SDN OpenFlow WORLD CONGRESS

EWSDN
European Workshop on
Software Defined Networks

A Host-based Performance Comparison of 40G NFV Environments Focusing on Packet Processing Architectures and Virtual Switches

R. Kawashima, S. Muramatsu, H. Nakayama
T. Hayashi, and H. Matsuo

Nagoya Institute of Technology
BOSCO Technologies

Contents

Network Functions Virtualization (NFV)

Traditional

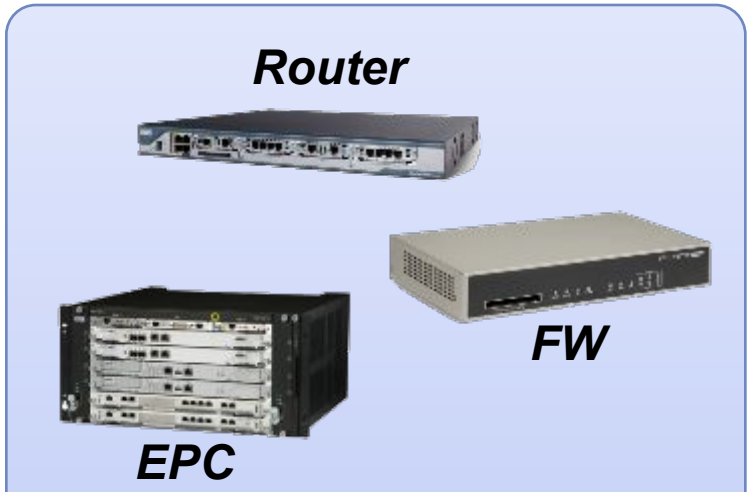

HW-based

- High cost
- Low flexibility
- Closed

NFV

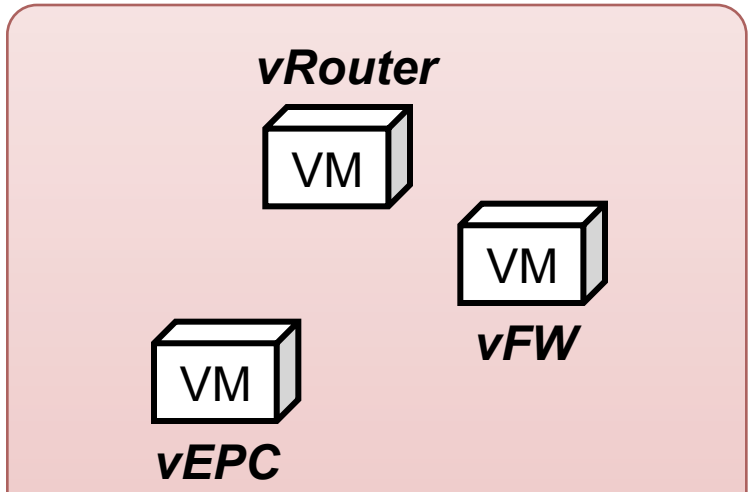

SW-based

- Low cost
- High flexibility
- Open

Host Architectures in NFV

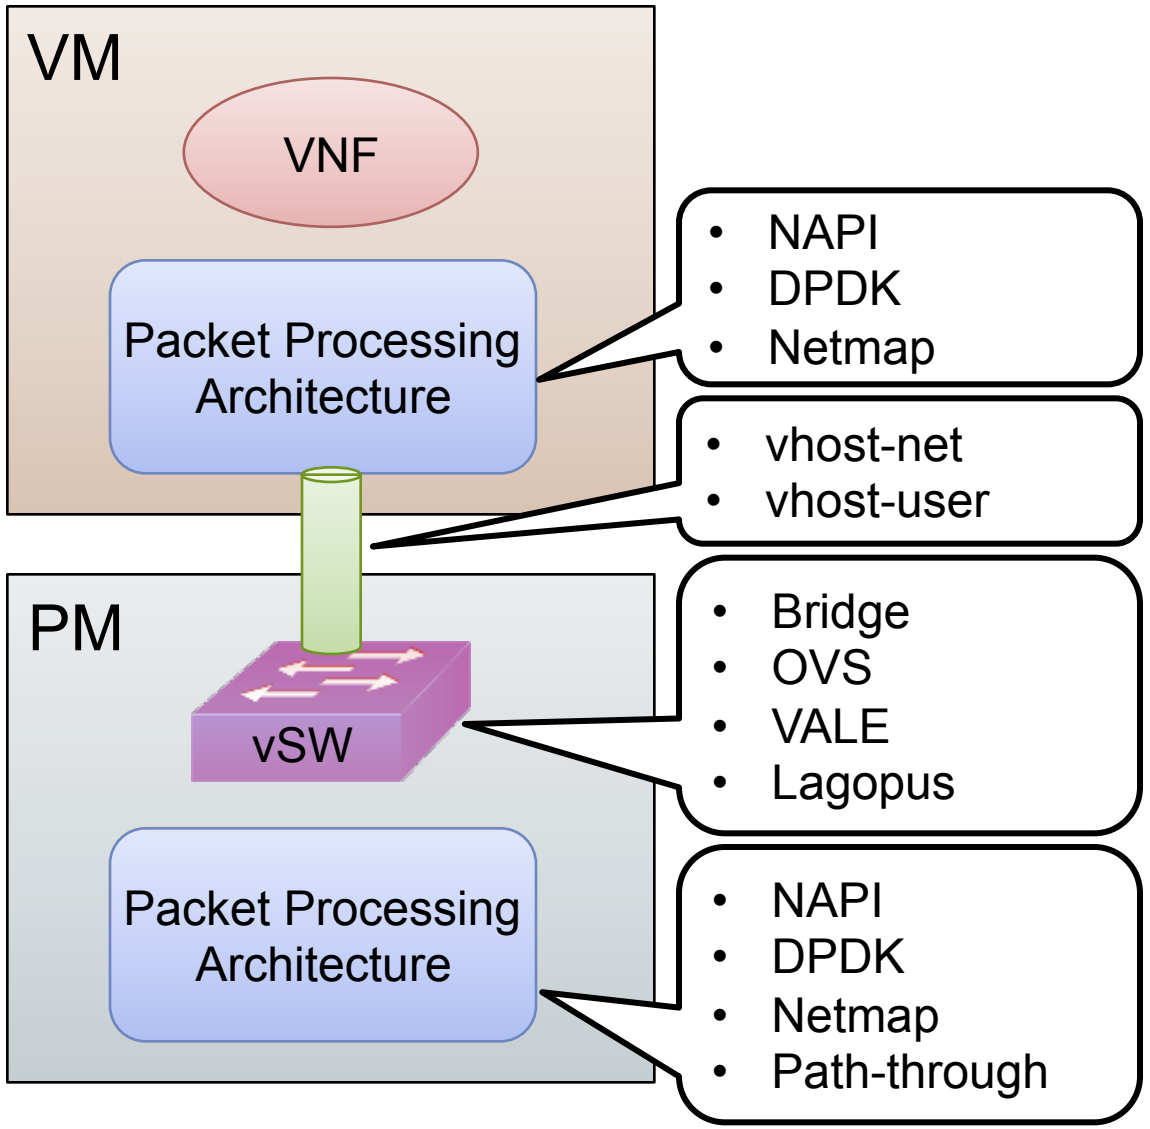

Which architecture or vswitch should be used for NFV hosts ?

How much performance differs depending on its architecture and vswitch ?

Related Work

- **Performance of baremetal servers [1][2]**
  - They focused on performance bottlenecks
- **Performance of KVM and Container-based virtualization [3]**

[1] P. Emmerich et al., "Assessing Soft- and Hardware Bottlenecks in PC-based Packet Forwarding Systems", Proc. ICN, pp. 78-83, 2015.

[2] S. Gallenmüller et al., "Comparison of Frameworks for High-Performance Packet IO", Proc. ANCS, pp. 29-38, 2015.

[3] R. Bonafiglia et al., "Assessing the Performance of Virtualization Technologies for NFV: a Preliminary Benchmarking", Proc. EWSDN, pp. 67-72, 2015.

Our Evaluation

Combination of PM and VM architectures and vswitches

Throughput and Latency/Jitter

Intel and Mellanox 40GbE NICs

SR-IOV

Contents

Packet Processing Architecture

- **A way to forward packets (NIC <=> Applications)**
  - Interruption or Polling
  - Single core or Multi cores
  - Kernel space or User space
  - Packet buffer structure
  - Packet buffer management

The architecture has a major effect on the performance !

Three Architectures

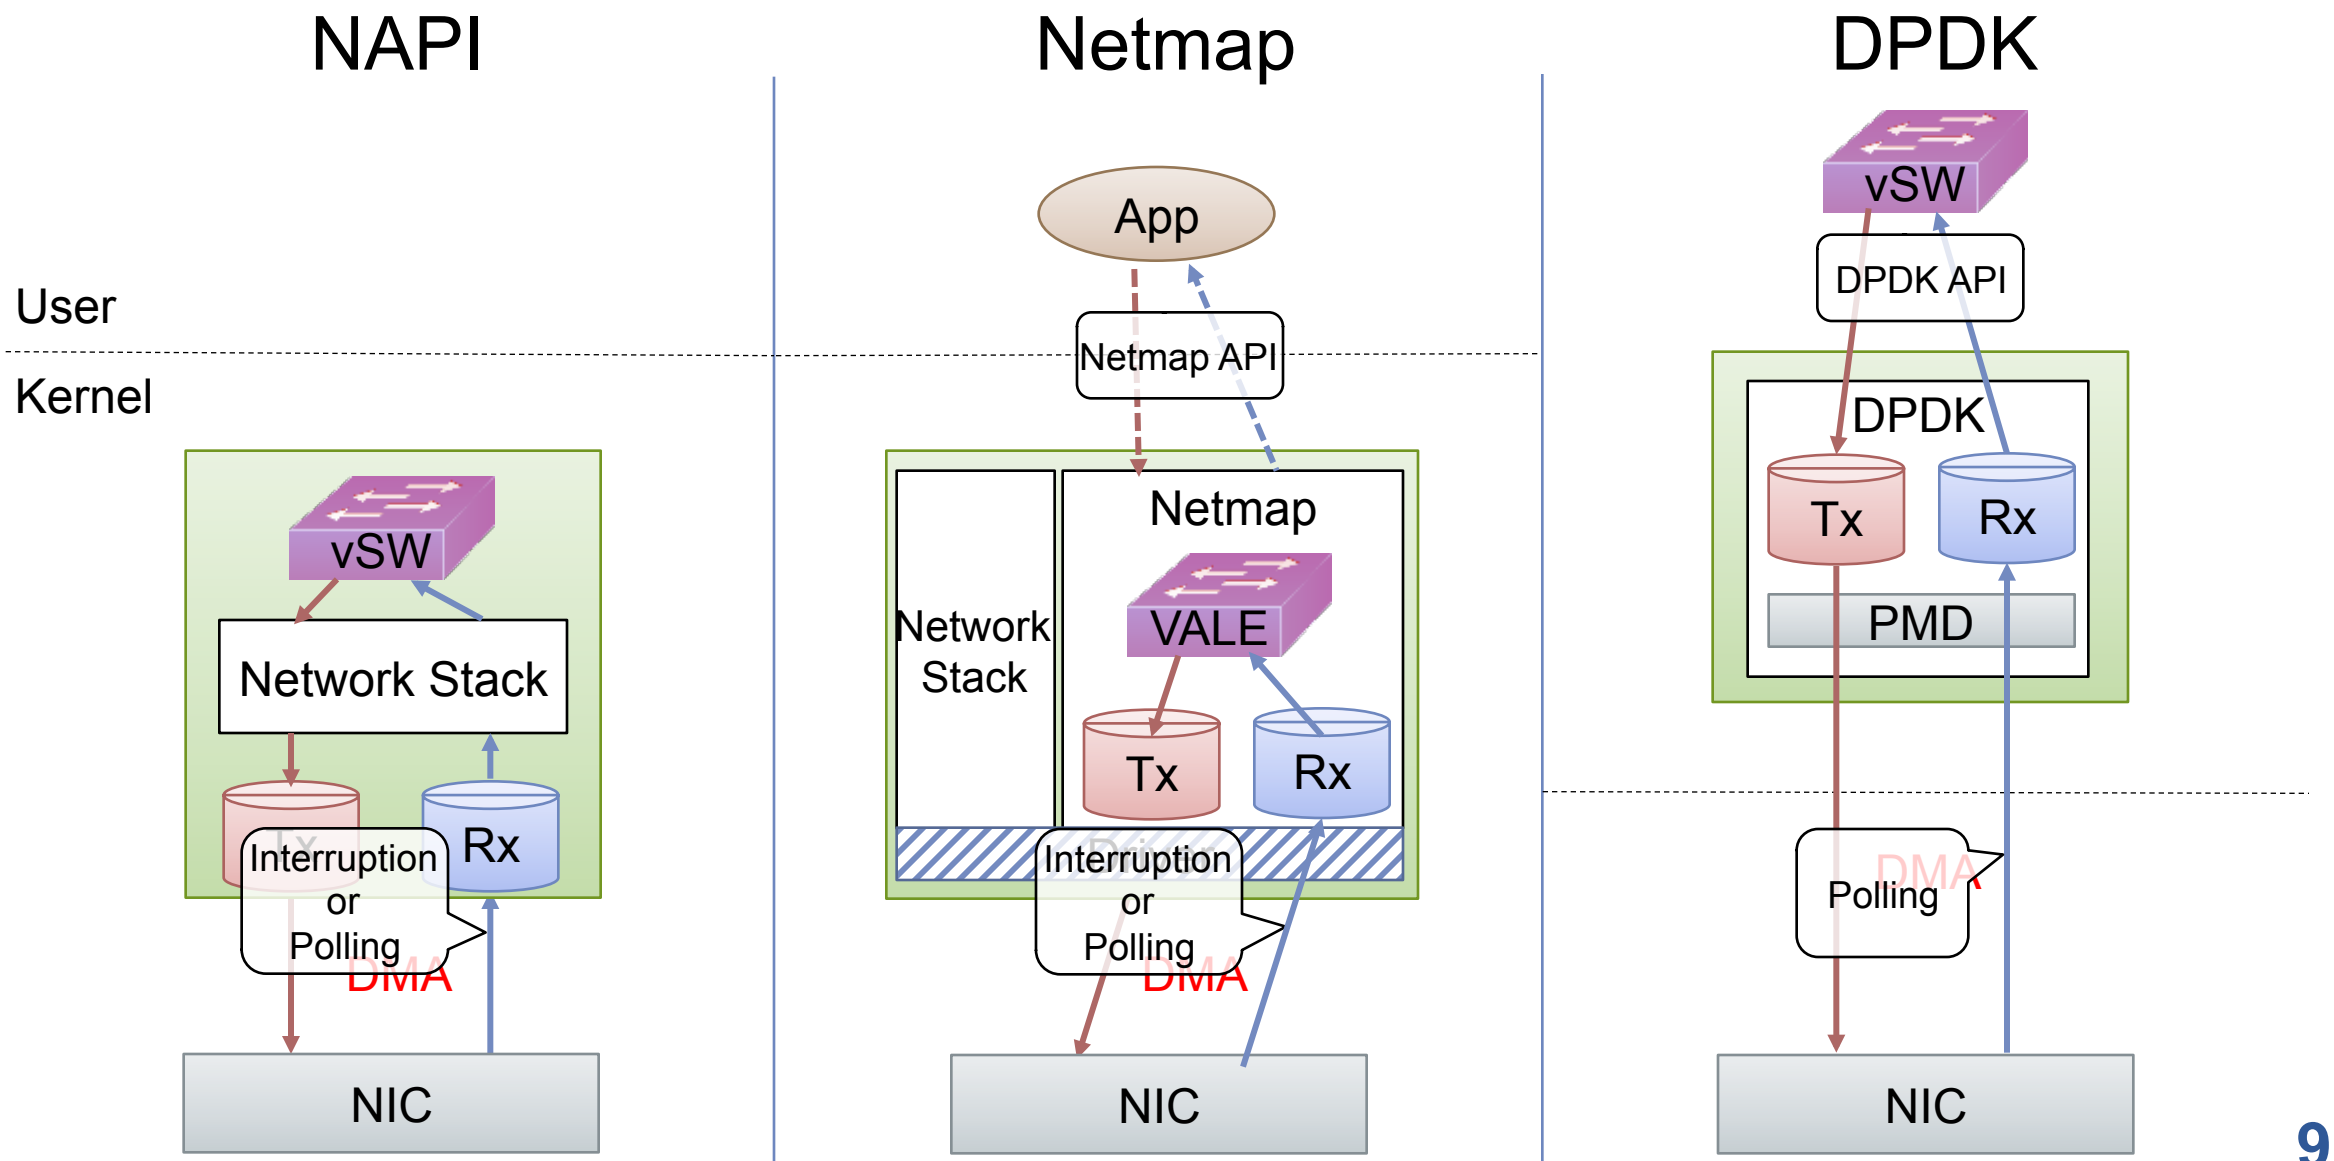

Contents

- **Backgrounds**
  - NFV
  - Related work
- **Packet Processing Architectures**
- **Virtual Switches**
- **Evaluation**
- **Conclusion**

Virtual Switch

- **A virtual switch bridges the host and VMs**

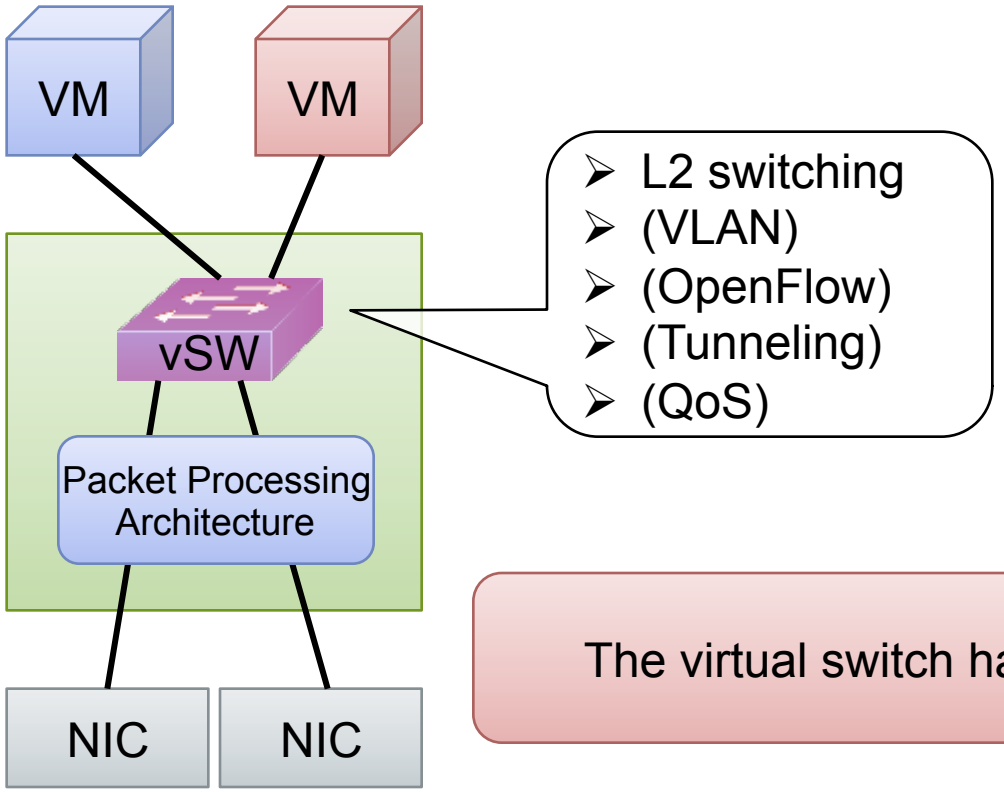

The virtual switch has an impact on performance too !

Six Virtual Switches

Name	Running Space	Architecture	Virtual I/O Support
Linux Bridge	Kernel	NAPI	TAP/vhost-net
OVS	Kernel	NAPI	TAP/vhost-net
VALE	Kernel	Netmap	(QEMU)
L2FWD-DPDK	User	DPDK	-
OVS-DPDK	User	DPDK	vhost-user
Lagopus	User	DPDK	(to be supported)

Contents

Goals

Clarify performance characteristics of existing systems

Propose appropriate NFV host environments

Find a proper direction for performance improvement

Experiment 1 (Baremetal)

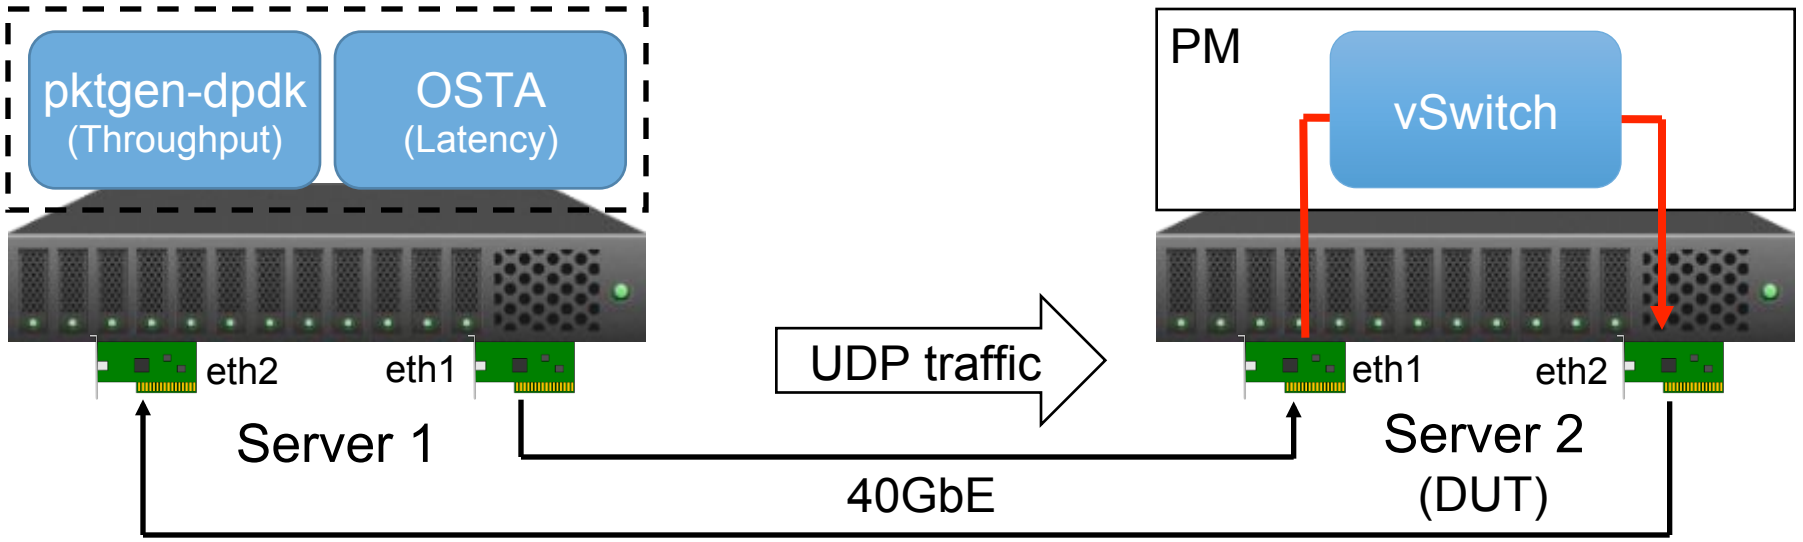

Physical	Server 1	Server 2 (DUT)
OS	CentOS 7.2	
CPU	Core i7-3770 3.40 GHz (4 cores with HT)	Core i7-6700K 4.00 GHz (4 cores with HT)
Memory	64 GB	32 GB
VMM	-	KVM
NIC	Intel Ethernet Converged Network Adapters XL710 Mellanox ConnectX-3 EN	

Throughput

Intel XL710

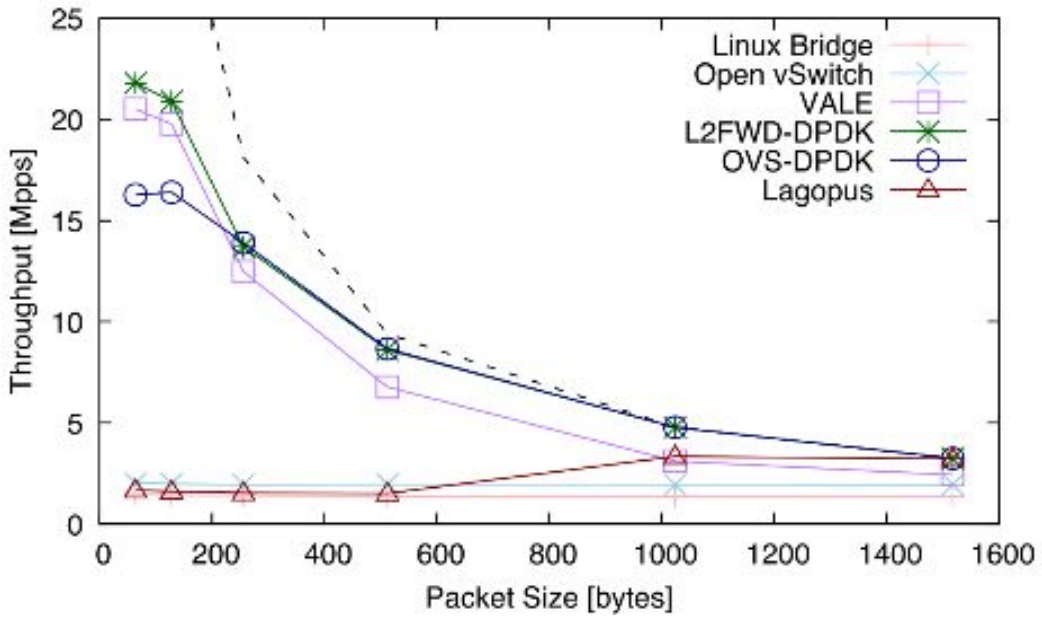

Mellanox
ConnectX-3 EN

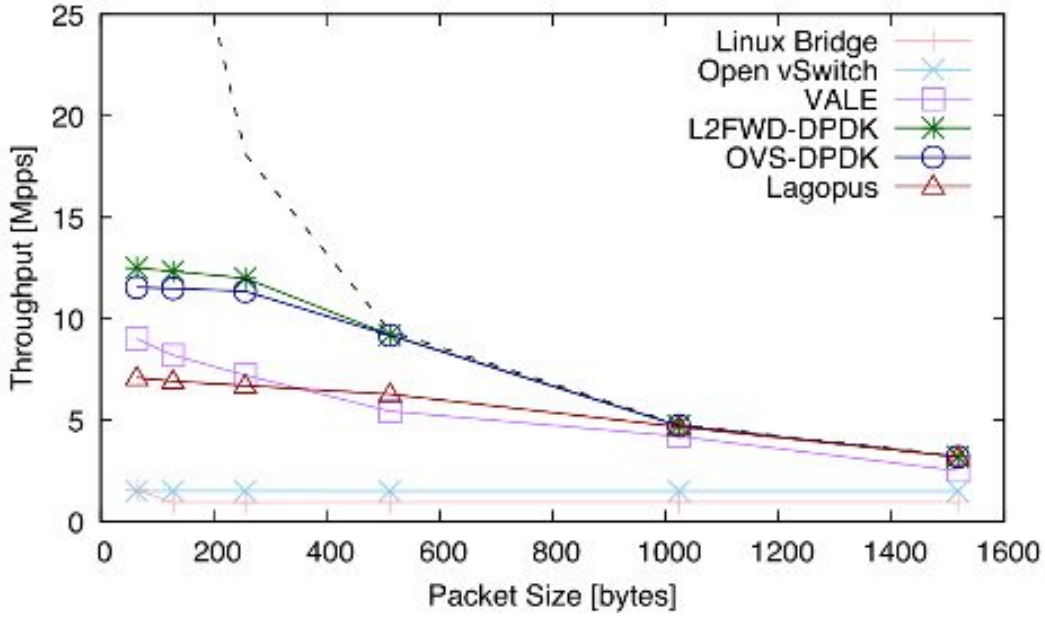

Throughputs differ depending on the NIC type

Throughputs with short packet sizes are far from wire rate

Latency

Intel XL710

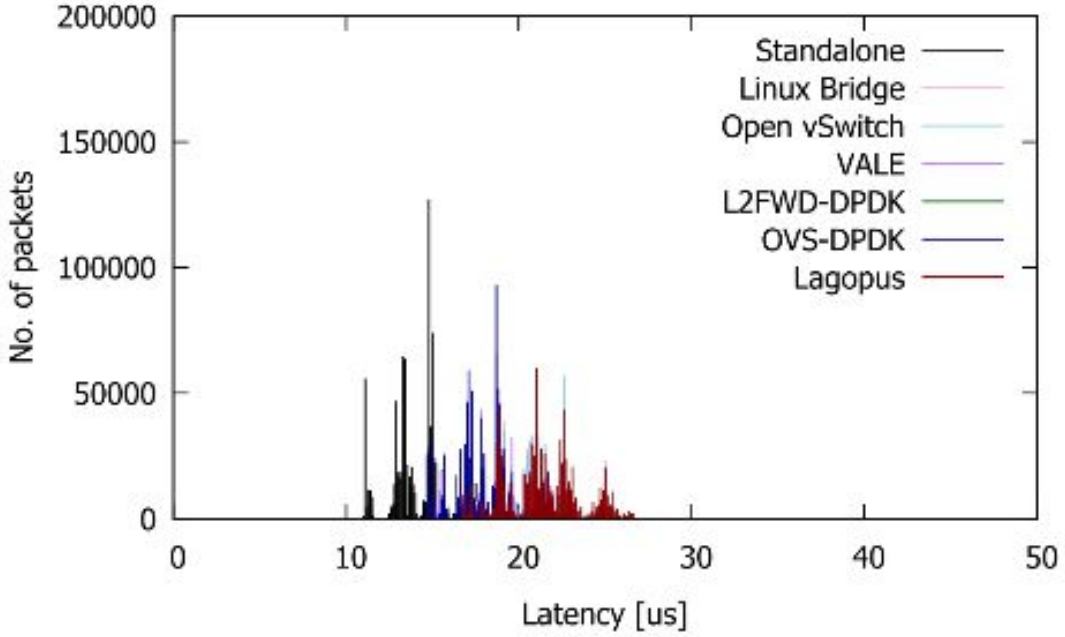

Mellanox ConnectX-3 EN

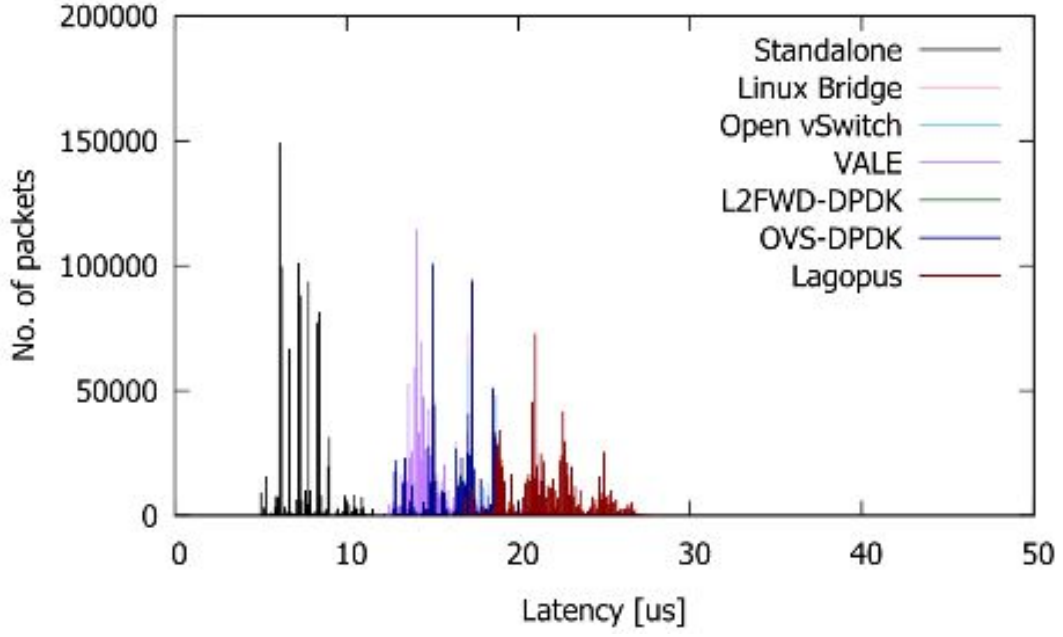

L2FWD-DPDK and Lagopus show worse latency

Jitter values are less than 10 µs

Experiment 2 (VM)

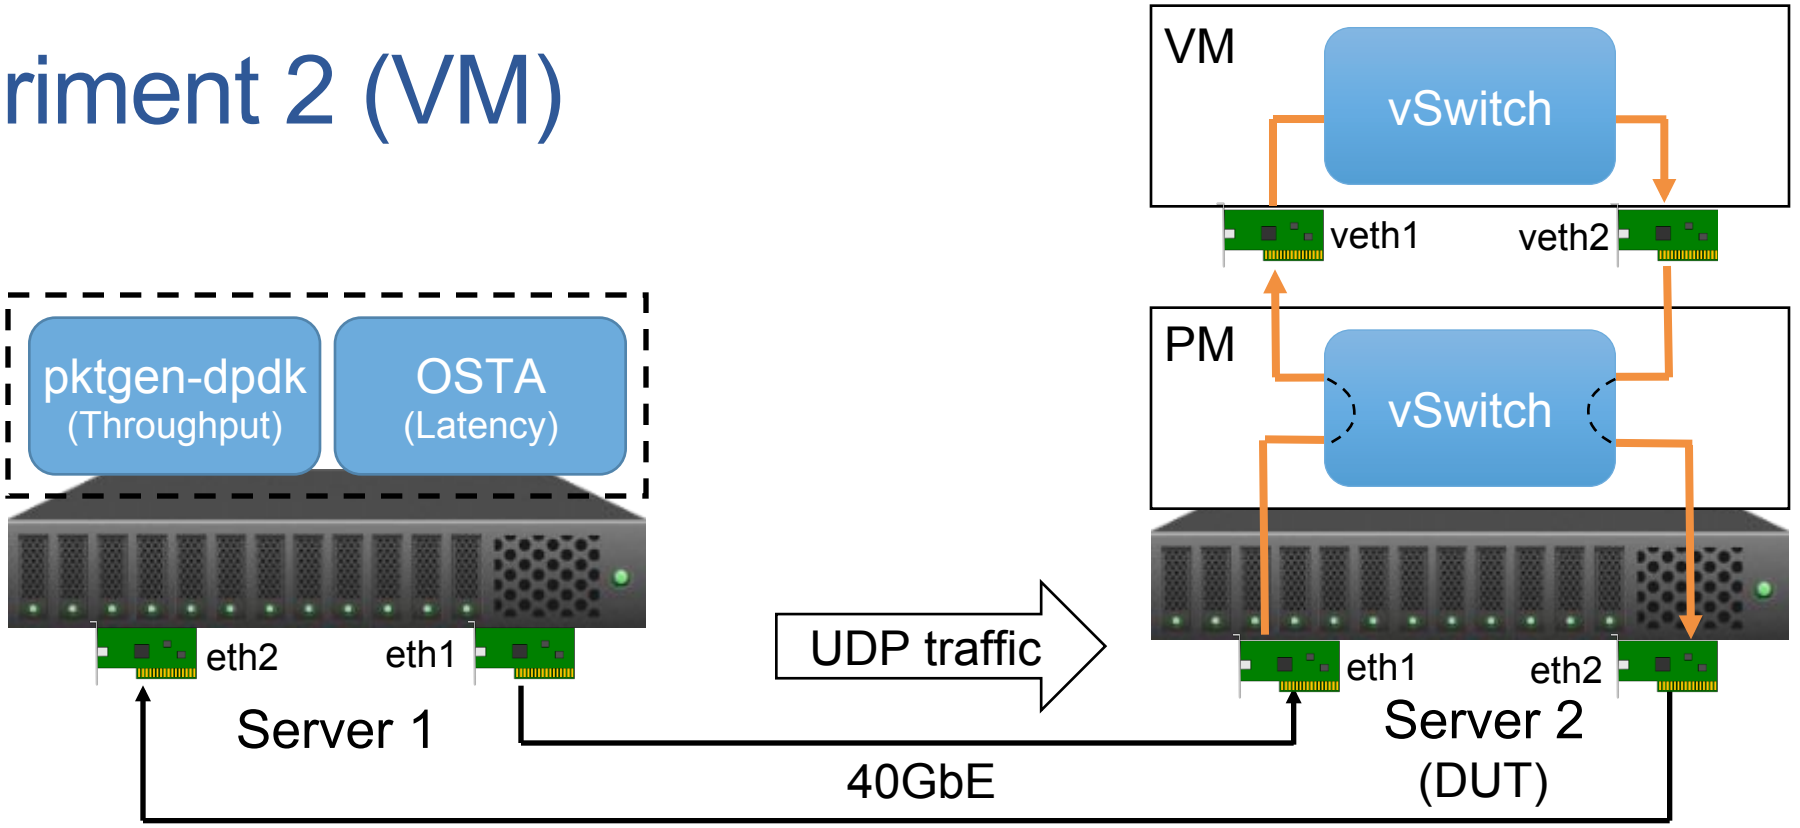

Physical	Server 1	Server 2 (DUT)
OS	CentOS 7.2	
CPU	Core i7-3770 3.40 GHz (4 cores with HT)	Core i7-6700K 4.00 GHz (4 cores with HT)
Memory	64 GB	32 GB
VMM	-	KVM
NIC	Intel Ethernet Converged Network Adapters XL710 Mellanox ConnectX-3 EN	

Virtual	VM (Experiment 2)
OS	CentOS 7.2
vCPU	2 cores
Memory	4 GB
vNIC	virtio-net

Throughput

Intel XL710

Mellanox ConnectX-3 EN

NAPI/ vhost-net

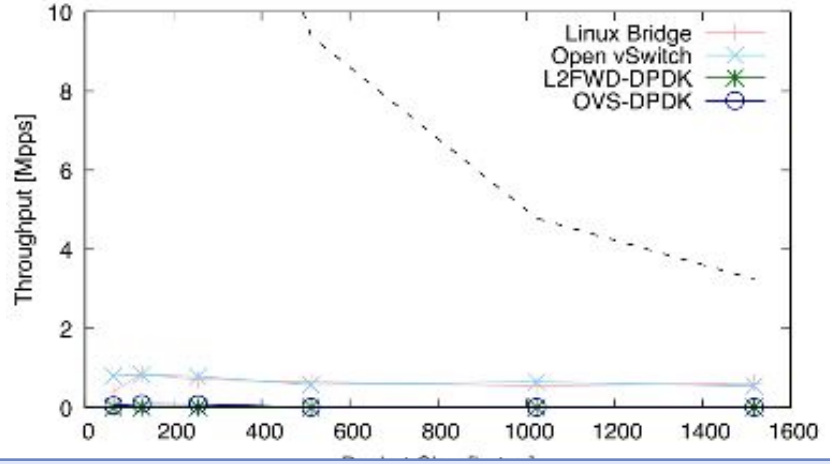

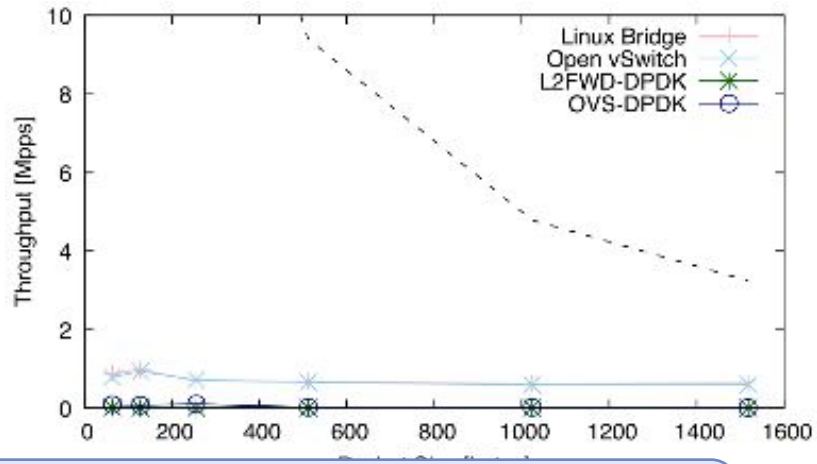

The virtualization overhead is fairly large

DPDK/ vhost-user

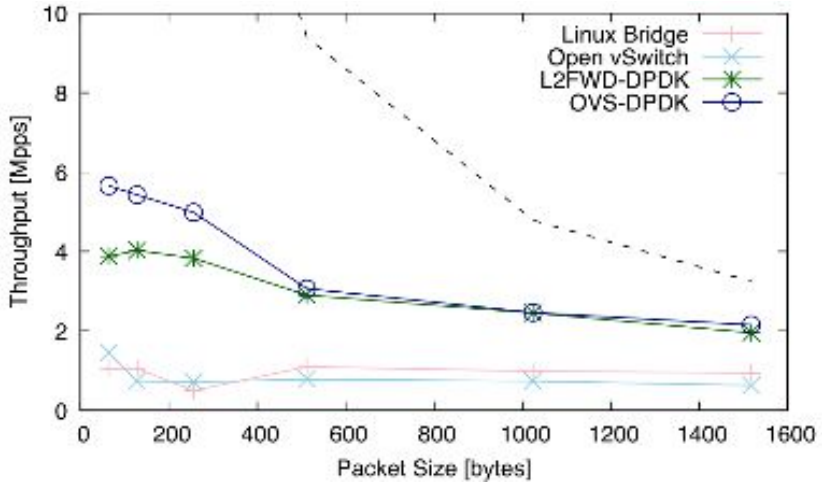

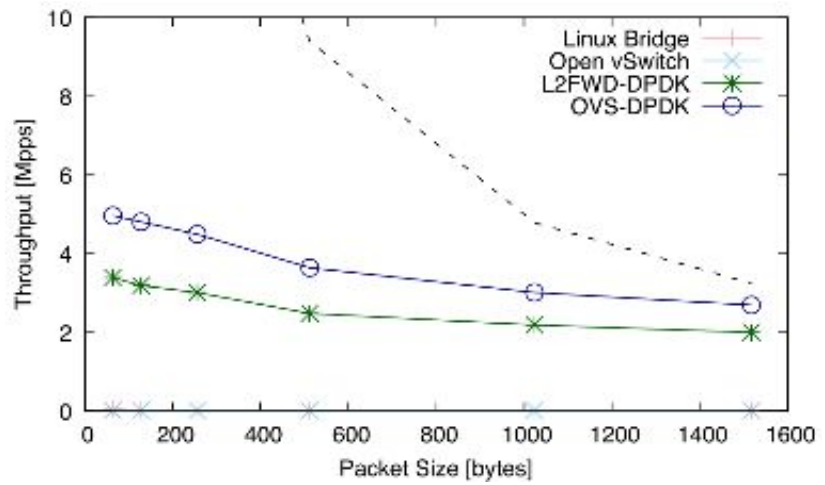

Latency

Intel XL710

Mellanox ConnectX-3 EN

NAPI/ vhost-net

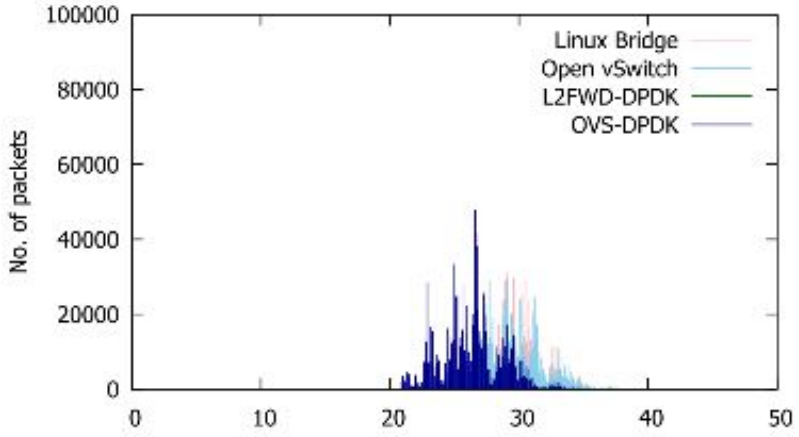

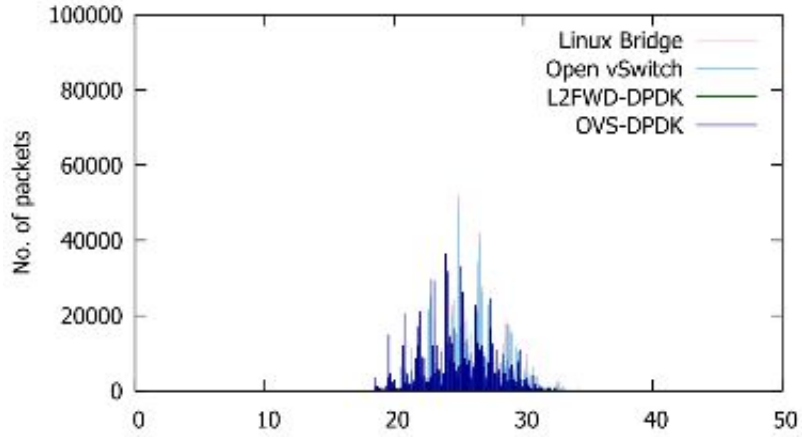

The virtualization amplifies jitters

DPDK/ vhost-user

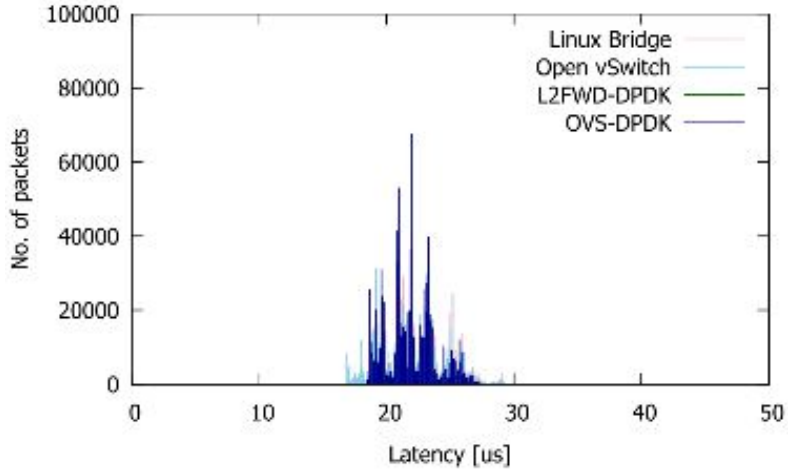

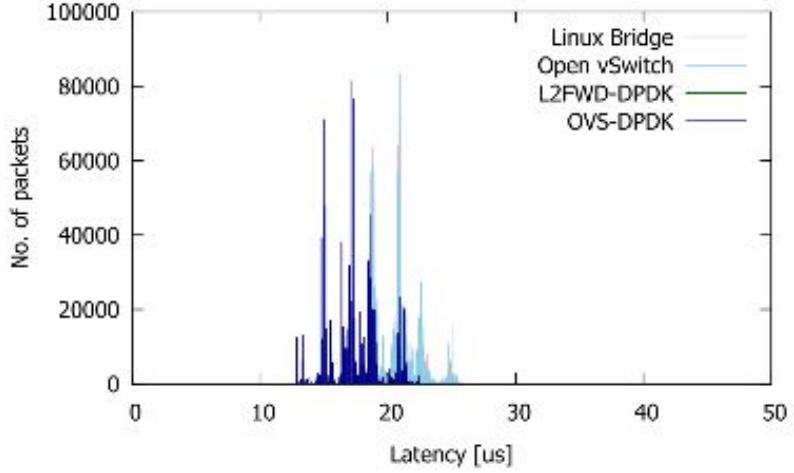

SR-IOV

Throughput

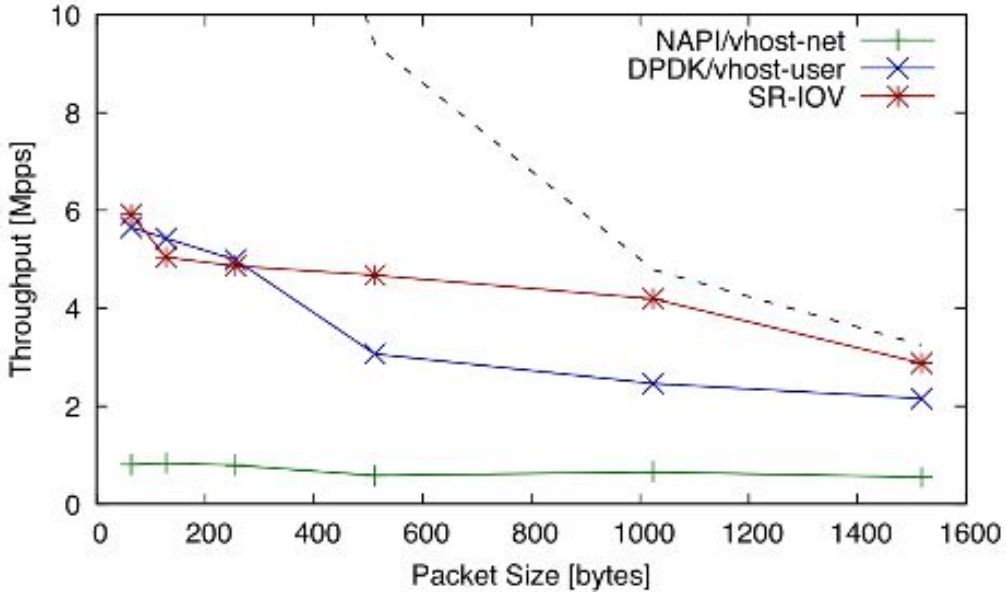

Latency

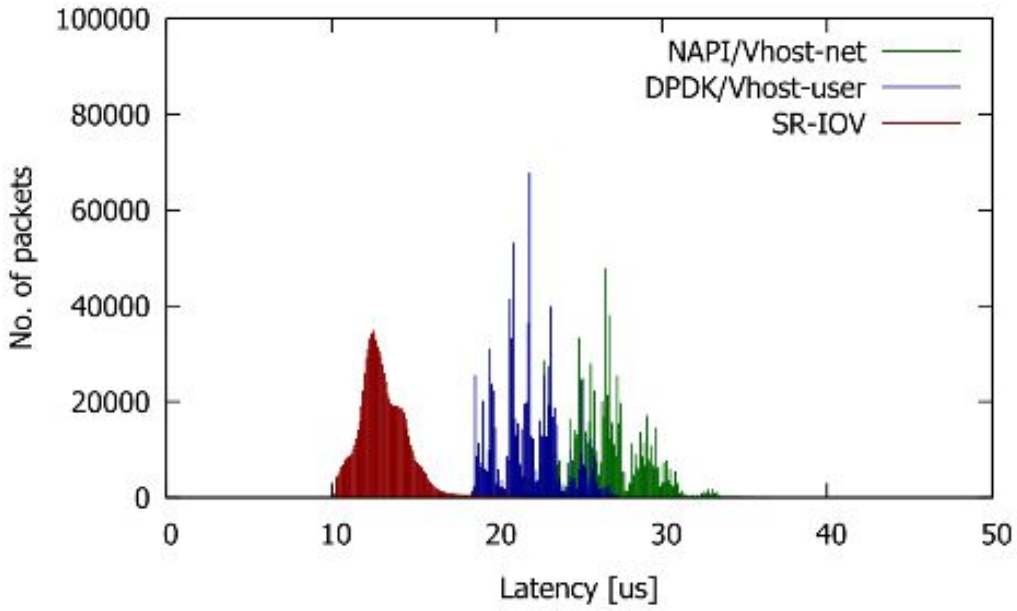

SR-IOV shows the best performance !

SR-IOV lacks flexibility of flow handling

Adequate NFV Host Environment

- Throughput
- Latency/Jitter
- Usability

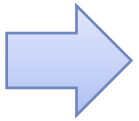

NIC	Architecture/ Virtualization	Virtual Switch
Intel XL710	NAPI/ vhost-net	Linux Bridge
Mellanox ConnectX-3	DPDK/ vhost-user	OVS
		L2FWD-DPDK
		OVS-DPDK

Contents

- **Backgrounds**
  - NFV
  - Related work
- **Packet Processing Architectures**
- **Virtual Switches**
- **Evaluation**
- **Conclusion**

Conclusion

❖ Summary

- We have evaluated NFV host environments with 40GbE
  - A NIC device affects performance characteristics
  - DPDK should be used for both the host and the guest
  - We cannot reach the wire rate with short packet sizes
  - Virtualization worsens both throughput and latency
  - SR-IOV showed better throughput and latency

❖ Future Work

- Further evaluations
  - VALE/Netmap based virtualization
  - VALE and Lagopus on the VM
  - Bidirectional and lots of flows