

Accelerate Distributed Deep Learning with a Fast Reconfigurable Optical Network

Wenzhe Li⁽¹⁾, Guojun Yuan^(1,*), Zhan Wang⁽¹⁾, Guangming Tan⁽¹⁾, Peiheng Zhang^(1,2), George N. Rouskas⁽³⁾

⁽¹⁾Institute of Computing Technology, CAS, Beijing, China

⁽²⁾Institute of Intelligent Computing Technology, CAS, Suzhou, China

⁽³⁾Department of Computer Science, North Carolina State University, Raleigh, NC

*yuanguojun@ncic.ac.cn

Abstract: We propose a fast-reconfigurable and scalable optical network architecture, which employs a flow-based transmit scheduling scheme to accelerate data parallelism in distributed deep learning. Experimental results demonstrate that the 4-node prototype achieves training times comparable to those of ideal electrical switching. © 2023 The Author(s)

1. Introduction

The success of deep learning is attributed to the large training datasets and large-scale models, which demand substantial computational resources [1]. Therefore, distributed deep learning (DDL) training has become a popular solution to improve model performance. Among the parallel mechanisms for DDL, data parallelism is a typical and widely employed one [2]. However, data-parallel training requires extensive synchronization of parameter updates, resulting in significant communication overhead, often accounting for as much as 90% of the computational time [3]. Thus, communication has emerged as a primary bottleneck in large-scale DDL frameworks. The availability of high-speed interconnects such as NVLink has allowed intra-node networks to achieve high throughput. Consequently, inter-node architectures have emerged as the main bottleneck of the performance of DDL due to their much lower bandwidth and higher latency. These shortcomings may be overcome by recent advancements in photonic technology which enable the provision of ample bandwidth for inter-node communication [4]. By introducing intelligent dynamic topology reconfiguration mechanisms, it becomes possible to fully utilize network resources and minimize latency overhead.

Traditional high-port optical circuit switches (OCS) are relatively slow to reconfigure, often used as supplementary elements in electrical networks [5, 6]. While OCS reduces congestion, its performance is constrained by electrical network bandwidth. Fine-grained transmit scheduling with system-level network reconfiguration times in the order of microseconds or nanoseconds, achieved through fast wavelength switching technology, overcomes this limitation. This paper introduces ODDL, an all-optical switching architecture for DDL, which improves performance with fine-grained topology reconfiguration. ODDL leverages fast optical switch technology for flow-level network reconfiguration, while leveraging slow reconfiguration for application-level reconfiguration. We constructed a four-node testbed, and the experimental results demonstrate that ODDL can achieve training times comparable to those of an ideal electrical switching network.

2. The ODDL Architecture

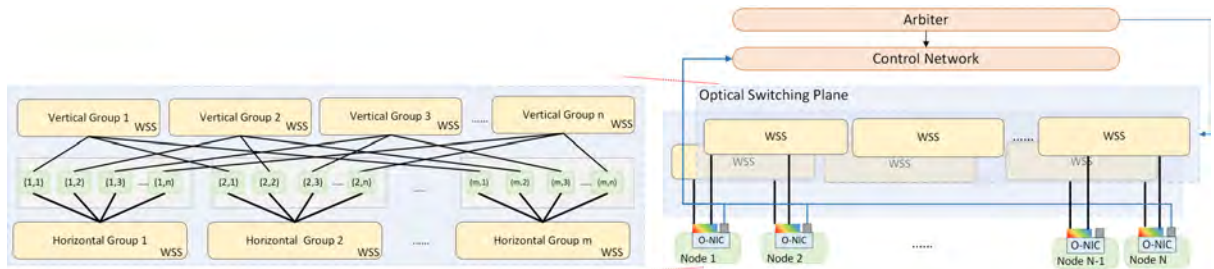


Fig. 1. The overall architecture of ODDL.

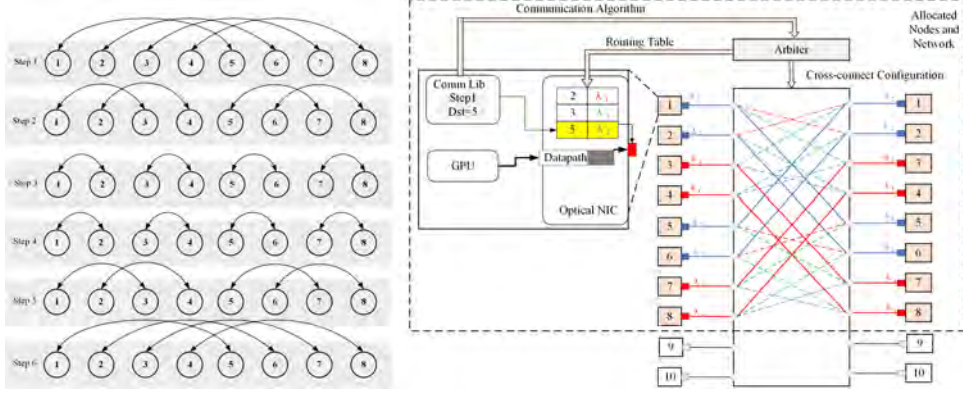


Fig. 2. Reconfiguration scheme.

Ring and Halving Doubling (HD) Allreduce are the most widely used Allreduce algorithms. Ring Allreduce necessitates a total of $2(N - 1)$ steps, where N is the number of nodes. This linear scaling results in suboptimal communication efficiency as the system expands. In contrast, as illustrated in Fig. 2, HD Allreduce exhibits a more efficient behavior, requiring only $2 \cdot \log_2 N$ steps. However, in conventional fat-tree networks, certain long-distance communications may encounter significant contention issues within multi-level electrical switches.

To tackle the challenges posed by this network bottleneck, we introduce the ODDL architecture in Figure 1. ODDL features a highly dynamic all-optical switching network designed for distributed training. In particular, it employs flow-based transmit scheduling and incorporates rapid optical link reconfiguration to ensure seamless scalability. Additionally, ODDL leverages a two-dimensional interconnection topology to enhance performance and improve scalability. Specifically, each node is individually connected to a horizontal switch and a vertical switch, and nodes in the same horizontal or vertical group are connected with one optical switch. Therefore, a 1024-node system requires 64 32-port WSSs and a 4096-node system can be implemented with 128 64-port WSSs.

In our ODDL framework, we employ Liquid Crystal on Silicon (LCoS) based WSS and tunable transceivers to facilitate dynamic topology reconfiguration at various granularities, as illustrated Fig. 2. When a new training task begins, the WSS is reconfigured under the control of an arbiter to interconnect all allocated nodes. After the reconfiguration of the WSS, nodes are linked using different wavelengths to create a subnetwork. To minimize reconfiguration overhead, the optical switch remains static until the training task is completed. Throughout training, the optical circuit is reconfigured by tuning transceiver wavelengths. In Fig. 2, Node 1 connects sequentially to Nodes 5, 3, 2, 3, and 5 by adjusting wavelengths from λ_2 to λ_3 , λ_1 , λ_3 , and λ_2 . Wavelength switching in nanoseconds fulfills the speed requirement for flow-based topology reconfiguration. With this flow-base topology reconfiguration, one-hop routing can be achieved. Furthermore, when the system is scaled to $n \times m$ nodes, where $n = 2^i$ and $m = 2^j$, all communication is completed with one-hop horizontal or vertical routing.

Due to the fact that ODDL reconfigures the topology on a per-flow basis, the control plane is able to respond rapidly to communication requests. At the same time, the control overhead should grow slower than linearly with system size. To this end, we developed a distributed control mechanism whereby each control unit manages its own transmission state and link state, and only exchanges control information with the destination nodes. The central arbiter only handles the network reconfiguration before training, which is the first stage of reconfiguration. Moreover, the control mechanism of wavelength reconfiguration (the second stage) is implemented in the control unit in each Optical-NIC (O-NIC).

3. Prototype and Preliminary Experimental Results

We built a 4-node prototype to validate the feasibility of ODDL and compare its performance with an *ideal* electrical switching network. We use ImageNet as the dataset and HD as the communication algorithm. As shown in Fig. 3, we train model Resnet50 on 4 servers equipped with Ascend 910 GPU, and each server is connected to one FPGA board, which serves as an optical network interface. One of the ports on the FPGA board is connected to WSS with a customized 100Gbps tunable transceiver, which is based on a laser array and fast optical switches to support fast wavelength switching. As shown in Fig. 4, the speed of photonic switches used for the tunable laser is measured, which determines the speed of wavelength switching of the transceiver. The switching delay is around 10.5 ns, which is fast enough to satisfy the flow-based reconfiguration requirement, further proving ODDL's feasibility. Another port on the FPGA board is connected to an Ethernet switch with a commodity 100Gbps

transceiver for the transmission of the control message. The function of the central arbiter is implemented using one of the FPGA boards. Two working wavelengths are $\lambda_1 = 1565.941$ nm and $\lambda_2 = 1546.5377$ nm. The one-tier electrical network connects all nodes with one 100 GE switch, representing an *ideal* network architecture.

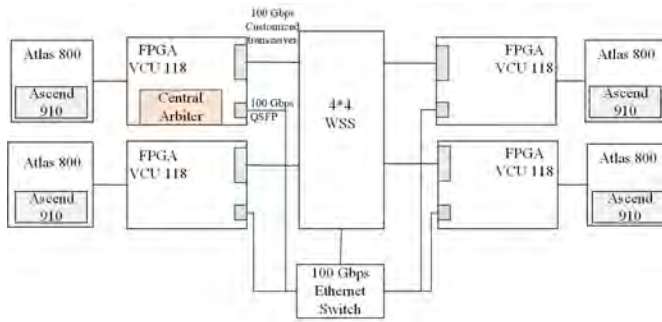


Fig. 3. The diagram of four-node ODDL prototype.




Fig. 4. The switching speed.

Table 1. The comparison of the four-node prototype.

Network	Training time (min)	Accuracy (%)
ODDL	97.8	80.2
One-tier electrical network	97.3	80.2

Table 1 shows the training time of 20 epochs of Resnet50 with different networks. ODDL achieves performance equivalent to that of an ideal electrical switching network, demonstrating that ODDL is efficient and feasible. This outstanding performance can be attributed to ODDL’s low network latency, which remains consistently low even with an increasing network size. Our ongoing work involves the development of a large-scale prototype, which aims to showcase the advantages of ODDL over state-of-the-art interconnection architectures.

4. Conclusion

In this paper, we presented a scalable and fast all-optical network for distributed training, along with a distributed control plane that enables fine-grained scheduling with low control overhead. The 4-node ODDL prototype matches the performance of an ideal electrical switching network, demonstrating the potential of ODDL.

Acknowledgement

This work is partly supported by the National Key Research and Development Program of China (No. SQ2021YFB3000082), NSFC (No.61972380), Jiangsu Science and Technology Project(No. BE2022051-2), and US National Science Foundation (CNS-1907142).

References

1. R. Mayer and H.-A. Jacobsen, “Scalable deep learning on distributed infrastructures: Challenges, techniques, and tools,” *ACM Comput. Surv. (CSUR)* **53**, 1–37 (2020).
2. G. Wang, S. Venkataraman, A. Phanishayee, J. Thelin, N. Devanur, and I. Stoica, “Blink: Fast and generic collectives for distributed ml,” *arXiv preprint arXiv:1910.04940* (2019).
3. D. Narayanan, A. Harlap, A. Phanishayee, V. Seshadri, N. R. Devanur, G. R. Ganger, P. B. Gibbons, and M. Zaharia, “Pipedream: Generalized pipeline parallelism for dnn training,” in *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, (Association for Computing Machinery, New York, NY, USA, 2019), SOSP ’19, p. 1–15.
4. R. Meade, S. Ardalan, M. Davenport, J. Fini, C. Sun, M. Wade, A. Wright-Gladstein, and C. Zhang, “Teraphy: A high-density electronic-photonic chiplet for optical i/o from a multi-chip module,” in *2019 Optical Fiber Communications Conference and Exhibition (OFC)*, (2019), pp. 1–3.
5. N. Farrington, G. Porter, S. Radhakrishnan, H. H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, and A. Vahdat, “Helios: a hybrid electrical/optical switch architecture for modular data centers,” in *Proceedings of the ACM SIGCOMM 2010 Conference*, (2010), pp. 339–350.
6. G. Michelogiannakis, Y. Shen, M. Y. Teh, X. Meng, B. Aivazi, T. Groves, J. Shalf, M. Glick, M. Ghobadi, L. Dennison, and K. Bergman, “Bandwidth steering in hpc using silicon nanophotonics,” in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, (ACM, New York, NY, USA, 2019), SC ’19.