

Improving the performance of distributed TensorFlow with RDMA

Chengfan Jia, Hong An

University of Science and Technology of China,
Hefei, Anhui 230026,
jcf94@mail.ustc.edu.cn

Abstract. TensorFlow is an open-source software library designed for Deep Learning using data flow graph computation. Thanks to the flexible architecture of TensorFlow, users can deploy computation to one or more CPUs or GPUs in a desktop, server, or mobile device with a single API. In a distributed TensorFlow work process, it uses gRPC to connect between different nodes. However, when deploying training tasks on high performance computing clusters, the performance of gRPC becomes a bottleneck of distributed TensorFlow system. HPC clusters are usually equipped with Infiniband network besides of traditional TCP/IP network. But open-sourced TensorFlow has not taken this advantage. We present a RDMA-capable design of TensorFlow. By porting the Tensor send/receive parts of TensorFlow into RDMA verbs, we finally get nearly 6x performance improvements over the original distributed TensorFlow based on gRPC. The TensorFlow system with RDMA support shows a great scalability among the training scale.

Keywords: Distributed TensorFlow, RDMA, Infiniband, Optimization

1 Introduction

In recent years, Machine Learning has driven advances in many different fields. Moreover, benefit from the rapid growth of computing power, Deep Learning has drawn wider and wider attention of researchers and engineers.

TensorFlow [1] is a numerical computation system designed for Machine Learning, which supports a variety of applications, with a focus on training and inference on deep neural networks. It can efficiently support large-scale training and inference as well as flexibly support experimentation and research into new models and system-level optimizations.

Deep Learning process usually requires large amount of training data and powerful computing recourses to deal with the data. With increasing size of the datasets, single compute node is highly restrictive.

Some studies about it build frameworks over single node TensorFlow to achieve good scalability, such as Spark based framework [2] or MPI based distributed system [3]. The official Distributed TensorFlow is open-sourced on April, 2016.

1.1 Problem Statement

Training a Neural Network with multiple hidden layers is usually a time consuming work. For example, the champion algorithm of 2012 ImageNet Large Scale Visual Recognition Challenge, which is called Alexnet [4] later, consists of five convolutional layers and three fully connected layers with nearly 60 million parameters. According to the paper, the training process for roughly 90 cycles through the training set of 1.2 million images took five to six days on two NVIDIA GTX 580 3GB GPUs. Despite of the sustained growing computing power, greater network will still show a press need in the future.

However, the open-source TensorFlow did not seem to show advantage in performance as we expected. A recent research from HKBU [5] tested several popular deep learning frameworks, TensorFlow did not perform well in their test results.

In order to evaluate the distributed TensorFlow, we wrote a distributed Alexnet benchmark based on the Alexnet CPU version provided by HKBU dl-bench project. Data scalar of this benchmark is as same as the original algorithm while the input data replaced by random pixels.

Test is performed in a 8-node CPU cluster. Each of it has dual Intel Xeon E5-2660 processors with 1000M Ethernet and Mellanox 40G QDR InfiniBand. The test count for 300 steps' elapsed time with each step of 24 RGB "pictures" of 224*224 pixels as input data.




Fig. 1. Test results of Alexnet in Distributed TensorFlow using TCP/IP

Figure 1 shows our test result. Single node TensorFlow costs 224 second, while the elapsed times of the same data size in the distributed version are much more than that. Besides, expanding from 2 workers to 3 workers (2 compute nodes to 3 compute nodes) shows a bad scalability. It is obvious that communication with gRPC through TCP/IP becomes a bottleneck of the distributed system.

1.2 Contributions

Aiming to improve the performance of distributed TensorFlow system and make use of the Infiniband network deployed in our HPC clusters, this paper presents a RDMA-capable design of TensorFlow, which, by porting the Tensor send/receive parts into RDMA verbs, achieves great performance improvement over the original gRPC version.

The rest of the paper is organized as follows:

Background material is presented in Section 2. Section 3 presents our contribution to TensorFlow. We first analyze the workflow of TensorFlow’s send/receive parts, and then design a RDMA protocol to fit into the distributed TensorFlow. Section 4 shows the evaluations of our work. Section 5 concludes our work.

2 Background Review

2.1 TensorFlow Architecture

TensorFlow uses dataflow graph to organize its computations and provides a friendly interface for users. Master translates user request into execution. Then the dataflow executor handles requests from the master, and schedules the execution of the kernels that comprise graph. The runtime of it contains over 200 standard operations, computation works can be efficiently done with the control of the executor.

Each operation resides on a particular device, such as a CPU or GPU in a particular task. A device is responsible for executing a kernel for each operation assigned to it.

Distributed TensorFlow works similar to multi-device execution [6]. After device placement, a subgraph is created per machine. Send/Receive node pairs that communicate across worker processes use remote communication mechanisms to move data across machine boundaries. Googles paper mentioned that their work both support TCP and RDMA while the open-source version we get now only implements gRPC to organize the communication.

Several studies mentioned in section 1 avoid the use of gRPC, they try to build a higher level scheduling framework and data gather/scatter framework (by Spark or MPI framework) over non-distributed TensorFlow.

2.2 gRPC

gRPC is an open-source framework from Google for handling remote procedure calls [7]. It is based on the HTTP/2 standard, and enables easy creation of highly performant, scalable APIs and micro services in many popular programming languages and platforms.

gRPC uses *“completion queue”* to manage the performing of multi-threaded task, and the execution of every task is associated with the *“grpc_exec_ctx”* closure mechanism. In order to asynchronously establish the connection and transport data, gRPC uses *“pollset”* to manage input and output file descriptors.

By polling the file descriptors of each "pollset" using *epoll*, once TCP socket start to connect in, data come in from TCP or remote TCP connection is going to shutdown, the callback functions which are combined with specific file descriptors will be called to handle these different conditions.

Our *previous work* (finished last year) tried to port gRPC into RDMA, aiming to improve the performance of gRPC and indirectly speed up the distributed TensorFlow. The result of it is about to double the performance of the original TCP version.

2.3 Infiniband and RDMA

The InfiniBand Architecture [8] defines a switched network fabric for interconnecting processing nodes and I/O nodes. It provides a communication and management infrastructure for inter-processor communication and I/O. In an InfiniBand network, hosts are connected to the fabric by Host Channel Adapters (HCAs). InfiniBand uses a queue-based model. A Queue Pair in InfiniBand consists of two queues: a send queue and a receive queue. The send queue holds instructions to transmit data and the receive queue holds instructions that describe where received data is to be placed. Communication operations are described in the Work Queue Requests (WQR), or descriptors, and submitted to the work queue. The completion of WQRs is reported through Completion Queues (CQs). Once a work queue element is finished, a completion entry is placed in the associated completion queue. Applications can check the completion queue to see if any work queue request has been finished.

InfiniBand Architecture supports both channel semantics and memory semantics. In channel semantics, send/receive operations are used for communication. In memory semantics, InfiniBand provides Remote Direct Memory Access (RDMA) operations, including RDMA Write and RDMA Read. RDMA operations are one-sided and do not incur software overhead at the remote side. This enables true application bypass message passing. The processor on the machine can continue its computation task without bothering about incoming messages. Thus, RDMA can positively impact the computation and communication overlap. Additionally, RDMA Write operation can gather multiple data segments together and write all data into a contiguous buffer at the receiver end. RDMA Write with Immediate data is also supported. With Immediate data, a RDMA Write operation consumes a receive descriptor and then can generate a completion entry to notify the remote node of the completion of the RDMA Write operation.

IB verbs is an implementation of the RDMA verbs for Infiniband (according to the Infiniband specifications). It handles the control path of creating, modifying, querying and destroying resources such as Protection Domains (PD), Completion Queues (CQ), Queue-Pairs (QP), Shared Receive Queues (SRQ), Address Handles (AH) and Memory Regions (MR). It also handles sending and receiving data posted to QPs and SRQs, getting completions from CQs using polling and completion events.

The control path is implemented through system calls to the ib verbs kernel module which further calls the low-level hardware driver. The data path is implemented through calls made to low-level hardware library which, in most cases, interacts directly with the hardware and provides kernel and network stack bypass (saving context/mode switches) along with zero copy and an asynchronous I/O model.

3 Porting TensorFlow into RDMA

3.1 Send/Receive process in TensorFlow

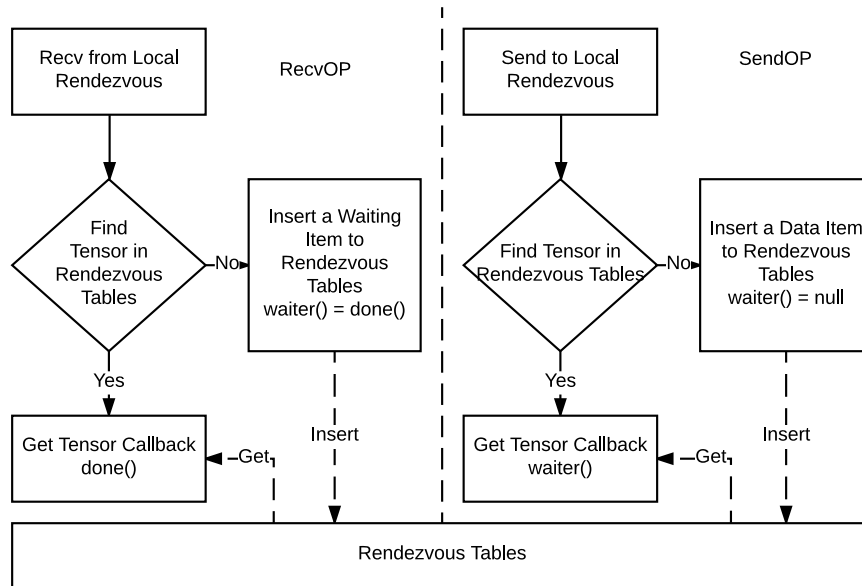


Fig. 2. Single node workflow of Send/Recv operations

Dataflow implementation simplifies distributed execution, as it makes communication between subcomputations explicit. TensorFlow runtime places operations on devices, subject to implicit or explicit constraints in the graph.

Once the operations in a graph have been placed, and the partial subgraph has been computed, TensorFlow partitions the operations into per-device subgraphs. A per-device subgraph for device d contains all of the operations that were assigned to d , with additional *Send* and *Recv* operations that replace edges across device boundaries. *Send* transmits its single input to a speied device as soon as the tensor is available, using a rendezvous key to name the value. *Recv* has a single output, and blocks until the value for a speied rendezvous key is available locally, before producing that value.

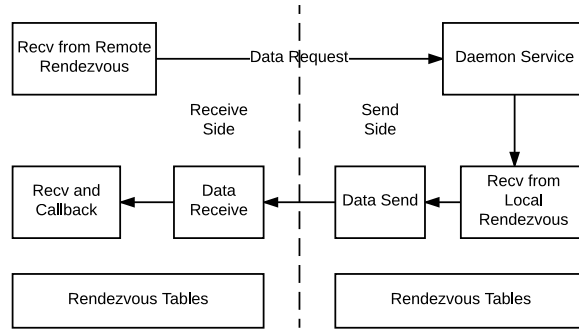


Fig. 3. Cross node workflow of Send/Recv operations

Figure 2 shows the workflow of *Send/Recv* operations within a same machine. The Tensor data transmitted by *Send/Recv* operations are indexed in an object called *Rendezvous*. *Recv* operation first check the rendezvous key of the value in *Rendezvous Tables*. If rendezvous key exists (inserted by a previous *Send* operation), just get the Tensor data and return it by a callback function. Otherwise, it inserts a waiting item to the *Rendezvous Tables*, which contains the requiring rendezvous key as well as saves the callback function as a waiter. When the *Send* operation is processed, it will check the *Rendezvous Tables*, too. In existing of the rendezvous key (inserted by a previous *Recv* operation), just get the waiting item and run the waiter function with the sending Tensor data. Otherwise, it inserts a data item to the *Rendezvous Tables*, which contains the sending rendezvous key and the sending Tensor data.

We should notice that *Send/Recv* operations within a single machine (run in the non-distributed TensorFlow) use memory (*Rendezvous* object) to perform the communication, while the network communications occur in distributed environment. Figure 3 shows such a workflow.

Communication starts by the receive side. *Recv* operation sends a *Data Request* to the send side asking for Tensor data. The waiting service of the send side performs a *Recv* operation to get data from local *Rendezvous* object. Once the Tensor is ready, it responses the data to the receive side. Receive side then runs the callback function to finish the remote *Recv* operation.

The *Data Request* and *Data Response* process is done by network communications, or more specifically gRPC here.

3.2 RDMA process design

Distributed TensorFlow uses rendezvous mechanism to organize its communication process.

Typically, in a Rendezvous Protocol, the sender and the receiver negotiate the buffer availability on both sides before the message transfer actually takes place. For achieving high performance message passing for large messages, it is critical that message copies are avoided. The Rendezvous Protocol provides a

way to achieve zero-copy message transfer because sender can know the location of the receivers buffer or vice-versa.

Several studies about RDMA have shown good performance in achieving computation/communication overlap with fast memory registration based RDMA Write Process [9] or RDMA Read Based Rendezvous Protocol [10].

To modify the communication process into RDMA, we design a RDMA read based protocol and a RDMA write based protocol to replace the original gRPC calls.

After the connection of Master Service and Worker Service in distributed TensorFlow cluster, the RDMA service of each side get connected. RDMA contexts (like lid, qpn, psn and so on) are exchanged through gRPC calls. The basic protocol is illustrated in Figure 4.

In the RDMA read protocol (Figure 4a), communication starts by the receive side. *Recv* operation sends a *Tensor Request* message to the send side asking for Tensor data. The waiting service of the send side performs a *Recv* operation to get data from local *Rendezvous* object, and at the same time, prepares a memory buffer for upcoming RDMA operations including buffer creations and memory bind registrations. Once the Tensor is ready, it responses the *data size*, *data address* and *rkey* to the receive side. Receive side then prepares a corresponding receive buffer and calls RDMA read to get the Tensor Data from remote memory. After the completion of RDMA read process, receive side then runs the callback function to finish the *Recv* operation as well as sends a *Buffer Release* message to inform the send side.

RDMA write protocol (Figure 4b) is similar to the process described above. Read protocol uses less interacting messages, while write protocol overlaps communication time and buffer register time in each side. In our evaluations, these two protocols acquire almost a same performance after further optimizations.

Just at the time we finish this paper, Yahoo submits a pull request about their RDMA version to TensorFlow and it is accepted in TensorFlow 1.2.0-rc0 on May 20th, 2017. Their design about IB verbs is similar to our RDMA write protocol.

3.3 Memory optimization

Studies about memory management of RDMA process [11] have shown the importance of hidden the memory manage costs.

One of the key differences between RDMA and TCP/IP communication is that the application has to explicitly manage the memory segments that will be used as communication buffers. The application has to preregister certain parts of its memory with the RDMA subsystem as source and/or destination buffers for the data transfers. During registration, the memory pages get pinned by the OS making sure they stay resident and cannot be swapped out to disk. The pinned pages are then registered with the RDMA enabled NIC (RNIC) so that it can access them using DMA operations which eliminates the need for OS

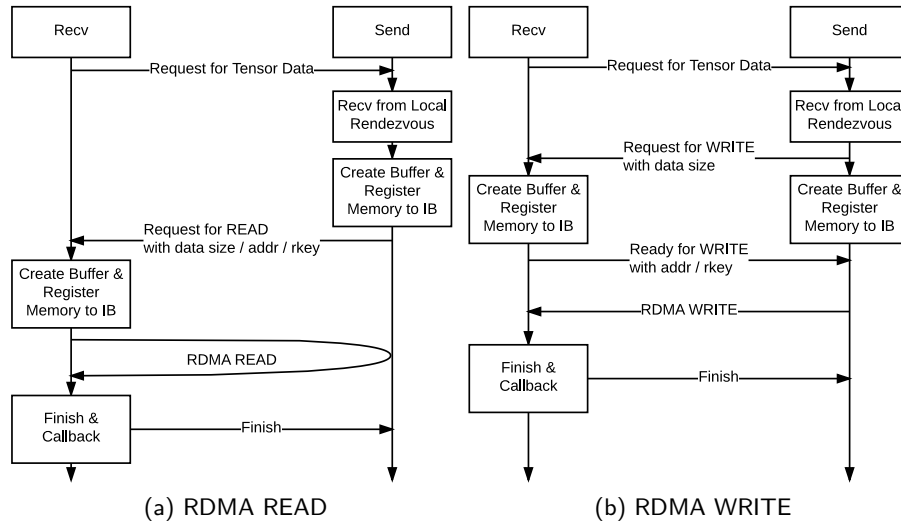


Fig. 4. RDMA protocol for distributed TensorFlow

callbacks and intermediate buffering during transfers. The RDMA syntax refers to these registered memory segments as Memory Regions (MR).

MR registration happens through the resource management path which requires kernel activity and therefore induces a delay as well as a non-negligible CPU load. Even though the expensive data copy operations are avoided with RDMA (zero copy), the explicit buffer management renders RDMA useless for applications that are not able to deal with their buffers efficiently.

Messages buffer reuse: Each side of the RDMA pair has a fixed message send buffer and message receive buffer, local send buffer is associated with the remote receive buffer when exchanging the RDMA context. The messages used during interactions have similar structures, so when sending another message we only rewrite the different parts of the message and make the same part of the buffer reused.

Tensor Data buffer reuse: During the Deep Learning training process or inferencing process, different input data goes through a same data path, which means a Tensor of the data flow graph will be sent and received many times during the whole process. MR registration is a time consuming operation in both RDMA read and write protocols, but these times can be cut off by memory reuse. Each time when a new block of memory is registered to MR, we store the information of it (address and rkey) to a memory table by Tensor name, Tensor source and Tensor Destination. The buffer stored in memory table will never be released when a pair of Send/Recv operations are completed. Next time when a same Tensor is sent or received, we can just get the Data buffer from the memory table instead of register a new block of memory.

Memory pool implementation: The Data buffer reuse method described above however has a disadvantage of costing too much memory. We then come up with a memory pool implementation to reuse buffer and cut down memory usage. After the connection of RDMA pair, we allocate a big block of memory from OS and register it to Infiniband. The memory is used as a memory pool. When *Send* operation or *Recv* operation require a data buffer, we can just get memory from the big block. Special "malloc" and "free" methods are used for allocating memory from the memory pool and releasing memory to the memory pool.

4 Evaluation

In this section we evaluate our proposed designs for the RDMA optimized distributed TensorFlow. Our platforms for evaluation is:

Cluster A: 8 SuperMicro nodes with dual Intel Xeon E5-2660 processors. Each node has 128GB of main memory. The nodes are connected with 1000M Ethernet and Mellanox 40G QDR InfiniBand.

Cluster B: 6 SuperMicro nodes with dual Intel Xeon E5-2698 processors. Each node has 128GB of main memory. The nodes are connected with Mellanox 100G EDR infiniband.

We expect that the RDMA optimization reduces the communication latency in the distributed TensorFlow system and improves its performance.




Fig. 5. Test results of Alexnet in Distributed TensorFlow in Cluster A

Figure 5 is the results of the Alexnet test case described in section 1.1. gRPC through ipoib shows a 3 to 4 times speed up comparing to the original TCP/IP version. This benefit comes from the hardware (Infiniband to Ethernet).

During the TCP/IP test, it can be observed clearly that CPU usages of the computer nodes are very low. Moreover, the CPU usages even periodically drop to less than 10% (when waiting for communication). This bottleneck is caused by communication latency of TCP/IP gRPC. When using ipolib or RDMA optimized version, CPU usages can be kept at a high level.

Our work over RDMA is about 30% to 40% faster than ipolib gRPC, and almost 6 times faster than TCP/IP gRPC.

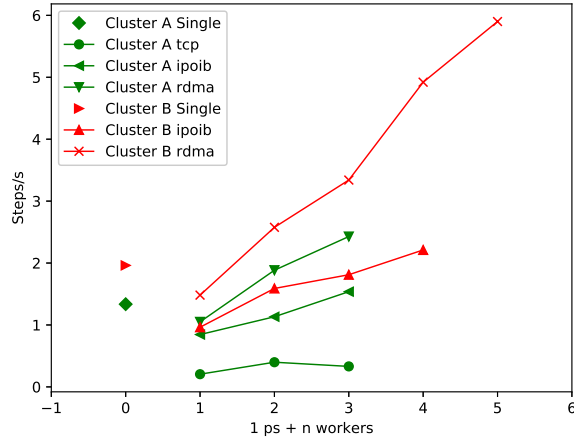


Fig. 6. Test results of Alexnet in Distributed TensorFlow in Cluster A

Further test in Figure 6 describes steps per second the TensorFlow system can train, which shows that our RDMA optimized distributed TensorFlow has a great scalability. Each step of Alexnet test case input 24 RGB pictures of 224*224 pixels, the data path includes inference, gradients computing and parameters updating. With the help of powerful CPU and Infiniband, RDMA optimized distributed TensorFlow acquires nearly linear growth in performance.

5 Conclusion

This paper presents a RDMA-capable design of TensorFlow, which, by porting the Tensor send/receive parts into RDMA verbs, achieves great performance improvement over the original TCP/IP gRPC version.

We analyse the workflow of distributed TensorFlow and especially focus on the send/receive parts of it. Then we design a RDMA read based and a RDMA write based rendezvous protocol to fit into the distributed TensorFlow workflow. Several optimizations are used. In our evaluation, the RDMA optimized

distributed TensorFlow shows a good performance over the original TCP/IP version as well as a good scalability with the increasing of compute nodes.

TensorFlow is a well designed framework which uses data flow model to organize its computation. Benefiting from the data flow model, we believe that it has the potential to gain more performance after further optimizations.

Acknowledgment

This research is conducted under Advanced Computer System Architecture (ACSA) Laboratory of University of Science and Technology of China.

References

1. Abadi M, Barham P, Chen J, et al. TensorFlow: A system for large-scale machine learning[J]. arXiv preprint arXiv:1605.08695, 2016.
2. Kim H, Park J, Jang J, et al. DeepSpark: A Spark-Based Distributed Deep Learning Framework for Commodity Clusters[J]. arXiv preprint arXiv:1602.08191, 2016.
3. Vishnu A, Siegel C, Daily J. Distributed TensorFlow with MPI[J]. arXiv preprint arXiv:1603.02339, 2016.
4. Krizhevsky A, Sutskever I, Hinton G E. Imagenet classification with deep convolutional neural networks[J]. Advances in neural information processing systems, 2012: 1097-1105.
5. Shi S, Wang Q, Xu P, et al. Benchmarking State-of-the-Art Deep Learning Software Tools[J]. arXiv preprint arXiv:1608.07249, 2016.
6. Abadi M, Agarwal A, Barham P, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems[J]. arXiv preprint arXiv:1603.04467, 2016.
7. Google Developers. Introducing gRPC, a new open source HTTP/2 RPC framework. <http://googledevelopers.blogspot.com/2015/02/introducing-grpc-new-open-source-http2.html>, 2015.
8. Pfister G F. An introduction to the infiniband architecture[J]. High Performance Mass Storage and Parallel I/O, 2001, 42: 617-632.
9. Ou L, He X, Han J. An efficient design for fast memory registration in RDMA[J]. Journal of Network and Computer Applications, 2009, 32(3): 642-651.
10. Sur S, Jin H W, Chai L, et al. RDMA read based rendezvous protocol for MPI over InfiniBand: design alternatives and benefits[C]//Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming. ACM, 2006: 32-39.
11. Frey P W, Alonso G. Minimizing the hidden cost of RDMA[C]//Distributed Computing Systems, 2009. ICDCS'09. 29th IEEE International Conference on. IEEE, 2009: 553-560.