Latency in Software Defined Networks: Measurement and Mitigation Techniques

Keqiang He, Junaid Khalid , Aaron Gember-Jacobson, Sourav Das, Chaithan Prakash, Aditya Akella, Li Erran Li, Marina Thottan

Latency in SDN

Timely interaction between an SDN controller and switches is crucial to many applications like MicroTE, fast Failover, mobility, etc. These applications assume that the latency in interacting with the network switches is constant and negligible. However our measurement studies shows that this latency is significant. Moreover, it varies with the switch platforms, type of operations performed, table occupancy and concurrent operations on the switches.

Using grey-box probing, we narrow down the key factors for these latencies to be TCAM organization, low power switch CPU and software implementation inefficiencies. To overcome the latencies and achieve responsive control, we develop a systematic framework leveraging both the logically central view and global control in SDN, and the dissection of latencies from our measurement study.

Elements of Latency

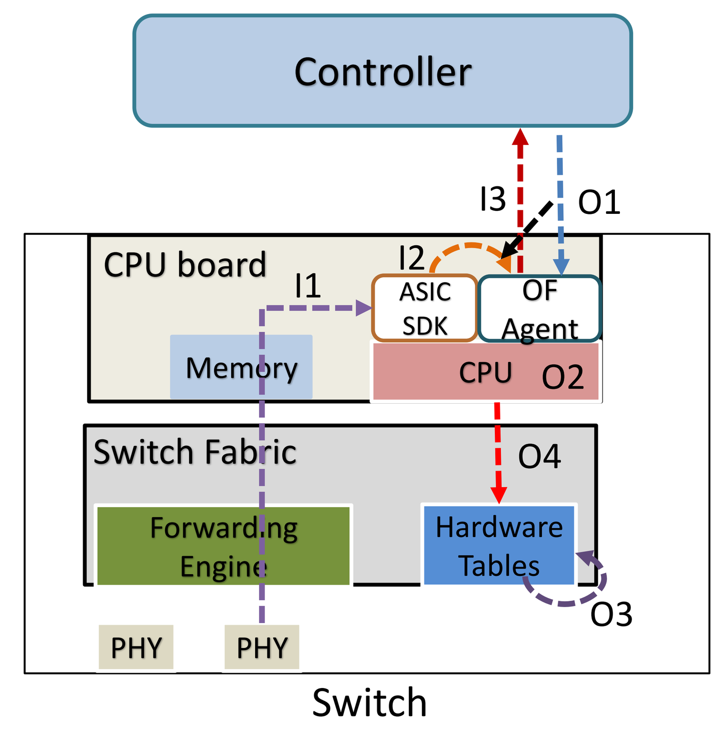

Inbound Latency

- **I1:** Send to ASIC SDK
- **I2:** Send to OF Agent
- **I3:** Send to Controller

Outbound Latency

- **O1:** Parse OF Message
- **O2:** Software schedules the rule
- **O3:** Reordering of rules in table
- **O4:** Rule is updated in table

Measurement Setup

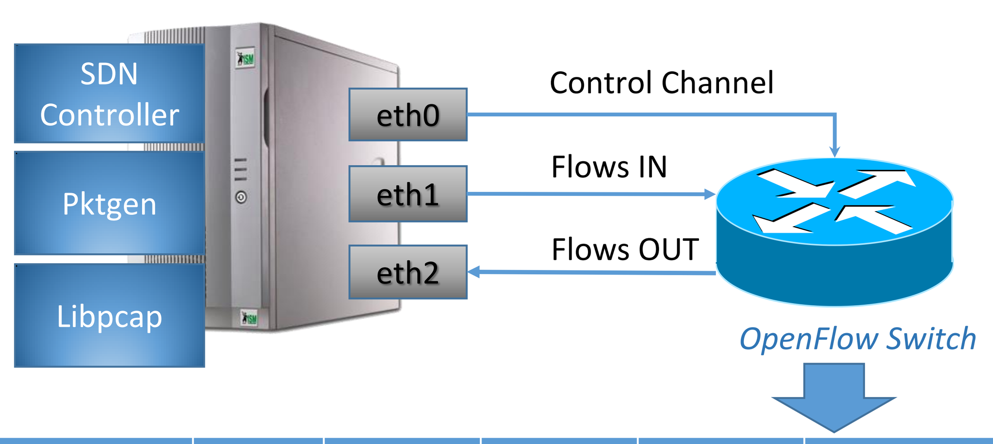

Switch	CPU	RAM	OF Version	Flow table size	Data Plane
Vendor A	2 Ghz	2 GB	1.0	4096	40*10G+4*40
Vendor B-1.0	1 Ghz	1 GB	1.0	896	14*10G+4*40G
Vendor B-1.3			1.3	1792 (ACL)	
Vendor C	?	?	1.0	750	48*10G+4*40G

Inbound Latency

- Increases with flow arrival rate
- Increases with interference from outbound msgs

Flow Arrival Rate (packets/sec)	Mean Delay per packet_in (msec)
100	3.32
200	8.33

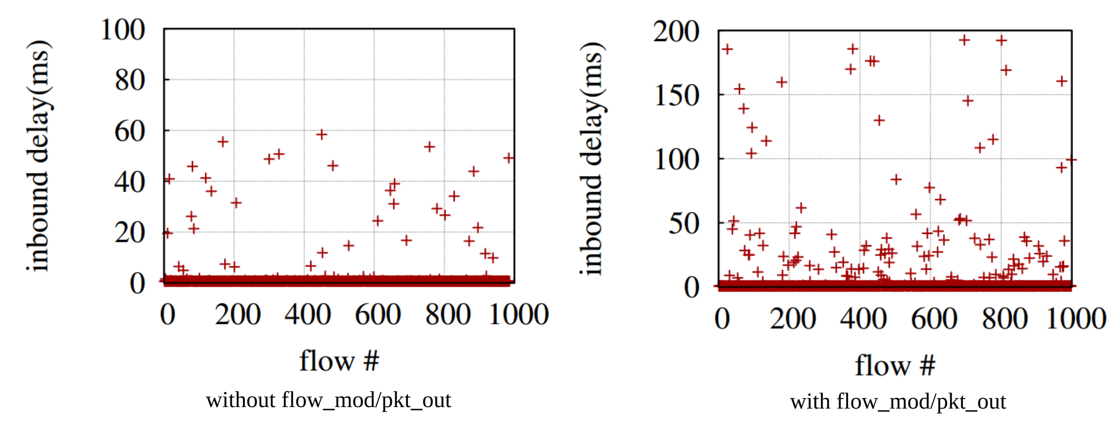

Outbound Latency - Insertion

- Affected by priority insertion patterns
- Affected by the table occupancy

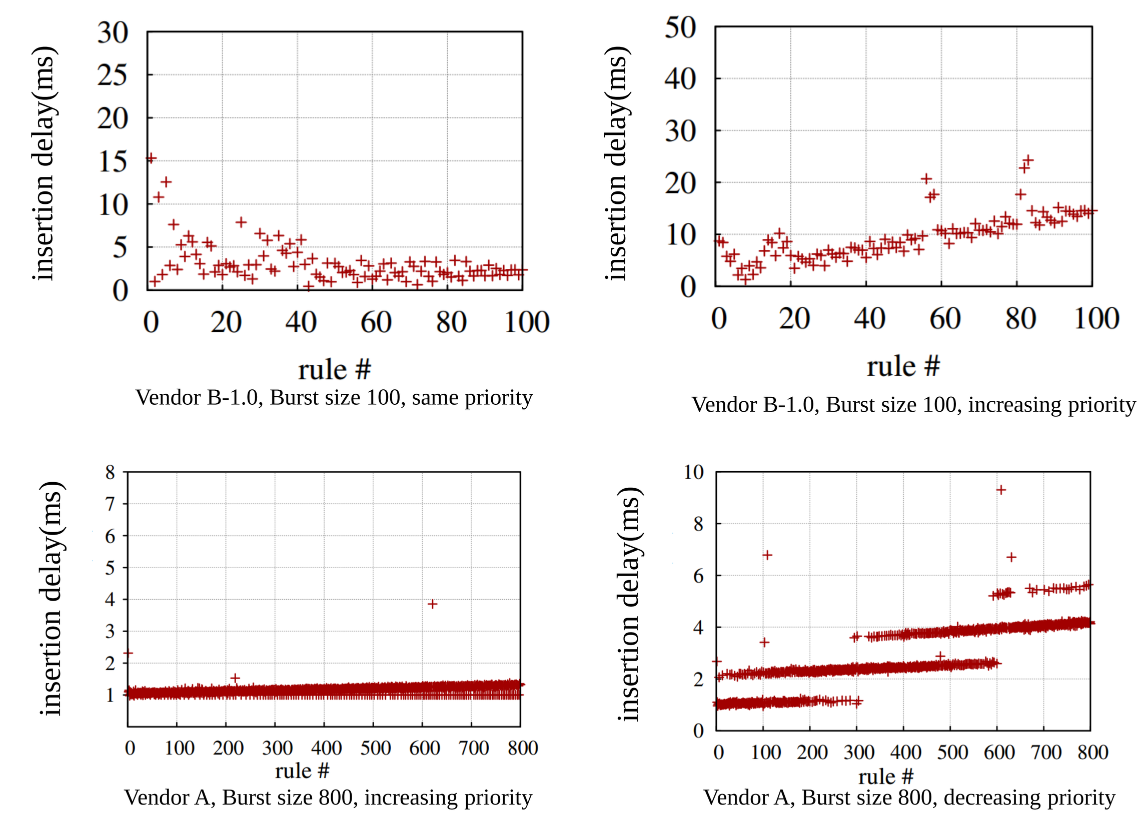

Outbound Latency – Modify/Delete

- Higher than Insertion latency
- Not affected by rule priority but affected by table occupancy

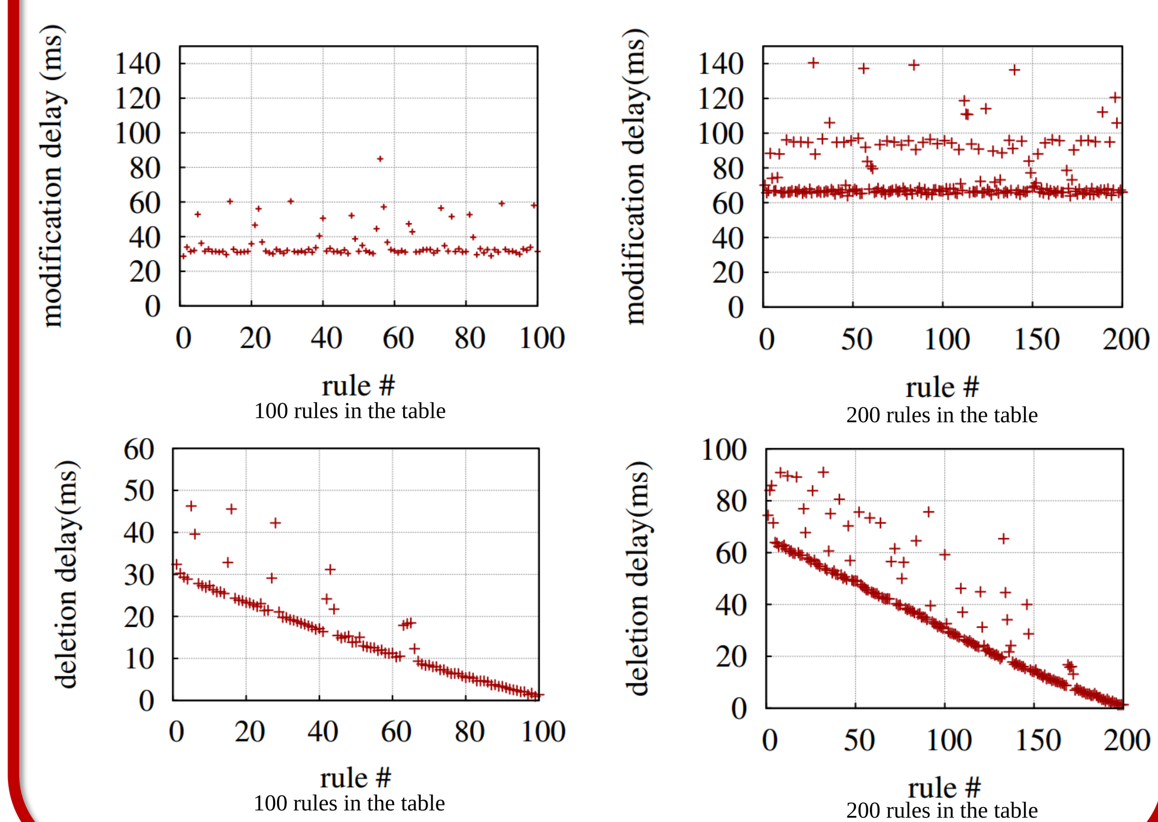

Solution Overview

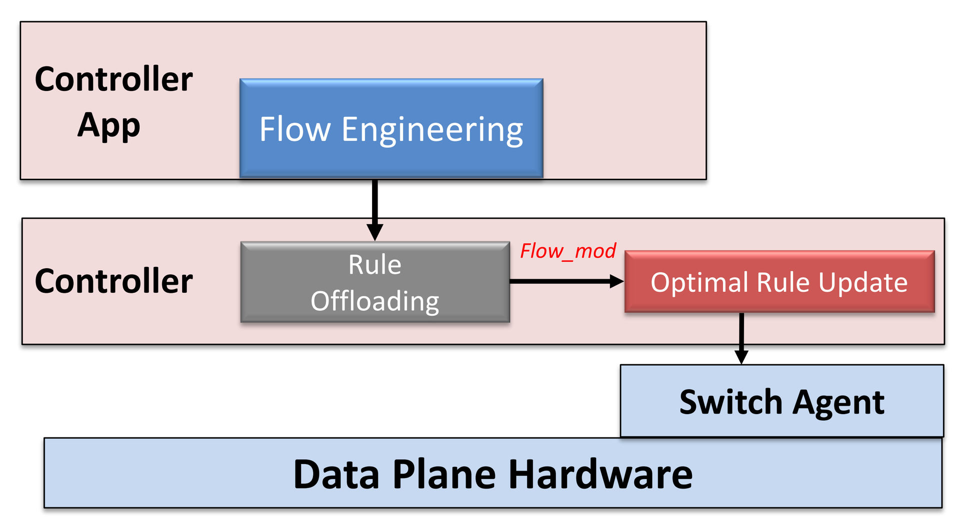

Optimal Rule Update

- Measurements show that *optimal order* of rule insertion *varies with switch platform*
- Goal: Control the actual rule insertion using the *pattern* that is *optimal* for the switch

Flow Engineering

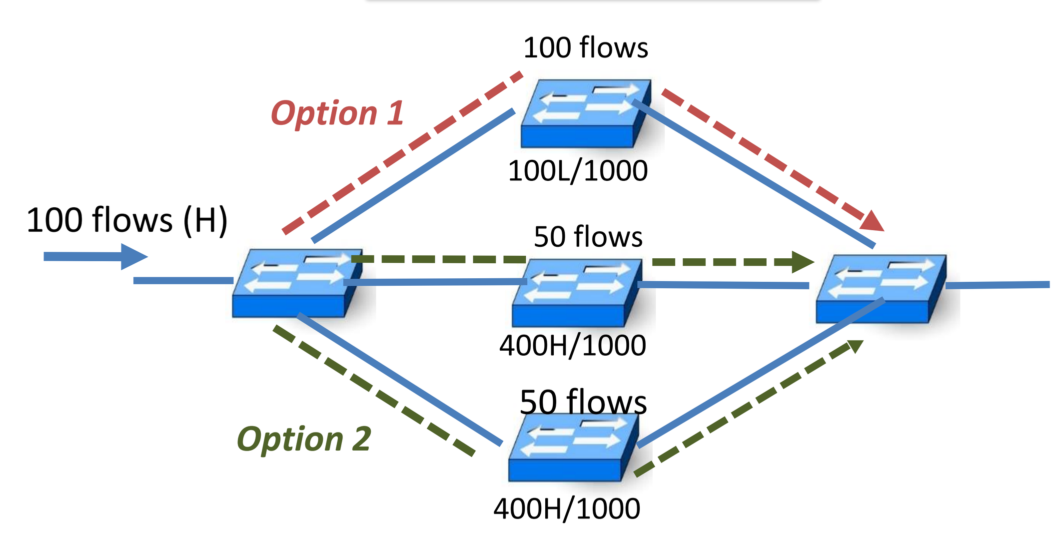

- Goal: Select paths across the network such that ***installation delay is minimized*** and *the network objective is satisfied*
- *Minimizes the aggregate impact* of both **rule displacement** in TCAM and **total number** of rules

Rule Offloading

- Networks with tunnels typically sees less *churn in forwarding state* in underlay network as compare to the end points
- Leverages this characteristic to *offload rules*
- Goal: Minimizes the installation latency by *offloading rules to* **underlay switches**

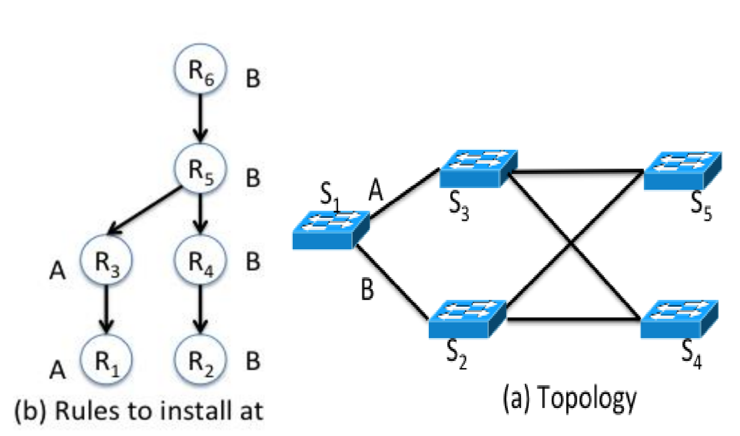

Switch	Rule
S1	R5
	R6
	RA
	RB
S3	R1
	R3
S2	R2
	R4

Performance

- Simulated failover scenario in a tunneled WAN Network
- *Topology:* Full mesh with 25 nodes
- *Traffic matrix:* Assign a popularity index to each node
- *Table occupancy:* Assume switches have some pre-installed rules
- *Workloads:* 6 workloads which have different table occupancies and traffic volumes