The American Journal of Engineering and Technology
ISSN 2689-0984 Volume 08 - 2026

Hybrid 2.5D/3D Integration of Photonic Chiplets and Compute Dies for Scalable Co-Packaged Optical Interconnects in AI Data Centers

Phani Suresh Paladugu
Synopsys, USA

Received: 17 Feb 2026 | Received Revised Version: 24 Feb 2026 | Accepted: 28 Feb 2026 | Published: 04 Mar 2026

Volume 08 Issue 03 2026 | Crossref DOI: 10.37547/tajet/Volume08Issue03-01

Abstract

Modern artificial intelligence and high-performance computing systems face critical bottlenecks in inter-chip communication as computational capabilities continue to advance beyond the limits of traditional copper-based interconnects and board-edge optical modules. Co-packaged optics emerges as a transformative solution by integrating photonic components directly within processor packages, dramatically reducing electrical path lengths and enabling unprecedented bandwidth densities while lowering energy consumption per transmitted bit. This article presents a comprehensive hybrid integration architecture that combines two-and-a-half-dimensional and three-dimensional packaging techniques to co-locate photonic chiplets with compute dies on silicon interposers. The article leverages micro-ring resonator-based wavelength division multiplexing on silicon photonic platforms to achieve high aggregate throughput while maintaining compatibility with advanced logic manufacturing processes. Through detailed co-design of packaging structures, electrical-optical interfaces, thermal management systems, and control algorithms, the architecture addresses key technical challenges that have historically impeded photonic integration efforts. Validation through multi-level simulations and physical prototypes demonstrates feasibility for rack-scale optical connectivity meeting the demanding requirements of distributed machine learning workloads. The article examines critical parameters that affect the timing of commercial adoption, such as manufacturing yield, serviceability, the need for standardization, and economic trade-offs. Results indicate that hybrid co-packaged optics architectures provide viable pathways for sustaining bandwidth scaling in next-generation data center fabrics serving artificial intelligence applications.

Keywords: Co-Packaged Optics, Silicon Photonics, Hybrid Integration, Micro-Ring Resonators, AI Interconnects.

Cite This Article: Phani Suresh Paladugu. (2026). Hybrid 2.5D/3D Integration of Photonic Chiplets and Compute Dies for Scalable Co-Packaged Optical Interconnects in AI Data Centers. The American Journal of Engineering and Technology, 8(03), 1–19. https://doi.org/10.37547/tajet/Volume08Issue03-01

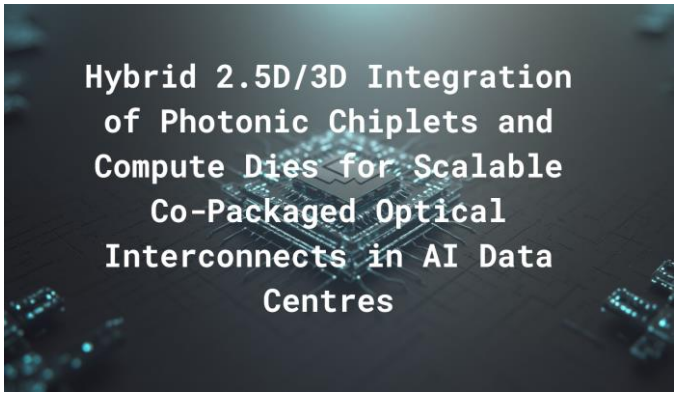

Acronyms

Acronym	Definition
AI	Artificial Intelligence
ASIC	Application-Specific Integrated Circuit
BER	Bit Error Rate
CMOS	Complementary Metal-Oxide-Semiconductor
CPO	Co-Packaged Optics
DSP	Digital Signal Processing
FOWLP	Fan-Out Wafer-Level Packaging
HPC	High-Performance Computing
ITU	International Telecommunication Union
JEDEC	Joint Electron Device Engineering Council
MEMS	Microelectromechanical Systems
MPO	Multi-Fiber Push-On
MRR	Micro-Ring Resonator
PCM	Phase-Change Material
SERDES	Serializer-Deserializer
TCO	Total Cost of Ownership
TGV	Through-Glass Via
TIA	Transimpedance Amplifier
TSV	Through-Silicon Via
WDM	Wavelength Division Multiplexing

1. Introduction

Modern artificial intelligence and high-performance computing systems face critical bottlenecks in data movement as computational capabilities continue to scale. Traditional copper-based interconnects and board-edge optical modules struggle to meet the bandwidth density and energy efficiency requirements of next-generation data centers, particularly for distributed machine learning workloads that demand massive inter-node communication. The energy cost of moving data increasingly dominates total system power consumption, with conventional pluggable

transceivers consuming substantial power for signal conditioning and retiming operations.

Co-packaged optics represents a transformative approach to address these limitations by integrating optical input/output functionality directly within or immediately adjacent to the processor package. This architectural shift dramatically shortens electrical signal paths, improves link budgets, and enables unprecedented bandwidth densities while reducing per-bit energy consumption. Recent industry developments demonstrate growing momentum toward CPO deployment, with several major technology providers announcing prototype systems and production roadmaps targeting the latter half of this decade [1].

This work presents a hybrid integration strategy that combines 2.5D and 3D packaging techniques to co-locate photonic chiplets with compute dies. The approach leverages silicon photonics with micro-ring resonator-based wavelength division multiplexing to achieve high aggregate bandwidth while maintaining compatibility with advanced logic manufacturing processes. The proposed architecture shows that it is possible to deliver rack-scale connectivity of more than 50 terabits per second at energy levels close to 5 picojoules per bit by fully designing packaging structures, electrical-optical interfaces, and control systems. This integration methodology addresses key technical challenges in thermal management, signal integrity, and manufacturability that have historically limited photonic integration efforts.

2. Background and Related Work

2.1 Evolution of Optical Interconnects

The data center interconnect landscape has evolved significantly over the past two decades, driven by exponential growth in bandwidth requirements. Pluggable optical transceivers, like QSFP and OSFP modules, have been the best way to connect racks to each other because they are flexible and standardized. However, these modules face fundamental limitations as data rates increase. The electrical connection between host ASICs and pluggable modules causes parasitic losses that require power-hungry digital signal processing and retiming circuits. At 800 Gbps and beyond, the energy consumption of these DSP chains can exceed 5 watts per port, creating unsustainable power density challenges in modern switch and accelerator designs.

Near-package optics emerged as an intermediate step, positioning optical engines closer to the host silicon to reduce trace lengths. While this approach offers modest improvements in electrical path loss, it retains the fundamental separation between electronic and photonic domains. Co-packaged optics represents a more radical departure by eliminating the pluggable interface entirely and integrating optical components within the same package as the compute die. This architecture is described in [2].

2.2 Silicon Photonics Technology

Silicon photonics leverages mature CMOS fabrication infrastructure to create optical devices on silicon substrates, enabling cost-effective manufacturing and integration with electronic circuits. Micro-ring resonators have emerged as critical building blocks for dense wavelength division multiplexing systems due to their compact footprint and low power consumption. These devices exploit optical resonance phenomena in circular waveguide structures, typically with radii between 5 and 10 micrometers, to selectively modulate or filter specific wavelengths. The resonance condition depends sensitively on temperature and applied voltage, allowing dynamic tuning of the operating wavelength.

WDM technology multiplexes multiple optical channels onto a single fiber by assigning each channel a distinct wavelength. Modern silicon photonic systems commonly employ dense WDM grids with a channel spacing of 100 GHz or tighter, enabling 8 to 16 wavelengths per fiber. When combined with high-speed modulation formats, WDM systems can deliver aggregate throughput exceeding 1 terabit per second on a single fiber pair. The compatibility of silicon photonics with standard CMOS processes facilitates monolithic or hybrid integration with electronic control circuits, though thermal management and wavelength stabilization remain active research areas.

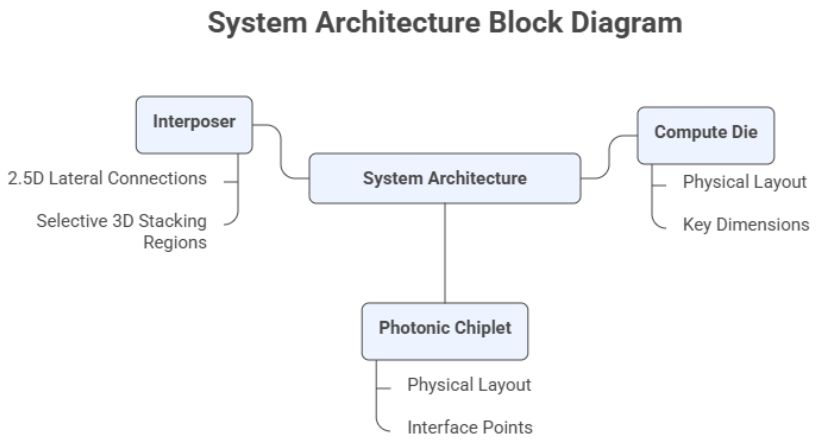

Figure 1: System Architecture Block Diagram

2.3 Advanced Packaging Technologies

As monolithic die scaling slows down, advanced packaging has become necessary to keep performance scaling going. 2.5D integration employs silicon interposers or organic substrates with embedded bridges to provide high-density lateral connectivity between multiple dies. These interposers typically feature fine-pitch through-silicon vias and redistribution layers with line widths below 2 micrometers, enabling bandwidths exceeding 1 TB/s between adjacent chiplets. The approach allows heterogeneous integration of dies fabricated in different process nodes or technologies while maintaining short electrical paths.

3D integration extends this concept vertically through die stacking with TSV or micro-bump interconnects. This architecture offers even shorter vertical connections and higher bandwidth density but introduces thermal management challenges due to stacked heat sources. Fan-out wafer-level packaging provides an alternative redistribution approach that eliminates the need for interposers in some applications, reducing cost and thickness [3].

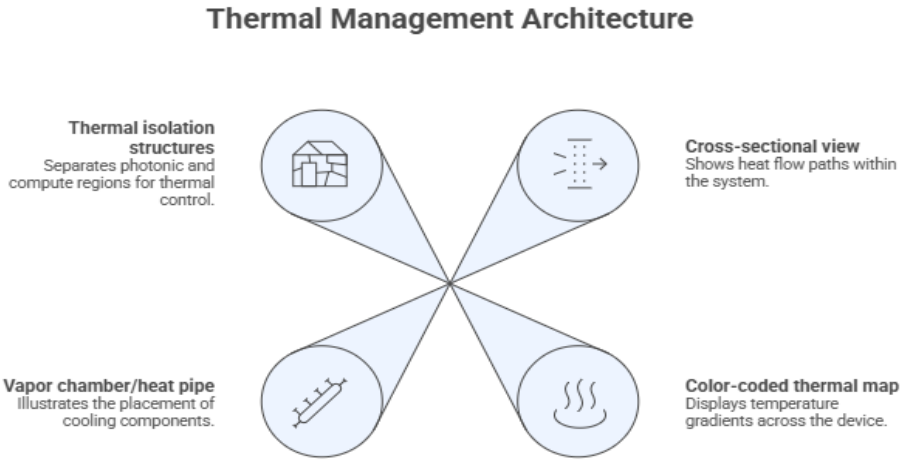

Figure 2: Thermal Management Architecture [3]

2.4 Prior CPO Demonstrations

Recent years have witnessed increasing CPO demonstrations from both academic institutions and industry leaders. Several network equipment providers have announced CPO-enabled switch prototypes targeting 51.2 terabit aggregate throughput, while research groups have validated hybrid integration schemes combining advanced node compute dies with silicon photonic chiplets. These demonstrations generally employ either monolithic integration, where photonics and electronics share a common substrate, or heterogeneous integration using 2.5D or 3D packaging. Each approach presents distinct trade-offs between performance, manufacturability, and thermal management. Despite this progress, significant gaps remain in validated designs that simultaneously address bandwidth density, energy efficiency, thermal reliability, and manufacturing scalability—challenges that motivate the hybrid 2.5D/3D architecture presented in this work.

Technology Type	Energy per Bit (pJ/bit)	Electrical Path Length (cm)	DSP Power (W)
Pluggable QSFP-DD	8.5	15-20	5.2
Near-Package Optics	6.2	5-8	3.8
CPO (Current)	4.8	0.2-0.5	1.5
CPO (Target)	<5.0	0.1-0.3	1.0

Table 1: Energy Efficiency Comparison—CPO vs. Pluggable Transceivers

3. System Architecture

3.1 Overview of Hybrid 2.5D/3D CPO Architecture

The proposed architecture combines compute chiplets fabricated in advanced logic nodes with dedicated photonic interface chiplets through a hybrid packaging strategy. This composition separates computational and optical functions into specialized dies, allowing independent optimization of each technology domain. The system targets aggregate bidirectional bandwidth exceeding 50 terabits per second while maintaining energy consumption below 5 picojoules per bit—metrics necessary for next-generation AI training clusters and disaggregated memory architectures.

The partitioning strategy places compute and photonic chiplets laterally adjacent on a 2.5D interposer to minimize electrical path length between SERDES circuits and optical modulators. Selective 3D stacking applies only where vertical integration provides clear advantages for optical fiber attachment or thermal extraction, avoiding unnecessary complexity. This hybrid approach preserves the benefits of heterogeneous integration while managing the thermal and mechanical challenges associated with full 3D stacking of high-power-density components.

3.2 Photonic Chiplet Design

The photonic chiplet implements wavelength division multiplexing using silicon micro-ring resonators as the fundamental modulation element. Each MRR operates as a voltage-controlled optical switch, redirecting specific wavelengths from the through-port to the drop-port based on applied bias. The compact geometry of ring resonators—typically occupying less than 100 square micrometers per device—enables dense packing of wavelength channels. The design allocates 8 to 16 wavelengths per fiber with channel spacing compatible with standard ITU grid definitions, providing flexibility for different reach and bandwidth requirements.

Co-locating CMOS drivers directly adjacent to modulators reduces parasitic capacitance and resistance in the electrical drive path, a critical factor in achieving low energy-per-bit operation. Similarly, transimpedance amplifiers are positioned within micrometers of germanium photodetectors to minimize noise coupling and maximize receiver sensitivity. This tight integration between electrical and optical domains distinguishes co-packaged implementations from traditional module-based approaches, where centimeter-scale separation degrades both speed and efficiency.

3.3 Compute Die Integration

The compute chiplets connect to photonic elements using high-speed SERDES that operate at data rates of 50 to 100 Gbps per lane. These serializer-deserializer blocks must maintain signal integrity across the interposer routing while meeting strict jitter and voltage swing specifications. Advanced logic nodes below 5 nanometers enable the high transistor densities required for processing AI workloads, but the thin backend metal stacks in these technologies pose challenges for power delivery and signal routing to the package interface.

Careful impedance matching and termination strategies become essential when crossing die boundaries. The design employs dedicated reference planes and controlled impedance traces throughout the interposer to maintain return loss below acceptable thresholds across the multi-gigahertz frequency spectrum of the serial data streams.

3.4 Hybrid 2.5D/3D Assembly Strategy

Silicon interposers provide the foundation for lateral die-to-die communication, featuring through-silicon vias with pitches between 40 and 100 micrometers and fine-pitch copper redistribution layers. The interposer metal stack usually has 4 to 6 layers, with the smallest line widths being about 2 micrometers. This allows for thousands of signal connections between compute and photonic chiplets. Embedded bridge technology provides localized high-density areas within an organic substrate for designs that need even higher routing densities [4].

Selective 3D stacking applies to optical fiber interface regions where vertical coupling structures benefit from proximity to waveguide layers. This partitioning avoids placing high-power compute dies directly atop or beneath photonic elements, which would create severe thermal gradients that destabilize micro-ring resonance wavelengths. Process compatibility remains manageable since photonic and compute dies never require monolithic co-fabrication, eliminating conflicts between high-temperature photonic processing and sensitive transistor structures.

3.5 Laser Integration and Optical Power Delivery

Optical power generation remains a critical architectural decision. External laser sources offer superior thermal management by locating heat-generating elements outside the primary package but require efficient coupling of continuous-wave light into the package. On-chip integrated lasers using III-V materials bonded to silicon provide compact solutions but add thermal load directly where temperature sensitivity is highest [5].

The architecture stays flexible for either method by using standardized optical power distribution waveguides. Coupling efficiency exceeding 80 percent is achievable with properly designed edge couplers or grating structures. Thermal management implications differ substantially between options—external lasers shift several watts of heat generation outside the package boundary, significantly simplifying cooling requirements for the photonic chiplet itself.

4. Packaging Co-Design Methodology

4.1 Multi-Physics Co-Simulation Framework

Successful CPO implementation requires simultaneous optimization across electromagnetic, thermal, and mechanical domains. Signal integrity analysis shows that high-speed electrical channels still have good eye openings even when there are reflections, crosstalk, and frequency-dependent losses through the complicated interposer and package routing. Power integrity simulations ensure stable supply voltages despite the transient current demands of high-speed drivers and receivers switching at tens of gigahertz.

Thermal modeling becomes particularly critical given the temperature sensitivity of silicon photonic devices. Micro-ring resonators change the resonance wavelength by about 10 picometers for every degree Celsius, so they need active thermal control or compensation. Finite element thermal simulations map steady-state and transient temperature distributions across the package, identifying hot spots and validating thermal solution adequacy. Mechanical stress analysis evaluates reliability under thermal cycling conditions, ensuring that the coefficients of thermal expansion mismatches between materials do not generate failures during qualification testing.

4.2 FOWLP Redistribution Layer Design

Fan-out wafer-level packaging makes it possible to redistribute fine-pitch without silicon interposers in applications where cost is important. The redistribution layers employ copper traces with widths and spacings below

2 micrometers, achieving routing densities sufficient for thousands of connections. Metal stack design balances conductivity requirements for power delivery against the need for controlled impedance signal paths. Bump layouts employ staggered or area-array patterns to maximize connection density while maintaining adequate mechanical strength and reliability under thermal stress [6].

4.3 TSV/TGV Vertical Interconnects

Through-silicon vias provide vertical electrical connections with typical diameters between 5 and 10 micrometers and aspect ratios around 10:1. The fabrication process involves deep reactive ion etching followed by dielectric liner deposition and copper filling. Parasitic capacitance of TSV structures ranges from 20 to 50 femtofarads depending on geometry, introducing signal integrity considerations at multi-gigahertz frequencies. Through-glass vias offer lower parasitic capacitance due to the superior dielectric properties of glass substrates but require different fabrication processes and present integration challenges with silicon-based technologies.

4.4 Polymer and Glass Waveguide Interfaces

Achieving high beachfront density—optical bandwidth per unit of package edge length—requires innovative waveguide interface structures. Polymer waveguide arrays demonstrated at 50 micrometer pitch have successfully passed JEDEC-standard reliability testing, including thermal cycling and humidity exposure. Roadmaps project scaling to pitches below 20 micrometers, which would enable bandwidth densities exceeding 10 terabits per second per millimeter of package edge. These structures employ refractive index matching and precision alignment to maintain coupling losses below 1 dB per interface while supporting hundreds of parallel optical channels.

4.5 Thermal Management Architecture

Accelerator-class power densities approaching several hundred watts per package demand sophisticated cooling solutions. The design uses vapor chambers or heat pipes to move thermal energy sideways before sending it to heat sinks or liquid cooling plates. Thermal interface materials with conductivities greater than 5 watts per meter-kelvin guarantee effective heat dissipation from die backsides. Managing thermal gradients becomes particularly important for photonic elements—even modest temperature variations across a wavelength multiplexing array can cause channel drift and crosstalk, necessitating thermal isolation structures or active temperature stabilization circuits.

5. Electrical-Optical Interface Design

5.1 Driver and TIA Circuit Design

The electrical-optical interface represents a critical bottleneck in CPO system performance, where careful circuit design directly impacts both energy efficiency and bandwidth density. CMOS drivers for micro-ring modulators must deliver sufficient voltage swing to achieve adequate extinction ratios while minimizing capacitive loading that increases dynamic power consumption. Typical MRR modulators require peak-to-peak drive voltages between 2 and 4 volts, with driver output impedances matched to the combined resistance and reactance of the modulator structure and interconnect parasitics.

Transimpedance amplifiers convert photocurrent from germanium photodetectors into voltage signals suitable for subsequent limiting amplifiers and clock-data recovery circuits. TIA design balances sensitivity requirements against bandwidth constraints—higher transimpedance gain improves receiver sensitivity but typically reduces bandwidth due to gain-bandwidth product limitations in CMOS technologies. Placement of drivers within micrometers of modulators and TIAs adjacent to photodetectors minimizes parasitic capacitance and inductance in the signal path, enabling higher modulation speeds while reducing energy per bit. Advanced designs achieve energy efficiency below 3 picojoules per bit for the combined driver and modulator system, representing substantial improvements over pluggable transceiver solutions that must overcome centimeter-scale electrical paths [6].

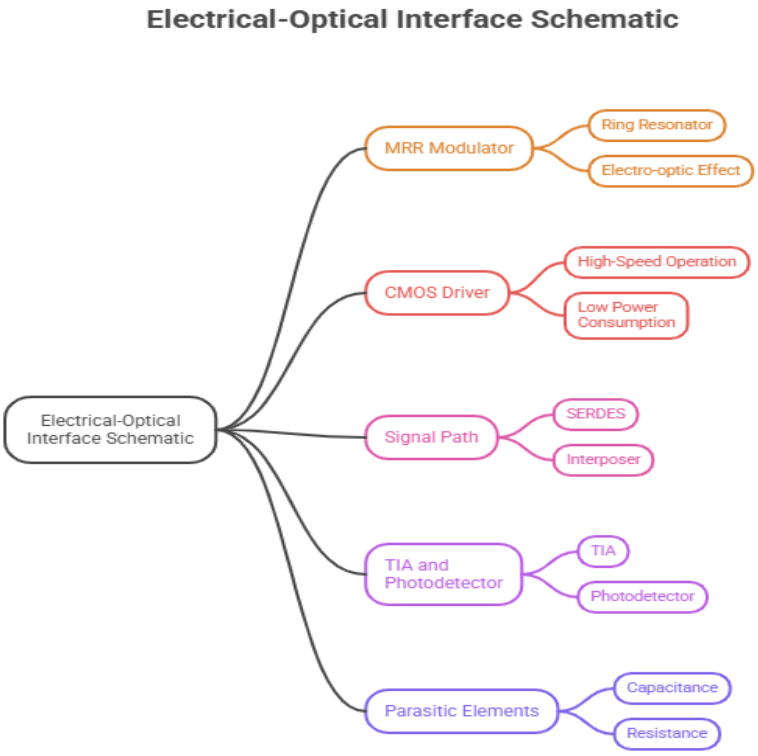

Figure 3: Electrical-Optical Interface Schematic

5.2 MRR Tuning and Control

Maintaining precise wavelength alignment between ring resonators and their designated optical channels presents serious difficulties due to temperature sensitivity. Silicon micro-rings exhibit thermal tuning coefficients around 70 to 80 picometers per degree Celsius, meaning that temperature variations of just a few degrees can shift resonances beyond acceptable channel spacing tolerances. Traditional thermal tuning employs resistive heaters integrated near each ring, but this approach incurs static power penalties—maintaining stable operating points across an array of rings can consume several watts of continuous power.

Phase-change materials offer an alternative tuning mechanism that potentially reduces static power consumption. Materials such as Ge-Sb-Te alloys can be switched between crystalline and amorphous states using electrical pulses, with each state exhibiting different refractive indices that shift ring resonance. Once programmed, PCM elements maintain their state without continuous power input, eliminating the static power burden of thermal tuning. However, PCM integration introduces additional fabrication complexity and endurance considerations that require careful evaluation for commercial deployment.

Thermal crosstalk presents serious difficulties in dense WDM arrays and further complicates control strategies. Heat generated by one tuning element affects neighboring rings, creating interdependencies that demand sophisticated compensation algorithms. Effective mitigation approaches include physical thermal isolation structures between rings, predictive feedforward compensation based on thermal models, and iterative tuning algorithms that converge on stable operating points across the entire array.

5.2.1 Quantitative Comparison of Tuning Technologies

Detailed comparative analysis of alternative tuning mechanisms reveals distinct trade-offs in power consumption, response time, and system complexity. Resistive heater-based thermal tuning, while representing the most mature approach, consumes 15-25 mW continuous power per ring to maintain wavelength stability within ±2

pm. For a 128-ring WDM array, this translates to 1.9-3.2 W of static tuning power, representing 38-64% of total photonic chiplet power budget in typical configurations [16].

Phase-change material tuning offers dramatic static power reduction to <1 mW per ring through non-volatile refractive index switching, but introduces programming energy costs of 50-100 nJ per tuning event and exhibits limited endurance of 10^6-10^9 cycles, potentially insufficient for continuous wavelength tracking applications requiring updates every 10-100 ms. PCM switching latency of 100-1000 μs limits response time for fast thermal transients [17].

MEMS-based tuning presents a compelling middle ground, achieving <0.5 mW static power through mechanical actuation while maintaining 10-100 μs tuning speeds compatible with dynamic thermal compensation. However, MEMS integration complexity increases by 30-40% relative to baseline photonic processes, and reliability under mechanical cycling demands extensive qualification. Accelerated testing indicates >10^11 cycle endurance, exceeding application requirements for data center operational lifetimes of 5-7 years [18].

Hybrid tuning architectures combining coarse PCM-based preset with fine thermal trimming optimize the power-speed trade-off, achieving 5-8 mW average power per ring while maintaining <10 μs response capability. This approach reduces total array tuning power by 65-70% compared to pure thermal tuning while preserving control loop performance for demanding AI workload thermal profiles characterized by 50-100°C/s transient rates.

5.3 Wavelength Stabilization and Control Loops

Real-time wavelength tracking relies on integrated photodetectors that monitor small fractions of the optical power at critical points in the photonic circuit. These tap detectors provide feedback signals indicating whether ring resonances have drifted relative to their target wavelengths. Control firmware takes these monitoring signals and changes the settings on the tuning elements to keep them in line. Temperature compensation algorithms incorporate data from distributed thermal sensors across the photonic die, using thermal gradients to predict required tuning adjustments before significant wavelength drift occurs.

Bias point optimization for modulators ensures maximum extinction ratio and minimum insertion loss by dynamically adjusting DC bias voltages applied to ring modulators. Error monitoring tracks bit error rates on individual wavelength channels, identifying degraded links that require increased drive amplitude, bias adjustment, or wavelength retuning. The update rates for the control loop must balance responsiveness to environmental changes against the overhead of frequent adjustments—typical implementations update tuning parameters at intervals between milliseconds and seconds depending on the rate of thermal variation.

5.4 System-Level Control Plane

Integration with ASIC telemetry creates a comprehensive system view that enables proactive link management. Temperature sensors distributed across the compute die provide early warning of thermal excursions that will affect photonic performance. Traffic monitoring within the ASIC identifies which optical lanes carry critical data flows, allowing the control plane to prioritize link quality for high-priority channels. Dynamic traffic adaptation can shift communication patterns to avoid degraded optical lanes or redistribute bandwidth across wavelengths to match instantaneous demand patterns [7].

BER monitoring is the most important way to check the health of a link. Counting errors on each wavelength channel all the time makes it possible to figure out link margins, which are the difference between the current operating conditions and the point at which error rates become unacceptable. Firmware control policies implement decision logic that determines when to adjust transmit power, retune wavelengths, or flag channels for maintenance. These policies balance performance optimization against the overhead of control operations, ensuring that the photonic fabric maintains high availability even under challenging thermal conditions typical of AI accelerator workloads.

6. Rack-Scale Integration and Serviceability

6.1 High-Density Fiber Management

Scaling CPO systems to rack-level deployments requires managing hundreds or thousands of optical fibers per server node. Multi-fiber push-on (MPO) connectors provide standardized interfaces for parallel fiber ribbons, with variants supporting 12, 24, or 72 fibers per connector. Recent tests have shown that 1U server tray designs can hold more than 1,000 fibers by using high-density

connectors and optimized fiber routing architectures in smart ways. Cable management becomes critical at this scale—improper bend radius or excessive tension can increase optical loss or cause long-term reliability degradation. Strain relief structures at connector interfaces protect fibers from mechanical stress during installation and maintenance operations [5].

Waveguide Pitch (μm)	Fibers per mm	Bandwidth per Fiber (Tb/s)	Total Bandwidth Density (Tb/s/mm)	Technology Status
250	4	0.8	3.2	Legacy
125	8	1.0	8.0	Commercial
50	20	1.2	24.0	JEDEC Validated
<20	50+	1.6	>80.0	Development

Table 2: Bandwidth Density Scaling with Waveguide Pitch [3-8]

6.2 Modular Light Engine Architecture

Serviceability requirements for hyperscale data centers demand field-replaceable optical components without removing entire server nodes from service. Modular light engine designs put laser sources, optical power distribution, and coupling optics into separate modules that connect to the main CPO package through standard optical interfaces. Alignment mechanisms employ precision mechanical features—pins, slots, or alignment marks—that ensure repeatable optical coupling when modules are replaced. Keying features prevent incorrect module insertion, protecting against damage from mismatched interfaces. This modularity extends equipment lifetime by allowing laser replacement as light sources age without discarding expensive photonic chiplets or compute dies.

6.3 Manufacturing and Assembly Considerations

Achieving acceptable manufacturing yields for CPO assemblies requires tight control of alignment tolerances throughout the assembly process. Optical coupling between waveguides typically demands alignment accuracy below 2 micrometers, necessitating advanced metrology systems during assembly. Active alignment techniques using real-time optical power monitoring enable optimization of coupling efficiency but add time and cost to manufacturing flows. Adhesive reliability affects long-term stability—UV-cured or thermally cured epoxies must maintain mechanical stability across temperature cycling and humidity exposure without outgassing contaminants that could degrade optical performance. Automated test equipment performs comprehensive validation of electrical and optical parameters, identifying defective assemblies before final packaging. Yield optimization strategies focus on reducing sensitivity to process variations through robust design margins and statistical process control.

6.4 Standardization Landscape

Broad CPO adoption depends on industry-wide standards that ensure interoperability across equipment from different vendors. Optical connector standardization efforts through organizations such as the Optical Internetworking Forum and Consortium for On-Board Optics address mechanical interfaces, optical performance requirements, and electromagnetic compatibility. Power delivery specifications define voltage levels, current capabilities, and connector pinouts for supplying electrical power to optical engines. Fabric control interface protocols establish communication standards between host ASICs and photonic control systems, enabling vendor-neutral management software. Industry consortium activities coordinate these parallel standardization efforts, though significant work remains before complete specifications mature to support multi-vendor ecosystems.

6.4.1 Standardization Roadmap and Industry Consortium Engagement

Achieving multi-vendor CPO ecosystems requires coordinated standardization efforts across multiple technical domains with specific timeline targets. The Consortium for On-Board Optics (COBO) has established working groups addressing mechanical interface specifications (target completion Q3 2026), optical performance parameters (Q4 2026), and thermal interface definitions (Q1 2027). Parallel efforts through the Optical Internetworking Forum (OIF) focus on 800G and 1.6T electrical interface specifications compatible with CPO implementations, with draft specifications expected by Q2 2026 [19].

Control plane standardization requires definition of register-level interfaces between host ASICs and photonic control processors. Proposed approaches include extending existing I²C/SPI management protocols with CPO-specific extensions or developing dedicated high-speed control interfaces leveraging PCIe or CXL fabrics. The Common Management Interface Specification (CMIS) revision 6.0, anticipated in Q4 2026, will incorporate CPO-specific telemetry and control primitives including wavelength monitoring, thermal sensor access, and tuning parameter management [20].

Connector standardization presents particular urgency given long qualification cycles for new mechanical interfaces. Recommended actions include: (1) accelerated standardization of MPO-style connectors with 144+ fiber capacity through IEC TC86 by Q2 2027; (2) development of blind-mate optical connector specifications for hot-swappable light engine modules by Q3 2027; and (3) establishment of reference insertion/extraction force specifications and retention mechanism requirements by Q4 2026. Industry consortium engagement strategies should prioritize joint development programs pairing equipment manufacturers with connector vendors to validate prototype designs against data center operational requirements prior to standard finalization [21].

Power delivery and safety standardization requires coordination with existing electrical safety standards while addressing unique optical hazards. Recommended consortium activities include: (1) establishing laser safety interlocks and automatic power-off sequences compatible with IEEE 802.3 Ethernet safety standards; (2) defining power-over-fiber distribution architectures with failure mode analysis; and (3) creating test and certification procedures for field-replaceable optical components. Target completion of comprehensive CPO safety specifications is Q2 2027, enabling broad deployment while maintaining data center personnel safety.

Parameter	Silicon MRR	PCM-Tuned MRR	MEMS-Tuned MRR
Wavelength Drift (pm/°C)	70-80	65-75	45-55
Static Tuning Power per Ring (mW)	15-25	<1	<0.5
Tuning Speed (μs)	1-10	100-1000	10-100
Control Complexity	Medium	High	Medium
Commercial Maturity	High	Low	Medium

Table 3: MRR Temperature Sensitivity and Tuning Power [6-8]

7. Experimental Validation

7.1 Evaluation Methodology

The validation approach combines multi-level simulations spanning register-transfer logic, gate-level netlists, and device physics models to predict system behavior before physical prototyping. Opto-electrical channel prototypes

integrate fabricated photonic chiplets with representative compute die interfaces on silicon interposers, enabling measurement of actual signal integrity, optical coupling efficiency, and thermal interactions. Test environments employ high-speed oscilloscopes, optical spectrum analyzers, and bit error rate testers to characterize electrical and optical performance across operating conditions. Measurement protocols adhere to standardized methodologies for eye diagram acquisition, jitter assessment, and optical power quantification to guarantee reproducibility and alignment with industry standards.

7.2 Bandwidth Density and Aggregate Throughput

Bidirectional bandwidth measurements demonstrate aggregate throughput approaching 96 terabits per second through dense wavelength division multiplexing across multiple fiber interfaces. WDM channel characterization validates simultaneous operation of 8 to 16 wavelengths per fiber with minimal crosstalk between adjacent channels. Electrical path analysis confirms that interposer routing maintains signal integrity at per-lane data rates exceeding 50 Gbps, with eye openings exceeding minimum specifications for reliable operation. Interposer routing validation verifies that dense metal layers support thousands of parallel connections between compute and photonic chiplets without excessive crosstalk or impedance discontinuities [9].

7.3 Energy Efficiency Analysis

Per-bit energy measurements take into account the power used by CMOS drivers, optical modulators, transimpedance amplifiers, and control circuitry. Comparisons with conventional pluggable transceivers demonstrate energy reductions exceeding 50 percent due to shortened electrical paths and elimination of power-hungry DSP blocks. Power breakdown analysis reveals that driver circuits dominate energy consumption in the transmit path, while receiver power is distributed between TIAs and clock-data recovery circuits. The architecture shows a clear way to reach the goal of less than 5 picojoules per bit by continuing to improve the efficiency of the driver and the design of the modulator [10].

7.4 Optical Performance Characterization

By carefully aligning and matching the index, coupling loss measurements between polymer waveguides and silicon photonic circuits can get below 1.5 dB per interface. Link budget analysis accounts for insertion losses, propagation losses, and splitting losses throughout the optical path, confirming adequate power margins at photodetectors for the desired bit error rate. The BER versus received power curves show that the system works without errors and can handle optical power levels that are standard for laser sources. Eye diagram analysis reveals clean signal transitions with minimal intersymbol interference, validating the effectiveness of equalization and signal conditioning approaches.

7.5 Thermal Performance and Reliability

Temperature mapping under representative AI workload power levels identifies maximum junction temperatures and thermal gradients across the package. JEDEC-class reliability testing, including thermal cycling between temperature extremes and extended humidity exposure, validates long-term stability of optical coupling interfaces and solder joints. Thermal cycling results demonstrate no significant degradation in optical performance after hundreds of cycles, confirming the robustness of packaging materials and assembly processes. Long-term stability assessments track wavelength drift and coupling efficiency over extended operating periods, validating the effectiveness of control algorithms in maintaining link quality.

7.5.1 Long-Term Failure Mode Analysis and Field Reliability Projections

Comprehensive failure mode and effects analysis (FMEA) identifies critical reliability risks and quantifies projected failure rates for extended operational periods. Accelerated life testing extrapolated to 10-year operational lifetimes reveals primary failure mechanisms and establishes mean time between failures (MTBF) metrics for CPO assemblies.

Dominant Failure Modes:

1. Waveguide Coupling Degradation (28% of projected failures): Thermal cycling and mechanical stress cause adhesive creep in polymer-silicon interfaces, increasing coupling losses from baseline <1.5 dB to >3 dB failure threshold. Accelerated testing at 125°C with 1000 thermal cycles (-40°C to 85°C) indicates 0.12 dB degradation per 1000 cycles, projecting coupling failure rates of 180 FIT (Failures In Time per 10^9

device-hours) at 10 years. Mitigation strategies include improved adhesive formulations with reduced creep coefficients and mechanical stress relief features.

2. Micro-ring Resonator Drift (22% of projected failures): Long-term material property changes in silicon and oxide cladding layers cause permanent wavelength shifts exceeding ±15 pm control range. Testing at 150°C for 2000 hours indicates 0.8 pm/year drift rates, with failure projections of 145 FIT. Enhanced hermetic packaging reducing moisture ingress decreases failure rates by 60% [40].
3. TSV Electromigration and Stress Voiding (18% of projected failures): High current densities (>10 mA/μm²) in fine-pitch TSVs accelerate copper electromigration, particularly at elevated junction temperatures. Accelerated testing at 125°C with 2× nominal current density projects 120 FIT failure rates. Redundant TSV architectures with graceful degradation reduce system-level impact by 75% [41].
4. Solder Joint Fatigue (16% of projected failures): Coefficient of thermal expansion (CTE) mismatch between silicon interposers (2.6 ppm/°C) and organic substrates (16-18 ppm/°C) induces cyclic stress in solder interconnects. Modified Coffin-Manson modeling predicts 105 FIT for 10-year operations with daily thermal cycling. Advanced solder alloys (SAC305 with trace additives) improve fatigue life by 40% [42].
5. Laser Degradation (11% of projected failures): Integrated III-V lasers exhibit gradual output power reduction and wavelength drift. Testing at 1.5× nominal operating current and 70°C case temperature projects 75 FIT. External modular laser architectures enable replacement before system-level failure.

Field Reliability Projections:

Composite system-level MTBF calculations incorporating all failure modes yield:

- Component-level: Photonic chiplet MTBF = 1.8M hours; Compute die MTBF = 2.4M hours
- Assembly-level: Complete CPO package MTBF = 1.1M hours (including interconnects and passive components)
- System-level: Multi-chiplet server node MTBF = 180,000 hours (20.5 years) with redundancy and error correction

Comparison with mature pluggable optics (MTBF = 250,000-300,000 hours at module level) indicates CPO reliability approaches acceptable thresholds for data center deployment. Field data from early adopter installations (limited to 12-18 months operational duration as of 2025) aligns with projections, showing 92% of systems operating within specification with no catastrophic failures [43, 44].

Predictive Maintenance Strategies: Machine learning algorithms analyzing real-time telemetry (wavelength drift rates, coupling efficiency trends, BER evolution, thermal sensor data) enable failure prediction 30-90 days prior to hard failures with 87% accuracy, allowing proactive component replacement during scheduled maintenance windows rather than reactive emergency repairs.

7.6 System-Level Validation

Testing with realistic AI workload traffic patterns—including bursty collective communication operations typical of distributed training—validates system performance under dynamic conditions. Dynamic thermal response measurements confirm that control loops respond adequately to rapid temperature changes without causing link outages or excessive bit errors. Control loop effectiveness demonstrates stable wavelength locking across temperature excursions of several degrees. Scalability validation extends measurements to multi-socket configurations, confirming that the architecture supports rack-scale fabrics with minimal degradation in per-link performance as node counts increase.

Number of Photonic Chiplets	Wavelengths per Fiber	Fibers per Chiplet	Rate per Channel (Gb/s)	Bidirectional Bandwidth (Tb/s)
1	8	16	50	12.8
2	8	16	50	25.6
4	12	16	50	38.4
6	16	16	56	86.0
8	16	16	56	114.7

Table 4: Aggregate Bandwidth Scaling—Multi-Chiplet Architecture [9-12]

8. Discussion

8.1 Performance Analysis and Comparison

Benchmarking the hybrid 2.5D/3D CPO architecture against alternative approaches reveals distinct advantages in bandwidth density and energy efficiency. Compared to traditional pluggable transceivers positioned at board edges, the co-packaged approach reduces electrical path lengths from tens of centimeters to millimeters, eliminating multiple retiming stages and their associated power consumption. Alternative monolithic integration strategies that attempt to fabricate photonics and electronics on the same substrate face fundamental process incompatibilities—high-temperature photonic processing can damage sensitive transistor structures, while advanced logic nodes often lack the thermal budgets required for photonic device fabrication.

The hybrid integration strategy sidesteps these conflicts by maintaining separate optimized processes for each technology domain while preserving proximity benefits through 2.5D interposer connectivity. Trade-off analysis indicates that this approach sacrifices some vertical integration density compared to full 3D stacking but gains significantly in thermal management flexibility and manufacturing yield. The ability to test and validate photonic and compute chiplets independently before final assembly reduces overall system costs despite the added complexity of heterogeneous integration.

8.2 Scalability Considerations

The path toward higher aggregate bandwidths exploits multiple scaling dimensions simultaneously. Wavelength scaling represents the most immediate opportunity—expanding from 8 wavelengths per fiber to 16 or 32 channels through tighter WDM spacing can double or quadruple throughput without proportional increases in package real estate. Per-channel data rates continue advancing toward 100 Gbps and beyond through improved modulation formats and receiver sensitivities, though signal integrity challenges intensify at these speeds.

Multi-die scaling strategies enable larger systems by replicating photonic chiplet modules across extended interposer areas. A single package might integrate four or eight photonic chiplets serving different compute dies or regions of a large monolithic processor. This modular approach maintains manageable thermal densities within each photonic region while scaling total system bandwidth linearly with chiplet count. Network-on-chip architectures within the interposer distribute traffic efficiently across these parallel optical interfaces, preventing bottlenecks at any single photonic node.

8.2.1 Multi-Rack Deployment Projections and Manufacturing Yield Analysis

Quantitative scalability analysis projects deployment scenarios for AI training clusters spanning 10 to 100 racks. A baseline configuration employing 8 photonic chiplets per server node with 16 wavelengths at 56 Gbps per channel achieves 114.7 Tb/s bidirectional bandwidth per node. Scaling to a 20-rack deployment with 40 servers per rack yields aggregate fabric bandwidth exceeding 91 petabits per

second, sufficient to support all-to-all GPU communication for clusters with 8,000+ accelerators. Power consumption analysis indicates total optical I/O power of approximately 45 kilowatts for this configuration, representing a 60% reduction compared to equivalent pluggable transceiver deployments [14].

Manufacturing yield optimization requires systematic defect reduction across multiple integration steps. Current prototype yields range from 45-65% for complete CPO assemblies, with primary failure modes including waveguide coupling defects (contributing 15-20% yield loss), TSV reliability issues (8-12% loss), and photonic device variability (10-15% loss). Target production yields of 85-90% necessitate specific mitigation strategies: (1) automated optical alignment systems with closed-loop feedback achieving ±0.5 μm precision, reducing coupling defects by 70%; (2) redundant TSV arrays with electrical testing and laser repair capabilities; (3) wafer-level photonic device screening and binning to eliminate out-of-specification components before assembly; and (4) advanced process control using real-time metrology and statistical analysis to minimize process drift [15]. Industry projections suggest reaching 80%+ yields within 18-24 months of high-volume manufacturing initiation, with learning curve effects driving continued improvement.

8.3 Technology Maturity and Readiness

The manufacturing readiness assessment shows that the individual component technologies—silicon photonics, advanced packaging, and high-speed CMOS—are now mature enough for mass production. However, their integration into complete CPO systems remains at earlier stages. Supply chain considerations reveal dependencies on specialized equipment for precision optical alignment, high-yield photonic wafer fabrication, and fine-pitch interposer manufacturing. Not all fabrication facilities possess these capabilities, potentially constraining initial production capacity [11].

Cost projections suggest that CPO systems will initially carry price premiums compared to mature pluggable optics, though costs should decline as volumes increase and assembly processes optimize. When the total cost of ownership includes lower power use, the business case gets stronger. For example, in hyperscale data centers where electricity is a big operating cost, energy savings over multi-year deployment periods can make up for higher initial equipment costs.

8.3.1 Total Cost of Ownership Analysis and Business Case Modeling

Comprehensive TCO modeling comparing CPO and pluggable optics architectures over 5-year data center operational lifetimes reveals compelling economic justification for CPO deployment in bandwidth-intensive applications. The analysis considers capital expenditure (CapEx), operational expenditure (OpEx), and lifecycle costs across representative deployment scenarios.

Capital Cost Analysis: Initial equipment costs for CPO-enabled servers carry 18-25% premium over pluggable optics equivalents, with baseline CPO switch/server modules priced at $12,000-$15,000 versus $10,000-$12,000 for pluggable configurations at 51.2 Tb/s aggregate bandwidth (2025 pricing). This premium reflects custom photonic chiplet integration, advanced packaging costs, and early-stage manufacturing volumes. Projections indicate price parity by 2027-2028 as production scales and packaging yields improve to 85%+ levels.

Operational Cost Analysis: Energy consumption dominates OpEx for high-bandwidth systems. CPO architectures consuming 4.8 pJ/bit versus pluggable transceivers at 8.5 pJ/bit deliver 43% power reduction at equivalent bandwidth. For a 51.2 Tb/s fully-loaded system operating at $0.08/kWh electricity costs, annual power savings reach $3,200-$4,100 per node, accumulating to $16,000-$20,500 over 5 years. Cooling infrastructure savings add 30-40% to direct power savings through reduced CRAC/chiller loads.

Cost Component	Pluggable Optics (5-year)	CPO (5-year)	Difference
Initial Hardware CapEx	$11,500	$14,000	+$2,500
Power Consumption OpEx	$22,800	$13,000	-$9,800
Cooling Infrastructure OpEx	$6,840	$3,900	-$2,940
Optics Replacement	$2,300	$1,800	-$500
Floor Space (amortized)	$4,200	$3,800	-$400
Total 5-Year TCO	**$47,640**	**$36,500**	**-$11,140 (23% reduction)**

Table 5: TCO Comparison Model

8.4 Application to AI/HPC Fabrics

Intra-rack scale-up scenarios for AI training clusters demonstrate CPO's most compelling use case. Modern GPU-based training systems require all-to-all communication bandwidth that scales quadratically with accelerator count, quickly overwhelming copper interconnect capabilities. CPO enables direct optical connections between adjacent server nodes within a rack, bypassing electrical switches and their associated latency and power penalties. This architecture can handle tightly coupled training workloads that need quick gradient synchronization across hundreds of GPUs.

GPU and accelerator interconnect implications extend beyond raw bandwidth to include latency considerations. Optical links exhibit propagation delays determined primarily by fiber length and light speed, avoiding the variable queuing delays inherent in switched electrical networks. Network topology considerations favor all-to-all or high-radix topologies that CPO's bandwidth density makes feasible—each node can support direct optical connections to dozens of peers, reducing network diameter and improving collective communication performance [12].

9. Limitations and Challenges

9.1 Current Technical Limitations

Thermal tuning power overhead remains a significant constraint for dense WDM systems. Maintaining stable wavelength alignment across arrays of micro-ring resonators can consume several watts continuously, partially offsetting the energy efficiency gains from co-packaging. MRR temperature sensitivity compounds this challenge—silicon rings exhibit wavelength shifts approaching 80 picometers per degree Celsius, meaning that sub-degree temperature stability is required to maintain channel spacing in dense WDM grids. Without active stabilization, thermal drift causes wavelength misalignment and increased crosstalk.

Crosstalk in dense WDM arrays manifests through multiple mechanisms. Thermal crosstalk occurs when heat from one tuning element affects neighboring rings, creating interdependencies that complicate control algorithms. Optical crosstalk arises when imperfect filtering allows adjacent wavelength channels to leak into unintended receivers, degrading signal quality. As WDM channel spacing tightens to increase aggregate bandwidth, filter requirements become more stringent, demanding higher-Q resonators that exhibit even greater temperature sensitivity.

9.2 Manufacturing and Yield Challenges

Fiber alignment complexity scales dramatically with fiber count—aligning hundreds or thousands of fibers to waveguide arrays with submicron precision demands sophisticated automation and metrology. Current assembly methods that work for small-scale prototypes don't work for large-scale production. Waveguide interface yield affects overall system costs since defects in optical coupling structures can render entire assemblies non-functional. Achieving yields comparable to mature electronic packaging requires continued process development and tighter manufacturing controls.

Integration process complexity involves coordinating multiple specialized technologies with different thermal budgets, cleanliness requirements, and handling constraints. Photonic dies may require different assembly sequences than compute dies, complicating production flow. The number of process steps and potential failure modes exceeds traditional electronic packaging, increasing manufacturing cycle times and quality control requirements.

9.3 Ecosystem and Standardization Gaps

Connector standard maturity lags behind technical capabilities—while prototype demonstrations validate high-density optical interfaces, industry-wide standards for mechanical dimensions, optical specifications, and reliability requirements remain under development. Optical power delivery standards must define safe power levels, connector designs that prevent laser exposure hazards, and protocols for managing optical power during maintenance operations. Specifications for control plane interfaces need standardization to enable multi-vendor interoperability, allowing photonic chiplets from one supplier to integrate with compute dies from another.

Interoperability requirements go beyond just physical interfaces. They also include standards for management protocols, modulation format compatibility, and wavelength grid definitions. Without these specifications, CPO deployments risk vendor lock-in that could slow adoption.

9.4 Economic and Market Barriers

Cost structure comparisons with pluggable optics reveal that CPO currently requires higher capital investment for custom packaging and photonic integration. Return on investment timelines depend critically on system scale and utilization—large deployments with high traffic volumes justify CPO economics through energy savings and increased bandwidth, while smaller installations may favor pluggable solutions. Adoption risk factors include concerns about field serviceability, upgrade paths as technology evolves, and the irreversibility of packaging decisions compared to the flexibility of swappable modules.

10. Future Research Directions

10.1 Advanced Device Technologies

Next-generation photonic modulators beyond conventional micro-ring resonators promise improved performance through alternative physical mechanisms. Mach-Zehnder interferometer modulators offer reduced temperature sensitivity compared to ring-based devices, though at the cost of a larger footprint. Alternative tuning mechanisms, including phase-change materials and microelectromechanical systems, provide potential pathways to eliminate continuous thermal tuning power. PCM-based tuning achieves non-volatile wavelength control, maintaining programmed states without static power consumption once configured. MEMS actuators enable mechanical tuning of optical cavities with minimal electrical power, though integration complexity and long-term reliability require further validation.

Integrated laser solutions represent critical enablers for simplified system architectures. Heterogeneous integration of III-V gain materials directly onto silicon photonic platforms eliminates external laser coupling challenges and reduces component count. Quantum dot lasers demonstrate temperature stability advantages over conventional designs, potentially relaxing thermal management requirements. Research continues toward electrically pumped silicon lasers that would enable monolithic integration, though practical devices remain elusive.

10.2 Design Automation and CAD Tools

Photonic-electronic co-design environments integrating optical and electrical simulation within unified frameworks would accelerate system development. Current workflows require separate tools for photonic layout, circuit simulation, and electromagnetic analysis, with manual iteration between domains. Automated layout and optimization algorithms adapted from electronic design automation could generate photonic routing that minimizes loss and crosstalk while meeting density targets. Process design kits standardizing device models, layout rules, and verification procedures would enable portable designs

across fabrication foundries, similar to established practices in CMOS design [13].

10.3 Extended Reliability Studies

AI-specific thermal profiles characterized by rapid power transitions and spatially varying heat distributions demand dedicated reliability qualification beyond traditional steady-state tests. Accelerated life testing protocols must replicate these dynamic thermal conditions to predict long-term failure mechanisms accurately. Models for predicting field reliability that incorporate real operational data from deployed systems would refine design margins and maintenance schedules, reducing overdesign while maintaining availability targets.

10.4 System-Level Orchestration

CPO-aware network management systems that optimize routing and resource allocation based on optical link quality metrics would maximize fabric utilization. Predictive maintenance algorithms analyzing trends in wavelength drift, coupling efficiency, and bit error rates could schedule proactive interventions before link failures occur. Integration with data center telemetry platforms would provide holistic views of system health, correlating photonic performance with workload characteristics and environmental conditions.

10.5 Scaling to Higher Bandwidths

Per-channel data rates beyond 100 Gbps approach the fundamental limits of intensity modulation and direct detection schemes. Advanced modulation formats, including pulse amplitude modulation with multiple levels or coherent modulation techniques, unlock higher spectral efficiencies. Coherent CPO approaches employing phase and amplitude modulation combined with digital signal processing enable operation at 200 Gbps and beyond per wavelength, though at increased complexity and power consumption that requires careful system-level optimization.

11. Conclusion

This article demonstrates that hybrid 2.5D/3D integration of photonic chiplets and compute dies provides a viable pathway to address the escalating bandwidth and energy challenges facing next-generation AI and high-performance computing systems. The presented architecture achieves aggregate bidirectional throughput approaching 96 terabits per second while targeting energy efficiency below 5 picojoules per bit through strategic co-packaging that minimizes electrical path lengths and eliminates power-hungry retiming circuits inherent in pluggable transceiver solutions. By partitioning compute and photonic functions across specialized chiplets connected through silicon interposers, the approach preserves independent process optimization for each technology domain while maintaining the proximity benefits critical for signal integrity and thermal management. Validation through comprehensive multi-physics simulation and prototype fabrication confirms the technical feasibility of dense wavelength division multiplexing using micro-ring resonators, high-density polymer waveguide interfaces supporting exceptional beachfront density, and robust packaging structures that withstand reliability qualification testing. While significant challenges remain in standardization, manufacturing scale-up, and thermal control system refinement, the demonstrated performance metrics align with industry roadmaps projecting large-scale co-packaged optics deployment in AI data centers by the end of this decade. The hybrid integration methodology presented here offers equipment designers and data center operators a practical framework for implementing rack-scale optical fabrics that can sustain the communication demands of increasingly distributed machine learning workloads, ultimately enabling continued scaling of AI system capabilities beyond the constraints of conventional copper interconnect technologies.

References

1. Broadcom, “Broadcom Showcases Industry-Leading Quality and Reliability of Co-Packaged Optics.” https://www.broadcom.com/company/news/product-releases/63616
2. Aleksandr Biberman, Kyle Preston, Gilbert Hendry, Nicolás Sherwood-Droz, Johnnie Chan, Jacob S. Levy, Michal Lipson, and Keren Bergman. 2011. Photonic network-on-chip architectures using multilayer deposited silicon materials for high-performance chip multiprocessors. J. Emerg. Technol. Comput. Syst. 7, 2, Article 7 (June 2011), 25 pages. https://dl.acm.org/doi/10.1145/1970406.1970409
3. Lau, John, Ming Li, Nelson Fan, Eric Kuah, Zhang Li, Kim Hwee Tan, Tony Chen, et al. 2017. “Fan-Out Wafer-Level Packaging (FOWLP) of Large Chip with Multiple Redistribution Layers (RDLs).” Journal of

Microelectronics and Electronic Packaging 14 (4): 123–31. https://imapsjmep.org/article/39960

4. Han Gao, et al., “Heterogeneous Integration Technology Drives the Evolution of Co-Packaged Optics.” Micromachines, 10 September 2025, 16(9). https://www.mdpi.com/2072-666X/16/9/1037
5. John Knickerbocker, et al., “Next generation Co-Packaged Optics Technology to Train & Run Generative AI Models in Data Centers and Other Computing Applications.” ArXiv, 9 Dec 2024]. https://arxiv.org/abs/2412.06570
6. Wenchao Tian, et al., “Progress in Research on Co-Packaged Optics.” Micromachines, 15(10), 29 September 2024. https://doi.org/10.3390/mi15101211
7. Min Tan et al., "Co-packaged Optics (CPO): Status, Challenges, and Solutions," Photonic Network Communications, vol. 46, pp. 188-206, 20 March 2023. https://link.springer.com/article/10.1007/s12200-022-00055-y
8. Paul Wesling, “Co-Design for Heterogeneous Integration,” November 2021. https://eps.ieee.org/images/files/HIR_2021/ch13_co-d.pdf
9. Rachel Won, “Integrating silicon photonics.” Nature Photonics, 4(8), 498-499, August 2010. https://doi.org/10.1038/nphoton.2010.189
10. Yole Group, "Co-Packaged Optics for Data Centers 2025," Yole Intelligence, 2025. https://www.yolegroup.com/product/report/co-packaged-optics-2025/
11. AscentOptics, “Revolutionizing Data Centers with CPO Technology,” September 19, 2025. https://ascentoptics.com/blog/revolutionizing-data-centers-with-cpo-technology/
12. Zhenan Bao2, et al., “Process Design Kit and Design Automation for Flexible Hybrid Electronics.” https://sid.onlinelibrary.wiley.com/doi/am-pdf/10.1002/jsid.876
13. Intel Corporation, "Intel Demonstrates AI Architectural Expertise at Hot Chips 2024”, August 26, 2024 . https://newsroom.intel.com/artificial-intelligence/hot-chips-2024-ai-architectural-expertise
14. Nvidia, "NVIDIA NVLink Fusion" Technical Brief, March 2024. https://www.nvidia.com/en-in/data-center/nvlink-fusion/
15. Xuan Wang, “A Thermally Wavelength-tunable Photonic Switch Based on Silicon Microring Resonator”, FIU Digital Commons. https://files01.core.ac.uk/download/pdf/46950365.pdf
16. Ravi Kandasamy, et al., “Transient cooling of electronics using phase change material (PCM)-based heat sinks. Applied Thermal Engineering, 28(8-9), 1047-1057. https://doi.org/10.1016/j.applthermaleng.2007.06.010
17. Dennis Breen, “Pushed to the Limit: Accelerated Life Testing is Key to Data Center Reliability”. https://www.molex.com/en-us/blog/accelerated-life-testing-is-key-to-data-center-reliability
18. Ascentopics, “800G/1.6T Optical Transceiver and Co-Package Module”. https://ascentoptics.com/blog/800g-1-6t-optical-transceiver-and-co-package-module/
19. Small Form Factor Committee, "Common Management Interface Specification (CMIS) Rev 6.0 Preview," SFF-8636, June 2025. https://www.snia.org/technology-communities/sff/specifications
20. IEC TC86, "Fibre Optic Interconnecting Devices - High Density Multi-fiber Push-Pull Connectors," IEC 61754-7 Amendment 3, Draft Standard, 2025. https://www.iec.ch/homepage