

Mobile-Kube: Mobility-aware and Energy-efficient Service Orchestration on Kubernetes Edge Servers

Saeid Ghafouri, Alireza Karami, Danial Bidekani Bakhtiarvand, Aliakbar Saleh Bigdeli, Sukhpal Singh Gill and Joseph Doyle

Abstract—In recent years Kubernetes has become the de facto standard in the realm of service orchestration. Despite its great benefits, there are still numerous challenges to make it compatible with decentralised cloud computing platforms. One of the challenges of mobile edge computing is that the location of the users is changing over time. This mobility will constantly alter the proximity of the users to their connected services. One solution to this problem is to regularly move services to computing nodes near the users. However, distributing the services in edge nodes only subject to user movements will result in the fragmentation of active nodes. This leads to having active nodes that do not use their full capacity. We have proposed a method called Mobile-Kube to reduce the latency of Kubernetes applications on mobile edge computing devices while maintaining energy consumption at a reasonable level. An experimental framework is designed on top of real-world Kubernetes clusters and real-world traces of mobile users' movements have been used to simulate the users' mobility. Experimental results show that Mobile-Kube can achieve similar energy consumption performance to a heuristic approach that focuses on reducing energy consumption only while reducing the latency of services by 43%.

Index Terms—Resource Management, Energy-Efficiency, Cloud Computing, Edge Computing, Reinforcement Learning



1 INTRODUCTION

NOWADAYS containerised services are ubiquitous. They provide a modular way of packing a single or sets of functionalities in separate isolated environments. The containerisation of services provides agile development of software with DevOps [1] practices by easing testing, modularisation and fully automating integration and delivery process. From the cloud provider point of view, containerisation provides lightweight and scalable deployment of cloud services. This interest in using containerised software has led to the implementation of production level container management and orchestration systems. Kubernetes has become the most popular orchestration system as it provides automated solutions to tasks that previously required numerous technologies working together with lots of human provisioning [2].

One of the main challenges with centralised cloud computing clusters is the distance of the infrastructure from the end users [3]. Time-sensitive, real time and location-aware services are some examples of the services that are hard to deploy on distant cloud servers. To address this challenge, the fog and edge computing paradigms [4] have been proposed to distribute the computation so that it is closer to the mobile devices. Mobile edge servers are typically smaller nodes deployed near base stations. Typically users connected to the services on edge nodes can access them directly rather than through the network core.

Services deployed on the edge are highly available and usually do not suffer from the communication overhead of centralised clouds [3]. Edge clusters are highly distributed and have limited resources compared to the centralised clouds. Therefore, their resource management is more difficult and requires more automated provisioning methods. Migrating services and virtual machines (VMs) subject to users' mobility is one of the objectives that has been the topic of some previous literature [5]–[10], [10]–[15]. This problem is studied in a subset of edge computing called **Mobile Edge Computing (MEC)** that deals with making edge computing paradigms more accessible to mobile device users. The mobile users are typically connected to the core network of fog devices through base stations close to them. The edge server is placed near the base station and can expose the service to users with lower latency. However, mobile users typically are not static users. Thus, the base station to which they have connected changes over time based on their mobility. Previous works in this area have mainly ignored the challenges of deploying their methods on real-world orchestration frameworks like Kubernetes.

The primary focus of architectures such as mobile edge computing [16], cloudlets [17] and fog computing [18] has been minimizing the latency or maximizing the throughput of services to mobile users without considering the energy consumption of the edge cloud and the associated environmental impact. Recently, however, cloud services with large providers such as Microsoft and Amazon have increased their focus on reducing their environmental impact by committing to carbon neutrality by 2030 and 2040 respectively [19], [20]. It is likely that the increasing pressure from government organizations will result in edge clouds also considering their operating energy consumption [21].

- S. Ghafouri, S.S. Gill and J. Doyle are with the School of Electronic Engineering and Computer Science, Queen Mary University of London, UK.
E-mail: {s.ghafouri, s.s.gill, j.doyle}@qmul.ac.uk.
- A. Karami, D.B. Bakhtiarvand and A.S. Bigdeli are independent researchers.

There has been considerable research on minimizing the energy consumption of centralized cloud services including work on consolidation [22], geographical load balancing [23], the management of workloads that do not have strict deadlines [24] and the choice of cloud architecture that should be used to support services [25]. Thus, we propose the Mobile-Kube system that considers both overall energy consumption and the latency of service users to provide a MEC platform with good performance for its services and a reduced environmental impact.

The *main contributions* of this research work are:

- We present the Mobile-Kube system that integrates with Kubernetes to reduce the latency of service users on edge clusters based on users' mobility while considering energy consumption. Mobile-Kube models the problem of service placement while maintaining a reasonable energy consumption as an optimization problem and solves it using a reinforcement learning algorithm.
- We compare several reinforcement learning methods to determine the best reinforcement learning algorithm to achieve a reasonable trade-off between the two aforementioned objectives. Our experiments show that IMPALA is the best Reinforcement Learning(RL) method for this scenario based on its data efficiency and fast convergence.
- We evaluate Mobile-Kube on the Google Cloud Platform using a mobile user movement emulator that is based on real-world traces. Our results show that Mobile-Kube can achieve similar energy consumption performance to a heuristic approach that focuses on reducing energy consumption only while reducing the latency of services by 43%.

2 MOTIVATING SCENARIO

In MEC networks, users are typically connected to a base station that connects them to one of the edge servers. In most scenarios users are typically connected to their closest base station. As the users of mobile edge cloud move around they might transfer from the one base station and connect to another base station that is closer as it provides a better signal. The problem is that the service that they are connected to may be closer to the former base station and this can result in increased latency. One solution to this is to move the service from their previous location to a place closer to a new base station that the users of that service are connected to. A greedy algorithm that just moves the service immediately to an empty server closest to the base station can be utilised. However, the problem with that approach is that the new placement might result in switching on a new node which will increase energy consumption. In this work, we try to learn a service placement algorithm that can achieve a balance between keeping the number of active nodes to a minimum while providing a reasonable Quality of Service (QoS) to the mobile users through moving services to nodes closer to the users. This problem is also similar to the problem of node autoscaling [26]. In node autoscaling in a Kubernetes cluster the number of nodes in clusters is up-scaled (nodes added) and down-scaled (nodes removed) based on the incoming workloads. We have avoided using

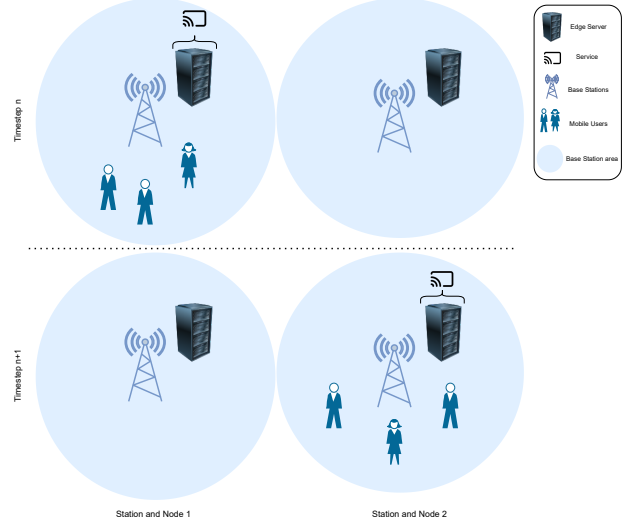


Fig. 1. In timestep n users are around the vicinity of the node and base station 1 and connected to their service residing in it. In timestep $n+1$ the users of the service move to the area covered by node and base station 2 therefore we move the service to node 2 to make it closer to its connected users.




Fig. 2. In timestep n there are three active nodes in the cluster. Moving services 7 and 8 to nodes 1 and 2 will leave node 3 empty and now we can turn off node 3 resulting in a reduction of energy consumption and costs.

the term autoscaling here because the migration of services is currently not something native to Kubernetes and node autoscaling happens based on incoming workloads rather than the service migration. You can see an example of moving services to a closer vicinity of the connected users in Figure 1. This is the first objective of our optimisation. In Figure 2 you can see an example of the binpacking objective that aims to have the minimum number of active distributed Kubernetes nodes in the Kubernetes cluster.

3 RELATED WORKS

Mobility Aware Service Orchestration at Edge With the advent of 5G and Long-Term Evolution (LTE) networks over the past few years, the need for mobility driven service placement methods has become more evident. Ouyang *et al.* [6] have proposed a method for migrating the services to nearby servers to the users while maintaining a reasonable migration cost using Lyapunov optimisation. While it is not evaluated experimentally they still have provided a strong theoretical analysis to achieve an optimised service placement subject to latency and a long term cost budget. Badri *et al.* [27] examined the trade-off between energy consumption and reducing the latency for QoS with service migration through a stochastic optimisation approach. The method is able to handle non-deterministic users movements. The cost of relocating a service is also considered in their simulations. Due to the Markovian nature of the service placement problem reinforcement learning and bandit based methods have also been used in recent years. Wang *et al.* [28] models the problem as a Markov decision process under different scenarios of 1-D and 2-D mobility. Real-world datasets have been used by Wang *et al.* in [11] to show the effectiveness of an offline reinforcement learning based approach to reduce the overall service delay. The most similar papers to our work are [8], [14] which have used different variations of RL for placement of cloud services subject to the user mobility while achieving a trade-off in energy consumption. Tang *et al.* [8] have used a variation of reinforcement learning namely Q-learning to model user movements as a Markov decision process and have considered migration costs into account. Despite their precise modeling of the cost and delay in edge node, their implementation is not yet on the real-world production-ready orchestration frameworks like Kubernetes. They have used Checkpoint/Restore In Userspace (CRIU) for implementation of the real-world containers migrations for comparison of migration time between VMs and containers. The resource consumption and delay are modeled and the RL algorithm can maintain a balance between the two objectives of reducing the delay and power consumption. However, their experiments have only been tested in a trace-driven simulation environment and no integration with real-world service orchestration systems is presented. To solve the large state space of the MEC placement problem Brandherm *et al.* and Liu *et al.* [9], [13] have used Multi-Agent version of the RL algorithms in distributed settings. However, they have used different Multi-agent algorithms. the former has used q-learning and the latter uses actor-critic networks.

Kubernetes Limitations Ignored in previous works Other than the lack of the real-world implementation of the container migration, there is another problem in the theoretical assumptions of the previous works when it comes to deploying them in real-world systems. Kubernetes resource models work through the request and limit model. The request is the amount of the resource reserved for a Kubernetes container in each node and the limit is the maximum resources that could be used from the cluster. The problem with some of the mentioned previous works [12], [27] is that the requested computational resource is reserved from the user side and a new container is started for each

users task. However, most of the time in reality a container could be serving many users at the same time. Therefore, the number of containers in the nodes for each service is determined by the cloud provider, not the cloud user. Our approach considers services with multiple connected users to address this issue.

Reinforcement Learning for Resource Management

Problems that can be expressed as a sequential decision making process can be modelled as a Markov Decision Processes (MDPs). MDPs are the core mathematical formalization that is used in most of the sequential decision ¹ making problems. An MDP is a model of sequential decisions that are an abstraction of an agents' behavior in a fully or partially observed environment. An agent is an entity that makes the actions. For example, in an Atari ² game the player who moves the joystick is the agent. At each step of decision-making, the agent receives a reward from the environment after the action has been taken. This reward indicates the value of the taken action. If a learning method is associated with the MDP then this reward is a measure that is used to learn to take better actions over the subsequent steps. The main property of the MDPs is that the action taken at each timestamp depends only on the current state of the environment and not any state(s) before that. Mathematically speaking, an MDP consists of a trajectory of states S , actions A and rewards R in the following order:

$$S_0, A_0, R_1, S_1, A_1, R_2, S_2, A_2, R_3, \dots \quad (1)$$

Dynamic Resource management can be considered as an example of sequential decision making since in almost all of its forms it involves deciding which task should be allocated to which resource. Therefore, it can be expressed as an MDP. This gives us robust modeling of the resource management problem that can be used alongside methods for solving the MDPs to allocate resources effectively.

After the success of the DeepRM [30] there is a line of research pursuing the problem of resource allocation with RL to optimize classic system problems. The successes of reinforcement learning in playing games and solving control problems [29], [31] have motivated many researchers to take it to their domain of interest. Some examples are device placement optimization [32], [33], video streaming bitrate adaptation [34], and Internet congestion control [35]. As noted before, the resource allocation problem can be expressed as an MDP. This gives rise to the idea of using reinforcement learning in this context which is naturally a learning model designed to solve MDPs. Other than the modeling capabilities of the RL it is adaptable toward the incoming workload. The policy can change over time according to the changes in the environment [30]. One of the drawbacks of RL methods is that they require a large amount of training data. Thus, the abundance of training historical data in resource allocation logs makes them a good fit for the problem [36].

TABLE 1
Nomenclature and Notational Conventions.

Key Concept	Definition
E	Edge Node Set
m	Number of Nodes
E_i	i^{th} node ($i \in [1, m]$)
E_i^l	Location of node i
E_i^C	CPU capacity of node i
E_i^M	Memory capacity of node i
p_i	Power consumed by node i
p_i^{idle}	Idle power of node i
p_i^{max}	Max power of node i
C	Container Set
n	Number of Containers
C_j	j^{th} container ($j \in [1, n]$)
t	Time interval
T	Total time interval in each training episode
$C_j^l(t)$	Location of container j at time t ($(c_j^l(t) \in \{F\})$)
$C_j^r(t)$	CPU requirement of container j at time t
$C_j^m(t)$	Memory requirement of container j at time t
$C_j^a(t)$	List of mobile applications on container j at time t
M	Mobile User set
o	Number of mobile users
M_k	k^{th} mobile user ($k \in [1, o]$)
M_k^l	Location of mobile user k at time t
M_k^r	Container request of mobile user k at time t
d_{M_k}	Distance of user M_k from its connected service
d_{total}	Total system delay
d_{net}	Network delay
d_q	Delay due to lack of available resources
p_{total}^r	Total Power consumption
u_i^t	Resource request of node i at time t
C_t	Total Cost

4 PROBLEM FORMULATION

Let $E = \{E_1, E_2, \dots, E_m\}$ be the set of edge nodes, $C = \{C_1, C_2, \dots, C_n\}$ be the set of containers which host mobile applications and $M = \{M_1, M_2, \dots, M_o\}$ be the set of mobile users which connect to these applications. Each edge node i has a location E_i^l , a CPU capacity E_i^C and a memory capacity E_i^M . Each container j has a location at time t which is denoted as $C_j^l \in \{F\}$. The resource requirements and allocation can change with time. Thus, let $C_j^r(t)$ and $C_j^m(t)$ represent the containers' CPU resource requirements, the containers' memory resource requirements respectively. We also need to consider the location of the mobile user which can change over time and that they may use different mobile applications at different times. Thus, let

1. In the context of reinforcement learning these decisions are referred to as actions.

2. After the DeepMind seminal paper [29] that used Atari games as a benchmark it has become a classic example and benchmark in the RL community research.

$M_k^l(t)$ represent the location of the mobile user k and $M_k^r(t)$ represent the application that the mobile user k is accessing.

The overall goal of the system is to minimise the energy consumption of edge cloud and maximise the performance of the mobile applications. Maximising the performance of the mobile applications is achieved by minimizing the total delay of the system and the migration of containers. It goal can be formulated as:

$$\text{minimise } C_t = w_1 d_{total} + w_2 p_{total} \quad (2)$$

$$\text{subject to } \sum_{C_j^l \in E_i} C_j^r \leq E_i^C \quad \forall i \quad (3)$$

$$\sum_{C_j^l \in E_i} C_j^m \leq E_i^M \quad \forall i \quad (4)$$

Where w_1 w_2 are relative pricing functions used to indicate the relative importance of each sub goal. They are frequently used in systems problems to balance the importance of factors in a trade off [37]. The conditions of the optimisation problem are used to represent the resource capacities of the edge servers.

The total delay is comprised of delays due to the network d_{net} and delays due to lack of available resources for computation d_q . It can be formulated as:

$$d_{total} = d_{net} + d_q \quad (5)$$

d_{net} will vary depending on the location of the containers and the mobile users.

d_{net} is computed by averaging each individual users' experienced latency:

$$d_{net} = \sum_{k=1}^o d_{M_k} / o \quad (6)$$

d_q will always be zero as the optimisation conditions prevent overloading an edge server. Further work will explore relaxing these constraints to achieve better service consolidation and how this affects service performance.

The power p_i consumed by node i is based upon the resource request of the at node $u_i(t)$. The power consumed by node i can be calculated as:

$$p_i(t) = \begin{cases} (p_i^{idle} + (p_i^{max} - p_i^{idle})) \times u_i(t), & \text{if } u_i(t) > 0 \\ 0, & \text{if } u_i(t) = 0 \end{cases} \quad (7)$$

If the utilization is zero then then the system switches the node off and it consumes no power which is why $p_i(t) = 0$ if $u_i(t) = 0$. The total power p_{total} can then be calculated as:

$$p_{total} = \sum_{i=1}^m p_i(t) \quad (8)$$

5 PROPOSED REINFORCEMENT LEARNING SOLUTION

As the user mobility and the movement of services in our problem are a sequential decision making problem preserving the Markov [8] property, deep reinforcement learning solutions are a good approach for them. In deep reinforcement learning an agent tries to learn a policy $\pi(s)$ that maximise the discounted reward received from the environment. To achieve this, it first tries to find the

value of the states which is the sum of the observed rewards from the starting state until the terminal state $v_\pi(s) = \mathbb{E}_\pi[R_{t+1} + \gamma R_{t+2} + \gamma^2 R_{t+3} + \dots | S_t = s]$. Saving all the state values in a table is not feasible in problems with large state-space. In deep reinforcement learning a neural network is used to save all the values of states and at each timestamp of the reinforcement learning algorithm, the weights of the neural network are updated based on the received reward from the environment. In value based reinforcement learning the neural network receives the state of the environment as the input and returns the state-value as the output. A maximisation or a greedy policy step is done afterward to choose the appropriate action from the output state values received from the neural network. However, in policy gradient methods the neural network returns the action directly to the output. [38] The general scheme of the used RL algorithms is shown in Algorithm 1. The $\pi_{\theta_{old/new}}$ on line 6 depends on whether the old or new policy is used based on whether the method is an on policy or off policy RL algorithm. In our case, PG and PPO are on policy and IMPALA is off policy.

Algorithm 1 Migration Algorithm

```

1: for  $iterations = 1, 2, \dots, N$  do
2:   for  $episodes = 1, 2, \dots, M$  do
3:      $R_b \leftarrow 0$ 
4:      $R_l \leftarrow 0$ 
5:     for  $timesteps = 1, 2, \dots, T$  do
6:       Run a trajectory  $\pi_{\theta_{old/new}}$  for  $T$  timesteps
7:       Compute the timestep latency reward  $r_l$ 
8:       Compute the timestep binpacking reward  $r_b$ 
9:     end for
10:     $R_b \leftarrow R_b + r_b$ 
11:     $R_l \leftarrow R_l + r_l$ 
12:  end for
13:   $R \leftarrow w_1 R_b + w_2 R_l$ 
14:  Optimize the RL Agent policy  $\pi$  based on  $R$ 
15:  Move the containers based on  $\pi$ 
16: end for

```

Agent We have used three different RL agents in our experiments:

- **PG** Vanilla Policy Gradient (PG) is considered as the basis of all policy gradient RL algorithms. At each iteration of the standard policy gradient method an episode τ (or a batch of episodes) is performed. Each timestep $t \in T$ of an episode is comprised of a state s_t , action a_t and a reward received from the environment for that state-action pair $r(s_t^i, a_t^i)$. The series of all these state, action and reward triplets constitute a full episode trajectory $\tau \sim \{s_1, a_1, r_1, \dots, s_{T-1}, a_{T-1}, r_T\}$. The sum of all the rewards in a sample episode $r(\tau) = \sum_t r(s_t^i, a_t^i)$ is used to update the policy network π sets of parameters θ . To reach this goal first a objective function is computed using the logarithm of the gradients of the policy neural network $\log \pi_\theta(\tau)$ based on the policy gradient theorem [38].

$$\nabla_\theta J(\theta) \approx \frac{1}{N} \sum_{i=1}^N \nabla_\theta \log \pi_\theta(\tau) r(\tau) \quad (9)$$

A gradient ascent step updates the weights of the policy network:

$$\theta \leftarrow \theta + \alpha \nabla_\theta J(\theta) \quad (10)$$

- **PPO** RL algorithms are of high variance by nature. This means that the output actions can change in different training iterations depending on the received random sample. Proximal Policy Optimization (PPO) algorithm solves this issue by constraining the next action within a certain range at each timestep. It does so by clipping the objective function and using a surrogate objective function instead of the equation 10 [39].
- **IMPALA** is a distributed deep RL method [40]. It proposes a distributed training method for one of the former Deep RL methods named A2C. The naive policy gradient method is a high variance method and it is not stable for environments with high fluctuation in the reward signal. This is because the sum of rewards $r(\tau)$ used in equation 9 are only sampled from a single episode. But due to environment dynamics, this single episode return might not be a good indicator of the states' worthiness. One of the techniques used for reducing the variance is to have another neural network called critic which acts as an estimator for the reward function. In each timestep, the parameters ϕ of another neural network is also trained for estimating a value function:

$$V^\pi(s_t) = \sum_{t'=t}^T E_{\pi_\theta} [r(s_{t'}, a_{t'}) | s_t] \quad (11)$$

This value function will be a better estimate since it is an approximation based on a series of episodes rather than a single episode roll out. substituting the $r(\tau)$ with $V^\pi(s_i, a_i)$ the equation 9 will update to:

$$\nabla_\theta J(\theta) \approx \sum_i \nabla_\theta \log \pi_\theta(a_i | s_i) V^\pi(s_i, a_i) \quad (12)$$

The update rule of the algorithm will stay the same as equation 10. At each timestep of the A2C, it distributes the computation of the advantage values into several learners. IMPALA improves the training procedure by distributing the learner into several learners. In former distributed versions of the A2C named A3C, the updating happens through passing gradient to a central learner, however, this can produce a large communication overhead. Instead of the gradients, IMPALA directly sends the trajectory of experiences received from multiple actors to the central learner. The distributed actors continuously update their policy with the latest updated policy in the learner and then send them to the learner. This approach increases the exploration and throughput rate. Each of the learners does the actions based on its version of the policy network not the latest updated version of the policy. To solve this problem, IMPALA uses a correction step using importance sampling called V-trace. Training another policy μ other than the policy π that the agent is using to act in the environment is referred to as off policy methods. To

fix the inconsistency of the actors' policy μ with the learner policy π , in IMPALA the value function in Eq 11 is substituted with another value function with importance sampling to make the learners' policy consistent with the actors' policy.

States Each state is derived by concatenation of two arrays 1. An array of size n , $U = \langle u_1, u_2, \dots, u_n \rangle$ which the indexes are the users' id and each entry value indicates the corresponding user connected station id and is a value from the set of all available stations $u_i \in [s_1, s_2, \dots, s_m]$. 2. Another array $C = \langle c_1, c_2, \dots, c_m \rangle$ which the indexes represent the containers and the items represent the node that the container is placed on, and the values are from $c_i \in [n_1, n_2, \dots, n_k]$. All the discrete values are then encoded using a one-hot encoding before being fed to the RL model. The final observation will be the concatenation of these two parts $O = U \cup C$.

Actions The action map is represented as an array of the size of the containers where each of the indices represents one of the hosts and each of the values at that index represents the id of the host that this container is placed at the next timestamp. This is exactly like the U part of the observations.

Reward The reward at each timestep is calculated per each objective according to the Equation 4. For minimizing the network latency we have $R_l = 1/d_{total}$, and for maximizing the number of empty servers we set the binpacking objective to $R_b = p_{total}$. Both of the values are then normalised according to the network size (see Section 7 for more information). The two values are then summed up according to two weights w_1 and w_2 .

$$R = w_1 R_b + w_2 R_l \quad (13)$$

RL helper Due to the large state space of the problem, the RL agent will face many illegal actions that try to place containers to hosts without enough space. If we want to end the training episode every time we face an illegal action then we stop the agent from learning longer episodes and it makes the training slow. To solve this problem during the training, we assigned a negative reward to the illegal actions but we continued the simulator traces to the next timestep. In the test phase on the real-world Kubernetes servers if an illegal action is received, we skip that action and keep the servers at their place until the next timesteps' action.

Policy Networks We have used a similar structure for the policy neural network across all the reinforcement learning agents. We have used a two-layer fully-connected network with 64 neurons at each layer. We chose this simple architecture as we did not see a meaningful difference in using more complicated architectures when we experimentally evaluated them.

6 SYSTEM DESIGN

We have designed a complete end to end experimental setup. The implementation of Kubernetes side is all done in real-world Google Cloud cluster. However, the user mobility side is simulation based on real-world data. This is due to the difficulty of having real-world experiments with real moving mobile users.

Kubernetes Internal Structure Kubernetes clusters are deployed with a number of nodes. *Nodes* can be considered as the equivalent of servers in other forms of computing clusters. A node could be a bare metal server or virtual machine of any kind. Once a computing resource is defined as a Kuberentes node, all the nodes of different types will look the same from the users' perspective. *Containers* in Kubernetes are the minimum isolation level of the applications, but the smallest Kuberentes abstraction for representing containers is another entity called a *pod*. Pods are the smallest deployable unit that presents a set of containers that share the same networking interface. All the containers inside a pod are always co-located in the same cluster node [41]. Pods are ephemeral objects and are replaced and rescheduled constantly during their lifetime. Kubernetes *services* are a way of building a consistent representation of the Kubernetes pods networking for accessing them. This is different from the concept of the service we have used in this work, it is just the Kubernetes internal networking tool that we have used to expose our applications which call a unit of them a service. The state of the cluster and all the internal Kubernetes communications are done through the Kubernetes API server. The API server exposes the Kubernetes API using a rest API. In order to be able to interact with the API server, there are many options like the Kuberntes CLI named kubectl or other client APIs available in multiple language. We have used the python client API for interacting with the API server.

Our definition of Services Usually, real world cloud services are made from a sets of containers and stateful and stateless microservices. For example, a streaming service consists of a database system, authentication system, video analytic service, and many other small decoupled modular objects. For experimental purposes we limit the definition of a service to a single pod containing one single container inside it with a Flask Python app [42] for generating load on the container CPU and RAM. The Flask Python app is exposed to the outside world using Kubernetes' service.

Load Generation Module To emulate artificial load on the services, we have used a model with two containerised Flask applications. One of them called the utilisation server is deployed as a single Flask application in one arbitrary server outside the cluster and the other one is an application running on all the services. At each timestamp, the utilisation server sends the resource usage of each application to them and the applications on the services put the load on the CPU and RAM of the application using a Linux tool named stress-ng.

Changes to the Kubernetes default scheduler Kubernetes schedules pods at the beginning of their lifetime. The scheduling of pods is a two step process: 1. Filtering and 2. Scoring. In the filtering phase, the Kubernetes filters out the nodes that are suitable for placing pods on them. The Kuberentes scheduler does this filtering based on a set of predicates. For example, PodFitsResources checks if the resource request of the pod is available in the node. Another example is NoDiskConflict which checks if the requested disk space is mounted. In the scoring phase, it scores the nodes that have passed the filtering phase and ranks them based on suitability for scheduling the pod on them. This is done through a set of priorities. For example, LeastRequestedPri-

we implemented a simulator that can mimic the dynamics of the real-world cluster during the training. The trained agent can be later used outside the box in the real-world cluster.

Complete System As it can be seen in Figure 3, our design consist of three main parts. The first part is the Kubernetes cluster containing our Kubernetes nodes and services. The mobility simulator is another part that is connected to the cluster through a simulated connection that is a Python script assigning the users to the nodes. The placement of the nodes and the location of the users is passed to the controller. The controller wraps the information received from both entities into a single environment OpenAI gym [46] environment. The gym environment is then used to calculate the reward based on the current observation from the environment. This information is then passed to the RL agent to decide the next placement of pods in the nodes. We have used the rllib [47] library for the implementation of the RL agents. This placement is then passed back to the Kubernetes using the Kubernetes Python API and the pods will be moved to a new node in the cluster. All the codebase of this project is available at <https://github.com/saeid93/mobile-kube>.

7 EXPERIMENTAL SETUP

We used the Google GKE service to deploy our Kubernetes cluster. Due to our computational budget, we performed the experiments on eight Kubernetes nodes. All the nodes were of e2-standard-4 type of the GKE platform with four cores and 16 GB of Memory. We used 16 stateless containerised services with the Flask app and the utilisation server explained in section 6 with some constant load on them³. We have used pods of guaranteed QoS Kubernetes class which have equal size requests and limits. All the services are of size 250 Mb RAM and 0.125 CPU. Due to the complexity of the problem we have used constant load on the services running on the nodes. This load is generated using the utilisation server explained before. However, as we discussed in section 6, for sensitivity analysis we conducted the experiments for 16, 32, 48, users. In each scenario, a service is serving one, two, and three users respectively. In our simulations, the station and nodes were co-located, therefore, we have eight stations proportional to the number of the nodes. The users move in the radius of 37.72 and 37.78 for latitude and -122.45 to -122.38 in longitude. In Figure 4 you can see the initial placement of nodes, servers, and users in the map. During the training phase of the RL algorithms we used the simulator explained in section 6. The number of user movements available per each user in the cabspotting dataset (explained in Section 6) was not sufficient for training the RL agent. Even with interpolation between the locations in the dataset location of the taxis, we ended up with 3453 locations for each taxi within five minutes intervals. Therefore, for the training phase, we generated a random user movements dataset within the same vicinity for the training with 100000 user movements. However, the final test results in Figures 9, 10, 11 and 12

3. As explained before the scheduling is based on the resource request and not on the load of the containers. However, we put some load on them to more closely emulate a real-world deployment.




Fig. 4. Placement of the users (16 users scenario) and stations/servers on the map, we considered co-located stations and nodes (servers).




Fig. 5. Comparison of scales of the two different reward.

are generated using the real world cabspotting dataset with 3453 timesteps.

The scaling of the training rewards between the two objectives is a challenging task as the two objectives need to be scaled so that the total values of the rewards in an episode are roughly the same scale. We were able to achieve this using experimental evaluation of different scaling methods. In Figure 5 you can see that scale of the rewards within a random 100 timestep episode is almost within the same range.

RL algorithms are very sensitive to variation of hyperparameters [36], therefore using them for new problems requires lots of hyperparameter tuning. By extensive simulation, we found out the following values as the optimal hyperparameters values for each of the RL algorithms. The values are presented in Table 2. For more detail about the meaning of each hyperparameter see [48].

In the testing phase, we averaged the results for 20 separate episodes for each of the experiments.

We compared the results of the learning algorithms with two greedy schedulers that focus on latency or energy consumption only. As there is no work with a similar setting to ours we chose two heuristic algorithms 1. Bestfit algorithm

Neural Network		
# Layers	Layers size	Activation funtion
2	64	Linear
PG		
Train_batch_size	Gamma	Learning rate
1000	0.99	0.0003
PPO		
SGD minibatch size	train_batch_size	learning rate
128	1000	0.0003
IMPALA		
Train batch size	Gamma	Learning rate
1000	0.99	0.0003

TABLE 2
Hyperparameters of the RL algorithms




Fig. 6. Average number of empty servers during the RL algorithms training

as the ground truth value for energy consumption and 2. A latency greedy algorithm that moves the services to the closest vicinity of the users as the ground truth value for latency. We compared IMPALA with two other RL algorithms that have been widely used in previous scheduling works, Vanilla Policy Gradient used in a paper with approximately similar setting [49] and PPO which is one of the state of the art algorithms previously used in systems research and Kubernetes scheduling [50].

8 RESULTS

We evaluated the result in terms of the two energy saving and latency objectives. During the training phase of the algorithm, we observed that among the three tested RL algorithms IMPALA showed the most stable convergence. As you can see in Figure 7 and Figure 8 the green line for IMPALA has fewer variations during the training. This variation is not evident in Figure 6 as all the algorithms converge to some optimal value very quickly. IMPALA also shows a better convergence in all scenarios to the optimal energy saving objective which is 6 empty servers. There was not a consistent performance difference between the Policy gradient and PPO during the training.

For the test results, we also compared the results with the latency and energy saving heuristic algorithms. From Figure 9 we can see that the PG and IMPALA can achieve the best average latency in the network among the three tested RL algorithms. However, the difference between the achieved results in both IMPALA and PG cases is very close.

In the case of the energy saving objective, the IMPALA was able to converge to the optimal result in all network sizes. This is illustrated in Figure 10. This was not the case for the other two RL algorithms.

Analysis of the timesteps (instead of the average of the servers) of a test episode can also confirm similar results




Fig. 7. Average network latency during the RL algorithms training




Fig. 8. Average episode rewards for different RL algorithms during the training




Fig. 9. Average latency of different number of users sizes for different algorithms




Fig. 10. Average number of empty servers of different number of users sizes for different algorithms




Fig. 11. Number of empty servers of different algorithms during a sample episode run




Fig. 12. Average network latency of different algorithms during a sample episode run

to averaged results of multiple episodes (previously shown test results). In Figures 11 and 12 we observe that during a sample 200 timestamp episode the ordering is the same as the averaged results.

We can also see that IMPALA is able to achieve similar performance with the heuristic method that focuses on the energy saving objective only. It can achieve this result while reducing the average latency of users by 43%. This demonstrates the efficacy of the Mobile-Kube system in providing MEC services in a sustainable fashion. From these results, we observed that IMPALA is able to achieve the best result in terms of achieving a trade off between the two aforementioned objectives.

9 CHALLENGES OF USING MACHINE LEARNING FOR RESOURCE MANAGEMENT

There are still significant challenges in using ML-based methods for resource management to make them fully compatible with real in-production systems. Firstly, the central existing ML-based methods are not scalable to datacenters with thousands of jobs waiting to be scheduled. One solution here would be to use multi-agent reinforcement learning [51] instead of a single agent central scheduler. A similar idea has been previously used in some literature in simulated environments [9], [52] but they have not been examined at scale experimentally. Another solution to the scalability problem would be to define a hierarchy in the resource allocation problem and use hierarchical reinforcement learning [53] to have access to different levels of the desired resource allocation objective. In [54] authors proposed the same scheme for resource allocation but the method was not ML-based.

Another challenge is that there are several resource allocation goals e.g. reducing energy consumption and reducing resource fragmentation. However, most of the previous works are only focused on optimising one resource allocation goals. For example, DeepRM and Decima focus on

minimising the average job slowdown. However, the same method could be used for minimizing other resource allocation objectives such as server consolidation. Other than the general objective of the resource allocation achieving the same objective when having multiple resources is yet another challenge. For example, the experiments on Decima were conducted with at most two types of resources at the same time (CPU and Memory) thus extending the same objective for different resources is yet another possible venue for future research. Multi-objective RL methods seems to be a good candidate for achieving such goals in the future.

One problem with most resource management problems is that the dynamics of the environment are not solely dependent on the agents' actions. For example, the job arrival of resources in DeepRM is based on a probability distribution that the agent does not have any knowledge of beyond those jobs that are already in the queue. This violation of Markovian property makes the design of the RL-based resource schedulers harder. The studies on applying reinforcement learning to non-Markovian environments is a longstanding research area [55], however, its applications in systems research and in particular resource management are not fully discovered. To the authors' knowledge the Mao et al. [56] paper is the only research that has tackled one of the challenges of a semi-Markovian environment MDP in the area.

10 CONCLUSIONS AND FUTURE WORK

In this paper, we presented a deep reinforcement learning solution for reducing energy consumption in containerised clusters. Our results suggest that commonly used heuristic algorithms like bin-packing can be replaced with learning-based methods to achieve similar performance for the targeted object while also improving performance in other areas. Some directions for future works are:

Checkpointing of stateful services In this work we only considered the case the case of stateless services that could be turned on and off anytime without the need of preserving their current state. However, this is not the case for many real-world use cases. To this end, a checkpointing mechanism that could preserve the current state of the service until it is restarted can be implemented in future works.

Kubernetes full implementation Currently we have our scheduling control loop completely out of the Kubernetes cluster. Although Mobile-Kube can be deployed in a cloud environment best practice would be to place the control loop inside the Kubernetes cluster. Kubernetes custom resource [57] and operators [58] can be used in future works to achieve this goal.

ACKNOWLEDGMENTS

We gratefully acknowledge Google for their support of this research through their GCP research credit program.

REFERENCES

- [1] C. Ebert, G. Gallardo, J. Hernantes, and N. Serrano, "Devops," *Ieee Software*, vol. 33, no. 3, pp. 94–100, 2016.
- [2] "10 trends in real-world container use," <https://www.datadoghq.com/container-report/>.

- [3] A. Yousefpour, C. Fung, T. Nguyen, K. Kadiyala, F. Jalali, A. Nikanlahiji, J. Kong, and J. P. Jue, "All one needs to know about fog computing and related edge computing paradigms: A complete survey," *Journal of Systems Architecture*, vol. 98, pp. 289–330, 2019.
- [4] O. C. A. W. Group et al., "Openfog reference architecture for fog computing, 2017," URL: https://www.openfogconsortium.org/wp-content/uploads/OpenFog_Reference_Architecture_2_09_17-FINAL.pdf.
- [5] M. Shojafar, N. Cordeschi, and E. Baccarelli, "Energy-efficient adaptive resource management for real-time vehicular cloud services," *IEEE Transactions on Cloud computing*, vol. 7, no. 1, pp. 196–209, 2016.
- [6] T. Ouyang, Z. Zhou, and X. Chen, "Follow me at the edge: Mobility-aware dynamic service placement for mobile edge computing," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 10, pp. 2333–2345, 2018.
- [7] T. Taleb, A. Ksentini, and P. A. Frangoudis, "Follow-me cloud: When cloud services follow mobile users," *IEEE Transactions on Cloud Computing*, vol. 7, no. 2, pp. 369–382, 2016.
- [8] Z. Tang, X. Zhou, F. Zhang, W. Jia, and W. Zhao, "Migration modeling and learning algorithms for containers in fog computing," *IEEE Transactions on Services Computing*, vol. 12, no. 5, pp. 712–725, 2018.
- [9] F. Brandherm, L. Wang, and M. Mühlhäuser, "A learning-based framework for optimizing service migration in mobile edge clouds," in *Proceedings of the 2nd International Workshop on Edge Systems, Analytics and Networking*, 2019, pp. 12–17.
- [10] D. Wang, Y. Peng, X. Ma, W. Ding, H. Jiang, F. Chen, and J. Liu, "Adaptive wireless video streaming based on edge computing: Opportunities and approaches," *IEEE Transactions on services Computing*, vol. 12, no. 5, pp. 685–697, 2018.
- [11] S. Wang, Y. Guo, N. Zhang, P. Yang, A. Zhou, and X. S. Shen, "Delay-aware microservice coordination in mobile edge computing: A reinforcement learning approach," *IEEE Transactions on Mobile Computing*, 2019.
- [12] T. Bahreini and D. Grosu, "Efficient algorithms for multi-component application placement in mobile edge computing," *IEEE Transactions on Cloud Computing*, 2020.
- [13] C. Liu, F. Tang, Y. Hu, K. Li, Z. Tang, and K. Li, "Distributed task migration optimization in mec by extending multi-agent deep reinforcement learning approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 32, no. 7, pp. 1603–1614, 2020.
- [14] S. Ghafouri, A. A. Saleh-Bigdeli, and J. Doyle, "Consolidation of services in mobile edge clouds using a learning-based framework," in *2020 IEEE World Congress on Services (SERVICES)*. IEEE, 2020, pp. 116–121.
- [15] K. Ray, A. Banerjee, and N. C. Narendra, "Proactive microservice placement and migration for mobile edge computing," in *2020 IEEE/ACM Symposium on Edge Computing (SEC)*. IEEE, 2020, pp. 28–41.
- [16] N. Abbas, Y. Zhang, A. Taherkordi, and T. Skeie, "Mobile edge computing: A survey," *IEEE Internet of Things Journal*, vol. 5, no. 1, pp. 450–465, 2017.
- [17] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies, "The case for vm-based cloudlets in mobile computing," *IEEE pervasive Computing*, vol. 8, no. 4, pp. 14–23, 2009.
- [18] C. Mouradian, D. Naboulsi, S. Yangui, R. H. Glitho, M. J. Morrow, and P. A. Polakos, "A comprehensive survey on fog computing: State-of-the-art and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 1, pp. 416–464, 2017.
- [19] Microsoft, <https://blogs.microsoft.com/blog/2020/01/16/microsoft-will-be-carbon-negative-by-2030/>, accessed June 2020.
- [20] Amazon, <https://sustainability.aboutamazon.com/>, accessed June 2020.
- [21] A. Shehabi, S. Smith, D. Sartor, R. Brown, M. Herrlin, J. Koomey, E. Masanet, N. Horner, I. Azevedo, and W. Lintner, "United states data center energy usage report," Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), Tech. Rep., 2016.
- [22] F. Farahnakian, T. Pahlkalla, P. Liljeberg, J. Plosila, N. T. Hieu, and H. Tenhunen, "Energy-aware vm consolidation in cloud data centers using utilization prediction model," *IEEE Transactions on Cloud Computing*, 2016.
- [23] X. Li, W. Yu, R. Ruiz, and J. Zhu, "Energy-aware cloud workflow applications scheduling with geo-distributed data," *IEEE Transactions on Services Computing*, 2020.
- [24] M. Xu, A. N. Toosi, and R. Buyya, "Ibrownot: an integrated approach for managing energy and brownout in container-based clouds," *IEEE Transactions on Sustainable Computing*, vol. 4, no. 1, pp. 53–66, 2018.
- [25] E. Ahvar, A.-C. Orgerie, and A. Lebre, "Estimating energy consumption of cloud, fog and edge computing infrastructures," *IEEE Transactions on Sustainable Computing*, 2019.
- [26] Kubernetes, "K8sNodeAutoscaling," <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler>.
- [27] H. Badri, T. Bahreini, D. Grosu, and K. Yang, "Energy-aware application placement in mobile edge computing: A stochastic optimization approach," *IEEE Transactions on Parallel and Distributed Systems*, vol. 31, no. 4, pp. 909–922, 2019.
- [28] S. Wang, R. Ugaonkar, M. Zafer, T. He, K. Chan, and K. K. Leung, "Dynamic service migration in mobile edge computing based on markov decision process," *IEEE/ACM Transactions on Networking*, vol. 27, no. 3, pp. 1272–1288, 2019.
- [29] V. Mnih, K. Kavukcuoglu, D. Silver, A. A. Rusu, J. Veness, M. G. Bellemare, A. Graves, M. Riedmiller, A. K. Fidjeland, G. Ostrovski et al., "Human-level control through deep reinforcement learning," *Nature*, vol. 518, no. 7540, pp. 529–533, 2015.
- [30] H. Mao, M. Alizadeh, I. Menache, and S. Kandula, "Resource management with deep reinforcement learning," in *Proceedings of the 15th ACM Workshop on Hot Topics in Networks*, 2016, pp. 50–56.
- [31] V. Mnih, K. Kavukcuoglu, D. Silver, A. Graves, I. Antonoglou, D. Wierstra, and M. Riedmiller, "Playing atari with deep reinforcement learning," *arXiv preprint arXiv:1312.5602*, 2013.
- [32] A. Mirhoseini, H. Pham, Q. V. Le, B. Steiner, R. Larsen, Y. Zhou, N. Kumar, M. Norouzi, S. Bengio, and J. Dean, "Device placement optimization with reinforcement learning," *arXiv preprint arXiv:1706.04972*, 2017.
- [33] Y. Gao, L. Chen, and B. Li, "Spotlight: Optimizing device placement for training deep neural networks," in *International Conference on Machine Learning*, 2018, pp. 1676–1684.
- [34] H. Mao, R. Netravali, and M. Alizadeh, "Neural adaptive video streaming with pensieve," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, 2017, pp. 197–210.
- [35] N. Jay, N. H. Rotman, P. Godfrey, M. Schapira, and A. Tamar, "Internet congestion control via deep reinforcement learning," *arXiv preprint arXiv:1810.03259*, 2018.
- [36] H. Mao, P. Negi, A. Narayan, H. Wang, J. Yang, H. Wang, R. Marcus, M. K. Shirkoobi, S. He, V. Nathan et al., "Park: An open platform for learning-augmented computer systems," in *Advances in Neural Information Processing Systems*, 2019, pp. 2494–2506.
- [37] J. Doyle, D. O'Mahony, and R. Shorten, "Server selection for carbon emission control," in *Proceedings of the 2nd ACM SIGCOMM workshop on Green networking*, 2011, pp. 1–6.
- [38] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 2018.
- [39] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [40] L. Espeholt, H. Soyer, R. Munos, K. Simonyan, V. Mnih, T. Ward, Y. Doron, V. Firoiu, T. Harley, I. Dunning et al., "Impala: Scalable distributed deep-rl with importance weighted actor-learner architectures," in *International Conference on Machine Learning*. PMLR, 2018, pp. 1407–1416.
- [41] "Kubernetes Documentation," <https://kubernetes.io/docs/concepts/overview/what-is-kubernetes/>.
- [42] "Flask Documentation," <https://flask.palletsprojects.com/en/2.0.x/>.
- [43] "Kubernetes Scheduler Plugin," <https://github.com/kubernetes/enhancements/blob/master/keps/sig-scheduling/624-scheduling-framework/README.md>.
- [44] "antennasearch," <http://antennasearch.com/>.
- [45] M. Piorkowski, N. Sarafjanovic-Djukić, and M. Grossglauser, "CRAWDAD dataset epfl/mobility (v. 2009-02-24)," Downloaded from <https://crawdad.org/epfl/mobility/20090224>, Feb. 2009.
- [46] G. Brockman, V. Cheung, L. Pettersson, J. Schneider, J. Schulman, J. Tang, and W. Zaremba, "Openai gym," *arXiv preprint arXiv:1606.01540*, 2016.
- [47] E. Liang, R. Liaw, R. Nishihara, P. Moritz, R. Fox, K. Goldberg, J. Gonzalez, M. Jordan, and I. Stoica, "Rllib: Abstractions for distributed reinforcement learning," in *International Conference on Machine Learning*. PMLR, 2018, pp. 3053–3062.
- [48] "rllib algorithms." [Online]. Available: [@inproceedings{binpacking,title={Binpackingproblem},url={https://en.wikipedia.org/wiki/Binpacking_problem}}](https://en.wikipedia.org/wiki/Binpacking_problem)

- [49] H. Mao, M. Schwarzkopf, S. B. Venkatakrishnan, Z. Meng, and M. Alizadeh, "Learning scheduling algorithms for data processing clusters," in *Proceedings of the ACM Special Interest Group on Data Communication*, 2019, pp. 270–288.
- [50] S. Li, L. Wang, W. Wang, Y. Yu, and B. Li, "George: Learning to place long-lived containers in large clusters with operation constraints," in *Proceedings of the ACM Symposium on Cloud Computing*, 2021, pp. 258–272.
- [51] L. Busoniu, R. Babuska, and B. De Schutter, "A comprehensive survey of multiagent reinforcement learning," *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, vol. 38, no. 2, pp. 156–172, 2008.
- [52] K. Haghshenas, A. Pahlavan, M. Zapater, S. Mohammadi, and D. Atienza, "Magnetic: Multi-agent machine learning-based approach for energy efficient dynamic consolidation in data centers," *IEEE Transactions on Services Computing*, 2019.
- [53] S. Li, R. Wang, M. Tang, and C. Zhang, "Hierarchical reinforcement learning with advantage-based auxiliary rewards," in *Advances in Neural Information Processing Systems*, 2019, pp. 1409–1419.
- [54] A. A. Bhattacharya, D. Culler, E. Friedman, A. Ghodsi, S. Shenker, and I. Stoica, "Hierarchical scheduling for diverse datacenter workloads," in *Proceedings of the 4th annual Symposium on Cloud Computing*, 2013, pp. 1–15.
- [55] J. Schmidhuber, "Reinforcement learning in markovian and non-markovian environments," in *Advances in neural information processing systems*, 1991, pp. 500–506.
- [56] H. Mao, S. B. Venkatakrishnan, M. Schwarzkopf, and M. Alizadeh, "Variance reduction for reinforcement learning in input-driven environments," *arXiv preprint arXiv:1807.02264*, 2018.
- [57] "Kubernetes Custom Resources," <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>.
- [58] "Operator pattern," <https://kubernetes.io/docs/concepts/extend-kubernetes/operator/>.

BIOGRAPHIES



Saeid Ghafouri received his B.Sc. degree in computer software engineering from Shahid Chamran University of Ahvaz, Ahvaz, Iran. He received his M.Sc. in the field of Artificial Intelligence from the Computer Engineering Department, K. N. Toosi, University of Technology, Tehran, Iran. He is currently working toward his PhD degree in computer science at Queen Mary University of London. His current research interests include cloud and edge computing, resource allocation in cloud and edge computing

and reinforcement learning.



Alireza Karami Motlagh received his B.Sc. in Software Engineering from Shahid Chamran University of Ahvaz in 2017. He is currently an M.Sc. student at Shiraz University majoring in Artificial Intelligence. His major research interests are Software Security, System Resource Management and Machine Learning. He has professional experience as a Site Reliability Engineer.



Danial Bidekani Bakhtiarvand earned a bachelor's degree in software engineering from Shahid Chamran University of Ahvaz. He graduated from K. N. Toosi University of Technology with an MSE in Artificial Intelligence in 2021. He has worked with financial institutions on a regular basis since then. From 2019 to 2020, he worked for blu Bank as a Data Engineer. He then began his career as a Site Reliability Engineer at Sadad Electronic Payment. His research focuses on information retrieval, user behavior

analysis and cloud computing.



Ali Akbar Saleh Bigdeli, received his bachelor's in Softwar Engineering from Persian Gulf University of Bushehr (Bushehr, Iran) in 2016 and the master's degree in Artificial Intelligence from Khaje Nasir Toosi University of Technology (Tehran, Iran) in 2019. Since then, he is actively working on AI products at the IT industry and tries to tackle the problems in IT industry by 12 contributing to AI research investigations.



Sukhpal Singh Gill is a Lecturer (Assistant Professor) in Cloud Computing at School of Electronic Engineering and Computer Science, Queen Mary University of London, UK. Prior to this, 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, Software Engineering, Internet of Things and Healthcare. For further information, please visit <http://www.ssgill.me>



Joseph Doyle graduated from Trinity College Dublin in 2009 with a B.A.I., B.A. degree in Computer and Electronic Engineering as a gold medalist. He was awarded a Ph.D in 2013 from Trinity College Dublin. He was a post-doctoral researcher in Trinity College Dublin and University College London from 2013 to 2014 and 2014 to 2016, respectively. He was Senior Lecturer in the University of East London, London, U.K. from 2016 to 2018. He is a co-founder of Dithen Ltd. (London, U.K.) and is also Lecturer at Queen

Mary University of London, London, U.K. He is an RSE cloud fellow. His research interests include cloud computing, virtual machine classification, green computing, and network optimization.