Proceedings of CONF-SEML 2026 Symposium: Multimodal Data Acquisition: Applications in Physiological and Behavioral Research
DOI: 10.54254/2755-2721/2026.AD32391

The Technology of 800G Optical Modules for AI Data Centers: Current Status and Challenges

Changcheng Fan

School of Communication & Information Engineering, Shanghai University, Shanghai, China
210609480@qq.com

Abstract. The rapid expansion of artificial intelligence (AI) workloads, particularly large language model training and inference, has driven unprecedented demand for high-bandwidth, low-latency interconnects within data centers. While 400G optical modules currently dominate the market, they are approaching their bandwidth limits, positioning 800G modules as a critical next-generation alternative. This paper presents a comprehensive review of 800G optical module technologies tailored for AI data center applications. Through a systematic literature review and comparative analysis of recent academic publications and industry specifications, three primary implementation approaches are examined: 8×100G PAM4, 4×200G PAM4, and 2×400G coherent detection. Furthermore, advances in core components—including thin-film lithium niobate modulators, high-extinction-ratio EMLs, and advanced DSP employing probabilistic constellation shaping—are discussed, alongside performance data extracted from recently demonstrated modules. Despite notable progress in modulator and DSP technologies, power consumption, thermal management, and standardization persist as formidable obstacles to large-scale deployment. By consolidating current knowledge, this review aims to offer insights for researchers and engineers developing next-generation AI data center networks.

Keywords: 800G optical module, AI data center, PAM4, EML

1. Introduction

Optical transceivers, which convert electrical signals to optical signals and vice versa, are fundamental components in optical communication systems [1]. In recent years, the rapid advancement of artificial intelligence (AI) has driven the evolution of data centers toward hyperscale architectures [2]. These AI workloads generate massive amounts of data that require efficient processing and communication, leading to an exponential increase in bandwidth demand within data centers [3]. While 400G optical modules currently dominate the market, they are approaching their bandwidth limits and struggling to meet the requirements of next-generation AI clusters [4]. Consequently, 800G optical modules have emerged as a key research focus, which are expected to be widely deployed in the near future. However, the 800G technology also faces significant challenges that hinder its large-scale adoption. Therefore, a systematic review consolidating current knowledge and comparing technical approaches is necessary to guide both academic research and industrial development.

This paper aims to provide a comprehensive review of 800G optical module technologies for AI data center applications. Through a systematic literature review and comparative analysis of recent academic publications and industry specifications, it examines three main implementation approaches (8×100G PAM4, 4×200G PAM4, and 2×400G coherent detection), core component advances (modulators, lasers, DSP), and performance data from demonstrated modules. The study also analyzes major deployment challenges, including power consumption, thermal management, signal integrity, manufacturing yield, and standardization. The findings are intended to offer insights for researchers and engineers developing next-generation AI data center networks.

2. Evolution of demand for optical interconnects in AI data centers

2.1. Characteristics of network bandwidth requirements for AI training and inference

The exponential growth of global data traffic has placed unprecedented demands on data center networks [5]. The International Telecommunication Union (ITU) projects that monthly mobile data traffic will reach 5 ZB by 2030, further intensifying the need for high-capacity infrastructure [5]. Within data centers, optical interconnects must satisfy three key objectives: ultra-high capacity to handle massive data throughput, ultra-low latency for real-time applications, and ultra-low power consumption for sustainable growth [5]. These demands are particularly acute in AI-driven environments, where large-scale model training and inference require rapid and efficient data exchange [2].

Network performance is critical in large-scale AI model training. Unlike general-purpose data center traffic, AI training workloads are characterized by "elephant flows"—large sustained data transfers from parameter synchronization across thousands of graphics processing units (GPUs) [6]. Traditional load balancing techniques often struggle with such traffic patterns, leading to suboptimal link utilization and increased flow completion times [6]. But optimized load balancing can improve bandwidth utilization by up to 38% and shorten model training tasks by over 3% compared to conventional approaches [6].

Meanwhile, AI inference applications such as autonomous driving, industrial IoT, and real-time analytics require ultra-low latency and deterministic response [7]. To meet these requirements, networks must support real-time traffic control, dynamic workload allocation, and continuous health monitoring [7]. Edge computing further amplifies these demands, requiring seamless integration between edge nodes and core data center infrastructure [7].

2.2. Technical features and limitations of current mainstream optical modules (400G)

Currently, 400G optical modules are the predominant solution in hyperscale data centers, with the 400G FR4 scheme being widely adopted [8]. This scheme typically employs 4×100 Gb/s 4-level Pulse Amplitude Modulation (PAM4) on the optical side and 8×50 Gb/s PAM4 on the electrical side, using a Quad Small Form Factor Pluggable-Double Density (QSFP-DD) or Quad Small Form Factor Pluggable (QSFP) form factor to provide eight electrical lanes [8]. To achieve the required performance, advanced technologies such as silicon photonics and digital signal processing (DSP) are integrated. Silicon photonics leverages mature Complementary Metal-Oxide-Semiconductor (CMOS) fabrication processes to deliver high bandwidth, high integration, low latency, and low power consumption [8]. DSP chips compensate signal distortion, reduce clock jitter, and perform forward error correction to ensure a low bit error rate [8]. More recent implementations replace Electro-absorption Modulated Laser (EML)-based transmitters with silicon photonic solutions—

employing separate continuous-wave lasers and silicon modulators—to further improve reliability and reduce power [8].

Despite these advances, 400G modules face inherent limitations when scaled to meet next-generation AI cluster demands. First, the bandwidth ceiling of the 4×100G architecture restricts per-port throughput, requiring multiple parallel modules to achieve the aggregate bandwidth needed by high-end AI servers such as NVIDIA DGX H100 [2]. A 400G-based switch fabric thus requires more physical ports and fiber connections than an 800G-based design. Second, the per-gigabit efficiency of 400G modules is lower than that of emerging 800G solutions, making them less attractive for hyperscale deployments where power and space are at a premium [4].

3. Key technologies and implementation schemes of 800G optical modules

3.1. Basic framework and classification of 800G optical modules

800G optical modules are primarily available in two form factors: QSFP-DD and Octal Small Form Factor Pluggable (OSFP). The QSFP-DD form factor supports 800G transmission through integration of multiple lanes. For instance, a thin-film lithium niobate (TFLN)-based 800G QSFP-DD transceiver incorporating two 4×100G DR4 chips has been demonstrated [1]. Compared to silicon photonics, this TFLN approach halves the number of laser sources, offering potential cost and power advantages [1]. The 800G OSFP DR8 silicon photonic module provides eight high-speed electrical channels for 8×100G PAM4 signals, achieving 800G throughput [4]. It comprises DSP, receiver optical sub-assembly (ROSA), transmitter optical sub-assembly (TOSA), and control circuit, using Chip-On-Board (COB) packaging to increase integration [4].

3.2. Comparison and analysis of different implementation schemes

For 10 km reach, three main 800G approaches have been proposed: 8×100G direct detection, 4×200G direct detection, and 1×800G coherent detection [4]. The 8×100G scheme, exemplified by the OSFP DR8 silicon photonic module, uses eight lanes of 100G PAM4 signals with LAN Wavelength Division Multiplexing (LWDM) and benefits from low DSP complexity [3,4]. This approach leverages silicon photonics technology and COB packaging to achieve high integration [3]. The 4×200G scheme halves the number of optical components but demands higher-bandwidth devices and more powerful forward error correction (FEC), increasing DSP complexity [4]. The 1×800G coherent approach offers better spectral efficiency and longer reach potential, but its technical requirements are significantly higher [4].

Another emerging 800G implementation is the 2×400G coherent scheme, targeting long-haul applications [9]. Unlike the 8×100G PAM4 solution designed for short-reach links, coherent solutions leverage advanced modulation techniques for extended distances. At 2,400 km and 5,000 km, the same module supports 600 Gb/s and 400 Gb/s respectively, demonstrating the flexibility of coherent technology [9].

3.3. Progress in core components: modulators, lasers, detectors, and DSP

Advanced modulator technologies are critical for 800G. Thin-film lithium niobate (TFLN) modulators have demonstrated 40 GHz bandwidth with S11 below −12 dB, enabling 100 Gb/s PAM4 per lane [1].

Electro-absorption modulated lasers (EMLs) are widely used due to monolithic integration of laser and modulator, offering superior speed and extinction ratio (ER) compared to directly modulated lasers [2]. Two competing technologies dominate: EMLs provide higher ER and lower drive voltage with proven reliability, while silicon photonics (SiPh) offers greater scalability and cost advantages in mass production [2]. State-of-the-art 100G PAM4 EMLs achieve threshold currents as low as 15 mA, output power >9 mW at 53 °C, and ER ≥8 dB at 1 Vpp, enabling transceiver-level bit error rate (BER) below 1×10^{-9} [2].

On the receiver side, high-speed p-type-intrinsic-n-type (PIN) photodiodes convert optical signals back to electrical signals; in silicon photonic modules, germanium-based detectors integrated on silicon offer high responsivity and bandwidth [8].

DSP is critical for both PAM4 and coherent modules. In PAM4-based 800G modules, DSP handles equalization, crosstalk cancellation, and FEC [3,4]. For coherent systems, DSP enables digital Nyquist subcarriers and probabilistic constellation shaping (PCS) to improve spectral efficiency and reach, as demonstrated in a 2×800G module using 7 nm DSP ASICs and InP-based photonic integrated circuits (PICs) [9].

3.4. Current product status and performance indicators of 800G optical modules

Among the implementations above, the 8×100G solution has the most mature industrial chain, with proven lasers and DSPs available. Its 100G electrical lane rate offers better cost and power efficiency than coherent alternatives, avoiding extreme signal integrity challenges of 200G lanes [4]. Therefore, the 800G QSFP-DD LR8 module based on 8×100G PAM4 is considered practical for near-term deployment [4]. An 800G PAM4 LR8 module using eight-channel LWDM and cooled EML lasers has been tested, achieving average output power ~3 dBm, ER 4-6 dB, optical modulation amplitude (OMA) ~3 dBm, and Transmitter and Dispersion Eye Closure for PAM4 (TDECQ) 2-2.5 dB [4]. Receiver sensitivity remained below –6.8 dBm after 10 km, satisfying 400G-LR4-10 specifications [4].

Coherent 800G modules have also reached commercial maturity. A two-channel digital coherent module based on InP PIC and 7 nm DSP ASIC has demonstrated 800 Gb/s over 1,000 km using 96 GBd PCS-64QAM, and can support 600 Gb/s over 2,400 km or 400 Gb/s over 5,000 km [9].

4. Challenges in large-scale deployment of 800G optical modules

Despite rapid technological advances, widespread adoption of 800G optical modules in AI data centers faces several key challenges.

Power consumption and thermal management remain critical. As data rates increase, thermal management becomes more difficult [3]. Emerging solutions such as Linear-drive Pluggable Optics (LPO) and liquid-cooled modules based on silicon photonics are being developed to address this challenge [3]. According to the LPO Multi-Source Agreement (MSA) 100G-DR-LPO specification, low power is a primary design objective for next-generation interconnects [10]. By removing the power-hungry DSP, linear architectures address this, though thermal management within existing form factors remains a concern.

Signal integrity and manufacturing yield pose additional engineering challenges. At 100G per lane, the LPO MSA mandates an effective return loss >10 dB to manage reflections and requires crosstalk consideration in testing [10]. It extends maximum electrical channel loss to 16 dB at 26.56 GHz, underscoring the need to mitigate inter-symbol interference in linear architectures lacking digital regeneration [10]. Tight constraints for the Medium Dependent Interface, referenced to International Electrotechnical Commission (IEC) standards, and budgeting of optical channel loss highlight the criticality of fiber-to-chip coupling precision, directly impacting yield in multi-lane modules [10].

Standardization and interoperability have progressed with the March 2025 publication of the LPO MSA 100G-DR-LPO specification. This standard defines 100 Gb/s per lane optical interfaces using PAM4 at 53.125 GBd, supporting configurations up to 800G via lane aggregation [10]. It builds upon Institute of Electrical and Electronics Engineers (IEEE) 802.3 and Optical Internetworking Forum (OIF) CEI-112G-LINEAR-PAM4, specifies normative test points (TP1a to TP4a) with detailed parameters, and mandates RS(544,514) FEC at the host with Common Management Interface Specification (CMIS) extensions [10]. These specifications enable multi-vendor interoperability for 800G LPO modules, directly addressing the industry's need for sustainable, high-bandwidth interconnects in AI data centers.

5. Conclusion

This paper has reviewed the current state of 800G optical module technologies for AI data center applications, focusing on key implementation schemes, core component advancements, and major deployment challenges. Three main conclusions can be drawn. First, among the various 800G solutions, the 8×100G PAM4 scheme, leveraging mature EML or silicon photonics technologies, is currently the most practical for short-reach intradatacenter links up to 10 km, while coherent 2×400G designs are better suited for longhaul interconnects. Second, advances in modulators (TFLN, EML) and DSP (e.g., probabilistic constellation shaping) have enabled the required performance, but power consumption and thermal management remain critical bottlenecks. Third, standardization efforts, notably the recent LPO MSA 100GDRLPO specification, are paving the way for interoperable, lowpower linear pluggable optics.

Nevertheless, this review is limited to publicly available literature and does not include experimental validation or detailed cost analysis. Future work should focus on practical deployment aspects, such as total cost of ownership and reliability testing under real AI workloads. Looking ahead, research will likely shift toward 1.6T modules, copackaged optics, and further reductions in power through advanced linear driver architectures and novel cooling techniques. These developments will be essential to sustain the exponential growth of AI computing infrastructure.

References

[1] H. Li et al., "800 Gbit/s QSFP-DD Transceiver Based on Thin-Film Lithium Niobate Photonic Integrated Circuit, " J. Light. Technol., vol. 41, no. 11, pp. 3462–3469, Jun. 2023, doi: 10.1109/JLT.2023.3238844.

[2] J. Jia-Sheng Huang et al., "100G PAM4 High-extinction Ratio EML Manufacturing for 400G and 800G Artificial Intelligence (AI) Optical Transmissions, " Current Optics and Photonics, Jul. 2025, doi: 10.17352/cop.000003.

[3] M. Tang, H. Zhu, Y. Zhu, and S. Wang, "Research of 800Gbit/s OSFP DR8 Silicon Photonics Optical Transceiver Module, " in 2024 4th International Conference on Electronic Information Engineering and Computer Communication (EIECC), Wuhan, China, 2024, pp. 187–191. doi: 10.1109/EIECC64539.2024.10929203.

[4] P. Xiang, X. Song, Y. Zheng, Y. Ye, C. G. Liu and L. Mu, "Research and design of 800Gbit/s PAM4 LR8 10km optical module, " in 2022 7th International Conference on Communication, Image and Signal Processing (CCISP), Chengdu, China, 2022, pp. 18-22, doi: 10.1109/CCISP55629.2022.9974495.

[5] Y. Gong, C. Yang, and T. Zhou, "Development and Application of 'Three-Super' Optical Interconnection Technology in Data Centers, " Study on Optical Communications, Nov. 2024. [Online]. Available: https://link.cnki.net/urlid/42.1266.TN.20241108.1736.004

[6] R. Zhang et al., "End-to-Network Load Balancing for AI Data Center Networks: A Convergence-Based Approach to Enhance Training Efficiency, " Transactions on Emerging Telecommunications Technologies, vol. 36, no. 10, p. e70249, 2025. doi: 10.1002/ett.70249.

[7] C. Mavani, H. K. Mistry, R. Patel, and A. Goswami, "Artificial Intelligence (AI) Based Data Center Networking, " International Journal on Recent and Innovation Trends in Computing and Communication, vol. 12, no. 2, pp. 508–518, 2024. [Online]. Available: https: //ijritcc.org/index.php/ijritcc/issue/view/27

[8] Z. Song, X. Yin, Z. Liu, and Q. Hao, "Design of 400G FR4 silicon photonics optical transceiver module, " Microprocess., 2023, doi: 10.19304/J.ISSN1000-7180.2023.0700.

[9] S. Porto et al., "Demonstration of a 2 × 800 Gb/s/wave Coherent Optical Engine Based on an InP Monolithic PIC, " J. Light. Technol., vol. 40, no. 3, pp. 664–671, Feb. 2022, doi: 10.1109/JLT.2021.3121284.

[10] LPO MSA, "Specification for 100 Gb/s per Lane Linear Pluggable Optics Single-Mode Optical Fiber Transmission, " LPO MSA Specification 100G-DR-LPO, Rev. 1.0, Mar. 2025. [Online]. Available: https: //www.lpo-msa.org/files/live/sites/lpomsa/files/specs/LPO_MSA_Specification_v1p01.pdf