

ANDREAS: Artificial intelligence training scheduler for accelerated resource clusters

Federica Filippini, Danilo Ardagna,
Marco Lattuada, Edoardo Amaldi

*Dipartimento di Elettronica, Informazione e Bioingegneria
Politecnico di Milano, Italy
name.lastname@polimi.it*

Michele Ciavotta

*Dipartimento di Informatica, Sistemistica e Comunicazione
Università degli studi di Milano-Bicocca, Italy
michele.ciavotta@unimib.it*

Maciek Riedl, Katarzyna Materka, Paweł Skrzypek
7bulls, Poland

{mriedl, kmaterka, pskrzypek}@7bulls.com

Fabrizio Magugliani, Marco Cicala
E4 Computer Engineering, Italy

{fabrizio.magugliani, marco.cicala}@e4company.com

Abstract—Artificial Intelligence (AI) and Deep Learning (DL) algorithms are currently applied to a wide range of products and solutions. DL training jobs are highly resource demanding and they experience great benefits when exploiting AI accelerators (e.g., GPUs). However, the effective management of GPU-powered clusters comes with great challenges. Among these, efficient scheduling and resource allocation solutions are crucial to maximize performance and minimize Data Centers operational costs. In this paper we propose ANDREAS, an advanced scheduling solution that tackles these problems jointly, aiming at optimizing DL training runtime workloads and their energy consumption in accelerated clusters. Experiments based on simulation demonstrate that we can achieve a cost reduction between 30 and 62% on average with respect to first-principle methods while the validation on a real cluster shows a worst case deviation below 13% between actual and predicted costs, proving the effectiveness of ANDREAS solution in practical scenarios.

Index Terms—Energy-aware hardware platforms; Scheduling; Deep Learning

I. INTRODUCTION

Today, artificial intelligence (AI) and deep learning (DL) methods are exploited in a wide range of products and solutions. DL models are usually trained on AI accelerators (e.g., GPU-powered clusters) achieving speed-up factors in the range of 5-40x compared to the CPU-only counterpart [12].

The ability to optimize the infrastructure utilization and process the workload with high efficiency under power constraints is critical in this context. Indeed, local operators/data centers/Cloud Service Providers (CSPs) are subject to power consumption quotas. Consequently, their financial results are strongly dependent on how efficiently they exploit the infrastructure. While there are advanced solutions to manage virtual workloads (viz. virtual machines and containers) [3], the ever-growing DL training workload demand on GPUs represents a business opportunity for cloud and data centers operators. However, opti-

mizing revenue requires keeping power consumption under control while maximizing the effectiveness of high-value assets like GPU-powered systems. The ability to manage the energy footprint in the design of new infrastructure, like distributed computing with GPUs located at the edge in 5G networks, is critical to telecommunication companies owning their infrastructure, to service companies managing infrastructure, and to SMEs developing and providing innovative solutions [13].

Despite the clear advantages of GPU-accelerated clusters in terms of performance, these systems are characterized by a large capital and operating expenditures and a significant energy footprint [1], [6].

This paper introduces ANDREAS, an advanced scheduling solution aimed at optimizing DL training run-time workloads and their energy consumption in accelerated clusters. Our framework, working in an online setting where jobs are expected to be continuously submitted by users in the form of Docker images, implements both a profiling system able to estimate the job processing time and a Random Greedy algorithm able to efficiently and effectively assign at run-time the GPUs of a cluster to the jobs considering energy costs and penalties due to delays in job completions. We compare our solution with first-principle approaches commonly used in online scheduling. The experiments performed in a simulated environment show that our method guarantees a cost reduction between 30 and 62% on average with respect to first-principle methods, while the validation on a real cluster highlights a worst case deviation between real and predicted costs below 13%, proving the effectiveness of ANDREAS solution in practical scenarios.

The remainder of this paper is organized as follows. In Section II, other literature proposals are reviewed. Section III introduces the main elements of the ANDREAS architecture while Section IV formally defines the problem and describes in detail the proposed algorithm. Section V compares ANDREAS with first-principle methods. Con-

clusions are finally drawn in Section VI.

II. RELATED WORK

The effective management of GPU-accelerated clusters comes with great challenges. Among these, designing efficient scheduling and resource allocation solutions is crucial to maximize performance and minimize operational costs of data centers. Even if an increasing effort has been put in recent years in tackling both problems, they are often addressed separately, focusing either on jobs scheduling or on resource allocation. In this section, we briefly review previous work in both scenarios.

A resource allocation strategy for DL training jobs is proposed in [16]. Authors determine the optimal jobs batch size according to their scaling efficiency and combine a FIFO scheduling with a dynamic programming algorithm to determine an effective solution. Similarly, *Optimus* [14] is a Kubernetes scheduler designed to manage DL jobs on a shared distributed containerized environment, which minimizes jobs training time by relying on online resource-performance models. *DL²* [15], also based on Kubernetes, dynamically allocates resources by exploiting Deep Reinforcement Learning instead of relying on analytical performance models and job profiling.

On the scheduling side, *Gandiva* [18] focuses on tasks recurrent behaviours during the sequence of mini-batches iterations and proposes a preemptive, time-sharing scheduler for DL jobs.

GPU cluster managers closer to ANDREAS are *Tire-sias* [8], *Philly* [11], and *FfdL* [10], that implement integrated frameworks supporting the full training process and dealing with jobs submissions and execution. However, they focus only on the scheduling of DL training jobs, while the resources to be assigned to all applications are specified by the users.

III. ANDREAS ARCHITECTURE

The high-level architecture of the ANDREAS framework is outlined in Figure 1. ANDREAS includes three main components: the *Job Manager*, the *Job Profiler* and the *Job Optimizer*. Training applications are submitted by users in form of Docker images. The Job Manager receives the users' requests and orchestrates the execution of the submitted jobs as follows. First of all, the Job Profiler is invoked to collect information about the expected per-epoch execution time of the submitted jobs on the different resources available in the system. One node (characterized by a given type and number of GPUs) is dedicated to collect these data, which are then stored in MariaDB. If the cluster is heterogeneous (i.e., it includes nodes with different CPU or GPU types), multiple nodes can be engaged in profiling activities; the profiling configuration can even be changed dynamically while production applications run.

The collected profiling data, together with a description of the system, with all available nodes, are provided to the

Job Optimizer, which is in charge of selecting the optimal deployment for the submitted jobs. In particular, the Job Optimizer exploits the algorithm described in the following section to determine which resource configuration should be used to deploy and run the different jobs. If the resources available in the system are insufficient to run all jobs concurrently, some of them are preempted, and their execution is postponed.

The Job Optimizer is invoked by the Job Manager periodically, or in reaction to re-scheduling events (viz., a new job submission or completion). The Job Manager deploys the jobs according to the optimization results, migrating the ones in progress among different GPUs if necessary.

IV. ANDREAS OPTIMIZER

This section presents the main ANDREAS assumptions and describes the problem; then, it focuses on the Randomized Greedy algorithm we propose to solve it.

A. Assumptions and Problem Description

This section addresses the capacity allocation problem solved by the Job Optimizer at every rescheduling point along with the main underpinning assumptions. The viewpoint of a CSP running a data center is considered. The data center is assumed to be composed by a set of heterogeneous nodes denoted by $N = \{n_1, n_2, \dots, n_N\}$, with $N = |N|$. Each node is characterized by a specific type and number of available GPUs. The number of GPUs available on node n is identified by a set $G_n = \{1, 2, \dots, G_n\}$. The node is characterized also by a cost per unit time c_{ng} , when g GPUs are in use. GPU homogeneity, however, is not preserved across different nodes.

Jobs are submitted continuously over time by several users, are profiled and executed by ANDREAS on a suitable node. Each node $n \in N$ can host multiple virtualized workloads, assigning in a dedicated fashion the available GPUs to jobs; a job, in turn, runs on a single node. In addition, since the number of available resources is limited, jobs that cannot be accommodated in the system are entered into a waiting set until the next rescheduling point. Let \mathcal{J} represent the set of submitted jobs. Each job $j \in \mathcal{J}$ is characterized by a due date d_j , a tardiness penalty p_j , and a penalty to pay for postponing a job β_j . Furthermore, each job is associated with an execution time t_{jng} that depends on the job, the number of GPUs g and their type, which in turn depends on the considered node n . The parameters and variables of the problem that are considered are reported in Table I.

For each rescheduling point, the problem to be solved is a capacity allocation problem where for each job $j \in \mathcal{J}$ it must be decided whether to execute it right away or to postpone it. In the former case, the system also has to determine in which node to host the job and the resources (GPUs) to assign to its execution. If jobs are preempted, they will restart from the last snapshot (under




Figure 1: ANDREAS framework. In the example, we consider five submitted jobs j_1, j_2, j_3, j_4, j_5 . Nodes n_1 to n_N are equipped with a variable number of GPUs, of possibly different type, and have different energy costs. Jobs j_1, j_2 and j_5 are deployed on n_1 , running on 1 GPU and 2 GPUs, respectively. The other jobs are sent back to the queue and no other nodes are selected.

Table I: Problem parameters and variables

Parameters	
J	set of submitted jobs
N	set of available nodes
G_n	set of available GPUs in each node $n \in N$
G_n	number of available GPUs in each node $n \in N$
c_{ng}	energy cost of each node $n \in N$ when g GPUs are used
d_j	due date of job $j \in J$
t_j	tardiness weight of job $j \in J$
t_{jng}	execution time of job $j \in J$ when running on g GPUs on node $n \in N$
	penalty coefficient for postponed jobs
Variables	
w_n	1 if node $n \in N$ is chosen
\hat{t}_j	tardiness of job $j \in J$
\tilde{t}_j	worst-case tardiness of job $j \in J$
j_n	operation cost of job $j \in J$ on node $n \in N$
j_n	1 if job $j \in J$ is the first ending job on node $n \in N$

the assumption that a snapshot of the DL model weight is taken every few epochs).

The primary objective of ANDREAS is to minimize costs, including tardiness-caused penalties and energy related execution costs (including, e.g., air conditioning, UPS overhead, etc.). Nonetheless, the Job Optimizer at each rescheduling point considers a problem with a much narrower scope making it challenging to estimate the impact of decisions on the overall cost. For this reason, the following proxy function, denoted by f_{OBJ} , is considered:

$$\min_{j \in J} (t_j + \hat{t}_j) + \sum_{j \in J, n \in N} j_n j_n \quad (1)$$

In particular, the proxy is expressed by two terms: the first one represents the job tardiness and the worst-case tardiness penalty costs (where \hat{t}_j is the tardiness a job may incur in if it is postponed to the next period).

The second term, in turn, takes into account operation costs of the fastest jobs on all nodes; j_n is 1 if job j is the first-ending job on node n while j_n represents its execution cost. This function has been designed with the precise aim of penalizing configurations with delayed jobs and, at the same time, incentivizing solutions with short-duration jobs. In this way, ANDREAS can provide a lower average job latency. In addition, since the Job Optimizer is invoked every time the fastest job finishes its execution, such a proxy function tends to promote higher system responsiveness and better control over the quality of the overall schedule.

For a matter of space, the complete Mixed Integer Non-Linear Programming (MINLP) formulation of the problem is reported in Appendix A of [4].

B. ANDREAS Randomized Greedy algorithm

The randomized greedy heuristic that we have developed for the capacity allocation problem is based on the following criteria and assumptions.

- Jobs are selected according to their pressure, which is an index that measures how close they are to the due date when executed with the fastest configuration. For each job $j \in J$, the pressure Δ_j is:

$$j = T_c + \min_{g \in G_n} \{t_{jng}\} - d_j, \quad (2)$$

where T_c denotes the current scheduling time.

- The best configuration $(n, g) \in N \times G_n$ for each selected job is either (i) the cheapest configuration such that the job is executed before its due date, if such a configuration exists, or (ii) the fastest available configuration if, independently from the selected setup, it is not possible to execute the job before

Algorithm 1 Randomized construction procedure

```
1: function RANDOMIZED_CONSTRUCTION( $\mathcal{J}$ ,  $\Delta$ , MAXITRG)
2:   iter = 0
3:    $S_{\text{best}}$  = empty schedule
4:   while iter < MAXITRG do
5:      $S$  = empty schedule            $S$ : current schedule
6:      $\mathcal{J}_s$  = SORT_JOBS_LIST( $\mathcal{J}$ ,  $\Delta$ )   : pressures of all jobs
7:     for all  $j \in \mathcal{J}_s$  do            $\mathcal{J}_s$ : sorted queue
8:        $D_j = \{(n, g) \text{ s.t. } t_{jng} + T_c < d_j\}$ 
9:        $(n, g) = \text{SELECT\_BEST\_CONFIGURATION}(j, D_j)$ 
10:      assigned = ASSIGN( $j, (n, g), N_O$ )
11:      if not assigned then
12:        ASSIGN_TO_SUBOPTIMAL( $j, S, N_O$ )
13:      else
14:         $S = S \cup (j, n, g)$ 
15:      end if
16:       $\mathcal{J}_s = \mathcal{J}_s \setminus \{j\}$ 
17:    end for
18:    if  $f_{\text{OBJ}}(S) < f_{\text{OBJ}}(S_{\text{best}})$  then
19:       $S_{\text{best}} = S$ 
20:    end if
21:    iter = iter + 1
22:  end while
23:  return  $S_{\text{best}}$ 
24: end function
```

its due date. Since in our framework the time unit execution cost of a job on any available configuration is always lower than the penalty incurred if the job due date is violated, it is reasonable to choose the cheapest configuration as long as the due date can be met, while, if it is violated, it is more convenient to complete the job execution as fast as possible to reduce the corresponding penalty.

- Execution costs increase linearly with the number of GPUs, as demonstrated by GPU-based servers energy profiles [10], while the speedup of the processing time is sublinear in the number of GPUs, as observed in GPU-based application benchmarks [5].

The randomized greedy algorithm, which is summarized in Algorithm 1, receives as input the set of jobs \mathcal{J} in the queue, their pressures and a maximum number of iterations, and returns a solution specifying for every job whether to execute it and the corresponding configuration. At each iteration, a candidate solution S is constructed via randomized greedy choices and is stored in S_{best} if $f_{\text{OBJ}}(S) < f_{\text{OBJ}}(S_{\text{best}})$.

At line 6, the set \mathcal{J} of submitted jobs is sorted, producing as output a new set denoted by \mathcal{J}_s . Note that job j precedes job k in \mathcal{J}_s if and only if $\Delta_j > \Delta_k$, i.e., job j is more likely to violate its due date. As a first randomization step, some jobs in \mathcal{J}_s can be swapped, with probability inversely proportional to their tardiness weight.

For all jobs j in the sorted list \mathcal{J}_s , the set of configurations (i.e., pairs made by a node and a certain number of GPUs) such that the job can be executed within its due date (determined by the condition $t_{jng} + T_c < d_j$) is defined, in line 8, as:

$$D_j = \{(n, g) \mid N \times G_n : T_c + t_{jng} < d_j\}.$$

The set D_j is used to select the best configuration for job j , according to the following rules. Since one of the main contributions to the total cost of the schedule is the penalty for due date violations, we always try to select configurations that guarantee to complete the jobs without tardiness (or with the least possible delay, if there is no possibility to match the required due date). Therefore, we define the optimal configuration as:

$$(n, g) = \begin{cases} \operatorname{argmin}_{D_j^*} \{t_{jng}c_{ng}\} & \text{if } D_j = \\ \operatorname{argmin}_{n, N, g, G_n} \{t_{jng}\} & \text{otherwise.} \end{cases}$$

In the first case, the set D_j is not empty, which entails that there exist configurations (n, g) such that the job can be executed before its due date. In this case, (n, g) is identified by the cheapest configuration, i.e., by the element in D_j with minimum cost (identified by the product between the time-unit cost c_{ng} and the expected execution time t_{jng}). In the second case, the job due date is violated with any available configuration, therefore we select the one that guarantees the minimum expected execution time, in order to reduce the tardiness as much as possible. The algorithm explores the space of possible configuration by introducing some randomness in the the configuration selection process, where, instead of choosing the configuration according to the rule defined above, it selects as candidate configuration for job j one of the (n, g) with lower cost, with probability inversely proportional to the cost itself. Having denoted by (n, g) the selected configuration for job j , the assignment proceeds as follows:

- 1) The algorithm tries to assign job j to the required node n , provided that the number of available GPUs in n is enough to assign g of them to job j .
- 2) If this assignment is not feasible, it tries to assign job j to a suboptimal configuration available in an open node. Iteratively, it loops over the elements in D_j until it finds a configuration that fits in any node $n \in N_O$.

V. EXPERIMENTAL ANALYSIS

This section describes the results obtained through an experimental campaign, aiming at assessing the effectiveness of ANDREAS framework in real-case scenarios.

The experimental campaign involved heterogeneous Deep Learning jobs based on different network types, namely multi-layer LSTM [9], EfficientNet [17], and simple Convolutional Neural Networks [7], referred to in the following as ConvolutionNet. All the jobs were implemented relying on TensorFlow 2 and during the training model weights snapshots are taken every epoch. The Job Optimizer has been configured with $\text{MaxIt}_{\text{RG}} = 1,000$. All experiments described in the following sections have been performed relying on ARMIDA (ARM Infrastructure for the Development of Applications) infrastructure, a test & development 8-node cluster located at E4's premises. ARMIDA configuration is detailed in Appendix B of [4]. The types and numbers of GPUs available are summarized in Table II.

Table II: Configurations available in ARMIDA

Node ID	GPU type	# GPU
armida-04	Nvidia Tesla V100-PCIe 32GB	1
armida-05	Nvidia Tesla V100-PCIe 32GB	1
armida-06	Nvidia Tesla V100-PCIe 32GB	2
armida-07	Nvidia Tesla T4 (TURING)	1

A. Profiling Analysis

The profiling analysis described in this section has been performed to measure the power consumption of ARMIDA nodes and GPUs, as well as the expected execution times of the considered applications on all the configurations available in the reference hardware system. The profiling has been performed through ExaMon (Exascale Monitoring) a highly scalable framework for the performance and energy monitoring of HPC servers [2].

The power consumption of ARMIDA nodes allowed to compute the hourly energy costs (starting from the base cost of 0.172 EUR/KWh) used by the optimizer to determine the best assignment for the different jobs, according to the different configurations. To evaluate also the energy consumption due to the air conditioning system, the UPS and other overheads of a data center environment, the energy costs has been inflated by considering a PUE (Power Usage Effectiveness) [6] equal to 1.33, which is the value measured in the ARMIDA setup.

B. Simulations

The results that can be achieved by exploiting the ANDREAS optimizer at a larger scale have been evaluated through a set of simulations, comparing the outcome of the Randomized Greedy (RG) algorithm against first-principle methods as First In First Out (FIFO), Earliest Deadline First (EDF) and Priority Scheduling (PS), as it is done in other literature proposals [1].

Different instances of the experiments have been generated by randomly selecting a set of DL jobs with different characteristics in terms of batch sizes and number of required training epochs. We have considered multiple scenarios, involving a number N of nodes varying from 10 to 100, and a number of jobs $J = 10N$. Jobs inter-arrivals are characterized by a *mixed* rate, generated as described in [16]. The characteristics of nodes are similar to those reported in Table II. Specifically, half of them are equipped with GPUs of type TeslaV100, while the others with GPUs of type TURING. We have considered two different cases: in the first one, nodes are characterized by two TeslaV100 GPUs or by one TURING; in the second case nodes with four TeslaV100 or two TURING GPUs are considered. The parameter for penalizing jobs postponements has been set to 100.

The results of the RG algorithm and of FIFO, EDF and PS methods in the first scenario are reported in Figure 3. The first plot reports the expected energy costs of jobs execution, while the second shows the total cost, which is the sum of the energy cost plus the penalties due to jobs' due dates violations. A significant reduction, of around

62% on average, is obtained considering the total costs. Indeed, first-principle methods never change the configuration assigned to a job once it has been started, which entails that, if jobs with higher priority are submitted after long-running jobs with lower priority, they have to wait for them to be completed before receiving resources. The third plot shows the expected time required to completely execute the training of all the J jobs submitted to the system. The reduction guaranteed by the RG algorithm helps in minimizing the probability that jobs due dates are violated, consequently reducing the corresponding penalty costs. Finally, the last plot reports the time required to the optimizer to solve one instance of the problem, answering the call of the Job Manager providing the schedule of all jobs. Even if our RG algorithm requires a longer time than first-principle methods to provide a solution, this time remains always below 0.1 seconds, which is significantly lower than the time required to reconfigure the system.

The results of the RG algorithm and of FIFO, EDF and PS methods in the second scenario are reported in Figure 2. Again, it is possible to notice that the average cost reduction guaranteed by the RG algorithm is of almost 30%, while the time required to determine a solution is lower than 0.01 seconds on average.


C. Validation in the ARMIDA setup

The results provided by the ANDREAS framework have been validated in the ARMIDA setup by considering a prototype system including 8 different jobs. Their characteristics are summarized in Table V of Appendix B [4]. The inter-arrival time has been set to 1200s. This is an accelerated scenario, in terms of submission frequency and average execution times of applications. Nonetheless, this does not affect the results effectiveness since the workload assigned to the available nodes is comparable to practical situations [16].


The system is composed by three nodes, which correspond to armida-05, armida-06 and armida-07 as described in Table II; the node armida-04 is not considered, since it is dedicated to profiling. The aim of this experiment was to assess the accuracy of our framework, by measuring the gap between the predicted energy costs and the actual costs measured in the ARMIDA nodes.

The sequence of jobs executions, with the relative configurations, is reported in Figure 4 of Appendix B in [4]. It is possible to observe both processes that characterize the decisions of the RG method: first of all, GPU sharing, that can be observed for instance when J10 is submitted and it is assigned to armida-06 together with J6. Then, preemption, that comes into place, for instance, when J2 is submitted: since it has a stricter due date than J7, it is immediately deployed on armida-07, while the execution of the latter is postponed to a future scheduling point.

Figure 5 in Appendix B of [4] reports the power consumption of ARMIDA nodes during the full experiment session. This has been used to compute the actual energy



(a) Energy costs and total costs



(b) Simulation time and avg. execution time per optimizer call

Figure 2: Average costs and execution times with the different methods - first scenario

(a) Energy costs and total costs

(b) Simulation time and avg. execution time per optimizer call

Figure 3: Average costs and execution times with the different methods - second scenario

costs incurred by the system during the whole experiment, including the reconfiguration costs that are not considered by our simulator. We have compared this result with our prediction, noticing that, in the accelerated framework we have considered, all jobs are completed within the due dates. Table III reports the deviations of the predicted cost with respect to the real cost (the energy cost is equal to the total cost, since no due date violations are observed). It is relevant to notice that the predicted cost is higher than the real cost, which makes our framework more reliable and validates its adoption in an industry setting.

Table III: Results of the validation on ARMIDA

	Energy cost
Real cost [€]	0.95111
Predicted cost [€]	1.08438
Deviation	12.29%
Average deviation for each optimizer call	10.81%

VI. CONCLUSIONS AND FUTURE WORK

The ANDREAS scheduling solution proposed in this work effectively tackles the scheduling and resource allocation problems in GPU-powered clusters, achieving a cost reduction between 30 and 62% on average with respect to first-principle methods. The deviation between the predicted and the real costs in a validation cluster is below 13%, supporting its adoption in an industry setting.

Future work will include the extension of our scheduling framework to disaggregated hardware architectures.

ACKNOWLEDGEMENTS

This work has been partially funded by the European Union’s Horizon 2020 research and innovation programme under the TETRAMAX grant agreement no 761349.

REFERENCES

- [1] Amaral, M., Polo, J., Carrera, D., Seelam, S., Steinder, M.: Topology-aware gpu scheduling for learning workloads in cloud environments. In: HPCNSA Proc. ACM (2017)
- [2] Bartolini, A., Beneventi, F., et al.: Paving the way toward energy-aware and automated datacentre. ICPP 2019
- [3] Bernstein, D.: Containers and cloud: From lxc to docker to kubernetes. IEEE Cloud Computing 1(3), 81–84 (2014)
- [4] Filippini, F., et al.: Andreas: Artificial intelligence training scheduler for accelerated resource clusters. ArXiv [abs/2105.05080](https://arxiv.org/abs/2105.05080) (2021)
- [5] Gianniti, E., Zhang, L., Ardagna, D.: Performance prediction of gpu-based deep learning applications. In: CLOSER (2019)
- [6] Google: E ciency in data centers. [Online]. Available: https://www.google.com/about/datacenters/e_ciency/
- [7] Gu, J., Wang, Z., et al.: Recent advances in convolutional neural networks. Pattern Recognition 77, 354–377 (2018)
- [8] Gu, J., et al.: Tiresias: A GPU cluster manager for distributed deep learning. In: 16th USENIX NSDI Symposium (2019)
- [9] Hochreiter, S., Schmidhuber, J.: Long short-term memory. Neural computation 9, 1735–80 (12 1997)
- [10] Jayaram, K.R., et al.: Ffdl. Proceedings of the 20th International Middleware Conference (2019)
- [11] Jeon, M., et al.: Analysis of large-scale multi-tenant GPU clusters for DNN training workloads. In: 2019 USENIX ATC
- [12] Madougou, S., Varbanescu, A., de Laat, C., van Nieuwpoort, R.: The landscape of GPGPU performance modeling tools. PC 56, 18–33 (2016)
- [13] Mishra, S., et al.: Energy-e cient deployment of edge datacenters for mobile clouds in sustainable iot. IEEE Access 6, 56587–56597 (2018)
- [14] Peng, Y., Bao, Y., Chen, Y., Wu, C., Guo, C.: Optimus: An e cient dynamic resource scheduler for deep learning clusters. In: EuroSys (2018)
- [15] Peng, Y., et al.: DI2: A deep learning-driven scheduler for deep learning clusters. ArXiv [abs/1909.06040](https://arxiv.org/abs/1909.06040) (2019)
- [16] Saxena, V., Jayaram, K.R., Basu, S., Sabharwal, Y., Verma, A.: E ective elastic scaling of deep learning workloads (2020)
- [17] Tan, M., Le, Q.V.: E cientnet: Rethinking model scaling for convolutional neural networks. CoRR [abs/1905.11946](https://arxiv.org/abs/1905.11946) (2019)
- [18] Xiao, W., Bhardwaj, R., et al.: Gandiva: Introspective cluster scheduling for deep learning. In: USENIX OSDI (2018)