


Guest Editorial

Scalability Issues and Solutions for Software Defined Networks

Oliver Hohlfeld, James Kempf, *Senior Member, IEEE*, Martin Reisslein , *Fellow, IEEE*, Stefan Schmid, and Nadir Shah

I. INTRODUCTION

A. Software-Defined Networks: A Success Story

SOFTWARE Defined Networking (in short SDN, which is also an acronym for Software Defined Network), has emerged as a response to the limitations and complexities of traditional network architectures. At the heart of SDN lies the idea to consolidate the control over network devices into a logically centralized (software) controller separated from the data plane. The separation of the control plane and the data plane is realized via an open programming interface between the data plane switches and the SDN controller. The decoupling allows the control plane to evolve independently of the data plane, which enables faster innovation since software often exceeds hardware in innovation speed. Furthermore, logical centralization has the potential to simplify network operation and management by providing a single focal point where the consequences of management actions can be assessed, and possibly rejected if they would lead to some violation of operational constraints. OpenFlow, the standard SDN protocol today, is based on a simple match-action paradigm which results in great flexibilities, e.g., in terms of traffic engineering, definition of flows, as well as in-band network control functionalities.

While originally proposed in an academic context, SDN has now achieved far-reaching impact in industry, with many companies, such as Google, Facebook, Yahoo, and Microsoft, promoting and adopting SDN through open standards development. Today, SDN is deployed in a wide range of network types: in enterprise and campus networks (where it originally started), in datacenter networks, in wide-area networks (for example, Google B4), as well as in Internet Exchange Points.

O. Hohlfeld is with the Network Architectures Group, COMSYS, RWTH Aachen University, 52074 Aachen, Germany (e-mail: oliver@comsys.rwth-aachen.de).

J. Kempf is with Ericsson Research, Santa Clara, CA 95054 USA (e-mail: james.kempf@ericsson.com).

M. Reisslein is with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287-5706 USA (e-mail: reisslein@asu.edu).

S. Schmid is with the Faculty of Computer Science, University of Vienna, 1090 Vienna, Austria (e-mail: stefan_schmid@univie.ac.at).

N. Shah is with the Department of Computer Science, COMSATS Institute of Information Technology, Wah Cantt 47040, Pakistan (e-mail: nadirshah82@gmail.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSAC.2018.2872214

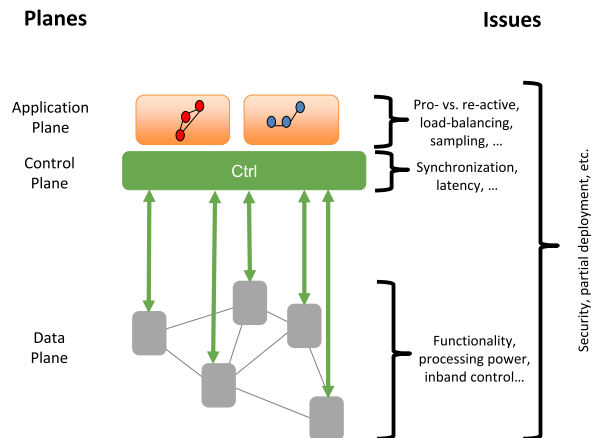


Fig. 1. Illustration of planes of an SDN and related scalability issues.

It is convenient to think of SDN in terms of its three main planes, which are illustrated in Fig. 1.

- **Data Plane:** The data plane consists of the network devices, such as the physical/virtual switches, routers, and access points. These data plane devices are managed by SDN controllers in the control plane.
- **Control Plane:** The SDN control plane consists of a set of software-based controllers that provide control functionality. Controllers have interfaces for communication with other controllers in the control plane, with network devices in the data plane, and with applications.
- **Application Plane:** End-user SDN applications interact with controllers to utilize an abstract view of the network for their internal decision making. Popular applications include traffic engineering and load balancing.

B. Scalability: The Next Frontier

Despite the large spectrum of SDN deployments, today's SDNs have in common that they are relatively small in scale. They are mostly limited to a small network or to a single administrative domain. There is a wide consensus that the next major challenge for SDN is to scale to massive numbers of routers and switches, as well as to address multi-domain networks.

We first have to note that while the centralized perspective offered by SDN is attractive, it was always meant to be *logically* centralized only, but physically *distributed*. Furthermore, there have already been a number of scalability related

improvements in OpenFlow specifications over the last few years, for example the introduction of a group table mechanism (OpenFlow version 1.1) which allows multiple flow table entries to point to (reference) the same group identifier: the controller only needs to update the referenced group table entry action instead of updating the action of all flow table entries. OpenFlow version 1.2 also allows to change the controller, enabling a switch to establish communication with multiple controllers in parallel.

Scalability involves generally a wide range of aspects and a system is only as scalable as its least scalable component, i.e., the component that first becomes a bottleneck as the network size scales up. The research addressing SDN scalability can be classified into the following main categories: data plane, control plane, applications, and cross-cutting issues.

1) *Data Plane Scalability*: Data plane scalability mainly depends on the processing power, memory/buffer capacity, and software implementation of the data plane devices [1]. One fundamental scalability question related to the data plane regards which functionality to include in the data plane in order to mitigate the controller load. A specialized data plane may be used to keep more traffic in the data plane, toward the goal of a more scalable and efficient control plane. Another scalability challenge arises in case of in-band control of data plane devices, where the control traffic shares paths (and competes for resources) with the data plane traffic, and where a reliable management of the data plane needs to be ensured.

2) *Control Plane Scalability*: The majority of the studies on SDN scalability addresses the control plane. The most prominent SDN control plane performance metrics are throughput (number of flow requests handled per second) and (flow setup) latency. A first set of scalability issues can result from the limited amount of requests that can be handled by a given controller per time. Distributed, flat, or hierarchical control plane designs can mitigate the controller bottleneck. However, such designs require inter-controller synchronization and data distribution among controller replicas. Further scalability issues arise from the separation of the control and data planes, as now, network devices are managed by a *remote* controlling mechanism. Indeed, the controller-switch communication delay introduces a delay, e.g., for flow setup. The situation is particularly problematic if the communication network between controllers and switches can be congested (e.g., in case of in-band management). In general, in order to reduce control latency in wide-area networks, data plane events should be handled close to where they occur, raising the question of where to place controllers.

3) *Application Scalability*: SDN applications themselves need to be scalable and support fine-grained and optimized resource utilization in wide-area networks. Load balancing was one of the first applications envisioned for SDN. One example technique to make load balancing applications scalable is to use wildcard-based rules to perform proactive load balancing. Such wildcard-based rules aggregate client requests, e.g., based on IP prefix ranges.

4) *Cross-Cutting Scalability Issues*: SDN will likely be deployed incrementally in large networks, which raises many other scalability related questions. For instance, which

switches should be upgraded first and how can an SDN operation be emulated with only a partial deployment of SDN devices [2].

Further cross-cutting scalability issues arise in the context of software-defined wireless networks. The wireless transmission medium introduces additional opportunities for control and optimization, i.e., presents additional tuning knobs that can be set with SDN control. However, fine-grained wireless transmission control can potentially give rise to novel bottlenecks in SDNs.

The larger the SDN, the more attractive it may become as a target of attacks. As an SDN is only as secure as its individual planes, strong yet scalable security mechanisms (e.g., for monitoring, detecting, and recovering from attacks) are required for the data plane, the control plane, and the application plane.

II. SUMMARY OF RESEARCH CONTRIBUTIONS IN THIS SPECIAL ISSUE

This Special Issue (SI) presents original state-of-the-art research studies on scalability issues and solutions for SDNs. The relevance of and interest in this topic area is reflected by the large number of submissions: we received a total of 70 submissions from authors distributed geographically over 21 countries and all 6 continents (19 from China, 15 from USA, 3 each from Canada, France, Germany, Korea, and Tunisia, 2 each from Japan, New Zealand, Poland, Spain, Taiwan, Turkey, and UK, and 1 each from Australia, Brazil, Hong Kong, Hungary, India, Italy, and Montenegro). All papers received at least three reviews and accepted papers went through at least one revision round. Out of the 70 papers, we eventually accepted 17 papers (acceptance ratio 24%).

The articles contained in this SI cover data plane, control plane, and application plane scalability issues and related solutions. The articles report on recent developments in architectural specifications, protocols, and application designs for achieving scalability in SDN. The papers in this SI employ a mixture of experimental, conceptual, and theoretical approaches to examine SDN scalability. The following subsections summarize the papers contained in this SI.

A. Data Plane

Patra *et al.* in *Towards a sweet spot of data plane programmability, portability and performance: On the scalability of multi-architecture P4 pipelines* examine the performance, portability, and scalability of the SDN data plane. Programming Protocol-Independent Packet Processors (P4) provide a high-level programming language for fine-grained programmability of the SDN data plane. Patra *et al.* argue that the P4 adoption is hindered by a lack of open source, protocol independent programmable data plane components. Patra *et al.* present a Multi-Architecture Compiler System for Abstract Data planes (MACSAD) concept. MACSAD builds on Application Programming Interfaces (APIs) from the OpenDataPlane (ODP) project to provide low-level hardware and software cross-platform programmability. Patra *et al.* identify three critical evaluation measures, namely performance,

portability, and scalability. They comprehensively evaluate the trade-offs between these three measures for a wide range of packet sizes, pipeline complexity levels, and numbers of cores across different platforms (e.g., x86 and ARM).

Oudin *et al.* in *OFLOPS-SUME and the art of switch characterization* introduce a novel hardware/software co-design for SDN switch characterization at 40 Gb/s and beyond. OFLOPS-SUME integrates the software based open framework for OpenFlow Switch Evaluation (OFLOPS) with the hardware based Open Source Network Tester (OSNT). The resulting OFLOPS-SUME can evaluate any hardware and software device with an OpenFlow control plane, including P4 programmable data plane devices. Oudin *et al.* conduct an extensive evaluation study of both software and hardware switches and identify the implications of different switch profiles on the overall network scalability and performance. Oudin *et al.* also release OFLOPS-SUME as an open source tool so as to facilitate the further refinement of the OFLOPS-SUME switch characterization tool and to support widespread reproducibility of SDN switch characterizations.

The next-generation 5G core will likely depend on a programmable data plane to support mobile edge computing, network slicing, and network function virtualization at scale. In order to support telcos and vendors to make the transition from the legacy mobile core to an SDN-based 5G-capable data plane, Lévai *et al.* in *The price for programmability in the software data plane: The vendor perspective* evaluate the performance and scalability of several programmable software switches for representative scenarios taken from real 5G deployments. In particular, Lévai *et al.* first introduce a taxonomy for data plane scalability together with a novel data-plane scalability benchmarking tool called TIPSy for Telco pIPeline benchmarking SYstem. Then, a set of ten standard telco pipelines is defined for benchmarking. Programmable software data planes based on eight different types of SDN programmable switches are then benchmarked. The evaluations indicate that most programmable switches result in significantly reduced performance, i.e., requiring significantly higher investments in hardware and energy (by up to roughly an order of magnitude) than conventional statically configured systems, i.e., there is a very high price to pay for the flexible configurability and management through programmable switches. Only one proprietary programmable switch has demonstrated the potential to match the performance of conventional systems.

Khalili *et al.* in *Flow setup latency in SDN networks* examine a path aggregation strategy to reduce the flow setup latency as SDN networks scale to large numbers of switches. Khalili *et al.* first measure flow setup latencies in conventional SDN networks and discover that the 99 percentiles of the latencies are on the order of half a second. In order to address these long latencies, Khalili *et al.* propose to pre-configure pipes that interconnect any pair of edge switches. The individual flows are processed at the ingress edge switches, from where they are forwarded through pre-configured fabric pipes to the egress edge switches. This approach reduces the number of switches that need to be configured for setting up a flow to typically two switches, decreasing the 99 percentile latencies to less than 10 milliseconds.

Xiao and Krunz in *Dynamic network slicing for scalable fog computing systems with energy harvesting* investigate SDN-based dynamic network slicing. Specifically, network slicing for a prescribed region of fog nodes is controlled by a regional orchestrator. The orchestrator coordinates with the region's fog nodes to dynamically distribute the resources of the network slices according to the local service demands and availability of harvested energy. Xiao and Krunz model the dynamic network slicing and the resource allocation problem as a stochastic overlapping coalition-formation game. The model reveals an increase in the overall computing capacity if fog nodes consider a belief function about the unknown states and private information of the other fog nodes. Xiao and Krunz optimize the dynamic network slicing through a belief-state partially observable Markov decision process. Numerical evaluations based on data from 400 base station locations in a real cellular network illustrate that cooperation of each fog node with its closest neighbor nearly doubles the processed workload.

Allybokus *et al.* in *Multi-path alpha-fair resource allocation at scale in distributed software defined networks* address the problem of fair bandwidth sharing for good network performance. In particular, the authors consider allocation problems for the important case that flows evolve over time. The authors propose a distributed algorithm to solve the multi-path fair resource allocation problem in a distributed SDN control architecture, overcoming the challenges of standard primal-dual decomposition methods. The authors then demonstrate the scalability of their approach on large instances with hundreds of nodes and thousands of requests and paths.

Bruyere *et al.* in *Rethinking IXPs' architecture in the age of SDN* propose to improve the scalability of the SDN control plane by handling the control traffic directly within the data plane. This approach is realized in the Umbrella switching fabric architecture and management approach. The Umbrella management approach improves the overall robustness by limiting the control plane dependency and is suitable for existing Internet eXchange Point (IXP) topologies. The scalability of Umbrella is demonstrated through the successful deployment at two commercial IXPs. Umbrella can be seen as a first step towards SDN architectures that are less dependent on the control plane and thus are better scalable. In such future architectures, the data plane components support the controller in its role of an intelligent supervisor, rather than as an active and critical decision point.

B. Control Plane

Yan *et al.* in *BigMaC: Reactive network-wide policy caching for SDN policy enforcement* address control and data plane scalability by proposing a reactive policy enforcement framework. BigMaC's main contribution is to provide a policy-consistency guarantee (i.e., packets should receive the same actions in the big switch abstraction and in the data plane switches with cached entries) across the network. The model presents layered views of the network with a big switch abstraction, a logical network, and a physical network. Yan *et al.* propose a network-wide "bucket"-based policy mapping and caching mechanism which guarantees policy-consistency, optimizes the flow table usage of physical network

devices, and reduces the churn incurred by policy updates. Trace-driven simulations indicate that BigMaC achieves better performance than traditional schemes, in particular, BigMaC saves table space and reduces the update complexity.

Görkemli *et al.* in *Dynamic control plane for SDN at scale* have been motivated by the scalability and reliability requirements of wide area networks and 5G core networks. In particular, the paper is concerned about the possibility of bottlenecks due to the controller CPU or the throughput of the in-band control channels. The authors propose a programmable distributed control plane architecture based on a dynamically managed in-band control network. Towards this end, the authors introduce a “control flow table” to manage in-band control flows, enabling for instance the offloading of congested controllers and congested in-band control channels.

Sakic and Kellerer in *Impact of adaptive consistency on distributed SDN applications: An empirical study* examine the scalability of the SDN control plane where state is replicated, e.g., for fault-tolerance. Sakic and Kellerer explore the effects of deployed consistency models on scalability and correctness, comparing strong and eventual consistency, and make the case for a novel adaptive consistency model. Sakic and Kellerer show how an adaptive consistency model offers scalability benefits in terms of the total request handling throughput and response time, in contrast to the strong consistency model. They also outline how the adaptive consistency model can provide correctness semantics, that are unachievable with the eventual consistency model. The paper further reports on an emulated testbed with a load balancer controller application.

Lyu *et al.* in *Multi-timescale decentralized online orchestration of software-defined networks* enhance SDN scalability with a hierarchical network of controllers, whereby controllers further down in the hierarchy can be switched off at different times depending on traffic conditions. Each switch in the network is associated with a controller at the first level, and if the controller is switched off, it forwards PACKET_IN messages from switches to the next level up. Lyu *et al.* develop an analytical formulation of a distributed online optimization problem which allows the control plane to determine the optimal set of active controllers. They test the algorithm in a MATLAB simulation platform with 320 switches and 14 controllers arranged in three levels: two root controllers, four intermediate controllers, and eight local controllers. They simulate three orchestration approaches: static orchestration, real-time dynamic orchestration, where the controller activation, request processing, and dispatching are carried out instantaneously at every time slot, and T -slot dynamic orchestration, where the request processing and dispatching are carried out per slot and the controller activation is carried out every T slots, with $T = 10$. The real-time orchestration achieves the lowest system cost, saving 73% over the static orchestration cost. The $T = 10$ slot orchestration approach achieves slightly less optimal performance than the real-time approach.

C. Scalable SDN Applications and Use Cases

Cheng and Jia in *Compressive traffic monitoring in hybrid SDN* investigate scalable load estimation, which can inform traffic engineering (an important application of SDN control).

In particular, Cheng and Jia propose a novel compressive traffic monitoring method for the accurate real-time collection of the load information of all links. The focus is on hybrid SDN networks and the main idea is to judiciously place a small number of SDN switches such that controllers only need to collect the load information of a small subset of important links. The loads of the other links are then estimated.

Moradi *et al.* in *Dragon: Scalable, flexible and efficient traffic engineering in software defined ISP networks* describe the Dragon SDN-based traffic engineering system for Internet Service Provider (ISP) networks that scales to large networks. To address the scalability challenge, Dragon consists of hierarchical and recursive traffic engineering algorithms that divide the flow optimization problems into subtasks and execute them in parallel. This approach finds an approximate solution much quicker than a full optimization problem. Dragon further enables ISPs to express diverse objectives for different parts of their network. Dragon also has optimizations to reduce the size of packets (avoiding excessive headers) and the Ternary Content-Addressable Memory (TCAM) size. Dragon has been evaluated for several network topologies.

Fu *et al.* in *Taming the wild: A scalable anycast-based CDN architecture (T-SAC)* utilize SDN to improve the scalability of anycast-based Content Distribution Networks (CDNs) by proposing to use OpenFlow and Network Function Virtualization (NFV) to complement current Domain Name System (DNS)-based redirection schemes. T-SAC achieves fine-grained control of redirections within the CDN through load-aware algorithms and setting a re-direction bit in the status/health check messages. T-SAC aims at fostering the collaboration of CDNs and ISPs by enabling more flexible routing decisions than a standard anycast architecture. The network itself measures the load imposed on CDN nodes and makes corresponding redirection decisions. At the same time, the CDN may influence the routing decision by signaling that certain CDN nodes are not capable of serving additional connections. The benefits of the T-SAC architecture are demonstrated by a real Amazon Web Services (AWS) based implementation.

Uddin *et al.* in *SDN-based multi-protocol edge switching for IoT service automation* propose an architecture called Muppet, for multiprotocol at the edge for IoT. The Muppet architecture consists of a collection of P4 Muppet switches at the edge that handle peer-to-peer non-IP traffic by tunneling over IP to other Muppet switches on the edge near the destination devices. This allows the IoT network to expand beyond a simple peer-to-peer network by using the IP backbone. The Muppet switches are programmed from a P4 SDN controller; the controller can install policy rules that allow traffic to be filtered depending on the content. The Muppet architecture is divided into three planes: a management plane which on-boards devices having specific link layer protocols, a data plane which supports packet inspection and processing using match-action rules, and a control plane that allows the P4 controller to program the switches. Uddin *et al.* demonstrate the Muppet architecture by supporting two different non-IP protocols, namely Bluetooth Low Energy (BLE) and Zigbee. The authors then describe a prototype implementation which

reduces the power consumption and increases the throughput compared to alternative systems (peer to peer with BLE, cloud with WiFi, and cloud with LTE). Muppet achieves these performance improvements by allowing the switches to filter device readings that do not match a specific policy (e.g., temperature readings that are above a threshold). They also tested the scalability using virtual BLE adaptors and found that the CPU usage of Muppet switches scales linearly with the number of ports, up to 1000 ports.

Zhang and Zhu in *Scalable virtualization and offloading based software-defined architecture for heterogeneous statistical QoS provisioning over 5G multimedia mobile wireless networks* propose an SDN architecture for offloading multimedia traffic to efficiently support multimedia traffic in 5G networks. The architecture consists of three virtual networks: a virtual network without offloading, a virtual network with WiFi offloading, and a virtual network with peer-to-peer offloading. The architecture achieves statistical QoS provisioning by different means on the different virtual networks. On the virtual networks without offloading, the authors analytically derive optimal wireless resource allocations to maximize the spectral efficiency. On virtual networks with WiFi offloading, an optimal power allocation scheme is developed to maximize the aggregate effective capacity. On virtual networks with peer-to-peer offloading, the performance improvements in terms of throughput, data content, and downloading delay are modelled. An SDN control plane makes the decision about what traffic to offload based on the characteristics of the traffic. Zhang and Zhu also estimate the scalability improvement of the overall integrated network. The advantages of the architecture are demonstrated through an analytical comparison with a standard base station cellular deployment.

Fawcett *et al.* in *TENNISON: A distributed SDN framework for scalable network security* describe a system for security monitoring and control based on a three layer architecture. At the bottom, in the Collection layer, existing tools collect data at L1 (sFlowRT), L2 (Open Network Operating System, ONOS), and L3 (Snort and Bro) to provide control and monitoring to higher layers. At the next layer, the Coordination layer, the Tennison controller provides coordination between different security applications. At the top layer, the Application layer, various security applications run and coordinate their activity southbound through the Tennison controller API. The Tennison controller design scales readily to multiple distributed instances so as to match the network topology. Fawcett *et al.* present emulation results from a Mininet simulation with 350 nodes connected to 19 partially connected switches, which is representative of a large-size business network, for four attack scenarios: Denial of Service (DoS), Distributed DOS (DDoS), scanning, and intrusion. They found that Tennison was able to detect an attack in less than 6 seconds, with the maximum amount of time needed for DDoS.

III. THE ROAD AHEAD: OPEN RESEARCH CHALLENGES

While the papers contained in this SI have achieved significant advances in making SDNs more scalable, there are still many important research challenges that need to be addressed to make large-scale SDNs a reality. We provide a perspective

of the main open research challenges towards scalable SDNs in this section. We first outline the main open research challenges in the data plane, the control plane, and the applications plane. We then outline several open research challenges that cut across these different planes and require a holistic approach to scalability.

A. Data Plane

1) *Flow Tables*: SDN switches are kept very simple by implementing only elementary data plane functionalities. The simple SDN switch design poses the challenge of efficiently using the limited SDN switch resources. For instance, how can the limited switch TCAM memory be used efficiently to store increasing numbers of active flows in the TCAM based switch flow table? How can the matching of an incoming packet with the patterns of large numbers of flow entries stored in the switch flow table be sped up? How can packets waiting for the controller commands be efficiently buffered at the switches? While a few approaches have been proposed to address these challenges, the scalability characteristics of flow table management require extensive future research.

2) *Adding Capabilities to Data Plane*: There is an emerging trend to offload more functionality to the data plane, e.g., through so-called smart Network Interface Controllers (NICs) that take over some of the control plane functions. The “smarter” data plane strives to offload the controller so as to make the controller more scalable. A smarter data plane can increase scalability also along other dimensions and positively affect applications. Aligned with this trend, and to give another example, there are studies evaluating under which conditions NICs and switches are suitable to work as CPU co-processors for neural networks, e.g., to accelerate Artificial Intelligence (AI).

3) *Consistent Data Plane Updates*: Some studies have addressed the problem of efficiently updating the flow tables at the switches to avoid inconsistencies. However, for large networks with frequent updates, the proposed protocols do not perform in real time and require extensive computation resources and network bandwidth. Thus, the consistent updating of the data plane in a scalable manner is still an open research challenge.

4) *Scalable Resilience Mechanism in Data Plane*: Links and nodes can fail in a network, affecting flows for critical applications [3]. To provide resilience against such failures, traditional networks employ distributed protocols. The logically centralized controller in SDNs can facilitate the efficient failure handling. The failure handling in SDNs poses unique challenges and several approaches have been proposed, e.g., [3]. In order to provide resilience against failures, the existing approaches install multiple paths for a given flow at the switches. The computational processing capacity and flow table memory of switches limit the scalability of the approach of installing multiple paths for a given flow. An open challenge is to examine approaches that compute the reliability levels of links and nodes, and then install the appropriate numbers of multiple paths based on the reliability levels. More specifically, if a path is highly stable, then we should not install multiple paths for a flow at the switches along the path.

Otherwise, we should install a number of multiple flow paths that is appropriate for compensating for the reliability levels of the paths.

B. Control Plane

1) *Multi-Controller Architectures*: A single controller can limit the scalability of an SDN. Multi-controller architectures have recently been proposed to increase the scalability of SDNs. However, there are several open challenges that still need to be addressed for multi-controller architectures. One main open challenge is the optimal arrangement of the controllers. Generally, multiple controllers can be arranged in different ways, such as isolated (in which a controller manages its own domain without communicating with other controllers), flat (in which multiple controllers are arranged in a flat structure and the controllers coordinate with each other to control and manage the network), and hierarchical (in which the controllers are arranged in layers and each controller has different roles, for example, local controller vs. global controller).

2) *Optimal Controller Placement*: Another open challenge is the optimal controller placement. The controllers should be placed to support efficient network operation. The controllers can be placed according to different operational objectives, e.g., each host should have the nearest controller within a prescribed maximum distance, or the overall flow setup latency should be minimized, or each controller should handle the same number of flows so as to balance the controller loads, or various combinations of these objectives.

3) *Reliable Distributed SDN Control*: The existing approaches for distributed SDN control limit some of the scalability aspects. For example, different links and switches have different reliability levels [3]. In order to address the reliability issue, the study [4] has advocated to place multiple controllers in the network using an in-band control plane. The in-band control plane conducts the communication between the switches and the controller over the communication links used for communication among the switches. In particular, the study [4] has advocated to connect the switches to the controllers so as to fairly balance the loads on the controllers, while the controllers back each other up in case of failures. Future research needs to consider the reliability levels of the paths from the switches to the controllers as well as the paths interconnecting the controllers in the switch to controller assignments. The reliability levels of a path will depend on the reliability of the bottleneck link of the path, i.e., the link with the lowest reliability level. More broadly, future research needs to ensure that the control plane does not become a performance bottleneck, nor a source for reliability problems as SDNs scale up in size and number of handled flows. Moreover, the control plane needs to efficiently and reliably support network virtualization, which requires hypervisors that intercept the data plane to controller traffic.

C. Scalable SDN Applications and Use Cases

1) *Scalable Approaches for Auto Network Debugging*: Recently, many approaches have been proposed for debugging

and verification of the network configurations and operations, such as RuleScope [5]. These approaches have major scalability problems: when the network scales, e.g., in terms of the number of switches and flow rules or in terms of the number of faulty rules, the execution time increases exponentially. Thus, as the network sizes scale up, it becomes increasingly harder to verify the network operations in real time. Therefore, the scalability of auto network debugging approaches is an important open research challenge.

2) *Network Updates*: Network maintenance and upgrades are vital tasks that require care and attention. Upgrades or maintenance of existing device hardware and software require temporarily the shifting of network traffic to other devices. Nevertheless, the overall network traffic service level should be maintained and congestion should be avoided. A range of mechanisms have been proposed for achieving these objectives in SDNs. However, these approaches suffer from scalability issues: when the network size increases in terms of the number of switches, the size of the traffic matrix, or the number flow rules, then the computation times of these approaches increase linearly. This can lead to network performance degradations, e.g., long computation times will cause congestion during the execution of the computations. Scalable network update mechanisms and corresponding resource allocation and dimensioning approaches are needed to ensure uninterrupted network service during upgrades and maintenance.

D. Cross-Cutting Challenges

1) *Hybrid SDNs*: Despite the benefits of SDN, many organizations hesitate to adopt SDN in practice due to a variety of reasons [2]. One main reason is that organizations often hesitate to invest large funds at once to completely replace traditional networks by SDNs. In order to reap the benefits of SDN while avoiding these concerns, researchers have proposed a new network architecture, referred to as a hybrid SDN. Hybrid SDNs enable the organizations to incrementally replace their traditional network devices with SDN devices. Due to their unique architecture, hybrid SDNs pose many scalability related research challenges.

These research challenges span all SDN planes. For instance, one data plane challenge is to optimize the locations of the SDN switches. An example of a control plane challenge is to design a scalable controller that can effectively control a large traditional network through a small set of SDN devices. Moreover, scalable protocols for communication among SDN devices and traditional network devices are an open research challenge. For example, Cheng and Jia in *Compressive traffic monitoring in hybrid SDN* in this SI achieve scalability by collecting the load levels of a small subset of important links and then use these load samples to infer the load levels of the remaining links. However, the problem of computing the subset of important links is NP-complete. Moreover, the problem of identifying the traditional routers to be replaced by SDN switches such that all important links are SDN-based is NP-hard.

A related open problem is the hybrid SDN control for large-scale hybrid SDNs. The recent Magneto study [6] has

proposed an SDN control mechanism to control traditional switches through SDN switches in a hybrid SDN. However, the performance evaluation of the Magneto control mechanism revealed that for a given fraction of SDN switches, a lower percentage of traditional switches is controlled by the SDN control mechanism in large-sized network topologies as compared to small-size and medium-size typologies [6, Fig. 7(c)]. Future research needs to develop effective hybrid SDN control mechanisms that scale well for large network topologies.

2) *Plane Interfaces*: The standardized SDN interfaces between the various planes are critical for the success of SDN. The so-called south-bound interface (Application Programming Interface, API), which is formally referred to as Data-Controller Plane Interface (D-CPI), interconnects the control plane with the data plane, while the so-called north-bound interface, formally referred to as Application-Controller Plane Interface (A-CPI), interconnects the control plane with the application plane. Moreover, the intermediate Controller Plane Interface (I-CPI) interconnects the SDN controllers of different network domains. Clearly, highly scalable SDNs require highly “scalable interfaces” that support efficient communication between the planes. More specifically, the interfaces need to be designed to efficiently handle growing flow numbers and network sizes. Growing network sizes, for instance, will likely require complex distributed control planes, which in turn complicate the south-bound interface. Accordingly, an interesting open question is whether today’s interfaces are indeed future proof, or whether improved programming interfaces need to be designed.

3) *Scalable Wireless SDN Networks*: Recently, SDN has been used in wireless networks for enhancing a wide range of wireless network functions, such as traffic offloading, making the wireless data plane more programmable, and managing node mobility more efficiently [7], [8]. Some studies have sought to make these wireless network enhancements scalable for large networks. For example, AeroFlux [7] has introduced two types of controllers, namely local controllers and global controllers. The local controller handles the local wireless network events while the globular controller handles the global events so as to improve the overall wireless network scalability. As another example, the SDN based smart GateWay (Sm-GW) [8] aggregates a large number of local wireless access nodes, e.g., numerous femto call base stations, to simplify the backhaul of the wireless traffic to and from the Internet at large.

Despite these and other recent advances towards scalable wireless SDNs, vast open research challenges remain. For instance, in the context of the emerging 5G wireless standard, thousands or even millions of Internet of Things (IoT) devices require connectivity at a fraction of the cost of today’s mobile devices. The question of which functionality to include in the data plane is particularly interesting for wireless SDNs due to the additional “tuning knobs” related to, e.g., power and rate control, as well as the inherent node mobility and changing wireless channel conditions. Additional challenges arise in multi-tenant wireless SDNs that require scalable isolation of the numerous wireless network slices and the abstraction of

the wireless network characteristics for the individual tenant controllers.

4) *Security*: SDNs need efficient and powerful security mechanisms to avoid security vulnerabilities across the data, control, and application planes. A number of security mechanisms have been proposed for SDNs. A security mechanism generally involves three phases, namely monitoring the network (this generates extra traffic both at the control and data planes), detecting the security breach (this takes some time for algorithm execution), and the recovery (once the attack is detected, this phase incurs both time delay and traffic overhead by taking the proper counter measures against the detected attack). The scalability of the existing security mechanisms is a big concern as the number of hosts, switches, controllers, flows, and attackers increases. The existing approaches typically attempt to achieve the scalability by improving the performance of an individual phase. For example, Fawcett *et al.* in *TENNISON: A distributed SDN framework for scalable network security* in this SI reduce the traffic overhead and the execution time of the monitoring phase. Similarly, the Athena approach [9] focuses on avoiding security vulnerabilities in the data plane. The main open research challenge is to develop holistic security approaches that improve the performance of all phases across all planes. These holistic approaches should reduce the execution time and increase the accuracy for increasing numbers of controllers and flows.

ACKNOWLEDGMENT

We are indebted to the authors that submitted papers for review for this SI. We are grateful to all the anonymous reviewers. This SI would not have been possible without the timely thorough reviews which have helped greatly to further refine and improve the papers in the revision round. We thank Muriel Médard and Raouf Boutaba, the former and present Editor-in-Chief of the IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS for their helpful guidance. A special thank you goes to Laurel Greenidge and Janine Bruttin, former and current Executive Editor of IEEE JSAC, who helped with countless logistical and procedural issues during the paper review and SI production process. We are also grateful to Prof. Moshe Zukerman, IEEE JSAC Senior Editor, for shepherding this SI.

REFERENCES

- [1] S. H. Yeganeh, A. Tootoonchian, and Y. Ganjali, “On scalability of software-defined networking,” *IEEE Commun. Mag.*, vol. 51, no. 2, pp. 136–141, Feb. 2013.
- [2] R. Amin, M. Reisslein, and N. Shah, “Hybrid SDN networks: A survey of existing approaches,” *IEEE Commun. Surveys Tuts.*, to be published.
- [3] P. Thorat, S. Jeon, and H. Choo, “Enhanced local detouring mechanisms for rapid and lightweight failure recovery in OpenFlow networks,” *Comput. Commun.*, vol. 108, pp. 78–93, Aug. 2017.
- [4] D. M. F. Mattos, O. C. M. B. Duarte, and G. Pujolle, “A resilient distributed controller for software defined networking,” in *Proc. IEEE Int. Conf. Commun. (ICC)*, May 2016, pp. 1–6.
- [5] X. Wen *et al.*, “RuleScope: Inspecting forwarding faults for software-defined networking,” *IEEE/ACM Trans. Netw.*, vol. 25, no. 4, pp. 2347–2360, Aug. 2017.
- [6] C. Jin, C. Lumezanu, Q. Xu, H. Mekky, Z.-L. Zhang, and G. Jiang, “Magneto: Unified fine-grained path control in legacy and OpenFlow hybrid networks,” in *Proc. ACM Symp. SDN Res.*, 2017, pp. 75–87.

- [7] N. A. Jagadeesan and B. Krishnamachari, "Software-defined networking paradigms in wireless networks: A survey," *ACM Comput. Surv.*, vol. 47, no. 2, pp. 27:1–27:11, Jan. 2015.
- [8] A. S. Thyagaturu, Y. Dashti, and M. Reisslein, "SDN-based smart gateways (Sm-GWs) for multi-operator small cell network management," *IEEE Trans. Netw. Service Manag.*, vol. 13, no. 4, pp. 740–753, Dec. 2016.
- [9] S. Lee, J. Kim, S. Shin, P. Porras, and V. Yegneswaran, "Athena: A framework for scalable anomaly detection in software-defined networks," in *Proc. IEEE/IFIP Int. Conf. Dependable Syst. Netw. (DSN)*, Jun. 2017, pp. 249–260.



Oliver Hohlfeld received the Ph.D. degree in computer science from TU Berlin in 2013. From 2008 to 2013, he was with TU Berlin/Deutsche Telekom Innovation Laboratories. He currently leads the Network Architectures Group, RWTH Aachen University. His research focuses on empirically understanding networked systems through large-scale measurements and subsequently optimizing their user experiences.



James Kempf (SM'18) received the Ph.D. degree in systems and industrial engineering, computer science minor, from the University of Arizona, Tucson, AZ, USA, in 1984. He promptly went to work in Silicon Valley. He spent three years at HP, 13 years at Sun Microsystems, primarily in research, and eight years at Docomo Innovations, Inc., Palo Alto, CA, USA, as a Research Fellow. He was with IETF for 10 years. He was the Chair of three working groups involved in developing standards for the mobile and wireless Internet and a member of the

Internet Architecture Board for two years. He holds 26 patents and authored many technical papers and three books, the latest of which, *Wireless Internet Security: Architecture and Protocols* (Cambridge University Press, 2008). Since 2008, he has been a Principal Researcher with Ericsson Research, Silicon Valley, where he was involved in software defined networking, network function virtualization, and cloud computing.



Martin Reisslein (S'96–M'98–SM'03–F'14) received the Ph.D. in systems engineering from the University of Pennsylvania, Philadelphia, PA, USA, in 1998. He is currently a Professor with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ, USA. He currently serves as an Associate Editor for the IEEE TRANSACTIONS ON MOBILE COMPUTING, the IEEE TRANSACTIONS ON EDUCATION, the IEEE ACCESS, and *Computer Networks*. He is an Associate Editor-in-Chief of the IEEE COMMUNICATIONS SURVEYS & TUTORIALS and the Co-Editor-in-Chief of *Optical Switching and Networking*. He chairs the Steering Committee of the IEEE TRANSACTIONS ON MULTIMEDIA.



Stefan Schmid received the Ph.D. (Dr.sc.) degree from ETH Zurich, Switzerland, in 2008. He is currently a Professor in computer science with the University of Vienna, Austria. He currently serves as an Editor of the *Distributed Computing Column of the Bulletin* of the EATCS and an Associate Editor for the IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT. He received the IEEE Communications Society ITC Early Career Award 2016.



Nadir Shah received the B.Sc. and M.Sc. degrees in computer science from Peshawar University, Peshawar, Pakistan, in 2002 and 2005, the M.S. degree in computer science from International Islamic University, Islamabad, Pakistan, in 2007, and the Ph.D. degree from the Sino-German Joint Software Institute, Beihang University, Beijing, China. He was a Lecturer with the Department of Computer Science, COMSATS Institute of Information Technology, Abbottabad, Pakistan, from 2007 to 2008. He is currently an Associate Professor with the COMSATS University Islamabad, Wah Campus. His current research interests include computer networks, distributed systems, and network security. He is serving in the Editorial Board of the IEEE SOFTWARE, AHWSN, and MJCS.