ASTRAEA: A Fair Deep Learning Scheduler for Multi-Tenant GPU Clusters

Zhisheng Ye[®], Peng Sun, Wei Gao, Tianwei Zhang[®], *Member, IEEE*, Xiaolin Wang, *Member, IEEE*, Shengen Yan, and Yingwei Luo, *Member, IEEE*

Abstract—Modern GPU clusters are designed to support distributed Deep Learning jobs from multiple tenants concurrently. Each tenant may have varied and dynamic resource demands. Unfortunately, existing GPU schedulers fail to thoroughly consider the fairness among the tenants and jobs, which can result in unbalanced resource allocation and unfair user experience. In this article, we present an efficient solution to provide strong fairness while maintaining high scheduling effectiveness in multi-tenant GPU clusters. First, we introduce a novel Long-Term GPU-time Fairness metric, which can comprehensively evaluate the fairness at both the tenant and job levels, based on both the temporal and spatial impacts of resource allocation. Second, we design a new and practical GPU scheduler, ASTRAEA, to enforce the desired fairness among tenants and jobs. Large-scale evaluations show that ASTRAEA can improve tenant fairness by up to 9.42× compared to state-of-the-art schedulers, without sacrificing the average job completion time.

Index Terms-Distributed systems, deep learning, GPU cluster scheduling

1 INTRODUCTION

DEEP learning (DL) has been utilized in a wide range of real-world applications, e.g., image recognition, natural language processing, recommendation systems [1]. State-ofthe-art DL models are trained from massive data samples, involving a large amount of computations. It is a common practice to leverage GPUs or other accelerators to speed up the training process [2], [3]. To deal with the ever-growing complexity of DL training (DLT) workloads and increased demands for computation resources, enterprises and institutes commonly set up large-scale GPU clusters.

A cluster is typically shared by multiple user groups (i.e., tenant), which can significantly improve resource utilization and reduce operational costs [4], [5]. It can concurrently

Manuscript received 29 Aug. 2021; revised 10 Nov. 2021; accepted 7 Dec. 2021. Date of publication 17 Dec. 2021; date of current version 23 May 2022.

(Corresponding author: Yingwei Luo.)

Recommended for acceptance by A. J. Peña, M. Si, and J. Zhai. Digital Object Identifier no. 10.1109/TPDS.2021.3136245 serve a large number of jobs with different features. Table 1 shows the running time distribution of DLT jobs from two production GPU clusters operated by Microsoft and Sense-Time. We observe that there are long-term jobs which take hours or days to complete. Besides, about 41% (Philly) and 67% (Venus) of the DLT jobs can be completed within 10 minutes. These short-term jobs generally adopt the feed-back-driven exploration method for the research and debugging purposes [6], [7].

The mixed workloads of long-term and short-term jobs call for fair resource allocation, and it is crucial to provide sharing incentive in a shared cluster. Here sharing incentive refers to the property that if N jobs share a cluster, then each job should not have worse performance compared to an independent cluster with 1/N of the resources [8]. On one hand, long-term jobs should not monopolize the entire cluster, which can significantly block the execution of shortterm jobs (a.k.a. head-of-line blocking [6], [7]). On the other hand, arbitrarily prioritizing short-term jobs can cause starvation for long-term jobs. It is necessary to balance all these types of jobs and satisfy every tenant. Unfortunately, most of the existing DLT schedulers, e.g., Slurm [9], Kubernetes [10], Yarn-CS [4], Tiresias [7], Gandiva [6], Hived [11] and Pollux [12] ignore the fairness issue in GPU clusters, which can result in severe discrimination of resource allocation across DLT jobs. Gandiva fair [13] targets the fairness of GPU clusters. However, it mainly considers fair share of GPU resources across tenants instead of DLT jobs.

This paper aims to *design new scheduling solutions to achieve fairness for DLT jobs in multi-tenant GPU clusters.* Past works have studied similar problems in the big-data cluster environment. The mainstream mechanism is *instantaneous resource-allocation fairness,* which can ensure instantaneous fair allocations across big data jobs (e.g., map-reduce tasks) by balancing the allocated resources among all the jobs in real-time [4], [5], [14], [15], [16]. Unfortunately, it is difficult

[•] Zhisheng Ye is with the School of Computer Science, Peking University, Beijing 100871, China, and also with Peng Cheng Laboratory, Shenzhen 518000, China, and also with SenseTime Research, Beijing 100080, China. E-mail: yezhisheng@pku.edu.cn.

Peng Sun and Shengen Yan are with SenseTime Research, Beijing 100080, China. E-mail: {sunpeng1, yanshengen}@sensetime.com.

Wei Gao and Tianwei Zhang are with the School of Computer Science and Engineering, Nanyang Technological University, Singapore 639798, and also with S-Lab, Nanyang Technological University, Singapore 639798. E-mail: (gaow0007, tianwei.zhang)@ntu.edu.sg.
 Xiaolin Wang and Yingwei Luo are with the School of Computer Science,

Xiaolin Wang and Yingwei Luo are with the School of Computer Science, Peking University, Beijing 100871, China, and also with the Peng Cheng Laboratory, Shenzhen 518000, China.
 E-mail: {wxl, lyw}@pku.edu.cn.

This work was supported in part by the National Science Foundation of China under Grants 62032001, 61672053, U1611461, and 62032008, in part by the RIE2020 Industry Alignment Fund - Industry Collaboration Projects (IAF-ICP) Funding Initiative, and in part by cash and in-kind contributions from the industry partner(s).

^{1045-9219 © 2021} IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See https://www.ieee.org/publications/rights/index.html for more information.

 TABLE 1

 Running Time Distributions of DLT Jobs in Two GPU Clusters:

 Philly From Microsoft and Venus From SenseTime

| Time | < 10 min | < 1 hour | < 6 hour | < 1 day | < 6 day |
|-----------------|----------------|----------------|--------------------|----------------|------------------|
| Philly Venus | 40.8% 67.3% | 71.6% 83.3% | $88.5\% \\ 90.6\%$ | 94.1% 96.2% | $98.5\%\ 99.6\%$ |

to achieve instantaneous fairness for DLT jobs in GPU clusters, due to their distinct characteristics. (1) Mainstream DL frameworks [17], [18] require gang scheduling, i.e., all the requested GPUs should be offered to the job in an all-ornothing manner. Under this restriction, it is infeasible to split partial resources from one job and give them to others for instantaneous fairness. (2) DLT jobs have high GPU affinity requirements. The performance of communicationintensive DLT jobs exhibits high sensitivity to inter-GPU topography. The instantaneous resource-allocation fairness mechanism cannot take this factor into consideration. (3) The preemption overhead of a DLT job is large for saving and loading the model. Frequent preempting DLT jobs for instantaneous fairness can cause huge performance penalty.

Hence, it is necessary to design a fairness solution dedicated to DLT workloads. Themis [8] made such an attempt by proposing the long-term finish-time fairness metric for DLT jobs. It considers the gang scheduling and GPU affinity requirements, and leverages a lease-based scheduling to handle the preemption overhead of DLT jobs. At each scheduling round, this mechanism tries to figure out the optimal resource allocation strategy based on the remaining time of each job. However, the metric only measures fairness in terms of completion time without considering different GPU requirements of DLT jobs. This can disincentive small DLT jobs with smaller GPU requirements from using the shared cluster. More seriously, a dishonest user may submit DLT jobs with overclaimed GPU demands to intentionally get higher throughputs with the same scheduling priority, which could lead to more serious waste of cluster resources and exacerbate the unfairness issue. We will give detailed analysis about these mechanisms in Section 4.1.

Motivated by the above limitations, we present a novel solution for fair GPU resource allocation across DLT workloads. We make the following two key contributions. First, we design a new metric, Long-Term GPU-Time Fairness (LTGF), to comprehensively characterize and evaluate the fairness of DLT jobs. The comprehensiveness of this metric is reflected from two aspects. (1) LTGF considers both the temporal and spatial impacts of a resource allocation on the fairness of DLT jobs. At a long-term scale, it quantifies the gap between the amount of resources allocated to the job and the amount of resources the job deserves to obtain under the sharing incentive. This can give more reasonable assessment than the finish-time fairness in Themis, which only considers the temporal factor. (2) LTGF considers fairness at both the tenant and job levels. At the tenant level, it distributes to different tenants fair amounts of GPU resources, which are proportional to their contributions (e.g., budget or GPU resources) over a period of time. At the job level, it ensures that GPU time is fairly allocated among concurrent DLT jobs in the long term within a tenant. This can provide guaranteed service for tenants, as well as high cluster utilization.

Second, we design ASTRAEA,¹ a practical and efficient scheduling system to achieve fairness in GPU clusters. It can simultaneously ensure fair share of GPU resources among tenants as well as jobs without bringing noticeable overhead for DLT jobs. Specifically, ASTRAEA adopts the LTGF metric to evaluate fairness and identify the optimal resource allocation decisions. It also utilizes a lease-based scheduling scheme, which can efficiently preempt running jobs to balance fairness and Job Completion Time (JCT). Evaluation shows that ASTRAEA outperforms state-of-the-art fair schedulers, and improves tenant-level fairness by up to $9.42 \times$ and job fairness by up to $10.3 \times$ without sacrificing the average JCT.

2 BACKGROUND

2.1 Multi-Tenant GPU Cluster

It becomes popular to operate the GPU cluster in a multitenant fashion. This can bring many benefits, such as reducing operational cost and improving resource utilization. A shared GPU cluster is divided into multiple Virtual Clusters (VCs), with each one assigned to a tenant. Based on a tenant's resource weight, its VC is associated with a static or dynamic resource quota in terms of the number of GPUs and other resources (e.g., CPU, RAM) [19]. Users of a tenant can only submit DLT jobs to their corresponding VC. A cluster scheduler is in charge of selecting DLT jobs from the pending job pool and placing them on GPUs for execution.

2.2 Deep Learning Training

A Deep Learning training (DLT) job aims to determine the optimal parameter values of a Deep Learning model from training data. Users submit their jobs to the cluster. Then a scheduler selects appropriate jobs from the pending queue and places them to the demanded GPU resources for execution. A running DLT job may be preempted by other jobs. There are three possible final states for a DLT job: (1) COM-PLETED: a DLT job successfully executes all the epochs of its training procedure. (2) CANCELED: users may manually early stop the execution of a DLT job socould be terminated by force due to programming or hardware issues [19]. DLT jobs have following features.

Iterative Computing. A DLT job is executed in an iterative fashion. At each iteration, a batch of training samples are selected as the input of two steps of computation: forward and backward propagation. A DLT job may conduct millions of iterations to achieve model convergence, resulting in long running time (e.g., tens of days) [19], [20]. It is common to leverage hardware accelerators (e.g., GPU) to improve the computation efficiency. *Feedback-Driven Exploration*. Users usually perform exploration tasks (e.g., searching different configurations or hyper-parameters, selecting the optimal network structures) in a trial-and-error manner before obtaining the final model [6], [7]. This exploration can be performed by hyperparameter-tuning frameworks [21], [22]. During this exploration process, poor-performing jobs will be killed in advance based on the intermediate results.

Gang Scheduling. Distributed DLT jobs could use multiple GPUs to train a single model using data-parallel or model-



Fig. 1. Distribution of requested GPU resources.



Fig. 2. Distribution of DLT job GPU service time.

parallel techniques. They usually require gang scheduling: all the requested GPUs should be offered in an all-or-nothing manner. While some new training solutions [23] with elastic GPU resources have been proposed, they are still at an early stage without broad deployment.

Placement Sensitivity. Distributed training procedure needs to exchange data between GPUs at every iteration. To guarantee training performance, communication-sensitive jobs require strict topology of allocated GPUs. For example, a job with 8 GPUs should be placed on a single 8-GPU node for the high performance of the nvlink-based communication. If this job is allocated with two nodes, the inter-node communication could be the performance bottleneck [3].

2.3 DLT Job Scheduling

A scheduler is indispensable in a cluster to manage the computing resources and schedule jobs. At runtime, the scheduler continuously receives jobs submitted by users with the explicit resource demands. Then it decides when and where to run by selecting the proper jobs to satisfy the scheduling objectives and placing them to the appropriate servers for execution [24]. Different scheduling systems may pursue various scheduling objectives, e.g., improving resource utilization [5], [20], [25], optimizing the performance of workloads [26], [27], maximizing scheduling efficiency [14], [28], guaranteeing service and user experience [29], [30], etc. Different from the traditional big-data or HPC workloads, DLT jobs exhibit unique characteristics (Section 2.2). This brings new challenges for GPU cluster scheduling, and it is necessary to design dedicated schedulers for DLT jobs, which will be discussed in Section 7.2.

3 A STUDY OF REAL-WORLD DLT JOB TRACE

In this section, we study Venus [20], a production DLT cluster in SenseTime. Venus contains 133 GPU servers connected by RDMA network. Each server is installed with 8 NVIDIA Volta GPUs. The cluster is shared among 16 tenants. Each tenant is allocated with a static number of GPUs. Venus uses Slurm [9] as the scheduler. We collect and analyze 109734 GPU-based DLT jobs over a period of 6 months.



Fig. 3. CDF of DLT job running time and GPU service time.



Fig. 4. GPU utilization of two VCs.

3.1 Characteristics of DLT Jobs in Shared Clusters

Unpredictable Training Time. In Venus, 52.2% of jobs are successfully completed, 27.9% are canceled, and 19.9% fail. Due to the high cancellation/failure ratio, it is difficult to predict the running time of DLT jobs just based on the remaining number of iterations. [31].

Various GPU Demands. Fig. 1 shows the distribution of DLT job in terms of requested GPU resources. About 52.5% of jobs use single GPU, 22.6% require 8 GPUs, and 10.3% need more than 8 GPUs. Jobs have high cancellation/failure ratio in all the cases. We define the *GPU service time* of a DLT job as the product of its GPU demand and running time. Fig. 2 shows such distribution. While single-GPU jobs account for the largest proportion of all the jobs, they only consume 4.7% of total GPU service time. The major GPU service time is consumed by DLT jobs with no less than 8 GPUs (85.7%).

Mixed Workload of Long-Term and Short-Term Jobs. Table 1 demonstrates that DLT job running time varies from minutes to days. Fig. 3 further shows the cumulative distributions of job running time and service time with different resource demands. We observe that DLT job running time follows the long-tail distribution. For example, for single-GPU jobs, 87.4% of jobs last for less than 1 hour, while 2.2% spend more than 1 day. In addition, DLT jobs with larger GPU demands tend to have longer running time.

Unbalanced Resource Demands. The GPU resource demands of VCs in a shared DLT cluster are unbalanced over time. Fig. 4 gives the VC utilization of two tenant. On June 15, Tenant-B has a high cluster GPU utilization of over 90%, but Tenant-A has an average utilization of less than 50%. On June 21, Tenant-A achieves near 100% utilization, while Tenant-B only uses 60% of its GPUs.



(a) CDF of job pending time. (b) CDF of job slowdown ratio.



Implication #1: The unpredictability of training time, and high variety in the GPU demands and running time significantly increase the difficulty of making optimal scheduling decisions and allocating resources.

3.2 Issues of DLT Clusters Without Fairness

Venus uses FIFO (First-In-First-Out) to schedule jobs without considering fair resource allocation across parallel DLT jobs. Long-term jobs may monopolize the whole cluster or VC, blocking the execution of pending jobs (a.k.a. head-of-line blocking). As shown in Fig. 5a, a portion of jobs have extremely long pending time. For example, 4.4% of single-GPU jobs wait for over 30 minutes to be scheduled. Jobs with 4 or more GPUs may have higher pending time. Blocked DLT jobs (especially for short-term jobs) may suffer from large execution slowdown ratio, which is defined as (*running_time* + *pending_time*)/*running_time*. As shown in Fig. 5b, a certain number of jobs have large (> 5) slowdown ratios.

Implication #2: Scheduling without considering fairness could worsen the job pending situation and result in unfair execution across different jobs.

4 LONG-TERM GPU-TIME FAIRNESS

In this section, we analyze the limitations of existing fairness mechanisms, and then present our new fairness metric. This will further inspire the design of our fair scheduler.

We consider two types of sharing incentive in this work.

- **Definition 1** (Sharing incentive for DLT jobs) Consider a M-GPU cluster hosting N DLT jobs with the same weight. We say there is sharing incentive for a job if its performance is no worse than the situation where it runs in an independent cluster of M/N GPUs.
- **Definition 2** (Sharing incentive for tenants) Consider a M-GPU cluster shared by T tenants with the same weight. We say there is sharing incentive for a tenant, if he could be allocated with at least the same amount of GPU service time as an independent cluster of M/T GPUs over a period of time, regardless of the workloads of other tenants.

4.1 Analysis of Existing Fairness Mechanisms

We give some illustrative examples to demonstrate the limitations of two common fairness mechanisms, as well as the innovation and superiority of our metric.



Fig. 6. Illustration of fair resource allocation. (a) Instantaneous resourceallocation fairness. (b) Long-term finish-time fairness. (c) Our proposed Long-term GPU-Time fairness.

4.1.1 Instantaneous Resource-Allocation Fairness

This mechanism is widely adopted by existing big-data clusters to achieve instantaneous fairness [4], [5], [14], [15], [16]. Consider there are N active jobs in the cluster for scheduling at one instant t. Each job i receives $r_i(t)$ resources from the scheduler. Then the goal of the instantaneous resource-allocation fairness is to figure out the optimal policy with the following objective:

maximize :
$$\min_{i \in N} r_i(t)$$
.

If a job does not fully utilize the allocated resources, these excess resources will be redistributed to other jobs.

Fig. 6a gives an example of this mechanism. Consider a cluster of 6 CPUs with three active jobs. The scheduler should allocate these CPUs equally to each job regardless of their types or demands. Following this principle, the scheduler is able to provide *instantaneous* fairness across all the jobs at any time.

Unfortunately, instantaneous resource-allocation fairness is hard to achieve for DLT jobs. As discussed in Section 1, DLT jobs require gang scheduling with high placement sensitivity and preemption overhead. These characteristics make it inflexible for the scheduler to adjust the resource allocation instantaneously to guarantee fairness.

4.1.2 Long-Term Finish-Time Fairness

This metric was proposed in [8] to address the fairness issue for Hyper-Parameter Optimization (HPO) workloads. It allocates resources to balance the finish time of each HPO job. It is defined as $\rho_i(t) = T_i(t)/(T_i^*(t) \cdot N)$, where $T_i(t)$ and $T_i^*(t)$ is the estimated time to finish job *i* using the actual allocated resources and the entire cluster, respectively. At each scheduling round, the scheduler tries to identify the optimal solution with the following objective:

minimize :
$$\max_{i \in N} \rho_i(t)$$
.

Fig. 6b gives an example when directly using this metric to schedule DLT jobs. Three DLT jobs are submitted to a 6-GPU cluster. *Job1* requires 6 GPUs while *Job2* and *Job3* require 3 GPUs. Each of them needs 40 minutes to complete. The duration of one scheduling round is 10 minutes. Guided by this long-term finish-time fairness principle, these jobs will be scheduled alternatively until their completion.

This metric only considers fairness in terms of completion time, while ignoring the amount of allocated resources. In Fig. 6b, for each period of 20 minutes, *Job1* gets 1/2 of the total GPU resources. In contrast, *Job2* and *Job3* only get 1/4 of the resources each. This cluster loses sharing incentive for these two jobs. Moreover, since jobs with higher GPU demands can get more GPU service time, tenants may overclaim excessive resources for their jobs to increase the throughput, resulting in a waste of GPU resources.

A fair scheduler should consider the spatial information (GPU demand) in addition to the temporal information (running time) of DLT jobs. Using the above example, the scheduler should give one round for *Job1* and two rounds for *Job2* and *Job3* (Fig. 6c). Then equal resources are allocated to each running job in any half hour. Our proposed metric can achieve this fairness, as described below.

4.2 Long-Term GPU-Time Fairness Definition

We propose a novel metric, Long-Term GPU-time Fairness (LTGF) for DLT scheduling. In a single-tenant cluster, LTGF achieves sharing incentive for DLT jobs. In a multi-tenant cluster, LTGF also provides sharing incentive for tenants.

4.2.1 LTGF in Single-Tenant Clusters

Consider a single-tenant cluster of M GPUs shared by N jobs from time t_1 to t_2 . Let $Demand_i^{job}$ be the GPU demand of job J_i , $Alloc_i^{job}(t)$ be the amount of GPUs allocated to J_i at time t. Due to the gang scheduling feature, we have

$$Alloc_i^{job}(t) = \begin{cases} 0, & \text{if } J_i \text{ is not running} \\ Demand_i^{job}, & \text{if } J_i \text{ is running} \end{cases}.$$
(1)

Over a period time of $[t_1, t_2]$, the accumulated GPU service time allocated to J_i by the scheduler is

$$\overline{Alloc}_i^{job}(t_1, t_2) = \int_{t_1}^{t_2} Alloc_i^{job}(t) \, dt.$$
⁽²⁾

We use $Active(J_i, t)$ to denote the state of Job J_i at time t: pending or running jobs have $Active(J_i, t) = 1$ while failed or completed jobs have $Active(J_i, t) = 0$. Given the share weight W_i^{job} of J_i , based on the max-min fairness, its instantaneous fair number of allocated GPUs at time t is

$$FairShare_i^{job}(t) = \frac{M \times W_i^{job}}{\sum_{j=1}^N (Active(J_i, t) \times W_j^{job})}.$$
 (3)

Due to gang scheduling, J_i could not use more GPUs than $Demand_i^{job}$ even if $FairShare_i^{job}(t)$ has a large value. The fair GPU service time of J_i over $[t_1, t_2]$ is

$$\overline{FairShare}_{i}^{job}(t_{1}, t_{2}) = \int_{t_{1}}^{t_{2}} \min(Demand_{i}^{job}, FairShare_{i}^{job}(t)) dt.$$
(4)

The fairness degree $\rho_i^{job}(t_1,t_2)$ for the job i from t_1 to t_2 is defined as

$$\rho_i^{job}(t_1, t_2) = \frac{\overline{Alloc}_i^{job}(t_1, t_2)}{\overline{FairShare}_i^{job}(t_1, t_2)}.$$
(5)

 $\rho_i^{job}(t_1, t_2) \ge 1$ implies good long-term fairness in terms of GPU service time for job J_i . In contrast, $\rho_i^{job}(t_1, t_2) < 1$

indicates unfairness since it violates the sharing incentive. For example, we assume all the jobs have the same weight and keep active over $[t_1, t_2]$. The GPU service time that J_i receives in the shared cluster is $\rho_i^{job}(t_1, t_2)$ $(t_2 - t_1)\min(Demand_i^{job}, M/N)$. Considering an independent cluster of M/N GPUs, the GPU service time that J_i acquires is $(t_2 - t_1)\min(Demand_i^{job}, M/N)$. It is noted that J_i could not run with M/N GPUs if $Demand_i^{job} > M/N$ due to the gang scheduling. The performance of J_i in a shared cluster is no worse than the situation where it runs in an independent cluster of M/N GPUs, if $\rho_i^{job}(t_1, t_2) \ge 1$.

4.2.2 LTGF in Multi-Tenant Clusters

Consider a cluster of M GPUs shared by T tenants. It is divided into T VCs, with each one assigned to a tenant. Each tenant j is associated with a weight W_j^{tenant} . Then the GPU quota of this tenant $Quota_j$ is calculated as

$$Quota_j = \frac{M \times W_j^{tenant}}{\sum_{k=1}^{T} (W_k^{tenant})}.$$
(6)

Let N_j be the number of jobs for tenant j, and $Demand_{i,j}^{job}$ be the GPU demand of $J_{i,j}$ from this tenant. The requested amount of GPUs by tenant j at any time is

$$Demand_{j}^{tenant} = \sum_{i=1}^{N_{j}} (Active(J_{i,j}, t) \times Demand_{i,j}^{job}).$$
(7)

If $Demand_j^{tenant}(t) < Quota_j(t)$, the tenant would not use more GPUs than its resource demands. The fair number of allocated GPUs for tenant tenant j at time t is

$$FairShare_{j}^{tenant}(t) = \min(Demand_{j}^{tenant}(t), Quota_{j}(t)).$$
(8)

The fair GPU service time that tenant j should receive over a period of $[t_1, t_2]$ is

$$\overline{FairShare}_{j}^{tenant}(t_{1}, t_{2}) = \int_{t_{1}}^{t_{2}} FairShare_{j}^{tenant}(t) dt.$$
(9)

As shown in Section 3, the total GPU demands of tenants are unbalanced and varying over time. The actual amount of GPUs allocated to tenant j at time t is

$$Alloc_{j}^{tenant}(t) = \sum_{i=1}^{N_{j}} (Alloc_{i,j}^{job}(t)),$$
(10)

where $Alloc_{i,j}^{job}(t)$ denotes the amount of GPUs allocated to job $J_{i,j}$ at time *t* (Equation (1)).

When $Alloc_j^{tenant}(t) < Quota_j$, the tenant does not fully utilize its resource quota due to insufficient workloads or GPU resource fragmentation. In contrast, $Alloc_j^{tenant}(t) > Quota_j$ means the tenant uses more resources than its quota to improve the overall cluster utilization.

Over a period of $[t_1, t_2]$, the actual GPU service time allocated to tenant *j* by the scheduler is

Authorized licensed use limited to: Nanyang Technological University. Downloaded on June 13,2022 at 02:15:35 UTC from IEEE Xplore. Restrictions apply.

$$\overline{Alloc}_{j}^{tenant}(t_1, t_2) = \int_{t_1}^{t_2} Alloc_{j}^{tenant}(t) \, dt. \tag{11}$$

The fairness degree $\rho_j^{tenant}(t_1, t_2)$ for tenant *j* from t_1 to t_2 is defined as follows:

$$\rho_j^{tenant}(t_1, t_2) = \frac{\overline{Alloc}_j^{tenant}(t_1, t_2)}{\overline{FairShare}_j^{tenant}(t_1, t_2)}.$$
 (12)

 $\rho_j^{tenant}(t_1, t_2) \ge 1$ implies good long-term fairness in terms of GPU service time for tenant T_i . It means the tenant can use more or at least the same amount of GPU service time than its fair share. In contrast, $\rho_i^{job}(t_1, t_2) < 1$ indicates unfairness since the tenant receives less GPU service time than a separate cluster of M/T GPUs.

We can also calculate the job fairness for $J_{i,j}$. Let $W_{i,j}^{job}$ be the share weight of $J_{i,j}$ in its corresponding VC. Its instantaneous fair resource share at time t should be

$$FairShare_{i,j}^{job}(t) = \frac{FairShare_{j}^{tenant}(t) \times W_{i,j}^{job}}{\sum_{i=1}^{N_{j}} (Active(J_{i,j}, t) \times W_{i,j}^{job})}.$$
 (13)

Its fair GPU service time over $[t_1, t_2]$ should be

$$\overline{FairShare}_{i,j}^{job}(t_1, t_2) = \int_{t_1}^{t_2} \min(Demand_{i,j}^{job}, FairShare_{i,j}^{job}(t)) dt.$$
(14)

The job fairness degree $\rho_{i,j}^{job}(t_1, t_2)$ for $J_{i,j}$ from t_1 to t_2 is defined as follows:

$$\rho_{i,j}^{job}(t_1, t_2) = \frac{\overline{Alloc}_{i,j}^{job}(t_1, t_2)}{\overline{FairShare}_{j\,i}^{job}(t_1, t_2)}.$$
(15)

4.2.3 Application

Our proposed LTGF can provide guidance for realizing fairness scheduling at both the tenant and job levels. To achieve tenant-level fairness, we can use the max-min approach to maximize the minimal ρ_j^{tenant} in Equation (12)

maximize :
$$\min_{j \in T} \rho_j^{tenant}(t_1, t_2).$$

Job level fairness can also be guaranteed in a similar way with the following objective:

$$\texttt{maximize}: \min_{i \in N_j} \rho_{i,j}^{job}(t_1,t_2)$$

Since these two parameters are used to regulate two different levels of fairness separately, we do not unify them. On one hand, our proposed two-phase scheduling algorithm (Section 5.2) enables to achieve both job-level and tenantlevel fairness simultaneously. On the other hand, these two parameters give the cluster administrators more flexibility to focus on different types of fairness based on their demands and situations.

Compared to the long-term finish-time fairness in Themis, LTGF can provide more comprehensive assessment for both temporal and spatial impacts at the tenant and job levels. Besides, LTGF computes the fairness of a job or tenant over a past period. In contrast, Themis needs to predict the remaining time of each job to calculate the finish-time fairness, which is not always accurate (Section 3.1), and can affect the subsequent scheduling decisions.

5 A FAIR SCHEDULING SYSTEM

In this section, we introduce ASTRAEA, a fair DLT job scheduler for GPU clusters. ASTRAEA achieves our proposed Long-Term GPU-time Fairness with two techniques. The first one is a lease-based training scheme (Section 5.1). It breaks a longterm job into a sequence of multiple sub-jobs, which enables balancing jobs with different running time, and rearranging job execution orders. The second innovation is a two-phase scheduler based on LTGF across tenants and jobs (Section 5.2). It maintains the two-level fairness using a max-min approach.

Fig. 7 shows the overview and workflow of ASTRAEA. Different tenants submit jobs to their pending queues individually (**0**), which are then processed by the two-phase scheduler. The scheduler first selects tenants based on the tenant-fairness metric (**0**), and then selects one job from the tenant's pending queue based on the job-fairness metric (**0**). The selected job will be allocated to the resource pool according to the placement strategy (**0**). A job will return to its pending queue for renewal when its lease is expired (**0**). Since job placement strategy is not included in our contributions, we use a common consolidated job placement strategy along with ASTRAEA, which deploys a job on as few nodes as possible, following [6], [7], [19]. Other state-of-the-art GPU placement algorithms determined by cluster administrators can be integrated into ASTRAEA as well.

To summarize, we leverage the lease-based training scheme for a more flexible arrangement of DLT jobs. We utilize the two-phase fairness scheduling algorithm to make scheduling decisions and achieve LTGF in a multi-tenant cluster. Below we detail each technique.

5.1 Lease-Based Training Scheme

To balance the jobs with different GPU demands and finishtime, we design a novel lease-based DL training scheme. As discussed in Section 2.2, a DLT job is usually a long iterative process, causing inflexibility to the scheduler. Thus, we propose to break a long training job into a series of periods with a fixed length, namely lease terms. As shown in Fig. 8, at every scheduling round, each pending job needs to request for a lease term. If this request is approved, the job will also receive the demanded resources, and perform the computation. At the end of this lease, the running job needs to proactively make a checkpoint, and then submit the renewal request for the next lease term. If the renewal is successful, the job can continue the execution. Otherwise, it will be preempted and its resources are released to other jobs and tenants. Jobs that can be finished within one term can release the resources immediately after the completion without the need for lease renewal.

The lease-based training scheme is essential to the scheduling process and is integrated with the two-phase scheduling algorithm tightly. Based on the scheduling decision made by the algorithm, the scheme is used to enforce resource re-



Fig. 7. Overview of ASTRAEA.

allocation between jobs. For the scheduler, at every cycle, it needs to decide whether the lease renewal request should be approved for each running job. If yes, the scheduler will send a notification of "successful renewal" to the job so it can continue the execution. Otherwise, the scheduler will suspend the job and reclaim the resources. In addition, the scheduler also needs to select pending jobs, grant lease terms and resources to them, and resume their computation. All these decisions are made by our two-phase scheduling algorithm with the LTGF metric. The concept of lease was adopted in various domains, e.g., networking [32], cloud computing [33], computer architecture [34], etc. Previous studies also leveraged similar ideas for resource management and DLT job scheduling [6], [7], [8]. Below we present some design details.

5.1.1 On-Demand Checkpoint

The lease-based training scheme brings imperative resource re-allocation and reclamation, introducing the need of ondemand checkpoint for DLT jobs. A job needs to make a checkpoint before the end of the lease term to avoid the loss of progress. We make some changes in the user-side DL library to achieve this function. A daemon is created when a job starts to execute. When the job losses its lease term and resources, the scheduler sends a POSIX signal to notify the daemon of saving the current checkpoint. However, a checkpoint has to be saved at the end of one iteration. There may be not enough time for the job to finish the current iteration before the resources are reclaimed. To solve this issue, ASTRAEA implements a last store technique. At every iteration, the job updates the model and other information as a snapshot. Once receiving the preemption signal, the daemon flushes the latest snapshot to the disk and waits for the current iteration to complete. An updated snapshot is made and stored if this iteration is finished before the job preemption. Otherwise, the previous flushed snapshot will be used for the future lease term.

5.1.2 Lease Renewal

There is minor modification over the conventional scheduling algorithm to accommodate the lease-based training



Fig. 8. Job execution under the lease-based training scheme.

scheme, along with other renewal mechanisms. At each scheduling cycle, the scheduler needs to handle the requests from both the running and pending jobs. ASTRAEA processes these requests with the following steps. (1) The scheduler treats all the jobs as in the pending state, and all the resources as available. It selects the jobs to be scheduled based on the fairness algorithm in Section 5.2. (2) For each running job in the current lease term, if it is selected by the scheduler, then its lease renewal is successful, and the execution will continue in the next lease with the current assigned resources. Otherwise, the job will be preempted and the resources are reclaimed. (3) After all the resources of the preempted jobs are freed, the scheduler performs the placement policy (consolidated placement strategy in our implementation) to allocate resources to the selected pending jobs, and then resume their execution.

We introduce a coordinator to reconcile the job renewal. This coordinator is responsible for not only triggering the lease renewal and reclaiming the resources of preempted jobs, but also resuming jobs with new assigned lease terms. The resumed jobs will continue running from the last checkpoint with the help of a network-attached storage system.

5.1.3 Lease Term Length

The length of a lease term is critical to determine the scheduling efficiency and training performance. A short lease term can increase the frequency of making checkpoints and training overhead, and cause heavy loads for the scheduler. A long lease term fails to appropriately adjust the job execution order, leading to poor fairness and JCT. Therefore, how to set a proper lease term is a challenging problem, as it depends on many configurations and behaviors of the cluster and DLT jobs.

We can heuristically identify the proper length for the target cluster based on its historical record. Specifically, we collect a trace of past jobs, simulate the scheduling behaviors, and measure the corresponding job fairness and performance under different lease term lengths. Based on the simulation results, we pick the optimal value that best balances the fairness and performance for the cluster. Based on the running time distribution of DLT jobs in Section 3, a majority of jobs have short duration (nearly 70% jobs can be finished within 10 minutes). Therefore, an appropriate length ensures that most jobs can be completed within one

Authorized licensed use limited to: Nanyang Technological University. Downloaded on June 13,2022 at 02:15:35 UTC from IEEE Xplore. Restrictions apply.

lease term. This significantly reduces the overhead. We can periodically collect the traces and simulate the scheduler to adjust the configuration of term length adaptively. Detailed results can be found in Section 6.2.

5.2 Two-Phase Fairness Scheduling Algorithm

To achieve the two-level LTGF defined in Section 4.2, we design a novel two-phase scheduling algorithm, as shown in Algorithm 1. The core of this algorithm is the Allocate-Job function, which is called in every scheduling round.

In this function, the scheduler first updates the fairness metric for each tenant based on Equation (12) (line 2), and for each job based on Equation (15) (line 3). These metrics are computed from the beginning t_{begin} to the current moment t_{now} . Then at Phase 1, the scheduler first maintains tenant-fairness by selecting the tenant p^* with the smallest fairness index (line 5). If this tenant does not have any pending jobs, the scheduler will skip it. Otherwise, it will continue with Phase 2, which selects the job with the highest priority from the tenant's pending queue (line 7). If there are enough resources and the job could be satisfied with the current placement strategy, the scheduler will grant the lease term and demanded GPUs to this selected job (line 10-15). Otherwise, it will skip to the next tenant to prevent starvation. Finally, it removes the job from its corresponding pending queue (line 13) and updates the tenant level fairness metrics as if the allocated job was running (line 14-15). The scheduler repeats the above two phases until there are no pending jobs or adequate resources.

Algorithm 1. Two-Phase Job Scheduling

Input: *T* tenants and *N* jobs, a cluster of *M* free GPUs.

1: Function AllocateJob

- 2: Update ρ^{tenant} for each tenant with Equation (12)
- 3: Update ρ^{job} for each job with Equation (15)
- 4: while there exist pending jobs ∈ N with demand r ≤ M ∧ satisfy placement strategy do // Phase 1: Tenant selection
 5: p* = arg min_{j∈T} ρ_j^{lenant}(t_{begin}, t_{now}) // Phase 2: Job selection
 6: J_{p*} ← pending job list of tenant p*
 7: i* = SelectJob(J_{p*})
 8: r_{i*} ← GPU demand of job i* // Apply current placement strategy

9:
$$\mathcal{A}_{i^*} \leftarrow \texttt{Placement}(M, i^*)$$

10: **if**
$$r_{i^*} \leq M \wedge A_{i^*} \neq \emptyset$$
 then

```
// Allocate r_{i^*} resources to job.
```

```
11: Allocate (\mathcal{A}_{i^*}, i^*)
```

```
12: M = M - r_{i^*}
```

```
13: J_{p^*} = J_{p^*} \setminus \{i^*\}
```

```
// Update the related metrics of tenant p^*.
```

14:
$$\overline{Alloc}_{j}^{tenant}(t) = \overline{Alloc}_{j}^{tenant}(t) + r_{i^{*}} * t_{lease}$$

15: Update
$$\rho_j^{tenant}$$
 according to Equation (12).

16: else

// Continue with the rest tenants.
17:
$$T = T \setminus \{p^*\}$$

5.2.1 Job Selection

In Phase 2, the scheduler selects a job from the tenant's pending queue for scheduling. Algorithm 2 shows the procedure of this function. Basically, it calculates the reward for each job based on the job-level fairness index. Then it selects the job with the maximal reward (e.g., smallest fairness index) as the candidate.

Note that this job selection process can be integrated with other job-level prioritization approaches. For instance, three types of jobs can be assigned with higher priority for selection if not scarifying the fairness. (1) Some jobs have ultrashort duration. If we prioritize the newly submitted jobs, these short jobs can be scheduled and completed promptly, which can reduce the job pending overhead. (2) Some jobs have been executed for a large number of lease terms, and more likely have a relatively long duration. Thus, special attention to the lease renewal of these jobs can reduce the total overhead caused by frequent preemption. (3) Users may specify some important and urgent jobs that need to meet certain deadlines. The scheduler can also increase the reward of these jobs during selection, to ensure they are not affected by lease expiration.

| Algorithm 2. Select a Candidate Job From a Job List | | | | |
|--|--|--|--|--|
| Input: A job_list J | | | | |
| Output: Selected job from <i>J</i> . | | | | |
| 1: Function SelectJob | | | | |
| 2: forall each job <i>i</i> in <i>J</i> do | | | | |
| // Calculate job reward. | | | | |
| 3: $i.reward = -\rho_i^{job}(t_{begin}, t_{now})$ | | | | |
| 4: $i^* = \arg \max_{i \in J} i.reward$ | | | | |
| 5: return i^* | | | | |
| | | | | |
| | | | | |

6 EVALUATION

6.1 Experimental Methodology

6.1.1 Implementation Details

We develop a trace-driven simulator to implement and evaluate our fairness metric and ASTRAEA system. We adopt several techniques to optimize the system implementation. First, to select a tenant at Phase 1, we check the gap between the resources a tenant can utilize in a scheduling round and his weighted quota. Specifically, we slightly modify Equation (8) for the tenant-level fairness as below:

$$FairShare_{j}^{tenant}(t) = Quota_{j}(t).$$
(16)

This can better reveal the discrimination of utilized resources with different resource weights of the tenants. The corresponding fairness metric (Equation (12)) illustrates the weighted amount of resources allocated to tenant p during $[t_1, t_2]$. This adjusted metric can motivate a tenant to *lend* his unused resources to others and reclaim more than his share when he has a burst demand in the future, which can achieve the sharing incentive more effectively. Moreover, although all tenants may remain fair on a long-time scale, it is possible that a burst request from a tenant may prevent him from submitting new jobs as long as he has a higher LTGF metric. Therefore, this modification can compensate for the problem that LTGF cannot be quickly adjusted on short-time scales.

Second, to select a job at Phase 2, we consider more joblevel prioritization, as discussed in Section 5.2.1. We prioritize the newly submitted job, which can help short jobs be scheduled and completed promptly. This achieves overall better results considering the fact that short jobs take the majority in GPU clusters.

6.1.2 Workloads

For the fidelity and generality of our simulation, we select two real-world DLT job traces and use a two-week snapshot of both two traces. The first one is Venus, collected from a production cluster of 1064 GPUs in SenseTime [20]. This trace contains 11304 jobs from 16 different tenants. The second one is Philly from the Microsoft cluster [19], containing 44329 jobs from 15 VCs with 2490 GPUs. Note that the Philly trace does not provide the size of each VC cluster, so we make an assumption that the size of a VC is proportional to the total number of GPUs requested by all the jobs in this VC.

6.1.3 Simulation Configurations

We set the lease term length as 900s, which can best balance the cluster efficiency and fairness. Identification of this value is detailed in Section 6.2. The scheduling interval is set as 10s to provide balanced response time for newly submitted jobs. These configurations are compatible with the default settings of real-world round-based schedulers [35]. Considering there exist CPU jobs in Venus and general under-utilization for production clusters, we follow [36] to scale down the cluster size to make usage more intensive: we set 100 nodes for Venus and 210 nodes for Philly, with each node comprised of 8 GPUs. We will fully investigate the effectiveness of ASTRAEA under different cluster sizes in Section 6.5. The duration and GPU demand of each job is unmodified in the workload, regardless of whether the submission time has been normalized. The cluster is initialized with all the nodes available and starts to schedule jobs submitted in the last scheduling round. We apply a consolidated placement strategy in the simulation.

6.1.4 Metrics

We use the following metrics to quantify the effectiveness of our proposed Astraea.

- *Fairness*. We measure the LTGF fairness at both the tenant and job levels (Section 4.2). For tenant-fairness, we calculate ρ^{tenant} at multiple fixed intervals for each tenant (Equation (12)), as well as the distribution. For job-fairness, we calculate ρ^{job} for every job (Equation (15)).
- Average slowdown. We use the average slowdown to quantify the normalized pending situation, which is calculated as the sum of pending time and running time divided by the running time (Section 3.2). This can represent users' experience: a small average slowdown indicates higher fairness and shorter waiting time, which is beneficial especially for short-term jobs.
- *Job Completion Time (JCT)*. We compute the average JCT for the cluster efficiency. An efficient cluster has better scheduling performance with lower average JCT.

6.1.5 Comparison With Other Baselines

We consider the following six baselines for comparisons. The first four are mainstream DLT job schedulers for different



Fig. 9. Fairness affected by different lease terms.

purposes (e.g., improving job performance and resource utilization, handling heterogeneous environments). The last two are state-of-the-art DLT schedulers specifically for joblevel fairness and tenant-level fairness, respectively.

- *YARN-CS* [4]: a static quota strategy is adopted to allocate resources according to users' weights.
- *HiveD* [11]: a dynamic quota is implemented in this scheduler, allowing users to submit spot jobs to other VCs with lower utilization.
- *Allox* [37]: this scheduler adopts the min-cost bipartite matching to solve the placement problem with heterogeneous resources, and select tenants with the lowest progress in each scheduling round.
- *Tiresias-L* [7]: it measures the aggregated GPU time each job receives, and uses the Least Attained Service to maintain the resource allocation.
- *Themis* [8]: it proposes the long-term finish-time fairness as a new metric to evaluate fairness. We implement this scheme by calculating the remaining iterations and predicting throughput with non-shared situation, with a reasonable estimation error rate.
- Gandiva_{fair} [13]: this scheduler utilizes a gangaware split stride, ticket-based mechanism to manage the resources in heterogeneous systems, focusing on the tenant-level fairness.²

6.2 Selection of Lease Term Length

Fairness and cluster efficiency highly depend on the lease term length. We evaluate the impact of lease term lengths on the scheduling system. Fig. 9 shows the average JCT (doted line) and preemption overhead (bars) with different lease term lengths. We observe the preemption overhead is relatively small, especially when the lease term becomes longer. When we increase the lease term length, the JCT is first decreased, as more jobs can be completed within one term without being preempted, and the frequency of job switching is reduced with smaller overhead. However, when the lease term becomes longer, the JCT also becomes larger because jobs have fewer chances to be adjusted. We can find a length that achieves strong fairness while not sacrificing the job performance. Fig. 9 also shows the relationship between the job-level fairness and lease term length.

^{2.} *Gandiva_{fair}* also introduced an automated trading mechanism to handle GPU heterogeneity and time sharing. This is not the focus of this paper and will not be included in the evaluation. This mechanism can be integrated with ASTRAEA as well.

Authorized licensed use limited to: Nanyang Technological University. Downloaded on June 13,2022 at 02:15:35 UTC from IEEE Xplore. Restrictions apply.



Fig. 10. Cumulative distributions of job-level fairness.

We observe that a longer term can reduce the fairness, as the frequency of job execution rearrangement is reduced either, making it less efficient to adjust the fairness via scheduling. Based on the above analysis, we select a proper length as 900s from our historical trace considering both fairness and efficiency, because it has similar performance while satisfactory fairness compared with the most efficient length (1200s). In practice, the lease term length could be adjusted by cluster administrators. We will use this configuration for the following evacuations. Scheduler designers can identify the lease term length for their clusters and jobs similarly.

6.3 Fairness Comparisons

Fig. 10 shows the cumulative distribution of job-level fairness under different schedulers. We observe more jobs maintain higher fairness with ASTRAEA than the other schedulers. If we treat jobs with fairness metric smaller than 0.95 as *sharing loss* jobs, then ASTRAEA only produces 7.1% *sharing loss* jobs in Venus. In contrast, Themis and YARN-CS schedulers give 19.7% and 73.6% *sharing loss* jobs, $2.8\times$ and $10.3\times$ of ASTRAEA. ASTRAEA exhibits even higher advantage in Philly, which has more uneven distributions of resources and jobs among tenants.

Figs. 11 shows the tenant-level fairness metrics with different schedulers for the Venus trace. The x-axis in each figure denotes tenants, and *y*-axis shows the distribution of tenantlevel fairness metric for each tenant over many periods. A fair scheduler should maintain a higher metric all the time. We observe that ASTRAEA can guarantee most tenants have the fairness metric larger than or equal to 1. The other schedulers give relatively worse tenant-level fairness, indicating they fail to meet the resource-as-you-contribute requirement and sharing incentive. Quantitatively, we calculate the ratio of cases whose metric value is smaller than 1 as the indicator that the tenant is experiencing sharing loss. ASTRAEA gives such an unfairness ratio of 5.2%. For the other schedulers, the best one is Allox (6.9%), similar with $Gandiva_{fair}$ (8.0%), and the worst are Tiresias-L (49.0%) and YARN-CS (44.6%). ASTRAEA outperforms these schedulers by $1.3 \times$ to $9.42 \times$. Similar conclusions can be drawn from the results for Philly trace.

6.4 Cluster Efficiency Comparisons

We show how effective our scheduler is in mitigating the queuing congestion and reducing the average JCT in the cluster. Fig. 12 shows the cumulative distributions of average slowdown. We observe that ASTRAEA can achieve the least pending overhead compared to other schedulers, indicating that it can provide prompt responses to short-term



Fig. 11. Distributions of tenant-level fairness in Venus.

jobs and reduce the waiting time of long-term jobs. A very small portion of jobs suffer from long pending overhead, because they have very short running time (e.g., 5 seconds) compared to one scheduling round. A shorter scheduling round can decrease these jobs' pending time at the cost of larger scheduling overhead.

Fig. 13 shows the average JCT of each scheduler. The fairness enforcement in ASTRAEA can ensure work conservation for short-time jobs and reduce the pending time for long-time jobs. It also enables the flexible adjustment of tenant quota. Due to these mechanisms, ASTRAEA has smaller average JCT compared to other schedulers.

Similar as other preemption-based schedulers, ASTRAEA introduces extra overhead when rearranging the job execution order. This overhead can be divided into two



Fig. 12. Cumulative distribution of average slowdown.

components: making checkpoints, and cold-start job. Checkpoints are made for every lease expiration in case of preemption, while cold start occurs only when the job lease renewal fails. Previous work has shown that cold start dominates the job preemption overhead[7]. Our lease-based training scheme enables the job to renew its lease and continue the execution without the cold start, which can improve the JCT. In general, this scheme only brings approximately 0.8% overhead to the average JCT with a lease term length of 900s (Fig. 9) in the Venus trace, which is negligible.

6.5 Impact of Cluster Sizes

The cluster loads can also affect the effectiveness of ASTRAEA and other scheduling algorithms. We quantitatively show the impact of cluster sizes on the average JCT and ratio of *sharing loss* jobs with a fairness metric < 0.95. We compare ASTRAEA with Tiresias-L, which outperforms other baseline methods according to our evaluations in previous sections. Figs. 14a and 14b give the comparison results with different cluster sizes (i.e., numbers of GPU nodes). We observe that ASTRAEA always beats Tiresias-L for both cluster efficiency and fairness, regardless of the cluster size. In addition, both ASTRAEA and Tiresias-L exhibit more advantages when the cluster becomes larger. This is because a larger cluster with more available resources can satisfy more jobs with reduced pending time, and these idle resources can be better utilized to achieve fairness for various jobs.

7 RELATED WORKS

7.1 Fairness Metrics

We review and analyze the fairness metrics for conventional CPU workload or DLT jobs in prior works. Some metrics consider the relative job performance (e.g., job execution time) between the shared and independent systems [8], [14].



Fig. 13. Average JCT for different schedulers.



Fig. 14. Impact of cluster sizes on ASTRAEA and Tiresias-L.

Some methods calculate the benefit gap between different tenants: a smaller gap indicates a fairer scheduler [38], [39]. Some methods utilize queuing theory to consider the job experience in the pending queue [40], [41]. A variety of works also proposed fairness solutions for OS processes [13], [42], [43], [44] or distributed workload scheduling [45], [46], [47], [48]. They cannot be applied to DLT job scheduling due to the unique job characteristics.

7.2 DLT Job Management

Over the past years, researchers have designed a quantity of scheduling algorithms and systems to optimize the execution of DLT jobs in GPU clusters from different perspectives. (1) Some works designed schedulers to maximize the cluster utilization [6], [12], [19], [49], [50]. Particularly, GPU sharing across jobs is a facilitating mechanism[51], [52] to improve GPU utilization, which is adopted in these works. (2) Some works aim to minimize the average JCT of DLT jobs [7], [31], [37], [53], [54]. (3) New schedulers were introduced to optimize the job performance [6], [7], [8], [11], [13], [19], [31], [37], [55]. Different from those works, our goal is to enhance both the job-level and tenant-level fairness while maintaining the job performance.

8 CONCLUSION

In this paper, we present a novel study about resource allocation fairness for DLT jobs in GPU clusters. We perform a quantitative analysis of a real-world job trace from a production cluster in SenseTime, to uncover the severe fairness problem in DLT scheduling. To mitigate this issue, we propose a new yet practical metric, LTGF, to accurately measure the fairness at both the job and tenant levels, without prediction of job remaining time and future throughput. We further design ASTRAEA, a new and efficient scheduler with a two-phased scheduling algorithm based on LTGF to enforce fairness. We evaluate ASTRAEA via large-scale simulations on real-world cluster traces. Experimental results show that ASTRAEA can guarantee stronger fairness without sacrificing the job performance or cluster utilization, compared to other state-of-the-art schedulers.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their valuable comments.

REFERENCES

 Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.

Authorized licensed use limited to: Nanyang Technological University. Downloaded on June 13,2022 at 02:15:35 UTC from IEEE Xplore. Restrictions apply.

- [2] Y. You, Z. Zhang, C.-J. Hsieh, J. Demmel, and K. Keutzer, "Imagenet training in minutes," in *Proc. 47th Int. Conf. Parallel Process.*, 2018, pp. 1–10.
- pp. 1–10.
 [3] P. Sun, Y. Wen, R. Han, W. Feng, and S. Yan, "Gradientflow: Optimizing network performance for large-scale distributed DNN training," *IEEE Trans. Big Data*, to be published, doi: 10.1109/TBDATA.2019.2957478.
- [4] V. K. Vavilapalli et al., "Apache hadoop yarn: Yet another resource negotiator," in Proc. 4th Annu. Symp. Cloud Comput., 2013, pp. 1–16.
- [5] B. Hindman *et al.*, "Mesos: A platform for fine-grained resource sharing in the data center," in *Proc. 8th USENIX Conf. Netw. Syst. Des. Implementation*, 2011, pp. 295–308.
- [6] W. Xiao et al., "Gandiva: Introspective cluster scheduling for deep learning," in Proc. 13th USENIX Conf. Oper. Syst. Des. Implementation, 2018, pp. 595–610.
- [7] J. Gu et al., "Tiresias: A GPU cluster manager for distributed deep learning," in Proc. 16th USENIX Conf. Netw. Syst. Des. Implementation, 2019, pp. 485–500.
- [8] K. Mahajan et al., "Themis: Fair and efficient GPU cluster scheduling," in Proc. 17th USENIX Symp. Netw. Syst. Des. Implementation, 2020, pp. 289–304.
- [9] A. B. Yoo, M. A. Jette, and M. Grondona, "Slurm: Simple Linux utility for resource management," in *Proc. Workshop Job Scheduling Strategies Parallel Process.*, 2003, pp. 44–60.
- [10] B. Burns, B. Grant, D. Oppenheimer, E. Brewer, and J. Wilkes, "Borg, omega, and kubernetes," *Commun. ACM*, vol. 59, no. 5, pp. 50–57, 2016.
- [11] H. Zhao *et al.*, "Hived: Sharing a GPU cluster for deep learning with guarantees," in *Proc. 14th USENIX Conf. Oper. Syst. Des. Implementation*, 2020, pp. 515–532.
 [12] A. Qiao *et al.*, "Pollux: Co-adaptive cluster scheduling for good-
- [12] A. Qiao et al., "Pollux: Co-adaptive cluster scheduling for goodput-optimized deep learning," in Proc. 15th USENIX Symp. Oper. Syst. Des. Implementation, 2021, pp. 1–18.
- [13] S. Chaudhary, R. Ramjee, M. Sivathanu, N. Kwatra, and S. Viswanatha, "Balancing efficiency and fairness in heterogeneous GPU clusters for deep learning," in *Proc. 15th Eur. Conf. Comput. Syst.*, 2020, pp. 1–16.
- [14] M. Isard, V. Prabhakaran, J. Currey, U. Wieder, K. Talwar, and A. Goldberg, "Quincy: Fair scheduling for distributed computing clusters," in *Proc. ACM SIGOPS 22nd Symp. Oper. Syst. Princ.*, 2009, pp. 261–276.
- [15] M. Zaharia, "Job scheduling with the fair and capacity schedulers," Hadoop Summit, vol. 9, Jun. 2009. [Online]. Available: http://people. csail.mit.edu/matei/talks/2009/hadoop_summit_fair_scheduler.pdf
- [16] A. Ghodsi, M. Zaharia, S. Shenker, and I. Stoica, "Choosy: Maxmin fair sharing for datacenter jobs with constraints," in *Proc. 8th* ACM Eur. Conf. Comput. Syst., 2013, pp. 365–378.
- [17] M. Abadi et al., "Tensorflow: A system for large-scale machine learning," in Proc. 12th USENIX Conf. Oper. Syst. Des. Implementation, 2016, pp. 265–283.
- [18] A. Paszke et al., "Pytorch: An imperative style, high-performance deep learning library," in Proc. 33rd Conf. Neural Inf. Process. Syst., 2019, pp. 1–12.
- [19] M. Jeon, S. Venkataraman, A. Phanishayee, J. Qian, W. Xiao, and F. Yang, "Analysis of large-scale multi-tenant GPU clusters for DNN training workloads," in *Proc. USENIX Conf. USENIX Annu. Tech. Conf.*, 2019, pp. 947–960.
- [20] Q. Hu, P. Sun, S. Yan, Y. Wen, and T. Zhang, "Characterization and prediction of deep learning workloads in large-scale gpu datacenters," in *Proc. Int. Conf. High Perform. Comput., Netw. Stor*age Anal., 2021, pp. 1–15.
- age Anal., 2021, pp. 1–15.
 [21] Q. Zhang et al., "Retiarii: A deep learning exploratory-training framework," in Proc. 14th USENIX Conf. Oper. Syst. Des. Implementation, 2020, pp. 919–936.
- [22] U. Misra *et al.*, "Rubberband: Cloud-based hyperparameter tuning," in *Proc. 16th Eur. Conf. Comput. Syst.*, 2021, pp. 327–342.
- [23] L. Xie et al., "Elan: Towards generic and efficient elastic training for deep learning," in Proc. IEEE 40th Int. Conf. Distrib. Comput. Syst., 2020, pp. 78–88.
- [24] D. G. Feitelson, "A survey of scheduling in multiprogrammed parallel systems," IBM T. J. Watson Res. Center, Oct. 1994. [Online]. Available: https://dominoweb.draco.res.ibm.com/ 052aab6422ad431f852565930064f913.html
- [25] M. Schwarzkopf, A. Konwinski, M. Abd-El-Malek , and J. Wilkes, "Omega: Flexible, scalable schedulers for large compute clusters," in *Proc. 8th ACM Eur. Conf. Comput. Syst.*, 2013, pp. 351–364.

- [26] M. Zaharia, D. Borthakur, J. Sen Sarma , K. Elmeleegy, S. Shenker, and I. Stoica, "Delay scheduling: A simple technique for achieving locality and fairness in cluster scheduling," in *Proc. 5th Eur. Conf. Comput. Syst.*, 2010, pp. 265–278.
- [27] W. Gao, Z. Ye, P. Sun, Y. Wen, and T. Zhang, "Chronus: A novel deadline-aware scheduler for deep learning training jobs," in *Proc. ACM Symp. Cloud Comput.*, 2021, pp. 609–623.
- [28] D. Narayanan et al., "Solving large-scale granular resource allocation problems efficiently with pop," in Proc. ACM SIGOPS 28th Symp. Oper. Syst. Princ., 2021, pp. 521–537.
- [29] J. W. Park, A. Tumanov, A. Jiang, M. A. Kozuch, and G. R. Ganger, "3sigma: Distribution-based cluster scheduling for runtime uncertainty," in *Proc. 13th EuroSys Conf.*, 2018, pp. 1–17.
- [30] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, "Dominant resource fairness: Fair allocation of multiple resource types," in *Proc. 8th USENIX Conf. Netw. Syst. Des. Implementation* 2011, pp. 323–336.
- [31] Y. Peng, Y. Bao, Y. Chen, C. Wu, and C. Guo, "Optimus: An efficient dynamic resource scheduler for deep learning clusters," in *Proc. 13th EuroSys Conf.*, 2018, pp. 1–14.
- [32] D. Banerjee, L. Fernandes, V. Vallabhaneni, and V. Jain, "Method, system and article for advance lease negotiation in DHCP," U.S. Patent App. 11/034,274, Jul. 13, 2006.
- [33] A. Adya, J. Dunagan, and A. Wolman, "Centrifuge: Integrated lease management and partitioning for cloud services," in *Proc. 7th* USENIX Symp. Netw. Syst. Des. Implementation, 2010, pp. 1–16. [Online]. Available: https://www.usenix.org/conference/nsdi10-0/centrifuge-integrated-lease-management-and-partitioningcloud-services
- [34] P. Li et al., "Beating OPT with statistical clairvoyance and variable size caching," in Proc. 24th Int. Conf. Architectural Support Program. Lang. Oper. Syst., 2019, pp. 243–256.
- [35] Kubernetes contributors, 2021. Accessed: Dec. 25, 2021. [Online]. Available: https://kubernetes.io/
- [36] A. Tumanov, T. Zhu, J. W. Park, M. A. Kozuch, M. Harchol-Balter, and G. R. Ganger, "Tetrisched: Global rescheduling with adaptive plan-ahead in dynamic heterogeneous clusters," in *Proc. 11th Eur. Conf. Comput. Syst.*, 2016, pp. 1–16.
- [37] T. N. Le, X. Sun, M. Chowdhury, and Z. Liu, "Allox: Compute allocation in hybrid clusters," in *Proc. 15th Eur. Conf. Comput. Syst.*, 2020, pp. 1–16.
- [38] R. Grandl, M. Chowdhury, A. Akella, and G. Ananthanarayanan, "Altruistic scheduling in multi-resource clusters," in *Proc. 12th USENIX Conf. Oper. Syst. Des. Implementation*, 2016, pp. 65–80.
- [39] D. Raz, H. Levy, and B. Avi-Itzhak, "A resource-allocation queueing fairness measure," ACM SIGMETRICS Perf. Eval. Rev., vol. 32, no. 1, pp. 130–141, 2004.
- [40] G. Sabin, G. Kochhar, and P. Sadayappan, "Job fairness in nonpreemptive job scheduling," in Proc. Int. Conf. Parallel Process., 2004, pp. 186–194.
- [41] G. Sabin and P. Sadayappan, "Unfairness metrics for space-sharing parallel job schedulers," in *Proc. Workshop Job Scheduling Strategies Parallel Process.*, 2005, pp. 238–256.
 [42] C. A. Waldspurger and W. E. Weihl, "Lottery scheduling: Flexible Description of the process."
- [42] C. A. Waldspurger and W. E. Weihl, "Lottery scheduling: Flexible proportional-share resource management," in *Proc. 1st USENIX Conf. Operating Syst. Des. Implementation*, 1994, pp. 1–es.
 [43] D. H. Epema, "An analysis of decay-usage scheduling in multi-
- [43] D. H. Epema, "An analysis of decay-usage scheduling in multiprocessors," ACM SIGMETRICS Perform. Eval. Rev., vol. 23, no. 1, pp. 74–85, 1995.
- [44] D. H. J. Epema, "Decay-usage scheduling in multiprocessors," ACM Trans. Comput. Syst., vol. 16, no. 4, pp. 367–415, 1998.
- [45] B. Chun *et al.*, "Planetlab: An overlay testbed for broad-coverage services," ACM SIGCOMM Comput. Commun. Rev., vol. 33, no. 3, pp. 3–12, 2003.
- [46] A. Mu'alem and D. Feitelson, "Utilization, predictability, workloads, and user runtime estimates in scheduling the IBM SP2 with backfilling," *IEEE Trans. Parallel Distrib. Syst.*, vol. 12, no. 6, pp. 529–543, Jun. 2001.
- [47] E. Shmueli and D. G. Feitelson, "Backfilling with lookahead to optimize the packing of parallel jobs," J. Parallel Distrib. Comput., vol. 65, no. 9, pp. 1090–1107, 2005.
- [48] S. Tang, B.-s. Lee, B. He, and H. Liu, "Long-term resource fairness: Towards economic fairness on pay-as-you-use computing systems," in *Proc. 28th ACM Int. Conf. Supercomput.*, 2014, pp. 251–260.

- [49] P. Yu and M. Chowdhury, "Fine-grained GPU sharing primitives for deep learning applications," in Proc. Mach. Learn. Syst., vol. 2, 2020, pp. 98-111
- [50] W. Xiao et al., "Antman: Dynamic scaling on GPU clusters for deep learning," in Proc. 14th USENIX Conf. Operating Syst. Des. Implementation, 2020, pp. 533-548.
- [51] L. Shi, H. Chen, J. Sun, and K. Li, "VCUDA: Gpu-accelerated highperformance computing in virtual machines," IEEE Trans. Comput., vol. 61, no. 6, pp. 804–816, Jun. 2012.
- [52] A. J. Pena, C. Reaño, F. Silla, R. Mayo, E. S. Quintana-Ortí, and J. Duato, "A complete and efficient cuda-sharing solution for hpc clusters," *Parallel Comput.*, vol. 40, no. 10, pp. 574–588, 2014.
- [53] S. Venkataraman, Z. Yang, M. Franklin, B. Recht, and I. Stoica, "Ernest: Efficient performance prediction for large-scale advanced analytics," in Proc. 13th USENIX Conf. Netw. Syst. Des. Implementation, 2016, pp. 363-378.
- [54] H. Zhang, L. Stafman, A. Or, and M. J. Freedman, "SLAQ: Quality-driven scheduling for distributed machine learning," in Proc. Symp. Cloud Comput., 2017, pp. 390-404.
- [55] D. Narayanan, K. Santhanam, F. Kazhamiaka, A. Phanishayee, and M. Zaharia, "Heterogeneity-aware cluster scheduling policies for deep learning workloads," in Proc. 14th USENIX Conf. Oper. Syst. Des. Implementation, 2020, pp. 481-498.



Zhisheng Ye received the BS degree in 2019 in computer science and technology from Peking University, China, where he is currently working toward the PhD degree with the School of Computer Science. His research interests include distributed systems, systems for machine learning, and resource management.



Peng Sun received the PhD degree in computer science from Nanyang Technological University, Singapore. He is currently a senior research scientist with SenseTime Group Limited. He was a research engineer with Nanyang Technological University, Baidu Institue of Deep Learning and Huawei 2012 Labs. His research interests include cloud computing, computer networking, data center, big data, and large-scale cluster computing systems for machine learning.



Wei Gao received the BS degree from Beihang University, Beijing, China, in 2019. He is currently working toward the PhD degree with Nanyang Technological University, Singapore. His research interests include distributed machine learning system, cluster resource management, and workload schedulinlg.



Tianwei Zhang (Member, IEEE) received the bachelor's degree from Peking University in 2011 and the PhD degree from Princeton University in 2017. He is currently an assistant professor with the School of Computer Science and Engineering, Nanyang Technological University. His research interests include computer system security, security threats and defenses in machine learning systems, autonomous systems, computer architecture, and distributed systems.



Xiaolin Wang (Member, IEEE) received the BS and PhD degrees from Peking University in 1996 and 2001 respectively. He is currently a full professor of computer science with Peking University. His research interests include system software, virtualization technologies, and distributed computing.



Shengen Yan received the BS degree from the Harbin Institute of Technology, Harbin, China, in 2009 and the PhD degree from the Institute of Software, Chinese Academy of Sciences, Beijing, China, in 2014. From June 2013 to February 2014, he was a visiting student with NC State University, Raleigh, North Carolina. From 2015 to 2017, he was a postdoctoral researcher with Multimedia Lab, Chinese University of Hong Kong, Hong Kong. He is currently the executive research director of Data and Computing Platform Department, Sense-

Time. He has authored or coauthored about 30 papers in the area of parallel computing and deep learning. His research interests interests include large scale deep learning and high-performance computing.



Yingwei Luo (Member, IEEE) received the BS degree from Zhejiang University in 1993, and the MS and PhD degrees from Peking University in 1996 and 1999, respectively. He is currently a full professor of computer science with Peking University. His research interests include system software, virtualization technologies, and distributed computing.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.