

MiCS: Near-linear Scaling for Training Gigantic Model on Public Cloud

Zhen Zhang*
Johns Hopkins University
zzhen1@jhu.edu

Justin Chiu
Amazon
justchiu@amazon.com

Mu Li
Amazon Web Services
mli@amazon.com

Shuai Zheng
Amazon Web Services
shzheng@amazon.com

George Karypis
Amazon Web Services
gkarypis@amazon.com

Xin Jin
Peking University
xinjinpku@pku.edu.cn

Yida Wang
Amazon Web Services
wangyida@amazon.com

Trishul Chilimbi
Amazon
trishulc@amazon.com

ABSTRACT

Existing general purpose frameworks for gigantic model training, i.e., dense models with billions of parameters, cannot scale efficiently on cloud environment with various networking conditions due to large communication overheads. In this paper, we propose MiCS, which Minimizes the Communication Scale to bring down communication overhead. Specifically, by decreasing the number of participants in a communication collective, MiCS can utilize heterogeneous network bandwidth, reduce network traffic over slower links, reduce the latency of communications for maintaining high network bandwidth utilization, and amortize expensive global gradient synchronization overhead. Our evaluation on AWS shows that the system throughput of MiCS is up to 2.89× that of the state-of-the-art large model training systems. MiCS achieves near-linear scaling efficiency, which is up to 1.27× that of DeepSpeed. MiCS allows us to train a proprietary model with 100 billion parameters on 512 GPUs with 99.4% weak-scaling efficiency, and it is able to saturate over 54.5% theoretical computation power of each GPU on a public cloud with less GPU memory and more restricted networks than DGX-A100 clusters.

PVLDB Reference Format:

Zhen Zhang, Shuai Zheng, Yida Wang, Justin Chiu, George Karypis, Trishul Chilimbi, Mu Li, and Xin Jin. MiCS: Near-linear Scaling for Training Gigantic Model on Public Cloud. PVLDB, 16(1): 37 - 50, 2022. doi:10.14778/3561261.3561265

1 INTRODUCTION

There is a growing body of research showing that large Deep Learning (DL) models deliver superior accuracy in areas such as natural language processing (NLP) [17, 54], speech recognition (SR) [11, 13, 70], and computer vision (CV) [15, 67, 68]. This has resulted in a more than 1000× increase in the size of the DL models that

are commonly trained with many of them having several hundred billion parameters. The high computational requirement associated with training DL models has led to effective and simple parallelization approaches based on data parallelism (DP) [26, 42, 43, 52, 57]. However, many of these approaches cannot be applied for training gigantic DL models, as their memory requirements exceed the amount of GPU memory.

A common way to train gigantic DL models is to use model-parallelism (MP) that decomposes the computation across the devices by partitioning the neural network architecture (i.e., the model). As a result of this network partitioning, the *model states* (i.e., the memory storing the model parameters, gradients, and optimizer states) are also partitioned across the devices, and as such it overcomes DP’s memory-related limitations. Unfortunately, existing MP frameworks require users to substantially modify the logic of their training code and add specific primitives [25, 28, 35, 53, 54]. In addition, many of the MP frameworks are specifically designed for certain types of neural network architectures [38, 54] and cannot be directly used for arbitrary architectures. However, the idea of partitioning the model states across different devices is essential for enabling large model training and was recently incorporated into DP by the development of ZeRO [47]. ZeRO, which is implemented in distributed systems DeepSpeed [47] and FairScale [18], evenly partitions the model states across the entire training cluster, enabling the training of very large models while retaining DP’s simplicity, ease of use, and generality.

ZeRO was designed for clusters using nodes based on NVIDIA’s DGX-2 or DGX-A100 multi-GPU systems [37, 45]. These nodes are connected via high-bandwidth low-latency InfiniBand leading to clusters whose intra- and inter-node GPU-to-GPU bandwidth is nearly balanced (intra-node bandwidth is about 3× faster than inter-node). ZeRO takes advantage of this balanced network to treat all GPU devices of the cluster equivalently and to partition the model states across the entire cluster. As a result, whenever during the forward or backward phase a parameter tensor is required for the computations, a collective communication operation needs to be performed that involves all devices of the entire cluster (§2.2). Training clusters in public cloud environment are not always equipped with high-speed InfiniBand networks as DGX

*Work partly done as an intern at Amazon Web Services.

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing info@vldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 16, No. 1 ISSN 2150-8097. doi:10.14778/3561261.3561265

nodes have. For example, cloud instances with V100 GPUs are typically paired with 100Gbps networks [3–5, 20], in which case the bandwidth is less balanced (intra-node bandwidth is about 24× faster than inter-node). In such scenarios, ZeRO is not well suited. By treating these devices equivalently and not accounting for the heterogeneous and hierarchical nature of the inter-node network, ZeRO fails to take advantage of the faster intra-node networks. Moreover, by partitioning the model states across the entire cluster, even when the model states can fit in the memory of a subcluster, ZeRO unnecessarily incurs high communication cost of collective communications, because of the low effective bandwidth caused by high algorithmic latency. And the communication overhead of ZeRO grows larger as the size of the cluster scales up (§2.3).

To surmount the aforementioned challenges, we propose MiCS, following a core design principle: to reduce the number of communicating participants, i.e., *communication scale*, as much as possible. By minimizing the scale, MiCS reduces the latency and the data volume transmitted over slow inter-node links. We design and implement three components to realize the design principle for reducing communication overheads.

- *Scale-aware model partitioning*. Instead of using all devices as a single group for holding the model states, MiCS divides all devices into partition groups. Each group holds a complete copy of the model states. Within each group, the model states are partitioned. Thus most frequent parameter gatherings are operated at the *scale* of each group (§3.2).
- *Hierarchical communication strategy*. Hierarchical communication allows us to parallelize multiple inter-node collective communications and reduce the scale of each collective communication. It reduces the aggregated traffic over the inter-node links, thus leading to lower communication cost (§3.3).
- *2-hop gradient synchronization*. Unlike ZeRO that synchronizes gradients over all devices for each micro-step, MiCS only synchronizes gradients within the partition group until the gradient accumulation boundary is reached. At the gradient accumulation boundary, gradients are synchronized across the partition groups. As a result, MiCS reduces the synchronization cost significantly by amortizing the cost to multiple micro-steps. (§3.4).

Our thorough evaluation shows significant system throughput and scaling efficiency improvement of MiCS on public clouds like AWS. On V100 GPU clusters with 100Gbps network, the system throughput of MiCS is 2.89× larger than that of DeepSpeed, which is the state-of-the-art DP framework for large model training. On A100 GPU clusters with 40GB memory per GPU and 400Gbps networks, MiCS is up to 2.74× as fast as DeepSpeed. Compared to Megatron-LM-3D, a state-of-the-art system specialized for training Transformer models, MiCS achieves up to 30.1% better throughput. MiCS gets near-linear (e.g., 99.4%) weak scaling efficiency in the cloud, which is up to 27% better than DeepSpeed. MiCS has been deployed to train a proprietary model with 100 billion (B) parameters, saturating over 170 TeraFLOP/s (TFLOPS) per A100 GPU with activation checkpointing at the scale of 512 GPUs.

In summary, this paper makes the following contributions.

- We identify the root problem—overwhelming communication overhead—that prevents DP-based model partitioning from efficiently scaling out on clusters with 100Gbps or 400Gbps network interfaces with relatively higher latency than InfiniBand [74].
- We design and implement a system MiCS that minimizes the communication scale to reduce the communication overhead.
- We evaluate MiCS thoroughly to justify the benefits of minimizing communication scale on clusters with up to 512 GPUs.

2 BACKGROUND AND MOTIVATION

In this section, we briefly review deep learning model training (§2.1) and how existing works tackle the large model training challenges (§2.2), and discuss its major limitation in the context of public clouds (§2.3). We then present the intuition that motivates our design (§2.4).

2.1 Model Training

Deep learning model training process mainly consists of three phases, i.e., forward computation, backward computation, and parameter updating. In order to train the model faster, we can harness the computing power of multiple machines. A gradient synchronization step is performed before updating the model parameters to ensure all workers will use the same set of parameters to evaluate the incoming new training samples.

Deep learning model training is memory consuming as it needs to hold the model states including model parameters, gradients from backward computation, and optimizer states for parameter updating. Because of the limited on-device memory resource, *activation checkpointing* and *gradient accumulation* are typically enabled. Activation checkpointing discards the activation outputs from the forward phase and requires activation recomputation in the backward phase. Gradient accumulation divides one large data batch into multiple small micro-batches to reduce the memory footprint of storing activation outputs. However, for models with billions of parameters, these two techniques alone are not sufficient. Many solutions targeting at gigantic model training are thus proposed.

2.2 Gigantic Model Training

In this paper, we use the term “gigantic model” to refer to those Deep Neural Network models that consist of billions of densely connected parameters, which means both the size of the model and the per-sample computation of the model, i.e., floating-point operations (FLOPs), are “gigantic”. Nowadays, the commonly adopted models that fall into this category are transformer-based models [9, 37, 44, 45, 56, 69] and latest wide computer vision models [67].

Traditionally, developers use model-parallel (MP) distributed training for gigantic model training. The basic idea is to distribute the model parameters and computations across multiple devices for each training sample. Thus, the memory for storing model states is also distributed across devices. This way of distributing computations comes with issues. Tensor model parallelism as one MP method requires lots of communications during computation [54]. On the other hand, pipeline MP strategy is advocated with smaller communication overheads, but it suffers from pipeline bubbles and causes under-utilization. Besides, MP solutions are not directly compatible with common frameworks like PyTorch or Tensorflow, and they require non-trivial engineering effort from the user side.

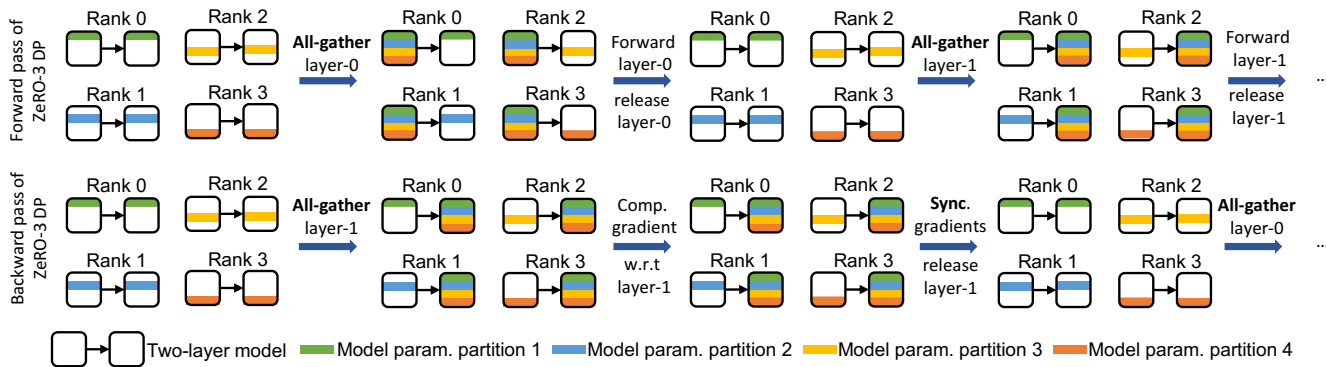


Figure 1: Forward and backward passes of ZeRO-3 Data Parallelism on a cluster with four devices; for brevity, only the model parameter states are shown in the figure; “Comp.”: compute, “Sync.”: synchronize, “param.”: parameter; communications are marked with bold font; a two-layer model is used here for illustration purposes.

Lastly, some of the MP designs [37, 38] are model-specific solutions, making them hard to generalize.

Compared to the MP solutions, ZeRO [45] powered data-parallel (DP) solutions are general to various models and do not require model refactoring. ZeRO partitions the *model states* onto all devices on the cluster to reduce the memory consumption on each device. ZeRO has three different stages, corresponding to three different levels of memory reduction: ZeRO-1 partitions optimizer states only; ZeRO-2 partitions gradients and optimizer states; ZeRO-3 partitions all three states, i.e., parameters, gradients and optimizer states, evenly across all devices on the training cluster. The full-fledged ZeRO allows us to train the extremely large models when we have a large enough cluster. However, we have to pay communication costs for gathering model parameters during both forward and backward. Figure 1 illustrates the forward and backward passes in ZeRO-3 powered DP, in which the parameters of each layer are partitioned across all the ranks. Here we use the same convention in the high-performance computing (HPC) community where we use a rank number to identify a computing device. Before computing the activations or gradient for a layer, all parameters of this layer are gathered back by all-gather communication. After computing the gradients on each rank with its own part of the data, the gradients are synchronized and partitioned across all ranks using reduce-scatter communication, which aggregates gradients among all ranks and partitions the gradients at the same time. The gradient partition is necessary for gigantic models with billions of parameters due to the limited memory on each rank.

ZeRO-Offload [49] and ZeRO-Infinity [46] are two extensions to ZeRO-3. These two systems offload model parameters, gradients, and optimizer states to CPU memory and NVMe SSDs. Both systems suffer from the same communication overheads as ZeRO-3, which will be discussed in the next subsection.

2.3 Communication Overhead

ZeRO’s model state partitioning mechanism results in the heavy use of collective communication for gathering model states, which is demonstrated in Figure 1. Specifically, ZeRO-3 transmits $3(n -$




Figure 2: Effective bandwidths measured with all-gather.

$1)M/n$ bytes [45] in forward and backward passes, where M denotes the size of the model in bytes, n denotes the number of devices. The transmitted data volume is as large as tens to hundreds of gigabytes for models with tens to hundreds of billions of parameters. The cost of transmitting these data crossing the entire cluster cannot be easily hidden via pipelining the communication and computation. Our timeline measurement shows that for a BERT model with 10B parameters, parameter gathering could take $2.85\times$ more time than computation in forward pass. Similar expensive communications also exist in the backward computation and gradient synchronization, which hurts the performance of ZeRO especially when the network bandwidth between devices is less preferable.

There are two main factors that contribute to the costly communications when using ZeRO-3 on the cloud. Firstly, at the hardware level, many of the available internode network bandwidths and latency of cloud-based GPU clusters are not as good as DGX systems [3, 5, 20, 74]. Moreover, unlike on-premise clusters, the network topology of cloud-based clusters is out of users’ control, which could negatively impact the network performance [32, 33]. Secondly, at the algorithmic side, the latency of collective algorithms for communication has a positive correlation with the communication scale and the startup time for transmission [10]¹. Therefore, as the scale grows, the latency becomes more significant and hurts

¹The latency of tree algorithms is bounded with $\lceil \log_2(p) \rceil \alpha$, and the ring algorithms have a latency term $2p\alpha$, where p denotes the number of participants and α denotes the startup time for transmission, §7.1.7 in [10]

the performance of communication at a large scale. In addition, previous studies [71] suggest that network bandwidth may not be the performance bottleneck of distributed model training. Based on our measurements, as the cluster size grows, we need larger message sizes to saturate the bandwidth. Figure 2 shows that small-size message such as 128MB obtains poor bandwidth utilization on 16 and 32 nodes.

In practice, we may not be able to always communicate large messages due to the memory constraint. Instead, it is better to control the communication scale to improve the bandwidth utilization, especially on the cloud with less favorable network conditions.

2.4 Motivation

As mentioned in §2.3, the frequent communication among all devices significantly hampers the training performance of ZeRO powered DP solutions. This motivates us to design a new system that reduces the cost of communications while preserving the generality and usability advantages. We found the communication overhead can be effectively reduced by shrinking the communication scale, i.e., reducing the number of participants in a collective communication. With the reduced communication scale, the majority of communications are restricted to a smaller group of devices. This allows us to maintain high bandwidth utilization in communications for various sized messages, as shown in Figure 2. In addition, because the transmitted data volume is positively correlated to the number of participants, reducing the communication scale reduces the data volume. We give detailed descriptions of our methodology in the next section.

3 MICS DESIGN

MiCS is designed for training large models on the public cloud. The overarching goal of MiCS’s design is to reduce the scale of communication. The reduced scale allows us to exploit heterogeneous network bandwidth, and to reduce the network traffic transmitted over slow links. To effectively reduce the communication scale, we propose three components named *small-scale model partitioning*, *hierarchical communication* and *2-hop gradient synchronization*. For each of them, we explain the motivation, the methodology, and the analysis of our design.

3.1 Notation

We define the notations used in this section as follows:

- n : Number of devices or ranks in the cluster.
- k : Number of devices on each computational node.
- M : Size of a model in bytes.
- p : Number of devices for holding a model replica.
- s : Number of micro-steps.
- B_g : Effective communication bandwidth among devices belonging to the group g . We define the effective communication bandwidth as the bandwidth measured using collective communication. Effective communication bandwidth takes algorithm latency into account. Thus it is smaller than the theoretical bandwidth of hardware specification. For a fixed message size, when the number of nodes increases, the effective bandwidth shrinks.
- C : Time cost.




Figure 3: Model training on n devices, with every 2 devices together holding a copy of the model states. Each *partition group* maintains a copy of the entire model states, and devices in a *replication group* hold the same part of model states.

3.2 Scale-aware Model Partitioning

Partitioning model states across all the devices causes significant communications overheads during training. Such communication overheads scale with the number of participants in a single collective communication (§2.3). To reduce the communication overheads, we consider distributing the model states over a subset of devices to reduce the scale of the communication. A modern single computing device typically has tens of gigabytes of memory, and tens of them provide sufficient memory for a model with tens of billions of parameters. For example, a model with 10 billion parameters takes about 160GB of memory when training with Adam optimizer using mixed-precision. Partitioning the model states across 8 V100 (32GB) GPUs is already more than enough. By using 8 V100 GPUs instead of all the devices for holding one model states replica, we can effectively reduce the scale of communication. If 8 GPUs are located on a single node, then we can leverage high-speed intra-node connections such as NVLink/NVSwitch to perform the most communications. Next, we give a general form of model states partitioning in our system and provide an analysis of the benefits.

In MiCS, we divide all the devices into multiple groups and partition the model states within each group. Every group has the same number of devices and holds a complete replica of the model states in training. We call these groups *partition groups*. Each device is tagged with a local group rank. Devices with the same local group rank form another type of group, named *replication group*, and they hold the same part of the model states. In Figure 3, we give an example that the model states are partitioned onto two devices. Thus every two devices with consecutive rank numbers form a partition group. The devices ranked with odd numbers and even numbers form two replication groups separately. During the training, when a parameter tensor is needed for either the forward or backward computation, MiCS invokes all-gather collective to gather the corresponding model parameters distributed within each partition group. After the gradients are computed on each device, MiCS uses all-reduce collective to aggregate the gradients, and then it partitions the gradients within each partition group.

Now we give the performance analysis of our partitioning strategy in terms of the cost of all-gather. We assume the iteration time is bounded by the communications, which is true based on our measurements (§2.3). The time cost of ZeRO-3’s partition-to-all strategy is $C_{\text{all}} = ((n - 1)M)/(nB_{\text{all}})$, where B_{all} denotes effective bandwidth among all devices. The time cost of MiCS is $C_{\text{MiCS}} = ((p - 1)M)/(pB_{\text{part}})$. Here we assume all partition groups




Figure 4: Differences between hierarchical communication and vanilla communication.

have the same intra-group bandwidth, and we denote this bandwidth by B_{part} . Because the value of function $(x - 1)/x$ increases when $x \geq 1$ and $p \leq n$, we have the following inequality for the ratio of two costs.

$$\frac{C_{\text{all}}}{C_{\text{MiCS}}} \geq \frac{B_{\text{part}}}{B_{\text{all}}}.$$

For models that can be partitioned to devices located on a single node, only local high-speed NVLinks connections are used for all-gather. Based on the measurement on 64 GPUs spread across 8 computational nodes, we get $B_{\text{part}} \approx 128\text{GB/s}$ and $B_{\text{all}} \approx 11\text{GB/s}$. Thus, the cost ratio can be as large as 11.6. For the case where there are 32 nodes in total and each partition group consists of 4 nodes, the cost ratio is ranging from 2.7 to 4.9 based on our measurements presented in §2.3. For the model states that can be partitioned within 4 nodes, we can expect about 63.6% to 91.3% time reduction for parameter gathering with our partitioning strategy.

3.3 Hierarchical Communication Strategy

When the model size grows, it requires more devices to hold the model states for the training. If the required devices span multiple computational nodes, inter-node communication is needed for parameter gathering. We can reduce the transmitted data volume over inter-node connections by reducing the scale. Assuming we want to all-gather a message with size M among p participants, the data traffic transmitted among participants is determined by $(p - 1)M/p$ [10]. This means we can split p participants into multiple small groups and perform independent communication within each group. We consider GPUs spanning across multiple nodes as a two-dimensional grid, in which we first aggregate the data across nodes in parallel and then merge the local data on each node.

For hierarchical communication to work properly, we first build communication channels for devices. Assuming each computational node has k devices, MiCS builds k communication channels for inter-node communication and a separate communication channel for intra-node communication (Figure 4 (right)). As a comparison, vanilla collective communication uses a single communication channel for devices spanning across nodes (Figure 4 (left)). In Figures 4, we illustrate the idea using two computational nodes, each of which has two devices, i.e., $p = 4$ and $k = 2$. Next, we introduce how hierarchical communication works for inter-node communication.

MiCS uses a three-stage algorithm for hierarchical communication. In the first stage, each device uses the inter-node communication channels to do all-gather with the devices that have the same local rank on respective nodes. The inter-node all-gather

operations are executed in parallel. In the second stage, the data chunks are rearranged to ensure correctness. In the third stage, we invoke batched intra-node all-gather. In general, for the model states partitioned onto p devices spanning p/k nodes, the number of batched all-gather calls is p/k in the third stage. An example of the algorithm running across two nodes is given in Figure 5. In the following, we explain why we have the second and third stages work as we described here.

The second and third stages are designed to fix the memory discontinuous issue. Otherwise, we would get the wrong output. We use the model states partitioned to two nodes with 4 GPUs for an explanation, shown in Figure 5. The final outputs of the hierarchical communication algorithm should place data C_0 and C_1 in the adjacent locations. However, the inter-node all-gather will gather C_0 and C_2 into a contiguous memory. Thus, if we directly launch an all-gather collective primitive on the output from the first stage, we will get the wrong memory layout $[C_0, C_2, C_1, C_3]$, while the correct one is $[C_0, C_1, C_2, C_3]$. To fix this, we add a data movement stage before intra-node all-gather to rearrange the data chunks. Then in the third stage, we launch $p/k = 2$ intra-node all-gather collectives in a batch, where each intra-node all-gather operation works on a subset of the data chunks. Launching multiple communications in a batch requires new communication API implementation to get good performance, which is detailed in §4.

The performance benefits of the hierarchical communication strategy depend on the scale of the model states partitioning. Assume the model is partitioned onto p devices, and p is divisible by k , where k is the number of devices on each computational node. With the vanilla communication strategy, the inter-node data traffic is $(p - 1)M/p$. Using the proposed hierarchical communication, the data volume transmitted over inter-node connections is reduced to $(p - k)M/p$. In this way, the communication volume over the slow inter-node links is reduced by

$$\frac{p - 1}{p - k}.$$

Given that $p \geq k \geq 1$, this ratio decreases monotonically and approaches 1 when p increases. Thus, the improvement is less when we have to use more devices to hold a model replica. In a typical setup, we would have $k = 8$. A 10B-50B parameter model typically requires $8 \leq p \leq 64$ number of workers for holding the model states. In this case, we will obtain 11.1% to 46.6% data volume reduction with hierarchical communication.

3.4 2-hop Gradient Synchronization

In the typical distributed training setting, we need to aggregate gradients across all the devices [21]. Gradient aggregation is an expensive synchronization step and its cost scales with the number of workers, detailed in §2.3. It ensures that all devices work on the same model states.

To improve the training efficiency, more recent works advocate large-batch training [21, 29, 65, 66, 73]. However, due to the limited device memory, practitioners have to resort to gradient accumulation that divides a large batch into multiple micro-batches and accumulates the gradient w.r.t. each micro-batch into a shared memory buffer. In the standard data parallel setting, the gradient synchronization is only needed at the accumulation boundary where all the




Figure 5: Hierarchical communication stages; C^* denotes a data chunk; the procedure of Node 0 is shown for brevity.

gradients have been computed. However, ZeRO requires additional gradient synchronization within each micro-step because of gradient partitioning. Since each device is only responsible for holding a part of the gradient, the gradient needs to be partitioned once it is computed. In order to avoid losing the gradient information, gradients have to be aggregated before the partitioning. This makes every gradient partitioning step become a global synchronization barrier among all devices. Since MiCS only partitions the model states into a small group of devices, we can restrict the gradient synchronization within the group for each micro-step and delay the global gradient synchronization to the accumulation boundary. This motivates the design of 2-hop gradient synchronization schedule without over-paying communication costs.

2-hop gradient synchronization performs gradient synchronization within each partition group for each micro-step. Only at the gradient accumulation boundary, global synchronization is performed among the devices that possess the same part of the model. Figure 6 gives an example of a model partitioned onto two devices. Every two consecutive ranks form a partition group. Ranks with odd number and even number indices form two different replication groups, respectively. For illustration purposes, we assume the number of gradient accumulation steps is $s = 4$. For each micro-step, MiCS uses reduce-scatter to synchronize gradients within each partition group. At the gradient accumulation boundary, an all-reduce operation is used within each replication group for synchronization. An alternative synchronization schedule is to use all-reduce for gradient synchronization at every micro-step and then partition the gradient on each device. During the partitioning, each device only keeps the part of the gradient that it is responsible for while discarding the rest. This alternative schedule is the default one implemented in DeepSpeed. However, this scheme is redundant and overpays the communication costs.

The performance benefits of the 2-hop gradient synchronization schedule depend on the number of micro-steps and the effective

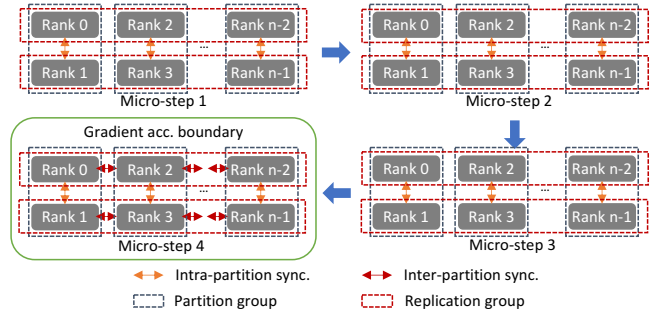


Figure 6: Gradient synchronization steps. A partition group consists of two nodes. The intra-group synchronization happens in every micro-step, while the inter-group synchronization only happens at gradient accumulation boundary.

communication bandwidths within partition groups and replication groups. For simplicity, we assume that every partition group has the same effective bandwidth B_{part} , and bandwidth within each replication group is B_{repl} . The time cost of 2-hop schedule is $C_{2-hop} = (sM(p-1))/(pB_{part}) + 2M(n-p)/(nB_{repl})$, while the time cost for the alternative schedule is $C_{alt} = 2sM(n-1)/(nB_{all})$. We take the ratio of two costs, and simplify the ratio using inequalities $(p-1)/p \leq (n-1)/n$ and $(n-p)/n \leq (n-1)/n$ when $n \geq p \geq 1$. In the following inequality, we can view the right-hand side as the lower bound for the improvement.

$$\frac{C_{alt}}{C_{2-hop}} \geq \frac{2s}{\frac{s}{B_{part}} + \frac{2}{B_{repl}}}.$$

Assuming $s = 4$, which is a reasonable setup for large batch training [45, 54], and assuming $B_{all} = B_{part} = B_{repl}$ for simplicity, we get the lower bound of the ratio at $4/3$. This means at least 25% cost reduction by using the 2-hop schedule. Taking heterogeneous bandwidth into consideration would further reduce the denominator on the right-hand side and helps achieve more gains. We notice that when $s = 1$, under the assumption that $B_{all} = B_{part} = B_{repl}$, the 2-hop synchronization is sub-optimal compared to the alternative schedule. However, given the heterogeneity of the effective bandwidths in a large cluster, e.g., having $B_{part} \approx B_{repl} > 1.5B_{all}$ (which is reasonable based on our measurement in §2.3), the 2-hop schedule typically costs less. Therefore, even for $s = 1$, in training large models with a large cluster, 2-hop is still preferred.

4 IMPLEMENTATION

The implementation of MiCS is based on DeepSpeed-v0.4.9 and PyTorch-v1.11. To efficiently implement our design, we make the following optimizations.

Fine-grained synchronization. Both parameter gathering and gradient synchronization involve a large number of communication kernel launches. Communication and computation operations are typically executed asynchronously from each other in their own CUDA streams. To maintain the data dependency correctly among these two types of operations, synchronization is required at proper

position. The synchronization mechanisms like device synchronization or stream synchronization used in DeepSpeed-v0.5.6 operate in a coarse granularity and hence lead to sub-optimal communication and computation overlapping, especially on lower bandwidth clusters. For example, if a communication operation *comm0* depends on the output from computation operation *comp0*, and the *comp0* is running with another computation *comp1* on the device, then using coarse-grained device synchronization would delay *comm0* until *comp1* is completed. Instead, MiCS follows the good practice in existing works, e.g., BytePS [26], that leverages much finer-grained `wait_event`, `wait_stream` and `record_stream` operations for synchronization, which allow us to maintain the relative order of computation and communication operations in different streams. Using this mechanism, *comm0* can kick off without waiting for *comp1* to complete. In addition, during the forward and backward passes, many complex decisions need to be made, such as which parameters should be fetched, predicting which parameters will be used next, which parameters may be reused soon and should be kept, and which can be released. We observe that making these decisions on-the-fly creates large computation and communication bubbles. We optimize this computation by precomputing and caching the decisions. The same decisions are reused throughout the training.

Coalesced communication APIs. MiCS’s hierarchical communication design introduces multiple communications over small messages. One way to improve bandwidth utilization is to batch communications. However, it is suboptimal to use existing `all_gather` and `reduce_scatter` operators in PyTorch to implement batched communication as we will have to explicitly use a custom interleaving scheme to copy the tensors into a shared buffer. MiCS introduces two coalesced communication APIs, `all_gather_coalesced` and `reduce_scatter_coalesced`. These APIs avoid the redundant buffer allocation and memory copy in `all-gather` and `reduce-scatter` API calls in PyTorch. MiCS leverages the group primitive in `nccl` to launch multiple communication operations at once, without extra data movement or allocation.

Memory defragmentation. Like DeepSpeed, MiCS also requires frequent memory allocation and deallocation operations as model states are frequently gathered and scattered. This results in serious memory fragmentation when using the dynamic allocation provided by PyTorch memory manager, causing out-of-memory errors when we try to allocate large contiguous memory buffers. DeepSpeed allocates contiguous memory buffers for holding gradients to mitigate the fragmentation issue. But it does not consider the fragmentation problems caused by operations related to partitioned parameters and gradients. MiCS’s memory management solves the memory fragmentation issue in a more comprehensive way. MiCS pre-allocates large contiguous memory buffers for holding partitioned parameters, partitioned gradients, and temporary small buffers ahead of the training. During training, MiCS reuses these buffers proactively, rather than relies on the memory management module in PyTorch.

5 EVALUATION

In this section, we evaluate the following three aspects.

- **Training performance:** Does MiCS provide better throughput than the existing solutions?
- **Effectiveness of the design:** How does each component of the system affect the performance?
- **Fidelity:** Is the system carefully implemented so that the training is converging correctly?

Setups. We conduct all experiments on AWS. Unless specified otherwise, we use Amazon EC2 p3dn.24xlarge instances for the evaluation. Each instance has 8 V100 (32GB) GPUs, which are interconnected via NVLink. The theoretical aggregated GPU interconnect bandwidth within the instance is 300 GB/s. For the inter-node communication, p3dn.24xlarge has a 100Gbps elastic fabric adaptor (EFA). In addition, we have also evaluated our system on Amazon EC2 p4d.24xlarge instances, which have 8 A100 (40GB) GPUs and a 400Gbps EFA network. The software environment includes CUDA-11.0, DeepSpeed-v0.5.6, PyTorch (customization from v1.11), Megatron-LM (git-hash d416968), and `nccl-v2.10.3`.

Metric and workloads. We use system throughput and TFLOPS as our main evaluation metrics. Unless specified otherwise, we use model variants based on the BERT model [17]. We vary the number of transformer layers and the size of each layer to get different model configurations. We also include two other popular language models, RoBERTa [31] and GPT2 [44]. Table 1 summarizes the detailed model configurations. Other than language models, we also evaluate the performance for training WideResNet to demonstrate the generality of our system. For the language models, the Wikipedia-en corpus is used as the training dataset. We fix sequence length to 512 for the training. For the WideResNet model, we use synthetic data with images sized $3 \times 224 \times 224$. By default, we use a micro-batch size of 8, global-batch size of 8192, mixed-precision, and activation checkpointing in training.

5.1 Performance

In this section, we demonstrate the performance advantages of MiCS. First, we show the scalability of MiCS against DeepSpeed, which is the state-of-the-art (SOTA) solution using DP with model states partitioning. The TFLOPS numbers are also reported to show the computation utilization of each GPU. In §5.1.2, we evaluate MiCS and DeepSpeed in a different network condition, i.e., 400Gbps network. In §5.1.5 we provide performance numbers of the 100B model training on a large scale. In §5.1.3, we show MiCS can outperform Megatron-LM-3D [37], which is a SOTA design for transformer-based language models that uses DP and MP.

5.1.1 Scalability in 100Gbps Network. In this subsection, we report the throughput performance and strong-scaling efficiency of MiCS and DeepSpeed in 100Gbps networks. The baselines include both ZeRO-2 and ZeRO-3 in DeepSpeed. ZeRO-1 is excluded because it is not runnable for the smallest model we consider. ZeRO-Offload [49] and ZeRO-Infinity [46] are not included, either. These two variants aim to utilize CPU memory and NVMe storage to support large models, which are orthogonal to MiCS. Instead, we focus on minimizing communication overhead to improve the training throughput. For both ZeRO-3 and MiCS, we use micro-batch size 8. But for ZeRO-2 we use a smaller micro-batch size 4, because ZeRO-2 does not perform parameter partitioning and uses more GPU memory for

Table 1: Structure of language models. BERT 10B means BERT with 10 Billion parameters, and similarly for other model names. We use a sequence length of 512 for all the models.

Model	Hidden size	Intermediate size	#layers	#Attention heads	Vocabulary size
BERT 10B	2560	10240	127	40	32008
BERT 15B	2560	10240	190	40	32008
BERT 20B	5120	20480	64	40	32008
BERT 50B	8192	32768	62	40	32008
RoBERTa 20B	5120	20480	62	40	50265
GPT2 20B	5120	20480	62	40	50265




Figure 7: Strong-scaling with different model sizes; × denotes “out-of-memory”; black rectangular denotes the linear-scaling efficiency.

the redundant model parameter replicas. We vary the number of computational nodes from 2 (resp. 16 GPUs) to 16 (resp. 128 GPUs). For the partition group size, we use the smallest number of nodes that allow us to train models with the selected batch size, i.e., 1 node for BERT 10B, 2 nodes for BERT 15B and 20B, 8 nodes for BERT 50B. All throughput numbers are averaged over 500 iterations.

As shown in Figure 7 and 8, the throughput of MiCS is significantly better than that of DeepSpeed. Our performance numbers show that the throughput of MiCS is up to 2.82× that of DeepSpeed for the BERT 15B model. MiCS achieves near-linear or super-linear scalability in all experiments. Here we define the linear-scaling as with respect to the smallest number of computational nodes that can hold the model states with the targeted micro-batch size, e.g., for BERT 50B the linear-scaling is with respect to 8 nodes. In most of the setups, ZeRO-2 has an out-of-memory (OOM) problem. Next, we explain the rationale of our results.

The performance improvements are different with respect to the different characteristics of the models. For the BERT 10B model, a single computational node has enough GPU memory to hold the model states so that we can leverage fast intra-node GPU interconnect to complete most of the communication. In this case, MiCS is 223% faster than ZeRO-3. And, larger micro-batch size allows MiCS to further achieve more gains over ZeRO-2. The performance gain of MiCS for BERT 15B is larger than that for BERT 20B model. The

difference is mainly due to the structural differences between the two models. As listed in Table 1, BERT 15B has narrower transformers layer but a larger number of layers. The narrower model leads to smaller computation and communication units, which allow finer grained overlapping of computation and communication. In BERT 20B experiments, we observe super-linear scaling. This is because we have to disable hierarchical communication on 16 GPUs due to the memory constraint. The all-reduce overhead among replication groups is amortized by multiple micro-batches (§3.4). The amortized overhead is relatively small, less than 1%, to the iteration time of each micro-step. Thus, MiCS can maintain near-linear scalability.

To compare the computation utilization, we calculate the TFLOPS performance based on system throughput. The TFLOPS numbers are shown in Figure 9. We follow the equation in [37] to calculate the total TFLOPS.

$$F = 96TlLh^2 \left(1 + \frac{l}{6h} + \frac{V}{16Lh} \right),$$

where V denotes vocabulary size, l is the sequence length, h is the hidden size, L refers to the number of layers, and T is throughput per second². As we can see, MiCS is better than ZeRO-3 by a large margin for all the model sizes. The maximum gain we observe is 223.7%. For the BERT 10B model, we achieve about 42% of the theoretical peak performance of V100. When the model size is over 10B, the performance dropping is mainly because of the cross node partitioning, which causes a larger communication overhead. However, the computation utilization we get is still on par with the numbers reported by DeepSpeed ZeRO [45] and Megatron-LM [54] on DGX-2 clusters, which have 800Gbps networking.

5.1.2 Scalability in 400Gbps Network. In this subsection, we evaluate MiCS on a GPU cluster with A100 GPUs and 400Gbps network (Amazon EC2 p4d.24xlarge instances, 8 GPUs per instance). We use the BERT 15B and BERT 20B models for the evaluation, and we fix the micro-batch size to 8 for all experiments. DeepSpeed ZeRO-3 is used as our baseline for comparison.

As shown in Figure 10, MiCS significantly outperforms DeepSpeed and achieves near-linear scaling. The throughput of MiCS is up to 2.21× that of ZeRO-3. The throughput gap enlarges as the scale of the cluster increases, demonstrating that MiCS can maintain near-linear scaling efficiency. In BERT 15B case, when we scale the cluster size from 16 GPUs to 64 GPUs, MiCS achieves 96.7% efficiencies with respect to 16 GPUs. In contrast, DeepSpeed ZeRO-3 only achieves 85.3% for BERT 15B. Compared to the results

²The derivation process of the formula is in the appendix of the paper [37].




Figure 8: Strong-scaling with other language models; × denotes “out-of-memory”; black rectangular denotes the linear scaling efficiency.




Figure 9: TFLOPS performance; BERT models with different sizes; × denotes “out-of-memory”.




Figure 10: Throughput comparison in 400Gbps network.

in Figure 7b, the performance gains are lower mainly because faster network bandwidth mitigates communication overheads.

5.1.3 Comparison to Megatron-LM-3D. We increase the number of layers to 128 while keeping the same hidden size and intermediate size as the BERT 10B model. This is because the pipeline parallelism of Megatron-LM-3D requires the number of layers to be divisible by the size of pipeline parallelism. We use micro-batch size 8 and global-batch size 4096 for this experiment. We follow the takeaways from Megatron-LM-3D [37] to tune the tensor parallel size and pipeline parallel size for better performance. Specifically, we avoid using tensor MP across nodes and use more pipeline MP than DP size if applicable. We report three reasonable setups of Megatron-LM-3D, as listed in table 2. In the table, we omit the DP size, because it depends on the size of the training cluster.

As shown in Figure 11a, the performance of Megatron-LM-3D is sensitive to model parallel configurations. We always restrict the tensor MP size to be lower than eight to make sure the tensor MP ranks only communicate through NVLink. But Megatron-LM-3D is still sensitive to the tuning of the MP sizes, e.g., configuration (3) is 38% better than configuration (1). This raises usability challenges to users. In contrast, MiCS does not have such complicate




Figure 11: Performance Comparison to Megatron-LM-3D.

configurations for different parallel sizes, because of the simplicity of data parallelism. And MiCS is up to 31% faster than the best results from Megatron-LM-3D. Our profiling shows the inefficiency of Megatron-LM-3D is mainly due to timeline bubbles in pipeline parallelism and communication overhead in tensor parallelism.

For some uncommonly structured models, Megatron-LM-3D could outperform MiCS marginally. We conducted some experiments to evaluate system performance with respect to the structural differences of models. The number of parameters of the model is fixed to 10B. Figure 11b presents the throughput of Megatron-LM-3D and MiCS, that are evaluated on a BERT model with wider transformer layers than a regular BERT 10B model. Specifically, the model consists of 80 transformer layers. The intermediate size of each transformer layer is equal to $8 \times$ hidden size. This kind of wider structure is used in the evaluation of GSPMD [63]. Usually, the intermediate size of a transformer layer is $4 \times$ that of the hidden size [37, 45, 54, 61]. The number of transformer layers, 80, is chosen to match the size of the regular BERT 10B model. The other training setups are the same as the previous experiment. In this experiment, Megatron-LM-3D with configuration (3) is slightly better than MiCS. The performance gaps are within 1.5%. For this specific setup, the wider structure produces larger intermediate activations and requires more memory for each transformer layer. Frequently allocating and releasing large memory chunks cause allocation failure and retry at the PyTorch allocator side [59], which impacts the efficiency of overlapping computation and communication in MiCS.

5.1.4 Performance on CV models. To show the performance improvements of MiCS generalize to other models, we report the training throughput of WideResNet [67], a computer vision model,

Table 2: Megatron-LM-3D configurations.

Configuration	Tensor MP size	Pipeline MP size
Megatron-LM-3D (1)	8	1
Megatron-LM-3D (2)	4	4
Megatron-LM-3D (3)	2	8

in Figure 11c. We compare MiCS against DeepSpeed (ZeRO-3). Note that Megatron-LM-3D cannot be applied to training this model. We scale up the size of WideResNet by enlarging the width and number of blocks of the network. In our setup, the WideResNet model has 3B parameters. It has 200 convolution layers, width factor 8, and its bottleneck block configuration is [6, 8, 46, 6]. We fix batch size 8 for each GPU, and use synthetic image data with size 224x224 for benchmarking. The training uses float32 and activation checkpointing is disabled. The model is not runnable under ZeRO-2 optimization. The system throughput of MiCS is up to 2.89x that of DeepSpeed (ZeRO-3).

5.1.5 Case Study: 52B and 100B Model Training. MiCS has been deployed to train proprietary models in distribution. Our training cluster consists of 128 A100 GPUs with 400Gbps networking. Our results show that we can achieve 179 and 171 TFLOPS per GPU for 52B and 100B parameter models, respectively. These are about 57% and 55% compute utilization of the peak half-precision performance of A100. The utilization results outperform the TFLOPS performance reported from Megatron-LM-3D [37] on DGX A100 clusters with 8 InfiniBand networking cards (1.6Tb/s) [40]. When we increase the number of GPUs from 128 to 512, we can obtain 170 TFLOPS per GPU for the 100B parameter model with 99.4% weak scaling efficiency, where the partition group size is 128 GPUs. When the cluster size equals the partition group size (128 GPUs), the performance improvements come from hierarchical communication and implementation optimizations. In contrast, DeepSpeed ZeRO-3 only achieves 62 TFLOPS per GPU for training a 100B parameter model on 512 GPUs with 72% weak-scaling efficiency. In this experiment, the size of each micro-batch is 16 and the number of micro-steps is 4. The TFLOPS performance of MiCS is 2.74x that of DeepSpeed ZeRO-3 on 512 GPUs.

5.2 Analysis of the Design

To understand the performance contribution of each component in MiCS, we conduct ablation tests in this section. We divide our studies into three subsections. Each subsection corresponds to one of the three components in §3. For each experiment, we present the setups followed by detailed results and takeaways.

5.2.1 Analysis of Partition Group Size. As the scale-aware model partitioning uses partition groups for storing model states replicas, it is natural to ask the relationship between the size of the partition group and the end-to-end performance. In this experiment, we use BERT 10B model, fix the micro-batch size to 8, and use 64 V100 GPUs in total. We vary the size of each group from 8 GPUs to 64 GPUs. If we use all the 64 GPUs for partitioning the model states, MiCS reduces to ZeRO-3. As shown in Figure 12, by increasing the partition group size, the end-to-end throughput trends down




Figure 12: Throughput change w.r.t. partition group sizes.




Figure 13: Benefits of hierarchical all-gather.

obviously. The throughput of partition group size 8 is 1.6x that of partition group size 64. Thus, it is preferable to partition the model states into a smallest possible group.

5.2.2 Analysis of Hierarchical Communication. Hierarchical communication plays an important role for good performance because it can reduce the transmitted data volume over the inter-node connections. In this subsection, we conduct performance analysis quantitatively to show its importance. We divide our experiments into two parts, micro-benchmark and end-to-end training throughput. In both experiments, we report normalized performance to baselines, i.e., vanilla all-gather and DeepSpeed ZeRO-3.

In the micro-benchmark experiment, we measure the elapsed time of vanilla all-gather and hierarchical all-gather operators for handling different message sizes. We use two Amazon EC2 p3dn.24xlarge instances. We cap the message size at 256MB, because a single parameter fetching typically gathers less data than it for better overlapping of computation and communication. In Figure 13a, we can see that the elapsed time of hierarchical communication operator is consistently lower than the baseline. For message size 128MB, hierarchical communication only uses about 72.1% of the time cost of vanilla all-gather.

For the end-to-end experiment, we use the BERT 15B model, which needs two computational nodes (i.e., 16 GPUs) to hold the model states for the training. For models that can be held by a single computational node (i.e., 8 GPUs), the hierarchical all-gather is not needed. We vary the cluster size from 16 to 128 GPUs and evaluate MiCS with and without hierarchical communication. We normalize throughput numbers to the results of DeepSpeed ZeRO-3. As shown in Figure 13b, MiCS with hierarchical communication is consistently better than the case where hierarchical communication is disabled. In particular, hierarchical communication improves the end-to-end training throughput by 30.6% to 38%.




Figure 14: Benefits of 2-hop gradient synchronization.




Figure 15: Improvements of implementation optimizations.

5.2.3 *Analysis of Synchronization Scheduling.* In this experiment, we report the throughputs of MiCS with 2-hop gradient synchronization enabled and disabled. We use the BERT 10B model for the experiments and fix the micro-batch size 8, global batch size 8192 for training. We partition model states on 8 GPUs. When the 2-hop synchronization is disabled, the system uses an alternative synchronization schedule that synchronizes the gradients across all devices at the end of each micro-step, explained in § 3.4. We can see the performance gaps between these two setups, Figure 14. When the cluster size increases to 128 GPUs, we get the max throughput gap. Numerical results indicate that the relative improvement ranges from 11% to 24.9%, when 2-hop synchronization is enabled.

5.3 Other Optimizations

We conduct experiments to analyze the performance improvements by using the optimization techniques described in § 4. We use the BERT 10B model for the evaluation. In the training, we use the default setup as mentioned at the beginning of § 5. When we turn off optimizations that are unique to MiCS and let the model states be partitioned over all devices, MiCS reduces to ZeRO-3 with the optimization techniques in § 4. We denote it as “MiCS (ZeRO-3)”. For comparisons, we report the throughput of DeepSpeed ZeRO-3.

Figure 15 shows the improvements of using the proposed system optimizations. MiCS (ZeRO-3) achieves 54.1% better system throughput than DeepSpeed ZeRO-3 when the cluster scales up to 128 GPUs, while the scaling efficiency of DeepSpeed ZeRO-3 drops when we scale out the cluster. In addition, MiCS still significantly outperforms MiCS (ZeRO-3), demonstrating the superiority of minimizing the communication scale.

5.4 Fidelity

In this section, we show that MiCS achieves consistent convergence as DeepSpeed, which validates the correctness of our system. We provide the training loss curves for training a 1.5B parameter model on the Wikipedia-en dataset. The model has 48 transformer layers, each of which is constructed with the hidden size 1,600 and intermediate size 6,400. The global batch size is 512. And the micro-batch




Figure 16: Fidelity of the implementation.

size is 8 (the number of gradient accumulation steps is 4). The loss validation process does not aim to produce exactly the same loss as DeepSpeed but to ensure the convergence behaviours are the same. We report the training losses on 1 million sequences. As shown in Figure 16, MiCS provides the same convergence as DeepSpeed.

6 RELATED WORK

Data parallelism. PyTorch-DDP [43], Horovod [52], ps-lite [42], Tensorflow-DDP [57], and BytePS [26] are distributed training frameworks using data parallelism. All of them place complete model states on each GPU for training. Thus, the supported model size is limited. Recently, ZeRO [45] has been proposed to address the memory limitation issue of the traditional data-parallel strategy, by partitioning the model states onto all GPUs. ZeRO-Offload [49] and ZeRO-Infinity [46] are two extensions to ZeRO that explore the possibility to extend the memory to hold the model from GPU memory to CPU memory and NVMe, which are orthogonal to MiCS. MiCS focuses on minimizing the communication overheads of the training system, addressing the challenges not solved in ZeRO.

Other parallelisms. ColocRL [34] formulates the distributed training as a placement optimization problem to maximize the throughput. FlexFlow [25] and OptCNN [24] use heuristic search for parallel strategies including tensor MP and device placement MP. Alpa [72] and Unity [60] use hierarchical search to jointly optimize within- and between-device to look for a good parallel strategy. These systems do not explicitly embed the memory constraints into their optimization objective, and are not directly verified to train models at the scale that MiCS trained. Megatron-LM-3D [37], GPipe [22], and DAPPLE [19] use pipeline parallelism to partition large models into multiple stages for the training in a synchronous manner. These solutions have resource under-utilization problems due to pipeline bubbles. PipeMare [64], PipeDream [35], and PipeDream-2BW [36] use asynchronous and bounded-staleness training for efficient resource utilization which, however, can affect the convergence quality [14, 16, 48]. The research direction of asynchronous methods are orthogonal to MiCS. Currently, MiCS uses synchronous training and it does not suffer from convergence issues. DLRM [38] and Megatron-LM [54] are specific designs for recommendation models and transformers, respectively. DLRM partitions the embedding table along row and column dimensions. Megatron-LM introduces tensor parallelism to parallelize the tensor computation on multiple devices. Megatron-LM-3D [37] integrates the pipeline parallelism into Megatron-LM for further scaling up the model size. Pipeline parallelism, tensor parallelism, and the mixture of multiple parallelisms require significantly additional efforts to program the

customized model implementation and tune the hyper-parameters for high performance. MiCS is orthogonal to this line of research. We compared MiCS against Megatron-LM-3D [37] in Section 5.1.3.

Communication optimizations. ByteScheduler [41] and P3 [23] overlap the computation with communication to hide the communication cost. SwitchML [50] and ATP [27] use the programmable switches as gradient aggregation servers to reduce the communication overheads. Lossy compression algorithms like 1bit-SGD [51] and DGC [30] compress the data transmitted over the network to improve the system performance. Those techniques are complementary to our system for further reducing the communication overheads. Blink [62] leverages multiple communication channels with optimized spanning trees to speed up the gradient synchronization. Plink [33] discovers and explores the locality of the distributed training cluster for better performance. Cloud Collective [32] reorders the ranks of cluster nodes to explore a better topology. Blueconnect [12] decomposes all-reduce primitive with pipelined reduce-scatter and all-gather. These techniques explore better locality or pipeline multiple communication primitives to speed up the synchronization. MiCS reduces communication overheads from a different perspective. In particular, our system reduces communication costs by reducing the scale of communications. Varuna [2] works on optimizing network jitter and instability among cheap “spot” instances [6] to lower the training cost. The objective of Varuna is orthogonal to MiCS. In high-performance computing, innovations [7, 8, 55] at the hardware level are critical to the efficiency of communications. On the other hand, researchers explore the relationship among collective communication algorithms, software implementations, and message sizes to optimize each individual communication primitives [1, 58]. These efforts are orthogonal to MiCS. Our system is built with GPU-aware library NCCL [39].

7 DISCUSSION AND FUTURE WORK

The optimality of the training throughput depends on the model structure, input data, and hardware. For the models used in the evaluation, we do not prove that MiCS is the optimal solution. MiCS is a pure data-parallel training system, which admittedly covers a limited space of parallelism strategies. Thus, for some less common model structures, e.g., wider feedforward layer in transformer blocks, Megatron-LM-3D could outperform MiCS in certain configurations marginally, shown in §5.1.3. It is worth noting that adapting tensor model parallelism and pipeline parallelism requires refactoring model implementations [45], thus is less favorable to practitioners. MiCS, as a pure DP solution, achieves state-of-the-art performance in training standard transformer-based models with billions of parameters and trades off performance for lower complexity in some less common cases.

Despite the model states replications created in MiCS, our system does not require additional hardware resources as compared to the existing ZeRO system. Partitioning one model states replication across all devices, as the existing ZeRO system does, underutilizes the memory of each device. As discussed in the first paragraph of §3.2, the memory capacity of eight V100 (32GB) GPUs are large enough for holding model states of a model with 10 billion (B) parameters. For a cluster with 16 or more V100 (32GB) GPUs, partitioning the 10B model to all devices consumes less than 32% memory

usage of each device for holding the model states. MiCS effectively leverages these spare memory resources for lowering communication costs (§3). For models that require all devices to hold the model states for training, MiCS still outperforms the ZeRO system because of the hierarchical communication module (§3.3).

In MiCS, the memory consumption of each device is controlled by the size of the partition group, which is configurable. MiCS uses a heuristic to pick the size for holding the model states, which is mentioned in §5.1.1. Compared to prior large model training systems, MiCS does not introduce extra issues in terms of system practicability. For ZeRO systems, users have to figure out the smallest size of the cluster for training, otherwise the system runs into out-of-memory issues. Similarly, the Megatron-LM-3D system requires users to configure the number of pipeline stages and the tensor parallelism size, so that the partitioned model components can fit into each GPU in a cluster. In MiCS, the way to figure out the partition group size is the same as figuring out the smallest size of the cluster for training in ZeRO systems.

To automate the configuration search for large model training, an accurate estimation of memory usage is needed. A profiling-based method can get relatively precise memory usage statistics. But once the training processing runs into the out-of-memory issue during configuration search, the dangling process can cause hanging and prevent successive configurations from launching. Estimating memory consumption from the model structure and input size is inaccurate due to the dynamic behavior of the memory management module in PyTorch runtime. Addressing challenges from estimating or predicting the memory usage of large models is beyond the scope of this paper, in which we focus on reducing the communication overheads in the ZeRO DP algorithms. We leave the configuration search for MiCS as future work.

8 CONCLUSION

In this paper, we present MiCS, a system that attains high training throughput and near-linear scalability on the cloud by only using data parallelism. The overarching goal of MiCS is to minimize the communication scale so as to reduce the expensive communication overhead rooted in parameter gathering and gradient synchronization. Specifically, we propose scale-aware model partitioning, hierarchical communication strategy, and 2-hop gradient synchronization to achieve this goal. We evaluate MiCS on various training workloads on large-scale clusters. MiCS outperforms DeepSpeed ZeRO by up to 2.89× and demonstrates near-linear scaling efficiency in various training setups.

ACKNOWLEDGMENTS

We sincerely thank the anonymous reviewers for their valuable feedback. We thank the Amazon Search M5 team for providing large clusters. Xin Jin and Shuai Zheng are the corresponding authors. Xin Jin is with the Key Laboratory of High Confidence Software Technologies (Peking University), Ministry of Education. Zhen Zhang is supported in part by NSF grants CNS-1813487 and CCF-1918757. Xin Jin is supported in part by National Natural Science Foundation of China under the grant number 62172008 and National Natural Science Fund for the Excellent Young Scientists Fund Program (Overseas).

REFERENCES

- [1] George Almási, Philip Heidelberger, Charles J Archer, Xavier Martorell, C Chris Erway, José E Moreira, Burkhard Steinmacher-Burow, and Yili Zheng. 2005. Optimization of MPI collective communication on BlueGene/L systems. In *Proceedings of the 19th annual international conference on Supercomputing*.
- [2] Sanjith Athlur, Nitika Saran, Muthian Sivathanu, Ramachandran Ramjee, and Nipun Kwatra. 2021. Varuna: Scalable, Low-cost Training of Massive Deep Learning Models. *arXiv preprint arXiv:2111.04007* (2021).
- [3] AWS-P3-Instances 2022. Amazon EC2 P3 Instances. <https://aws.amazon.com/ec2/instance-types/p3/>.
- [4] azure-gpu-ncv3-series 2022. Azure NCv3-series. <https://docs.microsoft.com/en-us/azure/virtual-machines/ncv3-series>.
- [5] azure-gpu-ndv2-series 2022. Azure Updated NDv2-series. <https://docs.microsoft.com/en-us/azure/virtual-machines/ndv2-series>.
- [6] Azure-Spot-VM 2022. Azure Spot Virtual Machines. <https://azure.microsoft.com/en-us/services/virtual-machines/spot/#overview>.
- [7] Kevin J Barker, Alan Benner, Ray Hoare, Adolfy Hoisie, Alex K Jones, Darren K Kerbyson, Dan Li, Rami Melhem, Ramakrishnan Rajamony, Eugen Schenfeld, et al. 2005. On the feasibility of optical circuit switching for high performance computing systems. In *SC'05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*. 16–16.
- [8] Aleksandr Biberman and Keren Bergman. 2012. Optical interconnection networks for high-performance computing systems. *Reports on Progress in Physics* (2012).
- [9] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [10] Ernie Chan, Marcel Heimlich, Avi Purkayastha, and Robert van de Geijn. 2007. Collective Communication: Theory, Practice, and Experience. *Concurrency and Computation: Practice and Experience* 19 (2007), 1749–1783.
- [11] William Chan, Daniel Park, Chris Lee, Yu Zhang, Quoc Le, and Mohammad Norouzi. 2021. SpeechStew: Simply mix all available speech recognition data to train one large neural network. *arXiv preprint arXiv:2104.02133* (2021).
- [12] Minsik Cho, Ulrich Finkler, and David Kung. 2019. BlueConnect: Novel hierarchical all-reduce on multi-tired network for deep learning. In *Conference on Machine Learning and Systems*.
- [13] Yu-An Chung, Yu Zhang, Wei Han, Chung-Cheng Chiu, James Qin, Ruoming Pang, and Yonghui Wu. 2021. W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training. *arXiv preprint arXiv:2108.06209* (2021).
- [14] Cody Coleman, Daniel Kang, Deepak Narayanan, Luigi Nardi, Tian Zhao, Jian Zhang, Peter Bailis, Kunle Olukotun, Chris Ré, and Matei Zaharia. 2019. Analysis of DAWN-Bench, a Time-to-Accuracy Machine Learning Performance Benchmark. *ACM SIGOPS Operating Systems Review* 53, 1 (2019), 14–25.
- [15] Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. 2021. CoAtNet: Marrying Convolution and Attention for All Data Sizes. *arXiv preprint arXiv:2106.04803* (2021).
- [16] Christopher M De Sa, Ce Zhang, Kunle Olukotun, and Christopher Ré. 2015. Taming the wild: A unified analysis of hogwild-style algorithms. *Advances in Neural Information Processing Systems* 28 (2015).
- [17] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [18] FairScale 2022. PyTorch extensions for high performance and large scale training. <https://github.com/facebookresearch/fairscale>.
- [19] Shiqing Fan, Yi Rong, Chen Meng, Zongyan Cao, Siyu Wang, Zhen Zheng, Chuan Wu, Guoping Long, Jun Yang, Lixue Xia, et al. 2021. DAPPLE: A pipelined data parallel approach for training large models. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*.
- [20] gcloud-gpu-bandwidths 2022. Google Cloud: Network bandwidths and GPUs. <https://cloud.google.com/compute/docs/gpus/gpu-network-bandwidth#vm-configurations>.
- [21] 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).
- [22] Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Chen, HyoukJoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, et al. 2019. Gpipe: Efficient training of giant neural networks using pipeline parallelism. In *Advances in Neural Information Processing Systems*, Vol. 32.
- [23] Anand Jayarajan, Jinliang Wei, Garth Gibson, Alexandra Fedorova, and Gennady Pekhimenko. 2019. Priority-based parameter propagation for distributed DNN training. *arXiv preprint arXiv:1905.03960* (2019).
- [24] Zhihao Jia, Sina Lin, Charles R Qi, and Alex Aiken. 2018. Exploring hidden dimensions in parallelizing convolutional neural networks. In *International Conference on Machine Learning (ICML)*. 2279–2288.
- [25] Zhihao Jia, Matei Zaharia, and Alex Aiken. 2018. Beyond data and model parallelism for deep neural networks. In *Conference on Machine Learning and Systems*, Vol. 1. 1–13.
- [26] Yimin Jiang, Yibo Zhu, Chang Lan, Bairen Yi, Yong Cui, and Chuanxiong Guo. 2020. A Unified Architecture for Accelerating Distributed {DNN} Training in Heterogeneous GPU/CPU Clusters. In *USENIX OSDI*. 463–479.
- [27] ChonLam Lao, Yanfang Le, Kshiteej Mahajan, Yixi Chen, Wenfei Wu, Aditya Akella, and Michael Swift. 2021. ATP: In-network Aggregation for Multi-tenant Learning. In *USENIX NSDI*. 741–761.
- [28] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668* (2020).
- [29] Conglong Li, Ammar Ahmad Awan, Hanlin Tang, Samyam Rajbhandari, and Yuxiong He. 2021. 1-bit LAMB: Communication Efficient Large-Scale Large-Batch Training with LAMB’s Convergence Speed. *arXiv preprint arXiv:2104.06069* (2021).
- [30] 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).
- [31] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692* (2019).
- [32] Liang Luo, Jacob Nelson, Arvind Krishnamurthy, and Luis Ceze. 2021. Cloud Collectives: Towards Cloud-aware Collectives for ML Workloads with Rank Reordering. *arXiv preprint arXiv:2105.14088* (2021).
- [33] Liang Luo, Peter West, Arvind Krishnamurthy, Luis Ceze, and Jacob Nelson. 2020. PLink: Discovering and Exploiting Datacenter Network Locality for Efficient Cloud-based Distributed Training. In *Conference on Machine Learning and Systems*, Vol. 2. 82–97.
- [34] Azalia Mirhoseini, Hieu Pham, Quoc V Le, Benoit Steiner, Rasmus Larsen, Yuefeng Zhou, Naveen Kumar, Mohammad Norouzi, Samy Bengio, and Jeff Dean. 2017. Device Placement Optimization with Reinforcement Learning. In *International Conference on Machine Learning*, Vol. 70. 2430–2439.
- [35] Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R Devanur, Gregory R Ganger, Phillip B Gibbons, and Matei Zaharia. 2019. PipeDream: generalized pipeline parallelism for DNN training. In *ACM SOSP*. 1–15.
- [36] Deepak Narayanan, Amar Phanishayee, Kaiyu Shi, Xie Chen, and Matei Zaharia. 2020. Memory-Efficient Pipeline-Parallel DNN Training. *arXiv preprint arXiv:2006.09503* (2020).
- [37] Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Anand Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, Amar Phanishayee, and Matei Zaharia. 2021. Efficient Large-Scale Language Model Training on GPU Clusters Using Megatron-LM. *arXiv preprint arXiv:2104.04473* (2021).
- [38] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G. Azzolini, Dmytro Dzhulgakov, Andrey Malleevich, Ilya Cherniavskii, Yinghai Lu, Raghuraman Krishnamoorthi, Ansha Yu, Volodymyr Kondratenko, Stephanie Pereira, Xianjie Chen, Wenlin Chen, Vijay Rao, Bill Jia, Liang Xiong, and Misha Smelyanskiy. 2019. Deep Learning Recommendation Model for Personalization and Recommendation Systems. *arXiv preprint arXiv:1906.00091* (2019).
- [39] NCCL 2022. NVIDIA Collective Communications Library (NCCL). <https://developer.nvidia.com/nccl>.
- [40] NVIDIA-DGX-A100 2022. NVIDIA DGX A100. <https://images.nvidia.com/aem-dam/Solutions/Data-Center/nvidia-dgx-a100-datasheet.pdf>.
- [41] Yanghua Peng, Yibo Zhu, Yangrui Chen, Yixin Bao, Bairen Yi, Chang Lan, Chuan Wu, and Chuanxiong Guo. 2019. A generic communication scheduler for distributed DNN training acceleration. In *ACM SOSP*. 16–29.
- [42] PS-lite 2022. lightweight implementation of the parameter server framework. <https://github.com/dmlc/ps-lite>.
- [43] PyTorch 2022. PyTorch. <https://pytorch.org/>.
- [44] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog* (2019).
- [45] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2019. Zero: Memory optimization towards training a trillion parameter models. *arXiv preprint arXiv:1910.02054* (2019).
- [46] Samyam Rajbhandari, Olatunji Ruwase, Jeff Rasley, Shaden Smith, and Yuxiong He. 2021. ZeRO-Infinity: Breaking the GPU Memory Wall for Extreme Scale Deep Learning. *arXiv preprint arXiv:2104.07857* (2021).
- [47] Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. DeepSpeed: System optimizations enable training deep learning models with over 100 billion parameters. In *ACM SIGKDD*. 3505–3506.
- [48] Benjamin Recht, Christopher Re, Stephen Wright, and Feng Niu. 2011. Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. *Advances in Neural Information Processing Systems* 24 (2011).

- [49] Jie Ren, Samyam Rajbhandari, Reza Yazdani Aminabadi, Olatunji Ruwase, Shuangyan Yang, Minjia Zhang, Dong Li, and Yuxiong He. 2021. Zero-offload: Democratizing billion-scale model training. *arXiv preprint arXiv:2101.06840* (2021).
- [50] Amedeo Sapio, Marco Canini, Chen-Yu Ho, Jacob Nelson, Panos Kalnis, Changhoon Kim, Arvind Krishnamurthy, Masoud Moshref, Dan RK Ports, and Peter Richtárik. 2019. Scaling distributed machine learning with in-network aggregation. *arXiv preprint arXiv:1903.06701* (2019).
- [51] 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 *INTERSPEECH*. 1058–1062.
- [52] Alexander Sergeev and Mike Del Balso. 2018. Horovod: fast and easy distributed deep learning in TensorFlow. *arXiv preprint arXiv:1802.05799* (2018).
- [53] Noam Shazeer, Youlong Cheng, Niki Parmar, Dustin Tran, Ashish Vaswani, Penporn Koanantakool, Peter Hawkins, HyoukJoong Lee, Mingsheng Hong, Cliff Young, et al. 2018. Mesh-tensorflow: Deep learning for supercomputers. *arXiv preprint arXiv:1811.02084* (2018).
- [54] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2020. Megatron-LM: Training Multi-Billion Parameter Language Models Using Model Parallelism. *arXiv preprint arXiv:1909.08053* (2020).
- [55] Ankit Singla, P Brighten Godfrey, and Alexandra Kolla. 2014. High throughput data center topology design. In *USENIX NSDI*. 29–41.
- [56] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990* (2022).
- [57] TensorFlow 2022. TensorFlow. <https://www.tensorflow.org/>.
- [58] Rajeev Thakur, Rolf Rabenseifner, and William Gropp. 2005. Optimization of collective communication operations in MPICH. *The International Journal of High Performance Computing Applications* 19, 1 (2005), 49–66.
- [59] Torch CUDA Memory Stats 2022. Torch CUDA Memory Stats. https://pytorch.org/docs/stable/generated/torch.cuda.memory_stats.html.
- [60] Colin Unger, Zhihao Jia, Wei Wu, Sina Lin, Mandeep Baines, Carlos Efrain Quintero Narvaez, Vinay Ramakrishnaiah, Nirmal Prajapati, Pat McCormick, Jamaludin Mohd-Yusof, Xi Luo, Dheevatsa Mudigere, Jongsoo Park, Misha Smelyanskiy, and Alex Aiken. 2022. Unity: Accelerating DNN Training Through Joint Optimization of Algebraic Transformations and Parallelization. In *USENIX OSDI*. Carlsbad, CA, 267–284.
- [61] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention Is All You Need. *arXiv preprint arXiv:1706.03762* (2017).
- [62] Guanhua Wang, Shivaram Venkataraman, Amar Phanishayee, Jorgen Thelin, Nikhil Devanur, and Ion Stoica. 2020. Blink: Fast and generic collectives for distributed ML. In *Conference on Machine Learning and Systems*, Vol. 2. 172–186.
- [63] Yuanzhong Xu, HyoukJoong Lee, Dehao Chen, Blake Hechtman, Yanping Huang, Rahul Joshi, Maxim Krikun, Dmitry Lepikhin, Andy Ly, Marcello Maggioni, et al. 2021. GSPMD: general and scalable parallelization for ML computation graphs. *arXiv preprint arXiv:2105.04663* (2021).
- [64] Bowen Yang, Jian Zhang, Jonathan Li, Christopher Ré, Christopher R. Aberger, and Christopher De Sa. 2020. PipeMare: Asynchronous Pipeline Parallel DNN Training. *arXiv preprint arXiv:1910.05124* (2020).
- [65] Yang You, Jing Li, Sashank Reddi, Jonathan Hseu, Sanjiv Kumar, Srinadh Bhojanapalli, Xiaodan Song, James Demmel, Kurt Keutzer, and Cho-Jui Hsieh. 2019. Large batch optimization for deep learning: Training bert in 76 minutes. *arXiv preprint arXiv:1904.00962* (2019).
- [66] Yang You, Zhao Zhang, Cho-Jui Hsieh, James Demmel, and Kurt Keutzer. 2018. Imagenet training in minutes. *arXiv preprint arXiv:1709.05011* (2018).
- [67] Sergey Zagoruyko and Nikos Komodakis. 2017. Wide Residual Networks. *arXiv preprint arXiv:1605.07146* (2017).
- [68] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. 2021. Scaling vision transformers. *arXiv preprint arXiv:2106.04560* (2021).
- [69] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068* (2022).
- [70] Yu Zhang, James Qin, Daniel S Park, Wei Han, Chung-Cheng Chiu, Ruoming Pang, Quoc V Le, and Yonghui Wu. 2020. Pushing the limits of semi-supervised learning for automatic speech recognition. *arXiv preprint arXiv:2010.10504* (2020).
- [71] Zhen Zhang, Chaokun Chang, Haibin Lin, Yida Wang, Raman Arora, and Xin Jin. 2020. Is network the bottleneck of distributed training?. In *Proceedings of the Workshop on Network Meets AI & ML*. 8–13.
- [72] Lianmin Zheng, Zhuohan Li, Hao Zhang, Yonghao Zhuang, Zhifeng Chen, Yanping Huang, Yida Wang, Yuanzhong Xu, Danyang Zhuo, Eric P. Xing, Joseph E. Gonzalez, and Ion Stoica. 2022. Alpa: Automating Inter- and Intra-Operator Parallelism for Distributed Deep Learning. In *USENIX OSDI*. 559–578.
- [73] Shuai Zheng, Haibin Lin, Sheng Zha, and Mu Li. 2020. Accelerated large batch optimization of bert pretraining in 54 minutes. *arXiv preprint arXiv:2006.13484* (2020).
- [74] Tobias Ziegler, Dwarakanandan Bindiganavile Mohan, Viktor Leis, and Carsten Binnig. 2022. EFA: A Viable Alternative to RDMA over InfiniBand for DBMSs?. In *Data Management on New Hardware*.