

Packaging and Interconnect Considerations in Neuromorphic Photonic Accelerators

Mohammadreza Sanadgol Nezami¹, Thomas Ferreira de Lima, Matthew Mitchell, Shangxuan Yu, Jing Wang², Simon Bilodeau³, *Graduate Student Member, IEEE*, Weipeng Zhang, *Graduate Student Member, IEEE*, Mohammed Al-Qadasi, Iman Taghavi, Alexander Tofini, Stephen Lin⁴, Bhavin J. Shastri, *Senior Member, IEEE*, Paul R. Prucnal, *Life Fellow, IEEE*, Lukas Chrostowski⁵, *Senior Member, IEEE*, and Sudip Shekhar⁶, *Senior Member, IEEE*

Abstract—Developing compute platforms capable of performing computations at high speed is essential for data processing in the next generation of data centers and edge devices. A neuromorphic photonic accelerator on a silicon photonic platform is a promising solution. Compared to silicon photonic data communication transceiver modules, neuromorphic photonic accelerators constitute a large number of active and passive components and optoelectronic devices to handle the parallel processing. Thus, an increased number of optical and electrical interconnects are required, making the packaging of such processors challenging. Moreover, thermal and electrical crosstalk can dramatically degrade the performance of such processors. Thus, packaging a neuromorphic photonic accelerator for efficient processing and data movement requires careful considerations at the chip, module, and board levels. This work investigates the challenges and potential solutions for optical coupling, optical and electrical interconnections, processor-memory communication, and thermal and electrical cross-talk to develop neuromorphic photonic accelerators.

Index Terms—Heterogeneous integration, optical computing, silicon photonics, co-packaging.

I. INTRODUCTION

PHOTONIC computing is a promising technology to address the explosive demand for cost-effective data movement

Manuscript received 1 June 2022; revised 17 August 2022; accepted 17 August 2022. Date of publication 22 August 2022; date of current version 1 September 2022. This research was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC). (*Corresponding author: Mohammadreza Sanadgol Nezami.*)

Mohammadreza Sanadgol Nezami, Matthew Mitchell, Shangxuan Yu, Jing Wang, Mohammed Al-Qadasi, Iman Taghavi, Alexander Tofini, Stephen Lin, Lukas Chrostowski, and Sudip Shekhar are with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver, BC V6T 1Z4, Canada (e-mail: mnezami@ece.ubc.ca; m.mitchell@ubc.ca; mathewyu@ece.ubc.ca; jing@ece.ubc.ca; alqadasi@ece.ubc.ca; staghavi3@ece.ubc.ca; alexanderjtofini@gmail.com; sl0804@ece.ubc.ca; lukasc@ece.ubc.ca; sudip@ece.ubc.ca).

Thomas Ferreira de Lima, Simon Bilodeau, Weipeng Zhang, and Paul R. Prucnal are with the Department of Electrical Engineering, Princeton University, Princeton NJ 08544 USA (e-mail: thomas.ferreira.de.lima@Princeton.EDU; sbilodeau@princeton.edu; weipengz@princeton.edu; prucnal@princeton.edu).

Bhavin J. Shastri is with the Department of Physics, Engineering Physics & Astronomy, Queen's University, Kingston, ON K7L 3N6, Canada, and also with the Vector Institute of Artificial Intelligence, Toronto, ON M5S 1G1, Canada (e-mail: bhavin.shastri@queensu.ca).

Color versions of one or more figures in this article are available at <https://doi.org/10.1109/JSTQE.2022.3200604>.

Digital Object Identifier 10.1109/JSTQE.2022.3200604

and processing in future data centers. Graphics processing units (GPUs) and tensor processing units (TPUs) enable energy-efficient parallel computation and hardware acceleration [1]. Together with CPUs, they provide a compute platform suitable for different kinds of neural networks. TPUs are especially suited to 8-bit precision computation with high energy efficiency. However, current and future datacenter inference and training demands require accelerators that drastically reduce latency and increase TOPS/W. Silicon photonic processors can exploit the inherent parallelism and large bandwidth at optical frequencies [2] to implement multiply-accumulate (MAC) operations with large throughput and low latency on-chip [3], [4], [5], [6], [7], [8].

Fig. 1(a) shows a building-block perceptron including multiple inputs, the sum of products, and a nonlinear activation function. Fig. 1(b) shows the concept of a silicon photonic neural network enabled by wavelength-division multiplexing (WDM). Ring resonators as tunable filters together with the detectors implement the vector-dot product (as the linear part of the neuron) with the ring resonator performing the weight tuning and elementwise multiplication, and the balanced photodetectors performing the summation (while enabling positive and negative weighting) and OE conversion. After current-to-voltage conversion, a ring modulator applies the nonlinear activation function to the electrical signal, performs EO conversion, and delivers the optical signal to the next layer or the feedback loop. Neuromorphic photonic accelerators consist of four fundamental elements: 1) A photonic engine for performing the processing and computation tasks, 2) CMOS circuits such as trans-impedance amplifiers (TIAs), digital-to-analog converters (DACs), and analog-to-digital-converters (ADCs) for electrical signal amplification, calibration and control of the neural network parameters, 3) laser sources to supply input optical signal at various wavelengths possibly aided by semiconductor optical amplifiers (SOAs) for maintaining the required optical signal-to-noise ratio and compensating for the optical propagation loss [9], and 4) a CMOS digital processor to handle the tasks not suited for the photonic engine. Fig. 2 shows the block diagram of a neuromorphic photonic accelerator. Interfacing the photonic engine with the CMOS circuits and optoelectronic components requires a significant amount of pads and interconnections, and hence the scalability of such processors is




Fig. 1. (a) Abstract model of an (n 'th) artificial neuron in a neural network. Incoming signals x_i with $i = 1, 2, \dots, N$ are weighted (w_i). Their sum, added to a trainable extra bias b_n , is nonlinearly transformed (ϕ) into an output signal y_n . (b) Silicon photonic neuron implementation, with electrical IO emphasized. Incoming signals $P_i(t)$ encoded onto the power envelope of wavelength-multiplexed signals are individually weighted by N spectral filters reconfigurable through, as an example, individual voltages V_{in}^w . Summing is performed by balanced photodetectors (each requiring a bias V_n^\pm), with the external bias captured by an added current I_n^b . Nonlinearity is achieved by remodulation onto the power envelope of an output carrier $P_n(t)$ through a transimpedance amplifier (TIA) and a high-speed reconfigurable spectral filter, whose wavelength can be separately adjusted on a slower timescale with, as an example, control voltage V_n .

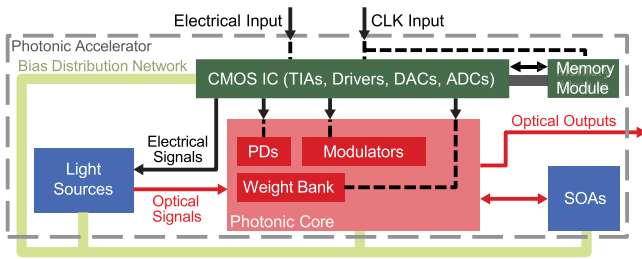


Fig. 2. Block diagram of a neuromorphic photonic accelerator and its electrical and optical interconnects. The bias distribution network delivers adequate DC signals to the light sources, active devices on the PIC, and SOAs.

a complex problem [9]. Indeed, one of the main challenges in developing advanced neuromorphic photonic accelerators is co-packaging of photonic integrated circuit (PIC) with CMOS electronics and opto-electronic components. This is primarily due to the number of electrical interconnects required, the size mismatch between CMOS and photonic chips, and the difficulty meeting the bandwidth requirements (since the photonic computing core requires serial data at tens of Gb/s for full utilization of its throughput capacity). Moreover, co-packaging CMOS and silicon photonic chips necessitates overcoming thermal, electrical, mechanical, and optical challenges [10]. This paper investigates potential packaging solutions for developing neuromorphic photonic accelerators from a laboratory prototype with a few I/Os to a large-scale module with many I/Os. Section II reviews existing technologies for developing neuromorphic photonic accelerators and provides an insight for packaging of large-scale neuromorphic photonic accelerators. It describes the challenges involved in photonic-electronic co-integration with




Fig. 3. Neuromorphic photonic accelerator packaging: (a) Top view of a conventional packaging technology consisting of a CMOS ASIC including SRAMs, TIAs, and drivers, and a neuromorphic photonic accelerator all on a PCB board which utilizes a PCIe interface to communicate with the main CPU on motherboard. (b) Prospect of a packaging solution consisting of a silicon PIC, CMOS chip, and III-V DFB lasers and SOAs integrated via photonic wire bonds. The CMOS ASIC includes memory controller, modulator drivers, and TIAs. BP: bond pad, EWB: electrical wire bond, PWB: photonic wire bond; The recess cavity, etched in the PIC surface, helps to vertically align the facet of the DFB laser with the waveguide facet at the edge of the PIC and hence facilitates the process of photonic wire bonding.

several optoelectronic components and optical and electrical interconnects. A scalability analysis for common packaging techniques is also presented. Section III reviews recent advances in memory-processor communication and proposes a hardware architecture for implementing efficient data movement between a neuromorphic photonic accelerator and memory modules.

II. PHOTONIC-ELECTRONIC INTEGRATION

Monolithic integration, multi-chip side-by-side integration through wire bonding on a printed circuit board (PCB) or flip-chip bonding on an interposer [11], and chip-stacked flip-chip bonding are common photonic-electronic integration approaches [12]. Dense photonic-electronic integration is possible using monolithic technology. However, considering the cost of fabrication and the lack of scaled CMOS transistors in existing monolithic CMOS-SOI photonic-electronic process, multi-chip integration is the most prevalent choice for high-volume production today. Multi-chip integration helps achieving the best functionality of photonic and CMOS ICs each fabricated in a separate technology node (PIC in 90 or 130 nm SOI, electronic IC in bulk FinFET or BiCMOS SiGe process) [13]. The photonic processor core should handle multiple optical and electrical signals such as the optical input signals, optical output signals, driving currents for distributed feedback (DFB) lasers or SOAs, and driving currents for weight control. To manage the signals, integration of multiple chips is essential - an SOI photonic chip as the processing core, and a CMOS chip in an advanced process node for high-speed analog to digital (A/D) and digital to analog (D/A) conversions, SRAM, memory controller, I/Os and digital cores. The low-speed controllers for the SOI photonic devices can be included in the CMOS chip, or included in the SOI photonic chip if the latter is on a monolithic CMOS-SOI photonic-electronic process ([14] as an example). DFB lasers or SOAs can be placed on the same carrier as the photonic chip or on the photonic chip inside a recess cavity, as illustrated in Fig. 3 and connected to the PIC using V-groove fiber coupling [15] or photonic wire bonding [16]. Fig. 3(a) shows an existing technology leveraging the Peripheral Component Interconnect Express (PCIe) interface to connect to a PCB incorporating a

neuromorphic photonic accelerator and a CMOS ASIC including SRAMs, TIAs, and drivers. In this approach, optical sources are placed outside the accelerator PCB on the motherboard. The Optical I/O ports provide an interface for the optical signals. Fig. 3(b) shows the prospect of co-packaging the memory as well. Flip-chip bonding is a feasible approach for integrating the CMOS application-specific integrated circuit (ASIC) and the memory module with the photonic processor core. More details are provided in Section III.A.

A. Optical Interconnects

1) *Chip-to-Fiber Connection:* Conventional light-coupling technologies such as edge coupling and vertical coupling require tedious active alignment [17]. V-grooves allow passive alignment, but occupy a large area on the PIC. Furthermore, a failed fiber attach process will waste a significant amount of good silicon in V-grooves [18]. Thus, it is essential to explore alternative methods for low-loss light coupling with passive alignment.

Photonic wire bonding (PWB) is a technique in which a three-dimensional free-form polymer waveguide is fabricated in-situ by two-photon polymerization [16], [19], [20]. This technology has enabled the demonstration of a highly compact hybrid multi-chip assembly of an optical transceiver [21]. PWB can also be employed to accommodate the complex packaging needs of neuromorphic photonic accelerator architectures for the following reasons: (1) PWB enables dense optical I/O interconnections with optical fibers by utilizing multi-core fibers [20]. (2) Chip-to-chip coupling can be realized with a pitch of $25\ \mu\text{m}$, and could be further reduced to $10\ \mu\text{m}$, which allows for 100 PWBs per millimeter of chip edge [21]. (3) Compared to out-of-plane coupling techniques, PWB interconnects enable the construction of flat packaging, which is helpful in miniaturizing the optical I/O to silicon photonic (SiP) chips via in-plane coupling to optical fibers [20]. (4) PWB allows for connecting components with disparate mode field diameters and material platforms, such as III-V lasers [22], SOAs [23], silicon-on-insulator (SOI) chips, and optical fibers [20]. (5) The insertion loss is low, with chip-to-chip loss of $0.7\ \text{dB} \pm 0.15\ \text{dB}$ [21], chip-to-laser loss of $0.4\ \text{dB} \pm 0.3\ \text{dB}$ [22], and chip-to-fiber loss of $1.6\ \text{dB} \pm 0.13\ \text{dB}$ [20]. (6) The PWB fabrication can be fully automated using existing 3D machine vision techniques [20] which enables high throughput, e.g., 30 s for each chip-to-chip bond [22]. (7) As the PWB process does not require precise ($< 1\ \mu\text{m}$) passive or active alignment of the components to be bonded, the assembly of components is greatly simplified. An illustrative example is shown in Fig. 4 as a laboratory prototype (all of the drivers, TIA, temperature controller, laser, and optical amplifiers are therefore external). The SOI PIC chip has its optical I/O coupled to the fiber array with PWBs. Electrical wire bonds connect the metal pads on the SOI chip and the pins on PCB, although flip-chip bonding must be used for a high pad-count implementation. A Peltier module and a negative temperature coefficient (NTC) temperature sensor is used for precise temperature control.

2) *Laser Integration:* To support parallel analog computing in the optical domain in large-scale neuromorphic photonic




Fig. 4. Assembly of a PIC with photonic wire bonding: (a) Isometric view of the assembly with PCB sub-mount, SOI chip, fiber array, negative temperature coefficient (NTC) temperature sensor, and the Peltier temperature control module. (b) The Peltier module is attached to the back side of the PCB. (c) An implementation of the assembly. (d) The photonic wire bonds couple light between the fiber array and the SOI chip.

accelerators, a scalable and cost-effective method for integrating multiple lasers with silicon photonic PICs is needed. Here we propose a heterogeneous integration approach utilizing wafer-scale SiP fabrication to produce integrated SiP circuits combined with high-quality III-V DFB lasers using photonic wire bonds. Heterogeneous integration of III-V components, specifically lasers, with SiP is a very active area of research, and many impressive techniques for integrating these components on-chip have been demonstrated. These strategies can be classified into two broad categories: 1) Heterogeneous integration of III-V materials on SOI, either through direct wafer bonding [24], or hetero-epitaxy [25], and 2) multi-chip integration, where laser (or SOA) dies are placed on an existing SiP chip [26], [27]. Broadly speaking, direct hetero-epitaxy of III-V lasers on Si can be further separated into two classes: selective epitaxy of III-V lasers on pre-patterned Si/SOI wafers, and blanket epitaxy of III-V lasers on III-V/Si compliant substrates [28]. While the blanket epitaxy approach is the most mature, and has demonstrated lasing properties approaching that on native substrates in the O-band [28], [29], [30], [31], [32], it suffers from the requirement of using thick buffer layers which separate the active gain region from the SiP circuit by several micrometers. This makes efficient coupling of light from the active region of quantum well lasers to the SiP circuit a challenge. Furthermore, quantum well lasers suffer from low laser efficiency. For example, quantum-well-based active region in a GaAs laser, epitaxially grown on silicon, does not provide a

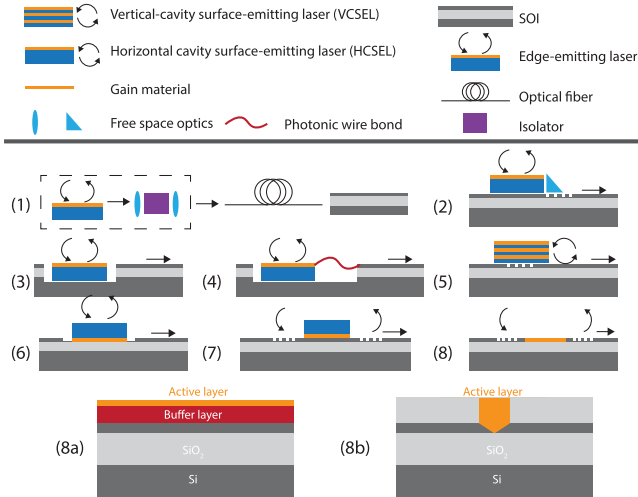


Fig. 5. Methods for integrating lasers with silicon photonics. (1) Traditional external laser approach utilizing free space optics and traditional Faraday isolators to realize high-performance lasers. Light is coupled to the SiP chip through fiber or free space coupling. (2) Laser die bonded on the surface of the SiP chip with free space optics to couple light to the Si layer. (3) Laser die placed p-side up in etched cavity where light is coupled via free space to the SiP circuit. This requires high placement accuracy or active alignment. (4) Laser is placed p-side up in an etched cavity but is connected to the SiP chip via a photonic wire bond. (5) VCSEL is placed on top of the SiP chip, and light is coupled to the SiP circuit via grating couplers. (6) Laser die is flip-chip bonded on a shallow etched region of the chip and coupled via free space to the SiP circuit. This requires high placement accuracy for low insertion loss coupling. (7) Heterogeneous integration approach where the gain material is wafer bonded to the SiP chip and the laser is co-fabricated with the SiP chip. (8) Direct heteroepitaxy where the gain material is grown on the SOI. This can be further subdivided into blanket epitaxy (8a) or selective epitaxy (8b), where the gain material is only grown in the desired region, and thick buffer layers can be avoided. (8a) and (8b) have been adapted from [28].

high laser efficiency [33]. In contrast, quantum dot-based active region in such a laser promises high laser efficiencies [34]. The slope efficiency of GaAs quantum well-based lasers and InAs quantum dot-based lasers are reported as approximately 10^{-5} W/A and 0.3 W/A, respectively [34]. Moreover, several potential promising methods for coupling light from a laser with a quantum dot-based active region to a silicon waveguide has been demonstrated [29]. Alternatively, the selective epitaxy approach solves this problem by growing III-V material in pre-patterned structures on the SiP chip [28], [35]. This enables the gain region to be close enough to the silicon device layer to enable efficient coupling to the SiP circuit. While this approach is promising, all selective epitaxy demonstrations to date have relied on optical excitation rather than electrical driving [28]. Fig. 5 summarizes the methods for integrating lasers with silicon photonics.

Multi-chip integration utilizes a pre-fabricated laser or SOA die placed on or next to the SiP chip. To couple light from the die to the SiP chip, complex fabrication and alignment techniques are required to achieve the sub-micrometer alignment accuracy required for efficient coupling from the laser to the SiP circuit. These challenges have prevented the inclusion of heterogeneously integrated III-V lasers on PICs available from SiP foundries that offer multi-project wafer (MPW) runs. As such, we have opted for the PWB approach detailed earlier in sub-Section II.A.1. [22], [36], which allows for the integration




Fig. 6. Conceptual schematic of multi-chip integration using photonic wire bonds.

of components without the need for high-precision/active alignment techniques [10], [37]. This enables connecting components with arbitrary mode fields in an automated fashion [38]. Integration of lasers and reflective SOAs with SiP chips via PWBs have been demonstrated with side-by-side assemblies [22], [23].

3) *SOA Integration*: One of the key bottlenecks in developing large-scale neuromorphic photonic accelerators is the optical power loss due to the aggregated insertion loss of the photonic components on the PIC and the optical coupling loss [8]. External optical amplifiers such as C-band erbium-doped fiber amplifier and O-band booster optical amplifier are widely used to mitigate propagation and coupling loss. However, external amplifiers are bulky and expensive. SOAs can be used to compensate for the optical propagation loss in large-scale neuromorphic photonic accelerators, by placing the SOA on the neuromorphic photonic accelerator PIC as illustrated conceptually in Fig. 6. SOAs with an on-chip gain of 39 dB and saturation power of 24 dBm have been demonstrated [39].

B. Electrical Interconnects

Neuromorphic photonic accelerator chips require many more low-frequency traces than high-frequency. That is because DC-controlled weights outnumber AC-coupled neurons in most neural networks. For example, if the chip contains a simple, fully connected two-layer neural network with N neurons in the first layer and M neurons in the second, this corresponds to $N + M$ neurons and at least $N \cdot M$ weights between them. To simplify calculations, if the number of neurons in the network is N , the number of weights scale with $\mathcal{O}(N^2)$. This is illustrated in Fig. 1. Consequently, configuring the neural network requires many independent, low-frequency signals, and a few high-frequency ones. The number of interconnects, however, depends on where the analog sources and control circuitry are placed: on-chip or in a multi-chip package. Regardless of the packaging structure, all data plane levels must be designed in tandem. Because integrated photonic and electronic design, packaging, mixed-signal circuit design, and PCB layout span a wide range of technical skills often distributed among multiple engineers, here we offer a few high-level considerations aimed at photonic engineers.

1) *Chip Level Considerations*: Each resonator weight, shown in Fig. 1.(b), takes at least one analog electrical input for weight tuning if the return pass is shared. Optionally, another electrical port is added for weight readback. As the number of neurons increases, the footprint dedicated to weights dominates the chip's overall area. For large networks implemented in

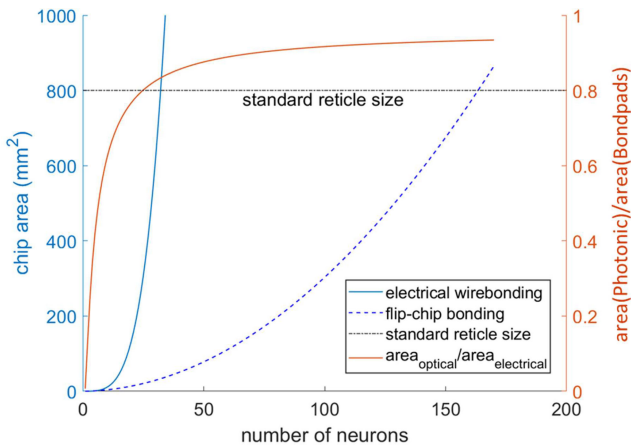


Fig. 7. Scaling analysis for electrical wire bonding and flip-chip bonding as two methods for electrical interconnects in neuromorphic photonic accelerators.

silicon photonics, it is only feasible to accommodate that many metal interconnects if the bump pads to these traces are mostly organized in a 2d-grid array for flip-chip bonding. Furthermore, placing pads next to the constituent elements could reduce the amount of electrical crosstalk. Fig. 7 shows a comparative study for the chip size as a function of number of neurons in two packaging approaches for implementing electrical interconnects, namely electrical wire bonding and flip-chip bonding. For the neurons displayed in Fig. 1, a $N \times N$ photonic network will consist of N^2 optical resonators each with two bond pads, N balanced photodetectors each with two pads, and N microring modulators with up to four bond pads (including return paths). In this study we assumed that most of the chip area on the photonic side is dedicated to the weight bank. Moreover, for wire bonding, we assumed that the maximum allowable number of bond pad rows for staggered wire bonding is two rows around the edge of the chip. To avoid crossing electrical wire bonds, maintaining a minimum length of wires, and observing the keep-out area required for the optical coupling, the bond pads on the photonic chip are placed around the photonic circuit close to the edge of the chip. Considering a pad size of $75 \mu\text{m}$ with a pitch of $100 \mu\text{m}$, microring resonator (MRR) radius of $10 \mu\text{m}$ with a spacing of $100 \mu\text{m}$, our analysis for wire bonding approach shows that with only 23 neurons, the chip size is at the limit of the standard reticle size of 800mm^2 . Flip-chip bonding, however, allows for spreading the metal contacts across the chip, which can accommodate up to 160 neurons considering the solder microbump diameter of $75 \mu\text{m}$ and a pitch of $100 \mu\text{m}$. It should be noted, recent advances in Cu-Cu flip-chip bonding allows for even smaller microbump and a finer pitch down to $10 \mu\text{m}$ [40]. This significantly helps accommodating larger number of neurons. In this analysis, the amount of area occupied by the traces in the electrical wire bonding approach has been neglected. Monolithic approaches with both electronics and photonics on the same die significantly relax the packaging concerns for weight control. Addressing schemes that include multiplexing (such as active matrices for time multiplexing) can further allow higher neuron densities for monolithic implementations, although on-chip drivers come with their own area and bond pad requirements [41].

Photonic devices, like modulators and detectors, can be operated at multiple gigahertz speeds. The neuron's modulator, for example, may require a high-speed electrical input for external modulation via a low-impedance metal trace from the edge of the chip to its location. In small silicon photonic circuits, this trace can be considered as a lumped element. However, the chip boundary creates an unavoidable impedance discontinuity, which causes reflections. The effect of wire bond parasitic inductance can be reduced via impedance-matching capacitors near the bond pads on the chip [42] or a matching network on the PCB [43].

The electrical interconnects must be placed far from the optical I/O's region. The reason is that the 3D-printing polymer epoxy for photonic wire bonds requires a large clearance from any active surface. This means that electrical bond pads placed on the chip's surface must be separated by about $800 \mu\text{m}$ from the tapered waveguides, despite tapered waveguides themselves being much smaller ($\sim 100 \mu\text{m}$). Fortunately, the scaling of the number of electrical interconnects, coupled with optical multiplexing via wavelength, enables a smaller number of optical I/Os than electrical pads. Hence, at the chip level, it is simpler to layout optical ports where convenient and route to the components as required.

Heat dissipation causes thermally-sensitive devices to interact with each other, imposing more space considerations. In silicon photonics, while free-carrier effects can be used to locally tune elements, the stronger thermo-optic effect is often used for index tuning. But because the chip acts as a thermal reservoir, this results in thermal crosstalk between components. This crosstalk can be mitigated by keeping heated components apart; for instance, a minimum of $10 \mu\text{m}$ is quoted for microdisk modulators [44]. Some processes offer extra thermal management by allowing deep trenches to be etched into the substrate, which also demonstrate up to $20\times$ improved heating efficiency [45].

The reconfiguration rate of the neuromorphic photonic accelerator circuit largely depends on the weight memory access and its update. Weights can be stored in the SRAM of the CMOS ASIC controller. The controller maintains a stable optical weight configuration by setting an analog voltage (current) to each weight-tuning device on the PIC for weight-stationary tasks. The setting depends on calibration parameters, and the controller must compensate for environmental fluctuations such as temperature variation. These fluctuations happen at a slow timescale, but as the weight count ($\mathcal{O}(N^2)$) increases, the algorithmic complexity for calibration and control also increases with $\mathcal{O}(N^2)$ for circuits [46]. The circuits require additional local sense ports on each weighting unit to provide a feedback signal that can be corrected dynamically by the controller, which makes the access slow and costs a large amount of energy per bit of control.

2) *Board-Level Considerations:* For rapid prototyping, where a CMOS ASIC controller is still not available, a PCB interposer board can be used to make low-noise electrical connections from the PIC to the control circuit (e.g. precision source-measure units such as the Keithley 2400). A typical interposer (as shown in Fig. 8) is designed with chip sockets for mounting the photonic chip, and electrical connectors for mating with cables linking

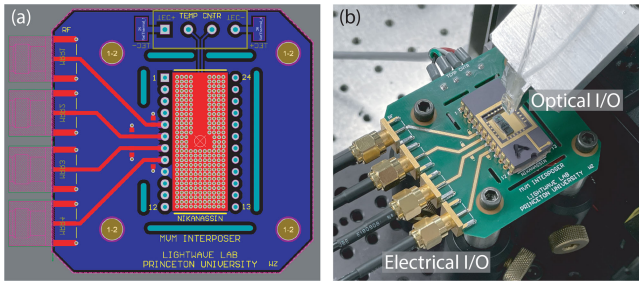


Fig. 8. Example of an interposer design. (a) Design view in the PCB CAD software. (b) Footage of the interposer mounted with a photonic chip and fully connected with optical and electrical I/O.

to other control circuits. Some interposers are also populated with a thermistor and thermal-electrical cooler for implementing the temperature control for the chip carrier. Since the major functionality of the interposer is transferring electrical control signals from the connectors to the chip pins, to assure the signal integrity, careful considerations on grounding and signals traces need to be carried out when designing the interposer.

As for grounding, a sufficiently wide ground plane is often a good practice. For robustness to electromagnetic interference (EMI), the ground plane should be allocated right under the signal traces and at the adjacent PCB layer. This helps to reduce the electrical interference since the electromagnetic fields of the coupled noise of the signal path and that of the corresponding return path in the ground plane tend to cancel each other in such a layout. The ground planes are usually allocated to multiple PCB layers, and it is also a good practice to apply via-stitching across the board that lowers the impedance between them. Finally, placing via arrays along the board edges further fosters the suppression of the EM radiation and coupling, helping with electromagnetic compatibility (EMC) specifications.

For applications where the MRR weights have to be updated at MHz or higher rates, dedicated signal pairs are recommended (one for signal and one for the ground) for signal integrity. Laying out signal traces in pairs helps lower the cross-talk between MRR channels and reduces EMI.

3) *Biasing Circuit*: The electrical parasitics between various components in a neuromorphic photonic accelerator must be accounted for carefully. For example, when supplying current to a thermo-optic heater, any common-mode resistance between multiple heaters, e.g., due to shared ground, will create crosstalk. This common-mode resistance makes sensing circuits more challenging to design since their sensing voltage will depend on other channels' actuation. This is a concern also for proof-of-concept designs where a PIC is controlled not using CMOS circuits but with benchtop equipment. Any resistance between the chip's shared ground and the PCB ground plane will cause a voltage differential between them. This can be mitigated through extra high-impedance sensing traces and more sophisticated control. Isolating ground planes between the controller and the power supply can significantly reduce parasitic voltage levels [47].

Some circuits in neuromorphic photonic accelerators require multi-point DC biasing of AC components. A good example is

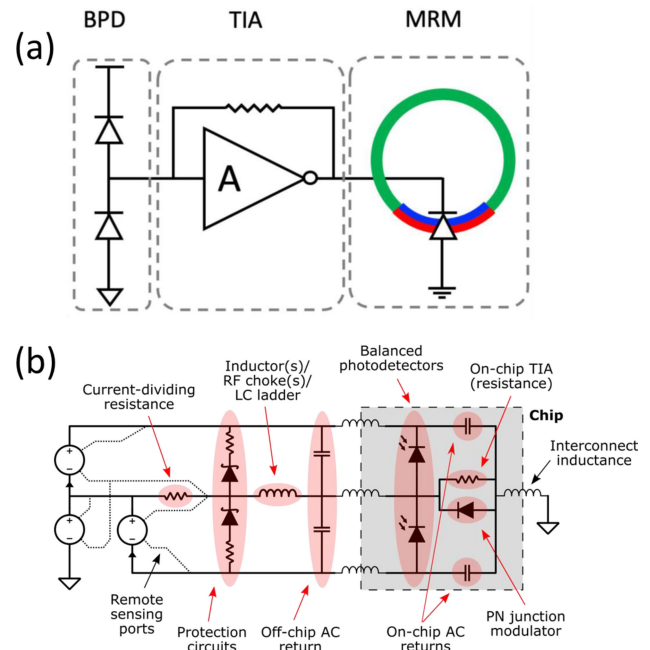


Fig. 9. Modulator photonic neuron alongside a DC-coupled biasing circuit: (a) Monolithic implementation; (b) Co-integration of PIC with ASIC. Inductive components and capacitors allow AC photocurrent generated on the chip to modulate a reverse-biased PN junction modulator by generating a voltage drop across a resistive on-chip TIA. At very low frequencies, a “current-dividing” resistance fulfills this role. Voltage sources use remote sensing to maintain a specified voltage across the photodiodes in the presence of the series resistance. A specific DC bias accounting for DC photocurrent can be imposed on the neuron. Schottky diodes in series with the photodiodes protect unbiased photodiodes from spurious forward biasing.

an optical-electrical-optical (OEO) unit [48]. When neurons in a recurrent neural network can respond to both low-speed and high speed inputs, it has the ability to process and recognize complex temporal patterns [49]. This can be achieved if the circuit has a flat frequency response from DC to some cutoff bandwidth. For a fully electronic-photonic single-chip monolithic implementation, the microring modulator (MRM) can be easily driven by an active TIA which buffers the differential current of the balanced photodetectors and provides the desired gain as shown in Fig. 9(a).

For a co-integrated solution with separate PIC and CMOS ASIC, the bias configuration is more involved. The optical signal's DC component impinging on the photodetectors should have the desirable effect of offsetting a configurable DC bias of the modulator. Setting this bias voltage electrically is tricky. An extra current or voltage source across the modulator will end up in series with the current source from the photodiodes in the transimpedance branch. A way around this is to monitor the optical DC power without such an extra source and make the extra source conditioned on this knowledge to impose the correct DC pedestal in addition to any desired bias, a form of current-controlled current source as shown in Fig. 9(b).

At multigigahertz frequencies, an efficient O-E conversion relies on making sure that the modulator is the lowest-impedance path with respect to the PD. This can be achieved by adding

RF chokes between bias voltage sources and the photodetector. Multiple such components, each operating over some range of frequency in a LC ladder operation (assisted by de-Q'ing with parallel resistors), can be employed to extend the frequency range of the inductive regime [50]. Short AC return paths provided by (small) on-chip and (large) off-chip capacitors at each LC ladder stage further improve circuit behaviour [48]. However, the footprint of these components on-chip and on-board would negatively impact neuron scalability, and it clearly demonstrates why a two-chip solution is difficult to scale.

C. Power Management

The total amount of power consumed for neuromorphic photonic accelerators can be estimated based on the speed and scale of the neuromorphic photonic accelerator units. The need to compute large matrices necessitates tiling several photonic units, where a photonic unit conducts a vector matrix multiplication. As an example, performing such an optical multiplication for an input vector with a size of 1024×1 requires tiling several photonic units. For a photonic unit with a size of $N = 85$ limited by the electrical power of the laser, 12 photonic tiles would be needed. For $N > 85$, the maximum rated laser power cannot compensate for the accrued amount of insertion loss in the network [8]. Therefore, for a photonic processor operating at a data rate (DR) of 10 GS/s with an energy efficiency (E) of $\sim 1pJ/OP$, the total power consumption can be estimated as $P = 12 \times 2N^2DR \times E(J/OP) \approx 1.7kW$ [8]. Assuming 2 sets of Peltier TEC coolers (driven by $\sim 10V$, sourcing $\sim 4A$), heat sinks and air fans ($\sim 30W$) to be used to maintain the temperature of the photonic cores and PCB electronics, a total cooling power of 140 W is dissipated, amounting to a total estimated power of $\sim 1.9kW$.

D. Thermal Management

Components of a neuromorphic photonic accelerator module such as the core photonic processor, DFB lasers, and SOAs are susceptible to thermal fluctuations [51], [52]. This means that a slight change in the refractive index induced by ambient temperature or unwanted thermal crosstalk between the adjacent components can significantly affect the performance of the photonic module. CMOS ASICs are relatively less sensitive to the changes in temperature. Lasers and CMOS ASICs generate significant heat and can influence the performance of the other components [10]. InP, as a constituent material of DFB lasers and SOAs, has a thermo-optic coefficient of $2.5 \times 10^{-4}/C$. [53]. These devices also show a heat flux of over $10^2 W/cm^2$ [54] which means that they can easily transfer heat to peripheral components. Apart from adequate spacing between the sensitive devices within the neuromorphic photonic accelerators [44], a temperature controller, a temperature sensor, and a Peltier thermoelectric cooler (TEC) is needed for photonic processors employing WDM [10]. A temperature control algorithm is needed on the CMOS ASIC to hold the state of the weight tuning components and compensate for the thermal fluctuations [55], [56]. A temperature control algorithm

is also implemented on the CMOS ASIC or micro-controller to sense the temperature from a temperature sensor (e.g., NTC) and trigger the TEC element to stabilize the temperature of the module within the allowed range similar to pluggable transceivers [57]. The choice of the module holder, heat sinks and thermal conductivity of the epoxy resin used for securing the components inside the module plays a vital role in the heat dissipation [51].

III. NETWORK AND MEMORY CONSIDERATIONS

Memory modules are crucial for inference and learning using neuromorphic photonic accelerators as the neural networks scale up to handle massive data flows and instructions [51]. One challenge in the practical applications of neuromorphic photonic accelerators is the memory wall which is the limited data rate at which the processor can communicate with the memory interface due to the gap between processing speed and the speed of accessing DRAM [58]. Although the inherent parallelism of photonics can greatly improve the computing speed, the performance of accelerators, either electronics [59] or photonics [60], is dominated by the interconnection bandwidth between their memory modules and processing units. This is because the advantage of the high-speed photonic computing core can only be realized if the high-speed serial data can be fed to it from the memory. The improvement in processing units alone could further enlarge the performance gap between the processing units and memory modules because of the latency to fetch data from memory modules and the limited interconnection bandwidth. Given the high throughput of SiP accelerators, memory interfacing circuits should be introduced to fetch and serialize data at the desired data rate. Moreover, the scalability of neuromorphic photonic accelerators is also limited by the available pins to memory channels. This is often called pin wall, which is encountered by electronic accelerators [61]. Simply increasing the number of channels and controllers for memory access could conflict with energy constraints [59]. Meanwhile, an energy-efficient high-speed data transfer between the photonic processor and the corresponding memory modules is a key to enable high-performance computing with massive data transmissions. 2.5D [62] and 3D [63] stacked memory modules using a silicon interposer [64] with embedded active SiP interconnects have been proposed to optically link the processor and its memory modules. Nevertheless, such schemes are yet to be realized.

One possible solution to ease the memory wall is to implement the processor-memory interconnection using a photonic link [65], as shown in Fig. 10. In this case, the data transfer speed no longer depends on the speed limitation of the DDR interface. The processed signal propagates through a waveguide or an optical fiber and is de-multiplexed in the ring-based WDM module. The photodetectors convert light into electrical signals which is written into the memory module. On the other hand, to retrieve data from the memory module, modulators convert electrical signals to optical signals which are later multiplexed in a ring-based WDM module and processed in the photonic processing units.




Fig. 10. An optical link connecting the memory to the SiP accelerator.

IV. CONCLUSION

The advances in silicon photonics have enabled the recent surge in neuromorphic photonic processors, which can potentially outperform electronic counterparts in terms of latency and bandwidth. Efficient data processing and data movement between a neuromorphic photonic accelerator and memory modules are directly linked to the packaging strategies. Configuring the weights and thermal stabilization of the PIC requires a CMOS controller with SRAM. Feeding the high-speed serial input data to the PIC requires efficient data movement from the DRAM. Hence, the neuromorphic photonic accelerator must accommodate many I/O pads for handling DC, low-speed and high-speed signals between the PIC and the CMOS ASIC, in addition to efficient DRAM access. Flip-chip packaging is the practical solution in the near-term before wafer-level integration becomes commercially viable. Monolithic CMOS-SOI photonic-electronic platforms significantly reduce the number of I/O pads required and hence are an attractive solution in the near term. Because a neuromorphic photonic accelerator must still include high-performance digital processors, a separate FinFET CMOS chip is still required. In the longer term, wafer-level integration of PICs with FinFET CMOS chips would significantly improve performance. Since the power dissipation exceeds kW, thermal management and reliability is the most crucial consideration. In the near term, the laser would likely be physically separated for thermal and ease of packaging considerations. Since V-grooves require a large area, photonic wire bonding is a viable solution for medium-volume production which is not limited by the position of the optical source. Advances in laser technology and heterogeneous/hybrid integration would help improve efficiency and reliability. Replacing electrical interconnects with optical interconnects could facilitate data movement and ease the pin wall and power wall issues associated with the scaling of neuromorphic photonic accelerators. New architectures and processor-memory communication schemes relying on optical links for data movement seems to be a promising path.

ACKNOWLEDGMENT

The authors would like to thank Mustafa Hammood for his help.

REFERENCES

- [1] N. P. Jouppi et al., "In-datacenter performance analysis of a tensor processing unit," in *Proc. 44th Annu. Int. Symp. Comput. Architecture*, 2017, pp. 1–12.
- [2] S. Shekhar, "Silicon photonics: A brief tutorial," *IEEE Solid-State Circuits Mag.*, vol. 13, no. 3, pp. 22–32, Summer 2021.
- [3] M. A. Nahmias et al., "Photonic multiply-accumulate operations for neural networks," *IEEE J. Sel. Topics Quantum Electron.*, vol. 26, no. 1, pp. 1–18, Jan./Feb. 2020.
- [4] 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. Lightw. Technol.*, vol. 32, no. 21, pp. 4029–4041, 2014.
- [5] A. N. Tait et al., "Neuromorphic photonic networks using silicon photonic weight banks," *Sci. Rep.*, vol. 7, no. 1, pp. 1–10, 2017.
- [6] F. Shokraneh, S. Geoffroy-Gagnon, M. S. Nezami, and O. Liboiron-Ladouceur, "A single layer neural network implemented by a 4×4 MZI-based optical processor," *IEEE Photon. J.*, vol. 11, no. 6, pp. 1–12, Dec. 2019.
- [7] Y. Shen et al., "Deep learning with coherent nanophotonic circuits," *Nature Photon.*, vol. 11, no. 7, pp. 441–446, 2017.
- [8] M. A. Al-Qadasi, L. Chrostowski, B. J. Shastri, and S. Shekhar, "Scaling up silicon photonic-based accelerators: Challenges and opportunities," *APL Photon.*, vol. 7, no. 2, 2022, Art. no. 020902.
- [9] B. J. Shastri et al., "Photonics for artificial intelligence and neuromorphic computing," *Nat. Photon.*, vol. 15, no. 2, pp. 102–114, 2021.
- [10] J. S. Lee et al., "Meeting the electrical, optical, and thermal design challenges of photonic-packaging," *IEEE J. Sel. Topics Quantum Electron.*, vol. 22, no. 6, pp. 409–417, Nov./Dec. 2016.
- [11] A. H. Ahmed et al., "A dual-polarization silicon-photonic coherent transmitter supporting 552 Gb/s/wavelength," *IEEE J. Solid-State Circuits*, vol. 55, no. 9, pp. 2597–2608, Sep. 2020.
- [12] W. Bogaerts and L. Chrostowski, "Silicon photonics circuit design: Methods, tools and challenges," *Laser Photon. Rev.*, vol. 12, no. 4, 2018, Art. no. 1700237.
- [13] A. H. Ahmed et al., "A 6 V swing 3.6% THD gt;40 GHz driver with $4.5 \times$ bandwidth extension for a 272 Gb/s dual-polarization 16-QAM silicon photonic transmitter," in *Proc. IEEE Int. Solid-State Circuits Conf.*, 2019, pp. 484–486.
- [14] J. E. Proesel et al., "A monolithically integrated silicon photonics 8×8 switch in 90 nm soi cmos," in *Proc. IEEE Symp. VLSI Circuits*, 2020, pp. 1–2.
- [15] T. Barwicz et al., "A novel approach to photonic packaging leveraging existing high-throughput microelectronic facilities," *IEEE J. Sel. Top. Quantum Electron.*, vol. 22, no. 6, pp. 455–466, Nov./Dec. 2016.
- [16] N. Lindenmann et al., "Photonic wire bonding: A novel concept for chip-scale interconnects," *Opt. Exp.*, vol. 20, no. 16, pp. 17667–17677, 2012.
- [17] A. Mekis et al., "A grating-coupler-enabled CMOS photonics platform," *IEEE J. Sel. Top. Quantum Electron.*, vol. 17, no. 3, pp. 597–608, May/Jun. 2010.
- [18] R. Mahajan et al., "Co-packaged photonics for high performance computing: Status, challenges and opportunities," *J. Lightw. Technol.*, vol. 40, no. 2, pp. 379–392, 2021.
- [19] S. Maruo, O. Nakamura, and S. Kawata, "Three-dimensional microfabrication with two-photon-absorbed photopolymerization," *Opt. Lett.*, vol. 22, no. 2, pp. 132–134, 1997.
- [20] N. Lindenmann et al., "Connecting silicon photonic circuits to multicore fibers by photonic wire bonding," *J. Lightw. Technol.*, vol. 33, no. 4, pp. 755–760, 2015.
- [21] M. Blaicher et al., "Hybrid multi-chip assembly of optical communication engines by in situ 3D nano-lithography," *Light Sci. Appl.*, vol. 9, no. 1, pp. 1–11, Dec. 2020.
- [22] M. R. Billah et al., "Hybrid integration of silicon photonics circuits and inp lasers by photonic wire bonding," *Optica*, vol. 5, no. 7, pp. 876–883, 2018.
- [23] Y. Xu et al., "Inp/silicon hybrid external-cavity lasers (ECL) using photonic wirebonds as coupling elements," in *Proc. Opt. Fiber Commun. Conf.*, Optical Society of America, 2020, pp. M4H–6.
- [24] D. Caimi, P. Tiwari, M. Sousa, K. E. Moselund, and C. B. Zota, "Heterogeneous integration of III–V materials by direct wafer bonding for high-performance electronics and optoelectronics," *IEEE Trans. Electron Devices*, vol. 68, no. 7, pp. 3149–3156, Jul. 2021.
- [25] Y. Han et al., "Bufferless $1.5 \mu\text{m}$ III–V lasers grown on si-photonics 220 nm silicon-on-insulator platforms," *Optica*, vol. 7, no. 2, pp. 148–153, 2020.
- [26] B. Haq et al., "Micro-transfer-printed III–V-on-silicon c-band soas with 17 db gain," in *Proc. IEEE Photon. Conf.*, 2019, pp. 1–2.

- [27] B. Haq et al., "Micro-transfer-printed III-V-on-silicon c-band distributed feedback lasers," *Opt. Exp.*, vol. 28, no. 22, pp. 32793–32801, 2020.
- [28] Y. Han and K. M. Lau, "III–V lasers selectively grown on (001) silicon," *J. Appl. Phys.*, vol. 128, no. 20, 2020, Art. no. 200901.
- [29] J. C. Norman, D. Jung, Y. Wan, and J. E. Bowers, "Perspective: The future of quantum dot photonic integrated circuits," *APL Photon.*, vol. 3, no. 3, 2018, Art. no. 030901.
- [30] D. Jung et al., "Impact of threading dislocation density on the lifetime of inas quantum dot lasers on Si," *Appl. Phys. Lett.*, vol. 112, no. 15, 2018, Art. no. 153507.
- [31] R. Jones, "Overview and future challenges on iii-v integration technologies in silicon photonics platform," in *Proc. Opt. Fiber Commun. Conf.*, Optical Society of America, 2021, pp. M5A–1.
- [32] H. Park, A. W. Fang, S. Kodama, and J. E. Bowers, "Hybrid silicon evanescent laser fabricated with a silicon waveguide and III-V offset quantum wells," *Opt. Exp.*, vol. 13, no. 23, pp. 9460–9464, 2005.
- [33] A. Y. Liu and J. Bowers, "Photonic integration with epitaxial III–V on silicon," *IEEE J. Sel. Topics Quantum Electron.*, vol. 24, no. 6, pp. 1–12, Nov/Dec. 2018.
- [34] A. Y. Liu, S. Srinivasan, J. Norman, A. C. Gossard, and J. E. Bowers, "Quantum dot lasers for silicon photonics," *Photon. Res.*, vol. 3, no. 5, pp. B1–B9, 2015.
- [35] J. Z. Li et al., "Defect reduction of gaas epitaxy on Si (001) using selective aspect ratio trapping," *Appl. Phys. Lett.*, vol. 91, no. 2, 2007, Art. no. 021114.
- [36] T. Hoose et al., "Multi-chip integration by photonic wire bonding: Connecting surface and edge emitting lasers to silicon chips," in *Proc. Opt. Fiber Commun. Conf.*, Optical Society of America, 2016, p. M21.7.
- [37] B. Pezeshki et al., "High performance MEMS-based micro-optic assembly for multi-lane transceivers," *J. Lightw. Technol.*, vol. 32, no. 16, pp. 2796–2799, 2014.
- [38] P.-I. Dietrich et al., "In situ 3D nanoprinting of free-form coupling elements for hybrid photonic integration," *Nat. Photon.*, vol. 12, no. 4, pp. 241–247, 2018.
- [39] S. Liu et al., "High efficiency, high gain and high saturation output power quantum dot soas grown on Si and applications," in *Proc. Opt. Fiber Commun. Conf.*, Optical Society of America, 2020, pp. T4H–3.
- [40] Y. Ma, A. Roshanghias, and A. Binder, "A comparative study on direct cu-cu bonding methodologies for copper pillar bumped flip-chips," *J. Mater. Sci.: Mater. Electron.*, vol. 29, no. 11, pp. 9347–9353, 2018.
- [41] V. Stojanović et al., "Monolithic silicon-photonic platforms in state-of-the-art CMOS SOI processes," *Opt. Exp.*, vol. 26, no. 10, pp. 13106–13121, 2018.
- [42] B. L. Meres and S. Khairi, "Broadband impedance matching for inductive interconnect in VLSI packages," in *Proc. Int. Conf. Comput. Des.*, 2005, pp. 683–688.
- [43] N. Sarmah et al., "A wideband fully integrated sige chipset for high data rate communication at 240 ghz," in *Proc. 11th Eur. Microw. Integr. Circuits Conf.*, 2016, pp. 181–184.
- [44] C. T. DeRose et al., "Thermal crosstalk limits for silicon photonic DWDM interconnects," in *Proc. Opt. Interconnects Conf.*, 2014, pp. 125–126.
- [45] A. Masood et al., "Comparison of heater architectures for thermal control of silicon photonic circuits," in *Proc. 10th Int. Conf. Group IV Photon.*, 2013, pp. 83–84.
- [46] A. N. Tait et al., "Feedback control for microring weight banks," *Opt. Exp.*, vol. 26, no. 20, pp. 26422–26443, Oct. 2018.
- [47] H. W. Ott, *Electromagnetic Compatibility Engineering*. Hoboken, NJ, USA: Wiley, Aug. 2009.
- [48] T. F. de Lima et al., "Noise analysis of photonic modulator neurons," *IEEE J. Sel. Topics Quantum Electron.*, vol. 26, no. 1, pp. 1–9, Jan./Feb. 2020.
- [49] S. C. Quax, M. D'Asaro, and M. A. J. van Gerven, "Adaptive time scales in recurrent neural networks," *Sci. Rep.*, vol. 10, no. 11360, pp. 1–14, Jul. 2020.
- [50] G. W. Johnson, "Wideband bias tee," Aug. 11, 2008. [Online]. Available: https://www.na60.com/main/Bias_Tee.html
- [51] K.-I. Kitayama et al., "Novel frontier of photonics for data processing-photonic accelerator," *APL Photon.*, vol. 4, no. 9, 2019, Art. no. 090901.
- [52] K. Padmaraju and K. Bergman, "Resolving the thermal challenges for silicon microring resonator devices," *Nanophotonics*, vol. 3, no. 4/5, pp. 269–281, 2014.
- [53] G. Gilardi, W. Yao, H. R. Haghighi, M. K. Smit, and M. J. Wale, "Substrate thickness effects on thermal crosstalk in inp-based photonic integrated circuits," *J. Lightw. Technol.*, vol. 32, no. 17, pp. 3061–3066, 2014.
- [54] R. Enright et al., "A vision for thermally integrated photonics systems," *Bell Labs Tech. J.*, vol. 19, pp. 31–45, 2014.
- [55] H. Jayatilaka, H. Shoman, L. Chrostowski, and S. Shekhar, "Photoconductive heaters enable control of large-scale silicon photonic ring resonator circuits," *Optica*, vol. 6, no. 1, pp. 84–91, 2019.
- [56] H. Jayatilaka et al., "Wavelength tuning and stabilization of microring-based filters using silicon in-resonator photoconductive heaters," *Opt. Exp.*, vol. 23, no. 19, pp. 25084–25097, 2015.
- [57] A. Brilliant, "Laser power and temperature control loops," *Digit. Analog Fiber Opt. Commun. CATV FTTx Appl.*, vol. 174, pp. 591–625, 2008.
- [58] W. A. Wulf and S. A. McKee, "Hitting the memory wall: Implications of the obvious," *ACM SIGARCH Comput. Architecture News*, vol. 23, no. 1, pp. 20–24, 1995.
- [59] M. A. Alves, P. C. Santos, F. B. Moreira, M. Diener, and L. Carro, "Saving memory movements through vector processing in the dram," in *Proc. Int. Conf. Compilers, Architecture Synth. Embedded Syst.*, 2015, pp. 117–126.
- [60] C. Cole, "Optical and electrical programmable computing energy use comparison," *Opt. Exp.*, vol. 29, no. 9, pp. 13153–13170, 2021.
- [61] P. Siegl, R. Buchty, and M. Berekovic, "Data-centric computing frontiers: A survey on processing-in-memory," in *Proc. 2nd Int. Symp. Memory Syst.*, 2016, pp. 295–308.
- [62] K. Wen et al., "Silicon photonic memory interconnect for many-core architectures," in *Proc. IEEE High Perform. Extreme Comput. Conf.*, 2016, pp. 1–7.
- [63] A. N. Udipi, N. Muralimanohar, R. Balasubramanian, A. Davis, and N. P. Jouppi, "Combining memory and a controller with photonics through 3d-stacking to enable scalable and energy-efficient systems," *ACM SIGARCH Comput. Architecture News*, vol. 39, no. 3, pp. 425–436, 2011.
- [64] Y. Arakawa, T. Nakamura, Y. Urino, and T. Fujita, "Silicon photonics for next generation system integration platform," *IEEE Commun. Mag.*, vol. 51, no. 3, pp. 72–77, Mar. 2013.
- [65] A. Hadke, T. Benavides, R. Amirtharajah, M. Farrens, and V. Akella, "Design and evaluation of an optical cpu-dram interconnect," in *Proc. IEEE Int. Conf. Comput. Des.*, 2008, pp. 492–497.



Mohammadreza Sanadgol Nezami received the Ph.D. degree in nanophotonics from the Department of Electrical and Computer Engineering, the University of Victoria, BC, Canada, in 2016, where he worked on the design, fabrication, and characterization of metasurfaces for the enhanced third-order nonlinear optical response, namely third-harmonic generation. In January 2017, he joined the Photonic DataCom team with McGill University, Montreal, QC, Canada, as a Postdoctoral Research Fellow. He was conducting research on energy-efficient high-speed optoelectronic receiver design and silicon photonics optical receivers for advanced data communication systems, and electronic-photonic co-designed integrated circuits for data communications. Furthermore, he was involved with the characterization of silicon photonics MZI-based optical processors. Moreover, he was involved in developing a comprehensive modeling platform for short-reach VCSEL-based optical interconnects in an electronic design automation tool, namely Cadence Virtuoso, which led to an industrial internship with Reflex Photonics Inc. (now Smiths Interconnect) in Montreal, QC. He is currently a Researcher with the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC, Canada. His research interests include neuromorphic photonics and optical computing.




Thomas Ferreira de Lima received the B.Sc. and the Ingénieur Polytechnicien master's degrees in 2016 from Ecole Polytechnique, Palaiseau, France, with a focus on physics for optics and nanosciences. He is currently working toward the Ph.D. degree in electrical engineering with the Lightwave Communications Research Lab, Department of Electrical Engineering, Princeton University, Princeton, NJ, USA. His research interests include real-time integrated photonic systems; nonlinear signal processing with photonic devices; spike-timing based processing, ultrafast cognitive computing, and dynamical light-matter neuro-inspired learning and computing.




Matthew Mitchell received the B.Sc. (with honors) and the Ph.D. degrees in physics from The University of Calgary, AB, Canada, in 2012 and 2020, respectively. He is currently a Postdoctoral Research Fellow with the Stewart Blusson Quantum Matter Institute, Vancouver, BC, Canada, The University of British Columbia, Vancouver. He was awarded the Governor General's Gold Medal for his Ph.D. thesis work on fabricating photonic devices from single-crystal diamond, which he utilized to demonstrate optomechanically mediated wavelength conversion and amplification of light. This work was an important step towards building a quantum computer based on solid-state qubits like the nitrogen vacancy center in diamond. He also co-founded the University of Calgary nanoGroup in 2013, which is a student run organization built to cultivate and support collaboration within the nanotechnology community on campus. He has authored or coauthored 12 journal articles and 20 refereed conference proceedings including two Postdeadline selections, and three invited talks. His research interests in cavity optomechanics, silicon photonics, photonic integrated circuits, optical packaging, and nanofabrication techniques,



Shangxuan Yu received the B.A.Sc and B.Eng. degrees in material engineering from the University of British Columbia, Vancouver, BC, Canada and Chongqing University, Chongqing, China, respectively, where he is currently working toward the M.A.Sc degree in electrical and computer engineering. His research interests include on-chip laser integration and compact silicon photonics packaging with photonics wire bonding.



Jing Wang received the B.A.Sc. degree in electrical engineering and the M.A.Sc. degree in mechanical engineering from The University of British Columbia, Vancouver, BC, Canada, in 2018 and 2020, respectively, where he is currently working toward the Ph.D. degree with the Department of Electrical Engineering, The University of British Columbia. He worked on the robotics vision and transfer learning during his M.A.Sc studies. His research interests include machine learning and silicon photonics.



Simon Bilodeau (Graduate Student Member, IEEE) received the B.Sc. and M.Sc. degrees in physics from McGill University, Montreal, QC, Canada, where he conducted nanoelectronics research with ultralow temperatures. He is currently working toward the Ph.D. degree in electrical and computer engineering with the Lightwave Laboratory, Princeton University, Princeton, NJ, USA. His research interests include nanophotonics, optoelectronics, integrated photonic systems, and neuromorphic computing.



Weipeng Zhang (Graduate Student Member, IEEE) received the B.S. degree in physics from Beihang University, Beijing, China, and the M.S. degree in optical engineering from Tsinghua University, Beijing, China. He was previously doing research on precision measurement utilizing the optical frequency comb, including ranging, surface metrology, and spectroscopy. He currently pivots himself towards a more exciting and promising field, silicon photonics. In Prof. Prucnal's lab, he is currently researching Blind Source Separation using the photonic integrated

chip.

Mohammed Al-Qadasi received the M.Sc. degree in electrical engineering electro-physics from KAUST University, Thuwal, Saudi Arabia, where he worked on the design of high speed, low power CMOS circuits and efficient successive approximation ADCs. He is currently working toward the Ph.D. degree with the Department of Electrical and Computer Engineering, The University of British Columbia, Vancouver, BC, Canada. His research interests include integrated silicon photonic circuits for neuromorphic computing and biosensing.


Iman Taghavi received the B.Sc. and M.Sc. degrees in electrical engineering from the Amirkabir University of Technology, Tehran, Iran, in 2006 and 2008, respectively. He is currently working toward the Ph.D. degree with the Department of Electrical Engineering, The University of British Columbia, Vancouver, BC, Canada. He joined the University of Illinois at Urbana-Champaign, Champaign, IL, USA, (2011–2012), where he worked on Transistor Lasers. He also worked on Organic Material modulators with Georgia Institute of Technology, Atlanta, GA, USA, (2015–2017). His main research interests include hybrid material and component platforms for optical computing.



Alexander Tofini received the B.A.Sc. degree in physics from the University of Victoria, Victoria, BC, Canada, in 2019. He is currently working toward the M.A.Sc. degree in electrical and computer engineering with the University of British Columbia, Vancouver, BC, Canada. His research mainly consists of automation of silicon photonic component design, with a focus on ring modulators. In addition to this research, he has investigated methods to improve the behavior of polymer modulators through the use of slotted-waveguide and characterized the effects of

optical cross-talk within silicon substrates in hopes of reducing the noise present in photodetectors and filters.

Stephen Lin received the MA.Sc. degree in electrical engineering from the University of British Columbia, Vancouver, BC, Canada, in 2020. He worked on creating a Computational Lithography Model for Silicon Photonic Design and has created libraries for PCM Spiral photonic devices that are included with the open-source SIEPIC-PDK. He is currently managing the Neuromorphic Photonics and Optical Computing projects with the University of British Columbia.



Bhavin J. Shastri (Senior Member, IEEE) received the Honours B.Eng. (with distinction), M.Eng., and the Ph.D. degrees in electrical engineering photonics from McGill University, Montreal, QC, Canada, in 2005, 2007, and 2012, respectively. He is currently an Assistant Professor of engineering physics with Queen's University, Kingston, ON, Canada, and a Faculty Affiliate with the Vector Institute for Artificial Intelligence, Toronto, ON. He was an NSERC and Banting Postdoctoral Fellow (2012–2016) and an Associate Research Scholar (2016–2018) with Princeton

University, Princeton, NJ, USA. He has authored or coauthored more than 70 journal articles and 100 conference proceedings, seven book chapters, and given more than 60 invited talks and lectures including five keynotes and three tutorials. He is a co-author of the book, *Neuromorphic Photonics* Taylor and Francis, CRC Press, 2017, a term he helped coin. His research interests include silicon photonics, photonic integrated circuits, neuromorphic computing, and machine learning. He was the recipient of the 2022 SPIE Early Career Achievement Award and the 2020 IUPAP Young Scientist Prize in Optics for his pioneering contributions to neuromorphic photonics from the ICO, the 2014 Banting Postdoctoral Fellowship from the Government of Canada, the 2012 D. W. Ambridge Prize for the top graduating Ph.D. student at McGill, an IEEE Photonics Society 2011 Graduate Student Fellowship, 2011 NSERC Postdoctoral Fellowship, 2011 SPIE Scholarship in Optics and Photonics, 2008 NSERC Alexander Graham Bell Canada Graduate Scholarship, including the Best Student Paper awards at the 2014 IEEE Photonics Conference, 2010 IEEE Midwest Symposium on Circuits and Systems, the 2004 IEEE Computer Society Lance Stafford Larson Outstanding Student Award, and the 2003 IEEE Canada Life Member Award. He is a Senior Member of Optica (formerly OSA).

Paul R. Prucnal (Life Fellow, IEEE) received the A.B. degree (summa cum laude) in mathematics and physics from Bowdoin College, Brunswick, ME, USA, the M.S., M.Phil., and the Ph.D. degrees in electrical engineering from Columbia University, New York, NY, USA. After his Doctorate, he joined the faculty with Columbia University, where, he is a Member of the Columbia Radiation Laboratory, he performed groundbreaking work in OCDMA and self-routed photonic switching. In 1988, he joined the faculty with Princeton University, Princeton, NJ, USA. His research on optical CDMA initiated a new research field in which more than 1000 papers have since been published, exploring applications ranging from information security to communication speed and bandwidth. In 1993, he invented the Terahertz Optical Asymmetric Demultiplexer, the first optical switch capable of processing terabit per second pulse trains. He is the author of the book *Neuromorphic Photonics* and the editor of the book *Optical Code Division Multiple Access: Fundamentals and Applications*. He was an Area Editor of IEEE TRANSACTIONS ON COMMUNICATIONS. He has authored or coauthored more than 350 journal articles and book chapters and holds 28 U.S. patents. He is a Fellow of the Optical Society of America and the National Academy of Inventors, and a Member of honor societies including Phi Beta Kappa and Sigma Xi. He was the recipient of the 1990 Rudolf Kingslake Medal for his paper entitled Self-routing photonic switching with optically-processed control, the Gold Medal from the Faculty of Mathematics, Physics and Informatics with the Comenius University, Bratislava, Slovakia, for Leadership in the field of Optics 2006 and has received multiple teaching awards at Princeton, including the E-Council Lifetime Achievement Award for Excellence in Teaching, the School of Engineering and Applied Science Distinguished Teacher Award, The President's Award for Distinguished Teaching. He has been instrumental in founding the field of Neuromorphic Photonics and developing the photonic neuron, a high-speed optical computing device modeled on neural networks, and integrated optical circuits to improve wireless signal quality by cancelling radio interference.

Lukas Chrostowski (Senior Member, IEEE) was born in Poland. He received the B.Eng. degree in electrical engineering from McGill University, Montreal, QC, Canada, and the Ph.D. degree in electrical engineering and computer science from the University of California, Berkeley, CA, USA. He is a Professor of electrical and computer engineering with the University of British Columbia, Vancouver, BC, Canada. He was the Co-Director of the University of British Columbia AMPEL Nanofabrication Facility between 2008 and 2017. He has authored or coauthored more than 275 journal and conference publications. He coauthored the book *Silicon Photonics Design* (Cambridge University Press, 2015). His research interests include silicon photonics, optoelectronics, high-speed laser (VCSEL) design, fabrication, and test, for applications in optical communication systems, biophotonics and quantum information. He was the Program Director with the NSERC CREATE Silicon Electronic-Photonic Integrated Circuits Research Training Program in Canada (2012–2018), and has been teaching numerous silicon photonics workshops and courses, since 2008. He was the recipient of the Killam Teaching Prize at the University of British Columbia, in 2014, the Natural Sciences and Engineering Research Council of Canada Discovery Accelerator Supplements Award, in 2015, for his research in silicon photonics integrated circuit design. He was an Elected Member of the IEEE Photonics 2014–2016 Society Board of Governors and was an Associate VP of Education.

Sudip Shekhar (Senior Member, IEEE) received the B.Tech. degree from the Indian Institute of Technology Kharagpur, Kharagpur, India, and the Ph.D. degree from the University of Washington, Seattle, WA, USA. From 2008 to 2013, he was with the Circuits Research Laboratory, Intel Corporation, Hillsboro, OR, USA, where he worked on high-speed I/O architectures. He is currently an Associate Professor of electrical and computer engineering with The University of British Columbia, Vancouver, BC, Canada. His research interests include circuits for electrical and optical interfaces, frequency synthesizers, and wireless transceivers. He was the recipient of the Schmidt Science Polymath Award and the UBC Killam Teaching Prize in 2022, the Young Alumni Achiever Award by IIT Kharagpur in 2019 and IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS DARTINGTON Best Paper Award in 2010 and a co-recipient of IEEE Radio-Frequency IC Symposium Student Paper Award in 2015. He is the Technical Program Committee of IEEE International Solid-State Circuits Conference (ISSCC) and the Optical Interconnects (OI) Conference, and as a Distinguished Lecturer for the IEEE Solid-State Circuits Society.