

The Road to SDN: An Intellectual History of Programmable Networks

Nick Feamster
Georgia Tech

feamster@cc.gatech.edu

Jennifer Rexford
Princeton University

jrex@cs.princeton.edu

Ellen Zegura
Georgia Tech

ewz@cc.gatech.edu

ABSTRACT

Software Defined Networking (SDN) is an exciting technology that enables innovation in how we design and manage networks. Although this technology seems to have appeared suddenly, SDN is part of a long history of efforts to make computer networks more programmable. In this paper, we trace the intellectual history of programmable networks, including active networks, early efforts to separate the control and data plane, and more recent work on OpenFlow and network operating systems. We highlight key concepts, as well as the technology pushes and application pulls that spurred each innovation. Along the way, we debunk common myths and misconceptions about the technologies and clarify the relationship between SDN and related technologies such as network virtualization.

1. Introduction

Computer networks are complex and difficult to manage. These networks have many kinds of equipment, from routers and switches to middleboxes such as firewalls, network address translators, server load balancers, and intrusion detection systems. Routers and switches run complex, distributed control software that is typically closed and proprietary. The software implements network protocols that undergo years of standardization and interoperability testing. Network administrators typically configure individual network devices using configuration interfaces that vary across vendors—and even across different products from the same vendor. Although some network-management tools offer a central vantage point for configuring the network, these systems still operate at the level of individual protocols, mechanisms, and configuration interfaces. This mode of operation has slowed innovation, increased complexity, and inflated both the capital and operational costs of running a network.

Software Defined Networking (SDN) is changing the way we design and manage networks. SDN has two defining characteristics. First, an SDN separates the *control plane* (which decides how to handle the traffic) from the *data plane* (which forwards traffic according to decisions that the control plane makes). Second, an SDN consolidates the control plane, so that a single software control program controls *multiple* data-plane elements. The SDN control plane exercises direct control over the state in the

network's data-plane elements (*i.e.*, routers, switches, and other middleboxes) via a well-defined Application Programming Interface (API). OpenFlow [51] is a prominent example of such an API. An OpenFlow switch has one or more tables of packet-handling rules. Each rule matches a subset of traffic and performs certain actions on the traffic that matches a rule; actions include dropping, forwarding, or flooding. Depending on the rules installed by a controller application, an OpenFlow switch can behave like a router, switch, firewall, network address translator, or something in between.

Over the past few years, SDN has gained significant traction in industry. Many commercial switches support the OpenFlow API. Initial vendors that supported OpenFlow included HP, NEC, and Pronto; this list has since expanded dramatically. Many different controller platforms have emerged [23, 30, 37, 46, 56, 64, 82]. Programmers have used these platforms to create many applications, such as dynamic access control [16, 52], server load balancing [39, 83], network virtualization [53, 68], energy-efficient networking [42], and seamless virtual-machine migration and user mobility [24]. Early commercial successes, such as Google's wide-area traffic-management system [44] and Nicira's Network Virtualization Platform [53], have garnered significant industry attention. Many of the world's largest information-technology companies (*e.g.*, cloud providers, carriers, equipment vendors, and financial-services firms) have joined SDN industry consortia like the Open Networking Foundation [55] and the Open Daylight initiative [58].

Although the excitement about SDN has become more palpable during the past few years, many of the ideas underlying SDN have evolved over the past twenty years (or more!). In some ways, SDN revisits ideas from early telephony networks, which used a clear separation of control and data planes to simplify network management and the deployment of new services. Yet, open interfaces like OpenFlow enable more innovation in controller platforms and applications than was possible on closed networks designed for a narrow range of telephony services. In other ways, SDN resembles past research on active networking, which articulated a vision for programmable networks, albeit with an emphasis on programmable data planes. SDN

also relates to previous work on separating the control and data planes in computer networks.

In this article, we present an intellectual history of programmable networks culminating in present-day SDN. We capture the evolution of key ideas, the application “pulls” and technology “pushes” of the day, and lessons that can help guide the next set of SDN innovations. Along the way, we debunk myths and misconceptions about each of the technologies and clarify the relationship between SDN and related technologies, such as network virtualization. Our history begins twenty years ago, just as the Internet takes off, at a time when the Internet’s amazing success exacerbated the challenges of managing and evolving the network infrastructure. We focus on innovations in the networking community (whether by researchers, standards bodies, or companies), although we recognize that these innovations were in some cases catalyzed by progress in other areas, including distributed systems, operating systems, and programming languages. The efforts to create a programmable network infrastructure also clearly relate to the long thread of work on supporting programmable packet processing at high speeds [5, 21, 38, 45, 49, 72, 74].

Before we begin our story, we caution the reader that any history is incomplete and more nuanced than a single storyline might suggest. In particular, much of the work that we describe in this article predates the usage of the term “SDN”, coined in an article [36] about the OpenFlow project at Stanford. The etymology of the term “SDN” is itself complex, and, although the term was initially used to describe Stanford’s OpenFlow project, the definition has since expanded to include a much wider array of technologies. (The term has even been sometimes co-opted by industry marketing departments to describe unrelated ideas that predated Stanford’s SDN project.) Thus, instead of attempting to attribute direct influence between projects, we instead highlight the evolution of and relationships between the ideas that represent the defining characteristics of SDN, regardless of whether or not they directly influenced specific subsequent research. Some of these early ideas may not have directly influenced later ones, but we believe that the connections between the concepts that we outline are noteworthy, and that these projects of the past may yet offer new lessons for SDN in the future.

2. The Road to SDN

Making computer networks more programmable enables innovation in network management and lowers the barrier to deploying new services. In this section, we review early work on programmable networks. We divide the history into three stages, as shown in Figure 1. Each stage has its own contributions to the history: (1) active networks (from the mid-1990s to the early 2000s), which introduced programmable functions in the network

to enable greater to innovation; (2) control and data plane separation (from around 2001 to 2007), which developed open interfaces between the control and data planes; and (3) the OpenFlow API and network operating systems (from 2007 to around 2010), which represented the first instance of widespread adoption of an open interface and developed ways to make control-data plane separation scalable and practical.

Network virtualization played an important role throughout the historical evolution of SDN, substantially predating SDN yet taking root as one of the first significant use cases for SDN. We discuss network virtualization and its relationship to SDN in Section 3.

2.1 Active Networking

The early- to mid-1990s saw the Internet take off, with applications and appeal that far outpaced the early applications of file transfer and email for scientists. More diverse applications and greater use by the general public drew researchers who were eager to test and deploy new ideas for improving network services. To do so, researchers designed and tested new network protocols in small lab settings and simulated behavior on larger networks. Then, if motivation and funding persisted, they took ideas to the Internet Engineering Task Force (IETF) to standardize these protocols. The standardization process was slow and ultimately frustrated many researchers.

In response, some networking researchers pursued an alternative approach of opening up network control, roughly based on the analogy of the relative ease of re-programming a stand-alone PC. Specifically, conventional networks are not “programmable” in any meaningful sense of the word. Active networking represented a radical approach to network control by envisioning a programming interface (or *network API*) that exposed resources (*e.g.*, processing, storage, and packet queues) on individual network nodes, and supported the construction of custom functionality to apply to a subset of packets passing through the node. This approach was anathema to many in the Internet community who advocated that simplicity in the network core was critical to Internet success.

The active networks research program explored radical alternatives to the services provided by the traditional Internet stack via IP or by Asynchronous Transfer Mode (ATM), the other dominant networking approach of the early 1990s. In this sense, active networking was the first in a series of *clean-slate* approaches to network architecture [14] subsequently pursued in programs such as GENI (Global Environment for Network Innovations) [33] and NSF FIND (Future Internet Design) [28] in the United States, and EU FIRE (Future Internet Research and Experimentation Initiative) [29] in the European Union.

The active networking community pursued two programming models:




Figure 1: Selected developments in programmable networking over the past 20 years, and their chronological relationship to advances in network virtualization (one of the first successful SDN use cases).

- the *capsule model*, where the code to execute at the nodes was carried in-band in data packets [84]; and
- the *programmable router/switch model*, where the code to execute at the nodes was established by out-of-band mechanisms (e.g., [8, 69]).

The capsule model came to be most closely associated with active networking. In intellectual connection to subsequent efforts, though, both models have some lasting legacy. Capsules envisioned installation of new data-plane functionality across a network, carrying code in data packets (as in earlier work on packet radio [90]) and using caching to improve efficiency of code distribution. Programmable routers placed decisions about extensibility directly in the hands of the network operator.

Technology push and use pull. The “technology pushes” that encouraged active networking included reduction in the cost of computing, making it conceivable to put more processing in the network, advances in programming languages such as Java that offered platform portability and some code execution safety, and virtual machine technology that protected the host machine (in this case the active node) and other processes from misbehaving programs [71]. Some active networking research projects also capitalized on advances in rapid code compilation and formal methods.

An important catalyst in the active networking ecosystem was funding agency interest, in particular the Active Networks program created and supported by the U.S. Defense Advanced Research Projects Agency (DARPA) from the mid-1990s into the early 2000s. Although not all

research work in active networks was funded by DARPA, the funding program supported a collection of projects and, perhaps more importantly, encouraged convergence on a terminology and set of active network components so that projects could contribute to a whole meant to be greater than the sum of the parts [14]. The Active Networks program placed an emphasis on demonstrations and project inter-operability, with a concomitant level of development effort. The bold and concerted push from a funding agency in the absence of near-term use cases may have also contributed to a degree of community skepticism about active networking that was often healthy but could border on hostility and may have obscured some of the intellectual connections between that work and later efforts to provide network programmability.

The “use pulls” for active networking described in the literature of the time [15, 75] are remarkably similar to the examples used to motivate SDN today. The issues of the day included network service provider frustration with the timescales necessary to develop and deploy new network services (so-called network ossification), third-party interest in value-added, fine-grained control to dynamically meet the needs of particular applications or network conditions, and researcher desire for a platform that would support experimentation at scale. Additionally, many early papers on active networking cited the proliferation of middleboxes, including firewalls, proxies, and transcoders, each of which had to be deployed separately and entailed a distinct (often vendor-specific) programming model. Active networking offered a vision of unified control over these middleboxes that could ultimately

replace the ad hoc, one-off approaches to managing and controlling these boxes [75]. Interestingly, the early literature foreshadows the current trends in network functions virtualization (NFV) [19], which also aims to provide a unifying control framework for networks that have complex middlebox functions deployed throughout.

Intellectual contributions. Active networks offered intellectual contributions that relate to SDN. We note three in particular:

- *Programmable functions in the network to lower the barrier to innovation.* Research in active networks pioneered the notion of programmable networks as a way to lower the barrier to network innovation. The notion that it is difficult to innovate in a production network and pleas for increased programmability were commonly cited in the initial motivation for SDN. Much of the early vision for SDN focused on *control-plane* programmability, whereas active networks focused more on *data-plane* programmability. That said, data-plane programmability has continued to develop in parallel with control-plane efforts [5, 21], and data-plane programmability is again coming to the forefront in the emerging NFV initiative. Recent work on SDN is exploring the evolution of SDN protocols such as OpenFlow to support a wider range of data-plane functions [11]. Additionally, the concepts of isolation of experimental traffic from normal traffic—which have their roots in active networking—also appear front and center in design documents for OpenFlow [51] and other SDN technologies (e.g., FlowVisor [31]).
- *Network virtualization, and the ability to demultiplex to software programs based on packet headers.* The need to support experimentation with multiple programming models led to work on network virtualization. Active networking produced an architectural framework that describes the components of such a platform [13]. The key components of this platform are a shared Node Operating System (NodeOS) that manages shared resources; a set of Execution Environments (EEs), each of which defines a virtual machine for packet operations; and a set of Active Applications (AAs) that work within a given EE to provide an end-to-end service. Directing packets to a particular EE depends on fast pattern matching on header fields and demultiplexing to the appropriate EE. Interestingly, this model was carried forward in the PlanetLab [61] architecture, whereby different experiments run in virtual execution environments, and packets are demultiplexed into the appropriate execution environment on their packet headers. Demultiplexing packets into different virtual execution environments has also been applied to the design of virtualized programmable hardware data planes [5].

- *The vision of a unified architecture for middlebox orchestration.* Although the vision was never fully realized in the active networking research program, early design documents cited the need for unifying the wide range of middlebox functions with a common, safe programming framework. Although this vision may not have directly influenced the more recent work on NFV, various lessons from active networking research may prove useful as we move forward with the application of SDN-based control and orchestration of middleboxes.

Myths and misconceptions. Active networking included the notion that a network API would be available to end-users who originate and receive packets, though most in the research community fully recognized that end-user network programmers would be rare [15]. The misconception that packets would necessarily carry Java code written by end users made it possible to dismiss active network research as too far removed from real networks and inherently unsafe. Active networking was also criticized at the time for not being able to offer practical performance and security. While performance was not a first-order consideration of the active networking research community (which focused on architecture, programming models, and platforms), some efforts aimed to build high-performance active routers [86]. Similarly, while security was under-addressed in many of the early projects, the Secure Active Network Environment Architecture project [2] was a notable exception.

In search of pragmatism. Although active networks articulated a vision of programmable networks, the technologies did not see widespread deployment. Many factors drive the adoption of a technology (or lack thereof). Perhaps one of the biggest stumbling blocks that active networks faced was the lack of an immediately compelling problem or a clear path to deployment. A significant lesson from the active networks research effort was that “killer” applications for the data plane are hard to conceive. The community proffered various applications that could benefit from in-network processing, including information fusion, caching and content distribution, network management, and application-specific quality of service [15, 75]. Unfortunately, although performance benefits could be quantified in the lab, none of these application areas demonstrated a sufficiently compelling solution to a pressing need.

Subsequent efforts, which we describe in the next subsection, were more modest in terms of the scope of problems they addressed, focusing narrowly on routing and configuration management. In addition to a narrower scope, the next phase of research developed technologies that drew a clear distinction and separation between the functions of the control and data planes. This separation ultimately made it possible to focus on innovations

in the control plane, which not only needed a significant overhaul but also (because it is commonly implemented in software) presented a lower barrier to innovation than the data plane.

2.2 Separating Control and Data Planes

In the early 2000s, increasing traffic volumes and a greater emphasis on network reliability, predictability, and performance led network operators to seek better approaches to certain network-management functions such as the control over the paths used to deliver traffic (a practice commonly known as traffic engineering). The means for performing traffic engineering using conventional routing protocols were primitive at best. Operators' frustration with these approaches were recognized by a small, well-situated community of researchers who either worked for or interacted regularly with backbone network operators. These researchers explored pragmatic, near-term approaches that were either standards-driven or imminently deployable using existing protocols.

Specifically, conventional routers and switches embody a tight integration between the control and data planes. This coupling made various network-management tasks, such as debugging configuration problems and predicting or controlling routing behavior, exceedingly challenging. To address these challenges, various efforts to separate the data and control planes began to emerge.

Technology push and use pull. As the Internet flourished in the 1990s, the link speeds in backbone networks grew rapidly, leading equipment vendors to implement packet-forwarding logic directly in hardware, separate from the control-plane software. In addition, Internet Service Providers (ISPs) were struggling to manage the increasing size and scope of their networks, and the demands for greater reliability and new services (such as virtual private networks). In parallel with these two trends, the rapid advances in commodity computing platforms meant that servers often had substantially more memory and processing resources than the control-plane processor of a router deployed just one or two years earlier. These trends catalyzed two innovations:

- an *open interface between the control and data planes*, such as the ForCES (Forwarding and Control Element Separation) [88] interface standardized by the Internet Engineering Task Force (IETF) and the Netlink interface to the kernel-level packet-forwarding functionality in Linux [66]; and
- *logically centralized control* of the network, as seen in the Routing Control Platform (RCP) [12, 26] and SoftRouter [47] architectures, as well as the Path Computation Element (PCE) [25] protocol at the IETF.

These innovations were driven by industry's demands for technologies to manage routing within an ISP network.

Some early proposals for separating the data and control planes also came from academic circles, in both ATM networks [10, 32, 80] and active networks [70].

Compared to earlier research on active networking, these projects focused on pressing problems in network management, with an emphasis on: innovation by and for network administrators (rather than end users and researchers); programmability in the control plane (rather than the data plane); and network-wide visibility and control (rather than device-level configuration).

Network-management applications included selecting better network paths based on the current traffic load, minimizing transient disruptions during planned routing changes, giving customer networks more control over the flow of traffic, and redirecting or dropping suspected attack traffic. Several control applications ran in operational ISP networks using legacy routers, including the Intelligent Route Service Control Point (IRSCP) deployed to offer value-added services for virtual-private network customers in AT&T's tier-1 backbone network [78]. Although much of the work during this time focused on managing routing within a single ISP, some work [25, 26] also proposed ways to enable flexible route control across multiple administrative domains.

Moving control functionality off of network equipment and into separate servers made sense because network management is, by definition, a network-wide activity. Logically centralized routing controllers [12, 47, 78] were enabled by the emergence of open-source routing software [9, 40, 65] that lowered the barrier to creating prototype implementations. The advances in server technology meant that a single commodity server could store all of the routing state and compute all of the routing decisions for a large ISP network [12, 81]. This, in turn, enabled simple primary-backup replication strategies, where backup servers store the same state and perform the same computation as the primary server, to ensure controller reliability.

Intellectual contributions. The initial attempts to separate the control and data planes were relatively pragmatic, but they represented a significant conceptual departure from the Internet's conventionally tight coupling of path computation and packet forwarding. The efforts to separate the network's control and data plane resulted in several concepts that have been carried forward in subsequent SDN designs:

- *Logically centralized control using an open interface to the data plane.* The ForCES working group at the IETF proposed a standard, open interface to the data plane to enable innovation in control-plane software. The SoftRouter [47] used the ForCES API to allow a separate controller to install forwarding-table entries in the data plane, enabling the complete removal of control functionality from the routers. Unfortunately, ForCES was not adopted by the major router vendors, which

hampered incremental deployment. Rather than waiting for new, open APIs to emerge, the RCP [12,26] used an *existing* standard control-plane protocol (the Border Gateway Protocol) to install forwarding-table entries in legacy routers, enabling immediate deployment. OpenFlow also faced similar backwards compatibility challenges and constraints: in particular, the initial OpenFlow specification relied on backwards compatibility with hardware capabilities of commodity switches.

- **Distributed state management.** Logically centralized route controllers faced challenges involving distributed state management. A logically centralized controller must be replicated to cope with controller failure, but replication introduces the potential for inconsistent state across replicas. Researchers explored the likely failure scenarios and consistency requirements. At least in the case of routing control, the controller replicas did not need a general state management protocol, since each replica would eventually compute the same routes (after learning the same topology and routing information) and transient disruptions during routing-protocol convergence were acceptable even with legacy protocols [12]. For better scalability, each controller instance could be responsible for a separate portion of the topology. These controller instances could then exchange routing information with each other to ensure consistent decisions [81]. The challenges of building distributed controllers would arise again several years later in the context of distributed SDN controllers [46, 56]. Distributed SDN controllers face the far more general problem of supporting *arbitrary* controller applications, requiring more sophisticated solutions for distributed state management.

Myths and misconceptions. When these new architectures were proposed, critics viewed them with healthy skepticism, often vehemently arguing that logically centralized route control would violate “fate sharing”, since the controller could fail independently from the devices responsible for forwarding traffic. Many network operators and researchers viewed separating the control and data planes as an inherently bad idea, as initially there was no clear articulation of how these networks would continue to operate correctly if a controller failed. Skeptics also worried that logically centralized control moved away from the conceptually simple model of the routers achieving distributed consensus, where they all (eventually) have a common view of network state (*e.g.*, through flooding). In logically centralized control, each router has only a purely local view of the *outcome* of the route-selection process.

In fact, by the time these projects took root, even the traditional distributed routing solutions already violated these principles. Moving packet-forwarding logic

into hardware meant that a router’s control-plane software could fail independently from the data plane. Similarly, distributed routing protocols adopted scaling techniques, such as OSPF areas and BGP route reflectors, where routers in one region of a network had limited visibility into the routing information in other regions. As we discuss in the next section, the separation of the control and data planes somewhat paradoxically *enabled* researchers to think more clearly about distributed state management: the decoupling of the control and data planes catalyzed the emergence of a state management layer that maintains consistent view of network state.

In search of generality. Dominant equipment vendors had little incentive to adopt standard data-plane APIs like ForCES, since open APIs could enable new entrants into the marketplace. The resulting need to rely on existing routing protocols to control the data plane imposed significant limitations on the range of applications that programmable controllers could support. Conventional IP routing protocols compute routes for destination IP address blocks, rather than providing a wider range of functionality (*e.g.*, dropping, flooding, or modifying packets) based on a wider range of header fields (*e.g.*, MAC and IP addresses, TCP and UDP port numbers), as OpenFlow does. In the end, although the industry prototypes and standardization efforts made some progress, widespread adoption remained elusive.

To broaden the vision of control and data plane separation, researchers started exploring clean-slate architectures for logically centralized control. The 4D project [35] advocated four main layers—the *data* plane (for processing packets based on configurable rules), the *discovery* plane (for collecting topology and traffic measurements), the *dissemination* plane (for installing packet-processing rules), and a *decision* plane (consisting of logically centralized controllers that convert network-level objectives into packet-handling state). Several groups proceeded to design and build systems that applied this high-level approach to new application areas [16, 87], beyond route control. In particular, the Ethane project [16] (and its direct predecessor, SANE [17]) created a logically centralized, flow-level solution for access control in enterprise networks. Ethane reduces the switches to flow tables that are populated by the controller based on high-level security policies. The Ethane project, and its operational deployment in the Stanford computer science department, set the stage for the creation of OpenFlow. In particular, the simple switch design in Ethane became the basis of the original OpenFlow API.

2.3 OpenFlow and Network OSEs

In the mid-2000s, researchers and funding agencies gained interest in the idea of network experimentation at scale, encouraged by the success of experimental in-

infrastructures (e.g., PlanetLab [6] and Emulab [85]), and the availability of separate government funding for large-scale “instrumentation” previously reserved for other disciplines to build expensive, shared infrastructure such as colliders and telescopes [54]. An outgrowth of this enthusiasm was the creation of the Global Environment for Networking Innovations (GENI) [33] with an NSF-funded GENI Project Office and the EU FIRE program [29]. Critics of these infrastructure-focused efforts pointed out that this large investment in infrastructure was not matched by well-conceived ideas to use it. In the midst of this, a group of researchers at Stanford created the Clean Slate Program and focused on experimentation at a more local and tractable scale: campus networks [51].

Before the emergence of OpenFlow, the ideas underlying SDN faced a tension between the vision of fully programmable networks and pragmatism that would enable real-world deployment. OpenFlow struck a balance between these two goals by enabling more functions than earlier route controllers and building on existing switch hardware, through the increasing use of merchant-silicon chipsets in commodity switches. Although relying on existing switch hardware did somewhat limit flexibility, OpenFlow was almost immediately deployable, allowing the SDN movement to be both pragmatic and bold. The creation of the OpenFlow API [51] was followed quickly by the design of controller platforms like NOX [37] that enabled the creation of many new control applications.

An OpenFlow switch has a table of packet-handling rules, where each rule has a *pattern* (that matches on bits in the packet header), a list of *actions* (e.g., drop, flood, forward out a particular interface, modify a header field, or send the packet to the controller), a set of *counters* (to track the number of bytes and packets), and a *priority* (to disambiguate between rules with overlapping patterns). Upon receiving a packet, an OpenFlow switch identifies the highest-priority matching rule, performs the associated actions, and increments the counters.

Technology push and use pull. Perhaps the defining feature of OpenFlow is its adoption in industry, especially as compared with its intellectual predecessors. This success can be attributed to a perfect storm of conditions between equipment vendors, chipset designers, network operators, and networking researchers. Before OpenFlow’s genesis, switch chipset vendors like Broadcom had already begun to open their APIs to allow programmers to control certain forwarding behaviors. The decision to open the chipset provided the necessary impetus to an industry that was already clamoring for more control over network devices. The availability of these chipsets also enabled a much wider range of companies to build switches, without incurring the substantial cost of designing and fabricating their own data-plane hardware.

The initial OpenFlow protocol standardized a data-plane model and a control-plane API by building on technology that switches already supported. Specifically, because network switches already supported fine-grained access control and flow monitoring, enabling OpenFlow’s initial set of capabilities on switch was as easy as performing a firmware upgrade—vendors did not need to upgrade the hardware to make their switches OpenFlow-capable.

OpenFlow’s initial target deployment scenario was campus networks, meeting the needs of a networking research community actively looking for ways to conduct experimental work on “clean-slate” network architectures within a research-friendly operational setting. In the late 2000s, the OpenFlow group at Stanford led an effort to deploy OpenFlow testbeds across many campuses and demonstrate the capabilities of the protocol both on a single campus network and over a wide-area backbone network spanning multiple campuses [34].

As real SDN use cases materialized on these campuses, OpenFlow began to take hold in other realms, such as data-center networks, where there was a distinct need to manage network traffic at large scales. In data centers, the cost of hiring engineers to write sophisticated control programs to run over large numbers of commodity switches proved to be more cost-effective than continuing to purchase closed, proprietary switches that could not support new features without substantial engagement with the equipment vendors. As vendors began to compete to sell both servers and switches for data centers, many smaller players in the network equipment marketplace embraced the opportunity to compete with the established router and switch vendors by supporting new capabilities like OpenFlow.

Intellectual contributions. Although OpenFlow embodied many of the principles from earlier work on the separation of control and data planes, the rise of OpenFlow offered several additional intellectual contributions:

- *Generalizing network devices and functions.* Previous work on route control focused primarily on matching traffic by destination IP prefix. In contrast, OpenFlow rules could define forwarding behavior on traffic flows based on any set of 13 different packet headers. As such, OpenFlow conceptually unified many different types of network devices that differ only in terms of which header fields they match, and which actions they perform. A router matches on destination IP prefix and forwards out a link, whereas a switch matches on source MAC address (to perform MAC learning) and destination MAC address (to forward), and either floods or forwards out a single link. Network address translators and firewalls match on the five tuple (of source and destination IP addresses and port numbers, and the transport protocol) and either rewrites address and port fields, or drops unwanted traffic. OpenFlow also generalized the rule-

installation techniques, allowing anything from *proactive* installation of coarse-grained rules (*i.e.*, with “wildcards” for many header fields) to *reactive* installation of fine-grained rules, depending on the application. Still, OpenFlow does not offer data-plane support for deep packet inspection or connection reassembly; as such, OpenFlow alone cannot efficiently enable sophisticated middlebox functionality.

- *The vision of a network operating system.* In contrast to earlier research on active networks that proposed a *node* operating system, the work on OpenFlow led to the notion of a *network* operating system [37]. A network operating system is software that abstracts the installation of state in network switches from the logic and applications that control the behavior of the network. More generally, the emergence of a network operating system offered a conceptual decomposition of network operation into three layers [46]: (1) a *data plane* with an open interface; (2) a *state management layer* that is responsible for maintaining a consistent view of network state; (3) *control logic* that performs various operations depending on its view of network state.
- *Distributed state management techniques.* Separating the control and data planes introduces new challenges concerning state management. Running multiple controllers is crucial for scalability, reliability, and performance, yet these replicas should work together to act like a single, logically centralized controller. Previous work on distributed route controllers [12, 81] only addressed these problems in the narrow context of route computation. To support arbitrary controller applications, the work on the Onix [46] controller introduced the idea of a *network information base*—a representation of the network topology and other control state shared by all controller replicas. Onix also incorporated past work in distributed systems to satisfy the state consistency and durability requirements. For example, Onix has a transactional persistent database backed by a replicated state machine for slowly-changing network state, as well as an in-memory distributed hash table for rapidly-changing state with weaker consistency requirements. More recently, the ONOS [56] system offers an open-source controller with similar functionality, using existing open-source software for maintaining consistency across distributed state and providing a network topology database to controller applications.

Myths and misconceptions. One myth concerning SDN is that the first packet of every traffic flow must go to the controller for handling. Indeed, some early systems like Ethane [16] worked this way, since they were designed to support fine-grained policies in small networks. In fact, SDN in general, and OpenFlow in particular, do not impose any assumptions about the granularity of rules or

whether the controller handles any data traffic. Some SDN applications respond only to topology changes and coarse-grained traffic statistics and update rules infrequently in response to link failures or network congestion. Other applications may send the first packet of some larger traffic aggregate to the controller, but not a packet from every TCP or UDP connection.

A second myth surrounding SDN is that the controller must be physically centralized. In fact, Onix [46] and ONOS [56] demonstrate that SDN controllers can—and should—be distributed. Wide-area deployments of SDN, as in Google’s private backbone [44], have many controllers spread throughout the network.

Finally, a commonly held misconception is that SDN and OpenFlow are equivalent; in fact, OpenFlow is merely one (widely popular) instantiation of SDN principles. Different APIs could be used to control network-wide forwarding behavior; previous work that focused on routing (using BGP as an API) could be considered one instantiation of SDN, for example, and architectures from various vendors (*e.g.*, Cisco ONE and JunOS SDK) represent other instantiations of SDN that differ from OpenFlow.

In search of control programs and use cases. Despite the initial excitement surrounding SDN, it is worth recognizing that SDN is merely a tool that enables innovation in network control. SDN neither dictates how that control should be designed nor solves any particular problem. Rather, researchers and network operators now have a platform at their disposal to help address longstanding problems in managing their networks and deploying new services. Ultimately, the success and adoption of SDN depends on whether it can be used to solve pressing problems in networking that were difficult or impossible to solve with earlier protocols. SDN has already proved useful for solving problems related to network virtualization, as we describe in the next section.

3. Network Virtualization

In this section, we discuss network virtualization, a prominent early “use case” for SDN. Network virtualization presents the abstraction of a network that is decoupled from the underlying physical equipment. Network virtualization allows multiple virtual networks to run over a shared infrastructure, and each virtual network can have a much simpler (more abstract) topology than the underlying physical network. For example, a Virtual Local Area Network (VLAN) provides the illusion of a single LAN spanning multiple physical subnets, and multiple VLANs can run over the same collection of switches and routers. Although network virtualization is conceptually independent of SDN, the relationship between these two technologies has become much closer in recent years.

We preface our discussion of network virtualization with three caveats. First, a complete history of network

virtualization would require a separate survey; we focus on developments in network virtualization that relate directly to innovations in programmable networking. Second, although network virtualization has gained prominence as a use case for SDN, the concept predates modern-day SDN and has in fact evolved in parallel with programmable networking. The two technologies are in fact tightly coupled: Programmable networks often presumed mechanisms for sharing the infrastructure (across multiple tenants in a data center, administrative groups in a campus, or experiments in an experimental facility) and supporting logical network topologies that differ from the physical network, both of which are central tenets of network virtualization. Finally, we caution that a precise definition of “network virtualization” is elusive, and experts naturally disagree as to whether some of the mechanisms we discuss (*e.g.*, slicing) represent forms of network virtualization. In this article, we define the scope of network virtualization to include any technology that facilitates hosting a virtual network on an underlying physical network infrastructure.

Network Virtualization before SDN. For many years, network equipment has supported the creation of virtual networks, in the form of VLANs and virtual private networks. However, only network administrators could create these virtual networks, and these virtual networks were limited to running the existing network protocols. As such, incrementally deploying new technologies proved difficult. Instead, researchers and practitioners resorted to running *overlay* networks, where a small set of upgraded nodes use tunnels to form their own topology *on top* of a legacy network. In an overlay network, the upgraded nodes run their own control-plane protocol, and direct data traffic (and control-plane messages) to each other by encapsulating packets, sending them through the legacy network, and decapsulating them at the other end. The Mbone (for multicast) [50], the 6bone (for IPv6) [43], and the X-Bone [77] were prominent early examples.

These early overlay networks consisted of dedicated nodes that ran the special protocols, in the hope of spurring adoption of proposed enhancements to the network infrastructure. The notion of overlay networks soon expanded to include any end-host computer that installs and runs a special application, spurred by the success of early peer-to-peer file-sharing applications (*e.g.*, Napster and Gnutella). In addition to significant research on peer-to-peer protocols, the networking research community reignited research on using overlay networks as a way to improve the network infrastructure, such as the work on Resilient Overlay Networks [4], where a small collection of communicating hosts form an overlay that reacts quickly to network failures and performance problems.

In contrast to active networks, overlay networks did not require any special support from network equipment or cooperation from the Internet Service Providers, mak-

ing them much easier to deploy. To lower the barrier for experimenting with overlay networks, researchers began building virtualized experimental infrastructures like PlanetLab [61] that allowed multiple researchers to run their own overlay networks over a shared and distributed collection of hosts. Interestingly, PlanetLab itself was a form of “programmable router/switch” active networking, but using a collection of servers rather than the network nodes, and offering programmers a conventional operating system (*i.e.*, Linux). These design decisions spurred adoption by the distributed-systems research community, leading to a significant increase in the role of experimentation with prototype systems in this community.

Based on the success of shared experimental platforms in fostering experimental systems research, researchers started advocating the creation of shared experimental platforms that pushed support for virtual topologies that can run custom protocols *inside* the underlying network [7, 62] to enable realistic experiments to run side-by-side with operational traffic. In this model, the network equipment itself “hosts” the virtual topology, harkening back to the early Tempest architecture [79] where multiple virtual ATM networks could co-exist on the same set of physical switches [79]; the Tempest architecture even allowed switch-forwarding behavior to be defined using software controllers, foreshadowing the work on control and data-plane separation.

The GENI [33, 60] initiative took the idea of a virtualized and programmable network infrastructure to a much larger scale, building a national experimental infrastructure for research in networking and distributed systems. Moving beyond experimental infrastructure, some researchers argued that network virtualization could form the basis of a future Internet that enables multiple network architectures to coexist at the same time (each optimized for different applications or requirements, or run by different business entities), and evolve over time to meet changing needs [27, 62, 73, 89].

Relationship of Network Virtualization to SDN. Network virtualization (an abstraction of the physical network in terms of a logical network) clearly does not require SDN. Similarly, SDN (the separation of a logically centralized control plane from the underlying data plane) does not imply network virtualization. Interestingly, however, a symbiosis between network virtualization and SDN has emerged, which has begun to catalyze several new research areas. SDN and network virtualization relate in three main ways:

- *SDN as an enabling technology for network virtualization.* Cloud computing brought network virtualization to prominence, because cloud providers need a way to allow multiple customers (or “tenants”) to share the same network infrastructure. Nicira’s Network Virtualization Platform (NVP) [53] offers this abstraction without

requiring any support from the underlying networking hardware. The solution is use overlay networking to provide each tenant with the abstraction of a single switch connecting all of its virtual machines. Yet, in contrast to previous work on overlay networks, each overlay node is actually an extension of the physical network—a software switch (like Open vSwitch [57,63]) that encapsulates traffic destined to virtual machines running on other servers. A logically centralized controller installs the rules in these virtual switches to control how packets are encapsulated, and updates these rules when virtual machines move to new locations.

- *Network virtualization for evaluating and testing SDNs.* The ability to decouple an SDN control application from the underlying data plane makes it possible to test and evaluate SDN control applications in a virtual environment before the application is deployed on an operational network. Mininet [41, 48] uses process-based virtualization to run multiple virtual OpenFlow switches, end hosts, and SDN controllers—each as a single process on the same physical (or virtual) machine. The use of process-based virtualization allows Mininet to emulate a network with hundreds of hosts and switches on a single machine. In such an environment, a researcher or network operator can develop control logic and easily test it on a full-scale emulation of the production data plane; once the control plane has been evaluated, tested, and debugged, it can then be deployed on the real production network.
- *Virtualizing (“slicing”) an SDN.* In conventional networks, virtualizing a router or switch is complicated, because each virtual component needs to run own instance of control-plane software. In contrast, virtualizing a “dumb” SDN switch is much simpler. The FlowVisor [68] system enables a campus to support a testbed for networking research on top of the same physical equipment that carries the production traffic. The main idea is to divide traffic flow space into “slices” (a concept introduced in earlier work on PlanetLab [61]), where each slice has a share of network resources and is managed by a different SDN controller. FlowVisor runs as a hypervisor, speaking OpenFlow to each of the SDN controllers and to the underlying switches. Recent work has proposed slicing control of home networks, to allow different third-party service providers (*e.g.*, smart grid operators) to deploy services on the network without having to install their own infrastructure [89]. More recent work proposes ways to present each “slice” of a software-defined network with its own logical topology [1, 22] and address space [1].

Myths and misconceptions. People often refer to supposed “benefits of SDN”—such as amortizing the cost of

physical resources or dynamically reconfiguring networks in multi-tenant environments—that actually come from network virtualization. Although SDN facilitates network virtualization and may thus make some of these functions easier to realize, it is important to recognize that the capabilities that SDN offers (*i.e.*, the separation of data and control plane, abstractions for distributed network state) do not directly provide these benefits.

Exploring a broader range of use cases. Although SDN has enjoyed some early practical successes and certainly offers much-needed technologies in support of the specific use case of network virtualization, more work is needed both to improve the existing infrastructure and to explore SDN’s potential to solve problems for a much broader set of use cases. Although early SDN deployments focused on university campuses [34], data centers [53], and private backbones [44], recent work explores applications and extensions of SDN to a broader range of network settings, including home networks, enterprise networks, Internet exchange points, cellular core networks, cellular and WiFi radio access networks, and joint management of end-host applications and the network. Each of these settings introduces many new opportunities and challenges that the community will explore in the years ahead.

4. Conclusion

This paper has offered an intellectual history of programmable networks. The idea of a programmable network initially took shape as active networking, which espoused many of the same visions as SDN, but lacked both a clear use case and an incremental deployment path. After the era of active networking research projects, the pendulum swung from vision to pragmatism, in the form of separating the data and control plane to make the network easier to manage. This work focused primarily on better ways to route network traffic—a much narrower vision than previous work on active networking.

Ultimately, the work on OpenFlow and network operating systems struck the right balance between vision and pragmatism. This work advocated network-wide control for a wide range of applications, yet relied only on the existing capabilities of switch chipsets. Backwards compatibility with existing switch hardware appealed to many equipment vendors clamoring to compete in the growing market in data-center networks. The balance of a broad, clear vision with a pragmatic strategy for widespread adoption gained traction when SDN found a compelling use case in network virtualization.

As SDN continues to develop, we believe that history has important lessons to tell. First, SDN technologies will live or die based on “use pulls”. Although SDN is often heralded as the solution to all networking problems, it is worth remembering that SDN is just a *tool* for solving network-management problems more easily. SDN merely

places the power in our hands to develop new applications and solutions to longstanding problems. In this respect, our work is just beginning. If the past is any indication, the development of these new technologies will require innovation on multiple timescales, from long-term bold visions (such as active networking) to near-term creative problem solving (such as the operationally focused work on separating the control and data planes).

Second, we caution that the balance between vision and pragmatism remains tenuous. The bold vision of SDN advocates a wide variety of control applications; yet, OpenFlow's control over the data plane is confined to primitive match-action operations on packet-header fields. We should remember that the initial design of OpenFlow was driven by the desire for rapid adoption, not first principles. Supporting a wide range of network services would require much more sophisticated ways to analyze and manipulate traffic (e.g., deep-packet inspection, and compression, encryption, and transcoding of packets), using commodity servers (e.g., x86 machines) or programmable hardware (e.g., FPGAs, network processors, and GPUs), or both. Interestingly, the renewed interest in more sophisticated data-plane functionality, such as Network Functions Virtualization, harkens back to the earlier work on active networking, bringing our story full circle.

Maintaining SDN's bold vision requires us to continue thinking "out of the box" about the best ways to program the network, without being constrained by the limitations of current technologies. Rather than simply designing SDN applications with the current OpenFlow protocols in mind, we should think about what kind of control we want to have over the data plane, and balance that vision with a pragmatic strategy for deployment.

Acknowledgments

We thank Mostafa Ammar, Ken Calvert, Martin Casado, Russ Clark, Jon Crowcroft, Ian Leslie, Larry Peterson, Nick McKeown, Vyas Sekar, Jonathan Smith, Kobus van der Merwe, and David Wetherall for detailed comments, feedback, insights, and perspectives on this article.

REFERENCES

[1] A. Al-Shabibi. Programmable virtual networks: From network slicing to network virtualization, July 2013. <http://www.slideshare.net/nvirts/virt-july2013meetup>.

[2] D. Alexander, W. Arbaugh, A. Keromytis, and J. Smith. Secure active network environment architecture: Realization in SwitchWare. *IEEE Network Magazine*, pages 37–45, May 1998.

[3] D. S. Alexander, W. A. Arbaugh, M. W. Hicks, P. Kakkar, A. D. Keromytis, J. T. Moore, C. A. Gunter, S. M. Nettles, and J. M. Smith. The SwitchWare active network architecture. *IEEE Network*, 12(3):29–36, 1998.

[4] D. G. Andersen, H. Balakrishnan, M. F. Kaashoek, and R. Morris. Resilient Overlay Networks. In *Proc. 18th ACM Symposium on Operating Systems Principles (SOSP)*, pages 131–145, Banff, Canada, Oct. 2001.

[5] B. Anwer, M. Motiwala, M. bin Tariq, and N. Feamster. SwitchBlade: A Platform for Rapid Deployment of Network

Protocols on Programmable Hardware. In *Proc. ACM SIGCOMM*, New Delhi, India, Aug. 2010.

[6] A. Bavier, M. Bowman, D. Culler, B. Chun, S. Karlin, S. Muir, L. Peterson, T. Roscoe, T. Spalink, and M. Wawrzoniak. Operating System Support for Planetary-Scale Network Services. In *Proc. First Symposium on Networked Systems Design and Implementation (NSDI)*, San Francisco, CA, Mar. 2004.

[7] A. Bavier, N. Feamster, M. Huang, L. Peterson, and J. Rexford. In VINI Veritas: Realistic and Controlled Network Experimentation. In *Proc. ACM SIGCOMM*, Pisa, Italy, Aug. 2006.

[8] S. Bhattacharjee, K. Calvert, and E. Zegura. An architecture for active networks. In *High Performance Networking*, 1997.

[9] BIRD Internet routing daemon. <http://bird.network.cz/>.

[10] 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. *Communications Magazine, IEEE*, 36(10):64–70, 1998.

[11] P. Bosshart, G. Gibb, H. Kim, G. Varghese, N. McKeown, M. Izzard, F. Mujica, and M. Horowitz. Forwarding metamorphosis: Fast programmable match-action processing in hardware for SDN. In *ACM SIGCOMM*, Aug. 2013.

[12] M. Caesar, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe. Design and implementation of a routing control platform. In *Proc. 2nd USENIX NSDI*, Boston, MA, May 2005.

[13] K. Calvert. An architectural framework for active networks (v1.0). <http://protocols.netlab.uky.edu/calvert/arch-latest.ps>.

[14] K. Calvert. Reflections on network architecture: An active networking perspective. *ACM SIGCOMM Computer Communications Review*, 36(2):27–30, 2006.

[15] K. Calvert, S. Bhattacharjee, E. Zegura, and J. Sterbenz. Directions in active networks. *IEEE Communications Magazine*, pages 72–78, October 1998.

[16] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker. Ethane: Taking control of the enterprise. In *ACM SIGCOMM '07*, 2007.

[17] M. Casado, T. Garfinkel, M. Freedman, A. Akella, D. Boneh, N. McKeown, and S. Shenker. SANE: A protection architecture for enterprise networks. In *Proc. 15th USENIX Security Symposium*, Vancouver, BC, Canada, Aug. 2006.

[18] B. Chun, D. Culler, T. Roscoe, A. Bavier, L. Peterson, M. Wawrzoniak, and M. Bowman. Planetlab: an overlay testbed for broad-coverage services. *ACM SIGCOMM Computer Communication Review*, 33(3):3–12, 2003.

[19] M. Ciosi et al. Network functions virtualization. Technical report, ETSI, Darmstadt, Germany, Oct. 2012. http://portal.etsi.org/NFV/NFV_White_Paper.pdf.

[20] S. da Silva, Y. Yemini, and D. Florissi. The NetScript active network system. *IEEE Journal on Selected Areas in Communications*, 19(3):538–551, 2001.

[21] M. Dobrescu, N. Egi, K. Argyraki, B.-G. Chun, K. Fall, G. Iannaccone, A. Knies, M. Manesh, and S. Ratnasamy. RouteBricks: Exploiting parallelism to scale software routers. In *Proc. 22nd ACM Symposium on Operating Systems Principles (SOSP)*, Big Sky, MT, Oct. 2009.

[22] D. Drutskey, E. Keller, and J. Rexford. Scalable network virtualization in software-defined networks. *IEEE Internet Computing*, March/April 2013.

[23] D. Erickson. The Beacon OpenFlow controller. In *Proc. HotSDN*, Aug. 2013.

[24] D. Erickson et al. A demonstration of virtual machine mobility in an OpenFlow network, Aug. 2008. Demo at *ACM SIGCOMM*.

[25] A. Farrel, J. Vasseur, and J. Ash. *A Path Computation Element (PCE)-Based Architecture*. Internet Engineering Task Force, Aug. 2006. RFC 4655.

[26] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and K. van der Merwe. The case for separating routing from routers. In *ACM SIGCOMM Workshop on Future Directions in Network Architecture*, Portland, OR, Sept. 2004.

- [27] N. Feamster, L. Gao, and J. Rexford. How to lease the Internet in your spare time. *ACM SIGCOMM Computer Communications Review*, 37(1):61–64, 2007.
- [28] NSF Future Internet Design. <http://www.nets-find.net/>.
- [29] EU Future Internet Research and Experimentation Initiative. <http://www.ict-fire.eu/>.
- [30] Floodlight OpenFlow Controller. <http://floodlight.openflowhub.org/>.
- [31] FlowVisor. <http://www.openflowswitch.org/wk/index.php/FlowVisor>.
- [32] A. Fraser. Datakit—A Modular Network for Synchronous and Asynchronous Traffic. In *Proc. Int. Conf. on Communications*, 1980.
- [33] GENI: Global Environment for Network Innovations. <http://www.geni.net/>.
- [34] GENI. Campus OpenFlow topology, 2011. <http://groups.geni.net/geni/wiki/OpenFlow/CampusTopology>.
- [35] 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. *ACM SIGCOMM Computer Communications Review*, 35(5):41–54, 2005.
- [36] K. Greene. TR10: Software-defined networking. *MIT Technology Review*, March/April 2009. <http://www2.technologyreview.com/article/412194/tr10-software-defined-networking/>.
- [37] N. Gude, T. Koponen, J. Pettit, B. Pfaff, M. Casado, N. McKeown, and S. Shenker. NOX: Towards an operating system for networks. *ACM SIGCOMM Computer Communication Review*, 38(3):105–110, July 2008.
- [38] S. Han, K. Jang, K. Park, and S. Moon. PacketShader: a GPU-accelerated software router. In *Proc. ACM SIGCOMM*, New Delhi, India, Aug. 2010.
- [39] N. Handigol, M. Flajslik, S. Seetharaman, N. McKeown, and R. Johari. Aster*x: Load-balancing as a network primitive. In *ACLD '10: Architectural Concerns in Large Datacenters*, 2010.
- [40] M. Handley, E. Kohler, A. Ghosh, O. Hodson, and P. Radoslavov. Designing extensible IP router software. In *Proc. Networked Systems Design and Implementation*, May 2005.
- [41] B. Heller, N. Handigol, V. Jeyakumar, B. Lantz, and N. McKeown. Reproducible network experiments using container based emulation. In *Proc. ACM CoNEXT*, Dec. 2012.
- [42] B. Heller, S. Seetharaman, P. Mahadevan, Y. Yiakoumis, P. Sharma, S. Banerjee, and N. McKeown. ElasticTree: Saving energy in data center networks. Apr. 2010.
- [43] R. Hinden and J. Postel. *IPv6 Testing Address Allocation*. Internet Engineering Task Force, January 1996. RFC 1897, obsoleted by RFC 2471 on "6bone Phaseout".
- [44] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hlzle, S. Stuart, and A. Vahdat. B4: Experience with a globally deployed software defined WAN. In *ACM SIGCOMM*, Aug. 2013.
- [45] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek. The Click modular router. *ACM Transactions on Computer Systems*, 18(3):263–297, Aug. 2000.
- [46] 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 *OSDI*, volume 10, pages 1–6, 2010.
- [47] T. V. Lakshman, T. Nandagopal, R. Ramjee, K. Sabnani, and T. Woo. The SoftRouter Architecture. In *Proc. 3rd ACM Workshop on Hot Topics in Networks (Hotnets-III)*, San Diego, CA, Nov. 2004.
- [48] B. Lantz, B. Heller, and N. McKeown. A network in a laptop: Rapid prototyping for software-defined networks (at scale!). In *Proc. HotNets*, Oct. 2010.
- [49] J. Lockwood, N. McKeown, G. Watson, G. Gibb, P. Hartke, J. Naous, R. Raghuraman, and J. Luo. NetFPGA: An open platform for gigabit-rate network switching and routing. In *IEEE International Conference on Microelectronic Systems Education*, pages 160–161, 2007.
- [50] M. R. Macedonia and D. P. Brutzman. Mbone provides audio and video across the internet. *Computer*, 27(4):30–36, 1994.
- [51] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. OpenFlow: Enabling innovation in campus networks. *ACM SIGCOMM Computer Communications Review*, Apr. 2008.
- [52] A. Nayak, A. Reimers, N. Feamster, and R. Clark. Resonance: Dynamic access control in enterprise networks. In *Proc. Workshop: Research on Enterprise Networking*, Barcelona, Spain, Aug. 2009.
- [53] Nicira. It's time to virtualize the network, 2012. <http://nicira.com/en/network-virtualization-platform>.
- [54] NSF Guidelines for Planning and Managing the Major Research Equipment and Facilities Construction (MREFC) Account. <http://www.nsf.gov/bfa/docs/mrefcguidelines1206.pdf>, Nov. 2005.
- [55] Open Networking Foundation. <https://www.opennetworking.org/>.
- [56] ON.Lab. ONOS: Open network operating system, 2013. <http://tinyurl.com/pjs9eyw>.
- [57] Open vSwitch. openvswitch.org.
- [58] Open Daylight. <http://www.opendaylight.org/>.
- [59] 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. *IEEE Computer*, 39(9):102–105, 2006.
- [60] 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. *IEEE Computer*, Sept. 2006.
- [61] L. Peterson, T. Anderson, D. Culler, and T. Roscoe. A blueprint for introducing disruptive technology into the Internet. In *Proc. 1st ACM Workshop on Hot Topics in Networks (Hotnets-I)*, Princeton, NJ, Oct. 2002.
- [62] L. Peterson, S. Shenker, and J. Turner. Overcoming the Internet impasse through virtualization. In *Proc. 3rd ACM Workshop on Hot Topics in Networks (Hotnets-III)*, San Diego, CA, Nov. 2004.
- [63] B. Pfaff, J. Pettit, K. Amidon, M. Casado, T. Koponen, and S. Shenker. Extending networking into the virtualization layer. In *Proc. HotNets*, Oct. 2009.
- [64] POX. <http://www.noxrepo.org/pox/about-pox/>.
- [65] Quagga software routing suite. <http://www.quagga.net/>.
- [66] J. Salim, H. Khosravi, A. Kleen, and A. Kuznetsov. *Linux Netlink as an IP Services Protocol*. Internet Engineering Task Force, July 2003. RFC 3549.
- [67] B. Schwartz, A. W. Jackson, W. T. Strayer, W. Zhou, R. D. Rockwell, and C. Partridge. Smart packets for active networks. In *IEEE Conference on Open Architectures and Network Programming*, pages 90–97. IEEE, 1999.
- [68] R. Sherwood, G. Gibb, K.-K. Yap, G. Appenzeller, M. Casado, N. McKeown, and G. Parulkar. Can the production network be the testbed? In *Proc. 9th USENIX OSDI*, Vancouver, Canada, Oct. 2010.
- [69] J. Smith et al. SwitchWare: Accelerating network evolution. Technical Report MS-CIS-96-38, CIS Department, University of Pennsylvania, 1996.
- [70] J. M. Smith, K. L. Calvert, S. L. Murphy, H. K. Orman, and L. L. Peterson. Activating networks: A progress report. 32(4):32–41, Apr. 1999.
- [71] J. M. Smith and S. M. Nettles. Active networking: One view of the past, present, and future. *IEEE Transactions on Systems, Man, and Cybernetics: Part C: Applications and Reviews*, 34(1), February 2004.
- [72] T. Spalink, S. Karlin, L. Peterson, and Y. Gottlieb. Building a robust software-based router using network processors. In *Proc. SOSP*, Dec. 2001.
- [73] D. Taylor and J. Turner. Diversifying the Internet. In *Proc. IEEE GLOBECOM*, Nov. 2005.

- [74] D. E. Taylor, J. S. Turner, J. W. Lockwood, and E. L. Horta. Dynamic hardware plugins: Exploiting reconfigurable hardware for high-performance programmable routers. *Computer Networks*, 38(3):295–310, 2002.
- [75] D. Tennenhouse, J. Smith, W. D. Sincoskie, D. Wetherall, and G. Minden. A survey of active network research. *IEEE Communications Magazine*, pages 80–86, January 1997.
- [76] D. L. Tennenhouse and D. J. Wetherall. Towards an Active Network Architecture. *ACM SIGCOMM Computer Communications Review*, 26(2):5–18, Apr. 1996.
- [77] J. Touch. Dynamic Internet overlay deployment and management using the X-Bone. *Computer Networks*, 36(2):117–135, 2001.
- [78] J. van der Merwe, A. Cepleanu, K. D’Souza, B. Freeman, A. Greenberg, et al. Dynamic connectivity management with an intelligent route service control point. In *ACM SIGCOMM Workshop on Internet Network Management*, Oct. 2006.
- [79] J. van der Merwe, S. Rooney, I. Leslie, and S. Crosby. The Tempest: A Practical Framework for Network Programmability. *IEEE Network*, 12(3):20–28, May 1998.
- [80] J. E. van der Merwe, S. Rooney, L. Leslie, and S. Crosby. The Tempest: A practical framework for network programmability. *IEEE Network*, 12(3):20–28, 1998.
- [81] P. Verkaik, D. Pei, T. Scholl, A. Shaikh, A. Snoeren, and J. van der Merwe. Wresting control from BGP: Scalable fine-grained route control. In *Proc. USENIX Annual Technical Conference*, June 2007.
- [82] A. Voellmy and P. Hudak. Nettle: Functional reactive programming of OpenFlow networks. In *Proc. Workshop on Practical Aspects of Declarative Languages*, pages 235–249, Jan. 2011.
- [83] R. Wang, D. Butnariu, and J. Rexford. OpenFlow-based server load balancing gone wild. In *Hot-ICE*, Mar. 2011.
- [84] D. Wetherall, J. Guttag, and D. Tennenhouse. ANTS: A toolkit for building and dynamically deploying network protocols. In *IEEE OpenArch*, April 1998.
- [85] B. White, J. Lepreau, L. Stoller, R. Ricci, S. Guruprasad, M. Newbold, M. Hibler, C. Barb, and A. Joglekar. An integrated experimental environment for distributed systems and networks. In *Proc. OSDI*, Dec. 2002.
- [86] T. Wolf and J. Turner. Design issues for high performance active routers. *IEEE Journal on Selected Areas of Communication*, pages 404–409, March 2001.
- [87] H. Yan, D. A. Maltz, T. S. E. Ng, H. Gogineni, H. Zhang, and Z. Cai. Tesseract: A 4D Network Control Plane. In *Proc. 4th USENIX NSDI*, Cambridge, MA, Apr. 2007.
- [88] L. Yang, R. Dantu, T. Anderson, and R. Gopal. *Forwarding and Control Element Separation (ForCES) Framework*. Internet Engineering Task Force, Apr. 2004. RFC 3746.
- [89] Y. Yiakoumis and N. McKeown. Slicing Home Networks. In *ACM SIGCOMM Workshop on Home Networking (Homenets)*, Toronto, Ontario, Canada, May 2011.
- [90] J. Zander and R. Forchheimer. Softnet: An approach to higher level packet radio. In *Proceedings, AMRAD Conference, San Francisco*, 1983.