

# DeepPlace: Learning to Place Applications in Multi-Tenant Clusters

Subrata Mitra  
Adobe Research

Shanka Subhra Mondal\*  
IIT Kharagpur

Nikhil Sheoran  
Adobe Research

Neeraj Dhake\*  
IIT Bombay

Ravinder Nehra\*  
IIT Roorkee

Ramanuja Simha\*  
Oakridge National Lab

## Abstract

Large multi-tenant production clusters often have to handle a variety of jobs and applications with a variety of complex resource usage characteristics. It is non-trivial and non-optimal to manually create placement rules for scheduling that would decide which applications should co-locate. In this paper, we present **DeepPlace**, a scheduler that learns to exploits various temporal resource usage patterns of applications using Deep Reinforcement Learning (Deep RL) to reduce resource competition across jobs running in the same machine while at the same time optimizing for overall cluster utilization.

## ACM Reference Format:

Subrata Mitra, Shanka Subhra Mondal\*, Nikhil Sheoran, Neeraj Dhake\*, Ravinder Nehra\*, and Ramanuja Simha\*. 2019. DeepPlace: Learning to Place Applications in Multi-Tenant Clusters. In *10th ACM SIGOPS Asia-Pacific Workshop on Systems (APSys '19)*, August 19–20, 2019, Hangzhou, China. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3343737.3343741>

## 1 Introduction

Today, large production environments often need to handle a large variety of applications, including but not limited to interactive (user-facing) services, latency sensitive applications, batch analytics jobs, stream processing, iterative computations, maintenance services, etc. The standard practice today is to deploy these applications as containers which are then managed by various container orchestration engines

---

### \* Work done while at Adobe Research.

---

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](mailto:permissions@acm.org). *APSys '19*, August 19–20, 2019, Hangzhou, China  
© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6893-3/19/08...\$15.00  
<https://doi.org/10.1145/3343737.3343741>

such as Docker-Swarm [16], YARN [43], Mesos [15], or Kubernetes [5]. These orchestration engines allocate resources (e.g., CPU and memory) to these jobs according to the estimated *resource limits* provided by the developers [5]. In a multi-tenant shared cluster, if multiple applications compete for the same shared resources, they slow each other down due to resource contention [24, 25, 45]. Thus to reduce the chances of contention, orchestration engines use developer specified *affinity*, and *anti-affinity* [5, 9] rules to place applications on different machines. For stateless-services, resource estimates can be a bit aggressive, such that the resources allocated to each of the deployed containers would be enough to make it run smoothly, while the load fluctuations can be handled through an autoscaling mechanism by increasing or decreasing the number of deployed containers on the fly. For stateful-services, autoscaling can be really tricky and reactive migration of containers across machines have high overheads [9]. Hence, the containers are usually deployed with very conservative estimates by specifying large *resource limits* so that they can sustain phases with substantial increase in the resource demands. However, periods with such high resource usages are rare and often span only a very short fraction of the life-cycle of application, leading to resource wastage during the comparatively idle times.

In most of the real production systems, not all the applications would require to use the peak resource at the same time, and not all phases of their execution would contend for resources in a similar manner [1, 37, 45]. For example, Figure 1 shows the CPU usage characteristics of three production services in our cluster, across various phases of their execution. Their temporal resource usages show several "*peaks*" and "*valleys*", some more regular than the others. For user-facing services, such temporal resource usages can have daily and seasonal patterns due to fluctuations in user-demands [30] (e.g. some services are mostly used during working hours, while some services are mostly used during major holidays seasons). Applications can have different resource usage patterns across algorithmic phases [31, 45, 48], e.g., between the map and reduce phases in map-reduce jobs. Because of these temporal variations in the resource usage, *deploying for peak* using developer-provided limits is inefficient from an overall




Figure 1: Examples of resource usages by three production services with several "peaks" and "valleys".

resource utilization perspective. The variety of applications and the complexity of their temporal resource usage patterns makes it infeasible for the developers to express the placement logic in terms of existing placement rules available in current schedulers, e.g., affinity and anti-affinity rules in Kubernetes [2].

Borg [44] partially addresses this problem by packing a mix of high and low priority jobs in each machine, so that high priority jobs can expand during load spikes whereas low priority jobs can take advantage during the idle periods of the high priority jobs. However, not all clusters see such a health mix of low priority jobs to effectively fill the valleys of the high priority jobs.

Along with the temporal usage patterns, some jobs might have dependent succeeding jobs that rely on the completion of the first job. These dependencies can be intra or inter services. For example, a customer might have a nightly recommendation model builder, post completion of which a service kicks in to generate a new set of recommendations. A job scheduler that is aware of such dependencies can further utilize this information to efficiently schedule the existing jobs while making room for the upcoming jobs. A central scheduler can even discover *serendipitous dependencies* between different jobs coming from completely different developer groups, opening up scopes for resource alignment among these jobs leading to improved utilization of the cluster.

In this paper, we introduce an early prototype of **DeepPlace**, a self-learning scheduler that can opportunistically place containerized applications such that their temporal resource usages are aligned, resource contentions are minimized, quality of service is maintained and overall utilization improved. **DeepPlace** uses deep reinforcement learning (Deep RL) to learn hidden patterns from historical data over time to improve its scheduling policy. Essentially, **DeepPlace** treats resource usages of the applications as a *multivariate timeseries* and learns how these timeseries can be placed across different machines so that their resource usages are better aligned. We show through some example cases, how **DeepPlace** can take non-trivial decisions by anticipating future placement requests in order to optimize the overall resource usage in the cluster. For stateful-services, **DeepPlace** helps by minimizing the chances of resource contention, without being overly conservative, leading to operational excellence. For stateless-services, the need for

*scaling-up* can be reduced by having a better placement to begin with. With a better placement, a small number of containers might be able to gracefully handle the load up to a certain extent without a need for scaling-up. However, when scaling-up does happen, *where to place* those new additional containers is another crucial question, as usually in container scaling new machines are not spawn off frequently, that can be answered by **DeepPlace**.

## 2 Background and Related Work

**Reinforcement Learning.** In RL, at a high-level, an agent interacts with a system and tries to learn an optimized policy. At each timestep  $t$ , the agent observes the *state* of the system  $s_t$ , and chooses to take an *action*  $a_t$  that changes the state to  $s_{t+1}$  at timestep  $t + 1$ , and the agent receives a reward  $r_t$ . The agent tries to maximize the received *reward* which would help it to learn an optimized policy. It is assumed that the state transitions and rewards are stochastic and the state transition probabilities and rewards depend only on the state of the environment  $s_t$  and the action taken by the agent  $a_t$  (i.e., show Markov property [41]).

The objective is to maximize the expected cumulative discounted reward:  $E[\sum_{t=0}^{\infty} \gamma^t r_t]$  where  $\gamma \in (0, 1)$  determines how much the future rewards contribute to the total reward. More details of theoretical background of RL can be found in [41] and [26]. Inspired by the recent trends in Deep Reinforcement Learning (DeepRL), in this paper, we use deep-neural-networks (DNNs) as a function approximator for the placement policy that **DeepPlace** wants to learn. The RL algorithm can perform gradient-descent on the parameters of this DNN so that it can maximize the expected cumulative discounted reward over the actions the RL-agent takes. The gradients are estimated by observing the trajectories of execution that are obtained by following the policy.

RL has been used in variety of scenarios including learning complex games [12, 20, 32], robotics [18, 19, 33], and very recently for video streaming [27], routing [4, 28, 42], device placement [29]. But, the application of RL to self-learning schedulers has not been thoroughly explored.

**Scheduling.** To the best of our knowledge, recently proposed DeepRM [26] is the only other self-learning scheduler that also attempts to learn novel scheduling policy using DeepRL. DeepRM has a very simplified view of the cluster and thus comes with several limitations.

- (1) DeepRM assumes that jobs will always take a fixed amount of resources. It does not capture their temporal variations. Users often overestimate resource requirements and there can be a significant difference in resource usage between a peak-load and off-peak loads (Figure 1). Thus, ignoring such temporal variations and using the user-specified resource-limits for placement is wasteful and leads to low utilization.
- (2) DeepRM models the resource capacity of the compute cluster as a single monolithic block. It does not have a machine specific view and during its scheduling decisions, it does not try to optimize for the set of job or services to be run together to avoid resource contention.
- (3) The single monolithic view of the total resource capacity of the cluster ignores the impacts of resource fragmentation (i.e., the total amount of available resource in the cluster is more than the requirement of a job, but no single machine has that much available resources left.)

Tetris [13] is another heuristic-based cluster scheduler that takes into account multiple resource dimensions as well as the alignment of the machine’s remaining usage with the job’s requirement for packing jobs to the machines.

A large body of work has focused on scheduling data-driven applications, long-running user-facing services, ML-services, etc. on multi-tenant commodity clusters covering various aspects such as fairness of resource sharing [10, 11, 13, 17, 35, 36], tail-latency optimization [8, 14, 21, 23, 38, 40, 46] and how to protect latency sensitive application while improving cluster utilization [3, 6, 7, 22, 34, 39, 44, 45, 47]. These are distinct from our work as none of these scheduler attempt to *self-learn* the best scheduling strategy by discovering hidden resource usage characteristics and dependence among applications, along multiple resource dimensions. However, some of the proposed techniques (e.g. cycle-per-instruction [49]) can be used with our technique to further fine-tune reward/penalty design.

### 3 Design

We now describe the design of **DeepPlace** explaining how it operates. **DeepPlace** observes temporal job behavior to optimize its policy, encoded in a DNN-based policy network, using RL. **DeepPlace** models the scheduling problem as an RL-environment where the compute cluster is composed of  $N$  machines on which the application services or jobs are to be scheduled. Each such machine has  $C_d$  amount of total physical resource capacity for resource dimension  $d$  (e.g., CPU, Memory, etc.). For a job or service  $j$ , **DeepPlace** observes the time-series of the resource usages denoted as  $r_{d_j}(t)$ , where  $r_d$  is the resource usage along the resource dimension  $d$ . **DeepPlace** also keeps track of the current placement map of which services or jobs are running on which machines as well as what are the incoming services or jobs that need to be scheduled in the cluster, as a queue. The purpose of the




Figure 2: Input space representation of DeepPlace




Figure 3: Workflow of DeepPlace

queue is to incorporate in the state representation, a view of the upcoming jobs thus allowing the scheduler to learn the arrival patterns and dependencies amongst the jobs. The complete workflow for **DeepPlace** is shown in Figure 3.

#### 3.1 State Space Representation

**DeepPlace**’s state space representation is inspired by [26]. Though DeepRM’s representation for scheduling is designed to answer: “*what job to schedule when*”, **DeepPlace** is designed primarily to answer: “*what job to schedule where*”. In extreme cases, **DeepPlace** can delay some scheduling decisions if no suitable placement exists. Thus, **DeepPlace** makes some key improvements in the input-space representation to capture the degree of competition for resources among the jobs sharing the same underlying resources of a machine and their temporal variations in resource usages.

Figure 2 illustrates the input-space representation.

- (1) State of each machine in the cluster is represented as a 2D matrix or an image with  $k \times C_d$  pixels for each of the resource dimension  $d$ , where  $k$  is the number of previous logical timesteps.
- (2) Within each machine, the vertical direction of the image (i.e., the matrix) represents the time axis and shows the utilization of jobs for up to  $k$  previous logical timesteps, and the horizontal direction represents the amount of resource used by each job/services (quantized into units of resources). This type of representation helps the DNN-based RL-agent to learn the temporal resource usage characteristics of each job.  $k$  is a configurable parameter that the user can choose. The value of  $k$  should be a number reasonably large enough w.r.t. scheduling time-scale so that it helps the agent to capture a significant overlap among applications as well as

temporal variations in the resource usage. However, larger  $k$  results in longer convergence time for the RL-agent.

- (3) For each machine, the number of pixels in the horizontal direction ( $C_d$ ) represents the resource capacity of that machine for resource dimension  $d$ .  $C_d$  is another configurable parameter that user can choose depending on the granularity of resource usage that needs to be tracked. Larger  $C_d$  results in longer convergence time for the RL-agent.
- (4) After each machine, there is a column representing the applications or jobs scheduled and waiting to be run in the machine. This representation is important for **DeepPlace** to take multiple decisions in the same logical timestep. Even if the machine representation is not showing the resource usage of the scheduled application (as time has not proceeded), the column will give an insight to the agent that in the next timestep the application will be running in that machine and hence helps in taking the next action in the same logical timestep.
- (5) The pixels of the image representing machine-states (i.e., the values in the matrix) are colored differently to denote how much of the available capacity of the machine is being used at what time by which job. To make **DeepPlace** scalable, we consider that **DeepPlace** will attempt to learn the characteristics of up to  $G$  types or equivalence-classes of applications, and each type of application is represented by a unique number between 0 and 1, both exclusive (this is analogous to a different color of the corresponding pixels in the image). The unused resources are marked with *white* color (or a value of 0 in the matrix). **DeepPlace** uses these colors (i.e. the numbers) to learn which type of applications when run together can potentially suffer from resource competition and for how long such competition might last.
- (6) There can be multiple instances of the same application type running in the same machine with an overlap in their duration (e.g. two instances of a face-detection service triggered by two different products). These different instances can potentially also create resource contention among themselves (e.g. when the application is highly CPU intensive) and therefore needs to be distinguished and captured by the RL-agent. We again assign different colors (i.e., floating point numbers) to each instance of the application that are unique but close-by within a small range to the original assigned color for that job type.
- (7) **DeepPlace** captures the state of individual machines and combines these machine-level state representations into a cluster-level state representation for creating a holistic input for the policy-network. **DeepPlace** does that by using a trick: (a) To clearly distinguish between the applications running in different machines, **DeepPlace** adds a different factor for different machines to the number (or color) assigned to the machine usage as well as the column containing the scheduled jobs. For example, if two instances of the

same type of application with assigned number representation 0.2, are running in two different machines (machine 1 and 2), then in the combined state-space representation, these tasks will be represented as **1.2** and **2.2** respectively.

- (8) Along with the combined representation of the machines, **DeepPlace** also keeps a *waiting-queue* in its state-space representation. This queue represents the tasks waiting to be scheduled. By observing the changes in the queue over time, the RL-agent learns some key dynamics about the arrival characteristics of the jobs, which type of and how many jobs come together, and the temporal dependency amongst them, as previously discussed.

### 3.2 Reward/Penalty Design

**DeepPlace** is driven by negative rewards (penalty) which has the following four components:

**Resource contention penalty.** To help **DeepPlace** learn a placement policy that results in better resource alignment (complementary) and avoid resource contention among tasks scheduled in the same machine we use a modified version of cross-correlation to penalize the RL-agent during its learning. Cross-correlation ( $Cr$ ) is calculated between all pairs of tasks  $i$  and  $j$  running on the same machine across resource dimension  $d$  as follows:

$$Cr(i, j, d) = \sum_{t=0}^{\min(T_i, T_j)} res\_usage(i, t, d) \times res\_usage(j, t, d)$$

where  $T_i$  is the length of task  $i$  and  $res\_usage(i, t, d)$  is the instantaneous resource demand across dimension  $d$  by task  $i$  at time  $t$ . Cross-correlation formula amplifies the effect of two peaks being scheduled together. The  $Cr$  for a particular state of the cluster is calculated by taking the sum of cross-correlation of each machine, which includes across all the resource dimensions (CPU or memory), the cross-correlation of each task with every other task in that machine.

**Resource over-utilization penalty.** To prevent scheduling of more tasks than that can be handled by a machine, there is a high penalty if the machine is not able to meet the resource requirement of tasks scheduled in that machine. It is calculated by adding a high negative factor each time a machine is unable to provide appropriate resources to the running tasks.

**Wait-time penalty.** To prevent the scheduler from holding jobs for a long time in search of a better place, we add a constant penalty proportional to the state of the waiting queue. It is equal to the number of waiting tasks in the queue multiplied by a negative constant at each time.

**Under-utilization penalty.** Since our goal is to improve overall utilization of the cluster by helping the scheduler learn how to achieve tighter packing and pack on less number of machines, if possible, we add a penalty proportional to the sum of unused resources in the *used machines*. White pixels

in our state-representations denote the number of unused resources at any given time.

## 4 Implementation

We use the modified version of REINFORCE algorithm as mentioned in [26]. The policy network consists of a single hidden layer of 20 neurons followed by output neurons equal to the number of actions (number of machines under consideration). We use a 36 core CPU server and python multiprocessing to create multiple workers (equal to  $batch\ size+1$ ) each operating on distinct examples, taking a fixed number of trajectories and accumulating gradients. The last worker is used to combine the gradients of each worker and send to the policy network for updating the parameters. This gives a major improvement in the training speed. The training time increases significantly as we increase the cluster load. It also depends significantly on the type of applications under consideration (For example, Long running vs Short running jobs). For the hidden layer, we use Relu activation function, while for the output layer we use softmax activation. We use Adam optimizer and a learning rate of 0.001. The number of trajectories taken by each worker is fixed at 20.

## 5 Evaluations

**[Workload.]** In our evaluation setup, jobs arrive online as a Poisson process. The average job arrival rate is calibrated to create three average cluster load scenarios: 30%, 50% and 80%. In our setup, 50% of the jobs are *long running* and the other half are *short running*. Each job has 2 dimensions of resource requirements: CPU and memory. The capacity of these two resources in each machine is denoted by  $\{1r, 1r\}$  For each job, dominant resource usage is randomly chosen to be either CPU or memory. The resource usage of the dominant resource is independently chosen from a uniform distribution between  $0.3r$  and  $0.5r$ . The non-dominant resource usage is also independently and uniformly varied between  $0.08r$  and  $0.16r$ . Thus there is no correlation between the CPU and memory usages. Temporal resource usage for each job varies as a square wave with period uniformly chosen between  $0.2t$  and  $0.5t$  and width as one-fourth of the period, where  $t$  denotes the job length. Total 50 such different jobs are used for training and 18 for testing. Our evaluation runs with a cluster of 10 machines.

**[Baselines.]** We compare **DeepPlace** with Tetris [13], which schedules jobs on machines based on how well job’s resource requirement aligns with the machine’s available resources balancing preferences for short jobs and packing in a combined score. We also compare it against *Best Fit* heuristic which allocates the job to the machine having the least units of the dominant resource of the job left.

*Note: It is not possible to directly compare DeepPlace with DeepRM [26] because DeepRM only specifies which job to be*

*scheduled next and does not say on which machine it should be scheduled. Thus DeepRM does not have any concepts of competition for resource usage among applications running in the same machine, resource fragmentation and machine-level over-utilization. Thus, a fair comparison with DeepRM with respect to our desired metrics is not possible.*

**[Learning Progress.]** We first show how **DeepPlace**’s learning converge across multiple iterations in Figure 4. It can be observed that roughly after 1000 iterations, **DeepPlace**’s policy learning starts to converge and does not see any further significant drop in the normalized penalty.

**[Improvement in Cluster Utilization.]** We measure average utilization of machines for each resource as:

$$\text{Avg Util} = \frac{\sum_{t=0}^T \text{Utilization across all machines at time } t}{T \times \max(\text{used machines}) \times \text{resource capacity}}$$

where  $T$  is the length of the observation period. Since the number of machines that are actually being actively used varies over time, in the denominator, we used maximum number of machines used at any point in time to normalize.

**[Comparing Scheduling Efficiency.]** Here in Figure 7, we show how **DeepPlace** optimizes for cluster utilization for both CPU and memory. We can see that **DeepPlace** can provide a 68-100% increase in average utilization compared to Tetris across different cluster-load conditions. This is primarily achieved by efficient packing that requires significantly less number of machines to be used compared to Tetris as shown in Figure 5. Further, it can be observed that the gap between **DeepPlace** and Tetris in terms of the number of machines required to accommodate the jobs increases with the increase of the cluster load. Although it looks like *BestFit* provides even higher utilization because it just packs the jobs into the machines without any knowledge of peak or future resource usages of the jobs and as a consequence, *BestFit* suffers from huge over-utilization of the resources as shown in Figure 6. On the other hand, over-utilization due to **DeepPlace**’s placement decisions are almost negligible. Tetris already includes peak resource usage information in its placement decision thus resulting in no resource over-utilization.

**[Improvement in Resource Fragmentation.]** Fragmentation score of a cluster at a high-level measure what part of all the available resource in a cluster are concentrated.

$$\text{Avg Frag} = 1 - \frac{\sum_{t=0}^T \max(\text{available space across all machines at } t)}{\text{Sum of available space over all machines at } t}$$

The lower the fragmentation score, the higher the ability of the cluster to schedule unanticipated large jobs. Hence, low resource fragmentation in the cluster is a desirable operational property. In Figure 8, we see **DeepPlace** provides 6-13% reduction in resource fragmentation compared to Tetris. **DeepPlace**’s intelligent placement which takes both temporal resource usage characteristics and job arrival patterns






Figure 4: Convergence of DeepPlace's training under 50% average load




Figure 5: Comparison of number of machines used






Figure 6: Comparison of over-utilization in the cluster



(a) CPU utilization

(b) Memory utilization


Figure 7: Comparison of average resource utilization in the cluster



(a) CPU fragmentation

(b) Memory fragmentation

Figure 8: Comparison of resource fragmentation level in the cluster



(a) Learned example 1

(b) Learned example 2

Figure 9: Examples of learned placement policies

leaves bigger room in the machines (i.e. less fragmentation score) to accommodate unanticipated large jobs.

## 6 Discussions

In this section, we discuss insights and applicability for real deployments.

**What DeepPlace learned?** Figure 9 illustrates how **DeepPlace** achieved better packing that ultimately resulted in higher overall utilization. Figure 9a shows how Job1 and Job2 were placed in the same machine because resource intensive parts of Job1 would finish before the resources are required by Job2. In Figure 9b, resource requirements for Job3 and Job4 alternate in such a manner that they do not exactly overlap with each other and thus placed in the same machine for better packing. All these patterns were learned by **DeepPlace** on its own without any guiding rules.

**Scheduling granularity for effectiveness.** **DeepPlace** looks at *where* to schedule an incoming application so that it can either improve the resource utilization or reduce the resource contention. However, *how often* such a placement decision needs to be made depends on the what kind of workload the cluster is handling. For a cluster handling short or medium-duration batch, cron or interactive

applications, frequent placement decisions need to be made and **DeepPlace** can be very useful. On the other hand, for long running services, typically new placement decisions are made less frequently, e.g., when the container for an upgraded service is being deployed, etc. However, if auto-scaling is enabled for these services, taking the decision on where the additional auto-scaled container should be placed in the cluster, can be suggested by **DeepPlace**.

**Cluster size.** Our input-space representation as well as action-space of the RL is proportional to the number of machines in the cluster. Hence, larger the size of the cluster, the more iterations and training examples it needs for its policy learning to converge.

**Bootstrapping learning in deployments.** **DeepPlace** uses historical time-series pattern of resource usages to learn what job is to be scheduled in which machine so that based on their resource usage characteristics, they either improve the overall utilization or avoid aggravating contention by using the same resource at the same time. If **DeepPlace** starts to learn from scratch, it can be long before it sees sufficient examples required for its learning to converge. An option to speed up learning by bootstrapping the RL-agent's policy is by replaying the time-series of historical resource usage through a simulation.

To conclude, in this paper we show an early design prototype of a self-learning scheduler that can exploit the temporal resource usage patterns and arrival dependencies of the jobs to provide a better placement policy and thus achieve better utilization without requiring any manually crafted rules or heuristics.

## References

- [1] AMVROSIADIS, G., PARK, J. W., GANGER, G. R., GIBSON, G. A., BASEMAN, E., AND DEBARDELEBEN, N. On the diversity of cluster workloads and its impact on research results. In *2018 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)* (2018), pp. 533–546.
- [2] AUTHORS, T. K. Affinity and anti-affinity. <https://kubernetes.io/docs/concepts/configuration/assign-pod-node/#affinity-and-anti-affinity>.
- [3] BOUTIN, E., EKANAYAKE, J., LIN, W., SHI, B., ZHOU, J., QIAN, Z., WU, M., AND ZHOU, L. Apollo: Scalable and coordinated scheduling for cloud-scale computing. In *OSDI* (2014), vol. 14, pp. 285–300.
- [4] BOYAN, J. A., AND LITTMAN, M. L. Packet routing in dynamically changing networks: A reinforcement learning approach. In *Advances in neural information processing systems* (1994), pp. 671–678.
- [5] BURNS, B., GRANT, B., OPPENHEIMER, D., BREWER, E., AND WILKES, J. Borg, omega, and kubernetes. *Queue* 14, 1 (2016), 10.
- [6] DELIMITROU, C., AND KOZYRAKIS, C. Paragon: Qos-aware scheduling for heterogeneous datacenters. In *ASPLOS* (2013).
- [7] DELIMITROU, C., AND KOZYRAKIS, C. Quasar: Resource-efficient and qos-aware cluster management. In *ASPLOS* (2014).
- [8] FERGUSON, A. D., BODIK, P., KANDULA, S., BOUTIN, E., AND FONSECA, R. Jockey: guaranteed job latency in data parallel clusters. In *Proceedings of the 7th ACM european conference on Computer Systems (Eurosys)* (2012), ACM, pp. 99–112.
- [9] GAREFALAKIS, P., KARANASOS, K., PIETZUCH, P. R., SURESH, A., AND RAO, S. Medea: scheduling of long running applications in shared production clusters. In *EuroSys* (2018), pp. 4–1.
- [10] GHODSI, A., ZAHARIA, M., HINDMAN, B., KONWINSKI, A., SHENKER, S., AND STOICA, I. Dominant resource fairness: Fair allocation of multiple resource types. In *Nsdi* (2011), vol. 11, pp. 24–24.
- [11] GHODSI, A., ZAHARIA, M., SHENKER, S., AND STOICA, I. Choosy: Max-min fair sharing for datacenter jobs with constraints. In *Proceedings of the 8th ACM European Conference on Computer Systems* (2013), ACM, pp. 365–378.
- [12] GIBNEY, E. Google ai algorithm masters ancient game of go. *Nature News* 529, 7587 (2016), 445.
- [13] GRANDL, R., ANANTHANARAYANAN, G., KANDULA, S., RAO, S., AND AKELLA, A. Multi-resource packing for cluster schedulers. *ACM SIGCOMM Computer Communication Review* 44, 4 (2015), 455–466.
- [14] HAQUE, M. E., HE, Y., ELNIKETY, S., BIANCHINI, R., MCKINLEY, K. S., ET AL. Few-to-many: Incremental parallelism for reducing tail latency in interactive services. In *ACM SIGPLAN Notices* (2015), vol. 50, ACM, pp. 161–175.
- [15] HINDMAN, B., KONWINSKI, A., ZAHARIA, M., GHODSI, A., JOSEPH, A. D., KATZ, R. H., SHENKER, S., AND STOICA, I. Mesos: A platform for fine-grained resource sharing in the data center. In *NSDI* (2011), vol. 11, pp. 22–22.
- [16] INC., D.
- [17] JOE-WONG, C., SEN, S., LAN, T., AND CHIANG, M. Multiresource allocation: Fairness-efficiency tradeoffs in a unifying framework. *IEEE/ACM Transactions on Networking (TON)* 21, 6 (2013), 1785–1798.
- [18] KAEHLING, L. P., LITTMAN, M. L., AND MOORE, A. W. Reinforcement learning: A survey. *Journal of artificial intelligence research* 4 (1996), 237–285.
- [19] KOBER, J., BAGNELL, J. A., AND PETERS, J. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research* 32, 11 (2013), 1238–1274.
- [20] LAMPLE, G., AND CHAPLOT, D. S. Playing fps games with deep reinforcement learning. In *AAAI* (2017), pp. 2140–2146.
- [21] LEVERICH, J., AND KOZYRAKIS, C. Reconciling high server utilization and sub-millisecond quality-of-service. In *EuroSys* (2014).
- [22] LO, D., CHENG, L., GOVINDARAJU, R., RANGANATHAN, P., AND KOZYRAKIS, C. Heracles: Improving resource efficiency at scale. In *ISCA* (2015).
- [23] MACE, J., BODIK, P., FONSECA, R., AND MUSUVATHI, M. Retro: Targeted resource management in multi-tenant distributed systems. In *NSDI* (2015), pp. 589–603.
- [24] MAJI, A. K., MITRA, S., AND BAGCHI, S. Ice: An integrated configuration engine for interference mitigation in cloud services. In *2015 IEEE International Conference on Autonomic Computing* (2015), IEEE, pp. 91–100.
- [25] MAJI, A. K., MITRA, S., ZHOU, B., BAGCHI, S., AND VERMA, A. Mitigating interference in cloud services by middleware reconfiguration. In *Proceedings of the 15th International Middleware Conference* (2014), ACM, pp. 277–288.
- [26] MAO, H., ALIZADEH, M., MENACHE, I., AND KANDULA, S. Resource management with deep reinforcement learning. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks* (2016), ACM, pp. 50–56.
- [27] MAO, H., NETRAVALI, R., AND ALIZADEH, M. Neural adaptive video streaming with pensieve. In *Proceedings of the ACM Special Interest Group on Data Communication* (2017), SIGCOMM '17.
- [28] MESTRES, A., RODRIGUEZ-NATAL, A., CARNER, J., BARLET-ROS, P., ALARCÓN, E., SOLÉ, M., MUNTÉS-MULERO, V., MEYER, D., BARKAI, S., HIBBETT, M. J., ET AL. Knowledge-defined networking. *ACM SIGCOMM Computer Communication Review* (2017).
- [29] MIRHOSEINI, A., PHAM, H., LE, Q. V., STEINER, B., LARSEN, R., ZHOU, Y., KUMAR, N., NOROUZI, M., BENGIO, S., AND DEAN, J. Device placement optimization with reinforcement learning. In *ICML* (2017).
- [30] MISHRA, A. K., HELLERSTEIN, J. L., CIRNE, W., AND DAS, C. R. Towards characterizing cloud backend workloads: insights from google compute clusters. *ACM SIGMETRICS Performance Evaluation Review* 37, 4 (2010), 34–41.
- [31] MITRA, S., GUPTA, M. K., MISAILOVIC, S., AND BAGCHI, S. Phase-aware optimization in approximate computing. In *2017 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)* (2017), IEEE, pp. 185–196.
- [32] MNIH, V., KAVUKCUOGLU, K., SILVER, D., GRAVES, A., ANTONOGLU, I., WIERSTRA, D., AND RIEDMILLER, M. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602* (2013).
- [33] MNIH, V., KAVUKCUOGLU, K., SILVER, D., RUSU, A. A., VENESS, J., BELLE-MARE, M. G., GRAVES, A., RIEDMILLER, M., FIDJELAND, A. K., OSTROVSKI, G., ET AL. Human-level control through deep reinforcement learning. *Nature* 518, 7540 (2015), 529.
- [34] NATHUJI, R., KANSAL, A., AND GHAFFARKHAH, A. Q-clouds: managing performance interference effects for qos-aware clouds. In *Proceedings of the 5th European conference on Computer systems* (2010), ACM, pp. 237–250.
- [35] PARKES, D. C., PROCACCIA, A. D., AND SHAH, N. Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation (TEAC)* 3, 1 (2015), 3.
- [36] POPA, L., KUMAR, G., CHOWDHURY, M., KRISHNAMURTHY, A., RATNASAMY, S., AND STOICA, I. Faircloud: sharing the network in cloud computing. *ACM SIGCOMM Computer Communication Review* 42, 4 (2012), 187–198.
- [37] REISS, C., TUMANOV, A., GANGER, G. R., KATZ, R. H., AND KOZUCH, M. A. Heterogeneity and dynamicity of clouds at scale: Google trace analysis. In *Proceedings of the Third ACM Symposium on Cloud Computing* (2012), ACM, p. 7.
- [38] REN, X., ANANTHANARAYANAN, G., WIERMAN, A., AND YU, M. Hopper: Decentralized speculation-aware cluster scheduling at scale. In *ACM SIGCOMM Computer Communication Review* (2015), vol. 45, ACM, pp. 379–392.
- [39] SCHWARZKOPF, M., KONWINSKI, A., ABD-EL-MALEK, M., AND WILKES, J. Omega: flexible, scalable schedulers for large compute clusters. In

- Proceedings of the 8th ACM European Conference on Computer Systems* (2013), ACM, pp. 351–364.
- [40] SURESH, P. L., CANINI, M., SCHMID, S., AND FELDMANN, A. C3: Cutting tail latency in cloud data stores via adaptive replica selection. In *12th USENIX Symposium on Networked Systems Design and Implementation* (2015), USENIX Association, pp. 513–527.
- [41] SUTTON, R. S., BARTO, A. G., ET AL. *Reinforcement learning: An introduction*. MIT press, 1998.
- [42] VALADARSKY, A., SCHAPIRA, M., SHAHAF, D., AND TAMAR, A. Learning to route. In *Proceedings of the 16th ACM Workshop on Hot Topics in Networks* (2017), ACM, pp. 185–191.
- [43] VAVILAPALLI, V. K., MURTHY, A. C., DOUGLAS, C., AGARWAL, S., KONAR, M., EVANS, R., GRAVES, T., LOWE, J., SHAH, H., SETH, S., ET AL. Apache hadoop yarn: Yet another resource negotiator. In *Proceedings of the 4th annual Symposium on Cloud Computing* (2013), ACM, p. 5.
- [44] VERMA, A., PEDROSA, L., KORUPOLU, M. R., OPPENHEIMER, D., TUNE, E., AND WILKES, J. Large-scale cluster management at Google with Borg. In *EuroSys* (2015).
- [45] XU, R., MITRA, S., RAHMAN, J., BAI, P., ZHOU, B., BRONEVETSKY, G., AND BAGCHI, S. Pythia: Improving datacenter utilization via precise contention prediction for multiple co-located workloads. In *Proceedings of the 19th International Middleware Conference* (2018), ACM, pp. 146–160.
- [46] XU, Y., MUSGRAVE, Z., NOBLE, B., AND BAILEY, M. Bobtail: Avoiding long tails in the cloud. In *NSDI* (2013).
- [47] YANG, H., BRESLOW, A., MARS, J., AND TANG, L. Bubble-flux: Precise online qos management for increased utilization in warehouse scale computers. In *International Symposium on Computer Architecture (ISCA)* (2013), pp. 607–618.
- [48] ZHANG, J., YOUSIF, M., CARPENTER, R., AND FIGUEIREDO, R. J. Application resource demand phase analysis and prediction in support of dynamic resource provisioning. In *Fourth International Conference on Autonomic Computing (ICAC'07)* (2007), IEEE, pp. 12–12.
- [49] ZHANG, X., TUNE, E., HAGMANN, R., JNAGAL, R., GOKHALE, V., AND WILKES, J. Cpi2: Cpu performance isolation for shared compute clusters. In *EuroSys* (2013).