Opto-electronic Convergence Technology for Achieving High Capacity and Low Power Consumption in Next-generation Data Centers

Authors: *Nobuo Ohata*, Mizuki Shirao**

**Information Technology R&D Center*

Abstract

With the emergence of generative AI represented by Large Language Models (LLM) in recent years, the amount of computation required for machine learning has continued to grow immensely. As computation capacity increases, the power consumption of AI data centers continues to increase, and has become a major social issue. As one solution to this challenge, opto-electronic convergence technology has been attracting attention. Co-packaged optics, which compactly integrate electrical ICs and optical devices, can deliver high-capacity communications and lower power consumption that are difficult to achieve with conventional optical transceivers, and are expected to be deployed in next-generation AI data centers. After outlining the evolution of backend networks in GPU clusters and the features of co-packaged optics, this report introduces our initiatives toward co-packaged optics.

1. Introduction

With the emergence of generative AI represented by LLM, the computation capacity required for machine learning has continued to grow immensely. Figure 1 plots the normalized computational performance (FLOPS) of processors, the network bandwidth for computing, and the number of parameters used in machine learning[(1)]. The number of parameters used in machine learning has been increasing at a pace of about 600-fold every two years; OpenAI's GPT-3[*1] uses 175 billion parameters, and Switch Transformer and GPT-4[*1] use over 1 trillion parameters. Meanwhile, the computational performance of processors increases about 2.6 times every two years, and the network bandwidth required for computing increases about 1.5 times every two years; improvements in computing performance are modest compared with the growth in parameters. Therefore, by interconnecting more than several thousand GPUs in parallel via optical communications and clustering them, computational performance is greatly enhanced for executing machine learning. However, the computations require many processors and optical communication equipment, leading to enormous power requirements at AI data centers. To address this challenge, we have begun exploring co-packaged optics that can balance power efficiency in optical communications with expanded network bandwidth. This report presents trends in backend networks required for machine learning and technological trends in co-packaged optics configurations. After introducing the characteristics of the high-speed Electro-absorption Modulator integrated Laser (EML) we have developed to date, we describe high-density integration technologies for applying them to co-packaged optics.

*1 GPT-3 and GPT-4 are registered trademarks of OpenAI OpCo, LLC.

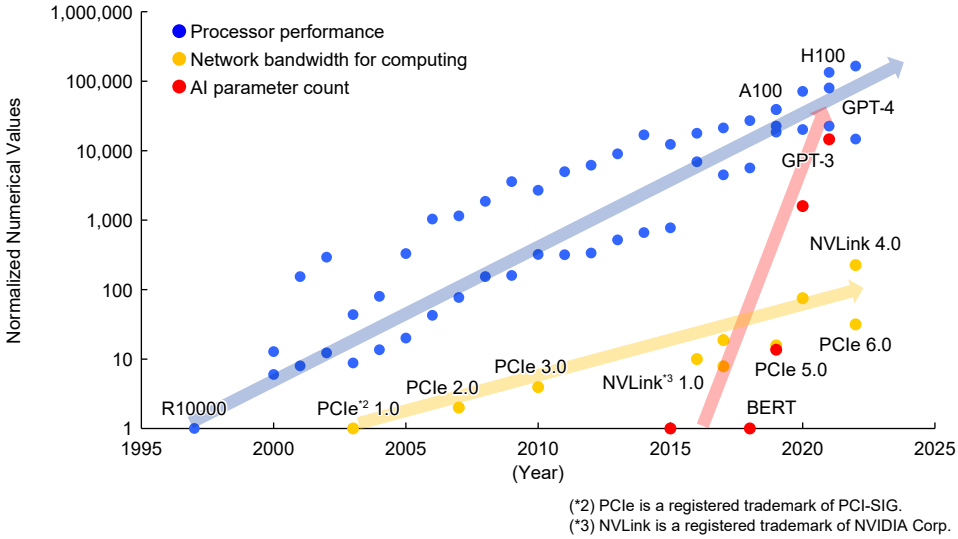

Fig. 1 Trends in processor performance, network bandwidth, and AI model parameters

2. Network Trends in AI Machine Learning and the Introduction of Co-packaged Optics

Figure 2 shows a conceptual diagram of the backend network. The configuration connects many GPUs in parallel through multiple switches. At NVIDIA, 72 GPUs are interconnected in parallel via NVSwitch using a proprietary high-speed network called NVLINK(2), but going forward, to enhance computational capacity, scaling is expected to progress from the current 72 GPUs to 576 and 1,152. It is expected that the number of optical transceivers used to connect GPUs and switches via optical communications will also increase, raising the issue of significantly higher power consumption. Furthermore, as the number of connections to GPUs increases, switch Application-Specific Integrated Circuits (ASIC) will also grow in capacity, and the signal bandwidth output from the switch ASIC will increase accordingly. Therefore, optical transceivers will need wider communication bandwidth, and the amount of signals that can be drawn per unit length—that is, edge density (Tb/s/mm)—will become an important metric going forward. Figure 3 shows the results of calculating edge density, referencing the trends shown in Fig. 1. Edge density will need to reach 1Tb/s/mm by 2028 and will increase exponentially to 2Tb/s/mm by 2030.

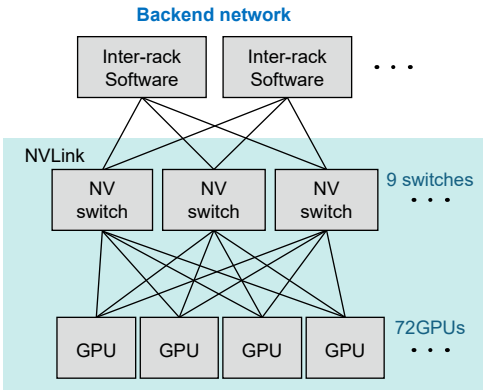

Fig. 2 Schematic diagram of the backend network

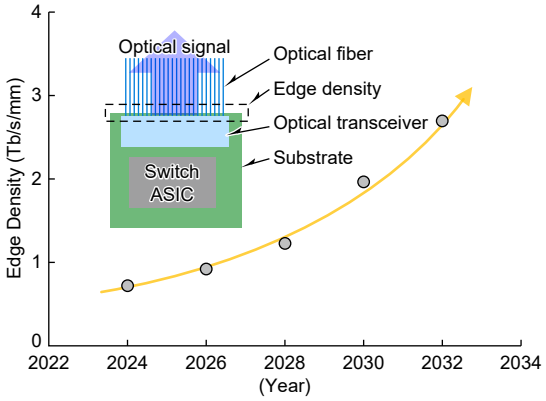

Fig. 3 Trends in bandwidth edge density

Because existing pluggable optical transceivers find it difficult to achieve low power and high edge density, co-packaged optics are expected to be a new technology that can solve these challenges. Table 1 compares pluggable optical transceivers and co-packaged optics. A pluggable optical transceiver is used by connecting to an electrical connector mounted on a printed circuit board. Therefore, the distance from logic ICs such as GPUs and switch ASIC to the transceiver is lengthy, requiring a DSP (Digital Signal Processor) to correct degraded electrical signals, which leads to high power consumption. Power efficiency is approximately 20pj/bit. Also, edge density is 0.1Tb/s/mm or lower due to wiring pitch constraints imposed by the connector. In contrast, co-packaged optics are mounted on the package substrate on which the logic

IC is installed, eliminating the need for a DSP and enabling power efficiency of 5pj/bit. Moreover, because they are implemented on substrates that support fine wiring, edge density of up to 2Tb/s/mm is achievable. Furthermore, in the future, configurations that mount on a silicon interposer are being considered; in that case, on the electrical L/S (line-and-space), edge density exceeding 2Tb/s/mm would be possible.

Table 1 Comparison of pluggable optical transceivers and co-packaged optics

	Implementation Diagram	Transceiver Top View	With or Without DSP	Power Efficiency	Bandwidth Edge Density
Pluggable optical transceivers	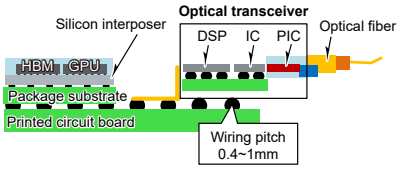	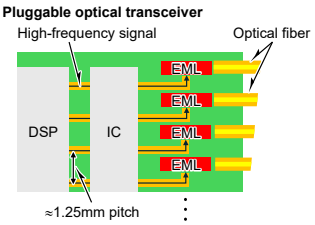	With	Approx. 20pj/bit	<0.1 Tb/s/mm
Co-packaged optics	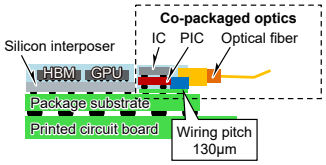	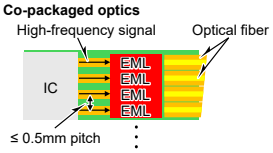	Without	Approx. 5pj/bit	<2 Tb/s/mm

We have begun examining co-packaged optics incorporating EML. An EML is a laser device that integrates a Distributed Feedback Laser Diode (DFB-LD) and an Electro-absorption Modulator (EAM), which absorbs light when a voltage is applied, on a single chip. Compared with Si photonics modulators, which excel at high-density integration, EML enables higher-speed operation at lower power, making it possible to achieve higher edge density. However, in co-packaged optics, EML must be placed in high densities, and it is difficult to achieve this with conventional assembly that connects the high-frequency traces to the EML with wires. We therefore considered adopting flip-chip bonding.

3. EML Implementation Technology for Low Power and High Edge Density

In this chapter, we introduce the ultra-high-speed EML[3][4] that we have developed, and describe EML implementation forms applicable to co-packaged optics and their characteristics.

3.1 Implementation configuration

Our EAM employs a unique high-mesa waveguide structure in which the optical absorption layer is sandwiched on both sides by low refractive index material (Fig. 4). This structure strongly confines light in the high-refractive-index optical absorption layer, enabling efficient optical absorption. In other words, sufficient optical absorption is achieved even with an EAM of low capacitance, enabling high-speed operation. Given that the EAM is located in front of the EML (on the light-output side), it is common to route high-frequency traces along the side of the EML and use gold wire for wire bonding to the EAM electrodes (Fig. 5). Wire bonding requires a large mounting footprint; in past developments it occupied a width of about 1.25mm per EML, limiting edge density. To achieve the high edge density required for co-packaged optics, we developed an implementation technology that does not use wire bonding.

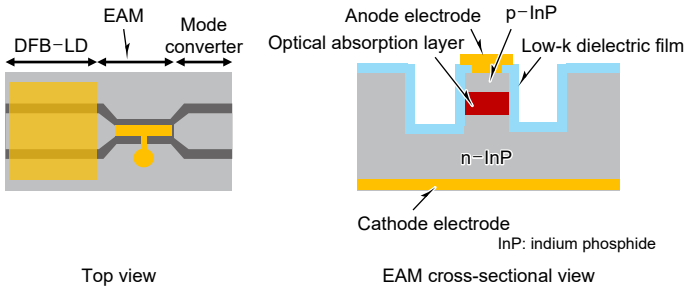

Fig. 4 EML employing high-mesa EAM

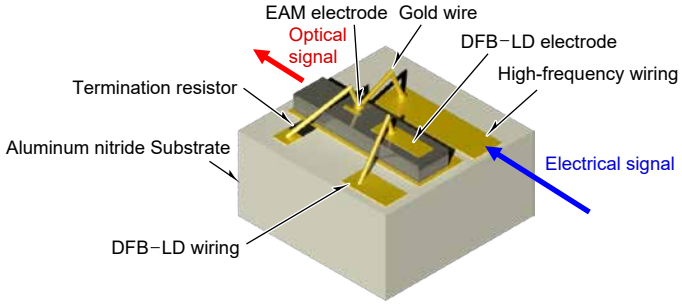

Fig. 5 Conventional assembly using wire bonding

3.2 Performance assessment

Figure 6 shows the configuration of the co-packaged optics. In co-packaged optics, the electronic IC and optical devices are mounted using flip-chip technology. In flip-chip bonding, the electrodes are directly joined using bumps, etc., so wire bonding as shown in Fig. 5 is not used. As a result, not only is a substantial increase in assembly density possible compared with conventional structures, but the parasitic inductance from wire bonding is eliminated, enabling higher operating speed. Simulations for flip-chip bonding indicate an operating speed of 145 GHz (Fig. 7).

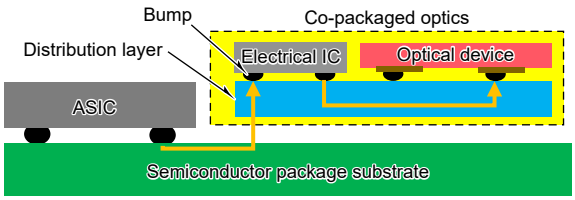

Fig. 6 Cross-sectional view of co-packaged optics using EMLs

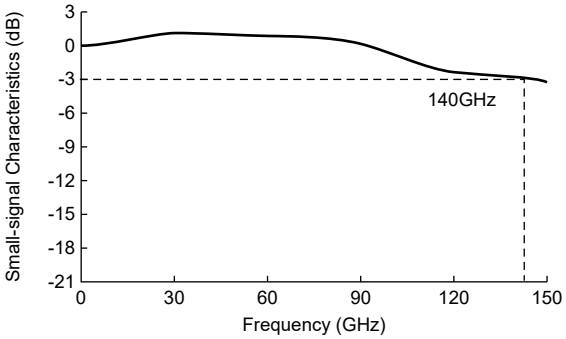

Fig. 7 Simulation results of co-packaged optics using EMLs

Experimental verification of flip-chip bonding onto an aluminum nitride substrate achieved an operating speed of 90GHz[5]. The small-signal characteristics are shown in Fig. 8. Ripples not seen in Fig. 7 are observed in the small-signal characteristics; these are believed to be due to impedance mismatches in the high-frequency interconnects. Going forward, if advances in manufacturing enable the arraying of EML, densification down to around 0.25mm—five times higher than before—can be expected.

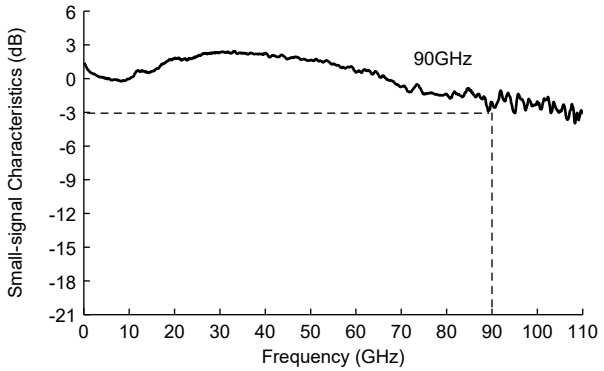

Fig. 8 Small-signal response of EML in co-packaged optics[5]

3.3 Future outlook

To increase the capacity of next-generation data centers, it will be necessary to optimize the device structure—including reducing EAM capacitance and arraying—to improve operating speed. In addition, using low-dielectric-constant substrate materials (e.g., quartz substrates) is expected to further improve operating speed. By achieving these and applying the next-generation modulation format, 200 Gbaud PAM4 (Pulse-Amplitude Modulation, 4 levels), an edge density of 1.6Tb/s/mm is anticipated.

4. Conclusion

This report presented trends in backend networks and co-packaged optics, introduced EML, and described our efforts toward co-packaged optics. Flip-chip bonding technology using EML with a high-mesa structure enables the high edge density required for co-packaged optics. With advances in AI and machine learning, the importance of opto-electronic convergence technology is expected to grow further, and, in addition to advances in optical device technology, advances in the packaging technologies for optoelectronics will be essential. We will continue development toward next-generation opto-electronic convergence technology.

References

(1) riselab: AI and Memory Wall
https://medium.com/riselab/ai-and-memory-wall-2cb4265cb0b8
(2) NVIDIA: NVLink and NVSwitch
https://www.nvidia.com/ja-jp/data-center/nvlink/
(3) Uchiyama, A., et al. : Demonstration of 155-Gbaud PAM4 and PAM6 Using a Narrow High-Mesa Electro-Absorption Modulator Integrated Laser for 400Gb/s Per Lane Transmission, Journal of Lightwave Technology, 43, No.4, 1868-1873 (2025)
(4) Shirao, M., et al. : High Speed EML and Assembly Techniques for GPU Cluster System, 2025 Optical Fiber Communications Conference and Exhibition, 1-3 (2025)
(5) Masuyama, K., et al. : EML Assembled with Flip-Chip Technology on AlN Sub-mount Operating at 212.5Gbps PAM4, OECC 2025, WG2-2 (2025)