

Software-Defined Networking: A Comprehensive Survey

Diego Kreutz, *Member, IEEE*, Fernando M. V. Ramos, *Member, IEEE*, Paulo Verissimo, *Fellow, IEEE*,
Christian Esteve Rothenberg, *Member, IEEE*, Siamak Azodolmolky, *Senior Member, IEEE*,
and Steve Uhlig, *Member, IEEE*

Abstract—The Internet has led to the creation of a digital society, where (almost) everything is connected and is accessible from anywhere. However, despite their widespread adoption, traditional IP networks are complex and very hard to manage. It is both difficult to configure the network according to pre-defined policies, and to reconfigure it to respond to faults, load and changes. To make matters even more difficult, current networks are also vertically integrated: the control and data planes are bundled together. Software-Defined Networking (SDN) is an emerging paradigm that promises to change this state of affairs, by breaking vertical integration, separating the network’s control logic from the underlying routers and switches, promoting (logical) centralization of network control, and introducing the ability to program the network. The separation of concerns introduced between the definition of network policies, their implementation in switching hardware, and the forwarding of traffic, is key to the desired flexibility: by breaking the network control problem into tractable pieces, SDN makes it easier to create and introduce new abstractions in networking, simplifying network management and facilitating network evolution.

In this paper we present a comprehensive survey on SDN. We start by introducing the motivation for SDN, explain its main concepts and how it differs from traditional networking, its roots, and the standardization activities regarding this novel paradigm. Next, we present the key building blocks of an SDN infrastructure using a bottom-up, layered approach. We provide an in-depth analysis of the hardware infrastructure, southbound and northbound APIs, network virtualization layers, network operating systems (SDN controllers), network programming languages, and network applications. We also look at cross-layer problems such as debugging and troubleshooting. In an effort to anticipate the future evolution of this new paradigm, we discuss the main ongoing research efforts and challenges of SDN. In particular, we address the design of switches and control platforms – with a focus on aspects such as resiliency, scalability, performance, security and dependability – as well as new opportunities for carrier transport networks and cloud providers. Last but not least, we analyze the position of SDN as a key enabler of a software-defined environment.

Index Terms—Software-defined networking, OpenFlow, network virtualization, network operating systems, programmable networks, network hypervisor, programming languages, flow-based networking, scalability, dependability, carrier-grade networks, software-defined environments.

I. INTRODUCTION

The distributed control and transport network protocols running inside the routers and switches are the key technologies that allow information, in the form of digital packets, to travel around the world. Despite their widespread adoption, traditional IP networks are *complex and hard to manage* [1]. To express the desired high-level network policies, network operators need to configure each individual network device separately using low-level and often vendor-specific commands. In addition to the configuration complexity, network environments have to endure the dynamics of faults and adapt to load changes. Automatic reconfiguration and response mechanisms are virtually non-existent in current IP networks. Enforcing the required policies in such a dynamic environment is therefore highly challenging.

To make it even more complicated, current networks are also *vertically integrated*. The control plane (that decides how to handle network traffic) and the data plane (that forwards traffic according to the decisions made by the control plane) are bundled inside the networking devices, reducing flexibility and hindering innovation and evolution of the networking infrastructure. The transition from IPv4 to IPv6, started more than a decade ago and still largely incomplete, bears witness to this challenge, while in fact IPv6 represented *merely* a protocol update. Due to the inertia of current IP networks, a new routing protocol can take 5 to 10 years to be fully designed, evaluated and deployed. Likewise, a clean-slate approach to change the Internet architecture (e.g., replacing IP), is regarded as a daunting task – simply not feasible in practice [2], [3]. Ultimately, this situation has inflated the capital and operational expenses of running an IP network.

Software-Defined Networking (SDN) [4], [5] is an emerging networking paradigm that gives hope to change the limitations of current network infrastructures. First, it breaks the vertical integration by separating the network’s control logic (the control plane) from the underlying routers and switches that forward the traffic (the data plane). Second, with the separation of the control and data planes, network switches become simple forwarding devices and the control logic is implemented in a *logically centralized* controller (or

D. Kreutz and F. Ramos are with the Department of Informatics of Faculty of Sciences, University of Lisbon, Lisbon, 1749-016 Portugal e-mail: kreutz@lasige.di.fc.ul.pt, fvramos@fc.ul.pt.

P. Verissimo is with the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, 4 rue Alphonse Weicker, L-2721 Luxembourg. e-mail: paulo.verissimo@uni.lu.

C. Esteve Rothenberg is with the School of Electrical and Computer Engineering (FEEC), University of Campinas, Brazil. e-mail: chesteve@dca.fee.unicamp.br.

S. Azodolmolky is with Gesellschaft für Wissenschaftliche Datenverarbeitung mbH Göttingen (GWDG), Am Faßberg 11, 37077 Göttingen, Germany. e-mail: siamak.azodolmolky@gwdg.de.

S. Uhlig is with Queen Mary University of London. is with Queen Mary, University of London, Mile End Road, London E1 4NS, United Kingdom. e-mail steve@eecs.qmul.ac.uk.

Manuscript received May 31, 2014.




Fig. 1. Simplified view of an SDN architecture.

network operating system¹), simplifying policy enforcement and network (re)configuration and evolution [6]. A simplified view of this architecture is shown in Figure 1. It is important to emphasize that a logically centralized programmatic model does not postulate a physically centralized system [7]. In fact, the need to guarantee adequate levels of performance, scalability, and reliability would preclude such a solution. Instead, production-level SDN network designs resort to physically distributed control planes [7], [8].

The separation of the control plane and the data plane can be realized by means of a well-defined programming interface between the switches and the SDN controller. The controller exercises direct control over the state in the data-plane elements via this well-defined application programming interface (API), as depicted in Figure 1. The most notable example of such an API is OpenFlow [9], [10]. An OpenFlow switch has one or more tables of packet-handling rules (flow table). Each rule matches a subset of the traffic and performs certain actions (dropping, forwarding, modifying, etc.) on the traffic. Depending on the rules installed by a controller application, an OpenFlow switch can – instructed by the controller – behave like a router, switch, firewall, or perform other roles (e.g., load balancer, traffic shaper, and in general those of a middlebox).

An important consequence of the software-defined networking principles is the *separation of concerns* introduced between the *definition* of network policies, their *implementation* in switching hardware, and the *forwarding* of traffic. This separation is key to the desired flexibility, breaking the network control problem into tractable pieces, and making it easier to create and introduce new abstractions in networking, simplifying network management and facilitating network evolution and innovation.

Although SDN and OpenFlow started as academic experiments [9], they gained significant traction in the industry over the past few years. Most vendors of commercial switches now include support of the OpenFlow API in their equipment. The

SDN momentum was strong enough to make Google, Facebook, Yahoo, Microsoft, Verizon, and Deutsche Telekom fund Open Networking Foundation (ONF) [10] with the main goal of promotion and adoption of SDN through open standards development. As the initial concerns with SDN scalability were addressed [11] – in particular the myth that logical centralization implied a physically centralized controller, an issue we will return to later on – SDN ideas have matured and evolved from an academic exercise to a commercial success. Google, for example, has deployed a software-defined network to interconnect its data centers across the globe. This production network has been in deployment for 3 years, helping the company to improve operational efficiency and significantly reduce costs [8]. VMware’s network virtualization platform, NSX [12], is another example. NSX is a commercial solution that delivers a fully functional network in software, provisioned independent of the underlying networking devices, entirely based around SDN principles. As a final example, the world’s largest IT companies (from carriers and equipment manufacturers to cloud providers and financial-services companies) have recently joined SDN consortia such as the ONF and the OpenDaylight initiative [13], another indication of the importance of SDN from an industrial perspective.

A few recent papers have surveyed specific architectural aspects of SDN [14], [15], [16]. An overview of OpenFlow and a short literature review can be found in [14] and [15]. These OpenFlow-oriented surveys present a relatively simplified three-layer stack composed of high-level network services, controllers, and the controller/switch interface. In [16], the authors go a step further by proposing a taxonomy for SDN. However, similarly to the previous works, the survey is limited in terms of scope and it does not provide an in-depth treatment of fundamental aspects of SDN. In essence, existing surveys lack a thorough discussion of the essential building blocks of an SDN such as the network operating systems, programming languages, and interfaces. They also fall short on the analysis of cross-layer issues such as scalability, security, and dependability. A more complete overview of ongoing research efforts, challenges, and related standardization activities is also missing.

In this paper, we present, to the best of our knowledge, the most comprehensive literature survey on SDN to date. We organize this survey as depicted in Figure 2. We start, in the next two sections, by explaining the context, introducing the motivation for SDN and explaining the main concepts of this new paradigm and how it differs from traditional networking. Our aim in the early part of the survey is also to explain that SDN is not as novel as a technological advance. Indeed, its existence is rooted at the intersection of a series of “old” ideas, technology drivers, and current and future needs. The concepts underlying SDN – the separation of the control and data planes, the flow abstraction upon which forwarding decisions are made, the (logical) centralization of network control, and the ability to program the network – are not novel by themselves [17]. However, the integration of already tested concepts with recent trends in networking – namely the availability of merchant switch silicon and the huge interest in feasible forms of network virtualization – are

¹We will use these two terms interchangeably.




Fig. 2. Condensed overview of this survey on SDN.

leading to this paradigm shift in networking. As a result of the high industry interest and the potential to change the status quo of networking from multiple perspectives, a number of standardization efforts around SDN are ongoing, as we also discuss in Section III.

Section IV is the core of this survey, presenting an extensive and comprehensive analysis of the building blocks of an SDN infrastructure using a bottom-up, layered approach. The option for a layered approach is grounded on the fact that SDN allows thinking of networking along two fundamental concepts, which are common in other disciplines of computer science: a) separation of concerns (leveraging the concept of abstraction) and b) recursion. Our layered, bottom-up approach divides the networking problem into eight parts: 1) hardware infrastructure, 2) southbound interfaces, 3) network virtualization (hypervisor layer between the forwarding devices and the network operating systems), 4) network operating systems (SDN controllers and control platforms), 5) northbound interfaces (to offer a common programming abstraction to the upper layers, mainly the network applications), 6) virtualization using slicing techniques provided by special purpose libraries or programming languages and compilers, 7) network programming languages, and finally 8) network applications. In addition, we also look at cross-layer problems such as

debugging and troubleshooting mechanisms. The discussion in Section V on ongoing research efforts, challenges, future work and opportunities concludes this paper.

II. STATE OF QUO IN NETWORKING

Computer networks can be divided in three planes of functionality: the data, control and management planes (see Figure 3). The data plane corresponds to the networking devices, which are responsible for (efficiently) forwarding data. The control plane represents the protocols used to populate the forwarding tables of the data plane elements. The management plane includes the software services, such as SNMP-based tools [18], used to remotely monitor and configure the control functionality. Network policy is defined in the management plane, the control plane enforces the policy, and the data plane executes it by forwarding data accordingly.

In traditional IP networks, the control and data planes are tightly coupled, embedded in the same networking devices, and the whole structure is highly decentralized. This was considered important for the design of the Internet in the early days: it seemed the best way to guarantee network resilience, which was a crucial design goal. In fact, this approach has been quite effective in terms of network performance, with a rapid increase of line rate and port densities.




Fig. 3. Layered view of networking functionality.

However, the outcome is a very complex and relatively static architecture, as has been often reported in the networking literature (e.g., [1], [3], [2], [6], [19]). It is also the fundamental reason why traditional networks are rigid, and complex to manage and control. These two characteristics are largely responsible for a vertically-integrated industry where innovation is difficult.

Network misconfigurations and related errors are extremely common in today's networks. For instance, more than 1000 configuration errors have been observed in BGP routers [20]. From a single misconfigured device may result very undesired network behavior (including, among others, packet losses, forwarding loops, setting up of unintended paths, or service contract violations). Indeed, while rare, a single misconfigured router is able to compromise the correct operation of the whole Internet for hours [21], [22].

To support network management, a small number of vendors offer proprietary solutions of specialized hardware, operating systems, and control programs (network applications). Network operators have to acquire and maintain different management solutions and the corresponding specialized teams. The capital and operational cost of building and maintaining a networking infrastructure is significant, with long return on investment cycles, which hamper innovation and addition of new features and services (for instance access control, load balancing, energy efficiency, traffic engineering). To alleviate the lack of in-path functionalities within the network, a myriad of specialized components and middleboxes, such as firewalls, intrusion detection systems and deep packet inspection engines, proliferate in current networks. A recent survey of 57 enterprise networks shows that the number of middleboxes is already on par with the number of routers in current networks [23]. Despite helping in-path functionalities, the net effect of middleboxes has been increased complexity of network design and its operation.

III. WHAT IS SOFTWARE-DEFINED NETWORKING?

The term SDN (Software-Defined Networking) was originally coined to represent the ideas and work around OpenFlow at Stanford University [24]. As originally defined, SDN refers to a network architecture where the forwarding state in the data plane is managed by a remote control plane decoupled from the former. The networking industry has on many occasions

shifted from this original view of SDN, by referring to anything that involves software as being SDN. We therefore attempt, in this section, to provide a much less ambiguous definition of software-defined networking.

We define an SDN as a network architecture with four pillars:

- 1) The control and data planes are *decoupled*. Control functionality is removed from network devices that will become simple (packet) forwarding elements.
- 2) Forwarding decisions are flow-based, instead of destination-based. A flow is broadly defined by a set of packet field values acting as a match (filter) criterion and a set of actions (instructions). In the SDN/OpenFlow context, a flow is a sequence of packets between a source and a destination. All packets of a flow receive identical service policies at the forwarding devices [25], [26]. The flow abstraction allows unifying the behavior of different types of network devices, including routers, switches, firewalls, and middleboxes [27]. Flow programming enables unprecedented flexibility, limited only to the capabilities of the implemented flow tables [9].
- 3) Control logic is moved to an external entity, the so-called SDN controller or Network Operating System (NOS). The NOS is a software platform that runs on commodity server technology and provides the essential resources and abstractions to facilitate the programming of forwarding devices based on a logically centralized, abstract network view. Its purpose is therefore similar to that of a traditional operating system.
- 4) The network is *programmable* through software applications running on top of the NOS that interacts with the underlying data plane devices. This is a fundamental characteristic of SDN, considered as its main value proposition.

Note that the logical centralization of the control logic, in particular, offers several additional benefits. First, it is simpler and less error-prone to modify network policies through high-level languages and software components, compared with low-level device specific configurations. Second, a control program can automatically react to spurious changes of the network state and thus maintain the high-level policies intact. Third, the centralization of the control logic in a controller with global knowledge of the network state simplifies the development of more sophisticated networking functions, services and applications.

Following the SDN concept introduced in [5], an SDN can be defined by three fundamental abstractions: (i) forwarding, (ii) distribution, and (iii) specification. In fact, abstractions are essential tools of research in computer science and information technology, being already an ubiquitous feature of many computer architectures and systems [28].

Ideally, the *forwarding abstraction* should allow any forwarding behavior desired by the network application (the control program) while hiding details of the underlying hardware. OpenFlow is one realization of such abstraction, which can be seen as the equivalent to a "device driver" in an operating system.




Fig. 4. SDN architecture and its fundamental abstractions.

The *distribution abstraction* should shield SDN applications from the vagaries of distributed state, making the distributed control problem a logically centralized one. Its realization requires a common distribution layer, which in SDN resides in the NOS. This layer has two essential functions. First, it is responsible for installing the control commands on the forwarding devices. Second, it collects status information about the forwarding layer (network devices and links), to offer a global network view to network applications.

The last abstraction is *specification*, which should allow a network application to express the desired network behavior without being responsible for implementing that behavior itself. This can be achieved through virtualization solutions, as well as network programming languages. These approaches map the abstract configurations that the applications express based on a simplified, abstract model of the network, into a physical configuration for the global network view exposed by the SDN controller. Figure 4 depicts the SDN architecture, concepts and building blocks.

As previously mentioned, the strong coupling between control and data planes has made it difficult to add new functionality to traditional networks, a fact illustrated in Figure 5. The coupling of the control and data planes (and its physical embedding in the network elements) makes the development and deployment of new networking features (e.g., routing algorithms) very hard since it would imply a modification of the control plane of all network devices – through the installation of new firmware and, in some cases, hardware upgrades. Hence, the new networking features are commonly introduced via expensive, specialized and hard-to-configure equipment (aka middleboxes) such as load balancers, intrusion detection systems (IDS), and firewalls, among others. These middleboxes need to be placed strategically in the network, making it even harder to later change the network topology, configuration, and functionality.

In contrast, SDN decouples the control plane from the network devices and becomes an external entity: the network




Fig. 5. Traditional networking versus Software-Defined Networking (SDN). With SDN, management becomes simpler and middleboxes services can be delivered as SDN controller applications.

operating system or SDN controller. This approach has several advantages:

- It becomes easier to program these applications since the abstractions provided by the control platform and/or the network programming languages can be shared.
- All applications can take advantage of the same network information (the global network view), leading (arguably) to more consistent and effective policy decisions while re-using control plane software modules.
- These applications can take actions (i.e., reconfigure forwarding devices) from any part of the network. There is therefore no need to devise a precise strategy about the location of the new functionality.
- The integration of different applications becomes more straightforward [29]. For instance, load balancing and routing applications can be combined sequentially, with load balancing decisions having precedence over routing policies.

A. Terminology

To identify the different elements of an SDN as unequivocally as possible, we now present the essential terminology used throughout this work.

Forwarding Devices (FD): Hardware- or software-based data plane devices that perform a set of elementary operations. The forwarding devices have well-defined instruction sets (e.g., flow rules) used to take actions on the incoming packets (e.g., forward to specific ports, drop, forward to the controller, rewrite some header). These instructions are defined by south-bound interfaces (e.g., OpenFlow [9], ForCES [30], Protocol-Oblivious Forwarding (POF) [31]) and are installed in the

forwarding devices by the SDN controllers implementing the southbound protocols.

Data Plane (DP): Forwarding devices are interconnected through wireless radio channels or wired cables. The network infrastructure comprises the interconnected forwarding devices, which represent the data plane.

Southbound Interface (SI): The instruction set of the forwarding devices is defined by the southbound API, which is part of the southbound interface. Furthermore, the SI also defines the communication protocol between forwarding devices and control plane elements. This protocol formalizes the way the control and data plane elements interact.

Control Plane (CP): Forwarding devices are programmed by control plane elements through well-defined SI embodiments. The control plane can therefore be seen as the “network brain”. All control logic rests in the applications and controllers, which form the control plane.

Northbound Interface (NI): The network operating system can offer an API to application developers. This API represents a northbound interface, i.e., a common interface for developing applications. Typically, a northbound interface abstracts the low level instruction sets used by southbound interfaces to program forwarding devices.

Management Plane (MP): The management plane is the set of applications that leverage the functions offered by the NI to implement network control and operation logic. This includes applications such as routing, firewalls, load balancers, monitoring, and so forth. Essentially, a management application defines the policies, which are ultimately translated to southbound-specific instructions that program the behavior of the forwarding devices.

B. Alternative and Broadening Definitions

Since its inception in 2010 [24], the original OpenFlow-centered SDN term has seen its scope broadened beyond architectures with a cleanly decoupled control plane interface. The definition of SDN will likely continue to broaden, driven by the industry business-oriented views on SDN – irrespective of the decoupling of the control plane. In this survey, we focus on the original, “canonical” SDN definition based on the aforementioned key pillars and the concept of layered abstractions. However, for the sake of completeness and clarity, we acknowledge alternative SDN definitions [32], including:

Control Plane / Broker SDN: A networking approach that retains existing distributed control planes but offers new APIs that allow applications to interact (bidirectionally) with the network. An SDN controller –often called orchestration platform– acts as a broker between the applications and the network elements. This approach effectively presents control plane data to the application and allows a certain degree of network programmability by means of “plug-ins” between the orchestrator function and network protocols. This API-driven approach corresponds to a hybrid model of SDN, since it enables the broker to manipulate and directly interact with the control planes of devices such as routers and switches. Examples of this view on SDN include recent standardization efforts at IETF (see Section III-C) and the design philosophy

behind the OpenDaylight project [13] that goes beyond the OpenFlow split control mode.

Overlay SDN: A networking approach where the (software- or hardware-based) network edge is dynamically programmed to manage tunnels between hypervisors and/or network switches, introducing an overlay network. In this hybrid networking approach, the distributed control plane providing the underlay remains untouched. The centralized control plane provides a logical overlay that utilizes the underlay as a transport network. This flavor of SDN follows a proactive model to install the overlay tunnels. The overlay tunnels usually terminate inside virtual switches within hypervisors or in physical devices acting as gateways to the existing network. This approach is very popular in recent data center network virtualization [33], and are based on a variety of tunneling technologies (e.g., STT [34], VXLAN [35], NVGRE [36], LISP [37], [38], GENEVE [39]) [40].

Recently, other attempts to define SDN in a layered approach have appeared [41], [16]. From a practical perspective and trying to keep backward compatibility with existing network management approaches, one initiative at IRTF SD-NRG [41] proposes a management plane at the same level of the control plane, i.e., it classifies solutions in two categories: control logic (with control plane southbound interfaces) and management logic (with management plane southbound interfaces). In other words, the management plane can be seen as a control platform that accommodates traditional network management services and protocols, such as SNMP [18], BGP [42], PCEP [43], and NETCONF [44].

In addition the broadening definitions above, the term SDN is often used to define extensible network management planes (e.g., OpenStack [45]), whitebox / bare-metal switches with open operating systems (e.g., Cumulus Linux), open-source dataplanes (e.g., Pica8 Xorplus [46], Quagga [47]), specialized programmable hardware devices (e.g., NetFPGA [48]), virtualized software-based appliances (e.g., Open Platform for Network Functions Virtualization - OPNFV [49]), in spite of lacking a decoupled control and data plane or common interface along its API. Hybrid SDN models are further discussed in Section V-G.

C. Standardization Activities

The standardization landscape in SDN (and SDN-related issues) is already wide and is expected to keep evolving over time. While some of the activities are being carried out in Standard Development Organizations (SDOs), other related efforts are ongoing at industrial or community consortia (e.g., OpenDaylight, OpenStack, OPNFV), delivering results often considered candidates for *de facto* standards. These results often come in the form of open source implementations that have become the common strategy towards accelerating SDN and related cloud and networking technologies [50]. The reason for this fragmentation is due to SDN concepts spanning different areas of IT and networking, both from a network segmentation point of view (from access to core) and from a technology perspective (from optical to wireless).

Table I presents a summary of the main SDOs and organizations contributing to the standardization of SDN, as well as

the main outcomes produced to date.

The Open Networking Foundation (ONF) was conceived as a member-driven organization to promote the adoption of SDN through the development of the OpenFlow protocol as an open standard to communicate control decisions to data plane devices. The ONF is structured in several working groups (WGs). Some WGs are focused on either defining extensions to the OpenFlow protocol in general, such as the Extensibility WG, or tailored to specific technological areas. Examples of the latter include the Optical Transport (OT) WG, the Wireless and Mobile (W&M) WG, and the Northbound Interfaces (NBI) WG. Other WGs center their activity in providing new protocol capabilities to enhance the protocol itself, such as the Architecture WG or the Forwarding Abstractions (FA) WG.

Similar to how network programmability ideas have been considered by several Working Groups (WGs) of the Internet Engineering Task Force (IETF) in the past, the present SDN trend is also influencing a number of activities. A related body that focuses on research aspects for the evolution of the Internet, the Internet Research Task Force (IRTF), has created the Software Defined Networking Research Group (SDNRG). This group investigates SDN from various perspectives with the goal of identifying the approaches that can be defined, deployed and used in the near term, as well as identifying future research challenges.

In the International Telecommunications Union's Telecommunication sector (ITU-T), some Study Groups (SGs) have already started to develop recommendations for SDN, and a Joint Coordination Activity on SDN (JCA-SDN) has been established to coordinate the SDN standardization work.

The Broadband Forum (BBF) is working on SDN topics through the Service Innovation & Market Requirements (SIMR) WG. The objective of the BBF is to release recommendations for supporting SDN in multi-service broadband networks, including hybrid environments where only some of the network equipment is SDN-enabled.

The Metro Ethernet Forum (MEF) is approaching SDN with the aim of defining service orchestration with APIs for existing networks.

At the Institute of Electrical and Electronics Engineers (IEEE), the 802 LAN/MAN Standards Committee has recently started some activities to standardize SDN capabilities on access networks based on IEEE 802 infrastructure through the P802.1CF project, for both wired and wireless technologies to embrace new control interfaces.

The Optical Internetworking Forum (OIF) Carrier WG released a set of requirements for Transport Software-Defined Networking. The initial activities have as main goal to describe the features and functionalities needed to support the deployment of SDN capabilities in carrier transport networks.

The Open Data Center Alliance (ODCA) is an organization working on unifying data center in the migration to cloud computing environments through interoperable solutions. Through the documentation of usage models, specifically one for SDN, the ODCA is defining new requirements for cloud deployment.

The Alliance for Telecommunication Industry Solutions (ATIS) created a Focus Group for analyzing operational issues and opportunities associated with the programmable capabilities

of network infrastructure.

At the European Telecommunication Standards Institute (ETSI), efforts are being devoted to Network Function Virtualization (NFV) through a newly defined Industry Specification Group (ISG). NFV and SDN concepts are considered complementary, sharing the goal of accelerating innovation inside the network by allowing programmability, and altogether changing the network operational model through automation and a real shift to software-based platforms.

Finally, the mobile networking industry 3GPP consortium is studying the management of virtualized networks, an effort aligned with the ETSI NFV architecture and, as such, likely to leverage from SDN.

D. History of Software-Defined Networking

Albeit a fairly recent concept, SDN leverages on networking ideas with a longer history [17]. In particular, it builds on work made on programmable networks, such as active networks [81], programmable ATM networks [82], [83], and on proposals for control and data plane separation, such as NCP [84] and RCP [85].

In order to present an historical perspective, we summarize in Table II different instances of SDN-related work prior to SDN, splitting it into five categories. Along with the categories we defined, the second and third columns of the table mention past initiatives (pre-SDN, i.e., before the OpenFlow-based initiatives that sprung into the SDN concept), and recent developments that led to the definition of SDN.

Data plane programmability has a long history. Active networks [81] represent one of the early attempts on building new network architectures based on this concept. The main idea behind active networks is for each node to have the capability to perform computations on, or modify the content of, packets. To this end, active networks propose two distinct approaches: programmable switches and capsules. The former does not imply changes in the existing packet or cell format. It assumes that switching devices support the downloading of programs with specific instructions on how to process packets. The second approach, on the other hand, suggests that packets should be replaced by tiny programs, which are encapsulated in transmission frames and executed at each node along their path.

ForCES [30], OpenFlow [9] and POF [31] represent recent approaches for designing and deploying programmable data plane devices. In a manner different from active networks, these new proposals rely essentially on modifying forwarding devices to support flow tables, which can be dynamically configured by remote entities through simple operations such as adding, removing or updating flow rules, i.e., entries on the flow tables.

The earliest initiatives on separating data and control signalling date back to the 80s and 90s. The network control point (NCP) [84] is probably the first attempt to separate control and data plane signalling. NCPs were introduced by AT&T to improve the management and control of its telephone network. This change promoted a faster pace of innovation of the network and provided new means for improving its efficiency,

TABLE I
OPENFLOW STANDARDIZATION ACTIVITIES

SDO	Working Group	Focus	Outcomes
ONF	Architecture & Framework	SDN architecture, defining architectural components and interfaces	SDN Architecture [51]
	Northbound Interfaces	Definition of standard NBIs for SDN controllers	
	Testing and Interoperability	Specification of OpenFlow conformance test suites	Conformance tests [52]
	Extensibility	Development of extensions to OpenFlow protocol, producing specifications of the OpenFlow switch (OF-WIRE) protocol	OF-WIRE 1.4.0 [53]
	Configuration & Management	OAM (operation, administration, and management) capabilities for OF protocol, producing specifications of the OF Configuration and Management (OF-CONFIG) protocol	OF-CONFIG 1.2 [54] OpenFlow Notifications Framework [55]
	Forwarding Abstractions	Development of hardware abstractions and simplification of behavioral descriptions mapping	OpenFlow Table Type Patterns [56]
	Optical Transport	Specification of SDN and control capabilities for optical transport networks by means of OpenFlow	Use cases [57] Requirements [58]
	Wireless & Mobile	Specification of SDN and control capabilities for wireless and mobile networks by means of OpenFlow	
	Migration	Methods to migrate from conventional networks to SDN-based networks based on OpenFlow	Use cases [59]
	Market Education	Dissemination of ONF initiatives in SDN and OpenFlow by releasing White Papers and Solution Briefs	SDN White Paper [60]
IETF	Application-Layer Traffic Optimization (ALTO)	Provides applications with network state information	Architectures for the coexistence of SDN and ALTO [61]
	Forwarding and Control Element Separation (ForCES)	Protocol specifications for the communication between control and forwarding elements.	Protocol specification [30]
	Interface to the Routing System (I2RS)	Real-time or event driven interaction with the routing system in an IP routed network	Architecture [62]
	Network Configuration (NETCONF)	Protocol specification for transferring configuration data to and from a device	NETCONF protocol [63]
	Network Virtualization Overlays (NVO3)	Overlay networks for supporting multi-tenancy in the context of data center communications (i.e., VM communication)	Control plane requirements [64]
IETF	Path Computation Element (PCE)	Path computation for traffic engineering and path selection based on constrains	ABNO framework [65] Cross stratum path computation [66]
	Source Packet Routing in Networking (SPRING)	Specification of a forwarding path at the source of traffic	OpenFlow interworking [67] SDN controlled use cases [68]
	Abstraction and Control of Transport Networks (ACTN) BoF	Facilitate a centralized virtual network operation	Virtual network controller framework [69]
IRTF	Software-Defined Networking Research Group (SDNRG)	Prospection of SDN for the evolution of Internet	SDN operator perspective [70] SDN Architecture [71] Service / Transport separation [72]
ITU-T	SG 11	Signalling requirements using SDN technologies in Broadband Access Networks	Q.Supplement-SDN [73] Q.SBAN [74]
	SG 13	Functional requirements and architecture for SDN and networks of the future	Recommendation Y.3300 [75]
	SG 15	Specification of a transport network control plane architecture to support SDN control of transport networks	
	SG 17	Architectural aspects of security in SDN and security services using SDN	
BBF	Service Innovation and Market Requirements	Requirements and impacts of deploying SDN in broadband networks	SD-313 [76]
MEF	The Third Network	Service orchestration in Network as a Service and NFV environments	
IEEE	802	Applicability of SDN to IEEE 802 infrastructure	
OIF	Carrier WG	Transport SDN networks	Requirements for SDN enabled transport networks [77]
ODCA	SDN/Infrastructure	Requirements for SDN in cloud environments	Usage model [78]
ETSI	NFV ISG	Orchestration of network functions, including the combined control of computing, storage and networking resources	NFV Architecture [79]
ATIS	SDN Focus Group	Operational aspects of SDN and NFV	Operation of SDN [80]

TABLE II
SUMMARIZED OVERVIEW OF THE HISTORY OF PROGRAMABLE NETWORKS

Category	Pre-SDN initiatives	More recent SDN developments
Data plane programmability	xbind [82], IEEE P1520 [86], smart packets [87], ANTS [88], SwitchWare [89], Calvert [90], high performance router [91], NetScript [92], Tennenhouse [93]	ForCES [30], OpenFlow [9], POF [31]
Control and data plane decoupling	NCP [84], GSMP [94], [95], Tempest [96], ForCES [30], RCP [85], Soft-Router [97], PCE [43], 4D [98], IRSCP [99]	SANE [100], Ethane [101], OpenFlow [9], NOX [26], POF [31]
Network virtualization	Tempest [96], MBone [102], 6Bone [103], RON [104], Planet Lab [105], Impasse [106], GENI [107], VINI [108]	Open vSwitch [109], Mininet [110], FlowVisor [111], NVP [112]
Network operating systems	Cisco IOS [113], JUNOS [114], ExtremeXOS [115], SR OS [116]	NOX [26], Onix [7], ONOS [117]
Technology pull initiatives	Open Signaling [118]	ONF [10]

by taking advantage of the global view of the network provided by NCPs. Similarly, other initiatives such as Tempest [96], ForCES [30], RCP [85], and PCE [43] proposed the separation of the control and data planes for improved management in ATM, Ethernet, BGP, and MPLS networks, respectively.

More recently, initiatives such as SANE [100], Ethane [101], OpenFlow [9], NOX [26] and POF [31] proposed the decoupling of the control and data planes for Ethernet networks. Interestingly, these recent solutions do not require significant modifications on the forwarding devices, making them attractive not only for the networking research community, but even more to the networking industry. OpenFlow-based devices [9], for instance, can easily co-exist with traditional Ethernet devices, enabling a progressive adoption (i.e., not requiring a disruptive change to existing networks).

Network virtualization has gained a new traction with the advent of SDN. Nevertheless, network virtualization also has its roots back in the 90s. The Tempest project [96] is one of the first initiatives to introduce network virtualization, by introducing the concept of switchlets in ATM networks. The core idea was to allow multiple switchlets on top of a single ATM switch, enabling multiple independent ATM networks to share the same physical resources. Similarly, MBone [102] was one of the early initiatives that targeted the creation of virtual network topologies on top of legacy networks, or overlay networks. This work was followed by several other projects such as Planet Lab [105], GENI [107] and VINI [108]. It is also worth mentioning FlowVisor [119] as one of the first recent initiatives to promote a hypervisor-like virtualization architecture for network infrastructures, resembling the hypervisor model common for compute and storage. More recently, Koponen et al. proposed a Network Virtualization Platform (NVP [112]) for multi-tenant datacenters using SDN as a base technology.

The concept of a network operating system was reborn with the introduction of OpenFlow-based network operating systems, such as NOX [26], Onix [7] and ONOS [117]. Indeed, network operating systems have been in existence for decades. One of the most widely known and deployed is the Cisco IOS [113], which was originally conceived back in the early 90s. Other network operating systems worth mentioning are JUNOS [114], ExtremeXOS [115] and SR OS [116]. Despite

being more specialized network operating systems, targeting network devices such as high-performance core routers, these NOSs abstract the underlying hardware to the network operator, making it easier to control the network infrastructure as well as simplifying the development and deployment of new protocols and management applications.

Finally, it is also worth recalling initiatives that can be seen as “technology pull” drivers. Back in the 90s, a movement towards open signalling [118] started to happen. The main motivation was to promote the wider adoption of the ideas proposed by projects such as NCP [84] and Tempest [96]. The open signalling movement worked towards separating the control and data signalling, by proposing open and programmable interfaces. Curiously, a rather similar movement can be observed with the recent advent of OpenFlow and SDN, with the lead of the Open Networking Foundation (ONF) [10]. This type of movement is crucial to promote open technologies into the market, hopefully leading equipment manufacturers to support open standards and thus fostering interoperability, competition, and innovation.

For a more extensive intellectual history of programmable networks and SDN we forward the reader to the recent paper by Feamster et al. [17].

IV. SOFTWARE-DEFINED NETWORKS: BOTTOM-UP

An SDN architecture can be depicted as a composition of different layers, as shown in Figure 6 (b). Each layer has its own specific functions. While some of them are always present in an SDN deployment, such as the southbound API, network operating systems, northbound API and network applications, others may be present only in particular deployments, such as hypervisor- or language-based virtualization.

Figure 6 presents a tri-fold perspective of SDNs. The SDN layers are represented in the center (b) of the figure, as explained above. Figures 6 (a) and 6 (c) depict a plane-oriented view and a system design perspective, respectively.

The following sections introduce each layer, following a bottom-up approach. For each layer, the core properties and concepts are explained based on the different technologies and solutions. Additionally, debugging and troubleshooting techniques and tools are discussed.




Fig. 6. Software-Defined Networks in (a) planes, (b) layers, and (c) system design architecture

A. Layer I: Infrastructure

An SDN infrastructure, similarly to a traditional network, is composed of a set of networking equipment (switches, routers and middlebox appliances). The main difference resides in the fact that those traditional physical devices are now simple forwarding elements without embedded control or software to take autonomous decisions. The network intelligence is removed from the data plane devices to a logically-centralized control system, i.e., the network operating system and applications, as shown in Figure 6 (c). More importantly, these new networks are built (conceptually) on top of open and standard interfaces (e.g., OpenFlow), a crucial approach for ensuring configuration and communication compatibility and interoperability among different data and control plane devices. In other words, these open interfaces enable controller entities to dynamically program heterogeneous forwarding devices, something difficult in traditional networks, due to the large variety of proprietary and closed interfaces and the distributed nature of the control plane.

In an SDN/OpenFlow architecture, there are two main elements, the controllers and the forwarding devices, as shown in Figure 7. A data plane device is a hardware or software element specialized in packet forwarding, while a controller is a software stack (the “network brain”) running on a commodity hardware platform. An OpenFlow-enabled forwarding device is based on a pipeline of flow tables where each entry of a flow table has three parts: (1) a matching rule, (2) actions to be executed on matching packets, and (3) counters that keep statistics of matching packets. This high-level and simplified model derived from OpenFlow is currently the most widespread design of SDN data plane devices. Nevertheless, other specifications of SDN-enabled forwarding devices are being pursued, including POF [31], [120] and the Negotiable Datapath Models (NDMs) from the ONF Forwarding Abstractions Working Group (FAWG) [121].

Inside an OpenFlow device, a path through a sequence of flow tables defines how packets should be handled. When a

new packet arrives, the lookup process starts in the first table and ends either with a match in one of the tables of the pipeline or with a miss (when no rule is found for that packet). A flow rule can be defined by combining different matching fields, as illustrated in Figure 7. If there is no default rule, the packet will be discarded. However, the common case is to install a default rule which tells the switch to send the packet to the controller (or to the normal non-OpenFlow pipeline of the switch). The priority of the rules follows the natural sequence number of the tables and the row order in a flow table. Possible actions include (1) forward the packet to outgoing port(s), (2) encapsulate it and forward it to the controller, (3) drop it, (4) send it to the normal processing pipeline, (5) send it to the next flow table or to special tables, such as group or metering tables introduced in the latest OpenFlow protocol.

As detailed in Table III, each version of the OpenFlow specification introduced new match fields including Ethernet, IPv4/v6, MPLS, TCP/UDP, etc. However, only a subset of those matching fields are mandatory to be compliant to a given protocol version. Similarly, many actions and port types are optional features. Flow match rules can be based on almost arbitrary combinations of bits of the different packet headers using bit masks for each field. Adding new matching fields has been eased with the extensibility capabilities introduced in OpenFlow version 1.2 through an OpenFlow Extensible Match (OXM) based on type-length-value (TLV) structures. To improve the overall protocol extensibility, with OpenFlow version 1.4 TLV structures have been also added to ports, tables, and queues in replacement of the hard-coded counterparts of earlier protocol versions.

Overview of available OpenFlow devices

Several OpenFlow-enabled forwarding devices are available on the market, both as commercial and open source products (see Table IV). There are many off-the-shelf, ready to deploy, OpenFlow switches and routers, among other appliances. Most of the switches available on the market have relatively small Ternary Content-Addressable Memory (TCAMs), with up to




Fig. 7. OpenFlow-enabled SDN devices

TABLE III

DIFFERENT MATCH FIELDS, STATISTICS AND CAPABILITIES HAVE BEEN ADDED ON EACH OPENFLOW PROTOCOL REVISION. THE NUMBER OF REQUIRED (REQ) AND OPTIONAL (OPT) CAPABILITIES HAS GROWN CONSIDERABLY.

OpenFlow Version	Match fields	Statistics	# Matches		# Instructions		# Actions		# Ports	
			Req	Opt	Req	Opt	Req	Opt	Req	Opt
v 1.0	Ingress Port	Per table statistics	18	2	1	0	2	11	6	2
	Ethernet: src, dst, type, VLAN	Per flow statistics								
	IPv4: src, dst, proto, ToS	Per port statistics								
	TCP/UDP: src port, dst port	Per queue statistics								
v 1.1	Metadata, SCTP, VLAN tagging	Group statistics	23	2	0	0	3	28	5	3
	MPLS: label, traffic class	Action bucket statistics								
v 1.2	OpenFlow Extensible Match (OXM)		14	18	2	3	2	49	5	3
	IPv6: src, dst, flow label, ICMPv6									
v 1.3	PBB, IPv6 Extension Headers	Per-flow meter	14	26	2	4	2	56	5	3
		Per-flow meter band								
v 1.4	—	—	14	27	2	4	2	57	5	3
		Optical port properties								

8K entries. Nonetheless, this is changing at a fast pace. Some of the latest devices released in the market go far beyond that figure. Gigabit Ethernet (GbE) switches for common business purposes are already supporting up to 32K L2+L3 or 64K L2/L3 exact match flows [122]. Enterprise class 10GbE switches are being delivered with more than 80K Layer 2 flow entries [123]. Other switching devices using high performance chips (e.g., EZchip NP-4) provide optimized TCAM memory that supports from 125K up to 1000K flow table entries [124]. This is a clear sign that the size of the flow tables is growing at a pace aiming to meet the needs of future SDN deployments.

Networking hardware manufacturers have produced various kinds of OpenFlow-enabled devices, as is shown in Table IV. These devices range from equipment for small businesses (e.g., GbE switches) to high-class data center equipment (e.g., high-density switch chassis with up to 100GbE connectivity for edge-to-core applications, with tens of Tbps of switching capacity).

Software switches are emerging as one of the most promising solutions for data centers and virtualized network infrastructures [147], [148], [149]. Examples of software-based OpenFlow switch implementations include Switch Light [145],

ofsoftswitch13 [141], Open vSwitch [142], OpenFlow Reference [143], Pica8 [150], Pantou [146], and XorPlus [46]. Recent reports show that the number of virtual access ports is already larger than physical access ports on data centers [149]. Network virtualization has been one of the drivers behind this trend. Software switches such as Open vSwitch have been used for moving network functions to the edge (with the core performing traditional IP forwarding), thus enabling network virtualization [112].

An interesting observation is the number of small, start-up enterprises devoted to SDN, such as Big Switch, Pica8, Cyan, Plexxi, and NoviFlow. This seems to imply that SDN is springing a more competitive and open networking market, one of its original goals. Other effects of this openness triggered by SDN include the emergence of so-called “bare metal switches” or “whitebox switches”, where the software and hardware are sold separately and the end-user is free to load an operating system of its choice [151].

B. Layer II: Southbound Interfaces

Southbound interfaces (or southbound APIs) are the connecting bridges between control and forwarding elements, thus

TABLE IV
OPENFLOW ENABLED HARDWARE AND SOFTWARE DEVICES

Group	Product	Type	Maker/Developer	Version	Short description
Hardware	8200zl and 5400zl [125]	chassis	Hewlett-Packard	v1.0	Data center class chassis (switch modules).
	Arista 7150 Series [126]	switch	Arista Networks	v1.0	Data centers hybrid Ethernet/OpenFlow switches.
	BlackDiamond X8 [127]	switch	Extreme Networks	v1.0	Cloud-scale hybrid Ethernet/OpenFlow switches.
	CX600 Series [128]	router	Huawei	v1.0	Carrier class MAN routers.
	EX9200 Ethernet [129]	chassis	Juniper	v1.0	Chassis based switches for cloud data centers.
	EZchip NP-4 [130]	chip	EZchip Technologies	v1.1	High performance 100-Gigabit network processors.
	MLX Series [131]	router	Brocade	v1.0	Service providers and enterprise class routers.
	NoviSwitch 1248 [124]	switch	NoviFlow	v1.3	High performance OpenFlow switch.
	NetFPGA [48]	card	NetFPGA	v1.0	1G and 10G OpenFlow implementations.
	RackSwitch G8264 [132]	switch	IBM	v1.0	Data center switches supporting Virtual Fabric and OpenFlow.
	PF5240 and PF5820 [133]	switch	NEC	v1.0	Enterprise class hybrid Ethernet/OpenFlow switches.
	Pica8 3920 [134]	switch	Pica8	v1.0	Hybrid Ethernet/OpenFlow switches.
	Plexxi Switch 1 [135]	switch	Plexxi	v1.0	Optical multiplexing interconnect for data centers.
	V330 Series [136]	switch	Centec Networks	v1.0	Hybrid Ethernet/OpenFlow switches.
	Z-Series [137]	switch	Cyan	v1.0	Family of packet-optical transport platforms.
Software	contrail-vrouter [138]	vrouter	Juniper Networks	v1.0	Data-plane function to interface with a VRF.
	LINC [139], [140]	switch	FlowForwarding	v1.4	Erlang-based soft switch with OF-Config 1.1 support.
	ofsoftswitch13 [141]	switch	Ericsson, CPqD	v1.3	OF 1.3 compatible user-space software switch implementation.
	Open vSwitch [142], [109]	switch	Open Community	v1.0-1.3	Switch platform designed for virtualized server environments.
	OpenFlow Reference [143]	switch	Stanford	v1.0	OF Switching capability to a Linux PC with multiple NICs.
	OpenFlowClick [144]	vrouter	Yogesh Mundada	v1.0	OpenFlow switching element for Click software routers.
	Switch Light [145]	switch	Big Switch	v1.0	Thin switching software platform for physical/virtual switches.
	Pantou/OpenWRT [146]	switch	Stanford	v1.0	Turns a wireless router into an OF-enabled switch.
XorPlus [46]	switch	Pica8	v1.0	Switching software for high performance ASICs.	

being the crucial instrument for clearly separating control and data plane functionality. However, these APIs are still tightly tied to the forwarding elements of the underlying physical or virtual infrastructure.

Typically, a new switch can take two years to be ready for commercialization if built from scratch, with upgrade cycles that can take up to nine months. The software development for a new product can take from six months to one year [152]. The initial investment is high and risky. As a central component of its design the southbound APIs represent one of the major barriers for the introduction and acceptance of any new networking technology. In this light, the emergence of SDN southbound API proposals such as OpenFlow [9] is seen as welcome by many in the industry. These standards promote interoperability, allowing the deployment of vendor-agnostic network devices. This has already been demonstrated by the interoperability between OpenFlow-enabled equipments from different vendors.

As of this writing, OpenFlow is the most widely accepted and deployed open southbound standard for SDN. It provides a common specification to implement OpenFlow-enabled forwarding devices, and for the communication channel between data and control plane devices (e.g., switches and controllers). The OpenFlow protocol provides three information sources for network operating systems. First, event-based messages are sent by forwarding devices to the controller when a link or

port change is triggered. Second, flow statistics are generated by the forwarding devices and collected by the controller. Third, packet-in messages are sent by forwarding devices to the controller when they do not know what to do with a new incoming flow or because there is an explicit “send to controller” action in the matched entry of the flow table. These information channels are the essential means to provide flow-level information to the network operating system.

Albeit the most visible, OpenFlow is not the only available southbound interface for SDN. There are other API proposals such as ForCES [30], OVSDB [153], POF [31], [120], OpFlex [154], OpenState [155], Revised OpenFlow Library (ROFL) [156], Hardware Abstraction Layer (HAL) [157], [158], and Programmable Abstraction of Datapath (PAD) [159]. ForCES proposes a more flexible approach to traditional network management without changing the current architecture of the network, i.e., without the need of a logically-centralized external controller. The control and data planes are separated but can potentially be kept in the same network element. However, the control part of the network element can be upgraded on-the-fly with third-party firmware.

OVSDB [153] is another type of southbound API, designed to provide advanced management capabilities for Open vSwitches. Beyond OpenFlow’s capabilities to configure the behavior of flows in a forwarding device, an Open vSwitch offers other networking functions. For instance, it allows the

control elements to create multiple virtual switch instances, set QoS policies on interfaces, attach interfaces to the switches, configure tunnel interfaces on OpenFlow data paths, manage queues, and collect statistics. Therefore, the OVSDB is a complementary protocol to OpenFlow for Open vSwitch.

One of the first direct competitors of OpenFlow is POF [31], [120]. One of the main goals of POF is to enhance the current SDN forwarding plane. With OpenFlow, switches have to understand the protocol headers to extract the required bits to be matched with the flow tables entries. This parsing represents a significant burden for data plane devices, in particular if we consider that OpenFlow version 1.3 already contains more than 40 header fields. Besides this inherent complexity, backward compatibility issues may arise every time new header fields are included in or removed from the protocol. To achieve its goal, POF proposes a generic flow instruction set (FIS) that makes the forwarding plane protocol-oblivious. A forwarding element does not need to know, by itself, anything about the packet format in advance. Forwarding devices are seen as white boxes with only processing and forwarding capabilities. In POF, packet parsing is a controller task that results in a sequence of generic keys and table lookup instructions that are installed in the forwarding elements. The behavior of data plane devices is therefore completely under the control of the SDN controller. Similar to a CPU in a computer system, a POF switch is application- and protocol-agnostic.

A recent southbound interface proposal is OpFlex [154]. Contrary to OpenFlow (and similar to ForCES), one of the ideas behind OpFlex is to distribute part of the complexity of managing the network back to the forwarding devices, with the aim of improving scalability. Similar to OpenFlow, policies are logically centralized and abstracted from the underlying implementation. The differences between OpenFlow and OpFlex are a clear illustration of one of the important questions to be answered when devising a southbound interface: where to place each piece of the overall functionality.

In contrast to OpFlex and POF, OpenState [155] and ROFL [156] do not propose a new set of instructions for programming data plane devices. OpenState proposes extended finite machines (stateful programming abstractions) as an extension (super-set) of the OpenFlow match/action abstraction. Finite state machines allow the implementation of several stateful tasks inside forwarding devices, i.e., without augmenting the complexity or overhead of the control plane. For instance, all tasks involving only local state, such as MAC learning operations, port knocking or stateful edge firewalls can be performed directly on the forwarding devices without any extra control plane communication and processing delay. ROFL, on the other hand, proposes an abstraction layer that hides the details of the different OpenFlow versions, thus providing a clean API for software developers, simplifying application development.

HAL [157], [158] is not exactly a southbound API, but is closely related. Differently from the aforementioned approaches, HAL is rather a translator that enables a southbound API such as OpenFlow to control heterogeneous hardware devices. It thus sits between the southbound API and the hardware device. Recent research experiments with HAL

have demonstrated the viability of SDN control in access networks such as Gigabit Ethernet passive optical networks (GEPONs) [160] and cable networks (DOCSIS) [161]. A similar effort to HAL is the Programmable Abstraction of Datapath (PAD) [159], a proposal that goes a bit further by also working as a southbound API by itself. More importantly, PAD allows a more generic programming of forwarding devices by enabling the control of datapath behavior using generic byte operations, defining protocol headers and providing function definitions.

C. Layer III: Network Hypervisors

Virtualization is already a consolidated technology in modern computers. The fast developments of the past decade have made virtualization of computing platforms mainstream. Based on recent reports, the number of virtual servers has already exceeded the number of physical servers [162], [112].

Hypervisors enable distinct virtual machines to share the same hardware resources. In a cloud infrastructure-as-a-service (IaaS), each user can have its own virtual resources, from computing to storage. This enabled new revenue and business models where users allocate resources on-demand, from a shared physical infrastructure, at a relatively low cost. At the same time, providers make better use of the capacity of their installed physical infrastructures, creating new revenue streams without significantly increasing their CAPEX and OPEX costs. One of the interesting features of virtualization technologies today is the fact that virtual machines can be easily migrated from one physical server to another and can be created and/or destroyed on-demand, enabling the provisioning of elastic services with flexible and easy management. Unfortunately, virtualization has been only partially realized in practice. Despite the great advances in virtualizing computing and storage elements, the network is still mostly statically configured in a box-by-box manner [33].

The main network requirements can be captured along two dimensions: network topology and address space. Different workloads require different network topologies and services, such as flat L2 or L3 services, or even more complex L4-L7 services for advanced functionality. Currently, it is very difficult for a single physical topology to support the diverse demands of applications and services. Similarly, address space is hard to change in current networks. Nowadays, virtualized workloads have to operate in the same address of the physical infrastructure. Therefore, it is hard to keep the original network configuration for a tenant, virtual machines can not migrate to arbitrary locations, and the addressing scheme is fixed and hard to change. For example, IPv6 cannot be used by the VMs of a tenant if the underlying physical forwarding devices support only IPv4.

To provide complete virtualization the network should provide similar properties to the computing layer [33]. The network infrastructure should be able to support arbitrary network topologies and addressing schemes. Each tenant should have the ability to configure both the computing nodes and the network simultaneously. Host migration should automatically trigger the migration of the corresponding virtual network

ports. One might think that long standing virtualization primitives such as VLANs (virtualized L2 domain), NAT (Virtualized IP address space), and MPLS (virtualized path) are enough to provide full and automated network virtualization. However, these technologies are anchored on a box-by-box basis configuration, i.e., there is no single unifying abstraction that can be leveraged to configure (or reconfigure) the network in a global manner. As a consequence, current network provisioning can take months, while computing provisioning takes only minutes [112], [163], [164], [165].

There is hope that this situation will change with SDN and the availability of new tunneling techniques (e.g., VXLAN [35], NVGRE [36]). For instance, solutions such as FlowVisor [166], [111], [167], FlowN [168], NVP [112], OpenVirteX [169], [170], IBM SDN VE [171], [172], RadioVisor [173], AutoVFlow [174], eXtensible Datapath Daemon (xDPd) [175], [176], optical transport network virtualization [177], and version-agnostic OpenFlow slicing mechanisms [178], have been recently proposed, evaluated and deployed in real scenarios for on-demand provisioning of virtual networks.

Slicing the network

FlowVisor is one of the early technologies to virtualize a SDN. Its basic idea is to allow multiple logical networks share the same OpenFlow networking infrastructure. For this purpose, it provides an abstraction layer that makes it easier to slice a data plane based on off-the-shelf OpenFlow-enabled switches, allowing multiple and diverse networks to co-exist.

Five slicing dimensions are considered in FlowVisor: bandwidth, topology, traffic, device CPU and forwarding tables. Moreover, each network slice supports a controller, i.e., multiple controllers can co-exist on top of the same physical network infrastructure. Each controller is allowed to act only on its own network slice. In general terms, a slice is defined as a particular set of flows on the data plane. From a system design perspective, FlowVisor is a transparent proxy that intercepts OpenFlow messages between switches and controllers. It partitions the link bandwidth and flow tables of each switch. Each slice receives a minimum data rate and each guest controller gets its own virtual flow table in the switches.

Similarly to FlowVisor, OpenVirteX [169], [170] acts as a proxy between the network operating system and the forwarding devices. However, its main goal is to provide virtual SDNs through both topology, address, and control function virtualization. All these properties are necessary in multi-tenant environments where virtual networks need to be managed and migrated according to the computing and storage virtual resources. Virtual network topologies have to be mapped onto the underlying forwarding devices, with virtual addresses allowing tenants to completely manage their address space without depending on the underlying network elements addressing schemes.

AutoSlice [179] is another SDN-based virtualization proposal. Differently from FlowVisor, it focuses on the automation of the deployment and operation of vSDN (virtual SDN) topologies with minimal mediation or arbitration by the substrate network operator. Additionally, AutoSlice targets

also scalability aspects of network hypervisors by optimizing resource utilization and by mitigating the flow-table limitations through a precise monitoring of the flow traffic statistics. Similarly to AutoSlice, AutoVFlow [174] also enables multi-domain network virtualization. However, instead of having a single third party to control the mapping of vSDN topologies, as is the case of AutoSlice, AutoVFlow uses a multi-proxy architecture that allows network owners to implement flow space virtualization in an autonomous way by exchanging information among the different domains.

FlowN [168], [180] is based on a slightly different concept. Whereas FlowVisor can be compared to a full virtualization technology, FlowN is analogous to a container-based virtualization, i.e., a lightweight virtualization approach. FlowN was also primarily conceived to address multi-tenancy in the context of cloud platforms. It is designed to be scalable and allows a unique shared controller platform to be used for managing multiple domains in a cloud environment. Each tenant has full control over its virtual networks and is free to deploy any network abstraction and application on top of the controller platform.

The compositional SDN hypervisor [181] was designed with a different set of goals. Its main objective is to allow the cooperative (sequential or parallel) execution of applications developed with different programming languages or conceived for diverse control platforms. It thus offers interoperability and portability in addition to the typical functions of network hypervisors.

Commercial multi-tenant network hypervisors

None of the aforementioned approaches is designed to address all challenges of multi-tenant data centers. For instance, tenants want to be able to migrate their enterprise solutions to cloud providers without the need to modify the network configuration of their home network. Existing networking technologies and migration strategies have mostly failed to meet both the tenant and the service provider requirements. A multi-tenant environment should be anchored in a network hypervisor capable of abstracting the underlying forwarding devices and physical network topology from the tenants. Moreover, each tenant should have access to control abstractions and manage its own virtual networks independently and isolated from other tenants.

With the market demand for network virtualization and the recent research on SDN showing promise as an enabling technology, different commercial virtualization platforms based on SDN concepts have started to appear. VMWare has proposed a network virtualization platform (NVP) [112] that provides the necessary abstractions to allow the creation of independent virtual networks for large-scale multi-tenant environments. NVP is a complete network virtualization solution that allows the creation of virtual networks, each with independent service model, topologies, and addressing architectures over the same physical network. With NVP, tenants do not need to know anything about the underlying network topology, configuration or other specific aspects of the forwarding devices. NVP's network hypervisor translates the tenants configurations and requirements into low level instruction sets to be installed on

the forwarding devices. For this purpose, the platform uses a cluster of SDN controllers to manipulate the forwarding tables of the Open vSwitches in the host’s hypervisor. Forwarding decisions are therefore made exclusively on the network edge. After the decision is made, the packet is tunneled over the physical network to the receiving host hypervisor (the physical network sees nothing but ordinary IP packets).

IBM has also recently proposed SDN VE [171], [172], another commercial and enterprise-class network virtualization platform. SDN VE uses OpenDaylight as one of the building blocks of the so-called Software-Defined Environments (SDEs), a trend further discussed in Section V. This solution also offers a complete implementation framework for network virtualization. Like NVP, it uses a host-based overlay approach, achieving advanced network abstraction that enables application-level network services in large-scale multi-tenant environments. Interestingly, SDN VE 1.0 is capable of supporting in one single instantiation up to 16,000 virtual networks and 128,000 virtual machines [171], [172].

To summarize, currently there are already a few network hypervisor proposals leveraging the advances of SDN. There are, however, still several issues to be addressed. These include, among others, the improvement of virtual-to-physical mapping techniques [182], the definition of the level of detail that should be exposed at the logical level, and the support for nested virtualization [29]. We anticipate, however, this ecosystem to expand in the near future since network virtualization will most likely play a key role in future virtualized environments, similarly to the expansion we have been witnessing in virtualized computing.

D. Layer IV: Network Operating Systems / Controllers

Traditional operating systems provide abstractions (e.g., high-level programming APIs) for accessing lower-level devices, manage the concurrent access to the underlying resources (e.g., hard drive, network adapter, CPU, memory), and provide security protection mechanisms. These functionalities and resources are key enablers for increased productivity, making the life of system and application developers easier. Their widespread use has significantly contributed to the evolution of various ecosystems (e.g., programming languages) and the development of a myriad of applications.

In contrast, networks have so far been managed and configured using lower level, device-specific instruction sets and mostly closed proprietary network operating systems (e.g., Cisco IOS and Juniper JunOS). Moreover, the idea of operating systems abstracting device-specific characteristics and providing, in a transparent way, common functionalities is still mostly absent in networks. For instance, nowadays designers of routing protocols need to deal with complicated distributed algorithms when solving networking problems. Network practitioners have therefore been solving the same problems over and over again.

SDN is promised to facilitate network management and ease the burden of solving networking problems by means of the logically-centralized control offered by a network operating system (NOS) [26]. As with traditional operating systems,

the crucial value of a NOS is to provide abstractions, essential services, and common application programming interfaces (APIs) to developers. Generic functionality as network state and network topology information, device discovery, and distribution of network configuration can be provided as services of the NOS. With NOSs, to define network policies a developer no longer needs to care about the low-level details of data distribution among routing elements, for instance. Such systems can arguably create a new environment capable of fostering innovation at a faster pace by reducing the inherent complexity of creating new network protocols and network applications.

A NOS (or controller) is a critical element in an SDN architecture as it is the key supporting piece for the control logic (applications) to generate the network configuration based on the policies defined by the network operator. Similar to a traditional operating system, the control platform abstracts the lower-level details of connecting and interacting with forwarding devices (i.e., of materializing the network policies).

Architecture and design axes

There is a very diverse set of controllers and control platforms with different design and architectural choices [7], [13], [183], [184], [185], [186]. Existing controllers can be categorized based on many aspects. From an architectural point of view, one of the most relevant is if they are centralized or distributed. This is one of the key design axes of SDN control platforms, so we start by discussing this aspect next.

Centralized vs. Distributed

A centralized controller is a single entity that manages all forwarding devices of the network. Naturally, it represents a single point of failure and may have scaling limitations. A single controller may not be enough to manage a network with a large number of data plane elements. Centralized controllers such as NOX-MT [187], Maestro [188], Beacon [186], and Floodlight [189] have been designed as highly concurrent systems, to achieve the throughput required by enterprise class networks and data centers. These controllers are based on multi-threaded designs to explore the parallelism of multi-core computer architectures. As an example, Beacon can deal with more than 12 million flows per second by using large size computing nodes of cloud providers such as Amazon [186]. Other centralized controllers such as Trema [190], Ryu NOS [191], Meridian [192], and ProgrammableFlow [193], [133] target specific environments such as data centers, cloud infrastructures, and carrier grade networks. Furthermore, controllers such as Rosemary [194] offer specific functionality and guarantees, namely security and isolation of applications. By using a container-based architecture called micro-NOS, it achieves its primary goal of isolating applications and preventing the propagation of failures throughout the SDN stack.

Contrary to a centralized design, a distributed network operating system can be scaled up to meet the requirements of potentially any environment, from small to large-scale networks. A distributed controller can be a centralized cluster of nodes or a physically distributed set of elements. While the first alternative can offer high throughput for very dense

data centers, the latter can be more resilient to different kinds of logical and physical failures. A cloud provider that spans multiple data centers interconnected by a wide area network may require a hybrid approach, with clusters of controllers inside each data center and distributed controller nodes in the different sites [8].

Onix [7], HyperFlow [195], HP VAN SDN [184], ONOS [117], DISCO [185], *yanc* [196], PANE [197], SMaRt-Light [198], and Fleet [199] are examples of distributed controllers. Most distributed controllers offer weak consistency semantics, which means that data updates on distinct nodes will *eventually* be updated on all controller nodes. This implies that there is a period of time in which distinct nodes may read different values (old value or new value) for a same property. Strong consistency, on the other hand, ensures that all controller nodes will read the most updated property value after a write operation. Despite its impact on system performance, strong consistency offers a simpler interface to application developers. To date, only Onix, ONOS, and SMaRtLight provide this data consistency model.

Another common property of distributed controllers is fault tolerance. When one node fails, another neighbor node should take over the duties and devices of the failed node. So far, despite some controllers tolerating crash failures, they do not tolerate arbitrary failures, which means that any node with an abnormal behavior will not be replaced by a potentially well behaved one.

A single controller may be enough to manage a small network, however it represents a single point of failure. Similarly, independent controllers can be spread across the network, each of them managing a network segment, reducing the impact of a single controller failure. Yet, if the control plane availability is critical, a cluster of controllers can be used to achieve a higher degree of availability and/or for supporting more devices. Ultimately, a distributed controller can improve the control plane resilience, scalability and reduce the impact of problems caused by network partition, for instance. SDN resiliency as a whole is an open challenge that will be further discussed in Section V-C.

Dissecting SDN Controller Platforms

To provide a better architectural overview and understanding the design a network operating system, Table V summarizes some of the most relevant architectural and design properties of SDN controllers and control platforms. We have focused on the elements, services and interfaces of a selection of production-level, well-documented controllers and control platforms. Each line in the table represent a component we consider important in a modular and scalable control platform. We observe a highly diversified environment, with different properties and components being used by distinct control platforms. This is not surprising, given an environment with many competitors willing to be at the forefront of SDN development. Note also that not all components are available on all platforms. For instance, east/westbound APIs are not required in centralized controllers such as Beacon. In fact, some platforms have very specific niche markets, such as telecom companies and cloud providers, so the requirements

will be different.

Based on the analysis of the different SDN controllers proposed to date (both those presented in Table V and others, such as NOX [26], Meridian [192], ForCES [30], and FortNOX [201]), we extract several common elements and provide a first attempt to clearly and systematically dissect an SDN control platform in Figure 8.

There are at least three relatively well-defined layers in most of the existing control platforms: (i) the application, orchestration and services; (ii) the core controller functions, and (iii) the elements for southbound communications. The connection at the upper-level layers is based on northbound interfaces such as REST APIs [202] and programming languages such as FML [203], Frenetic [204] and NetCore [205]. On the lower-level part of a control platform, southbound APIs and protocol plugins interface the forwarding elements. The core of a controller platform can be characterized as a combination of its base network service functions and the various interfaces.

Core controller functions

The base network service functions are what we consider the essential functionality all controllers should provide. As an analogy, these functions are like base services of operating systems, such as program execution, I/O operations control, communications, protection, and so on. These services are used by other operating system level services and user applications. In a similar way, functions such as topology, statistics, notifications and device management, together with shortest path forwarding and security mechanisms are essential network control functionalities that network applications may use in building its logic. For instance, the notification manager should be able to receive, process, and forward events (e.g., alarm notifications, security alarms, state changes) [55]. Security mechanisms are another example, as they are critical components to provide basic isolation and security enforcement between services and applications. For instance, rules generated by high priority services should not be overwritten with rules created by applications with a lower priority.

Southbound

On the lower-level of control platforms, the southbound APIs can be seen as a layer of device drivers. They provide a common interface for the upper layers, while allowing a control platform to use different southbound APIs (e.g., OpenFlow, OVSDB, ForCES) and protocol plugins to manage existing or new physical or virtual devices (e.g., SNMP, BGP, NetConf). This is essential both for backward compatibility and heterogeneity, i.e., to allow multiple protocols and device management connectors. Therefore, on the data plane, a mix of physical devices, virtual devices (e.g., Open vSwitch [142], [109], vRouter [206]) and a variety of device interfaces (e.g., OpenFlow, OVSDB, of-config [207], NetConf, and SNMP) can co-exist.

Most controllers support only OpenFlow as a southbound API. Still, a few of them, such as OpenDaylight, Onix and HP VAN SDN Controller, offer a wider range of southbound APIs and/or protocol plugins. Onix supports both the OpenFlow and OVSDB protocols. The HP VAN SDN Controller has other southbound connectors such as L2 and L3 agents.

TABLE V
ARCHITECTURE AND DESIGN ELEMENTS OF CONTROL PLATFORMS

Component	OpenDaylight	OpenContrail	HP VAN SDN	Onix	Beacon
Base network services	Topology/Stats/Switch Manager, Host Tracker, Shortest Path Forwarding	Routing, Tenant Isolation	Audit Log, Alerts, Topology, Discovery	Discovery, Multi-consistency Storage, Read State, Register for updates	Topology, device manager, and routing
East/Westbound APIs	—	Control Node (XMPP-like control channel)	Sync API	Distribution I/O module	<i>Not present</i>
Integration Plug-ins	OpenStack Neutron	CloudStack, OpenStack	OpenStack	—	—
Management Interfaces	GUI/CLI, REST API	GUI/CLI	REST API Shell / GUI Shell	—	Web
Northbound APIs	REST, REST-CONF [200], Java APIs	REST APIs (configuration, operational, and analytic)	REST API, GUI Shell	Onix API (general purpose)	API (based on OpenFlow events)
Service abstraction layers	Service Abstraction Layer (SAL)	—	Device Abstraction API	Network Information Base (NIB) Graph with Import/Export Functions	—
Southbound APIs or connectors	OpenFlow, OVSDB, SNMP, PCEP, BGP, NETCONF	—	OpenFlow, L3 Agent, L2 Agent	OpenFlow, OVSDB	OpenFlow




Fig. 8. SDN control platforms: elements, services and interfaces

OpenDaylight goes a step beyond by providing a Service Layer Abstraction (SLA) that allows several southbound APIs and protocols to co-exist in the control platform. For instance, its original architecture was designed to support at least seven different protocols and plugins: OpenFlow, OVSDB [153], NETCONF [44], PCEP [43], SNMP [208], BGP [42] and LISP Flow Mapping [13]. Hence, OpenDaylight is one of the few control platforms being conceived to support a broader integration of technologies in a single control platform.

Eastbound and Westbound

East/westbound APIs, as illustrated in Figure 9, are a special case of interfaces required by distributed controllers. Currently,

each controller implements its own east/westbound API. The functions of these interfaces include import/export data between controllers, algorithms for data consistency models, and monitoring/notification capabilities (e.g., check if a controller is up or notify a take over on a set of forwarding devices).

Similarly to southbound and northbound interfaces, east/westbound APIs are essential components of distributed controllers. To identify and provide common compatibility and interoperability between different controllers, it is necessary to have standard east/westbound interfaces. For instance, SDNi [209] defines common requirements to coordinate flow setup and exchange reachability information across multiple




Fig. 9. Distributed controllers: east/westbound APIs.

domains. In essence, such protocols can be used in an orchestrated and interoperable way to create more scalable and dependable distributed control platforms. Interoperability can be leveraged to increase the diversity of the control platform element. Indeed, diversity increases the system robustness by reducing the probability of common faults, such as software faults [210].

Other proposals that try to define interfaces between controllers include Onix data import/export functions [7], ForCES CE-CE interface [30], [211], ForCES Intra-NE cold-standby mechanisms for high availability [212], and distributed data stores [213]. An east/westbound API requires advanced data distribution mechanisms such as the Advanced Message Queuing Protocol (AMQP) [214] used by DISCO [185], techniques for distributed concurrent and consistent policy composition [215], transactional databases and DHTs [216] (as used in Onix [7]), or advanced algorithms for strong consistency and fault tolerance [213], [198].

In a multi-domain setup, east/westbound APIs may require also more specific communication protocols between SDN domain controllers [217]. Some of the essential functions of such protocols are to *coordinate flow setup* originated by applications, *exchange reachability information* to facilitate inter-SDN routing, *reachability update* to keep the network state consistent, among others.

Another important issue regarding east/westbound interfaces is heterogeneity. For instance, besides communicating with peer SDN controllers, controllers may also need to communicate with subordinate controllers (in a hierarchy of controllers) and non-SDN controllers [218], as is the case of Closed-Flow [219]. To be interoperable, east/westbound interfaces thus need to accommodate different controller interfaces, with their specific set of services, and the diverse characteristics of the underlying infrastructure, including the diversity of technology, the geographic span and scale of the network, and the distinction between WAN and LAN – potentially across administrative boundaries. In those cases, different information has to be exchanged between controllers, including adjacency and capability discovery, topology information (to the extent of the agreed contracts between administrative domains), billing information, among many others [218].

Lastly, a “SDN compass” methodology [220] suggests a finer distinction between eastbound and westbound horizontal interfaces, referring to westbound interfaces as SDN-to-SDN protocols and controller APIs while eastbound interfaces would be used to refer to standard protocols used to commu-

nicate with legacy network control planes (e.g., PCEP [43], GMPLS [221]).

Northbound

Current controllers offer a quite broad variety of northbound APIs, such as ad-hoc APIs, RESTful APIs [202], multi-level programming interfaces, file systems, among other more specialized APIs such as NVP NBAPI [7], [112] and SDMN API [222]. Section IV-E is devoted to a more detailed discussion on the evolving layer of northbound APIs. A second kind of northbound interfaces are those stemming out of SDN programming languages such as Frenetic [204], Nettle [223], NetCore [205], Procera [224], Pyretic [225], NetKAT [226] and other query-based languages [227]. Section IV-G gives a more detailed overview of the several existing programming languages for SDN.

Wrapping up remarks and platforms comparison

Table VI shows a summary of some of the existing controllers with their respective architectures and characteristics. As can be observed, most controllers are centralized and multi-threaded. Curiously, the northbound API is very diverse. In particular, five controllers (Onix, Floodlight, MuL, Meridian and SDN Unified Controller) pay a bit more attention to this interface, as a statement of its importance. Consistency models and fault tolerance are only present in Onix, HyperFlow, HP VAN SDN, ONOS and SMaRtLight. Lastly, when it comes to the OpenFlow standard as southbound API, only Ryu supports its three major versions (v1.0, v1.2 and v1.3).

To conclude, it is important to emphasize that the control platform is one of the critical points for the success of SDN [233]. One of the main issues that needs to be addressed in this respect is interoperability. This is rather interesting, as it was the very first problem that southbound APIs (such as OpenFlow) tried to solve. For instance, while WiFi and LTE (Long-Term Evolution) networks [234] need specialized control platforms such as MobileFlow [222] or SoftRAN [235], data center networks have different requirements that can be met with platforms such as Onix [7] or OpenDaylight [13]. For this reason, in environments where diversity of networking infrastructures is a reality, coordination and cooperation between different controllers is crucial. Standardized APIs for multi-controller and multi-domain deployments are therefore seen as an important step to achieve this goal.

E. Layer V: Northbound Interfaces

The North- and Southbound interfaces are two key abstractions of the SDN ecosystem. The southbound interface has already a widely accepted proposal (OpenFlow), but a common northbound interface is still an open issue. At this moment it may still be a bit too early to define a standard northbound interface, as use-cases are still being worked out [236]. Anyway, it is to be expected a common (or a *de facto*) northbound interface to arise as SDN evolves. An abstraction that would allow network applications not to depend on specific implementations is important to explore the full potential of SDN

The northbound interface is mostly a software ecosystem, not a hardware one as is the case of the southbound APIs.

TABLE VI
CONTROLLERS CLASSIFICATION

Name	Architecture	Northbound API	Consistency	Faults	License	Prog. language	Version
Beacon [186]	centralized multi-threaded	ad-hoc API	no	no	GPLv2	Java	v1.0
DISCO [185]	distributed	REST	—	yes	—	Java	v1.1
Fleet [199]	distributed	ad-hoc	no	no	—	—	v1.0
Floodlight [189]	centralized multi-threaded	RESTful API	no	no	Apache	Java	v1.1
HP VAN SDN [184]	distributed	RESTful API	weak	yes	—	Java	v1.0
HyperFlow [195]	distributed	—	weak	yes	—	C++	v1.0
Kandoo [228]	hierarchically distributed	—	no	no	—	C, C++, Python	v1.0
Onix [7]	distributed	NVP NBAPI	weak, strong	yes	commercial	Python, C	v1.0
Maestro [188]	centralized multi-threaded	ad-hoc API	no	no	LGPLv2.1	Java	v1.0
Meridian [192]	centralized multi-threaded	extensible API layer	no	no	—	Java	v1.0
MobileFlow [222]	—	SDMN API	—	—	—	—	v1.2
MuL [229]	centralized multi-threaded	multi-level interface	no	no	GPLv2	C	v1.0
NOX [26]	centralized	ad-hoc API	no	no	GPLv3	C++	v1.0
NOX-MT [187]	centralized multi-threaded	ad-hoc API	no	no	GPLv3	C++	v1.0
NVP Controller [112]	distributed	—	—	—	commercial	—	—
OpenContrail [183]	—	REST API	no	no	Apache 2.0	Python, C++, Java	v1.0
OpenDaylight [13]	distributed	REST, RESTCONF	weak	no	EPL v1.0	Java	v1.{0,3}
ONOS [117]	distributed	RESTful API	weak, strong	yes	—	Java	v1.0
PANE [197]	distributed	PANE API	yes	—	—	—	—
POX [230]	centralized	ad-hoc API	no	no	GPLv3	Python	v1.0
ProgrammableFlow [231]	centralized	—	—	—	—	C	v1.3
Rosemary [194]	centralized	ad-hoc	—	—	—	—	v1.0
Ryu NOS [191]	centralized multi-threaded	ad-hoc API	no	no	Apache 2.0	Python	v1.{0,2,3}
SMArtLight [198]	distributed	RESTful API	no	no	Apache	Java	v1.0
SNAC [232]	centralized	ad-hoc API	no	no	GPL	C++	v1.0
Trema [190]	centralized multi-threaded	ad-hoc API	no	no	GPLv2	C, Ruby	v1.0
<i>Unified Controller</i> [171]	—	REST API	—	—	commercial	—	v1.0
<i>yanc</i> [196]	distributed	file system	—	—	—	—	—

In these ecosystems, the implementation is commonly the forefront driver, while standards emerge later and are essentially driven by wide adoption [237]. Nevertheless, an initial and minimal standard for northbound interfaces can still play an important role for the future of SDN. Discussions about this issue have already begun [236], [237], [238], [239], [240], [241], [242], [243], and a common consensus is that northbound APIs are indeed important but that it is indeed too early to define a single standard right now. The experience from the development of different controllers will certainly be the basis for coming up with a common application level interface.

Open and standard northbound interfaces are crucial to promote application portability and interoperability among the different the control platforms. A northbound API can be compared to the POSIX standard [244] in operating systems, representing an abstraction that guarantees programming language and controller independence. NOSIX [245] is one of the first examples of an effort in this direction. It tries to define portable low-level (e.g., flow model) application interfaces, making southbound APIs such as OpenFlow look like “device

drivers”. However, NOSIX is not exactly a general purpose northbound interface, but rather a higher-level abstraction for southbound interfaces. Indeed, it could be part of the common abstraction layer in a control platform as the one described in Section IV-D.

Existing controllers such as Floodlight, Trema, NOX, Onix, and OpenDaylight propose and define their own northbound APIs [238], [246]. However, each of them has its own specific definitions. Programming languages such as Frenetic [204], Nettle [223], NetCore [205], Procera [224], Pyretic [247] and NetKAT [226] also abstract the inner details of the controller functions and data plane behavior from the application developers. Moreover, as we explain in Section IV-G, programming languages can provide a wide range of powerful abstractions and mechanisms such as application composition, transparent data plane fault tolerance, and a variety of basic building blocks to ease software module and application development.

SFNet [248] is another example of a northbound interface. It is a high-level API that translates application requirements into lower level service requests. However, SFNet has a limited scope, targeting queries to request the congestion state of

the network and services such as bandwidth reservation and multicast.

Other proposals use different approaches to allow applications to interact with controllers. The *yanc* control platform [196] explores this idea by proposing a general control platform based on Linux and abstractions such as the virtual file system (VFS). This approach simplifies the development of SDN applications as programmers are able to use a traditional concept (files) to communicate with lower level devices and sub-systems.

Eventually, it is unlikely that a single northbound interface emerges as the winner, as the requirements for different network applications are quite different. APIs for security applications are likely to be different from those for routing or financial applications. One possible path of evolution for northbound APIs are vertically-oriented proposals, before any type of standardization occurs, a challenge the ONF has started to undertake in the NBI WG in parallel to open-source SDN developments [50]. The ONF architectural work [218] includes the possibility of northbound APIs providing resources to enable dynamic and granular control of the network resources from customer applications, eventually across different business and organizational boundaries.

There are also other kind of APIs, such as those provided by the PANE controller [197]. Designed to be suitable for the concept of participatory networking, PANE allows network administrators to define module-specific quotas and access control policies on network resources. The controller provides an API that allows end-host applications to dynamically and autonomously request network resources. For example, audio (e.g., VoIP) and video applications can easily be modified to use the PANE API to reserve bandwidth for certain quality guarantees during the communication session. PANE includes a compiler and verification engine to ensure that bandwidth requests do not exceed the limits set by the administrator and to avoid starvation, i.e., other applications shall not be impaired by new resource requests.

F. Layer VI: Language-based Virtualization

Two essential characteristics of virtualization solutions are the capability of expressing modularity and of allowing different levels of abstractions while still guaranteeing desired properties such as protection. For instance, virtualization techniques can allow different views of a single physical infrastructure. As an example, one virtual “big switch” could represent a combination of several underlying forwarding devices. This intrinsically simplifies the task of application developers as they do not need to think about the sequence of switches where forwarding rules have to be installed, but rather see the network as a simple “big switch”. Such kind of abstraction significantly simplify the development and deployment of complex network applications, such as advanced security related services.

Pyretic [247] is an interesting example of a programming language that offers this type of high-level abstraction of network topology. It incorporates this concept of abstraction by introducing network objects. These objects consist of an

abstract network topology and the sets of policies applied to it. Network objects simultaneously hide information and offer the required services.

Another form of language-based virtualization is static slicing. This a scheme where the network is sliced by a compiler, based on application layer definitions. The output of the compiler is a monolithic control program that has already slicing definitions and configuration commands for the network. In such a case, there is no need for a hypervisor to dynamically manage the network slices. Static slicing can be valuable for deployments with specific requirements, in particular those where higher performance and simple isolation guarantees are preferable to dynamic slicing.

One example of static slicing approach is the Splendid isolation [249]. In this solution the network slices are made of 3 components: (a) *topology*, consisting of switches, ports, and links; (b) *mapping* of slice-level switches, ports and links on the network infrastructure; (c) *predicates on packets*, where each port of the slice’s edge switches has an associated predicate. The topology is a simple graph of the sliced nodes, ports and links. Mapping will translate the abstract topology elements into the corresponding physical ones. The predicates are used to indicate whether a packet is permitted or not to enter a specific slice. Different applications can be associated to each slice. The compiler takes the combination of slices (topology, mapping, and predicates) and respective programs to generate a global configuration for the entire network. It also ensures that properties such as isolation are enforced among slices, i.e., no packets of a slice A can traverse to a slice B unless explicitly allowed.

Other solutions, such as libNetVirt [250], try to integrate heterogeneous technologies for creating static network slices. libNetVirt is a library designed to provide a flexible way to create and manage virtual networks in different computing environments. Its main idea is similar to the OpenStack Quantum project [251]. While Quantum is designed for OpenStack (cloud environments), libNetVirt is a more general purpose library which can be used in different environments. Additionally, it goes one step beyond OpenStack Quantum by enabling QoS capabilities in virtual networks [250]. The libNetVirt library has two layers: (1) a generic network interface; and (2) technology specific device drivers (e.g., VPN, MPLS, OpenFlow). On top of the layers are the network applications and virtual network descriptions. The OpenFlow driver uses a NOX controller to manage the underlying infrastructure, using OpenFlow rule-based flow tables to create isolated virtual networks. By supporting different technologies, it can be used as a bridging component in heterogeneous networks.

Table VII summarizes the hypervisor and non-hypervisor based virtualization technologies. As can be observed, only libNetVirt supports heterogeneous technologies, not restricting its application to OpenFlow-enabled networks. FlowVisor, AutoSlice and OpenVirteX allow multiple controllers, one per network slice. FlowN provides a container-based approach where multiple applications from different users can co-exist on a single controller. FlowVisor allows QoS provisioning guarantees by using VLAN PCP bits for priority queues. SDN VE and NVP also provide their own provisioning methods for

TABLE VII
VIRTUALIZATION SOLUTIONS

Solution	Multiple controllers	Slicing	QoS “guarantees”	Multi-technology
AutoVFlow [174]	yes, one per tenant	flow space virtualization	no	no, OF only
AutoSlice [179]	yes, one per slice	VLAN tags	no	no, OF only
Compositional Hypervisor [181]	—	—	no	no, OF only
FlowVisor [166], [167]	yes, one per slice	virtual flow tables per slice	yes (VLAN PCP bits)	no, OF only
FlowN [168], [180]	no (contained applications)	VLAN tags	no	no, OF only
IBM SDN VE [171]	yes, a cluster of controllers	logical datapaths	yes (priority-based)	yes (VXLAN, OVS, OF)
libNetVirt [250]	no, one single controller	VLAN tags	no	yes (e.g., VPN, MPLS, OF)
NVP’s Hypervisor [112]	yes, a cluster of controller	logical datapaths	yes	no, OVS only
OpenVirteX [170]	yes, one per slice	virtual flow tables per slice	<i>unknown</i>	no, OF only
Pyretic [247]	no, one single controller	compiler time OF rules	no	no, OF only
Splendid Isolation [249]	no, one single controller	compiler time VLANs	no	no, OF only
RadioVisor [173]	yes, one per slice	3D resource grid slices	—	Femto API [252]
xDPd virtualization [176]	yes, one per slice	flow space virtualization	no	no, OF only

guaranteeing QoS.

G. Layer VII: Programming languages

Programming languages have been proliferating for decades. Both academia and industry have evolved from low-level hardware-specific machine languages, such as assembly for x86 architectures, to high-level and powerful programming languages such as Java and Python. The advancements towards more portable and reusable code has driven a significant shift on the computer industry [253], [254].

Similarly, programmability in networks is starting to move from low level machine languages such as OpenFlow (“assembly”) to high-level programming languages [204], [203], [223], [205], [224], [225], [112]. Assembly-like machine languages, such as OpenFlow [9] and POF [31], [120], essentially mimic the behavior of forwarding devices, forcing developers to spend too much time on low-level details rather than on the problem solve. Raw OpenFlow programs have to deal with hardware behavior details such as overlapping rules, the priority ordering of rules, and data-plane inconsistencies that arise from in-flight packets whose flow rules are under installation [204], [205], [255]. The use of these low-level languages makes it difficult to reuse software, to create modular and extensive code, and leads to a more error-prone development process [225], [256], [257].

Abstractions provided by high level programming languages can significantly help address many of the challenges of these lower-level instruction sets [204], [203], [223], [205], [224], [225]. In SDNs, high-level programming languages can be designed and used to:

- 1) create higher level abstractions for simplifying the task of programming forwarding devices;
- 2) enable more productive and problem-focused environments for network software programmers, speeding up development and innovation;
- 3) promote software modularization and code reusability in the network control plane;

- 4) foster the development of network virtualization.

Several challenges can be better addressed by programming languages in SDNs. For instance, in pure OpenFlow-based SDNs, it is hard to ensure that multiple tasks of a single application (e.g., routing, monitoring, access control) do not interfere with each other. For example, rules generated for one task should not override the functionality of another task [204], [255]. Another example is when multiple applications run on a single controller [225], [255], [201], [258], [259]. Typically, each application generates rules based on its own needs and policies without further knowledge about the rules generated by other applications. As a consequence, conflicting rules can be generated and installed in forwarding devices, which can create problems for network operation. Programming languages and runtime systems can help to solve these problems that would be otherwise hard to prevent.

Important software design techniques such as code modularity and reusability are very hard to achieve using low-level programming models [225]. Applications thus built are monolithic and consist of building blocks that can not be reused in other applications. The end result is a very time consuming and error prone development process.

Another interesting feature that programming language abstractions provide is the capability of creating and writing programs for virtual network topologies [247], [249]. This concept is similar to object-oriented programming, where objects abstract both data and specific functions for application developers, making it easier to focus on solving a particular problem without worrying about data structures and their management. For instance, in an SDN context, instead of generating and installing rules in each forwarding device, one can think of creating simplified virtual network topologies that represent the entire network, or a subset of it. For example, the application developer should be able to abstract the network as an atomic big switch, rather than a combination of several underlying physical devices. The programming languages or runtime systems should be responsible for generating and

installing the lower-level instructions required at each forwarding device to enforce the user policy across the network. With such kind of abstractions, developing a routing application becomes a straightforward process. Similarly, a single physical switch could be represented as a set of virtual switches, each of them belonging to a different virtual network. These two examples of abstract network topologies would be much harder to implement with low-level instruction sets. In contrast, a programming language or runtime system can more easily provide abstractions for virtual network topologies, as has already been demonstrated by languages such as Pyretic [247].

High-level SDN programming languages

High-level programming languages can be powerful tools as a mean for implementing and providing abstractions for different important properties and functions of SDN such as network-wide structures, distributed updates, modular composition, virtualization, and formal verification [29].

Low-level instruction sets suffer from several problems. To address some of these challenges, higher-level programming languages have been proposed, with diverse goals, such as:

- Avoiding low-level and device-specific configurations and dependencies spread across the network, as happens in traditional network configuration approaches;
- Providing abstractions that allow different management tasks to be accomplished through easy to understand and maintain network policies;
- Decoupling of multiple tasks (e.g., routing, access control, traffic engineering);
- Implementing higher-level programming interfaces to avoid low-level instruction sets;
- Solving forwarding rules problems, e.g., conflicting or incomplete rules that can prevent a switch event to be triggered, in an automated way;
- Addressing different race condition issues which are inherent to distributed systems;
- Enhancing conflict-resolution techniques on environments with distributed decision makers;
- Provide native fault-tolerance capabilities on data plane path setup;
- Reducing the latency in the processing of new flows;
- Easing the creation of stateful applications (e.g., stateful firewall).

Programming languages can also provide specialized abstractions to cope with other management requirements, such as monitoring [224], [204], [260], [227]. For instance, the runtime system of a programming language can do all the “laundry work” of installing rules, polling the counters, receiving the responses, combining the results as needed, and composing monitoring queries in conjunction with other policies. Consequently, application developers can take advantage of the simplicity and power of higher level query instructions to easily implement monitoring modules or applications.

Another aspect of paramount importance is the portability of the programming language, necessary so that developers do not need to re-implement applications for different control platforms. The portability of a programming language can be considered as a significant added value to the control

plane ecosystem. Mechanisms such as decoupled back-ends could be key architectural ingredients to enable platform portability. Similarly to the Java virtual machine, a portable northbound interface will easily allow applications to run on different controllers without requiring any modification. As an example, the Pyretic language requires only a standard socket interface and a simple OpenFlow client on the target controller platform [225].

Several programming languages have been proposed for SDNs, as summarized in Table VIII. The great majority propose abstractions for OpenFlow-enabled networks. The predominant programming paradigm is the declarative one, with a single exception, Pyretic, which is an imperative language. Most declarative languages are functional, while but there are instances of the logic and reactive types. The purpose – i.e., the specific problems they intend to solve – and the expressiveness power vary from language to language, while the end goal is almost always the same: to provide higher-level abstractions to facilitate the development of network control logic.

Programming languages such as FML [203], Nettle [223], and Procera [224] are functional and reactive. Policies and applications written in these languages are based on reactive actions triggered by events (e.g., a new host connected to the network, or the current network load). Such languages allow users to declaratively express different network configuration rules such as access control lists (ACLs), virtual LANs (VLANs), and many others. Rules are essentially expressed as allow-or-deny policies, which are applied to the forwarding elements to ensure the desired network behavior.

Other SDN programming languages such as Frenetic [204], Hierarchical Flow Tables (HFT) [255], NetCore [205], and Pyretic [225], were designed with the simultaneous goal of efficiently expressing packet-forwarding policies and dealing with overlapping rules of different applications, offering advanced operators for parallel and sequential composition of software modules. To avoid overlapping conflicts, Frenetic disambiguates rules with overlapping patterns by assigning different integer priorities, while HFT uses hierarchical policies with enhanced conflict-resolution operators.

See-every-packet abstractions and race-free semantics also represent interesting features provided by programming languages (such as Frenetic [204]). The former ensures that *all* control packets will be available for analysis, sooner or later, while the latter provides the mechanisms for suppressing unimportant packets. As an example, packets that arise from a network race condition, such as due to a concurrent flow rule installation on switches, can be simply discarded by the runtime system.

Advanced operators for parallel and sequential composition help bind (through internal workflow operators) the key characteristics of programming languages such as Pyretic [225]. Parallel composition makes it possible to operate multiple policies on the same set of packets, while sequential composition facilitates the definition of a sequential workflow of policies to be processed on a set of packets. Sequential policy processing allows multiple modules (e.g., access control and routing) to operate in a cooperative way. By using sequential composition complex applications can be built out of a combination of

TABLE VIII
PROGRAMMING LANGUAGES

Name	Programming paradigm	Short description/purpose
FatTire [261]	declarative (functional)	Uses regular expressions to allow programmers to describe network paths and respective fault-tolerance requirements.
Flog [257]	declarative (logic), event-driven	Combines ideas of FML and Frenetic, providing an event-driven and forward-chaining logic programming language.
FlowLog [256]	declarative (functional)	Provides a finite-state language to allow different analysis, such as model-checking.
FML [203]	declarative (dataflow, reactive)	High level policy description language (e.g., access control).
Frenetic [204]	declarative (functional)	Language designed to avoid race conditions through well defined high level programming abstractions.
HFT [255]	declarative (logic, functional)	Enables hierarchical policies description with conflict-resolution operators, well suited for decentralized decision makers.
Maple [262]	declarative (functional)	Provides a highly-efficient multi-core scheduler that can scale efficiently to controllers with 40+ cores.
Merlin [263]	declarative (logic)	Provides mechanisms for delegating management of sub-policies to tenants without violating global constraints.
nlog [112]	declarative (functional)	Provides mechanisms for data log queries over a number of tables. Produces immutable tuples for reliable detection and propagation of updates.
NetCore [205]	declarative (functional)	High level programming language that provides means for expressing packet-forwarding policies in a high level.
NetKAT [226]	declarative (functional)	It is based on Kleene algebra for reasoning about network structure and supported by solid foundation on equational theory.
Nettle [223]	declarative (functional, reactive)	Based on functional reactive programming principles in order to allow programmers to deal with streams instead of events.
Procera [224]	declarative (functional, reactive)	Incorporates a set of high level abstractions to make it easier to describe reactive and temporal behaviors.
Pyretic [225]	imperative	Specifies network policies at a high level of abstraction, offering transparent composition and topology mapping.

different modules (in a similar way as pipes can be used to build sophisticated Unix applications).

Further advanced features are provided by other SDN programming languages. FatTire [261] is an example of a declarative language that heavily relies on regular expressions to allow programmers to describe network paths with fault-tolerance requirements. For instance, each flow can have its own alternative paths for dealing with failure of the primary paths. Interestingly, this feature is provided in a very programmer-friendly way, with the application programmer having only to use regular expressions with special characters, such as an asterisk. In the particular case of FatTire, an asterisk will produce the same behavior as a traditional regular expression, but translated into alternative traversing paths.

Programming languages such as FlowLog [256] and Flog [257] bring different features, such as model checking, dynamic verification and stateful middleboxes. For instance, using a programming language such as Flog, it is possible to build a stateful firewall application with only five lines of code [257].

Merlin [263] is one of the first examples of unified framework for controlling different network components, such as forwarding devices, middleboxes, and end-hosts. An important advantage is backward-compatibility with existing systems. To achieve this goal, Merlin generates specific code for each type of component. Taking a policy definition as input, Merlin's compiler determines forwarding paths, transformation

placement, and bandwidth allocation. The compiled outputs are sets of component-specific low-level instructions to be installed in the devices. Merlin's policy language also allows operators to delegate the control of a sub-network to tenants, while ensuring isolation. This delegated control is expressed by means of policies that can be further refined by each tenant owner, allowing them to customize policies for their particular needs.

Other recent initiatives (e.g., systems programming languages [264]) target problems such as detecting anomalies to improve the security of network protocols (e.g., OpenFlow), and optimizing horizontal scalability for achieving high throughput in applications running on multicore architectures [262]. Nevertheless, there is still scope for further investigation and development on programming languages. For instance, one recent research has revealed that current policy compilers generate unnecessary (redundant) rule updates, most of which modify only the priority field [265].

Most of the value of SDN will come from the network managements applications built on top of the infrastructure. Advances in high-level programming languages are a fundamental component to the success of a prolific SDN application development ecosystem. To this end, efforts are undergoing to shape forthcoming standard interfaces (cf. [266]) and towards the realization of integrated development environments (e.g., NetIDE [267]) with the goal of fostering the development of a myriad of SDN applications. We discuss these next.

H. Layer VIII: Network Applications

Network applications can be seen as the “network brains”. They implement the control-logic that will be translated into commands to be installed in the data plane, dictating the behavior of the forwarding devices. Take a simple application as routing as an example. The logic of this application is to define the path through which packets will flow from a point A to a point B. To achieve this goal a routing application has to, based on the topology input, decide on the path to use and instruct the controller to install the respective forwarding rules in all forwarding devices on the chosen path, from A to B.

Software-defined networks can be deployed on any traditional network environment, from home and enterprise networks to data centers and Internet exchange points. Such variety of environments has led to a wide array of network applications. Existing network applications perform traditional functionality such as routing, load balancing, and security policy enforcement, but also explore novel approaches, such as reducing power consumption. Other examples include fail-over and reliability functionalities to the data plane, end-to-end QoS enforcement, network virtualization, mobility management in wireless networks, among many others. The variety of network applications, combined with real use case deployments, is expected to be one of the major forces on fostering a broad adoption of SDN [268].

Despite the wide variety of use cases, most SDN applications can be grouped in one of five categories: traffic engineering, mobility and wireless, measurement and monitoring, security and dependability and data center networking. Tables IX and X summarize several applications categorized as such, stating their main purpose, controller where it was implemented/evaluated, and southbound API used.

Traffic engineering

Several traffic engineering applications have been proposed, including ElasticTree [273], Hedera [275], OpenFlow-based server load balancing [336], Plug-n-Serve [284] and Aster*x [272], In-packet Bloom filter [276], SIMPLE [290], QNOX [285], QoS framework [287], QoS for SDN [286], ALTO [269], ViAggre SDN [291], ProCel [281], FlowQoS [274], and Middlepipes [27]. In addition to these, recent proposals include optimization of rules placement [337], the use of MAC as an universal label for efficient routing in data centers [338], among other techniques for flow management, fault tolerance, topology update, and traffic characterization [339]. The main goal of most applications is to engineer traffic with the aim of minimizing power consumption, maximizing aggregate network utilization, providing optimized load balancing, and other generic traffic optimization techniques.

Load balancing was one of the first applications envisioned for SDN/OpenFlow. Different algorithms and techniques have been proposed for this purpose [336], [272], [284]. One particular concern is the scalability of these solutions. A technique to allow this type of applications to scale is to use wildcard-based rules to perform proactive load balancing [336]. Wildcards can be utilized for aggregating clients requests based on the ranges of IP prefixes, for instance, allowing the distribution and directing of large groups of client requests without re-

quiring controller intervention for every new flow. In tandem, operation in reactive mode may still be used when traffic bursts are detected. The controller application needs to monitor the network traffic and use some sort of threshold in the flow counters to redistribute clients among the servers when bottlenecks are likely to happen.

SDN load-balancing also simplifies the placement of network services in the network [284]. Every time a new server is installed, the load-balancing service can take the appropriate actions to seamlessly distribute the traffic among the available servers, taking into consideration both the network load and the available computing capacity of the respective servers. This simplifies network management and provides more flexibility to network operators.

Existing southbound interfaces can be used for actively monitoring the data plane load. This information can be leveraged to optimize the energy consumption of the network [273]. By using specialized optimization algorithms and diversified configuration options, it is possible to meet the infrastructure goals of latency, performance, and fault tolerance, for instance, while reducing power consumption. With the use of simple techniques, such as shutting down links and devices intelligently in response to traffic load dynamics, data center operators can save up to 50% of the network energy in normal traffic conditions [273].

One of the important goals of data center networks is to avoid or mitigate the effect of network bottlenecks on the operation of the computing services offered. Linear bisection bandwidth is a technique that can be adopted for traffic patterns that stress the network by exploring path diversity in a data center topology. Such technique has been proposed in an SDN setting, allowing the maximization of aggregated network utilization with minimal scheduling overhead [275].

SDN can also be used to provide a fully automated system for controlling the configuration of routers. This can be particularly useful in scenarios that apply virtual aggregation [340]. This technique allows network operators to reduce the data replicated on routing tables, which is one of the causes of routing tables’ growth [341]. A specialized routing application [291] can calculate, divide and configure the routing tables of the different routing devices through a southbound API such as OpenFlow.

Traffic optimization is another interesting application for large scale service providers, where dynamic scale-out is required. For instance, the dynamic and scalable provisioning of VPNs in cloud infrastructures, using protocols such as ALTO [271], can be simplified through an SDN-based approach [269]. Recent work has also shown that optimizing rules placement can increase network efficiency [337]. Solutions such as ProCel [281], designed for cellular core networks, are capable of reducing the signaling traffic up to 70%, which represents a significant achievement.

Other applications that perform routing and traffic engineering include application-aware networking for video and data streaming [342], [343] and improved QoS by employing multiple packet schedulers [288] and other techniques [287], [285], [278], [344]. As traffic engineering is a crucial issue in all kinds of networks, upcoming methods, techniques and

TABLE IX
NETWORK APPLICATIONS

Group	Solution/Application	Main purpose	Controller	Southbound API
Traffic engineering	ALTO VPN [269]	on-demand VPNs	NMS [270], [271]	SNMP
	Aster*x [272]	load balancing	NOX	OpenFlow
	ElasticTree [273]	energy aware routing	NOX	OpenFlow
	FlowQoS [274]	QoS for broadband access networks	POX	OpenFlow
	Hedera [275]	scheduling / optimization	—	OpenFlow
	In-packet Bloom filter [276]	load balancing	NOX	OpenFlow
	MicroTE [277]	traffic engineering with minimal overhead	NOX	OpenFlow
	Middlepipes [27]	Middleboxes as a PaaS	middlepipe controller	—
	OpenQoS [278]	dynamic QoS routing for multimedia apps	Floodlight	OpenFlow
	OSP [279]	fast recovery through fast-failover groups	NOX	OpenFlow
	PolicyCop [280]	QoS policy management framework	Floodlight	OpenFlow
	ProCel [281]	Efficient traffic handling for software EPC	ProCel controller	—
	Pronto [282], [283]	Efficient queries on distributed data stores	Beacon	OpenFlow
	Plug-n-Serve [284]	load balancing	NOX	OpenFlow
	QNOX [285]	QoS enforcement	NOX	Generalized OpenFlow
	QoS for SDN [286]	QoS over heterogeneous networks	Floodlight	OpenFlow
	QoS framework [287]	QoS enforcement	NOX	OF with QoS extensions
	QoSFlow [288]	multiple packet schedulers to improve QoS	—	OpenFlow
	QueuePusher [289]	Queue management for QoS enforcement	Floodlight	OpenFlow
SIMPLE [290]	middlebox-specific “traffic steering”	Extended POX	OpenFlow	
ViAggre SDN [291]	divide and spread forwarding tables	NOX	OpenFlow	
Mobility & Wireless	AeroFlux [292], [293]	scalable hierarchical WLAN control plane	Floodlight	OpenFlow, Odin
	CROWD [294]	overlapping of LTE and WLAN cells	—	OpenFlow
	CloudMAC [295]	outsourced processing of WLAN MACs	—	OpenFlow
	C-RAN [296]	RAN [235] virtualization for mobile nets	—	—
	FAMS [297]	flexible VLAN system based on OpenFlow	ProgrammableFlow	OpenFlow
	MobileFlow [222]	flow-based model for mobile networks	MobileFlow	SDMN API
	Odin [298]	programmable virtualized WLANs	Floodlight	OpenFlow, Odin
	OpenRAN [299]	vertical programmability and virtualization	—	—
	OpenRoads [300]	control of the data path using OpenFlow	FlowVisor	OpenFlow
	SoftRAN [235]	load balancing and interference management	—	Femto API [252], [301]
Measurement & Monitoring	BISmark [6]	active and passive measurements	Procera framework	OpenFlow
	DCM [302]	distributed and coactive traffic monitoring	DCM controller	OpenFlow
	FleXam [303]	flexible sampling extension for OpenFlow	—	—
	FlowSense [304]	measure link utilization in OF networks	—	OpenFlow
	measurement model [305]	model for OF switch measurement tasks	—	OpenFlow
	OpenNetMon [306]	monitoring of QoS parameters to improve TE	POX	OpenFlow
	OpenSample [307]	low-latency sampling-based measurements	Floodlight	modified sFlow [308]
	OpenSketch [309]	separated measurement data plane	OpenSketch	“OpenSketch sketches”
	OpenTM [260]	traffic matrix estimation tool	NOX	OpenFlow
	PaFloMon [310]	passive monitoring tools defined by users	FlowVisor	OpenFlow
PayLess [311]	query-based real-time monitoring framework	Floodlight	OpenFlow	
Data Center Networking	Big Data Apps [312]	optimize network utilization	—	OpenFlow
	CloudNaas [313]	networking primitives for cloud applications	NOX	OpenFlow
	FlowComb [314]	predicts application workloads	NOX	OpenFlow
	FlowDiff [315]	detects operational problems	FlowVisor	OpenFlow
	LIME [316]	live network migration	Floodlight	OpenFlow
	NetGraph [317]	graph queries for network management	—	OpenFlow, SNMP
	OpenTCP [318]	dynamic and programmable TCP adaptation	—	—

TABLE X
NETWORK APPLICATIONS

Group	Solution/Application	Main purpose	Controller	Southbound API
Security & Dependability	Active security [319]	integrated security using feedback control	Floodlight	OpenFlow
	AVANT-GUARD [320]	DoS security specific extensions to OF	POX	OpenFlow
	CloudWatcher [321]	framework for monitoring clouds	NOX	OpenFlow
	Cognition [322]	cognitive functions to enhanced security mechanisms in network applications	—	—
	DDoS detection [323]	attacks detection and mitigation	NOX	OpenFlow
	Elastic IP & Security [324]	SDN-based elastic IP and security groups	NOX	OpenFlow
	Ethane [101]	flow-rule enforcement (match/action)	Ethane controller	first instance of OpenFlow
	FlowNAC [325]	flow-based network access control	NOX	OpenFlow
	FortNOX [201]	security flow rules prioritization	NOX	OpenFlow
	FRESCO [258]	framework for security services composition	NOX	OpenFlow
	LiveSec [326]	security policy enforcement	NOX	OpenFlow
	MAPPER [327]	fine-grained access control	—	—
	NetFuse [328]	protection against OF traffic overload	—	OpenFlow
	OF-RHM [329]	random host mutation (defense)	NOX	OpenFlow
	OpenSAFE [330]	direct spanned net traffic in arbitrary ways	NOX	OpenFlow
	OrchSec [331]	architecture for developing security apps	any	Flow-RT [332], OpenFlow
	Reliable multicasting [333]	reduce packet loss when failures occur	Trema	OpenFlow
	SANE [100]	security policy enforcement	SANE controller	SANE header (pre-OF)
	SDN RTBH [334]	DoS attack mitigation	POX	OpenFlow
VAVE [335]	source address validation with a global view	NOX	OpenFlow	

innovations can be expected in the context of SDNs.

Mobility & wireless

The current distributed control plane of wireless networks is suboptimal for managing the limited spectrum, allocating radio resources, implementing handover mechanisms, managing interference, and performing efficient load-balancing between cells. SDN-based approaches represent an opportunity for making it easier to deploy and manage different types of wireless networks, such as WLANs and cellular networks [298], [300], [294], [235], [345], [346]. Traditionally hard-to-implement but desired features are indeed becoming a reality with the SDN-based wireless networks. These include seamless mobility through efficient hand-overs [298], [347], [345], load balancing [298], [235], creation of on-demand virtual access points (VAPs) [298], [295], downlink scheduling (e.g., an OpenFlow switch can do a rate shaping or time division) [295], dynamic spectrum usage [295], enhanced inter-cell interference coordination [295], [345], device to device offloading (i.e., decide when and how LTE transmissions should be offloaded to users adopting the D2D paradigm [348]) [294], per client and/or base station resource block allocations (i.e., time and frequency slots in LTE/OFDMA networks, which are known as resource blocks) [235], [294], [346], control and assign transmission and power parameters in devices or in a group basis (e.g., algorithms to optimize the transmission and power parameters of WLAN devices, define and assign transmission power values to each resource block, at each base station, in LTE/OFDMA networks) [294], [235], simplified administration [298], [300], [235], easy management

of heterogenous network technologies [300], [235], [349], interoperability between different networks [349], [346], shared wireless infrastructures [349], seamless subscriber mobility and cellular networks [345], QoS and access control policies made feasible and easier [345], [346], and easy deployment of new applications [298], [235], [349].

One of the first steps towards realizing these features in wireless networks is to provide programmable and flexible stack layers for wireless networks [350], [235]. One of the first examples is OpenRadio [350], which proposes a software abstraction layer for decoupling the wireless protocol definition from the hardware, allowing shared MAC layers across different protocols using commodity multi-core platforms. OpenRadio can be seen as the “OpenFlow for wireless networks”. Similarly, SoftRAN [235] proposes to rethink the radio access layer of current LTE infrastructures. Its main goal is to allow operators to improve and optimize algorithms for better hand-overs, fine-grained control of transmit powers, resource block allocation, among other management tasks.

Light virtual access points (LVAPs) is another interesting way of improving the management capabilities of wireless networks, as proposed by the Odin [298] framework. In contrast to OpenRadio, it works with existing wireless hardware and does not impose any change to IEEE 802.11 standards. An LVAP is implemented as a unique BSSID associated with a specific client, which means that there is a one-to-one mapping between LVAPs and clients. This per-client access point (AP) abstraction simplifies the handling of client associations, authentication, handovers, and unified slicing of both the wired

and wireless portions of the network. Odin achieves control logic isolation between slices, since LVAPs are the primitive type upon which applications make control decisions, and applications do not have visibility of LVAPs from outside their slice. This empowers infrastructure operators to provide services through Odin applications, such as a mobility manager, client-based load balancer, channel selection algorithm, and wireless troubleshooting application within different network slices. For instance, when a user moves from one AP to another, the network mobility management application can automatically and proactively act and move the client LVAP from one AP to the other. In this way, a wireless client will not even notice that it started to use a different AP because there is no perceptible hand-off delay, as it would be the case in traditional wireless networks.

Very dense heterogeneous wireless networks have also been a target for SDN. These DenseNets have limitations due to constraints such as radio access network bottlenecks, control overhead, and high operational costs [294]. A dynamic two-tier SDN controller hierarchy can be adapted to address some of these constraints [294]. Local controllers can be used to take fast and fine-grained decisions, while regional (or “global”) controllers can have a broader, coarser-grained scope, i.e., that take slower but more global decisions. In such a way, designing a single integrated architecture that encompasses LTE (macro/pico/femto) and WiFi cells, while challenging, seems feasible.

Measurement & monitoring

Measurement and monitoring solutions can be divided in two classes. First, applications that provide new functionality for other networking services. Second, proposals that target to improve features of OpenFlow-based SDNs, such as to reduce control plane overload due to the collection of statistics.

An example of the first class of applications is improving the visibility of broadband performance [351], [6]. An SDN-based broadband home connection can simplify the addition of new functions in measurement systems such as BISmark [351], allowing the system to react to changing conditions in the home network [6]. As an example, a home gateway can perform reactive traffic shaping considering the current measurement results of the home network.

The second class of solutions typically involve different kinds of sampling and estimation techniques to be applied, in order to reduce the burden of the control plane with respect to the collection of data plane statistics. Different techniques have been applied to achieve this goal, such as stochastic and deterministic packet sampling techniques [352], traffic matrix estimation [260], fine-grained monitoring of wildcard rules [353], two-stage Bloom filters [354] to represent monitoring rules and provide high measurement accuracy without incurring in extra memory or control plane traffic overhead [302], and special monitoring functions (extensions to OpenFlow) in forwarding devices to reduce traffic and processing load on the control plane [355]. Point-to-point traffic matrix estimation, in particular, can help in network design and operational tasks such as load balancing, anomaly detection, capacity planning and network provisioning. With

information on the set of active flows in the network, routing information (e.g., from the routing application), flow paths, and flow counters in the switches it is possible to construct a traffic matrix using diverse aggregation levels for sources and destinations [260].

Other initiatives of this second class propose a stronger decoupling between basic primitives (e.g., matching and counting) and heavier traffic analysis functions such as the detection of anomaly conditions attacks [356]. A stronger separation favors portability and flexibility. For instance, a functionality to detect abnormal flows should not be constrained by the basic primitives or the specific hardware implementation. Put another way, developers should be empowered with streaming abstractions and higher level programming capabilities.

In that vein, some data and control plane abstractions have been specifically designed for measurement purposes. OpenSketch [309] is a special-purpose southbound API designed to provide flexibility for network measurements. For instance, by allowing multiple measurement tasks to execute concurrently without impairing accuracy. The internal design of an OpenSketch switch can be thought of as a pipeline with three stages (hashing, classification, and counting). Input packets first pass through a hashing function. Then, they are classified according to a matching rule. Finally, the match rule identifies a counting index, which is used to calculate the counter location in the counting stage. While a TCAM with few entries is enough for the classification stage, the flexible counters are stored in SRAM. This makes the OpenSketch’s operation efficient (fast matching) and cost-effective (cheaper SRAMs to store counters).

Other monitoring frameworks, such as OpenSample [307] and PayLess [311], propose different mechanisms for delivering real-time, low-latency and flexible monitoring capabilities to SDN without impairing the load and performance of the control plane. The proposed solutions take advantage of sampling technologies like sFlow [308] to monitor high-speed networks, and flexible collections of loosely coupled (plug-and-play) components to provide abstract network views yielding high-performance and efficient network monitoring approaches [307], [311], [353].

Security & Dependability

An already diverse set of security and dependability proposals is emerging in the context of SDNs. Most take advantage of SDN for improving services required to secure systems and networks, such as policy enforcement (e.g., access control, firewalling, middleboxes as middlepipes [27]) [100], [326], [335], [324], [27], DoS attacks detection and mitigation [323], [334], random host mutation [324] (i.e., randomly and frequently mutate the IP addresses of end-hosts to break the attackers’ assumption about static IPs, which is the common case) [329], monitoring of cloud infrastructures for fine-grained security inspections (i.e., automatically analyze and detour suspected traffic to be further inspected by specialized network security appliances, such as deep packet inspection systems) [321], traffic anomaly detection [352], [323], [334], fine-grained flow-based network access control [325], fine-grained policy enforcement for personal mobile applica-

tions [327] and so forth [100], [326], [323], [329], [321], [324], [335], [352]. Others address OpenFlow-based networks issues, such as flow rule prioritization, security services composition, protection against traffic overload, and protection against malicious administrators [201], [258], [320], [328], [199].

There are essentially two approaches, one involves using SDNs to improve network security, and another for improving the security of the SDN *itself*. The focus has been, thus far, in the latter.

Using SDN to improve the security of current networks. Probably the first instance of SDN was an application for security policies enforcement [100]. An SDN allows the enforcement to be done on the first entry point to the network (e.g., the Ethernet switch to which the user is connected to). Alternatively, in a hybrid environment, security policy enforcement can be made on a wider network perimeter through programmable devices (without the need to migrate the entire infrastructure to OpenFlow) [326]. With either application, malicious actions are blocked before entering the critical regions of the network.

SDN has been successfully applied for other purposes, namely for the detection (and reaction) against DDoS flooding attacks [323], and active security [319]. OpenFlow forwarding devices make it easier to collect a variety of information from the network, in a timely manner, which is very handy for algorithms specialized in detecting DDoS flooding attacks.

The capabilities offered by software-defined networks in increasing the ability to collect statistics data from the network and of allowing applications to actively program the forwarding devices, are powerful for proactive and smart security policy enforcement techniques such as Active security [319]. This novel security methodology proposes a novel feedback loop to improve the control of defense mechanisms of a networked infrastructure, and is centered around five core capabilities: protect, sense, adjust, collect, counter. In this perspective, active security provides a centralized programming interface that simplifies the integration of mechanisms for detecting attacks, by a) collecting data from different sources (to identify attacks), b) converging to a consistent configuration for the security appliances, and c) enforcing countermeasures to block or minimize the effect of attacks.

Improving the security of SDN itself. There are already some research efforts on identifying the critical security threats of SDNs and in augmenting its security and dependability [201], [258], [357]. Early approaches try to apply simple techniques, such as classifying applications and using rule prioritization, to ensure that rules generated by security applications will not be overwritten by lower priority applications [201]. Other proposals try to go a step further by providing a framework for developing security-related applications in SDNs [258]. However, there is still a long way to go in the development of secure and dependable SDN infrastructures [357]. An in-deep overview of SDN security issues and challenges can be found in Section V-F.

Data Center Networking

From small enterprises to large scale cloud providers, most of the existing IT systems and services are strongly dependent on highly scalable and efficient data centers. Yet,

these infrastructures still pose significant challenges regarding computing, storage and networking. Concerning the latter, data centers should be designed and deployed in such a way as to offer high and flexible cross-section bandwidth and low-latency, QoS based on the application requirements, high levels of resilience, intelligent resource utilization to reduce energy consumption and improve overall efficiency, agility in provisioning network resources, for example by means of network virtualization and orchestration with computing and storage, and so forth [358], [359], [360]. Not surprisingly, many of these issues remain open due to the complexity and inflexibility of traditional network architectures.

The emergence of SDN is expected to change the current state of affairs. Early research efforts have indeed showed that data center networking can significantly benefit from SDN in solving different problems such as live network migration [316], improved network management [316], [315], eminent failure avoidance [316], [315], rapid deployment from development to production networks [316], troubleshooting [316], [317], optimization of network utilization [317], [312], [314], [315], dynamic and elastic provisioning of middleboxes-as-a-service [27], minimization of flow setup latency and reduction of controller operating costs [361]. SDN can also offer networking primitives for cloud applications, solutions to predict network transfers of applications [312], [314], mechanisms for fast reaction to operation problems, network-aware VM placement [317], [313], QoS support [317], [313], realtime network monitoring and problem detection [317], [314], [315], security policy enforcement services and mechanisms [317], [313], and enable programmatic adaptation of transport protocols [312], [318].

SDN can help infrastructure providers to expose more networking primitives to their customers, by allowing virtual network isolation, custom addressing, and the placement of middleboxes and virtual desktop cloud applications [313], [362]. To fully explore the potential of virtual networks in clouds, an essential feature is virtual network migration. Similarly to traditional virtual machine migration, a virtual network may need to be migrated when its virtual machines move from one place to another. Integrating live migration of virtual machines and virtual networks is one of the forefront challenges [316]. To achieve this goal it is necessary to dynamically reconfigure all affected networking devices (physical or virtual). This was shown to be possible with SDN platforms, such as NVP [112].

Another potential application of SDN in data centers is in detecting abnormal behaviors in network operation [315]. By using different behavioral models and collecting the necessary information from elements involved in the operation of a data center (infrastructure, operators, applications), it is possible to continuously build signatures for applications by passively capturing control traffic. Then, the signature history can be used to identify differences in behavior. Every time a difference is detected, operators can reactively or proactively take corrective measures. This can help to isolate abnormal components and avoid further damage to the infrastructure.

Towards SDN App Stores

As can be observed in Tables IX and X, most SDN applications rely on NOX and OpenFlow. NOX was the first controller available for general use, making it a natural choice for most use-cases so far. As indicated by the sheer number of security-related applications, security is probably one of the killer applications for SDNs. Curiously, while most use cases rely on OpenFlow, new solutions such as SoftRAN are considering different APIs, as is the case of the Femto API [252], [301]. This diversity of applications and APIs will most probably keep growing in SDN.

There are other kinds of network applications that do not easily fit in our taxonomy, such as Avior [363], OESS [364], and SDN App Store [365], [366]. Avior and OESS are graphical interfaces and sets of software tools that make it easier to configure and manage controllers (e.g., Floodlight) and OpenFlow-enabled switches, respectively. By leveraging their graphical functions it is possible to program OpenFlow enabled devices without coding in a particular programming language.

The SDN App Store [365], [366], owned by HP, is probably the first SDN application market store. Customers using HP's OpenFlow controller have access to the online SDN App Store and are able to select applications to be dynamically downloaded and installed in the controller. The idea is similar to the Android Market or the Apple Store, making it easier for developers to provide new applications and for customers to obtain them.

I. Cross-layer issues

In this section we look at cross-layer problems such as debugging and troubleshooting, testing, verification, simulation and emulation.

Debugging and troubleshooting

Debugging and troubleshooting have been important subjects in computing infrastructures, parallel and distributed systems, embedded systems, and desktop applications [367], [368], [369], [370], [371], [372], [373]. The two predominant strategies applied to debug and troubleshoot are runtime debugging (e.g., gdb-like tools) and post-mortem analysis (e.g., tracing, replay and visualization). Despite the constant evolution and the emergence of new techniques to improve debugging and troubleshooting, there are still several open avenues and research questions [368].

Debugging and troubleshooting in networking is at a very primitive stage. In traditional networks, engineers and developers have to use tools such as ping, traceroute, tcpdump, nmap, netflow, and SNMP statistics for debugging and troubleshooting. Debugging a complex network with such primitive tools is very hard. Even when one considers frameworks such as XTrace [372], Nettleplay [374] and NetCheck [375], which improve debugging capabilities in networks, it is still difficult to troubleshoot networking infrastructures. For instance, these frameworks require a huge effort in terms of network instrumentation. The additional complexity introduced by different types of devices, technologies and vendor specific components and features make matters worse. As a consequence, these solutions may find

it hard to be widely implemented and deployed in current networks.

SDN offers some hope in this respect. The hardware-agnostic software-based control capabilities and the use of open standards for control communication can potentially make debug and troubleshoot easier. The flexibility and programmability introduced by SDN is indeed opening new avenues for developing better tools to debug, troubleshoot, verify and test networks [376], [377], [378], [379], [380], [381], [382], [383], [382].

Early debugging tools for OpenFlow-enabled networks, such as ndb [376], OFRewind [377] and NetSight [384], make it easier to discover the source of network problems such as faulty device firmware [376], inconsistent or non-existing flow rules [376], [377], lack of reachability [376], [377], and faulty routing [376], [377]. Similarly to the well-known gdb software debugger, ndb provides basic debugging actions such as *breakpoint*, *watch*, *backtrace*, *single-step*, and *continue*. These primitives help application developers to debug networks in a similar way to traditional software. By using ndb's postcards (i.e., a unique packet identifier composed of a truncated copy of the packet's header, the matching flow entry, the switch, and the output port), for instance, a programmer is able to quickly identify and isolate a buggy OpenFlow switch with hardware or software problems. If the switch is presenting abnormal behavior such as corrupting parts of the packet header, by analyzing the problematic flow sequences with a debugging tool one can find (in a matter of few seconds) where the packets of a flow are being corrupted, and take the necessary actions to solve the problem.

The OFRewind [377] tool works differently. The idea is to record and replay network events, in particular control messages. These usually account for less than 1% of the data traffic and are responsible for 95%-99% of the bugs [383]. This tool allows operators to perform fine-grained tracing of network behavior, being able to decide which subsets of the network will be recorded and, afterwards, select specific parts of the traces to be replayed. These replays provide valuable information to find the root cause of the network misbehavior. Likewise, NetRevert [385] also records the state of OpenFlow networks. However, the primary goal is not to reproduce network behavior, but rather to provide rollback recovery in case of failures, which is a common approach used in distributed systems for eliminating transient errors in nodes [386], [387].

Despite the availability of these debugging and verification tools, it is still difficult to answer questions such as: What is happening to my packets that are flowing from point A to point B? What path do they follow? What header modifications do they undergo on the way? To answer some of these questions one could recur to the *history* of the packets. A packet's history corresponds to the paths it uses to traverse the network, and the header modifications in each hop of the path. NetSight [384] is a platform whose primary goal is to allow applications that use the history of the packets to be built, in order to find out problems in a network. This platform is composed of three essential elements: (1) NetSight, with its dedicated servers that receive and process the postcards for building

the packet history, (2) the NetSight-SwitchAssist, which can be used in switches to reduce the processing burden on the dedicated servers, and (3) the NetSight-HostAssist to generate and process postcards on end hosts (e.g., in the hypervisor on a virtualized infrastructure).

`netwatch` [384], `netshark` [384] and `nprof` [384] are three examples of tools built over NetSight. The first is a live network invariant monitor. For instance, an alarm can be triggered every time a packet violates any invariant (e.g., no loops). The second, `netshark`, enables users to define and execute filters on the entire history of packets. With this tool, a network operator can view a complete list of properties of packets at each hop, such as input port, output port, and packet header values. Finally, `nprof` can be used to profile sets of network links to provide data for analyzing traffic patterns and routing decisions that might be contributing to link load.

Testing and verification

Verification and testing tools can complement debugging and troubleshooting. Recent research [381], [383], [380], [378], [382], [388], [389] has shown that verification techniques can be applied to detect and avoid problems in SDN, such as forwarding loops and black holes. Verification can be done at different layers (at the controllers, network applications, or network devices). Additionally, there are different network properties – mostly topology-specific – that can be formally verified, provided a network model is available. Examples of such properties are connectivity, loop freedom, and access control [29]. A number of tools have also been proposed to evaluate the performance of OpenFlow controllers by emulating the load of large-scale networks (e.g., Cbench [390], OFCBenchmark [391], PktBlaster [392]). Similarly, benchmarking tools for OpenFlow switches are also available (e.g., OFLOPS [379], FLOPS-Turbo [393]).

Tools such as NICE [378] generate sets of diverse streams of packets to test as many as possible events, exposing corner cases such as race conditions. Similarly, OFLOPS [379] provides a set of features and functions that allow the development and execution of a rich set of tests on OpenFlow-enabled devices. Its ultimate goal is to measure the processing capacity and bottlenecks of control applications and forwarding devices. With this tool, users are able to observe and evaluate forwarding table consistency, flow setup latency, flow space granularity, packet modification types, and traffic monitoring capabilities (e.g., counters).

FlowChecker [380], OFTEN [382], and VeriFlow [381] are three examples of tools to verify correctness properties violations on the system. While the former two are based on offline analysis, the latter is capable of online checking of network invariants. Verification constraints include security and reachability issues, configuration updates on the network, loops, black holes, etc.

Other formal modeling techniques, such as Alloy, can be applied to SDNs to identify unexpected behavior [388]. For instance, a protocol specification can be weak when it underspecifies some aspects of the protocol or due to a very specific sequence of events. In such situations, model checking techniques such as Alloy can help to find and correct unexpected

behaviors.

Tools such as FLOWGUARD [394] are specifically designed to detect and resolve security policy violations in OpenFlow-enabled networks. FLOWGUARD is able to examine on-the-fly network policy updates, check indirect security violations (e.g., OpenFlow’s `Set-Field` actions modification) and perform stateful monitoring. The framework uses five resolution strategies for real-time security policy violation resolution, flow rejecting, dependency breaking, update rejecting, flow removing, and packet blocking [394]. These resolutions are applied over diverse update situations in OpenFlow-enabled networks.

More recently, tools such as VeriCon [395] have been designed to verify the correctness of SDN applications in a large range of network topologies and by analyzing a broad range of sequences of network events. In particular, VeriCon confirms, or not, the correct execution of the SDN program.

One of the challenges in testing and verification is to verify forwarding tables in very large networks to find routing errors, which can cause traffic losses and security breaches, as quickly as possible. In large scale networks, it is not possible to assume that the network snapshot, at any point, is consistent, due to the frequent changes in routing state. Therefore, solutions such as HSA [396], Anteater [397], NetPlumber [398], VeriFlow [381], and assertion languages [399] are not suited for this kind of environment. Another important issue is related on how fast the verification process is done, especially in modern data centers that have very tight timing requirements. Libra [389] represents one of the first attempts to address these particular challenges of large scale networks. This tool provides the means for capturing stable and consistent snapshots of large scale network deployments, while also applying long prefix matching techniques to increase the scalability of the system. By using MapReduce computations, Libra is capable of verifying the correctness of a network with up to 10k nodes within one minute.

Anteater [397] is a tool that analyzes the data plane state of network devices by encoding switch configurations as boolean satisfiability problems (SAT) instances, allowing to use a SAT solver to analyze the network state. The tool is capable of verifying violations of invariants such as loop-free forwarding, connectivity, and consistency. These invariants usually indicate a bug in the network, i.e., their detection helps to increase the reliability of the network data plane.

Simulation and Emulation

Simulation and emulation software is of particular importance for fast prototyping and testing without the need for expensive physical devices. Mininet [110] is the first system that provides a quick and easy way to prototype and evaluate SDN protocols and applications. One of the key properties of Mininet is its use of software-based OpenFlow switches in virtualized containers, providing the exact same semantics of hardware-based OpenFlow switches. This means that controllers or applications developed and tested in the emulated environment can be (in theory) deployed in an OpenFlow-enabled network without any modification. Users can easily emulate an OpenFlow network with hundreds of

nodes and dozens of switches by using a single personal computer. Mininet-HiFi [400] is an evolution of Mininet that enhances the container-based (lightweight) virtualization with mechanisms to enforce performance isolation, resource provisioning, and accurate monitoring for performance fidelity. One of the main goals of Mininet-HiFi is to improve the reproducibility of networking research.

Mininet CE [401] and SDN Cloud DC [402] are extensions to Mininet for enabling large scale simulations. Mininet CE combines groups of Mininet instances into one cluster of simulator instances to model global scale networks. SDN Cloud DC enhances Mininet and POX to emulate an SDN-based intra-DC network by implementing new software modules such as data center topology discovery and network traffic generation. Recent emulation platform proposals that enable large scale experiments following a distributed approach include Max-iNet [403], DOT [404], and CityFlow [405]. The latter is a project with the main goal of building an emulated control plane for a city of one million inhabitants. Such initiatives are a starting point to provide experimental insights for large-scale SDN deployments.

The capability of simulating OpenFlow devices has also been added to the popular ns-3 simulator [406]. Another simulator is *fs-sdn*, which extends the *fs* simulation engine [407] by incorporating a controller and switching components with OpenFlow support. Its main goal is to provide a more realistic and scalable simulation platform as compared to Mininet. Finally, STS [408] is a simulator designed to allow developers to specify and apply a variety of test cases, while allowing them to interactively examine the state of the network.

V. ONGOING RESEARCH EFFORTS AND CHALLENGES

The research developments we have surveyed so far seek to overcome the challenges of realizing the vision and fulfilling the promises of SDN. While Section IV provided a perspective structured across the layers of the “SDN stack”, this section highlights research efforts we consider of particular importance for unleashing the full potential of SDN, and that therefore deserves a specific coverage in this survey.

A. Switch Designs

Currently available OpenFlow switches are very diverse and exhibit notable differences in terms of feature set (e.g., flow table size, optional actions), performance (e.g., fast vs. slow path, control channel latency/throughput), interpretation and adherence to the protocol specification (e.g., BARRIER command), and architecture (e.g., hardware vs. software designs).

Heterogenous Implementations

Implementation choices have a fundamental impact on the behavior, accuracy, and performance of switches, ranging from differences in flow counter behavior [416] to a number of other performance metrics [379]. One approach to accommodate such heterogeneity is through NOSIX, a portable API that separates the application expectations from the switch heterogeneity [245]. To do so, NOSIX provides a pipeline of multiple virtual flow tables and switch drivers. Virtual

flow tables are intended to meet the expectations of applications and are ultimately translated by the drivers into actual switch flow tables. Towards taming the complexity of multiple OpenFlow protocol versions with different sets of required and optional capabilities, a roadblock for SDN practitioners, tinyNBI [417] has been proposed as a simple API providing a unifying set of core abstractions of five OpenFlow protocol versions (from 1.0 to 1.4). Ongoing efforts to introduce a new Hardware Abstraction Layer (HAL) for non-OpenFlow capable devices [418] include the development of open source artifacts like ROFL (Revised OpenFlow Library) and the xDPd (eXtensible DataPath daemon), a framework for creating new OpenFlow datapath implementations based on a diverse set of hardware and software platforms. A related open source effort to develop a common library to implement OpenFlow 1.0 and 1.3 protocol endpoints (switch agents and controllers) is libfluid [419], winner of the OpenFlow driver competition organized by the ONF.

Within the ONF, the Forwarding Abstraction Working Group (FAWG) is pursuing another solution to the heterogeneity problem, through Table Type Patterns (TTPs) [121]. A TTP is a standards-based and negotiated switch-level behavioral abstraction. It consists of the relationships between tables forming a graph structure, the types of tables in the graph, a set of the parameterized table properties for each table in the graph, the legal `flow-mod` and `table-mod` commands for each flow table, and the metadata mask that can be passed between each table pair in the graph.

Flow Table Capacity

Flow matching rules are stored in flow tables inside network devices. One practical challenge is to provide switches with large and efficient flow tables to store the rules [420]. TCAMs are a common choice to hold flow tables. While flexible and efficient in terms of matching capabilities, TCAMs are costly and usually small (from 4K to 32K entries). Some TCAM chips today integrate 18 M-bit (configured as 500k entries * 36 bit per entry) into a single chip working at 133 Mhz [421], i.e., capable of 133M lookups per second. However, these chips are expensive and have a high-power consumption [422], representing a major power drain in a switching device [423]. These are some of the reasons why currently available OpenFlow devices have TCAMs with roughly 8K entries, where the actual capacity in terms of OpenFlow table size has a non-trivial relationship to the type of flow entries being used [424], [425]. OpenFlow version 1.1 introduced multiple tables, thereby adding extra flexibility and scalability. Indeed, OpenFlow 1.0 implied state explosion due to its flat table model [121]. However, supporting multiple tables in hardware is challenging and limited – yet another motivation for the ongoing ONF FAWG work on TTPs [121].

Some efforts focus on compression techniques to reduce the number of flow entries in TCAMs [426], [427], [428]. The Espresso heuristic [428] can be used to compress wild-cards of OpenFlow-based inter-domain routing tables, reducing the forwarding information base (FIB) by 17% and, consequently, saving up to 40,000 flow table entries [426]. Shadow MACs [427] propose label switching for solving two

TABLE XI
DEBUGGING, VERIFICATION AND SIMULATION

Group	Solution	Main purpose	Short description
Debugging	n:db [376]	<i>gdb</i> alike SDN debugging	Basic debugging primitives that help developers to debug their networks.
	NetSight [384]	multi purpose packet history	Allows to build flexible debugging, monitoring and profiling applications.
	OFRewind [377]	tracing and replay	OFRewind allows operators to do a fine-grained tracing of the network behavior. Operators can decide which subsets of the network will be recorded.
	PathletTracer [409]	inspect layer 2 paths	Allows to inspect low-level forwarding behavior through on-demand packet tracing capabilities.
	SDN traceroute [410]	query OpenFlow paths	Allows users to discover the forwarding behavior of any Ethernet packet and debug problems regarding both forwarding devices and applications.
Verification	Assertion language [399]	debug SDN apps	Enables assertions about the data plane on the apps with support to dynamic changing verification conditions.
	Cbench [390]	evaluate OpenFlow controllers	The Cbench framework can be used to emulate OpenFlow switches which are configured to generate workload to the controller.
	FLOVER [259]	model checking for security policies	FLOVER provides a provably correct and automatic method for verifying security properties with respect to a set of flow rules committed by an OF controller.
	FlowChecker [380]	flow table config verification	A tool used to verify generic properties of global behaviors based on flow tables.
	FLOWGUARD [394]	verify security policy	Provides mechanisms for accurate detection and resolution of firewall policy violations in OpenFlow-based networks.
	FlowTest [411]	verify network policies	Provides the means for testing stateful and dynamic network policies by systematically exploring the state space of the network data plane.
	NetPlumber [398]	real time policy checking	NetPlumber uses a set of policies and invariants to do real time checking. It leverages header space analysis and keeps a dependency graph between rules.
	NICE [378]	remove bugs in controllers	Its main goal is to test controller programs without requiring any type of modification or extra work for application programmers.
	OFCBenchmark [391]	evaluate OpenFlow controllers	creates independent virtual switches, making is possible to emulate different scenarios. Each switch has its how configuration and statistics.
	OFTEN [382]	catch correctness property violations	A framework designed to check SDN systems, analyzing controller and switch interaction, looking for correctness condition violation.
	OFLOPS [379]	evaluate OpenFlow switches	A framework with a rich set of tests for OpenFlow protocol, enabling to measure capabilities of both switch and applications.
	OFLOPS-Turbo [393]	evaluate OpenFlow switches	A framework that integrates OFLOPS with OSNT [412], a 10GbE traffic generation and monitoring system based on NetFPGA.
	PktBlaster [392]	emulation / benchmarking	Integrated test and benchmarking solution that emulates large scale software-defined networks.
	SDLoad [413]	evaluate OpenFlow controllers	A traffic generation framework with customizable workloads to realistically represent different types of applications.
	VeriCon [395]	verify SDN apps	Is a tool for verifying the correctness of SDN applications on large range of topologies and sequences of network events.
VeriFlow [381]	online invariant verification	It provides real time verification capabilities, while the network state is still evolving.	
Simulation & Emulation	DOT [404]	network emulation	Leverages VMs for large scale OpenFlow-based network emulations with resource allocation guarantees.
	<i>fs-sdn</i> [414]	fast simulation	Like Mininet, it provides a simulation environment, but with speed and scalability advantages.
	MaxiNet [403]	network simulation	Similarly to Mininet CE, it is a combination of Mininet tools for large scale simulation of network topologies and architectures.
	Mininet [110]	fast prototyping	It emulates and OpenFlow network using Open vSwitches to provide the exact same semantics of hardware devices.
	Mininet CE [401]	global network modeling	It is a combination of tools to create a Mininet cluster for large scale simulation of network topologies and architectures.
	Mininet-HiFi [400]	reproducible experimentation	Evolution of Mininet to enable repeatable and high fidelity experiments.
	ns-3 [406]	network simulation	The latest version of ns-3 simulator provides support to OpenFlow, enabling to create programmable network devices.
	SDN Cloud DC [402]	cloud data center emulation	The SDN Cloud DC solution allows users to evaluate the performance of their controllers at scale.
	STS [408]	troubleshooting	It simulates a network, allows to generate tricky test cases, and allows interactively examine the state of the network.
	VND-SDN [415]	simulation and analysis	Makes it easier to define experiments via GUI authoring of SDN scenarios and automatic generation of NSDL.

problems, consistent updates and rule space exhaustion, by using opaque values (similar to MPLS labels) to encode fine-grained paths as labels. A major benefit of fixed-size labels is relying on exact-math lookups which can be easily and cost-effectively implemented by simple hardware tables instead of requiring rules to be encoded in expensive TCAM tables.

Performance

Today, the throughput of commercial OpenFlow switches varies from 38 to 1000 `flow-mod` per second, with most devices achieving a throughput lower than 500 `flow-mod` per second [429], [430]. This is clearly a limiting factor that shall be addressed in the switch design process – support of OpenFlow in existing product lines has been more a retrofitting activity than a clean feature planning and implementation activity. Deployment experiences [431] have pointed to a series of challenges stemming from the limited embedded CPU power of current commercial OpenFlow switches. One approach to handle the problem consists of adding more powerful CPUs into the switches, as proposed in [432]. Others have proposed to rethink the distribution of control actions between external controllers and the OpenFlow agent inside the switch [416]. Our current understanding indicates that an effective way forward is a native design of SDN switches consistent with the evolution of the southbound API standardization activities [433], [121].

Evolving Switch Designs & Hardware Enhancements

As in any software/hardware innovation cycle, a number of advancements are to be expected from the hardware perspective to improve SDN capabilities and performance. New SDN switch designs are appearing in a myriad of hardware combinations to efficiently work together with TCAMs, such as SRAM, RDRAM, DRAM, GPU, FPGA, NPs, CPUs, among other specialized network processors [434], [435], [436], [437], [438], [439]. These early works suggest the need for additional efforts into new hardware architectures for future SDN switching devices. For instance, some proposals target technologies such as GPUs that have demonstrated 20 Gbps with flow tables of up to 1M exact match entries and up to 1K wildcard entries [436]. Alternatives to TCAM-based designs include new hardware architectures and components, as well as new and more scalable forwarding planes, such as the one proposed by the Rain Man firmware [440]. Other design solutions, such as parallel lookup models [441], can also be applied to SDN to reduce costs in switching and routing devices. Recent proposals on cache-like OpenFlow switch arrangements [442] shed some light on overcoming the practical limitations of flow table sizes with clever switching designs. Additionally, counters represent another practical challenge in SDN hardware implementations. Many counters already exist, and they could lead to significant control plane monitoring overhead [416]. Software-defined counters (SDC) [432] have been proposed to provide both scalability and flexibility.

Application-aware SDN architectures are being proposed to generalize the standard OpenFlow forwarding abstractions by including stateful actions to allow processing information from layers 4 to 7 [443]. To this end, application flow tables are proposed as data plane application modules that require only

local state, i.e., do not depend on a global view of the network. Those tiny application modules run inside the forwarding devices (and can be installed on-demand), alleviating the overhead on the control plane and augmenting the efficiency of certain tasks, which can be kept in the data plane. Similarly, other initiatives propose solutions based on pre-installed state machines. FAST (Flow-level State Transitions) [444] allows controllers to proactively program state transitions in forwarding devices, allowing switches to run dynamic actions that require only local information.

Other approaches towards evolving switch designs include Caching in Buckets (CAB), a reactive wildcard caching proposal that uses a geometric representation of the rule set, which is divided into small logical structures (buckets) [445]. Through this technique CAB is able to solve the rule dependency problem and achieve efficient usage of control plane resources, namely bandwidth, controller processing load, and flow setup latency.

New programmable Ethernet switch chips, such as XPliant Ethernet [446], are emerging into this new market of programmable networks. Its main aim is enabling new protocol support and the addition of new features through software updates, increasing flexibility. One example of such flexibility is the support of GENEVE [39], a recent effort towards generic network virtualization encapsulation protocols, and OpenFlow. The throughput of the first family of XPliant Ethernet chip varies from 880 Gbps to 3.2 Tbps, supporting up to 64 ports of 40 GbE or 50 GbE, for instance.

Microchip companies like Intel are already shipping processors with flexible SDN capabilities to the market [447]. Recent advances in general-purpose CPU technology include a data-plane development kit (DPDK) [448] that allows high-level programming of how data packets shall be processed directly within network interface cards. Prototype implementations of Intel DPDK accelerated switch shows the potential to deliver high-performance SDN software switches [439]. This trend is likely to continue since high-speed and specialized hardware is needed to boost SDN performance and scalability for large, real-world networks. Hardware-programmable technologies such as FPGA are widely used to reduce time and costs of hardware-based feature implementations. NetFPGA, for instance, has been a pioneering technology used to implement OpenFlow 1.0 switches [435], providing a commodity cost-effective prototyping solution. Another line of work on SDN data planes proposes to augment switches with FPGA to (remotely) define the queue management and scheduling behaviour of packet switches [449]. Finally, recent developments have shown that state-of-the-art System-on-chip (SoC) platforms, such as the Xilinx Zynq ZC706 board, can be used to implement OpenFlow devices yielding 88 Gbps throughput for 1K flow supporting dynamic updates [450].

Native SDN Switch Designs

Most of the SDN switch (re)design efforts so far follow an evolutionary approach to retrofit OpenFlow-specific programmable features into existing hardware layouts, following common wisdom on switch/router designs and consolidated technologies (e.g., SRAM, TCAM, FPGA). One departure

from this approach is the ongoing work on *forwarding metamorphosis* [433], a reconfigurable match table model inspired from RISC-like pipeline architecture applied to switching chips. This work illustrates the feasibility of realizing a minimal set of action primitives for flexible header processing in hardware, at almost no additional cost or power. Also in line with the core SDN goals of highly flexible and programmable (hardware-based) data planes, Protocol-Oblivious Forwarding (POF) [120] aims at overcoming some of the limitations of OpenFlow (e.g., expressiveness, support of user-defined protocols, memory efficiency), through generic flow instruction sets. Open-source prototypes are available [31] as well as evaluation results showing the line-speed capabilities using a network processing unit (NPU)-based [451] proof of concept implementation. In this line, we already mentioned OpenState [155], another initiative that aims to augment the capability and flexibility of forwarding devices. By taking advantage of eXtended Finite State Machines (XFSM) [452], [453], OpenState proposes an abstraction – as a super set of OpenFlow primitives – to enable stateful handling of OpenFlow rules inside forwarding devices.

In the same way as TTPs allow controllers to compile the right set of low-lever instructions known to be supported by the switches, a new breed of switch referred to as P4 (programmable, protocol-independent packet processor) [454] suggests an evolution path for OpenFlow, based on a high-level compiler. This proposal would allow the functionality of programmable switches (i.e., pipeline, header parsing, field matching) to be not only specified by the controller but also changed in the field. In this model, programmers are able to decide how the forwarding plane processes packets without caring about implementation details. It is then the compiler that transforms the imperative program into a control flow graph that can be mapped to different target switches.

B. Controller Platforms

In the SDN model, the controller platform is a critical pillar of the architecture, and, as such, efforts are being devoted to turn SDN controllers into high-performance, scalable, distributed, modular, and highly-available programmer-friendly software. Distributed controller platforms, in particular, have to address a variety of challenges. Deserving special consideration are the latency between forwarding devices and controller instances, fault-tolerance, load-balancing, consistency and synchronization, among other issues [7], [455], [456]. Operators should also be able to observe and understand how the combination of different functions and modules can impact their network [457].

As the SDN community learns from the development and operational experiences with OpenFlow controllers (e.g., Beacon [186]), further advancements are expected in terms of raw performance of controller implementations, including the exploitation of hierarchical designs and optimized buffer sizing [458]. One approach to increase the performance of controllers is the IRIS IO engine [459], enabling significant increases in the flow-setup rate of SDN controllers. Another way of reducing the control plane overhead is by keeping

a compressed copy of the flow tables in the controller’s memory [460].

Modularity & Flexibility

A series of ongoing research efforts target the modular and flexible composition of controllers. RAON [461] proposes a recursive abstraction of OpenFlow controllers where each controller sees the controllers below as OpenFlow switches. Open research issues include the definition of suitable interfaces between the different layers in such a hierarchy of controllers.

Other open issues to be further investigated in this context are the East/westbound APIs, and their use in enabling suitable hierarchical designs to achieve scalability, modularity, and security [218]. For instance, each level of a hierarchy of controllers can offer different abstractions and scopes for either intra- and inter-data center routing, thus increasing scalability and modularity. Similarly, from a security perspective, each hierarchical level may be part of a different trust domain. Therefore, east/westbound interfaces between the different layers of controllers should be capable of enforcing both intra- and inter-domain security policies.

Another important observation is that, currently, the lack of modularity in most SDN controllers forces developers to re-implement basic network services from scratch in each new application [29]. As in software engineering in general, lack of modularity results in controller implementations that are hard to build, maintain, and extend – and ultimately become resistant to further innovations, resembling traditional “hardware-defined” networks. As surveyed in Section IV-G, SDN programming abstractions (e.g., Pyretic [225]) introduce modularity in SDN applications and simplify their development altogether. Further research efforts (e.g., Corybantic [462]) try to achieve modularity in SDN control programs. Other contributions towards achieving modular controllers can be expected from other areas of computer science (e.g., principles from Operating System [196]) and best practices of modern cloud-scale software applications.

Interoperability and application portability

Similarly to forwarding device vendor agnosticism that stems from standard southbound interfaces, it is important to foster interoperability between controllers. Early initiatives towards more interoperable control platforms include portable programming languages such as Pyretic [225] and east/westbound interfaces among controllers, such as SDNi [209], ForCES CE-CE interface [30], [211], and ForCES Intra-NE mechanisms [212]. However, these efforts are yet far from fully realizing controller interoperability and application portability.

In contrast to Pyretic [247], PANE [197], Maple [262], and Corybantic [462], which are restricted to traffic engineering applications and/or impose network state conflict resolution at the application level (making application design and testing more complicated), Statesman [463] proposes a framework to enable a variety of loosely-coupled network applications to co-exist on the same control plane without compromising network safety and performance. This framework makes application development simpler by automatically and transparently resolving conflicts. In other words, Statesman allows a

safe composition of uncoordinated or conflicting application's actions.

Another recent approach to simplify network management is the idea of compositional SDN hypervisors [181]. Its main feature is allowing applications written in different languages, or on different platforms, to work together in processing the same traffic. The key integration component is a set of simple prioritized lists of OpenFlow rules, which can be generated by different programming languages or applications.

High-Availability

In production, SDN controllers need to sustain healthy operation under the pressure of different objectives from the applications they host. Many advances are called for in order to deal with potential risk vectors of controller-based solutions [357]. Certainly, many solutions will leverage on results from the distributed systems and security communities made over the last decade. For instance, recent efforts propose consistent, fault-tolerant data stores for building reliable distributed controllers [213], [198], [456].

Another possible approach towards building low latency, highly available SDN controllers is to exploit controller locality [455], [464]. Classical models of distributed systems, such as LOCAL and CONGEST [465], can be explored to solve this problem. Those models can be used to develop coordination protocols that enable each controller to take independent actions over events that take place in its local neighborhood [455].

Another core challenge relates to the fundamental trade-offs between the consistency model of state distribution in distributed SDN controllers, the consistency requirements of control applications, and performance [464]. To ease development, the application should ideally not be aware of the vagaries of distributed state. This implies a strong consistency model, which can be achieved with distributed data stores as proposed recently [213]. However, keeping all control data in a consistent distributed data store is unfeasible due to the inherent performance penalties. Therefore, hybrid solutions are likely to co-exist requiring application developers to be aware of the trade-offs and penalties of using, or not, a strong consistency model, a tenet of the distributed Onix controller [7].

High availability can also be achieved through improved southbound APIs and controller placement heuristics and formal models [466], [467], [468]. These aim to maximize resilience and scalability by allowing forwarding devices to connect to multiple controllers in a cost-effective and efficient way [467]. Early efforts in this direction have already shown that forwarding devices connecting to two or three controllers can typically achieve high availability (up to five nines) and robustness in terms of control plane connectivity [466], [468]. It has also been shown that the number of required controllers is more dependent on the topology than on network size [466]. Another finding worth mentioning is the fact that for most common topologies and network sizes fewer than ten controllers seem to be enough [466].

Delegation of control

To increase operational efficiency, SDN controllers can

delegate control functions to report state and attribute value changes, threshold crossing alerts, hardware failures, and so forth. These notifications typically follow a publish/subscribe model, i.e., controllers and applications subscribe (on-demand) to the particular class of notifications they are interested in. In addition, these subsystems may provide resilience and trustworthiness properties [469].

Some reasons for delegating control to the data plane include [218]:

- Low latency response to a variety of network events;
- The amount of traffic that must be processed in the data plane, in particular in large-scale networks such as data centers;
- Low level functions such as those (byte- or bit-oriented) required by repetitive SDH (Synchronous Digital Hierarchy) [470] multiplex section overhead;
- Functions well-understood and standardized, such as encryption, BIP [471], AIS [472] insertion, MAC learning, and CCM (Codec Control Messages) [473] exchanges;
- Controller failure tolerance, i.e., essential network functions should be able to keep a basic network operation even when controllers are down;
- Basic low-level functions usually available in data plane silicon, such as protection switching state machines, CCM counters and timers;
- All those functions that do not add any value when moved from the data to the control plane.

Strong candidates for execution in the forwarding devices instead of being implemented in the control platforms thus include OAM, ICMP processing, MAC learning, neighbor discovery, defect recognition and integration [218]. This would not only reduce the overhead (traffic and computing) of the control plane, but also improve network efficiency by keeping basic networking functions in the data plane.

C. Resilience

Achieving resilient communication is a top purpose of networking. As such, SDNs are expected to yield the same levels of availability as legacy and any new alternative technology. Split control architectures as SDN are commonly questioned [474] about their actual capability of being resilient to faults that may compromise the control-to-data plane communications and thus result in “brainless” networks. Indeed, the malfunctioning of particular SDN elements should not result in the loss of availability. The relocation of SDN control plane functionality, from inside the boxes to remote, logically centralized loci, becomes a challenge when considering critical control plane functions such as those related to link failure detection or fast reaction decisions. The resilience of an OpenFlow network depends on fault-tolerance in the data plane (as in traditional networks) but also on the high availability of the (logically) centralized control plane functions. Hence, the resilience of SDN is challenging due to the multiple possible failures of the different pieces of the architecture.

As noted in [475], there is a lack of sufficient research and experience in building and operating fault-tolerant SDNs. Google B4 [8] may be one of the few examples that have

proven that SDN can be resilient at scale. A number of related efforts [355], [476], [477], [261], [478], [479], [480], [481], [361] have started to tackle the concerns around control plane split architectures. The distributed controller architectures surveyed in Section IV-D are examples of approaches towards resilient SDN controller platforms with different tradeoffs in terms of consistency, durability and scalability.

On a detailed discussion on whether the CAP theorem [482] applies to networks, Panda et al. [477] argue that the tradeoffs in building consistent, available and partition-tolerant distributed databases (i.e., CAP theorem) are applicable to SDN. The CAP theorem demonstrates that it is impossible for datastore systems to simultaneously achieve strong consistency, availability and partition tolerance. While availability and partition tolerance problems are similar in both distributed databases and networks, the problem of consistency in SDN relates to the consistent application of policies.

Considering an OpenFlow network, when a switch detects a link failure (`port-down` event), a notification is sent to the controller, which then takes the required actions (re-route computation, pre-computed back-up path lookup) and installs updated flow entries in the required switches to redirect the affected traffic. Such reactive strategies imply (1) high restoration time due to the necessary interaction with the controller; and (2) additional load on the control channel. One experimental work on OpenFlow for carrier-grade networks investigated the restoration process and measured a restoration times in the order of 100 ms [476]. The delay introduced by the controller may, in some cases, be prohibitive.

In order to meet carrier grade requirements (e.g., 50 ms of recovery time), protection schemes are required to mitigate the effects of a separate control plane. Suitable protection mechanisms (e.g., installation of pre-established backup paths in the forwarding devices) can be implemented by means of OpenFlow group table entries using “fast-failover” actions. An OpenFlow fault management approach [355] similar to MPLS global path protection could also be a viable solution, provided that OpenFlow switches are extended with end-to-end path monitoring capabilities similarly to those specified by Bidirectional Forwarding Detection (BFD) [483]. Such protection schemes are a critical design choice for larger scale networks and may also require considerable additional flow space. By using primary and secondary path pairs programmed as OpenFlow fast failover group table entries, a path restoration time of 3.3 ms has been reported [484] using BFD sessions to quickly detect link failures.

On a related line of data plane resilience, SlickFlow [480] leverages the idea of using packet header space to carry alternative path information to implement resilient source routing in OpenFlow networks. Under the presence of failures along a primary path, packets can be rerouted to alternative paths by the switches themselves without involving the controller. Another recent proposal that uses in-packet information is INFLEX [481], an SDN-based architecture for cross-layer network resilience which provides on-demand path fail-over by having end-points tag packets with virtual routing plane information that can be used by egress routers to re-route by changing tags upon failure detection.

Similarly to SlickFlow, OSP [279] proposes a protection approach for data plane resilience. It is based on protecting individual segments of a path avoiding the intervention of the controller upon failure. The recovery time depends on the failure detection time, i.e., a few tens of milliseconds in the proposed scenarios. In the same direction, other proposals are starting to appear for enabling fast failover mechanisms for link protection and restoration in OpenFlow-based networks [485].

Language-based solutions to the data plane fault-tolerance problem have also been proposed [261]. In this work the authors propose a language that compiles regular expressions into OpenFlow rules to express what network paths packets may take and what degree of (link level) fault tolerance is required. Such abstractions around fault tolerance allow developers to build fault recovery capabilities into applications without huge coding efforts.

D. Scalability

Scalability has been one of the major concerns of SDNs from the outset. This is a problem that needs to be addressed in any system – e.g., in traditional networks – and is obviously also a matter of much discussion in the context of SDN [11].

Most of the scalability concerns in SDNs are related to the decoupling of the control and data planes. Of particular relevance are reactive network configurations where the first packet of a new flow is sent by the first forwarding element to the controller. The additional control plane traffic increases network load and makes the control plane a potential bottleneck. Additionally, as the flow tables of switches are configured in real-time by an outside entity, there is also the extra latency introduced by the flow setup process. In large-scale networks controllers will need to be able to process millions of flows per second [486] without compromising the quality of its service. Therefore, these overheads on the control plane and on flow setup latency are (arguably) two of the major scaling concerns in SDN.

As a result, several efforts have been devoted to tackle the SDN scaling concerns, including DevoFlow [416], Software-Defined Counters (SDCs) [432], DIFANE [487], Onix [7], HyperFlow [195], Kandoo [228], Maestro [188], NOX-MT [187], and Maple [262]. Still related to scalability, the notion of elasticity in SDN controllers is also being pursued [479], [361]. Elastic approaches include dynamically changing the number of controllers and their locations under different conditions [488].

Most of the research efforts addressing scaling limitations of SDN can be classified in three categories: data plane, control plane, and hybrid. While targeting the data plane, proposals such as DevoFlow [416] and Software-Defined Counters (SDC) [432] actually reduce the overhead of the control plane by delegating some work to the forwarding devices. For instance, instead of requesting a decision from the controller for every flow, switches can selectively identify the flows (e.g., elephant flows) that may need higher-level decisions from the control plane applications. Another example is to introduce more powerful general purpose CPUs

in the forwarding devices to enable SDCs. A general purpose CPU and software-defined counters offer new possibilities for reducing the control plane overhead by allowing software-based implementations of functions for data aggregation and compression, for instance.

Maestro [188], NOX-MT [187], Kandoo [228], Beacon [186], and Maple [262] are examples of the effort on designing and deploying high performance controllers, i.e., trying to increase the performance of the control plane. These controllers mainly explore well-known techniques from networking, computer architectures and high performance computing, such as buffering, pipelining and parallelism, to increase the throughput of the control platform.

The hybrid category is comprised of solutions that try to split the control logic functions between specialized data plane devices and controllers. In this category, DIFANE [487] proposes authoritative (intermediate) switches to keep all traffic in the data plane, targeting a more scalable and efficient control plane. Authoritative switches are responsible for installing rules on the remaining switches, while the controller is still responsible for generating all the rules required by the logic of applications. By dividing the controller work with these special switches, the overall system scales better.

Table XII provides a non-exhaustive list of proposals addressing scalability issues of SDN. We characterize these issues by application domain (control or data plane), their purpose, the throughput in terms of number of flows per second (when the results of the experiments are reported), and the strategies used. As can be observed, the vast majority are control plane solutions that try to increase scalability by using distributed and multi-core architectures.

Some figures are relatively impressive, with some solutions achieving up to 20M flows/s. However, we should caution the reader that current evaluations consider only simple applications and count basically the number of `packet-in` and `packet-out` messages to measure throughput. The actual performance of controllers will be affected by other factors, such as the number and complexity of the applications running on the controller and security mechanisms implemented. For example, a routing algorithm consumes more computing resources and needs more time to execute than a simple learning switch application. Also, current evaluations are done using plain TCP connections. The performance is very likely to change when basic security mechanisms are put in place, such as TLS, or more advanced mechanisms to avoid eavesdropping, man-in-the-middle and DoS attacks on the control plane.

Another important issue concerning scalability is data distribution among controller replicas in distributed architectures. Distributed control platforms rely on data distribution mechanisms to achieve their goals. For instance, controllers such as Onix, HyperFlow, and ONOS need mechanisms to keep a consistent state in the distributed control platform. Recently, experimental evaluations have shown that high performance distributed and fault-tolerant data stores can be used to tackle such challenges [213]. Nevertheless, further work is necessary to properly understand state distribution trade-offs [464].

E. Performance evaluation

As introduced in Section IV-A, there are already several OpenFlow implementations from hardware and software vendors being deployed in different types of networks, from small enterprise to large-scale data centers. Therefore, a growing number of experiments over SDN-enabled networks is expected in the near future. This will naturally create new challenges, as questions regarding SDN performance and scalability have not yet been properly investigated. Understanding the performance and limitation of the SDN concept is a requirement for its implementation in production networks. There are very few performance evaluation studies of OpenFlow and SDN architecture. Although simulation studies and experimentation are among the most widely used performance evaluation techniques, analytical modeling has its own benefits too. A closed-form description of a networking architecture paves the way for network designers to have a quick (and approximate) estimate of the performance of their design, without the need to spend considerable time for simulation studies or expensive experimental setup [431].

Some work has investigated ways to improve the performance of switching capabilities in SDN. These mainly consist of observing the performance of OpenFlow-enabled networks regarding different aspects, such as lookup performance [490], hardware acceleration [437], the influence of types of rules and packet sizes [491], performance bottlenecks of current OpenFlow implementations [416], how reactive settings impact the performance on data center networks [492], and the impact of configuration on OpenFlow switches [390].

Design choices can have a significant impact on the lookup performance of OpenFlow switching in Linux operating system using standard commodity network interface cards [490]. Just by using commodity network hardware the packet switching throughput can be improved by up to 25% when compared to one based on soft OpenFlow switching [490]. Similarly, hardware acceleration based on network processors can also be applied to perform OpenFlow switching. In such cases, early reports indicate that performance, in terms of packet delay, can be improved by 20% when compared to conventional designs [437].

By utilizing Intel's DPDK library [448], it has been shown that is possible to provide flexible traffic steering capability at the hypervisor level (e.g., KVM) without the performance limitations imposed by traditional hardware switching techniques [493], such as SR-IOV [494]. This is particularly relevant since most of the current enterprise deployments of SDN are in virtualized data center infrastructures, as in VMware's NVP solution [112].

Current OpenFlow switch implementations can lead to performance bottlenecks with respect to the CPU load [416]. Yet, modifications on the protocol specification can help reduce the occurrence of these bottlenecks. Further investigations provide measurements regarding the performance of the OpenFlow switch for different types of rules and packet sizes [491].

In data centers, a reactive setting of flow rules can lead to an unacceptable performance when only eight switches are handled by one OpenFlow controller [492]. This means that

TABLE XII
SUMMARY AND CHARACTERIZATION OF SCALABILITY PROPOSALS FOR SDNS.

Solution	Domain	Proposes	Main purpose	Flows/s	Resorts to
Beacon [186]	control plane	a multi-threaded controller	improve controller performance	12.8M	High performance flow processing capabilities using pipeline threads and shared queues.
Beacon cluster [489]	control plane	coordination framework	create clusters of controllers	6.2M	A coordination framework to create high-performance clusters of controllers.
DevoFlow [416]	data plane	thresholds for counters, type of flow detection	reduce the control plane overhead	—	Reduce the control traffic generated by counters statistics monitoring.
DIFANE [487]	control and data plane	authoritative specialized switches	improve data plane performance	500K	Maintain flows in the data plane reducing controller work.
Floodlight [186]	control plane	a multi-threaded controller	Improve controller performance	1.2M	High performance flow processing capabilities.
HyperFlow [195]	control plane	a distributed controller	distribute the control plane	—	Application on top of NOX to provide control message distribution among controllers.
Kandoo [228]	control plane	a hierarchical controller	distribute the control plane hierarchically	—	Use two levels of controller (local and root) to reduce control traffic.
Maestro [188]	control plane	a multi-threaded controller	improve controller performance	4.8M	High performance flow processing capabilities.
Maestro cluster [489]	control plane	coordination framework	create clusters of controllers	1.8M	A coordination framework to create high-performance clusters of controllers.
Maple [262]	control plane	programming language	scaling algorithmic policies	20M	Algorithmic policies and user- and OS-level threads on multicore systems (e.g., 40+ cores).
NOX [186]	control plane	a multi-threaded controller	improve controller performance	5.3M	High performance flow processing capabilities.
NOX-MT [187]	control plane	a multi-threaded controller	improve controller performance	1.8M	High performance flow processing capabilities.
NOX cluster [489]	control plane	coordination framework	create clusters of controllers	3.2M	A coordination framework to create high-performance clusters of controllers.
Onix [7]	control plane	a distributed control platform	robust and scalable control platform	—	Provide a programmable and flexible distributed NIB for application programmers.
SDCs [432]	data plane	Software-Defined Counters	reduce the control plane overhead	—	Remove counters from the ASIC to a general purpose CPU, improving programmability.

large-scale SDN deployments should probably not rely on a purely reactive “modus operandi”, but rather on a combination of proactive and reactive flow setup.

To foster the evaluation of different performance aspects of OpenFlow devices, frameworks such as OFLOPS [379], OFLOPS-Turbo [393], Cbench [187], and OFCBenchmark [391] have been proposed. They provide a set of tools to analyze the performance of OpenFlow switches and controllers. Cbench [187], [390] is a benchmark tool developed to evaluate the performance of OpenFlow controllers. By taking advantage of the Cbench, it is possible to identify performance improvements for OpenFlow controllers based on different environment and system configurations, such as the number of forwarding devices, network topology, overall network workload, type of equipments, forwarding complexity, and overhead of the applications being executed on top of controllers [187]. Therefore, such tools can help system designers make better decisions regarding the performance of devices and the network, while also allowing end-users to measure the device performance and better decide which one is best suited for the target network infrastructure.

Surprisingly, despite being designed to evaluate the performance of controllers, Cbench is currently a single-threaded tool. Therefore, multiple instances have to be started to utilize

multiple CPUs. It also only establishes one controller connection for all emulated switches. Unfortunately, this means little can be derived from the results in terms of controller performance and behavior or estimation of different bounds at the moment. For instance, aggregated statistics are gathered for all switches but not for each individual switch. As a result, it is not possible to identify whether all responses of the controller are for a single switch, or whether the capacity of the controller is actually shared among the switches. Flexible OpenFlow controller benchmarks are available though. OFCBenchmark [391] is one of the recent developments. It creates a set of message-generating virtual switches, which can be configured independently from each other to emulate a specific scenario and to maintain their own statistics.

Another interesting question to pose when evaluating the performance of SDN architectures is what is the required number of controllers for a given network topology and where to place the controllers [495], [467]. By analyzing the performance of controllers in different network topologies, it is possible to conclude that one controller is often enough to keep the latency at a reasonable rate [495]. Moreover, as observed in the same experiments, in the general case adding k controllers to the network can reduce the latency by a factor of k . However, there are cases, such as large scale networks and

WANs, where more controllers should be deployed to achieve high reliability and low control plane latency.

Recent studies also show that the SDN control plane cannot be fully physically centralized due to responsiveness, reliability and scalability metrics [464], [467]. Therefore, distributed controllers are the natural choice for creating a logically centralized control plane, while being capable of coping with the demands of large scale networks. However, distributed controllers bring additional challenges, such as the consistency of the global network view, which can significantly affect the performance of the network if not carefully engineered. Taking two applications as examples, one that ignores inconsistencies and another that takes inconsistency into consideration, it is possible to observe that optimality is significantly affected when inconsistencies are not considered and that the robustness of an application is increased when the controller is aware of the network state distribution [464].

Most of these initiatives towards identifying the limitations and bottlenecks of SDN architectures can take a lot of time and effort to produce consistent outputs due to the practical development and experimentation requirements. As mentioned before, analytic models can quickly provide performance indicators and potential scalability bottlenecks for an OpenFlow switch-controller system before detailed data is available. While simulation can provide detailed insight into a certain configuration, the analytical model greatly simplifies a conceptual deployment decision. For instance, a Network calculus-based model can be used to evaluate the performance of an SDN switch and the interaction of SDN switches and controllers [496]. The proposed SDN switch model captured the closed form of the packet delay and buffer length inside the SDN switch according to the parameters of a cumulative arrival process. Using recent measurements, the authors have reproduced the packet processing delay of two variants of OpenFlow switches and computed the buffer requirements of an OpenFlow controller. Analytic models based on queuing theory for the forwarding speed and blocking probability of current OpenFlow switches can also be used to estimate the performance of the network [490].

F. Security and Dependability

Cyber-attacks against financial institutions, energy facilities, government units and research institutions are becoming one of the top concerns of governments and agencies around the globe [497], [498], [499], [500], [501], [502]. Different incidents, such as Stuxnet [501], have already shown the persistence of threat vectors [503]. Put another way, these attacks are capable of damaging a nation's wide infrastructure, which represent a significant and concerning issue. As expected, one of the most common means of executing those attacks is through the network, either the Internet or the local area network. It can be used as a simple transport infrastructure for the attack or as a potentialized weapon to amplify the impact of the attack. For instance, high capacity networks can be used to launch large-scale attacks, even though the attacker has only a low capacity network connection at his premises.

Due to the danger of cyber-attacks and the current landscape of digital threats, security and dependability are top priorities

in SDN. While research and experimentation on software-defined networks is being conducted by some commercial players (e.g., Google, Yahoo!, Rackspace, Microsoft), commercial adoption is still in its early stage. Industry experts believe that security and dependability are issues that need to be addressed and further investigated in SDN [357], [504], [505].

Additionally, from the dependability perspective, availability of Internet routers is nowadays a major concern with the widespread of clouds and their strong expectations about the network [506]. It is therefore crucial to achieve high levels of availability on SDN control platforms if they are to become the main pillars of networked applications [466].

Different threat vectors have already been identified in SDN architectures [357], as well as several security issues and weaknesses in OpenFlow-based networks [507], [508], [509], [201], [510], [194], [511], [512]. While some threat vectors are common to existing networks, others are more specific to SDN, such as attacks on control plane communication and logically-centralized controllers. It is worth mentioning that most threats vectors are independent of the technology or the protocol (e.g., OpenFlow, POF, ForCES), because they represent threats on conceptual and architectural layers of SDN itself.

As shown in Figure 10 and Table XIII, there are at least seven identified threats vector in SDN architectures. The first threat vector consists of forged or faked traffic flows in the data plane, which can be used to attack forwarding devices and controllers. The second allows an attacker to exploit vulnerabilities of forwarding devices and consequently wreak havoc with the network. Threat vectors three, four and five are the most dangerous ones, since they can compromise the network operation. Attacks on the control plane, controllers and applications can easily grant an attacker the control of the network. For instance, a faulty or malicious controller or application could be used to reprogram the entire network for data theft purposes, e.g., in a data center. The sixth threat vector is linked to attacks on and vulnerabilities in administrative stations. A compromised critical computer, directly connected to the control network, will empower the attacker with resources to launch more easily an attack to the controller, for instance. Last, threat vector number seven represents the lack of trusted resources for forensics and remediation, which can compromise investigations (e.g., forensics analysis) and preclude fast and secure recovery modes for bringing the network back into a safe operation condition.

As can be observed in Table XIII, threat vectors 3 to 5 are specific to SDN as they stem from the separation of the control and data planes and the consequent introduction of a new entity in these networks – the logically centralized controller. The other vectors were already present in traditional networks. However, the impact of these threats could be larger than today – or at least it may be expressed differently – and as a consequence it may need to be dealt with differently.

OpenFlow networks are subject to a variety of security and dependability problems such as spoofing [507], tampering [507], repudiation [507], information disclosure [507], denial of service [507], [509], [510], elevation of privi-




Fig. 10. Main threat vectors of SDN architectures

TABLE XIII
SDN SPECIFIC VS. NON-SPECIFIC THREATS

Threat vectors	Specific to SDN?	Consequences in software-defined networks
Vector 1	no	Open door for DDoS attacks.
Vector 2	no	Potential attack inflation.
Vector 3	yes	Exploiting logically centralized controllers.
Vector 4	yes	Compromised controller may compromise the entire network.
Vector 5	yes	Development and deployment of malicious applications on controllers.
Vector 6	no	Potential attack inflation.
Vector 7	no	Negative impact on fast recovery and fault diagnosis.

leges [507], and the assumption that all applications are benign and will not affect SDN operation [194]. The lack of isolation, protection, access control and stronger security recommendations [508], [509], [201], [510], [194] are some of the reasons for these vulnerabilities. We will explore these next.

OpenFlow security assessment

There is already a number of identified security issues in OpenFlow-enabled networks. Starting from a STRIDE methodology [513], it is possible to identify different attacks to OpenFlow-enabled networks. Table XIV summarizes these attacks (based on [507]). For instance, information disclosure can be achieved through side channel attacks targeting the flow rule setup process. When reactive flow setup is in place, obtaining information about network operation is relatively easy. An attacker that measures the delay experienced by the first packet of a flow and the subsequent can easily infer that the target network is a reactive SDN, and proceed with a specialized attack. This attack – known as fingerprinting [509] – may be the first step to launch a DoS attack intended to exhaust the resources of the network, for example. If the SDN is proactive, guessing its forwarding rule policies is harder, but still feasible [507]. Interestingly, all reported threats and attacks affect all versions (1.0 to 1.3.1) of the OpenFlow

specification. It is also worth emphasizing that some attacks, such as spoofing, are not specific to SDN. However, these attacks can have a larger impact in SDNs. For instance, by spoofing the address of the network controller, the attacker (using a fake controller) could take over the control of the entire network. A smart attack could persist for only a few seconds, i.e., just the time needed to install special rules on all forwarding devices for its malicious purposes (e.g., traffic cloning). Such attack could be very hard to detect.

Taking counters falsification as another example, an attacker can try to guess installed flow rules and, subsequently, forge packets to artificially increase the counter. Such attack would be specially critical for billing and load balancing systems, for instance. A customer could be charged for more traffic than she in fact used, while a load balancing algorithm may take non-optimal decisions due to forged counters.

Other conceptual and technical security concerns in OpenFlow networks include the lack of strong security recommendations for developers, the lack of TLS and access control support on most switch and controller implementations [508], the belief that TCP is enough because links are “physically secure” [510], [508], the fact that many switches have listener mode activated by default (allowing the establishment of malicious TCP connections, for instance) [510] or that flow table verification capabilities are harder to implement when TLS is not in use [508], [259]. In addition, it is worth mentioning the high denial of service risk posed to centralized controllers [509], [259], the vulnerabilities in the controllers themselves [259], [357], bugs and vulnerabilities in applications [514], targeted flooding attacks [16], insecure northbound interfaces that can lead to security breaches [16], and the risk of resource depletion attacks [509], [510]. For instance, it has been shown that an attacker can easily compromise control plane communications through DoS attacks and launch a resource depletion attack on control platforms by exploiting a single application such as a learning switch [510], [509].

Another point of concern is the fact that current controllers, such as Floodlight, OpenDaylight, POX, and Beacon, have several security and resiliency issues [194]. Common application development problems (bugs), such as the sudden exit of an application or the continuous allocation of memory space, are enough to crash existing controllers. On the security perspective, a simple malicious action such as changing the value of a data structure in memory can also directly affect the operation and reliability of current controllers. These examples are illustrative that from a security and dependability perspective, there is still a long way to go.

Countermeasures for OpenFlow based SDNs

Several countermeasures can be put in place to mitigate the security threats in SDNs. Table XV summarizes a number of countermeasures that can be applied to different elements of an SDN/OpenFlow-enabled network. Some of these measures, namely rate limiting, event filtering, packet dropping, shorter timeouts, and flow aggregation, are already recommended in the most recent versions of the OpenFlow specification (version 1.3.1 and later). However, most of them are not yet supported or implemented in SDN deployments.

TABLE XIV
ATTACKS TO OPENFLOW NETWORKS.

Attack	Security Property	Examples
Spoofing	Authentication	MAC and IP address spoofing, forged ARP and IPv6 router advertisement
Tampering	Integrity	Counter falsification, rule installation, modification affecting data plane.
Repudiation	Non-repudiation	Rule installation, modification for source address forgery.
Information disclosure	Confidentiality	Side channel attacks to figure out flow rule setup.
Denial of service	Availability	Flow requests overload of the controller.
Elevation of privilege	Authorization	Controller take-over exploiting implementation flaws.

TABLE XV
COUNTERMEASURES FOR SECURITY THREATS IN OPENFLOW NETWORKS.

Measure	Short description
Access control	Provide strong authentication and authorization mechanisms on devices.
Attack detection	Implement techniques for detecting different types of attacks.
Event filtering	Allow (or block) certain types of events to be handled by special devices.
Firewall and IPS	Tools for filtering traffic, which can help to prevent different types of attacks.
Flow aggregation	Coarse-grained rules to match multiple flows to prevent information disclosure and DoS attacks.
Forensics support	Allow reliable storage of traces of network activities to find the root causes of different problems.
Intrusion tolerance	Enable control platforms to maintain correct operation despite intrusions.
Packet dropping	Allow devices to drop packets based on security policy rules or current system load.
Rate limiting	Support rate limit control to avoid DoS attacks on the control plane.
Shorter timeouts	Useful to reduce the impact of an attack that diverts traffic.

Traditional techniques such as access control, attack detection mechanisms, event filtering (e.g., controller decides which asynchronous messages he is not going to accept), firewalls, and intrusion detection systems, can be used to mitigate the impact of or to avoid attacks. They can be implemented in different devices, such as controllers, forwarding devices, middleboxes, and so forth. Middleboxes can be a good option for enforcing security policies in an enterprise because they are (in general) more robust and special purpose (high performance) devices. Such a strategy also reduces the potential overhead cause by implementing these countermeasures directly on controllers or forwarding devices. However, middleboxes can add extra complexity to the network management, i.e., increase the OPEX at the cost of better performance.

Rate limiting, packet dropping, shorter timeouts and flow aggregations are techniques that can be applied on controllers and forwarding devices to mitigate different types of attacks, such as denial-of-service and information disclosure. For instance, reduced timeouts can be used to mitigate the effect of an attack exploring the reactive operation mode of the

network to make the controller install rules that divert traffic to a malicious machine. With reduced timeouts, the attacker would be forced to constantly generate a number of forged packets to avoid timeout expiration, making the attack more likely to be detected. Rate limiting and packet dropping can be applied to avoid DoS attacks on the control plane or stop on-going attacks directly on the data plane by installing specific rules on the devices where the attacks is being originated.

Forensics and remediation encompass mechanisms such as secure logging, event correlation and consistent reporting. If anything wrong happens with the network, operators should be able to safely figure out the root cause of the problem and put the network to work on a secure operation mode as fast as possible. Additionally, techniques to tolerate faults and intrusions, such as state machine replication [515], proactive-reactive recovery [516], and diversity [210], can be added to control platforms for increasing the robustness and security properties by automatically masking and removing faults. Put differently, SDN controllers should be able to resist against different types of events (e.g., power outages, network disruption, communication failures, network partitioning) and attacks (e.g., DDoS, resource exhaustion) [357], [213]. One of the most traditional ways of achieving high availability is through replication. Yet, proactive-reactive recovery and diversity are two examples of crucial techniques that add value to the system for resisting against different kinds of attacks and failures (e.g., those exploring common vulnerabilities or caused by software aging problems).

Other countermeasures to address different threats and issues of SDN include enhancing the security and dependability of controllers, protection and isolation of applications [504], [357], [201], [194], trust management between controllers and forwarding devices [357], integrity checks of controllers and applications [357], forensics and remediation [504], [357], verification frameworks [517], [201], [518], and resilient control planes [519], [518], [357], [504]. Protection and isolation mechanisms should be part of any controller. Applications should be isolated from each other and from the controller. Different techniques such as security domains (e.g., kernel, security, and user level) and data access protection mechanisms should be put in place in order to avoid security threats from network applications.

Implementing trust between controllers and forwarding is another requirement for insuring that malicious elements cannot harm the network without being detected. An attacker can

try to spoof the IP address of the controller and make switches connect to its own controller. This is currently the case since most controllers and switches only establish insecure TCP connections. Complementarily, integrity checks on controller and application software can help to ensure that safe code is being bootstrapped, which eliminates harmful software from being started once the system restarts. Besides integrity checks, other things such as highly specialized malware detection systems should be developed for SDN. Third-party network applications should always be scanned for bad code and vulnerabilities because a malicious application represents a significant security threat to the network.

It is worth mentioning that there are also other approaches for mitigating security threats in SDN, such as declarative languages to eliminate network protocol vulnerabilities [264]. This kind of descriptive languages can specify semantic constraints, structural constraints and safe access properties of OpenFlow messages. Then, a compiler can use these inputs to find programmers' implementation mistakes on message operations. In other words, such languages can help find and eliminate implementation vulnerabilities of southbound specifications.

Proposals providing basic security properties such as authentication [520] and access control [521] are starting to appear. C-BAS [520] is a certificate-based AAA (Authentication, Authorization and Accounting) architecture for improving the security control on SDN experimental facilities. Solutions in the spirit of C-BAS can be made highly secure and dependable through hybrid system architectures, which combine different technologies and techniques from distributed systems, security, and fault and intrusion tolerance [522], [523], [524].

G. Migration and Hybrid deployments

The promises by SDN to deliver easier design, operation and management of computer networks are endangered by challenges regarding incremental deployability, robustness, and scalability. A prime SDN adoption challenge relates to organizational barriers that may arise due to the first (and second) order effects of SDN automation capabilities and "layer/domain blurring". Some level of human resistance is to be expected and may affect the decision and deployment processes of SDN, especially by those that may regard the control refactorization of SDN as a risk to the current chain of control and command, or even to their job security. This complex social challenge is similar (and potentially larger) to known issues between the transport and IP network divisions of service providers, or the system administrator, storage, networking, and security teams of enterprise organizations. Such a challenge is observable on today's virtualized data centers, through the shift in role and decision power between the networking and server people. Similarly, the development and operations (DevOps) movement has caused a shift in the locus of influence, not only on the network architecture but also on purchasing, and this is an effect that SDN may exacerbate. These changes in role and power causes a second order effect on the sales division of vendors that are required to adapt accordingly.

Pioneering SDN operational deployments have been mainly greenfield scenarios and/or tightly controlled single administrative domains. Initial roll-out strategies are mainly based on virtual switch overlay models or OpenFlow-only network-wide controls. However, a broader adoption of SDN beyond data center silos – and between themselves – requires considering the interaction and integration with legacy control planes providing traditional switching; routing; and operation, administration, and management (OAM) functions. Certainly, rip-and-replace is not a viable strategy for the broad adoption of new networking technologies.

Hybrid networking in SDN should allow deploying OpenFlow for a subset of all flows only, enable OpenFlow on a subset of devices and/or ports only, and provide options to interact with existing OAM protocols, legacy devices, and neighboring domains. As in any technology transition period where fork-lift upgrades may not be a choice for many, migration paths are critical for adoption.

Hybrid networking in SDN spans several levels. The Migration Working Group of the ONF is tackling the scenario where hybrid switch architectures and hybrid (OpenFlow and non-OpenFlow) devices co-exist. Hybrid switches can be configured to behave as a legacy switch or as an OpenFlow switch and, in some cases, as both simultaneously. This can be achieved, for example, by partitioning the set of ports of a switch, where one subset is devoted to OpenFlow-controlled networks, and the other subset to legacy networks. For these subsets to be active at the same time, each one having its own data plane, multi-table support at the forwarding engine (e.g., via TCAM partitioning) is required. Besides port-based partitioning, it is also possible to rely on VLAN-based (prior to entering the OpenFlow pipeline) or flow-based partitioning using OpenFlow matching and the LOCAL and/or NORMAL actions to redirect packets to the legacy pipeline or the switch's local networking stack and its management stack. Flow-based partitioning is the most flexible option, as it allows each packet entering a switch to be classified by an OpenFlow flow description and treated by the appropriate data plane (OpenFlow or legacy).

There are diverse controllers, such as OpenDaylight [13], HP VAN SDN [184], and OpenContrail [183], that have been designed to integrate current non-SDN technologies (e.g., SNMP, PCEP, BGP, NETCONF) with SDN interfaces such as OpenFlow and OVSDB. Nonetheless, controllers such as ClosedFlow [219] have been recently proposed with the aim of introducing SDN-like programming capabilities in traditional network infrastructures, making the integration of legacy and SDN-enabled networks a reality without side effects in terms of programmability and global network control. ClosedFlow is designed to control legacy Ethernet devices (e.g., Cisco 3550 switches with a minimum IOS of 12.2 SE) in a similar way an OpenFlow controller allows administrators to control OpenFlow-enabled devices. More importantly, ClosedFlow does not impose any change on the forwarding devices. It only takes advantage of the existing hardware and firmware capabilities to mimic an SDN control over the network, i.e., allow dynamic and flexible programmability in the data plane. The next step could be the integration of controllers

like ClosedFlow and OpenFlow-based controllers, promoting interoperability among controllers and a smooth transition from legacy infrastructures to SDN-enabled infrastructure with nearly all the capabilities of a clean-slate SDN-enabled infrastructure.

Furthermore, controllers may have to be separated into distinct peer domains for different reasons, such as scalability, technology, controllers from different vendors, controllers with different service functionality, and diversity of administrative domains [218]. Controllers from different domains, or with distinct purposes, are also required to be backwards compatible either by retrofitting or extending existing multi-domain protocols (e.g. BGP) or by proposing new SDN-to-SDN protocols (aka east/westbound APIs).

Some efforts have been already devoted to the challenges of migration and hybrid SDNs. RouteFlow [525] implements an IP level control plane on top of an OpenFlow network, allowing the underlying devices to act as IP routers under different possible arrangements. The Cardigan project [526], [50] has deployed RouteFlow at a live Internet eXchange now for over a year. LegacyFlow [527] extends the OpenFlow-based controlled network to embrace non-OpenFlow nodes. There are also some other early use cases on integrating complex legacy system such as DOCSIS [161], Gigabit Ethernet passive optical network and DWDM ROADM (Reconfigurable Optical Add/Drop Multiplexer) [157], [158]. The common grounds of these pieces of work are (1) considering hybrid as the coexistence of traditional environments of closed vendor's routers and switches with new OpenFlow-enabled devices; (2) targeting the interconnection of both control and data planes of legacy and new network elements; and (3) taking a controller-centric approach, drawing the hybrid line outside of any device itself, but into the controller application space.

Panopticon [528] defines an architecture and methodology to consistently implement SDN inside enterprise legacy networks through network orchestration under strict budget constraints. The proposed architecture includes policy configurations, troubleshooting and maintenance tasks establishing transitional networks (SDN and legacy) in structures called Solitary Confinement Trees (SCTs), where VLAN IDs are efficiently used by orchestration algorithms to build paths in order to steer traffic through SDN switches. Defying the partial SDN implementation concept, they confirm that this could be a long-term operational strategy solution for enterprise networks.

HybNET [529] presents a network management framework for hybrid OpenFlow-legacy networks. It provides a common centralized configuration interface to build virtual networks using VLANs. An abstraction of the physical network topology is taken into account by a centralized controller that applies a path finder mechanism, in order to calculate network paths and program the OpenFlow switches via REST interfaces and legacy devices using NETCONF [44].

More recently, frameworks such as ESCAPE [530] and its extensions have been proposed to provide multi-layer service orchestration in multi-domains. Such frameworks combine different tools and technologies such as Click [531], POX [230], OpenDaylight [13] and NETCONF [44]. In other words, those

frameworks integrate different SDN solutions with traditional ones. Therefore, they might be useful tools on the process of integrating or migrating legacy networking infrastructure to SDN.

Other hybrid solutions starting to emerge include Open Source Hybrid IP/SDN (OSHI) [532]. OSHI combines Quagga for OSPF routing and SDN capable switching devices (e.g., Open vSwitch) on Linux to provide backwards compatibility for supporting incremental SDN deployments, i.e., enabling interoperability with non-OF forwarding devices in carrier-grade networks.

While full SDN deployments are straightforward only in some green field deployments such as data center networks or by means of an overlay model approach, hybrid SDN approaches represent a very likely deployment model that can be pursued by different means, including [533]:

- Topology-based hybrid SDN: Based on a topological separation of the nodes controlled by traditional and SDN paradigms. The network is partitioned in different zones and each node belongs to only one zone.
- Service-based hybrid SDN: Conventional networks and SDN provide different services, where overlapping nodes, controlling a different portion of the FIB (or generalized flow table) of each node. Examples include network-wide services like forwarding that can be based on legacy distributed control, while SDN provides edge-to-edge services such as enforcement of traffic engineering and access policies, or services requiring full traffic visibility (e.g., monitoring).
- Class-based hybrid SDN: Based on the partition of traffic in classes, some controlled by SDN and the remaining by legacy protocols. While each paradigm controls a disjoint set of node forwarding entries, each paradigm is responsible for all network services for the assigned traffic classes.
- Integrated hybrid SDN: A model where SDN is responsible for all the network services, and uses traditional protocols (e.g., BGP) as an interface to node FIBs. For example, it can control forwarding paths by injecting carefully selected routes into a routing system or adjusting protocol settings (e.g., IGP weights). Past efforts on RCPs [85] and the ongoing efforts within ODL [13] can be considered examples of this hybrid model.

In general, benefits of hybrid approaches include enabling flexibility (e.g., easy match on packet fields for middleboxing) and SDN-specific features (e.g., declarative management interface) while partially keeping the inherited characteristics of conventional networking such as robustness, scalability, technology maturity, and low deployment costs. On the negative side, the drawbacks of hybridization include the need for ensuring profitable interactions between the networking paradigms (SDN and traditional) while dealing with the heterogeneity that largely depends on the model.

Initial trade-off analyses [533] suggest that the combination of centralized and distributed paradigms may provide mutual benefits. However, future work is required to devise techniques and interaction mechanisms that maximize such benefits while limiting the added complexity of the paradigm coexistence.

H. Meeting carrier-grade and cloud requirements

A number of carrier-grade infrastructure providers (e.g., NTT, AT&T, Verizon, Deutsche Telekom) are at the core of the SDN community with the ultimate goal of solving their long standing networking problems. In the telecom world, NTT can be considered one of the forefront runners in terms of investing in the adoption and deployment of SDN in all kinds of network infrastructures, from backbone, data center, to edge customers [268]. In 2013, NTT launched an SDN-based, on-demand elastic provisioning platform of network resources (e.g., bandwidth) for HD video broadcasters [534]. Similarly, as a global cloud provider with data centers spread across the globe [535], the same company launched a similar service for its cloud customers, who are now capable of taking advantage of dynamic networking provisioning intra- and inter-data centers [536]. AT&T is another telecom company that is investing heavily in new services, such as user-defined network clouds, that take advantage of recent developments in NFV and SDN [537]. As we mentioned before, SDN and NFV are complementary technologies that can be applicable to different types of networks, from local networks and data centers to transport networks [538], [539], [540], [541], [542], [543]. Recently, several research initiatives have worked towards combining SDN and NFV through Intel's DPDK, a set of libraries and drivers that facilitates the development of network-intensive applications and allows the implementation of fine-grained network functions [544]. Early work towards service chaining have been proposed by combining SDN and NFV technologies [545], [546], [547], [27], [548], and studies around the ForCES's [30] applicability to SDN-enhanced NFV have also come to light [538]. These are some of the early examples of the opportunities SDNs seem to bring to telecom and cloud providers.

Carrier networks are using the SDN paradigm as the technology means for solving a number of long standing problems. Some of these efforts include new architectures for a smooth migration from the current mobile core infrastructure to SDN [222], and techno-economic models for virtualization of these networks [549], [550]; carrier-grade OpenFlow virtualization schemes [551], [112], including virtualized broadband access infrastructures [552], techniques that are allowing the offer of network-as-a-service [553]; programmable GE-PON and DWDM ROADM [157], [158], [159], [160]; large-scale inter-autonomous systems (ASs) SDN-enabled deployments [554]; flexible control of network resources [555], including offering MPLS services using an SDN approach [556]; and the investigation of novel network architectures, from proposals to separate the network edge from the core [557], [558], with the latter forming the fabric that transports packets as defined by an intelligent edge, to software-defined Internet exchange points [559], [526].

Use-case analysis [560] of management functions required by carrier networks have identified a set of requirements and existing limitations in the SDN protocol toolbox. For instance, it has been pinpointed that OF-Config [54] needs a few extensions in order to meet the carrier requirements, such as physical resource discovery, logical link configuration,

logical switch instantiation, and device and link OAM configuration [560]. Similarly, OpenFlow extensions have also been proposed to realize packet-optical integration with SDN [561]. In order to support SDN concepts in large-scale wide area networks, different extensions and mechanisms are required, both technology-specific (e.g., MPLS BFD) and technology agnostic, such as: resiliency mechanisms for surviving link failures [484], failures of controller or forwarding elements; solutions for integrating residential customer services in different forms (i.e., support also current technologies); new energy-efficient networking approaches; QoS properties for packet classification, metering, coloring, policing, shaping and scheduling; and multi-layer aspects outlining different stages of packet-optical integration [562], [563], [561].

SDN technology also brings new possibilities for cloud providers. By taking advantage of the logically centralized control of network resources [564], [8] it is possible to simplify and optimize network management of data centers and achieve: (i) efficient intra-datacenter networking, including fast recovery mechanisms for the data and control planes [476], [565], [566], adaptive traffic engineering with minimal modifications to DCs networks [277], simplified fault-tolerant routing [567], performance isolation [568], and easy and efficient resource migration (e.g., of VMs and virtual networks) [476]; (ii) improved inter-datacenter communication, including the ability to fully utilize the expensive high-bandwidth links without impairing quality of service [8], [569]; (iii) higher levels of reliability (with novel fault management mechanisms, etc.) [567], [565], [476], [484]; and (iv) cost reduction by replacing complex, expensive hardware by simple and cheaper forwarding devices [570], [8].

Table XVI summarizes some of the carrier-grade network and cloud infrastructure providers' requirements. In this table we show the current challenges and what is to be expected with SDN. As we saw before, some of the expectations are already becoming a reality, but many are still open issues. What seems to be clear is that SDN represents an opportunity for telecom and cloud providers, in providing flexibility, cost-effectiveness, and easier management of their networks.

I. SDN: the missing piece towards Software-Defined Environments

The convergence of different technologies is enabling the emergence of fully programmable IT infrastructures. It is already possible to dynamically and automatically configure or reconfigure the entire IT stack, from the network infrastructure up to the applications, to better respond to workload changes. Recent advances makes on-demand provisioning of resources possible, at nearly all infrastructural layers. The fully automated provisioning and orchestration of IT infrastructures as been recently named Software-Defined Environments (SDEs) [171], [172], by IBM. This is a novel approach that is expected to have significant potential in simplifying IT management, optimizing the use of the infrastructure, reduce costs, and reduce the time to market of new ideas and products. In an SDE, workloads can be easily and automatically assigned to the appropriate IT resources based on application characteristics, security and service level policies, and the best-available

TABLE XVI
CARRIER-GRADE AND CLOUD PROVIDER EXPECTATIONS & CHALLENGES

What	Currently	Expected with SDN
Resource Provisioning	Complex load balancing configuration.	Automatic load balancing reconfiguration. [571], [8]
	Low virtualization capabilities across hardware platforms	NFV for virtualizing network functionality across hardware appliances. [570], [537]
	Hard and costly to provide new services.	Create and deploy new network service quickly. [570], [537]
	No bandwidth on demand.	Automatic bandwidth on demand. [550]
	Per network element scaling.	Better incremental scaling. [571], [565]
	Resources statically pre-provisioned.	Dynamic resource provisioning in response to load. [571], [8], [570], [549], [564]
Traffic Steering	All traffic is filtered.	Only targeted traffic is filtered. [571]
	Fixed only.	Fixed and mobile. [571]
	Per network element scaling.	Better incremental scaling. [549], [565]
	Statically configured on a per-device basis.	Dynamically configurable. [8], [550], [572]
Ad Hoc Topologies	All traffic from all probes collected.	Only targeted traffic from targeted probes is collected.
	Massive bandwidth required.	Efficient use of bandwidth. [8], [550]
	Per network element scaling.	Better incremental scaling. [571], [550]
	Statically configured.	Dynamically configured. [571], [573], [534]
Managed Router Services	Complex configuration, management and upgrade.	Simplified management and upgrade. [8], [571], [570], [550], [565]
	Different kinds of routers, such as changeover (CO).	No need for CO routers, reducing aggregation costs. [571], [570], [549]
	Manual provisioning.	Automated provisioning. [571], [550], [572]
	On-premises router deployment.	Virtual routers (either on-site or not). [550], [571], [549]
	Operational burden to support different equipments.	Reduced technology obsolescence. [549]
	Router change-out as technology or needs change.	Pay-as-you grow CAPEX model. [549]
	Systems complex and hard to integrate.	Facilitates simplified system integrations. [571], [570], [573]
Revenue Models	Fixed long term contracts.	More flexible and on-demand contracts. [550], [555]
	Traffic consumption.	QoS metrics per-application. [550], [565], [565], [574]
Middleboxes Deployment & Management	Composition of services is hard to implement.	Easily expand functionality to meet the infrastructure needs. [570]
	Determine where to place middleboxes a priori (e.g., large path inflation problems).	Dynamic placement using shortest or least congested path. [290], [574], [573]
	Excessive over-provisioning to anticipate demands.	Scale up to meet demands, and scale down to conserve resources (elastic middleboxes). [571], [549]
Other Issues	Energy saving strategies are hard to implement.	Flexible and easy to deploy energy saving strategies. [565]
	Complex and static control and data plane restoration techniques.	Automated and flexible restoration techniques for both control and data plane. [565]

resources to deliver continuous, dynamic optimization and reconfiguration to address infrastructure issues in a rapid and responsive manner. Table XVII summarizes the traditional approaches and some of the key features being enabled by SDEs [575], [576].

In an SDE the workloads are managed independently of the systems and underlying infrastructure, i.e., are not tied to a specific technology or vendor [172], [171]. Another characteristic of this new approach is to offer a programmatic access to the environment as a whole, selecting the best available resources based on the current status of the infrastructure, and enforcing the policies defined. In this sense, it shares much of the philosophy of SDN. Interestingly, one of the missing key pieces of an SDE was, until now, Software-Defined Networking.

The four essential building blocks of an SDE [172], [171], [576] are:

- Software-Defined Networks (SDN) [577], [578],

- Software-Defined Storage (SDS) [575],
- Software-Defined Compute (SDC) [171], and
- Software-Defined Management (SDM) [579].

In the last decade the advances in virtualization of compute and storage, together with the availability of sophisticated cloud orchestration tools have enabled SDS, SDC and SDM. These architectural components have been widely used by cloud providers and for building IT infrastructures in different enterprise environments. However, the lack of programmable network control has so far hindered the realization of a complete Software-Defined Environment. SDN is seen as the technology that may fill this gap, as attested by the emergence of cloud-scale network virtualization platforms based on this new paradigm [112].

The IBM SmartCloud Orchestrator is one of the first examples of an SDE [172], [171]. It integrates compute, storage, management and networking in a structured way. Figure 11 gives a simplified overview of an SDE, by taking

TABLE XVII
SDE PUSHING IT TO THE NEXT FRONTIER

Traditionally	Expected with SDEs
IT operations manually map the resources for apps for software deployment.	Software maps resources to the workload and deploys the workload.
Networks are mostly statically configured and hard to change.	Networks are virtualized and dynamically configured on-demand.
Optimization and reconfiguration to reactively address issues are manual.	Analytics-based optimization and reconfiguration of infrastructure issues.
Workloads are typically manually assigned to resources.	Workloads are dynamically assigned.




Fig. 11. Overview of an IT infrastructure based on a SDE.

the approach developed by IBM as its basis. The main idea of an SDE-based infrastructure is that the business needs that define the workloads trigger the reconfiguration of the global IT infrastructure (compute, storage, network). This is an important step towards a more customizable IT infrastructure that focuses on the business requirements rather than on the limitations of the infrastructure itself.

VI. CONCLUSION

Traditional networks are complex and hard to manage. One of the reasons is that the control and data planes are vertically integrated and vendor specific. Another, concurring reason, is that typical networking devices are also tightly tied to line products and versions. In other words, each line of product may have its own particular configuration and management interfaces, implying long cycles for producing product updates (e.g., new firmware) or upgrades (e.g., new versions of the devices). All this has given rise to vendor lock-in problems for network infrastructure owners, as well as posing severe restrictions to change and innovation.

Software-Defined Networking (SDN) created an opportunity for solving these long-standing problems. Some of the key ideas of SDN are the introduction of dynamic programmability in forwarding devices through open southbound interfaces, the decoupling of the control and data plane, and the global view of the network by logical centralization of the “network brain”. While data plane elements became dumb, but highly efficient and programmable packet forwarding devices, the control plane elements are now represented by a single entity, the controller or network operating system. Applications implementing the network logic run on top of the controller and are much easier to develop and deploy when compared to traditional networks. Given the global view, consistency

of policies is straightforward to enforce. SDN represents a major paradigm shift in the development and evolution of networks, introducing a new pace of innovation in networking infrastructure.

In spite of recent and interesting attempts to survey this new chapter in the history of networks [14], [16], [15], the literature was still lacking, to the best of our knowledge, a single extensive and comprehensive overview of the building blocks, concepts, and challenges of SDNs. Trying to address this gap, the present paper used a layered approach to methodically dissect the state of the art in terms of concepts, ideas and components of software-defined networking, covering a broad range of existing solutions, as well as future directions.

We started by comparing this new paradigm with traditional networks and discussing how academy and industry helped shape software-defined networking. Following a bottom-up approach, we provided an in-depth overview of what we consider the eight fundamental facets of the SDN problem: 1) hardware infrastructure, 2) southbound interfaces, 3) network virtualization (hypervisor layer between the forwarding devices and the network operating systems), 4) network operating systems (SDN controllers and control platforms), 5) northbound interfaces (common programming abstractions offered to network applications), 6) virtualization using slicing techniques provided by special purpose libraries and/or programming languages and compilers, 7) network programming languages, and finally, 8) network applications.

SDN has successfully managed to pave the way towards a next generation networking, spawning an innovative research and development environment, promoting advances in several areas: switch and controller platform design, evolution of scalability and performance of devices and architectures, promotion of security and dependability.

We will continue to witness extensive activity around SDN in the near future. Emerging topics requiring further research are, for example: the migration path to SDN, extending SDN towards carrier transport networks, realization of the network-as-a-service cloud computing paradigm, or software-defined environments (SDE). As such, we would like to receive feedback from the networking/SDN community as this novel paradigm evolves, to make this a “live document” that gets updated and improved based on the community feedback. We have set up a github page² for this purpose, and we invite our readers to join us in this communal effort.

²<https://github.com/SDN-Survey/latex/wiki>

ACKNOWLEDGMENT

The authors would like to thank the anonymous reviewers and a number of fellows that have contributed to this work. Jennifer Rexford for her feedback on an early version of this work and encouragement to get it finished. Srini Seetharaman for reviewing the draft and providing inputs to alternative SDN views. David Meyer for his thoughts on organizational challenges. Thomas Nadeau for his inputs on OpenDaylight. Luis Miguel Contreras Murillo for his contributions to SDN standardization. In addition, we would like also to acknowledge the several contributions from the community, namely from Aurel A. Lazar, Carmelo Cascone, Gyanesh Patra, Haleplidis Evangelos, Javier Ancieta, Joe Stringer, Kostas Pentikousis, Luciano de Paula, Marco Canini, Philip Wette, Ramon Fontes, Raphael Rosa, Regivaldo Costa, Ricardo de Freitas Gesuatto, Wolfgang John.

REFERENCES

- [1] T. Benson, A. Akella, and D. Maltz, "Unraveling the complexity of network management," in *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, ser. NSDI'09, Berkeley, CA, USA, 2009, pp. 335–348.
- [2] B. Raghavan, M. Casado, T. Koponen, S. Ratnasamy, A. Ghodsi, and S. Shenker, "Software-defined internet architecture: Decoupling architecture from infrastructure," in *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, ser. HotNets-XI. New York, NY, USA: ACM, 2012, pp. 43–48.
- [3] A. Ghodsi, S. Shenker, T. Koponen, A. Singla, B. Raghavan, and J. Wilcox, "Intelligent design enables architectural evolution," in *Proceedings of the 10th ACM Workshop on Hot Topics in Networks*, ser. HotNets-X. New York, NY, USA: ACM, 2011, pp. 3:1–3:6.
- [4] N. McKeown, "How SDN will Shape Networking," October 2011. [Online]. Available: http://www.youtube.com/watch?v=c9-K5O_qYgA
- [5] S. Schenker, "The Future of Networking, and the Past of Protocols," October 2011. [Online]. Available: <http://www.youtube.com/watch?v=YHeyuD89n1Y>
- [6] H. Kim and N. Feamster, "Improving network management with software defined networking," *Communications Magazine, IEEE*, vol. 51, no. 2, pp. 114–119, 2013.
- [7] T. Koponen, M. Casado, N. Gude, J. Stribling, L. Poutievski, M. Zhu, R. Ramanathan, Y. Iwata, H. Inoue, T. Hama, and S. Shenker, "Onix: a distributed control platform for large-scale production networks," in *Proceedings of the 9th USENIX conference on Operating systems design and implementation*, ser. OSDI'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 1–6.
- [8] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat, "B4: experience with a globally-deployed software defined wan," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 3–14.
- [9] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008.
- [10] ONF, "Open networking foundation," 2014. [Online]. Available: <https://www.opennetworking.org/>
- [11] S. Yeganeh, A. Tootoonchian, and Y. Ganjali, "On scalability of software-defined networking," *Communications Magazine, IEEE*, vol. 51, no. 2, pp. 136–141, 2013.
- [12] VMware, Inc., "VMware NSX Virtualization Platform," 2013. [Online]. Available: <https://www.vmware.com/products/nsx/>
- [13] OpenDaylight, "OpenDaylight: A Linux Foundation Collaborative Project," 2013. [Online]. Available: <http://www.opendaylight.org>
- [14] A. Lara, A. Kolasani, and B. Ramamurthy, "Network innovation using OpenFlow: A survey," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 1, pp. 493–512, First 2014.
- [15] B. Nunes, M. Mendonca, X.-N. Nguyen, K. Obraczka, and T. Turletti, "A survey of software-defined networking: Past, present, and future of programmable networks," *Communications Surveys Tutorials, IEEE*, vol. 16, no. 3, pp. 1617–1634, Third 2014.
- [16] Y. Jarraya, T. Madi, and M. Debbabi, "A survey and a layered taxonomy of software-defined networking," *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, pp. 1–1, 2014.
- [17] N. Feamster, J. Rexford, and E. Zegura, "The road to SDN," *Queue*, vol. 11, no. 12, pp. 20:20–20:40, Dec. 2013.
- [18] R. Presuhn, "Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)," RFC 3416 (INTERNET STANDARD), Internet Engineering Task Force, Dec. 2002. [Online]. Available: <http://www.ietf.org/rfc/rfc3416.txt>
- [19] J. Pan, S. Paul, and R. Jain, "A survey of the research on future internet architectures," *Communications Magazine, IEEE*, vol. 49, no. 7, pp. 26–36, July 2011.
- [20] N. Feamster and H. Balakrishnan, "Detecting BGP configuration faults with static analysis," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, ser. NSDI'05. Berkeley, CA, USA: USENIX Association, 2005, pp. 43–56.
- [21] R. Barrett, S. Haar, and R. Whitestone, "Routing snafu causes internet outage," *Interactive Week*, 1997.
- [22] K. Butler, T. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security issues and solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, Jan 2010.
- [23] J. Sherry and S. Ratnasamy, "A survey of enterprise middlebox deployments," EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2012-24, Feb 2012.
- [24] K. Greene, "MIT Tech Review 10 Breakthrough Technologies: Software-defined Networking," <http://www2.technologyreview.com/article/412194/tr10-software-defined-networking/>, 2009.
- [25] P. Newman, G. Minshall, and T. L. Lyon, "Ip switching—atm under ip," *IEEE/ACM Trans. Netw.*, vol. 6, no. 2, pp. 117–129, Apr. 1998.
- [26] N. Gude, T. Koponen, J. Pettit, B. Pfaff, M. Casado, N. McKeown, and S. Shenker, "NOX: towards an operating system for networks," *Comm. Rev.*, 2008.
- [27] H. Jamjoom, D. Williams, and U. Sharma, "Don't call them middleboxes, call them middlepipes," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 19–24.
- [28] H. Alkhatib, P. Faraboschi, E. Frachtenberg, H. Kasahara, D. Lange, P. Laplante, A. Merchant, D. Milojevic, and K. Schwan, "IEEE CS 2022 report (draft)," IEEE Computer Society, Tech. Rep., February 2014.
- [29] M. Casado, N. Foster, and A. Guha, "Abstractions for software-defined networks," *Commun. ACM*, vol. 57, no. 10, pp. 86–95, Sep. 2014.
- [30] A. Doria, J. H. Salim, R. Haas, H. Khosravi, W. Wang, L. Dong, R. Gopal, and J. Halpern, "Forwarding and Control Element Separation (ForCES) Protocol Specification," Internet Engineering Task Force, Mar. 2010. [Online]. Available: <http://www.ietf.org/rfc/rfc5810.txt>
- [31] H. Song, "Protocol-oblivious Forwarding: Unleash the power of SDN through a future-proof forwarding plane," in *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 127–132.
- [32] T. D. Nadeau and K. Gray, *SDN: software defined networks*, 1st ed. O'Reilly, 2013, vol. 1.
- [33] N. M. K. Chowdhury and R. Boutaba, "A survey of network virtualization," *Computer Networks*, vol. 54, no. 5, pp. 862 – 876, 2010.
- [34] B. Davie and J. Gross, "A Stateless Transport Tunneling Protocol for Network Virtualization (STT)," Internet Engineering Task Force, April 2014. [Online]. Available: <http://tools.ietf.org/html/draft-davie-stt-06>
- [35] M. Mahalingam, D. Dutt, K. Duda, P. Agarwal, L. Kreeger, T. Sridhar, M. Bursell, and C. Wright, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks," Internet Draft, Internet Engineering Task Force, November 2013. [Online]. Available: <http://www.ietf.org/id/draft-mahalingam-dutt-dcops-vxlan-06.txt>
- [36] M. Sridharan, A. Greenberg, Y. Wang, P. Garg, N. Venkataramiah, K. Duda, I. Ganga, G. Lin, M. Pearson, P. Thaler, and C. Tumuluri, "NVGRE: Network Virtualization using Generic Routing Encapsulation," Internet Draft, Internet Engineering Task Force, August 2013. [Online]. Available: <http://tools.ietf.org/id/draft-sridharan-virtualization-nvgre-03.txt>
- [37] F. Maino, V. Ermagan, Y. Hertoghs, D. Farinacci, and M. Smith, "LISP Control Plane for Network Virtualization Overlays," Internet Engineering Task Force, Oct 2013. [Online]. Available: <http://tools.ietf.org/html/draft-maino-nvo3-lisp-cp-03>
- [38] Y. Hertoghs, F. Maino, V. Moreno, M. Smith, D. Farinacci, and L. Iannone, "A Unified LISP Mapping Database for L2 and L3 Network Virtualization Overlays," Internet Engineering Task Force,

- Feb 2014. [Online]. Available: <http://tools.ietf.org/html/draft-herthoghs-nvo3-lisp-controlplane-unified-01>
- [39] J. Gross, T. Sridhar, P. Garg, C. Wright, and I. Ganga, "Geneve: Generic network virtualization encapsulation," Internet Draft, Internet Engineering Task Force, February 2014. [Online]. Available: <http://tools.ietf.org/id/draft-gross-geneve-00.txt>
- [40] R. Jain and S. Paul, "Network virtualization and software defined networking for cloud computing: a survey," *Communications Magazine, IEEE*, vol. 51, no. 11, pp. 24–31, 2013.
- [41] E. Haleplidis, S. Denazis, K. Pentikousis, S. Denazis, J. H. Salim, D. Meyer, and O. Koufopavlou, "SDN Layers and Architecture Terminology," Internet Draft, Internet Engineering Task Force, September 2014. [Online]. Available: <http://www.ietf.org/id/draft-irtf-sdnrg-layer-terminology-02.txt>
- [42] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)," RFC 4271 (Draft Standard), Internet Engineering Task Force, Jan. 2006. [Online]. Available: <http://www.ietf.org/rfc/rfc4271.txt>
- [43] J. Vasseur and J. L. Roux, "Path Computation Element (PCE) Communication Protocol (NETCONF)," RFC 5440 (Proposed Standard), Internet Engineering Task Force, Mar. 2009. [Online]. Available: <http://www.ietf.org/rfc/rfc5440.txt>
- [44] R. Enns, M. Bjorklund, J. Schoenwaelder, and A. Bierman, "Network Configuration Protocol (NETCONF)," RFC 6241 (Proposed Standard), Internet Engineering Task Force, Jun. 2011. [Online]. Available: <http://www.ietf.org/rfc/rfc6241.txt>
- [45] A. Corradi, M. Fanelli, and L. Foschini, "VM consolidation: A real case based on openstack cloud," *Future Generation Computer Systems*, vol. 32, no. 0, pp. 118 – 127, 2014.
- [46] A. Shang, J. Liao, and L. Du, "Pica8 Xorplus," 2014. [Online]. Available: <http://sourceforge.net/projects/xorplus/>
- [47] P. Jakma and D. Lamparter, "Introduction to the quagga routing suite," *Network, IEEE*, vol. 28, no. 2, pp. 42–48, March 2014.
- [48] "NetFPGA," 2014. [Online]. Available: <http://netfpga.org/>
- [49] Linux Foundation, "Open platform for NFV," <https://www.opnfv.org>, Sep 2014.
- [50] C. E. Rothenberg, R. Chua, J. Bailey, M. Winter, C. Correa, S. Lucena, and M. Salvador, "When open source meets network control planes," *IEEE Computer Special Issue on Software-Defined Networking*, November 2014.
- [51] ONF, "SDN architecture," Open Networking Foundation, Tech. Rep., June 2014. [Online]. Available: https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR_SDN_ARCH_1.0_06062014.pdf
- [52] ONF, "Conformance test specification for OpenFlow switch specification," Open Networking Foundation, Tech. Rep., June 2013. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-test/conformance-test-spec-openflow-1.0.1.pdf>
- [53] ONF, "OpenFlow switch specification," Open Networking Foundation, Tech. Rep., October 2013. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/openflow-spec-v1.4.0.pdf>
- [54] ONF, "OpenFlow management and configuration protocol (OF-CONFIG) v1.2," Open Networking Foundation, Tech. Rep., 2014. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-config/of-config-1.2.pdf>
- [55] ONF, "OpenFlow notifications framework OpenFlow management," October 2013. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-config/of-notifications-framework-1.0.pdf>
- [56] ONF, "OpenFlow table type patterns," Open Networking Foundation, Tech. Rep., August 2014. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/OpenFlow%20Table%20Type%20Patterns%20v1.0.pdf>
- [57] ONF, "Optical transport use cases," Open Networking Foundation, Tech. Rep., 2014. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/optical-transport-use-cases.pdf>
- [58] ONF, "Requirements analysis for transport OpenFlow/SDN," Open Networking Foundation, Tech. Rep., August 2014. [Online]. Available: https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/onf2014.227_Optical_Transport_Requirements.pdf
- [59] ONF - Migration Working Group, "Migration use cases and methods," Open Networking Foundation, Tech. Rep., 2014. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/use-cases/Migration-WG-Use-Cases.pdf>
- [60] ONF, "Software-defined networking: The new norm for networks," Open Networking Foundation, Tech. Rep., April 2012. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/white-papers/wp-sdn-newnorm.pdf>
- [61] H. Xie, T. Tsou, D. Lopez, and H. Yin, "Use cases for ALTO with software defined networks," Internet Draft, Internet Engineering Task Force, June 2012. [Online]. Available: <http://tools.ietf.org/html/draft-xie-alto-sdn-use-cases-01>
- [62] A. Atlas, J. Halpern, S. Hares, D. Ward, and T. Nadeau, "An architecture for the interface to the routing system," Internet Draft, Internet Engineering Task Force, July 2014. [Online]. Available: <https://tools.ietf.org/html/draft-ietf-i2rs-architecture-05>
- [63] R. Enns, "NETCONF configuration protocol," Internet Draft, Internet Engineering Task Force, December 2004. [Online]. Available: <http://tools.ietf.org/html/rfc4741>
- [64] L. Kreeger, D. Dutt, T. Narten, and D. Black, "Network virtualization NVE to NVA control protocol requirements," Internet Draft, Internet Engineering Task Force, April 2014. [Online]. Available: <http://tools.ietf.org/html/draft-ietf-nvo3-nve-nva-cp-req-02>
- [65] D. King and A. Farrel, "A PCE-based architecture for application-based network operations," Internet Draft, Internet Engineering Task Force, August 2014. [Online]. Available: <https://tools.ietf.org/html/draft-farrkingel-pce-abno-architecture-11>
- [66] D. Dhody, L. C. Y. Lee, O. Gonzalez, and N. Ciulli, "Cross stratum optimization enabled path computation," Internet Draft, Internet Engineering Task Force, July 2014. [Online]. Available: <http://tools.ietf.org/html/draft-dhody-pce-cso-enabled-path-computation-06>
- [67] B. K. F. Hu and H. Cankaya, "SPRING OpenFlow interworking requirements," Internet Draft, Internet Engineering Task Force, September 2014. [Online]. Available: <https://tools.ietf.org/html/draft-khc-spring-openflow-interworking-req-00>
- [68] E. P. S. Kim, J. Park and L. Contreras, "SPRING use cases for software-defined networking," Internet Draft, Internet Engineering Task Force, July 2014. [Online]. Available: <http://tools.ietf.org/html/draft-kim-spring-use-cases-00>
- [69] D. Ceccarelli, L. Fang, Y. Lee, and D. Lopez, "Framework for abstraction and control of transport networks," Internet Draft, Internet Engineering Task Force, February 2014. [Online]. Available: <https://xml.resource.org/html/draft-ceccarelli-actn-framework-01>
- [70] M. Boucadair and C. Jacquenet, "Software-defined networking: A perspective from within a service provider environment," Internet Draft, Internet Engineering Task Force, March 2014. [Online]. Available: <https://tools.ietf.org/html/rfc7149>
- [71] E. Haleplidis, K. Pentikousis, S. Denazis, J. Salem, D. Meyer, and O. Koufopavlou, "SDN layers and architecture terminology," Internet Draft, Internet Engineering Task Force, August 2014. [Online]. Available: <http://tools.ietf.org/html/draft-haleplidis-sdnrg-layer-terminology-07>
- [72] C. B. L. Contreras and D. Lopez, "Cooperating layered architecture for SDN," Internet Draft, Internet Engineering Task Force, August 2014. [Online]. Available: <http://tools.ietf.org/html/draft-contreras-sdnrg-layered-sdn-00>
- [73] Y. Cheng and C. Zhou, "Framework of signalling for SDN - working document," ITU-T, Tech. Rep., July 2014, work item: Q.Supplement-SDN. [Online]. Available: <http://www.itu.int>
- [74] ITU-T SG 11, "Scenarios and signalling requirements for software-defined BAN (SBAN) - working document," ITU-T, Tech. Rep., July 2014, work item: Q.SBAN. [Online]. Available: <http://www.itu.int>
- [75] ITU-T, "Framework of software-defined networking," ITU-T, Tech. Rep., June 2014, recommendation ITU-T Y.3300. [Online]. Available: <http://www.itu.int/rec/T-REC-Y.3300-201406-I/en>
- [76] Broadband Forum SD-313, "High level requirements and framework for SDN in telecommunication broadband networks," Broadband Forum, Tech. Rep., September 2014. [Online]. Available: <http://www.broadband-forum.org/technical/technicalwip.php>
- [77] OIF, "Requirements on transport networks in SDN architectures," Optical Interworking Forum, Tech. Rep., September 2013. [Online]. Available: http://www.oiforum.com/public/documents/OIF_Carrier_WG_Requirements_on_Transport_Networks_in_SDN_Architectures_Sept2013.pdf
- [78] Open Data Center Alliance, "Open Data Center Alliance: Software-Defined Networking Rev. 2.0," Open Data Center Alliance, Tech.

- Rep., 2014. [Online]. Available: http://www.opendatacenteralliance.org/docs/software_defined_networking_master_usage_model_rev2.pdf
- [79] ETSI GS NFV 002, "Network functions virtualization (NFV); architectural framework v1.1.1," ETSI, Tech. Rep., October 2013. [Online]. Available: http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.01.01_60/gs_NFV002v010101p.pdf
- [80] ATIS, "Operational opportunities and challenges of SDN/NFV programmable infrastructure," ATIS-I-0000044, Tech. Rep., October 2014. [Online]. Available: <https://www.atis.org/docstore/product.aspx?id=28143>
- [81] D. Tennenhouse, J. Smith, W. Sincoskie, D. Wetherall, and G. Minden, "A survey of active network research," *Communications Magazine, IEEE*, vol. 35, no. 1, pp. 80–86, 1997.
- [82] A. Lazar, K.-S. Lim, and F. Marconcini, "Realizing a foundation for programmability of atm networks with the binding architecture," *Selected Areas in Communications, IEEE Journal on*, vol. 14, no. 7, pp. 1214–1227, Sep 1996.
- [83] A. Lazar, "Programming telecommunication networks," *Network, IEEE*, vol. 11, no. 5, pp. 8–18, Sep 1997.
- [84] D. Sheinbein and R. P. Weber, "800 service using SPC network capability," *The Bell System Technical Journal*, vol. 61, no. 7, Sep. 1982.
- [85] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe, "Design and implementation of a routing control platform," in *Proceedings of the 2nd conference on Symposium on Networked Systems Design & Implementation - Volume 2*, ser. NSDI'05. Berkeley, CA, USA: USENIX Association, 2005, pp. 15–28.
- [86] J. Biswas, A. A. Lazar, J. F. Huard, K. Lim, S. Mahjoub, L. F. Pau, M. Suzuki, S. Torstensson, W. Wang, and S. Weinstein, "The IEEE P1520 standards initiative for programmable network interfaces," *Comm. Mag.*, vol. 36, no. 10, pp. 64–70, Oct. 1998.
- [87] B. Schwartz, A. Jackson, W. Strayer, W. Zhou, R. Rockwell, and C. Partridge, "Smart packets for active networks," in *Open Architectures and Network Programming Proceedings, 1999. OPENARCH'99. 1999 IEEE Second Conference on*, Mar 1999, pp. 90–97.
- [88] D. Wetherall, J. V. Guttag, and D. Tennenhouse, "Ants: a toolkit for building and dynamically deploying network protocols," in *Open Architectures and Network Programming, 1998 IEEE*, Apr 1998, pp. 117–129.
- [89] D. Alexander, W. Arbaugh, M. Hicks, P. Kakkar, A. Keromytis, J. Moore, C. Gunter, S. Nettles, and J. Smith, "The switchware active network architecture," *Network, IEEE*, vol. 12, no. 3, pp. 29–36, May 1998.
- [90] K. Calvert, S. Bhattacharjee, E. Zegura, and J. Sterbenz, "Directions in active networks," *Communications Magazine, IEEE*, vol. 36, no. 10, pp. 72–78, Oct 1998.
- [91] T. Wolf and J. Turner, "Design issues for high performance active routers," in *Broadband Communications, 2000. Proceedings. 2000 International Zurich Seminar on*, 2000, pp. 199–205.
- [92] S. da Silva, Y. Yemini, and D. Florissi, "The NetScript active network system," *IEEE J.Sel. A. Commun.*, vol. 19, no. 3, pp. 538–551, Mar. 2001.
- [93] D. L. Tennenhouse and D. J. Wetherall, "Towards an active network architecture," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 5, pp. 81–94, Oct. 2007.
- [94] P. Newman, W. Edwards, R. Hinden, E. Hoffman, F. C. Liaw, T. Lyon, and G. Minshall, "Ipsilon's General Switch Management Protocol Specification Version 1.1," RFC 1987 (Informational), Internet Engineering Task Force, Aug. 1996, updated by RFC 2297. [Online]. Available: <http://www.ietf.org/rfc/rfc1987.txt>
- [95] A. Doria and K. Sundell, "General Switch Management Protocol (GSMP) Applicability," RFC 3294 (Informational), Internet Engineering Task Force, Jun. 2002. [Online]. Available: <http://www.ietf.org/rfc/rfc3294.txt>
- [96] J. Van der Merwe, S. Rooney, I. Leslie, and S. Crosby, "The tempest-a practical framework for network programmability," *Network, IEEE*, vol. 12, no. 3, pp. 20–28, May 1998.
- [97] T. Lakshman, T. Nandagopal, R. Ramjee, K. Sabnani, and T. Woo, "The SoftRouter Architecture," in *Third ACM Workshop on Hot Topics in Networks (HotNets-III)*, San Diego, CA, November 2004.
- [98] A. Greenberg, G. Hjalmtysson, D. A. Maltz, A. Myers, J. Rexford, G. Xie, H. Yan, J. Zhan, and H. Zhang, "A clean slate 4D approach to network control and management," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 41–54, Oct. 2005.
- [99] J. Van der Merwe, A. Cepleanu, K. D'Souza, B. Freeman, A. Greenberg, D. Knight, R. McMillan, D. Moloney, J. Mulligan, H. Nguyen, M. Nguyen, A. Ramarajan, S. Saad, M. Satterlee, T. Spencer, D. Toll, and S. Zeligner, "Dynamic connectivity management with an intelligent route service control point," in *Proceedings of the SIGCOMM workshop on Internet network management*, ser. INM '06. New York, NY, USA: ACM, 2006, pp. 29–34.
- [100] M. Casado, T. Garfinkel, A. Akella, M. J. Freedman, D. Boneh, N. McKeown, and S. Shenker, "SANE: a protection architecture for enterprise networks," in *Proceedings of the 15th conference on USENIX Security Symposium - Volume 15*, ser. USENIX-SS'06, Berkeley, CA, USA, 2006.
- [101] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker, "Ethane: taking control of the enterprise," in *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM '07. New York, NY, USA: ACM, 2007, pp. 1–12.
- [102] M. Macedonia and D. Brutzman, "Mbone provides audio and video across the internet," *Computer*, vol. 27, no. 4, pp. 30–36, 1994.
- [103] R. Fink and R. Hinden, "6bone (IPv6 Testing Address Allocation) Phaseout," RFC 3701 (Informational), Internet Engineering Task Force, Mar. 2004. [Online]. Available: <http://www.ietf.org/rfc/rfc3701.txt>
- [104] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris, "Resilient overlay networks," *SIGOPS Oper. Syst. Rev.*, vol. 35, no. 5, pp. 131–145, Oct. 2001.
- [105] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman, "Planetlab: An overlay testbed for broad-coverage services," *SIGCOMM Comput. Commun. Rev.*, vol. 33, no. 3, pp. 3–12, Jul. 2003.
- [106] T. Anderson, L. Peterson, S. Shenker, and J. Turner, "Overcoming the internet impasse through virtualization," *Computer*, vol. 38, no. 4, pp. 34–41, Apr. 2005.
- [107] L. Peterson, T. Anderson, D. Blumenthal, D. Casey, D. Clark, D. Estrin, J. Evans, D. Raychaudhuri, M. Reiter, J. Rexford, S. Shenker, and J. Wroclawski, "Geni design principles," *Computer*, vol. 39, no. 9, pp. 102–105, Sept 2006.
- [108] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford, "In VINI veritas: realistic and controlled network experimentation," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 4, pp. 3–14, Aug. 2006.
- [109] B. Pfaff, J. Pettit, T. Koponen, K. Amidon, M. Casado, and S. Shenker, "Extending networking into the virtualization layer," in *Proc. of workshop on Hot Topics in Networks (HotNets-VIII)*, 2009.
- [110] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: rapid prototyping for software-defined networks," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, ser. Hotnets-IX. New York, NY, USA: ACM, 2010, pp. 19:1–19:6.
- [111] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar, "Can the production network be the testbed?" in *Proceedings of the 9th USENIX conference on Operating systems design and implementation*, ser. OSDI'10, Berkeley, CA, USA, 2010, pp. 1–6.
- [112] T. Koponen, K. Amidon, P. Balland, M. Casado, A. Chanda, B. Fulton, I. Ganichev, J. Gross, P. Ingram, E. Jackson, A. Lambeth, R. Lenglet, S.-H. Li, A. Padmanabhan, J. Pettit, B. Pfaff, R. Ramanathan, S. Shenker, A. Shieh, J. Stribling, P. Thakkar, D. Wendlandt, A. Yip, and R. Zhang, "Network virtualization in multi-tenant datacenters," in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*, Seattle, WA, Apr. 2014, pp. 203–216.
- [113] V. Bollapragada, C. Murphy, and R. White, *Inside Cisco IOS software architecture*, 1st ed. Cisco Press, Jul 2000.
- [114] Juniper Networks, "Junos OS Architecture Overview," 2012. [Online]. Available: http://www.juniper.net/techpubs/en_US/junos12.1/topics/concept/junos-software-architecture.html
- [115] Extreme Networks, "ExtremeXOS Operating System, Version 15.4," 2014. [Online]. Available: <http://learn.extremenetworks.com/rs/extreme/images/EXOS-DS.pdf>
- [116] Alcatel-Lucent, "SR OS," 2014. [Online]. Available: <http://www3.alcatel-lucent.com/products/sros/>
- [117] U. Krishnaswamy, P. Berde, J. Hart, M. Kobayashi, P. Radoslavov, T. Lindberg, R. Sverdlov, S. Zhang, W. Snow, and G. Parulkar, "ONOS: An open source distributed SDN OS," 2013. [Online]. Available: <http://www.slideshare.net/umeshkrishnaswamy/open-network-operating-system>
- [118] A. T. Campbell, I. Katzela, K. Miki, and J. Vicente, "Open signaling for atm, internet and mobile networks (opensig'98)," *SIGCOMM Comput. Commun. Rev.*, vol. 29, no. 1, pp. 97–108, Jan. 1999.
- [119] R. Sherwood, M. Chan, A. Covington, G. Gibb, M. Flajslik, N. Handigol, T.-Y. Huang, P. Kazemian, M. Kobayashi, J. Naous, S. Seetharaman, D. Underhill, T. Yabe, K.-K. Yap, Y. Yiakoumis, H. Zeng,

- G. Appenzeller, R. Johari, N. McKeown, and G. Parulkar, "Carving research slices out of your production networks with OpenFlow," *SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 1, pp. 129–130, Jan. 2010.
- [120] H. Song, J. Gong, J. Song, and J. Yu, "Protocol Oblivious Forwarding (POF)," 2013. [Online]. Available: <http://www.poforwarding.org/>
- [121] ONF, "Charter: Forwarding Abstractions Working Group," April 2014. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/working-groups/charter-forwarding-abstractions.pdf>
- [122] Centec Networks, "V350 - centec open SDN platform," 2013. [Online]. Available: <http://www.valleytalk.org/wp-content/uploads/2013/04/Centec-Open-SDN-Platform.pdf>
- [123] NEC, "Nec ProgrammableFlow UNIVERGE PF5820," 2013. [Online]. Available: http://www.nec.com/en/global/prod/pflow/images_documents/ProgrammableFlow_Switch_PF5820.pdf
- [124] NoviFlow, "NoviSwitch 1248 High Performance OpenFlow Switch," 2013. [Online]. Available: <http://205.236.122.20/gestion/NoviSwitch1248Datasheet.pdf>
- [125] HP, "HP 8200 ZL switch series," 2013. [Online]. Available: http://h17007.www1.hp.com/us/en/networking/products/switches/HP_8200_zl_Switch_Series/
- [126] Arista Networks, "7150 series," 2013. [Online]. Available: http://www.aristanetworks.com/media/system/pdf/Datasheets/7150S_Datasheet.pdf
- [127] Extreme Networks, "Blackdiamond x8," 2013. [Online]. Available: http://www.extremenetworks.com/libraries/products/DSBDX_1832.pdf
- [128] Huawei Technologies Co., Ltd., "Cx6000 metro services platform," 2013. [Online]. Available: http://www.huawei.com/ucmf/groups/public/documents/attachments/hw_132369.pdf
- [129] Juniper Networks, "Ex9200 ethernet switch," 2013. [Online]. Available: <http://www.juniper.net/us/en/local/pdf/datasheets/1000432-en.pdf>
- [130] I. Yokneam, "EZchip announces OpenFlow 1.1 implementations on its NP-4 100-Gigabit network processor," 2011. [Online]. Available: http://www.ezchip.com/pr_110713.htm
- [131] BROCADE, "Brocade MLX series," 2013. [Online]. Available: <http://www.brocade.com/products/all/routers/product-details/netiron-mlx-series/system-options.page>
- [132] IBM, "IBM System Networking RackSwitch G8264," 2013. [Online]. Available: <http://www-03.ibm.com/systems/networking/switches/rack/g8264/>
- [133] NEC, "Nec ProgrammableFlow family of products," 2013. [Online]. Available: <http://www.nec.com/SDN/>
- [134] Pica8, "Pica8 3920," 2013. [Online]. Available: <http://www.pica8.org/documents/pica8-datasheet-64x10gbe-p3780-p3920.pdf>
- [135] Plexxi, "Plexxi Switch 1," 2013. [Online]. Available: http://www.plexxi.com/wp-content/themes/plexxi/assets/pdf/Plexxi_Switch_1_Datasheet_Dec_2012.pdf
- [136] Centec Networks, "Centec v330 OpenFlow switch reference design," 2013. [Online]. Available: <http://www.centecnetworks.com/en/SolutionList.asp?ID=42>
- [137] Cyan, Inc., "Z-series," 2013. [Online]. Available: <http://www.cyaninc.com/en/our-solutions/z-series/>
- [138] Juniper Networks, Inc., "Contrail virtual router," 2013. [Online]. Available: <https://github.com/Juniper/contrail-vrouter>
- [139] FlowForwarding, "LINC-Switch," 2013. [Online]. Available: <http://www.flowforwarding.org/>
- [140] K. Rutka, K. Kaplita, S. Narayan, and S. Bailey, "LINC Switch," 2013. [Online]. Available: http://www.opennetsummit.org/pdf/2013/research_track/poster_papers/ons2013-final36.pdf
- [141] E. L. Fernandes and C. E. Rothenberg, "OpenFlow 1.3 software switch," In SBRC'2014, 2014. [Online]. Available: <https://github.com/CPqD/ofsoftswitch13>
- [142] "Open vSwitch," 2013. [Online]. Available: <http://vswitch.org/>
- [143] OpenFlow Community, "OpenFlow switching reference system," 2009. [Online]. Available: <http://www.openflow.org/wp/downloads/>
- [144] Y. Mundada, R. Sherwood, and N. Feamster, "An OpenFlow switch element for click," in *in Symposium on Click Modular Router*, 2009. [Online]. Available: http://www.cc.gatech.edu/~yogeshm3/click_symposium2009.pdf
- [145] Big Switch Networks, "Project Floodlight," 2013. [Online]. Available: <http://www.projectfloodlight.org/>
- [146] Y. Yiakoumis, J. Schulz-Zander, and J. Zhu, "Pantou : OpenFlow 1.0 for OpenWRT," 2011. [Online]. Available: http://www.openflow.org/wk/index.php/OpenFlow_1.0_for_OpenWRT
- [147] A. Weissberger, "VMware's Network Virtualization Poses Huge Threat to Data Center Switch Fabric Vendors," 2013. [Online]. Available: <http://viodi.com/2013/05/06/vmwares-network-virtualization-poses-huge-threat-to-data-center-switch-fabric-vendors/>
- [148] S. Shenker, "Stanford Seminar - Software-Defined Networking at the Crossroads," June 2013. [Online]. Available: <http://www.youtube.com/watch?v=WabdXYzCAOU>
- [149] M. Casado, "OpenStack and Network Virtualization," April 2013. [Online]. Available: <http://blogs.vmware.com/vmware/2013/04/openstack-and-network-virtualization.html>
- [150] Pica8 Open Networking, "Pica8's os for open switches," 2013. [Online]. Available: <http://www.pica8.org/open-switching/open-switching-overview.php>
- [151] ONIE, "Open Network Install Environment," 2013. [Online]. Available: <http://onie.org/>
- [152] T. Kato, M. Kawakami, T. Myojin, H. Ogawa, K. Hirono, and T. Hasegawa, "Case study of applying SPLE to development of network switch products," in *Proceedings of the 17th International Software Product Line Conference*, ser. SPLC '13. New York, NY, USA: ACM, 2013, pp. 198–207.
- [153] B. Pfaff and B. Davie, "The Open vSwitch Database Management Protocol," RFC 7047 (Informational), Internet Engineering Task Force, Dec. 2013. [Online]. Available: <http://www.ietf.org/rfc/rfc7047.txt>
- [154] M. Smith, M. Dvorkin, Y. Laribi, V. Pandey, P. Garg, and N. Weidenbacher, "OpFlex Control Protocol," Internet Draft, Internet Engineering Task Force, April 2014. [Online]. Available: <http://tools.ietf.org/html/draft-smith-opflex-00>
- [155] G. Bianchi, M. Bonola, A. Capone, and C. Cascone, "OpenState: Programming platform-independent stateful OpenFlow applications inside the switch," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 2, pp. 44–51, Apr. 2014.
- [156] M. Sune, V. Alvarez, T. Jungel, U. Toseef, and K. Pentikousis, "An OpenFlow implementation for network processors," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [157] D. Parniewicz, R. Doriguzzi Corin, L. Ogrodowczyk, M. Rashidi Fard, J. Matias, M. Gerola, V. Fuentes, U. Toseef, A. Zaalouk, B. Belter, E. Jacob, and K. Pentikousis, "Design and implementation of an OpenFlow hardware abstraction layer," in *Proceedings of the 2014 ACM SIGCOMM Workshop on Distributed Cloud Computing*, ser. DCC '14. New York, NY, USA: ACM, 2014, pp. 71–76.
- [158] B. Belter, D. Parniewicz, L. Ogrodowczyk, A. Binczewski, M. Stroinski, V. Fuentes, J. Matias, M. Huarte, and E. Jacob, "Hardware abstraction layer as an SDN-enabler for non-OpenFlow network equipment," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [159] B. Belter, A. Binczewski, K. Dombek, A. Juszczak, L. Ogrodowczyk, D. Parniewicz, M. Stroinski, and I. Olszewski, "Programmable abstraction of datapath," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [160] R. G. Clegg, J. Spencer, R. Landa, M. Thakur, J. Mitchell, and M. Rio, "Pushing software defined networking to the access," in *Proceedings of Third European Workshop on SDN*, 2014.
- [161] V. Fuentes, J. Matias, A. Mendiola, M. Huarte, J. Unzilla, and E. Jacob, "Integrating complex legacy systems under OpenFlow control: The DOCSIS use case," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [162] T. J. Bittman, G. J. Weiss, M. A. Margevicius, and P. Dawson, "Magic Quadrant for x86 Server Virtualization Infrastructure," Gartner, Tech. Rep., June 2013.
- [163] D. W. Cearley, D. Scott, J. Skorupa, and T. J. Bittman, "Top 10 Technology Trends, 2013: Cloud Computing and Hybrid IT Drive Future IT Models," February 2013. [Online]. Available: http://www.gartner.com/Gartner_top_10_technology_trends_2013_237716.pdf
- [164] C. Peng, M. Kim, Z. Zhang, and H. Lei, "VDN: virtual machine image distribution network for cloud data centers," in *INFOCOM, 2012 Proceedings IEEE*, March 2012, pp. 181–189.
- [165] Z. Zhang, Z. Li, K. Wu, D. Li, H. Li, Y. Peng, and X. Lu, "VMThunder: fast provisioning of large-scale virtual machine clusters," *Parallel and Distributed Systems, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2014.
- [166] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar, "FlowVisor: A Network Virtualization Layer," Deutsche Telekom Inc. R&D Lab, Stanford, Nicira Networks, Tech. Rep., 2009.
- [167] S. Azodolmolky, R. Nejabati, S. Peng, A. Hammad, M. P. Chanegowda, N. Efstathiou, A. Autenrieth, P. Kaczmarek, and D. Simeonidou, "Optical FlowVisor: An OpenFlow-based Optical Network Virtualization Approach," in *National Fiber Optic Engineers Conference*, ser. OSA Technical Digest. Optical Society of America, Mar. 2012.

- [168] D. A. Drutskey, "Software-Defined Network Virtualization with FlowN," Ph.D. dissertation, Department of Computer Science of Princeton University, Jun 2012.
- [169] A. Al-Shabibi, M. D. Leenheer, M. Gerolay, A. Koshibe, W. Snow, and G. Parulkar, "OpenVirteX: A Network Hypervisor," 2014. [Online]. Available: <http://ovx.onlab.us/wp-content/uploads/2014/04/ovx-ons14.pdf>
- [170] A. Al-Shabibi, M. De Leenheer, M. Gerola, A. Koshibe, G. Parulkar, E. Salvadori, and B. Snow, "OpenVirteX: Make your virtual SDNs programmable," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 25–30.
- [171] S. Racherla, D. Cain, S. Irwin, P. Ljungstrom, P. Patil, and A. M. Tarenzio, *Implementing IBM Software Defined Network for Virtual Environments*. IBM RedBooks, May 2014.
- [172] C. Li, B. Brech, S. Crowder, D. Dias, H. Franke, M. Hogstrom, D. Lindquist, G. Pacifici, S. Pappé, B. Rajaraman, J. Rao, R. Ratnaparkhi, R. Smith, and M. Williams, "Software defined environments: An introduction," *IBM Journal of Research and Development*, vol. 58, no. 2, pp. 1–11, March 2014.
- [173] A. Gudipati, L. E. Li, and S. Katti, "Radiovisor: A slicing plane for radio access networks," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 237–238.
- [174] H. Yamanaka, E. Kawai, S. Ishii, and S. Shimojo, "AutoVFlow: Autonomous virtualization for wide-area OpenFlow networks," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [175] Berlin Institute for Software Defined Networks (BISDN) GmbH, "The eXtensible OpenFlow datapath daemon (xdpd) bringing innovation into the fast path," 2014. [Online]. Available: <http://xdpd.org/>
- [176] R. Doriguzzi-Corin, E. Salvadori, M. Gerola, M. Sune, and H. Woesner, "A datapath-centric virtualization mechanism for OpenFlow networks," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [177] T. Szyrkowicz, A. Autenrieth, J.-P. Elbers, W. Kellerer, P. Kaczmarek, V. Lopez, L. Contreras, O. G. de Dios, J. P. Fernandez-Palacios, R. Munoz, R. Vilalta, R. Casellas, R. Martinez, A. Mayoral, M. Chanegowda, S. Peng, R. Nejabati, and D. Simeonidou, "Demonstration of SDN based optical network virtualization and multidomain service orchestration," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [178] D. Depaoli, R. Doriguzzi-Corin, M. Gerola, and E. Salvadori, "Demonstrating a distributed and version-agnostic OpenFlow slicing mechanism," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [179] Z. Bozakov and P. Papadimitriou, "AutoSlice: automated and scalable slicing for software-defined networks," in *Proceedings of the 2012 ACM conference on CoNEXT student workshop*, ser. CoNEXT Student '12. New York, NY, USA: ACM, 2012, pp. 3–4.
- [180] D. Drutskey, E. Keller, and J. Rexford, "Scalable network virtualization in software-defined networks," *Internet Computing, IEEE*, vol. 17, no. 2, pp. 20–27, 2013.
- [181] X. Jin, J. Rexford, and D. Walker, "Incremental update for a compositional SDN hypervisor," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 187–192.
- [182] S. Ghorbani and B. Godfrey, "Towards correct network virtualization," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 109–114.
- [183] Juniper Networks, "Opencontrail," 2013. [Online]. Available: <http://opencontrail.org/>
- [184] HP, "Hp SDN controller architecture," Hewlett-Packard Development Company, L.P., Tech. Rep., September 2013.
- [185] K. Phemius, M. Bouet, and J. Leguay, "DISCO: Distributed Multi-domain SDN Controllers," *ArXiv e-prints*, Aug. 2013.
- [186] D. Erickson, "The Beacon OpenFlow controller," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 13–18.
- [187] A. Tootoonchian, S. Gorbunov, Y. Ganjali, M. Casado, and R. Sherwood, "On controller performance in software-defined networks," in *Proceedings of the 2nd USENIX conference on Hot Topics in Management of Internet, Cloud, and Enterprise Networks and Services*, ser. Hot-ICE'12. Berkeley, CA, USA: USENIX Association, 2012, pp. 10–10.
- [188] Z. Cai, A. L. Cox, and T. S. E. Ng, "Maestro: A System for Scalable OpenFlow Control," Rice University, Tech. Rep., 2011.
- [189] "Floodlight is a Java-based OpenFlow controller," 2012. [Online]. Available: <http://floodlight.openflowhub.org/>
- [190] Y. Takamiya and N. Karanatsios, "Trema OpenFlow controller framework," 2012. [Online]. Available: <https://github.com/trema/trema>
- [191] Nippon Telegraph and Telephone Corporation, "Ryu Network Operating System," 2012. [Online]. Available: <http://osrg.github.com/ryu/>
- [192] M. Banikazemi, D. Olshefski, A. Shaikh, J. Tracey, and G. Wang, "Meridian: an SDN platform for cloud network services," *Communications Magazine, IEEE*, vol. 51, no. 2, pp. 120–127, 2013.
- [193] NEC, "Award-winning Software-defined Networking NEC ProgrammableFlow Networking Suite," September 2013. [Online]. Available: <http://www.necam.com/docs/?id=67c33426-0a2b-4b87-9a7a-d3cecc14d26a>
- [194] S. Shin, Y. Song, T. Lee, S. Lee, J. Chung, P. Porras, V. Yegneswaran, J. Noh, and B. B. Kang, "Rosemary: A robust, secure, and high-performance network operating system," in *Proceedings of the 21st ACM Conference on Computer and Communications Security (CCS)*, Nov. 2014, To appear.
- [195] A. Tootoonchian and Y. Ganjali, "HyperFlow: a distributed control plane for OpenFlow," in *Proceedings of the 2010 internet network management conference on Research on enterprise networking*, ser. INM/WREN'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 3–3.
- [196] M. Monaco, O. Michel, and E. Keller, "Applying Operating System Principles to SDN Controller Design," in *Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)*, College Park, MD, November 2013.
- [197] A. D. Ferguson, A. Guha, C. Liang, R. Fonseca, and S. Krishnamurthi, "Participatory networking: an API for application control of SDNs," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 327–338.
- [198] F. Botelho, A. Bessani, F. M. V. Ramos, and P. Ferreira, "On the design of practical fault-tolerant SDN controllers," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [199] S. Matsumoto, S. Hitz, and A. Perrig, "Fleet: Defending SDNs from malicious administrators," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 103–108.
- [200] A. Bierman, M. Bjorklund, K. Watsen, and R. Fernando, "RESTCONF protocol," Internet Draft, Internet Engineering Task Force, July 2014. [Online]. Available: <http://tools.ietf.org/html/draft-ietf-netconf-restconf-01>
- [201] P. Porras, S. Shin, V. Yegneswaran, M. Fong, M. Tyson, and G. Gu, "A security enforcement kernel for OpenFlow networks," in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 121–126.
- [202] L. Richardson and S. Ruby, *RESTful web services*. O'Reilly Media, Inc., 2008.
- [203] T. L. Hinrichs, N. S. Gude, M. Casado, J. C. Mitchell, and S. Shenker, "Practical declarative network management," in *Proceedings of the 1st ACM workshop on Research on enterprise networking*, ser. WREN '09. New York, NY, USA: ACM, 2009, pp. 1–10.
- [204] N. Foster, R. Harrison, M. J. Freedman, C. Monsanto, J. Rexford, A. Story, and D. Walker, "Frenetic: a network programming language," *SIGPLAN Not.*, 2011.
- [205] C. Monsanto, N. Foster, R. Harrison, and D. Walker, "A compiler and run-time system for network programming languages," *SIGPLAN Not.*, vol. 47, no. 1, pp. 217–230, Jan. 2012.
- [206] A. Singla and B. Rijtsman, "Contrail Architecture," Juniper Networks, Tech. Rep., 2013.
- [207] ONF, "OpenFlow management and configuration protocol (OF-Config 1.1.1)," March 2014. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-config/of-config-1-1-1.pdf>
- [208] D. Harrington, R. Presuhn, and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks," Internet Engineering Task Force, dec 2002. [Online]. Available: <http://www.ietf.org/rfc/rfc3411.txt>
- [209] H. Yin, H. Xie, T. Tsou, D. Lopez, P. Aranda, and R. Sidi, "SDNi: A Message Exchange Protocol for Software Defined Networks (SDNS) across Multiple Domains," Internet Draft, Internet Engineering Task Force, June 2012. [Online]. Available: <http://tools.ietf.org/id/draft-yin-sdn-sdni-00.txt>

- [210] M. Garcia, A. Bessani, I. Gashi, N. Neves, and R. Obelheiro, "Analysis of operating system diversity for intrusion tolerance," *Software: Practice and Experience*, vol. 44, no. 6, pp. 735–770, 2014.
- [211] Z. Wang, T. Tsou, J. Huang, X. Shi, and X. Yin, "Analysis of Comparisons between OpenFlow and ForCES," Internet Draft, Internet Engineering Task Force, December 2011. [Online]. Available: <http://tools.ietf.org/id/draft-wang-forces-compare-openflow-forces-00.txt>
- [212] K. Ogawa, W. M. Wang, E. Haleplidis, and J. H. Salim, "ForCES Intra-NE High Availability," Internet Draft, Internet Engineering Task Force, October 2013. [Online]. Available: <http://tools.ietf.org/id/draft-ietf-forces-ceha-08.txt>
- [213] F. A. Botelho, F. M. V. Ramos, D. Kreutz, and A. N. Bessani, "On the feasibility of a consistent and fault-tolerant data store for SDNs," in *Proceedings of the 2013 Second European Workshop on Software Defined Networks*, ser. EWSDN '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 38–43.
- [214] S. Vinoski, "Advanced message queuing protocol," *IEEE Internet Computing*, vol. 10, no. 6, pp. 87–89, Nov. 2006.
- [215] M. Canini, P. Kuznetsov, D. Levin, and S. Schmid, "Software transactional networking: concurrent and consistent policy composition," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 1–6.
- [216] A. Ghodsi, "Distributed k-ary system: Algorithms for distributed hash tables," Ph.D. dissertation, KTH-Royal Institute of Technology, 2006.
- [217] W. Stallings, "Software-defined networks and OpenFlow," *The Internet Protocol Journal*, vol. 16, no. 1, 2013.
- [218] ONF, "SDN architecture," June 2014. [Online]. Available: https://www.opennetworking.org/images/stories/downloads/sdn-resources/technical-reports/TR_SDN_ARCH_1.0_06062014.pdf
- [219] R. Hand and E. Keller, "ClosedFlow: OpenFlow-like control over proprietary devices," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 7–12.
- [220] M. Jarschel, T. Zinner, T. Hofffeld, P. Tran-Gia, and W. Kellerer, "Interfaces, attributes, and use cases: A compass for SDN," *IEEE Communications Magazine*, vol. 52, no. 6, pp. 210–217, 2014.
- [221] E. Mannie, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," RFC 3945 (Proposed Standard), Internet Engineering Task Force, Oct. 2004, updated by RFC 6002. [Online]. Available: <http://www.ietf.org/rfc/rfc3945.txt>
- [222] K. Pentikousis, Y. Wang, and W. Hu, "MobileFlow: Toward software-defined mobile networks," *Communications Magazine, IEEE*, vol. 51, no. 7, pp. 44–53, 2013.
- [223] A. Voellmy and P. Hudak, "Nettle: taking the sting out of programming network routers," in *Proceedings of the 13th international conference on Practical aspects of declarative languages*, ser. PADL'11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 235–249.
- [224] A. Voellmy, H. Kim, and N. Feamster, "Procera: a language for high-level reactive network control," in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 43–48.
- [225] C. Monsanto, J. Reich, N. Foster, J. Rexford, and D. Walker, "Composing software-defined networks," in *Proceedings of the 10th USENIX conference on Networked Systems Design and Implementation*, ser. nsdi'13. Berkeley, CA, USA: USENIX Association, 2013, pp. 1–14.
- [226] C. J. Anderson, N. Foster, A. Guha, J.-B. Jeannin, D. Kozen, C. Schlessinger, and D. Walker, "NetKAT: Semantic foundations for networks," *SIGPLAN Not.*, vol. 49, no. 1, pp. 113–126, Jan. 2014.
- [227] S. Narayana, J. Rexford, and D. Walker, "Compiling path queries in software-defined networks," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 181–186.
- [228] S. Hassas Yeganeh and Y. Ganjali, "Kandoo: A framework for efficient and scalable offloading of control applications," in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 19–24.
- [229] D. Saikia, "MuL OpenFlow controller," 2013. [Online]. Available: <http://sourceforge.net/projects/mul/>
- [230] M. McCauley, "POX," 2012. [Online]. Available: <http://www.noxrepo.org/>
- [231] H. Shimonishi and S. Ishii, "Virtualized network infrastructure using OpenFlow," in *Network Operations and Management Symposium Workshops (NOMS Wksp), 2010 IEEE/IFIP*, 2010, pp. 74–79.
- [232] G. Appenzeller, "SNAC," 2011. [Online]. Available: <http://www.openflowhub.org/display/Snac>
- [233] B. Casemore, "SDN controller ecosystems critical to market success," 2012. [Online]. Available: <http://nerdtwilight.wordpress.com/2012/06/05/sdn-controller-ecosystems-critical-to-market-success/>
- [234] R. Kwan and C. Leung, "A survey of scheduling and interference mitigation in lte," *JECE*, vol. 2010, pp. 1:1–1:10, Jan. 2010.
- [235] A. Gudipati, D. Perry, L. E. Li, and S. Katti, "SoftRAN: Software defined radio access network," in *Proceedings of the second workshop on Hot topics in software defined networks*, ser. HotSDN '13. New York, NY, USA: ACM, 2013.
- [236] J. Dix, "Clarifying the role of software-defined networking northbound APIs," May 2013. [Online]. Available: <http://www.networkworld.com/news/2013/050213-sherwood-269366.html>
- [237] I. GUIIS, "The SDN Gold Rush To The Northbound API," November 2012. [Online]. Available: <http://www.sdncentral.com/technology/the-sdn-gold-rush-to-the-northbound-api/2012/11/>
- [238] B. Salisbury, "The northbound API- a big little problem," 2012.
- [239] G. Ferro, "Northbound API, southbound api, east/north lan navigation in an OpenFlow world and an SDN compass," Aug. 2012.
- [240] B. Casemore, "Northbound API: The standardization debate," Sept. 2012. [Online]. Available: <http://nerdtwilight.wordpress.com/2012/09/18/northbound-api-the-standardization-debate/>
- [241] I. Pepelnjak, "SDN controller northbound API is the crucial missing piece," Sept. 2012. [Online]. Available: <http://blog.ioshints.info/2012/09/sdn-controller-northbound-api-is.html>
- [242] S. Johnson, "A primer on northbound APIs: Their role in a software-defined network," December 2012. [Online]. Available: <http://searchsdn.techtarget.com/feature/A-primer-on-northbound-APIs-Their-role-in-a-software-defined-network>
- [243] R. G. Little, "ONF to standardize northbound API for SDN applications?" October 2013. [Online]. Available: <http://searchsdn.techtarget.com/news/2240206604/ONF-to-standardize-northbound-API-for-SDN-applications>
- [244] Austin Common Standards Revision Group, "POSIX," 2014. [Online]. Available: <http://standards.ieee.org/develop/wg/POSIX.html>
- [245] M. Yu, A. Wundsam, and M. Raju, "NOSIX: A lightweight portability layer for the SDN OS," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 2, pp. 28–35, Apr. 2014.
- [246] R. Chua, "OpenFlow northbound API: A new olympic sport," 2012. [Online]. Available: <http://www.sdncentral.com/sdn-blog/openflow-northbound-api-olympics/2012/07/>
- [247] J. Reich, C. Monsanto, N. Foster, J. Rexford, and D. Walker, "Modular SDN Programming with Pyretic," *USENIX ;login*, vol. 38, no. 5, October 2013.
- [248] K.-K. Yap, T.-Y. Huang, B. Dodson, M. S. Lam, and N. McKeown, "Towards software-friendly networks," in *Proceedings of the first ACM asia-pacific workshop on Workshop on systems*, ser. APSys '10. New York, NY, USA: ACM, 2010, pp. 49–54.
- [249] S. Gutz, A. Story, C. Schlessinger, and N. Foster, "Splendid isolation: A slice abstraction for software-defined networks," in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 79–84.
- [250] D. Turull, M. Hidell, and P. Sjödin, "Evaluating OpenFlow in lib-netvirt," in *The 8th Swedish National Computer Networking Workshop 2012 (SNCNW 2012)*, Oct 2012.
- [251] Quantum Community, "OpenStack Networking ("Quantum")," 2012.
- [252] Small Cell Forum, "Femto APIs," 2013. [Online]. Available: <http://www.smallcellforum.org/developers/>
- [253] M. Guzdial, "Education: Paving the way for computational thinking," *Commun. ACM*, vol. 51, no. 8, pp. 25–27, Aug. 2008.
- [254] M. S. Farooq, S. A. Khan, F. Ahmad, S. Islam, and A. Abid, "An evaluation framework and comparative analysis of the widely used first programming languages," *PLoS ONE*, vol. 9, no. 2, 02 2014.
- [255] A. D. Ferguson, A. Guha, C. Liang, R. Fonseca, and S. Krishnamurthi, "Hierarchical policies for software defined networks," in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 37–42.
- [256] T. Nelson, A. D. Ferguson, M. J. Scheer, and S. Krishnamurthi, "Tierless Programming and Reasoning for Software-Defined Networks," in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. Seattle, WA: USENIX Association, Apr. 2014, pp. 519–531.
- [257] N. P. Katta, J. Rexford, and D. Walker, "Logic programming for software-define networks," in *ACM SIGPLAN Workshop on Cross-Model Language Design and Implementation*, ser. XLDI, 2012.
- [258] S. Shin, P. Porras, V. Yegneswaran, M. Fong, G. Gu, and M. Tyson, "FRESCO: Modular composable security services for software-defined networks," in *Internet Society NDSS*, Feb 2013.

- [259] S. Son, S. Shin, V. Yegneswaran, P. Porras, and G. Gu, "Model checking invariant security properties in OpenFlow," in *Communications (ICC), 2013 IEEE International Conference on*, June 2013, pp. 1974–1979.
- [260] A. Tootoonchian, M. Ghobadi, and Y. Ganjali, "OpenTM: traffic matrix estimator for OpenFlow networks," in *Proceedings of the 11th international conference on Passive and active measurement*, ser. PAM'10. Berlin, Heidelberg: Springer-Verlag, 2010, pp. 201–210.
- [261] M. Reitblatt, M. Canini, A. Guha, and N. Foster, "Fattire: Declarative fault tolerance for software defined networks," in *Proceedings of the second workshop on Hot topics in software defined networks*, ser. HotSDN '13. New York, NY, USA: ACM, 2013.
- [262] A. Voellmy, J. Wang, Y. R. Yang, B. Ford, and P. Hudak, "Maple: simplifying SDN programming using algorithmic policies," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 87–98.
- [263] R. Soule, S. Basu, R. Kleinberg, E. G. Sirer, and N. Foster, "Managing the Network with Merlin," in *Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)*, College Park, MD, November 2013.
- [264] C. Jasson Casey, A. Sutton, G. Dos Reis, and A. Sprintson, "Eliminating Network Protocol Vulnerabilities Through Abstraction and Systems Language Design," *ArXiv e-prints*, Nov. 2013.
- [265] X. Wen, C. Diao, X. Zhao, Y. Chen, L. E. Li, B. Yang, and K. Bu, "Compiling minimum incremental update for modular SDN languages," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 193–198.
- [266] P. Pereini, M. Kuzniar, and D. Kostic, "OpenFlow needs you! a call for a discussion about a cleaner OpenFlow API," in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 44–49.
- [267] F. Facca, E. Salvadori, H. Karl, D. Lopez, P. Aranda Gutierrez, D. Kostic, and R. Riggio, "NetIDE: First steps towards an integrated development environment for portable network apps," in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 105–110.
- [268] E. Reinecke, "Mapping the future of software-defined networking," 2014. [Online]. Available: <http://goo.gl/fQCvRF>
- [269] M. Scharf, V. Gurbani, T. Voith, M. Stein, W. Roome, G. Soprovich, and V. Hilt, "Dynamic VPN optimization by ALTO guidance," in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 13–18.
- [270] M. Stiernerling, S. Kiesel, S. Previdi, and M. Scharf, "ALTO Deployment Considerations," Internet Draft, Internet Engineering Task Force, February 2014. [Online]. Available: <http://tools.ietf.org/html/draft-ietf-alto-deployments-09>
- [271] R. Alimi, R. Penno, and Y. Yang, "ALTO Protocol," Internet Draft, Internet Engineering Task Force, March 2014. [Online]. Available: <https://datatracker.ietf.org/doc/draft-ietf-alto-protocol/>
- [272] N. Handigol, S. Seetharaman, M. Flajslik, A. Gember, N. McKeown, G. Parulkar, A. Akella, N. Feamster, R. Clark, A. Krishnamurthy, V. Brajkovic, and T. A. and, "Aster*x: Load-Balancing Web Traffic over Wide-Area Networks," 2009.
- [273] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown, "ElasticTree: saving energy in data center networks," in *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, ser. NSDI'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 17–17.
- [274] M. S. Seddiki, M. Shahbaz, S. Donovan, S. Grover, M. Park, N. Feamster, and Y.-Q. Song, "FlowQoS: QoS for the rest of us," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 207–208.
- [275] M. Al-Fares, S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat, "Hedera: dynamic flow scheduling for data center networks," in *Proceedings of the 7th USENIX conference on Networked systems design and implementation*, ser. NSDI'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 19–19.
- [276] C. Macapuna, C. Rothenberg, and M. Magalhaes, "In-packet bloom filter based data center networking with distributed OpenFlow controllers," in *GLOBECOM Workshops (GC Wkshps), 2010 IEEE*, 2010, pp. 584–588.
- [277] T. Benson, A. Anand, A. Akella, and M. Zhang, "MicroTE: Fine grained traffic engineering for data centers," in *Proceedings of the Seventh Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '11. New York, NY, USA: ACM, 2011, pp. 8:1–8:12.
- [278] H. Egilmez, S. Dane, K. Bagci, and A. Tekalp, "OpenQoS: An OpenFlow controller design for multimedia delivery with end-to-end quality of service over software-defined networks," in *Signal Information Processing Association Annual Summit and Conference (APSIPA ASC), 2012 Asia-Pacific*, 2012, pp. 1–8.
- [279] A. Sgambelluri, A. Giorgetti, F. Cugini, F. Paolucci, and P. Castoldi, "OpenFlow-based segment protection in Ethernet networks," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 5, no. 9, pp. 1066–1075, Sept 2013.
- [280] M. Bari, S. Chowdhury, R. Ahmed, and R. Boutaba, "PolicyCop: An autonomic QoS policy enforcement framework for software defined networks," in *Future Networks and Services (SDN4FNS), 2013 IEEE SDN for*, Nov 2013, pp. 1–7.
- [281] K. Nagaraj and S. Katti, "ProCel: Smart traffic handling for a scalable software epc," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 43–48.
- [282] P. Xiong and H. Hacigümüs, "Pronto: A software-defined networking based system for performance management of analytical queries on distributed data stores," *PVLDB*, vol. 7, no. 13, pp. 1661–1664, 2014.
- [283] P. Xiong, H. Hacigumus, and J. F. Naughton, "A software-defined networking based approach for performance management of analytical queries on distributed data stores," in *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, ser. SIGMOD '14. New York, NY, USA: ACM, 2014, pp. 955–966.
- [284] N. Handigol, S. Seetharaman, M. Flajslik, N. McKeown, and R. Johari, "Plug-n-Serve: Load-balancing web traffic using OpenFlow," 2009.
- [285] K. Jeong, J. Kim, and Y.-T. Kim, "QoS-aware network operating system for software defined networking with generalized OpenFlows," in *Network Operations and Management Symposium (NOMS), 2012 IEEE*, April 2012, pp. 1167–1174.
- [286] S. Sharma, D. Staessens, D. Colle, D. Palma, J. G. R. Figueiredo, D. Morris, M. Pickavet, and P. Demeester, "Implementing quality of service for the software defined networking enabled future internet," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [287] W. Kim, P. Sharma, J. Lee, S. Banerjee, J. Tourrilhes, S.-J. Lee, and P. Yalagandula, "Automated and scalable QoS control for network convergence," in *Proceedings of the 2010 internet network management conference on Research on enterprise networking*, ser. INM/WREN'10. Berkeley, CA, USA: USENIX Association, 2010, pp. 1–1.
- [288] A. Ishimori, F. Farias, E. Cerqueira, and A. Abelem, "Control of multiple packet schedulers for improving QoS on OpenFlow/SDN networking," in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 81–86.
- [289] D. Palma, J. Goncalves, B. Sousa, L. Cordeiro, P. Simoes, S. Sharma, and D. Staessens, "The QueuePusher: Enabling queue management in OpenFlow," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [290] Z. A. Qazi, C.-C. Tu, L. Chiang, R. Miao, V. Sekar, and M. Yu, "SIMPLE-fying middlebox policy enforcement using SDN," in *Proceedings of the Conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013.
- [291] P. Skoldstrom and B. C. Sanchez, "Virtual aggregation using SDN," in *2013 Second European Workshop on Software Defined Networks*, 2013, pp. –.
- [292] J. Schulz-Zander, N. Sarrar, and S. Schmid, "AeroFlux: A near-sighted controller architecture for software-defined wireless networks," in *Presented as part of the Open Networking Summit 2014 (ONS 2014)*. Santa Clara, CA: USENIX, 2014.
- [293] J. Schulz-Zander, N. Sarrar, and S. Schmid, "Towards a scalable and near-sighted control plane architecture for WiFi SDNs," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14, 2014.
- [294] H. Ali-Ahmad, C. Cicconetti, A. de la Oliva, M. Draxler, R. Gupta, V. Mancuso, L. Rouillet, and V. Sciancalepore, "CROWD: An SDN approach for densenets," in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 25–31.
- [295] J. Vestin, P. Dely, A. Kassler, N. Bayer, H. Einsiedler, and C. Peylo, "CloudMAC: towards software defined WLANs," *SIGMOBILE Mob. Comput. Commun. Rev.*, vol. 16, no. 4, pp. 42–45, Feb. 2013.
- [296] A. Dawson, M. K. Marina, and F. J. Garcia, "On the benefits of RAN virtualisation in C-RAN based mobile networks," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [297] Y. Yamasaki, Y. Miyamoto, J. Yamato, H. Goto, and H. Sone, "Flexible access management system for campus VLAN based on OpenFlow,"

- in *Applications and the Internet (SAINT), 2011 IEEE/IPSJ 11th International Symposium on*, 2011, pp. 347–351.
- [298] J. Schulz-Zander, L. Suresh, N. Sarrar, A. Feldmann, T. Hühn, and R. Merz, “Programmatic orchestration of WiFi networks,” in *2014 USENIX Annual Technical Conference (USENIX ATC 14)*, 2014.
- [299] M. Yang, Y. Li, D. Jin, L. Su, S. Ma, and L. Zeng, “OpenRAN: a software-defined ran architecture via virtualization,” in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM ’13. New York, NY, USA: ACM, 2013, pp. 549–550.
- [300] K.-K. Yap, M. Kobayashi, R. Sherwood, T.-Y. Huang, M. Chan, N. Handigol, and N. McKeown, “OpenRoads: empowering research in mobile networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 40, no. 1, pp. 125–126, Jan. 2010.
- [301] V. Chandrasekhar, J. Andrews, and A. Gatherer, “Femtocell networks: a survey,” *Communications Magazine, IEEE*, vol. 46, no. 9, pp. 59–67, September 2008.
- [302] Y. Yu, C. Qian, and X. Li, “Distributed and collaborative traffic monitoring in software defined networks,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 85–90.
- [303] S. Shirali-Shahreza and Y. Ganjali, “FlexAm: flexible sampling extension for monitoring and security applications in OpenFlow,” in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN ’13. New York, NY, USA: ACM, 2013, pp. 167–168.
- [304] C. Yu, C. Lumezanu, Y. Zhang, V. Singh, G. Jiang, and H. V. Madhyastha, “FlowSense: monitoring network utilization with zero measurement cost,” in *Proceedings of the 14th international conference on Passive and Active Measurement*, ser. PAM’13. Berlin, Heidelberg: Springer-Verlag, 2013, pp. 31–41.
- [305] L. Jose, M. Yu, and J. Rexford, “Online measurement of large traffic aggregates on commodity switches,” in *Proceedings of the 11th USENIX conference on Hot topics in management of internet, cloud, and enterprise networks and services*, ser. Hot-ICE’11. Berkeley, CA, USA: USENIX Association, 2011, pp. 13–13.
- [306] N. L. M. van Adrichem, C. Doerr, and F. A. Kuipers, “OpenNetMon: Network monitoring in openflow software-defined networks,” in *2014 IEEE Network Operations and Management Symposium, NOMS 2014, Krakow, Poland, May 5-9, 2014*, 2014, pp. 1–8.
- [307] J. Suh, T. Kwon, C. Dixon, W. Felter, and J. Carter, “OpenSample: A low-latency, sampling-based measurement platform for commodity SDN,” in *Distributed Computing Systems (ICDCS), 2014 IEEE 34th International Conference on*, June 2014, pp. 228–237.
- [308] sFlow.org Forum, “sFlow,” 2012. [Online]. Available: <http://www.sflow.org/>
- [309] M. Yu, L. Jose, and R. Miao, “Software defined traffic measurement with OpenSketch,” in *Proceedings of the 10th USENIX conference on Networked Systems Design and Implementation*, ser. nsdi’13. Berkeley, CA, USA: USENIX Association, 2013, pp. 29–42.
- [310] C. Argyropoulos, D. Kalogeras, G. Androulidakis, and V. Maglaris, “PaFloMon – a slice aware passive flow monitoring framework for OpenFlow enabled experimental facilities,” in *Software Defined Networking (EWSN), 2012 European Workshop on*, 2012, pp. 97–102.
- [311] S. R. Chowdhury, M. F. Bari, R. Ahmed, and R. Boutaba, “PayLess: A Low Cost Network Monitoring Framework for Software Defined Networks,” in *14th IEEE/IFIP Network Operations and Management Symposium (NOMS 2014)*, 2014.
- [312] G. Wang, T. E. Ng, and A. Shaikh, “Programming your network at run-time for big data applications,” in *HotSDN*. ACM, 2012.
- [313] T. Benson, A. Akella, A. Shaikh, and S. Sahu, “Cloudnaas: a cloud networking platform for enterprise applications,” in *Proceedings of the 2nd ACM Symposium on Cloud Computing*, ser. SOCC ’11. New York, NY, USA: ACM, 2011, pp. 8:1–8:13.
- [314] A. Das, C. Lumezanu, Y. Zhang, V. Singh, G. Jiang, and C. Yu, “Transparent and flexible network management for big data processing in the cloud,” in *Proceedings of the 5th USENIX conference on Hot Topics in Cloud Computing*, ser. HotCloud’13. Berkeley, CA, USA: USENIX Association, 2013.
- [315] A. Arefin, V. K. Singh, G. Jiang, Y. Zhang, and C. Lumezanu, “Diagnosing data center behavior flow by flow,” in *IEEE 33rd International Conference on Distributed Computing Systems*. Philadelphia, USA: IEEE, July 2013.
- [316] E. Keller, S. Ghorbani, M. Caesar, and J. Rexford, “Live migration of an entire network (and its hosts),” in *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, ser. HotNets-XI. New York, NY, USA: ACM, 2012, pp. 109–114.
- [317] R. Raghavendra, J. Lobo, and K.-W. Lee, “Dynamic graph query primitives for SDN-based cloudnetwork management,” in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN ’12. New York, NY, USA: ACM, 2012, pp. 97–102.
- [318] M. Ghobadi, “TCP Adaptation Framework in Data Centers,” Ph.D. dissertation, Graduate Department of Computer Science of University of Toronto, 2013.
- [319] R. Hand, M. Ton, and E. Keller, “Active Security,” in *Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)*, College Park, MD, November 2013.
- [320] S. Shin, V. Yegneswaran, P. Porras, and G. Gu, “AVANT-GUARD: Scalable and Vigilant Switch Flow Management in Software-Defined Networks,” in *Proceedings of the 2013 ACM conference on Computer and communications security*, ser. CCS ’13. New York, NY, USA: ACM, 2013.
- [321] S. Shin and G. Gu, “CloudWatcher: Network security monitoring using OpenFlow in dynamic cloud networks (or: How to provide security monitoring as a service in clouds?),” in *Proceedings of the 2012 20th IEEE International Conference on Network Protocols (ICNP)*, ser. ICNP ’12. Washington, DC, USA: IEEE Computer Society, 2012, pp. 1–6.
- [322] E. Tantar, M. Palattella, T. Avanesov, M. Kantor, and T. Engel, “Cognition: A tool for reinforcing security in software defined networks,” in *EVOLVE - A Bridge between Probability, Set Oriented Numerics, and Evolutionary Computation V*, ser. Advances in Intelligent Systems and Computing, A.-A. Tantar, E. Tantar, J.-Q. Sun, W. Zhang, Q. Ding, O. Schtze, M. Emmerich, P. Legrand, P. Del Moral, and C. A. Coello Coello, Eds. Springer International Publishing, 2014, vol. 288, pp. 61–78.
- [323] R. Braga, E. Mota, and A. Passito, “Lightweight DDoS flooding attack detection using NOX/OpenFlow,” in *Local Computer Networks (LCN), 2010 IEEE 35th Conference on*, Oct 2010, pp. 408–415.
- [324] G. Stabler, A. Rosen, S. Goasguen, and K.-C. Wang, “Elastic ip and security groups implementation using OpenFlow,” in *Proceedings of the 6th international workshop on Virtualization Technologies in Distributed Computing Date*, ser. VTDC ’12. New York, NY, USA: ACM, 2012, pp. 53–60.
- [325] J. Matias, J. Garay, A. Mendiola, N. Toledo, and E. Jacob, “FlowNAC: Flow-based network access control,” in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [326] K. Wang, Y. Qi, B. Yang, Y. Xue, and J. Li, “LiveSec: Towards Effective Security Management in Large-Scale Production Networks,” in *Distributed Computing Systems Workshops (ICDCSW), 2012 32nd International Conference on*, June 2012, pp. 451–460.
- [327] A. Sapio, M. Baldi, Y. Liao, G. Ranjan, F. Risso, A. Tongaonkar, R. Torres, and A. Nucci, “MAPPER: A mobile application personal policy enforcement router for enterprise networks,” in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [328] Y. Wang, Y. Zhang, V. Singh, C. Lumezanu, and G. Jiang, “NetFuse: Short-Circuiting Traffic Surges in the Cloud,” in *IEEE International Conference on Communications*, 2013.
- [329] J. H. Jafarian, E. Al-Shaer, and Q. Duan, “OpenFlow random host mutation: transparent moving target defense using software defined networking,” in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN ’12. New York, NY, USA: ACM, 2012, pp. 127–132.
- [330] J. R. Ballard, I. Rae, and A. Akella, “Extensible and scalable network monitoring using OpenSAFE,” in *Proceedings of the 2010 internet network management conference on Research on enterprise networking*, ser. INM/WREN’10. Berkeley, CA, USA: USENIX Association, 2010, pp. 8–8.
- [331] A. Zaalouk, R. Khondoker, R. Marx, and K. Bayarou, “OrchSec: An orchestrator-based architecture for enhancing network-security using network monitoring and SDN control functions,” in *Network Operations and Management Symposium (NOMS), 2014 IEEE*, May 2014, pp. 1–9.
- [332] Flow-RT group, “sFlow-RT,” 2014. [Online]. Available: <http://www.inmon.com/products/sFlow-RT.php>
- [333] D. Kotani, K. Suzuki, and H. Shimonishi, “A design and implementation of OpenFlow controller handling ip multicast with fast tree switching,” in *Applications and the Internet (SAINT), 2012 IEEE/IPSJ 12th International Symposium on*, 2012, pp. 60–67.
- [334] K. Giotis, G. Androulidakis, and V. Maglaris, “Leveraging SDN for efficient anomaly detection and mitigation on legacy networks,” in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [335] G. Yao, J. Bi, and P. Xiao, “Source address validation solution with

- OpenFlow/NOX architecture,” in *Network Protocols (ICNP), 2011 19th IEEE International Conference on*, 2011, pp. 7–12.
- [336] R. Wang, D. Butnariu, and J. Rexford, “OpenFlow-based server load balancing gone wild,” in *Proceedings of the 11th USENIX conference on Hot topics in management of internet, cloud, and enterprise networks and services*, ser. Hot-ICE’11. Berkeley, CA, USA: USENIX Association, 2011, pp. 12–12.
- [337] X.-N. Nguyen, D. Saucez, C. Barakat, and T. Turletti, “Optimizing rules placement in OpenFlow networks: Trading routing for better efficiency,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 127–132.
- [338] A. Schwabe and H. Karl, “Using MAC addresses as efficient routing labels in data centers,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 115–120.
- [339] I. F. Akyildiz, A. Lee, P. Wang, M. Luo, and W. Chou, “A roadmap for traffic engineering in SDN-OpenFlow networks,” *Comput. Netw.*, vol. 71, pp. 1–30, Oct. 2014.
- [340] H. Ballani, P. Francis, T. Cao, and J. Wang, “Making routers last longer with viaggre,” in *Proceedings of the 6th USENIX symposium on Networked systems design and implementation*, ser. NSDI’09. Berkeley, CA, USA: USENIX Association, 2009, pp. 453–466.
- [341] D. Meyer, L. Zhang, and K. Fall, “Report from the IAB Workshop on Routing and Addressing,” RFC 4984 (Informational), Internet Engineering Task Force, Sep. 2007. [Online]. Available: <http://www.ietf.org/rfc/rfc4984.txt>
- [342] M. Jarschel, F. Wamser, T. Hohn, T. Zinner, and P. Tran-Gia, “SDN-based application-aware networking on the example of youtube video streaming,” in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 87–92.
- [343] T. G. Edwards and W. Belkin, “Using SDN to facilitate precisely timed actions on real-time data streams,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 55–60.
- [344] H. Kumar, H. H. Gharakheili, and V. Sivaraman, “User control of quality of experience in home networks using SDN,” in *Advanced Networks and Telecommunications Systems (ANTS), 2013 IEEE International Conference on*, 2013.
- [345] L. Li, Z. Mao, and J. Rexford, “Toward software-defined cellular networks,” in *Software Defined Networking (EWSN), 2012 European Workshop on*, 2012, pp. 7–12.
- [346] X. Jin, L. Erran Li, L. Vanbever, and J. Rexford, “SoftCell: Scalable and Flexible Cellular Core Network Architecture,” in *Proceedings of the 9th international conference on Emerging networking experiments and technologies*, ser. CoNEXT ’13. New York, NY, USA: ACM, 2013.
- [347] P. Dely, A. Kasser, and N. Bayer, “OpenFlow for wireless mesh networks,” in *Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on*, 31 2011–aug. 4 2011, pp. 1–6.
- [348] M. J. Yang, S. Y. Lim, H. J. Park, and N. H. Park, “Solving the data overload: Device-to-device bearer control architecture for cellular data offloading,” *Vehicular Technology Magazine, IEEE*, vol. 8, no. 1, pp. 31–39, March 2013.
- [349] K.-K. Yap, R. Sherwood, M. Kobayashi, T.-Y. Huang, M. Chan, N. Handigol, N. McKeown, and G. Parulkar, “Blueprint for introducing innovation into wireless mobile networks,” in *Proceedings of the second ACM SIGCOMM workshop on Virtualized infrastructure systems and architectures*, ser. VISA ’10. New York, NY, USA: ACM, 2010, pp. 25–32.
- [350] M. Bansal, J. Mehlman, S. Katti, and P. Levis, “Openradio: a programmable wireless dataplane,” in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN ’12. New York, NY, USA: ACM, 2012, pp. 109–114.
- [351] S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè, “Broadband internet performance: a view from the gateway,” *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 134–145, Aug. 2011.
- [352] S. A. Mehdi, J. Khalid, and S. A. Khayam, “Revisiting traffic anomaly detection using software defined networking,” in *Proceedings of the 14th international conference on Recent Advances in Intrusion Detection*, ser. RAID’11. Berlin, Heidelberg: Springer-Verlag, 2011, pp. 161–180.
- [353] P. Wette and H. Karl, “Which flows are hiding behind my wildcard rule?: adding packet sampling to OpenFlow,” in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM ’13. New York, NY, USA: ACM, 2013, pp. 541–542.
- [354] S. Tarkoma, C. E. Rothenberg, and E. Lagerspetz, “Theory and practice of bloom filters for distributed systems,” *IEEE Communications Surveys and Tutorials*, vol. 14, no. 1, pp. 131–155, 2012.
- [355] J. Kempf, E. Bellagamba, A. Kern, D. Jocha, A. Takacs, and P. Skoldstrom, “Scalable fault management for OpenFlow,” in *Communications (ICC), 2012 IEEE International Conference on*, June 2012, pp. 6606–6610.
- [356] G. Bianchi, M. Bonola, G. Picierno, S. Pontarelli, and M. Monaci, “StreaMon: a data-plane programming abstraction for Software-defined Stream Monitoring,” *ArXiv e-prints*, Nov. 2013.
- [357] D. Kreutz, F. M. Ramos, and P. Verissimo, “Towards secure and dependable software-defined networks,” in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN ’13. New York, NY, USA: ACM, 2013, pp. 55–60.
- [358] K. Kant, “Data center evolution: A tutorial on state of the art, issues, and challenges,” *Computer Networks*, vol. 53, no. 17, pp. 2939–2965, 2009, virtualized Data Centers.
- [359] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, “The cost of a cloud: Research problems in data center networks,” *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 1, pp. 68–73, Dec. 2008.
- [360] M. Bari, R. Boutaba, R. Esteves, L. Granville, M. Podlesny, M. Rabhani, Q. Zhang, and M. Zhani, “Data center network virtualization: A survey,” *Communications Surveys Tutorials, IEEE*, vol. 15, no. 2, pp. 909–928, 2013.
- [361] A. Krishnamurthy, S. P. Chandrabose, and A. Gember-Jacobson, “Pratyastha: An efficient elastic distributed SDN control plane,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 133–138.
- [362] P. Callyam, S. Rajagopalan, A. Selvadurai, S. Mohan, A. Venkataraman, A. Berryman, and R. Ramnath, “Leveraging OpenFlow for resource placement of virtual desktop cloud applications,” in *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on*, 2013, pp. 311–319.
- [363] J. Parraga, “Avior,” 2013. [Online]. Available: <http://openflow.marist.edu/avior>
- [364] GlobalNOC, “OESS - Open Exchange Software Suite,” 2013. [Online]. Available: <http://globalnoc.iu.edu/sdn/oess.html>
- [365] C. Duckett, “Software Defined Networking: HP has an App Store for that,” 2013. [Online]. Available: <http://www.zdnet.com/software-defined-networking-hp-has-an-app-store-for-that-7000021365/>
- [366] Hewlett-Packard Company (HP), “SDN App Store,” 2013. [Online]. Available: <http://h17007.www1.hp.com/us/en/networking/solutions/technology/sdn/devcenter/#sdnAppstore>
- [367] B. H. Sigelman, L. A. Barroso, M. Burrows, P. Stephenson, M. Plakal, D. Beaver, S. Jaspan, and C. Shanbhag, “Dapper, a large-scale distributed systems tracing infrastructure,” Google, Inc., Tech. Rep., 2010.
- [368] L. Layman, M. Diep, M. Nagappan, J. Singer, R. Deline, and G. Venolia, “Debugging revisited: Toward understanding the debugging needs of contemporary software developers,” in *Empirical Software Engineering and Measurement, 2013 ACM / IEEE International Symposium on*, Oct 2013, pp. 383–392.
- [369] U. Erlingsson, M. Peinado, S. Peter, M. Budiu, and G. Mainar-Ruiz, “Fay: Extensible distributed tracing from kernels to clusters,” *ACM Trans. Comput. Syst.*, vol. 30, no. 4, pp. 13:1–13:35, Nov. 2012.
- [370] S. Tomaselli and O. Landsiedel, “Towards Lightweight Logging and Replay of Embedded, Distributed Systems,” in *Proceedings of Workshop ASCoMS (Architecting Safety in Collaborative Mobile Systems) of the 32nd International Conference on Computer Safety, Reliability and Security*, M. ROY, Ed., Toulouse, France, Sep. 2013.
- [371] J. Tan, S. Kavulya, R. Gandhi, and P. Narasimhan, “Visual, log-based causal tracing for performance debugging of mapreduce systems,” in *Distributed Computing Systems (ICDCS), 2010 IEEE 30th International Conference on*, June 2010, pp. 795–806.
- [372] R. Fonseca, G. Porter, R. H. Katz, S. Shenker, and I. Stoica, “X-trace: a pervasive network tracing framework,” in *Proceedings of the 4th USENIX conference on Networked systems design & implementation*, ser. NSDI’07. Berkeley, CA, USA: USENIX Association, 2007, pp. 20–20.
- [373] V. Trivedi, “Software development: Debugging and testing,” in *How to Speak Tech*. Apress, 2014, pp. 89–95.
- [374] A. Anand and A. Akella, “Netreplay: a new network primitive,” *SIGMETRICS Perform. Eval. Rev.*, vol. 37, no. 3, pp. 14–19, Jan. 2010.
- [375] Y. Zhuang, E. Gessiou, S. Portzer, F. Fund, M. Muhammad, I. Beschastnikh, and J. Cappos, “Netcheck: Network diagnoses from blackbox

- traces,” in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. Seattle, WA: USENIX Association, Apr. 2014, pp. 115–128.
- [376] N. Handigol, B. Heller, V. Jeyakumar, D. Mazières, and N. McKeown, “Where is the debugger for my software-defined network?” in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN ’12. New York, NY, USA: ACM, 2012, pp. 55–60.
- [377] A. Wundsam, D. Levin, S. Seetharaman, and A. Feldmann, “OFRewind: enabling record and replay troubleshooting for networks,” in *Proceedings of the 2011 USENIX conference on USENIX annual technical conference*, ser. USENIXATC’11. Berkeley, CA, USA: USENIX Association, 2011, pp. 29–29.
- [378] M. Canini, D. Venzano, P. Perešini, D. Kostić, and J. Rexford, “A NICE way to test OpenFlow applications,” in *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*, ser. NSDI’12. Berkeley, CA, USA: USENIX Association, Apr. 2012, pp. 127–140.
- [379] C. Rotsof, N. Sarrar, S. Uhlig, R. Sherwood, and A. W. Moore, “OFLOPS: an open framework for OpenFlow switch evaluation,” in *Proceedings of the 13th international conference on Passive and Active Measurement*, ser. PAM’12. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 85–95.
- [380] E. Al-Shaer and S. Al-Haj, “FlowChecker: configuration analysis and verification of federated OpenFlow infrastructures,” in *Proceedings of the 3rd ACM workshop on Assurable and usable security configuration*, ser. SafeConfig ’10. New York, NY, USA: ACM, 2010, pp. 37–44.
- [381] A. Khurshid, W. Zhou, M. Caesar, and P. B. Godfrey, “VeriFlow: verifying network-wide invariants in real time,” in *HotSDN*, 2012.
- [382] M. Kuzniar, M. Canini, and D. Kostic, “OFTEN Testing OpenFlow Networks,” in *Proceedings of the 1st European Workshop on Software Defined Networks (EWSN)*, 2012.
- [383] G. Altekar and I. Stoica, “Focus Replay Debugging Effort On the Control Plane,” Electrical Engineering and Computer Sciences University of California at Berkeley, Tech. Rep., May 2010.
- [384] N. Handigol, B. Heller, V. Jeyakumar, D. Mazières, and N. McKeown, “I know what your packet did last hop: Using packet histories to troubleshoot networks,” in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. Seattle, WA: USENIX Association, Apr. 2014, pp. 71–85.
- [385] Y. Zhang, N. Beheshti, and R. Manghirmalani, “NetRevert: Rollback recovery in SDN,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 231–232.
- [386] J. S. Plank, M. Beck, G. Kingsley, and K. Li, “Libckpt: Transparent checkpointing under unix,” in *Proceedings of the USENIX 1995 Technical Conference Proceedings*, ser. TCON’95. Berkeley, CA, USA: USENIX Association, 1995, pp. 18–18.
- [387] J. S. Plank, Y. Chen, K. Li, M. Beck, and G. Kingsley, “Memory exclusion: Optimizing the performance of checkpointing systems,” *Softw. Pract. Exper.*, vol. 29, no. 2, pp. 125–142, Feb. 1999.
- [388] N. Ruchansky and D. Proserpio, “A (not) nice way to verify the OpenFlow switch specification: formal modelling of the OpenFlow switch using alloy,” in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM ’13. New York, NY, USA: ACM, 2013, pp. 527–528.
- [389] H. Zeng, S. Zhang, F. Ye, V. Jeyakumar, M. Ju, J. Liu, N. McKeown, and A. Vahdat, “Libra: Divide and conquer to verify forwarding tables in huge networks,” in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. Seattle, WA: USENIX Association, Apr. 2014, pp. 87–99.
- [390] R. Sherwood and K.-K. Yap, “Cbench controller benchmarker,” 2011. [Online]. Available: <http://www.openflow.org/wk/index.php/Oflops>
- [391] M. Jarschel, F. Lehrieder, Z. Magyari, and R. Pries, “A flexible OpenFlow-controller benchmark,” in *Proceedings of the 2012 European Workshop on Software Defined Networking*, ser. EWSN ’12. Washington, DC, USA: IEEE Computer Society, 2012, pp. 48–53.
- [392] Veryx Technologies, “PktBlaster SDN controller test.” [Online]. Available: <http://sdn.veryxtech.com/>
- [393] C. Rotsof, G. Antichi, M. Bruyere, P. Owezarski, and A. W. Moore, “An open testing framework for next-generation OpenFlow switches,” in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [394] H. Hu, W. Han, G.-J. Ahn, and Z. Zhao, “FLOWGUARD: Building robust firewalls for software-defined networks,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 97–102.
- [395] T. Ball, N. Björner, A. Gember, S. Itzhaky, A. Karbyshev, M. Sagiv, M. Schapira, and A. Valadarsky, “Vericon: Towards verifying controller programs in software-defined networks,” *SIGPLAN Not.*, vol. 49, no. 6, pp. 282–293, Jun. 2014.
- [396] P. Kazemian, G. Varghese, and N. McKeown, “Header space analysis: Static checking for networks,” in *Proceedings of the 9th USENIX Conference on Networked Systems Design and Implementation*, ser. NSDI’12. Berkeley, CA, USA: USENIX Association, 2012, pp. 9–9.
- [397] H. Mai, A. Khurshid, R. Agarwal, M. Caesar, P. B. Godfrey, and S. T. King, “Debugging the data plane with anteater,” *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. 290–301, Aug. 2011.
- [398] P. Kazemian, M. Chang, H. Zeng, G. Varghese, N. McKeown, and S. Whyte, “Real time network policy checking using header space analysis,” in *Proceedings of the 10th USENIX conference on Networked Systems Design and Implementation*, ser. NSDI’13. Berkeley, CA, USA: USENIX Association, 2013, pp. 99–112.
- [399] R. Beckett, X. K. Zou, S. Zhang, S. Malik, J. Rexford, and D. Walker, “An assertion language for debugging SDN applications,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 91–96.
- [400] N. Handigol, B. Heller, V. Jeyakumar, B. Lantz, and N. McKeown, “Reproducible network experiments using container-based emulation,” in *Proceedings of the 8th international conference on Emerging networking experiments and technologies*, ser. CoNEXT ’12. New York, NY, USA: ACM, 2012, pp. 253–264.
- [401] V. Antonenko and R. Smelyanskiy, “Global network modelling based on Mininet approach,” in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN ’13. New York, NY, USA: ACM, 2013, pp. 145–146.
- [402] J. Teixeira, G. Antichi, D. Adami, A. Del Chiaro, S. Giordano, and A. Santos, “Datacenter in a box: Test your SDN cloud-datacenter controller at home,” in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 99–104.
- [403] P. Wette, M. Draxler, A. Schwabe, F. Wallaschek, M. Zahraee, and H. Karl, “Maxinet: Distributed emulation of software-defined networks,” in *Networking Conference, 2014 IFIP*, June 2014, pp. 1–9.
- [404] A. Roy, M. Bari, M. Zhani, R. Ahmed, and R. Boutaba, “Design and management of DOT: A distributed OpenFlow testbed,” in *Network Operations and Management Symposium (NOMS), 2014 IEEE*, May 2014, pp. 1–9.
- [405] A. Carter, D. Morris, S. Sharma, L. Cordeiro, R. Figueiredo, J. Goncalves, D. Palma, N. Johnson, and D. Staessens, “CityFlow: OpenFlow city experiment - linking infrastructure and applications,” in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [406] ns-3 project, “ns-3: OpenFlow switch support,” 2013. [Online]. Available: <http://www.nsnam.org/docs/release/3.13/models/html/openflow-switch.html>
- [407] J. Sommers, R. Bowden, B. Eriksson, P. Barford, M. Roughan, and N. Duffield, “Efficient network-wide flow record generation,” in *INFOCOM, 2011 Proceedings IEEE*, 2011, pp. 2363–2371.
- [408] ucb-sts, “STS - SDN troubleshooting simulator,” 2013. [Online]. Available: <http://ucb-sts.github.io/sts/>
- [409] H. Zhang, C. Lumezanu, J. Rhee, N. Arora, Q. Xu, and G. Jiang, “Enabling layer 2 pathlet tracing through context encoding in software-defined networking,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 169–174.
- [410] K. Agarwal, E. Rozner, C. Dixon, and J. Carter, “SDN traceroute: Tracing SDN forwarding without changing network behavior,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 145–150.
- [411] S. K. Fayaz and V. Sekar, “Testing stateful and dynamic data planes with FlowTest,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 79–84.
- [412] M. Shahbaz, G. Antichi, Y. Geng, N. Zilberman, A. Covington, M. Bruyere, N. Feamster, N. McKeown, B. Felderman, M. Blott, A. W. Moore, and P. Owezarski, “Architecture for an open source network tester,” in *Architectures for Networking and Communications Systems (ANCS), 2013 ACM/IEEE Symposium on*, Oct 2013, pp. 123–124.
- [413] N. Laurent, S. Vissicchio, and M. Canini, “SDLoad: An extensible framework for SDN workload generation,” in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN ’14. New York, NY, USA: ACM, 2014, pp. 215–216.

- [414] M. Gupta, J. Sommers, and P. Barford, "Fast, accurate simulation for SDN prototyping," in *Proceedings of the Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 31–36.
- [415] R. R. FONTES, A. L. C. OLIVEIRA, T. R. PINHEIRO, P. N. SAMPAIO, and R. A. FIGUEIRA, "Authoring of OpenFlow networks with visual network description (SDN version)," in *Proceedings of the Summer Computer Simulation Conference*, ser. SCSC '14, 2014.
- [416] A. R. Curtis, J. C. Mogul, J. Tourrilhes, P. Yalagandula, P. Sharma, and S. Banerjee, "DevoFlow: scaling flow management for high-performance networks," *Comput. Commun. Rev.*, vol. 41, no. 4, pp. 254–265, Aug. 2011.
- [417] C. J. Casey, A. Sutton, and A. Sprintson, "tinyNBI: Distilling an API from essential OpenFlow abstractions," *CoRR*, vol. abs/1403.6644, 2014.
- [418] L. Ogradowczyk et al., "Hardware abstraction layer for non-OpenFlow capable devices," in *The 30th Trans European Research and Education Networking Conference (TNC)*. TERENA, 2014.
- [419] A. Vidal, C. E. Rothenberg, and F. L. Verdi, "The libfluid OpenFlow driver implementation," in *SBRC*, 2014.
- [420] M. APPELMAN and M. D. BOER, "Performance Analysis of OpenFlow Hardware," University of Amsterdam, Tech. Rep., Feb 2012.
- [421] K. Kannan and S. Banerjee, "Compact TCAM: Flow entry compaction in TCAM for power aware SDN," in *Distributed Computing and Networking*, ser. Lecture Notes in Computer Science, D. Frey, M. Raynal, S. Sarkar, R. Shyamasundar, and P. Sinha, Eds. Springer Berlin Heidelberg, 2013, vol. 7730, pp. 439–444.
- [422] J. Liao, "SDN System Performance," June 2012. [Online]. Available: <http://pica8.org/blogs/?p=201>
- [423] B. Agrawal and T. Sherwood, "Modeling TCAM power for next generation network devices," in *Performance Analysis of Systems and Software, 2006 IEEE International Symposium on*, 2006, pp. 120–129.
- [424] B. Owens, "OpenFlow switching performance: Not all TCAM is created equal," February 2013. [Online]. Available: <http://packetpushers.net/openflow-switching-performance-not-all-tcam-is-created-equal/>
- [425] B. Salisbury, "TCAMs and OpenFlow - what every SDN practitioner must know," Jul. 2012. [Online]. Available: <http://www.sdncentral.com/technology/sdn-openflow-tcam-need-to-know/2012/07/>
- [426] W. Braun and M. Menth, "Wildcard compression of inter-domain routing tables for OpenFlow-based software-defined networking," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [427] K. Agarwal, C. Dixon, E. Rozner, and J. Carter, "Shadow MACs: Scalable label-switching for commodity ethernet," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 157–162.
- [428] R. Rudell and A. Sangiovanni-Vincentelli, "Multiple-valued minimization for pla optimization," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 6, no. 5, pp. 727–750, September 1987.
- [429] R. Bifulco and M. Dusi, "Reactive logic in software-defined networking: Accounting for the limitations of the switches," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [430] B. Stephens, A. Cox, W. Felcer, C. Dixon, and J. Carter, "Past: scalable ethernet for data centers," in *Proceedings of the 8th international conference on Emerging networking experiments and technologies*, ser. CoNEXT '12. New York, NY, USA: ACM, 2012, pp. 49–60.
- [431] M. Kobayashi, S. Seetharaman, G. Parulkar, G. Appenzeller, J. Little, J. van Reijndam, P. Weissmann, and N. McKeown, "Maturing of OpenFlow and software-defined networking through deployments," *Computer Networks*, vol. 61, no. 0, pp. 151 – 175, 2014, special issue on Future Internet Testbeds Part I.
- [432] J. C. Mogul and P. Congdon, "Hey, you darned counters!: Get off my asic!" in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 25–30.
- [433] P. Bosshart, G. Gibb, H.-S. Kim, G. Varghese, N. McKeown, M. Izzard, F. Mujica, and M. Horowitz, "Forwarding metamorphosis: fast programmable match-action processing in hardware for SDN," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 99–110.
- [434] O. Ferkouss, I. Snaiki, O. Mounaouar, H. Dahmouni, R. Ben Ali, Y. Lemieux, and C. Omar, "A 100gig network processor platform for openflow," in *Network and Service Management (CNSM), 2011 7th International Conference on*, 2011, pp. 1–4.
- [435] J. Naous, D. Erickson, G. A. Covington, G. Appenzeller, and N. McKeown, "Implementing an OpenFlow switch on the netfpga platform," in *Proceedings of the 4th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, ser. ANCS '08. New York, NY, USA: ACM, 2008, pp. 1–9.
- [436] G. Memon, M. Varvello, R. Laufer, T. Lakshman, J. Li, and M. Zhang, "FlashFlow: a GPU-based Fully Programmable OpenFlow Switch," University of Oregon, Tech. Rep., 2013.
- [437] Y. Luo, P. Cascon, E. Murray, and J. Ortega, "Accelerating OpenFlow switching with network processors," in *Proceedings of the 5th ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, ser. ANCS '09. New York, NY, USA: ACM, 2009, pp. 70–71.
- [438] A. Rostami, T. Jungel, A. Koepsel, H. Woesner, and A. Wolisz, "Oran: OpenFlow routers for academic networks," in *High Performance Switching and Routing (HPSR), 2012 IEEE 13th International Conference on*, 2012, pp. 216–222.
- [439] G. Pongracz, L. Molnar, and Z. Kis, "Removing roadblocks from SDN: OpenFlow software switch performance on Intel DPDK," in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 62–67.
- [440] B. Stephens, "Designing Scalable Networks for Future Large Datacenters," Ph.D. dissertation, Rice University, May 2012.
- [441] Y. Li, D. Zhang, K. Huang, D. He, and W. Long, "A memory-efficient parallel routing lookup model with fast updates," *Comput. Commun.*, vol. 38, pp. 60–71, Feb. 2014.
- [442] N. Katta, J. Rexford, and D. Walker, "Infinite CacheFlow in Software-Defined Networks," Princeton School of Engineering and Applied Science, Tech. Rep., October 2013.
- [443] H. Mekky, F. Hao, S. Mukherjee, Z.-L. Zhang, and T. Lakshman, "Application-aware data plane processing in SDN," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 13–18.
- [444] M. Moshref, A. Bhargava, A. Gupta, M. Yu, and R. Govindan, "Flow-level state transition as a new switch primitive for SDN," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 61–66.
- [445] B. Yan, Y. Xu, H. Xing, K. Xi, and H. J. Chao, "CAB: A reactive wildcard rule caching system for software-defined networks," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 163–168.
- [446] S. McGillicuddy, "XPliant ethernet chip sets new standard for programmability," 2014. [Online]. Available: <http://goo.gl/xE8K9B>
- [447] Intel Processors, "Software Defined Networking and Softwarebased Services with Intel Processors," Intel Corporation, 2012. [Online]. Available: <http://www.intel.com/content/dam/doc/white-paper/communications-ia-software-defined-networking-paper.pdf>
- [448] Intel Corporation, "Intel data plane development kit," 2014. [Online]. Available: <http://www.intel.com/content/dam/www/public/us/en/documents/guides/intel-dpdk-getting-started-guide.pdf>
- [449] A. Sivaraman, K. Winstein, S. Subramanian, and H. Balakrishnan, "No silver bullet: Extending SDN to the data plane," in *Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)*, College Park, MD, November 2013.
- [450] S. Zhou, W. Jiang, and V. Prasanna, "A programmable and scalable OpenFlow switch using heterogeneous SoC platforms," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 239–240.
- [451] S. Hauger, T. Wild, A. Mutter, A. Kirstaedter, K. Karras, R. Ohlendorf, F. Feller, and J. Scharf, "Packet processing at 100 Gbps and beyond - challenges and perspectives," in *Photonic Networks, 2009 ITG Symposium on*, May 2009, pp. 1–10.
- [452] G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, F. Gringoli, and I. Tinnirello, "MAClets: Active MAC protocols over hard-coded devices," in *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*, ser. CoNEXT '12. New York, NY, USA: ACM, 2012, pp. 229–240.
- [453] I. Tinnirello, G. Bianchi, P. Gallo, D. Garlisi, F. Giuliano, and F. Gringoli, "Wireless MAC processors: Programming MAC protocols on commodity hardware," in *INFOCOM, 2012 Proceedings IEEE*, March 2012, pp. 1269–1277.
- [454] P. Bosshart, D. Daly, M. Izzard, N. McKeown, J. Rexford, D. Talayco, A. Vahdat, G. Varghese, and D. Walker, "Programming protocol-independent packet processors," *CoRR*, vol. abs/1312.1719, 2013.
- [455] S. Schmid and J. Suomela, "Exploiting locality in distributed SDN control," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 121–126.

- [456] P. Berde, M. Gerola, J. Hart, Y. Higuchi, M. Kobayashi, T. Koide, B. Lantz, B. O'Connor, P. Radoslavov, W. Snow, and G. Parulkar, "ONOS: Towards an open, distributed SDN OS," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 1–6.
- [457] D. M. Volpano, X. Sun, and G. G. Xie, "Towards systematic detection and resolution of network control conflicts," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 67–72.
- [458] S. Azodolmolky, P. Wieder, and R. Yahyapour, "Performance evaluation of a scalable software-defined networking deployment," in *Software Defined Networks (EWSDN), 2013 Second European Workshop on*, Oct 2013, pp. 68–74.
- [459] S. H. Park, B. Lee, J. Shin, and S. Yang, "A high-performance IO engine for SDN controllers," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [460] Y. Zhang, S. Natarajan, X. Huang, N. Beheshti, and R. Manghirmalani, "A compressive method for maintaining forwarding states in SDN controller," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 139–144.
- [461] S. H. Park, B. Lee, J. You, J. Shin, T. Kim, and S. Yang, "RAON: Recursive abstraction of OpenFlow networks," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [462] A. AuYoung, S. Banerjee, J. Lee, J. C. Mogul, J. Mudigonda, L. Popa, P. Sharma, and Y. Turner, "Corybantic: Towards the Modular Composition of SDN Control Programs," in *Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII)*, College Park, MD, November 2013.
- [463] P. Sun, R. Mahajan, J. Rexford, L. Yuan, M. Zhang, and A. Arefin, "A network-state management service," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: ACM, 2014, pp. 563–574.
- [464] D. Levin, A. Wundsam, B. Heller, N. Handigol, and A. Feldmann, "Logically centralized? state distribution trade-offs in software defined networks," in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 1–6.
- [465] D. Peleg, *Distributed Computing: A Locality-sensitive Approach*. Philadelphia, PA, USA: Society for Industrial and Applied Mathematics, 2000.
- [466] F. J. Ros and P. M. Ruiz, "Five nines of southbound reliability in software-defined networks," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 31–36.
- [467] V. Daniel Philip and Y. Gourhant, "Cross-control: A scalable multi-topology fault restoration mechanism using logically centralized controllers," in *High Performance Switching and Routing (HPSR), 2014 IEEE 15th International Conference on*, July 2014, pp. 57–63.
- [468] M. Borokhovich, L. Schiff, and S. Schmid, "Provable data plane connectivity with local fast failover: Introducing OpenFlow graph algorithms," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 121–126.
- [469] D. Kreutz, A. Casimiro, and M. Pasin, "A trustworthy and resilient event broker for monitoring cloud infrastructures," in *Proceedings of the 12th IFIP WG 6.1 DAIS*. Springer-Verlag, 2012, pp. 87–95.
- [470] K. Tesink, "Definitions of managed objects for the synchronous optical network/synchronous digital hierarchy (SONET/SDH) interface type," Internet Draft, Internet Engineering Task Force, September 2003. [Online]. Available: <http://tools.ietf.org/html/rfc3592>
- [471] R. Prasanna, "BIP: Billing Information Protocol," Internet Draft, Internet Engineering Task Force, December 2002. [Online]. Available: <http://tools.ietf.org/html/draft-prasanna-bip-00>
- [472] G. Swallow, A. Fulignoli, M. Vigoureux, S. Boutros, and D. Ward, "MPLS Fault Management Operations, Administration, and Maintenance (OAM)," RFC 6427 (Proposed Standard), Internet Engineering Task Force, Nov. 2011. [Online]. Available: <http://www.ietf.org/rfc/rfc6427.txt>
- [473] D. Whiting, R. Housley, and N. Ferguson, "Counter with CBC-MAC (CCM)," RFC 3610 (Informational), Internet Engineering Task Force, Sep. 2003. [Online]. Available: <http://www.ietf.org/rfc/rfc3610.txt>
- [474] M. Desai and T. Nandagopal, "Coping with link failures in centralized control plane architectures," in *Communication Systems and Networks (COMSNETS), 2010 Second International Conference on*. IEEE, 2010, pp. 1–10.
- [475] H. Kim, J. Santos, Y. Turner, M. Schlansker, J. Tourrilhes, and N. Feamster, "Coronet: Fault tolerance for software defined networks," in *Network Protocols (ICNP), 2012 20th IEEE International Conference on*, Oct 2012, pp. 1–2.
- [476] S. Sharma, D. Staessens, D. Colle, M. Pickavet, and P. Demeester, "OpenFlow: Meeting carrier-grade recovery requirements," *Comput. Commun.*, vol. 36, no. 6, pp. 656–665, Mar. 2013.
- [477] A. Panda, C. Scott, A. Ghodsi, T. Koponen, and S. Shenker, "Cap for networks," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 91–96.
- [478] M. Kuźniar, P. Perešini, N. Vasić, M. Canini, and D. Kostić, "Automatic failure recovery for software-defined networks," in *Second ACM SIGCOMM Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 159–160.
- [479] A. Dixit, F. Hao, S. Mukherjee, T. Lakshman, and R. Kompella, "Towards an elastic distributed SDN controller," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 7–12.
- [480] R. Ramos, M. Martinello, and C. Esteve Rothenberg, "SlickFlow: Resilient source routing in data center networks unlocked by OpenFlow," in *Local Computer Networks (LCN), 2013 IEEE 38th Conference on*, Oct 2013, pp. 606–613.
- [481] J. T. Araújo, R. Landa, R. G. Clegg, and G. Pavlou, "Software-defined network support for transport resilience," in *IEEE NOMS*, 2014.
- [482] E. Brewer, "Pushing the cap: Strategies for consistency and availability," *Computer*, vol. 45, no. 2, pp. 23–29, Feb. 2012.
- [483] D. Katz and D. Ward, "Bidirectional Forwarding Detection (BFD)," RFC 5880 (Proposed Standard), Internet Engineering Task Force, Jun. 2010. [Online]. Available: <http://www.ietf.org/rfc/rfc5880.txt>
- [484] N. L. M. van Adrichem, B. J. van Asten, and F. A. Kuipers, "Fast recovery in software-defined networks," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [485] N. M. Sahri and K. Okamura, "Fast failover mechanism for software defined networking: Openflow based," in *Proceedings of The Ninth International Conference on Future Internet Technologies*, ser. CFI '14. New York, NY, USA: ACM, 2014, pp. 16:1–16:2.
- [486] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, ser. IMC '10. New York, NY, USA: ACM, 2010, pp. 267–280.
- [487] M. Yu, J. Rexford, M. J. Freedman, and J. Wang, "Scalable flow-based networking with difane," *SIGCOMM Comput. Commun. Rev.*, vol. 41, no. 4, pp. –, Aug. 2010.
- [488] M. F. Bari, A. R. Roy, S. R. Chowdhury, Q. Zhang, M. F. Zhani, R. Ahmed, and R. Boutaba, "Dynamic controller provisioning in software defined networks," in *9th International Conference on Network and Service Management*, ser. CNSM'13, 2013.
- [489] V. Yazici, M. O. Sunay, and A. O. Ercan, "Controlling a software-defined network via distributed controllers," in *Proceedings of the Conference on Implementing Future Media Internet Towards New Horizons*, ser. 2012 NEM SUMMIT. Heidelberg, Germany: Eurescom GmbH, Oct 2012, pp. 16–22.
- [490] M. Jarschel, S. Oechsner, D. Schlosser, R. Pries, S. Goll, and P. Tran-Gia, "Modeling and performance evaluation of an OpenFlow architecture," in *Telettraff Congress (ITC), 2011 23rd International*, Sept 2011, pp. 1–7.
- [491] A. Bianco, R. Birke, L. Giraud, and M. Palacin, "OpenFlow switching: Data plane performance," in *Communications (ICC), 2010 IEEE International Conference on*, may 2010, pp. 1–5.
- [492] R. Pries, M. Jarschel, and S. Goll, "On the usability of OpenFlow in data center environments," in *Communications (ICC), 2012 IEEE International Conference on*, June 2012, pp. 5533–5537.
- [493] J. Hwang, K. K. Ramakrishnan, and T. Wood, "Netvm: High performance and flexible networking using virtualization on commodity platforms," in *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14)*. Seattle, WA: USENIX Association, Apr. 2014, pp. 445–458.
- [494] Y. Dong, Z. Yu, and G. Rose, "Sr-io networking in xen: Architecture, design and implementation," in *Proceedings of the First Conference on I/O Virtualization*, ser. WIOV'08. Berkeley, CA, USA: USENIX Association, 2008, pp. 10–10.
- [495] B. Heller, R. Sherwood, and N. McKeown, "The controller placement problem," in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 7–12.

- [496] S. Azodolmolky, R. Nejabati, M. Pazouki, P. Wieder, R. Yahyapour, and D. Simeonidou, "An analytical model for software defined networking: A network calculus-based approach," in *IEEE GlobeCom 2013*, Oct 2013.
- [497] M. Marchetti, M. Colajanni, M. Messori, L. Aniello, and Y. Vigfusson, "Cyber attacks on financial critical infrastructures," in *Collaborative Financial Infrastructure Protection*, R. Baldoni and G. Chockler, Eds. Springer Berlin Heidelberg, 2012, pp. 53–82.
- [498] S. Amin and A. Giacomoni, "Smart grid, safe grid," *Power and Energy Magazine, IEEE*, vol. 10, no. 1, pp. 33–40, 2012.
- [499] A. Nicholson, S. Webber, S. Dyer, T. Patel, and H. Janicke, "SCADA security in the light of cyber-warfare," *Computers & Security*, vol. 31, no. 4, pp. 418 – 436, 2012.
- [500] K.-K. R. Choo, "The cyber threat landscape: Challenges and future research directions," *Computers & Security*, vol. 30, no. 8, pp. 719 – 731, 2011.
- [501] D. Kushner, "The Real Story of Stuxnet," *IEEE Spectrum*, Mar 2013. [Online]. Available: <http://spectrum.ieee.org/telecom/security/the-real-story-of-stuxnet>
- [502] R. Perez-Pena, "Universities face a rising barrage of cyberattacks," Jul. 2013. [Online]. Available: <http://www.nytimes.com/2013/07/17/education/barrage-of-cyberattacks-challenges-campus-culture.html>
- [503] C. Tankard, "Advanced persistent threats and how to monitor and deter them," *Network Security*, vol. 2011, no. 8, pp. 16 – 19, 2011.
- [504] S. Sorensen, "Security implications of software-defined networks," 2012. [Online]. Available: <http://www.fiercetelecom.com/story/security-implications-software-defined-networks/2012-05-14>
- [505] S. M. Kerner, "Is SDN Secure?" Mar 2013. [Online]. Available: <http://www.enterprisenetworkingplanet.com/netsec/is-sdn-secure.html>
- [506] A. Agapi, K. Birman, R. Broberg, C. Cotton, T. Kielmann, M. Millnert, R. Payne, R. Surton, and R. van Renesse, "Routers for the cloud: Can the internet achieve 5-nines availability?" *Internet Computing, IEEE*, vol. 15, no. 5, pp. 72–77, 2011.
- [507] R. Kloti, "OpenFlow: A security analysis," Master's thesis, Swiss Federal Institute of Technology Zurich (ETH), Zurich, Swiss, 2013.
- [508] M. Wasserman and S. Hartman, "Security analysis of the open networking foundation (onf) OpenFlow switch specification," Internet Engineering Task Force, Apr 2013. [Online]. Available: <https://datatracker.ietf.org/doc/draft-mrw-sdnsec-openflow-analysis/>
- [509] S. Shin and G. Gu, "Attacking software-defined networks: A first feasibility study," in *Proceedings of the second workshop on Hot topics in software defined networks*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 1–2.
- [510] K. Benton, L. J. Camp, and C. Small, "OpenFlow vulnerability assessment," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 151–152.
- [511] S. Scott-Hayward, G. O'Callaghan, and S. Sezer, "SDN security: A survey," in *Future Networks and Services (SDN4FNS), 2013 IEEE SDN for*, Nov 2013, pp. 1–7.
- [512] S. Sezer, S. Scott-Hayward, P. Chouhan, B. Fraser, D. Lake, J. Finnegan, N. Viljoen, M. Miller, and N. Rao, "Are we ready for SDN? implementation challenges for software-defined networks," *Communications Magazine, IEEE*, vol. 51, no. 7, pp. 36–43, July 2013.
- [513] S. Hernan, S. Lambert, T. Ostwald, and A. Shostack, "Uncover security design flaws using the STRIDE approach," *MSDN Magazine*, Nov. 2006.
- [514] B. Chandrasekaran and T. Benson, "Tolerating SDN application failures with LegoSDN," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 235–236.
- [515] W. J. Bolosky, D. Bradshaw, R. B. Haagens, N. P. Kusters, and P. Li, "Paxos replicated state machines as the basis of a high-performance data store," in *Symposium on Networked Systems Design and Implementation (NSDI)*, 2011, pp. 141–154.
- [516] P. Sousa, A. Bessani, M. Correia, N. Neves, and P. Verissimo, "Highly available intrusion-tolerant services with proactive-reactive recovery," *Parallel and Distributed Systems, IEEE Transactions on*, vol. 21, no. 4, pp. 452–465, April 2010.
- [517] R. Chua, "SDN security: Oxymoron? new interview with phil porras of SRI international," 2013. [Online]. Available: <http://www.sdncentral.com/technology/sdn-security-oxymoron-phil-porras-sri/2013/02/>
- [518] J. Korniak, "The GMPLS controlled optical networks as industry communication platform," *Industrial Informatics, IEEE Transactions on*, vol. 7, no. 4, pp. 671–678, Nov 2011.
- [519] P. Fonseca, R. Bennesby, E. Mota, and A. Passito, "A replication component for resilient OpenFlow-based networking," in *Network Operations and Management Symposium (NOMS), 2012 IEEE*, April 2012, pp. 933 –939.
- [520] U. Toseef, A. Zaalouk, T. Rothe, M. Broadbent, and K. Pentikousis, "C-BAS: Certificate-based AAA for SDN experimental facilities," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [521] F. Klaedtke, G. O. Karame, R. Bifulco, and H. Cui, "Access control for SDN controllers," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. New York, NY, USA: ACM, 2014, pp. 219–220.
- [522] D. Kreutz, A. Bessani, E. Feitosa, and H. Cunha, "Towards secure and dependable authentication and authorization infrastructures," in *Pacific Rim International Symposium on Dependable Computing*, ser. PRDC'14, Nov. 2014.
- [523] D. Kreutz and E. Feitosa, "Identity providers-as-a-service built as cloud-of-clouds: challenges and opportunities," in *Position Papers of the 2014 Federated Conference on Computer Science and Information Systems*, ser. Annals of Computer Science and Information Systems, M. P. M. Ganzha, L. Maciaszek, Ed., vol. 3. PTI, 2014, pp. pages 101–108.
- [524] P. Verssimo, N. Neves, and M. Correia, "Intrusion-tolerant architectures: Concepts and design," in *Architecting Dependable Systems*, ser. Lecture Notes in Computer Science, R. de Lemos, C. Gacek, and A. Romanovsky, Eds. Springer Berlin Heidelberg, 2003, vol. 2677, pp. 3–36.
- [525] C. E. Rothenberg, M. R. Nascimento, M. R. Salvador, C. N. A. Corrêa, S. Cunha de Lucena, and R. Raszuk, "Revisiting routing control platforms with the eyes and muscles of software-defined networking," in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 13–18.
- [526] J. P. Stringer, Q. Fu, C. Lorier, R. Nelson, and C. E. Rothenberg, "Cardigan: deploying a distributed routing fabric," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 169–170.
- [527] C. E. Rothenberg, A. Vidal, M. R. Salvador, C. Correa, S. Lucena, F. Farias, E. Cerqueira, and A. Abelem, "Hybrid networking towards a software defined era," in *Network Innovation through OpenFlow and SDN: Principles and Design book*, Taylor & Francis LLC, CRC Press., 2014.
- [528] D. Levin, M. Canini, S. Schmid, F. Schaffert, and A. Feldmann, "Panopticon: Reaping the benefits of incremental SDN deployment in enterprise networks," in *Proceedings of the 2014 USENIX Annual Technical Conference*, ser. USENIX ATC '14. USENIX Association, Jun 2014, pp. 333–345.
- [529] H. Lu, N. Arora, H. Zhang, C. Lumezanu, J. Rhee, and G. Jiang, "Hybnet: Network manager for a hybrid network infrastructure," in *Proceedings of the Industrial Track of the 13th ACM/IFIP/USENIX International Middleware Conference*, ser. Middleware Industry '13. New York, NY, USA: ACM, 2013, pp. 6:1–6:6.
- [530] A. Csoma, B. Sonkoly, L. Csikor, F. Nemeth, A. Gulyas, D. Jocha, J. Elek, W. Tavernier, and S. Sahhaf, "Multi-layered service orchestration in a multi-domain network environment," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [531] R. Morris, E. Kohler, J. Jannotti, and M. F. Kaashoek, "The click modular router," *SIGOPS Oper. Syst. Rev.*, vol. 33, no. 5, pp. 217–231, Dec. 1999.
- [532] S. Salsano, P. L. Ventre, L. Prete, G. Siracusano, M. Gerola, and E. Salvadori, "OSHI - open source hybrid IP/SDN networking (and its emulation on mininet and on distributed SDN testbeds)," *ArXiv e-prints*, Apr. 2014.
- [533] S. Vissicchio, L. Vanbever, and O. Bonaventure, "Opportunities and research challenges of hybrid software defined networks," *SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 2, pp. 70–75, Apr. 2014.
- [534] P. Bernier, "NTT Recognized with IBC Award for SDN-based HDTV Service," September 2013. [Online]. Available: <http://www.sdnzone.com/topics/software-defined-network/articles/353466-ntt-recognized-with-ibc-award-sdn-based-hdtv.htm>
- [535] NTT DATA, "Infrastructure Services," 2014. [Online]. Available: <http://www.nttdata.com/global/en/services/infrastructure/solution.html>
- [536] M. Wagner, "NTT Taps SDN to Enhance Cloud Flexibility," March 2014. [Online]. Available: <http://www.lightreading.com/ntt-taps-sdn-to-enhance-cloud-flexibility/d/d-id/708133>
- [537] AT&T Inc., "AT&T Introduces the "User-Defined Network Cloud": A Vision for the Network of the Future,"

- February 2014. [Online]. Available: <http://www.att.com/gen/press-room?pid=25274&cdvn=news&newsarticleid=37439&mapcode=>
- [538] E. Haleplidis, S. Denazis, O. Koufopavlou, D. Lopez, D. Joachimpillai, J. Martin, J. H. Salim, and K. Pentikousis, "ForCES applicability to SDN-enhanced NFV," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [539] Open Networking Foundation, "OpenFlow-enabled SDN and network functions virtualization," Feb 2014. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/solution-briefs/sb-sdn-nfv-solution.pdf>
- [540] I. Cerrato, A. Palesandro, F. Risso, T. Jungel, M. Sune, and H. Woesner, "User-specific network service functions in an SDN-enabled network node," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [541] M. Xia, M. Shirazipour, Y. Zhang, H. Green, and A. Takacs, "Sdn and optical flow steering for network function virtualization," in *Presented as part of the Open Networking Summit 2014 (ONS 2014)*. Santa Clara, CA: USENIX, 2014.
- [542] E. Haleplidis, J. Hadi Salim, S. Denazis, and O. Koufopavlou, "Towards a network abstraction model for SDN," *Journal of Network and Systems Management*, pp. 1–19, 2014.
- [543] A. Gember-Jacobson, R. Viswanathan, C. Prakash, R. Grandl, J. Khalid, S. Das, and A. Akella, "Opennf: Enabling innovation in network function control," in *Proceedings of the 2014 ACM Conference on SIGCOMM*, ser. SIGCOMM '14. New York, NY, USA: ACM, 2014, pp. 163–174.
- [544] I. Cerrato, M. Annarumma, and F. Risso, "Supporting fine-grained network functions through Intel DPDK," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [545] J. Ruckert, J. Blendin, N. Leymann, G. Schyguda, and D. Hausheer, "Demo: Software-defined network service chaining," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [546] J. Blendin, J. Ruckert, N. Leymann, G. Schyguda, and D. Hausheer, "Software-defined network service chaining," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [547] P. Skoldstrom, B. Sonkoly, A. Gulyas, F. Nemeth, M. Kind, F.-J. Westphal, W. John, J. Garay, E. Jacob, D. Jocha, J. Elek, R. Szabo, W. Tavernier, G. Agapiou, A. Manzalini, M. Rost, and N. S. adn Stefan Schmid, "Towards unified programmability of cloud and carrier infrastructure," in *Third European Workshop on Software Defined Networks*, 2014, pp. –.
- [548] W. John, K. Pentikousis, G. Agapiou, E. Jacob, M. Kind, A. Manzalini, F. Risso, D. Staessens, R. Steinert, and C. Meirosu, "Research directions in network service chaining," in *Future Networks and Services (SDN4FNS), 2013 IEEE SDN for*, Nov 2013, pp. 1–7.
- [549] B. Naudts, M. Kind, F. Westphal, S. Verbrugge, D. Colle, and M. Pickavet, "Techno-economic analysis of software defined networking as architecture for the virtualization of a mobile network," in *Software Defined Networking (EWSN), 2012 European Workshop on*, Oct 2012, pp. 67–72.
- [550] Open Networking Foundation, "Operator network monetization through OpenFlow-enabled SDN," Apr. 2013. [Online]. Available: <https://www.opennetworking.org/images/stories/downloads/sdn-resources/solution-briefs/sb-network-monetization.pdf>
- [551] P. Skoldstrom and W. John, "Implementation and evaluation of a carrier-grade OpenFlow virtualization scheme," in *Software Defined Networks (EWSN), 2013 Second European Workshop on*, Oct 2013, pp. 75–80.
- [552] H. H. Gharakheili and V. Sivaraman, "Virtualizing National Broad-band Access Infrastructure," in *Proceedings of the 9th international conference on Emerging networking experiments and technologies*, ser. CoNEXT '13. New York, NY, USA: ACM, 2013.
- [553] Pacnet Australia, "Pacnet offers first Pan-Asia Network-as-a-Service architecture," November 2013. [Online]. Available: <http://www.cmo.com.au/mediareleases/17701/pacnet-offers-first-pan-asia-network-as-a-service/>
- [554] R. Bennessy, P. Fonseca, E. Mota, and A. Passito, "An inter-AS routing component for software-defined networks," in *Network Operations and Management Symposium (NOMS), 2012 IEEE*, April 2012, pp. 138–145.
- [555] N. D. Corporation, "NTT DATA Advance in SDN Business Provides Highly-Flexible Control of Network by Software," June 2012. [Online]. Available: <http://www.nttdata.com/global/en/news-center/pressrelease/2012/060801.html>
- [556] S. Das, A. Sharafat, G. Parulkar, and N. McKeown, "MPLS with a simple OPEN control plane," in *Optical Fiber Communication Conference and Exposition (OFC/NFOEC), 2011 and the National Fiber Optic Engineers Conference*, 2011, pp. 1–3.
- [557] M. Casado, T. Koponen, S. Shenker, and A. Tootoonchian, "Fabric: a retrospective on evolving SDN," in *Proceedings of the first workshop on Hot topics in software defined networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 85–90.
- [558] M. Martinello, M. Ribeiro, R. de Oliveira, and R. de Angelis Vitoi, "Keyflow: a prototype for evolving SDN toward core network fabrics," *Network, IEEE*, vol. 28, no. 2, pp. 12–19, March 2014.
- [559] N. Feamster, J. Rexford, S. Shenker, R. Clark, R. Hutchins, D. Levin, and J. Bailey, "SDX: A software-defined internet exchange," IETF 86 Proceedings, Orlando, US, March 2013. [Online]. Available: <http://www.ietf.org/proceedings/86/slides/slides-86-sdnrg-6>
- [560] A. Devlic, W. John, and P. Skoldstrom, "A use-case based analysis of network management functions in the ONF SDN model," in *Software Defined Networking (EWSN), 2012 European Workshop on*, Oct 2012, pp. 85–90.
- [561] M. Shirazipour, W. John, J. Kempf, H. Green, and M. Tatipamula, "Realizing packet-optical integration with SDN and OpenFlow 1.1 extensions," in *Communications (ICC), 2012 IEEE International Conference on*, 2012, pp. 6633–6637.
- [562] W. John, A. Kern, M. Kind, P. Skoldstrom, D. Staessens, and H. Woesner, "SplitArchitecture: SDN for the carrier domain," *IEEE Communications Magazine*, Oct. 2014.
- [563] W. John, A. Devlic, Z. Ding, D. Jocha, A. Kern, M. Kind, A. Köpsel, V. Nordell, S. Sharma, P. Sköldström, D. Staessens, A. Takacs, S. Topp, F.-J. Westphal, H. Woesner, and A. Gladisch, "Split Architecture for Large Scale Wide Area Networks," *ArXiv e-prints*, Feb. 2014.
- [564] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer, "Achieving high utilization with software-driven WAN," in *Proceedings of the ACM SIGCOMM 2013 conference on SIGCOMM*, ser. SIGCOMM '13. New York, NY, USA: ACM, 2013, pp. 15–26.
- [565] D. Staessens, S. Sharma, D. Colle, M. Pickavet, and P. Demeester, "Software Defined Networking: Meeting Carrier Grade Requirements," in *Local Metropolitan Area Networks (LANMAN), 2011 18th IEEE Workshop on*, Oct 2011, pp. 1–6.
- [566] S. Sharma, D. Staessens, D. Colle, M. Pickavet, and P. Demeester, "A demonstration of automatic bootstrapping of resilient OpenFlow networks," in *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on*, 2013, pp. 1066–1067.
- [567] R. Niranjan Mysore, A. Pamboris, N. Farrington, N. Huang, P. Miri, S. Radhakrishnan, V. Subramanya, and A. Vahdat, "PortLand: A scalable fault-tolerant layer 2 data center network fabric," *SIGCOMM Comput. Commun. Rev.*, vol. 39, no. 4, pp. 39–50, Aug. 2009.
- [568] A. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta, "VL2: a scalable and flexible data center network," in *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, ser. SIGCOMM '09. New York, NY, USA: ACM, 2009, pp. 51–62.
- [569] A. Sadasivarao, S. Syed, P. Pan, C. Liou, I. Monga, C. Guok, and A. Lake, "Bursting data between data centers: Case for transport SDN," in *High-Performance Interconnects (HOTI), 2013 IEEE 21st Annual Symposium on*, 2013, pp. 87–90.
- [570] J. C. Tanner, "Taking SDN to transport and beyond," 2013. [Online]. Available: <http://www.telecomasia.net/content/taking-sdn-transport-and-beyond>
- [571] S. Elby, "Carrier Vision of SDN," 2012. [Online]. Available: <http://www.brighttalk.com/webcast/6985/58527>
- [572] B. Anwer, T. Benson, N. Feamster, D. Levin, and J. Rexford, "A slick control plane for network middleboxes," in *Proceedings of the second ACM SIGCOMM workshop on Hot topics in software defined networking*, ser. HotSDN '13. New York, NY, USA: ACM, 2013, pp. 147–148.
- [573] C. Gerlach and H.-M. Foisel, "OIF carrier WG requirements on transport networks in SDN architectures," Optical Internetworking Forum, The Optical Internetworking Forum, 48377 Fremont Blvd., Suite 117, Fremont, CA 94538, Tech. Rep., September 2013.
- [574] L. Velasco, A. Asensio, J. Berral, A. Castro, and V. Lopez, "Towards a carrier SDN: An example for elastic inter-datacenter connectivity," in *Optical Communication (ECOC 2013), 39th European Conference and Exhibition on*, 2013, pp. 1–3.
- [575] A. Alba, G. Alatorre, C. Bolik, A. Corrao, T. Clark, S. Gopisetty, R. Haas, R. Kat, B. Langston, N. Mandagere, D. Noll, S. Paddidri, R. Routray, Y. Song, C. Tan, and A. Traeger, "Efficient and agile storage management in software defined environments," *IBM Journal of Research and Development*, vol. 58, no. 2, pp. 1–12, March 2014.

- [576] W. Arnold, D. Arroyo, W. Segmuller, M. Spreitzer, M. Steinder, and A. Tantawi, "Workload orchestration and optimization for software defined environments," *IBM Journal of Research and Development*, vol. 58, no. 2, pp. 1–12, March 2014.
- [577] C. Dixon, D. Olshefski, V. Jain, C. DeCusatis, W. Felter, J. Carter, M. Banikazemi, V. Mann, J. Tracey, and R. Recio, "Software defined networking to support the software defined environment," *IBM Journal of Research and Development*, vol. 58, no. 2, pp. 1–14, March 2014.
- [578] IBM Systems and Technology Group, "IBM software defined network for virtual environments," IBM Corporation, Tech. Rep., January 2014.
- [579] IBM Systems, "Manage all workloads with an efficient, scalable software defined environment (SDE)," 2014. [Online]. Available: <http://www-03.ibm.com/systems/infrastructure/us/en/software-defined-environment/>

Diego Kreutz received his Computer Science degree, MSc degree in Informatics, and MSc degree in Production Engineering from Federal University of Santa Maria. Over the past 11 years he has worked as an Assistant Professor in the Lutheran University of Brazil and in the Federal University of Pampa, and as a researcher member of the Software/Hardware Integration Lab (LISHA) at Federal University of Santa Catarina. Out of the academia, he has also experience as an independent technical consultant on network operations and management for small and medium enterprises and government institutions. Currently, he is a PhD student at Faculty of Sciences of University of Lisbon, Portugal, involved in research projects related to intrusion tolerance, security, and future networks including the TRONE, and SecFuNet international projects. His main research interests are in network control platforms, software-defined networks, intrusion tolerance, system security and dependability, high performance computing, and cloud computing.

Fernando M. V. Ramos is an Assistant Professor in the University of Lisbon. Previous academic positions include those of Teaching Assistant (supervisor) in the University of Cambridge, in the ISEL and in the University of Aveiro. Over the past 12 years he has taught over a dozen courses: from physics (Electromagnetism) to EE (digital electronics, electric circuits, telecommunication systems and foundations) to CS (operating and distributed systems, computer networks, algorithms, programming languages). Periods outside academia include working as a researcher in Portugal Telecom and in Telefonica Research. He holds a PhD degree from the University of Cambridge where he worked on IPTV networks. His current research interests are: software-defined networking, network virtualization, and cloud computing, with security and dependability as an orthogonal concern.

Paulo Veríssimo is a Professor of the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg³. He is currently Chair of the IFIP WG 10.4 on Dependable Computing and Fault-Tolerance and vice-Chair of the Steering Committee of the IEEE/IFIP DSN conference. PJV is Fellow of the IEEE and Fellow of the ACM. He is associate editor of the Elsevier Int'l Journal on Critical Infrastructure Protection. Veríssimo is currently interested in distributed architectures, middleware and algorithms for: adaptability and safety of real-time networked embedded systems; and resilience of secure and dependable large-scale systems. He is author of over 170 peer-refereed publications and co-author of 5 books.

³This work was performed whilst this author was at the U. of Lisbon Faculty of Sciences

Christian Esteve Rothenberg is an Assistant Professor in the Faculty of Electrical and Computer Engineering at University of Campinas (UNICAMP), where he received his Ph.D. in 2010. From 2010 to 2013, he worked as Senior Research Scientist in the areas of IP systems and networking at CPqD Research and Development Center in Telecommunications (Campinas, Brazil), where he was technical lead of R&D activities in the field of OpenFlow/SDN such as the RouteFlow project, the OpenFlow 1.3 Ericsson/CPqD softswitch, or the ONF Driver competition. Christian holds the Telecommunication Engineering degree from Universidad Politécnica de Madrid (ETSIT - UPM), Spain, and the M.Sc. (Dipl. Ing.) degree in Electrical Engineering and Information Technology from the Darmstadt University of Technology (TUD), Germany, 2006. Christian holds two international patents and has over 50 publications including scientific journals and top-tier networking conferences such as SIGCOMM and INFOCOM. Since April 2013, Christian is an ONF Research Associate.

Siamak Azodolmolky received his Computer Engineering degree from Tehran University and his first MSc. degree in Computer Architecture from Azad University in 1994 and 1998 respectively. He was employed by Data Processing Iran Co. (IBM in Iran) as a Software Developer, Systems Engineer, and as a Senior R& D Engineer during 1992-2001. He received his second MSc. degree with distinction from Carnegie Mellon University in 2006. He joined Athens Information Technology (AIT) as a Research Scientist and Software Developer in 2007, while pursuing his PhD degree. In August 2010, he joined the High Performance Networks research group of the School of Computer Science and Electronic Engineering (CSEE) of the University of Essex as a Senior Research Officer. He received his PhD from Universitat Politècnica de Catalunya (UPC) in 2011. He has been the technical investigator of various national and EU funded projects. Software Defined Networking (SDN) has been one of his research interests since 2010, in which he has been investigating the extension of OpenFlow towards its application in core transport (optical) networks. He has published more than 50 scientific papers in international conferences, journals, and books. Software Defined Networking with OpenFlow is one of his recent books. Currently, he is with Gesellschaft für Wissenschaftliche Datenverarbeitung mbH Göttingen (GWDG) as a Senior Researcher and has lead SDN related activities since September 2012. He is a professional member of ACM and a senior member of IEEE.

Steve Uhlig obtained a Ph.D. degree in Applied Sciences from the University of Louvain, Belgium, in 2004. From 2004 to 2006, he was a Postdoctoral Fellow of the Belgian National Fund for Scientific Research (F.N.R.S.). His thesis won the annual IBM Belgium/F.N.R.S. Computer Science Prize 2005. Between 2004 and 2006, he was a visiting scientist at Intel Research Cambridge, UK, and at the Applied Mathematics Department of University of Adelaide, Australia. Between 2006 and 2008, he was with Delft University of Technology, the Netherlands. Prior to joining Queen Mary, he was a Senior Research Scientist with Technische Universität Berlin/Deutsche Telekom Laboratories, Berlin, Germany. Starting in January 2012, he is the Professor of Networks and Head of the Networks Research group at Queen Mary, University of London.