SiPhAI: A Reconfigurable Silicon Photonic Interposer Network for AI Acceleration

Mohammad Amin Mahdian, Ebadollah Taheri, Sudeep Pasricha and Mahdi Nikdast
Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523 USA

***Abstract*—We propose an energy-efficient and reconfigurable tree-based silicon photonic interposer network, adaptable to varying AI workloads. Our design is substantiated using an iPronics programmable photonics platform and demonstrates significant improvements in energy efficiency when compared to the state-of-the-art and with a different number of chiplets.**

I. Introduction and motivation

2.5D chiplet systems enhance traditional IC designs with better manufacturability, reusability, and customization. AI accelerators can benefit from this modular approach. By breaking down systems into smaller chiplets and reintegrating them through an interposer, 2.5D systems enable scalable, cost-effective AI accelerators. Multiply-accumulate processing element (MAC PE) chiplets provide essential computational power, while an SRAM-based global buffer ensures efficient data sharing, reducing costly off-chip communication. Conventional electronic interposers face performance issues due to high traffic and latency in ML tasks. Silicon photonic (SiPh) interposers offer a solution with higher bandwidth and lower latency through optical communication. SiPh interposers leverage hybrid integration techniques, where multiple chips from different material technologies are co-packaged together. This 2.5D integration enables alignment tolerances and addresses performance issues effectively [1], [2].

The novel contribution of this paper is on designing and evaluating SiPhAI, an energy-efficient silicon photonic (SiPh) interposer for AI accelerators. Our approach uses a reconfigurable SiPh interposer network based on interconnected Mach–Zehnder Switching elements (MZS) arranged in a tree architecture. This design addresses communication bottlenecks in traditional single-writer multiple-read interposer networks by eliminating accumulative losses associated with microring resonator (MRR) filters (see Fig. 1(a)). Additionally, SiPhAI enables signal broadcasting to multiple readers with an MZS arbitrary splitting ratio adaptable to varying AI workloads.

Our designed MZS devices include optimized Multi-Mode Interferometers (MMIs) with 0.17 dB experimentally demonstrated loss, to ensure a low-loss, reconfigurable connection for deep neural networks (DNNs), enabling efficient, low-latency communication among multiple chiplets. A proof of concept on an iPronics programmable photonic platform and simulations across six DNN models—DenseNet, ResNet, LeNet, VGG, MobileNet, and EfficientNet—show that SiPhAI enhances performance and energy efficiency in 2.5D AI accelerators by over 69% when compared to previous designs, paving the way for future advancements in chiplet-based acceleration architectures.

II. Proposed SiPhAI Design

In state-of-the-art 2.5D accelerators, there are two primary communication flows on the interposer: 1) MAC PEs that read inputs or layer weights from a global buffer (GLB), and 2) MAC PEs that write the layer's output to a GLB. This bidirectional communication pattern—*one-to-many and many-to-one*—motivates the development of a reconfigurable SiPh interposer network designed to support such interactions in AI hardware accelerators.

Previous chiplet system designs addressed scalability limitations by adopting a single-writer, multiple-reader (SWMR) communication paradigm [3]. This approach enhances scalability as the number of optical links at the sender (GLB) is independent of the number of receivers (MAC chiplets), as shown in Fig.1(a). However, as the number of wavelengths increases, new challenges arise. Data transmission from the GLB to the MAC chiplets involves shared waveguides. When an optical signal from the laser traverses multiple receivers, the cumulative through losses from the MRR filters at these receivers increase. Consequently, the laser must produce higher optical power to compensate, leading to increased energy consumption and reduced overall energy efficiency.

SiPhAI uses tree-based interconnections to minimize signal attenuation and enhance laser power efficiency. Tree-based sub-networks increase inter-chiplet bandwidth through parallel data communications, as shown in Fig.1(b). In this architecture, multiple GLB chiplets (writers) can communicate with several MAC chiplets (readers). Each of these chiplets is positioned along a bus waveguide at every port of the switch network. This strategic enhancement, compared to prior work (e.g., Trine [4]), increases SiPhAI flexibility, by allowing for multiple readers at the output of the GLB-to-MAC switch and multiple writers at the input of the MAC-to-GLB switch.

III. Implementation, Results, and Conclusion

To design the MZS in SiPhAI, we first designed and optimized an MMI coupler to reduce its optical loss and achieve an even 50-50 splitting ratio at the outputs. We introduced a figure of merit (FOM), defined as the power imbalance between the two MMI ports, and minimized it. Our FDTD simulations reveale that a single MMI coupler's worst-case transmission loss is 0.02 dB, with a power-splitting imbalance of less than 0.14 dB across the C-band. The designed MZS, based on this MMI coupler, achieves a worst-case loss of 0.12 dB in the Cross state and 0.5 dB in the Bar state using the electro-optical (E-O) tuning method. The MZS arms are 200 μm long, with a V_π of 1.3V and a switching time of 5.7 ns.

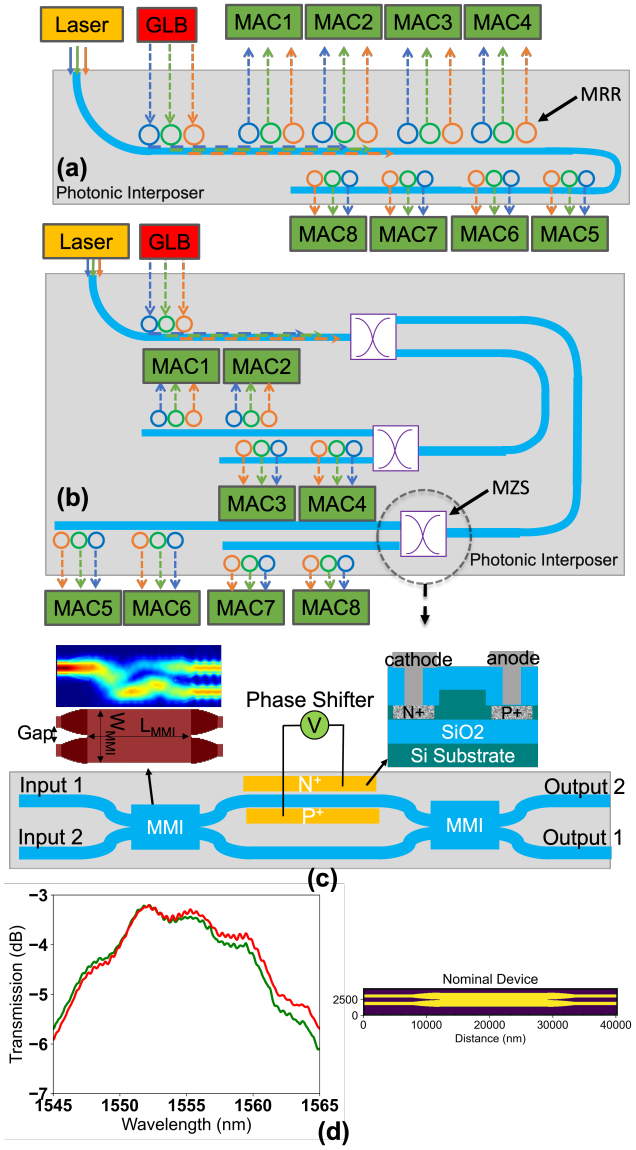

Fig. 1. (a) An SWMR design of an interposer network. (b) Proposed switch-enabled SiPh interposer: SiPhAI. (c) The building block of the designed MZS with optimized broadband MMI. (d) Response of each fabricated MMI output.

Our fabricated optimized MMI couplers exhibited a wide-band response over the C-band (1530–1565 nm), with a measured insertion loss of 0.17 dB (including loss due to an imperfect alignment during test), as shown in Fig. 1(d).

To emulate the performance of SiPhAI network, we used an iPronics programmable photonic platform, shown in Fig. 2(a). In our emulation, six MZS elements are used on each path to connect the input ports to the right-hand side outputs. Each MZS can be set to Bar, Cross, or a tunable state. Fig. 2(b) illustrates two different paths from input 1. On the green path, all MZSs are in the Bar state, directing the signal to output port 1. By setting one MZS to the 50-50 splitting mode in the red path, the signal is directed to both output ports 4 and 5, enabling signal broadcast. An optical switch controller manages the optical switches' state and path selection to optimize power consumption and minimize power loss [5]. This design adapts to different AI workloads, allowing dynamic reconfiguration

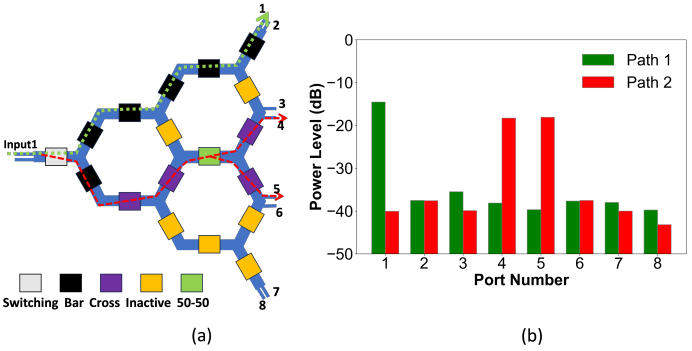

Fig. 2. (a) Emulated SiPhAI network with iPronics programmable platform. (b) The power level of each output port for paths 1 (green) and 2 (red) as shown in (a).

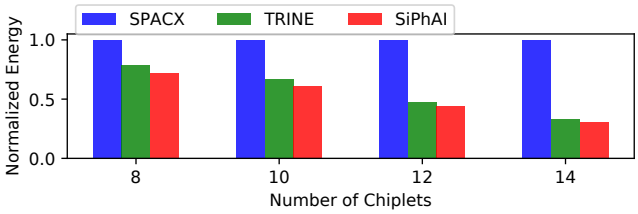

Fig. 3. Comparison of SiPhAI with prior related work.

to meet varying computational demands.

To showcase the scalability of our design, we compared its performance against two previous designs, SPACX [3] and TRINE [4]. SPACX is a bus-based communication architecture (Fig. 1(a)) and TRINE is an architecture similar to Fig. 1(b) but with only one reader at each output of the switch. Fig. 3 illustrates the normalized average energy consumption across six DNN models with varying numbers of chiplets. Our design consistently demonstrates significant reductions in energy consumption, surpassing both TRINE and SPACX. Even as the system scales up, these improvements remain notable, highlighting the scalability of our design. With 14 chiplets, our design achieves an impressive 69.5% improvement in energy efficiency compared to SPACX and 7.4% improvement in comparison with TRINE. Moreover, it reduces the number of switch stages by 16.7% compared to TRINE, which in turn improves the area overhead of the switch accordingly.

These substantial improvements highlight the promise of our design in maintaining energy efficiency as the system scales. The ability to offer low energy consumption with increasing number of chiplets makes our design particularly suitable for large-scale, high-performance computing applications, such as those required for advanced ML tasks. The scalability and energy efficiency not only reduce operational costs but also contribute to the development of more sustainable computing solutions, addressing the growing demand for energy-efficient technologies in the field of AI.

References

[1] S. Shekhar *et al.*, *Nature Communications*, p. 751, 2024.
[2] E. Taheri *et al.*, in *IEEE ICCAD*, 2022, pp. 1–9.
[3] Y. Li *et al.*, in *IEEE HPCA*, 2022, pp. 831–845.
[4] E. Taheri *et al.*, in *ACM NoCArc*, 2023, pp. 15–20.
[5] M. A. Mahdian *et al.*, in *IEEE PSC*, 2023, pp. 1–3.

This work was supported by the NSF under grants CNS-2046226 and CCF-2006788.