

Paper Review - ReCycle: Resilient Training of Large DNNs using Pipeline Adaptation

Summary

This paper introduces ReCycle, a system for resilient training of large deep neural networks that sustains throughput during hardware/software failures without warm spares and without affecting model accuracy. The motivation is clear: multi-week training on thousands of GPUs encounters frequent faults, existing approaches either pay for idle spare servers, suffer static slowdowns from redundant computation (Bamboo), or incur heavy reconfiguration and pipeline imbalance (Ooblock). ReCycle's key idea is to exploit functional redundancy across data-parallel pipelines- peers at the same pipeline stage hold identical weights and to mine pipeline "bubbles" (short idle gaps at the start (warm-up) and end (cool-down) of a pipeline iteration) in a common 1F1B pipeline schedule to absorb extra work. When a worker fails, Adaptive Pipelining re-routes that stage's micro-batches to peer stages. To prevent throughput loss, Decoupled BackProp splits the backward pass: input-gradients run immediately to unblock dependencies, while weight-gradients are deferred into cool-down bubbles. A Staggered Optimizer is used to shift per-stage optimizer steps so next-iteration warm-up bubbles become additional cool-down. An offline Planner which uses dynamic programming and Mixed-Integer Linear Programming (MILP - a mathematical optimizer that schedules tasks under dependencies and memory limits) precomputes adaptive schedules for varying failure counts. ReCycle is implemented on DeepSpeed to detect failures, reroute micro-batches to data-parallel peers, decouple backprop, and bypass cross-stage optimizer syncs. Experiments on GPT-3-style models show higher sustained throughput under failures, beating Ooblock by up to 1.46x and Bamboo by up to 1.64x.

Evaluation

This paper tackles an important problem: keeping large training jobs running despite frequent faults, cutting wasted GPU-hours and avoiding spare servers. The solution is novel and practical, using cross-pipeline peers for fast recovery without re-shuffling parameters, plus two smart schedulers (Decoupled BackProp and Staggered Optimizer) that turn warm-up and cool-down bubbles into useful work. The evaluation is strong. It compares against Ooblock and Bamboo on GPT-3 models up to 6.7B on a 32 A100 cluster and uses real failure traces. A validated simulator extends results to 18.4B, 39.1B, 76.1B, and 145.6B models across clusters up to 1,536 GPUs and 64 pipeline stages, focusing on steady-state throughput under 1%, 5%, and 10% GPU failures. The results clearly show that ReCycle is well-suited for training tasks in supercomputing-scale clusters with dynamic resource availability. The paper is well written and clear, and the ablation study is especially useful.

Main takeaways

1. ReCycle keeps training going by re-routing failed stage work to data-parallel peers that already hold the same weights. This avoids warm spares and big reconfigurations.
2. Decoupled BackProp is used selectively, mainly on later pipeline stages with spare memory, so extra work is hidden in bubbles without exceeding GPU memory limits.

Strengths

1. Decoupled BackProp and the Staggered Optimizer turn warm-up and cool-down "bubbles" into compute slots, hiding most or all failure overhead in common 1F1B schedules.
2. The offline Planner precomputes failure-count specific schedules, so runtime recovery is immediate and does not reshuffle parameters. The Planner can handle very large jobs (for example 2048 GPUs, up to 25% failures) in under an hour, which is under 0.1% of multi-week training time.

Weaknesses

1. The Planner's schedules depend on accurate latency/memory profiles and the MILP's plan quality. If runtimes drift (contention, firmware changes) or MILP picks a suboptimal plan, throughput can drop.
2. ReCycle is built on DeepSpeed with custom reroute operations, so using it in other stacks (like PyTorch pipelines, ZeRO variants) may require extra engineering and slow adoption.

Discussion topic

1. If latency/memory profiles change over time (new drivers, mixed workloads), how often should the Planner reprofile or refresh MILP plans? What lightweight signals could be used to trigger a replan without stalling training?