TELECHAT TO XING4: ADVANCING FRONTIER AGENTIC MODELS ON DOMESTIC CHIPS AND FRAMEWORKS

Xingchen AGI Lab, China Telecom

arXiv:2611.22115v1 [cs.AI] 17 Nov 2026

ABSTRACT

We present Xing 4, a series of large language models that extends the TeleChat lineage through coordinated advances in model architecture and training infrastructure on Ascend. While most open-source models are trained on GPUs, scaling on domestic chips and frameworks poses significant systems challenges. This report describes the design principles and systems engineering behind the series, spanning architecture selection, large-scale pretraining, and reasoning and agentic post-training. We study how sparse activation, attention design, model depth and width, and speculative decoding balance model capacity with computation, memory access, and execution overhead. Representative configurations combine MLA with DSA and cross-layer IndexShare for long-context serving, incorporate optimized MHC operators to enhance stability and computational efficiency, adopt a wider-and-shallower mixture-of-experts (MoE) topology, and pair a block-wise draft model with lightweight correction heads for speculative decoding. We have open-sourced a 29B MoE member, a model trained entirely on Ascend NPUs using MindSpore, activating only 4B parameters per token while supporting native 256K-token contexts. Our pretraining infrastructure builds on MindSpore HyperParallel, combining topology-aware hybrid sharding, efficient distributed optimizer execution, and coordinated management of long-sequence workloads and activation memory. CPU affinity and core isolation further reduce host scheduling interference. Together, these mechanisms support the pretraining of a 285B-parameter model on 8,192 Ascend 910C cards drawn from a large-scale production cluster, providing a concrete setting for the architectural and systems studies. Measurements show up to 4.68× prefill throughput over an MLA-only baseline, 2.67×–3.77× decoding throughput from speculative decoding, and a 16.1% throughput improvement from eliminating pipeline parallelism. For post-training, we introduce TeleGym, a unified agent execution platform that connects execution, evaluation, and trajectory construction to policy optimization. Disaggregated asynchronous execution, bounded oversampling, and partial rollout accommodate heterogeneous task durations while preserving trajectory continuity across policy updates. Together, these advances establish an Ascend-oriented training foundation for the continued development of Xing 4 across model scales and increasingly demanding reasoning and agent workloads. The open-sourced 29B member and ongoing trillion-parameter training further extend the series within the Ascend and MindSpore ecosystem.

CONTENTS

1 Introduction **2**

2 Model Architecture and Design Principles **4**

2.1 Design Objectives 5

2.2 Long-Context Attention 6

2.3 Depth-Width Trade-offs 7

2.4 Speculative Decoding 7

2.5 Representative Model Configurations 8

3 Training Infrastructure and Optimizations in HyperParallel **9**

3.1 Hybrid Sharded Data Parallelism (HSDP) 9

3.2 Owner-Based Execution for the Distributed Muon Optimizer 12

3.3 Fine-Grained Selective Recomputation 14

3.3.1 Object-Based Exclusion and Memory-Budgeted Selection 14

3.3.2 Reducing Activation Memory Inside Exclude Regions 15

3.4 DSA Long-Sequence Training Optimizations 16

3.4.1 Gradient Scheduling: Layer-Wise Indexer Backward 17

3.4.2 Q-Block Pairing and Layout Planning 17

3.4.3 KV Window: Document-Aware KV Compact 21

4 Post-Training Framework **22**

4.1 Overall Training Framework 22

4.2 TeleGym: Unified Agent Execution 23

4.3 Task Synthesis 24

4.4 Asynchronous Training with Partial Rollout 25

5 Training Cluster Performance Optimization **26**

5.1 CPU Affinity and Core Isolation 27

5.2 Training Throughput Evaluation 27

6 Conclusion **28**

A Author List **32**

1 INTRODUCTION

The TeleChat family has developed through successive advances in model capability and the infrastructure required to train at scale. TeleChat established an open foundation spanning pretraining, instruction tuning, and preference alignment (He et al., 2024; Wang et al., 2024). TeleChat2, TeleChat2.5, and T1 expanded this foundation with stronger reasoning, coding, and long-context capabilities (Wang et al., 2025). MoE paradigm has rapidly evolved into the prevailing architectural choice for scaling large language models Hu et al. (2026a;b). TeleChat3-MoE extended the family to

models with hundreds of billions to over a trillion parameters, accompanied by systematic work on numerical verification, distributed execution, and cluster performance (Liu et al., 2025b). Throughout this progression, development of the models has been supported by sustained investment in the domestic Ascend and MindSpore ecosystem.

Building on this foundation, **Xing 4** continues the TeleChat lineage as a series spanning tens to hundreds of billions of parameters, with ongoing expansion toward trillion-parameter scale. The series targets long-context reasoning, code generation, and agentic applications while advancing efficient training on Ascend. Its members adopt architectural configurations suited to their parameter budgets and execution requirements.

We have open-sourced **Xing4.0-29B-A4B**, a MoE member with 29B total parameters, 4B activated parameters per token, and native support for 256K-token contexts (XingChen-AGI, 2026). Its architecture combines Manifold-Constrained Hyper-Connections (mHC) (Xie et al., 2026) with Multi-head Latent Attention (MLA) and multi-token prediction (MTP), bringing sparse computation and long-context processing to a model intended for reasoning and agentic use.

This focus on execution efficiency reflects a broader direction in foundation-model development. Qwen3 and Tencent's Hunyuan-A13B couple model capability with control over reasoning cost and deployment efficiency (Yang et al., 2025a; Tencent Hunyuan Team et al., 2026). Kimi K3 and GLM-5 emphasize long-horizon reasoning and agentic execution (Kimi Team et al., 2026; GLM-5-Team et al., 2026), while DeepSeek-V4 makes efficient million-token attention a central architectural objective (DeepSeek-AI et al., 2026). As documents, tool outputs, and intermediate reasoning accumulate, efficient execution becomes necessary both within model calls and across post-training trajectories.

Developing Xing 4 across parameter scales requires coordinated architectural and systems design. Increasing capacity must be balanced against per-token computation and the cost of storing, communicating, and updating the full model. Long-context attention adds indexing, KV-access, and activation-memory costs; packed documents can produce unequal computational loads even when token counts match. Host scheduling delays can stall device execution, while variable task durations introduce further imbalance during reinforcement learning. These constraints motivate joint control of model-state placement, computation ownership, and work scheduling.

At the architectural level, we consider model capacity, long-context processing, and inference throughput jointly. Our studies examine compressed KV representations through MLA (DeepSeek-AI et al., 2024a), sparse access through DeepSeek Sparse Attention (DSA) (DeepSeek-AI et al., 2025), and cross-layer index reuse through IndexShare (Bai et al., 2026; Z.ai, 2026). Depth-width choices balance matrix computation against layer-wise routing and scheduling overhead. MTP provides an auxiliary training objective (DeepSeek-AI et al., 2024b), while DSpark offers block-wise speculative decoding (Cheng et al., 2026). These mechanisms address distinct execution costs and inform architecture selection under different parameter and workload budgets.

Our pretraining infrastructure is built on MindSpore HyperParallel (MindSpore Team, n.d.-a). Topology-aware hybrid sharded data parallelism (HSDP), building on fully sharded data parallelism (FSDP) (Zhao et al., 2023), aligns model-state distribution with the cluster's bandwidth hierarchy and enables execution without pipeline stages. Owner-based Muon execution assigns full-matrix orthogonalization to designated processes within supported shard domains, avoiding redundant computation while preserving sharded state (Liu et al., 2025a). For long sequences, cost-based workload assignment, reusable token layouts, and document-aware KV compaction address imbalance and memory access. Early indexer backward and selective recomputation control activation lifetimes. CPU affinity and core isolation reduce host contention that can delay task submission and collective participation.

For reasoning and agentic post-training, we introduce TeleGym, a unified execution platform for heterogeneous agent harnesses, environments, and persistent sessions. An intermediate server connects task interactions and evaluated outcomes to policy optimization through a shared trajectory representation. Task synthesis and validation reuse the same execution interfaces. Generation and optimization run asynchronously on separate device pools, with bounded oversampling accommodating variable task durations. Partial rollout allows unfinished trajectories to resume after policy updates. Sampling-time policy records and execution-level reward grouping preserve the information needed for learning as trajectories span multiple policy versions. Together, these mechanisms support mathematical

reasoning, code generation, and multi-turn agent tasks within a common reinforcement-learning workflow.

The report makes three main contributions:

- **Ascend-oriented architectural design for efficient training and inference.** We study how sparse activation, attention mechanisms, depth-width balance, and speculative decoding connect model capacity and long-context requirements to matrix utilization, memory access, and sequential execution costs. These studies make the design trade-offs behind Xing 4 configurations explicit.
- **Coordinated systems and cluster optimization for large-scale pretraining.** We develop a training infrastructure that combines topology-aware hybrid sharding and owner-based Muon execution with long-sequence workload balancing, KV compaction, and activation-memory management. CPU affinity and core isolation extend the optimization to host scheduling, coordinating device and host execution within MindSpore HyperParallel.
- **TeleGym and unified reinforcement learning for reasoning and agents.** We introduce TeleGym and connect its execution interfaces, task synthesis, and trajectory construction to a common policy-optimization pipeline. Disaggregated asynchronous execution, bounded oversampling, and partial rollout accommodate heterogeneous task durations while maintaining session continuity, sampling-time policy records, and execution-level reward grouping across policy updates.

The detailed engineering studies use a **285B-parameter MoE member** as the principal implementation example. It activates **17B parameters per token**, supports **1M-token contexts**, and is pretrained on **8,192 Ascend 910C cards**. The report combines architectural simulation, pretraining measurements, and operator-level analysis with a description of the post-training infrastructure; each evaluation specifies its model and hardware configuration.

Section 2 presents the architectural design principles and a representative configuration, and Section 3 details distributed pretraining. Section 4 describes the reinforcement-learning framework for reasoning and agentic post-training. Section 5 examines host scheduling and cluster throughput, and Section 6 concludes the report.

2 Model Architecture and Design Principles

Architectural design for Xing 4 brings model capacity, long-context execution, and hardware efficiency into a common set of trade-offs. Agentic workloads such as long-context reasoning, multi-turn tool use, complex task planning, and high-concurrency inference (Zhao et al., 2025; Zhang et al., 2024; Zou et al., 2026; Bian et al., 2026; Li et al., 2026; Xiong et al., 2025) make these requirements interdependent: longer histories increase attention cost, while repeated model calls expose decoding latency and sequential overhead.

Architectural choices vary with parameter budget. For example, the released 29B member combines mHC, which modifies residual connectivity (Xie et al., 2026), with MLA and MTP (XingChen-AGI, 2026). The studies below examine attention, depth-width balance, and speculative decoding as complementary choices for matching model configurations to their execution requirements.

As a concrete implementation, Figure 1 illustrates the 285B configuration. Its decoder-only MoE backbone combines MLA (DeepSeek-AI et al., 2024a), DSA (DeepSeek-AI et al., 2025), and IndexShare (Bai et al., 2026; Z.ai, 2026). The first 3 layers use Dense FFNs, while the remaining layers use MoE FFNs to expand capacity through sparse activation. MLA compresses KV representations; DSA and IndexShare reduce historical KV access and repeated index computation. DSpark supplies an inference-time speculative decoding component, using a block-wise draft model, a lightweight Markov head, and an acceptance confidence head to increase the number of tokens committed per verifier invocation.

From a hardware perspective, Ascend 910C provides strong Cube matrix-computation capability, while Vector compute and HBM bandwidth are relatively more constrained. These hardware characteristics guide our architectural choices:

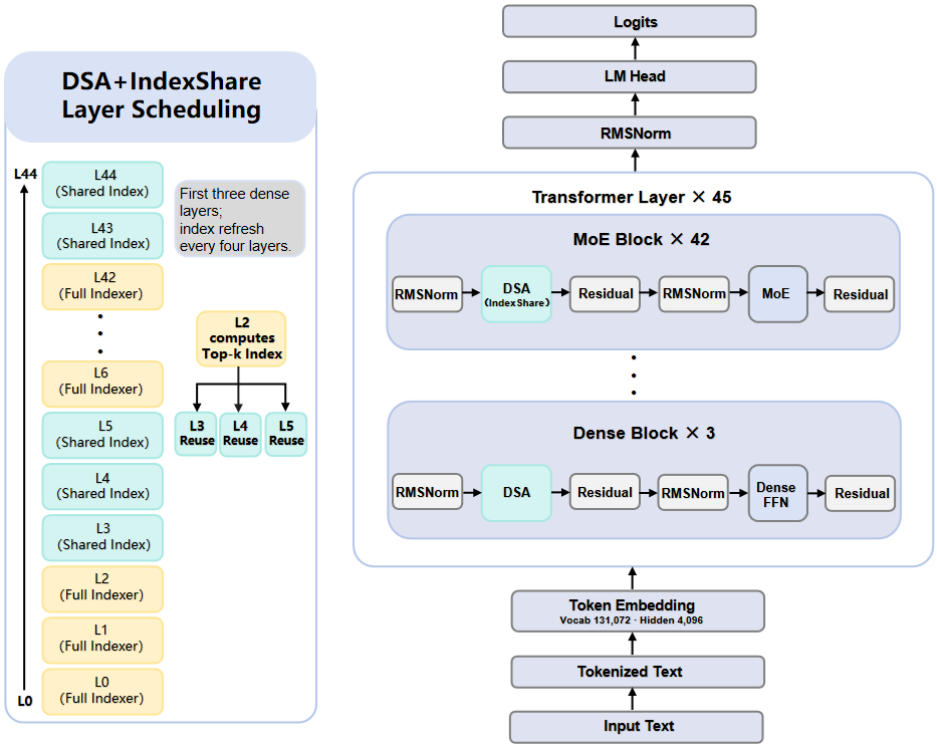


Figure 1: Architecture overview of the 285B representative model in the Xing 4 series.

Leverage Cube with matrix-dominant computation, mitigate Vector bottlenecks through structural optimization, and reduce HBM pressure through sparse memory access.

2.1 Design Objectives

We select model configurations by balancing total capacity, per-token computation, and execution cost on Ascend 910C. Parameter budgets determine the available design space, while long-context requirements and hardware characteristics guide the choice of attention, expert configuration, and network depth. The aim is to preserve sufficient capacity while improving long-context deployability and end-to-end execution efficiency.

The architecture mainly targets two bottlenecks:

1. **Long-context inference bottleneck**: MLA reduces KV Cache size, but in ultra-long-context scenarios, the main attention still needs to access a large number of historical KV states. As a result, decode latency increases with context length.
2. **Ascend hardware efficiency bottleneck**: Overly deep networks amplify non-matrix-computation overheads, including normalization, routing, indexing, scheduling, and data movement, which reduces effective Cube utilization on Ascend 910C.

Based on these observations, we organize the architectural studies around three choices:

- **Attention design**: combining MLA with DSA and IndexShare reduces historical KV access and repeated index computation in long-context inference.
- **Depth-width trade-off**: balancing layer count and expert width changes matrix-computation density and sequential overhead; controlled comparisons identify configurations suited to Ascend 910C.

- **Speculative decoding**: block-wise drafting with **DSpark** combines a draft model, a lightweight Markov head, and an acceptance confidence head to increase the number of accepted tokens per verifier invocation and improve output throughput.

Training objectives and inference-time decoding are configured separately. In the representative implementation, MTP remains an auxiliary training objective, and DSpark is paired with the main-model verifier for inference. Muon is applied to matrix-valued parameters during training; its distributed execution is discussed in Section 3.

2.2 Long-Context Attention

Long-context attention requires control over both KV storage and the work performed for each query. We examine MLA as the baseline and add DSA and IndexShare to address historical KV access and repeated indexing, yielding a **MLA + DSA + IndexShare** design.

MLA compresses KV representations and reduces KV-cache memory storage requirement. DSA uses a lightning indexer to select Top-K historical KV states for each query, so that the main attention only attends to a sparse subset of historical context. IndexShare further reuses sparse indices across layers, reducing repeated indexer computation in DSA and making sparse attention more suitable for long-context and high-concurrency decode workloads.

To evaluate these choices using attention variants of the 285B model on Ascend 910C, we primarily use simulation-based performance modeling based on Huawei's Zhanlu simulation framework. The simulation models the major computation, memory-access, and communication characteristics of the target Ascend 910C deployment, and is used to compare relative architectural trends under controlled conditions. Since simulation does not fully capture all production-serving overheads, the reported results may be more optimistic than full end-to-end deployment measurements. The comparison covers three variants: MLA, MLA+DSA, and MLA+DSA+IndexShare, under a prefill/decode(P/D) disaggregation deployment. On the prefill side, we focus on long-prompt throughput and time to first token(TTFT); on the decode side, we focus on time per output token(TPOT) and concurrent decode throughput.

Table 1: Inference evaluation setup for attention variants of the 285B model.

Item	Configuration
Inference environment	32 Ascend 910C cards
Deployment mode	P/D disaggregation
Prefill-side parallelism	TP4DP4
Decode-side parallelism	TP2DP8
Evaluation context length	256K

Table 2 compares the three variants. All results are normalized against the MLA baseline.

Table 2: Simulation-based inference efficiency of attention variants in the 285B study. All metrics are normalized against the MLA baseline. For access and latency metrics, lower is better; for throughput metrics, higher is better.

Architecture	KV HBM Access	Index Computation	Relative TTFT	Relative TPOT	Prefill Throughput	Decode Throughput
MLA	1.00×	None	1.00×	1.00×	1.00×	1.00×
MLA + DSA	0.85×	Per-DSA-layer	0.37×	0.84×	2.67×	1.18×
MLA + DSA + IndexShare	0.80×	Cross-layer	0.21×	0.79×	4.68×	1.27×

The results show that MLA + DSA substantially reduces TTFT and TPOT while increasing throughput in long-context scenarios. Adding IndexShare further reduces repeated index computation and provides additional performance gains. These results support the selection of **MLA + DSA + IndexShare** for the representative configuration.

2.3 Depth-Width Trade-offs

Model depth and expert width jointly determine the balance between matrix computation and layer-wise sequential work. In Xing 4 architecture selection, we consider the number of layers, hidden size, expert intermediate size, and the number of activated experts per token together, seeking configurations that fit both the parameter budget and Ascend 910C execution characteristics.

Deeper networks increase layer-wise sequential overhead, especially in non-matrix-computation components such as normalization, routing, and indexing. In contrast, a wider-and-shallower topology increases the scale of per-layer matrix multiplications, improves Cube utilization, and reduces cross-layer scheduling and communication overhead.

The released 29B member provides one concrete configuration: 40 layers, 64 routed experts with four selected per token, and one shared expert (XingChen-AGI, 2026). At larger parameter budgets, depth and expert width must be reconsidered together. The following controlled study compares candidate topologies for the 285B model while holding the main architectural components fixed.

Our depth-width study compares a baseline topology with a wider, shallower alternative. Both configurations share the same base architecture, consisting of MLA, DSA, IndexShare, a hybrid Dense/MoE backbone, and an MTP objective. The main differences are the number of layers, expert intermediate size, and the number of activated experts per token.

Table 3: Candidate configurations for the representative model's depth-width study.

Configuration	Layers	Hidden Size	Heads	Routed Experts	Expert Intermediate Size	Experts per Token	Shared Experts
Baseline	60	4096	128	256	1536	6	1
Wider-shallower	45	4096	128	256	2048	8	1

The inference-side comparison focuses on decode latency and throughput under long-context settings. Results are normalized against the baseline configuration.

Table 4: Normalized inference efficiency in the 285B depth-width study.

Configuration	Relative TTFT	Relative TPOT	Prefill Throughput	Decode Throughput	Decision
Baseline	1.00×	1.00×	1.00×	1.00×	Reference
Wider-shallower	0.83×	0.90×	1.20×	1.11×	Selected

Simulation and real-system measurements show that, under comparable total and activated parameter budgets, the 45-layer wider-FFN topology provides better inference latency and throughput than the 60-layer narrower-FFN topology. Large-scale training measurements also show that the wider-and-shallower topology improves overall training throughput by approximately 10% compared with the deeper baseline.

We therefore select the 45-layer wider-FFN topology. This design preserves model scale and quality while reducing cross-layer sequential overhead and increasing the proportion of matrix-dominant computation on Ascend 910C, making it more suitable for long-context, high-concurrency inference and large-scale training.

2.4 Speculative Decoding

We examine **DSpark** (Cheng et al., 2026), a speculative-decoding approach that combines a dedicated block-wise draft model, a lightweight Markov head, and an acceptance confidence head. This architecture aims to improve decoding throughput by increasing the number of tokens committed per verifier invocation while keeping drafting overhead low.

DSpark conditions a lightweight draft model on intermediate representations from multiple verifier layers, combining information from different depths to support block-wise token prediction. Its five-layer **DFlash** Chen et al. (2026) backbone computes representations for seven speculative positions in parallel, while the lightweight Markov head captures dependencies between adjacent proposed tokens through logit corrections. By consolidating drafting into parallel block-wise computation, **DSpark** better exploits Ascend 910C's strength in matrix-dominant computation. The confidence head estimates per-position acceptance probabilities using draft representations and previous-token information. The verifier evaluates the proposed block and determines the accepted prefix. Together, parallel drafting and lightweight sequential correction enable multiple tokens to be committed per verifier invocation, significantly improving decoding efficiency.

Table 5 compares the decoding throughput of the base configuration, **MTP3**, and **DSpark** on six mathematical-reasoning and code-generation benchmarks. **MTP3** achieves 1.80×–1.93× the throughput of the base configuration, while **DSpark** achieves 2.67×–3.77×. Across all six benchmarks, **DSpark** consistently outperforms **MTP3**, delivering 1.48×–2.10× its output throughput, with the largest gain observed on HumanEval. These results demonstrate the decoding efficiency advantage of **DSpark** over **MTP3** and support its adoption as the speculative decoding component in the representative configuration.

Table 5: Decoding throughput of the 285B model with the base configuration, MTP3, and DSpark. Throughput values are listed in the order Base / MTP3 / DSpark; ratios represent output throughput speedups.

Benchmark	Output Throughput (tokens/s)	MTP3/Base	DSpark/Base	DSpark/MTP3
GSM8K	35.4 / 63.7 / 115.8	1.80×	3.28×	1.82×
MATH-500	36.8 / 70.0 / 128.1	1.90×	3.49×	1.83×
HumanEval	35.7 / 64.1 / 134.7	1.80×	3.77×	2.10×
MBPP	37.1 / 70.0 / 132.0	1.89×	3.56×	1.89×
AIME25	37.0 / 71.5 / 109.2	1.93×	2.95×	1.53×
LiveCodeBench	35.6 / 64.5 / 95.2	1.81×	2.67×	1.48×

2.5 Representative Model Configurations

Table 6 summarizes two members of the Xing 4 series with different parameter budgets and architectural configurations. The 285B member serves as the principal implementation for the architectural and training-system studies in this report.

Table 6: Representative configurations of the Xing 4 series. The per-token expert count includes selected routed experts only; shared experts are listed separately.

Model	Layers	Hidden Size	Routed Experts	Expert Intermediate Size	Experts per Token	Shared Experts
Xing 4 (29B)	40	3584	64	1024	4	1
Xing 4 (285B)	45	4096	256	2048	8	1

Xing 4 (29B) corresponds to the open-sourced Xing4.0-29B-A4B. It activates 4B parameters per token, natively supports 256K-token contexts, and combines mHC with MLA and MTP (XingChen-AGI, 2026).

The 285B member activates approximately 17B parameters per token and supports contexts of up to one million tokens. Its 45-layer backbone uses Dense FFNs in the first 3 layers and MoE FFNs in the remaining layers. Attention combines MLA, DSA, and IndexShare; DSpark provides inference-time speculative decoding. MTP is retained as an auxiliary training objective, and Muon is adopted for optimization. This configuration satisfies the design budgets of fewer than 300B total parameters and fewer than 20B activated parameters.

3 TRAINING INFRASTRUCTURE AND OPTIMIZATIONS IN HYPERPARALLEL

The pretraining infrastructure for Xing 4 coordinates parallelization layout, optimizer execution, activation memory, and sparse-attention kernels on Ascend SuperPoD clusters. Built on MindSpore HyperParallel, the stack follows a two-tier network topology: high-bandwidth intra-SuperPoD links carry fine-grained weight gathering and sharded-gradient collectives, whereas coarse-grained gradient reduction crosses SuperPoD boundaries. This layout relieves memory pressure in the data-parallel dimension and enables execution without pipeline stages. The four subsections below address topology-aware sharded collectives, redundant full-matrix work in distributed Muon, activation residency under selective recomputation, and the computational cost of DSA sparse-attention inputs. The performance studies use the representative MoE configuration on a 1,024-NPU Ascend profiling cluster; production pretraining of the 285B member subsequently uses 8,192 Ascend 910C cards. Reported step times and operator-level breakdowns retain their profiling-scale scope. Relative to the HSDP baseline of $31.8\,\mathrm{s}$, owner-based Muon execution, fine-grained selective recomputation, and layer-wise indexer backward reduce the step time by 9.1%, 6.9%, and 15.2%, each measured independently against the same HSDP + Muon configuration.

3.1 HYBRID SHARDED DATA PARALLELISM (HSDP)

Training mixture-of-experts (MoE) models with hundreds of billions of parameters typically combines data, tensor, expert, and pipeline parallelism (DP/TP/EP/PP). While pipeline parallelism (PP) relieves memory pressure by partitioning model layers across stages, it inevitably introduces a structural pipeline bubble. For instance, with $p = 4$ pipeline stages and $m = 8$ micro-batches, the theoretical bubble fraction of the 1F1B schedule is $\frac{p-1}{m+p-1} = \frac{3}{11} \approx 27\%$. Although interleaved schedules, such as virtual pipeline parallelism (VPP), can mitigate this issue, they cannot eliminate the pipeline bubble entirely.

The primary benefit of PP lies in its layer-wise memory partitioning, which circumvents per-device memory bottlenecks. If the data-parallel dimension could instead provide this memory relief—ensuring that no single accelerator needs to keep the full parameter tensors of all layers resident simultaneously—PP would become unnecessary, thereby eliminating the pipeline bubble altogether.

Fully Sharded Data Parallel (FSDP) Zhao et al. (2023) adopts this philosophy. Model parameters are stored as distributed shards. Before each layer is executed, an AllGather operation reconstructs its full parameters; after the computation, the parameters are resharded and the temporary full tensors are released. This approach enables large-scale model training with minimal parameter redundancy.

However, when FSDP is scaled to clusters containing thousands or tens of thousands of NPUs, the sharding group may grow proportionally with the cluster size. Consequently, frequent AllGather and ReduceScatter operations must cross SuperPoD boundaries. These fine-grained collectives incur high communication latency, underutilize the relatively limited inter-SuperPoD bandwidth, and are difficult to overlap completely with computation. Moreover, the additional memory savings obtained from an excessively large sharding degree exhibit diminishing returns.

To address these challenges, we design and implement **HyperParallel-FSDP** Sun et al. (2026); MindSpore Team (n.d.-b), a hybrid sharded data-parallel (HSDP) scheme tailored to the Ascend SuperPoD architecture. Model parameters are FSDP-sharded within each SuperPoD and replicated across SuperPoDs. Frequent weight-gathering and sharded-gradient synchronization operations are therefore confined to the high-bandwidth intra-SuperPoD communication domain, while only coarse-grained gradient reduction traffic crosses SuperPoD boundaries. This design balances memory efficiency and communication overhead while enabling the removal of pipeline parallelism. The backward-pass workflow of HyperParallel-FSDP is illustrated in Figure 2.

The Ascend SuperPoD provides ultra-high intra-SuperPoD bandwidth and ultra-low intra-SuperPoD communication latency. HyperParallel-FSDP introduces two key designs that exploit these architectural properties.

The first design targets the additional data movement introduced by parameter flattening. To improve collective communication efficiency, PyTorch FSDP Zhao et al. (2023) flattens multiple parameters into a single contiguous buffer and then performs a fused collective operation, as shown in Figure 3(1). Although this design reduces the number of collective operations, it introduces additional data

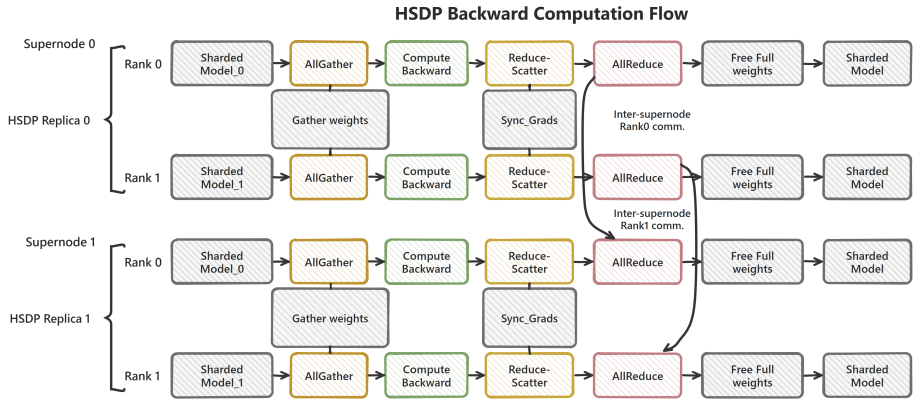


Figure 2: Backward-pass workflow of HyperParallel-FSDP.

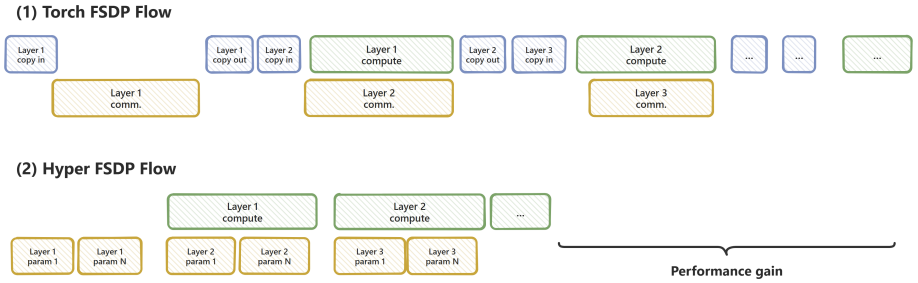


Figure 3: (1) Fused communication in PyTorch FSDP and (2) zero-copy communication in HyperParallel-FSDP.

movement. Before communication, individual parameter tensors must be copied into the flattened communication buffer (*copy-in*). After AllGather, the reconstructed parameters must be copied from the communication buffer back to their corresponding parameter tensors (*copy-out*). Copy-out lies on the critical execution path and cannot be overlapped with subsequent computation because the computation directly depends on the gathered parameters.

The Ascend SuperPoD provides sufficiently high intra-SuperPoD bandwidth for fine-grained collective operations to retain high communication efficiency. HyperParallel-FSDP therefore performs AllGather and ReduceScatter directly on individual parameter tensors, as illustrated in Figure 3(2). This per-parameter communication design removes the flattened intermediate buffer and eliminates the associated parameter packing, copy-in, and copy-out operations. Consequently, HyperParallel-FSDP achieves **zero-copy communication** while preserving efficient intra-SuperPoD collective communication.

The second design fuses inter-SuperPoD gradient reduction and overlaps multiple communication operations with backward computation. During backpropagation, the execution of each layer involves four major operations: intra-SuperPoD AllGather for parameter reconstruction, backward computation, intra-SuperPoD ReduceScatter for sharded-gradient synchronization, and inter-SuperPoD AllReduce for gradient reduction across replicas.

Figure 4(1) shows the forward-pass execution timeline. Before the forward computation of each layer, the required parameter shards are reconstructed through an intra-SuperPoD AllGather. Once the full parameters become available, the corresponding layer computation can proceed.

A straightforward backward implementation is shown in Figure 4(2). This baseline schedule overlaps part of the inter-SuperPoD AllReduce communication with backward computation, but it has two limitations. First, frequent fine-grained AllReduce operations are inefficient because inter-SuperPoD bandwidth is more limited than intra-SuperPoD bandwidth. Second, intra-SuperPoD ReduceScatter

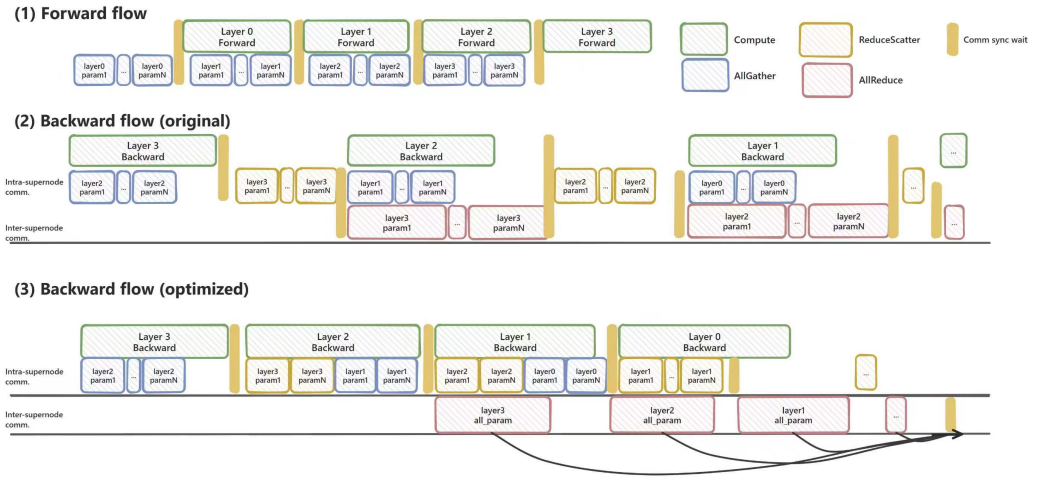


Figure 4: Execution timelines of HyperParallel-FSDP: (1) forward pass, (2) baseline backward pass, and (3) optimized backward pass with three-way communication overlap.

and communication synchronization remain partially exposed on the critical execution path, resulting in visible communication gaps between consecutive layer computations.

HyperParallel-FSDP addresses these limitations through three complementary techniques:

1. Intra-SuperPoD AllGather and ReduceScatter operations are asynchronously scheduled and overlapped with backward computation.
2. Inter-SuperPoD gradient reduction proceeds asynchronously once the corresponding gradient data become available. Synchronization is deferred until the end of the backward pass, allowing its latency to be hidden behind the backward computation of multiple layers.
3. Gradient shards participating in inter-SuperPoD reduction are aggregated and communicated through fused AllReduce operations, reducing communication frequency and collective-launch overhead.

To satisfy the contiguous-memory requirement of fused AllReduce, HyperParallel-FSDP pre-allocates a contiguous gradient buffer and configures each ReduceScatter operation to write its output directly into the corresponding region of this buffer. The gradient shards are therefore contiguous by construction, so the subsequent fused AllReduce requires no additional packing or memory-copy operations.

The optimized schedule is shown in Figure 4(3). During the backward computation of layer i, the system can simultaneously overlap the ReduceScatter of layer $i+1$, the AllGather of layer $i-1$, and inter-SuperPoD AllReduce communication. Communication synchronization is postponed until the end of the backward pass. This arrangement achieves **three-way communication overlap**, substantially reducing exposed communication latency.

To evaluate the effectiveness of the proposed design, we compare conventional pipeline-parallel configurations with HyperParallel-FSDP on the 1,024-NPU profiling cluster, using a 285-billion-parameter MoE model. The sequence length is set to 262,144, and the global batch size is 16. Tensor parallelism (TP), expert parallelism (EP), and context parallelism (CP) are fixed at 4, 64, and 32, respectively. Table 7 summarizes the experimental results.

Table 7: Training performance of the 285B model on the 1,024-NPU profiling cluster.

Scheme	Pipeline	DP (Replica × Shard)	mbs	Step Time	Speedup
PP-1F1B (Baseline)	$pp = 4,\ vpp = 1$	2×1	8	36.4 s	1.00×
PP-Interleaved	$pp = 4,\ vpp = 2$	2×1	4	34.2 s	1.06×
HSDP (**Ours**)	$pp = 1$	4×2	2	31.8 s	1.14×

Note: mbs denotes the micro-batch size.

The first two configurations employ $PP = 4$. Changing the execution schedule from 1F1B to interleaved pipeline parallelism ($vpp = 2$) reduces the step time from $36.4\,\mathrm{s}$ to $34.2\,\mathrm{s}$, a 6.0% reduction in step time ($1.06\times$). This result indicates that the pipeline bubble is a major structural bottleneck in this configuration. Although schedule interleaving reduces the bubble, it cannot eliminate it completely.

In contrast, the HSDP configuration eliminates pipeline parallelism by setting $PP = 1$ and instead relieves memory pressure through hybrid sharding in the data-parallel dimension, with $dp_replicate = 4$ and $dp_shard = 2$. This configuration reduces the step time to $31.8\,\mathrm{s}$, a 12.6% reduction and a $1.14\times$ end-to-end speedup over the PP-1F1B baseline.

These results demonstrate that zero-copy communication and three-way communication overlap sufficiently suppress the communication overhead introduced by HSDP, enabling large-scale training without pipeline parallelism. By removing PP, HyperParallel-FSDP eliminates the pipeline bubble entirely and outperforms conventional PP-based memory-relief strategies.

3.2 Owner-Based Execution for the Distributed Muon Optimizer

Muon orthogonalizes matrix-valued parameters with a Newton–Schulz (NS) iteration, which requires the full weight matrix. This conflicts with sharded training, where each rank holds only a slice of the parameters and optimizer state. In a straightforward distributed implementation, every rank that owns a slice of a two-dimensional weight updates its local momentum, reconstructs the full NS input, runs the same NS iteration, and keeps only the part of the resulting update that belongs to its local shard. The NS computation is therefore duplicated on every participating rank even though each rank uses only a fraction of the result. The redundancy lies in who performs the full-matrix step, not in how the state is sharded.

We remove this duplication by assigning each matrix a single owner rank within its shard group. Non-owner ranks send their input shards to the owner over point-to-point (P2P) transfers. The owner assembles the full input, runs NS, scales the update, keeps its own slice, and scatters the remaining slices back. Parameters and momentum stay in their original sharded layout; only the full-matrix computation moves to the owner (Figure 5). Ownership is per matrix rather than global, so replica-local shard groups can use different owners. Layouts that cannot use this P2P path fall back to the original execution path.

The scheduler combines owner assignment with batching. It groups matrices by shape, shard group, and estimated cost, then distributes the groups across eligible owners. Full inputs of compatible shapes can be stacked along a batch dimension and processed with batched matrix multiplication, which amortizes kernel-launch overhead. Each matrix is still normalized and orthogonalized independently; batching never concatenates different weights into one larger NS problem. The trade-off is that full inputs and NS work concentrate on the owners, so the achievable benefit depends on how well the assignment balances cost and on the distribution of matrix shapes.

Communication is issued in batches, and waits are placed only where the data is actually needed (Figure 6). Each rank prepares its local inputs, writes momentum back, and submits the gathers. Work that does not need the full matrix — for example NS on three-dimensional or purely local parameters — can run while the owner waits for its inputs. In the batched path, readiness is tracked with shared communication handles, so the wait is not necessarily per matrix. After the owner computation, a rank drains pending gathers and submits the scatter; local or three-dimensional parameter updates can follow immediately, while two-dimensional updates wait for the scatter to complete. These are per-rank data dependencies, not a global barrier. The ordering exposes overlap opportunities, but the overlap actually achieved depends on the communication backend, the workload, and how much independent work is available.

Temporary memory is kept in check by grouping work per phase and releasing references as soon as they are no longer needed. Preparation and parameter-application tasks are split according to estimated temporary-memory budgets, with a separate budget for local/expert NS groups. These budgets are heuristics rather than a global memory bound: the local NS budget does not cover gathered two-dimensional inputs or owner NS work. Momentum can be written back before preparation buffers

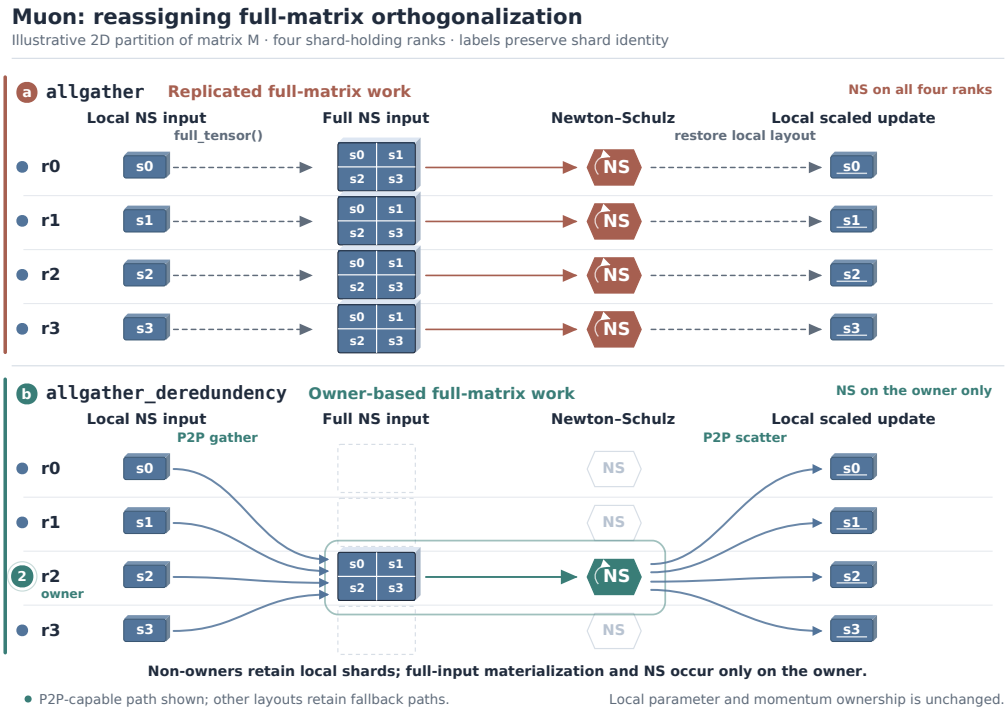


Figure 5: Reassignment of full-matrix Muon work. Four ranks hold illustrative shards of one matrix, with rank r2 selected as its owner. The baseline reconstructs and orthogonalizes the input on each participating rank. The supported P2P path reconstructs it at the owner and returns scaled-update shards to their original destinations. The owner's own input and output slices remain local. Faint outlines indicate omitted full-input materialization and NS, not missing local state. The shard arrangement is illustrative; arrows do not encode communication volume.

are freed, and consumed NS inputs can be released before the parameter update. Shorter reference lifetimes reduce peak usage, although the runtime may still retain buffers through outstanding aliases.

Compatibility rules and model-specific split/merge settings still decide which execution path is taken; in particular, compatibility mode disables cross-parameter NS batching. Reordering the execution and using batched kernels can change floating-point results, so precision settings remain part of the numerical configuration. Keeping per-matrix orthogonalization domains is therefore a structural property, not a guarantee of bitwise equivalence or unchanged convergence. Runtime, communication volume, memory, and numerical behavior must be measured under the actual training configuration.

We isolate the effect of owner-based execution by keeping the HSDP layout of Table 7 fixed and changing only the optimizer execution path (Table 8). The step time falls from 31.8 s to 28.9 s, a 9.1% reduction that corresponds to a 1.10× end-to-end speedup. Because the parameter and momentum layout is unchanged, the difference is attributable to the removal of duplicated full-matrix NS work rather than to a change in the distributed state.

Table 8: Effect of owner-based Muon execution under a fixed HSDP layout (1,024 NPUs, sequence length 262,144).

Configuration	Step Time	Speedup	Step-time reduction
HSDP	31.8 s	1.00×	—
HSDP + owner-based Muon	28.9 s	1.10×	9.1%

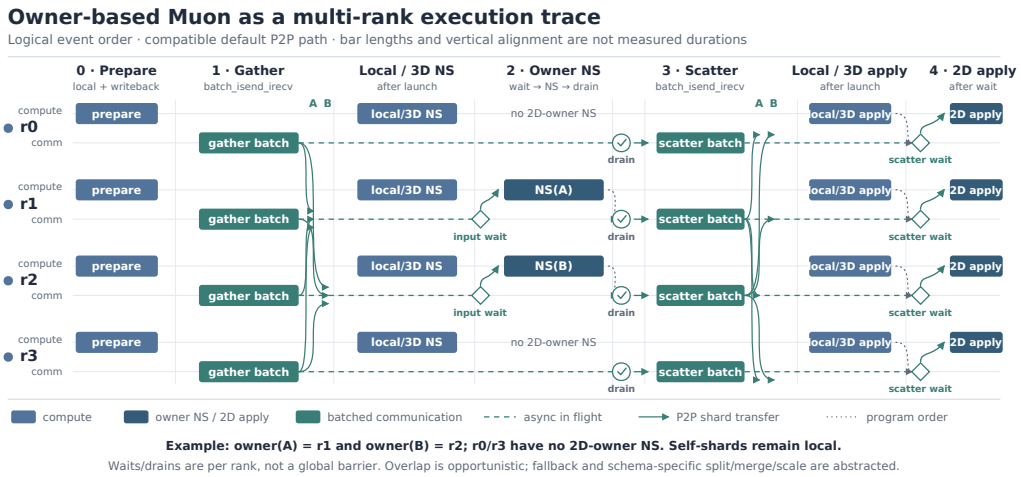


Figure 6: Logical execution order for a compatible owner-based P2P path. Matrix A is assigned to r1 and B to r2. Compute and communication lanes expose batched submission, input waits, gather draining, and scatter completion dependencies. Self-shards remain local. Local/3D work is shown where available; fallback paths and model-specific transformations are omitted. Bar lengths, dashed spans, and vertical alignment are schematic and do not represent measured durations or guaranteed overlap.

3.3 Fine-Grained Selective Recomputation

Layer-wise recomputation offers only a coarse memory–compute trade-off: saving or replaying an entire layer treats all of its operators alike. This is particularly inefficient for long-sequence DSA, where some operators are expensive to execute but produce relatively small outputs. Replaying these operators incurs substantial computation for limited activation-memory savings. Native operator-level selective checkpointing provides finer control through a dispatch mode, but every operator dispatch invokes a Python policy callback. This adds host work throughout the forward pass and can make execution host-bound when device kernels are short. Policies based on operator type also make it cumbersome to mark repeated instances of the same operator differently within a layer.

3.3.1 Object-Based Exclusion and Memory-Budgeted Selection

Our *exclude* scheme provides operator-level recomputation control through named model objects. An operator is assigned as an object member during initialization, and its member path can then be configured to select whether it is recomputed. Separate members distinguish different uses of the same operator type. Expensive operators are placed in exclude regions, while the surrounding computation remains eligible for replay. This does not require a framework-wide dispatch policy or a Python policy callback at every operator dispatch, preserving fine control while reducing host-side callback overhead.

The forward pass has two kinds of activation regions, illustrated in Figure 7. In a recompute region, activations that would otherwise be saved for backward are replaced by placeholders. In an exclude (save) region, the corresponding activations and operator outputs are retained. During backward replay, the former activations are regenerated and the excluded operators reuse their saved outputs. The choice therefore avoids replaying selected expensive operators without retaining every intermediate activation in their surrounding layer.

We choose excluded operators from profiling data. For each candidate operator i we record its forward execution time t_i together with its output shapes and data types, from which we estimate the output-memory cost m_i. Let M be the activation-memory budget remaining after the other training allocations, and let $x_i = 1$ mean that the operator output is retained and the operator is excluded from

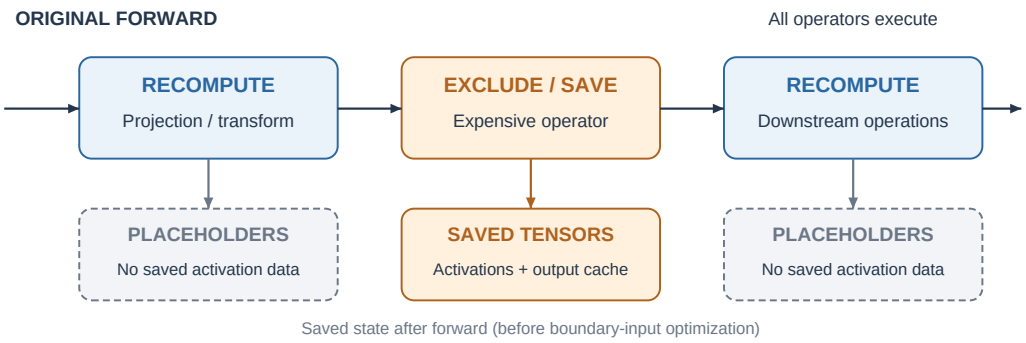


Figure 7: Activation handling in the original forward pass. Blue regions replace saved activations with placeholders; the orange exclude region retains activations and its output. Both regions execute normally in the forward pass.

replay. Selection is then a 0–1 knapsack problem:

$$\max_{x_i \in \{0,1\}} \sum_i x_i t_i \qquad \text{subject to} \qquad \sum_i x_i m_i \le M. \tag{1}$$

Retaining an output costs m_i bytes and removes t_i of replay computation; an operator that is not selected stays eligible for replay. The objective therefore spends a bounded amount of activation memory on the operators that are most expensive to execute, instead of applying one policy to an entire layer.

Table 9 applies this measurement to a single forward Transformer layer of the 256K-token run on rank 0, which issues 150 compute kernels summing to 115.1 ms. Execution time and output footprint are poorly correlated within the layer: the costliest kernel takes 38.32 ms and retains 193 MiB, whereas the largest output in the layer occupies 516 MiB yet costs only 13.47 ms. Ranking candidates by output size alone would therefore select the wrong operators, which is why the selection is formulated over measured time and measured footprint jointly.

Table 9: Compute kernels of one forward Transformer layer on rank 0 (256K-token run), aggregated by kernel name. Output size is the logical total derived from the output shapes and data types, not peak device memory.

Kernel	Duration (ms)	Output size (MiB)
SparseLightningIndexerGradKLLoss	38.32	193.0
LightningIndexer	21.51	48.0
AllGatherMatmul	15.90	656.0
SparseFlashAttention	13.47	516.0
MatmulReduceScatter	6.00	64.0

Solving Eq. 1 on this profile under the available activation budget yields an operator-level recomputation strategy: it retains the outputs that avoid the most replay time per unit of memory and leaves the remaining operators eligible for replay. The retained footprint of an excluded operator is smaller than its nominal output size, because the boundary treatment described in Section 3.3.2 releases inputs that would otherwise stay resident inside the save region.

Applying this strategy reduces the step time from 28.9 s to 26.9 s on top of the HSDP + Muon configuration, an end-to-end optimisation of 2.0 s **per step** (6.9%, 1.07×; Table 10).

3.3.2 Reducing Activation Memory Inside Exclude Regions

Saving an expensive operator's output does not by itself bound the memory retained inside its exclude region. The operator may also save large inputs for backward, and many of those inputs are outputs of the upstream recompute region. Keeping them alive inside the save region defeats the intended

Table 10: Effect of fine-grained selective recomputation, evaluated against the HSDP + Muon configuration (1,024 NPUs, sequence length 262,144).

Configuration	Step Time	Speedup	Step-time reduction
HSDP + Muon	28.9 s	1.00×	—
HSDP + Muon + selective recomputation	26.9 s	1.07×	6.9%

release of upstream activations. For example, the profiled SFA query has shape $16{,}384 \times 32 \times 512$ in BF16 and occupies 512 MiB; its positional-query input adds another 64 MiB. The DSA indexer's query projection has shape $16{,}384 \times 64 \times 128$ in BF16 and occupies 256 MiB. Retaining such inputs across many layers can consume substantial memory even when only a few operators are excluded.

We optimize the recompute-to-save boundary by recognizing the storage addresses of the activations crossing it. When the save region attempts to retain a tensor with a matching address, we store a placeholder instead of keeping the activation resident. Activations produced within the exclude region and the outputs needed to skip its execution are still saved. When backward triggers upstream recomputation, the newly generated boundary tensors restore the corresponding activations inside the save region. Figure 8 shows this process: the large boundary activation is released after forward and restored when needed, while the expensive excluded operator is not executed again. This removes the persistent boundary-input storage without adding another recomputation pass.

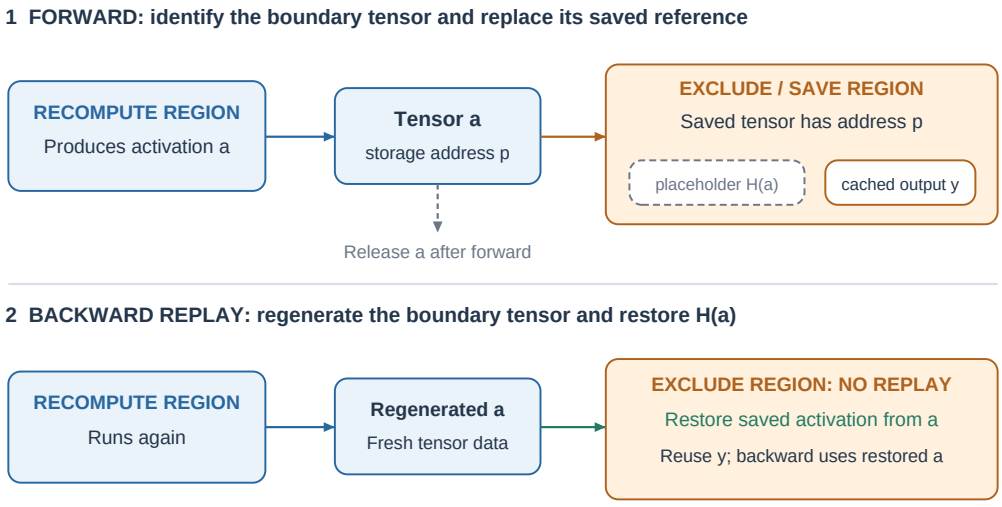


Figure 8: Boundary-activation memory optimization. Forward identifies a tensor crossing from recompute to save and replaces its saved reference with a placeholder when the storage address matches. Backward replay regenerates the boundary tensor and restores the save region's activation; the excluded operator reuses its cached output.

3.4 DSA Long-Sequence Training Optimizations

DeepSeek Sparse Attention (DSA) uses a trainable indexer to select Top-K keys before performing sparse attention. For packed TND long-sequence training, an equal number of tokens does not imply an equal amount of computation: documents excluded by the mask may still occupy physical KV inputs, and auxiliary indexer graphs may remain live for a long time. We optimize the MindFormers/HyperParallel implementation along three boundaries: gradient scheduling, Q-block pairing, and the KV window, which structure the three subsections below.

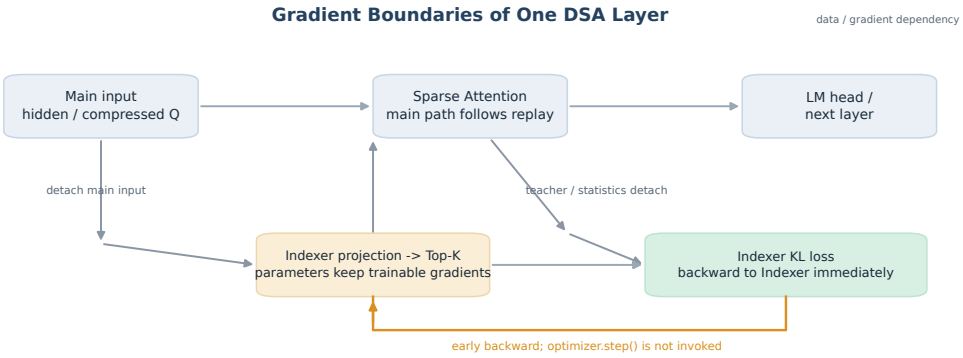


Figure 9: Indexer auxiliary-gradient boundaries: main-branch inputs and teacher statistics are isolated, the auxiliary KL loss is backpropagated early, and the optimizer still updates at the end of the step.

3.4.1 Gradient Scheduling: Layer-Wise Indexer Backward

The fused KL operator can return, already in the forward pass, a precomputed gradient with respect to the indexer projection output; an explicit backward then propagates it to the indexer projection parameters, and the auxiliary loss is not backpropagated along the teacher or the main-branch inputs. Without `IndexShare`, the auxiliary graph can be consumed layer by layer; when a shared graph exists, it is aggregated by `IndexShare` group and backpropagated once, avoiding repeated traversal of the shared graph. When main-branch recomputation is enabled, turning on layer-wise backward makes the framework automatically exclude the Indexer and `IndexerLoss` from recomputation and retain the inputs required by the early backward, so users do not have to set `exclude_op` for these two modules manually. Auxiliary activations therefore need not be retained until the main-branch LM backward finishes, which reduces duplicated computation and shortens their lifetime. However, Top-K outputs, the real inputs required by the early backward, and projections shared within a group may still occupy device memory, so this must not be described as a zero-memory-cost design.

Gradients are still accumulated according to the original gradient-accumulation (GAS) and loss-scaling rules, and the optimizer still updates at the end of the step. FSDP, recomputation input saving, and cache release cooperate to keep the early backward correct.

Table 11 measures the end-to-end effect against the same HSDP + Muon baseline used in Section 3.3. The step time falls from 28.9 s to 24.5 s, a 15.2% reduction and a 1.18× end-to-end speedup. This is the largest single-feature gain among the three optimizations, which is consistent with the mechanism: the indexer backward no longer waits for the main-branch language-model backward, so the auxiliary graph is retired one layer at a time instead of being kept alive to the end of the step.

Table 11: Effect of layer-wise indexer backward, evaluated against the HSDP + Muon configuration (1,024 NPUs, sequence length 262,144).

Configuration	Step Time	Speedup	Step-time reduction
HSDP + Muon	28.9 s	1.00×	—
HSDP + Muon + layer-wise indexer backward	24.5 s	1.18×	15.2%

3.4.2 Q-Block Pairing and Layout Planning

Q-side load balancing is applied at three levels: within a single sequence, across the layers of one replica, and across data-parallel replicas.

Under a lower-triangular causal mask, later queries can generally attend to more keys, so equal token counts do not necessarily correspond to equal computation. With a fixed Q-block size, head count, and feature dimension, the dominant theoretical workload of operators that scan without an index — LightningIndexer (which scans candidate keys before Top-K selection), FA, and the dense loss — grows approximately linearly with the number of effectively visible keys. This is why the causal pair

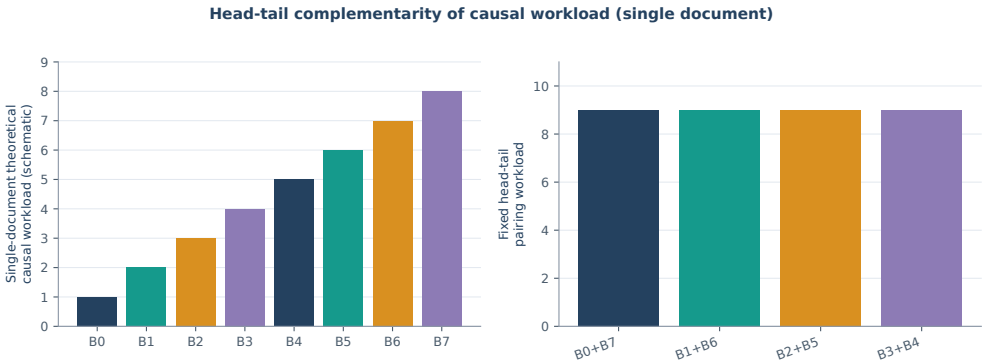


Figure 10: Head–tail complementarity of the theoretical causal workload for a single document. Values are abstract units and are not measured SFAG time.

count explains the head–tail complementarity of the balancing scheme. Linearity here refers to the dominant theoretical workload, not to a strictly linear kernel runtime.

SFA/SFAG already bound the number of selected keys through the index; their measured runtime is also affected by sparse memory access, document partitioning, physical K/V length, gradient initialization and write-back, and kernel tiling, so it cannot be predicted by a simple linear function of visible K. Document partitioning, sparse access, gradient write-back, and fixed kernel overhead can all affect runtime, so dynamic pairing should use measured costs rather than a simple linear balance over K or token counts.

Using the illustrative workloads 1 through 8 of the eight equal-length blocks in Figure 10, a contiguous split gives a maximum rank load of $7 + 8 = 15$ and a minimum of $1 + 2 = 3$. Pairing head and tail complementary blocks gives each rank an illustrative workload of $1 + 8 = 9$: the heaviest rank's load falls, while total workload is unchanged. This describes only the pairing geometry; whether P2P is needed and when folding/unfolding happens depend on implementation boundaries and must not be equated directly with the network-level input layout discussed below.

For the TND multi-document case, fixed folding still splits the sequence into $2R$ equal-length Q blocks and gives each rank a head–tail complementary pair. Causal positions reset at document boundaries, however, so cost no longer grows monotonically with global position, and the fixed head–tail rule cannot guarantee equal load or an optimal cost. It may still help, or it may leave substantial imbalance; the actual effect depends on the document distribution and on the overhead of layout rearrangement.

Dynamic pairing estimates cost on the same two-block layout using features such as in-document dense/sparse pair counts, logical KV, and active-document KV, and then pairs light with heavy blocks. It replaces the fixed block ownership only when the predicted peak decreases strictly; otherwise the fixed scheme is kept. Both schemes share the folding execution and gradient-restoration framework. In the network-level implementation of the sparse stage described here, the pairing table decides the selection of local input blocks at the model entry rather than renegotiating P2P communication at every operator or every layer. Inputs, labels, masks, and `position_ids` are selected together, and forward, backward, and recomputation share one immutable plan; non-per-token objectives preserve the original normalization semantics through equivalent weights using the original owner denominators.

Once the pairing is decided, the folded layout must be applied consistently across the network.

The reference operator-boundary folding exchanges or redistributes Q-side tensors at the operator input and immediately unfolds the output after the operator finishes; if Indexer, SFA, and KL each return to the original layout, the same folding boundary is crossed repeatedly. In the sparse stage of this work, folding is planned at the model entry instead: assuming the complete input is already available, the corresponding blocks of `input_ids`, `labels`, `loss_mask`, and `position_ids` are selected locally to form the folded token ownership before Embedding and the main-branch computation. Subsequent layers and all three DSA operators reuse this ownership; inside an operator

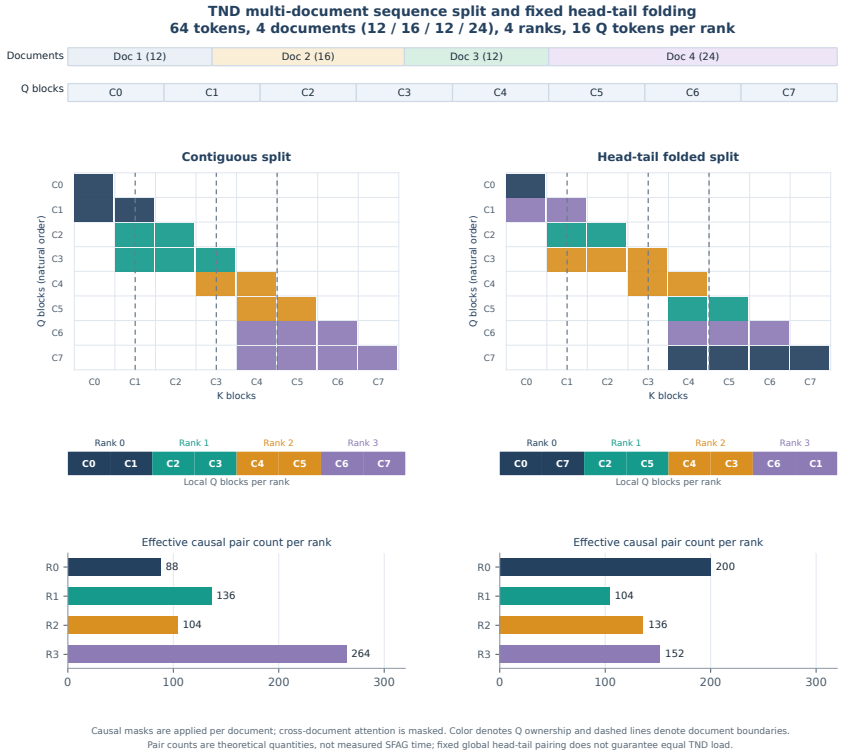


Figure 11: Fixed folding for TND multi-document sequences: causal positions reset at document boundaries, so head–tail complementarity no longer guarantees exactly equal load.

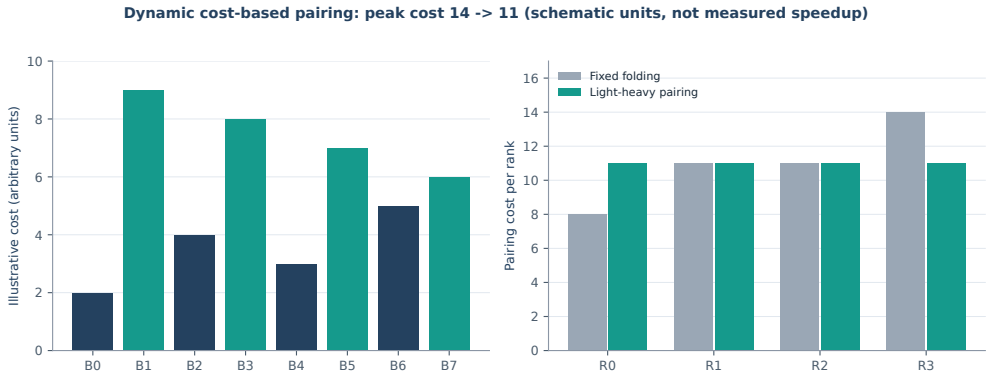


Figure 12: Illustration of dynamic cost-based pairing. $14 \rightarrow 11$ is an illustrative cost, not a measured speedup.

the two Q blocks are only split locally and the required KV is selected, so Q tail slices are no longer exchanged per operator solely for folding and per-Q outputs are no longer moved back to their original rank.

The CP and TP sequence layouts are further planned jointly: the local input of each CP rank is arranged according to the two-block combination of its TP ranks, and the existing TP sequence distribution/layout conversion directly yields complementary Q blocks, avoiding an extra round of TP folding exchange. What is removed is the folding-specific round trip of high-dimensional activations and the corresponding layout conversions in backward and recomputation; KV AllGather, gradient reduction, necessary TP head/sequence conversions, and FSDP communication are retained, so this must not be described as making all attention communication disappear. Positional encoding still uses

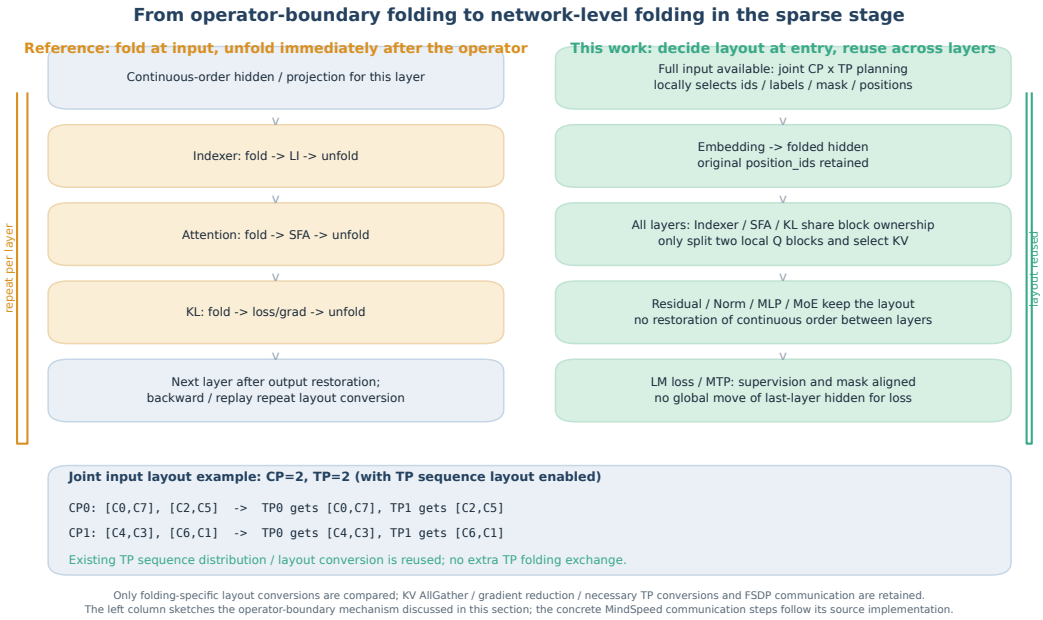


Figure 13: Operator-boundary immediate unfolding versus network-level folding in the sparse stage. Orange marks the extra boundary conversions of the reference path; green marks entry planning and folded-layout reuse. Existing KV/TP/FSDP communication is not removed by this change.

the original `position_ids`, MTP handles the shift per folded block, and the supervision mask and loss normalization are kept consistent, so last-layer hidden states need not be moved back globally to recover the loss.

This applies to the synchronous-CP, $PP = 1$ sparse DSA path of this work. During dense warm-up, the ordinary teacher FA keeps a natural-order main branch and folds only the Q-side tensors of the DSA operators within a layer; the same network-level hidden layout cannot be applied there. The left-hand side of Figure 13 sketches the operator-boundary mechanism discussed in this section; the concrete communication steps of the MindSpeed branch must be taken from its source code and should not be generalized to all MindSpeed paths.

Folding and dynamic pairing within CP/TP can only adjust Q-block load inside one replica; they cannot remove the workload differences that arise when different DP replicas receive different document combinations. Under TND, even when every sample has the same token count after padding, document lengths, causal scan volumes, and active KV may differ. Inter-DP data balancing estimates cost before model execution from lightweight document-boundary metadata: it can use the dense proxy $\sum_i L_i^2$, or the DSA cost model to estimate the cost of the heaviest sequence shard in a complete TND micro-batch; for $LBS > 1$, scoring uses the full layout with $T = LBS \times$ sequence length.

The current step-aware implementation first follows the framework's native shuffle and data-fetch order to obtain the global sample window of this step, then treats each original micro-batch as an indivisible unit, stably sorts them by estimated cost, and assigns them to DP replicas and micro slots in a SNAKE pattern. A candidate is checked against both the maximum cumulative replica cost $\max_d \sum_m C[d, m]$ and the sum of per-micro-slot maxima $\sum_m \max_d C[d, m]$; it is adopted only when neither proxy gets worse and at least one improves. The global sample set, each DP quota, and the composition inside a micro-batch remain unchanged per step, so re-packing cannot change the local denominators of CE/CP/MTP; prefetching only advances a virtual plan, and the state is committed after successful consumption.

This is a conservative heuristic, not a global optimum search: the baseline assignment is kept when each DP has only one micro-batch, when the window is incomplete, when an epoch boundary is crossed, or when the prediction shows no improvement. DP scheduling, document reordering inside a micro-batch, and CP/TP dynamic pairing are three distinct levels; `balance_enabled=True` does

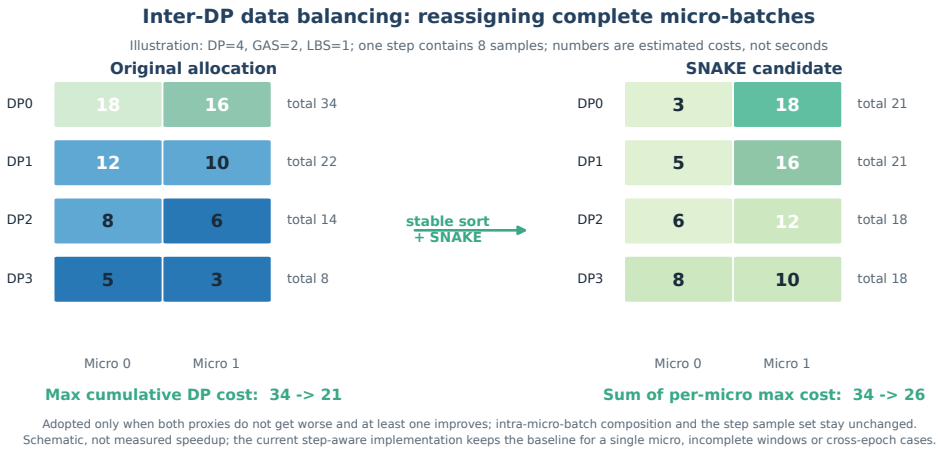


Figure 14: Scheduling illustration for DP=4 and GAS=2. The two cost proxies change from 34/34 to 21/26; the values are abstract estimated costs, not measured time or speedup.

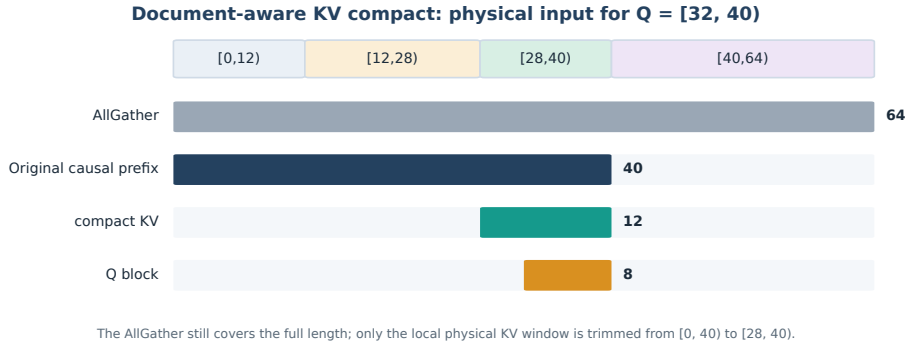


Figure 15: For $Q = [32, 40)$, the physical KV is reduced from $[0, 40)$ to $[28, 40)$; the AllGather still covers the full length.

not require `packed_sequence_reorder=True`. The current DSA data scoring still estimates cost under fixed folding and does not automatically call dynamic `paired_owners`, so when the two are combined the cost objectives should be checked for consistency; using the same coefficient set does not imply identical scoring. The older cross-window `BalanceScheduler` is a separate scheduling path that must be identified against the running version; its benefits must not be added twice or mixed.

3.4.3 KV Window: Document-Aware KV Compact

With a document-isolated causal mask, a Q block $B = [a, b)$ only needs the KV from the start $d(a)$ of its first active document up to b, i.e. $[d(a), b)$, rather than the original $[0, b)$ prefix. The implementation selects the intersecting chunks according to the physical arrangement of the current forward, trims boundaries, rewrites the Q/K cumulative lengths consistently, and removes empty segments and duplicated trailing padding. Indexer, SFA, and KL share the same coordinates. The backward sums gradients over overlapping intervals, then writes them back into the full KV and restores the saved arrangement so that the existing gradient-communication semantics are preserved.

To validate the combined pipeline, companion operator regressions verify document boundaries, Top-K coordinates, overlapping-gradient write-back, and interleaved plans; a weighted combined comparison on 8 devices / 8 layers / 16K with TP2/CP4 ran 8 steps in each configuration, with maximum absolute differences of 0.000530 and 0.000441 for the main loss and MTP loss.

4 POST-TRAINING FRAMEWORK

For reasoning and agentic capability development in Xing 4, we build a unified reinforcement-learning infrastructure. The framework connects task execution, experience collection, reward evaluation, and policy optimization in a common training loop. TeleGym provides the environment interface for heterogeneous task harnesses, while an intermediate server converts their interactions and outcomes into training-ready trajectories. This organization brings mathematical reasoning (Zhao et al., 2025), code generation (Yang et al., 2025b), and multi-turn agent (Bian et al., 2026; Li et al., 2026) tasks into a shared training workflow while preserving task-specific execution and scoring semantics.

The system is organized around two requirements: sustaining training-data production across tasks with different execution lengths, and preserving the relationship between sampled actions, policy versions, and rewards as the model evolves. Disaggregated asynchronous execution combines partial rollout with oversampling to decouple trajectory completion from policy updates. Consistent trajectory records and execution-level reward processing connect this scheduling design to the learning objective. The following sections describe the overall training framework, TeleGym's execution architecture, and asynchronous rollout and policy optimization.

4.1 OVERALL TRAINING FRAMEWORK

Agentic training must accommodate harnesses with different tool interfaces, software dependencies, and interaction patterns. We organize the framework into three cooperating layers, as illustrated in Figure 16. TeleGym provides a shared execution layer for these harnesses, managing task environments and agent sessions. An intermediate server connects their model requests to the policy service and converts the resulting interactions into training trajectories. Slime coordinates rollout scheduling and policy optimization. Together, these layers connect heterogeneous agent executions to a common reinforcement-learning loop.

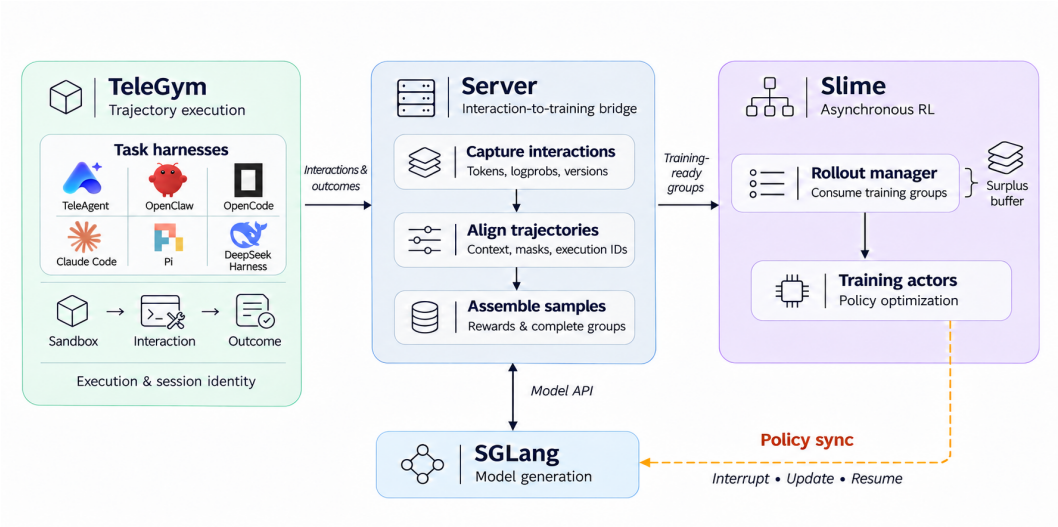


Figure 16: TeleGym executes agent tasks, the server converts interactions and outcomes into training groups, and Slime coordinates policy optimization. The dashed path denotes policy synchronization. Direct mathematical rollout is omitted for clarity.

For each task, the dispatcher supplies the prompt, workspace inputs, and execution configuration to TeleGym. The selected harness interacts with its environment and accesses SGLang generation through the server's model API. Code tasks are evaluated through executable tests; general agent tasks support automatic, model-based, or hybrid evaluation of transcripts and workspace artifacts. Mathematical tasks use direct rollout and answer grading.

The server records sampled tokens, behavior-policy log probabilities, and policy versions, and associates these records with task outcomes. It then aligns context and response tokens, constructs generated-token masks, and assembles scored executions into training groups. Request and execution

identities connect model interactions to asynchronously returned scores. These groups enter a shared policy-optimization pipeline, allowing task execution and evaluation to vary without changing the trainer's input representation.

For multi-turn tasks, one execution may produce several training segments but remains a single independent attempt. Group-relative reward processing therefore operates over executions of the same prompt before propagating the resulting training signal to their constituent segments. This prevents tasks with more model calls from being counted as additional attempts in group statistics. Only acknowledged task outcomes enter group assembly, with missing scores or evaluation failures distinguished from valid unsuccessful attempts.

4.2 TeleGym: Unified Agent Execution

Training across heterogeneous agents requires a common way to prepare task environments, run different harnesses, and maintain interactions over multiple turns. We use TeleGym as the execution and orchestration layer between task engines and the underlying model and sandbox services. Figure 17 presents the broader platform architecture through task access, agent orchestration, resource access, and sandbox management. Within this architecture, the RL execution path turns a dispatched task into a managed agent episode, connecting its workspace, harness session, and model interactions throughout execution. This provides a reusable execution substrate for data synthesis, evaluation, and reinforcement learning (Zou et al., 2026; Lv et al., 2026). We describe four aspects relevant to this training workflow below.

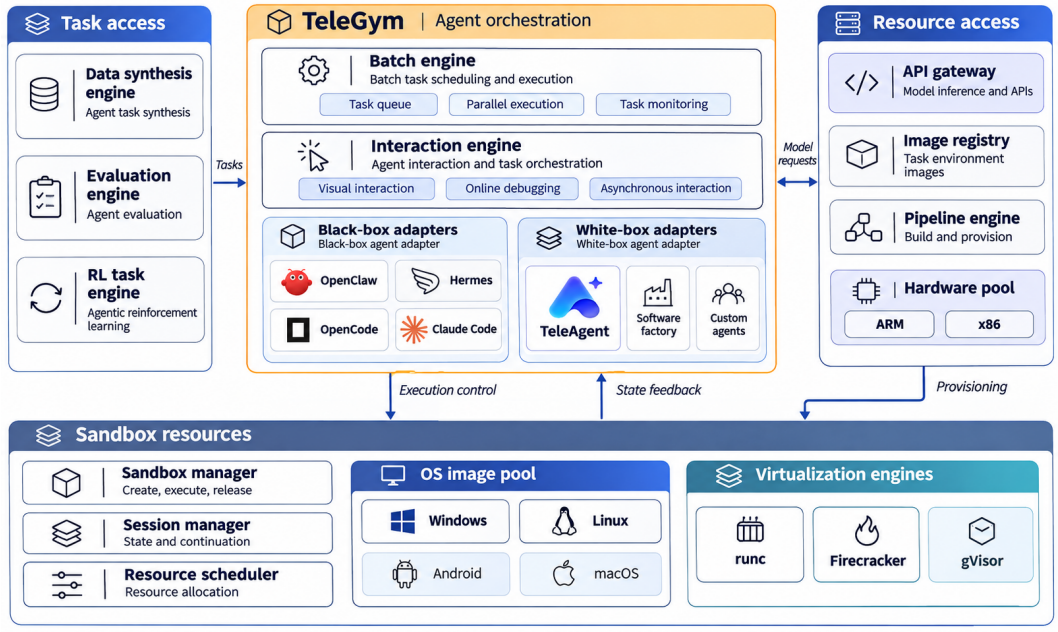


Figure 17: TeleGym platform architecture. Task engines connect to batch and interaction engines through the orchestration layer, which accommodates black-box and white-box agent adapters. Resource access connects model APIs, environment images, and deployment services to the sandbox layer. OS images, virtualization engines, and hardware pools describe the platform's integration scope; the RL execution path discussed here uses the common runtime and sandbox interfaces.

Task-facing execution. Task engines provide the problem specification, workspace inputs, and runtime configuration, while retaining control over task selection and evaluation. The platform distinguishes batch orchestration, organized around task queues, parallel execution, and monitoring, from interactive orchestration, which provides visual interaction, online debugging, and asynchronous interaction. At the runtime boundary, both use a common episode lifecycle for initialization, execution, continuation, and cleanup. Batch workflows invoke this interface across independent tasks; interactive

workflows advance an existing episode through successive actions. This allows task-specific synthesis, evaluation, or RL logic to reuse environment preparation and execution mechanisms without duplicating harness integration.

Harness-independent orchestration. TeleGym exposes a common episode interface while allowing each harness to retain its own agent loop and tool semantics. The platform distinguishes black-box integration of external harnesses from white-box integration of native or custom agents, including TeleAgent. Runtime adapters translate task actions into harness-specific invocations and normalize execution results. This accommodates different agents without requiring the task engine to implement their individual invocation and session protocols. An action may include several model calls and tool interactions, so the interface supports extended agent workflows as well as short executions. Harness adaptation and sandbox adaptation form separate interfaces, allowing either component to be extended without changing the task-facing execution contract.

Sandbox and session management. Sandbox adapters separate agent execution from deployment on local, containerized, or remote resources. For each active episode, TeleGym binds a runtime to a sandbox context, prepares workspace files and dependencies, and initializes the required tools. Image and deployment services supply the underlying execution resources, while the adapter exposes common command-execution and file-access operations. TeleGym retains the harness session across successive actions so that multi-turn execution proceeds in the same task context. Workspace artifacts remain accessible for evaluation before environment cleanup. Episode identity connects these operations and artifacts to the originating task, preserving a coherent execution context across model calls and intervening tool use.

Integration with the training loop. TeleGym configures each runtime to access the policy through the external model gateway and propagates episode-level association metadata. This connects environment-side execution to the model-side records collected by the server: TeleGym manages the agent session and workspace, while the server constructs the token-level trajectories consumed by the trainer. Keeping these responsibilities separate allows the training pipeline to support different harnesses through a shared data interface. Session continuity is maintained at the episode level; policy publication and interruption of active generation are coordinated by the rollout controller, as described next.

4.3 Task Synthesis

Data synthesis prepares executable tasks for general and code agents, each combining an instruction, an initial environment, and evaluation criteria. Multiple harnesses participate in both task design and selection: agents operating through different harnesses propose and refine candidate tasks, which are then independently attempted across harnesses. TeleGym provides a common runtime for these executions, linking design feedback and validation outcomes to each candidate task.

General agents. General-agent tasks are constructed around representative workflows involving information retrieval, document processing, and tool use. Instructions are paired with relevant files and tool interfaces, with variations in constraints and interaction steps to cover both single-turn and multi-turn scenarios. Feedback from multiple harnesses guides refinement of task instructions, tool requirements, and evaluation criteria. Validation runs start from equivalent task conditions and assess outputs or environment state against shared criteria, using programmatic checks or model-based assessment as appropriate.

Code agents. Code-agent tasks are constructed around software projects and programming requirements, including implementation and bug fixing. Agents using multiple harnesses contribute to task formulation and review of the associated tests. Candidate tasks are then attempted independently from the same initial code workspace and dependency configuration. Shared executable tests verify both the requested changes and preservation of existing functionality. These runs expose environment issues and mismatches between task descriptions and tests before selection for training.

Across both categories, task validation and deduplication remove unusable or repetitive examples. Accepted tasks enter the shared rollout pipeline through TeleGym's execution interface.

4.4 Asynchronous Training with Partial Rollout

Agentic rollouts vary in both generation length and environment execution time. Tool calls, workspace operations, and reward evaluation can extend a task beyond the completion of other samples dispatched with it. We combine disaggregated asynchronous execution, bounded oversampling, and partial rollout to accommodate this variability. The design separates three boundaries: completion of a task, assembly of a training batch, and publication of an updated policy. A long-running task can therefore span multiple policy updates without requiring all concurrent tasks to finish together.

Disaggregated execution. Generation and policy optimization run on separate device pools, with the inference service remaining resident while the trainer updates the policy. The rollout pipeline can prepare subsequent training data during optimization of the current batch. TeleGym manages environment-side execution, the server associates interactions with task outcomes, and the rollout manager supplies eligible groups to the trainer. Optimization starts once enough complete, scored groups are available. This producer–consumer organization allows model generation, tool execution, and reward processing to progress independently of individual optimizer steps. Policy publication remains an explicit coordination point between the training and inference pools.

Bounded oversampling. The scheduler authorizes more rollout work than the next training batch requires, providing additional candidates when some tasks finish late or their groups are removed by reward-based filtering. A finite lookahead window limits how far task production can advance beyond training consumption; separate bounds on pending and ready groups constrain unfinished work and accumulated training data. Together, these limits prevent a fast producer or a slow consumer from creating an unbounded backlog. Completed executions are assembled into query groups, and only eligible groups enter the ready buffer. Groups in excess of the current batch remain available for later updates, as illustrated in Figure 18. Retention carries unconsumed data forward; it does not by itself prescribe repeated training on previously consumed groups.

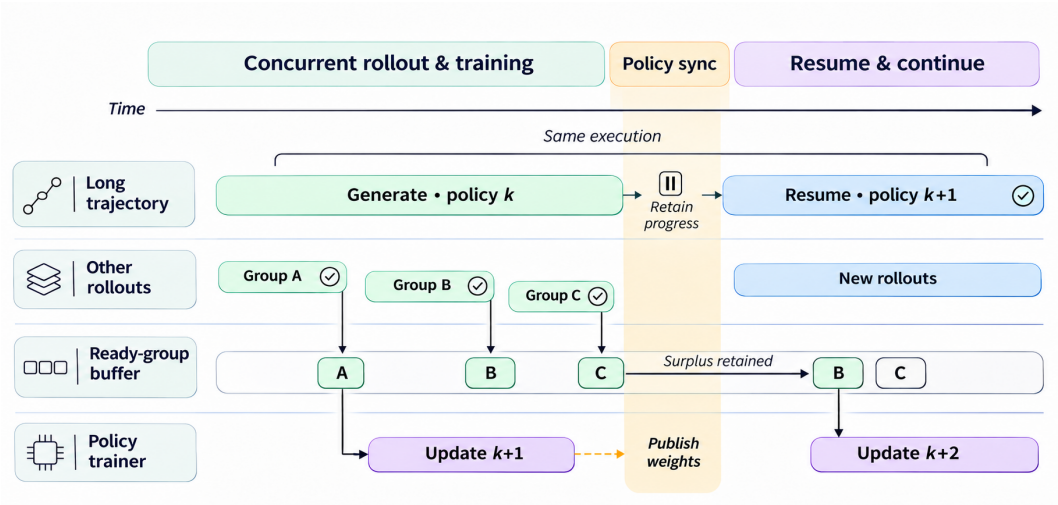


Figure 18: A long trajectory pauses at policy synchronization and resumes under the updated policy. Group A supports an earlier update; surplus groups B and C remain buffered, with B consumed later. Counts and durations are schematic.

Policy publication and trajectory continuation. The policy-publication interval is configurable and need not coincide with every rollout completion. At publication, the controller pauses new model-request admission, interrupts active generation, and waits for model execution to reach a safe synchronization boundary. It then installs the updated weights, advances the policy version, and reopens admission. Weight replacement cannot proceed while old model execution remains active. This ordering prevents an uncontrolled weight change within an executing generation segment while allowing the enclosing task to continue across policy versions.

Unfinished trajectories retain their execution identity and progress and resume after synchronization. For a multi-turn agent, this preserves the association between earlier responses, tool interactions, and subsequent model calls under the updated policy. The synchronization boundary concerns model execution rather than completion of the entire agent session, which may still require further environment interaction or evaluation. Partial rollout preserves work in progress: interruption alone neither assigns a terminal reward nor makes an unfinished prefix eligible for training.

Policy lag and sample eligibility. Resumed trajectories and buffered groups may contain samples generated by earlier policy versions. Each generated segment retains its sampling-time log probabilities, which provide the behavior-policy probabilities used in importance-ratio computation during optimization. Recomputing probabilities with updated weights does not replace this record. Optional stale-data filtering uses observed policy versions to identify old trajectories and removes the entire query group when a constituent trajectory is rejected, preserving the membership used for group-relative reward processing. This is a trajectory-level heuristic rather than a strict token-wise age bound. These mechanisms address policy lag separately from the scheduling limits that bound outstanding work.

Checkpoint boundaries. Checkpointing additionally coordinates the state of data production and consumption. At a save point, the scheduler stops extending the authorized task window and completes the outstanding window before persisting unconsumed groups, data-source progress, execution identities, and version metadata. Model and optimizer state are saved after this data state has been recorded. This differs from routine policy publication, which interrupts and resumes unfinished trajectories. On restart, the saved buffer and cursor preserve the outstanding data stream, while a new execution-generation namespace separates new work from callbacks associated with the previous run.

5 Training Cluster Performance Optimization

The preceding two sections describe distributed pretraining and reinforcement-learning workflows. We now examine how host CPU resources affect execution in the Xing 4 training infrastructure. Host CPUs handle data processing, task submission, and communication scheduling, while the accelerators perform most of the computation. When a critical host thread waits for CPU time, it may submit a compute task late or delay a process's arrival at a collective operation. The delay can then propagate to other training processes. Optimizing cluster performance therefore requires attention to host CPU contention as well as device computation and communication.

We use CPU pinning and core isolation to reduce scheduling delays and variability and to improve effective training throughput within a fixed CPU budget. Pinning sets each thread's CPU affinity, the set of logical CPUs on which it can run. In the coarse-grained scheme, threads in a training process share a CPU pool. The fine-grained scheme partitions that pool by thread role. Core isolation assigns selected system services and movable interrupt handlers to CPUs outside the training pool. Figure 5-1 shows how host contention can affect a training step and where each measure acts.

CPU contention	Critical threads wait	Work is submitted late	Steps take longer
Training threads compete for CPU time with each other and with external work.	Data processing and task scheduling stall.	Compute dispatch and participation in communication are delayed.	Throughput falls or varies across steps.

CPU pinning	Core isolation
Allocate training CPUs through a shared device pool or role-specific subsets.	Move selected system services and movable IRQ handling off training CPUs.

Figure 5-1: How host CPU contention can affect training and where the two optimizations act.

5.1 CPU Affinity and Core Isolation

Coarse-grained pinning assumes one training process per device. The job's CPU set T is divided into disjoint device pools T_i, where i is the device index. Threads belonging to a process share its device pool. A CPU ID identifies a logical CPU in the operating system. A pool of logical CPUs may include hardware threads that share a physical core. At launch, taskset can set CPU affinity, and numactl can set CPU and memory placement policies. After framework initialization, verify each thread's effective affinity mask, the bit mask that identifies the CPUs on which it may run. With load balancing enabled within the pool, the scheduler can move work among its CPUs as the load on different thread roles changes over the course of training. If multiple processes share a device, the unit of CPU allocation must be redefined (MindSpore Team, n.d.-c; Linux man-pages Project, 2026).

Fine-grained pinning divides the same device pool T_i by execution role. PyNative denotes the MindSpore execution mode and its execution threads. MindData covers data loading and preprocessing. Main is the training program's main thread, and runtime denotes the framework components that execute and schedule tasks. For device 0, coarse-grained pinning lets all relevant threads share 28 logical CPUs: {8, 10, 12, 14, 16–39}. Fine-grained pinning assigns {8, 10, 12, 14} to PyNative, 16–23 to MindData and 24–39 to Main. The proposed configuration explicitly places runtime and other training support threads in Main's 24–39 set and verifies their masks after thread creation or recreation. The existing samples do not establish whether this assignment was in use. The allocations of 4, 8 and 16 count logical CPUs. If a role receives too few CPUs, its threads may queue while CPUs assigned to another role sit idle. The benefit depends on how well the allocation matches the load.

Core isolation reserves a separate system CPU pool H, disjoint from the training set T. Selected services, logging and monitoring processes, and handlers for movable interrupt requests (IRQs) are restricted to H. Disabling isolation restores the recorded default placement. The configuration of interrupts that cannot be moved and of other kernel work is held constant and documented. The isolation policy covers only these configurable sources of interference. H must retain enough capacity to prevent a backlog of system work (Kernel Development Community, n.d.). Figure 5-2 compares the pinning schemes under the same CPU budget. Pinning controls resource allocation within training, while isolation limits external competition. Their combined effect must be measured in controlled experiments.

Training CPU set T → per-device pool T_i		
Device 0: {8, 10, 12, 14, 16–39} · 28 logical CPUs		
Coarse-grained pinning	**Fine-grained pinning**	**System CPU pool H**
Shared by all roles Main, PyNative, MindData, runtime and support threads use all 28 CPUs. Load balancing remains enabled within the pool.	**Pools by role** PyNative: 8, 10, 12, 14 MindData: 16–23 Main, runtime and support threads: 24–39 (shared)	**Services, logs, monitoring and movable IRQs** Isolation on: restrict to H. Isolation off: restore recorded defaults.

Figure 5-2: Thread placement and external interference control at a fixed per-device CPU budget.

5.2 Training Throughput Evaluation

We evaluate host resource management using the 285B pretraining workload with MindSpore HyperParallel. Production training runs on an A3 cluster with 8,192 Ascend 910C cards; initial validation used a 512-card A3 cluster. Linux applies CPU affinity per thread, and the framework can reset the mask inherited from the main thread. To enforce a process-wide CPU boundary, use a control group (cgroup) compatible with the resource manager and its cpuset controller for CPU and memory-node placement. Each thread's target set must be within the effective CPU set reported by cpuset.cpus.effective in cgroup v2. A CPU quota limits processor time and imposes no direct limit on thread count. Verify the non-uniform memory access (NUMA) topology and whether simultaneous multithreading (SMT) is enabled. NUMA access costs depend on the processor and memory node. With SMT, a physical core exposes multiple logical CPUs, so disjoint logical CPU sets can share a physical core. Memory binding must be configured separately from CPU affinity (Linux man-pages Project, 2026; Heo, 2015).

The experiment crosses three pinning policies with isolation off or on, giving six configurations (Figure 5-3). The baseline retains only the job-level set T: remove device- and role-level cpuset restrictions and restore the relevant threads' affinity to T. The coarse- and fine-grained configurations use the same device pool T_i and device-level resource limits. Only the division by thread role changes. Hold the overall CPU set and quota, process and thread counts, model and parallelism configuration, memory placement policy and external load constant. Keep load balancing enabled in the training pools and do not add single-CPU bindings for individual threads. Check the actual affinity masks and effective cpusets after framework initialization and after threads are recreated. Residual restrictions would change the baseline.

1. Fix conditions	**2. Set pinning**	**3. Verify placement**	**4. Run six configurations**	**5. Assess effects**
CPU budget, model, load and versions	Job, device or role level	Masks, cgroup and effective cpuset	Each policy with isolation off and on	Throughput, tail latency and scheduling delay

If verification fails, correct the constraints and repeat the check. If it passes, proceed to the six runs. Repeat tests to check attribution and retain results that show no improvement.

Baseline: job set T	**Coarse: device pool T_i**	**Fine: role subsets of T_i**
B0: isolation off	**C0: isolation off**	**F0: isolation off**
B1: isolation on	C1: isolation on	F1: isolation on

Figure 5-3: Validation procedure and six experimental configurations.

Relative to the throughput measured before optimization at each cluster size, enabling CPU pinning and core isolation together increased training throughput by 1–2% on the 512-accelerator cluster and by 5–10% on the 8,192-accelerator cluster (Table 5-1).

Table 5-1: Training throughput gains with CPU pinning and core isolation enabled together.

Optimization	**Cluster size**	**Throughput gain**	**Evaluation condition**
CPU resource isolation	Small (512 accelerators)	1–2%	Pinning and isolation combined
CPU resource isolation	Large (8,192 accelerators)	5–10%	Pinning and isolation combined

Both clusters achieved higher training throughput with CPU pinning and core isolation enabled together. The larger cluster showed the greater relative gain in these tests. The report attributes this difference to host scheduling: pinning restricts the CPUs on which training threads can run, and isolation moves selected background tasks and movable interrupt handling out of the training CPU pool. These changes reduced scheduling delays for critical threads, allowing processes that were late to become ready within a communication group to submit their data earlier. This shortened synchronization waits on the critical path of the training step and yielded a larger throughput gain on the larger cluster.

6 Conclusion

We have presented the architectural design principles and training infrastructure underpinning the Xing 4 series, extending the TeleChat lineage on Ascend. The open-sourced 29B member and the larger-scale studies in this report illustrate the range of configurations within the series. Across these efforts, model capacity, sparse computation, attention, depth-width balance, and decoding are considered together with compute, memory, and communication costs.

The pretraining system brings topology-aware hybrid sharding and owner-based Muon execution together with coordinated management of long-sequence workloads and activation memory. CPU

affinity and core isolation address host scheduling interference that can propagate through distributed execution. Implemented with MindSpore HyperParallel, these mechanisms align model-state placement, matrix computation, communication, and host resources with the Ascend platform. The large-scale pretraining run documented in this report demonstrates their integration on 8,192 Ascend 910C cards.

We have introduced TeleGym, a unified agent execution platform, and integrated it into the reinforcement-learning infrastructure for the Xing 4 series. Shared execution interfaces and trajectory construction connect task synthesis, environment interaction, and reward evaluation to policy optimization. Disaggregated asynchronous execution, bounded oversampling, and partial rollout accommodate heterogeneous task durations and allow trajectories to continue across policy updates. Sampling-time policy records and execution-level reward grouping retain the correspondence between generated actions and learning signals, supporting a common training workflow for reasoning, coding, and multi-turn agents.

Building on this work, we are currently training trillion-parameter models in the Xing 4 series on Ascend. This ongoing effort extends the series toward larger model capacity while investigating training efficiency and system scalability at that scale. Further development will address the interaction between model scale, longer contexts, and increasingly complex agentic workloads.

REFERENCES

Yushi Bai, Qian Dong, Ting Jiang, Xin Lv, Zhengxiao Du, Aohan Zeng, et al. Indexcache: Accelerating sparse attention via cross-layer index reuse, 2026. URL https://arxiv.org/abs/2603.12201.

Kai Bian, Haosi Mo, Xuebo Liu, Shuangyong Song, Jing Li, Yongxiang Li, et al. Agentvocab: Structure-aware vocabulary adaptation for efficient llm agents. In *Forty-third International Conference on Machine Learning*, 2026.

Jian Chen, Yesheng Liang, and Zhijian Liu. DFlash: Block diffusion for flash speculative decoding, 2026. URL https://arxiv.org/abs/2602.06036.

Xin Cheng, Xingkai Yu, Chenze Shao, Jiashi Li, Yunfan Xiong, Yi Qian, et al. DSpark: Confidence-scheduled speculative decoding with semi-autoregressive generation, 2026. URL https://arxiv.org/abs/2607.05147.

DeepSeek-AI, Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, et al. Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model, 2024a. URL https://arxiv.org/abs/2405.04434. arXiv preprint arXiv:2405.04434.

DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, et al. Deepseek-v3 technical report, 2024b. URL https://arxiv.org/abs/2412.19437. arXiv preprint arXiv:2412.19437.

DeepSeek-AI, Aixin Liu, Aoxue Mei, Bangcai Lin, Bing Xue, Bingxuan Wang, et al. Deepseek-v3.2: Pushing the frontier of open large language models, 2025. URL https://arxiv.org/abs/2512.02556.

DeepSeek-AI, Anyi Xu, Bangcai Lin, Bing Xue, Bingxuan Wang, Bingzheng Xu, et al. DeepSeek-V4: Towards highly efficient million-token context intelligence, 2026. URL https://arxiv.org/abs/2606.19348. arXiv preprint arXiv:2606.19348.

GLM-5-Team, Aohan Zeng, Xin Lv, Zhenyu Hou, Zhengxiao Du, Qinkai Zheng, et al. GLM-5: from vibe coding to agentic engineering, 2026. URL https://arxiv.org/abs/2602.15763. arXiv preprint arXiv:2602.15763.

Zhongjiang He, Zihan Wang, Xinzhang Liu, Shixuan Liu, Yitong Yao, Yuyao Huang, et al. Telechat technical report, 2024. URL https://arxiv.org/abs/2401.03804. arXiv preprint arXiv:2401.03804.

Tejun Heo. Control group v2. The Linux Kernel Documentation, 2015. URL https://docs.kernel.org/admin-guide/cgroup-v2.html. Continuously updated documentation. Accessed November 12, 2026.

Wentao Hu, Yanbo Zhai, Xiaohui Hu, Mingkuan Zhao, Xue Liu, Kaidong Yu, et al. Awakening dormant experts: Counterfactual routing to mitigate moe hallucinations. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 47257–47271, 2026a.

Wentao Hu, Mingkuan Zhao, Shuangyong Song, Xiaoyan Zhu, Xin Lai, and Jiayin Wang. Mosaic pruning: A hierarchical framework for generalizable pruning of mixture-of-experts models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 21885–21893, 2026b.

Kernel Development Community. CPU isolation. The Linux Kernel Documentation, n.d. URL https://docs.kernel.org/admin-guide/cpu-isolation.html. Accessed November 12, 2026.

Kimi Team, Tongtong Bai, Yifan Bai, Yiping Bao, M. C., Jianfeng Cai, et al. Kimi K3: Open frontier intelligence, 2026. URL https://arxiv.org/abs/2607.24653. arXiv preprint arXiv:2607.24653.

Dongfang Li, Zhuoen Chen, Bonian Jia, Zixuan Liu, Zhengxuan Lu, Can Deng, et al. Memory-centric continual learning for ai agents. *arXiv preprint arXiv:2610.29357*, 2026.

Linux man-pages Project. sched_setaffinity(2): Set and get a thread's CPU affinity mask. Linux man-pages 6.19, 2026. URL https://man7.org/linux/man-pages/man2/sched_setaffinity.2.html. Accessed November 12, 2026.

Jingyuan Liu, Jianlin Su, Xingcheng Yao, Zhejun Jiang, Guokun Lai, Yulun Du, et al. Muon is scalable for LLM training, 2025a. URL https://arxiv.org/abs/2502.16982. arXiv preprint arXiv:2502.16982.

Xinzhang Liu, Chao Wang, Zhihao Yang, Zhuo Jiang, Xuncheng Zhao, Haoran Wang, et al. Training report of TeleChat3-MoE, 2025b. URL https://arxiv.org/abs/2512.24157. arXiv preprint arXiv:2512.24157.

Haisheng Lv, Zikun Su, Jingbang Liu, Zilu Huang, Xinzhang Liu, Wenmin Deng, et al. Beyond sft-optimality: Cold-start data selection for reinforcement learning via gradient subspace coverage. In *Proceedings of the 2026 Conference on Empirical Methods in Natural Language Processing*, pp. 11769–11789, 2026.

MindSpore Team. HyperParallel: A SuperPod-affinity training acceleration framework. Official software repository, n.d.-a. URL https://github.com/mindspore-ai/hyper-parallel. Accessed November 8, 2026.

MindSpore Team. MindSpore. Official software repository, n.d.-b. URL https://github.com/mindspore-ai/mindspore. Accessed November 12, 2026.

MindSpore Team. msrun launching. MindSpore Tutorials, n.d.-c. URL https://www.mindspore.cn/tutorials/en/master/parallel/msrun_launcher.html. Accessed November 12, 2026.

Mo Sun, Yifan Yao, Yanwei Liu, Luobin Liu, Zhenzhang Yang, Kaisheng Wang, et al. HyperParallel-FSDP: Topology-aware fully sharded training with layout-driven muon on Ascend SuperPods. *arXiv preprint arXiv:2609.21594*, 2026. URL https://arxiv.org/abs/2609.21594.

Tencent Hunyuan Team, Ao Liu, Botong Zhou, Can Xu, Chayse Zhou, ChenChen Zhang, et al. Hunyuan-A13B technical report, 2026. URL https://arxiv.org/abs/2609.27284. arXiv preprint arXiv:2609.27284.

Zihan Wang, Yitong Yao, Li Mengxiang, Zhongjiang He, Chao Wang, Shuangyong Song, et al. Telechat: An open-source billingual large language model. In *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10)*, pp. 10–20, 2024.

Zihan Wang, Xinzhang Liu, Yitong Yao, Chao Wang, Yu Zhao, Zhihao Yang, et al. Technical report of telechat2, telechat2.5 and t1, 2025. URL https://arxiv.org/abs/2507.18013. arXiv preprint arXiv:2507.18013.

Zhenda Xie, Yixuan Wei, Huanqi Cao, Chenggang Zhao, Chengqi Deng, Jiashi Li, et al. mHC: Manifold-constrained hyper-connections, 2026. URL https://arxiv.org/abs/2512.24880. arXiv preprint arXiv:2512.24880.

XingChen-AGI. Xing4.0-29B-A4B. Official model card and weights, Hugging Face, 2026. URL https://huggingface.co/XingChen-AGI/Xing4.0-29B-A4B. Accessed November 14, 2026.

Sishi Xiong, Ziyang He, Zhongjiang He, Yu Zhao, Changzai Pan, Jie Zhang, et al. Tablezoomer: a collaborative agent framework for large-scale table question answering. *Vicinagearth*, 2(1):11, 2025.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, et al. Qwen3 technical report, 2025a. URL https://arxiv.org/abs/2505.09388. arXiv preprint arXiv:2505.09388.

Jian Yang, Wei Zhang, Shark Liu, Jiajun Wu, Shawn Guo, Yizhi Li, et al. From code foundation models to agents and applications: A practical guide to code intelligence. *arXiv preprint arXiv:2511.18538*, 2025b.

Z.ai. GLM-5.2: Built for long-horizon tasks. Official technical blog, 2026. URL https://z.ai/blog/glm-5.2.

Wei Zhang, Xiangyuan Guan, Lu Yunhong, Jie Zhang, Shuangyong Song, Xianfu Cheng, et al. Lemur: Log parsing with entropy sampling and chain-of-thought merging. *arXiv preprint arXiv:2402.18205*, 2024.

Deji Zhao, Donghong Han, Jia Wu, Zhongjiang He, Bo Ning, Ye Yuan, et al. Enhancing math reasoning ability of large language models via computation logic graphs. *Knowledge-Based Systems*, 325:113905, 2025.

Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, et al. PyTorch FSDP: Experiences on scaling fully sharded data parallel, 2023. URL https://arxiv.org/abs/2304.11277. arXiv preprint arXiv:2304.11277.

Lexiao Zou, Yitong Yao, Gongwei Chen, Wenmin Deng, Chao Wang, Shuangyong Song, et al. Group structural proximity shapes exploration-exploitation in tool-use reinforcement learning. In *Findings of the Association for Computational Linguistics: EMNLP 2026*, pp. 4682–4698, 2026.

A AUTHOR LIST

This work was initiated and led by Xingchen AGI Lab, China Telecom, in collaboration with Huawei and Harbin Institute of Technology (Shenzhen). The complete list of authors is as follows (authors are listed alphabetically by their first name):

Baotian Hu, Bingkai Yang, Boyang Wang, Changzai Pan, Chao Lyu, Chao Wang, Chen Li, Chenshuo Pan, Chenyu Lin, Chenyuan Zhang, Chongming Liu, Dawei Fan, Di Wu, Dong Guo, Dong Li, Dongfang Li, Enlin Liu, Fanye Xiao, Fengkai Li, Fubei Yao, Gang Lin, Han Gao, Hanming Wu, Haowei He, Huapeng Liu, Hua Wang, Huinan Xu, Jiahong Xiong, Jiaxin Peng, Jianxiang Ma, Jianxing Zhuge, Jie Zhang, Ji Xiong, Jinglin Chen, Jingqi Li, Jingxuan Li, Jiwang Ye, Junchen Li, Junbiao Li, Junyang Wu, Kaidong Yu, Kaipeng Jia, Kehai Chen, Kuai Li, Laobin Liu, Lei Li, Libo Qin, Lingling Shi, Luwen Pu, Menghao Xu, Mengjiao Bao, Mengqi Chen, Min Zhang, Mingrui Wang, Naying Wang, Peicong Han, Peiming Guo, Pengfei Li, Qiang Feng, Qi Wang, Qingjun Zhu, Qingli Li, Qitai Ding, Qiyi Xie, Quanchao Zhao, Rongbao Han, Rui Xue, Ruiyu Fang, Ruiting Song, Shijie Li, Shiquan Wang, Shuangyong Song, Sishi Xiong, Teng Su, Wankang Bao, Weijie Zhang, Wei Mao, Weiwei Fu, Weiyang Huang, Wenmin Deng, Wenqing Hu, Wujie Zhong, Xi Yu, Xiaohui Hu, Xiaomin Fang, Xiaomeng Huang, Xiaoqiong Shi, Xiaotong Shang, Xiaoyu Li, Xinglei Xu, Xin Jiang, Xin Wang, Xin Zhang, Xinzhang Liu, Xue Liu, Xuefeng Bai, Xuewei He, Yahan Li, Yanhan Zhang, Yanlei Xue, Yanzhang Wu, Yang Yu, Yi Liu, Yin Zhang, Yingyan Li, Yitong Yao, Yixuan Li, Yongqiang Xia, Yongxiang Li, Yinzhou Zhu, Yu Zhao, Yuhao Jiang, Yunpeng Zhang, Yuting Lin, Yuxuan Wu, Yuyao Huang, Yuyi Wang, Zebo Peng, Zhenwu Pei, Zhenxiong Li, Zhao Wang, Zhaoxi Zhang, Zhihao Yang, Zhiqiang Dong, Zhiwei Yan, Zhilong Lu, Zhonghao Guo, Zhongheng Wu, Zhongjiang He, Zhongqiu Li, Zhu Yuan, Zhuoru Jiang, Zhuoyi Wei, Zijia Guo, Zilu Huang, Zhipeng Guo, Zirui Li, Zihan Wang.