

Simulating Performance of ML Systems with Offline Profiling

Hongming Huang*, Peng Cheng[†], Hong Xu*, Yongqiang Xiong[†]

*Department of Computer Science, City University of Hong Kong

[†]Microsoft Research

ABSTRACT

We advocate that simulation based on offline profiling is a promising approach to better understand and improve the complex ML systems. Our approach uses operation-level profiling and dataflow based simulation to ensure it offers a unified and automated solution for all frameworks and ML models, and is also accurate by considering the various parallelization strategies in a real system.

1 INTRODUCTION

Machine learning (ML) systems are prevalently used to train the complicated deep neural network (DNNs) [5–7, 15] using a cluster of machines. In this context, characterizing the performance of an ML system becomes critical for optimizing training strategy [9, 10, 13] and system design [8, 12]. Intuitively, performance can be obtained by actually training a given DNN model over a given hardware cluster with a given strategy (e.g. parallelization, hyperparameters, etc.) and measuring its throughput. This *online profiling* approach is the de facto solution used by most existing work [9, 13].

Online profiling yields the most accurate performance. Yet it has several fundamental limitations that we believe make it ill-fitted for many practical scenarios.

Online profiling is very expensive and does not scale. While it is feasible to run a few iterations of training for a given setup, the exponential number of potential training strategies—a combination of hyperparameter setting, parallelization strategy, synchronization and pipelining method, etc.—makes it impractical to enumerate all possibilities in order to find the optimal one. Then if we have a new DNN, we have to do profiling all over again for this new model. Further, online profiling is inherently limited by the available hardware resources at our disposal. If we wanted to know the accurate system performance with a new computing or networking device, we would have to acquire this hardware first which takes significant time and financial investments. Finally, due to the complexity of ML systems, online profiling is typically done in an end-to-end fashion by treating the entire execution pipeline as a blackbox. Thus it remains difficult today to dissect and understand the impact of various aspects of the system (say computation vs communication) in different settings (GPUs, networking (PCI-e, NV-Link, RDMA, ...), architecture (PS [11], All-reduce [4]), ...), let alone how to improve the design.

We propose a different approach to address these issues. We rely on *offline profiling* to measure performance of basic execution units on different hardware platforms, and based upon it build a simulator to accurately characterize the system-level performance without actually training the DNN. Since we profile the basic execution units on tensors including computation (e.g. conv2d) and communication (e.g. allreduce) instead of the entire system, offline profiling is much more scalable and the results can naturally be reused. Different

users can easily contribute their profiling results on their hardware platforms, thus overcoming the hardware constraint. More importantly, instead of a blackbox approach, our simulator uses the detailed dataflow graph produced by all major ML systems to accurately estimate its performance for a given training strategy.

Offline profiling based simulation has promising potentials to become a foundation for many tasks in MLOps. It accelerates system design, since we can quickly identify the performance bottleneck of the complex ML system and project the potential gain of a certain optimization. It also helps many auto ML and performance engineering tasks, as systems like PipeDream [13] and FlexFlow [9] can use it to rapidly find the optimal parallelization strategy for any DNN, hardware, and hyperparameter settings without the high overheads of online profiling.

2 DESIGN

Building a performance simulator entails two basic questions: (1) how to profile the basic execution units of a DNN offline, and (2) how to simulate the system-wide performance using the profiling results? Answering them is challenging because we have to achieve three objectives at the same time: (1) the simulator should be *accurate* in order to be useful; (2) it should be *general* in order to cover the major dimensions of ML training: framework, library, hardware, training strategies, etc.; (3) it should also be *efficient* with minimal human intervention, preferably completely automated.

In the following we explain our key design decisions in answering the two questions and how they achieve our design goals.

Op-level profiling. There are two types of basic execution units for a ML system. One is operation (op) which is the atomic execution unit defined at the framework level (e.g. TensorFlow [3]). The other is kernel which is defined at the device level (e.g. CUDA for GPU). Op is a higher level abstraction and is implemented by one or more kernels. We perform profiling at the op-level that achieves a good trade-off between accuracy and overhead. We do not favor kernel-level profiling because it requires precise knowledge of how each op is implemented in the framework using different kernels. This can only be obtained by analyzing the framework’s source code which is extremely expensive and hard to automate.

Even for a single op, it may have many input arguments that makes it difficult to numerate all possible combinations. For example, the Conv2D op in TensorFlow has four arguments for input tensor shape, four for the filtering kernel shape, and several other arguments for other attributes of the 2-D convolution [2]. Thus we apply a machine learning approach here to reduce the complexity: for each input argument we profile a fixed number of values, and use these results to train a neural network to estimate the op performance. This is also easy to automate.

Dataflow based simulation. Another key design choice is to use the dataflow graph as the basis of our simulation. Dataflow graph is




Figure 1: System Overview.

widely used by all major ML frameworks [1] as the blueprint of execution and is generated automatically according to the user training program. It is a directed acyclic graph containing the ops as nodes and input/output data as edges between the nodes. It also represents various parallelism optimization and distributed execution strategies across devices and machines. Every node on the graph can be placed on a different machine, which is indicated in its “device” property, and the synchronization requirements are modeled by dependencies (i.e. edges) across computation and communication ops[14]. All this information is embedded in the dataflow graph which can be readily obtained with the framework’s API.

Thus our simulator essentially works by “replaying” the training execution based on the dataflow graph to calculate performance. Each independent device (CPU, GPU, or communication link) executes in parallel and maintains a job queue and its finish time. The simulator keeps a global ready list containing all nodes whose dependencies are fulfilled. The simulator runs in a loop: (1) It starts all nodes in the ready list by enqueueing them into their corresponding device’s job queues. (2) As soon as an op is finished on a device (using the profiling results), it updates all successor nodes’ dependency counter. If the counter becomes zero, the successor node is added into ready list. The system performance is obtained by looking at the finish time of the last device.

Overall System. With these key design choices, our simulator’s overall design is shown in Figure 1. As each framework expresses dataflow graph differently, the *preprocessing module* first transforms the dataflow graph extracted from the framework into a unified format. The *op estimator* estimates performance of each op in the dataflow graph by querying the *profiling database* with offline profiling results. This requires information about the training environment (e.g. hardware type, software library version, etc.) from a config file which is not in the graph. In case the graph has new ops not in the profiling database, we fall back to online profiling with the *new op profiler* and add the result to the database. Then the *simulation module* takes the augmented graph, and simulates the training execution accordingly and estimates performance as explained before. It also needs additional information about the training strategy from a config file, such as the number of replicas in data parallelism, and the pipelining setting for model parallelism which may not be available in the dataflow graph [13].

3 PRELIMINARY RESULTS

We present our preliminary evaluation results here.

Experiment Setup. We run experiments on a server with 2 Xeon CPU E5-2620 and 4 Tesla V100 GPU, running Ubuntu 16.04.5 LTS with CUDA 10.0, cuDNN 7.5, and TensorFlow 1.13.1.

Offline profiling. We build an automatic profiler with about 900 LoC in PyThon to profile computation ops. It constructs a dataflow




Figure 2: Performance of Conv2D on our testbed.

Scenario	QPI	Root Complex	PCI-e Switch
GPU to GPU uni-direct	10948.81	10270.59	13181.03
GPU to GPU bi-direct	16475.93	19325.81	25037.41
Host to GPU	11956.69	12027.22	12347.09
GPU to Host	13200.21	13201.87	13121.95
NCCL-AllReduce 2GPUs	6178.68	9162.42	11598.12
NCCL-AllReduce 4GPUs	—	—	8048.35

Table 1: Throughput (MB/s) of GPU communication on a single machine in different scenarios.

Model	TF . timeline (ms)	Simulation (ms)	Error
VGG_19	203.38	199.66	1.83%
ResNet_50	282.56	277.47	1.80%
ResNet_152	638.81	629.31	1.49%

Table 2: Per-iteration training time with batchsize 64 on CIFAR-10.

graph that only contains the input data nodes and 1000 identical computation nodes corresponding to the op. This is to amortize the constant overheads of launching the graph and input initialization on GPUs before training starts. As discussed in §2, we profile each input argument of the op with 16 possible values.

Our initial experiment profiles over 20 common ops for neural network building (Relu, Sigmoid, Conv2DBackpropFilter, etc.), linear algebra (MatMul), and element-wise mathematics. We observe that their performance is very stable with standard error lower than 1% of the mean, and has a strong linear relationship to the input shape. Figure 2 depicts the performance of Conv2D with varying number of input channels as an example. This verifies our design choice that we can model and accurately estimate op-level performance with offline profiling.

We also profile the GPU communication bandwidth in a single machine under different scenarios. Table 1 shows the result.

Simulation. We use three common CNN models to evaluate the accuracy of dataflow-based simulation. We use the TF . timeline to measure the actual training time. Since our profiling is not complete yet, we use the offline profiling results whenever applicable, and rely on TF . timeline to do online profiling for other ops.

Table 2 shows that our dataflow based simulation is accurate with <2% errors. The errors mainly come from memory copy and the time gap between ops that we have not considered.

4 FUTURE WORK AND CONCLUSION

We advocated offline profiling based simulation for ML systems in this work. To fully demonstrate its feasibility and potential, we are conducting extensive offline profiling to cover all operations in popular frameworks and common hardware, and investigating the ML approach for op-level performance estimation. We will also validate the effectiveness of our approach on distributed settings with multiple machines and various networking technologies. We aim to develop a fully automated simulator together with the profiling database and open source it for the community.

REFERENCES

- [1] [n.d.]. mmDNN. <https://www.microsoft.com/en-us/research/uploads/prod/2019/11/mmdnn.pdf>.
- [2] [n.d.]. tensorflow::ops::Conv2D. https://www.tensorflow.org/api_docs/cc/class/tensorflow/ops/conv2-d.
- [3] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*. 265–283.
- [4] Ammar Ahmad Awan, Khaled Hamidouche, Akshay Venkatesh, and Dhaleswar K Panda. 2016. Efficient large message broadcast using NCCL and CUDA-aware MPI for deep learning. In *Proceedings of the 23rd European MPI Users' Group Meeting*. ACM, 15–22.
- [5] Ciprian Abella, Tomas Mikolov, Mike Schuster, Qi Ge, Thorsten Brants, Phillip Koehn, and Tony Robinson. 2013. One billion word benchmark for measuring progress in statistical language modeling. *arXiv preprint arXiv:1312.3005* (2013).
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.
- [7] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4700–4708.
- [8] Zhihao Jia, Sina Lin, Charles R Qi, and Alex Aiken. 2018. Exploring hidden dimensions in parallelizing convolutional neural networks. *arXiv preprint arXiv:1802.04924* (2018).
- [9] Zhihao Jia, Matei Zaharia, and Alex Aiken. 2018. Beyond data and model parallelism for deep neural networks. *arXiv preprint arXiv:1807.05358* (2018).
- [10] Alex Krizhevsky. 2014. One weird trick for parallelizing convolutional neural networks. *arXiv preprint arXiv:1404.5997* (2014).
- [11] Mu Li, David G Andersen, Jun Woo Park, Alexander J Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J Shekita, and Bor-Yiing Su. 2014. Scaling distributed machine learning with the parameter server. In *11th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 14)*. 583–598.
- [12] Azalia Mirhoseini, Hieu Pham, Quoc V Le, Benoit Steiner, Rasmus Larsen, Yuefeng Zhou, Naveen Kumar, Mohammad Norouzi, Samy Bengio, and Jeff Dean. 2017. Device placement optimization with reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2430–2439.
- [13] Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R. Devanur, Gregory R. Ganger, Phillip B. Gibbons, and Matei Zaharia. 2019. PipeDream: Generalized Pipeline Parallelism for DNN Training. In *Proc. ACM SOSP*.
- [14] Yanghua Peng, Yibo Zhu, Yangrui Chen, Yixin Bao, Bairen Yi, Chang Lan, Chuan Wu, and Chuanxiong Guo. 2019. A generic communication scheduler for distributed DNN training acceleration. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. 16–29.
- [15] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 2818–2826.