



USENIX

THE ADVANCED COMPUTING
SYSTEMS ASSOCIATION

MuxTune: Efficient Multi-Task LLM Fine-Tuning in Multi-Tenant Datacenters via Spatial-Temporal Backbone Multiplexing

Chunyu Xue and Yi Pan, *Shanghai Jiao Tong University*;
Weihao Cui, *Shanghai Jiao Tong University and National University of Singapore*;
Quan Chen and Shulai Zhang, *Shanghai Jiao Tong University*; Bingsheng He,
National University of Singapore; Minyi Guo, *Shanghai Jiao Tong University*

<https://www.usenix.org/conference/nsdi26/presentation/xue-chunyu>

This paper is included in the Proceedings of the 23rd USENIX Symposium
on Networked Systems Design and Implementation.

May 4–6, 2026 • Renton, WA, USA

ISBN 978-1-939133-54-0

Open access to the Proceedings of the 23rd USENIX Symposium
on Networked Systems Design and Implementation is sponsored by



جامعة الملك عبد الله
للعلوم والتقنية
King Abdullah University of
Science and Technology

MuxTune: Efficient Multi-Task LLM Fine-Tuning in Multi-Tenant Datacenters via Spatial-Temporal Backbone Multiplexing

Chunyu Xue^{†*}, Yi Pan^{†*}, Weihao Cui^{†§}, Quan Chen[†], Shulai Zhang[†], Bingsheng He[§], Minyi Guo[†]
[†]Shanghai Jiao Tong University [§]National University of Singapore

Abstract

Parameter-Efficient Fine-Tuning (PEFT) is widely applied as the backend of fine-tuning APIs for large language model (LLM) customization in datacenters. Service providers deploy separate instances for individual PEFT tasks, giving rise to prominent resource inefficiencies, including (1) GPU underutilization from small-scale, PEFT-native operators and (2) device stalls from communication delays and data dependencies in parallelized execution. To address these issues, this paper presents MuxTune, a fine-tuning system that enables resource-efficient concurrent execution of multiple PEFT tasks. The key idea is to multiplex the backbone across independent tasks in a spatial-temporal manner for improved utilization and reduced stalls. Building on flexible, modularized backbone sharing via unified PEFT representations, MuxTune proposes hierarchical co-scheduling scheme with task, operator, and data-level optimizations. Specifically, it fuses tasks through a hybrid of spatial and temporal multiplexing, and orchestrates multi-task operator execution in two-tiered hybrid parallelism. Additionally, MuxTune employs chunk-based data alignment to mitigate inter-task ineffective tokens. Experimental results demonstrate that MuxTune achieves up to $2.33\times$ higher throughput and $5.29\times$ memory reduction compared to three state-of-the-art baselines.

1 Introduction

The paradigm of pretrain-finetune has surged in the realm of large language models (LLMs) [32, 32, 35, 37, 72, 87]. Enterprises and developers leverage domain-specific datasets to fine-tune a pretrained model into customized ones for more contextually relevant responses [7, 12, 15, 17, 75, 79]. With the evolution of emerging hardware and its rising cost, local fine-tuning has become increasingly prohibitive for many developers. In this context, LLM service providers offer public *fine-tuning APIs* [1–3, 6], empowering developers to fine-tune LLMs remotely in shared GPU datacenters. Given that these

*Equal contribution.




Figure 1: Workflow of submitting tasks and fine-tuning LLMs remotely. Each instance separately handles 1 task of diverse PEFT types, with its backbone parallelized on 2 GPUs.

APIs adopt a token-based pricing model [1, 6], optimizing fine-tuning efficiency with minimal resources is crucial.

To reduce service provisioning costs and accommodate diverse downstream domains, parameter-efficient fine-tuning (PEFT) is widely applied as one of the backends for these APIs [1, 2, 6, 43]. Figure 1 depicts an example workflow of fine-tuning services. A PEFT task attaches small-scale trainable *adapters* to targeted operators of a pretrained LLM *backbone*. Only the adapters are fine-tuned, while the backbone parameters remain frozen [47, 64]. As illustrated, developers create PEFT tasks using fine-tuning APIs, while the cluster scheduler *separately* dispatches tasks, deploys instances (hardware and the backbone), and launches PEFT programs.

However, existing PEFT frameworks (e.g., HuggingFace PEFT [18], NeMo [44]) suffer from prominent resource inefficiencies, stemming from two major reasons. First, PEFT inherently introduces small but non-negligible (in terms of latency) operators, such as LoRA down-projection [47], degrading average GPU utilization by up to $1.47\times$ (❶ in Figure 1). Some domain-specific PEFT corpora have shorter sequences and smaller batch sizes compared to pretraining [82, 88], further intensifying GPU underutilization. Second, PEFT adapters exacerbate device stalls arising from communication delays and data dependencies between operators, in both intra- [78] and inter-stage [49, 65] parallelism (❷ in Figure 1). Without backbone weight gradients, PEFT does not support stall-free pipeline, such as DeepSeek-V3 DualPipe [63] and ZeroBubble ZB-H2 [70]), which splits the backward pass into input

and weight gradients for finer scheduling. Communication overlapping techniques [26, 86], which decompose computation operators into smaller tiles to overlap with communication, are prone to GPU underutilization and ultimately inflate end-to-end latency (§2.2). As GPU compute capability grows faster than interconnect bandwidth, the above inefficiencies are being further aggravated — with a $3.18\times$ gap in Model FLOPs Utilization (MFU) [53] from NVIDIA H100 GPUs with NVLink to A40 GPUs with PCIe4.0 connection.

To improve GPU utilization and mitigate device stalls, in multi-tenant datacenters, we move beyond single-task optimization to multi-task co-scheduling by multiplexing the LLM backbone. We propose to address the inefficiencies via *spatial-temporal backbone multiplexing*, where a single backbone is shared across tasks by batching spatially and interleaving temporally. While some multiplexing strategies are adopted in pretraining or serving [33, 76, 95], they have critical limitations when directly reused in PEFT scenarios.

Specifically, batching-based spatial multiplexing adapted from multi-LoRA serving (e.g., SLoRA [76], PunicA [34]) can amplify device stalls and introduce ineffective computation, as PEFT workloads differ fundamentally from serving (§2.3). On the other hand, temporal multiplexing and overlapping techniques, which are inspired by recent pretraining and serving systems (e.g., Centauri [33], NanoFlow [95]), alleviate some communication bottleneck yet at the cost of even lower GPU utilization. Direct spatial or temporal multiplexing approaches not only fail to efficiently resolve these resource inefficiencies but may also conflict with one another.

We therefore propose **MuxTune**, a resource-efficient system for multi-task PEFT that enables spatial-temporal backbone multiplexing in datacenters. Nevertheless, several challenges arise across the system design and scheduling. At the task level, it is non-trivial to make the hybrid decision of spatial and temporal multiplexing across tasks to tame their inherent tradeoff. At the operator level, efficiently orchestrating intricate operator execution across spatial and temporal multiplexing remains challenging under hybrid parallelism. At the data level, efficiently aligning variable-length sequences from spatially batched tasks while mitigating inter-task ineffective tokens (i.e., zero padding) requires careful design.

The core idea of MuxTune is unifying spatial and temporal multiplexing within a hierarchical design, built on top of unified PEFT representations, to optimize utilization and device stalls jointly. This approach first modularizes diverse PEFT types into unified representations, enabling flexible sharing of the backbone with consistent training behavior (§3.2). Building on this foundation, MuxTune incorporates a hierarchical multi-task co-scheduling scheme with tailored solutions at each level. At the task level, MuxTune introduces a *hybrid task* (“hTask”) abstraction to navigate the spatial-temporal tradeoff: tasks within the same hTask are spatially batched to improve utilization, while different hTasks are temporally interleaved to hide stalls. Guided by a pipeline-based cost




Figure 2: Representative categories of PEFT algorithms.

model, it uses a dynamic programming (DP) algorithm to determine the optimal way to fuse tasks into multiple hTasks (§3.3). At the operator level, MuxTune orchestrates operator execution of hTasks on the shared, modularized backbone. By solving a two-tiered scheduling problem across inter- and intra-stage parallelism¹, it coordinates intricate dependencies under hybrid parallelism to generate a fine-grained, stall-free execution schedule (§3.4). At the data level, MuxTune aligns variable-length data across tasks within each spatially fused hTask via sequence packing and chunk-based partitioning (§3.5). The main contributions of this work are:

- We reveal the resource inefficiencies of PEFT workloads in datacenters, and identify the opportunities and challenges for efficient multi-task concurrent execution.
- We present MuxTune, an efficient and scalable fine-tuning system for multi-task PEFT models via hierarchical spatial-temporal backbone multiplexing.
- We propose modularized backbone sharing for flexible and scalable spatial-temporal multiplexing. Building on this, we devise fine-grained, hierarchical co-scheduling scheme that integrates task, operator, and data-level optimizations to systematically optimize resource efficiency.

We build MuxTune and evaluate it with four LLMs and three PEFT datasets on testbeds with NVIDIA A40 and H100 GPUs. Empirical results show that MuxTune achieves up to $2.33\times$ higher throughput and $5.29\times$ memory reduction compared to three state-of-the-art baselines. MuxTune is open-sourced at <https://github.com/sjtu-epcc/muxtune>.

2 Background and Motivation

2.1 Parameter-Efficient Fine-Tuning (PEFT)

Basics. PEFT employs trainable adapters to adapt pre-trained LLMs to specific domains cost-effectively, instead of training LLMs from scratch. As shown in Figure 2, it consists of three representative categories [43, 62]: (1) *Additive* (e.g., Adapter-Tuning [46]) that inserts adapters into specific positions of the model architecture. (2) *Selective* (e.g., Diff-Pruning [41]) that fine-tunes a subset of parameters via binary masks. (3) *Reparameterized* (e.g., LoRA [47]) that constructs low-rank transformation of the original model parameters.

¹Intra-stage refers to data parallelism (DP), tensor parallelism (TP) [78] and other variants [56, 57], inter-stage refers to pipeline parallelism (PP) [65].

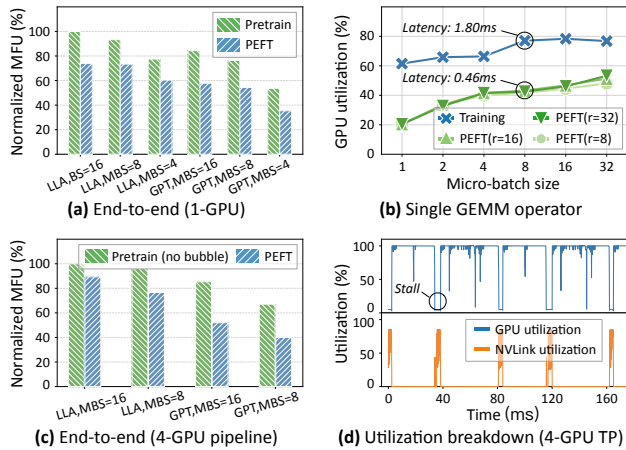


Figure 3: PEFT inefficiencies (MBS: micro-batch size, sequence length 128). (a) Single-GPU MFU of 8-layer models (global batch size 32, LLa: LLaMA7B, GPT: GPT2.7B). (b) Operator utilization (shape $[MBS, 128, 4096] \times [4096, r]$, $r=4096$ for pretraining). (c) Multi-GPU MFU of full models (global batch size 128). (d) GPU and NVLink utilization.

Characteristics. PEFT workloads in datacenters exhibit distinct characteristics. First, PEFT enables flexible configuration and attachment of adapters without affecting the backbone, as they manage independent parameters.

Second, many PEFT tasks are fine-tuned on the same backbone type, as there are thousands of developers [9] yet a few available backbones (e.g., 7 for OpenAI [1], 3 for Gemini [6]). These tasks retain independent and variable-length batches (i.e., *data heterogeneity*), while sharing backbone operators except for adapter-related ones (i.e., *backbone homogeneity*). Some PEFT-based multimodal works have also proposed fine-tuning multiple adapters on a single backbone using the same corpus, each as an independent task [42, 69].

Third, service providers have full access to model specifics, as they define, instantiate, and execute models when providing fine-tuning APIs. This differs from conventional infrastructure providers, which only allocate hardware and run scripts, with tasks internally optimized by their frameworks.

Lastly, compared with pretraining, some domain-specific PEFT corpora, such as sentiment analysis [82] and short text generation [88], feature shorter sequence lengths (e.g., 64) and smaller corpus sizes [4, 50]. To ensure model generalization, prior works [47, 80] advocate adopting smaller batch sizes for model fine-tuning (e.g., 128 for GPT-3 [47]). Moreover, sequence lengths vary significantly across different PEFT corpora due to their varying domain focuses [12, 15, 79].

2.2 Inefficiencies of PEFT Workloads

PEFT workloads typically suffer from resource inefficiencies on advanced GPUs (§5.1), as shown in Figure 3 with two metrics: (1) *GPU utilization*, which measures the occupancy of

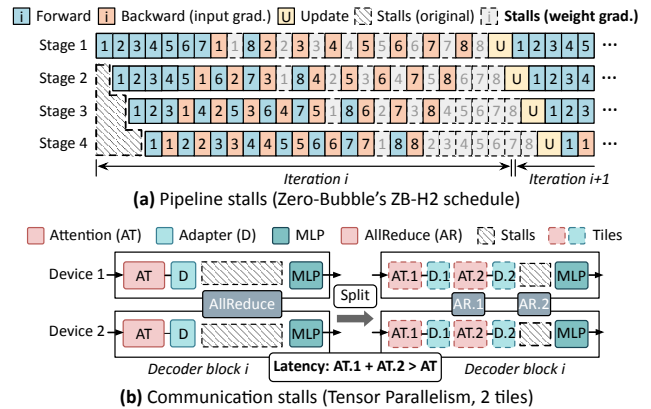


Figure 4: Device stalls in PEFT under model parallelism.

streaming multiprocessors (SMs) [24]; (2) *Model FLOPs Utilization* (MFU), which measures end-to-end efficiency [85].

Insufficient Utilization. Figure 3(a) shows that PEFT exhibits lower single-GPU MFU than model pretraining (up to $1.47\times$). The main reason is that PEFT omits compute-intensive backbone weight gradients and introduces small-scale operators, such as LoRA down-projection [47] and learnable vectors of Prefix-Tuning [60]. For example, LoRA rank (up to size 64) is $64.0\times$ smaller than the hidden size of LLaMA7B (size 4096). Figure 3(b) shows that these small operators incur non-negligible latency (e.g., 0.46ms vs. 1.80ms in pretraining) and GPU underutilization (a gap of up to 40.9%). Their attachment to non-continuous backbone operators (e.g., *qkv-proj* per decoder block [84]) prevents direct vertical or horizontal fusion due to stringent interdependencies. Moreover, since PEFT corpora typically feature shorter sequences and smaller batch sizes (§2.1), the input size remains limited and further undermines the efficiency of GPU parallel computing. Worse yet, our experiments across GPU architectures (configured as Figure 3(a)) show that the average MFU of PEFT on NVIDIA V100, A40, and RTX6000 is $0.84\times$, $0.68\times$, $0.59\times$ that of pretraining, demonstrating that the underutilization is exacerbated by higher-end hardware.

Device Stall. In model parallelism, device stalls arise from communication delays and data dependencies between parallelized operators. PEFT adapters exacerbate stalls by increasing stage latencies and intra-stage communication costs [76]. Worse still, advanced training optimizations, such as stall-free pipeline [63, 70] and communication overlapping [26, 86], cannot be directly reused in PEFT. This leads to further MFU degradation during multi-GPU execution, with a drop of up to $1.65\times$ (Figure 3(c), worse than 1-GPU case). Figure 3(d) shows GPU and NVLink utilization in 4-GPU tensor parallelism with significant stalls. Below, we demonstrate two stall types and discuss how PEFT aggravates them.

The first is *pipeline stalls* (also known as bubbles) arising from pipeline flushes and inter-stage dependencies (Figure 4(a)). Prior works (e.g., DualPipe [63], ZB-H2 [70]) reduce bubbles by splitting the backward pass into input

and weight gradients, finely scheduling for near-zero-bubble pipeline. PEFT inherently lacks support for this technique due to the absence of backbone weight gradients. As depicted, unlike original stalls, those stalls induced by omitted weight gradients grow linearly with the number of micro-batches, thus cannot be amortized. Directly adopting DualPipe in PEFT undermines throughput by $1.16\times$ compared to 1F1B [65].

The second is *communication stalls* in tensor parallelism caused by communication delays and operator dependencies (Figure 4(b)). Prior works [26, 86] reduce these stalls by decomposing computation into smaller tiles to overlap with communication. As shown, since PEFT inherently suffers underutilization, such decomposing intensifies the issue (a 24.5% utilization drop) and inflates overall latency by $1.17\times$ for GPT2.7B with 2 GPUs. Other stall types such as those from parameter synchronization [74, 92] also exist in PEFT.

2.3 Intuitive Multiplexing Approaches

Coarse-grained spatial multiplexing causes poor scalability and limited performance gains. A straightforward way to improve utilization is co-locating PEFT tasks via NVIDIA MPS [23] or multiple streams [21], similar to [48, 66]. However, the utility of this approach is limited by memory constraints (❶ in Figure 5). To illustrate, we profile the memory breakdown of a LoRA LLaMA7B (batch size 8, sequence length 128): the backbone parameters and activations consume 13.4GB and 4.3GB, respectively, while the total memory footprint is 18.1GB. With 4 A40 GPUs (48GB each), only 8 tasks can be co-located without parallelization, preventing scaling to larger models or more concurrent tasks. Additionally, the lack of fine-grained control over operator execution results in suboptimal stall reduction and potential inter-task interference (e.g., a $2.5\times$ performance drop reported in [95]).

Batching-based spatial multiplexing causes enlarged stalls and ineffective computation. Prior multi-LoRA serving systems [34, 76] simplistically batch multiple requests for concurrent computation. While effective for serving, this approach is misaligned with PEFT workloads (❷ in Figure 5). First, backward pass introduces pipeline stall concerns that are generally absent in forward-only serving [30]. Excessive batching exacerbates device stalls and undermines end-to-end performance (§3.3). Second, batching offers diminishing returns in compute-bound PEFT, compared to the memory-bound decoding phase of serving [29]. Lastly, continuous batching [90] used in serving avoids the need for padding, which remains essential yet computationally ineffective to align variable-length sequences in PEFT (§3.5).

Temporal multiplexing and overlapping result in even lower GPU utilization. The approach of temporally overlapping communication with computation from other tasks, inspired by recent pretraining and serving systems [33, 95], also proves counterproductive for PEFT (❸ in Figure 5). Specially, temporal multiplexing executes operators of each task




Figure 5: General views of intuitive multiplexing approaches.

sequentially, instead of exploiting batching opportunities to improve intra-operator utilization. Given the limited input sizes typical of PEFT workloads, it may also be impractical to excessively increase the micro-batch size for a single task.

2.4 Our Approach and Challenges

We identify an opportunity to address the limitations via spatial-temporal backbone multiplexing, i.e., sharing the backbone across tasks by batching spatially and interleaving temporally. This paradigm forms an intricate optimization space with multiple *coupled* dimensions, e.g., determining the optimal spatial-temporal combination while scheduling multi-task pipeline execution. To reduce optimization complexity, we propose *hierarchically decomposing* the space into three levels and tackling the unique challenges at each level:

Task Level: Navigating spatial-temporal tradeoff. There is an inherent tradeoff between spatial and temporal multiplexing (§3.3). This poses a primary challenge to design a dynamic scheduling policy that can intelligently batch tasks spatially to improve utilization while interleaving them temporally to hide the pipeline and communication stalls.

Operator Level: Stall-free multi-task orchestration. Coarse-grained spatial multiplexing fails to mitigate stalls without operator-level execution control. Under hybrid parallelism, operator orchestration across spatially and temporally multiplexed tasks becomes a complicated two-tiered problem (§3.4). It requires tailored algorithms and fine-grained coordination for stall-free execution and high GPU utilization.

Data Level: Mitigating inter-task ineffective computation. Spatially multiplexing tasks with variable-length sequences naturally requires aligning along sequence dimension. Naïve strategies such as zero padding or packing into long sequences either waste compute on ineffective tokens or degrade efficiency. Without careful alignment design, these effects can diminish the gains from multi-task multiplexing.

3 System Design

3.1 Overview

MuxTune is an efficient and scalable system for multi-task PEFT, serving as the backend for LLM fine-tuning APIs to enhance resource efficiency in multi-tenant GPU datacenters. The key idea of MuxTune is to multiplex the backbone in a spatial-temporal manner across independent PEFT tasks

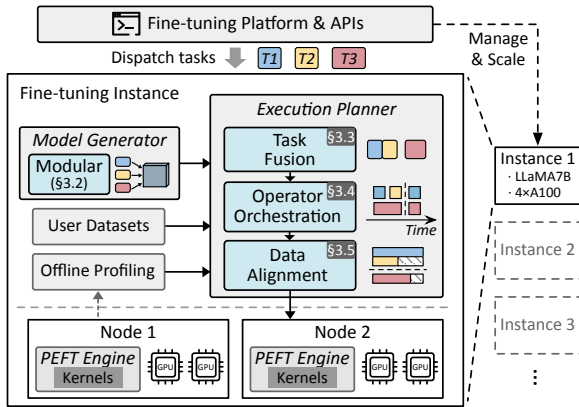


Figure 6: Architecture overview of MuxTune.

via flexible, modularized backbone sharing and fine-grained, hierarchical multi-task co-scheduling.

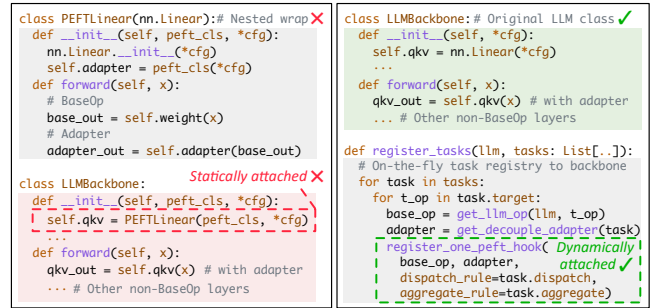
Figure 6 presents the architecture of MuxTune with three main modules: model generator, planner, and PEFT engine to address the challenges in §2.4. Initially, users configure their PEFT tasks (e.g., backbone type) and submit tasks via fine-tuning APIs. The cluster scheduler dispatches tasks with the same backbone to an in-flight instance or creates a new one based on scheduling policies (e.g., budget-based [19]). Given dispatched tasks, *Model Generator* builds a PEFT model with multi-task adapters based on modularization (§3.2). With the PEFT model, user datasets, and profiling data, *Execution Planner* fuses tasks into hybrid tasks by adaptively combining spatial and temporal multiplexing (§3.3). Then, the planner orchestrates the fine-grained operator execution of these hybrid tasks under two-tiered hybrid parallelism (§3.4). From the data perspective, data batches are loaded in a streaming manner and aligned across spatially batched tasks (§3.5). At runtime, *PEFT Engine* concurrently executes the dispatched PEFT tasks via efficient fused and overlapped kernels.

3.2 Backbone Sharing

MuxTune adopts flexible, modularized backbone sharing to co-locate PEFT tasks with diverse workloads. This subsection introduces PEFT modularization and sharing the backbone across multi-task sub-modules for efficient multiplexing.

PEFT Modularization. As shown in Figure 7(a), existing single-task frameworks [18, 44] *statically* inject adapters into the LLM backbone, i.e., directly treating adapters of a task as LLM layers. While intuitive, employing such a static implementation in multi-task PEFT scenarios cannot support dynamic workloads with shifting adapter numbers and types, as it requires from-scratch model reinitialization to tackle on-the-fly task arrival or completion events.

In contrast, MuxTune employs *decoupled adapters* to support flexible user customization and efficient scaling by the cluster scheduler. Specifically, by examining mainstream PEFT types in §2.1, we abstract the general PEFT workflow



(a) Static nested adapter implementation

(b) On-the-fly adapter attachment

Figure 7: Current static nested adapter implementation and our modularization-based dynamic adapter attachment.

into four sub-modules: (i) *BaseOp* is an operator of the backbone that has the possibility to be attached an adapter, such as QKV and Linear Projection (Attention is excluded). (ii) *Adapter* describes the PEFT algorithms (e.g., LoRA [47]) and is customized by users. (iii) *Dispatch* defines the multi-task data dispatching rules to prepare input tensors for *BaseOp* and *Adapter*. (iv) *Aggregate* defines the data aggregation rules to gather output tensors from *BaseOp* and *Adapter*.

Dynamic Multi-Task Backbone Sharing. As shown in Figure 7(b), unlike prior frameworks, MuxTune preserves non-intrusiveness to the LLM backbone while providing a `register_tasks()` API to handle on-the-fly task arrival events. This is the cornerstone of multi-task backbone sharing. When the cluster scheduler assigns new tasks to an in-flight instance, the local model generator reactively invokes this API to register on the multiplexed backbone without costly model reinitialization (§4). After that, MuxTune implicitly automates the task, operator, and data-level optimizations as introduced below, and efficiently executes the generated multi-task PEFT model at runtime.

Isolation and Convergence Guarantee. While improving efficiency via multi-task backbone sharing, MuxTune provides stringent guarantee for both inter-task isolation and consistent model convergence. Specifically, MuxTune safely instantiates the LLM backbone and user-defined adapters (§2.1), thereby preventing most runtime errors (e.g., semantic errors) described in [51]. Through fine-grained memory modeling and operator orchestration, MuxTune also mitigates inter-task performance interference and OOM risks.

In MuxTune, both *BaseOps* and *Adapters* across temporally interleaved tasks are naturally isolated in time, with no impact on convergence. For spatially batched tasks, *BaseOps* are concatenated along batch dimension and exhibit mathematical isolation. Taking a GEMM *BaseOp* with two *Adapters* as an example, the forward computation of *BaseOp* is:

$$BW_g = [B_1, B_2]_b W_g = [B_1 W_g, B_2 W_g]_b, \quad (\text{BaseOp fwd}) \quad (1)$$

where B_i is the batch of task i , B is the concatenated batch, and W_g is GEMM weight. In backward pass, each B_i is inde-

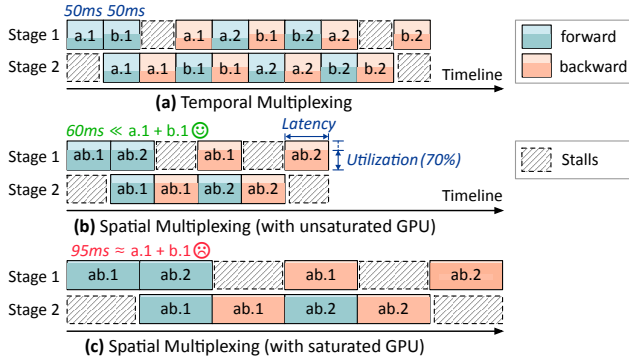


Figure 8: Tradeoff of spatial and temporal multiplexing. x_i represents i -th micro-batch of (hybrid) task x .

pendently computed a loss and backpropagated, while computations of the same *BaseOp* are batched as:

$$G^{in} = [G_1^{in}, G_2^{in}]_b \leftarrow [G_1^{out}, G_2^{out}]_b W_g^T, \quad (\text{BaseOp bwd}) \quad (2)$$

where G_i^{in} and G_i^{out} are the input and output gradient of task i . G^{in} is the transient gradient buffer of the input to *BaseOp*. Other non-batchable operators (e.g., adapters) are independently computed in a fused manner (§3.4.3). Such isolation preserves consistent convergence (e.g., 0.07 mean-square deviation) while avoiding numerical failure propagation (e.g., gradient NaN from overlarge learning rate) among tasks.

3.3 Task Fusion

Based on modularized representation and flexible interface, MuxTune is able to multiplex the backbone in different manners: (1) batching PEFT tasks spatially by executing batched *BaseOp*, *Dispatch*, fused *Adapter* and *Aggregate*; or (2) interleaving tasks temporally by sequentially executing *BaseOp*, *Dispatch*, *Adapter*, and *Aggregate* of each task.

Spatial-Temporal Tradeoff. Figure 8 illustrates these two multiplexing options. The *interleaving-based temporal multiplexing* computes tasks sequentially, overlapping stalls with computations across tasks, but risks GPU underutilization with separate *Adapters* and limited input sizes (§2.2). The *batching-based spatial multiplexing* batches the *BaseOp* and fuses the *Adapter* of independent tasks to improve utilization. However, excessive batching prolongs operator (or stage) latency and exacerbates stalls in parallelized execution.

The optimal multiplexing decision depends on task input sizes, PEFT parameters, and parallelism strategy. As shown in Figure 9(a), spatially batching tasks yields better performance when GPUs are unsaturated, while temporally interleaving is preferable at higher GPU utilization. This shift stems from the diminishing returns induced by excessive spatial batching beyond GPU saturation (Figure 9(b)). For example, ideally batching 8 tasks, each with micro-batch size 8 and sequence length 128, only improves throughput by 1.12 \times , far short of the expected 8 \times gain. Notably, the above analysis only




Figure 9: Quantified analysis of two options. (a) 2 tasks on 16-layer LLaMA7B with 4-GPU pipeline (4 micro-batches, seq-len 64). (b) 1 task on 8-layer LLaMA7B with 1 GPU.

involves two tasks, while scaling to more tasks exponentially elevates the complexity of such multiplexing decisions.

Hybrid Task Abstraction. Given the scheduling complexity, deriving the optimal multiplexing decision requires a hierarchical scheduling of all tasks. We introduce a *hybrid task abstraction* (“hTask” for short) to unify the two multiplexing options, which is fused from independent tasks. Within a hTask, tasks are fused and batched in a spatial manner; among hTasks, their computations are temporally interleaved.

The task fusion is framed as a bin-packing problem. Initially, M tasks $\mathcal{T} = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_M\}$ co-locates on the backbone (S stages, each has $N_g^{(s)}$ GPUs). As latency correlates positively with the input size (due to the backbone homogeneity described in §2.1), we sort \mathcal{T} by token count (denoted as n_i for the i -th task) in ascending order. Let $\mathcal{H}_{i \rightarrow j}$ be the hTask with task i to j . A unified number of micro-batches C for all tasks is set empirically. Sequences of each task are padded to a maximum length l_i , as detailed in §3.5.

Cost Model. We build a cost model to assess end-to-end latency and memory footprint under hybrid parallelism. Since forward and backward passes of the same stage share similar latency in PEFT (due to the absence of weight gradients), we model the latency of a hTask $\mathcal{H}_{i \rightarrow j}$ for the s -th stage ($St.$) as:

$$L^{(s)}(\mathcal{H}_{i \rightarrow j}) = \sum_{o \in St.} t_o \left(\sum_k n_k \right) / N_g^{(s)} \quad (\text{BaseOp Lat.}) \\ + \sum_{\{a\} \in St.} \max \left\{ \sum_k u_a \cdot t_a(n_k), \max_k t_a(n_k) \right\}, \quad (\text{Adapter Lat.}) \quad (3)$$

where the first line models the latency of *BaseOps* sharded across $N_g^{(s)}$ GPUs ($k \in [i, j]$). The $t_o(x)$ is the latency of computation operator o with x tokens (communication is overlapped as in §3.4.2). The second line estimates the latency of fused adapters $\{a\}$, where $u_a(x)$ and $t_a(x)$ are the GPU utilization and latency of *Adapter* a with x tokens. We use the weighted sum $u_a \cdot t_a(x)$ to estimate the overall latency after horizontal fusion (§3.4.3), while bounding it with the maximum per-adapter latency to avoid the bottleneck effect. In the pipeline with S stages, the end-to-end latency of $\mathcal{H}_{i \rightarrow j}$ is:

$$L(\mathcal{H}_{i \rightarrow j}) = 2 \sum_{s=1}^{S-1} \{L^{(s)}(\mathcal{H}_{i \rightarrow j})\} + 2C \max_{1 \leq s \leq S} L^{(s)}(\mathcal{H}_{i \rightarrow j}), \quad (4)$$

where the first term estimates the latency sum of warm-up and drain phases, while the second term models the overall latency of the steady phase with C micro-batches [52].

As for memory footprint, we find that it mainly consists of three parts in PEFT: (i) backbone parameters M_b , (ii) input gradients $M_g^{(i)}$ of task i , and (iii) activation $M_a^{(i)}$. For 1F1B pipeline, the maximum per-stage memory is estimated as:

$$M_{stage} = [M_b + \sum_{i=1}^M M_g^{(i)}] / S + \sum_{i=1}^M M_a^{(i)}(b_i, l_i), \quad (5)$$

where the first two terms are irrelevant to the input size (e.g., micro-batch size b_i). The third term is accumulated to at most S copies and is proportional to b_i and the sequence length l_i . In practice, M_g typically resues the allocated memory of M_a [28]. This memory model is used to assess whether a hybrid task (hTask) would cause Out-of-Memory (OOM) issues during the construction. Evaluated in §5.3, it precisely matches the scaling of the measured memory footprint.

Task Fusion with DP Algorithm. Based on cost modeling, we employ a dynamic programming (DP) algorithm to minimize the end-to-end pipeline latency when bin-packing M tasks into N hTasks. We derive the state transition equation, where $F(m, n)$ is the minimal end-to-end latency of packing the first m tasks into n hTasks ($m \in [1, M]$, $n \in [1, N]$, $m \geq n$):

$$F(m, n) = \min_{n-1 \leq i \leq m} \{F(i, n-1) + L(\mathcal{H}_{(i+1) \rightarrow k}) / S\}, \quad (6)$$

$$F(m', 1) = L(\mathcal{H}_{1 \rightarrow m'}), \quad \forall m' \in [1, M],$$

where the impact of $\mathcal{H}_{(i+1) \rightarrow k}$ over $F(m, n)$ is estimated via average per-stage latency to satisfy the optimal substructure of DP. This is because the steady phase typically dominates end-to-end latency, in which $\mathcal{H}_{(i+1) \rightarrow k}$ adds one forward-backward pass [61]. The optimal fusion plan is derived as $F^* = \min_{1 \leq N \leq M} \{F(M, N)\}$. Notably, minimizing end-to-end latency also balances the loads across hybrid tasks, as the pipeline is bottlenecked by its slowest stage [45]. While the time complexity is $O(M^2(S+M))$, the algorithm incurs minimal overhead with modest number of tasks per backbone (§5.3) and can be accelerated via parallel computing of N .

3.4 Operator Orchestration

After task fusion, MuxTune orchestrates hybrid task execution at the operator² granularity across spatial and temporal multiplexing in both intra- and inter-stage parallelism.

Disaggregating Two-Tiered Orchestration. Given the distinct characteristics of intra- and inter-stage parallelism, we disaggregate operator orchestration by grouping N hTasks $\mathcal{H} = \{\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_N\}$ into P buckets $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_P\}$. Those hTasks of the same bucket are interleaved within a pipeline clock (i.e., intra-stage parallelism), while different buckets are interleaved across pipeline clocks (Figure 10).

²To unify intra- and inter-stage parallelism, we use “operator” to represent both layer (e.g., GEMM) and stage-granularity computations.

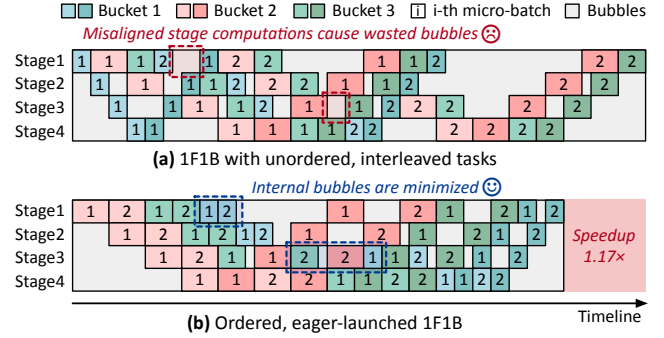


Figure 10: Example of inter-stage orchestration. Light/dark color represents forward/backward pass of hTask buckets.

Workload-Balanced Hybrid Task Grouping. To efficiently determine the optimal grouping strategy, we decouple hTask grouping from operator orchestration, based on the observation that *given a fixed P , balanced workloads lead to fewer internal bubbles and lower end-to-end latency*. We traverse P from 1 (grouping all hTasks) to N (one hTask in each bucket), optimizing \mathcal{G} by minimizing inter-bucket variance:

$$\mathcal{G}^*(P) = \arg \min_{\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_P\}} \sum_{j=1}^P |L^{(1)}(\mathcal{G}_j) - \overline{L^{(1)}(\mathcal{G})}|^2, \quad (7)$$

$$\text{s.t. } L^{(1)}(\mathcal{G}_j) = \sum_i L^{(1)}(\mathcal{H}_i), \quad \forall \mathcal{H}_i \in \mathcal{G}_j,$$

where $L^{(1)}(\mathcal{H}_i)$ is the first stage latency of hTask \mathcal{H}_i , which is used because this step focuses on balancing workloads across groups. We then invoke the orchestration process below to model the end-to-end latency for each fixed P , and select the optimal grouping strategy \mathcal{G}^* with the minimal latency.

3.4.1 Inter-Stage Orchestration

To optimize inter-stage execution with hTask buckets $\mathcal{G} = \{\mathcal{G}_1, \dots, \mathcal{G}_P\}$, the goal is to schedule the micro-batches of $\{\mathcal{G}_j\}_{1 \leq j \leq P}$ to minimize end-to-end pipeline latency.

Computation Homogeneity. Micro-batch scheduling involves a re-entrant flow shop problem [39], where $\{\mathcal{G}_j\}_{1 \leq j \leq P}$ are repeatedly processed by stages in forward and backward passes. Unlike prior training works with complex scheduling [52], we observe *computation homogeneity* in PEFT: (i) in each bucket \mathcal{G}_j , micro-batches retain a consistent shape across iterations (§3.5), while shapes may vary across buckets; (ii) for bucket \mathcal{G}_j , forward and backward passes share the same execution time due to the absence of weight gradients.

These features enable us to prune the scheduling space and reduce algorithm overhead. Specifically, with (i), MuxTune adopts a unified pipeline template to execute multi-task iterations in a *structured* manner, rather than relying on costly real-time scheduling per iteration. With (ii), as shown in Figure 10, MuxTune efficiently minimizes internal bubbles by interleaving forward and backward passes of those buckets with similar latencies (e.g., the lower blue block of Figure 10(b)).

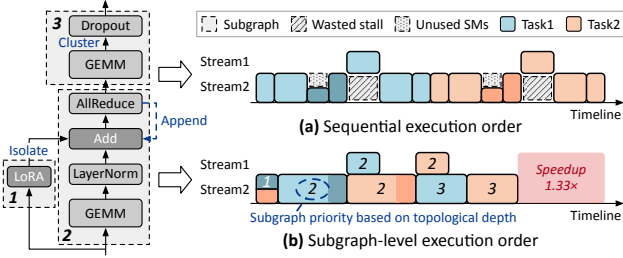


Figure 11: Example of intra-stage orchestration. Left is an example computational graph with darker blocks for adapters.

Structured Pipeline Template. The pipeline template is dynamically generated based on bucket information, such as stage latencies and memory footprint, for structured multi-task execution. Specifically, to mitigate bubbles, the template generation extends the 1F1B pipeline with three rules: (1) Sorting \mathcal{G} by $L^{(1)}(\mathcal{G}_j)$ in descending order to enable \mathcal{G}_j to fill the bubbles of $\mathcal{G}_{j-1}/\mathcal{G}_{j+1}$ (Figure 10(b)). (2) Keeping micro-batches of the same bucket consecutive, as they are perfectly matched in terms of latency. (3) Eagerly launching as many micro-batches as possible within memory limits (§3.3) to ensure sufficient pending batches per stage [96]. With these rules, MuxTune enhances efficiency by $1.17\times$ while maintaining structured execution. We have provided a detailed **optimality analysis** in Appendix §A that demonstrates this structured pipeline always mitigates internal bubbles at the last stage, thereby achieving the near-optimal execution.

3.4.2 Intra-Stage Orchestration

To optimize intra-stage execution within a hTask bucket \mathcal{G}_j , we schedule computation and communication operators of hTasks $\{\mathcal{H}_i\} \in \mathcal{G}_j$ to mitigate device stalls, which is equivalent to minimizing stage latency as illustrated in Figure 11.

Dependency-Aware Graph Construction. The computational graph of each hTask is a directed acyclic graph (DAG), naturally driving the intra-stage orchestration across hTasks as a multi-DAG scheduling problem [68]. MuxTune uses *subgraph* as the minimal orchestrating unit, because model execution is sequential, while a longer sequence of computation operators benefits fully overlapping communication.

The subgraph construction is operated in a *dependency-aware* manner. MuxTune first segments each DAG into subgraphs by clustering consecutive computation operators and appending each communication operator to its dependent operator (left part of Figure 11). Small-scale adapters are isolated as independent subgraphs. Then, MuxTune assigns a priority value to each subgraph according to its topological depth. This ensures interleaved execution while adhering to cross-graph dependencies created by adapter fusion (§3.4.3). Notably, the subgraph concept is compatible with vertical fusion-based kernels and CUDA graph techniques [10, 20].

Subgraph Scheduling. Given multiple segmented DAGs, MuxTune extends the Kahn algorithm [54] to a multi-DAG

Algorithm 1 Priority-Based Subgraph Scheduling

```

1: Input: hybrid tasks  $\mathcal{H} = \{\mathcal{H}_1, \dots, \mathcal{H}_n\}$ 
2: function SubgraphSchedule( $\mathcal{H}$ )
3:    $\mathcal{P} \leftarrow \text{PriorityQueue}()$ 
4:   for  $\mathcal{H}_i \in \mathcal{H}$  do ▷ Initialize queue
5:      $\mathcal{P}.\text{enqueue}(\text{GetZIDSubgraphs}(\mathcal{H}_i.\text{graph}))$ 
6:    $\text{launch\_schedule} \leftarrow 0; t \leftarrow 0$ 
7:   while  $\mathcal{P} \neq \emptyset$  do
8:      $\text{subgraph} \leftarrow \mathcal{P}.\text{dequeue}(\text{max\_lat})$  ▷ Highest priority
9:      $\text{DAG} \leftarrow \text{subgraph}.\text{parent\_graph}$  ▷ Parent DAG
10:     $\text{DAG}.\text{remove}(\text{subgraph})$ 
11:     $\mathcal{P}.\text{enqueue}(\text{GetZIDSubgraphs}(\text{DAG}))$ 
12:     $\text{launch\_schedule}.\text{record}(< \text{subgraph}, t >)$ 
13:     $t \leftarrow t + \text{subgraph}.\text{latency}$  ▷ Update timer
14:   return  $\text{launch\_schedule}$ 

```

variant in a latency-aware manner. Algorithm 1 outlines the scheduling process. Given initial hTasks, it maintains a priority queue \mathcal{P} to track zero in-degree subgraphs in each DAG (line 3-5). In each iteration, the algorithm filters for the highest-priority subgraphs, selecting the one with the longest cumulative latency (from internal operators) to maximize overlap with in-flight communication operators. Then, the subgraph is removed from its DAG, while new zero in-degree ones are enqueued (line 6-13). With *launch_schedule*, MuxTune efficiently coordinates the runtime execution of multi-task subgraphs within the pipeline stage (§4).

3.4.3 Adapter Fusion and Communication Overlapping

To improve GPU utilization of small-scale, PEFT-native operators, we adopt *horizontal adapter fusion* with fine-grained management of GPU resources (§4), given that adapters cannot be directly batched across tasks. The fusion strategy consists of three cases: (1) For spatially batched tasks within the same hTask, their adapters are fused; (2) For hTasks within the same bucket (i.e., interleaved in intra-stage parallelism), if all assigned a single task, their adapters are fused without impeding inter-task overlap. As shown in Figure 11, the LoRA operators of Task1 and Task2 are fusible because they are not in the same subgraph of AllReduce operator. Conversely, Add operators cannot be fused, because the fusion would introduce a global synchronization in prior to the AllReduce operators of Task1 and Task2. (3) No fusion occurs across buckets (i.e., interleaved in inter-stage parallelism).

When overlapping computation and communication across different streams in intra-stage orchestration (Figure 11), we observe that excessive CTA usage of communication kernels undermines compute kernel efficiency, while too few CTAs underutilize NVLink connection. To resolve this tradeoff, we adopt NVLink SHARP [27] to offload reductions into the NVSwitch, sustaining near-peak link bandwidth with a small CTA budget. As a result, the network kernel fully overlaps with computation of other tasks using only 8 CTAs.




Figure 12: Example of existing data alignment strategies and our chunk-based alignment (S_i : i -th sequence). Dependency represents KV cache reuse in causal attention computation.

3.5 Data Alignment

To execute a hybrid task (hTask), the data batches of its internal PEFT tasks need to be aligned along sequence dimension. One common strategy is to zero pad all sequences to a global maximum length (Figure 12(a)). This incurs substantial inter-task ineffective tokens (i.e., without semantic information), thereby wasting compute and memory resources. Another industrial-grade approach for pretraining is to pack sequences into longer ones. However, recent works [31, 52] observe that relying solely on packing degrades fine-tuning efficiency, as attention masks lead to wasted attention computation across sequences. Consequently, fine-tuning APIs often mandate sequence padding to maximum lengths, with intra-task zero-padded tokens billed to users [3]. However, in MuxTune, inter-task ineffective tokens arise from data alignment in co-scheduling and cannot be billed to users. Therefore, MuxTune focuses on *mitigating inter-task ineffective tokens*.

Reinventing Packing with Chunk-Based Alignment. To mitigate inter-task padding, MuxTune adopts a dual-step alignment strategy that maximizes efficiency without compromising model quality (Figure 12(c)). First, it adaptively packs sequences within a single global batch for each task, respectively, to ensure no impact on model convergence (e.g., Pack1/Pack2 for Task1 and Pack3 for Task2). This step transforms per-task sequences into a set of longer, denser packed ones. Second, MuxTune uniformly partitions packed sequences into equal-sized chunks (e.g., size of 4 in Figure 12(c)). For sequences longer than the chunk size (e.g., Pack1), MuxTune scatters them across multiple consecutive chunks with the dependency of KV cache reuse [61]. This step not only mitigates redundant cross-sequence attention computation but also breaks overlong packed sequences into shorter ones for more fine-grained pipeline, which benefits throughput and peak activation memory reduction.

The determination of chunk size involves a tradeoff between compute efficiency and padding reduction, as depicted in Figure 13. Smaller chunk size reduces padded tokens but risks underutilization and extra KV cache accesses; conversely, over-sized chunks hinder padding reduction and inflate stage latency. In practice, we set chunk size as the greatest power-of-2 divisor of all sequence lengths, with a minimum threshold (typically 64) to avoid underutilization.




Figure 13: Quantifying chunk-based alignment (1 task, 16-layer LLaMA7B, 4-GPU pipeline, sequence length 256).

4 Implementation

We build MuxTune with 14K LOCs of Python based on Megatron-LM [78] and PyTorch [28], and 2K LOCs of C++ and CUDA for kernel implementation. MuxTune supports deployment with hybrid parallelism, including pipeline (e.g., GPipe [49], 1F1B [65], interleaved-1F1B [67]), Megatron tensor parallelism [78], and data parallelism strategies (e.g., PyTorch DDP [59], FSDP [92]). MuxTune implements dynamic *Adapter* attachment to *BaseOp* via PyTorch *hook mechanism* [14], wrapping the logic of *Dispatch* and *Aggregate* sub-modules into the hooked *BaseOp* function at runtime.

Grouped Kernels. We implement MuxTune’s kernels based on NVIDIA CUTLASS and CuTe [81] for grouped computation across task adapters. In each kernel, we first assign thread blocks in proportion to the FLOPs and memory access of adapters for each task. Then, we decompose each adapter operator into tiles, assigning those that access the same partial of operator weight to the same allocated thread block. This implementation enhances load balancing across SMs while reducing memory read/write overhead.

Subgraph Execution. We use TorchFX [16] to trace the PEFT model into multi-task IR graphs and conduct subgraph-level orchestration (§3.4). Based on it, we adopt multiple streams and CUDA primitives (e.g., `Event.synchronize()`) to coordinate operator execution across independent tasks.

Offline Profiling and System Overhead. We conduct offline profiling across canonical operator configurations to build the cost model, given that hardware and backbones are pre-accessible. This is backed by PyTorch *dispatching mechanism* that ensures consistent kernel selection for identical input shapes, data types, and hardware [13]. We thus limit scheduling overhead to under 10 seconds by avoiding labor-intensive GPU operations, which is minimal compared to a fine-tuning task with a typical duration of 3-70 hours.

5 Evaluation

5.1 Experimental Setup

Testbeds. We evaluate MuxTune on three server setups: (1) Testbed-A: 1 node with 4 NVIDIA A40 GPUs (48GB), Intel




Figure 14: System throughput (in the number of processed tokens per second) across different global batch sizes, backbone models, and hardware configurations. Detailed workloads are presented above each subfigure (in grey blocks).

Table 1: Model configurations used in experiments. #GPUs denotes the number of GPUs for each model unless specified.

Model	#Layers	Hidden Dim	#Heads	#GPUs
GPT3-2.7B [32]	32	2560	32	2
LLaMA2-7B/13B [83]	32/40	4096/5120	32/40	4/8
OPT-30B [91]	48	7168	56	16

Xeon Silver 4310 CPU, and NVLink; (2) Testbed-B: 8 nodes, each with 2 NVIDIA A40 GPUs (48GB), Intel Xeon Gold 5318Y CPU, and Mellanox ConnectX-5 [22] (100Gb/s infiniband). (3) Testbed-C: 1 node with 8 NVIDIA H100 GPUs (80GB), Intel Xeon Platinum 8457C CPU, and NVLink.

Models and Datasets. We conduct experiments with four representative LLMs in Table 1. We have implemented three PEFT types: LoRA [47] (mainly used), Adapter Tuning [46] and Diff Pruning [41]. We use three datasets with varied sequence length, including SST2 [79], OpenBookQA [12], and RTE [15]. Since sequences of each task are padded to its maximum length, we analyze dataset distribution and pad (or truncate) sequences of SST2 to 64, QA to 128, and RTE to 256. In evaluation, we use two dataset combinations (§2.1): (1) *Uniform* shares the same dataset across tasks colocated on the same backbone; (2) *Non-uniform* adopts different datasets for these tasks. The testbed scales and workloads are sufficient to evaluate system performance, because PEFT consumes much less memory than pretraining while featuring limited input sizes (thus no large data parallelism is needed).

Baselines. We compare MuxTune with three baselines:

- (1) *HuggingFace PEFT* [18] (HF-PEFT): a user-friendly library for adapting various LLMs to downstream tasks, with support for memory optimizations like quantization.
- (2) *NeMo Megatron* [44] (NeMo): an AI framework built on Megatron-LM [78] that supports efficient kernels and scalable parallelism strategies to adapt LLMs with PEFT.
- (3) *SLoRA-PEFT* [76] (SL-PEFT): It adopts SLoRA’s techniques like backbone sharing and batching-only in PEFT, supporting fine-tuning multiple tasks concurrently.




Figure 15: Throughput on H100 GPUs across global batch sizes. Configurations are aligned with the above experiment.

Parallelism Selection. We grid-search the optimal parallelism for MuxTune and baselines across supported strategies (§4). For 2-GPU and 4-GPU cases, the search locates intra- and inter-stage parallelism, respectively. For more GPUs with heavier workloads, hybrid parallelism is selected (intra-stage parallelism within each node and pipeline across nodes).

5.2 End-to-End Performance

System Throughput. Figure 14 demonstrates the end-to-end throughput of MuxTune and baseline systems across various workloads. In the *Uniform* case, MuxTune improves throughput by up to 2.33 \times , 1.87 \times , and 1.64 \times over HF-PEFT, NeMo, and SL-PEFT, respectively. With lightweight workloads, MuxTune surpasses baselines by adaptively spatially fusing multiple tasks to improve utilization. The performance gains grow with more GPUs, benefited from the ability of MuxTune to overlap multi-task operators for stall reduction.

In the *Non-uniform* case, MuxTune achieves throughput improvements of 2.23 \times , 1.83 \times and 1.85 \times over the three baselines, respectively. For HF-PEFT and NeMo, the improvement remains consistent with *Uniform* case, as they execute tasks separately without extra zero-padded tokens. For SL-PEFT, MuxTune achieves higher improvement than in the *Uniform* case, because SL-PEFT incurs substantial zero-padded tokens that waste compute and memory resources.

Performance on Advanced Hardware. We further evaluate MuxTune against two baselines on H100 GPUs. As




Figure 16: Performance breakdown. TF, OO and CA represent task fusion, operator orchestration, and data alignment.

Table 2: Task workloads (WL) with random generated configurations (dataset, batch size) used in experiments.

Task Order	#1	#2	#3	#4	#5	#6	#7	#8
WL-A	SST2	QA	QA	SST2	SST2	SST2	QA	QA
WL-B	RTE	SST2	RTE	SST2	SST2	RTE	RTE	RTE
Batch Size	4	2	4	4	8	2	4	4


illustrated in Figure 15, MuxTune improves throughput by $5.29\times$ and $2.31\times$ over NeMo and SL-PEFT in the Uniform case (left), and $3.69\times$ and $1.94\times$ in the Non-uniform case (right). Beyond the above reasons, the more significant performance gains on H100 (compared to A40) stem from its superior compute power, which amplifies the underutilization inherent in single-task PEFT frameworks while unlocking more multi-task optimization potential for MuxTune.

5.3 Ablation Studies

Performance Breakdown. Figure 16 demonstrates the impact of each component in MuxTune using LLaMA7B, 4-GPU pipeline, and global batch size of 128. With lightweight workloads (Figure 16(a)), disabling three components undermines throughput by 36.1%, 30.3%, and 22.5%, respectively. This arises from improved GPU utilization (from computation fusion), reduced device stalls (from communication overlapping), and mitigated ineffective tokens across tasks (from chunk-based alignment). With heavier workloads (Figure 16(b)), data alignment increasingly dominates the overall throughput, causing a 34.3% decrease. Conversely, disabling task fusion only undermines throughput by 6.25% as GPU has been saturated. Since more micro-batches lead to fewer device stalls, disabling operator orchestration results in a 25.1% drop, slightly less than with lightweight workloads.


Memory Efficiency Analysis. We exclusively evaluate the memory efficiency of MuxTune against other baselines using two PEFT task workloads as listed in Table 2. As shown in Figure 17, we submit 32 PEFT tasks in a progressive manner, each with 1 micro-batch per iteration, by repeating the workloads four times. We employ the GPT2.7B backbone for WL-A and LLaMA7B for WL-B, which incur memory consumption of 5.2GB and 13.4GB, respectively.

In Figure 17(a), with tensor parallelism on 2 A40 GPUs (48GB each), MuxTune achieves up to $4.67\times$ and $1.44\times$



(a) GPT3-2.7B (2-GPU tensor parallel) (b) LLaMA2-7B (4-GPU pipeline)

Figure 17: Memory footprint with various number of tasks.



(a) NeMo (1 task) (b) Ours (4 tasks, no overlap) (c) Ours (4 tasks, overlap)

Figure 18: GPU and NVLink utilization of 1 layer with 4-GPU tensor parallelism. Tasks are interleaved in (b) and (c).

memory reduction over NeMo/HF-PEFT (OOM after 15 tasks) and SL-PEFT, respectively. Without memory constraints (i.e., scaling up to 32 tasks), MuxTune further reduces memory footprint by $5.29\times$ and $1.46\times$ compared to these baselines. The reasons are two-fold: (1) MuxTune flexibly shares the memory-intensive backbone across tasks, which consumes 13.4GB in contrast to 4.3GB of activation and 0.4GB of others (Figure 17(b)), while NeMo and HF-PEFT replicate one backbone per task; (2) MuxTune mitigates inter-task ineffective tokens and benefits fine-grained pipeline (§3.5), while SL-PEFT excessively batches tasks with substantial padding and higher peak activation memory. For larger backbone and more GPUs (Figure 17(b)), MuxTune reduces memory footprint by $3.57\times$ and $1.37\times$ over NeMo/HF-PEFT (OOM after 11 tasks) and SL-PEFT, demonstrating better scalability of MuxTune against baseline systems.

Efficiency of Operator Orchestration. We visualize the GPU compute and network bandwidth utilization of MuxTune with comparison to the NeMo baseline. Figure 18 uses Nsight toolkit [25] to profile the GPU and NVLink utilization under tensor parallelism. As shown in Figure 18(a), NeMo executes a single task with sequentially launched operators. Since computation is blocked by communication, the average GPU utilization remains at 82.5%, with a layer latency of 43.2ms. Figure 18(b) shows the results of 4 tasks with interleaved execution, but without overlap. The latency linearly increases to 172.5ms, while GPU utilization remains nearly constant at 84.7%. In Figure 18(c), MuxTune fully overlaps computation with communication. Without being blocked, the GPU utilization reaches 97.8%, achieving a $1.19\times$ improvement over the baseline, and the 4-task latency of the decoder layer is reduced to 156.2ms.

We also evaluate the end-to-end performance across different parallelism strategies and workloads, as shown in Fig-




Figure 19: Throughput of operator orchestration with varying number of tasks (LLaMA7B, sequence length 128, 64, 32). (a) 1 micro-batch of size 8, (b) 8 micro-batches of size 8.

ure 19 with only backbone sharing and operator orchestration enabled. Figure 19(a) illustrates the benefits of inter-task overlapping in tensor parallelism. With an increasing number of tasks, MuxTune delivers 1.20 \times , 1.22 \times , and 1.23 \times higher throughput than NeMo, owing to the overlap between computation and communication. Figure 19(b) demonstrates the effect of pipeline orchestration. Compared to NeMo, MuxTune improves throughput by 1.24 \times , 1.35 \times , and 1.36 \times , as the interleaved stage computations across tasks effectively mitigate pipeline bubbles. Notably, MuxTune is capable of achieving higher improvements (e.g., 1.59 \times with 4 micro-batches) with fewer micro-batches (more pipeline bubbles).

Effectiveness of Data Alignment. We assess chunk-based data alignment using the metric *effective throughput* [8], which measures the throughput of original tokens excluding inter-task zero-padded ones, reflecting the economic gains of service providers (§3.5). We use workloads in Table 2 to evaluate cases with and without intra-chunk zero-padding.

Figure 20 demonstrates the throughput of progressively adding tasks into a single hybrid task with one micro-batch. In Figure 20(a) with a chunk size of 64 (matching SST2), MuxTune avoids intra-chunk padding and achieves up to 2.33 \times higher throughput and 3.59 \times higher effective throughput than SL-PEFT. This is because MuxTune effectively mitigates inter-task padding without underutilizing GPUs. In Figure 20(b) with more inclined sequence lengths, MuxTune introduces intra-chunk zero-padding for SST2 tasks with a chunk size of 128. In this case, MuxTune still achieves up to 3.77 \times improvement on overall throughput, and 2.57 \times on effective throughput compared to SL-PEFT.

5.4 Scalability and Scheduling Study

Scalability Study. Figure 21(a) demonstrates the system scalability with two scaling strategies (LLaMA7B, global batch size 128, micro-batch size 8, n tasks for n -GPU): (1) “Up-only” (-UP): Scales up (i.e., increases allocated GPUs) for active instances as workloads increase; (2) “Up-then-out”: Scales up first, then scales out (i.e., replicates new instances) if workloads continue increasing. For the “up-only” strategy, despite sub-linear scaling (§3.3), MuxTune delivers 1.61 \times higher throughput than NeMo by efficiently improving utilization and mitigating device stalls. For the “up-then-out”




Figure 20: Throughput of 1 hybrid task with various number of tasks and datasets (LLaMA7B, 4-GPU pipeline). ZeroPad represents zero-padding all sequences as in SL-PEFT. -E denotes effective throughput, and overall if not marked.




Figure 21: System scalability of MuxTune and cluster-level performance under production-grade workloads.

strategy, both systems achieve near-linear scaling, while MuxTune still delivering up to 1.28 \times higher throughput.

Cluster-Level Performance. We further integrate MuxTune into cluster scheduling and evaluate it under production-grade workloads. In absence of public PEFT traces, we adapt a one-week Philly trace [51]. The average task duration and standard deviation are 372.6 min and 612.9 min, respectively, while the average arrival rate is 2.59 tasks/min. We replay PEFT workloads in a simulated 128-GPU cluster with a first-come, first-served scheduler, using both Uniform and Non-uniform combinations, LLaMA7B backbone, and randomly generated configurations for each task. As shown in Figure 21(b), in the Uniform case, MuxTune enhances cluster throughput by 1.61 \times , 1.51 \times , and 1.36 \times over HF-PEFT, NeMo, and SL-PEFT baselines, respectively. In the Non-uniform case, MuxTune further delivers 1.58 \times higher cluster throughput against SL-PEFT, which highlights the effectiveness of chunk-based data alignment when fusing multiple tasks under workloads with variable-length sequences.

6 Discussion and Future Work

Generality to Cluster Scheduling Policies. Beyond the first-come, first-served scheduler used in evaluation (§5.4), the system design of MuxTune is extensible as the fine-tuning backend for other cluster scheduling policies, such as budget-based [19], task priority-based [66], and SLO-aware scheduling [40]. For instance, with task priorities, the scheduler can colocate low-priority tasks to boost instance-level throughput while allocating dedicated resources for high-priority

ones to guarantee task-level latency. Moreover, tasks with the same backbone type could be colocated for backbone multiplexing while those with different types should be scheduled to different instances. We leave the in-depth exploration of “multiplexing-aware” scheduling as future work.

Extensibility to Performance Metric Optimizations. The optimizations of MuxTune focus on maximizing instance-level throughput by enhancing resource efficiency and reducing per-task memory consumption. Beyond that, MuxTune can achieve higher energy efficiency by mitigating wasted device stalls and lowering overall elapsed times of all colocated tasks [36]. Users can extend our techniques to optimize other performance metrics, such as integrating admission control (i.e., limiting the number of tasks dispatched to fine-tuning instances) to guarantee all colocated tasks can be completed within their user-specified SLOs. Moreover, users can adaptively scale the hardware frequencies while adhering to SLO requirements to further enhance energy efficiency [36, 71].

7 Related Work

Multi-Adapter LLM Systems. Recent years have witnessed significant progress in multi-adapter LLM systems [34, 38, 76, 89, 94]. PetS [94] proposes a unified framework for concurrent PEFT task serving. Punica [34], SLoRA [76], and dLoRA [89] focus on efficient scheduling and kernel implementation for multi-LoRA serving systems. Unlike these works, MuxTune targets multi-task execution optimization in PEFT scenarios, leveraging hierarchical co-scheduling to maximize resource utilization and minimize device stalls.

Techniques for Improving GPU Utilization. Several studies have explored GPU utilization optimizations through resource allocation and management across multiple ML applications, including temporal [5, 77] and spatial sharing [11, 21, 23] to multiplex the computational units. Other works [26, 58] horizontally or vertically fuse small kernels into larger ones to fully utilize hardware. However, they fail to be directly applied in PEFT, due to the memory-intensive backbone (which limits scalability), the lack of operator-level execution control (which incurs inter-task interference), and the interdependencies between adapters for the same task.

Frameworks for LLM Parallelization and Reducing Device Stalls. The field of LLM parallelization has been intensively studied in recent years, including pipeline [49, 65], tensor [78], and data parallelism [73, 74], as well as automatic parallelism optimizations [73, 93]. To resolve the issue of device stalls introduced by these parallelism strategies, DeepSeek-V3 [63] and ZeroBubble [70] split the backward pass to reduce pipeline bubbles. Overlapping-based methods [26, 86] decompose computations and overlap them with communication. TeraPipe [61] pipelines token-level computations to optimize long-sequence training. These techniques are unsuitable for PEFT, owing to the absence of weight gradient computation and the potential of GPU underutilization.

8 Conclusion

MuxTune is a resource-efficient system that optimizes concurrent multi-task PEFT execution in multi-tenant datacenters. The core of MuxTune is to multiplex the backbone across independent tasks in a spatial-temporal manner. MuxTune modularizes PEFT tasks for flexible backbone sharing, and devises hierarchical multi-task co-scheduling scheme with task, operator, and data-level optimizations to improve GPU utilization and reduce device stalls. Experimental results show that MuxTune achieves up to $2.33\times$ higher throughput and $5.29\times$ memory reduction compared to three baselines.

Acknowledgments

We would like to thank the anonymous reviewers and our shepherd, Jiarong Xing, for their valuable feedback. This work is supported by the National Natural Science Foundation of China (62232011) and the Natural Science Foundation of Shanghai Municipality (24ZR1430500). Quan Chen is the corresponding author of this paper (chen-quan@cs.sjtu.edu.cn).

References

- [1] Openai fine-tuning platform. <https://platform.openai.com/docs/guides/fine-tuning/>, 2023.
- [2] Openpipe fine-tuning quick start. <https://docs.openpipe.ai/features/fine-tuning/quick-start>, 2023.
- [3] Together ai fine-tuning cli. <https://docs.together.ai/docs/fine-tuning-cli>, 2023.
- [4] Train a large language model on a single amazon sagemaker gpu with hugging face and lora. <https://aws.amazon.com/cn/blogs/machine-learning/train-a-large-language-model-on-a-single-amazon-sagemaker-gpu-with-hugging-face-and-lora/>, 2023.
- [5] Alibaba cloud cgpu. <https://www.alibabacloud.com/help/en/elastic-gpu-service/latest/what-is-the-cgpu-service>, 2024.
- [6] Fine-tuning with the gemini api. <https://ai.google.dev/gemini-api/docs/model-tuning>, 2024.
- [7] Getting started with customizing a large language model (llm). <https://learn.microsoft.com/en-us/azure/ai-services/openai/concepts/customizing-llms>, 2024.
- [8] Goodput. <https://en.wikipedia.org/wiki/Goodput>, 2024.

- [9] Introducing vision to the fine-tuning api. <https://openai.com/index/introducing-vision-to-the-fine-tuning-api/>, 2024.
- [10] Nvidia apex extension. <https://nvidia.github.io/apex/>, 2024.
- [11] Nvidia multi-instance gpu. <https://www.nvidia.com/en-us/technologies/multi-instance-gpu/>, 2024.
- [12] Openbookqa dataset. <https://huggingface.co/datasets/allenai/openbookqa>, 2024.
- [13] Pytorch dispatcher. <http://blog.ezyang.com/2020/09/lets-talk-about-the-pytorch-dispatcher/>, 2024.
- [14] Pytorch hook mechanism. <https://pytorch.org/docs/stable/generated/torch.nn.Module.html>, 2024.
- [15] Recognizing textual entailment dataset. <https://huggingface.co/datasets/SetFit/rte>, 2024.
- [16] Torch fx. <https://pytorch.org/docs/stable/fx.html>, 2024.
- [17] Understanding custom llm models: A 2024 guide. <https://botpenguin.com/blogs/understanding-custom-llm-models>, 2024.
- [18] Huggingface peft library. <https://huggingface.co/docs/peft/index>, 2025.
- [19] Kubernetes. <https://kubernetes.io/>, 2025.
- [20] Nvidia cuda graph api. https://docs.nvidia.com/cuda/cuda-runtime-api/group__CUDA_RT__GRAPH.html, 2025.
- [21] Nvidia cuda stream management. https://docs.nvidia.com/cuda/cuda-runtime-api/group__CUDA_RT__STREAM.html, 2025.
- [22] Nvidia mellanox connectx-5. <https://www.nvidia.com/en-us/networking/ethernet/connectx-5/>, 2025.
- [23] Nvidia multi-process service. <https://docs.nvidia.com/deploy/mps/index.html>, 2025.
- [24] Nvidia nsight compute customization guide. <https://docs.nvidia.com/nsight-compute/CustomizationGuide/index.html>, 2025.
- [25] Nvidia nsight system. <https://developer.nvidia.com/nsight-systems>, 2025.
- [26] Nvidia transformer engine. <https://github.com/NVIDIA/TransformerEngine>, 2025.
- [27] Nvlink and nvswitch. <https://www.nvidia.com/en-us/data-center/nvlink/>, 2025.
- [28] Pytorch foundation. <https://pytorch.org>, 2025.
- [29] Amey Agrawal, Nitin Kedia, Ashish Panwar, Jayashree Mohan, Nipun Kwatra, Bhargav Gulavani, Alexey Tumanov, and Ramachandran Ramjee. Taming Throughput-Latency tradeoff in LLM inference with Sarathi-Serve. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, pages 117–134, Santa Clara, CA, July 2024. USENIX Association.
- [30] Reza Yazdani Aminabadi, Samyam Rajbhandari, Ammar Ahmad Awan, Cheng Li, Du Li, Elton Zheng, Olatunji Ruwase, Shaden Smith, Minjia Zhang, Jeff Rasley, et al. Deepspeed-inference: enabling efficient inference of transformer models at unprecedented scale. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15. IEEE, 2022.
- [31] Yushi Bai, Xin Lv, Jiajie Zhang, Yuze He, Ji Qi, Lei Hou, Jie Tang, Yuxiao Dong, and Juanzi Li. Longalign: A recipe for long context alignment of large language models. *arXiv preprint arXiv:2401.18058*, 2024.
- [32] 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.
- [33] Chang Chen, Xiuhong Li, Qianchao Zhu, Jiangfei Duan, Peng Sun, Xingcheng Zhang, and Chao Yang. Centauri: Enabling efficient scheduling for communication-computation overlap in large model training via communication partitioning. In *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, pages 178–191, 2024.
- [34] Lequn Chen, Zihao Ye, Yongji Wu, Danyang Zhuo, Luis Ceze, and Arvind Krishnamurthy. Punica: Multi-tenant lora serving. *Proceedings of Machine Learning and Systems*, 6:1–13, 2024.
- [35] Marcin Chrapek, Anjo Vahldiek-Oberwagner, Marcin Spoczynski, Scott Constable, Mona Vij, and Torsten Hoefler. Fortify your foundations: Practical privacy and security for foundation model deployments in the cloud, 2024.

- [36] Jae-Won Chung, Yile Gu, Insu Jang, Luoxi Meng, Nikhil Bansal, and Mosharaf Chowdhury. Reducing energy bloat in large model training. In *Proceedings of the ACM SIGOPS 30th Symposium on Operating Systems Principles*, SOSP '24, page 144–159, New York, NY, USA, 2024. Association for Computing Machinery.
- [37] 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.
- [38] Wei Gao, Weiming Zhuang, Minghao Li, Peng Sun, Yonggang Wen, and Tianwei Zhang. Ymir: A scheduler for foundation model fine-tuning workloads in datacenters. In *Proceedings of the 38th ACM International Conference on Supercomputing*, pages 259–271, 2024.
- [39] Stephen C Graves, Harlan C Meal, Daniel Stefek, and Abdel Hamid Zeghmi. Scheduling of re-entrant flow shops. *Journal of operations management*, 3(4):197–207, 1983.
- [40] Diandian Gu, Yihao Zhao, Yinmin Zhong, Yifan Xiong, Zhenhua Han, Peng Cheng, Fan Yang, Gang Huang, Xin Jin, and Xuanzhe Liu. Elasticflow: An elastic serverless training platform for distributed deep learning. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 266–280, 2023.
- [41] Demi Guo, Alexander M Rush, and Yoon Kim. Parameter-efficient transfer learning with diff pruning. *arXiv preprint arXiv:2012.07463*, 2020.
- [42] Avelina Asada Hadji-Kyriacou and Ognjen Arandjelovic. Context-peft: Efficient multi-modal, multi-task fine-tuning. *arXiv preprint arXiv:2312.08900*, 2023.
- [43] Zeyu Han, Chao Gao, Jinyang Liu, Sai Qian Zhang, et al. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*, 2024.
- [44] Eric Harper, Somshubra Majumdar, Oleksii Kuchaiev, Li Jason, Yang Zhang, Evelina Bakhturina, Vahid Noroozi, Sandeep Subramanian, Koluguri Nithin, Huang Jocelyn, Fei Jia, Jagadeesh Balam, Xuesong Yang, Micha Livne, Yi Dong, Sean Naren, and Boris Ginsburg. NeMo: a toolkit for Conversational AI and Large Language Models.
- [45] John L Hennessy and David A Patterson. *Computer architecture: a quantitative approach*. Morgan kaufmann, 2017.
- [46] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pages 2790–2799. PMLR, 2019.
- [47] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [48] Qinghao Hu, Meng Zhang, Peng Sun, Yonggang Wen, and Tianwei Zhang. Lucid: A non-intrusive, scalable and interpretable scheduler for deep learning training jobs. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2*, pages 457–472, 2023.
- [49] Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Chen, HyoukJoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, et al. Gpipe: Efficient training of giant neural networks using pipeline parallelism. *Advances in neural information processing systems*, 32, 2019.
- [50] Mathav Raj J, Kushala VM, Harikrishna Warriar, and Yogesh Gupta. Fine tuning llm for enterprise: Practical guidelines and recommendations, 2024.
- [51] Myeongjae Jeon, Shivaram Venkataraman, Amar Phanishayee, Junjie Qian, Wencong Xiao, and Fan Yang. Analysis of large-scale multi-tenant gpu clusters for dnn training workloads. In *USENIX Annual Technical Conference*, pages 947–960, 2019.
- [52] Chenyu Jiang, Zhen Jia, Shuai Zheng, Yida Wang, and Chuan Wu. Dynapipeline: Optimizing multi-task training through dynamic pipelines. In *Proceedings of the Nineteenth European Conference on Computer Systems*, pages 542–559, 2024.
- [53] Ziheng Jiang, Haibin Lin, Yinmin Zhong, Qi Huang, Yangrui Chen, Zhi Zhang, Yanghua Peng, Xiang Li, Cong Xie, Shibiao Nong, Yulu Jia, Sun He, Hongmin Chen, Zhihao Bai, Qi Hou, Shipeng Yan, Ding Zhou, Yiyao Sheng, Zhuo Jiang, Haohan Xu, Haoran Wei, Zhang Zhang, Pengfei Nie, Leqi Zou, Sida Zhao, Liang Xiang, Zherui Liu, Zhe Li, Xiaoying Jia, Jianxi Ye, Xin Jin, and Xin Liu. Megascall: Scaling large language model training to more than 10,000 gpus, 2024.
- [54] A. B. Kahn. Topological sorting of large networks. *Commun. ACM*, 5(11):558–562, November 1962.

- [55] Richard M Karp. Reducibility among combinatorial problems. In *50 Years of Integer Programming 1958-2008: from the Early Years to the State-of-the-Art*, pages 219–241. Springer, 2009.
- [56] Vijay Korthikanti, Jared Casper, Sangkug Lym, Lawrence McAfee, Michael Andersch, Mohammad Shoeybi, and Bryan Catanzaro. Reducing activation recomputation in large transformer models, 2022.
- [57] 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.
- [58] Ao Li, Bojian Zheng, Gennady Pekhimenko, and Fan Long. Automatic horizontal fusion for gpu kernels. In *2022 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, pages 14–27. IEEE, 2022.
- [59] Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, et al. Pytorch distributed: Experiences on accelerating data parallel training. *arXiv preprint arXiv:2006.15704*, 2020.
- [60] Xiang Lisa Li and Percy Liang. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- [61] Zhuohan Li, Siyuan Zhuang, Shiyuan Guo, Danyang Zhuo, Hao Zhang, Dawn Song, and Ion Stoica. Terapipe: Token-level pipeline parallelism for training large-scale language models. In *International Conference on Machine Learning*, pages 6543–6552. PMLR, 2021.
- [62] Vladislav Lialin, Vijeta Deshpande, and Anna Rumshisky. Scaling down to scale up: A guide to parameter-efficient fine-tuning. *arXiv preprint arXiv:2303.15647*, 2023.
- [63] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*, 2024.
- [64] Yongkang Liu, Yiqun Zhang, Qian Li, Shi Feng, Daling Wang, Yifei Zhang, and Hinrich Schütze. Hift: A hierarchical full parameter fine-tuning strategy. *arXiv preprint arXiv:2401.15207*, 2024.
- [65] Deepak Narayanan, Aaron Harlap, Amar Phanishayee, Vivek Seshadri, Nikhil R Devanur, Gregory R Ganger, Phillip B Gibbons, and Matei Zaharia. Pipedream: Generalized pipeline parallelism for dnn training. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*, pages 1–15, 2019.
- [66] Deepak Narayanan, Keshav Santhanam, Fiodar Kazhamiaka, Amar Phanishayee, and Matei Zaharia. Heterogeneity-aware cluster scheduling policies for deep learning workloads. In *Proceedings of the 14th USENIX Conference on Operating Systems Design and Implementation*, pages 481–498, 2020.
- [67] Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, et al. Efficient large-scale language model training on gpu clusters using megatron-lm. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–15, 2021.
- [68] Pál András Papp, Georg Anegg, and Albert-Jan N Yzelman. Dag scheduling in the bsp model. In *International Conference on Current Trends in Theory and Practice of Computer Science*, pages 238–253. Springer, 2025.
- [69] Jonas Pfeiffer, Aishwarya Kamath, Andreas Rücklé, Kyunghyun Cho, and Iryna Gurevych. Adapterfusion: Non-destructive task composition for transfer learning. *arXiv preprint arXiv:2005.00247*, 2020.
- [70] Penghui Qi, Xinyi Wan, Guangxing Huang, and Min Lin. Zero bubble pipeline parallelism. *arXiv preprint arXiv:2401.10241*, 2023.
- [71] Haoran Qiu, Weichao Mao, Archit Patke, Shengkun Cui, Saurabh Jha, Chen Wang, Hubertus Franke, Zbigniew Kalbarczyk, Tamer Başar, and Ravishankar K. Iyer. Power-aware deep learning model serving with μ -Serve. In *2024 USENIX Annual Technical Conference (USENIX ATC 24)*, pages 75–93, Santa Clara, CA, July 2024. USENIX Association.
- [72] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- [73] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE, 2020.
- [74] 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 *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pages 551–564, 2021.

- [75] Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, et al. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*, 2023.
- [76] Ying Sheng, Shiyi Cao, Dacheng Li, Coleman Hooper, Nicholas Lee, Shuo Yang, Christopher Chou, Banghua Zhu, Lianmin Zheng, Kurt Keutzer, et al. S-lora: Serving thousands of concurrent lora adapters. *arXiv preprint arXiv:2311.03285*, 2023.
- [77] Lin Shi, Hao Chen, Jianhua Sun, and Kenli Li. vcuda: Gpu-accelerated high-performance computing in virtual machines. *IEEE Transactions on Computers*, 61(6):804–816, 2011.
- [78] Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*, 2019.
- [79] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- [80] Andreas Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers, 2022.
- [81] Vijay Thakkar, Pradeep Ramani, Cris Cecka, Aniket Shivam, Honghao Lu, Ethan Yan, Jack Kosaian, Mark Hoemmen, Haicheng Wu, Andrew Kerr, Matt Nicely, Duane Merrill, Dustyn Blasig, Fengqi Qiao, Piotr Majcher, Paul Springer, Markus Hohnerbach, Jin Wang, and Manish Gupta. CUTLASS, January 2023.
- [82] Nirmalya Thakur, Vanessa Su, Mingchen Shao, Keshava A. Patel, Hongseok Jeong, Victoria Knieling, and Andrew Bian. A labelled dataset for sentiment analysis of videos on youtube, tiktok, and other sources about the 2024 outbreak of measles, 2024.
- [83] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutit Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [84] 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.
- [85] Borui Wan, Mingji Han, Yiyao Sheng, Zhichao Lai, Mofan Zhang, Junda Zhang, Yanghua Peng, Haibin Lin, Xin Liu, and Chuan Wu. Bytecheckpoint: A unified checkpointing system for llm development. *arXiv preprint arXiv:2407.20143*, 2024.
- [86] Shibo Wang, Jinliang Wei, Amit Sabne, Andy Davis, Berkin Ilbeyi, Blake Hechtman, Dehao Chen, Karthik Srinivasa Murthy, Marcello Maggioni, Qiao Zhang, et al. Overlap communication with dependent computation via decomposition in large deep learning models. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1*, pages 93–106, 2022.
- [87] Yu Wang. A tutorial on the pretrain-finetune paradigm for natural language processing. *arXiv preprint arXiv:2403.02504*, 2024.
- [88] Ziwen Wang, Jie Wang, Haiqian Gu, Fei Su, and Bojin Zhuang. *Automatic Conditional Generation of Personalized Social Media Short Texts*, page 56–63. Springer International Publishing, 2018.
- [89] Bingyang Wu, Ruidong Zhu, Zili Zhang, Peng Sun, Xuanzhe Liu, and Xin Jin. {dLoRA}: Dynamically orchestrating requests and adapters for {LoRA}{LLM} serving. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*, pages 911–927, 2024.
- [90] Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. Orca: A distributed serving system for {Transformer-Based} generative models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 521–538, 2022.
- [91] Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*, 2022.
- [92] Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, et al. Pytorch fsdp: experiences on scaling fully sharded data parallel. *arXiv preprint arXiv:2304.11277*, 2023.
- [93] Lianmin Zheng, Zhuohan Li, Hao Zhang, Yonghao Zhuang, Zhifeng Chen, Yanping Huang, Yida Wang, Yanzhong Xu, Danyang Zhuo, Eric P Xing, et al. Alpa: Automating inter-and intra-operator parallelism for distributed deep learning. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*, pages 559–578, 2022.

- [94] Zhe Zhou, Xuechao Wei, Jiejing Zhang, and Guangyu Sun. {PetS}: A unified framework for {Parameter-Efficient} transformers serving. In *2022 USENIX Annual Technical Conference (USENIX ATC 22)*, pages 489–504, 2022.
- [95] Kan Zhu, Yilong Zhao, Liangyu Zhao, Gefei Zuo, Yile Gu, Dedong Xie, Yufei Gao, Qinyu Xu, Tian Tang, Zihao Ye, et al. Nanoflow: Towards optimal large language model serving throughput. *arXiv preprint arXiv:2408.12757*, 2024.
- [96] Yonghao Zhuang, Hexu Zhao, Lianmin Zheng, Zhuohan Li, Eric P. Xing, Qirong Ho, Joseph E. Gonzalez, Ion Stoica, and Hao Zhang. On optimizing the communication of model parallelism, 2024.

A Optimality Analysis for Structured Pipeline Template

Theoretically proving the optimality of a pipeline schedule is non-trivial, not to mention that micro-batches are heterogeneous, i.e., with different micro-batch sizes and sequence lengths. Exhaustively enumerating all possible candidate pipeline schedules can be formulated as an integer linear programming (ILP) problem with $S \sum_i n_i$ constraints, which is proven to be NP-hard [55] with the time complexity of $e^{O(S \sum_i n_i)}$, where n_i is the number of micro batches of task i . However, in our PEFT scenarios, micro-batches of each hTask bucket retain a consistent shape, while its forward and backward passes share the same latency for each pipeline stage (§3.4.1). These characteristics offer us opportunities to reduce the complexity of the problem. Below, we theoretically discuss the optimality of our proposed pipeline template.

Our discussion focuses on how to achieve multi-task pipeline optimality (i.e., minimized end-to-end latency) in the paradigm of 1F1B [65], as it is one of the most widely used and efficient pipeline. The execution of 1F1B pipeline can be divided into three phases, including warm-up, steady, and drain phase. In multi-task 1F1B pipeline, we follow [52, 65] to first introduce a basic lemma:

Lemma 1 *The end-to-end pipeline latency is calculated by adding the latencies of three phases (T_{warm} , T_{steady} , T_{drain}), while T_{warm}/T_{drain} can be calculated as: $(S-1)t_1/(S-1)t_P$.*

In the above lemma, C is the number of micro-batches, P is the number of hybrid task (hTask) buckets, and S is the number of pipeline stages. t_i denotes the stage latency of the micro-batches for the i -th bucket. t_1 and t_P represent the forward (backward) stage latency of the first and last sorted hTask bucket, respectively. Then, as observed in Figure 22(b) and (c), we give the next lemma as follows:

Lemma 2 *In steady phase, all micro-batches are bound to go through the last stage for one forward and one backward pass, i.e., $T_{steady} \geq 2C \sum_{i=1}^P t_i$.*

After that, we can measure the latency ratio between T_{steady} and $T_{warmup} + T_{drain}$ as:

$$\frac{T_{steady}}{T_{warmup} + T_{drain}} \geq \frac{2C \sum_{i=1}^P t_i}{(S-1)(t_1 + t_P)}. \quad (8)$$

Since we construct hTask buckets in a workload-balanced manner (i.e., minimizing inter-bucket variance), here we simplify the equation by assuming $t_i = t_{i-1} = t$, $\forall i \in [2, P]$:

$$\frac{T_{steady}}{T_{warmup} + T_{drain}} \geq \frac{2C P t}{2(S-1)t} = \frac{CP}{S-1}. \quad (9)$$

In common practice, the number of micro-batches is set to much larger than S to reduce the pipeline bubble ratio (e.g.,

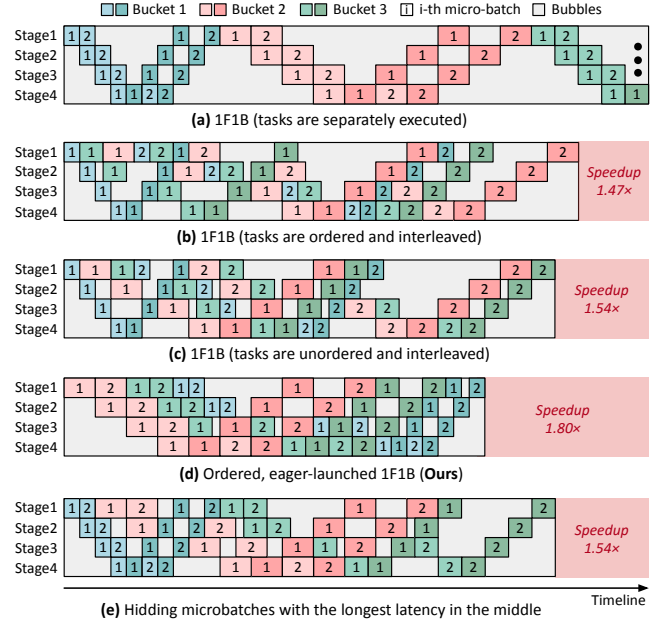


Figure 22: Various multi-task 1F1B pipeline schedules.

$4 \times$ in GPipe [49] and Alpa [93]). Besides, in our scenario of multi-task backbone sharing, the number of hTask buckets is also typically much larger than S . For example, in the “Memory Efficiency Analysis” of §5.3, 4 A40 GPUs can accommodate more than 32 PEFT tasks co-located on a single LLaMA7B backbone with 4-GPU pipeline. That is, the number of hTask buckets P can be set to up to 32 when $S = 4$. As a result, we give the theorem as follows:

Theorem 1 *In multi-task PEFT scenario, the latency of steady phase (T_{steady}) typically dominate the end-to-end pipeline latency.*

Therefore, to prove the near optimality of our proposed pipeline template, **we only need to prove that T_{steady} is minimized to $2C \sum_{i=1}^P t_i$, i.e., no pipeline internal bubbles exist in the last stage, as shown in Figure 22(d).**

To prove this statement, we notice that our template generation requires sorting hTask buckets $\mathcal{G} = \{\mathcal{G}_1, \mathcal{G}_2, \dots, \mathcal{G}_P\}$ by their latencies in descending order (micro-batches of the same bucket are consecutive), while advocating eagerly launching as many as micro-batches as possible. In this context, we give the third lemma as follows:

Lemma 3 *Before the backward pass of \mathcal{G}_j ’s last micro-batch in the last stage completes, the forward pass of \mathcal{G}_{j+1} ’s first micro-batch in the second-to-last stage always have been ready.*

As shown in Figure 22(d), this lemma is justified because: (1) forward and backward passes share identical latency, (2) the stage latency of \mathcal{G}_{j+1} is always shorter than that of \mathcal{G}_j , and (3) 1F1B pipeline always prioritizes the ready backward passes. With this lemma, we give the following theorem:

Theorem 2 *In our proposed pipeline template, once the first micro-batch begins its forward pass in the last stage, the last stage would “keep busy” until the backward pass of the last micro-batch completes.*

Therefore, **we have completed the near optimality proof for our proposed pipeline template.** It should be noted that the optimality proof is built on the premise that *device memory is always sufficient for our micro-batch eager launching scheme*, since the memory footprint is greatly reduce in multi-task PEFT scenarios (§5.3). In practice, we implement our pipeline template generation with memory limitations, i.e., delaying micro-batch launching if our memory cost model reports that at least one stage would incur OOM issues.

In Figure 22(e), we further illustrate the effect of hiding the longest micro-batches in the middle (rather than sorting in descending order). As observed, despite the reduction of T_{warmup} and T_{drain} , such modification disrupts Theorem 2 and thus leads to worse end-to-end pipeline latency.