OD-SGD: One-Step Delay Stochastic Gradient Descent for Distributed Training

YEMAO XU, DEZUN DONG, Yawei ZHAO, WEIXIA XU, and XIANGKE LIAO,
National University of Defense Technology, China

The training of modern deep learning neural network calls for large amounts of computation, which is often provided by GPUs or other specific accelerators. To scale out to achieve faster training speed, two update algorithms are mainly applied in the distributed training process, i.e., the Synchronous SGD algorithm (SSGD) and Asynchronous SGD algorithm (ASGD). SSGD obtains good convergence point while the training speed is slowed down by the synchronous barrier. ASGD has faster training speed but the convergence point is lower when compared to SSGD. To sufficiently utilize the advantages of SSGD and ASGD, we propose a novel technology named One-step Delay SGD (OD-SGD) to combine their strengths in the training process. Therefore, we can achieve similar convergence point and training speed as SSGD and ASGD separately.

To the best of our knowledge, we make the first attempt to combine the features of SSGD and ASGD to improve distributed training performance. Each iteration of OD-SGD contains a global update in the parameter server node and local updates in the worker nodes, the local update is introduced to update and compensate the delayed local weights. We evaluate our proposed algorithm on MNIST, CIFAR-10, and ImageNet datasets. Experimental results show that OD-SGD can obtain similar or even slightly better accuracy than SSGD, while its training speed is much faster, which even exceeds the training speed of ASGD.

CCS Concepts: • **Computing methodologies** → **Distributed computing methodologies; Distributed algorithms**;

Additional Key Words and Phrases: Overlap ratio, one-step delay, compensation, convergence speed, distributed performance

ACM Reference format:

Yemao Xu, Dezun Dong, Yawei Zhao, Weixia Xu, and Xiangke Liao. 2020. OD-SGD: One-Step Delay Stochastic Gradient Descent for Distributed Training. *ACM Trans. Archit. Code Optim.* 17, 4, Article 30 (September 2020), 26 pages.
<https://doi.org/10.1145/3417607>

1 INTRODUCTION

Deep learning has recently experienced success on various areas, ranging from computer vision to natural language processing and so on. Many prominent results in these areas are achieved through the utilization of deep learning networks [27]. The resurgence of deep learning mainly

This project is partially supported by the National Science and Technology Major Projects on Core Electronic Devices, High-end Generic Chips and Basic Software under Grant No. 2018ZX01028101.

Authors' address: Y. Xu, D. Dong (corresponding author), Y. Zhao, W. Xu, and X. Liao, National University of Defense Technology, 109 Deya Rd, Changsha, 410073, China; emails: {xuyemaovip, dong, zhaoyawei, xuweixia, xkliao}@nudt.edu.cn.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Association for Computing Machinery.

1544-3566/2020/09-ART30

<https://doi.org/10.1145/3417607>

Fig. 1. Mechanisms of SSGD and OD-SGD. Supposing the training and communication overheads are three time units, respectively. When training with SSGD, the next iteration task will not start until the communication process of current iteration ends, and the overhead per iteration is five time units. When training with OD-SGD, the next training task can be launched when the computation process of current iteration ends, and the overhead for one iteration is three time units.

results from the appearance of general and specific hardware accelerators (GPU [19], NPU [9], TPU [26], etc.), which brings about tremendous advances in computing power, and the public availability of versatile training datasets like ImageNet [17] and CIFAR [15] also contributes its development. However, with the fast-growing sizes of deep neural networks and datasets, the computing capability for training makes the bottleneck, it requires days or weeks to complete the training (for instance, the training of ResNet-50 [22] for 90 epochs takes 21 hours with 8 P100GPUs), which makes it impossible for real-time interaction in the development process. Under this circumstance, distributed training becomes the prevailing practice and it greatly improves the productivity of training deeper and larger neural networks [12, 31, 40, 46].

Stochastic Gradient Descent (SGD) is a widely used optimization algorithm for distributed training. For the training data of the same size, the computation time used in the forward-backward process can be reduced dramatically by increasing the training nodes and making use of data parallelism [29]. One can choose the synchronous SGD algorithm (SSGD) or asynchronous SGD algorithm (ASGD) in the training phase. For SSGD, devices in each local worker will figure out the gradients first with their own mini-batches of data, and then the gradients are added to the global models. SSGD can obtain good convergence speed while its training speed deteriorates due to the synchronous overhead. To circumvent the shortcoming of SSGD, practitioners have also resorted to ASGD, which focus on training speed and uses the staled information for computation. Although ASGD achieves better training performance by alleviating the waiting overhead in SSGD, it often suffers from poorer convergence accuracy because of the delayed gradients [7].

Based on the conclusions that SSGD achieves faster convergence speed while ASGD experiences faster training speed, the combination of these features would probably lead to better distributed training performance. Therefore, we propose a novel method named One-step Delay SGD algorithm (OD-SGD) to take advantage of the merits in both sides. The intrinsic motivation of this strategy is to start the next iteration training as soon as possible when current iteration ends. As illustrated in Figure 1, the computation and communication processes take three time units separately. The overhead for one iteration under SSGD is five time units. By launching the next training task immediately when the current iteration ends, the overhead under OD-SGD takes only three time units for one iteration.

OD-SGD is similar to SSGD in the sense that to ensure the convergence speed and training result, the synchronous mechanism is still contained. It differs from SSGD in that no workers need to wait for others, the weights used for the next iteration are delivered while the current iteration task is running. The similarity between OD-SGD and ASGD is that each local worker will update

the pulled weights with their local gradients. From Figure 1, we can notice OD-SGD has a higher overlap ratio, the communication overhead is overlapped by the computation overheads of two adjacent iterations. This new proposed method includes three main stages, namely the warm-up stage, the switching stage, and the training stage with one-step delay mechanism. The warm-up stage here means training with SSGD in the beginning, which is different from that mentioned in [18]. To facilitate the distinction between these two warm-up operations, the warm-up in OD-SGD is marked as warm-up while the latter is marked as WP. Warm-up stage is introduced in OD-SGD to make the global weights stable before the application of one-step delay mechanism, thus achieving a better convergence point. The switching stage is the last two iterations in the warm-up stage. In the penultimate iteration, the local workers will make a copy of the pulled weights. In the last iteration, each local worker will update the copied weights with the newly calculated gradients and launch the next training task immediately. In the meanwhile, the gradients are transmitted to the parameter servers. In the third stage, the parameter servers will update weights with the received gradients under synchronous mechanism, while the local workers will update the copied weights and deliver the gradients. It is noteworthy that OD-SGD is applicable to both Peer-to-Peer and Parameter Server structures.

We implement OD-SGD in MXNet [8], an efficient and flexible framework designed for neural network training. MXNet is parameter server structure based and one of the pivotal designs is the dependency engine, which is a library that takes a sequence of operations and schedules them according to the dependency pattern, potentially in parallel. It supports parallel executions of irrelevant operations in the training process. Although OD-SGD is implemented in MXNet, it is definitely applicable to end-to-end based learning platforms like Pytorch [33] and Caffe [25].

We emphasize the obstacles that need to be tackled for the sake of implementing the OD-SGD in MXNet as follows. First, how the computation and communication operations for training are scheduled in MXNet. Second, the communication mechanism for transmitting the gradients and updated parameters. Third, how does the dependency engine solve the data flow dependency problem, or a detailed understanding of the dependency engine. Finally, what operations should be added to implement the OD-SGD.

In order to settle the challenges above, we dive into the training mechanism of MXNet and solidly implement the OD-SGD, experimental results show the effectiveness of our OD-SGD. This article mainly makes the following contributions. In the first place, we briefly introduce the essential training mechanism of MXNet, emphasizing on communication pattern and dependency engine. Second, we design the OD-SGD and implement it on MXNet platform. Third, mathematical analysis is made on why OD-SGD works and its convergence property. Fourth, experimental results prove that OD-SGD delivers similar convergence accuracy as SSGD, or even slightly better, while it achieves faster training speed than ASGD. Finally, we make some analyses about warm-up sensitivity and the performance improvement under different cluster scales and bandwidths.

The rest of this article is structured as follows: We show the background in Section 2, then introduce the design and implementation of OD-SGD and provide mathematical proof in Section 3 and Section 4. Experiment results and related work are presented in Section 5 and Section 6. Finally, we conclude with insights we have learned through designing such an algorithm in Section 7.

2 BACKGROUND

2.1 Data-Parallelism SGD and Model-Parallelism

Data parallelism method is widely used in distributed deep learning training, in which the dataset is partitioned according to the number of local workers. Supposing m workers is involved, each worker is responsible for $1/M$ of the dataset in one epoch training task. In the data parallelism method, every worker makes a copy of the neural network and the weights (w_i). The training

Fig. 2. Data parallel and model parallel.

Fig. 3. Communication mechanism of MXNet.

process contains three parts: computing and delivering the local gradients, updating the global weights maintained in the master node, and pulling back the updated weights. The data parallelism strategy can be divided into two modes: synchronous mode and asynchronous mode, corresponding to SSGD and ASGD, whose advantages and disadvantages are discussed above. Unlike data parallelism, model parallelism method divides the neural network into several parts. Nevertheless, because of the limited input size (e.g., size of an image), the matrix operations are not large, the partition of a neural network does not need many machines [43] and the batch size for one iteration is limited because of the limited memory. Therefore, the data parallelism is often the choice of many state-of-art methods, Figure 2 provides examples of data and model parallelism.

2.2 Communication Mechanism of MXNet

To update the global weights, the parameter servers need to get gradients from all other workers. The communication process, which consists of two phases: (1) intra-node reduction and (2) inter-node point-to-point message passing, is completed through **Push/Pull** operations. In the first phase, gradients calculated by different devices of the local worker will be aggregated on one single device, as shown in Figure 3(a). In the second phase, each local worker will transmit the aggregated gradient to the server node via **Push/Pull** operations (Figure 3(b)). Every parameter corresponds to one pair of *<key, value>*, and the gradients of parameters are delivered one-by-one. Supposing a layer has two parameters, two pairs of **Push/Pull** operations will be launched.

2.3 Dependency Engine in MXNet

A dependency engine is a library that schedules a series of operations according to their dependency relationships between each other. An MXNet dependency engine is designed to make MXNet run faster and scale to larger datasets, with which we can parallelize computation and communication operations across devices. Both kinds of operations require reading or writing data from different variables. Each variable correspond to a maintained queue in the dependency engine and Figure 4 presents a sample which describes three different states of a queue. *write-1*, *read-2*, *read-3*, and *write-4* are operations to access the same variable, independent operations (*read-2* & *read-3*) can be parallelized to improve the execution efficiency. We come up with the OD-SGD by parallelizing more independent operations to improve the distributed training performance.

3 ONE-STEP DELAY SGD: ALGORITHM DESCRIPTION AND IMPLEMENTATION

This design aims at improving the training performance of distributed deep learning systems, and accelerating the computation and communication pipeline in the training engine. The main idea is to increase the overlap ratio between computation and communication process. OD-SGD is divided into three primary stages, as shown in Figure 5, warm-up stage, switching stage, and training stage (training under OD-SGD mechanism).

Fig. 4. A maintained queue sample in dependency engine.

Fig. 5. Overview of the three stages involved in OD-SGD. Weights (w'_{t+1}) copied at iteration $i-1$ is updated at iteration i (update: $(w'_{t+1}, \nabla w_{t+1})$) and used for training at iteration $i+1$ with the OD-SGD mechanism. Actually, the **pull** and **copy** operations in iteration i are conducted when the training task of iteration $i+1$ is executing, they can be parallelized.

ALGORITHM 1: OD-SGD: parameter server

Initialize: $count = 0, t = 0, w_0$ is initialized randomly, $m \in \{1, 2, \dots, M\}$;
repeat
 receive M gradients (gradient_m^t) from M different workers (worker_m);
 update w_t with aggregated gradient (gradient_t);
 $t \leftarrow t + 1$;
 while $count < M$ **do**
 if receive “pull request” from worker_m **then**
 Send w_t back to worker_m ;
 $count \leftarrow count + 1$;
 end
 end
 $count \leftarrow 0$;
until forever;

3.1 Algorithm Description

The flow of OD-SGD is shown in Algorithm 1 and Algorithm 2. Here we implement OD-SGD by using the parameter server framework (it can be implemented in other frameworks). According to Algorithm 1, the work process of parameter server keeps unchanged when compared to SSGD. It receives gradients from all other workers and update the global weights using the synchronous mechanism. Then it sends the updated weights back to corresponding worker on receiving its *pull request*. Algorithm 2 illustrates the work process of each worker node. The warm-up phase of OD-SGD differs from that mentioned in [18], which is increasing the learning rate from a relative small value to the target value. Warm-up stage here means the worker conducts training task under the SSGD in the beginning. It broadcasts the pulled weight to every local device to compute the gradient, then aggregate the gradients from local devices and push the aggregated one to the parameter server. In the second stage, the worker makes a copy of the pulled weight (w_{bak}^t) when the iteration number is $(wp - 1)$ and updates it in the next iteration. In the third stage, the worker directly broadcasts w_{bak}^{t-1} to its local devices without waiting for the completion of **Pull** operation in the last iteration, because there is no data dependency between these two operations. For

ALGORITHM 2: OD-SGD: worker m **Input:** Warm-up iteration number wp .**Initialize:** Pull initialized w_0 from parameter server, $num = 0$, $t = 0$;

```

while  $num \leq wp$  do
  Broadcast  $w_t$  to local devices ( $dev_{loc}^i$ );
  Computer gradient  $g_m$ ;
   $t \leftarrow t + 1$ ;
  Push  $g_m$  to the parameter server;
  if  $num == wp$  then
    Update  $w_{bak}^{t-1}$  with  $g_m$ ;
  end
  Pull  $w_t$  from the parameter server;
  if  $num == wp - 1$  then
     $w_{bak}^t = w_t$ ;
  end
   $num \leftarrow num + 1$ ;
end
repeat
  Broadcast  $w_{bak}^{t-1}$  to local devices ( $dev_{loc}^i$ );
   $w_{bak}^t = w_t$ ;
  Computer gradient  $g_m$ ;
  Push  $g_m$  to the parameter server;
  Update  $w_{bak}^t$  with  $g_m$ ;
   $t \leftarrow t + 1$ ;
  Pull  $w_t$  from the parameter server;
until forever;

```

instance, the last **Pull** operation in the *While* loop and the first **Broadcast** in the **repeat** part can be parallelized, so are the following adjacent **Pull** and **Broadcast** operations in the **repeat** part. The parallel execution of **Pull** and **Broadcast** is the key to increasing the overlap ratio, making it possible to launch the next iteration task as soon as the current iteration completes.

It is noteworthy that OD-SGD combines the features of SSGD and ASGD. OD-SGD, and SSGD are similar in that the parameter server works under synchronous update mechanism, and OD-SGD and ASGD are similar in that local workers can start the next iteration tasks quickly without waiting for the synchronous communication process. When compared to SSGD in implementation, no modification is required in the parameter server side. For the local worker side, OD-SGD calls for a copy of the local weights, which brings about extra memory usage. Therefore, we store the backup model weights (w_{bak}^t) in the CPU-side memory and will not occupy the memory of devices (GPUs, TPUs, etc.) used for computation task. In addition, OD-SGD has extra computational requirement on the local workers to update w_{bak}^t with newly calculated gradients. However, the additional computation only introduces a lightweight overhead and it can be partly overlapped by the **Push** operation. Because both the **Push** operation and the **local_update** operation have only *read dependency* on the newly calculated gradients, and they can be parallelized according to Figure 4. For the algorithm for **local_update**, we can utilize the existing SGD-based algorithms or custom algorithms defined by ourselves. In our experiments, we adopt the SGD, DC-ASGD-c, and DC-ASGD-a [45] algorithms for different neural network and cluster scales.

Fig. 6. Three cases of OD-SGD algorithm divided by the value of T_{cop} and T_{com} . T_{cop} is the computation overhead of forward and backward propagations while T_{com} is the pure communication cost in yellow color which excludes the overlaps.

3.2 Algorithm Analysis on Performance Improvement

This section introduces how to evaluate the performance improvement of OD-SGD, which is proposed for improving the distributed training performance and Figure 5 illustrates its principles briefly. Actually, situation described in Figure 5 can be further divided into three cases (Figure 6) according to the value of T_{cop} and T_{com} . We evaluate the training performance improvement rate with Equation (1), T_{org} is the time cost for completing one iteration with the SSGD algorithm while T_{new} represents the overhead of one iteration under OD-SGD algorithm, which can be calculated by Equation (2). When $T_{cop} < T_{com}$, T_{new} is the average of T_{cop} and T_{com} . From Figure 6(c), we can notice the total overhead for **iteration i** and **iteration (i + 1)** is the sum of T_{cop} and T_{com} . Besides, the overheads for different iterations are equal, which is the sum of T_{cop} and part of T_{com} . Therefore, we have the conclusion that T_{new} is the average of T_{cop} and T_{com} . It is noteworthy that the communication overhead of SSGD (T'_{com}) does not equal that of OD-SGD (T_{com}), because the lightweight overheads of the introduced **local update** operations may prolong the communication overhead by postponing the executions of the next **Pull** and **Push** operations. In addition, the degree of network congestion also differs as the communication traffic per second changes.

$$IMP_Rate = \frac{T_{org} - T_{new}}{T_{org}}; \quad (1)$$

$$T_{org} = T_{cop} + T'_{com}; \quad T_{new} = \begin{cases} T_{cop} & T_{com} \leq T_{cop} \\ \frac{T_{com} + T_{cop}}{2} & T_{cop} < T_{com} \end{cases} \quad (2)$$

Equation (1) and Equation (2) will lead to Equation (3) and we have the next observations based on this new equation: (1) When $T_{com} \leq T_{cop}$, a larger T'_{com} will lead to better performance improvement when OD-SGD is applied because more overhead can be overlapped, the T_{new} equals the T_{cop} and has nothing to do with T_{com} . (2) When $T_{com} > T_{cop}$, the improvement rate is able to reflect the influences of **local update** operations on the communication overhead. Because the value of T_{cop} is fixed, and the difference between T_{com} and T'_{com} will determine the value of **IMP_Rate**, which results from the additional computational overheads brought about by the **local update** operations. (3) The maximum value of performance improvement rate is 50%, whereas it cannot be attained as the value of T'_{com} is always less than T_{com} .

$$IMP_Rate = \begin{cases} 1 - \frac{T_{cop}}{T_{cop} + T'_{com}} & T_{com} \leq T_{cop} \\ 1 - \frac{T_{com} + T_{cop}}{2(T_{cop} + T'_{com})} & T_{cop} < T_{com} \end{cases} \quad (3)$$

3.3 Algorithm Implementation

In order to implement the OD-SGD, we focus on two things in this section, one of which is to define the local updater, and the other one is to adjust the execution sequences of operations involved in the third stage.

4 CONVERGENCE ANALYSIS

This section provides simple mathematical analysis of OD-SGD as well as its convergence property. Considering the simplicity of the proof process, we assume local workers utilize SGD algorithm for local update operations and the parameter server and local workers share the same learning rate. Nevertheless, different learning rates will also lead to the same proof result.

4.1 Update Rules

Parameter Server. The update rule of parameter server are consistent under SSGD and OD-SGD, the only difference is the value of gradients. Assuming the $(k-1)$ -th iteration is the last iteration in switching stage, then the parameter server conducts the update operation as Equation (7). After the switching stage, the parameter server will update global weights following Equations (8), where $\nabla F(\mathbf{w}'_{i,K}; \xi_i)$ is the gradient from the i th local worker, calculate with local weight $\mathbf{w}'_{i,K}$.

$$\mathbf{w}_K = \mathbf{w}_{K-1} - \frac{\eta}{N} \sum_{i=1}^N \nabla F(\mathbf{w}_{i,K-1}; \xi_i) \quad (7)$$

$$\mathbf{w}_{K+1} = \mathbf{w}_K - \frac{\eta}{N} \sum_{i=1}^N \nabla F(\mathbf{w}'_{i,K}; \xi_i) \quad (8)$$

$$\mathbf{w}_{K+2} = \mathbf{w}_{K+1} - \frac{\eta}{N} \sum_{i=1}^N \nabla F(\mathbf{w}'_{i,K+1}; \xi_i) \quad (9)$$

Local Worker. In the k th iteration, local workers calculate the weights of their respective versions with the local gradients, which are calculated using the global weight $\mathbf{w}_{K-1}$, as shown in Equation (10). Starting from the $(k+1)$ -th iteration, local workers update their weights according to Equation (11), $\mathbf{w}'_{i,K}$ is the local weights used for training in the previous iteration while $\mathbf{w}_K$ is the global weight retrieved from the parameter server.

$$\mathbf{w}'_K = \mathbf{w}_{K-1} - \eta \nabla F(\mathbf{w}_{K-1}; \xi_i) \quad (10)$$

$$\mathbf{w}'_{K+1} = \mathbf{w}_K - \eta \nabla F(\mathbf{w}'_{i,K}; \xi_i) \quad (11)$$

$$\mathbf{w}'_{K+2} = \mathbf{w}_{K+1} - \eta \nabla F(\mathbf{w}'_{i,K+1}; \xi_i). \quad (12)$$

4.2 Convergence Rate Proof

ASSUMPTION 1. Our analyses use the following assumptions:

- We assume that $\|\mathbf{w} - \mathbf{w}_*\| \leq R$ holds for the optimum $\mathbf{w}_*$ and any $\mathbf{w}$.
- We assume that $\|\nabla f(\mathbf{w})\| \leq G$ holds for any $\mathbf{w}$.
- We assume that $\mathbb{E} \|\nabla f(\mathbf{w}; \xi_{k,i}) - \nabla f(\mathbf{w})\| \leq \sigma$ holds for any $k \in [K]$, $i \in [N]$, and $\mathbf{w}$.
- We assume that f has L -Lipschitz gradient, that is, $\nabla f(\mathbf{u}) - \nabla f(\mathbf{v}) \leq L \|\mathbf{u} - \mathbf{v}\|$ holds for any $\mathbf{u}$ and $\mathbf{v}$.

THEOREM 1. Choose $\eta > 0$ in Algorithm 1 and $\eta' = \eta > 0$ in Algorithm 2. Under Assumption 1, we have

$$\mathbb{E} f\left(\frac{1}{K} \sum_{k=1}^K \mathbf{w}_k\right) - f(\mathbf{w}_*) \leq \frac{R^2}{2T\eta} + \frac{\eta}{2}(G + \sigma)^2 + 2\eta LR(G + \sigma). \quad (13)$$

PROOF. According to the update rule at the server, we have

$$\begin{aligned}
& \|\mathbf{w}_{k+1} - \mathbf{w}_*\|^2 \\
&= \left\| \mathbf{w}_k - \mathbf{w}_* - \frac{\eta}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}) \right\|^2 \\
&= \|\mathbf{w}_k - \mathbf{w}_*\|^2 - 2\eta \left\langle \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}), \mathbf{w}_k - \mathbf{w}_* \right\rangle + \eta^2 \left\| \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}) \right\|^2 \\
&\leq \|\mathbf{w}_k - \mathbf{w}_*\|^2 - 2\eta \left\langle \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}), \mathbf{w}_k - \mathbf{w}_* \right\rangle + \frac{\eta^2}{N} \sum_{i=1}^N \|\nabla f(\mathbf{w}'_{k,i}; \xi_{k,i})\|^2 \\
&\leq \|\mathbf{w}_k - \mathbf{w}_*\|^2 - 2\eta \left\langle \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}), \mathbf{w}_k - \mathbf{w}_* \right\rangle + \eta^2 (G + \sigma)^2.
\end{aligned}$$

Therefore, we have

$$\begin{aligned}
& \left\langle \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}), \mathbf{w}_k - \mathbf{w}_* \right\rangle \\
&\leq \frac{\|\mathbf{w}_k - \mathbf{w}_*\|^2}{2\eta} - \frac{\|\mathbf{w}_{k+1} - \mathbf{w}_*\|^2}{2\eta} + \frac{\eta}{2} (G + \sigma)^2.
\end{aligned}$$

Thus, we have

$$\begin{aligned}
& \mathbb{E} \langle \nabla f(\mathbf{w}_k), \mathbf{w}_k - \mathbf{w}_* \rangle \\
&= \mathbb{E} \left\langle \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}), \mathbf{w}_k - \mathbf{w}_* \right\rangle + \mathbb{E} \left\langle \nabla f(\mathbf{w}_k) - \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}), \mathbf{w}_k - \mathbf{w}_* \right\rangle \\
&= \mathbb{E} \left\langle \frac{1}{N} \sum_{i=1}^N \nabla f(\mathbf{w}'_{k,i}; \xi_{k,i}), \mathbf{w}_k - \mathbf{w}_* \right\rangle + \mathbb{E} \frac{1}{N} \sum_{i=1}^N \langle \nabla f(\mathbf{w}_k) - \nabla f(\mathbf{w}'_{k,i}), \mathbf{w}_k - \mathbf{w}_* \rangle \\
&\stackrel{\textcircled{1}}{\leq} \frac{\|\mathbf{w}_k - \mathbf{w}_*\|^2}{2\eta} - \frac{\|\mathbf{w}_{k+1} - \mathbf{w}_*\|^2}{2\eta} + \frac{\eta}{2} (G + \sigma)^2 + \frac{LR}{N} \sum_{i=1}^N \|\mathbf{w}_k - \mathbf{w}'_{k,i}\| \\
&\stackrel{\textcircled{2}}{\leq} \frac{\|\mathbf{w}_k - \mathbf{w}_*\|^2}{2\eta} - \frac{\|\mathbf{w}_{k+1} - \mathbf{w}_*\|^2}{2\eta} + \frac{\eta}{2} (G + \sigma)^2 + \frac{\eta LR}{N} \sum_{i=1}^N \left\| \nabla f(\mathbf{w}'_{k-1,i}; \xi_{k-1,i}) - \frac{1}{N} \sum_{j=1}^N \nabla f(\mathbf{w}'_{k-1,j}; \xi_{k-1,j}) \right\| \\
&\leq \frac{\|\mathbf{w}_k - \mathbf{w}_*\|^2}{2\eta} - \frac{\|\mathbf{w}_{k+1} - \mathbf{w}_*\|^2}{2\eta} + \frac{\eta}{2} (G + \sigma)^2 + \frac{\eta LR}{N^2} \sum_{i=1}^N \sum_{j=1}^N \|\nabla f(\mathbf{w}'_{k-1,i}; \xi_{k-1,i}) - \nabla f(\mathbf{w}'_{k-1,j}; \xi_{k-1,j})\| \\
&\leq \frac{\|\mathbf{w}_k - \mathbf{w}_*\|^2}{2\eta} - \frac{\|\mathbf{w}_{k+1} - \mathbf{w}_*\|^2}{2\eta} + \frac{\eta}{2} (G + \sigma)^2 + \frac{\eta LR}{N^2} \sum_{i=1}^N \sum_{j=1}^N (\|\nabla f(\mathbf{w}'_{k-1,i}; \xi_{k-1,i})\| + \|\nabla f(\mathbf{w}'_{k-1,j}; \xi_{k-1,j})\|) \\
&\leq \frac{\|\mathbf{w}_k - \mathbf{w}_*\|^2}{2\eta} - \frac{\|\mathbf{w}_{k+1} - \mathbf{w}_*\|^2}{2\eta} + \frac{\eta}{2} (G + \sigma)^2 + 2\eta LR (G + \sigma).
\end{aligned}$$

① holds because that f has L -Lipschitz gradient. ② holds due to the update rule at the worker. Therefore, we have

$$\begin{aligned}
& \mathbb{E} f\left(\frac{1}{K} \sum_{k=1}^K \mathbf{w}_k\right) - f(\mathbf{w}_*) \\
& \leq \frac{1}{K} \sum_{k=1}^K \mathbb{E} f(\mathbf{w}_k) - f(\mathbf{w}_*) \\
& \leq \frac{1}{K} \sum_{k=1}^K \mathbb{E} \langle \nabla f(\mathbf{w}_k), \mathbf{w}_k - \mathbf{w}_* \rangle \\
& \leq \frac{\|\mathbf{w}_1 - \mathbf{w}_*\|^2}{2T\eta} - \frac{\|\mathbf{w}_{K+1} - \mathbf{w}_*\|^2}{2T\eta} + \frac{\eta}{2}(G + \sigma)^2 + 2\eta LR(G + \sigma) \\
& \leq \frac{R^2}{2T\eta} + \frac{\eta}{2}(G + \sigma)^2 + 2\eta LR(G + \sigma).
\end{aligned}$$

Finally, we complete the proof. $\square$

COROLLARY 1. Choose $\eta = \frac{R}{\sqrt{T}(G+\sigma)}$ in Algorithm 1 and $\eta' = \eta > 0$ in Algorithm 2. Under Assumption 1, we have

$$\begin{aligned}
\mathbb{E} f\left(\frac{1}{K} \sum_{k=1}^K \mathbf{w}_k\right) - f(\mathbf{w}_*) & \leq \frac{R(G + \sigma)}{2\sqrt{T}} + \frac{R(G + \sigma)}{2\sqrt{T}} + \frac{2LR^2}{\sqrt{T}} \\
& \lesssim \frac{1}{\sqrt{T}}.
\end{aligned}$$

$a \leq b$ means a is less than b up to a constant. Therefore, our algorithm achieves at least $O(\frac{1}{\sqrt{T}})$ convergence rate.

5 EVALUATION

As is mentioned above, we come up with OD-SGD to improve the distributed deep learning training performance, and we evaluate its performance in this section. We first introduce the hardware configurations, datasets, and neural networks used in our experiments. Second, we demonstrate the performance comparisons with SSGD, ASGD, and DC-ASGD, during which ASGD and DC-ASGD are delayed-information-based algorithms like OD-SGD. Test accuracy of all involved neural networks is TOP-1 accuracy. Third, we assess the performance improvement of OD-SGD under various cluster scales and network bandwidths. Finally, we make some sensitivity analysis on warm-up period.

5.1 Experiments Configuration

The OD-SGD is implemented in MXNet framework and we conduct the evaluation experiments on two different GPU clusters. One of cluster is equipped with two K80 (dual GPUs) Tesla GPUs on each node, interconnected with both Ethernet (10 Gbps) and InfiniBand (56 Gbps), while the other cluster is equipped with four V100 Tesla GPUs on each node, interconnected with InfiniBand (56 Gbps) only. The topology of both clusters is the simple star network topology. Machines of the K80 GPU cluster are installed with Red Hat 4.8.3, CUDA 8.0 and cuDNN 6.0 while machines of

the V100 GPU cluster are installed with Centos 7.6, CUDA 10.0 and cuDNN 7.4.1. Because of the limited number of nodes, we set only one parameter server throughout our experiments.

Datasets. To prove the effectiveness of our algorithm, we utilize three datasets in our experiments: (1) MNIST [28], it is a database of handwritten digits, which consists of 60,000 training examples and 10,000 test samples; (2) CIFAR-10 [15], it contains 60,000 32x32 color images in 10 different classes, 50,000 pictures for training and 10,000 for testing; and (3) ImageNet ILSVRC2012 [17], it is a subset of ImageNet and contains 1.2 million pictures in 1,000 categories.

Neural Networks. Neural networks applied in our experiments comes with the MXNet framework, including MobileNet1.0 [23], ResNet18_V2, ResNet-20, ResNet-50 [22], ranking from low-complexity (MobileNet1.0) to high-complexity (ResNet-50). MobileNet1.0 is trained with MNIST dataset, ResNet18_V2, and ResNet-20 correspond to CIFAR-10 while ImageNet is for ResNet-50. For the sake of fairness, all these neural networks are trained from the same randomly initialized model.

We implement the OD-SGD and DC-ASGD on MXNet platform and set 200 iterations for the warm-up stage. In addition, the SSGD, ASGD have already been implemented, we compare the accuracy of these four algorithms to evaluate our proposed OD-SGD algorithm. With or without momentum makes almost no difference to the accuracy, therefore, the momentum of DC-ASGD is set to 0 in the experiments [45].

5.2 Experimental Results on MNIST

Training MobileNet1.0 with MNIST dataset is quick and efficient, it is unnecessary for large-scale distributed deployment. We conduct experiments on MobileNet1.0 to prove that OD-SGD algorithm is applicable to neural network of low-complexity. We perform the training task for 40 epochs, with a mini-batch size of 128, the training task is deployed on four worker nodes interconnected with Ethernet. The initial learning rate of the four different algorithms is set to 0.05, the momentum for ASGD and DC-ASGD is 0, while 0.9 for SSGD and the global updater of OD-SGD. We chose SGD algorithm as the local update function of OD-SGD, whose learning rate and momentum are the same as SSGD. In addition, we chose the constant version of DC-ASGD (DC-ASGD-c) for comparison with $\lambda = 0.04$.

In the first place, we demonstrate the learning curves under fixed number of training epochs as shown in Figure 9. From Figures 9(a) and 9(c) we can notice that the convergence speed of SSGD and OD-SGD-c is obviously faster than that of ASGD and DC-ASGD. SSGD and OD-SGD share similar curves while the learning patterns of ASGD and DC-ASGD are alike. The better performance of synchronous algorithms (SSGD, OD-SGD) mainly results from the low-complexity of neural network and the computing capacity differences of the workers. The low-complexity leads to low computing overhead for each iteration, coupled with the different computation capabilities, the delayed gradients problem is further deteriorated when training with asynchronous algorithms (ASGD, DC-ASGD). Besides, the training tasks for workers are launched at different times, this kind of delay also imposes a negative effect on the delayed gradients problem. Figures 9(b) and 9(d) demonstrate the further comparisons between SSGD and OD-SGD. The training accuracy curve of OD-SGD is similar to SSGD while its accuracy is slightly lower than that of SSGD in the first 25 epochs. This results from the fact that OD-SGD utilizes the features of SSGD and ASGD algorithms, the asynchronous mechanism has a negative influence on the accuracy. Nevertheless, OD-SGD introduces the SGD algorithm for local update to ensure the accuracy.

We then compare the convergence speed of these algorithms as shown in Figure 10. As the total training time for 40 epochs does not vary greatly of the algorithms, we present their test accuracy throughout the whole training process. From subgraph 10(a), we can see that it takes the longest

(a) Train accuracy under different algorithms

(b) Train accuracy comparison of SSGD and OD-SGD

(c) Test accuracy under different algorithms

(d) Test accuracy comparison of SSGD and OD-SGD

Fig. 9. Train accuracy and test accuracy of MobileNet1.0 under different epochs on MNIST dataset.

(a) Test accuracy under different algorithms

(b) Test accuracy comparison of SSGD and OD-SGD

Fig. 10. Test accuracy of MobileNet1.0 w.r.t wallclock time on MNIST dataset.

time for SSGD to complete the training task (≈ 2600 s) due to the synchronous barrier, followed by DC-ASGD, OD-SGD, and ASGD. It is obvious that convergence speed of synchronous algorithms is faster than that of asynchronous algorithms. Moreover, the ASGD achieves faster convergence speed than DC-ASGD algorithm, because the latter brings in extra computation overhead in the parameter server to compensate the delayed gradients. The delayed gradients problem may be alleviated, it also prolongs training time for one iteration, which is more distinct for neural network of low-complexity. Subgraph 10(b) provides further comparisons of convergence speed between SSGD and OD-SGD. It is noteworthy that their curves share similar variation trends, proving that the OD-SGD algorithm has the characteristics of SSGD. Furthermore, the convergence speed of OD-SGD is much faster than SSGD due to the feature of ASGD, it achieves the best test accuracy at about 700 seconds while SSGD fulfills the target at 1,200 seconds. Nevertheless, the SSGD algorithm experiences a drop after the peak accuracy and the curve of OD-SGD is more stable.

Table 1 demonstrates the training speed and classification accuracy on MNIST test set using MobileNet1.0. The following conclusions can be drawn from the table: (1) ASGD algorithm achieves the fastest training speed because there is no synchronous obstacle. DC-ASGD algorithm utilizes

Table 1. Training Speed and Classification Accuracy on MNIST Test Set

Algorithms	Speed (<i>samples/s</i>)	Train Acc (%)	Test Acc (%)
SSGD	236.63	100	98.592
ASGD	505.29	99.731	98.252
DC-ASGD	306.434	99.963	98.396
OD-SGD	401.395	100	98.756

Table 2. Test Accuracy on CIFAR-10 after the Moving Average Operation

# of workers	Test Accuracy (%)			
	SSGD	ASGD	DC-ASGD	OD-SGD
4	80.482	79.890	80.156	80.426
8	79.673	78.232	78.985	79.678
12	78.844	78.288	77.659	78.914
1	81.048			

the same asynchronous mechanism as ASGD, however, the introduced computational operations slow down its speed to process the training data. Training speed of OD-SGD is faster than DC-ASGD algorithm, because the introduced computational overhead for local update can be partly overlapped by the **Push** communication operation; (2) The training accuracy of OD-SGD can be as high as SSGD, both of their accuracy is 100% after training for 15 epochs; (3) OD-SGD achieves the best test accuracy and its final value is 98.756%. In addition, higher training accuracy may not ensure higher test accuracy. The accuracy of DC-ASGD is slightly higher than that of SSGD while its training accuracy is much less; and (4) Experiments on MNIST dataset with MobileNet1.0 prove that our proposed OD-SGD algorithm is suitable for training model of low-complexity, which keeps fast training speed and high test accuracy, its performance is slightly better than that of DC-ASGD algorithm.

5.3 Experimental Results on CIFAR-10

In this section, we investigate the training performance of those algorithms on CIFAR-10 dataset (without data augmentation), using the neural network ResNet18_v2. We conduct the training for 160 epochs, with a mini-batch size of 128. Training task is deployed on clusters with 4, 8, and 12 worker nodes (in addition to a parameter server node), corresponding to 16, 32, and 48 GPUs, respectively, and we make use of the InfiniBand network for communication process.

We conduct all the training for 160 epochs, with a mini-batch size of 128, and the learning rate is decreased by 10 times after training for 80 and 120 epochs. Hyper-parameters for DC-ASGD utilize the recommended value ($\eta = 0.5$, $\lambda = 0.04$ for DC-ASGD-c and $\eta = 0.5$, $\lambda = 2$, $m = 0.95$ for DC-ASGD-a) [45]. ASGD and SSGD share the same learning rates ($\eta_0 = 0.1$ ($M = 4$), $\eta_1 = 0.2$ ($M = 8$), $\eta_2 = 0.2$ ($M = 12$)), M means the number of workers in the training process. Momentums for DC-ASGD and ASGD are 0 while 0.9 for SSGD. For OD-SGD, SSGD is utilized for the global update operation in the parameter server. DC-ASGD-c and DC-ASGD-a are used for local update operations, corresponding to OD-SGD-c and OD-SGD-a. When training with four workers, the DC-ASGD-c is applied, DC-ASGD-a method can also be used while it brings about more additional computational overhead. However, when the cluster scale is increased to 8 or 12 worker nodes, local update operation with DC-ASGD-c cannot achieve the target accuracy and DC-ASGD-a is used. We compare the training performance and convergence speed of the above four algorithms (SSGD, ASGD, DC-ASGD-c, and OD-SGD). The DC-ASGD-c is chosen for comparison because its computational cost is lower than DC-ASGD-a and its test accuracy is higher than SSGD according to [45].

First, we present the training accuracy and test accuracy under different epochs (Figure 11). Subgraphs in the left column demonstrates the training accuracies of different clusters and we have the following observations: (1) All algorithms achieve 100% training accuracies after training for about 80 epochs, although they experience various convergence curves in the first 80 epochs. (2) The gradient problem deteriorates as the scale of cluster goes up. When $M = 4$, the ASGD,

Fig. 11. Train accuracy and test accuracy of ResNet18_V2 under different epochs on CIFAR-10 dataset.

SSGD and OD-SGD-c have similar convergence curves, while when $M = 12$, there exists obvious differences between their convergence curves. The OD-ASGD-c has the worst training curves in all three subgraphs and we attribute this to the improper hyper-parameters. Nevertheless, the hyper-parameters involved in OD-SGD-c for local update operation are the same to those of DC-ASGD-c. For the sake of fairness, we conduct experiments on ResNet-20 neural network later.

Subgraphs in the middle column displays the test accuracies under different clusters. Convergence curves in subgraphs 11(b), 11(e), and 11(i) are based on the raw test accuracy data, from which we draw the following conclusions. (1) OD-SGD (OD-SGD-a and OD-SGD-c) can achieve similar or slightly better test accuracy when compared with SSGD. (2) A larger number of workers leads to more pronounced vibration of test accuracy. When $M = 4$, the test accuracy curves of SSGD and OD-SGD-c are stable while the curves of ASGD and OD-SGD-c vibrate slightly. However, the accuracy curves of all algorithms vibrate significantly when $M = 12$. (3) A larger cluster scale expands the test accuracy gap between different algorithm. Because a larger cluster size leads to worse gradient delay problem. It is noteworthy that when $M = 12$, the test accuracy of OD-SGD-c is less than ASGD. Therefore, the compensation strategy based on ASGD may not be applicable when the batch size is large enough, extra optimization method is needed.

Subgraphs in the right column also illustrates the test accuracies. To get rid of the vibrations, we use moving average to smooth the data after the 91st epoch, with window size of 5. For instance, the test accuracy of the 92nd epoch is the average value of the 90th, 91st, 92nd, 93rd, and 94th epochs. Curves of these subgraphs are based on the average test accuracies, which provide clearer trends throughout the training process. From these subgraphs, we can apparently notice that OD-SGD is able to guarantee the test accuracy under different clusters, proving its effectiveness in distributed training.

Fig. 12. Train accuracy and test accuracy of ResNet18_V2 w.r.t wallclock time on CIFAR-10 dataset.

Table 2 displays the test accuracy after smoothing the data with moving average, from which we have the following observations: (1) Test accuracies of all algorithms decrease as the cluster size goes up in general. Take SSGD as an example, when $M = 4$, the test accuracy is 80.482%, which decreases to 78.844% when $M = 12$. This phenomenon is reasonable, ASGD and DC-ASGD suffers from the gradient delay problem which deteriorates as the cluster size goes up. The mini-batch size of each local worker keeps unchanged as the cluster size increase, and larger cluster size will lead to larger batch size. The enlarged batch size imposes negative influences on the test accuracy of neural network, which accounts for the accuracy decrease of SSGD and OD-SGD; (2) Training on a single local worker achieves the best test accuracy; (3) OD-SGD can achieve similar test accuracy when compared with SSGD. When $M = 4$, its accuracy is 0.056% lower than SSGD. Nevertheless, its test accuracy is higher than that of SSGD when $M = 8$ and $M = 12$.

The convergence speed of different algorithms is shown in Figure 12. From these subgraphs, we have the following observations: (1) OD-SGD-c and OD-SGD-a achieve the best training performance, which run faster than ASGD; (2) Although the convergence point of ASGD is lower than that of SSGD, its convergence speed is faster than that of SSGD; (3) The training speed of SSGD is slowed down by its global barrier, but it still runs faster than DC-ASGD-c. The additional computational overhead brought in by DC-ASGD-c significantly affect its throughput in the training process. When $M = 12$, the training speed of OD-SGD-c is too slow to show the training results in the subgraphs; and (4) Our proposed OD-SGD achieves both good accuracy and speed. On the one hand, its convergence speed is faster than that of ASGD. On the other hand, its convergence point is as good as, or even slightly better than that of SSGD.

Fig. 13. Train accuracy and test accuracy of ResNet-50 w.r.t epochs on ImageNet.

5.4 Experimental Results on ImageNet

For the sake of verifying our algorithm on large scale neural network, we conduct experiments on ImageNet dataset using ResNet-50 [22]. In this large scale experiment, we evaluate the performance of SSGD, DC-ASGD-a, and OD-SGD-a. For all algorithms in this experiment, we conduct the training task for 90 epochs (without data augmentation) on eight workers which are interconnected by InfiniBand network, with a mini-batch size of 256, and the learning rate is reduced by 10 times after training for 30, 60, and 80 epochs according to the practice in [18]. The initial learning rate for SSGD is 0.1 and the momentum value is 0.9, considering the batch size (2,048) and linear scaling rule, we adopt the linear WP strategy to adjust the learning rate to 0.8 after training for 5 epochs. The hyper-parameters of DC-ASGD-a follows the configure in [45] ($\eta = 0.1$, $\lambda = 2$, $m = 0$), the momentum value is 0. For OD-SGD-a, the global update in the parameter server follows the configuration of SSGD while local updates in the workers follow the configuration of OD-AGSD-a.

Table 3 displays the test accuracy while Figure 13 presents the train accuracy and test accuracy on ImageNet with regard to the epochs, from which we have the following observations: (1) After training for the same number of epochs, DC-ASGD-a always achieves higher training accuracy and its final value is similar to that of SSGD. On the contrary, training accuracy of OD-SGD-a is lower throughout the training process and the final accuracy is around 90%; (2) SSGD can obtain the best test accuracy (70.773%) and the accuracy of OD-SGD-a (69.847%) is slightly lower. Actually, we do not spare efforts to find the suitable hyper-parameters for OD-SGD-a and a suitable configuration may bring about similar accuracy to SSGD. Nevertheless, DC-ASGD-a (66.772%) is not able to get a good accuracy and the accuracy increase at epoch 60 is significantly lower. We attribute this to the large batch size and the configuration in [45] is not applicable to large batch size (512 vs 2,048) here; (3) Curves of SSGD and OD-SGD-a have similar patterns, proving that OD-SGD-a utilizes the feature of SSGD.

The convergence speed is illustrated in Figure 14, from which we can notice the speed of those algorithms are similar and this result is reasonable. The limited training speed mainly results from the low computing power of K80 GPU, each local worker can just handle about 50 samples per second. In addition, SSGD contains a synchronous barrier and DC-ASGD-a calls for frequent additional computation operations in the parameter server node. As to OD-SGD-a, the introduced local computation overheads in worker nodes prolong the communication overhead more or less. All these algorithms take more than 75 hours to complete the training process.

The long training time results from the low computation capability of K80 GPU and we conduct some supplementary experiments on a V100 GPU cluster. Because of the limited resource, we are only available to four worker nodes for parallel training. In the supplementary experiments,

Fig. 14. Train accuracy and test accuracy of ResNet-50 w.r.t wallclock time on ImageNet.

Fig. 15. Test accuracy of ResNet-50 w.r.t epochs on ImageNet. *SGD* means training on a single node, *OD(1.6+0.4)* means the global learning rate is 1.6 and the local learning rate for each worker is 0.4.

Fig. 16. Test accuracy w.r.t wallclock time on ImageNet, the mini-batch size is 512.

we adopt the data augmentation due to the high computing power of V100 GPU and the polynomial (power = 2) decay LR policy is used to adjust the learning rate. We compare the training performance of SSGD and OD-SGD, the mini-batch size of each node is 512 (V100 GPU has larger memory, 16G) with momentum = 0.9 and weight decay = 0.0001, the WP period for increasing the learning rate is 5 epochs.

Figure 15 demonstrates the test accuracy of SGD, SSGD, and OD-SGD-a. The accuracy of sequential SGD is the baseline and its learning rate is 0.4 with no WP. The initial learning rate of SSGD is 1.6 and we present accuracies of OD-SGD-a under different initial learning rates, as shown in Table 4. From Figure 15 we have the following observations: (1) When $M = 4$, the test accuracy for SSGD is 73.145%, which is better than sequential SGD 73.041%. This test performance lift might arise from the regularization effect brought by the variance introduced by parallel training; (2) OD-SGD-a achieves the best test accuracy when the global learning rate is 1.6 and local learning rate is 0.2, which is about 0.316% and 0.212% lower than that of SSGD and SGD, respectively; (3) It is necessary to search for the proper hyper-parameters like learning rate for OD-SGD via grid search, this figure shows that improper learning rates will lead to obvious accuracy drop. However, it is very time-consuming for a large-scale neural network to search for the best configuration. In addition to the learning rate, the hyper-parameters of DC-ASGD-a should also be taken into account. Figure 16 displays the convergence speed of SSGD and OD-SGD-a on the V100 GPU cluster. The sequential SGD takes about 35.6 hours to complete the training task and the convergence curve is not shown in this figure. The training time of SSGD is 10.5 hours while OD-SGD-a takes about

Table 3. Test Accuracy on ImageNet with ResNet-50 (K80 GPU Cluster)

Algorithms	Test Accuracy (%)
SSGD	70.773
DC-ASGD-a	66.772
OD-SGD-a	69.847

Table 4. Test Accuracy on ImageNet with ResNet-50 (V100 GPU Cluster)

Algorithm	SGD	SSGD	OD(1.6+1.6)	OD(1.6+0.4)
Accuracy(%)	73.041	73.145	72.439	72.671
Algorithm	OD(1.6+0.2)	OD(1.6+0.1)	OD(0.4+0.4)	OD(0.2+0.4)
Accuracy(%)	72.829	71.952	72.675	71.543

9.2 hours to accomplish the training workload, which achieves 12.38% improvement in training speed with only 0.317% convergence accuracy drop. Therefore, OD-SGD can be used for training models of large-scale.

We also evaluate the speedups of OD-SGD on different workloads. Figure 17 illustrates the normalized training speed of ResNet-50, AlexNet, and VGG11 under SSGD, ASGD, and OD-SGD, and the training speed of SSGD is set to 1. Numbers under the x axis are speedup percentages of OD-SGD as compared with SSGD. *Res - 50(128)* means the batch size of ResNet-50 for a single node is 128. From Figure 17 we have the following observations: (1) Training speed of OD-SGD is similar to that of ASGD when training ResNet-50. As ResNet-50 is a computation-intensive model, computation overhead of the delay step can overlap the communication efficiently, leading to the similar training speed. (2) The speedup percentage of *Res - 50(512)* is not equal to the above performance improvement (15.38% vs 12.38%). Because the speedup percentage in Figure 17 is evaluated with 4 adjacent epochs. However, when training ResNet-50 for 90 epochs, we need to restart the training task from the checkpoint once it is dumped, and different network streams of other tasks will occupy the network bandwidths at different times, resulting in the slightly different performance improvement ratio; and (3) Training speed of ASGD is faster than that of OD-SGD when training AlexNet while slightly slower when training VGG11. Both AlexNet and VGG11 are communication-intensive models, model size of VGG11 is about $2.7\times$ that of AlexNet while the amount of involved calculation is $23.25\times$ larger [41], more communication overhead can be overlapped when training VGG11, contributing to the faster training speed of OD-SGD when compared to ASGD.

5.5 Training Performance Under Different Network Bandwidths and Cluster Sizes

In this section, we evaluate the impacts of network bandwidth and cluster scale on training speed. Network bandwidth and cluster size are two main factors that affect the communication overhead in a distributed training process. Therefore, we are analyzing the influences of communication overhead on training performance.

$$GR_Rate_{od-sgd} = \frac{Speed_{od-sgd} - Speed_{ssgd}}{Speed_{ssgd}} * 100\%; \quad GR_Rate_{dc-sgd} = \frac{Speed_{dc-sgd} - Speed_{ssgd}}{Speed_{ssgd}} * 100\%. \quad (14)$$

We conduct the training task of ResNet18_V2 with CIFAR-10 again on clusters with different scales, while the communication process is completed through Ethernet instead of InfiniBand network. For the sake of saving the expenditure of renting machines, we do not run the training task to the end when the Ethernet is used because of the slow training speed. Table 5 presents the throughputs of all algorithms with various number of worker nodes. The throughput of SSGD is chosen as the baseline and we evaluate the performance improvement according to Equation (14). Figure 18 and Figure 19 illustrate the improvement ratio under Ethernet and InfiniBand network. From the table and figures, we have the following observations: (1) The throughput of each local

Fig. 17. Normalized training speed of different models.

Fig. 18. Performance improvement with Ethernet.

Fig. 19. Performance improvement with IB network.

Table 5. Training Speed Comparisons between SSGD, OD-SGD, and DC-ASGD under Different Clusters

# of Workers	Network Type	SSGD (samples/s)	ASGD (samples/s)	DC-ASGD (samples/s)	OD-SGD (samples/s)
4	Eth	95.785	157.887	151.011	149.517
	IB	236.78	244.19	215.03	309.782
8	Eth	74.685	107.35	98.26	98.76
	IB	193.43	220.63	149.83	270.08
12	Eth	57.62	83.00	74.82	78.37
	IB	155.47	163.40	88.47	196.94

The neural network and dataset are ResNet18_V2 and CIFAR-10 separately. **Eth** is short for Ethernet while **IB** means InfiniBand network.

worker decreases apparently as the cluster size goes up. This results from the increase of communication overhead, leading to growth of overhead for each iteration; (2) OD-SGD achieves training performance improvements with both Ethernet and InfiniBand networks; (3) When the Ethernet is used for communication, ASGD has the best performance and OD-SGD is slightly better than DC-ASGD when $M = 12$; and (4) When InfiniBand network is applied, the performance improvement ratio of OD-SGD is significantly higher than ASGD and OD-SGD. Furthermore, the DC-ASGD imposes negative impacts on the training speed and a larger cluster size leads to worse performance. Because the application of InfiniBand network significantly decreases the communication overhead and the additional computation overhead introduced by OD-SGD makes the bottleneck. The parameter server needs to conduct more computation operations with a larger cluster size.

5.6 Warm-Up Sensitivity Analysis

The warm-up stage introduced in OD-SGD is to make the weights more stable before the third stage, thus achieving a better convergence point. To analyze the influence from the warm-up stage, we conduct some training tasks on CIFAR-10 using the ResNet-20 model, using the V100 GPU cluster with four worker nodes. The warm-up period for OD-SGD-c (DC-ASGD-c is used for local update operation) is set 0 epoch, 1 epoch, 2 epochs, 3 epochs, 4 epochs, and 5 epochs, separately. The mini-batch size for a single node is 128, learning rates of SGD and SSGD are 0.1 and 0.4, the global learning for OD-SGD-c is 0.4 while the local learning rate is 0.5, no WP strategy is used.

Figure 20 demonstrates the test accuracy curves of SGD, SSGD, and OD-SGD-c under various warm-up epochs, the test accuracy is shown in Table 6. From the figure and table we have the following observations: (1) Test accuracy of SSGD is slightly better than that of sequential SGD,

Fig. 20. Test accuracy of ResNet-20 w.r.t different warm-up epochs on CIFAR-10, convergence curves of SGD and SSGD are also included.

Table 6. Test Accuracy on CIFAR-10 with ResNet-20 (V100 GPU Cluster), Different Epochs are Used for OD-SGD-c in the Warm-up Stage, **OD-1wp** Means the Warm-up Stage is 1 Epoch

Algorithm	SGD	SSGD	OD-0wp	OD-1wp	OD-2wp	OD-3wp	OD-4wp	OD-5wp
Accuracy(%)	92.027	92.125	91.780	91.643	92.050	92.136	92.177	92.183

which is the same as training ResNet-50 with ImageNet; (2) When the warm-up stage are 0 epoch and 1 epoch, the convergence accuracy is obviously lower than that of SGD, proving that warm-up stage in OD-SGD-c is necessary; (3) When the warm-up stage is 2 epochs, test accuracy of OD-SGD-c is slightly higher than that of SGD while lower than SSGD; and (4) When warm-up stage is 3 or more epochs, test accuracy of OD-SGD-c is comparable to that of SSGD (Figure 20(b)), we can also notice OD-SGD-c even achieves about 0.05% improvement in convergence accuracy according to Table 6.

Based on the observations above, we conduct the training task of ResNet-50 again using the V100 GPU cluster (four worker nodes) and ImageNet dataset, the warm-up period is set to 1 epoch. Considering the conclusion (Section 5.4) that hyper-parameters impose obvious influences on the convergence accuracy, we modify the mini-batch size from 512 to 128 as the configuration ($\eta = 0.1$, $\lambda = 2$, $m = 0$) mentioned in [45] is more suitable for batch size 128 (32 per GPU card). Nevertheless, the training speed is also greatly reduced as the batch size goes down. The learning rate of SSGD is 0.4 while the global learning of OD-SGD-a is 0.4 and the local learning rate is 0.1.

Figure 21 presents the test accuracy of ResNet-50 with regard to epochs, from which we have the following observations: (1) SSGD have better convergence rate than OD-SGD-a and its test accuracy is obviously higher than that of OD-SGD-a in the first 65 epochs; (2) It is noteworthy that OD-SGD-a achieves higher convergence accuracy than SSGD after training for 90 epochs, and their test accuracy is 73.864% and 73.786%, respectively. Figure 22 demonstrates the test accuracy with regard to time. Although SSGD has better convergence rate, its training speed is much slower than that of OD-SGD-a. With the same amount of training time, test accuracy of OD-SGD-a is significantly higher. It takes 39.30 hours for SSGD to complete the 90 epochs training task while 32.27 hours for OD-SGD-a, leading to 17.89% improvement in training speed.

Based on our experiments, the warm-up stage is necessary for OD-SGD and the length of warm-up period is much more like a empirical trick, which is difficult to provide an accurate number of epochs for this stage from a theoretical perspective. The required least number of epochs in this stage differs for different neural networks.

Fig. 21. Test accuracy of ResNet-50 w.r.t epochs on ImageNet, the mini-batch size is 128 and the warm-up period is 1 epoch.

Fig. 22. Test accuracy w.r.t wallclock time on ImageNet, the mini-batch size is 128.

6 RELATED WORK

This section describes the optimization research landscape of distributed deep learning training in recent years and we summary some works from four fields: overlapping communication and computation, reducing computation overhead, reducing the number of communication times, and reducing the communication overhead for each iteration.

6.1 Overlapping Communication and Computation

Several similar works like OD-SGD have been conducted previously, which improves distributed training performance via overlapping communication and computation. Agarwal and Duchi [2] analyze the convergence of gradient-based optimization algorithm, and results show that for smooth stochastic problems, delays due to asynchrony are asymptotically negligible. This algorithm is based on asynchronous mechanism and different nodes suffer from various delays according to their distances from the master node. However, OD-SGD utilizes the features of ASGD and SSGD, all local workers suffer from the same level of delay and it is suitable for more than smooth problems. Delay Compensated ASGD (DC-ASGD) proposed by Zheng et al. [45] is based on ASGD and the compensation operations are conducted in the master node to improve the convergence point. Different from DC-ASGD, compensation operations of OD-SGD are executed on each local worker node and there is a global synchronization in the parameter server, mechanism of OD-SGD combines the features of ASGD and SSGD. Assran et al. [5] also come up with the Stochastic Gradient Push (SGP) method for Peer-to-Peer structure based distributed training. A node will block if it does not receive messages from its in-neighbors after τ iterations. OD-SGD differs from SGP in that it applies to both Parameter Server and Peer-to-Peer structures. We implement OD-SGD in MXNet (Parameter Server based) and use convolution neural networks (non-smooth problems) for experiments, while SGP is designed for Peer-to-Peer structure framework and optimization algorithm in [2] is for smooth stochastic problems, therefore, we do not present the performance comparison between OD-SGD and them. However, based on our experimental results, performance of SGP should be better than OD-SGD and we need to make further optimizations to improve the performance OD-SGD.

6.2 Reducing Computation Overhead

A series of optimizations have been conducted to lower the computation time. One of the critical strategies is to improve the computing capacity of hardware, which is also the main driving force behind the momentum that deep learning has gained [6]. The GPU series ranges from Tesla K80 to Tesla T4 [20], and the TPU series ranges from TPUv1 to TPUv3 [38], the computing power is extremely improved. Besides, reducing the overhead for accessing main memory also contributes

to computation overhead reduction, Dadiannao [10] greatly accelerates the computing speed by introducing more on-chip storage, leading to much fewer accesses to main memory, while Prime [11] takes advantage of Processing-in-memory (PIM) to reduce the memory access overhead. Low-precision computation is also used for lowering the time cost of deep learning, and it can be implemented in terms of hardware or software. For the perspective of hardware, Gupta et al. designs an energy-efficient hardware accelerator to implement the low-precision (16-bit) neural network training with stochastic rounding [21], and the neural network accelerator named outlier-aware accelerator (OLAccel) is based on 4-bit multiply-accumulate (MAC) units [32], and so on. As to the term of software, De Sa et al. [16] introduce a training algorithm called High-Accuracy Low-Precision (HALP) and Micikevicius et al. [30] describe a methodology (Mixed Precision Training, MPT) for neural network training using half-precision floating point numbers.

6.3 Reducing Communication Times

Reducing communication times means a reduction in the number of iterations, and for fixed epochs, the way to reduce the iteration number is increasing the batch size. Goyal et al. [18] first adopt the warming up scheme and linear scaling rule to increase the batch size, and successfully train ResNet-50 ImageNet model in 1 hour. You et al. design the LARS [42] to ensure the training accuracy as the batch size goes up greatly, and they further increase the batch size from 8K to 32K. Sun et al. [36] increase the batch size to 64K and they complete the training of AlexNet ImageNet model in 1.5 minutes over 512 Volta GPUs. As a matter of fact, a series of research has been done to increase the batch size [3, 14, 24, 35], and the most important thing in this field is keeping the model accuracy as the batch size goes up. Therefore, for any strategy to improve the training performance, we should prevent the accuracy from dropping.

6.4 Reducing Communication Overhead

As the scale of cluster increases, the high communication overhead makes the bottleneck of distributed machine learning, especially for centralized deployment mode based frameworks such as MXNet [8] and Tensorflow [1]. The primary method for decreasing the communication overhead is reducing the communication traffic. Wen et al. [39] utilize ternary gradients to largely cut down the communication workload. Alistarh et al. [4] designs the Quantized SGD (QSGD), which is a family of compression schemes with convergence guarantees and good practical performance. Yu et al. [44] propose a gradient vector quantization technique called GradiVeQ by using the strong linear correlations between CNN gradients. In addition to these quantization strategies, Baidu [37] proposes the ring-based all-reduced algorithm and the PowerAI Distributed Deep Learning system of IBM [13] also mentioned a new all-reduced algorithm. In response to the disadvantage of the original version of ring-based all-reduced algorithm, Horovod [34] introduces the gradient fusion method to all-reduce algorithm, reducing tensor fragmentation and improving bandwidth utilization.

7 CONCLUSIONS

In this article, we proposed a novel algorithm named One-Step Delay SGD (OD-SGD), which combines the features of SSGD and ASGD, to improve the distributed deep learning training performance. We have evaluated OD-SGD on MNIST, CIFAR-10 and ImageNet datasets, and compare its performance with the ASGD, SSGD, and DC-ASGD. Experimental results demonstrate that OD-SGD can obtain similar or slightly better accuracy than SSGD with a much faster convergence speed, which is even faster than ASGD. When applied to large-scale neural network, OD-SGD experiences a light drop in convergence accuracy, which calls for further improvement. Therefore, as for future work, we plan to design a proper compensation algorithm for local update operation.

Further, we intend to extend the delay from one step to several steps to achieve higher communication and computation overlap ratio for performance improvement. Our work is open source on github (<https://github.com/CynthiaProtector/OD-SGD>).

ACKNOWLEDGMENTS

We sincerely thank all the reviewers for their constructive feedbacks and detailed advices.

REFERENCES

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467* (2016).
- [2] Alekh Agarwal and John C. Duchi. 2011. Distributed delayed stochastic optimization. In *Advances in Neural Information Processing Systems*. 873–881.
- [3] Takuya Akiba, Shuji Suzuki, and Keisuke Fukuda. 2017. Extremely large minibatch SGD: Training Resnet-50 on Imagenet in 15 minutes. *arXiv preprint arXiv:1711.04325* (2017).
- [4] Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. 2017. QSGD: Communication-efficient SGD via gradient quantization and encoding. In *Advances in Neural Information Processing Systems*. 1709–1720.
- [5] Mahmoud Assran, Nicolas Loizou, Nicolas Ballas, and Michael Rabbat. 2018. Stochastic gradient push for distributed deep learning. *arXiv preprint arXiv:1811.10792* (2018).
- [6] Ammar Ahmad Awan, Khaled Hamidouche, Jahanzeb Maqbool Hashmi, and Dhableswar K. Panda. 2017. S-Caffe: Co-designing MPI runtimes and Caffe for scalable deep learning on modern GPU clusters. In *Proceedings of the 22nd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. ACM, 193–205.
- [7] Jianmin Chen, Xinghao Pan, Rajat Monga, Samy Bengio, and Rafal Jozefowicz. 2016. Revisiting distributed synchronous SGD. *arXiv preprint arXiv:1604.00981* (2016).
- [8] Tianqi Chen, Mu Li, Yutian Li, Min Lin, Naiyan Wang, Minjie Wang, Tianjun Xiao, Bing Xu, Chiyuan Zhang, and Zheng Zhang. 2015. Mxnet: A flexible and efficient machine learning library for heterogeneous distributed systems. *arXiv preprint arXiv:1512.01274* (2015).
- [9] Yunji Chen, Tianshi Chen, Zhiwei Xu, Ninghui Sun, and Olivier Temam. 2016. DianNao family: Energy-efficient hardware accelerators for machine learning. *Commun. ACM* 59, 11 (2016), 105–112.
- [10] Yunji Chen, Tao Luo, Shaoli Liu, Shijin Zhang, Liqiang He, Jia Wang, Ling Li, Tianshi Chen, Zhiwei Xu, Ninghui Sun, and Olivier Temam. 2014. Dadiannao: A machine-learning supercomputer. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Computer Society, 609–622.
- [11] Ping Chi, Shuangchen Li, Cong Xu, Tao Zhang, Jishen Zhao, Yongpan Liu, Yu Wang, and Yuan Xie. 2016. Prime: A novel processing-in-memory architecture for neural network computation in reconfigurable main memory. In *ACM SIGARCH Computer Architecture News*, Vol. 44. IEEE Press, 27–39.
- [12] Trishul M. Chilimbi, Yutaka Suzue, Johnson Apacible, and Karthik Kalyanaraman. 2014. Project Adam: Building an efficient and scalable deep learning training system. In *OSDI*, Vol. 14. 571–582.
- [13] Minsik Cho, Ulrich Finkler, Sameer Kumar, David Kung, Vaibhav Saxena, and Dheeraj Sreedhar. 2017. Powerai DDL. *arXiv preprint arXiv:1708.02188* (2017).
- [14] Valeriu Codreanu, Damian Podareanu, and Vikram Saleetore. 2017. Scale out for large minibatch SGD: Residual network training on ImageNet-1K with improved accuracy and reduced time to train. *arXiv preprint arXiv:1711.04291* (2017).
- [15] The CIFAR-10 dataset. 2019. <https://www.cs.toronto.edu/kriz/cifar.html>.
- [16] Christopher De Sa, Megan Leszczynski, Jian Zhang, Alana Marzoev, Christopher R. Aberger, Kunle Olukotun, and Christopher Ré. 2018. High-accuracy low-precision training. *arXiv preprint arXiv:1803.03383* (2018).
- [17] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR'09)*. IEEE, 248–255.
- [18] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. 2017. Accurate, large minibatch SGD: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677* (2017).
- [19] CUDA GPUs. 2019. <https://developer.nvidia.com/cuda-gpus>.
- [20] Cuda GPUs. 2019. <https://developer.nvidia.com/cuda-gpus>.

- [21] Suyog Gupta, Ankur Agrawal, Kailash Gopalakrishnan, and Pritish Narayanan. 2015. Deep learning with limited numerical precision. In *International Conference on Machine Learning*. 1737–1746.
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 770–778.
- [23] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- [24] Xianyan Jia, Shutao Song, Wei He, Yangzihao Wang, Haidong Rong, Feihu Zhou, Liqiang Xie, Zhenyu Guo, Yuanzhou Yang, Liwei Yu, et al. 2018. Highly scalable deep learning training system with mixed-precision: Training imagenet in four minutes. *arXiv preprint arXiv:1807.11205* (2018).
- [25] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM International Conference on Multimedia*. ACM, 675–678.
- [26] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, Raminder Bajwa, Sarah Bates, Suresh Bhatia, Nan Boden, Al Borchers, et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Proceedings of the 44th Annual International Symposium on Computer Architecture*. ACM, 1–12.
- [27] Matthias Langer, Ashley Hall, Zhen He, and Wenny Rahayu. 2018. MPCA SGD—A method for distributed training of deep learning models on spark. *IEEE Transactions on Parallel and Distributed Systems* 29, 11 (2018), 2540–2556.
- [28] Yann LeCun. 1998. The MNIST database of handwritten digits. <http://yann.lecun.com/exdb/mnist/> (1998).
- [29] Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J. Dally. 2017. Deep gradient compression: Reducing the communication bandwidth for distributed training. *arXiv preprint arXiv:1712.01887* (2017).
- [30] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. 2017. Mixed precision training. *arXiv preprint arXiv:1710.03740* (2017).
- [31] Philipp Moritz, Robert Nishihara, Ion Stoica, and Michael I. Jordan. 2015. Sparknet: Training deep networks in spark. *arXiv preprint arXiv:1511.06051* (2015).
- [32] Eunhyeok Park, Dongyoung Kim, and Sungjoo Yoo. 2018. Energy-efficient neural network accelerator based on outlier-aware low-precision computation. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 688–698.
- [33] PyTorch. 2019. <https://pytorch.org/features>.
- [34] Alexander Sergeev and Mike Del Balso. 2018. Horovod: Fast and easy distributed deep learning in TensorFlow. *arXiv preprint arXiv:1802.05799* (2018).
- [35] Samuel L. Smith, Pieter-Jan Kindermans, Chris Ying, and Quoc V. Le. 2017. Don't decay the learning rate, increase the batch size. *arXiv preprint arXiv:1711.00489* (2017).
- [36] Peng Sun, Wansen Feng, Ruobing Han, Shengen Yan, and Yonggang Wen. 2019. Optimizing network performance for distributed DNN training on GPU clusters: ImageNet/AlexNet training in 1.5 minutes. *arXiv preprint arXiv:1902.06855* (2019).
- [37] Baidu. 2017. Bringing HPC Techniques to Deep Learning. <https://andrew.gibiansky.com/blog/machine-learning/baidu-allreduce/>.
- [38] TPU. 2019. <https://www.nextplatform.com/2018/05/10/tearing-apart-googles-tpu-3-0-ai-coprocessor/>.
- [39] Wei Wen, Cong Xu, Feng Yan, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. 2017. Terngrad: Ternary gradients to reduce communication in distributed deep learning. In *Advances in Neural Information Processing Systems*. 1509–1519.
- [40] Eric P. Xing, Qirong Ho, Wei Dai, Jin Kyu Kim, Jinliang Wei, Seunghak Lee, Xun Zheng, Pengtao Xie, Abhimanu Kumar, and Yaoliang Yu. 2015. Petuum: A new platform for distributed machine learning on big data. *IEEE Transactions on Big Data* 1, 2 (2015), 49–67.
- [41] Yemao Xu, Dezun Dong, Weixia Xu, and Xiangke Liao. 2019. SketchDLC: A sketch on distributed deep learning communication via trace capturing. *ACM Transactions on Architecture and Code Optimization (TACO)* 16, 2 (2019), 7.
- [42] Yang You, Igor Gitman, and Boris Ginsburg. 2017. Large batch training of convolutional networks. *arXiv preprint arXiv:1708.03888* (2017).
- [43] Yang You, Zhao Zhang, Cho-Jui Hsieh, James Demmel, and Kurt Keutzer. 2018. Imagenet training in minutes. In *Proceedings of the 47th International Conference on Parallel Processing*. ACM, 1.
- [44] Mingchao Yu, Zhifeng Lin, Krishna Narra, Songze Li, Youjie Li, Nam Sung Kim, Alexander Schwing, Murali Annavaram, and Salman Avestimehr. 2018. Gradiveq: Vector quantization for bandwidth-efficient gradient aggregation in distributed cnn training. In *Advances in Neural Information Processing Systems*. 5123–5133.

- [45] Shuxin Zheng, Qi Meng, Taifeng Wang, Wei Chen, Nenghai Yu, Zhi-Ming Ma, and Tie-Yan Liu. 2017. Asynchronous stochastic gradient descent with delay compensation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR.org, 4120–4129.
- [46] Martin Zinkevich, Markus Weimer, Lihong Li, and Alex J. Smola. 2010. Parallelized stochastic gradient descent. In *Advances in Neural Information Processing Systems*. 2595–2603.

Received September 2019; revised June 2020; accepted August 2020