# Understanding Latency in Software Defined Networks

**Junaid Khalid**

Keqiang He, Sourav Das, Aditya Akella

Li Erran Li, Marina Thottan

WISCONSIN
UNIVERSITY OF WISCONSIN–MADISON

Bell Labs

# Latency in SDN

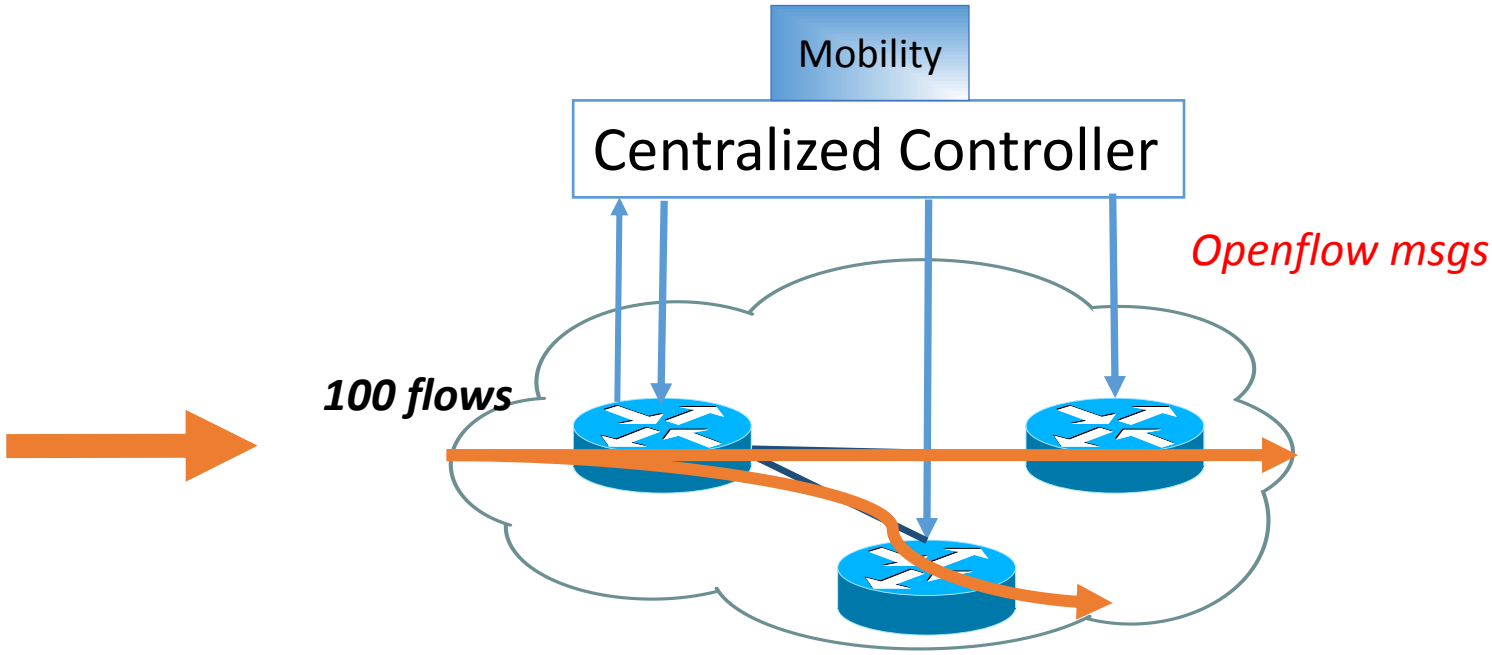


**Time taken to install 100 rules ?**

*Can be XXX msecs !!!*

# Do applications care about latency?

Latency is critical to many applications

- ***MicroTE*** *routes predictable traffic on time scales of 1-2s*

Limits the applicability of SDN applications

*Latency can undermine MicroTE's effectiveness*

Applications assume latency is low and constant

- *Reroute the affected flows quickly in face of failures*
- *Longer update time increases congestion and drops*

*Latency can inflate failover time by nearly 20s!*

# Factors contributing to Latency?

Robust Control Software Design and Distributed Controllers

Speed of Control Programs and Network Latency

Not received much attention

Latency in network switches

# Our Work

**Two contributions:**

- Systematic experiments to explore latencies in production switches
- Design a framework to overcome the impact of latencies

Latency in network switches

# Outline

- Motivation
- Elements of Latency
- Measurement Methodology
- Inbound Latency
- Outbound Latency
  - Insertion
  - Modification
  - Deletion

# Outline

- Motivation
- **Elements of Latency**
- Measurement Methodology
- Inbound Latency
- Outbound Latency
  - Insertion
  - Modification
  - Deletion

# Elements of Latency

1. Inbound Latency

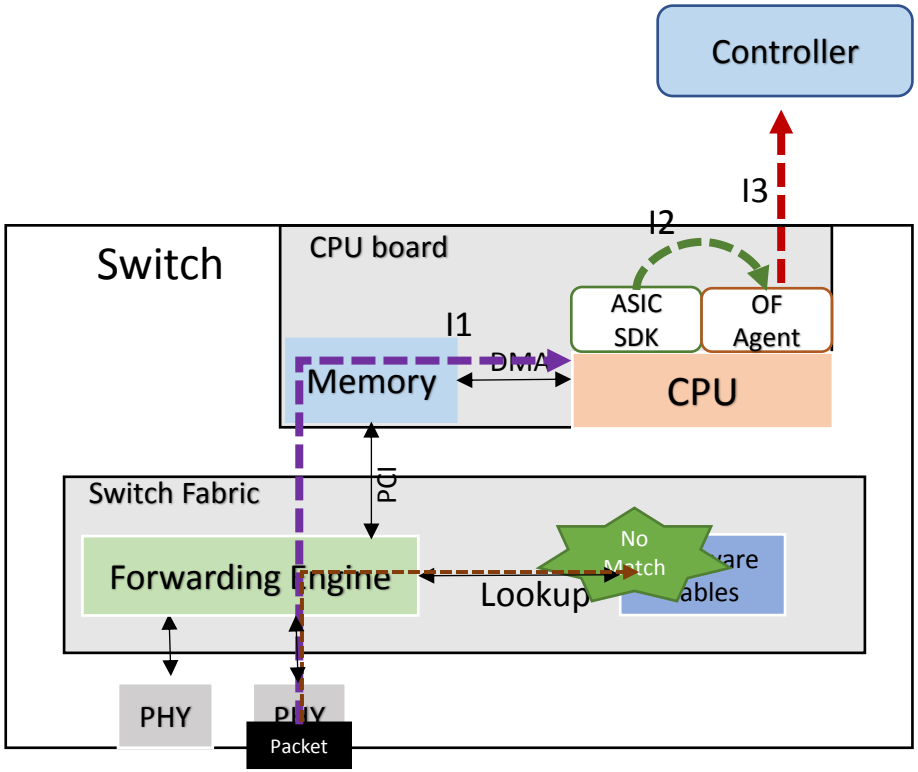


**Inbound Latency**

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

# Elements of Latency

1. Inbound Latency
2. Outbound Latency

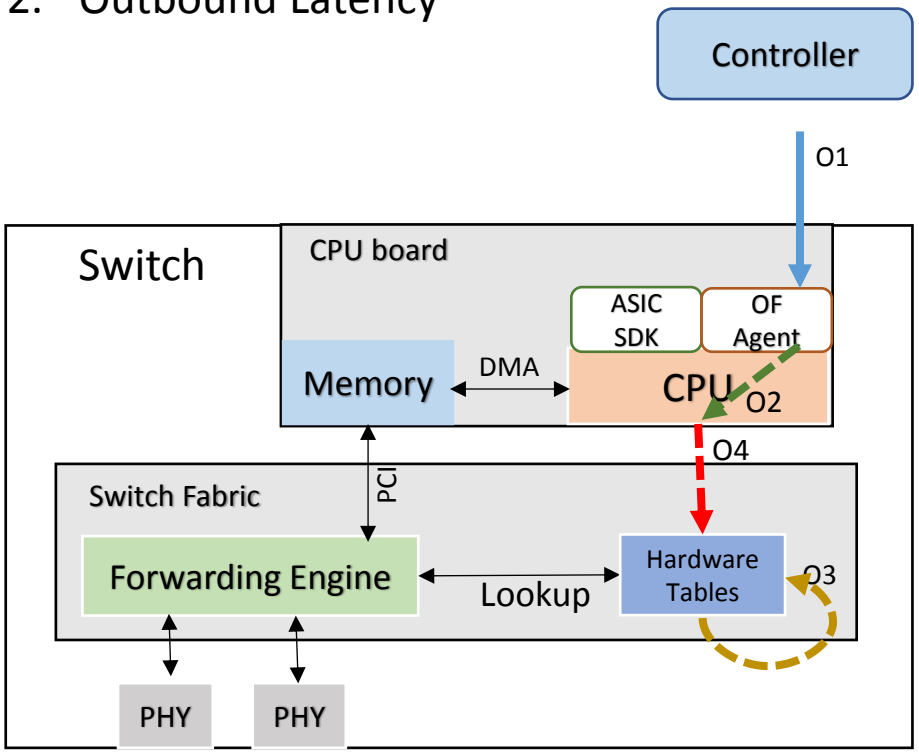


**Outbound Latency**

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

# Outline

- Motivation
- Elements of Latency
- **Measurement Methodology**
- Inbound Latency
- Outbound Latency
  - Insertion
  - Modification
  - Deletion

# Latency Measurements - Setup

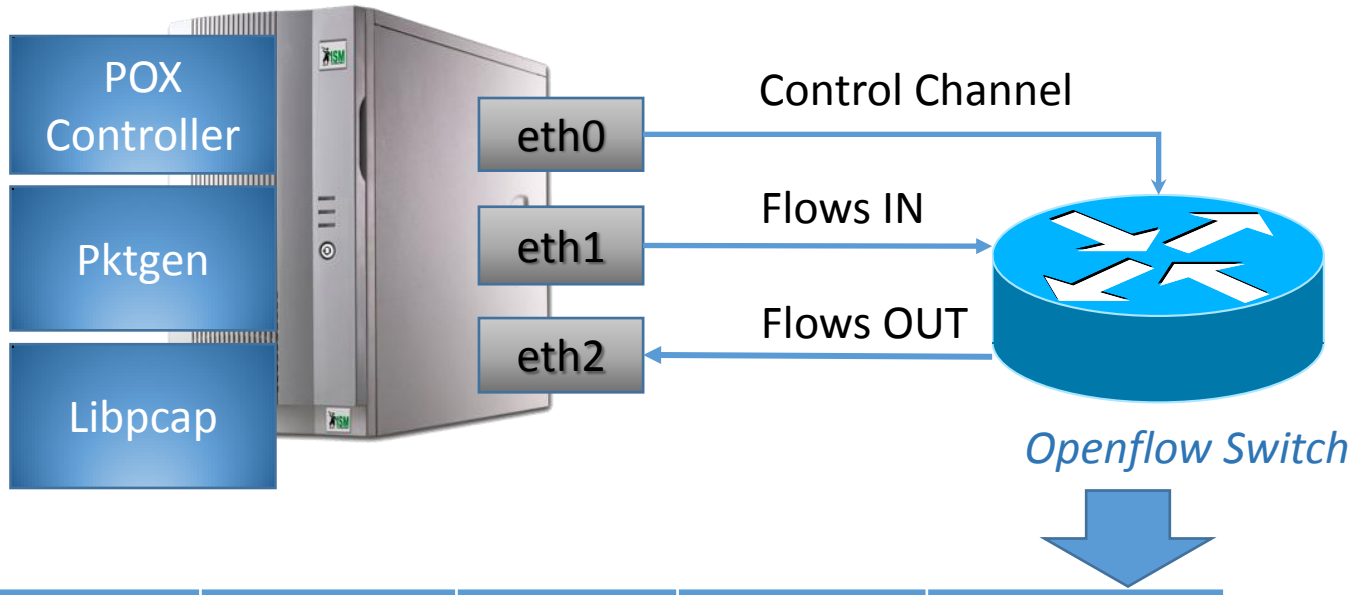


| Switch | CPU | RAM | Flow table size | Data Plane |
|---|---|---|---|---|
| **Vendor A** | 1 Ghz | 1 GB | 896 | 14*10Gbps + 4*40Gbps |
| **Vendor B** | 2 Ghz | 2 GB | 4096 | 40*10Gbps +4*40Gbps |

# Outline

- Motivation
- Elements of Latency
- Measurement Methodology
- **Inbound Latency**
- Outbound Latency
  - Insertion
  - Modification
  - Deletion

# Inbound Latency

- Increases with flow arrival rate
- CPU Usage is higher for higher flow arrival rates

| Low Power CPU | | |
|---|---|---|
| **Flow Arrival Rate (packets/sec)** | **Mean Delay per packet_in (msec)** | **CPU Usage ( % )** |
| 0 | - | 7.1 |
| 100 | 3.32 | 15.7 |
| 200 | 8.33 | 26.5 |

**vendor A** switch

# Inbound Latency

- Increases with interference from outbound msgs

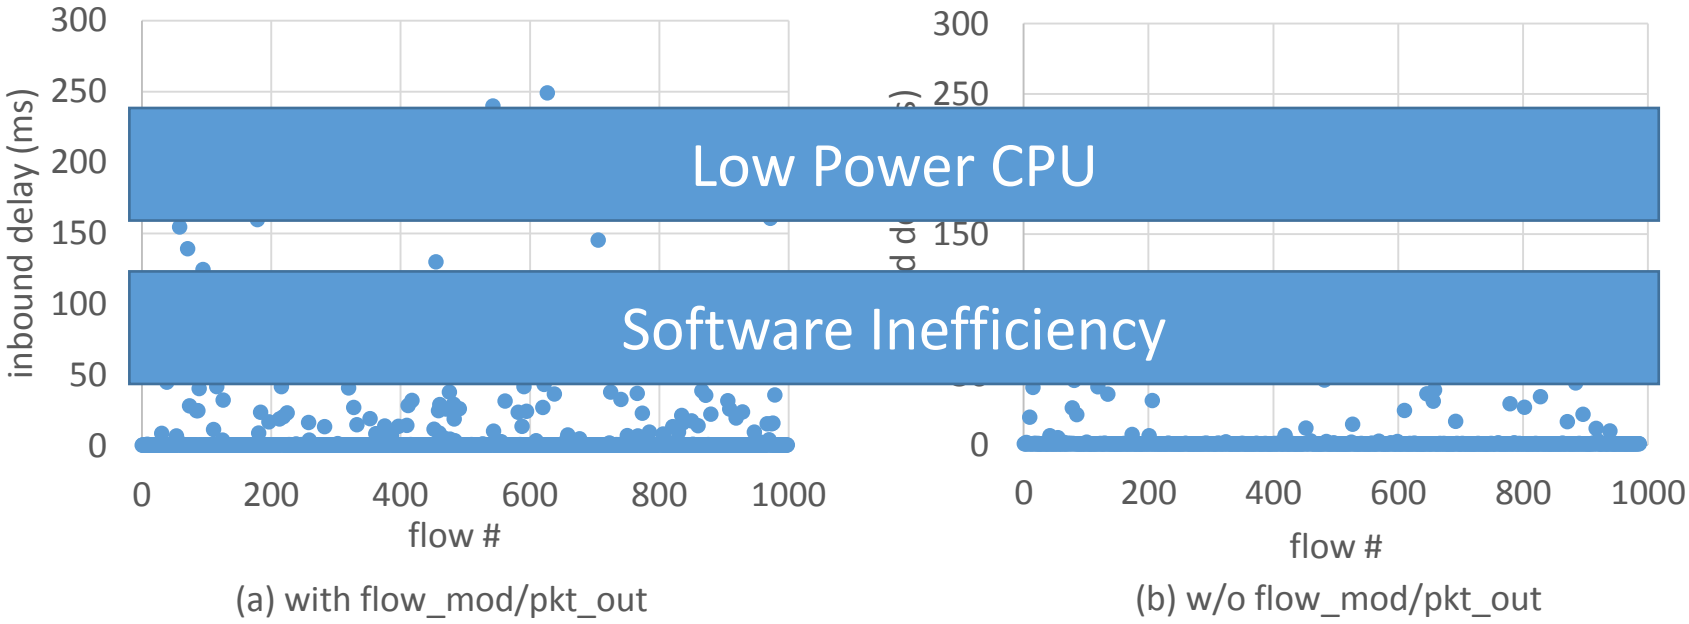


vendor A switch 200 flows/sec

# Outline

- Motivation
- Elements of Latency
- Measurement Methodology
- Inbound Latency
- **Outbound Latency**
  - Insertion
  - Modification
  - Deletion

# Outbound Latency

- Latency for three different flow_mod operations
  - Insertion
  - Modification
  - Deletion
- Impact of key factors on these latencies
  - Table occupancy
  - Rule priority structure

# Outline

- Motivation
- Elements of Latency
- Measurement Methodology
- Inbound Latency
- **Outbound Latency**
  - **Insertion**
  - Modification
  - Deletion

# Insertion Latency – Priority Effects

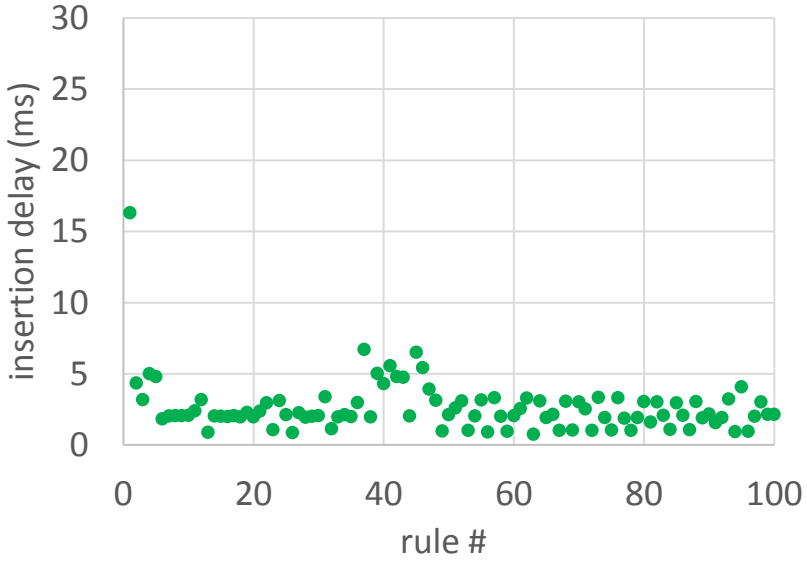


(a) Burst size 100, **same** priority

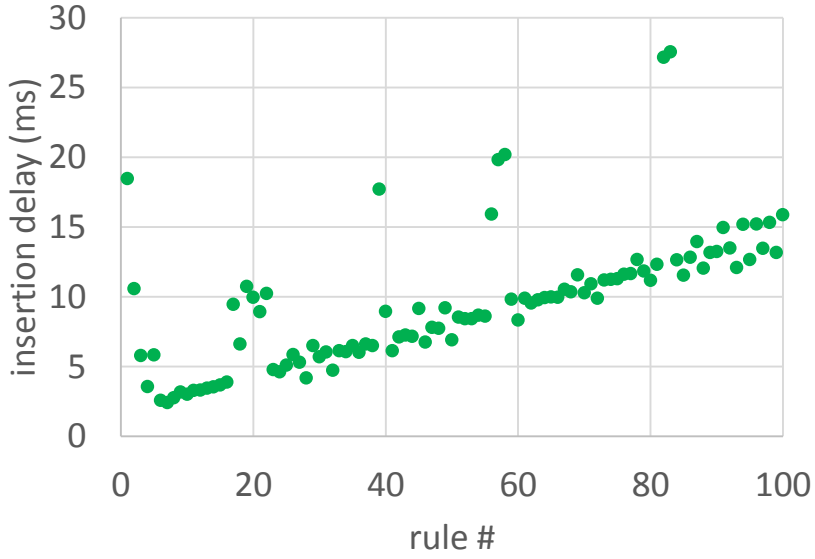


(b) Burst size 100, **increasing** priority

**Vendor B** switch

# Insertion Latency – Priority Effects

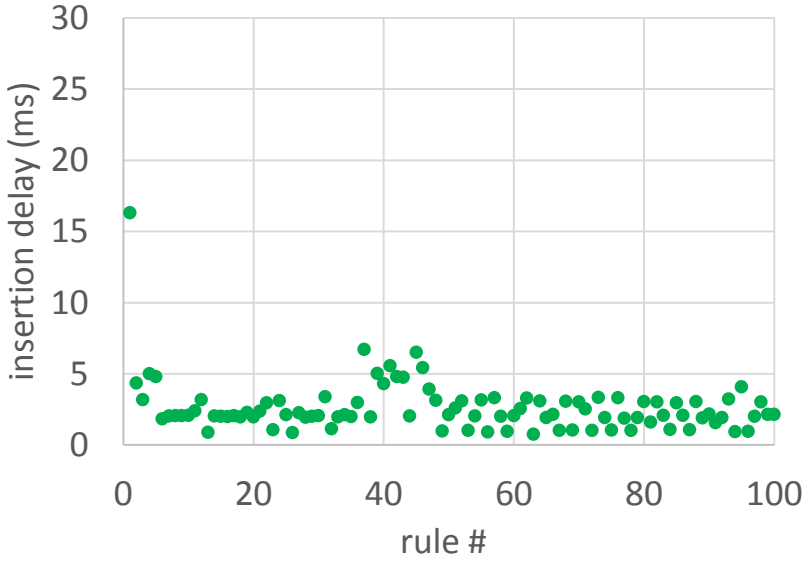


(a) Burst size 100, **same** priority

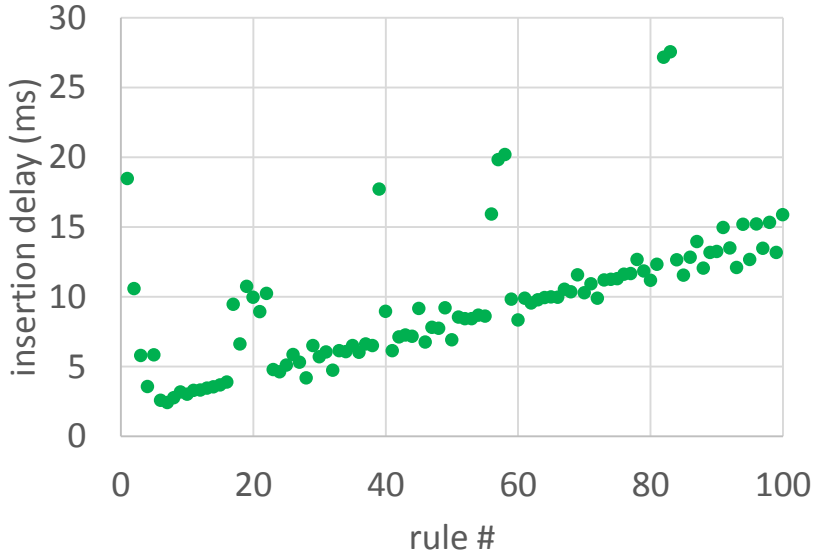


(b) Burst size 100, **increasing** priority

**Vendor B** switch

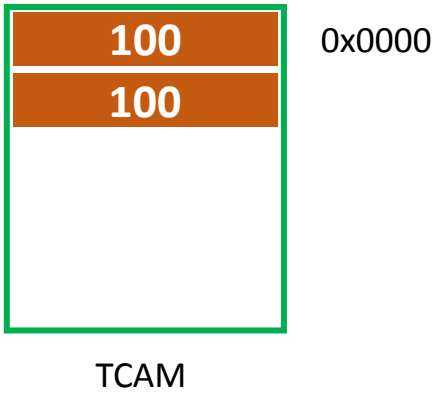

# Insertion Latency – Priority Effects

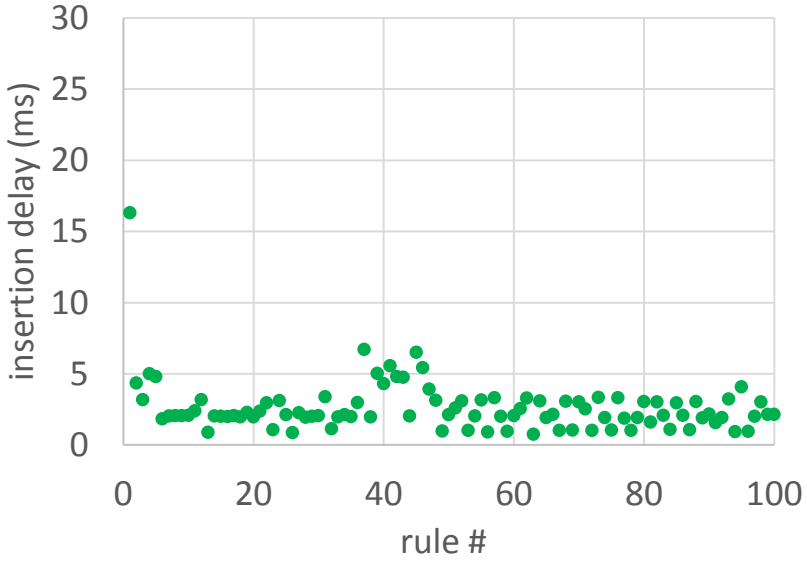


(a) Burst size 100, **same** priority

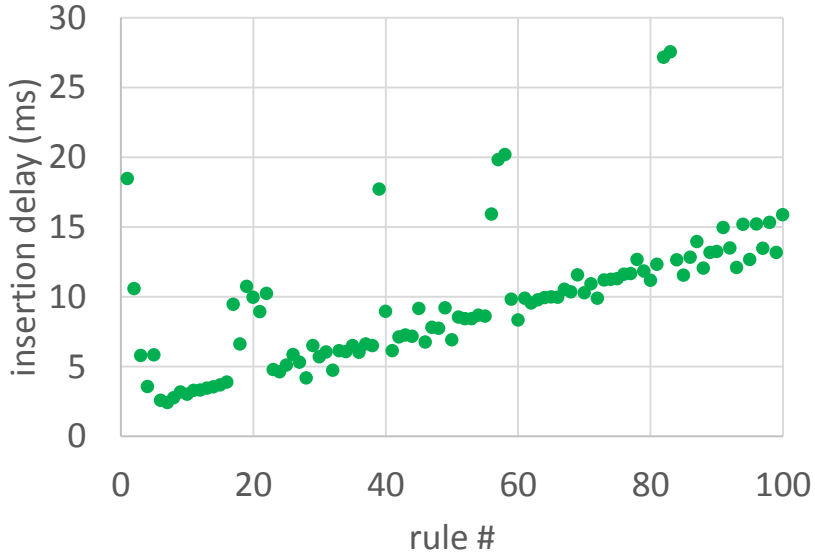


(b) Burst size 100, **increasing** priority

**Vendor B** switch

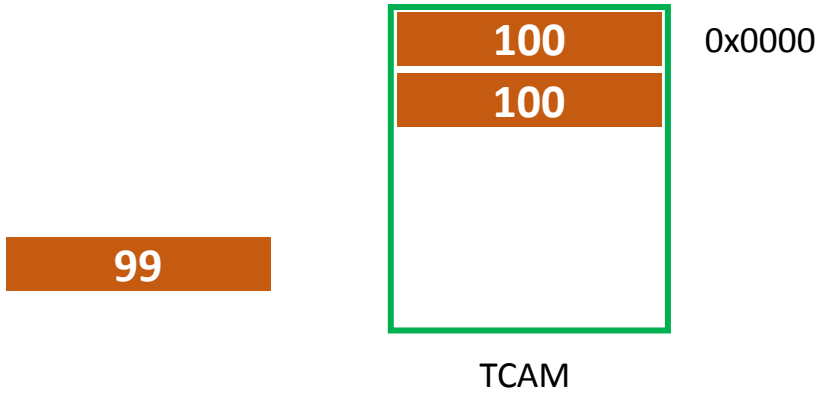

# Insertion Latency – Priority Effects

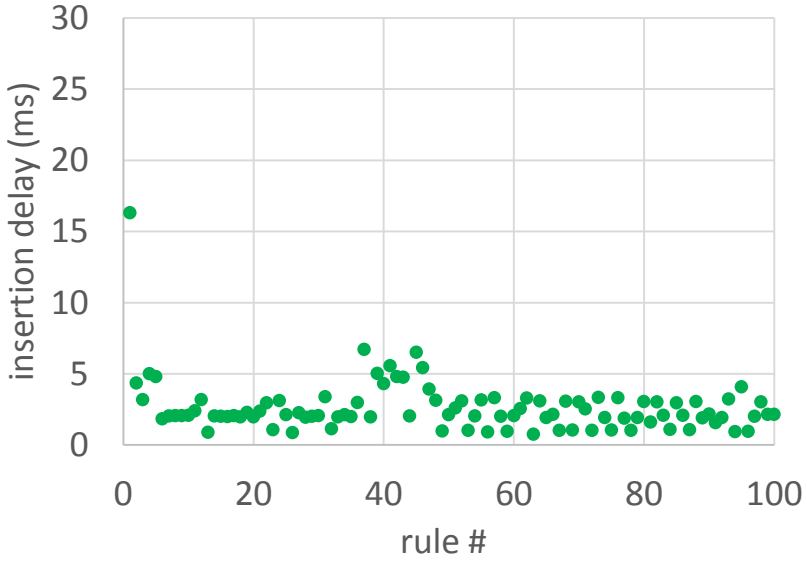


(a) Burst size 100, **same** priority

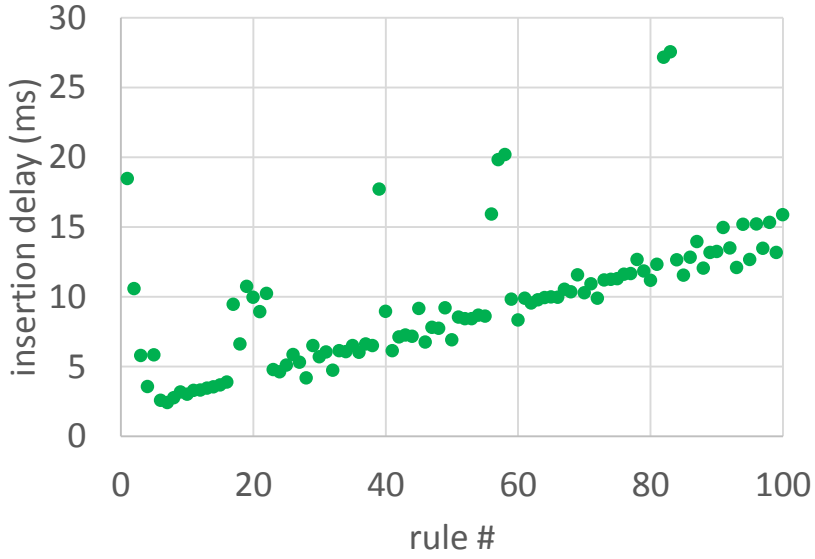


(b) Burst size 100, **increasing** priority

**Vendor B** switch

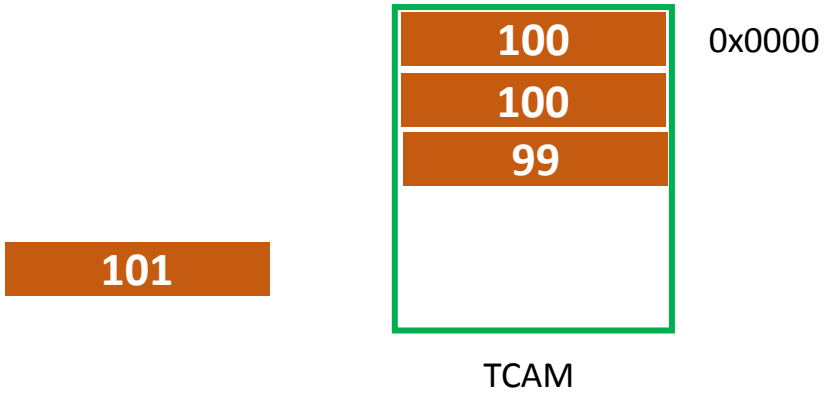

# Insertion Latency – Priority Effects

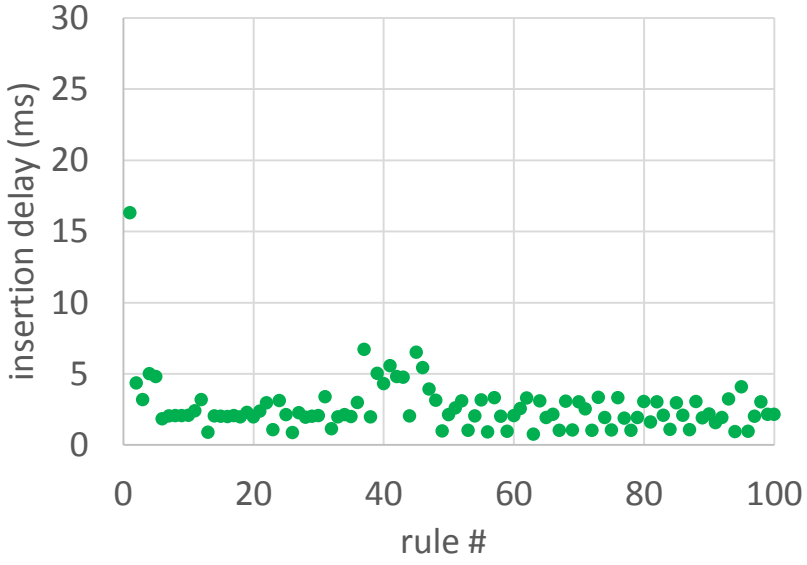


(a) Burst size 100, **same** priority

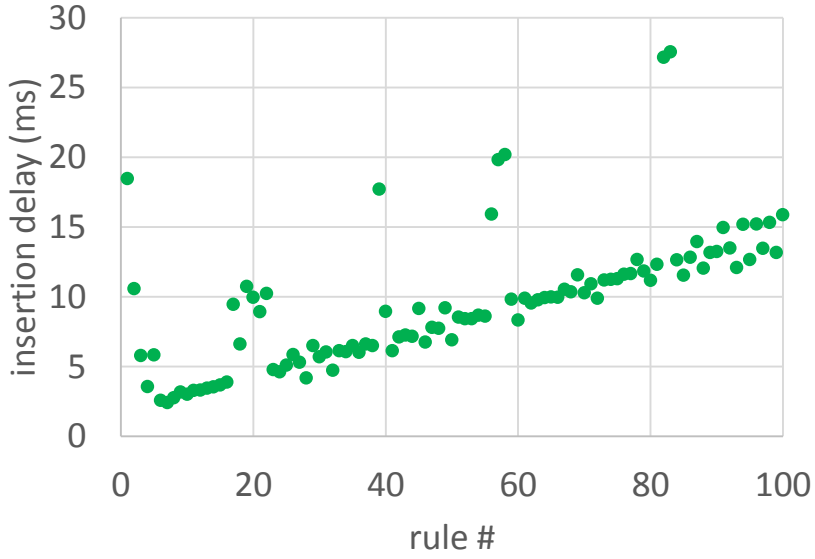


(b) Burst size 100, **increasing** priority

**Vendor B** switch

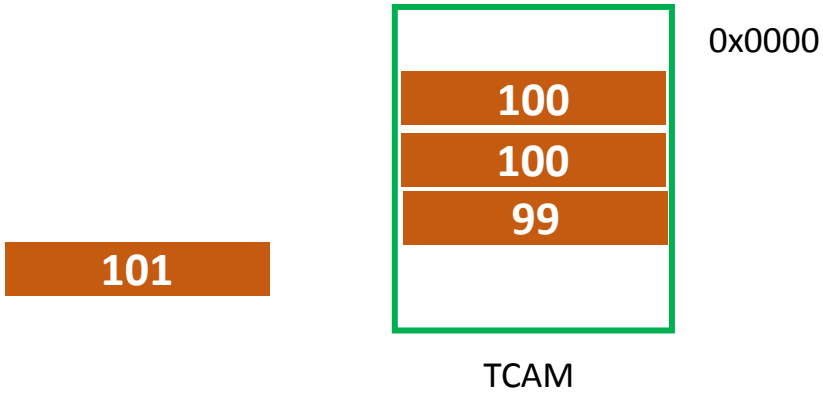

# Insertion Latency – Priority Effects

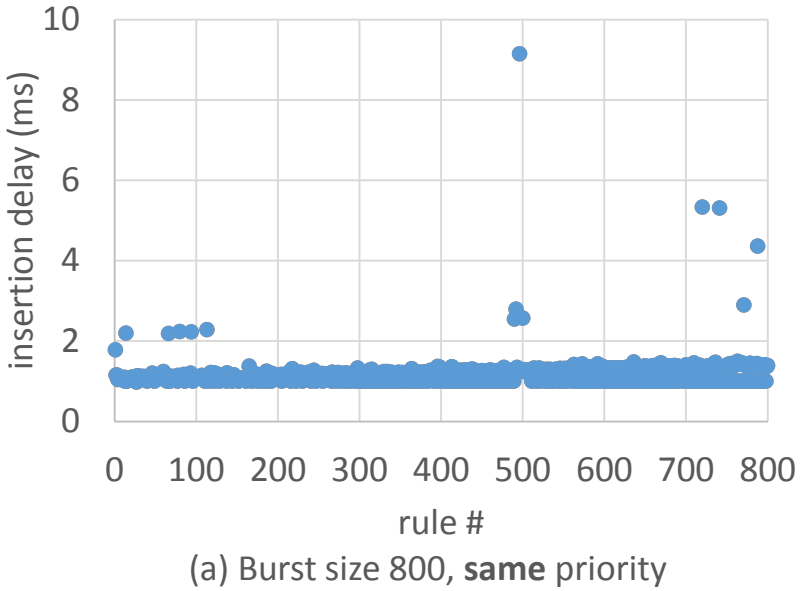


(a) Burst size 800, **same** priority

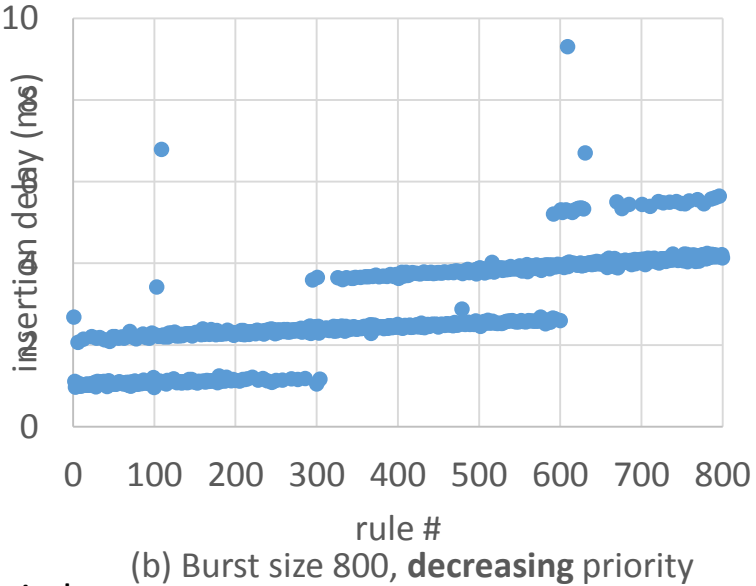


(b) Burst size 800, **decreasing** priority

**Vendor A** switch

# Insertion Latency – Priority Effects

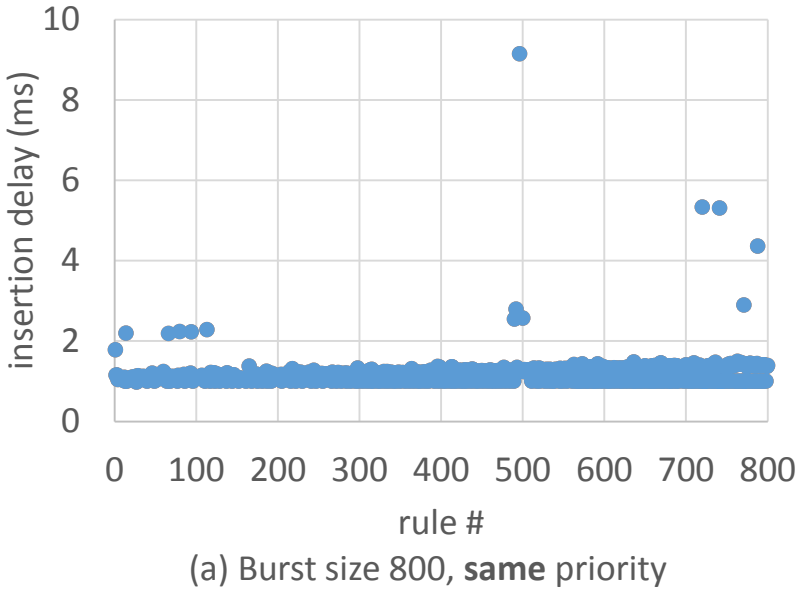


(a) Burst size 800, **same** priority

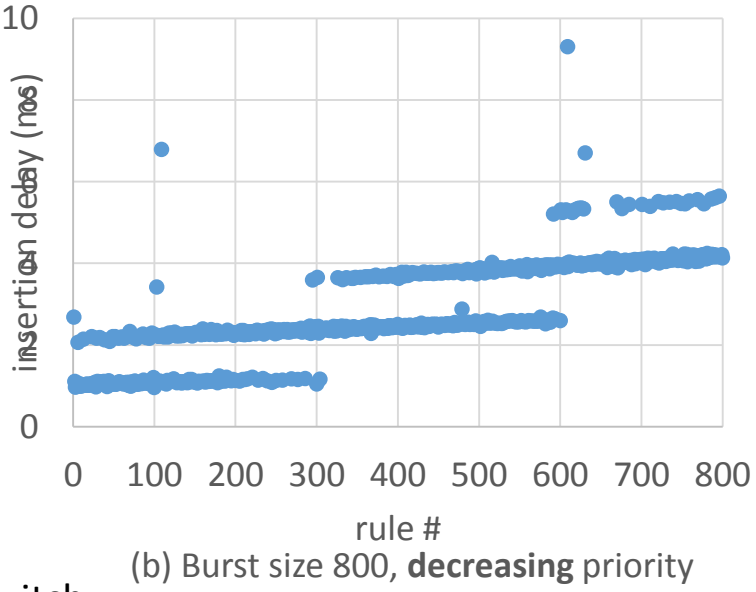


(b) Burst size 800, **decreasing** priority

**Vendor A** switch

# Insertion Latency – Priority Effects

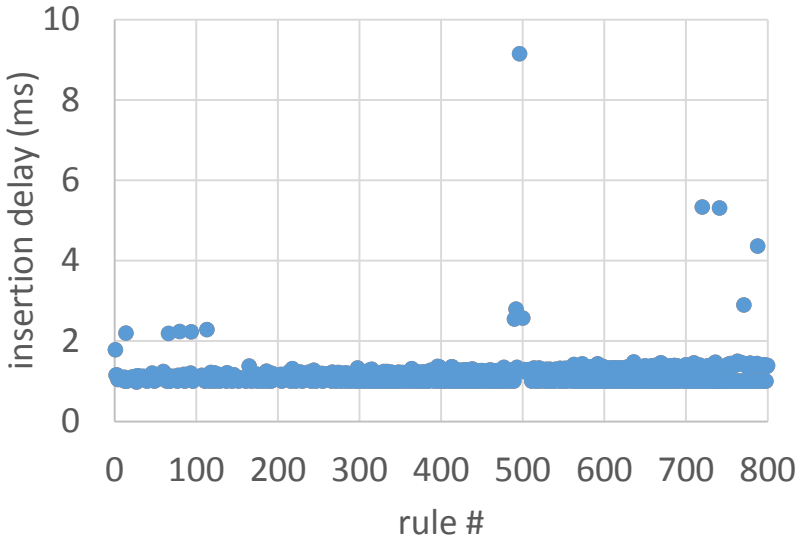


(a) Burst size 800, **same** priority

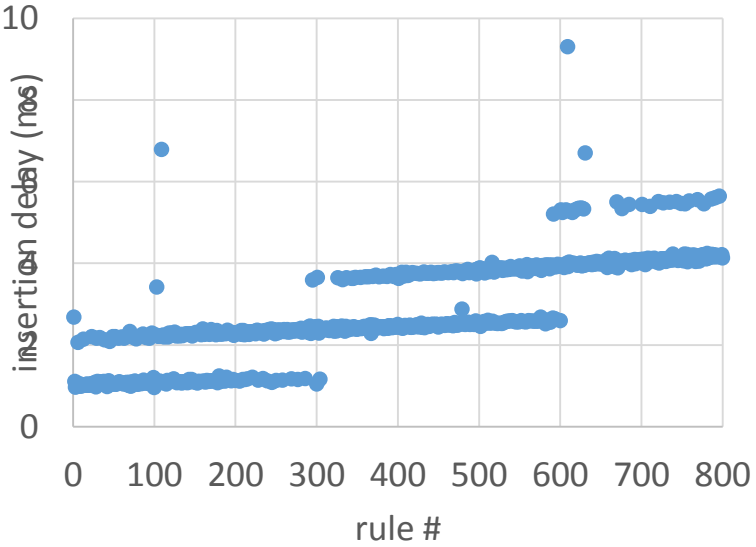


(b) Burst size 800, **decreasing** priority

**Vendor A** switch

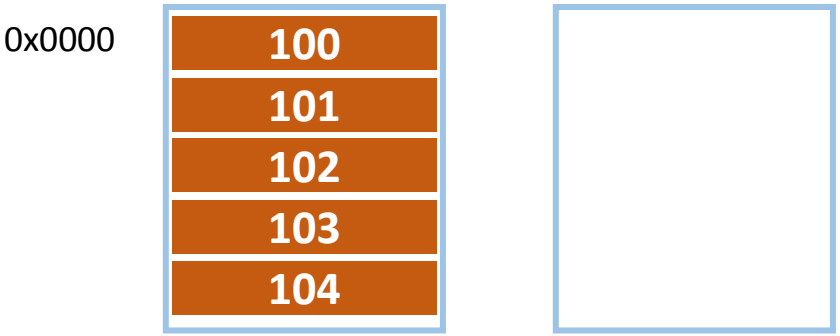

# Insertion Latency – Priority Effects

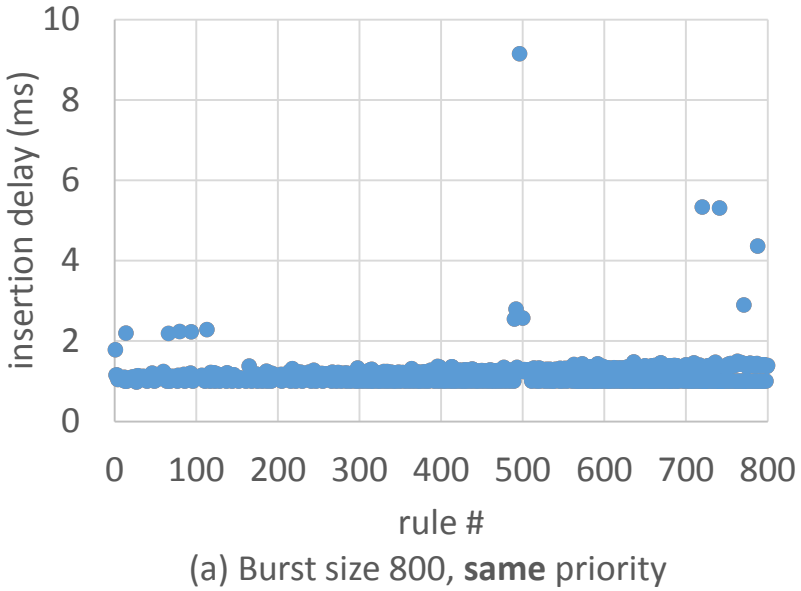


(a) Burst size 800, **same** priority

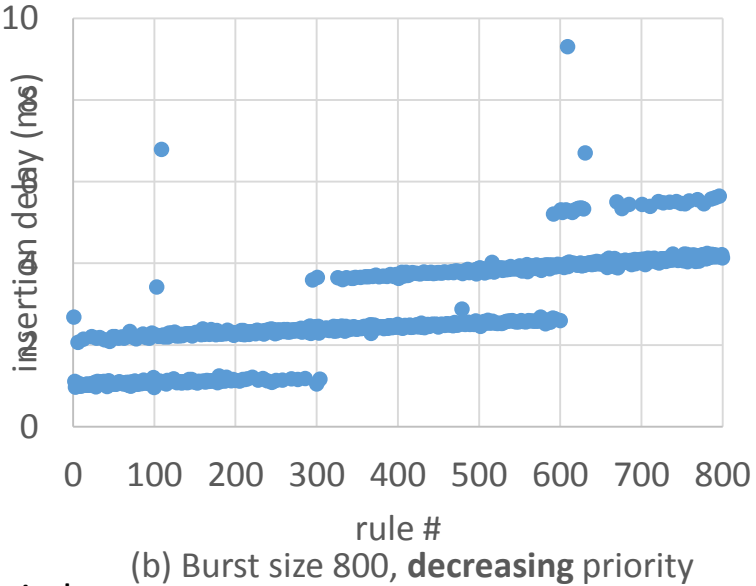


(b) Burst size 800, **decreasing** priority

**Vendor A** switch

# Insertion Latency – Priority Effects

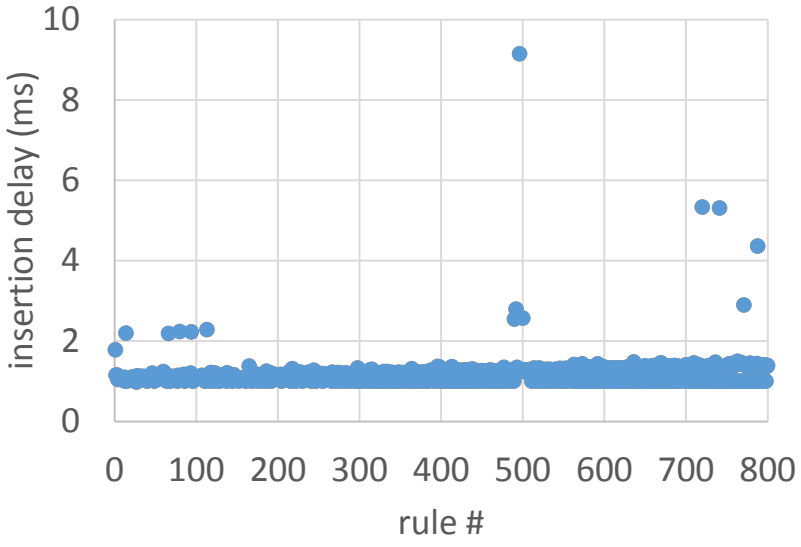


(a) Burst size 800, **same** priority

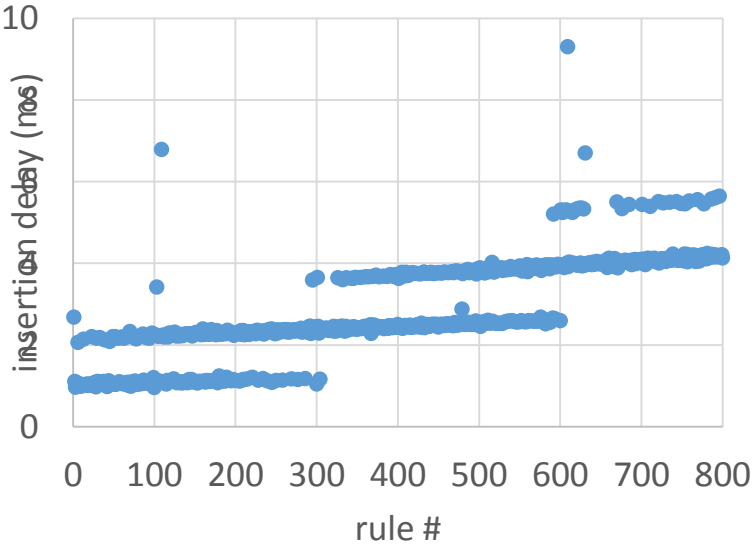


(b) Burst size 800, **decreasing** priority

**Vendor A** switch

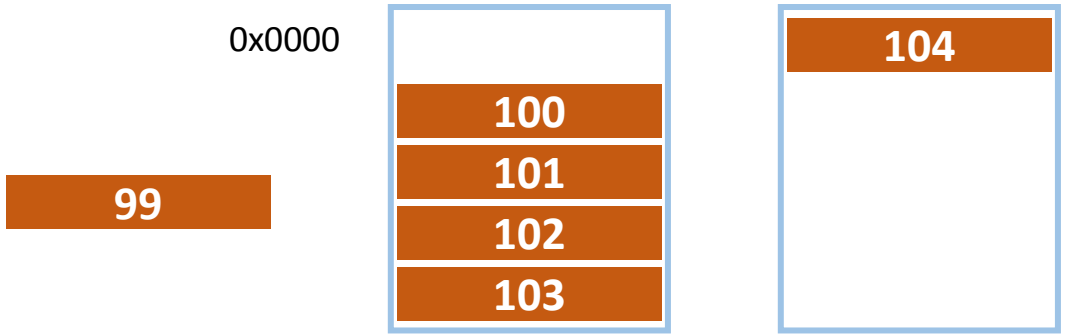

# Insertion Latency – Priority Effects

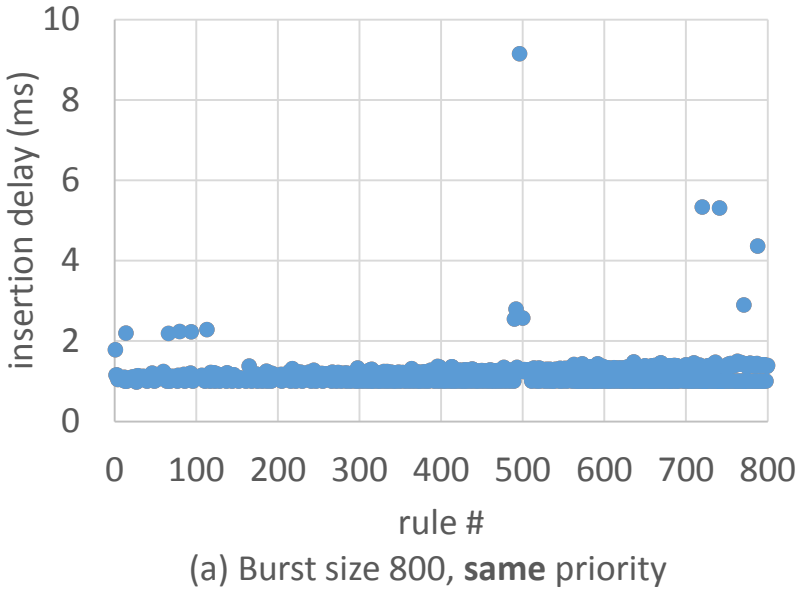


(a) Burst size 800, **same** priority

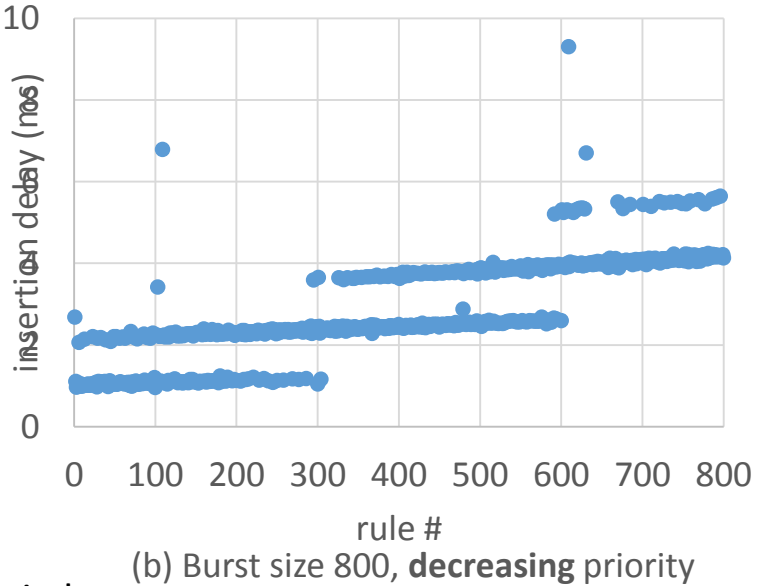


(b) Burst size 800, **decreasing** priority

**Vendor A** switch

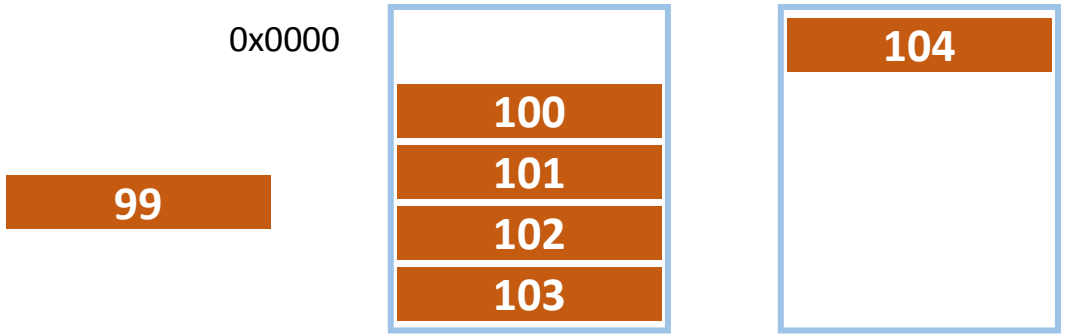

# Insertion Latency – Table occupancy Effects

Rule Priority Structure and Table Occupancy

TCAM Organization and # Of Slices

Switch Software Overhead

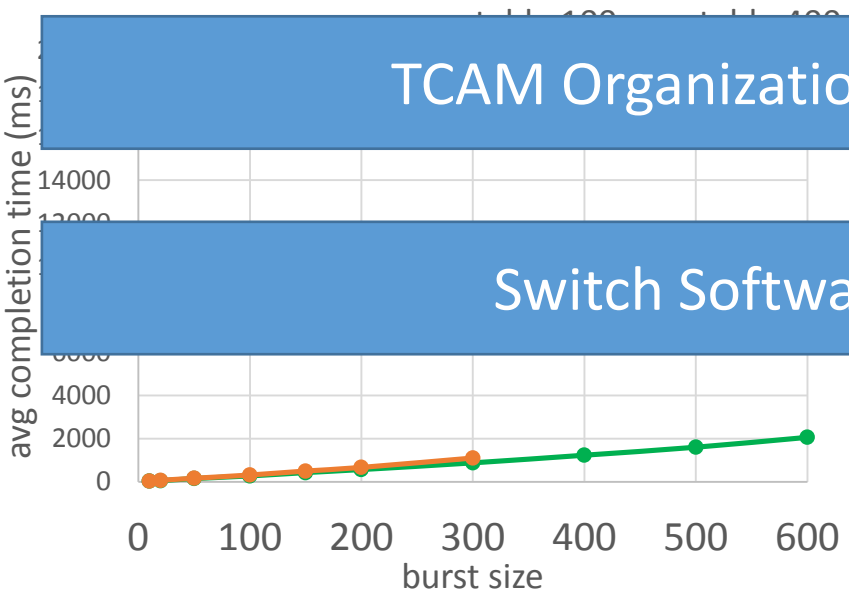


(a) low priority rules into a table with **high priority** existing rules

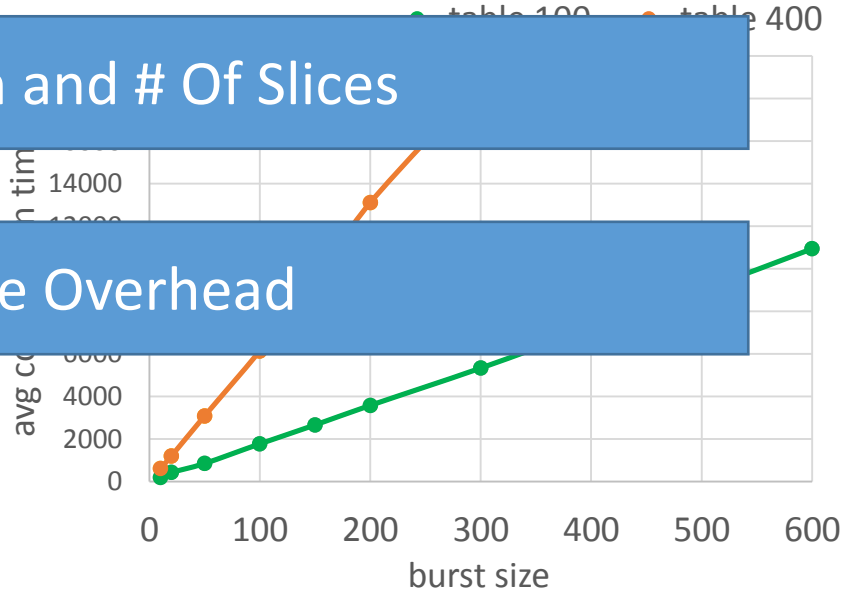


(b) high priority rules into a table with **low priority** existing rules

**Vendor B** switch

# Outline

- Motivation
- Elements of Latency
- Measurement Methodology
- Inbound Latency
- **Outbound Latency**
  - Insertion
  - **Modification**
  - Deletion

# Modification Latency

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


(a) 100 rules in the table

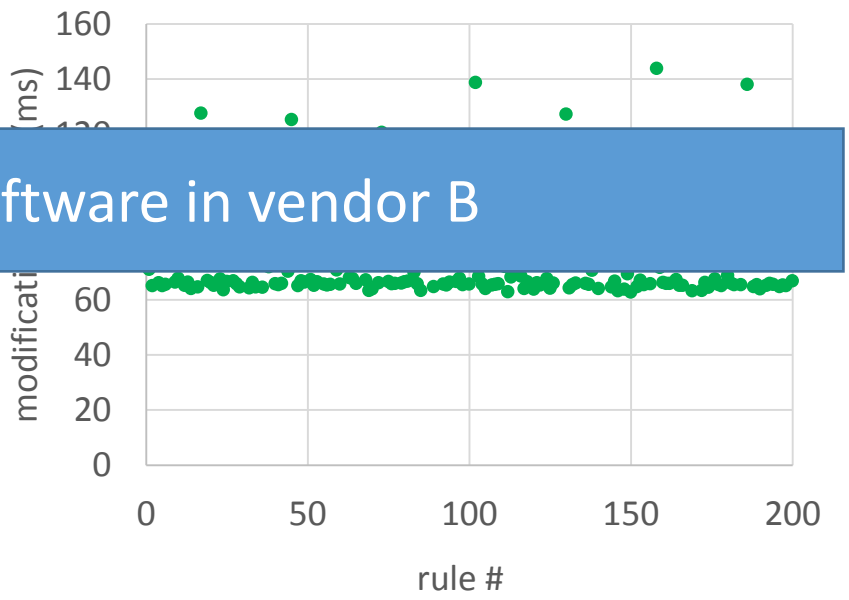


(b) 200 rules in table

Poorly optimized software in vendor B

**Vendor B** switch

# Outline

- Motivation
- Elements of Latency
- Measurement Methodology
- Inbound Latency
- **Outbound Latency**
  - Insertion
  - Modification
  - **Deletion**

# Deletion Latency

- Higher than Insertion latency for both vendor A and B
- Not affected by rule priority but affected by table occupancy

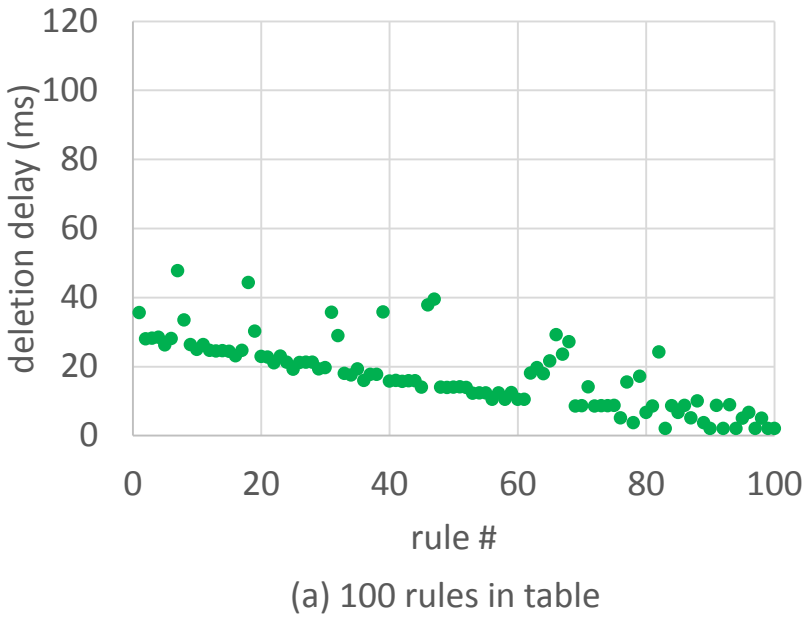


(a) 100 rules in table

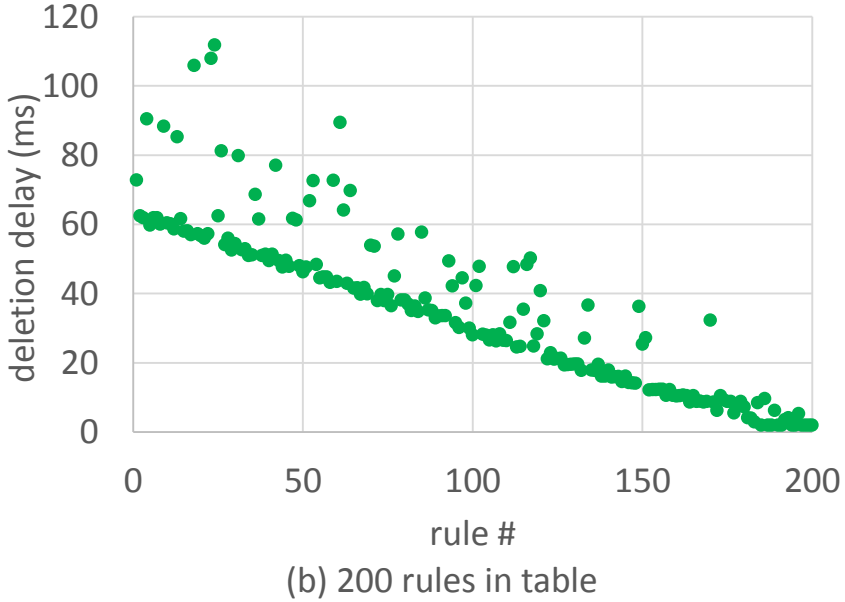


(b) 200 rules in table

**Vendor B** switch

# Deletion Latency

- Higher than Insertion latency for both vendor A and B
- Not affected by rule priority but affected by table occupancy

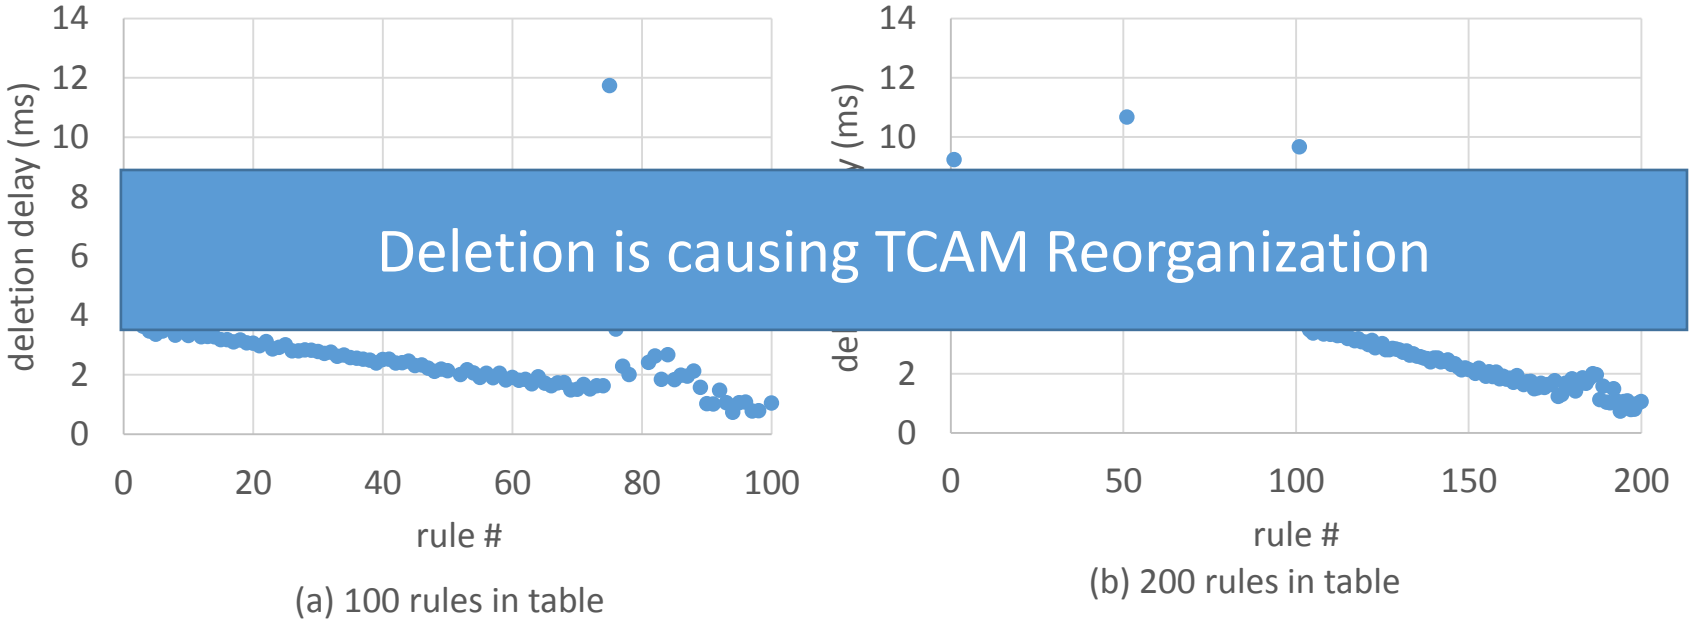


**Vendor A** switch

# Summary

- Latency in SDN is *critical* to many applications
- ~~Assumption: *Latency is small* or *constant*~~
- Latency is high and variable
- Varies with Platforms, Type of operations, Rule priorities, Table occupancy, Concurrent operations
- Key Factors: *TCAM Organization, Switch CPU* and *inefficient Software Implementation*

# Summary

- Latency in SDN is *critical* to many applications
- ~~Assumption: *Latency is small* or *constant*~~
- Latency is high and variable
- Varies with Platforms, Type of operations, Rule priorities, Table occupancy, Concurrent operations
- Key Factors: *TCAM Organization, Switch CPU* and *inefficient Software Implementation*

# Backup Slides

# Impact of Concurrent CPU jobs

- Impacts insertion delay. E.g. *Polling Statistics*
- Polling stats impacts more when table occupancy is higher

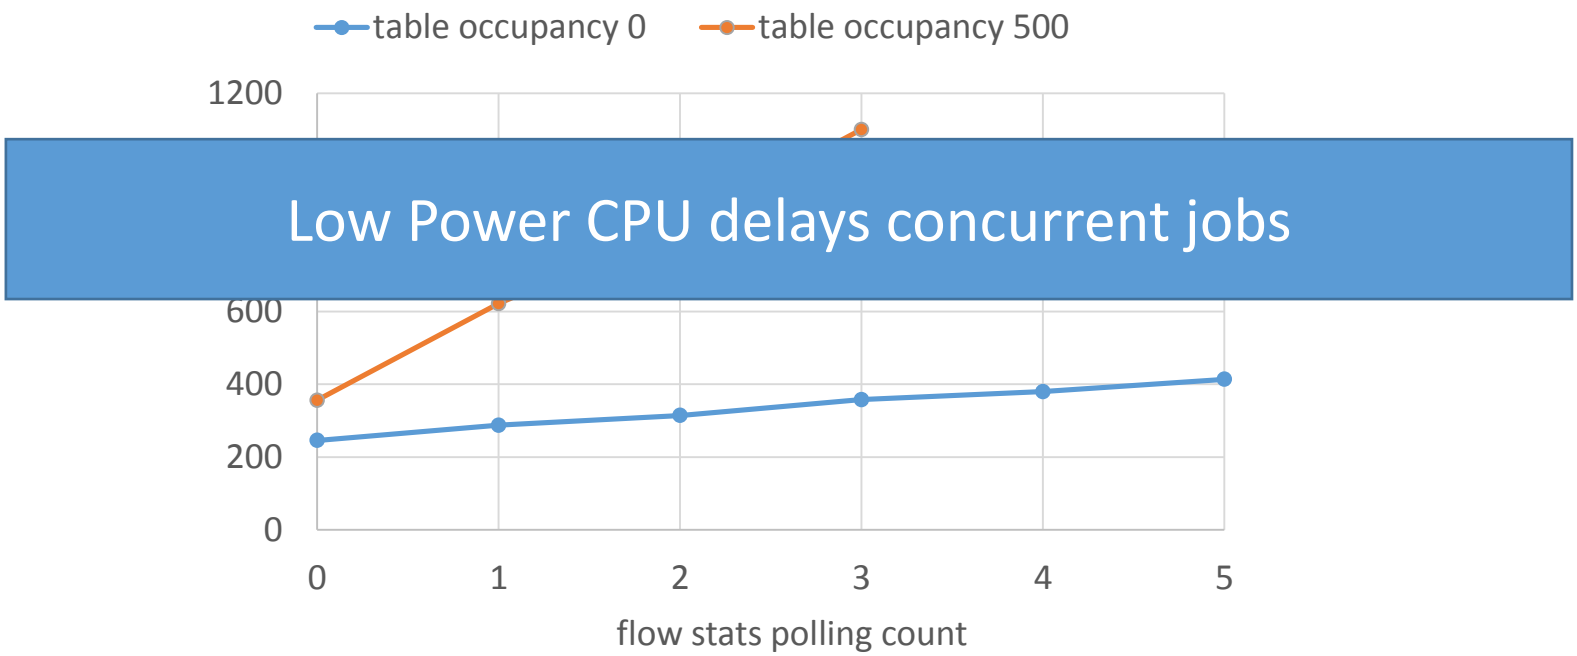


Vendor B