

Optimizing Controller Placement for Software-Defined Networks

Victoria Huang, Gang Chen, Qiang Fu
School of Engineering and Computer Science
Victoria University of Wellington
Wellington, New Zealand
guiying.huang, aaron.chen, qiang.fu@ecs.vuw.ac.nz

Elliott Wen
School of Computer Science
The University of Auckland
Auckland, New Zealand
jwen929@aucklanduni.ac.nz

Abstract—Controller placement problem (CPP) is a key issue for Software-Defined Networking (SDN) with distributed controller architectures. This problem aims to determine a suitable number of controllers deployed in important locations so as to optimize the overall network performance. In comparison to communication delay, existing literature on the CPP assumes that the influence of controller workload distribution on network performance is negligible.

In this paper, we tackle the CPP that simultaneously considers the communication delay, the control plane utilization, and the controller workload distribution. Due to this reason, our CPP is intrinsically different from and clearly more difficult than any previously studied CPPs that are NP-hard. To tackle this challenging issue, we develop a new algorithm that seamlessly integrates the genetic algorithm (GA) and the gradient descent (GD) optimization method. Particularly, GA is used to search for suitable CPP solutions. The quality of each solution is further evaluated through GD. Simulation results on two representative network scenarios (small-scale and large-scale) show that our algorithm can effectively strike the trade-off between the control plane utilization and the network response time.

Index Terms—Software-Defined Networking, Distributed Controller Architecture, Controller Placement, Genetic Algorithm, Gradient Descent

I. INTRODUCTION

Software-Defined Networking (SDN) is an emerging networking paradigm notable for decoupling the network control from the data plane and forming a logically centralized external control plane. It enables centralized network management and significantly speeds up network innovation. Traditionally, the control plane is equipped with one single controller. As the network scales up, architectures supporting distributed controllers (e.g., ONOS [9]) have been successfully developed to substantially increase the combined processing capacity of the control plane in order to handle the growing demand for traffic processing.

The increasing popularity of distributed controller architectures also gives rise to new research problems among which an essential one is the *Controller Placement Problem* (CPP) [20], [24], [46]. The CPP aims to identify a suitable number of controllers as well as their locations so as to optimize the network performance. For example, the CPP is first introduced by Heller *et al.* [20] in an attempt to optimize controller place-

ment by minimizing the communication delay among switches and controllers. Recently, several approaches [22]–[24] have been proposed with the aim of improving the resilience of the control plane with respect to unexpected network failures. For example, the research work [22] developed a reliable controller placement strategy taking both the node reliability and link quality into consideration.

Nevertheless, a majority of existing works on the CPP aim to improve network performance without explicitly considering the impact of varied workload distributions among controllers. Particularly, the workload on each controller is usually oversimplified to be identical [21]. In some extreme cases [20], the controller workload is completely ignored. However, in reality, the workload distribution can bring significant changes to the controller performance. Specifically, when the controller workload reaches a certain level, the processing time can increase substantially [46]. In this case, the overall response time of a controller mainly depends on the processing time while the impact of communication delay becomes trivial.

In this paper, we will investigate the CPP based on our recently proposed BindingLess Architecture for Distributed Controllers (BLAC) [25]. BLAC introduces a scheduling plane to allow dynamic association between any controllers and any switches. In other words, every switch enjoys the flexibility of passing its processing requests to arbitrary controllers available in the network. This is vital for us to explicitly control the workload distribution across all controllers.

Based on BLAC, the CPP can be addressed in a systematic manner. Specifically, a queuing model will be established in this paper to simultaneously consider the impact of controller workload distribution, communication delay and control plane utilization on network performance. Unlike existing works restricted by the assumption of equal workload over all controllers, we consider a more realistic scenario in which the workload distribution matches closely with the controller location, capacity and other network settings. For example, in a heterogeneous network setting (i.e., controllers with different processing capacities are located in various geographic locations, leading to significantly different network latency), dispatching workload evenly is unlikely to achieve reasonable performance. Alternatively, sending more requests to nearby controllers without overloading them can be a better option.

Our new model of the CPP quantifies the benefits of the latter option and describes the CPP in the form of a constrained optimization problem with the objective of simultaneously reducing the packet response time and improving the control plane utilization.

Given the fact that CPPs without taking controller workload distribution into account are already NP-hard [10], [20], [46]. The CPP considered in this paper is also NP-hard since it generalizes previous CPPs by supporting dynamically changing controller workload. To tackle this problem, a new direct optimization method is designed by combining the use of both the genetic algorithm (GA) and gradient descent (GD) optimization to obtain near-optimal solutions. Particularly, GA is used to optimize both the number and locations of controllers in the network, according to which GD is further exploited to optimize the corresponding workload distribution over all deployed controllers.

In previous works [5], [6], [32], GA and GD are combined to build a memetic algorithm. In such an algorithm, evolutionary search is used to explore new solutions while GD-based local search is used for improving existing solutions. For example, *Li et al.* [32] solved a microphone array placement problem by applying GA to find a candidate solution which was further improved by GD. Different from these works, GD in our algorithm is used purely for fitness evaluation (no existing solutions will be improved through GD). Although we realize that similar approaches have been spotted lately in designing artificial neural networks [35], [39], our algorithm, however, aims to solve a very different problem in SDN. To the best of our knowledge, we are the first to simultaneously use GA and GD to address such networking problems.

To evaluate the performance of our proposed method, extensive simulations have been conducted under two network settings in terms of controller specification and network sizes. Simulation results show that our approach can effectively strike the trade-off between the controller utilization and response time.

II. RELATED WORK

Existing studies on the CPP can be generally classified into two categories: uncapacitated CPP (UCPP) and capacitated CPP (CCPP). Literature [20], [23], [24], [37] in UCPP barely considers the workload distribution among controllers when solving the CPP. For example, when the CPP is first proposed [20], the problem is assumed to be static and formulated as a K-median problem with the aim of minimizing the average communication delay among switches and controllers, which is further showed as NP-hard. *Obadia et al.* [37] formulated the CPP as the control plane overhead minimization problem and solved by a greedy heuristic. Note that these works all assume that the controller processing time is negligible compared with the communication delay. However, this assumption is not always valid. For example, when the controller workload reaches a certain level, the processing time can increase substantially which turns out to be a dominating factor in the packet response time [46].

Recently, an increasing number of research works [21], [30], [46] realize the importance of considering the workload distribution among deployed controllers. *Yao et al.* [46] redefined the CPP as a capacitated controller placement problem which considered the controller capacity as a constraint. With the aim of minimizing the load diversity among controllers, *Hu et al.* [10] applied GA to divide the network into a given number of domains so as to balance the workload in each domain. Similar research can also be found in [21], [29], [43]. Despite their promising performance, these methods tend to oversimplify the workload distribution. For instance, *Hock et al.* [21] considered that the workload of a controller equaled the number of connected switches. In other words, all switches are assumed to receive the same amount of network traffic. *Huque et al.* [43] set the controller workload as a fixed value corresponding to the maximum requests it can receive from the switches. However, these assumptions are unlikely to remain valid since the network traffic is highly dynamic and unpredictable. Therefore, how to model the impact of workload distribution in the CPP remains unsolved.

Apart from that, existing works on the CPP are based on static-binding-based controller architectures which can easily lead to controller overloading when associated switches generate high volume of requests within a short time frame [15]. Although dynamic-binding-based controller architectures [15] have been proposed to alleviate this issue by migrating switches from overloaded controllers to underloaded ones, the requests from one switch can still only be processed by one controller. Thus, the granularity of workload distribution is still limited to switch level. In this case, the destination controller is prone to be overloaded since its workload will increase dramatically if the switch migrated to it has accumulated a large number of pending requests.

To enable fine-grained workload distribution, we will investigate the CPP based on our recently proposed architecture BLAC. BLAC introduces bindingless switch-controller association so that requests from one switch can be flexibly processed by different controllers. The detailed design of BLAC will be demonstrated in Section III. On top of BLAC, the CPP is formulated as an optimization problem aiming at simultaneously reducing the response time while maintaining high control plane utilization. Different from previous works that tend to ignore or oversimplify the workload distribution, the relationship among the network response time, the workload distribution and communication delay will be systematically captured using a queuing model. Apart from that, we investigate the combined use of GA and GD to simultaneously identify both the CPP solution and the optimal workload distribution.

III. BINDINGLESS ARCHITECTURE FOR DISTRIBUTED CONTROLLERS

In this section, we briefly review BLAC and its design. After that, key reasons for choosing BLAC in this study will be elaborated.

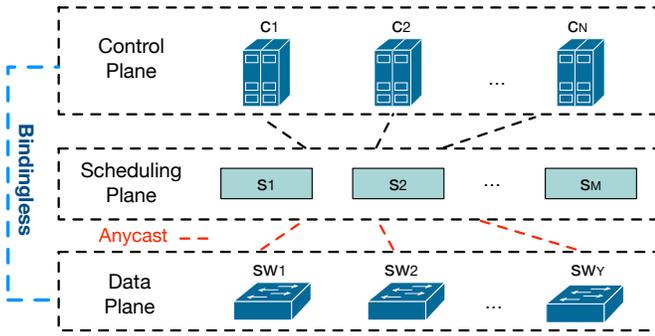


Fig. 1: System Architecture.

A. Architecture Design

As shown in Fig. 1, BLAC is comprised of three components: the data plane made up of numerous switches and communication links, the scheduling plane consisting of multiple schedulers, and the control plane composed of distributed controllers that jointly provide consistent network control.

The switches, denoted as $sw = \{sw_1, sw_2, \dots, sw_Y\}$ in the data plane, are in charge of packet forwarding based on well-defined flow rules. Specifically, upon receiving a packet, sw_y will try to match the packet against any rules stored in its Ternary Content Addressable Memory (TCAM) [2]. Only if no matching rule can be identified, a routing request will be forwarded to the scheduling plane for subsequent processing.

The scheduling plane is supported by a group of M schedulers, i.e., $s = \{s_1, s_2, \dots, s_M\}$. Each scheduler communicates with switches through the anycast technique [38], which guarantees low network latency and simplifies network configuration. Specifically, the switch will automatically connect to the nearest scheduler, reducing the network communication delay. Besides, all schedulers listen on a unique anycast address for the requests from nearby switches, enabling simple switch configuration. Note that the schedulers are low-cost and easy-to-deploy NFV entities. They can be strategically placed in many locations of a network to minimize communication delay without incurring too much cost. Thus, the scheduling plane should never become the bottleneck of the architecture. It is therefore not important for us to treat scheduler placement as a separate problem.

On top of the scheduling plane in Fig. 1 is the control plane, where N distributed controllers, denoted as $c = \{c_1, c_2, \dots, c_N\}$, will be established according to the controller placement decisions. Specifically, the processing capacities of the N controllers can be captured through vector $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_N) \in \mathbb{R}^N$ where α_n represents the maximum number of routing requests that controller c_n can process within a unit of time (i.e., a second). Each controller is free to communicate with any scheduler in the scheduling plane. The communication delay between s and c is collectively expressed through a matrix $D \in \mathbb{R}^{M \times N}$ where d_{mn} is the communication delay between s_m and c_n . Every time when a scheduler s_m receives a packet (or routing request), it must forward it to one of the controllers for processing. Such

a controller c_n will be selected according to a probability denoted as p_{mn} . For concise discussion, we jointly present the selection probabilities for all controllers in the form of a matrix $P \in \mathbb{R}^{M \times N}$.

B. Using BLAC for the Controller Placement Problem

In this research, we decide to use BLAC to address the CPP because of two main reasons:

First, BLAC enables dynamic and transparent controller placement without incurring time-consuming switch migration. Other distributed controller architectures such as ONOS [9] and ElastiCon [15] support either static or dynamic switch-controller binding. Therefore, when the controllers are relocated, switches need to migrate from previous controllers to new ones. This process can be time-consuming and even introduce network disruption. On the contrary, BLAC introduces bindingless switch-controller association. In BLAC, switches are connected directly to the scheduling plane. Thus, controllers can be flexibly and transparently relocated without interrupting the switch connections.

Second, BLAC enables fine-grained controller load control. Due to the switch-controller binding constraint, the majority of distributed controller architectures can only manipulate the controller workload at switch level. In other words, requests from one switch can only be processed by one controller. However, BLAC can dispatch requests from one switch to different controllers, achieving flow-level load control. Thus, the architecture improves the flexibility of distributing the workload among controllers.

Apart from providing flexible and explicit control over the controller workload distribution, the complete separation between the control plane and the data plane also enables us to formulate the CPP without the switch-controller binding constraint. With the help of BLAC, the CPP is effectively formulated in Section IV-A.

IV. CONTROLLER PLACEMENT

After reviewing the architecture design, we now consider how to solve the CPP based on BLAC. In this section, we address the CPP for the purpose of improving network response time and control plane utilization. Particularly, we can formulate the CPP as a constrained optimization problem using a queuing model. In line with the problem formulation, we will further investigate several alternative ways to solve the CPP, ranging from simple heuristic approaches to more advanced search methods.

A. Problem Formulation

Using BLAC introduced in Section III, within a reasonable period of time, we assume that in the queuing model, the requests arriving at s_m ($m = 1, \dots, M$) follow a Poisson distribution with a stable arrival rate denoted as λ_m . This is a common assumption in the literature [7], [18], [44], enabling us to easily analyze the performance of controllers. Although our formulation is based on this assumption, we measured the simulation results under different traffic distribution and

found that the simulated performance did not depend heavily on this assumption. As previous work [25], [44], we also assume that the complexity for processing each request by the same controller is roughly identical since the complexity of routing decisions depends on the network size that remains stable within a long time span [14]. While the complexity of routing decisions being identical, it should also be noted that if controllers have different capacities, the processing time would vary significantly from one controller to another.

Given a group of candidate controllers c with fixed locations selected based on network constraints (e.g. bandwidth), the job for the CPP is to select a subset of locations (controllers) to meet the performance demand from the data plane. In view of this, the CPP can be converted to a controller location selection problem. Note that different from previous works [20], [24], [46] which assume the number of selected controllers is given, the CPP to be tackled here requires the number of deployed controllers to be determined too. In particular, the solution of the CPP can be represented as a binary vector $\mathbf{x} = (x_1, x_2, \dots, x_N) \in \mathbb{R}^N$, where each dimension of \mathbf{x} can be represented as

$$x_n = \begin{cases} 1, & \text{if } c_n \text{ is selected} \\ 0, & \text{otherwise} \end{cases} \quad (1)$$

and the number of selected controllers can be calculated as $\sum_{n=1}^N x_n$.

Similar to recent works [36], [44] which try to understand the controller's behavior using a queuing model, we consider each controller as an independent M/M/1 queue. Unlike these works which model the controller traffic with the switch-controller binding constraint, we consider that each controller maintains a queue for pending requests arriving from any scheduler/switch. This assumption facilitates simple and efficient implementation of schedulers (no queuing management required). Accordingly, c_n 's workload can be determined as:

$$\theta_n = \sum_{m=1}^M \lambda_m p_{mn} \quad (2)$$

According to Little's Law [34], the average processing time of c_n with capacity α_n is:

$$v_n = \frac{1}{\alpha_n - \theta_n} \quad (3)$$

Since the network we consider here may span large geographic areas, the communication delay d_{mn} between s_m and c_n must be taken into account when computing the response time. The average communication delay of c_n is:

$$l_n = \frac{\sum_{m=1}^M \lambda_m p_{mn} d_{mn}}{\theta_n} \quad (4)$$

In consideration of both the processing time and the communication delay, the average response time of c_n is:

$$t_n = v_n + 2l_n \quad (5)$$

Based on the design of BLAC, it is easy to see that the communication delay between a switch and the scheduler

responsible for handling its requests is generally fixed. We can also see that the time required for a scheduler to dispatch a request to any controller is negligible (provided that a time-efficient scheduling algorithm is used). Therefore in our CPP, we only need to consider t_n that measures the response time between schedulers and controllers.

To simplify our discussion, our model takes no consideration of synchronization costs among controllers. The simplification is justifiable due to the following aspects: (1) State-of-the-art distributed controller architectures/frameworks are designed to incorporate a large number of controllers such that multiple optimization techniques (e.g., the anti-entropy mechanism [9]) have been applied to achieve almost negligible synchronization costs. (2) Even if the costs cannot be ignored, their impact can be easily incorporated into our model by adjusting the controller capacities accordingly. Based on the above arguments, a majority of existing research works [20], [24], [46] tend to safely ignore the synchronization costs.

The average response time over all requests generated by the data plane can be decided from below:

$$t = \frac{\sum_{n=1}^N \theta_n t_n}{\sum_{n=1}^N \theta_n} \quad (6)$$

Given the request arrival rate and the controller selection decision \mathbf{x} , the controller utilization can be determined as:

$$u = \frac{\sum_m^M \lambda_m}{\sum_n^N \alpha_n x_n} \quad (7)$$

In this study, we choose to use t in (6) as the metric for control plane performance. In fact, by selecting controller locations properly, the performance of the control plane can be improved, resulting directly in reduced t . Meanwhile, to maintain reasonable operational cost, the average controller utilization u must be kept at a high level to avoid controller over-provisioning. In consideration of both requirements, an intuitive approach is to formulate the CPP as below:

$$\underset{\mathbf{x} \in \{0,1\}^N}{\text{minimize}} \quad \underset{\mathbf{P}}{\text{minimize}} \quad \mu_t t(\mathbf{x}, \mathbf{P}) - \mu_u u(\mathbf{x}) \quad (8)$$

Note that both $t(\mathbf{x}, \mathbf{P})$ and $u(\mathbf{x})$ are direct functions of placement vector \mathbf{x} and workload distribution probability \mathbf{P} . μ_t and μ_u represent the weight coefficients which strike the trade-off between the response time and utilization. However, this formulation may not be suitable for two reasons. First, it requires a high level of domain expertise to select the appropriate value for the weights. Second, (8) requires both objective values to be normalized.

In consideration of the above reasons, the CPP is formalized as below:

$$\underset{\mathbf{x} \in \{0,1\}^N}{\text{minimize}} \quad \underset{\mathbf{P}}{\text{minimize}} \quad \frac{t(\mathbf{x}, \mathbf{P})}{u(\mathbf{x})} \quad (9)$$

$$\text{subject to} \quad \theta_n \leq \beta_n \alpha_n, \quad \forall n \quad (10)$$

$$0 \leq p_{mn} \leq 1, \quad \forall m, n \quad (11)$$

$$\sum_{n=1}^N p_{mn} = 1, \quad \forall m \quad (12)$$

Inequality (10) guarantees that the workload of each controller chosen by \mathbf{x} will not exceed their capacity. Specifically, through properly adjusting $\beta = (\beta_1, \beta_2, \dots, \beta_N)$, sufficient capacity can be reserved at each controller to cope with unexpected ephemeral traffic bursts. Both inequality (11) and equality (12) jointly ensure that the packet distribution probabilities p_{mn} are well-defined.

B. Placement Optimization

The CPP can be considered as a variant of the facility location problem, which is known as NP-hard [20]. Thus, finding the optimal solution is almost impossible for large networks. To solve this problem efficiently, we will study and compare several alternative approaches.

1) **Random Approach:** One simple and widely used heuristic is the random selection. It selects controllers randomly and uniformly from the given set of candidate controllers \mathcal{C} until total capacity of all selected controllers reaches a given level (e.g., $\gamma \|\lambda\|_1$, where γ is a provisioning factor and $\|\lambda\|_1$ is the L_1 norm of request arrival rate λ).

Due to its simplicity, the random approach is widely utilized in real-world networks [27], [33]. However, the effectiveness of the random approach relies on the assumption that all controllers have identical capacities and propagation latencies. It is unlikely for this assumption to remain valid with SDN controllers hosted by hybrid data centers.

2) **Capacity-based Greedy Approach:** To address the limitations of the random approach, a greedy approach is considered. Specifically, the controller selection is performed in an iterative manner. In each iteration, the controller with the highest capacity among existing unselected controllers is chosen. The iteration terminates when the capacity of all selected controllers reaches a certain level.

Although the greedy approach is simple and efficient, it only considers the controller capacity, deliberately ignoring the utilization of the control plane and the communication delays between controllers and schedulers. As a result, it may lead to high operational costs and long communication delays.

3) **K-median Approach:** Currently, one of the most widely-adopted controller placement strategies is called K-median introduced by *Heller et al.* [20]. K-median aims to find the controller locations that minimize the average communication delay. Note that the number of controllers is generally assumed to be given or can be easily obtained in the conventional K-median approach. However, finding the number of controllers that can closely match the network traffic in a wide-area network equipped with controllers with various capacities can be complicated. Thus, in this paper, we drop the assumption and adapt K-median to solve the CPP by adding one constraint, i.e., the total capacity of all selected controllers must reach a

given level. The adapted K-median can be formulated as:

$$\text{minimize}_{x_n \in \{0,1\}} \frac{\sum_{m,n} d_{mn} x_n}{\sum_n x_n} \quad (13)$$

$$\text{subject to} \quad \sum_n x_n \alpha_n \geq \gamma \|\lambda\|_1 \quad (14)$$

where γ and $\|\lambda\|_1$ are the provisioning factor and total request arrival rate described in random approach.

To solve this problem, an n -approximation algorithm called forward greedy approach is adapted [12]. It starts by assigning a location x_n with smallest $d_{m,n}$ to be 1. Then at each iteration, the controller with the least average communication delay among existing unselected controllers is chosen until the total capacity constraint is satisfied.

Compared with previous two heuristics, K-median approach focuses on minimizing the communication delay which plays an important role in network response time. Therefore, the K-median approach is expected to be more effective than previous heuristics, especially in wide-area networks. However, the K-median approach still cannot effectively manage the utilization of the control plane, which may result in additional operating costs.

4) **Direct Optimization Approach:** In literature, evolutionary computation (EC) algorithms are often exploited to find near-optimal solutions to NP-hard problems [13], [16], [26]. They stand for a promising alternative approach to tackle the CPP formulated in (9). Specifically, GA with a binary solution representation is employed by us for controller placement. GA is an EC approach inspired by Darwin's theory of evolution. It solves the optimization problem by simulating the process of natural selection to find highly fit solutions to a given problem. In this paper, we follow the GA framework introduced in [45]. Each individual in a GA population is a binary vector \mathbf{x} that directly represents a solution to our CPP, as explained in Section IV-A.

Given any possible controller placement solution \mathbf{x} , it must satisfy all constraints from (10) to (12). To enforce constraint satisfaction during the evolution, we transform the constrained optimization problem in Section IV-A into an unconstrained one by introducing several penalty terms [40]. Consequently, the fitness function for GA to determine the goodness of solution \mathbf{x} becomes:

$$f(\mathbf{x}) = \text{minimize}_{\mathbf{P}} \left\{ \begin{aligned} & \frac{t(\mathbf{x}, \mathbf{P})}{u(\mathbf{x})} \\ & - \mu_1 \sum_{n=1}^N \min(0, \beta_n a_n - \theta_n) \\ & - \mu_2 \sum_{m=1}^M \min(0, 1 - \sum_{n=1}^{N-1} p_{mn}) \\ & - \mu_3 \sum_{m=1}^M \sum_{n=1}^N \min(0, p_{m,n}) \end{aligned} \right\} \quad (15)$$

where μ_1 , μ_2 , and μ_3 are penalty coefficients used to quantify the imposed penalty terms. Usually, the value of penalty

coefficients can be set as a monotonically decreasing function of the generation number in GA so that infeasible solutions are greatly penalized and the existing solutions in a GA population are forced to satisfy the constraints.

Algorithm 1 GD-based Fitness Evaluation

- 1: Initialize the workload distribution \mathbf{P}
 - 2: **for** iteration = 1, 2, ... **do**
 - 3: Calculate the gradient \mathbf{g} using Theano: $\mathbf{g} \leftarrow \frac{\partial f}{\partial \mathbf{P}}$
 - 4: Update \mathbf{P} based on (17)
 - 5: **end for**
 - 6: **return** Fitness value of a given solution $\mathbf{x} \leftarrow f(\mathbf{x}, \mathbf{P})$
-

According to (15), the fitness of \mathbf{x} is obtained through minimizing (15) with respect to \mathbf{P} . Such minimization can be approached straightforwardly through a gradient descent (GD) approach, as shown in Algorithm 1. In each iteration of GD, given candidate solution \mathbf{x} for our CPP, the gradient \mathbf{g} of the fitness function $f(\mathbf{x}, \mathbf{P})$ with respect to \mathbf{P} is calculated via Theano [3], an efficient gradient computation tool for large-scale optimization. Instead of updating the parameter \mathbf{P} using a constant learning rate l , we adaptively adjust l from l_H (a higher threshold) to l_L (a lower threshold) using equation (16):

$$l = l_H - \frac{(l_H - l_L)_i}{N_I} \quad (16)$$

where N_I is the total number of iterations and i represents the i^{th} iteration.

Based on l in (16) and the gradient \mathbf{g} calculated by Theano, \mathbf{P} can be updated as:

$$\mathbf{P}_{i+1} = \mathbf{P}_i - l\mathbf{g} \quad (17)$$

The whole process iterates until the stopping criteria are reached. As a result of minimizing (15) through iterative updating \mathbf{P} in GD, the fitness of \mathbf{x} can be obtained eventually.

Algorithm 2 Direct Optimization Algorithm (GA+GD)

- 1: Initialize the population
 - 2: Evaluate the population using Algorithm 1
 - 3: **for** generation = 1, 2, ... **do**
 - 4: Generate new individuals via Crossover and Mutation
 - 5: Evaluate the new individuals using Algorithm 1
 - 6: Select individuals to form a new population
 - 7: **end for**
 - 8: **return** The best individual
-

Based on the fitness values obtained from Algorithm 1, the best performing individuals are selected as the basis for the new generation that is produced via genetic operators (e.g., mutation and crossover). This evolution process repeats over many iterations until a sufficiently fit solution is found. The overall procedure is summarized in Algorithm 2. The complexity of our solution is $\mathcal{O}(M^3 + MN)$ where M and N are the numbers of schedulers and controllers respectively.

V. EVALUATION

In this section, we first present the evaluation setup. Then the performance of our algorithm will be examined using two representative networks: a small-scale network and a large-scale network. To ease the discussion, we start with a small-scale network. To generalize our simulation setting, we consider that the network is equipped with a set of heterogeneous SDN controllers. Such a setting is frequently seen especially for small data centers hosted on the premises of many organizations in the process of hardware upgrading. After that, we consider a large-scale network that supports a hybrid data center with controller resources scattered around the globe. In this case, the CPP becomes extremely challenging (i.e., NP-hard) making simple heuristics ineffective. On the other hand, GA is expected to address this complex problem effectively.

A. Evaluation Setup

During the evaluation, the network topology we used is the fat-tree topology which is widely used in real-world data centers (e.g., Facebook [1] and Google [4]) due to its simplicity and efficiency.

Note that the communication delay highly depends on the geographic distance covered by the network. We simulate the communication delay differently in two networks. In particular, for the small-scale network, we consider that all network devices are placed within the data center. Therefore the communication delay can be safely ignored (below 1ms) [17]. Nonetheless, to make the simulation more accurate, we simulate the latency sampled from a distribution observed in a real-world data center [17]. For the large-scale network which may span large geographic areas, the communication delay becomes significant. According to existing studies, the communication delay can vary from 10 ms to 200 ms [11], [19], [31]. In our simulations, the communication delay is simulated by sampling from a distribution measured in [31].

Regarding the parameter tuning, we mainly follow the typical settings in existing works [41], [42]. Specifically, a chromosome in GA is a binary vector representing a CPP solution. The mutation and crossover rates are 0.1 and 1 respectively. For fitness evaluation, GD is performed for 15 iterations so as to balance the algorithm efficiency and performance effectiveness. As for the penalty coefficients $\boldsymbol{\mu} = (\mu_1, \mu_2, \mu_3)$ in (15), we set them to be 10 for all simulations. In fact, different settings of $\boldsymbol{\mu}$ ranging from 1 to 100 have been tested and no significant performance impact was noticed. We therefore believe our algorithm can work robustly with varied settings of $\boldsymbol{\mu}$. Besides, the decay factor β in (15) for each controller is set to be 0.83. Correspondingly, the provisioning factor γ for the heuristics is set to be 1.2.

B. Small-scale Network

The topology consists of 80 switches at a scale comparable to many university and enterprise data centers [8]. We set up a group of 10 candidate controllers with different capacities. Since the network we used in this setting is small, we only

TABLE I: Controller placement performance comparison in a small-scale network with heterogeneous controllers.

CPP Algorithm	Random Approach	Capacity_based Greedy Approach	K-median Approach	GA+GD
Average Response Time (ms)	0.3928	0.2807	0.2816	0.2816
Control Plane Utilization (%)	80.00	66.67	72.72	80.00
Controller Number	4	4	5	4

consider controllers with two different capacities (i.e., 45,000 and 30,000 pkt/s). The total number of requests generated by the entire data plane is averaged to be 120,000 pkt/s.

We measured and compared the network performance with different numbers of schedulers. The results show that there is no significant difference in terms of the network performance. This finding confirms our belief that schedulers are not the performance bottleneck. Given the relatively small network size, we decide to report results when only one scheduler is used. In the small-scale network setting, we first demonstrate the influence of workload distribution on network performance. After that, we evaluate the performance of different placement methods mentioned in Section IV-B using BLAC.

To demonstrate that it is important to consider workload distribution, comparison experiments are conducted. Specifically, given a placement solution, we distribute requests using the workload distribution optimized by GD and measure the network response time. In comparison, without the help of GD, we consider that each switch is statically connected to its nearest controller and sends all requests to that controller. In terms of CPP strategies, we run experiments using different heuristics (e.g., greedy and K-median) and all of them show similar performance patterns. Due to the space limit, we only report results with K-median.

As we expected, the average response time we measured without considering workload distribution is 0.65 ms in K-median which is over 2 times larger than 0.28 ms in K-median+GD. This is because K-median solely optimizes the communication delay without considering the controller workload distribution. Note that in a small-scale network with negligible communication delay, it is preferable to dispatch requests to controllers so that the utilization is balanced among all controllers. Since the switches and network traffic are not evenly distributed in the network, without considering workload distribution can easily lead to uneven control plane utilization. To verify this, we measure the controller utilization as the total requests received by the controller divided by the controller capacity. As shown in Fig. 2, with the help of GD, the controller utilization is relatively even and less than 80%. On the other hand, we also notice that when only K-median is used, 2 controllers are highly-loaded (over 90% utilization) while 2 controllers remain under-utilized (i.e., less than 60% utilization), rendering comparatively high response time.

To compare the performance of different placement methods mentioned in Section IV-B, two performance metrics including the average response time defined in (6) and control plane utilization defined in (7) are evaluated. In particular, regardless which algorithm is used to deploy controllers in the following simulations, the workload distribution over all deployed controllers will be consistently optimized through GD.

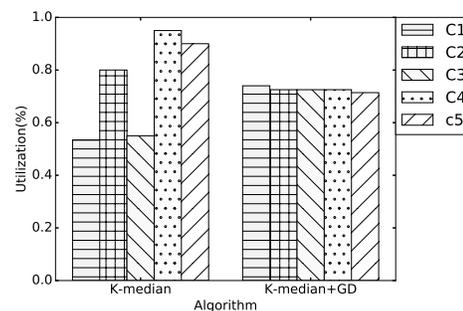


Fig. 2: Controller utilization comparison between K-median and K-median with GD.

Intuitively, given the same request arrival rate, controllers with more capacities tend to achieve less response time. Consequently, we expect the capacity-based greedy heuristic to be highly effective. Our simulation result as shown in Table I clearly demonstrates the usefulness of this greedy heuristic with 0.2807 ms response time. However, we notice that GA provides an alternative selection which provides similar response time (0.2816 ms) with higher control plane utilization (80%). Although choosing the most powerful controllers as capacity-based greedy approach can clearly achieve best response time, it has resulted in low control plane utilization (e.g., 66.7% in this example), incurring additional energy and operational costs. With regard to GA, it achieves clearly better trade-off between response time and controller utilization, thus is more suitable for controller placement.

C. Large-scale Network

To evaluate the performance of GA, we simulate a network with a total number of 720 switches, at a scale comparable to many large commercial data centers [8]. Besides, 50 controller candidates are provided for the CPP with capacities spanning from 15,000 pkt/s to 90,000 pkt/s [28].

Table II shows the evaluation result using different placement methods with the combined request arrival rate over all switches in the network reaching 720,000 pkt/s. Significantly different from Section V-B, capacity-based greedy approach has the highest response time (26.1 ms) among all methods. This is understandable because the communication delay plays a major part in the response time in a large-scale network while capacity-based greedy heuristic chooses controllers solely based on their capacities. Thus, remote controllers may be selected, resulting in large response time. We also notice that the K-median approach is a good competitor to GA with respect to response time (8.5 ms for both GA and K-median). However, in terms of controller utilization, GA with 82.75% utilization outperforms the K-median approach with 76.19% utilization thanks to our design of fitness function in (15) that

TABLE II: Controller placement performance comparison in a large-scale network with heterogeneous controllers.

CPP Algorithm	Random Approach	Capacity_based Greedy Approach	K-median Approach	GA+GD
Average Response Time (ms)	16.3	26.1	8.5	8.5
Control Plane Utilization (%)	81.36	80.00	76.19	82.75
Controller Number	18	10	20	16

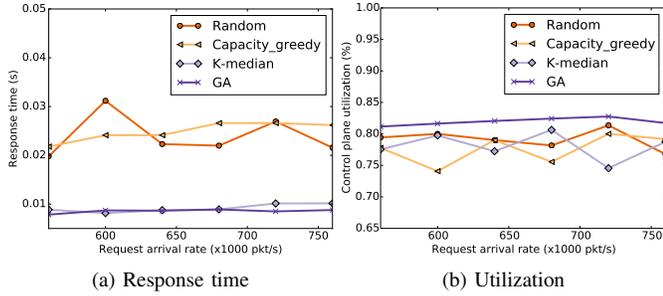


Fig. 3: The performance comparison of different methods for controller placement in a large scale network with heterogeneous controller setting.

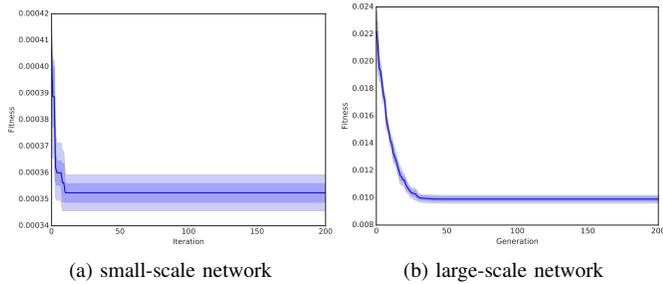


Fig. 4: The convergence of GA for controller placement problem in a network with heterogeneous controller setting.

enables GA to optimize both the response time and controller utilization simultaneously.

To further demonstrate the effectiveness of GA, we run different placement methods at various request arrival rates. Fig. 3 depicts the performance comparison in terms of the packet response time and control plane utilization. It can be seen from Fig. 3(a) that the response time obtained by using both GA and K-median remain steadily below 10 ms. In certain situations, GA even performs better. In comparison, the response time of the random approach shows large fluctuation ranging from 20 ms to 32 ms due to the randomly chosen remote controllers. In terms of the control plane utilization, we notice that GA achieves the highest utilization which is around 82% (closely matches the decay factor $\beta = 0.83$ we set in Section V-A) and remains nearly constant regardless the traffic changes. In comparison, the utilization of other methods varies dramatically with different packet arrival rates. For example, the lowest utilization which both capacity-based greedy and K-median methods reach is 74% while their highest utilization can be 80%. This is mainly because these methods choose controllers solely based on their capacities or communication delay, without carefully considering or even completely ignoring the control plane utilization. Thus, the

utilization of the control plane can vary dramatically as the request arrival rate changes.

We also investigate the convergence speed of GA by running the algorithm for 30 times with two different network settings (i.e., Section V-B and Section V-C) respectively. In each run, GA is set with 200 generations with a population size 50 for the small-scale network and 200 for the large-scale network. Fig. 4 depicts the change of fitness value across generations with both 68% and 95% confidence bands. From Fig. 4, we find that GA converges quickly regardless of the problem size. In the small-scale network, GA can converge within 10 generations. In the large network, the algorithm can still converge within 30 generations.

Considering that GA can be potentially computational intensive, several approaches can be considered to improve the practical usefulness of our algorithm. For example, as demonstrated in many recent studies, GA can outperform many competing optimization methods in terms of computation speed through parallel evaluation of all candidate solutions in a population. Apart from that, since network traffic often exhibits diurnal and weekly patterns, we can predict upcoming traffic change based on historical data in order to plan ahead and get the required placement ready in time for expected traffic spikes and long-term changes.

VI. CONCLUSIONS

When deploying distributed controllers in an SDN network, one essential issue is the controller placement problem (CPP). Existing works mainly focus on reducing the communication delay while undermining the significance of the workload distribution among controllers. However, previous works show that the performance of the CPP can be highly related to the workload distribution. In this paper, a queuing model is built which systematically measures the relationship among the network response time, the workload distribution and the communication delay. Meanwhile, to maintain reasonable operation cost, the control plane utilization must be kept at a high level. Motivated by this, we formulate the CPP as a constrained optimization problem using the built queuing model, which simultaneously optimizes the response time and the control plane utilization. Several alternative ways are investigated to solve the CPP, ranging from simple heuristic approaches to more advanced search methods. In particular, a new algorithm combining the use of GA and GD is proposed. The performance of our algorithm is analyzed in detail via a series of simulation featuring different network settings. It is shown that our algorithm achieved highest control plane utilization and competitively low response time compared to the widely-used heuristic methods.

REFERENCES

- [1] "Introducing data center fabric, the next-generation Facebook data center network," <https://code.facebook.com/posts/360346274145943/>.
- [2] "TCAMs and OpenFlow : What Every SDN Practitioner Must Know," <https://www.sdxcentral.com/articles/contributed/sdn-openflow-tcam-need-to-know/2012/07/>.
- [3] "Theano," <http://deeplearning.net/software/theano/>.
- [4] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 4. ACM, 2008, pp. 63–74.
- [5] A. Alberó Ortíz, J. Monzó Cabrera, A. B. Díaz Morcillo, M. E. Requena Pérez *et al.*, "Combined use of genetic algorithms and gradient descent optimization methods for accurate inverse permittivity measurement," 2006.
- [6] A. Arab and A. Alfi, "An adaptive gradient descent-based local search in memetic algorithm applied to optimal controller design," *Information Sciences*, vol. 299, pp. 117–142, 2015.
- [7] W. Bai, L. Chen, K. Chen, D. Han, C. Tian, and H. Wang, "Pias: Practical information-agnostic flow scheduling for commodity data centers," *IEEE/ACM Transactions on Networking (TON)*, vol. 25, no. 4, pp. 1954–1967, 2017.
- [8] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 2010, pp. 267–280.
- [9] P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow *et al.*, "Onos: towards an open, distributed sdn os," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*. ACM, 2014, pp. 1–6.
- [10] H. Bo, W. Youke, W. Chuan'an, and W. Ying, "The controller placement problem for software-defined networks," in *Computer and Communications (ICCC), 2016 2nd IEEE International Conference on*. IEEE, 2016, pp. 2435–2439.
- [11] S. Choy, B. Wong, G. Simon, and C. Rosenberg, "The brewing storm in cloud gaming: A measurement study on cloud to end-user latency," in *Proceedings of the 11th annual workshop on network and systems support for games*. IEEE Press, 2012, p. 2.
- [12] M. Chrobak, C. Kenyon, and N. E. Young, "The reverse greedy algorithm for the metric k-median problem," in *International Computing and Combinatorics Conference*. Springer, 2005, pp. 654–660.
- [13] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: Nsga-ii," *IEEE transactions on evolutionary computation*, vol. 6, no. 2, pp. 182–197, 2002.
- [14] E. W. Dijkstra, "A note on two problems in connexion with graphs," *Numerische mathematik*, vol. 1, no. 1, pp. 269–271, 1959.
- [15] A. A. Dixit, F. Hao, S. Mukherjee, T. Lakshman, and R. Kompella, "Elasticon: an elastic distributed sdn controller," in *Proceedings of the Tenth ACM/IEEE Symposium on Architectures for Networking and Communications Systems*. ACM, 2014, pp. 17–28.
- [16] A. E. Eiben and J. Smith, "From evolutionary computation to the evolution of things," *Nature*, vol. 521, no. 7553, p. 476, 2015.
- [17] C. Guo, L. Yuan, D. Xiang, Y. Dang, R. Huang, D. Maltz, Z. Liu, V. Wang, B. Pang, H. Chen *et al.*, "Pingmesh: A large-scale system for data center network latency measurement and analysis," in *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4. ACM, 2015, pp. 139–152.
- [18] J. Guo, F. Liu, T. Wang, and J. C. Lui, "Pricing intra-datacenter networks with over-committed bandwidth guarantee," in *Proc. USENIX ATC*, 2017.
- [19] K. He, A. Fisher, L. Wang, A. Gember, A. Akella, and T. Ristenpart, "Next stop, the cloud: Understanding modern web service deployment in ec2 and azure," in *Proceedings of the 2013 conference on Internet measurement conference*. ACM, 2013, pp. 177–190.
- [20] B. Heller, R. Sherwood, and N. McKeown, "The controller placement problem," in *Proceedings of the first workshop on Hot topics in software defined networks*. ACM, 2012, pp. 7–12.
- [21] D. Hock, M. Hartmann, S. Gebert, M. Jarschel, T. Zinner, and P. Tran-Gia, "Pareto-optimal resilient controller placement in sdn-based core networks," in *Teletraffic Congress (ITC), 2013 25th International*. IEEE, 2013, pp. 1–9.
- [22] T. Hu, J. Zhang, L. Cao, and J. Gao, "A reliable controller deployment strategy based on network condition evaluation in sdn," in *Software Engineering and Service Science (ICSESS), 2017 8th IEEE International Conference on*. IEEE, 2017, pp. 367–370.
- [23] Y. Hu, W. Wang, X. Gong, X. Que, and S. Cheng, "On reliability-optimized controller placement for software-defined networks," *China Communications*, vol. 11, no. 2, pp. 38–54, 2014.
- [24] Y. Hu, W. Wendong, X. Gong, X. Que, and C. Shiduan, "Reliability-aware controller placement for software-defined networks," in *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on*. IEEE, 2013, pp. 672–675.
- [25] V. Huang, Q. Fu, G. Chen, E. Wen, and J. Hart, "BLAC: A bindingless architecture for distributed sdn controllers," in *2017 IEEE 42nd Conference on Local Computer Networks (LCN)*. IEEE, 2017, pp. 146–154.
- [26] A. Jaskiewicz, "On the performance of multiple-objective genetic local search on the 0/1 knapsack problem—a comparative experiment," *IEEE Transactions on Evolutionary Computation*, vol. 6, no. 4, pp. 402–412, 2002.
- [27] C. Joo and N. B. Shroff, "Performance of random access scheduling schemes in multi-hop wireless networks," *IEEE/ACM Transactions on Networking*, vol. 17, no. 5, pp. 1481–1493, 2009.
- [28] D. Kreutz, F. M. Ramos, P. E. Verissimo, C. E. Rothenberg, S. Azodolmoly, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015.
- [29] A. Ksentini, M. Baga, T. Taleb, and I. Balasingham, "On using bargaining game for optimal placement of sdn controllers," in *Communications (ICC), 2016 IEEE International Conference on*. IEEE, 2016, pp. 1–6.
- [30] S. Lange, S. Gebert, T. Zinner, P. Tran-Gia, D. Hock, M. Jarschel, and M. Hoffmann, "Heuristic approaches to the controller placement problem in large scale sdn networks," *IEEE Transactions on Network and Service Management*, vol. 12, no. 1, pp. 4–17, 2015.
- [31] A. Li, X. Yang, S. Kandula, and M. Zhang, "Cloudcmp: comparing public cloud providers," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*. ACM, 2010, pp. 1–14.
- [32] Z. Li, K. F. C. Yiu, and Z. Feng, "A hybrid descent method with genetic algorithm for microphone array placement design," *Applied Soft Computing*, vol. 13, no. 3, pp. 1486–1490, 2013.
- [33] C. Liang, Y. Guo, and Y. Liu, "Is random scheduling sufficient in p2p video streaming?" in *Distributed Computing Systems, 2008. ICDCS'08. The 28th International Conference on*. IEEE, 2008, pp. 53–60.
- [34] J. D. Little and S. C. Graves, "Little's law," in *Building intuition*. Springer, 2008, pp. 81–100.
- [35] J. Liu, M. Gong, Q. Miao, X. Wang, and H. Li, "Structure learning for deep neural networks based on multiobjective optimization," *IEEE transactions on neural networks and learning systems*, vol. 29, no. 6, pp. 2450–2463, 2018.
- [36] K. Mahmood, A. Chilwan, O. Østerbø, and M. Jarschel, "Modelling of openflow-based software-defined networks: the multiple node case," *IET Networks*, vol. 4, no. 5, pp. 278–284, 2015.
- [37] M. Obadia, M. Bouet, J.-L. Rougier, and L. Iannone, "A greedy approach for minimizing sdn control overhead," in *Network Softwarization (Net-Soft), 2015 1st IEEE Conference on*. IEEE, 2015, pp. 1–5.
- [38] C. Partridge, T. Mendez, and W. Milliken, "RFC 1546: Host anycasting service," *InterNet Network Working Group*, 1993.
- [39] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, J. Tan, Q. Le, and A. Kurakin, "Large-scale evolution of image classifiers," *arXiv preprint arXiv:1703.01041*, 2017.
- [40] T. P. Runarsson and X. Yao, "Stochastic ranking for constrained evolutionary optimization," *IEEE Transactions on evolutionary computation*, vol. 4, no. 3, pp. 284–294, 2000.
- [41] M. Srinivas and L. M. Patnaik, "Adaptive probabilities of crossover and mutation in genetic algorithms," *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 24, no. 4, pp. 656–667, 1994.
- [42] —, "Genetic algorithms: A survey," *computer*, vol. 27, no. 6, pp. 17–26, 1994.
- [43] M. T. I. ul Huque, G. Jourjon, and V. Gramoli, "Revisiting the controller placement problem," in *Local Computer Networks (LCN), 2015 IEEE 40th Conference on*. IEEE, 2015, pp. 450–453.
- [44] T. Wang, F. Liu, J. Guo, and H. Xu, "Dynamic sdn controller assignment in data center networks: Stable matching with transfers," in *Proc. of INFOCOM*, 2016.
- [45] D. Whitley, "A genetic algorithm tutorial," *Statistics and computing*, vol. 4, no. 2, pp. 65–85, 1994.
- [46] G. Yao, J. Bi, Y. Li, and L. Guo, "On the capacitated controller placement problem in software defined networks," *IEEE Communications Letters*, vol. 18, no. 8, pp. 1339–1342, 2014.