Optical Computing with Co-Packaged Electronics for Space-Based Environments

Adam H. Slater, Steven M. Bowers
University of Virginia

***Abstract*—Size, weight, and power constraints are significant limiting factors to space-based computing performance. Optical computing with wavelength-division multiplexing is a promising approach for achieving low-energy computation in a small low-weight package. Co-packaging electronics and photonics is a critical pathway to this approach. Additionally, this new paradigm of computing will require engineering advances in subsystem circuits, namely the electro-optic modulator driver. This work details the driver design and co-packaged integration techniques, demonstrating state-of-the-art efficiencies for 4-level pulse-amplitude modulation optical drivers with 56fJ/bit at 2Gb/s.**

I. Introduction

VARIOUS space missions have begun to explore the capabilities of artificial intelligence in space [1]. Space-tested machine learning hardware includes only a handful of processors including the Google Edge Tensor Processing Unit (TPU), the Qualcomm Snapdragon 885, and the NVIDIA H100 [2] [3]. Recent advances in space computing represent large leaps forward with the launch of the Three-Body Computing Constellation and the Starcloud [4]. However, neither of those projects are adopting emerging technologies that the research community is discovering. Computing speed and energy usage for various processors are shown in Fig. 1 and Fig. 2 indicating that optical computing has the potential to provide comparable speed performance with a much lower energy usage [5]. However, speed and energy usage metrics are not yet standardized across the optical computing research community; for example, the optical outlier in 2026 only includes optical power [6]. This work evaluates how much electrical energy is needed to support generalized optical computing systems, identifies the electrical transmitter as a critical pathway and presents potential improvements to electro-optic transmitters that arise from the unique requirements of optical computing.

Optical computing offers two fundamental advantages: (1) optical mediums are operational over a bandwidth thousands of times greater than electrical mediums and (2) optical processing delay is limited only by propagation velocity, not by the RC delay and transit times of conventional logic gates. For this work we aim to enable rapid adoption. This leads to a target system with digital inputs and outputs, which requires electrical data converters, eliminating the gains of optical processing delay. Thus, we focus on systems for ultra-wideband single-shot computing of matrix multiplications, a fundamental operation in modern AI neural networks [11].

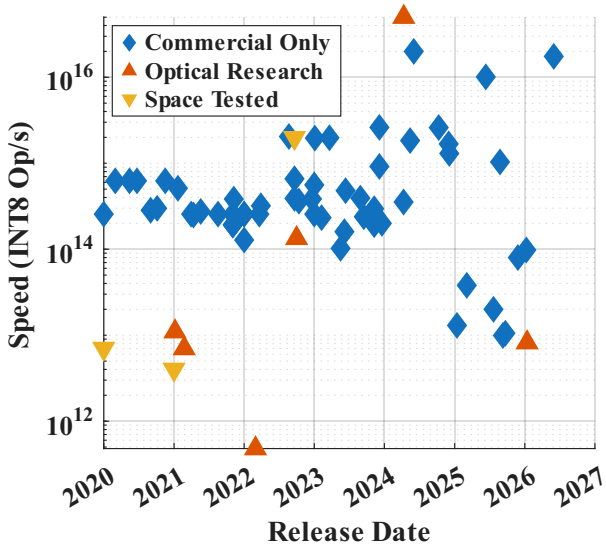

Fig. 1: Speed of commercial, space-tested, and optical research processors. [5] [6] [7] [8] [9] [10]

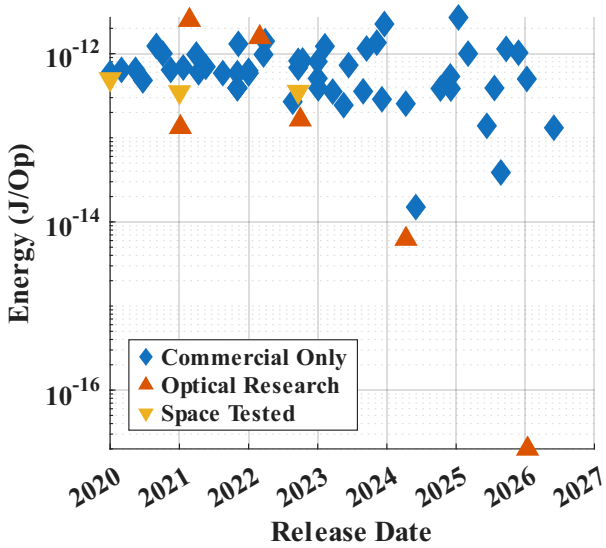

Fig. 2: Energy usage of commercial, space-tested, and optical research processors. [5] [6] [7] [8] [9] [10]

This work builds towards achieving high-speed, low-energy computation by (1) maximizing compute performance through photonic computing architecture design, (2) reducing power consumption by application-specific design of interface electronics, and (3) increasing the computational performance

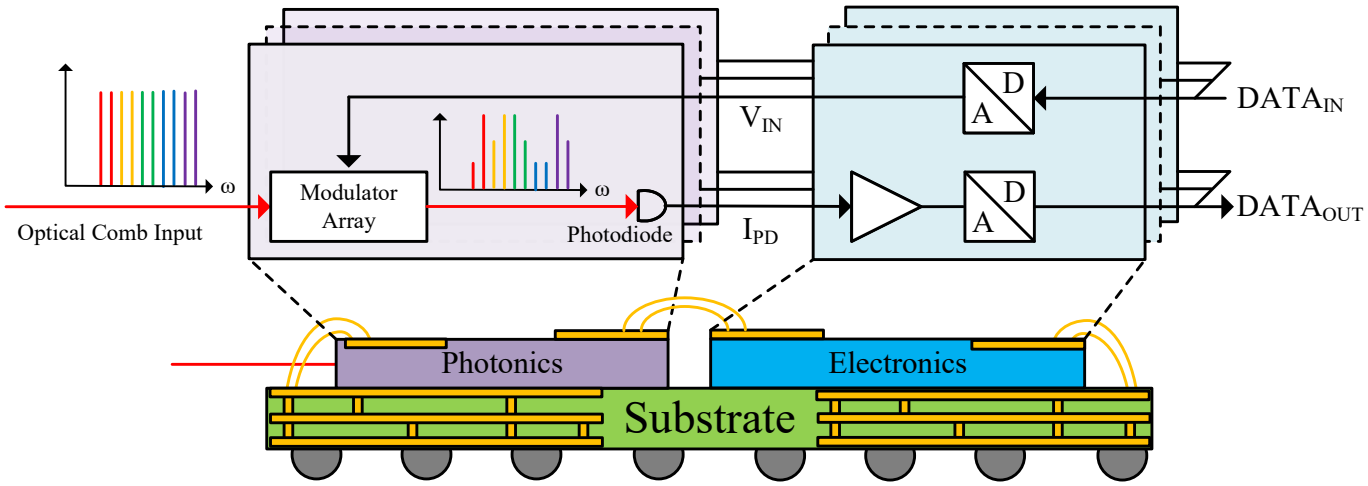

Fig. 3: Photonic tensor accelerator design concept illustrating co-packaged integration of electronics and photonics.

per unit of volume through photonic-electronic co-packaged integration.

To decrease size, weight, power, and cost while improving performance we will be leveraging co-packaging techniques of photonic and electronic integrated circuits. The hardware system concept for this work is shown in Fig. 3. This includes an electronic chip and a photonic chip assembled within the same package using chip-to-chip wirebonding. The photonic chip is the compute core and the electronic chip provides an interface to the larger system.

II. Photonic MAC using WDM

A generalized architecture of the photonic multiply and accumulate (MAC) processor using wavelength division multiplexing (WDM) is presented here to establish the potential speed and energy efficiency with the goal of integrating it into a general-purpose machine learning application that can be deployed in size constrained and power constrained environments.

The primary goal of this subsystem is to maximize the number of matrix multiplication operations per second and minimize the energy consumed. An operation is defined here as any multiplication or addition. For example, a single vector-vector-multiplication (VVM) of two three-element vectors would contain five total operations (three multiplications and two additions), as follows for a weight matrix W and and an input matrix U:

$$W \cdot U = \begin{bmatrix} w_1 & w_2 & w_3 \end{bmatrix} \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} = w_1 \cdot u_1 + w_2 \cdot u_2 + w_3 \cdot u_3 \quad (1)$$

Expanding this to matrix-vector multiplication (MVM) results in a total number of operations of $(N + (N-1)) \cdot M$ where $W \in \mathbb{R}^{M \times N}$, $U \in \mathbb{R}^{N \times 1}$. If $M = N$, the number of operations increases approximately as $2N^2$. To achieve a speed of 10 peta-operations per second (POPS) with a fourth generation high-bandwidth memory interface rate of 8GHz requires N = 800.

In the proposed system, N input data elements are encoded onto N unique optical wavelengths. The encoded signal is split M times and routed through the weight matrix multiplier then to a photodiode such that there is one photodiode for each element of the output vector as shown in Fig. 4. In this example the encoding is performed with microring modulators (MRMs), which only resonate at the wavelength λ_n.

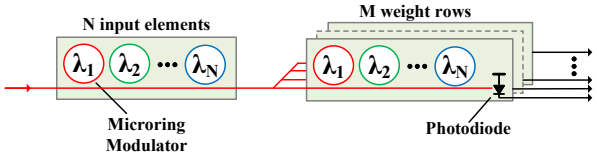

Fig. 4: Generalized photonic circuit for performing matrix-vector multiplication. Each input element is encoded using microring modulators onto a unique optical wavelength. These optical signals are split and multiplied by the corresponding rows of the weight matrix. Photodiodes convert the optical signal to the electrical domain, and in the process perform all of the addition operations at once.

High-speed digital-to-analog converters (DACs) are required for each element of the input encoder which updates with every cycle. Compute hardware is typically limited compared to the size of the neural network layers and matrix dimension, meaning high-speed DACs are also needed for each element of the weight matrix. Similarly, a high-speed analog-to-digital converter (ADC) is required for each output element. This results in $N(1+M)$ DACs and M ADCs.

A. Practical Photonics Challenges

Perhaps the greatest challenge in space-based photonics is high-energy radiation. The material structures that make photonic circuits are affected by radiation. High-energy radiation causes imperfections in the device material which causes attenuation and eventually loss of function. Surrounding these circuits in centimeters of lead works, but is not ideal from a weight perspective. Other mitigation techniques have been proposed: predictive counter-design, post-radiation annealing, higher doping concentrations, and thicker features [12].

The proposed system is highly generalized and does not consider many of the practical considerations of photonics design such as off-resonant loss, high frequency splitting ratios, and resonance synchronization. These practical challenges have architectural solutions that could be implemented today. However, the best performance will be achieved with component advancements, which are expected in the near future as photonic devices advance [13]. In this work, we will continue to consider the photonic hardware accelerators at the system level with an emphasis on the electronics required.

III. Electronic Interface

The electronics interface contains many input drivers and output photoreceivers. As the optical communications field begins to adopt MRMs as the electro-optic modulator (EOM), the driver is conventionally designed, not as a DAC, but as a 1-bit digital buffer with an extended voltage swing. The output photoreceiver generally contains a transimpedance amplifier (TIA) followed by an ADC. If we consider the power consumed by each circuit using state-of-the-art values, as shown in Table I, we see that the receiver consumes a significant amount of energy. However, there are many more drivers. Thus, to optimize the system for efficiency requires optimizing for driver efficiency. Additionally, the ADC has been exhaustively researched for decades while drivers for MRMs are a new area of research with many opportunities for advancement.

TABLE I: Electronics Comparison for MVM

Component	Example Circuit	Energy	Number Used
EOM Driver	Nature '25 [14]	50 fJ/bit	$N(1+M)$
Photoreceiver TIA	JSSC '22 [15]	87.5 fJ/bit	M
Photoreceiver ADC	JSSC '22 [16]	4.3 pJ/bit	M

A. EOM Driver Circuit

To reduce power and minimize area, the circuit architecture is designed to eliminate the need for power hungry multiplexing circuits and large impedance matching networks. Operating at gigahertz frequencies eliminates the need for parallel-to-serial time division multiplexing. Decreasing the distance to the load to 1/10 of a wavelength or less effectively eliminates reflections, removing the need for impedance matching circuits.

Most low-energy optical transmitters are 1-bit with not-return-to-zero (NRZ) encoding. As the number of amplitude levels (N) increases, the work done to transition between voltage levels is divided. A pulse amplitude modulation (PAM) order of four is chosen for this design to balance the tradeoff between higher EOM electrical efficiency and the efficiency reduction in the driver due to the added circuitry. The proposed circuit is the supply-multiplexing PAM-4 DAC shown in Fig. 5.

This circuit switches the output between ground and one of three power supplies: V01, V10, and VDD. A logic circuit, implemented with inverters and transmission gates, determines which power supply is connected to the output given the state of the inputs (LSB and MSB). All devices are intended

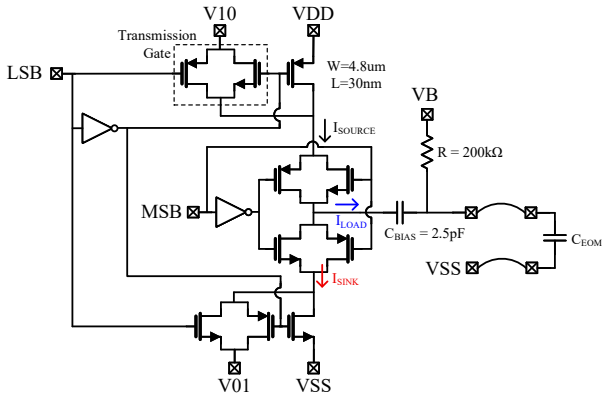

Fig. 5: Supply-multiplexing PAM-4 DAC for driving high-impedance photonic modulators.

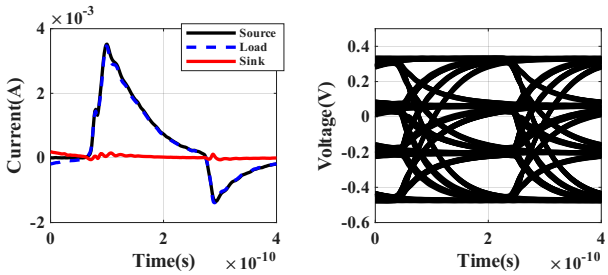

Fig. 6: Simulation results showing minimal shoot-through current and open eyes at 10Gb/s.

to operate in the large signal regime. To ensure the devices enter triode, transmission gates were implemented. Otherwise transitions to or from levels '01' and '10' exhibit large RC time constants and limit operational speed.

Simulated performance reached 10Gb/s consuming 20fJ/bit using a 250fF EOM, 10fF pad capacitance and 100pH pillar inductance achieving the eye diagram shown in Fig. 6. Low-power design techniques included a transistor size optimization to minimize shoot through current also shown in Fig. 6.

This PAM-4 driver circuit was designed into a 1x4 array and fabricated in 28nm CMOS as shown in the die photo, Fig. 7. Each driver has an active area of 0.0012mm^2 including the 2.5pF bias tee capacitor. Energy efficiency and open eye diagrams were measured over single channel data rates of 500Mb/s to 3Gb/s and a supply voltage range of 0.75V to 0.9V. Results are summarized in Table II.

To test this driver standalone, a test structure was designed in which the output signal is observed using an off-chip high-impedance buffer to allow the use of 50Ω measurement

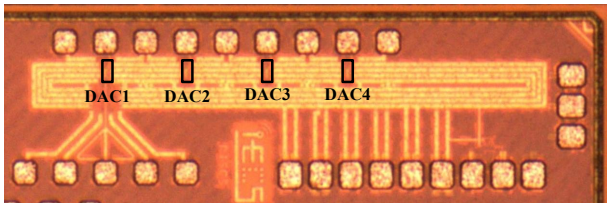

Fig. 7: 28nm CMOS die photo showing 1x4 array of supply-multiplexing PAM-4 DAC test structures.

TABLE II: Transmitter Comparison

Work	**This Work**	JSSC '25 [17]	Nature '25 [14]	RFIC '25 [18]
DAC Topology	**Supply Mux**	4xVDD Driver	2xVDD Driver	Distributed
Technology	**28nm CMOS**	45nm CLO	28nm CMOS	180nm SiGe
Modulation	**PAM-4**	NRZ	NRZ	PAM-4
Speed/Ch	**3Gb/s (10Gb/s)**†	32Gb/s	10Gb/s	224Gb/s
TX Area	**0.0012mm^2**	0.002mm^2 *	0.0018mm^2	2mm^2
Speed Density	**2.5Tb/s/mm^2**	13.25Tb/s/mm^2	5.3Tb/s/mm^2	0.112Tb/s/mm^2
Energy	**56fJ/bit (34fJ/bit)****	328fJ/bit	50fJ/bit	3620fJ/bit
Load	**800fF**	11fF	198fF	50Ω

† Simulated performance using an EOM load capacitance of 250fF
* Area includes MRM
** When normalized to a 200fF load.

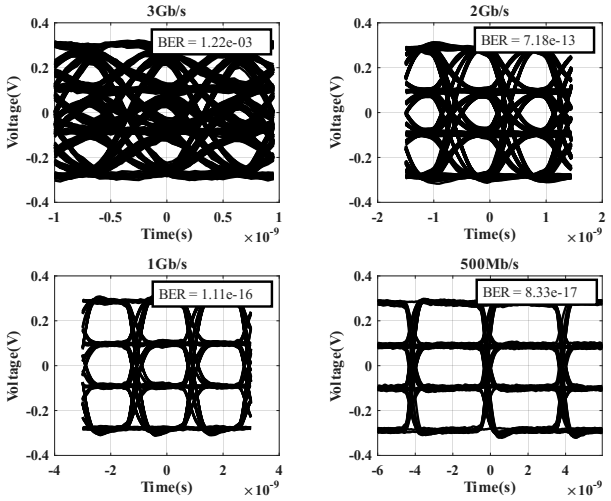

Fig. 8: Measured eye diagrams for data rates ranging from 500Mb/s to 3Gb/s.

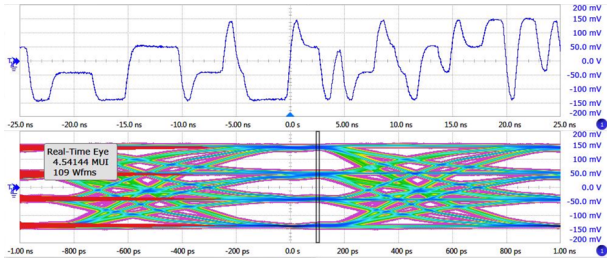

Fig. 9: Measured 2Gb/s eye diagram illustrating the sampling window and histogram data (in red) used for statistical analysis of BER.

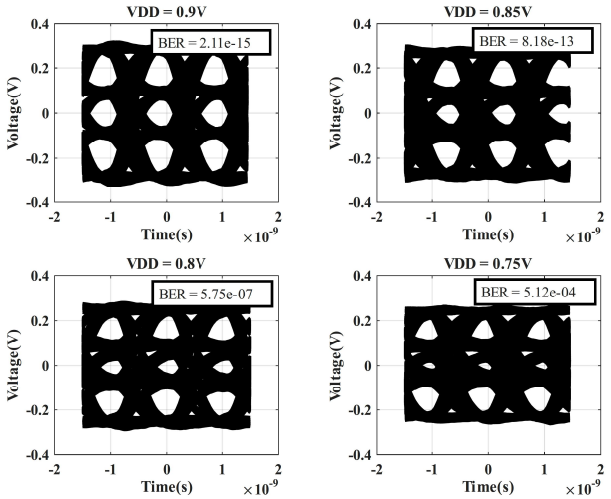

Fig. 10: Measured eye diagrams at 2Gb/s as supply voltage is varied from VDD = 0.75V to VDD = 0.9V.

equipment. One channel is driven from two off-chip sources and the other three channels are driven by six separate on-chip PRBS15 generators.

Measured eye diagrams are shown in Fig. 8. This device works for supply voltages down to VDD = 0.75V, indicating a potential use with future MRMs which are designed to reduce the RF voltage swing for lower dynamic power consumption. As supply voltage is varied, V01 and V10 are always 1/3 and 2/3 of VDD, respectively. At 2Gb/s, the minimum energy consumption over supply voltage is measured at 56fJ/bit. Eye diagrams for various supply voltages are shown in Fig. 10 demonstrating open eyes down to 0.75V. The high-impedance observation buffer (OPA859) used in measurement has a large-signal cutoff frequency around 400MHz and a single-ended input capacitance of 800fF limiting the speed and efficiency. This results in an attenuated output signal around 600mV$_{pp}$.

For a fair comparison of driver energy consumption across various load capacitances, the theoretical energy required to charge the capacitor is evaluated. The average energy added to the modulator capacitor when transmitting random PAM data is derived here as eq.2, where C_{EOM} is the capacitance of the EOM and N is the number of amplitude levels.

$$\overline{E_{added}} = \frac{1}{2}C_{EOM} \cdot \frac{1}{N^2} \sum_{i=1}^{N-1} \sum_{j=0}^{i-1} \left[\left(\frac{v_{pp} \cdot i}{N-1} \right)^2 - \left(\frac{v_{pp} \cdot j}{N-1} \right)^2 \right]$$

$$\overline{E_{added}} = \frac{1}{2}C_{EOM} \cdot v_{pp}^2 \cdot \frac{N+1}{6N} \tag{2}$$

Approximated bit error rate (BER) is determined using a statistical noise analysis assuming a gaussian distribution. Fig. 9 shows a 2Gb/s eye diagram capture and the histogram for an example sampling window used for BER approximation. The sampling window is chosen at the peak of the eye opening to remove receiver clock jitter from the evaluation of the driver. The measured data within each amplitude section is fit to a gaussian distribution and the probability of an error is determined as the likelihood of a sample crossing the decision threshold placed at the mean between peaks.

B. *Co-packaging*

This driver was co-packaged with a tunable microring resonator using chip-to-chip wirebonding as shown in Fig. 11. The microring array transmission coefficient was measured across wavelength and control voltage. Initial results show a thermal resonance tuning range of >1.25nm as shown in the Fig. 11 inset. Due to a low microring bandwidth of 10Hz, no data throughput tests were performed.

IV. Conclusion

An optical computing architecture for accelerating matrix multiplication has been presented enabling high-speed, low-energy, space-based artificial intelligence processors. A digital-in digital-out optical computing paradigm is described which allows for rapid adoption of analog WDM optical computing.

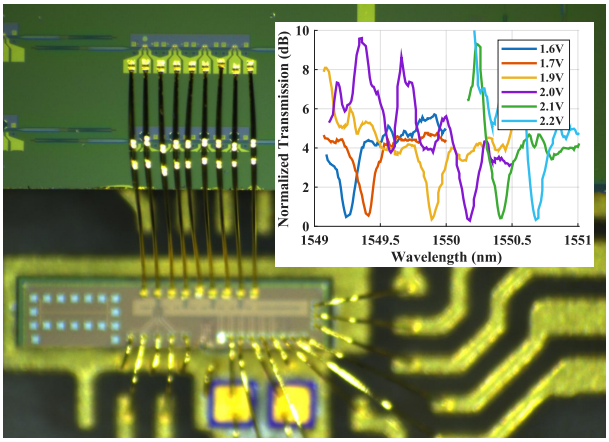

Fig. 11: Chip-to-chip wirebonding integration of electrical driver and a 1x4 array of microring resonators. (Inset) Microring resonator tuning range.

The electronic interface is identified as a critical path towards low-energy operation. Thus, a low-energy PAM-4 EOM driver circuit is presented, demonstrating the most efficient PAM-4 optical driver to date at 56fJ/bit. Co-packaged integration is demonstrated using chip-to-chip wirebonding.

ACKNOWLEDGMENTS

The authors would like to thank the Virginia Space Grant Consortium for financially supporting this research. This material is based on research sponsored, in part, by the DARPA NaPSAC program under contract N660012424000.

REFERENCES

[1] J. Goodwill, C. Wilson, and J. MacKinnon, "Chapter 16 - current ai technology in space," in *Precision Medicine for Long and Safe Permanence of Humans in Space*, C. Krittanawong, Ed. Academic Press, 2025, pp. 239–250. [Online]. Available: https://www.sciencedirect.com/science/article/pii/B9780443222597000060

[2] J. Goodwill, G. Crum, J. Mackinnon, C. Brewer, M. Monaghan, T. Wise, and C. Wilson, "NASA SpaceCube Edge TPU SmallSat Card for Autonomous Operations and Onboard Science-Data Analysis," *Proceedings of the Small Satellite Conference*, no. SSC21-VII-08, Aug 2021.

[3] J. Swope, F. Mirza, E. Dunkel, Z. Towfic, S. Chien, D. Russell, J. Sauvageau, D. Sheldon, M. Fernandez, and C. Knox, "Benchmarking remote sensing image processing and analysis on the snapdragon processor onboard the international space station," in *IGARSS 2022 - 2022 IEEE International Geoscience and Remote Sensing Symposium*, 2022, pp. 5305–5308.

[4] T. Kacprzak, "AI on Modern HPC Platforms In Space," in *Proceedings of SIAPS: Swiss Innovation for Astrophysics and Planetary Science*. Bern, Switzerland: SIAPS, Jan 2026, 12-minute contributed talk, Session: Big-data handling, embedded software.

[5] Z. Xu, T. Zhou, M. Ma, C. Deng, Q. Dai, and L. Fang, "Large-scale photonic chiplet Taichi empowers 160-TOPS/W artificial general intelligence," *Science*, vol. 384, no. 6692, pp. 202–209, 2024. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.adl1203

[6] C. Luan, R. Davis III, Z. Chen, D. Englund, and R. Hamerly, "Single-shot matrix-matrix photonic processor based on spatial-spectral hyper-multiplexed parallel diffraction," *Nature Communications*, vol. 17, p. 484, 2026.

[7] G. Mourgias-Alexandris, M. Moralis-Pegios, A. Tsakyridis, N. Passalis, M. Kirtas, A. Tefas, T. Rutirawut, F. Y. Gardes, and N. Pleros, "Channel response-aware photonic neural network accelerators for high-speed inference through bandwidth-limited optics," *Opt. Express*, vol. 30, no. 7, pp. 10 664–10 671, Mar 2022. [Online]. Available: https://opg.optica.org/oe/abstract.cfm?URI=oe-30-7-10664

[8] X. Xu, M. Tan, B. Corcoran, J. Wu, A. Boes, T. G. Nguyen, S. T. Chu, B. E. Little, D. G. Hicks, R. Morandotti, A. Mitchell, and D. J. Moss, "11 TOPS photonic convolutional accelerator for optical neural networks," *Nature*, vol. 589, no. 7840, pp. 44–51, 2021.

[9] J. Feldmann, N. Youngblood, M. Karpov, H. Gehring, X. Li, M. Stappers, M. L. Gallo, X. Fu, A. Lukashchuk, A. S. Raja, J. Liu, C. D. Wright, A. Sebastian, T. J. Kippenberg, W. H. P. Pernice, and H. Bhaskaran, "Parallel convolutional processing using an integrated photonic tensor core," *Nature*, vol. 589, no. 7840, pp. 52–58, 2021.

[10] L. De Marinis, A. Catania, P. Castoldi, G. Contestabile, P. Bruschi, M. Piotto, and N. Andriolli, "A codesigned integrated photonic electronic neuron," *IEEE Journal of Quantum Electronics*, vol. 58, no. 5, pp. 1–10, 2022.

[11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," 2023. [Online]. Available: https://arxiv.org/abs/1706.03762

[12] Q. Du, "High energy radiation damage on silicon photonic devices: a review," *Opt. Mater. Express*, vol. 13, pp. 403–412, 2023.

[13] S. Shekhar, W. Bogaerts, L. Chrostowski *et al.*, "Roadmapping the next generation of silicon photonics," *Nature Communications*, vol. 15, no. 751, 2024.

[14] S. Daudlin, A. Rizzo, S. Lee *et al.*, "Three-dimensional photonic integration for ultra-low-energy, high-bandwidth interchip data links," *Nature Photonics*, vol. 19, pp. 502–509, 2025. [Online]. Available: https://doi.org/10.1038/s41566-025-01633-0

[15] S. Daneshgar, H. Li, T. Kim, and G. Balamurugan, "A 128 Gb/s, 11.2 mW Single-Ended PAM4 Linear TIA With 2.7 μArms Input Noise in 22 nm FinFET CMOS," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 5, pp. 1397–1408, 2022.

[16] J. Shen, W.-H. Chen, E. Burlinqame, S. Weinreich, M. Elliott, S. McCracken, J. Kenney, J. Brunsilius, A. Korkmaz, E. A. Fontecilla, N. Rakuljic, U. Mehta, B. Sullivan, J. Scuteri, B. B. Luu, M. Nichols, D. Martin, R. Sullivan, D. DeBolt, and R. Kapusta, "24.8 A 12GS/s 9b 16× Time-Interleaved SAR ADC in 16nm FinFET," in *2025 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 68, 2025, pp. 442–444.

[17] K. Omirzakhov, H. Hao, A. Pirmoradi, and F. Aflatouni, "Energy Efficient Monolithically Integrated 256 Gb/s Optical Transmitter With Autonomous Wavelength Stabilization in 45 nm CMOS SOI," *IEEE Journal of Solid-State Circuits*, vol. 60, no. 7, pp. 2522–2531, 2025.

[18] H. Liu, R. Deng, Z. Dong, G. Li, J. Liu, N. Wu, W. F. Cops, T. Chen, L. Liu, and N. Qi, "A 224-Gb/s PAM-4 Linear Distributed Driver for Silicon-Photonic Modulators in SiGe BiCMOS," in *2025 IEEE Radio Frequency Integrated Circuits Symposium (RFIC)*, 2025, pp. 387–390.