



DLRover-RM: Resource Optimization for Deep Recommendation Models Training in the Cloud

Qinlong Wang*
Independent Researcher
wangql1201@outlook.com

Tingfeng Lan*
Yinghao Tang
Sichuan University
tafflan2001@gmail.com
yinghaotang2001@gmail.com

Bo Sang
Independent Researcher
xuanying2019@gmail.com

Ziling Huang
Yiheng Du
Sichuan University
youxingling1@gmail.com
yihengdu42@gmail.com

Haitao Zhang
Jian Sha
Independent Researcher
t12345@gmail.com
nash635@gmail.com

Hui Lu
The University of Texas at Arlington
hui.lu@uta.edu

Yuanchun Zhou
Chinese Academy of Science
zyc@cnic.cn

Ke Zhang
Independent Researcher
wang_ke@126.com

Mingjie Tang
Sichuan University
tangrock@gmail.com

ABSTRACT

Deep learning recommendation models (DLRM) rely on large embedding tables to manage categorical sparse features. Expanding such embedding tables can significantly enhance model performance, but at the cost of increased GPU/CPU/memory usage. Meanwhile, tech companies have built extensive cloud-based services to accelerate training DLRM models at scale. In this paper, we conduct a deep investigation of the DLRM training platforms at ANTGROUP and reveal two critical challenges: *low resource utilization* due to suboptimal configurations by users and *the tendency to encounter abnormalities* due to an unstable cloud environment. To overcome them, we introduce DLROVER, an elastic training framework for DLRMs designed to increase resource utilization and handle the instability of a cloud environment. DLROVER develops a resource-performance model by considering the unique characteristics of DLRMs and a three-stage heuristic strategy to automatically allocate and dynamically adjust resources for DLRM training jobs for higher resource utilization. Further, DLROVER develops multiple mechanisms to ensure efficient and reliable execution of DLRM training jobs. Our extensive evaluation shows that DLROVER reduces job completion times by 31%, increases the job completion rate by 6%, enhances CPU usage by 15%, and improves memory utilization by 20%, compared to state-of-the-art resource scheduling frameworks. DLROVER has been widely deployed at ANTGROUP and processes thousands of DLRM training jobs on a daily basis. DLROVER is open-sourced and has been adopted by 10+ companies.

PVLDB Reference Format:

Qinlong Wang, Tingfeng Lan, Yinghao Tang, Bo Sang, Ziling Huang, Yiheng Du, Haitao Zhang, Jian Sha, Hui Lu, Yuanchun Zhou, Ke Zhang, Mingjie Tang. DLRover-RM: Resource Optimization for Deep Recommendation Models Training in the Cloud. PVLDB, 17(12): 4130 - 4144, 2024.

* Qinlong Wang and Tingfeng Lan contributed equally to this work. Jian Sha and Mingjie Tang are the corresponding authors.

doi:10.14778/3685800.3685832

PVLDB Artifact Availability:

The source code, data, and/or other artifacts have been made available at <https://github.com/intelligent-machine-learning/dlrover>.

1 INTRODUCTION

Deep learning based recommendation models (DLRM) are prevalent in recommendation scenarios [18, 26, 33, 46, 51, 79]. For example, Meta uses DLRMs for advertisement recommendation to optimize ad content for individual users, aiming to maximize click-through rates and advertising revenue [41]. The training of DLRMs at Meta, Amazon, Alibaba, and ANTGROUP can account for over 50% of the total AI training cycles in cloud data centers [8, 15, 20, 36].

A typical DLRM uses *embedding tables* to manage sparse categorical features (e.g., User IDs) and several *deep neural networks* (DNNs) to improve the generalization of the models (§2.1). As the accuracy of a DLRM often improves with larger embeddings, which incorporate more feature data points, the size of DLRM embeddings has been steadily expanding, reaching up to terabytes with billions of embedding vectors [8, 35, 78]. Tech companies build extensive cloud-based services to accelerate training these models at scale, e.g., with thousands of computing nodes [48, 61, 66]. Unfortunately, we observed from our cloud-based cluster that the resource utilization of over 80% DLRM training jobs is under 50%, indicating a significant *underutilization* and waste of computation resources. Moreover, we observed that *high instability* in cloud environments leads DLRM training to: 1) experience a high failure rate and 2) frequently encounter abnormalities (e.g., stragglers) (§2.2).

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@vlldb.org. Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment. Proceedings of the VLDB Endowment, Vol. 17, No. 12 ISSN 2150-8097. doi:10.14778/3685800.3685832




Figure 1: (a) The operator’s time proportion in multiple DLRM training jobs. (b) The memory demand of one job.

In this paper, we focus on developing a highly resource-efficient and reliable *DLRM training system*, especially in a cloud environment, where failures are common and resource availability varies dynamically. Such a training system should be capable of training a multitude of DLRMs (e.g., 1,000s) concurrently with the following key goals: maximizing resource utilization (e.g., CPU and memory), achieving rapid training speeds, and ensuring fault tolerance. Achieving these stringent goals requires the training system to accurately *allocate* computational resources to individual DLRM training jobs and *schedule* these jobs in the cloud elastically and robustly. However, the unique characteristics of DLRMs, in combination with the dynamic nature of the cloud, make resource allocation and scheduling for DLRM training jobs extremely challenging.

Unlike traditional compute-intensive deep learning (DL) models used in computer vision (CV) and natural language processing (NLP), DLRM training incurs *massive I/O operations* in addition to its compute-intensive operations (e.g., matrix multiplication for DNNs). These I/Os are largely due to frequent *lookups* to embedding tables, consuming 30-48% of the training time (see Fig. 1(a)). Existing schedulers [45, 57, 58], without considering such a unique blend of I/O and computation operations in the DLRM training, fall short in ensuring optimal resource utilization and training efficiency.

Another unique characteristic of DLRMs lies in that their embedding tables are notably *memory-intensive*. DLRMs can easily demand tens of terabytes of memory (§2.1). As user-targeted applications evolve, the size of embedding tables keeps increasing [40, 68]. For example, the memory usage of a typical DLRM job can surge to over 2.3TB within just 15 hours (see Fig. 1(b)). Consequently, there is a significant risk of hitting out-of-memory (OOM) for a DLRM job if the allocated resources cannot quickly adjust to its increased memory demand. We observed that, in the production environment of ANTGROUP, thousands of jobs (5% - 8%) had been derailed due to OOM, leading to compromised user satisfaction and suboptimal cluster performance.

The high *instability* of the cloud environment necessitates frequent scaling resources to adapt to the ever-changing cloud [19, 34, 75]. To achieve this, traditional DL schedulers, like [45, 57], require stopping a job and then restarting it with adjusted resources. This stop-and-restart process often takes up to tens of minutes (§2.2),

highlighting the need for a more efficient approach. Additionally, model training with elastic training frameworks may result in inconsistent model accuracies due to stale gradients being submitted [16] or disruption of training data (§2.2).

To tackle these challenges, we introduce DLROVER-RM, a cloud-based deep learning training system designed for DLRMs. DLROVER-RM takes the runtime training information into account for accurately allocating and elastically scheduling resources for training jobs along with a bunch of novel mechanisms, including *dynamic data sharding*, *flash-checkpoint*, *seamless migration*, and pre-adjustment-based, *OOM prevention*. Together, DLROVER-RM attains exceptional throughput, high resource utilization, and robust fault tolerance. In summary, we have made the following contributions:

- 1) We build a *resource-performance* model by considering I/O overhead and computation demands during DLRM training. With this model, we design a three-stage algorithm that can dynamically allocate resources during the whole cycle of DLRM training and significantly reduce the job completion time.
- 2) We invent a *dynamic data sharding* mechanism to maintain the model quality when scaling or a job failure happens in the cloud. We further develop *seamless migration* and *flash-checkpoint* strategies to reduce the overhead of scaling jobs. We also develop an *OOM-prediction* mechanism to prevent OOM.
- 3) We implement DLROVER-RM as a cloud-native auto-configuration service and open-source all technical implementations. Thereby, end-users can train DLRM jobs in the production environment without concern for resource configuration and job failures.
- 4) We thoroughly evaluate DLROVER-RM with thousands of jobs collected from months of various DLRM training workloads in a production environment equipped with more than 62K CPU and 3.24PB memory. The evaluation shows that DLROVER-RM improves the CPU utilization by 21.0% to 27.6% and memory utilization by 17.3% to 31.6%, and reduces job completion time by 30.9% without compromising model accuracy.

2 BACKGROUND AND MOTIVATION

In this section, we briefly introduce DLRMs and the DLRM training platform at ANTGROUP (§2.1). We then discuss the key issues when training DLRMs and the challenges to address them through investigating the unique characteristics of DLRMs and sharing the observations from our DLRM system deployed at ANTGROUP (§2.2).

2.1 DLRM Training at AntGroup

At AntGroup, DLRMs are extensively applied to scenarios such as service/content search, marketing vouchers, Tab3 video recommendation, and advertising [23, 74]. Within our cloud-based cluster, DLRM training jobs account for more than 70% of the total training jobs daily, consuming a significant amount of the cluster resources. Optimizing DLRM training holds substantial importance for the effective utilization of the cluster resources.

Overview of DLRM. Fig. 2 illustrates a typical model architecture of DLRMs: First, it integrates fully connected deep neural networks (DNNs) to capture continuous dense features like timestamps. Further, it uses embedding tables to transform various categorical sparse features, such as user and video IDs, into low-dimensional dense representations. These embedding tables represent the *sparse*

part of DLRMs, while the rest of the models (e.g., DNNs) represent the *dense part*. DLRMs take both sparse categorical and dense features as inputs for model training.

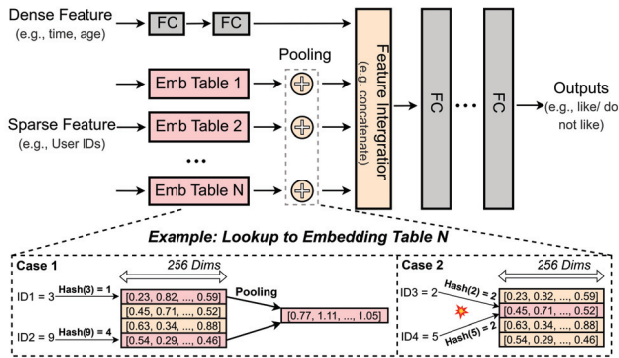


Figure 2: A typical DLRM architecture consists of neural networks, which make up the dense part, combined with memory-intensive embedding tables, forming the sparse part. The portion within the dashed box highlights examples of embedding table lookup in forward propagation of DLRM.

As illustrated in Fig. 2, each categorical feature has its own embedding table. A data point, or an instance, for a categorical feature is mapped to a specific row in this table where its embedding vector is stored. For instance, consider a user with ID 3, representing an instance of the "User IDs" feature. The embedding for this user is located at the index $(hash(3) \bmod M)$ within the "User IDs" embedding table, which consists of M rows. Embedding vectors that are being accessed undergo an integration process into a singular dense vector through element-wise *pooling* operations. These operations typically use either the sum or maximum value of the vectors. This pooled embedding vector is further concatenated with the intermediate output derived from the dense features. This combined output forms the input for the subsequent DNN layers. In practice, a DLRM might comprise thousands of embedding tables, with some tables having millions of rows [27].

Table 1: DLRM training cost comparison between the hybrid(GPU-CPU) and CPU-only approach on AWS.

Model	Device	Time	Unit Price	Samples/ Price	CPU Util	GPU Util
Wide& Deep	CPU	1.41h	0.53usd/h	3.4m/usd	≈ 33%	—
	Hybrid	0.98h	3.59usd/h	1.9m/usd	≈ 26%	≈ 3%
DeepFM	CPU	1.53h	0.53usd/h	3.1m/usd	≈ 34%	—
	Hybrid	0.95h	3.59usd/h	2.1m/usd	≈ 28%	≈ 4%

At AntGroup, we use the parameter server (PS) architecture [31] for DLRM training as it serves as a *de facto* framework for this purpose [8, 21, 25, 38, 77].

Hardware Selection: CPU-only. In practice, there are two approaches in training recommendation models: 1) a CPU-only approach by exclusively using CPUs; 2) a CPU-GPU hybrid approach

by employing CPUs for handling the embedding data while GPUs for executing data-parallel neural networks [9]. While the CPU-GPU approach is the preferred choice for companies like Meta [8, 27], we lean toward the CPU-only approach due to the following reasons: First, CPU resources are more readily accessible and cost-effective compared to GPUs in a production environment – our clusters have an abundance of CPU resources, e.g., over 200k cores, whereas housing fewer than 1,000 high-performance GPU cards. Further, given the scarcity of GPU resources, it is desired to maximize their utilization. However, DLRM training involves intensive I/O operations, including 1) transferring embedding data between CPUs and GPUs (up to 22% [9] of the total training time), and 2) conducting a myriad of lookup operations to the embedding tables (over 30% of the total training time). Such massive I/O operations render GPUs underutilized. As shown in Table 1, when training the two most common DLRM models at ANTGROUP [11, 54, 74, 74]: 1) the average GPU utilization under the CPU-GPU hybrid approach is lower than 3%, while 2) the CPU-only approach can train more data at a unit price. Therefore, we use CPU to train DLRMs and base the remaining discussion on the CPU-only approach.

Cloud Environment: Workload Consolidation. Table 2 shows that at AntGroup, different types of jobs (i.e., training, serving, and stream processing jobs) are running in the same cluster and sharing the resources [52]. For isolation and security reasons, the DLRM system is unaware of the resource usage by other services, as well as the overall resource consumption of the cloud. That said, the DLRM system has no direct control over the cluster resources and has to request resources from the cluster resource scheduler when scheduling or scaling out a DLRM training job.

Table 2: Statistic of Jobs at ANTGROUP

Job Type	Count	vCPU	CPU Util	MEM
Training	62K	600K	20%	0.9PB
Stream Processing	43K	450K	15%	0.63PB
Inference Service	3K	300K	10%	0.41PB
Search Service	0.9K	200K	15%	1.2PB
Other	2K	50K	10%	0.1PB

2.2 Challenges of DLRM Training at AntGroup

By collecting significant training task data from the largest machine learning platform at AntGroup, we identified two primary issues with DLRM training: *low resource utilization* and *high cloud instability*. In this section, we provide a detailed analysis of these two problems and highlight the challenges in addressing them.

Low Resource Utilization. As depicted in Fig. 3, over 80% of the jobs in our cluster had CPU and memory utilization rates below 50% back in 2021, resulting in a significant waste of cluster resources. The core factor causing this is *suboptimal configurations by users*. Specifically, in a typical cloud environment, cloud users need to specify a fixed amount of resources before deploying their cloud-based services [2, 57, 60]. Similarly, our previous training system running in a cloud-based cluster also needed such inputs, i.e., resource configurations, from system users (e.g., ML engineers or data scientists). Such resource configurations were used to guide the

training system for resource allocation during DLRM training. Users typically resorted to a time-consuming *trial-and-error* approach to determine these configurations – by manually (re-)running their jobs multiple times with varying resource configurations in search of the "optimal" one. Oftentimes, these user-provided configurations tended to ask for "more-than-needed" resources to avoid job failures during training, resulting in inefficient use of resources.

To overcome this, instead of relying on user-provided suboptimal resource configurations, we need an approach allowing the DLRM system to automatically allocate and dynamically adjust resources for training jobs for high resource utilization. Note that for distributed training, resource adjustment includes changing the number of nodes (horizontal scaling) and the resources of each node (vertical scaling) [60]. This is nontrivial due to two main challenges:

- **Timely Meeting Memory Demands.** The memory requirement for storing embedding tables can surge up to 10s of terabytes [27] in a short period. As shown in Fig. 1(b), the memory usage of embedding tables in a typical DLRM model can spike to more than 2.3TB within 15 hours. This renders DLRMs vulnerable to out-of-memory (OOM) issues if memory allocation cannot timely meet the model’s demands – we observed 5%-8% of jobs suffering from OOM in our production environment, greatly affecting the overall cluster resource efficiency.
- **Precisely Allocating CPU Resources.** In DLRM training, extensive I/O operations impact both model training efficiency and job CPU utilization. As shown in Fig. 1(a), the lookup operations can account for 30%-48% of the training duration in a single iteration. Conventional deep learning resource schedulers [45, 57, 58] fall short of handling this unique training process, often overlooking the lookup latency and allocating inappropriate CPU resources for training jobs.

High Cloud Instability. Unlike a dedicated cluster (for a single-purpose service), the cloud environment has a much higher job failure rate [39, 69]. Statistically, we observed that the daily failure rate for a simple job (e.g., hosted in a single Kubernetes pod) in our cloud-based cluster is 1.5% due to network errors, node malfunction, etc. The failure rate increases exponentially for a more complex distributed job with hundreds of components. For example, the daily failure rate for a job with 50 pods increases dramatically to $1 - (1 - 0.015)^{50} = 53.03\%$. Moreover, in our cloud-based cluster, different services co-exist, sharing the same cloud resources (§2.1). Compared to other higher-priority services, e.g., online services, DLRM training is typically labeled with a lower priority. When higher-priority services encounter workload spikes, the cluster scheduler preempts resources allocated to the DLRM system, resulting in the failure of DLRM training jobs or the emergence of stragglers (e.g., slow workers) due to insufficient resources.

To address this, our system needs the capability to 1) frequently scale up/down training jobs to adapt to the changing cloud environment and 2) detect failed nodes and recover them swiftly. However, this is also not trivial due to the following two main challenges:

- **Ensuring Consistent Model Quality.** Elastic training frameworks can enhance training throughput by dynamically scaling training jobs up or down. However, elasticity operations (e.g., increasing/decreasing the number of worker nodes and/or adding/shrinking computational resources to a worker) can also




Figure 3: DLRM jobs’ resource utilization and pending time derived from cluster traces in ANTGROUP.

lead to inconsistent job configurations (e.g., batch size and the number of parallel workers) and/or changed data sequences. For example, some slow workers may submit too many stale gradients to Pses, causing instability in parameter updates; some workers might miss specific data batches due to failures, or the training data sequence could be disrupted during scaling operations (§5.1). These inconsistent configurations and disruptions could further compromise the consistency of model training quality [16, 32], especially in asynchronous training [12].

- **Providing Fast Elasticity.** Swift scaling operations are essential for accelerating jobs (e.g., by allocating more resources) and managing instability (e.g., by addressing slow workers in a job group). Conventional DL schedulers [29, 45] involve a *stop-and-restart* operation to scale up/down a job – by saving the job’s checkpoint (to hard disks) and restarting the job with adjusted resources/configurations, e.g., reallocating training data (to scale workers) or re-partitioning the model (to scale Pses). The *stop-and-restart* operation is very *costly*: First, checkpointing a job to remote disk storage (RDS) typically takes 5-10 minutes [65]. Further, the scheduler takes another 5-10 minutes to complete the necessary preparation before restarting, including submitting a new job YAML, requesting resources for the new pods, pulling images from the registry, and re-establishing the code environment. Under conditions of resource scarcity (e.g., daytime [19]), the duration can extend beyond 30 minutes. Last, loading the checkpoint from RDS and restarting the training takes another 5-10 minutes. Altogether, the whole process could consume tens of minutes, introducing high overhead for DLRM training (§5.2).

3 OVERVIEW OF DLROVER

In this section, we present the architecture overview of DLROVER¹ and highlight its key design objectives.

Design Objectives. DLROVER focuses on efficiently training a multitude of DLRMs simultaneously in a dynamic, shared cloud environment. It dynamically schedules computational and memory resources for DLRM training jobs to optimize training throughput and resource utilization while mitigating the job failure rate. The design objectives of DLROVER are to answer the two key questions:

- **Automated Resource-Performance Optimization:** How can DLROVER accurately allocate resources for DLRM training jobs –

¹In this work, "DLRover" and "DLRover-RM" are used interchangeably.

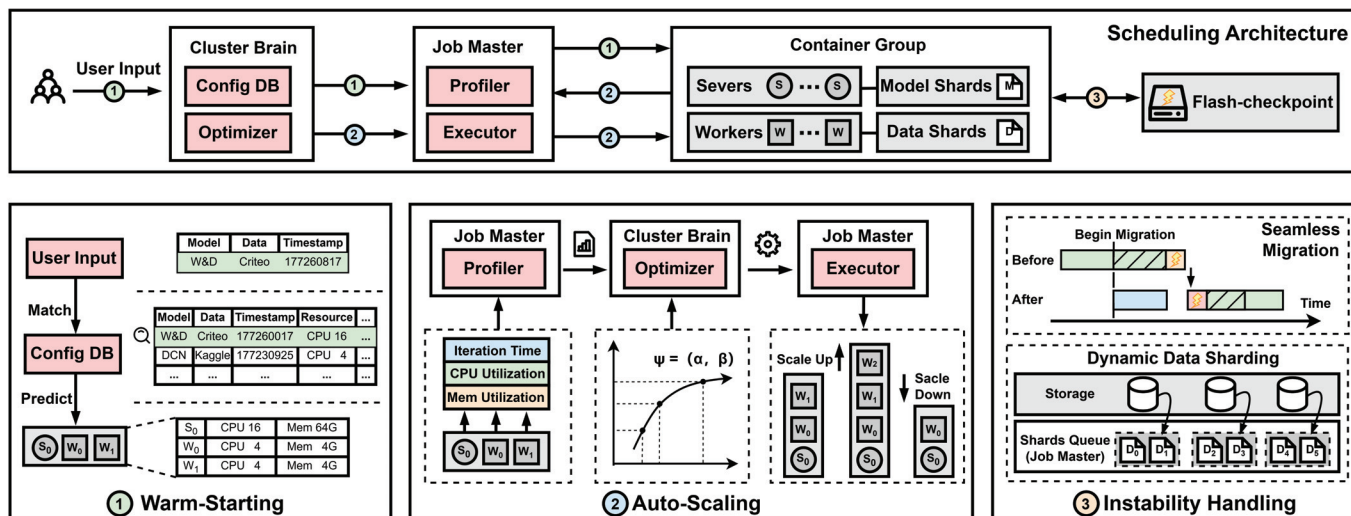


Figure 4: Overview of DLROver and Model Training Workflow

without user-handcrafted configurations – to maximize training throughput and minimize resource cost (§4)

- **High Stability Assurance:** How does DLROVER overcome the dynamic nature of the cloud environment to achieve robust execution of DLRM training jobs with low job failure rate and high fault tolerance (§5)

Architecture Overview. DLROVER is based on the parameter server architecture (see §2.1) in our production cloud environment. As illustrated in Fig. 4, DLROVER consists of two main components: 1) a cluster-level central coordinator, called *cluster brain*; and 2) a group of job-level distributed training agents, called *job master*:

- The **cluster brain** comprises two subcomponents: the *optimizer* and *config database* (*config DB*). The *optimizer* receives the runtime profiles (e.g., CPU and memory utilization) of training jobs from each *profiler* periodically. With such information, the *optimizer* creates resource plans and sends them to the corresponding *executors*. Meanwhile, the *config DB* stores the information as the historical job traces.
- Each **job master** also comprises two subcomponents: the *profiler* and *executor*. The *profiler* monitors and collects runtime information for each job (i.e., from its workers and PSes) in a fixed interval and reports it to the *optimizer* of the *cluster brain*. The *executor* feeds data shards (e.g., a slice of training data) to the job’s workers (e.g., hosted in pods) for training.

Life Cycle of Training. As detailed in Fig. 4, upon submission of a job by the user, the *cluster brain* quickly learns the job’s characteristics – by leveraging relevant historical data from the *config DB* – and then generates an initialization (*warm-starting*) resource plan with the relative configuration (e.g., the number of CPU for each worker/PS) and similarity information (e.g., time series information)(①). Note that, at this moment, we choose a reasonable configuration near the optimal configuration (hence, with fewer scaling operations and shorter scaling times for auto-scaling) instead of

pursuing an optimal configuration. Subsequently, the *cluster brain* sends the *warm-starting* resource plan to the respective *job master* for job initialization.

During job running, the *profiler* profiles the job’s runtime statistics and reports them back to the *optimizer* periodically. With such updated runtime information, the *optimizer* can generate a refined resource plan, upon which the *executor* dynamically adjusts the number of workers and/or PSes, and their resource configurations accordingly, i.e., the execution plan (②).

DLROVER further provides a set of reliable instability handling mechanisms to ensure the robust execution of training jobs (③). For failed/slow workers, DLROVER implements a *dynamic data sharding* mechanism to redistribute missed data and rebalance workloads between workers (§5.1). For failed/slow PSes, DLROVER devises a *seamless migration* with in-memory checkpoint, named *flash-checkpoint*, to minimize the overhead in failure recovery and job migration (§5.2). We will look closer at this three-stage job life cycle of DLROVER (① - ③) in our three-stage algorithm (§4.3).

4 EXPLORING OPTIMAL CONFIGURATIONS

In this section, we present how DLROVER accurately allocates resources for DLRM training. It first builds a resource-performance model for DLRM training (§4.1). Then it formulates the optimizing objective (§4.2). Finally, it proposes a novel three-stage algorithm to guide resource allocation (§4.3).

4.1 Resource-Performance Modeling

Throughput Modeling. The throughput of a DLRM training job represents the number of samples processed per unit of time. To model and predict the throughput, we divide one iteration time T_{iter} into two parts: computation time T_{comp} and communication time T_{comm} . Let w denote the number of workers. Each worker

consumes a mini-batch of data with size m per iteration. Then, we formally model the throughput, denoted as Ψ_{thp} , as follows:

$$\Psi_{thp} = \frac{w \cdot m}{T_{comp} + T_{comm}} \quad (1)$$

Note that the batch size m remains unchanged during training.

Computation Time Modeling. The computation time T_{comp} includes two parts: the workers compute gradients (T_{grad}) and then the PSes update parameters using corresponding gradients (T_{upd}).

In each iteration, the gradient computation workload is proportional to the number of samples processed (m). The gradient computation rate is proportional to the number of parallel computing CPU cores in each worker (λ_w). Since the gradient update time (i.e., T_{grad}) can be calculated as the workload divided by the rate, we formulate T_{grad} as:

$$T_{grad}(\alpha_{grad}, \beta_{grad}) = \alpha_{grad} \cdot \frac{m}{\lambda_w} + \beta_{grad} \quad (2)$$

where α_{grad} and β_{grad} are learnable parameters representing how T_{grad} scales with m and λ_w linearly.

The parameter updating workload is proportional to the number of workers (w) as each worker computes one copy of the gradient and submits it to the PSes. The workload is also inversely proportional to the number of PSes (p) as all the PSes share these gradients. The update rate is proportional to the number of parallel computing CPU cores in each PS (λ_p). Therefore, we formulate T_{upd} as follows:

$$T_{upd}(\alpha_{upd}, \beta_{upd}) = \alpha_{upd} \cdot \frac{w}{p \cdot \lambda_p} + \beta_{upd} \quad (3)$$

where α_{upd} and β_{upd} are learnable parameters representing how T_{upd} scales with p and λ_p linearly.

Communication Time Modeling. Communication time T_{comm} includes two parts: 1) Workers pull parameters from PSes and push gradients to PSes to synchronize parameters (i.e., T_{sync}); 2) Workers lookup embeddings from PSes for gradient computation (i.e., T_{emb}).

For parameter synchronization, the network traffic (the amount of network communication data) between workers and PSes is twice the size of the model parameters (M) because both pulling and pushing operations transfer one copy of the data with size M . The network bandwidth B is shared by w workers. The pushing-and-pulling workload is divided by the number of PSes p as the model parameters and gradients are distributed across all PSes. Therefore, we formulate T_{sync} as:

$$T_{sync}(w, p) = \alpha_{sync} \cdot \frac{M/p}{B/w} + \beta_{sync} \quad (4)$$

where α_{sync} and β_{sync} are learnable parameters representing how T_{sync} scales with w and p linearly. Here, the model size M and network bandwidth B are considered constants during a short time.

For embedding lookups, the network traffic is proportional to the samples m and the dimensions of the embedding table D , and it is shared by all PSes as the embedding table are distributed across PSes. Therefore, we formulate T_{emb} as:

$$T_{emb}(\alpha_{emb}, \beta_{emb}) = \alpha_{emb} \cdot \frac{m \cdot D}{p} + \beta_{emb} \quad (5)$$

where α_{emb} and β_{emb} are learnable parameters representing how T_{emb} scales with m and p linearly. Here, the embedding table dimension D is fixed (hence being a constant) in the initialization of the embedding table.

Finally, we formally model throughput Ψ_{thp} with a function \mathcal{F} represented by the tuple of learnable parameters as follows:

$$\Psi_{thp} = \mathcal{F}(\alpha_{grad}, \beta_{grad}, \alpha_{upd}, \beta_{upd}, \alpha_{emb}, \beta_{emb}, \alpha_{sync}, \beta_{sync}) \quad (6)$$

4.2 Optimization Formulation

Given the resource-performance model in Eqn. 6, we formulate our optimization objective based on the "Resource Cost" (i.e., for additional allocated resources) and the "Throughput Gain" (i.e., from additional allocated resources) when scaling DLRM training jobs. The goal of our optimization is to minimize the "Resource Cost" while maximizing the "Throughput Gain".

Resource Cost Function (RC). The resource scaling set, denoted as A , represents the additional allocated resources to speed up a training job (e.g., the number of CPUs). Each type of resource in the set is denoted as a_r (i.e., $A = \{a_0, a_1, \dots, a_r\}$). Let $Money(a_r)$ denote the expense a user should spend to allocate resource a_r . The "Resource Cost" can be formulated as the sum of all resources' expenses (e.g., CPU and memory) :

$$RC(A) = \sum_{a_r \in A} a_r \times Money(a_r) \quad (7)$$

Throughput Gain Function (TG). We denote "Throughput Gain" as the increased throughput that benefits from additional allocated resources. To illustrate, if we add 2 CPUs to a worker, the throughput might increase by 10 samples per second, namely the *throughput gain*. However, increasing workers' resources, especially in cloud environments, practically comes with overheads (e.g., the time required to start a new worker equipped with 32 CPUs and 128GB of memory). Therefore, we formulate the *Throughput Gain* as:

$$TG(A) = \Delta\Psi_{thp} - Overhead(A) \quad (8)$$

Here, $\Delta\Psi_{thp}$ represents the ideal increase in throughput if we neglect scaling overheads. $Overhead(A)$ is the wasted training time caused by scaling the job with A . This is estimated through statistical analysis based on the resource information of historical jobs within the cluster (e.g., the time required to start a worker/PS).

Multi-Objective Optimization. Given the two functions Eqn. 7 & 8, our goal is to find an optimal resource allocation set A , which minimizes the "Resource Cost" while maximizing the "Throughput Gain". We formulate the optimization problem as follows:

$$\text{Objective: } \arg \min_A (RC(A), \frac{1}{TG(A)}) \quad (9)$$

Given that Equ.9 is neither linear nor convex, the problem cannot be solved using linear programming or convex optimization techniques and is NP-hard in general. To address this, we develop a heuristic *auto-scaling* algorithm (see §4.3).

4.3 3-Stage Auto-Scaling Optimization

Intuition. Given a DLRM training job, we first profile its runtime information for fitting the resource-performance model (Eqn. 6). Based on the model, DLROVER generates an optimal resource plan with an *auto-scaling* algorithm, aiming to achieve the optimization objective (Eqn. 9) (i.e., Scaling Stage ② in Fig. 4).

To make the job training more robust in practice, we design a pre-scaling stage (stage ①) to *warm-starting* the job and a post-scaling stage (stage ③) to handle the cloud instability. Specifically, compared to scaling the training job from scratch (i.e., cold start), users hope to see the submitted job performing well upon submission rather than waiting through a prolonged scaling process. Thus, we introduce a pre-scaling stage to allocate suitable start-up configurations. On the other hand, even provided with optimal resources, training jobs still encounter performance degradation (e.g., stragglers) due to cloud instability. Consequently, we introduce a post-scaling stage to ensure smooth training in the cloud.

Algorithm 1: Warm-Starting

Input: Historical Configurations \mathcal{D} , Incoming Job \mathcal{J} ,
Exponential Smoothing Function \mathcal{E} ,
Smoothing Factor μ ($0 < \mu < 1$)

Output: Warm-Starting Resource Allocation \bar{A}^{k-1}

Identify Top-K Similar Jobs of \mathcal{J} With MetaData:

1 $\{A^0, A^1, \dots, A^{k-1}\} \leftarrow$ top- k similar job configuration in \mathcal{D} ;

Initialize Smoothing for Configuration:

2 Rank $\{A^0, A^1, \dots, A^{k-1}\}$ with similarity;

3 Initialize smoothed configuration: $\bar{A}^0 = A^0$;

4 **for** $i = 1$ to $k - 1$ **do**

5 Apply \mathcal{E} to get the smoothed configuration \bar{A}^i :

6 $\bar{A}^i \leftarrow \mu \times A^i + (1 - \mu) \times \bar{A}^{i-1}$;

7 **end**

8 **return** \bar{A}^{k-1} ;

① Pre-scaling Stage: Warm-Starting.

As shown in Algorithm 1, we adopt a *warm-starting* algorithm to identify a suitable start-up resource configuration. We first use the job's features (e.g., model metadata) to collect top- k similar jobs. Specifically, given historical job configurations stored in the *Configuration Database* \mathcal{D} , the algorithm first calculates the similarity in each type of feature and then gets the top- k similar job configurations where the configuration of i -th similar job is denoted as A^i (A^{k-1} is the job configuration with highest similarity). We first initialize the target configuration set as $\bar{A}^0 = A^0$. Subsequently, we use the *Exponential Smoothing Function* \mathcal{E} to generate the smoothed configuration \bar{A}^i in an iterative manner. Formally, for each A^i , we calculate \bar{A}^i as follows:

$$\mathcal{E} : \bar{A}^i = \mu \times A^i + (1 - \mu) \times \bar{A}^{i-1} \quad (10)$$

where the smoothing factor μ balances the influence of historical configurations, determining the weight between the job's configuration A^i and the result of last iteration \bar{A}^{i-1} . Lastly, we use the final iteration result \bar{A}^{k-1} as the start-up job configuration.

② Scaling Stage: Auto-Scaling.

We aim to auto-scale training jobs in this stage according to our resource-performance model (Eqn. 6). Auto-scaling includes three steps: 1) online model fitting, 2) job-level resource plan candidate generation, and 3) cluster-level weighted greedy selection.

Online Model Fitting. As detailed in §4.1, we represent the throughput of a training job by a group of α and β parameters (e.g., α_{grad} , β_{grad}). To build our resource-performance model, these parameters can be fitted based on runtime profiles. DLROVER continuously monitors the time taken for each iteration T_{iter} to measure the throughput Ψ_{thp} as in Eqn.1. At a fixed interval, DLROVER refines the groups of α and β using the accumulated data, minimizing the root mean squared logarithmic error (RMSLE) between the theoretical model and the actual data (by employing Non-Negative Least Squares (NNLS)[4]). Note that all parameters (α, β) are bound to remain non-negative.

Job-Level Resource Plan Candidates Generation. After model fitting, we learn the function \mathcal{F} as Eqn. 6 representing the relation between resource allocation A and throughput Ψ_{thp} . We utilize NSGA-II [3] to generate resource allocation plans that meet the Pareto Frontier. The Pareto Frontier represents the set of all optimal allocations that cannot be improved on one dimension without worsening another. For example, we can not increase "Throughput Gain" without increasing the "Resource Cost". NSGA-II is an evolutionary algorithm known for its rapid convergence to the Pareto Frontier in low-dimensional multi-objective problems.

Cluster-Level Weighted Greedy Selection. With all the optimization plan candidates for each job, we employ *weighted greedy selection* to determine the final execution plan for each job. We denote the "Resource Efficiency" for allocating resources A^j to a specific job j as $RE(A^j)$. $RE(A^j)$ is a function of the "Throughput Gain" we get from the allocation, normalized by its associated "Resource Cost", mathematically represented as:

$$RE(A^j) = \frac{TG(A^j)}{RC(A^j)} \quad (11)$$

To determine a set of efficient cluster-wide resource allocations, we maximize the weighted benefit sum of each job:

$$\text{Weighted Greedy: } \arg \max_{A^j} \sum_{j \in J} RE(A^j) \cdot WG(A^j) \quad (12)$$

$$\text{subject to: } \sum_{j \in J} A^j \leq S \quad (13)$$

Here, J represents the jobs needing reallocation; S denotes the total resources; $WG(A^j)$ denotes the priority value determined by a range of priority algorithms tailored to the cluster's preference. In our cluster, we use the remaining time for each training job (represented by the remaining samples divided by the throughput) to calculate the weight $WG(A^j)$ as follows:

$$WG(A^j) = \frac{1}{(\Phi_{sp}^j / \Psi_{thp}^j + \epsilon)^\rho} \quad (14)$$

Here, Φ_{sp}^j represents the number of the remaining samples to be trained in the job j . As $\rho \rightarrow 0$, $WG(A^j)$ smoothly approaches 1 for every job j , which means we consider the weights of all jobs

to be equal. When $\rho \rightarrow \infty$, $WG(A^j)$ can prioritize jobs with a shorter completion time. In contrast, as $\rho \rightarrow -\infty$, $WG(A^j)$ prioritize jobs with longer completion time. A cluster operator may select a suitable value for ρ , based on practical priorities. At ANTRGROUP, we choose $\rho = 2.5$ as a reasonable value to complete shorter jobs quicker and release the resources. Additionally, ϵ denotes a very small value used to prevent division by zero.

3 Post-scaling Stage: Instability Handling.

During the auto-scaling phase (stage 2), we assume that the job is interference-free. Yet, in practical cloud environments, this assumption might not hold as discussed in §2.2. Therefore, DLROVER involves many techniques to handle various cloud instabilities in the post-scaling stage (stage 3). We highlight some key instability issues and the techniques adopted by DLROVER to address them in the following (with more details in §5).

Worker Stragglers. In heterogeneous clusters, certain worker pods may be assigned to physical machines with slow hardware (i.e., low-frequency CPU and/or low-speed memory) or be hindered by high-priority pods due to resource contention. These worker stragglers can result in submitting stale model gradients to the parameter servers, leading to decreased model accuracy and longer training time [16]. To address this, DLROVER implements a *dynamic data sharding* mechanism to minimize the discrepancy in iteration rounds among workers (see §5.1).

PS Stragglers. Our DLRM jobs run on the TensorFlow framework, which determines parameter allocation based on *tensors* or multi-dimensional arrays. The size of tensor-based parameters assigned to PSeS can differ substantially, resulting in unbalanced workloads [7]. Consequently, PSeS performing large matrix multiplications with more allocated parameters experience significantly higher CPU loads, resulting in the PS stragglers. To address this, we adopt DeepRec [6] to ensure that the embedding parameters are evenly distributed across the new set of PS nodes.

Scaling Overhead. As formulated in Eqn. 8, scaling overhead plays an important role in workload migration. To speed up the migration process, DLROVER introduces a *seamless migration* mechanism to mitigate the scaling overhead (see §5.2).

Out of Memory Problem. Uneven allocation and high memory consumption from large embedding tables can lead to out-of-memory problems in PSeS. To address this, DLROVER invents a *OOM prediction* mechanism (see §5.3).

5 HANDLING INSTABILITY

DLROVER’s design draws upon the collective insights and practices from the realm of distributed/training systems. In this section, we focus on the key mechanisms developed by DLROVER to enhance the performance and reliability of the DLRM training system.

5.1 Dynamic Data Sharding

DLROVER introduces a *dynamic data sharding* mechanism that enables fine-grained data serving by partitioning training data into *numerous small shards of various sizes*. These data shards can be on-demand, dynamically assigned/reassigned to 1) slow workers (i.e., stragglers) to balance their paces of data processing and model updates with their peers for consistent model quality; and 2) new/healthy workers for fast elasticity or fault tolerance.

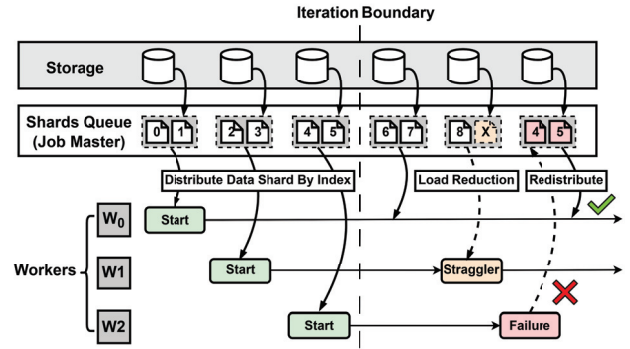


Figure 5: Dynamic data sharding.

Data Sharding of DLROVER. Traditionally, training data are (statically) partitioned and distributed among workers at the beginning of job training. In contrast, DLROVER splits the dataset into *numerous, much smaller, and variably-sized shards* (e.g., 64, 128, or 256 data batches), each labeled with a unique index; DLROVER then manages the data serving by delivering such fine-grained data shards to the corresponding worker *on-demand* during a worker’s life cycle.

Specifically, as illustrated in Fig. 5, data shards are organized within a *shards queue*. Upon initiating job training, a worker fetches the required data shard (by its index) from this queue. Note that workers are initially assigned comparable workloads, as determined by the quantity of data shards and the number of data samples within them over a fixed time interval. Throughout the training process, workers dispatch heartbeat packets at regular intervals to the job master (Fig. 4), which include the number of data samples they have processed, named *progress offset*. These heartbeat packets and the progress offsets serve three critical functions: 1) They signal that the worker is operational and active. Conversely, this mechanism is employed to identify workers that have failed. 2) The job master uses the *progress offset* to gauge the worker’s training progress and determine if it is falling behind – i.e., identifying any straggler if the offset is noticeably lesser than its peers. 3) Upon completion of the designated data shard, the worker sends its final *progress offset* to the job master, indicating the completion of that shard’s processing. Subsequently, the worker acquires a new shard from the shards queue to proceed with the training.

Failure Recovery. Once the job master does not receive the heartbeat package from a worker for a reasonably long time, it is identified as a failure event. In the event of a worker failure (highlighted in red in Fig. 5), the job master re-joins the unfinished data shard(s) of the failed worker to the *shards queue*, awaiting redistribution to another healthy worker. This mechanism simply ensures that the training job consumes the training data *without* any data omission or duplication, guaranteeing the model’s consistent quality.

Handling Stragglers. Throughout job training, the job master also keeps track of the *progress offsets* each worker provides. A worker is labeled a straggler if it lags significantly behind its peers (highlighted in orange in Fig. 5). In such cases, the system mitigates the issue by reassigning the slower worker a smaller workload, such as providing a data shard with fewer batches (e.g., scaling

down from a shard with 256 batches to one with 128). This way, the system can dynamically tailor the volume of data samples a slower worker has to process before it submits its gradients to the PSEs. Consequently, it enables the straggling worker to align its gradient submission rate with others, preventing the submission of stale gradients and maintaining consistent model quality.

Fast Elasticity. Partitioning data into smaller data shards and adopting the *shards queue* ultimately gives rise to the benefit of enabling fast elasticity because any new worker – i.e., after a stop-and-restart with adjusted CPU/memory resources – can simply retrieve a data shard from the *shards queue* without the process of data re-partitioning and redistribution among all workers.

5.2 Seamless Migration

DLROVER develops a *seamless migration* mechanism to minimize resource scaling overhead during training *by strategically overlapping the scaling progress with ongoing training activities*. This approach effectively reduces training downtime associated with worker/PS initialization and delays in allocating new resources. To further reduce the scaling overhead, DLROVER invents a *flash-checkpoint* mechanism, which accelerates checkpointing through the use of shared memory and asynchronous data persistence.

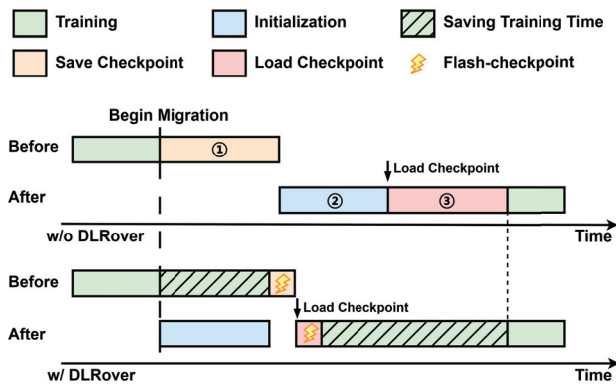


Figure 6: Seamless Migration.

Seamless Migration. Scaling resources up/down conventionally involves a *stop-and-restart* operation: As shown in Fig. 6, the migration progress (w/o DLROVER) ① stops the workers/PSEs and checkpoints the model parameters in (remote) persistent storage (e.g., the RDS at ANTGROUP); ② deploys and initializes new PSEs and workers based on the updated resource plans – including adjusting the number of worker/PS pods and/or their allocated resources, pulling images of worker/PS pods, and launching the new worker/PS pods; and ③ finally loads the checkpoints from the persistent storage and resumes the training job. While being straightforward, the *stop-and-restart* operation results in nontrivial pending time (e.g., up to tens of minutes), especially in the cloud environment (§2.2).

We observe that such a synchronous *stop-and-restart* operation can be decoupled into two parallel ones. Initially, rather than beginning with step ①, DLROVER starts initializing and deploying new workers and PSEs (i.e., ②) while allowing the "old" workers/PSEs

to proceed with the ongoing training job. Once all the new workers/PSEs are ready to use, DLROVER performs steps ① and ③, during which training jobs must be paused. To further speed up the critical path (① and ③), DLROVER invents the *flash-checkpoint* approach.

Flash-checkpoint. In the production environment of ANTGROUP, the RDS services are commonly shared among various internal services, each allocated a limited bandwidth. This, unfortunately, prolongs the critical path of DLROVER’s migration (① and ③). Unlike checkpointing for fault tolerance, which focuses on data persistence, checkpointing during the migration (of workers and PSEs in DLROVER) should emphasize speed over strong persistence. This insight motivates DLROVER to leverage a distributed caching service (at AntGroup [5]) to enhance migration efficiency, called *flash-checkpoint*. During migration, DLROVER checkpoints (①) or loads (③) model parameters via the caching service. As the bandwidth and access speed of the caching service is much faster than the RDS, the checkpointing process takes much less time (e.g., less than 1 second for a 20GB model). Further, the caching service facilitates data sharing between new and old workers/PSEs when they are located on the same physical node, eliminating network transmission.

Note that DLROVER has a separate RDS-based checkpointing mechanism for job-level and system-level fault tolerance (i.e., to recover a job or the system from failures). The *flash-checkpoint* mechanism, though separately, enhances the RDS-based checkpointing by flushing model parameters from the caching system to the RDS asynchronously (i.e., more update-to-date checkpoints).

5.3 OOM Prevention

DLROVER employs a prediction mechanism for preventing out-of-memory (OOM) via modeling the dynamics of memory usage. After modeling the memory usage, DLROVER can use it to predict memory usage and check if PSEs would exceed the memory capacity before the job completion (e.g., 10,000 steps to complete). If an OOM error is estimated to occur within the job completion step, DLROVER scales the PSEs with larger memory capacity. We omit the thorough analysis due to space limitations, please refer to our technical report [62] for more details.

6 EVALUATION

Evaluation Environments. We thoroughly evaluated DLROVER using both a small-scale (and more controlled) cluster and production environment at ANTGROUP. 1) The small-scale cloud-based cluster has 20 CPU servers, each equipped with two 16-core Intel Xeon E5-2682 @2.5GHz CPU and 192GB RAM. All experiments were conducted on a Kubernetes-managed cluster. 2) We further evaluated DLROVER by deploying it in our production environment. The cloud-based cluster provides ~1.6 million CPU cores, 3.24 PB of memory, and 344 PB of disk storage. Different types of jobs (e.g., training, serving, and stream processing) are sharing these cloud-based resources.

Workloads. We employed three representative DLRM models at ANTGROUP [11, 54, 74, 74]: 1) **Model-X:** Wide & Deep [13]; 2) **Model-Y:** xDeepFM[33], and 3) **Model-Z:** DCN[63]. We measured the performance of DLROVER on Criteo dataset [1]. All models were implemented in TensorFlow 1.13, and each job had the same batch size of 512 with 200,000 training steps.

Comparison Baselines. The evaluation covers two primary baselines: 1) *w/o DLROVER*: the general distributed framework (e.g., Kubeflow) used in the cloud [2] without DLROVER support. In this baseline, each job’s resources needed to be manually configured. To show the effectiveness of DLROVER, we first well-tuned the resource configuration of this baseline and compared it with the support of DLROVER (i.e., *w/ DLROVER*). 2) The state-of-the-art deep learning training job schedulers for *CPU-only* scenarios (i.e., as discussed in §2, DLROVER is designed for CPU-only hardware), including Elastic Scheduler (short as **ES**) [43] and **Optimus** [45]. These schedulers are well-designed for traditional deep learning models in NLP and CV. Note that, both **ES** and **Optimus** add or remove a fixed number of nodes each time, while **ES** only modulates workers and **Optimus** adjusts PS or workers.

Metrics. We focus on four kinds of metrics. 1) *Job Completion Time (JCT)*: the end-to-end model training time (the shorter is better); 2) *Job Completion Rate (JCR)*: the proportion of successfully completed jobs to the total jobs submitted within a defined timeframe. A higher JCR indicates efficient job processing and a lower fault ratio; 3) *CPU Utilization Rate (CUR)*: the workload processed by the CPU within a specified time period; and 4) *Memory Utilization Rate (MUR)*: the volume of system memory consumed during operations.

6.1 End-to-End System Performance

In this section, we demonstrate the end-to-end performance of DLROVER within the small-scale cluster, which eliminates the impact of cloud instability and ensures a more fair comparison.

DLROVER Nears Well-Tuned Configurations. To show the effectiveness of DLROVER, we manually tuned the resource configurations for jobs without the support of DLROVER until they almost reached the best throughput. Fig. 7 shows that DLROVER yields comparable JCTs to well-tuned settings. For instance, the JCT for model-X with DLROVER is 27.74 minutes – 1.4% higher than the well-tuned counterpart. Note that manual tuning is a time-consuming *trial-and-error* approach. For example, for Model-X, we re-run the job for more than 10 times to achieve its optimal configuration. The results demonstrate that DLROVER’s three-stage algorithm (§4.3) can accurately capture/allocate the resource demands for different types of DLRM training jobs.

DLROVER Beats Traditional Schedulers. Fig. 7 depicts that jobs under DLROVER consistently achieve shorter JCT than ES and Optimus. On average, there is a 17.7% or 28.5% improvement compared with ES or Optimus. The results show that by considering the unique lookup operations and memory demands in DLRMs, DLROVER can effectively schedule DLRM training jobs to improve the training efficiency with shorter JCT.

DLROVER Preserves Model Convergence. In this experiment, we used 90% of the Criteo dataset as the training set and the remaining as the test set. As shown in Fig. 8, DLROVER does *not* compromise the model performance (i.e., convergence and accuracy) compared to jobs whose resource configurations are well-tuned. This verifies that the proposed *dynamic data sharding* mechanism prevents inconsistent model quality caused by elastic operations (§5.1).

In summary, DLROVER enhances the training efficiency of various DLRM models by reducing JCT while maintaining model convergence. In the following section, we conduct an in-depth analysis of the separate effectiveness of the main components in DLROVER.




Figure 7: DLROVER achieves comparable JCT to well-tuned configurations and reduces JCT compared with competitors.




Figure 8: DLROVER preserves the model performance (i.e., accuracy and loss) on different DLRM training jobs.

6.2 Ablation Study

Warm-starting. To demonstrate the effectiveness of the *warm-starting* algorithm (algorithm 1), we collected one month’s job training data from a user within the production cluster. Fig. 9 shows that DLROVER, with the *warm-starting* algorithm, provided initial resource allocation very close to the final configuration. On average, the accuracy of DLROVER’s initial and final configurations for workers and PSeS are respectively 92% and 85%. It is because, based on similarity analysis, DLROVER’s *warm-starting* algorithm extracts the most matching jobs from the users’ historical task data to serve as guidance; With a good initial resource configuration, DLROVER reduces the number of scaling. Thus, we reduce the time of scaling jobs from initial resource allocation to optimal resource allocation. Based on the statistics from the production cluster logs, compared to a cold-start approach (adjusting resources from zero), the scaling time was reduced by an average of 26%.

Auto-Scaling. To evaluate the effectiveness of auto-scaling strategies, we trained DLRMs from scratch (i.e., cold-started, thus removing the effect of DLROVER’s *warm-starting*) with different schedulers. Every three minutes, schedulers adjusted the resources of PSeS or workers based on the runtime information. Fig. 10 shows that, compared to ES and Optimus, DLROVER can achieve higher throughput within the same time period. For instance, for model-X,

DLROVER achieves a throughput of 250 steps/second after running for 12 minutes, while others still stay at the throughput of 100-150 steps/second. This is because DLROVER considers unique lookup operations of DLRM training (in Eqn. 5).




Figure 9: DLROVER’s warm-starting strategy provides a resource allocation close to the final configuration.




Figure 10: DLROVER’s scheduling algorithm achieves higher throughput in less time.

To further verify this, we validate the proposed throughput prediction model (Eqn. 1). We sampled a set of training data points under different resource setups (different (p, w, C_w, C_p)). Then, we use NNLS [4] to find α s and β s that best fit the collected data points. As shown in Fig. 11, our model can closely describe the relationship between the training throughput and resource configurations. We report coefficients in Eqn. 1 as: $\alpha_{grad} = 3.48$, $\alpha_{upd} = 2.36$, $\alpha_{lookup} = 2.45$, $\alpha_{sync} = 0.68$, and 2.45 for the sum of β .

Instability Handling. In this study, we verify the capability of DLROVER to handle instability in the cloud (e.g., straggler and hot PSes). To simulate the worker/PS straggler cases, we randomly selected a worker or PS and set the CPU cores to 3% of that in the well-tuned resource configuration.

As shown in Fig. 12, for the hot PS case, we observe that: 1) DLROVER can reduce the JCT by 36.4% and 27.6% compared to the "no intervention" and "traditional migration" methods, respectively. 2) Unlike "traditional migration", when stragglers are detected, DLROVER enables continuous training (i.e., *seamless migration*), as DLROVER restarts the job using an asynchronous approach that does not interrupt the job’s training (§5.2) – saving about 5 minutes of the training time. 3) With the *flash-checkpoint* mechanism, DLROVER saves 3 minutes in saving and loading checkpoints due to DLROVER using shared memory to store checkpoints instead of

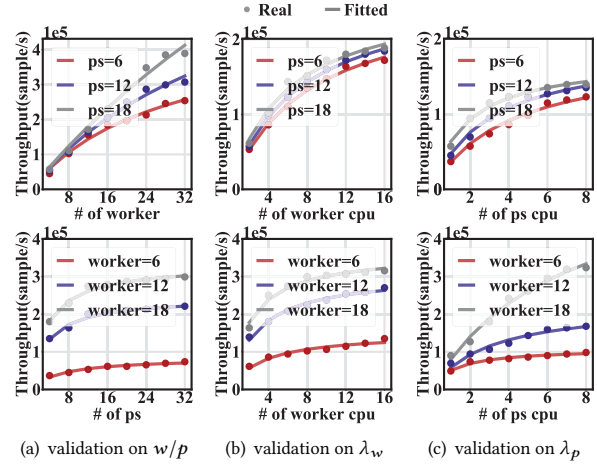


Figure 11: Sampled data points and the fitted curves of the throughput prediction model. Under the setups with varying numbers of workers and PS, DLROVER accurately predicts the throughput when adjusting resource variables.

communicating with RDS – significantly reducing the communication overhead (§5.2). As shown in Fig. 13, for the worker straggler case, we observe that 1) DLROVER can shorten the JCT by 48.5% and 37% compared to the "no intervention" and "traditional handling" approaches, respectively. 2) Compared to "traditional handling", DLROVER can also rapidly handle the straggler (within 1 minute) and recover the healthy training instead of restarting the job. This is achieved by redistributing less data shard(s) to the straggler pod, with the support of dynamic data sharding mechanism (§5).




Figure 12: When a hot PS occurs, three distinct strategies result in different JCT: The "no intervention" approach continues training under an unhealthy state. The "traditional migration" method employs a *stop-and-restart* strategy, while DLROVER utilizes the mechanism of *seamless migration* and *flash-checkpoint*, significantly reducing overhead (§5.2).

6.3 Production Adoption and Evaluation

DLROVER has been deployed in the cloud-based production cluster at ANTGROUP since June 2022. We conduct a comprehensive

production evaluation to reveal the benefits of production adoption of DLROVER at ANTGROUP. We omit this study due to space limitations, please refer to our technical report [62] for more details.

7 RELATED WORK

Training of Deep Learning Recommendation. Existing DLRM training systems focus on accelerating training speed and addressing memory pressure. For example, AIBox [77] and HierPS [76] overlap training execution on CPUs and GPUs, while using SSDs to store massive parameters of the model. Ekko [56] accelerates DLRM training over wide-area networks. Adnan et al. [9] place highly accessed embeddings on GPU memory to reduce communication time. Gupta et al. [21] compresses parameter gradients during model synchronization and squeezes activations and gradients across subnetworks during the forward and backward propagations. TTRec [70] adopts tensor train decomposition to mitigate memory consumption. AdaEmbed [27] identifies the embedding rows with larger importance to improve model accuracy. AutoShard [72], DreamShard [73] and [14] seek the optimal embedding table sharding strategy to mitigate the lookup imbalances across devices. In contrast, the data sharding mechanism (§5.1) in DLROVER focuses on training data serving. FSDP[44] provides an industry-grade solution for large model data parallel training. Borg[60], Twine[57], and especially Akkio[10] apply similar sharding implementations to achieve good latency on data serving by leveraging data locality. In contrast, DLROVER focuses on scheduling resources and provides robust fault tolerance for model training in a cloud environment, where failures are common and resource availability varies.

Automatic Resource Configuration. Automatic resource management is widely used in distributed data processing jobs on Hadoop [55] and Spark [71], and machine learning jobs on Spark MLlib [37]. For example, Huang et al. [24] introduce an approach to configure the memory of large-scale ML on Spark automatically. Angel-PTM [42] focuses on memory management optimization for large-scale transformer models. Eigen and fastflow [30, 59] explores resource efficiency optimization in large-scale public-cloud production environments. Additionally, numerous endeavors have been dedicated to the automated configuration of the number of workers and parameter servers, such as Optimus, Pollux, and Tiresias [17, 45, 49]. Pollux [49] dynamically adjusts the number of workers and learning rate to improve the throughput for synchronous SGD. Both Pollux and Tiresias must re-deploy all workers when adjusting resources, resulting in a long transition time. To minimize it, [43] starts new all-reduce operations only when new workers are ready and proposes a heuristic scaling to search the optimal number of workers. For asynchronous training with the parameter server architecture, Optimus [45] dynamically adds one worker or parameter server each time to maximize the cluster’s performance without considering the transition time of elasticity. Regarding a large-scaling training job, the transition cost is not trivial because the number of workers and parameters is huge. DLROVER is designed to conduct elasticity in a more effective way with little overhead. Furthermore, we plan to identify better configurations for training jobs using LLM-based optimization techniques[28].

Elastic Deep Learning Training. There are deep learning frameworks that support elastic training. For instance, for asynchronous

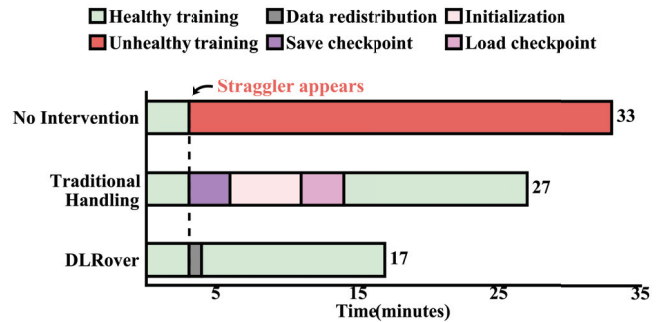


Figure 13: When a straggler occurs, three distinct strategies result in different JCT: The No Intervention approach persists in training under an unhealthy state. The Traditional Handling method employs a ‘stop-and-restart’ strategy, while DLROver redistributes the data shard for the straggler with the support of *dynamic data sharding* mechanism (§5.1), avoiding the restart of the job.

training, the PS training of TensorFlow [7] supports scaling workers at runtime. Using checkpoint, TensorFlow enables the elasticity for parameter servers. For synchronous training, there are Elastic DL [67], PyTorch-Elastic [47], and elastic Horovod [53]. These systems typically restart the job by relaunching all pods [45, 49]. Users need to implement their dynamic data partition policy. The re-partitioning may result in inconsistency of sample iterations if the dataset is huge. In contrast, DLROVER has a *dynamic data sharding* service and does not need to re-partition during elasticity. DLROVER also ensures consistency when adjusting parameter servers by checkpointing unused data shards and model parameters leveraging memory storage like Gemini[64].

Straggler Mitigation. For a distributed asynchronous DLRM training job using parameter servers, the straggler could be a PS or worker because of hardware heterogeneity [50] or unbalanced data/parameter distribution [7]. Existing works simply replace the slowest node with a new node to mitigate stragglers [22, 43, 45]. However, this introduces additional overhead. Existing frameworks like Kubeflow [2] can only set the same CPU and memory for the workers or PSes. In contrast, DLROVER can adjust the workload and resources for various types of components on the fly based on updated elasticity decisions.

8 CONCLUSION

We have presented DLROVER, a cloud-based DLRM training system. In designing DLROVER, we have considered the unique characteristics of DLRMs and the practical challenges in a cloud environment. DLROVER builds an accurate resource-performance model incorporating various runtime training information and develops a three-stage scheduling algorithm for elastic resource allocation and adjustment for DLRM training jobs. Moreover, DLROVER offers a bunch of novel mechanisms to handle high cloud instability. Our evaluation demonstrates the effectiveness of DLROVER in reduced job completion time and increased resource utilization.

REFERENCES

- [1] 2014. Criteo. <http://labs.criteo.com/downloads/2014-kaggle-display-advertising-challenge-dataset/>
- [2] 2023. Kubeflow: The Machine Learning Toolkit for Kubernetes. <https://www.kubeflow.org/>
- [3] 2023. Pymoo NSGA-II. <https://pymoo.org/algorithms/moo/nsga2.html>
- [4] 2023. SciPy NNLS. <https://docs.scipy.org/doc/scipy/reference/generated/scipy.optimize.nnls.html>
- [5] 2024. Accelerate Enterprise AI Path to Production. <https://www.alluxio.io/>
- [6] 2024. DeepRec. <https://github.com/DeepRec-AI/DeepRec>
- [7] Martín Abadi and Ashish Agarwal et al. 2015. TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. <https://www.tensorflow.org/> Software available from tensorflow.org.
- [8] Bilge Acun, Matthew Murphy, Xiaodong Wang, Jade Nie, Carole-Jean Wu, and Kim Hazelwood. 2021. Understanding training efficiency of deep learning recommendation models at scale. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 802–814.
- [9] Muhammad Adnan, Yassaman Ebrahimzadeh Maboud, Divya Mahajan, and Prashant J Nair. 2021. Accelerating recommendation system training by leveraging popular choices. *arXiv preprint arXiv:2103.00686* (2021).
- [10] Muthukaruppan Annamalai, Kaushik Ravichandran, Harish Srinivas, Igor Zinkovskiy, Luning Pan, Tony Savor, David Nagle, and Michael Stumm. 2018. Sharding the Shards: Managing Datastore Locality at Scale with Akkio. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*. USENIX Association, Carlsbad, CA, 445–460. <https://www.usenix.org/conference/osdi18/presentation/annamalai>
- [11] Chaochao Chen, Jun Zhou, Bingzhe Wu, Wenjing Fang, Li Wang, Yuan Qi, and Xiaolin Zheng. 2020. Practical privacy preserving POI recommendation. *ACM Transactions on Intelligent Systems and Technology (TIST)* 11, 5 (2020), 1–20.
- [12] Jianmin Chen, Xinghao Pan, Rajat Monga, Samy Bengio, and Rafal Jozefowicz. 2017. Revisiting Distributed Synchronous SGD. *arXiv:1604.00981* [cs.LG]
- [13] Heng-Tze Cheng, Levent Koc, Jeremiah Harmsen, Tal Shaked, Tushar Chandra, Hrishikesh Aradhya, Glen Anderson, Greg Corrado, Wei Chai, Mustafa Ispir, Rohan Anil, Zakaria Haque, Lichan Hong, Vihang Jain, Xiaobing Liu, and Hemal Shah. 2016. Wide & Deep Learning for Recommender Systems. In *Proceedings of the 1st Workshop on Deep Learning for Recommender Systems, DLRS@RecSys 2016, Boston, MA, USA, September 15, 2016*. Alexandros Karatzoglou, Balázs Hidasi, Domonkos Tikk, Oren Sar Shalom, Haggai Roitman, Bracha Shapira, and Lior Rokach (Eds.). ACM, 7–10. <https://doi.org/10.1145/2988450.2988454>
- [14] Runxiang Cheng, Chris Cai, Selman Yilmaz, Rahul Mitra, Malay Bag, Mrinmoy Ghosh, and Tianyin Xu. 2023. Towards GPU Memory Efficiency for Distributed Training at Scale. In *Proceedings of the 2023 ACM Symposium on Cloud Computing (SoCC '23)*. Association for Computing Machinery, New York, NY, USA, 281â€297. <https://doi.org/10.1145/3620678.3624661>
- [15] Michael Chui, James Manyika, Mehdi Miremadi, Nicolaus Henke, Rita Chung, Pieter Nel, and Sankalp Malhotra. 2018. Notes from the AI frontier: Insights from hundreds of use cases. *McKinsey Global Institute 2* (2018).
- [16] Wei Dai, Yi Zhou, Nanqing Dong, Hao Zhang, and Eric P Xing. 2018. Toward understanding the impact of staleness in distributed machine learning. *arXiv preprint arXiv:1810.03264* (2018).
- [17] Juncheng Gu, Mosharaf Chowdhury, Kang G. Shin, Yibo Zhu, Myeongjae Jeon, Junjie Qian, Hongqiang Liu, and Chuanxiong Guo. 2019. Tiresias: A GPU Cluster Manager for Distributed Deep Learning. In *Proceedings of the 16th USENIX Conference on Networked Systems Design and Implementation* (Boston, MA, USA) (NSDI'19). USENIX Association, USA, 485â€500.
- [18] Huifeng Guo, Ruiming Tang, Yunming Ye, Zhenguo Li, and Xiuqiang He. 2017. DeepFM: A Factorization-Machine based Neural Network for CTR Prediction. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, Carles Sierra (Ed.). ijcai.org, 1725–1731. <https://doi.org/10.24963/ijcai.2017/239>
- [19] Jing Guo, Zihao Chang, Sa Wang, Haiyang Ding, Yihui Feng, Liang Mao, and Yungang Bao. 2019. Who limits the resource efficiency of my datacenter: An analysis of alibaba datacenter traces. In *Proceedings of the International Symposium on Quality of Service*. 1–10.
- [20] Udit Gupta, Carole-Jean Wu, Xiaodong Wang, Maxim Naumov, Brandon Reagen, David Brooks, Bradford Cottle, Kim Hazelwood, Mark Hempstead, Bill Jia, et al. 2020. The architectural implications of facebook’s dnn-based personalized recommendation. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 488–501.
- [21] Vipul Gupta, Dhruv Choudhary, Peter Tang, Xiaohan Wei, Xing Wang, Yuzhen Huang, Arun Kejariwal, Kannan Ramchandran, and Michael W Mahoney. 2021. Training recommender systems at scale: Communication-efficient model and data parallelism. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 2928–2936.
- [22] Aaron Harlap, Henggang Cui, Wei Dai, Jinliang Wei, Gregory R. Ganger, Phillip B. Gibbons, Garth A. Gibson, and Eric P. Xing. 2016. Addressing the straggler problem for iterative convergent parallel ML. In *Proceedings of the Seventh ACM Symposium on Cloud Computing, Santa Clara, CA, USA, October 5-7, 2016*, Marcos K. Aguilera, Brian Cooper, and Yanlei Diao (Eds.). ACM, 98–111. <https://doi.org/10.1145/2987550.2987554>
- [23] Zhaoxin Huan, Ke Ding, Ang Li, Xiaolu Zhang, Xu Min, Yong He, Liang Zhang, Jun Zhou, Linjian Mo, Jinjie Gu, et al. 2023. AntM²C: A Large Scale Dataset For Multi-Scenario Multi-Modal CTR Prediction. *arXiv preprint arXiv:2308.16437* (2023).
- [24] Botong Huang, Matthias Boehm, Yuanyuan Tian, Berthold Reinwald, Shirish Tatikonda, and Frederick R. Reiss. 2015. Resource Elasticity for Large-Scale Machine Learning. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data, Melbourne, Victoria, Australia, May 31 - June 4, 2015*, Timos K. Sellis, Susan B. Davidson, and Zachary G. Ives (Eds.). ACM, 137–152. <https://doi.org/10.1145/2723372.2749432>
- [25] Yuzhen Huang, Xiaohan Wei, Xing Wang, Jiyan Yang, Bor-Yiing Su, Shivam Bharuka, Dhruv Choudhary, Zewei Jiang, Hai Zheng, and Jack Langman. 2021. Hierarchical training: Scaling deep recommendation models on large cpu clusters. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*. 3050–3058.
- [26] Shristi Shakya Khanal, PWC Prasad, Abeer Alsadoon, and Angelika Maag. 2020. A systematic review: machine learning based recommendation systems for e-learning. *Education and Information Technologies* 25 (2020), 2635–2664.
- [27] Fan Lai, Wei Zhang, Rui Liu, William Tsai, Xiaohan Wei, Yuxi Hu, Sabin Devkota, Jianyu Huang, Jongsoo Park, Xing Liu, et al. 2023. {AdaEmbed}: Adaptive Embedding for {Large-Scale} Recommendation Models. In *17th USENIX Symposium on Operating Systems Design and Implementation (OSDI 23)*. 817–831.
- [28] Jiale Lao, Yibo Wang, Yufei Li, Jianping Wang, Yunjia Zhang, Zhiyuan Cheng, Wanghu Chen, Mingjie Tang, and Jianguo Wang. 2024. GPTuner: A Manual-Reading Database Tuning System via GPT-Guided Bayesian Optimization. *Proc. VLDB Endow.* 17, 8 (may 2024), 1939â€1952. <https://doi.org/10.14778/3659437.3659449>
- [29] Jiamin Li, Hong Xu, Yibo Zhu, Zherui Liu, Chuanxiong Guo, and Cong Wang. 2023. Lyra: Elastic Scheduling for Deep Learning Clusters. In *Proceedings of the Eighteenth European Conference on Computer Systems (Rome, Italy) (EuroSys '23)*. Association for Computing Machinery, New York, NY, USA, 835â€850. <https://doi.org/10.1145/3552326.3587445>
- [30] Ji You Li, Jiachi Zhang, Wenchao Zhou, Yuhang Liu, Shuai Zhang, Zhuoming Xue, Ding Xu, Hua Fan, Fangyuan Zhou, and Feifei Li. 2023. Eigen: End-to-end Resource Optimization for Large-Scale Databases on the Cloud. *Proceedings of the VLDB Endowment* 16, 12 (2023), 3795–3807.
- [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 *11th USENIX Symposium on Operating Systems Design and Implementation, OSDI '14, Broomfield, CO, USA, October 6-8, 2014*, Jason Flinn and Hank Levy (Eds.). USENIX Association, 583–598. https://www.usenix.org/conference/osdi14/technical-sessions/presentation/li_mu
- [32] Mingzhen Li, Wencong Xiao, Hailong Yang, Biao Sun, Hanyu Zhao, Shiru Ren, Zhongzhi Luan, Xianyan Jia, Yi Liu, Yong Li, Wei Lin, and Depei Qian. 2023. EasyScale: Elastic Training with Consistent Accuracy and Improved Utilization on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (<conf-loc>, <city>Denver</city>, <state>CO</state>, <country>USA</country>, </conf-loc>) (SC '23)*. Association for Computing Machinery, New York, NY, USA, Article 55, 14 pages. <https://doi.org/10.1145/3581784.3607054>
- [33] Jianxun Lian, Xiaohuan Zhou, Fuzheng Zhang, Zhongxia Chen, Xing Xie, and Guangzhong Sun. 2018. xDeepFM: Combining Explicit and Implicit Feature Interactions for Recommender Systems. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, Yike Guo and Faisal Farooq (Eds.). ACM, 1754–1763. <https://doi.org/10.1145/3219819.3220023>
- [34] Qixiao Liu and Zhibin Yu. 2018. The elasticity and plasticity in semi-containerized co-locating cloud workload: a view from alibaba trace. In *Proceedings of the ACM Symposium on Cloud Computing*. 347–360.
- [35] Michael Lui, Yavuz Yetim, Özgür Özkan, Zhuoran Zhao, Shin-Yeh Tsai, Carole-Jean Wu, and Mark Hempstead. 2021. Understanding capacity-driven scale-out neural recommendation inference. In *2021 IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*. IEEE, 162–171.
- [36] JP Mangalindan. 2012. Amazonâ€™s recommendation secret. *CNN Money* <http://tech.fortune.cnn.com/2012/07/30/amazon-5> (2012).
- [37] Xiangrui Meng, Joseph K. Bradley, Burak Yavuz, Evan R. Sparks, Shivaram Venkataraman, Davies Liu, Jeremy Freeman, D. B. Tsai, Manish Amde, Sean Owen, Doris Xin, Reynold Xin, Michael J. Franklin, Reza Zadeh, Matei Zaharia, and Ameet Talwalkar. 2016. MLlib: Machine Learning in Apache Spark. *J. Mach. Learn. Res.* 17 (2016), 34:1–34:7. <http://jmlr.org/papers/v17/15-237.html>

- [38] Dheevatsa Mudigere, Yuchen Hao, Jianyu Huang, Zhihao Jia, Andrew Tulloch, Srinivas Sridharan, Xing Liu, Mustafa Ozdal, Jade Nie, Jongsoo Park, et al. 2022. Software-hardware co-design for fast and scalable training of deep learning recommendation models. In *Proceedings of the 49th Annual International Symposium on Computer Architecture*. 993–1011.
- [39] Mukosi Abraham Mukweho and Turgay Celik. 2018. Toward a smart cloud: A review of fault-tolerance methods in cloud systems. *IEEE Transactions on Services Computing* 14, 2 (2018), 589–605.
- [40] Maxim Naumov, John Kim, Dheevatsa Mudigere, Srinivas Sridharan, Xiaodong Wang, Whitney Zhao, Serhat Yilmaz, Changkyu Kim, Hector Yuen, Mustafa Ozdal, et al. 2020. Deep learning training in facebook data centers: Design of scale-up and scale-out systems. *arXiv preprint arXiv:2003.09518* (2020).
- [41] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G Azzolini, et al. 2019. Deep learning recommendation model for personalization and recommendation systems. *arXiv preprint arXiv:1906.00091* (2019).
- [42] Xiaonan Nie, Yi Liu, Fangcheng Fu, Jinbao Xue, Dian Jiao, Xupeng Miao, Yangyu Tao, and Bin Cui. 2023. Angel-PTM: A Scalable and Economical Large-scale Pre-training System in Tencent. arXiv:2303.02868 [cs.LG]
- [43] Andrew Or, Haoyu Zhang, and Michael J. Freedman. 2020. Resource Elasticity in Distributed Deep Learning. In *Proceedings of Machine Learning and Systems 2020, MLSys 2020, Austin, TX, USA, March 2–4, 2020*, Inderjit S. Dhillon, Dimitris S. Papaliopoulos, and Vivienne Sze (Eds.). mlsys.org. <https://proceedings.mlsys.org/book/314.pdf>
- [44] Adam Paszke and Gross et al. 2019. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*, H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Eds.). Curran Associates, Inc., 8024–8035.
- [45] Yanghua Peng, Yixin Bao, Yangrui Chen, Chuan Wu, and Chuanxiong Guo. 2018. Optimus: An Efficient Dynamic Resource Scheduler for Deep Learning Clusters. In *Proceedings of the Thirteenth EuroSys Conference (Porto, Portugal) (EuroSys '18)*. Association for Computing Machinery, New York, NY, USA, Article 3, 14 pages. <https://doi.org/10.1145/3190508.3190517>
- [46] Ivens Portugal, Paulo Alencar, and Donald Cowan. 2018. The use of machine learning algorithms in recommender systems: A systematic review. *Expert Systems with Applications* 97 (2018), 205–227.
- [47] PyTorch. 2020. Pytorch with Elastic. <https://pytorch.org/elastic/0.1.0rc2/overview.html>.
- [48] Ling Qian, Zhiguo Luo, Yujian Du, and Leitao Guo. 2009. Cloud computing: An overview. In *Cloud Computing: First International Conference, CloudCom 2009, Beijing, China, December 1–4, 2009. Proceedings 1*. Springer, 626–631.
- [49] Aurick Qiao, Sang Keun Choe, Suhas Jayaram Subramanya, Willie Neiswanger, Qirong Ho, Hao Zhang, Gregory R. Ganger, and Eric P. Xing. 2021. Pollux: Co-adaptive Cluster Scheduling for Goodput-Optimized Deep Learning. In *15th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2021, July 14–16, 2021*, Angela Demke Brown and Jay R. Lorch (Eds.). USENIX Association. <https://www.usenix.org/conference/osdi21/presentation/qiao>
- [50] Charles Reiss, Alexey Tumanov, Gregory R. Ganger, Randy H. Katz, and Michael A. Kozuch. 2012. Heterogeneity and dynamics of clouds at scale: Google trace analysis. In *ACM Symposium on Cloud Computing, SOCC '12, San Jose, CA, USA, October 14–17, 2012*, Michael J. Carey and Steven Hand (Eds.). ACM, 7. <https://doi.org/10.1145/2391229.2391236>
- [51] Pradeep Kumar Roy, Sarabjeet Singh Chowdhary, and Rocky Bhatia. 2020. A Machine Learning approach for automation of Resume Recommendation system. *Procedia Computer Science* 167 (2020), 2318–2327.
- [52] Bo Sang, Shuwei Gu, Xiaojun Zhan, Mingjie Tang, Jian Liu, Xuan Chen, Jie Tan, Haoyuan Ge, Ke Zhang, Ruoyi Ruan, et al. 2023. Cougar: A General Framework for Jobs Optimization In Cloud. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*. IEEE, 3417–3429.
- [53] Alexander Sergeev and Mike Del Balso. 2018. Horovod: fast and easy distributed deep learning in TensorFlow. *CoRR* abs/1802.05799 (2018). [arXiv:1802.05799](http://arxiv.org/abs/1802.05799)
- [54] Qitao Shi, Ya-Lin Zhang, Longfei Li, Xinxing Yang, Meng Li, and Jun Zhou. 2020. Safe: Scalable automatic feature engineering framework for industrial tasks. In *2020 IEEE 36th International Conference on Data Engineering (ICDE)*. IEEE, 1645–1656.
- [55] Konstantin Shvachko, Hairong Kuang, Sanjay Radia, and Robert Chansler. 2010. The Hadoop Distributed File System. In *IEEE 26th Symposium on Mass Storage Systems and Technologies, MSST 2012, Lake Tahoe, Nevada, USA, May 3–7, 2010*, Mohammed G. Khatib, Xubin He, and Michael Factor (Eds.). IEEE Computer Society, 1–10. <https://doi.org/10.1109/MSST.2010.5496972>
- [56] Chijun Sima, Yao Fu, Man-Kit Sit, Liyi Guo, Xuri Gong, Feng Lin, Junyu Wu, Yongsheng Li, Haidong Rong, Pierre-Louis Aublin, et al. 2022. Ekko: A {Large-Scale} deep learning recommender system with {Low-Latency} model update. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. 821–839.
- [57] Chunqiang Tang, Kenny Yu, Kaushik Veeraraghavan, Jonathan Kaldor, Scott Michelson, Thawan Kooburat, Aravind Anbudurai, Matthew Clark, Kabir Gogia, Long Cheng, et al. 2020. Twine: A unified cluster management system for shared infrastructure. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. 787–803.
- [58] Alexey Tumanov, Timothy Zhu, Jun Woo Park, Michael A. Kozuch, Mor Harchol-Balter, and Gregory R. Ganger. 2016. TetriSched: Global Rescheduling with Adaptive Plan-Ahead in Dynamic Heterogeneous Clusters. In *Proceedings of the Eleventh European Conference on Computer Systems (London, United Kingdom) (EuroSys '16)*. Association for Computing Machinery, New York, NY, USA, Article 35, 16 pages. <https://doi.org/10.1145/2901318.2901355>
- [59] Taegeon Um, Byungsoo Oh, Byeongchan Seo, Minhyeok Kweun, Goeun Kim, and Woo-Yeon Lee. 2023. Fastflow: Accelerating deep learning model training with smart offloading of input data pipeline. *Proceedings of the VLDB Endowment* 16, 5 (2023), 1086–1099.
- [60] Abhishek Verma, Luis Pedrosa, Madhukar R. Korupolu, David Oppenheimer, Eric Tune, and John Wilkes. 2015. Large-scale cluster management at Google with Borg. In *Proceedings of the European Conference on Computer Systems (EuroSys)*. Bordeaux, France.
- [61] Luping Wang, Lingyun Yang, Yinghao Yu, Wei Wang, Bo Li, Xianchao Sun, Jian He, and Liping Zhang. 2021. *Morphling: Fast, Near-Optimal Auto-Configuration for Cloud-Native Model Serving*. Association for Computing Machinery, New York, NY, USA, 639a–653. <https://doi.org/10.1145/3472883.3486987>
- [62] Qinlong Wang, Tingfeng Lan, Yinghao Tang, Ziling Huang, Yiheng Du, Haitao Zhang, Jian Sha, Hui Lu, Yuanchun Zhou, Ke Zhang, and Mingjie Tang. 2024. DLRouter-RL: Resource Optimization for Deep Recommendation Models Training in the Cloud. arXiv:2304.01468 [cs.DC] <https://arxiv.org/abs/2304.01468>
- [63] Ruoxi Wang, Bin Fu, Gang Fu, and Mingliang Wang. 2017. Deep & cross network for ad click predictions. In *Proceedings of the ADKDD'17*. 1–7.
- [64] Zhuang Wang, Zhen Jia, Shuai Zheng, Zhen Zhang, Xinwei Fu, TS Eugene Ng, and Yida Wang. 2023. Gemini: Fast failure recovery in distributed training with in-memory checkpoints. In *Proceedings of the 29th Symposium on Operating Systems Principles*. 364–381.
- [65] Zhuang Wang, Zhen Jia, Shuai Zheng, Zhen Zhang, Xinwei Fu, T. S. Eugene Ng, and Yida Wang. 2023. GEMINI: Fast Failure Recovery in Distributed Training with In-Memory Checkpoints. In *Proceedings of the 29th Symposium on Operating Systems Principles*. ACM, Koblenz Germany, 364–381. <https://doi.org/10.1145/3600006.3613145>
- [66] Qizhen Weng, Wencong Xiao, Yinghao Yu, Wei Wang, Cheng Wang, Jian He, Yong Li, Liping Zhang, Wei Lin, and Yu Ding. 2022. {MLaaS} in the wild: Workload analysis and scheduling in {Large-Scale} heterogeneous {GPU} clusters. In *19th USENIX Symposium on Networked Systems Design and Implementation (NSDI 22)*. 945–960.
- [67] Yidi Wu, Kaihao Ma, Xiao Yan, Zhi Liu, Zhenkun Cai, Yuzhen Huang, James Cheng, Han Yuan, and Fan Yu. 2022. Elastic Deep Learning in Multi-Tenant GPU Clusters. *IEEE Transactions on Parallel and Distributed Systems* 33, 1 (2022), 144–158. <https://doi.org/10.1109/TPDS.2021.3064966>
- [68] Jie Amy Yang, Jianyu Huang, Jongsoo Park, Ping Tak Peter Tang, and Andrew Tulloch. 2020. Mixed-precision embedding using a cache. *arXiv preprint arXiv:2010.11305* (2020).
- [69] Guangshun Yao, Xiaoping Li, Qian Ren, and Rubén Ruiz. 2022. Failure-aware Elastic Cloud Workflow Scheduling. *IEEE Transactions on Services Computing* (2022).
- [70] Chunxing Yin, Bilge Acun, Carole-Jean Wu, and Xing Liu. 2021. Tt-rec: Tensor train compression for deep learning recommendation models. *Proceedings of Machine Learning and Systems* 3 (2021), 448–462.
- [71] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauly, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing. In *9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. USENIX Association, San Jose, CA, 15–28.
- [72] Daochen Zha, Louis Feng, Bhargav Bhusanam, Dhruv Choudhary, Jade Nie, Yuandong Tian, Jay Chae, Yinbin Ma, Arun Kejariwal, and Xia Hu. 2022. Autoshard: Automated embedding table sharding for recommender systems. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 4461–4471.
- [73] Daochen Zha, Louis Feng, Qiaoyu Tan, Zirui Liu, Kwei-Herng Lai, Bhargav Bhusanam, Yuandong Tian, Arun Kejariwal, and Xia Hu. 2022. Dreamshard: Generalizable embedding table placement for recommender systems. *Advances in Neural Information Processing Systems* 35 (2022), 15190–15203.
- [74] Kai Zhang, Hao Qian, Qing Cui, Qi Liu, Longfei Li, Jun Zhou, Jianhui Ma, and Enhong Chen. 2021. Multi-interactive attention network for fine-grained feature learning in ctr prediction. In *Proceedings of the 14th ACM international conference on web search and data mining*. 984–992.
- [75] Yongkang Zhang, Yinghao Yu, Wei Wang, Qiukai Chen, Jie Wu, Zuwei Zhang, Jiang Zhong, Tianchen Ding, Qizhen Weng, Lingyun Yang, et al. 2022. Workload consolidation in alibaba clusters: the good, the bad, and the ugly. In *Proceedings*

- of the 13th Symposium on Cloud Computing. 210–225.
- [76] Weijie Zhao, Deping Xie, Ronglai Jia, Yulei Qian, Ruiquan Ding, Mingming Sun, and Ping Li. 2020. Distributed hierarchical gpu parameter server for massive scale deep learning ads systems. *Proceedings of Machine Learning and Systems 2* (2020), 412–428.
- [77] Weijie Zhao, Jingyuan Zhang, Deping Xie, Yulei Qian, Ronglai Jia, and Ping Li. 2019. AIBox: CTR prediction model training on a single node. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management*. 319–328.
- [78] Guorui Zhou, Na Mou, Ying Fan, Qi Pi, Weijie Bian, Chang Zhou, Xiaoqiang Zhu, and Kun Gai. 2019. Deep interest evolution network for click-through rate prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 33. 5941–5948.
- [79] Guorui Zhou, Xiaoqiang Zhu, Chenru Song, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2018. Deep interest network for click-through rate prediction. In *Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining*. 1059–1068.