
LLM INFERENCE BEYOND A SINGLE NODE: FROM BOTTLENECKS TO MITIGATIONS WITH FAST ALL-REDUCE COMMUNICATION

Prajwal Singhania¹ Siddharth Singh¹ Lannie Dalton Hough¹ Akarsh Srivastava¹ Harshitha Menon²
Charles Fredrick Jekel² Abhinav Bhatele¹

ABSTRACT

As large language models (LLMs) continue to grow in size, distributed inference has become increasingly important. Model-parallel strategies must now efficiently scale not only across multiple GPUs but also across multiple nodes. In this work, we present a detailed performance study of multi-node distributed inference using LLMs on GPU-based supercomputers. We conduct experiments with several state-of-the-art inference engines alongside YALIS, a research-oriented prototype engine designed for controlled experimentation. We analyze the strong-scaling behavior of different model-parallel schemes and identify key bottlenecks. Since all-reduce operations are a common performance bottleneck, we develop NVRAR, a hierarchical all-reduce algorithm based on recursive doubling with NVSHMEM. NVRAR achieves up to 1.9x-3.6x lower latency than NCCL for message sizes between 128 KB and 2 MB on HPE Slingshot and InfiniBand interconnects. Integrated into YALIS, NVRAR achieves up to a 1.72x reduction in end-to-end batch latency for the Llama 3.1 405B model in multi-node decode-heavy workloads using tensor parallelism.

Code: <https://github.com/axonn-ai/yalis>, <https://github.com/hpcgroup/nvra>

1 INTRODUCTION

As large language models (LLMs) grow in size and with wider adoption, inference costs are rising rapidly (Maslej et al., 2025; Bick et al., 2024; International Energy Agency, 2025). As practitioners pursue higher-quality outputs through larger models, longer sequence lengths, compute-intensive reasoning architectures, and test-time scaling (Wei et al., 2022; Jaech et al., 2024; Snell et al., 2024), improving inference performance is becoming critical for reducing energy consumption and operational costs.

With increasing LLM sizes, memory footprints often exceed the capacity of a single GPU, requiring parallel execution across multiple devices. On most clusters, a single node, typically with 4 to 8 GPUs, is insufficient to host large models such as Llama 3.1 405B (Grattafiori et al., 2024). To enable inference using such models, inference engines use *model parallelism* (Shoeybi et al., 2020; Huang et al., 2019) strategies, which partition the model parameters across GPUs. While model parallelism for inference on a single GPU and within a single node has been extensively studied and optimized (Aminabadi et al., 2022; Li et al., 2024b), inference

in *multi-node* settings is comparatively underexplored. In this paper, we systematically study, comparatively evaluate and optimize multi-node inference workloads.

Multi-node inference introduces new challenges such as higher inter-node latencies compared to faster within-node NVLink connections. As a result, parallelization strategies that perform well within a node can experience substantial degradation in multi-node settings due to increased communication overheads. Moreover, the optimal choice of parallelism strategy often depends on the specific inference workload characteristics and the efficiency of underlying communication libraries. Consequently, it remains unclear which model parallel schemes are more suited for multi-node inference and how to optimize them further.

In this work, we address the following research questions:

- How do different model parallel schemes (tensor and hybrid parallelism) scale across multiple nodes in distributed environments for specific inference workloads?
- What performance bottlenecks arise in these model parallel schemes under the workloads studied?
- Can we optimize collective communication, which appears to be a common performance bottleneck in multi-node inference?

To investigate the questions above, we systematically study the performance of two popular model-parallel schemes:

¹Department of Computer Science, University of Maryland

²Lawrence Livermore National Laboratory. Correspondence to: Prajwal Singhania <prajwal@umd.edu>, Abhinav Bhatele <bhatele@cs.umd.edu>.

tensor parallelism (TP) and hybrid tensor-pipeline parallelism (HP), in multi-node distributed environments. We evaluate two state-of-the-art inference engines, vLLM (Kwon et al., 2023) and SGLang (Zheng et al., 2024), alongside YALIS, an inference engine developed as a research vehicle to facilitate easy instrumentation of multi-node experiments on Slurm-based environments. We evaluate the performance of these engines on a specific set of inference workloads, and study the scaling behavior of both model-parallel schemes. We identify bottlenecks via detailed performance breakdowns for each of the schemes.

Based on the results of our performance study, we identify that workloads that perform better with TP suffer from significant communication overheads arising from all-reduce operations. To address this, we propose *NVRAR*: a hierarchical recursive all-reduce implementation built using NVSHMEM (NVIDIA, 2020b), and optimized for message sizes occurring in inference workloads. We evaluate *NVRAR* against NCCL’s all-reduce (NVIDIA, 2020a) on multiple HPC interconnects, and observe up to $1.9\times$ better performance on HPE Slingshot-11 and $2.8\times$ on InfiniBand networks in the 256 KB to 2 MB message size range. Integrating *NVRAR* into YALIS yields up to a $1.9\times$ improvement in multi-node inference performance for the Llama 3.1 405B model in decode-heavy regimes.

The main contributions of this work are as follows:

- We systematically study the performance of model-parallel inference schemes for specific inference workloads in multi-node settings, producing detailed performance breakdowns. To facilitate this study, we develop YALIS, an inference engine designed for easier experimentation in multi-node HPC environments.
- Based on our performance analysis, we characterize how tensor parallelism compares to pipeline parallelism for different multi-node workloads and across inference phases. We identify bottlenecks in both parallelism schemes.
- To address a communication bottleneck in multi-node TP inference, we develop *NVRAR*, a custom all-reduce implementation optimized for the small-message regime characteristic of decode-heavy workloads. *NVRAR* delivers up to $1.92\times$ faster multi-node TP inference for the Llama 3.1 405B model.

2 BACKGROUND

This section provides an overview of LLM inference, model parallelization strategies used in inference, and a discussion of communication primitives used to in them.

LLM inference consists of two phases: *prefill* and *decode*. In prefill, the model processes all prompt tokens in parallel to generate the first output token and is typically compute-

bound due to large matrix multiplications. In decode, it generates subsequent tokens sequentially, becoming memory-bandwidth-bound because of smaller matrix multiplications and frequent parameter/KV-cache accesses.

2.1 Model Parallelism for Inference

LLMs that exceed the memory capacity of a single GPU require distributing model parameters and computations across multiple GPUs. This is broadly termed as *model parallelism*, which can be implemented in several ways. In *pipeline parallelism (PP)*, contiguous groups of layers are assigned to P processing units (pipeline stages), forming a sequential dependency chain with point-to-point communications. It achieves high utilization by splitting a batch of prompts into pipelined micro-batches. In *tensor parallelism (TP)*, the computation of each layer is partitioned across GPUs by splitting the underlying matrix multiplications. TP has no sequential dependency between GPUs, but aggregation of partial results incurs high communication overheads due to per-layer *all-reduce* operations.

2.2 Algorithms for All-reduce

NCCL (NVIDIA, 2020a) is the default communication library for AI workloads on NVIDIA GPUs and primarily implements two all-reduce algorithms: *Ring* and *Tree* (Hu et al., 2025). Other variants, such as *CollNet*, depend on specialized DGX hardware and are out of scope for this study. We model the performance of Ring and Tree all-reduce using the α - β communication model (Hockney, 1994). Consider a system with N nodes, each containing G GPUs. The inter-node network has latency α_{inter} and bandwidth β_{inter} , while the intra-node interconnect has latency α_{intra} and bandwidth β_{intra} , where $\alpha_{intra} < \alpha_{inter}$ and $\beta_{intra} > \beta_{inter}$. Let M denote the input message of size $|M|$ bytes.

Ring all-reduce: NCCL’s Ring all-reduce performs a reduce-scatter followed by an all-gather over a flat ring topology where all links are active each step. Inter-node links dominate the cost, and the communication time is modeled as:

$$T_{ring} = 2(NG - 1)\alpha_{inter} + 2\frac{NG - 1}{NG} \left(\frac{|M|}{\beta_{inter}} \right) \quad (1)$$

For small messages, bandwidth term can be neglected, giving

$$T_{ring} \approx 2(NG - 1)\alpha_{inter} \quad (2)$$

Tree all-reduce: The Tree all-reduce performs a reduction followed by a broadcast using a double binary tree topology (NVIDIA, 2019b) for inter-node communication and a simple intra-node chain. The communication time is mod-

eled as:

$$T_{\text{tree}} = 2(G - 1) \alpha_{\text{intra}} + 2 \log_2(N) \alpha_{\text{inter}} + 2 \frac{N - 1}{N} \left(\frac{|M|}{\beta_{\text{inter}}} \right) \quad (3)$$

For simplicity, only the inter-node bandwidth term is considered in the above expression ($\beta_{\text{intra}} \gg \beta_{\text{inter}}$). On further simplification for latency-bound small messages, we get:

$$T_{\text{tree}} \approx 2(G - 1) \alpha_{\text{intra}} + 2 \log_2(N) \alpha_{\text{inter}} \quad (4)$$

NVSHMEM: NVSHMEM is NVIDIA’s OpenSHMEM-based (Chapman et al., 2010) communication library providing host and device level APIs for one-sided put, get, and collective communication operations. It allows implementation of GPU-initiated communication kernels and supports multiple transports, including Slingshot and InfiniBand.

3 STUDYING PERFORMANCE OF MULTI-NODE INFERENCE

This section presents a performance study of multi-node LLM inference. Our objective is to evaluate different model-parallelism strategies, understand their scaling behavior and identify bottlenecks for specific workloads. We first introduce YALIS, a prototype inference engine developed as a research tool for controlled and instrumented performance studies. We then detail our benchmarking methodology and present our experimental results, discussing the performance of YALIS and existing state-of-the-art inference engines in multi-node settings.

3.1 YALIS: Yet Another LLM Inference System

YALIS is an open-source inference engine built as a research vehicle to study multi-node LLM inference. It is intended to be performant, easy to instrument, and more amenable to Slurm-based HPC environments. These properties allow for detailed analysis of multi-node LLM inference performance on HPC systems. Its design is centered around three key components: (1) a unified model definition layer, adapted from LitGPT (Lightning AI, 2023), providing compatibility with a wide range of model architectures; (2) an execution layer utilizing Torch Compile (Meta, 2023) for kernel fusion and optimization, and CUDA Graphs (NVIDIA, 2019a) for minimal kernel-launch overheads; and (3) tensor model parallelism implemented via AxoNN (Singh & Bhatele, 2022; Singh et al., 2024a;b), both within and across nodes. YALIS supports several modern inference optimizations including multiple attention backends, paged KV caching (Kwon et al., 2023), and speculative decoding (Leviathan et al., 2023).

3.2 Benchmarking Methodology

Hardware and Models: Our scaling experiments use the Perlmutter system (NERSC) (Table 1) (80 GB nodes unless otherwise specified). We evaluate on two dense LLMs - Llama 3.1 70B (Instruct) 405B (Instruct) (Grattafiori et al., 2024), run in bf16 precision.

Table 1. Details of the HPC systems used in our experiments.

System	GPU	GPUs/Node	Interconnect
Perlmutter	A100 (40/80 GB)	4	Intra-Node: 3 rd gen NVLink Inter-Node: Slingshot-11
Vista	GH200 (96 GB)	1	Inter-Node: InfiniBand

Table 2. Details of workloads evaluated in our experiments.

Workload	Prompt Length	Decode Length	NumPrompts (#P)
Prefill-heavy	2363	128	8, 32
Decode-heavy	1426	3072	8, 32

Workloads and Metrics: Table 2 lists the workload configurations used in our experiments. We define *NumPrompts* (#P) as the number of prompts provided to the inference engine in a single user batch. For brevity, a subset of results are presented in the main text, with additional results in the appendix.

In our workload characterization, one batch of prompts runs through the inference engine to completion before the next batch is submitted. This mirrors real-world settings such as GRPO (Shao et al., 2024; Guo et al., 2025), where a single batch of prompts require completions for the next policy update. We report the total time to completion for a single batch of prompts in our experiments.

In our strong scaling experiments (fixed workload across GPUs), the 70B model is scaled from 4 GPUs (1 node) to 32 GPUs (8 nodes), and the 405B model is scaled from 16 GPUs (4 nodes) to 128 GPUs (32 nodes). Each run has two warm-up generations and up to three timed generations. Runs are repeated three times and the average performance is reported. For performance breakdowns, we run one trial with two warm-up generations and one profiled generation.

Software Stack and Parallelism Schemes:

Table 3 lists the inference engines we use in our experiments for different model-parallel schemes. All experiments use PyTorch 2.8 (Paszke et al., 2019) and CUDA 12.9. For vLLM, we use the V0 engine for HP, as the V1 engine exhibited persistent hangs on Slurm-based systems. For performance breakdowns, we use Nsight Systems (NVIDIA) to collect traces and Pipit (Bhatele et al., 2023) to analyze them.

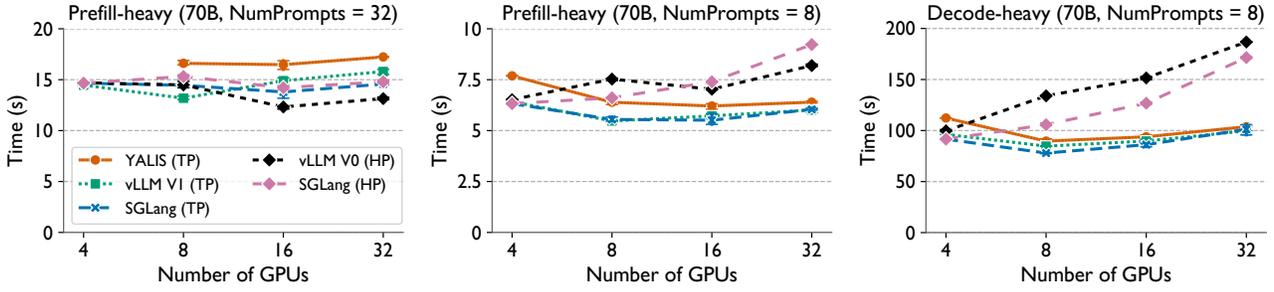


Figure 1. Strong scaling performance of different inference engines on Perlmutter for Llama 3.1 70B Instruct. The Y-axis shows the end-to-end latency per batch in seconds and the X-axis shows the number of GPUs.

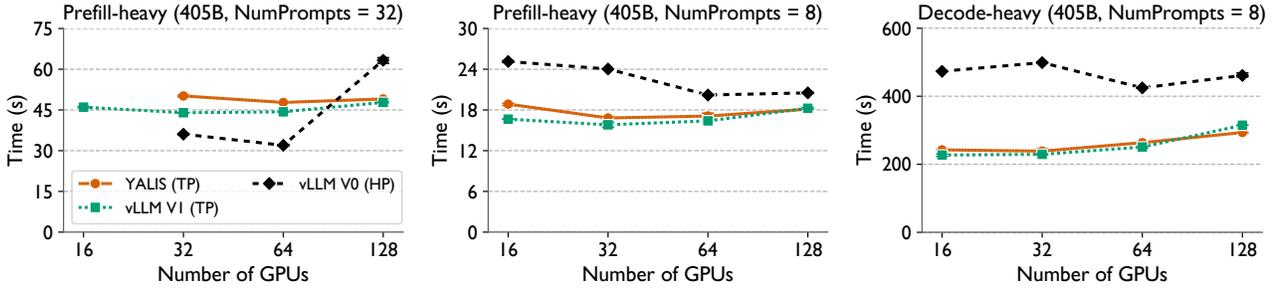


Figure 2. Strong scaling performance of different inference engines on Perlmutter for Llama 3.1 405B Instruct. The Y-axis shows the end-to-end latency per batch in seconds and the X-axis shows the number of GPUs.

Table 3. Evaluated parallelism schemes and inference frameworks.

Parallelism	Intra-Node	Inter-Node	Frameworks
Tensor Parallelism (TP)	TP	TP	YALIS
			vLLM V1 (v0.11.0)
			SGLang (v0.5.1)
Hybrid Parallelism (HP)	TP	PP	vLLM V0 (v0.10.0) SGLang (v0.5.1)

3.3 Scaling Multi-node LLM Inference

Figures 1 and 2 show the time-to-completion for a batch of prompts across all inference engines for Llama 3.1 70B and 405B (Instruct) on Perlmutter, respectively. From left to right, the workloads transition from compute-bound to increasingly memory-bound regimes.

We first observe that YALIS (orange line) achieves performance comparable to state-of-the-art engines, particularly in more memory-bound workloads. For the 70B model, YALIS is within 5-16% of vLLM V1 (TP) at 8 GPUs and beyond. For the 405B model, it remains within 8% across all GPU counts. The only noticeable deviation occurs for the 70B model’s prefill-heavy workload at 16 GPUs. Importantly, YALIS exhibits scaling trends consistent with other frameworks, validating its suitability as a research vehicle for studying multi-node LLM inference. The few missing data points correspond to OOM failures from `torch.compile`, which we plan to address in fu-

ture. Across all models and inference engines, both model-parallelism schemes exhibit poor strong scaling, where the time to solution does not scale inversely with GPU count.

Focusing on the 70B model (Figure 1), vLLM V1 (TP) (green line) latencies decrease from a 4 GPUs (single node) to 8 GPUs (two nodes), with more noticeable improvements for the decode-heavy workload (right-most plot). However, thereafter, the latency remains almost constant or increases with each doubling of the number of GPUs. This trend is consistent for TP, across all engines and models.

When using HP, a different trend is observed. In the prefill-heavy regime for the 70B model (Figure 1), vLLM V1 (HP) (black line) latencies remain nearly constant with a smaller number of prompts (middle plot), but decrease up to 16 GPUs with a larger number of prompts (left plot). For the 405B model, latencies decrease initially for both small and large number of prompts, before increasing or flattening out. SGLang (HP) (pink line) exhibits a similar trend. For decode-heavy workloads, however, HP latencies increase significantly with each doubling of node count for both vLLM and SGLang.

Comparing the two schemes, HP outperforms TP for the most compute-bound and prefill-heavy workload (Figure 1 left), but TP starts to outperform HP as workloads become more memory-bound and decode-heavy, across engines (Figure 1 middle, right). Similar trends are observed

for the 405B model.

Observation 1

For the workloads studied, TP and HP do not scale ideally. HP shows modest advantages in compute-bound regimes, whereas TP is better for memory-bound and decode-heavy cases.

3.4 Identifying Performance Bottlenecks

To better understand the scaling behaviors of TP and HP across prefill- and decode-heavy workloads, we look at performance breakdowns of YALIS (HP) and vLLM V0 (HP) on 8 and 16 GPUs for the 70B model (Figure 3). The total time is decomposed into four components: *Matmul* (time spent in matrix multiplications), *Other Comp.* (time spent in other computations), *Comm.* (time spent in communication), and *Idle* (GPU idle time).

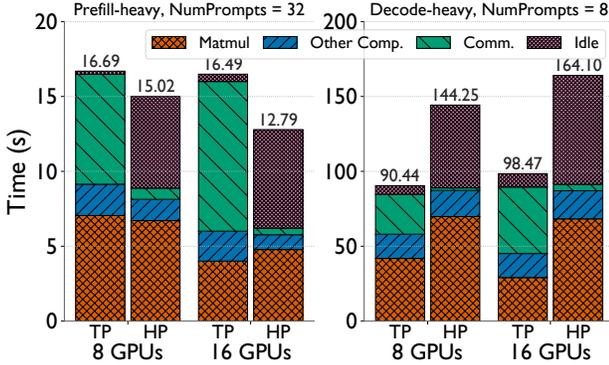


Figure 3. Performance breakdown of TP (using YALIS) and HP (using vLLM V0) for the prefill-heavy and decode-heavy workloads on Perlmutter for the 70B Llama model.

For the prefill-heavy workload (Figure 3, left), both YALIS (TP) and vLLM V0 (HP) reduce computation time going from 8 to 16 GPUs, with vLLM V0 (HP) achieving lower overall latency due to reduced communication overhead. However, vLLM V0 (HP) exhibits unexpectedly high GPU idle time. We hypothesize that this arises from repeated kernel launches on small micro-batches, but more analysis is needed to confirm the root cause.

For decode-heavy workloads, HP fails to reduce the time spent in matrix multiplications, unlike TP. This partially explains why, despite lower communication costs, HP does not scale as well for such workloads. To further isolate this behavior, we run a synthetic GEMM benchmark using two representative matrix sizes: Prefill-GEMM ($M=32768, N=8192, K=57344$) and Decode-GEMM ($M=32, N=8192, K=57344$). The former models large- M prefill matmuls

(batch size \times prompt length), while the latter models small- M decode matmuls (batch size \times 1). Figure 4 shows the performance of both when either M is halved (micro-batching in the PP phase of HP) or K is halved (TP). For Prefill-GEMM, time decreases in both cases, but for Decode-GEMM, time decreases only when K is halved. This behavior likely arises due to tiling in GEMM kernels, where decreasing M below the tile size yields no performance gain. While TP outperforms HP for these workloads, it still incurs significant communication overhead. Figure 3 (right) highlights that the communication time in YALIS (TP) increases by $\sim 1.6\times$, offsetting the gains from reduced computation time, when going from 8 to 16 GPUs.

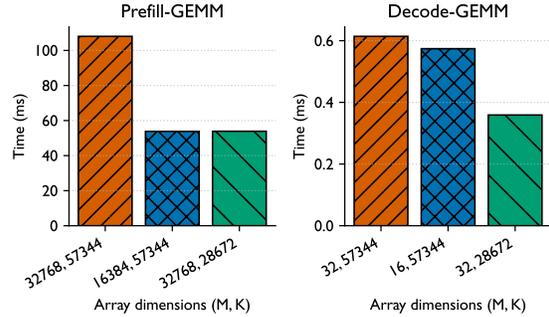


Figure 4. Synthetic GEMM benchmarks modeling Prefill (left) and Decode (right) matrix multiplications in the MLP layer of the 70B Llama model.

Observation 2

For prefill-heavy workloads, both TP and PP reduce computation time, with PP achieving lower overall latency due to its reduced communication overhead. For decode-heavy workloads, PP does not reduce matrix multiplication time, while TP suffers from significant communication overhead.

3.5 Communication Issues in Tensor Parallelism

The primary communication collective in TP is all-reduce, which, in the decode-heavy regime, is dominated by small messages of size $B \times H$, where B is the batch size and H the hidden dimension. For the 70B model with $B=8$ and $H=8192$, this message size is 128 KB. Across our workloads, message sizes range from 128 KB to 1 MB.

To further analyze communication bottlenecks, we benchmarked NCCL all-reduce against GPU-aware Cray-MPICH on Perlmutter (40 GB nodes), focusing on small messages. We used the `nccl-tests` (NVIDIA, 2017) and OSU benchmark (University) suites and report average latency over 10 trials (200 warm-up and 10,000 timed iterations). Figure 6 presents the results, showing that NCCL is much

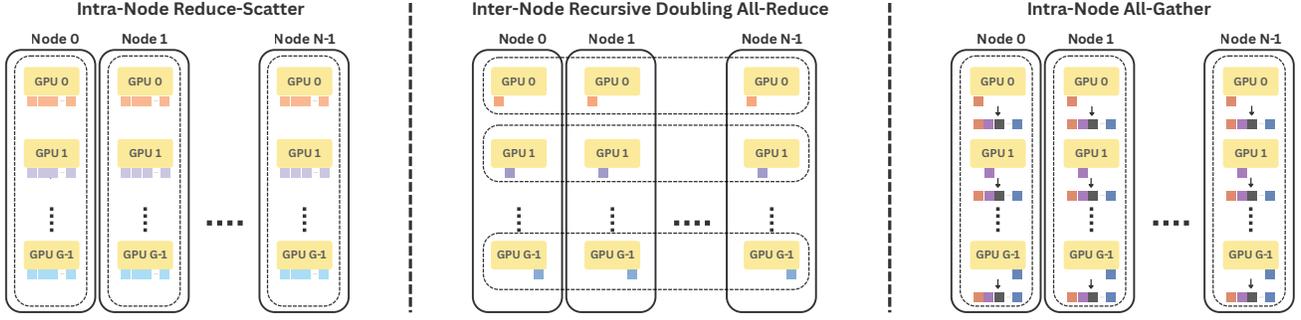


Figure 5. Three-phase NVRAR design: (1) intra-node reduce-scatter, (2) inter-node recursive-doubling all-reduce, (3) intra-node all-gather.

faster within a node, but that its latency increases sharply across nodes and scales poorly. For 512 KB-1 MB messages, NCCL is 1.5-2 \times slower than MPI, with latency growing faster with message size at any given scale. While Cray-MPICH’s implementation is proprietary, the open-source MPICH (Gropp et al., 2023) library typically employs a latency optimal recursive-doubling algorithm (Thakur & Gropp, 2003) for these regimes. We hypothesize that this can explain the observed performance gap.

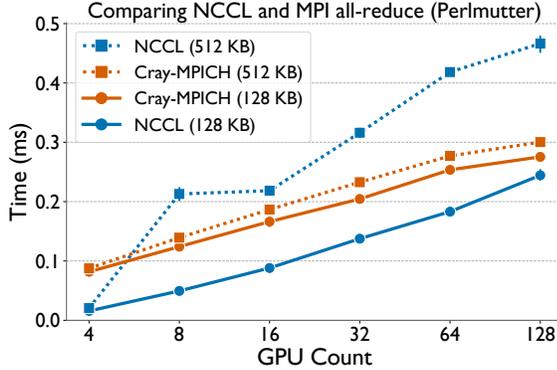


Figure 6. Scaling performance of NCCL and MPI all-reduce for a range of message sizes on Perlmutter.

Observation 3

For small message sizes, typical in the decode phase, NCCL all-reduce exhibits poor scaling across nodes and can at times be slower than MPI.

4 OPTIMIZED MULTI-NODE ALL-REDUCE

Having established the usefulness of TP for decode-heavy workloads, we now focus on optimizing its communication bottlenecks. One potential approach to address NCCL’s all-reduce performance issues is to use MPI when it’s faster than NCCL. However, standard MPI implementations are ill-suited for inference workloads, due to lack of CUDA Graph support and under-optimized NVLink communication paths.

To overcome these limitations, we propose NVRAR: an NVSHMEM-based hierarchical recursive all-reduce implementation, optimized for small-message inter-node communication. This section first describes the three-phase design of NVRAR, followed by optimizations for reducing synchronization overheads and data pipelining. Finally, we reason about the theoretical performance of the algorithm.

4.1 Three-Phase Hierarchical All-reduce Design

Algorithm 1 NVRAR

Require: Message M ; GPUs/node G ; number of nodes N ; chunk size C ; sequence number seq ; GPU rank (within node) r_g ; node rank r_n , pre-allocated send/receive buffers B_{send}, B_{recv}
Ensure: M reduced in-place

```

1: function NVRAR( $M, G, N, C, seq, r_g, r_n$ )
2:    $M' \leftarrow \text{REDUCE-SCATTER}_{intra}(M, G)$ 
3:    $seq \leftarrow seq + 1$ 
4:   for  $i = 0$  to  $\log_2 N - 1$  do
5:      $peer_i \leftarrow (r_n \oplus 2^i, r_g)$ 
6:     WAIT( $peer_i, seq$ ) {Synchronize Sequence Number}
7:   end for
8:    $B_{send}[0] \leftarrow \text{PACKDATAANDSEQNUM}(M', seq)$ 
9:    $B_{out} \leftarrow \text{RD}_{inter}(B_{send}, B_{recv}, N, C, seq, r_g, r_n)$ 
10:   $M' \leftarrow \text{UNPACKDATAANDSEQNUM}(B_{out}, seq)$ 
11:   $M \leftarrow \text{ALL-GATHER}_{intra}(M', G)$ 
12: end function
13:
14: function  $\text{RD}_{inter}(B_{send}, B_{recv}, N, C, seq, r_g, r_n)$ 
15:   $Q \leftarrow \lceil |B_{send}[0]|/C \rceil$  {Number of chunks}
16:  for  $\ell = 0$  to  $\log_2 N - 1$  do
17:     $peer_\ell \leftarrow (r_n \oplus 2^\ell, r_g)$ 
18:    for  $q = 0$  to  $Q - 1$  do
19:       $src_q \leftarrow B_{send}[\ell][q]$ ;  $dst_q \leftarrow B_{recv}[\ell][q]$ 
20:      NON-BLOCKING-PUT :  $src_q \rightarrow peer_\ell$ 's  $dst_q$ 
21:      Wait until  $flag(dst_q) == seq$  {Wait for data}
22:       $B_{send}[\ell + 1][q] \leftarrow dst_q + src_q$ 
23:    end for
24:  end for
25:  return  $B_{send}[\ell]$ 
26: end function
    
```

NVRAR consists of three phases: (1) an intra-node reduce-scatter, (2) an inter-node recursive-doubling all-reduce, and (3) an intra-node all-gather. Figure 5 illustrates this design

for an N -node system with G GPUs per node. Algorithm 4.1 describes the NVRAR algorithm in detail.

Reduce-scatter Phase: In the first phase (Line 2 of Algorithm 4.1), GPUs within a node perform a local reduce-scatter operation. For input message M of size $|M|$ bytes, each GPU holds $\frac{|M|}{G}$ bytes of locally reduced data after this phase. We implement this using the `nvshmemx_TYPENAME_sum_reducescatter` API, which internally calls NCCL reduce-scatter.

Inter-Node Recursive-Doubling Phase: In the second phase (Line 9 of Algorithm 4.1), corresponding GPUs across nodes perform a recursive-doubling all-reduce. Each node is identified by its rank $r_n \in [0, N - 1]$, and each GPU within a node by its local rank $r_g \in [0, G - 1]$, so each GPU is uniquely identified by the pair (r_n, r_g) . This phase completes in $\log_2 N$ steps. At each step $0 \leq i < \log_2 N$, GPU (r_n, r_g) exchanges data with its 2^i -th logical peer, $(r_n \oplus 2^i, r_g)$, where \oplus denotes bitwise XOR. Thus, GPUs with the same local rank communicate across nodes. In systems with multiple NICs per GPU, these exchanges can occur in parallel across local ranks. Upon receiving data, each GPU performs a local reduction with the received buffer before proceeding to the next step. After all $\log_2 N$ steps, each GPU holds $\frac{|M|}{G}$ bytes of the globally reduced data. This phase is implemented as a custom NVSHMEM kernel using non-blocking `put_nbi`-based RMA primitives.

All-gather Phase: In the third and final phase (Line 11 of Algorithm 4.1), the GPUs within a node perform a local all-gather operation to combine their $\frac{|M|}{G}$ fraction of the globally reduced data into a single tensor. Similar to the reduce-scatter phase, this is implemented using NVSHMEM’s host API. After completion, every GPU holds the full globally reduced tensor, completing the all-reduce.

4.2 Performance Optimizations

The inter-node phase of NVRAR contributes most to the overall all-reduce runtime. Apart from choosing the algorithmically optimal recursive doubling approach, we make three key optimizations for increased efficiency and lower latency: (1) chunked non-blocking communication, (2) fused data-flag payloads for per-step synchronization, and (3) sequence number based global.

4.2.1 Chunked Non-Blocking Communication

To utilize the GPU SMs efficiently and increase concurrency, we divide the message into disjoint data blocks, processed independently by B thread blocks. Data blocks are further subdivided into smaller chunks of size C bytes (Lines 15-21 in Algorithm 4.1) by the thread block. Each chunk is then transmitted to the corresponding peer GPU using a non-blocking, block-cooperative NVSHMEM primitive (Line

20). Upon receiving the peer’s chunk, the data is locally reduced, and the thread block advances to the next chunk.

This design enables different thread blocks to progress through distinct stages of the all-reduce operation in parallel, creating coarse-grained overlap between computation and communication across SMs in the GPU. The non-blocking communication allows sending and waiting for peer data to proceed asynchronously, while per-block chunking offers tunable control for the granularity of network injection. B and C can be tuned once for a given message size and node-count. In practice, we observe significant performance impact of these hyperparameters.

To further enable asynchronous progress across inter-node steps, we use pre-allocated send and receive buffers per step. This allows data to be sent to the next peer, before the receiving peer has completed its previous step. The extra memory overhead is negligible for small messages and logarithmic recursion depth.

4.2.2 Fused Payloads for Step Synchronization

At each recursive-doubling step, synchronization is required between peers to ensure completion of remote data receipt before its local reduction. A naive approach is to use NVSHMEM’s `put_with_signal` and `wait_until` primitives. However, we found that these explicit signaling primitives introduce non-trivial latency overheads, particularly on Slingshot. The root cause lies in NVSHMEM’s current libfabric implementation, where `put_with_signal` relies on software fences instead of the hardware fences and message ordering available on Slingshot.

To avoid explicit signaling, we adopt NCCL’s low-latency LL protocol design, fusing data and synchronization flags into a single 8B payload (4B data + 4B flag). This granularity ensures atomic and ordered delivery of each data word and its flag, on both Slingshot and InfiniBand. Fused payloads also allow reductions to begin immediately upon receipt (at a warp level), enabling fine-grained progress and synchronization without extra communication overhead.

4.2.3 Sequence Number Based Global Synchronization

When issued in succession, it is important to ensure that previous all-reduce operations have completed to safely reuse the intermediate buffers. NVSHMEM provides `quiet` and `fence` primitives to achieve this, but they add significant latency overheads. Instead, we use a sequence number based approach where each all-reduce operation is assigned a unique sequence number (Line 2 of Algorithm 4.1), and each rank waits for its peers to reach the same sequence number (Lines 4-6 of Algorithm 4.1) before sending data in the inter-node communication phase.

Crucially, each rank only waits for its peers and not all ranks

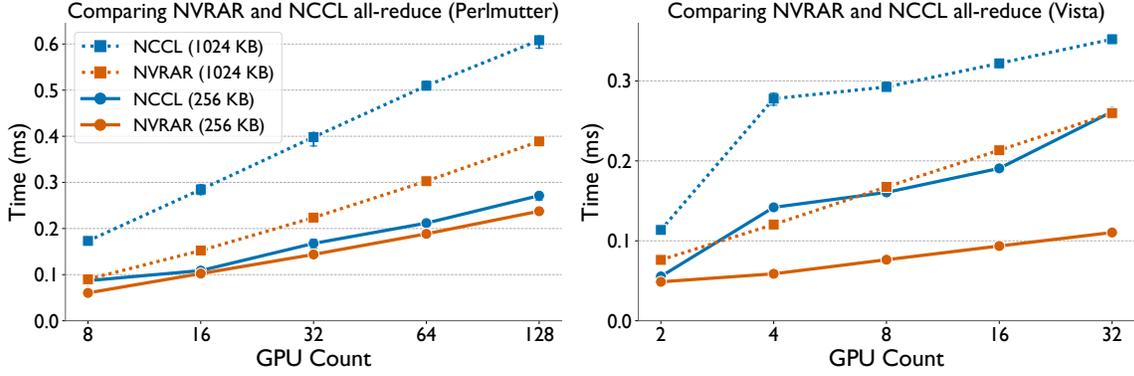


Figure 7. Performance comparison of NVRAR and NCCL all-reduce for 256 KB and 1024 KB input sizes, across varying GPU counts on Perlmutter (A100, Slingshot-11) (left) and Vista (GH200, InfiniBand) (right).

at the same time (avoiding a barrier-like synchronization). This check is performed at the beginning of the all-reduce operation, allowing for a rank to finish its all-reduce and use the data immediately, without waiting for peer ranks to finish. The waiting only occurs when the next all-reduce is issued. This is implemented using NVSHMEM’s atomics.

4.3 Performance Model for NVRAR

To reason about the performance of NVRAR, we model its communication time using the α - β model, assuming the same notations as in Section 2.

Reduce-scatter Phase: Within a node, NVRAR uses NCCL’s reduce-scatter operation (ring algorithm). Thus, the communication time is given by:

$$T_{RS} = (G - 1)\alpha_{\text{intra}} + \frac{G - 1}{G} \left(\frac{|M|}{\beta_{\text{intra}}} \right) \quad (5)$$

Inter-Node Recursive-Doubling Phase: The inter-node phases proceeds in $\log_2(N)$ steps across N nodes, with a message size of $|M|/G$. Packing data and flag leads to a $1 < \eta < 2$ factor increase in the message size. Each step requires a single exchange between peers and thus, the communication time is:

$$T_{RD} = \log_2(N)\alpha_{\text{inter}} + \frac{N - 1}{N} \left(\frac{\eta|M|}{G\beta_{\text{inter}}} \right) \quad (6)$$

All-gather Phase: Finally, results are aggregated within each node by NCCL’s all-gather (Ring), modeled as:

$$T_{AG} = (G - 1)\alpha_{\text{intra}} + \frac{G - 1}{G} \left(\frac{|M|}{\beta_{\text{intra}}} \right) \quad (7)$$

Total Communication Time: The total time for the hierarchical NVRAR algorithm is the sum of the three phases:

$$T_{\text{NVRAR}} = 2(G - 1)\alpha_{\text{intra}} + \log_2(N)\alpha_{\text{inter}} + \frac{|M|}{G} \left[\frac{2(G - 1)}{\beta_{\text{intra}}} + \frac{(N - 1)\eta}{N\beta_{\text{inter}}} \right] \quad (8)$$

If we ignore the bandwidth terms, for small latency-bound messages, we can approximate the total time as:

$$T_{\text{NVRAR}} \approx 2(G - 1)\alpha_{\text{intra}} + \log_2(N)\alpha_{\text{inter}} \quad (9)$$

Compared to Ring all-reduce (Equation 2), NVRAR scales logarithmically with the number of nodes, rather than linearly. Relative to Tree all-reduce (Equation 4), NVRAR also exhibits $\mathcal{O}(\log_2(N))$ scaling, but with a lower inter-node latency coefficient—since each recursive-doubling step involves a single exchange. Thus, theoretically, NVRAR achieves the same asymptotic behavior as Tree all-reduce, but with reduced latency.

5 RESULTS

This section describes the experimental setup and presents a detailed performance evaluation of NVRAR against NCCL, both as an independent collective primitive and within the context of tensor-parallel inference workloads.

Additional Setup Details: We use the Perlmutter and Vista machines (Table 1) for our evaluations. To isolate collective performance, we run a microbenchmark that executes NCCL all-reduce and NVRAR, each within a CUDA Graph, for 100 consecutive iterations. It replays the captured graph 1000 times (200 warm-up). The average all-reduce time is reported. CUDA Graphs help mimic inference workloads more accurately. NCCL 2.27.3 and PyTorch 2.8 are used for all experiments. We modify YALIS to use NVRAR for all-reduce operations in TP and benchmark end-to-end performance for decode-heavy workloads with both NCCL and NVRAR, following the Section 3 methodology.

5.1 Comparison of NVRAR and NCCL

Figure 7 shows all-reduce microbenchmark performance on Perlmutter and Vista. On Perlmutter (left), NVRAR (orange) scales linearly with GPU count on a logarithmic X-

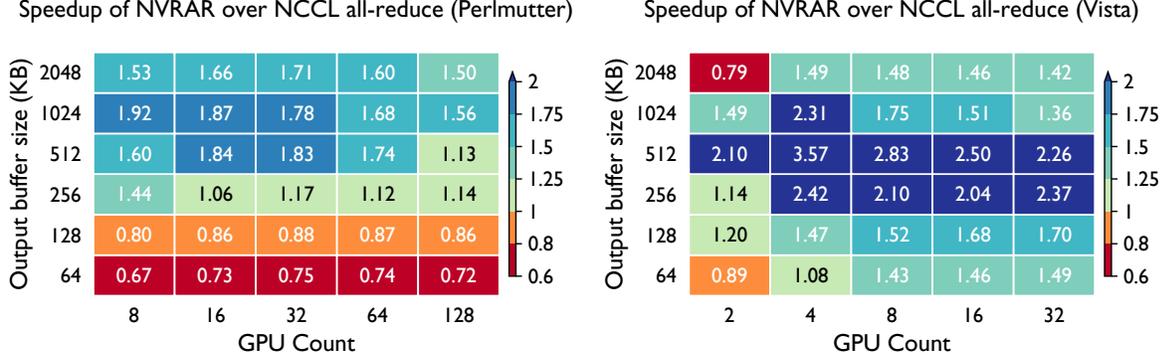


Figure 8. Heatmaps showing the speedup of NVRAR over NCCL all-reduce in the standalone microbenchmark on Perlmutter and Vista.

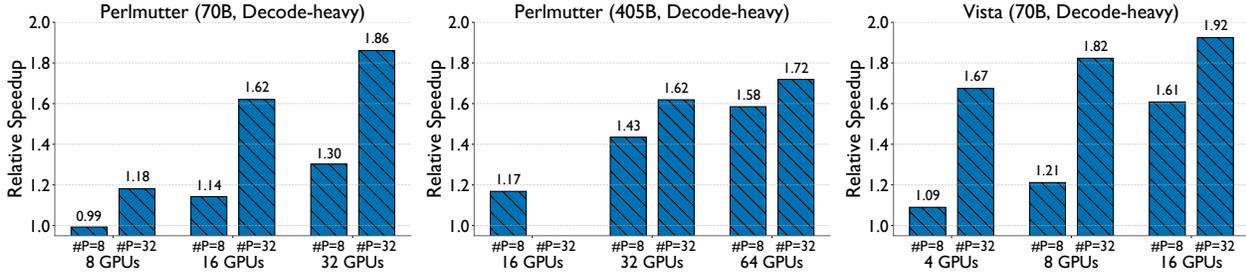


Figure 9. Relative Speedup of YALIS (TP) using NVRAR all-reduce over Yalis (TP) using NCCL all-reduce for the decode-heavy workload on Perlmutter and Vista, across different models and NumPrompts (#P = 8 and #P = 32).

axis, consistent with our theoretical model (Eq. 9), for both 256 KB and 1024 KB messages. NCCL (blue) exhibits similar scaling for 1024 KB messages, as it consistently uses the Tree algorithm (LL protocol) across all node counts. This provides a good point of comparison where both algorithms scale logarithmically. NVRAR achieves better performance due to its lower latency coefficients, consistent with our model. For 256 KB messages, however, NCCL switches from the Ring to the Tree algorithm beyond 16 GPUs, complicating the analysis. Nevertheless, with respect to empirical performance, NVRAR outperforms NCCL at most node counts for both message sizes. On Vista (right), similar trends are observed.

Figure 8 shows the relative speedup of NVRAR over NCCL all-reduce for a range of message sizes and node counts. On Perlmutter (left), NVRAR is ~ 0.7 - 0.8 x as fast as NCCL for 64 KB and 128 KB messages, which we attribute partly to kernel launch overheads introduced by its three-phase design. Profiling also revealed that the intra-node all-gather phase takes an unexpectedly long time in the microbenchmark setup, but not in real end-to-end workloads (discussed in the next section). For 256 KB messages, we observe modest speedups of 1.06-1.44x with NVRAR. Beyond 256 KB, NVRAR outperforms NCCL by 1.13-1.92x. On Vista (right), the speedups are considerably higher. For GPU

counts greater than 4, NVRAR outperforms NCCL by 1.08-1.70x for 64 KB and 128 KB messages. Between 256 KB and 1 MB, NVRAR can achieve up to 3.5x speedups over NCCL. We attribute these higher speedups to the underlying architecture of Vista, where each node has a single GPU. As a result, NVRAR only executes the inter-node recursive-doubling phase, reducing kernel launch overheads and avoiding previously mentioned all-gather issues.

5.2 Improving Multi-Node TP Inference Performance

We now evaluate the performance improvements achieved by integrating NVRAR into YALIS for multi-node TP inference on decode-heavy workloads (Table 2).

Figure 9 shows the relative speedup of YALIS (TP) using NVRAR over NCCL all-reduce for 70B and 405B models, on Perlmutter and Vista. On Perlmutter, for the 70B model (left), NVRAR accelerates YALIS (TP) by 1.3x for NumPrompts (#P) = 8 and 32 GPUs. For this configuration, the all-reduce message size is 128 KB, for which our microbenchmark setup showed slowdowns with NVRAR in the previous section (Figure 8). To investigate this discrepancy, we profiled NVRAR in both setups, and found that the intra-node all-gather phase is 4-5x slower in the microbenchmark than in YALIS (see appendix). Although the root cause re-

mains unclear, we conclude that our microbenchmark setup does not perfectly capture the workload nuances, with factors such as cache performance and back-to-back kernel launches likely contributing to the difference. Nevertheless, overall speedup trends are consistent. NVRAR achieves a higher 1.86x speedup for NumPrompts = 32 on 32 GPUs (message size 512 KB). For the 405B model (middle plot), speedups range from 1.17x to 1.72x, again higher due to the more favorable message sizes (256 KB and 1024 KB) compared to the 70B model. On Vista (right), NVRAR accelerates YALIS by up to 1.92x for the 70B model with NumPrompts = 32 and 16 GPUs. Speedups on Vista are consistently higher than on Perlmutter, in agreement with our microbenchmarking.

Finally, we analyze the performance breakdowns of YALIS (TP) using NVRAR and NCCL all-reduce on 16 GPUs of Perlmutter for the 70B model (Figure 10). For NumPrompts = 8 and 32, we observe that communication time using NVRAR is lower than NCCL all-reduce. The decrease is more pronounced for NumPrompts = 32 compared to NumPrompts = 8, due to the more favorable message size. Idle time is marginally higher for NVRAR, but not enough to offset the overall performance gains. We plan to investigate and reduce idle time in future versions of NVRAR.

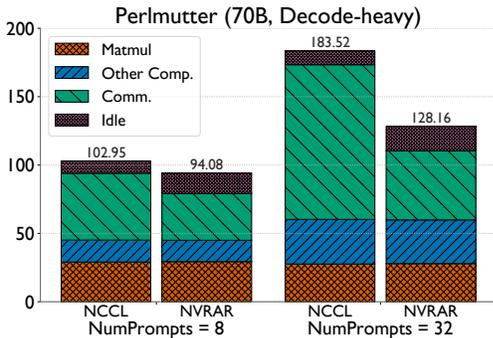


Figure 10. Performance breakdown of YALIS using NVRAR and NCCL all-reduce on 16 GPUs of Perlmutter for the 70B model.

Our evaluation demonstrates that NVRAR achieves strong performance improvements over NCCL all-reduce for small message sizes, characteristic of decode-heavy workloads. Integrating NVRAR leads to significant gains in multi-node TP performance. In the future, we will explore the use of NVRAR in other latency-bound applications.

6 RELATED WORK

Model Parallel Performance Studies: Prior work investigates model-parallel scaling on multiple nodes, often in tandem with optimization efforts (Xu et al., 2025; Zhang et al., 2025a; Su et al., 2025; Alvarez, 2025; Zhu et al., 2025a; Spector et al., 2025). vLLM benchmarks Llama 3.1

405B on up to 16 GPUs on Infiniband and non-Infiniband networks, exposing TP as weak with slow inter-node interconnects (vLLM Team, 2024). Google scales up to 64-way TP inference, but their solution is specific to TPUs (Pope et al., 2022). To the best of our knowledge, no prior work systematically studies performance breakdowns of the distinct components of inference for both TP and PP, for large models in large multi-node settings.

Collective Communication Optimization: Recent studies examine collective optimizations as a means to improve distributed inference. Some (Li et al., 2024a; Hansen-Palmus et al., 2024) leverage compression to alleviate communication bottlenecks, but these methods are typically not exact and are orthogonal to our work. A different approach taken is to hide communication by overlapping it with computation. ISO (Xiao & Su, 2024) achieves overlap in pre-fill, but not in decode, where our approach is most beneficial. In other cases (Zhang et al., 2025a), overlap is achieved with model architecture changes, which our approach avoids. Similar to our work, StragglAR (Devraj et al., 2025) introduces a novel all-reduce collective, focusing on reducing stragglers in bandwidth-bound regimes, as opposed to latency-bound regimes that we target. Other recent works (Zhang et al., 2025b; Aimuyo et al., 2025; Zhu et al., 2025b) explore NVSHMEM-based collective optimizations for expert parallelism in mixture-of-experts (MoE) inference, unlike our work which focuses on TP.

7 CONCLUSION

In this work, we conduct a detailed performance study of model-parallelism schemes for multi-node LLM inference workloads. We compare the performance characteristics of Tensor Parallelism (TP) and Hybrid Parallelism (TP + PP) across different workloads and identify the scaling bottlenecks for each strategy. Focusing on workloads that benefit from TP, we observe severe communication bottlenecks arising from poor NCCL all-reduce performance beyond a single node. To address this, we propose NVRAR, a hierarchical recursive all-reduce implementation built using NVSHMEM. We make several key optimizations in NVRAR to reduce latencies for small-message all-reduce operations. When evaluated against NCCL all-reduce, NVRAR achieves 1.06-1.92x speedups on Slingshot and 1.14-3.57x on InfiniBand for 256 KB-2 MB message sizes. Integrating NVRAR into YALIS, a prototype inference engine developed for this work, achieves up to 1.92x faster multi-node tensor-parallel inference performance in decode-heavy workloads.

ACKNOWLEDGMENTS

This work was performed in part under the auspices of the U.S. Department of Energy by Lawrence Livermore

National Laboratory under Contract DE-AC52-07NA27344 (LLNL-CONF-2013350).

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility, operated under Contract No. DE-AC02-05CH11231 using NERSC award DDR-ERCAP0034262 and ALCC-ERCAP0034775. This research is supported by the National Artificial Intelligence Research Resource (NAIRR) Pilot and used the Delta advanced computing and data resource which is supported by the NSF (award NSF-OAC 2005572) and the State of Illinois, and the Vista supercomputing resource at the Texas Advanced Computing Center (TACC) at The University of Texas at Austin. The authors acknowledge the University of Maryland supercomputing resources made available for conducting the research reported in this paper. This work was supported by a grant from the Swiss National Supercomputing Centre (CSCS) under project ID lp98 on Alps.

REFERENCES

- Aimuyo, O. J., Oh, B., and Singh, R. Flashdmoe: Fast distributed moe in a single kernel, June 2025. URL <https://arxiv.org/abs/2506.04667>.
- Alvarez, E. Analyzing the impact of tensor parallelism configurations on LLM inference performance. March 2025. URL <https://rocm.blogs.amd.com/artificial-intelligence/tensor-parallelism/README.html>.
- Aminabadi, R. Y., Rajbhandari, S., Awan, A. A., Li, C., Li, D., Zheng, E., Ruwase, O., Smith, S., Zhang, M., Rasley, J., et al. DeepSpeed-inference: enabling efficient inference of transformer models at unprecedented scale. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–15. IEEE, 2022.
- Bhatele, A., Dhakal, R., Movsesyan, A., Ranjan, A. K., and Cankur, O. Pipit: Scripting the analysis of parallel execution traces, 2023.
- Bick, A., Blandin, A., and Deming, D. J. The rapid adoption of generative ai. Technical report, National Bureau of Economic Research, 2024.
- Chapman, B., Curtis, T., Pophale, S., Poole, S., Kuehn, J., Koelbel, C., and Smith, L. Introducing openshmem: Shmem for the pgas community. In *Proceedings of the Fourth Conference on Partitioned Global Address Space Programming Model, PGAS '10*, New York, NY, USA, 2010. Association for Computing Machinery. ISBN 9781450304610. doi: 10.1145/2020373.2020375. URL <https://doi.org/10.1145/2020373.2020375>.
- Devraj, A., Ding, E., Vijaya Kumar, A., Kleinberg, R., and Singh, R. Accelerating AllReduce with a persistent straggler, May 2025. URL <https://arxiv.org/abs/2505.23523>.
- Grattafiori, A., Dubey, A., Jauhri, A., et al. The llama 3 herd of models, 2024. URL <https://arxiv.org/abs/2407.21783>.
- Gropp, W., Lusk, E. R., Thakur, R., Balaji, P., Gillis, T., Guo, Y., Latham, R., Raffanetti, K., and Zhou, H. Mpi. [Computer Software] <https://doi.org/10.11578/dc.20200514.13>, jun 2023. URL <https://doi.org/10.11578/dc.20200514.13>.
- Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*, 2025.
- Hansen-Palmus, J., Le, M. T., Hausdörfer, O., and Verma, A. Communication compression for tensor parallel llm inference. *arXiv preprint arXiv:2411.09510*, 2024.
- Hockney, R. W. The communication challenge for mpp: Intel paragon and meiko cs-2. *Parallel Comput.*, 20(3): 389–398, March 1994. ISSN 0167-8191. doi: 10.1016/S0167-8191(06)80021-9. URL [https://doi.org/10.1016/S0167-8191\(06\)80021-9](https://doi.org/10.1016/S0167-8191(06)80021-9).
- Hu, Z., Shen, S., Bonato, T., Jeaugey, S., Alexander, C., Spada, E., Dinan, J., Hammond, J., and Hoefler, T. Demystifying nccl: An in-depth analysis of gpu communication protocols and algorithms, 2025.
- Huang, Y., Cheng, Y., Bapna, A., Firat, O., Chen, M. X., Chen, D., Lee, H., Ngiam, J., Le, Q. V., Wu, Y., and Chen, Z. GPipe: Efficient training of giant neural networks using pipeline parallelism, 2019.
- International Energy Agency. Energy and ai, April 2025. URL <https://www.iea.org/reports/energy-and-ai>. Licence: CC BY 4.0.
- Jaech, A., Kalai, A., Lerer, A., Richardson, A., El-Kishky, A., Low, A., Helyar, A., Madry, A., Beutel, A., Carney, A., et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.
- Kwon, W., Li, Z., Zhuang, S., Sheng, Y., Zheng, L., Yu, C. H., Gonzalez, J. E., Zhang, H., and Stoica, I. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*, 2023.
- Leviathan, Y., Kalman, M., and Matias, Y. Fast inference from transformers via speculative decoding, 2023. URL <https://arxiv.org/abs/2211.17192>.

- Li, Q., Zhang, B., Ye, L., Zhang, Y., Wu, W., Sun, Y., Ma, L., and Xie, Y. Flash communication: Reducing tensor parallelization bottleneck for fast large language model inference, December 2024a. URL <https://arxiv.org/abs/2412.04964>.
- Li, Q., Zhang, B., Ye, L., Zhang, Y., Wu, W., Sun, Y., Ma, L., and Xie, Y. Flash communication: Reducing tensor parallelization bottleneck for fast large language model inference, 2024b. URL <https://arxiv.org/abs/2412.04964>.
- Lightning AI. Litgpt. <https://github.com/Lightning-AI/litgpt>, 2023.
- Maslej, N., Fattorini, L., Perrault, R., Gil, Y., Parli, V., Kariuki, N., Capstick, E., Reuel, A., Brynjolfsson, E., Etchemendy, J., Ligett, K., Lyons, T., Manyika, J., Niebles, J. C., Shoham, Y., Wald, R., Walsh, T., Hamrah, A., Santarlasci, L., Lotufo, J. B., Rome, A., Shi, A., and Oak, S. Artificial intelligence index report 2025, 2025. URL <https://arxiv.org/abs/2504.07139>.
- Meta. Torch compile. https://docs.pytorch.org/tutorials/intermediate/torch_compile_tutorial.html, 2023.
- NERSC. Perlmutter system architecture. <https://docs.nersc.gov/systems/perlmutter/architecture/>.
- NVIDIA. Nvidia nsight systems. <https://developer.nvidia.com/nsight-systems>.
- NVIDIA. Nccl tests. <https://github.com/NVIDIA/nccl-tests>, 2017.
- NVIDIA. Cuda graphs. <https://developer.nvidia.com/blog/cuda-graphs/>, 2019a.
- NVIDIA. Massively scale your deep learning training with nccl 2.4. <https://developer.nvidia.com/blog/massively-scale-deep-learning-training-nccl-2.4/>, 2019b.
- NVIDIA. Nccl. <https://docs.nvidia.com/deeplearning/nccl/user-guide/docs/overview.html>, 2020a.
- NVIDIA. Nvshmem. <https://developer.nvidia.com/nvshmem>, 2020b.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., and Garnett, R. (eds.), *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>.
- Pope, R., Douglas, S., Chowdhery, A., Devlin, J., Bradbury, J., Levskaya, A., Heek, J., Xiao, K., Agrawal, S., and Dean, J. Efficiently scaling transformer inference, 2022.
- Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., and Catanzaro, B. Megatron-lm: Training multi-billion parameter language models using model parallelism. Technical report, 2020.
- Singh, S. and Bhatele, A. AxoNN: An asynchronous, message-driven parallel framework for extreme-scale deep learning. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '22*. IEEE Computer Society, May 2022.
- Singh, S., Singhanian, P., Ranjan, A., Kirchenbauer, J., Geiping, J., Wen, Y., Jain, N., Hans, A., Shu, M., Tomar, A., Goldstein, T., and Bhatele, A. Democratizing AI: Open-source scalable LLM training on GPU-based supercomputers. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '24*, November 2024a.
- Singh, S., Singhanian, P., Ranjan, A. K., Sating, Z., and Bhatele, A. A 4d hybrid algorithm to scale parallel training to thousands of gpus, 2024b.
- Snell, C., Lee, J., Xu, K., and Kumar, A. Scaling llm test-time compute optimally can be more effective than scaling model parameters, 2024. URL <https://arxiv.org/abs/2408.03314>.
- Spector, B., Juravsky, J., Sul, S., Lim, D., Dugan, O., Arora, S., and Ré, C. We Bought the Whole GPU, So We're Damn Well Going to Use the Whole GPU, sep 2025. URL <https://hazyresearch.stanford.edu/blog/2025-09-28-tp-llama-main>. Hazy Research Blog.
- Su, Q., Zhao, W., Li, X., Andoorveedu, M., Jiang, C., Zhu, Z., Song, K., Giannoula, C., and Pekhimenko, G. Seesaw: High-throughput llm inference via model re-sharding. *arXiv preprint arXiv:2503.06433*, 2025.

- Thakur, R. and Gropp, W. D. Improving the performance of collective operations in mpich. In Dongarra, J., Laforenza, D., and Orlando, S. (eds.), *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, pp. 257–267, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. ISBN 978-3-540-39924-7.
- University, O. S. Osu micro-benchmarks 5.8. <http://mvapich.cse.ohio-state.edu/benchmarks/>.
- vLLM Team. Announcing Llama 3.1 support in vLLM. *vLLM Blog*. <https://blog.vllm.ai/2024/07/23/llama31.html>, July 2024. URL <https://blog.vllm.ai/2024/07/23/llama31.html>. Accessed: 2025-10-30.
- Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q. V., Zhou, D., et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837, 2022.
- Xiao, B. and Su, L. ISO: Overlap of computation and communication within sequence for LLM inference, September 2024. URL <https://arxiv.org/abs/2409.11155>.
- Xu, L., Suresh, K. K., Anthony, Q., Alnaasan, N., and Panda, D. K. Characterizing communication patterns in distributed large language model inference. *arXiv preprint arXiv:2507.14392*, 2025.
- Zhang, M., Mishra, M., Zhou, Z., Brandon, W., Wang, J., Kim, Y., Ragan-Kelley, J., Song, S. L., Athiwaratkun, B., and Dao, T. Ladder-residual: parallelism-aware architecture for accelerating large model inference with communication overlapping. *arXiv preprint arXiv:2501.06589*, 2025a.
- Zhang, S., Zheng, N., Lin, H., Jiang, Z., Bao, W., Jiang, C., Hou, Q., Cui, W., Zheng, S., Chang, L.-W., Chen, Q., and Liu, X. Comet: Fine-grained computation-communication overlapping for mixture-of-experts. *ArXiv*, abs/2502.19811, 2025b. URL <https://api.semanticscholar.org/CorpusID:276647637>.
- Zheng, L., Yin, L., Xie, Z., Sun, C. L., Huang, J., Yu, C. H., Cao, S., Kozyrakis, C., Stoica, I., Gonzalez, J. E., et al. Sglang: Efficient execution of structured language model programs. *Advances in neural information processing systems*, 37:62557–62583, 2024.
- Zhu, K., Gao, Y., Zhao, Y., Zhao, L., Zuo, G., Gu, Y., Xie, D., Ye, Z., Kamahori, K., Lin, C.-Y., et al. {NanoFlow}: Towards optimal large language model serving throughput. In *19th USENIX Symposium on Operating Systems Design and Implementation (OSDI 25)*, pp. 749–765, 2025a.
- Zhu, R., Jiang, Z., Jin, C., Wu, P., Stuardo, C. A., Wang, D., Zhang, X., Zhou, H., Wei, H., Cheng, Y., Xiao, J., Zhang, X., Liu, L., Lin, H., Chang, L.-W., Ye, J., Yu, X., Liu, X., and Jin, X. Megascale-infer: Serving mixture-of-experts at scale with disaggregated expert parallelism, April 2025b. URL <https://arxiv.org/abs/2504.02263>.

A EXTENDED PERFORMANCE STUDY RESULTS

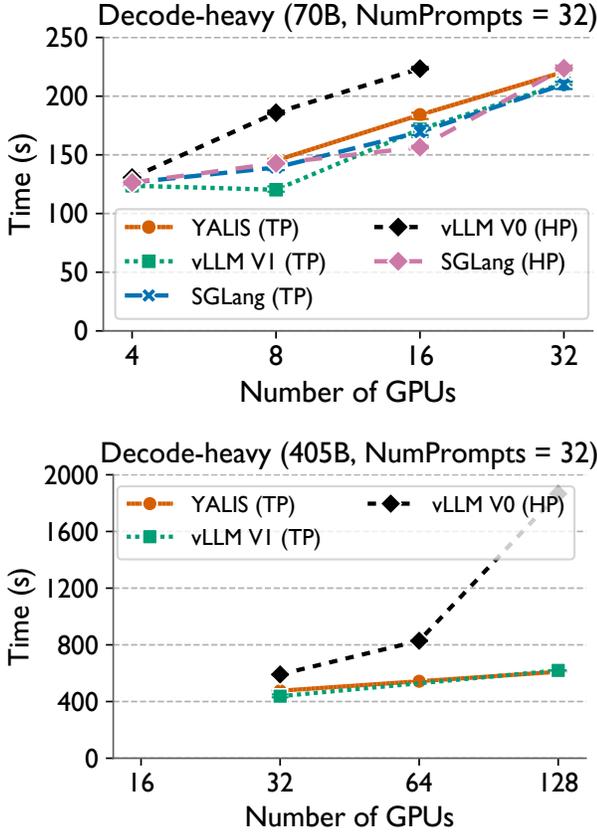


Figure 11. Strong scaling performance of different inference engines on Perlmutter for the Llama 3.1 70B (top) and 405B (bottom) models, for the Decode-Heavy with NumPrompts = 32. The Y-axis shows the time to completion for a batch of prompts in seconds and the X-axis shows the number of GPUs.

Figure 11 shows the strong scaling performance of different inference engines on Perlmutter for the Llama 3.1 70B and 405B model, for the decode-heavy workload with NumPrompts = 32. We observe that vLLM V0 (HP) (black line) scales poorly for both the 70B and 405B model. Both TP and HP have poor scaling with increasing total time with increasing number of GPUs for the 70B model (left). Interestingly, SGLang (HP) (pink line) has times closer to all the TP configurations for the 70B model (left), unlike vLLM V0 (HP) (black line). Nevertheless, the scaling performance is still poor. This closely matches our observation in the main text.

Table 4. Different NVRAR hyper-parameter configurations for an all-reduce message size of 1024 KB on 16 GPUs.

Block Size (B)	Chunk Size (C)	Time (ms)
32	32768	0.1522 ms
32	4096	0.2271 ms
8	16384	0.1891 ms
8	131072	0.1655 ms

B IMPACT OF CHUNK SIZE AND BLOCK SIZE ON NVRAR PERFORMANCE

In this section, we show that the hyper-parameters: Block Size (B) and Chunk Size (C) - have a significant impact on the performance of NVRAR. For tuning these hyper-parameters, we run NVRAR with different values of B and C for different message sizes and node counts. Table 4 shows the performance of NVRAR with 4 different hyper-parameter configurations for an all-reduce message size of 1024 KB on 16 GPUs. We observe that the performance is significantly better impacted by changing the Chunk Size (C) than the Block Size (B). This validates our design choice of keeping these hyper-parameters tunable. We also see that chunking big messages within a block is helpful for improving performance, as smaller chunk size runs have better performance than some larger chunk size runs. In future work, we plan to heuristically tune these hyper-parameters for different message sizes and node counts.

C ALL-GATHER PERFORMANCE VARIABILITY IN MICRO-BENCHMARK VS. END-TO-END WORKLOAD

Micro-benchmark	YALIS
ncclDevKernel_ReduceScatter...	ncclDevKernel_ReduceScatter... 10.304 μ s
bump_seq(unsigned int *, unsi...	bump_seq(unsigned int *, unsi... 2.528 μ s
void pack_payloads_kernel<_n...	void pack_payloads_kernel<_n... 2.208 μ s
void recursive_allreduce_kerne...	void recursive_allreduce_kerne... 28.480 μ s
void unpack_payloads_kernel<_...	void unpack_payloads_kernel<_... 4.576 μ s
ncclDevKernel_AllGather_RING...	ncclDevKernel_AllGather_RING... 9.216 μ s

Figure 12. NSYS profile snapshots of the NVRAR kernel in the microbenchmark setup (left) and YALIS (right).

Figure 12 shows the NSYS profile snapshots of NVRAR kernel in the microbenchmark setup (top) and YALIS (bottom). We observe that the time spent in the all-gather phase differs significantly in the two setups. For the microbenchmark setup, 45 μ s are spent in this phase, which can even be higher than the inter-node reduce phase. In YALIS, however, only \sim 10 μ s are spent in this phase, which is similar to the time spent in the intra-node reduce scatter phase, as expected. This suggests that the microbenchmark does not perfectly capture end-to-end workload nuances.