# Co-Packaged Optics Integration for Hyperscale Networking

Sandeep Razdan[1], Matt Traverso, and Anthony Torza,
Cisco Systems, Inc.
[1]Senior Member IEEE

**Abstract**

Co-packaged Optics (CPO) is an emerging technology that integrates high bandwidth optical engines next to a compute chip on the same substrate. The technology's key advantage lies in the ability to provide high bandwidth [1, 10] and lower latency [18] while reducing overall system power consumption compared to traditional pluggable optical transceivers [1]. By leveraging silicon photonics-based optical engines, CPO achieves high level of optical and electrical device integration, using proven semiconductor fabrication technologies and design processes, thereby achieving scale, reliability, and cost reduction. However, critical challenges in assembly, optical coupling, mechanical and thermal design must be addressed to ensure reliable and scalable manufacturing. Demonstrating long term reliability, serviceability and interoperability are critical to wide scale adoption and deployment of CPO in hyperscale networks and high-performance computing applications.

**Introduction**

Web scale data centers and cloud service providers today handle massive amounts of data traffic that is growing at a rapid pace due to high bandwidth applications such as AI/ML (Artificial Intelligence/Machine Learning), high-resolution video streaming and AR/VR (Augmented Reality/Virtual Reality). Handling such massive amounts of data is expensive, complex, and requires lots of orchestration between networking, server/compute, and memory. This, in turn, is driving increasing demands on network & memory bandwidth, resulting in doubling the aggregate bandwidth of Switch systems/ASICs (Application Specific Integrated Circuits) and Optics every two to three years. Figure 1 below depicts the relative trends in ASIC and Ethernet optics bandwidth scaling over last several years [1, 2].

However, unlike Moore's law scaling, the increase in bandwidth density results in increasing power requirements for each successive generation. Typical data centers are built with fixed electrical power budgets & energy use forecasts, hence this scaling in bandwidth needs to be within strict power envelopes or require expensive upgrades in infrastructure.

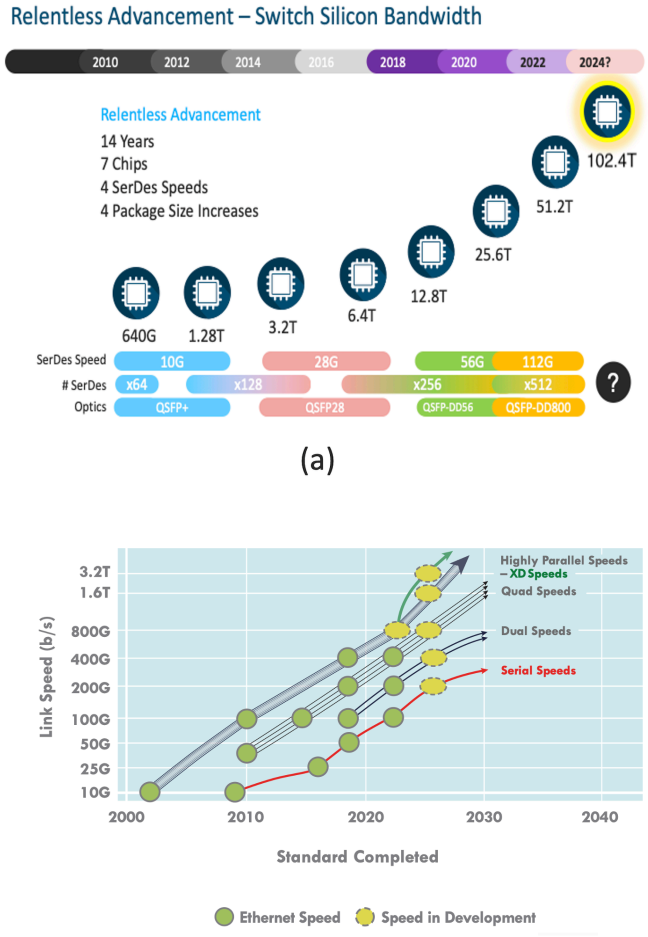


*Figure 1 ASIC and Ethernet bandwidth evolution with time depicting (a) ASIC bandwidth doubling roughly every 2 years, and (b) Ethernet speeds scaling with ASIC bandwidth.*

While the compute space drives a lot of efficiencies in silicon architecture, design and software, the networking bottlenecks remain thermal & power optimization (as well as cost reduction) with each successive technology generation.

Figure 2a and 2b show these trends for high bandwidth switch systems, with significant contributions from optics toward total system power as well as cost [1, 3].

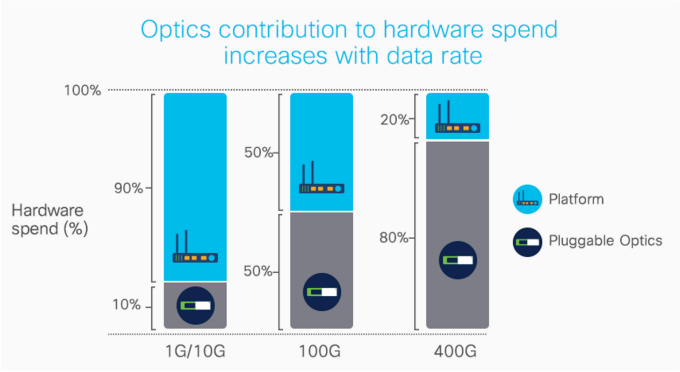


(a)

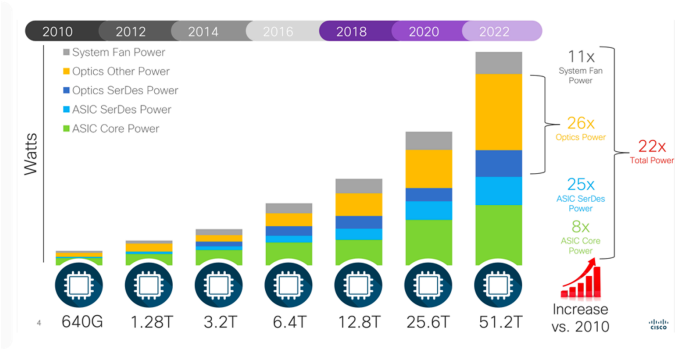


(b)

*Figure 2 (a) Optical Interconnects represent ~80% of total switch system cost at higher data rates [3] and (b) Optical Systems represent ~50% of total system power for 51.2T systems [1].*

As is evident, a large amount of system power and cost goes into moving bits from one part of the data center to another [1]. Even for AI/ML clusters, system bandwidth and latency between GPU/CPU, Memory and Networking interconnects forms a big part of system performance, power, and cost [4, 5, 18].

There are various efforts ongoing in the industry to optimize interconnect & Serdes (Serialization-Deserialization) power across switch systems – these include power efficient Serdes design by shrinking silicon nodes, low loss PCB materials, low loss flyover cables as interconnects, linear pluggable optics, as well as Co-Packaged Optics (CPO) [6-8].

CPO optimizes power consumption by bringing Optical transceivers closer to the ASIC chips, thereby eliminating retimers and Optics DSPs, and enabling use of low power Serdes options (e.g., XSR & VSR Serdes) [9, 10]. Figure 3 below depicts various architectural choices available for reducing power consumption for Serdes, with CPO approach being one of the most efficient in terms of system power optimization.

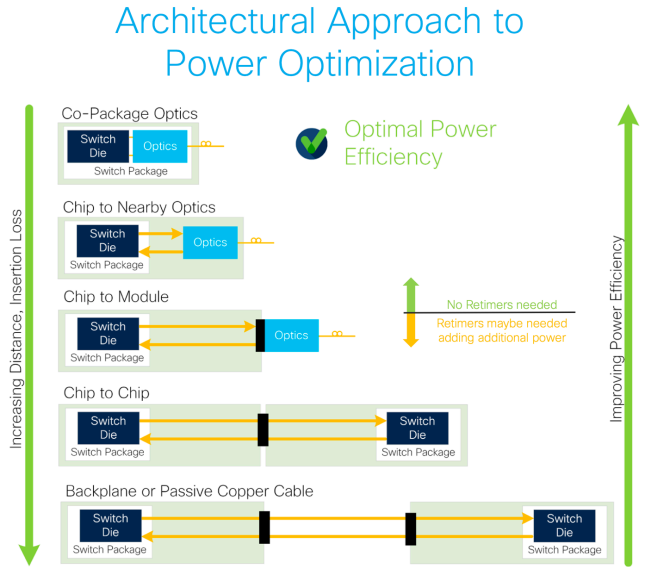


*Figure 3 Switch system power optimization by reducing interconnect distances between ASIC and Optical modules [1].*

## Co-Packaged Optics Development

Co-Packaged Optics (CPO) is a paradigm shift in building switch systems that involves packaging Optical engines on the same substrate or interposer as the Switch ASIC (or in CPU/GPU ICs for AI/ML applications). This allows the ASIC to drive the optical devices directly without using a DSP (Digital Signal Processor) in the optical engine. Figure 4 below shows a block diagram for a CPO architecture. The system depicts a direct drive interface between the ASIC and the Optical TIA, while using low power XSR (Extra Short Reach) interface on the transmit side.

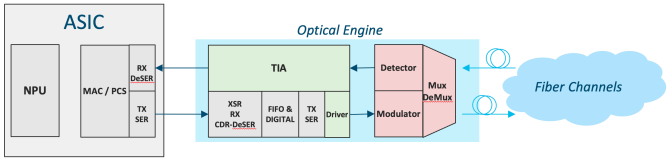


*Figure 4 Simplified block diagram for a CPO system*

By packaging optical engines next to the ASIC and allowing the ASIC to drive optical devices, a significant reduction in overall system power can be achieved – by some estimates up to 25-30% [11, 12]. This is due to elimination of a power-hungry DSP in the Optics, as well

as using lower power Serdes on the ASIC (XSR vs. LR). Using CPO also helps reduce overall system cost, by eliminating need for expensive components such as Optical DSPs, electrical connectors, high speed compatible PCBs or flyover cables [11]. Figure 5 below depicts such a power comparison for 51.2T Switch system between traditional pluggable Optics vs CPO [11].

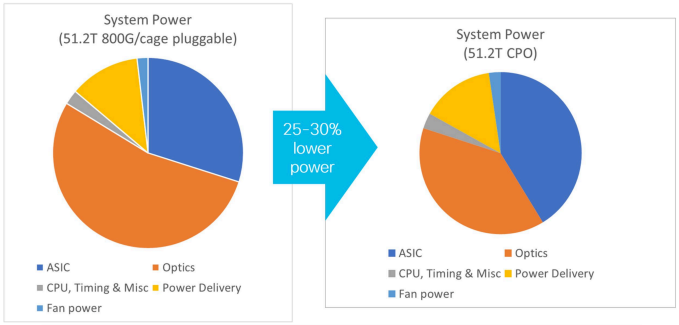


*Figure 5 System power comparison between pluggable optics vs. CPO for 51.2T switch system*

CPO implementation involves system scale integration of various components – including Optical Engines, ASIC package, fiber arrays, large body substrates, power management components, thermal assembly as well as routed fiber assemblies. CPO systems also consist of external laser sources that provide optical power at a low loss to the optical engine modules. An example implementation of a CPO system is shown in Figure 6 below.

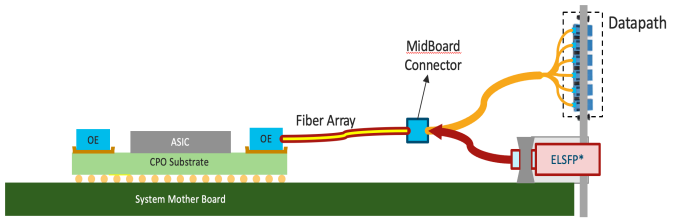


*Figure 6 Schematic drawing of a typical CPO implementation*

The key components of a CPO system are the Optical Engine and switch system assembly, which we will discuss in the next few sections. Designing a low loss, high signal-to-noise ratio (SNR) channel between the ASIC and the Optical Engine is critical for enabling power reduction in a CPO system. To achieve this, photonics & analog device design, package design, and choice of materials systems (substrate dielectric materials and advanced fabrication technologies) all play a critical role toward enabling a low loss, high-speed interconnect that maximizes power savings at the system level.

**Silicon Photonics Based Optical Engines**

Silicon Photonics technology utilizes standard complementary metal-oxide-semiconductor (CMOS) fabrication processes to enable dense integration of multiple optical and electrical devices on a single silicon chip. These devices include optical modulators, photodetectors, low loss optical waveguides, optical coupling structures, thermistors, sensors, capacitors [13] and in some cases, even analog electrical devices such as Transimpedance Amplifier (TIA) and Driver [14]. This allows for cost effective, scalable mass production of photonics interconnects that can also enable co-packaging with ASICs (or CPUs/GPUs) using semiconductor packaging & assembly techniques.

Using advanced silicon photonics fabrication techniques, Cisco has been able to develop highly scalable PAM-4 Optical interconnects with critical features such as low loss modulators [13], low loss optical coupling structures [16] and integrated Mux-Demux (Multiplexing-Demultiplexing) in a 400G FR4 architecture [15]. This allows for miniaturization of high bandwidth optical engines that can be efficiently co-packaged with ASICs on the same substrate.

To achieve high levels of integration between the Photonics Integrated Circuit (PIC), Electrical Integrated Circuits (EICs) & other components, as well as to optimize signal & power integrity for the optical system, 2.5D or 3D packaging architectures provide a distinct advantage – both in terms of high level of IC integration on a small footprint, as well as providing low loss, short interconnects with efficient power delivery both within and external to the Optical Engine. Package architectures such as Fan-out packaging or TSVs (through-silicon vias) enable tight integration of PIC & EIC components in a compact footprint and allow use of well-established semiconductor assembly processes to optics co-packaging. Co-packaging optics using 2.5D or 3D packaging also enable efficient thermal dissipation path by use of thermal vias and separate thermal heat spreader lids for optical engines vs. the ASIC to avoid thermal crosstalk. Figure 7 below shows an example of a 3.2T Optical Engine from Cisco based on Silicon Photonics platform, highlighting the miniaturization of a high channel density, high bandwidth optical device.

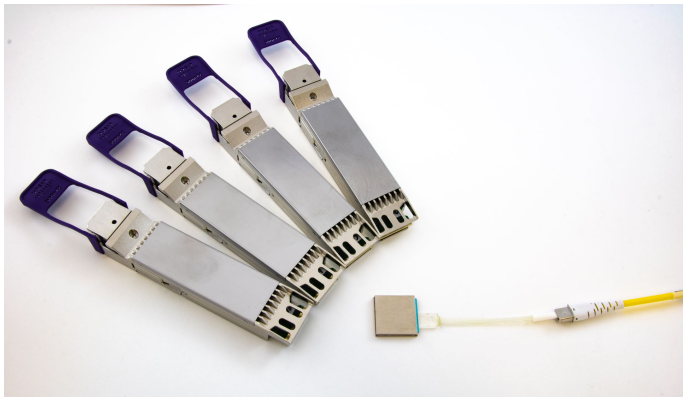

*Figure 7 4x OSFP800 vs. Cisco 3.2T CPO module*

Fiber coupling to optical engines used for CPO is another big challenge. Depending on the CPO application, optical engines can have anywhere from 24-72+ fiber channels (depending on 2xFR4 or DR8 configuration, assuming 100G-PAM4 & 3.2T bandwidth), including PM fibers used to couple remote laser optical power to the silicon photonics IC. For such a large array of fibers, chip or package warpage plays a major role in determining optical coupling efficiency and reliability of the coupling solution. Figure 8 below shows Monte Carlo simulation results for a 24-channel Fiber array coupled to a PIC die edge, depicting die edge warpage correlated with different fiber channels.

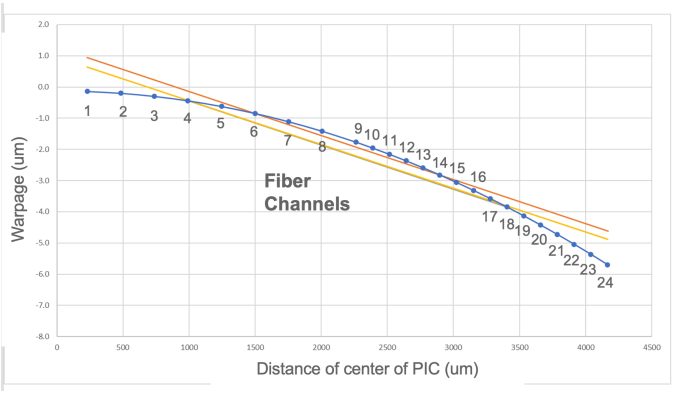


*Figure 8 Monte Carlo simulation results for die edge warpage with respect to 24-channel fiber array.*

Optical coupling of large channel density fiber arrays thus required optimization of overall optical package warpage, fiber alignment axis, and efficiency of the coupling structures used in the optical assembly.

**System Integration for Co-Packaged Optics**

Co-packaging optical engines on an ASIC platform requires complex chip-level and system-level assembly. This requires several innovations in co-package mechanical & thermal design, power delivery architecture (delivering sufficient power to both ASIC & optical tiles in a small form factor), package warpage management and assembly material selection.

Large body substrates (> 80 x 80mm) are typically needed to house both Switch ASIC as well as 8-16x Optical tile assemblies. This gives rise to high warpage and high stress in the package owing to CTE mismatch between various assembled components. Warpage management using right combination of materials, assembly techniques as well as use of stiffener materials with optimized CTE is highly critical to achieve a high performing, reliable co-package. Advanced substrate materials and fabrication technologies are also critical to enable low loss channels between ASIC and the optical tiles, where reducing channel discontinuities, lowering channel loss, and enabling short channel lengths are all critical parameters.

Thermal management for CPO becomes critical as both the high-power ASIC and various Optical Engines share the same real estate, thereby increasing thermal density at the center of the system. Prevention of thermal crosstalk between the ASIC and the optical tiles is critical for performance and reliability of the system. Design and strategic placement of enhanced surface area heat sink systems that separate heat dissipation paths for ASIC and the optics, along with advanced air-cooling mechanisms help efficiently dissipate heat and prevent temperature related performance degradation. Other factors used to enhance thermal dissipation include use of advanced thermal interface materials (TIMs) that fill gaps in highly warped packages, as well integration of design features such as thermally conductive vias in the package help distribute heat more uniformly across the system.

To make a CPO system more reliable and robust, adding redundancy in both electrical and optical power sources is very critical. For this reason, the laser sources are developed as a remote laser package that is assembled separately from the ASIC package, owing to concerns around thermal crosstalk and lack of repairability options in case of field failures. Cisco, along with several other industry partners drove the development of standard multi-source agreement (MSA) for ELSFP (External Laser Small Form Factor Pluggable) modules, with an objective toward cost efficiency, manufacturing scalability and standardization of optical power sources for CPO. These allow lasers to be passively cooled (and thus run with more efficiency and reliability [12]) as well

as being replaceable in the field. Figure 9 below shows ELSFP design as proposed by OIF committee [17].

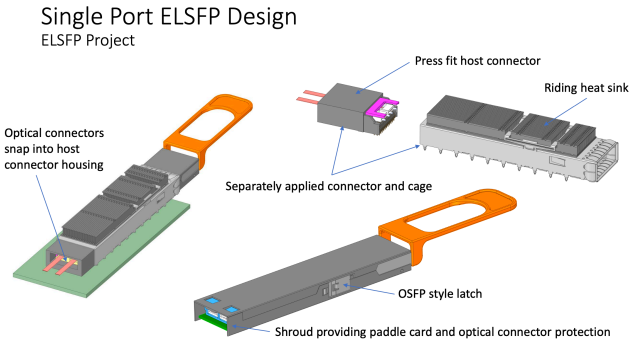


*Figure 9 ELSFP design proposed by OIF [17]*

For typical 25.6T/51.2T switch system-based CPO implementations, these usually consist of high-power lasers (various power classes are defined in the OIF document) packaged together in a pluggable form factor, with blind mate optical connectors enabling "hot swap" replacement capabilities. Having redundancy in number of lasers and laser channels is critical to ensure reliability and continuity of operation.

Figure 10 shows a 25.6T CPO Switch system demonstrated by Cisco at OFC 2023 [11], that consisted of 8x 3.2T Optical Engines (64x400G FR4) co-packaged with a 25.6T Silicon One G100 ASIC on a common substrate. The demonstration also included using ELSFP modules from multiple vendors used as a reliable optical power source.

*Figure 10 Cisco's 25.6T CPO System demonstrated in OFC 2023*

**Challenges in Scaling & Deployment**

There are a variety of challenges associated with CPO system integration, especially signal integrity & high-speed channel optimization, packaging, thermal as well as reliability of the overall system. Minimizing thermal crosstalk between a high thermal dissipation power ASIC and Optics, as well as managing stress and warpage for the overall package also present unique challenges and require innovative thermal, opto-mechanical, and electrical design.

Modeling and demonstrating reliability of various components is critical to the success and adoption of CPO across data center networking and AI/ML applications. Since majority of optical engine components (except laser sources) are inside the switch box, field repair or replacement are not viable options in case of failure of any optical components. Thus, the reliability and Failure-in-Time (FIT) rates for optical engine components must be proven to be acceptable to data center use case applications, and in certain cases may need to be proven equal or better than switch ASIC system components. Silicon Photonics based optical engines use reliable CMOS technologies with well understood failure models and field use data that, along with partitioning of remote laser sources away from the co-package, provide a reliable path to large scale deployment for these systems.

**Summary**

Co-packaged optics enables significant power and cost reduction for hyperscale networking applications, including high performance computing and AI/ML applications. As the bandwidth density for these applications grows and cost & power optimization become significant bottlenecks for data center, CPO technology offers a viable, scalable solution to reduce overall system power and cost for advanced interconnect technologies such as 224Gbps per channel bandwidth and switch systems at 102.4T aggregate bandwidth and higher.

Enabling CPO requires tight integration in design and assembly of various optical and electrical components, especially high bandwidth optical engines and ASIC/CPU/GPU packages. Integration challenges related to optical assembly, mechanical design & assembly, and thermal solutions need to be addressed for viable deployment of CPO in the data center environment.

Silicon Photonics enables miniaturization of optical & electrical components and allows co-packaging of optics

and ASICs together using standard semiconductor processes. This is critical to enable high density optical/electrical channel bandwidth in small form factors and provide a path to reliable and cost-effective systems.

Several technical solutions now exist to mitigate integration, yield and reliability challenges and several implementations of the CPO technology have been demonstrated by Cisco and other industry partners. Key challenges in wide scale adoption of CPO however remain due to concerns around field reliability, repairability, possible system downtime as well as demonstrating significant cost and performance advantages over comparable pluggable optics solutions. Despite these challenges, co-packaged optics technology shows immense potential to reduce overall system power, interconnect latency, and enable environmental sustainability and governance goals for next generation data centers and high-performance computing applications.

**References**


[1] R. Chopra, "Looking beyond 400G - a system vendor perspective," https://www.ieee802.org/3/B400G/public/21_02/chopra_b400g_01_210208.pdf (2020).

[2] "2022 Ethernet Roadmap," Ethernet Alliance, https://ethernetalliance.org/wp-content/uploads/2023/03/EthernetRoadmap-2023-Website-REV-March-17.pdf.

[3] https://blogs.cisco.com/news/optics-are-the-link-to-more-sustainable-reliable-networks

[4] A. Kumar, L. Chang, G. Tellez, L. Clevenger, J. Burns, "System performance: from enterprise to AI.", Proceedings of 2018 IEEE International Electron Devices Meeting (IEDM). pp.28.1.1–28.1.4 (2018).

[5] P. Siegl, R. Buchty, M. Berekovic, "Data-centric computing frontiers: a survey on processing-in-memory", Proceedings of the Second International Symposium on Memory Systems. pp.295–308 (2016).

[6] H. Mekawey, M. Elsayed, Y. Ismail, M. A. Swillam, "Optical Interconnects Finally Seeing the Light in Silicon Photonics: Past the Hype", *Nanomaterials (Basel)*, 12(3), p. 485, Feb 2022.

[7] D.A. Miller, "Optical interconnects to electronic chips", *Appl. Opt.,* 49, F59–F70, 2010.

[8] H.J. Caulfield, S. Dolev, "Why future supercomputing requires optics", *Nat. Photon.,* 4, pp. 261–263, 2010.

[9] Optical Internetworking Forum: *Common electrical I/O (CEI) 112G*, OIF (2017).

[10] C. Minkenberg, R. Krishnaswamy, A. Zilkie, D. Nelson, "Co-Packaged datacenter optics: Opportunities and challenges", *IET Optoelectronics*, 15(2), pp. 77-91, Apr 2021.

[11] https://blogs.cisco.com/sp/cisco-demonstrates-co-packaged-optics-cpo-system-at-ofc-2023

[12] J. E. Johnson, K. Bacher, R. Schaevitz and V. Raghunathan, "Performance and Reliability of Advanced CW Lasers for Silicon Photonics Applications," *2022 Optical Fiber Communications Conference and Exhibition (OFC)*, San Diego, CA, USA, 2022, pp. 1-27.

[13] M. Traverso et al., "Integrated Silicon Photonics Transmitters," *2022 27th OptoElectronics and Communications Conference (OECC) and 2022 International Conference on Photonics in Switching and Computing (PSC)*, Toyama, Japan, pp. 1-3, 2022.

[14] K. W. Ang et al., "Silicon photonics technologies for monolithic electronic-photonic integrated circuit (EPIC) applications: Current progress and future outlook," 2009 IEEE International Electron Devices Meeting (IEDM), Baltimore, MD, USA, pp. 1-4, 2009.

[15] Welch et al, Towards 400G BASE-FR4 Baseline Proposal, IEEE March 2019

[16] R. S. Tummidi and M. Webster, "Multilayer Silicon Nitride-Based Coupler Integrated into a Silicon Photonics Platform with <1 dB Coupling Loss to a Standard SMF over O, S, C and L Optical Bands," 2020 Optical Fiber Communications Conference and Exhibition (OFC), San Diego, CA, USA, pp. 1-3, 2020.

[17] OIF Co-Packaging Interop Demo (OFC 2023) https://www.oiforum.com/wp-content/uploads/OFC2023_OIF_Co-Packaging_Interop_Demo.pdf

[18] M Tan et al. "Co-packaged optics (CPO): status, challenges, and solutions", Frontiers of Optoelectronics, 16(1), Mar 20, 2023.