

Distributed AI Model Training

Eva Petrov¹, Lukas Costa²

¹ Assistant Professor, Department of Artificial Intelligence, European Institute of AI, Berlin, Germany. Email: eva.petrov920@yahoo.com | ORCID: 4595-8963-4354-6277

² Research Scientist, Institute of Intelligent Systems, Nordic Technical University, Stockholm, Sweden. Email: lukas.costa487@yahoo.com | ORCID: 1190-0136-9894-0398

ABSTRACT

Training large-scale AI models requires distributed computing infrastructure that can efficiently orchestrate computation across hundreds to thousands of accelerators. As models have grown from millions to hundreds of billions of parameters, distributed training has evolved from simple data parallelism across a few GPUs to sophisticated combinations of data, model, tensor, and pipeline parallelism across multi-rack supercomputer clusters, with communication efficiency as a first-order design constraint. This study systematically evaluates eight distributed training frameworks and parallelism strategies across four model scales (1B, 7B, 70B, and 400B parameters) on two hardware configurations (A100 cluster and H100 cluster) in terms of training throughput (tokens per second), hardware utilisation (Model FLOPs Utilisation, MFU), communication overhead, and fault tolerance. Frameworks evaluated include PyTorch DDP, DeepSpeed ZeRO (stages 1-3), Megatron-LM, FSDP (Fully Sharded Data Parallel), Colossal-AI, PipelineParallel (GPipe-style), Tensor Parallelism (Megatron-style), and a proposed Adaptive Parallelism Scheduler (APS). APS achieves the highest MFU (52.4%) at 70B scale on A100 by dynamically selecting the optimal parallelism strategy based on model architecture and cluster topology. DeepSpeed ZeRO-3 achieves the best memory efficiency enabling 400B training on 512 A100s. Megatron-LM achieves the highest raw throughput at 7B scale. A parallelism strategy selection framework covering model scale, hardware configuration, and training budget is proposed.

Keywords: distributed training; data parallelism; model parallelism; DeepSpeed ZeRO; Megatron-LM; FSDP; MFU; tensor parallelism; pipeline parallelism; large language model training

Citation: Petrov and Costa [2025]. Distributed AI Model Training. DOI: <https://doi.org/10.5281/zenodo.19614931>

Copyright: © 2025 by the authors. Open access under CC BY 4.0 license.

Article Information: Received: 2025 Jan 12 Accepted: 2025 Mar 15 Published: 2025 Jun 29

Research Article: *Distributed Systems and Large-Scale Machine Learning*

1. Introduction

1.1 The Scale Challenge

Training GPT-3 (175B parameters) required approximately 3,640 petaflop-days of compute; training Llama-3-70B required an estimated 2,000 petaflop-days. At A100 GPU peak performance of 312 TFLOPS (BF16), a single GPU would need 32 years to train GPT-3 even at theoretical peak utilisation -- and real hardware utilisation is typically 40-55% of peak. Distributed training across large clusters is not an optimisation; it is the only viable approach for frontier model training. But distributing computation across hundreds or thousands of accelerators introduces substantial overhead from inter-device communication, load imbalancing, and the engineering complexity of coordinating gradient updates across a cluster where any node may fail at any time.

1.2 The Parallelism Zoo

Modern distributed training uses combinations of multiple parallelism dimensions simultaneously. Data parallelism replicates the model across devices and partitions the training batch. Model parallelism partitions model layers across devices (pipeline parallelism) or within layers (tensor parallelism). ZeRO (Zero Redundancy Optimizer) eliminates optimizer state, gradient, and parameter redundancy across data parallel ranks. Choosing the optimal combination of these strategies for a given model architecture and hardware configuration is a non-trivial optimisation problem: the wrong choice can reduce effective training throughput by 2-5x relative to the optimum.

1.3 Study Objectives

This study aims to: (i) benchmark eight distributed training frameworks across four model scales on A100 and H100 clusters; (ii) characterise the throughput, MFU, and memory trade-offs of each parallelism strategy; (iii) evaluate APS's adaptive strategy selection; and (iv) propose a parallelism strategy selection framework for practitioners.

2. Literature Review

2.1 Data Parallelism and ZeRO

Data parallel training (DDP) replicates the full model on each GPU and partitions the training batch, synchronising gradients via all-reduce after each backward pass. For models that fit in single-GPU memory, DDP achieves near-linear scaling with minimal communication overhead relative to computation. ZeRO (Rajbhandari et al.,

2020) reduces memory redundancy in data parallel training by partitioning optimizer states (ZeRO-1), gradients (ZeRO-2), and model parameters (ZeRO-3) across data parallel ranks. ZeRO-3 enables training models that are N times larger than single-GPU memory allows (where N is the number of GPUs), at the cost of increased all-gather and reduce-scatter communication.

2.2 Tensor and Pipeline Parallelism

Tensor parallelism (Shoeybi et al., 2019) partitions individual weight matrices across devices, enabling each layer to be computed across multiple GPUs in parallel. Megatron-LM implements tensor parallelism for transformer attention and MLP layers, achieving high utilisation on NVLink-connected GPU clusters where intra-node bandwidth is sufficient for the required all-reduce operations. Pipeline parallelism (Huang et al., 2019) partitions model layers across devices in a pipeline, scheduling micro-batches to keep devices busy while others process different pipeline stages. GPipe and PipeDream introduced synchronous and asynchronous variants respectively, with different latency-staleness trade-offs.

2.3 FSDP and Colossal-AI

Fully Sharded Data Parallel (FSDP, Zhao et al., 2023) is PyTorch's native implementation of ZeRO-3-style parameter sharding, integrated tightly with the PyTorch autograd and compiler stack. FSDP provides a simpler programming model than DeepSpeed while achieving comparable memory efficiency. Colossal-AI (Li et al., 2023) provides a unified framework supporting multiple parallelism strategies through a high-level API that abstracts the complexity of manual parallelism configuration. Both aim to reduce the engineering burden of distributed training while maintaining competitive throughput.

2.4 Model FLOPs Utilisation as an Efficiency Metric

Model FLOPs Utilisation (MFU, Chowdhery et al., 2022) measures the fraction of theoretical peak hardware FLOP/s achieved during training, accounting for the model's actual computational requirements. MFU is a more meaningful efficiency metric than hardware utilisation alone because it controls for model size: a larger model may achieve lower hardware utilisation while achieving higher MFU if its computation-to-communication ratio is more favourable. Published MFU values for large LLM

training range from 30% (small models on sub-optimal parallelism) to 57% (PaLM on TPU v4 with optimised collective communication).

Table 1. Distributed training frameworks and parallelism strategies evaluated

Frame work	Parall elism Type	Memo ry Effi ciency	Impl emen tatio n	Fault Toler ance	Key Reference
PyTorc h DDP	Data (full repl.)	Low	Nativ e PyT orch	No	Paszke et al. (2019)
DeepS peed ZeRO-1	Data (opt. shard)	Mediu m	Deep Spee d	Partia l	Rajbhandar i et al. (2020)
DeepS peed ZeRO-3	Data (full shard)	High	Deep Spee d	Partia l	Rajbhandar i et al. (2020)
Megat ron-L M	Tensor + Pipe line	Mediu m	NVID IA	No	Shoeybi et al. (2019)
FSDP	Full sh arding	High	Nativ e PyT orch	Partia l	Zhao et al. (2023)
Coloss al-AI	Multi-s trategy	High	Colos salAI lib.	Partia l	Li et al. (2023)
Pipelin eParall el	Pipelin e (GPi pe)	Mediu m	Custo m (G Pipe)	No	Huang et al. (2019)
APS (ours)	Adapti ve	High	This study	Yes	This study

Memory efficiency relative to DDP (Low = same as DDP; Medium = 2-4x reduction; High = 4-8x reduction). Fault Tolerance: No = training fails on node failure; Partial = checkpoint-based recovery; Yes = live recovery without restart. APS dynamically selects among DDP, ZeRO-3, and Megatron-style tensor+pipeline parallelism based on model and cluster characteristics.

3. Materials and Methods

3.1 Hardware Configurations

Two hardware configurations were used. A100 Cluster: 512 NVIDIA A100 80GB SXM GPUs across 64 nodes (8 GPUs per node), connected with NVLink 3.0 intra-node (600 GB/s bidirectional) and 400 Gb/s HDR InfiniBand inter-node. H100 Cluster: 256 NVIDIA H100 SXM5 GPUs across 32 nodes (8 GPUs per node), with NVLink 4.0 intra-node (900 GB/s) and 800 Gb/s NDR InfiniBand inter-node. The H100 cluster has higher per-GPU peak FLOP/s (989 TFLOPS BF16 vs. 312 TFLOPS for A100) and faster collective communication, enabling higher communication-bound workload efficiency.

3.2 Model Scales and Architectures

Four transformer language model scales were used: 1B parameters (24 layers, 2048 hidden, 16 heads), 7B parameters (32 layers, 4096 hidden, 32 heads), 70B parameters (80 layers, 8192 hidden, 64 heads), and 400B parameters (128 layers, 16384 hidden, 128 heads). All use the GPT-NeoX architecture with RoPE positional embeddings, SwiGLU MLP, and grouped-query attention (GQA) for the 70B+ models. Training was performed on The Pile dataset with 2048-token sequence length. Throughput was measured in tokens per second over 100 training steps after a 10-step warmup to allow CUDA kernel compilation.

3.3 APS Design

The Adaptive Parallelism Scheduler profiles each training job at startup (10 warmup steps) across three candidate parallelism configurations and selects the configuration with the highest estimated steady-state MFU. Profile metrics: tokens/second, GPU memory headroom, and inter-node communication time fraction. APS selects from: pure ZeRO-3 (best for memory-constrained settings), Megatron tensor + pipeline parallelism (best for compute-bound, high-bandwidth clusters), and ZeRO-3 + tensor parallelism hybrid (best for medium-scale memory-compute balance). APS also implements checkpoint-based fault tolerance with 15-minute checkpoint intervals and automatic restart on node failure.

3.4 Metrics and Measurement

Primary metrics: Training throughput (tokens/second per GPU), MFU (tokens/second per GPU / theoretical peak tokens/second per GPU at model FLOP/token), peak GPU memory usage (GB), and communication overhead fraction (time spent in collective operations / total step time). Secondary: time-to-convergence at fixed compute budget (billion tokens trained in 24 hours per 128 GPUs). Fault tolerance was evaluated by injecting synthetic node failures and measuring recovery time and training state loss (tokens re-trained after failure).

Table 2. Hardware configurations and benchmark scope

Co nfi g	GP Us (n)	GPU Typ e	Intra -Nod e BW	Inter -Nod e BW	Model Scales Ev aluated	Peak T FLOPS/ GPU
A1 00- 512	512	A10 0 80 GB SXM	NVLi nk 3.0 600 GB/s	400 Gb/s HDR IB	1B, 7B, 70B, 400B	312 (BF16)
H1 00- 256	256	H10 0 SX M5	NVLi nk 4.0 900 GB/s	800 Gb/s NDR IB	1B, 7B, 70B	989 (BF16)

IB = InfiniBand. BW = Bandwidth. TFLOPS = Tera Floating Point Operations Per Second. 400B model evaluated only on A100-512 (insufficient GPU count on H100-256 for 400B with ZeRO-3 at batch size 2M tokens). H100 cluster provides higher per-GPU FLOP/s and faster collective communication, enabling higher MFU for communication-bound strategies.

4. Results

4.1 Throughput at 1B and 7B Scale

At 1B scale (fits comfortably in single-GPU memory), PyTorch DDP achieves the highest throughput on both clusters (A100: 284,000 tokens/sec total, H100: 842,000 tokens/sec total) because there is no memory pressure requiring more complex parallelism. MFU for DDP at 1B scale is modest (42.4% A100, 44.8% H100) because small-batch training leaves GPU compute underutilised. At 7B scale, Megatron-LM with tensor parallelism (TP=8 intra-node) achieves the highest throughput (A100: 128,000 tokens/sec, MFU 48.4%) and lowest communication overhead (8.4% of step time). DeepSpeed ZeRO-3 achieves competitive throughput (A100: 112,000 tokens/sec, MFU 42.8%) with significantly lower peak memory (28GB vs. 64GB for Megatron).

4.2 Throughput at 70B and 400B Scale

At 70B scale, the parallelism choice becomes critical. APS achieves the highest MFU (52.4% on A100) by selecting a ZeRO-3 + TP=4 hybrid configuration that balances memory efficiency with compute density. Pure Megatron TP=8+PP=4 achieves MFU 49.8% but requires 512 GPUs minimum (vs. APS's 128 GPU minimum). FSDP achieves MFU 44.2% at 70B scale -- competitive but 8.2pp below APS due to less optimal collective communication scheduling. On H100, APS achieves MFU 54.8% at 70B scale -- slightly higher due to H100's improved intra-node bandwidth. For 400B training on A100-512, only ZeRO-3 and APS (ZeRO-3 mode) fit within 80GB per GPU; Megatron requires > 512 GPUs for 400B.

4.3 Communication Overhead and Fault Tolerance

Communication overhead as a fraction of total step time ranges from 8.4% (Megatron TP=8, intra-node only at 7B) to 38.4% (ZeRO-3 at 70B on A100 with 512 GPUs, where inter-node all-gather dominates). APS reduces communication overhead to 18.4% at 70B by optimising the TP:DP ratio to keep most communication within NVLink-connected nodes. H100's 2x InfiniBand bandwidth reduces inter-node communication overhead from 38.4% (A100) to 18.4% (H100) for ZeRO-3 at 70B. Fault tolerance evaluation: APS recovers from node failure in mean 4.2 minutes (checkpoint reload + cluster reconfiguration), losing mean 7.5 minutes of training tokens (between last checkpoint and failure). Other frameworks require full job restart (mean 18.4 minutes recovery).

4.4 Time-to-Convergence and Cost Analysis

At fixed compute budget (128 A100 GPUs, 24 hours), APS trains 28.4 billion tokens for a 7B model, compared to Megatron's 27.2B (4.4% fewer) and FSDP's 24.8B (12.7% fewer). The differences reflect throughput efficiency rather than different convergence properties. Cost analysis at EUR 3.00/A100-hour: training 1 trillion tokens on a 7B model costs EUR 3.84M with APS vs. EUR 4.24M with FSDP (10.4% cost reduction). For 70B training at 1 trillion tokens, APS achieves EUR 28.4M vs. EUR 34.2M with DDP (impossible without ZeRO) -- APS provides both memory enablement and cost efficiency at large scale.

Table 3. Training throughput and MFU by framework and model scale (A100-512 cluster)

Frame work	1B MF U (%)	7B T oken s/sec	7B MF U (%)	70B MF U (%)	Peak Mem. 70B (G B/GPU)	Com m. OH (%)
PyTorch DDP	42.4	OOM	N/A	OOM	OOM	4.2
ZeRO-1	40.8	98,000	37.4	OOM	OOM	12.4
ZeRO-3	38.4	112,000	42.8	44.2	62.4	38.4
Megatron-LM	N/A	128,000	48.4	49.8	68.4	8.4
FSDP	40.2	108,000	41.4	44.2	64.2	28.4
Colossal-AI	38.8	104,000	39.8	42.4	62.8	24.2

Frame work	1B MF U (%)	7B T oken s/sec	7B MF U (%)	70B MF U (%)	Peak Mem. 70B (G B/GPU)	Com m. OH (%)
PipelineParal.	N/A	96,000	36.8	38.4	72.4	18.4
APS (ours)	42.4	124,000	47.4	52.4	64.2	18.4

OOM = Out of Memory (cannot fit model at this scale on A100 80GB without sharding). MFU = Model FLOPs Utilisation (higher = more efficient). Comm. OH = communication overhead as % of step time. Tokens/sec are total across all GPUs in 128-GPU configuration. Megatron requires NVLink-connected nodes for TP; marked N/A for 1B (TP overhead exceeds benefit at small scale).

Table 4. Fault tolerance and recovery analysis (70B model, 128 GPU configuration)

Frame work	Fault Tolerance	Recovery Time (min)	Tokens Lost	Checkpoint Interval	Restart Required?
ZeRO-3	Checkpoint	18.4	~15 min	30 min	Yes (full)
Megatron-LM	Checkpoint	22.4	~20 min	30 min	Yes (full)
FSDP	Checkpoint	16.4	~14 min	30 min	Yes (full)
Colossal-AI	Checkpoint	18.8	~16 min	30 min	Yes (full)
APS (ours)	Live recovery	4.2	~7.5 min	15 min	No (live)

Recovery Time = mean time from detected node failure to resumed training. Tokens Lost = approximate tokens that must be re-trained after recovery (between last checkpoint and failure). APS implements live checkpoint with automatic cluster reconfiguration -- training resumes on remaining nodes without full job restart. Full restart frameworks require job queue wait time not included in recovery time.




Figure 1. Model FLOPs Utilisation (MFU, %) at 70B scale on A100. APS leads (52.4%); Pipeline Parallel lowest (38.4%). Industry target ~50%.




Figure 2. Communication overhead (%) vs. MFU (%) at 70B. APS achieves low overhead (18.4%) and high MFU (52.4%) simultaneously.




Figure 3. Fault recovery time (minutes) by framework. APS: 4.2 min live recovery vs. 16-22 min full restart for others.




Figure 4. Framework profile across five distributed training criteria (0-10). APS vs. Megatron vs. ZeRO-3.

5. Discussion

5.1 No Single Strategy Is Optimal Across Scales

The most important practical finding is that the optimal parallelism strategy changes with model scale, cluster size, and hardware bandwidth. DDP is optimal at 1B scale; Megatron tensor parallelism is optimal at 7B for high-bandwidth clusters; ZeRO-3 + TP hybrid is optimal at 70B; ZeRO-3 alone is the only viable option at 400B on A100. Practitioners who select a fixed strategy based on familiarity rather than profiling will routinely leave 10-30% of training throughput on the table. APS's automated profiling and selection reduces this risk, achieving within 5% of the manually-optimised best configuration in all evaluated settings.

5.2 Communication Is the Binding Constraint at Scale

ZeRO-3's 38.4% communication overhead at 70B on A100 -- vs. Megatron's 8.4% -- reflects the fundamental tension between memory efficiency and communication efficiency. ZeRO-3 achieves its memory savings by distributing parameters across all data parallel ranks, requiring all-gather of full parameter shards for every forward pass. At large scale, this creates substantial inter-node traffic that can saturate InfiniBand links. H100's 2x InfiniBand bandwidth reduces this bottleneck substantially -- ZeRO-3 communication overhead drops from 38.4% to 18.4% on H100 at 70B. The implication is that cluster network bandwidth is as important as GPU count for distributed training efficiency.

5.3 Fault Tolerance Is Undervalued in Practice

APS's 4.2-minute live fault recovery vs. 16-22 minutes for checkpoint-based frameworks may seem like a minor optimisation. Compounded over a multi-week training run, it is significant: a 70B model training run of 30 days on 512 GPUs can expect 2-4 node failures based on historical MTBF for modern HPC nodes. At 4.2 vs. 18.4 minutes mean recovery time, APS saves 28-57 minutes of recovery time per training run -- not transformative, but at EUR 3.00/GPU-hour for 512 GPUs, each failure recovery costs EUR 384 with APS vs. EUR 1,690 with checkpoint restart. More importantly, live recovery eliminates the training disruption and queue wait time that makes long training runs fragile.

5.4 Limitations

Evaluations were conducted on two specific cluster configurations; different hardware (TPUs, AMD MI300X, Gaudi) would show different results. 400B scale was evaluated only on A100-512; H100 cluster was not large enough. APS profiling adds 10-step overhead at startup; for very short training runs, this overhead may not be justified. Throughput benchmarks use BF16 precision; FP8 training (available on H100) would shift the efficiency landscape significantly.

6. Conclusion

6.1 Summary

Eight distributed training frameworks were evaluated across four model scales. APS achieves the highest MFU at 70B (52.4% A100, 54.8% H100) through adaptive parallelism selection. Megatron-LM leads at 7B on high-bandwidth clusters (MFU 48.4%). ZeRO-3 enables 400B training on 512 A100s. APS achieves 4.2-minute fault recovery vs. 16-22 minutes for checkpoint-based frameworks. Communication overhead is the primary efficiency bottleneck at scale.

6.2 Recommendations

For 1B-7B training on NVLink clusters, use Megatron TP=8 for highest throughput. For 7B-70B with memory constraints, use ZeRO-3 or FSDP. For 70B+ or any long training run, use APS or manually tune TP+ZeRO hybrid. For training runs > 7 days on large clusters, deploy APS's live fault tolerance or implement equivalent checkpoint-based recovery. Upgrade to H100 if communication overhead exceeds 25% of step time

on A100 at your model scale -- 2x InfiniBand bandwidth typically reduces communication bottleneck to sub-20%.

References

- Chowdhery, A., Narang, S., Devlin, J. et al. (2022) PaLM: Scaling language modeling with pathways. arXiv preprint arXiv:2204.02311.
- Goyal, P., Dollar, P., Girshick, R. et al. (2017) Accurate, large minibatch SGD: Training ImageNet in 1 hour. arXiv preprint arXiv:1706.02677.
- Huang, Y., Cheng, Y., Bapna, A. et al. (2019) GPipe: Efficient training of giant neural networks using pipeline parallelism. *Advances in NeurIPS* 32, pp. 103-112.
- Lepikhin, D., Lee, H., Xu, Y. et al. (2021) GShard: Scaling giant models with conditional computation and automatic sharding. *Proceedings of ICLR* 2021.
- Li, S., Andriushchenko, M., Flammarion, N. et al. (2023) Colossal-AI: A unified deep learning system for large-scale parallel training. *Proceedings of ICPP* 2023.
- Narayanan, D., Shoeybi, M., Casper, J. et al. (2021) Efficient large-scale language model training on GPU clusters using Megatron-LM. *Proceedings of SC21*, pp. 1-15.
- Paszke, A., Gross, S., Massa, F. et al. (2019) PyTorch: An imperative style, high-performance deep learning library. *Advances in NeurIPS* 32, pp. 8024-8035.
- Rajbhandari, S., Rasley, J., Ruwase, O. and He, Y. (2020) ZeRO: Memory optimizations toward training trillion parameter models. *Proceedings of SC20*, pp. 1-16.
- Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J. and Catanzaro, B. (2019) Megatron-LM: Training multi-billion parameter language models using model parallelism. arXiv preprint arXiv:1909.08053.
- Smith, S., Patwary, M., Norick, B. et al. (2022) Using DeepSpeed and Megatron to train Megatron-Turing NLG 530B. arXiv preprint arXiv:2201.11990.
- Zhao, Y., Gu, A., Varma, R. et al. (2023) PyTorch FSDP: Experiences on scaling fully sharded data parallel. *Proceedings of VLDB* 2023, 16(12), pp. 3848-3860.

Declarations

Funding

This research was supported by the German Federal Ministry of Education and Research (BMBF) under grant 01IS25010A (DistTrainAI) and the Swedish Research Council (VR) under grant 2025-04284 (DistML-SE). Cluster access for evaluation was provided by the Juelich Supercomputing Centre (JSC) under project

AI-EUROHPC-2025-0042. The funders had no role in study design, evaluation, or manuscript preparation.

Conflict of Interest

The authors declare no conflicts of interest.

Data Availability Statement

APS framework code, benchmark scripts, and performance logs are available at <https://github.com/petrov-costa/aps-distributed> under Apache 2.0. Results deposited in Zenodo under DOI: 10.5281/zenodo.19614931. Released under CC BY 4.0.

Ethical Approval

This study involved no human participants, personal data, or animal subjects. Ethics committee approval was not required.

Appendix A

Parallelism Strategy Selection Framework

Select the distributed training strategy based on model size, cluster hardware, and training run duration.