Characterization of SDN Switch Response Time in Proactive Mode

Rob Edwards, Yinghua Ye, Anuj Kaul, James Yu
CEF Technology and Innovation, Nokia Solutions and Networks,
380 N. Bernardo Ave, Mountain View, CA 94043

Abstract

Reconfiguration of the Fronthaul network allows the mobile Radio Access Network (RAN) to adapt to varying user profiles and traffic loads. Network performance optimization and energy savings can be achieved. This paper examines whether SDN enabled commercial switches can be dynamically reconfigured with bounded and predictable response times. A methodology is presented for measuring response time in proactive mode. Measurement results of two commercial switches are summarized.

1. Background and Motivation

To date, SDN work has focused on traditional data center applications. However, disaggregation and centralization of the network control plane can also apply to mobile networks. Two trends are anticipated in SDN enabled mobile networks. First, portions of the mobile network will migrate into the data center cloud. For instance, virtualization of EPC (Evolved Packet Core) and partial mobile base stations, such as Baseband Units (BBUs) in RAN, have become important use cases in Network Function Virtualization (NFV) [1]. Second, portions of the mobile network such as Fronthaul and Radio Access Units (RAUs) will likely remain outside the cloud and its functionality can be optimized through SDN. Here, Fronthaul refers to the transport network between RAUs and BBUs.

The existing commercial Fronthaul architecture supports only one to one mapping between BBU and RAU. Such inflexibility leads to suboptimal performance [2]. [2] showed that dynamic reconfiguration of the Fronthaul is needed to optimize network bandwidth efficiency and energy savings. Reconfiguration adapts the network to varying user profiles and traffic loads. However, the proof of concept implementation in [2] has some limitations: 1) a Coarse Wavelength Division Multiplexing (CWDM) customized optical switch was required to achieve re-configurability; 2) the traffic profiles were semi-static. To overcome these limitations, we conduct feasibility studies on whether commercial SDN Ethernet switches can be used to support Fronthaul. Both dynamic and low latency reconfiguration is needed in the Fronthaul to support varying user profiles, varying traffic loads and some LTE Advanced features. The switch response time plays a key role in determining supported reconfiguration rates and reconfiguration latencies. Our studies provide a systematic approach to characterize reconfiguration response time of SDN switches.

The contributions of this paper include the following recognitions: SDN support of low latency Fronthaul reconfiguration will be likely needed as mobile networks evolve; SDN switch response time is a key factor in Fronthaul reconfiguration latency and rate. However, too little attention has been given to SDN switch response latency. This paper also contributes a methodology for SDN switch response characterization to help switch vendors optimize reconfiguration response time.

2. Related work

SDN switches are enabled by the OpenFlow (OF) protocol where an SDN controller can program switch entries dynamically. Flexible mapping between BBU and RAU can be supported since OF is capable of doing various flow table manipulations. However, the challenge is to reconfigure with low latency and minimal traffic disruption. Flow reconfiguration can be done in two modes, i.e. reactive mode and proactive mode. [3] has characterized the response time of SDN switches from three vendors in reactive mode. Unfortunately, none of these switches can achieve less than 50ms response time. [3] also cannot separate egress packet delay from ingress packet delay due to methodology limitations. Operation in proactive mode reduces configuration latency when the control plane initiates configuration changes triggered by high level system policy. This is in contrast to reconfiguration triggered by low level events like missing flow table entries. We are exploring whether fast reconfiguration response time can be met if proactive mode is employed.

3. Measurement Setup

We have devised a methodology to characterize SDN switch response time in proactive mode. The test-bed, shown in Figure 1, is designed to measure latency for both control plane and data plane traffic. Measurements are based on hardware timestamps applied to both control and data packets. POX 1.0 is used as OpenFlow controller. A switch and mirror port send duplicate OF packets to both switch and IXIA. The IXIA also sends and receives the data stream through the switch. Returned data and mirrored OF packets are captured and time stamped. OF commands alternately enable and disable the data stream via flow table updates. Analysis of captured time stamps reveals the latency between arrival of OF command packets and their effect on the data stream.

This methodology provides higher fidelity compared with existing approaches [3-4] since it can pinpoint how fast the SDN switch updates its flow entry table with low ambiguity.

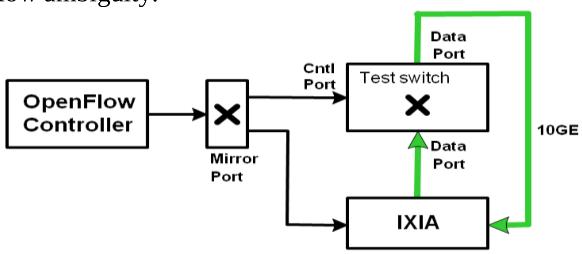

Figure 1: Measurement Setup.

4. Performance Evaluation

Measurements on two commercial switches from different vendors have been performed. In this paper, we name a pure SDN switch from vendor A as switch A, and a hybrid switch from vendor B as switch B. Since we can observe control plane and data plane traffic separately, our characterization methodology offers some insight into SDN switch implementation.

We have observed some common SDN switch implementation problems such as how packet-in events are mishandled on two studied switches and how this issue impacts switch response time. When a new flow arrives and there is no flow table entry, the normal behavior should be: the switch forwards the first packet to controller and buffers the rest of arrived packets from the same flow until the flow table gets updated. However, switch A forwards all packets to the controller until the flow table is updated. Under the same scenario, switch B floods all data packets to all ports while forwarding all packets to the controller until the flow table is updated. For both switches, OF response time is significantly increased by the switch's forwarding packets to the controller. To work around this issue for switch A, we set an entry that drops all unknown traffic. For switch B, the packet-in event is disabled. To prevent packet flooding, we block the flow first and then enable the flow, so we can easily measure the switching response time in proactive mode.

Our measurements showed switching delay at switch A is from 1.88*us* to 2*us* while switching delay at switch B varies from 3.94 ~ 3.98 *us*. The preliminary results for switch response time distribution in proactive mode are shown in Figure 2 and Figure 3 for switch A and switch B respectively. For each switch, 98 samples have been measured. Switch response time generally falls within a 1ms to 5ms range for Switch A. However, there are seven outliers with maximum value 69ms. Switch B has better performance. Its response time mostly lies between 1ms and 2ms but it also has four outliers with maximum value of 20ms. This makes response time quite non-deterministic and it would be challenging to use these switches in applications requiring low deterministic response time.

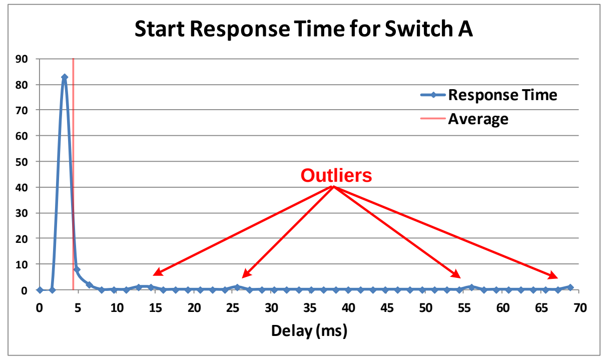

Figure 2: Switch A response time distribution.

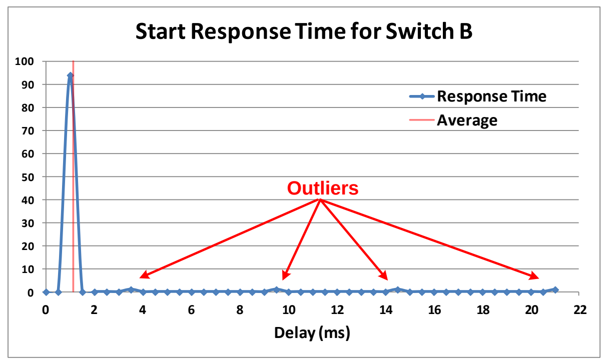

Figure 3: Switch B response time distribution

5. Conclusion

Commercial switch OF agents are being developed and deployed without sufficient attention to response time determinism. Existing embedded OF agents may not adequately support future SDN applications requiring low deterministic response times. Further work is needed to identify response time requirements and promulgate those to switch silicon and platform vendors.

6. References

[1] ETSI specification: Network Functions Virtualisation (NFV) -Use Cases, Oct. 2013, http://www.etsi.org/deliver/etsi_gs/NFV/001_099/001/01.01.01_60/gs_NFV001v010101p.pdf

[2] C. Liu, K. Sundaresan, M. Jiang, S. Rangarajan, G. Chang, "The Case for Re-configurable Backhaul in Cloud-RAN based Small Cell Networks," in Proceedings IEEE INFOCOM, 2013.

[3] D.Huang, K. Yocum, A. Snoeren, "High-Fidelity Switch Models for Software-Defined Network Emulation," in ACM SIGCOMM Workshop on Hot Topics in SDN, 2013

[4] D . Jin, D. Nicol, and M. Caesar. Efficient gigabit Ethernet switch models for large-scale simulation. In IEEE Workshop on Principles of Advanced and Distributed Simulation, pages 1 –10, May 2010.