nature electronics

Review article

https://doi.org/10.1038/s41928-026-01681-6

Co-packaged optics for high-performance computing and artificial intelligence

Received: 18 November 2025

Accepted: 8 July 2026

Published online: 19 August 2026

Check for updates

Byungsoo Kim[1,2,3,10], Soo Ho Choi[4,10], Georgii Zograf[5,6,10], Young Jin Yoo[7,10], Yongmin Baek[1,10], Seokho Kim[1,10], Jongchan Kim[2,3] ✉, Jeehwan Kim[7] ✉, Sang Hoon Chae[5,6] ✉, Hyunseok Kim[4] ✉, Seunghoon Hong[8] ✉ & Kyusang Lee[1,9] ✉

The large data movement required in high-performance computing and artificial intelligence workloads has exposed the fundamental limits of electrical interconnects, where resistive losses, capacitive loading and frequency-dependent distortion increasingly constrain bandwidth, latency and energy efficiency. Optical compute interconnects, which replace electrical links with co-packaged photonic channels, could provide low propagation loss, high bandwidth and superior signal integrity. Here, we examine the development of optical chip-to-chip interconnects and co-packaged optics for high-performance computing and artificial intelligence. We analyse the key domains, including electrical subsystems, electro–optical and opto–electronic conversion interfaces, and optical transmission networks, that determine system-level performance across bandwidth, energy and latency metrics. We also provide a technology roadmap from two-dimensional (2D) co-packaged optics, 2.5D interposer-based integration and 3D heterogeneous stacking, identifying critical challenges in thermal management, manufacturability and standardization that will need to be addressed to establish optical compute interconnects as a foundational communication technology for high-performance computing infrastructure.

Modern high-performance computing (HPC) and artificial intelligence (AI) systems require the generation and movement of large amounts of data. In large-scale model training and other data-intensive workloads, the key constraint on overall efficiency is not computational throughput but the latency and energy overhead to move, access and store data[1–3]. Consequently, the classical compute bottleneck has evolved into a communication-limited regime, where incremental improvements in transistor scaling or logic throughput will yield only marginal system-level benefits unless the interconnect architecture is fundamentally redesigned[4,5].

High-bandwidth memory (HBM) has emerged as an important technology to mitigate data transfer bottlenecks. By vertically stacking multiple dynamic random-access memory (DRAM) dies and connecting them with through-silicon vias (TSVs), HBM can achieve multi-terabits-per-second bandwidths with energy efficiencies in the picojoules-per-bit range[6]. However, the gains are confined to

[1]Department of Electrical and Computer Engineering, University of Virginia, Charlottesville, VA, USA. [2]School of Electrical and Electronic Engineering, Yonsei University, Seoul, Republic of Korea. [3]Department of Integrated Display Engineering, Yonsei University, Seoul, Republic of Korea. [4]Department of Electrical and Computer Engineering, University of Illinois Urbana-Champaign, Urbana, IL, USA. [5]School of Electrical and Electronic Engineering, Nanyang Technological University, Singapore, Singapore. [6]School of Materials Science and Engineering, Nanyang Technological University, Singapore, Singapore. [7]Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA, USA. [8]AI Infra Optimization Team, SK hynix Inc., Seoul, Republic of Korea. [9]Department of Material Science and Engineering, University of Virginia, Charlottesville, VA, USA. [10]These authors contributed equally: Byungsoo Kim, Soo Ho Choi, Georgii Zograf, Young Jin Yoo, Yongmin Baek, Seokho Kim. ✉e-mail: jck@yonsei.ac.kr; jeehwan@mit.edu; Sanghoon.chae@ntu.edu.sg; hyunseok@illinois.edu; seunghoon.hong@sk.com; kl6ut@virginia.edu

https://doi.org/10.1038/s41928-026-01681-6

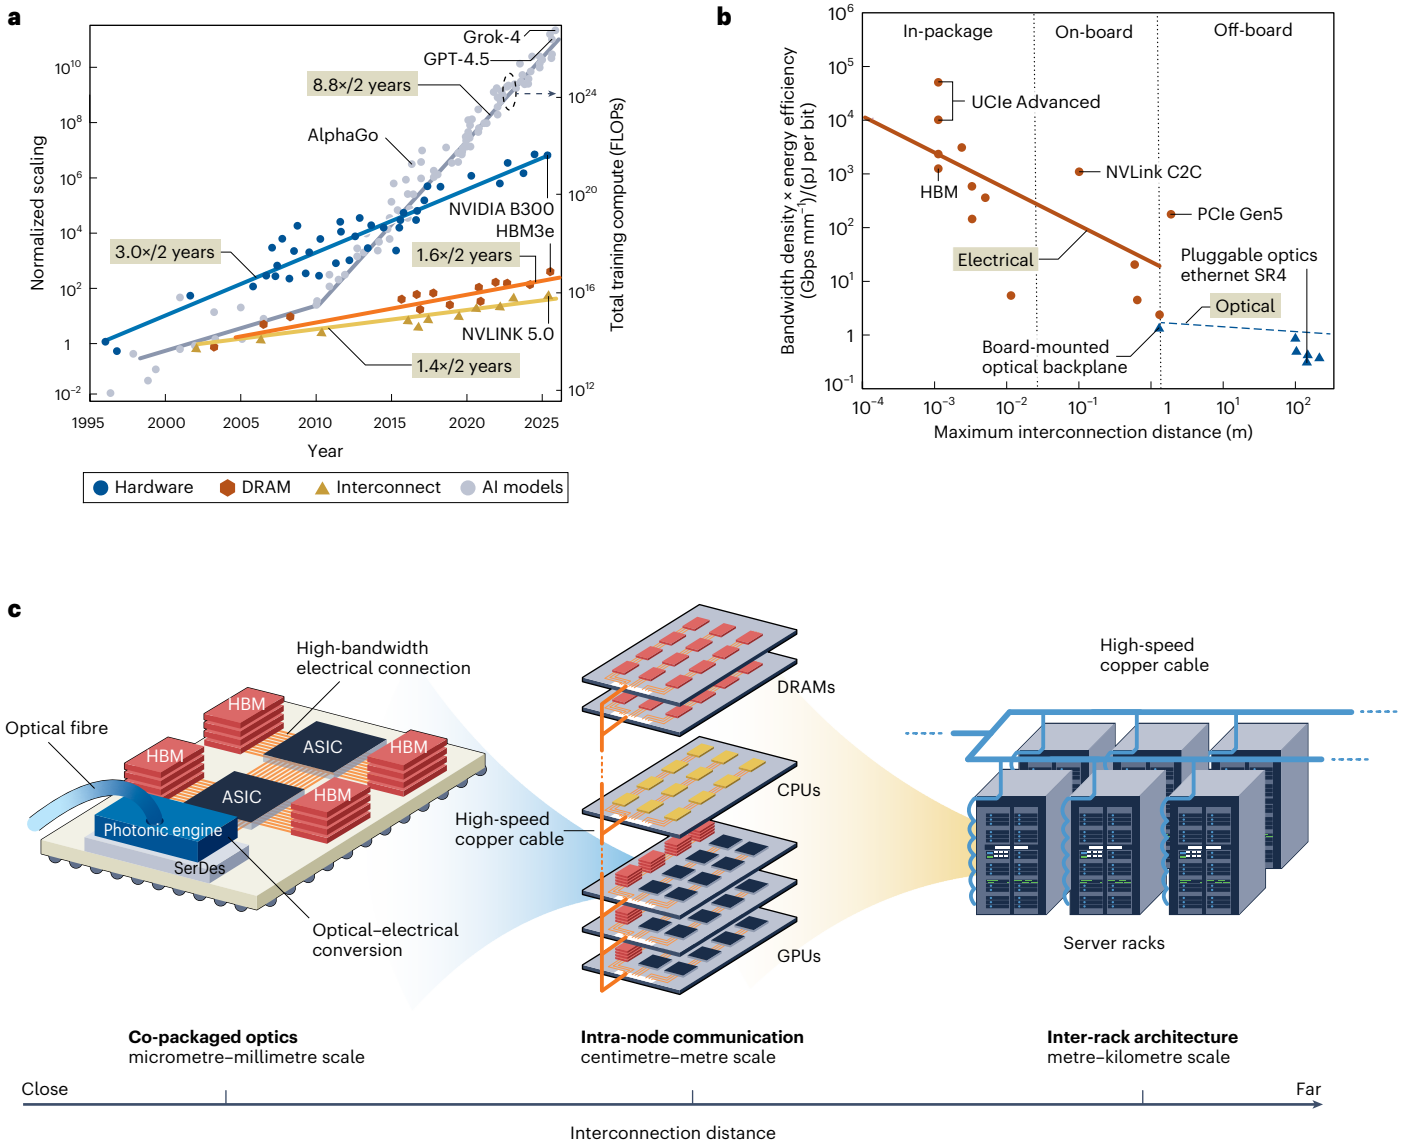

Fig. 1 | Optical interconnect technologies for HPC. a, Historical scaling of computational throughput (blue), DRAM bandwidth (red), interconnect bandwidth (yellow) and AI model training compute (grey). The comparison shows the growing disparity between compute capability and communication performance. The hardware, DRAM and interconnect series are each normalized to their own earliest reported value and shown on the left *y* axis. The underlying metrics are floating-point operations per second for hardware computational throughput and gigabytes per second for DRAM and interconnect bandwidth. AI models are represented by their total training compute in floating-point operations (FLOPs) on the right *y* axis. **b**, Figure of merit, defined as bandwidth density/energy consumption, as a function of interconnect distance for in-package, on-board and off-board domains. The red solid line indicates the projected scaling trend for electrical interconnects and serves as the benchmark of 1 Tbps mm^{-1} bandwidth density at 1 pJ per bit energy efficiency. The blue dashed line indicates the projected scaling trend for optical interconnects and the target performance regime for CPO. **c**, Hierarchical deployment of optical interconnects across spatial scales. At the millimetre scale, optical modules and a network interface controller integrate electronic and photonic interfaces. In an intra-node distributed computing architecture, optical links extend connectivity between units within a rack over centimetre- to metre-scale distances. At the rack and data-centre levels, metre- to kilometre-class optical communication provides low-latency, high-bandwidth links across large-scale computing infrastructures.

intra-package communication; once data goes beyond the memory–logic boundary, electrical interconnects, such as through printed-circuit-board traces, board connectors and direct attach cables, incur resistive and capacitive losses that greatly increase with frequency and distance. Even with advanced packaging approaches such as chip-on-wafer-on-substrate, embedded multi-die interconnect bridge and high-density silicon interposers, the intrinsic energy–distance trade-off of inter-package electrical interconnects remains a fundamental barrier[7,8].

While compute throughput has historically followed a trend of tripling every 2 years, interconnect bandwidth has only advanced by around 1.4-fold over the same period (Fig. 1a), creating a 'bandwidth wall'. Electrical interconnects remain indispensable for short-reach communication owing to their manufacturing maturity and compatibility with silicon complementary metal–oxide–semiconductor (CMOS) technology, but their bandwidth density and energy efficiency are restricted by the intrinsic properties of copper traces, which degrade rapidly as communication distance increases (Fig. 1b). Optical compute interconnects (OCIs) offer an alternative approach for high-bandwidth and energy-efficient interconnects, providing low propagation loss, wavelength-division multiplexing (WDM) for parallelism, and immunity to resistive and electromagnetic degradation. Optical links already dominate long-distance and data-centre-scale communication, and their migration into the compute domain, by

means of chip-to-chip links, in-package optics and on-chip photonic fabrics, could transform HPC and AI infrastructure.

Silicon photonics offers modulators with over 100 GHz bandwidth, photodetectors operating at picojoules-per-bit energy levels and heterogeneous integration of III–V light sources on silicon[9–11]. These advances mean that optical links can combine low latency with high energy efficiency, and optical interfaces have moved increasingly closer to the computational core, from rack-pluggable and on-board optics to fully co-packaged architectures (Fig. 1c). Co-packaged optics (CPO) integrate photonic and electronic circuits within a single package at the micrometre-to-millimetre scale, substantially reducing parasitic losses and latency[12]. Intra-node optical interconnects extend this approach to the centimetre-to-metre regime, connecting processors, memory and accelerators within a compute node, whereas inter-rack optical links can enable high-speed communication across data-centre networks[13–15].

In this Review, we examine the architecture and component hierarchy of OCIs, including the electrical subsystems, electro–optical and opto–electronic conversion interfaces, and optical transmission domain, focusing on the technologies that govern performance and scalability. We benchmark device-level trade-offs and analyse co-packaging strategies from two-dimensional (2D) and 2.5-dimensional (2.5D) interposer-based configurations, as well as fully three-dimensional (3D) heterogeneous stacking. Finally, we provide a roadmap for the development of OCIs, which highlights key challenges in thermal management, manufacturability and ecosystem standardization, and offers a route to establishing the technology as the foundational infrastructure for communication in high-performance and data-intensive computing systems.

Fundamental limitations of electrical interconnects

As data transmission rates exceed hundreds of gigabits per second and device dimensions shrink to the nanometre regime, conventional electrical interconnects encounter intrinsic physical constraints that fundamentally limit bandwidth, energy efficiency and scalability. These limitations originate from a combination of resistive–capacitive (RC) delay, frequency-dependent attenuation, signal distortion, electromagnetic interference and power dissipation, all of which hinder long-distance and ultrahigh-bandwidth data transmission[16,17].

At the chip level, RC delay represents the primary bottleneck. As interconnect linewidths shrink, electrical resistance increases nonlinearly due to enhanced surface and grain-boundary scattering. Simultaneously, the reduction in spacing between adjacent conductors leads to larger parasitic capacitance. The resulting RC time constant introduces a propagation delay that becomes more pronounced with continued scaling. At high frequencies, the skin effect confines current flow to the surface of the channel, thus reducing the effective conductive area and increasing alternating-current resistance that scales roughly with the square root of frequency[18]. This effect, compounded by surface roughness and line-edge irregularities, further increases the resistive losses. In parallel, dielectric loss caused by lagging polarization response in insulating materials introduces additional frequency-dependent attenuation. These effects cause insertion loss, signal jitter and intersymbol interference that degrade signal integrity and limit usable bandwidth[19].

Crosstalk and electromagnetic interference further constrain performance as interconnect density increases. Closely spaced conductors couple both capacitively and inductively, thus the electromagnetic field from one channel induces phase and amplitude noise in neighbouring lines. Such coupling effects intensify in multi-lane and high-speed parallel communication schemes. Although shielding or increased spacing can suppress these couplings, such measures compromise routing density.

At the system level, these physical effects manifest as power dissipation and degraded thermal reliability. Resistive losses in copper traces convert electrical energy into heat, leading to high static and dynamic power consumption. Power loss scales with current density and resistance, both of which increase as interconnect dimensions shrink or signal frequency rises. The resultant local heating combined with the limited thermal conductivity of the surrounding dielectric materials intensifies temperature gradients that increase bit-error rates and accelerate device performance degradation.

These limitations become increasingly prohibitive as communication distances and signalling frequency rise. For terabits-per-second data rates, on-board copper traces exceeding a few centimetres incur attenuation losses of tens of decibels. Signal restoration techniques such as equalization, pre-emphasis, retimers and forward error correction (FEC) can partially recover fidelity; however, each adds circuit complexity, latency and power overhead. For short-distance communication within a rack, direct attach cables, based on twinaxial copper wiring with pluggable connectors and supporting data rates up to 800 Gbps, are often used alongside optical fibre cables to interconnect adjacent boards. As data movement continues to scale across HPC and AI workloads, these cumulative inefficiencies cause electrical interconnects to dominate the overall energy and latency budget. Overcoming these constraints demands a fundamental transition from resistive electrical links to low-loss interconnects that are capable of sustaining bandwidth and efficiency.

Required performance metrics for interconnects in HPC

Modern AI training has fundamentally transformed interconnect requirements. Large-scale models require distributed training where gradient synchronization scales proportionally with model complexity. Three key metrics of bandwidth, energy efficiency and latency define interconnect performance and determine whether computational resources operate seamlessly or remain idle while waiting for data. Meeting all three requirements simultaneously has become essential for scalable HPC.

Bandwidth

The demand for exceeding terabits-per-second-scale interconnect bandwidth arises primarily from the growth of AI model size. Trainable parameters have expanded beyond trillion scale, with end-to-end training requiring on the order of 2×10^{25} floating-point operations[20]. In parallel training, each iteration demands exchanging gradient tensors proportional to the model size, leading to terabit-scale communication per step for the synchronized operations across multiple accelerators. To sustain submillisecond iteration times in distributed AI model training systems, and thereby fully utilize the computational capability of state-of-the-art graphics processing units (GPUs), on-board interconnect bandwidth must exceed hundreds of terabits per second[21,22]. At the component level, modulators exceeding 100 GHz and serializer–deserializer (SerDes) reaching 224 Gbps per lane have been demonstrated. In the optical input/output (I/O) chiplet space, Intel's OCI delivers 4 Tbps bidirectional (32 Gbps × 64 lanes, 8 dense WDM wavelengths) at over 4× the shoreline bandwidth density of PCIe Gen6[23], Ayar Labs's third-generation TeraPHY achieves 8 Tbps bidirectional (32 Gbps × 128 channels) with UCIe die-to-die compatibility[24] and Lightmatter's Passage L200 targets 32–64 Tbps per 3D CPO chiplet[25]. These per-chiplet bandwidths already surpass what electrical die-to-die interfaces can sustain beyond a few millimetres of reach, while extending chip-to-chip communication to >100 m over a single-mode fibre. However, achieving over 100 Tbps aggregate bandwidth per compute node, required for large-scale distributed training, demands tiling multiple optical chiplets per accelerator package with coordinated co-design of packaging, fibre routing and thermal management.

Energy efficiency

State-of-the-art HPC racks typically consume between 30 kW and 100 kW depending on configuration and workload while delivering aggregate bandwidths exceeding 100 Tbps (ref. 26). To constrain interconnect power dissipation to below 10% of the total system power budget, the energy cost of data transmission must be reduced to the subpicojoules-per-bit regime. At the component level, silicon photonic transmitters have demonstrated 0.7 pJ per bit, whereas 112 Gb s^{-1} SerDes still operate in the 1–2.3 pJ per bit range[10,27]. Among optical I/O chiplets, Intel's OCI consumes about 5 pJ per bit with a target of <3 pJ per bit (ref. 23), Ayar Labs' TeraPHY operates at <5 pJ per bit (ref. 24) and Lightmatter's Passage demonstrates 4.6 pJ per bit wall-plug efficiency (2.6 pJ per bit for photonics and light source, about 2 pJ per bit for co-packaged SerDes)[25]. While these represent substantial improvements over pluggable transceivers (15–30 pJ per bit), all remain well above the subpicojoules-per-bit target. The gap is dominated by residual digital signal processing (DSP) and FEC processing, laser wall-plug efficiency, thermal regulation and SerDes power, overheads that do not diminish with photonic device scaling alone and require coordinated system-level co-design[28]. Achieving this target emerges as a fundamental requirement for sustaining system scalability, thermal reliability and operational efficiency in HPC computing infrastructures.

Latency

A sub-10-ns chip-to-chip latency target emerges from the synchronization requirements of large-scale distributed training clusters[29]. As interconnect bandwidth scales into the multi-terabits-per-second regime, end-to-end latency becomes increasingly dominated by fixed overheads. Among optical I/O chiplets, Intel's OCI targets <10 ns chip-to-chip latency plus time-of-flight (about 5 ns m^{-1} in fibre)[23], and Ayar Labs' TeraPHY achieves <10 ns chiplet-to-chiplet latency without requiring FEC[24]. These represent >15× improvements over pluggable baselines. However, these consider only the optical link contribution. In full system deployments—incorporating multi-hop switch traversals, host-side SerDes, residual FEC and network protocol overhead—end-to-end round-trip latency is in the microsecond regime, a threshold at which synchronization primitives dominate the critical path and consume more than half of the total iteration time in distributed training across thousands of accelerators. Therefore, large-scale distributed training exhibits major degradation when the switch latency within the rack exceeds microseconds. To constrain synchronization overhead within acceptable envelopes as systems scale to thousands of GPUs in an HPC system, the chip-to-chip latency must be reduced to below 10 ns, enabling the submicrosecond switching latency required for sustaining near-linear scalability, minimizing synchronization overhead and enabling efficient operations in distributed AI systems.

Core components in OCIs

As illustrated schematically in Fig. 2a, chip-to-chip OCIs comprise three tightly coupled functional domains: the electrical, electro–optical (E–O)/opto–electronic (O–E), and optical layers. The electrical domain governs data processing and flow management, including serialization, synchronization, buffering and switching across chiplets. The E–O/O–E interface forms the transduction boundary between electronic and photonic domains where electrical signals are converted into modulated optical carriers via transmit-side modulators or emitters, and optical data are reverted into electrical currents at the receiver side through photodetectors. The optical domain then transports information through low-loss dielectric waveguides, leveraging WDM and coupling elements to achieve high-bandwidth, low-latency communication. These domains establish the foundation for scalable, high-bandwidth and energy-efficient OCI architectures.

Electrical subsystems

Electrical subsystems constitute the control and timing backbone of OCIs through serialization, synchronization, buffering and switching operations. At the link level, the transmitter incorporates data encoding, serialization, phase-locked-loop-based clock multiplication and a differential output driver with pre-emphasis to place high-speed data streams onto the channel. SerDes architectures perform parallel-to-serial and serial-to-parallel conversion that enables high-aggregate bandwidth with fewer transmission lanes (Fig. 2b). The serialized electrical signals subsequently drive electro–optic modulators or directly modulate integrated light sources. Differential output drivers with pre-emphasis boost the amplitude of high-frequency components to compensate for frequency-dependent signal attenuation in the transmission channel. Multi-level modulation formats such as four-level pulse amplitude modulation (PAM-4) improve spectral efficiency but reduce noise margins, necessitating complex DSP and FEC circuits (Fig. 2c).

On the receiving end, adaptive equalization restores distorted waveforms by mitigating intersymbol interference and compensating frequency-dependent losses. The deserializer converts the high-speed serial stream back into its native parallel format. Clocking and synchronization circuits, including delay-locked loops, phase-locked loops, and clock and data recovery units, are indispensable for retiming and jitter suppression, but they contribute substantially to area and power consumption.

Despite advances in transmitter and receiver architectures, electrical links typically consume energy in the multi-picojoules-per-bit range and have cumulative latencies that exceed the subnanosecond targets demanded by distributed AI and HPC systems[27]. As increased packaging density exacerbates thermal constraints and limits further parallelization, co-optimization of electrical subsystems with integrated photonic architectures is essential to sustain performance scaling and energy efficiency.

Electro–optical and opto–electronic interfaces

The E–O/O–E interfaces define the critical transduction boundaries between electrical and optical domains. At the transmitter, digital data are encoded onto continuous-wave optical carriers using optical sources and electro–optic modulators (EOMs). Optical sources provide coherent carriers that drive data transmission. Distributed-feedback (DFB) lasers deliver narrow-linewidth emission ideal for dense WDM systems while vertical-cavity surface-emitting lasers offer cost-effective and low-power solutions. Traditionally, externally coupled laser sources dominate industrial CPO deployments because they decouple the source thermal environment from the switch application-specific integrated circuit (ASIC), ease laser replacement and leverage mature pluggable-module packaging. However, external sourcing carries intrinsic penalties that compound with channel count. Each chip-to-fibre interface contributes 1–3 dB of coupling loss[30], polarization-maintaining fibre and submicrometre alignment margin increase packaging complexity and cost, and fibre-burn risk grows with per-channel optical power. External-source architecture sustains a 1–5 dB link-budget penalty relative to integrated source architecture at a given bandwidth[31]. These penalties become acute as WDM channel counts extend towards the tens-of-channels regime. Links built from external DFB arrays require independently stabilized and wavelength-locked cavities across all transmit channels, and the aggregate footprint of thermoelectric coolers, control electronics and polarization-maintaining fibre bundles scales in direct proportion to the number of channels[31,32]. However, heterogeneous III–V/silicon integration approaches are becoming more widespread, which can be realized via wafer bonding, direct epitaxy or micro-transfer printing, combining superior optical gain with scalability. However, challenges persist in thermal management[33,34], yield[35,36] and spectral stability[37,38].

EOMs are used to modulate optical carriers through amplitude, phase or frequency modulation (Fig. 2d) and can be broadly classified into four categories based on their underlying mechanism: electro-absorption modulators, Mach–Zehnder modulators,

https://doi.org/10.1038/s41928-026-01681-6

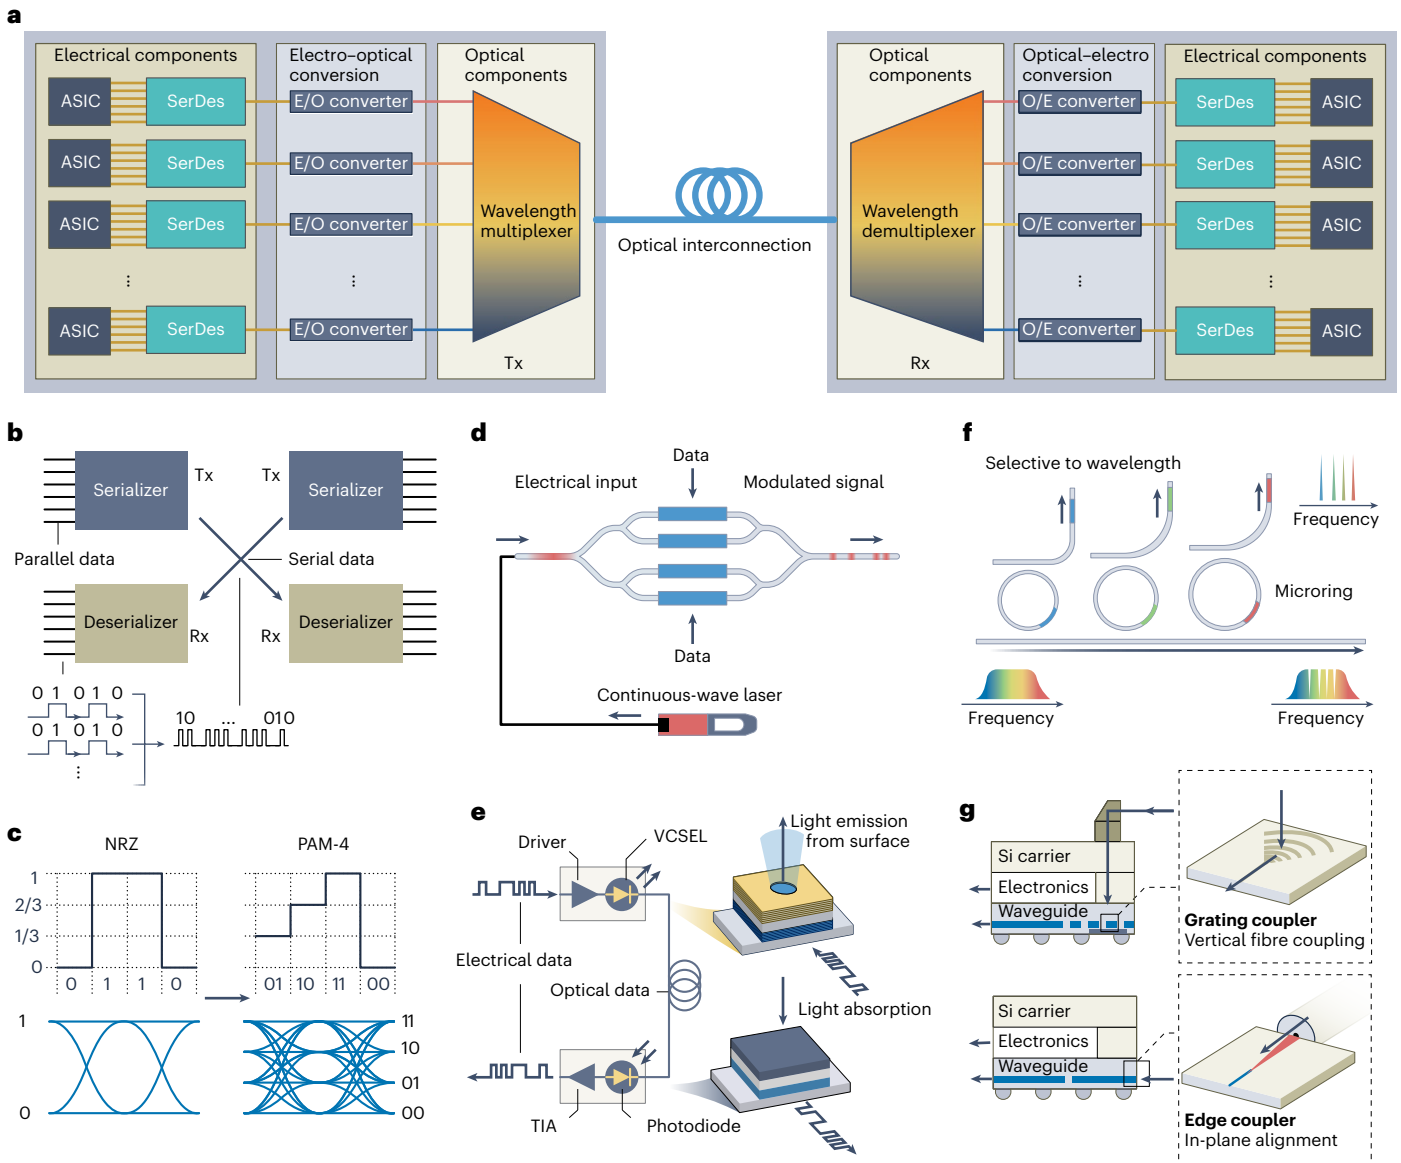

Fig. 2 | Core technologies enabling OCIs. a, Schematic of end-to-end high-bandwidth optical data transmission. Electrical signals from computing units are serialized and converted to optical carriers through E–O transmitters (Tx), multiplexed across distinct wavelengths, transmitted through optical channels and recovered at the receiver (Rx) through O–E conversion. **b**, SerDes architectures that perform parallel-to-serial and serial-to-parallel conversion for bandwidth aggregation across wide electrical buses. **c**, Representative modulation formats including non-return-to-zero (NRZ) and 4-level pulse amplitude modulation (PAM-4), where PAM-4 doubles the spectral efficiency but reduces the noise margin. **d**, Operation of a Mach–Zehnder modulator that uses interference between its two optical paths to encode electrical data as intensity-modulated carriers. A continuous-wave laser, provided externally, typically supplies the input illumination. **e**, Active optoelectronic components comprising vertical-cavity surface-emitting lasers (VCSELs) driven by electrical signals for optical signal emission (top) and photodiodes with TIAs for optical-to-electrical conversion at the receiver (bottom). **f**, Wavelength-selective microring resonators that provide on-chip optical routing, filtering and multiplexing. **g**, Chip-to-fibre and fibre-to-chip coupling strategies including grating couplers for vertical coupling and edge couplers for in-plane, low-loss interfacing with silicon photonic waveguides.

carrier-depletion silicon modulators and Pockels-effect modulators in ferroelectric materials. Electro-absorption modulators exploit field-induced absorption changes through the Franz–Keldysh or quantum-confined Stark effects, offering subpicojoules-per-bit switching energy and compact footprint suitable for dense interconnects, although their operating wavelength range is narrow and temperature sensitive. Mach–Zehnder modulators use interferometric phase modulation between two parallel optical paths to achieve highly linear operation with bandwidths exceeding 100 GHz, at the cost of a larger footprint that limits integration density. Carrier-depletion silicon photonic modulators using p–n or p–i–n junctions offer CMOS compatibility and scalability, although the weak plasma dispersion effect in silicon constrains their efficiency and typically requires longer phase shifters and higher drive voltages.

EOMs based on the Pockels effect in ferroelectric materials such as lithium niobate ($LiNbO_3$, LN) and barium titanate ($BaTiO_3$, BTO) achieve subvolt drive operation and bandwidths over 100 GHz. Thin-film LN (TFLN) has emerged as the current benchmark for integrated high-speed modulation, especially in hybrid Si/LN platforms, whereas BTO offers even larger electro–optic coefficients suitable for ultra-compact and low-power modulators[39–42]. However, heterogeneous integration of these ferroelectrics on silicon has reached markedly different stages of maturity. Wafer-scale LN-on-insulator substrates fabricated via smart-cut technology have shown rapid

maturation of TFLN modulator technology[43]. Direct argon-based etching of the LN layer can achieve maximal mode confinement and electro–optic efficiency[44], while heterogeneous bonding of unpatterned TFLN onto pre-patterned Si or silicon nitride (SiN) waveguides avoids LN etching and the associated Li^+ contamination concerns to enable foundry compatibility[45]. BTO integration, by contrast, remains at a pre-commercial stage, as high-quality films require epitaxial growth on strontium titanate ($SrTiO_3$)-buffered silicon at substrate temperatures exceeding CMOS back-end-of-line thermal budgets, although recent low-temperature wafer-bonded BTO on 200 mm photonic platforms offers a scalable route[46,47]. Two-dimensional materials such as graphene and transition metal dichalcogenides are emerging alternatives owing to their strong light–matter interactions and broadband operation, although challenges in uniformity, efficiency, insertion loss and manufacturability remain[48,49]. Across all platforms, thermal stability is a major constraint. Refractive index fluctuations as small as 10^{-4} can induce gigahertz-scale wavelength drift, necessitating athermal designs or active temperature compensation for stable operation.

At the receiver, the O–E interfaces form the core of signal conversion, where photodiodes convert optical carriers back into electrical currents (Fig. 2e). Germanium-on-silicon p–i–n photodiodes dominate current OCI implementations owing to their telecom-band responsivity and CMOS compatibility, while avalanche photodiodes (APDs) enhance sensitivity via internal gain at the cost of excess noise and higher bias[50,51]. Transimpedance amplifiers (TIAs) convert photocurrents to voltage signals with gain-bandwidth product, which governs receiver speed, sensitivity and noise performance. Monolithic or co-packaged photodiode–TIA integration reduces parasitic capacitance and enhances overall energy efficiency, but demands strict thermal and impedance matching. In practice, p–i–n photodiodes with CMOS TIAs are widely used in short-reach links, while APD-based frontends and coherent receiver architectures are being explored to extend sensitivity and reach[52].

Optical transmission domain

The optical domain provides the transmission backbone of OCIs, transporting data through low-loss dielectric media. Central to this domain is WDM, which combines multiple wavelength channels into a single waveguide to scale bandwidth without proportional increase in footprint or power consumption (Fig. 2f).

Microring resonators and arrayed waveguide gratings are the primary multiplexing technologies. Microring resonators offer exceptional compactness and high-integration density but are thermally sensitive, while arrayed waveguide gratings provide robust and thermally stable operation with larger footprints. Coarse WDM offers wide channel spacing and high tolerance to fabrication and thermal variations, whereas dense WDM achieves higher spectral density at the cost of tighter process control and active wavelength stabilization. Emerging technologies, such as microresonator frequency combs and inverse-designed (de)multiplexers, offer scalable, on-chip multiwavelength generation, balancing integration density with spectral precision[53–56].

Optical waveguides provide the physical transmission paths of the link. Silicon-on-insulator platforms enable ultra-dense routing but suffer from nonlinear absorption at high optical intensities. Silica provides ultralow propagation losses below 0.001 dB cm^{-1}, but with limited integration density, and SiN offers broadband transparency, back-end CMOS compatibility and propagation losses below 0.1 dB cm^{-1} (ref. 57). TFLN and BTO platforms introduce active electro–optic tunability, enabling dynamic routing, nonlinear signal processing and hybrid microwave–optical transduction.

Optical couplers interface on-chip waveguides with external fibres, playing a pivotal role in power efficiency and packaging scalability (Fig. 2g). Edge couplers achieve subdecibel insertion losses but require micrometre-scale alignment precision, whereas grating couplers support wafer-scale vertical coupling with simplified packaging at the expense of polarization dependence and limited spectral bandwidth. On-chip multimode interference and directional couplers extend coupling functionality for channel splitting, combining and routing within dense WDM systems[58].

At the packaging level, optical interposers integrate waveguides, couplers and multiplexers directly within the substrate, enabling dense optical I/O and minimal electrical routing overhead. Recent implementations have achieved subdecibel inter-chip coupling losses and broadband operation across O- and C-bands. Remaining challenges include thermal management, alignment precision and wafer-level fabrication yield, which must be addressed to enable scalable CPO systems[59–61].

Benchmarking and integration readiness

The scalability of OCI technologies is determined not solely by record device metrics but also by the balance between intrinsic performance, process maturity and integration readiness. Representative benchmarks delineate how these factors diverge across materials and device classes, defining the trade-offs that govern system-level scalability (Table 1). The viability of different OCI architectures thus depends on coordinated progress across materials, devices and packaging. Bridging the gap between device-level performance and manufacturable implementation, particularly for integrated light sources, high-speed modulators and wafer-scale coupling architectures, will be critical to achieve subpicojoules-per-bit energy efficiency, sustain terabit-scale bandwidth density, and enable large-scale deployment in HPC and AI infrastructures.

Co-packaged optics

The transition from monolithic system-on-chip designs to modular chiplet-based architectures has reshaped the landscape of HPC and AI workloads. In this paradigm, logic, memory, accelerators and photonic functions are disaggregated into specialized dies, each fabricated using process nodes optimized for their respective functionalities. Interconnecting these heterogeneous chiplets demands high-density, low-latency communication links that maintain the bandwidth uniformity within a unified system-on-package[62]. CPO represents a pivotal interconnect technology advancement towards this goal through seamless convergence of photonic and electronic integration, enabled by minimizing the physical separation between optical and electrical domains.

In conventional pluggable optics, transceivers are mounted at the periphery of boards, where centimetre-scale copper traces introduce substantial resistive loss, frequency-dependent signal distortion and power dissipation as data rates approach and exceed 800 Gbps (Fig. 3a). On-board optics mitigate these effects by co-locating photonic integrated circuits (PICs) with electronic integrated circuits (EICs) directly on board, thereby reducing electrical paths and parasitic delay (Fig. 3b). CPO extends this approach further by integrating PICs and EICs within a unified package through a 2.5D interposer-based platform (Fig. 3c) or fully 3D-integrated architectures (Fig. 3d). These configurations achieve interconnect lengths on the micrometre-to-millimetre scale with minimal parasitic loss while enabling aggregate bandwidth and energy efficiencies beyond the reach of conventional electrical interconnects. Recent implementations of CPO have demonstrated up to 50% lower power consumption compared with pluggable transceivers with scalability beyond 12.8 Tbps per-link throughput[22].

2D and 2.5D integration

Early generations of CPO used a 2D configuration in which photonic and electronic dies are mounted side by side on a shared substrate (Fig. 3e). Electrical interconnects are provided through wire bonding or flip-chip soldering that are fully compatible with established CMOS packaging flows. However, as network switch fabrics expand to hundreds of ports and per-lane data rates exceed 100 gigabaud, extended electrical trace lengths inherent to planar layouts introduce substantial

https://doi.org/10.1038/s41928-026-01681-6

Table 1 | Benchmarking key components for hybrid optical interconnect links

Component		Benchmark metrics		Current state of the art	Integration readiness	Coupling/ packaging	Thermal/ operation	References
		Output power	**Linewidth**					
Laser	III–V bonding	>5 mW	~40 kHz	Ultralow-noise DFB on silicon	Wafer bonding advancing	Alignment sensitive	Mismatch risk	105
	Transfer printing	~1–5 mW	<1 MHz	High-yield bonded microlasers	Lab-level promising	Flexible, tolerant	Continuous-wave output thermally limited (adhesive)	106
	Direct epitaxy	<1 mW	~1–5 MHz	Quantum dots on silicon; nano-ridge CMOS	Early stage, defect limited	–	Drift risk	107,108
		Bandwidth	**Energy per bit**					
Modulator	Silicon carrier depletion	>50 GHz	Multi-picojoule	Slow-light silicon modulators	CMOS foundry ready	Edge/grating compatible	Tuning overhead	109
	Thin-film LN	>100 GHz	Subpicojoule	Hybrid LN/Si modulators	Wafer bonding maturing	Fibre attach challenging	Mismatch risk	44,110
	Plasmonic/ ferroelectric	>100 GHz	<1 fJ per bit	BTO-on-silicon, plasmonic modulators	Early stage	On-chip integration	Stability issues	111,112
		Bandwidth	**Responsivity**					
Detector	Germanium-on-silicon photodetector	265 GHz	0.3 A W^{-1}	CMOS-compatible germanium photodetectors	Wafer scale	Butt/edge coupling	Absorption trade-off	51
	Silicon APD	~40 GHz	~0.4 A W^{-1}	Wafer-scale APD arrays	Wafer scale	Edge coupling	Excess noise	113
		Propagation loss	**Footprint**					
Waveguide	Silica	<0.1 dB m^{-1}	Low density	Planar lightwave circuit platforms	Mature	Fibre coupling matched	Large footprint	114
	Silicon-on-insulator	1–3 dB cm^{-1}	High density	CMOS roadmap	CMOS foundry	Edge/grating coupling	Tuning overhead	115
	SiN	~0.1 dB cm^{-1}	High density	Ultralow-loss circuits	Back-end compatible	Subdecibel edge coupling	Heater drift	116
		Insertion loss	**Alignment**					
Coupler	Edge	<1 dB	±2 µm	SiN edge couplers	CMOS packaging	High yield	Alignment sensitive	117
	Grating	3–6 dB	±1 µm	Silicon photonics gratings	Wafer-scale integration	Scalable	Polarization dependence	118
	Freeform	1–2 dB	–	Freeform coupling	Early stage	Flexible	Process complexity	119

parasitic capacitance and inductance that degrade signal integrity and increase energy consumption. Frequency-dependent wire-bond impedance distortion and capacitive loading limit achievable frequency to below 30 GHz in high-speed applications[63]. Moreover, I/O density at the die periphery restricts the number of available channels and routing flexibility[22].

To address these constraints, 2.5D integration introduces an interposer based on silicon or glass that provides short and dense electrical and optical interconnections between heterogeneous chiplets while maintaining modularity. Flip-chip bonding with micro-bumps (5–40 µm pitch) and TSVs (0.2–5.0 µm diameter) enables low-loss signal transmission and reduced power dissipation. TSVs, in particular, offer low-resistance vertical connections with minimized inductance and reduced skin-effect losses, supporting transmission beyond 100 Gbps. By leveraging established foundry processes, 2.5D integration achieves superior bandwidth density and lower insertion loss compared with 2D configurations, offering a practical near-term solution for scalable CPO deployment[22]. Within silicon-based approaches, chip-on-wafer-on-substrate uses large silicon interposers to interconnect chiplets and HBM, delivering exceptional data throughput. However, this approach necessitates substantial manufacturing costs and incurs increased thermal management complexity[64]. In contrast, embedded multi-die interconnect bridge architectures embed localized silicon bridges only where high-speed connectivity is required, reducing costs and improving design flexibility while preserving performance[65].

Glass interposers are emerging as a compelling alternative for 2.5D integration. Their intrinsic optical transparency allows direct embedding of waveguides, gratings and couplers within the substrate, eliminating separate cladding layers and simplifying assembly. Glass interposers exhibit propagation losses below 0.05 dB cm^{-1}, an order of magnitude lower than silicon. Their low dielectric constant minimizes capacitance loading and RC delays, while superior thermal isolation mitigates photonic crosstalk in densely integrated systems[66]. However, challenges persist, including managing thermal expansion mismatch, ensuring mechanical stability, and achieving submicrometre

https://doi.org/10.1038/s41928-026-01681-6

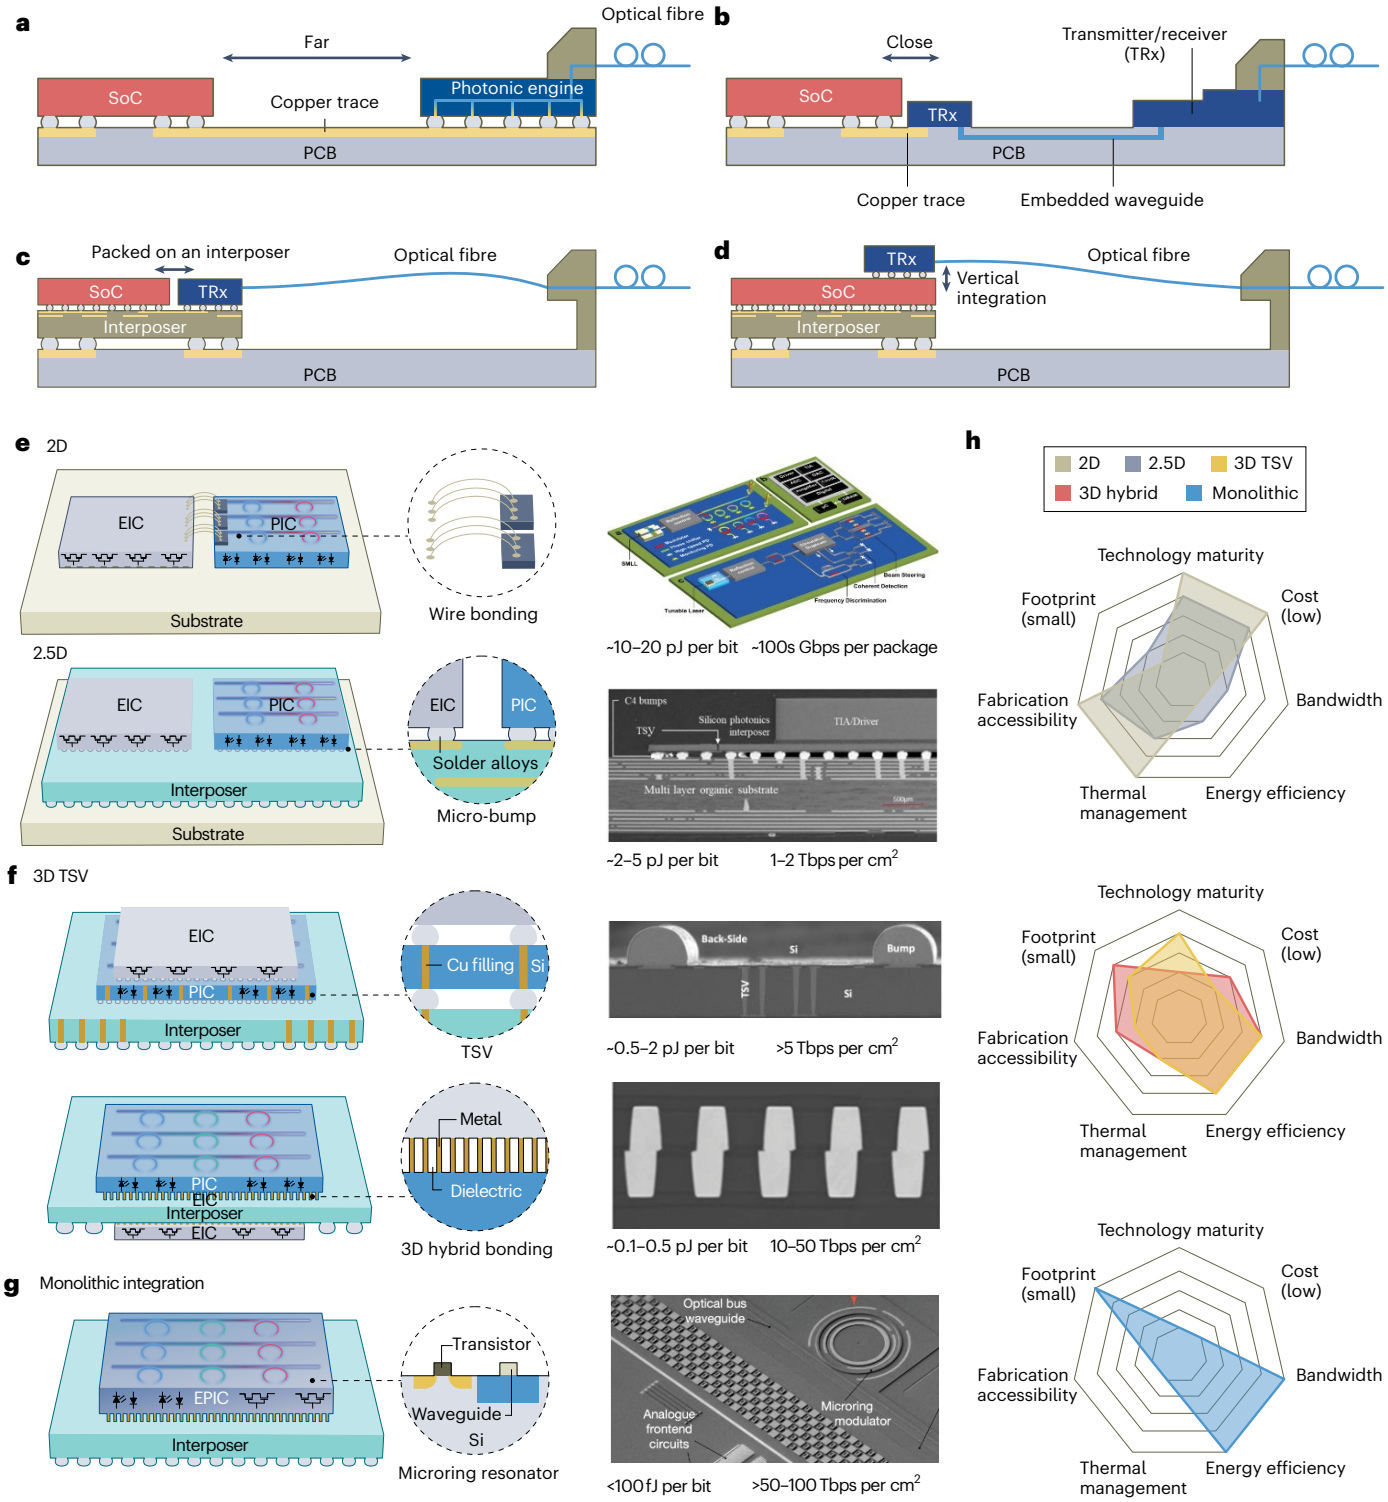

Fig. 3 | Co-packaging architectures for electronic–photonic integration. **a**, Conventional pluggable optics in which interchangeable transceiver modules interface with computing units through copper traces on printed circuit boards (PCBs). **b**, On-board optics that position optical transceivers (TRx) adjacent to EICs on a shared substrate, thereby reducing interconnect length and parasitic losses. **c**, CPO where photonic and electronic dies are integrated on a common interposer to minimize latency, signal loss and power consumption. **d**, Heterogeneous integration platforms using vertical assembly and interposer-level routing to achieve high-density electro–optical interconnects. **e**–**g**, Evolution of integration strategies from 2D to fully monolithic architectures. **e**, Two-dimensional side-by-side die placement and 2.5D interposer-based interconnects. **f**, Three-dimensional stacking enabled by TSVs and Cu–Cu hybrid bonding. **g**, Monolithic co-fabrication of EICs and PICs on a common substrate. Cross-sectional illustrations (left) and scanning electron microscopy images (right) show each integration tier, annotated with representative bandwidth density and energy-per-bit metrics. **h**, Multidimensional comparison of integration strategies across technological maturity, footprint, fabrication accessibility, thermal management, energy efficiency, achievable bandwidth and cost, illustrating the inherent trade-offs among system-level scalability, complexity and performance. SoC, system-on-chip. Panels reproduced from: **e** (top right), ref. 91 under a Creative Commons licence CC BY 4.0; **e** (bottom right), ref. 124 under a Creative Commons licence CC BY NC ND. Panels adapted with permission from: **f** (top right), ref. 125, IEEE; **f** (bottom right), ref. 126, IEEE; **g** (right), ref. 127, Springer Nature Limited.

patterning precision across large areas. Realizing consistent dimensional precision in embedded optical features will require advances in high-resolution lithography, etching and alignment control on glass substrates[67].

Recent progress has substantially narrowed the gap between glass and silicon interposers for 2.5D CPO, but three challenges require quantitative assessment. The first is the coefficient of thermal expansion (CTE) mismatch. Standard soda-lime glass at ~9 ppm K^{-1} is incompatible with silicon (~2.6 ppm K^{-1}), but display-grade borosilicate substrates exhibit CTE values of 3.2–3.6 ppm K^{-1} and are now used as 2.5D interposers, reducing die–substrate shear stress by more than threefold relative to soda-lime platforms[66]. The residual mismatch is accommodated through system-level thermal design, including a liquid and two-phase cooling loop that limits ΔT-driven strain at the die–interposer interface[67]. The second is mechanical stability and through-glass-via (TGV) fabrication. Laser-assisted selective etching enables wafer-scale TGVs with diameter below 30 µm, aspect ratio exceeding 5:1 and pitches approaching 100 µm, with demonstrated flip-chip integration of high-speed drivers and electro-absorption modulated lasers on glass at >50 GHz operating bandwidth[68,69]. Multilayer redistribution remains constrained by panel-level warpage, which scales with metallization thickness and asymmetric stress. Current practice mitigates this through symmetric stack-up and intermediate annealing steps[70]. The third is submicrometre patterning over a large area. Stepper-based lithography on 300 mm glass wafers has demonstrated ~0.5 µm feature resolution for embedded waveguides and coupling structures while panel-scale platforms remain in the 1–2 µm regime and are currently limited to lower-density optical layers[71]. Active-device integration introduces additional thermal and mechanical constraints[72], and recent 102.4 Tbps data-centre switch demonstrations on glass substrates confirm the platform's system-level viability at scale[73]. The remaining barriers to industrial deployment are therefore yield, long-term reliability under thermal cycling and standardized process design kits for active-on-glass integration.

3D integration

Three-dimensional integration extends CPO beyond planar and interposer-based designs towards true vertically stacked architectures (Fig. 3f). Two building blocks support these stacks. TSVs provide vertical electrical routing through silicon dies to expose signals and power on their back sides, and die-to-die bonding forms the physical and electrical interface between stacked dies. TSVs with diameters of 1–10 µm are now routinely fabricated in advanced packaging flows and are indispensable for any 3D stack that requires back-side I/O, including HBM and photonic–electronic co-stacking.

The choice of bonding interface, however, sets the upper bound on 3D stack bandwidth density and energy efficiency. Conventional stacks use solder micro-bumps, which limit die-to-die pitch to typically >10 µm and introduce interface parasitic capacitance on the order of tens of femtofarads per bump[74]. Hybrid bonding replaces the micro-bump stack with direct Cu–Cu contacts embedded in a dielectric interface[75], reducing the interconnect length to tens of micrometres and enabling submicrometre pitch; wafer-to-wafer demonstrations have reached ~400 nm, more than an order of magnitude below micro-bump pitches. This bump-less interface yields higher bandwidth density through pitch scaling and superior energy efficiency by eliminating the micro-bump parasitic capacitance that dominates CV^2 switching energy at the die boundary (where C is capacitance and V is voltage). Combined with TSV-based vertical routing, hybrid bonding therefore enables ultrashort and low-parasitic interconnects that support >5 Tbps mm^{-2} bandwidth densities in 3D photonic–electronic stacks. Such ultrashort interconnects through vias minimize high-frequency signal degradation, eliminate the need for retimers and complex DSP circuits, and thereby substantially improve overall power efficiency[76]. As photonic device performance is often constrained by the parasitic capacitance of their electronic driver and receiver circuits, minimizing the interconnect length directly enhances modulation bandwidth, photodetection sensitivity and system-level signal-to-noise ratio.

Three-dimensional architectures also enable unique spatial configurations in which detectors, modulators and logic blocks are vertically aligned above photonic waveguides, which enables unprecedented interconnect density and latency approaching subpicosecond scales. Early 3D photonic–electronic stacks have achieved bandwidth densities exceeding 5 Tb s^{-1} mm^{-2}, outperforming state-of-the-art 2.5D architectures[77]. Nevertheless, challenges persist in thermal management, bonding yield and process uniformity. Localized self-heating, thermal stress and nanoscale alignment tolerances remain limiting factors for reliable scaling of vertically integrated photonic–electronic systems.

Monolithic integration

Monolithic integration represents the ultimate convergence of electronic and photonic functionality in which all active and passive components including waveguides, modulators, detectors and transistors are fabricated within a unified process flow on a single substrate (Fig. 3g). By eliminating inter-die interfaces and post-assembly alignment, monolithic integration minimizes parasitic losses, maximizes integration density and ensures precise temporal synchronization between optical and electronic operations. The absence of bonding interfaces and micro-bumps suppresses impedance mismatch and reflection losses across the optoelectronic interface, resulting in seamless signal transfer between optical and electronic domains.

Among existing technologies, CMOS-compatible silicon photonics have achieved the highest process maturity, supporting co-fabrication of germanium photodetectors and high-speed silicon modulators alongside logic transistors. Beyond silicon, group-IV alloys and 2D semiconductors are being investigated to extend functionality towards nonlinear optics, quantum photonics and CMOS-compatible light sources[78]. However, material incompatibility, thermal constraints and process co-optimization continue to limit yield and increase complexity.

Each integration approach, including 2D, 2.5D, 3D and monolithic, exhibits distinct trade-offs in maturity, performance, cost and scalability (Fig. 3h). Two-dimensional packaging offers simplicity and low cost but with limited bandwidth. 2.5D interposers offer modest density, manufacturability and cost, while 3D architecture pushes the limits of bandwidth density and energy efficiency at the expense of yield and thermal complexity. Monolithic integration promises the highest performance and energy efficiency but is contingent on breakthroughs in material integration and process co-optimization that can unify photonic and electronic systems within a seamless fabrication ecosystem. These architectural trade-offs are reflected in deployed industrial platforms, which Table 2 summarizes across bandwidth, latency, energy per bit and integration scheme.

Challenges

Despite substantial advances in individual photonic and electronic components, the large-scale deployment of OCIs hinges on cohesive system integration rather than isolated device breakthroughs. Record performance in bandwidth, energy efficiency and latency has yet to translate into manufacturable, reliable and interoperable systems suitable for industrial use. Realizing this transition requires coordinated progress across multiple fronts, including co-design of memory, logic and photonic subsystems, preservation of signal integrity and thermal stability under dense integration while minimizing power consumption, and the establishment of cost-efficient high-yield manufacturing and packaging strategies. Equally essential is the development of unified design frameworks and standardized process interfaces that can bridge device-level innovations with system-level architectures, ensuring cross-foundry interoperability and accelerating ecosystem-wide adoption.

Table 2 | Representative industrial CPO and optical I/O platforms

Company	Product/platform	Architecture	Bandwidth	Latency	Energy	Integration	Reach	Reference
Broadcom	Bailly (TH5)	Switch CPO + ASIC	51.2 Tbps	Sub-10 ns per hop	~6.9 pJ per bit (5.5 W/800G port)	2.5D co-packaged	Rack scale	[a]
Broadcom	Davisson (TH6)	Switch CPO + ASIC	102.4 Tbps	Sub-10 ns per hop	~5–7 pJ per bit (projected)	2.5D co-packaged	Rack scale	[b]
NVIDIA	Quantum-X Photonics	Switch CPO (COUPE 3D SiP)	115.2 Tbps	Sub-10 ns per hop	~5.6 pJ per bit (9 W/1.6T port)	3D co-packaged	Rack scale	120
Ayar Labs	TeraPHY	Optical I/O chiplet (UCIe)	8 Tbps (bidirectional)	~10 ns chiplet-hop	<5 pJ per bit	Chiplet (SiP)	mm–km	121
Intel	OCI	Optical I/O chiplet (PIC+EIC)	~4 Tbps (bidirectional)	<10 ns + time of flight	~5 pJ per bit (end to end)	Co-packaged + laser	Rack scale	23
Ranovus	Scale-up engines	Optical engines (CPO)	Multi Tbps	Low	Low (laser)	Co-packaged	Rack scale	122
Lightmatter	Passage M1000	3D photonic interposer	114 Tbps (4,000 mm^2)	Sub-10 ns	4.6 pJ per bit wall-plug	Chip-on-wafer (3D)	10 m–km	25
Celestial AI	Photonic Fabric (PFLink)	Chiplet + system fabric	14.4 Tbps	Sub-10 ns (reported)	–	Chiplet + SiP	>50 m	123
Avicena	LightBundle	Micro-LED parallel transceiver	~1 Tbps (304 channel × 3.3 Gbps)	Micro-LED native	~1 pJ per bit	CMOS (16 nm) + micro-LED array	<50 m	100
Microsoft	MOSAIC	Micro-LED parallel transceiver	~1 Tbps (460 channel × 2 Gbps)	Few ns (no DSP/FEC)	~6.6 pJ per bit (5.3 per end)	CMOS + micro-LED array (pluggable)	<50 m	101

[a]Broadcom Inc. Bailly (Tomahawk 5) 51.2-Tbps CPO ethernet switch platform. Product announcement, March 2024. [b]Broadcom Inc. Davisson (Tomahawk 6) 102.4-Tbps CPO ethernet switch platform. Product announcement, October 2025.

Signal integrity, thermal stability and power consumption

The performance and reliability of OCIs are governed by three tightly coupled factors: signal integrity, thermal stability and power consumption. Even modest degradation in signal quality necessitates sophisticated DSP and FEC, which introduce latency, area and power overheads that erode the intrinsic advantage of optical links[79]. SerDes circuits remain among the most power-intensive components, and as data rates scale into the multi-terabits-per-second regime, the complexity of DSP and FEC scales rapidly. For instance, coherent equalization stages scale nearly with the number of taps, with power consumption increasing from 1 pJ per bit for a 4-tap configuration to over 5 pJ per bit for 16-tap configurations[80]. While modest per channel, the cumulative DSP overheads become prohibitive in multi-terabit systems, diminishing the efficiency of interconnects.

Thermal instability further compounds these challenges. Temperature-induced spectral drift is one of the most persistent barriers to stable operation. Semiconductor laser emission wavelengths typically shift by tens-of-picometres per kelvin[81], severely affecting alignment in WDM systems. This effect is particularly pronounced in CPO, where lasers operate within the same thermal envelope as ASICs or accelerators. Local temperature fluctuations reduce quantum efficiency and drive wavelength drift, creating a feedback loop that destabilizes both optical power and spectral coherence[82]. Microring modulators, resonators and WDM filters also exhibit very high temperature sensitivities; typical silicon photonic microrings shift by tens of gigahertz per kelvin[83]. Thermal gradients also produce local variations in resistance and capacitance that distort the timing and signal integrity. Non-uniform RC delays in the driver and receiver circuits introduce jitter and synchronization errors, increasing bit-error rate and compromising overall system reliability, particularly in hyperscale data centres[84]. At terabits-per-second data rates, such impairments are compounded by parasitic phenomena including crosstalk, reflections and chromatic dispersion.

In parallel, the total power consumption of CPO systems continues to grow, driven not only by optical sources but also by the control logic, transceiver DSP and thermal regulation subsystems. This increasing energy footprint creates an adverse thermal–power feedback loop. As control and cooling units consume more power, additional heat is generated, which in turn demands more aggressive active cooling power. Such thermal–power coupling ultimately limits scalability, as system-level energy efficiency degrades faster than link bandwidth scales.

Addressing these interdependent challenges requires innovations across materials, devices, packaging and systems. Advanced thermal management techniques, such as microfluidic cooling, localized thermoelectric coolers and high thermal-conductivity interposers, have shown effective heat dissipation under high-density packaging[85–87]. Emerging thermal photonic approaches and metastructured heat spreaders provide additional control over nanoscale heat transport[88,89]. Athermal resonator designs based on SiN or TFLN reduce temperature-induced refractive index drift[90], while polarization-insensitive modulators and distributed monitoring circuits enable adaptive compensation for link degradation. However, embedding these features without incurring prohibitive penalties in cost, area or power remains an ongoing challenge. Ultimately, stable OCI operation will require co-design strategies that integrate device engineering, packaging thermomechanics and adaptive control algorithms, to preserve link fidelity and predictable performance across large-scale deployments.

Barriers to scalable industrialization

The industrial adoption of OCI depends critically on scalable manufacturing processes that ensure high yield, low cost and robust long-term reliability. Although recent demonstrations of CPO validate the technical feasibility of tightly integrated electronic–photonic subsystems, they also revealed the gap between prototype and industrial

https://doi.org/10.1038/s41928-026-01681-6

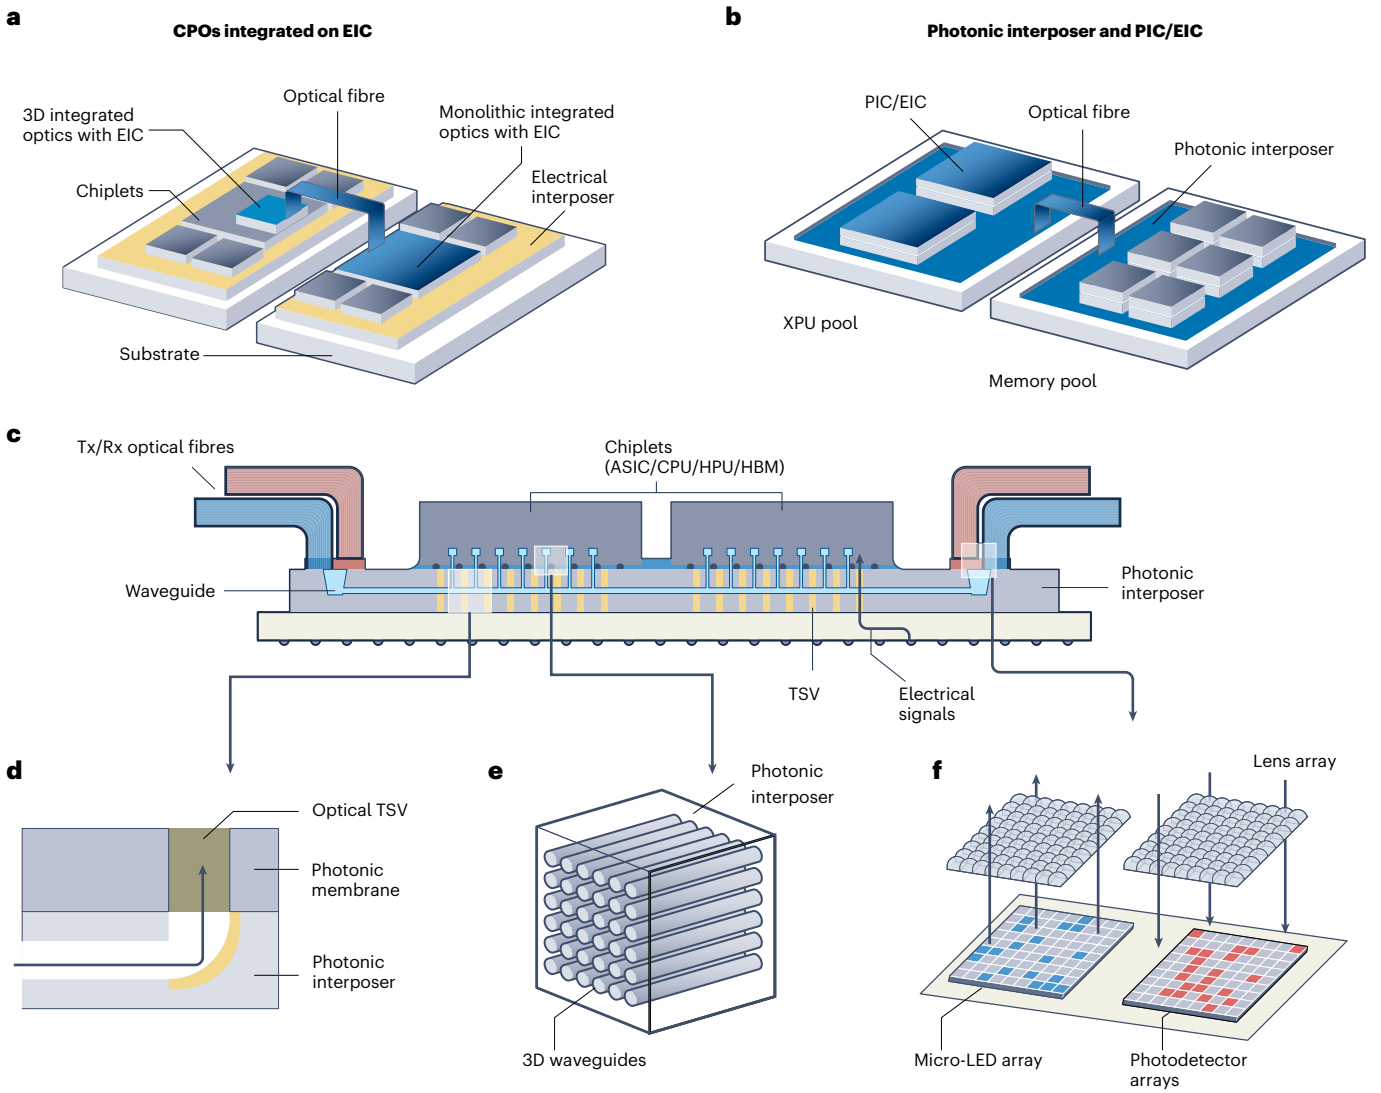

Fig. 4 | CPO and photonic-interposer technologies for high-bandwidth OCIs. **a**, Co-packaged architectures in which photonic components, implemented through advanced 3D integration or monolithic co-fabrication, are tightly integrated with electronic logic dies on a shared interposer. **b**, Photonic interposer-based networks that provide chip-to-chip and node-to-node connectivity, linking compute and memory chiplets across intra- and inter-node clusters. **c**, Cross-sectional schematic of a multi-chip module, comprising chiplets such as ASICs, CPUs, high-bandwidth processing units (HPUs) and HBM, interconnected through a photonic interposer embedding optical waveguides for low-loss and high-bandwidth chip-to-chip communication. **d**, Optical TSVs enabling vertical routing of optical signals between stacked PICs and EICs. **e**, Three-dimensional waveguide networks embedded within photonic interposers, supporting high-density wavelength-multiplexed optical channels with minimal crosstalk and scalable spatial routing. **f**, Massively parallel optical interconnects using microlens arrays to couple emission from micro-LED transmitters to photodetector arrays that enable optical communications without reliance on SerDes.

manufacturability. A primary barrier is the integration of optical gain blocks, such as lasers and semiconductor optical amplifiers within silicon-based platforms. External laser modules remain a practical interim solution, but are limited by large form factors, stringent alignment tolerance and high coupling losses[91]. Wafer- or die-scale bonding of III–V gain media onto silicon has emerged as the most mature integration route towards on-chip light sources; however, yield loss, defect density and process complexity continue to impede scalability. Micro-transfer printing offers a complementary heterogeneous integration strategy with improved material utilization and the prospect of parallel, wafer-scale assembly of non-native optoelectronic components. The INSPIRE project exemplifies this direction[92], demonstrating InP-on-SiN photonic integration through wafer-scale micro-transfer printing and thereby combining the optical gain of III–V materials with the low propagation loss of SiN passive circuits. The approach extends beyond light sources. InGaAs/InP photodiodes have been transfer-printed onto SiN photonic platforms, confirming that transfer printing supports receiver-side integration on low-loss passive photonic platforms[93].

Packaging complexity introduces further barriers. Efficient optical coupling demands submicrometre alignment accuracy that must be maintained during high-throughput assembly. Advances such as passive wafer-level alignment, beam-shaping couplers and inverse-designed photonic interfaces have expanded coupling tolerances and improved insertion efficiency. However, repeatable, low-loss, high-throughput packaging solutions remain elusive. In addition, mechanical and thermal stresses arising from CTE mismatches can lead to delamination and long-term performance drift. The reliability of densely integrated modules under continuous multi-terabits-per-second traffic and elevated thermal loads remains insufficiently demonstrated. Ensuring industrial-grade robustness will require accelerated lifetime testing, a redundancy mechanism and system-level fault-tolerance strategies.

Beyond technical challenges, ecosystem fragmentation remains a substantial obstacle. Photonic foundry platforms differ widely in material sets, device libraries, process parameters and packaging compatibility. This heterogeneity complicates multi-sourcing and inhibits interoperability. Furthermore, photonic process rules often diverge from established CMOS manufacturing standards, impeding seamless electronic–photonic co-integration. This contrasts with the EIC domain, where standardized process nodes, unified electronic design automation tools and long-term roadmaps have enabled predictable scaling. The current lack of convergence across wafer-level packaging standards, electronic–photonic design automation and reliability qualification protocols severely delays the transition of CPO from prototypes to the foundational infrastructures of HPC systems[94]. Establishing such a unified ecosystem will be crucial for the widespread deployment of OCIs into HPC and AI platforms.

Outlook

The future development of CPO will require a series of architectural and material innovations. In the near term, progress in OCIs will be driven by chiplet-based architectures that integrate heterogeneous photonic and electrical components via advanced packaging technologies (Fig. 4a). These architectures assemble central processing units (CPUs), GPUs, AI accelerators and HBM onto shared interposer platforms with dense optical I/O interfaces. Transferable membrane-based photonic components, typically a few hundred nanometres to several micrometres thick, offer a promising route to overcome the lattice mismatch and thermal incompatibility issues that have prevented monolithic growth of these materials directly on silicon. Heterogeneous integration of single-crystal III–V semiconductors, BTO and LN membranes via remote epitaxy and layer transfer techniques exemplifies this trajectory[95,96]. Their high crystallinity, strong optical confinement and low thermal mass can provide a large electro–optic coefficient, enhanced modulation bandwidth and reduced switching energy. Recent demonstrations of wafer-scale fabrication and transfer-printed membrane photonic devices using CMOS-compatible process flows highlight their potential for scalable and manufacturable OCI deployment[97–99].

The longer-term trajectory of OCI points towards optics-centric architectures that reshape how information is transmitted, routed and processed (Fig. 4b). Photonic interposers (Fig. 4c) provide the structural foundation for this transition by embedding low-loss waveguides that route optical signals between chiplets at the intra-package scale. Fibre array units provide both on- and off-board optical connectivity for high-capacity inter-package routing and signal distribution. Through-silicon optical vias (Fig. 4d) expand this architecture to the vertical dimension, incorporating photonic components that enable low-loss optical communication between stacked layers, overcoming the bandwidth density and thermal limits of conventional electrical TSVs. However, realizing this at scale will require addressing challenges in optical alignment, loss control, thermal management and manufacturing yield.

The emergence of volumetric optical architectures, such as multilayer waveguide networks embedded within photonic interposers (Fig. 4e), could lead to ultra-dense routing topologies capable of supporting thousands of wavelength channels with minimal crosstalk and loss. Alternatively, SerDes-free optical interconnect arrays based on micro-light-emitting diode (micro-LED) and photodetector arrays integrated with microlens systems (Fig. 4f) are a complementary route towards massively parallel and reconfigurable chip-to-chip communication. Recent demonstrations have quantified this trade-off. Avicena's LightBundle transceiver integrated circuit integrates a 304-channel micro-LED array on a 16 nm CMOS chip, delivering around 1 Tbps aggregate throughput at 3.3 Gbps per lane and 1 pJ per bit (ref. 100), and Microsoft's MOSAIC platform offers a 100-channel prototype operating error free at 2 Gbps per lane over 20 m and predicts scaling to 800 Gbps via 460-channel arrays[101].

At the link level, an 800 Gbps micro-LED-based transceiver is projected to consume 3.1–5.3 W per end, a 56–68% reduction relative to mainstream 800 Gbps laser-based active optical cables (9.8–12 W per end), and further scaling to 1.6–3.2 Tbps could be achieved through denser arrays and 4–8 Gbps per lane operation. These advantages are offset by chromatic dispersion from the broader micro-LED emission spectrum (>10 nm versus <1 pm for lasers), Lambertian beam profiles that demand custom micro-optical coupling and reach limited to tens of metres. However, these constraints are largely addressed within co-packaged architectures where the short chip-to-chip distances match the micro-LED reach limit and the native modulation speed of micro-LEDs eliminates the high-speed electrical conversion required by laser-based links.

Such systems could ultimately achieve multi-terabits-per-second inter-chip bandwidths and femtojoules-per-bit energy efficiencies, which are performance metrics unattainable with conventional architectures[102,103]. Beyond high-bandwidth communication, OCI architectures could enable co-optimization of memory hierarchies, logic circuits and photonic fabrics, supporting emerging technologies such as in-network computing and photonic memory acceleration[104]. Achieving this vision will, however, require breakthroughs in three fundamental domains, which are optical memory interfaces capable of preserving coherence and minimizing latency, low-power and high-speed photonic switching for wavelength- and space-division multiplexing, and distributed optical signal regeneration schemes that maintain fidelity across multi-hop optical networks. Progress across these areas, combined with advances in materials, heterogeneous integration and system-level co-design, could transform photonics from a communication medium into a core computational fabric for data-centric and HPC systems.

References

1. Xie, M. et al. Monolithic 3D integration of 2D transistors and vertical RRAMs in 1T–4R structure for high-density memory. *Nat. Commun.* **14**, 5952 (2023).
2. Meng, L. et al. Dual-gated single-molecule field-effect transistors beyond Moore's law. *Nat. Commun.* **13**, 1410 (2022).
3. Li, Q. et al. 3D stacked IGZO 2T0C DRAM array with multibit capability for computing in memory applications. *Sci. Adv.* **11**, eadu4323 (2025).
4. Datta, S. et al. Toward attojoule switching energy in logic transistors. *Science* **378**, 733–740 (2022).
5. Salahuddin, S. et al. The era of hyper-scaling in electronics. *Nat. Electron.* **1**, 442–450 (2018).
6. Shulaker, M. M. et al. Three-dimensional integration of nanotechnologies for computing and data storage on a single chip. *Nature* **547**, 74–78 (2017).
7. Borsoi, F. et al. Shared control of a 16 semiconductor quantum dot crossbar array. *Nat. Nanotechnol.* **19**, 21–27 (2024).
8. Gholami, A. et al. AI and memory wall. *IEEE Micro* **44**, 33–39 (2024).
9. Alloatti, L. et al. 100 GHz silicon-organic hybrid modulator. *Light Sci. Appl.* **3**, e173 (2014).
10. Li, K. et al. An integrated CMOS–silicon photonics transmitter with a 112 gigabaud transmission and picojoule per bit energy efficiency. *Nat. Electron.* **6**, 910–921 (2023).
11. Hu, Y. et al. III/V-on-Si MQW lasers by using a novel photonic integration method of regrowth on a bonding template. *Light Sci. Appl.* **8**, 93 (2019).
12. Ahmed, S. R. et al. Universal photonic artificial intelligence acceleration. *Nature* **640**, 368–374 (2025).
13. Margalit, N. et al. Perspective on the future of silicon photonics and electronics. *Appl. Phys. Lett.* **118**, 220501 (2021).
14. Zhang, Y. et al. Scalable 3D silicon photonic electronic integrated circuits and their applications. *IEEE J. Sel. Top. Quantum Electron.* **26**, 1–10 (2020).

15. Zou, K. et al. High-capacity free-space optical communications using wavelength- and mode-division-multiplexing in the mid-infrared region. *Nat. Commun.* **13**, 7662 (2022).
16. Kim, C. et al. Robust Co alloy design for Co interconnects using a self-forming barrier layer. *Sci. Rep.* **12**, 12291 (2022).
17. Deutsch, A. et al. Frequency-dependent losses on high-performance interconnections. *IEEE Trans. Electromagn. Compat.* **43**, 446–465 (2001).
18. Williamson, I. A. D. et al. Suppression of the skin effect in radio frequency transmission lines via gridded conductor fibers. *Appl. Phys. Lett.* **108**, 083502 (2016).
19. Wang, T.-Y. et al. Polymeric insulating materials characteristics for high-voltage applications. *Nat. Rev. Electr. Eng.* **1**, 516–528 (2024).
20. Wang, W. et al. Rail-only: a low-cost high-performance network for training LLMs with trillion parameters. In *2024 IEEE Symposium on High-Performance Interconnects (HOTI)* 1–10 (IEEE, 2024).
21. Tian, W. et al. Progress in research on co-packaged optics. *Micromachines* **15**, 1211 (2024).
22. Priyadarshi, S. Advancing AI scalability and performance with optical interconnects. *IEEE Commun. Mag.* **63**, 22–27 (2025).
23. Fathololoumi, S. 4 Tb/s optical compute interconnect chiplet for XPU-to-XPU connectivity. In *2024 IEEE Hot Chips 36 Symposium (HCS)* 1–18 (IEEE, 2024).
24. Lou, L. et al. From silicon to system: recent advances in wireline transceiver ICs and integration applications for next-generation AI infrastructure. In *2025 IEEE 8th International Conference on Integrated Circuits, Technologies and Applications (ICTA)* 151–157 (IEEE, 2025).
25. Lightmatter announces Passage L200, the fastest co-packaged optics for AI. Lightmatter https://lightmatter.co/press-release/lightmatter-announces-passage-l200-the-fastest-co-packaged-optics-for-ai/ (31 March 2025).
26. Ahmed, K. M. U. et al. A review of data centers energy consumption and reliability modeling. *IEEE Access* **9**, 152536–152563 (2021).
27. Ye, B. et al. A 2.29-pJ/b 112-Gb/s wireline transceiver with RX four-tap FFE for medium-reach applications in 28-nm CMOS. *IEEE J. Solid State Circuits* **58**, 19–29 (2023).
28. Das Sharma, D. et al. Universal Chiplet Interconnect Express (UCIe): an open industry standard for innovations with chiplets at package level. *IEEE Trans. Compon. Packaging Manuf. Technol.* **12**, 1423–1431 (2022).
29. Das, A. et al. Chip and package-scale interconnects for general-purpose, domain-specific, and quantum computing systems—overview, challenges, and opportunities. *IEEE J. Emerg. Sel. Top. Circuits Syst.* **14**, 354–370 (2024).
30. Wang, X. et al. Perspectives of active Si photonics devices for data communication and optical sensing. *J. Appl. Phys.* **138**, 060901 (2025).
31. Buscaino, B. et al. External vs. integrated light sources for intra-data center co-packaged optical interfaces. *J. Lightwave Technol.* **39**, 1984–1996 (2021).
32. Tan, M. et al. Co-packaged optics (CPO): status, challenges, and solutions. *Front. Optoelectron.* **16**, 1 (2023).
33. Koscica, R. et al. Heterogeneous integration of a III–V quantum dot laser on high thermal conductivity silicon carbide. *Opt. Lett.* **48**, 2539–2542 (2023).
34. Coenen, D. et al. Thermal scaling analysis of large hybrid laser arrays for co-packaged optics. *IEEE J. Sel. Top. Quantum Electron.* **31**, 1501009 (2025).
35. Ghosh, S. et al. Scalable transfer printing approach to heterogeneous integration of InP lasers on silicon-on-insulator waveguide platform. *Appl. Phys. Lett.* **125**, 081104 (2024).
36. Hughes, E. T. et al. Gradual degradation in InAs quantum dot lasers on Si and GaAs. *Nanoscale* **16**, 2966–2973 (2024).
37. Shang, C. et al. High-temperature reliable quantum-dot lasers on Si with misfit and threading dislocation filters. *Optica* **8**, 749–754 (2021).
38. Shi, Y. et al. Exploring the feedback limits of quantum dot lasers for isolator-free photonic integrated circuits. *Light Sci. Appl.* **15**, 96 (2026).
39. Geler-Kremer, J. et al. A ferroelectric multilevel non-volatile photonic phase shifter. *Nat. Photon.* **16**, 491–497 (2022).
40. Chelladurai, D. et al. Barium titanate and lithium niobate permittivity and Pockels coefficients from megahertz to sub-terahertz frequencies. *Nat. Mater.* **24**, 868–875 (2025).
41. Reynaud, M. et al. Enhancement of electro-optic response in $BaTiO_3$ films integrated on Si via heating and cooling rate control. *Commun. Mater.* **6**, 176 (2025).
42. Xu, Z. et al. Ferroelectric-based Pockels photonic memory. *Nat. Commun.* **16**, 8329 (2025).
43. Hu, Y. et al. Integrated electro-optics on thin-film lithium niobate. *Nat. Rev. Phys.* **7**, 237–254 (2025).
44. Wang, C. et al. Integrated lithium niobate electro-optic modulators operating at CMOS-compatible voltages. *Nature* **562**, 101–104 (2018).
45. Weigel, P. O. et al. Bonded thin film lithium niobate modulator on a silicon photonics platform exceeding 100 GHz 3-dB electrical modulation bandwidth. *Opt. Express* **26**, 23728–23739 (2018).
46. Eltes, F. et al. A $BaTiO_3$-based electro-optic pockels modulator monolithically integrated on an advanced silicon photonics platform. *J. Lightwave Technol.* **37**, 1456–1462 (2019).
47. Guo, W. et al. Epitaxial integration of $BaTiO_3$ on Si for electro-optic applications. *J. Vac. Sci. Technol. A* **39**, 030804 (2021).
48. Datta, I. et al. Low-loss composite photonic platform based on 2D semiconductor monolayers. *Nat. Photon.* **14**, 256–262 (2020).
49. Phare, C. T. et al. Graphene electro-optic modulator with 30 GHz bandwidth. *Nat. Photon.* **9**, 511–514 (2015).
50. Michel, J. et al. High-performance Ge-on-Si photodetectors. *Nat. Photon.* **4**, 527–534 (2010).
51. Lischke, S. et al. Ultra-fast germanium photodiode with 3-dB bandwidth of 265 GHz. *Nat. Photon.* **15**, 925–931 (2021).
52. Sun, W. et al. Ultralow-LO-power silicon photonic coherent receiver with balanced Ge/Si avalanche photodiodes. *ACS Photon.* **12**, 1793–1801 (2025).
53. Sun, A. et al. Edge-guided inverse design of digital metamaterial-based mode multiplexers for high-capacity multi-dimensional optical interconnect. *Nat. Commun.* **16**, 2372 (2025).
54. Sideris, C. et al. Foundry-fabricated grating coupler demultiplexer inverse-designed via fast integral methods. *Commun. Phys.* **5**, 68 (2022).
55. Yang, K. Y. et al. Multi-dimensional data transmission using inverse-designed silicon photonics and microcombs. *Nat. Commun.* **13**, 7862 (2022).
56. Zhou, X. et al. Silicon photonics for high-speed communications and photonic signal processing. *npj Nanophoton.* **1**, 27 (2024).
57. Komljenovic, T. et al. Heterogeneous silicon photonic integrated circuits. *J. Lightwave Technol.* **34**, 20–35 (2016).
58. Xu, Y. et al. Methods and applications of on-chip beam splitting: a review. *Front. Phys.* **10**, 985208 (2022).
59. Rao, A. et al. Towards integrated photonic interposers for processing octave-spanning microresonator frequency combs. *Light Sci. Appl.* **10**, 109 (2021).
60. Zhang, Y. et al. Low-loss and broadband wafer-scale optical interposers for large-scale heterogeneous integration. *Opt. Express* **32**, 40 (2024).
61. Liang, D. et al. An energy-efficient and bandwidth-scalable DWDM heterogeneous silicon photonics integration platform. *IEEE J. Sel. Top. Quantum Electron.* **28**, 6100819 (2022).

62. Larocque, H. et al. Tunable quantum emitters on large-scale foundry silicon photonics. *Nat. Commun.* **15**, 5781 (2024).
63. Zahran, S. R. et al. Bondwire integration challenges in E-band systems: from PCB to die level. In *2023 53rd EuMC* 846–849 (IEEE, 2023).
64. Lau, J. H. Recent advances and trends in chiplet design and heterogeneous integration packaging. *J. Electron. Packag.* **146**, 010801 (2024).
65. Mahajan, R. et al. Scaling EMIB—a comprehensive architecture portfolio for localized interconnects. *IEEE Trans. Electron Devices* **72**, 3911–3914 (2025).
66. Zhao, J. et al. Development of high performance 2.5D packaging using glass interposer with through glass vias. *J. Mater. Sci. Mater. Electron.* **34**, 1790 (2023).
67. Mirshafiei, M. et al. Glass interposer for short reach optical connectivity. *Opt. Express* **24**, 12375 (2016).
68. Yu, C. et al. Application of through glass via (TGV) technology for sensors manufacturing and packaging. *Sensors* **24**, 171 (2024).
69. Ge, C. et al. High-speed wafer-level TGV interposer for 2.5D CPO. *Opt. Commun.* **579**, 131517 (2025).
70. Li, X. et al. Die-embedded glass interposer with minimum warpage for 5G/6G applications. In *2023 IEEE 73rd Electronic Components and Technology Conference (ECTC)* 2247–2254 (IEEE, 2023).
71. Yoon, S. et al. Stitching-loss-tolerant silicon nitride waveguides on glass for panel-level optical interconnects. *Opt. Lett.* **51**, 1903 (2026).
72. Lee, B. G. et al. Beyond CPO: a motivation and approach for bringing optics onto the silicon interposer. *J. Lightwave Technol.* **41**, 1152–1162 (2023).
73. Yeary, L. et al. Co-packaged optics on glass substrates for 102.4 Tb/s data center switches. In *2023 IEEE Electronic Components and Technology Conference (ECTC)* 224–227 (IEEE, 2023).
74. Stojanovic, V. et al. High-density 3D electronic-photonic integration. In *2015 Fourth Berkeley Symposium on Energy Efficient Electronics Systems (E3S)* 1–2 (IEEE, 2015).
75. Xiang, C. et al. Building 3D integrated circuits with electronics and photonics. *Nat. Electron.* **7**, 422–424 (2024).
76. Chew, S. A. et al. Wafer-to-wafer hybrid bonding at 400-nm interconnect pitch. *Nat. Rev. Electr. Eng.* **1**, 71–72 (2024).
77. Daudlin, S. et al. Three-dimensional photonic integration for ultra-low-energy, high-bandwidth interchip data links. *Nat. Photon.* **19**, 502–509 (2025).
78. You, J. et al. Hybrid/integrated silicon photonics based on 2D materials in optical communication nanosystems. *Laser Photon. Rev.* **14**, 2000239 (2020).
79. Alvarado, A. et al. On the impact of optimal modulation and FEC overhead on future optical networks. *J. Lightwave Technol.* **34**, 2339–2352 (2016).
80. Lundberg, L. et al. Power consumption of a minimal-DSP coherent link with a polarization multiplexed pilot-tone. In *ECOC 2016: 42nd European Conference on Optical Communication* 1–3 (VDE, 2016).
81. Wang, R. et al. Heterogeneously integrated III–V-on-silicon 2.3x μm distributed feedback lasers based on a type-II active region. *Appl. Phys. Lett.* **109**, 221111 (2016).
82. Wu, S. et al. Simulation and experimental investigation of liquid-cooling thermal management for high-bandwidth co-packaged optics. *Front. Optoelectron.* **18**, 11 (2025).
83. Parsons, R. et al. Highly uniform thermally undercut silicon photonic devices in a 300 mm CMOS foundry process. *Sci. Rep.* **15**, 29906 (2025).
84. Chen, X. et al. Improving the thermal reliability of photonic chiplets on multicore processors. *Integration* **86**, 9–21 (2022).
85. Lin, C. et al. Miniaturized Mg_3Bi_2-based thermoelectric cooler for localized electronic thermal management. *Nat. Commun.* **16**, 7779 (2025).
86. Zhong, Y. et al. Heterogeneous integration of diamond-on-chip-on-glass interposer for efficient thermal management. *IEEE Electron Device Lett.* **45**, 448–451 (2024).
87. van Erp, R. et al. Co-designing electronics with microfluidics for more sustainable cooling. *Nature* **585**, 211–216 (2020).
88. Li, Y. et al. Transforming heat transfer with thermal metamaterials and devices. *Nat. Rev. Mater.* **6**, 488–507 (2021).
89. Yang, S. et al. Nonreciprocal thermal photonics. *Nat. Photon.* **18**, 412–424 (2024).
90. Ling, J. et al. Athermal lithium niobate microresonator. *Opt. Express* **28**, 21682 (2020).
91. Shekhar, S. et al. Roadmapping the next generation of silicon photonics. *Nat. Commun.* **15**, 751 (2024).
92. Heck, M. J. INSPIRE: InP on SiN photonic integrated circuits realized through wafer-scale micro-transfer printing. In *Integrated Photonics Platforms II Proc. SPIE* 12148 (SPIE, 2022).
93. Qin, S. et al. Micro-transfer printing of O-band InGaAs/InP photodiodes on a silicon nitride photonic platform. *Opt. Express* **33**, 34153–34169 (2025).
94. Ranno, L. et al. Integrated photonics packaging: challenges and opportunities. *ACS Photon.* **9**, 3467–3485 (2022).
95. Kim, H. et al. Remote epitaxy. *Nat. Rev. Methods Primers* **2**, 41 (2022).
96. Meng, Y. et al. Photonic van der Waals integration from 2D materials to 3D nanomembranes. *Nat. Rev. Mater.* **8**, 498–517 (2023).
97. Shim, J. et al. Room-temperature waveguide-integrated photodetector using bolometric effect for mid-infrared spectroscopy applications. *Light Sci. Appl.* **14**, 125 (2025).
98. Shin, J. et al. Vertical full-colour micro-LEDs via 2D materials-based layer transfer. *Nature* **614**, 81–87 (2023).
99. Qi, L. et al. Monolithic full-color active-matrix micro-LED micro-display using InGaN/AlGaInP heterogeneous integration. *Light Sci. Appl.* **12**, 258 (2023).
100. Pezeshki, B. et al. 304 channel microLED based CMOS transceiver IC with aggregate 1 Tbps and sub-pJ per bit capability. In *Proceedings Optical Fiber Communication Conference (OFC) 2024* Technical Digest Series, Paper M3A.1 (Optica Publishing Group, 2024).
101. Benyahya, K. et al. MOSAIC: breaking the optics versus copper trade-off with a wide-and-slow architecture and microLEDs. In *Proceeding of the ACM SIGCOMM 2025 Conference* 234–247 (ACM, 2025).
102. Farmakidis, N. et al. Integrated photonic neuromorphic computing: opportunities and challenges. *Nat. Rev. Electr. Eng.* **1**, 358–373 (2024).
103. Shastri, B. J. et al. Photonics for artificial intelligence and neuromorphic computing. *Nat. Photon.* **15**, 102–114 (2021).
104. Bai, B. et al. Microcomb-based integrated photonic processing unit. *Nat. Commun.* **14**, 66 (2023).
105. Xiang, C. et al. 3D integration enables ultralow-noise isolator-free lasers in silicon photonics. *Nature* **620**, 78–85 (2023).
106. Roelkens, G. et al. Present and future of micro-transfer printing for heterogeneous photonic integrated circuits. *APL Photon.* **9**, 010901 (2024).
107. Dong, B. et al. Turnkey locking of quantum-dot lasers directly grown on Si. *Nat. Photon.* **18**, 669–676 (2024).
108. De Koninck, Y. et al. GaAs nano-ridge laser diodes fully fabricated in a 300-mm CMOS pilot line. *Nature* **637**, 63–69 (2025).
109. Han, C. et al. Slow-light silicon modulator with 110-GHz bandwidth. *Sci. Adv.* **9**, eadi5339 (2023).
110. He, M. et al. High-performance hybrid silicon and lithium niobate Mach–Zehnder modulators for 100 Gbit s^{-1} and beyond. *Nat. Photon.* **13**, 359–364 (2019).

111. Eppenberger, M. et al. Resonant plasmonic micro-racetrack modulators with high bandwidth and high temperature tolerance. *Nat. Photon.* **17**, 360–367 (2023).
112. Abel, S. et al. Large Pockels effect in micro- and nanostructured barium titanate integrated on silicon. *Nat. Mater.* **18**, 42–47 (2019).
113. Peng, Y. et al. An 8 × 160 Gb s^{-1} all-silicon avalanche photodiode chip. *Nat. Photon.* **18**, 928–934 (2024).
114. Lee, H. et al. Ultra-low-loss optical delay line on a silicon chip. *Nat. Commun.* **3**, 867 (2012).
115. Siew, S. Y. et al. Review of silicon photonics technology and platform development. *J. Lightwave Technol.* **39**, 4374–4389 (2021).
116. Liu, J. et al. High-yield, wafer-scale fabrication of ultralow-loss, dispersion-engineered silicon nitride photonic circuits. *Nat. Commun.* **12**, 2236 (2021).
117. Zhu, X. et al. Low-loss fiber-to-chip edge coupler for silicon nitride integrated circuits. *Opt. Express* **31**, 10525–10532 (2023).
118. Thomson, D. et al. Roadmap on silicon photonics. *J. Opt.* **18**, 073003 (2016).
119. Wang, H. et al. Ultralow-loss optical interconnect enabled by topological unidirectional guided resonance. *Sci. Adv.* **10**, eadn4372 (2024).
120. Samanta, A., Lee, S.-H., Cheng, C.-Y., Palermo, S. & Yoo, S. J. B. 3D electronic–photonic heterogeneous interconnect platforms enabling energy-efficient scalable architectures for future HPC systems. *IEEE J. Sel. Top. Quantum Electron.* **32**, 1–14 (2026).
121. Stojanovic, V. A UCIe optical I/O retimer chiplet for AI scale-up. In *2025 IEEE Hot Chips 37 Symposium (HCS)* 1–22 (IEEE, 2025).
122. Cao, R. et al. Thermomechanical and compression analyses for large-scale co-packaged optics (CPO) assembly. *IEEE Trans. Compon. Packaging Manuf. Technol.* **14**, 2079–2087 (2024).
123. Winterbottom, P. Photonic interconnect for accelerated computing celestial AI photonic fabric module—the world's first SoC with in-die optical IO. In *2025 IEEE Hot Chips 37 Symposium (HCS)* 1–13 (IEEE, 2025).
124. Nagarajan, R. et al. 2.5 D heterogeneous integration for silicon photonics engines in optical transceivers. *IEEE J. Sel. Top. Quantum Electron.* **29**, 8200209 (2023).
125. Li, H. Y. et al. The cost study of 300mm through silicon interposer (TSI) with BEOL interconnect. In *2013 IEEE 15th Electronics Packaging Technology Conference (EPTC 2013)* 664–668 (IEEE, 2013).
126. Lau, J. H. Current advances and outlooks in hybrid bonding. *IEEE Trans. Compon. Packaging Manuf. Technol.* **15**, 651–681 (2025).
127. Atabaki, A. H. et al. Integrating photonics with silicon nanoelectronics for the next generation of systems on a chip. *Nature* **556**, 349–354 (2018).

Acknowledgements

This work was supported by the Air Force Office of Scientific Research Young Investigator Program (YIP) (FA9550-23-1-0159), the National Science Foundation (NSF) under Electrical, Communications and Cyber Systems (ECCS) (ECCS-2332060 and ECCS-2328839), under Industry-University Cooperative Research Centers (IUCRC) (EEC-2231625), and under Materials Research Science and Engineering Centers (MRSEC) (DMR-2309037), the National Research Foundation of Korea (NRF) grant funded by the Korea government (Ministry of Science and ICT; RS-2025-24535263 and RS-2025-25442127), the Global–Learning & Academic research institution for Master's–PhD students and Postdocs (LAMP) Program of the NRF funded by the Ministry of Education (RS-2024-00442483), the Korea Institute for Advancement of Technology (KIAT) grant funded by the Korean government (MOTIE; RS-2025-02263458), the MOE AcRF Tier 3 grant (MOE-MOET32023-0003) 'Quantum Geometric Advantage' and the A*STAR, Singapore, Advanced Manufacturing and Engineering (AME) Individual Research Grant (IRG) under the Project M23M6c0109.

Author contributions

B.K., S. H. Choi, G.Z., Y.J.Y., Y.B. and S.K. contributed equally to this review article. B.K., S. H. Choi, G.Z., Y.J.Y., Y.B. and K.L. wrote the initial draft and researched data for this review article. B.K., Y.B., S.K., Jongchan Kim, Jeehwan Kim, S. H. Chae, H.K., S.H. and K.L. contributed to reviewing and editing before submission.

Competing interests

S.H. is an employee of SK hynix Inc. The other authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to Jongchan Kim, Jeehwan Kim, Sang Hoon Chae, Hyunseok Kim, Seunghoon Hong or Kyusang Lee.

Peer review information *Nature Electronics* thanks Jianji Dong and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© Springer Nature Limited 2026