

# LANCET: ACCELERATING MIXTURE-OF-EXPERTS TRAINING VIA WHOLE GRAPH COMPUTATION-COMMUNICATION OVERLAPPING

Chenyu Jiang<sup>1\*</sup> Ye Tian<sup>1\*</sup> Zhen Jia<sup>2</sup> Shuai Zheng<sup>3†</sup> Chuan Wu<sup>1</sup> Yida Wang<sup>2</sup>

## ABSTRACT

The Mixture-of-Expert (MoE) technique plays a crucial role in expanding the size of DNN model parameters. However, it faces the challenge of extended all-to-all communication latency during the training process. Existing methods attempt to mitigate this issue by overlapping all-to-all with expert computation. Yet, these methods frequently fall short of achieving sufficient overlap, consequently restricting the potential for performance enhancements. In our study, we extend the scope of this challenge by considering overlap at the broader training graph level. During the forward pass, we enable non-MoE computations to overlap with all-to-all through careful partitioning and pipelining. In the backward pass, we achieve overlap with all-to-all by scheduling gradient weight computations. We implement these techniques in Lancet, a system using compiler-based optimization to automatically enhance MoE model training. Our extensive evaluation reveals that Lancet significantly reduces the time devoted to non-overlapping communication, by as much as 77%. Moreover, it achieves a notable end-to-end speedup of up to 1.3 times when compared to the state-of-the-art solutions.

## 1 INTRODUCTION

Recent research has prompted a continuous trend of constructing larger DNN models across application domains. However, directly adopting wider or deeper network architecture typically leads to a proportional increase in computation. In contrast, Mixture of Experts (MoE) (Shazeer et al., 2017; Lepikhin et al., 2020) has the ability to increase the parameter size without escalating the total computation. It has enabled scaling model parameters to the trillion-level (Yang et al., 2021; Lin et al., 2021; Fedus et al., 2022; Nie et al., 2022), showcasing the superior performance compared to dense counterparts (Fedus et al., 2022; Hwang et al., 2023; Rasley et al., 2020).

Efficient parallelization of MoE models requires assigning distinct experts to separate accelerator devices (Lepikhin et al., 2020). Yet, distributing input samples to these scattered experts demands resource-intensive all-to-all communication (Fig. 1). High communication volume in all-to-all operations significantly hampers the training speed of MoE models (up to 40% of training time).

For non-MoE models, communication scheduling (Jayaraman et al., 2019; Peng et al., 2019) is an effective way to

\*Work done during internship at AWS. †Work done while at AWS. <sup>1</sup>The University of Hong Kong, Hong Kong <sup>2</sup>Amazon Web Services, USA <sup>3</sup>Boson AI, USA. Correspondence to: Chenyu Jiang <jchenyu@connect.hku.hk>.

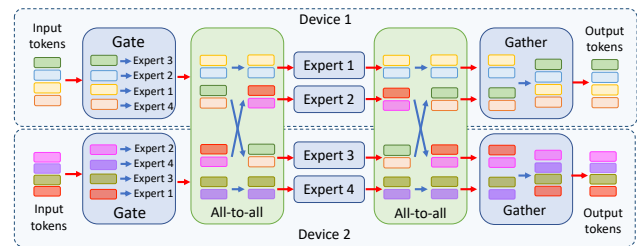


Figure 1. An example MoE layer with 4 experts scattered on 2 devices. Assume top-1 gating is used. Blue (green) boxes represent computation (communication) operators. Data dependency between operators are highlighted by red arrows. The Gate assigns each input token to an expert. All-to-alls fetch expert input/output from other devices. Gather restores the received tokens back to their original order, matching the input sequence.

overlap the communication (for synchronizing model parameters) and backward propagation. However, they are inapplicable for MoE models, which have a direct data dependency between all-to-all and other computations (experts and non-MoE computation like self-attention), as in Fig. 1. For MoE models, existing studies (Hwang et al., 2023; He et al., 2022; Wang et al., 2022; Li et al., 2023b) focused on alleviating this problem by partitioning operators into finer-grained ones and overlapping communication with computation between different partitions. Nonetheless, their focus region is limited to encompass only the all-to-all communication and expert computation. In this paper, we define the focus region as the subset of operators within the training graph responsible for concurrent (overlapping)

computation and communication. We have observed that the all-to-all communication time is usually much longer than expert computation time, thus the overall execution time is still bounded by the all-to-all communication despite overlapping (Fig. 2). The small focus region considered in existing works limits the overlapping possibilities and thus results in the sub-optimal performance.

In this paper, we extend the focus region to the whole training graph and identify two more types of operators to overlap: 1) weight gradient computation in backward pass, which does not depend on all-to-all communication and thus is able to overlap with it directly. 2) non-MoE model computation in forward pass, which has dependency with all-to-all but can perform overlapping by properly partitioning. However, extending the focus region also raises new challenges: 1) Extending overlapping to non-MoE computation requires partitioning along batch dimension. A direct partition may cause mathematical in-equivalency since the routing decision of many gating methods can be affected by batch size. 2) Partitioning introduces more smaller operators, thereby incurring GPU kernel launching overhead and under-utilization of streaming multiprocessors. Over-partitioning computations can lead to excessive overhead, negating the benefits of overlapping. Conversely, insufficient partitioning hinders the full utilization of potential overlap with all-to-all communication. Additionally, gating methods limit the types of operators that can be partitioned. Hence, extending the focus region introduces complexity in establishing the best partitioning range, which refers to the number of computation operators preceding and succeeding an all-to-all communication operator that need to be partitioned (and overlapped with the all-to-all communication).

To overcome those challenges, we propose Lancet, a system designed to enhance the throughput of MoE model training by considering the entire training graph as focus region. Lancet leverages a compiler-based approach, providing us with increased flexibility for controlling operator partitioning and scheduling. Distinct mechanisms are applied for the forward and backward passes during training. In the forward pass, where nearly all computations rely on all-to-all dependencies, it becomes necessary to partition both computation and all-to-all operators to achieve efficient overlaps. In the backward pass, we employ scheduling to ensure the weight gradient computation overlaps with all-to-all operations. The rationale behind this approach lies in the backward pass, where there are an ample number of weight gradient computation operators that can be scheduled to enable near-complete overlap with all-to-all operations. As a result, there is no need to explore partitioning solutions, as is required in the forward pass. The method we designed to overlap all-to-all with entire training graph does not conflict with non-MoE model communication scheduling strategies (Jayarajan et al., 2019; Peng et al., 2019). And

all transformations (scheduling and partitioning) maintain mathematical equivalence (i.e., the model accuracy remains unaffected by the optimizations) and can be kept transparent to users.

In summary, our contributions include:

- ▷ For the first time, we expand the focus region to encompass the entire training graph, mitigating the prolonged all-to-all communication’s impact on MoE model training. This extension enables us to discover new operators that can be overlapped with all-to-all communication.
- ▷ We adopt a greedy algorithm to schedule each weight gradient computation operator to overlap with the appropriate all-to-all.
- ▷ We devise a partitioning scheme for MoE layers that allows for the extension of partitioning to non-MoE computations while maintaining mathematical equivalency.
- ▷ We apply a dynamic programming based algorithm to identify the optimal range of non-MoE computation for partitioning and overlapping.

Comprehensive evaluations demonstrate that Lancet can decrease non-overlapping communication time by as much as 77% and deliver an up to 1.3x end-to-end speedup when compared to state-of-the-art solutions, including DeepSpeed (Rasley et al., 2020) and Tutel (Hwang et al., 2023).

## 2 BACKGROUND AND MOTIVATION

### 2.1 Mixture of Experts (MoE)

Most MoE models (Shazeer et al., 2017; Lepikhin et al., 2020) replace the feed-forward module in every two Transformer layers with multiple independent sub-networks (experts), each activated by a subset of input data. Different experts are usually placed on distinct devices for efficient parallelization. In this work, we assume non-MoE parts of the model are replicated across the devices while receiving different partitions of the training data (i.e., data parallelism). The assignment (routing) of inputs to experts is decided by a gating function at runtime.

**Expert-parallelism** Once the gating function determines expert assignments, all-to-all communication transmits inputs to the respective devices. After expert processing, another all-to-all operation sends their output back to the original devices. Due to dynamic expert assignment at runtime, token distribution among experts varies. To maintain static tensor shapes (essential for certain frameworks/hardware like XLA/TPU (Lepikhin et al., 2020)) and ensure balanced computation across experts, a common approach is to restrict the maximum tokens assigned to each expert (expert capacity,  $C$ ) on each device (Lepikhin et al., 2020; Fedus et al., 2022). Any excess tokens assigned to an expert are

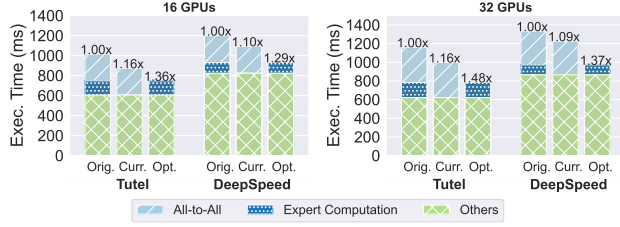


Figure 2. Breakdown of execution time when running a GPT-2 model with MoE layers using Tutel and DeepSpeed on Amazon EC2 p3dn instances. *Orig.*: unoptimized execution time. *Curr.*: performance upper-bound when optimized using current overlapping methods (expert computation completely hidden by all-to-all). *Opt.*: ideal execution time (all-to-all fully overlapped by computation).

discarded, while experts receiving fewer than  $C$  tokens are zero padded.

**Routing algorithms** Routing is often computed by assigning a gating score for each expert using a trainable linear layer, and choosing  $k$  experts with highest scores (top-k routing) (Lepikhin et al., 2020; Fedus et al., 2022). Recent works also propose other routing methods, such as hash-based (Roller et al., 2021) or random expert assignment (Zuo et al., 2022; Chen et al., 2023) and expert-choice routing (Zhou et al., 2022) (experts choose top-k tokens with highest scores). Routing algorithms significantly influence MoE model training, impacting expert balance (Zhou et al., 2022), communication costs (He et al., 2022) and more. The upcoming sections also demonstrate how routing algorithms affect feasible optimizations.

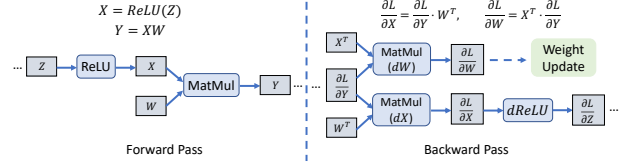
## 2.2 Overlapping all-to-all and experts

Existing techniques aim to mitigate long latency all-to-all impacts on training by overlapping it with expert computation (Hwang et al., 2023; He et al., 2022). This involves partitioning all-to-all and experts along the capacity dimension and forming a communication-computation pipeline with (only) all-to-all and experts (Fig. 4b). As shown in Fig. 2, we often observe the all-to-all time significantly surpasses that of the experts (up to 3.36x). Therefore, these techniques can only conceal the execution time of experts, while the bottleneck execution time for all-to-all communication remains unaffected.

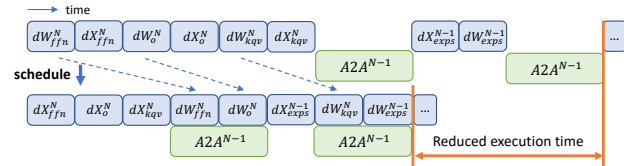
## 2.3 Opportunities and Challenges

By extending the focus region to the whole training graph, we identify more opportunities that can overlap with all-to-all communications.

**Opportunity 1: Weight gradient computation.** Computation of the weight gradient in layer  $N$ , which is essential



(a) Forward and backward pass of  $Z = \text{ReLU}(X)$ ;  $Y = XW$ .  $dX$ : activation gradient computation;  $dW$ : weight gradient computation. Gray block: tensors; Blue block: operators. Note that  $d\text{ReLU}$  (and further back-propagation of  $\frac{\partial L}{\partial Z}$ , if needed) does not depend on  $dW$ .



(b) Overlapping all-to-all communication by scheduling weight gradient computation ( $dW$ ). Superscripts on operators indicate the layer number, with the first layer after the MoE layer (during forward) numbered layer  $N$ . Subscripts indicate the type of operators:  $\text{ffn}$ : non-expert feed-forward;  $\text{kv}$ : output projection in self-attention;  $\text{kv}$ : key, query and value projection;  $\text{exps}$ : experts. Other operators are ignored. While the figure shows overlapping all-to-all in layer  $N - 1$  with  $dW$ 's in layer  $N$ , in general the all-to-all can be overlapped with any  $dW$ 's in layer  $N + k$ ,  $k \geq 0$ .

Figure 3. Scheduling weight gradient computation to overlap with all-to-all.

for updating model parameters, is independent of all-to-all communication of previous layers  $N - 1$ ,  $N - 2$ ,  $\dots$ , 1 in the backward pass. Consequently, it can be scheduled to overlap with all-to-all, allowing for flexibility in optimizing the training time (Fig. 3).

**Opportunity 2: Non-MoE computation.** In existing systems (e.g., (Hwang et al., 2023; He et al., 2022)), computation before (e.g., self-attention) and after (e.g., the following Transformer layer) the MoE layer does not overlap with all-to-all communication since their limited focus region. Nevertheless, if we partition non-MoE computations and integrate them into the computation-communication pipeline, we can create additional opportunities to overlap operations with the all-to-all communication (Fig. 4c, 4d).

**Challenge 1: How to perform mathematically equivalent partition.** Consider feeding a tensor with dimension  $B \times S$  ( $B$  as batch size,  $S$  as sequence length) into an MoE layer (in Fig. 5). The tokens are re-arranged according to their target experts, undergoing an all-to-all with shape  $E \times C$  ( $E$  as the total number of experts,  $C$  as the expert capacity) for distributing to the corresponding device. Each expert processes the  $C$  received tokens, followed by a reciprocal all-to-all (not shown in the figure). Reverting tokens to their original order yields the MoE layer's output. The existing

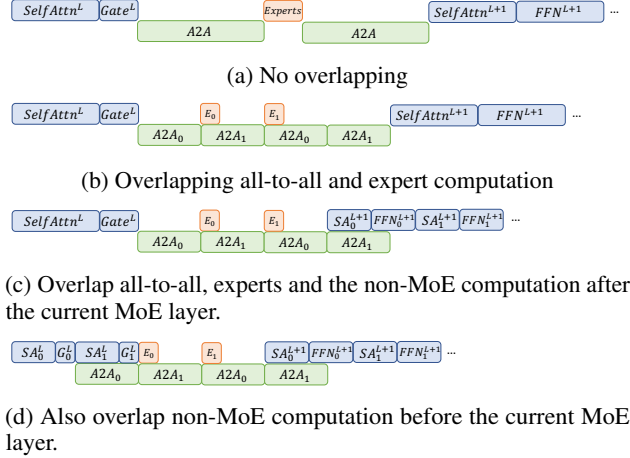


Figure 4. Performance gain of different overlapping types.

methods all partition the all-to-all and experts at the capacity dimension ( $C$ ), and thus the tokens in same partition appear in irregular locations in the re-arranged MoE output, e.g., belong to different sequences across the batch (Fig. 5a). Therefore, the following computation must wait until all partitions finish execution, interrupting the pipeline.

In order to overlap forward pass non-MoE computation with all-to-all, the MoE input and output must be partitioned along the batch dimension. However, directly partitioning the input (micro-batching) may result in extra token dropping since the expert capacity also drops accordingly. For example, consider a input batch (with corresponding expert capacity  $C$ ) partitioned into two micro-batches, each processed with expert capacity  $\frac{1}{2}C$ . Assume the first micro-batch contains  $\frac{3}{4}C$  tokens for an expert, and the second contains  $\frac{1}{4}C$  tokens for that expert. If not partitioned, then all tokens can fit into expert capacity  $C$  thus no token will be dropped. When directly partitioned,  $\frac{1}{4}C$  tokens will be dropped from the first micro-batch since it now only has expert capacity  $\frac{1}{2}C$  (Fig. 5b). Such change in mathematical equivalency is undesirable as it may affect model performance.

To avoid this effect, we implement special gating operators that pass capacity information between partitions (e.g., when the first partition (micro-batch) uses  $\frac{3}{4}C$  capacity, the second partition will adjust its remaining capacity to  $\frac{1}{4}C$ ), preserving the exact token-to-expert mapping and token dropping as the un-partitioned case. This however implies that any partition can send any amount of token (ranging from 0 to  $C$ ) to an expert (while tokens sent from all partitions add up to  $C$ ) (Fig. 5c). We implement irregular-shaped all-to-all to efficiently handle such a dynamic communication pattern (details discussed in Sec. 6).

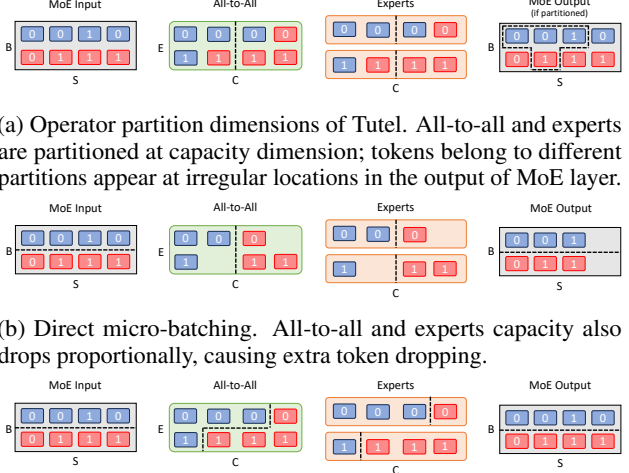
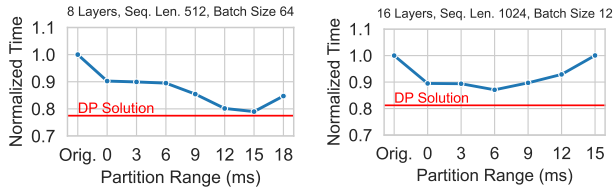


Figure 5. Operator partitioning scheme in an MoE layer. Number in each token shows their assigned expert. Tokens of the same color belong to the same sequence.

Figure 5. Operator partitioning scheme in an MoE layer. Number in each token shows their assigned expert. Tokens of the same color belong to the same sequence.

**Challenge 2: How to determine the optimal partition range for non-MoE operators.** GPU kernel launches involve startup overhead (Rotem et al., 2018), which occurs with each launch. Partitioned computation operators deal with smaller input tensors, potentially leading to GPU core under-utilization. Similarly, smaller communication operators might not fully utilize network bandwidth. So it is not always optimal to partition the entire Transformer layer before and after MoE layer, which may increase training time due to partition overheads. Fig. 6 shows this phenomena. The optimal partition range (a set of computation ops around all-to-all) depends on model specification, input size, the underlying computation power and also network bandwidth. Our extension of the focus region to whole training graph makes the decision more challenging.

Furthermore, the gating methods limit the partitioning opportunities. Some gating methods assign target experts based on the information calculated over the entire batch of tokens. For example, Batch-prioritized Routing (Riquelme et al., 2021) sorts tokens in a batch first by their “importance score” (the sum of top-k largest gating scores) and then assigns tokens to experts. So tokens with lower scores would be dropped first. Splitting along batch dimension would thus cause differences in token dropping. For such gating methods, we can only extend partitioning after the MoE layer (Fig. 4c). For other gating methods whose expert assignment can be decided from partial batches (e.g., Switch (Fedus et al., 2022) or Random (Zuo et al., 2022; Chen et al., 2023) gating), we can extend partitioning to both after and before the MoE layer (Fig. 4d).



(a) Less layers, large batch size (b) More layers, small batch size

Figure 6. Effect of partition range on GPT-2 MoE model forward time on 16 A100 GPUs (32 experts). X axis shows how many ops (measured in their execution time) before and after the MoE layer is included in the partition. *Orig.*: no partitioning. *0*: only partition all-to-all and experts (as in Tutel).

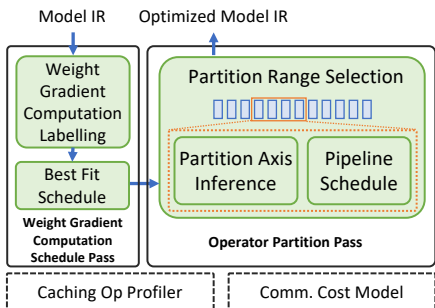


Figure 7. Overview of Lancet modules.

### 3 LANCET OVERVIEW

To address the opportunities and challenges mentioned above, we devised Lancet, a compiler-based solution designed to optimize MoE model training. The key advantage of this approach is the explicit extraction of a model’s computation and communication through compilers’ intermediate representation (IR), granting precise control over operator execution and simplifying the implementation of weight gradient computation scheduling and the partition of operators.

Lancet adopts compiler passes to enhance MoE model training and resolve associated issues. At a higher level, it encompasses two primary optimization passes: weight gradient computation scheduling and operator partitioning, modifying the backward and forward pass of the model, respectively. Fig. 7 gives an overview of Lancet.

1) **Weight Gradient Computation Schedule Pass (§4)** takes the model IR as input, which is a sequence of instructions, and re-orders the instructions corresponding to weight gradient computation operators to overlap with all-to-alls during backward propagation. Dependency analysis is first performed to identify the weight gradient computation instructions that can be overlapped with each all-to-all (§4.1). Then for each all-to-all op, we employ a best-fit greedy algorithm to choose a set of weight gradient computation ops with comparable total execution time to maximize overlap

(§4.2).

2) **Operator Partition Pass (§5)** receives the IR with weight gradient computation scheduled and further optimizes the all-to-alls in the forward pass through partitioning and pipelining. A dynamic programming algorithm is employed to find the optimal partition range for non-MoE ops (§5.1). During this process, a partition axis inferencer (§5.2) employs a constraint programming algorithm to deduce the partition axis for each instruction, facilitating partitioning of IR. Then, a pipeline scheduler (§5.3) estimates the cost of the resulting computation-communication pipeline, guiding the dynamic programming algorithm.

These optimizations are supported by a **Caching Op Profiler**, which profiles and caches the execution time of all ops in the model IR. Profiling is done once for each (partitioned) operation with the same shape; the cached execution time can be subsequently reused. Communication costs (e.g., partitioned all-to-all) are estimated by a **Communication Cost Model**. The communication cost model is built by profiling communication operations across various input sizes (e.g., 1KB, 2KB, 4KB, ..., up to the maximum possible communication used in models), and the cost is linearly interpolated among these points. Since Lancet uses irregular-shaped all-to-alls (Fig. 10), their execution time depend on the combination of actual amount of data to be communicated, which is not known at compilation time. Therefore, we resort to a static-shape approximation: the cost of an  $n$ -partitioned all-to-all with original capacity  $C$  is obtained by querying the profiled (uniform-shaped) cost model at capacity  $C/n$ . We observe that such approximation suffices to produce a good prediction of overall iteration time during our experiments (Fig. 14).

### 4 WEIGHT GRADIENT COMPUTATION SCHEDULE PASS

The weight gradient computation schedule pass takes the model IR describing the training iteration as input and re-orders the instructions to overlap weight gradient computation with all-to-alls. The IR is represented as a sequence of instructions  $\mathcal{I} = [I_1, I_2, \dots, I_N]$ ; each instruction is characterized by its input tensors  $\mathbf{x}$ , output tensors  $\mathbf{y}$ , and operator  $f$ :  $I_n = (\mathbf{x}^n, \mathbf{y}^n, f^n)$ , representing the operation  $y_1^n, y_2^n, \dots, y_{|\mathbf{y}^n|}^n = f^n(x_1^n, x_2^n, \dots, x_{|\mathbf{x}^n|}^n)$ .

#### 4.1 Weight Gradient Computation Labelling

Due to the fine-grained nature of instructions, identifying weight gradient computation instructions becomes challenging. While there is no direct dependency between weight gradient computation and all-to-alls, the weight gradient computations must adhere to the constraints imposed by the chain rule, which imposes scheduling restrictions. There-

fore, we first identify the set of weight gradient computation instructions that can be overlapped with each all-to-all by analyzing instruction dependencies. Consider a dependency graph  $\mathbf{G} = (\mathcal{I}, \mathcal{E})$ , where each directed edge  $E_{i,j}$  asserts that  $I_j$  depends on  $I_i$ , i.e.,  $I_j$  consumes the output of  $I_i$  thus must be executed after it. Then, a weight gradient computation instruction  $I_i$  can be overlapped with an all-to-all instruction  $I_a$  if and only if there is no directed path between  $I_i$  and  $I_a$  in  $\mathbf{G}$ . Such paths can be discovered by a simple Depth- or Breadth-first Search algorithm. For each all-to-all instruction  $I_a$ , we compute the set of instructions that can overlap with it as  $\mathbf{W}^{I_a}$ , which is used in the scheduling algorithm.

## 4.2 Weight Gradient Computation Scheduling

We then optimize the scheduling of labelled weight gradient computation operations to minimize the overall training time. Determining the schedule of weight gradient computation is equivalent to deciding an assignment of each weight gradient computation operator to an all-to-all with which it will overlap. Let  $\mathcal{I}^W$  be the sub-sequence of  $\mathcal{I}$  containing all weight gradient computation instructions, and  $\mathcal{I}^a$  be the sub-sequence containing all all-to-alls. Let variable  $x_{i,j} = 1$  if  $I_i^W$  (the  $i$ th weight gradient computation instruction) is assigned to  $I_j^a$  (the  $j$ th all-to-all), and otherwise  $x_{i,j} = 0$ . The execution time of  $I_i^W$  ( $I_i^a$ ) is  $t_i^W$  ( $t_i^a$ ). Then maximization of total overlapped all-to-all execution time can be formulated as the following integer program:

$$\begin{aligned} \max_{\mathbf{x}} \quad & \sum_{j=1}^{|\mathcal{I}^a|} \min\{t_j^a, \sum_{i=1}^{|\mathcal{I}^W|} t_i^W \cdot x_{i,j}\} \\ \text{s.t.} \quad & \sum_{j=1}^{|\mathcal{I}^a|} x_{i,j} \leq 1, \quad \forall i \in [1, |\mathcal{I}^W|] \quad (1) \\ & x_{i,j} = 0, \quad \forall I_i^W \notin \mathbf{W}^{I_j^a} \quad (2) \end{aligned}$$

The  $\min\{t_j^a, \sum_{i=1}^{|\mathcal{I}^W|} t_i^W \cdot x_{i,j}\}$  gives the amount of overlapped time in each all-to-all. Constraint (1) states that each weight gradient computation instruction can only be used to overlap with at most one all-to-all. (2) restricts the assignment based on instruction dependency calculated during weight gradient computation labelling.

Such a problem is a generalized assignment problem (GAP) with non-linear objective and additional constraints (2). Since GAP is already known to be NP-hard (Martello & Toth, 1990), we resort to a greedy heuristic. We sequentially iterate through  $\mathcal{I}^a$ : for each  $I_i^a$ , weight gradient computation instructions are greedily chosen from  $\mathbf{W}^{I_i^a}$ , that are not already used to overlap with other all-to-alls and minimize the absolute difference between the all-to-all execution time and sum of all weight gradient computation to be overlapped with it. We proceed to the next all-to-all when the current

---

### Algorithm 1 Weight Gradient Computation Schedule Pass

---

**Input:**  $\mathcal{I}$  - a sequence of instructions

**Output:**  $\mathcal{I}'$  - scheduled instructions

```

1:  $\mathbf{G} \leftarrow \text{CreateDependencyGraph}(\mathcal{I})$ 
2: /* Weight gradient computation labelling */
3:  $\mathcal{I}^a \leftarrow [I_i \in \mathcal{I} \mid f^i \text{ is all-to-all}]$ 
4:  $\mathbf{W}^{I_i} \leftarrow \{\}$  for each  $I_i \in \mathcal{I}^a$ 
5: for  $I_i \in \mathcal{I}^a, I_j \in \mathcal{I}, I_i \neq I_j$  do
6:   if no directed path between  $I_i$  and  $I_j$  then
7:      $\mathbf{W}^{I_i}.\text{insert}(I_j)$ 
8:   end if
9: end for
10: /* Weight gradient computation scheduling */
11:  $\mathbf{t}^a, \mathbf{t}^W \leftarrow \text{GetInstrExecTime}(\mathcal{I})$ 
12:  $\mathbf{W}^{\text{used}} \leftarrow \{\}$ 
13:  $\mathbf{Asg} \leftarrow \{\}$  /* map recording the assignment results */
14: for  $i \in [1, |\mathcal{I}^a|]$  do
15:    $t_u \leftarrow \mathbf{t}_i^a$  /* unoverlapped time of all-to-all i */
16:   while  $t_u > 0$  and  $\mathbf{W}^{I_i^a} \cap (\mathcal{I} - \mathbf{W}^{\text{used}}) \neq \emptyset$  do
17:     /* Find available instr that best matches  $t_u$  */
18:      $j_{\min} \leftarrow \text{argmin}_j \{|t_u - t_j^W| \mid I_j^W \in \mathbf{W}^{I_i^a}, I_j^W \notin \mathbf{W}^{\text{used}}\}$ 
19:      $t_u \leftarrow t_u - t_{j_{\min}}^W$ 
20:      $\mathbf{W}^{\text{used}}.\text{insert}(I_{j_{\min}}^W)$ 
21:      $\mathbf{Asg}.\text{insert}(\{I_{j_{\min}}^W : I_i^a\})$ 
22:   end while
23: end for
24:  $\mathcal{I}' \leftarrow \text{ReorderInstrs}(\mathbf{Asg})$ 

```

---

one is fully overlapped.

After deciding the assignment of weight gradient computation, we reorder the instructions, placing them right after their overlapping all-to-all instructions. This ensures the weight gradient computation start execution immediately following the launch of all-to-all communication. Alg. 1 presents the entire weight gradient computation scheduling process.

## 5 OPERATOR PARTITION PASS

With the scheduled instructions from the weight gradient computation scheduling pass, we next hide all-to-alls in the forward pass through extensive operator partitioning.

### 5.1 Partition Range Selection

Selecting a proper range of non-MoE computation to partition is crucial to maximize overlap and minimize overheads. Different gating functions affect the type of non-MoE operators we can partition (only ops after the MoE layer, or both before and after the MoE layer). The number of partitions (how many parts each operator is partitioned into) also affects model performance. We introduce a dynamic programming-based algorithm to optimize the aforementioned decisions.

Given an instruction sequence (for forward pass)  $\mathcal{I} = [I_1, I_2, \dots, I_N]$ , let  $T(n)$  denote the end-to-end execution time (after considering overlapping) of instructions 1 to  $n$  when we have optimally partitioned these instructions. Then, we have

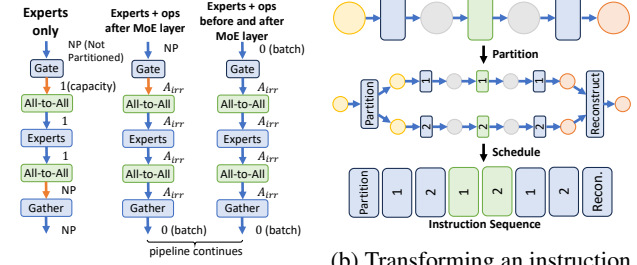
$$T(n) = \min_{1 \leq i \leq n-1} \{T(i) + \min_{1 \leq k \leq K} P(i, n, k)\}$$

where  $P(i, n, k)$  is the end-to-end execution time of instructions  $i$  to  $n$  if they are partitioned into  $k$  parts and their computation and communication components (all-to-all) arranged to overlap with each other.  $K$  is the maximum allowed number of partitions which is a hyper-parameter. If there is no valid way to partition instructions  $i$  to  $n$  (e.g., if unsupported gating function is included), we let  $P(i, n, k) = \infty$ . The partition axes (dimensions through which the input and output tensors will be split) of each instruction  $i, \dots, n$  is determined by our *partition axis inferrer*. Then the partitioned instructions are scheduled to overlap each other (form a computation-communication pipeline) by our *pipeline scheduler*, which reports the end-to-end execution time after partition as  $P(i, n, k)$ . The optimal end-to-end execution of the entire forward pass of the model is thus  $T(N)$ .

The dynamic programming algorithm requires  $O(N^2K)$  evaluations of cost  $P(i, n, k)$  in total, since there are  $N$   $T(n)$ s to evaluate and each  $T(n)$  requires  $K(n-1)$  evaluations of  $P(i, n, k)$ . In practice, the number of partitions  $k$  is limited by the size of the partitioned dimension (e.g., if we are partitioning along the batch dimension and the batch size is 4, we can have at most 4 partitions). Partition overhead also limits very fine-grained partitioning (in our experiments, we never observed the optimal number of partitions exceeding 4). Therefore,  $K$  can be safely set to a relatively small value (e.g., 8). To further reduce optimization time, we group several consecutive instructions together based on execution time (e.g., total execution time sum up to 2ms) and perform dynamic programming on these groups instead. Due to partition overhead, an optimal communication-computation pipeline would likely be not very long. Therefore we can also limit the range of  $i$  (i.e., set a maximum length limit on the partition ranges). Suppose there are  $N'$  instruction groups in total and the maximum partition range is  $G$  groups. Then the algorithm requires an  $O(N'GK)$  number of  $P(i, n, k)$  evaluations in total.

## 5.2 Partition Axis Inference

To identify the partition axis of each instruction's input and output, we formulate a constraint satisfaction problem. For the  $n$ th instruction in the input sequence, let  $a_{x_i}^n$  represent the partition axis of its  $i$ th input, and  $a_{y_i}^n$  for its  $i$ th output. For each different operator ( $f$  in the instructions), we define function  $F_{\mathcal{Z}}^f : \mathbf{a}_x \times \mathbf{a}_y \mapsto \mathcal{Z}$  which takes the input and out-



(a) Partition axis of data tensors in different partition types. Orange arrow indicates pipeline input (output) tensors. Blue begin and end locations, where extra partition/reconstruction instructions are needed. (b) Transforming an instruction sequence to form a pipeline. Yellow (orange) circles denote input (output) tensors. Blue rectangle: computation instruction; Green rectangle: communication instructions are needed.

Figure 8. Operator partitioning.

put axes of an instruction  $I$  as input and returns a constraint  $\mathcal{Z}^I$  (a boolean expression). Such a constraint specifies how the input and output axes of the instruction should relate for a valid partition (i.e., the original output can be reconstructed from the partitioned ones). Take matrix multiplication  $Y = X \cdot W$  as an example: we can split  $X$  along the row (1st) dimension and not change  $W$ , resulting in  $Y$  partitioned in the row axis ( $\begin{bmatrix} X_1 \\ X_2 \end{bmatrix} W = \begin{bmatrix} X_1 W \\ X_2 W \end{bmatrix}$ ); or we can keep  $X$  and split  $W$  along the column (2nd) axis, partitioning  $Y$  in the column axis ( $X \begin{bmatrix} W_1 & W_2 \end{bmatrix} = \begin{bmatrix} X W_1 & X W_2 \end{bmatrix}$ ). To capture the above possible partition axes combinations, we have the following constraint:

$$(a_{x_1} = 0 \wedge a_{x_2} = -1 \wedge a_{y_1} = 0) \vee (a_{x_1} = -1 \wedge a_{x_2} = 1 \wedge a_{y_1} = 1)$$

$a_{x_1}, a_{x_2}, a_{y_1}$  are partition axes for  $X$  (the 1st input),  $W$  (the 2nd input) and  $Y$  (the 1st output) respectively (dimension index starting from 0; -1 means not partitioned). We also introduce a special partition axis  $A_{irr}$  for each MoE-related operator, to represent the irregular partition of all-to-all and experts in our extended computation-communication pipeline (Fig. 5c). The constraints for all-to-alls and experts are written to accept partition at capacity axis if the partition range ( $i, n$ ) only covers the all-to-all and experts, and  $A_{irr}$ , otherwise. Correspondingly, we write  $F_{\mathcal{Z}}$  of the MoE gather operator to only allow its input to be partitioned at  $A_{irr}$  but not the capacity axis, and  $F_{\mathcal{Z}}$  of the gating function (if it can be partitioned) to allow batch-partitioned inputs and generate  $A_{irr}$  partitioned outputs (Fig. 8a).

If the constraints of all instructions are satisfied, every original tensor can be reconstructed from the partitioned ones, asserting correctness of the partition. We also require that the partition axes of the same tensor cannot be changed, since switching the partition axes requires data from other partitions thus interrupting the computation-communication pipeline. Putting the above together, we have the following constraint satisfaction problem:

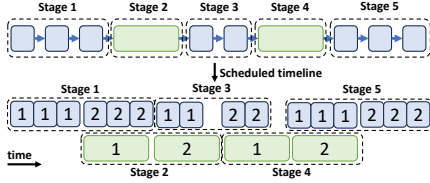


Figure 9. Pipeline schedule by stages. Blue rectangle: computation instructions; Green: communication. Numbers indicate partition index (i.e., the  $n$ th partition).

find  $\mathbf{a}$

$$\begin{aligned} \text{s.t. } & F_{\mathcal{Z}}^{f^i}(\mathbf{a}_{\mathbf{x}}^i, \mathbf{a}_{\mathbf{y}}^i) = 1, \quad \forall i \in [1, N] \\ & a_{y_j}^i = a_{x_l}^k, \quad \forall (i, j, k, l) \in \mathcal{D} \end{aligned}$$

where  $\mathcal{D}$  describes tensor dependency between operators: indices  $(i, j, k, l) \in \mathcal{D}$  if the  $j$ th output of instruction  $i$  is fed to the  $l$ th input of instruction  $k$ .

Solving the problem (e.g., using an off-the-shelf solver like OR-Tools (Perron & Furnon, 2019)) gives the partition axes of all input/output tensors of the partitioned operators. The corresponding partitioned instructions are generated and sent to the pipeline scheduler (Fig. 8b).

### 5.3 Pipeline Scheduling

To organize the partitioned instructions into a computation-communication pipeline, the instructions in each partition are divided into stages. Each stage contains all computation or communication that can be consecutively executed. Instructions in each stage of a partition are always scheduled together. Within each stage, instructions from the different partition are ordered by partition index (e.g., the first partition always gets scheduled first, and then the second and so on). The resulting schedule is demonstrated in Fig. 9.

To obtain the end-to-end (pipelined) execution time of partitioned operators  $P(i, n, k)$ , we simulate the execution timeline by calculating the start and end time of each instruction relative to pipeline start. Specifically, each instruction’s start time is the maximum over (i) the end time of all instructions that it depends on and (ii) the end time of the previous computation/communication instruction (of the same type) in the scheduled order.  $P(i, n, k)$  is thus the end time of the last instruction, which is reported back to guide the dynamic programming procedure.

## 6 IMPLEMENTATION

Lancet is generally applicable to any deep learning compiler for training. We adopt RAF (Yu et al., 2023), an open-source compiler extended from Apache TVM (Chen et al., 2018), as our underlying compiler, which provides a comprehensive compilation of DL models. We implement Lancet with

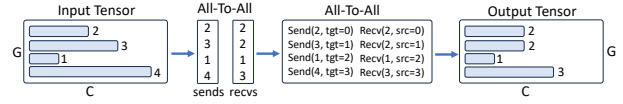


Figure 10. Implementation of irregular all-to-all.  $G$ : number of GPUs participating in the all-to-all,  $G = E/E_l$  ( $E_l$ : the number of experts per GPU). On each device, an input and output buffer of fixed shape ( $G \times C$ ) is allocated. Number in the Input/Output Tensors indicate the actual size of the data to be sent/received on the GPU. The first All-to-All communicates the data sizes to be exchanged; the second All-to-All communicates the actual data. Send/Recv( $\mathbf{x}$ , tgt/src= $\mathbf{y}$ ) indicates an NCCL send/recv primitive that sends/receives a data chunk of size  $\mathbf{x}$  to/from  $\mathbf{y}$ .

13K LoC in C++. Communication primitives such as all-to-all are implemented based on NCCL (NVIDIA, 2021). Lancet also implements partition constraints ( $F_{\mathcal{Z}}$ ) for all computation operators in common Transformer-based models. The MoE dispatching ops are implemented based on Tutel’s (Hwang et al., 2023) kernel.

Since Lancet is fully implemented in two optimization passes as IR transformations, users only need to enable them in RAF’s optimization pass manager, without any modification to the existing code-base. The three hyper-parameters for speeding up the optimization process (i.e.,  $\rho$ , the maximum number of partitions;  $\gamma$ , the group size;  $\iota$ , maximum partition range in dynamic programming) can be set through environment variables.

**Irregular all-to-all** (all-to-allv in MPI (Message Passing Interface Forum, 2021) terminology) sends different amounts of data to different target devices. In MoE layers, the amount of data to send to each device depends on the gating function and is only known at runtime (Fig. 5c). To implement such dynamic communication scheme in a static-shaped system like Lancet, we allocate the input and output tensors based on the maximum amount of data to be sent (i.e., capacity of each expert). As shown in Fig. 10, at runtime, the input buffer is only partially filled based on the result of the gating function. A first all-to-all is performed to exchange the amount of data to be sent and received across devices, followed by a second all-to-all only sending and receiving the required amount of data. The all-to-alls are implemented via a grouped NCCL communication consisting of NCCLSends and NCCLRecvs.

## 7 EVALUATION

**Experiment Setup** We evaluate Lancet on an Amazon EC2 p4de.24xlarge cluster and a p3dn.24xlarge cluster, each with 8 nodes. Each p4de node has 8 NVIDIA A100 80GB GPUs and 4x100 Gbps NICs. Each p3dn node has 8 NVIDIA V100 GPUs and one 100 Gbps NIC. We refer to the cluster of p4de.24xlarge and



p3dn.24xlarge nodes as A100 and V100 respectively, for the rest of the paper. All nodes run in the same docker environment where we used Ubuntu 20.06 with CUDA 11.3 and NCCL 2.12.12 with PXN enabled.

**Benchmark Models and Datasets** We conduct our evaluations on MoE versions of the GPT-2 (Radford et al., 2019) model (from Huggingface transformers (Wolf et al., 2020) version 4.18.0). The base models are enhanced by replacing every other Transformer block’s feed-forward layer with an MoE layer. Two variants of the model are used: the smaller model (GPT2-S-MoE) has 12 layers with hidden dimension size 768; the larger one (GPT2-L-MoE) has 24 layers with hidden size 1024. In all experiments, we scale the number of experts along with the number of GPUs: each GPU always hosts two experts. The SGD optimizer (with momentum) is used for training the model.

For all experiments, we use the WikiText (Merity et al., 2016) dataset as model inputs. We fix the input sequence length to 512 and use the largest batch size that can fit into the GPU memory for each model: on A100, we use batch size 24 per GPU for GPT2-S-MoE and 48 for GPT2-L-MoE. On V100, we use batch size 16 for GPT2-S-MoE and 8 for GPT2-L-MoE.

**Baselines** We compare Lancet’s training performance with DeepSpeed (version 0.5.8, without Tutel’s kernels) (Rasley et al., 2020) and Tutel (version 0.3) (Hwang et al., 2023). Tutel implements overlapping between all-to-all and expert computation. For each experiment with Tutel, we search through the overlapping degree (the number of partitions) of 1, 2, 4 and 8 and report the best result. Tutel and DeepSpeed are both built on PyTorch (Paszke et al., 2019), whose performance on computation ops may be different from RAF (Yu et al., 2023). Therefore, we also include results of RAF without Lancet’s modifications for comparison.

**Hyper Parameters** We set the maximum number of partitions  $\rho$  to 8, except when excessive partitions cause out-of-memory (OOM) errors. In that case, we reduce it to 4 (and 2 if still OOMs). We set the group size  $\gamma$  according to the model execution time so that there are 5 groups between each MoE layer. The maximum partition range  $\iota$  is set to be the execution time between two MoE layers, so one pipeline will be formed per MoE layer.

## 7.1 Throughput

We compare Lancet’s training throughput against baselines using different numbers of GPUs. We do weak scaling, i.e., keep the local batch size fixed at each GPU while the effective total batch size of the model scales linearly. Since gating method constraints the available pipeline range, we run the

experiments with two different gating methods: Switch (Fedus et al., 2022) gate which allows overlapping with computation both before and after the MoE layer (Fig. 4d) and Batch Prioritized (Riquelme et al., 2021) gate which only allows overlapping with computation after the MoE layer (Fig. 4c).

Fig. 11 shows that Lancet achieves up to 1.21x (1.17x on average) speed up compared to the baselines on the A100 cluster, and up to 1.3x (1.22x on average) on V100 cluster when using Switch gate. We find DeepSpeed exhibits slightly higher memory requirements than other frameworks, leading to OOM on A100 when running the GPT2-S-MoE model (OOM does not happen on V100 since a smaller batch size is used, i.e., 24 v.s. 16). When using Batch Prioritized gate (Fig. 12), we observed up to 1.24x (1.17x on average) speed up on the A100 cluster, and up to 1.24x (1.21x on average) on V100 cluster. Despite more constraint pipeline range, the achieved speed up for Batch Prioritized gate is overall similar to that of the Switch gate. This is because despite only pipelining with computation after the MoE layer, significant amount of overlapping can still happen. Our  $dW$  scheduling is also unaffected by the gating methods. The maximum achieved speed up on V100 is lower when using Batch Prioritized gate though, indicating that partitioning may have a larger impact on V100.

As shown in Fig. 13, Lancet achieves a higher level of computation-communication overlapping than baselines, reducing non-overlapped communication time by up to 69% (A100) and 83% (V100) compared to RAF, 66% (A100) and 77% (V100) compared to Tutel. The trade-off of applying partition-pipeline is also clearly shown in Fig. 13. While Lancet’s optimizations decrease the end-to-end execution time, the total execution time of computation (Non-overlapped Computation + Overlapped) ops can be higher than that of RAF, due to partition overheads. Since Lancet implements irregular all-to-alls and do not transmit any padding tokens between experts, the overall communication time (Non-overlapped Communication + Overlapped) can be lower than baselines.

## 7.2 Accuracy of cost model

Fig. 14 shows the accuracy of Lancet’s cost model, used to predict the iteration time after applying each optimization. The prediction error is very small (3.83%). Such an accurate cost model provides useful information to guide our weight gradient computation scheduling and DP-based operator partitioning algorithms.

## 7.3 Optimization Time

Fig. 15 shows the time taken to optimize the models in our experiments. Optimization time is dominated by the operator partition pass (Sec. 5) since weight gradient computation

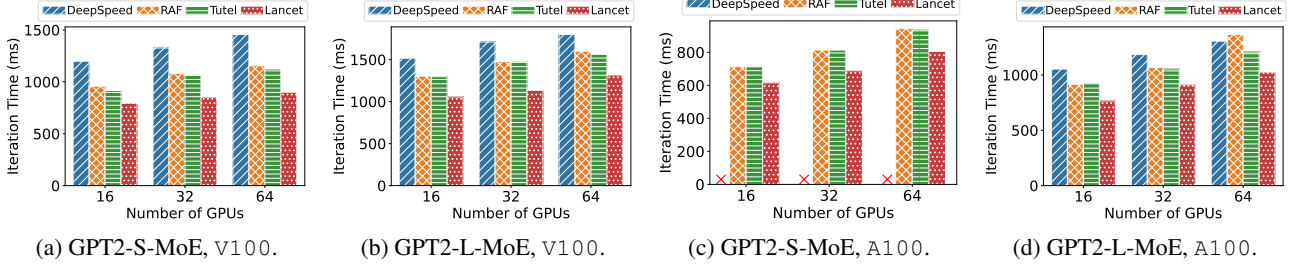


Figure 11. Training iteration time when using Switch gate. Red cross indicates out-of-memory.

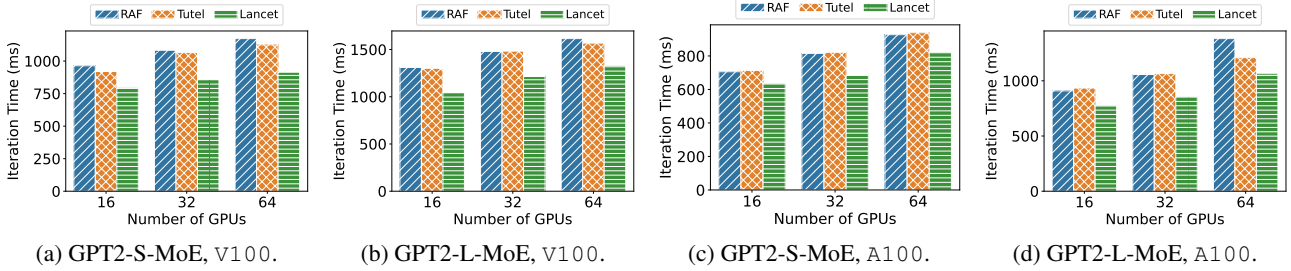


Figure 12. Training iteration time when using Batch Prioritized gate.

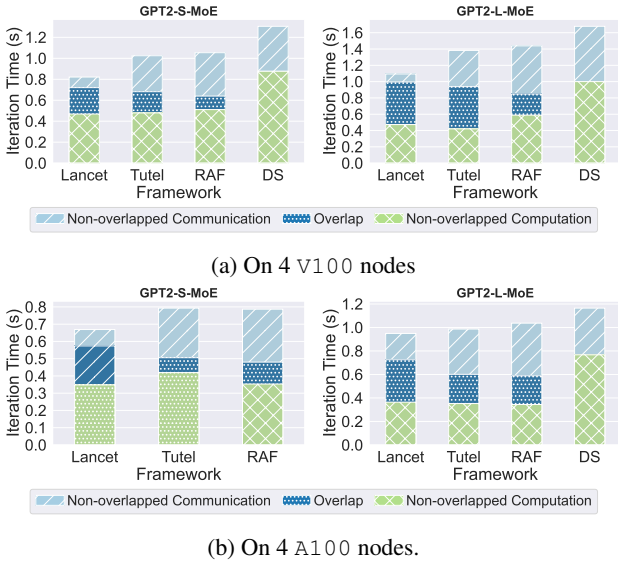


Figure 13. Iteration time decomposition. DS: DeepSpeed.

schedule (Sec. 4) uses a fast greedy algorithm. Since every device shares the same computation graph, the optimization time is less affected by the number of GPUs used and more by the number of layers in the model. The optimization time of most models bench-marked is below 20 minutes. Our optimization also only requires one GPU to run (for bench-marking execution time of partitioned computation ops).

#### 7.4 Ablation Study

To show the effects of weight gradient computation scheduling and pipelining separately, we conduct an ablation study

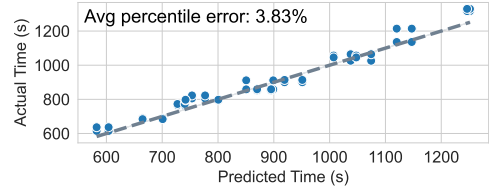


Figure 14. Prediction accuracy of Lancet’s cost model. Data aggregated from all models bench-marked on all clusters during our experiments.

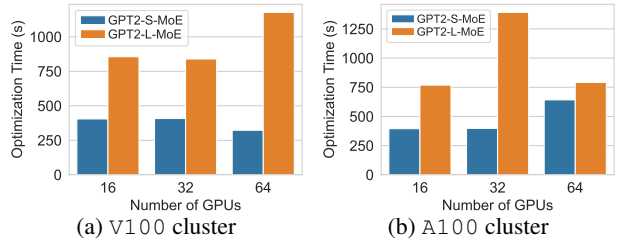


Figure 15. Lancet’s optimization time when using Switch gate.

on 4 A100 and V100 nodes. In Fig. 16, the relative speed-up is computed by dividing the training throughput under each scheme by that of RAF without any Lancet optimizations. For both models, applying only scheduling or only pipelining yields a lower speedup compared to using them together. On both clusters, GPT2-L-MoE is affected more by disabling weight gradient computation scheduling, while the two optimizations have more similar performance gain on GPT2-S-MoE. This is because GPT2-L-MoE has more parameters and layers while using a smaller batch size, thus having higher partition overheads, rendering weight gradient computation scheduling more effective compared to operator partition.

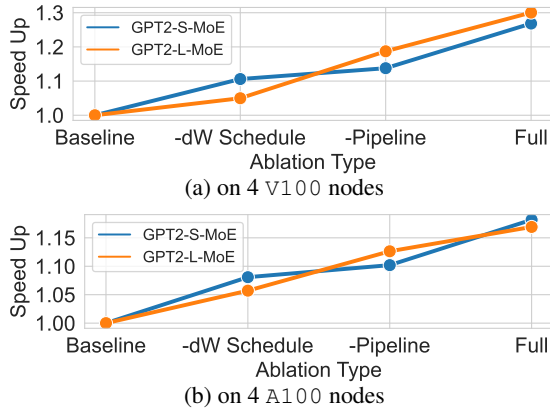


Figure 16. Ablation study on 4 A100 and V100 nodes. *dW*: weight gradient computation.

## 8 DISCUSSION AND RELATED WORKS

### Compatibility with other large-scale training techniques

While Lancet is evaluated with data and expert parallelism, the techniques are in principle compatible with most other commonly used training optimizations. Weight gradient scheduling only utilizes operator dependency during backward propagation, thus unaffected by most distributed training sharding techniques. Some techniques introduce extra communication which may interfere with partition-based all-to-all overlapping. FSDP/ZeRO3 (Rajbhandari et al., 2020) inserts additional all-gather communication in the forward passes, which may require additional scheduling to avoid interference with overlapped all-to-all. Tensor parallelism (Shoeybi et al., 2019) requires all-reduce communication after self-attention; Ring-attention (sequence parallelism) (Liu et al., 2023) communicates the key-value blocks during the attention process. If different devices or communication channels are used for expert and tensor/sequence parallelism (e.g., inter-node vs. intra-node), the overlapped all-to-all communication can be arranged to execute concurrently with tensor/sequence parallelism traffic. Investigating the efficient orchestration and overlapping of communication arising from various sharding techniques, particularly the intricate patterns generated by automatic sharding (Zheng et al., 2022), remains future work.

### Optimizing irregular communication and expert computation

Lancet’s partition produces irregular-shaped all-to-alls and expert computation. While we use a simple NCCL based implementation (Fig. 10), better communication implementations targeting such dynamic workload may further improve the performance. Similarly, the shape irregularity in expert computation may cause extra computation due to padding. Block-sparse expert kernels (e.g., MegaBlocks (Gale et al., 2023)) can be further applied to accelerate the computation.

MoE architectures that facilitate overlapping PR-MoE (Rajbhandari et al., 2022) and DeepSeek-MoE (Dai et al., 2024) use a shared expert which all tokens are routed to. The all-to-all communication (for non-shared experts) can also be overlapped with the computation of such shared expert. Lancet’s approach can be applied to a wider-range of MoE models that use traditional architectures, e.g., (Jiang et al., 2024).

### Other MoE training optimization techniques

Tutel (Hwang et al., 2023) and FasterMoE (He et al., 2022) are two popular frameworks optimizing for MoE models. Both frameworks support overlapping all-to-all and expert computation. Tutel (Hwang et al., 2023) also implements fast dispatching kernels, better all-to-all algorithm, and adaptive parallelism switching for dynamic workloads. FasterMoE (He et al., 2022) proposes techniques to handle imbalanced expert selection and to select experts based on network topology. These optimizations are orthogonal to ours and can potentially be used in conjunction. (Zhang et al., 2022) proposes to run two copies of the model on the same device, overlapping computation and communication between different model replicas. However, splitting the input among the two model replicas may result in mathematical in-equivalence (e.g., due to extra token dropping). (Li et al., 2023a) optimizes MoE training by prioritizing all-to-all traffic over all-reduce traffic, avoiding bandwidth contention and improving all-to-all latency. This method can also be used in conjunction with Lancet.

## 9 CONCLUSION

This paper presents Lancet, a system to automatically optimize MoE model training. We extend the optimization space of current methods and seek whole-training-graph-level opportunities to overlap all-to-all communication. In the forward pass, we overlap all-to-all with both expert and non-MoE computation through proper partitioning and pipelining. The optimal partition range is determined by a dynamic programming algorithm. In the backward pass, we schedule weight gradient computation to overlap all-to-all using an best-fit greedy algorithm. Experimental evaluation shows that Lancet reduces non-overlapped communication time by up to 77%, and achieves up to 1.3x end-to-end speed up compared to state-of-the-art solutions.

## 10 ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their valuable feedback. This work was supported by an Amazon Research Award (ARA) on AWS AI and grants from Hong Kong RGC under the contracts HKU 17208920, 17204423 and C7004-22G (CRF).

## REFERENCES

- Chen, T., Moreau, T., Jiang, Z., Zheng, L., Yan, E., Shen, H., Cowan, M., Wang, L., Hu, Y., Ceze, L., Guestrin, C., and Krishnamurthy, A. TVM: An automated end-to-end optimizing compiler for deep learning. In *Proc. of OSDI*, pp. 578–594, 2018.
- Chen, T., Zhang, Z., Jaiswal, A. K., Liu, S., and Wang, Z. Sparse moe as the new dropout: Scaling dense and self-slimmable transformers. In *Proc. of ICLR*, 2023.
- Dai, D., Deng, C., Zhao, C., Xu, R. X., Gao, H., Chen, D., Li, J., Zeng, W., Yu, X., Wu, Y., Xie, Z., Li, Y. K., Huang, P., Luo, F., Ruan, C., Sui, Z., and Liang, W. DeepSeekMoE: Towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint 2401.06066*, 2024.
- Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022.
- Gale, T., Narayanan, D., Young, C., and Zaharia, M. MegaBlocks: Efficient Sparse Training with Mixture-of-Experts. *Proc. of MLSys*, 5, 2023.
- He, J., Zhai, J., Antunes, T., Wang, H., Luo, F., Shi, S., and Li, Q. FasterMoE: modeling and optimizing training of large-scale dynamic pre-trained models. In *Proc. of PPOPP*, pp. 120–134, 2022.
- Hwang, C., Cui, W., Xiong, Y., Yang, Z., Liu, Z., Hu, H., Wang, Z., Salas, R., Jose, J., Ram, P., et al. Tutel: Adaptive mixture-of-experts at scale. *Proc. of MLSys*, 5, 2023.
- Jayarajan, A., Jinliang, W., Gibson, G., Fedorova, A., and Pekhimenko, G. Priority-based parameter propagation for distributed dnn training. In *Proc. of MLSys*, 2019.
- Jiang, A. Q., Sablayrolles, A., Roux, A., Mensch, A., Savary, B., Bamford, C., Chaplot, D. S., de las Casas, D., Hanna, E. B., Bressand, F., Lengyel, G., Bour, G., Lample, G., Lavaud, L. R., Saulnier, L., Lachaux, M.-A., Stock, P., Subramanian, S., Yang, S., Antoniak, S., Scao, T. L., Gervet, T., Lavril, T., Wang, T., Lacroix, T., and Sayed, W. E. Mixtral of experts. *arXiv preprint 2401.04088*, 2024.
- Lepikhin, D., Lee, H., Xu, Y., Chen, D., Firat, O., Huang, Y., Krikun, M., Shazeer, N., and Chen, Z. GShard: Scaling giant models with conditional computation and automatic sharding. In *Proc. of ICLR*, 2020.
- Li, J., Jiang, Y., Zhu, Y., Wang, C., and Xu, H. Accelerating distributed MoE training and inference with Lina. In *Proc. of ATC*, pp. 945–959, 2023a.
- Li, S., Lai, Z., Hao, Y., Liu, W., Ge, K., Deng, X., Li, D., and Lu, K. Automated tensor model parallelism with overlapped communication for efficient foundation model training. *arXiv preprint 2305.16121*, 2023b.
- Lin, J., Yang, A., Bai, J., Zhou, C., Jiang, L., Jia, X., Wang, A., Zhang, J., Li, Y., Lin, W., et al. M6-10t: A sharing-delinking paradigm for efficient multi-trillion parameter pretraining. *arXiv preprint 2110.03888*, 2021.
- Liu, H., Zaharia, M., and Abbeel, P. Ring attention with blockwise transformers for near-infinite context. *arXiv preprint 2310.01889*, 2023.
- Martello, S. and Toth, P. *Knapsack problems: algorithms and computer implementations*. John Wiley & Sons, Inc., 1990.
- Merity, S., Xiong, C., Bradbury, J., and Socher, R. Pointer sentinel mixture models. *arXiv preprint 1609.07843*, 2016.
- Message Passing Interface Forum. *MPI: A Message-Passing Interface Standard Version 4.0*, June 2021. URL <https://www.mpi-forum.org/docs/mpi-4.0/mpi40-report.pdf>.
- Nie, X., Zhao, P., Miao, X., and Cui, B. HetuMoE: An efficient trillion-scale mixture-of-expert distributed training system. *arXiv preprint 2203.14685*, 2022.
- NVIDIA. NCCL, 2021. <https://developer.nvidia.com/nccl>.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. PyTorch: An imperative style, high-performance deep learning library. In *Proc. of NeurIPS*, pp. 8024–8035, 2019.
- Peng, Y., Zhu, Y., Chen, Y., Bao, Y., Yi, B., Lan, C., Wu, C., and Guo, C. A generic communication scheduler for distributed dnn training acceleration. In *Proc. of SOSPP*, pp. 16–29, 2019.
- Perron, L. and Furnon, V. OR-Tools. <https://developers.google.com/optimization/>, 2019.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., and Sutskever, I. Language Models are Unsupervised Multi-task Learners, 2019. <https://openai.com/blog/better-language-models/>.
- Rajbhandari, S., Rasley, J., Ruwase, O., and He, Y. ZeRO: Memory optimizations toward training trillion parameter models. In *Proc. of SC*, pp. 1–16, 2020.

- Rajbhandari, S., Li, C., Yao, Z., Zhang, M., Aminabadi, R. Y., Awan, A. A., Rasley, J., and He, Y. DeepSpeed-MoE: Advancing mixture-of-experts inference and training to power next-generation AI scale. In *Proc. of ICML*, volume 162, pp. 18332–18346, 2022.
- Rasley, J., Rajbhandari, S., Ruwase, O., and He, Y. DeepSpeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proc. of KDD*, pp. 3505–3506, 2020.
- Riquelme, C., Puigcerver, J., Mustafa, B., Neumann, M., Jenatton, R., Susano Pinto, A., Keyzers, D., and Houlsby, N. Scaling vision with sparse mixture of experts. *Proc. of NeurIPS*, 34:8583–8595, 2021.
- Roller, S., Sukhbaatar, S., Weston, J., et al. Hash layers for large sparse models. In *Proc. of NeurIPS*, volume 34, pp. 17555–17566, 2021.
- Rotem, N., Fix, J., Abdulrasool, S., Catron, G., Deng, S., Dzhabarov, R., Gibson, N., Hegeman, J., Lele, M., Levenstein, R., et al. Glow: Graph lowering compiler techniques for neural networks. *arXiv preprint 1805.00907*, 2018.
- Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q. V., Hinton, G. E., and Dean, J. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *Proc. of ICLR*, 2017.
- Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., and Catanzaro, B. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint 1909.08053*, 2019.
- Wang, S., Wei, J., Sabne, A., Davis, A., Ilbeyi, B., Hechtman, B., Chen, D., Murthy, K. S., Maggioni, M., Zhang, Q., et al. Overlap communication with dependent computation via decomposition in large deep learning models. In *Proc. of ASPLOS*, pp. 93–106, 2022.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. Transformers: State-of-the-art natural language processing. In *Proc. of EMNLP*, pp. 38–45, 2020.
- Yang, A., Lin, J., Men, R., Zhou, C., Jiang, L., Jia, X., Wang, A., Zhang, J., Wang, J., Li, Y., et al. M6-t: Exploring sparse expert models and beyond. *arXiv preprint 2105.15082*, 2021.
- Yu, C. H., Fan, H., Huang, G., Jia, Z., Liu, Y., Wang, J., Zheng, Z., Zhou, Y., Shen, H., Shao, J., et al. RAF: Holistic compilation for deep learning model training. *arXiv preprint 2303.04759*, 2023.
- Zhang, S., Diao, L., Wu, C., Wang, S., and Lin, W. Accelerating large-scale distributed neural network training with SPMD parallelism. In *Proc. of SoCC*, pp. 403–418, 2022.
- Zheng, L., Li, Z., Zhang, H., Zhuang, Y., Chen, Z., Huang, Y., Wang, Y., Xu, Y., Zhuo, D., Gonzalez, J. E., et al. Alpa: Automating inter- and intra-operator parallelism for distributed deep learning. In *Proc. of OSDI*, pp. 559–578, 2022.
- Zhou, Y., Lei, T., Liu, H., Du, N., Huang, Y., Zhao, V. Y., Dai, A. M., Chen, Z., Le, Q. V., and Laudon, J. Mixture-of-experts with expert choice routing. In *Proc. of NeurIPS*, 2022.
- Zuo, S., Liu, X., Jiao, J., Kim, Y. J., Hassan, H., Zhang, R., Gao, J., and Zhao, T. Taming sparsely activated transformer with stochastic experts. In *Proc. of ICLR*, 2022.