Silicon photonics enabling high bandwidth density, energy efficient interconnects

M. Glick, Y.-H. Hung, K. Bergman, Columbia University

The bottleneck in today's high performance computing (HPC) and data center systems has shifted from computation to communication. This trend is increasing [1] and is expected to accelerate with the expansion of compute time spent devoted to machine learning and data analytics. As machine learning models and data sets become larger the total training time becomes dominated by the communication time which results in diminishing returns with additional workers.

Over the past few decades, silicon photonics, leveraging silicon electronics manufacturing infrastructure, has emerged as a promising technology to both increase bandwidth density and reduce energy per bit in all layers of the system [2]. The ability to share much of the same infrastructure as silicon processes used for conventional complementary metal oxide semiconductor (CMOS) logic enables a path towards scalable manufacturing, reducing cost and co-integration of optical with electrical components,

Although silicon photonics can be exploited to significantly increase interconnect bandwidth density in the near term, silicon photonics can have substantially more impact in concert with architectural changes enabling dynamical adjustment of connectivity. Silicon photonics and in particular silicon photonic switches [3] can be utilized to dynamically reconfigure the network (bandwidth steering) to match the dynamics of traffic patterns and thereby improve performance and also improve resource utilization. For example, a key advantage of bandwidth steering would be to enable operators to aggressively oversubscribe top-layer bandwidth in hierarchical tree structures without performance penalties. Figure 1 [1] shows results of a testbed experiment with execution time performance savings of 69%. Similar execution time savings have been shown with different network topologies (dragonfly, HyperX).

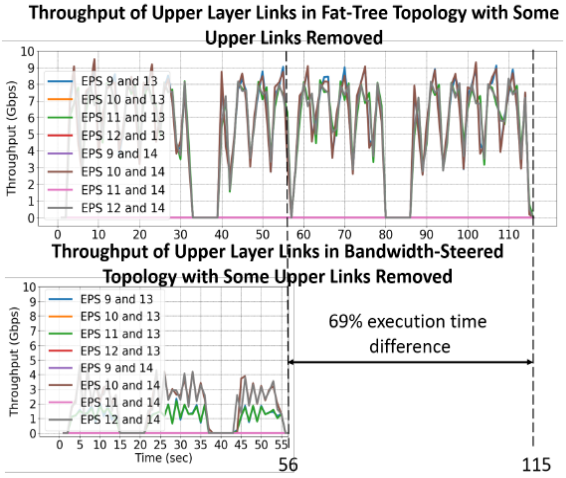

Figure 1. From Ref. [1]

Photonic interconnects using bandwidths steering can also enable disaggregated networks to reconfigure compute and memory resources to meet application requirements more efficiently than those using fixed resource configurations. Memory resources can be pooled and connected to other resources using reconfigurable optical switch fabrics improving performance and increasing scalability. The feasibility of integrating photonic switched optically connected memory into processing systems to address memory challenges in deep learning has been investigated, showing dynamic allocation of additional memory to the processing system and a constant reconfiguration latency. The experimental testbed demonstrates real memory transactions between the processing system and remote memory nodes. [4]

Silicon photonics has great potential to enhance performance, power efficiency and resource utilization in computer networks. Many research challenges remain for optimal device and network integration. to implement the technology and make it a commercial reality.

References

1. G. Michelogiannakis et al., "Bandwidth Steering in HPC using Silicon Nanophotonics", SC19.
2. Shen et al., "Silicon Photonics for Extreme Scale Systems," in Journal of Lightwave Technology, vol. 37, no. 2, pp. 245-259, 15 Jan.15, 2019.
3. Cheng, Qixiang, et al. "Recent advances in optical technologies for data centers: a review." Optica 5.11 (2018): 1354-1370.
4. Z. Zhu et al., "Photonic Switched Optically Connected Memory: An Approach to Address Memory Challenges in Deep Learning," accepted to Journal of Lightwave Technology