On the Design and Development of Emulation Platforms for NFV-based Infrastructures

Vinicius Fulber-Garcia*, Giovanni Venâncio de Souza, Elias Procópio Duarte Junior

Department of Informatics,
Federal University of Paraná,
Paraná, PR, BR
E-mail: {vfgarcia, gvsouza, elias}@inf.ufpr.br

Thales N. Tavares, Leonardo da C. Marcuzzo, Carlos R. P. dos Santos

Department of Applied Computing,
Federal University of Santa Maria,
Santa Maria, RS, BR
E-mail: {tntavares, lmarcuzzo, csantos}@inf.ufsm.br

Muriel Figueredo Franco, Lucas Bondan, Lisandro Zambenedetti Granville, Alberto Egon Schaeffer-Filho

Institute of Informatics,
Federal University of Rio Grande do Sul,
Porto Alegre, RS, BR
E-mail: {mffranco, lbondan, granville, alberto}@inf.ufrgs.br

Filip De Turck

INTEC,
Ghent University,
Ghent, PR, VB
E-mail: {filip.deturck}@ugent.be

*Corresponding author

Abstract: Network Functions Virtualization (NFV) presents several advantages over traditional network architectures, such as flexibility, security, and reduced CAPEX/OPEX. In traditional middleboxes, network functions are usually executed on specialized hardware (e.g., firewall, DPI). Virtual Network Functions (VNFs) on the other hand, are executed on commodity hardware, employing Software Defined Networking (SDN) technologies (e.g., OpenFlow, P4). Although platforms for prototyping NFV environments have emerged in recent years, they still present limitations that hinder the evaluation of NFV scenarios such as fog computing and heterogeneous networks. In this work, we present NIEP: a platform for designing and testing NFV-based infrastructures and VNFs. NIEP consists of a network emulator and a platform for Click-based VNFs development. NIEP provides a complete NFV emulation environment, allowing network operators to test their solutions in a controlled scenario prior to deployment in production networks.

Keywords: NFV; VNF; Emulation; Platform; Infrastructure; Click; Mininet; Network

Reference to this paper should be made as follows: Garcia, V.F., de Souza, G.V., Duarte Junior, E.P., Tavares, T.N., Marcuzzo, L.D.C., dos Santos, C.R.P., Franco, M.F., Bondan, L., Granville, L.Z., Schaeffer-Filho, A.E. and De Turck, F. (2020) On the design and development of emulation platforms for NFV-based infrastructures, Int. J. Grid and Utility Computing, Vol. 11, No. 2, pp.230242.

Biographical notes:

Vinicius Fulber-Garcia is a Ph.D. student in Computer Science at the Department of Informatics of the Federal University of Paran (UFPR - Brazil) under the supervision of Prof. Dr. Elias Procpio Duarte Jnior. He holds a Computer Science degree from Federal University of Santa Maria (UFSM - Brazil) and a Master degree in Computer Science from UFSM Post-Graduate Program in Computer Science. His research interests include, but not limited to, network functions virtualization and information theory.

Giovanni Venncio de Souza is a Ph.D. student in Computer Science at the Department of Informatics of the Federal University of Paran (UFPRBrazil) under the supervision of Prof. Dr. Elias Procpio Duarte Jnior. Giovanni holds an MSc (2017) in Computer Science and a Computer Science degree (2016) at the same institution. His research interests include Network Function Virtualization and Fault-Tolerant Distributed Systems.

Elias Procpio Duarte Junior is a Full Professor at Federal University of Parana, Curitiba, Brazil, where he is the leader of the Computer Networks and Distributed Systems Lab (LaRSis). His research interests include Computer Networks and Distributed Systems, their Dependability, Management, and Algorithms. He has published more than 200 peer-reviewer papers and has supervised more than 130 students both on the graduate and undergraduate levels. Prof. Duarte is currently Associate Editor of the IEEE Transactions on Dependable and Secure Computing, and has served as chair of more than 20 conferences and workshops in his fields of interest. He received a Ph.D. degree in Computer Science from Tokyo Institute of Technology, Japan, 1997, M.Sc. degree in Telecommunications from the Polytechnical University of Madrid, Spain, 1991, and both BSc and MSc degrees in Computer Science from Federal University of Minas Gerais, Brazil, 1987 and 1991, respectively. He chaired the Special Interest Group on Fault Tolerant Computing of the Brazilian Computing Society (2005-2007); the Graduate Program in Computer Science of UFPR (2006-2008); and the Brazilian National Laboratory on Computer Networks (2012-2016). He is a member of the Brazilian Computing Society and a Senior Member of the IEEE.

Thales N. Tavares is a graduate of the course on Computer Network Technology at the Federal University of Santa Maria (2016) in Brazil. He is currently a substitute lecturer at the polytechnic school of the same university. Has knowledge in the area of Computing, with emphasis on Computer Networks. Research interests in network management, software networks and virtualization of network functions.

Leonardo da C. Marcuzzo holds a degree in Computer Science from Federal University of Santa Maria (UFSM) and currently is a M.Sc. Candidate in Computer Science at the same institution. His research interests include network functions virtualization and operating systems.

Carlos R. P. dos Santos is Adjunct Professor of Computer Science at the Department of Applied Computing of the Federal University of Santa Maria (UFSM), Brazil. He holds Ph.D. (2013) and M.Sc. (2008) degrees in Computer Science, both received from the Federal University of Rio Grande do Sul (UFRGS), where he was also Postdoctoral Research Fellow from October 2013 to September 2014. From May 2010 to April 2011 he was a visiting researcher at the IBM T.J. Watson Research Center - Hawthorne, where he developed projects on IT Service Management and Security Management. His current research interests focus on design and management of Future Networks and Technologies, including aspects such as network virtualization, quality of service management, network programmability, and security management.

Muriel Figueredo Franco is pursuing his Ph.D. under the supervision of Prof. Dr. Burkhard Stiller at the University of Zurich (UZH). He is also a Research Assistant at the Communication Systems Group (CSG). Muriel holds an M.Sc (2017) in Computer Science from the Federal University of the Rio Grande do Sul (UFRGS) under the supervision of Prof. Dr. Lisandro Granville and obtained a B.Sc (2014) in Computer Science from the Federal University of Pelotas (UFPEL). His research topics include Network Functions Virtualization, Information Visualization, and Blockchain.

Lucas Bondan is a Ph.D. student in Computer Science at the Institute of Informatics of the Federal University of Rio Grande do Sul (UFRGS) in Brazil and an R&D Project Manager at the Brazilian National Research and Educational Network (RNP). From 2016 to 2018 he was a Ph.D. student fellow at the Department of Information Technology of Ghent University in Belgium, working with security-related areas of Network Functions Virtualization (NFV). He has a Computer Engineering degree from Pontifcia Universidade Catlica do Rio Grande do Sul and a Master Degree in Computer Science from UFRGS. His research interests include network functions virtualization, network management and orchestration, service function chaining, cognitive networks, and wireless communication systems.

Lisandro Zambenedetti Granville is Full Professor of Computer Science at the Institute of Informatics of the Federal University of Rio Grande do Sul (UFRGS), Brazil. He holds Ph.D. (2001) and M.Sc. (1998) degrees in Computer Science, both received from UFRGS. From September 2007 to August 2008 he was a visiting researcher at the University of Twente, The Netherlands, with the Design and Analysis of Communication Systems group. He is a member of the Computer Networks Group, where he develops research projects on network and service management. As a Full Professor, he is also involved with supervision and education activities on undergraduate and graduate courses in both Computer Science and Computer Engineering.

Alberto Egon Schaeffer-Filho holds a Ph.D. in Computer Science (Imperial College London, 2009) and is Associate Professor at Federal University of Rio Grande do Sul (UFRGS), Brazil. From 2009 to 2012 he worked as a research associate at Lancaster University, UK. Dr. Schaeffer-Filho is a CNPq-Brazil Research Fellow and his areas of expertise are network/service management, network virtualization and software-defined networks, policy-based management, and security and resilience of networks. He has authored over 60 papers in leading peer-reviewed journals and conferences related to these topics, and also serves as TPC member for important conferences in these areas, including: IFIP/IEEE IM (2019), NetSoft (2019), CNSM (2018), and IEEE/IFIP NOMS (2018). He is the general chair for SBRC 2019, co-chair for IEEE ICC 2018 CQRM Symposium, and demo co-chair for IFIP/IEEE IM 2017.

Filip De Turck is professor in the department of Information Technology (Intec) of Ghent University with expertise in network software and research interests in adaptive large-scale data processing and software systems for healthcare, anomaly detection, and resilience of ICT infrastructures and services. In this research area, he is involved in several research projects with industry and academia, serves as Chair of the IEEE Technical Committee on Network Operations and Management (CNOM), chair of the Future Internet Cluster of the European Commission, and is on the TPC of many network and service management conferences and workshops and serves in the editorial board of several network and service management journals. Together with a team of PhD students and postdoctoral researchers, novel techniques and algorithms are designed, and validated by means of large scale evaluation studies, together with partners from industry and academia.

1 Introduction

Network Functions Virtualization (NFV) is a novel networking paradigm that fosters innovation and supports the creation of disruptive network services [MSG+16]. In NFV, the network functions are decoupled from the associated hardware and executed in commodity servers (*i.e.*, commercial off-the-shelf servers) by using virtualization technologies. This shift provides significant advancements in how the networks are designed, maintained, and managed while improving the flexibility, scalability, and cost-benefit of networked environments [LLF+15].

All those advantages have brought NFV to the attention of the industry, academia, and standardization bodies. Several efforts have been conducted for the development of new architectures, systems, and applications for NFV [GMUJ16]. Despite a large number of results that have already appeared in this field, several challenges are still open. One of those challenges is to develop a *de facto* approach for predicting the impact of deploying novel Virtualized Network Functions (VNFs) in production environments. In this context, VNF emulation is a promising method that can support the design and evaluation of NFV-based scenarios.

Emulation has proved to be an effective method to evaluate network-based environments, systems and applications [ISH10] [SVZ+14]. In the same way, providing comprehensive emulation tools that support the specific NFV elements (*e.g.,* Virtualized Infrastructure Manager (VIM) and VNF Manager (VNFM)) is of paramount importance for network operators, researchers, and developers. However, despite its inherent benefits, solutions for NFV emulation are still scarce, limited (*e.g.,* due to low portability or lack of support for heterogeneous environments), usually they are not intuitive, and involve a steep learning curve before they can be fully adopted.

In this paper, we present the NFV Infrastructure Emulation Platform (NIEP) [1], a novel platform based on Click-on-OSv [dCMGC+17] and Mininet [LHM10] that allows VNF evaluation by the emulation of diverse NFV scenarios. NIEP allows operators to rapidly create heterogeneous NFV emulated scenarios. These scenarios are portable because of the full virtualization strategy adopted by NIEP. We also show the feasibility of NIEP in a case study considering a Fog computing and Virtual Customer Premises Equipment (vCPE) scenario. We expect that NIEP will effectively assist network operators in the offline analysis of the functionality and performance of VNF deployments. Pre-tested configurations can be evaluated and optimal configurations may be established before actual VNFs are deployed in the network infrastructure.

The rest of this paper is organized as follows. In Section II, the background and related work are reviewed. We then present the simulation/emulation requirements and the proposed NIEP architecture in Section III. In Section IV, we discuss the data model employed to specify the network topologies. In Section V, we describe a case study to demonstrate the feasibility of the platform. Finally, the conclusions follow in Section VI, along with a discussion of future work directions.

2 Background and Related Work

In this section, we present an overview of NFV and network virtualization technologies. After reading this section, it should be clear that NFV brings multiple advantages in comparison with traditional network architectures that are based on middleboxes often deployed on specialized hardware. However, it should be also clear that there are challenges for the successful deployment of NFV-based solutions in production networks. This section also presents issues related to NFV prototyping and evaluation are discussed, highlighting the pros and cons of existing frameworks.

2.1 Network Functions Virtualization (NFV)

NFV technology was first proposed and standardized by the European Telecommunications Standards Institute (ETSI) as a paradigm that decouples network functions from dedicated hardware allowing their implementation using virtualization technology that can be executed on Commercial Off-The-Shelf (COTS) hardware. The fact that NFV does not require specialized hardware and is deployed as a virtual infrastructure enables the development and management of network function in an easy, cost-effective, and flexible way [CW+12]. Furthermore, NFV allows the fast creation of new network functions that can be combined to provide complex network services. NFV together with other technologies based on virtualization has solved the Network Ossification phenomenon [Han06].

The multiple advantages of NFV technology include: *(i)* NFV is cheap, in particular in terms of capital/operational expenditures (CAPEX/OPEX) as general purpose hardware can be used; *(ii)* NFV is fast to deploy, configure and update; *(iii)* NFV is flexible, as virtual functions can dynamically migrate and by using elasticity technology they can be scaled up and down according to the demand; and *(iv)* NFV opens up the market, allowing new players to develop for the computer networks market. NFV is often employed with other technologies, such as Software-Defined Networking (SDN) [HGJL15], which allows the substrate network to be more easily customized to fulfill the specific needs of customers.

Individual VNFs may be combined to execute complex network services. Service Function Chains (SFC) [HP15] are composed of multiple and independent VNFs that can be executed on different virtualization environments. The IETF envisions even multi-domain SFCs, which should employ a hierarchy of orchestrators [BCV+18]. Figure 1 shows an SFC example with three servers running each a hypervisor and set of VNFs which are interconnected forming multiple SFCs.

Despite the multiple advantages, NFV increases the complexity and it is undeniable that before it is deployed major changes are required to the existing network infrastructures. Therefore, new technologies and tools for VNF evaluation are needed. Although tools such as *tcpdump*, *ping*, and *traceroute* can still be used to identify problems in virtualized networks, in NFV, while those tools can still be useful, new network components must be monitored and evaluated to identify problems such as bottlenecks, failures, misconfiguration, or implementation bugs.

Mininet [Tea12] is a simple yet powerful tool that allows the evaluation of Software Defined Network technology. Mininet can be used with an external SDN controller for running experiments. However, Mininet does not offer support for experimentation with VNFs. As a consequence, new tools have been proposed that integrate Mininet with other software components that can be used to deploy and manage VNFs. In the next subsection, we present and discuss some of those efforts that have been proposed for NFV experimentation.

2.2 NFV Experimentation Frameworks

A number of tools and frameworks have been recently proposed for the experimental evaluation of NFV technology. EsCAPE, MeDICINE, SONATA, and Maxinet are among

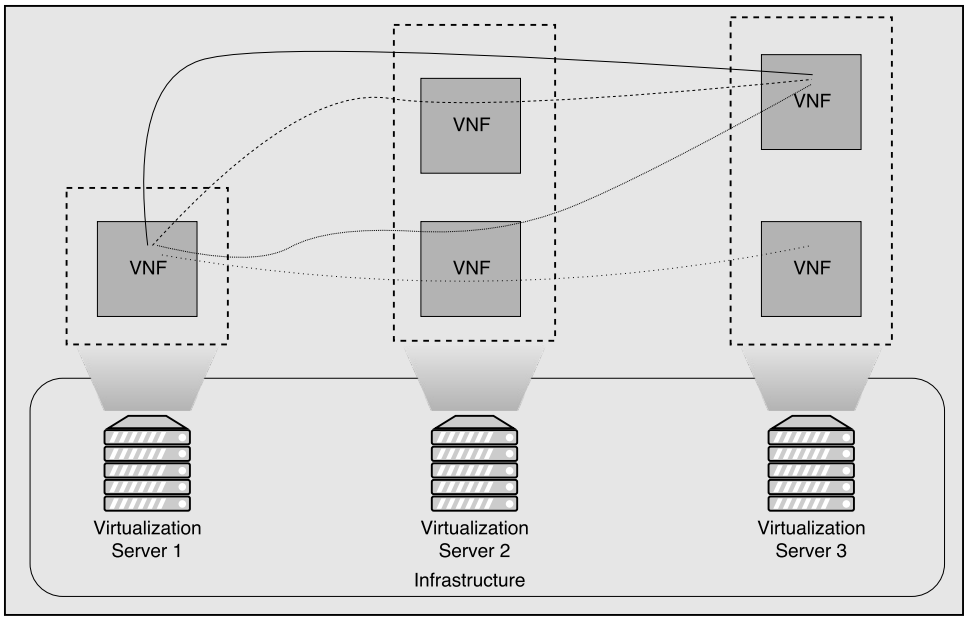

Figure 1 A SFC example

some of the most important of those tools and frameworks. These four platforms are described next.

- **EsCAPE**[SCS+15]: Extensible Service Chain Prototyping Environment (EsCAPE) is a prototyping framework developed in the context of the UNIFY architecture, consisting of three abstraction layers: Service Layer, Orchestrator Layer, and Infrastructure Layer. EsCAPE provides a common platform that enables users to prototype and orchestrate SFCs whose VNFs are deployed as containers running Click[MKJK99]. EsCAPE also features a built-in VNF catalog with basic virtual functions. EsCAPE's network infrastructure is based on Mininet with OpenVSwitches[PPK+15] connected to an external SDN controller (POX) responsible for steering traffic between VNFs. ExCAPE also supports the development and test of orchestration components, extending Mininet to work with NETCONF. The focus of EsCAPE is thus on the creation and management of SFCs, although it can be used to prototype and evaluate other technologies as well.

- **MeDICINE**[PKvR16]: the Multi Datacenter service Chain Emulator (MeDICINE) is an NFV prototyping platform that was designed to emulate multi-PoP environments in which virtual functions are executed on containers. MeDICINE is based on ContainerNET[2], which extends the Mininet framework to support container-based VNFs. Links between complex multi-PoP environments are established using the Mininet API, allowing the specification of multiple requirements such as delay, bandwidth, and the packet loss rate. Docker[3] is used in MeDICINE to deploy VNFs on these PoPs. MeDICINE also provides end-points for each PoP, enabling the interconnection of the elements also to other elements of the ETSI architecture.

- **SONATA**[KDP+16]: SONATA is a tool for NFV composition, testing, and orchestration. It contains an emulation platform based on ContainerNet [PKvR16] which allows developers to prototype network services in end-to-end multi-PoP scenarios. The platform also provides APIs for integration with other components and systems based on the ETSI specifications.

- **Maxinet**[WDS+14]: Maxinet is an extension of Mininet that can be executed in a distributed fashion, and in this way supports the emulation of networks. Maxinet works as an abstraction layer connecting multiple Mininet instances running on distinct hosts connected on a network.

EsCAPE, MeDICINE and SONATA use containers for deploying and executing VNFs. Although container technology should be enough for most NFV use cases[NFV13], container-based virtualization presents some issues for specific NFV scenarios. For example, in comparison to hypervisor-based virtualization, containers do not provide multi-platform compatibility and their life-cycle management is certainly more expensive[MKK15]. Moreover, as opposed to Virtual Machines (VMs), containers increase the vulnerability in terms of security threats [MBD16], since each operating system image has its own set of vulnerabilities and share the same kernel. In scenarios with heterogeneous networks, multiple hosts with different operating systems form the infrastructure substrate, such as vCPE, virtual Evolved Packet Core (vEPC) and Fog Computing. Any given VNF can be deployed and migrated anywhere in the infrastructure.

As for MaxiNet, although it also supports virtual nodes in the same way that Mininet does, not all virtualization technology is available. The main purpose is to run Mininet instances in a distributed way.

In the next section, we introduce NIEP, a framework that integrates a minimal VNF platform (Click-on-OSv) with Mininet, allowing diverse NFV scenarios to be prototyped and evaluated. NIEP fills a gap in terms of the lack of experimental frameworks for evaluating heterogeneous NFV scenarios, as the focus of existing platforms is on SFCs and multi-pop environments. We believe our solution is the first to provide a prototyping framework for the emulation of NFV technology in a variety of scenarios.

3 NIEP: NFV Infrastructure Emulation Platform

In this section, we describe NIEP by first, in Subsection 3.1 discussing a set of requirements identified that must be satisfied by the proposed platform. Next, in Subsection 3.2 the NIEP architecture is presented, with a description of all modules of which it is composed. Next, in Subsection 3.3, the interactions between the modules are described.

3.1 Emulation Requirements

Emulation plays an important role in the design, development, and analysis of VNFs, especially for innovative functions and services. The emulation of a system should represent the system as accurately as possible, allowing the execution of real live tests on the emulation environment [CS03]. The use of these environments has increased significantly in recent years, as they are so convenient for the evaluation of large-scale systems, allowing deep analysis of the system under realistic conditions before it is actually deployed.

An emulation platform for NFV technology should satisfy a set of fundamental requirements. We identified the following requirements raised by Varga and Horing [VH08], Baumgat, Heep and Krause [BHK07], and Schaeffer-Filho *et al.* [SFMH+13] as the basis on which a novel platform for the emulation of NFV-based infrastructures should be designed.

- **Scalability**: the platform must be able to involve a large number of nodes when emulating the NFV-based system;
- **Flexibility**: it should be straightforward to define and update the emulation process, and the user should have a choice of network topologies to specify and use. The building blocks (*e.g.,* hosts, switches, VNFs) with which the system is specified must be generic enough to be reused in a range of scenario definitions;
- **Remodeling**: the definition of evaluation scenarios must be simple, dynamic, allowing fast prototyping; the network topology must be easily modified as needed;
- **Software Execution**: the building blocks provided must reflect those of the corresponding actual system in production, thus providing reliable experimental results.

3.2 NIEP Architecture

NIEP is based on the integration of existing tools for VNF design (Click-on-OSv), VM management (KVM hypervisor), and network emulation (Mininet), plus a core element, which is the orchestration module. The architecture is shown in Figure 2.

We start the description of the NIEP modules with Click-on-OSv [dCMGC+17], an NFV system based on the single-process operating system OSv. Click-on-OSv leverages the Click Modular Router [KMC+00] to create and execute virtual functions and provides a Representational State Transfer (REST) interface for controlling the underlying operations (*e.g.,*., monitoring and lifecycle management). Click-on-OSv itself is a complete virtual machine, it simplifies the control and provisioning processes due to its independence from the host operating system. Moreover, it is possible to remotely create VMs on a set of heterogeneous hosts that run VNF functions in a distributed way.

NIEP is based on a KVM hypervisor, which is a virtual VM manager that implements full virtualization, to support the execution of multiple VMs running images of different operating systems. The Virsh tool[4] is used by the NIEP orchestrator to manage the KVM virtual machines. It is a Command Line Interface (CLI) that enables VM control with system calls. We highlight that KVM provides better performance for Click-on-OSv due to VirtIO[5]. These virtualization optimizations make packet processing by Click running on OSv faster than other hypervisors (*e.g.* VirtualBox, Xen).

Mininet [Tea12], as mentioned above, is a widely used network emulator that relies on process-level virtualization. This lightweight virtualization strategy is used to emulate guest machines as isolated processes, with the proper share of memory, CPU and network resources, enabling the simulation of large-scale network environments. In NIEP, Mininet hosts are used to representing servers and clients, OpenFlow switches and controllers.

Network topologies in NIEP are specified with JSON (JavaScript Object Notation), which also simplifies the infrastructure deployment process when compared to Mininet. Thus, users can configure a Mininet topology with hosts, switches, and controllers also defining other useful information such as resource allocation for VNFs and their interconnections forming SFCs.

The NIEP-Orchestrator provides the user interface. The topology is entered as input, and all required actions are executed to instantiate system. The NIEP-Orchestrator consists of four elements, described next.

- **VNF Repository**: this module is responsible for storing VNFs, which are implemented as Click scripts. As VNFs can be deployed in a distributed fashion across multiple hosts, the repository must be universally accessible. Therefore it works as a marketplace where users can share, publish and obtain VNFs;
- **Virtualized Elements Manager (VEM)**: this module both controls the execution of VNFs and provides communication interfaces (*e.g.,*., network bridges). The VEM is composed of two functional blocks, the Network Functions blocks, which directly controls Click-on-OSv instances using a REST interface, and the Infrastructure, that controls the KVM hypervisor execution using the Virsh CLI;
- **Topology Manager**: this module allows the creation and initialization of the Mininet topology. It creates all the specified elements (*e.g.,*., hosts, switches, controllers) through the Mininet API which later run user operations;
- **Interpreter**: this component is responsible for validating the topology specifications and handling user requests (*e.g.,*., specifying a new network topology or obtaining monitoring data). The input consists of NIEP topology specification, and the output consists of request results.

3.3 Module Interactions

The modules of the NIEP architecture presented in the previous subsection are integrated by the Orchestrator, which as the name implies orchestrates the VNF emulation on the specified Mininet network topology. Initially, the Orchestrator receives the topology specification, forwards the topology to the Interpreter module which does the required validation, checking for mandatory elements and evaluating the configuration correctness. The Interpreter then organizes

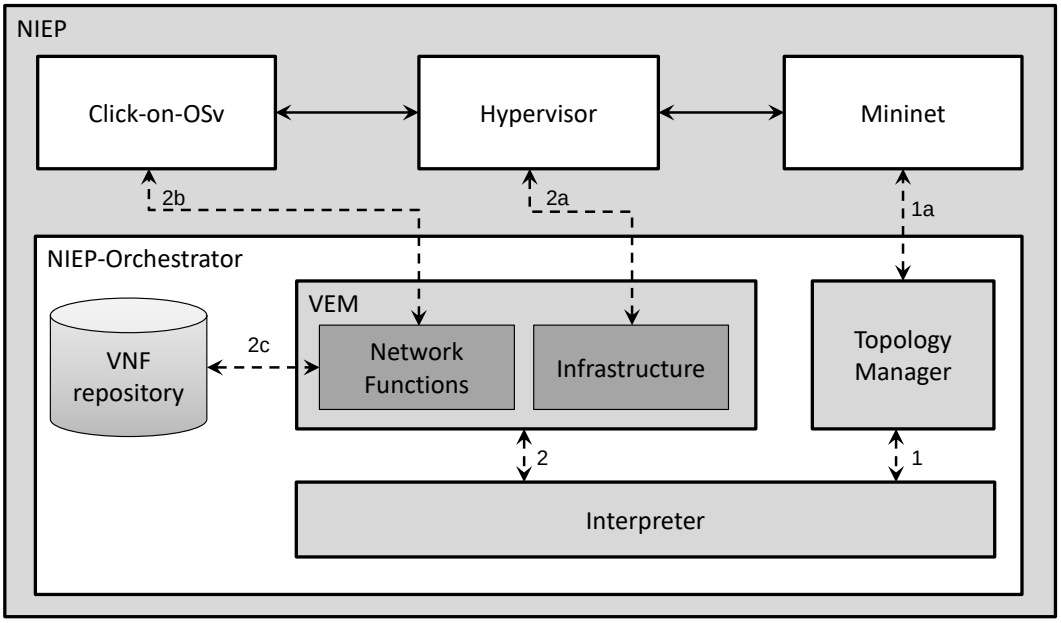

Figure 2 NIEP: The Architecture

the information in two sets: one with information on the Mininet network topology (*e.g.,* hosts, switches, controllers) and the other set with information related to the VNFs to be executed (*e.g.,* memory, CPU, interfaces, plus the Click network function itself).

The first information set of computed by the Interpreter is sent to the Topology Manager – labeled with (1) in Fig. 2, that processes the data of the Mininet environment. The Topology Manager, after processing the information to create the requested topology, triggers the Mininet emulator (1a). The VEM element receives the second set of information from by the Interpreter (2). It checks actions to be executed, which are forwarded to the Network Functions and Infrastructure blocks. The Infrastructure block (2a) executes first, creating virtual machines using the Hypervisor and the communication links to the network topology in Mininet using a bridge interface.

At the end of this process, the user can make requests to the Functions Virtualization block (2b) to start Click-on-Osv by fetching the user-defined Click function from the VNF Repository (2c). This process can also be used to deploy a new Click function on the same system instance, by re-uploading and restarting the Click-on-OSv service without having to restart the VM or the entire topology. Figure 3 shows the high-level communication steps to run a NIEP topology.

In synthesis, NIEP is a platform that allows rapid prototyping and evaluation of large-scale NFV scenarios. In this way, it can be a valuable tool for network operators, by allowing the assessment of the functionality and performance of individual VNFs as well as SFCs before their actual deployment in production networks.

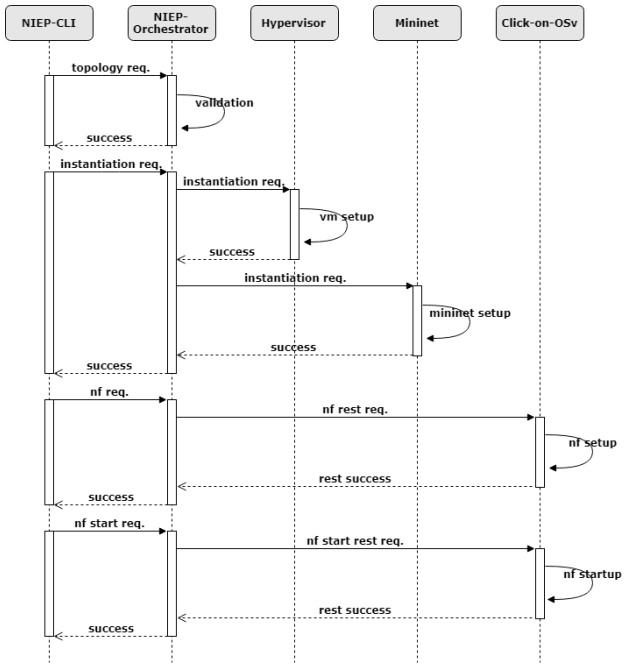

Figure 3 NIEP Topology Instantiation

4 NIEP Data Model

In this section, we present the NIEP data model for defining network topologies. The data model includes definitions of VNFs and SFCs and the network topology, including hosts, switches, and controllers. The elements are described in separate JSON files and are described next.

A VNF is described as a JSON object with five properties, as shown in Figure 4. The unique ID is the key to identify a VNF instance and is used along the execution of the emulation execution by the orchestrator to access the required VNFs as it executes tasks including monitoring, deployment, and VNF lifecycle operations. The other properties are used for VM configuration: memory requirements (MEMORY), number of virtual CPUs (VCPU), and network interfaces (MANAGEMENT_MAC and INTERFACES). The MANAGEMENT_MAC corresponds to a dedicated interface that is employed for the sole purpose of sending and receiving data from the NIEP-Orchestration. The INTERFACES property is used to connect the emulated VNFs and hosts and contains the MAC address and the

```
{
  "type": "object",
  "properties": {
    "ID": {"type": "string"},
    "MEMORY": {"type": "integer"},
    "VCPU": {"type": "integer"},
    "MANAGEMENT_MAC": {"type": "string"},
    "INTERFACES": {"type": "array",
      "items": {
        "type": "object",
        "properties": {
          "MAC": {"type": "string"},
          "ID": {"type": "string"}
        }
      }
    }
  }
}
```

Figure 4 VNF Simplified JSON Schema

virtual connection (*i.e.,* network bridge) ID from which data is received.

SFCs are also specified with a JSON file, Figure 5, composed of five attributes that represent the Service Function Chain. The attributes are as follows. The ID is unique and is used to identify the SFC as a whole, thus making possible to monitor and configure the lifecycle of all the VNFs composing the service. The VNFS attribute represents the set of VNFs that compose the SFC. The VNFS contains the identifier of the VNF in the context of the SFC, as well as a path for the VNF JSON file (created using the schema presented in Figure 4). The VNFs are connected each with a single Incoming Point (IP) and one or more Outgoing Points (OP). The boundary nodes (IP/OP) represent the first (IP) and last (OP) point of a service chain. Both IP and OP are represented by an identifier and a virtual connection is used for the connection of the elements of the SFC.

In the specification of a VNF in a SFC Descriptor, the INTERFACE attribute is omitted and the CONNECTIONS attribute is employed instead. The CONNECTIONS attribute consists of four elements: Input Logical Link (ILL), Output Logical Link (OLL), and the associated MAC addresses (when necessary). The network traffic is delivered to a VNF from an ILL and, after being processed, the traffic is forwarded to the next VNF or to an OP (through an OLL connection). The OLL and ILL elements are specified either by an existing VNF ID or boundary node ID. In the case of an existing VNF, one of its interfaces is employed (which indicated in the corresponding MAC field). In the case of boundary nodes, no MAC is defined because the sender and receiver hosts are outside NIEP.

The complete topology is represented with a third schema, shown in Figure 6. The five attributes of this description carry information about VNFs and SFCs, plus the Mininet emulated network infrastructure. The NIEP topology is identified with a unique ID. A NIEP instance is responsible for the execution of a topology, thus after the topology is deployed the ID also represents the NIEP process itself.

The virtual functions and function chains of a given NIEP topology are defined by the VNFS and SFCS properties. These attributes specify the location of the corresponding description files, created according to the schemas presented above. Note that the VNFs used in an SFC specified with the SFCS attribute should no be explicit in the VNFS attribute. Whenever a duplicated request is required, no ID should be replicated, (*i.e.*, the internal SFC ID must be different from the VNF ID).

The Mininet network is described within a JSON object within the MININET attribute. This "sub-object" contains four attributes, each specifying an operational element of the emulation. The HOST attribute is an array that keeps virtual host IDs; the SWITCHES attribute is an array with the IDs

```
{
  "type": "object",
  "properties": {
    "ID": {"type": "string"},
    "VNFS": {"type": "array",
      "items": {
        "type": "object",
        "properties": {
          "ID": {"type": "string"},
          "PATH": {"type": "string"}
        }
      }
    },
    "IP": {"type": "object",
      "properties": {
        "ID": {"type": "string"},
        "LINK": {"type": "string"}
      }
    },
    "OPS": {"type": "array",
      "items": {
        "type": "object",
        "properties": {
          "ID": {"type": "string"},
          "LINK": {"type": "string"}
        }
      }
    },
    "CONNECTIONS": {"type": "array",
      "items": {
        "type": "object",
        "properties": {
          "OLL": {"type": "string"},
          "ILL": {"type": "string"},
          "OLL_MAC": {"type": "string"},
          "ILL_MAC": {"type": "string"}
        }
      }
    }
  }
}
```

Figure 5 SFC Simplified JSON Schema

of switches; The other two attributes, CONTROLLERS and OVSWITCHES, refer to the OpenFlow SDN network. The CONTROLLERS attribute contains a list of controllers; each object of this list consists of the ID, controller IP address and the PORT the controller uses to communicate. The configuration and initialization of the OpenFlow controller are out of the scope of NIEP, which is controller-agnostic, even though POX [KSG14] is used as default. The last attribute OVSWITCHES is another object array, which specifies the IDs and connections of OpenFlow switches. Any ID employed in the system must be unique.

Finally, the CONNECTIONS attribute is used to specify the interconnections of Mininet components among themselves and with VNFs and SFCs. The CONNECTION JSON object has two mandatory components: IN/OUT and OUT/IN. The IN/OUTIFACE and OUT/INIFACE indicate the virtual interface where the VNF will set up a connection. In the context of SFCs, the connection is specified in terms of its VNFs, typically one connection is defined for input and one for output. The connections between the VNFs of an SFC are specified in the SFC description file.

5 Case Study and Experimental Evaluation

In this section, we describe a case study executed to obtain empirical results to evaluate the effectiveness of NIEP. First, in Subsection 5.1, the case study described. Next, in Subsection 5.2, results are presented and discussed. Finally, a qualitative evaluation is discussed in Subsection 5.3. The main objective of these experiments is to investigate

```
{
  "type": "object",
  "properties": {
    "ID": {"type": "string"},
    "VNFS": {"type": "array",
      "items": {"type": "string"}
    },
    "SFCS": {"type": "array",
      "items": {"type": "string"}
    },
    "MININET": {
      "type": "object",
      "properties": {
        "HOSTS": {"type": "array",
          "items": {
            "type": "object",
            "properties": {
              "ID": {"type": "string"},
              "IP": {"type": "string"}
            }
          }
        },
        "SWITCHES": {"type": "array",
          "items": {"type": "string"}
        },
        "CONTROLLERS": {"type": "array",
          "items": {
          "type": "object",
            "properties": {
              "ID": {"type": "string"},
              "IP": {"type": "string"},
              "PORT": {"type": "string"}
            }
          }
        },
        "OVSWITCHES": {"type": "array",
          "items": {
          "type": "object",
            "properties": {
              "ID": {"type": "string"},
              "CONTROLLER": {"type": "
                  string"}
            }
          }
        }
      }
    },
    "CONNECTIONS": {"type": "array",
      "items": {
        "type": "object",
        "properties": {
          "IN/OUT": {"type": "string"},
          "OUT/IN": {"type": "string"},
          "IN/OUTIFACE": {"type": "string"}
              ,
          "OUT/INIFACE": {"type": "string"}
        }
      }
    }
  }
}
```

Figure 6 Topology Simplified JSON Schema

whether our platform can efficiently and effectively emulate heterogeneous NFV scenarios deployed on different network topologies. We also assess whether/how NIEP meets the requirements defined in Section 3.

5.1 Case Study: Description

The experimental setup defined is composed of two locations: the Customer Premises (CP) and an Internet Service Provider (ISP), as shown in Fig. 7. In the CP, a Mininet host acting as a client is connected to a VNF with limited resources (1 core, 192MB RAM) running a static router to emulate Customer Premises Equipment (CPE) connected to an ISP. At the ISP side, a VNF with more resources (2 cores, 2GB RAM) running a firewall is connected to a Mininet topology with a virtual OpenFlow switch (OpenVSwitch), which in turn is connected to an SDN Controller and a host acting as an application server.

Four different configurations were used to evaluate how the number of customers connected to the ISP impacts the performance of NIEP. We varied the number of CPEs from 2, 4, 8 and 16. In addition, a second configuration was tested, in which two VNFs implementing a firewall were deployed on the ISP side with the purpose of balancing the load imposed by the CPEs.

The experiments were executed on the following system. The CP instances were executed on an Intel Core i7-6700k@4.00GHz server, with 8GB RAM DDR4, 4 cores, and running CentOS 7. The ISP, in turn, was executed on an Intel Xeon E3-1220v6@3.00GHz, 8GB RAM DDR4, 4 cores, running Ubuntu 14.04. The hosts were connected on a 1Gbps Ethernet network. We employed the KVM hypervisor to deploy both the Mininet VM and VNFs in both hosts.

5.2 Results

Each experiment was repeated 30 times, results are presented with a confidence interval of 99%. NIEP can be used to emulate multiple NFV scenarios, providing fine-grained control over several configuration parameters. The network topology can be easily changed, for example, to test different network paths, or to add, remove, and reconfigure hosts, or also to change link properties. The boot time is an important metric to be evaluated since it can impact the time it takes to evaluate a configuration.

The NIEP components were instrumented to send the boot time to a centralized server. The longest boot time corresponds to the VNFs on the CP side, as their number grows this increases the load on the processor. At the same time, the number of VNFs running on the processor used to emulate the ISP does not change as a single experiment is executed, it only changes from one experiment to another. The VNFs are instantiated at the same time, and Mininet is started before that. Thus, the boot time is the time Mininet takes to initialize plus the average boot time of the VNFs of the CP side. To make it clear, this is the time it takes for the entire emulation setup to be ready to execute. Results are shown in Figure 8.

In the second experiment, two VNFs implementing firewalls were deployed instead of a single one. Because of this, the boot time is slightly higher on the CPs VNFs, as they need to do some additional configuration to send traffic to the two different firewalls. The exception is when two CP instances are deployed (as shown in the first two bars of Figure 3) – the first setup takes 529 ms plus/minus 3 ms, and the second setup 2 takes 532 ms plus/minus of 2.8 ms. This happens because with only two CP instances there are still free processor cores left that are used exclusively by the hypervisor and operating system to execute tasks related to the configuration and deployment of virtual machines.

The evaluation of the performance of NIEP under heavy load allows the identification of bottlenecks caused by different factors on the emulation, due for instance to CPU and memory limitations. Thus if for example, physical resources are not sufficient, this would be reflected in the low throughput between the CPs and the ISP, due to the limited packet processing capacity. In this case, the throughput gets much lower than link speed (1Gbps). *iperf* [TQD+05] was employed in the experiments; all CPs send traffic at the same time to the server running at the ISP side. The values obtained from each CP in each experiment were aggregated since CPs share the same 1Gbps link with which they are connected to the ISP.

As shown in Figure 9, in all experiments executed for both setups the bottleneck was always the link between the CPs and the ISP. In this way, increasing the number of instances does not affect link throughput. We can conclude that NIEP scales well as required.

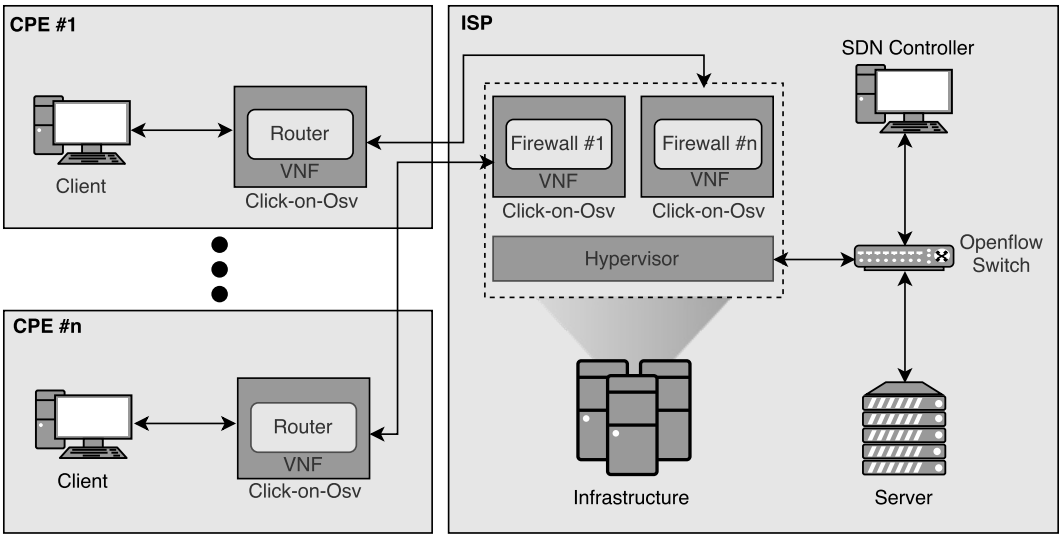

Figure 7 Experimental Evaluation: Environment

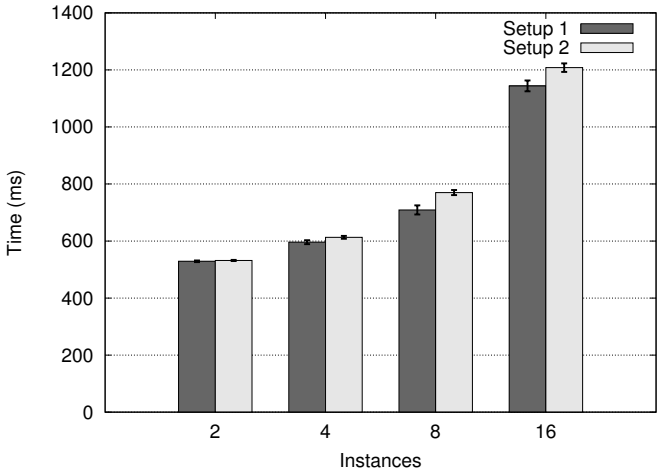

Figure 8 NIEP evaluation: average boot time.

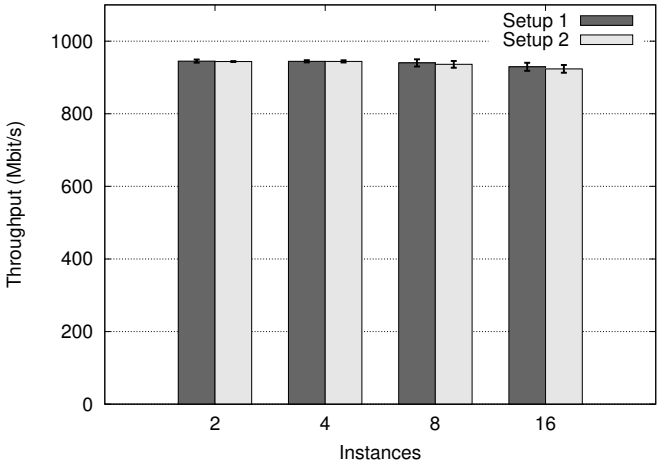

Figure 9 NIEP evaluation: throughput.

5.3 Discussion

Regarding the scalability, NIEP proved to be able to emulate complex scenarios, with increasing numbers of hosts, switches, and VNFs, as well as the overall topology. The good scalability can be explained because in NIEP the VNFs do not run within Mininet. Instead, the VNFs are connected through external bridges, which allow their execution remotely. On the other hand, although EsCAPE can also simulate complex scenarios, VNFs are deployed on the same host, since containers are defined within Mininet and connections must be established locally.

Different from Mininet, the topologies NIEP employs are defined at a high level with JSON, which simplifies the process for deploying the infrastructure. Users can define, in addition to hosts, switches, and controllers, different types of VNFs, different amounts of resources that are allocated for the VNFs, and the connections among them, including the capacity of creating SFCs.

The use of hypervisor-based virtualization for deploying VNFs enables the emulation of heterogeneous network infrastructures since a VNF can be directly deployed on any server and operating system running a compatible hypervisor (*e.g.,* KVM, Xen, and VirtualBox) without requiring any change to the VNF source code. This can be a serious limitation for the use of the other platforms discussed in Section 2.2, which rely on container-based virtualization.

Finally, NIEP is more secure that the other platforms as NIEP is based on VMs while the others are based on containers which have more vulnerabilities.

6 Conclusion and Future Work

The Network Function Virtualization (NFV) paradigm replaces traditional middleboxes with virtual functions that are executed on general purpose hardware. NFV brings multiple advantages in terms of cost and flexibility, but it also brings new challenges. In this work, we presented NIEP, an NFV Infrastructure Emulation Platform to emulate VNFs. Emulation platforms provide a realistic alternative to execute VNFs. This NIEP allows VNFs to be tested and evaluated before they are deployed in production networks. Most existing NFV emulation platforms are based on containers or process virtualization. Furthermore, they do not provide native support for the distribution of the emulation, which should be based on processes running and communicating on different machines. Thus the emulation is limited to a single machine, which has obvious scalability limitations.

NIEP is based on Click-on-OSv and Mininet. NIEP is based on VMs, which provides higher security guarantees that containers. NIEP allows the emulation of different NFV scenarios and VNF design and evaluation, supporting the emulation of heterogeneous infrastructures and scenarios. The evaluation of the performance of NIEP included the boot time and the throughput of VMs running CP and ISP sites. The results show that the boot time of VNFs increases almost linearly, indicating almost no impact of NIEP on VNF instantiation. Moreover, increasing the number of VNFs does not reduce throughput.

Future work includes the design of a user-friendly web interface for network operators, allowing information about VNFs and their operation to be obtained with a REST API. Furthermore, it is important to extend the tool to support different VNF technologies, such as nginx and BRO.

Acknowledgements

This research was performed partially within the project "Federated Ecosystem for Offering, Distribution, and Execution of Virtual Network Functions" (GT-FENDE). The authors would like to thank Rede Nacional de Ensino e Pesquisa (RNP), for their support to the GT-FENDE project.

References

[BCV+18] Carlos Jsus Bernardos, Luis M. Contreras, Ishan Vaishnavi, Robert Szabo, Xi Li, Francesco Paolucci, Andrea Sgambelluri, Barbara Martini, Luca Valcarenghi, Giada Landi, Dmitriy Andrushko, and Alain Mourad. Multi-domain Network Virtualization. Internet-Draft draft-bernardos-nfvrg-multidomain-04, Internet Engineering Task Force, March 2018. Work in Progress.

[BHK07] Ingmar Baumgart, Bernhard Heep, and Stephan Krause. Oversim: A flexible overlay network simulation framework. In *IEEE Global Internet Symposium, 2007*, pages 79–84. IEEE, 2007.

[CS03] Mark Carson and Darrin Santay. Nist net: a linux-based network emulation tool. *ACM SIGCOMM Computer Communication Review*, 33(3):111–126, 2003.

[CW+12] M Chiosi, S Wright, et al. Network functions virtualisation (nfv). *ETSI NFV ISG, White Paper*, 1, 2012.

[dCMGC+17] Leonardo da Cruz Marcuzzo, Vinicius F Garcia, Vitor Cunha, Daniel Corujo, Joao P Barraca, Rui L Aguiar, Alberto E Schaeffer-Filho, Lisandro Z Granville, and Carlos RP dos Santos. Click-on-osv: A platform for running click-based middleboxes. In *Integrated Network and Service Management (IM), 2017 IFIP/IEEE Symposium on*, pages 885–886. IEEE, 2017.

[GMUJ16] J. Garay, J. Matias, J. Unzilla, and E. Jacob. Service description in the nfv revolution: Trends, challenges and a way forward. *IEEE Communications Magazine*, 54(3):68–74, March 2016.

[Han06] M. Handley. Why the internet only just works. *BT Technology Journal*, 24(3):119–129, July 2006.

[HGJL15] Bo Han, V. Gopalakrishnan, Lusheng Ji, and Seungjoon Lee. Network function virtualization: Challenges and opportunities for innovations. *IEEE Communications Magazine*, 53(2), 2015.

[HP15] Joel Halpern and Carlos Pignataro. Service function chaining (sfc) architecture. RFC 7665, 2015.

[ISH10] M. Imran, A. M. Said, and H. Hasbullah. A survey of simulators, emulators and testbeds for wireless sensor networks. In *2010 International Symposium on Information Technology*, volume 2, pages 897–902, June 2010.

[KDP+16] Holger Karl, Sevil Dräxler, Manuel Peuster, Alex Galis, Michael Bredel, Aurora Ramos, Josep Martrat, Muhammad Shuaib Siddiqui, Steven Van Rossem, Wouter Tavernier, et al. Devops for network function virtualisation: an architectural approach. *Transactions on Emerging Telecommunications Technologies*, 27(9):1206–1215, 2016.

[KMC+00] Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, and M Frans Kaashoek. The click modular router. *ACM Transactions on Computer Systems (TOCS)*, 18(3):263–297, 2000.

[KSG14] Sukhveer Kaur, Japinder Singh, and Navtej Singh Ghumman. Network programmability using pox controller. In *ICCCS International Conference on Communication, Computing & Systems, IEEE*, volume 138, 2014.

[LHM10] Bob Lantz, Brandon Heller, and Nick McKeown. A network in a laptop: rapid prototyping for software-defined networks. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, page 19. ACM, 2010.

[LLF+15] K. Lu, S. Liu, F. Feisullin, M. Ersue, and Y. Cheng. Network function virtualization: opportunities and challenges. *IEEE NETWORK*, 29(3):4–5, May 2015.

[MBD16] A. A. Mohallel, J. M. Bass, and A. Dehghantaha. Experimenting with docker: Linux container and base os attack surfaces. In *2016 International Conference on Information Society (i-Society)*, pages 17–21, Oct 2016.

[MKJK99] Robert Morris, Eddie Kohler, John Jannotti, and M. Frans Kaashoek. The click modular router. *ACM SIGOPS Operating Systems Review*, 33(5):217–231, 1999.

[MKK15] R. Morabito, J. Kjllman, and M. Komu. Hypervisors vs. lightweight virtualization: A performance comparison. In *2015 IEEE International Conference on Cloud Engineering*, pages 386–393, March 2015.

[MSG+16] Rashid Mijumbi, Joan Serrat, Juan-Luis Gorricho, Niels Bouten, Filip De Turck, and Raouf Boutaba. Network function virtualization: State-of-the-art and research challenges. *IEEE Communications Surveys & Tutorials*, 18(1):236–262, 2016.

[NFV13] GS NFV. 001: Network functions virtualisation (nfv); use cases, v 1.1. 1. *ETSI, December*, 2013.

[PKvR16] M. Peuster, H. Karl, and S. van Rossem. Medicine: Rapid prototyping of production-ready network services in multi-pop environments. In *2016 IEEE Conference on Network Function Virtualization and Software Defined Networks (NFV-SDN)*, pages 148–153, Nov 2016.

[PPK+15] Ben Pfaff, Justin Pettit, Teemu Koponen, Ethan J Jackson, Andy Zhou, Jarno Rajahalme, Jesse Gross, Alex Wang, Joe Stringer, Pravin Shelar, et al. The design and implementation of open vswitch. In *NSDI*, pages 117–130, 2015.

[SCS+15] Balázs Sonkoly, János Czentye, Robert Szabo, Dávid Jocha, János Elek, Sahel Sahhaf, Wouter Tavernier, and Fulvio Risso. Multi-domain service orchestration over networks and clouds: A unified approach. In *Proceedings of the 2015 ACM Conference on Special Interest Group on Data Communication*, SIGCOMM '15, pages 377–378, New York, NY, USA, 2015. ACM.

[SFMH+13] Alberto Schaeffer-Filho, Andreas Mauthe, David Hutchison, Paul Smith, Yue Yu, and Michael Fry. Preset: A toolset for the evaluation of network resilience strategies. In *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on*, pages 202–209. IEEE, 2013.

[SVZ+14] D. Salopek, V. Vasi, M. Zec, M. Mikuc, M. Vaarevi, and V. Konar. A network testbed for commercial telecommunications product testing. In *2014 22nd International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 372–377, Sept 2014.

[Tea12] Mininet Team. Mininet: An instant virtual network on your laptop (or other pc), 2012.

[TQD+05] Ajay Tirumala, Feng Qin, Jon Dugan, Jim Ferguson, and Kevin Gibbs. Iperf: The tcp/udp bandwidth measurement tool. *http://iperf.fr*, 2005.

[VH08] András Varga and Rudolf Hornig. An overview of the omnet++ simulation environment. In *Proceedings of the 1st international conference on Simulation tools and techniques for communications, networks and systems & workshops*, page 60. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2008.

[WDS+14] Philip Wette, Martin Draxler, Arne Schwabe, Felix Wallaschek, Mohammad Hassan Zahraee, and Holger Karl. Maxinet: Distributed emulation of software-defined networks. In *Networking Conference, 2014 IFIP*, pages 1–9. IEEE, 2014.