



DEFT: Exploiting Gradient Norm Difference between Model Layers for Scalable Gradient Sparsification

Daegun Yoon
Ajou University
Suwon, Republic of Korea
kljp@ajou.ac.kr

Sangyoon Oh
Ajou University
Suwon, Republic of Korea
syoh@ajou.ac.kr

ABSTRACT

Gradient sparsification is a widely adopted solution for reducing the excessive communication traffic in distributed deep learning. However, most existing gradient sparsifiers have relatively poor scalability because of considerable computational cost of gradient selection and/or increased communication traffic owing to gradient build-up. To address these challenges, we propose a novel gradient sparsification scheme, DEFT, that partitions the gradient selection task into sub tasks and distributes them to workers. DEFT differs from existing sparsifiers, wherein every worker selects gradients among all gradients. Consequently, the computational cost can be reduced as the number of workers increases. Moreover, gradient build-up can be eliminated because DEFT allows workers to select gradients in partitions that are non-intersecting (between workers). Therefore, even if the number of workers increases, the communication traffic can be maintained as per user requirement.

To avoid the loss of significance of gradient selection, DEFT selects more gradients in the layers that have a larger gradient norm than the other layers. Because every layer has a different computational load, DEFT allocates layers to workers using a bin-packing algorithm to maintain a balanced load of gradient selection between workers. In our empirical evaluation, DEFT shows a significant improvement in training performance in terms of speed in gradient selection over existing sparsifiers while achieving high convergence performance.

CCS CONCEPTS

• Computing methodologies → Distributed computing methodologies; Machine learning.

KEYWORDS

distributed deep learning, gradient sparsification, scalability

ACM Reference Format:

Daegun Yoon and Sangyoon Oh. 2023. DEFT: Exploiting Gradient Norm Difference between Model Layers for Scalable Gradient Sparsification. In *52nd International Conference on Parallel Processing (ICPP 2023)*, August 07–10, 2023, Salt Lake City, UT, USA. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3605573.3605609>

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 the author(s) 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.

ICPP 2023, August 07–10, 2023, Salt Lake City, UT, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0843-5/23/08...\$15.00
<https://doi.org/10.1145/3605573.3605609>

1 INTRODUCTION

Scaling out the GPU cluster size is crucial for faster training of deep neural network (DNN) models [15, 26, 36]. However, the scalability of distributed training is affected by communication bottlenecks, where the bandwidth is insufficient to timely transmit gradients [31]. To alleviate the communication problem, gradient sparsification [14, 23, 30, 37] was introduced to reduce the number of transmitting gradients. Most of the gradient sparsification methods that have been proposed thus far [8, 9, 12, 20, 24, 38] share the same objective: selecting and transmitting large gradients. Because stochastic gradient descent (SGD) aims to minimise the objective function of DNN models, large gradients have a greater impact on updates than small gradients [27]. Therefore, the selection of larger gradients is preferred in gradient sparsification. Given this, Top-k gradient sparsification [10, 11, 32–35] has become the most basic approach.

However, the implementation of Top-k gradient sparsification is expensive. To identify the top k elements from n elements, a large amount of computational resources are required [16] for computational complexity of $O(n \log k)$ [29]. The computational cost remains constant even if we increase the number of workers because every worker must find their top k gradients from the entire gradient vector.

Furthermore, although every worker searches the entire gradient vector (i.e., an index selected by a worker may not be selected by another worker), the indices of the selected gradients by the workers do not fully overlap because every worker selects their top k gradients based on their locally computed gradients. Therefore, to compute the average gradients, the selected indices must be collected from all the workers [21, 34, 39, 40]. Consequently, the number of globally selected indices increases, whereas the number of locally selected indices remains fixed at k . This phenomenon is called gradient build-up [13], and may increase the communication traffic by up to n times when the number of workers is n in the worst case. Figure 1 shows that the measurement of density¹ increases through the gradient build-up. In Figure 1, the density measured in each case was significantly higher than user-set value 0.01, although it should be retained at 0.01. Therefore, gradient build-up can become a negative constraint on communication in terms of scalability and must be addressed to scale-out the cluster.

Recently, Chen et al. [13] proposed the cyclic local top-k (CLT-k) to eliminate gradient build-up. In the proposed scheme, the top k gradient selection is delegated to one worker who becomes the

¹Density ($d = \frac{k}{n_g}$) is defined as the ratio of the selected gradients (k) to all gradients (n_g), and it is the inverse of compression ratio which is the term used in several papers. For example, $d = 0.01$ indicates that 1% of the gradients are selected from among all gradients. In this case, the compression ratio is $100\times$.




Figure 1: Gradient build-up problem of Top-k gradient sparsification by scale-out of a cluster in distributed training. The measurement was conducted on ResNet-18 with CIFAR-10, and the batch size was 256 throughout. For each case, the actual density increases as the number of workers increases even though the density was set to 0.01 by the user.

leader of the selection in that iteration. The leader worker then broadcasts the indices of the selected gradients to the other workers after each iteration. However, the CLT-k scheme has limitations in terms of scalability of the distributed training. Herein, while a leader worker finds the top k gradients, other workers must wait for the result while being idle. As we scale-out the cluster, resource wastage increases. In addition, the computational cost of finding the top k gradients cannot be reduced even if the number of workers increases; that is, the task of finding the top k gradients is not divided and performed concurrently.

In this study, we propose a scheme: distributed execution of fragmented top-k (DEFT) to address this issue. DEFT partitions the gradient vector of the entire model into multiple tensors and exclusively allocates them to workers to reduce the search space for gradient selection. The allocation of exclusive parts of the gradients to workers leads to a reduction in the computational cost of the gradient selection of each worker as the number of workers increases (i.e., parallelisation). Although it would be natural to equally divide the gradient vector by the number of workers, DEFT divides the vector into multiple tensors based on the model layers². Because every layer of a DNN model has a different size and gradient norm [41], DEFT assigns different numbers of gradients to be selected for each layer based on the gradient norm of that layer. We claim that a sparsifier has a higher probability of selecting large gradients from layers with large gradient norms.

Upon computation of k of each layer, the layers are allocated to workers based on the computational cost of each layer using the bin-packing algorithm. Consequently, the entire gradient selection load is evenly distributed among the workers. That is, the computational cost of gradient selection can be reduced as the cluster scales out. Furthermore, gradient build-up can be eliminated

²Each tensor may be a gradient vector that corresponds to weights or biases of one layer in deep learning frameworks such as PyTorch. For simplicity, we refer to each partitioned tensor as a layer in this paper.

Algorithm 1 Distributed SGD with error-feedback

Input: $G(\cdot)$: stochastic gradients
 Sparsify(\cdot): gradient sparsifier
 n_g : number of gradients in model

```

1: for worker  $i \in [0, n - 1]$  in parallel do
2:   Initialise model  $x_0 \in \mathbb{R}^{n_g}$ ;
3:   Initialise local error  $e_{i,0} \leftarrow 0^{n_g}$ ;
4:   for iteration  $t \geq 0$  do
5:      $acc_{i,t} \leftarrow e_{i,t} + \eta_t G_{i,t}(x_t)$ ;
6:      $idx_{i,t} \leftarrow \text{Sparsify}(acc_{i,t})$ ;
7:      $idx_t \leftarrow \text{AllGather}(idx_{i,t})$ ;
8:      $g_{i,t} \leftarrow acc_{i,t}[idx_t]$ ;
9:      $g_t \leftarrow \text{AllReduce}(g_{i,t}, \text{SUM})$ ;
10:     $x_{t+1} \leftarrow x_t - \frac{1}{n} g_t$ ;
11:     $acc_{i,t}[idx_t] \leftarrow 0$ ;
12:     $e_{i,t+1} \leftarrow acc_{i,t}$ ;

```

because workers use different parts of the gradient vector, and the collected indices are combined without overlap. Thus, the density of gradient sparsification is retained regardless of the scale-out.

This study makes the following contributions:

- **Comparably high accuracy.** The accuracy of the proposed DEFT is comparable to that of non-sparsified distributed training. This achievement is mainly due to our approach, which selects large gradients from layers with large gradient norms.
- **Low computational cost.** DEFT incurs a significantly lower computational cost for gradient selection than the existing sparsifiers because of the load-balanced partitioning of the gradient vector.
- **High scalability of distributed training.** DEFT achieves a high speedup of gradient selection while retaining high accuracy when the cluster scales out.

The rest of this paper is organised as follows. Section 2 presents the preliminaries for this paper. Section 3 clarifies the limitations of state-of-the-art gradient sparsification methods. Section 4 proposes the DEFT designed to overcome the problems stated in this paper. Section 5 verifies our contributions by empirical comparisons between DEFT and existing gradient sparsifiers. Section 6 concludes the paper.

2 PRELIMINARIES

Gradient sparsification is in the family of lossy algorithms, because only a subset of gradients is selected while a majority of gradients are discarded. Discarding gradients leads to a difference between the trained model of gradient sparsification and that of non-sparsified distributed training. To improve the convergence rate of sparsified model, this difference must be reduced. In this section, we introduce the method that reduces the difference between sparsified- and non-sparsified- models.

Table 1: Strengths and weaknesses of state-of-the-art gradient sparsifiers and the proposed DEFT.

Sparsifier	Gradient build-up	Unpredictable density	Hyperparameter tuning	Worker idling	Gradient selection cost	Additional overhead
Top-k [11]	Yes	Yes	No	No	Very high	No
CLT-k [13]	No	No	No	Yes	Very high	No
Hard-threshold [27]	Yes	Yes	Yes	No	Very low	No
SIDCo [24]	Yes	Yes	No	No	Very low	Very high
DEFT	No	No	No	No	Low	Very low

We assume that a distributed SGD is scheduled at the learning rate: $\eta_t > 0$. Subsequently, the gradient sparsification is denoted by

$$x_{t+1} = x_t - \eta_t \frac{1}{n} \sum_{i=0}^{n-1} \tilde{G}_{i,t}(x_t), \quad (1)$$

where $\tilde{G}_{i,t}(x_t) = \text{Sparsify}(G_{i,t}(x_t))$ for worker i and iteration t . Function $\text{Sparsify}(x_t)$ selects a small fraction of the gradients from the entire gradient vector. For a gradient vector, $x \in n_g$, the output of $\text{Sparsify}(x_t) \in n_g$ is a very sparse vector because the unselected gradients become zero in the vector and are discarded.

However, discarding unselected gradients is inefficient because computing gradients using backward propagation consumes considerable computing resources during distributed training. Moreover, some layers of the DNN model have larger gradients than the rest of the layers, and a sparsifier is likely to select most of the gradients only in particular layers during most of the training iterations. Although this behaviour trains a model with a fast convergence rate, the SGD may converge at an undesired point because the sparsifier conducts a biased selection of gradients. To make unselected gradients contribute to model updates, Seide et al. [28] introduced error-feedback that accumulates unselected gradients in the local memory of each worker. The number of unselected gradients increases owing to accumulation over the iterations. Thus, smaller gradients, when sufficiently large for selection, can also contribute to updating the model.

The sum of each gradient accumulation for each worker is referred to as the local error. In this study, the local error is denoted by $\|e_{i,t}\|$ using the L2-norm for the magnitude of the error. Then, the error is denoted as

$$\|e_t\| = \frac{1}{n} \sum_{i=0}^{n-1} \|e_{i,t}\|. \quad (2)$$

The error represents the update difference between models with and without sparsification. We consider that a sparsifier conducts a significant gradient selection if it effectively reduces the error.

Algorithm 1 presents the pseudocode of distributed SGD to which gradient sparsification with error-feedback is applied. In line 5, the computed gradients are accumulated to local error. In line 6, the sparsifier selects the gradients based on its policy and returns the indices of the selected gradients. From lines 7 to 9, the workers collect the globally selected indices and compute the averages of the selected gradients. From lines 10 to 12, the model is updated using the averaged gradients and the accumulated value

of each selected gradient is initialised to zero. If the employed sparsifier selects large gradients effectively, the error, $\|e_t\|$, approaches zero as the iterations proceed.

3 LIMITATIONS OF STATE-OF-THE-ART METHODS

In this section, we discuss the limitations of the state-of-the-art methods. Table 1 lists the strengths and weaknesses of the state-of-the-art methods.

3.1 Sorting-based Sparsifiers

To determine top k gradients, the entire gradient vector must be sorted. In Table 1, Top-k and CLT-k [13] correspond to sparsifiers that require gradient-vector sorting. Owing to sorting, these methods suffer from extremely high computational costs, thereby limiting their practicality. Moreover, this computational cost remains the same regardless of the number of workers.

In terms of communication, Top-k incurs a gradient build-up that becomes more severe as the number of workers increases, making the actual density unpredictable. In contrast, CLT-k eliminates gradient build-up by delegating the top k gradient selection to one worker. Thus, the actual density is always maintained at the user-set value. However, as a side effect, CLT-k causes other workers to remain idle, which is not the case for Top-k. Therefore, Top-k and CLT-k have limitations in terms of scalability.

3.2 Threshold-based Sparsifiers

To overcome the high computational cost of sorting, threshold-based sparsifiers [24, 27] have been proposed; hard-threshold sparsifier and SIDCo are few examples of this type. Threshold-based sparsifiers require only $O(n_g)$ for gradient selection, because they check whether each gradient is larger than the threshold. Thus, the computational cost of the gradient selection in the hard-threshold sparsifier and SIDCo is considerably low. The difference between a hard-threshold sparsifier and SIDCo lies in the determination of the threshold. For hard-threshold sparsifier, the threshold should be defined before training and should not change during training. Therefore, strict hyperparameter tuning is required for each model and dataset. In contrast, SIDCo determines the threshold based on a statistical analysis of the gradient distribution at every iteration. Thus, the threshold changes according to the gradients of the model. However, SIDCo incurs a very high computational overhead owing to the statistical analysis at every iteration.

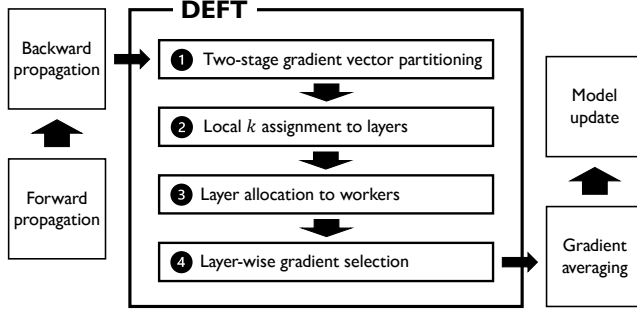


Figure 2: Overview of DEFT.

Despite the threshold estimation of each sparsifier, there is a large difference between the number of locally selected gradients of each worker and the number that the user wants to select. This is because of inaccurate threshold estimation. An inaccurate threshold estimation limits the practicality of threshold-based sparsifiers. Moreover, they exhibit gradient build-up because each worker selects different gradient indices. Therefore, it is difficult to predict the actual density of threshold-based sparsifiers. Consequently, they must be further studied to satisfy the fundamental purpose of gradient sparsification.

The limitations of the state-of-the-art methods reveal that it is very challenging to satisfy all the criteria in Table 1. In this study, we address this challenge using a novel sparsifier with negligible overhead.

4 DEFT DESIGN

We designed the DEFT as a gradient sparsifier that determines gradient indices by partitioning the gradient vector. DEFT comprises the following sequence: 1) two-stage gradient vector partitioning, 2) local k assignment to layers, 3) layer allocation to workers, and 4) layer-wise gradient selection. Figure 2 presents an overview of the DEFT. In the process of distributed training, the DEFT is located between the backward propagation and communication for gradient averaging. The following subsections present a detailed discussion of each process in the DEFT.

4.1 Two-Stage Gradient Vector Partitioning

To reduce the computational cost of the Top- k selection in the entire gradient vector, DEFT performs two-stage partitioning. First, in DEFT, the vectors are divided into layers. Because each layer comprises a different number of parameters, some layers have a much larger number of gradients than others. Thus, DEFT equally divides a layer that has more than $\frac{n_g}{n}$ into n fractions. This two-stage partitioning is required to balance the computational load between the layers. For simplicity, we refer to all partitioned fractions as layers.

Algorithm 2 presents the pseudocode for the two-stage gradient vector partitioning. In line 1, in the first stage of partitioning, a gradient vector is partitioned into multiple layers. From lines 7 to 26, the second stage of partitioning is performed to partition the larger layers into smaller layers. Owing to the gradient vector partitioning, the number of gradients in each layer is less than $\frac{n_g}{n}$. Two-stage partitioning contributes to a balanced computational

Algorithm 2 Two-stage gradient vector partitioning

Input: g : gradient vector
 n_g : number of gradients in model
 $n_workers$: number of workers
Output: layers: list of partitioned layers

```

1: tensors  $\leftarrow$  partition_by_layers( $g$ );
2: thre_part  $\leftarrow n_g / n\_workers$ ;
3: layers  $\leftarrow []$ ;
4: alloc_pos  $\leftarrow 0$ ;
5: for tensor in tensors do
6:   if tensor.size > thre_part then
7:     quotient  $\leftarrow$  tensor.size /  $n\_workers$ ;
8:     remainder  $\leftarrow$  tensor.size %  $n\_workers$ ;
9:     for i in range( $n\_workers$ ) do
10:      size_layer  $\leftarrow$  quotient;
11:      if remainder > 0 then
12:        size_layer  $\leftarrow$  size_layer + 1;
13:        remainder  $\leftarrow$  remainder - 1;
14:      st_layer  $\leftarrow$  alloc_pos;
15:      alloc_pos  $\leftarrow$  alloc_pos + size_layer;
16:      end_layer  $\leftarrow$  alloc_pos;
17:      layer  $\leftarrow$  make_layer(st_layer, end_layer);
18:      layers.append(layer);
19:   else
20:     st_layer  $\leftarrow$  alloc_pos;
21:     alloc_pos  $\leftarrow$  alloc_pos + tensor.size;
22:     end_layer  $\leftarrow$  alloc_pos;
23:     layer  $\leftarrow$  make_layer(st_layer, end_layer);
24:     layers.append(layer);

```

load between workers during layer allocation because each worker cannot take the majority of the gradients.

4.2 Local k Assignment to Layers

To achieve high accuracy of a DNN model, a sparsifier should select large gradients that can provide significant updates to the model. Each layer has a different gradient norm; consequently, some layers have a larger gradient norm than others. Thus, such layers have a strong possibility of comprising large gradients. Thus, DEFT is designed to select more gradients from layers with large gradient norms. In other words, the density of each layer varies.

To apply different densities to each layer, DEFT computes the gradient norm of each layer and assigns priority to the layers in the order of the norm magnitude. Thus, high priority is assigned to the layer with a large gradient norm. Based on the given priority, DEFT executes an algorithm that assigns a local k to the layers. Algorithm 3 presents the pseudocode for the gradient-norm-based local k assignment. The algorithm begins by initialising k_remain and $norm_remain$ to k and the sum of the gradient norms of all layers, respectively. The assignment is processed in the order of high priority. DEFT assesses the significance of each layer based on the ratio of the gradient norm to $norm_remain$. As shown in line 7,

Algorithm 3 Gradient norm-based local k assignment

Input: layers: list of layers in order of high priority
 n_g : number of gradients in model
 d : density

Output: layers: list of layers with local k

```

1:  $k\_remain \leftarrow n_g \cdot d$ ;
2:  $norm\_remain \leftarrow 0$ ;
3: for layer in layers do
4:    $norm\_remain \leftarrow norm\_remain + layer.norm$ ;
5: for layer in layers do
6:   if  $norm\_remain > 0$  then
7:      $k\_temp \leftarrow k\_remain \cdot \frac{layer.norm}{norm\_remain}$ ;
8:   else
9:      $k\_temp \leftarrow 0$ ;
10:  if  $layer.size < k\_temp$  then
11:     $layer.k \leftarrow layer.size$ ;
12:  else
13:     $layer.k \leftarrow \max(1, k\_temp)$ ;
14:   $k\_remain \leftarrow k\_remain - layer.k$ ;
15:   $norm\_remain \leftarrow norm\_remain - layer.norm$ ;

```

DEFT assigns $k_remain \cdot \frac{layer.norm}{norm_remain}$ as local k to the layer. Whenever each layer is assigned its local k , k_remain and $norm_remain$ are deducted by the assigned k and the layer's gradient norm, respectively. The algorithm is terminated when all layers are assigned their local k .

Because the local k of each layer is determined by the magnitude of its gradient norm, the size of k can be large, although the number of gradients in the layer is small. Similarly, the converse is also possible. Therefore, DEFT can select larger gradients from layers with large gradient norms.

4.3 Layer Allocation to Workers

In this study, we consider the computational cost of finding the top k gradients from n_g -sized gradient vector as $n_g \log k$ [29]. When we expand this concept to partitioned layers, the computational cost of layer x is defined as follows:

$$c_x = n_{g,x} \log k_x, \quad (3)$$

where $n_{g,x}$ and k_x are the number of gradients (vector size) and local k , respectively. Because each layer has a different local k and vector size, the computational cost differs between layers. Consequently, if the same number of layers is distributed to the workers, the computational load will be unbalanced among the workers. Even if only one worker consumes a considerable amount of time for the gradient selection, the entire gradient selection process is delayed. To balance the load between workers, DEFT allocates layers to workers by considering the computational costs of the layers using a bin-packing algorithm. In DEFT, the number of items and bins are the number of layers and workers, respectively. As a bin-packing policy, the layer with the highest cost is allocated to the worker with minimum sum of costs of the layers at that time.

Algorithm 4 Bin-packing-based layer allocation

Input: layers: list of layers
rank: identifier of worker
iteration: current iteration number
 $n_workers$: number of workers

Output: alloc_part: index list of allocated layers

```

1: cycle  $\leftarrow$  iteration %  $n\_workers$ ;
2: curr_part  $\leftarrow$  (cycle + rank) %  $n\_workers$ ;
3: if rank == cycle then
4:   size_bin  $\leftarrow$  zeros( $n\_workers$ );
5:   alloc_bin  $\leftarrow$  empty( $n\_workers$ );
6:   cost  $\leftarrow$  empty(layers.size);
7:   for  $i$  in range(layers.size) do
8:      $cost[i] \leftarrow layers[i].size \cdot \log(layers[i].k)$ ;
9:   for  $i$  in range(layers.size) do
10:    val_max  $\leftarrow$  max(cost);
11:    idx_max  $\leftarrow$  cost.index(max(cost));
12:    bin_min  $\leftarrow$  size_bin.index(min(size_bin));
13:    alloc_bin[bin_min].append(idx_max);
14:    size_bin[bin_min]  $\leftarrow$  size_bin[bin_min] + val_max;
15:    cost[idx_max]  $\leftarrow$  -1;
16: broadcast(alloc_bin, cycle);
17: alloc_part  $\leftarrow$  alloc_bin[curr_part];

```

Algorithm 4 presents the pseudocode for the bin-packing-based layer allocation. While allocating layers to workers, DEFT delegates the decision making to one worker at each iteration. Each worker becomes a decision-making worker once within every n iterations in a cyclic order. The delegated worker computes computational cost c_x for each layer x . Based on the computed computational costs, the delegated worker allocates layers to bins of workers according to the defined bin-packing policy. When all the layers are allocated to bins, the list of bins is broadcast to other workers and each worker accepts the allocated layers from its bin.

DEFT incurs an overhead from layer allocation to workers because the delegated worker requires broadcast communication to share the bins. Nevertheless, the overhead is negligible because the payload is only $4L$ bytes, where L is the number of layers³. We verify how much runtime is occupied by this overhead through experiments detailed in Section 5.

As a result of exclusive layer allocation, each worker selects gradient indices that do not overlap with those of the other workers. Although the local k of each layer differs slightly between workers, the total number of selected gradients is approximately k when the indices are collected. This is because all workers have the same state of the model, and the gradient computed from the same parameter does not vary significantly between workers. Therefore, DEFT can maintain the actual density as the setup density, and a gradient build-up does not occur.

³ $L = \sum_{i=0}^{n-1} n_{l,i}$, where $n_{l,i}$ is the number of layers allocated to worker i .

Algorithm 5 Layer-wise gradient selection

Input: layers: list of layers
 alloc_part: index list of allocated layers
Output: indices: indices of gradients selected by worker
 k: number of gradients selected by worker

```

1:  $k \leftarrow 0$ ;
2: indices  $\leftarrow []$ ;
3: for part in alloc_part do
4:   part_idx  $\leftarrow \text{topk}(\text{layers}[\text{part}].\text{grads}.\text{abs}(), \text{layers}[\text{part}].k)$ ;
5:   part_idx  $\leftarrow \text{part\_idx} + \text{layers}[\text{part}].\text{st\_layer}$ ;
6:   indices.extend(part_idx);
7:    $k \leftarrow k + \text{layers}[\text{part}].k$ ;

```

4.4 Layer-Wise Gradient Selection

To exploit the gradient norm difference between layers, each worker executes an individual Top-k operation in each layer. This method is referred to as layer-wise gradient selection, and its pseudocode is presented in Algorithm 5. According to line 4, each worker executes only the Top-k operations in the layers allocated to them.

To discuss the computational cost of the layer-wise gradient selection of DEFT, we formulated the computational cost for worker i as follows:

$$C_i = \sum_{x=0}^{n_{l,i}-1} c_x = \sum_{x=0}^{n_{l,i}-1} n_{g,x} \log k_x, \quad (4)$$

where $n_{l,i}$ denotes the number of layers allocated to worker i . Because the runtime is determined by the worker that lastly finishes the gradient selection, the computational cost of DEFT is formulated as follows:

$$C(n) = \max_{i \in [0, n-1]} C_i = \max_{i \in [0, n-1]} \sum_{x=0}^{n_{l,i}-1} n_{g,x} \log k_x. \quad (5)$$

The computational cost of the Top-k sparsifier and CLT-k [13] for gradient selection is $n_g \log k$. Therefore, the speedup of DEFT over Top-k sparsifier and CLT-k is denoted by

$$f(n) = \frac{n_g \log k}{C(n)} = \frac{n_g \log k}{\max_{i \in [0, n-1]} \sum_{x=0}^{n_{l,i}-1} n_{g,x} \log k_x}. \quad (6)$$

However, the speedup, $f(n)$, is nondeterministic because the result of the layer allocation is nontrivial. Thus, for the theoretical analysis of the speedup, we consider a trivial case of layer allocation. Assume that a gradient vector is divided into n partitions and each worker is allocated one partition. Then, the computational cost of the gradient selection is as follows:

$$C_{\text{trivial}}(n) = \frac{n_g}{n} \log \frac{k}{n}. \quad (7)$$

Similar to $f(n)$, the speedup of the trivial case over the Top-k sparsifier and CLT-k [13] is denoted by

$$f_{\text{trivial}}(n) = \frac{n_g \log k}{C_{\text{trivial}}(n)} = \frac{n_g \log k}{\frac{n_g}{n} \log \frac{k}{n}}. \quad (8)$$

It should be noted that the partitioning policy of the trivial case is more coarse-grained than that of DEFT. Thus, because $C(n) \leq C_{\text{trivial}}(n)$ is satisfied, we also confirm that $f(n) \geq f_{\text{trivial}}(n)$.

Table 2: Description of each DNN application, where B_l is local batch size and n_e is number of epochs.

Application	Model	Dataset	B_l	n_e
Computer vision	ResNet-18	CIFAR-10	2^5	200
Language modelling	LSTM	WikiText-2	2^5	90
Recommendation	NCF	MovieLens-20M	2^{16}	30

Similarly, because $f_{\text{trivial}}(n) \geq n$ is satisfied, we confirm that the following inequality holds:

$$f(n) \geq f_{\text{trivial}}(n) \geq n. \quad (9)$$

From (9), in terms of gradient selection, we can confirm that the theoretical speedup of DEFT over the Top-k sparsifier and CLT-k [13] exceeds the linear speedup. Therefore, DEFT is massively scalable owing to its outstanding speedup. We will provide the verification for (9) through experiments detailed in Section 5.

5 EVALUATION**5.1 Methodology**

System configuration. All experiments in this study were conducted on a cluster equipped with 32 GPUs; the number of GPUs used in each experiment was different. A cluster comprises 8 nodes, and each has 4 NVIDIA Tesla V100 GPUs, 2 20-core Intel Xeon Gold 6230 @ 2.10 GHz CPUs, and 384 GB DDR4 memory. Each worker process was executed using mpirun⁴ of Open MPI 4.0.5 [6] and run on one GPU with CUDA 10.1 [5].

Models and datasets. We evaluate the performance of DEFT and other sparsifiers (CLT-k and Top-k sparsifier) for three types of DNN applications: 1) computer vision via ResNet-18 [17] on CIFAR-10 [22], 2) language modelling via long short-term memory (LSTM) [19] on WikiText-2 [25], and 3) recommendation system via neural collaborative filtering (NCF) [18] on MovieLens-20M [1]. Table 2 lists the descriptions of the DNN applications used in our evaluation.

Implementation. We implemented DEFT and other approaches on top of the deep learning framework PyTorch 1.5 [7]. The communication routine for distributed training was implemented using the distributed communication package of PyTorch, and NCCL 2.4 [4] was adopted as backend to support the multi-GPU and multi-node communication primitives optimised for NVIDIA GPUs and networking: broadcast, all-gather, and all-reduce for distributed GPU training. The source code includes everything required to reproduce the results of this study, and is available publicly at <https://github.com/kljp/deft>.

Metrics. The metrics used for each type of performance evaluation are as follows:

- Convergence performance: Test accuracy on the computer vision application, test perplexity on language modelling application, and best Hit Rate at 10 (hr@10) [2, 3] on the recommendation system application.

⁴In all experiments, MPI is only used for multi-process execution to automatically assign an exclusive rank to each worker.




Figure 3: Convergence performance of sparsifiers on 16 GPUs.




Figure 4: Sparsification performance of sparsifiers on 16 GPUs. The Y-axis indicates the actual density measured over training iterations.

- Sparsification performance: The actual density was measured to evaluate whether each sparsifier could maintain the density set by the user.
- Error minimisation performance: The error was measured to evaluate the significance of the gradient selection of each sparsifier.

5.2 Performance Evaluation

Convergence performance. Fig 3 shows the convergence performance of each sparsifier in the DNN applications. In every application, all sparsifiers attain the convergence point of the non-sparsified distributed training, which is a distributed training without sparsification. However, the convergence rates of the sparsifiers are different. Among the three sparsifiers, the Top-k sparsifier showed the fastest convergence rate, which was close to the non-sparsified distributed training, whereas DEFT and CLT-k showed rates lower than those of the Top-k sparsifier.

Sparsification performance. Despite the fast convergence rate, the Top-k sparsifier showed weakness in terms of sparsification performance owing to the gradient build-up. Fig 4 presents the sparsification performance of the sparsifiers on DNN applications. The Top-k sparsifier shows an excessively high actual density, whereas DEFT and CLT-k maintain the actual density at the user-set value. In this evaluation, all sparsifiers were tested on 16 GPUs, and the actual density of the Top-k sparsifier attained values of 13.59, 14.18,

and 5.31 times of the setup density for the computer vision, language modelling, and recommendation system applications, respectively. In Figure 4c, the actual density of the Top-k sparsifier in the recommendation system is higher than that of the other sparsifiers, but not significantly higher as in the other two applications. This is because the Top-k sparsifier selected 53.57% of all gradients, which was sufficiently high for the workers' gradient indices to overlap. Although DEFT showed fluctuations in the actual density in this application, it was sufficiently stable because the average density was 0.1003, which is approximate to the setup density.

Thus, it is evident that the convergence rate of the Top-k sparsifier was accomplished by a very high actual density caused by the gradient build-up. The fundamental purpose of gradient sparsification is to reduce the communication traffic. Thus, because of high actual density, Top-k sparsifier may not be an effective method despite the fast convergence rate. In contrast, DEFT can attain the convergence point of the non-sparsified distributed training at a rate similar to that of CLT-k while maintaining the actual density as the setup density.

Error Minimisation Performance. Figure 5 presents the error-minimisation performance of the sparsifiers in DNN applications. Owing to the gradient build-up, the Top-k sparsifier shows a considerably low error compared to DEFT and CLT-k during all training iterations. This implies that the Top-k sparsifier can minimise the objective function of the DNN model faster than DEFT and CLT-k.




Figure 5: Error minimisation performance of sparsifiers on 16 GPUs. The Y-axis indicates the error, which is the average of local errors of workers.




Figure 6: Comparison of error minimisation performance between DEFT and Top-k sparsifier on 16 GPUs. The density of DEFT was set to a level similar to the actual density of the Top-k sparsifier. The density of DEFT was set to (a) $d = 0.1$ and (b) $d = 0.01$.

Consequently, DEFT and CLT-k terminate with higher errors than the Top-k sparsifier at the end of model training. In Figure 5a, the sudden decrease in error after iteration 14,600 is because of learning rate decay.

As noted, the low error of the Top-k sparsifier during the training iterations is accomplished by a large number of selected gradients. To fairly compare the DEFT and Top-k sparsifiers with similar levels of actual density, we conducted additional experiments on the computer vision and language modelling applications. In these experiments, we scaled up the setup density of the DEFT to 10 times that in the previous experiments: $d = 0.1$ and $d = 0.01$ for the computer vision and language modelling applications, respectively. It should be noted that these setup densities are still lower than the actual densities of the Top-k sparsifier owing to the gradient build-up. Figure 6 presents the error minimisation performance between the DEFT and Top-k sparsifiers. By setting the density of the DEFT comparable to that of the Top-k sparsifier, the DEFT showed an error minimisation performance similar to that of the Top-k sparsifier.

Although DEFT selects a different number of gradients by layers, it achieves a similar level of error as the other sparsifiers. That is,




Figure 7: Training time breakdown of sparsifiers on LSTM with WikiText-2. The experiments were conducted on 16 GPUs. The training time is the average wall-clock time for one iteration.

DEFT can select significant gradients in the partitioned gradient vectors by considering the gradient norm of each layer, as if the CLT-k and Top-k sparsifiers find gradients in one Top-k operation in the entire gradient vector. Thus, our insight that a sparsifier should select more gradients from layers with large gradient norms is significant. Furthermore, DEFT attains the convergence point of the non-sparsified distributed training while maintaining actual density regardless of the number of workers by eliminating gradient build-up. Therefore, we conclude that DEFT has a similar level of convergence, and error minimisation performances compared to those of the other sparsifiers while achieving sparsification performance of CLT-k.

5.3 Training Efficiency

Training time breakdown. Figure 7 shows a breakdown of the training time for one iteration. For each sparsifier, the wall-clock time of each iteration was measured by the slowest worker of that iteration, and the average time was calculated across all iterations. This experiment was repeated 4 times using different seeds for each




Figure 8: Convergence performance of DEFT by different setup of density on LSTM with WikiText-2. The experiments were conducted on 16 GPUs.

sparsifier, and the average wall-clock time in Figure 7 was obtained by averaging the values from 4 executions.

In Figure 7, the measured training time comprises times for forward propagation, backward propagation, gradient selection, and communication; however, DEFT also includes the partitioning overhead. DEFT requires a significantly less time for gradient selection than the CLT-k and Top-k sparsifiers. This speedup was achieved using gradient vector partitioning with a load-balanced distribution of the partitioned layers.

The communication in DEFT also reduced compared with that of the CLT-k and Top-k sparsifiers. This result can be explained by the communication time cost of the Top-k sparsifier [34]. The communication time cost of the Top-k sparsifier is $\log(n)\alpha + 2(n-1)k\beta$, where α and β are the system-dependent constants. In DEFT, owing to gradient partitioning, k in the cost can be reduced to $\max_{i \in [0, n-1]} \sum_{x=0}^{n_{l,i}-1} k_x$. Accordingly, the communication time cost of DEFT is derived as $\log(n)\alpha + 2(n-1)(\max_{i \in [0, n-1]} \sum_{x=0}^{n_{l,i}-1} k_x)\beta$. This time cost is evidence that DEFT achieves a speedup over CLT-k and Top-k sparsifiers in communication.

Furthermore, as discussed in Section 4.3, the overhead from the layer allocation of DEFT is negligible because it is a very small fraction of the runtime in one iteration.

Convergence performance by density. Figure 8 presents the convergence performance of DEFT for different densities. The results reveal that the convergence rate was slightly slower when the density was set to 0.001; however, there was no significant difference after epoch 45. Thus, DEFT can achieve reasonable convergence performance with robustness to varying densities.

5.4 Scalability

Computational speedup by scale. Figure 9 presents the computational speedup of the layer-wise gradient selection of DEFT over the Top-k gradient selection by scaling out the cluster. In Figure 9, if the number of workers is one, this case corresponds to the Top-k gradient selection adopted in the CLT-k and Top-k sparsifiers. The results reveal that the speedup of DEFT over the Top-k gradient selection increases as the number of workers increases.




Figure 9: Computational speedup of layer-wise gradient selection of DEFT over Top-k gradient selection by scale-out on LSTM with WikiText-2.




Figure 10: Convergence performance of DEFT by scale-out on LSTM with WikiText-2.

It is noteworthy that the speedup of DEFT is higher than the linear speedup. As discussed in (9) in Section 4.4, the theoretical speedup of trivial partitioning ($f_{trivial}(n)$) is higher than the linear speedup. Because $f_{trivial}(n)$ is the worst-case speedup of DEFT, the speedup of DEFT can be higher than the linear speedup. Furthermore, it can be seen that the rate of increase of the DEFT speedup becomes significantly higher as the number of workers increases because $\frac{f_{trivial}(n)}{n}$ increases as n increases. Therefore, the layer-wise gradient selection of DEFT has a tremendous advantage in the scale-out of the cluster.

Convergence performance by scale. Figure 10 shows the convergence performance of the DEFT by a scale-out of the cluster. All the experiments were conducted with a setup density 0.001. Although the convergence rate differed between cases, the test perplexity of every case eventually reached the convergence point of the non-sparsified distributed training. Because DEFT shows a high speedup over Top-k gradient selection-based sparsifiers while

achieving a similar convergence rate, it is evident that the consistent convergence performance regardless of the scale of the cluster is a tremendous advantage in terms of the scalability of DEFT.

6 CONCLUSION

In this study, we proposed DEFT, which partitions the gradient vector into layers and assigns different computational loads to each layer, considering the gradient norm of the layer. The design of DEFT comprises a two-stage gradient vector partitioning, local k assignment to layers, layer allocation to workers, and layer-wise gradient selection. Using these components, DEFT can achieve high performance in terms of convergence, sparsification, and error minimisation. Additionally, it provides a remarkable training efficiency owing to its reduced computational and communication costs. In particular, it exhibits a high speedup of gradient selection over the CLT-k and Top-k sparsifiers. Furthermore, it has a remarkable scalability owing to the speedup of gradient selection over linear speedup, while achieving high convergence performance.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their insightful feedback. This work was jointly supported by the Korea Institute of Science and Technology Information (KSC-2022-CRE-0406), BK21 FOUR program (NRF5199991014091), and Basic Science Research Program (2021R1F1A1062779) of National Research Foundation of Korea.

REFERENCES

- [1] 2015. MovieLens 20M dataset. <https://grouplens.org/datasets/movielens/20m/>.
- [2] 2018. Neural collaborative filtering source code. <https://github.com/yihongchen/neural-collaborative-filtering/>.
- [3] 2021. Hard-threshold sparsifier source code. <https://github.com/sands-lab/rethinking-sparsification/>.
- [4] 2023. NCCL. <https://developer.nvidia.com/nccl/>.
- [5] 2023. NVIDIA CUDA. <https://developer.nvidia.com/cuda-toolkit/>.
- [6] 2023. Open MPI. <https://www.open-mpi.org/>.
- [7] 2023. PyTorch. <https://pytorch.org/>.
- [8] Lusine Abrahamyan, Yiming Chen, Giannis Bekoulis, and Nikos Deligiannis. 2021. Learned gradient compression for distributed deep learning. *IEEE Transactions on Neural Networks and Learning Systems* 33, 12 (2021), 7330–7344.
- [9] Saurabh Agarwal, Hongyi Wang, Kangwook Lee, Shivaram Venkataraman, and Dimitris Papailiopoulos. 2021. Adaptive gradient communication via critical learning regime identification. *Proceedings of Machine Learning and Systems* 3 (2021), 55–80.
- [10] Alham Fikri Aji and Kenneth Heafield. 2017. Sparse communication for distributed gradient descent. *arXiv preprint arXiv:1704.05021* (2017).
- [11] Dan Alistarh, Torsten Hoefer, Mikael Johansson, Nikola Konstantinov, Sarit Khirirat, and Cédric Renggli. 2018. The convergence of sparsified gradient methods. *Advances in Neural Information Processing Systems* 31 (2018).
- [12] Chia-Yu Chen, Jungwook Choi, Daniel Brand, Ankur Agrawal, Wei Zhang, and Kailash Gopalakrishnan. 2018. Adacomp: Adaptive residual gradient compression for data-parallel distributed training. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 32.
- [13] Chia-Yu Chen, Jiamin Ni, Songtao Lu, Xiaodong Cui, Pin-Yu Chen, Xiao Sun, Naigang Wang, Swagath Venkataramani, Vijayalakshmi Viji Srinivasan, Wei Zhang, and Kailash Gopalakrishnan. 2020. Scalecom: Scalable sparsified gradient compression for communication-efficient distributed training. *Advances in Neural Information Processing Systems* 33 (2020), 13551–13563.
- [14] Nikoli Dryden, Tim Moon, Sam Ade Jacobs, and Brian Van Essen. 2016. Communication quantization for data-parallel training of deep neural networks. In *2016 2nd Workshop on Machine Learning in HPC Environments (MLHPC)*. IEEE, 1–8.
- [15] Luciano Floridi and Massimo Chiriatti. 2020. GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines* 30, 4 (2020), 681–694.
- [16] Anil Gaihre, Da Zheng, Scott Weitz, Lingda Li, Shuaiwen Leon Song, Caiwen Ding, Xiaoye S Li, and Hang Liu. 2021. Dr. Top-k: delegate-centric Top-k on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. 1–14.
- [17] 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.
- [18] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [19] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation* 9, 8 (1997), 1735–1780.
- [20] Samuel Horvóth, Chen-Yu Ho, Ludovik Horvath, Atal Narayan Sahu, Marco Canini, and Peter Richtárik. 2022. Natural compression for distributed deep learning. In *Mathematical and Scientific Machine Learning*. PMLR, 129–141.
- [21] Nikita Iykin, Daniel Rothchild, Enayat Ullah, Ion Stoica, Raman Arora, et al. 2019. Communication-efficient distributed SGD with sketching. *Advances in Neural Information Processing Systems* 32 (2019).
- [22] Alex Krizhevsky. 2009. Learning multiple layers of features from tiny images. *Technical report, University of Toronto* (2009).
- [23] 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).
- [24] Ahmed M Abdelmoniem, Ahmed Elzanaty, Mohamed-Slim Alouini, and Marco Canini. 2021. An efficient statistical-based gradient compression technique for distributed training systems. *Proceedings of Machine Learning and Systems* 3 (2021), 297–322.
- [25] Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843* (2016).
- [26] Corby Rosset. 2020. Turing-NLG: A 17-billion-parameter language model by Microsoft. *Microsoft Blog* (2020).
- [27] Atal Sahu, Aritra Dutta, Ahmed M Abdelmoniem, Trambak Banerjee, Marco Canini, and Panos Kalnis. 2021. Rethinking gradient sparsification as total error minimization. *Advances in Neural Information Processing Systems* 34 (2021), 8133–8146.
- [28] Frank Seide, Hao Fu, Jasha Droppo, Gang Li, and Dong Yu. 2014. 1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns. In *Fifteenth annual conference of the international speech communication association*.
- [29] Anil Shanbhag, Holger Pirk, and Samuel Madden. 2018. Efficient top-k query processing on massively parallel hardware. In *Proceedings of the 2018 International Conference on Management of Data*. 1557–1570.
- [30] Shaohuai Shi, Xiaowen Chu, Ka Chun Cheung, and Simon See. 2019. Understanding top-k sparsification in distributed deep learning. *arXiv preprint arXiv:1911.08772* (2019).
- [31] Shaohuai Shi, Zhenheng Tang, Xiaowen Chu, Chengjian Liu, Wei Wang, and Bo Li. 2020. A quantitative survey of communication optimizations in distributed deep learning. *IEEE Network* 35, 3 (2020), 230–237.
- [32] Shaohuai Shi, Zhenheng Tang, Qiang Wang, Kaiyong Zhao, and Xiaowen Chu. 2019. Layer-wise adaptive gradient sparsification for distributed deep learning with convergence guarantees. *arXiv preprint arXiv:1911.08727* (2019).
- [33] Shaohuai Shi, Qiang Wang, Xiaowen Chu, Bo Li, Yang Qin, Ruihao Liu, and Xinxiao Zhao. 2020. Communication-efficient distributed deep learning with merged gradient sparsification on GPUs. In *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 406–415.
- [34] Shaohuai Shi, Qiang Wang, Kaiyong Zhao, Zhenheng Tang, Yuxin Wang, Xiang Huang, and Xiaowen Chu. 2019. A distributed synchronous SGD algorithm with global top-k sparsification for low bandwidth networks. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*. IEEE, 2238–2247.
- [35] Shaohuai Shi, Kaiyong Zhao, Qiang Wang, Zhenheng Tang, and Xiaowen Chu. 2019. A Convergence Analysis of Distributed SGD with Communication-Efficient Gradient Sparsification. In *IJCAI*. 3411–3417.
- [36] Mohammad Shoeybi, Mostafa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053* (2019).
- [37] Nikko Ström. 2015. Scalable distributed DNN training using commodity GPU cloud computing. (2015).
- [38] Hanlin Tang, Yao Li, Ji Liu, and Ming Yan. 2021. ErrorCompensatedX: error compensation for variance reduced algorithms. *Advances in Neural Information Processing Systems* 34 (2021), 18102–18113.
- [39] Thijs Vogels, Sai Praneeth Karimireddy, and Martin Jaggi. 2019. PowerSGD: Practical low-rank gradient compression for distributed optimization. *Advances in Neural Information Processing Systems* 32 (2019).
- [40] Mingchao Yu, Zhifeng Lin, Krishna Narra, Songze Li, Youjie Li, Nam Sung Kim, Alexander Schwing, Murali Annamalai, and Salman Avestimehr. 2018. GradiVec: Vector quantization for bandwidth-efficient gradient aggregation in distributed cnn training. *Advances in Neural Information Processing Systems* 31 (2018).
- [41] Chiyuan Zhang, Samy Bengio, and Yoram Singer. 2019. Are all layers created equal? *arXiv preprint arXiv:1902.01996* (2019).