ONE-SIDED EXPERT-PARALLEL ALL-TO-ALL FOR MIXTURE-OF-EXPERTS INFERENCE OVER NVLINK

Bo Li [1] **Chulian Zhang** [1] **Or Biri** [1] **Anthony Chang** [1] **Anton Korzh** [1] **Noam Duvdevani** [1] **XingFei Xi** [1]
Daniel Stokes [1] **Zhenhuan Chen** [1] **Ludwig Schneider** [1] **Artem Y. Polyakov** [1] **Julien Demouth** [1] **Chandler Zhou** [1]

arXiv:2611.13188v1 [cs.DC] 9 Nov 2026

ABSTRACT

Expert parallelism (EP) shards the experts of a Mixture-of-Experts (MoE) layer across GPUs, so every MoE layer requires two all-to-all communications, namely dispatch and combine. Compared with training, inference processes relatively few tokens per GPU at a time, and at small batch sizes fixed costs that do not grow with the number of tokens, rather than data movement, dominate the latency. We present `NVLinkOneSided`, a one-sided EP all-to-all for GPUs within a single NVLink domain. Each source rank pushes its tokens, without first exchanging metadata, into rank-major buffers on the target GPUs, and the grouped GEMMs read them in place instead of from a per-expert copy. We provide two protocols for the synchronization that the communication requires, namely FENCE and SENTINEL, for high-throughput and low-latency scenarios respectively. On Blackwell GPUs, `NVLinkOneSided` completes dispatch and combine 1.2–3.0× faster than DeepEP, with the largest speedups at small batch sizes. The resulting MoE layer is up to 2.4× faster than DeepEP with DeepGEMM. It even outperforms the fused DeepGEMM MegaMoE at up to 128 tokens per GPU, making it still competitive for decoding. We expose `NVLinkOneSided` through a modular interface, so that it is agnostic to the activation quantization format, decoupled from the MoE compute backend and easy to integrate into inference frameworks.

1 INTRODUCTION

Mixture-of-Experts (MoE) models activate only a few of their many experts for each token (Shazeer et al., 2017; Lepikhin et al., 2021; Fedus et al., 2022). They are commonly served with data-parallel (DP) attention and expert-parallel (EP) MoE layers, a deployment strategy popularized by DeepSeek-V3 (DeepSeek-AI, 2024). In the attention layers, every rank serves its own requests and processes only their tokens. In the MoE layers, every rank holds a subset of the experts, so every MoE layer contains two all-to-all communications. *Dispatch* sends each token to the ranks that own its experts, and *combine* returns the expert outputs to the source rank of the token.

Compared with training, inference processes few tokens per GPU at a time. Prefill is usually split into chunks of thousands to tens of thousands of tokens (Agrawal et al., 2024; DeepSeek-AI, 2026a), and decoding must meet a service-level objective on the time per output token, which bounds its batch to a few hundred tokens per rank. DeepSeek's DeepInfer (DeepSeek-AI, 2026a), for example, decodes DeepSeek-V4.1-Flash at about 250 tokens per second per user while serving 64 requests per rank; each request adds three draft tokens, so a step processes 256 tokens per rank.

At large batches, dispatch and combine are bound by the communication bandwidth, so the goal is to use as much of it as possible and to eliminate unnecessary data movement. At small batches, however, every microsecond matters (Shen et al., 2026). At a few tokens per rank, the bytes a rank sends would take less than a microsecond at the theoretical NVLink bandwidth, yet measured round trips of dispatch and combine take tens of microseconds (§5) in every MoE layer: the time goes to fixed costs that do not grow with the number of tokens, such as synchronization between producers and consumers on different GPUs. In common designs such as NCCL (NVIDIA, a) and DeepEP (Zhao et al., 2025), the sender writes data into an intermediate buffer, from which the receiver copies or regroups it into the layout that the next kernel needs; this extra pass over the data adds latency at small batches and consumes memory bandwidth at large ones.

Within one NVLink domain, symmetric memory lets a rank load from and store to a peer's memory directly, so data can be placed exactly where its consumer reads it. We present `NVLinkOneSided`, a one-sided EP all-to-all built on this direct access. The routing decisions of a rank alone determine where each of its tokens goes, so dispatch can push the tokens into the memory of the receivers without first sending routing metadata such as token counts or expert

[1]NVIDIA. Correspondence to: Bo Li <lbo@nvidia.com>, Chulian Zhang <chulianz@nvidia.com>.

Preprint.

ids. The dispatch *workspace* of each rank, a buffer in symmetric memory, is divided into one region per source rank, so that different ranks do not interfere when writing to the same target rank; within its region, a source rank places each token in one slot, however many experts of the target rank the token selects. We call this layout *rank-major*. The grouped GEMMs of the MoE layer read the tokens of each expert directly from this workspace through slot indices, without materializing the tokens into per-expert order. Dispatch and combine still need synchronization between producers and consumers, for which we provide two protocols targeting high-throughput and low-latency scenarios respectively. FENCE publishes per-peer flags after a memory fence. SENTINEL uses the data itself as the flag. We expose `NVLinkOneSided` through the general EP communication interface of FlashInfer, `MoEEpCommunication`, which makes it agnostic to the activation quantization format and easy to integrate with different inference frameworks and MoE backends.

The rest of the paper is organized as follows. Section 3 presents the one-sided communication and the rank-major layout that MoE compute reads. Section 4 explains the synchronization protocols and implementation details. Section 5 evaluates `NVLinkOneSided` and the MoE layer built on it on Blackwell GPUs, analyzes the two protocols and compares both with existing implementations.

2 BACKGROUND

2.1 Mixture-of-Experts and expert parallelism

An MoE layer replaces the dense feed-forward network with E experts and a router (Shazeer et al., 2017; Lepikhin et al., 2021; Fedus et al., 2022), which selects k experts for each token and weights their outputs. Under EP the experts are split over P GPUs (*ranks*), $E_\ell = E/P$ per rank, while attention runs in DP, so each rank holds its own T tokens. The layer then takes three steps. *Dispatch* sends each token to the ranks that own its experts, each rank runs its local experts on the tokens it received, and *combine* returns the weighted expert outputs to the source rank of each token, which sums them. Dispatch is a multicast whose destinations depend on the router's choices, and combine is the reverse, a gather with a reduction. Every rank may send tokens to every other rank, and the amounts depend on the routing, so both are all-to-all exchanges with variable sizes. Figure 1 shows the example used throughout the paper.

2.2 NVLink domains and symmetric memory

A Blackwell GPU has 18 fifth-generation NVLink links of 100 GB/s each, 900 GB/s per direction or 1.8 TB/s in total. NVLink switches fully connect the eight GPUs of an HGX B300 board, and a GB200 or GB300 NVL72 rack forms one

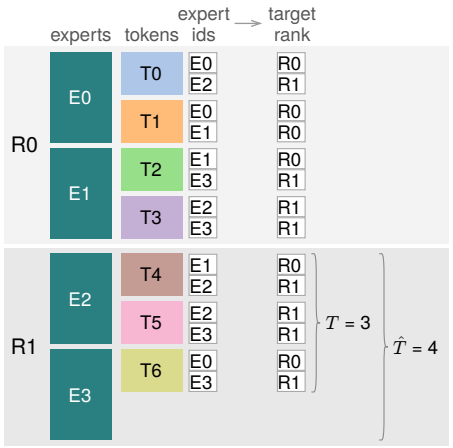


Figure 1. Example used throughout the paper: two ranks, experts E0, E1 on R0 and E2, E3 on R1, and top-2 routing. R0 holds tokens T0–T3 and R1 holds T4–T6. Beside each token are its two expert ids and, through the routing (arrow), the target rank of each: both experts of T1 are on R0, both experts of T3 and T5 are on R1, and every other token reaches both ranks. The braces on R1 mark its number of tokens, $T = 3$, and the largest number of tokens of any rank, $\hat{T} = 4$ (Table 2).

72-GPU NVLink domain. Beyond the domain, each GPU communicates through one RDMA NIC, a ConnectX-8 of 800 Gb/s, that is, 100 GB/s per direction, roughly an order of magnitude less (Aubrey & Stam, 2025).

Within the domain, GPUs address each other's memory directly. With CUDA's Virtual Memory Management, each rank allocates physical memory, exports a handle, and maps every peer's allocation into its own address space; the result is a *symmetric memory* region in which a kernel accesses peer memory with ordinary loads, stores and atomics (NVIDIA, 2025; PyTorch Contributors). RDMA offers remote access through a different interface: the initiator, a CPU thread or, with GPU-initiated networking, a GPU thread (Markthub et al., 2022), posts work requests to the NIC, names registered regions by key, and is notified of completion through completion queues or a write with immediate data (InfiniBand Trade Association, 2020); transfers are efficient for messages of kilobytes. Table 1 lists the differences that matter for an all-to-all of small, data-dependent messages. `NVLinkOneSided` moves data into and out of peer memory with 16-byte stores and loads from many threads and detects completion through the GPU memory model (§4).

A remote store does not wait for the peer: the thread moves on at once, and the data reach the peer later, possibly out of order with other stores. GPU memory is weakly ordered (Lustig et al., 2019): a consumer on another GPU that sees a producer's write is guaranteed to see the producer's earlier writes only if the producer ordered them with

Table 1. Remote access within an NVLink domain and over RDMA (Blackwell, per GPU).

	NVLink domain	RDMA NIC
Initiator	any thread, load/store	work request to the NIC
Transfer size	store ($\leq$32 B) or TMA copy	message, KB scale
Ordering	PTX fences, system scope	per queue pair
Arrival in a message	whole 32-bit words	unspecified

Table 2. Notation and the corresponding parameter names in the implementation.

Symbol	Meaning	Name in the implementation
P	EP ranks	`ep_size`
k	experts per token	`top_k`
E	experts of the layer	`num_experts`
E_ℓ	experts per rank	`num_local_experts`
T	tokens of a rank in the step	`local_num_tokens`
$\hat{T}$	largest T across ranks	`runtime_max_tokens_per_rank`

a system-scope release and the consumer reads with acquire semantics. In addition, each aligned 32-bit word stored to peer memory is observed whole or not at all (NVIDIA, 2026d). NCCL's LL protocol relies on a stronger guarantee: an 8-byte store of 4 bytes of data and a 4-byte flag arrives whole (NVIDIA, a; Hu et al., 2025). RDMA offers no counterpart: the order of placement within an RDMA Write is undefined, and the responder may rely on the data only after a later completion, such as a write with immediate data (InfiniBand Trade Association, 2020; Recio et al., 2007).

3 One-Sided Communication with Rank-Major Layout

This section describes one-sided communication, the rank-major layout and how MoE compute reads it. Table 2 summarizes the notation.

3.1 One-sided communication

Each rank routes its own tokens, so it knows locally which target ranks each token must reach. In dispatch, the source rank *pushes* each token directly, at a slot that it chooses itself, into the dispatch workspace of every target rank, which lies in symmetric memory, and records the target rank and slot of each selection (Figure 2). No metadata, such as token counts or expert ids, travels ahead of the data. MoE compute leaves the partial results in a second workspace, the combine workspace, from which combine returns them.

3.2 Rank-major layout

Source ranks push into a target rank concurrently, so their stores must not interfere. Slots drawn from one shared pool would need atomics on a remote counter, or offsets precomputed from the counts of all source ranks, which would require a metadata exchange. Instead, each dispatch workspace has one region per source rank, which only that source rank writes: it takes the next slot by atomically incrementing its *slot counter* for that target rank, and its slots form a dense prefix of the region. We call this layout *rank-major*. A region has $\hat{T}$ slots, which every rank knows because inference engines with DP attention share the batch size information every step; the dispatch workspace is sized at initialization for the maximum number of tokens a rank can process. A token goes to each target rank once, however many experts of that rank it selects, with all k expert ids and router weights in its slot. Some grouped GEMM implementations, such as DeepGEMM (DeepSeek-AI, 2025), read the tokens of each expert contiguously, so the corresponding dispatch delivers an *expert-major* layout instead, with one slot per selected expert.[1] The rank-major layout has two benefits over it:

De-duplication. A token is sent at most $\min(k, P)$ times instead of k times as in a per-expert transfer; for example, with $k = 16$ on eight GPUs, it sends at most half the bytes, whatever the routing. The saving is larger when P is small relative to k, or when a token chooses several experts located on the same rank.

Smaller workspaces. The rank-major dispatch and combine workspaces each hold $P\hat{T}$ slots. An expert-major buffer must hold every selection that a rank may receive. If every local expert may receive every token, as in the low-latency kernels of DeepEP V1, it needs E_ℓ times as many slots. Since a token selects at most $\min(k, E_\ell)$ experts of one rank, DeepEP V2.5 still needs $\min(k, E_\ell)$ times as many. For DeepSeek-V4.1-Flash on eight GPUs with up to 4096 tokens per rank, the hidden states take 160 MiB in our dispatch workspace (8×4096 slots of 5120 bytes in MXFP8) and 320 MiB in our combine workspace (BF16), besides a few MiB for scaling factors, expert ids and router weights. DeepEP V2.5 with DeepGEMM needs a 242 MiB communication buffer, shared by dispatch and combine, and an expert-major GEMM input whose hidden states alone take 1020 MiB, around $k = 6$ times the 160 MiB of our dispatch workspace (the excess is due to alignment padding). All these buffers grow linearly with the EP size, so memory becomes more pressing at large EP.

[1] The low-latency kernels of DeepEP V1 push each token once per selected expert into an expert-major buffer. DeepEP V2.5 pushes each token once per target rank into a rank-major communication buffer; with `do_expand`, which DeepGEMM requires, it then copies each token once per selected local expert into an expert-major GEMM input.

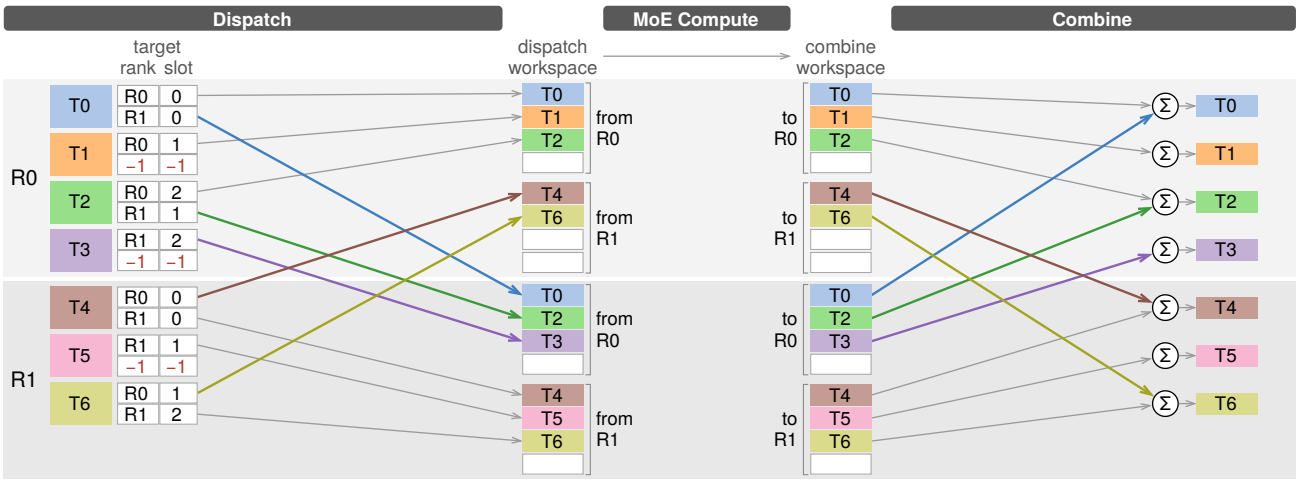


Figure 2. Data path for the routing of Figure 1, with the FENCE synchronization protocol. Dispatch: each source rank stores its tokens into its own region of the dispatch workspaces of the target ranks and records the target rank and slot of each selection; T1 selects E0 and E1, both on R0, so it takes one slot there and its second selection records no slot (-1). MoE compute turns the tokens from each source rank into partial results at the same slots of the combine workspace. Combine: each source rank loads the partial results from the target ranks and slots recorded in dispatch and adds them. White slots are empty.

3.3 MoE compute on the rank-major layout

An MoE layer on the rank-major layout runs dispatch, the MoE compute kernels (an indexing kernel, the grouped GEMMs FC1 and FC2, and finalize) and combine (Figure 3). The tokens stay at their slots of the dispatch workspace, and the grouping by expert is expressed as slot indices.

Indexing. A small indexing kernel selects the ids of the local experts, skips the slots marked invalid (details in §4), and builds per-expert lists of slot indices, plus the inverse map, from each slot to the FC2 output rows of its token, for finalize. FC1 gathers its input through these indices, so no separate pass permutes the received tokens before FC1.

FC1 and FC2. FC1 applies the up and gate projections and the activation of each expert to its tokens, and FC2 applies the down projection of the expert to the result. FC1 loads its input by index directly from the dispatch workspace: the Prims-TS grouped GEMM of FlashInfer (FlashInfer Contributors, 2026) gathers it with TMA `gather4` instructions, each of which loads 128-byte slices of four non-contiguous slots, and never loads empty slots. FC1 writes its output grouped by expert, and FC2 reads it in that order.

Finalize and combine. Finalize gathers the FC2 outputs of each token through the inverse map, scales them by the router weights, sums them over the local experts of the token, and writes the result into the combine workspace that the communication layer exposes, at the slot of the token. The finalized partial results are thus back in the rank-major layout of dispatch, so combine needs no metadata either: the source rank *pulls* the partial results of each token from the combine workspaces of its target ranks, at the slots recorded in dispatch, and adds them.

4 IMPLEMENTATION AND SYNCHRONIZATION PROTOCOLS

Kernels in one stream are ordered: a kernel sees every write to local memory made by the kernels before it. Remote memory accesses have no such guarantee: the stores of a source rank reach a peer while the kernels of the peer run, in no particular order (§2.2). One-sided dispatch and combine therefore need two producer–consumer synchronizations to prevent races:

- **End of dispatch.** Before a target rank starts MoE compute, every source rank must have finished storing its tokens into the dispatch workspace of the target rank.
- **Start of combine.** Before a source rank reads the partial results of its tokens, every target rank must have finished MoE compute and written the results into its combine workspace.

`NVLinkOneSided` provides two synchronization protocols: FENCE for high-throughput and SENTINEL for low-latency scenarios. They differ in how a source rank assigns slots and how a target rank learns that its slots are complete (Figure 4).

4.1 FENCE synchronization

Each CUDA thread block handles the dispatch and the combine of one token, and the blocks of a rank run in paral-

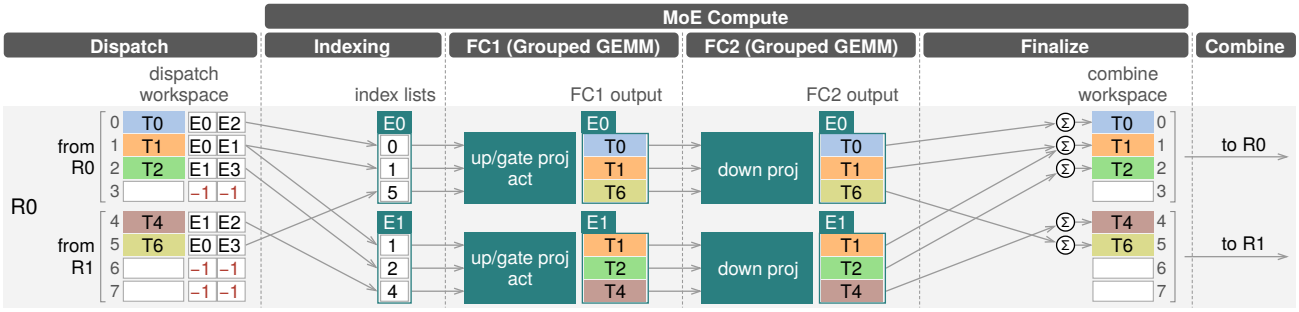


Figure 3. The MoE layer on R0, detailing Figure 2; MoE compute comprises indexing, FC1, FC2 and finalize, and slots are numbered 0–7 across the two regions. The indexing kernel turns the expert ids of the received slots into per-expert lists of slot indices (empty slots 3, 6 and 7 carry invalid ids); FC1 gathers tokens through these lists and multiplies them by the expert weights, FC2 multiplies the FC1 output by its own weights, and finalize adds the outputs of each token into its slot of the combine workspace.

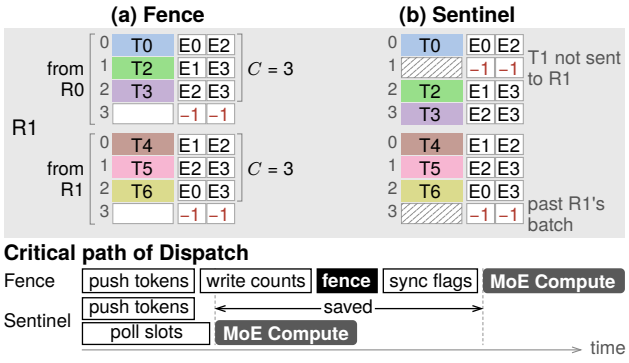


Figure 4. The dispatch workspace of R1 at the end of dispatch. (a) FENCE gathers the tokens of each source rank into a prefix of its region, so each source rank sends its slot count C; the slots past the count get invalid expert ids (-1). (b) SENTINEL pushes token i of each source rank into slot i of its region, so every slot signals its own arrival; a slot without a token keeps the sentinel payload (hatched) and gets invalid expert ids. Bottom: the critical path of dispatch until MoE compute starts; SENTINEL shortens it, which pays off at small batches.

lel (Algorithm 1). A block takes the slot of its token in each target rank by atomically incrementing the slot counter for that rank. When a block has sent its token, it increments another local counter of finished blocks, and the last block to finish synchronizes for the whole rank, so the rank executes one system-scope fence per step. This block publishes the slot counts to their target ranks, executes a system-scope release fence (`fence.release.sys`), and writes a flag to every peer rank. It spins until the flags of all peer ranks have arrived, which ends the dispatch kernel. A small kernel marks the expert ids past each published count invalid, so that the slots past the prefix take no part in MoE compute. Combine synchronizes in the same way, except that the first block alone writes the flags while every block spins on the flags of all peers and executes an acquire fence (`fence.acquire.sys`) before it pulls the finalized partial results of its token from the target ranks and slots recorded in dispatch.

4.2 SENTINEL synchronization

FENCE ends every dispatch with a serial chain: atomic counting for the slot indices, a system-scope fence, and the publishing of and spinning on completion flags. SENTINEL shortens the critical path by letting the data signal their own arrival. Source ranks push the tokens into landing slots, each 32-bit word of which is initialized with a sentinel value, a bit pattern that the data do not take or that the producer avoids: BF16 payload words, for example, use `0xFFFFFFFF`, two BF16 NaNs, and the producer stores such a word as two `0x7FFF` NaNs instead. A target rank polls the words of a landing slot; since each word arrives whole (§2.2), a slot whose words all differ from the sentinel is complete without a fence. The target rank copies it into the same slot of its dispatch workspace, which MoE compute reads as with FENCE, and resets it to the sentinel. Unlike NCCL's LL protocol (NVIDIA, a), which interleaves flag words with the data, SENTINEL adds no flag words: the data themselves are the flags.[2]

The routing decides which slots receive a token, so a target rank must tell a slot still in flight from one that will never be written. Token i of source rank s therefore always lands in slot i of region s, and for each of its $\hat{T}$ slots every source rank writes expert ids to every rank, invalid ones (-1) where no token goes (Figure 4b): every expert-id slot arrives and tells whether a payload follows. Expert-id slots use the sentinel `0x88888888`, which differs from -1.

[2]Several code bases call this Lamport synchronization, after Lamport's happened-before relation (Lamport, 1978): the arrival of the data itself orders the producer's stores before the consumer's loads.

Algorithm 1 `NVLinkOneSided` FENCE protocol

Thread block: token i on source rank s
$W_d[r]$, $W_c[r]$: dispatch and combine workspaces of rank r
$C[r][s]$, $F[r][s]$: slot count and flag of source rank s on rank r
$cnt[t]$: local slot counter for target rank t, initially 0
$done$: local counter of finished blocks, initially 0
f: flag value of the step

Dispatch
$D \leftarrow$ target ranks of token i
for all $t \in D$ **do**
  $j_t \leftarrow \text{atomicAdd}(cnt[t], 1)$ ▷ pick slot on t
  $W_d[t][s][j_t] \leftarrow$ token i ▷ push token
end for
if $\text{atomicAdd}(done, 1) = T - 1$ **then**
  $C[r][s] \leftarrow cnt[r]$ for all r ▷ write slot counts
  `fence.release.sys`
  $F[r][s] \leftarrow f$ for all r ▷ sync via completion flags
  wait until all $F[s][r] = f$
end if
for all r, $j \geq C[s][r]$ **do** ▷ sanitize expert ids
  $W_d[s][r][j].expert_ids \leftarrow -1$
end for

Combine
if first block **then**
  $F[r][s] \leftarrow f + 1$ for all r ▷ sync via completion flags
end if
wait until all $F[s][r] = f + 1$
`fence.acquire.sys`
$p_t \leftarrow W_c[t][s][j_t]$ for all $t \in D$ ▷ pull tokens in parallel
output of token $i \leftarrow \sum_{t \in D} p_t$ ▷ reduction

Algorithm 2 `NVLinkOneSided` SENTINEL protocol

Thread block: token slot $i < \hat{T}$ on source rank s
$L[r]$, $I[r]$, $M[r]$: landing, expert-id and combine slots of rank r
$W_d[r]$, $W_c[r]$: dispatch and combine workspaces of rank r
σ: sentinel, initial value of L, I and M

Dispatch
$D \leftarrow$ target ranks of token i, $\emptyset$ if $i \geq T$
$L[t][s][i] \leftarrow$ token i for all $t \in D$ ▷ push token
$I[r][s][i] \leftarrow$ expert ids of token i (-1 if $r \notin D$) for all r
for all source ranks r **do**
  wait until no word of $I[s][r][i]$ equals σ ▷ poll expert ids
  if $I[s][r][i] \neq -1$ **then**
    wait until no word of $L[s][r][i]$ equals σ ▷ poll token
    $W_d[s][r][i] \leftarrow L[s][r][i]$; $L[s][r][i] \leftarrow \sigma$ ▷ copy, reset
  end if
end for

Combine
$M[r][i][s] \leftarrow W_c[s][r][i]$ for all r with $I[s][r][i] \neq -1$ ▷ push partial results
$I[s][r][i] \leftarrow \sigma$ for all r ▷ reset expert-id slots
for all $t \in D$ **do** ▷ poll partial results
  wait until no word of $M[s][i][t]$ equals σ
  $p_t \leftarrow M[s][i][t]$; $M[s][i][t] \leftarrow \sigma$ ▷ load, reset
end for
output of token $i \leftarrow \sum_{t \in D} p_t$ ▷ reduction

Combine pushes too: once the MoE compute of a target rank has finished, it stores the partial result of slot i of region s into a combine slot on source rank s without waiting for its peers, and the block of token i polls these slots like landing slots and adds the partial results.

The kernel that polls a slot consumes it and resets it to the sentinel. A rank writes a landing or combine slot on a peer again only after receiving data that the peer sends in a later kernel, the partial result of the landing slot or the expert ids of the next step, so one set of these slots suffices. Expert-id slots also reach ranks that send nothing back, so they alternate between two sets: a rank finishes dispatch only after the expert ids of every rank have arrived, so no rank starts a step before every rank has finished the step before last.

SENTINEL's costs grow with the batch: each rank writes, polls and resets $P\hat{T}$ expert-id slots per step, and every received word is checked, copied and reset, which shrinks the advantage of SENTINEL (§5.2).

4.3 Kernel implementation details

`NVLinkOneSided` is implemented as a set of CUDA kernels behind FlashInfer's `MoEEpCommunication` interface. Dispatch moves tokens after quantization: MoE compute follows dispatch, so the quantized tokens naturally take part in the computation. Dispatch therefore supports every quantization format of the MoE compute backends (e.g., BF16, FP8, MXFP8 and NVFP4) and sends the payloads of each token: the quantized hidden states, their scaling factors, and the expert ids and router weights from the routing. MoE compute outputs BF16, so combine usually transfers BF16. The kernels with the FENCE protocol can also quantize the partial results to FP8 before combine and dequantize them before the reduction, which saves combine bandwidth at a potential loss of precision.

Both kernels keep as many bytes in flight over NVLink as they can, through two kinds of parallelism:

- **Data-level parallelism (DLP).** Each thread uses the widest vectorized loads and stores that the payload alignment allows, 16 bytes.
- **Instruction-level parallelism (ILP).** In dispatch, a thread loads each vector from local memory once and stores it to all target ranks of the token in an unrolled loop, so the remote stores of the token are in flight together. In combine, a thread first loads the vectors of all target ranks of the token into registers and then adds them pairwise in a tree, which keeps the dependency chain short.

Table 3. Configuration of the MoE layers of the tested models, with the bytes per token and target rank.

	DeepSeek-V4.1-Flash	Kimi-K3
Hidden size	5120	3584[a]
Experts, top-k	384, 6	896, 16
Expert intermediate size	2304	3072
Dispatch bytes (MXFP8)	5328	3824
Combine bytes (BF16)	10240	7168

[a] Latent MoE; the original hidden size is 7168.

The size of the rank-major workspaces depends only on the number of tokens, so the whole MoE layer can be captured in a CUDA graph, in which dispatch, the MoE compute kernels and combine use Programmatic Dependent Launch (PDL) (NVIDIA, 2026a) to overlap each kernel with its predecessor. Every synchronization is a point where a failed peer stalls the others, so every wait has a timeout and traps with an error message instead of hanging the engine.

5 Evaluation

We measure the all-to-all alone (§5.2), the whole MoE layer with it (§5.3), and both in a larger NVLink domain (§5.4).

5.1 Setup

Sections 5.2–5.3 run on eight NVLink-connected B300 GPUs and §5.4 on a GB200 NVL72 rack to evaluate large EP. We test the routed-expert layers of DeepSeek-V4.1-Flash (DeepSeek-AI, 2026b) and Kimi-K3 (Moonshot AI, 2026) (Table 3) at 1 to 4096 tokens per rank, with *perfect* routing: the k selections of each token go to consecutive ranks and spread evenly over their experts.

Our MoE layer pairs `NVLinkOneSided` with the Prims-TS grouped GEMMs of §3.3, with MXFP4 weights and MXFP8 activations; where a single number stands for `NVLinkOneSided`, it is that of the faster protocol. The baselines are DeepEP V2.5 (Zhao et al., 2025) with the settings of its own tests; DeepGEMM MegaMoE (DeepSeek-AI, 2025), one persistent kernel that overlaps NVLink transfers with computation; and DeepEP + DeepGEMM (DE+DG), the unfused baseline of MegaMoE's test. Appendix B describes the measurement.

5.2 All-to-all performance

`NVLinkOneSided` is faster than DeepEP at every batch, most where fixed costs dominate (Figure 5): by 1.75–3.00× up to 128 tokens per rank and by 1.17–1.63× from 256. At one token per rank, the round trip on DeepSeek-V4.1-Flash takes 17.4 µs with SENTINEL, 27.6 µs with FENCE and 47.9 µs with DeepEP; at 4096 tokens, FENCE reaches 68–71% of the 900 GB/s unidirectional NVLink bandwidth.

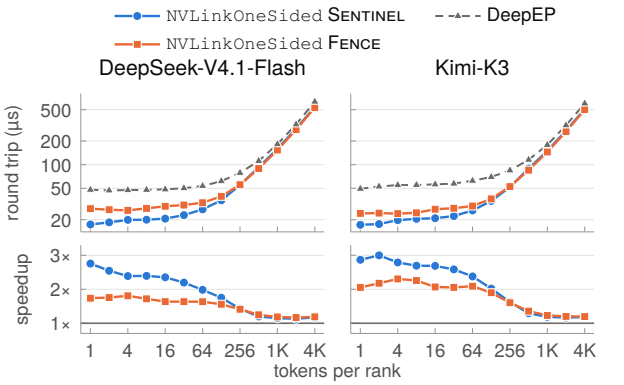


Figure 5. `NVLinkOneSided` completes dispatch and combine faster than DeepEP at every batch. Round trip on eight B300 GPUs; bottom: speedup over DeepEP.

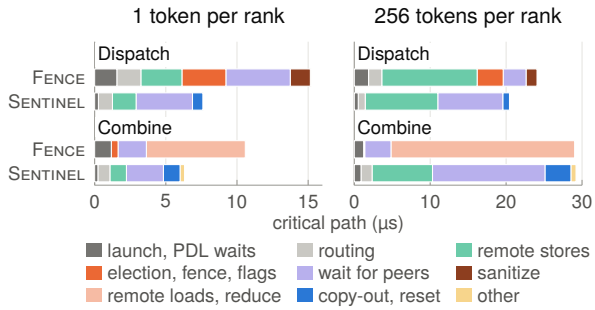


Figure 6. SENTINEL shortens the work around the wait for peers, not the wait. Critical path inside the dispatch and combine kernels, from IKET timestamps; DeepSeek-V4.1-Flash on eight B300 GPUs. Launch, PDL waits include FENCE's counter-reset kernel; election, fence, flags: FENCE's serial chain; sanitize: FENCE's kernel that invalidates the expert ids past the slot counts, mostly overlapped with dispatch; copy-out, reset: SENTINEL's handling of received slots; other: kernel exit and SENTINEL's local reduction.

FENCE or SENTINEL. SENTINEL is faster at decoding batches and FENCE at large ones, with the crossover at 256 tokens per rank (Figure 5, bottom). In-Kernel Event Tracing (IKET) (NVIDIA, 2026b) shows why (Figure 6). At one token per rank, both protocols wait about equally long for their peers, and SENTINEL saves the work around the wait: its dispatch skips FENCE's serial chain and extra kernels and takes half as long, and its combine pushes each partial result when it is ready instead of waiting for all peers and pulling with remote loads. As the batch grows, the transfers dominate with either protocol, while SENTINEL's per-slot work, copying out and resetting every received slot and running more combine blocks than fit on the GPU at once, cancels its saving.

Payload precision. Low-precision payloads pay off at large batches, where FENCE is used (Figure 7): at 4096

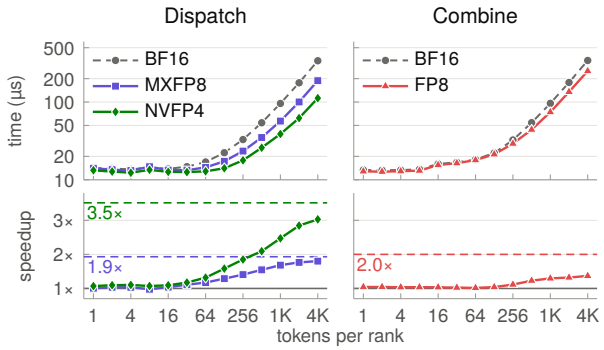


Figure 7. Low-precision formats accelerate dispatch and combine at large batches; dashed: the byte ratio, their theoretical limit. `NVLinkOneSided` with FENCE, DeepSeek-V4.1-Flash on eight B300 GPUs; top: latency of each phase; bottom: speedup over BF16.

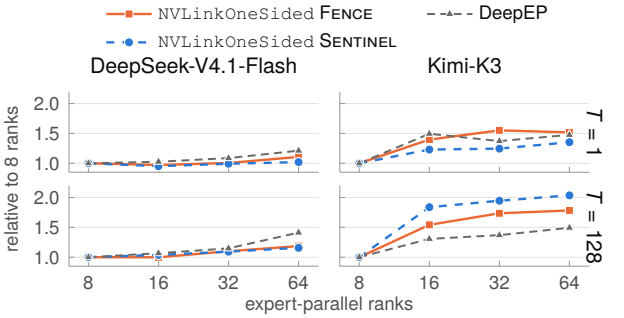


Figure 8. At one token per rank, SENTINEL's round trip grows least from 8 to 64 GPUs. All-to-all round trip on NVL72 relative to the same implementation on 8 GPUs, at a fixed number of tokens per rank.

tokens per rank, NVFP4 makes dispatch 3.0× faster than BF16 for a byte ratio of 3.5×, and MXFP8 1.8× for 1.9×, while at one token all formats take the same time. FP8 combine halves the bytes but is only 1.4× faster, as a kernel must first read and quantize each BF16 partial result; with per-element E4M3, it shows no statistically significant accuracy loss on downstream tasks.

5.3 MoE layer

Table 4 lists the layer and its stages. Although built from separate kernels, our layer is faster than the fused MegaMoE up to 128 tokens per rank, by up to 30%; from 256 tokens MegaMoE leads, by up to 35% at 4096, as it hides the communication behind the grouped GEMMs. Our layer is faster than DeepEP + DeepGEMM at every batch, by up to 2.4×, mostly through dispatch and combine.

Skewed routing. Real routing is imbalanced, and the rank that holds the popular experts limits the layer. We model this with Zipf routing (Table 5): each token draws k distinct experts with probability proportional to $i^{-\alpha}$, where i is the expert's position in a random order ($\alpha = 0$: uniform). At one token per rank, the grouped GEMMs are bound by loading expert weights, so skew, which activates fewer experts, can even make the layers faster. At 1024 tokens per rank, the busiest rank computes up to 1.9 times as many token–expert routing pairs as a rank under perfect routing (load imbalance in Table 5), and its grouped GEMMs slow all three layers by 50–70%. Our communication grows less: a token reaches a target rank once, however many of the rank's experts it selects, so the bytes a rank receives follow its tokens, at most $P/\min(k, P)$ times as many as under perfect routing, rather than its token–expert routing pairs. At $\alpha = 1.25$, dispatch and combine of the busiest rank take 160 and 135 µs, against 267 and 400 µs with DeepEP + DeepGEMM, which copies each token once per selected local expert.

5.4 Large EP scaling

A larger EP group gives each GPU fewer experts, which frees memory and, at small batches, where the grouped GEMMs are bound by loading weights, reduces the weights each GPU loads. We scale EP from 8 to 64 GPUs on one NVL72 rack at a fixed number of tokens per rank.

The all-to-all slows only slightly (Figure 8). At one token per rank on DeepSeek-V4.1-Flash, SENTINEL's round trip stays at 16.6–17.8 µs, as it has no per-peer handshake, while FENCE's grows from 43.0 to 47.6 µs, as it publishes and awaits one flag per peer. Kimi-K3 scales worse, as from 16 GPUs on its tokens ($k = 16$) reach 16 ranks instead of eight.

The MoE layer gains at intermediate batches (Table 6): on 64 GPUs, a rank holds six experts of DeepSeek-V4.1-Flash instead of 48, and at 128 tokens per rank our layer takes 103 instead of 182 µs. At one token per rank, a rank activates about k experts at any EP size, and at large batches it processes Tk token–expert routing pairs at any EP size, so large EP does not reduce its work in either case. Our layer stays 23–26% below MegaMoE at one token per rank, while MegaMoE gains more at intermediate batches and leads from 32 tokens per rank on 64 GPUs instead of from 128 on eight.

6 DISCUSSION

Advantages and limitations. `NVLinkOneSided` performs well across the board: its all-to-all is faster than DeepEP at every batch, and its layer, built from separate kernels, is faster than the fused MegaMoE up to 128 tokens per rank on eight GPUs (§5). Keeping the all-to-all separate from MoE compute makes it adaptable: it carries any

Table 4. MoE layer latency (µs) on eight B300 GPUs: layer time and dispatch / MoE compute / combine stages (MegaMoE: one kernel); fastest layer in bold.

Tokens/rank	`NVLinkOneSided` FENCE MoE layer	`NVLinkOneSided` FENCE stages	`NVLinkOneSided` SENTINEL MoE layer	`NVLinkOneSided` SENTINEL stages	DE+DG MoE layer	DE+DG stages	MegaMoE MoE layer
DeepSeek-V4.1-Flash							
1	46.0	12.8 / 19.6 / 14.4	**36.7**	8.3 / 23.9 / 9.8	74.7	21.4 / 26.0 / 27.2	49.8
16	79.6	13.3 / 49.9 / 17.2	**70.5**	8.5 / 54.7 / 12.9	122.0	22.0 / 68.7 / 31.4	90.3
128	188.5	16.3 / 150.0 / 23.1	**183.3**	13.5 / 150.6 / 21.5	235.9	25.8 / 168.7 / 41.4	188.7
256	208.1	22.2 / 153.4 / 33.6	206.5	20.5 / 154.9 / 33.8	251.2	29.8 / 170.6 / 50.8	**194.1**
1024	353.8	55.2 / 201.3 / 99.2	356.9	60.8 / 206.0 / 96.2	409.6	72.9 / 215.8 / 121.3	**255.3**
4096	1185.2	187.2 / 647.9 / 351.6	1198.6	197.4 / 665.7 / 343.3	1290.2	240.1 / 612.1 / 429.6	**766.1**
Kimi-K3							
1	47.7	13.1 / 23.5 / 15.1	**42.1**	8.8 / 28.1 / 13.5	99.7	26.5 / 37.8 / 35.3	60.4
16	128.4	13.4 / 100.9 / 15.0	**122.0**	9.4 / 106.5 / 12.5	191.2	23.5 / 129.4 / 38.2	148.6
128	351.1	17.2 / 313.4 / 21.5	**346.4**	13.7 / 314.8 / 21.0	430.8	27.0 / 346.9 / 56.8	358.5
256	396.1	22.8 / 342.2 / 32.0	393.2	20.3 / 343.4 / 32.6	482.1	31.2 / 374.0 / 76.7	**370.2**
1024	719.5	54.9 / 561.6 / 102.9	732.1	61.5 / 580.7 / 98.2	976.3	78.1 / 678.5 / 222.3	**535.8**
4096	2553.6	181.9 / 2023.0 / 360.2	2601.5	195.8 / 2084.3 / 345.3	3091.1	269.8 / 2006.4 / 817.7	**1901.9**

Table 5. MoE layer under Zipf routing with exponent α ($\alpha = 0$: uniform) on eight B300 GPUs, latency in µs. Imba. (load imbalance) is the ratio of the token–expert routing pairs that the busiest rank computes to those that each rank computes under perfect routing; 1 means balanced. Communication is the dispatch and combine of the busiest rank. Ours: `NVLinkOneSided` with the faster protocol + Prims-TS; DE+DG: DeepEP + DeepGEMM.

		DeepSeek-V4.1-Flash						Kimi-K3					
			Communication		MoE layer				Communication		MoE layer		
Tokens/rank	α	Imba.	Ours	DeepEP	Ours	DE+DG	MegaMoE	Imba.	Ours	DeepEP	Ours	DE+DG	MegaMoE
1	0	1.50	16.1	47.6	52.2	97.5	68.3	1.38	22.4	58.0	84.6	148.0	114.2
	0.5	1.50	15.8	49.4	57.8	105.2	73.0	1.25	28.2	68.5	81.0	146.1	106.8
	1	1.50	15.8	52.6	52.4	99.3	67.6	1.38	30.0	65.4	71.1	133.4	96.6
	1.25	1.67	15.9	51.6	52.3	97.9	69.2	1.62	22.1	62.2	66.4	129.4	95.4
1024	0	1.02	119.5	178.4	414.9	461.1	255.6	1.01	143.0	291.7	717.6	968.6	538.9
	0.5	1.29	125.7	199.8	445.4	499.5	303.9	1.19	131.3	306.7	852.3	1096.7	638.9
	1	1.69	146.4	242.3	557.0	626.7	374.4	1.55	134.2	373.0	1075.3	1352.6	814.0
	1.25	1.91	160.4	267.1	619.4	696.0	411.5	1.70	134.7	400.0	1184.5	1477.0	918.4

quantization format and pairs with various MoE compute backends behind one interface, and it leaves room for features of the communication alone, such as fault tolerance. Its main limitation is that the EP group must fit in one NVLink domain: both protocols rely on direct loads and stores to peer memory and on 32-bit words that arrive whole, which RDMA does not offer (§2.2). Its dispatch workspace holds one region per rank of the group, so its memory grows with the EP size.

MegaMoE. DeepGEMM MegaMoE fuses dispatch, both grouped GEMMs and combine into one persistent kernel that exchanges token indices first and then pulls the tokens into expert-major blocks sized for its GEMM tiles. Its strength is overlap: within one kernel, NVLink transfers proceed during Tensor Core work, which separate kernels cannot do, so MegaMoE leads from 256 tokens per rank (§5.3). Its weaknesses are the metadata exchange ahead of the data and its coupling: layout, transfer granularity and formats follow the GEMM's tiles, so a new grouped GEMM or precision needs a new implementation. In future work, we plan to carry the one-sided design into such a fused operator, pushing tokens without metadata into a rank-major layout and detecting their arrival from the data, so that communication also overlaps with computation at large batches.

7 Related Work

MoE systems. Expert parallelism was introduced with GShard (Lepikhin et al., 2021) and Switch Transformer (Fedus et al., 2022). DeepSpeed-MoE (Rajbhandari et al., 2022) co-designs the model with an inference system that combines expert, tensor and data parallelism with hierarchical all-to-all; FasterMoE (He et al., 2022) counters load imbalance by replicating popular experts and overlaps communication through fine-grained scheduling; Tutel (Hwang

Table 6. MoE layer latency (µs) of DeepSeek-V4.1-Flash on one GB200 NVL72 rack, by EP size; Ours: the faster protocol. Fastest layer in bold.

		EP size (GPUs)			
T		8	16	32	64
1	Ours	**35.4**	**35.4**	**35.7**	**36.9**
	DE+DG	78.1	78.0	79.1	82.3
	MegaMoE	45.7	46.3	47.2	49.7
16	Ours	**66.4**	**51.9**	**47.7**	**49.0**
	DE+DG	124.5	100.6	88.9	90.9
	MegaMoE	83.4	63.5	58.6	56.0
128	Ours	181.7	124.7	120.9	102.6
	DE+DG	237.7	169.3	136.2	126.8
	MegaMoE	**179.3**	**121.0**	**91.2**	**86.6**

et al., 2023) switches parallelism and pipelining at runtime; and MegaBlocks (Gale et al., 2023) computes experts as block-sparse operations, so that no token is dropped. MegaScale-Infer (Zhu et al., 2025) places attention and experts on separate GPUs, possibly of different types, hides the traffic between them with a ping-pong pipeline over micro-batches, and carries it with an RDMA library from the attention nodes to the expert nodes.

Expert-parallel communication. DeepEP V1 (Zhao et al., 2025) has two kernel families. Its high-throughput kernels, for training and prefill, first exchange token counts and then forward the tokens between the NVLink and RDMA domains; its low-latency kernels, for decoding, send each token over RDMA, later also over NVLink, into per-expert receive buffers without a separate count exchange, and are compatible with CUDA graphs. DeepEP V2 rebuilds both on NCCL's GPU-Initiated Networking (GIN) (Hamidouche et al., 2025) behind one buffer interface, with an expanded per-expert output layout for grouped GEMMs. Perplexity's kernels (Licker et al., 2025b;a) split sending and receiving into separate kernels and deliver per-expert outputs, through NVSHMEM across nodes and direct stores within a node. NCCL EP (Goldman et al., 2026) offers both modes on NCCL's device API; its low-latency mode can also return the tokens in a rank-major layout. UB-X (NVIDIA, 2026c), an experimental NVLink library in NCCL, computes dispatch offsets from a routing matrix that every rank holds and provides a combine that pushes with sentinel synchronization.

Collective protocols. NCCL's LL and LL128 protocols (NVIDIA, a; Hu et al., 2025) and MSCCL++'s LL packets (Shah et al., 2025) pair data with flag words so that receivers detect arrival without fences, at the cost of transferring the flags; NVSHMEM's put-with-signal instead pairs each transfer with a separate signal word (NVIDIA, b). Shen et al. (2026) measure memory barriers at over a microsecond each on GB200 and build barrier-free AllReduce from LL and sentinel synchronization on NCCL's device API. In dispatch and combine, sentinel synchronization meets two new problems: quantized payloads may take the sentinel's bit pattern, and the slots that receive tokens on a rank depend on the routing (§4.2). In-band completion over one-sided writes also appears in RDMA systems, which validate remote objects or messages by per-cache-line versions or by the arrival of their last bytes (Dragojević et al., 2014; Kalia et al., 2014), and HPC separates completion from data transfer through MPI-3 epochs (Mes, 2015) and Notified Access (Belli & Hoefler, 2015).

Overlap and fusion. Comet (Zhang et al., 2025) overlaps expert computation with communication at tile granularity, and TileLink (Zheng et al., 2025) generates overlapped kernels from tile-level primitives. FlashDMoE (Aimuyo et al., 2025) and DeepGEMM MegaMoE (DeepSeek-AI, 2025) fuse dispatch, expert computation and combine into one persistent kernel, and FarSkip-Collective (Anonymous, 2026) changes the model so that blocking collectives overlap with computation. These reduce how much of the communication is exposed; `NVLinkOneSided` reduces the cost of the communication itself, which matters most in decoding, where there is little computation to hide it behind.

8 Conclusion

`NVLinkOneSided` implements dispatch and combine as one-sided communication within an NVLink domain: each source rank pushes its tokens into rank-major workspaces using only its own routing information, MoE compute reads them through slot indices, and two protocols, Sentinel for small batches and Fence for large ones, synchronize the exchange. On Blackwell GPUs, `NVLinkOneSided` is faster than DeepEP at every batch, and the resulting MoE layer, built from separate kernels, achieves better performance than DeepGEMM MegaMoE at decoding batches. `NVLinkOneSided` is available in FlashInfer[3] and is integrated into vLLM (Kwon et al., 2023), SGLang (Zheng et al., 2024) and TensorRT-LLM (NVIDIA, 2026e).

Acknowledgements

We thank Itamar Rosenfeld Rauch, Chanan Maoz, Samuel Nordmann, Nikita Korobov, Yicheng Qiang, Ke Wen, Xiaofan Li, Jack Yang, Hao Lu and Ao Tang for their help and discussions.

[3] https://docs.flashinfer.ai/generated/flashinfer.moe_ep.NVLinkOneSidedAlltoAll.html

REFERENCES

Agrawal, A., Kedia, N., Panwar, A., Mohan, J., Kwatra, N., Gulavani, B. S., Tumanov, A., and Ramjee, R. Taming throughput-latency tradeoff in LLM inference with Sarathi-Serve. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, 2024.

Aimuyo, O. J., Oh, B., and Singh, R. FlashDMoE: Fast distributed MoE in a single kernel. *arXiv preprint arXiv:2506.04667*, 2025.

Anonymous. FarSkip-Collective: Unhobbling blocking communication in mixture of experts models. In *Proceedings of Machine Learning and Systems (MLSys)*, 2026. arXiv:2511.11505.

Aubrey, K. and Stam, N. Inside NVIDIA Blackwell Ultra: The chip powering the AI factory era. NVIDIA Technical Blog, https://developer.nvidia.com/blog/inside-nvidia-blackwell-ultra-the-chip-powering-the-ai-factory-era/, August 2025.

Belli, R. and Hoefler, T. Notified access: Extending remote memory access programming models for producer-consumer synchronization. In *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2015.

DeepSeek-AI. DeepSeek-V3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.

DeepSeek-AI. DeepGEMM. https://github.com/deepseek-ai/DeepGEMM, 2025.

DeepSeek-AI. DeepInfer: High-performance, large-scale inference framework for DeepSeek models. https://github.com/deepseek-ai/DeepInfer, 2026a. Commit 5537415, decode benchmarks in README and doc/inference_deployment_qa.md.

DeepSeek-AI. DeepSeek-V4.1-Flash. https://huggingface.co/deepseek-ai/DeepSeek-V4.1-Flash, 2026b.

Dragojević, A., Narayanan, D., Hodson, O., and Castro, M. FaRM: Fast remote memory. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 2014.

Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.

FlashInfer Contributors. FlashInfer Prims-TS fused MoE API. https://docs.flashinfer.ai/api/fused_moe.html#prims-ts-fused-moe, 2026. Accessed 2026-10-30.

Gale, T., Narayanan, D., Young, C., and Zaharia, M. MegaBlocks: Efficient sparse training with mixture-of-experts. In *Proceedings of Machine Learning and Systems (MLSys)*, 2023.

Goldman, A., Boker, N., Sheraizin, M., Admoni, N., Polyakov, A., Bhattacharya, S., Yu, F., Sun, K., Theodorakis, G., Yin, H.-C., Gootzen, P.-J., Shafi, A., Ravid, A., Di Girolamo, S., Dinan, J., Li, X., Gorentla Venkata, M., and Bloch, G. NCCL EP: Towards a unified expert parallel communication API for NCCL. arXiv:2603.13606, 2026.

Hamidouche, K., Bachan, J., Markthub, P., Gootzen, P.-J., Agostini, E., Jeaugey, S., Shafi, A., Theodorakis, G., and Gorentla Venkata, M. GPU-initiated networking for NCCL. arXiv:2511.15076, 2025.

He, J., Zhai, J., Antunes, T., Wang, H., Luo, F., Shi, S., and Li, Q. FasterMoE: Modeling and optimizing training of large-scale dynamic pre-trained models. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, 2022.

Hu, Z. et al. Demystifying NCCL: An in-depth analysis of GPU communication protocols and algorithms. *arXiv preprint arXiv:2507.04786*, 2025.

Hwang, C., Cui, W., Xiong, Y., Yang, Z., Liu, Z., Hu, H., Wang, Z., Salas, R., Jose, J., Ram, P., Chau, H., Cheng, P., Yang, F., Yang, M., and Xiong, Y. Tutel: Adaptive mixture-of-experts at scale. In *Proceedings of Machine Learning and Systems (MLSys)*, 2023.

InfiniBand Trade Association. *InfiniBand Architecture Specification, Volume 1, Release 1.4*, 2020.

Kalia, A., Kaminsky, M., and Andersen, D. G. Using RDMA efficiently for key-value services. In *ACM SIGCOMM*, 2014.

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 *ACM Symposium on Operating Systems Principles (SOSP)*, 2023.

Lamport, L. Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7): 558–565, 1978.

Lepikhin, D., Lee, H., Xu, Y., Chen, D., Firat, O., Huang, Y., Krikun, M., Shazeer, N., and Chen, Z. GShard: Scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations (ICLR)*, 2021.

Licker, N., Chen, L., and Hoang, P. Efficient and portable mixture-of-experts communication. Perplexity Research Blog, https://www.perplexity.ai/hub/blog/efficient-and-portable-mixture-of-experts-communication, 2025a.

Licker, N., Hu, K., Zaytsev, V., and Chen, L. pplx-kernels: Perplexity MoE kernels. https://github.com/perplexityai/pplx-kernels, 2025b.

Lustig, D., Sahasrabuddhe, S., and Giroux, O. A formal analysis of the NVIDIA PTX memory consistency model. In *International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2019.

Markthub, P., Dinan, J., Potluri, S., and Howell, S. Improving network performance of HPC systems using NVIDIA Magnum IO NVSHMEM and GPUDirect Async. NVIDIA Technical Blog, https://developer.nvidia.com/blog/improving-network-performance-of-hpc-systems-using-nvidia-magnum-io-nvshmem-and-gpudirect-async/, November 2022.

MPI: A Message-Passing Interface Standard, Version 3.1. Message Passing Interface Forum, 2015.

Moonshot AI. Kimi-K3. https://huggingface.co/moonshotai/Kimi-K3, 2026.

NVIDIA. NCCL: NVIDIA collective communications library. https://github.com/NVIDIA/nccl, a.

NVIDIA. NVSHMEM. https://developer.nvidia.com/nvshmem, b.

NVIDIA. Enabling fast inference and resilient training with NCCL 2.27. https://developer.nvidia.com/blog/enabling-fast-inference-and-resilient-training-with-nccl-2-27/, 2025.

NVIDIA. CUDA programming guide: Programmatic dependent launch. https://docs.nvidia.com/cuda/cuda-programming-guide/04-special-topics/programmatic-dependent-launch.html, 2026a. Accessed 2026-10-30.

NVIDIA. IKET profiling (In-Kernel event tracing). NVIDIA CUTLASS documentation 4.7.0, https://docs.nvidia.com/cutlass/4.7.0/media/docs/pythonDSL/guides/iket_profiling.html, 2026b. Accessed 2026-10-30.

NVIDIA. UB-X: Ultra bandwidth—eXperimental NVLink collectives. https://github.com/NVIDIA/nccl/tree/v2.32.3-1/contrib/nccl_ubx, 2026c. NCCL contrib; first released with NCCL 2.30.7.

NVIDIA. *Parallel Thread Execution ISA, Version 9.4*, 2026d. https://docs.nvidia.com/cuda/parallel-thread-execution/.

NVIDIA. TensorRT-LLM. https://github.com/NVIDIA/TensorRT-LLM, 2026e.

PyTorch Contributors. PyTorch symmetric memory. https://pytorch.org/docs/stable/symmetric_memory.html.

Rajbhandari, S., Li, C., Yao, Z., Zhang, M., Aminabadi, R. Y., Awan, A. A., Rasley, J., and He, Y. DeepSpeed-MoE: Advancing mixture-of-experts inference and training to power next-generation AI scale. In *International Conference on Machine Learning (ICML)*, 2022.

Recio, R., Metzler, B., Culley, P., Hilland, J., and Garcia, D. A remote direct memory access protocol specification. RFC 5040, IETF, 2007.

Shah, A. et al. MSCCL++: Rethinking GPU communication abstractions for cutting-edge AI applications. *arXiv preprint arXiv:2504.09014*, 2025.

Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., and Dean, J. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations (ICLR)*, 2017.

Shen, S., Korzh, A., Bachan, J., Chen, T., Goel, A., Schneider, L., Kousha, P., He, Z., Jeaugey, S., Iskra, K., Chandawala, N., Hammond, J. R., and Hoefler, T. Every microsecond matters: Achieving near speed-of-light latency in GPU collectives. arXiv:2607.16100, 2026.

Zhang, S. et al. Comet: Fine-grained computation-communication overlapping for mixture-of-experts. In *Proceedings of Machine Learning and Systems (MLSys)*, 2025.

Zhao, C., Zhou, S., Zhang, L., Deng, C., Xu, Z., Liu, Y., Yu, K., Li, J., and Zhao, L. DeepEP: An efficient expert-parallel communication library. https://github.com/deepseek-ai/DeepEP, 2025. V1: tag v1.2.1; V2.5: branch v2.5, 2026-09-29.

Zheng, L., Yin, L., Xie, Z., Sun, C., Huang, J., Yu, C. H., Cao, S., Kozyrakis, C., Stoica, I., Gonzalez, J. E., Barrett, C., and Sheng, Y. SGLang: Efficient execution of structured language model programs. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

Zheng, S. et al. TileLink: Generating efficient compute-communication overlapping kernels using tile-centric primitives. In *Proceedings of Machine Learning and Systems (MLSys)*, 2025.

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., Jin, X., and Liu, X. MegaScale-Infer: Efficient mixture-of-experts model serving with disaggregated expert parallelism. In *ACM SIGCOMM Conference (SIGCOMM)*, 2025.

AI USAGE DISCLOSURE

AI coding assistants were used to help draft and edit the text and to write the plotting and benchmark scripts. All technical content, the design, the experiments and the reported numbers were produced or verified by the authors. No experimental data was generated by AI.

A LATENCY TABLES

Tables A.1 and A.2 list dispatch, combine and round trip of the all-to-all (§5.2) for SENTINEL, FENCE and DeepEP from 1 to 4096 tokens per rank. Tables A.3 and A.4 list the MoE layer (§5.3) with each protocol, MegaMoE and DeepEP + DeepGEMM; Table A.5 gives the payload formats (Figure 7), and Table A.6 the layer under Zipf routing. On one GB200 NVL72 rack (§5.4), Table A.7 lists the all-to-all and Table A.8 the layer with perfect routing.

B MEASUREMENT DETAILS

The B300 node runs CUDA 13.5, and the tokens are quantized to MXFP8 before the timed span. Ten warm-up and 100 timed iterations run back to back in one CUDA graph, with L2 flushed before each; we report the span from the first to the last kernel, averaged over iterations and ranks. All-to-all outputs match a reference bitwise, and every layer matches one to a relative L2 error of at most 5.7×10^{-3}. We time phases from CUPTI kernel records in back-to-back CUDA graph replays, with no barrier or event between phases: a barrier before every phase hides the waiting a protocol does in steady state, and events in the graph keep Programmatic Dependent Launch (PDL) from overlapping the kernels around them. The round trip of the all-to-all is the sum of dispatch and combine. Each stage of the MoE layer (Table 4) spans its first kernel's start to its last kernel's end; PDL lets stages overlap slightly, so they add up to more than the layer. The breakdown in Figure 6 comes from IKET timestamps that every warp records when it enters and leaves each step, traced back from the block that finishes each phase last and averaged over ranks and iterations; instrumentation changes the round trip by at most 6%.

Table A.1. All-to-all latency (µs), DeepSeek-V4.1-Flash, 8×B300: dispatch / combine / round trip.

	SENTINEL			FENCE			DeepEP		
Tokens/rank	disp.	comb.	total	disp.	comb.	total	disp.	comb.	total
1	8.2	9.1	17.4	14.2	13.3	27.6	23.2	24.7	47.9
2	9.0	9.5	18.5	13.5	13.2	26.8	21.5	25.6	47.0
4	8.3	11.6	19.9	13.0	13.3	26.2	21.4	26.1	47.5
8	8.3	11.7	20.0	14.2	13.6	27.8	21.4	26.5	47.8
16	8.5	12.1	20.6	13.6	16.0	29.6	21.7	26.8	48.5
32	9.3	13.5	22.9	13.9	16.8	30.7	22.0	28.2	50.2
64	10.8	16.3	27.0	14.6	18.2	32.8	23.0	30.6	53.6
128	13.5	21.7	35.2	17.4	22.2	39.6	25.2	36.5	61.7
256	20.6	34.5	55.1	23.2	32.5	55.7	30.7	48.0	78.7
512	37.2	55.7	92.9	35.1	54.1	89.2	42.1	69.3	111.4
1024	60.4	98.0	158.4	57.4	96.1	153.5	68.9	113.3	182.2
2048	106.1	180.8	286.9	99.9	178.6	278.5	120.3	205.0	325.3
4096	197.0	341.7	538.7	188.2	342.4	530.6	227.5	404.1	631.6

Table A.2. All-to-all latency (µs), Kimi-K3, 8×B300: dispatch / combine / round trip.

	SENTINEL			FENCE			DeepEP		
Tokens/rank	disp.	comb.	total	disp.	comb.	total	disp.	comb.	total
1	8.6	8.6	17.2	12.7	11.3	24.0	23.2	26.2	49.3
2	8.9	8.7	17.6	13.0	11.2	24.2	23.8	29.0	52.8
4	8.7	11.1	19.7	12.4	11.5	23.9	25.1	29.9	55.0
8	9.0	11.5	20.5	13.2	11.3	24.4	24.4	30.8	55.2
16	9.2	11.7	20.9	13.1	14.0	27.2	25.0	31.2	56.2
32	9.4	12.8	22.2	13.2	14.7	27.9	25.1	32.3	57.4
64	10.7	15.6	26.2	13.1	16.7	29.8	26.5	35.9	62.4
128	13.8	20.7	34.5	16.1	20.7	36.7	29.1	40.7	69.7
256	19.9	32.1	52.0	21.8	30.9	52.7	33.6	51.0	84.6
512	38.7	50.6	89.2	33.6	51.6	85.1	43.9	71.6	115.5
1024	60.7	89.1	149.8	54.8	90.2	145.0	66.9	111.4	178.3
2048	105.7	165.2	270.8	96.0	167.2	263.3	117.4	199.1	316.4
4096	191.5	316.1	507.6	180.1	321.1	501.1	223.9	376.3	600.3

Table A.3. MoE layer latency (µs), DeepSeek-V4.1-Flash, 8×B300, MXFP4 weights and MXFP8 activations: layer time and dispatch / MoE compute / combine, each stage from its first kernel's start to its last kernel's end.

	`NVLinkOneSided` FENCE		`NVLinkOneSided` SENTINEL		MegaMoE	DE+DG	
Tokens/rank	MoE layer	stages	MoE layer	stages	MoE layer	MoE layer	stages
1	46.0	12.8 / 19.6 / 14.4	36.7	8.3 / 23.9 / 9.8	49.8	74.7	21.4 / 26.0 / 27.2
2	47.9	12.8 / 21.2 / 14.6	39.2	8.9 / 26.6 / 9.8	52.4	80.5	21.8 / 28.4 / 30.0
4	51.2	12.3 / 25.5 / 14.2	44.6	8.1 / 29.7 / 12.1	57.8	85.8	21.6 / 34.4 / 29.8
8	59.5	13.1 / 32.4 / 14.7	52.1	8.5 / 37.2 / 11.9	68.1	98.4	21.9 / 46.1 / 30.4
16	79.6	13.3 / 49.9 / 17.2	70.5	8.5 / 54.7 / 12.9	90.3	122.0	22.0 / 68.7 / 31.4
32	113.3	13.5 / 83.2 / 17.8	105.8	9.1 / 88.7 / 14.2	123.0	160.2	22.7 / 104.4 / 33.2
64	177.9	13.5 / 146.6 / 18.9	171.8	10.4 / 147.1 / 16.2	184.7	227.4	23.6 / 168.1 / 35.7
128	188.5	16.3 / 150.0 / 23.1	183.3	13.5 / 150.6 / 21.5	188.7	235.9	25.8 / 168.7 / 41.4
256	208.1	22.2 / 153.4 / 33.6	206.5	20.5 / 154.9 / 33.8	194.1	251.2	29.8 / 170.6 / 50.8
512	259.3	33.2 / 170.5 / 57.0	260.5	37.5 / 171.5 / 53.9	214.9	300.3	44.7 / 181.0 / 74.4
1024	353.8	55.2 / 201.3 / 99.2	356.9	60.8 / 206.0 / 96.2	255.3	409.6	72.9 / 215.8 / 121.3
2048	637.9	99.6 / 355.5 / 184.5	645.1	106.2 / 363.3 / 179.8	428.3	672.3	127.8 / 320.2 / 224.3
4096	1185.2	187.2 / 647.9 / 351.6	1198.6	197.4 / 665.7 / 343.3	766.1	1290.2	240.1 / 612.1 / 429.6

Table A.4. MoE layer latency (µs), Kimi-K3, 8×B300, MXFP4 weights and MXFP8 activations: layer time and dispatch / MoE compute / combine, each stage from its first kernel's start to its last kernel's end.

	`NVLinkOneSided` FENCE		`NVLinkOneSided` SENTINEL		MegaMoE	DE+DG	
Tokens/rank	MoE layer	stages	MoE layer	stages	MoE layer	MoE layer	stages
1	47.7	13.1 / 23.5 / 15.1	42.1	8.8 / 28.1 / 13.5	60.4	99.7	26.5 / 37.8 / 35.3
2	52.7	13.0 / 27.8 / 16.0	47.2	8.9 / 32.6 / 14.9	67.5	109.2	26.0 / 45.5 / 37.4
4	61.8	13.1 / 37.7 / 15.5	57.4	8.9 / 42.6 / 15.2	78.6	120.6	24.2 / 58.3 / 38.2
8	84.7	14.1 / 59.4 / 15.5	80.7	9.2 / 65.2 / 16.0	106.7	146.2	23.4 / 85.0 / 37.8
16	128.4	13.4 / 100.9 / 15.0	122.0	9.4 / 106.5 / 12.5	148.6	191.2	23.5 / 129.4 / 38.2
32	211.9	13.5 / 183.5 / 15.9	206.4	9.7 / 189.4 / 13.9	231.2	281.8	24.4 / 216.4 / 41.0
64	334.5	14.8 / 303.2 / 17.6	328.9	10.7 / 309.4 / 16.3	350.9	416.0	25.5 / 344.2 / 46.3
128	351.1	17.2 / 313.4 / 21.5	346.4	13.7 / 314.8 / 21.0	358.5	430.8	27.0 / 346.9 / 56.8
256	396.1	22.8 / 342.2 / 32.0	393.2	20.3 / 343.4 / 32.6	370.2	482.1	31.2 / 374.0 / 76.7
512	515.6	33.8 / 419.3 / 58.8	515.3	38.9 / 423.8 / 54.1	418.4	620.6	46.2 / 451.5 / 121.5
1024	719.5	54.9 / 561.6 / 102.9	732.1	61.5 / 580.7 / 98.2	535.8	976.3	78.1 / 678.5 / 222.3
2048	1342.9	97.7 / 1060.3 / 189.7	1367.5	109.6 / 1093.9 / 179.3	1007.5	1789.6	138.8 / 1232.6 / 422.6
4096	2553.6	181.9 / 2023.0 / 360.2	2601.5	195.8 / 2084.3 / 345.3	1901.9	3091.1	269.8 / 2006.4 / 817.7

Table A.5. FENCE with other payload formats on eight B300 GPUs: bytes per token and target rank, and round trip (µs). The last two rows of each model include quantization to MXFP8, on the source rank or, after a BF16 dispatch, on the target rank.

Dispatch	Combine	Bytes	Tokens per rank			
		disp. / comb.	1	128	1024	4096
DeepSeek-V4.1-Flash						
BF16	BF16	10288 / 10240	29.0	46.3	194.3	684.5
MXFP8	BF16	5328 / 10240	28.7	40.7	154.9	532.7
FP8 (1×128)	BF16	5328 / 10240	28.6	41.1	155.3	532.7
NVFP4	BF16	2928 / 10240	28.1	37.6	137.4	456.0
MXFP8	FP8	5328 / 5120	28.0	39.9	131.4	438.9
MXFP8, at the source		5328 / 10240	30.4	43.4	158.9	547.7
BF16, at the target		10288 / 10240	30.6	48.8	214.2	768.1
Kimi-K3						
BF16	BF16	7296 / 7168	25.6	42.7	183.5	644.0
MXFP8	BF16	3824 / 7168	25.0	37.7	146.5	502.0
FP8 (1×128)	BF16	3824 / 7168	25.1	37.6	146.6	502.2
NVFP4	BF16	2144 / 7168	24.8	35.2	128.8	430.1
MXFP8	FP8	3824 / 3584	24.3	36.5	124.2	419.9
MXFP8, at the source		3824 / 7168	26.4	40.3	149.7	514.5
BF16, at the target		7296 / 7168	25.2	42.7	196.8	701.8

Table A.6. MoE layer under Zipf routing with exponent α on eight B300 GPUs: layer time (µs) of Ours, DE+DG and MegaMoE. Ratio is ours relative to MegaMoE, and Imba. is the load imbalance of Table 5. Ours: `NVLinkOneSided` with the faster protocol + Prims-TS (†: FENCE); DE+DG: DeepEP + DeepGEMM.

		DeepSeek-V4.1-Flash					Kimi-K3				
Routing	T	Ours	DE+DG	MegaMoE	Ratio	Imba.	Ours	DE+DG	MegaMoE	Ratio	Imba.
Perfect	1	36.7	74.7	49.8	0.74	1.00	42.1	99.7	60.4	0.70	1.00
	16	70.5	122.0	90.3	0.78	1.00	122.0	191.2	148.6	0.82	1.00
	128	183.3	235.9	188.7	0.97	1.00	346.4	430.8	358.5	0.97	1.00
	1024	353.8†	409.6	255.3	1.39	1.00	719.5†	976.3	535.8	1.34	1.00
Zipf, $\alpha = 0$	1	52.2	97.5	68.3	0.76	1.50	84.6	148.0	114.2	0.74	1.38
	16	154.1	208.8	167.3	0.92	1.09	302.9	387.8	328.7	0.92	1.23
	128	188.8†	233.7	188.8	1.00	1.03	347.5	432.3	359.0	0.97	1.03
	1024	414.9	461.1	255.6	1.62	1.02	717.6†	968.6	538.9	1.33	1.01
Zipf, $\alpha = 0.5$	1	57.8	105.2	73.0	0.79	1.50	81.0	146.1	106.8	0.76	1.25
	16	154.6	209.0	165.1	0.94	1.31	295.5	376.3	320.2	0.92	1.36
	128	204.2	250.2	196.1	1.04	1.25	386.7	456.4	384.9	1.00	1.20
	1024	445.4†	499.5	303.9	1.47	1.29	852.3†	1096.7	638.9	1.33	1.19
Zipf, $\alpha = 1$	1	52.4	99.3	67.6	0.77	1.50	71.1	133.4	96.6	0.74	1.38
	16	142.2	179.5	140.5	1.01	1.62	257.5	322.8	273.7	0.94	1.59
	128	233.8	279.2	221.1	1.06	1.67	445.8	539.1	449.1	0.99	1.56
	1024	557.0	626.7	374.4	1.49	1.69	1075.3†	1352.6	814.0	1.32	1.55
Zipf, $\alpha = 1.25$	1	52.3	97.9	69.2	0.76	1.67	66.4	129.4	95.4	0.70	1.62
	16	132.1	163.2	129.7	1.02	1.93	218.0	274.0	233.9	0.93	1.70
	128	240.0	289.7	225.8	1.06	1.91	447.1	550.0	453.5	0.99	1.71
	1024	619.4	696.0	411.5	1.51	1.91	1184.5†	1477.0	918.4	1.29	1.70

Table A.7. All-to-all round trip (µs, dispatch + combine) on one GB200 NVL72 rack with perfect routing, by expert-parallel ranks (on ranks / 4 compute trays) and tokens per rank.

		DeepSeek-V4.1-Flash			Kimi-K3		
Ranks	T	FENCE	SENTINEL	DeepEP	FENCE	SENTINEL	DeepEP
8	1	43.0	17.5	44.7	34.6	17.5	53.6
	16	43.1	17.8	44.2	36.0	18.2	57.6
	128	48.9	31.7	57.8	44.1	31.5	70.9
	1024	160.1	151.6	173.9	150.6	144.0	179.0
16	1	41.8	16.6	45.9	48.4	21.5	80.0
	16	41.8	18.4	47.1	50.6	23.2	63.5
	128	48.8	33.0	61.7	68.1	57.9	92.6
	1024	171.5	160.2	186.2	295.0	308.3	335.9
32	1	43.2	17.3	48.6	53.7	21.8	73.5
	16	44.3	18.2	49.5	57.9	25.4	65.1
	128	53.8	34.7	66.4	76.6	61.3	97.1
	1024	183.3	169.7	203.1	310.0	320.9	341.6
64	1	47.6	17.8	54.2	52.5	23.6	79.1
	16	46.7	19.1	55.7	56.0	26.1	90.0
	128	58.0	36.7	81.6	78.7	64.2	105.9
	1024	191.4	179.8	213.7	320.8	340.6	378.9

Table A.8. MoE layer latency (µs) on one GB200 NVL72 rack with perfect routing, by expert-parallel ranks and tokens per rank: `NVLinkOneSided` (the faster protocol; †: FENCE), DeepEP + DeepGEMM and MegaMoE. –: not tested, as the Prims-TS grouped GEMMs need at least k experts per rank (Kimi-K3 has 14 on 64 GPUs).

		DeepSeek-V4.1-Flash			Kimi-K3		
Ranks	T	Ours	DE+DG	MegaMoE	Ours	DE+DG	MegaMoE
8	1	35.4	78.1	45.7	40.7	116.5	57.3
	16	66.4	124.5	83.4	117.5	200.6	143.8
	128	181.7	237.7	179.3	341.5	429.9	344.7
	256	208.4	256.9	188.1	491.5	517.4	373.3
	512	348.2	317.2	222.0	624.3	728.2	473.1
	1024	438.1	481.8	294.1	881.4	1152.7	658.0
	4096	1378.8	1457.3	902.4	3123.6†	3537.7	2283.5
16	1	35.4	78.0	46.3	43.6	109.5	54.4
	16	51.9	100.6	63.5	84.7	151.3	105.4
	128	124.7	169.3	121.0	286.0	300.8	214.3
	256	200.3	197.2	140.1	395.6	419.1	250.3
	512	252.7	272.2	172.1	560.1†	652.5	330.3
	1024	440.0	431.4	276.9	1013.3†	1113.8	595.3
	4096	1446.6†	1432.2	818.1	3372.0†	3520.2	1885.1
32	1	35.7	79.1	47.2	43.7	107.8	53.1
	16	47.7	88.9	58.6	66.8	142.9	86.7
	128	120.9	136.2	91.2	234.0	258.2	167.0
	256	153.0	171.0	114.2	322.0	396.8	207.7
	512	265.3	264.8	173.4	570.9†	637.0	340.1
	1024	463.8	445.6	265.0	996.3†	1013.9	580.5
	4096	1560.7†	1481.1	824.5	3604.2†	3425.8	2018.7
64	1	36.9	82.3	49.7	–	110.3	56.7
	16	49.0	90.9	56.0	–	122.8	74.6
	128	102.6	126.8	86.6	–	242.1	143.2
	256	172.8	165.3	112.6	–	347.6	196.4
	512	309.0†	253.3	164.1	–	567.4	302.3
	1024	487.0	422.5	250.9	–	937.8	524.1
	4096	1759.0†	1473.1	789.9	–	3345.2	1799.4