APNOMS 2017

SDN@NIT
The Power of Softwarization

Finding the Right Way for High-Performance NFV

Ryota Kawashima

Nagoya Institute of Technology, Japan

Agenda

1. Backgrounds
2. D-Plane Technologies
3. Performance
4. Research Directions

Virtualizing Everything

Client Computers

1st-gen virtualization (Before NFV)

Edge Functions

2nd-gen virtualization (Early-stage NFV)

Core Functions

3rd-gen virtualization (Future NFV)

Core Network Traffic

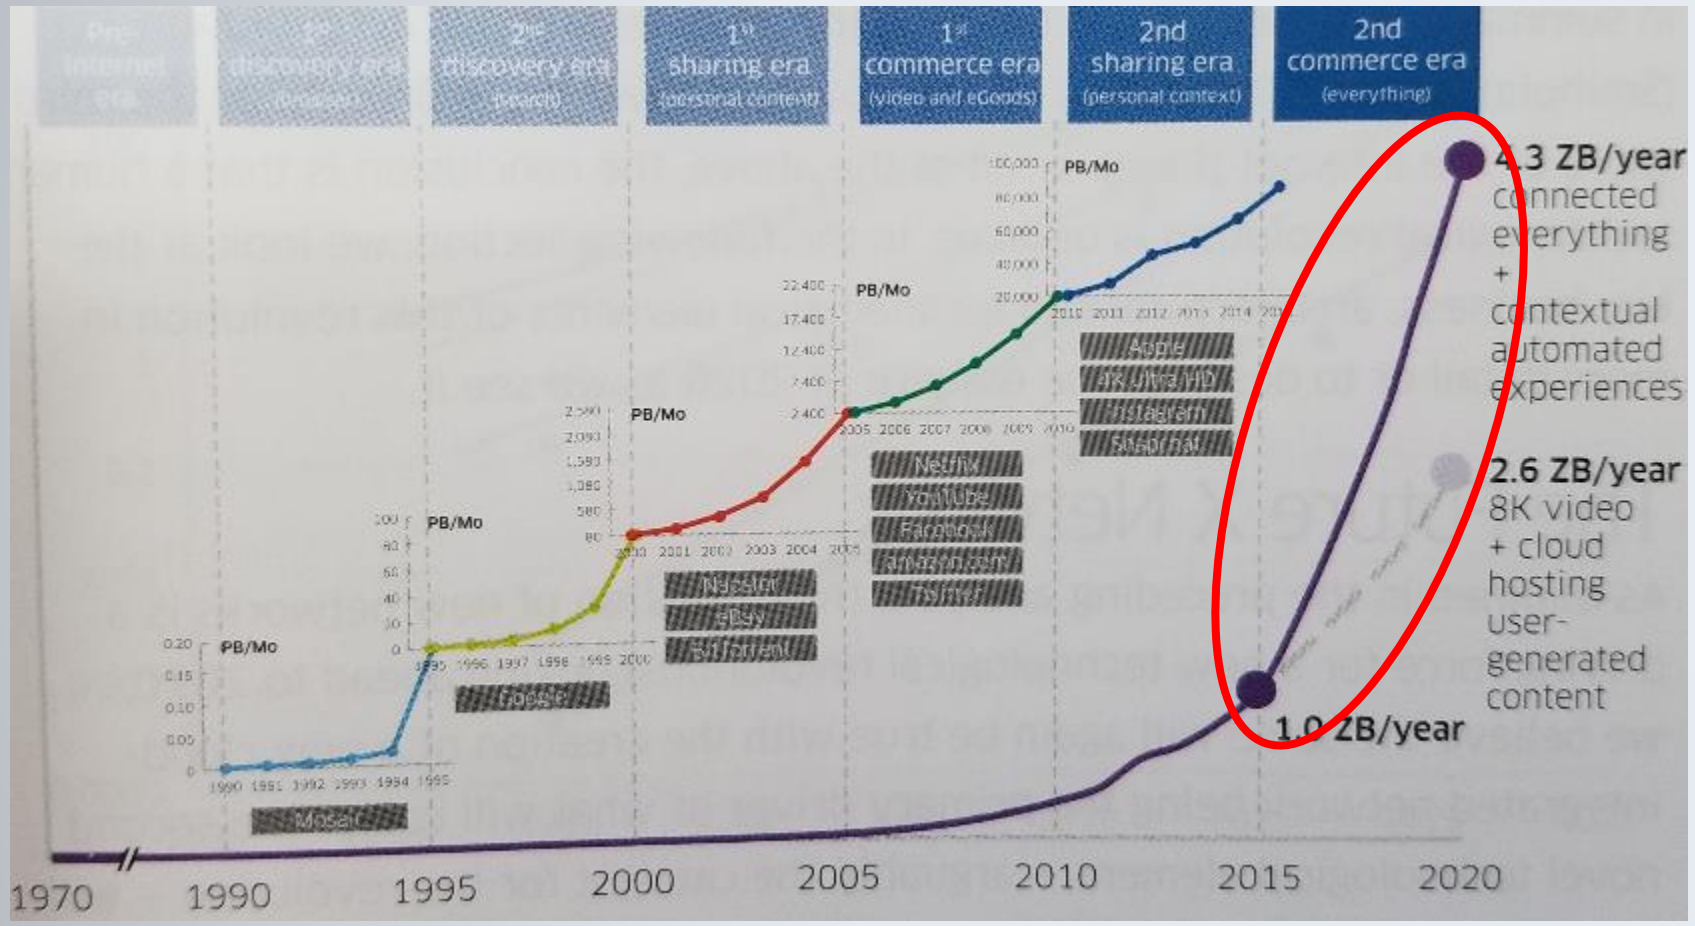

Marcus K. Weldon, The Future X Network

Hardware

400G Ethernet
(IEEE 802.3bs)

Software

Next-gen. VNF?

Agenda

1. Backgrounds
2. D-Plane Technologies
3. Performance
4. Research Directions

D-Plane Technologies in NFV-nodes

SPP NetVM openNetVM

FastClick The Universal Dataplane OpenFastPath

OpenDataPlane.org

vhost-user SR-IOV

DPDK DATA PLANE DEVELOPMENT KIT Netmap XDP

OvS Open vSwitch lagopus BESS

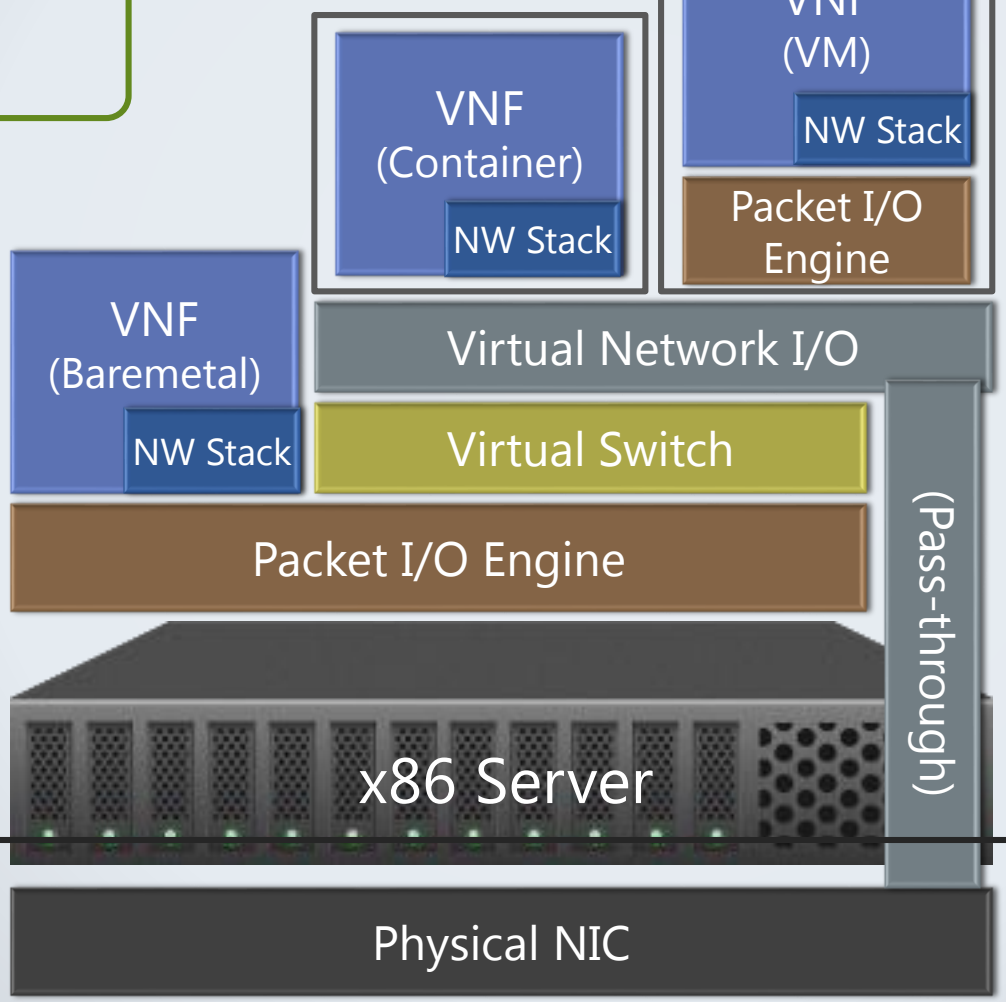

intel Mellanox TECHNOLOGIES

Chelsio Communications Accelerate QLOGIC NETRONOME

x86 Server Architecture (Intel Xeon Scalable Processor)

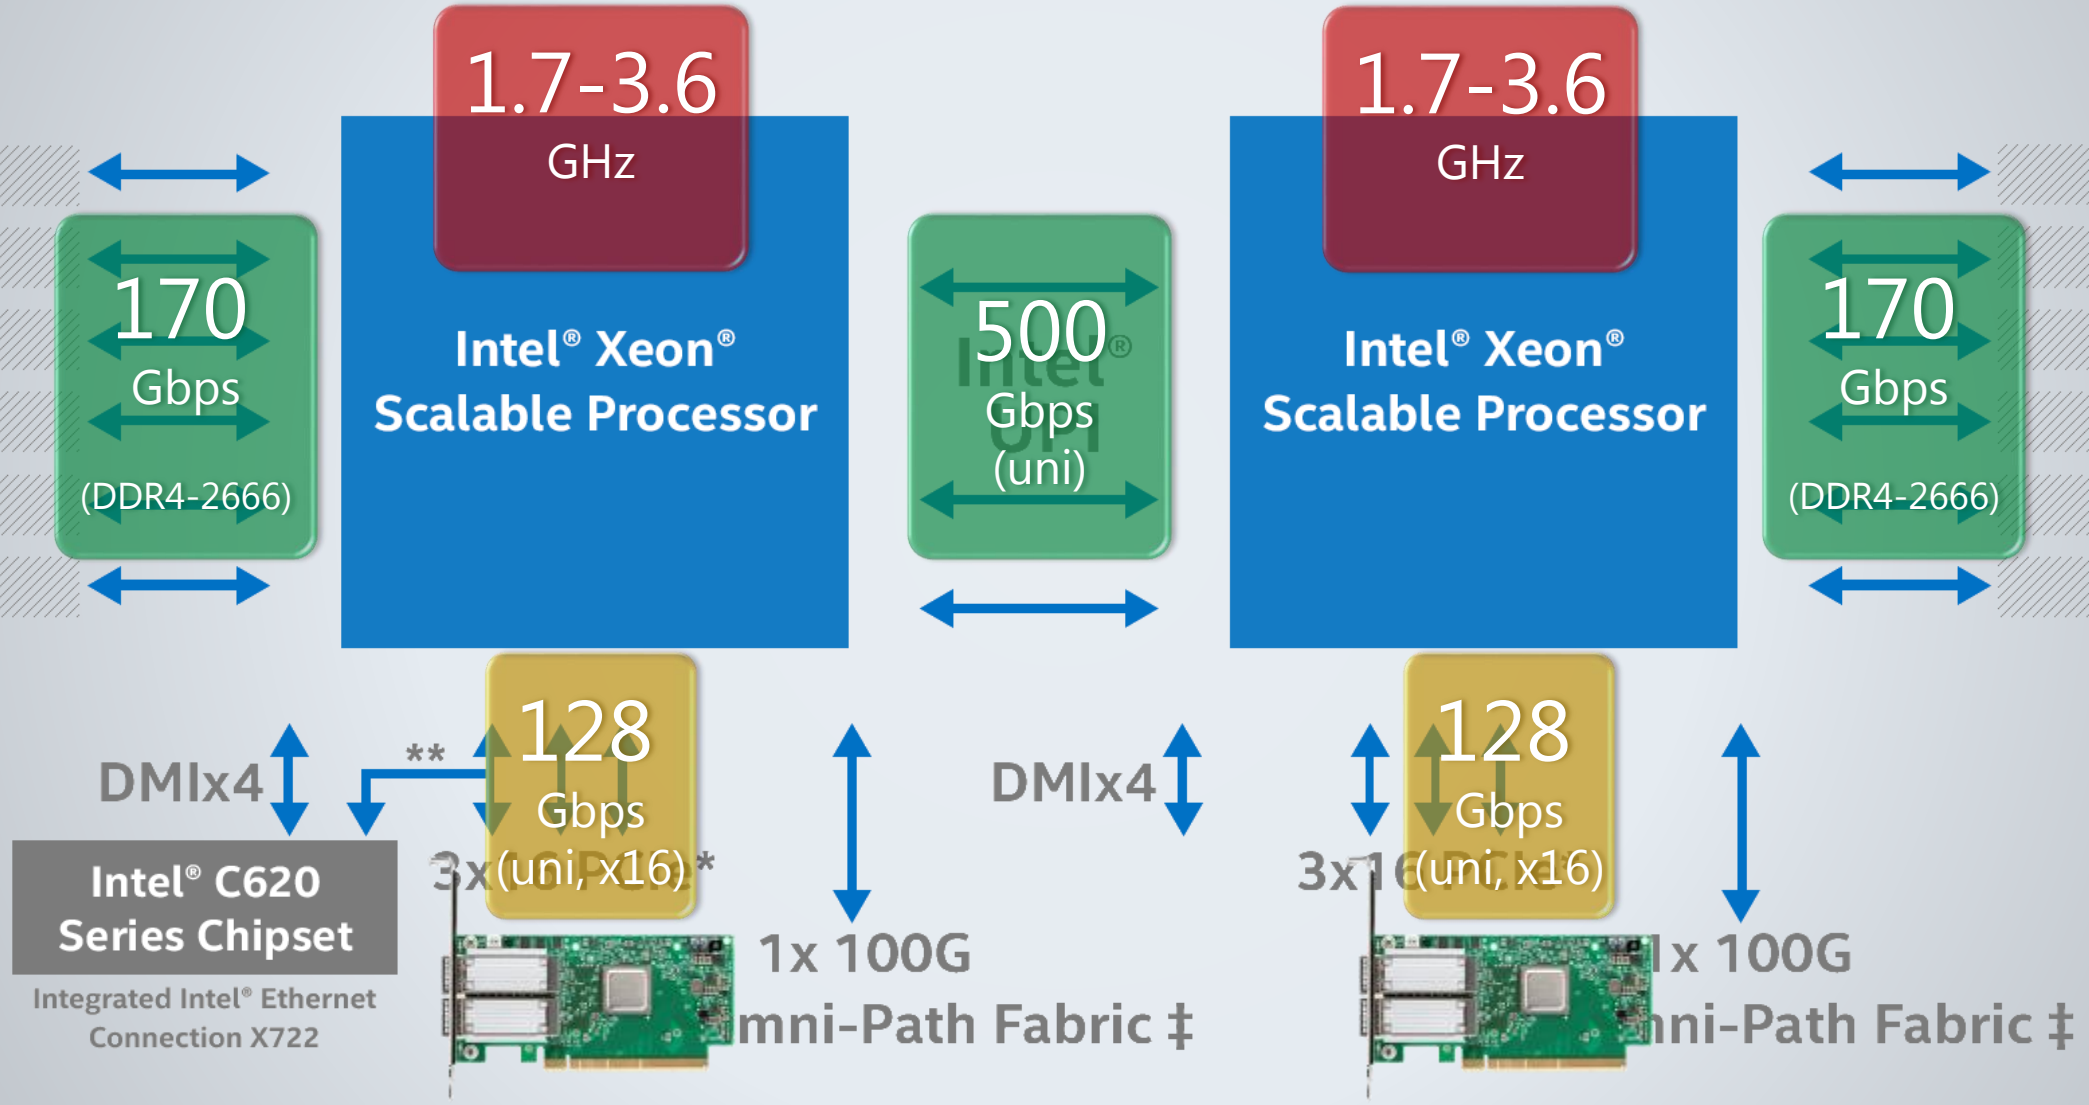

https://www.intel.com/content/www/us/en/design/products-and-solutions/processors-and-chipsets/purley/intel-xeon-scalable-processors.html

Fast Packet I/O Engine

DPDK
DATA PLANE DEVELOPMENT KIT

- Pure Polling
- Pre-allocated Memory
- User-space DMA
- Packet Batching

Bridging VNFs and the Host

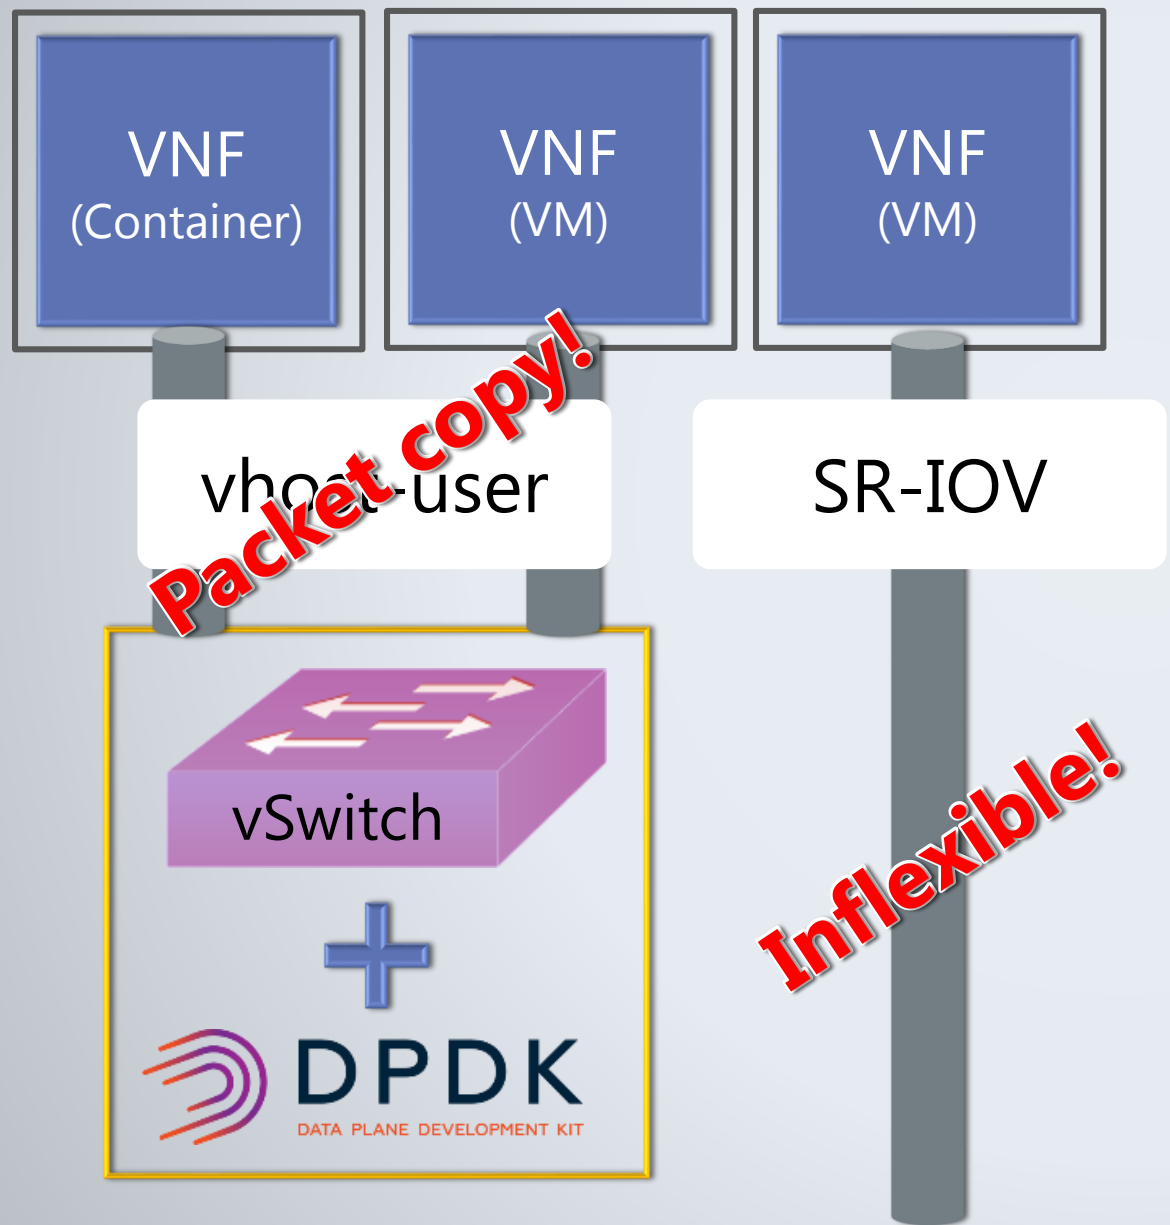

User-space Network Stack

Vector Packet Processing (VPP)

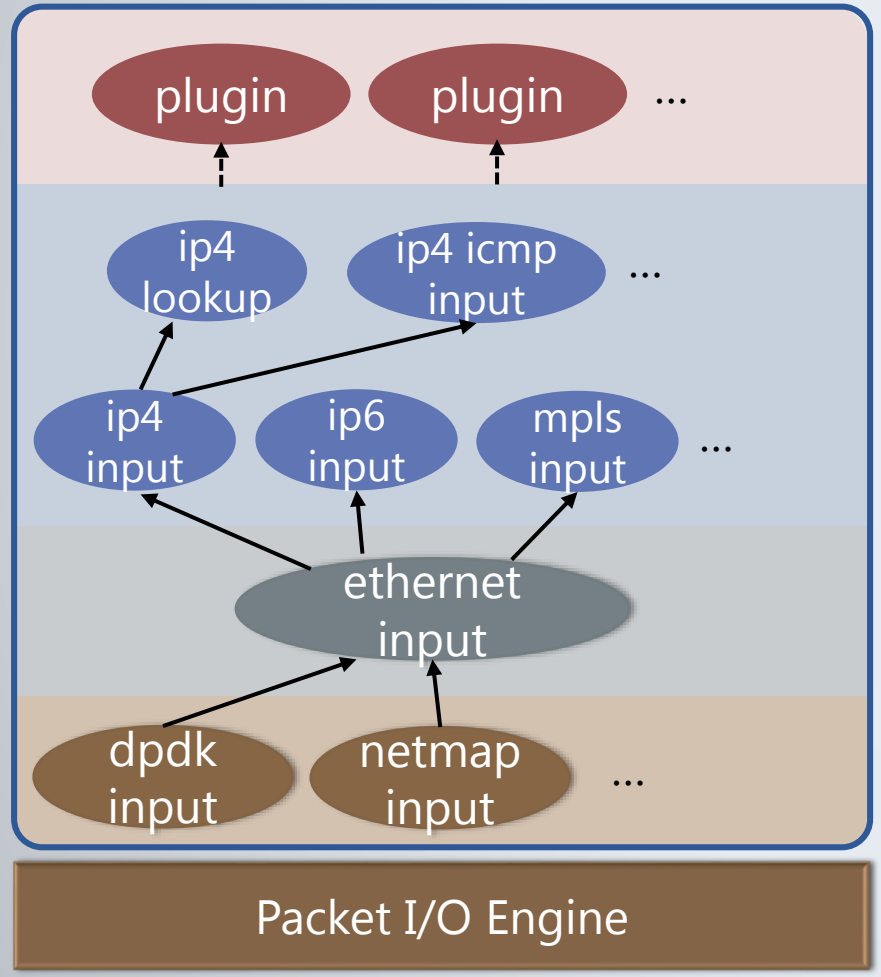

- Full network stack
- Click-like modular design
- Vectorized processing

Packet Vector

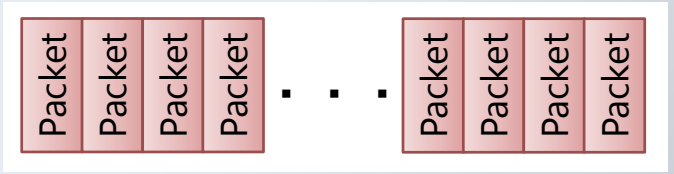

(N: 4 - 256)

Agenda

1. Backgrounds
2. D-Plane Technologies
3. Performance
4. Research Directions

Evaluation

Environment

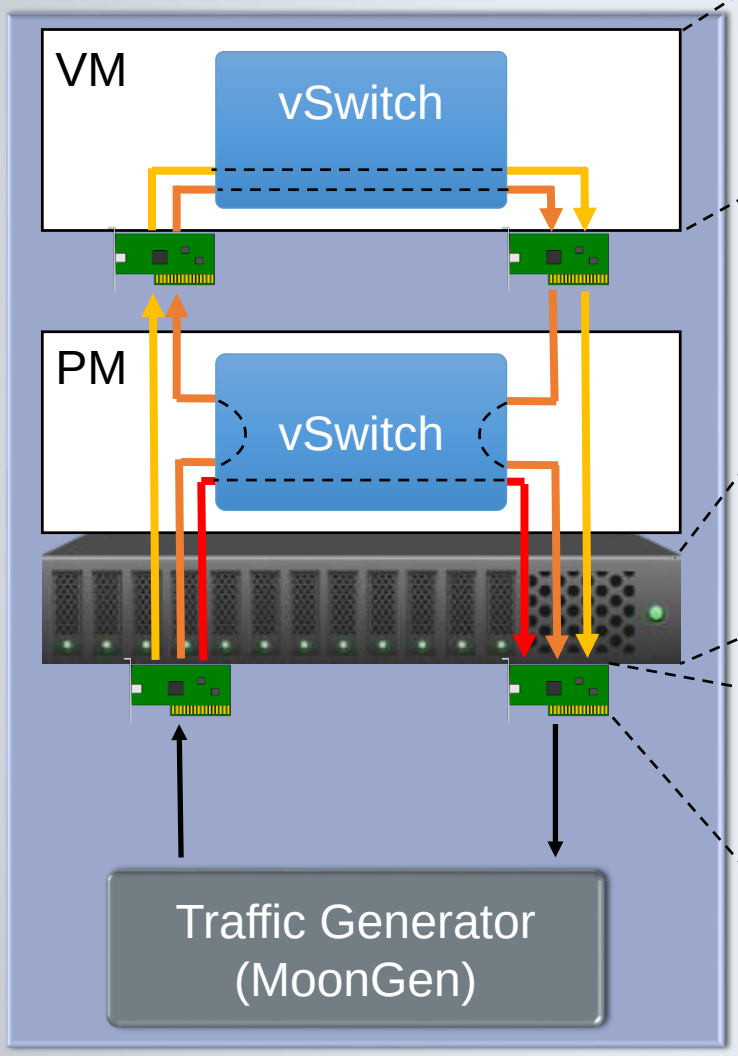

VM
3 cores
4 GB

Machine 1	Machine 2
Core i7 6700K 4.0 GHz (4 cores w HT)	Xeon E5-2630 v4 2.2 GHz (2x10 cores)
32 GB (DDR4-2133)	128 GB (DDR4-2133)

NIC
Intel XL710
40 GbE
PCIE 3.0 (x8)

R. Kawashima et al., "Evaluation of Forwarding Efficiency in NFV-nodes toward Predictable Service Chain Performance", IEEE TNSM, vol. 14, no. 4, 2017 (preprint)

LFMT (Lock-free Multi-threading)

Run-to-Completion

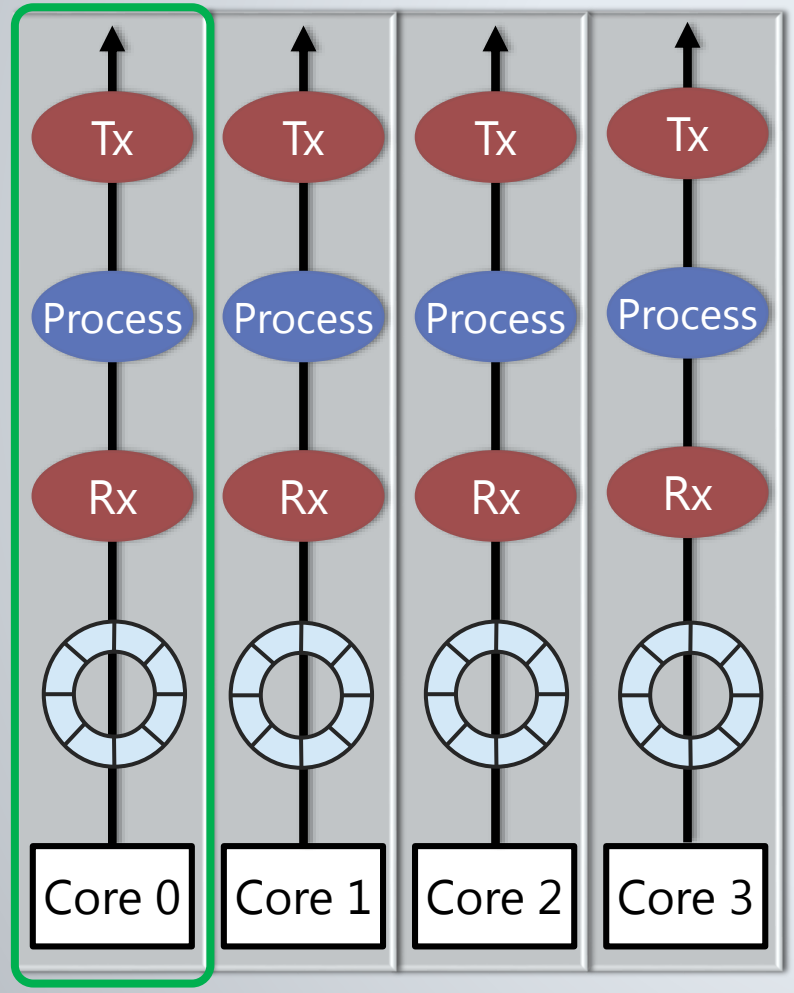

Lock free!

Scale out!

Pipeline

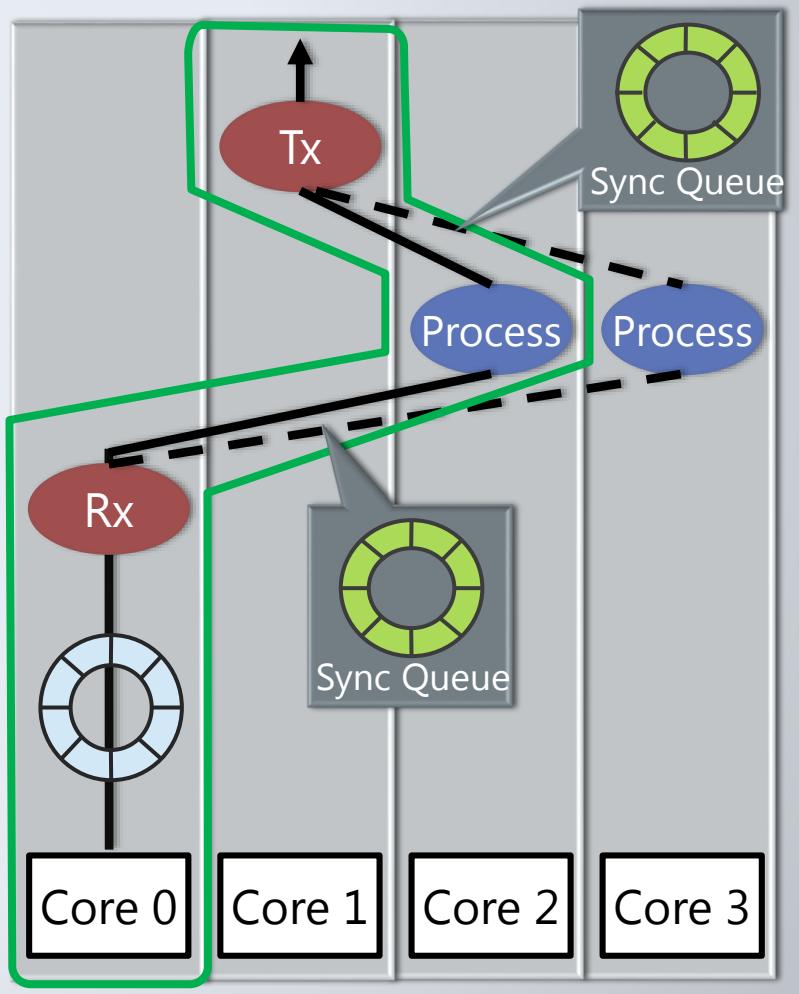

Phy-to-Phy Throughput (Baremetal)

Machine 1 (4.0 GHz)

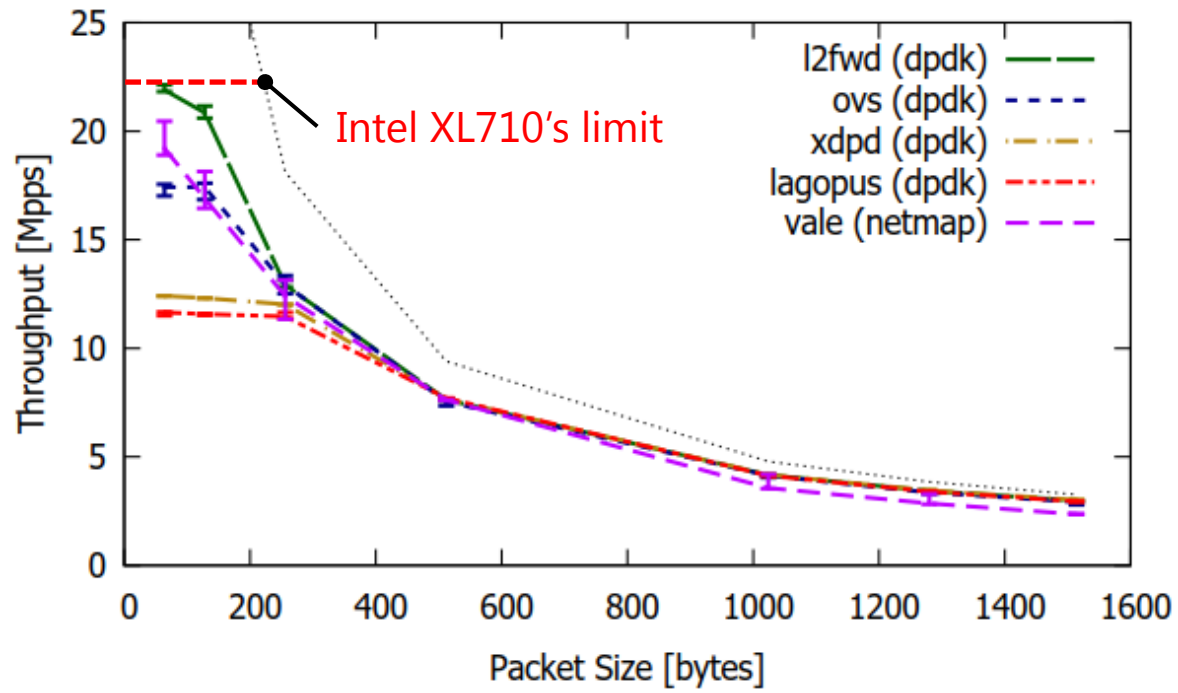

Machine 2 (2.2 GHz)

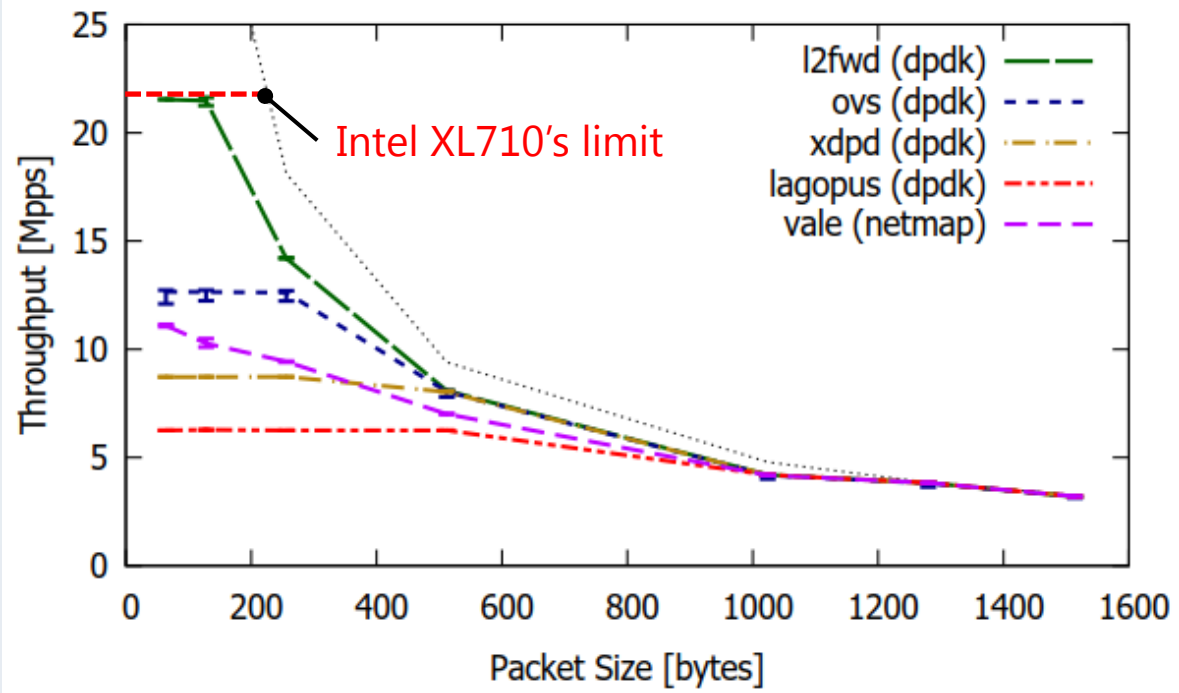

Poor throughput for short packet traffic

R. Kawashima et al., "Evaluation of Forwarding Efficiency in NFV-nodes toward Predictable Service Chain Performance", IEEE TNSM, vol. 14, no. 4, 2017 (preprint)

(Ref.) Intel XL710 vs. Mellanox ConnectX-5

Intel XL710 (40GbE)

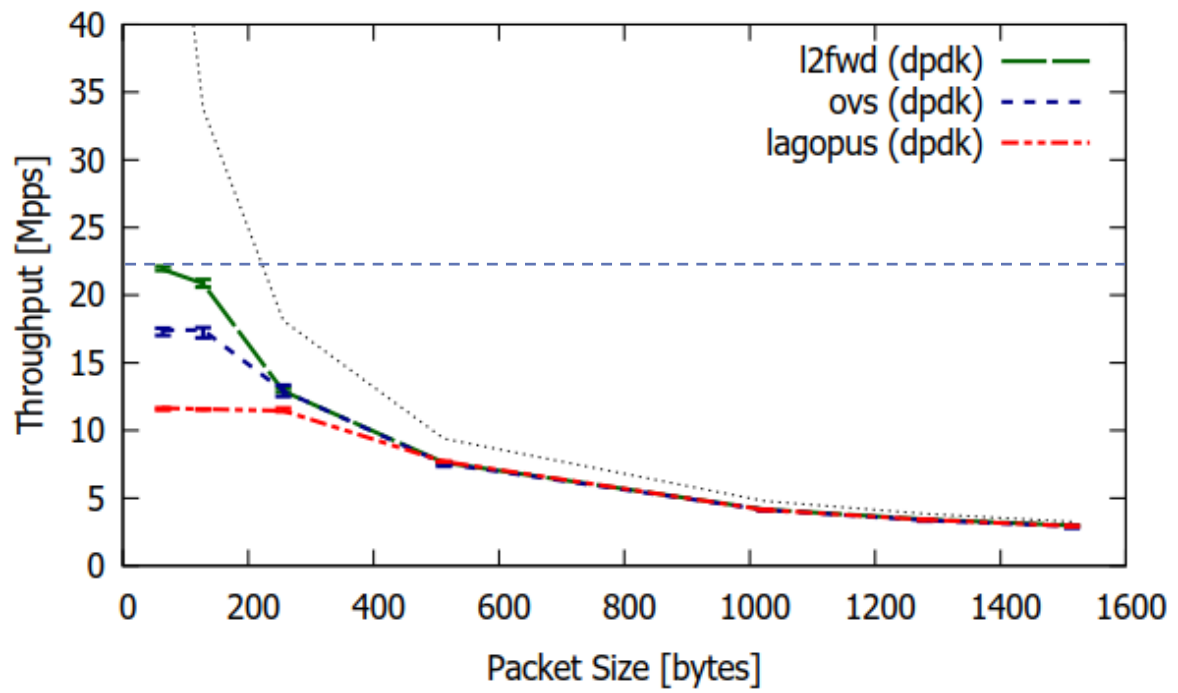

Mellanox ConnectX-5 (100GbE)

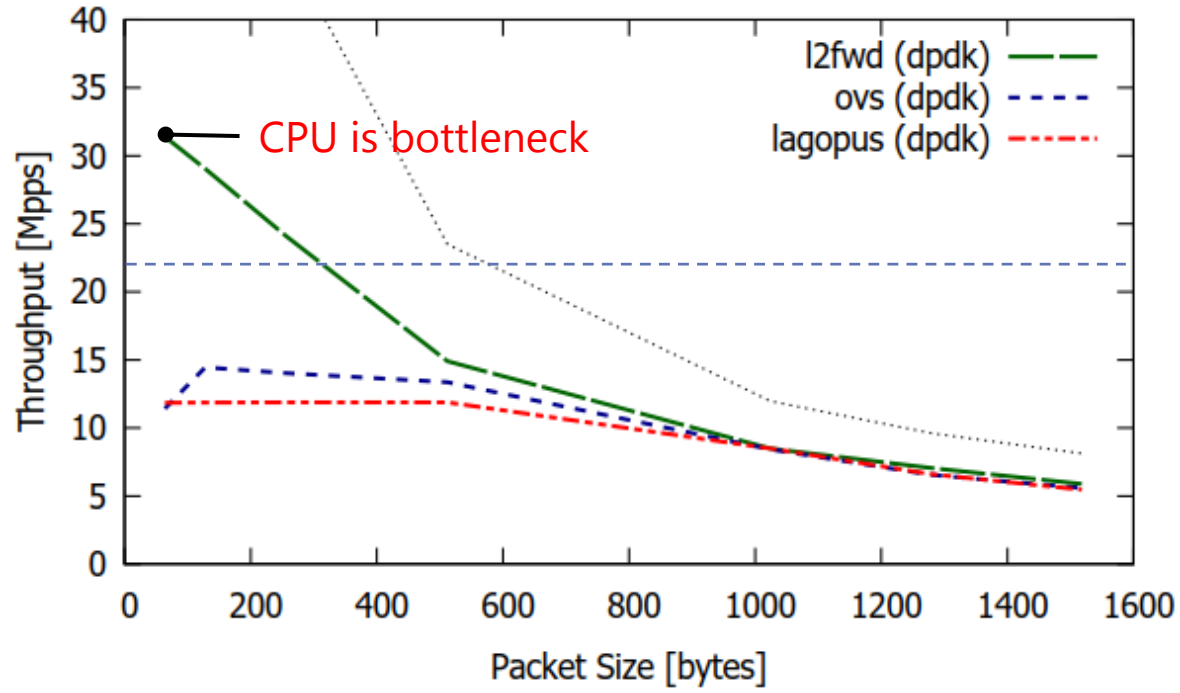

Single core can support a mere 30 Mpps

Phy-to-Phy Latency/Jitter (Baremetal)

Machine 1 (4.0 GHz)

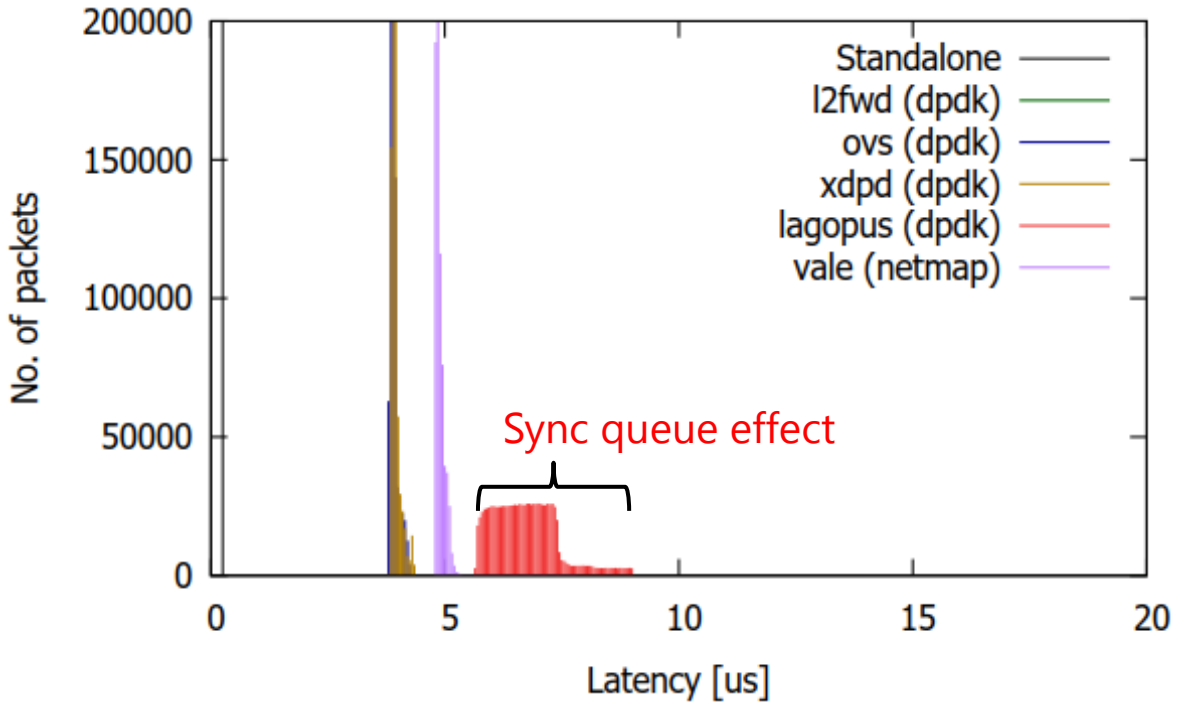

Machine 2 (2.2 GHz)

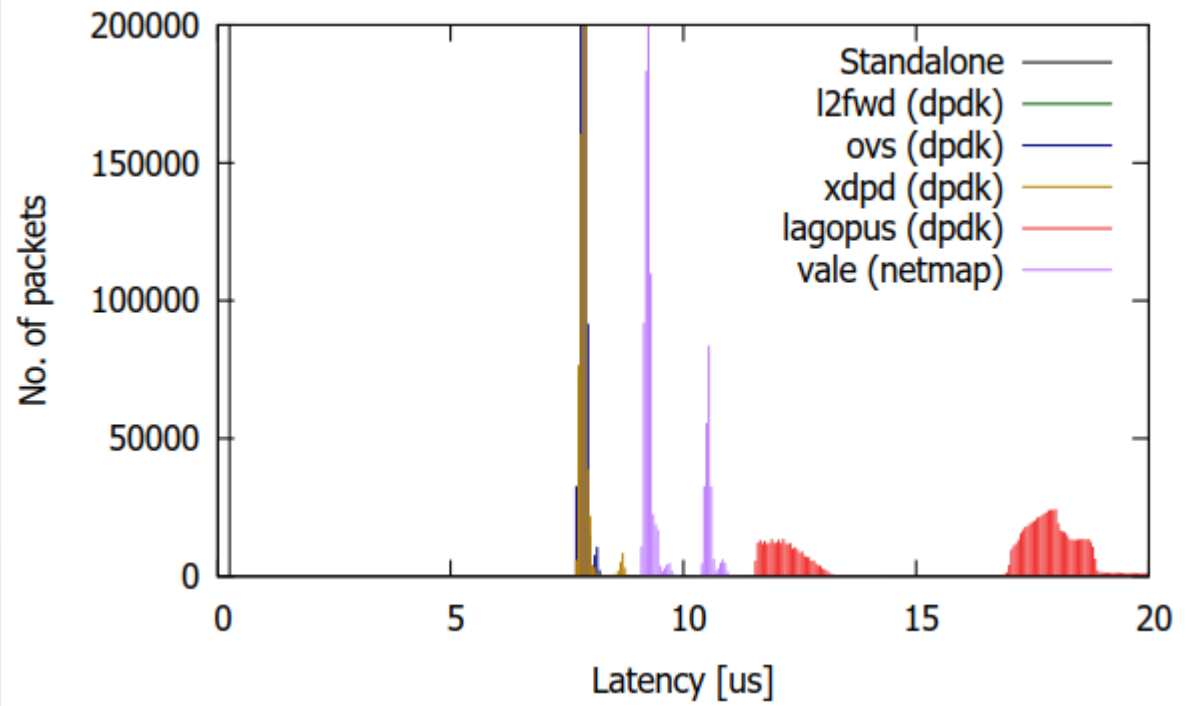

10x latency/jitter to hardware switches

R. Kawashima et al., "Evaluation of Forwarding Efficiency in NFV-nodes toward Predictable Service Chain Performance", IEEE TNSM, vol. 14, no. 4, 2017 (preprint)

Evaluation Summary

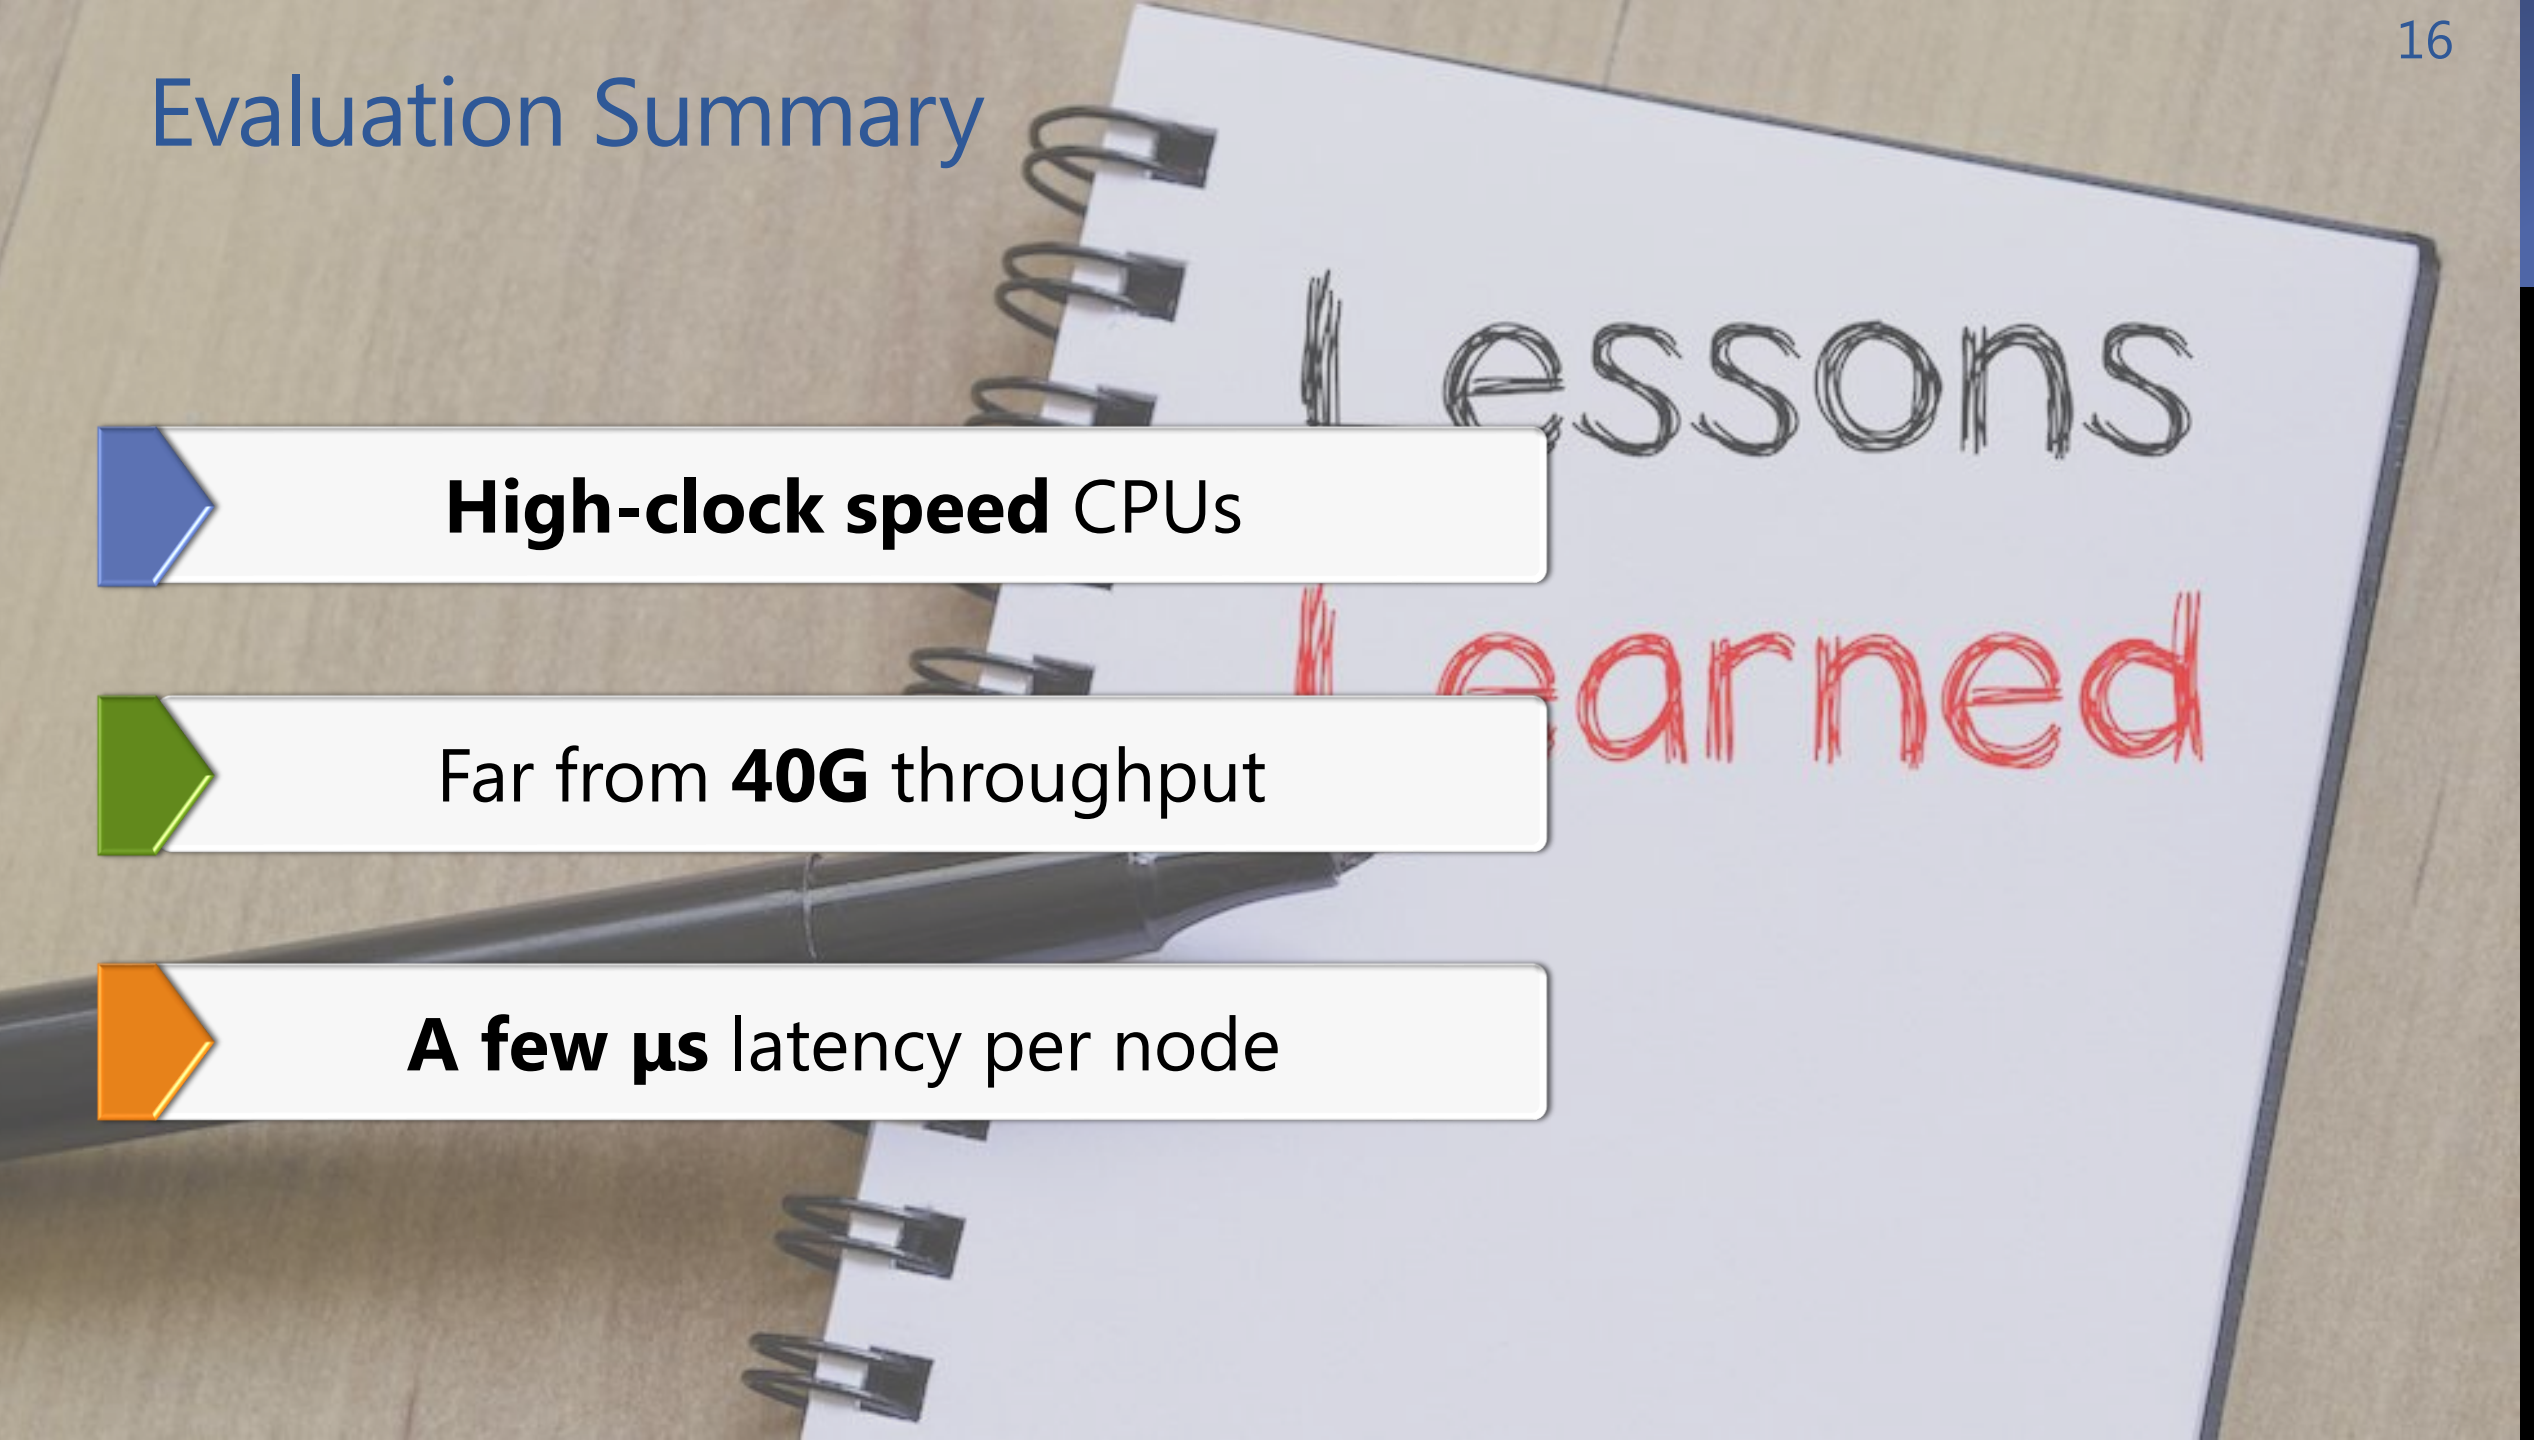

- **High-clock speed** CPUs
- Far from **40G** throughput
- **A few µs** latency per node

Agenda

1. Backgrounds
2. D-Plane Technologies
3. Performance
4. Research Directions

Research Direction 1

Scale-out Approach

- ✅ Aggregate Throughput
- Per-datapath Throughput
- Latency/Jitter
- ❌ Power Consumption

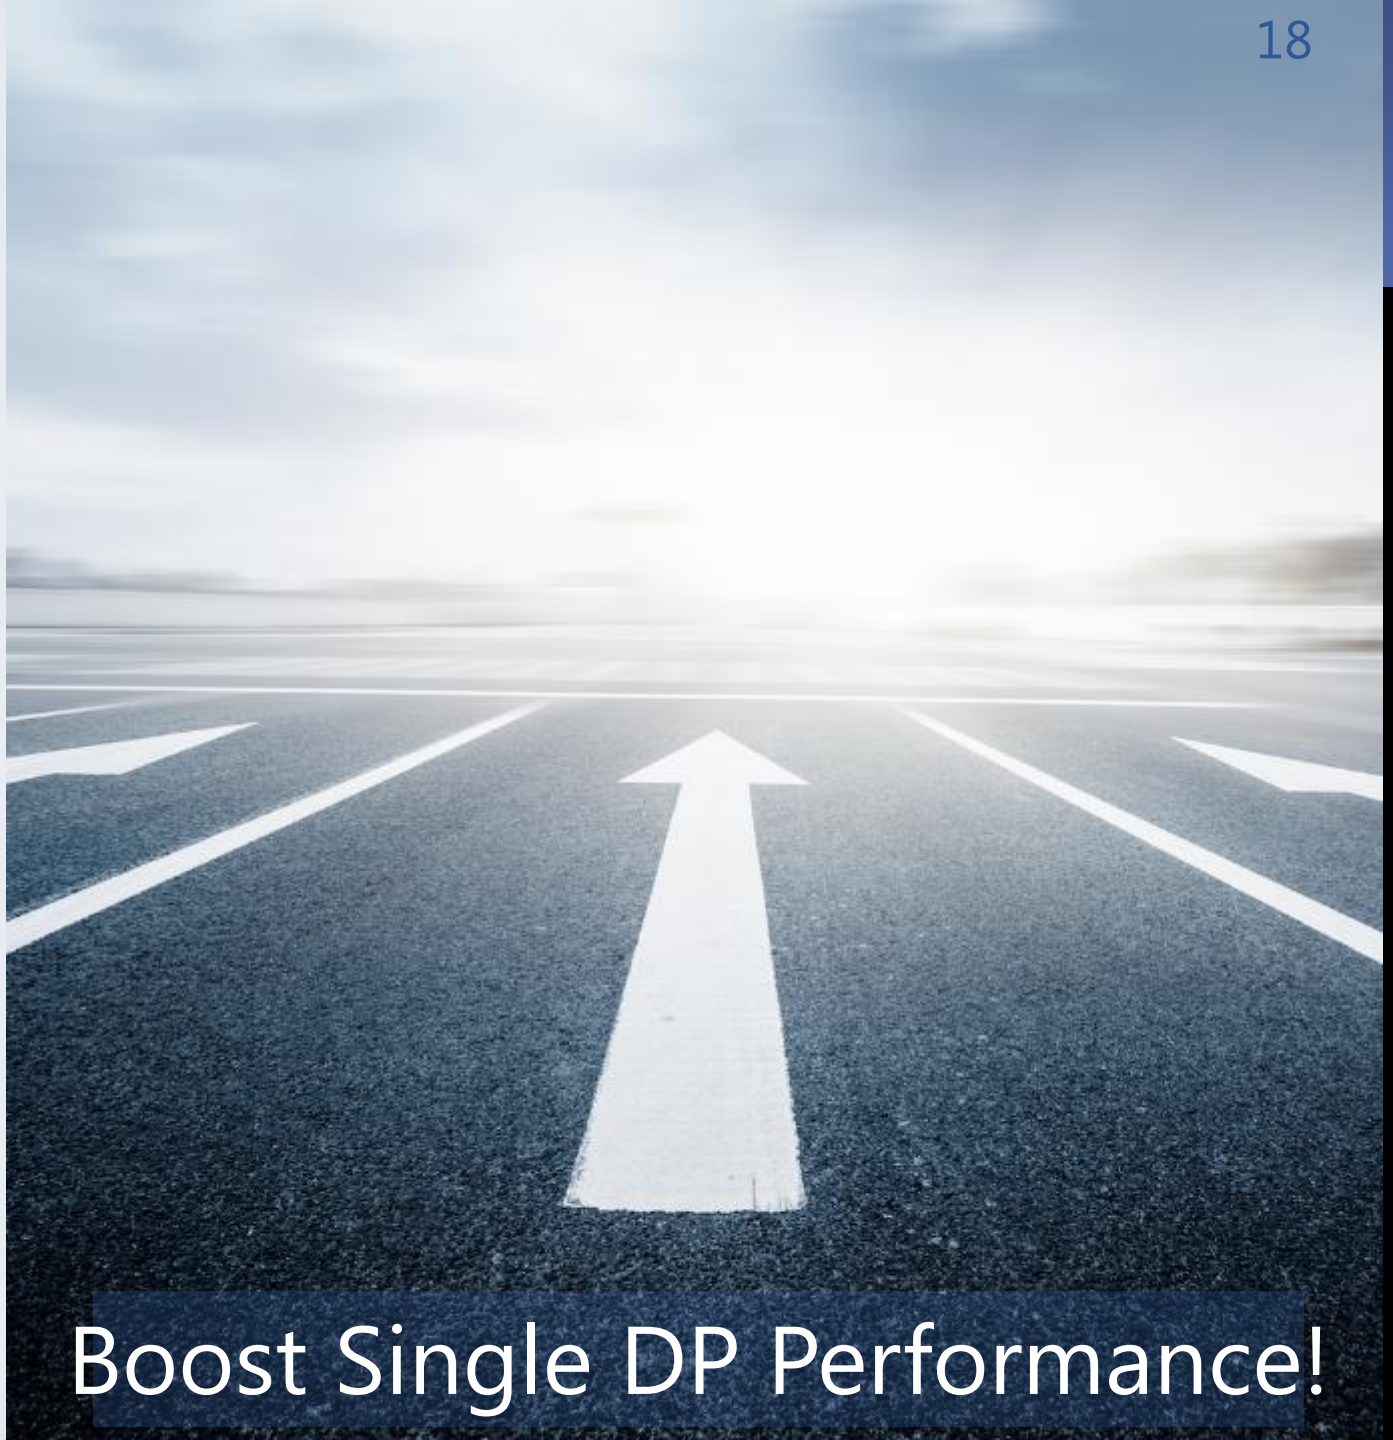

Boost Single DP Performance!

Approach 1: Hardware Offloading

NETRONOME (Smart NIC)

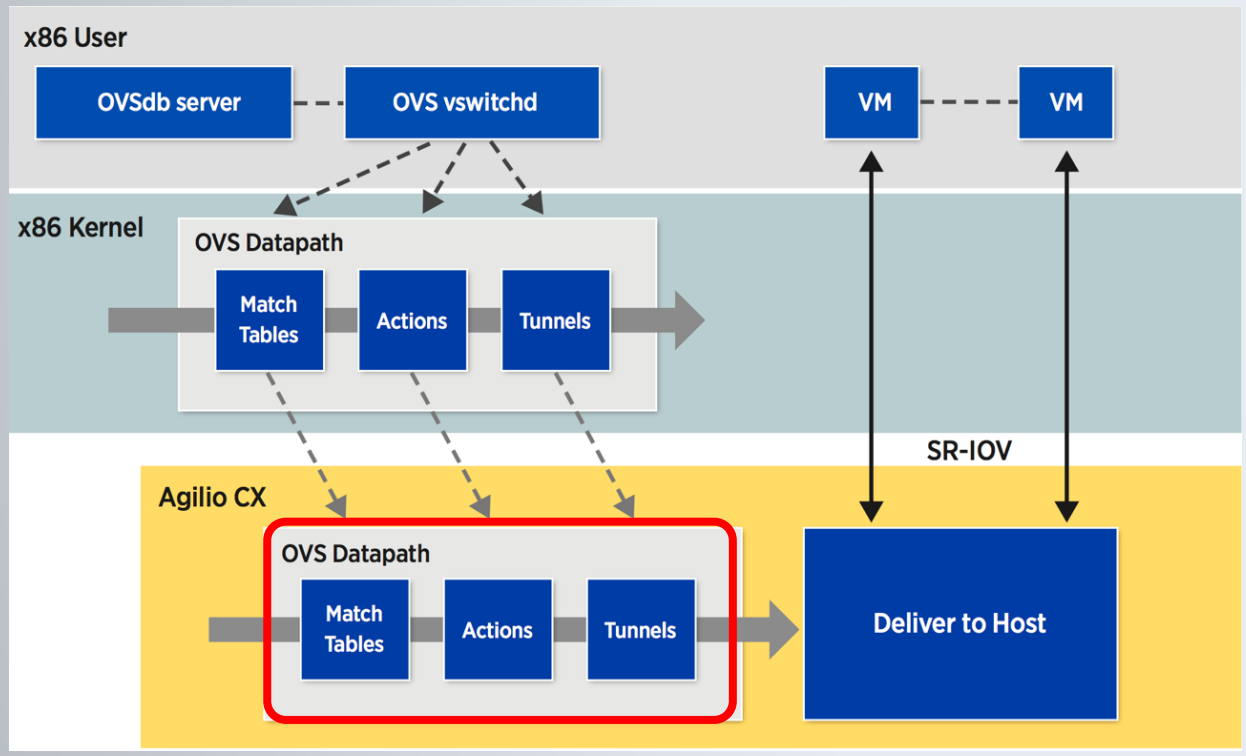

PHY-to-VM Throughput (40G)

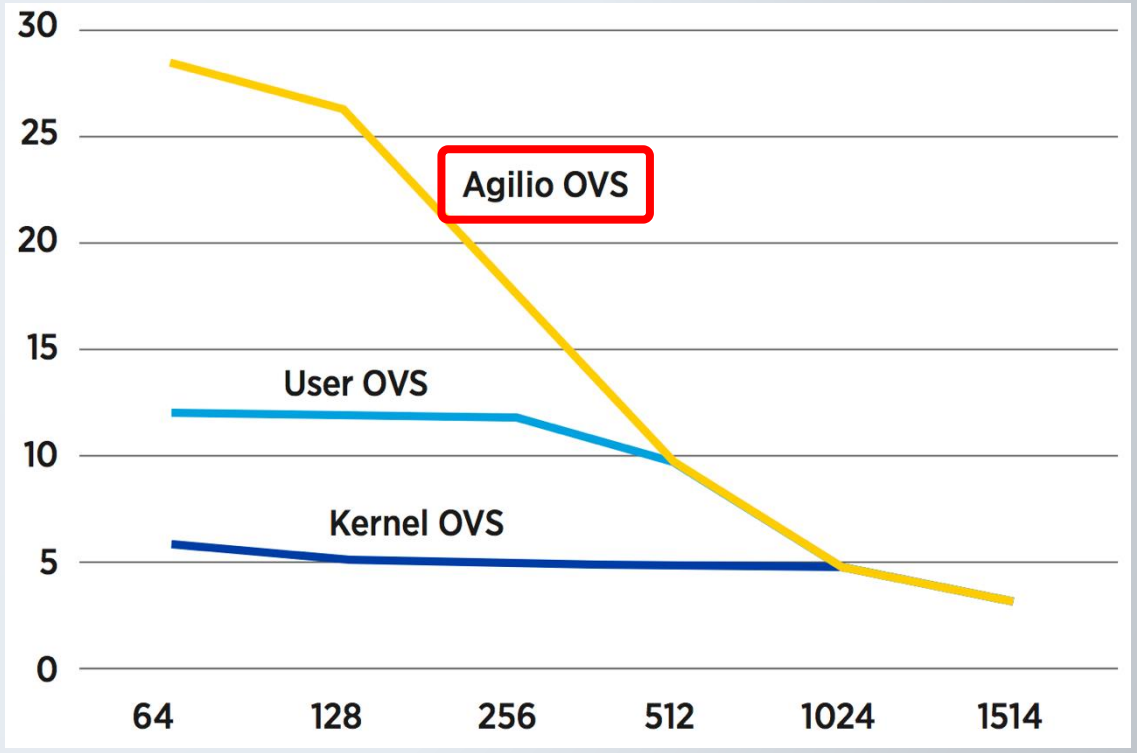

https://www.netronome.com/

Approach 2: Packet Aggregation

PA-Flow†

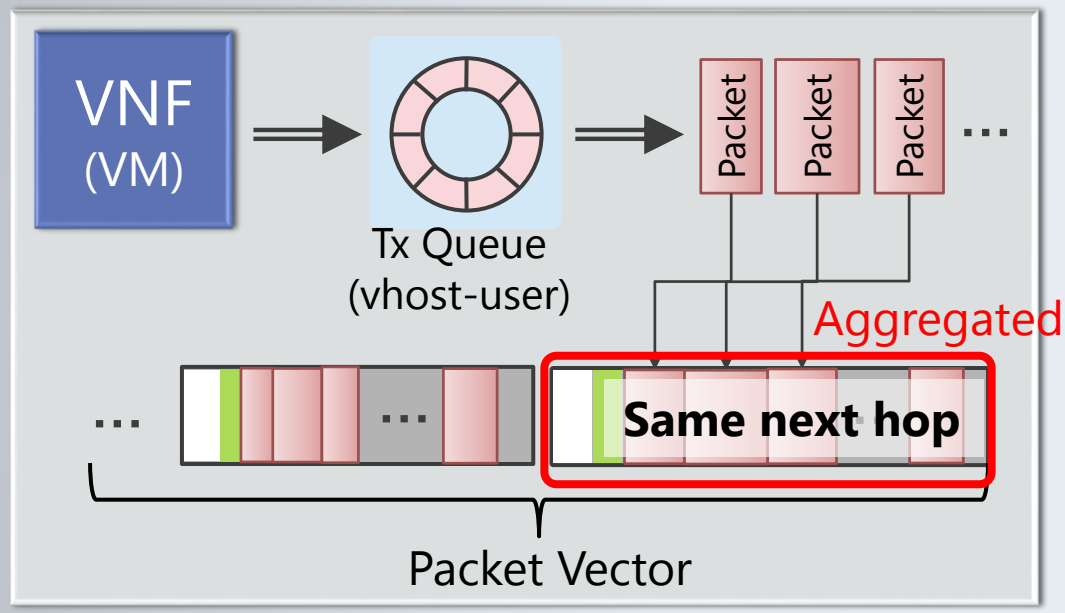

Traffic amount increases …

↳ Aggregation rate increases!

✅ Datapath Throughput

70% up!

👍 Datapath Latency

0.2µs decreases!

† Y. Taguchi et al., "Efficient NFV using Gradual Packet Aggregation Approach Focusing on Network I/O Processing between VMs and a Host", IEICE Tech. Rep., ICM2017-10, pp. 33-38, July 2017.

Research Direction 2

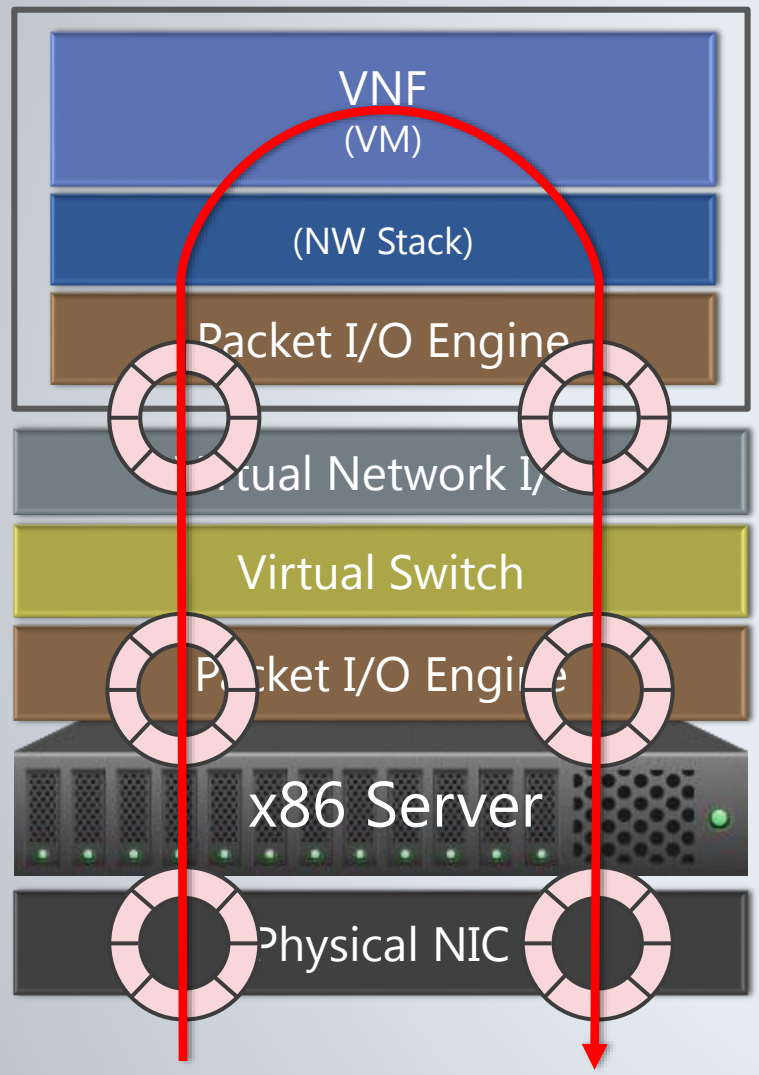

What about Service Chaining?

Reduce No. Enqueue Times!

Approach: Local Service Chaining

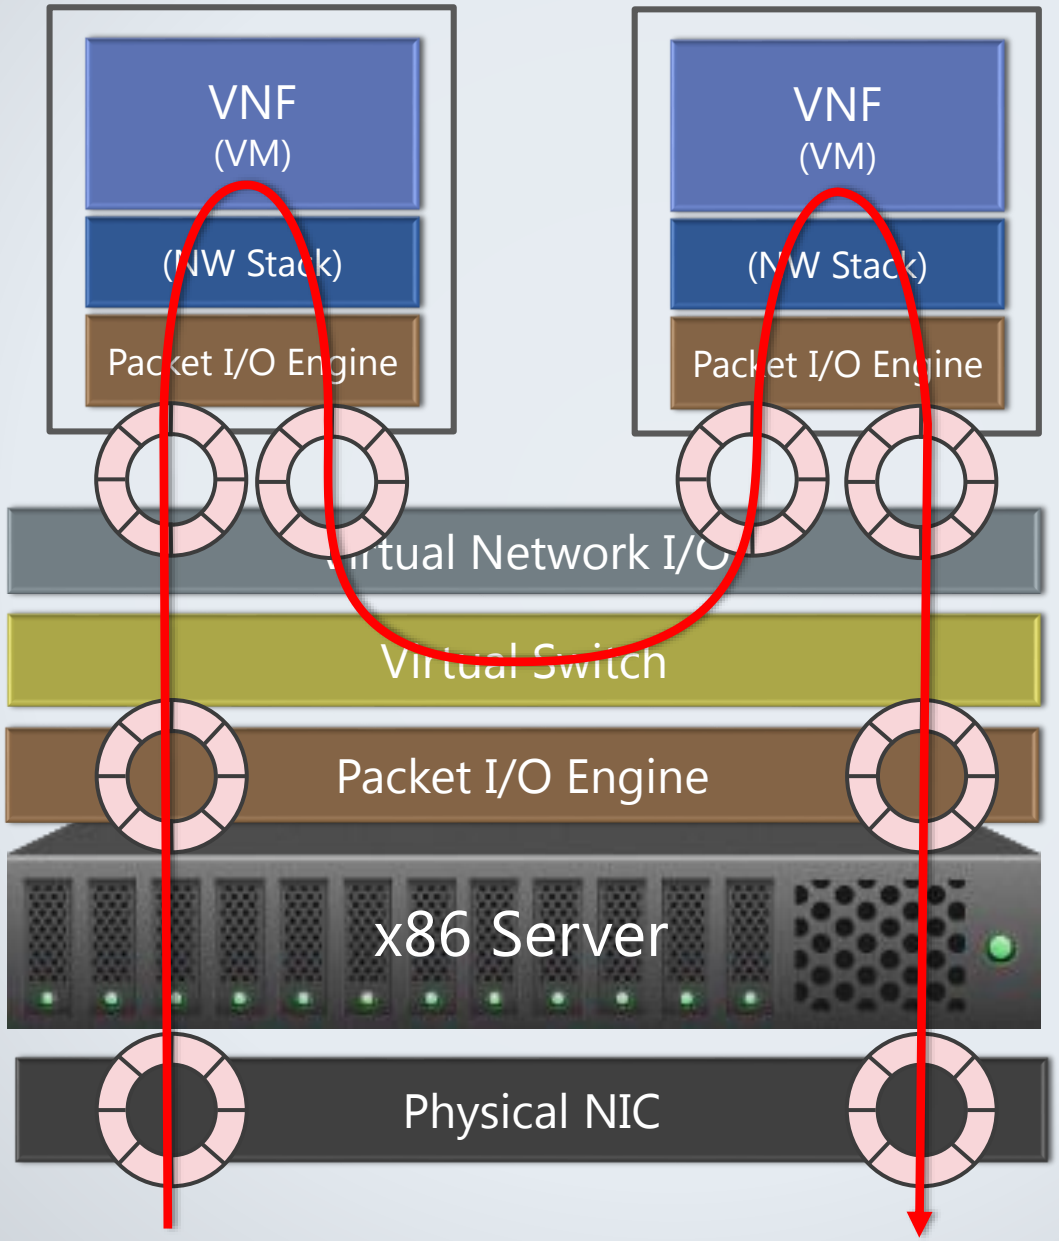

Fast Inter-VNF Communication

Soft Patch Panel (SPP)

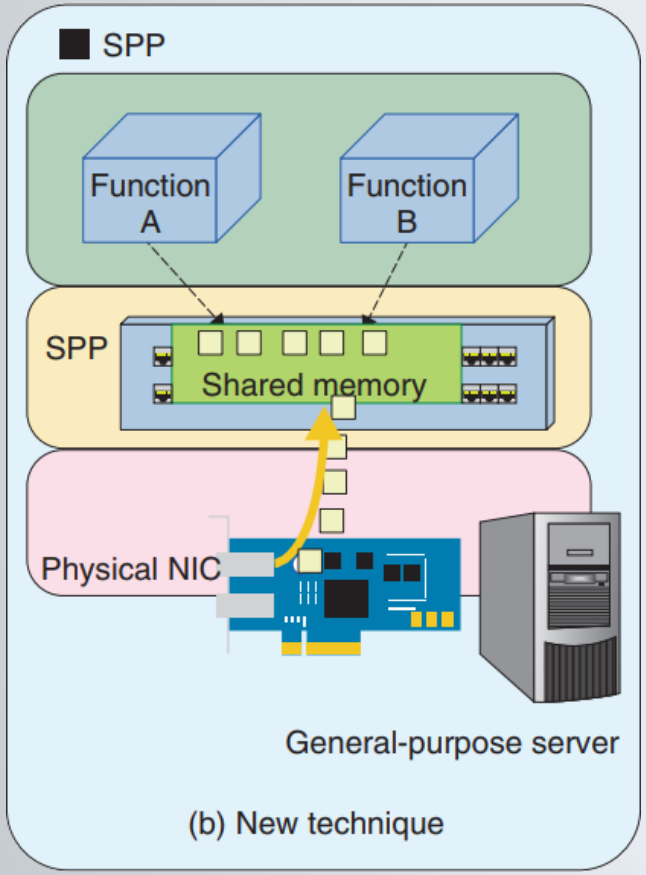

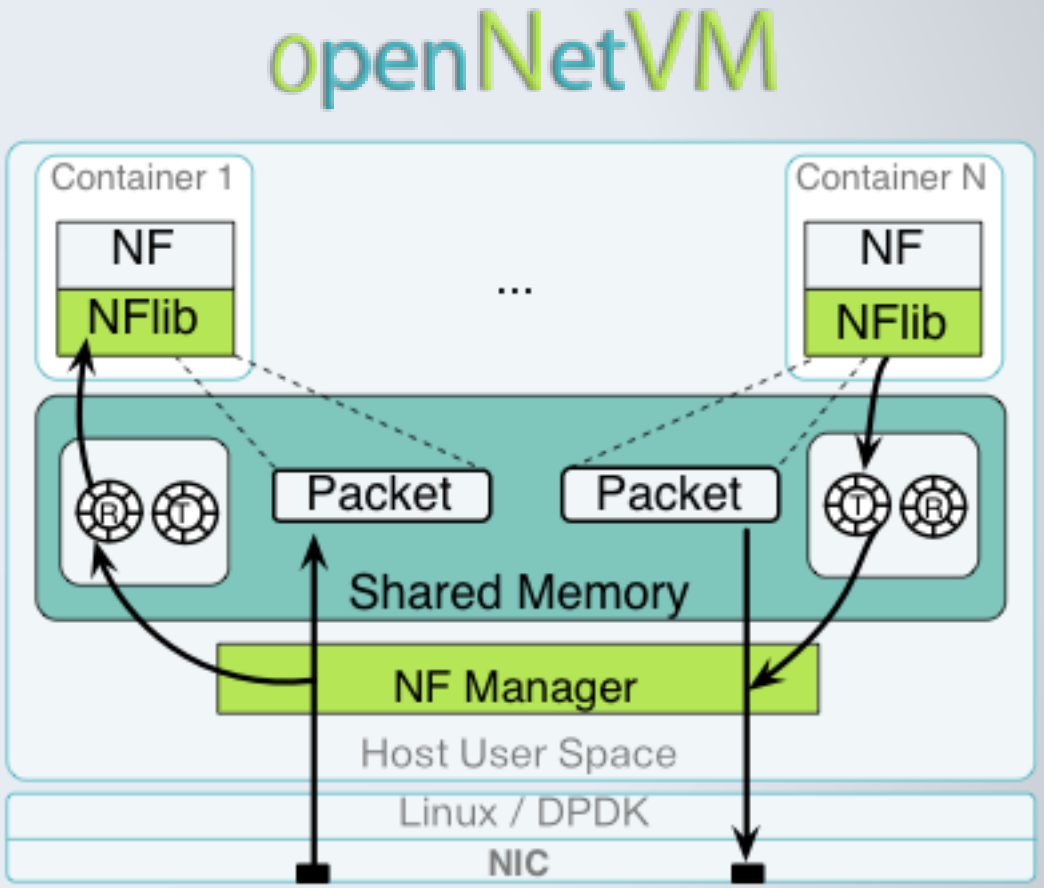

T. Nakamura et al., "MAGONIA (Soft Patch Panel): High-speed Inter-function Technique", NTT Tech. Rev., Vol. 14 No. 10 Oct. 2016

W. Zhang et al., "OpenNetVM: A Platform for High Performance Network Service Chains", Proc. HotMiddlebox, pp. 26-31, Aug. 2016

Dynamic VNF Fusion (ultimate)

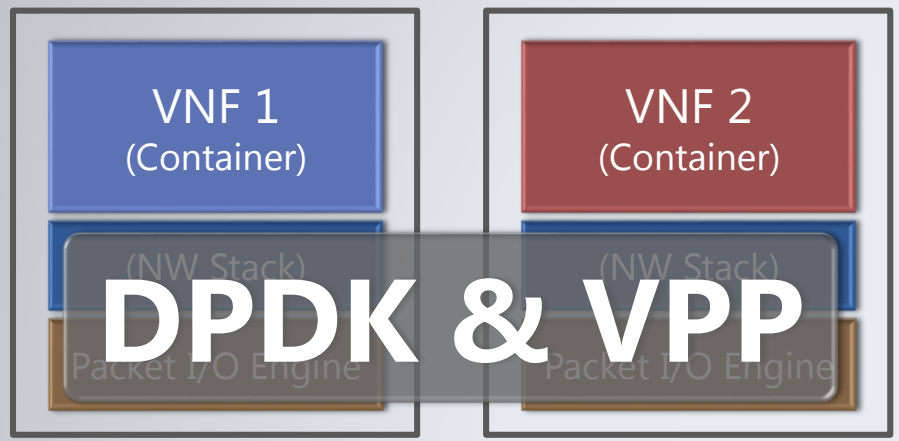

Dynamic Fusion

- Fusing processes
- Merging VPP graphs

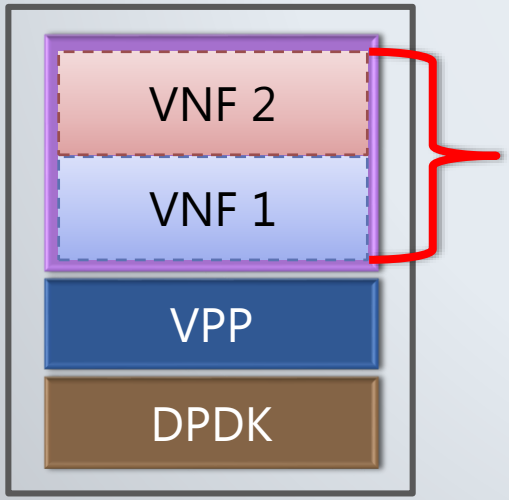

Single Datapath
(no sync queue)

Advantages

- ✅ Single DP Performance
- ✅ Modularity

Challenges

- 😓 Process Fusion Tech.

Summary

(Hardware Offloading)

Packet Aggregation

Local Service Chaining

Dynamic VNF Fusion

Boost Single DP Performance!

Reduce No. Enqueue Times!