

Speculative Container Scheduling for Deep Learning Applications in a Kubernetes Cluster

Ying Mao, *Member, IEEE*, Yuqi Fu, Wenjia Zheng, *Student Member, IEEE*,
Long Cheng, Qingzhi Liu, and Dingwen Tao *Member, IEEE*

Abstract—In the past decade, we have witnessed a dramatically increasing volume of data collected from varied sources. The explosion of data has transformed the world as more information is available for collection and analysis than ever before. To maximize the utilization, various machine and deep learning models have been developed, e.g. CNN [1] and RNN [2], to study data and extract valuable information from different perspectives. While data-driven applications improve countless products, training models for hyperparameter tuning is still a time-consuming and resource-intensive process. Cloud computing provides infrastructure support for the training of deep learning applications. The cloud service providers, such as Amazon Web Services [3], create an isolated virtual environment (virtual machines and containers) for clients, who share physical resources, e.g., CPU and memory. On the cloud, resource management schemes are implemented to enable better sharing among users and boost the system-wide performance. However, general scheduling approaches, such as spread priority and balanced resource schedulers, do not work well with deep learning workloads. In this project, we propose *SpeCon*, a novel container scheduler that is optimized for short-lived deep learning applications. Based on virtualized containers, such as Kubernetes [4] and Docker [5], *SpeCon* analyzes the common characteristics of training processes. We design a suite of algorithms to monitor the progress of the training and speculatively migrate the slow-growing models to release resources for fast-growing ones. Specifically, the extensive experiments demonstrate that *SpeCon* improves the completion time of an individual job by up to 41.5%, 14.8% system-wide and 24.7% in terms of makespan.

Index Terms—Container Management; Docker Swarm; Kubernetes; Tensorflow; Pytorch; Apache Yarn and Spark;

I. INTRODUCTION

With the dramatic growth of data volume from different sources, such as Edge, Cloud and Internet of Things [6]–[10], tons of businesses have transformed themselves to embrace data driven models. A number of deep learning and data mining backed applications are reforming our daily life. For instance, Uber [11] and Lyft [12] utilize the data to predicate the traffic and demand, which benefits clients with sharing economies. Another example is Apple FaceID, which uses computer vision to provide hands-free unlocking ability.

Y. Mao, Y. Fu, and W. Zheng are with the Department of Computer and Information Science at Fordham University in the New York City. E-mail: {ymao41, yfu81, wzhang33}@fordham.edu

L. Cheng is with the School of Computing, Dublin City University, Ireland. E-mail: long.cheng@dcu.ie

Q. Liu is with Information Technology Group, Wageningen University, The Netherlands. E-mail: qingzhi.liu@wur.nl

D. Tao is with School of Electrical Engineering and Computer Science at Washington State University, Pullman, Washington. E-mail: dingwen.tao@ieee.org

Leveraging the power of data, practitioners and researchers have proposed various learning architectures and models. Generative Adversarial Networks [13], for example, are taught to create words similar to given data sets in any domain: images, music, speech, and prose. A trained Generative Adversarial Network is a robot artist in a sense, who can compose impressive work.

To obtain a well-trained model, an iterative optimization algorithm is usually involved in hyperparameter tuning in order to find the best parameter combination, which can optimize a pre-defined evaluation function. This function mathematically determines the performance of a model, such as classification accuracy and inception score [14]. In each iteration, the model learns a batch of samples (training data). Based on the training results, the parameters in the model will be updated according to the evaluation function.

Despite the excellent performance achieved by learning models, the training process is both time-consuming and resource intensive. Depending on the hardware types, training the database of ImageNet [15] may take hours or even days [16]. Consequently, the process of training models is usually done on the cloud, where big companies, like Amazon Web Services [3], Microsoft Azure [17] and Google Cloud Platform [18], provide infrastructures support. The service providers create an isolated computing environment with traditional virtual machines and/or more recently, virtual containers, such as Kubernetes [4] and Docker [5]. In the cloud, resources, such as CPU and memory, are shared by many users. Therefore, cloud providers implement scheduling algorithms from different perspectives to ensure a better sharing of resources. For example, Kubernetes provides multiple options [19] that can be adopted to the system scheduler, such as Balanced Resource Allocation, which favors nodes with balanced resource usage, and Selector Spread Priority, which spreads the workload across hosts.

These general scheduling approaches, however, do not work well for training processes of deep learning applications. First of all, they fail to consider characteristics of evaluation functions, where the loss reduces dramatically at the beginning of the process, then converges and becomes stable. Although the resource usage may stay at the same level, the gain of the training process varies over the time. Secondly, the schedulers, in current system, migrate or terminate a container when it consumes much more resources than others and overwhelms the physical machine. With deep learning workloads, a converged model may hurt other training jobs on the same machine since the scheduler fails to distinguish them.

With respects to the features of deep learning applications,

we develop *SpeCon*, a Speculative Container scheduler, to accelerate the training jobs in a containerized cluster shared by multiple users. Based on the growth of evaluation functions, *SpeCon* categorizes a training job into different phases, which reflects the dynamics of the efficiency. According to the categories, *SpeCon* decides when a deep learning model starts its convergence. For a converged job, *SpeCon* first gathers the information of the system-wide resource usages as well as the existing jobs on each node and then, migrates this job to a most desirable host at the moment. The container reallocation in *SpeCon* releases resources to fast-growing jobs and boost the overall system performance. When all containers are converged, *SpeCon* keeps monitoring the cluster and rebalancing the workloads. The main contributions of this paper are summarized as follows.

- We analyze the characteristics of training processes for the various deep learning models and conduct experiments to study the overhead of saving and resuming jobs in a Kubernetes cluster.
- We propose *SpeCon* with a suite of algorithms that monitor the evaluation functions of deep learning models on the worker side and on the manager side. It collects system-wide information to calculate a weighted score for each worker to select a most desirable host for the migration. Furthermore, it rebalances the workload when all containers are converged.
- We implement *SpeCon* on top of a leading container orchestration platform, Kubernetes, and evaluate it with popular containerized deep learning applications. With intensive experiments on the cloud, we demonstrate that *SpeCon* achieves significant improvement in the time of completion. It reduces an individual job by as large as 41.5%, 14.8% in average, 24.7% for system makespan.

II. RELATED WORK

With prevalence of data-driven businesses, learning from big data draws an increasing attention in both industry and academia. In this domain, various data mining algorithms as well as deep learning frameworks have been proposed to maximize the utilization. On the back-end side, the cloud computing powered all those platforms by providing on-demand storage, computing and analytics services [20]–[23].

The cloud providers create a virtualized and isolated computing environment for each of the clients, who share resources on physical machines, such as CPU and memory. Traditionally, technologies of virtual machines are used for creating the environment [24]. To start a virtual machine, the system is required to boot its guest operating system, which consumes lots of resources and significantly increases the provisioning delay [25]. In the literature, various approaches focused on improving the provision time for virtual machines [26]–[31]. While many efforts have been made in this field, reducing the provision cost remains a challenge task due to the requirement of guest operating systems. After the provision, it is a common scenario that a live virtual machine need to be migrated to another node in the cluster to the accommodate the real-time network topology and workload [32]–[34]. However, the

heavy-weight design of storage system in virtual machines made the migration process an expensive operation. It usually cost a few tens of seconds for a regular virtual machine [35].

To address the limitations, containers are designed for deploying and running distributed applications, where multiple isolated containers are sharing the host operating system and physical resources. A comparative study in [36] demonstrated that containers, when compared to virtual machines, require far less provisioning and migration time. As the popularity increases, container scheduling has become an emerging research topic. While Stratus [37] exploited cloud properties and runtime estimates to reduce the dollar cost of cluster jobs executed on public clouds, PIVOT [38] investigated a cloud-agnostic system and the proposed scheduler enables data-intensive applications to scale across clouds instantly in a cost-efficient manner. Focusing on heterogeneities of a cluster, ECSched [39] enabled high-quality and fast placement decisions for concurrent deployment requests on heterogeneous clusters. The ECSched mapped scheduling decisions to a graphic data structure and model it as minimum cost flow problem. Besides CPU-based clusters, authors in [40] explored a GPU rich datacenter and proposed Knots, a GPU-aware resource orchestration layer and integrate it with Kubernetes container orchestrator.

Despite the efforts promoted in this field, limited number of them focused on scheduling containers for training processes of deep learning applications. While FlowCon [41] tried to set limits to container's resource requests, the evaluation focused on the single machine setting. Furthermore, Gandiva [42] was developed as a new cluster scheduling framework that utilizes domain-specific knowledge to improve latency and efficiency of training models. However, it does not take into account the properties of evaluation functions of learning models. Additionally, ProCon [43] tried to solve container placement in Kubernetes, but it fails to take container migration into consideration and SLAQ [44] schedules concurrent machine learning training jobs based on quality improvement for resource usage, by allocating cluster resources iteratively. However, it fails to re-allocate the resources when workload updates.

We present *SpeCon*, in this project, which keeps tracking the convergence of the deep learning models based on their pre-defined evaluation functions. In *SpeCon*, scheduling decisions are made according to the individual gains of each model in a given period. The scheduler considers reallocation of a container when it becomes a slow-growing one.

III. BACKGROUND AND MOTIVATION

In this section, we study the training process of containerized deep learning applications and motivate our work with an example.

A. Training a Deep Learning Model

Given a deep learning algorithm, training the model is the parameters tuning process to approach the global optimum of a predefined evaluation function. Initially, parameters are randomly selected for the model. This model is iteratively fed with mini-batches of training data, where each mini-batch




Figure 1: Training progress of three deep learning models

contains features and labels. It is, then, evaluated with the evaluation function that informs users how far off does it from the target. The returned value of this function is recorded along with the model parameters. Together, they propagate backward through the neural network architecture. In the next round of the training process, the algorithm updates and adjusts the parameters aiming to achieve a better performance. Finally, a well-trained model is generated with a parameter combination that produces an optimized evaluation function.

Fig. 1 presents the training process of three deep learning models, Variational Auto-Encoders(VAE) [45], Bidirectional Recurrent Neural Networks(Bi-RNN) [46] and Gated Recurrent Unit(GRU) [47]. The Y-axis is the normalized values of evaluation functions (the lower the better) and the X-axis is the cumulative time. As we can see from, considering the first 20% of the total training time, 65%, 98%, and 90% of the maximum reduction has been achieved for VAE, Bi-RNN and GRU, respectively. Since they are trained individually on a physical machine, resources are fully occupied by each model without competitions. Therefore, given the same resource amount, the training gain and efficiency is decreasing as the time goes on.




Figure 2: Overhead of saving and resuming VAE at different progress stage




Figure 3: Overhead of saving and resuming deep learning models at 30% and 70%

B. Saving and Resuming a Container

As analyzed above, a learning model, at any given iteration, contains the neural network design (the algorithm itself) and values of the network parameters that it has trained. That is to say the model progress can be saved during and after training. When necessary, the saved model can be resumed from exactly where they left off.

We experimentally study the overhead of saving and resuming learning models and compare to the nonstop training process. The experiments are conducted under Tensorflow [48] and Pytorch [49], where the training jobs are running inside a Kubernetes container. Fig. 2 shows the experiments of

training a VAE model, where we save the model at a given percentage according to the total number of iterations and resume it immediately after saving completes. Comparing to a nonstop training, denoted as *NS* on the figure, we find that the overhead is minimum with all less than 5 seconds and consistent with a standard deviation value of 1.43 across different percentages during the train process.

When saving a model in Tensorflow, there are four files generated internally.

- Meta graph file: This is a protocol buffer which saves the complete Tensorflow graph.
- Index file: It stores the list of variable names and shapes saved.
- Data file: This is the file which contains all the values of the weights, biases, gradients and all the other variables saved.
- Checkpoint file: It keeps a record of latest checkpoint files saved.

When restoring the model, the system read these files and restart on where it stopped.

Next, we investigate the overhead with different deep learning models. Fig. 3 illustrates the comparison between nonstop training and the ones with saving/resuming. The results of completion time are similar as shown on the figure that the overhead is minimum and consistent across different models.

IV. SPECON SYSTEM DESIGN

SpeCon is implemented on top of Kubernetes, a dominant container-orchestration platform designed to automate the deployment, scaling, and management of containerized applications. In this section, we present the system architecture of SpeCon in details, including its key modules and their functionalities as well as design logics.

A. Framework of Kubernetes

Generally, a cluster of containers comprises managers and workers. Managers are responsible for system management, such as allocating the resources and handling the failures. Workers are in charge of hosting the containers and executing the workloads. In a Kubernetes cluster of containers, a pod is a group of one or more containers, which are hosted on the same worker and share the lifecycle as well as storage resources.

Typically, there are 6 basic units in a Kubernetes cluster. The manager nodes consist of API Server, Controller Manager and Default Scheduler and etcd. The Kuberlet and Service Proxy resident in workers.

- API Server: It is the main management point of the entire cluster and processes REST operations, validates them, and updates the corresponding objects in storage.
- Controller Manager: It runs controllers, which are the background threads that handle routine tasks.
- Default Scheduler: It watches for newly created Pods that have no node assigned and is responsible for placement of pods on workers.
- etcd: It is a distributed data storage solution that stores all the data, e.g. configuration, state, and metadata.

- **Kuberlet**: It is responsible for maintaining a set of pods, which are composed of one or more containers, on a local system.
- **Service Proxy**: It maintains network rules on nodes, e.g. implementing a form of virtual IP for services.

B. SpeCon System Architecture

Fig. 4 draws the SpeCon architecture in a 3-node cluster, which consists of 1 manager and 2 workers. Since SpeCon is built on top of Kubernetes, Fig. 4 includes the above mentioned 6 components in light blue. Besides, it contains 4 key modules of SpeCon, **Worker Monitor** and **SpeCon Scheduler** on the manager side, **Container Monitor** on the worker side, **SpeCon Messenger** that works in between. Particularly,

- **Container Monitor**: It runs on the worker side, which provides the resources, such as CPU and memory, for containers to execute the training jobs. It keeps tracking the progress of each model from two perspectives, resource usages and current returns of their evaluation functions. Based on the collected data, it categorizes training jobs into different stages and active the corresponding algorithm (details in Section V).
- **Worker Monitor**: It resides on manager side, which is responsible for the overall system management. It monitors all the workers in the cluster, such as the number of running jobs and the resource availabilities on each of them. Moreover, it processes the messages from workers and communicates with SpeCon scheduler for further actions.
- **SpeCon Scheduler**: It works on the manager side. When the worker monitor decides that a training job needs to be migrated to another worker node, it utilizes the data that gathered by worker monitor and execute the algorithm to calculate a score for each of the worker, and then, select a most desirable node to host this migrated training job (details in Section V).
- **SpeCon Messenger**: It generates heartbeat messages that transfer between managers and workers. For example, workers, periodically, send a message to notify the manager such as the number of training job hosted on themselves and their categories. Additionally, whenever a job needs to be migrated, workers make use of this channel to notify the manager immediately.

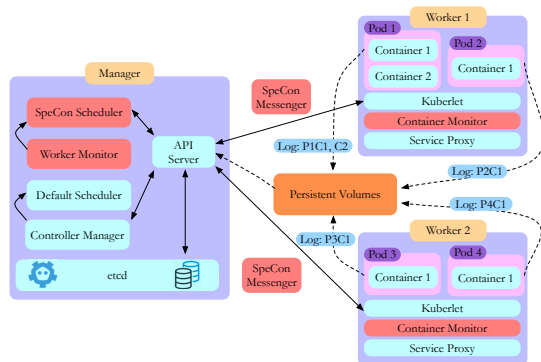


Figure 4: SpeCon System Modules

The training jobs in SpeCon write their progress information into the log system in the system. Particularly, SpeCon utilizes the Persistent Volume (PV [50]), which is a piece of storage in the cluster and has a lifecycle independent of any individual Pod that uses the PV, to store the logs. This provides the stability to the SpeCon system.

V. SPECON SOLUTION DESIGN

SpeCon aims to improve the system performance by migrating the slow-growing jobs to other workers to release resources for the fast-growing ones. It achieves the improvement through three steps, (1) identify slow-growing jobs; (2) select the most desirable host; (3) rebalance the workers. The Table I summarizes the parameters and functions that are used in SpeCon.

Table I: Notations

$w_i \in W$	The worker ID in the worker set
$c_i \in C$	The container ID in the container set
$E_{w_i, c_j}(t)$	The value of evaluation function of c_j on w_i at time t
$G_{w_i, c_j}(t)$	The growth of evaluation function of c_j on w_i at time t
$R(w_i, t)$	Resource consumption of w_i at time t
PC	The progressing category set
WC	The watching category set
CC	The converged category set
α	Categorization threshold
$t_m - t_{m-1}$	Categorization interval
T	Candidate set
S_{w_i}	The scoring set that stores weighted scores for each worker
bf	The balance factor for an uniform distribution

A. Identifying the slow-growing jobs

In SpeCon, we consider a cluster of containers that hosted on workers. Inside the cluster, we have managers (denoted as M). In addition, we denote $c_i \in C$ to be the set of containers and $w_i \in W$ to be the set of workers. Therefore, the relation, $c_m \in w_n$, represents that a container, c_m , is running on a worker w_n .

In our problem setting, training jobs are running inside containers, where each container hosts one specific job. Consequently, each container, c_i , can be resolved as one particular training job in our setting. As analyzed in the previous sections, each job has a predefined evaluation function. During the whole training process, it is a time series function, where at the end of each iteration the return value updates. When querying in the middle of an iteration, the previous value will be returned. Based on the query time, Equation 1 presents the growth of the training job in a given interval, $t_2 - t_1$.

$$G_{w_i, c_j}(t_2) = E_{w_i, c_j}(t_2) - E_{w_i, c_j}(t_1) \quad (1)$$

, where c_j is a container that runs a training job on w_i with $E_{w_i, c_j}(t)$ as the evaluation function.

According to the values of $G_{w_i, c_j}(t)$, SpeCon classifies c_i into 3 different categories.

- **Progressing Containers (PC)**: The jobs inside these containers are still in the fast growing stage such that their evaluation function is progressing quickly.
- **Watching Containers (WC)**: When a training job join WC category, it indicates that this job is slowing down

Algorithm 1 Container Monitor: c_j on w_i

```

1: System Parameters:
   Categorization Interval :  $t_m - t_{m-1}$ 
   Categorization Threshold :  $\alpha$ 
   Evaluation Function :  $E_{w_i, c_j}(t)$ 
2: Initial Settings:
    $\forall c_{id} \in w_i, PC_{w_i}.insert(c_{id})$ 
    $\forall c_{id} \in w_i, c_{id}.migrated = \text{False}$ 
3: for  $c_j \in w_i$  &  $c_j.migrated = \text{False}$  do
4:   At time  $t_m$ : Calculate  $E_{c_j, w_i}(t_m)$ 
5:    $G_{w_i, c_j}(t_m) = |E_{w_i, c_j}(t_m) - E_{w_i, c_j}(t_{m-1})|$ 
6:   if  $G_{w_i, c_j}(t_m) < G_{w_i, c_j}(t_{m-1})$  &  $G_{w_i, c_j}(t_m) < \alpha$ 
       then
7:     if  $c_j \in PC_{w_i}$  &  $c_j \notin WC_{w_i}$  &  $c_j \notin CC_{w_i}$  then
8:        $PC_{w_i}.remove(c_j)$ 
9:        $WC_{w_i}.insert(c_j)$ 
10:    if  $c_j \notin PC_{w_i}$  &  $c_j \in WC_{w_i}$  &  $c_j \notin CC_{w_i}$  then
11:       $PC_{w_i}.remove(c_j)$ 
12:       $WC_{w_i}.insert(c_j)$ 
13:    if  $c_j \notin PC_{w_i}$  &  $c_j \notin WC_{w_i}$  &  $c_j \in CC_{w_i}$  then
14:       $c_j$  remains in  $CC_{w_i}$ 
15:      if  $|PC| + |WC| > 1$  then
16:        SpeCon-Messenger( $w_i, c_j$ ).Request-Reallocation
17:    else if  $G_{w_i, c_j}(t_m) \geq G_{w_i, c_j}(t_{m-1})$  &  $G_{w_i, c_j}(t_m) < \alpha$ 
        then
18:       $c_j$  stays in the current category
19:    else if  $G_{w_i, c_j}(t_m) > \alpha$  then
20:       $PC_{w_i}.insert(c_j)$ 
21:       $WC_{w_i}.remove(c_j)$ 
22:       $CC_{w_i}.remove(c_j)$ 

```

on the gain of the evaluation function. In this zone, the gain starts jittering and, depending on the model itself, it may speedup the progress again.

- Converged Containers (CC): If training jobs in WC continue slowing down, they are inserted into a Converged Container category, which suggests that they are slow-growing on their evaluation functions and unlikely to speedup again.

Based on the above analysis, we develop the Algorithm 1 in SpeCon to keep tracking the progress of training jobs on each worker. Please note that, in order to avoid the unnecessary network traffic and maintain scalability, the containers in SpeCon will be reallocated at most once.

As shown in Line 1, the Algorithm 1 on each worker first fetches key parameters from the system. The categorization interval ($t_m - t_{m-1}$) determines how frequent a worker checks containers' log files, which store return values of evaluation functions. The categorization threshold, α , is a percentage value that SpeCon uses to decide whether the progress is growing quickly. Then, in Line 2, the algorithm initializes parameters such that all containers are inserted into PC when they join a worker and their migration indicator is set to false at the beginning.

At time t_m , it reads the current value of the evaluation func-

tion and calculates the gain during the previous categorization interval, $t_m - t_{m-1}$ (Line 3-5). If the gain is less than the previous round and the threshold α , SpeCon updates the category of this container based on the following condition (Line 6-16).

- If the container is currently in the category PC , but not WC and CC , it will be removed from PC and inserted to WC . It means that this job starts slowing down, but not stable yet.
- If the container is currently in the category WC , but not PC and CC , it will be removed from WC and inserted to CC . It indicates that this job has started its convergence process.
- If the container is currently in the category CC , but not PC and WC , it suggests that the value of the gain during the interval continues decreasing and it has converged. In this condition, it will remain in the CC category and the worker will call SpeCon messenger to send a migrated request of (c_j, w_i) to the manager.

In a scenario that the current gain is less than α , but larger than the previous gain value, this container stays in the current category. This is due to the fact that when a training job goes on, the model randomly selects and updates parameters at each iteration that leads to an unstable trends (bouncing growth values). In this stage, the container remains in the same category and waiting for next round (Line 17-18).

At the moment that the gain becomes larger than the threshold, we reset the container's category to PC , which is utilized to accommodate a sudden change of the evaluation function and possible errors from the previous rounds (Line 19-22).

B. Scheduling on slow moving jobs

The container monitor that runs each worker keeps tracking the evaluation functions and collects data that stores in the persistent volume. Whenever receives a request from workers, the manager responses the reallocation message, e.g. (w_i, c_j), by executing the Algorithm 2. The main objective of Algorithm 2 is to select the most desirable worker node to host this container, e.g. c_j .

SpeCon utilizes a weighted scoring algorithm to rank worker nodes. As shown in Line 1 of Algorithm 2, it first initializes parameters, which include a candidate set (T) that uses to store targeted workers and it is set to empty initially. Additionally, SpeCon maintains predefined weights for each category as well as resource consumption functions.

For each worker node in the cluster, it calculates a score based on the number of containers in each category and its weights. SpeCon aims to improve the efficiency by allocating more resources to fast-growing jobs and limit resources for slow-growing ones. With this objective in mind, the values of weights have a relationship of $w_{pc} > w_{wc} > w_{cc}$, where priorities are given to progressing containers (Line 2 - 3).

Given the scores of the workers, the algorithm finds out workers that have the minimum score. Those workers are inserted into the candidate set. If multiple workers have the minimum score, it results in more than one candidates (Line

Algorithm 2 Speculative Scheduling of c_j in w_i on Manager

```

1: Parameter Initialization:
   Candidate_Set :  $T = \phi$ 
   Weights :  $w_{pc}, w_{wc}, w_{cc}$ 
   Resource Consumption on  $w_i$  at time  $t$  :  $R(w_i, t)$ 

2: for  $w_j \in W$  do
3:    $S_{w_j} = |PC_{w_j}| \times w_{pc} + |WC_{w_j}| \times w_{wc} + |CC_{w_j}| \times w_{cc}$ 
4: for  $w_j \in W$  do
5:   Find Min( $S$ )
6:   if  $S_{w_j} = \text{Min}(S_{w_j \in W})$  then
7:      $T.\text{insert}(w_j)$ 
8:   if  $w_i \in T$  then
9:      $c_{id}.\text{migrated} = \text{True}$ 
10:     $c_j$  remains on  $w_i$ 
11: else
12:   for  $w_j \in T$  do
13:     if  $|T| = 1$  then
14:       Save  $c_j$  on  $w_i$ 
15:        $c_{id}.\text{migrated} = \text{True}$ 
16:       Restore  $c_j$  on  $w_j$ ;
17:     else if  $|T| > 1$  then
18:       Rank  $w_j \in T$  with  $R(w_j, t)$ 
19:       Save  $c_j$ 
20:        $c_{id}.\text{migrated} = \text{True}$ 
21:       Restore  $c_j$  on  $w_j$  with lowest  $R(w_j, t)$ ;

```

4-7). The algorithm checks whether w_i , the current host of c_j , is in the set or not. Then the following logic will be executed.

- If w_i in candidate set, worker w_i will be returned, which means that container c_j should continue running on w_i . In this case, the algorithm marks the container as migrated (Line 8-10). This manner reduces unnecessary overhead caused by migration.
- If w_i is not in the candidate set, the algorithm takes two branches. (1) If there is only 1 worker in T , SpeCon select this worker for migration. (2) If $|T| > 1$, SpeCon ranks candidates with their resource usages and selects the one with least resource usage as the new host (Line 11-21).

C. Rebalancing workloads in the cluster

Together with Algorithm 1 and 2, SpeCon distributes the workload based on real-time returns from evaluation functions of containers as well as current resource consumption on workers. At this stage, however, it could make inaccurate decisions due to missing information of finishing time. For instance, if $\forall c_i \in W, c_{i}.\text{migrated} = \text{True}$ that means all jobs have converged, therefore, the workload distribution in the cluster is fixed, which would result in an imbalanced cluster. As intuitive example, when all 4 jobs become converged in a 2-node cluster, W_1 hosts c_1 and c_2 and W_2 contains c_3 and c_4 . In a scenario that c_1 and c_2 completes before c_3 and c_4 , W_1 runs without any workloads.

In a real cloud environment, it is challenge and costly to obtain an accurate finish time due to the various implementations,

dynamic workloads and resource competition. In our solution, SpeCon incorporates the converged duration, which is time length between when a container is marked as converged and current time to compare active containers and re-distribute the workload in a cluster by using Algorithm 3.

The algorithm prepares the required parameters in Line 1. Then, it calculates the sum of active jobs on workers across the cluster and computes the converged duration for each container c_j on each worker w_i . The each converged duration, d_{c_j} , is inserted into a set D (Line 2-6). If there is no active jobs on a particular worker w_i , its id would be inserted to the candidate set T , which stores the workers that can take more workload (Line 7-8). Given t_j and $|W|$, the balanced factor, bf , is obtained in Line 9. The value of bf is based on uniform distribution.

Depending on the number of workers in the candidate set T , Algorithm 3 takes the following two branches.

- If T is nonempty, it suggests that, at least, one worker is idling. For each idling worker w_i in T , the algorithm assigns c_j , which has the smallest d_{c_j} to w_i if the number of active containers on it is less than bf and marks c_j to be rebalanced (Line 10-18).
- If T is empty, it means that every worker runs some active containers. Then, SpeCon enumerates all workers and finds the ones that host less jobs than $bf - 1$, which indicates they can hold, at least, one more job. These workers are inserted into T (Line 19-22). Then, the algorithm finds the c_j with the smallest d_{c_j} and, in the meanwhile, c_j runs on a worker w_j such that $w_j \notin T$ (Line 23-25). It basically avoids the scenario that c_j assigns to its current host. With a nonempty T , SpeCon assigns the previously found c_j to w_i which has room for an additional container and marks this container to be rebalanced (Line 26-31).

VI. EVALUATION

In this section, we evaluate the performance of SpeCon through intensive cloud-executed experiments.

A. Experimental Setup

1) *Implementation, Testbed and Workloads*: SpeCon is integrated into Kubernetes 1.17 as plug-in modules that reside on both managers and workers. It receives tasks from the manager, directs the given tasks to workers, monitors the evaluation functions and balances the workloads.

We build a testbed on a NSF Cloudlab [51] datacenter that hosts at University of Utah. Specifically, we use multiple M510 as our physical machines that contain two 8-core Intel Xeon D-1548 at 2.0 GHz, 64GB ECC Memory, and 256 GB NVMe flash storage. Two clusters have been built to evaluate SpeCon: **Cluster 1**: 1 Manager and 4 Workers **Cluster 2**: 1 Manager and 8 Workers

SpeCon focuses on accelerating the training process of deep learning models. We evaluate it with two popular open-source platforms, Pytorch (P) [49] and Tensorflow (T) [48]. As shown in Table II, we choose 5 different model on those two platforms as our workloads that execute inside containers.

Algorithm 3 Rebalance Active Container in the Cluster

```

1: Parameter Initialization:
   Candidate_Set :  $T = \phi$ 
   Resource Consumption on  $w_i$  at time  $t$ :  $R(w_i, t)$ 
    $|w_i|$ : The number of active containers on  $w_i$ 
    $t_{c_j}$ : The timestamp when  $c_j$  is marked as migrated
    $t, D$ : The current timestamp, duration set  $D$ 
    $bf, |W|$ : The balance factor, total number of workers

2: for  $w_i \in W$  do
3:    $sum = sum + |w_i|$ 
4:   for  $c_j \in w_i$  do
5:      $d_{c_j} = t - t_{c_j}$ 
6:      $D.insert(d_{c_j})$ 
7:     if  $|w_i| = 0$  then
8:        $T.insert(w_i)$ 
9:    $bf = \lfloor sum \div |W| \rfloor$ 
10: if  $|T| > 0$  then
11:   for  $w_i \in T$  do
12:     if  $|w_i| < bf$  then
13:       Find  $c_j$  with  $Min(d_{c_j})$  in  $D$ 
14:        $D.remove(d_{c_j})$ 
15:       Migrate  $c_j$  to  $w_i$ 
16:        $c_j.rebalanced = True$ 
17:     else
18:       Continue;
19:   else if  $|T| = 0$  then
20:     for  $w_i \in W$  do
21:       if  $|w_i| < bf - 1$  then
22:          $T.insert(w_i)$ 
23:     for  $w_i \notin T$  do
24:        $\forall c_j \in w_i$ 
25:         Find  $c_j$  with  $Min(d_{c_j})$  in  $D$ 
26:     for  $w_j \in T$  do
27:       if  $|w_j| < bf - 1$  then
28:         Migrate  $c_j$  to  $w_j$ 
29:          $c_j.rebalanced = True$ 
30:       else
31:         Continue;

```

Table II: Tested Deep Learning Models

Model	Loss Function	Platform
Variational Autoencoders (VAE)	Recon. Loss	P/T
Gated Recurrent Unit (GRU)	Quadratic Loss	P/T
Bidirectional-RNN	Softmax	P/T
Recurrent Network (RNN)	Softmax	P/T
Dynamic RNN	Softmax	P/T

2) *Evaluation Metrics*: The training processes of deep learning applications are computational intensive. They are more sensitive to CPU powers than memory spaces and network bandwidth. The following evaluation metrics are considered in our experiments.

- **Completion time**: The completion time of each individual job in the cluster.
- **Average completion time**: The average completion time of among all the workload in the cluster.

- **Makespan**: The makespan is the time between the initialization of the first job and the completion of the last one across all workers.

When multiple models are training on the same worker, the jobs create a highly dynamic computing environment in terms of resource competition. Therefore, we design the following submission schedules to ensure a comprehensive evaluation.

- **Fixed Schedule**: The time to launch containers follows a fixed interval. It simulates an administrator controlled cloud environment.
- **Random Schedule**: The time to launch a job is randomly selected within an interval. It simulates an user-specified, first-come-first-serve cloud environment.

As for the weight of each category that utilized by Algorithm 2, we use $w_{pc} = 2, w_{wc} = 1.5$ and $w_{cc} = 1$. Due to the page limit, we omit the analysis of weights in this paper. Additionally, we compare SpeCon with the default scheduler in Kubernetes, which noted as **DS** in the rest of the evaluation section. For a fair comparison, we utilize a fixed initial container placement, where workload are evenly distributed based on the number of containers, for both *DS* and SpeConso that only the algorithms in SpeCon can affect the performance.

B. Fixed schedule

Firstly, we conduct experiments with a fixed schedule such that containers with the same model image are submitted to the system one by one with a fixed interval. In this experiment, we use VAE model and 50 second interval. Furthermore, we set $\alpha = 0.01, |t_m - t_{m-1}| = 30$ for SpeCon.




Figure 5: Fixed Schedule with 20 VAE jobs and 50s interval ($\alpha = 0.01, |t_m - t_{m-1}| = 30$)

As shown on Fig. 5, we can clearly see that SpeCon outperforms DS in terms of completion time. The largest gain is found on Job-17, which reduces from 348.7s to 217.6s, that is a 37.6% reduction. There are 11 jobs out of 20 get a reduction the completion time. The average reduction in those 11 jobs is 83.6s (26.7%). For those 9 jobs that fail to obtain improvement, the average increase is 18.3s (8.1%). Overall, the average completion time of SpeCon reduces from 275.2s to 234.7s (14.7%). With speculative scheduling on slow-growing containers, SpeCon is able to migrate them to other workers and release resources for fast-growing jobs,

which leads to reduction on completion time. As for the makespan of the system, *SpeCon* reduces the value from 1260s to 1190s.

C. Random schedule

1) *Single model*: There are two key parameters used by *SpeCon*, α that determines the category of each job and $|t_m - t_{m-1}|$ that decides when to start categorization. We conduct experiments to study the performance of *SpeCon* under various parameter combinations. In these experiments, we randomly submit jobs (single VAE model) within the interval [0, 300] to the cluster.

Table III presents the results when comparing to *DS* and $\alpha = 0.05$, $|t_m - t_{m-1}| = 20, 25, 30$ as well as $|t_m - t_{m-1}| = 30$, $\alpha = 0.01, 0.05, 0.1$. The data includes, the percentage of jobs that achieve reduction on completion time (Reduced), the overall average completion time improvement among all the jobs (Overall), the reduction of the best job (Best), and the improvement of makespan of the whole system (Makespan). The results demonstrate that *SpeCon* clearly outperforms *DS* from different perspectives. When comparing to the fixed schedule, degrees of improvements that achieved by *SpeCon*, with random submissions, reduced. The reason lies in the fact that random submissions creates challenges categorizing the running containers since there is delay when we monitoring the evaluation function and the single model generates difficulties to select most desirable worker since the same training pattern has a higher probability leads to the same score for workers.

Table III: Parameters in *SpeCon*

Parameters	Reduced	Overall	Best	Makespan
0.05, 20s	60.0%	5.7%	29.1%	6.1%
0.05, 25s	60.0%	6.8%	27.8%	7.5%
0.05, 30s	55.0%	3.2%	21.3%	4.0%
0.01, 30s	65.0%	9.6%	31.0%	8.2%
0.05, 30s	55.0%	3.2%	21.3%	4.0%
0.10, 30s	60.0%	6.4%	23.9%	3.1%

Remarks: Based on the experimental results and analysis in Section VI-B and VI-C, we can conclude that *SpeCon* can improve both the system makespan and the completion time of individual learning jobs with different parameter settings, compared to the original Kubernetes system. In the meantime, different parameter settings do affect the degree of improvement. Since values of α and $t_m - t_{m-1}$ indicate how quickly a given container enters *CC* category in Algorithm 1, they controls how many containers can be considered for migration and frequency of Algorithm 2. Therefore, the values are proportional to overhead, e.g. calculating the data for algorithms and cost for migrating a container, that *SpeCon* brings to the system. Consequentially, as the frequency of running Algorithm 1 and 2 decreases, the rooms scheduling and adjusting the system will be reduced. Therefore, the best α setting depends on the number of active containers in the system, the deep learning model in each container and the corresponding datasets. Due to the page limit, we omit the theoretical analysis of parameters and, in the rest of the evaluation, the experiments utilize $|t_m - t_{m-1}| = 30$, $\alpha = 0.01$.

2) *Multiple models*: Next, we conduct experiments that employ a random selection of both models and platforms in Table II. Fig. 6 plots the same cluster with 20 randomly selected jobs submitted to the system within [0, 300]. Overall, there are 15 out of 20 jobs (75%) record a reduced completion time. The largest gain is obtained on Job-14, whose completion time reduced from 1167.8s to 799.2s (31.6%). The average improvement of those 15 jobs is 165.7s. For the other 5 jobs, the average increase is 110.8s. The increase is cause by migrating slow growing containers to the same worker lead to more resource competition. *SpeCon* sacrifices a small portion of jobs to improve others significantly. With a system-wide consideration, *SpeCon* achieves 13.6% reduction on the average completion time. In terms of makespan in the cluster, *SpeCon* achieves 24.7% reduction from 1421s to 1070s.




Figure 6: Random schedule [0, 300] with 20 random jobs

The previous experiments contain only single model (VAE). Next, we conduct experiments that employ a random selection of both models and platforms in Table II. In this experiment, we randomly launch 30 jobs within the interval of [0, 600]. Fig. 7 shows the completion time of these jobs. Comparing with *DS*, 18 out of 30 jobs (60.0%) finish training faster with *SpeCon*. For those 18 jobs, the average improvement is 15.1% and for reminder 12 jobs, the average increase is 10.6%. Among the workloads, Job-5 gains the largest, 28.8%, reduction from 1539.6s to 1095.5s (detailed explanations below). Overall, the average completion time improves from 865.7s to 800.2s, 7.6% and makespan reduces from 1760s to 1498s, 14.8%.




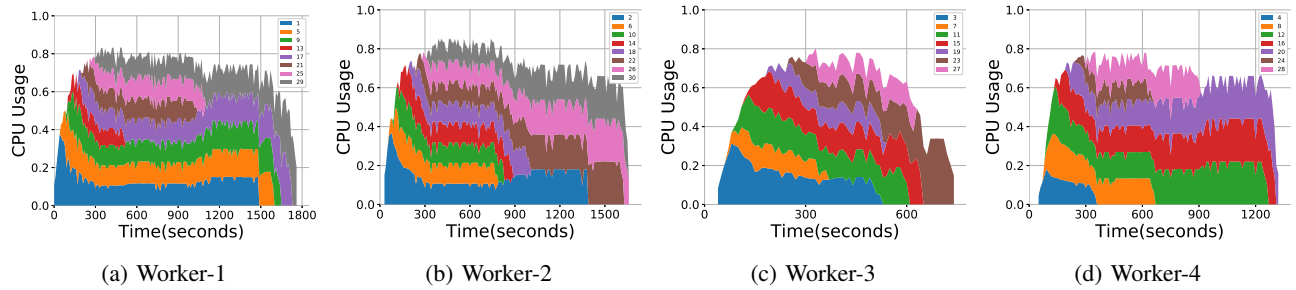
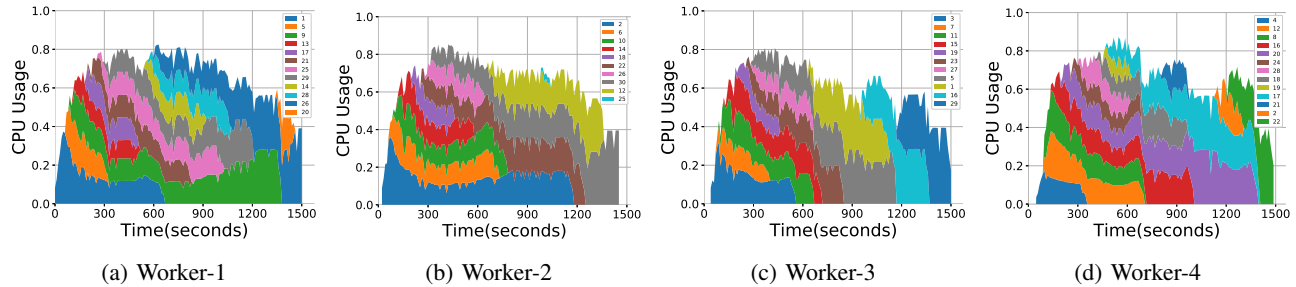


Figure 7: Random schedule [0, 600] with 30 random jobs

Next, we dig into the details of the experiment. Fig. 8 and Fig. 9 illustrate the container distribution over time on each

Figure 8: Workload distribution for 30-job experiment with *DS*Figure 9: Workload distribution for 30-job experiment with *SpeCon*Figure 10: CPU usage for 30-job experiment with *DS*Figure 11: CPU usage for 30-job experiment with *SpeCon*

of the workers. Comparing Fig. 8 with Fig. 9, we have two observations. (1) As described in the experiment setup, both *DS* and *SpeCon* use a fixed initial container placement for fairness. Therefore, container 1, 2, 3, 4, as shown on the figures, is hosted by Worker-1,-2,-3,-4 respectively. Then, container 5, 6, 7, 8, rotate to another around of workers. (2) We can clearly see that none of the containers get migrated. With the default scheduler (*DS*), the migration only happens when the hosting worker is out of resource or other unhealthy condition. *DS* fails to react based on the job's progress. Although, the fixed initial container placement ensures a balance at the beginning, the default scheduler fails to take characteristics of loss functions and resource utilization into account. After the initial placement, it lacks a mechanism to discover and readjust the system.

With *SpeCon*, however, the containers can be migrated according to their training progresses and system-wide workload distribution. In Fig. 9, we mark the job that is migrated due to its slow-growing progress as “Job-ID-M” and “Job-ID-R” indicates that this job is migrated as a reaction of system-wide rebalancing. For example, Job-14 was originally running on Worker-2 and it was added to *CC* category at 537.3s by Algorithm 1. Then, it is delivered to Algorithm 2 to decide its new host. At this particular moment, the scores of Worker-1,-2,-3,-4 were 63, 88, 92, 72. Therefore, Job-14 was migrated to Worker-1 and the released resource can be utilized by other jobs on this worker. This migration, together with Job-26's migration (at time 592.1s), benefit Job-2, Job-6, Job-10, Job-30, which all achieve reduction on the completion time. In addition, for Job-5 on Worker-1, it was added to *CC* at 286.8s




Figure 12: Random schedule [0, 1200s] with 50 random jobs

and reassigned to Worker-3, which has 7 active jobs, but 4 of them are in CC already. On Worker-3, several of them finished (Job-3, -7, -11, -23, -27) and released resources so that Job-5 can utilize that results in a faster completion time.

Furthermore, when the system is unbalanced, Algorithm 3 is executed to re-distribute the workload. For example, Job-1 was originally running on Worker-1 and, at time 620s, the system was unbalanced due to finished job, such as Job-3, Job-7, Job-21. SpeCon captured this moment and redistributed Job-1 on Worker-1 to Worker-3. This migration ease the resource competition on Worker-1, which has 8 active containers at the time. It benefits both Job-1, which has been moved to a worker with less workload, and others that can utilize the released resources.

Fig. 10 and Fig. 11 plot the comparison of CPU usage between DS and SpeCon. The resource usage tops at around 80% since Kubernetes reserves part of them for the system itself. The resource utilization rates are similar between them. Due to migration and re-distribution, SpeCon presents more jobs on the figures.

D. Larger cluster

Finally, we conduct experiment with a larger cluster, 1 manager and 8 workers. The more workers give more options for SpeCon, which leads to challenges when reallocating containers. Fig. 12 present the experiment with 50 random jobs running in the larger cluster. Overall, SpeCon records 30 out of 50 jobs (60.0%) are completing faster than DS . The largest gain is found on Job-50, which reduces from 467.5s to 273.4s that is a 41.5% reduction. The average completion time reduces 7.2% from 596.3s to 553.0s. Additionally, the makespan value reduces from 2070s to 1840s, 11.1%.

VII. CONCLUSION

We propose SpeCon in this work. It aims to facilitate speculative container scheduling for deep learning training jobs at runtime. We presented the detailed system design of SpeCon, and conducted extensive experiments with 5 different deep learning models on 2 frameworks, Pytorch and Tensorflow, in a cloud computing environment. Our experimental results have proven the effectiveness of SpeCon. Specifically, compared to a default scheduler, it has achieved significant performance boost in the presence of various deep learning workloads, by up to 41.5% reduction in completion time for an individual job and 14.8% in average. Furthermore, SpeCon achieves improvement on makespan for up to 24.7%.

REFERENCES

- [1] "Convolutional neural network," https://en.wikipedia.org/wiki/Convolutional_neural_network.
- [2] "Recurrent neural network," https://en.wikipedia.org/wiki/Recurrent_neural_network.
- [3] "Amazon web services," <https://aws.amazon.com/>.
- [4] "Kubernetes," <https://kubernetes.io/>.
- [5] "Docker," <https://docker.com/>.
- [6] A. Acharya, Y. Hou, Y. Mao, and J. Yuan, "Edge-assisted image processing with joint optimization of responding and placement strategy," in *2019 International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCoM) and IEEE Smart Data (SmartData)*. IEEE, 2019, pp. 1241–1248.
- [7] A. Acharya, Y. Hou, Y. Mao, M. Xian, and J. Yuan, "Workload-aware task placement in edge-assisted human re-identification," in *2019 16th Annual IEEE International Conference on Sensing, Communication, and Networking (SECON)*. IEEE, 2019, pp. 1–9.
- [8] H. H. Harvey, Y. Mao, Y. Hou, and B. Sheng, "Edos: Edge assisted offloading system for mobile devices," in *2017 26th International Conference on Computer Communication and Networks (ICCCN)*. IEEE, 2017, pp. 1–9.
- [9] Y. Mao, J. Oak, A. Pompili, D. Beer, T. Han, and P. Hu, "Drops: Dynamic and resource-aware placement scheme for docker containers in a heterogeneous cluster," in *2017 IEEE 36th International Performance Computing and Communications Conference (IPCCC)*. IEEE, 2017, pp. 1–8.
- [10] Y. Mao, J. Wang, J. P. Cohen, and B. Sheng, "Pasa: Passive broadcast for smartphone ad-hoc networks," in *2014 23rd International Conference on Computer Communication and Networks (ICCCN)*. IEEE, 2014, pp. 1–8.
- [11] "Uber," <https://www.uber.com/>.
- [12] "Lyft," <https://www.lyft.com/>.
- [13] "Generative adversarial networks," https://en.wikipedia.org/wiki/Generative_adversarial_network.
- [14] "Model evaluation," <https://heartbeat.fritz.ai/introduction-to-machine-learning-model-evaluation-fa859e1b2d7f>.
- [15] "Image-net," <http://www.image-net.org/>.
- [16] Y. You, Z. Zhang, C.-J. Hsieh, J. Demmel, and K. Keutzer, "Imagenet training in minutes," in *Proceedings of the 47th International Conference on Parallel Processing*. ACM, 2018, p. 1.
- [17] "Microsoft azure," <https://azure.microsoft.com/en-us/>.
- [18] "Goole cloud platform," <https://cloud.google.com/>.
- [19] "Kubernetes scheduler," <https://kubernetes.io/docs/concepts/scheduling/kube-scheduler/>.
- [20] A. Saiyeda and M. A. Mir, "Cloud computing for deep learning analytics: A survey of current trends and challenges," *International Journal of Advanced Research in Computer Science*, vol. 8, no. 2, 2017.
- [21] Y. Mao, Y. Fu, S. Gu, S. Vhaduri, L. Cheng, and Q. Liu, "Resource management schemes for cloud-native platforms with computing containers of docker and kubernetes," 2020.
- [22] X. Chen, L. Cheng, C. Liu, Q. Liu, J. Liu, Y. Mao, and J. Murphy, "A woa-based optimization approach for task scheduling in cloud computing systems," *IEEE Systems Journal*, 2020.
- [23] W. Zheng, Y. Song, Z. Guo, Y. Cui, S. Gu, Y. Mao, and L. Cheng, "Target-based resource allocation for deep learning applications in a multi-tenancy system," in *2019 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 2019, pp. 1–7.
- [24] G. Galante and L. C. E. d. Bona, "A survey on cloud computing elasticity," in *Proceedings of the 2012 IEEE/ACM Fifth International Conference on Utility and Cloud Computing*. IEEE Computer Society, 2012, pp. 263–270.
- [25] J. Zhang, F. Ren, R. Shu, T. Huang, and Y. Liu, "Guaranteeing delay of live virtual machine migration by determining and provisioning appropriate bandwidth," *IEEE Transactions on Computers*, vol. 65, no. 9, pp. 2910–2917, 2015.
- [26] B. Xavier, T. Ferreto, and L. Jersak, "Time provisioning evaluation of kvm, docker and unikernels in a cloud platform," in *2016 16th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*. IEEE, 2016, pp. 277–280.
- [27] Q. Wu, F. Ishikawa, Q. Zhu, and Y. Xia, "Energy and migration cost-aware dynamic virtual machine consolidation in heterogeneous cloud datacenters," *IEEE transactions on Services Computing*, 2016.

- [28] V. R. Reguri, S. Kogotam, and M. Moh, "Energy efficient traffic-aware virtual machine migration in green cloud data centers," in *2016 IEEE 2nd International Conference on Big Data Security on Cloud (BigDataSecurity), IEEE International Conference on High Performance and Smart Computing (HPSC), and IEEE International Conference on Intelligent Data and Security (IDS)*. IEEE, 2016, pp. 268–273.
- [29] U. Mandal, P. Chowdhury, M. Tornatore, C. U. Martel, and B. Mukherjee, "Bandwidth provisioning for virtual machine migration in cloud: Strategy and application," *IEEE Transactions on Cloud Computing*, vol. 6, no. 4, pp. 967–976, 2016.
- [30] B. Zhang, Y. Al Dhuraibi, R. Rouvoy, F. Paraiso, and L. Seinturier, "Cloudgc: Recycling idle virtual machines in the cloud," in *2017 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 2017, pp. 105–115.
- [31] J.-W. Sung, S.-J. Han, and J.-w. Kim, "Virtual machine pre-provisioning for computation offloading service in edge cloud," in *2019 IEEE 12th International Conference on Cloud Computing (CLOUD)*. IEEE, 2019, pp. 490–492.
- [32] R. W. Ahmad, A. Gani, S. H. A. Hamid, M. Shiraz, A. Yousafzai, and F. Xia, "A survey on virtual machine migration and server consolidation frameworks for cloud data centers," *Journal of network and computer applications*, vol. 52, pp. 11–25, 2015.
- [33] M. Noshay, A. Ibrahim, and H. A. Ali, "Optimization of live virtual machine migration in cloud computing: A survey and future directions," *Journal of Network and Computer Applications*, vol. 110, pp. 1–10, 2018.
- [34] T. He, A. N. Toosi, and R. Buyya, "Performance evaluation of live virtual machine migration in sdn-enabled cloud data centers," *Journal of Parallel and Distributed Computing*, vol. 131, pp. 55–68, 2019.
- [35] F. Zhang, G. Liu, X. Fu, and R. Yahyapour, "A survey on virtual machine migration: Challenges, techniques, and open issues," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 2, pp. 1206–1243, 2018.
- [36] Q. Zhang, L. Liu, C. Pu, Q. Dou, L. Wu, and W. Zhou, "A comparative study of containers and virtual machines in big data environment," in *2018 IEEE 11th International Conference on Cloud Computing (CLOUD)*. IEEE, 2018, pp. 178–185.
- [37] A. Chung, J. W. Park, and G. R. Ganger, "Stratus: cost-aware container scheduling in the public cloud," in *Proceedings of the ACM Symposium on Cloud Computing*. ACM, 2018, pp. 121–134.
- [38] F. Jiang, K. Ferriter, and C. Castillo, "Pivot: Cost-aware scheduling of data-intensive applications in a cloud-agnostic system."
- [39] Y. Hu, H. Zhou, C. de Laat, and Z. Zhao, "Ecsched: Efficient container scheduling on heterogeneous clusters," in *European Conference on Parallel Processing*. Springer, 2018, pp. 365–377.
- [40] P. Thinakaran, J. Raj, B. Sharma, M. T. Kandemir, and C. R. Das, "The curious case of container orchestration and scheduling in gpu-based datacenters," in *Proceedings of the ACM Symposium on Cloud Computing*. ACM, 2018, pp. 524–524.
- [41] W. Zheng, M. Tynes, H. Gorelick, Y. Mao, L. Cheng, and Y. Hou, "Flowcon: Elastic flow configuration for containerized deep learning applications," in *Proceedings of the 48th International Conference on Parallel Processing*, ser. ICPP 2019. New York, NY, USA: ACM, 2019, pp. 87:1–87:10.
- [42] W. Xiao, R. Bhardwaj, R. Ramjee, M. Sivathanu, N. Kwatra, Z. Han, P. Patel, X. Peng, H. Zhao, Q. Zhang *et al.*, "Gandiva: Introspective cluster scheduling for deep learning," in *13th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 18)*, 2018, pp. 595–610.
- [43] Y. Fu, S. Zhang, J. Terrero, Y. Mao, G. Liu, S. Li, and D. Tao, "Progress-based container scheduling for short-lived applications in a kubernetes cluster," in *2019 IEEE International Conference on Big Data (Big Data)*. IEEE, 2018.
- [44] H. Zhang, L. Stafman, A. Or, and M. J. Freedman, "Slaq: quality-driven scheduling for distributed machine learning," in *Proceedings of the 2017 Symposium on Cloud Computing*, 2017, pp. 390–404.
- [45] "Variational auto-encoders," <http://kvfrans.com/variational-autoencoders-explained/>.
- [46] "Bidirectional recurrent neural networks," https://en.wikipedia.org/wiki/Bidirectional_recurrent_neural_networks.
- [47] "Gated recurrent unit," https://en.wikipedia.org/wiki/Gated_recurrent_unit.
- [48] "Tensorflow," <https://www.tensorflow.org/>.
- [49] "Pytorch," <https://pytorch.org/>.
- [50] "Persistent volume," <https://kubernetes.io/docs/concepts/storage/persistent-volumes/>.
- [51] "Nsf cloudlab," <https://www.cloudlab.us/>.