

# Beyond the Beachfront: Integration of Silicon Photonic I/Os under a High-Power ASIC

Subal Sahni, Abhijit Abhyankar, Ankur Aggarwal, Nikos Bamiedakis, Zoltan Bekker, Mohamed Benromdhane, Nadav Bergstein, Ties Bos, Christopher Davies, Andrew Gimlett, Xiaoping Han, Kelin Lee, Kavya Mahadevaiah, Hakki Ozguc, Kevin Park, Jeremy Plunkett, Sujit Ramachandra, Jason Redgrave, Ajmer Singh, Matteo Staffaroni, Angelina Totovic, Saurabh Vats, Phil Winterbottom, Darren Woodhouse, Waleed Younis, Shifeng Yu and David Lazovsky

*Celestial AI, 2962 Bunker Hill Ln, Suite 200, Santa Clara, CA 95054, USA*

*ssahni@celestial.ai*

**Abstract:** We present a photonics platform targeting optical connectivity at the point of compute in high-power ASICs. The platform uses bias-controlled electro-absorption modulators and is differentiated by broad temperature stability coupled with high bandwidth density. © 2024 The Authors

## 1. Motivation

The current trajectory of Artificial Intelligence (AI)-empowered services calls for future-proof hardware solutions able to support exponential scaling of model sizes and computational complexity [1]. The so-called memory-wall problem has been well documented [2], wherein the limited capacity and bandwidth of local memory packaged with the processor is significantly impacting AI compute performance. Additionally, fixed memory-to-compute ratios are failing to meet the demands of different workloads, leaving up to 30% of memory stranded and CPUs/GPUs frequently operating at only 30-50% capacity, implying a strong need for system disaggregation and memory pooling [3,4]. Optical networks can facilitate this disaggregation and the transition from pluggable transceivers to Co-Packaged Optics (CPO) and, lately, to 2.5D/3D integration of Photonic Integrated Circuits (PICs) and Application Specific ICs (ASICs) on a shared interposer has brought significant performance improvements to Silicon Photonics solutions [5-10]. Nonetheless, all these options are limited to the bandwidth accessible from the beachfront of the ASIC, putting a fundamental ceiling on the total data throughput that can be provided to and from the chip package. The next step in the evolution of these networks is to go beyond the beachfront and bring photonics to the point-of-compute in the ASIC, unlocking the full potential of optical interconnects as enablers of the next generation of AI models. The key differentiators of a photonics platform that can support such networks are high areal density and exceptional temperature stability in the face of 100s of Watts of dynamically switching power in the ASIC.

## 2. High-bandwidth, low-power dissipation, high density, temperature-resilient photonic platform

For the past several years, the modulation technology of choice for dense photonic-electronic co-integration have been Micro-Ring/Disk Resonators (MRRs/MDRs) owing to their compact footprint of no more than few 10s of  $\mu\text{m}$  in diameter and inherent compatibility with multichannel systems [5-9]. Their major drawback, however, is the ultra-narrow sub-1 nm optical bandwidth, making them very sensitive to environmental conditions, particularly thermal gradients created by high power ASICs operating in relatively close proximity. As shown in [6], this results in the need for extremely precise ring stabilization, since even local temperature variations of less than  $0.5^\circ\text{C}$  can adversely impact performance. If MRRs are kept a few mm away from the edge of the ASIC, then thermal stabilization is conceivable within a reasonable timeframe, to cope with temperature deltas created by varying workloads in the IC. But coming any closer to the edge of the ASIC, or residing underneath it, increases the slope of temperature change with time ( $\partial T/\partial t$ ) to an extent where MRR tracking can hardly be achieved without extraordinary control system complexity and penalty, if at all.

For dense temperature-gradient resilient Silicon photonic platforms, both in temporal and spatial terms, another modulation technology offers a significantly more stable, optically broadband alternative, that still retains the very high speed and extremely small footprint of MRRs. Germanium-Silicon (GeSi) Electro-Absorption Modulators (EAMs) have been known for their exceptional electro-optic bandwidth, supporting data-rates  $> 100$  Gbps, and 1-dB optical bandwidth beyond 30 nm [11, 12]. Combined with typical coefficient of peak wavelength shift with temperature of  $\partial\lambda/\partial T = 0.8$  nm/K, these devices can comfortably operate over a wide range of temperatures even without feedback control [13], making them an ideal candidate for point-of-compute integration with an ASIC.

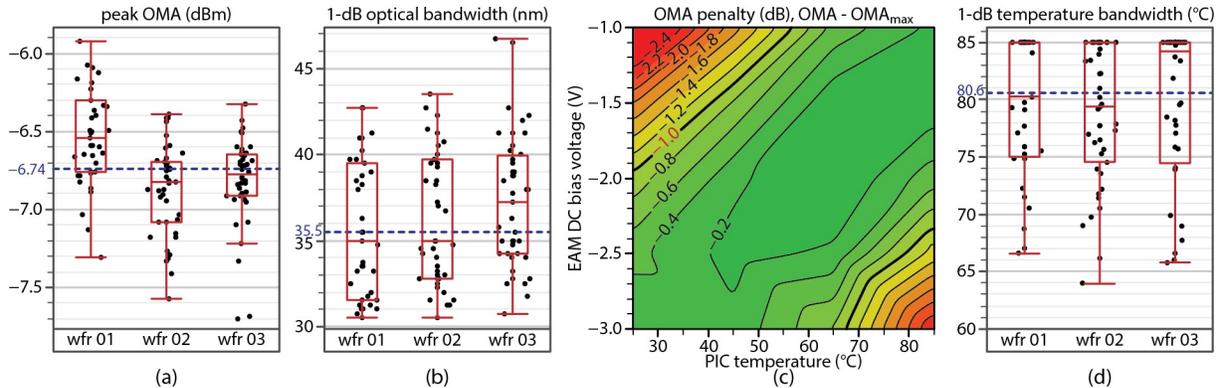


Fig. 1. (a-b) Variability charts of the (a) peak OMA in dBm and (b) 1-dB optical bandwidth in nm for EAMs over 3 sample wafers. (c) Heatmap of OMA penalty in dB for a sample EAM device, calculated as the difference between OMA at the operating point and its maximum value, due to changes in PIC temperature and EAM DC bias voltage. (d) Variability chart of the two-sided 1-dB temperature bandwidth given in °C for DC bias voltage range  $[-1, -2.5]$  V collected over 3 sample wafers.

Figure 1 shows characterization results over 3 sample wafers of a  $49.6 \mu\text{m}$  long EAM with GeSi composition targeting L-band operation at a high temperature corner of  $85^\circ\text{C}$ . Variability charts of the peak Optical Modulation Amplitude (OMA), Fig. 1(a), and 1-dB Optical Bandwidth (OBW), Fig. 1(b), reveal that under a DC bias voltage of  $-1$  V and a voltage swing of  $2 V_{pp}$ , the EAM achieves median peak OMA of  $-6.74$  dBm and an OBW consistently surpassing  $30$  nm, with a median value of  $35.5$  nm. As the temperature changes, the OMA spectral response shifts at a rate of  $\partial\lambda/\partial T = 0.77 \text{ nm}/^\circ\text{C}$  resulting in a modest reduction compared to its peak value. This yields a minimum 1-dB Temperature Bandwidth (TBW) of nearly  $40^\circ\text{C}$  without active control, which is comfortably large enough to accommodate instantaneous temperature swings created by workload changes in a co-located ASIC that may be burning up to a kW. Ideally this TBW should also extend to cover the entire range of expected temperatures beneath the ASIC, which can be as high as  $70^\circ\text{C}$  due to additional changes in the ambient. To compensate for the OMA loss at the targeted wavelength due to wider temperature shifts, the EAM DC bias voltage can be used as a control signal. A drop in EAM absorption due to a reduction in temperature can be compensated by an increase in reverse bias and vice versa.

The combined effect of temperature and DC bias voltage on OMA for one of the characterized devices is captured in the heatmap in Fig. 1(c), which shows the OMA loss with respect to its maximum value at a fixed operating wavelength. Here, the EAM is optimized for high temperature operation and so as the temperature drops the OMA loss is recovered by increasing the reverse DC voltage, revealing a very broad 1-dB operating range with maximum temperature stability. Fig. 1(d) shows the result of such temperature compensation on the EAM operating window measured across 3 sample wafers, with the reverse bias on each device varied between  $-1$  V and  $-2.5$  V depending on the temperature. We see a median TBW extended out to  $80.6^\circ\text{C}$ , with the bulk of the distribution capped at  $85^\circ\text{C}$  due to the temperature range limitation of the measurements themselves. This shows that the EAMs can readily function in the harsh thermal environments that will be encountered when photonics I/Os are brought right into the envelope of a very high-power ASIC.

Having resolved the critical challenge of dense temperature-stable modulation technology, we combine the GeSi EAM with a GeSi photodetector, forming the backbone of an E/O/E link that can be integrated below a high-power ASIC, fed by external light sources. This is the basis for the Photonic Fabric<sup>TM</sup>, which is our photonic interconnectivity platform. Optically broadband EAMs allow multiple wavelength channels to be used, increasing the link capacity, whereas a low-loss custom Process Design Kit (PDK) combined with high-power lasers enables feeding of multiple links with a single light source. The high-speed transmit and receive electronics that interface with the photonics through fine-pitch 2.5D/3D integration are designed using state of the art CMOS nodes, significantly driving down power consumption and allowing straightforward IP and packaging integration with the target ASIC.

### 3. Test-vehicle for integration with advanced-node ASIC

We have fabricated and assembled an integration test-vehicle, which is a single-tile functional demonstrator of Photonic Fabric. The ASIC here (Fig. 2(a)) is a 4nm CMOS chip manufactured at TSMC<sup>®</sup> that is designed to support 16 bi-directional data lanes, each operating at 56Gbps NRZ. The IC contains all the high-speed and control blocks needed to support E-O-E links in conjunction with the photonics platform described above, including 64:1

serialization/de-serialization, clock and data recovery units, modulator drivers and transimpedance amplifiers. To fully utilize the high density of the photonics, the ASIC design is also optimized for area, with the entire 16-lane test chip occupying only  $\sim 1\text{mm} \times 1\text{mm}$ . In this test-vehicle the ASIC is flip-chipped onto a co-designed PIC using Copper-pillar micro-bumps at  $40\ \mu\text{m}$  minimum pitch and electrical connections to the substrate are achieved with wire-bonds. Finally, a 24-ch Fiber Array Unit (FAU) is attached to the top of the PIC to provide optical connections through grating couplers. Fig. 2(b) shows the final assembly mounted on a test-board and Fig. 2(c) shows a representative Tx eye-diagram for a 56 Gbps NRZ signal, recorded by a high-speed sampling oscilloscope.

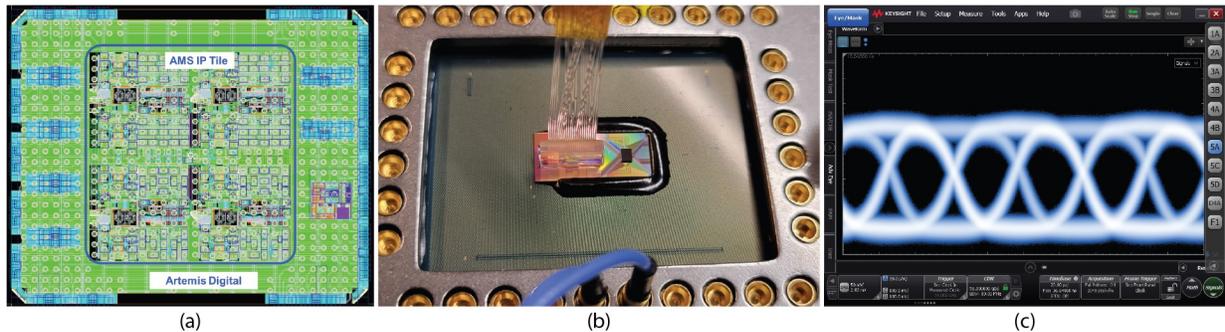


Fig. 2. (a) Layout of the 4nm CMOS test-chip enabling full E-O-E Silicon photonic links at 56Gbps. The analog/mixed signal blocks are in the middle, with digital control and interface circuitry on the outside. (b) Test-chip flip-chipped onto a matched PIC containing 16 complete optical lanes, with an FAU bonded to a grating-coupler array to provide optical access. (c) Representative 56Gbps NRZ eye with a  $49.6\ \mu\text{m}$  long EAM and 1.8V peak-peak swing from the modulator driver.

#### 4. Conclusion

We have presented a Silicon Photonics platform that is specifically targeted at bringing the extremely high bandwidth and low latency of optical connectivity directly to the point of compute in the high-power XPUs that are at the heart of the innovation currently happening in AI, thereby breaking through the beachfront limitation on package bandwidth. This level of integration is possible only if the photonics can withstand the extraordinarily dynamic thermal conditions that come with such integration. Additionally, the entire photonics macro, including the driving electronics, should be implementable with very high areal density so that any region of the target ASIC can freely access the optical connectivity without over-allocating extremely valuable chip area to the I/O. In Silicon Photonics, the modulator technology largely dictates compliance with these requirements, and we believe that our platform based on EAMs is uniquely suited to deliver this next step in optical integration. We have demonstrated very compact modulators that can operate over temperature ranges of more than  $80^\circ\text{C}$  and have built a full E-O-E link prototype vehicle based on this technology, which is designed to deliver aggregate bandwidths of 1.8Tbps in little over a square mm of Silicon.

#### 5. References

- [1] C.-J. Wu *et al.*, “Sustainable AI: Environmental implications, challenges and opportunities,” in *Proceedings of Machine Learning and Systems*, vol. 4, D. Marculescu, Y. Chi, and C. Wu, Eds. (2022), pp. 795–813.
- [2] A. Gholami, “AI and Memory Wall,” <https://medium.com/riselab/ai-and-memory-wall-2cb4265cb0b8> (2021). Accessed: Oct 16<sup>th</sup>, 2023.
- [3] D. S. Berger *et al.*, “Design tradeoffs in CXL-based memory pools for public cloud platforms,” *IEEE Micro* 43, 30–38 (2023).
- [4] J. Li, G. Michelogiannakis, B. Cook, D. Cooray, and Y. Chen, “Analyzing resource utilization in an HPC system: A case study of NERSC’s Perlmutter,” in *High Performance Computing*, A. Bhatele, J. Hammond, M. Baboulin, and C. Kruse, Eds. (Springer Nature Switzerland, 2023), pp. 297–316.
- [5] F. Sunny, E. Taheri, M. Nikdast, and S. Pasricha, “Machine learning accelerators in 2.5D chiplet platforms with silicon photonics,” in *Design, Automation & Test in Europe Conference & Exhibition (DATE)*, (2023)
- [6] B. G. Lee, N. Nedovic, T. H. Greer, and C. T. Gray, “Beyond CPO: A motivation and approach for bringing optics onto the silicon interposer,” *J. Light. Technol.* 41, 1152–1162 (2023).
- [7] S. Daudlin *et al.*, “3D photonics for ultra-low energy, high bandwidth-density chip data links,” arXiv preprint arXiv:2310.01615 (2023).
- [8] Lightelligence, “Hummingbird™ processor,” <https://www.lightelligence.ai> (2023). Accessed: Oct 16<sup>th</sup>, 2023.
- [9] J. Howard, “The First Direct Mesh-to-Mesh Photonic Fabric,” in *IEEE Hot Chips 35 Symposium (HCS)*, Palo Alto, CA, USA, (2023)
- [10] N. Harris, “Introducing enwise, idiom and passage — next generation AI compute, compile and interconnect platforms,” <https://medium.com/lightmatter/> (2023). Accessed: Oct 16<sup>th</sup>, 2023.
- [11] A. Rahim, A. Hermans, B. Wohlfeil, D. Petousi, B. Kuyken, D. V. Thourhout, and R. G. Baets, “Taking silicon photonics modulators to a higher performance level: state-of-the-art and a review of new technologies,” *Adv. Photonics* 3, 024003 (2021).
- [12] J. Verbist *et al.*, “100 Gb/s DAC-less and DSP-free transmitters using GeSi EAMs for short-reach optical interconnects,” in *Optical Fiber Communications Conference and Exposition (OFC)*, (2018)
- [13] D. Coenen *et al.*, “Electro-absorption modulator thermo-optical self-heating analysis,” *J. Light. Technol.* 41, 6000–6006 (2023).

Trademarks: Celestial AI and Photonic Fabric (patents pending) are trademarks of Celestial AI, TSMC is a registered trademark of Taiwan Semiconductor Manufacturing Company, Ltd.