption of the laser and high-speed drivers is amortized with a limited network scalability. On the other hand, MRR-based topologies provide better energy efficiencies due to their small footprint and, thus, lower modulation energy consumption. The overall energy efficiency for either of the implementations experiences major degradation mainly due to the inefficiency in lasers and phase shifters, the insertion and excess losses of the optical components, and the optical to electrical and electrical to optical conversion overhead.

However, an order of magnitude higher operating speed and the inherent parallelism in conducting multiplication for analog signals render them attractive for reducing delay and enhancing throughput in comparison to digital CMOS implementations. With the emerging technologies in SiP, e.g., NOEMS and LCOS, low energy tuning schemes have the potential to significantly improve the energy efficiency of the photonic accelerators. Nonetheless, thermal PS with insulation is still an efficient weight tuning option, which is attractive for mass production. Low voltage-swing modulators with heterogeneous integration of polymers also promise improvements in energy efficiency due to the significant reduction in modulator and CMOS driver power consumption.

Scaling SiP accelerators to larger network sizes is limited by the rated output optical power of the laser and the SNR required for a given signal resolution,  $n_{i/p}$ . MRM-based networks can scale to larger values than their MZM-based counterparts due to the lower loss associated with cascading microrings in a WDM implementation. Incorporating high-power multi-wavelength lasers will be crucial. However, MRR-based networks are more sensitive to temperature and often require temperature control between the photonic IC and the laser. The size of MRR-based networks that have been demonstrated in prototype hardware has been limited to eight modulators<sup>122</sup> or  $16 \times 16$  switch.<sup>7</sup> Until larger MRM-based networks are demonstrated in hardware, the adoption of MZM-based networks will continue to be favored.

Heterogeneous integration of low-noise SOAs in SiP is a possible way to increase the network size, but the high-power consumption of SOAs degrades the energy efficiency significantly. Higher responsivity and low-noise APDs will also prove beneficial in scaling up the network sizes. The size can be further scaled up by reducing the insertion losses of contributors, such as directional couplers and phase shifters (if lossy). The splitters remain a significant limitation for scaling. To make efficient use of the limited optical network sizes, general matrix multiplication (GEMM) algorithms must be adopted.

Enhancing the energy efficiency can be achieved by adopting modulators with low static and dynamic power consumption, high-responsivity APDs along with TIAs with high sensitivity, and efficient SOAs and lasers with high WPE. To better address the need for high resolution, multi-level signaling significantly beyond PAM4 and PAM8 must be implemented. Controlling the temperature of the chip also helps with maintaining high resolution. In addition, the crosstalk and distortion of SOAs should be further investigated.

## ACKNOWLEDGMENTS

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). Access to CAD tools and technology is facilitated by CMC Microsystems. The authors would like to thank Dr. Alex Tait of Queen's University and Avilash Mukherjee of UBC for their technical comments.

## AUTHOR DECLARATIONS

### Conflict of Interest

The authors declare no conflicts of interest.

## DATA AVAILABILITY

The data that support the findings of this study are available within the article.

## APPENDIX: SCALING UP USING SOAs

This section illustrates the feasibility of incorporating SOAs in SiP networks in Mach-Zehnder and microring based implementations. Figures 17 and 18 illustrate the number of SOAs in terms of the network resolution and scale for MZM- and MRM-based implementations, respectively. It can be observed that the resolution at which a SiP network is desired to operate is dependent on the network scale.

It can be inferred that an MZM-based network can support signal resolutions of  $4b$  for a network size of  $N_{td} = 55$  with a single SOA. In comparison, incorporating SOAs into MRM implementations increases the network limited scale to  $N_{td} = 94$  for  $n_{ip} = 4b$ , as shown in Fig. 18. The number of SOAs that can be added to a network is limited to 1 or 2.

## REFERENCES

- V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proc. IEEE* **105**, 2295–2329 (2017).
- Kung, "Why systolic architectures?," *Computer* **15**, 37–46 (1982).
- A. H. Ahmed, A. E. Moznine, D. Lim, Y. Ma, A. Rylyakov, and S. Shekhar, "A dual-polarization silicon-photonics coherent transmitter supporting 552 Gb/s/wavelength," *IEEE J. Solid-State Circuits* **55**, 2597–2608 (2020).
- A. H. Ahmed, A. Sharkia, B. Casper, S. Mirabbasi, and S. Shekhar, "Silicon-photonics microring links for datacenters—Challenges and opportunities," *IEEE J. Sel. Top. Quantum Electron.* **22**, 194–203 (2016).
- S. Moazeni, S. Lin, M. Wade, L. Alloatti, R. J. Ram, M. Popović, and V. Stojanović, "A 40-Gb/s PAM-4 transmitter based on a ring-resonator optical DAC in 45-nm SOI CMOS," *IEEE J. Solid-State Circuits* **52**, 3503–3516 (2017).
- M. W. AlTaha, H. Jayatilaka, Z. Lu, J. F. Chung, D. Celo, D. Goodwill, E. Bernier, S. Mirabbasi, L. Chrostowski, and S. Shekhar, "Monitoring and automatic tuning and stabilization of a  $2 \times 2$  MZI optical switch for large-scale WDM switch networks," *Opt. Express* **27**, 24747–24764 (2019).
- H. Jayatilaka, H. Shoman, L. Chrostowski, and S. Shekhar, "Photoconductive heaters enable control of large-scale silicon photonic ring resonator circuits," *Optica* **6**, 84–91 (2019).
- P. Dong, A. Melikyan, and K. Kim, "Commercializing silicon microring resonators: Technical challenges and potential solutions," in *2018 Conference on Lasers and Electro-Optics (CLEO)* (Optica Publishing Group, 2018), pp. 1–2.
- K. Ikeda, K. Suzuki, R. Konoike, S. Namiki, and H. Kawashima, "Large-scale silicon photonics switch based on 45-nm CMOS technology," *Opt. Commun.* **466**, 125677 (2020).
- M. K. Bhaskar, R. Riedinger, B. Machielse, D. S. Levonian, C. T. Nguyen, E. N. Knall, H. Park, D. Englund, M. Lončar, D. D. Sukachev, and M. D. Lukin, "Experimental demonstration of memory-enhanced quantum communication," *Nature* **580**, 60–64 (2020).
- A. N. Tait, T. Ferreira de Lima, M. A. Nahmias, H. B. Miller, H.-T. Peng, B. J. Shastri, and P. R. Prucnal, "Silicon photonic modulator neuron," *Phys. Rev. Appl.* **11**, 064043 (2019).
- Y. Shen, N. C. Harris, S. Skirlo, M. Prabhu, T. Baehr-Jones, M. Hochberg, X. Sun, S. Zhao, H. Larochelle, D. Englund, and Marin Soljačić, "Deep learning with coherent nanophotonic circuits," *Nat. Photonics* **11**, 441–446 (2017).
- T. W. Hughes, M. Minkov, Y. Shi, and S. Fan, "Training of photonic neural networks through *in situ* backpropagation and gradient measurement," *Optica* **5**, 864 (2018).
- V. Bangari, B. A. Marquez, H. B. Miller, A. N. Tait, M. A. Nahmias, T. F. de Lima, H.-T. Peng, P. R. Prucnal, and B. J. Shastri, "Digital electronics and analog photonics for convolutional neural networks (DEAP-CNNs)," *IEEE J. Sel. Top. Quantum Electron.* **26**, 7701213 (2019); [arXiv:1907.01525](https://arxiv.org/abs/1907.01525) [eess.SP].
- S. Shekhar, "Silicon photonics: A brief tutorial," *IEEE Solid-State Circuits Mag.* **13**, 22–32 (2021).
- C. Ramey, "Silicon photonics for artificial intelligence acceleration: HotChips 32," in *2020 IEEE Hot Chips 32 Symposium (HCS)*, 2020.
- X. Li, G. Zhang, H. H. Huang, Z. Wang, and W. Zheng, "Performance analysis of GPU-based convolutional neural networks," in *2016 45th International Conference on Parallel Processing (ICPP)* (IEEE, 2016), pp. 67–76.
- C. Huang, V. J. Sorger, M. Miscuglio, M. Al-Qadasi, A. Mukherjee, L. Lampe, M. Nichols, A. N. Tait, T. Ferreira de Lima, B. A. Marquez, J. Wang, L. Chrostowski, M. P. Fok, D. Brunner, S. Fan, S. Shekhar, P. R. Prucnal, and B. J. Shastri, "Prospects and applications of photonic neural networks," *Adv. Phys.: X* **7**, 1981155 (2022).
- M. J. Filipovich, Z. Guo, M. Al-Qadasi, B. A. Marquez, H. D. Morison, V. J. Sorger, P. R. Prucnal, S. Shekhar, and B. J. Shastri, "Monolithic silicon photonic architecture for training deep neural networks with direct feedback alignment," [arXiv:2111.06862](https://arxiv.org/abs/2111.06862) [cs.LG] (2021).
- D. Pérez, I. Gasulla, and J. Capmany, "Field-programmable photonic arrays," *Opt. Express* **26**, 27265–27278 (2018).
- W. Zhang and J. Yao, "Photonic integrated field-programmable disk array signal processor," *Nat. Commun.* **11**, 406 (2020).
- N. Tezak, T. Van Vaerenbergh, J. S. Pelc, G. J. Mendoza, D. Kielpinski, H. Mabuchi, and R. G. Beausoleil, "Integrated coherent Ising machines based on self-phase modulation in microring resonators," *IEEE J. Sel. Top. Quantum Electron.* **26**, 5900115 (2020).
- X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O'Gara, L. Kling, G. D. Marshall, R. Santagati, T. C. Ralph *et al.*, "Large-scale silicon quantum photonics implementing arbitrary two-qubit processing," *Nat. Photonics* **12**, 534–539 (2018).
- C. L. Lawson and R. J. Hanson, *Solving Least Squares Problems* (Society of Industrial and Applied Mathematics, 1995).

- <sup>25</sup>G. H. Golub and C. Reinsch, "Singular value decomposition and least squares solutions," *Numer. Math.* **14**, 403–420 (1970).
- <sup>26</sup>M. Reck, A. Zeilinger, H. J. Bernstein, and P. Bertani, "Experimental realization of any discrete unitary operator," *Phys. Rev. Lett.* **73**, 58–61 (1994).
- <sup>27</sup>W. R. Clements, P. C. Humphreys, B. J. Metcalf, W. S. Kolthammer, and I. A. Walsmley, "Optimal design for universal multiport interferometers," *Optica* **3**, 1460–1465 (2016).
- <sup>28</sup>S. Pai, I. A. D. Williamson, T. W. Hughes, M. Minkov, O. Solgaard, S. Fan, and D. A. B. Miller, "Parallel programming of an arbitrary feedforward photonic network," *IEEE J. Sel. Top. Quantum Electron.* **26**(5), 6100813 (2020).
- <sup>29</sup>C. Taballione, T. A. W. Wolterink, J. Lugani, A. Eckstein, B. A. Bell, R. Grootjans, I. Visscher, D. Gekus, C. G. H. Roeloffzen, J. J. Renema *et al.*, " $8 \times 8$  reconfigurable quantum photonic processor based on silicon nitride waveguides," *Opt. Express* **27**, 26842 (2019).
- <sup>30</sup>S. Shekhar, L. Chrostowski, S. Mirabbasi, S. Nayak, M. W. Alta, A. Naguib, A. S. Ramani, and H. Jayatilleka, "Silicon electronics-photonics integrated circuits for datacenters," in *2016 IEEE Compound Semiconductor Integrated Circuit Symposium (CSICS)* (IEEE, 2016), pp. 1–4.
- <sup>31</sup>C. Sorace-Agaskar, J. Leu, M. R. Watts, and V. Stojanovic, "Electro-optical co-simulation for integrated CMOS photonic circuits with VerilogA," *Opt. Express* **23**, 27180–27203 (2015).
- <sup>32</sup>P.-I. Dietrich, M. Blaicher, I. Reuter, M. Billah, T. Hoose, A. Hofmann, C. Caer, R. Dangel, B. Offrein, U. Troppenz *et al.*, "In situ 3D nanoprinting of free-form coupling elements for hybrid photonic integration," *Nat. Photonics* **12**, 241–247 (2018).
- <sup>33</sup>H. Honmou, R. Ishikawa, H. Ueno, and M. Kobayashi, "1.0 dB low-loss coupling of laser diode to single-mode fibre using a planoconvex graded-index rod lens," *Electron. Lett.* **22**, 1122–1123 (1986).
- <sup>34</sup>R. Won, "Wire-bonding assembly," *Nat. Photonics* **12**, 500 (2018).
- <sup>35</sup>Technology for photonic multi-chip integration—photonic wire bonding.
- <sup>36</sup>M. R. Billah, M. Blaicher, J. N. Kemal, T. Hoose, H. Zwickel, P. Dietrich, U. Troppenz, M. Moehrl, F. Merget, A. Hofmann, J. Witzens, S. Randel, W. Freude, and C. Koos, "8-channel 448 Gbit/s silicon photonic transmitter enabled by photonic wire bonding," in *2017 Optical Fiber Communications Conference and Exhibition (OFC)* (IEEE, 2017), pp. 1–3.
- <sup>37</sup>N. Lindenmann, G. Balthasar, D. Hillerkuss, R. Schmogrow, M. Jordan, J. Leuthold, W. Freude, and C. Koos, "Photonic wire bonding: A novel concept for chip-scale interconnects," *Opt. Express* **20**, 17667 (2012).
- <sup>38</sup>N. Lindenmann, S. Dottermusch, M. L. Goedecke, T. Hoose, M. R. Billah, T. P. Onanuga, A. Hofmann, W. Freude, and C. Koos, "Connecting silicon photonic circuits to multicore fibers by photonic wire bonding," *J. Lightwave Technol.* **33**, 755–760 (2015).
- <sup>39</sup>L. Chrostowski, H. Shoman, M. Hammood, H. Yun, J. Jhoja, E. Luan, S. Lin, A. Mistry, D. Witt, N. A. F. Jaeger, S. Shekhar, H. Jayatilleka, P. Jean, S. B.-d. Villers, J. Cauchon, W. Shi, C. Horvath, J. N. Westwood-Bachman, K. Setzer, M. Aktary, N. S. Patrick, R. J. Bojko, A. Khavasi, X. Wang, T. Ferreira de Lima, A. N. Tait, P. R. Prucnal, D. E. Hagan, D. Stevanovic, and A. P. Knights, "Silicon photonic circuit design using rapid prototyping foundry process design kits," *IEEE J. Sel. Top. Quantum Electron.* **25**, 8201326 (2019).
- <sup>40</sup>Y. P. Li and C. H. Henry, "Silicon optical bench waveguide technology," in *Optical Fiber Telecommunications*, 3rd ed. (Academic Press, 1997), Chap. 8, pp. 319–376, an optional note.
- <sup>41</sup>A. Samani, V. Veerasubramanian, E. El-Fiky, D. Patel, and D. V. Plant, "A silicon photonic PAM-4 modulator based on dual-parallel Mach-Zehnder interferometers," *IEEE Photonics J.* **8**, 7800610 (2016).
- <sup>42</sup>D. Dai, K. Ma, and H. Wu, "Mode/polarization manipulation in silicon photonics," *J. Phys.: Conf. Ser.* **844**, 012039 (2017).
- <sup>43</sup>D. González-Andrade, C. Lafforgue, E. Durán-Valdeiglesias, X. Le Roux, M. Berciano, E. Cassan, D. Marris-Morini, A. V. Velasco, P. Cheben, L. Vivien, and C. Alonso-Ramos, "Polarization- and wavelength-agnostic nanophotonic beam splitter," *Sci. Rep.* **9**, 3604 (2019).
- <sup>44</sup>Z. Lu, H. Yun, Y. Wang, Z. Chen, F. Zhang, N. A. F. Jaeger, and L. Chrostowski, "Broadband silicon photonic directional coupler using asymmetric-waveguide based phase control," *Opt. Express* **23**, 3795–3808 (2015).
- <sup>45</sup>S. Chen, Y. Shi, S. He, and D. Dai, "Low-loss and broadband  $2 \times 2$  silicon thermo-optic Mach-Zehnder switch with bent directional couplers," *Opt. Lett.* **41**, 836–839 (2016).
- <sup>46</sup>R. Hui, *Introduction to Fiber-Optic Communications* (Academic Press, 2020).
- <sup>47</sup>C. Li, S. Xu, X. Huang, Z. Feng, C. Yang, K. Zhou, J. Gan, and Z. Yang, "High-speed frequency modulated low-noise single-frequency fiber laser," *IEEE Photonics Technol. Lett.* **28**, 1692–1695 (2016).
- <sup>48</sup>K. Giewont, K. Nummy, F. A. Anderson, J. Ayala, T. Barwicz, Y. Bian, K. K. Dezfulian, D. M. Gill, T. Houghton, S. Hu, B. Peng, M. Rakowski, S. Rauch, J. C. Rosenberg, A. Sahin, I. Stobert, and A. Stricker, "300-mm monolithic silicon photonics foundry technology," *IEEE J. Sel. Top. Quantum Electron.* **25**, 8200611 (2019).
- <sup>49</sup>L. Szilagyi, R. Henker, D. Hame, and F. Ellinger, "2.2-pJ/bit 30-Gbit/s Mach-Zehnder modulator driver in 22-nm-FDSOI," in *2018 IEEE/MTT-S International Microwave Symposium: IMS* (IEEE, 2018), pp. 1530–1533.
- <sup>50</sup>M. Jacques, A. Samani, E. El-Fiky, D. Patel, Z. Xing, and D. V. Plant, "Optimization of thermo-optic phase-shifter design and mitigation of thermal crosstalk on the SOI platform," *Opt. Express* **27**, 10456–10471 (2019).
- <sup>51</sup>N. C. Harris, Y. Ma, J. Mower, T. Baehr-Jones, D. Englund, M. Hochberg, and C. Galland, "Efficient, compact and low loss thermo-optic phase shifter in silicon," *Opt. Express* **22**, 10487–10493 (2014).
- <sup>52</sup>S. Nakamura, S. Takahashi, I. Ogura, J. Ushida, K. Kurata, T. Hino, H. Takeshita, A. Tajima, M. Yu, and G. Lo, "High extinction ratio optical switching independently of temperature with silicon photonic  $1 \times 8$  switch," in *OFC/NFOEC* (Optica Publishing Group, 2012), pp. 1–3.
- <sup>53</sup>F. Y. Liu, D. Patil, J. Lexau, P. Amberg, M. Dayringer, J. Gainsley, H. F. Moghadam, X. Zheng, J. E. Cunningham, A. V. Krishnamoorthy, E. Alon, and R. Ho, "10-Gbps, 5.3-mW optical transmitter and receiver circuits in 40-nm CMOS," *IEEE J. Solid-State Circuits* **47**, 2049–2067 (2012).
- <sup>54</sup>Y. Lee and W. Chen, "A 20-Gb/s, 2.4 pJ/bit, fully integrated optical receiver with a baud-rate clock and data recovery," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)* (IEEE, 2018), pp. 1–4.
- <sup>55</sup>T. Takemoto, F. Yuki, H. Yamashita, S. Tsuji, T. Saito, and S. Nishimura, "A 25 Gb/s  $\times$  4-channel 74 mW/ch transimpedance amplifier in 65 nm CMOS," in *IEEE Custom Integrated Circuits Conference 2010* (IEEE, 2010), pp. 1–4.
- <sup>56</sup>H. Morita, K. Uchino, E. Otani, H. Ohtorii, T. Ogura, K. Oniki, S. Oka, S. Yanagawa, and H. Suzuki, "8.2 A  $12 \times 5$  two-dimensional optical I/O array for 600Gb/s chip-to-chip interconnect in 65nm CMOS," in *2014 IEEE International Solid-State Circuits Conference Digest of Technical Papers (ISSCC)* (IEEE, 2014), pp. 140–141.
- <sup>57</sup>C. L. Schow, F. E. Doany, C. W. Baks, Y. H. Kwark, D. M. Kuchta, and J. A. Kash, "A single-chip CMOS-based parallel optical transceiver capable of 240-Gb/s bidirectional data rates," *J. Lightwave Technol.* **27**, 915–929 (2009).
- <sup>58</sup>K. Park, B. Yoo, M. Hwang, H. Chi, H. Kim, J. Park, K. Kim, and D. Jeong, "A 10-Gb/s optical receiver front-end with 5-mW transimpedance amplifier," in *2010 IEEE Asian Solid-State Circuits Conference* (IEEE, 2010), pp. 1–4.
- <sup>59</sup>K. R. Lakshminikumar, A. Kurylak, M. Nagaraju, R. Booth, R. K. Nandwana, J. Pampanin, and V. Bocuzzi, "A process and temperature insensitive CMOS linear TIA for 100 Gb/s  $\lambda$  PAM-4 optical links," *IEEE J. Solid-State Circuits* **54**, 3180–3190 (2019).
- <sup>60</sup>Y. Frans, J. Shin, L. Zhou, P. Upadhyaya, J. Im, V. Kireev, M. Elzeftawi, H. Hedayati, T. Pham, S. Asuncion, C. Borrelli, G. Zhang, H. Zhang, and K. Chang, "A 56-Gb/s PAM4 wireline transceiver using a 32-way time-interleaved SAR ADC in 16-nm FinFET," *IEEE J. Solid-State Circuits* **52**, 1101–1110 (2017).
- <sup>61</sup>J. Cao, M. Green, A. Momtaz, K. Vakilian, D. Chung, K.-C. Jen, M. Caresosa, X. Wang, W.-G. Tan, Y. Cai, L. Fujimori, and A. Hairapetian, "OC-192 transmitter and receiver in standard 0.18- $\mu$ m CMOS," *IEEE J. Solid-State Circuits* **37**, 1768–1780 (2002).
- <sup>62</sup>S. Tanaka, T. Simoyama, T. Aoki, T. Mori, S. Sekiguchi, S.-H. Jeong, T. Usuki, Y. Tanaka, and K. Morito, "Ultralow-power (1.59 mW/Gbps), 56-Gbps PAM4 operation of Si photonic transmitter integrating segmented PIN Mach-Zehnder modulator and 28-nm CMOS driver," *J. Lightwave Technol.* **36**, 1275–1280 (2018).
- <sup>63</sup>A. Michard, J.-F. Carpentier, N. Michit, P. Le Maître, P. Bénabès, and P. M. Ferreira, "A sub-pJ/bit, low-ER Mach-Zehnder-based transmitter for chip-to-chip optical interconnects," *IEEE J. Sel. Top. Quantum Electron.* **26**, 8301910 (2020).



- <sup>64</sup>S. Nayak, A. H. Ahmed, A. Sharkia, A. S. Ramani, S. Mirabbasi, and S. Shekhar, "A 10-Gb/s  $-18.8$  dBm sensitivity 5.7 mW fully-integrated optoelectronic receiver with avalanche photodetector in 0.13  $\mu\text{m}$  CMOS," *IEEE Trans. Circuits Syst. I* **66**, 3162–3173 (2019).
- <sup>65</sup>B. Moons, B. De Brabandere, L. Van Gool, and M. Verhelst, "Energy-efficient ConvNets through approximate computing," in *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)* (IEEE, 2016), pp. 1–8.
- <sup>66</sup>M. A. Nahmias, T. F. de Lima, A. N. Tait, H.-T. Peng, B. J. Shastri, and P. R. Prucnal, "Photonic multiply-accumulate operations for neural networks," *IEEE J. Sel. Top. Quantum Electron.* **26**, 7701518 (2020).
- <sup>67</sup>S. Agarwal, T.-T. Quach, O. Parekh, A. H. Hsia, E. P. DeBenedictis, C. D. James, M. J. Marinella, and J. B. Aimone, "Energy scaling advantages of resistive memory crossbar based computation and its application to sparse coding," *Front. Neurosci.* **9**, 484 (2016).
- <sup>68</sup>S. Gudaparthi, S. Narayanan, R. Balasubramanian, E. Giacomini, H. Kam-balasubramanyam, P.-E. Gaillardon, "Wire-aware architecture and dataflow for CNN accelerators," in *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO '52)* (IEEE 2019).
- <sup>69</sup>D. Benedikovic, L. Viro, G. Aubin, J.-M. Hartmann, F. Amar, B. Szelag, X. Le Roux, C. Alonso-Ramos, P. Crozat, É. Cassan, D. Marriss-Morini, C. Baudot, F. Boeuf, J.-M. Fédéli, C. Kopp, and L. Vivien, "Silicon-germanium p-i-n photodiodes with double heterojunction: High-speed operation at 10 Gbps and beyond," in *2020 European Conference on Integrated Optics* (published online 2020), hal-02539654.
- <sup>70</sup>A. N. Tait, M. A. Nahmias, B. J. Shastri, and P. R. Prucnal, "Broadcast and weight: An integrated network for scalable photonic spike processing," *J. Lightwave Technol.* **32**, 4029–4041 (2014).
- <sup>71</sup>H. Li, G. Balamurugan, T. Kim, M. N. Sakib, R. Kumar, H. Rong, J. Jaussi, and B. Casper, "A 3-D-integrated silicon photonic microring-based 112-Gb/s PAM-4 transmitter with nonlinear equalization and thermal control," *IEEE J. Solid-State Circuits* **56**, 19–29 (2021).
- <sup>72</sup>R. Konoike, K. Suzuki, S. Namiki, H. Kawashima, and K. Ikeda, "Ultra-compact silicon photonics switch with high-density thermo-optic heaters," *Opt. Express* **27**, 10332–10342 (2019).
- <sup>73</sup>X. Zheng, E. Chang, P. Amberg, I. Shubin, J. Lexau, F. Liu, H. Thacker, S. S. Djordjevic, S. Lin, Y. Luo, J. Yao, J.-H. Lee, K. Raj, R. Ho, J. E. Cunningham, and A. V. Krishnamoorthy, "A high-speed, tunable silicon photonic ring modulator integrated with ultra-efficient active wavelength control," *Opt. Express* **22**, 12628–12633 (2014).
- <sup>74</sup>F. Gan, T. Barwicz, M. A. Popovic, M. S. Dahlem, C. W. Holzwarth, P. T. Rakich, H. I. Smith, E. P. Ippen, and F. X. Kartner, "Maximizing the thermo-optic tuning range of silicon photonic structures," in *2007 Photonics in Switching* (IEEE, 2007), pp. 67–68.
- <sup>75</sup>A. Masood, M. Pantouvaki, G. Lepage, P. Verheyen, J. Van Campenhout, P. Absil, D. Van Thourhout, and W. Bogaerts, "Comparison of heater architectures for thermal control of silicon photonic circuits," in *10th International Conference on Group IV Photonics* (IEEE, 2013), pp. 83–84.
- <sup>76</sup>A. N. Tait, A. X. Wu, T. F. de Lima, E. Zhou, B. J. Shastri, M. A. Nahmias, and P. R. Prucnal, "Microring weight banks," *IEEE J. Sel. Top. Quantum Electron.* **22**, 312–325 (2016).
- <sup>77</sup>H. Jayatilaka, K. Murray, M. Caverley, N. A. F. Jaeger, L. Chrostowski, and S. Shekhar, "Crosstalk in SOI microring resonator-based filters," *J. Lightwave Technol.* **34**, 2886–2896 (2016).
- <sup>78</sup>Q. Xu, D. Fattal, and R. G. Beausoleil, "Silicon microring resonators with 1.5- $\mu\text{m}$  radius," *Opt. Express* **16**, 4309–4315 (2008).
- <sup>79</sup>A. Biberman, M. J. Shaw, E. Timurdogan, J. B. Wright, and M. R. Watts, "Ultralow-loss silicon ring resonators," in *9th International Conference on Group IV Photonics (GFP)* (IEEE, 2012), pp. 39–41.
- <sup>80</sup>H. Shoman, H. Jayatilaka, A. H. K. Park, A. Mistry, N. A. F. Jaeger, S. Shekhar, and L. Chrostowski, "Compact wavelength- and bandwidth-tunable microring modulator," *Opt. Express* **27**, 26661–26675 (2019).
- <sup>81</sup>M. Milanizadeh, M. Petrini, F. Morichetti, and A. Melloni, "FSR-free filter with hitless wavelength across C+L telecom band," in *OSA Advanced Photonics Congress (AP) 2020 (IPR, NP, NOMA, Networks, PVLED, PSC, SPPCom, SOF)* (Optical Society of America, 2020), p. IM3A.5.
- <sup>82</sup>N. Eid, R. Boeck, H. Jayatilaka, L. Chrostowski, W. Shi, and N. A. F. Jaeger, "FSR-free silicon-on-insulator microring resonator based filter with bent contra-directional couplers," *Opt. Express* **24**, 29009–29021 (2016).
- <sup>83</sup>T. Barwicz, Y. Taira, T. W. Lichoules, N. Boyer, Y. Martin, H. Numata, J.-W. Nah, S. Takenobu, A. Janta-Polczynski, E. L. Kimbrell, R. Leidy, M. H. Khater, S. Kamapurkar, S. Engelmann, Y. A. Vlasov, and P. Fortier, "A novel approach to photonic packaging leveraging existing high-throughput microelectronic facilities," *IEEE J. Sel. Top. Quantum Electron.* **22**, 455–466 (2016).
- <sup>84</sup>S. Fatholouloumi, K. Nguyen, H. Mahalingam, M. Sakib, Z. Li, C. Seibert, M. Montazeri, J. Chen, J. Doyle, H. Jayatilaka, C. Jan, J. Heck, R. Venables, H. Frish, R. Defrees, R. Appleton, S. Hollingsworth, S. McCargar, R. Jones, and L. Liao, "1.6Tbps silicon photonics integrated circuit for co-packaged optical-IO switch applications," in *Optical Fiber Communication Conference (OFC) 2020* (Optica Publishing Group, 2020), p. T3H.1.
- <sup>85</sup>B. Shi, N. Calabretta, and R. Stabile, "Deep neural network through an InP SOA-based photonic integrated cross-connect," *IEEE J. Sel. Top. Quantum Electron.* **26**, 7701111 (2020).
- <sup>86</sup>Z. Huang, C. Li, D. Liang, K. Yu, C. Santori, M. Fiorentino, W. Sorin, S. Palermo, and R. G. Beausoleil, "25 Gbps low-voltage waveguide Si-Ge avalanche photodiode," *Optica* **3**, 793–798 (2016).
- <sup>87</sup>J. Joo, S. Kim, I. G. Kim, K.-S. Jang, and G. Kim, "High-sensitivity 10 Gbps Ge-on-Si photoreceiver operating at  $\lambda$  1.55  $\mu\text{m}$ ," *Opt. Express* **18**, 16474–16479 (2010).
- <sup>88</sup>S. Shekhar, J. E. Jaussi, F. O'Mahony, M. Mansuri, and B. Casper, "Design considerations for low-power receiver front-end in high-speed data links," in *2013 IEEE Custom Integrated Circuits Conference (CICC)* (IEEE, 2013), pp. 1–8.
- <sup>89</sup>P. J. Lim, A. Y. C. Tzeng, H. L. Chuang, and S. A. St. Onge, "A 3.3-V monolithic photodetector/CMOS-preamplifier for 531 Mb/s optical data link applications," in *1993 IEEE International Solid-State Circuits Conference Digest of Technical Papers* (IEEE, 1993), pp. 96–97.
- <sup>90</sup>A. J. Zilkie, P. Seddighian, B. J. Bijlani, W. Qian, D. C. Lee, S. Fatholouloumi, J. Fong, R. Shafiiha, D. Feng, B. J. Luff, X. Zheng, J. E. Cunningham, A. V. Krishnamoorthy, and M. Asghari, "Power-efficient III-V/silicon external cavity DBR lasers," *Opt. Express* **20**, 23456–23462 (2012).
- <sup>91</sup>S. Tanaka, S.-H. Jeong, S. Sekiguchi, T. Kurahashi, Y. Tanaka, and K. Morito, "High-output-power, single-wavelength silicon hybrid laser using precise flip-chip bonding technology," *Opt. Express* **20**, 28057–28069 (2012).
- <sup>92</sup>J. Lee, J. Bovington, I. Shubin, Y. Luo, J. Yao, S. Lin, J. E. Cunningham, K. Raj, A. V. Krishnamoorthy, and X. Zheng, "12.2% waveguide-coupled wall plug efficiency in single mode external-cavity tunable Si/III-V hybrid laser," in *2015 IEEE Optical Interconnects Conference (OI)* (IEEE, 2015), pp. 142–143.
- <sup>93</sup>C. R. Doerr, N. Dupuis, and L. Zhang, "Optical isolator using two tandem phase modulators," *Opt. Lett.* **36**, 4293–4295 (2011).
- <sup>94</sup>H. Shoman, N. Jaeger, C. Mosquera, H. Jayatilaka, M. Ma, H. Rong, S. Shekhar, and L. Chrostowski, "Stable and reduced-linewidth laser through active cancellation of reflections without a magneto-optic isolator," *J. Lightwave Technol.* **39**, 6215 (2021).
- <sup>95</sup>P. Dong, W. Qian, H. Liang, R. Shafiiha, D. Feng, G. Li, J. E. Cunningham, A. V. Krishnamoorthy, and M. Asghari, "Thermally tunable silicon racetrack resonators with ultralow tuning power," *Opt. Express* **18**, 20298–20304 (2010).
- <sup>96</sup>X. Zheng, E. Chang, I. Shubin, G. Li, Y. Luo, J. Yao, H. Thacker, J. Lee, J. Lexau, F. Liu, P. Amberg, K. Raj, R. Ho, J. E. Cunningham, and A. V. Krishnamoorthy, "A 33mW 100Gbps CMOS silicon photonic WDM transmitter using off-chip laser sources," in *2013 Optical Fiber Communication Conference and Exposition and the National Fiber Optic Engineers Conference (OFC/NFOEC)* (Optica Publishing Group, 2013), pp. 1–3.
- <sup>97</sup>C. Sun, M. Wade, M. Georgas, S. Lin, L. Alloati, B. Moss, R. Kumar, A. H. Atabaki, F. Pavanello, J. M. Shainline, J. S. Orcutt, R. J. Ram, M. Popović, and V. Stojanović, "A 45 nm CMOS-SOI monolithic photonics platform with bit-statistics-based resonant microring thermal tuning," *IEEE J. Solid-State Circuits* **51**, 893–907 (2016).
- <sup>98</sup>P. Alipour, A. H. Atabaki, M. Askari, A. Adibi, and A. A. Eftekhar, "Robust postfabrication trimming of ultracompact resonators on silicon on insulator with relaxed requirements on resolution and alignment," *Opt. Lett.* **40**, 4476–4479 (2015).



- <sup>99</sup>X. Chen, M. M. Milosevic, X. Yu, B. Chen, A. F. J. Runge, A. Z. Khokhar, S. Mailis, D. J. Thomson, A. C. Peacock, S. Saito, O. L. Muskens, and G. T. Reed, "Germanium implanted photonic devices for post-fabrication trimming and programmable circuits," *Proc. SPIE* **10823**, 108230U (2018).
- <sup>100</sup>Y. Feng, D. J. Thomson, G. Z. Mashanovich, and J. Yan, "Performance analysis of a silicon NOEMS device applied as an optical modulator based on a slot waveguide," *Opt. Express* **28**, 38206–38222 (2020).
- <sup>101</sup>Y. Xing, T. Ako, J. P. George, D. Korn, H. Yu, P. Verheyen, M. Pantouvaki, G. Lepage, P. Absil, A. Ruocco, C. Koos, J. Leuthold, K. Neyts, J. Beeckman, and W. Bogaerts, "Digitally controlled phase shifter using an SOI slot waveguide with liquid crystal infiltration," *IEEE Photonics Technol. Lett.* **27**, 1269–1272 (2015).
- <sup>102</sup>Q. Zhang, Y. Zhang, J. Li, R. Soref, T. Gu, and J. Hu, "Broadband nonvolatile photonic switching based on optical phase change materials: Beyond the classical figure-of-merit," *Opt. Lett.* **43**, 94–97 (2018).
- <sup>103</sup>N. Dhingra, J. Song, G. J. Saxena, E. K. Sharma, and B. M. A. Rahman, "Design of a compact low-loss phase shifter based on optical phase change material," *IEEE Photonics Technol. Lett.* **31**, 1757–1760 (2019).
- <sup>104</sup>C. Ríos, M. Stegmaier, P. Hosseini, D. Wang, T. Scherer, C. D. Wright, H. Bhaskaran, and W. H. P. Pernice, "Integrated all-photonic non-volatile multi-level memory," *Nat. Photonics* **9**, 725–732 (2015).
- <sup>105</sup>User Guide NVIDIA Docs, Optimizing Linear/Fully-Connected Layers, <https://docs.nvidia.com/deeplearning/performance/pdf/Optimizing-Linear-Fully-Connected-Layers-User-Guide.pdf>, 2021.
- <sup>106</sup>M. J. Connelly, *Semiconductor Optical Amplifiers* (Springer, 2011).
- <sup>107</sup>V. Sasikala and K. Chitra, "All optical switching and associated technologies: A review," *J. Opt.* **47**, 307–317 (2018).
- <sup>108</sup>D. M. Baney, P. Gallion, and R. S. Tucker, "Theory and measurement techniques for the noise figure of optical amplifiers," *Opt. Fiber Technol.* **6**, 122–154 (2000).
- <sup>109</sup>T. Numai, "Semiconductor optical amplifiers," in *Laser Diodes and Their Applications to Communications and Information Processing* (John Wiley & Sons, Ltd., 2010), Chap. 9, pp. 233–245.
- <sup>110</sup>C-band semiconductor optical amplifier: SOA1530S, InP Photonic Integrated Circuits Foundry, 2021, <http://www.zwphotronics.com/en/>.
- <sup>111</sup>M. Chauhan, O. P. Vyas, and S. Bhandari, "'Cross gain modulation' effect of SOAs for different modulation formats," *Int. J. Eng. Res. Technol.* **2**, 357–360 (2018).
- <sup>112</sup>S. Xu and J. B. Khurgin, "A dispersion management scheme for reducing SOA-induced crosstalk in WDM links," *J. Lightwave Technol.* **22**, 417–422 (2004).
- <sup>113</sup>B. J. Shastri, A. N. Tait, T. Ferreira de Lima, W. H. P. Pernice, H. Bhaskaran, C. D. Wright, and P. R. Prucnal, "Photonics for artificial intelligence and neuromorphic computing," *Nat. Photonics* **15**, 102–114 (2021).
- <sup>114</sup>D. E. Tamir, N. T. Shaked, P. J. Wilson, and S. Dolev, "High-speed and low-power electro-optical DSP coprocessor," *J. Opt. Soc. Am. A* **26**, A11–A20 (2009).
- <sup>115</sup>W. Zhang, C. Huang, S. Bilodeau, A. Jha, E. Blow, T. F. D. Lima, B. J. Shastri, and P. Prucnal, "Microring weight banks control beyond 8.5-bits accuracy," [arXiv:2104.01164](https://arxiv.org/abs/2104.01164) [physics.app-ph] (2021).
- <sup>116</sup>I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Quantized neural networks: Training neural networks with low precision weights and activations," [arXiv:1609.07061](https://arxiv.org/abs/1609.07061) [cs.NE] (2016).
- <sup>117</sup>E. H. Lee, D. Miyashita, E. Chai, B. Murmann, and S. S. Wong, "LogNet: Energy-efficient neural networks using logarithmic computation," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (IEEE, 2017), pp. 5900–5904.
- <sup>118</sup>S. K. Esser, P. A. Merolla, J. V. Arthur, A. S. Cassidy, R. Appuswamy, A. Andreopoulos, D. J. Berg, J. L. McKinstry, T. Melano, D. R. Barch, C. di Nolfo, P. Datta, A. Amir, B. Taba, M. D. Flickner, and D. S. Modha, "Convolutional networks for fast, energy-efficient neuromorphic computing," *Proc. Natl. Acad. Sci. U. S. A.* **113**, 11441–11446 (2016).
- <sup>119</sup>C. Cole, "Optical and electrical programmable computing energy use comparison," *Opt. Express* **29**, 13153–13170 (2021).
- <sup>120</sup>X. Wu, B. Dama, P. Gothoskar, P. Metz, K. Shastri, S. Sunder, J. V. d. Spiegel, Y. Wang, M. Webster, and W. Wilson, "A 20Gb/s NRZ/PAM-4 1V transmitter in 40nm CMOS driving a Si-photonic modulator in 0.13- $\mu$ m CMOS," in *2013 IEEE International Solid-State Circuits Conference Digest of Technical Papers* (IEEE, 2013), pp. 128–129.
- <sup>121</sup>L. Carroll, J.-S. Lee, C. Scarcella, K. Gradkowski, M. Duperron, H. Lu, Y. Zhao, C. Eason, P. Morrissey, M. Rensing, S. Collins, H. Y. Hwang, and P. O'Brien, "Photonic packaging: Transforming silicon photonic integrated circuits into photonic devices," *Appl. Sci.* **6**, 426 (2016).
- <sup>122</sup>M. Wade, "TeraPHY: A chiplet technology for low-power, high-bandwidth in-package optical I/O," in *2019 IEEE Hot Chips 31 Symposium (HCS)* (IEEE, 2019), pp. i–xlvi.