

Max Progressive Network Update

Yang Chen and Jie Wu

Temple University, USA

Email: {yang.chen, jjewu}@temple.edu

Abstract—Datacenter network always needs to be updated so that it can adapt to frequent traffic changes for high network utilization. However, the quality of service can be easily reduced by unexpected transient congestion, which may be caused by variable link transmission latency during flow migration. Furthermore, when network administrators reallocate bandwidth resource without taking flow path information into consideration, deadlocks among flows and link resources unavoidably block the updating process. In this paper, we prove it is NP-hard to find a feasible updating plan with the minimum number of rate-limiting flows when deadlocks occur. We propose a buffer-assisted schedule, Max Progressive Updating Method (MAPUM), to update network in a consistent way. Extensive simulations prove that MAPUM accomplishes network updating in a lossless and efficient way.

Index Terms—Network updates, SDN, consistent, congestion.

I. INTRODUCTION

Temporary network updating disruptions may occur with many unexpected events like packet loss and blackholes. The reasons for disruptions include straggling of switch rule updating, controller disconnection, and imperfection in clock synchronization. To prevent anomalies, researchers pay extensive attention to the network updating problem in SDNs. There are three key challenges for network updating issues: *optimality*, *consistency*, and *swiftness* [1]. An optimal migration requires that the final routing configuration after the update should be the same as the given target state, which is derived by solving the multi-commodity flow problem [2]. In SDNs, a centralized controller with the global view can guarantee optimality by directly establishing the targeted network configuration. A consistent migration of flows ensures the congestion-free, loop freedom and no temporary demand reduction properties. A swift update prevents the new routing setup from becoming obsolete due to frequent changing network conditions. Swiftness is more critical in data centers, where network updating serves as usual maintenance.

However, current update schedules suffer different downsides. We illustrate the unfeasibility of the current methods in Fig. 1. There are three flows f_1 , f_2 , and f_3 . Each flow's size is 1 unit, and each link has a capacity of 1 unit. We need to migrate the flows' paths from the initial state in 1a to the final state in 1b. Only when a flow's final path has enough available bandwidth, it can be consistently migrated. However, it is obvious none of these flows can be updated to their final paths because other flows' initial paths are occupying their resources. When various flows' initial paths are occupying other flows' final state paths, none of them can be updated, we call this a deadlock. Flows f_1 , f_2 and f_3 form a deadlock.




Fig. 1: A network update example

Current works neglect handling this situation, or just randomly stub out flows to vacate competing links. For example, SWAN [3] can only accomplish the update when leisure capacity exists along all links of the flow's path. In this example, all the links with flows going through are saturated and have no free bandwidth. SWAN will come up with no solution. Dionysus [4] generates a dependency graph and finds all SCCs (strong connected components). When deadlocks exist, it tries to randomly rate limit a few flows until all deadlocks break. This solution causes packet loss, which will reduce the quality of service. Additionally, the link capacity reservation method [3] has so little flexibility that bandwidth utilization is always low. The intermediate state involvement method [5] increases the number of updating steps to avoid link congestion due to flow mixture in different network states.

In this paper, we focus primarily on how to generate a consistent updating plan of breaking all of the deadlocks with the help of the switch buffer, a kind of resource ignored in previous researches. We present a heuristic algorithm, Max Progressive Updating Method (MAPUM), to address this problem. With the help of adequate amount of buffer, the deadlocks where interwoven flows compete for bandwidth can be broken and we can update the network in a lossless manner. Deliberately deploying buffers in switches further shortens the time that it takes to migrate the flows. We also investigate the complexity of the deadlock-breaking problems mentioned above and demonstrate that they are NP-hard.

The main contributions of our work are in the following:

- 1) We summarize current network updating methods and analyse their advantages and disadvantages.
- 2) We propose a buffer-assisted strategy in order to update network in a consistent way even when there are complex deadlocks. We also prove that it is NP-hard

to allocate buffers' sizes and locations to resolve all deadlocks with minimum effect on the network. Then a heuristic solution is introduced to efficiently arrange buffers in the switches.

- 3) We demonstrate the significant advantages of our approach in simulations compared to current works.

The remainder of this paper is organized as follows. Section II surveys related works. Section III describes the model, and problem formulation. Section IV talks about generating resource dependency graph. Section V analyses the problem and proposes our MAPUM solution. Section VI includes the experiments. Finally, Section VII concludes the paper.

II. RELATED WORK

When it comes to the routing reconfiguration problem, there are two basic mainstream methods: ordering [6, 7] and two phase [8–10]. The ordering method usually updates the forwarding table in the switches one-by-one in a specified order. The order is carefully calculated to preserve certain required properties, for example, loop-free, and blackhole-free. However, this order might not guarantee both forwarding and policy demand. The two-phase updating scheme installs both the initial and final rules on all switches, and tags packets to signal, which rule should be applied. This method ensures the success of an updating, but it doubles the number of rules on every switch, which wastes expensive and power-hungry TCAM memory resource. In this paper, we perform the two-phase commit using version numbers for updating rules to maintain packet coherence.

Due to constantly changing traffic demands, data center network updates occur frequently, whether triggered by the operators, applications, or sometimes, even failures. The inherent asynchrony will lead to over-utilization of links, inducing congestion and packet loss. One might just ignore these effects, hoping things will get better on their own, but delays and loss of data are not desirable, especially in real-time applications. Hence, recent development has considered consistent migration, i.e., congestion-free and without temporary demand reduction [1, 3, 5, 11]. The paper [12] does not strictly require that the updating process be lossless. Rather, it aims to obtain time-efficiency when the loss of packets is minimized.

We can roughly classify updating strategies into three categories: link reservation [3], intermediate state involvement [5], and time-awareness [1, 13–15]. The approach of SWAN [3] has two parts: First, if a fraction s of capacity (*slack*) is guaranteed to be free on each link for the old and new flow, the network can be updated in $\lceil 1/s \rceil - 1$ steps. Second, in order to solve the problem optimally, SWAN uses linear programming to check if a solution with x steps exists. However, when there is no slack on some edges, the algorithm is unlikely to halt in certain steps with a high computation complexity. ZUpdate [5] attempts to compute and execute a sequence of steps to progressively meet the end requirements from an initial traffic matrix and a traffic distribution. To safely migrate flows, these representative update scenarios need to be guaranteed that there will be no congestion. But applying




Fig. 2: Scheduling Process

intermediate states will stretch the update time, which means the traffic chaos caused by the migration will last longer. Time-awareness consistent updating strategy [13] utilizes time-triggered network updates to achieve consistency by either ordering or two-phase commit method. This scheme asks too much of time synchronization. With one straggling switch, the whole process can cause serious congestion or other problems.

III. MODEL AND FORMULATION

Our paper studies the network updating problem: given the initial and final network states and the properties that must be preserved according to the network administrator, we need to find a feasible and efficient solution to consistently migrate flows with the help of switch buffer, especially when the network has a lot of deadlocks among flows in different states. Assume the initial and final states are both valid, and the flows are unsplitable. Consistency requires that the updating be loop-free, drop-free, and congestion-free. We propose an efficient approach, MAPUM, which focuses on handling deadlocks with the help of proper buffer location and size allocation. Our scheduling process is shown in Fig. 2.

A. A Motivating Example

Consider the same example in Fig. 1. As discussed above, there does not exist a consistent update plan. But with the proper allocation of buffer and careful selection of flows that will be temporarily buffered in switches, we can update the network to its final state consistently. Now let's think of the most trivial plan. Assume that the migration is flow-based; this means that the flow is allowed to update only when all the links along its path are available. We only rate limit the flows in their source switches' buffer. And rate limiting means to reduce the flow's rate to 0. We list all the feasible updating plans, shown in Tab. I. Take rate-limiting flow f_1 as an example. At time slot 0, we rate limit f_1 . After 2 slots, f_2 can migrate because of the available bandwidth of AB . At slot 4, f_3 can move to its final path. At time slot 5, f_1 migrates and the updating is finished. The buffer size is 5. There are, in total, five methods to consistently break all the deadlocks. The minimum number of rate-limiting flows is only one. The minimum time is 4 and the minimum buffer size is 5.

Intuitively, if a flow's initial path include too many congested link resources, or its final path consists of less bottleneck links, it should be buffered in order to allow others to migrate. It is also better to limit the flows that are engaged in

TABLE I: Rate-limiting Plan

Combination of flows	f_1	f_1 and f_2	f_1 and f_3	f_2 and f_3	f_1, f_2 and f_3
# rate-limiting flows	1	2	2	2	3
Update time	5	4	5	7	3
Buffer size	5	5	9	12	7

a higher number of deadlocks, because it is more beneficent in resolving all deadlocks. Furthermore, we can utilize buffer to accelerate the updating. Consider the following situation. If flow f_1 's initial path occupies f_2 's final path resources, f_2 occupies f_3 's, ..., and f_{n-1} occupies f_n 's, we need to first update f_1 , then f_2 , ..., and finally, f_n , which definitely costs a lot of time to finish the update. With the help of buffer, we can select a middle flow to be buffered so that those two parts of the updating process can be done concurrently, significantly shortening the time needed.

B. Network Model

The network G consists of a set of switches S and a set of directed links L . A flow f is from a source switch s_i to a destination switch s_j with traffic volume t_f . Links along the path of f are defined as $L_f = \{l_{i,j} | l_{i,j} \in p_f\}$, where p_f is the forwarding path of flow f . The network state N is the combined state of all flows, i.e., $N = \{f | f \in F\}$. The network state includes each flow's path and its allocated bandwidth. We use two phase commit to update each switch's routing table information. At every hop, the switch checks the packet's header to decide its outgoing port. The controller needs to manage the time to tag the packets with the new version number.

Network, graph, capacity: We define a network as a simple directed graph with edge capacities.

Definition 1: Let $G = (V, E)$ be a simple connected directed graph with $n = |V|$ nodes representing the switches, and $m = |E|$ edges representing the links. We denote the set of outgoing edges (u, v) of a node $v \in V$ by $out(v)$ and the set of incoming edges (u, v) by $in(v)$. A network N is a pair (G, c) , with $c : E \rightarrow \mathbb{R}^+$ being a function assigning each edge $e \in E$ a capacity of $c(e)$.

Flow: Next, we define an unsplittable flow according to the buffer-involved flow constraints, i.e., demand satisfaction, flow conservation, and capacity constraints. As is common in this context, we only consider self cycle-free flows in this paper.

Definition 2: Let $N = (G, c)$ be a network. A map $f : E \rightarrow \mathbb{R}^+$ is called an unsplittable flow from s to t if it is cycle-free and satisfies the buffer-involved flow constraints, i.e., $\forall v \in V \setminus \{s, t\}$:

$$f(e)_{e \in in(v)} = f(e)_{e \in out(v)} + buffer(v, f), \quad (1)$$

$$f(e)_{out} = d_f = f(e)_{in}, \quad (2)$$

$$\forall e \in E : \sum_f f(e) \leq c(e) \quad (3)$$




Fig. 3: Resource Dependency Graph

with d_f being the size of f and the edges with $f(e) > 0$ forming the path of f from s to t . All these unsplittable flows form the set $F = \{f_i, i = 1, 2, \dots, m\}$ if there are m flows.

Initial and final network state: In this paper we only consider network updates for flows, i.e., given a set of old forwarding rules for some flow f , we want to change into a set of new forwarding rules for another flow f' .

Definition 3: Let N be a network, and let F and F' be flows in N . A network update is a triple (N, F, F') .

Atomicity of network updates: We assume that the change from a flow f to a flow f' , both from s to t , is performed as an atomic operation on the source router s . In practice, this can be achieved by a two-phase protocol [9]. All forwarding rules in the network for F' are installed by the SDN controller first. When these installations are confirmed, the ingress router s will start tagging all packets from F to F' . Note that with this method, if the latency from s to some node in the network is the same for F and F' , the flow F' will arrive after F has already departed.

IV. DEPENDENCY GRAPH GENERATION

The primary challenge is to tractably explore relationships among flows and resources in the initial and final states. In this paper, we leverage a dependency graph to describe the update orderings. As shown in Fig. 2, the dependency graph generator takes the initial state NS_i , the final state NS_f , and the consistency property as input. The dependencies fundamentally limit how quickly the SDN can be updated. The goal of dependency graph generation is to inter-link these operations based on the consistency property. There are two types of nodes in our dependency graph: link nodes and flow nodes. Link nodes represent link resources and are labelled with the amount of current residue capacity. Flow nodes correspond to different flows marked with demands. An edge from a link node to a flow node means that the flow needs this link resource to update, and the weight of the edge shows the delay after the resource becomes available. The link transmission delay is taken into consideration.

To simplify, we measure the delay in the units of the updating step. $late_{l_j, f_i} = m$ if the link l_j is the m_{th} link that flow f_i passes in the final state. $late_{f_i, l_j} = m$ if the link l_j is the m_{th} link that flow f_i passes in the final state. Algorithm 1 shows the details of generating a dependency graph. Take Fig. 3, the dependency graph of Fig. 1, as an example. There are three deadlocks in the dependency graph. We can clearly

Algorithm 1 Dependency Graph Generator

Input: The initial and final network State N_{ini} and N_{fin} , each flow i 's demand dem_{f_i} , and each link's capacity $c_j \in C$;

Output: Dependency Graph G ;

- 1: **for** each link l_j in the network **do**
 - 2: add a new node with its residue capacity $res_{l_j} = c_{l_j} - \sum dem_{f_i}$ in the initial state N_{ini} , p_{f_i} passes link l_j ;
 - 3: **for** each flow $f_i \in F$ **do**
 - 4: add a new flow node with its demand dem_{f_i} ;
 - 5: **for** each link along the initial (and final) path of flow f_i **do**
 - 6: add a directed edge from the flow node to the link node (in the final path, the direction is reversed);
 - 7: label the edge with the transmission latency $late_{f_i, l_j}$ (or $late_{l_j, f_i}$ in final st) in units of step;
 - 8: **return** Dependency Graph G ;
-

see that f_1 involves all these three deadlocks, which means that rate limiting f_1 will make the updating feasible. So the minimum number of flows to be rate limited is 1. However, this does not demonstrate that this method is the fastest plan or that it needs the least amount of buffer size.

V. CONSISTENT UPDATING STRATEGY WITH BUFFER ASSISTANCE

A. The hardness of the buffer-assisted schedule

To ensure a consistent updating plan, we introduce buffer to strategically break the cycles in the dependency graph. Buffer scheduling is a resource allocation problem: how can we allocate fewer buffer resources to accomplish a consistent network update.

Theorem 1: Without limitations on buffer size or location, a consistent network update can be accomplished in one step.

Proof: First, we block all the flows and wait for a sufficient period of time. When there is no residue traffic in the network, we can migrate flows to their new paths with adequate bandwidths. Because the initial and final states are valid, there will not be any inconsistencies like loops, blackholes, or congestions. \square

Buffer is expensive and has limited size. The above situation is an extreme condition. Actually, it is possible to strategically migrate flows in a specific order or to limit a few of them to achieve the consistent update. As a result, we are required to choose certain flows to be rate limited from all the flows that need to be migrated. Moreover, we ought to decide where to begin buffering and what percentage of the flow to be buffered would be best according to the Quality of Service (QoS) requirement. After making a scheduling decision, we need to select the next flow until all the deadlocks have been resolved. This is a complex resource allocation problem. We can prove the complexity of handling deadlocks in network update with the help of buffer.




Fig. 4: Deadlocks in dependency graph

Theorem 2: Even in the absence of partial flow limiting, finding a feasible update schedule with a minimum number of rate limited flows with deadlocks is NP-complete.

Proof: Given a network's initial and final states, there exists several deadlocks among links in the dependency graph. Each deadlock is a cycle consisting of multiple flows. Each flow can be involved in several deadlocks, because it must wait for multiple link resources. Suppose there are n deadlocks in the graph G and the general collection of deadlocks is U . In total, m flows become involved in G and the i^{th} flow ($i = 1, \dots, m$) is able to break the set $S_i \subseteq U$. The goal is to break all the deadlocks by cutting off the fewest number of flows using buffer. So we reduce the original problem into the so-called set cover problem, an NP-complete problem that covers all elements using the fewest number of sets.

B. Rate Limiting Flow Selection

Because of the NP-complete nature of the problem, we propose a heuristic algorithm to find a consistent updating plan by selectively limiting flows as seldom as possible.

1) *Flow Priority:* In order to limit fewer flows, we are required to carefully sort the flows engaged in deadlocks, and to pick up the "more beneficial" flows to be buffered first. From the example above, we can draw a conclusion that the priority of a flow is related to the out-degree and in-degree of the simplified dependency graph, the deadlocks it involves, and the whole latency of its engaged deadlocks. So we formulate a flow's priority in Eqn. 4 to evaluate the importance of breaking deadlocks and fast updating.

$$pri_{f_i} = \frac{degree(out)}{degree(in)} * \max delay(cycle_k), \quad k = 1, \dots, cyc_{f_i} \quad (4)$$

cyc_{f_i} is the number of deadlocks where flow i is involved in the dependency graph. In the example's resource dependency graph, we can see that f_2 's in-degree and out-degree difference is the largest. However, f_1 appears in more deadlocks, and the sum of deadlocks delays is biggest. Using Eq. 4, we sort the importance of these three flows as f_1 , f_2 , and f_3 . Table I verifies the efficiency of our formulation.

2) *MAPUM Scheduling:* The specific updating process of deadlocks is exhibited in Algorithm 2, where time is measured in units of steps. After constructing the dependency graph G using the above strategy, we use Dionysus [4] to consistently and efficiently update the flows that are not involved in any deadlock. We simplify the dependency graph G by deleting the

Algorithm 2 MAPUM

Input: The dependency graph G and the deadlock set L ;
Output: Updating plan of flow and buffer allocation;

```
1: for every new time step do
2:   In order, update every flow in the next link of its initial
   and final state;
3:   if there is still a deadlock in  $L$  then
4:     Select the current highest priority flow  $f_{highest-pri}$ ;
5:     if  $f_{highest-pri}$  can break down any deadlock in  $L$ 
       then
6:       Rate-limit  $f_{highest-pri}$  and add buffer at the first
       blocked link's import switch along its path;
7:       Delete all deadlocks that limiting the flow
        $f_{highest-pri}$  can solve in the deadlock set  $L$ ;
8:     else
9:       The updating arrangement is finished;
10:    Break out of the for-circulation;
11: return ;
```

non-blocked flows and their occupied resources. Next, [16] is utilized to find all the elementary circuits in the current G . Each flow's priority is calculated by Eqn. 4 and arranged in a descending order. In every time step, we rate limit the highest priority flow and wait for $T_{interval}$ to make sure that no chaotic situations arise. We repeat the process until there is no deadlock in the G .

It takes some time to transfer a packet along a link; this is known as the link latency. By using time step to reserve enough interval, we can efficiently avoid accident congestions. Moreover, due to the difference of each link's latency and the impossibility of time synchronization, it is difficult for the controller to decide the exact updating time to send messages to updating switches. It also takes different amount of time for switches to enforce the updating plans, and for the same switch to handle different kinds of modifications, such as rule table insertion or deletion. We reserve enough time in each time step to make sure that the interval is long enough for even the slowest switch to accomplish the controller's updating demand. The interval can be estimated as follows:

$$T_{interval} = \max\{\tau(s_i), i = 1, 2, \dots, n\} \quad (5)$$

$$\tau(s_i) = \max\{t_i(\gamma), \gamma \in \Gamma\} \quad (6)$$

Γ is the general set of all kinds of updating demands from the controller. We yield to the longest time to guarantee the congestion-free property.

C. Schedule Improvemet

We propose a heuristic algorithm to solve the consistent network updating problem. We also find some improvements of our algorithm that help to achieve a better performance. Every link has its transmission latency, so it takes some time for the flow to pass the link. As a result, we can begin to migrate the flow from its source switch when it can obtain every link resource before arriving at its destination. This




Fig. 5: Topology

saves the updating time, but will have a more complicated scheduling strategy. Furthermore, if the updating is slow, and some buffer is released before, it is flexible to reuse the free buffer, which improves the utilization of switch buffer. We also observe that if a deadlock consists of many flows and has a large update delay, it is necessary to rate limit a certain number of flows instead of just one. The divided parts of the deadlock can update in parallel, which saves significant time. The selection of rate-limiting flows should be balanced to achieve the min max update time of each branch.

VI. EVALUATION

Simulated experiments are conducted to evaluate the performances of the proposed algorithms. After presenting the network and flow settings, the results are shown from different perspectives to provide insightful conclusions.

A. Network and Flow Setting

We do simulations in two realistic topologies. The first is Microsoft's inter-data center WAN topology [4], consisting of 8 switches connected as shown in Fig. 5a. Each link is two-way and has a capacity of 10-Gbps. The second is a fat-tree topology [17] for the data center network scenario shown in Fig. 5b. There are 10 core switches, 20 aggregation switches, and 40 edge switches in this network. Each edge switch connects 2 hosts. Each switch has 4 10-Gbps ports, resulting in a full bisection bandwidth network. There are several flows generated randomly with different capacities, and we change flow numbers to simulate traffic variations. We assume the initial and final states of the network and flows are all valid, meaning that the path has no loop and the load on every link is within the capacity.

B. Benchmark Schemes

We compare our MAPUM with three schemes: RS (random flow updating order); DELS (delay-consideration); DEGS (degree-consideration).

C. Performance

We study the updating time and buffer size needed to finish the flow migration using four methods.

WAN: Fig. 6a shows the updating time under different numbers of flows. As there are more flows added, the updating time increases more quickly. When the traffic load is heavier, the network is likely to have more deadlocks. Then it becomes harder to solve deadlocks, and the advantage of our scheme, MAPUM, is more obvious (shown as the purple line in the

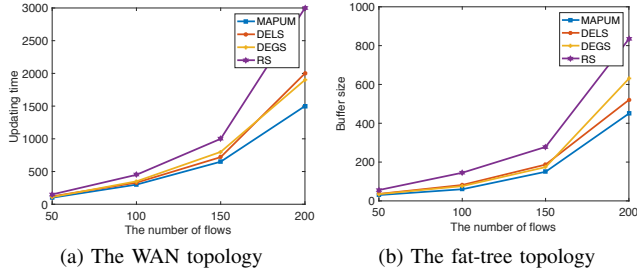


Fig. 6: WAN simulation results




Fig. 7: Data center simulation results

picture). This tendency demonstrates the effectiveness of our algorithm. DELS's and DEGS's performances are roughly the same, but not as good as MAPUM's, because they only consider a part of the elements influencing the flow deadlocks. RS is used as a baseline. Fig. 6b shows the buffer size needed to accomplish the updating process. In order to migrate flows in a lossless way, we store the upcoming flows in the switch buffer to vacate their rooms to other flows of the final state. However, buffer is expensive and rare, so we would like to use it as seldom as possible. PAPUM not only achieves the best performance in updating time, but also uses the least buffer.

Fat-tree: We also do large scale simulations in a fat-tree data center topology. Though there are more switches, deadlocks are less likely to happen under this topology. The flow paths are configured neatly, so the traffic load balances much better than in the WAN. Resource in the data centers is expensive. It is highly recommended to leverage the network in a high-utilization way. Even a small congestion that lasts for only a short time will lead to a great loss and should be avoided to the best of the network administrators' abilities. Fig. 7a and 7b show that MAPUM performs the best of all four methods. It migrates flows with the shortest updating time and the least buffer. The difference shows more clearly when the traffic load is heavier.

VII. CONCLUSION

MAPUM achieves consistent and swift network update, no matter how interwoven among flows in different states. We tactfully select several switches to limit bypass flows' rates with careful allocation of properly sized buffer. We prove that the buffer deployment problem is NP-hard. Furthermore, we propose a heuristic greedy algorithm to efficiently choose

rate-limiting flows to accomplish different system constraints. We also make some improvements of our algorithm. Our simulation results evaluate the correctness, feasibility, and effectiveness of our approach.

VIII. ACKNOWLEDGMENT

This research was supported in part by NSF grants CNS1629746, CNS 1564128, CNS 1449860, CNS 1461932, CNS 1460971, CNS 1439672, CNS 1301774, and ECCS 1231461.

REFERENCES

- [1] N. W. S. Tseng, C. Lim and A. Tang, "Time-aware congestion-free routing reconfiguration," in *The IFIP Networking 2016 Conference (NETWORKING 2016)*, 2016, pp. 202–208.
- [2] L. Fratta, M. Gerla, and L. Kleinrock, "The flow deviation method: An approach to store-and-forward communication network design," *Networks*, vol. 3, no. 2.
- [3] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven wan," ser. SIGCOMM '13.
- [4] X. Jin, H. H. Liu, R. Gandhi, S. Kandula, R. Mahajan, M. Zhang, J. Rexford, and R. Wattenhofer, "Dynamic scheduling of network updates," ser. SIGCOMM '14, pp. 539–550.
- [5] H. H. Liu, X. Wu, M. Zhang, L. Yuan, R. Wattenhofer, and D. Maltz, "zupdate: Updating data center networks with zero loss," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, pp. 411–422, Aug. 2013.
- [6] L. Vanbever, S. Vissicchio, C. Pelsser, P. Francois, and O. Bonaventure, "Seamless network-wide igp migrations," ser. SIGCOMM '11.
- [7] J. McClurg, H. Hojjat, P. Černý, and N. Foster, "Efficient synthesis of network updates," in *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI 2015.
- [8] M. Reitblatt, N. Foster, J. Rexford, and D. Walker, "Consistent updates for software-defined networks: Change you can believe in!" in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, ser. HotNets-X, 2011.
- [9] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker, "Abstractions for network update," ser. SIGCOMM '12, pp. 323–334.
- [10] N. P. Katta, J. Rexford, and D. Walker, "Incremental consistent updates," ser. HotSDN '13.
- [11] R. W. S. Brandt, K. Foerster, "On consistent migration of flows in sdns," in *IEEE International Conference on Computer Communications (INFOCOM 2016)*, 2016.
- [12] J. Zheng, H. Xu, G. Chen, and H. Dai, "Minimizing transient congestion during network update in data centers," in *2015 IEEE 23rd International Conference on Network Protocols (ICNP)*.
- [13] T. Mizrahi, E. Saat, and Y. Moses, "Timed consistent network updates in software-defined networks," *IEEE/ACM Transactions on Networking*, 2016.
- [14] T. Mizrahi, O. Rottenstreich, and Y. Moses, "Timeflip: Scheduling network updates with timestamp-based tcam ranges," in *2015 IEEE Conference on Computer Communications (INFOCOM)*, 2015, pp. 2551–2559.
- [15] J. Zheng, H. Xu, G. Chen, H. Dai, and J. Wu, "Congestion-minimizing network update data centers," in *IEEE Transactions on Services Computing*, 2016.
- [16] D. B. Johnson, "Finding all the elementary circuits of a directed graph," *SIAM Journal on Computing*.
- [17] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," *SIGCOMM Comput. Commun. Rev.*