Proceedings of the 2017 Winter Simulation Conference
W. K. V. Chan, A. D'Ambrogio, G. Zacharewicz, N. Mustafee, G. Wainer, and E. Page, eds.

REPRODUCIBLE NETWORK RESEARCH WITH A HIGH-FIDELITY SOFTWARE-DEFINED NETWORK TESTBED

Xiaoliang Wu
Qi Yang
Xin Liu
Dong Jin

Illinois Institute of Technology
10 West 31st Street
Chicago, IL, USA

Cheol Won Lee

National Security Research Institute
1559, Yuseong-daero, Yuseong-gu
Daejeon, SOUTH KOREA

ABSTRACT

The transformation of innovative research ideas to production systems is highly dependent on the capability of performing realistic and reproducible network experiments. In this work, we present a hybrid testbed to advocate high fidelity and reproducible networking experiments, which consists of network emulators, a distributed control environment, physical switches, and end-hosts. The testbed (1) offers functional fidelity through unmodified code execution on an emulated network, (2) supports large-scale network experiments using lightweight OS-level virtualization techniques and capable of running across distributed physical machines, (3) provides the topology flexibility, and (4) enhances the repeatability and reproducibility of network experiments. We validate the fidelity of our hybrid testbed through extensive experiments under different network conditions and compare the results with the benchmark data collected on physical devices. We also use the testbed to reproduce key results from published network experiments, such as Hedera, a scalable and adaptive network traffic flow scheduling system.

1 MOTIVATION

Software-Defined Networking (SDN) is a novel network architecture which aims to simplify network management by decoupling the control logic from the underlying network infrastructure and offers online and direct network programmability to enable innovation. It draws massive attention from both academic and industrial network research community. Thus, it is critical to have a testing and evaluation framework to enable researchers to conduct high fidelity analysis and evaluation of their research ideas before the real system deployment. An ideal testbed ought to offer desired fidelity, scalability, flexibility and reproducibility for network experiments. Currently, there are three common types of SDN testbeds, i.e., simulation, emulation, and physical testbeds, and each type has different advantages and disadvantages in terms of the aforementioned properties. In particular, physical network testbeds offer high fidelity but are often technically challenging and economically infeasible to perform large-scale experiments, and it is also hard to reproduce experimental results. Simulators, on the other hand, offer good scalability and flexibility to set up network experiments but have relatively low fidelity due to model abstraction and simplification. Virtual-machine-based emulation testbeds, such as Mininet (Handigol et al. 2012), offer a balanced solution: they allow unmodified code execution for high functional fidelity, and the lightweight OS-level virtualization also offers reasonably good scalability. However, virtual machines share the underlying physical resources, which may violate timing realism due to resource contention, scheduling serialization, and background system load. Thus, we aim to build a hybrid SDN testbed which improves the existing testbeds by enabling distributed emulators and hardware integration.

978-1-5386-3428-8/17/$31.00 ©2017 IEEE

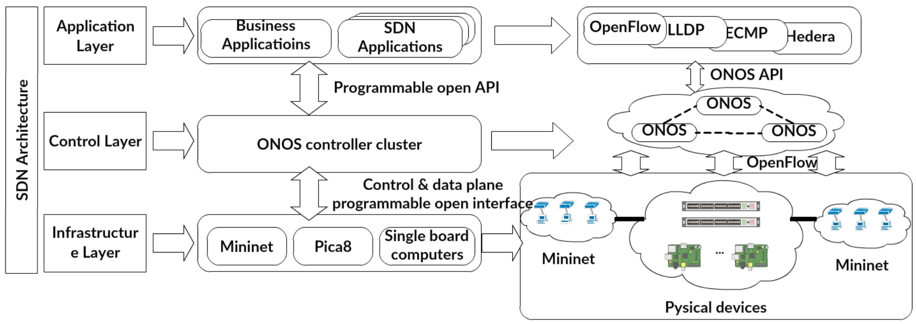

Figure 1: Testbed architecture.

2 SYSTEM DESIGN

Figure 1 presents the testbed architecture. The application layer contains runs the target network application prototypes that researchers aim to evaluate. The control layer consists of the Operating Network Operating System (ONOS) distributed controller over a cluster of commodity servers, which significantly increases the size of network experiments that one can emulate.

The infrastructure layer is built on the container-based network emulation, which enables high fidelity analysis by allowing real networking applications to run in the network emulator and interact with the controller. We used Mininet as the network emulator to provide flexible configuration and direct network programmability as an inherent advantage. To enhance the testbed scalability, we enabled the distributed emulation feature in our testbed by connecting multiple instances of Mininet on different physical machines. In addition, the network layer also supports physical hardware integration to enhance the testbed fidelity. We have integrated Pica8 physical switches as well as Raspberry Pi physical end-hosts. The testbed can run network experiments in a hybrid mode concurrently supporting both virtual and physical devices. Besides SDN network experiments, we also extended Mininet with the ability to model traditional network routers by incorporating Quagga software router (Jakma and Lamparter 2014) into it.

3 VALIDATION EXPERIMENTS AND CASE STUDY

We have performed extensive experiments to evaluate and validate our testbed. We conducted three sets of experiments to demonstrate the improvements of measured throughputs of (1) distributed v.s. standalone testbed, (2) hardware integrated v.s. non-hardware testbed and (3) different placements of the hardware in the emulated network. The experiment results concluded that our testbed provides the realism, scalability, and flexibility as required features of an evaluation testbed. As a case study, we investigated Hedera (Al-Fares, Radhakrishnan, Raghavan, Huang, and Vahdat 2010) traffic flow scheduling system and compared it with Equal Cost Multipath (ECMP) algorithm. In the original paper, Hedera always generates higher bisection throughput (around 12 Mbps) than ECMP (around 10 Mbps) using two emulation modes. Using our testbed, we conducted the same experiments and reproduced the same results as in the original paper.

REFERENCES

Al-Fares, M., S. Radhakrishnan, B. Raghavan, N. Huang, and A. Vahdat. 2010. "Hedera: Dynamic Flow Scheduling for Data Center Networks". In *Proceedings of the 7th USENIX Conference on Networked Systems Design and Implementation*, NSDI'10, 19–19. Berkeley, CA, USA: USENIX Association.

Handigol, N., B. Heller, V. Jeyakumar, B. Lantz, and N. McKeown. 2012. "Reproducible network experiments using container-based emulation". In *Proceedings of the 8th International Conference on Emerging Networking Experiments and Technologies*, 253–264. ACM.

Jakma, P., and D. Lamparter. 2014, 03. "Introduction to the Quagga Routing Suite". 28:42–48.