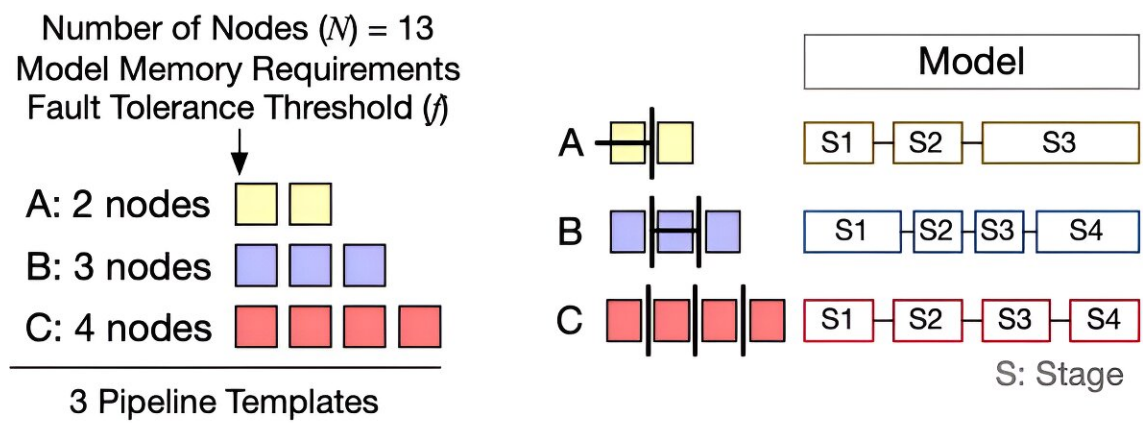
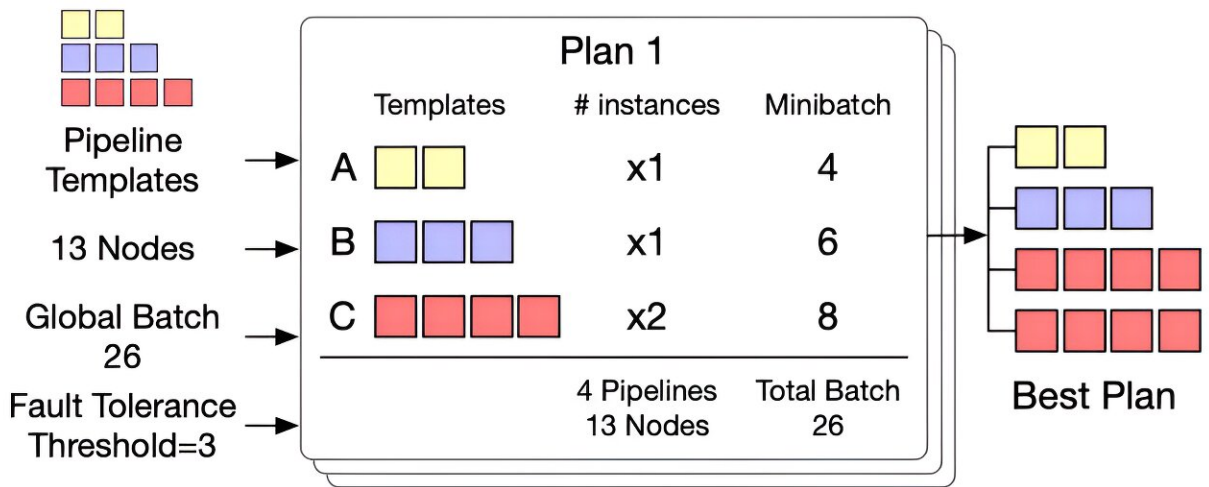


Open-source training framework increases the speed of large language model pre-training when failures arise

December 18 2023, by Patricia DeLacey



(a) Generating Pipeline Templates



(b) Pipeline Instantiation

Oobleck's planning algorithm overview. a) First, it generates a set of pipeline

templates, a combination of which can utilize all available nodes. b) Then, pipelines are instantiated following the fastest plan after checking all possible plans. Credit: Insu Jang, University of Michigan

As the demand for technologies that enable generative AI continues to skyrocket, processing capacities must keep pace to accommodate model training and fault tolerance. University of Michigan researchers designed a solution specific to modern AI workloads.

A research team developed Oobleck, an open-source large-model training framework, using the concept of [pipeline](#) templates to provide fast and guaranteed fault recovery without training throughput degradation.

The results were presented in October 2023 in the [Proceedings of the 29th Symposium on Operating Systems Principles](#) in Koblenz, Germany.

"Oobleck is a general-purpose solution to add efficient resilience to any large model pre-training. As a result, its impact will be felt in foundation model pre-training for the entire range of their applications from big tech and [high-performance](#) computing to science and medical fields," said Mosharaf Chowdhury, an associate professor of electrical engineering and computer science and corresponding author of the paper.

Large language models require massive GPU clusters for large durations during pre-training, and the likelihood of experiencing failures increases with the training's scale and duration. When failures do occur, the synchronous nature of large language model pre-training amplifies the issue as all participating GPUs idle until the failure is resolved.

Existing frameworks have little systemic support for [fault tolerance](#) during large language model pre-training. Current solutions rely on checkpointing or recomputation to recover from failures, but both methods are time-consuming and cause cluster-wide idleness during recovery with no formal guarantees of fault tolerance.

Pipeline templates are at the core of Oobleck's design. A pipeline template, a specification of training pipeline execution for a given number of nodes, is used to instantiate pipeline replicas. All pipeline templates are logically equivalent (i.e., can be used together to train the same model) but physically heterogeneous (i.e., use different numbers of nodes).

"Oobleck is the first work that exploits inherent redundancy in [large language models](#) for fault tolerance while combining pre-generated heterogeneous templates. Together, this framework provides high throughput with maximum utilization, guaranteed fault tolerance, and fast recovery without the overheads of checkpointing- or recomputation-based approaches," said Insu Jang, a doctoral student in computer science and engineering and first author of the paper.

Given a training job starting with the number of maximum simultaneous failures to tolerate, f , Oobleck's execution engine instantiates at least $f + 1$ heterogeneous pipeline from the generated set of templates. The fixed global batch is distributed proportionally to the computing capability of pipeline replicas to avoid having stragglers in training synchronization.

Upon failures, Oobleck simply re-instantiates pipelines from the precomputed pipeline templates, avoiding the demanding analysis of finding a new optimal configuration at runtime. It is provably guaranteed that using the precomputed set of pipeline templates always enables Oobleck to recover from f or fewer failures.

Resilience to unpredictable events is a classic problem in computer science. Instead of addressing problems after they happen, which is slow, or planning for all possible scenarios, which is practically impossible, pipeline templates strike a balance between speed and effectiveness in resilient distributed computing.

"Oobleck gives the first demonstration of the effectiveness of this idea, but it can potentially be applied to any distributed computing system where the same dichotomy exists. Going forward, we want to apply pipeline templates to improve the resilience of all facets of GenAI applications, starting with inference serving systems," said Chowdhury.

Oobleck is open-source and [available on GitHub](#).

More information: Insu Jang et al, Oobleck: Resilient Distributed Training of Large Models Using Pipeline Templates, *Proceedings of the 29th Symposium on Operating Systems Principles* (2023). [DOI: 10.1145/3600006.3613152](#)

Provided by University of Michigan College of Engineering

Citation: Open-source training framework increases the speed of large language model pre-training when failures arise (2023, December 18) retrieved 28 April 2024 from <https://techxplore.com/news/2023-12-open-source-framework-large-language-pre-training.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.