High Satisfaction and Fair Allocation of Resources in Software-Defined Data Center Networks

Chuanji Zhang and Douglas M. Blough
School of Electrical and Computer Engineering, Georgia Institute of Technology

***Abstract*— This paper studies how to fairly and efficiently allocate the two limited resources in software-defined data center networks (DCNs), namely the flow table and the control channel bandwidth. The problem considers routing path selection together with allocation of flow table entries and control channel bandwidths. The objective is to maximize the satisfaction ratio for flow groups in the network in terms of the two aforementioned resources, with the routing path optimized and different fairness constraints enforced. Our approach is to aggregate individual flows into flow groups and then find the optimal routing paths and the corresponding resource allocation vectors for each flow group. We study different fairness models in this work and also include a mechanism to relax the fairness constraint, which produces a range of solutions that permits a trade-off between total demand satisfaction and fairness.**

I. Introduction

Software-defined networking (SDN) has become a popular candidate for data center networks (DCNs). It enables flexible network resource management with the controller's global view. In SDN, the controller communicates with the SDN switches through the switch-to-controller link via southbound protocols and configures the network behaviors by installing flow entries on the flow tables of switches. However, the resource allocation problem in software-defined DCN exposes new design challenges. One of the main challenges is to cope with the limited resources brought by the unique architecture of SDN. The primary resources are the flow tables, which restrict the number of flow entries that can be stored on switches, and the capacity of the switch-to-controller links (henceforth referred to as control channels), which limits the message exchange rate between switches and the controller. As an illustrative example, in a typical data center, the number of active flows at the edge switch can reach as many as 8000 [1], but the flow table size in OpenFlow switches is limited to a few thousand entries due to power, cost and chip constraints [2]. As for the control channel capacity, an edge switch in DCN can see 100k new flows per second [1], which corresponds to a large traffic load on the control channel. However, the switch only supports limited control channel bandwidth due to its low internal control path bandwidth and limited CPU performance [3]. Thus, new resource allocation algorithms need to be investigated to consider these two unique constraints.

In general, high demand satisfaction and fairness are two fundamental objectives in resource allocation that cannot be maximized simultaneously. This motivates the investigation of inherent tradeoffs between the two objectives, where a common approach is to maximize demand satisfaction subject to some fairness constraints.

In this paper, we address the resource allocation problem in software-defined DCNs, with the objective to maximize the total satisfaction ratio of the aforementioned two SDN resources, subject to different fairness constraints. Our approach is to aggregate individual flows into flow groups and then find the optimal routing paths and the corresponding resource allocation vectors for each flow group satisfying our objective. Our main contributions are as follows:

- We address the resource allocation problem coupled with satisfaction maximization, fairness constraints, and routing path selection in software-defined DCNs. Instead of fixing the routing and allocating the resources subject to fairness constraints, we include routing path selection into the optimization problem to better meet the resource demands. The evaluation results show that deviating from simple shortest-path routing can improve the satisfaction ratio by 25%-30% while maintaining very good fairness.
- We investigate different fairness models, including classical max-min fairness, a simple max-min fairness, and an equal share fairness model. The fairness models are also modified to cooperate with our two-resource satisfaction maximization case. Moreover, to accommodate various network requirements, we introduce a relaxation parameter $\delta \in [0, 1]$ into these fairness models. It allows the network operator to control the trade-off between total demand satisfaction and fairness.
- We consider joint optimization of the two SDN resources, which leads to better network utilization compared with single-resource optimization.

II. Related work

Some research works propose to maximize the network flows or network throughput in SDN with limited flow table size [4], [5], without considering any fairness constraints among flows. The authors in [6] design a heuristic algorithm to maximize the minimum throughput satisfaction. That work considers only a simple max-min fairness model and does not allow tradeoffs between fairness and satisfaction. The aforementioned algorithms only consider the flow table size constraint while the control channel bandwidth is another important constraint for software-defined DCNs. The controller processing capacity constraint is added to the SDN throughput maximization problem in [5], but the control channel bandwidth limitation is still omitted. Moreover, these works focuses on maximizing network throughput with limited flow

978-1-5386-8088-9/19/$31.00 ©2019 IEEE

table space. However, in current DCNs, the data plane link utilization level is relatively low [1], and thus the throughput demands of users can be easily satisfied. On the other hand, the resource constraints brought by the unique architecture of SDN are likely to become the bottleneck. Optimization problems with the objective to improve SDN resource usage have not been well explored. In this paper, we present the fairness guaranteed satisfaction ratio maximization problem in terms of both the flow table usage and the control channel usage in software-defined DCNs.

III. System Model

We consider a scenario where the software-defined DCN is modeled as a graph $G = (V, E)$ and the flows from each server are aggregated into flow groups at the source switches based on their destination switches. The set of all aggregated flow groups $F = [f_1, f_2, \cdots, f_N]$ and f_n is defined as:

$$f_n = (s_n, t_n, D_b(n), D_{fe}(n), D_{sc}(n)),$$

where s_n and t_n is the source and destination switch of f_n, and $D_b(n), D_{fe}(n)$, and $D_{sc}(n)$ are the total data plane bandwidth demand, flow table demand, and control channel bandwidth demand requested by f_n, respectively. The measured statistical traffic characteristics over time or the total amount of resource purchased by the users can be used to form the demands of flow groups. The flow aggregation step will not only reduce the computation complexity of the optimization problem significantly, but also help reduce the flow table usage. While SDN enables individual flow management, this fine-grained control will cause flow table overflow problem. Therefore, coarse-grained control is used in some cases. For example, in the software-defined WAN deployed by Google, the flows are aggregated to groups defined by $\{src, dst, QoS\}$ for scalability [7]. In this work, we assume that with the flow aggregation, some flow rules (e.g., rules for the flows with the same source and destination address, and QoS class, or some mice flows that do not need full control) can be compressed together. The number of flow rules demanded after aggregation depends on the number of QoS classes enforced in the network, and level of control granularity the service providers prefer. The parameter $D_{sc}(n)$ is the control channel bandwidth required to set up the flow entries, collect the statistics, and maintain the flows. It can be estimated based on previous traffic statistics, the controller's functionality, or the users' purchased share.

Simple shortest-path routing may lead to unfavorable cases for resource sharing in the network. Besides, with SDN deployed, optimized routing becomes possible. Thus in this work, we explore the resource allocation problem with routing path selections to improve the optimization results. The number of paths connecting the sources and destinations can be of exponential size in network, but in practice, potential paths for routing the traffic are limited to a specific set and provided ahead of operation time. We thus consider a practical case where a set of pre-generated paths p for flow group f_n are provided as inputs. Let P_n be the set of p for f_n and the size of P_n is k.

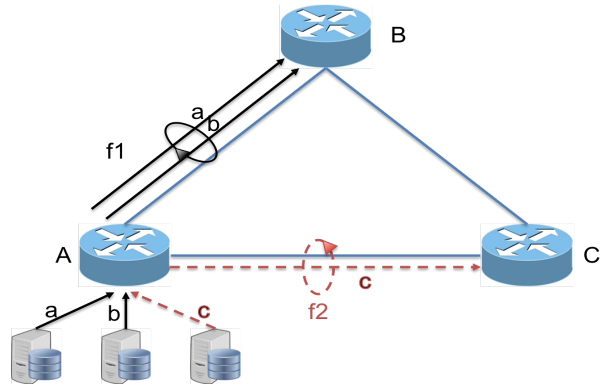

Fig. 1: Simple network model

Next, we present a simple example as illustration. Consider the scenario in Fig. 1, where three flows a, b and c are aggregated at switch A. Since flow a and b are destined at switch B, flow a and b are aggregated as flow group:

$$f_1 = (A, B, D_b(1), D_{fe}(1), D_{sc}(1)),$$

where $D_b(1), D_{fe}(1)$, and $D_{sc}(1)$ are the aggregated data plane link bandwidth, flow table size, and control channel bandwidth requested by flow a and b. Assuming the number of flow entries required by a and b is 1, if both flow a and b belong to the same QoS class and can be aggregated, $D_{fe}(1) = 1$. On the other hand, if a and b are from different QoS classes, $D_{fe}(1) = 2$. $D_{sc}(1)$ is the control channel bandwidth required to maintain flow a and b. Similarly flow c is aggregated as flow group:

$$f_2 = (A, C, D_b(2), D_{fe}(2), D_{sc}(2)).$$

Let $k = 2$, which is the maximum number in this case, $P_1 = \{(A, B), (A, C, B)\}$, and $P_2 = \{(A, C), (A, B, C)\}$. Our algorithm should return the optimal routing paths for f_1 and f_2 on P_1 and P_2, and the amount of resource allocated to each group to meet our objective. A feasible solution should satisfy the following capacity constraints:

- For each flow group f_n, the total resources assigned to it should be equal to or less than its demands.
- For each data plane link e, the total amount of flows going through it should not exceed $C_b(e)$, where $C_b(e)$ is the bandwidth capacity for link e
- For each node v, the total number of flow entries assigned to the flow groups should not exceed $C_{fe}(v)$, where $C_{fe}(v)$ is the maximum flow table size at v
- For each node v, the total traffic going through its control channel should not exceed its capacity $C_{sc}(v)$

IV. Satisfaction Maximization and Fairness

A. Satisfaction Maximization

Since maximizing the actual amount of resources might lead to unfair solutions where a high-demanded group only gets a small fraction of resources, we ensure a first level of fairness in the allocation problem by targeting on maximizing the ratio between allocation and demand for each group.

The satisfaction ratio of the flow table demand for flow group f_n, $X_n \in [0, 1]$, is defined as the ratio between the allocated flow table space for f_n on the selected path, and its demand $D_{fe}(n)$. As in our previous simple example, if $D_{fe}(1) = 2$ and the selected path for f_1 is (A, C, B), its maximum satisfaction ratio of the flow table demand, X_1, is 0.5 when the available flow table size on switch C is one entry. In this case, the controller can decide to install the entry with higher priority or other kinds of admission control mechanism can be used. The satisfaction ratio of control channel bandwidth demand, $Y_n \in [0, 1]$, is defined in a similar way. The satisfaction ratio vectors/resource allocation vectors of the two resources are $\vec{X} = [X_1, X_2, \cdots, X_N]$, and $\vec{Y} = [Y_1, Y_2, \cdots, Y_N]$. The overall satisfaction ratio of a flow group is the sum of the two individual ratio $(X_n + Y_n)$. We target on maximizing the total overall satisfaction ratio of all the flow groups. Additional constraints on the value of X_n and Y_n are set to avoid highly unbalanced solution as discussed in next section. Since the data plane link utilization level is relatively low in DCNs [1], we assume that the data plane bandwidth can always be well satisfied, and so we do not include it in the maximization objective.

Let the variable $x_n^p \in \{0, 1\}$ denote whether flow f_n goes through path $p \in P_n$ or not, and variable $y_e^p, z_v^p \in \{0, 1\}$ denote whether path p goes through link e and node v or not. The capacity constraints can be presented as:

$$\sum_{n \le N} \sum_{p \in P_n} x_n^p y_e^p \cdot 1 \cdot D_b(n) \le C_b(e) \qquad \forall e \in E \quad (1)$$

$$\sum_{n \le N} \sum_{p \in P_n} x_n^p z_v^p \cdot X_n \cdot D_{fe}(n) \le C_{fe}(v) \qquad \forall v \in V \quad (2)$$

$$\sum_{n \le N} \sum_{p \in P_n} x_n^p z_v^p \cdot Y_n \cdot D_{sc}(n) \le C_{sc}(v) \qquad \forall v \in V \quad (3)$$

$$\sum_{p \in P_n} x_n^p = 1 \qquad \forall n \le N \quad (4)$$

$D_{fe}(n)$ is defined as the number of flow entries required by the f_n, which should be an integer. Since we are considering aggregated flow groups in this work, the value of $D_{fe}(n)$ is at the level of hundreds and the flow table capacity is at the level of thousands. Allowing the value of $x_n^p X_n D_{fe}(n)$ to be fractional when solving the optimization problem, and rounding it to its nearest integer when deploying the network will not affect the performance significantly, but will improve the computation efficiency by a substantial amount.

We refer to the Pure Maximization case as the optimization problem with the objective to maximize the total satisfaction ratio in the network for all flow groups subject to the above resource constraints. However, since each flow group represents all the flows generated from a single switch, fairness should also be considered to avoid severely biased cases where some flow groups are poorly satisfied in the Pure Maximization problem. In order to address this issue, we consider different fairness models as constraints.

B. Fairness Models

1) Simple max-min fairness: We first consider a simple max-min fairness model. The allocation vectors $\vec{X}^*$ and $\vec{Y}^*$ is said to be simple max-min if $min(\vec{X}^*)$ and $min(\vec{Y}^*)$ is maximized among all the possible routings and allocations. Our objective is to seek the maximum satisfaction ratio among all the simple max-min guaranteed allocations. This optimization can be formulated as a mathematical problem composing of two sub-problems: the Max-Min problem to calculate the maximized minimum satisfaction for all flow groups, and the total Satisfaction Maximization problem. The Max-Min problem is defined as:

$$\begin{aligned} \text{maximize} \quad & \alpha + \beta \\ \text{subject to} \quad & Eq(1) - Eq(4) \\ & X_n \ge \alpha \quad \forall n \le N \\ & Y_n \ge \beta \quad \forall n \le N, \end{aligned}$$

and the Satisfaction Maximization problem is:

$$\begin{aligned} \text{maximize} \quad & \sum_{n \le N} X_n + \sum_{n \le N} Y_n \\ \text{subject to} \quad & Eq(1) - Eq(4) \\ & C1 : X_n \ge \delta \cdot \alpha^* \quad \forall n \le N \\ & C2 : Y_n \ge \delta \cdot \beta^* \quad \forall n \le N. \end{aligned}$$

Instead of setting $X_n + Y_n \ge \alpha + \beta$, we tighten the constraints by setting $X_n \ge \alpha$ and $Y_n \ge \beta$. This avoids the highly suboptimal solutions of very large X_n with very small Y_n or vice versa. Assume the maximized minimum satisfaction ratio obtained in the Max-Min problem $(\alpha + \beta)^*$ is equal to $\alpha^* + \beta^*$. Constraints $C1$ and $C2$ in the Satisfaction Maximization problem enforce the fairness constraint by making sure that each flow group has a minimum overall satisfaction ratio at $\delta(\alpha+\beta)^*$. The fairness relaxation parameter $\delta \in [0, 1]$ allows controlled tradeoffs between satisfaction and fairness. When $\delta = 1$, the solution leads to perfect simple max-min fairness with the minimum satisfaction ratio maximized. The fairness constraint is relaxed as δ decreases and when $\delta = 0$, the optimization problem corresponds to the Pure Maximization case with no fairness constraint.

2) Classical max-min fairness: The classical max-min fairness is said to be achieved by an allocation if and only if the allocation is feasible and an attempt to increase the allocation of any participant necessarily results in the decrease in the allocation of some other participant with an equal or smaller allocation. The classical max-min fair allocation can be obtained by a progressive filling algorithm where the allocation for each user starts from 0. Since our objective is to achieve fairness in terms of the satisfaction ratio with two resources, we propose a slightly different algorithm than the conventional progressive filling method as described in Algorithm 1. The algorithm seeks classical max-min for the flow table satisfaction ratio and control channel satisfaction ratio separately. The evaluation results show that this method can actually achieve a high and fair allocation for both resources at the same time. In Algorithm 1, parameters R_{fe}, R_{sc} denote the unsaturated switches in terms of flow table and control channel, and Q_{fe}, Q_{sc} denote the non-bottlenecked flows for the two resources. $U_{fe}(v), U_{sc}(v)$ and $T_{fe}(v), T_{sc}(v)$ are the current utilization level and total demands of non-bottlenecked flows on switch v, respectively. Since our optimization objective is the satisfaction ratio, we define the most bottlenecked switch v_{fe}, v_{sc} based on $T_{fe}(v), T_{sc}(v)$, instead of the number of flows on switch v as in the conventional progressive filling algorithm. Besides, the increment of X_n, Y_n is also based on satisfaction ratio.

Algorithm 1 Classical Max-min Fairness Algorithm with Two Resources

Input: Feasible routing path assignment x_n^p
Output: $\vec{X}, \vec{Y}$

Initialize :$\vec{X} \leftarrow \vec{0}, \vec{Y} \leftarrow \vec{0}, R_{fe} \leftarrow V, R_{sc} \leftarrow V, Q_{fe} \leftarrow F, Q_{sc} \leftarrow F, U_{fe} \leftarrow \vec{0}, U_{sc} \leftarrow \vec{0}$;
Initialize: $T_{fe}(v) = \sum_{f_n \in Q_{fe}} \sum_{p \in P_n} x_n^p z_v^p D_{fe}(n)$
Initialize: $T_{sc}(v) = \sum_{f_n \in Q_{sc}} \sum_{p \in P_n} x_n^p z_v^p D_{sc}(n)$
while $Q_{fe} \neq \emptyset$ or $Q_{sc} \neq \emptyset$ **do**
  $v_{fe} = \arg\min_v \frac{C_{fe}(v) - U_{fe}(v)}{T_{fe}(v)}, \forall v \in R_{fe}$
  $v_{sc} = \arg\min_v \frac{C_{sc}(v) - U_{sc}(v)}{T_{sc}(v)}, \forall v \in R_{sc}$
  $X_n + = min(\frac{C_{fe}(v_{fe}) - U_{fe}(v_{fe})}{T_{fe}(v_{fe})}, min_{\forall f_n \in Q_{fe}}(1 - X_n))$ $, \forall f_n \in Q_{fe}$
  $Y_n + = min(\frac{C_{sc}(v_{sc}) - Usc(v_{sc})}{T_{sc}(v_{sc})}, min_{\forall f_n \in Q_{sc}}(1 - Y_n))$, $\forall f_n \in Q_{sc}$
  $U_{fe}(v) = \sum_{n \leq N} \sum_{p \in P_n} x_n^p z_v^p \cdot X_n \cdot D_{fe}(n)$
  $U_{sc}(v) = \sum_{n \leq N} \sum_{p \in P_n} x_n^p z_v^p \cdot Y_n \cdot D_{sc}(n)$
  $R_{fe} \leftarrow \{v | C_{fe}(v) - U_{fe}(v) > 0\}$
  $R_{sc} \leftarrow \{v | C_{sc}(v) - U_{sc}(v) > 0\}$
  $Q_{fe} \leftarrow \{f_n | f_n$ spans only on $v \in R_{fe}$ and $X_n < 1\}$
  $Q_{sc} \leftarrow \{f_n | f_n$ spans only on $v \in R_{sc}$ and $Y_n < 1\}$

Furthermore, in our case, the resource allocation is coupled with routing path selection and thus the solution cannot be obtained using only the progressive filling algorithm, where the routing paths should be provided as inputs. To solve this problem, a brute-force method which evaluates all possible routing paths and selects the best one can return the optimal solution, but since brute-force is not scalable in a large network, we propose a heuristic approach to approximate the brute-force method. With a proper initial routing path, the total satisfaction is calculated subject to the classical max-min fairness constraint. Next, rerouting of the bottlenecked flows is considered. The possible reroutings are evaluated greedily by starting with the flows at the most bottlenecked switch. To reduce the computation complexity, we only consider rerouting each flow once and set a limit on how many reroutings will be explored in each run. Finally, the routing path selections x_n^{p*} and the allocation vectors $\vec{X}^*, \vec{Y}^*$, which generate the highest total satisfaction ratio with classical max-min fairness, are selected.

To relax the fairness constraint, after we get the optimal routing path selection x_n^{p*} and the allocation vectors $\vec{X}^*, \vec{Y}^*$, we allow the actual satisfaction ratio for f_n (X_n and Y_n) to be in the range of $[\delta X_n^*, \frac{X_n^*}{\delta}]$ and $[\delta Y_n^*, \frac{Y_n^*}{\delta}]$, where $\delta \in [0, 1]$. The maximum satisfaction ratio and the corresponding routing paths can be obtained with the following optimization problem:

$$\begin{aligned}
\text{maximize} \quad & \sum_{n \leq N} X_n + \sum_{n \leq N} Y_n \\
\text{subject to} \quad & Eq(1) - Eq(4) \\
& \delta X_n^* \leq X_n \leq \frac{X_n^*}{\delta} \quad \forall n \leq N \\
& \delta Y_n^* \leq Y_n \leq \frac{Y_n^*}{\delta} \quad \forall n \leq N.
\end{aligned}$$

3) Equal Share: Another fairness model considered in this work is strictly equal share, where all the flow groups should achieve the same level of satisfaction, which leads to $X_1 = X_2 = \cdots X_N = \alpha$ and $Y_1 = Y_2 = \cdots Y_N = \beta$. This allocation will achieve optimal Jain's fairness index [8]. The maximum satisfaction ratio with equal share fairness guaranteed can be obtained by solving the following optimization problem:

$$\begin{aligned}
\text{maximize} \quad & \alpha + \beta \\
\text{subject to} \quad & Eq(1) - Eq(4) \\
& C3 : \delta\alpha^* \leq X_n \leq \frac{\alpha^*}{\delta} \quad \forall n \leq N \\
& C4 : \delta\beta^* \leq Y_n \leq \frac{\beta^*}{\delta} \quad \forall n \leq N,
\end{aligned}$$

and setting δ to be 1. Parameter $\delta \in [0, 1]$ allows control of how tightly the solution achieves the fairness objective.

V. Performance Evaluation

The network topology used for performance evaluation is a fat-tree topology [9], which has been widely used in DCNs. The topology contains 4 core switches, 8 aggregation switches, and 8 edge switches, which can represent a typical scale of a campus data center [1]. Each edge switch is connected to hundreds of servers and the flows are aggregated into groups such that, for each edge switch, there is one flow group destined to every other edge switch. The bandwidth of the data plane links are 10Gbps and the flow tables of switches are limited to 2000 entries. Multiple connections between switch and controller are proposed in OpenFlow 1.3 and analyzed in [10]. Thus, we assume the control channel bandwidth is 2Gbps for the core switches, and 1Gbps for the aggregation and edge switches. The individual demands are generated randomly according to a uniform distribution and all the results are obtained based on eight random trials.

The optimization problems are solved using Gurobi 8.0.1 on a 4-core 3.2GHz Intel i5-6500 processor with 7.7 GB memory. The average running times for the joint optimization cases in Section V-A with simple max-min and equal share fairness model are $0.93s$ and $0.20s$, respectively. Due to multiple iterations in the classical max-min fairness algorithm, its average running time is $70.2s$. The traffic in a DCN is relatively stable on the timescale of a few seconds up to a few minute [11]. In addition, only the amounts of the three demands need to be collected for calculation, which will incur only a very small overhead in the network. Given these conditions, our algorithm can be run periodically in an online fashion to optimize the network resource allocation.

A. Satisfaction Ratio

We first compare the satisfaction ratio in terms of the two SDN resources under different optimization scenarios as shown in Fig. 2. The total demands for the two SDN resources are set to be larger than the network capacity with

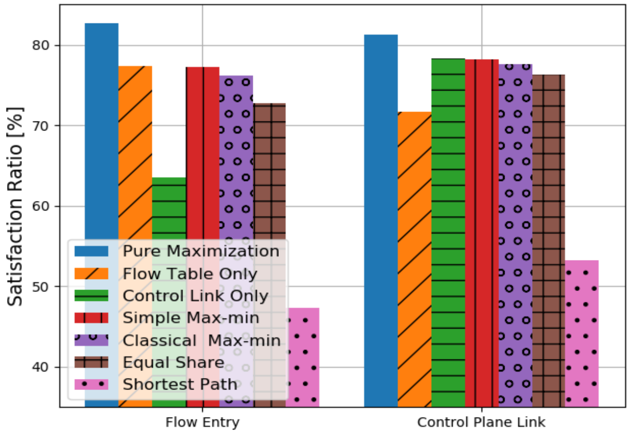

Fig. 2: Satisfaction Ratio Comparison

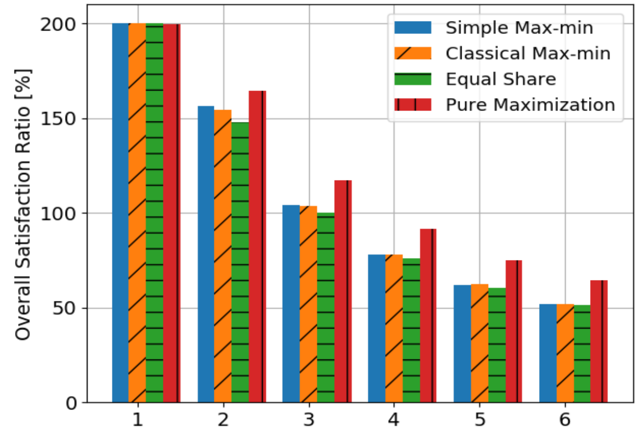

Fig. 3: Satisfaction Ratio with Varying Demand Scaling

the average demand being 200 flow entries and 100Mbps bandwidth per group. The optimization scenarios include: 1) the Pure Maximization case with no fairness constraint, 2) the optimization cases for a single resource (Flow Table Only or Control Link Only) with simple max-min fairness, 3) the jointly optimized cases for both resources with different fairness constraints, and 4) the optimization case with Shortest Path routing and simple max-min fairness.

First, as expected, the Pure Maximization case achieves the highest satisfaction ratio for both resources. However, it will generate undesirable solutions where some of the flow groups get very low satisfaction ratio. On average, 7% of the flow groups get 0% satisfaction ratio for flow table space and 13% of the flow groups get 0% satisfaction ratio for control channel. To make things worse, the flow groups receiving no flow entry resource are not identical to the flow groups receiving no control plane resource. Thus, some resources are wasted with pure maximization since a flow group cannot be routed successfully with only one SDN resource.

Second, utilizing optimized routing shows substantial improvement compared with the Shortest Path case. Shortest path routing is 30% and 25% worse compared with the optimized routing case with simple max-min fairness enforced, in terms of flow table and control channel satisfaction.

Lastly, when the allocation of a single resource is set as the optimization target (Flow Table Only and Control Link Only), the satisfaction ratio of the targeted resource is slightly better than the corresponding jointly optimized case, but it causes around 10% satisfaction degradation on the other resource. This fact emphasizes the need to jointly optimize the allocation of the two resources. The jointly optimized cases with all three fairness objectives produce well-balanced allocation of the two resources, and achieve almost as good a satisfaction ratio as the targeted resource in the single-resource optimization cases.

Next, we will elaborate on the difference between the fairness models and the impacts of fairness relaxation. Since the satisfaction ratios of the two SDN resources can be optimized jointly with all three fairness models as shown in Fig. 2, we only show the overall satisfaction ratio (the sum of the two individual ratios) in the following analyses.

B. Comparison of Fairness Models

As the traffic in DCNs has grown dramatically in recent years and there is an increasing need to manage the traffic with finer-grained control, the SDN resources might not be sufficient to satisfy the demands of all flow groups. We scale up the demands from where the network resources are sufficient to where resources are insufficient and resource allocation optimization is required, and study the performance using the different fairness models. The overall satisfaction ratios achieved under different demand scaling are shown in Fig. 3. A demand scaling of x means x times the baseline demand (100 flow entries and 50Mbps bandwidth per group on average). When demand scaling is low, all models achieve the maximum satisfaction ratio of 200%. As demand scaling increases, the overall satisfaction ratio decreases, and the pure maximization model achieves higher satisfaction than the other approaches. However, as discussed before, the solution to the pure maximization is undesirable due to unfairness and resource waste. For the other three fairness models, the simple max-min model yields the highest satisfaction ratio while the equal share model generates the lowest at first, but as the demand scaling keeps increasing, the benefits of the simple max-min over the classical max-min diminishes and finally, all the three models achieve similar satisfaction ratios. Based on the individual satisfaction ratio of the flow groups, we find that with high demands in the network, the three fairness models tend to generate the same allocation as discussed below.

To compare the fairness approaches more closely, we investigate the satisfaction ratios for each flow group when demand scaling is 3. The allocation for flow table resource is similar as for control channel resource, so we only show the allocation of flow table satisfaction in Fig 4. Pure maximization causes a severe bias where some groups get 0% satisfaction, which means the access to the network for these flow groups is denied, while other groups get perfect satisfaction. With equal share model, all flow groups obtain equal satisfaction ratio for the two resources as expected. For the two max-min models, the majority of flow groups obtain the same satisfaction ratio as with equal share fairness. A few groups get higher satisfaction than others and thus a

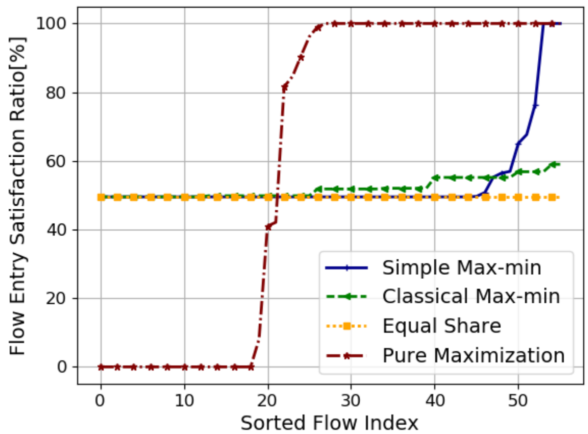

Fig. 4: Flow Table Satisfaction Ratio Allocation

TABLE I: Jain's Fairness Index of Flow Table Demand

δ	0	0.2	0.4	0.6	0.8	1
Simple max-min	0.53	0.54	0.55	0.57	0.60	0.66
Classical max-min	0.53	0.54	0.56	0.59	0.68	0.77
Equal share	0.53	0.54	0.55	0.60	0.70	1.0

higher total satisfaction ratio is achieved than with the equal share model. With the classical max-min model, the flow groups are partitioned to several clusters (usually the number of clusters is very small) and within each cluster, all flow groups achieve exactly the same satisfaction. As the demand scaling increases, the allocation of both the max-min models converges to the equal share model.

C. Fairness Relaxation

Finally, the impacts of the fairness relaxation parameter δ are studied. When $\delta = 0$, the optimization problems with different fairness objectives all become a pure maximization problem. We show the overall satisfaction ratio degradation compared with the pure maximization when demand scaling is 3 with varying δ in Fig. 5. As the fairness requirement is tightened, the achieved satisfaction ratio decreases. When $\delta < 0.6$, the degradation of the three models is the same. As δ keeps increasing, simple max-min fairness tends to achieve the highest satisfaction ratio. With perfect fairness, the performance degradations of simple max-min, classical max-min, and equal share are 10%, 12%, and 17%, respectively.

We also investigate Jain's fairness index of the flow table satisfaction ratio as displayed in Table I. Jain's fairness index of the control channel resource follows a similar trend so we don't show it here. When the fairness requirement is relaxed ($\delta < 0.6$), the three models all produce the same fairness index. Combined with the satisfaction ratio results, we can infer that with relaxed fairness requirement, the three models produce identical results. As δ increases, while the simple max-min model leads to the highest satisfaction ratio, it also induces the lowest fairness index. The equal share model will reach the optimal Jain's index when $\delta = 1.0$. In summary, performance and fairness tradeoff can be achieved by choosing the proper δ and fairness model. A service provider can determine the best fairness possible for a given minimum performance threshold or determine the best performance for a given minimum fairness threshold.

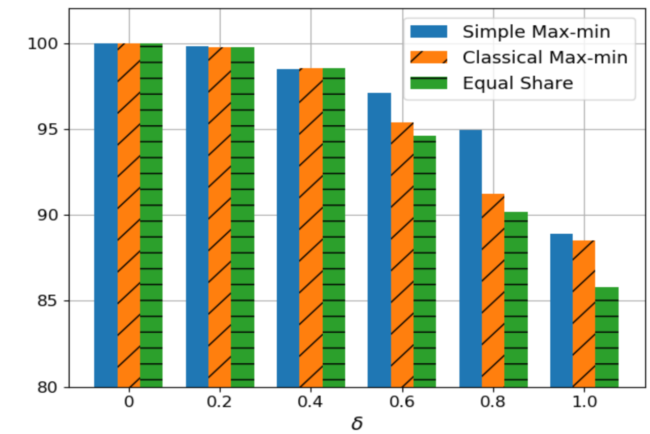

Fig. 5: Achieved Overall Satisfaction Ratio Compared with Pure Maximization [%]

VI. Conclusion

In this paper, we addressed the high satisfaction and fair allocation of resources in software-defined DCNs, including routing path selection. We jointly optimized the allocation of flow table and control channel resources with different fairness constraints. The fairness constraints were compared and a mechanism to relax these constraints was studied. Our future work will consider the problem of delay-guaranteed fair allocation of resources in software-defined DCN.

References

[1] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th SIGCOMM Conference on Internet Measurement*. ACM, 2010, pp. 267–280.

[2] B. Stephens, A. Cox, W. Felter, C. Dixon, and J. Carter, "Past: Scalable Ethernet for data centers," in *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*. ACM, 2012, pp. 49–60.

[3] A. R. Curtis, J. C. Mogul, J. Tourrilhes, P. Yalagandula, P. Sharma, and S. Banerjee, "Devoflow: Scaling flow management for high-performance networks," *ACM SIGCOMM Computer Communication Review*, vol. 41, no. 4, pp. 254–265, 2011.

[4] R. Cohen, L. Lewin-Eytan, J. S. Naor, and D. Raz, "On the effect of forwarding table size on SDN network utilization," in *INFOCOM, 2014 Proceedings IEEE*. IEEE, 2014, pp. 1734–1742.

[5] G. Zhao, L. Huang, Z. Yu, H. Xu, and P. Wang, "On the effect of flow table size and controller capacity on SDN network throughput," in *Communications(ICC), IEEE Int'l Conf. on*. IEEE, 2017, pp. 1–6.

[6] J. Zhang, D. Zeng, L. Gu, H. Yao, and Y. Fan, "On rule placement for multi-path routing in software-defined networks," in *International Conference on Collaborative Computing: Networking, Applications and Worksharing*. Springer, 2015, pp. 59–71.

[7] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu *et al.*, "B4: Experience with a globally-deployed software defined WAN," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 3–14, 2013.

[8] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Eesign, Measurement, Simulation, and Modeling*. John Wiley & Sons, 1990.

[9] 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.

[10] C. Zhang, H. Yang, G. Riley, and D. Blough, "Queueing analysis of auxiliary-connection-enabled switches for software-defined networks," to appear in Computing, Networking and Communications (ICNC), 2019 Int'l Conf. on. IEEE, 2019.

[11] A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren, "Inside the social network's (datacenter) network," in *ACM SIGCOMM Computer Communication Review*, vol. 45, no. 4. ACM, 2015, pp. 123–137.