

Mitigating Cold Start Problem in Serverless Computing: A Reinforcement Learning Approach

Parichehr Vahidinia¹, Bahar Farahani², Fereidoon Shams Aliee¹

Abstract—Serverless computing has revolutionized the world of cloud-based and event-driven applications with the introduction of Function-as-a-Service (FaaS) as the latest cloud computing model. This computational model increases the level of abstraction from the infrastructure and breaks the program into small units called functions. Thus, it brings benefits such as ease of development, saving resources, and reducing product launch time for enterprises and developers. Thanks to the scale-to-zero feature of this computational model, idle functions with no traffic will be depreciated from memory. However, this cost-saving approach adversely impacts delay leading to the cold start problem. Unfortunately, the existing solutions to alleviate the cold start delay are not resource-efficient as they follow a fixed policy over time. Thereby, this paper proposes a novel two-layer adaptive approach to tackle this issue. The first layer utilizes a holistic reinforcement learning algorithm to discover the function invocation patterns over time for determining the best time to keep the containers warm. The second layer is designed based on a Long Short-Term Memory (LSTM) to predict the function invocation times in the future to determine the required pre-warmed containers. The experimental results on the Openwhisk platform show that the proposed approach reduces the memory consumption by 12.73% and improves the execution invocations on pre-warmed containers by 22.65% compared to the Openwhisk platform.

Index Terms—Serverless Computing, Function-as-a-Service (FaaS), Cold Start Delay, Serverless Platforms, Memory Consumption, Openwhisk, Reinforcement Learning

I. INTRODUCTION

With the advent of the digital revolution and the entry into the information age in the mid-20th century, traditional organizations changed, e-commerce emerged, and computing and communications were based on the Internet. Therefore, the use of information technology became inevitable, and organizations needed infrastructure, network-based technologies, computing resources, and storage. This path was followed by the advent of client/server architecture and centralized computing. However, implementation costs, human resources, infrastructure maintenance concerns, and growing data volumes led to the idea of cloud computing.

In cloud computing, infrastructure, computing and storage resources, platforms and services are shared virtually, dynamically and on-demand over the Internet. Cloud computing has three models of services. These models are Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS) [1].

In 2014, Amazon introduced the latest cloud computing model called serverless computing. This computational model

reduces operational costs and the complexity of system development, and also increases agility in an enterprise. Because any responsibility for preparing, implementing, and managing infrastructure for product development is on the cloud provider. By reducing these responsibilities, the enterprise can focus on its business. It also allows developers to focus only on writing the core logic of their functions and finally place them on a serverless platform [2], [3], [4].

Despite all the benefits of serverless computing, it faces some challenges. These challenges briefly divided into three categories: system, performance, and software engineering challenges. One of the most critical performance challenges is the cold start delay, which is considered as a delay in preparation of the function execution environment [5]. The most common methods used by popular platforms to reduce cold start delay are using a pool of warm containers, reusing the containers, and calling functions periodically. However, these methods waste resources such as memory, increase cost and do not have any insight into function invocation patterns over time. In other words, these methods provide fixed mechanisms to reduce cold start delay, and they are not suitable for dynamic cloud environment.

The convergence of serverless computing and the Internet of Things (IoT), particularly for edge applications, is an emerging paradigm that has attracted researchers' attention in recent years. Many IoT devices are resource-constrained by processing and storage capacity as well as power consumption requirements. Serverless computing and the corresponding serverless functions ensure that the number of resources consumed by an application running on a resource-constrained IoT device can be controlled on-demand and dynamically. In addition, despite the scale-to-zero feature, edge device resources such as energy sources are saved. Another main advantage that serverless computing can provide to the edge is the auto scalability feature, as load balancing and scalability are handled by the cloud provider. Moreover, serverless computing can play an important role in deploying IoT applications. Due to the lightweight properties of serverless functions, they can be easily deployed and placed independently anywhere in the cloud, edge, or fog [6].

Although serverless computing alleviates some of the major challenges of IoT, still these converging technologies suffer from specific limitations such as cold start delay that must be approached holistically. Indeed, real-time interaction, prompt processing, and low-latency response are the primary needs for a wide variety of context-sensitive IoT applications from healthcare to online games. Thereby, tackling the cold start delay is the key to the successful integration of IoT and

¹Faculty of Computer Science and Engineering, Shahid Beheshti University

²Cyberspace Research Institute, Shahid Beheshti University

Corresponding author: Bahar Farahani (b_farahani@sbu.ac.ir)

serverless computing [7].

In this paper, in order to reduce cold start delay and consider resource consumption, we propose an intelligent approach that determines the best policy for keeping the containers warm, according to the function invocations over time.

The rest of the paper is organized as follows. In section II, the basic concepts and features of serverless computing are explained. In Section III, the related works for reducing cold start delay are categorized and compared. Section IV describes the proposed two-layer approach to balance the number of cold start occurrences and memory consumption. Section V presents and discusses the experimental results. Finally, section VI concludes the paper.

II. SERVERLESS COMPUTING

Defining serverless computing can be a bit tricky. However, one way to define it is to consider the amount of control the developer has over resources. The IaaS model is where the developer has the most control over both the application and the cloud operating infrastructure. In the PaaS model, the developer is not aware of the infrastructure, therefore has no control over it. But has access to a ready platform to manage the entire software stack. Finally, SaaS hosts applications as a service in the cloud, and users across the Internet can use them without installing them on the system. However, serverless computing is somewhere between PaaS and SaaS. As shown in Fig. 1, this computational model increases the granularity of applications. According to it, an application is broken down into independent units called functions and introduces a new service model called Function as a Service (FaaS). In other words, serverless computing is a general concept, and one way to perform it is Function as a Service.

Each function is a small, short-lived, stateless, and task-specific piece of code that is executed in response to events. In other words, the functions are event-driven. These events include the application interface requests (such as an HTTP request), changes to the database, and events from IoT devices [5], [8], [9].

Serverless functions are deployed on containers. As soon as a function is triggered, a new container is created, and the function with all its dependencies is deployed on it. It should be noted that if several invocations are occurred simultaneously to execute a function, new instances of the function are created according to the number of requests, and each instance is deployed and executed independently on a separate container. According to Fig. 2, when an invocation occurs from an event source, the following steps must be performed to execute the function [10], [11]:

- 1) Initialize a new container that includes preparing the container environment such as setting environment variables, connecting to the database, user authentication, etc.
- 2) Allocate memory and computing resources.
- 3) Load the function and required libraries.
- 4) Execute the function.
- 5) Return the results to the user and save the logs in the database.

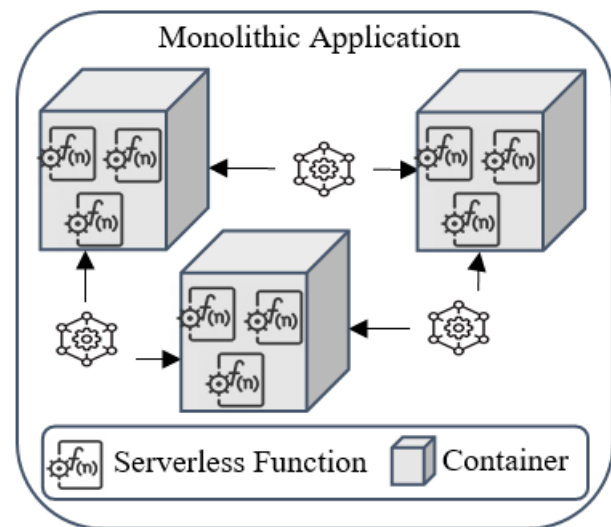


Fig. 1: Increasing the Granularity of Applications by Serverless Computing.

The first three steps take time and leading to a delay in the overall response time. This delay is called cold start. At the end of the execution, the container with all its resources will be released. This is called scale-to-zero that is an important feature of serverless computing and leads to cost reduction. According to this, users do not pay for times when the functions are idle. On the other hand, it also leads to cold start challenge. Because when new invocations occur, all previous steps must be repeated. On the other hand, if the number of simultaneous invocations exceeds the number of available containers, several instances of a function must be created, and each instance must go through the above steps separately. Because as mentioned, each instance of the function is placed on a separate container. Therefore, the functions have a cold start delay at the first invocation and when the scale increases [10], [12], [13].

III. RELATED WORK

To reduce the cold start delay, two general approaches can be considered. These approaches include: Reduce cold start delay and occurrence.

A. Reduce Cold Start Delay

The reduce cold start delay approach includes solutions that focus on reducing the delay in each step of the function execution mentioned in Section I.

- Reduce container preparation delay:
 - The Openwhisk platform uses pre-warmed stem cell containers which categorizes the containers according to the memory and programming language. By this approach, the preparation time of the container and consequently the cold start delay is reduced. By default, there are two of them with Node.js execution environment. Each time one of them is assigned to a request, another pre-warmed container is created immediately [14].

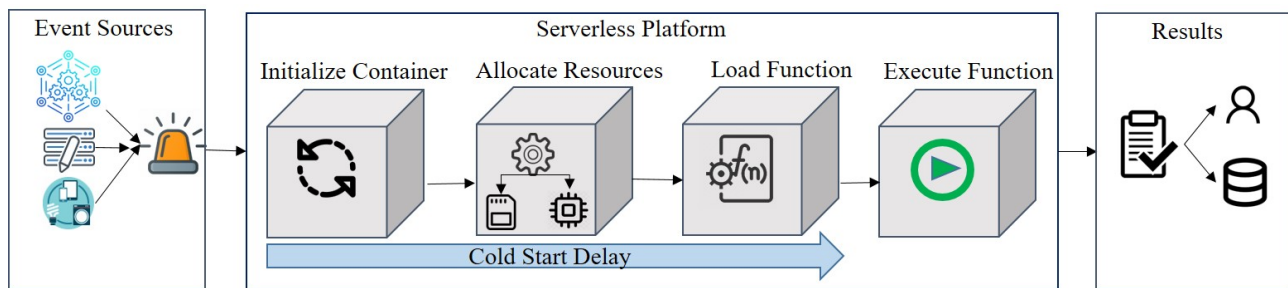


Fig. 2: Function-as-a-Service Execution Workflow.

- Akkus et al. in [15] reduce container preparation delay by providing a platform called SAND that performs logic separation at the application level. Furthermore, Oakes et al. in [16] provide a platform called SOCK to deal with performance bottlenecks of Linux kernels used in containers. Moreover, Silva et al. in [17] introduce a method called Prebaking that reuses the snapshots created for the previous functions.
 - Xu et al. in [18] use Adaptive Container Pool Scaling Strategy (ACPS). According to it, there is a pool of pre-launched containers that are classified according to different programming languages. Moreover, the number of cold starts for each category is predicted to determine the number of container instances adaptively.
 - Loading function code and other dependencies into memory:
 - The application-level separation method introduced by Akkus et al. in [15] places all the functions of an application on the same container. In addition, Oakes et al. in [16] use the Zygote technique for pre-importing required libraries.
 - To reduce the delay of calling functions by each other, Akkus et al. in [15] also use a messaging queue to reduce the delay of internal events.
- B. Reduce Cold Start Occurrence**
- The aim of reducing the cold start occurrence approach is to reduce the number of times the function suffers from this delay. For this approach, two general solutions are presented in the literature: Prepare function containers and container reuse.
- Prepare function containers: To reduce the cold start occurrence, one approach is to prepare a ready container before the request for a function is entered.
 - Tools such as CloudWatch [19], Thundra.io [20], Dashbird.io [21], and Lambda Warmer [22] were used to set rules to periodically invoke functions and warm-up them (for example, every 5 minutes).
 - Xu et al. in [18] present a strategy called Adaptive Warm-UP that uses the Long Short Term Memory (LSTM) network to predict future invocation times. In addition, Gunasekaran et al. in [23] provide a resource management framework called Fifer that predicts future workloads using the LSTM network.
 - Agarwal et al. in [24] uses the Q-learning algorithm to determine the number of function instances dynamically. The number of available function instances and the discretize CPU consumption of each function, are considered as environmental states, and the action is scaled up or down of function instances.
 - Container reuse: According to this approach, containers of the previous execution are reused for new request.
 - In some platforms like Google Cloud Functions, AWS Lambda, Microsoft Azure, Open Lambda, and Openwhisk, after completing a function execution, the container and its resources are paused for a few minutes. Thus, if a request is received during this period, the container will be reused and it will have a warm start. More specifically, in the openwhisk platform, the container stays running or hot for 50 milliseconds after an execution. After this time, the container and its resources are paused, and after 10 minutes they are completely released. This time is called the idle-container window. This method can be effective in reducing delay. Because restarting the paused container has less delay than allocating a new container [25].
 - OpenFaaS gives the developer the option to use an always-running container for a function [26]. Knative and fission platforms have a pool of warm containers [27], [28]. Lloyd et al. in [29] also use an approach called Keep-alive to keep containers warm at all times. McGrath et al. in [30] use a queue of warm containers per function. In addition, Solaiman et al. in [31] present a container management architecture called WLEC that has warm and template queues. The template queue keeps a copy for each warm container.
 - The AWS Lambda platform provides Provisioned Concurrency mechanism that the developer can determine the number of required warm function instances for future requests [32].
 - Shahrad et al. in [33] present a resource management policy called the Hybrid Histogram Policy that determines the optimal value of the idle-container window according to the characteristics of the functions.
 - Gias et al. in [34] perform capacity planning for serverless functions using a queueing-based approach and considering the effect of cold start delay. Therefore, the optimal

TABLE I: Comparison of Related Work.

Reference	Platform/Method Name	Reduce Container Preparation Delay			Reduce Cold Start Occurrence		Analysis	
		Reduce Container Preparation Delay	Function Chain Calling	Loading Dependencies into Memory	Prepare Function Containers	Container Reuse	Optimal Resource Consumption	Adaptability
[14]	Openwhisk	✓	x	x	x	✓	x	x
[15]	SAND	✓	✓	✓	x	x	✓	x
[16]	SOCK	✓	x	✓	x	x	✓	x
[17]	Prebaking	✓	x	x	x	x	x	x
[18]	Adaptive Warm-up and Adaptive Container Pool Scaling Strategies	✓	x	x	✓	x	✓	✓
[19-22]	CloudWatch, Thundra.io, Dashbird.io, Lambda Warmer	x	x	x	✓	x	x	x
[23]	Fifer	x	x	x	✓	x	✓	✓
[24]	-	x	x	x	✓	x	✓	✓
[25]	Open Lambda	x	x	x	x	✓	x	x
[26]	OpenFaaS	x	x	x	x	✓	x	x
[27]	Knative	x	x	x	x	✓	x	x
[28]	Fission	x	x	x	x	✓	x	x
[29]	Keep-alive	x	x	x	x	✓	x	x
[30]	-	x	x	x	x	✓	x	x
[31]	WLEC	x	x	x	x	✓	x	x
[32]	AWS Lambda	x	x	x	x	✓	x	x
[33]	Hybrid Histogram Policy	x	x	x	x	✓	✓	✓
[34]	COCOA	x	x	x	x	✓	✓	x
[35]	ETAS	x	x	x	x	✓	x	x
[36]	CAS	x	x	x	x	✓	x	x
-	Two-layer Proposed Approach	✓	x	x	x	✓	✓	✓

idle time of a container is determined by considering the resource consumption.

- Banaei et al. in [35] provide a function scheduler called ETAS which predicts the execution time of a invocation and assigns the best worker node to it based on the availability of a warm container or the availability of the requested resources.
- Wu et al. in [36] propose a container lifecycle-aware scheduling called CAS. This approach has 3 strategies that monitor the state of the containers and select the best container to execute the input invocation to prevent the creation of new containers and cold start occurrence.

C. Comparison and Analysis of Related Work

In this section, we analyze the related work in terms of resource consumption and adaptability.

- Resource consumption analysis: Resource consumption analysis aims to determine whether reducing cold start delay in a method leads to higher resource consumption. Note that processes such as keeping the containers warm, pausing and reusing it, creating a pool of warm containers, and calling functions waste resources. The reasons are as follows:
 - Container reuse: This method has no threshold to determine the best time to release the container after completing a function. Therefore, if subsequent requests are received long after the container has been released, keeping the container warm will only lead to additional resources and costs.

- Pool of warm containers: These containers are always running and occupy memory and computing resources in the presence or absence of requests.
- Calling functions periodically: This method does not have any insight into the time and number of future requests. There may not be a request in a while, thus we do not need to restart the function container. On the other hand, the requests may be bursty, and we have not considered enough instances for all of them.
- Adaptability: The purpose of adaptability analysis is to determine whether a method for reducing cold start delay considers the pattern of invocations using machine learning algorithms. Because given the dynamics of the cloud environments, prediction of the function invocations over time cannot be ignored. This proves the importance of using machine learning algorithms. Methods that provide fixed mechanisms to reduce cold start delay may not respond to such a changing environment. Due to the variability of function invocation times and the dynamic nature of the cloud environment, we need a self-adapting strategy that can provide automated decisions by monitoring the execution environment. Therefore, reinforcement learning-based algorithms seem to be suitable for this area. Because they do not need prior knowledge of the environment and learn the optimal action when the system is running. On the other hand, classical reinforcement learning algorithms, such as Q-learning, have a discrete state and action space. However, when it comes to reducing cold start delay, environmental states such as memory and CPU

consumption and invocation times are continuous, therefore it is best to use deep reinforcement learning algorithms. A comparison of the methods is given in Table I. Each row of this table represents the features and information of a method. Table I contains five columns. The first column is the reference number of the presented method. The second column includes the name of the platform. The third and fourth columns indicate that the purpose of this method is to reduce the cold start occurrence or to reduce the cold start delay. The fifth column compares the methods in terms of resource consumption and adaptability. According to Table I, most methods only either reduce the cold start delay or reduce its occurrence. However, the Openwhisk platform and Xu et al. in [18] provide methods to reduce the cold start occurrence and delay. On the other hand, according to the efficiency of keeping the containers warm in reducing the cold start delay, there is a need to determine the value of keeping the container warm according to the invocation patterns. Because the advantage of these methods is reusing the existing containers instead of creating new containers.

IV. THE POPOSED TWO-LAYER APPROACH

The most common way to reduce cold start delay among popular platforms is to keep the containers warm for a fixed period of time after completing a function. This is called the idle-container window. However, among these platforms, only Openwhisk has provided solutions to reduce both the time and number of cold start delays. However, since the pattern of function invocations varies over time, considering a constant value for this window is not an efficient method. This proves the importance of using machine learning algorithms to predict future function invocation times. Therefore, the Openwhisk platform method is considered the baseline of this paper, and in this section, a two-layer approach is proposed to make the platform adaptable.

The first layer aims to determine the best time for the idle-container window to reduce the number of cold start delays and resource consumption using the reinforcement learning algorithm. Because there is a tradeoff between the number of cold start occurrences and resource consumption. Therefore, the longer this window is, the shorter cold start delays and more consumption of resources.

The purpose of the second layer is to reduce the cold start delay time. When requests are received bursty, there may not be a warm container for each one, and some requests will inevitably be delayed. Therefore in this layer, the number of requests that will be delayed in the future is predicted using LSTM and based on them, the number of pre-warmed containers required in the future is determined. Fig. 3 shows an overview of this approach.

A. Cold Start Occurrence Reduction Layer

As mentioned, the purpose of this layer is to reduce the number of cold start occurrence over a period of time. Given that the state and action space is continuous, there is a need to receive immediate rewards at each step instead of

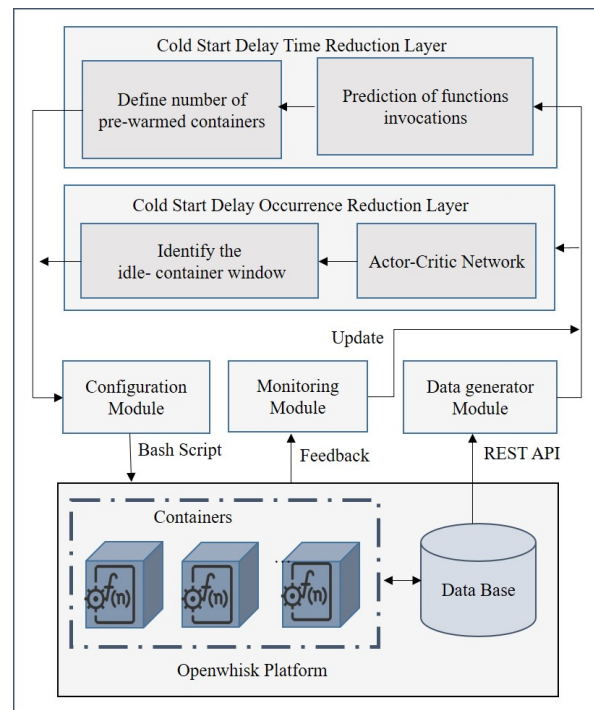


Fig. 3: Two-layer Proposed Approach.

each episode, and also considering that the approach used to optimize the agent behavior should be on-policy, the Temporal Difference (TD) Advantage Actor-Critic algorithm, which is a deep reinforcement learning algorithm, is used. Therefore, the intelligent agent learns the invocation pattern of the function by observing the time intervals between invocations in the past and determines the value of the idle-container window in the future. The steps of this algorithm are as follows:

- 1) Determine the state of the environment: The function execution log is received from the platform and is considered as state space. Therefore, the state space includes the time intervals between the arrival of invocations and whether it had a cold or warm start.
- 2) Determine appropriate action: The action is to determine the value of idle-container window. In each step, we give the state as input to the Actor neural network. This neural network then returns two values of μ and σ as the mean and standard deviation at the output. μ and σ are functions of state s . Since we are interested in the case where the action is a continuous variable, we used a stochastic policy $\pi = P(a|s)$, where $P(a|s)$ is the probability distribution of taking action a given state s . Thereby, P will be a probability density. We use a Gaussian probability distribution to represent our stochastic policy, where μ is the mean, σ is the standard deviation of the Gaussian (normal) distribution abbreviated as $N(a|\mu, \sigma)$, and a is a continuous action drawn from. Therefore, the action is randomly selected from a normal distribution with mean μ and standard deviation σ .
- 3) Take action on the environment: The serverless computing

platform is then updated with a new value for the idle-container window parameter by the reconfiguration module.

- 4) Receive rewards: at the end of a time period, and based on the information received from the monitoring module, the reward function is calculated based on Equation (1). As mentioned, if the size of the idle-container window is large, the number of delays is reduced. However, resource consumption is increased because the idle container wastes memory. Therefore, the reward function is based on the ratio of the number of cold start delays to total invocations and a penalty based on the amount of memory wastage when keeping the container warm. If the number of delays is high, the wasted memory is reduced. However, the reward is a high negative number. Conversely, if the number of delays is low, the wasted memory will increase, and the reward will still be a high negative number. Hence, this algorithm will balance the number of delays and memory wastes. The number of cold start occurrences, the total number of invocations, and the penalty are indicated by Cold, N, and P, respectively.

$$Reward = -(Cold)/N - P \quad (1)$$

B. Cold Start Delay Reduction Layer

A cold start delay is inevitable. Because some invocations occur at long intervals, keeping the containers warm will result in memory wastage. Therefore, some invocations that enter outside the idle-container window will suffer a cold start delay. Moreover, invocations may be occurring simultaneously, and there may not be enough warm containers for all of them. The purpose of this layer is to adapt the number of stem cell containers. This layer predicts the maximum number of concurrent invocations over a period of time, and accordingly determines the number of pre-warmed containers. This is a time series problem because the number of concurrent invocations in the future depends on the pattern of the previous invocations. Thus, in this approach we will use the LSTM.

V. EVALUATION

In this section, the efficiency of the proposed approach is evaluated through the simulations. The following is a description of the Experimental Setup and Performance Evaluation.

A. Experimental Setup

This section includes the description of the Dataset, Computing Platform, and Training Setup.

- **Dataset:** We considered two types of datasets. One includes sequential invocations, and the other includes concurrent invocations. To evaluate the first layer of the proposed approach, a dataset containing sequential invocations is required. According to the dataset simulation introduced by Gias et al. in [34], we expect invocations to enter based on the Poisson process. According to this process, the distribution of invocations entry time is random and exponentially distributed. We created 10,000 entry times with different 5,

10, and 20 average entry rates of invocations per hour. These entry times represent the time interval between invocations. To evaluate the second layer of the proposed approach, a dataset containing concurrent invocations is required. We randomly created 10,000 invocations, including concurrent invocations, to determine how this concurrency affects the idle-container window and the number of pre-warmed containers, as well.

- **Computing Platform:** As mentioned in Section III, the Openwhisk platform method is used as the baseline for this paper. The purpose of this paper is using Reinforcement Learning algorithm to make this method adaptable in order to reduce memory consumption. The Openwhisk platform is implemented on an Ubuntu server by a Kubernetes cluster. Then, the configuration parameters corresponding to the parameters of the proposed approach on the platform should be determined. This platform uses a parameter called idle-container to keep the containers warm. However, by default this parameter cannot be configured to the desired value. Therefore, in this step, we first configured this parameter. In addition, this platform has an adjustable configuration file for the number of its pre-warmed containers.
- **Training Setup:** Each layer of the proposed approach has its own hyperparameters. As mentioned in section IV, in the first layer of the proposed approach, the TD Advantage Actor-Critic algorithm is used to determine the idle-container window. To implement the actor network, three hidden layers have been used. Each layer contains 32 neurons, a ReLU activator function, and a uniform weight function. This network also has two output layers for μ and σ as mean and standard deviation, which have Softplus activation function. Moreover, the Critic neural network has two layers with 32 and 16 neurons and a ReLU activation function. The output layer also has a linear activation function to calculate the value of the next state. It should be noted that, by adding a regularizer, the training will stop faster, and the loss function will fluctuate. Fig. 4a and Fig. 4b are the actor loss function and critic loss function for training without regularization, respectively. In addition, Fig. 5a and Fig. 5b are the actor loss function and critic loss function for training with regularization, respectively. The LSTM network has five hidden layers, each with 32 neurons and a ReLU activation function. In addition, to prevent overfitting, a Dropout layer with a value of 0.5 is used. The network also has one output and a linear activations function. Note that the Mean Squared Error (MSE) is also used as a loss function. Note that, the training was done in Google Colab using a TPU v2 node with 8 cores and 64 GB of total memory, and it took 2 hours and 20 minutes. We have considered a monitoring module that activates the trigger to retrain the model if the model accuracy decreases or the data distribution changes. In addition, we can utilize continuous learning to adapt to change.

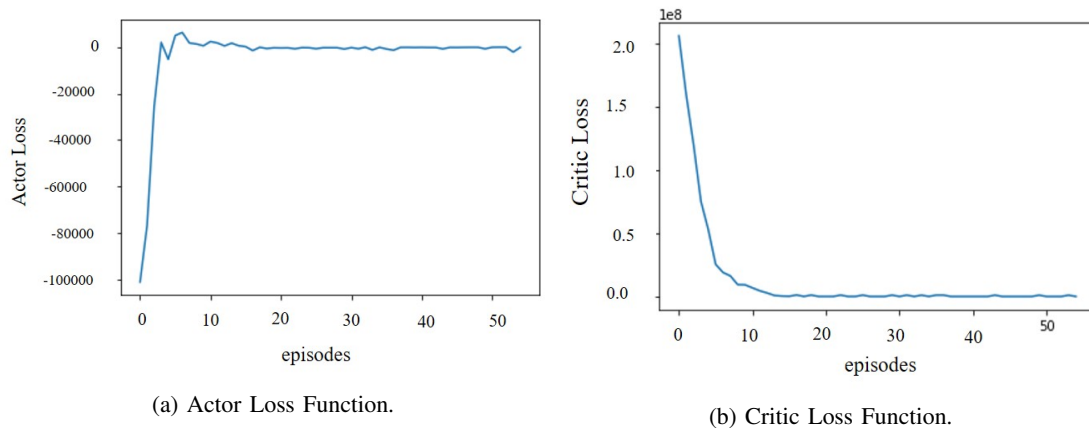


Fig. 4: Training Loss Functions without Regularizer.

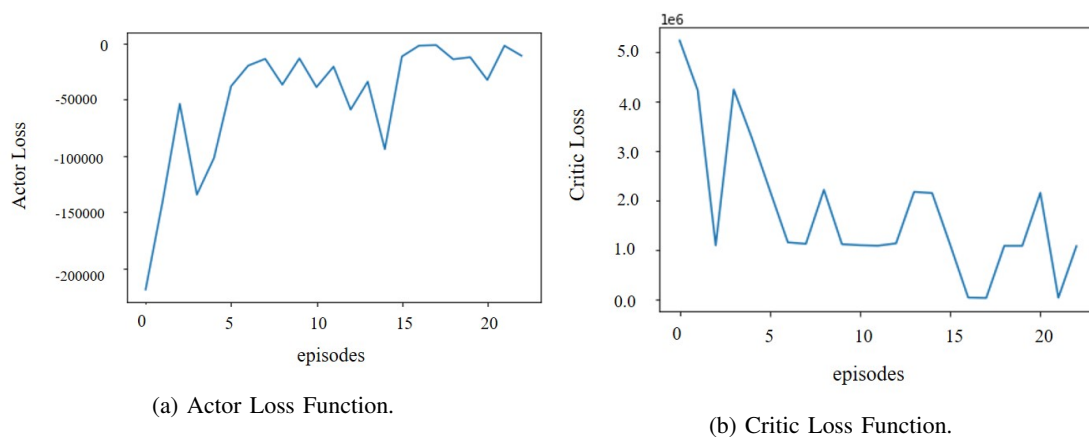


Fig. 5: Training Loss Functions with Regularizer.

TABLE II: The Number of Cold Start Occurrence.

Average Invocation Rate per Hour	Proposed Approach	Openwhisk Platform	Memory Improvement
5	123	126	11.11%
10	84	89	12.73%
20	53	55	4.05%

B. Performance Evaluation

In order to perform evaluation, we used an I/O bound function that sends an HTTP request to a web page and receives a text file. The response time of this function is about 6 seconds. Then, according to the simulated intervals, asynchronous requests are sent to the function through the REST API of the Openwhisk platform. The result of each request will be recorded in the Openwhisk platform log service. Two simulations are performed: 1) Using the Openwhisk platform with its default parameters. 2) Setting new parameters for Openwhisk obtaining from the proposed approach. In the second simulation, the results of each layer of the proposed approach are compared to the results of the first simulation. Note that test dataset includes 200 requests.

- **Cold Start Occurrence Reduction Layer:** Comparison

metrics in this layer are the value of the idle-container window, number of cold start occurrences, and memory consumption.

- **Idle-container window:** The idle-container window of the Openwhisk platform is fixed at 10 minutes. However, according to the proposed approach, this window has different values for each average entry rate at different time intervals. Fig. 6 compares the idle-container window of the proposed approach and the Openwhisk platform over time for three different average invocations per hour (5, 10, 20). The horizontal axis shows the number of time intervals of 200 invocations according to each average entry rate. For example, when we have 5 average invocations per hour, 200 requests have 40 time intervals (200/5). Also, the idle-container window varies from 3 to 13 minutes. It should be noted that, this window varies from 7 to 15 minutes, and 8 to 13 minutes for average invocations per hour of 10 and 20, respectively. According to the figures, the idle-container window is determined according to the invocations entry pattern using the proposed Reinforcement Learning algorithm. Fig. 7 compares the Cumulative Distribution Function

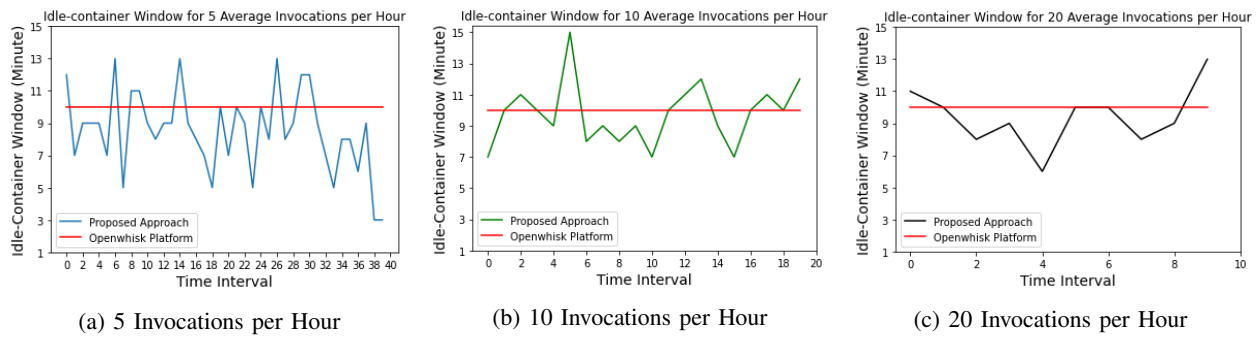


Fig. 6: Idle-Container Window According to Proposed Approach and Openwhisk Platform.

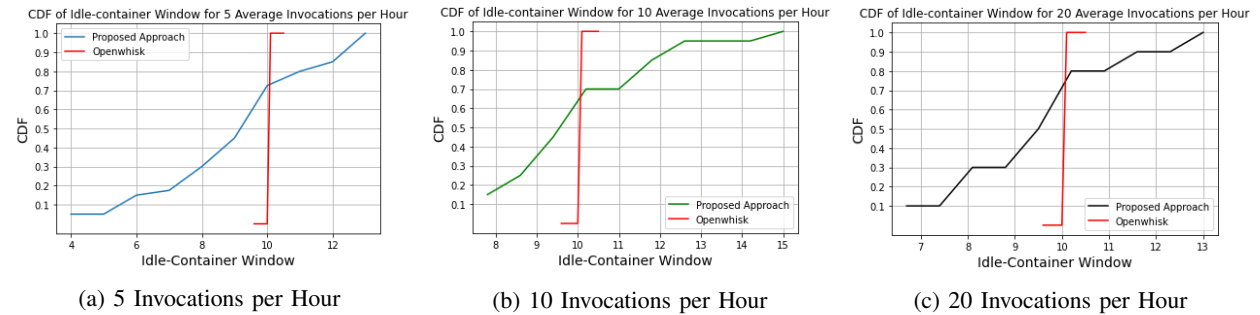


Fig. 7: CDF of Idle-Container Window According to Proposed Approach and Openwhisk Platform.

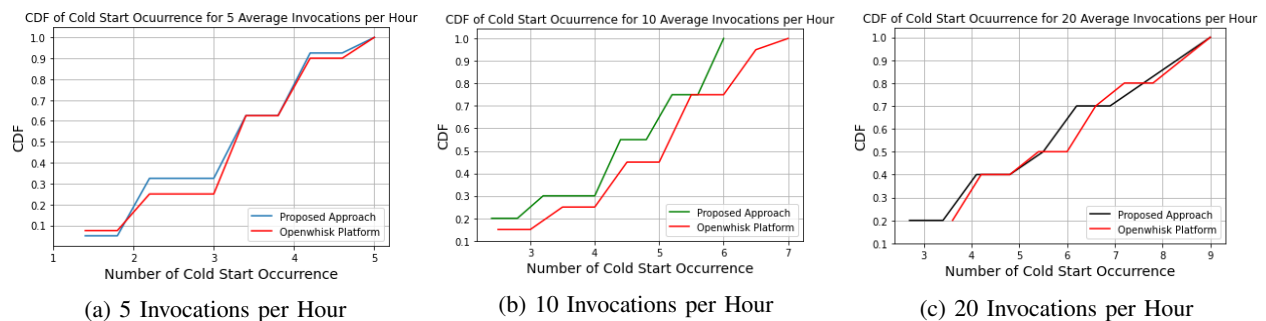


Fig. 8: CDF of Number of Cold Start Occurrence According to Proposed Approach and Openwhisk Platform.

(CDF) plot of the idle-container window determined by the proposed approach and the Openwhisk platform for three different average invocations per hour (5, 10, 20). According to Fig. 7a, 73% of invocations have less than 10 minutes idle-container window. This value is equal to 67% and 76% for average invocations per hour of 10 and 20, respectively. Therefore, most of the time a container does not need to stay warm for up to 10 minutes. In order to reduce the cold start occurrence, different values can be considered for the idle-container window depending on the different invocation patterns. It can be interpreted that the proposed approach expands the window for times when the intervals are longer and decreases it when the intervals are shorter.

– Number of Cold Start Occurrences: Fig. 8 compares the CDF of number of the cold start occurrences for the

proposed approach and the Openwhisk platform. This simulation performed for three different average invocations per hour of 5, 10 and 20. As shown in Fig. 8a, for both CDFs, 75% of invocations have less than 4 cold start occurrences within 1 hour when the average invocations per hour is 5. According to Fig. 8b, when the average invocations per hour is 10, in the proposed approach, 65% of invocations have less than 5 cold start occurrences. This value equals to 47% for the Openwhisk platform. Finally, according to Fig. 8c, for both CDFs, 42% of invocations have less than 5 cold start occurrences within 1 hour when the average invocations per hour is 20. Table II compares the number of cold start delays for the proposed approach and the Openwhisk platform for the average invocations per hour of 5, 10 and 20. As shown in Table II, the proposed approach results in a negligible improvement

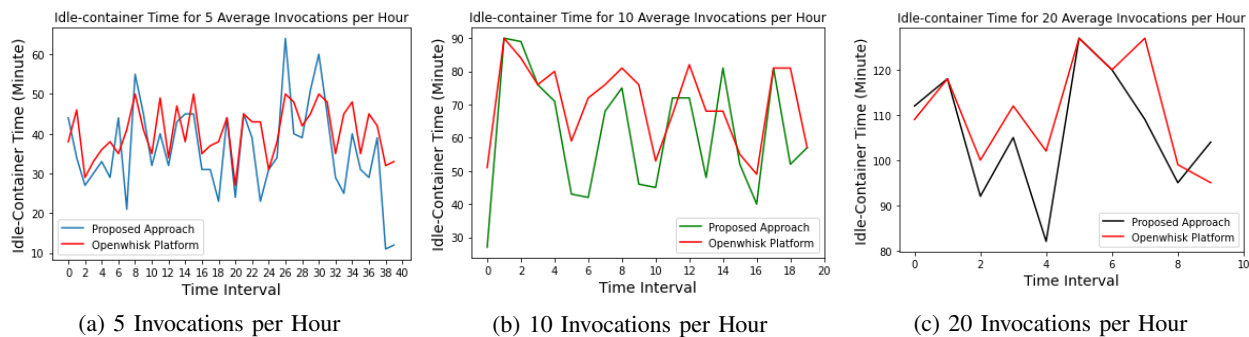


Fig. 9: Idle-Container Time According to Proposed Approach and Openwhisk Platform.

in the number of cold start occurrences. However, the proposed approach and the Openwhisk platform have a similar impact on reducing cold start occurrences. Note that this result can be concluded from Fig. 7 as well. The improvement in the number of cold start delays is low due to memory control. If we want to reduce delay more than the Openwhisk platform, we have to choose a larger idle-container window. However this is contrary to the main purpose of the paper (i.e., control memory consumption while reducing cold start delay.)

- Memory consumption: We have compared the memory consumption according to the idle-container time. That is, the length of time interval a container waits until the next invocation arrives. If the next invocation is warm, this time is equal to the idle-container window according to the proposed approach and 10 minutes according to the Openwhisk platform. Fig. 9 compares the idle-container time of the proposed approach and the Openwhisk platform for average invocations per hour of 5, 10 and 20. As shown in the Table II, our proposed approach has 11.11%, 12.73%, and 4.05% improvements in memory consumption time for average invocation per hour of 5, 10, and 20, respectively.

- **Results for Cold Start Delay Reduction Layer:** The purpose of the Cold Start Delay Reduction Layer is to determine the number of pre-warmed containers based on the number of concurrent invocations. Fig. 10 compares the number of invocations executed on the pre-warmed containers for the proposed approach and the Openwhisk platform. The test dataset has 200 invocations, and the maximum number of concurrent invocations per hour is considered. The dataset has a total of 106 concurrent invocations. Based on the results, the proposed approach was able to correctly predict 56 invocations and execute them on the pre-warmed containers. However, the Openwhisk platform, which has two pre-warmed containers at all of the time, has executed 32 requests on pre-warmed containers. Therefore, it is clear that the proposed approach has 22.65% improvement in execution of invocations on pre-warmed containers.

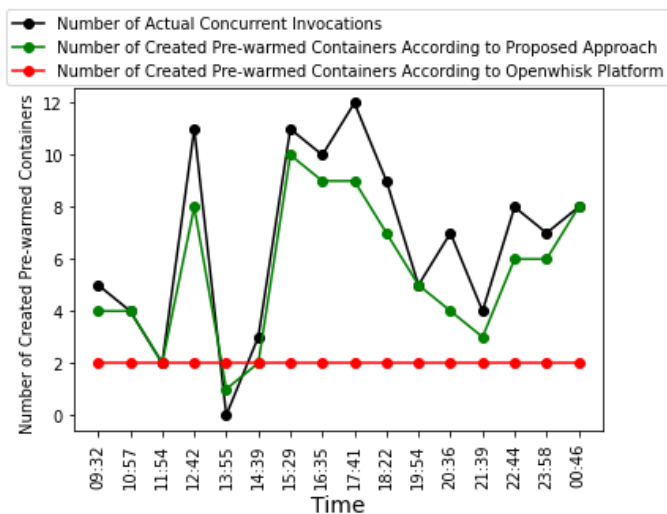


Fig. 10: Number of Invocations Executed on the Pre-warmed Containers.

VI. CONCLUSION

Cold start delay is one of the most important challenges in serverless computing. Reducing this delay in real-time applications is critical because it directly affects performance and customer satisfaction. However, the most common solutions used by serverless platforms, such as keeping the containers warm, have a fixed policy and lead to memory waste. Therefore, we need a method that determines the time required to keep the containers warm by discovering the invocation pattern of the functions. It should also balance reducing cold start latency and memory consumption. To meet these goals, in this paper, we propose a two-layer approach. In the first layer, the Reinforcement Learning algorithm is used to discover the function invocation pattern and determine the idle-container window. In the second layer, according to the prediction of the next invocation time by LSTM, the required number of pre-warmed containers is determined. Evaluations prove that it is possible to determine the amount of idle-container window by discovering the pattern of invocations. This reduces the number of cold start occurrences and controls memory consumption. In addition, the number of cold start delays that

occurred according to the proposed approach is approximately equal to the Openwhisk platform results. Moreover, the Cold Start Occurrence Reduction Layer of the proposed approach has 11.11%, 12.73%, and 4.05% improvement in memory consumption time for the average invocations per hour of 5, 10 and 20, respectively. Finally, the Cold Start Delay Reduction Layer of the proposed approach has 22.65% improvement for execution invocations on pre-warmed containers.

REFERENCES

- [1] V. Arutyunov, "Cloud computing: Its history of development, modern state, and future considerations," *Scientific and Technical Information Processing*, vol. 39, no. 3, pp. 173–178, 2012.
- [2] N. Kratzke, "A brief history of cloud application architectures," *Applied Sciences*, vol. 8, no. 8, p. 1368, 2018.
- [3] I. Baldini, P. Castro, K. Chang, P. Cheng, S. Fink, V. Ishakian, N. Mitchell, V. Muthusamy, R. Rabbah, A. Slominski *et al.*, "Serverless computing: Current trends and open problems," in *Research advances in cloud computing*. Springer, 2017, pp. 1–20.
- [4] M. Sewak and S. Singh, "Winning in the era of serverless computing and function as a service," in *2018 3rd International Conference for Convergence in Technology (I2CT)*. IEEE, 2018, pp. 1–5.
- [5] E. Van Eyk, A. Iosup, C. L. Abad, J. Grohmann, and S. Eismann, "A spec rg cloud group's vision on the performance challenges of faas cloud architectures," in *Companion of the 2018 ACM/SPEC International Conference on Performance Engineering*, 2018, pp. 21–24.
- [6] S. R. Chaudhry, A. Palade, A. Kazmi, and S. Clarke, "Improved qos at the edge using serverless computing to deploy virtual network functions," *IEEE Internet of Things Journal*, vol. 7, no. 10, pp. 10673–10683, 2020.
- [7] V. Kjorveziroski, S. Filiposka, and V. Trajkovic, "Iot serverless computing at the edge: Open issues and research direction," *Computers*, vol. 10, p. 130, 2021.
- [8] E. Van Eyk, L. Toader, S. Talluri, L. Versluis, A. Uä, and A. Iosup, "Serverless is more: From paas to present cloud computing," *IEEE Internet Computing*, vol. 22, no. 5, pp. 8–17, 2018.
- [9] A. Prez, G. Molt, M. Caballer, and A. Calatrava, "Serverless computing for container-based architectures," *Future Generation Computer Systems*, vol. 83, pp. 50–59, 2018.
- [10] A. Mohan, H. Sane, K. Doshi, S. Edupuganti, N. Nayak, and V. Sukhomlinov, "Agile cold starts for scalable serverless," in *11th {USENIX} Workshop on Hot Topics in Cloud Computing (HotCloud 19)*, 2019.
- [11] J. Manner, M. Endreß, T. Heckel, and G. Wirtz, "Cold start influencing factors in function as a service," in *2018 IEEE/ACM International Conference on Utility and Cloud Computing Companion (UCC Companion)*. IEEE, 2018, pp. 181–188.
- [12] C. Denninart and M. A. Salehi, "Efficiency in the serverless cloud computing paradigm: A survey study," *arXiv preprint arXiv:2110.06508*, 2021.
- [13] S. Eismann, J. Scheuner, E. van Eyk, M. Schwinger, J. Grohmann, N. Herbst, C. L. Abad, and A. Iosup, "Serverless applications: Why, when, and how?" *IEEE Software*, vol. 38, no. 1, pp. 32–39, 2021.
- [14] Apache, "Apache/openwhisk: Apache openwhisk is an open source serverless cloud platform." [Online]. Available: <https://github.com/apache/openwhisk>
- [15] I. E. Akkus, R. Chen, I. Rimac, M. Stein, K. Satzke, A. Beck, P. Aditya, and V. Hilt, "{SAND}: Towards high-performance serverless computing," in *2018 {Usenix} Annual Technical Conference ({USENIX}{ATC} 18)*, 2018, pp. 923–935.
- [16] E. Oakes, L. Yang, D. Zhou, K. Houck, T. Harter, A. Arpacı-Dusseau, and R. Arpacı-Dusseau, "{SOCK}: Rapid task provisioning with serverless-optimized containers," in *2018 {USENIX} Annual Technical Conference ({USENIX}{ATC} 18)*, 2018, pp. 57–70.
- [17] P. Silva, D. Fireman, and T. E. Pereira, "Prebaking functions to warm the serverless cold start," in *Proceedings of the 21st International Middleware Conference*, 2020, pp. 1–13.
- [18] Z. Xu, H. Zhang, X. Geng, Q. Wu, and H. Ma, "Adaptive function launching acceleration in serverless computing platforms," in *2019 IEEE 25th International Conference on Parallel and Distributed Systems (ICPADS)*. IEEE, 2019, pp. 9–16.
- [19] "Amazon CloudWatch - Application and Infrastructure Monitoring." [Online]. Available: <https://aws.amazon.com/cloudwatch/>
- [20] "Thundra | Monitoring for your CI Pipelines & Tests." [Online]. Available: <https://www.thundra.io>
- [21] "Serverless Observability and Intelligence Platform." [Online]. Available: <https://dashbird.io/>
- [22] "A module to optimize aws lambda function cold starts." [Online]. Available: <https://github.com/jeremydaly/lambd warmer>
- [23] J. R. Gunasekaran, P. Thinakaran, N. C. Nachiappan, M. T. Kandemir, and C. R. Das, "Fifer: Tackling resource underutilization in the serverless era," in *Proceedings of the 21st International Middleware Conference*, 2020, pp. 280–295.
- [24] S. Agarwal, M. A. Rodriguez, and R. Buyya, "A reinforcement learning approach to reduce serverless function cold start frequency," in *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. IEEE, 2021, pp. 797–803.
- [25] S. Shillaker, "A provider-friendly serverless framework for latency-critical applications," in *12th Eurosys Doctoral Workshop*, 2018, p. 71.
- [26] "Openfaas makes it simple to deploy both functions and existing code to kubernetes." [Online]. Available: <https://www.openfaas.com/>
- [27] P.-M. Lin and A. Glikson, "Mitigating cold starts in serverless platforms: A pool-based approach," *arXiv preprint arXiv:1903.12221*, 2019.
- [28] "Open source, kubernetes-native, serverless framework." [Online]. Available: <https://docs.fission.io/>
- [29] W. Lloyd, M. Vu, B. Zhang, O. David, and G. Leavesley, "Improving application migration to serverless computing platforms: Latency mitigation with keep-alive workloads," in *2018 IEEE/ACM International Conference on Utility and Cloud Computing Companion (UCC Companion)*. IEEE, 2018, pp. 195–200.
- [30] G. McGrath and P. R. Brenner, "Serverless computing: Design, implementation, and performance," in *2017 IEEE 37th International Conference on Distributed Computing Systems Workshops (ICDCSW)*. IEEE, 2017, pp. 405–410.
- [31] K. Solaiman and M. A. Adnan, "Wlec: A not so cold architecture to mitigate cold start problem in serverless computing," in *2020 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 2020, pp. 144–153.
- [32] "Aws lambda announces provisioned concurrency." [Online]. Available: <https://aws.amazon.com/about-aws/whats-new/2019/12/aws-lambda-announces-provisioned-concurrency/>
- [33] M. Shahrad, R. Fonseca, Í. Goiri, G. Chaudhry, P. Batum, J. Cooke, E. Laureano, C. Tresness, M. Russinovich, and R. Bianchini, "Serverless in the wild: Characterizing and optimizing the serverless workload at a large cloud provider," in *2020 {USENIX} Annual Technical Conference ({USENIX}{ATC} 20)*, 2020, pp. 205–218.
- [34] A. U. Gias and G. Casale, "Cocoa: Cold start aware capacity planning for function-as-a-service platforms," in *2020 28th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS)*. IEEE, 2020, pp. 1–8.
- [35] A. Banaei and M. Sharifi, "Etas: predictive scheduling of functions on worker nodes of apache openwhisk platform," *The Journal of Supercomputing*, pp. 1–36, 2021.
- [36] S. Wu, Z. Tao, H. Fan, Z. Huang, X. Zhang, H. Jin, C. Yu, and C. Cao, "Container lifecycle-aware scheduling for serverless computing," *Software: Practice and Experience*, vol. 52, no. 2, pp. 337–352, 2022.



Parichehr Vahidinia received the M.S. degree in the field of Information Technology from Shahid Beheshti University, Tehran, Iran in 2021. Her current research focuses on Serverless Computing and Artificial Intelligence.



Dr. Bahar Farahani received her Ph.D. and Post-doctoral degrees in Computer Engineering from University of Tehran, and Shahid Beheshti University, respectively. She is an assistant professor at Cyberspace Research Institute, Shahid Beheshti University. She authored several peer-reviewed Conference/Journal papers as well as book chapters on IoT, Big Data, and AI. Dr. Farahani has served as a Guest Editor of several journals, such as IEEE Internet of Things Journal (IEEE IoT-J), IEEE Transactions on Very Large Scale Integration Systems (IEEE TVLSI), IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (IEEE TCAD), Elsevier Future Generation Computer Systems (FGCS), Elsevier Microprocessors and Microsystems (MICPRO), Elsevier Journal of Network and Computer Applications (JNCA), and Elsevier Information Systems. Besides, she has also served in the Technical Program Committee (TPC) of many international conferences/workshops on AI/IoT/eHealth as well as the Technical Chair of the IEEE COINS conference.



Dr. Fereidoon Shams has received his Ph.D. in Software Engineering from the Department of Computer Science, Manchester University, UK, in 1996 and his M.S. from Sharif University of Technology, Tehran, Iran, in 1990. His major interests are Software Architecture, Enterprise Architecture, Service Oriented Architecture, Agile Methodologies, Ultra-Large-Scale (ULS) Systems and Ontological Engineering. He is currently a Professor of Software Engineering Department, Shahid Beheshti University of Iran. Also, he is heading two research groups namely SOEA (Service Oriented Enterprise Architecture Reference Lab) (soea.sbu.ac.ir) and ISA (Information Systems Architecture) (isa.sbu.ac.ir) at Shahid Beheshti University.