Petr Velan
petr.velan@cesnet.cz

High-Density Network Flow Monitoring

IM2015

12 May 2015, Ottawa

Motivation

What is high-density flow monitoring?

- Monitor high traffic in as little rack units as possible

Why do we want high-density flow monitoring?

- Flow monitoring is deployed on many lines
- Number of flow probes is growing
- Management and operational costs are growing
- One probe per link does not scale

Our Approach

- Use our custom made network interface cards to monitor multiple 10G links
- See what throughput can be achieved in one machine
- Test how advanced features of our NICs help the monitoring
- Identify performance bottlenecks

Monitoring Setup

The Testbed

The Server

- Dell PowerEdge R720 server (2 rack units)
- 2× E5-2670 v1 CPUs (8 cores, 3 GHz)
- 64GB DDR3 RAM (1600 MHz)
- Scientific Linux 6.5 with 2.6.32-41 kernel
- 2× COMBO-80G cards

COMBO-80G NIC

- FPGA based programmable hardware
- Two QSPF+ interfaces in 40 G or 4× 10 G Ethernet mode
- **80 G** per card
- PCI-Express gen3 x8 bus (64 Gb/s)
- Additional features:
  - Accurate timestamps
  - Hash based packet distribution
  - Packet trimming
  - Packet feature extraction into Unified Header

Our setup allows to monitor **16**× **10 G** links

COMBO-80G NIC

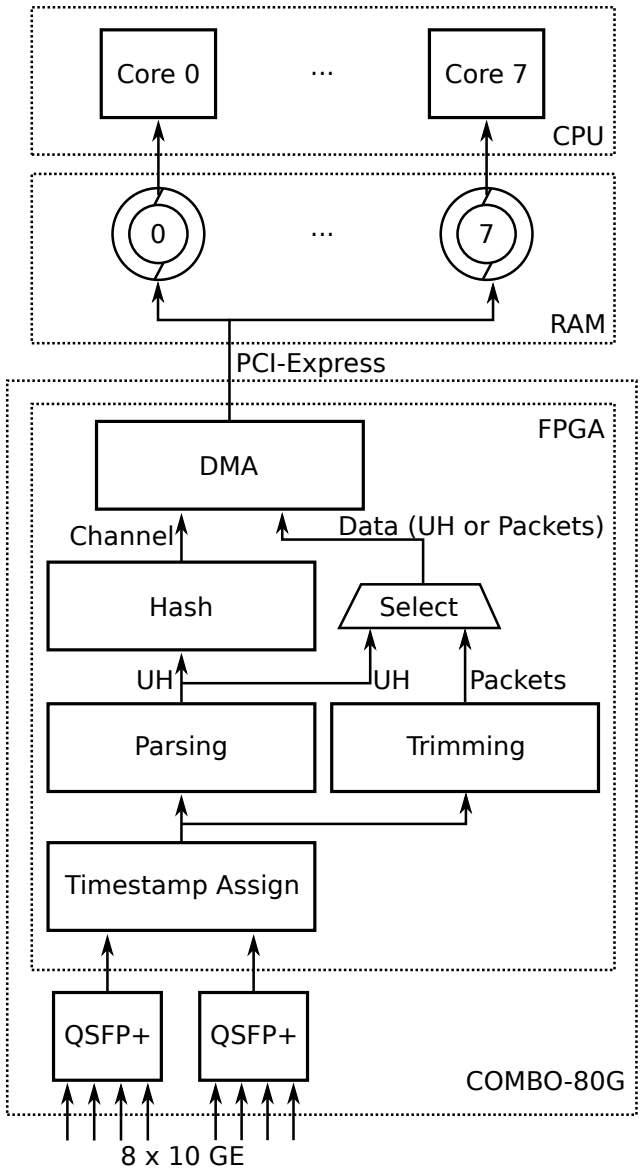

Flow Exporter Architecture

- Multi-threaded design
- Utilizes $2N + 1$ CPU cores where N is number of ring buffers

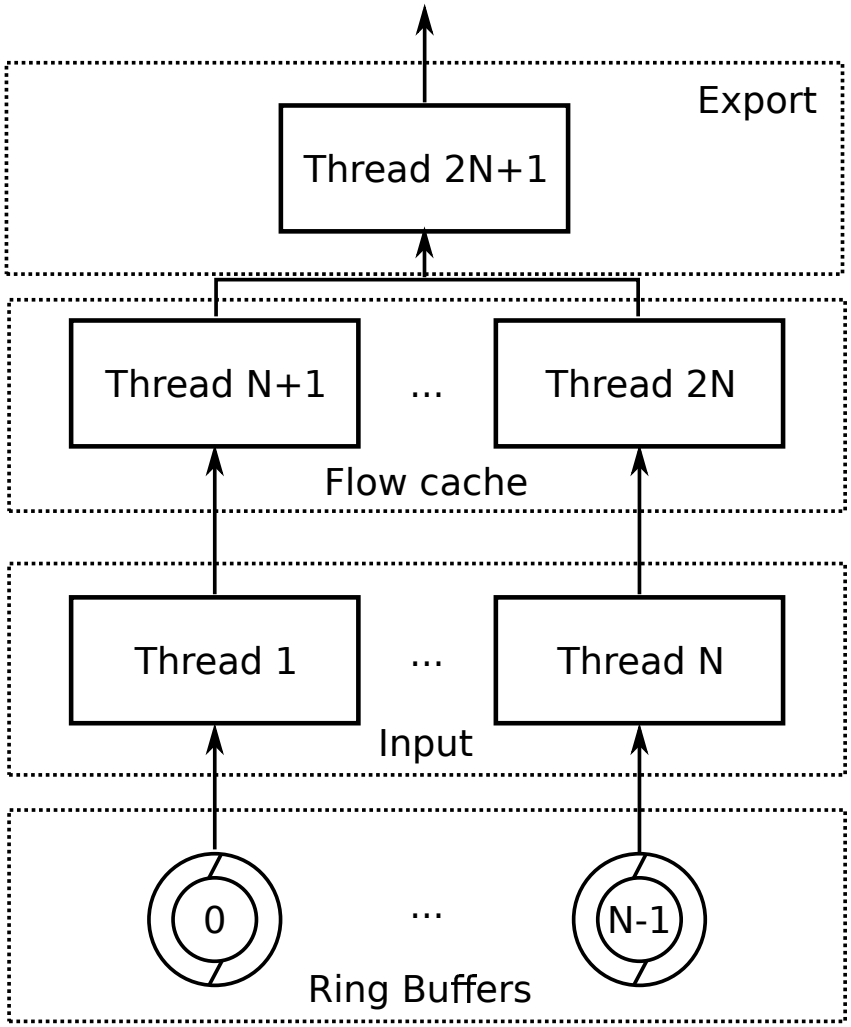

Data Generation

- Spirent TestCenter hardware
- 1× 10 G repeated to all 16 interfaces
- IPv4 UDP packets
- Packet sizes 64 B, 128 B, 256 B, 512 B
- Flow counts 2^{11}, 2^{18}, 2^{21}

Results

The Measurement

Basic Performance on Full Packets

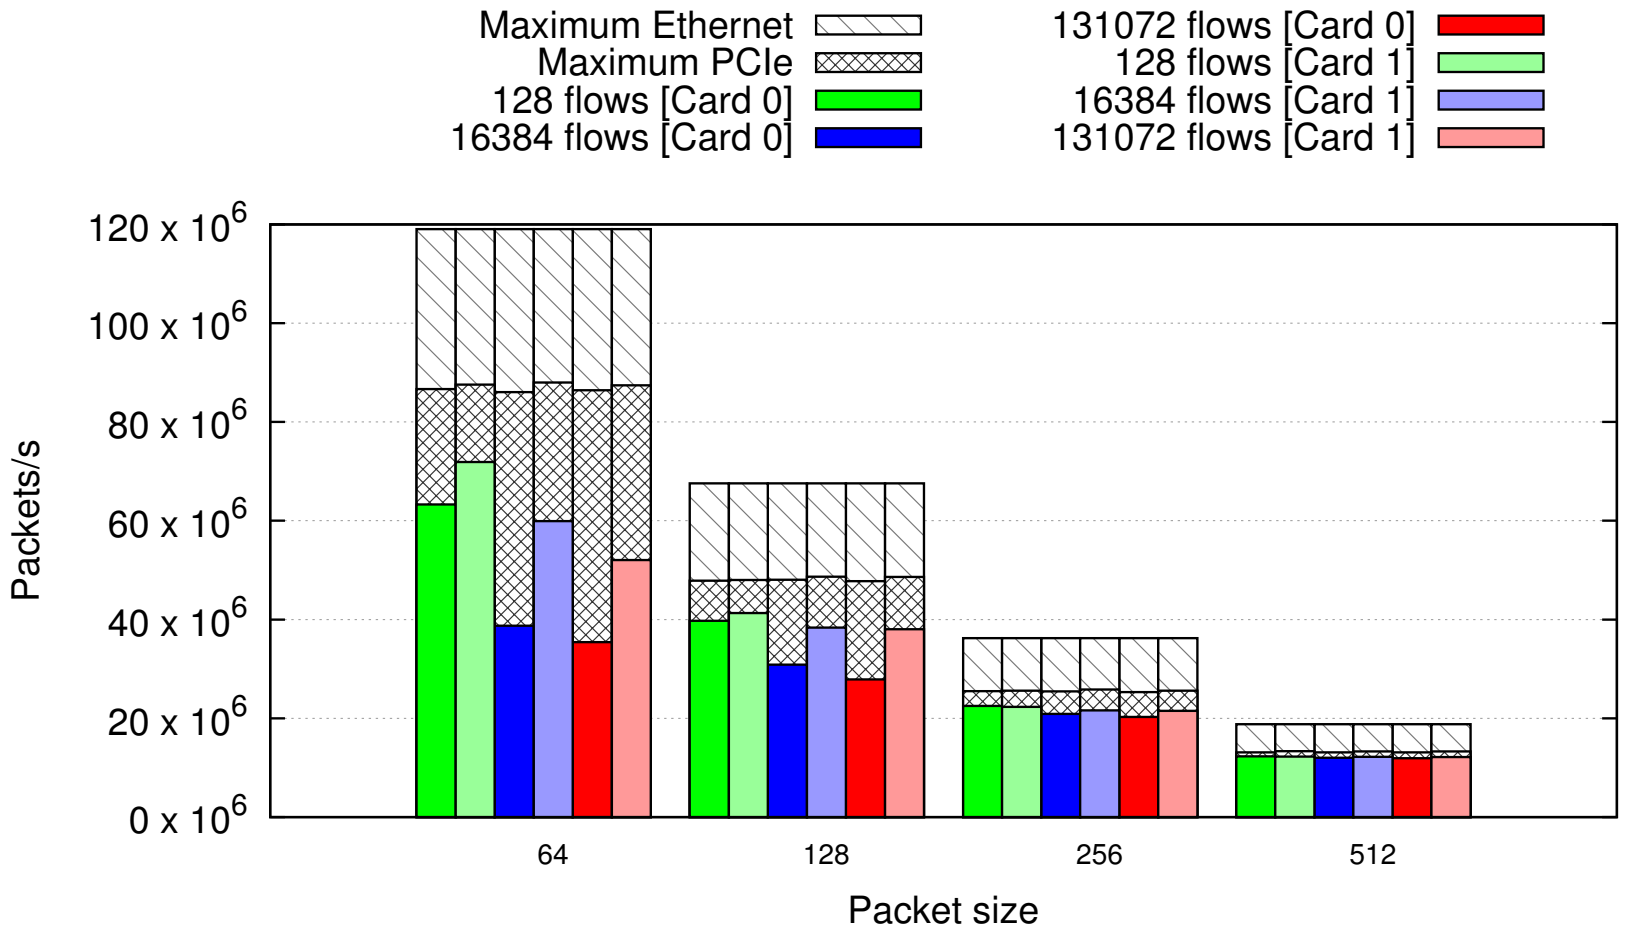

Full packet processing performance in packets/s.

Basic Performance on Full Packets

- PCI-Express limit is reached only for the longest packets
- NUMA architecture affects the performance
- Number of flows has large impact mainly on short packets

Hardware Accelerated Performance

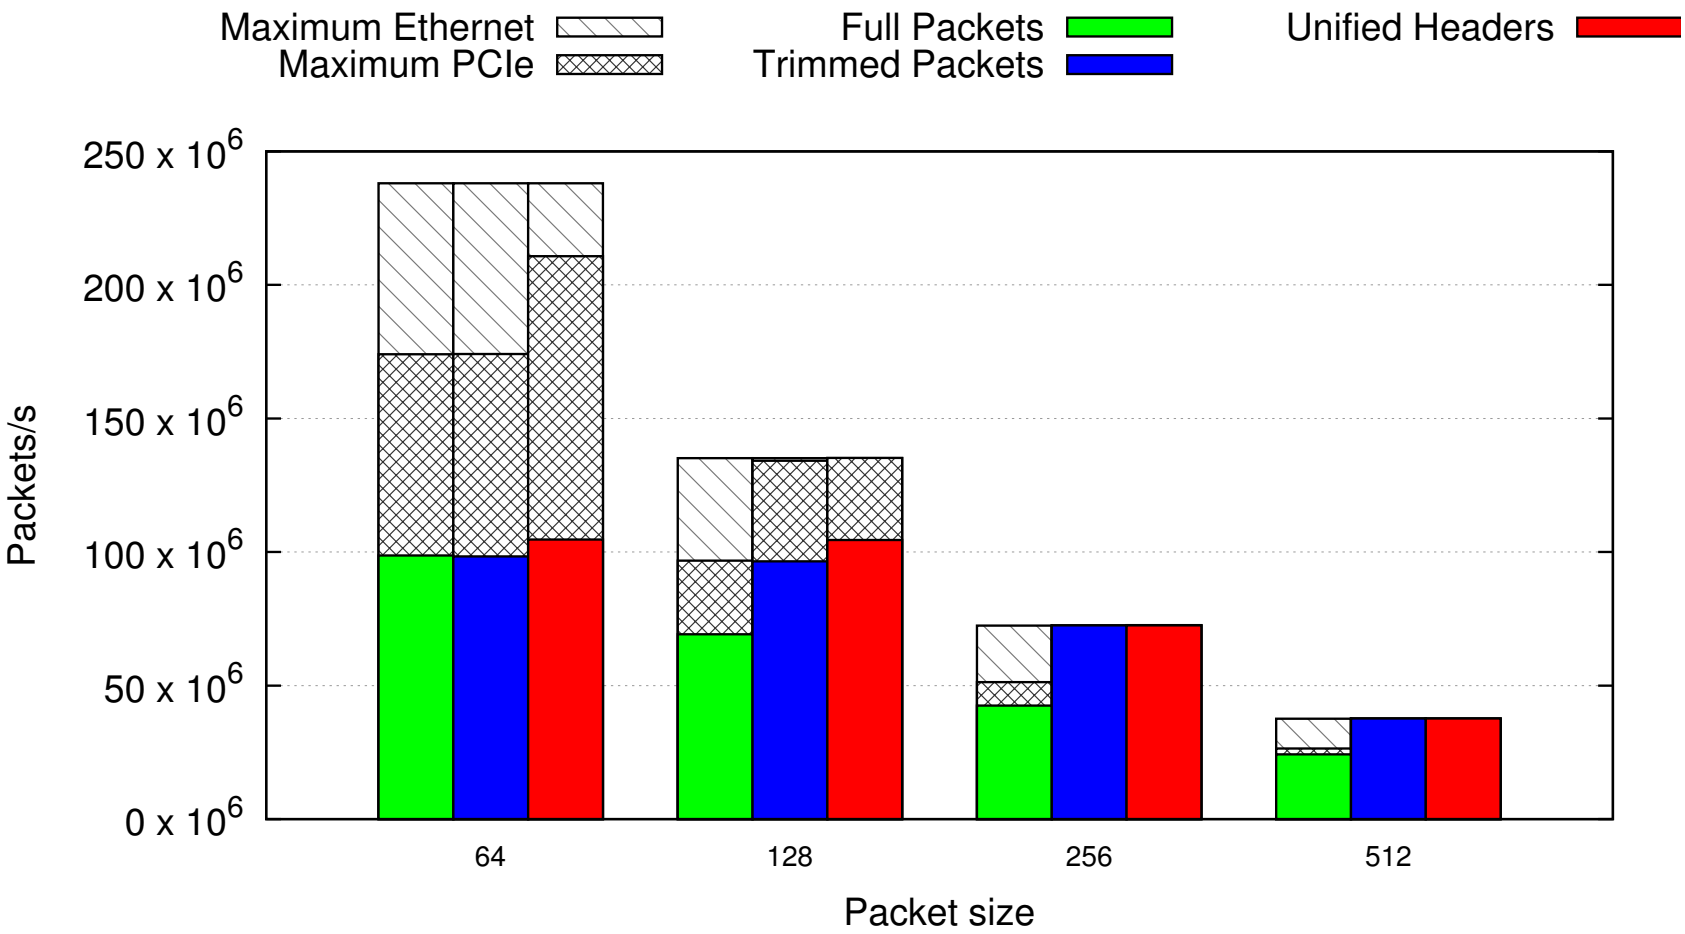

Packet processing performance in packets/s for 2^{18} flows.

Hardware Accelerated Performance

- Packet trimming ad Unified Headers help significantly
- Full throughput monitoring for 256 B and longer packets
- Packet trimming and Unified Headers solve the problem of PCI-Express throughput

Impact of CPU Choice

- Comparison of the E5-2670 with E5-2620
- Only 6 cores, 2 Ghz frequency

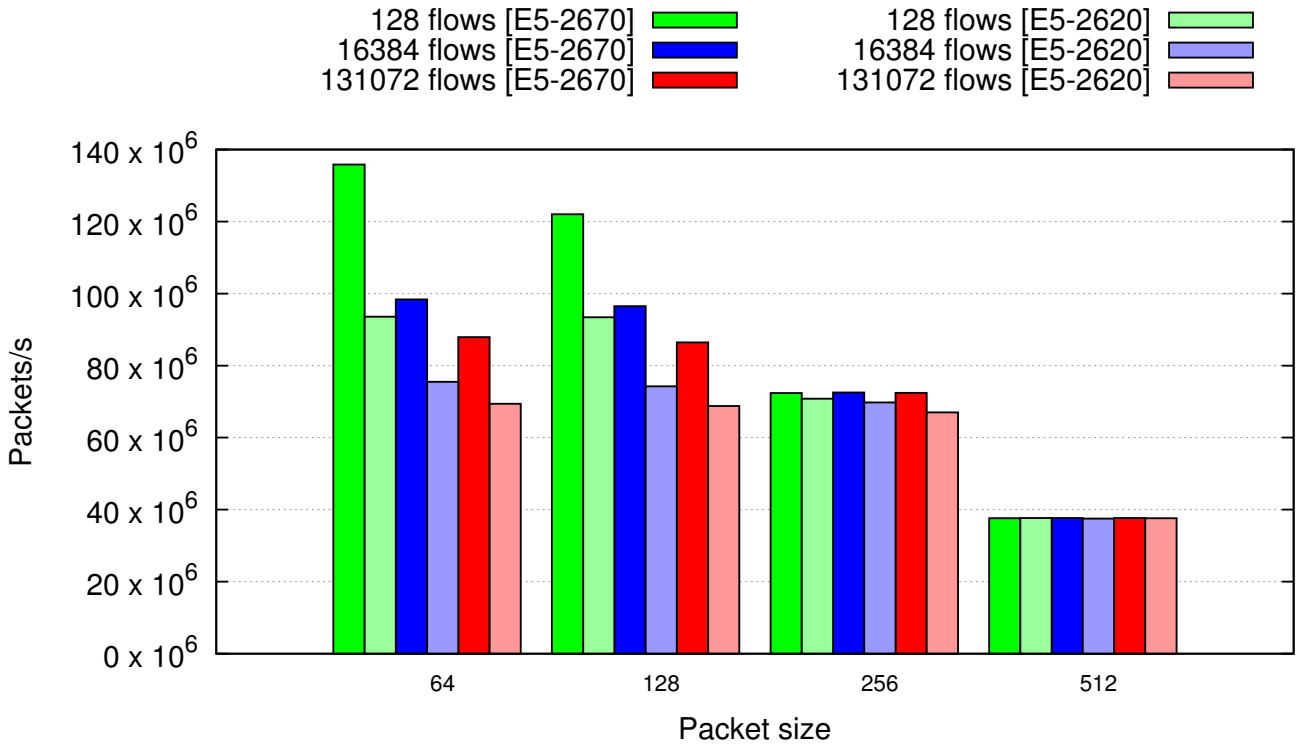

Comparison for two different CPUs on trimmed packets

Impact of CPU Choice

- Faster CPU helps greatly for less flows
- Large number of flows has greater impact on memory bus utilization
- Both CPU are doing well for longer packets

Conclusions

& Future Work

Conclusions

- It is possible to monitor 16× 10 G links in one 2U box
- Hardware acceleration can significantly help to improve the performance
- PCI-Express can be limiting for commodity cards
- NUMA architecture must be taken into consideration
- Number of flows has significant impact on performance

Future Work

- Test the performance with COMBO-100G cards (PCIe gen3 x16)
- High-speed expriment with application flow measurements
- Build better framework for measurements
  - Different flow lengths
  - More complex packets and flows
  - Different packets in one flow

Thanks for your attention

High-Density Network Flow Monitoring

IM2015

Petr Velan
petr.velan@cesnet.cz