DOCTORAL THESIS

Communication optimizations for distributed deep learning

Shi, Shaohuai

Date of Award:
2020

Abstract

With the increasing amount of data and the growing computing power, deep learning techniques using deep neural networks (DNNs) have been successfully applied in many practical artificial intelligence applications. The mini-batch stochastic gradient descent (SGD) algorithm and its variants are the most widely used algorithms in training deep models. The SGD algorithm is an iterative algorithm that needs to update the model parameters many times by traversing the training data, which is very time-consuming even using the single powerful GPU or TPU. Therefore, it becomes a common practice to exploit multiple processors (e.g., GPUs or TPUs) to accelerate the training process using distributed SGD. However, the iterative nature of distributed SGD requires multiple processors to iteratively communicate with each other to collaboratively update the model parameters. The intensive communication cost easily becomes the system bottleneck and limits the system scalability.

In this thesis, we study the communication-efficient techniques for distributed SGD to improve the system scalability and thus accelerate the training process. We identify the performance issues in distributed SGD through benchmarking and modeling and then propose several communication optimization algorithms to address the communication issues.

First, we build a performance model with a directed acyclic graph (DAG) to modeling the training process of distributed SGD and verify the model with extensive benchmarks on existing state-of-the-art deep learning frameworks including Caffe, MXNet, TensorFlow, and CNTK. Our benchmarking and modeling point out that existing optimizations for the communication problems are sub-optimal, which we need to address in this thesis.

Second, to address the startup problem (due to the high latency of each communication) of layer-wise communications with wait-free backpropagation (WFBP), we propose an optimal gradient merging solution for WFBP, named MG-WFBP, that exploits the layer-wise property to well overlap the communication tasks with the computing tasks and can be adaptive to the training environments. Experiments are conducted on dense-GPU clusters with Ethernet and InfiniBand, and the results show that MG-WFBP can well address the startup problem in distributed training

of layer-wise structured DNNs.

Third, to make the high computing-intensive training tasks be possible in GPU clusters with low-bandwidth interconnect, we investigate the gradient compression techniques in distributed training. The top-k sparsification can well compress the communication traffic with little impact on the model convergence, but it suffers from a linear communication complexity to the number of workers so that top-k sparsification cannot scale well in large-scale clusters. To address the problem, we propose a global top-k (gTop-k) sparsification algorithm that reduces the communication complexity to be logarithmic to the number of workers. We also provide detailed theoretical analysis for the gTop-k SGD training algorithm, and the theoretical results show that our gTop-k SGD has the same order of convergence rate with SGD. Experiments are conducted on up to 64-GPU cluster to verify that gTop-k SGD significantly improves the system scalability with only a slight impact on the model convergence.

Lastly, to enjoy the both benefits of the pipelining technique and the gradient sparsification algorithm, we propose a new distributed training algorithm, layer-wise adaptive gradient sparsification SGD (LAGS-SGD), which supports layer-wise sparsification and communication, and we theoretically and empirically prove that the LAGS-SGD preserves the convergence properties. To further alliterate the impact of the startup problem of layer-wise communications in LAGS-SGD, we also propose the optimal gradient merging solution for LAGS-SGD, named OMGS-SGD, and theoretical prove its optimality. The experimental results on a 16-node GPU cluster connected 1Gbps Ethernet show that OMGS-SGD can always improve the system scalability while the model convergence properties are not affected.

Keywords: Deep Learning, Distributed Deep Learning, Communication Optimizations, SGD, Distributed SGD, Gradient Sparsification

Table of Contents