

Reducing idleness in financial cloud services via multi-objective evolutionary reinforcement learning based load balancer

Peng YANG^{1,2*}, Laoming ZHANG^{2,3}, Haifeng LIU⁴ & Guiying LI^{2,5}

¹Department of Statistics and Data Science, Southern University of Science and Technology, Shenzhen 518055, China;

²Guangdong Provincial Key Laboratory of Brain-Inspired Intelligent Computation,
Department of Computer Science and Engineering, Southern University of Science and Technology,
Shenzhen 518055, China;

³Academy for Advanced Interdisciplinary Studies, Southern University of Science and Technology, Shenzhen 518055, China;

⁴Guangdong OPPO Mobile Telecommunications Corp., Ltd, Shenzhen 518052, China;

⁵Research Institute of Trustworthy Autonomous Systems, Southern University of Science and Technology,
Shenzhen 518055, China

Received 27 April 2023/Revised 4 July 2023/Accepted 21 September 2023/Published online 25 January 2024

Abstract In recent years, various companies have started to shift their data services from traditional data centers to the cloud. One of the major motivations is to save on operational costs with the aid of cloud elasticity. This paper discusses an emerging need from financial services to reduce the incidence of idle servers retaining very few user connections, without disconnecting them from the server side. This paper considers this need as a bi-objective online load balancing problem. A neural network based scalable policy is designed to route user requests to varied numbers of servers for the required elasticity. An evolutionary multi-objective training framework is proposed to optimize the weights of the policy. Not only is the new objective of idleness reduced by over 130% more than traditional industrial solutions, but the original load balancing objective itself is also slightly improved. Extensive simulations with both synthetic and real-world data help reveal the detailed applicability of the proposed method to the emergent problem of reducing idleness in financial services.

Keywords evolutionary reinforcement learning, evolutionary multi-objective optimization, load balance, cloud computing

1 Introduction

Elasticity forms a cornerstone of cloud computing that has helped it succeed in the past decade [1]. In many fields, the number of concurrent users of a cloud service system may vary significantly over time, while the cloud system can reduce the idle servers elastically when the users become fewer, and vice versa [2]. The concept of idle servers refers to those servers with few user tasks, in which most of the hardware resources, e.g., the central processing unit (CPU), are running for no purpose. Despite the environmental issues, according to [3], the energy consumption cost may amount to 30%–50% of the operational cost of those large-scale data centers built by companies such as Google, Microsoft, and Facebook, which is a huge waste of money. Those costs will eventually be transferred to the tenants who pay for and use the cloud servers [4]. In financial scenarios (especially in stock trading), it is common for the number of concurrent users at the close of the market to be over 10 times smaller than that at the opening [5]. Traditional financial services focus more on stability than elasticity, and now the cloud native offers a new possibility for achieving both in the same architecture [6].

Unfortunately, no industrial solution has been available to readily reduce idle servers in the financial cloud services as it is still no trivial task. Ideally, if a server serves zero user connections, it can be shut down to reduce idleness. In practice, when the user concurrency decreases markedly, the servers that become idle may still retain a small number of user connections. Other services such as the entertainment

* Corresponding author (email: yangp@sustech.edu.cn)

and office sectors can simply disconnect all these users and re-connect them to the busy servers quickly. In this regard, these idle servers become effectively empty and thus, can be shut down to save costs. However, in financial services, the availability of the connections is extremely sensitive, as even a millisecond of disconnection may cause a heavy loss of investment and damage to the financial market [7]. This means that an idle server in the financial services sector must wait passively for all its users (even though very few) to disconnect and then can be shut down afterward, while the waiting time is a direct waste of the idle servers.

We propose that the above issue is mainly caused by the improper routing decisions of the user requests over the servers. Specifically, given that a server can only be shut down when it serves no connection, we know that the shutdown time of a server is decided by the latest disconnection time among all its user connections. Once we can route the user connections with a similar disconnection time in the same server, this server can be shut down in time and the wasted idleness can be greatly reduced. Based on this, the present paper focuses on the load balancer module within the cloud computing architecture, as it is responsible for routing user requests to different servers to construct connections.

As its name suggests, the traditional routing algorithms inside the industrial load balancers mainly focus on minimizing the load imbalance, i.e., the standard deviation of the workloads among multiple servers. The performance of those servers with heavy workloads can be very poor and will impact the whole system negatively [8]. Furthermore, overloaded servers may fail much more easily. In the case of servers with light workloads, the hardware resources are thereby wasted [9]. Thus, the load balance has been a major problem in cloud computing [10,11]. By bearing the workloads, such as user connections, the hardware resources of the servers will be consumed and, thus, considered the indicators of the load balance. For modern computing systems, four basic types of resources are the CPU and memory (random access memory, RAM) for computing, a disk for storage, and bandwidth (BW) for network communication. In this paper, we consider all of them to make the proposed method more general in terms of hardware resources. The representative routing algorithms are heuristics such as round robin, random routing, IP Hash method, the least connection, and their weighted variants [8, 12–14]. These algorithms are based on either some degree of randomness or some hardware indicators of the servers. None of these considers the behavior of users, i.e., the user connection duration. In this regard, while the users were routed to the servers for load balance, the remaining user connections at the close of the market will also be distributed randomly among the servers. This explains why, in real cases, the idle servers usually retain a few connections when the user concurrency decreases significantly. Therefore, although they have contributed greatly to enhancing the capacity (concurrent users) and the stability of the cloud services by balancing workloads, they may perform poorly on routing the user requests concerning further reducing idle servers without killing connections, as is required in financial services. Recently, intelligent routing algorithms have also been developed based on meta-heuristics and reinforcement learning (RL) [15–18]. Unfortunately, their targeted scenarios were quite different from the financial cloud and thus, did not encounter the idleness issue.

To address this issue, this paper aims to automatically learn the optimal policy for routing incoming user requests over the available servers so that the load imbalance and the idle time of all servers can be minimized simultaneously. For this purpose, when routing an incoming user request, we not only reorganize the traditional indicators of the servers along the time axis but also explicitly exploit the historical connection duration of that user as new features. Specifically, it is assumed that the connection duration of each user on different days follows a similar distribution. Thus, the connection duration of an incoming user can be predicted statistically. This assumption is reasonable because most of the users of financial services are from financial companies who usually have regular work schedules. Different degrees of the predictability of users are simulated in the experiments to support this assumption positively. Due to the online decision-making nature of routing incoming requests, the whole routing process is modeled as an RL problem with two objectives, namely, the load imbalance and the idle time of all available servers. Then, a parameter-sharing neural-network-based scalable routing policy is designed for two reasons: (1) to learn about the non-linear relationship between the statistically estimated connection duration and the states of the servers; (2) to deal with the elastic scenarios where the number of servers may change due to the autoscaling or servers' failure. Considering that it is difficult to set proper weights for aggregating the two objectives, an evolutionary multi-objective algorithm-based training framework is proposed to optimize the parameters of the neural-network-based policy. Lastly, an action mask operator is designed to help train the policy stably. As a result, a multi-objective evolutionary reinforcement learning-based load balancer (MERL-LB) is proposed.

Extensive simulations show that an MERL-LB can outperform comparable algorithms significantly on the emerging task of reducing idleness without disconnecting user connections. The diverse Pareto optimal produced by the evolutionary multi-objective training framework offers various options for the users. Among the diverse policies, the idleness objective can be reduced by over 130% in using the traditional methods while the load balancing objective itself is still slightly improved. Meanwhile, it has been shown that the evolutionary multi-objective training framework facilitates a much faster convergence rate over policy gradients [19]. The proposed scalable routing policy has also been verified successfully against a variety of servers and request loads, with the number of users ranging from 600 to 7500. Finally, an emerging sawtooth pattern is observed from the process of MERL-LB-based policies, which has been studied by comparing it to the patterns of traditional methods in detail to explain its rationality and may shed light on designing novel heuristics in the future. In a nutshell, MERL-LB is shown empirically to be a powerful solution to the emerging problem of reducing idleness and load imbalance in financial services.

The remainder of this work is as follows. Section 2 reviews the related work. Section 3 describes the problem scenarios. Section 4 first models this problem as an RL problem and then presents the scalable neural-network-based policy and its evolutionary multi-objective algorithm-based training framework. The simulations are conducted in Section 5 to verify the effectiveness of the proposed MERL-LB. The conclusions drawn from this work are discussed in the last section.

2 Related work

The load balancer works between the users' applications and the servers of the backend system. In cloud data services, the users' requests come up in an online manner. The load balancer is expected to route each incoming user's request to one of the available servers so that any single server in the data center will not be overloaded. And once the request is routed successfully, a connection is constructed between the user's application and the routed server to bi-directionally transmit data.

Industrial routing algorithms mostly make their decisions with certain degrees of randomness or based on the states of the servers to keep load balance. Typical static routing algorithms include random routing, round robin, weighted round robin, and IP Hash [20]. Round robin and weighted round robin distribute requests according to a specific probabilistic distribution among servers [14]. IP Hash maps the IP address to the corresponding server using a Hash algorithm. Static routing algorithms are generally simple to implement but lack the ability to adapt their probabilistic distributions to the changing characteristics of the network traffic. The other type of dynamic load balancing algorithms considers the servers' states as input [21]. Examples are least connection [12] and the throttled method [13]. Among them, the least connection dynamically assigns tasks based on the already served number of connections of each server. The throttled method consists of an index table of available virtual machines and their states. Requests are allocated to the virtual machine that is available and has sufficient resources [22]. However, those methods are designed for the single objective of load balancing and cannot address the problem of reducing idleness without disconnecting user connections.

There are also some intelligent load balancers. Meta-heuristics based load balancers [15] are generally derived from the behavior of natural evolution and iteratively search for the optimal solution in the policy space. Unfortunately, those methods can hardly be applied to the underlying online real-time routing scenarios, as their iterative optimization process usually takes a relatively long time to find an optimal routing solution [17]. RL-based load balancers are also warmly studied for the purpose of real-time routing. The witnessed studies mainly focus on the scenarios of software defined networks [16] and Internet of Things networks [18]. However, as the elasticity is not the major concern of those networks, those works do not encounter the same difficulty as in the financial scenarios of keeping connections alive while reducing idleness.

For cloud computing scenarios, the general load balance problems have been more intensively studied in the context of job shop scheduling [23–25] in the scheduler module in the backend. Honestly, the job scheduling problems and the connection routing problems generally follow the same vector bin-packing mathematical model [10, 26], where a set of vector items are expected to be distributed among multiple bins so that their loads can be balanced. Unfortunately, we discuss as follows that those job scheduling methods cannot readily address the underlying problem of this work. The two major differences between them emerge from practice.

(1) A computing job is basically a sequence of predefined computing steps where any intermediate data can be easily saved as checkpoints with established techniques. In this regard, killing a job on one server and restarting it on another server can lose no data with extra techniques like live migration [27]. However, a user connection does not know what data will be delivered through it at a specific time, which means a disconnection will lose data that can hardly be recovered. In a nutshell, those job scheduling methods do not explicitly keep connections alive while reducing idleness.

(2) Though some online job scheduling works also consider objectives similar to reducing idleness, like minimizing slowdown [28], makespan [29] or job completion time [30], they need the running duration of the jobs explicitly in advance. While in the connection routing problem, the online duration of any user connection is unknown when the request comes up.

To summarize, to our best knowledge, existing related works are not dedicated to solving the underlying bi-objective problem, i.e., load balancing and reducing idleness without actively disconnecting users. Thus, directly applying them may not be satisfactory.

3 Problem description

In typical financial data services, e.g., the stock market data services, the user connections usually follow a quite regular mode. First, the number of user connections increases very rapidly at the market opening in the morning, and then decreases gradually when the o'clock is approaching the market closing in the afternoon. Second, the duration of different user connections may vary significantly. Consequently, as shown in Figure 1, the users usually disconnect from the data services at different times, leaving the servers running with a very low utilization rate but cannot be shutdown due to few alive connections on them, e.g., the duration from t_2 to t_3 . This economic waste resulting from the server idleness can be significantly scaled up as the financial data services usually run on cloud computing infrastructures with large-scale servers.

Fortunately, the duration of each user connection in different days often varies slightly as most users are staff from financial industries who often have regular schedules [31]. This simply suggests that the disconnection time of the same user on different days can be predictive to some extent and maybe statistically modeled following a probabilistic distribution. With this special feature in the financial services, this work describes the problem of routing user connections to reduce idleness as follows.

As shown in Figure 2, each incoming user request has five types of features, including the demand of CPU, memory (RAM), storage (hard disk drive, HDD), BW and the historical distribution of the connection duration of this user. The first four features are hardware requirements indicating how many resources the system needs to serve that connection and its related data services. The last feature emerges from the financial services and is, for the first time, considered in the load balancer. Each time a new request arrives, it first enters the ready queue and waits to be routed to one of the available servers by the routing policy. Once the user request is successfully routed to a server, a user connection between the user and the server can be constructed, and data services can be provided to the user through the connection. Otherwise, those requests that cannot be routed due to insufficient hardware resources will be tentatively placed in the block queue until the next round of routing. The routing policy works by optimally distributing the users' requests to minimize the load imbalance of 4 types of resource utilization among servers while simultaneously reducing the idle time of the servers.

In this paper, the above mentioned two objective functions are defined as follows. Suppose the load balancer continues serving for a period of T timesteps. For the objective of load balancing, it is defined as the standard deviation of the utilization among 4 resource types:

$$F_{\text{balance}} = \frac{1}{4 \cdot T} \sum_{t=0}^T \sum_{r=1}^4 \sqrt{\frac{\sum_{i=1}^N (x_{ri}^t - \mu_r^t)^2}{N}}, \quad (1)$$

where x_{ri}^t indicates the utilization of r -th resource type of the i -th server at the t -th timestep, and μ_r^t denotes the average of the r -th type resource utilization among all N servers at the t -th timestep. For the objective of reducing idle time, instead of using the commonly-seen makespan of all the servers, we calculate the average remaining duration of the servers over all timesteps. The reason is that we want to

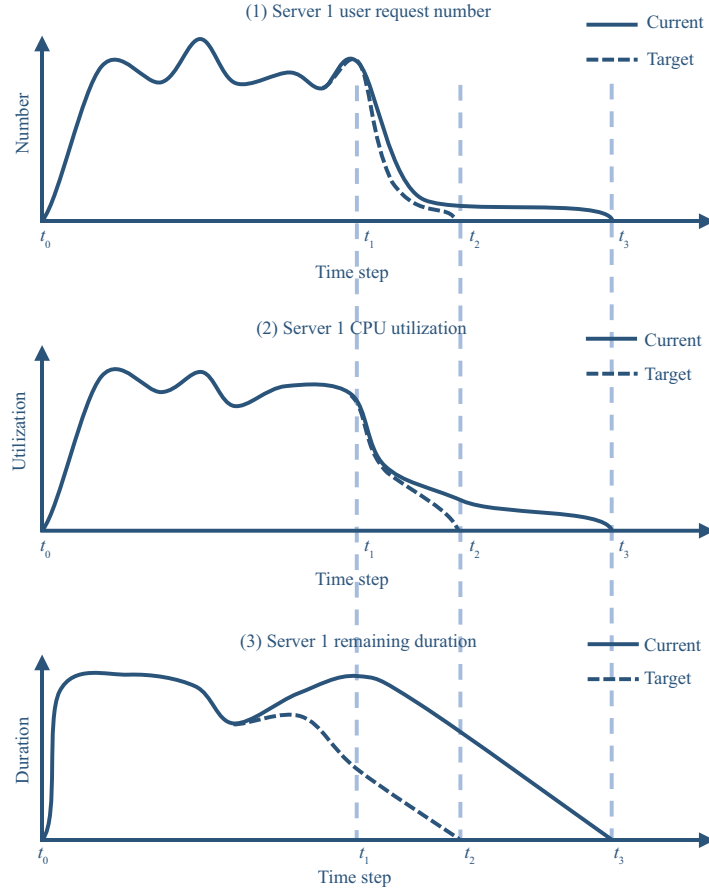


Figure 1 (Color online) A typical phenomenon of financial services that there exist very few connections from t_2 to t_3 , which are targeted to be reduced in this work.

achieve idle reduction at any time, not just at the last moment.

$$F_{\text{idle}} = \frac{1}{T \cdot N} \sum_{t=0}^T \sum_{i=1}^N d_i^t, \quad (2)$$

where d_i^t is maximal remaining duration among all connections of the i -th server at the t -th timestep. Note that, the remaining duration of any connection cannot be known exactly before its disconnection. On the other hand, as discussed previously, according to the observation that each user's connection duration on different official days changes slightly due to the regular schedule, we can predict the duration of each user at the current t -th timestep as the average of the user's historical duration. This paper also empirically assesses the different impacts on the load balancer caused by different degrees of predictability. Also note that, in the whole problem description, the system is not allowed to actively disconnect the user connections. Only passive disconnection from the user's application is permitted.

4 Method

In this section, the RL-based problem formulation is first described, including the representation of actions, the RL states, and the rewards. Then, a neural network is designed to learn the non-linear relationship between the statistically estimated connection duration and the available servers. The architecture of this neural network is also designed for elastic scenarios with changing numbers of servers. After that, an evolutionary multi-objective algorithms based training framework is introduced for this scalable policy, in terms of both minimal load imbalance and server idleness. The action mask operator for helping stably train the policy is detailed lastly.




Figure 2 (Color online) The routing process of load balancer in financial services.




Figure 3 (Color online) This work considers two types of states for RL. (a) 5-dimensional states of each user request; (b) $41N$ -dimensional server states include N servers with 4 types of resources looking forward 10 timesteps into the future, and 1 maximal remaining duration.

4.1 RL based formulation

In this RL model, each of the user's requests in the ready queue is represented together with the utilization of all servers as the RL states. Based on the states, the policy outputs the action of routing the request to one of the servers to construct a user connection.

As shown in Figure 3, the RL states are a conjunction of both the states of one user's request and the states of N servers. The state of each request is represented as a 5-dimensional vector, including the demands of 4 types of hardware resources and 1 scalar value of the predicted duration, i.e., the average of the historical duration of that user. The state of each server is represented as a $(4h + 1)$ -dimensional vector, i.e., the utilization of 4 types of resources along h looking forward timesteps and 1 scalar value of the maximal remaining duration among all connections in the future h -th timestep. Here the remaining duration of each connection is calculated as its predicted duration minus h timesteps, and the utilization of each resource in some future timestep is calculated as the summation of all alive connections at that moment. By using this feature along h looking forward timesteps, the future states of the servers are involved to help the policy learn the optimal distribution of the requests for reducing idleness in the future. In the empirical studies, we simply set h to 10. As a result, the total length of the RL states




Figure 4 (Color online) Inference process of the scalable policy network.

vector is $5 + 41N$.

The RL model has $N + 1$ actions, i.e., routing the request to one of the N servers or the blocking queue if no proper server is available. For the reward function, as we employ the evolutionary multi-objective algorithm to train the policy, it is unnecessary to derive the intermediate rewards of each routing action. Instead, the two objectives in (1) and (2) are used as the accumulative rewards for each policy after the whole simulation process [32].

4.2 The scalable policy

To learn the optimal distribution of connections over servers, the neural network based policy is preferred due to its strong learning ability. The input of the neural network will be the above described $(5 + 41N)$ -dimensional RL states, and the output of the neural network will be the $(N + 1)$ -dimensional RL actions. Therefore, the sizes of the input and output of the policy change with the number of available servers. Additionally, due to the elasticity of cloud computing, the number of available servers usually changes over time for a given data service. Besides, it is also important to generalize the trained policy to different data services that may be supported by different numbers of servers. For these purposes, the employed policy should be able to scale with the number of servers automatically. To this end, this paper proposes a parameter-sharing based neural network architecture for the underlying routing tasks.

As shown in Figure 4, the policy network achieves the scalability with the global state management and the sharing parameters among N copies, each of which corresponds to one available server. The input of this policy network is divided into three components: the request state, the individual server state, and the global state. The global state refers to the mean and the standard deviation among all individual server states, either of which is basically a 40-dimensional vector with 4 resources looking forward for 10 timesteps and 1 remaining duration. There are N fully connected networks jointly for the decision making. All networks share the same weight parameters and can receive the shared global state of servers as well as the 5-dimensional request state. In addition, each i -th network will individually receive the 41-dimensional states of the i -th server. For the decision making, each i -th network outputs a score. A higher score indicates better load balance and predicted idle time reduction. At last, the action with the highest score will be finally selected and the corresponding server will be chosen to construct the connection. By this architecture, the whole policy is able to scale with the number of servers and can intuitively be generalized to different data services. The update of the parameter-sharing networks is simple. Each network targets a server. No matter how main servers there are, their corresponding networks share the same weights. And the weights are represented as an individual in the evolutionary multi-objective training framework.

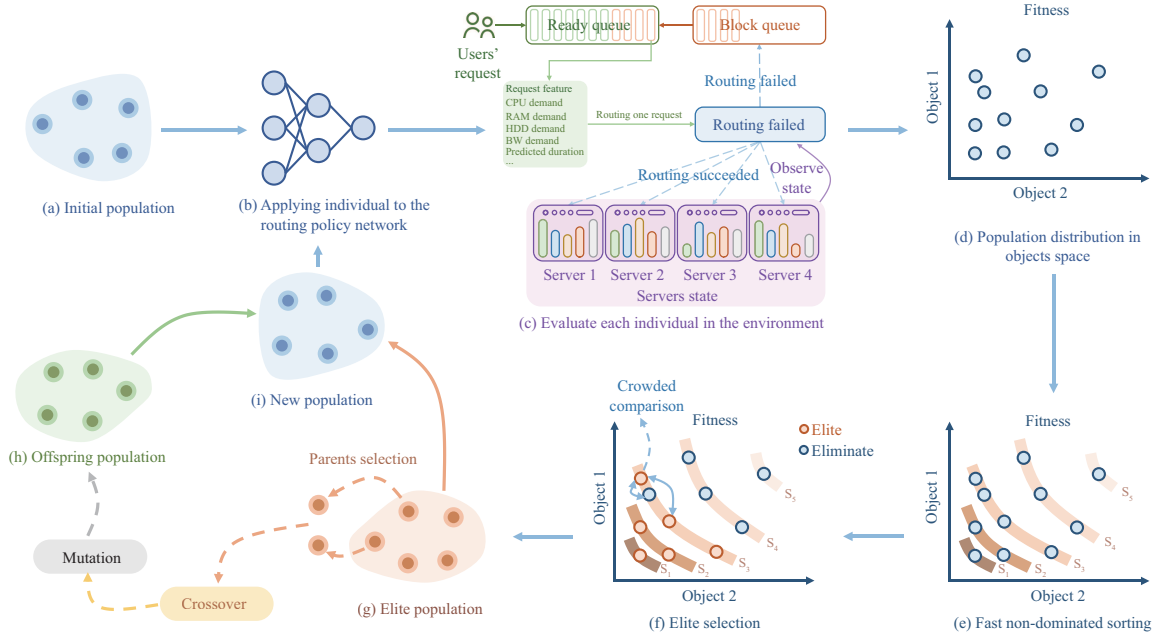


Figure 5 (Color online) Overview of the evolutionary multi-objective training framework.

4.3 Evolutionary multi-objective training framework

As discussed above, the underlying load balancer has two objectives (see (1) and (2)), i.e., the long term rewards. In traditional RL policy training, it is common to aggregate these two rewards into a single one by weights sum. Unfortunately, due to the different scales of those two objectives, it is quite difficult to set proper weights, and the two objectives might not easily be satisfied.

In recent years, evolutionary algorithms based on RL policy training methods have been warmly studied, and form a new topic called evolutionary RL [33–35]. This type of new methods enjoys several advantages over the policy-gradient methods which use gradient descent to optimize the policy. Apart from the ability of global search [36] and learning from sparse reward [37], one major benefit of evolutionary RL is the flexibility of incorporating the evolutionary multi-objective framework [38] to train the policy. Inspired by this progress, we propose to employ the evolutionary multi-objective algorithms [39,40] to optimize the weights of the policy directly for three reasons. First, evolutionary multi-objective algorithms have been the mainstream methods for multi-objective optimization problems and have successfully shown their power in many real-world applications [41–44]. Second, it does not need to set the aggregation weights manually but directly compare solutions with the domination relation. Third, it suffices to output a set of optimal solutions called the non-dominated Pareto set, which is equivalent to the optima with different weights sum aggregation [45,46], thus offering the users the flexibility to select the policy online. Here we employ the well-established NSGA-II, i.e., almost the most successful evolutionary multi-objective algorithm developed by Deb et al [47]. in 2002, to construct the evolutionary multi-objective training framework for the policy. For more details on NSGA-II, please refer to [47]. It is worth to mention that several advanced evolutionary multi-objective RL methods have been proposed recently [48–52]. Though they are not developed for the underlying routing problem, their algorithmic insights can shed light on the further improvement of training the proposed policy network.

The proposed NSGA-II based training framework shown in Figure 5 is an iterative process. At the beginning of the training (step (a)), a population of the routing policy network is randomly initialized. Second, we evaluate each individual policy in the simulation environment (steps (b) and (c)). Here, the objective functions in (1) and (2) are directly used for evaluation, as evolutionary algorithms usually measure the accumulated performance of a candidate policy over the whole simulation period [37]. It is important to note that the settings of the simulation environment, e.g., the requests sequence and the servers' initial states, are identical for the evaluation of all individual policies. After the evaluation, we obtain the bi-objective fitness values for all individuals in the current population, which distributes in a 2-dimensional objective space (step (d)). In this space, the closer to the origin, the better the individual will be.




Figure 6 (Color online) Process of generating offspring.

Then, the individuals in the current population are sorted by the fast non-dominated sorting of NSGA-II, and L non-dominated solution sets S_1, S_2, \dots, S_L are obtained (step (e)). Here, a non-dominated solution set means that an individual within the solution set cannot dominate any other individuals within the same solution set, but is dominated by the sets prior to it, e.g., S_1 dominates S_2, \dots, S_L . More specifically, an individual A dominates an individual B means that A is not inferior to B in all objectives and is superior to B in at least one objective. The elite individuals are selected from the current population based on the precedence of the non-dominated solution sets and the crowding comparison (step (f)), where the crowding comparison algorithm is used to avoid selecting similar individuals and thus ensure the diversity of the population. Finally, based on the elite population, the offspring population is generated with typical operators of evolutionary algorithms (step (h)) and forms. It forms the new population for the next iteration together with the elite population (steps (g)–(i)). And then the iteration continues until some stop-criteria is met, and the individuals in S_1 are output as the trained policies.

In the above framework, each policy is represented as a vector of weight parameters of the network described in Subsection 4.2. The offspring generation step is to generate two new vectors based on two parent vectors. As shown in Figure 6, the two parent vectors crossover at a uniformly randomly chosen point to produce two tentative vectors. Then, based on the two tentative vectors, two offspring vectors are generated by the commonly used Gaussian mutation operator. Specifically, for each i -th weight parameter of the policy network, with a probability of γ , it will be mutated by sampling from the Gaussian distribution $N(\theta_i, \beta)$, where θ_i denotes the i -th weight parameter of a tentative vector, and β denotes the magnitude of the mutation.

4.4 Actions mask operator

Note that, while simulating the individual policies within step (c) of Figure 5, the policies will not be updated unless the simulation ends. This might be a major difference between evolutionary RL and traditional policy gradient based RL. Though it has been revealed that this evolutionary RL training paradigm has merits for those scarce reward problems [53], it indeed causes an issue in the load balancer. That is, the policy may be trapped in the situation of continuously routing the incoming requests to the same server, making that server fully utilized quickly. As a result, the subsequent requests cannot be connected to that server and thus have to be put into the block queue. Once the block queue overflows, the simulation will be stopped, and the evaluated objective fitness of that policy will be very poor.

As the initial policies are randomly generated, they are highly likely to face this issue. And the training may easily fail. To stabilize the training process, a simple action mask operator is proposed. That is, after N sharing networks in a policy outputting their scores, those servers that cannot satisfy the hardware demand of the current request will be masked and the score of the corresponding sharing networks will be set to 0. After that, the server with the highest score will be selected. By doing so, the above issue has been addressed and the action mask operator can be viewed as a specific kind of diversity enhancement.

5 Experiments

The experimental studies are designed to answer the following research questions.

(1) How does MERL-LB perform on the task of reducing idleness without disconnecting user connections, in comparisons with traditional methods?

Table 1 Configurations of the simulated virtual data center

Parameter	Value
time_step	12 (s)
server_num	10
cpu_capacity	500 (units)
ram_capacity	500 (units)
hdd_capacity	500 (units)
bw_capacity	500 (units)
ready_queue_size	200
block_queue_size	200
predicted_range	120 (min)
future_sample	10

Table 2 Simulation data generation parameter setting

Parameter	Value
data_time	120 (min)
time_step	12 (s)
mean_req_num	3
min_res_req	0 (units)
max_res_req	10 (units)
min_user_duration	1 (min)
max_user_duration	120 (min)

(2) How does the proposed evolutionary multi-objective framework facilitate the training of RL-based routing policy, in terms of the convergence rate and the diversified user options?

(3) How is the stability of MERL-LB against different loads?

(4) How is the scalability of the proposed parameter-sharing routing policy against different numbers of user requests and servers?

(5) How does MERL-LB tolerate to the different variances of the user connection time?

5.1 Simulation environment

The experiment simulates a commonly encountered scenario of financial data services. The simulation starts at the opening of the stock market and ends after all users disconnect from the servers naturally, which will be later than the close of the stock market. For the major experiments, we simulate around 1500 user requests coming up gradually to 10 servers. The coming up duration simulates between the opening and close time of the stock market, which is set to 2 h. Later, for the stability and scalability tests, the number of user requests varied up to 7500, and the number of available servers increased from 10 to 50.

All simulations are conducted on a discrete event simulation based virtual data center implemented with Python¹⁾. The configurations of the virtual data center are shown in Table 1. The number of servers in the data center is set to 10; each server has four resources (CPU, RAM, HDD, and BW), and the capacity of each resource is set to 500 units; at the beginning of the simulation, all the resource utilization is initialized to zero. The queue sizes of the ready queue and block queue in the data center are fixed to 200 connections, respectively. The event of the simulation process is triggered five times every minute. That is, the minimum simulation interval is set to 12 s. The predicted range for the future states of the data center is 120 min, among which $h = 10$ timesteps are sampled to represent the prediction of the system's future state. The parameters setting for generating synthetic data is listed in Table 2.

5.2 Compared algorithms

Four classic rule-based heuristic routing methods and two RL based algorithms were used for comparisons. As discussed in Section 2, three industrial algorithms, like random routing, round robin, and the least connection are selected as they merely consider load balancing. Another algorithm named least duration gap is designed based on the least connection to only consider the reduction of idleness, ignoring the load

1) <https://github.com/zlaom/MERL-LB>.

balancing. The core idea of the least duration gap is to greedily route the user requests with similar predicted remaining duration to the same server with the following rule:

$$a = \arg \min_i (|d_s^i - d_r|), \quad (3)$$

where d_s^i indicates the maximum predicted remaining duration of the i -th server, while d_r is the predicted duration of the request.

For the two RL based methods, they are designed to show the effectiveness of the proposed evolutionary multi-objective training framework. For this purpose, the proximal policy optimization (PPO) [19] and the independent-input policy gradient (IPG) [54] are employed to train the proposed parameter-sharing policy. The former is a well-established RL method that has been successfully applied in many fields. The latter is a variant of the original IPG method for job scheduling problems that are similar to the traditional load balancing problem. Both of them follow the policy gradient training framework. PPO calculates the gradient based on the temporal-difference error by an incomplete sequence of interactions, while IPG calculates the gradient with a complete sequence of interactions. The immediate rewards of a routing action at the specific t -th timestep can be simply derived from (1) and (2) by eliminating the summation over t , as shown in (4) and (5). And the two rewards are linearly aggregated with fixed weights to form the training signal, as shown in (6).

$$\text{reward}_{\text{balance}}^t = \frac{1}{4} \sum_{r=1}^4 \sqrt{\frac{\sum_{i=1}^N (x_{ri}^t - \mu_r^t)^2}{N}}, \quad (4)$$

$$\text{reward}_{\text{idle}}^t = \frac{1}{N} \sum_{i=1}^N d_i^t, \quad (5)$$

$$\text{reward}^t = w_1 \cdot \text{reward}_{\text{balance}}^t + w_2 \cdot \text{reward}_{\text{idle}}^t. \quad (6)$$

5.3 Implementation details

The architecture of the parameter-sharing neural network is introduced as follows. The network is basically a multi-layer perceptron with a 126-dimensional input layers, a 32-dimensional hidden layers, and a 1-dimensional output layer. The Relu is used as the activation function of the hidden layer. The 126-dimensional input refers to the 5-dimensional resource state of each incoming request, the 41-dimensional individual server state of the corresponding server, and the 80-dimensional global state with half for the average and half for the standard deviation of all individual servers on the 4 types of hardware indicators along 10 steps forward. The 1-dimensional output is the score of routing the incoming request to the corresponding server. The network with the highest output score among all N networks is found, and the corresponding server is chosen to be routed. The number of weight parameters of the network is 4097, and all N networks share the same weights²⁾. All three RL-based methods train this network through the experiments.

The major hyperparameters of the PPO-LB and IPG-LB are described as follows. For the major simulation settings, both PPO-LB and IPG-LB share the same parameters. The maximum number of training simulations is set to 750000 for both algorithms, which is kept the same for MERL-LB. The learning rate of the network is set to 0.001, the discount factor γ is 0.99, and the aggregation weights of the load balancing reward and the idle time reward are 0.5, respectively, for both algorithms. The network of PPO-LB is updated every 256 steps of interaction. Each interaction sequence is reused for 5 epochs. The clipping threshold ϵ of PPO-LB is set to 0.2. The weight α_1 of the loss of the value network in the loss function is 0.5, and the weight α_2 of the entropy term coefficient is 0.01. The batch size during training is set to 512 for IPG-LB, and each data instance randomly samples 10 interaction sequences using the same policy. In general, those parameters are all suggested by the original paper of PPO [19] and IPG [54]. The 4 rule-based heuristics also do not involve hyperparameters.

For the major hyperparameters of the proposed MERL-LB, we set the population size to 50, the number of elites selected from the population per iteration to 25, and the number of offspring generated by elites to 25. The probability of mutation on each variable is 0.25, and the magnitude of mutation is 0.05. The time budget of the simulation is the same as PPO-LB and IPG-LB for fairness.

2) The codes can be found <https://github.com/zlaom/MERL-LB>.

Table 3 Performance of different algorithms for two objectives on a validation set of 10 servers with a 75% load

Method	F_{balance} (units) ↓	F_{idle} (min) ↓
Random	30.15 ± 2.55	106.26 ± 0.83
Round robin	22.15 ± 1.96	107.12 ± 0.85
Least connection	16.27 ± 1.04	106.98 ± 0.95
Least duration gap	66.89 ± 6.12	69.91 ± 1.56
IPG-LB	17.05 ± 1.12	79.75 ± 1.41
PPO-LB	16.23 ± 1.14	80.67 ± 1.34
MERL-LB-1	15.70 ± 1.17	88.68 ± 1.56
MERL-LB-4	16.88 ± 1.06	76.75 ± 1.25
MERL-LB-6	23.42 ± 1.50	72.44 ± 1.26
MERL-LB-11	66.58 ± 4.93	61.63 ± 1.26
MERL-LB-25	178.77 ± 4.88	46.48 ± 1.73

5.4 Overall comparisons

Table 3 reports the overall comparison results between MERL-LB and the compared algorithms, and is divided into three parts. The top part shows the performance of the heuristic-based algorithms, the middle part is about the two policy gradient based RL load balancers, and the bottom part shows representative solutions generated by the proposed MERL-LB algorithm, where the suffix indicates different individuals in the last population.

First of all, the three RL-based methods generally perform better than the heuristic-based methods, especially on the load balancing objective. Furthermore, MERL-LB has been able to find the best solution for either objective and also offers different options for various balances between the two objectives.

It can be seen that MERL-LB-1 focuses more on the load balancing objective and performs the best among all algorithms. Compared with the least connection method, which merely targets such an objective and performs the best among the compared heuristics, MERL-LB-1 reduces the idle time by 17.11% while achieving better load balancing performance. MERL-LB-25 focuses more on the idle time objective and can reduce up to 33.51% over the least duration gap method, which is specifically designed for reducing idleness.

The other 3 solutions of MERL-LB successfully show how evolutionary multi-objective optimization can offer users different options. MERL-LB-4 not only achieves a better load balancing objective than IPG-LB but reduces the idle time by 3.76%. MERL-LB-6 achieves similar idleness to the least duration gap method and can improve the load balancing objective by 64.99%. MERL-LB-11 performs similarly to the least duration gap in terms of the load balance while successfully reducing the idleness by 11.87%. To summarize, the Pareto optima output by MERL-LB can effectively balance the two objectives automatically.

Figure 7 shows the distribution of the optimal policies generated by different algorithms in the objective space. The figure clearly shows that compared with PPO-LB and IPG-LB, MERL-LB can generate a set of diverse policies that are non-dominated by each other on both objectives. That means any one of the policies is a good choice for a specific need. Looking at the error bars of the policies in both directions (i.e., both objectives) in the figure shows that the policies generated by MERL-LB are mostly stable, especially compared with the least duration gap method. Also note that MERL-LB generalizes well as the testing instances are different from the training instances, while the Pareto policies trained by MERL-LB still distribute well on the objective space of the testing instances.

To further demonstrate the effectiveness of the proposed evolutionary multi-objective training framework, Figure 8 depicts how the three RL-based methods converge during the first 300000 simulations in the training process. There are two parts of the training overhead, i.e., the algorithmic part and the simulation part. We repeat the training processes of IPG-LB, PPO-LB, and MERL-LB for 50 runs and the average of the computational time of each algorithm is recorded. It shows that the total computational costs at each iteration (one complete simulation) of IPG-LB, PPO-LB, and MERL-LB are 14.64, 9.3498, and 11.3127 s, indicating that MERL-LB is competitive in terms of the wall-clock training time³⁾. Furthermore, it has been found that the main overhead of training lies in the simulation process, as it costs 9.3426 s to simulate 1 thousand users. Therefore, we mainly use the number of simulations to measure

3) Experiments run on a workstation with 2 Intel(R) Xeon(R) Gold 6240 CPU (2.60 GHz), in total with 72 threads, 377 GB memory, and 4 TITANX GPUs.

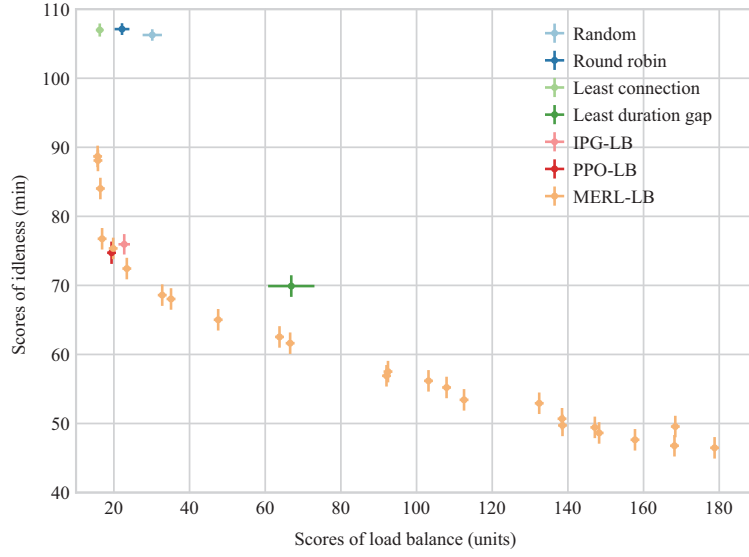


Figure 7 (Color online) Performance of different algorithms on two objectives. The cross-like error bars indicate the standard deviation on both objectives.




Figure 8 (Color online) The convergence curves of RL-based load balancers on two objectives in the training process.

the training cost of the algorithm in this paper. Each curve is composed of the scores over the training simulation. For each score, it is an average of two min-max normalized objective values. In short, all three methods have converged, and MERL-LB achieves a better convergence rate. One major reason is that, as the optimization process goes on, the conflict between two objectives becomes more significant. This makes PPO-LB and IPG-LB more difficult to express the non-dominated relationship between solutions with the weighted sum method. The simpler objective may impact the weighted objective more heavily and thus lead to significant bias in the search. As a result, PPO-LB and IPG-LB bias to the load balance objective and eventually converge slower with the indicator of the normalized average score in Figure 8.

5.5 Impacts on different numbers of user requests

Figure 9 shows how different numbers of user requests may impact the performance of the algorithms on the two objectives. Through this group of experiments, the number of available servers is fixed to 10. The number of user requests is simulated from 600 to 2100 by increasing the mean value of the employed Poisson distribution. With a fixed 2 h of the simulated opening of the market, not only does the total number of user requests increase but the concurrently incoming user requests at each time interval are enlarged, which leads to the increased loads of servers. As a result, this group of experiments assesses the algorithms for different numbers of user requests and different loads of servers in the same way.

For the load balancing objective, the performance of random, round robin, the least connection, PPO-




Figure 9 (Color online) Performance of algorithms on two objectives under different loads, i.e., different numbers of user requests. (a) Load balance objective; (b) idleness objective.

LB, and MERL-LB-6 gradually deteriorates as the load increases. This is mainly because higher resource utilization rates may lead to a greater load imbalance among servers, as the standard deviation among servers can be larger. With the increase of loads, the load balancing performance of LDG, MERL-LB-11, and MERL-LB-25 all first deteriorated and then gradually improved. This is mainly because those algorithms all pay more attention to reducing idleness. When the load is at a low level, it leads to a greater imbalance among servers. However, when the load increases, even servers with lower loads will be assigned more requests, which will gradually reduce the differences in resource utilization rates among servers.

For the idleness objective, as the load increases, the idle time of all algorithms increases accordingly. This is mainly because an increase in the load means an increase in the number of long connection requests, which would need more servers to serve the connections, while the number of servers is fixed at 10.

In general, MERL-LB-1 performs similarly to the least connection method at all loads and is able to reduce the idle time by 10%–20%. MERL-LB-6 achieves a similar idleness objective value to LDG at 75%–105% loads while it improves the load balancing objective by about 50%–60%. Compared to PPO-LB and IPG-LB, MERL-LB can generate diverse policies and consistently outperform the two at different loads. Therefore, MERL-LB performs more robustly overall when the number of requests or loads varies.

5.6 Impacts on different numbers of available servers

Figure 10 shows the performance of the algorithms on two objectives under different numbers of servers, while the average load is fixed at 75%. Thus, by varying the number of servers from 10 to 50, the number of user requests also changes proportionally, up to around 7500 requests in total for the case of 50 servers. For the load balancing objective, it can be observed that as the number of servers increases, the load balance of each algorithm does not deteriorate too much, except for the least duration gap method, which is the only algorithm that never considers the load balance objective. Similarly, for the idleness objective, the performance of the least duration gap method improves most significantly as the number of servers increased, though it still cannot outperform MERL-LB-25 in all tested cases. The other algorithms perform rather stably, regardless of the changed number of servers. Overall, this suggests that




Figure 10 (Color online) Performance of algorithms on two objectives with different numbers of servers. (a) Load balance objective; (b) idleness objective.

the proposed scalable policy network can be flexibly applied to scenarios with different numbers of servers and can maintain stable routing capabilities on the two objectives.

5.7 Impacts on the predictability of user duration

This work assumes that we can predict the duration of a user by statistically analyzing their connection time over a recent period (for example, the average value of the user's recent connection duration). In real cases, the actual duration of users may vary over time, and is highly unlikely identical to the predicted value. Thus, this group of experiments aims to assess how different degrees of the predictability may impact the performance of the algorithms, reflecting how effective MERL-LB can behave in a real-world noisy environment.

For this purpose, we generate new testing data sequences based on the original 50 testing instances by adding a random noise to each user's connection time. More specifically, the random noise follows a Gaussian distribution $N(\mu, \sigma)$, where μ is the duration of the original data, and σ represents the standard deviation of the noise, resulting in different predictability of the user connection time. The random noise is truncated at a 3σ level.

Figure 11 shows the impacts of different user duration predictability on the algorithms for two objectives respectively. In general, for the three heuristics that do not consider the idleness objective, their performance does not improve when the predictability increases. For the least duration gap method, which focuses on the idleness objective but does not rely on the prediction of user connections, its performance on the idleness objective improves near 20 min. For the other RL-based policies that share the same scalable routing network, the prediction of the user connections has been involved in their training processes. As can be seen, the various predictability does not deteriorate their performance unexpectedly. That is, as the predictability increases, their performance on the idleness objective becomes better similar to the least duration gap method. This shows that with up to a standard deviation of 30 min error on the prediction of the user connection time (maximally 120 min), the proposed scalable routing network works quite stable.

On the other hand, the policies that tend to focus more on the idleness objective deteriorate on the load balance objective as the predictability increases. The reason is that as the predictability of the user connection increases, the randomness of the data as well as the routing decreases. And the randomness usually forms a source of load balancing, as suggested by the Random method.




Figure 11 (Color online) Performance of algorithms on two objectives under different user predictability. (a) Load balance objective; (b) idleness objective.

5.8 Understanding the trade-off between two objectives

To understand how the load balancing objective and the idleness objective are balanced during the learning of MERL-LB, we show how the two objectives change over time for the least duration gap method, the least connection method, and MERL-LB-6. For the least duration gap method, only the idleness objective is considered, while for the least connection method, only the load balancing objective is considered. MERL-LB-6 is selected here as it balances the two objectives well. We depict the curves of the idle time and the CPU resource utilization in Figure 12. For clarity, the simulated opening period is enlarged to a full range of 6 h to make the observation easier. As shown in Figures 12(a)–(c), each curve in the figure represents the change of the remaining duration of each server per minute (known from the ground truth of the testing data). In Figures 12(d)–(f), each curve in the figure represents the change of CPU resource utilization of each server per minute.

During the opening period, the number of users continues to come up, so the CPU utilization of the 10 servers rises up and fluctuates, and becomes stable after the number of concurrent alive users is stabilized. Then, after the close of the market, no new user request comes up, and the number of connections as well as the CPU utilization, continues to decrease. The closer the utilization rate among servers per minute, the better the load balance is. In Figure 12(e), each server's CPU resource utilization of MERL-LB-6 fluctuates about half of that in Figure 12(d), so MERL-LB-6 has better load balancing performance than the least duration gap. In Figure 12(f), the least connection's CPU utilization fluctuates even smaller than MERL-LB-6, so its load balancing performance is better correspondingly.

An interesting phenomenon can be observed by comparing Figures 12(a)–(c) that the least duration gap and MERL-LB-6 minimize the idleness very differently. In Figure 12(a), the curve of the remaining duration of each server is completely separated and even shows a layered distribution. In this regard, the least duration gap actually minimizes the idleness by clustering each incoming user request with all the alive connections based on their predicted remaining durations. On the other hand, as shown in Figure 12(b), the curves of MERL-LB-6 overlap with each other and show a sawtooth pattern. This pattern may trade off the load balance better and contribute to better idleness as all curves drop to 0 in the shortest expectation time. In either above-mentioned case, the spectrum of the curves is much larger than that of the least connection, indicating that both the least duration gap and MERL-LB-6 have a better trade-off on both objectives.

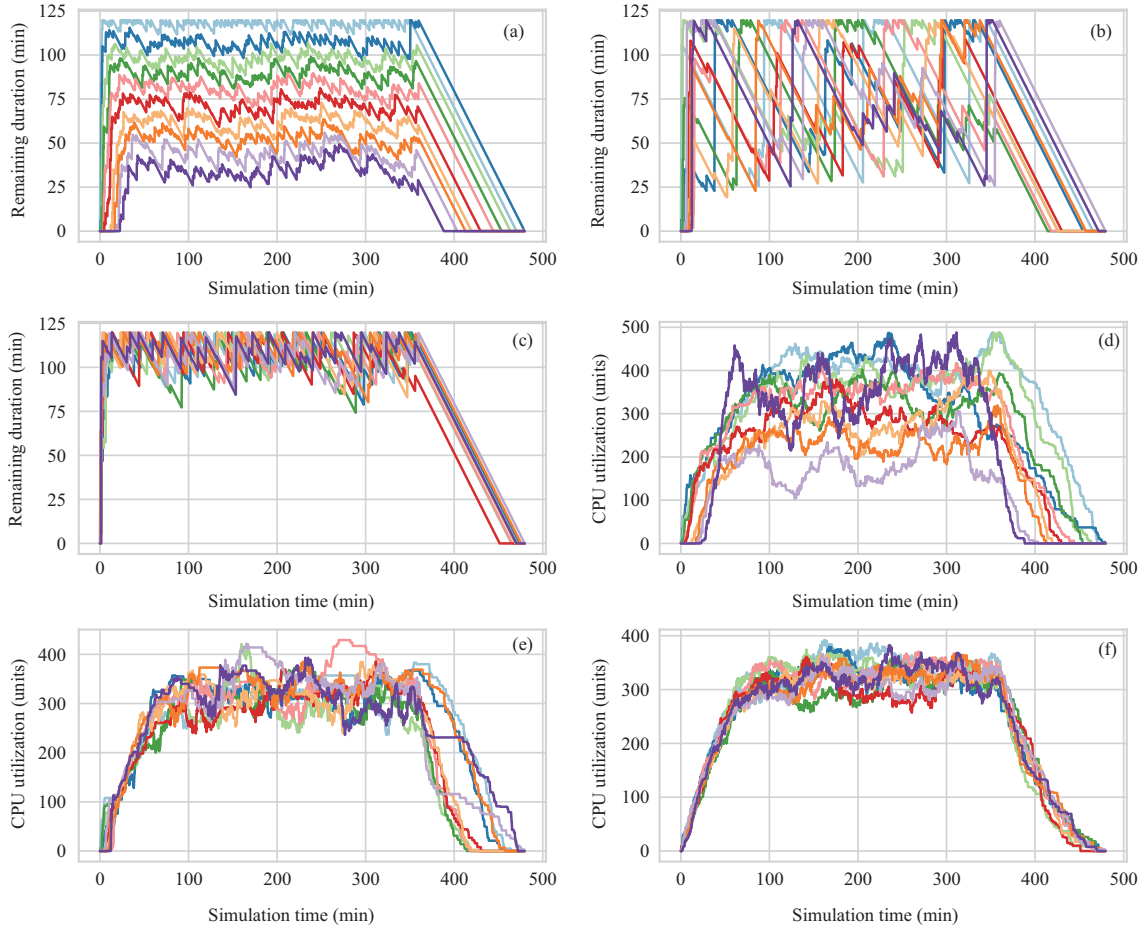


Figure 12 (Color online) Comparison of least duration gap, MERL-LB-6, and least connection characteristics on 10 servers. The remaining duration of each server in (a) the least duration gap, (b) MERL-LB-6, and (c) the least connection. The CPU utilization of each server in (d) the least duration gap, (e) MERL-LB-6, and (f) the least connection.

5.9 Understanding the sawtooth pattern of MERL-LB

To further understand the above sawtooth pattern of MERL-LB, the utilization rates of 4 types of resources of MERL-LB-6 over time are visualized in Figure 13. In each figure, 4 curves are depicted, i.e., the corresponding resource utilization of a server, the range of the corresponding resource utilization among all servers, the number of alive connections, and the ground truth remaining duration of this server.

It can be observed that the number of alive connections generally increases though keeps fluctuating. This indicates that new user requests are routed to the server continuously and some existing connections become disconnected coincidentally. With the increase of the number of connections, the remaining duration does not go up accordingly. In fact, its sawtooth pattern suggests that the requests routed to this server are mostly with remaining durations smaller than that of this server. Thus, the remaining duration of this server gradually drops as time goes by.

If we look at the place where the remaining duration of the server goes up significantly, it is usually found that the utilization of one resource type drops to the bottom of the corresponding range. If the utilization curve goes below its range, it means that the load balance in this resource type is enlarged. And it is necessary to route a request with a larger remaining duration to this server so that the resource utilization can be enlarged for a long time. On the other hand, once the utilization of one resource type goes to the top of the range, the number of alive connections will not increase. This is the opposite strategy learned by the policy to minimize the load balance by forcing the resource utilization not to exceed the range.

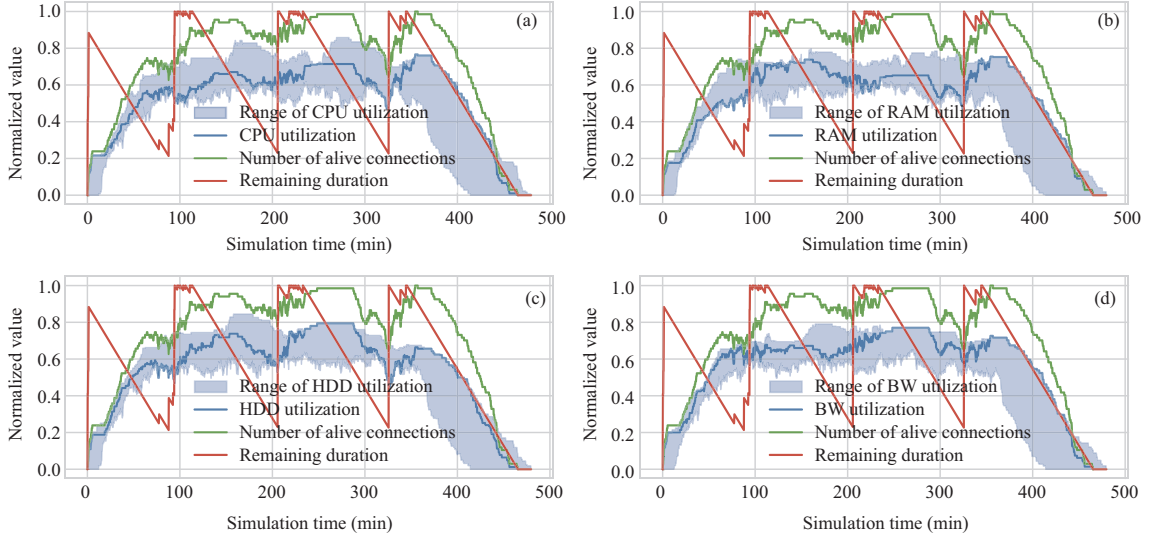


Figure 13 (Color online) Analysis of the sawtooth pattern of MERL-LB-6 on a server. The ground truth remaining duration changes with (a) CPU, (b) RAM, (c) HDD, (d) BW utilization over time.

5.10 A real-world application on the Alibaba Cluster-trace dataset

The simulations have also been carried out in a real-world scenario with the well-known Alibaba Cluster-trace-v2017 dataset. This dataset was released by Alibaba Cloud in September 2017⁴⁾, consisting of detailed statistics of 11089 online service jobs and 12951 batch jobs co-locating on 1300 machines over 12 h. The dataset has been found with four characteristics [11]. (1) Heterogeneous resource utilization across machines and workloads. (2) Greatly time-varying and multi-dimensional resource usages per workload and machine. (3) Imbalanced resource demands and runtime statistics (duration and task number) between online service and offline batch jobs. Thus, we believe this dataset is well suited to verify the proposed method.

In this experiment, we randomly picked 500 thousands user requests with their hardware demands and duration from the online service jobs. The picked requests were used as one simulation scenario for training the routing policies. We randomly generated 50 testing simulation scenarios by disturbing the picked data. Thus, each of the 50 testing simulation scenarios also consisted of 500 thousands user requests that were different from the picked one. All the compared algorithms were involved in the comparison and the experimental settings kept the same to the synthetic simulation cases.

The results in Table 4 and Figure 14 clearly show that our MERL-LB can still beat the compared algorithms in the similar pattern with the synthetic scenarios. The general performance of each algorithm does not surprisingly change. First, the Pareto solutions obtained by MERL-LB still distribute well. Second, there are always solutions of MERL-LB can significantly outperform the compared algorithms on either objective. Third, for each of the compared algorithm, there is always at least one solution of MERL-LB that dominates its output on both objectives. This verifies the effectiveness of the proposed MERL-LB in real-world datasets.

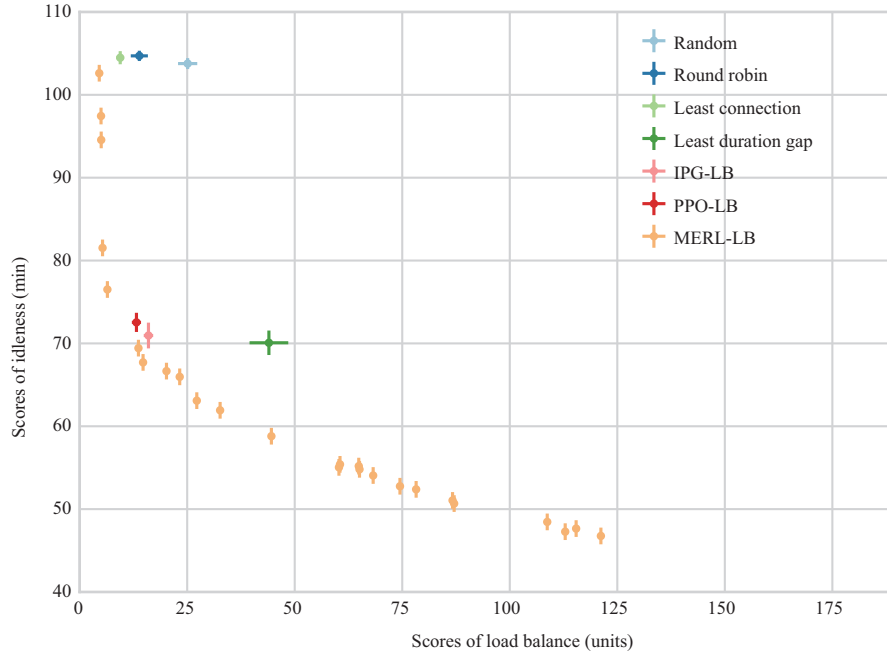
6 Conclusion

This paper first introduced the new issue of reducing idleness in the financial servers where the user connections were only allowed to be naturally disconnected by users. This paper targeted the issue as a constrained online routing problem and identified the limitations of existing load balancers in not considering user connection duration information. To address this issue, this paper proposes to take into account the user connection duration information and model the problem as a bi-objective RL problem, i.e., the load balancing objective and the idleness objective. With this modeling, a parameter-sharing routing policy is designed for varied numbers of servers, and an evolutionary multi-objective training framework is proposed based on NSGA-II to train the policy.

4) <https://github.com/alibaba/clusterdata>.

Table 4 Solution distribution of different algorithms in the objective space

Method	F_{balance} (units) ↓	F_{idle} (min) ↓
Random	25.11 ± 2.23	103.77 ± 0.66
Round robin	13.87 ± 1.99	104.70 ± 0.62
Least connection	9.43 ± 0.91	104.48 ± 0.79
Least duration gap	43.99 ± 4.50	70.08 ± 1.48
IPG-LB	15.99 ± 1.09	70.95 ± 1.55
PPO-LB	13.20 ± 1.06	72.54 ± 1.16
MERL-LB-1	4.56 ± 0.39	102.61 ± 1.01
MERL-LB-5	6.45 ± 0.54	76.51 ± 1.74
MERL-LB-6	13.67 ± 1.29	69.43 ± 0.99
MERL-LB-12	44.58 ± 4.28	58.80 ± 1.93
MERL-LB-25	121.19 ± 6.73	46.76 ± 2.89

**Figure 14** (Color online) Performance of different algorithms on two objectives. The cross-like error bars indicate the standard deviation on both objectives.

Extensive experimental studies have been conducted to reveal the advantages of the proposed method in detail. It is found that the proposed method generally beats comparable algorithms, including the traditional heuristics and advanced RL-based ones. The evolutionary multi-objective training framework not only facilitates policy training with a much faster convergence rate but also offers a set of diverse non-dominated policies for users' options. To consolidate the experiments, we tested the algorithms with different numbers of user requests and different numbers of servers. It was observed that the proposed method performed in a quite stable manner due to the proposed scalable routing policy. Then, the proposed method generated a sawtooth pattern of decision-making strategy, which has been analyzed to be beneficial in balancing the two objectives. Finally, the proposed method is verified on a real-world dataset released by Alibaba Cloud. The comparisons again support the effectiveness of MERL-LB.

We look forward to applying this policy to real-world financial cloud systems soon. In this regard, it is important to study the more robust and efficient evolutionary training method for MERL-LB in a realistic noisy environment [55] and expensive simulation environment [56]. Moreover, novel heuristics or solvers [57] inspired by the effective sawtooth pattern can be studied with sound theoretic guarantees.

Acknowledgements This work was supported in part by National Natural Science Foundation of China (Grant Nos. 62272210, 62331014, 62250710682), Guangdong Provincial Key Laboratory (Grant No. 2020B121201001), Program for Guangdong Introducing Innovative and Entrepreneurial Teams (Grant No. 2017ZT07X386), CCF-Tencent Open Fund (Grant No. RAGR20220110), and Market Data Cloud Joint Laboratory between Shenzhen Securities Information Co., Ltd. and Southern University of Science and Technology (Grant No. XB11ZC20210189).

References

- 1 Al-Dhuraibi Y, Paraiso F, Djarallah N, et al. Elasticity in cloud computing: state of the art and research challenges. *IEEE Trans Serv Comput*, 2018, 11: 430–447
- 2 Qu C, Calheiros R N, Buyya R. Auto-scaling web applications in clouds: a taxonomy and survey. *ACM Comput Surv*, 2019, 51: 1–33
- 3 Qureshi A, Weber R, Balakrishnan H, et al. Cutting the electric bill for internet-scale systems. In: *Proceedings of the ACM SIGCOMM 2009 Conference on Data Communication*, 2009. 123–134
- 4 Shastri S, Rizk A, Irwin D. Transient guarantees: maximizing the value of idle cloud capacity. In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, 2016. 992–1002
- 5 Stoll H R. Electronic trading in stock markets. *J Economic Perspect*, 2006, 20: 153–174
- 6 Li F. Cloud-native database systems at Alibaba. *Proc VLDB Endow*, 2019, 12: 2263–2272
- 7 Easley D, de Prado M M L, O'Hara M. The microstructure of the “flash crash”: flow toxicity, liquidity crashes, and the probability of informed trading. *J Portfolio Manage*, 2011, 37: 118–128
- 8 Mirobi G J, Arockiam L. Dynamic load balancing approach for minimizing the response time using an enhanced throttled load balancer in cloud computing. In: *Proceedings of International Conference on Smart Systems and Inventive Technology (ICSSIT)*, 2019. 570–575
- 9 Kansal N J, Chana I. Cloud load balancing techniques: a step towards green computing. *IJCSI Int J Comput Sci Issues*, 2012, 9: 238–246
- 10 Kumar P, Kumar R. Issues and challenges of load balancing techniques in cloud computing. *ACM Comput Surv*, 2019, 51: 1–35
- 11 Lu C, Ye K, Xu G, et al. Imbalance in the cloud: an analysis on Alibaba cluster trace. In: *Proceedings of IEEE International Conference on Big Data*, 2017. 2884–2892
- 12 Shafiq D A, Jhanjhi N Z, Abdullah A. Load balancing techniques in cloud computing environment: a review. *J King Saud Univ-Comput Inf Sci*, 2022, 34: 3910–3933
- 13 Johora F T, Ahmed I, Shajal M A I, et al. A load balancing strategy for reducing data loss risk on cloud using remodified throttled algorithm. *Int J Electr Comput Eng*, 2022, 12: 3217
- 14 Marinescu D C, Paya A, Morrison J P, et al. An approach for scaling cloud resource management. *Cluster Comput*, 2017, 20: 909–924
- 15 Sim K M, Sun W H. Ant colony optimization for routing and load-balancing: survey and new directions. *IEEE Trans Syst Man Cybern A*, 2003, 33: 560–572
- 16 Gures E, Shayea I, Ergen M, et al. Machine learning-based load balancing algorithms in future heterogeneous networks: a survey. *IEEE Access*, 2022, 10: 37689–37717
- 17 Brar G K, Chhabra A. Meta-heuristics based load balancing algorithms in grid and clouds — a review. In: *Proceedings of International Conference on Electrical, Electronics, and Optimization Techniques (ICEEOT)*, 2016. 2938–2943
- 18 Farag H, Stefanović Ć. Congestion-aware routing in dynamic IoT networks: a reinforcement learning approach. In: *Proceedings of IEEE Global Communications Conference (GLOBECOM)*, 2021. 1–6
- 19 Schulman J, Wolski F, Dhariwal P, et al. Proximal policy optimization algorithms. 2017. ArXiv:1707.06347
- 20 Mehta S. Speckle noise reduction using hybrid wavelet packets-wiener filter. *Int J Comput Sci Eng*, 2017, 5: 95–99
- 21 Alakeel A M. A guide to dynamic load balancing in distributed computer systems. *Int J Comput Sci Inform Secur*, 2010, 10: 153–160
- 22 Chen S L, Chen Y Y, Kuo S H. CLB: a novel load balancing architecture and algorithm for cloud services. *Comput Electr Eng*, 2017, 58: 154–160
- 23 Afzal S, Kavitha G. Load balancing in cloud computing — a hierarchical taxonomical classification. *J Cloud Comp*, 2019, 8: 22
- 24 Carrión C. Kubernetes scheduling: taxonomy, ongoing issues and challenges. *ACM Comput Surv*, 2022, 55: 1–37
- 25 Kashani M H, Mahdipour E. Load balancing algorithms in fog computing. *IEEE Trans Serv Comput*, 2023, 16: 1505–1521
- 26 Trivella A, Pisinger D. The load-balanced multi-dimensional bin-packing problem. *Comput Oper Res*, 2016, 74: 152–164
- 27 Basu D, Wang X, Hong Y, et al. Learn-as-you-go with Megh: efficient live migration of virtual machines. *IEEE Trans Parallel Distrib Syst*, 2019, 30: 1786–1801
- 28 Zhu J, Yang R, Sun X, et al. QoS-aware co-scheduling for distributed long-running applications on shared clusters. *IEEE Trans Parallel Distrib Syst*, 2022, 33: 4818–4834
- 29 Zhang S, Guo Y, Guo Z, et al. SMAF: a secure and makespan-aware framework for executing serverless workflows. *Sci China Inf Sci*, 2023, 66: 139105
- 30 Wang G G, Gao D, Pedrycz W. Solving multiobjective fuzzy job-shop scheduling problem by a hybrid adaptive differential evolution algorithm. *IEEE Trans Ind Inf*, 2022, 18: 8519–8528
- 31 Kandel S, Paepcke A, Hellerstein J M, et al. Enterprise data analysis and visualization: an interview study. *IEEE Trans Visual Comput Graph*, 2012, 18: 2917–2926
- 32 Salimans T, Ho J, Chen X, et al. Evolution strategies as a scalable alternative to reinforcement learning. 2017. ArXiv:1703.03864
- 33 Qian H, Yu Y. Derivative-free reinforcement learning: a review. *Front Comput Sci*, 2021, 15: 156336
- 34 Bai H, Cheng R, Jin Y. Evolutionary reinforcement learning: a survey. *Intell Comput*, 2023, 2: 0025
- 35 Jianye H, Li P, Tang H, et al. ERL-Re²: efficient evolutionary reinforcement learning with shared state representation and individual policy representation. In: *Proceedings of the 11th International Conference on Learning Representations*, 2022
- 36 Yang P, Yang Q, Tang K, et al. Parallel exploration via negatively correlated search. *Front Comput Sci*, 2021, 15: 155333
- 37 Wang Y, Xue K, Qian C. Evolutionary diversity optimization with clustering-based selection for reinforcement learning. In: *Proceedings of International Conference on Learning Representations*, 2022
- 38 Li B, Li J, Tang K, et al. Many-objective evolutionary algorithms: a survey. *ACM Comput Surv*, 2015, 48: 1–35
- 39 Qian C, Liu D X, Feng C, et al. Multi-objective evolutionary algorithms are generally good: maximizing monotone submodular functions over sequences. *Theor Comput Sci*, 2023, 943: 241–266
- 40 Wang S, Li B, Zhou A. A regularity augmented evolutionary algorithm with dual-space search for multiobjective optimization. *Swarm Evolary Comput*, 2023, 78: 101261
- 41 Hong W, Qian C, Tang K. Efficient minimum cost seed selection with theoretical guarantees for competitive influence maximization. *IEEE Trans Cybern*, 2020, 51: 6091–6104
- 42 Liu S, Lu N, Hong W, et al. Effective and imperceptible adversarial textual attack via multi-objectivization. 2021.

ArXiv:2111.01528

- 43 Liu S, Peng F, Tang K. Reliable robustness evaluation via automatically constructed attack ensembles. In: Proceedings of the 37th AAAI Conference on Artificial Intelligence, Washington, 2023. 8852–8860
- 44 Li B, Zhang Y, Yang P, et al. A two-population algorithm for large-scale multi-objective optimization based on fitness-aware operator and adaptive environmental selection. *IEEE Trans Evol Comput*, 2023, doi: 10.1109/TEVC.2023.3296488
- 45 Chen L, Xin B, Chen J. Interactive multiobjective evolutionary algorithm based on decomposition and compression. *Sci China Inf Sci*, 2021, 64: 202201
- 46 Falcón-Cardona J G, Coello C A C. Indicator-based multi-objective evolutionary algorithms: a comprehensive survey. *ACM Comput Surv*, 2020, 53: 35
- 47 Deb K, Pratap A, Agarwal S, et al. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comput*, 2002, 6: 182–197
- 48 Shen R, Zheng Y, Hao J, et al. Generating behavior-diverse game ais with evolutionary multi-objective deep reinforcement learning. In: Proceedings of the 29th International Joint Conference on Artificial Intelligence, 2020. 3371–3377
- 49 Song F, Xing H, Wang X, et al. Evolutionary multi-objective reinforcement learning based trajectory control and task offloading in UAV-assisted mobile edge computing. *IEEE Trans Mobile Comput*, 2023, 22: 7387–7405
- 50 Xu J, Tian Y, Ma P, et al. Prediction-guided multi-objective reinforcement learning for continuous robot control. In: Proceedings of International Conference on Machine Learning, 2020. 10607–10616
- 51 Abels A, Roijers D, Lenaerts T, et al. Dynamic weights in multi-objective deep reinforcement learning. In: Proceedings of International Conference on Machine Learning, 2019. 11–20
- 52 van Moffaert K, Nowé A. Multi-objective reinforcement learning using sets of Pareto dominating policies. *J Mach Learn Res*, 2014, 15: 3483–3512
- 53 Kaushik R, Chatzilygeroudis K, Mouret J B. Multi-objective model-based policy search for data-efficient learning with sparse rewards. In: Proceedings of the 2nd Conference on Robot Learning, 2018. 839–855
- 54 Mao H, Venkatakrishnan S B, Schwarzkopf M, et al. Variance reduction for reinforcement learning in input-driven environments. In: Proceedings of International Conference on Learning Representations, 2019
- 55 Bian C, Qian C, Yu Y, et al. On the robustness of median sampling in noisy evolutionary optimization. *Sci China Inf Sci*, 2021, 64: 150103
- 56 Hao H, Zhang X, Zhou A. Enhancing SAEAs with unevaluated solutions: a case study of relation model for expensive optimization. 2023. ArXiv:2309.11994
- 57 Liu S, Zhang Y, Tang K, et al. How good is neural combinatorial optimization? A systematic evaluation on the traveling salesman problem. *IEEE Comput Intell Mag*, 2023, 18: 14–28