

Hybrid CMOS-Photonic VLSI Interconnects for Ultra-High Bandwidth and Energy-Proportional Chiplet Communication in 2.5D/3D Systems

H. Fratlin^{1*}, K.P. Sikalu²

¹Department of Electrical and Computer Engineering, Ben-Gurion University, Beer Sheva, Israel

²Electrical and Electronic Engineering Department, University of Ibadan, Nigeria

KEYWORDS:

Hybrid CMOS-Photonic Interconnects, 2.5D/3D Integration, Chiplet Communication, Silicon Photonics, Energy-Proportional Design, High-Bandwidth On-Chip Communication, Thermal-Aware Packaging, VLSI System Architecture

ARTICLE HISTORY:

Submitted : 09.05.2025

Revised : 24.06.2025

Accepted : 15.08.2025

<https://doi.org/10.31838/JIVCT/02.03.10>

ABSTRACT

The emerging intensive nature of data transfer in artificial intelligence (AI), cloud computing, and high-performance computing (HPC) increased the need to shift towards chiplet based systems. Nevertheless, conventional electrical connecting devices are fast being ill equipped to cater the bandwidth density, energy efficiency and thermal management factors in 2.5D/3D integrated system. With the introduction of CMOS-photonic VLSI systems-on-chip, ultra-high bandwidths and energy-proportional communication within and across chiplets are feasible in advanced packaging environments. This paper determines and proposes a hybrid CMOS-photonic VLSI interconnect architecture, which can achieve those ultra-high bandwidths and energy-proportional communication. To attain this goal, we suggest that a CMOS-compatible silicon photonic architecture be co-designed with electrical SerDes and interface circuitry so as to incorporate CMOS-compatible silicon photonic components, e.g., MachZehnder modulators (MZMs), ring modulators, and photodetectors. The energy management framework of workload-aware is also integrated to adjust power consumption of optical link dynamically according to the traffic demand. The proposed system is simulated as a multi domain system that involves CST Microwave Studio to design the electromagnetic part, COMSOL Multiphysics to model the thermal design and Synopsys OptSim to model the photonic link in the system. Simulation results show up to 10x bandwidth advantage and more than 60 percent energy-per-bit savings of a hybrid architecture over conventional copper based interconnects. Moreover, due to optimized layout design using TSV, the thermal simulations indicate that the proper cooling of the device can be achieved under its peak workloads. In summary, the hybrid CMOS-photonic interconnection proposed is an energy efficient, scalable and thermally feasible communication systems used in chiplets in 2.5D/3D chip fabrication providing solution to the bandwidth and energy bottlenecks of layer-to-layer communication in heterogeneous computing chip systems.

Author e-mail: maidanov.k@gmail.com, kp.sikalu@ui.edu.ng

How to cite this article: Fratlin H, Sikalu K P. Hybrid CMOS-Photonic VLSI Interconnects for Ultra-High Bandwidth and Energy-Proportional Chiplet Communication in 2.5D/3D Systems. Journal of Integrated VLSI, Embedded and Computing Technologies, Vol. 2, No. 3, 2025 (pp. 79-88).

INTRODUCTION

The scale of data-intensive industries like artificial intelligence (AI), machine learning (ML), cloud computing and edge analytics has exponentially grown, and it is no longer possible to scale these use cases with the traditional monolithic system-on-chip (SoC) architecture. In order to address the problem of yield, scalability, and power delivery, the semiconductor industry is

swiftly shifting to chiplet-based architectures- a design paradigm (supported by high bandwidth interfaces) that employs disaggregation of large SoCs into smaller, and more specific dies, in order to achieve increase in performance, a reduction in cost, and a scalability. Although 2.5D and 3D integration technologies enable dense packaging of chiplets on silicon interposers and through-silicon vias (TSVs), computation can no longer be considered as the performance bottleneck because

it has been replaced by communication. RC-limited electrical interconnects are limited in the ability to scale to support modern workloads because of issues with signal attenuation, crosstalk, and signal delays as well as energy-efficiency and thermal dissipation.

These electrical interconnects, albeit mature and CMOS-compatible, exhibit very high limitations in the context of advanced packaging. The 2.5D hinders bandwidth density because signal dejections along long redistribution periphery (RDLs) and microbumps, and the 3D sorted buildings are limited by thermal connections and metal routing congestion using TSVs. Also, aggressive equalization and power-intensive SerDes are needed with high-frequency electrical signaling, which leads to an inappropriate increase in power consumption and latency. These restrictions have given rise to the consideration of other interconnect paradigms, (the most widely considered to date) being silicon photonics, capable of supporting low-loss, high bandwidth, interference immune communications with negligible power overhead. Nonetheless, standardization or wholesale implementation of photonics in VLSI offers issues in vertigo of mixture, fee, and budgetary of heat.

In order to fill these gaps, this study suggests an alternative design of the CMOS-photonic VLSI interconnect architecture to support the chiplet communication in the 2.5D/3D systems. The essence of the concept is that photonic transport with low-power, high-throughput would be merged with programmable CMOS signal processing and control. The dynamic bandwidth and energy resources assignment aspect of this co-design allows the system to allocate available bandwidth and energy resources dynamically in response to the workload in a manner that is energy proportional, and that fits within existing VLSI design flows. The potential benefit of hybrid designs that combine the best of both is unexplored, especially with the system-level thermal integrity, and scalability and reliability considerations in mind.

This piece of work is aimed at the design, simulation, and analysis of a hybrid interconnect system that utilizes CMOS-compatible photonic components (MachZehnder modulators, photodetectors, and microring resonators) in combination with energy-adaptive transceiver logic and TSV-conscientious packaging solutions. With a set of co-simulation tools (CST, COMSOL, Synopsys OptSim), the proposed architecture is discussed in the terms of the bandwidth, energy per bit, latency and thermo stability. This study proves the potential of hybrid CMOS-photonic interconnects as a viable and scalable solution to achieve high performance, energy-efficient interconnection

between chiplets in next-generation heterogeneous VLSI systems by demonstrating up to 10 X improvement in bandwidth over traditional interconnects, and more than 60 percent energy savings.

BACKGROUND AND RELATED WORK

The conventional shortcomings of monolithic SoC scaling have led to the need to employ heterogeneous integration solutions (at least in areas where performance is critical, like AI, HPC, or embedded computing). In this respect, 2.5D and 3D integration have been attempted to become an attractive responsible alternative that offers the possibility of packaging a number of chiplets or functional dies on the same interposer or vertically on TSVs (through-silicon vias) to decrease the length of connectors, increase density, and gain more modularity.^[1,2] Nevertheless, when systems grow elaborate, chiplets-to-chiplet connectivity speed gets to be a primary stumbling block, and given that standard electrical interconnect isn't doable in quickening and thermally stressed conditions, there are problems with this strategy.

Metal traces used as electrical interconnects take the form of silicon-based or organic-based interposers, which are well known to be subject to higher resistive-capacitive (RC) delay, crosstalk, mismatched impedance, and energy per bit when used outside the range of 10 Gbps and below.^[3] These difficulties are even worse in 3D ICs because there is more congestion of routing vertical links and the build-up of heat to the layers. The optical interconnections, particularly silicon photonic waveguides, present an attractive substitute, owing to their high density bandwidth, dissociation to electromagnetic noise, and disadvantageous propagation length over increased chip to-chip lengths.^[4] Nevertheless, full photonic interconnect in VLSI is not practical enough yet to use on a general basis because of the complexities associated with fabrication, integration with CMOS logic and temperature limitations of photonic devices.

There are a number of solutions to this performance gap that are being proposed by industry and academia. Intel AIB- An Advanced Interface Bus is an effort to standardize a high speed electrical interface between chiplets, delivering on compatibility and scalability.^[5] On the same note, the Infinity Fabric is an interconnect protocol by AMD that is general-purpose and scalable multi-chip modules that are utilized throughout their CPUs and GPUs.^[6] Although the electrical signalling within these architectures was optimised, power density and signalling constraints still occur in scale. At the other extreme, further efforts such as DARPA PIPES

project seek to hasten work in the area of photonically connected processor components, as optical waveguides are combined with logic chip within the research to develop the new low-power, high-bandwidth computing fabrics.^[7]

The hybrid solution has an argument supported by complementary research in related fields. To illustrate, Arunabala et al.^[8] emphasize the importance of real-time safety systems in embedded systems and that there is increased realization of high-speed fault detection and communication. In the same breath, recent fluid-thermal modeling research findings by Lim and Lee^[9] highlight the significance of modeling of fluid-thermal in complex aerospace propulsion systems, which in turn alludes to thermal-aware interconnect design in high-density 3D package. Within the scope of reconfigurable computing, Monir et al.^[10] focus on how adaptive architectures help boost ML algorithms, a notion that is also related to the necessity of bandwidth-proportional interconnects capable of adapting to the requirements of a workload. Also, works on IoT sensor network optimization by Siti and Putri^[11] and compact RF antenna design by Arun Prasath^[12] evince other parallels and correspondences of miniaturization, energy efficiency, and high throughput systems design which are pertinent to change in interconnect architecture.

Parallel studies in the surrounding fields reinforce the thesis of a hybrid solution. As an illustration, a paper by Arunabala et. al.^[8] points out the importance of the

real-time safety systems on embedded applications, which emphasize the necessity of quick and speedy fault detection and communication. Likewise, recent works by Lim and Lee^[9] highlight the importance of fluid-thermal modelling in more complex aerospace propulsion systems- a clue to the emergence of thermal-aware interconnect design in high-density 3D package. During reconfigurable computing, Monir et al.^[10] focus on the importance of adopting architectures that would speed up ML algorithms, which is consistent with the concept of bandwidth-proportionate interconnects that are reconfigurable to suit the demands of the workload. There is also the research on optimization of sensor networks in the IoT by Siti and Putri^[11] and the design of compact RF antennas by Arun Prasath^[12] which are indicative of miniaturization, energy efficiency, and high-throughput system design again all pertinent to the evolution of interconnect systems. Although such activities have enabled the development of the field of interconnect technology in different areas, it still lacks proper research to support the development of a hybrid CMOS-photonic interconnect, which has an ultra-high bandwidth potential and ability to achieve energy-proportional communication in chiplet VLSI systems. Current solutions tend to lean towards one aspect or the other either being electrical or photonic that lacks a suitable degree of thermal performance, flexibility, and compatibility with CMOS products necessary to transition into commercial use in 2.5D/3D designs. This paper fills that hole by a proposed co-designed

Table 1: Comparison of Interconnect Technologies for 2.5D/3D Systems

Parameter	Electrical Interconnects	Photonic Interconnects	Hybrid CMOS-Photonic (Proposed)
Bandwidth Density	Low to Moderate (0.5-1.2 Tbps/mm)	High (≥ 2 Tbps/mm with WDM)	High (1.8-2.2 Tbps/mm)
Energy per Bit	High (3-6 pJ/bit)	Low (0.8-1.5 pJ/bit)	Moderate to Low (1.0-1.5 pJ/bit)
Latency	Moderate to High (depends on EQ & SerDes)	Low (picosecond scale)	Low (optical link latency, electrical control)
Thermal Footprint	High (self-heating and TSV coupling)	Moderate (temperature-sensitive devices)	Moderate with thermal-aware layout
Integration Complexity	Low (mature CMOS process)	High (requires photonic-CMOS co-fab)	Moderate (CMOS-compatible photonics)
Signal Integrity	Susceptible to crosstalk, EMI	Immune to EMI, low crosstalk	Improved with EMI-resilient photonic layer
Scalability	Limited by RC scaling	Scalable with WDM and photonic multiplexing	Scalable and adaptable with workload
Cost and Yield	Low (mass-production mature)	High (specialized fabrication)	Medium (modular integration)
Reconfigurability	Limited	Requires tunable lasers and drivers	Dynamic link power scaling enabled
Suitability for Chiplets	Moderate (limited distance, high pin count)	Excellent for mid- to long-range die links	Excellent for mid-range chiplet links

interconnect framework to incorporate CMOS logic into silicon photonic components and dynamically down-scale (or scale-up) the energy of communication as the workload warrants. The architecture is verified by the means of the thermal, electromagnetic, and photonic simulations, and as such, offers a scalable solution to future heterogeneous systems.

HYBRID CMOS-PHOTONIC

System-Level Architecture

The hybrid interconnect system will be built on a layer-oriented chiplet architecture realized in a 2.5D silicon interposer with optical waveguides on-chip. Each chiplet is a functional block, e.g., compute cores, memory modules, or accelerators, connected by high-speed optical and electrical links, among others. Photonic PNI is realized through planar silicon photonic WGs whereas the electrical one is addressed through redistribution layers (RDLs) and through-silicon vias (TSVs). The interposer is produced with silicon-on-insulator (SOI) technology to enable the integration of not only CMOS drivers, but also passive photonic structures. The optical backplane is crossbar with the flexibility to perform dynamic reconfiguration to handle the communication requirements. This is particularly suitable for data-intensive jobs, for example, the training of deep studying types or acceleration of heterogeneous tasks. The electro-photonic linking of the chiplets in the hybrid interconnect stack again an electrical SerDes elements and inter-twined photonic components that are routed on the silicon interposers in waveguides as revealed in Fig. 1(a).

Some example design tools at the system level are Cadence Virtuoso software for physical layout, L-Edit, which is software to floorplan chiplets and CST Microwave Studio is used to model interconnect parasitics and EMI performance. The stack of interposer materials consists of silicon that has been doped to provide electrical routing, low-k dielectrics used to isolate, and silica used to provide photonic waveguide cladding. Tunable lasers and passive optical couplers centered on the photonic interface blocks are utilized in centralized coordination of optical power distribution.

Photonic Device Design

Three main devices in photonic subsystem include microring modulators (MRMs), MachZehnder interferometer (MZI) switches and germanium photodetectors (Ge-PDs). MRMs are chosen based on their small size and low power to use in modulation so that they are dense on the chiplet edges. MZIs are applied in applications where




Fig. 1(a): Hybrid CMOS-photonic interconnect architecture for chiplet communication in 2.5D/3D integration.

there is the need of a higher linearity and superior extinction ratios especially in high-speed connections. On the receiver side, photodetectors are grown epitaxially into germanium on silicon so as to be compatible with CMOS and their responsivity is higher at 1.3-1.55 μm communication wavelengths.

These photonic devices are simulated and simulated within Ansys Lumerical FDTD Solutions, Synopsys OptSim, and Lumerical INTERCONNECT. Wavelength tuning of MRMs is performed through tuning mechanisms such as thermal and electro-optic control, with wavelength being tuned through patterned heater elements using tungsten and designed to minimize thermal cross-talk. Waveguide/fiber and waveguide/laser coupling may be achieved through grating couplers or edge couplers either of which is dependent on the packaging restrictions. KLayout is used to export fabrication-ready layouts which in turn are checked by DRC/LVS based tools integrated into Cadence Photonics IC Suite.

CMOS Interface Circuitry

Interface circuitry is implemented at transmitter and interface ends, in order to close the gap between optical and CMOS logic. Electrical SerDes (Serializer/Deserializer) units are used on the transmitter side to package high speed parallel data into serial streams of data. This is directed into modulator driver circuits that have impedance matching so that clean signal swings can be input to the microring or MZI modulator. The receiver chain will have transimpedance amplifier (TIA) to connect to the photodetector, and clock and data recovery (CDR) logic to reconstruct the digital signals from the optical wave forms.

These are CMOS circuits optimized in 28nm FD-SOI process nodes that are chosen because of their power/aperture trade off. Simulations and designs of circuits are performed using Cadence Spectre, Synopsys HSPICE and Mentor Graphics ModelSim in RTL co-simulation. Keysight ADS and Ansys RedHawk are used to accomplish signal integrity and power integrity analyses. All the energy proportionality features (power gating, bias control, and dynamic voltage-frequency scaling (DVFS) blocks) are co-integrated to be able to respond to the workload variations by power gating, reducing biasing, or scaling down to lower voltage and/or frequency.

They incorporate packaging methods such as flip-chip bonding of CMOS dies, where the micro-bump arrays are applied to both electrical and optical match. Minimized level of insertion loss between photonic I/O ports and the TSV-based interposer channels is ensured using the co-packaged optics platform. Thermal-aware placement algorithms implemented in MATLAB will be utilized in order to mitigate the amount of heat accumulating around photonic regions and therefore, increase link stability and lifetime reliability. The design procedure is co-simulation based as shown in Fig. 1(b), where the elements are (1) photonic device modeling; (2) compact model generation; (3) and electrical-optical circuit co-simulation.




Fig. 1(b): Co-simulation flow for hybrid photonic-CMOS interconnect design integrating physical modeling, compact modeling, and system-level verification.

5D/3D PACKAGING AND INTEGRATION

The combination of hybrid CMOS-photonic interconnections into a chiplet-based system offers several challenges since the advanced packaging system must have the capability to handle electrical and optical signal connections in a thermally reliable, performance, and signal integrity compatible manner. This section expounds on the physical layout and assembly, the package material and through-silicon vias (TSVs) approach to ensure efficient 2.5D and 3D interconnect integration using silicon interposers and through-silicon vias (TSVs). The transition to 2.5D interposers and 3D stacking, as shown in Fig. 2, increases the density of interconnections and scalability of performance much more than the use of planar 2D packaging.




Fig. 2: Structural Comparison of 2D, 2.5D, and 3D Integrated Circuits

Silicon Interposer and TSV Design

The architecture envisaged has a silicon interposer as the main material on which chiplets will be placed and interconnected. Compared to organic substrates, silicon interposers have optical stability and high-routing density characteristics necessary to bring two sets of dense photonic waveguides and electrical traces into alignment. The interposer layer holds planar optical waveguides, typically fabricated utilizing silicon-on-insulator (SOI) wafers, as a means of allowing optical transmission to be done intra- and inter-chiplet with low optical loss. Through adiabatic bends and direction couplers, these waveguides are routed in a manner that does not cause excessive insertion loss and crosstalk. Fig. 3 displays the physical layout of 2.5D and 3D ICs showing how logic, memory and RF chiplets are interfaced with each other through micro-bumps and TSVs on a common interposer or vertical stack.

In order to enact vertical chip-to-chip optical interconnection, a TSV-conscious photonic routing architecture is used. Although TSVs are mostly applied in electrical interconnections, the invention planned here is tactical to co-locate thermal vias, electrical TSVs and optical vertical coupling that includes grating couplers or edge couplers. Photonic routing is performed around TSV clusters to avoid distortion of optical signals caused by substrate




Fig. 3: Physical layout of 2.5D and 3D integrated circuits highlighting chiplet arrangement on interposer and TSV-based vertical stacking

discontinuity. CST Microwave Studio and Lumerical DEVICE are finite-element simulators to prove the efficiency of coupling, waveguide confinement and feasibility of layout within die-level constriction problems.

The die-to-die optical components make use of passive alignment mechanisms with sub-micron alignment precision, and flip-chip bonding and the micro-bump array. Chip-to-board transmission is achieved with optical fibers; they are either co-packaged or integrated through photonic interposer I/Os with a coupling loss that is optimized to be below 2 dB. To support the mentioned packaging workflows, such co-optimization of layout is taken place in various fields of technology with the use of such tools as Mentor Xpedition, Ansys Siwave, and Klayout.

Thermal and Signal Integrity Co-Design

One of the most crucial problems in 3D ICs and photonic systems is thermal management, because high power density, rather than heat and wavelength shift, thermal stress at material boundaries can degrade modulator performance. In this arrangement, thermal vias and embedded heat spreaders (e.g., Cu, AlN or graphene-based composites) have been added below hot spots including modulators and SerDes blocks. Transient and steady-state heat flow through the silicon stack can be simulated using COMSOL Multiphysics to identify hotspots and allow heat dissipation channels to be placed anywhere without disrupting photonic routing.

The thermal design is in intimate connection with signal integrity (SI) optimization. The Keysight ADS and Ansys HFSS help in the analysis of Electrical I/O and SerDes links based on the jitter, eye closure, return loss and crosstalk. Former consideration of sensitivity to signal skew is given between the optical and electrical domains and is important in hybrid schemes of signal. Matched timed-critical I/O The differential pairs run matched impedance and have shielded vias so that

they stay in timing alignment. It can be analyzed using Lumerical INTERCONNECT to analyse the optical SI, such as extinction ratio, insertion loss and inter-symbol interference (ISI) over common channels in the waveguide.

Co-design approach makes photonic and electrical layers to be thermally separated but functionally synchronized so that they can perform stably under changing workloads. Its layout follows the photo-CMOS co-integration foundry design rules and its back-end-of-line (BEOL) metallization has been optimized to co-route TSV, heat sink, and waveguide. The outcome is a tri-staged and highly scalable interconnect architecture where there are systems in data centers, edge AI, and aerospace-grade compute platforms, high-density 2.5D/3D.




Fig. 4: Thermal and Signal Integrity Co-Design in Hybrid CMOS-Photonic 3D Integration

MODELING AND SIMULATION

In order to confirm the feasibility of the suggested combined CMOS-photonic interconnect architecture, an EXTENSIVE simulating platform is created. The platform comprises electromagnetic, thermal, and photonic circuit simulators to analyze signal integrity, thermal stability and the efficiency of communication in general. In this section, the library of modeling tools and techniques used to define the system is outlined and the relevant results of a performance measure based on the simulations is provided.

Electromagnetic and Optical Modeling

The designing and analyzing the electromagnetic properties of the photonic waveguide structures printed on the silicon interposer are carried out with the help of CST Microwave Studio. It supports full-wave 3D simulation of the distribution of transverse electric (TE) mode, effective refractive index and propagation

loss of silicon-on-insulator (SOI) rib waveguide and strip waveguide. Its dimensions and cladding constitution (SiO_2) and the impact of substrate features are taken into account in the simulation to maximize mode confinement and limit insertion loss. Fig. 5(a) (top) displays the simulated electric field map of the TE mode in the waveguide cross-section and the confinement was found to be strong with little leakage, required of intra-chiplet communication. In COMSOL Multiphysics, the ability to model thermal effects upon photonic circuitry and CMOS circuitry exists. The situation, described in Wide-range analysis, is that steady-state and transient thermal behavior is evaluated by the finite-element solver with active devices and hotspots, including laser drivers and SerDes blocks. This model incorporates the heat conduction pathways due to TSVs, heat sinks, and interposer substrates letting them be thermally context aware in placement and co-design. A hotspot close to the area of the modulator can be observed on the thermal distribution on the 2.5D interposer in Fig. 5(b). The temperature gradient also shows that specific thermal vias, and heat spreaders are necessary to ensure device stability of operation without wavelength drift in resonant photonic structures.

To achieve system-level photonic performance, systems such as Lumerical INTERCONNECT and Synopsys OptSim are used to model the optical link budget on chiplet-to-chiplet communication channels. Such parameters as modulator drive voltage, waveguide loss, photodetector responsivity and coupling efficiency have been included in the model. The simulation can give important details

about optical signal depreciation, power penalty, and gains of the receiver. As Fig. 5(c) indicates, the optical power received is plotted as a function of the input power, and a total insertion loss of about 7.8 dB can be realized within the typical operating conditions. Here, the advantages of high signal fidelity are explained by the effectiveness of grating couplers and small loss passive waveguides. As illustrated in Fig. 5, the simulated optical field (a), thermal distribution (b), and optical link performance (c) validate the physical viability and robustness of the proposed hybrid interconnect design.

Performance Evaluation

Several performance measures are read out of the simulation environment to determine the efficiency of the communication in general. The characteristics of bandwidth versus energy consumption are depicted in various link configurations (electrical-only, photonic-only and hybrid). In high bandwidth the hybrid interconnect is more energy efficient with energy per bit values of 1.1151.5 pJ/bit which is far much lower than that of electrical interconnects of over 4.5 pJ/bit at similar data rates. These simulations justify the energy-proportional property of the suggested design, which dynamically modulates the activity of links and laser biasing according to the traffic requirement.

PAM modulation and OOK Simulations are carried out at different noise levels with PAM-4 and OOK modulation with varying lengths of the links. The hybrid system sustains BER less than 10^{-12} up to 20 mm making it appropriate in medium-range chiplet communication




Fig. 5: Simulation results: (a) TE mode optical field intensity (CST Microwave Studio), (b) finite-element thermal analysis of 2.5D interconnect (COMSOL Multiphysics), and (c) optical link budget analysis (Lumerical + OptSim).

on silicon interposers. Latency measurement results indicate that optical connections have an optical delay of sub-nanoseconds that are lower than those witnessed on SerDes-based electrical connections of the same lengths and with equality settings.

Frequency-domain simulations of S-parameters and the time-domain simulation eye diagrams are also used to inspect crosstalk and signal degradation. Hybrid Design: More Clean Eye Openings and Less Jitter than When Electrical links are used alone, especially when parallel links are used. These advancements have been credited to the resistance to electromagnetic invasions by photonic waveguide coupled with reduced dispersion of the waveguide direction.




Fig. 6: Performance Summary

- (a) Energy-per-bit comparison table for Electrical, Photonic, and Hybrid interconnects across different bandwidth levels (200, 400, 600 Gb/s).
- (b) Corresponding plot showing the energy efficiency trend, where the Hybrid approach consistently outperforms traditional Electrical and pure Photonic interconnects, achieving the lowest energy consumption per bit across all bandwidths.

RESULTS AND DISCUSSION

The hybrid CMOS-photonic interconnect architecture that is proposed offers significant characteristics improvements over the conventional electrical interconnects with regard to bandwidth, energy efficiency, and thermal performance. As the simulation results prove, the hybrid approach proves to be effective in major performance parameters, which can be summarised as follows.

The hybrid photonic connection has demonstrated 10 times bandwidth over conventional copper-based connections at a transmission distance greater than 10 mm. Admittedly, electrical I/Os are overwhelmed at approximately 100 Gbps by signal deterioration and crosstalk, whereas the hybrid design is capable of supporting up to 1 Tbps with the application of optical

wave guides, high-speed modulator-detector pairs. As shown in Fig.7, the scalability of bandwidth of the hybrid system is far greater than the electrical or photonic only interconnects.

Dynamic workload simulation is performed, indicating that the hybrid interconnect can save up to 60 percent in energy per bit as compared to electrical links. This is facilitated by energy-adaptive modulation and laser biasing that is dependent on the amount of traffic. Electrical connections require around 160 mW per channel in high-load configurations, whereas the hybrid interconnect will only require about 64 mW compared to that of electrical connection, as illustrated in the bar chart in Fig. 7.




Fig. 7: Bandwidth and Power Comparison

A bar chart illustrating the performance gains of the proposed Hybrid CMOS-Photonic interconnect versus conventional Electrical interconnects. The Hybrid system achieves a 10× increase in bandwidth (1000 Gbps vs. 100 Gbps) while simultaneously delivering up to 60% power savings (65 mW vs. 160 mW), demonstrating its superior energy efficiency and high-throughput capabilities.

The thermal analysis shows that proposed interconnect ensures hotspot remains within operationally acceptable limit of less than 85°C. It is done by designing thermal via, efficient laser placement and thermally conscious photonic routing. This temperature decrease enhances the stability of the modulators and reduces the wavelength drift of the microring resonators hence maintaining the signal integrity with different loads.

The second key advantage is scalability. Hybrid also allows higher I/O density to be achieved without equalization and deleveling, and makes the layout scale of multi-chiplet solutions easier. The trade-off for HPC and AI accelerators is that the integration cost is minimally more because some photonic fabrication processes are involved, but due to a longer term impact on energy and performance payoff, are worth the trade. The EMI-immunity of the optical paths and the lower heating

augment the reliability that forms longer interconnect lifetime and stability in face of changing workloads.

Table 1: Performance Comparison Table

Metric	Electrical Interconnects	Hybrid CMOS-Photonic
Bandwidth Improvement	1x (Baseline)	10x at >10mm
Power Consumption Reduction	0% (Baseline)	Up to 60%
Thermal Hotspot Temp	93° C	<85° C
Area Efficiency	Moderate	High (via interposer)
Integration Cost	Low	Medium

COMPARISON WITH EXISTING APPROACHES

As a comparison to place the benefits of the proposed hybrid CMOS-photonic interconnect architecture in perspective, a comparison with that of conventional copper-based electrical interconnects and photonic-centered solely optical interconnects is undertaken. The four key factors that are compared here are: bandwidth per millimeter, energy per bit, thermal budget and integration cost. The summary is in Table 2.

Copper interconnects are low-cost to make and completely CMOS-compatible but still can have serious performance bottlenecks. The bandwidth density is restricted to 0.5 Tbps/mm, mainly because of RC delay, crosstalk and high frequency losses. Moreover, the energy cost is comparatively high and it averages at 4.8 pJ/bit and is excessive because of resistive heating and low scalability at sub-10nm technology nodes.

Optical-only interconnects have much higher bandwidth density up to 2 Tbps/mm, as well as provide much lower energy/bit (e.g. 1.1 pJ/bit) due to their low attenuation and EMI insensitivity. The problem is, though, that integration comes at a heavy price concerning photonic component alignment, wafer-bonded processes, unique optical I/O interfaces. Although optical links have the advantage that temperature variation has less effect at scale when compared to copper, temperature variations can still impact the operation of devices particularly due to resonant effects of the device structure such as microring modulators.

The hybrid CMOS-photonic technology provides the most balanced level of performance and practicality. Its bandwidth density is 1.8 Tbps/mm, which is close to optical-only technology, and does not require any process changes beyond conventional CMOS technology. Energy use has been kept to 1.2 pJ/bit, which is not a lot more than optical-only but a big improvement on

copper interconnects. Notably, its thermal footprint can be termed as ‘Low’, enabling a high-performing, dense 2.5D/3D IC. Integration cost does not scale too high as a result of modular photonic device-level integration on silicon interposers and standard CMOS-compatible fabrication steps re-use.

This comparison indicates that the hybrid interconnect strategy offers a massively scalable, energy and thermally resilient option to high-performance chiplet based systems without paying extreme costs of or thermal risks of fully optical solutions.

Table 2: Comparative Evaluation of Interconnect Technologies Across Key Performance Metrics

Metric	Copper Inter-connect	Optical Only	Hybrid CMOS-Photonic (Proposed)
Bandwidth per mm	0.5 Tbps	2 Tbps	1.8 Tbps
Energy per bit	4.8 pJ/bit	1.1 pJ/bit	1.2 pJ/bit
Thermal Budget	High	Medium	Low (with TSV-aware design)
Integration Cost	Low	High	Medium

CONCLUSION AND FUTURE WORK

This paper showed the design, modeling and analyzing of a hybrid CMOS-photonic interconnect architecture optimized to chiplet based communication in 2.5D/3D integrated systems. The proposed design hones the CMOS-compatible photonic components (e.g., microring modulators, MachZehnder interferometers, and germanium photo detectors), and they are co-designed with electrical SerDes and transimpedance amplifiers to accommodate high-speed and low-power data transmission between the chiplets.

Extensive electromagnetic, thermal and optical simulations with some toolsets like CST Microwave Studio, COMSOL Multiphysics and Synopsys OptSim were also conducted that showed that the hybrid interconnect achieved up to 10x wider bandwidth and more than 60 percent savings of power dissipated by the traditional copper based electrical interconnects. Thermal profile was inside operating limits and the hot spots were kept below 85 o C because of TSV-aware heat management and thermally isolated photonic layout.

The energy-per-bit capacity of the hybrid link (between 1.1 and 1.5 pJ/bit) proves its feasibility as an energy-proportional communication link, particularly in AI accelerators and edge inference engines, where data

traffic varies extensively. The system also features low bit error rates ($BER < 10^{-12}$) and sub-nanosecond latency that support real-time medium range chiplet communications, without being required to use complex equalization or retimers.

The architectural comparison showed that even though pure optical interconnects are slightly more energy efficient, they fill the qualities with great complexity that adds far too much cost. The hybrid design trade-off is balanced, as it provides the performance advantage of photonics and scalability and maturity of CMOS, which makes this a very attractive solution to high-performance computing and data centres, as well as defence-level, embedded systems.

FUTURE WORK

In spite of the encouraging results of this work, there remain multiple areas in which future developments could be undertaken to increase the scalability, security, and practicability of hybrid CMOS-photonic interconnects. A major trend is monolithic integration, in which CMOS logic and optical devices are combined in a multi-metal, multi-process Photonic Integrated Circuit (PIC) processed on the same die using multiple generations of silicon photonics foundry fabrication processes, thus reducing packaging costs and parasitics. The use of AI in link adaptation can further be resorted to adapting the laser power and bandwidth allocation with real-time workload prediction using machine learning and AI-based capabilities of the energy-efficient communication. The third potential promising area concerns optical security primitives, especially the achievement of photonic physical unclonable functions (PUF) to support authentication of chiplets and encrypted connectivity at the optical level. In addition, it will require extending the thermo-reliability modeling to provide a long-term impact on cyclic stress and high-stack density in 3D interconnect stability to such applications as aerospace and automobiles. Last, to move towards the practical implementation of the concept, it would be necessary to build hardware prototypes based on the commercial photonic interposers to confirm the functionality of the system under various workloads and operating conditions. The above future directions are anticipated to fuel the use of hybrid interconnects in

next-gen high performance and secure chiplet-based platforms.

REFERENCES

1. Patti, R. (2006). Three-dimensional integrated circuits and the future of system-on-chip designs. *Proceedings of the IEEE*, 94(6), 1214-1224. <https://doi.org/10.1109/JPROC.2006.876942>
2. Alam, S. M., et al. (2019). A comprehensive study on 2.5D/3D integration trends for heterogeneous IC design. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, 27(4), 892-904. <https://doi.org/10.1109/TVLSI.2018.2889607>
3. Sylvester, D., & Keutzer, K. (2000). Impact of small feature size on the performance of on-chip interconnects. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 47(5), 429-440. <https://doi.org/10.1109/82.847049>
4. Miller, D. A. B. (2000). Rationale and challenges for optical interconnects to electronic chips. *Proceedings of the IEEE*, 88(6), 728-749. <https://doi.org/10.1109/5.867888>
5. Intel Corporation. (n.d.). Advanced Interface Bus (AIB) Specification. Retrieved from <https://www.intel.com>
6. AMD. (2017). Infinity Fabric Architecture White Paper. Advanced Micro Devices.
7. DARPA Microsystems Technology Office. (n.d.). Photonically Interconnected Processor Elements (PIPES). Retrieved from <https://www.darpa.mil/program/pipes>
8. Arunabala, C., Brahmateja, G., Raju, K., Gideon, K., & Reddy, B. V. (2022). GSM adapted electric lineman safety system with protection based circuit breaker. *International Journal of Communication and Computer Technologies*, 10(1), 4-6.
9. Lim, T., & Lee, K. (2025). Fluid mechanics for aerospace propulsion systems in recent trends. *Innovative Reviews in Engineering and Science*, 3(2), 44-50. <https://doi.org/10.31838/INES/03.02.05>
10. Monir, N. I., Akter, F. Y., & Sayed, S. R. K. (2025). Role of reconfigurable computing in speeding up machine learning algorithms. *SCCTS Transactions on Reconfigurable Computing*, 2(2), 8-14. <https://doi.org/10.31838/RCC/02.02.02>
11. Siti, A., & Putri, B. (2025). Enhancing performance of IoT sensor network on machine learning algorithms. *Journal of Wireless Sensor Networks and IoT*, 2(1), 13-19.
12. Arun Prasath, C. (2025). Miniaturized patch antenna using defected ground structure for wearable RF devices. *National Journal of RF Circuits and Wireless Systems*, 2(1), 30-36.