

Distributed Large-Scale Model Training and Fine-Tuning System Optimization with Error-Bounded Lossy Compression

Dr. Zhaorui ZHANG, Research Assistant Professor
The Hong Kong Polytechnic University



Abstract:

In recent years, with the rapid evolution of large-scale machine learning models, such as GPTs, Computer Vision models, and AlphaFold in life science, large-scale models have demonstrated strong performance across a wide range of domains and tasks, especially for emerging generative tasks. Along with these unprecedented capabilities, training such large-scale models comes with unique challenges. Communication has been identified as the main bottleneck when training large-scale models due to large-volume model parameters and gradient transmission across public networks with limited bandwidth (often less than 1 Gbps). Most previous works focus on gradient compression, while limited work tries to compress parameters that can not be ignored and extremely affect communication performance during the training. To bridge this gap, we propose an adaptive parameter update strategy for accelerating large-scale model training under limited network bandwidth across regions or countries. At its core are the state-of-the-art parameter compression library SZ3 and an adaptive compression error-bound calculation algorithm. It can explore the appropriate compression error bound adaptively during the training. It uses idle CPU resources to train clients to compress the parameters and reduce the communication overhead for each training round. Through extensive evaluation of multiple DNN models with different training settings on a GPU cluster with 65 GPUs and comparable network bandwidth to the real-world WAN (about 1 Gbps), we demonstrate that it can achieve the same model accuracy as vanilla approach while reducing up to $7.39\times$ and $288\times$ communication overhead for parameters and gradients, respectively, which can save the training time when the number of participant clients larger than 32 for the model with 4Gb size.

Biography:

Zhaorui Zhang is currently a research assistant professor in the Department of Computing at The Hong Kong Polytechnic University. She received her Ph.D. from the Department of Computer Science at The University of Hong Kong, Hong Kong, SAR, and her BSc degree in computer science from Xi'an Jiaotong University. Before joining HKU as a PhD student, she worked as a research assistant in the Department of Electrical Engineering at the City University of Hong Kong. Her research interests include distributed machine learning training and inference systems, distributed systems, high-performance computing, cloud computing, and data reduction.