Enhancing Distributed Computing with Optical Interconnects in Data Centers and High-Performance Computing (HPC) Systems

Abidur Rahman, Yuyang Wang, Brian Wu, Alex Meng
Lightwave Research Laboratory
Department of Electrical Engineering, Fu Foundation School of Engineering and Applied Science, Columbia University, New York, NY, USA

Introduction: Optical photonics interconnects are emerging as a highly promising solution for distributed deep learning (DDL) due to their ability to address the critical bottlenecks in communication among distributed computing units (CUs). As deep learning models like OpenAI's GPT-3 and GPT-4, Meta's Llama 2 and Llama 3, and Google's Gemini grow in size and complexity, the need for distributed training across multiple CUs has become essential. However, traditional electronic interconnects often struggle to provide the necessary bandwidth and low latency required for efficient communication in these distributed systems [1],[2]. Optical interconnects provide high-bandwidth, low-latency communication paths make silicon photonics an ideal solution for overcoming the communication bottlenecks in distributed deep learning. As a result, these technologies are poised to significantly improve the performance and scalability of deep learning models, enabling faster and more efficient training of increasingly complex neural networks.

Methods

We use the Roofline model to compare optical interconnect workloads with the latest GPU models on the market. The Roofline model will help us determine how much we deviate from the latest GPU models in terms of performance and operational intensity. The Roofline model is a visual performance model designed to provide insight into the performance limitations of multicore architectures [3].

Results

From the figure below, we can see that the Columbia Lightwave Research Laboratory (LRL) optical interconnect has a much higher link bandwidth for data centers and high-performance computing (HPC) systems. LRL has a 6000 GB/s bandwidth link speed, which is much higher than the existing optical link bandwidths in the market [4]. The roofline model has three components: GPU, link bandwidth, and workloads. The dotted horizontal lines represent GPUs, the solid lines represent link bandwidths, the dotted vertical lines represent Convolutional Neural Network (CNN) workloads, and all the crosses represent Large Language Model (LLM) workloads. If a workload is located on the right side of the bandwidth links, then that workload is compute-bound, which means that to improve the performance of the workload, we need better GPUs, CPUs, or memory. However, if a workload is located on the left side of the bandwidth links, we can say that the workload is link-bound, which means we need better link bandwidth to improve the performance of the workload [3].

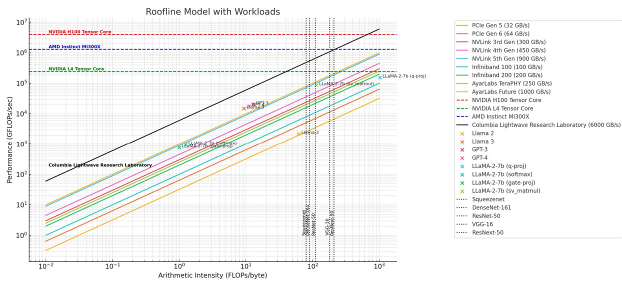

Figure 1: Roofline Model

Conclusions

Based on our research, Optical interconnects offer significant advantages for data centers and high-performance computing (HPC) systems due to their superior bandwidth, latency, and scalability characteristics. The transition to optical interconnects in data centers and HPC systems addresses critical challenges related to bandwidth, latency, scalability, and energy efficiency, making them an optimal choice for supporting the next generation of high-performance computing and AI workloads.

References

[1] Zhenguo Wu, et al., SiP Architecture For Accelerating Collective Communication in Distributed Deep Learning, *Optica Publishing Group*, 2023.

[2] Zhenguo Wu, et al., Peta-Scale Embedded Photonics Architecture for Distributed Deep Learning Applications, *Journal Of Lightwave Technology*, vol. 41, no. 12, June 15, 2023.

[3] Samuel Williams, et al., Roofline: An Insightful Visual Performance Model for Multicore Architectures, *Communications of the ACM*, vol. 52, no. 4, Apr. 2009. https://doi.org/10.1145/1498765.1498785

[4] Y. Wang *et al*., Silicon Photonics Chip I/O for Ultra High-Bandwidth and Energy-Efficient Die-to-Die Connectivity, *2024 IEEE Custom Integrated Circuits Conference (CICC)*, Denver, CO, USA, 2024, pp. 1-8, doi: 10.1109/CICC60959.2024.10529018.

Acknowledgments

I would like to thank Professor Alex Meng, Postdoctoral Researcher Yuyang Wang, and PhD candidate Brian Wu for their continued support throughout the Columbia-Amazon SURE program. I would also like to thank Columbia University for giving me the opportunity to conduct undergraduate research at the Lightwave Research Laboratory.