Renaissance: A Self-Stabilizing Distributed SDN Control Plane

Marco Canini[1]   Iosif Salem[2]   Liron Schiff[3]   Elad M. Schiller[2]   Stefan Schmid[4]

[1] Université catholique de Louvain   [2] Chalmers University of Technology   [3] GuardiCore Labs   [4] University of Vienna

Abstract—By introducing programmability, automated verification, and innovative debugging tools, Software-Defined Networks (SDNs) are poised to meet the increasingly stringent dependability requirements of today's communication networks. However, the design of fault-tolerant SDNs remains an open challenge. This paper considers the design of dependable SDNs through the lenses of *self-stabilization*—a very strong notion of fault-tolerance. In particular, we develop algorithms for an in-band and distributed control plane for SDNs, called *Renaissance*, which tolerate a wide range of (concurrent) controller, link, and communication failures. Our self-stabilizing algorithms ensure that after the occurrence of an arbitrary combination of failures, (i) every non-faulty SDN controller can eventually reach any switch in the network within a bounded communication delay (in the presence of a bounded number of concurrent failures) and (ii) every switch is managed by at least one non-faulty controller. We evaluate *Renaissance* through a rigorous worst-case analysis as well as a prototype implementation (based on OVS and Floodlight), and we report on our experiments using Mininet.

I. INTRODUCTION

Context and Motivation. Software-Defined Networks (SDNs) have emerged as a promising alternative to the error-prone and manually configured traditional communication networks. In particular, by outsourcing and consolidating the control over the data plane elements, SDNs support a programmatic verification and enable new debugging tools.

However, while the literature articulates well the benefits of the separation between control and data plane and the need for distributing the control plane (e.g., for performance and fault-tolerance), the question of how connectivity between these two planes is maintained (i.e., the communication channels from controllers to switches and between controllers) has not received much attention. Providing such connectivity, is critical for ensuring the availability and robustness of SDNs.

Guaranteeing that each switch is managed, at any time, by at least one controller is challenging especially if control is *in-band*, i.e., if control and data traffic is forwarded along the same links and devices and hence arrives at the same ports. In-band control is desirable as it avoids the need to build, operate, and ensure the reliability of a separate out-of-band management network. Moreover, in-band management can in principle improve the resiliency of a network, by leveraging a higher path diversity (beyond connectivity to the management port).

The goal of this paper is the design of a highly fault-tolerant distributed and in-band control plane for SDNs. In particular, we aim to develop a self-stabilizing software-defined network: An SDN which recovers from controller, switch, and link failures, as well as a wide range of communication failures (such as packet omissions, duplications, or reorderings). As such, our work is inspired by Radia Perlman's pioneering work [23]: Perlman's work envisioned a self-stabilizing Internet and enabled today's link state routing protocols to be robust, scalable, and easy to manage. Perlman also showed how to modify the ARPANET routing broadcast scheme, so that it becomes self-stabilizing [24], and provided a self-stabilizing spanning tree algorithm for interconnecting bridges [25]. Yet, while the Internet core is "conceptually self-stabilizing", Perlman's vision remains an open challenge, especially when it comes to recent developments in computer networks, such as SDNs, for which we propose self-stabilizing algorithms.

Fault Model. We consider (i) fail-stop failures of controllers, which failure detectors can observe, (ii) link failures, and (iii) communication failures, such as packet omission, duplication, and reordering. In particular, our fault model includes up to κ link failures, for some parameter $\kappa \in \mathbb{Z}^+$. In addition, to the failures captured in our model, we also aim to recover from *transient faults*, i.e., any temporary violation of assumptions according to which the system and network were designed to behave, e.g., the corruption of the packet forwarding rules or malicious changes to the availability of links, switches, and controllers. We assume that (an arbitrary combination of) these transient faults can corrupt the system state in unpredictable manners. In particular, when modeling the system, we assume that these violations bring the system to an arbitrary state (while keeping the program code intact). Starting from an arbitrary state, the correctness proof of self-stabilizing systems [11], [13] has to demonstrate the return to

correct behavior within a bounded period, which brings the system to a *legitimate state*.

The Problem. This paper answers the following question: How can all non-faulty controllers maintain bounded (in-band) communication delays to any switch as well as to any other controller? We interpret the requirements for provable (in-band) bounded communication delays to imply (i) the absence of out-of-band communications or any kind of external support, and yet (ii) the possibility of fail-stop failures of controllers and link failures, as well as (iii) the need for guaranteed bounded recovery time after the occurrence of an arbitrary combination of failures. The studied problem also considers the possibility of any transient violation of the assumptions according to which the system was designed to behave, which we call transient faults.

Our Contributions. We present an important module for dependable networked systems: a self-stabilizing software-defined network. In particular, we provide a (distributed) self-stabilizing algorithm for decentralized SDN control planes that, relying solely on in-band communications, recover (from a wide spectrum of controller, link, and communication failures as well as transient faults) by re-establishing connectivity in a robust manner. Concretely, we present a system, henceforth called *Renaissance*[1], which, to the best of our knowledge, is the first to provide:

(1) A robust efficient and decentralized control plane: We maintain short, $O(D)$-length control plane paths in the presence of controller and link (at most κ many) failures, as well as, communication failures, where $D \leq N$ is the (largest) network diameter (when considering any possible network topology changes over time) and N is the number of nodes in the network. More specifically, suppose that throughout the recovery period the network topology was $(\kappa+1)$-edge-connected and included at least one (non-failed) controller. We prove that starting from a legitimate state, i.e., after recovery, our self-stabilizing solution can: *(i) Deal with fail-stop failures of controllers:* These failures require the removal of stale information (related to unreachable controllers) from the switch configurations. Cleaning up stale information avoids inconsistencies and having to store large amounts of history data. *(ii) Deal with link failures:* Starting from a legitimate system state, the controllers maintain an $O(D)$-length path to all nodes (switches and other controllers), as long as at most κ links fail. That is, after the recovery period the communication delays are bounded.

(2) Recovery from transient faults: We show that our control plane can even recover after the occurrence of transient faults. That is, starting from an *arbitrary*

state, the system recovers within time $O(D^2N)$ to a legitimate state. Note that $O(D^2N)$ depicts the absolute worst case, since recovery can be much faster depending on the actual failures that occur in the network. In a legitimate state, the number of packet forwarding rules per switch is at most N_C times the optimal, where N_C is (an upper bound on) the number of controllers.

While we are not the first to consider the design of self-stabilizing systems which maintain redundant paths also beyond transient faults, the challenge and novelty of our approach comes from the specific restrictions imposed by SDN (and in particular the switches). In this setting not all nodes can compute and communicate, and in particular, SDN switches can merely forward packets according to the rules that are decided by other nodes, the controllers. This not only changes the model, but also requires different proof techniques, e.g., regarding the number of resets and illegitimate rule deletions.

In order to validate and evaluate our model and algorithms, we implemented a prototype of *Renaissance* in Floodlight using Open vSwitch (OVS), complementing our worst-case analysis. Our experiments in Mininet demonstrate the feasibility of our approach, indicating that in-band control can be bootstrapped and maintained efficiently and automatically in the presence of failures.

Organization. We give an overview of our system and the components it interfaces in Section II. The model appears in Section III. Our algorithm is presented (Section IV), analyzed (Section V), and validated (Section VI). We discuss related work (Section VII) and then conclude (Section VIII). Due to the page limit, the proof details appear in [10].

II. THE SYSTEM IN A NUTSHELL

We provide an overview of the network architecture. The network includes a set $P_C = \{p_1, \ldots, p_{N_C}\}$ of N_C *(remote) controllers*, and a set $P_S = \{p_{N_C+1}, \ldots, p_{N_C+N_S}\}$ of the N_S *(packet forwarding) switches*, where i is the unique identifier of node $p_i \in P = P_C \cup P_S$. Each switch $p_i \in P_S$ stores a set of rules that the controllers install in order to define which packets have to be forwarded to which ports. In the out-of-band control scenario, a controller communicates the forwarding rules via a dedicated management port to the *control module* of the switch. In contrast, in an in-band setting, the control traffic is interleaved with the data plane traffic, which is the traffic between *hosts* (as opposed to controller-to-controller and controller-to-switch traffic): switches can be connected to hosts through data ports and may have additional rules installed to correctly forward their traffic. We do not assume anything about the hosts' network service, except for that their traffic may traverse the network.
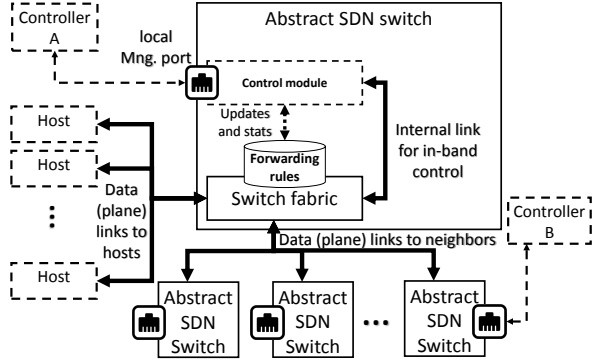
---

[1]The word renaissance means 'rebirth' (French) and it symbolizes the ability to recover after the occurrence of transient faults.

Fig. 1: Abstract SDN switch illustration.

In an in-band setting, control and data plane traffic arrive through the same ports at the switch, which implies a need for being able to *demultiplex* control and data plane traffic: switches need to know whether to forward (data) traffic out of another port or (control) traffic to the control module. Thus, control plane packets need to be logically distinguished from data plane traffic by some tag (or another deterministic discriminator).

Figure 1 illustrates the switch model considered in this paper. Our self-stabilizing control plane considers a proposal for *abstract switches* that do not require the extensive functionality that existing SDN switches provide. An abstract switch can be managed either via the management port or in-band. It stores forwarding (match-action) rules. These rules are used to forward data plane packets to ports leading to neighboring switches, or to forward control packets (e.g., instructing the control module to change existing rules) to the local control module. Rules can also drop all the matched packets. The match part of a rule can either be exact match or optionally include wildcards.

Maintaining the forwarding rules with in-band control is the key challenge addressed in this paper: for example, these rules must ensure (in a self-stabilizing manner) that control and data packets are demultiplexed correctly (e.g., using tagging). Moreover, it must be ensured that we do not end up with a set of misconfigured forwarding rules that drop *all* arriving (data plane and control plane) packets: in this case, a controller will never be able to manage the switch anymore.

In the following, we will assume a local topology discovery mechanism which reports to the controllers the availability of their direct neighbors. Also, we assume access to self-stabilizing protocols for the link layer (and the transport layer) [14] that provide reliable, bidirectional FIFO-communication channels over unreliable media that is prone to packet omission, reordering, and duplication.

A. Switches and Rules

Let p_j be a node and $N_c(j) \subseteq P$ (communication topology) be the set of directly attached neighboring nodes of p_j. At any given time, and for any given node $p_i \in P$, the set $N_o(j) \subseteq N_c(j)$ (operational topology) refers to p_j's directly connected nodes for which ports are currently available for packet forwarding.

Suppose that $p_i \in P_S$ is a switch that receives a packet with $p_{src} \in P_C$ and $p_{dest} \in P$ as the packet source, and destination respectively. We refer to a *rule* (for packet forwarding at the switch) by a tuple $\langle k, i, src, dest, prt, j, metadata \rangle$, where p_k is the controller that created this rule, $prt \in \{0, \ldots, n_{prt}\}$: $n_{prt} \geq \kappa + 1$ is a priority that p_k assigns to this rule, κ is a bound on the number of concurrently failing links, $p_j \in N_c(i)$ is a port on which the packet can be sent whenever $p_j \in N_o(i)$, and $metadata$ is an (optional) opaque data value. Our abstract switch considers only rules that are installed on the switches indefinitely, i.e., until a controller *explicitly* requests to delete them, rather than setting up rules with expiration *timeouts*.

Configuration Queries (via a Direct Neighbor): As long as the system rules and operational links support (bidirectional) packet forwarding between controller p_i and switch p_j, the abstract switch allows p_i to access p_j's configuration remotely, i.e., via the interfaces $manager(j)$ (query and update), $rules(j)$ (query and update) as well as $N_c(j)$ (query-only), where $manager(j) \subseteq P_C$ is p_j's set of assigned managers and $rules(j)$ is p_j's rule set. Also, a switch p_j, upon arrival of a query of a controller p_i, responds to p_i with the tuple $\langle j, N_c(j), manager(j), rules(j) \rangle$.

The abstract switch also allows controller p_i to query node p_j via p_j's direct neighbor, p_k as long as p_i knows p_k's local topology. In case p_j is a switch, p_i can also modify p_j's configuration (via p_j's abstract switch) to include a flow to p_i (via p_k) and then to add itself as a manager of p_j. We refer to this as the *query (and modify)-by-neighbor* functionality.

The Switch Memory Management: The number of rules and managers that each switch can store is bounded by $maxRules$ and $maxManagers$, respectively. The abstract switch has a way to deal with clogged memory by storing the rules and managers in a FIFO manner (say, using local counters that serve as timestamps in the meta-information ($metadata$) part of each rule). Whenever a controller accesses a switch, that switch refreshes these timestamps, i.e., all switch configuration items related to this controller. When the switch memory has more than $maxRules$ rules, the switch removes the rule that has the earliest timestamp so that a new rule can be added. This mechanism prioritizes newer rules (and manager information) that controllers install. Note that, as long as a switch has

sufficient memory to store the rules of all controllers in P_C, the above mechanism does not need to remove any rule of controller $p_i \in P_C$ after the first time that p_i has refreshed its rules on that switch. Similarly, we assume that whenever the number of managers that a switch stores exceeds $maxManagers$, the last to be stored (or access) manager is removed so that a new manager can be added.

Keeping the Switch Accessible: In order to render the network self-stabilizing, we assume that the abstract switch is self-stabilizing in a way that has the ability to ensure that initial misconfigurations are removed eventually. A switch may initially store rules which block any arriving packet, potentially making it unmanageable. The algorithm of the abstract switch can resolve this situation by interpreting any arriving packet as a control packet by default (if not matched otherwise). We leverage the limited size of the switch memory and disallow wildcarding on a specific field used for control traffic, rendering it impossible to disallow all control traffic given the limited number of rules. Even if some rules block control packets, the abstract switch algorithm can iterate over all possible header values to find ones that are not blocked. Note that this also prevents controller applications and user-level policies (e.g., a firewall) from blocking all control to a switch.

B. Building Blocks

Our architecture relies on a fault-tolerant mechanism for topology discovery. We use such a mechanism as an external building block. Moreover, we require a notion of resilient flows. We next discuss both these aspects.

1) Topology Discovery: The local topology information in $N_o(i)$ (operational topology) is liable to rapidly change without notice. We consider a system that uses an (ever running) failure detection mechanism, such as the self-stabilizing Θ-failure detector [6]: it discovers the switch neighborhood by identifying the failed/non-failed status of its attached links and neighbors. This mechanism reports the set of nodes $N_c(i) \subseteq P$ (communication topology) which are directly connecting node $p_i \in P$ and node p_j, i.e., $p_j \in N_c(i)$.

2) Fault-resilient Flows: We use κ-fault-resilient flows which are a reminiscent of the flows in [21]. The idea is that the network can forward the data packets along the shortest routes, and use alternative routes in the presence of link failures, based on conditional forwarding rules [7]; these failover rules provide a backup for every link. The flows in [21] are computed by constructing a directed acyclic graph for every destination in the network topology, such that all edges point to a neighbor that is closer to the destination. The authors propose a loop-free routing algorithm, that bases on shortest paths according to node identifiers.
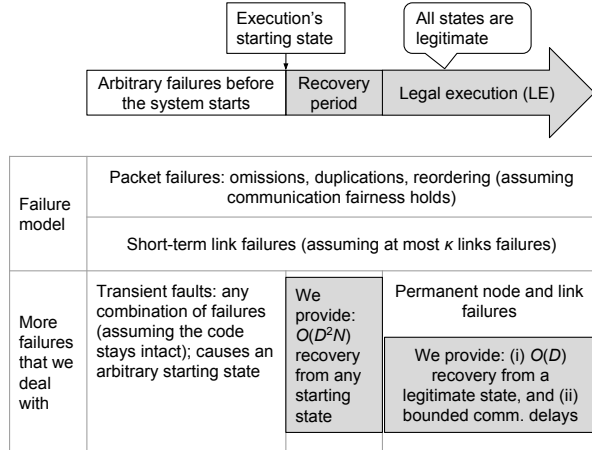

	Execution's starting state		All states are legitimate
	Arbitrary failures before the system starts	Recovery period	Legal execution (LE)

Failure model	Packet failures: omissions, duplications, reordering (assuming communication fairness holds)		
	Short-term link failures (assuming at most κ links failures)		
More failures that we deal with	Transient faults: any combination of failures (assuming the code stays intact); causes an arbitrary starting state	We provide: $O(D^2N)$ recovery from any starting state	Permanent node and link failures
			We provide: (i) $O(D)$ recovery from a legitimate state, and (ii) bounded comm. delays

Fig. 2: Failures (white boxes) and recovery guarantees (gray boxes) of the proposed self-stabilizing SDN solution.

III. THE MODEL

We model the control plane as a message passing system that has no notion of clocks (nor timeout mechanisms), as in the Paxos model [6], [20]. We borrow from [6, Section 6] a technique for local link monitoring (Section II-B1), which also addresses the need to detect fail-stop failures. This technique assumes that between the time in which a node completes one message round-trip with a (non-failing) directly connected neighbor, it cannot complete more than $\Theta-1$ round-trips with all the other directly connected (non-failing) neighbors, where Θ is a constant. Apart from this failure detector (which implies synchrony), we consider the control plane as an asynchronous system. We model the nodes as automata with input events that are either a packet reception or a periodic timer (at an unknown rate), as well as output events that send messages.

The *state* of a node consists of all its local variables as well as all the packets in the incoming communication channels. A *system state* is the collection of the states of all the nodes. A *step* starts with a node's local computations and ends with a single packet send or receive operation. We refer to an alternating sequence of states and steps as a (system) *execution*. We denote the operational and connected communication topology as $G_o = (P, E_o)$, and respectively, as $G_c = (P, E_c)$, where for $x \in \{o, c\}$, $E_x = \{(p_i, p_j) \in P \times P : p_j \in N_x(i)\}$. We assume that, during the system run, there are no more than κ link failures. We model as a *transient fault* the events of a failure of more than κ links (or the addition of new links) as well as switch fail-stop failures (or the addition of switches to the network). Moreover, the fail-stop failure of node p_j is a transient fault that results in the removal of (p_i, p_j) from the network

and p_j from $N_c(i)$. A transient fault can also corrupt the state of the nodes or the communication channels. We model the rare occurrence of transient faults by considering system executions in which transient faults never happen. Our model permits any finite number of transient faults before the system start. Namely, during the system run, G_c does not change and it is $(\kappa + 1)$-edge connected (cf. Figure 2).

Suppose that a κ-fault-resilient flow from p_i to p_j is installed in the network. The term *primary path* refers to the path along which the network forwards packets from p_i to p_j *in the absence of failures*. We assume that $myRules()$ maintains rules for κ-fault-resilient flows; their primary paths are also the shortest paths in G_c.

We define the system's task by a set of system runs called *legal executions* (LE) in which the task's requirements hold. I.e., each controller p_i constructs a κ-fault-resilient flow to every node $p_j \in P$ (either a switch or controller). We say that a system state c is *legitimate*, when every execution R (run) that starts from c is in LE. A system is *self-stabilizing* [13] with relation to task LE, when every (unbounded) system execution (with an arbitrary starting system state) reaches a legitimate state with relation to LE (Figure 2).

IV. A Self-Stabilizing SDN Control Plane

We present a self-stabilizing SDN control plane, called *Renaissance*, that enables each controller to discover the network, remove any stale information in the configuration of the discovered unmanaged switches (e.g., rules of failed controllers), and construct a κ-fault-resilient flow to any other node that it discovers.

Algorithm 1 creates an iterative process of topology discovery that, first, lets each controller identify the set of nodes that it is directly connected to; from there, it finds the nodes that are directly connected to them; and so on. This network discovery process is combined with another process for bootstrapping communication between any controller and any node in the network, i.e., connecting each controller to its direct neighbors, and then to their direct neighbors, and so on, until it is connected to the entire reachable network.

A. Variables, Building Blocks, Interfaces

Before presenting our algorithm, we introduce some notation, interfaces, and building blocks.

Local Variables. Each controller's state includes $responses$ (line 2), which is the set of the most recent query replies, and the tags (line 3), which are p_i's current ($currTag$) and previous ($prevTag$) synchronization round tags. Node p_j's response $m(j) : p_j \in P$ has the form $\langle j, N_c(j), manager(j), rules(j) \rangle$. The code denotes by $N_c(j)$ the neighborhood of p_j,

by $manager(j) \subseteq P_C$ the controllers of p_j, and by $rules(j) \subseteq \{\langle k, j, src, dest, prt, z, tag \rangle : (p_k, p_j, p_z, p_{dest} \in P) \wedge (p_{src} \in P_C) \wedge prt \in \{0, \ldots, n_{prt}\} \wedge tag \in tagDomain\}$ the rule set of p_j (cf. Section II-A). We assume that the size of $responses$ is bounded by $maxResponses \geq 2(N_C + N_S)$.

The Round Synchronization Mechanism. An SDN controller accesses the abstract switch in synchronized rounds. Each round has a unique tag that distinguishes the given round from its predecessors. We assume access to a self-stabilizing algorithm that generates unique *tags* of bounded size from a finite domain of tags, $tagDomain$. The algorithm provides a function called $nextTag()$ that, during a legal execution, returns a unique tag. That is, immediately before calling $nextTag()$ there is no tag anywhere in the system that has the returned value from that call. The function $nextTag()$ depends only on local variables, thus we require that the round synchronization mechanism is self-stabilizing to guarantee that $nextTag()$ returns a unique and maximum tag during a legal execution. Given two tags, t_1 and t_2, we require that $t_1 = t_2$ holds if, and only if, they have identical values. Note that there are self-stabilizing algorithms [4] that during fair executions provide unique tags. They recover within a bounded number of synchronization rounds (whenever the execution is legal with respect to the self-stabilizing end-to-end protocol and in which all failing links and nodes have been detected).

We use these tags for synchronizing the rounds in which the controllers perform configuration updates and queries. Namely, in the beginning of a round, controller $p_i \in P_C$ generates a new tag and stores that tag in the variable $currTag \leftarrow nextTag()$. Controller p_i then attempts to install at every reachable switch $p_j \in P_S$ a special meta-rule $\langle i, j, \bot, \bot, n_{prt}, \bot, t_{metaRule} \rangle$, which includes, in addition to p_i's identity, the tag $t_{metaRule} = currTag$ and has the lowest priority (before making any configuration update on that switch). It then sends a query to all (possibly) reachable nodes in the network and combines that query with the tag $t_{query} = currTag$. The response to that query from other controllers $p_j \in P_C$ includes the query tag, t_{query}. The response to the query from the switch $p_k \in P_S$ includes the tag $t_{metaRule}$ of the most recently installed meta-rule that p_k has in its configuration. Once p_i has received a response with $currTag$ from all reachable nodes, it ends that round. In Section IV-B we explain how Algorithm 1 treats packets with outdated tags.

Interfaces. Controller p_i can send requests or *queries* to any other node p_j. The controllers send command batches, which are sequences of commands. The special metadata command $\langle `newRound', t_{metaRule} \rangle$ is always the first command. We use it for starting a new round (where $t_{metaRule} = t$ is the round's tag). This starting

5

Algorithm 1: Self-stabilizing SDN, code for controller p_i.

---

1 **Symbols and operators:** '•' stands for 'any sequence of values', () is the empty sequence, ∘ (binary) is the sequence concatenation operator and ◯ (unary) concatenates a set's items in an arbitrary order.

2 **Local state:** $responses \subseteq \{m(j) : p_j \in P\}$ has the most recently received query replies $m(j)$, $p_j \in P$, where $m(j) := \langle j, N_c(j),$ $manager(j), rules(j)\rangle$, $N_c(j)$ is p_j's neighborhood, $manager(j) \subseteq P_C$ has p_j's controllers, and $rules(j) \subseteq \{\langle k, j, src, dest,$ $prt, z, tag\rangle : (p_k, p_j, p_z, p_{dest} \in P) \wedge p_{src} \in P_C \wedge prt \in \{0, \ldots, n_{prt}\} \wedge tag \in tagDomain\}$ is p_j's rule set;

3 $currTag$ and $prevTag$ are p_i's current, and respectively, previous synchronization round tags;

4 **Interfaces (sections II-B2 and IV-A):** $myRules(G, j, tag)$: creates p_i's rules at switch p_j according to G with tag tag;

5 **Macros:** $res(x) = \{\langle \bullet, rules(j)\rangle \in responses : \forall_{r \in rules(j)} r = \langle \bullet, x\rangle\} \cup \{\langle i, N_c(i), \emptyset, \emptyset\rangle\}$;

6 $G(S) := (\{p_k : \exists_{\langle j, N_c(j), \bullet\rangle \in S} : (k = j \vee p_k \in N_c(j))\}, \{(j, k) : \exists_{\langle j, N_c(j), \bullet\rangle \in S} : (p_k \in N_c(j))\})$;

7 $fusion := res(currTag) \cup \{\langle k, \bullet, prevTag\rangle \in res(prevTag) : \langle k, \bullet, currTag\rangle \notin res(currTag)\}$;

8 $p_j \rightarrow_G p_k :=$ true if there is a path from p_j to p_k in G;

9 **do forever begin**

  /* Remove replies from unreachable senders or not from round $prevTag$ or $currTag$. */

10 $\quad responses \leftarrow \{\langle k, \bullet, rules\rangle \in responses : k \neq i \wedge (\exists_{x \in \{currTag, prevTag\}} \langle k, \bullet, rules\rangle \in res(x) \wedge p_i \rightarrow_{G(res(x))}$
$\quad p_k \wedge \langle i, \bullet, x\rangle \in rules)\} \cup \{\langle i, N_c(i), \emptyset, \emptyset\rangle\}$;

11 $\quad$ **let** $(newRound, msg) := (false, \emptyset)$; /* $newRound$ and msg get their default values */

  /* a new round with a new tag; remove responses with tag $currTag$ */

12 $\quad$ **if** $(\forall p_\ell : p_i \rightarrow_{G(res(currTag))} p_\ell \implies \langle \ell, \bullet\rangle \in res(currTag)$ **then**

13 $\quad\quad (newRound, prevTag) \leftarrow (true, currTag)$; $currTag \leftarrow nextTag()$;

14 $\quad\quad responses \leftarrow responses \setminus \{\langle j, \bullet\rangle \in res(currTag) : p_j \in P\}$;

  /* The reference tag, $referTag$, is $currTag$ when a topology change is discovered */

15 $\quad$ **if** $G(fusion) = G(res(prevTag))$ **then let** $referTag := prevTag$ **else let** $referTag := currTag$;

16 $\quad$ **foreach** $p_j \in P_S : \langle j, Ngb, Mng, Rul\rangle \in res(referTag)$ **do** /* **manage switch p_j's rules** */

  /* p_i is a manager; remove unreachable managers on new rounds and nodes with no rules */

17 $\quad\quad$ **let** $M := \{p_k \in Mng : (\exists_{r \in Rul} r = \langle k, \bullet\rangle) \wedge (\neg newRound \vee p_i \rightarrow_{G(res(prevTag))} p_k)\} \cup \{p_i\}$;

18 $\quad\quad msg \leftarrow msg \cup \{(p_j, \langle 'delMngr', k\rangle) : p_k \in (Mng \setminus M)\} \cup \{(p_j, \langle 'addMngr', i\rangle)\}$;

  /* Remove any p_j's rule that is associated with an unreachable node, p_k */

19 $\quad\quad msg \leftarrow msg \cup \{(p_j, \langle 'delAllRules', k\rangle) : (\exists_{r \in Rul} r = \langle k, \bullet\rangle) \wedge p_k \notin M\}$;

  /* p_i refreshes all of its rules at switch p_j according to $referTag$ */

20 $\quad\quad msg \leftarrow msg \cup \{(p_j, \langle 'updateRules', myRules(G(res(referTag)), j, currTag)\rangle)\}$;

21 $\quad$ **foreach** $p_j : p_i \rightarrow_{G(fusion)} p_j$ **do send** $(\langle 'newRound', currTag\rangle) \circ [\bigcirc_{m:(p_j, m) \in msg}(m)] \circ (\langle 'query', currTag\rangle)$ **to** p_j;

22 **upon query reply** $m := \langle j, \bullet, rls\rangle$ **from** p_j **begin**

23 $\quad$ **if** $|responses \cup \{m\}| > maxResponses$ **then** $responses \leftarrow \{\langle i, N_c(i), \emptyset, \emptyset\rangle\}$; /* **C-reset** */

24 $\quad$ **if** $(\exists_{r \in rls} r = \langle \bullet, currTag\rangle)$ **then** $responses \leftarrow (responses \setminus \{\langle j, \bullet\rangle\}) \cup \{m\}$;

25 **upon arrival of** $(\bullet \circ (\langle 'query', tag\rangle))$ **from** p_j **do send** $\langle i, N_c(i), \perp, \{\langle j, i, \perp, \perp, \perp, \perp, tag\rangle\}\rangle$ **to** p_j

---

command could be followed by a number of commands, such as $\langle 'delMngr', k\rangle$ for the removal of controller p_k from the management of switch p_j, $\langle 'addMngr', k\rangle$ for the addition of controller p_k from the management of switch p_j, and $\langle 'delAllRules', k\rangle$ for the deletion of all of p_k's rules from p_j's configuration, where $p_k \in P_C \setminus \{p_i\}$. The rules' update, done via $\langle 'updateRules', newRules\rangle$, replaces all of p_i's rules at switch p_j (except for the special meta-rule).

These commands are to be followed by the round's query $\langle 'query', t_{query}\rangle$, where $t_{query} = t$ is the query's tag. The switch p_j replies to a query by sending $m = \langle j, N_c(j), manager(j), rules(j)\rangle$ to p_i, such that the rule set includes also the special meta-rule $\langle i, \bullet, t\rangle \in rules(j)$. Whenever $p_j \in P_C$ is another controller, the response to a query is simply $\langle i, N_c(i), \perp, \{\langle j, i, \perp, \perp, \perp, \perp, t_{query}\rangle\}\rangle$ (line 25). Note that controller p_j ignores all other types of commands. We use the interface function $myRules(G, j, tag)$ (Section II-B2) for creating the packet forwarding rules that controller p_i installs at switch p_j when p_i's current view

on the network topology is G in round tag (line 4).

B. Description of Algorithm 1

Each controller associates independently each iteration with a unique tag that synchronizes a round in which the controller performs configuration updates and queries. Controller p_i maintains the variables $currTag$ and $prevTag$ (line 3) of the round synchronization procedure, which starts when p_i queries all reachable nodes and ends when it receives replies from all of these nodes (cf. lines 6 and 8, 12–14, as well as, Section III). When a query response arrives at p_i, before the update of the response set (line 24), p_i checks that there is sufficient storage space for the arriving response (line 23). If space is lacking, p_i performs what we call a 'C-reset'. Note that p_i stores responses only for the current round, $currTag$. Controller p_i replies to other controllers' queries in line 25.

Controller $p_i \in P_C$ keeps a local state of query responses (cf. Section II-A) from other nodes (line 2). The

6

responses accumulate information about the network topology according to which the switch configurations are updated in each round. Algorithm 1's do-forever loop (lines 9–21) provides these functionalities.

Establishing communication between every controller and every other node. A controller $p_i \in P_C$ can communicate and manage a switch $p_j \in P_S$ only after p_i has installed rules at all the switches on a path between p_i and p_j. This, of course, depends on whether there are no permanent link failures on the path. In order to discover these link failures, we use local mechanisms for link state monitoring (cf. Section II-B1). The algorithm considers any permanent link failure as a transient fault and we assume that Algorithm 1 starts running only after the last occurrence of any transient fault (cf. Figure 2). Thus, as soon as there is a flow installed between p_i and p_j and there are no permanent failures on the primary path (Section III), p_i and p_j can exchange messages eventually.

The above iterative process of network topology discovery and the process of rule installation consider κ-fault-resilient flows (cf. Section II-B2 and $myRules()$ function in Section III). These flows are computed through the interface $myRules(G, j, tag)$ (line 4). Once the entire topology is discovered, Algorithm 1 guarantees the installation of a κ-fault-resilient flow between p_i and p_j. Thus, once the system is in a legitimate state, the availability of these flows implies that the system is resilient to the occurrence of at most κ temporary link failures (and recoveries) and p_i can communicate with any node in the network within a bounded time.

Topology discovery and dealing with unreachable nodes. Algorithm 1 lets the controllers connect to each other via κ-fault-resilient flows. Moreover, Algorithm 1 can detect situations in which controller $p_k \notin P_C$ is not reachable from controller p_i (line 10). The reason is that p_i is guaranteed to (i) discover the entire network eventually, and (ii) communicate with any node in the network. This means that p_i eventually gets a response from every node. Once that happens, the set of nodes that respond to p_i equals to the set of nodes that were discovered by p_i (line 12) and thus p_i can restart the process of discovering the network (lines13–14).

The start of a new (rediscovery) round, allows p_i to also remove information at the switches that is related to any unreachable controller $p_k \in P_C$ (assuming that it has succeeded in discovering the network and bootstrapped communication). During new rounds, p_i removes information related to p_k from any switch p_j (lines 17–19); whether this information is a rule or p_k's membership in p_j's management set. This stale information clean-up eventually brings the system to a legitimate state, as we will prove in Section V. The do-forever loop of Algorithm 1 completes by sending

rule (line 20) and manager updates to every switch that has a reply in $responses$, as well as querying every reachable node, with the current synchronization round's tag (lines 21–21). Note that each of these configuration updates are done via a single message that aggregates all commands for a given destination.

V. CORRECTNESS PROOF

We prove Algorithm 1's correctness by showing that (i) when the system starts in an arbitrary state, it reaches a legitimate state within a bounded period (Theorem 1), i.e., Figure 2's recovery period is bounded. Also, (ii) when starting from a legitimate state and letting a bounded number of failures to occur, the system returns to a legitimate state within a bounded period. Due to the page limit, we present a proof sketch for Theorem 1. The detailed proof appears in [10].

Refined model. We measure the recovery period of Algorithm 1 in terms of frames. The first *(asynchronous) frame* of run R is the shortest prefix R' of R in which every controller starts and ends at least one complete round-trip with every node of its discovered topology. The second frame in R is the first frame in R'', which is R's suffix that starts after R', and so on. We denote by Δ the recovery time of an end-to-end channel [14] and the self-stabilizing algorithm for unique label generation [4].

Definition 1 (Legitimate System State). *We say that $c \in R$ is Algorithm 1's legitimate state when: (1) the controllers have the correct information about the nodes in the network, (2) any controller is the manager of every switch and only these controllers can be the managers of any switch, (3) the rules installed in the switches encode κ-fault-resilient flows between all controllers and nodes, and (4) the end-to-end and round synchronization protocols are in a legitimate state.*

Theorem 1. *Within $((\Delta+2)D+1)[(\Delta D+1)N_S+N_C+1]$ frames in run R of Algorithm 1, the system reaches a state $c_{safe} \in R$ that is legitimate (Definition 1).*

Proof sketch: The proof follows by three claims. In Claim 1 we show that each switch needs to store a bounded number of rules and each controller needs to store a bounded number of responses and performs at most one C-reset (line 22). Then, we remark that the system cannot reach a legitimate state before it removes stale information from every switch configuration. Note that failing controllers cannot remove stale information that is associated with them and therefore non-failing controllers have to remove this information for them. Due to transient faults, it could be the case that a controller removes erroneously information about another non-failing controller. We refer to these 'mistakes' as illegitimate deletions of rules and note that they

Fig. 3: Bootstrap time for Telstra (T), EBONE (E) and Exodus (X) for 1 to 7 controllers.

occur when the (stale) information that a controller has about the network topology differs from the actual network topology, G_c. Due to stale information in the communication channels, any given controller might aggregate (possibly stale) information about the network more than once and thus instruct the switches again to delete illegitimately rules of other controllers. Thus, in Claim 2 we show that the number of steps in which a controller instructs the switches to perform illegitimate deletions (due to stale information) is bounded by $D \cdot (\Delta + 1) + 1$. Finally, in Claim 3 we show that Algorithm 1 recovers from transient faults, i.e., within $((\Delta + 2)D + 1)[(\Delta D + 1) \cdot N_S + N_C + 1]$ frames, R has a period of $(\Delta + 2)D + 1$ frames in which there are no C-resets and no illegitimate deletions and thus the system reaches c_{safe}. ∎

Our proof also shows that, when starting from a legitimate state and then letting a single link in the network to be added or removed from G_c, the system recovers within $O(D)$ frames. The arguments here consider the number of frames it takes for each controller to notice the change and to update all the switches. By similar arguments, we show that, within $O(D)$ frames, the system recovers after the addition or removal of at most $N_C - 1$ controllers in a legitimate system state.

VI. Evaluation

We evaluate our approach and study *Renaissance*'s performance. We implemented a prototype using Open vSwitch (OVS) and Floodlight. The prototype source code, once ready for release, will be made available [1]. We note that a comparison with the performance of existing in-band control planes would not be meaningful, since our approach is the first to consider self-stabilization. If the control plane is not self-stabilizing, then it cannot automatically recover from transient faults (e.g., arbitrarily corrupted states). In this case, recovery depends on external intervention.

Setup. The experiments are conducted using PCs with Ubuntu 16.04.1 OS, Intel Core i7-2600K CPU at 3.40GHz (8CPUs) with 16GB RAM. The link status detectors (for switches and controllers) are parametrized with frequency $\theta = \Delta(G) \cdot 3$, where $\Delta(G)$ is the maximum node degree. If not stated otherwise, the controllers issue requests and install flows once per second. Paths are computed according to Breadth First Search (BFS) and we use OpenFlow fast-failover groups for backup paths. The hosts for traffic and RTT evaluation are placed such that the distance between them is as large as the network diameter.

How efficiently can *Renaissance* bootstrap an SDN (resp. handle transient faults)? We first study how fast we can establish a stable network starting from empty switch configurations. Towards this end, we measure how long it takes until all controllers in the network reach a legitimate state in which each controller can communicate with any other node in the network (by installing packet-forwarding rules on the switches). For the smaller networks (B4 [17] and Clos [3]), we use 3 controllers, and for the Rocketfuel networks [30], [31] (Telstra, EBONE, and Exodus), we use up to 7 controllers. We note that figures 3, 5, 6, 7 are violin plots [16]. A violin plot illustrates the first and third quartile with a thick black line connecting them, the median with a white dot, the boundary values at the endpoints of a thin black line, and the kernel density estimation with the plot's vertical boundary (same on both sides). The horizontal boundary only closes the surface, which is colored for clarity.

We are indeed able to bootstrap in *any* of the configurations studied in our experiments. In terms of performance, as expected, the stabilization time is proportional to the network diameter and also depends on the number of controllers (Figure 3): more controllers result in slightly longer stabilization times as there will be more flows needed for each node in the network.

Note that the shown stabilization times only provide qualitative insights: they are, up to a certain point, proportional to the frequency at which controllers request configurations and install flows (Figure 4).

How efficiently does *Renaissance* recover in the presence of benign failures (link and node failures)? We consider different types of benign failures: controller fail-stop, permanent switch-failure, and permanent link-failure. The experiments start from a legitimate system state, to which we apply these failures.

In the *fail-stop* failure experiments (figures 5 and 6), we disconnect a single controller that is initially chosen at random, and measure the recovery time (from benign failures). The procedure is repeated for the same controller for each measurement. We also measure disconnecting 1-6 controllers simultaneously for the

Fig. 4: **Bootstrap time for Telstra, EBONE and Exodus using 7 controllers, as a function of query intervals. Very frequent queries cause congestion, hence the peak in 0.01.**


Fig. 5: **Recovery time after fail-stop failure of a controller.**

Fig. 6: **Recovery time after fail-stop failure for 1-6 controllers in Telstra (T), EBONE (E) and Exodus (X).**

Fig. 7: **Recovery time after multiple (2, 4 and 6) permanent link-failures at random for B4 (B), Clos (C), Telstra (T), EBONE (E) and Exodus (X).**

Rocketfuel (Telstra, EBONE and Exodus) networks, while running 7 controllers. The multiple controllers chosen for disconnection are also initially chosen at random, and the same ones are used when repeating the procedure for each measurement.

For the experiments for *permanent link-failures* (Figure 7), we disconnect either a single link that has maximal distance from all the controllers in the network (the procedure is repeated for the same link for each measurement), or, in case of multiple link failures, we choose failed links at random (making sure we do not disconnect the entire network). We find that the recovery time (after a fail-stop failure of a controller) is roughly linear in the number of nodes (Figure 5). The diameter also affects the time, but only to a smaller extent. For example, B4, which has a larger network diameter but is smaller compared to Clos, has a smaller average recovery time. The number of failed controllers (in Figure 6) does not affect the time much, as expected.
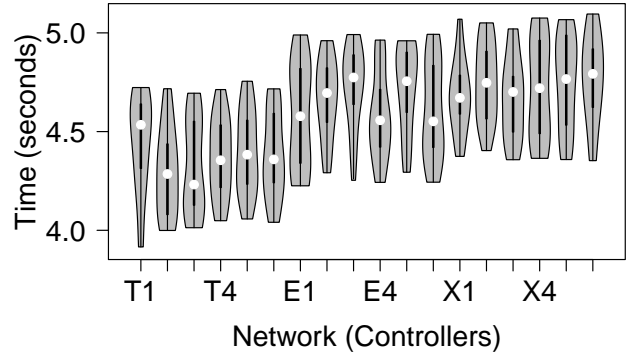
Performance and transient behavior. Besides connectivity, we are also interested in performance metrics such as throughput and message loss *during stabilization*, that is, recovery from transient faults that bring the system to an arbitrary starting state. In particular, while the recovery time (from benign failures) is quite fast, involving the control plane is time-consuming and can lead to packet reorderings and congestion. Accordingly, we employ local fast failover mechanisms.

In the following, we measure the TCP throughput between two hosts (placed at maximal distance from each other), in the presence of a link-failure located as close to the middle of the primary path. To generate traffic, we use Iperf. A specific link to fail is chosen such that it enables a backup path between the hosts.

The maximum link bandwidth is set to 1000 Mbits/s. We conduct throughput measurements during a period of 30 seconds. The link-failure occurs after 10 seconds, and we expect a throughput drop due to the traffic being rerouted to a backup path, which causes TCP's congestion control mechanism to reduce the

Fig. 8: Throughput for the different networks. A single link failure causes the drop after the 10^{th} second.

transmission rate when packet loss or reorderings occur. We note that our prototype utilizes packet tagging for incremental update [27]. This helps to avoid another drop in throughput as a result of the new paths being installed in order to repair flows while destroying the old backup paths.

We can see in Figure 8 that *one* throughput valley occurs indeed (to around 750 - 800 mbits/s). This is interesting, because, as we confirmed in additional experiments (not shown here), a naive approach which does not account for the multiple control planes and recovery from benign failures, results in repeated rerouting and hence repeated performance drops: The throughput increases after the traffic is rerouted to a backup path but drops once again after a few seconds, when the network recovers from benign failures. Using per-packet consistent paths and tagging forces the packets to only use the new primary path once all the necessary rules have been installed on the switches, which can reduce packet loss and re-transmissions of packets.

VII. RELATED WORK

To the best of our knowledge, our paper is the first to present a comprehensive model and rigorous approach for the design of in-band decentralized control planes providing self-stabilizing properties. As such, our approach complements much ongoing, often more applied, related research. In particular, our control plane can be used together with and support distributed systems, such as ONOS [5], ONIX [19], ElastiCon [12], Beehive [32], Kandoo [15], STN [9]. Our paper also provides missing links for the interesting work by Akella and Krishnamurthy [2], whose switch-to-controller and controller-to-controller communication mechanisms rely on strong primitives, such as consensus protocols, consistent snapshot and reliable flooding, which are not currently available in OpenFlow switches. We also note that our approach is not limited to a specific technology, but offers flexibilities and can be configured with additional robustness mechanisms, such as warm backups, local fast failover [26], or alternatives spanning trees [8], [22].

Furthermore, there exists interesting work on boot-

strapping connectivity in an OpenFlow network [18], [29] (that does not consider self-stabilization). In contrast to our paper, Sharma et al. [29] do not consider how to support multiple controllers nor how to establish the control network. Moreover, their approach relies on switch support for traditional STP and requires modifying DHCP on the switches. We do consider multiple controllers and establish an in-band control network in a self-stabilizing manner. Katiyar et al. [18] suggest bootstrapping a control plane of SDN networks, supporting multiple controller associations and also non-SDN switches. However, the authors do not consider fault-tolerance. We provide a very strong notion of fault-tolerance, which is self-stabilization.

We are not the first to consider self-stabilization in the presence of faults that are not just transient faults (see [13], Chapter 6). Thus far, self-stabilizing algorithms consider networks in which all nodes can compute and communicate. In the context of the studied problem, some nodes (the switches) can merely forward packets according to rules that are decided by other nodes (the controllers). To the best of our knowledge, we are the first to demonstrate a rigorous proof for the existence of self-stabilizing algorithms for an SDN control plane. This proof uses a number of techniques that are unique to the area, such as the one for assuring a bounded number of resets and illegitimate deletions. We reported on preliminary insights in a short paper on *Medieval* [28]. However, Medieval is not self-stabilizing because its design depends on the presence of non-corrupted configuration data.

VIII. CONCLUSION

We understand our paper as a first step, and believe it opens several interesting directions for future research. In particular, while we have deliberately focused on the more challenging in-band control scenario only, we anticipate that our approach can also be used in networks which combine both in-band and out-of-band control, e.g., depending on the network sub-regions. Moreover, while fundamental, our model is still simple and could be extended, e.g., to account for the dynamics of control and data plane traffic, e.g., by adjusting the failure detector model accordingly or to establish the backup routing paths for control traffic by considering the data traffic dynamics. Finally, while our prototype experiments demonstrate feasibility of our approach and show promising results, it remains to conduct a more rigorous practical evaluation. In order to facilitate future research, we will release the prototype source code together with this paper.

REFERENCES

[1] http://www.renaissance-sdn.net/.

[2] Aditya Akella and Arvind Krishnamurthy. A Highly Available Software Defined Fabric. In *HotNets*, 2014.

[3] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. A scalable, commodity data center network architecture. *ACM SIGCOMM Computer Communication Review*, 38(4):63–74, 2008.

[4] Noga Alon, Hagit Attiya, Shlomi Dolev, Swan Dubois, Maria Potop-Butucaru, and Sébastien Tixeuil. Practically stabilizing SWMR atomic memory in message-passing systems. *J. Comput. Syst. Sci.*, 81(4):692–701, 2015.

[5] Pankaj Berde, Matteo Gerola, Jonathan Hart, Yuta Higuchi, Masayoshi Kobayashi, Toshio Koide, Bob Lantz, Brian O'Connor, Pavlin Radoslavov, William Snow, and Guru M. Parulkar. ONOS: towards an open, distributed SDN OS. In Aditya Akella and Albert G. Greenberg, editors, *Proceedings of the third workshop on Hot topics in software defined networking, HotSDN '14, Chicago, Illinois, USA, August 22, 2014*, pages 1–6, 2014.

[6] Peva Blanchard, Shlomi Dolev, Joffroy Beauquier, and Sylvie Delaët. Practically self-stabilizing paxos replicated state-machine. In *Proc. 2nd International Conference on Networked Systems (NETYS)*, pages 99–121, 2014.

[7] Michael Borokhovich, Liron Schiff, and Stefan Schmid. Provable data plane connectivity with local fast failover: introducing openflow graph algorithms. In *Proc. 3rd Workshop on Hot Topics in Software Defined Networking (HotSDN)*, pages 121–126, 2014.

[8] Michael Borokhovich, Liron Schiff, and Stefan Schmid. Provable Data Plane Connectivity with Local Fast Failover: Introducing OpenFlow Graph Algorithms. In *HotSDN*, 2014.

[9] Marco Canini, Petr Kuznetsov, Dan Levin, and Stefan Schmid. A Distributed and Robust SDN Control Plane for Transactional Network Updates. In *INFOCOM*, 2015.

[10] Marco Canini, Iosif Salem, Liron Schiff, Elad Michael Schiller, and Stefan Schmid. Renaissance: Self-stabilizing distributed SDN control plane. *CoRR*, abs/1712.07697, 2017.

[11] Edsger W. Dijkstra. Self-stabilizing systems in spite of distributed control. *Commun. ACM*, 17(11):643–644, 1974.

[12] Advait Dixit, Fang Hao, Sarit Mukherjee, T.V. Lakshman, and Ramana Kompella. Towards an Elastic Distributed SDN Controller. In *HotSDN*, 2013.

[13] Shlomi Dolev. *Self-Stabilization*. MIT Press, 2000.

[14] Shlomi Dolev, Ariel Hanemann, Elad Michael Schiller, and Shantanu Sharma. Self-stabilizing end-to-end communication in (bounded capacity, omitting, duplicating and non-fifo) dynamic networks. In *Proc. International Symposium on Stabilization, Safety, and Security of Distributed Systems (SSS)*, pages 133–147, 2012.

[15] Soheil Hassas Yeganeh and Yashar Ganjali. Kandoo: A Framework for Efficient and Scalable Offloading of Control Applications. In *Proc. ACM HotSDN*, 2012.

[16] Jerry L Hintze and Ray D Nelson. Violin plots: a box plot-density trace synergism. *The American Statistician*, 52(2):181–184, 1998.

[17] Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, Jon Zolla, Urs Hölzle, Stephen Stuart, and Amin Vahdat. B4: Experience with a Globally-Deployed Software Defined WAN. In *SIGCOMM*, 2013.

[18] Rohit Katiyar, Prakash Pawar, Abhay Gupta, and Kotaro Kataoka. Auto-configuration of SDN switches in sdn/non-sdn hybrid network. In *Proceedings of the Asian Internet Engineering Conference, AINTEC 2015, Bangkok, Thailand, November 18-20, 2015*, pages 48–53. ACM, 2015.

[19] Teemu Koponen, Martin Casado, Natasha Gude, Jeremy Stribling, Leon Poutievski, Min Zhu, Rajiv Ramanathan, Yuichiro Iwata, Hiroaki Inoue, Takayuki Hama, and Scott Shenker. Onix: A Distributed Control Platform for Large-scale Production Networks. In *OSDI*, 2010.

[20] Leslie Lamport. The part-time parliament. *ACM Trans. Comput. Syst.*, 16(2):133–169, 1998.

[21] Junda Liu, Baohua Yang, Scott Shenker, and Michael Schapira. Data-driven network connectivity. In *Proc. ACM HotNets*, page 8, 2011.

[22] Merav Parter. Dual failure resilient BFS structure. In *PODC*, pages 481–490. ACM, 2015.

[23] Radia Perlman. *Interconnections (2Nd Ed.): Bridges, Routers, Switches, and Internetworking Protocols*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2000.

[24] Radia J. Perlman. Fault-tolerant broadcast of routing information. *Computer Networks*, 7:395–405, 1983.

[25] Radia J. Perlman. An algorithm for distributed computation of a spanningtree in an extended LAN. In William Lidinsky and Bart W. Stuck, editors, *SIGCOMM '85, Proceedings of the Ninth Symposium on Data Communications, British Columbia, Canada, September 10-12, 1985*, pages 44–53. ACM, 1985.

[26] Mark Reitblatt, Marco Canini, Arjun Guha, and Nate Foster. FatTire: Declarative Fault Tolerance for Software-defined Networks. In *HotSDN*, 2013.

[27] Mark Reitblatt, Nate Foster, Jennifer Rexford, and David Walker. Consistent updates for software-defined networks: change you can believe in! In Hari Balakrishnan, Dina Katabi, Aditya Akella, and Ion Stoica, editors, *Tenth ACM Workshop on Hot Topics in Networks (HotNets-X), HOTNETS '11, Cambridge, MA, USA - November 14 - 15, 2011*, page 7. ACM, 2011.

[28] Liron Schiff, Stefan Schmid, and Marco Canini. Ground control to major faults: Towards a fault tolerant and adaptive sdn control network. In *Proc. IEEE/IFIP DSN Workshop on Dependability Issues on SDN and NFV (DISN)*, 2016.

[29] Sachin Sharma, Dimitri Staessens, Didier Colle, Mario Pickavet, and Piet Demeester. In-Band Control, Queuing, and Failure Recovery Functionalities for OpenFlow. *IEEE Network*, 30(1), January 2016.

[30] N. Spring, R. Mahajan, and T. Andersonr. Quantifying the causes of path inflation. In *Proc. SIGCOMM*, 2003.

[31] Neil Spring, Ratul Mahajan, David Wetherall, and Thomas Anderson. Measuring ISP topologies with rocketfuel. *IEEE/ACM Trans. Netw.*, 12(1):2–16, 2004.

[32] Soheil Hassas Yeganeh and Yashar Ganjali. Beehive: Simple Distributed Programming in Software-Defined Networks. In *Proc. ACM SOSR*, 2016.