Lita: Accelerating Distributed Training of Sparsely Activated Models

Jiamin Li
City University of Hong Kong

Yimin Jiang
ByteDance

Yibo Zhu
ByteDance

Cong Wang
City University of Hong Kong

Hong Xu
The Chinese University of Hong Kong

Abstract

Scaling model parameters usually improves model quality, but at the price of high computation overhead. Sparsely activated models, usually in the form of Mixture of Experts (MoE) architecture, have constant computation cost over their dense counterparts, thus providing opportunities to train and serve a large model at a reasonable cost. However, the distributed training of an MoE model is prone to low efficiency, mainly due to the interleaved all-to-all communication during model computation.

This paper makes three main contributions. First, we systematically analyze the all-to-all overhead in distributed training of MoE. Second, we propose a new communication scheduling scheme based on tensor partitioning that prioritizes the all-to-all operations over other communication, due to its blocking nature. Third, we introduce expert packing that reduces the all-to-all transfer size and incorporates optimizations to mitigate its overheads. Both techniques effectively tackle the all-to-all bottleneck, and we integrate them into a new system called Lita. Experiments on an A100 GPU testbed show that Lita improves the training step time of popular NLP models by up to 1.73x over the state-of-the-art.

1 Introduction

Recent advances in deep learning have shown that a model's quality is typically improved as its number of parameters increases [12, 15, 17, 33]. Based on this, many new frontiers in Computer Vision (CV) and Natural Language Processing (NLP) have been explored using large models [16, 27, 32]. While effective in terms of model quality, the computation cost of training and serving large-scale models is extremely high, which hinders them from being more widely adopted.

Following the basic idea of using massive model parameters while preserving constant computation cost, *sparsely activated* models have recently been introduced [10, 17, 32]. The *Mixture-of-Experts* (MoE) structure is now one of the most popular way to implement sparse activation [10, 11, 32].

For each input, instead of using all parameters, an MoE model intentionally selects just a few of them, i.e. *experts*, for computation. This leads to sub-linear scaling of FLOPS needed with model size. Recent literature [8, 16, 20, 27, 36] has proven the potential of MoE models. For instance, Google develops a family of language models named GLaM using MoE [16]. Compared to GPT-3 with 175 billion parameters, the largest GLaM has 1.2 trillion parameters while only consuming 1/3 of the energy for training. Meanwhile, GLaM still achieves better zero-shot and one-shot performance than GPT-3. Microsoft also reports that their MoE-based language models achieve a 5x training cost reduction compared to a dense model with the same model quality [27].

Given the uptake of MoE, there have been several software systems for efficient MoE training and serving, including Google's Mesh TensorFlow [31], Meta's FairScale [9], Microsoft's DeepSpeed [2] and Tutel [7], etc. They provide APIs for users to replace the conventional dense layers with MoE layers with minimal code changes. They adopt both data parallelism and expert parallelism to accelerate the training process. That is, each device (e.g. GPU) is assigned with a unique expert, and uses all-to-all to receive inputs from other devices and then sends the gradients back to them accordingly. Allreduce is then used to aggregate non-expert gradients in the backward pass.

We focus on the efficiency of MoE training in this work. As some [22, 29] has shown, the all-to-all operation is the main bottleneck in distributed MoE training. We find three main reasons from our empirical analysis. First, all-to-all blocks the subsequent computation operations and needs to be invoked four times for one iteration of a single MoE layer. Second, all-to-all and allreduce often contend for network bandwidth when they overlap in the backward pass, leading to a prolonged blocking period to the computation. Lastly, the transfer size of all-to-all is large and grows linearly with the number of experts. More discussion is presented in §2.2. As efforts are being undertaken to tackle the other time-consuming aspects of MoE training such as the einsum operation in the gating network [7], all-to-all becomes a more pressing issue.

We are thus motivated to systematically tackle the all-to-all bottleneck which has received little attention thus far. Our solution is Lita, an MoE training system that combines priority-based micro-op communication scheduling with pipeline-driven expert packing. We outline our two key ideas below.

First, we prioritize all-to-all over allreduce in order to improve its bandwidth. Existing MoE systems launch separate CUDA streams for the expert-parallel and data-parallel process groups which correspond to all-to-all for expert and allreduce for non-expert parameters, respectively. As there is no coordination between these streams, all-to-all and allreduce can overlap and fair-share the network bandwidth. Unlike allreduce, all-to-all is blocking and cannot be made parallel with the computation process. Thus, prioritizing all-to-all in the backward pass and avoid concurrent allreduce is crucial to reducing the blocking period.

To efficiently prioritize all-to-all, we adopt tensor partitioning which breaks down a tensor into smaller chunks, each of which forms a micro-op. With micro-ops, simple priority scheduling can be applied to guarantee full network bandwidth for all-to-all while allowing allreduce micro-ops to make progress when all-to-all is not present. In addition, micro-ops allow the expert computation to be pipelined with all-to-all, the gain of which becomes more substantial with our next idea.

Second, we consider packing multiple experts on a single device to reduce the all-to-all transfer size and time, because now fewer samples need to be sent over the network to other experts on other devices. One expert per device is commonly adopted today because it requires minimal effort to coordinate interaction across devices. We design mechanisms to optimize the efficiency of expert packing. Given the low computation demand and GPU utilization of experts, we exploit parallel execution of experts to minimize the additional computation time. We also use DRAM offloading to mitigate the additional GPU memory footprint. More interestingly, as more experts colocate, computation time grows while all-to-all time shrinks, and pipelining them becomes more beneficial with smaller "bubbles" in between.

We build Lita based on DeepSpeed MoE [2] and PyTorch, and evaluate it on a cluster with up to 16 Ampere A100 GPUs with 40GB memory and 100Gbps InfiniBand. Results show that Lita accelerates all-to-all by at least 2.21x, and achieves on average 1.57x speedup in overall training step time compared to state-of-the-art system DeepSpeed. Prioritizing all-to-all with micro-op communication scheduling contributes to a 1.38x speed up and expert packing brings a further 1.19x speed up by reducing the all-to-all transfer size.

Our contributions can be summarized as follows:

- We present an in-depth empirical analysis of distributed MoE training to show three main causes for all-to-all to be the performance bottleneck.
- We propose to prioritize all-to-all over allreduce in order to improve its bandwidth and reduce its blocking

period. Lita's scheduler incorporates tensor partitioning and pipelining to perform micro-op scheduling.
- We introduce expert packing to further reduce the all-to-all transfer size, with optimizations including parallel execution and DRAM offloading to mitigate the additional overheads. Lita iteratively searches for the optimal number of experts per device that maximizes pipelining efficiency.
- We implement a concrete prototype system and conduct comprehensive testbed experiments to demonstrate the benefits of our design for various NLP models in a realistic GPU cluster setting.

2 Background and Motivation

We start with an introduction on MoE and on how an MoE model is trained in a distributed manner in §2.1. Then, we motivate our idea by analyzing the performance bottleneck (i.e. all-to-all) in MoE training in §2.2.

2.1 A Primer on MoE

Mixture-of-Experts (MoE) has been adapted to different types of DNN models, and in particular exhibits great potential in improving the performance of language models. GShard [22] and Switch Transformer [17] are two seminal works on scaling Transformer-based language models with MoE layers. We focus on MoE in Transformer-based models in this work, while most of our design also applies to MoE in other models directly (see §8 for more discussion).

Transformer-based models normally use an MoE layer to replace the feed-forward network (FFN) layer. The MoE layer consists of multiple FFNs as *experts* and a gating network (Figure 1a). Every expert FFN is a fully-connected two-layer network using ReLU activation but with different parameters. The gating network dispatches each token to a small number of experts (usually one or two) with weights. The output of the MoE layer is thus the weighted average of outputs from the selected expert(s). This way experts are trained to specialize on certain tokens (i.e., words). The sparsity nature of MoE improves the model scaling in size without increasing the training cost. During training, a load balancing loss is added to ensure the gating network evenly distribute processing burden across experts [32].

Distributed training of MoE. Training MoE models in a distributed manner is necessary due to the tremendous compute requirement of large-scale language models. For efficiency, both data parallelism and MoE-specific *expert parallelism* (as a form of model parallelism) are applied [17, 22]. Existing MoE systems [2, 7, 9, 17, 22, 31] invariably adopt a simple strategy of allocating one unique compute device (e.g., GPU) for each expert in expert parallelism. An all-to-all communication is then needed to send the tokens to their experts selected by the gating network, and another all-to-all is needed to send
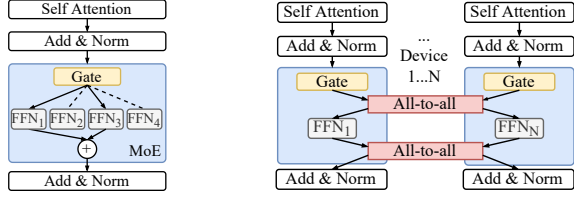
(a) There are four experts and the gate selects two experts.

(b) Distributed training. Data parallelism and expert parallelism are used.

Figure 1: MoE layer in Transformer-based models.

# Experts (GPUs)	Model (# Layers & Parameters)	All-to-all (ms)	Step Time (ms)	Ratio (%)
4	12L + 117M	259	722	36.7
	24L + 233M	589	1684	35.4
	36L + 349M	1479	3894	38.2
16	12L + 419M	333	854	39.5
	24L + 838M	715	1934	37.6
	36L + 1.2B	1545	4293	36.8

Table 1: The completion time of all-to-all and training step time of the Transformer-XL [14] language model in different model sizes. Each FFN layer is replaced with MoE and the number of experts is equal to the number of GPUs similar to the common practice [17]. A100 GPUs with 40GB memory and 100Gb/s InfiniBand are used. We use the MoE implementation in DeepSpeed.

tokens back to the device they belong to in data parallelism to finish the rest of the forward pass as shown in Figure 1b. This simplifies implementation and deployment as no coordination is needed for communication operations. The devices form one data-parallel group, which is also the expert-parallel group. Thus, non-expert gradients are aggregated with allreduce and there is no aggregation on expert gradients among devices.

2.2 MoE Training Bottleneck

Much prior work has identified the training bottleneck of MoE models to be the all-to-all communication brought by expert parallelism [7, 29, 36]. Table 1 shows the completion time of all-to-all operations versus the training step time of various language models in our GPU cluster. On average, all-to-all takes 37.4% – a significant fraction of the step time. In the following, we motivate our work by dissecting three main causes for the cost of all-to-all communication in MoE. Our analysis is based on the common scenario where the number of experts is the same as the number of (GPU) devices.

Synchronous communication. All-to-all is a synchronous operation and blocks the computation process. Figure 2 shows an empirical timeline view of the forward pass of MoE training in our cluster. We observe that expert FFN computation and the combine operation only happens when all-to-all operation completes. During this period, GPU is mostly idle: We use the PyTorch Profiler [6] to profile the GPU activities for 20 steps in each experiment in Table 1, and find that the
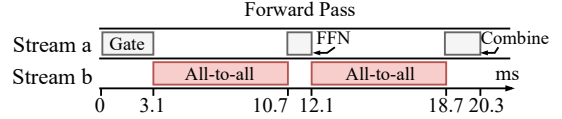

Figure 2: Timeline of forward propagating an MoE layer. We simplify the presentation by bundling GPU kernels here: The computation kernels are grouped by their roles in the MoE layer into Gate, FFN and Combine. The Combine operation involves reshaping the tensors and computing the weighted output. The timeline is taken from a sample run of training the 419M-parameter model in Table 1.
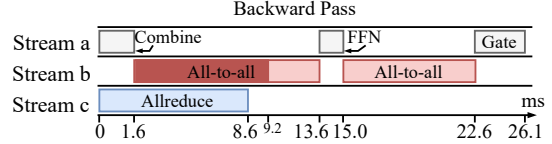

Figure 3: Timeline of backward propagating an MoE layer under hybrid parallelism. The first all-to-all is prolonged by the allreduce operation in Stream b. The shadowed part is its original completion time.

average GPU SM efficiency during all-to-all is 3.7%. Note that FFN computation takes much less time than all-to-all (on average 3.3% of step time compared to 37.4% for all-to-all; more on this in §5.2), which means pipelining all-to-all alone does not resolve this bottleneck. As explained in §2.1, there are two all-to-all operations in the forward pass, one for routing tokens from the previous Add & Norm layer to the chosen experts, and another for sending them back to their original GPU after the expert computation. Therefore a complete training step for one MoE layer incurs four all-to-all operations including two from the backward pass. This aggravates the inefficiency of MoE training.

Hybrid parallelism. The negative impact of all-to-all becomes more salient when we adopt hybrid parallelism [22] in training. Data parallelism requires allreduce to aggregate gradients, and expert parallelism requires all-to-all to exchange tokens. Since the two operations control their own process groups respectively, two dedicated CUDA streams are launched concurrently. This is demonstrated in Figure 3 with the timeline of backward processing in a sample run of MoE training. When the two operations overlap, they share the network bandwidth and their completion time is prolonged. However, we find that the slowdown factor is non-deterministic and varies a lot. We collect the completion times of 1,500 all-to-all operations and plot the CDF of the slowdown factor they endure with allreduce in Figure 4. We observe that the median slowdown is over 1.83x and the worst is 4.14x.

Large data transfer. The second reason for all-to-all to be a bottleneck in MoE is its large data transfer size. Since the expert's FFN architecture ensures that its input data size is the same as the output data size, the data transfers in the two all-to-all operations of the forward pass have the same size. The position-wise expert processing is in token granularity. Thus, the data transfer size is determined by five parameters:
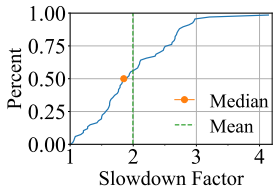
Figure 4: CDF of how much all-to-all is prolonged when it overlaps with allreduce operation. We also plot the median and average slowdown factor.
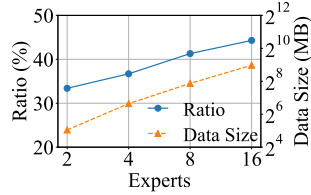
Figure 5: The proportion of all-to-all's completion time over training step time when the number of experts grows. Dashed line plots the data size in one all-to-all operation.

the batch size per GPU *bs*, number of experts (or GPUs) *N*, sequence length *seq_len*, number of selected experts per token *k* as in top-*k*, and the number of expected features in the encoder/decoder inputs *d_model*. That is, the data size (byte) transmitted per GPU in one all-to-all when using single-precision floating-point (i.e. float32) is

$$T = bs \times seq\_len \times d\_model \times k \times \frac{N-1}{N} \times 4, \quad (1)$$

as tokens on each GPU are split into N partitions and $N-1$ partitions are sent to other GPUs (ensured by load balancing). Then, TN bytes of data are transmitted across all GPUs in one all-to-all and in total $4TN$ bytes for one MoE layer during training. The data transfer size grows linearly with the number of experts. Figure 5 presents the empirical evidence of all-to-all's transfer size as the number of experts grows from 2 to 16 (128). With the increasing number of experts, the time taken by all-to-all grows from 33.4% to 44.5% of the step time.

3 Design Overview

Now we introduce our key ideas and Lita's architecture.
Key ideas. Lita is a system designed to accelerate the distributed training of MoE models. Based on the bottleneck analysis in §2.2, we tackle this problem from two aspects.

First, we aim to improve the *network bandwidth of all-to-all* in order to reduce the blocking period of the computation operations. Our key idea here is to *prioritize all-to-all* so it does not fair-share bandwidth with concurrent allreduce (§4). This is achieved using tensor-partitioning. We partition all-to-all and allreduce tensors into small chunks, each of which then forms a *micro-op*. Lita schedules an allreduce micro-op only when there is no all-to-all waiting or ongoing so that all-to-all is guaranteed the full network bandwidth during its lifetime. Without prior information, tensor-partitioning and micro-ops can ensure that in most cases all-to-all can launch as soon as it arrives and allreduce is not deferred excessively.

In addition, micro-ops allow expert FFN computation to be pipelined with all-to-all in both forward and backward pass to further reduce training time. Although the benefit is small due to FFN's much shorter duration compared to all-to-all as

mentioned in §2.2, pipelining becomes more important as we introduce our second idea now.

Our second idea aims to reduce all-to-all's *network traffic by expert packing*. Different from the common practice that places one expert per device, in §5 we propose to pack multiple experts to a single device to reduce the amount of tokens that need to be transmitted across the network. We design optimizations to improve feasibility and efficiency of expert packing, including DRAM offloading for mitigating the additional GPU memory footprint, and parallel execution of experts for reducing the additional expert running time. More interestingly, expert computation time grows while the all-to-all time shrinks as more experts colocate on one device. As such, pipelining them becomes more beneficial with smaller "bubbles" in between.
Lita's workflow. Lita is an MoE-specific system built upon general-purpose ML frameworks such as DeepSpeed [2]. When the training begins, its communication scheduler partitions the incoming operations into micro-ops based on the pre-determined partition size (a hyperparameter in Lita) and prioritizes those of all-to-all. FFN computation is pipelined with all-to-all in both forward and backward passes. Then Lita starts to pack experts when the training process is stabilized after several steps. It increases the number of experts per device iteratively until the pipeline bubbles are minimized without overflowing GPU memory. DRAM offloading and parallel expert execution are used whenever possible.

4 Prioritizing All-to-All in Backward Pass

We have shown that all-to-all is slowed down significantly if it overlaps with allreduce in the backward pass in MoE training. Lita partitions the communication operations into small micro-ops and schedule them strategically in order to prioritize all-to-all without impeding allreduce and the computation process. We introduce the design challenges of prioritizing all-to-all in §4.1. In §4.2, we present Lita's communication scheduler that uses tensor partitioning and pipeline execution to improve the training efficiency.

4.1 Design Challenge

Intuitively, Lita can prioritize all-to-all and avoid concurrent execution with allreduce with strict priority scheduling. All-to-all is always dispatched first if both are present in the queue, and subsequent operations have to wait until the running one finish to make sure allreduce does not share the bandwidth.

It turns out that simply prioritizing all-to-all is not as efficient as one may expect. For work-conservation, when an allreduce arrives first, it should be launched immediately. The problem is when an all-to-all arrives later, though ideally one would preempt the allreduce due to priority scheduling, this is not possible in current multi-GPU communication libraries such as NCCL [3]. The communication primitives are highly
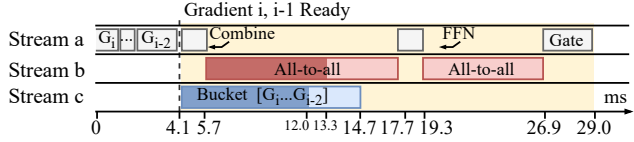
(a) Baseline. Shadowed all-to-all and allreduce are their completion times without concurrent operations. Computing the entire MoE layer's gradients ends at 29.0ms.
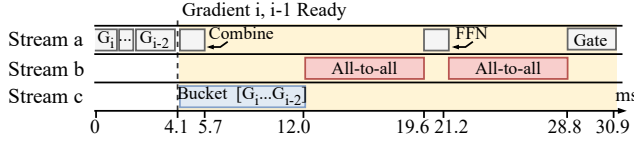
(b) Naively prioritize all-to-all without concurrent transmission can lead to worse results; computing the MoE layer's gradients ends at 30.9ms.
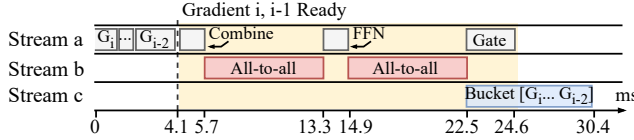
(c) Deferring allreduce to after the second all-to-all leads to better training efficiency; computing the MoE layer's gradients ends at 24.6ms.
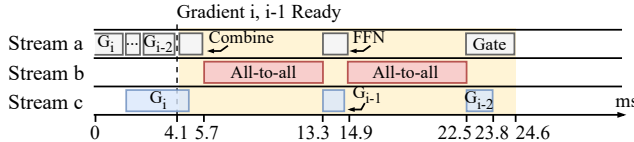

(d) Scheduling results if the arrival time and running time of communication operations are known a priori. The allreduce completes much faster than (c).

Figure 6: Backward pass of MoE training with hybrid parallelism. The yellow background marks the time period of computing the gradients of one MoE layer. Stream a is responsible for the computation process and streams b and c are for communication. This timeline is extracted from a real run of the 419M-parameter benchmark model in Table 1.

optimized and upon being called, their complete transmission strategies are settled and pushed to the CUDA streams. There is no control knob inside each primitive to adjust how it shares resources (e.g. CUDA cores, network bandwidth) with others. Thus, as the example in Figure 6b shows, naively prioritizing all-to-all without concurrent allreduce actually leads to a longer completion time for the first all-to-all and training step time compared to the baseline in Figure 6a.

A potential solution is to obtain the arrival time and running time of the upcoming all-to-all and allreduce, and orchestrate them accordingly to maximize the efficiency. Assuming we know that the allreduce for gradient i can complete before all-to-all and the completion time of gradient $i-1$'s allreduce is shorter than FFN computation. Then we can schedule gradient $i-1$'s allreduce to the gap between the two all-to-all operations at 13.3ms as depicted in Figure 6d. Obtaining the precise knowledge of arrival and running times is, however, a rather daunting if not impossible task. ML frameworks such as PyTorch fuse gradients into buckets based on a user-defined bucket size to optimize allreduce efficiency. Yet in
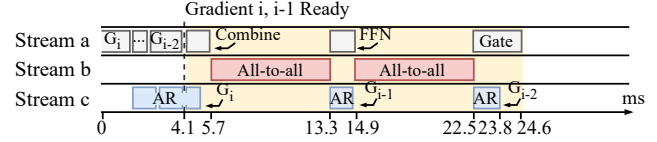
(a) Prioritize all-to-all and partition allreduce tensors. Instead of bucketing gradients, we partition gradient i into three chunks when it is computed.
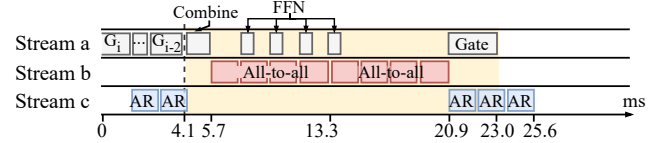
(b) Tensor partitioning for all-to-all and pipeline the FFN computation.

Figure 7: We show the scheduling results of the setup from Figure 6a with tensor partitioning. All-to-all and allreduce micro-ops are of the same size.

large Transformer-based models, gradient sizes are also large; since bucketing is done on the gradient level, the actual bucket size for allreduce varies wildly [5]. Moreover, the implementation details of allreduce make it difficult to acquire a reliable running time estimate as prior work has found out [13].

The other design choice is to blindly defer allreduce until an even number of all-to-all finish as there should be a larger gap between the backpropagation of two MoE layers relative to FFN's backward computation. Figure 6c shows the best scheduling result based on the baseline in Figure 6a. In this case allreduce can be launched when the second all-to-all finishes and completes before the first all-to-all of the next MoE layer (not shown in the figure). Yet, in other (worse) cases, allreduce may still block the all-to-all of the upcoming MoE layer if it takes relatively longer. In the extreme case, no allreduce can be launched until all four all-to-all operations of the current step finish. Since devices have to wait for allreduce before moving onto the optimization phase, this incurs more delay and is undesirable for wait-free backward pass [35].

4.2 Tensor Partitioning and Micro-Ops

To resolve the above challenges, we propose tensor partitioning that breaks down individual communication operations into micro-ops, which can be easily prioritized with higher efficiency.

Tensor partitioning. Unlike tensor bucketing which fuses multiple gradients for an allreduce, Lita partitions each gradient tensor into equal-sized small chunks and executes individual allreduce *micro-ops* independently. This brings two advantages. First, it resolves the varying bucket size problem for allreduce since each micro-op is uniform in size now. Second, micro-ops naturally make better use of bandwidth [25] without causing too much delay to allreduce under priority scheduling. Consider the same setup from Figure 6a, in Figure 7a we partition gradients into five chunks. Before the first all-to-all arrives, Lita launches three allreduce micro-

ops; after the first all-to-all ends, it starts another micro-op to opportunistically make use of the expert computation time. Compared to the scheduling result without micro-ops in Figure 6c, allreduce for gradient $i-2$ now completes 6.6ms or 21.7% faster without prolonging all-to-all. Tensor partitioning does incur overhead due to the partition and concatenation operations before and after an allreduce, but it is mild: the overall overhead in Figure 7a's case is 764us. §7.3 has more details of the overhead analysis.

Pipelining the micro-ops. Intuitively, we can also partition all-to-all which provides an opportunity to pipeline the expert FFN and further reduce the time that computation is blocked. Specifically, we can pipeline the expert computation and all-to-all micro-ops as in Figure 7b. Since the FFN computation is in token granularity, the expert can start computing with a subset of the tokens after one all-to-all micro-op. With pipelining, we can eliminate the FFN time which is 1.6ms in this example.

All-to-all partitioning and pipelining can also be adapted in the forward pass, where the gain actually becomes more salient as we now explain in §5.

5 Expert Packing

The previous section focuses on the backward pass. Here we present Lita's design, where we propose to colocate multiple experts on a single device as a new angle to improve training efficiency. We start with presenting this idea and its potential challenges in §5.1. Then we explain in §5.2 Lita's design in tackling these challenges and maximizing the efficiency gain of expert packing. Expert packing is used throughout the training, and we use the forward pass as the context hereafter.

5.1 Packing Multiple Experts

Key benefits. We propose to place multiple experts on each device to reduce the transfer size and time of all-to-all. In Lita, each device hosts E experts and sends only tokens that need to be handled by experts on other devices through the network. Figures 8 show a 4-expert MoE model with different packing decisions. Existing approach (Figure 8a) places one expert per GPU and forms a 4-GPU expert-parallel group for all-to-all. Now with two experts per device (Figure 8b), we can form two smaller expert-parallel groups (GPU 1 and 2, GPU 3 and 4) and reduce the all-to-all transfer size per GPU by 1/3.

More generally, consider the all-to-all transfer size when each device has E experts. Now that each device only sends out tokens that cannot be handled locally (to peers in the same expert-parallel group), following the same setup in §2.2, the transfer size per GPU is reduced to

$$M = bs \times seq\_len \times d\_model \times k \times \frac{N/E - 1}{N/E} \times 4. \quad (2)$$
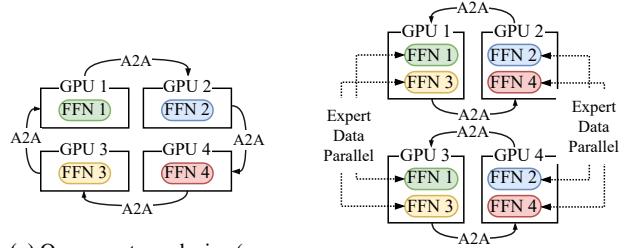
(a) One expert per device (existing approach). All-to-all is done among all four GPUs and there is no allreduce on expert gradients.

(b) Two experts per device. All-to-all is done within two expert-parallel groups, GPU 1 and 2, and 3 and 4; in the backward pass, additional allreduce is needed to aggregate expert gradients between GPUs 1 and 3, and 2 and 4.

Figure 8: Expert packing strategies for a 4-expert MoE model.

Compared to Equation 1, transfer size is reduced to $\frac{N-E}{N-1}$ of the original size.

The parallelism strategy does become more complex with expert packing. For the example in Figure 8b, there are two new data-parallel groups (GPU 1 and 3, GPU 2 and 4) who hold the same experts in addition to the two expert-parallel groups. During the backward pass, additional allreduce is needed to aggregate expert gradients within the two data-parallel groups. This can be pipelined with computation operations and does not incur penalty to training time (§7.4).

Challenges. The central design question is then, how many experts should be packed on a device for optimal performance? Answering this simple question entails various challenges.

First, feasibility. One has to consider the additional GPU memory required now to hold the extra expert parameters which is an imminent and very common problem for large MoE models [30], given that GPU memory is still scarce. Consider a model with L MoE layers, with input dimension d_model and FFN hidden layer dimension d_inner. In total, $L \times d\_model \times d\_inner \times 2 \times (E-1)$ more parameters need to be loaded in GPU memory compared with the current one-expert-per-device approach. For example, two experts per device increases the memory footprint by 19.8% for the 419M-parameter model in Table 1.

Second, efficiency. Placing more experts per device is essentially trading the expert computation time for reducing all-to-all time. Computation execution on GPU is sequential by default, and expert computation time would grow almost linearly with the number of experts in this case. The communication-bounded MoE model would eventually become computation-bounded.

5.2 Design

Lita exploits DRAM offloading to reduce the GPU memory footprint, and uses parallel execution of experts to minimize the marginal increase of computation time (Figure 9). With
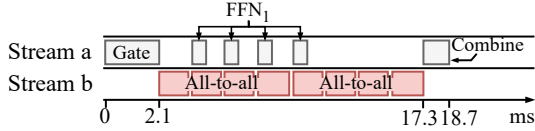
(a) Pipelining with one expert per device in forward pass.
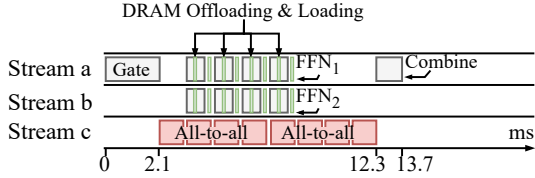
(b) Pipelining with expert packing and DRAM offloading in forward pass. With expert packing all-to-all becomes smaller.

Figure 9: We show how Lita executes packed experts and exploit DRAM offloading techniques to swap expert parameters between host and device.

these optimizations, Lita finds the optimal number of experts per device when the resulting expert FFN's execution time equals that of all-to-all so pipelining them is most efficient.

DRAM offloading. We leverage host DRAM to offload one layer of the two-layer FFNs during expert computation [30]. The two layers in an FFN expert have an identical number of parameters as the first layer transforms the token embeddings from *d_model* dimension to *d_inner* dimension, while the second one does the reverse. We load the first layers of experts on GPU memory during initialization. After receiving a subset of tokens, Lita starts computing their results of the first layer (Figure 9b). Then the first layers are offloaded to DRAM and the second layers are then loaded to GPU memory for computation. Upon completion, the location of the two layers are again swapped to get ready for the next subset of tokens. The number of offloading and loading operations are proportional to the number of micro-op created by Lita with the communication scheduler. We provide overhead analysis in §7.4.

Parallel execution. Lita also exploits parallel execution of the FFNs to minimize the additional running time caused by packing multiple experts. Figure 9b shows the parallel execution of two FFN experts on one device. We empirically find that the expert FFNs are particularly suitable for parallel execution. One expert per GPU leaves many SMs under-utilized with modern GPUs: as shown in Table 2, the average achieved occupancy [1] for four different FFNs is around 35% during the execution period, implying that many warps are not utilized. Table 3 further shows the time of concurrently executing multiple experts. When two experts are executed in parallel on the same GPU, the running time is only 1.04x that of single expert computation; with four experts in parallel, the computation time increases to at most 1.67x for large experts. The average achieved occupancy improves by 66% and 72% for two and four concurrent FFNs, respectively. Thus parallel execution provides salient efficiency gains.

Overall pipelining efficiency. With the two optimizations,

d_model	d_inner	Completion Time (ms)	Mean Achieved Occupancy* (%)
512	1024	1.248	32
512	2048	2.317	34
1024	2048	4.263	34
1024	4096	8.367	36

*: Achieved occupancy indicates the fraction of active warps at the moment on the SMs. Mean achieved occupancy reports the average values of all warp schedulers during the kernel execution period; the larger the better [1].

Table 2: SM's achieved occupancy [1] of the FFN expert computation in different configurations on our A100 GPUs.

d_model	d_inner	Concurrent FFN completion time			
		2		4	
512	1024	1.02x	1.272ms	1.12x	1.835ms
512	2048	1.02x	2.363ms	1.33x	3.313ms
1024	2048	1.04x	4.433ms	1.64x	6.991ms
1024	4096	1.07x	8.952ms	1.67x	13.972ms

Table 3: We measure the end-to-end completion times of concurrent FFNs in different configurations and compare it to single FFN computation.

Lita determines the optimal number of experts per device by examining the overall pipelining efficiency of computation and communication micro-ops. Specifically, we notice that a single FFN micro-op takes much less time than its corresponding all-to-all micro-op (Figure 9a). In Table 1, FFN only takes on average 12.4% of the time compared to all-to-all that takes 37.4%. In other words, experts have to wait on all-to-all in the pipeline. Ideally, the FFN and all-to-all micro-ops should take a similar time so that both compute capacity and network bandwidth are fully utilized without any bubbles in the pipeline. Thus, Lita adopts the following approach: starting with one expert, it iteratively increases the number of experts per device in powers of two, until the FFN computation plus DRAM offloading time exceeds that of the all-to-all micro-op, or until GPU memory is exhausted, whichever comes first. Figure 9b presents an example with expert packing. Compared to Figure 9a, the MoE layer's completion time is reduced from 18.7ms to 13.7ms (32.1%).

6 Implementation

We have implemented Lita on DeepSpeed MoE and PyTorch using C++ and Python. PyTorch 1.10, CUDA 11, and NCCL 2.10 are used. We modify the existing implementation of distributed training in PyTorch to support prioritized communication operations and implement our dynamic expert packing strategy of MoE layer in DeepSpeed. The current implementation has 3216 lines of code in total.

6.1 Communication Scheduler

Lita's communication scheduler runs a single thread, handling events including enqueuing micro-ops, and scheduling

and launching them based on priority. Each device runs an instance of the scheduler. Since all the devices perform the same tasks and follow the same scheduling policy, there is no communication among the scheduler instances.

The scheduler maintains a priority queue and each micro-op task has a status and a priority rank. The status can be queuing, running, and completed. Tasks that are completed would be removed from the queue. To avoid concurrent communication operations, the scheduler stops launching micro-ops if the leading task in the queue is running. Moreover, the scheduler stops launching allreduce micro-ops if the computation process has entered the gating computation in forward pass or the combining computation in the backward pass, since this implies all-to-all is imminent.

The micro-op size is passed in as a hyperparameter. Lita partitions the data accordingly using the built-in APIs `chunk` and `cat` in LibTorch to chunk the transfer data and concatenate the data chunks. The micro-ops are assigned with a priority rank (i.e. 0 for all-to-all and 1 for allreduce). We avoid putting the data chunks from two gradients into the same micro-op. This simplifies the subsequent concatenation operation.

To enable pipelining in the MoE layer, we partition the all-to-all transfer data in the token dimension (while respecting the partition size) so that each all-to-all micro-op sends out a subset of complete tokens. Thus, FFN is ready to start once part of the tokens are transferred.

6.2 Expert Packing Implementation

Expert packing coordinator. We embed a packing controller in the MoE model and it runs a single thread. The expert packing is dynamically adjusted after the training process is stable. In the forward pass, the controller records the completion times of all-to-all micro-ops and FFN micro-ops.

When the FFN micro-ops are shorter than all-to-all, the controller starts to pack experts in three steps. First, we initialize the new data-parallel and expert-parallel process groups based on the packing decisions. Second, the controller inserts a one-time synchronous all-to-all to exchange expert parameters between packed devices that would be invoked at the upcoming iteration. Now each device holds the packed expert parameters. Finally, both the forward and backward pass would adopt parallel execution when the size of the expert size is larger than one.

Parallel execution of CUDA kernels. We implement parallel expert computation in the forward and backward pass as follows. We use CUDA Runtime to create and destroy the same number of CUDA streams as the number of experts per device at the beginning of a training step. Inside the MoE computation kernel, we reserve these streams and dispatch each expert to a unique stream. In each stream, we use the cuBLAS API `cublasXgemm` to perform matrix multiplication inside the FFN experts. Overall, we wrap these

up into two new operations `parallel_ffn_forward` and `parallel_ffn_backward` to replace the original one.

Dynamically-adjusted parallelism. Packing multiple experts per device requires more delicate coordination of communication operations. With Lita, the parallel process group changes along with the expert packing decision. In the existing approach, this can only be achieved by restarting the training model with new process groups. In Lita, we modify the existing implementation to support dynamic parallelism without interrupting the training process.

Allreduce operations in `DistributedDataParallel` API are predefined before the training starts. We register the allreduce hook to every gradient in the model regardless of its type (i.e. expert, non-expert). When allreduce of expert gradients are pushed to the queue, we let the communication scheduler to decide whether it should be launched and the process groups it operates on. In one-expert-per-device case, allreduce of expert gradients are skipped. Otherwise, they are launched on the expert parallel groups as the packing decision suggests. For all-to-all, we just need to replace the process group input before its execution.

7 Evaluation

We evaluate Lita with Transformer-based language models. The highlights of our findings are:
- Lita improves the end-to-end training step time by an average of 1.57x across all models. The largest gain is 1.73x in 8-expert MoE models.
- Training step time is improved by an average of 1.37x with Lita's communication scheduler. The benefit is mainly from the mitigation of period blocked by all-to-all in the backward pass. The average GPU utilization is improved by an average of 30% over Baseline.
- Expert packing effectively improves the pipelining efficiency by up to 2.61x in 16-expert models. The step time is improved by an average of 1.41x of Baseline for the Transformer-XL model in different configurations.
- The overhead incurred by Lita's communication scheduler is 1.02% of the step time and the packing overhead is mostly hidden by threading and pipelining.

7.1 Setup

Testbed setup. Our testbed has four worker nodes. Each node has 4 Ampere A100 GPUs with 40GB memory and is equipped with 100Gbps InfiniBand.

DNN models. We convert three common Transformer-based dense language models to MoE models.
- Transformer-XL [14]: a 24-layer encoder model.
- GPT-2 [28]: a 12-layer decoder model.
- BERT2GPT2 [34]: a 12-layer encoder-decoder model.

All the FFN layers in these models are converted to MoE layers. We vary the number of experts of an MoE layer from

2, 4, 8, to 16. We adopt a top-2 gate as suggested by [17]. Load balancing at the gate is enabled so that the volume of tokens received by each expert is the same.

Metrics. We consider five metrics to evaluate Lita.

- Training step time: It measures a complete step of distributed MoE training, including the forward and backward pass.
- MoE layer completion time: We measure the completion time of an MoE layer in both forward and backward directions averaged over all MoE layers in a model. It consists of the gating network, all-to-all, expert computation, and the combine operation that averages the expert output. Note that the allreduce over MoE layer gradients is not included since it can be fully pipelined with the computation.
- All-to-all completion time: We record the completion time of all-to-all in both forward and backward pass. This is obtained by summing up the communication kernels in the profiled timeline.
- GPU cycle utilization: We also monitor the GPU utilization during model training and SM occupancy during the FFN expert computation.
- Pipelining efficiency: It is used to assess whether expert computation can be fully pipelined with all-to-all. This is only defined for Lita with tensor partitioning and pipelining. We measure the pipelining efficiency to be the fraction of non-idle time in the computation CUDA stream during the total all-to-all duration. DRAM offloading overhead counts into the non-idle period as it happens in the same GPU stream.

Throughout the evaluation we use PyTorch Profiler to profile the training process and obtain the execution times of CUDA kernels and GPU activities. All the reported numbers are averaged over 50 training steps after a 10-step warm-up period. Since the optimization introduced by Lita does not affect the precision of model parameters, model accuracy is unchanged and we omit its evaluation.

Lita configurations. Lita's micro-op communication scheduler adopts a tensor partition size of 30MB, which can minimize the period blocked by all-to-all in most cases. We show how the partition size affects the training step time in §7.3. Expert packing is launched at the 10-th step of each training task and is adjusted every four steps. We present how packing decisions are made along the training process in §7.4.

Baselines. We use the vanilla DeepSpeed [2] as the Baseline. Each GPU handles one unique expert. We also provide a brief comparison with the open-source version of Tutel [7], which shows very similar performance to DeepSpeed.

7.2 Overall Performance

We evaluate Lita's overall performance in this section. Note that Lita is evaluated when the expert packing decision is stabilized; all settings here use 2 experts per device as the
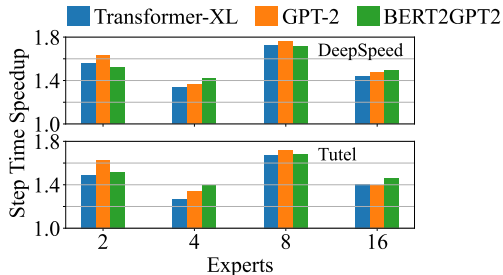

Figure 10: Speedup of training step time against two Baselines.

best strategy except Transformer-XL with 16 experts, which uses 4 experts per device. The number of GPUs is equal to the number of experts per MoE layer in both Baseline and Lita. Analysis of the two key ideas behind Lita is separately presented in §7.3 and §7.4.

Training step time. Figure 10 shows Lita's speedup in step time over Baseline and Tutel using the three models with different numbers of experts. All other aspects of the models stay the same (e.g. sequence length, hidden states dimension, etc.). Compared to Baseline (DeepSpeed), step time is reduced by an average of 1.37x and 1.47x for the 4-expert and 16-expert cases, respectively, and by an average of 1.71x and 1.73x for 2-expert and 8-expert models, respectively. The 2-expert and 8-expert cases see more significant gains as Lita's packs two experts per device as mentioned before. The 2-expert case thus converts to pure data parallelism and eliminates all the all-to-all operations; the 8-expert models avoid inter-node all-to-all as our servers have 4 GPUs each. Lita's speedup over Tutel is slightly smaller than that of DeepSpeed. Thus in the following we only use DeepSpeed as the baseline.

MoE layer completion time. We specifically seek to understand Lita's gain in the MoE layer in both the forward and backward pass. As Figures 11 and 12 show, similar to step time, the gain here in the 2-expert and 8-expert cases is the largest. The forward and backward pass of MoE layers in the 2-expert case are accelerated by 1.84x and 2.41x, and in the 8-expert case by 1.89x and 2.32x, respectively. Meanwhile, since backward pass in Baseline suffers from the interference of allreduce while the forward pass does not, the improvement in the backward pass is more significant.

All-to-all completion time. We then compare the completion time of all-to-all between Lita and Baseline in the backward pass. Lita's communication scheduler prioritizes all-to-all and avoids concurrent execution with allreduce. Moreover, expert packing reduces the all-to-all transfer size. Figure 13 shows an average improvement of 2.21x, 2.39x, and 2.31x in 4-, 8-, and 16-expert cases, respectively. As discussed before in the 2-expert case Lita packs both experts of an MoE layer on each device which leads to pure data parallelism. Therefore, no all-to-all operation is required.

GPU cycle utilization. To see Lita's impact on resource efficiency, we monitor the GPU activities throughout the training process. Table 4 shows the average GPU utilization in the
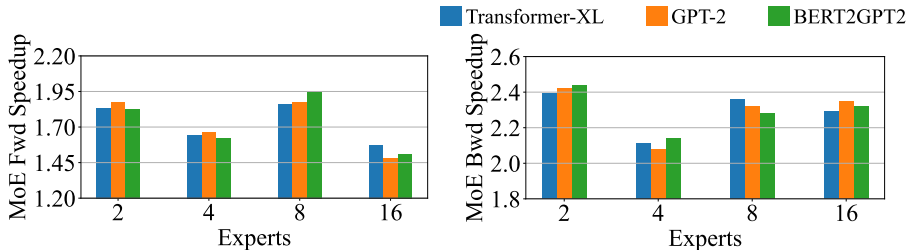
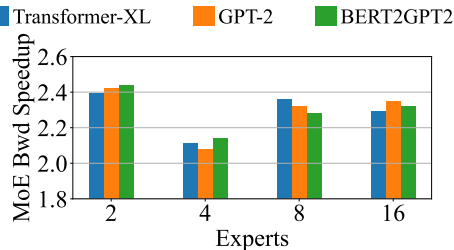**Figure 11:** Speedup of MoE layer's forward pass completion time.

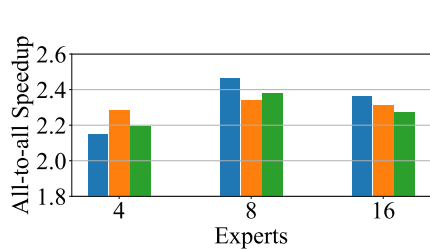**Figure 12:** Speedup of MoE layer's backward pass completion time.

Figure 13: Speedup of all-to-all completion time in both forward and backward pass.

Expert	Model	Avg. GPU utilization[*]		Mean achieved occupancy (FFN)	
		Baseline	Lita	Baseline	Lita
16	Transformer-XL	69%	85%	36%	67%
	GPT-2	65%	84%	33%	73%
	BERT2GPT2	68%	85%	35%	66%

[*]: GPU utilization indicates the fraction of time that at least one SM is active.

Table 4: Average GPU utilization and SM's mean achieved occupancy of expert FFN computation throughout the training process.

MoE layer and SM's achieved occupancy in FFN computation across models with 16 experts. Average GPU utilization is improved by at least 16% as the period blocked by all-to-all is minimized with Lita's communication scheduler. SM's achieved occupancy is also increased by an average of 1.98x as we adopt parallel execution in the expert computation.

7.3 Deep Dive: Communication Scheduler

We now present an in-depth analysis of Lita's priority-based micro-op scheduler, aiming to understand how tensor partitioning and pipelining contribute to the training step time gain. We also investigate the impact of partition size. For fairness all experiments here are obtained without expert packing in Lita, i.e. there is only one expert per device as in Baseline.

Tensor partitioning and pipelining. To justify our communication scheduler's design, we incrementally add the key design choices to Baseline and see their corresponding gain: first priority scheduling, then tensor partitioning, and lastly pipelining. Besides, we consider a fixed scheduling strategy where allreduce is always scheduled between pairs of all-to-all operations (i.e. two MoE layers) with tensor fusion enabled in PyTorch's `DistributedDataParallel` by default (same as Baseline).

Figure 14 shows the step time comparison. We make several interesting observations here. First, using priority brings about 10%–30% gain over Baseline in most cases, with an average of 24%. Priority scheduling in general presents more benefit when more devices and nodes are used in training. The main reason is that all-to-all's slowdown due to sharing bandwidth with allreduce is more severe as training scales out. Second, tensor partitioning significantly improves the benefit of prioritizing all-to-all: the step time is improved over Baseline by 1.36x, 1.36x, 1.41x and 1.42x in 2-expert,

4-expert 8-expert and 16-expert cases, respectively on average. On the other hand, pipelining's gain is limited as expected, since expert computation takes much less time than all-to-all without expert packing (recall §4.2). Overall, all three design choices can effectively reduce all-to-all's completion time.

We also observe that the relative benefit of priority scheduling and tensor partitioning is model-specific: GPT-2 enjoys much more gain from priority compared to tensor partitioning while the other two models do not exhibit such clear pattern. This is likely due to the degree of overlapping of all-to-all and allreduce: most allreduce can fit in between all-to-all operations in GPT-2, and as a result using priority scheduling alone is very beneficial.

Finally, the fixed scheduling strategy leads to the smallest gains in almost all cases. This is because (1) all-to-all is not prioritized and still has to fair-share bandwidth with allreduce when it arrives, and (2) tensors are not partitioned which aggravates the impact of allreduce. This demonstrates again the effectiveness of our design in prioritizing all-to-all with smaller tensors instead of using fixed heuristics that cannot opportunistically maximize the scheduling efficiency.

Partition size. We also evaluate the impact of partition size on the communication scheduler. Figure 15 shows the step time of 16-expert models when we gradually increase the partition size from 10MB to 100MB. We find that a partition size beyond 50MB slows down Transformer-XL and BERT2GPT2 compared with 30 MB. As long as the period blocked by all-to-all is minimized, step time would be minimum. Therefore, for each model and setting, there are multiple optimum partition sizes. Theoretically, the scheduler can control the operations more precisely with a smaller partition size. In practice, a smaller partition size (below 10MB) may cause a heavy transmission overhead in each micro-op, thus degrading the overall performance [26].

Overhead analysis. We provide a brief analysis of the overhead incurred by Lita's communication scheduler. First, the preprocessing and postprocessing, including tensor partitioning and concatenation, take an average 1.02% of the step time. Second, we measure the transmission overhead of micro-ops. We sum up running times of all the communication micro-ops in Lita and compare against that without partitioning in Baseline. The average completion time is lengthened by 1.7%,

(a) Transformer-XL.
(b) GPT-2.
(c) BERT2GPT2.

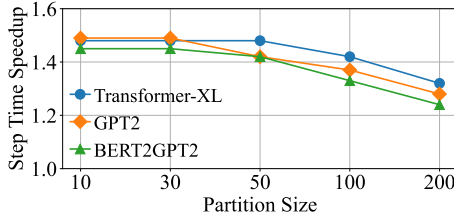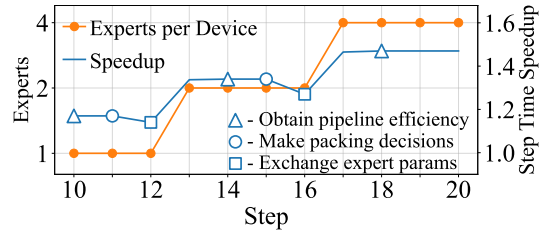**Figure 14:** Speedup of training step time over Baseline with different design choices of the communication scheduler.

Figure 15: Step time speedup when the partition size in Lita increases from 10MB to 200MB. We use the 16-expert models here.

which is negligible.

7.4 Deep Dive: Expert Packing

We then evaluate the effectiveness and overhead of Lita's expert packing design. Here Lita's communication scheduler is enabled so we can discuss the gain from pipelining.

Packing and step time. Figure 16 shows how the packing decision evolves and its effect on step time during the training process of the 16-expert Transformer-XL in different *d_model* and *d_inner* configurations. Recall Lita iteratively searches the best packing decision. In each round, it first profiles the current pipelining efficiency of the MoE layers (steps 10, 14, and 18 in Figure 16a). At the next step Lita updates the packing decision if necessary. The models experience a one-off slowdown (steps 12 and 16 in Figure 16a) as an additional synchronous all-to-all is launched to exchange packed expert parameters among devices at the next step. After this, the new packing decision becomes effective and step time improves. Thus it takes 4 steps to complete one round of adjustment for the packing decision.

We observe that Lita can quickly find model-specific packing decision. In Figures 16a and 16b, the optimal decision is to pack 4 experts per device, and in Figure 16c 2 experts per device is the best. This is because the model in Figure 16c has more parameters and the expert FFN computation becomes longer than all-to-all with 4 experts on one device.

Pipelining efficiency. We calculate the pipelining efficiency of Lita before and after adopting expert packing in Table 5. The average improvement is 2.14x and 2.43x in 4-expert and 16-expert cases. The pipelining efficiency without packing in the 4-expert case is higher than that in 16-expert case because 4-expert case only involves relatively short intra-node

(a) *d_model* is 512 and *d_inner* is 1024.
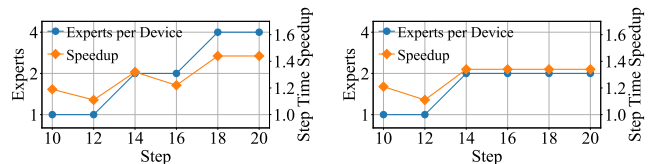
(b) *d_model* is 512 and *d_inner* is 2048.

(c) *d_model* is 1024 and *d_inner* is 2048.

Figure 16: The expert packing decision process of the 16-expert Transformer-XL with different *d_model* and *d_inner* configurations.

Expert	Model	Pipelining Efficiency		
		w/o Packing	w/ Packing	(Experts per Device)
4	Transformer-XL	37%	84%	2
	GPT-2	38%	87%	2
	BERT2GPT2	44%	83%	2
16	Transformer-XL	33%	86%	4
	GPT-2	36%	85%	2
	BERTGPT2	34%	79%	2

Table 5: Pipelining efficiency comparison with and without expert packing. The FFN parameters (*d_model*, *d_inner*) for Transformer-XL, GPT-2, and BERTGPT2 are (512, 2048), (768, 3072) and (768, 3072), respectively.

all-to-all. The expert FFN micro-op computation time is thus closer to the completion time of all-to-all micro-op. When the packing decision is four experts per device in the 16-expert case, pipelining efficiency is improved by 2.61x for Transformer-XL. We find that two experts per device can achieve the best pipelining efficiency in most cases. Packing more experts slows down computation and degrades the pipelining efficiency.

Overhead analysis. Expert packing incurs two types of overhead. (1) Data swapping between host DRAM and device. The average overhead per MoE layer is 22.3% of the computation time of one expert. Since it is pipelined with the long

all-to-all operations, its impact on end-to-end step time is little. In fact Table 5 shows that the expert computation time including DRAM offloading overhead is fully masked by all-to-all, since otherwise the pipelining efficiency would be 100% by our definition in §7.1. (2) Expert parameter exchanging via all-to-all, which is already presented in Figure 16. Though it prolongs the step time, its impact is limited to a single step. Besides, the time for making packing decisions has a negligible impact on the overall performance as it is running as a thread concurrently with the main computation process.

8 Discussion

We discuss some potential issues of Lita in this section.

Pipelining gate computation. Lita's current design pipelines all-to-all with expert computation in both forward and backward pass. We find that the einsum and cumsum operations in the gating network are time-consuming [22]. In principle, all-to-all can also be pipelined with the gate computation to further reduce the training step time. There are two critical issues here. First, in the forward pass, pipelining changes the gating results since the top-k selection is performed in a smaller pool of samples now. This may affect the model convergence and accuracy. Second, pipelining in the backward pass requires heavy engineering effort to modify the PyTorch's `Autograd` package. We need to implement a customized backward function that allows an input with shape different from the output shape in the forward function. Moreover, we find that Tutel [7] has introduced optimized kernels for the einsum operation of the gate, which enhances the efficiency of the gate computation. Therefore, the gain of pipelining gate would be diminished.

Parallel execution of CUDA kernels. In Lita, we propose to adopt parallel execution for FFN computation and customize a forward and a backward operation to improve the computation efficiency and the utilization of the device resources. Existing work [23, 24] has identified such problem in a more general scope and introduces a compiler-level optimization technique to exploit inter-operation and intra-operation parallelism. These solutions could further improve the parallel FFN computation efficiency in MoE systems.

9 Related Work

MoE systems and optimization. An increasing number of general ML frameworks have incorporated support for MoE. Google's Mesh TensorFlow [31], Meta's FairScale [9] and Microsoft's DeepSpeed [2] develop MoE-specific APIs for existing ML frameworks.

Recent literature has also proposed some MoE-specific optimization techniques. DeepSpeed [29] leverages flexible combinations of parallelism strategies and enables distributed training for MoE models. DeepSpeed also introduces an MoE inference system that optimizes the all-to-all communication to reduce latency. It adopts a hierarchical all-to-all design to address the inter/intra-server GPU topology.

Tutel [7] is designed to optimize the computation kernels used in the MoE layer. Off-the-shelf DNN operations provided by PyTorch [4] incur significant performance overheads, especially for the einsum operation which is called at least four times in one MoE layer. Tutel replaces them with a fast cumsum-minus-one operation using explicit data layout transformations. It also uses kernel-fusion to fuse the gating function into a single kernel and a dense token-to-expert mapping table to represent the gating decisions.

FasterMoE [19] proposes a performance model to analyze the end-to-end performance of MoE systems. Guided by this model, they propose a dynamic shadowing approach to handle the imbalanced load among experts. Instead of sending tokens to the experts, they replicate popular experts on each device to reduce the transfer size. They also design a topology-aware expert selection strategy that relieves network congestion by sending tokens to experts with a lower latency.

Compared to these efforts, Lita focuses on optimizing all-to-all in MoE training which is complementary.

Communication acceleration in distributed training. Our community has proposed several communication schedulers for generic distributed training [18, 26]. The objective is to better overlap the communication and computation operations in the backward pass and prioritize the communication of former layers over latter layers in the model. In Lita, we leverage the domain-specific insight that all-to-all should be prioritized over allreduce in MoE training, which is different from prior work. BytePS [21] proposes to reduce the communication traffic by utilizing the heterogeneous GPU/CPU resources in a training clusters. Lita can also benefit from this idea, since more available bandwidth can be left to all-to-all operations. The acceleration techniques of these existing work can all be integrated to MoE distributed training for aggregating gradients with allreduce.

10 Conclusion

We present Lita, a new system that accelerates all-to-all in distributed training of large MoE models. Through a systematic analysis, we build Lita upon two key ideas: first to prioritize all-to-all over allreduce using tensor partitioning and pipelining to improve its bandwidth, and second to pack multiple experts into one device to reduce the transfer size of all-to-all. We implement Lita over DeepSpeed and perform extensive testbed evaluation using A100 GPUs and 100Gbps InfiniBand to show that Lita significantly improves efficiency of training MoE layers and thus the overall training step time.

We plan to extend our work in two directions. First, accelerating all-to-all from the bottom up. All-to-all can be decomposed into multiple send and receive operations. A fine-grained implementation of all-to-all could eliminate redundant operations and further elevate the efficiency when gate decisions follow a certain traceable pattern. Second, we plan to also study MoE inference. Large MoE inference models are inevitably deployed in a distributed manner, where

communication adds to the response time. Minimizing infer-
ence latency is critical to the practicality of MoE systems.

References

[1] Achieved Occupancy. `https://docs.nvidia.com/gameworks/content/developertools/desktop/analysis/report/cudaexperiments/kernellevel/achievedoccupancy.htm`.

[2] DeepSpeed. `https://github.com/microsoft/DeepSpeed`.

[3] NCCL. `https://github.com/NVIDIA/nccl`.

[4] PyTorch. `https://pytorch.org`.

[5] PyTorch Distributed Data Parallel. `https://pytorch.org/docs/stable/notes/ddp.html`.

[6] PyTorch Profiler. `https://pytorch.org/blog/pytorch-profiler-1.9-released/`.

[7] Tutel. `https://github.com/microsoft/tutel`.

[8] Mikel Artetxe, Shruti Bhosale, Naman Goyal, Todor Mihaylov, Myle Ott, Sam Shleifer, Xi Victoria Lin, Jingfei Du, Srinivasan Iyer, Ramakanth Pasunuru, et al. Efficient large scale language modeling with mixtures of experts. *arXiv preprint arXiv:2112.10684*, 2021.

[9] Mandeep Baines, Shruti Bhosale, Vittorio Caggiano, Naman Goyal, Siddharth Goyal, Myle Ott, Benjamin Lefaudeux, Vitaliy Liptchinsky, Mike Rabbat, Sam Sheiffer, Anjali Sridhar, and Min Xu. Fairscale: A general purpose modular pytorch library for high performance and large scale training. `https://github.com/facebookresearch/fairscale`.

[10] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. Conditional computation in neural networks for faster models. *arXiv preprint arXiv:1511.06297*, 2015.

[11] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

[12] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.

[13] Minsik Cho, Ulrich Finkler, David Kung, and Hillery Hunter. BlueConnect: Decomposing All-Reduce for Deep Learning on Heterogeneous Network Hierarchy. In *Proc. MLSys*, 2019.

[14] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-XL: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.

[15] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.

[16] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. GLaM: Efficient Scaling of Language Models with Mixture-of-Experts. *arXiv preprint arXiv:2112.06905*, 2021.

[17] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*, 2021.

[18] Sayed Hadi Hashemi, Sangeetha Abdu Jyothi, and Roy Campbell. Tictac: Accelerating distributed deep learning with communication scheduling. In A. Talwalkar, V. Smith, and M. Zaharia, editors, *Proc. MLSys*, 2019.

[19] Jiaao He, Jidong Zhai, Tiago Antunes, Haojie Wang, Fuwen Luo, Shangfeng Shi, and Qin Li. FasterMoE: modeling and optimizing training of large-scale dynamic pre-trained models. In *Proc. ACM SIGPLAN PPoPP*, pages 120–134, 2022.

[20] Xianyan Jia, Le Jiang, Ang Wang, Jie Zhang, Xinyuan Li, Wencong Xiao, Yong Li, Zhen Zheng, Xiaoyong Liu, Wei Lin, et al. Whale: Scaling deep learning model training to the trillions. *arXiv preprint arXiv:2011.09208*, 2020.

[21] Yimin Jiang, Yibo Zhu, Chang Lan, Bairen Yi, Yong Cui, and Chuanxiong Guo. A unified architecture for accelerating distributed {DNN} training in heterogeneous {GPU/CPU} clusters. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*, pages 463–479, 2020.

[22] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.

[23] Lingxiao Ma, Zhiqiang Xie, Zhi Yang, Jilong Xue, Youshan Miao, Wei Cui, Wenxiang Hu, Fan Yang, Lintao Zhang, and Lidong Zhou. Rammer: Enabling Holistic Deep Learning Compiler Optimizations with rTasks. In *Proc. USENIX OSDI*, 2020.

[24] Wei Niu, Jiexiong Guan, Yanzhi Wang, Gagan Agrawal, and Bin Ren. Dnnfusion: Accelerating deep neural networks execution with advanced operator fusion. In *Proc. ACM SIGPLAN PLDI*, 2021.

[25] Kay Ousterhout, Aurojit Panda, Joshua Rosen, Shivaram Venkataraman, Reynold Xin, Sylvia Ratnasamy, Scott Shenker, and Ion Stoica. The Case for Tiny Tasks in Compute Clusters. In *Proc. USENIX HotOS*, 2013.

[26] Yanghua Peng, Yibo Zhu, Yangrui Chen, Yixin Bao, Bairen Yi, Chang Lan, Chuan Wu, and Chuanxiong Guo. A generic communication scheduler for distributed dnn training acceleration. In *Proc. ACM SOSP*, 2019.

[27] Andrey Proskurin. DeepSpeed: Advancing MoE inference and training to power next-generation AI scale. https://www.microsoft.com/en-us/research/blog/deepspeed-advancing-moe-inference-and-training-to-power-next-generation-ai-scale.

[28] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. 2019.

[29] Samyam Rajbhandari, Conglong Li, Zhewei Yao, Minjia Zhang, Reza Yazdani Aminabadi, Ammar Ahmad Awan, Jeff Rasley, and Yuxiong He. DeepSpeed-MoE: Advancing Mixture-of-Experts Inference and Training to Power Next-Generation AI Scale. *arXiv preprint arXiv:2201.05596*, 2022.

[30] Jie Ren, Samyam Rajbhandari, Reza Yazdani Aminabadi, Olatunji Ruwase, Shuangyan Yang, Minjia Zhang, Dong Li, and Yuxiong He. ZeRO-Offload: Democratizing Billion-Scale Model Training. In *Proc. USENIX ATC*, 2021.

[31] Noam Shazeer, Youlong Cheng, Niki Parmar, Dustin Tran, Ashish Vaswani, Penporn Koanantakool, Peter Hawkins, HyoukJoong Lee, Mingsheng Hong, Cliff Young, Ryan Sepassi, and Blake Hechtman. Mesh-TensorFlow: Deep learning for supercomputers. In *Proc. ACM NeurIPS*, 2018.

[32] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*, 2017.

[33] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

[34] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2020.

[35] Hao Zhang, Zeyu Zheng, Shizhen Xu, Wei Dai, Qirong Ho, Xiaodan Liang, Zhiting Hu, Jinliang Wei, Pengtao Xie, and Eric P Xing. Poseidon: An efficient communication architecture for distributed deep learning on GPU clusters. In *Proc. USENIX ATC*, 2017.

[36] Barret Zoph, Irwan Bello, Sameer Kumar, Nan Du, Yanping Huang, Jeff Dean, Noam Shazeer, and William Fedus. Designing Effective Sparse Expert Models. *arXiv preprint arXiv:2202.08906*, 2022.