

# SDN Controllers: Benchmarking & Performance Evaluation

Liehuang Zhu, *Member, IEEE*, Md Monjurul Karim, Kashif Sharif, *Member, IEEE*, Fan Li, *Member, IEEE*, Xiaojiang Du, *Senior Member, IEEE*, and Mohsen Guizani, *Fellow, IEEE*

**Abstract**—Software Defined Networks offer flexible and intelligent network operations by splitting a traditional network into a centralized control plane and a programmable data plane. The intelligent control plane is responsible for providing flow paths to switches and optimizes network performance. The controller in the control plane is the fundamental element used for all operations of data plane management. Hence, the performance and capabilities of the controller itself are extremely important. Furthermore, the tools used to benchmark their performance must be accurate and effective in measuring different evaluation parameters. There are dozens of controller proposals available in existing literature. However, there is no quantitative comparative analysis for them. In this article, we present a comprehensive qualitative comparison of different SDN controllers, along with a quantitative analysis of their performance in different network scenarios. More specifically, we categorize and classify 34 controllers based on their capabilities, and present a qualitative comparison of their properties. We also discuss in-depth capabilities of benchmarking tools used for SDN controllers, along with best practices for quantitative controller evaluation. This work uses three benchmarking tools to compare nine controllers against multiple criteria. Finally, we discuss detailed research findings on the performance, benchmarking criteria, and evaluation testbeds for SDN controllers.

**Index Terms**—Software Defined Network, SDN Controller, Benchmarking, Performance.

## I. INTRODUCTION

SOFTWARE Defined Networks (SDN) have seen tremendous growth and deployment in different types of networks in recent times. They are being actively used in datacenter networks [1], [2], wireless & Internet of Things (IoT) networks [3], [4], wide area & cellular networks [5], as well as security and privacy of domains [6]. Compared to traditional networks it decouples the control logic from network layer devices, and centralizes it for efficient traffic forwarding and flow management across the domain. This multi-layered architecture, as shown in Figure 1, has data forwarding devices at the bottom in data plane, which are programmed by controllers in the control plane. The high level application or management plane interacts with control layer to program the whole network and enforce different policies.

L. Zhu, M. Karim, K. Sharif, and F. Li are with Beijing Engineering Research Center for Massive Language Information Processing & Cloud Computing Application, and School of Computer Science and Technology, Beijing Institute of Technology, Beijing, China. Email: {liehuangz, mkarim, kashif, fli}@bit.edu.cn

Xiaojiang Du is with Department of Computer and Information Sciences, Temple University, Philadelphia, USA. Email: dux@temple.edu

Mohsen Guizani is with Department of Electrical and Computer Engineering, University of Idaho, Moscow, ID, USA. Email: mguizani@uidaho.edu

K. Sharif is the corresponding author.

The interaction among these layers is done through interfaces which work as communication/programming protocols.

Traditional networks suffer from a number of limitations, mainly due to diverse service requirements and the scale of the network. Some of these are related to traffic engineering, flow management, policy enforcement, security, and virtualization [7]–[11]. SDN presents a simplified, centralized, and efficient solution to these, by decoupling the data plane forwarding and control plane intelligence. Hence, the network switched become simple forwarding devices, which route data traffic based on instruction from a softwarized controller. This centralized entity provides a programmatic control of whole network and enables real-time control of underlying devices. By using SDN, network management becomes straightforward and helps in removing rigidity from the network.

Some of the well known controllers are NOX [12], POX [13], Floodlight [14], OpenDaylight (ODL) [15], Open Network Operating System (ONOS) [16] and RYU [17]. However, a number of other controllers and flavors are available in the literature. From a practical implementation perspective, it is very difficult to determine which controller will perform best in any given type of network. Hence, the qualitative

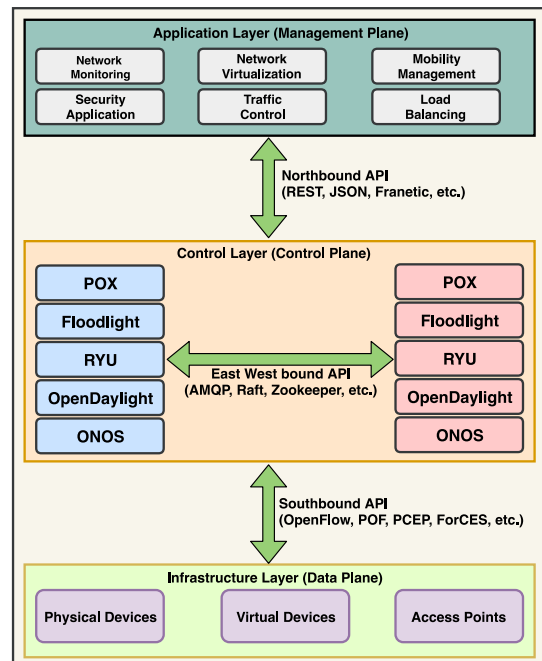


Fig. 1: Elements in a layered structure of SDN.

and quantitative comparative analysis of these controllers is very important. To the best of our knowledge, there is no such work which compares the controllers for their properties and evaluates their performance. Although a number of surveys have been done for SDN in general, there are none which provide a comprehensive controller evaluation. Works in [18]–[31] present some quantitative comparison, however, most of them either feature a specific application or simple environment to execute multiple experiments. In this work, we have adopted a different method by using a number of different benchmarking tool specifically developed for controller evaluations. The contributions of this work are multi-fold:

- We present the generic architecture of SDN controller and the evolution of modern SDN controllers.
- We present a qualitative comparative analysis of 34 different controllers for their properties and capabilities. We also discuss the different use cases for these controllers and the enhancements done to improve their performance by other works.
- We present a comprehensive study of benchmarking techniques and tools for SDN controllers. This includes the existing works & approaches used for evaluation, capabilities of benchmarking tools, and most importantly the details of metrics which should be used for quantitative evaluations.
- We conduct quantitative analysis of 9 different controllers using 3 different benchmarking tools for a variety of metrics. The results presented show the actual performance of controllers.
- We present comprehensive discussion on research findings not only for controller behavior but also for the metrics and tools used.

The rest of the paper is organized as follows: Section II gives an overview of SDN controllers, followed by comparison and classification of controllers in Section III. Benchmarking metrics and existing efforts are detailed in Section IV. Benchmarking tools and their properties are evaluated in Section V. Experimental results and research findings are detailed in section VI and VII respectively. Section VIII concludes the paper.

## II. SDN CONTROLLERS

A controller is the core component of any SDN infrastructure, as it has the global view of entire network including data plane SDN devices. It connects these resources with management applications, and performs flow actions dictated by application policy among the devices. In this section, we present the generic architecture of the controllers, and the evolution towards modern controllers. We also present the classification, comparison, and use case enhancements for 34 different controllers.

### A. Architecture of SDN Controllers

The controller in a software defined network, also referred as Network Operating Systems (NOS), is the core and critical component responsible for making decisions on managing

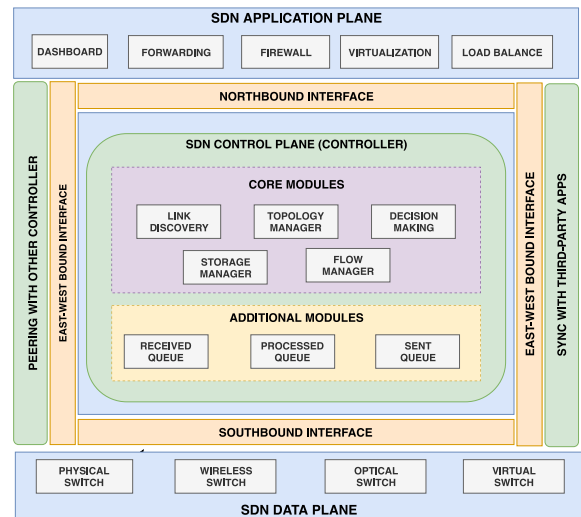


Fig. 2: General Overview of SDN Controller

traffic in underlying network. The proposals put forth for different controllers in literature do not modify the basic controller architecture, rather they differ in terms of modules and capabilities. Hence, we find that presenting individual architectures to be less useful for the reader. Here, we present the general architecture as shown in Figure 2, and discuss its different modules.

**Controller Core:** The core functions of the controller are mainly related to topology and traffic flow. The link discovery module regularly transmits inquiries on external ports utilizing packet\_out messages. These inquiry messages return in the form of packet\_in messages, which allows the controller to build the topology of network. The topology itself is maintained by the topology manager. This provides the decision making module to find optimal paths between nodes of the network. The paths are built such that the different QoS policies or security policies can be enforced during path installation. In addition, the controller may also have dedicated statistics collector/manager and queue manager for collecting performance information and management of different incoming and outgoing packet queues, respectively. Flow manager is one of the major modules which directly interacts with data plane’s flow entries and flow tables. It utilizes southbound interface for this purpose.

**Interfaces:** The core controller is surrounded by different interfaces for interaction with other layers and devices. Southbound Interface (SBI) defines a set of processing rules that enable packet forwarding between forwarding devices and controllers. SBI helps the controller to provision physical and virtual network devices intelligently. OpenFlow (OF) [32] is the most commonly used SBI and is a de-facto standard for industry. The fundamental responsibility of OF is to define flows and classify network traffic based on a predefined rule set. On the opposite end, the controller uses Northbound Interface (NBI) to allow developers to integrate their applications with controller and data plane devices. Controllers support a number of northbound APIs, but most of them are based on REST API. For inter controller communication, West Bound

Interface (WBI) is used. There is no standard communication interface for this purpose, hence different controllers use different mechanisms. Moreover, heterogeneous controllers do not usually communicate with each other. East Bound API (EBI) extends the capability of controller to interact with legacy routers. BGP [33] is the most commonly used protocol for this purpose.

### B. Evolution of SDN Controllers

Modern SDN controllers and SDN design is not the first attempt at centralizing the network control. From mid-2000s, several attempts have been made to separate the control logic from the data plane.

SoftRouter [34] and ForCES [35] were introduced in a single network device to separate control elements (CEs) from forwarding elements (FEs). However, they were limited to packet modification functionalities, as most of the routers (at the time) were limited in computing intelligence or network awareness to perform required operations. Routing Control Platforms (RCP) [36] was proposed as an intra-AS (Autonomous System) platform to implement an expandable control platform for BGP. However, the solution is for heterogeneous networks and prone to single point of failure. Path Computation Engine (PCE) [37] was presented to enable clients to execute path computations in routers but lacks dedicated centralized path computation engine and fails to provide cooperation among different entities. Although Intelligent Route Service Control Point (IRSCP) [38] introduces path allocation module in an external router and provides dynamic connectivity feature to enhance traffic flows throughout a network, it was limited to single ISP service. On the other hand, 4D project [39] was intended as a clean-state solution to introduce a control plane for topology discovery and to provide traffic forwarding logic and rule sets. However, there is no practical implementation of this approach. The SANE project [40] was developed by National Science Foundation (NSF) to enable traffic forwarding and access control policies using logically centralized server within enterprise networks. Ethane [41] is the successor of the SANE project that brings a more improved and practical control management module, aware of global network and performs routing operations based on pre-defined flows. Both, SANE and Ethane fail to acknowledge the network components as an overall representation. Besides, they also lack flow-level control over traditional routing approaches.

The control plane of these earlier proposals is missing a broad range of matching header fields and also lacks a wide range of functionalities. As a result, SDN has become mainstream with the introduction of OpenFlow [32] which is a data-plane Application Programming Interface (API), and a robust centralized controller named NOX [12]. OpenFlow is different from previous solutions as it is an open protocol to favor software developers to build applications on different switches that support flow tables with an extensible range of header fields. SDN brings flexibility and agility by allowing virtualization of the servers, rapid response to network changes, deployment of policies, and centralized control over complete network.

## III. CLASSIFICATION AND COMPARISON OF SDN CONTROLLERS

In order to compare different SDN controllers, we have performed an extensive search of proposals not only in academic literature, but also in commercial domain. Here, we first present the possible classification criteria of controllers, followed by the comparative analysis, and then different use case specific enhancements.

### A. Classification & Selection Criteria

The working of controllers is more or less same across all the proposals listed in Table I. After analysis of 34 controllers we conclude that the working, role, and responsibilities of majority of these do not present any classification basis. Perhaps the only classification criteria that can be used is the *deployment architecture*. The initial aim of SDN was to centralize the control plane, hence most of the controllers utilized a single controller, however, this created single point of failure and scalability challenges. The distributed architecture allows usage of multiple controllers inside a domain, working in a flat or hierarchical formation.

In this work, we have not limited the selection of controllers to any specific criteria. Rather we have collected all possible controllers from literature and other documented projects. To the best of our knowledge, there is no other work that collects and compares such a large number of controllers.

### B. Qualitative Comparison

Table I presents a comprehensive view of different properties of the controllers. In the interest of space and the fact that not all proposals provide extensive details about their inner-workings, we do not discuss each controller individually. Rather we present the properties and design choices of controllers.

**Programming Language:** Controllers have been written using different programming languages, such as C, C++, Java, JavaScript, Python, Ruby, Haskell, Go, and Erlang. In some cases, the entire controller is built using a single language. While in many other controllers multiple languages are used in their core and modules, so that they can offer efficient memory allocation, can be executed on multiple platforms, or most importantly achieve higher performance under certain conditions.

**Architecture:** The major design decision of a controller is its architecture, which can be centralized or distributed. Centralized controllers are mostly used in small scale networks, whereas distributed controllers are able to span across multiple domains. They can further be classified into flat, where all controller instances have equal responsibilities, or hierarchical, where a root controller is present.

**Programmable Interface (API):** Generally, Northbound API (NBI) allows the controller to facilitate applications like topology monitoring, flow forwarding, network virtualization, load-balancing, and intrusion detection based on the network events which are generated by data plane devices. On the other hand, low-level API like Southbound API (SBI) is responsible

TABLE I: SDN Controller's Feature Comparison Table

Name	Programming Language	Architecture	Northbound API	Southbound API	EastWestbound API	Supported Platform	Interface	License	Multithreading	Modularity	Consistency	Documentation
Beacon [42]	Java	Centralized	ad-hoc	OpenFlow 1.0	-	Linux, MacOS, Windows	CLI, Web UI	GPL 2.0	Yes	Fair	No	Fair
Beehive [43]	Go	Distributed Hierarchical	REST	OpenFlow 1.0, 1.2	-	Linux	CLI	Apache 2.0	Yes	Good	Yes	Limited
DCFabric [44]	C, Javascript	Centralized	REST	OpenFlow 1.3	-	Linux	CLI, Web UI	LGPL 3.0	Yes	Good	Yes	Fair
Disco [45]	Java	Distributed Flat	REST	OpenFlow 1.0	AMQP	-	-	Proprietary	-	Good	No	Limited
Faucet [46]	Python	Centralized	-	OpenFlow 1.3	-	Linux	CLI, Web UI	Apache 2.0	Yes	-	Yes	Good
Floodlight [14]	Java	Centralized	REST, Java RPC, Quantum	OpenFlow 1.0, 1.3	-	Linux, MacOS, Windows	CLI, Web UI	Apache 2.0	Yes	Fair	Yes	Good
FlowVisor [47]	C	Centralized	JSON RPC	OpenFlow 1.0, 1.3	-	Linux	CLI	Proprietary	-	-	No	Fair
HyperFlow [48]	C++	Distributed Flat	-	OpenFlow 1.0	Publish and subscribe messages	-	-	Proprietary	Yes	Fair	No	Limited
Kandoo [49]	C, C++, Python	Distributed Hierarchical	Java RPC	OpenFlow 1.0-1.2	Messaging Channel	Linux	CLI	Proprietary	Yes	High	No	Limited
Loom [50]	Erlang	Distributed Flat	JSON	OpenFlow 1.3-1.4	-	Linux	CLI	Apache 2.0	Yes	Good	No	Good
Maestro [51]	Java	Centralized	ad-hoc	OpenFlow 1.0	-	Linux, MacOS, Windows	Web UI	LGPL 2.1	Yes	Fair	No	Limited
McNettle [52]	Haskell	Centralized	-	OpenFlow 1.0	-	Linux	CLI	Proprietary	Yes	Good	No	Limited
Meridian [53]	Java	Centralized	REST	OpenFlow 1.0, 1.3	-	Cloud-based	Web UI	-	Yes	Good	No	Limited
Microflow [54]	C	Centralized	Socket	OpenFlow 1.0-1.5	-	Linux	CLI, Web UI	Apache 2.0	Yes	-	No	Limited
NodeFlow [55]	JavaScript	Centralized	JSON	OpenFlow 1.0	-	Node.js	CLI	Cisco	-	-	No	Limited
NOX [12]	C++	Centralized	ad-hoc	OpenFlow 1.0	-	Linux	CLI, Web UI	GPL 3.0	Yes (Nox-MT)	Low	No	Limited
Onix [56]	C++	Distributed Flat	Onix API	OpenFlow 1.0, OVSDB	Zookeeper	-	-	Proprietary	Yes	Good	No	Limited
ONOS [16]	Java	Distributed Flat	REST, Neutron	OpenFlow 1.0, 1.3	Raft	Linux, MacOS, Windows	CLI, Web UI	Apache 2.0	Yes	High	Yes	Good
OpenContrail [57]	C, C++, Python	Centralized	REST	BGP, XMPP	-	Linux	CLI, Web UI	Apache 2.0	Yes	High	Yes	Good
OpenDaylight [15]	Java	Distributed Flat	REST, RESTCONF, XMPP, NETCONF	OpenFlow 1.0, 1.3	Akka, Raft	Linux, MacOS, Windows	CLI, Web UI	EPL 1.0	Yes	High	Yes	Good
OpenIRIS [58]	Java	Distributed Flat	REST	OpenFlow 1.0-1.3	Custom Protocol	Linux	CLI, Web UI	Apache 2.0	Yes	Fair	No	Limited
OpenMul [59]	C	Centralized	REST	OpenFlow 1.0, 1.3, OVSDB, Netconf	-	Linux	CLI	GPL 2.0	Yes	High	No	Good
PANE [60]	Haskell	Distributed Flat	PANE API	OpenFlow 1.0	Zookeeper	Linux, MacOS	CLI	BSD 3.0	-	Fair	No	Fair
POF Controller [61]	Java	Centralized	-	OpenFlow 1.0, POF-FIS	-	Linux	CLI, GUI	Apache 2.0	-	-	No	Limited
POX [13]	Python	Centralized	ad-hoc	OpenFlow 1.0	-	Linux, MacOS, Windows	CLI, GUI	Apache 2.0	No	Low	No	Limited
Ravel [62]	Python	Centralized	ad-hoc	OpenFlow 1.0	-	Linux	CLI	Apache2.0	-	-	Yes	Fair
Rosemary [63]	C	Centralized	ad-hoc	OpenFlow 1.0, 1.3, XMPP	-	Linux	CLI	Proprietary	Yes	Good	No	Limited
RunOS [64]	C++	Distributed Flat	REST	OpenFlow 1.3	Maple	Linux	CLI, Web UI	Apache2.0	Yes	High	Yes	Fair
Ryu [17]	Python	Centralized	REST	OpenFlow 1.0-1.5	-	Linux, MacOS	CLI	Apache 2.0	Yes	Fair	Yes	Good
SMaRtLight [65]	Java	Distributed Flat	REST	OpenFlow 1.3	BFT-SMaRt	Linux	CLI	Proprietary	-	-	No	Limited
TinySDN [66]	C	Centralized	-	OpenFlow 1.0	-	Linux	CLI	BSD 3.0	No	-	No	Limited
Trema [67]	C, Ruby	Centralized	ad-hoc	OpenFlow 1.0	-	Linux	CLI	GPL 2.0	-	Good	No	Fair
Yanc [68]	C, C++	Distributed Flat	REST	OpenFlow 1.0-1.3	yanc File System	Linux	CLI	Proprietary	-	-	No	Limited
ZeroSDN [69]	C++	Distributed Flat	REST	OpenFlow 1.0, 1.3	ZeroMQ	Linux	CLI, Web UI	Apache 2.0	-	High	Yes	Fair

for enabling the communication between a controller and SDN enabled switches or routers. Additionally, east-west API (EWBI) is used by multiple controllers from different domains to form peering with each other in a distributed or hierarchical environment. Not all controllers provide all APIs, and only select few have customized them for their own specific use.

**Platform and Interface:** These properties describe the implementation of controller to be compatible with specific operating system. Majority of controllers are built on top of Linux distributions. Moreover, in order to configure and view statistical information, some controllers provide graphical or web based interfaces to the administrators.

**Threading and Modularity:** A single-threaded controller is more suitable for lightweight SDN deployments. In contrast, multi-threaded controllers are suitable for commercial purposes such as, 5G, SDN-WAN, and optical networks. On the other hand, a controller's modularity allows the integration of different applications and functionalities. High modularity allows a controller to perform faster task execution in a distributed environment.

**License, Availability, and Documentation:** Most of the controllers discussed in this article are licensed as Open-

Source. However, a few have a proprietary license which means they are only available through special request or for research purpose. Regular maintenance of these controllers is also a challenging task for the developers which is why a number of them do not receive regular updates. Nevertheless, the source code is available online which allows anyone to make further changes according to the requirements. While accessing them online, we have found that the majority of them lack proper documentation. On the contrary, the ones which are updated on a regular basis feature detailed and updated documentation for all the available version and also include community-based support.

### C. Use case Specific Enhancements to SDN Controllers

The adoption of different controllers and SDN in general, has also triggered enhancements and use case specific improvements for different controllers. Here, we have grouped these enhancements into different categories, and summarize how they improve the capabilities of controllers.

1) *Network Monitoring:* Network monitoring has become one of the most vital use cases of SDN controllers. SDN controller can take advantage of the global view of topology

and proactively query the performance. OpenTM [70] was proposed by as a module for NOX, one of the earliest open-source OpenFlow controller. This monitoring scheme evaluates Traffic Matrix (TM) of OpenFlow switches with a consistent polling rate. However, this also leads to higher monitoring overhead. Adrichem et al. [71] presented OpenNetMon, a Python-based module for POX controller to monitor end-to-end per flow QoS metrics like throughput, delay, packet loss, etc. From the statistical analysis results, the approach for monitoring throughput is excellent, although continuous polling of information make cause overhead on the controller. Flow monitoring is limited to edge switches only. On the other hand, Payless [72] implemented over Floodlight controller is another query-based monitoring framework that can request the desired QoS metrics using a set of well-defined RESTful APIs. However, some trade-off between accuracy and overhead can lead to slight performance degradation for different polling intervals. SDN Interactive Manager [73] and OFMon [74] are two recent implementation of network monitoring modules that have been built over Floodlight and ONOS controller respectively.

**2) Load Balancing:** SDN controller plays an important role to enable load balancing in distributed systems by optimizing resource allocation, minimizing response time, and maximizing throughput of that system. Without rewriting IP addresses, Handigol et al. [75] implemented a method where NOX controller can be used along with OpenFlow switch reactively to reduce response time for load balancing of multiple web servers. Contrarily, Uppal et al. [76] used address rewriting techniques for NOX-based load balancer which cuts down cost and brings flexibility. Another NOX-based proactive load balancer was proposed by Wang et al. [77] which uses OpenFlow wild card rules that can achieve faster adaptation with new load balancing weights and to redistribute the existing weight more efficiently. Based on switch migration technique Liang et al. in [78] presented a dynamic load balancing method that has been implemented over cluster OpenDaylight controller [15]. However, this method may fail in large scale networks due to coordinator node's recurring load collection issue.

**3) Network Virtualization & Cloud Orchestration:** With addition of Network virtualization (NV) techniques SDNs have gained a new dimension. This has allowed network slicing and multi-tenant hosting on existing physical network resources. FlowVisor [47] is the most popular SDN based implementation to utilize virtual networks by leveraging OpenFlow functionality to abstract the underlying hardware. VeRTIGO [79] is an extension of FlowVisor that provides the controllers to choose the depth of virtual network abstraction required. This extension increases more flexibility in provisioning SDNs, however at the cost of hypervisor complexity. In order to reduce complexity of network management, Xingtao et al. [80] presented an SDN controller built on docker [81] to improve the deployment speed with expanded mobility. In [82] the flexibility of NOX controller has been used as a container-based controller virtualization module to effectively cache and manage mappings between virtual networks and physical switches. HyperFlex [83] proposes a control plane virtualization model which largely aims at achieving scalabil-

ity, privacy, and extensibility. In this architecture, FlowVisor and Ryu controllers have been combined to provide the core hypervisor functions and to control the hypervisor network respectively.

Cloud orchestration defines the integration of SDN controllers with a cloud based resource manager, such as OpenStack [84] to enable dynamic interworking between data centers, wide area networks, transport network, and other enterprise networks. In [85], OpenDaylight is integrated with OpenStack Havana [86] to evaluate the effectiveness of SDN in a cloud-based architecture where multiple data centers (DC) are located in different domains. In this architecture, the controller communicates with Havana using its REST NBI to perform critical tasks such as building, removal, and migration of virtual instances which are located in inter-DC and intra-DC environments.

**4) Policy Enforcement:** To enhance the security and flexible network management, an SDN controller has the capability to assign different policy decisions by implementing flow-based forwarding rules. Hinrichs et al. [87] implemented NOX as an application to provide access control, external authentication, and to enable policy enforcement along with network isolation. PANE [60] presents an API to allow administrators to install policies for bandwidth allocation, access control, and path control. Additionally, the API provides the capability to query the state of network or to provide information to SDN controller regarding future traffic characteristics. PolicyCop [88] based on Floodlight controller, is an autonomic QoS policy enforcement architecture, that presents an interface for specifying QoS requirements in Service Layer Arguments and implementation through the OpenFlow API. Besides, it can monitor different policies so that control plane rules can be modified with changing traffic conditions autonomously. An extra module of ONOS controller has been extended to implement a policy-based secure framework in [89]. The authors allowed an end-to-end SDN services across various domains including inter and intra domain, using a wild card based policy language which includes a group of entities and services. Associated action such as acceptance or denial of a request is executed when a policy statement is satisfied.

#### IV. BENCHMARKING PROCESS & METRICS

Theoretical comparison based on features and properties do not reflect the actual performance of any controller. Hence, real deployment and benchmarking is necessary for true evaluation. In this section, we first present and overview on the necessity and importance of evaluating controllers. Following it, we discuss existing efforts for benchmarking along with important lessons learned. Finally, we present a list of performance metrics, which should be used in benchmarking of controllers.

##### A. Why Benchmark a Controller?

Prior to executing SDN-based operations, network administrators are required to verify whether available components can match their requirements to perform necessary tasks. Hence, evaluations related to data plane (vSwitches, links, etc.) may include tasks such as measurement of flow table

TABLE II: Comparative analysis of different benchmarking studies.

Reference	Testbed Specifications	Evaluation Tool Used	Controller(s) Evaluated	Evaluation Metrics	Optimization Objectives	Lessons Learned
[18]	1 × Quad-core & 1 × Octa-Core Server 2 Gbps Link Speed	CBench	NOX, NOX-MT, Beacon, Maestro	Throughput, Latency	Batching I/O Boost Async I/O	Number of switches impact the controller performance.
[19]	1 × Cluster with 2 Separate Xeon Servers 8 Gbps Link Speed	CBench	NOX-MT, Beacon, Maestro, Floodlight	Throughput, Latency, Threading Scalability, Delay Sensitivity	Switch Partitioning Packet Batching Task Batching	Switch partitioning & switch batching impacts throughput. Packet batching & task batching impacts delay sensitivity.
[20]	2 separate Xeon Servers 10 Gbps Link Speed	CBench Hcprobe	NOX, POX, Floodlight, Ryu, Mul, Beacon, Maestro	Throughput, Latency, Reliability, Security	Flow Modification Customized Workload	Scalability of controller depends on the number of cores. Not every controllers can handle heavy workload.
[21]	Single testbed with 4 servers (dual core) 100 Mbps Link Speed	OFCBenchmark	NOX, Floodlight, Maestro	Round Trip Time Send and Response Rate Packet Processing Rate	Implement Boost Libraries to handle Threads	Transmitting larger flows helps in detecting congestion in networks.
[22]	Not Specified	OFCProbe	NOX and Floodlight	Impact of Fat-tree Topology Load Balancing	Java library is used to handle OpenFlow connections	Topology has an impact on flow processing time. Efficient handling of switch depends on the characteristic of controller.
[23]	5 × Server with Core i5 CPU	CBench	Floodlight and OpenDaylight	Throughput, Latency, Failure	Not Specified	Custom profile is proposed for CBench. Controllers may suffer from memory leakages.
[24]	Not Specified	Analytic Hierarchy Process (AHP)	POX, Floodlight, OpenDaylight, Ryu and Trema	Virtual Switch Support, Modularity, Documentation, API Compatibility	Not Specified	Evaluation Method is Subjective Testing process may effect the outcome.
[25]	Single Testbed with Quad-Core Xeon Server	CBench, Open vSwitch	NOX, POX, Floodlight, Ryu, Beacon	Throughput, Latency, Threading Capability, Python Interpretation	Python Interpreter, Hyper-Threading (HT)	HT offers performance improvement for java-based controllers. Reliability, Trustworthiness, Usability, and Scalability should be considered equally.
[26]	1 × Quad-core, 1 × One Octa-core Testbed	Mininet, Open vSwitch, Indigo vSwitch	POX	CPU Utilization, Topology Impact, Ping Delay	Not Specified	Number of switches impact the flow installation time Mininet utilizes maximum system memory. Initial Ping Delay is larger than average Ping Delay.
[27]	1 × Multi-Core, 1 × Many-Core Testbed 10 Gbps Link Speed	CBench	NOX-MT, Floodlight, Beacon, Maestro	Latency, Throughput, Energy Consumption, I/O Threading Impact	Floodlight Learning Switch CBench Delay Parameter Maestro Config File Modification	Number of Switches and cores impact NOX's performance. CPU types and system architecture impact scalability.
[28]	Single Testbed with Octa-core CPU 10 Gbps Link Speed	CBench	NOX, POX, Floodlight, OpenDaylight, ONOS, Ryu, IRIS, Beacon, Maestro	Latency, Throughput	Not Specified	Controller's SBI allows additional support for future Internet architecture
[29]	Dual Core Virtual Testbed	Open vSwitch, Cluster Testbed, HTTP Generator, REST Client	OpenDaylight, ONOS	Flow Installation Rate Flow Reading Rate Failover Time	Controllers are customized for WAN environment	Size of a cluster has impact on flow installation rate. Failover Time of a controller depends on number of devices. Latency has significant impact on large-scale WAN.
[30]	Not Specified	Mininet, Open vSwitch, Traffic Generator	POX and Floodlight	Round Trip Delay, Average Throughput	Not Specified	Simple controllers better suited for configuration-related tasks. Feature-based controllers are good for performance-based tasks.
[31]	2 Xeon Testbeds	OFCProbe	ONOS	Topology Discovery Time, Path Provision Time, ASYN. Msg. Process Time	Not Specified	Number of links has equal impact as number of switches regarding performance. Reactive path provisioning time relies on length of corresponding path.

capacity, progressing times of OpenFlow messages, and bandwidth utilization, etc. Similarly, for the control plane it is equally essential to evaluate whether the controller is capable to efficiently manage the complete network, and utilize the capabilities of data plane to its maximum capacity. Although the fundamental function of a controller is flow management and installation, a number of different performance metrics can be used for its benchmarking. As there are numerous controllers available with different architectures and properties, it becomes extremely important to have a standard benchmarking criteria for evaluation.

In this regard, there are two basic requirements: a) a set of benchmarking metrics, and b) an efficient tool for benchmarking test. In [90], authors have presented a basic list of tests which should be conducted to evaluate the performance of a controller. However, there can be a number of other metrics which should also be used when benchmarking different controllers. Similarly, the tool used to perform the test in an emulated environment is critical.

### B. Existing Works & Lessons Learned

Prior to this article, [18]–[31] use multiple techniques, tools, and testbeds to evaluate the performance of several SDN controllers including scalability, reliability efficiency, and robustness.

In Table II, we compile most of the existing works associated with the evaluation of the controller performance and the major findings. Majority of these works use CBench [91], to evaluate the performance based on latency and throughput.

In most cases, throughput mainly correlates with threading capability of a controller, regarding the number of flows it can process in a specified time slot. Some other works extend CBench to integrate support with the operating system's kernel and compilers like Java and Python. The aim is to improve threading scalability of a controller regarding system's I/O modules. Some works include simulation-based environments where hosts and vSwitches are virtualized to evaluate the impact of topology on the performance of a controller. In these experiments, the load balancing functionality is extensively tested. Moreover, some works evaluate the reliability of the controller by generating vulnerable flows. Energy consumption has also been evaluated using fat-tree or data-center topologies. Below we give brief description of some of the notable works.

Authors in [18] present CBench [91] tool for evaluation of different controllers. They perform multiple flow-based experiments using it to compare the effectiveness and performance of NOX-MT, a multi-threaded adoption of NOX controller with other controllers like NOX, Beacon, and Maestro. Despite showing a notable improvement in performance, NOX-MT fails to identify some of the limitations of NOX such as massive utilization of Dynamic memory allocation and redundant representation of multiple requests.

In [19], authors compare four multi-threaded controllers (NOX-MT, Floodlight, Beacon, and Maestro) for architectural features like multi-core availability, controller impact on OF switch, packet batching, and task processing. Authors use CBench to compare these controllers based on their throughput and latency performance. In throughput mode, two scenarios

are considered including a fixed amount of switches with an increasing number of threads and fixed threads with an increasing number of switches. Beacon shows better performance in these two scenarios due to its ability to use the multi-core and multi-threading functionalities. Besides, the dynamic changing of packet sizes allows Maestro to perform better in latency test.

Work in [20] presented a framework named HCprobe to compare seven different SDN controllers: NOX, POX, Floodlight, Beacon, Ryu, MUL and Maestro. To compare the effectiveness of these controllers, the authors performed some additional measurements like scalability, reliability, and security along with latency and throughput. The testbed analysis presents some security vulnerabilities along with the reliability issues with MUL and Maestro controllers. On the other hand, Beacon, MUL, and Floodlight obtained minimum latency while Beacon performed relatively well in the throughput test.

Analytic Hierarchy Process (AHP) is used in [24] to analyze POX, Floodlight, OpenDaylight, Ryu, and Trema based on multiple standards like virtual switch support, modularity, documentation, programming language compatibility and availability of user interface. According to calculation, Ryu was elected to be the most suitable controller based on requirements as mentioned earlier. However, the AHP method is subjective and changing of measurements or scenarios may lead to a different outcome.

In [25], authors use multi-core and many-core testbeds to evaluate NOX, Maestro, Floodlight, Beacon on the aspect of multi-core utilization efficiency, performance scalability, and energy consumption regarding data center environments. The work emphasizes on existing controllers limitation in taking advantages of the concurrency in modern hardware.

In [28] the performance of well-known centralized and distributed SDN controllers has been studied using CBench. The results show that both MUL and Libfluid MSG (written in C) achieved the highest throughput under an increasing number of switches whereas python-based Ryu and POX obtained better score in latency mode. However, with the increasing number of threads, both Beacon and MUL performed better while python-based controllers failed to show satisfying performance.

### C. Benchmarking Metrics and their Impact

In this section we present a detailed list of performance metrics that can be used to benchmark SDN controllers. Table III outlines the grouping and description of each of these metrics. Some of these have also been identified by [90], however, we have extended this list and grouped them to eliminate the confusion regarding terminology. Generic terms such as, throughput and latency can have significantly different meaning depending on measurement process. Additionally, there can be other metrics to evaluate a controller, e.g. security, reliability, etc. However, we refer to them as non-measurable parameters which are more subjective in nature. We leave their classification as future work. The measurable parameters are grouped as following.

1) *Throughput Metrics*: Throughput is usually measured as a rate for processing flow requests by the controller. The important thing to note, is that it is not the flow installation

time (path provisioning). From the test tools perspective, it is the number of packet\_in messages sent and the corresponding packet\_out messages received per unit time. These requests could be synchronous or asynchronously coming from the vSwitches in real environment.

2) *Latency Metrics*: This group of metrics is measured in time units. Similar to throughput it only deals with the time between packets sent to controller and response received at the vSwitch. A number of factors can effect the latency of a controller, including computation time require by the controller and link delay.

3) *Flow Related Metrics*: These metrics deal with the complete path provisioning and flow installation. The primary difference between this and throughput is the complete path. Throughput only measures the rate from vSwitch to controller and back to vSwitch. However, complete flow installation requires installation of flow entries at other vSwitches along the path. We group both rate and time variants of these parameters in the same category, along with load balancing capability of the controller.

4) *Topology Based Metrics*: The ability to detect or determine a topology including its type (single, linear, overlay and tree), size and number of integrated nodes altogether represent a vital aspect to evaluate the efficiency of a controller. Interaction with its southbound interface also plays a significant role in these metrics.

5) *Threading & Session Metrics*: This set of metrics identifies controller competence with respect to utilizing the system architecture, hardware capabilities, and I/O units. Optimization of thread-based capabilities like multi-threading offers several advantages of task batching, event scheduling, process flows as groups and most importantly increases controller's flow processing time and rate.

6) *Miscellaneous Metrics*: Here we group other parameters which can also be used for evaluating the controllers. Some of these can be crucial in specialized scenarios. For example energy consumption in mobile environments where controllers are deployed on devices which are energy constrained. Similarly, in situations where hardware failure is a concern, the failover time needs to be reduced so that backup controllers can takeover as quickly as possible.

## V. TOOLS FOR CONTROLLER BENCHMARKING

Evaluating or benchmarking the performance of a controller can be done either through simulation/emulation or by using a hardware based testbed. Although, hardware testbeds provide measurements which are closer to actual values in production environment, however their cost is significant for research community. Hence emulation based evaluations are common practice. However, for benchmarking of SDN controllers, the software tool used has to be extremely efficient and precise. In this section, we present a number of well known tools available for benchmarking, followed by analysis for their properties and benchmarking capabilities.

TABLE III: Classification of Benchmarking Metrics and Tool Capabilities

Group	Measurable Metrics Parameters	Description	Benchmarking Tools		
			CBench	PktBlaster	OFNet
Throughput	Async Message Processing Rate	Determines number of flow requests a controllers can process per unit time. A processed request does not mean a successfully installed flow.	✓	✓	○
	Sync Message Processing Rate		✓	✓	○
	Send and Response Rate		×	○	✓
Latency	Async Message Processing Time	Denotes the delay or time duration between request from the vSwitch and response received back.	✓	✓	✓
	Sync Message Processing Time		✓	✓	✓
	Round Trip Time		×	○	✓
Flow Related	Path Provision Time (Proactive/Reactive)	Determines the efficiency of a controller to install flows, or measures which include communication between a source and destination.	×	✓	✓
	Path Provision Rate (Proactive/Reactive)		×	✓	○
	Flow Reading Rate		×	×	○
	Flow Installation Time		✓	✓	✓
	Load Balancing		×	○	×
Topology	Topology Discovery Time/Size	Measures the capability to discover topology or change in topology.	×	✓	✓
	Topology Change Time	This also indirectly measure the SBI performance.	×	✓	✓
Threading	Thread Capability	Indicates the utilization efficiency of a controller regarding the OS and physical hardware resources.	✓	✓	×
	I/O Impact		✓	✓	×
	Control Session Capability		×	○	×
	vSwitch CPU Utilization		×	×	✓
Others	Forwarding Table Capacity	Miscellaneous parameters which can be measured for specific scenarios.	×	✓	×
	Ping Delay Time		×	×	×
	Energy Consumption		×	×	×
	Network Re-provisioning Time		×	×	×
	Controller Failover Time		×	○	×

### A. Benchmarking Tools

Following are some of the commonly used tools for benchmarking. Table IV provides a comparative analysis of the three main tools used for evaluation in this work.

**CBench** [91] is one of the fundamental benchmarking tools with open-source license. It is designed explicitly for evaluating the performance of OpenFlow SDN controllers which support OpenFlow 1.0 and 1.3. However, due to compatibility limitation, controllers with OpenFlow 1.3 may experience performance issues. There are two basic evaluation metrics in CBench, i.e., Latency and Throughput. To measure Latency, the vSwitch forwards a single packet\_in message towards the controller and waits for a response. Tests can be repeated several times to obtain the average performance. The total number of acknowledgments obtained in a test period is used to compute the average latency. As for throughput measurement, each vSwitch continuously sends as many packet\_in messages as possible, to estimate the capability of the controller.

**HCprobe** [20] is an open-source extension of CBench, developed with the combination of Python and Shell scripts, to provide additional performance evaluation capabilities, such as reliability and scalability. The emulated switch can send vulnerable OpenFlow messages to controllers to check for resiliency and trustability. Besides, the test engine utilizes a Linux kernel, which allows customizable and scalable tuning of CPU threading. This allows the tester to obtain more accurate performance statistics of an SDN controller.

**WCBench** [92] is another variants of CBench built in Python and utilizes the core library module of CBench. Compared to CBench, feature set of this tool goes beyond latency and throughput, and offers additional aspects of automated evaluation with detailed and graphical statistics. Although it extends the support of OpenFlow to version 1.3, the compatibility of WCBench is still limited to specific versions of ODL controller.

**OFCBenchmark** [21] is built using C++ and Boost library to address some of the limitations of CBench. The components of this benchmarking tool include a graphical dashboard (built with Delphi), virtualized scalable vSwitch which is the core module, and includes a client that can administer evaluation tests. The tool offers distributed benchmarking by allowing clients to run in multiple instances, and offers extensible

benchmarking such as Round Trip Time (RTT), flow installation rate, and CPU utilization, etc.

**OFCProbe** [93] is an upgraded version of OFCBenchmark which concentrates on maximizing the flexibility of SDN controllers by emulating a significant amount of OpenFlow switches in a large scale environment. It is re-designed using Java to make it a platform-independent tool and also to overcome the virtualization overhead caused by SDN emulation tool like Mininet [94]. The core competence of this tool is to analyze the impact of the network topology during the evaluation executed by the client component.

**PktBlaster** [95] is a unified test solution that emulates large scale SDN networks including network infrastructure and orchestration layers of SDN controllers. The free version with limited capabilities offers features such as, latency and throughput measurement with different testing profiles, i.e. TCP, UDP, ARP\_Request, and ARP\_Reply. A throughput test determines the rate at which the controller configures the flows in the switches. The latency test gives the exact time in milliseconds which the controller takes to process a flow in the switch. Although the free version is limited to 16 switches and 64 MAC address, it offers additional properties like Flow tables, Group tables, Meter tables, size of the Switch Buffer, and maximum entries per flow table.

**OFNet** [96] is a combined approach to integrate OpenFlow network emulation with performance monitoring and visual debugging of SDN controllers. OFNet can be deployed in a system to generate different types of topologies. The inbuilt traffic generator produces different types of network traffic. It is capable to measure performance characteristics of the controller such as flow generations, flow failures, CPU utilization, flow table entries, average RTT, latency of flow setup, etc.

### B. Benchmarking Capabilities

In this work we use the three of the tools, i.e. CBench, PktBlaster, and OFNet, to evaluate different controllers. It is important to note that none of the tools available can measure all performance statistics. In most of the previous works and the output of tools, the metrics are rather simplified. For example, the throughput of a controller can be interpreted in a number of different ways. Similarly, as shown in Table III, the latency can be determined using different metrics. The



TABLE IV: Comparison of Benchmarking Tools.

Tool	Advantages	Limitations	License	Availability	User Interface
<b>CBench</b>	Faster Analysis Execution Platform Independent Source Code is available	vSwitches limited to 256 Supports only OpenFlow 1.0 Flow Length is Limited Supports only IP-based traffic Lacks User-Interface	Open-Source	Yes	CLI
<b>PktBlaster</b>	1000 Emulated Switches Customized Switch Groups Detailed Statistical Results Accuracy is better than CBench	No Customized Topology No Application-based Traffic Free Edition lacks Deep Analysis	Open-Source Proprietary	Yes	Web UI
<b>OFNet</b>	In-depth Performance Analysis Self-defined Topology Various Traffic Profiles Flow Event Syntax Traffic Generator	Benchmarks Relies on Topology Slower Test Duration	Open-Source	On Request	GUI

TABLE V: Parameters used in evaluation setup.

Tool	Parameter Values	
<b>CBench</b>	Number of Switch	2, 4, 8, 16
	Number of Test Loops	20
	Test Duration	300 sec
	MAC Addresses per Switch (Hosts)	64
	Delay between Test Intervals	2 sec
<b>PktBlaster</b>	Number of Switch	2, 4, 8, 16
	Test Duration	300 sec
	Number of Iterations	5
	Traffic Profile	TCP
	Ports per Switch (Hosts)	64
	Flow Counts per Table	65536 (Default)
<b>OFNet</b>	Packet Length	64 bytes
	Number of Hosts	20
	Number of Switches	7
	Desired Traffic Rate	100 flow/sec
	Flow measured by	Packet-out & Flow-Mod
Total Test Duration	300 sec	

columns on right side of table shows each individual metric which can be directly measured, indirectly measured, or not measurable by a specific tool.

## VI. EVALUATION AND BENCHMARKING OF CONTROLLERS

This section discusses performance of 9 different controllers using previously described benchmarking tools. To the best of our knowledge, no previous work has compared such a large number of controllers, and performed cross comparison using different tools. The controllers evaluated are NOX, POX, Floodlight, ODL, ONOS, Ryu, OpenMUL, Beacon, and Maestro. The reason to select these out of previously discussed 34, is the availability of controller source code or implementation. The 3 benchmarking tools used are CBench, PktBlaster, and OFNet. We use a virtualized environment to emulate the controller and tools in separate virtual machines, running on a 2.10 GHz i7-3612QM processor with 12 GB of DDR3 RAM. Ubuntu 16.04.03 LTS is the base operating system and 1 Gbps link connects the VMs.

It is important to note that all results are plotted as bar graphs. This is done to increase visual understanding of the reader. Overlapping nine different controller outputs in a single plot were not only visually confusing, but also made it difficult to infer any meaningful information.

### A. Evaluation Setup

Table V shows the different parameters for evaluation setup. It is important to note that the programmable parameters in different tools are not identical, hence, we have tried to

best possible extent to make them similar. However, once the parameters are set, all controllers use the same values.

**CBench** tests the performance by sending asynchronous messages. For latency the messages are in series, i.e. it send a packet\_in message to the emulated switch and waits for a response before sending the next one. We execute 20 iterations with varying number of emulated switches to observe the impact of switches on the controller. On the other hand, with same parameters we test the throughput of the running controller. However, the packets are not sent in series, and requests are sent without waiting for a response. One execution, CBench outputs the flow messages a controller can handle per second. The results presented here are an average of number of responses per second from all switches in that execution.

**PktBlaster** utilizes the in-built TCP-based traffic emulation profile that creates an OpenFlow session between the emulated switch and the controller. Due to free edition of tool the number of iterations is limited to 5. The nine controllers are evaluated based on latency (flow installation rate) and throughput (flow processing rate).

**OFNet** uses a custom tree-based topology consisting of 7 emulated switches and 20 virtual hosts. We limit the number of hosts and switches due to limited resources available on emulating machines. Inbuilt traffic generator is used, which initiates and transfers multiple types of traffic, such as DNS, Web, Ping, NFS, Multi-cast, Large-send, FTP and Telnet among hosts in the emulated network much like Mininet Emulation environment. Host 2, 12 and 20 act as DNS, NFS and Multicast server respectively. We analyze metrics such as, Round Trip Time, average flow setup latency, vSwitch CPU utilization, number of flows missed by the controllers, number of flows sent and received. OFNet provides analysis against time, hence the average of 10 iteration is plotted against a 300 seconds simulation.

### B. Latency Performance

1) *CBench*: We observe two different effects on latency using CBench tool. First we observe the latency against number of switches in topology, from 2 to 16. Figure 3a shows that there are two distinct groups, one with high latency, and one with significantly lower. An interesting observation is the Ryu controller which has negligible impact on its latency performance. Similarly, NOX and POX also show minimal change in latency as the switches increase. However, less latency does not translate to out-right winner, as the capabilities of controller itself must also be considered. In this regard, ODL, consistently performs in the middle and offers a number of other feature as listed in Table I.

The second experiment observes the effect of tool's own performance on latency measurement. Here we change the number of iterations while the number of switches is fixed at 16. Interestingly, the pattern in Figure 3b shows most of the controllers to change their latency as the results are averaged-out over a larger set of repetitions. The basic take-away from this is that the setup environments effect on measurements should never be disregarded. It may positively or negatively impact the obtained results with the same parameters.

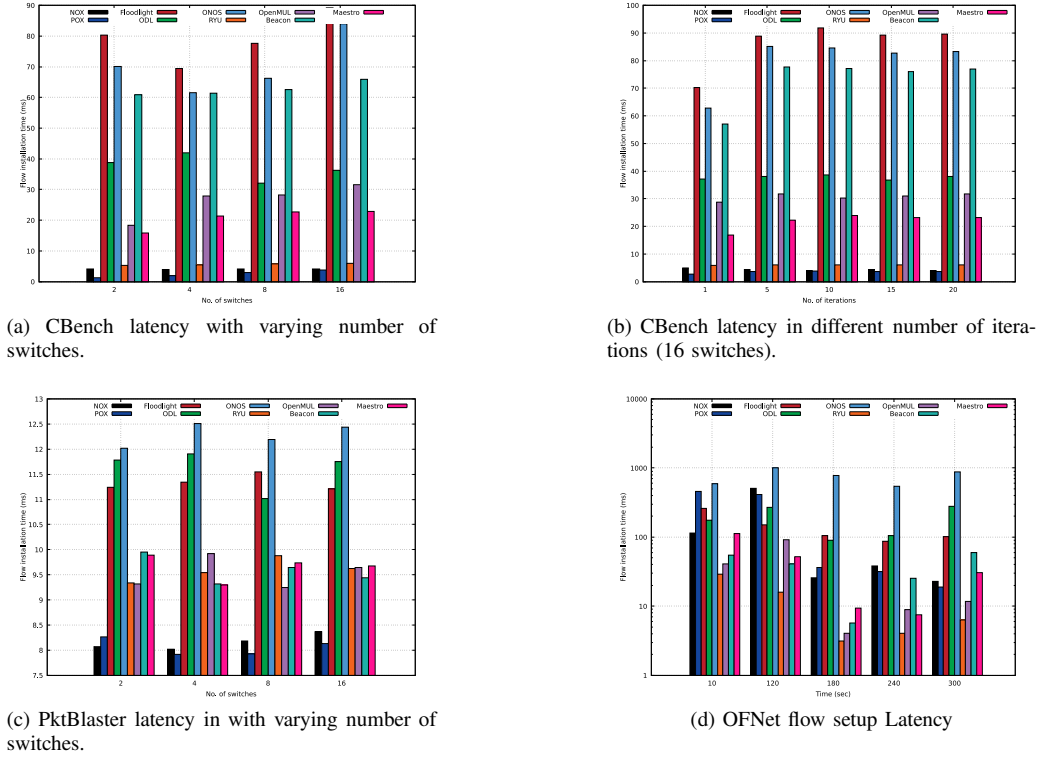


Fig. 3: Latency performance for CBench, PktBlaster, and OFNet.

2) *PktBlaster*: Latency calculation using PktBlaster is also done against increasing number of switches. Figure 3c shows three distinct groups of controllers. NOX and POX show minimum latency, while Floodlight, ODL, and ONOS have the highest latency in this test. Ryu, OpenMUL, Maestro, and Beacon are in the middle. The important factor to note here is that the number of switches does not have any significant impact on the latency calculation. We again emphasize the fact, that the measurement process should reflect the metric being measured. Here latency is more closer to RTT between observing node and controller. On the other hand, flow installation time (path provisioning) would include multiple switches, hence increasing the time.

3) *OFNet*: Unlike CBench and PktBlaster, OFNet has a different evaluation and reporting method, where it simulates the SDN network much like Mininet. The output values are reported against time, instead of a specific value. Figure 3d shows the averaged result of 10 iterations on a time line of 300 seconds. It can be observed that there is no specific pattern over time followed by any given controller. The overall effect that we observe is that less time is required to install flows as the simulation progresses. The dip and rise in latency at around 180 sec mark is due to traffic generation artifact, where some types of traffic are generated later in the simulation, hence requiring more flows.

4) *Cross-Tool Analysis*: One of the contributions of this article is to demonstrate the difference in outcome for same metric under potentially similar network environments. As can be seen from Figure 3 the Y-axis scale varies extensively in all three tools. For CBench the measured latency is in the

orders of tens of milliseconds, where as in PktBlaster the same controllers perform under 10ms. In a total contrast the latency measurements on OFNet are in the order of hundreds of milliseconds. Controllers which performed the best in one simulator, are the worst performs in the other. Although OFNet has a different topological setup, however there is no correlation in the observed results.

### C. Throughput Performance

This metric is measured using CBench and PktBlaster only as shown in Figure 4. OFNet does not provide direct measurement of flow processing, however, indirect measurement can be done through sent and received flow messages, which is discussed in later section.

1) *CBench*: In throughput mode, CBench switches send as many packets as possible at once, and does not wait for a reply. Figure 4a shows the comparison based on increasing number of switches. It is observed that NOX, POX, and RYU remain the lowest performers, while controllers like ODL, Beacon and Maestro have up to 100 responses per milliseconds. Although both OpenMUL and Floodlight performed consistently well around 150 flows/ms, the flow response rate of ONOS is significantly higher around 400 flows/ms to 500 flows/ms.

2) *PktBlaster*: The measurements of throughput shown in Figure 4b present minimal effect from change in number of switches when testing with PktBlaster. The performance of Floodlight, ODL, and ONOS is the best among all the controllers compared, while NOX and POX are at the lower end.

A minor (insignificant) decrease in throughput was observed as the number of switches increased for NOX, POX, and Ryu. However, after running 5 iterations each, the change remains insignificant.

3) *Cross-Tool Analysis*: Similar to earlier analysis, the tools differ in throughput metric also, however the change is not too drastic. All the controllers tend to perform better in PktBlaster evaluations as compared to CBench. Specifically, ODL and Floodlight show significant gain in the performance.

#### D. OFNet Specific Measurements

In this set of experiments, we focus specifically on the performance metrics offered by OFNet.

1) *Average Round Trip Time*: RTT evaluation is an important factor to consider when identifying the location of controller deployment. It identifies the communication delay between the controller and the switch. If the controller and switches are physically far apart, the increased RTT will contribute to increased latency. Similarly, the time complexity of packet processing at controller affects the overall performance. Based on our tree topology, Figure 5a shows that ONOS has high RTT that starts with 100 ms and goes past 1000 ms during the simulation. On the other hand, Ryu & OpenMUL have least RTTs, mostly because of less complex algorithms involved at the controller. However, less complex does not translate to better, rather, they may be attributed to less number of controller capabilities.

2) *CPU Utilization of vSwitch Daemon*: Here we use the OFNet's in-built traffic emulation application to transmit various packets and to identify the CPU usage of the vSwitch process while the vSwitch is interacting with a controller. While running a single-threaded controller like NOX, POX, and RYU, the CPU utilization in Figure 5b of vSwitch daemon remains under 30% to 40%. On the contrary, CPU utilization is remarkably higher at 90% in the case of the multi-threaded controller like ONOS. Besides, the CPU usage remains under 70% rest of the controllers including Floodlight and ODL. One major factor in high throughput performance of ONOS is the multi-threading capabilities. However, they can be limited by the capabilities of the vSwitches.

3) *Missed Flows*: Here we measure the number of flows that the controller misses while the test is ongoing. Typically the traffic generator initiates flow requests to the vSwitches, which in-turn sends requests to the controllers and waits for the response. In this testing environment, vSwitch transmits reactive flows to benchmark the SDN controllers. Figure 6a depicts that, ONOS, ODL and Floodlight miss the least number of flows as opposed to NOX, POX and RYU. This again is attributed to the multi-threading capabilities of the controllers, which allows them to perform comparatively better than the single-threaded ones.


4) *Flow Messages Sent & Received*: This experiment calculates the number of flow messages that have been sent to the controller by vSwitch and the received flow messages from the controller. Although, both CBench and PktBlaster use the term "Packet\_in" to send flows towards controller to evaluate latency and throughput, OFNet instead sends flow

messages continuously at a specific duration to determine the flow acceptance efficiency of the controller. Figure 6b shows that least amount of OF messages have been sent to NOX, POX, RYU, and OpenMul compared to others while a significant amount of messages has been transmitted to ONOS, ODL and Floodlight controllers. Figure 6c depicts that, the flow reception rate is higher from the controllers like NOX, POX, and RYU as these controllers have less computational time. On the contrary, flow reception rate of multi-threaded controllers such as Floodlight, ONOS, and ODL is less than the single-threaded ones, which is due to the distributed nature of these controllers. As the received messages are coming from a specific instance of the controller, hence the plot reflects a lesser value.


## VII. RESEARCH FINDINGS

Based on the qualitative analysis of controllers, properties & capabilities of benchmarking tools, and the evaluation of controllers using them, we have summarized the main findings below.

- Considering latency and throughput, multi-threaded controllers including centralized ones (Floodlight, OpenMul, Beacon, Maestro) and distributed ones (OpenDaylight and ONOS) perform significantly better than centralized and single-threaded controllers like NOX, POX, and Ryu. However, they also require more physical resources in order to perform efficiently.
- Majority of the controllers proposed in literature have no implementation available and the details available are not sufficient for third person to code it. Hence, other than theoretical comparison, it is not possible to evaluate them.
- Placement of controller in physical topology, directly impacts a number of performance parameters. In this regard, we plan to conduct an extensive study with different topological setups (datacenter, WAN, mobile, etc.) to compare distributed controllers.
- Limitations of tools also directly effect the benchmarking. For CBench and PktBlaster we only utilized a specified number of the emulated switches due to available hardware resources and in-built traffic profiles. Therefore, physical resource and modification of compiler (or interpreter) may have some noticeable impact on the collected results.
- We also noticed that some of the available features of tools, such as packet length, vSwitch buffer size, etc. impact the performance of the controller. However it is important to note that the outputs given by any tool also indicate the performance of components used in complete topology. Isolating the performance of controller from the results is not possible.
- Utilization of benchmarking tool like OFNet allows us to define custom topology with a variety of traffic profiles. We observed that single-threaded centralized controller can still perform better in simplified topologies while multi-threaded controllers are more suitable for complex environments.




(a) Cbench throughput with varying number of switches.




(b) PktBlaster throughput with varying number of switches.

Fig. 4: Throughput performance for Cbench and PktBlaster.




(a) Average RTT Measurement.




(b) CPU utilization of vSwitch Daemon.


Fig. 5: RTT and CPU performance for OFNet.



(a) Missing Flows.



(b) Flows Sent to Controller.



(c) Flows Received from Controller.

Fig. 6: Flow measurements for OFNet.

## VIII. CONCLUSION

Benchmarking the performance of a controller is a challenging task. In this work we qualitatively compare 34 controllers, and then perform benchmarking and evaluation in quantitative terms for 9 controllers. During this process, we have also categorized and classified the different metrics which should be used for controller benchmarking. Moreover, we conduct an analysis of tools which can be used in the benchmarking process. Based on the observations, we find that very few controllers comply to OpenFlow 1.3 (or higher version) and provide enough information for actual deployment. Most of the evaluations done previously are based on simple metrics, with specific optimization objectives. Moreover, the tools used vary

significantly in features and capabilities. It is impractical to compare results of one tool with another. Simulation/emulation based evaluation can give only an indication of performance at best, and may significantly differ from actual production environment evaluation.

## REFERENCES

- [1] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski *et al.*, "B4: Experience with a globally-deployed software defined WAN," in *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, 2013, pp. 3–14.
- [2] C.-Y. Hong, S. Mandal, M. Al-Fares, M. Zhu *et al.*, "B4 and After: Managing Hierarchy, Partitioning, and Asymmetry for Availability and Scale in Google's Software-defined WAN," in *Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication*, 2018, pp. 74–87.

- [3] S. Bera, S. Misra, and A. V. Vasilakos, "Software-Defined Networking for Internet of Things: A Survey," *IEEE Internet of Things Journal*, vol. 4, no. 6, pp. 1994–2008, 2017.
- [4] I. T. Haque and N. Abu-Ghazaleh, "Wireless Software Defined Networking: A Survey and Taxonomy," *IEEE Communications Surveys Tutorials*, vol. 18, no. 4, pp. 2713–2737, 2016.
- [5] V. Nguyen, A. Brunstrom, K. Grinnemo, and J. Taheri, "SDN/NFV-Based Mobile Packet Core Network Architectures: A Survey," *IEEE Communications Surveys Tutorials*, vol. 19, no. 3, pp. 1567–1602, 2017.
- [6] L. Zhu, X. Tang, M. Shen, X. Du, and M. Guizani, "Privacy-preserving ddos attack detection using cross-domain traffic in software defined networks," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 3, pp. 628–643, 2018.
- [7] X. Du, M. Guizani, Y. Xiao, and H.-H. Chen, "A routing-driven elliptic curve cryptography based key management scheme for heterogeneous sensor networks," *Trans. Wireless. Comm.*, vol. 8, no. 3, pp. 1223–1229, Mar. 2009.
- [8] Y. Xiao, V. K. Rayi, B. Sun, X. Du, F. Hu, and M. Galloway, "A survey of key management schemes in wireless sensor networks," *Comput. Commun.*, vol. 30, no. 11–12, pp. 2314–2341, Sep. 2007.
- [9] M. G. X. Du, Y. Xiao and H. H. Chen, "An effective key management scheme for heterogeneous sensor networks," *Ad Hoc Networks*, vol. 5, no. 1, pp. 24–34, January 2007.
- [10] Y. Xiao, X. Du, J. Zhang, F. Hu, and S. Guizani, "Internet protocol television (iptv): The killer application for the next-generation internet," *IEEE Communications Magazine*, vol. 45, no. 11, pp. 126–134, November 2007.
- [11] X. Du and H. Chen, "Security in wireless sensor networks," *IEEE Wireless Communications*, vol. 15, no. 4, pp. 60–66, Aug 2008.
- [12] N. Gude, T. Koponen, J. Pettit *et al.*, "Nox: Towards an Operating System for Networks," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 3, p. 105, 2008.
- [13] "POX Controller Manual Current Documentation." [Online]. Available: <https://noxrepo.github.io/pox-doc/html/>
- [14] Big Switch Networks, "Project Floodlight." [Online]. Available: <http://www.projectfloodlight.org/floodlight/>
- [15] "OpenDaylight: A Linux Foundation Collaborative Project." [Online]. Available: <https://www.opendaylight.org/>
- [16] P. Berde, M. Gerola, J. Hart *et al.*, "Onos: Towards an open, distributed sdn os," in *Proceedings of the Third Workshop on Hot Topics in Software Defined Networking*, ser. HotSDN '14. ACM, 2014, pp. 1–6.
- [17] Ryu SDN Framework Community, "Ryu Controller." [Online]. Available: <https://osrg.github.io/ryu/index.html>
- [18] A. Tootoonchian, S. Gorbunov, Y. Ganjali, M. Casado, and R. Sherwood, "On controller performance in software-defined networks." *Hot-ICE*, vol. 12, pp. 1–6, 2012.
- [19] S. A. Shah, J. Faiz, M. Farooq *et al.*, "An Architectural Evaluation of SDN Controllers," in *IEEE International Conference on Communications*, 2013, pp. 3504–3508.
- [20] A. Shalimov, D. Zuiikov, and D. a. Zimarina, "Advanced Study of SDN/OpenFlow Controllers," in *Proceedings of the Central Eastern European Software Engineering Conference*. ACM, 2013, pp. 1–6.
- [21] M. Jarschel, F. Lehrieder, Z. Magyari, and R. Pries, "A Flexible OpenFlow-Controller Benchmark," in *2012 European Workshop on Software Defined Networking*, 2012, pp. 48–53.
- [22] M. Jarschel, C. Metter, T. Zinner, S. Gebert, and P. Tran-Gia, "Ofcprobe: A platform-independent tool for openflow controller analysis," in *IEEE International Conference on Communications and Electronics*, 2014, pp. 182–187.
- [23] Z. K. Khattak, M. Awais, and A. Iqbal, "Performance evaluation of OpenDaylight SDN controller," *Proceedings of the International Conference on Parallel and Distributed Systems*, pp. 671–676, 2014.
- [24] R. Khondoker, A. Zaalouk, R. Marx, and K. Bayarou, "Feature-based comparison and selection of Software Defined Networking (SDN) controllers," in *World Congress on Computer Applications and Information Systems*, 2014, pp. 1–7.
- [25] Y. Zhao, L. Iannone, and M. Riguidel, "On the performance of SDN controllers: A reality check," in *IEEE Conference on Network Function Virtualization and Software Defined Network*, 2015, pp. 79–85.
- [26] P. Isaia and L. Guan, "Performance benchmarking of SDN experimental platforms," in *IEEE NetSoft Conference and Workshops*, 2016, pp. 116–120.
- [27] S. Mallon, V. Gramoli, and G. Jourjon, "Are today's sdn controllers ready for primetime?" in *IEEE Conference on Local Computer Networks*, 2016, pp. 325–332.
- [28] O. Salman, I. H. Elhajj, A. Kayssi, and A. Chehab, "SDN controllers: A comparative study," in *Mediterranean Electrotechnical Conference*, 2016, pp. 1–6.
- [29] D. Suh, S. Jang, S. Han *et al.*, "Toward Highly Available and Scalable Software Defined Networks for Service Providers," *IEEE Communications Magazine*, vol. 55, no. 4, pp. 100–107, 2017.
- [30] I. Z. Bholebawa and U. D. Dalal, "Performance analysis of sdn/openflow controllers: Pox versus floodlight," *Wireless Personal Communications*, vol. 98, no. 2, pp. 1679–1699, 2018.
- [31] A. Nguyen-Ngoc, S. Raffeck, S. Lange *et al.*, "Benchmarking the ONOS Controller with OFCProbe," in *IEEE Seventh International Conference on Communications and Electronics*, 2018, pp. 367–372.
- [32] N. McKeown, T. Anderson, H. Balakrishnan *et al.*, "OpenFlow: Enabling Innovation in Campus Networks," *ACM SIGCOMM Computer Communication Review*, vol. 38, no. 2, p. 69, 2008.
- [33] Y. Rekhter, T. Li, and S. Hares, "A Border Gateway Protocol 4 (BGP-4)." [Online]. Available: <https://tools.ietf.org/html/rfc4271>
- [34] T. V. Lakshman, T. Nandagopal, R. Ramjee, K. Sabnani, and T. Woo, "The SoftRouter Architecture," in *In ACM HOTNETS*, 2004.
- [35] L. Yang, T. A. Anderson, R. Gopal, and R. Dantu, "Forwarding and Control Element Separation (ForCES) Framework," RFC 3746, 2004.
- [36] N. Feamster, H. Balakrishnan, J. Rexford *et al.*, "The Case for Separating Routing from Routers," in *Proceedings of the ACM SIGCOMM Workshop on Future Directions in Network Architecture*, 2004, pp. 5–12.
- [37] A. Farrel, J.-P. Vasseur, and J. Ash, "A Path Computation Element (PCE)-Based Architecture," RFC 4655, Internet Engineering Task Force, 2006.
- [38] J. Van der Merwe, A. Cepleanu, K. D'Souza *et al.*, "Dynamic Connectivity Management with an Intelligent Route Service Control Point," in *Proceedings of the SIGCOMM Workshop on Internet Network Management*, 2006, pp. 29–34.
- [39] A. Greenberg, G. Hjalmytsson, D. A. Maltz *et al.*, "A Clean Slate 4D Approach to Network Control and Management," *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 5, pp. 41–54, 2005.
- [40] M. Casado, T. Garfinkel, A. Akella *et al.*, "SANE: A Protection Architecture for Enterprise Networks," in *USENIX fSecurity Symposium*, vol. 49, 2006, pp. 137–151.
- [41] M. Casado, M. J. Freedman, J. Pettit, J. Luo, N. McKeown, and S. Shenker, "Ethane: Taking Control of the Enterprise," *SIGCOMM Comput. Commun. Rev.*, vol. 37, no. 4, pp. 1–12, 2007.
- [42] D. Erickson, "The beacon openflow controller," *Proceedings of the second ACM SIGCOMM workshop*, pp. 13–18, 2013.
- [43] S. H. Yeganeh and Y. Ganjali, "Beehive: Towards a Simple Abstraction for Scalable Software-Defined Networking," *Proceedings of the ACM Workshop on Hot Topics in Networks - HotNets-XIII*, pp. 1–7, 2014.
- [44] GitHub, "An Open Source SDN Controller for Cloud Computing Data Centers." [Online]. Available: <https://github.com/China863SDN/DCFabric>
- [45] K. Phemius, M. Bouet, and J. Leguay, "DISCO: Distributed multi-domain SDN controllers," *IEEE/IFIP Network Operations and Management Symposium: Management in a Software Defined World*, 2014.
- [46] J. Bailey and S. Stuart, "Faucet: Deploying SDN in the Enterprise," *ACM Queue*, no. October, pp. 1–15, 2016.
- [47] R. Sherwood, G. Gibb, K.-k. Yap *et al.*, "FlowVisor: A Network Virtualization Layer," *Network*, p. 15, 2009.
- [48] A. Tootoonchian and Y. Ganjali, "Hyperflow: a distributed control plane for openflow," *Proceedings of the 2010 internet network management conference on Research on enterprise networking*, pp. 3–3, 2010.
- [49] S. Hassas Yeganeh and Y. Ganjali, "Kandoo: A Framework for Efficient and Scalable Offloading of Control Applications Soheil," *Proceedings of the first workshop on Hot topics in software defined networks*, p. 19, 2012.
- [50] A. Kazarez, "Loom Github Page." [Online]. Available: <https://github.com/FlowForwarding/loom>
- [51] Z. Cai, A. Cox, and E. T. S. Ng, "Maestro: A System for Scalable OpenFlow Control," *Cs.Rice.Edu*, p. 10, 2011. [Online]. Available: <http://www.cs.rice.edu/~eugeneng/papers/TR10-11.pdf>
- [52] A. Voellmy, B. Ford, Y. R. Yang *et al.*, "Scaling Software-Defined Network Controllers on Multicore Servers," *Proceedings of the ACM SIGCOMM Conference on Applications, technologies, architectures, and protocols for computer communication*, pp. 289–290, 2012.
- [53] M. Banikazemi, D. Olshefski, A. Shaikh *et al.*, "Meridian: An SDN platform for cloud network services," *IEEE Communications Magazine*, vol. 51, no. 2, pp. 120–127, 2013.
- [54] GitHub, "MicroFlow: The light-weighted, lightning fast OpenFlow SDN controller." [Online]. Available: <https://github.com/PanZhang/Microflow>

- [55] "NODEFLOW: An openflow controller node style." [Online]. Available: <http://garyberger.net/?p=537>
- [56] T. Koponen, M. Casado, N. Gude, and other, "Onix: A Distributed Control Platform for Large-Scale Production Networks," *USENIX Conference on Operating Systems Design and Implementation*, pp. 1–6, 2010.
- [57] "OpenContrail An open-source network virtualization platform for the cloud." [Online]. Available: <http://www.opencontrail.org/>
- [58] B. Lee, S. H. Park, J. Shin, and S. Yang, "IRIS: The Openflow-based Recursive SDN controller," *International Conference on Advanced Communication Technology*, pp. 1227–1231, 2014.
- [59] "OpenMUL SDN Platform." [Online]. Available: <http://www.openmul.org/openmul-controller.html>
- [60] A. D. Ferguson, A. Guha, C. Liang *et al.*, "Participatory Networking: An API for Application Control of SDNs," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 4, pp. 327–338, 2013.
- [61] S. Li, D. Hu, W. Fang, S. Ma *et al.*, "Protocol Oblivious Forwarding (POF): Software-Defined Networking with Enhanced Programmability," *IEEE Network*, vol. 31, no. 2, pp. 58–66, 2017.
- [62] A. Wang, X. Mei, J. Croft *et al.*, "Ravel: A Database-Defined Network," *Proceedings of the Symposium on SDN Research*, pp. 1–7, 2016.
- [63] S. Shin, Y. Song, T. Lee *et al.*, "Rosemary: A Robust, Secure, and High-Performance Network Operating System Seungwon," *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pp. 78–89, 2014.
- [64] GitHub, "RunOS OpenFlow Controller." [Online]. Available: <https://github.com/ARCCN/runos>
- [65] F. Botelho, A. Bessani, F. M. V. Ramos, and P. Ferreira, "SMaRtLight: A Practical Fault-Tolerant SDN Controller," *Arxiv preprint*, pp. 1–7, 2014. [Online]. Available: <http://arxiv.org/abs/1407.6062>
- [66] B. Trevizan De Oliveira, L. Batista Gabriel, and C. Borges Margi, "TinySDN: Enabling multiple controllers for software-defined wireless sensor networks," *IEEE Latin America Transactions*, vol. 13, no. 11, pp. 3690–3696, 2015.
- [67] Y. Takamiya and N. Karanatsios, "Trema OpenFlow controller framework," 2018. [Online]. Available: <https://github.com/trema/trema>
- [68] M. Monaco, O. Michel, and E. Keller, "Applying Operating System Principles to SDN Controller Design," in *Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks*. ACM, 2013, p. 2.
- [69] F. Dürr, T. Kohler, J. Grunert, and A. Kutzleb, "Zerosdn: A message bus for flexible and light-weight network control distribution in SDN," *CoRR*, vol. abs/1610.04421, 2016. [Online]. Available: <http://arxiv.org/abs/1610.04421>
- [70] A. Tootoonchian, M. Ghobadi, and Y. Ganjali, "Opentm: Traffic matrix estimator for openflow networks," in *Proceedings of the Conference on Passive and Active Measurement*, 2010, pp. 201–210.
- [71] N. L. M. van Adrichem, C. Doerr, and F. A. Kuipers, "Opennetmon: Network monitoring in openflow software-defined networks," in *IEEE Network Operations and Management Symposium*, 2014, pp. 1–8.
- [72] S. R. Chowdhury, M. F. Bari, R. Ahmed, and R. Boutaba, "Payless: A low cost network monitoring framework for software defined networks," in *IEEE Network Operations and Management Symposium*, 2014, pp. 1–9.
- [73] P. H. Isolani, J. A. Wickboldt, C. B. Both, J. Rochol, and L. Z. Granville, "Sdn interactive manager: An openflow-based sdn manager," in *IFIP/IEEE International Symposium on Integrated Network Management*, 2015, pp. 1157–1158.
- [74] W. Kim, J. Li, J. W. K. Hong, and Y. J. Suh, "Ofmon: Openflow monitoring system in onos controllers," in *2016 IEEE NetSoft Conference and Workshops (NetSoft)*, 2016, pp. 397–402.
- [75] N. Handigol, S. Seetharaman, M. Flajslik *et al.*, "Plug-n-serve: Load-balancing web traffic using openflow," *ACM Sigcomm Demo*, vol. 4, no. 5, p. 6, 2009.
- [76] H. Uppal and D. Brandon, "Openflow based load balancing," *CSE561: Networking Project Report*, University of Washington, 2010.
- [77] R. Wang, D. Butnariu, J. Rexford *et al.*, "Openflow-based server load balancing gone wild," *Hot-ICE*, vol. 11, pp. 12–12, 2011.
- [78] C. Liang, R. Kawashima, and H. Matsuo, "Scalable and crash-tolerant load balancing based on switch migration for multiple open flow controllers," in *International Symposium on Computing and Networking*, 2014, pp. 171–177.
- [79] R. D. Corin, M. Gerola, R. Riggio, F. D. Pellegrini, and E. Salvadori, "Vertigo: Network virtualization and beyond," in *European Workshop on Software Defined Networking*, 2012, pp. 24–29.
- [80] L. Xingtao, G. Yantao, W. Wei, Z. Sanyou, and L. Jiliang, "Network virtualization by using software-defined networking controller based docker," in *IEEE Information Technology, Networking, Electronic and Automation Control Conference*, 2016, pp. 1112–1115.
- [81] "Docker overview." [Online]. Available: <https://docs.docker.com/engine/docker-overview/>
- [82] D. Drutskey, E. Keller, and J. Rexford, "Scalable network virtualization in software-defined networks," *IEEE Internet Computing*, vol. 17, no. 2, pp. 20–27, March 2013.
- [83] A. Blenk, A. Basta, and W. Kellerer, "Hyperflex: An sdn virtualization architecture with flexible hypervisor function allocation," in *IFIP/IEEE International Symposium on Integrated Network Management*, 2015, pp. 397–405.
- [84] "OpenStack:Open source software for creating private and public clouds." [Online]. Available: <https://www.openstack.org/software/>
- [85] A. Mayoral, R. Vilalta, R. Muoz *et al.*, "Sdn orchestration architectures and their integration with cloud computing applications," *Optical Switching and Networking*, vol. 26, pp. 2 – 13, 2017.
- [86] "OpenStack Havana Release." [Online]. Available: <https://www.openstack.org/software/havana/>
- [87] T. Hinrichs, N. Gude, M. Casado, J. Mitchell, and S. Shenker, "Expressing and enforcing flow-based network security policies," *University of Chicago, Tech. Rep.*, vol. 9, 2008.
- [88] M. F. Bari, S. R. Chowdhury, R. Ahmed, and R. Boutaba, "Policycop: An autonomic qos policy enforcement framework for software defined networks," in *IEEE SDN for Future Networks and Services*, 2013, pp. 1–7.
- [89] V. Varadharajan, K. K. Karmakar, and U. Tupakula, "Securing communication in multiple autonomous system domains with software defined networking," in *IFIP/IEEE Symposium on Integrated Network and Service Management*, 2017, pp. 195–203.
- [90] B. Vengainathan, A. Basil, M. Tassinari *et al.*, "Benchmarking Methodology for Software-Defined Networking (SDN) Controller Performance," RFC 8456, 2018.
- [91] R. Sherwood and K.-K. Yap, "Cbench Controller Benchmark," 2011. [Online]. Available: <https://github.com/andi-bigswitch/oflops/tree/master/cbench>
- [92] GitHub, "WCBench:CBench, Wrapped in stuff that makes it Useful." [Online]. Available: <https://github.com/dfarrell07/wcbench>
- [93] "OFCProbe: A Platform Independent Tool for OpenFlow Controller Analysis." [Online]. Available: <https://www3.informatik.uni-wuerzburg.de/research/ngn/ofcprobe/ofcprobe.shtml>
- [94] M. Team, "Mininet: An Instant Virtual Network on your Laptop (or other PC)." [Online]. Available: <https://mininet.org/>
- [95] "PktBlaster SDN Datasheet," Veryx Technologies, Tech. Rep., 2016. [Online]. Available: <http://www.veryxtech.com/wp-content/uploads/2015/10/Datasheet-PktBlaster-SDN-Cor>
- [96] G. H. Shankar, "OFNet Quick User Guide." [Online]. Available: <http://sdninsights.org/>