

Design and Operation of Shared Machine Learning Clusters on Campus

Kaiqiang Xu
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
kxuar@cse.ust.hk

Decang Sun
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
dsunak@cse.ust.hk

Hao Wang
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
hwangdv@connect.ust.hk

Zhenghang Ren
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
zrenak@cse.ust.hk

Xinchen Wan
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
xinchen.wan@connect.ust.hk

Xudong Liao
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
xliaoaf@connect.ust.hk

Zilong Wang
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
zwangfb@connect.ust.hk

Junxue Zhang
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
jzhangcs@connect.ust.hk

Kai Chen
Hong Kong University of Science and
Technology
Hong Kong, Hong Kong
kaichen@cse.ust.hk

Abstract

The rapid advancement of large machine learning (ML) models has driven universities worldwide to invest heavily in GPU clusters. Effectively sharing these resources among multiple users is essential for maximizing both utilization and accessibility. However, managing shared GPU clusters presents significant challenges, ranging from system configuration to fair resource allocation among users.

This paper introduces SING, a full-stack solution tailored to simplify shared GPU cluster management. Aimed at addressing the pressing need for efficient resource sharing with limited staffing, SING enhances operational efficiency by reducing maintenance costs and optimizing resource utilization. We provide a comprehensive overview of its four extensible architectural layers, explore the features of each layer, and share insights from real-world deployment, including usage patterns and incident management strategies.

As part of our commitment to advancing shared ML cluster management, we open-source SING's resources to support the development and operation of similar systems.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ASPLOS '25, March 30-April 3, 2025, Rotterdam, Netherlands

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0698-1/25/03

<https://doi.org/10.1145/3669940.3707266>

CCS Concepts: • Computer systems organization → Distributed architectures.

Keywords: Shared GPU Cluster; Multi-tenant Cluster Operations; Resource Management

ACM Reference Format:

Kaiqiang Xu, Decang Sun, Hao Wang, Zhenghang Ren, Xinchen Wan, Xudong Liao, Zilong Wang, Junxue Zhang, and Kai Chen. 2025. Design and Operation of Shared Machine Learning Clusters on Campus. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1 (ASPLOS '25), March 30-April 3, 2025, Rotterdam, Netherlands*. ACM, New York, NY, USA, 16 pages. <https://doi.org/10.1145/3669940.3707266>

1 Introduction

In the age of large machine learning (ML) models [6, 38], there has been a significant increase in the need for GPU clusters. This surge in demand has been met with substantial research funding and human capital investments in universities throughout the world [28]. However, not all campus infrastructures are equal, and many universities do not have the expertise to make the best use of shared GPU clusters. In most cases, campus IT (or graduate students) pick a tool that they are familiar with (e.g., Slurm [41]) and use it as-is with sub-optimal settings. Consequently, they end up underutilizing the already limited resources, putting a significant barrier to the democratization of access to ML.

Effective management of a shared GPU cluster requires addressing a hierarchical set of system challenges, including cluster management frameworks [41, 50], scheduling strategies [34, 47], network transport protocols [1, 52], topology

design [7, 27], and other system configurations [3, 26]. A brief online search yields numerous inquiries [32, 36, 37] for advice on managing shared GPU clusters. Our observations indicate that fellow researchers invest significant effort in exploring software solutions, and frequently struggle with improving the clusters' *usability*, *stability*, and *performance* to meet expectations.

Admittedly, a wealth of information is available online offering recommendations [16, 40] for managing shared clusters. However, effectively leveraging these solutions for a research-focused, shared GPU cluster demands deep expertise and continuous operational effort for evaluation, deployment, and maintenance.

Consider multi-user resource allocation as an example: A simple and common approach to GPU sharing is to provide direct shell access to users, which can result in system instability arising from users concurrently running computing processes, competing for memory resources, potentially disrupting each other's runtime dependencies, and requiring manual coordination and careful cleanup. On the flip side, Kubernetes [50], a widely adopted cluster manager, is known to be effort-demanding in deployment. Originally designed for internet backend services, Kubernetes features a complex architecture with service-hosting capabilities like load balancing and rollbacks but lacks native support for user-access control and job queuing for resource sharing. These complexities create engineering challenges and overhead when applied to ML model training jobs in campus ML clusters. Details of existing solutions are discussed in §2.

In this paper, we present our experience from the *design* and *operation* of SING, a shared GPU cluster tailored for ML research in academic institutions.

Design: SING streamlines ML job processing into four system layers: job profile, adapter, scheduling, and execution (§3). Each layer utilizes a low-maintenance software stack that efficiently handles tasks like job scheduling and user isolation, while minimizing burdens in operation. We employ effective resource allocation strategies to prevent low resource utilization and improve fairness during allocation.

Operation: Launched in early 2021, SING has become a cornerstone of our institution's AI research infrastructure, managing over 160 GPUs and serving over 480 active users. These users have submitted over 28,000 ML jobs, which generated experiment results that contribute to about 40 peer-reviewed ML papers. We present a comprehensive analysis of SING's usage patterns, job characteristics, and our incident handling experiences (§5). We will also release job trace data to facilitate further research in this area (§6).

At its core, SING places a strong emphasis on simplicity and stability when it comes to selecting software stacks. In line with the constraints of limited manpower within research teams, our approach to designing and choosing SING's components is guided by the principle of achieving *minimal*

viability while still fulfilling the essential requirements for running ML jobs. To facilitate this design goal, SING adopts a 4-layer architectural abstraction, connecting its components that handle the compilation, scheduling, and execution of user-submitted jobs within the cluster. The detailed exploration of SING's interface, internal workflow, and system configurations can be found in §3.

SING offers a user-friendly interface from the local command line, allowing users to submit jobs with code, configurations, and specified resources. This workflow is also referred to as Machine-Learning-as-a-Service (MLaaS) [55]. For maximum usability, SING supports two job formats simultaneously: scripting (submitting a script, which is widely used with HPC clusters and adopted into ML computing) and containerizing (submitting a Dockerfile [10] describing a container that runs the job). Similar to cloud services providers, SING automates resource provisioning for job execution, but with additional ML-specific features including dependency setup, job queuing and scheduling, port forwarding, interactive debugging, and access to logs and outputs.

SING is designed with extensibility, allowing cluster operators to employ alternative options in the layers of SING to achieve new capabilities. Beyond the default software selections that prioritize architectural simplicity and operational efficiency, we also discuss scaling opportunities such as expanding the executor backend to cloud platforms.

During developing and operating SING, balancing user-friendliness with operational efficiency was challenging. We refined our design choices by monitoring user interactions and analyzing anonymous usage data, while addressing performance bottlenecks and resolving instances of failures and incidents. We delve into these operational insights, focusing on usage patterns, job characteristics, and incident handling and prevention in §5.

This paper contributes to the field in two main aspects:

- **Operation-Efficient Design (§3):** SING stands out as an accessible solution that manages a shared GPU cluster and achieves operational efficiency (i.e., low maintenance cost and high resource utilization), specifically addressing the challenges posed by limited manpower in research teams. We describe seven core features of SING's design that contribute to its operational efficiency.
- **Operational Insights (§5):** We track and analyze usage patterns and job characteristics and also share our experiences in incident handling during the operation of SING. These data and experiences offer valuable insights into the design and operation of similar campus ML clusters.

This paper represents our ongoing effort aimed at improving the management of shared ML clusters in academic institutions. We make the source code, software stack configurations, and SING job traces (§6) available to assist in deploying and operating similar facilities.

Table 1. Comparison of existing solutions for shared ML clusters. The analysis considers each solution’s core features without extensive customization. References to #1-7 correspond to SING’s features described in §3.

Feature Category	Direct Shell Access	Traditional HPC	Container Platforms	SING
Resource Management				
Job Resource Access	SSH & Linux Permissions	Login Node Submission	Deploy Container via UI/CLI	User Local Submission (#2)
Resource Scheduling	Manual Coordination	FCFS, Queuing	Quota-based	Optimized FCFS (#4)
Job Execution & Monitoring				
Work Description	Manual Command	Job Script	Container File	Multi-Format Support (#1) & Early-Init & Caching (#5)
Dependency Setup	Manual Command	Script Execution	Container Initialization	
Job Monitoring	Server Login	Offline Debug with Output Log		Interactive Monitoring (#3)
User Experience and Complexity				
User Learning Curve	Simple - Basic Shell	Moderate - Job Script	Complex - Container Skills	Moderate - Intuitive Workflow (#6)
Operator Setup Cost	Low - Basic Linux	Moderate - Install Software	High - Complex Architecture	Low to Moderate Cost (§4) & with Executor Extensibility (#7)
Maintenance Cost	High - Stability / Conflicts	Moderate - Stability Issues	High - Limited ML Job Support	

2 Background and Motivations

The growing demand for GPUs has led to significant financial and human resource investments. However, effectively managing shared GPU clusters presents complex challenges, including resource and environment provisioning, usability and stability issues, and resource allocation efficiency.

2.1 Existing Solutions for Shared ML Clusters

In this section, we analyze the workflows of three common solutions for managing shared ML clusters. In Table 1, we summarize their advantages and disadvantages.

2.1.1 Direct Shell Access Sharing. A common approach is to grant users native shell access to GPU-equipped machines through SSH, allowing direct interaction with GPU resources. However, due to the lack of resource isolation and management mechanisms, this method raises concerns about system stability and security. Without careful coordination, it can result in conflicts between competing user sessions, degrading system performance and making it difficult to ensure fair resource allocation.

Users typically follow the following workflow:

Step 1. User Connection: Users connect to GPU nodes via SSH, gaining direct access to GPU resources.

Step 2. Manual Environment Setup: Users are responsible for manually configuring their runtime environments, including installing necessary libraries and dependencies.

Step 3. Resource Allocation: Users must manually coordinate GPU resources through other channels, to avoid potential conflicts or performance degradation.

Step 4. Job Execution and Cleanup: Users initiate ML tasks by running their scripts directly on GPU nodes, requiring them to manually manage job execution and clean up the environment after the job finishes.

Analysis. While this approach offers simplicity in terms of user access, it presents significant drawbacks.

First, users are required to manually configure their runtime environments, including the installation of necessary

libraries and dependencies. This manual setup process can be error-prone and time-consuming, adding to the operational burden and potentially risking system stability.

Furthermore, users must proactively take responsibility for coordinating GPU sharing, initiating ML tasks, and cleaning up their environment afterward. It can potentially introduce conflicts between concurrent user sessions, which not only degrades system performance but also makes it challenging to ensure fair resource allocation.

2.1.2 Traditional HPC Solution. Universities often adopt cluster workload management solutions from traditional High-Performance Computing (HPC) clusters, with the Slurm workload manager [41] being a popular choice. Slurm provides job queueing and scheduling capabilities in clusters.

In Slurm, users follow a more structured workflow:

Step 1. Job Submission: Users submit ML job requests from a login node to the Slurm, with resource request specified.

Step 2. Job Queueing: Slurm is responsible for job queueing and scheduling, using built-in schedulers.

Step 3. Runtime Environment: Users generally need to set up the runtime environments within the job script, as there is a lack of support for managing ML dependencies in Slurm.

Step 4. Execution: The system bootstraps and executes the ML job on allocated nodes, and releases resources afterwards.

Analysis. Slurm lacks support for provisioning the complex dependency environments required by ML jobs, posing engineering difficulties to individual users. Job preparation and submission are required to be done on the login nodes of a Slurm cluster, and the login node often becomes a single point of failure [15, 46], due to unexpected user actions on the node (see detailed discussion in §3.1). Furthermore, Slurm’s coarse-grained resource allocation mechanisms may lead to resource fragmentations and underutilization under certain submission orders and job characteristics [11].

Opportunity. Originally tailored for MPI applications, Slurm stood out for its simplicity and performance. ML training jobs in research often resemble MPI applications, using advanced MPI techniques [25, 43]. This similarity brings ML

job scheduling within Slurm’s scope. In this paper, our solution involves augmenting Slurm with ML-specific features that transform it into an effective tool for shared ML clusters.

2.1.3 Container Platforms. Container orchestration platforms such as Kubernetes [50] have garnered widespread adoption in the industry for managing large-scale distributed applications. Recently, Kubernetes has also been used to manage shared ML clusters, especially in production clusters with user-facing ML applications.

In containerized solutions such as Kubernetes, users follow a workflow centered around container orchestration:

Step 1. Containerization: Users package their ML program and runtime dependencies into containers by writing a description file in a specific format.

Step 2. Container Deployment: Users submit container images on the Kubernetes cluster, specifying resource requests.

Step 3. Job Scheduling: Kubernetes schedules pods on available nodes, managing resource allocation.

Step 4. Job Execution: Containers are initiated to execute ML tasks, with Kubernetes handling creation and releasing.

Analysis. Designed for internet services, Kubernetes has a complex architecture with features like service discovery, load balancing, and feature rollouts and rollbacks. In the context of ML research, these complexities are unnecessary and introduce engineering overhead.

More importantly, Kubernetes lacks native support for user-level access control and job queuing essential for shared ML clusters, each requiring a separate set of tools and configurations [30] to be added on top of Kubernetes. Therefore, deploying and maintaining a Kubernetes setup demands considerable effort from cluster operators, including integration, testing, and ongoing maintenance.

Meanwhile, for users *without prior experience*, creating containers can be error-prone [39] with Dockerfile, and replicating a cluster environment locally for error reproduction is challenging and often requires operator assistance.

Overall, both users and operators face a steep learning curve with Kubernetes, requiring skill development to manage their inherent complexities effectively [30].

Opportunity. Container platforms like Kubernetes are powerful server operation tools, and SING’s design does not conflict with them. SING uses an adapter layer that can operate over any execution layer, including container platforms or computing clouds, allowing resource extensibility while remaining transparent to the users.

2.2 Cluster Schedulers and Computing Frameworks

This section provides a brief overview of related systems that also operate in cluster environments. A more detailed discussion can be found in §8.

ML Cluster Schedulers. State-of-the-art schedulers, such as Sia [47], Pollux [34], and Gavel [24], optimize scheduling

strategies to enhance cluster efficiency, such as reducing average job completion time. These schedulers can be integrated into cluster managers like Slurm, Kubernetes, and SING, which manage the entire ML job lifecycle.

Computing Frameworks. Distributed computing frameworks such as Ray [23] and Spark [59] parallelize large jobs for efficient distributed execution. These distributed systems are programming frameworks used directly in user code, which can be submitted as a part of the user’s job to cluster managers. These systems also support an optional “cluster mode” similar to a cluster manager and we discuss this in §8.

2.3 Motivations and Requirements

Current ML cluster management solutions reveal limitations in usability and operational efficiency in multi-tenant environments. SING’s design is motivated by the need for a simple, stable, and efficient solution for managing shared ML clusters in research institutions.

An effective solution should address three objectives:

1. **Reduce Operational Complexity:** Leverage a stable and low-maintenance architecture to reduce operational workload for cluster operators compared to existing solutions.
2. **Improve User Accessibility:** Provide a clear and accessible interface that simplifies ML job submission and monitoring, reducing the learning curve for users.
3. **Enable Fair and Efficient Resource Allocation:** Ensure fair resource allocation to maximize utilization while maintaining system responsiveness.

SING addresses these requirements with a full-stack architecture designed for ML workflows, introducing features such as multi-format job support, optimized first-come-first-serve (FCFS) scheduling, and environment caching. SING aims to create a more stable and user-friendly environment for shared ML cluster operations.

3 Design and Implementation

SING’s core architecture design follows a four-layer abstraction to define, construct, and execute user-submitted ML jobs within a shared ML cluster. As illustrated in Figure 1, SING’s design is structured as follows:

The *Job Profile Layer* allows users to define job and environment requirements with a self-contained profile syntax, ensuring unambiguous execution of jobs.

The *Adapter Layer* is the interface between users and the cluster, facilitating job submission and monitoring.

The *Scheduler Layer* manages resource allocation and job scheduling, employing an optimized FCFS with a backfill mechanism to prioritize and utilize resources efficiently.

The *Executor Layer* is responsible for executing ML jobs on cluster nodes, featuring network file systems and customizable execution backends for scalability.

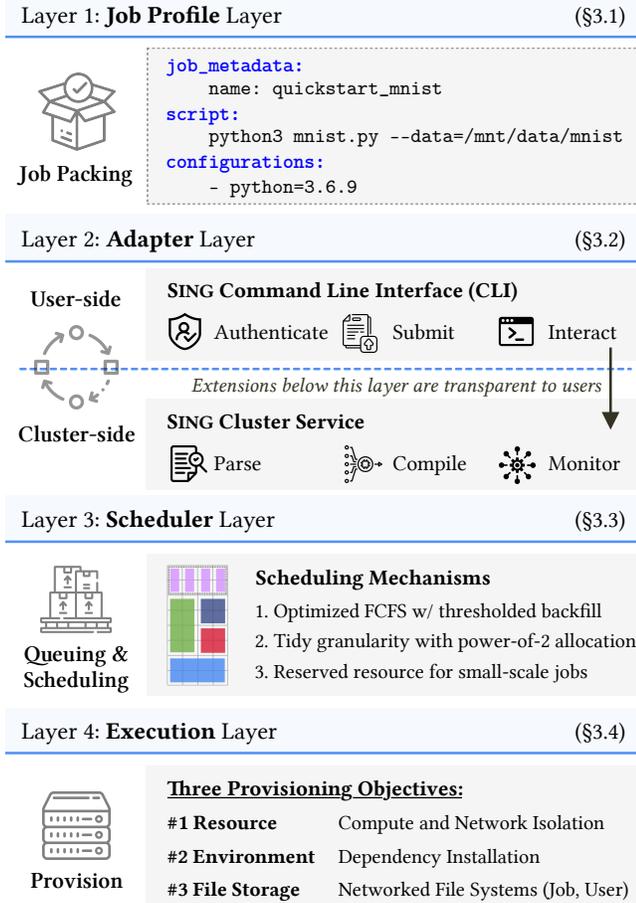


Figure 1. SING uses 4-layer architecture to enhance the overall efficiency and flexibility in operation.

This abstraction offers a key advantage by decoupling the ML job execution workflow to independent input and output in each layer, creating a versatile design framework where different components can be combined to accomplish the cluster management objective.

3.1 Layer 1: Job Profile Layer

Jobs submitted to SING are defined using job profiles, and an example showing the syntax of these profiles is in Figure 2. These profiles consist of the following sections:

1. Metadata and Resource Request: users specify resource requirements, including the selected resource group, number of nodes (which can be a range), the number of GPUs required on each node, and the format of the job code.
2. Job Code: this section points to the entry point of the user’s job code, which can be a script or a container file.
3. Configurations: users specify the job’s configurations, such as runtime libraries like CUDA version and programming dependencies like Python packages.

Feature #1: Job Format Flexibility. SING supports script-format (common in HPC) and container-format job profiles simultaneously. Users can specify their preferred format by indicating it in the *job* section and providing either an execution script or a container file (Dockerfile) in the *entrypoint* section of the job profile. This versatility in format minimizes migration costs during job submissions across different computing environments. Subsequently, the job profile is compiled into executable scripts by the adapter layer (§3.2) before being submitted to the scheduler layer (§3.3) for execution.

Comparing to traditional HPC clusters. Having a profile layer eliminates the need for setting up a server from where users submit jobs via Slurm toolchains. In traditional HPC clusters, this server often called a *login node*, is a common single point of failure. Some users might unintentionally execute their job scripts directly on the login node, either for debugging purposes or due to workflow misunderstandings. Running scripts locally on the login node can fully occupy its computing or memory resources, rendering it unresponsive and preventing other users from submitting new jobs. This can disrupt the system and potentially lead to a cluster-wide outage. Despite warnings from HPC cluster operation teams [32, 46], such incidents still happen as they cannot be completely prevented [15], necessitating manual recovery by the operation team which can be time-consuming.

Another advantage of using the job profile is its self-contained and repeatable nature. Users can save and reuse their job profiles for recurring tasks, simplifying job submission. Moreover, these profiles can be shared with others to replicate experiments, fostering collaboration and reproducibility within the research community.

Job Profile Example: quickstart_mnist.conf

```

job_metadata:
  name: quickstart_minst
  resource: sing_gpu_1 # Select target resource group.
  nodes: 4 # Peer IPs provided during runtime.
  gpu: 2 # Number of GPUs requested per node.
  format: script # Job format: script or container.

script: quickstart.sh # The submitted script
container: quickstart.dockerfile # or, a container file.

configurations: # For script-format only.
  - python=3.6.9 # Dependencies are configured
  - cudatoolkit=11.1.0 # through Conda and bash scripts.
  - pytorch=1.9.0

```

Figure 2. Example job profile where users specify job requirements, code entry point, and configurations.

3.2 Layer 2: Adapter Layer

The adapter layer takes the user’s job profile from the previous layer and submits it to the ML cluster. It also offers

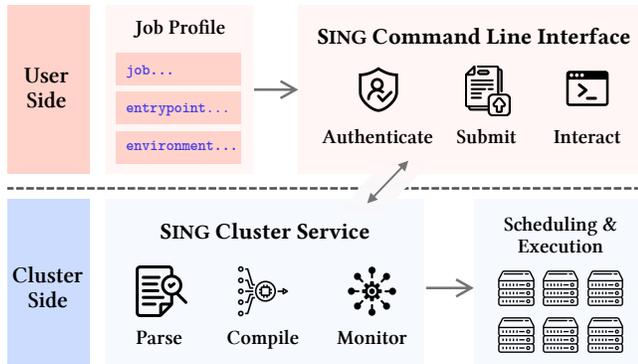


Figure 3. Workflow and functions of the adapter layer.

user interfaces for functions like monitoring running jobs, retrieving output files, and streaming logs.

The adapter layer is referred to as such because it is programmed against the job schedulers in the layer below. It abstracts the complexity of the cluster manager’s functionalities and provides a user-friendly interface for job submission and monitoring.

Our implementation, illustrated in Figure 3, comprises two components: SING-cli and SING-service, running locally on user computers and on the cluster, respectively.

SING-cli (SING Command Line Interface), is a local software distributed to SING users. It stores cluster configurations, such as the cluster’s IP and login credentials, and connects with the cluster. Beyond authentication, the *cli* acts as a thin layer that wraps around the services provided by SING-service, providing commands that facilitate job submission and access to other SING-service functionalities.

SING-service operates on a cluster node and receives instructions from the *cli*. When a job is submitted, the service compiles the job files by parsing the job profile and converting it into the scheduler’s recognized format. This process includes mapping configuration parameters and packaging code files before passing them to the scheduler layer. SING-service also provides job monitoring functions, such as fetching job status and accessing output files and logs.

Connection Security. SING-cli uses remote command execution over an SSH tunnel to invoke functions in SING-service. To authenticate users, their public key—created during registration—is added to the dedicated cluster node that runs SING-service. This setup does not allow users to log in directly to the cluster node, thereby reducing the risk of unexpected user actions on the node that could disrupt operations (as discussed in §3.1). A more secure approach that offers better isolation between the client and cluster is to implement SING-service functions as HTTP APIs, connected over HTTPS with password authentication. However, we opt for remote command execution as it offers more flexibility (e.g., file uploads) with less engineering effort. We discuss potential security concerns related to this solution in §7.

We discuss two key features of the adapter layer and provide details of our implementation as follows.

Feature #2: Multi-format Job Compilation. Job compilation in SING enables jobs in various formats to be transformed into a standardized output script that the backend can interpret and execute. Since this compilation process is handled on the cluster side (Figure 3), the same SING-cli can connect to different clusters running SING-service without requiring user code modifications. This design supports extensibility across various execution backends. While our SING implementation uses Slurm [41] as the execution backend, community extensions (§6) have enabled integration with alternative backends—or even multiple backends simultaneously based on a routing strategy—without adding usability burdens on the user side.

For *script-format jobs*, the SING-service identifies required dependencies and injects scripts to prepare the environment. This, for example, includes setting the `LD_LIBRARY_PATH` to switch between pre-installed CUDA versions [58] on cluster nodes and initializing Conda [3] for Python dependencies. We optimize Conda initialization by caching environments using MD5 hashes of sorted dependency lists, enabling instant retrieval for recurring jobs with identical requirements.

For *container-format jobs*, which are self-contained packages, the job is passed to the scheduler layer for resource allocation. After that, the compute node invokes container runtimes (such as Docker or Containerd [8]) to execute the provided container file directly. Optionally, script-format jobs can also be encapsulated into a container file to simplify execution, albeit with an added layer of abstraction.

After compilation, the job is submitted to the specified resource group and cluster. This layer provides extensibility by allowing the SING-cli to connect to SING-service on different computation clusters (varying in location, speed, or cost) *without* user code modification. Additionally, the SING-service compiles the job profile into different output scripts suitable for its respective cluster.

Feature #3: Interactive Job Monitoring. Many research projects running on our cluster involve ongoing development and exploratory experiments. Users therefore need to interact with their jobs in real time, especially when debugging or monitoring the training process. Hence, SING provides the following interactive features.

1. *Remote port forwarding.* For real-time debugging tools such as Jupyter Notebook or TensorBoard, SING-service dynamically edits port forwarding rules on the cluster gateway server, which runs HAProxy [17]. This setup exposes the container’s port through a randomly assigned port on the gateway. Users can then access the service using the gateway’s public IP and the forwarded port. Once the job is stopped, the executor backend automatically triggers [42] the removal of the port forwarding rule. Note that this approach opens a public port on the gateway server, which

could allow unauthorized access if the service behind the forwarded port lacks authentication. We discuss this security risk and a potential solution in §7.

2. *Remote SSH access.* Users can gain SSH access to the remote host (either a server or a container) where their job is running to perform debugging tasks. To facilitate this connection, SING-service generates a one-time SSH key pair and adds the public key to the authorized keys file on the remote host. The service then exposes the SSH port of the remote host on the gateway server using the previously mentioned port forwarding function. The private key is securely transmitted back to SING-cli via the cli-to-service connection (§3.2). SING-cli then use this private key to establish an SSH connection to the remote host through the gateway. After the job stops, both the key pair and the port forwarding rule are removed.

3. *Retrieving log and output files.* SING-cli can download or stream remote files (similar to `tail -follow`) in the job’s runtime directory to access output or log.

In practice, to prevent low resource utilization resulting from remote job debugging, we restrict the usage of remote functions to jobs that require two or fewer GPUs. This encourages users to initially test their jobs with fewer resources and scale up once the testing phase is complete.

3.3 Layer 3: Scheduler Layer

The scheduling layer is responsible for making decisions on the allocation of resources. ML cluster scheduling is a popular research area. However, as many approaches improve cluster efficiency by evaluating jobs with performance factors such as estimated job completion time, resource utilization, and statistical efficiency, they may not be suitable for shared ML clusters in research institutions because:

- They determine the order of job execution based on optimization algorithms targeting cluster-wide metrics. This approach may sometimes result in unfair resource allocation from the perspective of individual users.
- They often require modifications to user code, including model hyperparameters like batch size, to exploit optimization opportunities. In the context of ML research, such modifications are less desirable, if not unacceptable, as echoed in recent literature [18].

SING incorporates the following mechanisms to improve cluster-wide resource utilization while being fair to individual users, as illustrated in Figure 4. Resource abuse prevention is discussed in §5.3.

Feature #4: Fair FCFS with Backfill. SING’s scheduling approach is *analogous* to the way universities manage their public facilities, recognizing that ML clusters function as university-wide shared resources. Fairness in this context means that resource requests are generally processed on a first-come, first-served (FCFS) basis to ensure equal access for all users. However, to enhance overall cluster efficiency,

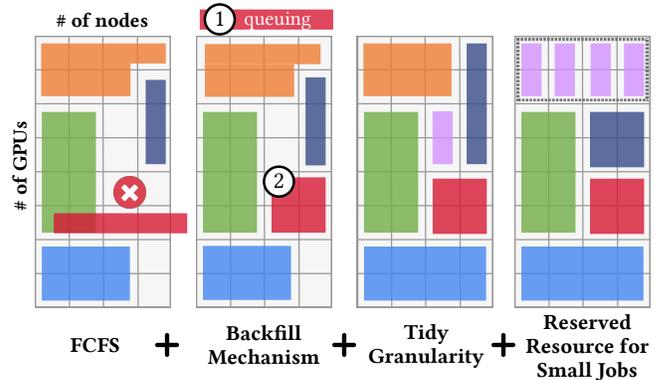


Figure 4. SING’s scheduler layer implements mechanisms to improve cluster-wide resource utilization.

certain requests may be given higher priority than earlier ones, as detailed below, in order to improve cluster-wide resource utilization.

SING uses a scheduling strategy that combines optimized FCFS with a backfill mechanism, as depicted in Figure 4. When a job is submitted, it is placed in a queue, awaiting execution. The scheduler engages in a bin-packing process to determine if there are sufficient available resources to start the next job. In cases where the available resources cannot satisfy the first job in the queue, the backfill mechanism will help avoid head-of-line blocking due to large jobs in the front of the queue: it temporarily skips the first job and evaluates the subsequent jobs to identify the ones that can be executed with available resources.

Next, we introduce strategies that further enhance cluster utilization by minimizing fragmentation and prioritizing smaller jobs. The usage statistics in §5 validate the effectiveness of these strategies.

Threshold for Backfilled Requests. In some cases, the backfill mechanism may lead to the starvation of the first job, especially if the first job request is large while the subsequent request is small. To address this, we restrict the total number of allocated GPUs for backfilled requests with a threshold (in SING, we set the threshold to the size of the first job request). Once this limit is reached, priority is given to the first job to make sure it will be scheduled next. Algorithm 1 illustrates the pseudocode of this scheduling strategy.

Tidy Granularity. SING accepts two-dimensional resource requests: the number of nodes and the number of GPUs per node. To prevent fragmentation in the bin-packing process, we have established a limited granularity for the quantity of GPUs in each node type, offering preset choices of 1/2/4/8-GPU nodes. Additionally, when users request multiple nodes, they must select nodes of the same type. These predefined granularities, when carefully scheduled with bin-packing algorithms [51], can enhance resource utilization as they minimize small fragments during scheduling. From a user

Algorithm 1: Optimized FCFS with Backfill

```

Input:
  • Job Queue  $Q$ , Backfill Threshold  $\mathcal{T}$ 
1 begin
  /* Decide the next job to run          */
2   $job = Q.peek()$ 
  /* Check the first job in queue       */
3  if  $available(job)$  then
4     $Q.pop()$ 
5     $\mathcal{T} = Q.peek().GPUs$ 
6    return  $job$ 
  /* Try backfilling jobs                */
7  for  $backfill\_job$  in  $Q$  do
8    if  $available(backfill\_job)$  and
       $\mathcal{T} \geq backfill\_job.GPUs$  then
9      /* Reduce backfill threshold    */
       $\mathcal{T} = \mathcal{T} - backfill\_job.GPUs$ 
10     return  $backfill\_job$ 
11 end
12 end

```

perspective, these choices of granularities align well with typical ML experiment settings (i.e., powers of two). The allocation of CPU and memory is proportional to the number of GPUs in each node type, which facilitates data processing during GPU training. This approach is generally sufficient based on our operational experience, as our GPU nodes are equipped with 40-core CPUs and 256GB of memory (which exceeds the total GPU memory on a node). For tasks that require only CPU or memory resources, we recommend using CPU clusters.

Reserved Resource for Small Jobs. Another effective strategy we implemented is the prioritization for smaller jobs (illustrated in Figure 4). We partition some 8-GPU nodes into 1/2/4-GPU nodes with isolated computation and network environments with SR-IOV [29], which allows a single physical NIC to appear as multiple virtual NICs with minimal performance loss [22]. These nodes are exclusively allocated to serve small requests. This strategy offers priority incentives that encourage users to request fewer resources when they only need to debug or quickly test for a few epochs, reducing the waiting time for other users. This optimization also implicitly resembles the effects of a short-job-first strategy, which reduces the waiting time for queueing jobs.

Feature #5: Early-Initialization and Caching. When jobs are queued and awaiting execution, SING accelerates the bootstrap stage by *pre-initializing* Conda environments (for script-format jobs) or building container images (for container-format jobs) as immutable directories stored in distributed

network storage. This network storage is accessible to all nodes (as described in Feature #6 in §3.4).

To further reduce job startup time, SING caches the dependency environment and container images locally on each node. When a job is started on compute nodes, SING first checks if the resource is available without accessing the nodes' network storage. If not, the environment is fetched to the local node for job execution and cached there for subsequent executions, unless purged due to local storage space constraints under the LRU cache replacement policy.

3.4 Layer 4: Execution Layer

The executor layer is where job profiles become instances and bootstrap for execution, and it provides necessary architecture support for parallel computation and resource isolations for multi-job processing. Figure 5 illustrates the topology and architecture of SING that supports distributed job execution. SING utilizes Slurm [41] as the default backend and this section discuss its capability and our enhancements to Slurm to make it a better fit for ML cluster management. §6 discusses a community work that utilizes Kubernetes for executor backend in SING's architecture. §8 discusses the reason why we choose Slurm as default backend.

Stage 1: Job Profile and Submission

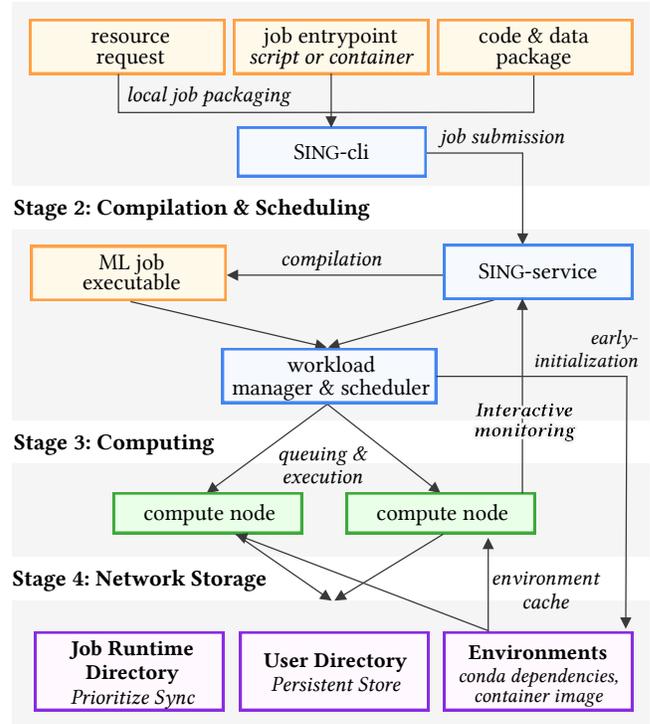


Figure 5. An end-to-end workflow of job processing in SING, from job submission to execution.

Feature #6: Distributed Job Bootstrapping. Each cluster node runs a Slurm agent that manages job processes. For parallel jobs spanning multiple nodes, an environment variable with assigned IPs facilitates peer connections. The agent executes the compiled job script from the job profile in the adapter layer (in §3.2), starting the initialized environment (Conda or container) and calling the job script within it. This environment is pre-initialized and cached in the network file system, as described in Feature #5 (§3.3).

We deploy a network file system with RDMA support using GlusterFS [13] for file synchronization across nodes and to prevent file duplication. There are two user-level file spaces: the job runtime directory and the user directory. Both directories are accessible during job runtime or using SING-cli but offer different Quality of Service (QoS) guarantees. The user directory provides persistent storage shared across multiple job submissions, prioritizing stable storage over synchronization speed. In contrast, the job runtime directory offers faster storage that only persists for a short period after job completion unless the user submits a new job. Otherwise, it is automatically purged.

Feature #7: Heterogeneous Executor Backends. SING’s executor backend, Slurm, operates under a “controller-agent” model. In this setup, the controller acts as a puppet master, managing resource groups, while agents run on compute nodes to execute jobs. The controller can partition resources into different groups (e.g., CPU-only, consumer GPU, data-center GPU), allowing users to specify the cluster and optionally the resource group to meet heterogeneous hardware requirements, while SING enforces user resource quotas.

Additionally, resources can be added or removed from the cluster simply by starting or stopping the agent. This model enables dynamic scaling of resources across the cluster, allowing it to adapt to varying workloads and user demands. Additionally, cloud-based compute nodes can be scaled through an agent running in the cloud, which connects back to the controller. We utilize a self-built cloud instance image (e.g., Amazon Machine Images [2]) that embeds the agent and starts it automatically at instance boot, to simplify the process. This approach ensures scalability and flexibility, adding additional resources to the cluster while remaining transparent to users. Future implementations may fully automate scaling operations based on cluster occupancy and cloud resource budgets.

3.5 SING Implementation

We implement the scheduling and execution layer of SING based on Slurm [41], which is extended to support the scheduling policies in §3.3. We selected Slurm for its popularity and simplicity in academic cluster management (further discussion in §8). We note that Kubernetes [50] can provide similar functionality and can also serve as the executor backend for SING, which has been developed in the Kubernetes community and is discussed in §6.

Table 2. Breakdown of estimated human efforts for building the SING solution in a new cluster.

Tasks / Steps	Est. Person-Days	
	Ground-up	Adopt
1. Basic OS and GPU Driver Configs	2	2
2. Cluster Softwares: Slurm and NFS	4	3
3. Integrate SING-cli and -service	8	2
4. Testing & Documentation	3	1

4 Usability and Accessibility

This section discusses the usability of SING’s design from the perspective of cluster operators and users, reflected by its operational workload and user experience.

4.1 Operation and Maintenance Cost

Table 2 compares the estimated human effort needed to build SING from scratch (17 person-days), or adopt the open-source SING configuration and code (8 person-days). The difference is mainly due to SING-cli and SING-service development and figuring out software configurations. See our open-source plan in §6.

Installing System Software. The effort required for the initial deployment varies depending on the operation team’s familiarity with Linux software configurations. For instance, setting up GlusterFS for the networked file system and Slurm for the scheduler is relatively straightforward if the team efficiently understands and follows their well-documented setup procedures. Meanwhile, they also have a large community where solutions to most deployment issues can be readily found.

Service Requests and Maintenance. The operation team’s day-to-day workload is primarily on responding to user requests and maintaining the cluster’s stability:

- To reduce basic service requests, our team provide clear job examples under various distribution settings and programming frameworks, helping users quickly get started with these templates. We also outlines detailed steps and best practices for job debugging. As a result, over the past 12 months, our cluster has seen an average of one service request for every 940 jobs submitted.
- System failures in SING are rare due to the minimalism in software architecture, which minimizes the risks of software failure. Hardware and networking issues, such as network file system delays, are discussed in §5.

4.2 User Experience

Our design strives to improve users experience for ML clusters, yielding several key advantages, as discussed below.

User-Friendly Features. SING-cli offers a suite of functions tailored to typical ML research workflows, allowing users to easily access remote resources. These functions include remote file uploads/downloads, attaching bash environment, and remote port forwarding.

Simplified Environment Management. Users no longer need to maintain two sets of scripts for local and remote executions. The job profile encapsulates all environment dependencies as well as scripts to start the ML process, and the steps to execute jobs locally and remotely are identical.

Documentation. The user documentation for SING offers a detailed collection of clear and intuitive examples for job submissions, ranging from basic workflow testing to hybrid parallel training. These examples help users to quickly understand the interface and functionality of SING, providing easy-to-follow solutions while reducing the operational team’s workload in addressing user inquiries.

Summary. Cluster operation includes both problem *prevention* and *resolution*. We prioritize *prevention* by adhering to design principles that reduce architectural complexity and use foolproof interfaces. For *resolution*, please see §5.3.

5 Operation and Analysis

We analyze SING’s performance and usage over 2.5 years to offer insights into academic ML cluster operation. The cluster consists of 128 Nvidia GTX 3090 GPUs on 16 nodes.

We first analyze cluster-wide trends including temporal patterns (§5.1) and job-specific characteristics (§5.2). Then, we discuss the experiences and actions taken to handle resource bottlenecks and hardware failures (§5.3). Finally, we present the schema of the SING cluster trace, a more comprehensive trace compared to prior work in §6.

5.1 Usage Statistics: Temporal Patterns

SING exhibited various temporal patterns during operation. These patterns are instrumental in planning for capacity and system maintenance schedules. Unless otherwise specified, the period under review in this section is for the 30 months from July 2021 to December 2023.

Number of Users. Figure 6 shows periodic surges in new user registrations for SING, where spikes are visible around September and February, aligning with the distribution of email newsletters at the beginning of each semester to introduce new researchers to the platforms. Outside these periods, user growth is steady and gradual, attributed to word-of-mouth referrals. Additionally, the data reflects an increasing trend in new user registrations over time, showing the growing popularity of the SING within our institution.

Task Submissions. Figure 7a shows the task submission statistics of SING, demonstrating a consistent growth in task submissions, with over 200 weekly average submissions. Notable submission spikes are observed around January 2022,

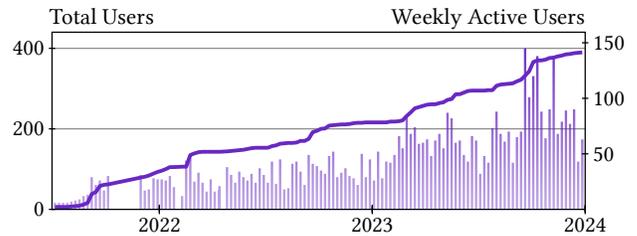
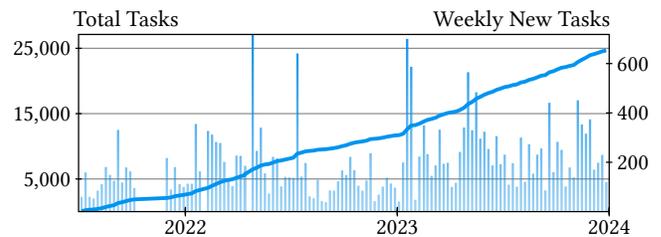
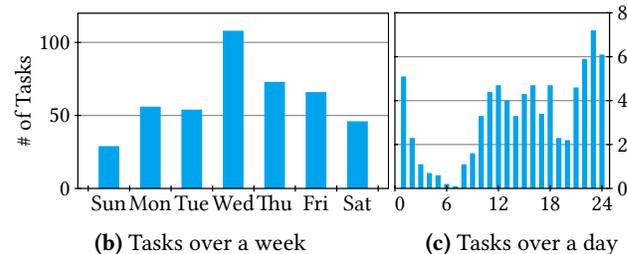


Figure 6. User growth and weekly active users.



(a) Number of total tasks and weekly new tasks from 2021 to 2023.



(b) Tasks over a week

(c) Tasks over a day

Figure 7. Task submissions over a year in monthly, daily, and hourly submission patterns.

late April to early May 2022, and late January 2023, possibly corresponding to periods with paper submission deadlines.

Figure 7b shows a mid-week peak in task submissions, particularly on Wednesdays, and gradually decreases towards the weekends, with the lowest on Sundays. Figure 7c illustrates a daily trend where submissions are minimal in the early morning, gradually increasing and peaking in late evenings, especially between 10pm through 1am the next day. This trend may reflect researchers’ late-night work habits and highlights the importance of a low-maintenance cluster, given the limited availability of operators during late hours.

GPU Occupancy Patterns. Figure 8 illustrates the GPU occupancy rate over a busy week, with the x-axis representing hours from 0 to 167, corresponding to Sunday 0am to Saturday 11pm. During the sampled period, SING experienced job oversubscription for about 30% of the time, with peaks lasting around two hours. We calculate the GPU assignment rate as the sum of assigned GPUs at each hour divided by the sum of the minimum value between the number of requested GPUs and the cluster capacity at each hour. This approach reflects utilization adjusted for demand. SING maintained a GPU assignment rate of over 96%, reflecting the effectiveness of our backfilling scheduling mechanism (§3.3). Less than 4%

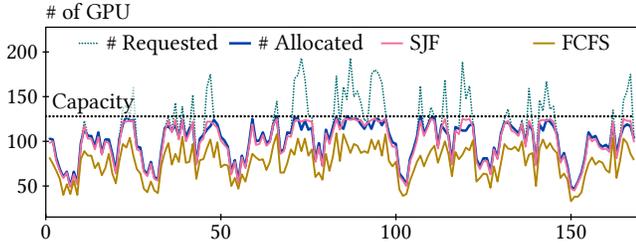


Figure 8. GPU occupancy over a week, from Sunday 0am to Saturday 11pm (168 hours).

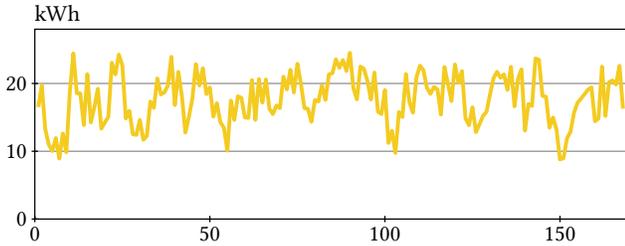


Figure 9. Cluster-wide power draw over the same time period as in Figure 8, showing the GPU power is dominant.

of GPUs are idle during the oversubscription period due to resource fragmentation, optimized by pre-defined granularity and the bin-packing algorithm (§3.3).

Comparing with Alternative Scheduling Policies. Using the job trace shown in Figure 8, we conducted a mini-evaluation of alternative scheduling policies to determine job priorities and compare their performance with SING. We replay the job trace in a simulator. The simulator functions as a step-based discrete-time system with a 1-minute granularity, resulting in 10,080 steps for this evaluation. Each step performs three tasks:

1. Calculate job progress to identify completions.
2. Check the job trace and add new arrivals to the queue.
3. Apply scheduling policy to start new jobs.

We compared two baseline policies: (1) the First-Come, First-Served (FCFS) policy, which is the default in Slurm, and (2) the Short-Job-First (SJF) policy used in Lucid [18], where we provided the exact running time to simulate an ideal scenario. Like in SING, this evaluation did not include job resource scaling or preemption. Additionally, all jobs were subject to a 24-hour rule (§5.3), consistent with the policy applied in SING from which the job trace was collected.

For the evaluation metrics, we examined job occupancy rates and queuing times. The GPU assignment times for SING, FCFS, and SJF were 96%, 75.1%, and 96.5%, respectively, as illustrated in Figure 8. The average queuing times for SING, FCFS, and SJF were 0.81, 1.32, and 0.64 hours, respectively. SING outperformed FCFS due to its backfill mechanism, while SJF performed better than SING in reducing waiting times (note that SING’s dedicated resource pool for small jobs

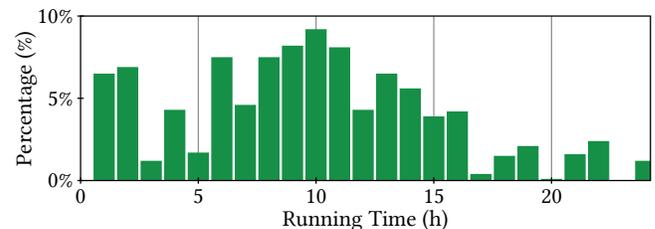
contributes to the low average queuing time). However, in terms of tail (maximum) queuing times, SING, FCFS, and SJF recorded 7.2, 7.4, and 10.2 hours, respectively, indicating that SJF may lead to starvation for longer jobs.

Power Consumption Metrics. Figure 9 illustrates the power consumption over the same time frame as the GPU occupancy graph. Max power draw is crucial for cluster planning. The primary reasons for the remaining differences are as follows: (1) the base power draw, which includes the idle power of the CPU, GPU, and other components, is relatively constant and independent of the GPU occupancy rate; and (2) the GPU utilization rate also influences power draw [57], meaning that two jobs occupying the same number of GPUs may exhibit different power consumption if one job is more computationally intensive than the other.

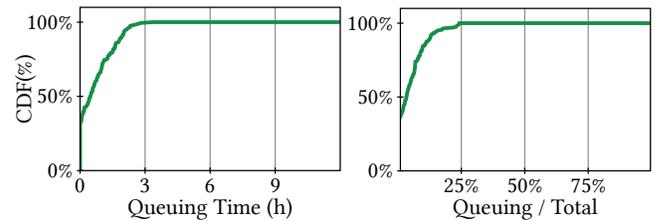
5.2 Usage Statistics: Job-level Characteristics

We present job-level characteristics of how long jobs run, how long they wait in the queue, and how efficiently they use resources, collected during the operation of SING.

Job Running-Time. Figure 10a illustrates the job durations within the GPU cluster. The majority of job lengths fall below 2 hours or between 6 to 11 hours. This pattern is a result of both user preferences and our scheduling approach. As discussed in §3.2, SING supports interactive usage of cluster resources, such as port forwarding and remote bash access, but restricts these features to smaller job requests. This encourages users to initially debug their jobs with smaller resource requests. During peak hours, SING automatically imposes a maximum runtime of 24 hours *per submission* to



(a) Distribution of job run-times in the GPU cluster, with significant clusters observed at 1-2 hours and 8-11 hours.



(b) Queuing-Time Distribution **(c)** Queuing / Run-Time %

Figure 10. Analysis of job run-times (a) and queuing times (b and c) in the GPU cluster, illustrating the distribution of job durations and queuing times.

ensure fair usage (§5.3). Users with longer-duration needs are advised to checkpoint their jobs around the 12-hour mark and submit a new job to continue after checkpointing. These mechanisms explain the prevalence of jobs running for less than 2 hours and between 6 to 11 hours.

Job Queuing-Time. Figures 10b and 10c displays the cumulative distribution function (CDF) of job queuing times in hours and the queuing time as a percentage of the total job time. These figures highlight that over 40% of jobs have no waiting time, and more than 90% of jobs wait less than 1 hour or 10% of their total runtime. These statistics are influenced by SING’s design. Short jobs receive prioritized allocation with dedicated resource and backfilling mechanism [54] to improve cluster-wide efficiency. Longer jobs are required to be checkpointed within 24 hours and resubmitted to prevent excessive occupation during peak hours.

Job Sizes Distribution. Figure 11 shows the top 10 job sizes by resource requests, with about 40% being small jobs requesting fewer than four GPUs. This is influenced by SING’s design, which favors smaller jobs by prioritizing their execution (§3.3) and limits interactive debugging tools (§3.2) to only small jobs. This approach motivates users to start with smaller resource requests for initial testing and debugging, enhancing overall resource utilization. Additionally, we make note that job sizes are restricted to powers of two (§3.3). This restriction aids in reducing resource fragmentation in the scheduling’s bin-packing process and does not conflict with typical ML research experiment settings.

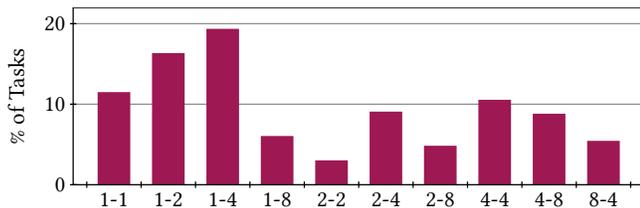


Figure 11. Distribution of job sizes, measured by the number of GPUs requested (e.g., 2-4 is two nodes with 4 GPUs each).

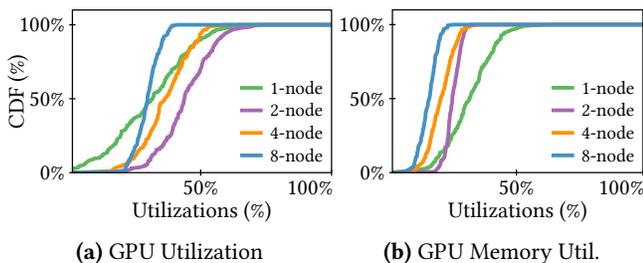


Figure 12. CDFs of job-level GPU and memory utilization, illustrating resource consumption across various job sizes (i.e., number of nodes requested).

GPU Utilization and Memory Utilization. Figure 12 illustrates GPU and GPU memory utilization based on the number of nodes requested by jobs. The utilization statistics only account for periods when a GPU is occupied. Generally, as the number of nodes in a job increases, its utilization decreases, suggesting possible synchronization overheads. For 8-node jobs, the average GPU utilization is 25% and the average GPU memory utilization is 15%. This may indicate that communication-efficient training methods may not be widely adopted in current ML experiments.

Table 3. Summary of issues and types of solutions.

Type	System Health Issues	Type of Solutions
Resource	Excessive Storage Usage	Policy Enforcement
Resource	Low GPU Utilization	User Guidance
Resource	Long Job Running Time	Policy Enforcement
Failure	Network File System Delay	Incidents Handling
Failure	System Configuration Lost	Incidents Handling

5.3 System Availability & Maintenance

This section shares our operational experiences with SING. In Table 3, we summarize the actions that address availability challenges into three types: user guidance, policy enforcement, and specific incident handling.

5.3.1 Resource Bottlenecks. Running into resource bottlenecks and limitations is the top feedback we hear from cluster users. We employ the following mechanism to promote fair usage of resources.

Efficient Storage Usage. Training ML models often involves accessing large model parameters and datasets. While we have over 1PB storage space available, the rapid accumulation of files from a large user base can exhaust our disk space. To ensure efficient storage utilization, we have implemented the following guidance and policies:

- **Public Dataset Directory.** In addition to individual user directories, we create a public directory where commonly used large datasets [20] are stored. Users can soft link these datasets directly to their directories, removing the need for duplication in downloads or copies. This approach not only prevents duplicate files but also saves download time and bandwidth.
- **File Expiration Policy.** Some users store large files in their directories that are not immediately necessary for their job. To free up storage space and promote responsible resource usage, we implement a file expiration policy: files that have not been accessed for more than 60 days are subject to deletion. To prevent accidental data loss, this expiration process is not real-time: We initiate the removal process at the beginning of each month and send email notifications seven days prior to deletion.

GPU Utilization Monitoring. We closely monitor GPU utilization in SING to identify cases of underutilized resources, which are often caused by non-computation-intensive tasks, inefficient job implementations, or interactive usage (only allowed on small nodes). Although we do not immediately intervene, we collect data and send email notifications to users with significantly low GPU utilization, and also offer help for job optimization when needed.

Job Running Time Limit. In SING, we limit job running time to 24 hours during peak times, defined as when there is a pending job queued for over an hour (other definitions may apply here). Users requiring longer runtimes should use checkpoints and resubmit their jobs. Experienced users can automate this process for continuous execution.

This policy is also *analogous* to the way universities manage their public resources, similar to borrowing a library book: you need to return it after a certain period, but you can re-borrow if it is not reserved by someone else. The policy is in place to prevent excessively long or unattended jobs and ensure fair allocation, particularly during peak periods.

5.3.2 Incidents Handling. We share lessons and recommendations from unexpected incidents.

Network File System Delay. We observed substantial delays and inconsistencies in the network file system during peak periods of concurrent remote I/O, impacting job initialization and causing runtime errors. We addressed this issue by switching to RDMA connections for the network file system. Users are also advised to utilize collective communications directly over the network and avoid file synchronization during job execution.

System Configuration Lost Due to Hardware Upgrade. While upgrading the physical network interface cards (NICs) in SING, we encountered an issue due to the tight coupling of the container’s PCI-E setting and specific NICs. When we switched NICs, it triggered a failure that rendered the container inoperable. We could not directly modify the PCI-E setting inside of the container to use the new NIC because the container was not able to start at all. In the end, we spent about 3 days (as per Table 2) recreating the environment.

6 SING Open Source

This paper represents our ongoing effort aimed at improving the management of shared ML clusters in academic institutions. Since SING is currently deployed in a production cluster, we are unable to fully explore all potential design choices and optimization strategies.

To facilitate the research and operations of similar infrastructures, we provide SING’s source code, configurations, and our cluster’s job traces as open-source resources [48]. Researchers interested in hands-on exploration of SING are welcome to access our on-premise deployment hosted at [49].

Job Trace. The SING trace offers more detailed job characteristics and runtime measurements compared to previous traces like Microsoft Philly [21] and Alibaba PAI [55].

- **Job Submission History (36 months):** ① timestamps at job submission, start, and completion; ② resource requests; ③ assigned node IDs; ④ *list of software dependencies* (indicative of the ML model/framework used). ⑤ *user name and job name* (anonymized through hashing, yet still useful for inferring relationships between jobs, such as resubmissions);
- **Node Status (aggregated for every 5/10/30 minutes):** ① node ID; ② average/maximum CPU and host memory utilization rates; ③ average/maximum GPU and memory utilization rates; ④ *inbound/outbound traffic in bytes*; ⑤ *networked file storage IO in bytes*.
- **Node Specifications:** ① Node ID; ② GPU model; ③ number of GPUs; ④ total memory and other hardware specifications.

The *italicized items* are unique data points that are not available in previous traces.

Software Stack. Our software stack includes the source code and configuration scripts for the cluster workload manager, scheduler, and monitoring tools. This includes:

- SING-cli and SING-service (§3.2) code and makefiles.
- Scheduler implementation (§3.3) as a Slurm plugin.
- Compute node (§3.4) configs (e.g., SR-IOV, GlusterFS).

Executor Backend Extension in Kubernetes Community. SING’s execution layer is designed to be extensible, enabling support for new cluster managers, with Slurm as the default. Recently, developers from DaoCloud [9], a leading Kubernetes contributor and software vendor, submitted a merge request to SING’s open-source repository [12]. This update adds support for job execution via Kubernetes and ports the scheduler to use the Kubernetes API. The development was presented in a report at KubeCon 2024 [31], a Linux Foundation community event, as a step toward integrating cloud-native infrastructure with AI.

7 Limitations

SING’s design has the following known limitations:

Resource Isolation. Slurm isolates resources assigned to different jobs on the same host using Linux cgroups [5]. We currently depend on Slurm’s built-in mechanisms for isolating resources assigned to various jobs. For container-format jobs, the container runtime is also invoked by Slurm within this isolation mechanism. The performance and security implications—specifically, whether user resources are entirely separated—require further research.

Connection Between SING-cli and SING-service. Currently, SING-cli establishes a connection to SING-service via SSH to execute remote commands. The user’s SSH public

key is added to the host running SING-service, associated with a new Linux user created during registration. While this connection is secure, there remains a potential vulnerability where a malicious user could exploit it to execute unauthorized commands on the SING-service host. Although we have restricted user permissions on the host system to allow only specific operations (such as invoking SING-service), a more secure approach would be to expose SING-service functionalities through a secure API with authentications.

Port Forwarding for User Jobs. To enable users to access interactive services (e.g., TensorBoard) running within their jobs, we employ port forwarding on the cluster gateway (§3.2). This approach exposes a port on the SING-service host, which may allow unauthorized access if the service lacks authentication. A more secure alternative is utilizing SSH tunneling [45]. This method creates a secure tunnel between the job host and the user’s local machine via the gateway, restricting access exclusively to the user.

Cluster Failure. SING does not automatically recover from cluster failures, such as node failures. When a node fails, it and any associated user jobs become disconnected from the cluster, and require manual recovery by the cluster operator. Consequently, the job loses progress (up to the last checkpoint) and must be resubmitted.

8 Related Work and Discussion

We discuss the distributed frameworks and schedulers for ML workloads. SING can extend with these capabilities, leveraging advancements in storage, networking, and hardware for efficient, flexible cluster management.

Distributed Computing Frameworks. Distributed computing frameworks such as Ray [23] and Spark [59] perform computations at a granular level, allowing them to parallelize large jobs for efficient distributed execution. For example, Ray employs an actor and task-based model for more precise task management and asynchronous execution. These frameworks are integrated directly into user code, and submitted as part of the user’s job to cluster managers. We note that these frameworks can optionally operate in cluster mode [35, 44], functioning similarly to cluster managers by managing and allocating cluster resources to jobs. However, when running in cluster mode, submitted jobs need to use their computation APIs. This approach differs from traditional cluster managers like Slurm, Kubernetes, or SING, which allow users to submit code using any programming framework.

Cluster Schedulers. Cluster schedulers for ML, including the state-of-the-art Pollux [34] and Sia [47], focus on job scheduling to improve cluster-wide performance such as job completion time. These schedulers can be integrated into cluster managers like Slurm, Kubernetes, and SING to manage the entire ML job lifecycle. However, as discussed in our scheduler design in §3.3, SING generally processes

requests in a first-come, first-served (FCFS) manner to ensure equal access for any individual users.

ML System Optimizations. SING is extensible to leverage advancements in system optimizations for ML, such as faster distributed storage [33] and more efficient networking [4, 53, 61]. These optimizations typically operate at layers orthogonal to SING and can enhance ML job performance transparently. Additionally, SING’s scheduling policies can be extended to accommodate user-defined requirements or job profiling information. This will enable SING to more effectively optimize job placement across heterogeneous nodes, including those with specialized hardware (e.g., TPU [14] and FPGA [19, 60]), as well as in geographically distributed environments (§3.4).

Using Slurm as Executor Backend. Slurm, originally designed for MPI applications in scientific computing, is valued for its simplicity and performance. While the MPI model has gained less focus with the rise of big data and internet services, ML training jobs share system-level similarities with MPI, leveraging advanced MPI techniques [25, 43]. As described in this paper, by adding ML workflow features to Slurm, we leverage Slurm’s architecture and extend it to manage shared ML clusters. Kubernetes can also serve as the executor backend (§6).

9 Conclusion

This paper addresses the pressing challenges faced by universities worldwide in effectively managing shared GPU clusters for ML research. Our experience with the design and operation of SING, a campus ML cluster manager, offers an operation-efficient solution, reducing the deployment and maintenance burden on operations teams while achieving high resource utilization and user satisfaction.

Acknowledgments

We thank our shepherd, Malte Schwarzkopf, and the anonymous reviewers from ASPLOS 2025, as well as reviewers from our previous submissions, for their valuable feedback. We also thank the Kubernetes community contributors, Peter Pan (@panpan0000) and Xiao Zhang (@wawa0210) from DaoCloud [9], for their insightful discussions and valuable contributions to the implementation of SING.

SING was developed, deployed, and operated on the TACC cluster [49] at The Hong Kong University of Science and Technology. The statistics and analysis reported in this paper are also based on SING’s deployment and operation on TACC. An arXiv paper [56] detailing the TACC architecture, which includes a bottom-up approach to efficient AI computing with networking, frameworks, and algorithms, forms the foundation and inspiration of this work.

This work is supported in part by the Hong Kong RGC TRS T41-603/20R, GRF 16213621, NSFC 62062005, 62402407. Kai Chen is the corresponding author.

References

- [1] Mohammad Alizadeh, Albert G. Greenberg, David A. Maltz, Jitendra Padhye, Parveen Patel, Balaji Prabhakar, Sudipta Sengupta, and Murari Sridharan. 2010. Data center TCP (DCTCP). In *Proceedings of the ACM SIGCOMM 2010 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, New Delhi, India, August 30–September 3, 2010*. ACM, 63–74. <https://doi.org/10.1145/1851182.1851192>
- [2] Amazon Web Services. 2024. *Amazon Machine Images in Amazon EC2*. Retrieved December 13, 2024 from <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIs.html>
- [3] Anaconda. 2024. *User guide - conda documentation*. Retrieved December 13, 2024 from <https://docs.conda.io/projects/conda/en/latest/user-guide/index.html>
- [4] Wei Bai, Kai Chen, Hao Wang, Li Chen, Dongsu Han, and Chen Tian. 2015. Information-Agnostic Flow Scheduling for Commodity Data Centers. In *12th USENIX Symposium on Networked Systems Design and Implementation, NSDI 15, Oakland, CA, USA, May 4-6, 2015*. USENIX Association, 455–468. <https://www.usenix.org/conference/nsdi15/technical-sessions/presentation/bai>
- [5] Bill Brophy, Martin Perry, Moe Jette, Yiannis Georgiou, and Matthieu Hautreux. 2014. *Slurm Processes Isolation*. Retrieved December 13, 2024 from https://slurm.schedmd.com/SUG14/process_isolation.pdf
- [6] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*. <https://proceedings.neurips.cc/paper/2020/hash/1457c0d6bfc4967418bfb8ac142f64a-Abstract.html>
- [7] Kai Chen, Ankit Singla, Atul Singh, Kishore Ramachandran, Lei Xu, Yueping Zhang, Xitao Wen, and Yan Chen. 2012. OSA: An Optical Switching Architecture for Data Center Networks with Unprecedented Flexibility. In *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2012, San Jose, CA, USA, April 25-27, 2012*. USENIX Association, 239–252. https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/chen_kai
- [8] Containerd. 2024. *An industry-standard container runtime with an emphasis on simplicity, robustness and portability*. Retrieved December 13, 2024 from <https://containerd.io/>
- [9] DaoCloud. 2024. <https://github.com/daocloud>.
- [10] Docker. 2024. *Dockerfile overview*. Retrieved December 13, 2024 from <https://docs.docker.com/build/concepts/dockerfile/>
- [11] Paul Edmon. 2022. *Cluster Fragmentation*. Retrieved December 13, 2024 from <https://www.rc.fas.harvard.edu/blog/cluster-fragmentation/>
- [12] GitHub. 2024. *GitHub Pull Request: Submit kubernetes job without changing tuxiv.conf format*. Retrieved December 13, 2024 from <https://github.com/turingacloud/tcloud-sdk/pull/2>
- [13] GlusterFS. [n. d.]. *Gluster Documentation*. Retrieved December 13, 2024 from <https://docs.gluster.org/en/main/>
- [14] Google. [n. d.]. *Get started | Cloud TPU | Google Cloud*. Retrieved December 13, 2024 from <https://cloud.google.com/tpu/docs/quick-starts>
- [15] Google Groups. 2022. *What is an easy way to prevent users from running programs on the master/login node*. Retrieved December 13, 2024 from <https://groups.google.com/g/slurm-users/c/HKMTjPQN910>
- [16] Raz Haleva. 2021. *The Challenges of Sharing GPUs and How to Solve Them*. Retrieved December 13, 2024 from <https://developer.hpe.com/blog/the-challenges-of-sharing-gpus-and-how-to-solve-them/>
- [17] HAProxy. 2024. *HAProxy, The Reliable, High Performance TCP/HTTP Load Balancer*. Retrieved December 13, 2024 from <https://www.haproxy.org/>
- [18] Qinghao Hu, Meng Zhang, Peng Sun, Yonggang Wen, and Tianwei Zhang. 2023. Lucid: A Non-intrusive, Scalable and Interpretable Scheduler for Deep Learning Training Jobs. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2, ASPLOS 2023, Vancouver, BC, Canada, March 25-29, 2023*. ACM, 457–472. <https://doi.org/10.1145/3575693.3575705>
- [19] Qijing Huang, Dequan Wang, Zhen Dong, Yizhao Gao, Yaohui Cai, Tian Li, Bichen Wu, Kurt Keutzer, and John Wawrzynnek. 2021. CoDeNet: Efficient Deployment of Input-Adaptive Object Detection on Embedded FPGAs. In *FPGA '21: The 2021 ACM/SIGDA International Symposium on Field Programmable Gate Arrays, Virtual Event, USA, February 28–March 2, 2021*. ACM, 206–216. <https://doi.org/10.1145/3431920.3439295>
- [20] Hugging Face. 2024. *Datasets - Hugging Face*. Retrieved December 13, 2024 from https://huggingface.co/datasets?sort=most_rows
- [21] Myeongjae Jeon, Shivaram Venkataraman, Amar Phanishayee, Junjie Qian, Wencong Xiao, and Fan Yang. 2019. Analysis of Large-Scale Multi-Tenant GPU Clusters for DNN Training Workloads. In *Proceedings of the 2019 USENIX Annual Technical Conference, USENIX ATC 2019, Renton, WA, USA, July 10-12, 2019*. USENIX Association, 947–960. <https://www.usenix.org/conference/atc19/presentation/jeon>
- [22] Microsoft. 2024. *Overview of Single Root I/O Virtualization (SR-IOV)*. <https://learn.microsoft.com/en-us/windows-hardware/drivers/network/overview-of-single-root-i-o-virtualization--sr-iov->
- [23] Philipp Moritz, Robert Nishihara, Stephanie Wang, Alexey Tumanov, Richard Liaw, Eric Liang, Melih Elibol, Zongheng Yang, William Paul, Michael I. Jordan, and Ion Stoica. 2018. Ray: A Distributed Framework for Emerging AI Applications. In *13th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2018, Carlsbad, CA, USA, October 8-10, 2018*. USENIX Association, 561–577. <https://www.usenix.org/conference/osdi18/presentation/nishihara>
- [24] Deepak Narayanan, Keshav Santhanam, Fiodar Kazhemiaka, Amar Phanishayee, and Matei Zaharia. 2020. Heterogeneity-Aware Cluster Scheduling Policies for Deep Learning Workloads. In *14th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2020, Virtual Event, November 4-6, 2020*. USENIX Association, 481–498. <https://www.usenix.org/conference/osdi20/presentation/narayanan-deepak>
- [25] NVIDIA. 2020. *Overview of NCCL - NCCL documentation*. Retrieved December 13, 2024 from <https://docs.nvidia.com/deeplearning/nccl/user-guide/docs/overview.html>
- [26] NVIDIA. 2024. *About CUDA | NVIDIA Developer*. Retrieved December 13, 2024 from <https://developer.nvidia.com/about-cuda>
- [27] NVIDIA. 2024. *Data Center Networking Concepts*. Retrieved December 13, 2024 from <https://docs.nvidia.com/networking-ethernet-software/guides/EVPN-Network-Reference/Data-Center-Networking-Concepts/>
- [28] NVIDIA. 2024. *NVIDIA Solutions for Higher Education and Research*. Retrieved December 13, 2024 from <https://www.nvidia.com/en-us/industries/higher-education-research/>
- [29] NVIDIA. 2024. *Single Root IO Virtualization (SR-IOV) - NVIDIA Docs*. Retrieved December 13, 2024 from [https://docs.nvidia.com/networking/display/mlnxofedv24010331/single+root+io+virtualization+\(sr-iov\)](https://docs.nvidia.com/networking/display/mlnxofedv24010331/single+root+io+virtualization+(sr-iov))
- [30] Michael Packard, Joe Stubbs, Justin A. Drake, and Christian Garcia. 2021. Real-World, Self-Hosted Kubernetes Experience. In *PEARC '21: Practice and Experience in Advanced Research Computing, Boston, MA, USA, July 18-22, 2021*. ACM, 16:1–16:5. <https://doi.org/10.1145/3437359.3465603>
- [31] Peter Pan and Kaiqiang Xu. 2024. *Breaking Boundaries: TACC as an Unified Cloud-Native Infra for AI + HPC*. Retrieved December 13, 2024 from <https://www.classcentral.com/course/youtube-breaking-boundaries-tacc-as-an-unified-cloud-native-infra-for-ai-hpc-peter-pan-kaiqiang-xu-312983>

- [32] Princeton Research Computing. [n. d.]. *Top 10 Mistakes to Avoid on the Research Computing Clusters*. Retrieved December 13, 2024 from <https://researchcomputing.princeton.edu/get-started/mistakes-avoid>
- [33] Yingjin Qian, Marc-André Vef, Patrick Farrell, Andreas Dilger, Xi Li, Shuichi Ihara, Yinjin Fu, Wei Xue, and André Brinkmann. 2024. Combining Buffered I/O and Direct I/O in Distributed File Systems. In *22nd USENIX Conference on File and Storage Technologies, FAST 2024, Santa Clara, CA, USA, February 27-29, 2024*. USENIX Association, 17–33. <https://www.usenix.org/conference/fast24/presentation/qian>
- [34] 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*. USENIX Association. <https://www.usenix.org/conference/osdi21/presentation/qiao>
- [35] Ray. [n. d.]. *Ray Clusters Overview*. Retrieved December 13, 2024 from <https://docs.ray.io/en/latest/cluster/getting-started.html>
- [36] Reddit 2021. Research Group GPU Sharing: some open-source tool. Retrieved December 13, 2024 from https://www.reddit.com/r/MachineLearning/comments/k07jn9/d_research_group_gpu_sharing_some_opensource_tool/
- [37] Reddit 2022. *What is the best way to manage GPU server for multi-users?* Retrieved December 13, 2024 from https://www.reddit.com/r/MachineLearning/comments/ve987y/d_what_is_the_best_way_to_manage_gpu_server_for/
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-Resolution Image Synthesis with Latent Diffusion Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*. IEEE, 10674–10685. <https://doi.org/10.1109/CVPR52688.2022.01042>
- [39] Giovanni Rosa, Federico Zappone, Simone Scalabrino, and Rocco Oliveto. 2024. Fixing Dockerfile smells: an empirical study. *Empir. Softw. Eng.* 29, 5 (2024). <https://doi.org/10.1007/S10664-024-10471-7>
- [40] Run:ai 2024. *Simplify GPU Sharing in Multi-GPU Environments*. Retrieved December 13, 2024 from <https://www.run.ai/guides/multi-gpu/simplify-gpu-sharing-part-1>
- [41] SchedMD 2021. *Slurm Workload Manager - Overview*. Retrieved December 13, 2024 from <https://slurm.schedmd.com/overview.html>
- [42] SchedMD 2024. *Slurm Workload Manager - strigger*. Retrieved December 13, 2024 from <https://slurm.schedmd.com/strigger.html>
- [43] 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>
- [44] Spark. [n. d.]. *Spark - Cluster Mode Overview*. Retrieved December 13, 2024 from <https://spark.apache.org/docs/latest/cluster-overview.html>
- [45] SSH Academy [n. d.]. *SSH Tunneling: Client Command and Server Configuration*. Retrieved December 13, 2024 from <https://www.ssh.com/academy/ssh/tunneling-example>
- [46] Stanford Research Computing Center. 2024. *Running Jobs - Sherlock*. Retrieved December 13, 2024 from <https://www.sherlock.stanford.edu/docs/user-guide/running-jobs/>
- [47] Suhas Jayaram Subramanya, Daiyaan Arfeen, Shouxu Lin, Aurick Qiao, Zhihao Jia, and Gregory R. Ganger. 2023. Sia: Heterogeneity-aware, goodput-optimized ML-cluster scheduling. In *Proceedings of the 29th Symposium on Operating Systems Principles, SOSP 2023, Koblenz, Germany, October 23-26, 2023*. ACM, 642–657. <https://doi.org/10.1145/3600006.3613175>
- [48] TACC Authors. 2024. *Open Source Artifacts: TACC*. Retrieved December 13, 2024 from <https://tacc.ust.hk/#opensource>
- [49] TACC Authors. 2024. *Scalable AI Infrastructure Designed for Evolving Machine Learning Research*. Retrieved December 13, 2024 from <https://tacc.ust.hk/>
- [50] The Kubernetes Authors 2024. *Overview | Kubernetes*. Retrieved December 13, 2024 from <https://kubernetes.io/docs/concepts/overview/>
- [51] Vladimir Vondrus. 2022. *Zero-waste single-pass packing of power-of-two textures — blog.magnum.graphics*. Retrieved December 13, 2024 from <https://blog.magnum.graphics/backstage/pot-array-packing/>
- [52] Hao Wang, Han Tian, Jingrong Chen, Xinchun Wan, Jiacheng Xia, Gaoxiong Zeng, Wei Bai, Junchen Jiang, Yong Wang, and Kai Chen. 2024. Towards Domain-Specific Network Transport for Distributed DNN Training. In *21st USENIX Symposium on Networked Systems Design and Implementation, NSDI 2024, Santa Clara, CA, April 15-17, 2024*. USENIX Association. <https://www.usenix.org/conference/nsdi24/presentation/wang-hao>
- [53] Zilong Wang, Layong Luo, Qingsong Ning, Chaoliang Zeng, Wenxue Li, Xinchun Wan, Peng Xie, Tao Feng, Ke Cheng, Xiongfei Geng, Tianhao Wang, Weicheng Ling, Kejia Huo, Pingbo An, Kui Ji, Shideng Zhang, Bin Xu, Ruiqing Feng, Tao Ding, Kai Chen, and Chuanxiong Guo. 2023. SRNIC: A Scalable Architecture for RDMA NICs. In *20th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2023, Boston, MA, April 17-19, 2023*. USENIX Association, 1–14. <https://www.usenix.org/conference/nsdi23/presentation/wang-zilong>
- [54] Ahuva Mu'alem Weil and Dror G. Feitelson. 2001. Utilization, Predictability, Workloads, and User Runtime Estimates in Scheduling the IBM SP2 with Backfilling. *IEEE Trans. Parallel Distributed Syst.* 12, 6 (2001), 529–543. <https://doi.org/10.1109/10.1109/71.932708>
- [55] 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 2022, Renton, WA, USA, April 4-6, 2022*. USENIX Association, 945–960. <https://www.usenix.org/conference/nsdi22/presentation/weng>
- [56] Kaiqiang Xu, Xinchun Wan, Hao Wang, Zhenghang Ren, Xudong Liao, Decang Sun, Chaoliang Zeng, and Kai Chen. 2021. TACC: A Full-stack Cloud Computing Infrastructure for Machine Learning Tasks. *CoRR abs/2110.01556* (2021). arXiv:2110.01556 <https://arxiv.org/abs/2110.01556>
- [57] Jie You, Jae-Won Chung, and Mosharaf Chowdhury. 2023. Zeus: Understanding and Optimizing GPU Energy Consumption of DNN Training. In *20th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2023, Boston, MA, April 17-19, 2023*. USENIX Association, 119–139. <https://www.usenix.org/conference/nsdi23/presentation/you>
- [58] Yushan. 2024. *Managing Multiple CUDA + cuDNN Installations*. Retrieved December 13, 2024 from <https://medium.com/@yushantripleseven/managing-multiple-cuda-cudnn-installations-ba9cdc5e2654>
- [59] 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 *Proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2012, San Jose, CA, USA, April 25-27, 2012*. USENIX Association, 15–28. <https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/zaharia>
- [60] Junxue Zhang, Xiaodian Cheng, Wei Wang, Liu Yang, Jinbin Hu, and Kai Chen. 2023. FLASH: Towards a High-performance Hardware Acceleration Architecture for Cross-silo Federated Learning. In *20th USENIX Symposium on Networked Systems Design and Implementation, NSDI 2023, Boston, MA, April 17-19, 2023*. USENIX Association, 1057–1079. <https://www.usenix.org/conference/nsdi23/presentation/zhang-junxue>
- [61] Junxue Zhang, Chaoliang Zeng, Hong Zhang, Shuihai Hu, and Kai Chen. 2022. LiteFlow: towards high-performance adaptive neural networks for kernel datapath. In *SIGCOMM '22: ACM SIGCOMM 2022 Conference, Amsterdam, The Netherlands, August 22 - 26, 2022*. ACM, 414–427. <https://doi.org/10.1145/3544216.3544229>