

InfinitePOD: Building Datacenter-Scale High-Bandwidth Domain for LLM with Optical Circuit Switching Transceivers

Chenchen Shou^{1,2,3} Guyue Liu^{1,†} Hao Nie^{2,†} Huaiyu Meng^{3,†} Yu Zhou²
Yimin Jiang⁴ Wenqing Lv³ Yelong Xu³ Yuanwei Lu² Zhang Chen³
Yanbo Yu² Yichen Shen³ Yibo Zhu² Daxin Jiang²

¹Peking University ²StepFun ³Lightelligence Pte. Ltd. ⁴Unaffiliated

ABSTRACT

Scaling Large Language Model (LLM) training relies on multi-dimensional parallelism, where High-Bandwidth Domains (HBDs) are critical for communication-intensive parallelism like Tensor Parallelism (TP) and Expert Parallelism (EP). However, existing HBD architectures face fundamental limitations in scalability, cost, and fault resiliency: switch-centric HBDs (e.g., NVL-72) incur prohibitive scaling costs, while GPU-centric HBDs (e.g., TPUv3/Dojo) suffer from severe fault propagation. Switch-GPU hybrid HBDs such as TPUv4 takes a middle-ground approach by leveraging Optical Circuit Switches, but the fault explosion radius remains large at the cube level (e.g., 64 TPUs).

We propose *InfinitePOD*, a novel transceiver-centric HBD architecture that *unifies connectivity and dynamic switching at the transceiver level* using Optical Circuit Switching (OCS). By embedding OCS within each transceiver, *InfinitePOD* achieves reconfigurable point-to-multipoint connectivity, allowing the topology to adapt into variable-size rings. This design provides: i) datacenter-wide scalability without cost explosion; ii) fault resilience by isolating failures to a single node, and iii) full bandwidth utilization for fault-free GPUs. Key innovations include a Silicon Photonic (SiPh) based low-cost OCS transceiver (**OCSTrx**), a reconfigurable k-hop ring topology co-designed with intra-/inter-node communication, and an HBD-DCN orchestration algorithm maximizing GPU utilization while minimizing cross-ToR datacenter network traffic. The evaluation demonstrates that *InfinitePOD* achieves **31%** of the cost of NVL-72, **near-zero** GPU waste ratio (over one order of magnitude lower than NVL-72 and TPUv4), **near-zero** cross-ToR traffic when node fault ratios under 7%, and improves Model FLOPs Utilization by **3.37x** compared to NVIDIA DGX (8 GPUs per Node).

1 INTRODUCTION

Large-scale Large Language Models (LLMs) training rely on various parallelism strategies [67, 72], such as Tensor Parallelism (TP), Expert Parallelism (EP), Data Parallelism (DP), Pipeline Parallelism (PP), Context Parallelism (CP) and Sequence Parallelism (SP). These strategies communicate over two types of AI datacenter compute fabrics, each with distinct bandwidth requirements. First, Datacenter Networks (DCNs) provide hundreds of Gbps per GPU and primarily handle DP, PP, CP, and SP traffic, which has lower communication demands. Second, High-Bandwidth Domains (HBDs) offer Tbps-level bandwidth, which is crucial for communication-intensive TP and EP. Efficient HBD design can reduce communication overhead, thereby improving Model FLOPs Utilization (MFU)-a key performance metric for LLM training.

The community has made significant advancements in designing DCNs for LLM training [24, 53, 65, 79]. However, scaling HBD to optimize MFU in LLM training remains a challenging problem. Existing HBD architectures [33, 34, 56, 74, 75] take important steps but still suffer from fundamental limitations in scalability, cost, and fault resiliency.

- **Switch-centric HBDs**, such as NVIDIA NVL-72 [56], build multilayer nonconvergent networks for HBD with switch chips. However, the switch fabric incurs superlinear cost growth as it scales, constraining the number of GPUs per HBD. This limitation prevents optimal large TP and EP and causes severe *resource fragmentation* when the size of TP/EP group increases. For instance, with 2 HBDs (32 GPUs each), 30 GPUs are wasted for TP-16 jobs if each HBD has a single GPU failure. This waste reduces to 14 GPUs if the two HBDs are combined into a 64-GPU unit.
- **GPU-centric HBDs**, such as Dojo [74], NVIDIA V100 [11], TPUv3 [34], and SiP-Ring [35], adopts low-cost GPU-to-GPU links to construct large-scale ring or mesh topologies, forwarding traffic directly through GPUs. However,

[†]Guyue Liu, Hao Nie, and Huaiyu Meng contributed equally to this work and share the corresponding authorship.

these architectures suffer from a large *fault explosion radius*, where a single GPU failure degrades bandwidth for a group of adjacent GPUs, compromising the entire topology. For example, in SiP-Ring, one single GPU failure breaks the ring and reforms the topology a line.

- **Switch-GPU Hybrid HBDs**, TPUv4 [33] alleviates the limitation of large fault explosion radius via OCS-based switch¹: each set of 64 TPUs is connected as a cube, with these cubes connected to multiple OCS-based switches to isolate faults within the cubes. However, it cannot resolve it fundamentally, as the fault explosion radius remains large at the cube level (64 TPUs).

In this paper, we take a first-principles approach to re-designing HBD for LLM training workloads. Through a top-down analysis of parallelism strategies (§2.3) for maximizing MFU, we show that increasing TP size yields the most significant MFU gains for both dense and sparse LLM models. For large dense models [42], the optimal TP size scales from 16 to 64 as the number of GPUs increases. For sparse MoE models [12, 30, 73], enlarging TP improves MFU more effectively than EP, particularly when considering the expert imbalance problem [41].

These findings lead to two key design principles: i) HBD should be optimized exclusively for TP Ring-Allreduce communication with large message sizes, which communicates with only logical neighboring nodes, eliminating the need for EP and the associated any-to-any communication; ii) Supporting large and adaptable TP is essential, as different GPU numbers and model sizes require varying TP configurations to maximize MFU.

Based on these principles, we propose *InfinitePOD*, a scalable and fault-resilient HBD architecture designed for optimizing TP communication. Our key insight is *unifying connectivity and dynamic switching at the transceiver level* using OCS. By embedding OCS in each transceiver, we achieve reconfigurable point-to-multipoint connectivity. This marks a departure from traditional designs, where transceivers support only point-to-point connections and rely on high-radix switches for routing. We call this new design **transceiver-centric HBD architecture**. This transceiver-centric architecture offers two key benefits: i) It enables the flexible construction of arbitrarily large ring topologies by intra-node loopback mechanism. This can support optimal TP group sizes for different models, while effectively minimizing resource fragmentation; ii) When one node fails, its neighboring transceivers dynamically reconfigure connections to reroute traffic, significantly reducing the fault explosion radius and improving system resilience.

¹In this paper, "OCS" specifically denotes *optical circuit switching capability*, while OCS-based switch denotes *optical circuit switch*.

We realize the transceiver-centric HBD architecture in production by combining the following key ideas:

- **Silicon Photonics based OCS transceiver (OCSTrx)**: To design a cost-effective low-power transceiver with OCS support, we leverage the current advances of Silicon Photonics (SiPh) technology. Compared to MEMS [76, 85] technology which has been widely used to realize OCS, SiPh offers simpler structures, lower cost and power consumption. We build OCS with Mach-Zehnder interferometer (MZI) matrix [82], taped out with 65nm CMOS processes. The chip size is less than 136.5mm² while the chip power consumption is 3.2Watts, which can be integrated into commercial QSFP-DD 800Gbps transceiver [48] with sub-1ms path reconfiguration latency.
- **Reconfigurable K-Hop Ring Topology**: While OCSTrx offers reconfigurable connections at the transceiver level, constructing adaptive-size rings that maximize GPU utilization remains a challenge. For example, a naive full-mesh topology built with OCSTrx would impose strict limits on TP size, while also resulting in significant bandwidth waste and fragmentation. To address this, we propose a reconfigurable K-Hop Ring, where each node connects to all other nodes within $\leq K$ hops via OCSTrx. The intra-node *loopback* mechanism enables dynamic ring construction, while the inter-node *backup link* bypasses faulty nodes, ensuring high fault tolerance.
- **HBD-DCN Orchestration Algorithm**: While an optimal HBD topology is critical, end-to-end training performance also depends on efficient HBD-DCN coordination. For example, the orchestration of TP groups in HBD directly determines DP traffic distribution, which impacts congestion in DCN, ultimately governing training performance. Unfortunately, existing approaches lack mechanisms to jointly coordinate DCN and HBD to mitigate congestion and optimize communication efficiency. To address this, we propose a new orchestration algorithm that minimizes cross-ToR traffic, thereby minimizing congested traffic.

To the best of our knowledge, *InfinitePOD* is the first HBD design capable of scaling to datacenter scale while avoiding cost explosion and increased failure-induced waste. We evaluated *InfinitePOD* with the real 160-day fault trace from our 3K GPU cluster². When executing TP32 jobs with the trace, it demonstrates 0.44% GPU waste ratio - 22x and 23x lower than NVL-72 (9.67%) and TPUv4 (10.30%). It achieves 3.24x and 1.59x cost reductions compared to NVIDIA NVL-72 and Google TPUv4 respectively. Through the orchestration algorithm, it maintains near-zero cross-ToR traffic under 7%

²Details in Appendix §A. We will open-source the trace later.

node failure rates. Its dynamic ring formation capability enables 3.37x higher MFU than NVIDIA DGX systems [63] (8 GPUs/node).

2 BACKGROUND AND MOTIVATION

In this section, we first introduces LLM training in AI datacenters (DCs) (§2.1). Then, we examine existing High-Bandwidth Domain (HBD) architectures and discuss their limitations (§2.2). Finally, we summarize key design principles of HBD for LLM training (§2.3).

2.1 LLM Training in AI DC

LLM training parallelism and communication. LLM training jobs employ various parallelism strategies to efficiently utilize GPUs distributed across AI DCs [67, 71]. Based on communication loads, parallelism can be categorized into two types. The first type is *communication-intensive parallelism* which involves high communication load. Tensor Parallelism (TP) splits the model across multiple GPUs and synchronizes via AllReduce. The ring algorithm for AllReduce is theoretically optimal [61], making ring-based topologies ideal for TP. Expert Parallelism (EP), designed for Mixture of Experts (MoE) models [12, 30, 73], assigns experts to different GPUs and relies on AlltoAll communication, requiring topologies with high bisection bandwidth (e.g., Full-Mesh). In contrast, parallelism strategies such as Data Parallelism (DP), Pipeline Parallelism (PP), Context Parallelism (CP), and Sequence Parallelism (SP) introduce lower communication overhead, placing less demands on network performance.

Compute fabric. Compute fabric in AI DC interconnects GPUs to efficiently transmit model gradients and parameters. It consists of two primary components: Datacenter Network (DCN) and High-Bandwidth Domain (HBD). DCN provides communication across the entire AI DC via Ethernet or Infiniband, the bandwidth is around 200 ~ 800Gbps. Widely used DCN architectures include Fat-Tree [2] and Rail-Optimized [53]. In comparison, HBD offers Tbps-level throughput, and is more suitable for TP/EP. However, its scale is typically constrained by interconnection costs and fault tolerance considerations. For example, NVL-72 [56] only interconnects 72 GPUs per HBD.

Faults and fault explosion radius. As revealed by current advances of AI DCs [25, 65], training jobs experience a variety of faults, such as GPU faults, optical transceiver faults, switch faults, and link faults. We quantify the fault impact using the *fault explosion radius*, defined as *the number of GPUs degraded by a single fault event*. The fault explosion radius varies depending on both the system architecture and the fault component. For example, if a switch fails, the bandwidth of all devices connected to it will degrade, illustrating the switch-level fault explosion radius.

HBD fragmentation. When the number of GPUs in the HBD cannot be evenly divided by the size of the parallel group (i.e., TP size), the remaining GPUs become unusable, leading to resource waste. The GPU waste ratio for each HBD can be expressed by the formula $\{(HBD_{size} - N_{fault}) \bmod TP_{size}\} / HBD_{size}$. In AI DCs with small-scale HBDs, GPU waste due to fragmentation is significant because each HBD experiences independent fragmentation. This issue worsens as the TP group size increases with model scale. For example, for NVL-36 shown in Figure 1a, running TP-16 causes $\geq 11\%$ GPU waste ratio.

2.2 Limitations of Existing HBDs

Existing HBD architectures for LLM training can be categorized into three types, based on the key components that provide connectivity. A summary is shown in Table 1.

Switch-centric HBD. This type architecture leverages switch chips to interconnect GPUs, as shown in Figure 1a. A prominent example is NVIDIA, which utilizes NVLink and NVLink Switch [52, 58], e.g. DGX H100 [63] with 8-GPU and GB200 NVL-36, NVL-72, and NVL-576 [56]. These architectures offer high-performance any-to-any communication. However, switch-centric HBDs have several drawbacks: i) They require a large number of switch chips due to their limited per-chip throughput; ii) They are vulnerable to a switch-level fault explosion radius—when a switch chip fails, all connected nodes experience bandwidth degradation; iii) High interconnect costs constrain the scale of HBDs, leading to significant fragmentation when serving large models.

GPU-centric HBD. GPU-centric HBD architectures construct the HBD using direct GPU-to-GPU connections, eliminating the need for switch chips. As a result, cost scales linearly with HBD size. A representative example is SiP-Ring [35], shown in Figure 1b, where GPUs are organized into fixed-size rings. However, this design imposes a strict limitation: the TP group size must remain fixed. To enable communication at dynamic scales and support a wider range of workloads, more complex topologies are adopted (e.g., Dojo [74], NVIDIA V100 [11], TPUv3 [34], and AWS Trainium [75]), which support dynamic scaling by allowing jobs to execute on topology subsets of varying sizes. As shown in Figure 1c, Dojo [74] connects GPUs via mesh-like topologies and employ GPUs to forward traffic. While GPU-centric architectures mitigate cost explosion and can support various scales, they suffer from a large fault explosion radius. A single GPU failure can disrupt the entire HBD by altering its connectivity, degrading communication performance even for healthy GPUs—such as the yellow GPUs in Figure 1c.

Switch-GPU Hybrid HBD. This architecture interconnects GPUs via a combination of direct GPU-to-GPU connections and switch links. A typical example is TPUv4 [33], which

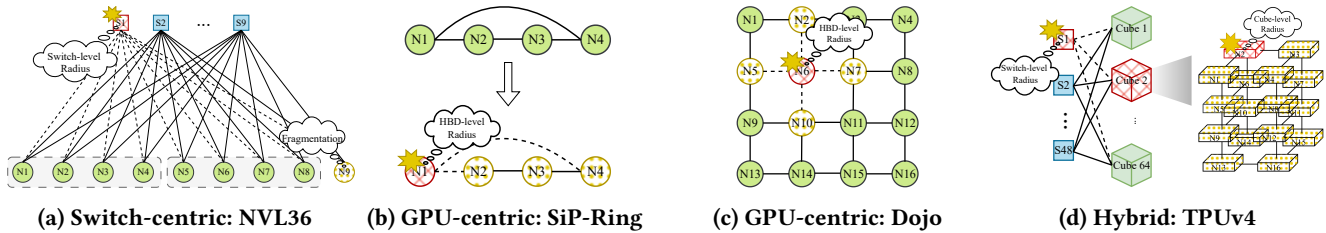


Figure 1: Illustrative examples of HBD architectures. N represents Node, and S represents Switch. Red (with cross hatch) represents fault device and yellow (with dots) represents unavailable or downgraded GPU.

Architecture	Type	Scalability	Collective Primitives	Fault Explosion Radius		Interconnect Cost	Fragmentation
				Node-Side	Switch-Side		
NVL	Switch-centric	Low	Full CCL	Node-level	Switch-level	High	Many
Dojo, TPUv3, SiP-Ring	GPU-centric	High	Ring-Allreduce	HBD-level	✗	Low	Few
TPUv4, TPUv5p	Switch-GPU Hybrid	Moderate	Ring-Allreduce	Cube-level	Switch-level	Moderate	Few
<i>InfinitePOD</i>	Transceiver-centric	High	Ring-Allreduce	Node-level	✗	Low	Few

Table 1: Comparative analysis of HBD architectures.

organizes TPUs into 4^3 TPU cubes and connect them via centralized OCS-based switches (Figure 1d). TPUv4 scales up to 4,096 TPUs, with its expansion primarily limited by the port count of the OCS-based switch. Furthermore, it suffers from a cube-level fault explosion radius—a failure in any single TPU affects the entire 64-TPU cube, leading to significant performance degradation. Furthermore, OCS-based switches face challenges of high costs and manufacturing complexity, which undermines the cost-effectiveness of TPUv4. TPUv5p cluster [26] is similar to TPUv4 but can scale out to 8,960 TPUs.

2.3 Key Attributes of An Ideal HBD

GPU	TP	PP	DP	MFU	MFU _{TP-8}	Improve
1024	16	4	16	0.5236	0.5217	1.0036
4096	16	8	32	0.4668	0.4282	1.0901
8192	32	8	32	0.4247	0.3512	1.2093
16384	32	16	32	0.3756	0.2584	1.4536
32768	32	16	64	0.3090	0.1690	1.8284
65536	64	16	64	0.2493	0.0999	2.4955
131072	64	16	128	0.1851	0.0550	3.3655

Table 2: Optimal parallelism strategy for maximum MFU of Llama 3.1-405b, compared to the baseline MFU for TP-8 (e.g., in widely-deployed NVLink architectures), when GPU number varies.

Existing HBD architectures face fundamental limitations in interconnection cost, resource utilization, and failure resiliency when scaling. To guide a better design, we analyze existing training workloads and explore two key questions without the limitations imposed by current HBD: i) What is the optimal group size that HBD should support? ii) What traffic patterns should HBD accommodate?

Large and adaptable TP size is critical for dense models. The optimal LLM training parallelism depends on model

architectures and cluster configurations. For example, as illustrated by previous work [86, 87]. We evaluate the Model FLOPs Utilization (MFU) for Llama 3.1-405B [46] using our in-house LLM training simulator (§6.3) and report the results in Table 2. MFU and TP/PP/DP columns denote the optimal MFU when TP size is unconstrained and the corresponding parallelism strategies respectively. MFU_{TP-8} column denotes the optimal MFU when TP size is limited to 8. As we increase the number of GPUs, the optimal TP size grows from 16 to 64, a trend we observe across other large dense models. In this case, the HBD scale restricts the maximum size of TP, which affects training performance as a result.

Parallelism	Operation	Traffic Load
TP	AllReduce	$2bsh \cdot \frac{n-1}{n}$
EP	AllToAll	$2bsh \cdot \frac{n-1}{n} \cdot \frac{k}{n}$

Table 3: Communication load of TP and EP on a single MoE layer. b : batch size; s : sequence length; h : hidden dim; k : topK of MoE router; n : parallel size. Assume each expert is assigned equal number of tokens.

MoE can also be efficient with large-size TP. Beyond widely used dense models, we also examine sparse MoE models, which are trending toward larger scales (e.g., 1T parameters [15]). The distributed training for MoE can be achieved through TP or EP (or a combination of them)³ [41], both TP and EP are communication-intensive [36], making them heavily reliant on HBD.

Our production training experience on a 1T MoE model in production brings the following insights into the pros and cons of TP and EP. On the one hand, EP is more communication-efficient than TP. Table 3 compares the communication volume of TP and EP. Clearly, EP is better if $k < n$, which is

³For TP, each expert is equally sharded to GPUs. For EP, each expert is indivisible and allocated to one GPU in the EP group.

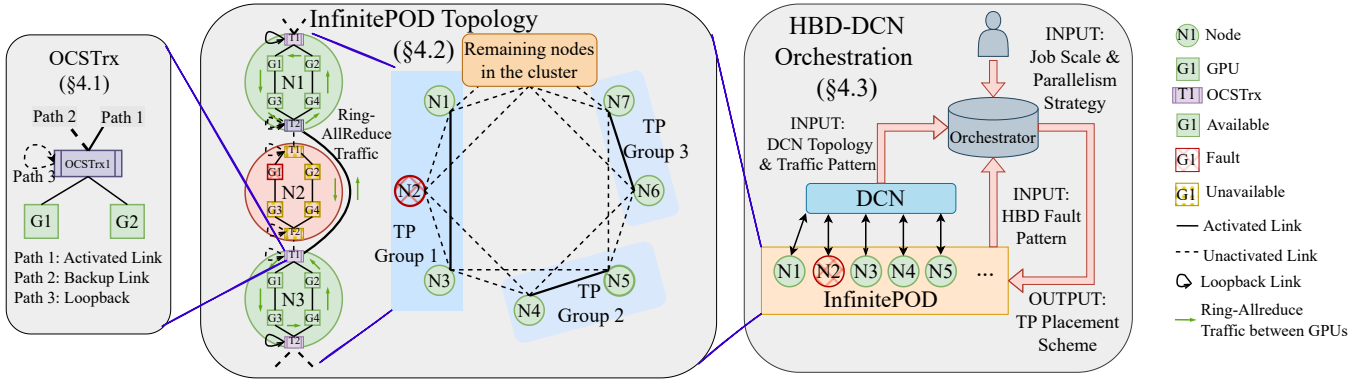


Figure 2: *InfinitePOD* overview.

	TP		EP		
imbalance coef	-	0%	10%	20%	30%
MFU (%)	31.2	31.5	30.5	29.8	28.8

Table 4: Performance comparison of TP and EP when training GPT-MoE.

common [39] because existing models often choose small k for higher computation sparsity. On the other hand, EP suffers from the well-known expert imbalance problem [41], especially when the MoE routers use the no-token-left-behind algorithm [13, 23, 39]. This will result in non-equivalent number of tokens that each expert will receive, which hence causes straggler nodes that waste GPU cycles of other nodes. Table 4 shows the simulated result of training GPT-MoE with 1.1T parameters (details in Appendix §B) under different expert imbalance coefficients⁴. When $coef = 0$, EP is better than TP due to smaller communication overhead. As $coef$ increases, the MFU drops because of the straggler issue.

Key findings. These experiments provide us two key findings for HBD design. First, larger HBD size is increasingly needed for rapidly scaling LLMs (i.e., more than 1T parameters). Second, with larger HBD enabled, using TP is more favorable than EP to train MoE, because TP shards the computation equally across GPUs and hence bypasses the expert imbalance problem.

These findings reveal two key design principles for HBD: i) HBD must inherently support large and adaptable TP sizes, which fundamentally requires the scalability of HBD architecture; ii) the HBD designs need to ensure the effective support for the Ring-AllReduce communication. Given the demonstrated efficiency of TP in MoE training, ensuring support for Ring-AllReduce support is sufficient for mainstream LLM training scenarios; iii) small fault explosion radius. Thus, **we propose designing a large and adaptable HBD architecture tailored for ring-based TP communication to optimize LLM parallelism strategies.**

⁴Calculated as $\frac{max-min}{max}$, where max and min represent the maximum and minimum tokens allocated to each expert respectively.

3 DESIGN OVERVIEW

In this section, we first present our new HBD architecture *InfinitePOD* guided by the design principles outlined above. We then provide an overview of its key components.

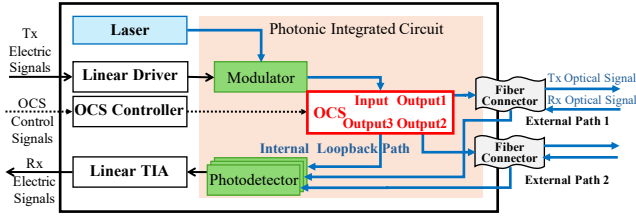
Transceiver-centric HBD architecture. As discussed in §2.2 and summarized in Table 1, existing architectures face a fundamental tradeoff among scalability, cost, and fault isolation. The GPU-centric architecture offers high scalability and low cost connectivity but suffers from a large fault explosion radius. In contrast, the switch-centric architecture improves fault isolation by leveraging centralized switches to confine failures to the node level. However, this comes at the cost of reduced scalability and higher connection overhead. The GPU-switch hybrid architecture takes a middle-ground approach but still suffers from significant fault propagation. As a result, no existing architecture fully meets all requirements.

Our key insight is that *connectivity and dynamic switching can be unified at the transceiver level* using Optical Circuit Switching (OCS). By embedding OCS within each transceiver, we can enable reconfigurable point-to-multipoint connectivity, effectively combining both connectivity and switching at the optical layer. This represents a fundamental departure from conventional designs, where transceivers are limited to static point-to-point links and rely on high-radix switches for dynamic switching. We refer to this novel design as the *transceiver-centric HBD architecture*.

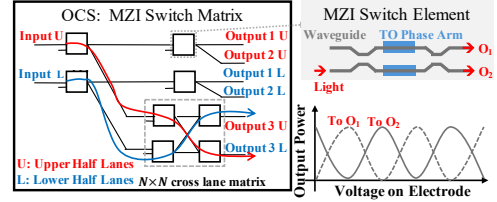
We realize this design with *InfinitePOD*, which has three key components as shown in Figure 2.

Design 1: Silicon Photonics based OCS transceiver (OCSTrx) (§4.1). To enable large-scale deployment, we require a low-cost, low-power transceiver with Optical Circuit Switching (OCS) support. Unlike prior high-radix switches solutions that rely on MEMS-based switching [76, 85], we leverage advances in Silicon Photonics (SiPh), which offer a simpler structure, lower cost, and reduced power consumption—making them well-suited for commercial transceivers.

Our SiPh-based OCS transceiver (OCSTrx), shown on the left of Figure 2, provides two types of communication paths:



(a) Components of OCS transceivers.



(b) Zoom into OCS MZI switch matrix.

Figure 3: Design of OCS Transceivers. The core component is OCS integrated in transceivers.

i) *Cross-lane loopback path (path 3)*, enabling direct GPU-to-GPU communication within the node, which can be used to construct dynamic size topologies; ii) *Dual external paths (path 1&2)*, connecting to external nodes. All these paths utilize time-division bandwidth allocation, featuring sub-1ms switching latency. With this capability, our *OCSTrx* allows dynamic reallocation of full GPU bandwidth to an active external path rather than splitting bandwidth across multiple paths. This eliminates redundant link waste—for instance, activating one external path completely disables the other, ensuring efficient bandwidth utilization.

Design 2: Reconfigurable K-Hop Ring topology (§4.2). With *OCSTrx* that provides reconfigurable connections at the transceiver, the next challenge is designing the topology. A naive starting point is the full-mesh topology [60] which can provide full connectivity among all nodes using *OCSTrx*. However, full-mesh design requires $O(N^2)$ links, inducing prohibitive complexity and cost. To reduce costs while maintaining near-ideal fault tolerance and performance, we prune the full-mesh topology into a K-Hop Ring topology based on traffic locality and fault non-locality (Details in §4.2). Combining the reconfigurability of *OCSTrx*, we propose a *reconfigurable K-Hop Ring topology*, shown in the middle of Figure 2, which consists of two key parts:

i) *Intra-node topology*: dynamic GPU-granular ring construction is enabled by activating loopback paths. For example, while N_1 - N_3 physically form a line topology, activating loopback paths creates a ring between N_1 's GPUs (1-4) and N_3 's GPUs (1-4). This mechanism allows for the construction of arbitrary-sized rings at any location, supporting optimal TP group sizes for different models while effectively minimizing resource fragmentation.

ii) *Inter-node fault isolation*: dual external paths connect to primary and secondary neighbors (e.g., 2-Hop Ring). When a node fails (e.g., N_2), its neighbor (N_1) activates the backup path (N_1 - N_3) to bypass the fault while maintaining full bandwidth, approaching node-level fault explosion radius. §4.2 generalizes this design to $K > 2$.

Design 3: HBD-DCN Orchestration Algorithm (§4.3). Designing an optimal HBD topology is crucial, but end-to-end training performance also depends on the efficient coordination between HBD and DCN. For instance, improper

orchestration of TP groups can cause DP traffic to span across ToRs, resulting in DCN congestion. However, existing methods lack the ability to jointly optimize HBD and DCN coordination to alleviate congestion and enhance communication efficiency. To address this, we propose the HBD-DCN Orchestrator, as shown on the right side of Figure 2. The orchestrator takes three inputs: the user-defined job scale and parallelism strategy, the DCN topology and traffic pattern, and the real-time HBD fault pattern. It then generates the TP placement scheme, which maximizes GPU utilization and minimizes cross-ToR communication within the DCN.

4 INFINITEPOD DESIGN

This section first introduces the innovative design of OCS transceivers (*OCSTrx*) based on Silicon Photonics (SiPh) chips (§4.1), a key enabler for *InfinitePOD*, providing both cost efficiency and reconfigurability. Next, we present the DC-scale *InfinitePOD* topology design (§4.2) based on *OCSTrx*. Finally, we outline the HBD-DCN orchestration algorithm (§4.3), designed to optimize communication efficiency for training jobs.

4.1 SiPh-based OCS transceiver (*OCSTrx*)

The *OCSTrx* is designed for reconfigurable point-to-multipoint connectivity. It incorporates a compact OCS-based switch with three Rx/Tx paths, utilizing the MZI switch [82] micro-structure with thermo-optic (TO) effect [29] phase arms. This OCS-based switch is seamlessly integrated into the Photonic Integrated Chip (PIC) of the transceiver, serving as the MZI switch matrix within the Tx light path, and providing photodetector (PD) modules for each Rx paths.

SiPh-Based OCS. Currently, there are two predominant technological approaches for OCS. Micro Electromechanical systems (MEMS) [76, 85] are attractive for commercial adoption due to supporting large port radix, up to a 320×320 matrix [7]. Another option is SiPh-based OCS. Its structure is simpler and cheaper to manufacture, the limitation is its radix due to optical losses in multistage light path selector. Given that the locality of traffic and external paths number of *OCSTrx* is only two, SiPh-based OCS offers greater advantages.

So we choose the design of MZI micro-structure [82] based SiPh-Based OCS. The basic mechanism of MZI switch elements is controlling the phase difference between light paths

in two phase arms, then directs the output light to specific ports through interference at the output combiner. TO effect is utilized for phase arm control, for better switching latency compared to MEMS.

OCS Micro-Structure Design. As shown in Figure 3a, the initial routing decision is made by two MZI switch elements, determining whether to direct the signal through external output 1&2, or the internal loopback path. Subsequently, an internal $N \times N$ MZI switch matrix is incorporated to facilitate the cross-lane loopback mechanism, exemplified by the blue and red paths. Notably, this design can reduce stages count and light attenuation of output 1&2, while ensures consistent light attenuation for them. The design is implemented on the Photonic Integrated Circuit (PIC) chip.

Transceiver Design. In *OCSTrx*, Tx electrical signal is amplified by linear driver and converted to optical signal by modulators as in Figure 3. One laser is coupled into the PIC as optical source. On the receiving end, multiple photodetectors capture the Rx optical signal from all available paths separately. The output from the activated photodetector is then amplified by a linear transimpedance amplifier (TIA). *OCSTrx* offers significant benefits, including high compactness, low power consumption, and cost-effective mass production.

4.2 InfinitePOD Topology

In this section, we present the *InfinitePOD* topology design (Figure 2) integrating *OCSTrx* that allows all GPUs within datacenter to be connected in a *reconfigurable K-Hop Ring topology*, while supporting dynamic ring construction and high fault tolerance.

Intra-node Topology. The intra-node topology is designed for *dynamic ring construction* and complies with the OCP UBB 2.0 standard [59]. As shown in Figure 4, one node equipped with R GPUs can support R bundles of *OCSTrx*. Each *OCSTrx* bundle is connected to a pair of GPUs, with one GPU linking to the upper-half SerDes and the other to the lower-half. For one group of nodes connected as one line, the two GPU pairs at each end can interconnect with the *OCSTrx* internal loopback path, forming a GPU-level ring. As shown in Figure 2, nodes N_1 and N_3 are connected in a line, where $OCSTrx_1(N_1)$ and $OCSTrx_2(N_3)$ activate the cross-lane loopback path, creating a ring between the 8 GPUs of N_1 and N_3 . During ring construction, only two *OCSTrx* bundles per node are utilized, while the remaining *OCSTrx* operate in loopback mode. These idle *OCSTrx* can be replaced with direct connections, such as DAC links, offering a trade-off between cost and reliability. Figure 5(a,b) shows a 4-GPU node with varying numbers of *OCSTrx* bundles. Note that the topology design in this section utilizes a 4-GPU node as an example, it can be easily scaled for 8-GPU nodes.

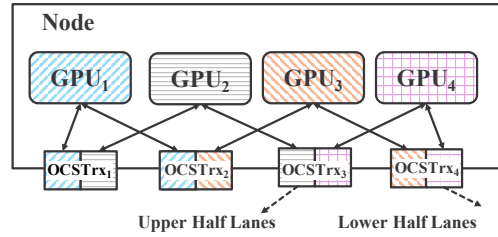


Figure 4: *OCSTrx* connection within nodes. Each block contains multiple *OCSTrx* as one bundle, .e.g, $8 \times 800Gbps$ *OCSTrx* for a 6.4Tbps GPU.

Inter-node Topology. We construct the inter-node topology by pruning the full-mesh design, based on two key observations: i) *Traffic locality*: TP Ring-AllReduce in HBD exhibits neighbor communication patterns, eliminating the need for distant connections; ii) *Fault non-locality*: node-side failures typically occur independently at the node level, meaning consecutive multi-node failures follow an exponentially decaying probability. Each node provides up to $2R$ external paths, allowing us to construct a DC-scale reconfigurable K -Hop Ring topology ($K \leq R$) by connecting them to nodes at $\pm 1, \dots, \pm K, K \leq R$. For AllReduce communication, only two out of the $2K$ links are activated once, with the others serving as backup links for fault isolation. For example (Figure 2), if N_2 fails, $OCSTrx_2(N_1)$ and $OCSTrx_1(N_3)$ can switch to backup links, maintaining connectivity between N_1 and N_3 while isolating N_2 's fault. As K increases, the probability of encountering an unbypassed failure rapidly decreases, which is nearly negligible for $K = 3$ (detailed analysis in Appendix §C). Thus, this architecture typically achieves a node-level explosion radius. Moreover, the K -Hop Ring can be broken into the K -Hop line topology, with the trade-off of reduced fault tolerance of $2K$ nodes at two ends.

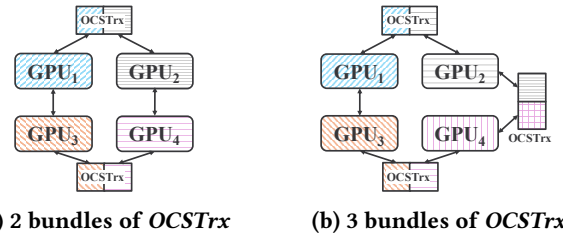


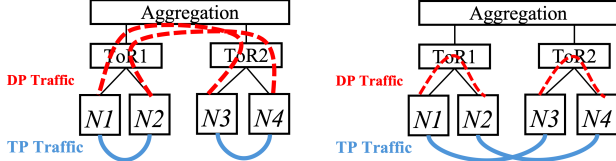
Figure 5: 4-GPU node with *OCSTrx*.

4.3 HBD-DCN Orchestration Algorithm

InfinitePOD is designed to work with arbitrary DCN, including Rail-Optimized [53, 65] and Fat-Tree [2]. This section co-optimizes communication performance for both HBD and DCN in *InfinitePOD*.

Problem Statement. In *InfinitePOD*, GPUs communicate without routing traffic, preventing congestion at any scale. In contrast, DCNs experience inevitable congestion, leading

to performance degradation. To mitigate this, we leverage traffic locality to orchestrate nodes, minimizing cross-ToR traffic. Given a job J requiring N nodes from an available pool of M ($M \geq N$), we must select and order N nodes to satisfy two requirements: (1) nodes in the same TP group should communicate via *InfinitePOD*, and (2) other parallel traffic should minimize congestion. Ideally, communication remains within the same ToR, confining congestion to switch-to-node links.



(a) Orchestration scheme 1. (b) Orchestration scheme 2.
Figure 6: Illustration for problem statement of node orchestration.

A naive approach is sorting nodes based on deployment order in *InfinitePOD*, fulfilling the first requirement but not the second. As shown in Figure 6a, this method places (N_1, N_2) in the same TP group and (N_1, N_3) in the same DP group, forcing DP traffic across ToRs. A better scheme (Figure 6b) eliminates cross-ToR traffic and congestion. However, considering failures and multiple parallel dimensions complicates orchestration, necessitating an efficient method.

Our key insight is to arrange nodes in *InfinitePOD* based on DCN traffic locality, prioritizing appropriate network distances over minimal ones. For example, in Figure 6b, N_1 's *InfinitePOD* neighbor is N_3 , despite a 3-hop network distance in DCN. We propose a two-phase solution: (1) a deployment phase defining physical connections in DCN and *InfinitePOD*, and (2) a runtime phase using an algorithm to orchestrate nodes for arbitrary-scale jobs.

Algorithm 1: Orchestration For Fat-Tree

Input: Topology of DCN and HBD G , Faulty Node Set F , Job Information J .

Output: Placement scheme that satisfies job scale and minimizes cross-ToR traffic.

Create graph $G_{deploy} = \langle S_{deploy}, E_{deploy} \rangle =$
 Deployment-Strategy(G);

Initialize $high = n_{allconstraints}$, $low = 0$, $placement = \{\}$;

while $low \leq high$ **do**

$mid = \lfloor \frac{low+high}{2} \rfloor$;

$placement =$ Placement-Fat-Tree(G_{deploy}, mid, F, J);

if $placement$ satisfies job J **then**

$low = mid + 1$;

else

$high = mid - 1$;

return $placement$

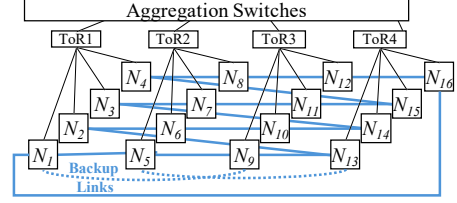


Figure 7: Illustration of the deployment phase, showing only two backup links for simplicity.

Deployment Phase. Figure 7 shows node deployment in HBD and DCN. *InfinitePOD* connects nodes at a network distance of 3 (i.e., cross-ToR). In a DCN with r nodes per ToR, node N_n connects to $N_{n\pm r}$ as main links, while backup links connect to $N_{n\pm 2r}$. For $1 < n \leq r$, N_n connects to $N_{D+n-r-1}$, where D is the total node count (e.g., N_3 connects to N_{14}). Additionally, N_1 may link to the last node, forming a ring.

Runtime Phase. Without considering DCN topology, *InfinitePOD* orchestrates nodes in three steps: (1) identifying cluster faults and modeling healthy nodes as a graph, (2) using Depth-First Search to find connected components, and (3) sequentially placing TP groups within these components. Due to *InfinitePOD*'s topology, each TP group forms a ring.

For real-world DCNs, topology constraints refine step (2) and (3). In Fat-Tree networks, congestion arises when (1) a TP group spans multiple Aggregation-Switch domains, or (2) GPUs within a ToR have mismatched TP group ranks, forcing DP, CP, PP, SP traffic across ToRs. Thus, we aim to localize TP groups within the same Aggregation-Switch domain and align ranks within each ToR. Our scheduling algorithm minimizes cross-ToR traffic while meeting job scale requirements via a binary search over constraint variables. Algorithm 1 outlines the approach, with full details in Appendix §D.

5 IMPLEMENTATION

OCSTrx: We have successfully built a test board featuring the OCS Controller chip and a pre-release Photonic Integrated Circuit (PIC) module without the MZI switch matrix, as shown in Figure 8. The Controller Chip, measuring $4mm \times 4mm$, is manufactured using a 28nm process, while the PIC, sized at $10.5mm \times 13mm$, uses a 65nm CMOS process. The evaluation board supports 8 pairs of TX/RX SerDes at each end and has been validated for compatibility with various link layer protocols, including PCIe (32Gbps, 64Gbps) and Ethernet (56Gbps, 112Gbps). We assessed the power consumption of the peripheral circuitry using the test board. For an $8 \times 112G$ configuration, the power consumption was 8.5 watts. With the addition of 3.2 watts for the MZI switch matrix, the overall consumption totals approximately 12 watts, meeting the QSFP-DD 800Gbps standard[51].

Notably, the full-featured version of the PIC chip has successfully completed tape-out and is currently in the packaging and testing phase. It will be available for evaluation prior to final publication.

Small-scale Cluster: We constructed a small-scale cluster to evaluate the communication performance of the ring topology. Using 32 experimental GPUs equipped with inter-host HBD support (96 lanes on PCIe 4 protocol), we formed a physical ring utilizing fixed optical modules. This mini-cluster was manually reconfigured for both 32-GPU and 16-GPU ring topology. The communication latency and AllReduce performance is evaluated. For small packets, direct GPU-to-GPU links reduced latency by approximately 13% compared to the NVLink switch design. For large packets, the 16-GPU AllReduce utilized 77.11% of the ring bandwidth, with the utilization rate increasing to 77.26% for the 32-GPU configuration, showing minimal degradation with scaling. In comparison, the NVIDIA H100 8-GPU machine achieves an 81.77% utilization rate without SHARP. After deployment of *OCSTrx*, the size of communication group can be reconfigured within 1ms, while maintaining maximum throughput.

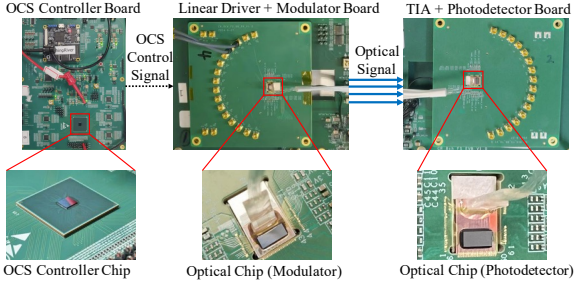


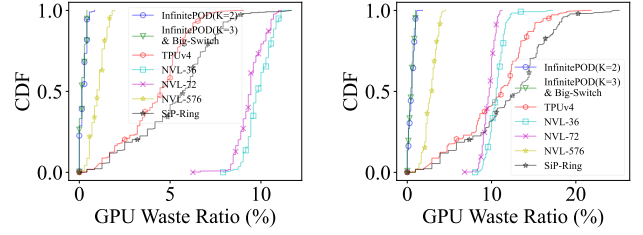
Figure 8: Evaluation board for components of *OCSTrx*.

6 LARGE-SCALE SIMULATION

We begin by outlining the experimental methodology and setup (§6.1). Next, we assess fault tolerance across different HBD architectures (§6.2), followed by end-to-end simulations to evaluate training performance under varying parallelism and GPU resource allocations (§6.3). We then examine the improvements in communication efficiency achieved by our orchestration algorithm (§6.4). Finally, we present a comparative cost and power analysis of different HBD architectures (§6.5). The simulations demonstrate that *InfinitePOD* outperforms other architectures across all metrics.

6.1 Methodology and Setup

An in-house simulator dedicated for LLM training is used to evaluate *InfinitePOD* comprehensively. The simulator supports end-to-end simulations of both model training performance and hardware faults, with the HBD-DCN orchestration algorithm seamlessly integrated into the system.



(a) TP-16.

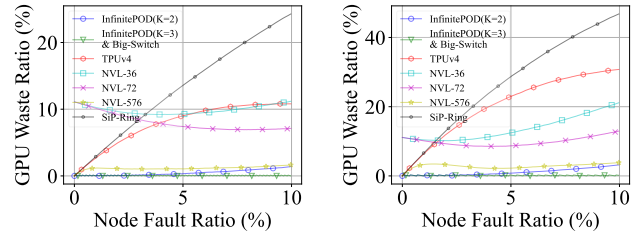
(b) TP-32.

Figure 9: CDF of GPU waste ratio over 4-GPU node based on production fault trace.

GPU and network specification. The NVIDIA H100 [57] (989 TFLOPS, 80GiB) is used for the configuration of GPU in simulation. And the HBD bandwidth of GPU is set as 6.4Tbps, which is the sum of 8 QSFP-DD *OCSTrx*. The DCN bandwidth is configured the same as NVIDIA ConnectX-7 (400Gbps). Since the simulation primarily focuses on HBD, the DCN is configured as a Fat-Tree topology [2]. Several HBD architectures are then evaluated, including:

- **Big-Switch:** The ideal HBD design, featuring a large centralized switch with no forwarding latency that connects all nodes, as the theoretical upper limit of communication performance and fault resilience.
- **InfinitePOD:** Two configurations are evaluated: the *OCSTrx* bundle is set to either $K = 2$ or $K = 3$ (§4.2), constructing 2/3-Hop Ring respectively.
- **NVL-36, NVL-72, NVL-576** [56]: HBDs with 36, 72, or 576 GPUs, GPU are interconnected via NVLink Switches.
- **TPUv4** [33]: Centralized OCS capable of scheduling with a 4^3 TPU cube granularity.
- **SiP-Ring** [35]: All nodes are connected in a series of static rings with fix sizes equal to the TP sizes.

GPU count per node. The simulation aligns with both 4-GPU node (e.g. NVIDIA GB200 NVL-36/72/576 [56] and TPUv4 [33]) and 8-GPU node design (NVIDIA H100, AMD MI300X [3], Intel Gaudi3 [10], and UBB 2.0 standard[59]).



(a) TP-16.

(b) TP-32.

Figure 10: GPU wastes ratio over the 4-GPU node with different GPU fault ratio based on fault model.

Parallelism strategy. Since *InfinitePOD* is primarily designed for TP, the key variable is the TP size. TP-8, TP-16,

TP-32, and TP-64 are tested to evaluate the fault resilience of various HBD architectures (§6.2). Additionally, other parallelism strategies, such as PP and DP, are used to simulate cross-ToR traffic and evaluate the orchestration algorithm (§6.4).

Fault patterns. The fault trace used in the simulation was collected from an 8-GPU node cluster with approximately 3,000 GPUs over a span of 160 days. On average, the ratio of faulty 8-GPU nodes is 3.83%, with the P99 value as 7.22%, more details in Appendix §A. In some simulations, fault traces generated based on this trace statistics are also derived.

6.2 HBD Fault Resilience

This section evaluates the fault resilience of various HBD architectures, focusing on GPU waste ratio, job fault-waiting time, and the maximum job scale supported by the cluster. The main text presents the key results, with more detailed results provided in Appendix §E.

GPU waste. Apart from faulty GPUs, issues such as fragmentation, topology disconnections, and bandwidth degradation can render healthy GPUs wasted. The GPU waste ratio quantifies the number of wasted GPUs under different fault scenarios. Figure 9 illustrates GPU waste ratios over production trace, while Figure 10 depicts the GPU waste ratio as node fault ratio vary.

GPU Num	TP	DP	PP	EP	MFU
1024	16	16	4	1	0.4276
2048	16	16	8	1	0.4140
4096	32	16	8	1	0.3894
8192	32	16	16	1	0.3656
16384	64	16	16	1	0.3116

Table 5: Optimal parallelism strategies for maximize MFU of GPT-MoE under varying GPU numbers.

In these scenarios, *InfinitePOD* ($K = 3$) achieves near-zero GPU waste ratio, and outperforming all other architectures. Especially, the waste ratio for *InfinitePOD* ($K = 2$) remains almost identical to *InfinitePOD* ($K = 3$), allowing one bundle of *OCSTrx* to be saved for clusters with low fault rates. NVL-36 and NVL-72 typically experience an 11% waste ratio for TP sizes of 16 or larger, as 1/9 of GPUs are reserved for redundant backups. NVL-576 has less fragmentation, benefiting from its larger size. TPUv4 performs well at low fault ratios and small TP sizes, but significantly degrades with larger TP sizes due to its coarse 4^3 cube-based resource management, which amplifies the fault explosion radius. To sum up, *InfinitePOD* demonstrates the strongest fault resilience among all architectures.

Maximum job supported. In fixed-size clusters, large job must pause when the available GPUs drop below the required count. Faced with same fault rate, cluster with lower GPU waste ratio can support larger job scales. Figure 11 shows

the maximum job scale supported for various HBD architectures cluster with 2880-GPU, simulated with the fault traces normalized for 4-GPU nodes. *InfinitePOD* ($K = 2$ or $K = 3$) and NVL-576 lead in performance, and SiP-Ring exhibits declining efficiency as TP size increases.

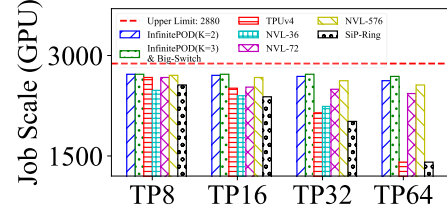


Figure 11: Maximal job scale supported by 2880 GPUs.

Job fault-waiting time. Large job must wait for the repairing when GPU availability falls below the required threshold. This simulations assume the average recovery time in the fault trace as a fixed repair duration. The total wasted time during 160 days is evaluated (Figure 12). For smaller TP sizes (TP-8/TP-16), NVL-36/NVL-72 exhibit the weakest resilience due to their 11% backup overhead. For larger TP sizes (TP-32/TP-64), SiP-Ring and TPUv4 perform worst.

6.3 Training Performance

This section analyzes the training performance of two representative large models, Llama 3.1-405B [42] and GPT-MoE (configuration detailed in Appendix §B), under various GPU resource configurations and parallelism strategies. The simulation results validate the practical applicability of the *InfinitePOD* architecture. In simulations, we model practical TP and EP behaviors: For TP, increasing parallelism splits GEMMs into smaller, less efficient tasks, reducing hardware efficiency [54]; for EP, we practically set expert imbalance coefficient at 20%.

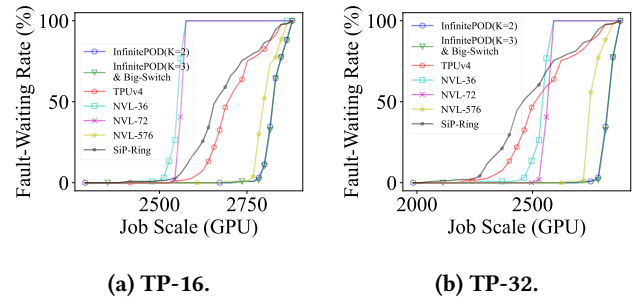


Figure 12: Job fault-waiting time over the 4-GPU node with different levels of job-scale.

LLama 3.1-405B⁵. The model adopts a classical decoder-only Transformer architecture. The simulation employs the

⁵To support larger-scale TP parallelism, we simplified the GQA [1] architecture of Llama 3.1-405B to a traditional MHA architecture.

conventional 3D parallelism strategy⁶, which combines TP, DP, and PP for performance analysis. Table 2 presents the optimal parallelism strategies and their corresponding MFU for LLama 3.1-405B under varying GPU resources. As GPU resources increase, the optimal TP size also increases. When the number of GPUs exceeds 8192, the traditional 8-GPU HBD architecture within a single node begins to limit training efficiency. As the cluster size expands, larger TP sizes become increasingly optimal.

GPT-MoE. The model utilizes the Mixture-of-Experts (MoE) architecture, with $EP \in \{1, 2, 4, 8\}$ introduced in the simulation. Table 5 shows the optimal parallelism strategy and the corresponding MFU for GPT-MoE under various GPU resources. The optimal EP value is 1, suggesting that MoE can also achieve high efficiency with TP.

6.4 Communication Efficiency

This section examines the impact of orchestration algorithms on DCN communication efficiency. Experiments were performed on a Fat-Tree architecture, like the setup in [25]. As shown in Figure 13a, the algorithm is not sensitive to cluster size. Therefore, the evaluation is based on TP-32 operations on *InfinitePOD* with 8192 GPUs.

- **Baseline:** A greedy algorithms, which randomly select nodes from the cluster and use the first permutation that meets the requirements.
- **Optimized:** The HBD-DCN orchestration algorithm proposed in §4.3.

Figure 13b illustrates the impact of job-scale ratios (job size/total cluster GPUs) on cross-ToR traffic, where node fault ratio is 5%. Baseline consistently results in approximately 10% cross-ToR traffic. In contrast, the Optimized algorithm significantly outperforms the Baseline, reducing cross-ToR traffic to just 1.72% even at a 90% job-scale ratio. Figure 13c explores the sensitivity to node faults, with the job scale ratio fixed at 85%. The Baseline shows a linear increase of cross-ToR traffic, while the Optimized algorithm sustains near-zero cross-ToR traffic for fault ratios under 7%.

6.5 Cost and Power Analysis

To evaluate the interconnect costs of HBD architectures, we gather the cost and power information with the following methodologies:

- For standard components (DAC cables, optical transceivers, fibers), pricing is sourced from official retailer websites [17, 18, 50] with a 60% wholesale discount validated against internal data.
- For components with scarce public pricing information, such as Google Palomar OCS, NVIDIA NVLink Switch, 1.6

⁶ $TP \in \{1, 2, 4, 8, \dots, 128\}$, $DP \in \{1, 2, 4, 8, \dots, 1024\}$, $PP \in \{1, 2, 4, 8, 16\}$, $bsz = 2048$

Tbps ACC cables/optical transceivers, the data is amalgamated from multiple sources [68–70] to enhance accuracy.

- Public power consumption data is available for most components, though for NVLink Switch, multiple sources are combined to estimate a reasonable value.

The breakdown analysis of each architecture is provided in the Appendix §F. Based on this, the cost and power consumption are normalized according to GPU count and per-GPU bandwidth. As depicted in Table 6, *InfinitePOD* exhibits the lowest interconnect cost per GPU per GBps. Under the $K = 2$ configuration, its cost is only 62.84% of Google TPUv4 and 30.86% of the NVIDIA GB200 NVL-36/72, with minimal power consumption. This efficiency is primarily attributed to the avoidance of centralized switches. TPUv4 ranks second in interconnect cost and lowest in power consumption, achieved by reducing optical module use and per-port OCS costs. The NVL series has higher interconnect costs and power consumption due to its fully-connected topology and high-cost NVLink Switches. Notably, NVL-576 incurs the highest cost and power consumption due to its multilayer nonconvergent topology, which increases optical module expenses and requires more NVLink Switches.

Architecture	Per-GPU		Per-GPU Per-GBps	
	Cost	Watts	Cost	Watts
TPUv4	1567.20	19.39	5.22	0.06
NVL-36	9563.20	75.95	10.63	0.08
NVL-72	9563.20	75.95	10.63	0.08
NVL-36x2	17924.00	150.33	19.92	0.17
NVL-576	30417.60	413.45	33.80	0.46
<i>InfinitePOD</i> ($K = 2$)	2626.80	48.10	3.28	0.06
<i>InfinitePOD</i> ($K = 3$)	3740.60	72.05	4.68	0.09

Table 6: Interconnect cost (\$) and power (watts).

Beyond interconnect costs, fault resilience variations also affect aggregate costs. The aggregate cost is defined as:

$$Cost_{GPU} \times (N_{Wasted-GPU} + N_{Faulty-GPU}) + Cost_{Interconnect}$$

Simulations on a 11,520-GPU cluster using the TP-32 configuration evaluate GPU availability under varying fault ratios across different architectures. The variation in aggregate cost for different HBD architectures under varying node fault ratios is illustrated in Figure 13d. *InfinitePOD* consistently exhibits the lowest aggregate cost. Furthermore, when the fault ratio is below 12.1%, the aggregate cost of *InfinitePOD* ($K = 2$) is less than that of *InfinitePOD* ($K = 3$), suggesting that ($K = 2$) is the optimal design for most scenarios.

7 DISCUSSION

AllToAll communication. Ring topology in *InfinitePOD* struggles with AllToAll communication (e.g., EP), exhibiting poor performance at $O(p^2)$, where p is the group size. This can be improved by linking backup lines to nodes indexed at

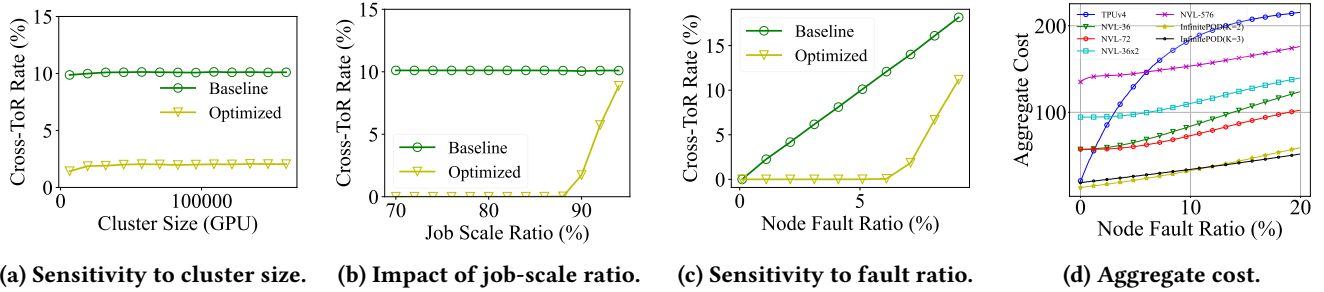


Figure 13: DCN traffic optimization analysis and aggregate normalized cost varies across different architectures under different fault ratios.

$n \pm 2^i$ instead of $n \pm i$ and Applying the Binary Exchange algorithm, and reduce time complexity to $O(p \log_2 p)$. During the algorithm, *OCSTrx* need to connect to different GPUs with runtime switching, since *OCSTrx* switches in under 1ms, reconfiguration can be overlapped with computation. For $K = 2$ *InfinitePOD* designs, performance matches the ideal Bruck algorithm [6] when $p < 8$. However, this design introduces complexities in construction, failover, and orchestration, and necessitates GPU routing capabilities. So it is not applied.

Simulation Scale. Simulations using real 3,000 GPU fault traces, as detailed in §6.2, were conducted on a cluster comprising 2880 GPUs. The reason is that the simulation’s GPU count must be less than the total GPUs in the fault trace, and 2880 is the largest number divisible by 576 and less than 3,000. This configuration allows the entire cluster to be divided into five NVL-576 units for the simulation. Larger scales are simulated in other scenarios.

Multi-dimension parallelism. *InfinitePOD* is optimized for single-dimension parallelism. To support multi dimensional communication, two approaches are viable. 1) *Independent Interconnects*: Each *OCSTrx* bundle includes multiple *OCSTrx* units (e.g., 4 or 8), then link each of the unit to a separate inter-host topology. This isolates parallel dimensions but results in fixed bandwidth per dimension, leading to inefficiencies. 2) *Time-Division multi-dimension*: Main and backup lines of *OCSTrx* can be used to form separated inter-host topology, rapidly switching between them can support multi-dimension parallelism. However, this introduces complexity in managing multi-dimensional overlap and reduces the fault tolerance of *InfinitePOD*.

Single-Job vs. Multi-Job. Existing studies explore multi-job scheduling in GPU clusters [37, 81]. Deploying certain small jobs, such as inference tasks, can mitigate GPU fragmentation. However, given the shortage of GPU in LLM training, any idle GPU—whether repurposed for small job or not—is undesirable. Thus, *InfinitePOD* prioritizes single-job execution for simplicity.

OCS vs. EPS. *OCSTrx* enables multi-path selection, a feature also achievable with Electronic Packet Switching (EPS). For

example, inter-host topology of *InfinitePOD* can be implemented using UBB 2.0-based servers by adding external optical interfaces to switches in server. However, this would require twice optical modules and numerous high-throughput switching chips for the entire system, significantly increasing cost and power consumption compared to *OCSTrx*.

8 RELATED WORK

HBD Architectures. HBDs are crucial for enabling communication intensive parallelism strategies (TP/EP) for LLM training. NVIDIA DGX SuperPOD [55] and GB200 NVL series [56] use any-to-any electrical switching, delivering high performance but suffering from high costs, scalability limitations, and fragmentation. In contrast, direct interconnect HBDs like Dojo [74], TPUv3 [34], and SiP-Ring [35] improve scalability but have a large fault explosion radius. TPUv4 [33] and TPUv5p [26] attempts a middle ground but still lacks full node-level fault isolation. *InfinitePOD* introduces a novel architecture that reduces cost, improves scalability, minimizes fragmentation, enhances fault isolation, and dynamically supports TP.

AI DCN Architectures. MegaScale [32] and Meta’s [25] AI DC use Clos-based topologies, while Rail-Optimized [53] and Rail-Only [79] architectures optimize for LLM traffic patterns. Alibaba HPN [65] enhances fault tolerance with a dual-plane design. *InfinitePOD* is compatible with all of them on LLM-training.

OCS Technologies. OCS enables dynamic topology reconfiguration in datacenters [33, 38, 77]. MEMS OCS-based switch supports high port counts [7, 76], while silicon photonics (SiPh) achieves lower latency and cost [29]. This work proposes a SiPh-based OCS transceiver (*OCSTrx*), constructing an interconnect fabric without centralized switches.

Reconfigurable Networks. Traditional studies [4, 5, 8, 9, 14, 28, 44, 45, 64, 78, 83] focus on generic DCN architectures without optimizing for LLM training traffic, leading to sub-optimal topologies. Recent advancements like SiP-ML [35], TopoOpt [80], and mFabric [38] introduce dedicated training optimizations but still underutilize optical network reconfigurability for better fault tolerance and GPU utilization.

AI Job Schedulers. Schedulers such as [27, 31, 43, 62, 66, 84] aim to improve GPU utilization. However, they exhibit dual limitations: their designs are premised on non-reconfigurable network, while also failing to consider job scheduling within HBD for optimizing traffic patterns in DCN. This work proposes a HBD-DCN orchestration algorithm based on reconfigurable networks to address these limitations.

9 CONCLUSION

In this paper, we propose *InfinitePOD*, a novel HBD design that supports datacenter scale, dynamic TP group size and near-ideal fault explosion radius. *InfinitePOD* is built upon a novel design of optical transceivers interaged with SiPh-based *OCSTrx*, a reconfigurable K-Hop Ring topology and a HBD-DCN orchestration algorithm to leverage the capabilities of the new hardware. Using real fault trace of 3K GPU cluster and the in-house simulator, we demonstrate that *InfinitePOD* achieves GPU utilization close to the ideal model during faults, delivers superior cost and energy efficiency compared to existing designs, and provides effective control over cross-ToR DCN traffic. We believe *InfinitePOD* provides an efficient scaling solution for HBD, which offers new insights for the infrastructure of the next generation Trillion-Parameters LLM training.

REFERENCES

- [1] Joshua Ainslie, James Lee-Thorp, Michiel de Jong, Yury Zemlyanskiy, Federico Lebrón, and Sumit Sanghai. 2023. GQA: Training Generalized Multi-Query Transformer Models from Multi-Head Checkpoints. (2023). arXiv:cs.CL/2305.13245 <https://arxiv.org/abs/2305.13245>
- [2] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. 2008. A scalable, commodity data center network architecture. In *Proceedings of the ACM SIGCOMM 2008 Conference on Data Communication (SIGCOMM '08)*. Association for Computing Machinery, New York, NY, USA, 63–74. <https://doi.org/10.1145/1402958.1402967>
- [3] AMD. 2025. AMD Instinct™ MI300 Series Microarchitecture. (2025). <https://rocm.docs.amd.com/en/latest/conceptual/gpu-arch/mi300.html>
- [4] Daniel Amir, Nitika Saran, Tegan Wilson, Robert Kleinberg, Vishal Shrivastav, and Hakim Weatherspoon. 2024. Shale: A Practical, Scalable Oblivious Reconfigurable Network. In *Proceedings of the ACM SIGCOMM 2024 Conference (ACM SIGCOMM '24)*. <https://doi.org/10.1145/3651890.3672248>
- [5] Hitesh Ballani, Paolo Costa, Raphael Behrendt, Daniel Cletheroe, Istvan Haller, Krzysztof Jozwik, Fotini Karinou, Sophie Lange, Kai Shi, Benn Thomsen, et al. 2020. Sirius: A flat datacenter network with nanosecond optical switching. In *Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication*. 782–797.
- [6] J. Bruck, Ching-Tien Ho, S. Kipnis, E. Upfal, and D. Weathersby. 1997. Efficient algorithms for all-to-all communications in multiport message-passing systems. *IEEE Transactions on Parallel and Distributed Systems* 8, 11 (1997), 1143–1156. <https://doi.org/10.1109/71.642949>
- [7] Calient.AI. 2024. (2024). <https://www.calient.net/>.
- [8] Kai Chen, Ankit Singla, Atul Singh, Kishore Ramachandran, Lei Xu, Yueping Zhang, Xitao Wen, and Yan Chen. 2012. OSA: An Optical Switching Architecture for Data Center Networks with Unprecedented Flexibility. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. USENIX Association.
- [9] Li Chen, Kai Chen, Zhonghua Zhu, Minlan Yu, George Porter, Chunming Qiao, and Shan Zhong. 2017. Enabling Wide-Spread Communications on Optical Fabric with MegaSwitch. In *14th USENIX Symposium on Networked Systems Design and Implementation (NSDI 17)*. USENIX Association, Boston, MA, 577–593. <https://www.usenix.org/conference/nsdi17/technical-sessions/presentation/chen>
- [10] Intel Corporation. 2023. Intel® Gaudi® 3 AI Accelerator White Paper. (2023). <https://www.intel.com/content/www/us/en/content-details/817486/intel-gaudi-3-ai-accelerator-white-paper.html>
- [11] NVIDIA Corporation. 2018. *Accelerated Computing and the Democratization of Supercomputing: Technical Overview*. Technical Report. NVIDIA Corporation.
- [12] DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Haowei Zhang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Li, Hui Qu, J. L. Cai, Jian Liang, Jianzhong Guo, Jiaqi Ni, Jiashi Li, Jiawei Wang, Jin Chen, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Junxiao Song, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Lei Xu, Leyi Xia, Liang Zhao, Litong Wang, Liyue Zhang, Meng Li, Miaojun Wang, Mingchuan Zhang, Minghua Zhang, Minghui Tang, Mingming Li, Ning Tian, Panpan Huang, Peiyi Wang, Peng Zhang, Qiancheng Wang, Qihao Zhu, Qinyu Chen, Qiushi Du, R. J. Chen, R. L. Jin, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, Runxin Xu, Ruoyu Zhang, Ruyi Chen, S. S. Li, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shaoqing Wu, Shengfeng Ye, Shengfeng Ye, Shirong Ma, Shiyu Wang, Shuang Zhou, Shuiping Yu, Shunfeng Zhou, Shuting Pan, T. Wang, Tao Yun, Tian Pei, Tianyu Sun, W. L. Xiao, Wangding Zeng, Wanbiao Zhao, Wei An, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, X. Q. Li, Xiangyue Jin, Xianzu Wang, Xiao Bi, Xiaodong Liu, Xiaohan Wang, Xiaojin Shen, Xiaokang Chen, Xiaokang Zhang, Xiaosha Chen, Xiaotao Nie, Xiaowen Sun, Xiaoxiang Wang, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xingkai Yu, Xinnan Song, Xinxia Shan, Xinyi Zhou, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, Y. K. Li, Y. Q. Wang, Y. X. Wei, Y. X. Zhu, Yang Zhang, Yanhong Xu, Yanhong Xu, Yanping Huang, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Li, Yaohui Wang, Yi Yu, Yi Zheng, Yichao Zhang, Yifan Shi, Yiliang Xiong, Ying He, Ying Tang, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yu Wu, Yuan Ou, Yuchen Zhu, Yudian Wang, Yue Gong, Yuheng Zou, Yujia He, Yukun Zha, Yunfan Xiong, Yunxian Ma, Yuting Yang, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Z. F. Wu, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhen Huang, Zhen Zhang, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhibin Gou, Zhicheng Ma, Zhigang Yan, Zhihong Shao, Zhipeng Xu, Zhiyu Wu, Zhongyu Zhang, Zhuoshu Li, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin Li, Ziwei Xie, Ziyang Song, Ziyi Gao, and Zizheng Pan. 2024. DeepSeek-V3 Technical Report. (2024). arXiv:cs.CL/2412.19437 <https://arxiv.org/abs/2412.19437>
- [13] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*. PMLR, 5547–5569.
- [14] Nathan Farrington, George Porter, Sivasankar Radhakrishnan, Hamid Hajabdolali Bazzaz, Vikram Subramanya, Yeshaihu Fainman, George Papen, and Amin Vahdat. 2010. Helios: a hybrid electrical/optical switch architecture for modular data centers. In *Proceedings of the ACM SIGCOMM 2010 Conference (SIGCOMM '10)*. Association for Computing Machinery. <https://doi.org/10.1145/1851182.1851223>
- [15] William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research* 23, 120 (2022), 1–39.
- [16] fibermall. 2024. OSFP-400G-FR4 400G FR4 OSFP PAM4 CWDm4 2km LC SMF FEC Optical Transceiver Module. (2024). <https://www.fibermall.com/sale-459190-osfp-400g-fr4-cwdm4-2km.htm>
- [17] FIBERMALL.COM. 2024. (2024). <https://www.fibermall.com/>.
- [18] FS. 2024. (2024). <https://www.fs.com/>.
- [19] FS. 2024. 1.5m (5ft) Generic Compatible 400G OSFP Flat Top Passive Direct Attach Copper Twinax Cable. (2024). <https://www.fs.com/products/219579.html>.
- [20] FS. 2024. 1.5m (5ft) NVIDIA/Mellanox MCP1650-V01AE30 Compatible 200G QSFP56 Ethernet Passive Direct Attach Copper Twinax Cable. (2024). <https://www.fs.com/products/155618.html>.
- [21] FS. 2024. 1m (3ft) Generic Compatible 1.6T OSFP Close Top Passive Direct Attach Copper Twinax Cable. (2024). <https://www.fs.com/products/244361.html>.
- [22] FS. 2024. 50m (164ft) Fiber Patch Cable, LC UPC to LC UPC, Duplex, 2 Fibers, Single Mode (OS2), Riser (OFNR), 2.0mm, Tight-Buffered, Yellow. (2024). <https://www.fs.com/products/177394.html>.
- [23] Trevor Gale, Deepak Narayanan, Cliff Young, and Matei Zaharia. 2023. Megablocks: Efficient sparse training with mixture-of-experts. *Proceedings of Machine Learning and Systems* 5 (2023), 288–304.
- [24] Adithya Gangidi, Rui Miao, Shengbao Zheng, Sai Jayesh Bondu, Guilherme Goes, Hany Morsy, Rohit Puri, Mohammad Riftadi, Ashmitha Jeevaraj Shetty, Jingyi Yang, Shuqiang Zhang, Mikel Jimenez Fernandez, Shashidhar Gandham, and Hongyi Zeng. 2024. RDMA

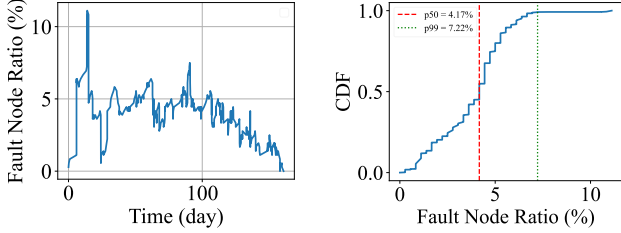
- over Ethernet for Distributed Training at Meta Scale. In *Proceedings of the ACM SIGCOMM 2024 Conference (ACM SIGCOMM '24)*. Association for Computing Machinery, New York, NY, USA, 57–70. <https://doi.org/10.1145/3651890.3672233>
- [25] Adithya Gangidi, Rui Miao, Shengbao Zheng, Sai Jayesh Bondu, Guilherme Goes, Hany Morsy, Rohit Puri, Mohammad Riftadi, Ashmitha Jeevaraj Shetty, Jingyi Yang, Shuqiang Zhang, Mikel Jimenez Fernandez, Shashidhar Gandham, and Hongyi Zeng. 2024. RDMA over Ethernet for Distributed Training at Meta Scale. In *Proceedings of the ACM SIGCOMM 2024 Conference (ACM SIGCOMM '24)*. Association for Computing Machinery, New York, NY, USA, 57–70. <https://doi.org/10.1145/3651890.3672233>
- [26] Google. 2024. TPUv5p. (2024). <https://cloud.google.com/tpu/docs/v5p>.
- [27] Juncheng Gu, Mosharaf Chowdhury, Kang G. Shin, Yibo Zhu, Myeongjae Jeon, Junjie Qian, Hongqiang Liu, and Chuanxiong Guo. 2019. Tiresias: A GPU Cluster Manager for Distributed Deep Learning. In *16th USENIX Symposium on Networked Systems Design and Implementation (NSDI 19)*. USENIX Association, Boston, MA, 485–500. <https://www.usenix.org/conference/nsdi19/presentation/gu>
- [28] Navid Hamedazimi, Zafar Qazi, Himanshu Gupta, Vyas Sekar, Samir R Das, Jon P Longtin, Himanshu Shah, and Ashish Tanwer. 2014. Firefly: A reconfigurable wireless data center fabric using free-space optics. In *Proceedings of the 2014 ACM conference on SIGCOMM*. 319–330.
- [29] Ralf Hauffe and Klaus Petermann. 2006. Thermo-Optic Switching. In *Optical Switching*, Tarek S. El-Bawab (Ed.). Springer US, Boston, MA, 111–139. https://doi.org/10.1007/0-387-29159-8_4
- [30] Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, L el io Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Th eophile Gervet, Thibaut Lavril, Thomas Wang, Timoth ee Lacroix, and William El Sayed. 2024. Mixtral of Experts. (2024). [arXiv:cs.LG/2401.04088](https://arxiv.org/abs/2401.04088) <https://arxiv.org/abs/2401.04088>
- [31] Yimin Jiang, Yibo Zhu, Chang Lan, Bairen Yi, Yong Cui, and Chuanxiong Guo. 2020. A Unified Architecture for Accelerating Distributed DNN Training in Heterogeneous GPU/CPU Clusters. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. USENIX Association, 463–479. <https://www.usenix.org/conference/osdi20/presentation/jiang>
- [32] Ziheng Jiang, Haibin Lin, Yinmin Zhong, Qi Huang, Yangrui Chen, Zhi Zhang, Yanghua Peng, Xiang Li, Cong Xie, Shibiao Nong, Yulu Jia, Sun He, Hongmin Chen, Zhihao Bai, Qi Hou, Shipeng Yan, Ding Zhou, Yiyao Sheng, Zhuo Jiang, Haohan Xu, Haoran Wei, Zhang Zhang, Pengfei Nie, Leqi Zou, Sida Zhao, Liang Xiang, Zherui Liu, Zhe Li, Xiaoying Jia, Jianxi Ye, Xin Jin, and Xin Liu. 2024. MegaScale: Scaling Large Language Model Training to More Than 10,000 GPUs. (2024). [arXiv:cs.LG/2402.15627](https://arxiv.org/abs/2402.15627) <https://arxiv.org/abs/2402.15627>
- [33] Norm Jouppi, George Kurian, Sheng Li, Peter Ma, Rahul Nagarajan, Lifeng Nai, Nishant Patil, Suvinay Subramanian, Andy Swing, Brian Towles, Clifford Young, Xiang Zhou, Zongwei Zhou, and David A Patterson. 2023. TPU v4: An Optically Reconfigurable Supercomputer for Machine Learning with Hardware Support for Embeddings. In *Proceedings of the 50th Annual International Symposium on Computer Architecture (ISCA '23)*. Association for Computing Machinery, New York, NY, USA, Article 82, 14 pages.
- [34] Norman P. Jouppi, Doe Hyun Yoon, George Kurian, Sheng Li, Nishant Patil, James Laudon, Cliff Young, and David Patterson. 2020. A domain-specific supercomputer for training deep neural networks. *Commun. ACM* (Jun 2020), 67–78. <https://doi.org/10.1145/3360307>
- [35] Mehrdad Khani, Manya Ghobadi, Mohammad Alizadeh, Ziyi Zhu, Madeleine Glick, Keren Bergman, Amin Vahdat, Benjamin Klenk, and Eiman Ebrahimi. 2021. SiP-ML: high-bandwidth optical network interconnects for machine learning training. In *Proceedings of the 2021 ACM SIGCOMM 2021 Conference (SIGCOMM '21)*. Association for Computing Machinery, New York, NY, USA, 657–675. <https://doi.org/10.1145/3452296.3472900>
- [36] Jiamin Li, Yimin Jiang, Yibo Zhu, Cong Wang, and Hong Xu. 2023. Accelerating distributed {MoE} training and inference with lina. In *2023 USENIX Annual Technical Conference (USENIX ATC 23)*. 945–959.
- [37] Jiamin Li, Hong Xu, Yibo Zhu, Zherui Liu, Chuanxiong Guo, and Cong Wang. 2023. Lyra: Elastic Scheduling for Deep Learning Clusters. In *Proceedings of the Eighteenth European Conference on Computer Systems (EuroSys '23)*. Association for Computing Machinery, New York, NY, USA, 835–850. <https://doi.org/10.1145/3552326.3587445>
- [38] Xudong Liao, Yijun Sun, Han Tian, Xinchun Wan, Yilun Jin, Zilong Wang, Zhenghang Ren, Xinyang Huang, Wenxue Li, Kin Fai Tse, Zhizhen Zhong, Guyue Liu, Ying Zhang, Xiaofeng Ye, Yiming Zhang, and Kai Chen. 2025. mFabric: An Efficient and Scalable Fabric for Mixture-of-Experts Training. (2025). [arXiv:cs.NI/2501.03905](https://arxiv.org/abs/2501.03905) <https://arxiv.org/abs/2501.03905>
- [39] Aixun Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437* (2024).
- [40] Hong Liu, Ryohei Urata, Kevin Yasumura, Xiang Zhou, Roy Bannon, Jill Berger, Pedram Dashti, Norm Jouppi, Cedric Lam, Sheng Li, Erji Mao, Daniel Nelson, George Papan, Mukarram Tariq, and Amin Vahdat. 2023. Lightwave Fabrics: At-Scale Optical Circuit Switching for Datacenter and Machine Learning Systems. In *Proceedings of the ACM SIGCOMM 2023 Conference (ACM SIGCOMM '23)*. Association for Computing Machinery, New York, NY, USA, 499–515.
- [41] Juncai Liu, Jessie Hui Wang, and Yimin Jiang. 2023. Janus: A unified distributed training framework for sparse mixture-of-experts models. In *Proceedings of the ACM SIGCOMM 2023 Conference*. 486–498.
- [42] Llama Team, AI @ Meta. 2024. The Llama 3 Herd of Models. (2024). [arXiv:cs.AI/2407.21783](https://arxiv.org/abs/2407.21783) <https://arxiv.org/abs/2407.21783>
- [43] Kshiteej Mahajan, Arjun Balasubramanian, Arjun Singhvi, Shivaram Venkataraman, Aditya Akella, Amar Phanishayee, and Shuchi Chawla. 2020. Themis: Fair and Efficient GPU Cluster Scheduling. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. USENIX Association, Santa Clara, CA, 289–304. <https://www.usenix.org/conference/nsdi20/presentation/mahajan>
- [44] William M Mellette, Rajdeep Das, Yibo Guo, Rob McGuinness, Alex C Snoeren, and George Porter. 2020. Expanding across time to deliver bandwidth efficiency and low latency. In *17th USENIX Symposium on Networked Systems Design and Implementation (NSDI 20)*. 1–18.
- [45] William M Mellette, Rob McGuinness, Arjun Roy, Alex Forencich, George Papan, Alex C Snoeren, and George Porter. 2017. Rotornet: A scalable, low-complexity, optical datacenter network. In *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*. 267–280.
- [46] Introducing Llama 3.1: Our most capable models to date. (2024). <https://ai.meta.com/blog/meta-llama-3-1/>.
- [47] OSFP MSA. 2022. The Next Generation of Pluggable Optical Module Solutions from the OSFP MSA. (2022). https://osfpmasa.org/assets/pdf/OSFP1600_and_OSFP-XD.pdf.
- [48] QSFP-DD MSA. 2024. QSFP-DD/QSFP-DD800/QSFP-DD1600 Hardware Specification. (2024). <http://www.qsfp-dd.com/wp-content/uploads/2024/07/QSFP-DD-Hardware-Rev7.1.pdf>.
- [49] NADDOD. 2024. N9500-128QC, 128x400G QSFP112 Ethernet L3 4U Managed Switch, 51.2Tbps, Broadcom Tomahawk 5, Support RoCEv2,

- for AI/ML/Cloud Data Center/HPC. (2024). <https://www.naddod.com/products/102323.html>.
- [50] NADDOD.COM. 2024. (2024). <https://www.naddod.com/>.
- [51] Mark Nowell, Cisco Attila Aranyosi, Vu Le, Jeffery J Maki, Juniper Networks Scott Sommers, Tom Palkert, and Weiming Chen. 2018. QSFDD: Enabling 15 Watt Cooling Solutions. (2018).
- [52] NVIDIA. 2018. NVIDIA NVSwitch Technical Overview. (2018). <https://images.nvidia.com/content/pdf/nvswitch-technical-overview.pdf>.
- [53] NVIDIA. 2021. Doubling all2all Performance with NVIDIA Collective Communication Library 2.12. (2021). <https://developer.nvidia.com/blog/doubling-all2all-performance-wit-h-nvidia-collective-communication-library-2-12/>
- [54] NVIDIA. 2023. Matrix Multiplication Background User’s Guide. (2023). <https://docs.nvidia.com/deeplearning/performance/dl-performance-matrix-multiplication/index.html>.
- [55] NVIDIA. 2024. NVIDIA DGX SuperPOD. (2024). <https://www.nvidia.com/en-us/data-center/dgx-superpod>.
- [56] NVIDIA. 2024. NVIDIA GB200 NVL72. (2024). <https://www.nvidia.com/en-us/data-center/gb200-nvl72/>.
- [57] NVIDIA. 2024. NVIDIA H100 Tensor Core GPU. (2024). <https://www.nvidia.com/en-us/data-center/h100/>.
- [58] NVIDIA. 2024. NVLink and NVLink Switch. (2024). <https://www.nvidia.com/en-us/data-center/nvlink>.
- [59] Open Accelerator Infrastructure (OAI). 2023. Universal Baseboard (UBB) Base Specification r2.0 v1.0. (2023). <https://www.opencompute.org/documents/oai-ubb-base-specification-r2-0-v1-0-20230919-pdf>.
- [60] Deepika Pandey. 2012. Comparative Analysis of Different Topologies Based On Network-on-Chip Architectures. *International Journal of Engineering Research and Development* 1, 11 (2012), 71–76.
- [61] Pitch Patarasuk and Xin Yuan. 2009. Bandwidth optimal all-reduce algorithms for clusters of workstations. *J. Parallel and Distrib. Comput.* 69, 2 (2009), 117–124.
- [62] Yanghua Peng, Yibo Zhu, Yangrui Chen, Yixin Bao, Bairen Yi, Chang Lan, Chuan Wu, and Chuanxiong Guo. 2019. A Generic Communication Scheduler for Distributed DNN Training Acceleration. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles (SOSP ’19)*. Association for Computing Machinery, New York, NY, USA, 16–29. <https://doi.org/10.1145/3341301.3359642>
- [63] NVIDIA DGX Platform. 2024. (2024). <https://www.nvidia.com/en-us/data-center/dgx-platform/>.
- [64] George Porter, Richard Strong, Nathan Farrington, Alex Forencich, Pang Chen-Sun, Tajana Rosing, Yeshaihu Fainman, George Papen, and Amin Vahdat. 2013. Integrating microsecond circuit switching into the data center. *ACM SIGCOMM Computer Communication Review* 43, 4 (2013), 447–458.
- [65] Kun Qian, Yongqing Xi, Jiamin Cao, Jiaqi Gao, Yichi Xu, Yu Guan, Binzhang Fu, Xuemei Shi, Fangbo Zhu, Rui Miao, Chao Wang, Peng Wang, Pengcheng Zhang, Xianlong Zeng, Eddie Ruan, Zhiping Yao, Ennan Zhai, and Dennis Cai. 2024. Alibaba HPN: A Data Center Network for Large Language Model Training. In *Proceedings of the ACM SIGCOMM 2024 Conference (ACM SIGCOMM ’24)*. 691–706.
- [66] Aurick Qiao, Sang Keun Choe, Suhas Jayaram Subramanya, Willie Neiswanger, Qirong Ho, Hao Zhang, Gregory R. Ganger, and Eric P. Xing. 2021. Pollux: Co-adaptive Cluster Scheduling for Goodput-Optimized Deep Learning. In *15th USENIX Symposium on Operating Systems Design and Implementation (OSDI 21)*. USENIX Association, 1–18. <https://www.usenix.org/conference/osdi21/presentation/qiao>
- [67] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. ZeRO: Memory Optimizations Toward Training Trillion Parameter Models. (2020). arXiv:cs.LG/1910.02054 <https://arxiv.org/abs/1910.02054>
- [68] SemiAnalysis. 2023. Google OCS Apollo: The >\$3 Billion Game-Changer in Datacenter Networking. (2023). <https://semianalysis.com/2023/03/17/google-apollo-the-3-billion-game/>.
- [69] SemiAnalysis. 2024. GB200 Hardware Architecture – Component Supply Chain & BOM. (2024). <https://semianalysis.com/2024/07/17/gb200-hardware-architecture-and-component/>.
- [70] SemiAnalysis. 2024. NVIDIA’s Blackwell Reworked – Shipment Delays & GB200A Reworked Platforms. (2024). <https://semianalysis.com/2024/08/04/nvidias-blackwell-reworked-shipment/>.
- [71] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2020. Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism. (2020). arXiv:cs.CL/1909.08053 <https://arxiv.org/abs/1909.08053>
- [72] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, Elton Zhang, Rewon Child, Reza Yazdani Aminabadi, Julie Bernauer, Xia Song, Mohammad Shoeybi, Yuxiong He, Michael Houston, Saurabh Tiwary, and Bryan Catanzaro. 2022. Using DeepSpeed and Megatron to Train Megatron-Turing NLG 530B, A Large-Scale Generative Language Model. (2022). arXiv:cs.CL/2201.11990 <https://arxiv.org/abs/2201.11990>
- [73] Xingwu Sun, Yanfeng Chen, Yiqing Huang, Ruobing Xie, Jiaqi Zhu, Kai Zhang, Shuaipeng Li, Zhen Yang, Jonny Han, Xiaobo Shu, Jiahao Bu, Zhongzhi Chen, Xueming Huang, Fengzong Lian, Saiyong Yang, Jianfeng Yan, Yuyuan Zeng, Xiaoqin Ren, Chao Yu, Lulu Wu, Yue Mao, Tao Yang, Suncong Zheng, Kan Wu, Dian Jiao, Jinbao Xue, Xipeng Zhang, Decheng Wu, Kai Liu, Dengpeng Wu, Guanghui Xu, Shaohua Chen, Shuang Chen, Xiao Feng, Yigeng Hong, Junqiang Zheng, Chengcheng Xu, Zongwei Li, Xiong Kuang, Jianglu Hu, Yiqi Chen, Yuchi Deng, Guiyang Li, Ao Liu, Chenchen Zhang, Shihui Hu, Zilong Zhao, Zifan Wu, Yao Ding, Weichao Wang, Han Liu, Roberts Wang, Hao Fei, Peijie She, Ze Zhao, Xun Cao, Hai Wang, Fusheng Xiang, Mengyuan Huang, Zhiyuan Xiong, Bin Hu, Xuebin Hou, Lei Jiang, Jiajia Wu, Yaping Deng, Yi Shen, Qian Wang, Weijie Liu, Jie Liu, Meng Chen, Liang Dong, Weiwen Jia, Hu Chen, Feifei Liu, Rui Yuan, Huilin Xu, Zhenxiang Yan, Tengfei Cao, Zhichao Hu, Xinhua Feng, Dong Du, Tinghao She, Yangyu Tao, Feng Zhang, Jianchen Zhu, Chengzhong Xu, Xirui Li, Chong Zha, Wen Ouyang, Yinben Xia, Xiang Li, Zekun He, Rongpeng Chen, Jiawei Song, Ruibin Chen, Fan Jiang, Chongqing Zhao, Bo Wang, Hao Gong, Rong Gan, Winston Hu, Zhanhui Kang, Yong Yang, Yuhong Liu, Di Wang, and Jie Jiang. 2024. Hunyuan-Large: An Open-Source MoE Model with 52 Billion Activated Parameters by Tencent. (2024). arXiv:cs.CL/2411.02265 <https://arxiv.org/abs/2411.02265>
- [74] Tesla. 2024. Tesla Dojo. (2024). https://en.wikipedia.org/wiki/Tesla_Dojo.
- [75] AWS Trainium. 2024. (2024). <https://aws.amazon.com/ai/machine-learning/trainium/>.
- [76] Ryohei Urata, Hong Liu, Kevin Yasumura, Erji Mao, Jill Berger, Xiang Zhou, Cedric Lam, Roy Bannon, Darren Hutchinson, Daniel Nelson, Leon Poutievski, Arjun Singh, Joon Ong, and Amin Vahdat. 2022. Mission Apollo: Landing Optical Circuit Switching at Datacenter Scale. (2022). arXiv:cs.NI/2208.10041
- [77] Ryohei Urata, Hong Liu, Kevin Yasumura, Erji Mao, Jill Berger, Xiang Zhou, Cedric Lam, Roy Bannon, Darren Hutchinson, Daniel Nelson, Leon Poutievski, Arjun Singh, Joon Ong, and Amin Vahdat. 2022. Mission Apollo: Landing Optical Circuit Switching at Datacenter Scale. (2022). arXiv:cs.NI/2208.10041 <https://arxiv.org/abs/2208.10041>
- [78] Guohui Wang, David G Andersen, Michael Kaminsky, Konstantina Papagiannaki, TS Eugene Ng, Michael Kozuch, and Michael Ryan. 2010. c-Through: Part-time optics in data centers. In *Proceedings of the ACM SIGCOMM 2010 Conference*. 327–338.

- [79] Weiyang Wang, Manya Ghobadi, Kayvon Shakeri, Ying Zhang, and Naader Hasani. 2024. Rail-only: A Low-Cost High-Performance Network for Training LLMs with Trillion Parameters. (2024). arXiv:cs.NI/2307.12169 <https://arxiv.org/abs/2307.12169>
- [80] Weiyang Wang, Moein Khazraee, Zhizhen Zhong, Manya Ghobadi, Zhihao Jia, Dheevatsa Mudigere, Ying Zhang, and Anthony Kewitsch. 2023. TopoOpt: Co-optimizing Network Topology and Parallelization Strategy for Distributed Training Jobs. In *20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 23)*. USENIX Association, Boston, MA, 739–767.
- [81] Qizhen Weng, Wencong Xiao, Yinghao Yu, Wei Wang, Cheng Wang, Jian He, Yong Li, Liping Zhang, Wei Lin, and Yu Ding. 2022. MLaaS in the Wild: Workload Analysis and Scheduling in Large-Scale Heterogeneous GPU Clusters. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. USENIX Association, Renton, WA, 945–960. <https://www.usenix.org/conference/nsdi22/presentation/weng>
- [82] Wikipedia. 2024. Mach–Zehnder interferometer. (2024). https://en.wikipedia.org/wiki/Mach-Zehnder_interferometer.
- [83] Yiting Xia, Mike Schlansker, TS Eugene Ng, and Jean Tourrilhes. 2015. Enabling Topological Flexibility for Data Centers Using {OmniSwitch}. In *7th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 15)*.
- [84] Wencong Xiao, Romil Bhardwaj, Ramachandran Ramjee, Muthian Sivathanu, Nipun Kwatra, Zhenhua Han, Pratyush Patel, Xuan Peng, Hanyu Zhao, Quanlu Zhang, Fan Yang, and Lidong Zhou. 2018. Gandiva: Introspective Cluster Scheduling for Deep Learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. USENIX Association, Carlsbad, CA, 595–610. <https://www.usenix.org/conference/osdi18/presentation/xiao>
- [85] Tze-Wei Yeow, K.L.E. Law, and A. Goldenberg. 2001. MEMS optical switches. *IEEE Communications Magazine* 39, 11 (2001), 158–163.
- [86] Zili Zhang, Yinmin Zhong, Ranchen Ming, Hanpeng Hu, Jianjian Sun, Zheng Ge, Yibo Zhu, and Xin Jin. 2024. DistTrain: Addressing model and data heterogeneity with disaggregated training for multimodal large language models. *arXiv preprint arXiv:2408.04275* (2024).
- [87] Yinmin Zhong, Zili Zhang, Bingyang Wu, Shengyu Liu, Yukun Chen, Changyi Wan, Hanpeng Hu, Lei Xia, Ranchen Ming, Yibo Zhu, et al. 2024. RLHFuse: Efficient rlhf training for large language models with inter-and intra-stage fusion. *arXiv preprint arXiv:2409.13221* (2024).

A PRODUCTION FAULT TRACE

The production fault trace was collected from an 8-GPU node pretrain cluster with 2880 GPUs over a period of 160 days. The trace includes details such as fault start time, fault end time, and the ID of the faulty node. Figure 14a and Figure 14b provide a macro-level overview of the production fault trace. On average, the ratio of faulty 8-GPU nodes at any given time is 3.83%, with a p99 value of 7.22%.



(a) Fault Node Ratio Trace. (b) Cumulative Distribution.

Figure 14: Fault node trace in the production AI DC.

Since most of failure events are GPU faults, we normalized the trace of 8-GPU nodes to generate 4-GPU nodes trace. Assuming that the fault rates of GPUs are i.i.d. with a fault probability of p for each GPU, and considering that a node is deemed faulty if any GPU within it fails, the fault rate of an 8-GPU node is calculated as:

$$P_{fault}(8\text{-GPU}) = 1 - (1 - p)^8 = 3.83\%.$$

From this, we derive $p = 0.49\%$. The fault rate for a 4-GPU node is then:

$$P_{fault}(4\text{-GPU}) = 1 - (1 - p)^4 = 1.93\%.$$

The fault event of 4-GPU node is generate with Bayesian Equation, as:

$$\begin{aligned} P_{fault}(4\text{-GPU} | 8\text{-GPU}) &= \frac{P_{fault}(8\text{-GPU} | 4\text{-GPU})P_{fault}(4\text{-GPU})}{P_{fault}(8\text{-GPU})} \\ &= \frac{1 \times 1.93\%}{3.83\%} = 50.39\% \end{aligned}$$

Thus, whenever a fault occurs in an 8-GPU node in the original trace, each of the two corresponding 4-GPU nodes at the same location has a 50.39% probability of fault. This method is used to convert the traces.

As node faults are i.i.d., the simulator linearly maps the fault trace to different network architectures.

B GPT-MOE ARCHITECTURE

This model is a mixture-of-experts (MoE) model with the following configuration:

Model Configuration:

- **Number of Layers:** 192
- **Inner Layer Dimension:** 49152
- **Embedding Dimension:** 12288
- **Hidden Dimension:** 12288
- **Vocabulary Size:** 64000
- **Number of Attention Heads:** 128
- **Maximum Sequence Length:** 2048
- **Number of Experts:** 8
- **MoE Layer Ratio:** 0.5
- **Top-K Experts:** 2

Runtime Configuration:

- **Virtual Pipeline Parallelism:** 3
- **Micro Batch Size:** 1
- **Global Batch Size:** 1536
- **Max Sequence Length:** 2048

C THEORETICAL ANALYSIS OF WASTED GPU RATIO FOR *INFINITEPOD*

The count of backup lines as $2K-2$ will significantly influence the fault tolerance of *InfinitePOD*. We use the expectation of waste ratio caused by GPU failure and fragmentation problem to evaluate this design, the result is shown in Table 7.

For one single working server in the middle of line, the count of breakpoints B on its two sides has the expectation as:

$$E_B(\eta = 1, middle) = 2(P_s^K + P_s^{2K})$$

Where P_s is the fail probability of GPU server, and η is count of servers. The expectation of breakpoints count is:

Once the distance between one server and the tail of line is $\alpha < K$, it will connect to all servers between itself and the last one, so there will be no breakpoints on this side, and the expectation of breakpoints count is less than servers in the middle of line. Then, for any server in the line topology:

$$E_B(\eta = 1) \leq E_B(\eta = 1, middle)$$

When the distance between two servers is $\beta \geq K$, the breakpoints among them can be calculated as independent. Once the distance $\beta < K$, as all servers in this range are connected to these two servers, there will be no breakpoints between them. So, the expectation is less than two independent servers. Then,

$$E_B(\eta = 2) < E_B(\eta = 2, \beta \geq K) = 2E(\eta = 1)$$

$$E_B(\eta = N_s) \leq N_s E_B(\eta = 1)$$

For a LLM job which require a ring communication size (TP .etc) as N_t , *InfinitePOD* will cut the whole line topology

into several sub lines with the length of N_t/R . Once *InfinitePOD* is cutting a new sub line from the remaining servers in the line, all N_t GPU will be wasted when one break point exist in the middle of this sub line required, shown in Figure 15. Then the expectation for waste GPU caused by one single break point is:

$$E_W(B = 1) = N_t R \cdot (1 - (N_t/R)^{-1}) = R(N_t - R)$$

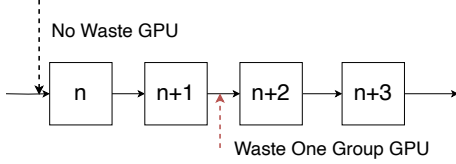


Figure 15: Break point can cause server waste compare to ideal situation.

As the influence between two break points only reduce the expectation of wasted GPUs, we can have this for X break points:

$$E_W(B = X) \leq X E_W(B = 1) = X R(N_t - R)$$

So the expectation of wasted GPU for a servers cluster with N_s GPU servers is:

$$\begin{aligned} E_W(\eta = N_s) &\leq \sum P(B = X, \eta = N_s) \cdot X \cdot E_W(B = 1) \\ &= E_B(\eta = N_s) \cdot E_W(B = 1) \\ &\leq \lim_{P_s \rightarrow 0} 2N_s \cdot R \cdot (N_t - R) P_s^K \end{aligned}$$

The final expectation of GPUs waste ratio is (1):

$$E_{WR}(\eta = N_s) = \frac{E_W(\eta = N_s)}{N_g} \leq 2(N_t - R)(P_s)^K \quad (1)$$

In our trace for a 160 days long pre-train job on 10K-GPU, the p99 failure rate for 8-card machines is 7%. If a TP32 jobs is running on *InfinitePOD*, we can get the upper bond for waste ratio expectation for various configuration in Table 7.

	$K = 2$	$K = 3$	$K = 4$
R=4	7.35%	0.26%	9.00×10^{-4}
R=8	27.4%	1.92%	0.13%

Table 7: Upper bond for waste ratio expectation of GPU, where GPU failure rate is 0.875% and X is 32

As shown in the table, for 4 GPU server ($R = 4$) 3 bundles ($K = 3$) design, the additional waste of GPU is less than 0.26%, while the waste ratio for $R = 8, K = 4$ is less than 0.13%. This is sufficient for production clusters.

D ORCHESTRATION FOR FAT-TREE

In this section, we introduce the orchestration algorithm under Fat-Tree DCN in detail.

Notations To ensure rigorous mathematical reasoning, we introduce the following notations:

- n : number of nodes in the data-center.
- K : OCSTrx bundle (see §4.2).
- S_{all} : ordered set, represents all nodes numbered from 1 according to their physical connection order in DCN fabric. $|S_{all}| = n$.
- S : ordered subset, represents nodes, $\forall u \in S, u \in S_{all}$. Adjacent elements in S are also adjacent from the perspective of the *InfinitePOD* topology.
- E : The set of edges across S , should be equal to $\{(S_i, S_j) \mid 1 \leq i < j \leq n, j - i \leq K\}$, representing the connections between nodes, including both primary and backup links, and $O(|E|) = O(K|S|)$.
- $InfHBD = \langle S, E \rangle$: the topology of *InfinitePOD* as an undirected graph.
- F : faulty nodes.
- $HealthyHBD = \langle H, HE \rangle$: healthy node subgraph where the set of healthy nodes $H = S - F$ and the edge set $HE = \{(u, v) \mid u \in H \text{ and } v \in H \text{ and } (u, v) \in E\}$.
- t : TP size, number of GPUs in one TP Group.
- r : GPU ranks per node.
- $m = t/r$: number of nodes in a TP group.
- s : job scale, number of GPUs required for the job.
- d : Aggregation-Switches Domain size. Number of nodes under coverage of one group of Aggregation-Switches.
- $n_{constrains}$: number of applied constraints in binary-search-based orchestration algorithm.
- p : number of nodes under each ToR.
- l : shortest sub-line length under fat-tree orchestration.
- $n_{maxsubline} = \lfloor \frac{nd}{p} \rfloor$: max number of sub-lines.
- $G_{deploy} = \langle S_{deploy}, E_{deploy} \rangle$: deployed topology. After applying the deployment strategy, the topology from the perspective of *InfinitePOD* is described as follows: S_{deploy} is an ordered set where adjacent elements correspond to adjacent nodes in *InfinitePOD*, and E_{deploy} represents the connections between nodes.

The orchestration algorithm (Algorithm 2) without considering DCN has the overall time complexity $3 \cdot O(|H| + |HE|) = O(|S| + |E|) = O((K + 1)|S|) = O(|S|)$.

Fat-Tree topology is another common data center topology. A typical training strategy for this topology aims to maximize the bandwidth utilization under ToR (Top of Rack) Switches. Using Meta’s two-stage clos topology[24] as a reference, it can be observed that there is an attempt to run CP under ToR.

Deployment Strategy: Assuming there are p nodes under each ToR, nodes with the same index under each ToR are

Algorithm 2: Orchestration-DCN-Free

Input: InfHBD = $\langle S, E \rangle, F, m$
Output: Placement scheme maximizing GPU utilization
Initialize $H = S - F$;
Initialize $HE = \{(u, v) \mid u \in H \text{ and } v \in H \text{ and } (u, v) \in E\}$;
Create subgraph $HealthyHBD = \langle H, HE \rangle$;
Initialize $component_list = []$;
Initialize $visited = \{\}$;
Initialize $placement_scheme = \{\}$;
for each node s in H **do**
 if s not in $visited$ **then**
 $component = Connected - Component - DFS(s, HealthyHBD, visited)$;
 Add $component.sortedinHBD()$ to $component_list$;
for each component in $component_list$ **do**
 while $component.size() \geq m$ **do**
 Add $component.pop(m)$ to $placement_scheme$;
return $placement_scheme$

Algorithm 3: Deployment-Strategy

Input: Node ordered set S , OCSTrx direction K , parallel factor p
Output: Deployment topology
 $G_{deploy} = \langle S_{deploy}, E_{deploy} \rangle$
Initialize ordered set $S_{deploy} = []$;
Initialize $l = \lfloor \frac{|S|}{p} \rfloor$;
for i in $0 \dots p - 1$ **do**
 for j in $0 \dots l - 1$ **do**
 Add $i + j \cdot p$ to S_{deploy} ;
Create $E_{deploy} = \{(S_{deploy}^i, S_{deploy}^j) \mid 1 \leq i \leq j \leq |S_{deploy}|, j - i \leq K\}$;
return $G_{deploy} = \langle S_{deploy}, E_{deploy} \rangle$

deployed along the same parallel sub-line, and the p sub-lines are connected end-to-end, as shown in Figure 7. The training strategy involves running CP p across the sub-lines and running TP within them.

Orchestration Constraints. To maximize the utilization of ToR bandwidth and minimize cross-ToR traffic, the fat-tree topology introduces two constraints:

- **Aggregation-Switches Domain Constraint:** The coverage domain of a group of Aggregation Switches is limited, meaning that TP groups spanning across Aggregation Switches domains would result in cross-rail traffic, which should be avoided as much as possible.
- **TP Group Alignment Constraint:** A CP Group consists of TP Groups across parallel sub-lines. To keep CP traffic within the ToR, the TP Groups must be aligned. If

a node fails under one ToR, all nodes under that ToR are considered failed, expanding the failure radius by a factor of p .

Binary-Search-Based Orchestration Algorithm. Based on the constraints and deployment strategy, we develop a binary search orchestration algorithm (see Algorithm 5) that adjusts the number of satisfied constraints. The binary search first relaxes the TP Group alignment constraints within the Aggregation-Switches Domain and then relaxes the TP Group crossing constraints between Aggregation-Switch domains (see Algorithm 4). This process is monotonic.

The time complexity of Algorithm 2 is $O(|S|)$, and the complexity of Algorithm 4 is

$$\sum_{i=1}^{n_{subline}} O(|S_{subline}|) = O\left(\sum_{i=1}^{n_{subline}} |S_{subline}| \right) = O(|S_{all}|) = O(n)$$

Thus, the overall time complexity of Algorithm 5 is $O(n \log n)$.

Algorithm 4: Placement-Fat-Tree

Input: $G_{deploy} = \langle S_{deploy}, E_{deploy} \rangle, n_{constraints}, F, l, m, n_{maxsubline}, d, p$
Output: Placement scheme
Initialize $placement_scheme = \{\}$;
Initialize $n_{align} = \max(0, n_{constraints} - n_{maxsubline})$,
 $n_{subline} = \min(n_{maxsubline}, n_{constraints})$;
for i in $0 \dots n_{align} - 1$ **do**
 for j in $1 \dots d$ **do**
 $sid = i * d + j$;
 if $sid \in F$ **then**
 $F \cup \{\lfloor \frac{sid-1}{p} \rfloor \cdot p + 1 \dots (\lfloor \frac{sid-1}{p} \rfloor + 1) \cdot p\}$;
for i in $1 \dots n_{subline}$ **do**
 $S_{subline} = S_{deploy}.pop(l)$;
 $E_{subline} = \{(u, v) \mid u \in S_{subline} \text{ and } v \in S_{subline} \text{ and } (u, v) \in E_{subline}\}$;
 $F_{subline} = F \cap S_{subline}$;
 $placement_scheme = placement_scheme \cup Orchestration-Ideal(\langle S_{subline}, E_{subline} \rangle, F_{subline}, m)$;
 $E_{res} = \{(u, v) \mid u \in S_{deploy} \text{ and } v \in S_{deploy} \text{ and } (u, v) \in E_{deploy}\}$;
 $F_{res} = F \cap S_{deploy}$;
 $placement_scheme = placement_scheme \cup Orchestration-Ideal(\langle S_{deploy}, E_{res} \rangle, F_{res}, m)$;
return $placement_scheme$

E ADDITIONAL SIMULATION RESULTS FOR FAULT RESILIENCE

This section presents additional simulation results related to §6.2. Figure 16 shows the variation of the GPU waste ratio

Algorithm 5: Orchestration-Fat-Tree

Input: S, r, p, F, t, s, d, K .

Output: Placement scheme that satisfies job scale and minimizes cross-rail traffic.

Initialize $m = t/r, n = |S|, l = \lfloor \frac{d}{p} \rfloor, n_{domain} = \lfloor \frac{n}{d} \rfloor$,

$n_{maxsubline} = \lfloor \frac{nd}{p} \rfloor$;

Create graph $G_{deploy} = \langle S_{deploy}, E_{deploy} \rangle =$
Deployment-Strategy(S, K, p);

Initialize $high = n_{domain} + n_{maxsubline}$;

Initialize $low = 0$;

Initialize $placement_scheme = \{\}$;

while $low \leq high$ **do**

$mid = \lfloor \frac{low+high}{2} \rfloor$;

$placement_scheme =$

 Placement-Fat-Tree($G_{deploy}, mid, F, l, m, n_{maxsubline}, d, p$);

if $|placement_scheme| \cdot m \cdot r \geq s$ **then**

$low = mid + 1$;

else

$high = mid - 1$;

if $|placement_scheme| \cdot m \cdot r \geq s$ **then**

return $placement_scheme$

else

return *None*

over time under the production fault trace. Figure 17 presents the CDF data for the GPU waste ratio. Figure 18 illustrates the waste GPU ratio for different HBD architectures under various node failure rates, including the results for TP-8 to TP-64. Figure 19 shows the proportion of job-fault waiting time relative to total time for different job scales. All the aforementioned experiments include results for TP-8, TP-16, TP-32, and TP-64 configurations.

F DETAILED COST AND POWER CONSUMPTION ANALYSIS

In this section, Table 8 provides a detailed description of the quantity, cost, bandwidth, and power consumption of the interconnect components in various network architectures, including Google TPUv4 [33], NVIDIA GB200 NVL series [56], Alibaba HPN[65], and *InfinitePOD*.

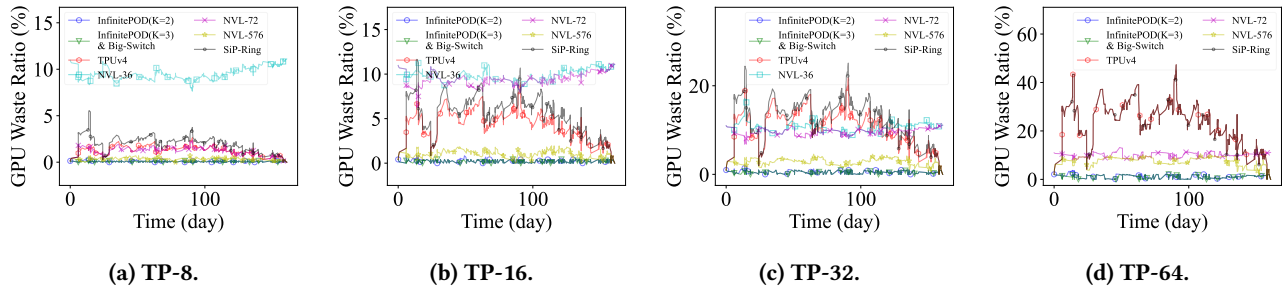


Figure 16: GPU waste ratio over production fault trace, 4 GPU node.

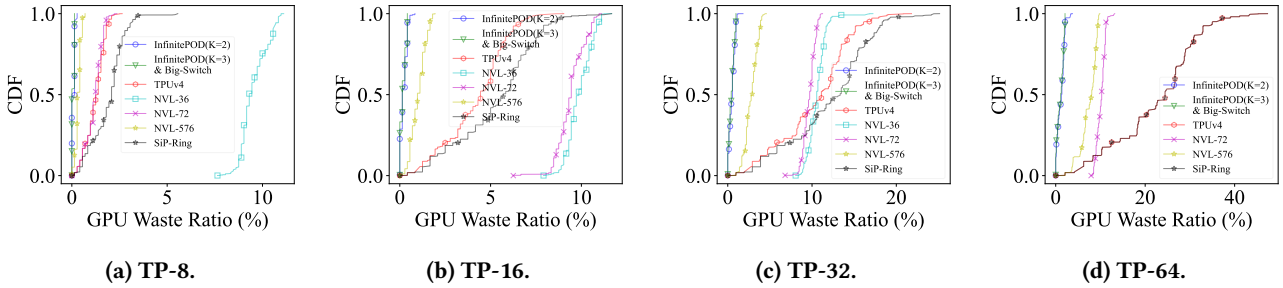


Figure 17: CDF of GPU waste ratio over production fault trace, 4 GPU node.

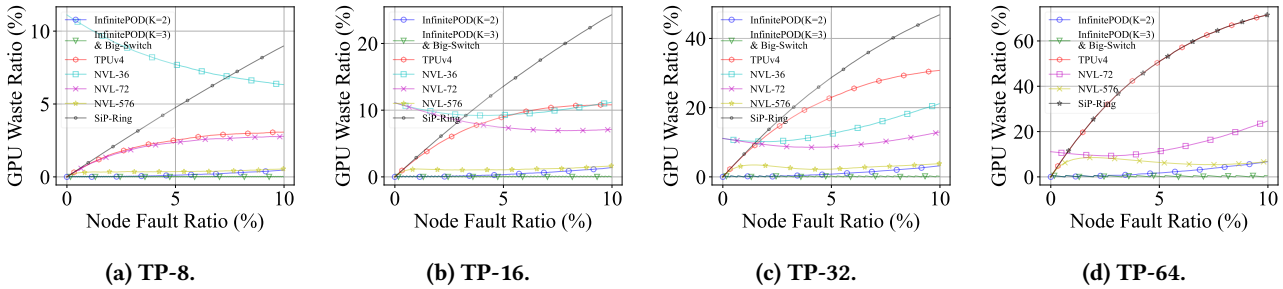


Figure 18: GPU wastes ratio with different GPU fault ratio, 4-GPU node.

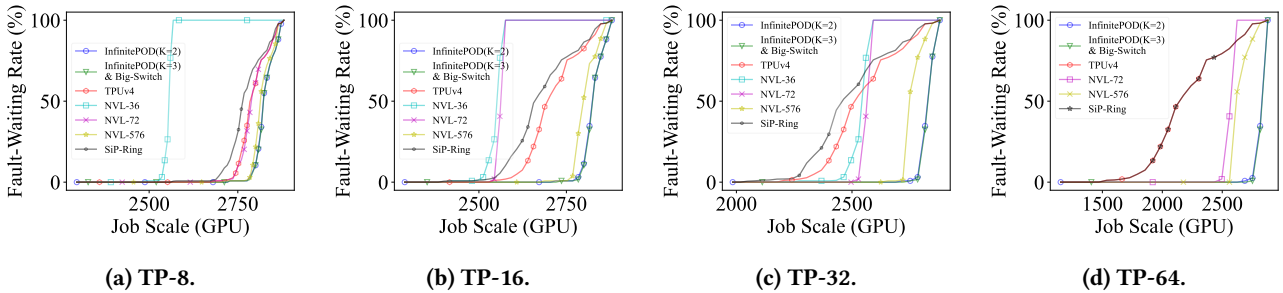


Figure 19: Job fault-waiting duration with different levels of job-scale, 4 GPU node

Component	Quantity	Unit Cost (\$)	Unit Bandwidth (GBps)	Unit Power (W)
Google TPUv4[33] with 4096 GPU, bandwidth 300GBps/GPU				
OCS[40]	48	80000	6400	108
DAC Cable[19]	5120	63.60	50	0.1
Optical Module[16]	6144	360	50	12
Fiber[22]	6144	6.80	50	0
NVIDIA GB200 NVL-36[69] with 36 GPU, bandwidth 900GBps/GPU				
NVLink Switch[70]	9	28000	3600	275
DAC Cable[20]	2592	35.60	25	0.1
NVIDIA GB200 NVL-72[56][69] with 72 GPU, bandwidth 900GBps/GPU				
NVLink Switch[70]	18	28000	3600	275
DAC Cable[20]	5184	35.60	25	0.1
NVIDIA GB200 NVL-36x2[69] with 72 GPU, bandwidth 900GBps/GPU				
NVLink Switch[70]	36	28000	3600	275
DAC Cable[20]	6480	35.60	25	0.1
ACC Cable[70]	162	320	200	2.5
NVIDIA GB200 NVL-576[69] with 576 GPU, bandwidth 900GBps/GPU				
NVLink Switch[70]	432	28000	3600	275
DAC Cable[20]	41472	35.60	25	0.1
Optical Module[47]	4608	850	200	25
Fiber[22]	4608	6.80	200	0
Alibaba HPN[65] with 16320 GPU, bandwidth 50GBps/GPU				
EPS[49]	360	14960	6400	3145
DAC Cable[20]	32640	35.60	25	0.1
Optical Module[16]	28800	360	50	12
Fiber[22]	14400	6.80	50	0
InfinitePOD($K = 2$) with 4 GPU, bandwidth 800GBps/GPU				
DAC Cable[21]	4	199.60	200	0.1
dOCS Module	16	600	100	12
Fiber[22]	16	6.80	100	0
InfinitePOD($K = 3$) with 4 GPU, bandwidth 800GBps/GPU				
DAC Cable[21]	2	199.60	200	0.1
dOCS Module	24	600	100	12
Fiber[22]	24	6.80	100	0

Table 8: Interconnect cost and power consumption of components used in different network architectures.