

# ANGEL-PTM: A Scalable and Economical Large-scale Pre-training System in Tencent

Xiaonan Nie<sup>†</sup>  
Peking University  
xiaonan.nie@pku.edu.cn

Jinbao Xue  
Tencent Inc.  
jinbaoxue@tencent.com

Yangyu Tao  
Tencent Inc.  
brucetao@tencent.com

Yi Liu  
Tencent Inc.  
callbackliu@tencent.com

Dian Jiao  
Tencent Inc.  
focusjiao@tencent.com

Bin Cui<sup>†‡</sup>  
Peking University  
bin.cui@pku.edu.cn

Fangcheng Fu<sup>†</sup>  
Peking University  
ccchengff@pku.edu.cn

Xupeng Miao  
Carnegie Mellon University  
xupeng@cmu.edu

## ABSTRACT

Recent years have witnessed the unprecedented achievements of large-scale pre-trained models, especially Transformer models. Many products and services in Tencent Inc., such as WeChat, QQ, and Tencent Advertisement, have been opted in to gain the power of pre-trained models. In this work, we present ANGEL-PTM, a productive deep learning system designed for pre-training and fine-tuning Transformer models. ANGEL-PTM can train extremely large-scale models with hierarchical memory efficiently. The key designs of ANGEL-PTM are a fine-grained memory management via the *Page* abstraction and a unified scheduling method that coordinates computations, data movements, and communications. Furthermore, ANGEL-PTM supports extreme model scaling with SSD storage and implements a lock-free updating mechanism to address the SSD I/O bottlenecks. Experimental results demonstrate that ANGEL-PTM outperforms existing systems by up to 114.8% in terms of maximum model scale as well as up to 88.9% in terms of training throughput. Additionally, experiments on GPT3-175B and T5-MoE-1.2T models utilizing hundreds of GPUs verify our strong scalability.

### PVLDB Reference Format:

Xiaonan Nie, Yi Liu, Fangcheng Fu, Jinbao Xue, Dian Jiao, Xupeng Miao, Yangyu Tao, Bin Cui. ANGEL-PTM: A Scalable and Economical Large-scale Pre-training System in Tencent. PVLDB, 16(12): 3781-3794, 2023.  
doi:10.14778/3611540.3611564

## 1 INTRODUCTION

Large-scale pre-trained models, such as Transformer models, have achieved remarkable advancements in various fields such as computer vision [13, 32, 51], natural language processing [8, 11, 45, 46],

speech recognition [12, 60], and generative AI [1, 52] in recent years, outperforming traditional ML models and becoming the SOTA approach. For example, Chat-GPT [1] is capable of generating human-like text and performing various NLP tasks with impressive accuracy. The success of Transformer models can be attributed to their ability to automatically learn and extract hierarchical representations of data, making them highly suited for complex tasks [7].

With the hope of achieving better performance, efforts are made to increase the scale of models — the flagship NLP model size has been increasing at a rate of 240× for every 2 years [20], and Kaplan et al. [28] suggested that the trend of increasing model size will continue for better model quality. Such an explosive growth in model scale inevitably increases the computational and memory cost, and training large-scale Transformer models becomes extremely expensive. For instance, Microsoft totally adopts 4480 A100 80G GPUs for training Megatron-Turing NLG 530B [53], which would cost almost 70 million dollars for only purchasing these computing resources.

Undoubtedly, in order to simultaneously enable the superior ability of Transformer models in real-world applications and meet the rapid evolution of model scales, it is necessary for companies to re-think and re-design the productive deep learning systems in this era. Regarding the use cases and demands in Tencent Inc., we would like identify two key characteristics of deep learning systems.

- **Easy-to-use and easy-to-scale.** In real-world productive applications, most users are deep learning researchers or data scientists that are good at designing task-specific model architectures. In contrast, they usually lack the expert knowledge or experiences in deploying and accelerating the model training process in the distributed manner. Consequently, the deep learning system should require only a few lines' modifications to parallelize the training tasks. Moreover, since productive clusters are multi-tenant by nature, the available hardware resources would be varying. Thus, we seek for seamless scalability. In other words, when users wish to tune the amount of resources for their tasks, there should be no need to re-configure their parallel schemes.
- **Efficient and cost-effective.** Due to the stunning scale of models and datasets, training large-scale models is extremely time-consuming. Therefore, how to achieve a better training efficiency

<sup>†</sup>School of Computer Science & Key Lab of High Confidence Software Technologies (MOE), Peking University,

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

Proceedings of the VLDB Endowment, Vol. 16, No. 12 ISSN 2150-8097.  
doi:10.14778/3611540.3611564

is vital to the practical deployment of pre-trained models. In addition to the running time, we should also improve the hardware usage ratio with best efforts. It is undoubtedly that it makes pre-trained models more economical if we can support the efficient training of large-scale models using as few resources as possible.

We observe that existing deep learning systems, such as Megatron-LM [38] and DeepSpeed [47, 50], fail to achieve these characteristics. First, many systems involve complex parallelism strategies (such as tensor parallelism and pipeline parallelism) to scale the training tasks, which require the users to have a thorough understanding about how to split the models across the available accelerators (e.g., GPU devices) to optimize the distributed computing efficiency. This incurs extra efforts to deploy and scale the training tasks. Second, we find that existing systems adopt a coarse memory management to allocate and release the memory during training, which makes them unavoidably suffer from memory fragments and low resource usage ratio when training large-scale Transformer models, as we will analyze more in-depth in Section 3.

In this work, we develop ANGEL-PTM, a brand new deep learning system to support the booming applications of Transformer models in Tencent. The main contributions are summarized as follows:

- We analyze the characteristics and requirements of large-scale model training tasks in Tencent and propose the underlying designs of ANGEL-PTM to address these requirements, which integrates data parallelism, parameter sharding, and hierarchical memory to gain the convenience of use and the transparency of scaling to various numbers of GPUs.
- To reduce the memory fragments and fully utilize the memory and bandwidth, we propose the *fine-grained Page abstraction* and manage the model states at the page level, including allocation, release, movement, and communication. Furthermore, we design the unified scheduler together with an *fine-grained life-time based scheduling* method to dynamically manage these operations in a holistic manner for efficient training.
- To support enlarging models to an extreme scale, we integrate the SSD storage and design the *Lock-Free Updating Mechanism* to eliminate the bottleneck of SSD I/O bandwidth.
- We have conducted evaluations on various representative large-scale Transformer models. Results show that ANGEL-PTM achieves up to 114.8% improvement in maximum supported model scale and up to 88.9% improvement in throughput performance compared to existing systems. Experiments also verify the near-linear scalability of ANGEL-PTM when training on hundreds of GPUs.

ANGEL-PTM has been deployed in Tencent for around one year, facilitating the training of foundational models used across a diverse spectrum of products and services, including WeChat, QQ, Tencent Games, Tencent Advertising, and Tencent Cloud. In addition, ANGEL-PTM has also contributed to the training of the HunYuan series models, notably aiding the HunYuan-1T model in securing first place in the overall rankings of the CLUE benchmark [2].

## 2 BACKGROUND

### 2.1 Memory Management in Deep Learning

As the size of the models increases, GPU memory management becomes a critical factor that affects the performance and scalability

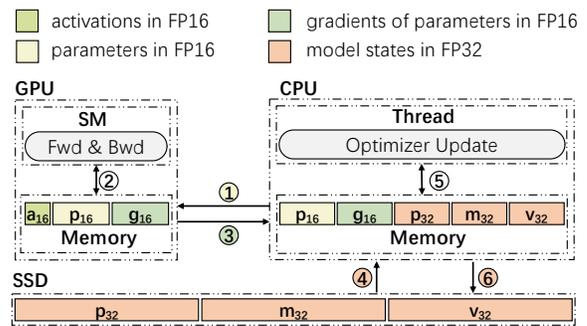


Figure 1: The workflow of training on hierarchical memory.

of deep learning models [8, 28, 39]. In this section, we will provide a brief overview of the memory consumption during training and how they are managed in current deep learning frameworks.

**Deep Learning Training.** The deep learning training can be represented as a computation graph, where each node stands for an operation, such as matrix multiply, and each edge is a tensor or dependency [4]. To achieve a satisfactory model quality, the training involves numerous forward and backward propagation passes. During the forward pass, training data is fed through the computation graph, utilizing the model parameters to produce each layer’s activation. The final outputs are compared to the expected values using a loss function. During the backward pass, the error values are propagated back to compute the gradients of activations and parameters, respectively, where the gradients of parameters are further used by the optimizer to update the model parameters. In summary, the memory during training is primarily consumed by parameters and their gradients, activations and their gradients, and the optimizer states. Among them, the parameters and optimizer states will be preserved during training, while activations, gradients of activations, and gradients of parameters will be dynamically generated and released. In the rest of this work, “model states” is used to denote parameters and optimizer states for simplicity.

**Mixed Precision Training.** To reduce the computation and memory requirements without sacrificing model quality, Micikevicius et al. [37] proposed the *mixed precision* techniques for training. As shown in Figure 1, the parameters are cast to the half-precision format (i.e., FP16 or the BF16 variant) before computation, so that activations and both types of gradients will be calculated in the FP16 fashion. Meanwhile, the model states are stored in the single-precision format (i.e., FP32) to preserve model quality. With the rapid growth of model size, *mixed precision* has become a de facto paradigm for large-scale model training and deployment.

**Memory Management in Existing Frameworks.** Existing deep learning frameworks (e.g., PyTorch [43], TensorFlow [4], and Hetu [34]) employ a general memory management method, which manages the GPU memory in a separate memory pool that responds to the executor’s requests, such as allocation and de-allocation. For example, TensorFlow [4] utilizes the best-fit allocation algorithm, i.e., BFC, to manage GPU memory and minimize memory fragmentation caused by frequent allocations and de-allocations. By allocating only the necessary amount of memory, this algorithm tries best to reduce the waste in GPU memory. Compared to other

allocation algorithms, the BFC algorithm may take longer to find an available block, but it is well-suited for systems with limited GPU memory. Additionally, Chen et al. [9] proposes the recomputation technique for memory savings, which releases partial activations in the forward pass and regenerates them with extra computation in the backward pass.

**Hierarchical Memory for Training.** To accommodate the memory demands of training large models, many frameworks attempt to incorporate the hierarchical memory within GPU servers. To be specific, to address the memory consumption of Transformer-based pre-trained models, researchers have proposed shifting the optimizer states and computations to the CPU memory and the SSD storage [14, 49, 50]. We illustrate the workflow of training on hierarchical memory in Figure 1. The GPU (1) fetches the parameters from the CPU, (2) performs forward and backward computations on the GPU, and then (3) sends the calculated gradients back to the CPU. The CPU (4) loads optimizer states from the SSD storage, (5) performs optimizer updating on CPU, and (6) stores the optimizer states on the SSD storage.

## 2.2 Memory Footprints of Transformer

A Transformer layer is stacked by a self-attention network and a position-wise feed-forward network (FFN), and it employs a residual connection on each of these two sub-layers, followed by a normalization layer [6]. In the following, we will approximately formulate the memory footprints of each component within a Transformer layer under the popular mixed precision training with the Adam optimizer scenario, where the input data is  $X \in \mathbb{R}^{b \times s \times d_m}$ . Specifically,  $b$  is batch size,  $s$  is sequence length,  $d_m$  is hidden size of embeddings and  $d_{ffn}$  is hidden size of FFN. Results of footprints are summarized in Table 1, and we ignore the small tensors for simplicity when calculating the total size, such as *params.* of LayerNorm.

**Self-Attention.** The attention block [54] could capture the dependencies between tokens in the sequence, and is effective in sequence modeling. As shown in Equation 1, it first linearly projects the input  $X$  into queries ( $Q$ ), keys ( $K$ ) and values ( $V$ ) with three linear functions respectively, where  $\{W_Q, W_K, W_V\} \in \mathbb{R}^{d_m \times d_m}$  and the total footprint of *Params* in this layer is  $3(Q, K, V) \times 2(\text{forward and backward}) \times 2(\text{FP16, 2 bytes}) \times d_m \times d_m = 12d_m^2$ . Similarly, the *Acts* is  $\{Q, K, V\} \in \mathbb{R}^{b \times s \times d_m}$  and their footprint is  $3(Q, K, V) \times 2(\text{forward and backward}) \times 2(\text{FP16, 2 bytes}) \times b \times s \times d_m = 12bsd_m$ . The footprint of model states (*Optims*) is  $3(Q, K, V) \times 3(\text{master parameter, momentum, variance}) \times 4(\text{FP32, 4 bytes}) \times d_m \times d_m = 12d_m^2$ . The same calculation method is also applicable to other layers, we will directly give the results in the table 1 because of the limited space. After this linear layer, it will compute the attention scores (shape:  $b \times s$ ) between each query-key ( $Q$ - $K$ ) pair, which is obtained by performing the dot-product (MatMul) operation as well as the ScaledMaskSoftmax operation. And the three operations, including Scale, Mask(opt.) and Softmax, are always fused into the ScaledMaskSoftmax operation for time-efficiency and memory-saving via kernel fusion techniques [56]. The attention vectors (shape:  $b \times s \times d_m$ ) are computed by weighted summation between the attention scores and values  $V$ . The self-attention layer finally produces the output (shape:  $b \times s \times d_m$ ) by applying a linear transformation ( $W \in \mathbb{R}^{d_m \times d_m}$ ) to

**Table 1: Memory footprints of a single Transformer layer under the mixed-precision training with Adam optimizer.**

Block	Layer Name	Parameters. (B)	Activations. (B)	Optimizers. (B)
Attn	Linear(Q,K,V)	$12d_m^2$	$12bsd_m$	$36d_m^2$
	MatMul	-	$4bs^2$	-
	Scaled-Softmax	-	$4bs^2$	-
	MatMul	-	$4bsd_m$	-
	Linear	$4d_m^2$	$4bsd_m$	$12d_m^2$
	Add	-	$4bsd_m$	-
	LayerNorm	$4d_m$	$4bsd_m$	$12d_m$
FFN	Linear	$4d_m d_{ffn}$	$4bsd_{ffn}$	$12d_m d_{ffn}$
	GeLU	-	$4bsd_{ffn}$	-
	Linear	$4d_m d_{ffn}$	$4bsd_m$	$12d_m d_{ffn}$
	Add	-	$4bsd_m$	-
	LayerNorm	$4d_m$	$4bsd_m$	$12d_m$
Total		$16d_m^2 + 8d_m d_{ffn}$	$40bsd_m + 8bs^2 + 8bsd_{ffn}$	$48d_m^2 + 24d_m d_{ffn}$

the attention vectors.

$$\text{Attn}(X) = \text{Softmax} \left( \text{Mask} \left( \frac{(XW_Q)(XW_K)^T}{\sqrt{d_k}} \right) \right) (XW_V) \quad (1)$$

**Add & LayerNorm.** A residual connection is employed between the input (shape:  $b \times s \times d_m$ ) and output (shape:  $b \times s \times d_m$ ) of the self-attention layer. Afterwards, a normalization layer is applied to the output of the Add operation (shape:  $b \times s \times d_m$ ). These same transformations are also performed on the FFN block, and we formulate them in Equation 2. Moreover, the size of parameters and their gradients of the LayerNorm layer is  $4d_m$ , one for weights and one for bias, which can be ignored compared to other parts.

$$y = \text{LayerNorm}(f(x) + x), \quad \text{where } f \in \{\text{Attention, FFN}\} \quad (2)$$

**Feed-Forward Networks.** As formulated in Equation 3, the feed-forward network (FFN) layer applies linear transformations to the inputs with two fully-connected (FC) layers separated by a GeLU activation function [22]. Specifically, the first FC layer ( $W_1 \in \mathbb{R}^{d_m \times d_{ffn}}$ ) projects the input into a new space with higher dimension, which allows the model to capture more complex relationships within a single token, while the second FC layer ( $W_2 \in \mathbb{R}^{d_{ffn} \times d_m}$ ) shrinks the dimension back to original, which helps to ensure that the output of the model is well-behaved.

$$\text{FFN}(x_s) = W_2 \cdot \text{GeLU}(W_1 \cdot x_s) \quad (3)$$

**Memory Usage Analysis.** According to Table 1, we can estimate the memory usage of any decoder-only Transformer models, where we do not take the embedding\_lookup and loss function into consideration. For the GPT-3 175B [8], the *Params*, *Acts* and *Optims* consumes 648GB, 162GB, and 1944GB, respectively, when batch size ( $b$ ) is 1, sequence length ( $s$ ) is 2048, hidden size of embeddings ( $d_m$ ) is 12288 and hidden size of FFN ( $d_{ffn}$ ) is 49152. To satisfy the memory requirement of large model training, multiple GPUs will be involved for distributed training with parallelism strategies.

## 2.3 Distributed Training

By partitioning model parameters as well as their computation among multiple GPUs, the training can be significantly accelerated, enabling researchers to train large-scale Transformer models in a shorter amount of time. In this section, we will provide a comprehensive overview on existing parallelism strategies that have been widely adopted for Transformer models.

**Data Parallelism and Zero Redundancy Optimization.** In data parallelism (DP), training samples are partitioned while model parameters and optimizer states are duplicated across multiple devices [31]. Each device executes the forward and backward propagation on its local mini-batch data to obtain its parameter gradients, and the gradients are averaged through a synchronization across all devices (e.g., by all-reduce). Eventually, each device updates the model parameters and optimizer states individually via the synchronized gradients. However, the vanilla data parallelism requires each device to maintain a full copy of model states, which is memory-inefficient for large-scale Transformer models. To reduce the memory consumption, Rajbhandari et al. [48] proposed the Zero Redundancy Optimization (ZeRO) technique, which evenly partitions the model states across all devices. To be specific, when training with  $N$  devices, each device only stores and updates  $1/N$  of the model states. However, in each training iteration, an extra round of all-gather communication is needed to ensure each device gets the full updated parameters in order to accomplish the propagation. In short, ZeRO-powered data parallelism improves the memory efficiency at the cost of extra communication overheads.

**Model Parallelism and Hybrid Parallelism.** Model parallelism splits the model across multiple devices and performs the forward and backward propagation in a distributed manner. Megatron-LM [38] proposed tensor parallelism (TP), which partitioned the queries, keys and values matrices of the attention network in a row- or column-parallel fashion, which exploits the inherent parallelism of the multi-head attention. In pipeline parallelism (PP), the model is partitioned into a sequence of stages and each stage is executed on a separate device. Huang et al. [23] proposed the batch-splitting pipeline algorithm and achieved almost linear speedup over multiple GPUs. Hybrid parallelism refers to the combination of two or more parallelism strategies to improve the training efficiency, which must consider the three aspects of computation, communication, and storage simultaneously. Zheng et al. [63] constructed a large amounts of model parallelism execution plans by exploiting both inter-operator and intra-operator parallelism in a hierarchical manner. They also designed many compilation passes to automatically derive efficiency plans at each parallelism level.

## 3 MOTIVATIONS AND SYSTEM DESIGN

### 3.1 Use Cases in Tencent

Through collecting a significant amount of training task data from the machine learning platform in Tencent, we have identified two main categories of common use cases, including pre-training and fine-tuning. The majority of these tasks involve the large-scale training of Transformer models, as the Transformer architecture has revolutionized natural language processing and enabled efficient handling of sequential data with long-range dependencies.

Each task category has its unique characteristics and primary objectives, and we provide a detailed analysis of these categories in the following section. Additionally, we briefly introduce corresponding system optimizations to improve their efficiency and effectiveness.

**Pre-Training.** Pre-training refers to the process of training a large-scale Transformer model on vast amounts of unlabelled data to learn rich and diverse features, which then can be useful for a wide range of downstream tasks, such as question answering and machine translation. Given the scale of models and datasets, pre-training tasks are extremely time-consuming and memory-hungry. Meanwhile, Kaplan et al. [28] suggested that the model quality of pre-trained models scales as a power-law with data size, model size, and the amount of computation, which further increases the demand of pre-training tasks.

After analyzing the log information of the platform, we find that although pre-training tasks have to use hundreds or even thousands of GPUs to train for several weeks, they account for only about 10% of the total number of tasks. This is because researchers in Tencent prefer to jointly train a large-scale Transformer model as their shared base model. And we observe that there exist two main characteristics associated with pre-training:

- **Low-Efficiency on Scalability.** During the running of pre-training tasks, users may pause and request more GPUs to obtain experimental results faster. However, in many cases the GPU utilization and training throughput decrease after more GPUs are involved, depending on the distributed training strategy of the model. Take a real case as the example. Training a 64-layer GPT model with the hybrid parallelism strategy of Megatron-LM on 72 GPUs is slower than that on 64 GPUs. As discussed in Section 1, many researchers in Tencent are not familiar with distributed computation and cannot adjust the parallelism strategy correspondingly.
- **Failure and Recovery.** When more GPUs are involved, the Mean Time To Failure (MTTF) is shortened accordingly. Given the large amount of GPUs and the long training time, pre-training tasks would encounter GPU failure with a high probability, and should be restarted after failure.

**Fine-Tuning.** Fine-tuning refers to taking the pre-trained model and adapting it to a specific downstream task with domain-specific data, such as the supervised fine-tuning (SFT) phase in Instruct-GPT [42]. After fine-tuning the pre-trained model for a specific downstream task, researchers aim to deploy the resulting model in a real-world product. Therefore, they need to carefully tune hyperparameters and iterate on experiments to achieve the best possible performance. This phase involves a trial-and-error progress and may require several iterations until a satisfactory model is obtained.

After analyzing the log information of the platform, we notice that the fine-tuning tasks account for about 90% of the total number of tasks. Each fine-tuning task also requires a large number of GPUs, but the running time is shorter than pre-training tasks (usually in hours). And we find that there exists three main characteristics associated with fine-tuning:

- **Low-Efficiency on GPU Utilization.** Smaller batch sizes are often used in fine-tuning tasks due to smaller downstream datasets and to avoid overfitting. This leads to a disproportionate amount

Table 2: Distribution of tensor sizes within one layer of GPT3.

Tensor Size (MB)	3072	2304	1152	768	576
Counts	4	6	4	20	12
Tensor Size (MB)	288	0.375	0.046875	0.0234375	
Counts	8	4	6	4	

of time spent on distributed communication, reducing the utilization of expensive GPU computing units.

- Long Response Time.** The task queue is frequently populated with a large number of fine-tuning tasks, each requiring lots of GPU resources. Due to limited cluster resources, most fine-tuning tasks face long wait times, sometimes amounting to several hours, although most of them typically require only a few hours.

### 3.2 System Design

**Lessons from Tencent.** To effectively address the two prevalent use cases within Tencent, we aim to design a system that can not only scale efficiently across hundreds of GPUs but also facilitate the fine-tuning of large models within the constraints of limited GPU resources. Unlike existing systems, such as DeepSpeed, we strive to avoid modifications to user code, making it easier for wider adoption. Additionally, we aim to eliminate system complexity, focusing on user needs to provide the most lightweight implementation, rather than integrating nearly all optimization strategies as DeepSpeed does. Therefore, we’ve incorporated the following three strategic elements into the foundational design of our system:

**Data Parallelism.** During the pre-training process, the number of GPUs used for training may vary for each task, which requires our distributed strategy to be easy-to-scale and have good scalability. While model parallelism might offer superior throughput performance in certain instances, its design necessitates specialized knowledge and presents migration challenges across varying degrees of parallelism. As such, we’ve opted for data parallelism as our foundational distributed solution.

**Parameter Sharding.** With data parallelism(DP), each GPU needs to hold the complete model states, posing challenges for training larger models. To address this, we have adopted the parameter sharding technique as proposed by ZeRO [48]. This approach splits each parameter evenly across multiple GPUs and when a parameter needs to be calculated, a complete parameter is fetched through an all-gather operation. This approach dramatically reduces each GPU’s memory requirements, enabling the training of substantially larger models.

**Hierarchical Memory.** The issues of long response time and low resource utilization are significantly due to the large number of fine-tuning tasks as well as their excessive numbers of GPUs w.r.t. the relatively small batch sizes. To tackle these problems, we incorporate hierarchical storage within GPU servers to meet the memory requirements for fine-tuning tasks, thereby reducing the number of GPUs. Additionally, leveraging hierarchical storage also allows pre-training tasks to train larger models with the same numbers of GPU servers.

Upon establishing the foundational principles of our system, we recognized that simply integrating these strategies does not optimally exploit resource utilization as evidenced in existing systems [47]. As an initial step, we examined the distribution of tensor

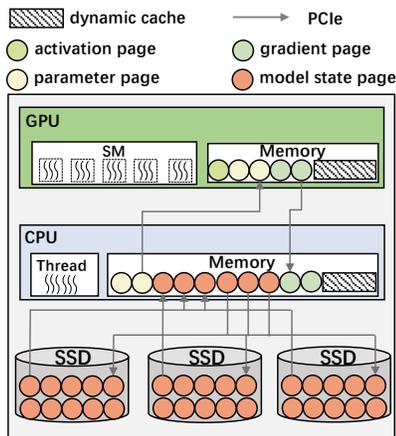


Figure 2: System architecture of ANGEL-PTM.

sizes within one Transformer layer of the GPT3-175B model during training, employing the formulation detailed in Table 1. The findings, outlined in Table 2, illustrate a significant variation in tensor sizes, ranging from 3072MB to a meager 0.02MB. We identified three main inefficiencies stemming from this disparity, including *inefficient memory utilization*, *inefficient bandwidth utilization* and *inefficient GPU utilization*.

To cope with these inefficiencies, we first propose the *Page-based memory organization* in Section 4.1 to improve the memory usage and then design the *Unified Scheduler* in Section 4.2 to fully utilize the bandwidth. Moreover, we also design a *Lock-Free Updating Mechanism* in Section 4.3 to address the insufficiency in GPU usage when enlarging model to an extreme scale.

## 4 ANGEL-PTM

Our system, namely ANGEL-PTM, is designed for researchers and developers to design and experiment large-scale Transformer models in Tencent. Figure 2 illustrates the system architecture, where we enable the fine-grained memory management at the *Page* level and adopt the dynamic cache techniques for sufficient memory usage. In the following subsections, we will analyze the limitations of existing systems and present our solutions to address these issues.

### 4.1 Page-Based Memory Organization

**Inefficient Memory Utilization.** Existing systems suffer from memory fragments due to their coarse memory management. For instance, DeepSpeed uses the original memory management of PyTorch for offloading and recomputing, which frequently allocates and releases tensors, leading to space fragments because the sizes of these tensors are not uniform as discussed in Section 3. Patrick-Star [14] manages GPU memory in chunks rather than tensors, where the chunk size must be larger than the largest tensor used in model training. This would also result in memory fragments within each chunk as well as the in-efficiency of the overlapping between communication and computation.

To reduce the fragments of memory space organization and improve the efficiency of memory movements, we introduce the *Page* abstraction, which works as the minimum unit of memory

```

1 struct Page {
2     /* Page Information */
3     void* data_ptr;
4     size_t total_bytes;
5     size_t available_bytes;
6     // device_map: {0: GPU, 1: CPU, 2: SSD}
7     size_t device_index;
8     // ids for tensors in this page
9     size_t tensor_id[2];
10    // occupied bytes for each tensor
11    size_t tensor_bytes[2];
12
13    /* Page Interface */
14    // allocate required bytes for id-th Tensor
15    void allocate(size_t required_bytes, size_t id);
16    // release space of id-th Tensor
17    void release(size_t id);
18    // move this page to target device asynchronously
19    void move(size_t target_device_index);
20    // send this page to id-th server asynchronously
21    void send(size_t id);
22    // receive contents from id-th server
23    void receive(size_t id);
24 };

```

Figure 3: The Page Abstraction.

operations for heterogeneous storage, including allocation, release, movement, and remote communication. In our proposed model, each tensor is comprised of multiple Pages. Given the fine-grained level of memory operations, a single tensor may span over several discontinuous Pages. To mitigate such discontinuity, we initially leverage the iterative training characteristic to schedule the placement of each Page. If necessary, we further employ an additional merge operation to arrange these Pages into a contiguous buffer.

**Page Abstraction.** As illustrated in Figure 3, the *Page* abstraction includes several key pieces of information, including a pointer to the actual memory data, the total number of bytes in the page, the number of bytes that are available for the next allocation, and the index of the device where the page is currently located. Additionally, each *Page* can be associated with one or more tensors, with unique identifiers and information about the amount of memory occupied by each tensor. In order to simplify the memory management, we decide to limit each page to contain information about a maximum of two tensors at any given time. The reasoning behind this decision will be analyzed in conjunction with the discussion of the optimal page size in detail below.

Moreover, *Page* also provides several interfaces for accessing and manipulating the data stored within it. These interfaces enable developers to perform a wide range of operations on *Page* objects, including allocating and releasing memory for specific tensors, moving pages between the heterogeneous memory, and sending/receiving pages across different servers.

**Optimal Page Size.** The selection of an optimal *Page* size represents a crucial balancing act between memory management efficiency and overall throughput. A larger *Page* size can result in numerous tensors cohabiting within a single page, subsequently raising management complexity and leading to wasted space. Conversely, a smaller *Page* size could lead to increased overhead due to under-utilized bandwidth during data movement. In our approach, we identify the smallest *Page* size that fully leverages the PCIe

```

1 struct Tensor {
2     /* Tensor Information */
3     size_t id;
4     vector<Page> page_list;
5     size_t dtype;
6     size_t* shape;
7     size_t device_index;
8
9     /* Tensor Interface */
10    void allocate(size_t* shape, size_t dtype);
11    void release();
12    void move(size_t target_device_index);
13    void merge();
14 };

```

Figure 4: The Tensor structure in our system.

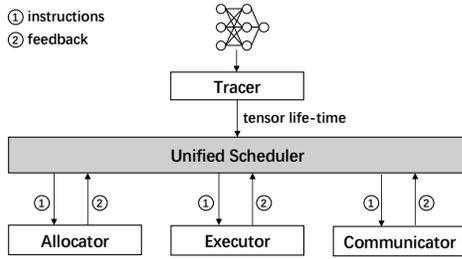
bandwidth as optimal for our system — this corresponds to a size of 4MB. Our observations during training show that the vast majority of model states exceed this 4MB size, a finding that’s also supported by the data in Table 1. By arranging these tensors strategically, we can limit the association to a maximum of two tensors per page, substantially reducing management complexity. The rationale behind the maximum limit of two tensors per page is rooted in our approach, which allows a tensor to begin on a partially filled page, span several complete pages to form its body, and may end on producing another partially filled page. For tensors smaller than 4MB, we permit each to occupy a separate page for simplicity, given they represent only a small fraction of the total memory usage.

**Tensor Management.** The *Tensor* structure is a fundamental data structure in our system that represents multi-dimensional arrays of numerical data, composed of at least one page. As shown in Figure 4, it contains several crucial pieces of information, such as a unique id associated with the tensor, data type, shape, and the current device index. Meanwhile, we set the device index as -1 when the tensor is not ready for computation (i.e., some of its pages need to be fetched from hierarchical memory or other servers). The *Tensor* structure also provides a set of interfaces for memory management, such as allocating a certain shape tensor or releasing its data. Additionally, it offers an explicit interface for moving data between heterogeneous memory. Since the space of different pages may not be contiguous, the *Tensor* structure provides the merge interface to make them contiguous.

In short, the *Tensor* structure, which is user-facing, incorporates a variety of common operations, thus providing a simple and intuitive interface for users. On the other hand, the *Page* structure, invoked by the system runtime, encapsulates all the fine-grained memory operations. This separation of duties ensures efficient training, with the *Page* structure acting as a workhorse that manages intricate memory tasks behind the scenes, allowing users to interact with the more accessible *Tensor* structure.

## 4.2 Unified Scheduler

**Inefficient Bandwidth Utilization.** In the context of this hierarchical storage-based distributed training, model states must be transferred among various levels of storage via PCIe, while parameter gradients need to be synchronized via network. In terms of hardware design, GPU computations are inherently capable of overlapping with both PCIe and network. Therefore, deciding what needs



**Figure 5: Illustration of the Unified Scheduler and associated components in ANGEL-PTM.**

to be communicated and when is crucial to the overall throughput efficiency of the system. However, existing systems, such as DeepSpeed, distribute model states statically across multi-level storage, for instance, relegating all optimizer states to CPU memory, while adopting static scheduling strategies such as prefetching  $k$ -layers in advance. These systems make decisions without considering the actual demands of workloads, markedly impeding the system’s throughput and prompting unnecessary data movements.

To efficiently utilize the complex hierarchical network resources within GPU servers, we devise a unified scheduling method built upon the *Page* abstraction for efficient training. As illustrated in Figure 5, once the user defines the model, the Tracer obtains the tensor access pattern and lifetime for each tensor (as detailed in Section 5). Subsequently, the Unified Scheduler, armed with these statistics, produces a static schedule for each operation during training. The fundamental principle here is that deep learning model training is iterative in nature, allowing us to schedule the operators based on the tensor access pattern. Specifically, by analyzing the lifetime and size of each tensor, the Unified Scheduler determines the most appropriate time to execute each operation in each iteration. This includes calling the Allocator to move tensors, instructing the Executor to perform GPU computations, and prompting the Communicator for network communication.

**Tensor Life-Time.** The life-time of a tensor refers to the duration from its first access time to its last access time within a training iteration. A fundamental characteristic of deep learning training is its iterative nature, meaning that computation operations are sequentially and iteratively inserted into the computing stream. This allows us to optimize the sequence of data movements and communications relative to computation, ensuring that the required data is available at the right time for each computation. This reduces idle time and improves memory efficiency. Moreover, as computation is performed at the tensor level, tensor allocation and release can also be conducted at this level. The use of logical IDs instead of real-time for lifetime tracking simplifies the scheduling process. By utilizing tensor lifetime information, we can optimize the scheduling of computations, movements, and communications in a comprehensive manner.

**Unified Scheduler.** The Unified Scheduler is the crucial component in ANGEL-PTM, responsible for orchestrating the activities of other components. To be more precise, it assimilates life-time information and access patterns as provided by the Tracer, subsequently generating a task queue to establish the operation schedule.

---

### Algorithm 1: Fine-grained Life-time based Scheduling

---

**Input:**  $model = \{l_0, \dots, l_{n-1}\}$ : list of layers, where  $l_i$  is  $i$ -th layer  
 $traces$ : List of trace for each tensor  
 $gpuMemory$ : total GPU memory

**Ouput:**  $S$ : Queue of tasks, each is  $\{operation, page, trigger\_id\}$

```

1 /* Phase 1: Prioritize move_to_gpu tasks */
2  $wait\_stack = \{\}$ ;
3 for  $l_i \in model$  do
4   for  $page \in l_i.param.page\_list$  do
5      $S.append(\{move\_to\_gpu, page, 0\})$ ;
6 for  $l_i \in model$  do
7   while  $get\_available\_memory(S, traces) < size(l_i)$  do
8      $task \leftarrow pop$  the last movement task from  $S$ ;
9      $wait\_stack.push\_back(task.page)$ ;
10  for  $page \in l_i.param.page\_list$  do
11     $S.append(\{all\_gather, page, i\})$ ;
12     $S.append(\{compute, l_i, i\})$ ;
13    while  $wait\_stack.non\_empty()$  and
14       $get\_available\_memory(S, traces) > page\_size$  do
15       $page \leftarrow wait\_stack.pop\_back()$ ;
16       $S.append(\{move\_to\_gpu, page, i\})$ ;
17 /* Phase 2: Advance all_gather tasks to overlap them with
18    previous computation if no out-of-memory (OOM) */
19 for  $task \in \{t | t \in S, t.operation == all\_gather\}$  do
20    $task \leftarrow pop$   $task$  from  $S$ ;
21    $id \leftarrow get$  the earliest possible id to trigger  $task$  without OOM;
22    $S \leftarrow insert \{all\_gather, task.page, id\}$  into  $S$ ;
23 return  $S$ ;

```

---

The use of a queue, rather than a priority queue, is deemed sufficient as we have incorporated priority considerations within Algorithm 1 and have also implemented a roll-back mechanism.

First, we must determine an effective way to distribute the model states to various memory devices, such as GPU memory, CPU memory and SSD storage, and distribute computation tasks to various computing devices, such as CPUs and GPUs, which is a complex NP planning problem and cannot be solved directly for large-scale models. To tackle this problem, we develop a heuristic method and incorporate empirical information to aid in the design of our strategies. To be specific, the forward and backward computations of the transformer models are mainly composed of FP16 matrix multiplication, which is rather compute-intensive and requires less memory. The optimizer update computations, to the contrary, are composed of FP32 matrix addition, which is memory-intensive and takes less time to compute. Given the fact that GPUs have high computing capabilities but limited memory, we prioritize putting the forward and backward computations into the GPU, and the optimizer update computations into the CPU. The SSD storage will also be involved to store the FP32 optimizer states for scaling model size. With this approach, it is necessary to transfer the FP16 parameters from the CPU to the GPU, and subsequently transfer the FP16 gradients from the CPU back to the GPU. These data movements need to be carefully scheduled to overlap with computations and avoid introducing extra overhead. Meanwhile, we utilize the recomputation technique [9] to further alleviate the GPU memory

pressure, where some activations are released in the forward pass and then are regenerated in the backward pass by re-executing their forward computation.

Second, we leverage caching techniques to fully exploit the high-speed memory and computation capabilities of the GPUs after previous assignments. For instance, if sufficient space is available, we reserve a portion of the GPU memory as the cache to store a segment of the CPU’s optimizer states. Additionally, we move the relevant computations to GPUs, which reduces memory transfers and accelerates computation, leading to an overall improvement in training throughput. Moreover, we dynamically make cache size decisions for each model based on its tensor lifetime information, ensuring training without encountering GPU out-of-memory errors. Such a caching technique allows us to maximize GPU memory utilization and minimize data transfer overheads.

The Algorithm 1 consists of two phases. In the first phase, we prioritize the data movements by inserting `move_to_gpu` tasks for each page of each tensor at the beginning of the schedule (lines 3-5), which is based on our prior knowledge that the speed of CPU-GPU data transfer (32GB/s) is slower than that of GPU-GPU communication (200GB/s). We then iterate over each layer to insert `all_gather` tasks and compute tasks on demand (lines 10-12). If the current available memory is not sufficient for training, we will pop the last movement task until the memory requirement is satisfied (lines 7-9). If the current available memory is sufficient for the movement of the next page, we will schedule it immediately (lines 13-15). In second phase, we advance `all_gather` tasks as early as possible to overlap them with previous computation tasks. Specifically, for each `all_gather` task, we gradually try to shift its scheduled time earlier (i.e., decrease its trigger id by 1). In each trial, the trace information, tensor sizes, and the schedules after shifting will be utilized to measure the peak memory usage in order to determine whether there will be an out-of-memory (OOM) error after the shifting. And this progress ends at the smallest trigger id that does not encounter OOM error. Eventually, the `all_gather` task will be inserted back to the schedule, along with the new trigger id (Line 20-21). This approach contributes to more efficient use of available memory and better overlapping of tasks on the GPU, ultimately leading to improved performance.

### 4.3 Lock-Free Updating Mechanism

**Inefficient GPU Utilization.** When incorporating heterogeneous memory or multiple GPUs into the training process, GPUs computations often face delays due to data transfer. Primarily, GPUs are left waiting for the updated parameters to be transferred for computations, or computed gradients to be cleared to free up space. Additionally, GPUs must wait for gradient synchronization through remote communication to update local parameters. This, inevitably, results in the wastage of GPU resources and a decrease in overall efficiency. Even more problematic is the integration of SSD storage into training within existing systems, which further reduces GPU utilization due to the slow I/O speed of SSDs. To illustrate, consider the A100 Server from Tencent, which demonstrates I/O speeds for GPU memory access, CPU-GPU transfer, and SSD-CPU transfer at 600 GB/s, 32 GB/s, and 3.5 GB/s, respectively. Our observations have shown that after integrating CPU memory and SSD storage,

---

#### Algorithm 2: Lock-Free Updating Mechanism

---

**Input:**  $model = \{l_0, \dots, l_{n-1}\}$ : list of layers, where  $l_i$  is  $i$ -th layer  
 $p_{32}(l_i)$ : parameters for  $l_i$  in FP32  
 $m_{32}(l_i)$ : first moments of gradients for  $l_i$  in FP32  
 $v_{32}(l_i)$ : second moments of gradients for  $l_i$  in FP32  
 $g_{16}(l_i)$ : gradients for  $l_i$  in FP16  
 $p'_{16}(l_i)$ : buffered parameters for  $l_i$  in FP16  
 $g'_{16}(l_i)$ : buffered gradients for  $l_i$  in FP16

```

1 /* Updating thread on CPU */
2 while there are uncleared buffered gradients do
3   for  $l_i \in reverse(model)$  do
4     Fetch  $p_{32}(l_i), m_{32}(l_i), v_{32}(l_i)$  from SSD storage;
5     Update  $p_{32}(l_i), m_{32}(l_i), v_{32}(l_i)$  via  $g'_{16}(l_i)$ ;
6     Pass  $p_{32}(l_i)$  to the buffering thread;
7     Offload  $p_{32}(l_i), m_{32}(l_i), v_{32}(l_i)$  to SSD storage;
8
9 /* Buffering thread on CPU */
10 while not finished do
11   if received  $p_{32}(l_i)$  from the updating thread then
12     Clear buffered gradients:  $g'_{16}(l_i) \leftarrow 0$ ;
13     Update buffered parameters:  $p'_{16}(l_i) \leftarrow cast(p_{32}(l_i), FP16)$ ;
14   else if received  $g_{16}(l_i)$  from the GPU then
15     Accumulate buffered gradients:  $g'_{16}(l_i) \leftarrow g'_{16}(l_i) + g_{16}(l_i)$ ;
16
17 /* Computation on GPU */
18 while not finished do
19   for  $l_i \in model$  do
20     Fetch buffered parameters  $p'_{16}(l_i)$  from CPU memory;
21     Perform forward computation with  $p'_{16}(l_i)$ ;
22   for  $l_i \in reverse(model)$  do
23     Perform backward computation with  $p'_{16}(l_i)$  and get  $g_{16}(l_i)$ ;
24     Offload  $g_{16}(l_i)$  to CPU memory (inform the buffering thread);

```

---

nearly 80% of iteration time is spent idle, whereas this figure is a mere 10% when only CPU memory is introduced.

To mitigate this training bottleneck, we design a novel *Lock-Free Updating Mechanism*, which decouples GPU computation from CPU optimizer operations through a novel asynchronous consistency control protocol. Algorithm 2 illustrates the details of our proposed mechanism. The essential idea is to employ two buffers in CPU memory to store the FP16 parameters and gradients respectively, and leverage an auxiliary buffering thread to maintain the buffers. During training, each GPU fetches the FP16 parameters of each layer from the CPU buffer and perform forward and backward computations (Line 19-23). Then, the generated gradients are offloaded to CPU memory (Line 24), and eventually accumulated into the CPU buffer by the buffering thread (Line 15). During optimizer updating, the updating thread on CPU reads the FP32 parameters and optimizer states from SSD, which are then updated according to the buffered, accumulated gradients (Line 4-5). Subsequently, the buffering thread receives a signal to individually update the buffers of each layer. This process includes purging the buffered gradients and converting the updated parameters into the buffered parameters (Line 12-13). These steps occur concurrently with offloading the updated parameters and optimizer states to the SSD (Line 7). Given that the optimizer’s updates to parameters and the GPU’s retrieval of parameters operate on a layer-level granularity

for consistency control, rather than on a full model scale, we’ve designed our update algorithm to be lock-free.

Although our updating mechanism incurs extra memory overhead, it is acceptable since both the buffered parameters and gradients are stored in FP16, requiring small memory footprints. Another side effect of the lock-free mechanism is that it may introduce staleness into the parameters given the fact that the updating thread on CPU runs slower than the GPU due to the limited SSD I/O bandwidth. Nevertheless, existing studies have verified that deep learning model training can well tolerate such staleness. In Section 6, we will empirically show that the convergence is not harmed while enjoying the efficiency improvement brought by the lock-free mechanism. Last but not least, since the optimizer updates are element-wise, the data movement and CPU computations can be scheduled at the *Page* level, which further improves the efficiency.

## 5 IMPLEMENTATION

ANGEL-PTM offers a comprehensive solution for efficient deep learning model training in industrial settings. It leverages some key techniques [33, 36, 41] from Hetu [34], gets implemented over PyTorch [43], and features the *Page* abstraction for memory efficiency and a unified scheduling method for resource utilization. Furthermore, ANGEL-PTM has undergone extensive optimization on A100 servers, enabling it to take full advantage of hardware capabilities for deep learning tasks. We would like to briefly go through the implementation of the key components in this section.

**Tracer.** The Tracer in ANGEL-PTM is responsible for tracking the usage of each tensor and summarizing a tensor access pattern for the given model as a list of following elements:

- *tensor\_id*: The logical ID of this tensor.
- *first\_id*: The logical ID when first accessing this tensor.
- *end\_id*: The logical ID when last accessing this tensor.
- *cpu\_time*: The time for producing this tensor on CPU.
- *gpu\_time*: The time for producing this tensor on GPU.

To assign a unique tensor ID to each parameter, we modify the `__init__` method of the Parameter class to use a global variable *id*. Then, we track the first and last use of each parameter during an iteration by registering hook functions, recording them as *first\_id* and *end\_id*, respectively. To capture the generation time of tensors on both CPUs and GPUs, we use the `time.time()` and `CudaEvent` interfaces respectively to accurately measure the CPU and GPU time for each tensor.

**Unified Scheduler.** The Unified Scheduler is responsible for coordinating the activities of three components, including Allocator, Executor, and Communicator. Sending instructions by the message passing will bring severe overheads into training, thus we adopt the event-driven programming techniques to implement our Unified Scheduler. For example, computations will be launched into threads only if the events of modifying its input tensor are completed.

**Allocator.** The Allocator in ANGEL-PTM is responsible for managing tensors at the *Page* level in the hierarchical memory resources. It provides three memory operations for each page, including `allocate`, `release` and `move`. To reduce the overhead of requesting memory space and take advantages of the iterative nature of training, we pre-allocate space from the hierarchical memory of the system, including GPU memory, CPU pinned memory, and

**Table 3: Overview of hardware environments.**

The Configuration of one GPU server	
CPU	4 × AMD EPYC 7K62 48-Core Processor
CPU Memory	32 × 32GB 2933MHz DDR4
GPU	8 × NVIDIA Tesla A100 Tensor Core GPU
GPU Memory	8 × 40GB HBM2
SSD	11TB

SSD memory. To enable fine-grained memory operations, we divide the pre-allocated memory into pages of fixed size, where each page can be allocated, released and moved independently. Moreover, we utilize `cudaMemcpyAsync()` and the DeepNVMe library [49] for asynchronous GPU-CPU and CPU-SSD data movements, respectively.

**Executor.** The Executor in ANGEL-PTM is responsible for scheduling the computation of Tensors on computational devices such as CPUs and GPUs on the server. Meanwhile, it maintains a separate stream for each of these computational devices, including a CPU stream and a GPU stream. By receiving instructions from the unified scheduler, it inserts computations into the corresponding stream and schedules them to the computation threads in the order of insertion. When all the inputs for the computation are ready, the computation begins, and feedback is sent back to the unified scheduler after the computation is complete.

**Communicator.** The Communicator in ANGEL-PTM is responsible for scheduling communication between different network devices, including NIC and NVLink. We implement the Communicator by using the NCCL library [24], which provides efficient communication primitives for multi-GPU systems. These primitives include collective operations such as `AllReduce`, `AllGather`, and `ReduceScatter`, which are essential for exchanging model parameters and gradients between GPUs during training. The Communicator also maintains a queue to store communication tasks and schedules them for execution based on instructions from the Unified Scheduler, thus it enables reordering the tasks in the queue to improve the overlap between computation and communication.

**Efficient Movement on Distributed Servers.** GPU servers typically have a complex interconnect topology, such as A100 servers [3] that contain two CPUs, four PCIe switches, and eight A100 GPUs. These GPUs can communicate with the CPU memory in parallel, providing efficient data movement in distributed training. To take full advantage of this hardware feature, we evenly partition the model parameters across GPUs to parallelize the movement of parameters between the CPU and GPUs, which is similar to ZeRO-Infinity [49]. This further accelerates data movement during training and achieves excellent scalability.

## 6 EXPERIMENTAL EVALUATION

### 6.1 Experiment Setup

**Machine environment.** Ours experiments are conducted on a production-grade GPU cluster in Tencent, composed of 96 servers. Each server is equipped with 8 NVIDIA Tesla A100-40GB GPUs, with additional details provided in Table 3. Within each server, GPUs are connected via NVLink-3.0, while servers are networked using RoCE and supported by 16 NICs, yielding a total bandwidth

**Table 4: Models for Evaluation.**

Model	#Layer	#Head	$d_{\text{Model}}$	$d_{\text{FFN}}$	#Expert
GPT3-1.7B	24	24	2304	9216	-
GPT3-13B	40	40	5140	20506	-
GPT3-28B	26	128	8192	32768	-
GPT3-30B	64	36	8192	32768	-
GPT3-55B	68	128	8192	32768	-
GPT3-120B	64	96	12288	49152	-
GPT3-175B	70	112	14336	57344	-
T5-1.4B	16	16	1024	16384	-
T5-27B	28	64	4096	16384	-
T5-58B	60	64	4096	16384	-
T5-MoE-1.2T	16	16	1024	16384	2304

of 200GB/s ( $16 \times 12.5$  GB/s). The system’s PCIe bandwidth stands at 32GB/s with the SSD peak bandwidth reaching 3.5GB/s.

**Benchmarks.** We conduct evaluations on three large-scale Transformer models, namely GPT-3 [8], T5 [46], and T5-MoE [15], to validate the effectiveness and scalability of our proposed system. To achieve different model sizes, we experiment with varying numbers of Transformer blocks, hidden dimensions, and experts, and the specific model configurations are presented in Table 4. We train all of these models using the mixed precision technique as introduced in Section 2, which stores the model states in Float32 while computes in BFloat16.

**Baselines.** Prior to ANGEL-PTM, DeepSpeed [50] and Megatron-LM [38], due to their widespread adoption, are the two pre-training systems incorporated in the Taiji Machine Learning Platform of Tencent. Therefore, we choose these two systems as our baselines to evaluate the effectiveness of ANGEL-PTM. DeepSpeed is a heterogeneous training solution that currently achieves state-of-the-art performance, and we use the official examples to ensure a fair comparison. For Megatron-LM, we search the best parallelism strategy for each experimented model, which results in a hybrid parallelism solution that combines data parallelism, model parallelism, and pipeline parallelism.

We conduct a series of experiments to present the effectiveness and scalability of ANGEL-PTM in training large-scale Transformer models. Specifically, we evaluate the maximum model scale supported by our system on a single server in Section 6.2 and compare the throughput of different models on two GPU settings in Section 6.3. Additionally, we demonstrate the scalability of our system on hundreds of GPUs for training both billion-scale dense models and trillion-scale sparse models in Section 6.4. Furthermore, in Section 6.5, we evaluate the convergence performance of our proposed *Lock-Free Updating Mechanism* to further validate the effectiveness of our system when introducing SSD storage to scale the model size to an extreme level. It is worth noting that, except for experiments in Section 6.5, which utilizes SSD storage for training, all other sections utilize the memory of CPUs and GPUs by default.

## 6.2 Model Scale

We first conduct evaluations on a single server to test the maximum model size and corresponding maximum throughput that can be supported by ANGEL-PTM and DeepSpeed, where we increase the

**Table 5: Max Supported Model Scale on a Single Server.**

Model	System	#Params	#Batch	GPU Mem	Samples/s
GPT	DeepSpeed	28B	1	18	0.404
		28B	36	40	7.61
	AngelPTM	28B	38	39	<b>10.99</b>
		<b>55B</b>	1	33	0.464
T5	DeepSpeed	27B	1	20	0.317
		27B	32	39	7.31
	AngelPTM	27B	50	40	<b>14.38</b>
		<b>58B</b>	1	38	0.432
		<b>58B</b>	4	40	3.37
		<b>58B</b>	4	40	3.37

number of transformer blocks and fix other model settings. Both systems partition the model evenly across multiple GPUs using ZeRO-3 [48], which enables efficient linear scaling for model size.

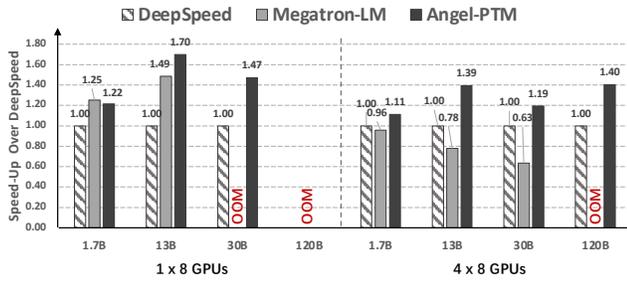
Results are summarized in Table 5. For GPT models, we set the number of heads as 128, the embedding dimension as 8192, and the FFN hidden size as 32768. DeepSpeed can support a maximum model scale of 28B with 26 layers, while ANGEL-PTM can further scale it up to 55B with 68 layers, which is a 96.4% improvement. Note that despite each GPU having 22GB of memory available, DeepSpeed fails to scale to a larger model size. The reason is that since DeepSpeed statically partitions the model states across GPUs and CPUs, the maximum model scale will be limited by the CPU memory. In contrast, to fully exploit this available memory, ANGEL-PTM uses the dynamic memory management that moves partial model states into GPU memory to achieve larger model scale. Regarding training efficiency, specifically the samples/s column in Table 5, the maximum throughput of DeepSpeed is 7.61 samples/s, while that of ANGEL-PTM is 10.99 samples/s, which is a 44% improvement. Furthermore, the training efficiency of ANGEL-PTM for the largest supported GPT model, 55B, is 3.34 samples/s. These analyses can also be adapted to T5 models, where ANGEL-PTM achieves 114.8% improvement in terms of max model scale as well as 96.7% improvement in terms of throughput performance.

In summary, compared with DeepSpeed, ANGEL-PTM can (1) support larger scale of models using the same hardware resources and (2) achieve higher training efficiency for the same model.

## 6.3 Throughput

To verify the training efficiency of ANGEL-PTM, we assess the throughput of each competitor. Specifically, we trained a series of GPT models with the maximum batch size on 8 GPUs and 32 GPUs respectively, and the sizes of GPT models range from 1.7B to 120B. To provide a clear comparison between the systems, we normalize the throughput of each system w.r.t. DeepSpeed. The results are presented in Figure 6.

In the configuration of 1x8 GPUs, ANGEL-PTM consistently outperforms the other systems in terms of training efficiency, except for the 1.7B models. This is because the 1.7B model is small enough to be accommodated by a single GPU, and therefore the vanilla data parallelism (without ZeRO) achieves the best performance, which is also the strategy adopted by Megatron-LM. Since ANGEL-PTM involves extra overhead on memory management even when memory movement is not needed, it runs slightly slower than Megatron-LM



**Figure 6: Compare ANGEL-PTM with other famous pre-trained systems on a series of GPT models.**

(a 2.4% slowdown). On the 1.7B and 13B models, both ANGEL-PTM and Megatron-LM outperform DeepSpeed. This is reasonable as DeepSpeed statically partitions models states between the CPU memory and the GPU memory, leading to redundant memory movements. However, as the model size increased to 30B, Megatron-LM fails with the out-of-memory error due to the limited GPU memory and ANGEL-PTM still outperforms DeepSpeed because of our fine-grained life-time based scheduling method, which partitioned the model states to CPU on demand and scheduled the movements at the right time.

In the configuration of 4x8 GPUs, the performance of ANGEL-PTM is still the best. With more GPUs, Megatron-LM is able to support the 30B model, while DeepSpeed and ANGEL-PTM are further able to support the 120B model thanks to the ZeRO technique. Both ANGEL-PTM and DeepSpeed outperform Megatron-LM because they can train with larger micro batch sizes.

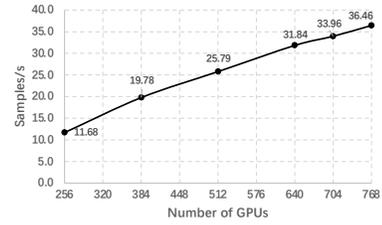
In summary, the experimental results demonstrate that ANGEL-PTM consistently outperforms both DeepSpeed and Megatron-LM in terms of throughput performance. Specifically, ANGEL-PTM achieves an average of 35.4% and up to 70% improvement over the state-of-the-art hierarchical training system, DeepSpeed. It also outperforms the state-of-the-art hybrid parallelism training system, Megatron-LM, by an average of 38.9% and up to 88.9%.

## 6.4 Scalability

To verify the scalability of ANGEL-PTM, we conduct evaluations on two extremely large-scale Transformer models, which are GPT3-175B and T5-MoE-1.2T.

**Evaluations on GPT3-175B.** The GPT3-175B model is first proposed by OpenAI [8], which is trained on an enormous amount of text data and performs a range of NLP tasks, such as language translation and question answering. It also acts as the foundation model for ChatGPT [1], which is a groundbreaking achievement in the field of artificial intelligence. Therefore, it is crucial to verify our system’s ability and scalability to support this model.

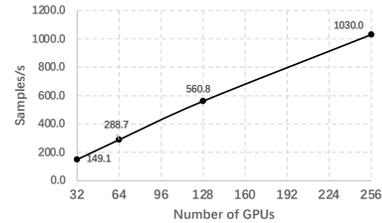
The configuration of GPT-175B is detailed in Table 4 and we illustrate the throughput of training this model on different number of GPUs in Figure 7. Our results demonstrate that ANGEL-PTM achieves super-linear scalability when training the GPT-175B model. We observe a throughput of 11.68 samples/s on 32 nodes (256 GPUs), which increases to 36.46 samples/s on 96 nodes (768 GPUs), resulting in a 3.12x speed-up. As the number of GPUs increases, the model states are distributed across more GPUs. This allows us to increase



**Figure 7: Scalability on training GPT3-175B models.**

the global batch size, which in turn fully utilizes the available GPU memory. Moreover, the optimizer updating process is parallelized across more CPUs, and data movements are parallelized across more PCIe. These factors attribute to the super-linear scalability and higher training throughput.

The results indicate that our system can take full advantage of the increasing number of GPUs and achieve efficient parallelization of data movements on PCIe, which is critical for training large models in a timely and cost-effective manner.



**Figure 8: Scalability on training T5-MoE models.**

**Evaluations on T5-MoE-1.2T.** The T5-MoE model is another well-known line of large models that employs a sparse MoE architecture and was first proposed by Google as Switch-Transformer [15, 30]. MoE models can significantly reduce training cost by routing input data to only a small number of expert networks. ANGEL-PTM trained T5-MoE models using expert parallelism [40], where expert parameters within an MoE layer are sharded among all GPUs while non-MoE parameters are duplicated.

The T5-MoE-1.2T model has 2304 experts per MoE layer, and the number of experts per GPU per MoE layer is fixed at 9 to achieve different model sizes when varying the number of GPUs. For example, the T5-MoE model trained on 128 GPUs has 1152 experts per MoE layer. The detailed configuration is presented in Table 4. We summarize the results in Figure 8, which indicates that ANGEL-PTM has near-linear scalability when training the T5-MoE model. With more servers involved in training, more input data will be feed into the all-to-all communication of the MoE layer, which can result in throughput degradation. Thus, the scalability on T5-MoE-1.2T is lower than that on GPT3-175B.

## 6.5 Advancing Support for Extreme Model Scale

As noted by Kaplan et al. [28], larger models tend to outperform smaller ones, which has motivated many researchers to increase the size of their models continuously. In this study, we evaluate the ability of ANGEL-PTM in supporting extreme-scale models, using a scaled-up version of the T5-MoE model with 10T parameters.

**Table 6: Training on Large-scale T5-MoE model with SSD .**

System	#Params	#GPUs	Samples/s	Valid Loss↓
AngelPTM	1T	64	37.26	1.124
	10T	576	317.82	<b>0.853</b>
+ Lock-Free	10T	576	<b>942.31</b>	0.861

We conduct pre-training of T5-MoE-1T and T5-MoE-10T based on an industrial text dataset. To support extreme model scales, we utilize the SSD storage into training. The experimental results are presented in Table 6, including training efficiency in samples/s and model quality in terms of validation loss. ANGEL-PTM achieves a throughput of 37.26 samples/s using 64 GPUs on the T5-MoE-1T model. When we scale up the model to 10T by increasing the number of experts, ANGEL-PTM achieves a throughput of 317.72 samples/s with 64 GPUs, demonstrating the near-linear scalability.

Moreover, when introducing SSD training, the I/O bandwidth of SSD significantly slows down the overall training speed. By leveraging the *Lock-Free Updating Mechanism* in Section 6.5 to perform asynchronous updates between CPUs and GPUs, ANGEL-PTM significantly improves the overall training throughput. To be specific, for the T5-MoE-10T model, the training throughput increased from 317.82 samples/s to 962.31 samples/s when the lock-free mechanism is enabled, achieving a speedup of 2.96×. Meanwhile, experimental results on the validation loss verify that this mechanism has little impact to the model quality.

## 7 RELATED WORK

**Distributed Training System.** Many well-known systems has been designed and implemented for large-scale model training. DeepSpeed proposed the ZeRO optimization on data parallelism, which evenly partitions model states across all devices with different optimization levels, including optimizer states (e.g., 32-bit parameter, the first and second moments of Adam [29]) in stage 1, gradients in stage 2 and 16-bit parameters in stage 3 [47]. However, this approach introduces extra communication to obtain the full parameters, leading to more time overheads for higher optimization levels, although they result in more memory reductions. Megatron-LM proposed a novel approach called tensor-parallelism, which is easy to implement and efficient in executing on highly connected GPUs such as A100 with NVLink-3.0 [38]. Galvatron proposed an efficient algorithm for finding the optimal strategy to combine data-parallelism, model-parallelism, pipeline-parallelism and ZeRO optimization together [35, 59]. ZeRO can not be used for scaling large models due to the limited GPU memory and the model parallelism techniques, including Megatron-LM and Galvatron, are difficult to deploy in industry settings due to their complexity. ANGEL-PTM addresses these challenges by integrating ZeRO with hierarchical memory, which achieves flexibility as well as good scalability.

**Heterogeneous Training System.** With the evolution of deep learning, the type of models being trained has also evolved from early CNN-based models [21] to current transformer-based models [8, 11, 45, 46]. For CNN-based models, the memory usage is mainly occupied by activations, and research is mainly focused on optimizing single-GPU memory. Wang et al. [58] proposed evicting some activations by either recomputing them or offloading them

to CPU based on cost analysis. Nie et al. [41] proposed the tensor-splitting optimization for fine-grained memory management, which can help break memory bottlenecks and lead to more efficient execution plans. Transformer-based models have high memory usage due to their large model states, and research efforts have focused on optimizing GPU memory for distributed training. Ren et al. [50] proposed offloading optimizer computations and model states to the CPU to save GPU memory, and further combined this with a unique optimal offloading strategy and ZeRO-powered data parallelism. Rajbhandari et al. [49] introduced SSD storage into training and proposed a bandwidth-centric partitioning algorithm to distribute the model among all devices. Fang et al. [14] dynamically managed the model states during training via a chunk-based memory manager. Different from existing systems, ANGEL-PTM employs a fine-grained memory management via *Page* to reduce the memory fragments and improve the training efficiency.

**GPU Memory Management.** Many memory management optimizations have been proposed to pre-allocate most GPU memory and then manage the memory themselves, including paging [5], replacement caching [57] and memory pool [62]. Mosaic [5] provided application-transparent support for multiple page sizes to page-in and page-out. MultiQx-GPU [57] designed a cost-driven replacement policy for efficient executions of concurrent queries in GPU databases. Zhang et al. [62] proposed a memory pool, *CNMeM*, which utilizes lifetime semantics to reduce memory fragments and designed a heuristic algorithm to simplify the optimization problem. However, studies on paging, replacement caching, and unified memory address are not designed for deep learning training and do not utilize the special nature of tensor access patterns, while others about memory pool do not consider CPU memory. ANGEL-PTM utilized the life-time information to improve overall performance, which reduced the fragments and improve the overlap between different resources.

## 8 CONCLUSION

This work introduced ANGEL-PTM, an easy-to-use and highly-efficient deep learning systems for pre-training and fine-tuning tasks in Tencent. To be user-friendly and seamlessly scalable, we designed ANGEL-PTM with the basis of data parallelism, parameter sharding, and hierarchical memory. To fully utilize the memory and bandwidth during training, a *Page* abstraction was introduced to enable the fine-grained memory management along with a unified scheduling method was proposed to holistically manage the key operations during training. Moreover, we integrated the SSD storage to boost the pre-trained models to extreme scale and developed a lock-free updating mechanism to address the SSD I/O bandwidth bottleneck. Empirical results showed that ANGEL-PTM outperforms existing systems in terms of both maximum supported model scale and training throughput.

## 9 ACKNOWLEDGMENTS

This work is supported by the National Key Research and Development Program of China (No. 2020AAA0105200), the National Natural Science Foundation of China (No. 61832001 and U22B2037) and PKU-Tencent joint research Lab. Fangcheng Fu and Bin Cui are the corresponding authors.

## REFERENCES

- [1] 2022. ChatGPT: Optimizing Language Models for Dialogue. <https://openai.com/blog/chatgpt/>.
- [2] 2022. CLUE Benchmark. <https://www.cluebenchmarks.com/rank.html>.
- [3] 2022. Overall DGX A100 System Architecture. <https://www.microway.com/hpc-tech-tips/dgx-a100-review-throughput-and-hardware-summary/>.
- [4] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, Manjunath Kudlur, Josh Levenberg, Rajat Monga, Sherry Moore, Derek G. Murray, Benoit Steiner, Paul Tucker, Vijay Vasudevan, Pete Warden, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. 2016. TensorFlow: A System for Large-Scale Machine Learning. In *USENIX Symposium on Operating Systems Design and Implementation*. USENIX Association, 265–283.
- [5] Rachata Ausavarungnirun, Joshua Landgraf, Vance Miller, Saugata Ghose, Jayneel Gandhi, Christopher J. Rossbach, and Onur Mutlu. 2017. Mosaic: a GPU memory manager with application-transparent support for multiple page sizes. In *Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, 136–150.
- [6] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer Normalization. *arXiv preprint arXiv:1607.06450* (2016).
- [7] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258* (2021).
- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [9] Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174* (2016).
- [10] Ali Davoudian, Liu Chen, Hongwei Tu, and Mengchi Liu. 2021. A workload-adaptive streaming partitioner for distributed graph stores. *Data Science and Engineering* 6 (2021), 163–179.
- [11] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *The North American Chapter of the Association for Computational Linguistics*. Association for Computational Linguistics, 4171–4186.
- [12] Linhao Dong, Shuang Xu, and Bo Xu. 2018. Speech-transformer: a no-recurrence sequence-to-sequence model for speech recognition. In *IEEE international conference on acoustics, speech and signal processing*. IEEE, IEEE, 5884–5888.
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiuhua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*. OpenReview.net.
- [14] Jiarui Fang, Zilin Zhu, Shenggui Li, Hui Su, Yang Yu, Jie Zhou, and Yang You. 2022. Parallel Training of Pre-Trained Models via Chunk-Based Dynamic Memory Management. *IEEE Transactions on Parallel and Distributed Systems* 34, 1 (2022), 304–315.
- [15] William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch Transformers: Scaling to Trillion Parameter Models with Simple and Efficient Sparsity. *Journal of Machine Learning Research* 23, 120 (2022), 1–39.
- [16] Fangcheng Fu, Yuzheng Fu, Yihan He, Jiawei Jiang, Yingxia Shao, Ce Zhang, and Bin Cui. 2020. Don't Waste Your Bits! Squeeze Activations and Gradients for Deep Neural Networks via TinyScript. In *International Conference on Machine Learning*, Vol. 119. PMLR, 3304–3314.
- [17] Fangcheng Fu, Xupeng Miao, Jiawei Jiang, Huanran Xue, and Bin Cui. 2022. Towards Communication-efficient Vertical Federated Learning Training via Cache-enabled Local Update. *Proceedings of the VLDB Endowment* 15, 10 (2022), 2111–2120.
- [18] Fangcheng Fu, Yingxia Shao, Lele Yu, Jiawei Jiang, Huanran Xue, Yangyu Tao, and Bin Cui. 2021. VF<sup>2</sup>Boost: Very Fast Vertical Federated Gradient Boosting for Cross-Enterprise Learning. In *International Conference on Management of Data*. ACM, 563–576.
- [19] Jia-Ke Ge, Yan-Feng Chai, and Yun-Peng Chai. 2021. WATuning: a workload-aware tuning system with attention-based deep reinforcement learning. *Journal of Computer Science and Technology* 36, 4 (2021), 741–761.
- [20] Amir Gholami, Zhewei Yao, Sehoon Kim, Michael W Mahoney, and Kurt Keutzer. 2021. AI and Memory Wall.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *IEEE conference on computer vision and pattern recognition*. 770–778.
- [22] Dan Hendrycks and Kevin Gimpel. 2016. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415* (2016).
- [23] Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Chen, Hyoukjoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, et al. 2019. Gpipe: Efficient training of giant neural networks using pipeline parallelism. *Advances in neural information processing systems* 32 (2019), 103–112.
- [24] Sylvain Jeaugey. 2017. Nccl 2.0. In *GPU Technology Conference (GTC)*, Vol. 2.
- [25] Zhong Ji, Kexin Chen, Yuqing He, Yanwei Pang, and Xuelong Li. 2022. Heterogeneous memory enhanced graph reasoning network for cross-modal retrieval. *Science China Information Sciences* 65, 7 (2022), 172104.
- [26] Jiawei Jiang, Fangcheng Fu, Tong Yang, and Bin Cui. 2018. SketchML: Accelerating Distributed Machine Learning with Data Sketches. In *International Conference on Management of Data*. 1269–1284.
- [27] Youhe Jiang, Fangcheng Fu, Xupeng Miao, Xiaonan Nie, and Bin Cui. 2023. OSDP: Optimal Sharded Data Parallel for Distributed Deep Learning. *arXiv preprint arXiv:2209.13258* (2023).
- [28] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361* (2020).
- [29] Diederik P. Kingma and Jimmy Ba. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- [30] Dmitry Lepikhin, Hyoukjoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2021. GShard: Scaling Giant Models with Conditional Computation and Automatic Sharding. In *International Conference on Learning Representations*.
- [31] Mu Li, David G Andersen, Jun Woo Park, Alexander J Smola, Amr Ahmed, Vanja Josifovski, James Long, Eugene J Shekita, and Bor-Yiing Su. 2014. Scaling distributed machine learning with the parameter server. In *USENIX Symposium on Operating Systems Design and Implementation*. 583–598.
- [32] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 10012–10022.
- [33] Xupeng Miao, Xiaonan Nie, Yingxia Shao, Zhi Yang, Jiawei Jiang, Lingxiao Ma, and Bin Cui. 2021. Heterogeneity-Aware Distributed Machine Learning Training via Partial Reduce. In *International Conference on Management of Data*. ACM, 2262–2270.
- [34] Xupeng Miao, Xiaonan Nie, Hailin Zhang, Tong Zhao, and Bin Cui. 2023. Hetu: A highly efficient automatic parallel distributed deep learning system. *Science China Information Sciences* 66, 1 (2023), 1–2.
- [35] Xupeng Miao, Yujie Wang, Youhe Jiang, Chunan Shi, Xiaonan Nie, Hailin Zhang, and Bin Cui. 2022. Galvatron: Efficient Transformer Training over Multiple GPUs Using Automatic Parallelism. *arXiv preprint arXiv:2211.13878* (2022).
- [36] Xupeng Miao, Hailin Zhang, Yiming Shi, Xiaonan Nie, Zhi Yang, Yangyu Tao, and Bin Cui. 2021. HET: scaling out huge embedding model training via cache-enabled distributed framework. *Proceedings of the VLDB Endowment* 15, 2 (2021), 312–320.
- [37] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. Mixed Precision Training. In *International Conference on Learning Representations*. OpenReview.net.
- [38] Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prithvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, et al. 2021. Efficient large-scale language model training on gpu clusters using megatron-lm. In *International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 1–15.
- [39] Xiaonan Nie, Xupeng Miao, Shijie Cao, Lingxiao Ma, Qibin Liu, Jilong Xue, Youshan Miao, Yi Liu, Zhi Yang, and Bin Cui. 2021. Evomoe: An evolutionary mixture-of-experts training framework via dense-to-sparse gate. *arXiv preprint arXiv:2112.14397* (2021).
- [40] Xiaonan Nie, Xupeng Miao, Zilong Wang, Zichao Yang, Jilong Xue, Lingxiao Ma, Gang Cao, and Bin Cui. 2023. FlexMoE: Scaling Large-scale Sparse Pre-trained Model Training via Dynamic Device Placement. *Proceedings of the ACM on Management of Data* 1, 1 (2023), 1–19.
- [41] Xiaonan Nie, Xupeng Miao, Zhi Yang, and Bin Cui. 2022. TSPLIT: Fine-grained GPU Memory Management for Efficient DNN Training via Tensor Splitting. In *International Conference on Data Engineering*. IEEE, 2615–2628.
- [42] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155* (2022).
- [43] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* 32 (2019).
- [44] Yun Peng, Byron Choi, and Jianliang Xu. 2021. Graph learning for combinatorial optimization: a survey of state-of-the-art. *Data Science and Engineering* 6, 2 (2021), 119–141.
- [45] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language Models are Unsupervised Multitask Learners. *OpenAI blog* 1, 8 (2019), 9.
- [46] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the Limits of Transfer Learning with a Unified Text-to-Text Transformer. *Journal of*

- Machine Learning Research* 21 (2020), 1–67.
- [47] Samyam Rajbhandari, Conglong Li, Zhewei Yao, Minjia Zhang, Reza Yazdani Aminabadi, Ammar Ahmad Awan, Jeff Rasley, and Yuxiong He. 2022. DeepSpeed-MoE: Advancing Mixture-of-Experts Inference and Training to Power Next-Generation AI Scale. In *International Conference on Machine Learning*. PMLR, 18332–18346.
- [48] Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: Memory optimizations toward training trillion parameter models. In *International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 1–16.
- [49] Samyam Rajbhandari, Olatunji Ruwase, Jeff Rasley, Shaden Smith, and Yuxiong He. 2021. Zero-infinity: Breaking the gpu memory wall for extreme scale deep learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 1–14.
- [50] Jie Ren, Samyam Rajbhandari, Reza Yazdani Aminabadi, Olatunji Ruwase, Shuangyan Yang, Minjia Zhang, Dong Li, and Yuxiong He. 2021. ZeRO-Offload: Democratizing Billion-Scale Model Training. In *USENIX Annual Technical Conference*. USENIX Association, 551–564.
- [51] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. 2021. Scaling vision with sparse mixture of experts. *Advances in Neural Information Processing Systems* 34 (2021), 8583–8595.
- [52] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. IEEE, 10684–10695.
- [53] Shaden Smith, Mostofa Patwary, Brandon Norick, Patrick LeGresley, Samyam Rajbhandari, Jared Casper, Zhun Liu, Shrimai Prabhumoye, George Zerveas, Vijay Korthikanti, et al. 2022. Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. *arXiv preprint arXiv:2201.11990* (2022).
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [55] Danshi Wang, Chunyu Zhang, Wenbin Chen, Hui Yang, Min Zhang, and Alan Pak Tao Lau. 2022. A review of machine learning-based failure management in optical networks. *Science China Information Sciences* 65, 11 (2022), 211302.
- [56] Guibin Wang, YiSong Lin, and Wei Yi. 2010. Kernel fusion: An effective method for better power efficiency on multithreaded GPU. In *International Conference on Green Computing and Communications*. IEEE, IEEE, 344–350.
- [57] Kaibo Wang, Kai Zhang, Yuan Yuan, Siyuan Ma, Rubao Lee, Xiaoning Ding, and Xiaodong Zhang. 2014. Concurrent analytical query processing with GPUs. *Proceedings of the VLDB Endowment* 7, 11 (2014), 1011–1022.
- [58] Linnan Wang, Jinmian Ye, Yiyang Zhao, Wei Wu, Ang Li, Shuaiwen Leon Song, Zenglin Xu, and Tim Kraska. 2018. Superneurons: Dynamic GPU memory management for training deep neural networks. In *ACM SIGPLAN symposium on principles and practice of parallel programming*. ACM, 41–53.
- [59] Yujie Wang, Youhe Jiang, Xupeng Miao, Fangcheng Fu, Xiaonan Nie, and Bin Cui. 2023. Improving Automatic Parallel Training via Balanced Memory Workload Optimization. *arXiv preprint arXiv:2307.02031* (2023).
- [60] Yongqiang Wang, Abdelrahman Mohamed, Due Le, Chunxi Liu, Alex Xiao, Jay Mahadeokar, Hongzhao Huang, Andros Tjandra, Xiaohui Zhang, Frank Zhang, et al. 2020. Transformer-based acoustic modeling for hybrid speech recognition. In *International Conference on Acoustics, Speech and Signal Processing*. IEEE, IEEE, 6874–6878.
- [61] Hua-Peng Wei, Ying-Ying Deng, Fan Tang, Xing-Jia Pan, and Wei-Ming Dong. 2022. A Comparative Study of CNN-and Transformer-Based Visual Style Transfer. *Journal of Computer Science and Technology* 37, 3 (2022), 601–614.
- [62] Junzhe Zhang, Sai Ho Yeung, Yao Shu, Bingsheng He, and Wei Wang. 2019. Efficient memory management for gpu-based deep learning systems. *arXiv preprint arXiv:1903.06631* (2019).
- [63] Lianmin Zheng, Zhuohan Li, Hao Zhang, Yonghao Zhuang, Zhifeng Chen, Yanping Huang, Yida Wang, Yuanzhong Xu, Danyang Zhuo, Eric P Xing, et al. 2022. Alpa: Automating Inter-and {Intra-Operator} Parallelism for Distributed Deep Learning. In *USENIX Symposium on Operating Systems Design and Implementation*. USENIX Association, 559–578.