Survey of Consistent Software-Defined Network Updates

Klaus-Tycho Foerster, Stefan Schmid, and Stefano Vissicchio

Abstract—Computer networks have become a critical infrastructure. In fact, networks should not only meet strict requirements in terms of correctness, availability, and performance, but they should also be very flexible and support fast updates, e.g., due to policy changes, increasing traffic, or failures. This paper presents a structured survey of mechanisms and protocols to update computer networks in a fast and consistent manner. In particular, we identify and discuss the different desirable consistency properties that should be provided throughout a network update, the algorithmic techniques which are needed to meet these consistency properties, and the implications on the speed and costs at which updates can be performed. We also explain the relationship between consistent network update problems and classic algorithmic optimization ones. While our survey is mainly motivated by the advent of Software-Defined Networks (SDNs) and their primary need for correct and efficient update techniques, the fundamental underlying problems are not new, and we provide a historical perspective of the subject as well.

I. INTRODUCTION

Computer networks such as datacenter networks, enterprise networks, carrier networks etc. have become a critical infrastructure of the information society. The importance of computer networks and the resulting strict requirements in terms of availability, performance, and correctness however stand in contrast to today's ossified computer networks: the techniques and methodologies used to build, manage, and debug computer networks are largely the same as those used in 1996 [1]. Indeed, operating traditional computer networks is often a cumbersome and error-prone task, and even tech-savvy companies such as GitHub, Amazon, GoDaddy, etc. frequently report issues with their network, due to misconfigurations, e.g., resulting in forwarding loops [2], [3], [4], [5]. An anecdote reported in [1] illustrating the problem, is the one by a Wall Street investment bank: due to a datacenter outage, the bank was suddenly losing millions of dollars per minute. Quickly the compute and storage emergency teams compiled a wealth of information giving insights into what might have happened. In contrast, the networking team only had very primitive connectivity testing tools such as ping and traceroute, to debug the problem. They could not provide any insights into the actual problems of the switches or the congestion experienced by individual packets, nor could the team create any meaningful experiments to identify, quarantine and resolve the problem [1].

K.-T. Foerster and S. Schmid are with University of Vienna, Vienna, Austria.
E-mail: klaus-tycho.foerster@univie.ac.at and stefan_schmid@univie.ac.at.
S. Vissicchio is with University College London, London, United Kingom.
E-mail: s.vissicchio@s.ucl.ac.uk.

Software-defined networking is an interesting new paradigm which allows to operate and verify networks in a more principled and formal manner, while also introducing flexibilities and programmability, and hence foster innovations. In a nutshell, a Software-Defined Network (SDN) outsources and consolidates the control over the forwarding or routing devices (located in the so-called *data plane*) to a logically centralized controller software (located in the so-called *control plane*). This decoupling allows to evolve and innovate the control plane independently from the hardware constraints of the data plane. Moreover, OpenFlow, the de facto standard SDN protocol today, is based on a simple match-action paradigm: the behavior of an OpenFlow switch is defined by a set of forwarding rules installed by the controller. Each rule consists of a match and an action part: all packets matched by a given rule are subject to the corresponding action. Matches are defined over Layer-2 to Layer-4 header fields (e.g., MAC and IP addresses, TCP ports, etc.), and actions typically describe operations such as forward to a specific port, drop, or update certain header fields. In other words, in an SDN/OpenFlow network, network devices become simpler: their behavior is defined by a set of rules installed by the controller. This enables formal reasoning and verification, as well as flexible network update, from a logically centralized perspective [6], [7]. Moreover, as rules can be defined over multiple OSI layers, the distinction between switches and routers (and even simple middleboxes [8]) becomes blurry.

However, the decoupling of the control plane from the data plane also introduces new challenges. In particular, the switches and controllers as well as their interconnecting network form a complex asynchronous distributed system. For example, a remote controller may learn about and react to network events slower (or not at all) than a hardware device in the data plane: given a delayed and inconsistent view, a controller (and accordingly the network) may behave in an undesirable way. Similarly, new rules or rule updates communicated from the controller(s) to the switch(es) may take effect in a delayed and asynchronous manner: not only because these updates have to be transmitted from the controller to the switches over the network, but also the reaction time of the switches themselves may differ (depending on the specific hardware, data structures, or concurrent load).

Thus, while SDN offers great opportunities to operate a network in a correct and verifiable manner, there remains a fundamental challenge of how to deal with the asynchrony inherent in the communication channel between controller and switches as well as in the switches themselves. Accordingly, the question of how to update network behavior and configurations correctly yet efficiently has been studied intensively over the

last years. However, the notions of correctness and efficiency significantly differ across the literature. Indeed, what kind of correctness is needed and which performance aspects are most critical often depends on the context: in security-critical networks, a very strong notion of correctness may be needed, even if it comes at a high performance cost; in other situations, however, short transient inconsistencies may be acceptable, as long as at least some more basic consistency guarantees are provided (e.g., loop-freedom).

We observe that not only is the number of research results in the area growing very quickly, but also the number of models, the different notions of consistency and optimization objectives, as well as the algorithmic techniques: Thus, it has become difficult to keep an overview of the field even for active researchers. Moreover, many of the underlying update problems are not entirely new or specific to SDN: rather, techniques to consistently update legacy networks have been studied in the literature, although they are based on the (more restrictive) primitives available in traditional protocols (e.g., IGP weights).

We therefore believe that it is time for a comprehensive survey of the subject.

A. The Network Update Problem

Any dependable network does not only need to maintain a range of static invariants, related to correctness, availability, and performance, but also needs to be flexible and support reconfigurations and updates. Reasons for updating a network are manifold, including:

1) *Security policy changes:* For example, in enterprise networks, traffic from a specific subnetwork may have to be rerouted via a firewall upon certain events (suspicious traffic). Also, in wide-area networks, the set of countries via which it is safe to route sensitive traffic may change over time.

2) *Traffic engineering:* In order to improve traffic engineering metrics (e.g., minimizing the maximal link load), a system administrator or operator may decide to reroute (parts of) the traffic along different links. For example, many Internet Service Providers switch between multiple routing patterns during the day, depending on the expected load. These patterns may be precomputed offline, or may be computed as a reaction to an external change (e.g., due to a policy change of a Content Distribution Provider).

3) *Maintenance work:* Also maintenance work may require the update of network routes. For example, in order to replace a faulty router, or to upgrade an existing router, it can be necessary to temporarily reroute traffic.

4) *Link and node failures:* Failures happen quite frequently and unexpectedly in today's computer networks, and typically require a fast reaction. Accordingly, fast network monitoring and update mechanisms are required to react to such failures, e.g., by determining a failover path.

5) *Adding or removal of services:* A wide range of services are typically run over networks, from in-network packet processing functions (e.g., virtualized middleboxes) to applications (like data storage or application servers).

Addition, removal or relocation of any of those services would require the update of network routes to bring the intended traffic to the affected service.

Despite these changes, it is often desirable that the network maintains certain consistency properties *throughout the update*. Those properties may include per-packet path consistency (a packet should be forwarded along the old or the new route, but never a mixture of both), waypoint enforcement (a packet should never bypass a firewall), or at least correct packet delivery (at no point in time packets should be dropped or trapped in a loop).

While the above reasons for network updates are general and relevant independently of the adopted paradigm, the decoupling of control- and data-plane, as well as the flexibility allowed by the SDN architecture are likely to increase the frequency of network updates, e.g., for supporting more fine-grained and frequent optimization of traffic paths [9].

B. Our Contributions

This paper presents a comprehensive survey of the consistent network update problem in Software-Defined Networks (SDNs). The goal of our survey is to (1) provide active researchers in the field with an overview of the state-of-the-art literature, but also to (2) help researchers who only recently became interested in the subject to bootstrap and learn about open research questions.

We identify and compare both the different notions of consistency and the different performance objectives considered in the literature. In particular, we provide an overview of the algorithmic techniques required to solve specific classes of network update problems, and discuss inherent limitations and tradeoffs between the achievable level of consistency and the speed at which networks can be updated. In fact, as we will see, some update techniques are not only less efficient than others, but with them, it can even be impossible to consistently update a network.

We also present a historical perspective, surveying the consistency notions provided in traditional networks and discussing the corresponding techniques accordingly.

Moreover, we put the algorithmic problems into perspective and discuss how these problems relate to classic optimization and graph theory problems, such as multi-commodity flow problems or maximum acyclic subgraph problems.

C. Paper Organization

The remainder of this paper is organized as follows. §II presents a historical perspective and reviews notions of consistency and techniques both in traditional computer networks as well as in Software-Defined Networks. §III then presents a classification and taxonomy of the different variants of the consistent network update problems. §IV, §V, and §VI review models and techniques for connectivity consistency, policy consistency, and capacity consistency related problems, respectively. §VII discusses proposals to further relax consistency guarantees by introducing tighter synchronization. In §VIII, we identify practical challenges. After highlighting future research directions in §IX, we conclude our paper in §X.

II. THE NETWORK UPDATE PROBLEM FROM THE ORIGINS TO SDN

Any computer network must guarantee some consistency properties for the configured forwarding rules and paths. For example, forwarding loops must be avoided, as they can quickly deplete switch buffers and harm the availability and connectivity provided by a network.

It is eminently desirable to preserve consistency properties *during network updates* – i.e., while changing packet-processing rules on network devices. In fact, early studies on consistent network updates date back long before the advent of software-defined networking. In this section, we provide a historical perspective on the many research contributions that can be considered as the main precursors of the state of the art for SDN updates.

We first discuss update problems and techniques in traditional networks (§II-A-II-B). In those networks, forwarding rules are computed by routing protocols that run standardly-defined distributed algorithms, whose output is influenced by both physical topology (e.g., active links) and routing configurations (e.g., logical link costs). Pioneering update works aimed at avoiding transient inconsistencies due to modified topology or configurations, mainly focusing on the Interior Gateway Protocols (IGPs) that are commonly used to control forwarding within a single network. A first set of contributions tried to modify IGP protocol definitions, mainly to provide forwarding consistency guarantees upon link or node failures. Progressively, the research focus has shifted to a more general problem of finding sequences of IGP configuration changes that modify forwarding while guaranteeing forwarding consistency, e.g., for service continuity (§II-A). More recent works have also considered reconfigurations of protocols different or deployed in addition to IGPs, mostly generalizing previous techniques while keep focusing on forwarding consistency (§II-B).

Subsequently (§II-C), we discuss update problems in the context of SDNs. Those networks are based on a clear separation between controller (implementing the control logic) and dataplane elements (applying controller's decision on packets). This separation indisputably provides new flexibility and opens new network design patterns, for example, enabling security requirements to be implemented by careful path computation (done by the centralized controller). This also pushed network update techniques to consider additional consistency properties like policies and performance.

Throughout the section, we rely on the generic example shown in Fig. 1 for illustration. The figure shows the intended forwarding changes to be applied for a generic network update. Observe that possible forwarding loops can occur when we update nodes one by one, because links $(v1, v2)$ and $(v2, v3)$ are traversed in opposite directions before and after the update.

A. IGP Reconfigurations

In traditional (non-SDN) networks, forwarding paths are computed by distributed routing protocols. Among them, link-state IGPs are typically used to compute forwarding paths within a network owned by the same administrative entity. They are based on computing shortest-paths on a logical


Fig. 1. A network update example, where forwarding paths have to be changed from the Surpassed (Fig. 1a) to the Down (Fig. 1b) state. Arrows represent paths on which traffic (e.g., from s to d) is forwarded, while (gray) undirected links between nodes represent unused links.

view of the network, that is, a weighted graph which is shared across routers. Parameters influencing IGP computations, like link weights, can be set by operators by editing router configurations.

As an illustration, Fig. 2 shows a possible IGP implementation for the network states shown in Fig. 1. In particular, Fig. 2 reports the IGP graph (consistent with the physical network topology) with explicit mention of the configured link weights. Based on those weights, for each destination (e.g., d in this example), all routers independently compute the shortest paths, and forward the corresponding packets to the next-hops on those paths. Consequently, the IGP configurations in Figs. 2a and 2b respectively produce the forwarding paths depicted in Figs. 1a and 1b.


Fig. 2. Possible implementation of pre- and post-update forwarding paths for the update in Fig. 1 in a traditional, IGP-based network. Numbers close to network links represent the corresponding IGP weights.

When the IGP graph is modified (e.g., because of a link failure, a link-weight change or a router restart), messages are propagated by the IGP itself from node to node, so that all nodes rebuild a consistent view of the network: This process is called *IGP convergence*. However, IGPs do not provide any guarantee on the timing and ordering in which nodes receive messages about the new IGP graphs. This potentially triggers transient forwarding disruptions due to temporary state inconsistency between a set of routers. For example, assume that we simply remove link $(v3, d)$ from the IGP graph shown in Fig. 2a. This will eventually lead us to the configuration presented in Fig. 2b. Before the final state is reached, the notification that $(v3, d)$ is removed has to be propagated to all routers. If $v3$ receives such notification before $v2$ (e.g., because closer to the removed link), then $v3$ would recompute its next-hop based on the new information, and starts forwarding packets for d to $v2$ (see Fig. 1b). Nevertheless, $v2$ keeps forwarding packets to $v3$ as it considers that $(v3, d)$ is still up. This creates a loop between $v3$ and $v2$: The loop remains until $v2$ is notified about the removed link. A similar loop can occur between $v2$ and $v1$.

Guaranteeing disruption-free IGP operations has been considered by research and industry since almost two decades. We now briefly report on the main proposals in the area, which share the focus on support for planned operations.

Protocol extensions. Early contributions focused on the modification of IGPs, mainly to avoid forwarding inconsistencies. Among them, protocol extensions have been proposed [10], [11], [12] to gracefully restart a routing process, that is, to avoid forwarding disruptions (e.g., blackholes) during the software upgrade of a router. Other works focused on avoiding forwarding loops during configuration changes. For example, François and Bonaventure [13] propose oFIB, an IGP extension that guarantees the absence of forwarding loops upon manually managed topological changes, e.g., to propagate information about a link or a router that has to be shut down for maintenance. The key intuition behind oFIB is to use explicit synchronization between routers to constrain the order in which each router changes its forwarding entries. Namely, each router is forced not to update its forwarding entry for a given destination until all its final next-hops switched to their respective final next-hops for that destination. Consider again Fig. 2, assuming that oFIB is deployed. To prepare the shutdown of link $(v3, d)$, oFIB ensures that $v1$ is the only router changing its forwarding entry to d at first: this is safe because $v1$'s final next-hop is directly the destination d. All the other routers (e.g., $v2$) do not update yet until their final next-hops (e.g., $v1$) use their respective final paths. In fact, after $v1$ starts forwarding traffic through the $(v1, d)$ link, it notifies its neighbors about its updated state. At that point, $v2$ and s can update their respective forwarding entries for d. The whole network is eventually updated by iterating this process.

oFIB inspired a number of variants, aiming at applying explicit notification to more generic updates. Fu *et al*. [14] generalize the approach by defining a loop-free ordering of IGP-entry updates for arbitrary forwarding changes. Shi *et al*. [15] also extend the reconfiguration mechanism to avoid traffic congestion in addition to forwarding incorrectness. A broader overview of loop avoidance and mitigation techniques mostly inspired by oFIB is reported in [16].

Modifying protocol specifications may seem the most straightforward solution to deal with reconfigurations in traditional networks, but it actually has practical limitations. First, this approach cannot accommodate custom reconfiguration objectives. For instance, ordered forwarding changes generally work only on a per-destination basis [13], which can make the reconfiguration process slow if many destinations are involved – while one operational objective could be to exit transient states as quickly as possible. Second, protocol modifications are targeted to specific reconfiguration cases (e.g., single-link failures), since it is intrinsically hard to predict the impact of any possible configuration change on forwarding paths. Finally, protocol extensions are not easy to implement in practice, because of the reluctance of vendors to change their proprietary router software, as well as the additional complexity and potential overhead (e.g., load) induced on routers.

Limited practicality of protocol modifications quickly motivated new approaches, based on coordinating operations readily available in deployed routers, at a per-router level.

Coarse-grained operation scheduling. Planned reconfigurations may encompass several coarse-grained operations. Consider the case where in Fig. 2a, the weight of link $(s, v1)$ has to be set to 70 in addition to removing the link $(v3, d)$. The link reweighting might be desirable to improve load balancing across the network, e.g., adapting to a permanently increased volume of traffic from s to d. Such a reconfiguration effectively consists of two macro operations: removing $(v3, d)$ and changing the weight of $(s, v3)$. Even assuming that each coarse-grained operation is atomic, the order in which distinct operations are performed can have an impact on how much the network traffic is disrupted during the reconfiguration. Assume that link $(s, v3)$ can sustain no more than 50% of the $s - d$ traffic volume. Reweighting $(s, v1)$ before removing $(v3, d)$ forces all the $s - d$ traffic on the path $(s, v3, d)$ which overloads link $(s, v3)$, while removing $(v3, d)$ first would not cause congestion.

Several works propose to use optimization techniques to compute the order of macro-operations so as to guarantee given consistency properties. As a first approach, Keralapura *et al*. [17] formalized the problem of finding the optimal order in which nodes can be added to a network, one by one, so as to minimize an objective function modeling typical costs of connectivity and traffic disruptions in Internet Service Providers. THe following contributions encompass additional operations. In 2009 [18] and 2011 [19], for instance, Raza *et al*. propose a theoretical framework to schedule link weight changes in a way that minimizes a generic disruption function. This approach enables to formulate our reconfiguration example as a formal optimization problem, where constraints enforce that the reweighting of links $(s, v1)$ and $(v3, d)$ (from 10 to 70 and from 10 to ∞, respectively) are both scheduled, and the objective functions aggregates the cost of every step in the schedule. The works also describe two algorithms to solve the formalized problems, one based on dynamic programming and the other on an ant colony optimization heuristic.

The approaches just described basically spread coarse-grained operations over time. This is not sufficient to deal with many update scenarios. Fig. 2 displays one of such scenarios: since the reconfiguration includes a single coarse-grained operation (link removal), previous approaches cannot prevent forwarding loops possibly occurring when that single operation is performed.

Progressive link reweighting. Intermediate link weights can be set to avoid disruptions during a reconfiguration, even if it includes a single link weight change. Consider again the example in Fig. 2, and let the final weight for link $(v3, d)$ conventionally be ∞. In this case, the forwarding loops potentially triggered by the IGP reconfiguration can be provably prevented by using two intermediate weights for link $(v3, d)$, as illustrated in Fig. 3. The first of those intermediate weights (see Fig. 3b) is used to force $v1$ and only $v1$ to change its next-hop, from $v2$ to d: Intuitively, this prevents the loop between $v2$ and $v1$. The second intermediate weight (see Fig. 3c) similarly guarantees that the loop between $v3$ and $v2$ is avoided, i.e., by forcing $v2$ to use its final next-hop before $v3$.

Of course, computing intermediate weights that guarantee

Fig. 3. Intermediate IGP weights that enable a loop-free reconfiguration for the example in Fig. 1.

the absence of disruptions becomes trickier when multiple destinations are involved.

Such a technique can be straightforwardly applied to real routers. For example, an operator can progressively change the weight of $(v3, d)$ to 31 by editing the configuration of $v3$ and d, then check that the all IGP routers have converged on the paths in Fig. 3b, repeat similar operations to reach the state in Fig. 3c, and finally remove the link safely. Even better, François *et al.* [20] have proved that it is always possible to compute a sequence of intermediate link weights that provably avoids all transient loops when a single link has to be reweighted. Obviously, the weight of multiple links can be changed in a loop-free way, by safely reweighting links one by one.

Additional research contributions focused on minimizing the number of intermediate weights that ensure loop-free reconfigurations. Surprisingly, the problem is *not* computationally hard, despite the fact that all destinations have potentially to be taken into account when changing link weights. Polynomial-time algorithms have been proposed to support planned operations at the per-link [21], [20] (e.g., single-link reweighting) and at a per-router [22], [23] (e.g., router shutdown/addition) granularity.

Ships-in-the-Night (SITN) techniques. To improve the update speed in the case of simultaneous link weight changes and to deal with additional reconfiguration scenarios (from changing routing parameters like OSPF areas to replacing an IGP with another), both industrial best practices and research works often rely on a technique commonly called Ships-in-the-Night [24]. This technique builds upon the capability of traditional routers to run multiple routing processes at the same time. Thanks to this capability, both the initial and final configurations can be installed (as different routing processes) on all nodes at the same time. When multiple configurations are installed on the same node, only one of them is preferred and used. Fig. 4 shows the setup for a Ships-in-the-Night reconfiguration for the reconfiguration case in Fig. 2.

In SITN, the reconfiguration process then consists in swapping the preference between the initial and the final configurations on every node, typically one by one. Configuration preference can be swapped at a per-destination granularity. This means that (1) for each destination, every node either forwards packets to its initial next-hops or its final ones; (2) at any time during the reconfiguration, distinct nodes can use different configurations; hence, (3) inconsistencies may arise from the mismatch between the configurations used by distinct nodes.


Fig. 4. Ships-in-the-Night setup: All routers run two routing processes, one with the initial configuration (blue, solid lines) and the other with the final configuration (green, dashed and dotted segments).

Because of those potential inconsistencies, the Ships-in-the-Night approach opens a new algorithmic problem, that is, to decide a safe order in which to swap preferences on a per-router basis. For example, if the configuration preference is swapped on $v3$ before doing the same on $v2$ in Fig. 4, we end up with a loop between $v2$ and $v3$. In contrast, Fig. 5 shows a SITN-based safe reconfiguration that mimics the progressive link weight increment depicted in Fig. 3.

Naive algorithms for swapping configuration preferences cannot guarantee disruption-free reconfigurations. For example, replacing the initial configuration with the final one on all nodes at once provides no guarantee on the order in which new preferences are applied by nodes, hence they potentially trigger packet losses and service disruptions (in addition to massive control-plane message exchanges). Such an approach will also leave the network in an inconsistent, disrupted and hard-to-troubleshoot state if any reconfiguration command is lost or significantly delayed. Industrial best practices (e.g., [24], [25]) only provide rules of thumb which do not apply in the general case, and do not guarantee lossless reconfiguration processes.

Hence, the problem of computing a safe per-router reconfiguration order within SITN called for new research contributions. Prominently, [26], [27] show that no SITN-based update order guarantees the absence of forwarding loops in some cases, and it is NP-complete to even assess whether a loop-free order exists. Those papers also propose two algorithms to compute a loop-free order. The first algorithm is based on applying traditional optimization algorithms to solve a Linear Program (LP) that models the input update problem. Such an LP is derived from enumerating all possible loops, and encoding all the possibilities to avoid every loop as LP constraints (that must be alternatively satisfied). For the example in Fig. 4, this algorithm would enumerate the two potential loops (between $v1$ and $v2$, and between $v2$ and $v3$), and formulate LP constraints

Fig. 5. Ships-in-the-Night reconfiguration that mimics the progressive link reweighting shown in Fig. 3. In each figure, the colors of router names indicate their respective control-plane preferences at the represented reconfiguration step.

forcing the control-plane preference to be swapped at $v1$ before $v2$ (to break the first loop) and at $v2$ before $v3$ (to break the second loop). The second algorithm is a heuristic based on reconfiguring routers according to the final paths (e.g., $v1$ before $v2$ and s in our example, because $v1$ is closer to d), aimed to avoid the scalability problems of the LP-based approach. Finally, [26], [27] describe a comprehensive system to carry out loop-free SITN-based reconfigurations. The system computes the operational sequence to perform an input reconfiguration, and directly interacts with the routers to modify their configurations, and check when every operation in the sequence is completed. The system was envisioned to work semi-automatically, waiting for an explicit confirmation from the operator before performing the next operation in the computed sequence.

B. Generalized Routing Reconfigurations in Traditional Networks

Research contributions have been devoted to reconfigurations in more realistic settings, including other protocols in addition to an IGP.

Enterprise networks, with several routing domains. As a first example, the Ships-in-the-Night framework has been used to carry out IGP reconfigurations in enterprise networks. Those networks typically use *route redistribution* [28], a mechanism enabling the propagation of information from one routing domain (e.g., running an IGP) to another (e.g., running another IGP). Route redistribution may be responsible for both routing (inability to converge to a stable state) and forwarding (e.g., loop) anomalies [28]. SITN-based update procedures have been proposed in [29] to avoid transient anomalies while (i) reconfiguring a specific routing domain, and/or (ii) arbitrarily changing the size and shape of routing domains.

Internet Service Providers (ISPs), with BGP and MPLS. In ISP networks, the Border Gateway Protocol (BGP) and often the Multi-Protocol Label Switching (MPLS) protocol are pervasively used to manage transit traffic, for which both the source and the destination is external to the network. Vanbever *et al.* [30] showed that even techniques guaranteeing safe IGP reconfigurations can cause transient forwarding loops in those settings, because of the interaction between IGP and BGP. They also proved conditions to avoid those BGP-induced loops during IGP reconfigurations, by leveraging the presence

of MPLS or carefully configuring BGP (according to some guidelines).

In parallel, a distinct set of techniques aimed at supporting BGP reconfigurations. François *et al.* [31] propose a solution to avoid churn and loss of connectivity due to planned BGP session shutdown: This solution is based on modifying BGP to distribute alternate routes and move traffic on them *before* the target BGP session is actually removed. Wang *et al.* [32] present an approach, based on extending virtual machine migration techniques, to quickly transfer virtual routers from one physical device to another. Keller *et al.* [33] address the more general problem of fastly migrating parts of the BGP configuration (e.g., transferring a BGP session from one router to another), with a technique that takes care of moving the BGP state to the new route and reduce the impact of the migration on both BGP peers and other routers. Vissicchio *et al.* [34] describe a framework that enables radical re-organizations of BGP sessions (e.g., changing several of them, or transforming a full-mesh into a route reflector topology) while guaranteeing the absence of forwarding and routing anomalies: The framework is based on implementing Ships-in-the-Night in BGP (with a minimal extension to existing routers), and tagging packets so that routers can uniformly apply a single BGP configuration to every packet.

Finally, Internet-level problems, like maintaining global connectivity upon failures, have also been explored (see, e.g., [35]).

Protocol-independent reconfiguration frameworks. By design, all the above approaches are dependent on the considered (set of) protocols and even on their implementation.

Protocol-independent reconfiguration techniques have been studied as well in the literature. Mainly, [36] generalizes SITN by proposing a new design for the internal router architecture. This re-design would allow routers not only to run multiple configurations simultaneously, and to select the configuration to apply for every packet on the basis of a specific bit in the packet header. The work also describes a commit protocol to support the switch between configurations without creating forwarding loops.

General mechanisms for consensus routing have also been explored in [37].

C. Software-Defined Networks

Recently, software-defined networking has grown in popularity, thanks to its promises to spur abstractions, mitigate

compelling management problems and avoid network ossification [38]. SDN is currently used or discussed in a wide range of contexts [39], e.g., to improve network virtualization in datacenters, generalize traffic engineering in the wide-area network, or enable slicing in emerging 5G applications, to just name a few.

In pure SDN networks, rather than having devices (switches and routers) run their own distributed control logic, the controller computes (according to operators' input) and installs (on the controlled devices) the rules to be applied to packets traversing the network: No message exchange or distributed computation are needed anymore on network devices. This is very different from the existing decentralized control planes typically used in traditional networks (as well as in many Ad Hoc and P2P networks) and allows, e.g., to overcome the notoriously slow reaction to changes (e.g., link failures) and rerouting of flows in those networks: one of the key reasons behind Google's move to SDN [40].

Fig. 6 depicts an example of an SDN network, configured to implement the initial state of our update example (see Fig. 1). Beyond the main architectural components, the figure also illustrates a classic interaction between them. Indeed, the dashed lines indicate that the SDN controller instructs the programmable network devices, typically switches [38]), on how to process (e.g., forward) the traversing packets. An example command sent by the controller to switch s is reported next to the dashed line connecting the two: This command instructs s to use $v1$ as next-hop for any packet destined to d.


Fig. 6. Implementation of the surpassed state in Fig. 1 in an SDN network.

The SDN architecture is expected to make network updates more frequent and more critical than in traditional networks. On the one hand, controllers are often intended to support several different requirements, including performance (like optimal choice of per-flow paths), security (like firewall and proxy traversal) and packet-processing (e.g., through the optimized deployment of virtualized network functions) ones. On the other hand, devices cannot provide any reaction (e.g., to topological changes) like in traditional networks. In turn, this comes at the risk of triggering inconsistencies, e.g., creating traffic blackholes during an update, that are provably impossible to trigger by reconfiguring current routing protocols [41]. As a consequence, the controller has to carry out a network update for every event (from failures to traffic surges and requirement modification) that can impact the forwarding rules installed on the switches; additionally, it should perform such updates while typically supporting more critical consistency guarantees (e.g., security-related ones) and performance objectives (e.g., for prompt reaction to failures) than in traditional networks.

An extended corpus of SDN update techniques have already been proposed in the literature, following up on the large interest raised by SDN in the last few years. This research effort nicely complements approaches to specify [42], compile [43], [44], and check the implementation of [45], [6] network requirements that operators may want to implement in their (SDN) networks.

The first cornerstone of SDN updates is represented by the work by Reitblatt *et al.* in 2011 [46] and 2012 [47]. This work provides a first analysis of the additional (e.g., security) requirements to be considered for SDN updates, extending the scope of consistency properties from forwarding to policy ones. In particular, it focuses on per-packet consistency property, imposing that packets have to be forwarded either on their initial or on their final paths (never a combination of the two), throughout an update.

The technical proposal is centered around the 2-phase commit technique, which relies on tagging packets at the ingress so that either all initial rules or all final ones can be consistently applied network-wide. Initially, all packets are tagged with a given "old label" (e.g., no tag) and rules matching the old label are pre-installed on the switches. In a first step, the controller then instructs the internal switches to apply the final forwarding rules to packets carrying a "new label" – even if no packet carries such label at this step. After the internal switches have confirmed the successful installation of these new rules, the controller then changes the tagging policy at the ingress switches, requiring them to tag packets with the "new label". As a result, packets are immediately forwarded along the new paths. Finally, the internal switches are updated (to remove the old rules), and an optional cleaning step can be applied to remove all tags from packets. Fig. 7 shows the operational sequence produced by the 2-phase commit technique for the update case in Fig. 3.

Several works have been inspired by the 2-Phase technique presented in [47]. One first line of research focuses on providing additional guarantees, e.g., congestion-freedom (from [48] to [49], [50], [51], [52], [53], [54], [55]). In a second line of research, several algorithms [27], [56], [57], [58], [59] to compute a set of ordered rule replacements have been proposed to deal with specific SDN update cases (e.g., where only forwarding consistency is needed), avoiding the need for additional rules and hence enabling more resource-efficient approaches (e.g., TCAM memory slots are expensive and precious).

In the following sections, we detail most of those contributions and the insights on different update problems that globally emerge from them.

For a tutorial description of a few major works in the area, mainly [52], [60], we refer to the recent article of Li *et al.* [61].

III. TAXONOMY

We now present a general formulation of network update problem (§III-A), which abstracts from assumptions and settings considered in the literature. This formulation enables us to classify research contributions on the basis of the proposed techniques (e.g., simultaneous usage of multiple

Fig. 7. Application of the 2-phase commit technique for carrying out our update example (see Fig. 3). A final (optional) step consists in cleaning the configuration by removing packet tags, i.e., reverting tagging at $v3$ and s as enforced by Step 2.

configurations on nodes or not) and algorithms, independently of their application to traditional and SDN networks (§III-B).

A. Generalized Network Update Problems

In order to compare and contrast research contributions, we first provide a generalized statement for network update problems. We use again Fig. 1 for illustration.

Basic Problem. Generally speaking, a network update problem consists in computing a sequence of operations that changes the packet-processing rules installed on network devices. Consider any communication network: It is composed by a given set of inter-connected devices, that are able to process data packets (e.g., forwarding them to a next-hop) according to rules installed on them. We refer to the set of rules installed on all devices at a given time as network state at that time. Given an initial and final state, a network update consists in passing from the initial state to the final one by applying operations (i.e., adding, removing or changing rules) on different devices. In Fig. 1, the initial state forces packets from source s to destination d along the path $(s, v1, v2, v3, d)$; the final state forwards the same packets over $(s, v1, d)$, and packets from $v3$ to d on $(v3, v2, v1, d)$. The network update problem consists in replacing the initial rules with the final ones, so that the paths for d are updated from $(s, v1, v2, v3, d)$ to $(s, v1, d)$ and $(v3, v2, v1, d)$.

Operations. To perform a network update, a sequence of operations has to be computed. By operation, we mean a (direct or indirect) modification of packet-processing rules installed on one or more devices. As an example, an intuitive and largely-supported operation on network devices is rule replacement, which consists in instructing a device (e.g., $v3$) to replace an initial rule (e.g., forward the $s - d$ packet flow to $v2$) with the corresponding final one (e.g., forward the $s - d$ flow to d). Operations can be coarse-grained and indirect, as IGP link reweighting or configuration swapping in legacy networks that imply multiple rule replacements at distinct devices (see §II).

Consistency. The difficulty in solving network update problems is that some form of consistency must be guaranteedly preserved *during* the update, for practical purposes (e.g., avoiding service disruptions and packet losses). Preserving consistency properties, in turn, depends on the order in which operations appear in the computed sequence and are executed by network devices. For example, if $v3$ replaces its initial rule with its final one before $v2$ in Fig. 1, then the operational sequence triggers a forwarding loop between $v2$ and $v3$ that interrupts the connectivity from s to d. In §III-B, we provide an overview of consistency properties considered in the literature.

The practical need for guaranteeing consistency has two main consequences (as shown in §II). First, it forces network updates to be performed incrementally, i.e., appropriately scheduling operations over time so that the installed intermediate states are provably disruption-free. Second, it requires a careful computation of operational sequences, implementing specific reasoning in the problem-solving algorithms (e.g., to avoid replacing $v3$'s rule before $v2$'s one in the previous example).

Performance. Another algorithmic challenge consists in optimizing network-update performance. As an example, minimizing the time to complete an update is commonly considered among those optimization objectives. Indeed, carrying out an update incrementally requires to install intermediate configurations, and in many cases it is practically desirable to minimize the time spent in such intermediate states. We provide a broader overview of performance goals considered by previous works in §III-B.

Final Operational Sequences. Generally, the solution for an update problem can be represented as a sequence of *steps* or *rounds*, that both (i) guarantees consistency properties and (ii) optimizes update performance. Each step is a set of operations that can be started at the same time. Note that this does not mean that operations in the same step are assumed to be executed simultaneously on the respective devices; Rather, all operations in the same step can be started in parallel because the target consistency properties are guaranteed independently of the relative order in which those operations are executed. Examples of operational sequences, computed by different techniques, are reported in §II (see Figs. 3 and 5).

B. Update Techniques

In this section, we provide an overview of the problem space and classify existing models and techniques. Previous contributions have indeed considered several variants of the generalized network update problem as we formulated in §III-A. Those variants basically differ in terms of the update problem on which they focus. Update problems in turn define both (1) the network setting, including admitted rule granularity and operations, (2) consistency properties, and (3) performance goals. An overview of the main update problems considered previously is depicted in Fig. 8, where we skipped the
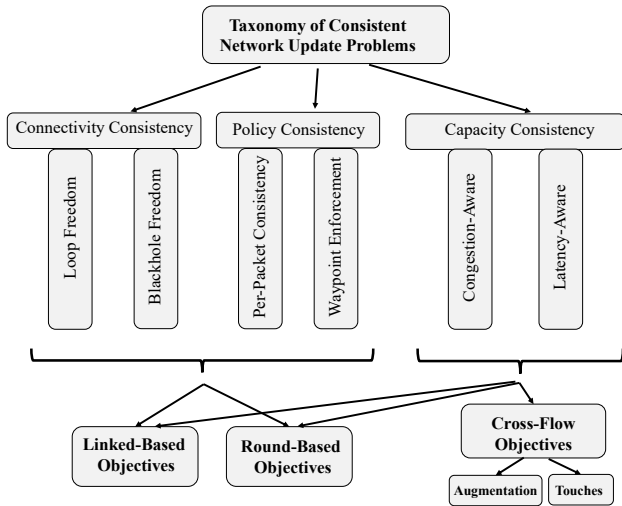
Fig. 8. Types of consistent network update problems, defined independently of the network setting (rule granularity and supported operations).

orthogonal network setting dimensions, rule granularity and supported operations, for clarity.

Rule Granularity. Network update techniques assume that the underlying devices support rules that implement one of the two alternative routing models: *destination-based* and *per-flow* routing.

1) **Destination-based Routing:** In destination-based routing, routers forward packets based on the destination only. An example for destination-based routing is IP routing, where routers forward packets based on the longest common IP destination prefix. In particular, destination-based routing describes confluent paths: once two flows from different sources destined toward the same destination intersect at a certain node, the remainder (suffix) of their paths will be the same. In destination-based routing, routers store at most one forwarding rule per specific destination.

2) **Per-flow Routing:** In contrast, according to *per-flow* routing, routes are not necessarily confluent: the forwarding rules at the routers are defined per-flow, i.e., they may depend not only on the destination but for example also on the source. In traditional networks, flows and per-flow routing could for example be implemented using MPLS: packets belonging to the same equivalence class respectively packets with the same MPLS tag are forwarded along the same path.

Operations. Techniques to carry out network updates can be classified in broad categories, depending on the operations that they consider.

1) **Rule replacements:** A first class of update techniques is based on computing an order in which initial rules are replaced by the corresponding final ones on the devices. Depending on the target setting (e.g., legacy networks or SDNs), such replacement can be admitted at different granularity, that is, on a per-rule and per-device basis (as

in OpenFlow networks) or on a per-group of rules and devices (as for link reweighting in IGP).

2) **Rule additions:** A second class of network update algorithms is based on adding rules to guarantee consistency during the update. The following two main variants of this approach have been explored so far.

   a) **2-Phase commit:** In this case, both the initial and the final rules are installed on all devices in the central steps of the updates. Packets are tagged at the border of the network to enforce that the internal devices either (i) all use the initial rules, or (ii) all use the final rules. See Fig. 7 for an example.

   b) **Helper rules:** Some techniques introduce additional rules, which do not belong neither to the old state nor to the new one, in some intermediate update step. These rules allow to divert the traffic temporarily to other parts of the network, and are called *helper rules*.

3) **Mixed:** Recently, some update techniques combine rule replacements and additions, in order to reduce the update overhead (especially in terms of device-memory consumption) while keeping the flexibility provided by adding rules.

Consistency properties. Update techniques typically target to preserve one (or more) of the following consistency properties. This dimension is also reflected in the structure of this survey.

1) **Connectivity consistency:** The most basic form of consistency regards the capability of the network to keep delivering packets to their respective destinations, throughout the update process. This boils down to guaranteeing two correctness properties: absence of blackholes (i.e., paths including routers that cannot forward the packets further) and absence of forwarding loops (i.e., packets bouncing back and forth on a limited set of routers, without reaching their destinations).

2) **Policy consistency:** Paths used to forward packets may be selected according to specific forwarding policies, for example, security ones imposing that given traffic flows must traverse specific waypoints (firewalls, proxies, etc.). In many cases, those policies have to be preserved during the update. Generally speaking, policy consistency properties impose constraints on which paths can be installed during the update. For example, an already-mentioned policy consistency property (see §II) is *per-packet consistency*, requiring that packets are always forwarded along either the pre-update or the post-update paths, but never a combination of the two.

3) **Capacity consistency:** A third class of consistency properties takes into account the actual availability and limits of network resources. For instance, many techniques account for traffic volumes and corresponding constraints raised by the limited capacity of network links: Those techniques aim at respecting link-capacity constraints in each update step, e.g., to avoid *transient congestion* during updates.

Performance goals. We can distinguish between three broad classes of performance goals.

1) **Link-based:** A first class of consistent network update protocols aims to make new links available as soon as possible, i.e., to maximize the number of switch rules which can be updated simultaneously without violating consistency.

2) **Round-based:** A second class of consistent network update protocols aims to minimize the total makespan, by computing a schedule of rounds or steps, each consisting of switch rules that are safe to update simultaneously.

3) **Cross-Flow Objectives:** A third class of consistent network update protocols targets objectives arising in the presence of multiple flows.

   a) **Augmentation:** Minimize the extent to which link capacities are oversubscribed during the update (or make the update entirely congestion-free).

   b) **Touches:** Minimize the number of interactions with the switch, i.e., the sent messages.

Link-based and round-based objectives are usually considered for node-ordering algorithms and for weak-consistency models. Congestion-based objectives are naturally considered for capacitated consistency models.

IV. CONNECTIVITY CONSISTENCY

In this section, we focus on update problems where the main consistency property to be guaranteed concerns the delivery of packets to their respective destinations. Packet delivery can be disrupted during an update by forwarding loops or blackholes transiently present in intermediate states. We separately discuss previous results on how to guarantee loop-free and blackhole-free network updates. We start from the problem of avoiding forwarding loops during updates, because they are historically the first update problems considered – by works on traditional networks (see §II). This is also motivated by the fact that blackholes cannot be created by reconfiguring current routing protocols, as proved in [41]. We then shift our focus on avoiding blackholes during arbitrary (e.g., SDN) updates.

A. Loop-Freedom

Loop-freedom is a most basic consistency property and has hence been explored intensively already in the network update literature. So far in this work, we presented the notion of loop-freedom in the following framework: 1) routing based on the destination, and 2) avoiding all transient loops. Current research extends the loop-freedom model in both dimensions, introduced next, beginning with the routing model.

1) Definitions: We distinguish between flow-based and destination-based routing: in the former, we can focus on a single (and arbitrary) path from s to d: forwarding rules stored in the switches depend on both s and d, and flows can be considered independently. In the latter, switches store a single forwarding rule for a given destination: once the paths of two different sources destined to the same destination intersect, they will be forwarded along the same nodes in the rest of their route: the routes are confluent.


(a) Initial

(b) After updating s

Fig. 9. Example to illustrate the differences between strong and relaxed loop-freedom. The current rules are drawn in solid blue, the new rules are drawn in dash-dotted green. In the left initial state, the task is to update all nodes to use the new forwarding rules. The right figure shows the network state after updating the node s. Note that the current rules for $v1$, $v2$, $v3$ are no longer on a path from s to t, and are hence drawn dotted. Thus, under RLF, $v1$, $v2$, $v3$ can be updated next, as possible loops are not connected to the source. Using SLF, $v3$ must be updated *after* $v2$, which in turn must be updated *after* $v1$. Hence, by updating $v4$ last, the RLF schedule length is just three, even when the construction is extended from 4 nodes to up to vx nodes. On the other hand, SLF requires at least x rounds. Hence, RLF can yield a speedup linear in the number of nodes.

Moreover, one can distinguish between two different definitions for loop-free network updates: *Strong Loop-Freedom (SLF)* and *Relaxed Loop-Freedom (RLF)* [59]. SLF requires that at any point in time, the forwarding rules stored at the switches should be loop-free. RLF only requires that forwarding rules stored by switches *along the path from a source s to a destination d* are loop-free: only a small number of "old packets" may temporarily be forwarded along loops. RLF can significantly speed up the consistent migration process, as illustrated in Fig. 9.

2) Algorithms and Complexity: Two performance objectives are investigated in the literature, node-based und round-based, to be discussed in turn. Node-based objectives were studied first in the literature: the goal is to update as many nodes/links at a time as possible. Round-based objectives can be seen as more intuitive, aiming at minimizing the number of (controller-switch interaction) rounds. Node-based approaches are also used in round-based contexts, with the following intuition: by updating as many nodes as possible, the total makespan is hopefully minimized – coining the notion of *greedy approaches*. However, it has been shown that node- and round-based objectives can conflict and lead to vastly different schedules.

Node-based objective ("greedy approach"). Mahajan and Wattenhofer [60] initiated the study of destination-based (strong) loop-free network updates. In particular, the authors show that by scheduling updates based on combinatorial dependencies, consistent update schedules can be derived which do not require any packet tagging, and which allow some updated links to become available earlier. The authors also present a first algorithm that quickly updates routes in a transiently loop-free manner: based on the current state, the controller greedily attempts to update as many nodes as possible. E.g., in Fig. 9a, their algorithm picks the nodes s and $v1$, as $v2$-4 have unresolved dependencies. The study of this model has been refined in [57], [51], where the authors also establish hardness results, see Table I. A related variant using so-called proof labeling schemes [62] was proposed in [63].

TABLE I
OVERVIEW OF RESULTS FOR LOOP-FREEDOM.

Model	NP-hard	Polynomial time	Remarks
# Rounds, strong LF	Is there a 3-round loop-free update schedule? [59] *For 2-destination rules and sublinear x: Is there a x-round loop-free update schedule? [51]*	Is there a 2-round loop-free update schedule? [59]	In the worst case, $\Omega(n)$ rounds may be required. [59], *[57]*. $O(n)$-round schedules always exist *[60]*. Both applies to flow-based & *destination-based* rules.
# Rounds, relaxed LF	No results known.	$O(\log n)$-round update schedules always exist. [59]	It is not known whether $o(\log n)$-round schedules exist (in the worst case). No approximation algorithms are known.
# Links, strong LF	Is it possible to update x nodes in a loop-free manner? [64], *[51]*	Polynomial-time optimal algorithms are known to exist in the following cases: A maximum SLF update set can be computed in polynomial-time in trees with two leaves. [64]	The optimal SLF schedule is 2/3-approximable in polynomial time in scenarios with exactly three leaves. For scenarios with four leaves, there exists a polynomial-time 7/12-approximation algorithm. [64] Approximation algorithms from maximum acyclic subgraph [64] and minimum feedback arc set *[57]* apply.
# Links, relaxed LF	Is it possible to update x nodes in a loop-free manner? [64]	Polynomial-time optimal algorithms are known to exist in the following cases: A maximum RLF update set can be computed in polynomial-time in trees with two leaves. [64]	No approximation results known. [64]

Note: Results/references in *italics* are in the destination-based model.

Ludwig et al. in [65], extended in [59], and [66] initiated the study of arbitrary route updates: routes which are not necessarily destination-based. The authors show that the update problem in this case boils down to an optimization problem on a very simple directed graph: initially, before the first update round, the graph simply consists of two connected paths, the old and the new route. In particular, every network node which is not part of both routes can be updated trivially, and hence, there are only three types of nodes in this graph: the source s has out-degree 2 (and in-degree 0), the destination d has in-degree 2 (and out-degree 0), and every other node has in-degree and out-degree 2, as shown in Fig. 9. The authors also observe that loop-freedom can come in two flavors, strong and relaxed loop-freedom [59].

Despite the simple underlying graph, however, Amiri *et al.* [64] show that the node-based optimization problem is NP-hard, both in the strong and the relaxed loop-free model (SLF and RLF). In the example of Fig. 9a the maximization problem is easy though, clearly no more than two nodes (s and $v1$) can be updated initially. As selecting a maximum number of nodes to be updated in a given round (i.e., the node-based optimization objective) may also be seen as a heuristic for optimizing the number of update rounds (i.e., the round-based optimization objective), the authors refer to the node-based approach as the "greedy approach". Amiri *et al.* [64] also present polynomial-time optimal algorithms for specific scenarios, and both[1] [64], [51] provide further insights into approximability properties, see Table I.

Round-based objective. Ludwig *et al.* [59], [65] initiate the study of consistent network update schedules which minimize the number of interaction rounds with the controller: *How many communication rounds k are needed to update a network in a (transiently) loop-free manner?* The authors show that

answering this question is difficult in the strong loop-free case. In particular, they show that while deciding whether a k-round schedule exists is trivial for $k = 2$, it is already NP-complete for $k = 3$. Moreover, the authors show that there exist problem instances which require $\Omega(n)$ rounds, where n is the network size, see Fig. 9. Furthermore, the authors show that the greedy approach, aiming to "greedily" update a *maximum* number of nodes in each round, may result in $\Omega(n)$-round schedules in instances which actually can be solved in $O(1)$ rounds; even worse, a *single* greedy round may inherently delay the schedule by a factor of $\Omega(n)$ more rounds.

However, fast schedules exist for *relaxed loop-freedom*: the authors present a deterministic update scheduling algorithm which completes in $O(\log n)$-round in the worst case.

Other objectives. Dudycz *et al.* [56] initiated the study of how to update multiple policies simultaneously, in a loop-free manner. In their approach, the authors aim to minimize the number of so-called *touches*: the total number of update messages the switches receive from the controller. The number of touches can be reduced if controllers *bundle* the updates of multiple flows to a given switch into a single message. However, consistency requirements impose limits on the extent to which updates can be bundled: e.g., in order to preserve loop-freedom, the update of a flow f_1 needs to take place at node $v1$ *before* $v2$, while the update of a flow f_2 needs to take place at node $v1$ *after* $v2$. The authors establish connections to the *Shortest Common Supersequence (SCS)* and *Supersequence Run* problems [69], and show NP-hardness already for two policies, each of which can be updated in two rounds, by a reduction from *Max-2SAT* [70].

Notwithstanding, Dudycz *et al.* [56] also present optimal polynomial-time algorithms to combine consistent update schedules computed for individual policies (e.g., using any existing algorithm, e.g., [59], [60]), into a global schedule guaranteeing a minimal number of touches. This optimal merging algorithm is not limited to loop-free updates, but applies to any consistency property: if the consistency property holds for individual policies, then it also holds in the joint schedule minimizing the number of touches.

3) Related Optimization Problems: The link-based optimization problem, the problem of maximizing the number of links

---

[1]We note that there exists a subtle difference between the approximation results by Foerster et al. and Ludwig et al.: the former authors usually aim to minimize the number of links which *cannot* be updated [51] (a feedback arc set problem), while Ludwig et al. [64] consider the dual problem variant and aim to maximize the links which *can* be updated in the given round (the maximum acyclic subgraph problem). The approximation guarantees of the two problems differ: for the former model, the best known approximation bound is $O(\log n \log \log n)$ [67] while for the latter, constant approximation results exist [68].

(or equivalently nodes) which can be updated simultaneously, is an instance of the maximum acyclic subgraph problem; equivalently, the dual problem of minimizing the number of links which cannot be updated is a minimum feedback arc set problem. For showing NP-hardness, reductions from the *minimum hitting set problem* [64] and the *feedback arc set problem* [51] are used.

The problem of reconfiguring routes in a network can be seen as a special case of *combinatorial reconfiguration theory*: an abstract reconfiguration framework to transform a feasible solution of a problem (e.g., shortest path routing) into another solution of the same problem, e.g., while ensuring shortest path routing during the update [71]).

4) Concluding Remarks and Open Problems: Loop-free network updates still pose several open problems. Before going into more details regarding open problems below, we summarize current hardness and algorithmic results in Table I, denoting results of the destination-based model in *italics*, whereas the remaining entries refer to route updates which are not necessarily destination-based.

Regarding the node-based objective, Amiri *et al.* [64] conjecture that update problems on bounded directed path-width graphs may still be solvable efficiently: none of the negative results for bounded degree graphs on graphs of bounded directed treewidth seem to be extendable to digraphs of bounded directed pathwidth with bounded degree. More generally, the question of on which graph families network update problems can be solved optimally in polynomial time in the node-based objective remains open. Regarding the round-based objective, it remains an open question whether strong loop-free updates are NP-hard for any $k \geq 3$ (but smaller than n): so far only $k = 3$ has been proved to be NP-hard. More interestingly, it remains an open question whether the relaxed loop-free update problem is NP-hard, e.g., are 3-round update schedules NP-hard to compute also in the relaxed loop-free scenario? Moreover, it is not known whether $\Omega(\log n)$ update rounds are really needed in the worst-case in the relaxed model, or whether the problem can always be solved in $O(1)$ rounds. Some brute-force computational results presented in [59] indicate that if it is constant, the constant must be large.

B. Blackhole-Freedom

Another consistency property is blackhole freedom, i.e., a switch should always have a matching rule for any incoming packet, even when rules are updated (e.g., removed and replaced). This property is easy to guarantee by implementing some default matching rule which is never updated, which however could in turn induce forwarding loops. A straightforward mechanism, if there is currently no blackhole for any destination, is to install new rules with a higher priority, and then delete the old rules [57], [60]. Nonetheless, in the presence of memory limits and guaranteeing loop-freedom, finding the fastest blackhole-free update schedule is NP-hard [57].

V. POLICY CONSISTENCY

Modern requirements go often beyond connectivity. For example, operators may want to ensure that packets traverse a

(a) Extended Surpassed state   (b) Novel state

Fig. 10. A WPE-consistent update example where forwarding paths have to be changed from the Surpassed (Fig. 10a) to the New (Fig. 10b) state, while preserving traversal of the waypoint $v2$ (highlighted in the figure) at any time during the update.

given middlebox (e.g., a firewall) for security reasons, or a chain of middleboxes (e.g., encoder and decoder) for performance reasons; and/or they might like to enforce that paths comply with Service Level Agreements (e.g., in terms of delay). In this section, we discuss studied problems and proposed techniques aiming at preserving such requirements during network updates.

A. Definitions

Requirements on forwarding paths additional to connectivity can be modeled by routing *policies*, that is, links, nodes or sub-paths that have to be traversed by certain traffic flows at any moment in time.

Over the years, several contributions have targeted updates focusing on preserving specific policies. Historically, the first policy considered during SDN updates is *per-packet consistency (PPC)*, which ensures that every packet travels either on its initial or on its final paths, never on intermediate ones. PPC seems a natural choice to comply with high-level network requirements. Assume indeed that both the initial and the final paths accommodate requirements like security, performance, and SLA compliance. The most straightforward way to guarantee that those requirements are not violated is to constrain all paths installed during the update to always be either initial paths or final ones.

Nonetheless, guaranteeing PPC may be an unnecessarily strong requirement in practice. Not always it is strictly needed that transient paths must coincide with either the initial or the final ones. For example, in some cases (e.g., for enterprise networks), security may be the only major concern, and it may translate into simply enforcing that some flows traverse a firewall. Fig. 10 shows an example of this case, where the flow from $s1$ to d has to traverse $v2$. We refer to the property of forced traversal of a given node (waypoint) as *waypoint enforcement (WPE)*.

Policies more complex than WPE (but less constraining than PPC) may also be needed in general. For example, it may be desirable in large Internet Service Providers that specific traffic flows follow certain sub-paths (e.g., with low delay for video streaming and online gaming applications) or are explitictly denied to pass through other sub-paths (e.g., because of political or economical constraints). Such *arbitrary policies* are also considered in recent SDN update works.

TABLE II
OVERVIEW OF RESULTS FOR POLICY-PRESERVING UPDATES.

Ref.	Approach	Guarantees	Computation	Remarks
[46], [47]	2-phase commit	PPC	Constant	Always applicable (if switches have free memory slots); requires packet tagging and additional rules on internal switches
[72]	incremental 2-phase commit	PPC	Exponential	Always applicable (if switches have at least one free memory slot); spreads switch-memory overhead over time
[73]	packet storing protocol	PPC	Constant	Dedicated protocol, based on storing packets at the controller
[74]	per-switch update protocol	PPC	Polynomial	Dedicated protocol, based on the translation of updates into logic circuits
[75]	anti-tampering update protocol	PPC	Polynomial	Dedicated protocol, based on updating switches according to paths; Robust to tampering and dropping attacks
WayUp [66]	ordered rule replacements	WPE	Polynomial	Finishes in 4 rounds; Does not guarantee connectivity (e.g., for loops)
MIP of [66]	ordered rule replacements	WPE	Exponential	Optimizes update time; Guarantees absence of loops
MIP of [58]	ordered rule replacements	WPE chains	Exponential	Optimizes update time; Guarantees absence of loops
[76]	ordered rule replacements	arbitrary	Exponential	Update synthesis based on linear temporal logic and model checking
[77]	ordered rule replacements	PPC	Polynomial	Algorithm based on necessary conditions for PPC-preserving switch updates; Minimizes update rounds
GPIA [78]	ordered rule replacements	PPC	Polynomial	Greedy algorithm; Minimizes update rounds; Applicable to hybrid SDNs (becomes exponential)
GPIA+FED [78]	mixed	PPC	Polynomial	Applies restricted 2-phase commit after rule replacement sequence; Aims at reducing # of update rounds; Applicable to hybrid SDNs (becomes exponential)
[79]	mixed	arbitrary	Exponential	Optimizes the interleaving of rule replacements and additions; Aims at reducing # of update rounds

Note: Contributions are grouped by approach and year.

B. Algorithms and Complexity

Table II overviews solving algorithm and complexity of policy-preserving update problems, further discussed in the following.

2-Phase commit techniques. As described in §III-B, 2-phase commit techniques carry out updates by setting an initial or final tag on packets at ingress devices (e.g., on $s1$ and $s2$ in Fig. 10), and maintaining two forwarding rules at internal nodes (e.g., $v1$, $v2$, $v3$, $v4$, and $v5$ in Fig. 10) so that each packet is forwarded over either its initial or final paths, according to the carried tag. This approach guarantees per-packet consistency by design.

A framework to implement 2-phase commit in traditional networks has been proposed by Alimi *et al.* [36] (see also §II). It requires invasive modification of router internals, to manage tags and run arbitrary routing processes in separate process spaces. Such modifications are not needed in SDNs, where data-plane devices exhibit finer-grained programmability. Beyond presenting an implementation of 2-phase commit in OpenFlow, pioneering SDN update works [46], [47] argue for the criticality of ensuring PPC in the SDN case, and of providing programmatic support for consistent updates within SDN controllers.

A downside of 2-phase commit techniques is that they require internal switches (e.g., $v1$, $v2$, $v3$, and $v5$ in Fig. 10) to maintain an additional rule every time their initial and final rules differ. This overhead requires all switches to have free memory slots (whose number depends on the update case), generally wastes memory resources, and might hamper other applications to properly work during updates (e.g., fast rerouting or security ones that might need to also install new flow rules in reaction to sudden traffic changes). To mitigate those problems, an incremental version of the original technique has been studied in [72]. This work proposes to divide the input update into sub-updates that can be carried out one after the other. Consider for example a variant of Fig. 10 where d is replaced by a set $\mathcal{D}$ of N replicas $d_1, \ldots, d_N$, and the update consists in changing the

paths for every $d_i \in \mathcal{D}$ as shown in the figure. The incremental 2-phase commit technique enables to break down such an update into a sequence of sub-updates, where each sub-update modifies the paths of a distinct set of destinations $\mathcal{D}_j \subset \mathcal{D}$. This will limit the rules added to every switch at any time, at the price of a longer update completion time.

Ultimately, switch-memory consumption remains a fundamental limitation of of 2-phase commit techniques, since tag-matching rules must be added to internal switches sooner or later during an update (or sub-updates). Both the original and the incremental techniques also exhibit other limitations, like the need for packet-header space, the tagging overhead, and complications with middleboxes modifying packet headers and tags [80], [81].

SDN-based update protocols. McGeer [73], [74] presented two protocols to carry out network updates defined on top of OpenFlow. The first update protocol [73] is based on sending packets to the controller during updates, so that the controller can locally store packets until the final flow rules are installed on all the switches. As a result, switch resources (especially, TCAM entries) are saved, at the cost of adding delay on packet delivery, consuming network bandwidth, and requiring the controller to temporarily store packets. The second protocol [74] implements sequences of per-switch rule updates that guarantee PPC: It updates one switch at the time, ensuring that PPC is preserved at every step. In addition, Hua *et al.* [75] initiate a study on how to support PPC in an adversarial setting. They present FOUM, an update protocol where the update sequence is encoded in a packet signed by the controller, sent to one switch, and passed among switches at runtime. FOUM is shown to be robust to packet-tampering and packet dropping attacks.

All those works assume dedicated protocols that are not supported by devices out of the box.

Rule replacement ordering. Further works explore which policies can be supported by only relying on carefully-computed sequences of rule replacements, so as to (i) introduce

Fig. 11. Sequence generated by the FLIP algorithm proposed in [79]. Note that the loop between $v1$ and $v2$ in Step 2 and 3 does not break connectivity, as it is traversed *only once* by packets. This is because $v1$ matches the tag set by $v2$, hence it forwards packets that have already traversed $v2$ directly to d.

no memory overhead, and (ii) be readily supported by all (traditional and SDN) devices, without tag-related issues.

Initial contributions mainly focused on WPE consistency. Prominently, [66] studies how to compute short sequences of rule replacements, for quick updates ensuring that any given flow traverse a single waypoint. The authors propose WayUp, an algorithm that computes WPE-preserving updates spanning 4 rounds. However, they also show that it is not possible to guarantee both WPE and loop-freedom in some cases. Fig. 10 actually shows one case in which any sequence of rule replacements either causes a loop or a WPE consistency violation. Those infeasibility results have been extended to chains of waypoints in [58]. The latter work shows that flexibility in ordering and placing virtualized functions specified by a chain do not make the update problem always solvable. Those two works also prove that it is NP-hard to even decide if there exists a sequence of rule replacements preserving both loop-freedom and WPE (or waypoint chain traversal). Mixed Integer Program (MIP) formulations to find a safe sequence of rule replacements (when any exists) are proposed and evaluated in both cases.

The more general problem of preserving arbitrary policies defined by operators is tackled in [76]. This paper describes an approach to (i) model update-consistency properties as Linear Temporal Logical formulas, and (ii) automatically synthesize SDN updates that preserve input properties. Such a synthesis is performed by an algorithm based on counterexample-guided search and incremental model checking. Experimental results are provided about the scalability of the algorithm (up to networks with 1,000 nodes).

More recent works finally consider the problem of guaranteeing PPC by ordering rule replacements. Vissicchio *et al.* [78] show that this problem can be solved efficiently; they prove that a polynomial-time greedy algorithm called GPIA finds the sequence of per-switch rule replacements that does not violate PPC while updating the maximal number of switches and allowing the maximal parallelism between per-switch updates. The algorithm is based on iteratively simulating the update of every switch which has not been updated yet. For instance, in Fig. 10, GPIA would initially simulate the replacement of next-hops on every node. It would then define the first round of the update under computation by collecting all the switches ($v4$ and $v5$ in this example) whose update does not violate PPC. Then, it would iterate on the remaining nodes, discovering that $s2$ can be updated in the second round without violating PPC. GPIA terminates when there are no switches that can be safely updated. Cerný *et al.* [77] describe a refined version of this

algorithm that avoids the simulation of switch updates thanks to the identification of necessary conditions for safely updating the switches.

Opportunities and limitations of rule replacement ordering for PPC-preserving updates have also been evaluated in [78], by simulating realistic update scenarios, on real network topologies. Results show that ordered rule replacements can rarely complete an update (for example, it could not in Fig. 10, but they can safely update many switches (typically, even more than the 3 out of 7 that it would update in Fig. 10). Those results further motivate rule-replacement algorithms tailored to a more restricted family of policies (like WPE-preserving ones) as well as mixed approaches (employing both rule replacements and duplication, as described below).

Mixed approaches. Following up on their experimental results, Vissicchio *et al.* [78] argue for sequentially computing a safe rule replacement ordering, and applying a scope-limited 2-phase commit variant afterwards. In the example of Fig. 10, this approach would then lead to concatenating the ordered rule replacements on $v4$, $v5$ and $s2$ with the application of a 2-phase commit technique to the sub-network of non-updated switches (i.e., $s1$, $v1$, $v2$ and $v3$). This combination ensures the possibility to always perform the update (contrary to pure rule-replacement approaches) while reducing the update overhead (in switch memory and data-plane) with respect to the original 2-phase commit technique. The same work also generalizes this strategy to hybrid SDNs, potentially running any number both traditional and/or SDN control-planes – a setting in which a brute force (exponential) algorithm might be needed, depending on the nature of control-planes involved in the update.

FLIP, a different algorithm implementing a mixed approach, is described in [79] and detailed in [82]. FLIP jointly optimizes the interleaving of rule replacements and additions (for matching packet tags) so as to preserve arbitrary policies, including PPC. For instance, in Fig. 10, FLIP would compute a sequence of operations (illustrated in Fig. 11) where only $v1$ matches a tag set by $v2$: That is, FLIP needs only 1 additional rule when 2-phase commit would add 4, and sequentially combining replacements and additions would result in 3 additional rules. This implies that FLIP is strictly more powerful (i.e., solves a higher number of update cases) than only relying on rule replacements, exclusively using 2-phase commit, and sequential combining those two approaches. However, it is unclear how FLIP can be used in hybrid SDNs, with more than one control-plane. Also, its time complexity is not polynomial – even if the experiments suggest that FLIP quickly computes short update sequences in realistic networks.

Fig. 12. In this flow migration example, all links have unit bidirectional capacity, and both orange and green flows have unit size as well. The task is to move both the green and orange flows from their initial paths in Fig. 12a to their final ones shown in Fig. 12d. Updating both flows together could lead to the green flow being moved first, inducing congestion, see Fig. 12b. However, this can be avoided by using succinct updates, first moving the orange flow as in Fig. 12c, then the green flow.

C. Related Optimization Problems

Many policy-preserving algorithms face generalized versions of the optimization problems associated to connectivity-preserving updates (see §IV): While the most common objective remains the maximization of parallel operations (to speed-up the update), policy consistency requires that all possible intermediate paths comply with certain regular expressions in addition to being simple (that is, loop-free) paths. Mixed policy-preserving approaches focus on even more general problems where (i) different operations can be interleaved in the output operational sequence (which provides more degrees of freedom in solving the input problems), and (ii) multiple optimization objectives are considered at the same time (typically, maximizing the update parallelism while also minimizing the consumed switch memory).

D. Concluding Remarks and Open Problems

Unsurprisingly, preserving policies requires more sophisticated update techniques, since it is generally harder to extract policy-induced constraints and model the search space. Two major families of solutions have been explored so far. On the one hand, 2-phase commit techniques and update protocols sidestep the algorithmic challenges, at the cost of relying on specific primitives (packet tagging and tag matching) that comes with switch memory consumption. On the other hand, ordering-based techniques directly deal with problem complexities, at the cost of algorithmic simplicity and impossibility to always solve update problems. Finding the best balance between those two extremes is an interesting research direction. Some initial work has started in this direction, with the proposal of algorithms that can interleave different kinds of operations within the computed sequence (see mixed approaches in §V-B). However, many research questions are left open. For example, the computational complexity of solving update problems while mixing rule additions (for packet tagging and matching) with replacements is unknown. Moreover, it is unclear whether the proposed algorithms can be improved exploiting the structure of specific topologies or the flexibility of new devices (like those implementing P4 [83]) – for example, to achieve a better trade-off between switch memory consumption and update speed.

VI. CAPACITY-AWARE CONSISTENCY

Computer networks are inherently capacitated, and respecting resource constraints is hence another important aspect of consistent network updates. Congestion is known to significantly impact throughput and increase latency, therefore negatively impacting user experience and even leading to unpredictable economic loss.

A. Definitions

The capacitated update problem is to migrate from a multi-commodity flow $\mathcal{F}_{old}$ to another multi-commodity flow $\mathcal{F}_{new}$, where consistency is defined as not violating any link capacities and not rate-limiting any flow below its demand in $\min(\mathcal{F}_{old}, \mathcal{F}_{new})$. In few works, e.g., [49], $\mathcal{F}_{new}$ is only implicitly specified by its demands, but not by the actual flow paths. Some migration algorithms will violate consistency properties to guarantee completion, as a consistent migration does not have to exist in all cases.

Typically, four different variants are studied in the literature: First, individual flows may either only take one path (unsplittable) or they may follow classical flow-theory, where the incoming flow at a switch must equal its outgoing flow (splittable). Secondly, flows can take any paths via helper rules in the network during the migration (intermediate paths), or may only be routed along the old or the new paths (no intermediate paths).

To exactly pinpoint congestion-freedom, one would need to take many detailed properties into account, e.g., buffer sizes and ASIC computation times. As such, the standard consistency model does not take this fine-grained approach, but rather aims at avoiding ongoing bandwidth violations and takes a mathematical flow-theory point of view. Introduced by [48], consistent flow migration is captured in the following model: No matter if a flow is using the rules before the update or after the update, the sum of all flow sizes must respect the link capacity.

B. Algorithms and Complexity

1) Algorithms: Current algorithms for capacitated updates of network flows use the seminal work by Reitblatt *et al.* [47] as an update mechanism. Analogously to *per-packet consistency* (cf. §V), one can achieve *per-flow consistency* by a 2-phase commit protocol. While this technique avoids many congestion

TABLE III
COMPACT OVERVIEW OF FLOW MIGRATION ALGORITHMS.

Ref.	Approach	(Un-)splittable model	Interm. paths	Computation	# Updates	Complete (decides if consistent migration exists)
[47]	Install old and new rules, then switch from old to new	Both, move each flow only once	No	Polynomial	1	No bandwidth guarantees
[48]	Partial moves according to free slack capacity s	Splittable	No	Polynomial	$\lceil 1/s \rceil - 1$	Requires slack on flow links
[52]	Greedy traversal of dependency graph	Both, move each flow only once	No	Polynomial	Linear	No (rate-limit flows to guarantee completion)
[54]	MIP of [52]	Both, move each flow only once	No	Exponential	Linear	Yes
[55]	Fix number of x intermediate states ahead of time, optimize via LP	Both	No	Polynomial	Any $x \in \mathbb{N}$	For a given number of intermediate states x, approximate minimum transient congestion (if > 0) by $\log n$ factor
			Yes	*Exponential*		
[55]	... via MIP	Both	Both	*Exponential*	Any $x \in \mathbb{N}$	For any given x yes, but not in general
[48]	Binary search of intermediate states via LP	Splittable	Yes	Polynomial in # of updates	Unbounded	Cannot decide if cons. migration possible
[50]	Create slack with intermediate states, then use partial moves of [48]	Splittable	Yes	Polynomial	Unbounded	Yes
[51]	Split flows along old and new paths	2-Splittable	No	Polynomial	Unbounded	Yes
[49]	Use augmenting flows to find updates	Split., 1 dest., paths not fixed	Yes	Polynomial	Linear	Yes
[84]	Dynamic programming	Unsplittable	Both	*Exponential*	Exponential	Yes
Further practical extensions						
[53]	Extends approach of *SWAN* [48] in a data center setting					
[85]	Extends approach of *Dionysus* [52] with local dependency resolving					
[86]	Extends approach of *Dionysus* [52] with circuit nodes for optical wavelengths					
[87]	Extends approach of *Dionysus* [52] by using switch buffers to break deadlocks					
[88]	Extends approach of *Dionysus* [52] by allowing multiple target states					
[89]	Considers reconfiguration for dynamic flow arrivals					
[90]	Allows (un-)splittable flow migration (move once) with user-spec. deadlines & requirements via MIP (LP heuristic)					
[91]	Does not require tagging of flows in the packet header, flows may take a mix of the old and new paths. For a constant number of flows on directed acyclic graphs (DAGs), a linear-time (fixed parameter tractable) algorithm is provided.					

problems, is not sufficient for bandwidth guarantees: When updating the flows in Fig. 12, if the green flow moves up before orange flow is on its new path, congestion occurs.

Mizrahi and Moses [92] prove that flow swapping is necessary for throughput optimization in the general case, as thus algorithms are needed that do not violate any capacity constraints during the network update, beyond simple flow swapping as well.

An overview of current algorithmic approaches can be found in Table III. In particular, we briefly describe the technique used, e.g., partial moves using slack capacity or dependency graphs. Furthermore, we categorize the algorithm techniques according to their model assumptions (splittability, helper rules), their complexity (computation and # updates), and if they can decide the underlying decision problem under their model assumptions, see also the later Table IV. Note that small changes in the model can lead to different complexities. For example, by not allowing intermediate paths as in, e.g., [52], [54], flow migration is easier to handle, but is also less powerful. In the following, we focus on selected works that introduce general techniques or model ideas.

Slack. The seminal work by Hong *et al.* [48] on *SWAN* introduces the current standard model for capacitated updates. Their algorithmic contribution is two-fold, and also forms the basis for *zUpdate* [53]: First, the authors show that if all flow links have free capacity *slack s*, consistent migration is possible using $\lceil 1/s \rceil - 1$ updates: E.g., if the free capacity is 10%, 9 updates are required, always moving 10% of the links'

capacity to the new flow paths. If the network contains non-critical background traffic, free capacity can be generated for a migration by rate-limiting this background traffic temporarily, see Fig. 13: removing some background traffic allows for consistent migration.


Fig. 13. In this network the task is to migrate consistently from the initial to the final state. If all flows and links have unit size, no consistent migration is possible: the destination has just incoming links of size two. If the flows just have a size of 2/3, one can migrate consistently in $\lceil 1/(1/3) \rceil - 1 = 2$ updates by moving half of the flow size of 1/3 each time in parallel.

LP-formulation. Second, Hong *et al.* [48] present an LP-formulation for splittable flows which provides a consistent migration schedule with x updates, if one exists. By performing a binary search over the number of updates, the number of necessary updates can be minimized. This approach allows for intermediate paths, where the flows can be re-routed anywhere in the network. E.g., consider the example in Fig. 13 with all flows and links having unit size. If there was an additional third route to d, the orange flow could temporarily use this intermediate path: we can then switch the green flow, and

TABLE IV
TABLE SUMMARIZING DECISION PROBLEM RESULTS FOR FLOW MIGRATION.

Flow migration problem	Intermediate paths	Memory restrictions	Decision problem hardness
Unsplittable	Yes	Yes	NP-hard [50],
		No	EXPTIME [84]
	No	Yes	NP-hard [51],
		No	EXPTIME [84]
Unit size	Yes	Yes	NP-hard [50],
		No	EXPTIME [84]
	No	Yes	Open (also for integer size splitting)
		No	
Splittable	Yes	Yes	NP-hard [52]
		No	P [50]
	No	Yes	NP-hard [52]
		No	Open
Move every flow only once	Yes	Yes	Not allowed (model)
		No	
	No	Yes	NP-complete [52]
		No	NP-complete [51]
Node-ordering	Mix of old and new	Yes	NP-hard [91]
		No	

Note: In general, it is unknown if flow migration is in NP if flows can be moved more than once, except for the case of splittable flows without memory restrictions. We note that if a problem is NP-hard without memory restrictions, it is also NP-hard with memory restrictions, as providing sufficient memory is a special case of memory restrictions.

eventually the orange flow could be moved to its desired new path.

Spread flows over the network. Brandt *et al.* [50] prove that splittable flow migration is always decidable in polynomial time, by providing an algorithm that attempts to create slack capacity on all links. The fundamental idea is to keep splitting flows along new paths, until slack is obtained such that the algorithm of Hong *et al.* [48] is applicable. For example in Fig. 12a, one can proceed as follows: 1) split the orange flow equally along the old and new path, and afterwards 2) route a quarter of the green flow via either $s2$ or $d2$. The correctness of their approach relies on an augmenting flow techniques, we refer the reader to [50] for the intricate details.

Dynamicity & dependency graphs. Jin *et al.* [52] also consider the variable update times of switches in the network. To this end, inspired by [60], they build a dependency graph of the individual updates, greedily sending out updates once the respective pre-conditions are satisfied. For example, assume a flow f_1 can be migrated to its new path when least one of the flows f_2, f_3 was moved, but f_2, f_3 take different (unknown) time-spans to complete the move: the fastest method is to dynamically wait until either f_2, f_3 moves, any pre-computed schedule will have a longer makespan under adversarial conditions. When this greedy traversal of the dependency graph results in a deadlock, flows are rate-limited to guarantee progress. We note that it is not clear how to extend the dependency graph idea to intermediate paths.

Jointly optimize migration & new paths. Brandt *et al.* [49] introduce the idea that the new flow paths should not be part of the problem input, but rather be computed jointly with the migration schedule. This idea aims at 1) speeding up the migration process and 2) allowing more problem instances to be solved. As an illustration, consider the problem in Fig. 13: from a pure admission perspective, no migration is needed, both flow demands are already satisfied in the initial state. Interestingly, for a single destination (but multiple commodities), consistent migration is *always* possible in this model, as long as there is some way to admit all flows without violating capacities. A framework for general multi-commodity flows is still missing.

Ghandi *et al.* [88] also observe that flows can have multiple migration options, as networks are often built with redundancy in mind. To this end, they first compute multiple choices for new flow paths, optimizing for close-to-optimal path properties and few stages in the resulting dependency graph. This allows them to dynamically speed up the execution of the consistent network updates, depending on the runtime conditions.

Node-ordering instead of 2-phase commit Amiri *et al.* [91] propose to identify flows only by their source and destination, removing flow version numbers from the packet header ("tagging"). Their approach reduces complexity overhead, but does not permit the use of 2-phase commit techniques. Conceptually, each node has an old and a new forwarding rule for each flow, where the challenge is how to *order* these updates, without inducing congestion or forwarding loops. For an intuition, recall Fig. 5d, and let the old rules be marked in solid blue, with the new rules being in dash-dotted green: over multiple rounds, the routing rules converge to the new state.

2) Complexity: The complexity of capacitated updates can roughly be summarized as follows: Problems involving splittable flows can be decided in polynomial time, while restrictions such as unsplittable flows or memory limits turn the problem NP-hard, see Table IV. For unsplittable flows, an exponential time algorithm exists. In a way, the capacitated update problems differs from related network update problems in that it is not always solvable in a consistent way. On the other hand, e.g., per-packet/flow consistency can always be maintained by a 2-phase commit, and loop-free updates for a single destination can always be performed in a linear number of updates.

One standard approach in recent work for flow migration is linear (splittable flows) or integer programing (unsplittable flows): With the number of intermediate configurations x as an input, it is checked if a consistent migration with x intermediate states exists. Should the answer be yes, then one can use a binary search over x to find the fastest schedule. This idea originated in *SWAN* [48] for splittable flows, and was later

TABLE V
COMPACT OVERVIEW OF FLOW MIGRATION HARDNESS TECHNIQUES AND RESULTS.

Ref.	Reduction via	(Un-)splittable model	Interm. paths	Memory limits	Decision prob. in general	Optimization problems/remarks
[52]	Partition	Splittable	No	Yes	NP-hard	NP-complete if every flow may only move once
[52]	Partition	Splittable	No	No	–	NP-hard (fewest rule modifications)
[50]	–	Splittable	Yes	No	P	Fastest schedule can be of unbounded length, LP for new reachable demands if cannot migrate
[51]	–	2-Splittable	No	No	P	studies slightly different model
[50]	(MAX) 3-SAT	Unsplittable	Yes	No	NP-hard (also for unit size)	NP-hard to approx. additive error of flow removal for consistency better than $7/8 + \varepsilon$
[55]	Partition	Unsplittable	Yes/No	No	–	NP-hard (fastest schedule)
[51]	Partition	Unsplittable	No	No	NP-hard	Stronger consistency model, but proof carries over
[90]	Partition & Subset sum	Unsplittable	No	No	–	NP-hard for 3-update schedule
[84]	Disjoint paths	Unsplittable	Yes	No	NP-hard (also for unit size)	NP-hard for already 2 unit size flows
[91]	3-SAT	Node-ordering	Mix of old & new	No	NP-hard (also for unit size)	NP-hard on DAGs, NP-hard on general graphs for already 2 unit size flows

extended to other models, cf. Table III.

However, the LP-approach via binary search (likewise for the integer one) suffers from the drawback that it is only complete if the model is restricted: If x is unbounded, then one can only decide whether a migration with x updates exists, but not whether there is no migration schedule with y steps, for some $y > x$. Additionally, it is not even clear to what complexity class the general capacitated update problem belongs to, cf. the decision problem hardness column of Table IV.

The only exception arises in case of splittable flows without memory restrictions, where either an (implicit) schedule or a certificate that no consistent migration is possible, is found in polynomial time [50]. The authors use a combinatorial approach not relying on linear programming. Adding memory restrictions turns this problem NP-hard as well [52].

If the model is restricted to allow every flow only to be moved once (from the old path to the new path), then the capacitated update problem becomes NP-complete [51], [52]: Essentially, as the number of updates is limited by the number of flows, the problem is in NP. In this specific case, one can also approximate the minimum congestion for unsplittable flows in polynomial time by randomized rounding [55].

Hardly any (in-)approximability results exist today, and most work relies on reductions from the Partition problem, cf. Table V. The only result that we are aware of is via a reduction from MAX 3-SAT, which also applies to unit size flows [50].

C. Related Optimization Problems

In a practical setting, splitting flows is often realized via deploying multiple unsplittable paths, which is an NP-hard optimization problem as well, both for minimizing the number of paths and for maximizing k-splittable flows, cf. [93], [94]. Another popular option is to split the flows at the routers using hash functions; other major techniques are flow(let) caches and round-robin splitting, cf. [95]. Nonetheless, splitting flows along multiple paths can lead to packet reordering problems, which need to be handled by further techniques, see, e.g., [96].

Many of the discussed flow migration works rely on linear programming formulations: Even though their runtime is polynomial in theory, the timely migration of large networks

with many intermediate states is currently problematic in practice [48]. If the solution takes too long to compute, the to-be solved problem might no longer exist, a problem only made worse when resorting to (NP-hard) integer programming for unsplittable flows. As such, some tradeoff has to be made between finding an optimal solution and one that can actually be deployed.

Orthogonal to the problem of consistent flow migration is the approach of scheduling flows beforehand, not changing their path assignments in the network during the update. We refer to the recent works by Kandula *et al.* [97] and Perry *et al.* [98] for examples. Game-theoretic approaches have also been considered, e.g., [99]. Lastly, the application of model checking does not cover bandwidth problem restrictions yet [76], [100].

D. Concluding Remarks and Open Problems

The classification of the complexity of flow migration still poses many questions, cf. Table V: If every flow can only be moved once, then the migration (decision) problem is clearly in NP. However, what is the decision complexity if flows can be moved arbitrarily often, especially with intermediate paths? Is the "longest" fastest update schedule for unsplittable flows: linear, polynomial or exponential? In other words, is the problem complete in NP, PSPACE, or EXPTIME? Related questions are also open for flows of unit or integer size in general.

The problem of migrating splittable flows without memory limits and without intermediate paths is still not studied either: It seems as if the methods of [50] and [51] also apply to this case, but a formal proof is missing. Another open issue which researchers recently started to consider concerns how to exploit traffic engineering flexibilities and "helper rules" to jointly optimize update scheduling and route selection [101].

Lastly, it would be interesting to compare the power and performance of the node-ordering approach introduced by Amiri *et al.* [91] to using the 2-phase commit of Reitblatt *et al.* [47].

VII. Relaxing Safety Guarantees

So far we studied network updates assuming that consistency in the respective model must be maintained, e.g., no forwarding

loops must appear at any time. In situations where the computation is no longer tractable or the consistency property cannot be maintained at all, some works proposed to break consistency in a controlled manner.

A first approach in this direction consists in trying to minimize the time spent in an inconsistent state, with underlying protocols being able to correct the induced problems (e.g., dropped packets are re-transmitted), as done in a production environment in Google's *B4* [9].

An alternative is to synchronize the clocks in the switches so that network updates can be performed simultaneously: With perfect clock synchronization, lossless communications and switch execution behavior, loop freedom could be maintained. As the standard Network Time Protocol (NTP) does not have sufficient synchronization behavior, the Precision Time Protocol (PTP) was adapted to SDN environments in [102], [103], achieving microsecond accuracy in experiments. This obviously comes with additional message overhead for time synchronization in the whole network. An introduction and overview of so-called timed consistent updates is provided in [104]. Zheng *et al.* [105] study the use of timed consistent updates in order to prevent congestion in the context of flow migration, in combination with latency considerations.

Nonetheless, in some situations synchronized updates can be considered optimal: E.g., consider the case in Fig. 13 where two unsplittable flows need to be swapped [92], with no alternative paths in the network available for the final links. Then, synchronizing the new flow paths can minimize the induced congestion [106]. Synchronized updates cannot guarantee packet consistency on their own, as packets that are currently en-route may still encounter changed forwarding rules at the next switch. Time can also be used similarly to a 2-phase commit though, by analogously using timestamps in the packet header as tags during the update [107], with [107] also showing an efficient implementation using timestamp-based TCAM ranges. Additional memory, as in the 2-phase commit approach of Reitblatt *et al.* [47], will be used for this method, but packets only need to be tagged implicitly by including the timestamp (where often 1 bit suffices [108], [107]). In [109] some additional methods are discussed on how to guarantee packet consistency by temporarily storing traffic at the switches.

Despite all those advantages, the proposed clock synchronization approaches do not prevent unpredictable variations of command execution time on network switches [52], motivating the need for prediction-based scheduling methods [110], [111]. Even worse, failures have an intrinsic, unavoidable cost in this approach. If a switch fails to update at all, the network can stay in an inconsistent state until the controller is notified and takes appropriate actions (e.g., computing another update). The same risk of inconsistencies holds if controller-to-switch messages are delayed or lost. In contrast, techniques based on sequential approaches can verify the application of sent update commands one by one, possibly moving forward (to the next update) or back (if a command is not received or not yet applied) with no risk of incurring safety violations.

VIII. FROM THEORY TO PRACTICE

As a complement to the previously-described theoretical and algorithmic results, we now provide an overview on practical challenges to ensure consistent network updates. We also describe how previous works tackled those challenges in order to build automated systems that can automatically carry out consistent updates. A brief overview indicating the current status is presented in Table VI.

1) **Ensuring basic communication with network devices:** Automated update systems classically rely on a logically-centralized coordinator, which must interact with network devices to instruct them to apply operations (in a given order). Such a device-coordinator interaction requires a communication channel. Update coordinators in traditional networks typically exploit the command line interface of devices, as noted in, e.g., [112], [26]. For SDNs, the interaction is simplified by their very architecture, since the coordinator is typically embodied by the SDN controller which must be already able to program (e.g., through OpenFlow [38] or similar protocols) and monitor (e.g., thanks to a Network Information Base [113]) the controlled devices.

2) **Applying operational sequences, step by step:** Both devices and the device-coordinator communication are not necessarily reliable. For example, messages sent by the coordinator may be lost or not be applied by all devices upon reception [52], [114]. Those possibilities are typically taken into account in the computation of the update sequence (see §III). However, an effective update system must also ensure that operations are actually applied as in the computed sequences, e.g., that all operations in one update step are actually executed on the switches before sending operations in the next step. To this end, a variety of strategies are applied in the literature, from dedicated monitoring approaches (based on available network primitives like status-checking commands and protocols [112], lower-level packet cloning mechanisms [26], or active probing packets [115]) to acknowledgement-based protocols implemented by SDN devices [114].

3) **Working around device limitations:** Applying carefully-computed operational sequences ensures update consistency but not necessarily performance (e.g., speed), as the latter also depends on device efficiency in executing operations. This aspect has been analyzed by several works, especially focused on SDN updates which are more likely to be applied in real-time (e.g., even to react to a failure). It has been pointed out that current SDN device limitations impact update performance in two ways. First, SDN switches are not yet fast to change their packet-processing rules, as highlighted by several measurement studies. For example, in the Devoflow [116] paper, the authors showed that the rate of statistics gathering is limited by the size of the flow table and is negatively impacted by the flow setup rate. In 2015, He *et al.* [117] experimentally demonstrated the high rule installation latency of four different types

TABLE VI
TABLE INDICATING THE CURRENT STATUS OF PRACTICAL CHALLENGES TO ENSURE CONSISTENT UPDATES.

Challenge	Approach / Measurement	Status
(1) Ensuring basic communication	A: command line interface of devices	Available in (traditional) networks [112], [26]
	A: SDN controller programs and monitors	Deployed via, e.g., OpenFlow [38] and Network Information Bases [113]
(2) Applying operational sequences	M: coordinator messages lost respectively not applied by all devices	Measured on commodity hardware [52], [114]
	A: status-checking commands and protocols	Evaluated in testbed [112]
	A: lower-level packet cloning mechanisms	Evaluated in simulations [26], data set at https://inl.info.ucl.ac.be/softwares
	A: active probing packets	Evaluated in (small) testbed [115]
	A: acknowledgement-based protocols	Evaluated in small testbed [114]
(3) Working around device limitations	M: flow table size and setup rate limits statistics gathering	Measured on a common OpenFlow implementation on a switch [116]
	M: high rule installation latency	Measured on various types of (SDN) switches [117], [118], [119], [52]
	A: adapt schedules dynamically	Evaluated in testbed [52]
	A: eliminate redundant updates	Evaluated in Mininet [120], algorithms [56]
(4) Multiple control-plane conflicts	A: pro-actively specifying computation of final rules	Implementation of [44] available at http://www.frenetic-lang.org
	A: implementing coordination and locking primitives on switches	Implementation of [121] available at https://github.com/lironsc/of-sync-lib
	A: reactively detecting and possibly resolving conflicts	Algorithms [122]
	A: meta-algorithms	General theory [41]
(5) Updating the control-plane	A: hypervisor maintains history	Evaluated in Mininet [123]
	A: explicit state transfer	Evaluated in Mininet [124]
(6) Events occurring during updates	M: impact of link failures (IGP)	Evaluated in simulations [27], data set at https://inl.info.ucl.ac.be/softwares
	A: enforcing per-packet consistency against packet-tampering adversary	Evaluation in small testbed [75]

of production SDN switches. This confirmed the results of independent studies [118], [119] providing a more in-depth look into switch performance across various vendors. Second, rule installation time can highly vary over time, independently on any switch, because it is a function of runtime factors like already-installed rules and data-plane load. The measurement campaign on real OpenFlow switches performed in Dionysus [52] indeed shows that rule installation delay can vary from seconds to minutes. Update systems are therefore engineered to mitigate the impact of those limitations – despite not avoiding per-rule update bottlenecks. Prominently, Dionysus [52] significantly reduces multi-switch update latency by carefully scheduling operations according to dynamic switch conditions. In addition, CoVisor [120] and [56] minimize the number of rule updates sent to switches through eliminating redundant updates.

4) **Avoiding conflicts between multiple control-planes:** For availability, performance, and robustness, network control-planes are often physically-distributed, even when logically centralized (as in the case of SDNs with replicated controllers). For updates of traditional networks, the control-plane distribution is straightforwardly taken into account, since it is encompassed in the update problem definition (see §II). In contrast, additional care must be applied to SDN networks with multiple controllers: if several controllers try to update network devices at the same time, one controller may override rules installed by another, impacting the correctness of the update (both during and after the update itself). This requires to solve potential conflicts between controllers, e.g., by pro-actively specifying how the final rules have to be computed (e.g., [44]), by implementing coordination and locking primitives on switches (e.g., [121]), or by reactively detecting and possibly resolving conflicts (e.g., [122]). A generalization of the above setting consists in considering multiple control-planes that may

be either all distributed, all centralized, or hybrid (some distributed and some centralized). Potential conflicts and general meta-algorithms to ensure consistent updates in those cases are described in [41].

5) **Updating the control-plane:** In traditional networks, data-plane changes can only be enforced by changing the configuration of control-plane protocols (e.g., IGPs). In contrast, the most studied case for SDN updates considers an unmodified controller that has to change the packet-processing rules on network switches. Nevertheless, a few works also considered the problem of entirely replacing the SDN controller itself, e.g., upgrading it to a new version or replacing an old controller with a newer one. In particular, HotSwap [123] describes an architecture that enables the replacement of an SDN controller, by relying on a hypervisor that maintains a history of network events. As an alternative, explicit state transfer is used to design and implement the Morpheus controller platform in [124].

6) **Dealing with events occurring during an update:** Operational sequences computed by update algorithms forcedly assume stable network conditions. In practice, however, unpredictable events, like failures, can modify the network behavior concurrently and independently from the operations performed to update the network. While concurrent events can be very unlikely (especially for fast updates), by definition they cannot be prevented. A few contributions assessed the impact of such unpredictable events on the update safety. For instance, the impact of link failures on SITN-based IGP reconfigurations is experimentally evaluated in [27]. A more systematic approach is taken by the recent FOUM work [75], that aims at guaranteeing per-packet consistency in the presence of an adversary able to perform packet-tampering and packet-dropping attacks.

IX. FUTURE RESEARCH DIRECTIONS

While we have already identified specific open research questions in the corresponding sections, we now discuss more general areas which we believe deserve more attention by the research community in the future.

1) **Charting the complexity landscape:** Researchers have only started to understand the computational complexities underlying the network update problem. In particular, many NP-hardness results have been derived for general problem formulations for all three of our consistency models: connectivity consistency, policy consistency, and capacity consistency. So far, only for a small number of specific models polynomial-time optimal algorithms are known. Even less is known about approximation algorithms. Hence, more research efforts would be needed to chart a clearer picture of the complexity landscape for network update problems. We expect that some of these insights will also have interesting implications on classic optimization problems.

2) **Refining and tailoring our models:** While today's network models capture well the *fundamental* constraints and tradeoffs in consistent network update problems, these models are relatively simple. Today, SDN is used and discussed in various contexts and for different reasons, from network virtualization in datacenters, over traffic engineering in the wide-area network to network slicing in emerging 5G applications [125], [126], [127]. These use cases come with different *specific* requirements on the consistency (e.g., whether per-packet consistency is strictly needed) and performance (e.g., tolerable number of rounds), and also differ in terms of the available *knobs* which can be used for the network update (e.g., helper rules may undesirable in the context of network slicing). Accordingly, our models need to be refined and tailored towards specific use cases.

3) **Considering new update problems:** We expect future update techniques to ensure consistency of higher-level network requirements (like NFV, path delay, etc.), the same way as recent SDN controllers are supporting them. More dynamic and stateful applications also introduce new consistency criteria (e.g., [128], [129]) that might be interesting to take into account in future update techniques.

4) **Dealing with distributed control planes:** We believe that researchers have only started to understand the design and implication of SDN control planes which are distributed (not only vertically but also horizontally [130], [131], [132]), e.g., for dependability and performance purposes. In particular, distributed SDN control-planes introduce additional challenges in terms of consistent network updates and controller coordination.

X. CONCLUSION

The purpose of this survey was to provide researchers active in or interested in the field of network update problems (and in particular Software-Defined Networks) with an overview of the state-of-the-art, including models, techniques, impossibility results as well as practical challenges. We also presented a historical perspective and discussed the fundamental new challenges introduced in Software-Defined Networks, also relating them to classic graph-theoretic optimization problems. Finally, we have identified open questions for future research.

REFERENCES

[1] Barefoot Networks, "The world's fastest and most programmable networks (white paper)," https://barefootnetworks.com/white-paper/the-worlds-fastest-most-programmable-networks/, 2016.

[2] M. Imbriaco, "Network problems last friday," GitHub: https://github.com/blog/1346networkproblemslastfriday, December 2012.

[3] J. Jackson, "Godaddy blames outage on corrupted router tables," PCWorld: http://www.pcworld.com/article/262142/godaddy_blames_outage_on_corrupted_router_tables.html, September 2012.

[4] R. Mohan, "Storms in the cloud: Lessons from the amazon cloud outage," Security Week: http://www.securityweek.com/storms-cloud-lessons-amazon-cloud-outage, June 2011.

[5] United, "United Airlines Restoring Normal Flight Operations Following Friday Computer Outage," http://newsroom.united.com/news-releases?item=124170, June 2011.

[6] P. Kazemian, G. Varghese, and N. McKeown, "Header space analysis: Static checking for networks," in *Proc. USENIX NSDI 2*, 2012. [Online]. Available: https://www.usenix.org/conference/nsdi12/technical-sessions/presentation/kazemian

[7] A. Khurshid, W. Zhou, M. Caesar, and B. Godfrey, "Veriflow: verifying network-wide invariants in real time," *Computer Communication Review*, vol. 42, no. 4, pp. 467–472, 2012. [Online]. Available: http://doi.acm.org/10.1145/2377677.2377766

[8] N. Feamster, J. Rexford, and E. W. Zegura, "The road to SDN: an intellectual history of programmable networks," *Computer Communication Review*, vol. 44, no. 2, pp. 87–98, 2014. [Online]. Available: http://doi.acm.org/10.1145/2602204.2602219

[9] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, "B4: Experience with a globally-deployed software defined wan," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, pp. 3–14, Aug. 2013. [Online]. Available: http://doi.acm.org/10.1145/2534169.2486019

[10] J. Moy, P. Pillay-Esnault, and A. Lindem, "Graceful OSPF Restart," RFC 3623, 2003.

[11] A. Shaikh, R. Dube, and A. Varma, "Avoiding instability during graceful shutdown of multiple OSPF routers," *IEEE/ACM Trans. Netw.*, vol. 14, no. 3, pp. 532–542, 2006. [Online]. Available: http://doi.acm.org/10.1145/1143396.1143403

[12] M. Shand and L. Ginsberg, "Restart Signaling for IS-IS," RFC 5306, 2008.

[13] P. François and O. Bonaventure, "Avoiding transient loops during the convergence of link-state routing protocols," *IEEE/ACM Trans. Netw.*, vol. 15, no. 6, pp. 1280–1292, 2007. [Online]. Available: http://doi.acm.org/10.1145/1373476.1373482

[14] J. Fu, P. Sjödin, and G. Karlsson, "Loop-free updates of forwarding tables," *IEEE Trans. Network and Service Management*, vol. 5, no. 1, pp. 22–35, 2008. [Online]. Available: http://dx.doi.org/10.1109/TNSM.2008.080103

[15] L. Shi, J. Fu, and X. Fu, "Loop-free forwarding table updates with minimal link overflow," in *Proc. IEEE ICC*, 2009. [Online]. Available: http://dx.doi.org/10.1109/ICC.2009.5199146

[16] M. Shand and S. Bryant, "A Framework for Loop-Free Convergence," IETF, RFC 5715, January 2010.

[17] R. Keralapura, C.-N. Chuah, and Y. Fan, "Optimal strategy for graceful network upgrade," in *Proc. ACM INM*, 2006. [Online]. Available: http://doi.acm.org/10.1145/1162638.1162652

[18] S. Raza, Y. Zhu, and C. Chuah, "Graceful network operations," in *Proc. IEEE INFOCOM*, 2009. [Online]. Available: http://dx.doi.org/10.1109/INFCOM.2009.5061932

[19] ——, "Graceful network state migrations," *IEEE/ACM Trans. Netw.*, vol. 19, no. 4, pp. 1097–1110, 2011. [Online]. Available: http://dx.doi.org/10.1109/TNET.2010.2097604

[20] P. François, M. Shand, and O. Bonaventure, "Disruption free topology reconfiguration in OSPF networks," in *Proc. IEEE INFOCOM*, 2007. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2007.19

[21] F. Clad, P. Mérindol, J. Pansiot, P. François, and O. Bonaventure, "Graceful convergence in link-state IP networks: A lightweight algorithm ensuring minimal operational impact," *IEEE/ACM Trans. Netw.*, vol. 22, no. 1, pp. 300–312, 2014. [Online]. Available: http://dx.doi.org/10.1109/TNET.2013.2255891

[22] F. Clad, P. Mérindol, S. Vissicchio, J. Pansiot, and P. François, "Graceful router updates in link-state protocols," in *Proc. IEEE ICNP*, 2013. [Online]. Available: http://dx.doi.org/10.1109/ICNP.2013.6733587

[23] F. Clad, S. Vissicchio, P. Mérindol, P. François, and J. Pansiot, "Computing minimal update sequences for graceful router-wide reconfigurations," *IEEE/ACM Trans. Netw.*, vol. 23, no. 5, pp. 1373–1386, 2015. [Online]. Available: http://dx.doi.org/10.1109/TNET.2014.2332101

[24] G. Herrero and J. van der Ven, *Network Mergers and Migrations: Junos Design and Implementation.* Wiley, 2010.

[25] I. Pepelnjak, "Changing the Routing Protocol in Your Network," ipSpace: http://blog.ipspace.net/2007/10/change-routing-protocol-in-your-network.html, October 2007.

[26] L. Vanbever, S. Vissicchio, C. Pelsser, P. François, and O. Bonaventure, "Seamless network-wide IGP migrations," in *Proc. ACM SIGCOMM*, 2011. [Online]. Available: http://doi.acm.org/10.1145/2018436.2018473

[27] ——, "Lossless migrations of link-state igps," *IEEE/ACM Trans. Netw.*, vol. 20, no. 6, pp. 1842–1855, 2012. [Online]. Available: http://dx.doi.org/10.1109/TNET.2012.2190767

[28] F. Le, G. G. Xie, D. Pei, J. Wang, and H. Zhang, "Shedding light on the glue logic of the internet routing architecture," in *Proc. ACM SIGCOMM*, 2008. [Online]. Available: http://doi.acm.org/10.1145/1402958.1402964

[29] S. Vissicchio, L. Vanbever, L. Cittadini, G. G. Xie, and O. Bonaventure, "Safe routing reconfigurations with route redistribution," in *Proc. IEEE INFOCOM*, 2014. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2014.6847940

[30] L. Vanbever, S. Vissicchio, L. Cittadini, and O. Bonaventure, "When the cure is worse than the disease: The impact of graceful IGP operations on BGP," in *Proc. IEEE INFOCOM*, 2013. [Online]. Available: http://dx.doi.org/10.1109/INFCOM.2013.6567025

[31] P. François, O. Bonaventure, B. Decraene, and P. Coste, "Avoiding disruptions during maintenance operations on BGP sessions," *IEEE Trans. Network and Service Management*, vol. 4, no. 3, pp. 1–11, 2007. [Online]. Available: http://dx.doi.org/10.1109/TNSM.2007.021102

[32] Y. Wang, E. Keller, B. Biskeborn, J. van der Merwe, and J. Rexford, "Virtual routers on the move: Live router migration as a network-management primitive," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 4, pp. 231–242, Aug. 2008. [Online]. Available: http://doi.acm.org/10.1145/1402946.1402985

[33] E. Keller, J. Rexford, and J. E. van der Merwe, "Seamless BGP migration with router grafting," in *Proc. USENIX NSDI*, 2010. [Online]. Available: http://www.usenix.org/events/nsdi10/tech/full_papers/keller.pdf

[34] S. Vissicchio, L. Vanbever, C. Pelsser, L. Cittadini, P. François, and O. Bonaventure, "Improving network agility with seamless BGP reconfigurations," *IEEE/ACM Trans. Netw.*, vol. 21, no. 3, pp. 990–1002, 2013. [Online]. Available: http://dx.doi.org/10.1109/TNET.2012.2217506

[35] N. Kushman, S. Kandula, D. Katabi, and B. M. Maggs, "R-BGP: staying connected in a connected world," in *Proc. USENIX NSDI*, 2007. [Online]. Available: http://www.usenix.org/events/nsdi07/tech/kushman.html

[36] R. Alimi, Y. Wang, and Y. R. Yang, "Shadow configuration as a network management primitive," in *Proc. ACM SIGCOMM*, 2008. [Online]. Available: http://doi.acm.org/10.1145/1402958.1402972

[37] J. P. John, E. Katz-Bassett, A. Krishnamurthy, T. E. Anderson, and A. Venkataramani, "Consensus routing: The internet as a distributed system," in *Proc. USENIX NSDI*, 2008. [Online]. Available: http://www.usenix.org/events/nsdi08/tech/full_papers/john/john.pdf

[38] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008. [Online]. Available: http://doi.acm.org/10.1145/1355734.1355746

[39] D. Kreutz, F. M. V. Ramos, P. J. E. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015. [Online]. Available: https://doi.org/10.1109/JPROC.2014.2371999

[40] A. Vahdat, D. Clark, and J. Rexford, "A purpose-built global network: Google's move to sdn," *Queue*, vol. 13, no. 8, pp. 100:100–100:125, Oct. 2015. [Online]. Available: http://doi.acm.org/10.1145/2838344.2856460

[41] S. Vissicchio, L. Cittadini, O. Bonaventure, G. G. Xie, and L. Vanbever, "On the co-existence of distributed and centralized routing control-planes," in *Proc. IEEE INFOCOM*, 2015. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2015.7218413

[42] N. Foster, A. Guha, M. Reitblatt, A. Story, M. J. Freedman, N. P. Katta, C. Monsanto, J. Reich, J. Rexford, C. Schlesinger, D. Walker, and R. Harrison, "Languages for software-defined networks," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 128–134, 2013. [Online]. Available: http://dx.doi.org/10.1109/MCOM.2013.6461197

[43] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker, "Ethane: taking control of the enterprise," in *Proc. ACM SIGCOMM*, 2007. [Online]. Available: http://doi.acm.org/10.1145/1282380.1282382

[44] C. Monsanto, J. Reich, N. Foster, J. Rexford, and D. Walker, "Composing software defined networks," in *Proc. USENIX NSDI*, 2013. [Online]. Available: https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/monsanto

[45] P. Kazemian, M. Chan, H. Zeng, G. Varghese, N. McKeown, and S. Whyte, "Real time network policy checking using header space analysis," in *Proc. USENIX NSDI*, 2013. [Online]. Available: https://www.usenix.org/conference/nsdi13/technical-sessions/presentation/kazemian

[46] M. Reitblatt, N. Foster, J. Rexford, and D. Walker, "Consistent updates for software-defined networks: change you can believe in!" in *Proc. ACM HotNets*, 2011. [Online]. Available: http://doi.acm.org/10.1145/2070562.2070569

[47] M. Reitblatt, N. Foster, J. Rexford, C. Schlesinger, and D. Walker, "Abstractions for network update," in *Proc. ACM SIGCOMM*, 2012. [Online]. Available: http://doi.acm.org/10.1145/2342356.2342427

[48] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven wan," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 4, pp. 15–26, Aug. 2013. [Online]. Available: http://doi.acm.org/10.1145/2534169.2486012

[49] S. Brandt, K.-T. Foerster, and R. Wattenhofer, "Augmenting flows for the consistent migration of multi-commodity single-destination flows in sdns," *Pervasive and Mobile Computing*, vol. 36, pp. 134–150, 2017. [Online]. Available: https://doi.org/10.1016/j.pmcj.2016.09.012

[50] ——, "On consistent migration of flows in sdns," in *Proc. IEEE INFOCOM*, 2016. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2016.7524332

[51] K.-T. Foerster and R. Wattenhofer, "The power of two in consistent network updates: Hard loop freedom, easy flow migration," in *Proc. ICCCN*. IEEE, 2016. [Online]. Available: http://dx.doi.org/10.1109/ICCCN.2016.7568583

[52] X. Jin, H. H. Liu, R. Gandhi, S. Kandula, R. Mahajan, M. Zhang, J. Rexford, and R. Wattenhofer, "Dynamic scheduling of network updates," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 539–550, Aug. 2014. [Online]. Available: http://doi.acm.org/10.1145/2740070.2626307

[53] 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. [Online]. Available: http://doi.acm.org/10.1145/2534169.2486005

[54] L. Luo, H. Yu, S. Luo, and M. Zhang, "Fast lossless traffic migration for SDN updates," in *Proc. IEEE ICC*, 2015. [Online]. Available: http://dx.doi.org/10.1109/ICC.2015.7249247

[55] J. Zheng, H. Xu, G. Chen, and H. Dai, "Minimizing transient congestion during network update in data centers," in *Proc. ICNP*. IEEE, 2015. [Online]. Available: http://dx.doi.org/10.1109/ICNP.2015.33

[56] S. Dudycz, A. Ludwig, and S. Schmid, "Can't touch this: Consistent network updates for multiple policies," in *Proc. IEEE/IFIP DSN*, 2016. [Online]. Available: http://dx.doi.org/10.1109/DSN.2016.21

[57] K.-T. Foerster, R. Mahajan, and R. Wattenhofer, "Consistent updates in software defined networks: On dependencies, loop freedom, and blackholes," in *Proc. IFIP Networking*. IEEE, 2016. [Online]. Available: http://dx.doi.org/10.1109/IFIPNetworking.2016.7497232

[58] A. Ludwig, S. Dudycz, M. Rost, and S. Schmid, "Transiently secure network updates," in *Proc. ACM SIGMETRICS*, 2016. [Online]. Available: http://doi.acm.org/10.1145/2896377.2901476

[59] K. T. Foerster, A. Ludwig, J. Marcinkowski, and S. Schmid, "Loop-free route updates for software-defined networks," *IEEE/ACM Transactions on Networking*, vol. PP, no. to appear, pp. 1–14, 2018.

[60] R. Mahajan and R. Wattenhofer, "On consistent updates in software defined networks," in *Proc. ACM HotNets*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2535771.2535791

[61] D. Li, S. Wang, K. Zhu, and S. Xia, "A survey of network update in SDN," *Frontiers of Computer Science*, vol. 11, no. 1, pp. 4–12, 2017. [Online]. Available: https://doi.org/10.1007/s11704-016-6125-y

[62] A. Korman, S. Kutten, and D. Peleg, "Proof labeling schemes," *Distributed Computing*, vol. 22, no. 4, pp. 215–233, 2010. [Online]. Available: https://doi.org/10.1007/s00446-010-0095-3

[63] K.-T. Foerster, T. Luedi, J. Seidel, and R. Wattenhofer, "Local checkability, no strings attached: (a)cyclicity, reachability, loop free updates in sdns," *Theor. Comput. Sci.*, vol. 709, pp. 48–63, 2018. [Online]. Available: https://doi.org/10.1016/j.tcs.2016.11.018

[64] S. A. Amiri, A. Ludwig, J. Marcinkowski, and S. Schmid, "Transiently consistent SDN updates: Being greedy is hard," in *Proc. SIROCCO*, 2016. [Online]. Available: https://doi.org/10.1007/978-3-319-48314-6_25

[65] A. Ludwig, J. Marcinkowski, and S. Schmid, "Scheduling loop-free network updates: It's good to relax!" in *Proc. ACM PODC*, 2015. [Online]. Available: http://doi.acm.org/10.1145/2767386.2767412

[66] A. Ludwig, M. Rost, D. Foucard, and S. Schmid, "Good network updates for bad packets: Waypoint enforcement beyond destination-based routing policies," in *Proc. ACM HotNets*, 2014. [Online]. Available: http://doi.acm.org/10.1145/2670518.2673873

[67] G. Even, J. Naor, B. Schieber, and M. Sudan, "Approximating minimum feedback sets and multicuts in directed graphs," *Algorithmica*, vol. 20, no. 2, pp. 151–174, 1998. [Online]. Available: http://dx.doi.org/10.1007/PL00009191

[68] R. Hassin and S. Rubinstein, "Approximations for the maximum acyclic subgraph problem," *Information Processing Letters*, vol. 51, no. 3, pp. 133 – 140, 1994. [Online]. Available: http://www.sciencedirect.com/science/article/pii/0020019094000867

[69] M. Middendorf, "Supersequences, runs, and cd grammar systems," *Developments in Theoretical Computer Science*, vol. 6, pp. 101–114, 1994.

[70] M. Lewin, D. Livnat, and U. Zwick, "Improved rounding techniques for the MAX 2-sat and MAX DI-CUT problems," in *Proc. IPCO*, 2002. [Online]. Available: http://link.springer.de/link/service/series/0558/bibs/2337/23370067.htm

[71] M. Kaminski, P. Medvedev, and M. Milanic, "Shortest paths between shortest paths," *Theor. Comput. Sci.*, vol. 412, no. 39, pp. 5205–5210, 2011. [Online]. Available: https://doi.org/10.1016/j.tcs.2011.05.021

[72] N. P. Katta, J. Rexford, and D. Walker, "Incremental consistent updates," in *Proc. ACM HotSDN*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2491185.2491191

[73] R. McGeer, "A safe, efficient update protocol for openflow networks," in *Proc. ACM HotSDN*, New York, NY, USA, 2012. [Online]. Available: http://doi.acm.org/10.1145/2342441.2342454

[74] ——, "A correct, zero-overhead protocol for network updates," in *Proc. ACM HotSDN*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2491185.2491217

[75] J. Hua, X. Ge, and S. Zhong, "FOUM: A flow-ordered consistent update mechanism for software-defined networking in adversarial settings," in *Proc. IEEE INFOCOM*, 2016. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2016.7524499

[76] J. McClurg, H. Hojjat, P. Cerný, and N. Foster, "Efficient synthesis of network updates," in *Proc. ACM SIGPLAN PLDI*, 2015. [Online]. Available: http://doi.acm.org/10.1145/2737924.2737980

[77] P. Cerný, N. Foster, N. Jagnik, and J. McClurg, "Optimal consistent network updates in polynomial time," in *Proc. DISC*, 2016. [Online]. Available: http://dx.doi.org/10.1007/978-3-662-53426-7_9

[78] S. Vissicchio, L. Vanbever, L. Cittadini, G. G. Xie, and O. Bonaventure, "Safe update of hybrid sdn networks," *IEEE/ACM Transactions on Networking*, vol. 25, no. 3, pp. 1649–1662, June 2017.

[79] S. Vissicchio and L. Cittadini, "FLIP the (flow) table: Fast lightweight policy-preserving SDN updates," in *Proc. IEEE INFOCOM*, 2016. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2016.7524419

[80] S. K. Fayazbakhsh, V. Sekar, M. Yu, and J. C. Mogul, "Flowtags: enforcing network-wide policies in the presence of dynamic middlebox actions," in *Proc. ACM HotSDN*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2491185.2491203

[81] Z. A. Qazi, C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu, "Simple-fying middlebox policy enforcement using SDN," in *Proc. ACM SIGCOMM*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2486001.2486022

[82] S. Vissicchio and L. Cittadini, "Safe, efficient, and robust SDN updates by combining rule replacements and additions," *IEEE/ACM Trans. Netw.*, vol. 25, no. 5, pp. 3102–3115, 2017. [Online]. Available: http://doi.ieeecomputersociety.org/10.1109/TNET.2017.2723461

[83] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "P4: programming protocol-independent packet processors," *Computer Communication Review*, vol. 44, no. 3, pp. 87–95, 2014. [Online]. Available: http://doi.acm.org/10.1145/2656877.2656890

[84] K.-T. Foerster, "On the consistent migration of unsplittable flows: Upper and lower complexity bounds," in *Proc. IEEE NCA*, 2017. [Online]. Available: https://doi.org/10.1109/NCA.2017.8171348

[85] W. Wang, W. He, J. Su, and Y. Chen, "Cupid: Congestion-free consistent data plane update in software defined networks," in *Proc. IEEE INFOCOM*, 2016. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2016.7524420

[86] X. Jin, Y. Li, D. Wei, S. Li, J. Gao, L. Xu, G. Li, W. Xu, and J. Rexford, "Optimizing bulk transfers with software-defined optical WAN," in *Proc. ACM SIGCOMM*, 2016. [Online]. Available: http://doi.acm.org/10.1145/2934872.2934904

[87] Y. Chen and J. Wu, "Max progressive network update," in *Proc. IEEE ICC*, 2017. [Online]. Available: https://doi.org/10.1109/ICC.2017.7997201

[88] R. Gandhi, O. Rottenstreich, and X. Jin, "Catalyst: Unlocking the power of choice to speed up network updates," in *Proc. CoNEXT*. ACM, 2017. [Online]. Available: http://doi.acm.org/10.1145/3143361.3143397

[89] S. Paris, A. Destounis, L. Maggi, G. S. Paschos, and J. Leguay, "Controlling flow reconfigurations in SDN," in *Proc. IEEE INFOCOM*, 2016. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2016.7524330

[90] S. Luo, H. Yu, L. Luo, and L. Li, "Arrange your network updates as you wish," in *Proc. IFIP Networking*, 2016. [Online]. Available: http://dx.doi.org/10.1109/IFIPNetworking.2016.7497214

[91] S. A. Amiri, S. Dudycz, S. Schmid, and S. Wiederrecht, "Congestion-free rerouting of flows on dags," *CoRR*, vol. abs/1611.09296, 2016. [Online]. Available: http://arxiv.org/abs/1611.09296

[92] T. Mizrahi and Y. Moses, "On the necessity of time-based updates in SDN," in *Proc. ONS*. USENIX Association, 2014. [Online]. Available: https://www.usenix.org/conference/ons2014/technical-sessions/presentation/mizrahi

[93] G. Baier, E. Köhler, and M. Skutella, "The k-splittable flow problem," *Algorithmica*, vol. 42, no. 3-4, pp. 231–248, 2005. [Online]. Available: http://dx.doi.org/10.1007/s00453-005-1167-9

[94] T. Hartman, A. Hassidim, H. Kaplan, D. Raz, and M. Segalov, "How to split a flow?" in *Proc. IEEE INFOCOM*, 2012. [Online]. Available: http://dx.doi.org/10.1109/INFCOM.2012.6195830

[95] J. He and J. Rexford, "Toward internet-wide multipath routing," *IEEE Network*, vol. 22, no. 2, pp. 16–21, 2008. [Online]. Available: http://dx.doi.org/10.1109/MNET.2008.4476066

[96] S. Kandula, D. Katabi, S. Sinha, and A. Berger, "Dynamic load balancing without packet reordering," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 2, pp. 51–62, 2007. [Online]. Available: http://doi.acm.org/10.1145/1232919.1232925

[97] S. Kandula, I. Menache, R. Schwartz, and S. R. Babbula, "Calendaring for wide area networks," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 515–526, Aug. 2014. [Online]. Available: http://doi.acm.org/10.1145/2740070.2626336

[98] J. Perry, A. Ousterhout, H. Balakrishnan, D. Shah, and H. Fugal, "Fastpass: A centralized "zero-queue" datacenter network," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 307–318, Aug. 2014. [Online]. Available: http://doi.acm.org/10.1145/2740070.2626309

[99] M. Hoefer, V. S. Mirrokni, H. Röglin, and S. Teng, "Competitive routing over time," *Theor. Comput. Sci.*, vol. 412, no. 39, pp. 5420–5432, 2011. [Online]. Available: http://dx.doi.org/10.1016/j.tcs.2011.05.055

[100] W. Zhou, D. K. Jin, J. Croft, M. Caesar, and P. B. Godfrey, "Enforcing customizable consistency properties in software-defined networks," in *Proc. NSDI*. USENIX Association, 2015. [Online]. Available: https://www.usenix.org/conference/nsdi15/technical-sessions/presentation/zhou

[101] H. Xu, Z. Yu, X. Li, C. Qian, L. Huang, and T. Jung, "Real-time update with joint optimization of route selection and update scheduling for sdns," in *Proc. ICNP*. IEEE, 2016. [Online]. Available: https://doi.org/10.1109/ICNP.2016.7784436

[102] T. Mizrahi and Y. Moses, "Reverseptp: A clock synchronization scheme for software-defined networks," *Int. Journal of Network Management*, vol. 26, no. 5, pp. 355–372, 2016. [Online]. Available: http://dx.doi.org/10.1002/nem.1942

[103] ——, "Using reverseptp to distribute time in software defined networks," in *Proc. ISPCS*, 2014.

[104] T. Mizrahi, E. Saat, and Y. Moses, "Timed consistent network updates in software-defined networks," *IEEE/ACM Trans. Netw.*, vol. 24, no. 6, pp. 3412–3425, 2016. [Online]. Available: https://doi.org/10.1109/TNET.2016.2529058

[105] J. Zheng, G. Chen, S. Schmid, H. Dai, and J. Wu, "Scheduling congestion-and loop-free network update in timed sdns," *IEEE JSAC Special Issue on Emerging Technologies in Software-Driven Communication*, 2017.

[106] T. Mizrahi and Y. Moses, "Software defined networks: It's about time," in *Proc. IEEE INFOCOM*, 2016. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2016.7524418

[107] T. Mizrahi, O. Rottenstreich, and Y. Moses, "Timeflip: Using timestamp-based TCAM ranges to accurately schedule network updates," *IEEE/ACM Trans. Netw.*, vol. 25, no. 2, pp. 849–863, 2017. [Online]. Available: https://doi.org/10.1109/TNET.2016.2608441

[108] T. Mizrahi and Y. Moses, "The case for data plane timestamping in SDN," in *Proc. IEEE INFOCOM SWFAN*, 2016. [Online]. Available: http://dx.doi.org/10.1109/INFCOMW.2016.7562197

[109] ——, "Time-based updates in software defined networks," in *Proc. ACM HotSDN*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2491185.2491214

[110] ——, "OneClock to rule them all: Using time in networked applications," in *Proc.IEEE/IFIP NOMS*. IEEE, 2016. [Online]. Available: http://dx.doi.org/10.1109/NOMS.2016.7502876

[111] ——, "Time Capability in NETCONF," in *RFC 7758*, 2016.

[112] X. Chen, Z. M. Mao, and J. E. van der Merwe, "PACMAN: a platform for automated and controlled network operations and configuration management," in *Proc. ACM CoNEXT*, 2009. [Online]. Available: http://doi.acm.org/10.1145/1658939.1658971

[113] T. Koponen, M. Casado, N. Gude, J. Stribling, L. Poutievski, M. Zhu, R. Ramanathan, Y. Iwata, H. Inoue, T. Hama, and S. Shenker, "Onix: A distributed control platform for large-scale production networks," in *Proc. USENIX OSDI*, 2010. [Online]. Available: http://www.usenix.org/events/osdi10/tech/full_papers/Koponen.pdf

[114] M. Kuzniar, P. Peresíni, and D. Kostic, "Providing reliable FIB update acknowledgments in SDN," in *Proc. ACM CoNEXT*. ACM, 2014. [Online]. Available: http://doi.acm.org/10.1145/2674005.2675006

[115] P. Peresíni, M. Kuzniar, and D. Kostic, "Rule-level data plane monitoring with monocle," in *Proc. ACM SIGCOMM*. ACM, 2015. [Online]. Available: http://doi.acm.org/10.1145/2785956.2790012

[116] A. R. Curtis, J. C. Mogul, J. Tourrilhes, P. Yalagandula, P. Sharma, and S. Banerjee, "Devoflow: scaling flow management for high-performance networks," in *Proc. ACM SIGCOMM*, 2011. [Online]. Available: http://doi.acm.org/10.1145/2018436.2018466

[117] K. He, J. Khalid, A. Gember-Jacobson, S. Das, C. Prakash, A. Akella, L. E. Li, and M. Thottan, "Measuring control plane latency in sdn-enabled switches," in *Proc. ACM SOSR*, 2015. [Online]. Available: http://doi.acm.org/10.1145/2774993.2775069

[118] D. Y. Huang, K. Yocum, and A. C. Snoeren, "High-fidelity switch models for software-defined network emulation," in *Proc. ACM HotSDN*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2491185.2491188

[119] C. Rotsos, N. Sarrar, S. Uhlig, R. Sherwood, and A. W. Moore, "OFLOPS: an open framework for openflow switch evaluation," in *Proc. PAM*, 2012. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-28537-0_9

[120] X. Jin, J. Gossels, J. Rexford, and D. Walker, "Covisor: A compositional hypervisor for software-defined networks," in *Proc. USENIX NSDI*, 2015. [Online]. Available: https://www.usenix.org/conference/nsdi15/technical-sessions/presentation/jin

[121] L. Schiff, S. Schmid, and P. Kuznetsov, "In-band synchronization for distributed SDN control planes," *Computer Communication Review*, vol. 46, no. 1, pp. 37–43, 2016. [Online]. Available: http://doi.acm.org/10.1145/2875951.2875957

[122] M. Canini, P. Kuznetsov, D. Levin, and S. Schmid, "A distributed and robust SDN control plane for transactional network updates," in *Proc. IEEE INFOCOM*, 2015. [Online]. Available: http://dx.doi.org/10.1109/INFOCOM.2015.7218382

[123] L. Vanbever, J. Reich, T. Benson, N. Foster, and J. Rexford, "Hotswap: correct and efficient controller upgrades for software-defined networks," in *Proc. ACM HotSDN*, 2013. [Online]. Available: http://doi.acm.org/10.1145/2491185.2491194

[124] K. Saur, J. M. Collard, N. Foster, A. Guha, L. Vanbever, and M. W. Hicks, "Safe and flexible controller upgrades for sdns," in *Proc. SOSR*. ACM, 2016. [Online]. Available: http://doi.acm.org/10.1145/2890955.2890966

[125] X. Zhou, R. Li, T. Chen, and H. Zhang, "Network slicing as a service: enabling enterprises' own software-defined cellular networks," *IEEE Communications Magazine*, vol. 54, no. 7, pp. 146–153, 2016. [Online]. Available: https://doi.org/10.1109/MCOM.2016.7509393

[126] I. da Silva, G. Mildh, A. Kaloxylos, P. Spapis, E. Buracchini, A. Trogolo, G. Zimmermann, and N. Bayer, "Impact of network slicing on 5g radio access networks," in *Proc. EuCNC*. IEEE, 2016. [Online]. Available: https://doi.org/10.1109/EuCNC.2016.7561023

[127] J. Ordonez-Lucena, P. Ameigeiras, D. Lopez, J. J. Ramos-Muñoz, J. Lorca, and J. Folgueira, "Network slicing for 5g with SDN/NFV: concepts, architectures, and challenges," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 80–87, 2017. [Online]. Available: https://doi.org/10.1109/MCOM.2017.1600935

[128] S. Ghorbani and B. Godfrey, "Towards correct network virtualization," in *Proc. ACM HotSDN*, 2014. [Online]. Available: http://doi.acm.org/10.1145/2620728.2620754

[129] J. McClurg, H. Hojjat, N. Foster, and P. Cerný, "Event-driven network programming," in *Proc. ACM SIGPLAN PLDI*. ACM, 2016. [Online]. Available: http://doi.acm.org/10.1145/2908080.2908097

[130] S. Schmid and J. Suomela, "Exploiting locality in distributed SDN control," in *Proc. ACM HotSDN*. ACM, 2013. [Online]. Available: http://doi.acm.org/10.1145/2491185.2491198

[131] K. Phemius, M. Bouet, and J. Leguay, "DISCO: distributed SDN controllers in a multi-domain environment," in *Proc. IEEE NOMS*, 2014. [Online]. Available: http://dx.doi.org/10.1109/NOMS.2014.6838273

[132] T. D. Nguyen, M. Chiesa, and M. Canini, "Decentralized consistent updates in SDN," in *Proc. SOSR*. ACM, 2017. [Online]. Available: http://doi.acm.org/10.1145/3050220.3050224