HyScale: Scaling Long-Context Training for Learned Sparse Attention

Yangshen Deng[1], **Xinjing Huang**[2], **Yulong Wang**[2], **Yunzhuo Liu**[2], **Jun Hu**[2], **Shuyong Tan**[2], **Chun-Chih Tseng**[2], **Mingji Han**[2], **Ray Ying**[2], **Key Zhang**[2], **Luo Mai**[1]

[1]University of Edinburgh [2]Tencent Hunyuan

✉ yangshen.deng@ed.ac.uk

arXiv:2610.32057v3 [cs.AI] 4 Nov 2026

Abstract

We present HyScale, a long-context training system for learned sparse attention. Learned sparse attention is a promising architecture for million-token contexts, widely adopted in production since DeepSeek-V3.2. Selection, however, grows into the training bottleneck. For each query, it scores every previous token whereas sparse attention reads only k; worse, current systems materialize a full score table and repeat selection during recomputation. Our key insight is that the output of selection is not an activation, but rather a set of indices. Building on it, HyScale introduces *FlashTopK* and *selection replay*. *FlashTopK* finds each query's exact top-k on the fly without a score table, while *selection replay* reuses forward-pass selection in recomputation. Together with holistic system optimizations tackling end-to-end challenges, HyScale speeds up GLM-5.1 training on 1M-token sequences by 2.07× over Megatron-LM on 512 H20 GPUs. It also speeds up the CSA layers of DeepSeek-V4-Flash and the QSA layers of Qwen3.8-Flash-Next by 2.01× and 1.57× at 1M tokens on 32 H800 GPUs.

1 Introduction

Long context has become a core capability of large language models, driving training sequence lengths from 128K to 1M tokens across long documents, code repositories, and agent trajectories. At these lengths, full attention dominates training cost: its computation grows quadratically with context length because each query attends to every previous token. Learned sparse attention (Yuan et al., 2025; DeepSeek-AI et al., 2025) reduces this cost by splitting attention into two steps. In *selection*, a lightweight indexer scores every previous token for each query and keeps the top-k. In *sparse attention*, each query attends only to the selected tokens. Token-level designs, pioneered by DeepSeek Sparse Attention (DSA) (DeepSeek-AI et al., 2025), are widely adopted in production models such as GLM-5.1 (GLM-5-Team et al., 2026) and Hy4 preview (Tencent-Hunyuan, 2026). Block-level variants select contiguous token blocks instead, including CSA in DeepSeek-V4 (DeepSeek-AI et al., 2026), QSA in Qwen3.8-Flash-Next (Qiu et al., 2026), and MSA in MiniMax-M3 (Lai et al., 2026).

Selection, however, grows into the new training bottleneck. While sparse attention reads only k tokens per query, selection still scores every previous token. Worse, current systems amplify this cost in two ways. First, they *materialize* a full score table across all previous tokens, even though top-k keeps only k entries per query. Second, they *repeat* selection during activation recomputation (Chen et al., 2016; Korthikanti et al., 2023) by rerunning the entire layer from scratch. In Megatron-LM (Shoeybi et al., 2019), selection consequently consumes 47% of a GLM-5.1 training step at 1M tokens, excluding the pipeline wait—6.6× the time of sparse attention (§2.3).

Both inefficiencies stem from treating the output of selection as a standard activation. Our key insight is that **the output of selection is not a continuous activation, but rather a discrete set of indices**. Unlike standard activations, top-k selection is non-differentiable: the backward pass requires neither the intermediate score table nor a rerun of top-k search. Specifically, sparse attention uses only the selected indices, and the indexer loss differentiates through only the k selected scores. Consequently, training can find these indices on the fly without materializing the full score table, and directly reuse them during recomputation instead of repeating selection.

Realizing this insight in end-to-end training raises three challenges. First, on-the-fly top-k selection must decide whether to retain or discard each score immediately upon computation. Existing sampling-based top-k operators (Yin et al., 2026; Meng & Fang, 2026) estimate a score threshold for inference, but break down in training: their neighbor sampling is incompatible with context parallelism (CP), their sampling and key pruning both fail in early training, and selection becomes non-deterministic. Second, existing selective recomputation policies (Kor-

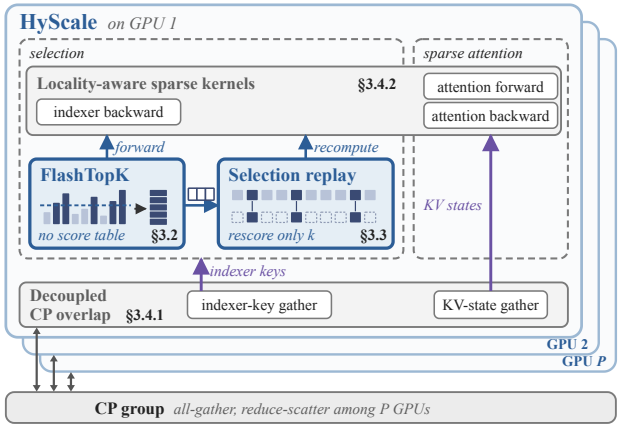


Figure 1: HyScale architecture.

thikanti et al., 2023; PyTorch, 2026) operate at the granularity of whole operators, whereas the indexer loss requires only the k selected scores. Consequently, recomputation must either retain the full score table or rescore all previous tokens merely to recover k values per query. Third, once selection is accelerated, CP communication and inefficient sparse kernels become exposed. Existing systems gather indexer keys and KV states synchronously before selection—stalling selection on KV transfers it never reads—while their sparse kernels assign each thread block a single query, repeating the loads and updates of entries adjacent queries share.

To address these challenges, we present HyScale (Figure 1), a long-context training system for learned sparse attention. We make three contributions:

(1) FlashTopK, an on-the-fly top-k selection operator via optimistic thresholding. FlashTopK selects top-k entries while scoring, eliminating the score table entirely. The key insight is that threshold-based filtering can tolerate optimistic thresholds as long as failures are rare and cheap to recover from. Built with optimistic strided sampling and adaptive key pruning, FlashTopK delivers fast, exact, and deterministic top-k selection with full CP compatibility.

(2) Selection replay, a fine-grained recomputation policy for selection. Rather than operating at the operator level, selection replay records forward-pass selection results and replays only the k selected entries during recomputation. This fine-grained replay eliminates repeated selection with negligible memory overhead, reducing recomputation cost by orders of magnitude over existing selective policies.

(3) HyScale, a holistic training system for learned sparse attention. Built on Megatron-LM, HyScale integrates FlashTopK and selection replay with end-to-end system optimizations: decoupled CP overlap hides the CP collectives of sparse attention behind selection, while locality-aware sparse kernels process two adjacent queries together so that they load and update their shared entries once. HyScale supports diverse learned sparse attention architectures, including DSA, CSA, and QSA.

Across 512 H20 GPUs, HyScale accelerates end-to-end GLM-5.1 training on 1M-token sequences by 2.07× over Megatron-LM, with speedups scaling further as context length grows. Beyond token-level DSA, HyScale generalizes seamlessly to block-level architectures: at 1M tokens on 32 H800 GPUs, it accelerates the CSA layers of DeepSeek-V4-Flash by 2.01× and the QSA layers of Qwen3.8-Flash-Next by 1.57×. HyScale is deployed in production, powering large-scale LLM training at million-token contexts.

2 Background and Motivation

2.1 Learned Sparse Attention

Learned sparse attention lets each query attend to only a few entries. An entry is an earlier token or a block of earlier tokens. Each layer runs in two steps. *Selection* uses a learned indexer to score every earlier entry of a query and keeps the k highest. *Sparse attention* then reads only the KV states of the selected entries (Appendix A). Selection thus costs up to $O(L)$ in computation per query for L tokens, while sparse attention costs only $O(k)$.

In training, gradients likewise reach only the selected entries. Top-k is not differentiable. The indexer therefore learns from a loss of its own, defined on its k selected scores. This loss pulls these scores toward the attention distribution over the selected entries. Sparse attention also back-propagates only to the selected entries (Appendix B).

Designs differ mainly in what an entry is and how it is scored. DSA (DeepSeek-AI et al., 2025) scores single tokens with a multi-head lightning indexer. All indexer heads share one indexer key per token. Its sparse attention uses multi-head latent attention (MLA) (DeepSeek-AI et al., 2024). All attention heads likewise share one compressed KV latent per token. Block-level designs score blocks of tokens instead. CSA in DeepSeek-V4 (DeepSeek-AI et al., 2026) and QSA in Qwen3.8-Flash-Next (Qiu et al., 2026) compress or mean-pool every four tokens into one entry. MSA in MiniMax-M3 (Lai et al., 2026) takes the maximum token score within each block of 128 tokens.

2.2 Long-Context Training

Context parallelism. To train on sequences that are too long for one GPU, context parallelism (CP) splits a sequence of length L across a group of P GPUs, so each GPU holds only $n = L/P$ queries. Ring Attention (Liu et al., 2024) and DeepSpeed-Ulysses (Jacobs

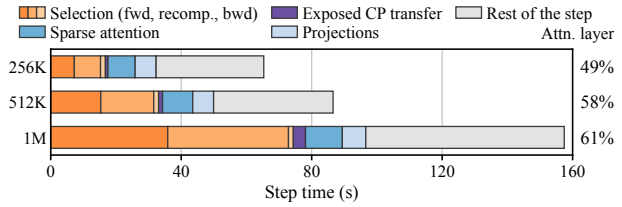


Figure 2: One training step of GLM-5.1 in Megatron-LM on 512 H20 GPUs with full recomputation and 8K tokens per GPU, first pipeline stage. Step times exclude the pipeline wait. The attention layer is in color.

et al., 2023), the widely-used CP methods designed for full attention, do not fit learned sparse attention (§3.4.1). Existing systems instead reuse the KV all-gather method of Llama 3 (Grattafiori et al., 2024). Before selection, the forward pass synchronously all-gathers the indexer keys and the KV states over the CP group. The backward pass reduce-scatters their gradients.

Activation recomputation. The backward pass needs the activations of the forward pass. At long context that storage can exceed GPU memory, so training often uses activation recomputation (Chen et al., 2016). It discards some activations during the forward pass and recomputes them when the backward pass needs them, trading compute for memory. Under full activation recomputation, each layer keeps only its input and reruns its forward pass right before the backward pass. Selective recomputation (Korthikanti et al., 2023; PyTorch, 2026) keeps more activations and reruns only parts of the layer, such as the core attention. It works at the granularity of operators, either keeping an operator's output or rerunning the operator in full.

2.3 The Selection Bottleneck

Although sparse attention is cheap, the attention layer remains the bottleneck of long-context training, and selection accounts for most of it. Figure 2 profiles GLM-5.1 training in Megatron-LM, a state-of-the-art training system, on 512 H20 GPUs with 8K tokens per GPU. Excluding the pipeline wait, the attention layer takes 49% of the step at 256K tokens and 61% at 1M. At 1M tokens, selection alone takes 47% of the step, 6.6× as long as sparse attention.

Figure 3a shows how Megatron-LM executes one learned sparse attention layer in a training step and marks two inefficiencies in selection (① and ②).

Inefficiency 1: Materialized scores. Selection materializes every score before selecting any (①). Current systems run it as two kernels. The scoring kernel scores the n local queries against all L keys and writes the scores to GPU memory as an $n \times L$ table. The top-k kernel then reads the table back and keeps k entries per query. In the 1M-token step (n=8K, k=2048), the table of one layer takes 32 GiB on each GPU, whereas the selected indices take only 64 MiB. Since top-k keeps only 0.2% of the scores, materializing them wastes GPU memory and memory traffic, a cost that grows with the context.

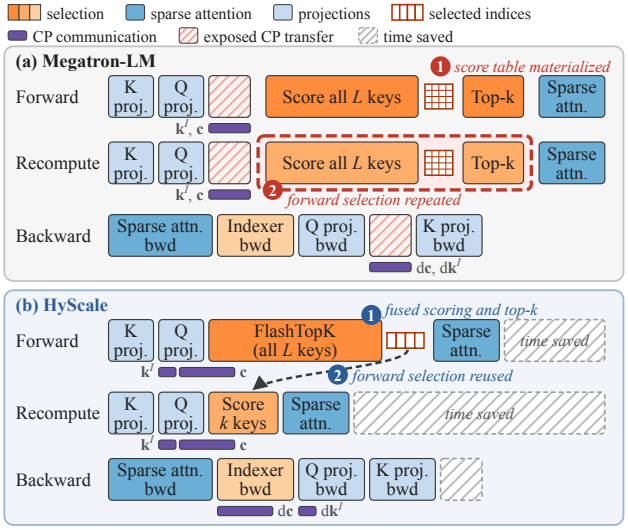


Figure 3: One learned sparse attention layer on one GPU under CP and full recomputation (not to scale). Red numbers in (a) mark the score table (①) and the repeated selection (②). Blue numbers in (b) mark where FlashTopK (①) and selection replay (②) remove them. HyScale also overlaps the CP collectives with computation. $\mathbf{k}^I$ denotes the indexer keys, and $\mathbf{c}$ the KV states.

Inefficiency 2: Repeated selection. Under full activation recomputation, selection runs twice in every training step (②). Before each layer's backward pass, the rerun of its forward pass scores all L keys and selects the top-k again. As Figure 2 shows, this rerun takes about as long as selection in the forward pass. At 1M tokens, it thus accounts for half of the 47% of the step that selection takes, and like the forward selection, its cost grows with the context.

2.4 Key Insight

Both inefficiencies stem from the same root cause. Current systems handle the output of selection, the selected indices, like a regular activation. They compute it with two dense operators, scoring and top-k, so the score table becomes an intermediate tensor in GPU memory (①). Recomputation regenerates it by rerunning both operators (②).

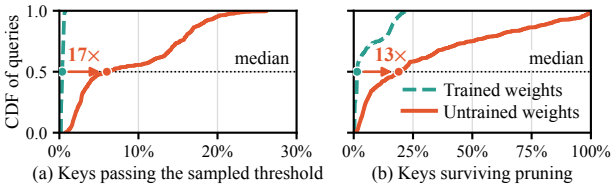
Our key insight is that **the output of selection is not a continuous activation, but rather a discrete set of indices**. Two properties of learned sparse attention support this view. First, the score table serves only to find these indices. Sparse attention reads only the selected entries, and

the indexer loss reads only their scores, so nothing reads the scores of the other keys. Second, the backward pass needs only these indices and their scores. Top-k is not differentiable, so no gradient flows through selection, and the backward pass never needs the score table or any other intermediate result of selection (Appendix B).

This insight motivates our design in two ways. First, selection need not materialize the score table. Since only the scores of the selected entries are used, the scoring kernel can keep or drop each score as soon as it computes it (FlashTopK, §3.2). Second, recomputation need not repeat selection. Since the backward pass needs only the selected indices and their scores, recomputation can reuse the indices from the forward pass and rescore only the selected entries instead of all entries (selection replay, §3.3).

Table 1: Design of fused sampling-based top-k selection.

	LiteTopK	DoSA	**FlashTopK**
Sample from	Neighbors	Neighbors	**Own context**
Threshold	Guaranteed	Guaranteed	**Optimistic**
Works under CP	✗	✗	✓
Key pruning	✗	Always	**Adaptive**



(a) Keys passing the sampled threshold (b) Keys surviving pruning

Figure 4: Sampling and pruning efficiency of DoSA (GLM-5.1, 1M). (b) uses the exact k-th score as threshold.

3 HyScale Design

3.1 Overview

HyScale builds on Megatron-LM and changes only its learned sparse attention layers. Every GPU in the CP group runs one copy of HyScale (Figure 1). Two techniques apply the insight of §2.4 to remove both inefficiencies (Figure 3b). FlashTopK (§3.2) selects each query's top-k while scoring, without writing the score table (①). Selection replay (§3.3) reuses the forward indices and rescores only their k entries in the recomputation (②). Accelerating selection exposes two remaining costs, CP communication and the sparse kernels. Decoupled CP overlap (§3.4.1) and locality-aware sparse kernels (§3.4.2) reduce them respectively.

3.2 FlashTopK

Fused sampling-based top-k selection. Current systems run unfused selection with two kernels. The scoring kernel writes all $n \times L$ scores to GPU memory, and a top-k kernel reads them back to keep k per query (①). Fused selection merges the two with a sampled threshold. It first estimates the threshold from a sample of keys. The scoring kernel then compares every score with it on chip and writes only the passing keys to a small buffer per query. If every key of the true top-k passes, selecting from these candidates is exact. The score table thus never reaches GPU memory.

The performance of this design depends on the quality of the sample. A poor sample would give a threshold far below the true k-th score, which lets many keys through, adds memory traffic and selection work, and can overflow the buffer. LiteTopK (Yin et al., 2026) pioneered this design for inference. Its threshold is the k-th largest score among the keys that nearby queries selected, as nearby queries tend to select similar keys. This guaranteed lower bound never exceeds the true k-th score and lets every key of the true top-k pass. If too many keys still pass, it raises the threshold during the scan so that the candidates fit a fixed-size buffer. DoSA (Meng & Fang, 2026) further speeds up this design with key pruning: it first computes a cheap upper bound of each key's score, and skips the full scoring of keys whose bound falls below the threshold. Both designs, however, build their sampling and pruning around inference prefill and break down in three ways during training (Table 1).

Challenges in training. *C1: Neighbor sampling breaks under CP.* LiteTopK and DoSA use *neighbor sampling*: they sample the keys that nearby queries selected. This sample covers most of the true top-k, keeping their guaranteed lower bound tight. Inference has these selections at hand, as it processes a long prompt chunk by chunk on one GPU (Agrawal et al., 2024). Under CP, however, other GPUs score nearby queries at the same time (§2.2). Their selections are thus not yet available. Without them, a guaranteed lower bound must come from some other sample of the query's context. Such a bound can be so loose that far too many keys pass, slowing down selection.

C2: Neighbor sampling and key pruning fail in early training. They both rely on a trained indexer, but training starts from untrained weights. First, with random weights, nearby queries no longer select similar keys. Neighbor sampling thus sets the threshold far too low: 17× more keys pass at the median than with trained weights (Figure 4a). Second, key pruning skips far fewer keys: even with the exact k-th score as threshold, 13× more keys survive at the median (Figure 4b). In our practice, as a result, DoSA can run even slower than unfused selection with untrained weights.

C3: Training requires deterministic selection. Training needs bitwise reproducible kernels to replay a run and debug loss spikes or hardware faults (DeepSeek-AI et al., 2026; Sul et al., 2026). Existing sampling-based operators break this, as they append the passing keys to each query's buffer with atomics. The key order thus varies across runs and changes the output in two ways. First, when several

keys tie at the k-th score, the order decides which are kept. Second, it decides the floating-point rounding of the sum in sparse attention. Sorting each query's candidates by position fixes both, but costs even more than selecting the top-k itself.

FlashTopK. To bridge the gap between sampling-based top-k selection and training, we propose FlashTopK (Table 1). Its key insight is that *the threshold need not be a guaranteed lower bound, as long as it rarely fails and each failure is cheap to recover from*. The threshold then only has to be close to the k-th score, yet selection stays exact. FlashTopK therefore estimates the threshold from each query's own context (C1, C2). It also prunes keys only when a sample predicts a gain (C2), and makes its selection deterministic (C3). Appendix C details its kernels.

Optimistic strided sampling (C1, C2). FlashTopK gives up the guarantee and estimates each query's threshold from a strided sample of its own context. It thus needs no other query's selection and works under CP. First, FlashTopK samples once for all queries on a GPU: it takes the middle key of each of $M = \min(\max(2048, L/128), 8192)$ equal segments of the L keys. Second, FlashTopK scores this sample for all queries in one matrix multiplication. Query t uses the m_t samples within its context of L_t keys, each standing for about L_t/m_t keys. Third, FlashTopK sets the threshold to the $\lceil 2.5km_t/L_t \rceil$-th best of them, an estimate of the $2.5k$-th best score. About $2.5k$ candidates thus pass, $2.5\times$ the k keys to select. The sample spans the whole context, so the threshold rarely rises above the k-th best score. For the few queries that still end with fewer than k candidates or overflow the buffer, the GPU redoes their selection over all keys without host synchronization.

Adaptive key pruning (C2). FlashTopK uses DoSA's score bound but prunes adaptively, unlike DoSA's always-on pruning. In early training, the bound loosens and pruning may not pay off (C2). FlashTopK thus decides before each selection whether to prune. It first estimates the share of keys the bound keeps, on one query in 16 against 256 sampled keys. If this share is at most 20%, a pruning pass marks keys that any of 8 adjacent queries could select, and the scan scores only these keys. Otherwise, it scans every key.

Deterministic selection (C3). FlashTopK keeps the fast atomic writes in its kernels and restores determinism afterwards. It orders keys only within the smallest sets whose order affects the output. First, when more keys share the k-th highest score than there are places left in the selection, the select kernel runs a radix select (Alabi et al., 2012) over the positions of just these tied keys and keeps those with the smallest positions. The selected set then depends only on the scores. Second, a radix sort (Merrill & Grimshaw, 2011) orders the k selected entries of each query by score, breaking ties by position, so sparse attention sums over them in the same order in every run. Two runs on the same inputs thus produce the same exact top-k, bit for bit.

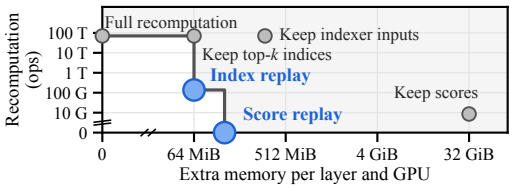


Figure 5: Costs of selection recomputation policies for one GLM-5.1 layer at 1M tokens on one GPU. Appendix D derives them.

3.3 Selection Replay

Under activation recomputation, current systems rerun selection before the backward pass (② in Figure 3a). They score all L keys of every query and select the top-k again, although the layer input has not changed and the selection is the same, which doubles the cost of selection.

Our insight is that this rerun is unnecessary, because the output of selection is only indices. Top-k is not differentiable, so gradients reach the k selected entries but never pass through the selection itself. The backward pass thus needs only the selected indices and the k scores at them (§2.1), not the scoring of all L keys. This insight motivates us to keep the indices from the forward pass and, in the recomputation, rescore only the entries they point to, cutting selection recomputation from $O(L)$ to $O(k)$ per query.

The challenge is that existing selective recomputation (Korthikanti et al., 2023; PyTorch, 2026) cannot rescore only these keys. It decides for each operator whether to keep the operator's output or to rerun the operator in full. Even if it keeps the indices that top-k outputs, it cannot rerun the scoring operator for only the k entries they point to. Every policy it can express therefore wastes either memory or computation (Figure 5). Keeping all activations or the score table avoids rescoring but wastes memory, because both grow with the context. Keeping nothing, the indexer's inputs, or only the top-k indices wastes computation, because the recomputation still scores all L keys to obtain the k scores that the backward pass needs.

Selection replay. To address this challenge, we design *selection replay*, a recomputation policy for selection that decides what to keep and what to rerun per entry rather than per operator. The forward pass keeps the selection of each query, and the recomputation computes at most the k scores that the backward pass reads. Replay comes in two variants that differ in what the forward pass keeps. *Index replay* keeps only the k indices of each query and rescores only the k entries they point to. *Score replay* also keeps the k scores and rescores nothing.

Cost comparison. Table 2 in Appendix D derives the cost of every policy, and Figure 5 plots these costs for one GLM-5.1 layer at 1M tokens with 8K tokens per GPU. Both variants of replay sit in the lower-left corner, while every operator-level policy in the figure costs at least two orders of magnitude more in memory or in recomputation. We use index replay by default, since its rescoring costs under 1% of a training step at 1M and it saves about 5 GiB per GPU over score replay in GLM-5.1.

Implementation of index replay. We describe *index replay* in detail. Figure 3b shows index replay in one layer of HyScale (②). First, the forward pass scores all L keys and selects the top-k as in §3.2, and keeps the k selected indices of each query. Second, the recomputation does not score all keys or select again. It loads the kept indices, scores only the k entries they point to with a gather-and-score kernel, and runs sparse attention on the same indices. Third, the backward pass runs exactly as without replay.

To save GPU memory, the kept indices can be offloaded to pinned host memory with nearly zero overhead. The forward pass copies them to the host on a separate stream right after selection, overlapping with sparse attention. The recomputation copies them back when it enters the attention layer, overlapping with the projections.

Generality. Selection replay applies to any learned sparse attention that selects by top-k, whether its entries are tokens as in DSA or blocks as in CSA and QSA. Each such layer learns its selection only through the k selected scores (§2.1), so replay rescores only these entries (§4.5). Replay is also compatible with existing recomputation policies. It keeps the interface of the operator that scores and selects, so policies that decide per module what to rerun, such as Megatron-LM's, work with replay unchanged.

3.4 Holistic System Implementation

The two techniques cut the work of selection but leave two other costs of learned sparse attention. First, CP communication runs synchronously before selection and sparse attention. Second, sparse kernels incur heavy global memory traffic. HyScale addresses each cost with a property of learned sparse attention. Selection and sparse attention consume different CP data. Decoupled CP overlap therefore hides each transfer behind computation independent of it (§3.4.1). Adjacent queries also select largely the same keys. Locality-aware sparse kernels thus process two adjacent queries in one thread block (§3.4.2).

3.4.1 Decoupled CP Overlap for Selection and Attention

Bottleneck of CP communication. Under CP, each GPU holds only its own queries but needs the indexer keys and the KV states of all earlier tokens. Existing systems therefore all-gather both before selection, in the forward pass and again in the recomputation, and reduce-scatter their gradients in the backward pass (Figure 3a). These collectives are large and synchronous. Figure 2 shows the resulting cost when Megatron-LM trains GLM-5.1: the exposed CP transfer grows 4.8× from 256K tokens (CP32) to 1M (CP128).

Limits of existing CP optimizations. Existing CP optimizations designed for full attention do not carry over to learned sparse attention. Ring Attention (Liu et al., 2024) overlaps each transfer with the attention that consumes it, but sparse attention reads only a few selected entries per query and has too little computation to hide the transfer. DeepSpeed-Ulysses (Jacobs et al., 2023) saves communication by splitting attention across GPUs by heads, so that each GPU gathers only the KV of its own heads. In DSA, however, all attention heads read the same KV latent and all indexer heads read the same indexer key, so each GPU would still gather the whole latent and key of every token, exactly what all-gather transfers.

Opportunity from disjoint inputs. To optimize transfers in learned sparse attention, we propose decoupled CP overlap. Its key insight is that the two steps of learned sparse attention consume different CP data: selection reads only the indexer keys, and sparse attention, which runs after selection, reads only the KV states.

Design. HyScale decouples the CP communication of the two steps, sending the indexer keys and KV states separately, and overlaps each transfer with computation that does not consume it (Figure 3b). Each transfer is issued as soon as its data is ready and waited for only before its first consumer. In the forward pass and recomputation, the indexer-key gather overlaps the query projections, and the KV-state gather overlaps the query projections and selection. In the backward pass, the reduce-scatter of the KV-state gradient overlaps the indexer backward, and that of the indexer-key gradient overlaps the query projections' backward.

3.4.2 Locality-Aware Sparse Kernels

Bottleneck of sparse kernels. Besides selection, learned sparse attention adds three sparse kernels to training: the sparse attention forward, its backward, and the indexer backward (Figure 3). Existing implementations (Li et al., 2025; NVIDIA, 2026; Wang et al., 2025a) assign each thread block a single query (Figure 6a). This per-query split costs L2 traffic in two forms. First, *per-query loads* (①): the forward kernel loads each query's selected entries separately. Second, *per-query atomics* (②): the backward kernels load entries the same way. They also add each

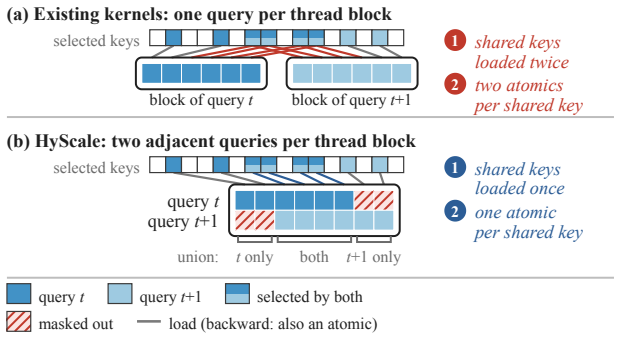


Figure 6: Locality-aware sparse kernels. Existing kernels (a) assign each thread block a single query; HyScale (b) processes two adjacent queries in one thread block over the union of their selections. Each row in a block marks the keys that one query uses.

query's gradient to its entries with separate atomics. Profiling the backward kernels for GLM-5.1 at 512K tokens on H800 confirms both costs. The L2 cache runs at 53–65% of its peak throughput, and warps stall on memory in 40–44% of their cycles. The tensor cores are only 21–32% busy.

Opportunity from selection locality. To address these bottlenecks, we propose locality-aware sparse kernels. Their key insight is *selection locality*: adjacent queries select largely the same keys, as prior work has observed (Yin et al., 2026; Liu et al., 2026a; Meng & Fang, 2026). In GLM-5.1, for example, two adjacent queries share 71% of their top-2048 keys on average, measured over 21 of its 78 layers on real documents of 128K and 192K tokens; the union of their selections is thus only 1.29 times one selection. Locality-aware kernels therefore process two adjacent queries in one thread block over the union of their selections (Figure 6b), which reduces both bottlenecks. First, each kernel loads a shared key once for both queries (①). Second, the backward kernels sum both queries' gradients of a shared key before adding them with one atomic (②).

Design. Once per layer, a lightweight preprocessing step builds the unions of all adjacent query pairs in one batch. All three sparse kernels reuse these unions. Each thread block reads the union of its own pair. The union lists the keys selected only by the first query, then the shared keys, then the keys selected only by the second (Figure 6b). The first query uses the first two parts. The second query uses the last two. Each query masks out the remaining part, the hatched cells, so outputs and gradients stay exact. This design spends otherwise idle tensor-core cycles on the masked part to save L2 traffic on the shared keys.

3.4.3 Implementation Details

HyScale is implemented on top of Megatron-LM with about 3.4K lines of Python and 14.7K lines of GPU kernels (CUDA, Triton, CuTe DSL). Besides the two optimizations above, four techniques target two overheads in Megatron-LM's attention layers: host synchronizations and wasted GPU work. The first two remove host synchronizations. (1) *CP layout cache*: Megatron-LM rebuilds the CP layout with its masks in every layer and pass, synchronizing with the host for every CP rank. HyScale builds it once per micro-batch. (2) *Sequence boundaries cache*: Megatron-LM's RoPE also synchronizes in every call to read the sequence boundaries, which HyScale copies to the host once per micro-batch. The other two remove wasted GPU work. (3) *Kernel fusion*: HyScale fuses chains of elementwise operators, such as the indexer's KL loss, normalizations, and RoPE. (4) *Permutation-free KV gather*: Megatron-LM reorders the gathered KV states into token order in every pass, so that sparse attention can index them by token. HyScale instead adjusts the indices to the gathered order.

4 Evaluation

4.1 Setup

Testbed. Our experiments run on an H20 cluster and an H800 cluster. Each H20 node has eight NVIDIA H20 GPUs with 96 GB of memory, two 96-core AMD EPYC 9K84 CPUs, and 2.2 TiB of DRAM. Each H800 node has eight NVIDIA H800 GPUs with 80 GB of memory, two 48-core Intel Xeon Platinum 8476C CPUs, and 2 TiB of DRAM. Nodes in both clusters connect through eight 400 Gb/s RoCE NICs (NVIDIA ConnectX-7).

Baseline and configuration. We compare HyScale with Megatron-LM, which runs Megatron Core 0.19.2 with the cuDNN DSA backend. HyScale builds on the same code base, ensuring the two systems differ only in the attention layers. All runs use BF16, zigzag all-gather CP, and full activation recomputation unless noted otherwise. In HyScale, index replay and the offloading of its indices are enabled.

4.2 End-to-End Training

We train the full GLM-5.1 from its released weights on 512 H20 GPUs, fixing the global batch at 4M tokens and the tokens per GPU at 8K. Each sample is a single sequence of 256K, 512K, or 1M tokens from real documents. The runs use PP8×CP32×DP2×EP64 at 256K, PP8×CP64×EP64 at 512K, and PP4×CP128×EP128 at 1M.

HyScale trains much faster with similar memory. HyScale is up to 2.07× faster than Megatron-LM, and the

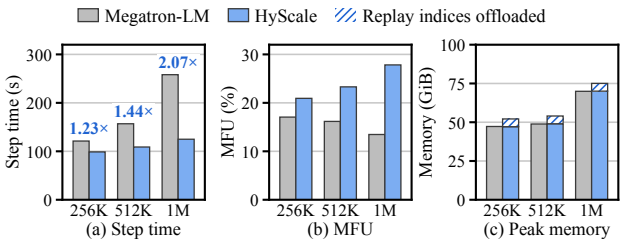


Figure 7: Full GLM-5.1 on 512 H20 GPUs. Both systems force balanced MoE routing and use momentum-free SGD with Megatron-LM's distributed optimizer. Step times average steps 4 to 10.

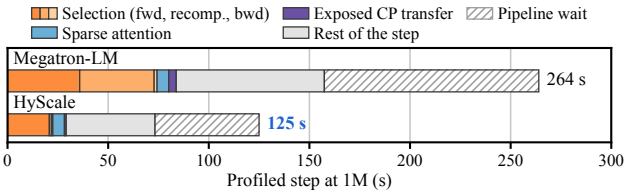


Figure 8: Breakdown of one full GLM-5.1 training step at 1M.

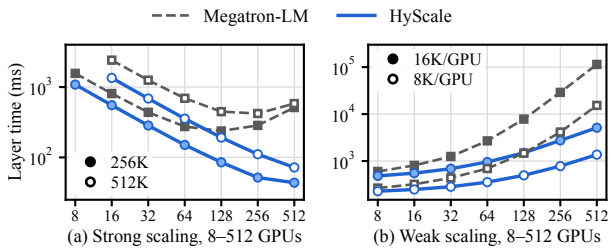


Figure 9: (a) Strong and (b) weak CP scaling of one GLM-5.1 attention layer on H20. Layer time covers the forward, recomputation, and backward passes on the slowest rank.

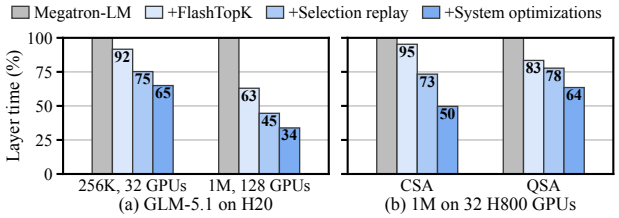


Figure 10: Cumulative ablation of one attention layer, as a share of Megatron-LM's layer time: (a) GLM-5.1 on H20 at 8K queries per GPU, (b) CSA and QSA on 32 H800 GPUs.

speedup grows with the context, from 1.23× at 256K to 2.07× at 1M (Figure 7a). HyScale achieves this speedup with about the same peak GPU memory (Figure 7c). Even without offloading the replay indices (§3.3), it would only use at most 11% more.

HyScale becomes more efficient at longer contexts. From 256K to 1M, with the same 4M tokens in every step, HyScale's step time grows only 1.27×, and its MFU even rises, from 20.9% to 27.8% (Figure 7b). Over the same range, Megatron-LM's step time grows 2.13×, and its MFU falls from 17.1% to 13.5%.

HyScale's advantage stems mainly from faster selection, the main bottleneck that grows with the context. On the first pipeline stage at 1M (Figure 8), FlashTopK cuts the forward selection from 35.9 s to 20.8 s, and selection replay cuts the recomputed selection from 37.0 s to 1.0 s. Faster attention layers also halve the pipeline wait, from 106.6 s to 51.7 s.

4.3 Context-Parallel Scalability

We benchmark a single GLM-5.1 attention layer on 8 to 512 H20 GPUs, with the CP size equal to the GPU count, under two standard settings (Jacobs et al., 2023) (Figure 9). First, strong scaling fixes the context length and adds GPUs. Second, weak scaling fixes the queries per GPU and grows the context with the GPUs.

HyScale keeps speeding up with more GPUs. At both 256K and 512K context, HyScale runs faster every time the GPU count doubles (Figure 9a). Up to 256 GPUs, each doubling cuts the layer time by 1.65 to 1.96×. Beyond 256 GPUs, the gain drops to 1.18 to 1.54×, because each GPU holds too few queries to hide the CP transfers. HyScale's speedup over Megatron-LM grows with the GPU count, from 1.44 to 11.7× at 256K and from 1.79 to 8.1× at 512K. In fact, Megatron-LM stops improving as early as 128 GPUs, and even slows down beyond 128 GPUs. This slowdown mainly stems from rebuilding the CP layout in every layer, incurring host synchronizations that grow linearly with the GPU count. HyScale avoids this cost with its CP layout cache (§3.4.3). The cache alone does not explain HyScale's speedup, however. At 1M context on 128 GPUs, all system optimizations together, the cache included, save only 11% of Megatron-LM's layer time (§4.4). Even if the cache saved all of it, HyScale would still run 2.6× faster than Megatron-LM.

HyScale's CP extends the context gracefully. HyScale is faster than Megatron-LM at all context lengths, and its layer time grows far more slowly as the context grows (Figure 9b). With 16K queries per GPU, from 128K to 8M context, HyScale's layer time grows only 11×, while Megatron-LM's grows 190×. HyScale is 1.24× faster than Megatron-LM at 128K context, and by 8M, this speedup climbs to 22.3×. The runs with 8K queries per GPU show the same trend, with an 11.2× speedup at 4M.

4.4 Ablation Study

We measure the layer time of one GLM-5.1 attention layer at 256K context on 32 GPUs and at 1M on 128 GPUs, with the CP size equal to the GPU count and 8K queries per GPU (Figure 10a). Starting from Megatron-LM, we add FlashTopK (§3.2), selection replay (§3.3), and the system optimizations of §3.4 one at a time. Each step is measured on top of the previous ones. We report each step's saving as a share of Megatron-LM's layer time.

FlashTopK and selection replay bring most of the speedup. At 1M context, HyScale saves 66% of Megatron-LM's layer time (Figure 10a). The two techniques deliver 84% of this saving. FlashTopK saves 37% of the layer time, because it removes the score table from both the forward pass and the recomputation. Selection replay saves another 18%, because it skips the recomputed selection. The system optimizations save the remaining 11%. FlashTopK's share also grows fastest, from 8% at 256K to 37% at 1M, because it replaces Megatron-LM's unfused selection, whose cost grows faster than the context. Megatron-LM cannot hold the 32 GiB score table at 1M, so it splits the queries into chunks whose score tables fit a fixed memory limit. Longer contexts leave fewer queries per chunk, too few at 1M to keep the GPU busy. FlashTopK avoids this trade-off: it writes no score table, so it scores all queries at once. Selection replay keeps 16–18% at both contexts, because FlashTopK already accelerates the selection it skips.

4.5 Generality

This experiment tests whether HyScale carries over to block-level architectures. We implement the CSA layer of DeepSeek-V4-Flash (DeepSeek-AI et al., 2026) and the QSA layer of Qwen3.8-Flash-Next (Qiu et al., 2026) in Megatron-LM from their papers. We repeat §4.4's ablation on 32 H800 GPUs at 1M context with CP32 (Figure 10b).

HyScale generalizes across learned sparse attention. HyScale runs CSA 2.01× and QSA 1.57× faster than Megatron-LM. Every ablation step speeds up both layers, showing that HyScale can also accelerate block-level architectures. FlashTopK's gain on block-level layers is smaller than on token-level DSA (§4.4), since they score 4× fewer keys per query, making selection less significant.

4.6 FlashTopK Efficiency

Operator efficiency. We compare FlashTopK with Megatron-LM's cuDNN selection on one H20 GPU. We benchmark all selections captured from full GLM-5.1 training at 256K, 512K, and 1M context, on all 78 layers and CP ranks. Each selection picks $k = 2048$ keys for

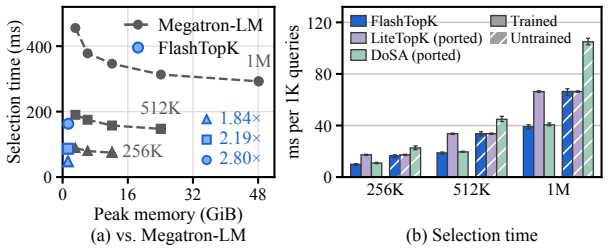


Figure 11: FlashTopK efficiency for GLM-5.1 on one H20 GPU.

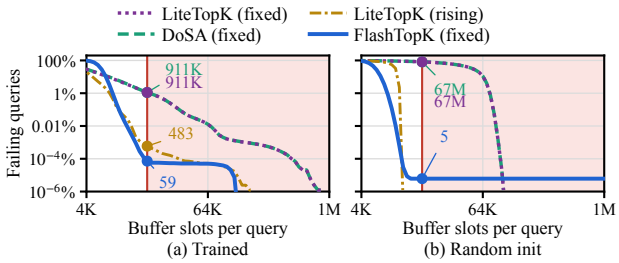


Figure 12: Sampling efficiency: share of failed queries at each buffer size. Labels count failures at FlashTopK's default 16K-slot candidate buffer (red line), 1 GiB per GPU.

8K queries. Among fused sampling-based methods, only FlashTopK supports CP and determinism (§3.2). Moreover, LiteTopK and DoSA release kernels only for FP8 on Blackwell GPUs. We thus reimplement both in BF16 on H20 following their code, and run them without CP. Both are also non-deterministic, while FlashTopK keeps its deterministic sort.

FlashTopK removes the memory–time trade-off of selection. FlashTopK uses the least memory and runs the fastest at every context (Figure 11a). Including its buffers, it peaks at only 1.3–1.4 GiB per layer. Megatron-LM must instead trade memory for time. With its default 1 GiB limit on the score table, it peaks at 3.1 GiB but runs 1.84–2.80× slower than FlashTopK. With the whole score table, it needs 12–48 GiB and still runs 55–80% slower. The limit saves memory, but it splits the queries into chunks too small to keep the GPU busy (§4.4). FlashTopK avoids both costs because it never writes the score table.

FlashTopK is fast with both trained and untrained weights. It matches or outperforms LiteTopK and DoSA in both settings (Figure 11b). With trained weights, FlashTopK runs 1.70–1.79× faster than LiteTopK and 1.04–1.10× faster than DoSA. FlashTopK and DoSA both outpace LiteTopK mainly through pruning. Pruning cuts FlashTopK's time by 39–42%, while LiteTopK never prunes. FlashTopK further outpaces DoSA because neighbor sampling forces DoSA's query blocks to run one after another. With untrained weights, FlashTopK runs 1.34–1.58× faster than DoSA and 0.99–1.06× as fast as

LiteTopK. The lead over DoSA is large because DoSA always prunes, and pruning adds large overhead once the bound loosens (C2 in §3.2). With neighbor sampling and a fixed threshold, 20% of DoSA's queries also overflow at 1M. FlashTopK stays on par with LiteTopK because pruning rarely pays off, yet adaptive pruning still turns on in some calls. Overall, optimistic strided sampling and adaptive pruning keep FlashTopK fast and robust.

Sampling efficiency. Next, we compare FlashTopK's optimistic strided sampling with the neighbor sampling of LiteTopK and DoSA. We count the candidates of all 81.8M queries in one GLM-5.1 training step at 1M, with the trained weights and with a randomly initialized indexer. LiteTopK raises its threshold during the scan (§3.2), while DoSA and FlashTopK keep theirs fixed. We show LiteTopK both ways.

Optimistic thresholds outperform guaranteed lower bounds. With fixed thresholds and trained weights, FlashTopK fails on over 15,000× fewer queries than LiteTopK and DoSA at its 16K-slot candidate buffer (Figure 12a). With random weights, both baselines fail on 82% of the queries, and FlashTopK on only 5 of the 81.8M (Figure 12b). Raising the threshold during the scan cuts LiteTopK's failures to 483 with trained weights, still 8× more than FlashTopK, and to none with random weights. FlashTopK's lightweight recovery rescores only the failed queries on the GPU, with near-zero overhead.

These results reflect the limits of neighbor sampling in training. Its bound comes from the keys that nearby queries selected. Their k-th best can thus lie deep in this query's ranking and let many extra keys pass. With random weights, nearby queries also barely share keys (C2 in §3.2). Under CP, their selections are not even available (C1). FlashTopK instead samples evenly from each query's own context. Combined with lightweight recovery, optimistic strided sampling thus suits training better.

5 Related Work

Selection acceleration for learned sparse attention. HPC-Ops Top-K (Liu et al., 2026b), GVR (Cheng et al., 2026), DeepSelect (Qian et al., 2026), and Prof-K (Dziarmaga et al., 2026) accelerate the top-k step but leave scoring in a separate kernel. The score table thus still makes a round trip through GPU memory. Fusing scoring with top-k selection removes it. LiteTopK (Yin et al., 2026) pioneered this fusion with sampling-based selection, taking its threshold from nearby queries' selections. DoSA (Meng & Fang, 2026) further skips part of the scoring with cheap per-key bounds. Both, however, build their sampling around prefill and face challenges in training (§3.2). FlashTopK addresses them with optimistic strided sampling, adaptive key pruning, and deterministic selection. Orthogonally, PIVOT (Liu et al., 2026a), IndexCache (Bai et al., 2026), HISA (Xu et al., 2026), and MISA (Zhou et al., 2026) cut selection work by reusing selections or scoring fewer entries or heads.

Activation recomputation. Activation recomputation saves memory by rerunning the forward pass during backpropagation (Chen et al., 2016), at the cost of extra computation. Prior work cuts this cost by choosing what to recompute. Checkmate (Jain et al., 2019) and DTR (Kirisame et al., 2020) pick tensors offline or at runtime; selective recomputation (Korthikanti et al., 2023; PyTorch, 2026) reruns only cheap operators with large outputs. DeepSeek-V4 (DeepSeek-AI et al., 2026) further makes this choice per tensor by extending automatic differentiation. All of them, however, rerun whole operators. In learned sparse attention, this wastes memory or computation (§3.3). Selection replay avoids both by rescoring only the selected entries.

Context parallelism. The basic CP design all-gathers all KV states before attention (Grattafiori et al., 2024), leaving this communication exposed. Two schemes optimize it. Ring Attention (Liu et al., 2024), Striped Attention (Brandon et al., 2023), and DistFlashAttn (Li et al., 2023) hide it behind attention computation by passing KV blocks around a ring. DeepSpeed-Ulysses (Jacobs et al., 2023) reduces it by splitting attention heads across GPUs. USP (Fang & Zhao, 2024) and LoongTrain (Gu et al., 2024) further combine both for hierarchical GPU networks. Both schemes, however, fall short in learned sparse attention (§3.4.1). Ring designs lack the computation to hide transfers (Bu et al., 2026), and Ulysses cannot split KV states shared across heads, such as in DSA. HyScale instead hides the KV communication behind selection. Orthogonally, MagiAttention (Tao & Huang, 2025), ByteScale (Ge et al., 2025), WLB-LLM (Wang et al., 2025c), and FlexSP (Wang et al., 2025b) balance the CP workload across GPUs.

6 Conclusion

We presented HyScale, a long-context training system for learned sparse attention. Its key insight is that selection outputs only indices, not activations. FlashTopK thus selects the top-k without writing the score table, and selection replay skips repeated selection in recomputation. HyScale trains GLM-5.1 on 1M-token sequences 2.07× faster than Megatron-LM on 512 H20 GPUs and generalizes across learned sparse attention architectures. Future work includes extending FlashTopK to inference, other precisions such as FP8, and newer GPUs such as Blackwell.

Acknowledgements

We thank Tian Zheng from NVIDIA and Yeqi Huang from the University of Edinburgh for insightful discussions and helpful feedback.

References

Agrawal, A., Kedia, N., Panwar, A., Mohan, J., Kwatra, N., Gulavani, B., 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 24)*, pp. 117–134. USENIX Association, 2024. URL https://www.usenix.org/conference/osdi24/presentation/agrawal.

Alabi, T., Blanchard, J. D., Gordon, B., and Steinbach, R. Fast k-selection algorithms for graphics processing units. *ACM Journal of Experimental Algorithmics*, 17, 2012. doi: 10.1145/2133803.2345676. URL http://dx.doi.org/10.1145/2133803.2345676.

Bai, Y., Dong, Q., Jiang, T., Lv, X., Du, Z., Zeng, A., Tang, J., and Li, J. IndexCache: Accelerating sparse attention via Cross-Layer index reuse. *arXiv preprint arXiv:2603.12201*, 2026. URL https://arxiv.org/abs/2603.12201.

Brandon, W., Nrusimha, A., Qian, K., Ankner, Z., Jin, T., Song, Z., and Ragan-Kelley, J. Striped attention: Faster ring attention for causal transformers. *arXiv preprint arXiv:2311.09431*, 2023. URL https://arxiv.org/abs/2311.09431.

Bu, T., Wang, Q., Zeng, B., Sun, H., Huang, Y., Cao, C., and Xu, J. Long-Context attention benchmark: From kernel efficiency to distributed context parallelism. In *The Fourteenth International Conference on Learning Representations*, 2026. URL https://openreview.net/forum?id=W7sVYFJAEp.

Chen, T., Xu, B., Zhang, C., and Guestrin, C. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016. URL https://arxiv.org/abs/1604.06174.

Cheng, L., Zhao, R., Liu, T., Li, M., Qiao, X., Duan, K., Chen, Y.-J., Chen, X., Rouhani, B. D., and Yang, J. Guess-Verify-Refine: Data-Aware Top-K for Sparse-Attention decoding on Blackwell via temporal correlation. *arXiv preprint arXiv:2604.22312*, 2026. URL https://arxiv.org/abs/2604.22312.

Dao, T., Fu, D., Ermon, S., Rudra, A., and Ré, C. FlashAttention: Fast and Memory-Efficient exact attention with IO-Awareness. In *Advances in Neural Information Processing Systems 35*, pp. 16344–16359. Neural Information Processing Systems Foundation, Inc. (NeurIPS), 2022. doi: 10.52202/068431-1189. URL http://dx.doi.org/10.52202/068431-1189.

DeepSeek-AI, Liu, A., Feng, B., Wang, B., Wang, B., Liu, B., Zhao, C., Dengr, C., Ruan, C., Dai, D., et al. DeepSeek-V2: A strong, economical, and efficient Mixture-of-Experts language model. *arXiv preprint arXiv:2405.04434*, 2024. URL https://arxiv.org/abs/2405.04434.

DeepSeek-AI, Liu, A., Mei, A., Lin, B., Xue, B., Wang, B., Xu, B., Wu, B., Zhang, B., Lin, C., et al. DeepSeek-V3.2: Pushing the frontier of open large language models. *arXiv preprint arXiv:2512.02556*, 2025. URL https://arxiv.org/abs/2512.02556.

DeepSeek-AI, Xu, A., Lin, B., Xue, B., Wang, B., Xu, B., Wu, B., Zhang, B., Lin, C., Dong, C., et al. DeepSeek-V4: Towards highly efficient Million-Token context intelligence. *arXiv preprint arXiv:2606.19348*, 2026. URL https://arxiv.org/abs/2606.19348.

Dziarmaga, T., Sikora, W., Struski, Ł., Tabor, J., and Mazur, M. Prof-K: Probabilistic One-Pass filtering for efficient top-k selection. *arXiv preprint arXiv:2608.12573*, 2026. URL https://arxiv.org/abs/2608.12573.

Fang, J. and Zhao, S. USP: A unified sequence parallelism approach for long context generative AI. *arXiv preprint arXiv:2405.07719*, 2024. URL https://arxiv.org/abs/2405.07719.

Ge, H., Feng, J., Huang, Q., Fu, F., Nie, X., Zuo, L., Lin, H., Cui, B., and Liu, X. ByteScale: Efficient scaling of LLM training with a 2048K context length on more than 12,000 GPUs. *arXiv preprint arXiv:2502.21231*, 2025. URL https://arxiv.org/abs/2502.21231.

GLM-5-Team, Zeng, A., Lv, X., Hou, Z., Du, Z., Zheng, Q., Chen, B., Yin, D., Ge, C., Huang, C., et al. GLM-5: from vibe coding to agentic engineering. *arXiv preprint arXiv:2602.15763*, 2026. URL https://arxiv.org/abs/2602.15763.

Grattafiori, A., Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Vaughan, A., et al. The Llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024. URL https://arxiv.org/abs/2407.21783.

Gu, D., Sun, P., Hu, Q., Huang, T., Chen, X., Xiong, Y., Wang, G., Chen, Q., Zhao, S., Fang, J., Wen, Y., Zhang, T., Jin, X., and Liu, X. LoongTrain: Efficient training of Long-Sequence LLMs with Head-Context parallelism. *arXiv preprint arXiv:2406.18485*, 2024. URL https://arxiv.org/abs/2406.18485.

Jacobs, S. A., Tanaka, M., Zhang, C., Zhang, M., Song, S. L., Rajbhandari, S., and He, Y. DeepSpeed ulysses: System optimizations for enabling training of extreme long sequence transformer models. *arXiv preprint*

arXiv:2309.14509, 2023. URL https://arxiv.org/abs/2309.14509.

Jain, P., Jain, A., Nrusimha, A., Gholami, A., Abbeel, P., Keutzer, K., Stoica, I., and Gonzalez, J. E. Checkmate: Breaking the memory wall with optimal tensor rematerialization. *arXiv preprint arXiv:1910.02653*, 2019. URL https://arxiv.org/abs/1910.02653.

Kirisame, M., Lyubomirsky, S., Haan, A., Brennan, J., He, M., Roesch, J., Chen, T., and Tatlock, Z. Dynamic tensor rematerialization. *arXiv preprint arXiv:2006.09616*, 2020. URL https://arxiv.org/abs/2006.09616.

Korthikanti, V. A., Casper, J., Lym, S., McAfee, L., Andersch, M., Shoeybi, M., and Catanzaro, B. Reducing activation recomputation in large transformer models. In *Proceedings of Machine Learning and Systems*, volume 5, pp. 341–353, 2023. URL https://proceedings.mlsys.org/paper_files/paper/2023/file/80083951326cf5b35e5100260d64ed81-Paper-mlsys2023.pdf.

Lai, X., Xu, W., Yang, Y., Chen, Q., Xu, Y., Zeng, L., Li, X., Sun, H., Zhu, H., Zhang, V., Hu, J., Li, J., Gao, R., Li, Z., Zhu, S., Zhou, J., and Zhao, P. MiniMax sparse attention. *arXiv preprint arXiv:2606.13392*, 2026. URL https://arxiv.org/abs/2606.13392.

Li, D., Shao, R., Xie, A., Xing, E. P., Ma, X., Stoica, I., Gonzalez, J. E., and Zhang, H. DISTFLASHATTN: Distributed memory-efficient attention for long-context LLMs training. *arXiv preprint arXiv:2310.03294*, 2023. URL https://arxiv.org/abs/2310.03294.

Li, J., Liu, S., and Sun, Y. FlashMLA: Efficient multi-head latent attention kernels. https://github.com/deepseek-ai/FlashMLA, 2025.

Liu, H., Zaharia, M., and Abbeel, P. RingAttention with blockwise transformers for Near-Infinite context. In *The Twelfth International Conference on Learning Representations*, 2024. URL https://openreview.net/forum?id=WsRHpHH4s0.

Liu, H., Cheng, Y., Niu, L., Su, Y., Xue, Y., Liu, A., Yu, G., and Zhu, J. PIVOT: Efficient Query-Group indexing for Token-Level sparse attention. *arXiv preprint arXiv:2607.24593*, 2026a. URL https://arxiv.org/abs/2607.24593.

Liu, S., Xue, Y., Tang, T., Shao, C., Cheng, Q., Ren, H., Jiang, D., Ming, H., Ding, L., Lin, Z., Wei, S., Liu, W., Liu, K., and Zhu, J. Sample-Guided exact Top-K selection for Long-Context sparse attention. *arXiv preprint arXiv:2609.08450*, 2026b. URL https://arxiv.org/abs/2609.08450.

Meng, F. and Fang, T. DoSA: Dots sparse attention. *arXiv preprint arXiv:2610.13990*, 2026. URL https://arxiv.org/abs/2610.13990.

Merrill, D. and Grimshaw, A. High performance and scalable radix sorting: A case study of implementing dynamic parallelism for GPU computing. *Parallel Processing Letters*, 21(02):245–272, 2011. doi: 10.1142/s0129626411000187. URL http://dx.doi.org/10.1142/s0129626411000187.

NVIDIA. DeepSeek sparse attention (DSA). NVIDIA cuDNN, https://docs.nvidia.com/deeplearning/cudnn/latest/fe-oss-apis/dsa.html, 2026.

PyTorch. torch.utils.checkpoint. PyTorch main documentation, https://docs.pytorch.org/docs/main/utils/checkpoint.html, 2026.

Qian, Y., Liu, S., and Li, Y. DeepSelect: High-Performance TopK kernels for DeepSeek sparse attention and sampling. https://github.com/deepseek-ai/DeepSelect, 2026.

Qiu, Z., Wang, Z., Li, X., Li, Y., Xu, Y., Wang, Y., Zhang, H., Men, R., Mao, B., Zhang, C., Zhou, F., Luo, H., Huang, H., Lian, H., Huang, H., Chen, H., Zhang, J., Xu, J., Wang, J., Chen, L., Wang, L., Jiang, L., Yuan, M., Sun, M., Jin, P., Zhang, S., Wang, S., Ren, X., Wang, Y., Zhang, Y., Dong, Y., Cao, Y., Ma, Y., Mao, Y., Zheng, B., and Liu, D. On the design of Qwen3.8-Next architecture: Evaluation, efficiency, and training stability. *arXiv preprint arXiv:2608.30320*, 2026. URL https://arxiv.org/abs/2608.30320.

Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., and Catanzaro, B. Megatron-LM: Training Multi-Billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019. URL https://arxiv.org/abs/1909.08053.

Sul, S. H., Brown, N., Wildermuth, H., Lin, W., Cassano, F., and Ré, C. Mixture-of-Kittens: MoE megakernel for NVL72s. *arXiv preprint arXiv:2609.36070*, 2026. URL https://arxiv.org/abs/2609.36070.

Tao, Z. and Huang, Y. MagiAttention: A distributed attention towards linear scalability for Ultra-Long context, heterogeneous mask training. https://github.com/SandAI-org/MagiAttention, 2025.

Tencent-Hunyuan. Hy4-preview. https://github.com/Tencent-Hunyuan/Hy4-preview, 2026.

Wang, L., Cheng, Y., Shi, Y., Tang, Z., Mo, Z., Xie, W., Ma, L., Xia, Y., Xue, J., Yang, F., and Yang, Z. TileLang: A composable tiled programming model for AI systems. *arXiv preprint arXiv:2504.17577*, 2025a. URL https://arxiv.org/abs/2504.17577.

Wang, Y., Wang, S., Zhu, S., Fu, F., Liu, X., Xiao, X., Li, H., Li, J., Wu, F., and Cui, B. FlexSP: Accelerating large language model training via flexible sequence parallelism. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming*

Languages and Operating Systems, Volume 2, pp. 421–436. ACM, 2025b. doi: 10.1145/3676641.3715998. URL http://dx.doi.org/10.1145/3676641.3715998.

Wang, Z., Cai, A., Xie, X., Pan, Z., Guan, Y., Chu, W., Wang, J., Li, S., Huang, J., Cai, C., Hao, Y., and Ding, Y. WLB-LLM: Workload-Balanced 4D parallelism for large language model training. In *19th USENIX Symposium on Operating Systems Design and Implementation (OSDI 25)*, pp. 785–801. USENIX Association, 2025c. URL https://www.usenix.org/conference/osdi25/presentation/wang-zheng.

Xu, Y., Meng, F., Jiang, F., Wang, Y., Zhou, R., Wang, Z., Wu, J., Pan, Z., Tang, X., Pei, W., Liu, T., Yin, D., Sun, X., and Zhang, M. HISA: Efficient hierarchical indexing for Fine-Grained sparse attention. *arXiv preprint arXiv:2603.28458*, 2026. URL https://arxiv.org/abs/2603.28458.

Yin, Z., Gao, J., Yin, P., Li, J., and Cong, G. LiteTopK: Exploiting the curse of dimensionality for a fused Indexer-TopK kernel in Long-Context sparse attention. *arXiv preprint arXiv:2607.11976*, 2026. URL https://arxiv.org/abs/2607.11976.

Yuan, J., Gao, H., Dai, D., Luo, J., Zhao, L., Zhang, Z., Xie, Z., Wei, Y., Wang, L., Xiao, Z., Wang, Y., Ruan, C., Zhang, M., Liang, W., and Zeng, W. Native sparse attention: Hardware-Aligned and natively trainable sparse attention. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 23078–23097. Association for Computational Linguistics, 2025. doi: 10.18653/v1/2025.acl-long.1126. URL http://dx.doi.org/10.18653/v1/2025.acl-long.1126.

Zhou, R., Meng, F., Xu, Y., Liu, T., Lu, G., Zhang, M., and Pei, W. MISA: Mixture of indexer sparse attention for Long-Context LLM inference. *arXiv preprint arXiv:2605.07363*, 2026. URL https://arxiv.org/abs/2605.07363.

A Formulation of Learned Sparse Attention

Consider a sequence of L tokens. For the query at position t, let $\mathcal{U}_t$ be the set of entries it can select from. *Selection* scores every entry and keeps the k highest,

$$\mathbf{I}_t = \big(\mathrm{Score}(\mathbf{q}_t^I, \mathbf{k}_u^I)\big)_{u\in\mathcal{U}_t}, \quad \mathcal{T}_t = \text{Top-k}(\mathbf{I}_t), \tag{1}$$

where $\mathbf{q}_t^I$ is the indexer's query, $\mathbf{k}_u^I$ is the indexer's key of entry u, $\mathbf{I}_t \in \mathbb{R}^{|\mathcal{U}_t|}$ is the score vector, and $\mathcal{T}_t \subseteq \mathcal{U}_t$ is the index set of size k. *Sparse attention* then attends from query $\mathbf{q}_t$ to the selected entries alone,

$$\mathbf{o}_t = \mathrm{Attn}\big(\mathbf{q}_t, \{\mathbf{c}_u\}_{u\in\mathcal{T}_t}\big), \tag{2}$$

where $\mathbf{c}_u$ is what sparse attention reads for entry u, e.g., a KV latent, and $\mathbf{o}_t$ is the output. The indexer learns from a loss of its own,

$$\mathcal{L}^I = \lambda \sum_t \mathrm{KL}\big(\mathbf{p}_t \,\big\|\, \mathrm{Softmax}(\mathbf{I}_{t,\mathcal{T}_t})\big), \tag{3}$$

where λ is a scalar weight, $\mathbf{p}_t \in \mathbb{R}^k$ is the attention distribution of query t over $\mathcal{T}_t$ averaged over heads, and $\mathbf{I}_{t,\mathcal{T}_t} \in \mathbb{R}^k$ holds the indexer's scores on those entries. Both $\mathbf{p}_t$ and the indexer's inputs are detached. This loss thus updates only the indexer, through the k selected scores alone.

DSA scores each token with a lightning indexer of H_I heads of dimension d_I,

$$\mathrm{Score}(\mathbf{q}_t^I, \mathbf{k}_u^I) = \sum_{h=1}^{H_I} w_{t,h}\,\mathrm{ReLU}\big(\mathbf{q}_{t,h}^{I\top}\,\mathbf{k}_u^I\big), \tag{4}$$

where $\mathbf{q}_{t,h}^I \in \mathbb{R}^{d_I}$ is the query of head h, the heads share one key $\mathbf{k}_u^I \in \mathbb{R}^{d_I}$ per token, and $w_{t,h}$ weighs each head. In MLA, all heads share one compressed KV latent per token. This latent serves as $\mathbf{c}_u$.

B Backward Pass of Learned Sparse Attention

This appendix derives the backward pass of one learned sparse attention layer in the notation of Appendix A, for DSA with MLA. The backward passes of selection and sparse attention use only the selected entries.

Selection. With $z_{t,h,s} = \mathbf{q}_{t,h}^{I\top}\mathbf{k}_s^I$, Equation (4) reads $I_{t,s} = \sum_h w_{t,h}\,\mathrm{ReLU}(z_{t,h,s})$. Let $\hat{\mathbf{p}}_t = \mathrm{Softmax}(\mathbf{I}_{t,\mathcal{T}_t})$ be the indexer's distribution over the selection. The target $\mathbf{p}_t$ is detached, and $\mathcal{T}_t$ is piecewise constant in the scores. The indexer loss of Equation (3) thus has the score gradient

$$g_{t,s} = \frac{\partial \mathcal{L}^I}{\partial I_{t,s}} = \begin{cases} \lambda\,(\hat{p}_{t,s} - p_{t,s}) & s \in \mathcal{T}_t, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

Let $\delta_{t,h,s} = g_{t,s}\,w_{t,h}\,\mathbf{1}[z_{t,h,s} > 0]$. The chain rule gives

$$\frac{\partial \mathcal{L}^I}{\partial w_{t,h}} = \sum_{s\in\mathcal{T}_t} g_{t,s}\,\mathrm{ReLU}(z_{t,h,s}), \tag{6}$$

$$\frac{\partial \mathcal{L}^I}{\partial \mathbf{q}_{t,h}^I} = \sum_{s\in\mathcal{T}_t} \delta_{t,h,s}\,\mathbf{k}_s^I, \tag{7}$$

$$\frac{\partial \mathcal{L}^I}{\partial \mathbf{k}_s^I} = \sum_{t\,:\,s\in\mathcal{T}_t} \sum_{h=1}^{H_I} \delta_{t,h,s}\,\mathbf{q}_{t,h}^I. \tag{8}$$

Every sum runs over the k selected entries of a query. The backward pass thus needs $z_{t,h,s}$ only for $s \in \mathcal{T}_t$, or kH_I per-head scores per query. The forward scan instead computes all LH_I. Selection replay recomputes exactly these scores (§3.3). The indexer's inputs are detached. The gradients thus reach only the indexer's projection weights.

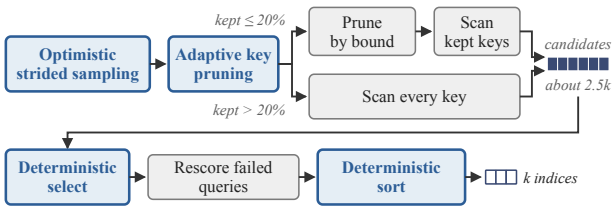


Figure 13: One selection call of FlashTopK. All steps run on the GPU without host synchronization.

Sparse attention. In MLA, head h of query t attends to the selected latents with logits $\ell_{t,h,s} = \mathbf{q}_{t,h}^\top \mathbf{c}_s / \sqrt{d}$ for $s \in \mathcal{T}_t$. Together with its positional part, $\mathbf{c}_s$ serves as the key. Its latent part $\mathbf{c}_s^V$ serves as the value. Let $a_{t,h,s} = \mathrm{softmax}_{s\in\mathcal{T}_t}(\ell_{t,h,s})$ be the attention weights and $\mathbf{o}_{t,h} = \sum_{s\in\mathcal{T}_t} a_{t,h,s}\, \mathbf{c}_s^V$ the output. The standard attention backward (Dao et al., 2022) restricted to $\mathcal{T}_t$ gives

$$\mathrm{d}\ell_{t,h,s} = a_{t,h,s}\big(\mathrm{d}\mathbf{o}_{t,h}^\top \mathbf{c}_s^V - \mathrm{d}\mathbf{o}_{t,h}^\top \mathbf{o}_{t,h}\big), \tag{9}$$

$$\mathrm{d}\mathbf{q}_{t,h} = \tfrac{1}{\sqrt{d}} \sum_{s\in\mathcal{T}_t} \mathrm{d}\ell_{t,h,s}\, \mathbf{c}_s, \tag{10}$$

$$\mathrm{d}\mathbf{c}_s = \sum_{t\,:\,s\in\mathcal{T}_t} \sum_{h=1}^{H} \Big(\tfrac{1}{\sqrt{d}}\, \mathrm{d}\ell_{t,h,s}\, \mathbf{q}_{t,h} + a_{t,h,s}\, \mathrm{d}\mathbf{o}_{t,h} \Big). \tag{11}$$

The first term of $\mathrm{d}\mathbf{c}_s$ comes from $\mathbf{c}_s$ as a key. The second comes from $\mathbf{c}_s$ as a value and lands on $\mathbf{c}_s^V$. The target $\mathbf{p}_t$ averages $a_{t,h,s}$ over heads and is detached. The indexer loss thus sends no gradient into sparse attention.

Context parallelism. Under CP, $\mathbf{k}_s^I$ and $\mathbf{c}_s$ are all-gathered. Each GPU thus computes partial $\partial\mathcal{L}^I/\partial\mathbf{k}_s^I$ and $\mathrm{d}\mathbf{c}_s$ from its own n queries for every position s. The backward reduce-scatter of §2.2 sums these partial gradients. Each GPU then receives the gradients of its own positions.

C FlashTopK Implementation

Figure 13 shows one selection call of FlashTopK. Optimistic strided sampling first sets each query's threshold (§3.2). Adaptive key pruning then estimates the share of keys that DoSA's bound keeps. If this share is at most 20%, FlashTopK prunes by the bound and scans only the kept keys. Otherwise, it scans every key. Either scan writes each query's passing keys to its 16K-slot candidate buffer. About $2.5k$ keys pass on average. The deterministic select then keeps each query's k best candidates and breaks ties at the k-th score by position. Recovery rescores the failed queries. The deterministic sort finally orders the k indices by score, then by position. Each branch reads a flag on the GPU. No step thus synchronizes with the host.

Optimistic strided sampling. The scan kernel scores the sampled keys of §3.2 with its own arithmetic. A sample's score thus has the same bits as in the scan. Each query reads only the samples inside its window. FlashTopK keeps LiteTopK's bucketed threshold. Its 256 buckets span the range of a query's sample scores.

Adaptive key pruning. FlashTopK estimates the kept share on one query in 16 against 256 sampled keys. Every call launches both paths. The other path sees empty windows and returns at once.

Prune by bound. FlashTopK adapts DoSA's upper bound on the score of key $\mathbf{k}$ for query t,

$$U = \textstyle\sum_{j\in R} \tilde{w}_j\, \mathrm{ReLU}(\mathbf{q}_j \cdot \mathbf{k}) - \mathrm{ReLU}(\mathbf{a} \cdot \mathbf{k}) + E\, \|\mathbf{k}\|.$$

The set R holds the 14 positive heads with the largest $w_h\|\mathbf{q}_h\|$. Each other positive head is projected onto its most aligned head in R. Its residual joins the allowance E. The allowance also covers rounding. The aggregate $\mathbf{a}$ sums the negative heads active on at least half of the keys. These heads can only lower the score. A key with a bound below the threshold thus cannot pass it. Each query's bound takes 16 columns, 14 for the retained heads and two for the negative aggregate. The pruning kernel takes 16 queries per tile. It writes one bit per key for each group of 8 queries.

Scan kept keys. The scan locates the kept keys with a prefix sum and a binary search. It packs them into dense tiles and scores them with the same arithmetic as the scan of every key.

Rescore failed queries. A check kernel marks a query as failed if its candidates overflowed the buffer or number fewer than $\min(k, L_t)$. Recovery kernels then rescore up to 64 failed queries over their windows and reselect them in place. They run after every selection with a fixed grid. Each returns at once when no query failed. With more than 64 failures, every query is rescored. This never occurred in our runs. At 1M, one failed query adds 0.3% to the selection time, and 64 add 2.6%.

Deterministic select and sort. The scans write candidates with atomics. Their order thus varies between runs. Sometimes more candidates tie at the k-th score than there are places left. The select kernel then runs a radix select over the positions of just these tied keys. It keeps the tied keys with the smallest positions. The selected set thus depends only on the scores. A radix sort then orders each query's k indices by score and breaks ties by position (§3.2). At 1M, the sort adds 0.8% to the selection time.

D Costs of Selection Recomputation

Table 2 gives the cost of every policy in Figure 5, in general and for GLM-5.1.

Table 2: Selection recomputation costs per layer and GPU. Memory excludes the layer input. It counts 4 B per index or score and 2 B per indexer input element. Recomputation counts 2 ops per multiply-add of scoring and 1 per score for top-k. GLM-5.1 has L = 1M, n = 8K, k = 2048, H_I = 32, and d_I = 128.

	Memory (B)		Recomputation (ops)	
Policy	General	GLM-5.1	General	GLM-5.1
Operator-level policies				
No recomp.	all	all	0	0
Full recomp.	0	0	$2nLH_Id_I+nL$	70.4 T
Keep scores	$4nL$	32 Gi	nL	8.6 G
Keep top-k	$4nk$	64 Mi	$2nLH_Id_I$	70.4 T
Keep inputs	$2(L+nH_I)d_I$	320 Mi	$2nLH_Id_I+nL$	70.4 T
Selection replay				
Index replay	$4nk$	64 Mi	$2nkH_Id_I$	137 G
Score replay	$8nk$	128 Mi	0	0