CoScal: Multi-faceted Scaling of Microservices with Reinforcement Learning

Minxian Xu, Member, IEEE, Chenghao Song, Shashikant Ilager, Sukhpal Singh Gill, Juanjuan Zhao, Kejiang Ye, Member, IEEE, and Chengzhong Xu, Fellow, IEEE

Abstract—The emerging trend towards moving from monolithic applications to microservices has raised new performance challenges in cloud computing environments. Compared with traditional monolithic applications, the microservices are lightweight, fine-grained, and must be executed in a shorter time. Efficient scaling approaches are required to ensure microservices' system performance under diverse workloads with strict Quality of Service (QoS) requirements and optimize resource provisioning. To solve this problem, we investigate the trade-offs between the dominant scaling techniques, including horizontal scaling, vertical scaling, and brownout in terms of execution cost and response time. We first present a prediction algorithm based on gradient recurrent units to accurately predict workloads assisting in scaling to achieve efficient scaling. Further, we propose a multi-faceted scaling approach using reinforcement learning called CoScal to learn the scaling techniques efficiently. The proposed CoScal approach takes full advantage of data-driven decisions and improves the system performance in terms of high communication cost and delay. We validate our proposed solution by implementing a containerized microservice prototype system and evaluated with two microservice applications. The extensive experiments demonstrate that CoScal reduces response time by 19%-29% and decreases the connection time of services by 16% when compared with the state-of-the-art scaling techniques for Sock Shop application. CoScal can also improve the number of successful transactions with 6%-10% for Stan's Robot Shop application.

Index Terms—Cloud Computing, workload prediction, microservices, reinforcement learning, brownout, scalability

I. INTRODUCTION

The cloud computing paradigm needs to satisfy strict performance requirements for diverse users hosting heterogeneous workloads. These cloud workloads are driven by applications belonging to various domains, including enterprise businesses and government services that require uninterrupted, reliable service delivery [1] [2]. The ever-increasing complexity of these large applications has recently stipulated moving towards microservice-based application development and deployment. The microservice paradigm has shifted the traditional monolithic application design into decomposed and selfcontained standalone application components, which generally

Email: czxu@um.edu.mo. (Corresponding author: Kejiang Ye.) communicate through a RESTful Application Programming Interface (API) [3] [4]. The features of microservices include lightweight design, flexible development, continuous deployment, and independent management. These attractive features have promoted cloud computing providers, including Amazon, Google, and Alibaba, to adopt microservice-based application service delivery models and platforms.

The users request on-demand services from cloud service providers under the specified Quality of Service (QoS) requirements with the pay-as-you-go model [5]. If the QoS is unsatisfied, service providers may suffer revenue loss due to a Service Level Agreement (SLA) violation with their customers. The QoS is defined in terms of different performance metrics, including resource availability and latency. Satisfying the QoS of microservices-based applications is more challenging than traditional monolithic applications since the performance of microservices is more sensitive to latency. This is further complicated by dynamic workload levels. Thus, service providers aim to allocate sufficient resources to application workloads to avoid QoS degradation and performance bottlenecks and prevent huge revenue loss.

Over-provisioning is an effective approach to ensure the QoS of microservices. The over-provisioning method aims to allocate additional resources for user workloads to provide guaranteed performance [6]. However, provisioning as many resources as possible is cost-ineffective due to the finite hardware resources, energy budget, and operational costs of the cloud data centers [7]. Maintaining a large number of physical machines with lower utilization can lead to higher costs for service providers. Therefore, resource scaling approaches have been applied to optimize resource planning and provisioning in microservice-based applications. The resource scaling methods are significant for infrastructure providers as they can contribute to cost reduction, increase resource utilization, and simultaneously improve the QoS parameters of microservice applications.

The dominant resource scaling approaches for microservices can be classified into two categories, horizontal scaling and vertical scaling [8]. Horizontal scaling adjusts the provisioned resource quantities by adding or removing microservice replicas to improve resource usage and QoS parameters (e.g, system availability, latency). In contrast, vertical scaling adjusts the capabilities of existing provisioned resources by increasing or decreasing the amount of CPU, memory, and network resources. These two approaches have both been validated to be effective for resource scaling.

Currently, Kubernetes (promoted by Google) has become

M. Xu, C. Song, J. Zhao, K. Ye are with Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences, China. Email: {mx.xu, ch.song, jj.zhao, kj.ye}@siat.ac.cn.

S. Ilager is with Vienna University of Technology, Austria. Email: shashikant.ilager@tuwien.ac.at.

S. S. Gill is with School of Electronic Engineering and Computer Science, Queen Mary University of London, London, UK. Email: s.s.gill@qmul.ac.uk. C. Xu is with State Key Lab of IoTSC, University of Macau, Macau, China.

the most popular platform for container-based microservice application deployment in Cloud [9]. Kubernetes supports horizontal scaling based on a runtime usage threshold. However, such a threshold-based scaling method only works well in simple cases, and achieves sub-optimal solutions in complex workload conditions (e.g. dynamic workload conditions). To address this issue, many studies have proposed different autoscaling algorithms based on the default Kubernete's autoscaling algorithms [10]. Nevertheless, these algorithms primarily focus on a single type of resource, e.g., CPU, memory, or bandwidth which are infeasible for many workload scenarios. In addition, Horizontal and vertical scaling has many practical limitations when applied individually. For instance, the communication overhead of horizontal scaling are nontrivial, which directly impacts application QoS, especially in microservice-based applications due to their shorter execution time and increased latency sensitivity. Although vertical scaling can increase the capacity of a provisioned physical machine's resource, upgrading the resource capacity in runtime is expensive. A recent study has considered combining these two scaling approaches together [8], although it does not address the complex real world scenarios such as infrequent bursty workloads that leads to overloading of the whole system, thus, affecting overall infrastructure and application OoS [11].

Some recent methods, such as brownout techniques, [12] have explored an alternative scaling method complementing both the horizontal and vertical scaling. Brownout is a selfadaptive approach for managing application components by dynamically activating and deactivating optional application components to be adaptive to the variance of workloads. The brownout can effectively address the limitations of vertical scaling (limited resource capacity of local machines) and horizontal scaling (the resource usage of replicating the microservice to other machines) in overloaded conditions. In the microservice-based application, a brownout can temporarily deactivate some of the optional microservices to reduce resource usage while ensuring the necessary functionalities of microservice applications. Thus, comprehensive scaling techniques addressing the limitations of existing individualistic approaches are required for adaptive and efficient scaling in microservice environments.

A. Existing Challenges

The efficient scaling of microservices poses several significant challenges. Firstly, the micro-service based cloud workloads have high-variance in resource usage and are sensitive to resource types (CPU, network, memory). It is difficult to predict the accurate amount of workloads in a specific time period. An accurate workload prediction is an essential element for auto-scaling approaches in microservices, e.g., in horizontal scaling, the predicted workload can enable service providers to boot up and deploy applications beforehand to avoid startup cost and QoS degradation. Secondly, assigning resources to microservices is an NP-hard problem considering the multi-dimensional resources. Due to the various run-time parameters of microservices and hardware configurations, it is time-consuming to find the optimal results given the large solution space [13]. Thirdly, there are trade-offs considering when and how to trigger the auto-scaling algorithms to handle the situations with high dynamics while ensuring the QoS requirement. We propose multi-faceted, data-driven, and adaptive auto-scaling approaches for microservices to address the above discussed challenges.

B. Our Contributions

In this paper, we investigate the promising approach called **CoScal** that combines horizontal scaling, vertical scaling, and brownout techniques to address the resource scaling problem in microservice-based cloud computing environments. The proposed approach exploits the advantages of abovementioned individual techniques, including the high availability of horizontal scaling, fine-grained control of vertical scaling, and self-adaptability of brownout. This combined auto-scaling approach is more effective under diverse and complex workload scenarios derived from Alibaba traces [14]. The key **contributions** of this paper are:

- A performance analysis of horizontal scaling, vertical scaling, and Brownout investigates the trade-offs in execution costs and performance of microservices.
- A Reinforcement Learning (RL) based scaling algorithm for decision making to optimize the performance of microservices, such as response time, connection time, and the number of failed requests.
- A prototype system implementing the proposed approach is evaluated with realistic traces and extensive experiments.

The rest of the paper is organized as follows: Section II discusses the related work for auto-scaling microservices in the cloud computing environment. The motivation and performance analysis of the applied techniques in this work are introduced in Section III. Section IV depicts the system model of our proposed approach. The proposed algorithm based on RL is detailed introduced in Section V. The Section VI illustrates the details of our experiments conducted using a dataset derived from realistic traces and demonstrates the feasibility of our approach to improve the resource scaling of cloud data centers. Finally, conclusions along with the future directions are given in Section VII.

II. RELATED WORK

In this section, we discuss existing research works in microservices scaling. The current scaling approaches for performance optimization in microservices can be mainly categorized into three categories: the performance model, resource orchestration, and prototype systems.

A. Performance Model

The performance model is applied for modeling the resource usage of the system and the performance of cloud services. Thus, resource management policies can be designed to fit the performance requirements. Nishtala et al. [24] [25] proposed scalable QoS-aware resource management approaches for latency-critical and co-located services in traditional cloud data centers. The proposed approaches are based on heuristic and reinforcement learning techniques to achieve energyefficient goals, which require no service or system-specific information. Liu et al. [15] analyzed the bottlenecks of the mainstream microservice applications and applied multiple machine learning models to schedule resources. The first model is used for finding the resource area that satisfies the microservice performance. The second model investigates the trade-offs between the QoS and allocated resources, and the third model dynamically adjusts resources according to the system states. Chang et al. [26] designed an automatic resource planning approach and modeled application performance based on an empirical model. Khazaei et al. [27] constructed a performance model via analyzing microservice platforms. This approach applies the Markov chain model to represent microservice resources, virtual machine (VM) resources, and physical machine resources. The number of tasks, number of microservices, and number of VMs have all been considered as the system states. Although the performance model is a popular approach for performance optimization, the model needs to be reconfigured or retrained when the service or environment changes dramatically. Gan et al. [28] exploited the prediction method for QoS violations, which utilizes a set of machine learning models and large-scale history data to locate the microservices that lead to QoS violations. The approach can relieve the QoS degradation by reallocating hardware resources. Qiu et al. [10] proposed a fine-grained control framework to relieve resource competition and optimize resource utilization and analyzed the microservice dependency relationship, microservice chains, and key call paths. Kannan et al. [16] modeled the multiple-stage tasks as Directed Acyclic Graph (DAG) and used DAG to estimate the task completion time. Yu et al. [17] proposed Microscaler framework by using service mesh to record the resource usage behaviors and applied online learning and heuristic approaches to obtain the near-optimal solutions for resource demands.

Scaling commands must be configured manually based on the performance model in all these approaches. Thus, they are far slower to react to load variations. We apply neural networkbased workload prediction to achieve high adaptation to load variances to solve this challenge.

B. Resource Orchestration

Other works apply vertical scaling or horizontal scaling that are not required to be configured manually based on the performance model, which utilizes resource orchestration to optimize resource provisioning for microservices by allocating and managing resources efficiently. Suresh et al. [29] investigated the overload control mechanism for a microservice cluster with complex dependency. The proposed approach considers the resource sharing problem under a multi-tenancy scenario, in which all the tasks are modeled as DAG. The scheduling is processed at the workflow level and requests level separately. Zhou et al. [18] pointed out that monitoring and collecting the data of each microservice under largescale and high-dynamicity scenarios is not feasible. Therefore, authors proposed a workload control approach to maintain a self-adaptive threshold and algorithm for each microservice. Each microservice can shed the loads independently with quite limited communication costs. Rzdaca et al. [19] designed Auto-pilot as an auto-scaling approach to scale based on the change of workloads. Auto-pilot combines time-series analysis and scaling the number of microservices with related CPU and memory amount. Kwan et al. [8] designed an approach that combines vertical and horizontal scaling methods, which deploy microservices to suitable hardware to efficiently reduce service delay due to burst requests. Hou et al. [20] introduced a power-aware and latency-aware scheduling approach to scale resources from micro and macro perspectives. They also apply the decision tree and tagging method to achieve fast resources matching. He et al. [30] took advantage of genetic and heuristic algorithms to find the optimized microservice deployment place under an edge-cloud environment. Zhang et al. [21] proposed a predictive RL algorithm to horizontally scale containers based on the Autoregressive Integrated Moving Average (ARIMA) model and neural network model, which can ensure the predictability and accuracy of the scaling process. Rossi et al. [22] introduced RL-based approaches to control the horizontal and vertical scaling for containers to increase system flexibility under varying workloads, which also accelerate the learning process by exploiting different degrees of knowledge about the environment. However, this work was only evaluated with synthetic workloads. Gias et al. [23] proposed a model-driven scaling approach, named ATOM, for microservices via analyzing a layered queueing network of applications. This approach can dynamically adjust the number of replications via horizontal scaling.

In resource orchestration approaches, limited hybrid scaling for microservices has been proposed, and most of them focused on a single scaling technique, either vertical or horizontal, but not on both. The existing hybrid approaches are threshold-based or heuristic that requires heavy manual configurations. Unlike these efforts, we apply a hybrid scaling technique that combines vertical, horizontal, and brownout together, making adaptive decisions via RL.

C. Prototype Systems and Tools

Kubernetes [31] and Docker Swarm have become the dominant systems for managing microservices. However, the scaling techniques in these systems are primarily thresholdbased or static policies. Kiss et al. [32] implemented a generic microservice orchestration platform for heterogeneous cloud clusters, which can auto-scale resources without binding to specific applications. Zhou et al. [33] analyzed and compared the interaction patterns of open-source microservicebased applications. They found that the investigated open source application called TrainTicket. Xu et al. [12] proposed a prototype system for managing co-located interactive and batch microservices based on the brownout approach and workloads deferral to achieve an energy-efficient cloud data center.

In these prototype systems and tools, threshold-based heuristic algorithms are applied popularly. However, heuristic

Anneach	S	caling Technic	ues	workloa	d prediction	Scl	neduling Policy
Арргоасн	Vertical	Horizontal	Brownout	Linear	Non-Linear	Heuristic	Supervised Learning
Liu et al. [15]	\checkmark				\checkmark		\checkmark
Qiu et al. [10]							\checkmark
Kannan et al. [16]	\checkmark			\checkmark		\checkmark	
Yu et al. [17]		\checkmark			\checkmark	\checkmark	
Zhou et al. [18]				\checkmark		\checkmark	
Rzdaca et al. [19]	\checkmark	\checkmark		\checkmark		\checkmark	\checkmark
Kwan et al. [8]						\checkmark	
Hou et al. [20]	\checkmark			\checkmark		\checkmark	
Zhang et al. [21]		\checkmark		\checkmark			\checkmark
Rossi et al. [22]	\checkmark						
Gias et al. [23]		\checkmark					\checkmark
Xu et al. [12]		\checkmark	\checkmark	\checkmark		\checkmark	
CoScal (Our Approach)	\checkmark	\checkmark	\checkmark		$\overline{\mathbf{v}}$		$\overline{\mathbf{v}}$

TABLE I: Comparison of related work

algorithms can find a solution quickly, the performance needs of manually tuning various configurations, especially in an environment with high dynamics [22]. Unlike these approaches, we utilize the RL-based approach with an adaptive nature to learn and make good decisions via interactions with the environment.

D. Critical Analysis

This paper contributes to the growing body of research related to the microservice area. The Table I compares the proposed approach (**CoScal**) with related works based on the scaling techniques, workload prediction techniques, and type of scheduling algorithms. Given the contributions of the existing works, it is important to highlight the key difference between our work and the prior ones. To the best of our knowledge, the proposed work is the first to offer a multifaceted scaling approach based on vertical scaling, horizontal scaling, and brownout. Prior works only applied included one or two techniques. We also utilize gradient recurrent units for accurate workload prediction and apply an RL-based algorithm for resource management to offer flexible adaption for a highly dynamic environment.

III. MOTIVATION: PERFORMANCE TRADE-OFFS IN INDIVIDUAL SCALING TECHNIQUES

In this section, we conduct experiments to investigate the effects of individual scaling techniques (vertical, horizontal, and brownout) on important performance metrics in microservicebased computing environments. The results illustrate that applying individual scaling techniques such as horizontal scaling or vertical scaling does not yield desired results due to tradeoffs.

A. Use Case Setup

In this motivational study, we have used a microservice application named Sock Shop¹, which is an online e-commerce website for products sale. The Sock Shop application consists of multiple microservices representing different components of the application, including front-end User Interface (UI), catalog, carts, and some other supporting microservices, which are deployed on two Docker Swarm nodes.

	Request Distribution				
Microservice Type	Home	Orders	Catalog	Others	
Percentage	83.8%	7.2%	5.4%	3.6%	

In this study, we consider the browsing workload for the Sock Shop application. The users can view the information of the items, add or remove items in the carts, pay for the items. The used microservices in the Sock Shop application and the request distribution are summarized in Table II. We have investigated the high load use case representing the system's over-utilized (when utilization is above a predefined threshold) condition from Alibaba's sample trace [14]. The sample trace contains 1200 lines of data representing resource utilization in 6000 seconds, where the data is collected every 5 seconds. We use JMeter² toolkit to generate workloads for the Sock Shop application, which simulates the scenario of users visiting the website for shopping. JMeter is an open-source and completely Java-based application. It is mainly used to conduct stress tests of the target server and validate test-case functionalities of service endpoints. Only one node is initially deployed with microservices in our setup, and the second node is kept in an idle state. In the horizontal scaling, the replicates of Sock Shop microservices are deployed on the second node. In the vertical scaling, the CPU is configured to be scaled to a maximum of 8 cores, and the memory can be scaled to a maximum of 3.5 GB. In the brownout approach, the optional microservices (3 out of 14, i.e., the recommendation engine) are temporarily deactivated (due to space limitation, please see the more detailed descriptions in Section V-A.).

We apply the individual scaling techniques separately to evaluate system performance regarding connection time, response time, and success rate of requests. *Connection time* represents the time required to establish a connection between the user and the target microservice server. *Response time* represents the total time from the start of the request to the reception of the response. *Success rate* represents the percentage of successful requests (i.e., expected response is received for request) received by the user.

B. Comparison between Scaling Techniques

Fig. 1a represents the connection time achieved by three different scaling techniques. The collected data is obtained



Fig. 1: Performance comparison of different scaling techniques when applied separately

every 200 ms, and the average value is calculated based on every 5-minute interval. We choose 5-minute interval based on Alibaba's practice where 5-minute minimum time interval is considered for scheduling decisions. A smaller time interval can lead to frequent scaling costs, while a longer time interval cannot respond to the system change promptly. As this is a motivational example, we only conduct the experiments one time. We can observe that the vertical scaling converges faster to a stable connection time than horizontal and brownout scaling. However, at the initial period of scaling, its performance for connection time is not as good as brownout. The brownout approach initially achieves the best connection time as it deactivates optional parts that help provide additional resources for the front-end connections. After some periods, the connection time keeps growing because all the available optional components have been deactivated. Horizontal scaling can achieve better connections than brownout since more resources are provided regarding the number of nodes. However, the extra communication costs involving two replicas of services deployed on different hosts induce higher connection time in horizontal scaling than vertical scaling, as the communication in vertical scaling is negligible as it remain on the same host.

Figs. 1b and 1c represent the response time and success rate achieved by different scaling methods. The vertical scaling achieves the best performance with the lowest response time and the highest success rate. The reason is that the vertical scaling on the local machine can improve the system performance without incurring additional overheads. In contrast, Brownout can incur some delays induced by its resource optimization technique. Moreover, the horizontal scaling leads to poor performance in response time and success rate of user requests due to its associated overhead, such as boot-up cost of new instances and frequent scaling operations, which is crucial in highly time-sensitive microservice environments. Although vertical scaling achieves the best performance when compared to the other two techniques, it is not effective when the local machine is overloaded and constrained by the maximum resource capacity of an individual physical machine.

To summarize, individual scaling techniques may not yield the desired result in complex conditions. A multi-faceted scaling can solve the limitations of individual techniques. For instance, when a limited capacity is left on a hosted machine (i.e., vertical scaling hits resource limitations), the brownout technique can be triggered to maintain the connection time by relieving resource overloading. Meanwhile, the horizontal scaling can be applied when enough resources are available to host additional workload, and optional components can be reactivated further. However, the horizontal scaling requires more time to take effect compared to brownout as it requires significant time to activate host and initialize new container instances. Thus, under the high-variable workloads, the choice and combination of scaling techniques play an important role in delivering reliable services, especially in latency-sensitive microservice environments. Therefore, to overcome the limitations of individual scaling techniques, a joint, data-driven, and adaptive auto-scaling framework is necessary to provide optimal scaling decisions based on the current infrastructure and workload conditions.

IV. SYSTEM MODEL

In this section, we describe our proposed auto-scaling system named as *CoScal* that is composed of three modules, including *Workload Analyzer*, *Workload Predictor*, and *RL-based Resource Scaler*, as shown in Fig. 2. It is noted that this is a general system model, and the modules can be extended to new workload analysis techniques and more advanced RL-based resource scheduling solutions. The details of these modules are given in the following subsections:

A. Workload Analyzer

The Workload Analyzer module analyzes system workload and processes the raw data generated by log files. In this module, the *Requests Handler* handles the user-generated requests, which allocates the requests to hosts. The *Preprocessor* extracts the required meta information of workloads (e.g., time series and resource utilization metrics). Later, the *Workload Predictor* module takes our inputs to perform its stated operations (described in Section IV-B).

Fig. 3 shows a schematic diagram of load generation based on JMeter. In this case, The JMeter Master and Workers (Slaves) are regarded as part of the *User Requests* and *Requests Handler* components. Web servers are the infrastructure to provision resources. Based on our practice, when the number



Fig. 2: The CoScal system model

of requests is huge, a single JMeter server mode may not function well due to overloads, then the master-worker will be used collaboratively. With each test case, the data of collected metrics will be stored locally, such as timestamps, response time, connection time, and the number of failed requests.



Fig. 3: Schematic diagram of requests generation by JMeter

B. Workload Predictor

The Workload Predictor module in Fig. 2 is responsible for estimating the expected load in the next scheduling window, which guides scaling methods to decide the amount of resources to be added or removed. This module can be constituted by different predictive models such as *Multivariate Time Series Forecaster* [34], which accepts the pre-processed information of workloads from the *Workload Preprocessor* and predicts the future workload based on suitable predictive models, such as long short-term memory (LSTM) or gated recurrent unit (GRU) model. The responsibility of the *Tensor-Flow Server* is to manage the *Trained Model*, including the training process of the models and updating the trained models if necessary, and finally deploying trained models. As shown in Fig. 2, the communications between *Workload Analyzer* and *Workload Predictor* transfer the inputs with past system data to the *TensorFlow Server* (the left-bottom corner of this module). It provides feedback such as predicted workload level to *Requests Handler*.

To generate user requests for the prototype system, the load generator toolkits, e.g., JMeter, can be utilized. And the typical microservice application, like Sock Shop [35] as introduced in Section III, can be deployed as benchmark application. The profiling approach based on the stress test can establish a model to represent the relationship between the number of requests and resource utilization (see Section V-A for more details.)

C. RL-based Resource Scaler

The *CoScal Controller* is the core component in this module which makes decisions on scaling strategies based on the RL framework. Compared with static performance models and heuristic-based approaches that suffer from model reconstructions and retraining problems, the RL-based approach is well suited for learning resource scaling policies to address dynamic system status [36]. The *RL-based Resource Scaler* in Fig. 2 receives the user requests from *Request Handler* and information like predicted workloads from *Workload Predictor* in our system. After executing the scaling techniques, the *CoScal Controller* collects the data of performance metrics, such as response time, and sends the data back to the users. The data is the response to the user request to the application.

In the *RL-based Resource Scaler* module, the supported techniques work in different manners. Vertical scaling is applied to the local machine with multiple resources, e.g., CPU, memory, and network capacity. The vertical scaling is faster compared with other techniques as shown in Section III, it can be leveraged initially. The brownout technique is also applied to the local machine, where the optional microservices can be dynamically deactivated to relieve the overloaded situation that vertical scaling cannot handle alone. Horizontal scaling is applied at the node level by adding or removing additional nodes into the system. Considering the communication costs of horizontal scaling, this technique can be applied when vertical scaling and brownout cannot handle the workloads level and keep the system in a normal state.

V. CoScal: A MULTI-FACETED AUTO-SCALING IN MICROSERVICE ENVIRONMENTS

This section introduces the key elements to realize our system model, including the performance profiling, neural network-based workload prediction and RL-based policy for scaling microservices.

A. Performance Profiling

To estimate the approximate resource usage corresponding to a different amount of workloads with the Workload Analyzer module in Fig.2, we use a deep neural network model to profile the performance of machines. To simulate the resource utilization in a realistic environment, we use the data derived from Alibaba traces, including workload traces of 4000 machines' containing resource usage data for 8 days [14].

The detailed profiling procedures are as follows: we consider the scheduling interval as 5 minutes, as 5 minutes are the minimum monitoring interval for data collection. A short scheduling interval can make too frequent scaling operations to influence system stability, while a long scheduling interval can delay the adaptive scaling decisions. Firstly, we apply stress test by gradually increasing the number of requests over 5 minutes with 200 requests in each test case and sending these requests to the host. As the number of requests increases, the host's resource utilization also increases. Then we use the nmon [37] performance monitoring toolkit to record the change of utilization in the host as per the number of requests sent. In this way, a detailed mapping relationship between resource utilization and the number of requests is obtained. Based on the profiled data, we apply a Multi-Layer Perceptron (MLP) with three layers to establish a model to represent the relationship efficiently. Compared with the traditional regression-based approaches, the MLP-based approaches can capture a more accurate relationship between workloads and utilization. Finally, we can convert the host utilization into the number of requests dispatching to our system for any utilization level with the trained model.

As shown in Fig. 4, we demonstrate the results between CPU utilization and the number of requests based on the tested host. We can observe the CPU utilization change under the different number of requests. The host consumes about 25% CPU utilization when it does not handle any requests. This is due to the resource consumption by the operating system and deployed applications. The CPU utilization increases gradually with the increased number of requests. In this experiment, a host can accept at most 9000 requests. We can also obtain the relationship between other resource types and the number of requests. Based on our experiments for the Sock Shop application, the CPU is a dominant resource type indicating that the deployed application is mainly compute-intensive.



Fig. 4: The number of requests and the corresponding CPU utilization

TABLE III: Defined Workload Levels and Corresponding Meanings

Workload Level	Meaning
Level 0	No loads are allocated to the host, and the hosts can be switched into the low-power mode or turned off, i.e., the resource utilization is 0.
Level 1	The host is under light load, with low latency and low CPU and memory usage, i.e., the resource utilization ranges from 0% to 25%.
Level 2	The host is at a medium load level, the delay is low, and the QoS is not significantly affected, i.e., utilization is between 25% and 50%.
Level 3	The host is potentially under heavy load, the CPU or memory usage is high, and QoS may have affected; thus, the scaling techniques may be triggered if needed, i.e., utilization is between 50% to 75%.
Level 4	The host is overloaded, the QoS will be significantly impacted, and the effective scaling approaches should be performed, i.e., utilization is between 75% to 100%.

B. Neural Network-based Workload Prediction

To address the prediction of resource usage levels in the Workload Predictor module in Fig. 2, we realize a neural network-based workload prediction approach. To reduce the state space of the RL model, we consider dividing the workloads into several levels (the exact number of levels configured) representing different levels of utilization. Moreover, this also helps to apply similar scaling techniques for workloads at the same level. This helps to significantly reduce the state-action space in the RL-based approach. We divide the workloads into five levels that can represent the degree of overloads, as shown in Table III. The overloaded threshold is configured as 75%, as it has been evaluated in our previous work [12] that it can achieve good trade-offs between resource utilization and QoS.

To accurately predict the server load status in the next time interval, we apply the Multivariate Time Series forecasting (MTFS) method, which converts multivariate time series forecasting into supervised learning. The MTFS algorithm can be applied to any time-related dataset, and it is highly correlated to temporal aspects and contains all the data from the previous time intervals. The MTFS makes each generated supervised learning sequence to have sample labeled datasets in the algorithm.

The Algorithm 1 shows the pseudocode of the MTFS algorithm, which generates labeled time series data for supervised learning. With the original time series dataset E, the algorithm first generates an empty matrix S to store the supervised learning sequence. After the labeled data is generated for supervised learning, to achieve an accurate prediction for workloads, we apply the Gated Recurrent Unit (GRU) [38] derived from Recurrent Neural Network (RNN) [39], which has been validated to have better performance in time series related prediction than traditional RNNs. Although RNN can use its memory to process a set of inputs sequentially, it is inefficient to learn long-term memory dependencies due to gradient vanishing, while GRU can overcome this limitation by merging the data processing elements (gates).

Table IV shows the mean square errors (MSE) of actual utilization and predicted utilization for Alibaba workloads during different periods. The MSE has lower values about

Algorithm 1: CoScal: workload preprocessing

```
Input : Multivariate time series dataset E, time intervals T, k
            time-related variables, and a dataset F needs to be forecast
   Output: Supervised learning dataset S, each row of it has 2k + 3
            data
1 Initialize an empty matrix S to record supervised time series data
2 for t from 1 to T do
       Record data in current time interval: C(t) \leftarrow E(t)
3
       if t = 1 then
4
            Record NONE
5
6
       else
            Record data in last time interval: L(t-1) \leftarrow E(t-1)
7
8
       end
       if t = n - 1 then
9
10
            Record NONE
11
       else
            Record data in next time interval: F(t+1) \leftarrow E(t+1)
12
       end
13
14
       Merge L(t-1), C(t) and F(t+1) together into S(t):
15
       S(t) \leftarrow \{L(t-1), C(t), F(t+1)\}
       L(t-1), C(t) are marked as samples in supervised learning
16
       F(t+1) are marked as labels in supervised learning
17
       if S(t) contains NONE then
18
19
            Remove S(t)
       Update S
20
21 end
```

TABLE IV: MSE of our workload prediction algorithm

Time (minutes)	100	200	300	400	500
MSE (10^{-3})	3.3	2.8	4.2	4.8	6.2

 3×10^{-3} to 7×10^{-3} (in [40], the RNN-based and LSTMbased approaches have the MSE values about 4×10^{-2}), which demonstrates that our workload prediction algorithm has a good performance in utilization prediction and validates that the neural network based approach is suitable for cloud workloads prediction.

C. RL-based Resource Scaling

The RL-based approach solves sequential decision-making problems by modeling the problem as Markov Decision Process [41]. At each time interval t, the system is at a state $s_t \in S$, and performs an action $a_t \in A$ based on policy π_{θ} , where θ are parameters configured in model. The statespace S is mapped with action space A. In the following time intervals, the current system state can reach another state by taking actions and obtaining reward $r_t \in R$, calculated via reward function $r(s_t, a_t)$. The next state will only rely on the current state and the performed actions on it. The rewards represent the benefits that can be achieved by transiting the state from s_t to s_{t+1} . When transferring states, there is also a transition probability of presenting the possibility to take different actions. The objective of RL is to maximize the expected cumulative reward by optimizing policy. Qlearning [42] is a typical type of RL to maximize the value function $Q_{\pi_{\theta}}(s, a)$. The value function estimates the expected cumulative reward of state s with action a under policy π_{θ} . Consider action a_t is selected at time interval t, and at time interval t + 1 with reward $r(s_t, a_t)$, the value of Q function can be can be updated as:

$$Q(s_t, a_t) = Q(s_t, a_t) + \alpha[r(s_t, a_t)] + \gamma max_{a'}Q(s_{t+1}, a') - Q(s_t, a_t)]$$
(1)

where $\alpha \in (0,1]$ is the learning rate and $\gamma \in [0,1]$ is the discount factor. The Equation defines a mapping table containing states with actions and their expected value. The learning process happens in the form of S * A - > R over time to achieve optimized results via iterative trials. We consider the state S as workload level $S = \{0, 1, 2, \ldots, W\}$, where $W \in Z$ represents the maximum level of workloads and is a non-negative value.

To illustrate our RL-based multi-faceted scaling framework (CoScal), let us assume that we have a set of physical machines $P = (pm_1, pm_2, \dots, pm_K)$ as infrastructure to provision resources for microservices. Each pm_k can be represented with a tuple $U_k = (u_k^1, u_k^2, \dots, u_k^I)$, where u_k^i represents resource utilization of type i with total I types on physical machines. For each pm_k , the actions performed on it, which are denoted as $a_k^i = \{h_k, v_k^i, b_k\}$, where $h_k \in [-n, n]$ presents the $n \in \mathbb{Z}$ number of replicates in horizontal scaling, $v_k^i \in [-m, m]$, where $m \in R$, represents the amount of resources via vertical scaling for resource type i, and $b_k \in \{0,1\}$ represents the whether brownout is triggered $(b_k = 1)$ or not $(b_k = 0)$. The positive and negative values of h_k and v_k^i represent more resources are added and removed respectively, and value 0 means no change will be performed. Considering the total number of physical machines is K, the final set of actions is the Cartesian product of the sub-action sets as: $A = \prod_{k=1}^{K} \prod_{i=1}^{I} a_k^i.$

The objective of our technique is to find the suitable configuration of resources by dynamically adjusting the provisioned resources to adapt to changes in the environments, e.g., the load fluctuations. However, the amount of scaled resources is limited by the available resources on physical machine pm_k and the minimum resources allocated to microservices. To avoid unnecessary vertical scaling, we consider adding an action a_q to make decisions from the global view. Let us consider the scenario that when the system with $P = \{pm_1, pm_2\}$ is running at the normal states that vertical scaling is sufficient for adjusting resources. However, when unpredictable workloads arrive, the physical machines are overloaded. To handle such bursts, horizontal scaling must be performed on the system. If both pm_1 and pm_2 are completed with vertical scaling, two more physical machines are added. However, the system may only need one more physical machine to keep the system at the normal state. Therefore, to optimize resource usage, the action a_q is required to make scaling decisions based on a global view.

CoScal incorporates the offline training and online training approach together to achieve optimized actions. Once the load change is identified, CoScal can select an action in response to the QoS or performance degradation of deployed applications. Given s_t as the observed state, a policy exploits the knowledge of previous decisions (offline). It evaluates the performance of new selected actions (online), improving the mapping relationship between states and actions. We apply the ϵ -greedy policy, [36] a standard policy to balance exploration

Algorithm	2:	CoScal:	General	Procedure
	<i></i>	CODCar.	OUTUR	TIOCCUUI

Inp	ut : Table $Q(s, a)$ contains all state/action pairs from
	experience pool by offline training, time intervals T ,
	probability of random action ϵ , learning rate α , discount
	factor γ
1 Init	ialize system status, monitoring model;
2 for	t from 1 to T do
3	$W_{t-1} \leftarrow$ Workloads level at time interval $t-1$;
4	$W_t^p \leftarrow$ Predicted workload level according to Algorithm 1;
5	$U_k^t \leftarrow \text{Resource utilization of } pm_k \text{ at time interval } t;$
6	if $W_t^p - W_{t-1} \neq 0$ then
7	Choose a action from action set A with ϵ probability, or
	select an action with the $\max(Q(s_t, a_t))$;
8	Execute a_t according to Algorithm 3;
9	if online training is triggered then
10	$s_{t+1} \leftarrow$ system state at time interval $t+1$;
11	$r_t \leftarrow$ reward calculation by Equation (4);
12	Update Q value: $Q(s_t, a_t) = Q(s_t, a_t) +$
	$\alpha[r(s_t, a_t)] + \gamma max_{a'}Q(s_{t+1}, a') - Q(s_t, a_t)];$
13	end
14	Store transition $(s_t, a_t r_t, s_{t+1})$ in experience pool;
15	end
16 end	

and exploitation [43]. In ϵ -greedy policy, a random action with probability equals ϵ is selected. Otherwise, it chooses the action with the maximum Q value.

The final objective of *CoScal* is to improve the QoS of services and utilization of physical machines. The reward is modeled based on this objective which composes of two parts. For the QoS, we choose to use response time as the metric, representing the latency from submitting the request to completion. We model the reward of response time $R_{qos}(rt)$ based on the maximum acceptable response time with RT_{max} . As shown in Equation (2), when the system is working at the normal status, the reward is 1. However, when the system performance violates the RT_{max} , the reward converges to 0, punishing the actions that lead to overloaded situations.

$$R_{qos}(rt) = \begin{cases} e^{-\left(\frac{rt - RT_{max}}{RT_{max}}\right)^2} & , rt > RT_{max} \\ 1 & , rt \le RT_{max} \end{cases}$$
(2)

As for the reward of resource utilization (please note that the utilization can also be referred as resource costs in our model), we model it as shown in Equation (3). Here, U_k^{max} defines the maximum utilization threshold of pm_k , which is also the highest utilization of all resource types (CPU and memory utilization are both considered). u_k is the current utilization of pm_k . Higher utilization without violating the threshold can contribute positively to the reward, while utilization above the threshold can undermine the reward.

$$R_{util}(u_k) = \begin{cases} \frac{\sum_{k=1}^{K} U_k^{max} - u_k}{K} + 1 & , u_k \le U_k^{max} \\ \frac{\sum_{k=1}^{K} u_k - U_k^{max}}{K} + 1 & , u_k > U_k^{max} \end{cases}$$
(3)

Equation (4) shows the final reward value based on response time and resource utilization, in which the higher values of R_{aos} and lower values of R_{util} can increase the total reward.

$$r(s_t, a_t) = \frac{R_{qos}(rt)}{R_{util}(ut)}$$
(4)

Algorithm 3: CoScal: Action Execution

Input : Time interval t, action sets $a_t = \{a_k^i(t) | i \in \{0, 1, \dots, I\},\$ $k \in \{1, 2, \dots, K\}\}, a_k^i(t) = \{h_k(t), v_k^i(t), b_k(t)\},\$ selected action at time t with horizontal scaling operation $h_k(t) = n$, vertical scaling operation $v_k^i(t) = m$, and brownout operation $b_k(t) \in \{0, 1\}$ 1 for k from 1 to K do /*Vertical scaling*/ 2 for i from 1 to I do 3 if m > 0 then 4 Add m resources of type i on local machine; 5 else 6 Remove m resources of type i on local machine; 7 8 end end 9 10 end 11 /*Brownout control*/ 12 if $b_k(t) > 0$ then Deactivate microservices via brownout; 13 14 else Brownout will not be triggered; 15 16 end /*Horizontal scaling*/ 17 if n > 0 then 18 Add *n* microservices replicates ; 19 20 else 21 Remove n microservices replicates; 22 end

Algorithm 2 shows the general procedure of *CoScal*. The algorithm initializes the monitoring model to collect system status for the RL process (line 1), including workloads level, utilization, and relevant metrics at each time interval (lines 4-5). These observations constitute a state in the RL framework. If the workload level changes (line 6), scaling approaches should be applied to ensure QoS or optimize resource usage. Thus, the *Q*-learning process will be started by choosing actions from the experience pool and transiting the current system state to another one (line 7). The actions are executed based on the Algorithm 3 (line 8). *CoScal* also supports online training, after *CoScal* transits the state s_t to s_{t+1} , *CoScal* firstly stores the transition (s_t, a_t, r_t, s_{t+1}) into the experience pool. The reward r_t is applied to evaluate the goodness of (s_t, a_t) (lines 9-14).

Algorithm 3 shows the action execution process of CoScal. The system states and corresponding actions are provided by the Algorithm 2, which includes the decisions of horizontal scaling, vertical scaling, and brownout. The algorithm will first check the decision for vertical scaling. If the vertical scaling should be performed, then resources of the specific type will be added or removed (lines 1-10). After that, the brownout will be examined. If the brownout is triggered, the optional microservices are deactivated in this time interval (lines 11-16). And finally, the horizontal scaling will be checked and executed by increasing or decreasing the number of replications. We consider the execution sequence based on the execution costs of different scaling techniques as discussed in Section III, as the vertical scaling can be completed within the shortest time while the horizontal scaling brings much longer processing time and communication costs (inter-process communication costs between processes on different hosts for user authentication, database access, remote function call, image retrieval and etc.) than the other two techniques as shown in the motivational example in Section III.

Decision-making complexity of CoScal: After the RL model is trained based on historical data, the system states and corresponding actions can be stored in a lookup table. At each iteration to access and modify the lookup table, CoScal applies the open addressing technique to resolve hash collisions to ensure the access and modification operations to take negligible time. The open addressing technique has the computational complexity as O(1).

Notes on model re-training: To be noted, when microservices are updated and new optional components are added, the RL model can be re-trained to improve the performance, as the brownout part in *CoScal* is influenced by the identification of the optional parts. If the model is not re-trained, the added components will be regarded as the mandatory ones.

D. Auto-adapter for Unexpected Workloads Change

Although our workload prediction algorithm can achieve high accuracy, fluctuations of cloud workload level during the stable period (workload level remain unchanged) are usually unpredictable. For instance, Fig. 5a shows a sample of workload changes within 500 minutes. It can be observed that during the first 100 minutes, the workloads change from workload level 3 to level 4 and return to level 3 within a short time (as the workload level is the average level during a period of time, the changes represent the total amount of workloads change significantly). This often happens when there are bursts in the system. Such workload changes are difficult to be predicted due to randomness.

Microservices are more sensitive to resource fluctuations than traditional applications [10]. Therefore, bursts can cause QoS violation, resource wastage, and system overload. To address this challenge and complement the prediction algorithm, we introduce a component named *auto-adapter* that detects and adapts to these changes. The auto-adapter is integrated into the *CoScal Controller* module in Fig. 2. Auto-adapter collects the information at the first minute of each time interval and examines whether the actual workload level (collected via performance counters) equals the predicted one (via workload predictor module). If not equal, auto-adapter applies vertical scaling to optimize resource usage by adding or removing resources according to the predicted and actual workload difference.

The auto-adapter needs to be lightweight to avoid excessive resource usage incurred by the auto-adapter. We measure the resource usage in one hour caused by auto-adapter as shown in Fig. 5b, which shows that auto-adapter only costs about a maximum of 3% extra CPU and 1% extra memory resources. This additional limited resource usage is acceptable considering its optimization effects on resource usage.

VI. PERFORMANCE EVALUATION

To evaluate the performance of our proposed approach for scaling microservices, we conduct experiments in the container-based prototype system. We first present the experimental settings in Section VI-A and then introduce the baselines in Section VI-B. The results and analysis are presented in SectionVI-C.



(a) The fluctuations of workloads (b) Resource usage (auto-adapter)



A. Experimental Setup

We built a cluster with heterogeneous nodes for the microservice-based application deployment. The cluster consists of four physical machines, including a machine with an Intel Core i7-9700 CPU and 16GB of RAM, two machines with Intel Core i7-4790 and 4GB of RAM, and a machine with two Intel Xeon E5-2660 CPUs and 48GB of RAM. The containers or microservices are managed through Docker Swarm, a container orchestrating tool. We deployed the Sock Shop application, a microservice-based application. The computing environment has been configured with Java (version 1.8), Python (version 3.7), and TensorFlow (version 2.2.0) development toolkits. We use the cluster-trace-v2018 of Alibaba dataset³ that contains workloads of machines. And we use 500-minute data for evaluations.

To enable our scaling algorithms to support the brownout mechanism, we categorize the individual microservices in Sock Shop into two categories: mandatory microservices and optional microservices. Mandatory microservices include microservices that are are the key components (e.g. database-related microservices) of the system, and the optional microservices (e.g., recommendation engine) can be dynamically activated or deactivated, In our testbed, the CPU can be scaled from one core to a maximum of 8 cores, memory can be scaled from 2 GB to 3.75 GB, and for horizontal scaling, the maximum number of hosts is set to 3.

B. Baselines

We compare the performance of *CoScal* with several stateof-the-art algorithms for scaling microservices. Some baselines have been used in dominant microservice platforms, including Docker Swarm and Kubernetes.

DoScal [22]: it is derived from the resource scaling approach implemented in native Docker Swarm, which is mainly based on vertical scaling. For instance, when one of the microservice instances is overloaded, DoScal can reallocate the microservices equipped with sufficient resources.

KuScal [31]: this scaling approach has been used in Kubernetes, which is mainly based on horizontal scaling that can dynamically increase or decrease the number of replicates. KuScal decides on how many replicates should be added or

³https://github.com/alibaba/clusterdata

removed based on the resource fluctuations, e.g., CPU and memory, on the current servers.

HyScal [8]: it is a hybrid scaling approach for microservices. Apart from scaling CPU utilization, the HyScal algorithm also scales memory resources. The HyScal algorithm applies a similar approach to the KuScal, allocating resources and monitoring resource utilization.

C. Experiment Analysis

The results are collected from JMeter log files, including connection time, response time, and a number of failed requests. The lower values of connection time and response time represent better QoS performance. A lower number of failed requests is desired, indicating a higher number of successful requests processed by the deployed applications.

We conducted 500-minute experiments to accurately observe and analyze the advantages and disadvantages of the scaling approaches in different periods. We demonstrated the average value of each metric with an interval of 100 minutes (i.e., 100 minutes, 200 minutes, to 500 minutes) to represent the statistical change, and the experiments are repeated 3 times to avoid randomness.

Connection Time: Fig. 6a shows the comparison of the average value of connection time during 500 minutes. The results of the baselines vary between 0.75 ms and 1.25 ms. More specifically, HyScal keeps the connection time at a high level, from 1.1 ms to 1.25 ms, and finally reaches around 1.25 ms at peak since the number of requests increases during the observed period. KuScal can achieve a better performance than HyScal with about 1.1 ms. The reason lies in that KuScal can scale resources more sufficiently than HyScal that pre-configures some limitations on resources allocated to microservices. The connection time achieved by DoScal is 1.1 ms. CoScal can achieve the best performance compared with the baselines, and it can reach around 0.8 ms within 500 minutes. Although the results of the CoScal vary slightly, the connection time is still 19.3% to 29.3% lower than the average value of the other baselines.

Response Time: Fig. 6b illustrates the comparison of the average value of response time in the first 500 minutes. The average response time values for all four scaling approaches show a decreasing trend. DoScal remains at the highest value but shows an overall decreasing trend from 419.9 ms to 242.1 ms. KuScal and HyScal show irregular fluctuations, which are caused by the different number of requests during different time periods, and their values worse than *CoScal*. For *CoScal*, it maintains the lowest response time throughout the observed time and finally drops to 198.4 ms. *CoScal* can reduce 16.1% response time compared with the baselines, which shows that combining different techniques can achieve better performance.

Number of Failed Requests: Fig. 6c demonstrates the comparison of failed requests during the experiments. The results show that *CoScal* can significantly decrease the failed requests. For instance, after 100 minutes experiments, DoScal has 31574 failed requests, HyScal brings 28666 failed requests, KuScal reduces the number to 22756, and our *CoScal* only

TABLE V: Connection time (ms) comparison based on CDF percentile

Percentile	DoScal	KuScal	HyScal	CoScal
20th (ms)	0.758	0.780	0.782	0.752
40th (ms)	0.805	0.812	0.828	0.775
60th (ms)	0.847	0.887	0.936	0.807
80th (ms)	1.208	1.3	1.241	0.896
95th (ms)	2.469	2.129	2.651	1.407
99th (ms)	5.336	4.437	4.614	1.550

TABLE VI: Response time (ms) comparison based on CDF percentile

Percentile	DoScal	KuScal	HyScal	CoScal
20th (ms)	3.26	3.22	3.22	3.38
40th (ms)	4.25	4.23	4.30	4.37
60th (ms)	378.46	353.99	336.84	342.46
80th (ms)	378.46	476.36	464.52	370.56
95th (ms)	567.20	583.55	561.84	412.70
99th (ms)	629.21	641.63	640.09	425.51

has 2345 failed requests. For other observations with longer time, *CoScal* always has the lowest number of failed requests with 92% to 96% reduction compared with other baselines. In conclusion, *CoScal* can ensure more requests are processed successfully.

Transaction Per Second (TPS): it is an important metric to that shows the number of requests per second a system can handle. The higher TPS value represents the better performance of the algorithm. As DoScal has achieved better performance than KuScal and HyScal, we compare DoScal and *CoScal* in Fig. 7. We can observe that the TPS value of CoScal is greater than DoScal. *CoScal* can handle requests stably under heavy loads, and its TPS value fluctuates slightly.

As experimental results of connection time and response time during first 200 minutes are quite close, to better demonstrate the differences between the values of the four scaling approaches, we use the Cumulative Distribution Function (CDF) as the metric which is the integral of the probability density function representing the sum of the probability of occurrence of all values less than or equal to x. The formula of the CDF function is $F_X(x) = P(X \le x)$.

Fig. 8 shows the CDF curves of connection time and response time. There is a clear gap between the *CoScal* and the baselines. Almost all requests can establish the connection within 2 ms by *CoScal*, while other baselines can only have about 90% requests in this range. *CoScal* can respond within 400 ms for most of the requests, while other baselines need more than 500 ms. The increase in response time around 0.43 results from the change of workload level. Therefore, we can conclude that *CoScal* achieves better performance overall compared to the other baselines.

Tables V and VI also demonstrate the connection time and response time at different percentiles. From Table V, we can notice that *CoScal* outperforms other baselines during the observed period, for instance, the connection time of 95% requests falls into 1.407 ms in *CoScal*, while other baselines require 2.129 to 2.651 ms. When comparing the response time in Table VI, we can observe that although *CoScal* is not the



Fig. 6: Performance Comparison of DoScal, KuScal, HyScal, and CoScal: (a) Average Connection Time. (b) Average Response Time. (c) Number of Failed Requests



Fig. 7: Comparison of TPS for DoScal and CoScal



Fig. 8: CDF of connection time and response time curves

best one before 60th percentile, it performs the best at 80th, 90th and 99th percentile. For example, the response time of 80% requests are 370.56 ms in *CoScal*, while other baselines need 464.53 to 476.36 ms.

RL Decisions Analysis: to investigate the essential reasons for resource optimization by applying different techniques in *CoScal*, we summarize the adopted techniques by our RL agent under different workload levels as shown in Table VII, which are collected by the log files that record the adopted operations. The first row represents the original workload level and the first column represents the transferred workload level. The contents in other cells represent the techniques that have been applied, where V stands for vertical scaling, B stands for brownout, H stands for horizontal scaling, and NA means no action is conducted. For example, if the workload level needs

TABLE VII: Multi-faceted scaling techniques adopted under different loads by *CoScal* (the first row represents the current workload level, and the first column represents the next workload level)

Current Level	1	2	3	4
1	NA	Н	Н	NA
2	V	NA	Н	NA
3	V+B	V	NA	NA
4	V+B+H	V+B+H	V+B	NA

to be transferred from level 1 to 4, vertical scaling, horizontal scaling, and brownout should be applied together, since the workloads change significantly. While if the workloads are increased from level 2 to 3, only vertical scaling is required. When the workload level is 4, the response time is high thus the reward function in Equation (4) will not guide horizontal scaling to reduce available resources. It indicates that our RL agent in *CoScal* effectively learns the environment complexity and workload characteristic and accordingly takes the suitable combinations of scaling decisions to increase application QoS and optimize resource usage of physical infrastructure.



Fig. 9: Reward value during learning process

RL Convergence Analysis: to observe the convergence behavior of our RL-based approach, we train the RL model following the same steps of resource scaling according to Section IV-C. Fig. 9 shows that the approach updates its

TABLE VIII: Experimental results with Stan's Robot Shop

Metrics	DoScal	KuScal	HyScal	CoScal
TPS	909	915	882	980
Successful TPS	825	841	815	892
No. of Failures per Second	74	74	68	88
Response Time (ms)	45.1	43.8	61.4	47.2

scaling policy as the training progresses and improves the total reward value. As the RL agent spends initially in the exploration phase, it exponentially increases the reward. After the 16522 epochs, it converges to a stable reward.

D. Experimental Analysis with Stan's Robot Shop

To further evaluate the effectiveness of *CoScal*, we also conduct experiments with another Docker-based microservice application, named Stan's Robot Shop⁴, which is an online system composed of 12 microservices for robots sale. We use Locust⁵, a Python-based testing tool similar to JMeter, to generate requests based on Alibaba's trace [14], and send the requests to the homepage of Stan's Robot Shop. The experiments are conducted in a Docker Swarm cluster with 5 homogeneous machines equipped with Intel Xeon CPU E5-2630 v3 and 64GB RAM.

We compare several metrics in Table VIII collected from Locust, including TPS, successful TPS, response time and the number of failed requests, which have also been evaluated in our previous experiments with Sock Shop application. The final results of each approach are the average value of 300 minutes, and the experiments have been repeated 3 times to avoid randomness. The results show that CoScal can improve 7%-11% TPS compared with the other three baselines, and CoScal can achieve the highest successful TPS. In terms of the number of failures, CoScal is very close to the best result of HyScal. Moreover, CoScal can obtain the lower response time than HyScal, and achieve close results with KuScal. In conclusion, the experiments with Stan's Robot Shop application also validate the effectiveness of our proposed approach in improving system performance. CoScal can achieve good performance as it is adaptive to the change of environments and it searches large solution space.

VII. CONCLUSIONS AND FUTURE WORK

The microservice-based applications have been widely adopted in cloud computing environments, converting monolithic applications into lightweight, flexible, and looselycoupled application components. Such a modular design of the application helps achieve CI/CD (continuous integration, continuous delivery, and continuous deployment) of the application life cycle in cloud environments. However, efficient algorithms for resource scaling in cloud infrastructure are required to ensure the sustainable development of microservice technology. In this paper, we proposed a multi-faceted scaling approach, named *CoScal*, for scaling resources for microservices that combine the three techniques, including horizontal scaling, vertical scaling, and brownout. It ensures cloud service providers optimize their resource usage while guaranteeing QoS. CoScal utilizes deep learning approaches for workload prediction to predict load more accurately compared to traditional regression-based prediction approaches. Furthermore, leveraging the RL framework takes decisions on scaling strategies, adapting to the load changes, and allocating the appropriate resources. To evaluate the performance of CoScal, we deployed our algorithms on a prototype system with representative microservices application. The results are compared with popular state-of-the-art algorithms from research and industry domains with two microservice applications. The experimental results have demonstrated that CoScal can outperform the baseline algorithms in connection time, response time and number of failed requests for Sock Shop application, and improve the number of successful transactions with 6%-10% for Stan's Robot Shop application with slight degradation in response time.

As for future work, we would like to investigate network resource scaling to extend to the suitable scenarios of our scaling approach. As the size of Q-table can increase significantly when the state-action space increases largely, we would also like to investigate Deep Q-Learning with workloads prediction to improve the speed and accuracy of inference, which can also address the limitation in Q-Learning based approach that heavily relies on Q-table. Moreover, we plan to migrate our system to a Kubernetes-based platform to evaluate mediumscale microservice applications, such as TrainTicket and Hotel Reservation applications.

ACKNOWLEDGMENTS

This work is supported by National Key R&D Program of China (No. 2021YFB3300200), National Natural Science Foundation of China (No. 62102408, 62072451), Shenzhen Science and Technology Program (Grant No. RCBS20210609104609044), Youth Innovation Promotion Association CAS (2019349), and CCF-Huawei Innovative Research Plan. The authors would like to thank Dr. Zhicheng Cai for his suggestions to improve this work.

SOFTWARE AVAILABILITY

The source codes of *CoScal* project can be found at the GitHub repository: https://github.com/Kminassch/CoScal

REFERENCES

- M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica *et al.*, "A view of cloud computing," *Communications of the ACM*, vol. 53, no. 4, pp. 50–58, 2010.
- [2] M. Kumar, S. Sharma, A. Goel, and S. Singh, "A comprehensive survey for scheduling techniques in cloud computing," *Journal of Network and Computer Applications*, vol. 143, pp. 1–33, 2019. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1084804519302036
- [3] S. Newman, Building Microservices. "O'Reilly Media, Inc.", 2015.
- [4] F. Al-Doghman, N. Moustafa, I. Khalil, Z. Tari, and A. Zomaya, "Aienabled secure microservices in edge computing: Opportunities and challenges," *IEEE Transactions on Services Computing*, pp. 1–1, 2022.
- [5] M. Kumar, A. Kishor, J. Abawajy, P. Agarwal, A. Singh, and A. Zomaya, "Arps: An autonomic resource provisioning and scheduling framework for cloud platforms," *IEEE Transactions on Sustainable Computing*, pp. 1–1, 2021.

⁴https://www.instana.com/blog/stans-robot-shop-sample-microserviceapplication/

⁵https://locust.io/

- [6] J. Son and R. Buyya, "Sdcon: Integrated control platform for softwaredefined clouds," *IEEE Transactions on Parallel and Distributed Systems*, vol. 30, no. 1, pp. 230–244, 2019.
- [7] M. S. Hasan, F. Alvares, T. Ledoux, and J.-L. Pazat, "Investigating energy consumption and performance trade-off for interactive cloud application," *IEEE Transactions on Sustainable Computing*, vol. 2, no. 2, pp. 113–126, 2017.
- [8] A. Kwan, J. Wong, H.-A. Jacobsen, and V. Muthusamy, "Hyscale: Hybrid and network scaling of dockerized microservices in cloud data centres," in 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS). IEEE, 2019, pp. 80–90.
- [9] L. Toka, G. Dobreff, B. Fodor, and B. Sonkoly, "Machine learning-based scaling management for kubernetes edge clusters," *IEEE Transactions* on Network and Service Management, vol. 18, no. 1, pp. 958–972, 2021.
- [10] H. Qiu, S. S. Banerjee, S. Jha, Z. T. Kalbarczyk, and R. K. Iyer, "FIRM: An intelligent fine-grained resource management framework for slo-oriented microservices," in *14th USENIX Symposium* on Operating Systems Design and Implementation (OSDI 20). USENIX Association, Nov. 2020, pp. 805–825. [Online]. Available: https://www.usenix.org/conference/osdi20/presentation/qiu
- [11] L. Tomás, C. Klein, J. Tordsson, and F. Hernández-Rodríguez, "The straw that broke the camel's back: safe cloud overbooking with application brownout," in *Proceedings of the 2014 IEEE International Conference on Cloud and Autonomic Computing*, 2014, pp. 151–160.
- [12] M. Xu, A. N. Toosi, and R. Buyya, "A self-adaptive approach for managing applications and harnessing renewable energy for sustainable cloud computing," *IEEE Transactions on Sustainable Computing*, vol. 6, no. 4, pp. 544–558, 2021.
- [13] Z. Zhong, M. Xu, M. A. Rodriguez, C. Xu, and R. Buyya, "Machine learning-based orchestration of containers: A taxonomy and future directions," ACM Comput. Surv., jan 2022, just Accepted. [Online]. Available: https://doi.org/10.1145/3510415
- [14] W. Chen, K. Ye, Y. Wang, G. Xu, and C. Xu, "How does the workload look like in production cloud? analysis and clustering of workloads on alibaba cluster trace," in 2018 IEEE 24th International Conference on Parallel and Distributed Systems (ICPADS), 2018, pp. 102–109.
- [15] L. Liu, "Qos-aware machine learning-based multiple resources scheduling for microservices in cloud environment," arXiv preprint arXiv:1911.13208, 2019.
- [16] R. S. Kannan, L. Subramanian, A. Raju, J. Ahn, J. Mars, and L. Tang, "Grandslam: Guaranteeing slas for jobs in microservices execution frameworks," in *Proceedings of the Fourteenth EuroSys Conference 2019*, ser. EuroSys '19. New York, NY, USA: Association for Computing Machinery, 2019. [Online]. Available: https://doi.org/10.1145/3302424.3303958
- [17] G. Yu, P. Chen, and Z. Zheng, "Microscaler: Automatic scaling for microservices with an online learning approach," in 2019 IEEE International Conference on Web Services (ICWS), July 2019, pp. 68–75.
- [18] H. Zhou, M. Chen, Q. Lin, Y. Wang, X. She, S. Liu, R. Gu, B. C. Ooi, and J. Yang, "Overload control for scaling wechat microservices," in *Proceedings of the ACM Symposium on Cloud Computing*, ser. SoCC '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 149–161. [Online]. Available: https://doi.org/10.1145/3267809.3267823
- [19] K. Rzadca, P. Findeisen, J. Swiderski, P. Zych, P. Broniek, J. Kusmierek, P. Nowak, B. Strack, P. Witusowski, S. Hand, and J. Wilkes, "Autopilot: Workload autoscaling at google," in *Proceedings of the Fifteenth European Conference on Computer Systems*, ser. EuroSys '20. New York, NY, USA: Association for Computing Machinery, 2020. [Online]. Available: https://doi.org/10.1145/3342195.3387524
- [20] X. Hou, C. Li, J. Liu, L. Zhang, Y. Hu, and M. Guo, "Ant-man: towards agile power management in the microservice era," in 2020 SC20: International Conference for High Performance Computing, Networking, Storage and Analysis (SC). IEEE Computer Society, 2020, pp. 1098– 1111.
- [21] S. Zhang, T. Wu, M. Pan, C. Zhang, and Y. Yu, "A-sarsa: A predictive container auto-scaling algorithm based on reinforcement learning," in 2020 IEEE International Conference on Web Services (ICWS), 2020, pp. 489–497.
- [22] F. Rossi, M. Nardelli, and V. Cardellini, "Horizontal and vertical scaling of container-based applications using reinforcement learning," in 2019 IEEE 12th International Conference on Cloud Computing (CLOUD), 2019, pp. 329–338.
- [23] A. U. Gias, G. Casale, and M. Woodside, "Atom: Model-driven autoscaling for microservices," in 2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS), 2019, pp. 1994–2004.

- [24] R. Nishtala, P. Carpenter, V. Petrucci, and X. Martorell, "Hipster: Hybrid task manager for latency-critical cloud workloads," in 2017 IEEE International Symposium on High Performance Computer Architecture (HPCA). IEEE, 2017, pp. 409–420.
- [25] R. Nishtala, V. Petrucci, P. Carpenter, and M. Sjalander, "Twig: Multiagent task management for colocated latency-critical cloud services," in 2020 IEEE International Symposium on High Performance Computer Architecture (HPCA). IEEE, 2020, pp. 167–179.
- [26] M. A. Chang, A. Panda, Y.-C. Tsai, H. Wang, and S. Shenker, "Throttlebot-performance without insight," arXiv preprint arXiv:1711.00618, 2017.
- [27] H. Khazaei, C. Barna, and M. Litoiu, "Performance modeling of microservice platforms considering the dynamics of the underlying cloud infrastructure," arXiv preprint arXiv:1902.03387, 2019.
- [28] Y. Gan, Y. Zhang, K. Hu, D. Cheng, Y. He, M. Pancholi, and C. Delimitrou, "Seer: Leveraging big data to navigate the complexity of performance debugging in cloud microservices," in *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 19–33. [Online]. Available: https://doi.org/10.1145/3297858.3304004
- [29] L. Suresh, P. Bodik, I. Menache, M. Canini, and F. Ciucu, "Distributed resource management across process boundaries," in *Proceedings of the 2017 Symposium on Cloud Computing*, ser. SoCC '17. New York, NY, USA: Association for Computing Machinery, 2017, p. 611–623. [Online]. Available: https://doi.org/10.1145/3127479.3132020
- [30] X. He, Z. Tu, X. Xu, and Z. Wang, "Re-deploying microservices in edge and cloud environment for the optimization of user-perceived service quality," in *International Conference on Service-Oriented Computing*. Springer, 2019, pp. 555–560.
- [31] B. Burns, J. Beda, and K. Hightower, Kubernetes: up and running: dive into the future of infrastructure. O'Reilly Media, 2019.
- [32] T. Kiss, P. Kacsuk, J. Kovács, B. Rakoczi, Á. Hajnal, A. Farkas, G. Gesmier, and G. Terstyanszky, "Micado—microservice-based cloud application-level dynamic orchestrator," *Future Generation Computer Systems*, vol. 94, pp. 937–946, 2019.
- [33] X. Zhou, X. Peng, T. Xie, J. Sun, C. Xu, C. Ji, and W. Zhao, "Poster: Benchmarking microservice systems for software engineering research," in 2018 IEEE/ACM 40th International Conference on Software Engineering: Companion (ICSE-Companion). IEEE, 2018, pp. 323–324.
- [34] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng, "Tensorflow: A system for large-scale machine learning," in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*. Savannah, GA: USENIX Association, Nov. 2016, pp. 265–283.
- [35] "Sock shop: a microservices demo application. available:https://microservices-demo.github.io/," 2017.
- [36] S. Kardani-Moghaddam, R. Buyya, and K. Ramamohanarao, "Adrl: A hybrid anomaly-aware deep reinforcement learning-based resource scaling in clouds," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 3, pp. 514–526, 2021.
- [37] A. Mendoza, "Using muon to monitor sas applications on aix servers," in SAS Global Forum. Citeseer, 2008, pp. 16–19.
- [38] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio, "Empirical evaluation of gated recurrent neural networks on sequence modeling," in NIPS 2014 Workshop on Deep Learning, December 2014, 2014.
- [39] T. Mikolov, S. Kombrink, L. Burget, J. Černocký, and S. Khudanpur, "Extensions of recurrent neural network language model," in 2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2011, pp. 5528–5531.
- [40] Z. Chen, J. Hu, G. Min, A. Y. Zomaya, and T. El-Ghazawi, "Towards accurate prediction for high-dimensional and highly-variable cloud workloads with deep learning," *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 4, pp. 923–934, April 2020.
- [41] S. Wang, Y. Guo, N. Zhang, P. Yang, A. Zhou, and X. Shen, "Delayaware microservice coordination in mobile edge computing: A reinforcement learning approach," *IEEE Transactions on Mobile Computing*, vol. 20, no. 3, pp. 939–951, 2021.
- [42] G. Yin, C. Z. Xu, and L. Y. Wang, "Q-learning algorithms with random truncation bounds and applications to effective parallel computing," *Journal of optimization theory and applications*, 2008.
- [43] S. Levine, A. Kumar, G. Tucker, and J. Fu, "Offline reinforcement learning: Tutorial, review, and perspectives on open problems," arXiv preprint arXiv:2005.01643, 2020.



Minxian Xu is currently an Associate Professor at Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences. He received the BSc degree in 2012 and the MSc degree in 2015, both in software engineering from University of Electronic Science and Technology of China. He obtained his PhD degree from the University of Melbourne in 2019. His research interests include resource scheduling and optimization in cloud computing. He has co-authored 40+ peer-reviewed papers published in prominent international journals and conferences,

such as ACM CSUR, ACM TOIT, IEEE TSUSC, IEEE TCC, IEEE TASE, IEEE TGCN, JPDC, JSS and ICSOC. His Ph.D. Thesis was awarded the 2019 IEEE TCSC Outstanding Ph.D. Dissertation Award. More information can be found at: minxianxu.info.



Juanjuan Zhao received her Ph.D degree from Shenzhen College of Advanced Technology, University of Chinese Academy of Sciences in 2017, and received the M.S. degree from the Department of Computer Science, Wuhan University of Technology in 2009. She is an Associate Professor at Shenzhen Institutes of Advanced Technology, Chinese Academy of Sciences. Her research topics include data-driven urban systems, mobile data collection, cross-domain data fusion, heterogeneous model integration.



Chenghao Song received his BSc degree from University of Electronic Science and Technology of China. Now he is a master student at University of Melbourne, he is also a visiting student at Shenzhen Institutes of Advanced Technology, Chinese Academy of Science. His research interest includes deep learning for cloud resource optimization.



Kejiang Ye received his BSc and PhD degree in Computer Science from Zhejiang University in 2008 and 2013, respectively. He was also a joint PhD student at The University of Sydney from 2012 to 2013. After graduation, he works as Post-Doc Researcher at Carnegie Mellon University from 2014 to 2015 and Wayne State University from 2015 to 2016. He is currently a Professor at Shenzhen Institutes of Advanced Technology, Chinese Academy of Science. His research interests focus on the performance, energy, and reliability of cloud computing

and network systems.



Shashikant Ilager is a postdoctoral researcher at Vienna University of Technology (TU Wien), Austria. He completed his PhD at the University of Melbourne in 2021. His research interests covers the boundaries of large scale Distributed Systems and Machine Learning. He published several papers in leading journals and conferences including IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Mobile Computing, and IEEE/ACM International Symposium on Cluster, Cloud and Internet Computing (CCGRID). He is a

recipient of Best Paper award from CCGRID 2020 conference.



Sukhpal Singh Gill is a Lecturer (Assistant Professor) in Cloud Computing at the School of Electronic Engineering and Computer Science, Queen Mary University of London, UK. Prior to his present stint, Dr. Gill has held positions as a Research Associate at the School of Computing and Communications, Lancaster University, UK and also as a Postdoctoral Research Fellow at CLOUDS Laboratory, The University of Melbourne, Australia. Dr. Gill is serving as an Associate Editor in Wiley ETT and IET Networks Journal. His research interests include Cloud

Computing, Fog Computing, Internet of Things and Energy Efficiency. For further information, please visit: http://www.ssgill.me.



Chengzhong Xu (Fellow, IEEE) is the Dean of Faculty of Science and Technology and the Interim Director of Institute of Collaborative Innovation, University of Macau, and a Chair Professor of Computer and Information Science. Dr. Xu's main research interests lie in parallel and distributed computing and cloud computing, in particular, with an emphasis on resource management for system's performance, reliability, availability, power efficiency, and security, and in big data and data-driven intelligence applications in smart city and self-driving vehicles.

He published two research monographs and more than 300 peer-reviewed papers in journals and conference proceedings; his papers received about 10K citations with an H-index of 59. He serves or served on a number of journal editorial boards, including IEEE Transactions on Computers (TC), IEEE Transactions on Cloud Computing (TCC), IEEE Transactions on Parallel and Distributed Systems (TPDS) and Journal of Parallel and Distributed Computing (JPDC). Dr. Xu has been the Chair of IEEE Technical Committee on Distributed Processing (TCDP) since 2015. He obtained BSc and MSC degrees from Nanjing University in 1986 and 1989 respectively, and a PhD degree from the University of Hong Kong in 1993, all in Computer Science and Engineering.