Dynamic Colocation Policies with Reinforcement Learning

YUHAO LI, DAN SUN, and BENJAMIN C. LEE, Duke University

We draw on reinforcement learning frameworks to design and implement an adaptive controller for managing resource contention. During runtime, the controller observes the dynamic system conditions and optimizes control policies that satisfy latency targets yet improve server utilization. We evaluate a physical prototype that guarantees 95th percentile latencies for a search engine and improves server utilization by up to 70%, compared to exclusively reserving servers for interactive services, for varied batch workloads in machine learning.

CCS Concepts: • **Information systems applications** → **data centers**; • **Computer systems organization** → *Self-organizing autonomic computing*; Cloud computing; • **Reinforcement learning** → Sequential decision making;

Additional Key Words and Phrases: Resource contention, adaptive control, machine learning

ACM Reference format:

Yuhao Li, Dan Sun, and Benjamin C. Lee. 2020. Dynamic Colocation Policies with Reinforcement Learning. *ACM Trans. Archit. Code Optim.* 17, 1, Article 1 (March 2020), 25 pages.

<https://doi.org/10.1145/3375714>

1 INTRODUCTION

Most datacenter servers are under-utilized. Average server utilization is 30% to 50% in warehouse-scale datacenters and much less in typical datacenters [8, 11]. We can explain low utilization, in part, by the prevalence of interactive workloads. These workloads often compute alone, avoiding shared resource contention and meeting performance targets, but cannot fully utilize the server.

In this article, we study strategies for increasing server utilization by colocating interactive services with batched machine learning workloads. Machine learning has emerged as not only a driving force that enhances the user experience but also an important consumer of server resources in datacenters [16]. Machine learning workloads comprise a new breed of batch workloads with three unique characteristics and challenges.

First, machine learning algorithms are iterative with diverse, recurring computational phases. Colocation managers should monitor the batch workload to determine when and which resources should be allocated. This approach departs from prior work that focuses on monitoring and

This work was supported by National Science Foundation grants CNS-1822085 (IUCRC), CCF-1149252, CCF-1337215, SHF-1527610, and AF-1408784, and was partly supported by the Semiconductor Research Corporation's Global Research Collaboration (GRC) program under task 2821.001. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of these sponsors.

Authors' addresses: Y. Li and D. Sun, Duke University, 213 Hudson Hall, Durham, NC 27708; emails: {yuhao.li, dan.sun}@duke.edu; B. C. Lee, Duke University, 210 Hudson Hall, Box 90291, Durham, NC 27708; email: benjamin.c.lee@duke.edu.

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.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

1544-3566/2020/03-ART1

<https://doi.org/10.1145/3375714>

managing the interactive workload [21, 25, 26, 32, 46]. Indeed, we find that neglecting the batch dynamics leads to conservative colocation policies and poor server utilization.

Second, machine learning requires fine-grained, responsive colocation management. Computational phases span seconds, and when the workload transitions from phase to phase, the colocation manager should respond quickly. These short time scales differ from prior studies that manage pairwise colocations that persist until workloads complete [14, 24].

Finally, machine learning applications compute on diverse datasets, which determine memory intensity and shared resource contention. This diversity precludes offline profiling, popular in prior studies [14, 24, 29], because profiling every machine learning library for every dataset is intractable. Colocation managers should analyze dynamics within the batch workload online, updating models and policies as the server runs.

We address these challenges with reinforcement learning, which analyzes data to make a sequence of optimized decisions [38]. We use reinforcement learning to build a lightweight, online controller. The controller modulates batch computation to balance service targets and server utilization. It uses Q-learning to optimize its control policy as the server computes. The optimized policy assesses system conditions and dynamically scales the voltage and frequency for cores running machine learning workloads.

We show that reinforcement learning can respond quickly to severe, short-term variability in machine learning workloads that jeopardize interactive service targets. The colocation controller learns policies that balance short- and long-term objectives, meeting interactive performance targets and increasing server utilization. Moreover, the controller requires no offline profiling and learns to adapt its control policies as resource demands from batch workloads evolve. In summary, we make the following contributions:

- We characterize machine learning codes and identify colocation challenges that arise from their diverse computational phases and variable resource demands.
- We use reinforcement learning to determine, during runtime, optimal policies for setting the batch workload's processor core frequencies.
- We implement a controller that balances service quality and server utilization. It learns policies online and adapts quickly as workload demands evolve.
- We prototype the controller and find that it enforces a target on 95th percentile latency, improves server utilization by up to 70% compared to exclusively reserving servers for interactive service, and reduces datacenter energy consumption by 50% through colocation.

2 UNDERSTANDING BATCH DYNAMICS

We study datacenter servers that colocate interactive and batch machine learning workloads. As these workloads share server resources, software mechanisms, such as Linux Cgroup [2] and Docker containers [30], manage the allocation of cores, memory capacity, and network I/Os to each process. However, these software mechanisms are unable to regulate the competition for hardware-managed resources such as last-level cache capacity and main memory bandwidth. In this work, we modulate batch computation to manage contention in the shared memory subsystem and avoid interference with the interactive service.

Variability in batch workloads. Maintaining both service quality and server utilization is difficult because of workload dynamics. Researchers have characterized interactive workloads and their variations in service time, queue length, and query loads [21]. Such variability lengthens the tail distribution for query latency. However, there has been no corresponding study of workload dynamics for batch workloads. In this article, we find that distributed machine learning workloads

Table 1. Basic Data Operators in Machine Learning Workloads Implemented with Spark MLlib

Algorithm	Type	Data Operation
Linear Regression	Regression	Data sampling, gradient calculation, gradient aggregation
Support Vector Machine	Classification	Data sampling, gradient calculation, gradient aggregation
Decision Tree	Classification/Regression	Tree split (sort, reduction)
Gradient Boost Tree	Classification	Sampling, tree split (sort, reduction)
Naive Bayesian	Classification	Filtering, extracting features, data aggregation
Correlation	Statistics	Map, data aggregation

(a) Last-level cache misses for linear regression

(b) Memory transfers for linear regression

(c) 95th-percentile latency for search latency

Fig. 1. (a–c) Colocation dynamics as varying memory activity for batch workload affects tail latency for the interactive workload. Search engine and linear regression are colocated. Data collected every 50 ms.

exhibit diverse, recurring phases that cause contention in the last-level cache and memory channel to vary across time.

We survey machine learning algorithms used in production machine learning pipelines, as shown in Table 1. It includes two regression algorithms, three classification algorithms, and one statistical computation. We find that these algorithms are built upon basic data operations, which repeatedly operate on the input data and thus produce rich variations in machine analytics’ resource demand.

Figure 1 illustrates variations in batch machine learning and its effect on interactive service quality. We colocate Spark linear regression, a representative machine learning workload, with a search engine. Regression’s cache misses and memory transfers exhibit phase behavior with periodic spikes in activity. Variations in the regression workload produce variations in search query latency. The search engine fails to meet its target 95th percentile latency when memory activity spikes and the engine satisfies its target otherwise.

For this example, previously proposed colocation mechanisms are reactive and ineffective. They rely only on the interactive workload’s profile to throttle batch computation when interactive service quality is poor [21, 25, 26, 46]. Because these mechanisms lack knowledge of the batch workload’s phases and periodic demands for memory, they would conservatively halt batch computation when interactive latency spikes. Server utilization would suffer.

Phase behaviors in machine learning. We can understand oscillations between high and low memory activity by examining the structure of machine learning algorithms. Many algorithms rely on iterative training, which repeatedly updates model parameters until predictive accuracy for the input data reaches some target.

For example, Figure 2(b) presents phases that arise from iterative training in Linear Regression. In each iteration, Linear Regression performs three stages of computation. Data

(a) **Gradient Boosting Tree.** Phases: (I) data pre-processing; (II) data sampling to determine split candidates; (III) recursively choose best split for each sub partition.

(b) **Linear Regression.** Phases: (I) data sampling; (II) gradient calculation; (III) gradient aggregation.

(c) **Correlation.** Phases: (I) calculating coefficients (map operation); (II) aggregating feature values.

(d) **Naive Bayesian.** Phases: (I) data preprocessing; (II) filtering; (III) extracting features; (IV) aggregating values.

Fig. 2. (a–d) Computational phases for representative machine learning workloads. Data collected every 50 ms.

sampling suffers the highest cache miss rates because it accesses data and memory randomly. Gradient calculation also suffers high cache miss rates because it must multiply large matrices. Gradient aggregation is relatively compute intensive and generates little contentious behavior in the memory system.

In another example, Figure 2(b) presents phases from the iterative construction of a gradient boosting tree. Data sampling determines threshold candidates for splitting the dataset. Next, the algorithm recursively splits the data and constructs the decision tree. As the recursion progresses and the algorithm approaches the tree’s leaf nodes, the amount of data split decreases and the cache miss rate falls.

Some workloads must compute a large number of parameters. This computation cannot be parallelized within the server, and the same operation must be repeated for different values in the dataset. Figure 2(c) visualizes dynamics for Correlation, which determines the Pearson correlation between each feature of the input and output variable. The workload aggregates feature values and then calculates the correlation coefficient. Although calculating Pearson for one feature is short and simple, the dataset may possess a large number of features (e.g., more than 1 million in KDDA 2010 [5]), and computation may be serialized for subsets of the features.

Table 2. Linear Regression–Gradient Aggregation

Data Size	Duration	Peak LLC Miss Count
1 GB	0.52 s	3.0e6
512 MB	0.30 s	1.8e6
256 MB	0.17 s	0.8e6
128 MB	0.08 s	0.3e6
64 MB	<0.05 s	0.2e6

Dataset size and its effect on experimental duration, last-level cache misses. Results for the KDDA 2010 dataset [5].

Table 3. Gradient Boosting Tree–Tree Splitting

Dataset	Feature Count	Operation Duration	Peak LLC Miss Count
phishing	68	0.72 s	1.60e6
madelon	500	2.60 s	1.50e6
gisette	5,000	5.01 s	1.60e6
rcv1	47,236	41.00 s	1.65e6
kdda2010	1,163,024	>300.00 s	1.70e6

Number of dataset features and its effect on experimental duration, last-level cache misses. Results for the KDDA 2010 dataset.

For completeness, Figure 2(d) presents Naive–Bayesian, an application that does not operate on data repeatedly. Naive–Bayesian performs parallel tasks for each class of input. Because the dataset has two classes, the algorithm uses two threads to train in parallel.

Implications for management. The rich variations in behavior across computational phases have several implications for colocation management. First, the manager should determine the batch workload’s current phase, from hardware performance counters, and respond accordingly. Second, because only some phases and spikes in memory activity cause an interactive workload to violate its service targets, the framework should only throttle the batch workload during the most contentious phases.

Third, the manager should learn an effective throttling strategy online as the server computes. Online analysis is imperative because offline profiling is intractable. Machine learning workloads’ behavior is highly sensitive to the dataset size and the number of features. An offline profile for a particular dataset cannot generalize and collecting offline profiles for every representative dataset is impossible. For example, Tables 2 and 3 show the impact of dataset size and the number of features on experimental duration and the number of last-level cache misses.

In this article, we propose reinforcement learning for an online colocation controller that manages resource contention. This controller learns and adapts to phases in batch workloads. It also predicts and safeguards the service quality of interactive workloads. Our approach departs from related work, which focuses on the interactive workload’s tail latency alone and neglects batch workload dynamics [15, 29].

3 CONTROLLER DESIGN

3.1 Control Mechanisms

We consider a colocation scenario in which interactive services are assigned to a set of cores and the rest are allocated to batch machine learning. We aim to control batch workloads at fine time granularities to respond to rapid changes in analytics phases and interactive services’ query loads.

Fig. 3. Enforce service quality as a decision-making process. At each step in the sequence, the controller monitors system conditions affecting tail latencies and the performance of interactive services. This information is translated into the state description and reward signal that guide the controller to make decisions on the best effort cores' new DVFS level.

For fine-grained control, we use per-core dynamic voltage and frequency scaling (DVFS) to control resource contention and enforce service quality.¹ DVFS can respond in microseconds for mission-critical workloads [13, 18, 21, 46]. In this project, we propose to use DVFS as a throttling mechanism to control the behavior of batch analytic applications.

3.2 Reinforcement Learning

Managing service quality can be viewed as a sequential decision-making problem, illustrated in Figure 3. At each step in the sequence, the controller monitors system conditions that may produce long tail latencies. Measurements may include software-level queuing statistics, such as distributions on service time and queue length, and hardware-level performance counters. These measured features define a state, which maps to a decision—the voltage and frequency setting for the batch workload's processor cores. The decision changes the system state and tail latency.

The controller defines its reward signal based on tail latency. In each step of the sequence, its goal is maximizing the reward it will receive in subsequent steps. Pursuing this goal is particularly difficult when system conditions vary on short timescales. The controller models system variations and uncertainty with stochastic models such as Markov decision processes (MDPs). Traditional approaches that use stochastic models would take two separate, offline steps: one to model the system and another to derive control policies.

In contrast, reinforcement learning is dynamic and online. Reinforcement learning constructs an optimal, decision-making policy using trial and error in a stochastic environment. It is a powerful framework for solving problems when environmental features that define the state space follow unknown probability distributions. We provide an overview of reinforcement learning and then detail its application to enforcing service quality.

Reinforcement learning places an agent in a stochastic environment. At every timestep t , it perceives the state of the environment s_t and issues an action a_t that yields a reward r_t . The decision process is defined by a five-tuple where

- **S** is a finite set of states that describe the environment.
- **A** is a finite set of actions that an agent may perform. The action at time t affects the state at time $t + 1$.

¹Our method can be generalized to coarse-grained control mechanisms, such as core reallocation, with minimal effort.

- $P_a(s, s') = \Pr(s_{t+1} = s' | s_t = s, a_t = a)$ is the probability that action a taken at time t changes the state to s' . The framework of reinforcement learning assumes that $P_a(s, s')$ depends only on the current state s_t and the current action a_t , and is conditionally independent of previous states and actions.
- $R(s, a)$ is the reward received after action a is taken at state s .
- $\gamma \in [0, 1]$ is the discount factor, which represents how much less the agent values reward at $t + 1$ compared to reward at time t .

A policy π determines the action an agent takes in a given state. The optimal policy maximizes the expected long-term cumulative reward $E(\sum_{t=0}^{\infty} \gamma^t R(s_t, a_t))$.

This measure of long-term reward aligns with our system goals. We wish to improve server utilization while satisfying the 95th percentile latency. Temporarily violating tail latency may reduce rewards now but permit sustained rewards in the future as a short-term burst in the batch workload's resource intensity subsides.

Q-Learning. We optimize the policy with Q-learning, which avoids the difficulty of analyzing transition probabilities in an MDP. Q-Learning is an attractive technique for policy optimization, especially for systems management, because it is model free and lightweight. The model-free approach means that we need not explicitly model the relationship between contention and latency penalties. The lightweight approach to updating models means that the agent can respond quickly to changes in system conditions. For these reasons, Q-learning has already found application in server power management and memory controller design [20, 23].

The Q-value $Q^\pi(s, a)$ for policy π is based on expected future value when taking a particular action in a particular state.

$$Q^\pi(s, a) = E_\pi \left\{ \sum_{k=0}^{\infty} \gamma^k r_{t+k+1} | s_t, a_t \right\}$$

Deriving the optimal policy is equivalent to learning the maximum Q-value for each state-action pair. In state s , the best action a is the one that maximizes the Q-value.

The learning procedure updates Q-values as the controller visits states, takes actions, and observes rewards. The update procedure accounts for the current reward r , the discounted Q-value under the best action a' , and the previous Q-value. The learning rate of the algorithm is dictated by α .

$$Q(s, a) \leftarrow Q(s, a) + \alpha(r + \gamma \max_{a'} Q(s', a') - Q(s, a)) \quad (1)$$

The Q-table tracks Q-values for all state-action pairs. Tabular Q-learning initializes all Q-table entries with the same value. During runtime, the agent repeatedly and greedily selects the action that maximizes Q-value in its current state. The agent updates the Q-value for state-action (s, a) with Equation (1) when it experiences the transition $(s, a) \rightarrow (s', r)$.

3.3 Learning for Colocation Management

We formulate the contention management problem as a sequential decision-making problem, illustrated in Figure 3. We model the problem as an MDP and solve it with Q-learning. We carefully design the action space, state space, and reward signal to satisfy Markovian properties in the reinforcement learning framework.

Action space. We define actions for the control agent to be the available core frequencies in the system. The optimal action depends on the current state of the system and workload.

State space. We design a state space with five features. We choose features that characterize short-term uncertainty in the workload and system. We restrict the space to five features to

constrain state size and Q-table storage overheads. The features include software statistics that describe the interactive workload's performance, hardware counters that describe the system's resource contention, and a single value that describes the voltage and frequency setting for the batch workload's cores.

We use two software statistics—maximum queue length and maximum service time—to characterize uncertainty and variation in the interactive workload. The maximum queue length measures the largest number of pending requests, which is measured when each request enters the system, in the past time interval. The maximum service time measures the longest time required to execute a request in the past time interval. The control agent learns the most effective core frequency for the batch workload for given queue lengths and service times by trying different settings and observing the interactive workload's tail latency.

We also characterize resource contention with two hardware performance counters: batch cores' last-level cache misses and server-wide memory bandwidth usage. We measure cache misses to assess the competition for cache lines between co-runners. We count memory transactions to characterize the contention for shared memory bandwidth. These profiles account for system conditions and phases that vary across time within batch workloads. Finally, we include the current clock frequency for the batch workload's cores. Together, hardware counters and core frequency describe server resources that have been allocated to the batch workload.

Features in our state space are either continuous or unbounded, but Q-tables require discrete representations. First, we bound queue length, service time, and memory throughputs with cut-off values such that any value greater than the cut-off is assumed to cause performance violations in the interactive workload. Second, we divide each continuous domain into 20 bins with equal ranges. The discretized space includes $O(10^5)$ states.

Reward function. We define the reward as a function of interactive latency and system utilization. The function must encode the control agent's two optimization objectives. First, the agent seeks to minimize the difference between measured and target tail latency for the interactive workload. Second, the agent seeks to maximize task throughput for the batch workload. We realize these objectives with the following function.

$$R = \begin{cases} = \text{Batch Frequency} & \text{Latency} \leq \text{Target} \\ = -2 \cdot \text{Max Frequency} & \text{Latency} > \text{Target} \end{cases} \quad (2)$$

The reward function balances competing objectives for the interactive and batch workloads. The control agent is incentivized to meet interactive latency targets with a positive reward when the target is met and a negative reward otherwise. The agent is incentivized to sustain high system utilization with a reward that is proportional to the batch workload's processor core frequency.

Penalties for violating latency targets are large and set at twice the core's maximum frequency. The magnitude of this penalty gives the agent flexibility to balance long-term reward and short-term penalty. For example, the agent may increase the batch workload's core frequency, which leads to an immediate penalty but sustained rewards in subsequent management intervals. Such tradeoffs are encoded in the Q-value for each state-action pair.

3.4 Improving the Learning Mechanism

Although Q-learning is guaranteed to converge in theory, two difficulties arise from tabular Q-learning in practice. First, the agent may rarely observe some states in the table and suffers from randomly chosen actions in those states. Updates to these states are not only rare but also slow. The agent updates the Q-value for the state-action pair (s, a) only after it has taken the action a at state s . There is a one-step delay between updating the Q-value and using the new Q-value.

In practice, the agent may observe some states only once and needs a mechanism to improve on actions taken in rarely visited states by drawing on Q-values from other states.

Second, baseline Q-learning uses ϵ -greedy algorithms to avoid local minima. This algorithm selects a random action with probability ϵ and follows a greedy strategy with probability $1 - \epsilon$. In practice, choosing a good value of ϵ is difficult, depends on hand-tuned values for a specific control problem, and is hard to justify. Moreover, a fixed ϵ causes the agent to react slowly to changes in system conditions. We propose two remedies to address these challenges: neighborhood value approximation and adaptive exploration.

Neighborhood approximation. One way to avoid the one-delay between updating and utilizing the Q-value is to exploit similarities between adjacent states in the Q-table. Because nearby states are similar, we can approximate an unvisited state's Q-values with the average of neighboring Q-values [10]. Approximation extends Q-values of visited states to rarely visited states, avoiding the poor outcomes from choosing actions randomly. We use this approximation to choose an action when a state has been visited fewer than k times and use the state's own Q-value otherwise. In our implementation, we set k to be 4 and the neighborhood radius to be 2.

Adaptive exploration. Balancing exploration and exploitation is one of the most challenging parts in online reinforcement learning. Excessive exploration prevents the controller algorithm from pursuing short-term rewards, causing the system to violate latency targets frequently. Exploration is as likely to choose a poor action as it is to discover a good one because the ϵ -greedy algorithm selects an action uniformly at random from the space. However, exploiting a developed policy may prevent the agent from discovering greater long-term rewards. The controller may stay with a suboptimal policy that often throttles the batch workload.

We propose adaptive exploration with a mechanism that combines softmax exploration [7] and value-difference based exploration (VDBE) [40]. Softmax exploration uses the Boltzmann distribution function to determine the probability that an agent selects action a in state s . T is the temperature factor in the Boltzmann distribution. As $T \Rightarrow 0$, the agent explores less.

$$P(s, a) = \frac{\exp(Q(s, a)/T)}{\sum_{i=1}^m \exp(Q(s, a_i)/T)}$$

VDBE extends the original ϵ -greedy algorithm by assigning independent $\epsilon[s]$ for each state and adjusting $\epsilon[s]$ over time. The desired outcome is that the agent aggressively explores when it knows little about the environment and follows the developed policy when it understands the environment. Knowledge of the environment is reflected by the difference δ between existing and updated Q-values. According to this intuition, we adapt $\epsilon[s]$ with the following rule whenever we update the Q-value for (s, a) .

$$\epsilon_{t+1}[s] = \delta \cdot \frac{1 - \exp(-|\alpha \cdot \delta|)}{1 + \exp(-|\alpha \cdot \delta|)} + (1 - \delta) \cdot \epsilon_t[s]$$

3.5 Deploying the Algorithm Online

Algorithm 1 summarizes our control algorithm. The algorithm has two tunable parameters: the discount factor γ and the learning rate α . First, the discount factor reflects how much less the agent values expected future rewards compared to the current reward. A small discount (i.e., large γ) motivates agents to increment the batch workload's core frequencies to achieve higher task throughput in subsequent management intervals. In contrast, a large discount (i.e., small γ) says that rewards and interactive performance now is worth much more than rewards and batch throughput later. Second, the learning rate represents the agent's expectations for changes in the environment. In our system, $\gamma = 0.9$ and $\alpha = 0.5$.

Fig. 4. The architecture of the Q-learning controller and its interaction with the interactive service and the batch workload. It communicates with the interactive service via a shared memory region. It outputs a DVFS level for the batch workload at each step.

Fig. 5. Online training of the Q-learning agent. Online, the experience buffer is updated with a state-action transition at the beginning of a time interval and trained with a batch of recorded transitions sampled from the experience buffer during each time interval.

ALGORITHM 1: Q-Learning for controlling the batch workload's core frequency

```

1: Initialize Q-table with  $Q(s, a) = \frac{-\max\_penalty}{1-\gamma}$ 
2: Initialize  $\epsilon[s]$  to be 0.1 for all state  $s$ 
3: Initialize BatchDVFS to be lowest
4: while True do
5:   sleep(50 ms)
6:   (LLCMiss, Mem)  $\leftarrow$  Get_Batch_Stage()
7:   (maxService, maxQueue, maxLatency)  $\leftarrow$ 
8:   Get_Req_Statistics()
9:    $s_t \leftarrow$  (maxService, maxQueue, LLCMiss, Mem, BatchDVFS)
10:   $r_t \leftarrow$  reward(maxLatency, BatchDVFS)
11:  if rand()  $< \epsilon[s_t]$  then
12:     $a_t \leftarrow$  Softmax( $s_t$ )
13:  else
14:     $a_t \leftarrow$  SelectwihMaxQ( $s_t$ )
15:  end if
16:  BatchDVFS  $\leftarrow a_{t+1}$ 
17:   $Q(s_{t-1}, a_{t-1}) \leftarrow Q(s_{t-1}, a_{t-1}) + \alpha \cdot (\gamma \cdot (r_t +$ 
18:   $\max_a Q(s_t, a) - Q(s_{t-1}, a_{t-1})))$ 
19:   $\epsilon[s_t] \leftarrow$  VDBE( $\epsilon[s_t]$ )
20: end while
21: return

```

4 CONTROLLER IMPLEMENTATION

We implemented the controller with more than 1,000 lines of Python and C++ codes. The controller runs as a daemon process. It periodically polls queue statistics from the interactive workload and reads hardware counters from the system. It supplies these measurements to the Q-learning agent,

which returns an action from the Q-table. The controller uses this action to update the frequency setting on the batch workload's processor cores.

Figure 4 presents the controller architecture. The controller and interactive workloads use inter-process communication and share a memory page [4]. The page contains a list of query service times, query latencies, and queue lengths at the time of query arrivals. The list is updated by the interactive workload whenever a query completes. The list is read and flushed by the controller periodically. The flush ensures that information read by the controller is independent in time, which guarantees the Markovian assumptions of the Q-learning algorithm. Additionally, the controller integrates performance counter monitors and uses them to read hardware counter values. Finally, the controller relies on Q-learning and Q-tables to update and issue actions.

Experience replay. The controller implements experience replay to accelerate training [9], increasing the learning algorithm's performance by using its measured data more effectively. Figure 5 shows the training time line. At the beginning of time window t , the controller measures the state description s_t , generates a transition $(s_{t-1}, a_{t-1}) \rightarrow (s_t, r_t)$, updates the Q-table with the transition, and updates an experience list that holds the past n state-action-state transitions. The controller obtains an action from the Q-table and sets the batch workload's core frequency.

During the time window t , the controller waits for the next round of software statistics. As it waits, it randomly samples a batch of transitions from the experience list and uses them to train the Q-table. Experience replay strengthens the agent's perception of the environment by repeating previously observed transitions. At the end of this window and the beginning of the next, the controller will receive new information and perform another transition.

Control interval. We implement a control interval of 50 ms. The control frequency represents a tradeoff between hardware and software profiling. With a longer control period, the interactive workload observes a larger population of queries, which mitigates variations in query length and complexity. However, with a shorter control period, the hardware counters provide a more accurate snapshot of microarchitectural activity and contention.

5 EXPERIMENTAL METHODS

Server configuration. We evaluate our controller on a single-socket server with an eight-core E5-2620-v4 chip multiprocessor and the Linux 3.13.0 kernel. The multiprocessor has a 20-MB last-level cache, and the server has 16 GB of memory. The nominal processor frequency is 2.1 GHz and provides nine frequency steps. We use four of these frequency settings with a step size of 0.3 GHz. We control each core's frequency setting with the cpufreq governor [3]. In addition to DVFS, the controller can temporarily stop batch computation, and we denote this action as "0 GHz."

Interactive workloads. We colocate interactive and batch workloads. We use the Xapian search engine and masstree in-memory, key-value store [22], setting the 95th percentile latency target as the interactive workload's performance under maximum sustainable query load. We determine this load by increasing query arrival rates and measuring the inflection point at which latency increases rapidly. On two cores, the search engine sustains 600 queries per second with a 95th percentile latency of 5.8 ms. The key-value store sustains 6,000 queries per second with a 95th percentile latency of 1.5 ms.

Batch workloads. Table 4 lists batch applications and their execution time in our experiments. We use Spark for distributed machine learning on KDDA 2010 datasets [6]. As we observed in Section 2 and Figure 2, these workloads exhibit rich, diverse phase behaviors in the memory subsystem. Linear Regression (LR) and SVM employ stochastic gradient descent, which exhibits three recurring phases that sample data, calculate gradients, and aggregate gradients. Gradient

Table 4. Distributed Machine Learning Workloads (Spark) and Their Demands for Memory Bandwidth

Workload	Memory Bandwidth (GB/s)	Execution Time (minutes)
Correlation (CR)	25.05	2.5
Decision Tree (DT)	21.03	>15
Gradient Boost Tree (GBT)	21.06	>15
Support Vector Machine (SVM)	14.59	5
Linear Regression (LR)	14.66	8
Bayesian (BY)	23.44	2

Table 5. Normalized Average LLC Miss and BW Read/Write for Six Batch Applications at Four DVFS Levels

		2.1 GHz	1.8 GHz	1.5 GHz	1.2 GHz
Naive Bayesian	LLC miss	1.0	0.89	0.85	0.70
	BW	1.0	0.90	0.82	0.75
Correlation	LLC miss	1.0	0.59	0.34	0.31
	BW	1.0	0.65	0.45	0.36
Decision Tree	LLC miss	1.0	0.85	0.74	0.56
	BW	1.0	0.82	0.72	0.55
Gradient Boosting Tree	LLC miss	1.0	0.63	0.55	0.43
	BW	1.0	0.64	0.55	0.45
Linear Regression	LLC miss	1.0	0.92	0.86	0.74
	BW	1.0	0.96	0.75	0.70
SVM	LLC miss	1.0	0.75	0.62	0.55
	BW	1.0	0.71	0.64	0.47

Boosting Tree (GBT) and Decision Tree (DT) determine how best to construct a tree, which exhibits two recurring phases that sample data and recursively evaluate split points. Similarly, Naive Bayesian (BY) and Correlation (CR) exhibit several distinct phases. Recurring memory activity interferes with the interactive workload and tests the controller’s ability to learn and adapt online.

Colocation. We allocate two physical cores to the interactive workload and the remainder to batch workloads. We turn off hyper-threading on the cores that run interactive workloads to avoid intra-core resource contention. Moreover, we assign interactive workloads to the highest priority to avoid performance variations due to the operating system scheduler. Finally, we configure Spark jobs with sufficient thread-level parallelism to fully utilize the allocated cores.

Controller effectiveness. We verify the effectiveness of our control actions by measuring and reporting memory contention at proposed DVFS levels in Table 5. Reducing frequency by one level can reduce memory system contention by 10% to 30%.

6 EVALUATION

Figures 6 and 7 show how reinforcement learning can colocate batch and interactive workloads while ensuring the latter’s service quality. The batch workloads, running at nominal frequencies and without controlling resource contention, cause the interactive workloads to violate service targets even at low query loads. In contrast, when the controller applies its policy from reinforcement learning, batch workloads adapt their core frequencies and interactive workloads satisfy their service obligations at every load level. Indeed, the controller successfully eliminates service

Fig. 6. Search latency without (a) and with (b) the Q-learning controller (5.8-ms target).

Fig. 7. Key-value store latency without (a) and with (b) the reinforcement learning controller (1.5-ms target).

Fig. 8. Visualization of reinforcement learning policy for representative colocations: search regression (a) and search correlation (b). Grids show average, normalized reward for actions given varied LLC miss ratios reported by batch workload. Higher rewards indicate more likely actions in a given state.

violations even as interactive query load approaches the maximum load that can be sustained by the server.

6.1 Understanding Learned Policies

Figures 8 and 9 visualize reinforcement learning's results: the Q-tables. The Q-table estimates the reward from acting in a particular state. Visualizing the table provides insight into the policies learned by the controller. In the figures, grids illustrate the state (x -axis), the action (y -axis), and the corresponding reward (colormap). In each grid, we consider a particular dimension in the state space and assess the reward from an action averaged across the other dimensions. This approach projects a multidimensional state space into a single state to aid the visualization.

Figure 8 illustrates how the control policy responds to LLC miss ratios reported by the batch workload. In general, as the miss ratio increases, the batch workload is more likely to interfere with

Fig. 9. Visualization of reinforcement learning policy for representative colocations, search regression (a) and key-value regression (b). Grids show average, normalized reward for actions given queue lengths reported by interactive workload. Higher rewards indicate more likely actions in a given state.

the interactive workload, which in turn penalizes expected rewards. To mitigate such penalties, the controller throttles batch computation and favors lower core frequencies.

The controller tailors this general policy to the specific behaviors from each batch workload. For example, compare and contrast policies when colocating search with correlation and with regression. The controller learns a more conservative policy for correlation. When the batch workload reports an LLC miss ratio of 0.1, the optimal frequency is 2.1 GHz for regression but only 1.2 GHz for correlation. The controller arrives at this policy by observing that correlation’s phases are relatively short and its memory activity is relatively contentious even when the reported LLC miss ratio is low.

Similarly, Figure 9 illustrates how the control policy responds to queue lengths reported by the interactive workload. Longer queues increase waiting time for computation and degrade service quality, which in turn penalizes expected rewards. The controller mitigates such penalties by lowering the core frequency for batch workloads as queue occupancy increases for the interactive workload.

Again, the controller tailors policies to specific service requirements. The key-value store’s target for tail latency is more stringent than that of the search engine. In addition, its latency is more sensitive to the number of pending queries. In response, the controller learns a more conservative policy that throttles the batch workload more often when colocated with the key-value store. For example, the optimal core frequency for the batch workload is 2.1 GHz when the search engine’s queue occupancy is 1 but is only 1.8 GHz when the key-value store’s queue occupancy is 1.

6.2 Increasing Utilization and Efficiency

The ultimate goal of workload colocation is increasing a server’s utilization, thereby amortizing its fixed power costs over more computation. We measure reinforcement learning’s effect on utilization even as it ensures service quality for the interactive workload. Server utilization is often reported relative to its peak, but determining the normalizing factor can be complicated for colocated workloads. We assess utilization in three steps.

First, we normalize interactive throughput by the server’s maximum sustainable throughput. If the search engine’s two cores support up to 600 queries per second without violating service targets, a search engine that completes 300 queries per second leads to 50% utilization of its cores.

Second, we normalize batch throughput by the workload’s task completion rate when running alone at nominal frequencies. If the Spark runtime reports 20 tasks per second under nominal operation, the batch computation that completes 15 tasks per second leads to 75% utilization of its cores. Note that tasks count toward batch throughput only when the interactive workload meets its performance target.

Fig. 10. (a–b) Colocations with reinforcement learning increase server utilization, a weighted average of interactive and batch workloads’ utilization.

Fig. 11. Colocations with reinforcement learning increase server efficiency by amortizing fixed power costs over more (batch) computation. Power measured for interactive (left: search; right: memcached), batch workloads in Figure 10.

Finally, we calculate the average of interaction and batch workloads’ throughputs weighted by their share of the server’s cores. When the system assigns two and six cores to interactive and batch workloads, throughputs are weighted by 0.25 and 0.75, respectively. In our system, the interactive workloads, the search engine, and the key-value store cannot use more than two cores because additional thread parallelism induces contention for shared data structures such as the search engine’s index table and the key-value store’s in-memory database.

Figure 10 presents server utilization when the search engine and key-value store are colocated with batch machine learning workloads. In the baseline system, the lack of contention control causes any colocation to violate latency targets and utilization depends only on the interactive workload’s query arrival rate. With reinforcement learning, contention control permits colocations that do not compromise interactive latencies, increasing utilization from 40% to 80% of the server’s peak capacity. The controller achieves similar utilization levels for batch workloads with similar workload phases (e.g., regression and svm). Batch workloads that do not achieve high utilization, such as correlation, exhibit large oscillations in memory intensity that cause the controller to learn conservative policies and throttle the batch workload’s processor cores. The controller is similarly conservative when colocating batch workloads with the key-value store, which is sensitive to memory system contention.

Figure 11 illustrates the server’s energy disproportionality and the effect on efficiency, presenting data for every experiment in Figure 10. We define efficiency as utilization divided by power, measured relative to the server’s peak power. We find that idle power, at 30 W, is a dominant share of the total because we must turn off processor cores’ C-states to guarantee the interactive service’s performance. Under these conditions, normalized efficiency is (approximately) inversely proportional to batch throughput. As throughput increases, servers’ large fixed power costs are

Fig. 12. (a–d) Reinforcement learning responds to new batch workloads while interactive service runs at peak query loads. At 200 seconds, gradient boosting tree ends and linear regression begins.

amortized over more work and dissipated for shorter durations as jobs complete faster. Thus, by permitting colocation, reinforcement learning increases efficiency by up to 50%.

6.3 Measuring Controller Responsiveness

Reinforcement learning requires time to accumulate experience and optimize policies. This online training time determines the controller’s responsiveness to a dynamic system. The interactive service’s query load may change. Batch workloads may depart or arrive. As the system evolves, the controller must collect data and update its Q -tables.

We assess training overheads by measuring how much time the controller needs to stabilize tail latencies and meet performance targets. We define a time series of tail latency to be stable if 99% of the series is below the target tail latency. We measure two types of training overhead. First, reinforcement learning incurs cold-start overheads when the server is turned on and the controller constructs a new Q -table. Second, learning incurs transition overheads when a batch workload completes, a different one begins, and the controller updates a previously trained Q -table.

Figure 12 indicates that reinforcement learning responds quickly to new batch workloads despite transition overheads. The controller can restabilize the interactive workload’s latency shortly after the batch workload changes from gradient boosting tree to linear regression at 200 seconds in the time series. The transition requires only tens of online profiling intervals.

Figure 13 illustrates how quickly reinforcement learning responds when the batch workload changes. Transition overheads are less than cold-start overheads because the controller benefits from learned policies and a complete Q -table even if they were learned for different, prior workloads. Q -tables for different batch workloads often share common characteristics, as we saw in

Fig. 13. (a–f) Online overhead for adapting to the arrival of batch applications. The interactive service used in this set of experiments is search. The x-axis in each figure represents the original batch application in the system. The “cold” data represents the cold-start scenario in which no batch application is in the system before the arrival of new batch application.

Fig. 14. (a–b) Reinforcement learning learns from peak at 40 seconds and responds to peak at 140 seconds.

Figure 8. By exploiting these similarities, the controller spends less time transitioning between workloads.

Transition overheads decrease with workload similarity. The gradient boosting tree and decision tree perform similar computation to construct their trees, which leads to negligible transition overheads in Figure 13(b) and (c). The linear regression and SVM both require gradient descent, which leads to low transition overheads in Figure 13(e) and (f).

Similarly, Figure 14 shows that reinforcement learning responds quickly to variations in interactive query load. When the load peaks for the first time at 40 seconds, the search engine violates the targeted latency for 5 seconds as the controller explores new states with greater queue lengths. When the load peaks again at 140 seconds, the search engine meets its performance targets as the controller implements decisions with newly learned policies and updated Q-tables.

ALGORITHM 2: StatThreshold controls batch core's frequency

```

1: while True do
2:   sleep(50ms)
3:   (InstantQPS,TailLatency)  $\leftarrow$  Get_Req_Statistics()
4:   slack  $\leftarrow$  (Target - TailLatency)/Target
5:   if slack < 0 then
6:     action  $\leftarrow$  StopBatch
7:   else if slack <  $T$  then
8:     action  $\leftarrow$  DecreaseFreq
9:   else if InstantQPS >  $Q$  then
10:    action  $\leftarrow$  KeepFreq
11:   else
12:    action  $\leftarrow$  IncreaseFreq
13:   end if
14: end while
15: return

```

6.4 Comparing with Feedback Control

We compare reinforcement learning against popular closed-loop methods. First, we consider static control thresholds [25, 26]. Algorithm 2 sets two thresholds—one on slack and another on query load—to control the batch workload's core frequency. Slack measures the difference between target and measured latency for the interactive workload. If no slack exists, batch workloads are halted. If slack is modest and less than some threshold T , batch workloads compute at lower core frequencies. If slack is abundant but interactive query load is high and exceeds some threshold Q , batch workloads continue at current frequencies. Otherwise, batch computation proceeds at higher frequencies. We optimize parameters T and Q for each colocated batch workload to ensure that 95th percentile latency targets are met.

Second, we consider PID control to set batch workloads' core frequencies [17, 35]. Error $e(t)$ at time t is the difference between target and measured latency for the interactive workload. In each time interval, the controller computes the proportional error $p = e(t)$, the net error $i = \sum_{i=0}^t e(i)$, and the derivative error $d = e_t - e_{i-1}$. The weighted sum of these terms $u = k_p \cdot p + k_i \cdot i + k_d \cdot d$ is used to correct the output in next interval by increasing, decreasing, or maintaining the batch workload's core frequency.

We determine optimal control parameters by profiling the complete execution of batch applications at different frequency levels and use profiles to construct a closed-loop transfer function between input signals (batch frequency level, past tail latency, past LLC counters, past bandwidth counters) and the output signal (future tail latency). We use standard root locus analysis to optimize control parameters [17].

These profiling overheads for setting PID control parameters are particularly significant for cloud and datacenter computing in which new batch applications arrive frequently. These applications consist of diverse phase patterns that arise from different combinations of different data operations. For example, a machine learning pipeline submitted to the cloud could include operations for data preprocessing, variable selection, and concatenation of different learning models. The system cannot simply use control parameters developed for a subset of batch applications to effectively control all batch workloads. Given these challenges, online learning is a compelling alternative that requires significantly less offline profiling.

Figure 15 compares server utilization from reinforcement learning and previously proposed closed-loop methods. The figure reports utilization averaged across five interactive query load

Fig. 15. Comparisons between reinforcement learning and closed-loop methods [17, 25, 26, 35]. Data shown for search engine (a) and key-value store (b).

Fig. 16. (a–f) Static cache partitioning performance. Batch applications are colocated with web search at the highest load (i.e., 500 QPS). For comparison, the horizontal stars and dashes show the Q-learning controller's interactive latency and batch throughput.

levels. *StatThreshold* achieves less utilization than the other two methods because it is being over-conservative.

Control with reinforcement learning is typically 10% to 30% higher than control with static thresholds. First, static thresholds pause batch computation whenever interactive performance targets are violated. Because machine learning codes can oscillate between high and moderate memory activity, simply pausing batch computation will over-throttle. Second, mechanisms with static thresholds gradually increase frequency after resuming computation, finding a suitable decision and under-utilizing the server during this time. Finally, static thresholds lack adaptivity and cannot respond to variations in workload behavior.

6.5 Comparing with Static Cache Partition

Figure 16 compares dynamic Q-learning with static cache partitioning based on Intel's cache allocation technology [1]. The figure reports interactive latency and batch throughput as the number

Table 6. Overhead Measurement for Three Key Operations

Function Blocks		Overhead
PCM Read	4.4 ms	± 1.4 ms
Shared Memory Read	6 us	± 0.19 us
Q-Learning Inference	0.3 ms	± 0.04 ms

of allocated cache ways varies from 4 to 18. When batch applications are colocated with web search at the highest query load (i.e., 500 QPS), increasing the batch workload's cache allocation at the expense of the interactive workload's allocation increases batch throughput but harms interactive latency. Static cache allocation fails to balance batch and interactive performance and performs worse than Q-learning for nearly all applications. Q-learning's advantage arises from detecting phase changes in the batch workload and throttling its cache allocation during contentious phases.

6.6 Characterizing Controller Overhead

Table 6 reports the mean and the standard deviation of measured overheads for three of the Q-learning controller's operations. First, *PCM Read* accesses performance counter monitors for the number of last-level cache misses and the number of bytes read and written by the memory controller. Second, *Shared Memory Read* obtains the interactive applications' performance metrics from a shared memory page. Third, *Q-Learning Inference* generates control actions. We measure elapsed time of these operations, using Time module in Python 2.7, for all pairings of two interactive services and six batch applications. *PCM Reads* account for most of the runtime overhead—approximately 8.5% of the control interval, on average—and the other two operations incur less than 1% of the control interval.

6.7 Analyzing Hyperparameter Sensitivity

We analyze the Q-learning controller's sensitivity to two key hyperparameters. First, the punishment coefficient in the reward signal (Equation (2)) represents the controller's tolerance to violations of the target interactive latency. Practically, a more severe coefficient leads to more frequent punishments for the colocated batch workload, which may be necessary when collocations involve workloads with less predictable phase patterns. The default coefficient value is 2, and we assess sensitivity to values in the range of [0.5, 10]. Figure 17 indicates that increasing the punishment coefficient reduces interactive tail latency and batch throughput. The trend flattens out as the punishment coefficient approaches the upper end of the range.

Second, the latency percentile used to specify target latency describes how strictly performance requirements are enforced. A higher percentile measure reflects more stringent requirements and leads to more conservative control policies. The default percentile is 95, and we measure the effect of percentiles set at 90 and 99 as well. Figure 18 shows that the controller cannot meet performance targets consistently given the strictest latency percentile. Tail latency exceeds the target by 0.7 to 0.8 ms. Note that we obtain these measurements with the default punishment coefficient, and practical controller configurations should increase punishments as latency requirements become more stringent.

6.8 Assessing Q-Learning Optimization Techniques

We evaluate the effect of three proposed Q-learning optimization techniques: neighborhood approximation, adaptive exploration, and experience replay. Experience replay increases exploration speed. Neighborhood approximation and adaptive exploration balance tradeoffs between

Fig. 17. (a–f) Controller's sensitivity to punishment coefficient in reward function. Batch applications are colocated with web search at the highest load (i.e., 500 QPS).

Fig. 18. (a–f) Controller's sensitivity to interactive latency percentile in performance target. Batch applications are colocated with web search at the highest load (i.e., 500 QPS).

exploration and exploitation. Figure 19 shows reductions in tail latency when using optimized Q-learning rather than the naive baseline. Adaptive exploration improves performance least and reduces tail latency by only 3% to 8%. Neighborhood approximation improves latency more and reduces tail latency by 15% to 20%. Finally, experience replay improves performance the most and is the most significant optimization technique for online Q-learning.

Fig. 19. (a–b) Q-learning optimization techniques. Batch workloads are colocated with search and memcached running at highest loads (i.e. 500 QPS/5K QPS). We run each colocation 20 times, reporting average and standard deviation.

7 RELATED WORK

Managing service quality. Many studies predict the interactive workload’s quality of service (QoS) on a shared chip multiprocessor, with different degrees of offline profiling and heuristics. In contrast, Q-learning develops policies online. Bubble-Up and Bubble-Flux [29, 42] microbenchmarks batch workloads’ contributions to memory system contention and predict latency with a calibration model. SMiTe [45] profiles a workload’s sensitivity to shared core resources and predict tail latency with heuristics. Dirigent [46] monitors execution and compares progress with offline profiles to estimate completion time. Paragon and Quasar estimate interference with recommendation algorithms [14, 15]. These frameworks classify each application by its resource demands and sensitivity to resource contention, and recommend the best co-runner. Q-Clouds and DeepDive [31, 33] address resource allocation and interference in virtual machines to ensure QoS. Hipster [32] applies a combination of heuristics and reinforcement learning to manage QoS of latency-critical applications in heterogeneous chip multiprocessors. Our work focus on the effect of rapidly changing phases in the batch applications and does not rely on heuristic methods. Sparrow [34] schedules fine-grained tasks that complete in hundreds of milliseconds. It supports strict priority-based resource allocation policies but fails to consider resource contention between colocated tasks. In addition, Duplexity [47] uses an aggressive multithreading technique to hide the latency in the order of microseconds.

Isolating resources in multiprocessors. Recent microarchitectural advances motivate significant work in shared cache partitioning. Effective partitions can support varied latency, throughput, and fairness goals [12, 21, 31, 36, 37]. CPI² [44] develops a CPI-based software mechanism that identifies the application that causes severe contention and throttles it if necessary. Heracles [26] combines a hardware and software mechanism that manages processor, memory, cache, and network usage to mitigate the interference from batch application. It uses threshold-based feedback control to dynamically adjust the resource allocation. However, static thresholds may be conservative and leave throughput unexploited.

Reinforcement learning for systems. The most significant application of reinforcement learning in computer architecture is a self-optimizing memory controller [20], which uses Q-learning to schedule memory transactions. In addition, reinforcement learning has also been applied to data-center resource allocation [19, 27, 28, 39, 43].

Improving energy efficiency in servers. Prior work reduces energy consumed by latency-critical workloads. Q-learning is also suitable for this setting because it can exploit and adapt to short-term variations in queueing delay. We could adapt our controller with slight modifications to the state and action space. Pegasus [13] improves datacenter efficiency by dynamically adjusting server power caps in response to changes in load at coarse timescales. Rubik [21] and TimeTrader [41] save energy by adjusting frequency at fine timescales in response to short-term variations in queueing delay. Adrenaline [18] adjusts the frequency on a per-query basis. μ DPM [48] focuses on efficient dynamic power management for micro-services with services time in the order of hundreds of microseconds.

8 CONCLUSION

We present a reinforcement learning framework for dynamically optimizing resource management policies in shared servers. Servers can improve efficiency by colocating interactive and batch workloads. But batch computation, especially modern machine learning applications, can exhibit severe, short-term variations that jeopardize interactive latency targets. We show that reinforcement learning can respond quickly to these variations, satisfying latency targets while increasing server utilization. The controller requires no offline profiling and learns to adapt its control policies as resource demands from batch workloads evolve.

REFERENCES

- [1] Intel. (n.d.). Cache Monitoring Technology and Cache Allocation Technology—Intel. Retrieved November 17, 2017 from <https://www.intel.com/content/www/us/en/communications/cache-monitoring-cache-allocation-technologies.html>.
- [2] Paul Manag (2004). Cgroups—The Linux Kernel Archives. Retrieved November 17, 2017 from <https://www.kernel.org/doc/Documentation/cgroup-v1/cgroups.txt>.
- [3] Dominik Brodowski. (n.d.). CPUFreq Governor—The Linux Kernel Archives. Retrieved January 29, 2020 from <https://www.kernel.org/doc/Documentation/cpu-freq/governors.txt>.
- [4] David Marshall. (1999). IPC:Shared Memory. Retrieved November 20, 2017 from <https://users.cs.cf.ac.uk/dave/C/node27.html>.
- [5] National Taiwan University. (n.d.). LIBSVM Data: Classification (Binary Class). Retrieved January 29, 2020 from <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/binary.html>.
- [6] Apache Spark. (2018). MLlib. Retrieved November 17, 2017 from <https://spark.apache.org/mllib/>.
- [7] Marcello Restelli. (2015). Reinforcement Learning: Exploration vs Exploitation. Retrieved November 17, 2017 from <http://home.deib.polimi.it/restelli/MyWebSite/pdf/rl5.pdf>.
- [8] Alex Benik. (2013). The sorry state of server utilization and the impending post-hypervisor era. *GigaOm*. Retrieved January 29, 2020 from <https://gigaom.com/2013/11/30/the-sorry-state-of-server-utilization-and-the-impending-post-hypervisor-era/>.
- [9] Sander Adam, Lucian Busoniu, and Robert Babuska. 2012. Experience replay for real-time reinforcement learning control. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 42, 2 (2012), 201–212.
- [10] Martin Allen and Phil Fritzsche. 2011. Reinforcement learning with adaptive Kanerva coding for Xpilot game AI. In *Proceedings of the 2011 IEEE Congress of Evolutionary Computation (CEC'11)*. IEEE, Los Alamitos, CA, 1521–1528.
- [11] Luiz André Barroso and Urs Hölzle. 2007. The case for energy-proportional computing. *Computer* 40, 12 (2007), 33–37.
- [12] Sangyeun Cho and Lei Jin. 2006. Managing distributed, shared L2 caches through OS-level page allocation. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE, Los Alamitos, CA, 455–468.
- [13] Ewa Deelman, James Blythe, Yolanda Gil, Carl Kesselman, Gaurang Mehta, Sonal Patil, Mei-Hui Su, Karan Vahi, and Miron Livny. 2004. Pegasus: Mapping scientific workflows onto the grid. In *Grid Computing*. Springer, 131–140.
- [14] Christina Delimitrou and Christos Kozyrakis. 2013. Paragon: QoS-aware scheduling for heterogeneous datacenters. *ACM SIGPLAN Notices* 48 (2013), 77–88.
- [15] Christina Delimitrou and Christos Kozyrakis. 2014. Quasar: Resource-efficient and QoS-aware cluster management. *ACM SIGPLAN Notices* 49 (2014), 127–144.
- [16] Kim Hazelwood, Sarah Bird, David Brooks, Soumith Chintala, Utku Diril, Dmytro Dzhulgakov, Mohamed Fawzy, et al. 2018. Applied machine learning at Facebook: A datacenter infrastructure perspective. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA'18)*. IEEE, Los Alamitos, CA, 620–629.

- [17] Joseph L. Hellerstein, Yixin Diao, Sujay Parekh, and Dawn M. Tilbury. 2004. *Feedback Control of Computing Systems*. John Wiley & Sons.
- [18] Chang-Hong Hsu, Yunqi Zhang, Michael A. Laurenzano, David Meisner, Thomas Wenisch, Jason Mars, Lingjia Tang, and Ronald G. Dreslinski. 2015. Adrenaline: Pinpointing and reining in tail queries with quick voltage boosting. In *Proceedings of the IEEE 21st International Symposium on High Performance Computer Architecture (HPCA'15)*. IEEE, Los Alamitos, CA, 271–282.
- [19] Markus C. Huebscher and Julie A. McCann. 2008. A survey of autonomic computing? Degrees, models, and applications. *ACM Computing Surveys* 40, 3 (2008), 7.
- [20] Engin Ipek, Onur Mutlu, José F. Martínez, and Rich Caruana. 2008. Self-optimizing memory controllers: A reinforcement learning approach. In *Proceedings of the 35th International Symposium on Computer Architecture (ISCA'08)*. IEEE, Los Alamitos, CA, 39–50.
- [21] Harshad Kasture, Davide B. Bartolini, Nathan Beckmann, and Daniel Sanchez. 2015. Rubik: Fast analytical power management for latency-critical systems. In *Proceedings of the 48th International Symposium on Microarchitecture*. ACM, New York, NY, 598–610.
- [22] Harshad Kasture and Daniel Sanchez. 2016. Tailbench: A benchmark suite and evaluation methodology for latency-critical applications. In *Proceedings of the IEEE International Symposium on Workload Characterization (IISWC'16)*. IEEE, Los Alamitos, CA, 1–10.
- [23] Wei Liu, Ying Tan, and Qinru Qiu. 2010. Enhanced Q-learning algorithm for dynamic power management with performance constraint. In *Proceedings of the Conference on Design, Automation, and Test in Europe*. 602–605.
- [24] Qiuyun Llull, Songchun Fan, Seyed Majid Zahedi, and Benjamin C. Lee. 2017. Cooper: Task colocation with cooperative games. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA'17)*. IEEE, Los Alamitos, CA, 421–432.
- [25] David Lo, Liqun Cheng, Rama Govindaraju, Luiz André Barroso, and Christos Kozyrakis. 2014. Towards energy proportionality for large-scale latency-critical workloads. *ACM SIGARCH Computer Architecture News* 42 (2014), 301–312.
- [26] David Lo, Liqun Cheng, Rama Govindaraju, Parthasarathy Ranganathan, and Christos Kozyrakis. 2015. Heracles: Improving resource efficiency at scale. *ACM SIGARCH Computer Architecture News* 43 (2015), 450–462.
- [27] Tania Lorido-Botran, Jose Miguel-Alonso, and Jose A. Lozano. 2014. A review of auto-scaling techniques for elastic applications in cloud environments. *Journal of Grid Computing* 12, 4 (2014), 559–592.
- [28] Hongzi Mao, Mohammad Alizadeh, Ishai Menache, and Srikanth Kandula. (2016). Resource management with deep reinforcement learning. In *Proceedings of the 15th ACM Workshop on Hot Topics in Networks (HotNets'16)*. 50–56.
- [29] Jason Mars, Lingjia Tang, Robert Hundt, Kevin Skadron, and Mary Lou Soffa. 2011. Bubble-Up: Increasing utilization in modern warehouse scale computers via sensible co-locations. In *Proceedings of the 44th Annual IEEE/ACM International Symposium on Microarchitecture*. ACM, New York, NY, 248–259.
- [30] Dirk Merkel. 2014. Docker: Lightweight Linux containers for consistent development and deployment. *Linux Journal* 2014, 239 (2014), 2.
- [31] Ripal Nathuji, Aman Kansal, and Alireza Ghaffarkhah. 2010. Q-Clouds: Managing performance interference effects for QoS-aware clouds. In *Proceedings of the 5th European Conference on Computer Systems*. ACM, New York, NY, 237–250.
- [32] Rajiv Nishtala, Paul Carpenter, Vinicius Petrucci, and Xavier Martorell. 2017. Hipster: Hybrid task manager for latency-critical cloud workloads. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA'17)*. IEEE, Los Alamitos, CA, 409–420.
- [33] Dejan Novakovic, Nedeljko Vasic, Stanko Novakovic, Dejan Kostic, and Ricardo Bianchini. 2013. DeepDive: Transparently identifying and managing performance interference in virtualized environments. In *Proceedings of the 2013 USENIX Annual Technical Conference*.
- [34] Kay Ousterhout, Patrick Wendell, Matei Zaharia, and Ion Stoica. 2013. Sparrow: Distributed, low latency scheduling. In *Proceedings of the 24th ACM Symposium on Operating Systems Principles*. ACM, New York, NY, 69–84.
- [35] Vinicius Petrucci, Michael A. Laurenzano, John Doherty, Yunqi Zhang, Daniel Mosse, Jason Mars, and Lingjia Tang. 2015. Octopus-Man: QoS-driven task management for heterogeneous multicores in warehouse-scale computers. In *Proceedings of the IEEE 21st International Symposium on High Performance Computer Architecture (HPCA'15)*. IEEE, Los Alamitos, CA, 246–258.
- [36] Moinuddin K. Qureshi and Yale N. Patt. 2006. Utility-based cache partitioning: A low-overhead, high-performance, runtime mechanism to partition shared caches. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-39)*. IEEE, Los Alamitos, CA, 423–432.
- [37] Daniel Sanchez and Christos Kozyrakis. 2011. Vantage: Scalable and efficient fine-grain cache partitioning. *ACM SIGARCH Computer Architecture News* 39 (2011), 57–68.
- [38] Gerald Tesauro. 2007. Reinforcement learning in autonomic computing: A manifesto and case studies. *IEEE Internet Computing* 11, 1 (2007), 22–30.

- [39] Gerald Tesauro. 2005. Online resource allocation using decompositional reinforcement learning. In *Proceedings of the 20th National Conference on Artificial Intelligence (AAAI'05)*, Vol. 2. 886–891.
- [40] Michel Tokic and Günther Palm. 2011. Value-difference based exploration: Adaptive control between epsilon-greedy and softmax. *KI 2011: Advances in Artificial Intelligence*. Lecture Notes in Computer Science, Vol. 7006. Springer, 335–346.
- [41] Balajee Vamanan, Hamza Bin Sohail, Jahangir Hasan, and T. N. Vijaykumar. 2015. TimeTrader: Exploiting latency tail to save datacenter energy for online search. In *Proceedings of the 48th International Symposium on Microarchitecture*. ACM, New York, NY, 585–597.
- [42] Hailong Yang, Alex Breslow, Jason Mars, and Lingjia Tang. 2013. Bubble-Flux: Precise online QoS management for increased utilization in warehouse scale computers. *ACM SIGARCH Computer Architecture News* 41 (2013), 607–618.
- [43] Jingling Yuan, Xing Jiang, Luo Zhong, and Hui Yu. 2012. Energy aware resource scheduling algorithm for data center using reinforcement learning. In *Proceedings of the 5th International Conference on Intelligent Computation Technology and Automation (ICICTA'12)*. IEEE, Los Alamitos, CA, 435–438.
- [44] Xiao Zhang, Eric Tune, Robert Hagmann, Rohit Inagal, Vrigo Gokhale, and John Wilkes. 2013. CPI 2: CPU performance isolation for shared compute clusters. In *Proceedings of the 8th ACM European Conference on Computer Systems*. ACM, New York, NY, 379–391.
- [45] Yunqi Zhang, Michael A. Laurenzano, Jason Mars, and Lingjia Tang. 2014. Smiter: Precise QoS prediction on real-system SMT processors to improve utilization in warehouse scale computers. In *Proceedings of the 47th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-47)*. IEEE, Los Alamitos, CA, 406–418.
- [46] Haishan Zhu and Mattan Erez. 2016. Dirigent: Enforcing QoS for latency-critical tasks on shared multicore systems. *ACM SIGPLAN Notices* 51, 4 (2016), 33–47.
- [47] Amirhossein Mirhosseini, Akshitha Sriraman, and Thomas F. Wenisch. 2019. Enhancing server efficiency in the face of killer microseconds. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA'19)*. IEEE, 185–198.
- [48] Chih-Hsun Chou, Laxmi N. Bhuyan, and Daniel Wong. 2019. μ dpm: Dynamic power management for the microsecond era. In *Proceedings of the IEEE International Symposium on High Performance Computer Architecture (HPCA'19)*. IEEE, 120–132.

Received May 2019; revised October 2019; accepted December 2019