cesnet

CONFIGURABLE FPGA PACKET PARSER FOR TERABIT NETWORKS WITH GUARANTEED WIRE-SPEED THROUGHPUT

Jakub Cabal et al.
CESNET

2018/02/27
FPGA, Monterey, USA

INTRODUCTION

- **Packet parsing**
  - It is among basic operations that are performed at all points of each network.
- **Different requirements on parser performance:**
  - **Latency** – critical for high-frequency trading
  - **Area** – important especially for embedded devices
  - **Throughput** – backbone networks and data centers

OUR PREVIOUS PARSER

- **We have a lot of experience with parsers for FPGA.**
- **Previous generation of our packet parser:**
  - Introduced by Kekely et al. in ANCS 2012.
  - Raw throughput over 100Gbps, but not wire-speed.
  - Processing of at most one packet per clock cycle.
  - Modular structure for simple addition of new protocols.

NEXT GEN PARSER

- **We propose a next generation of packet parser for FPGA.**
- **The contribution of our next-gen parser:**
  - The raw throughput of over 1Tbps.
  - Processing of multiple packets per clock cycle that leads to sufficient effective throughput for wire-speed processing.
  - The modular structure supporting automatic generation of HDL implementation from a high-level P4 description.

PARSER STRUCTURE

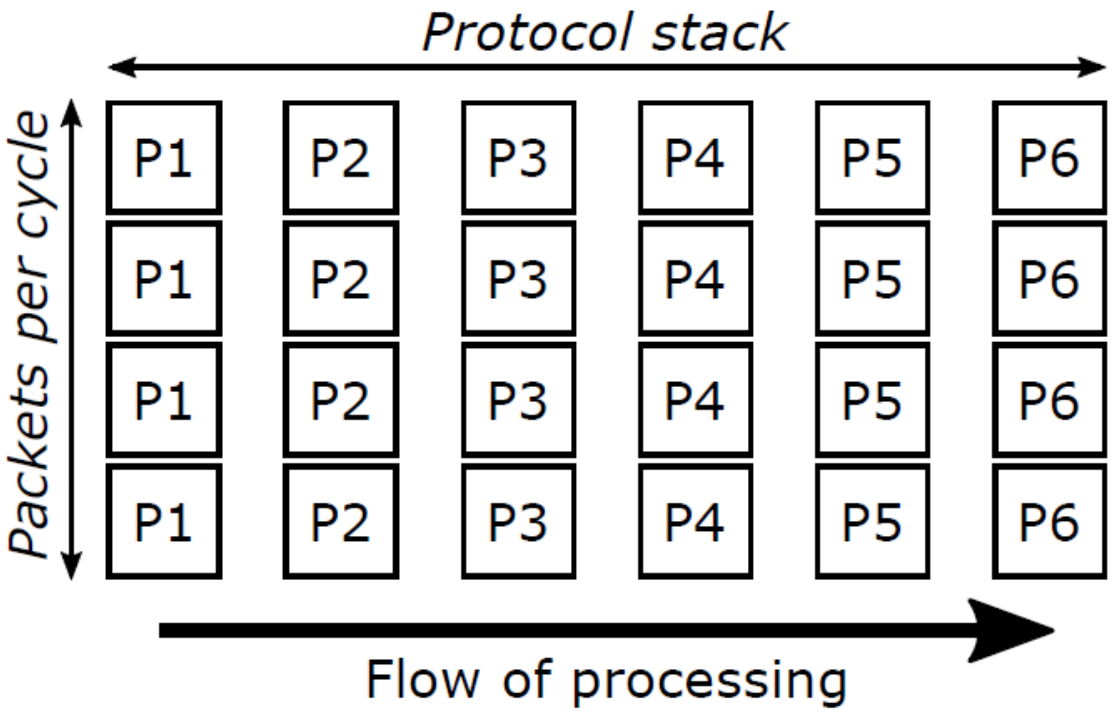

- Packet parsing is divided into multiple simple analyzers.
- Each analyzer can process only one packet per clock cycle and parse only one protocol header.

SIMPLE PROTOCOL ANALYZERS

- **The input information necessary to parse a single protocol header consists of:**
  - Packet data from the data bus.
  - Offset of the current word counted from the last start of packet.
  - Type of the expected protocol header.
  - Offset of the current header start.

SIMPLE PROTOCOL ANALYZERS

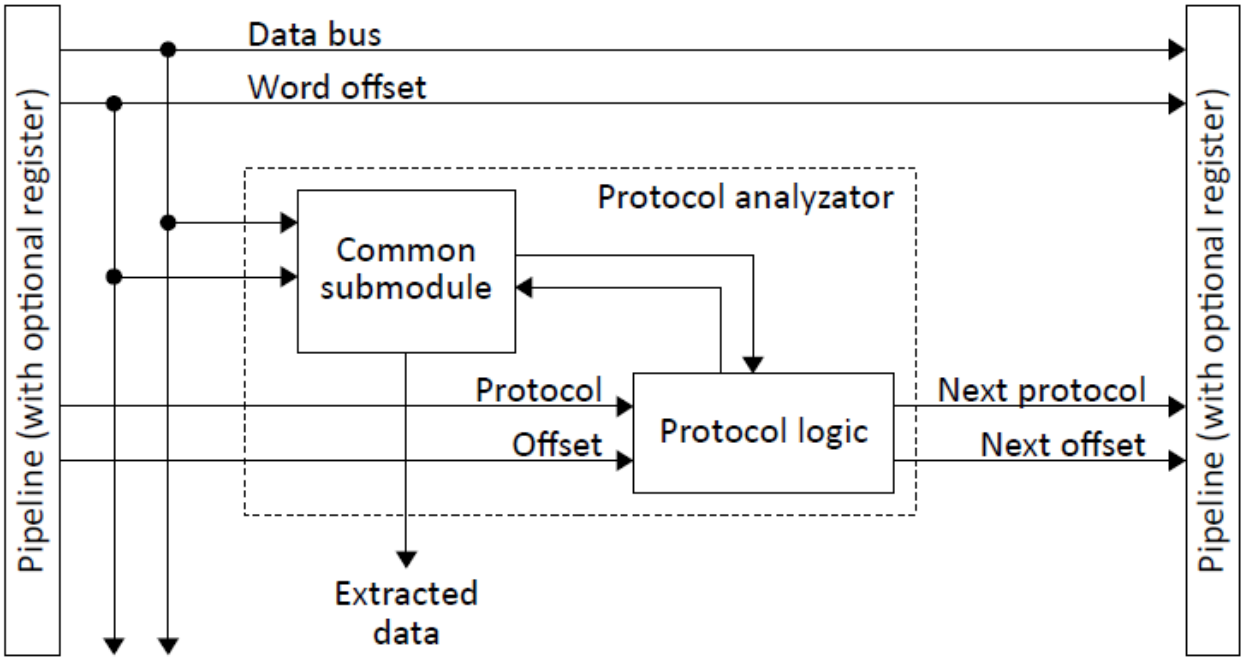

- **The output information includes:**
  - Type of the next expected protocol.
  - Offset of the next header start.
  - Extracted data from the current protocol header.

PARSER TOP-LEVEL DESIGN

- **Our parser uses pipelining to achieve high working frequency and thus high throughput.**
  - The input pipeline stage generates the initial control data for the first protocol analyzer stage.
  - Each pipeline stage contains optional registers and a logic for the distribution of control data.

PARSER TOP-LEVEL DESIGN

- The example of parser pipeline support of Ethernet, VLAN, IPv4 and TCP parsing in up to 4 packets per clock cycle.

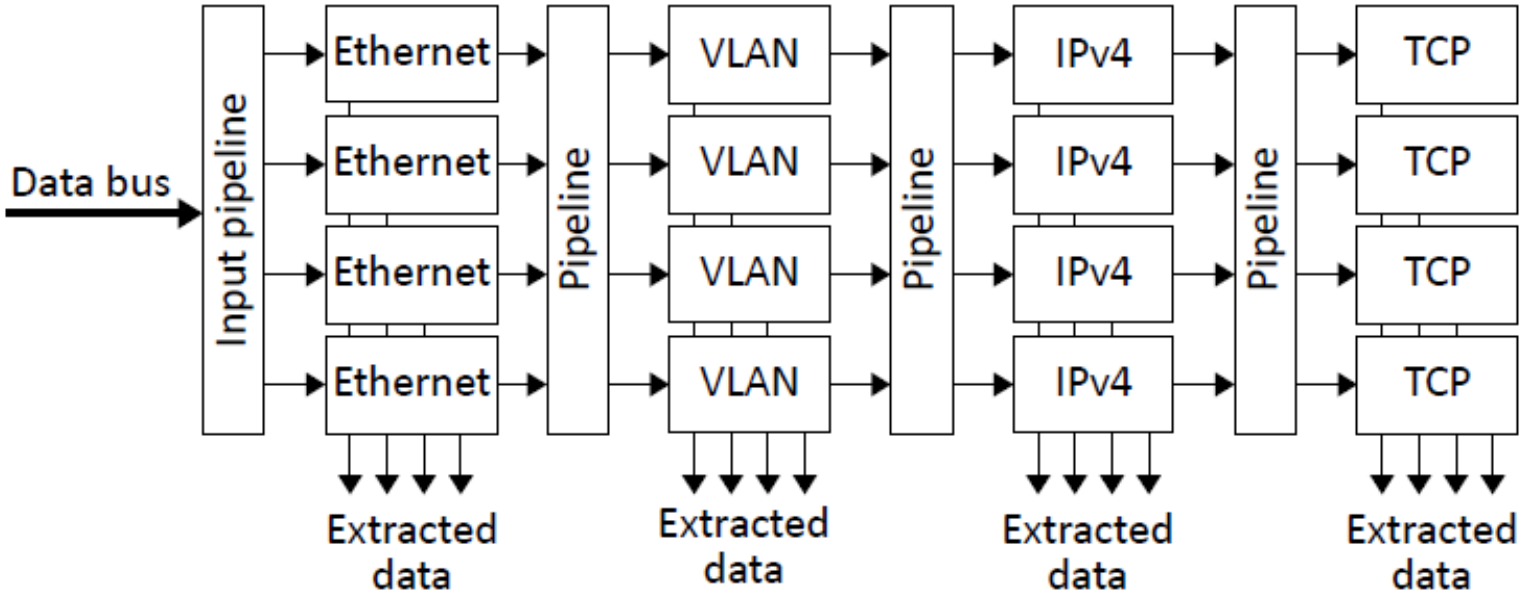

HDL GENERATION FROM P4 DESCRIPTION

- **The transformation process from $P4_{14}$ to VHDL was introduced by Benáček et al. in FCCM 2016.**
  - The architecture of our parser is designed to be compatible with the parser generator by Benaček.
  - We integrated our parser into the generator, enabling it to create parsers from a P4 description with much higher effective throughput.

RESULTS CONFIGURATIONS

- **Results of the parser for 2 protocol stacks:**
  - Full – Ethernet, 4xVLAN, 4xMPLS, IPv4/v6 (+2xExt), TCP/UDP
  - Simple L2 – Ethernet, IPv4/v6 (+2xExt), TCP/UDP
- **Variable parameters for testing:**
  - Data width - powers of 2 from 512 to 4096 bits
  - Pipelining - every possible configuration of pipeline placement for simple L2 protocol stack

RESULTS CONFIGURATIONS

- **Creates very large solution space:**
  - e.g. 2^9 configurations of the simple L2 parser.
  - e.g. 2^{17} configurations of the Full parser (measured only 1% of configurations, random).
- **Test results for:**
  - Xilinx Virtex-7 XCVH870T FPGA
  - Xilinx UltraScale+ XCVU7P FPGA

RESULTS
MEASURED DATA

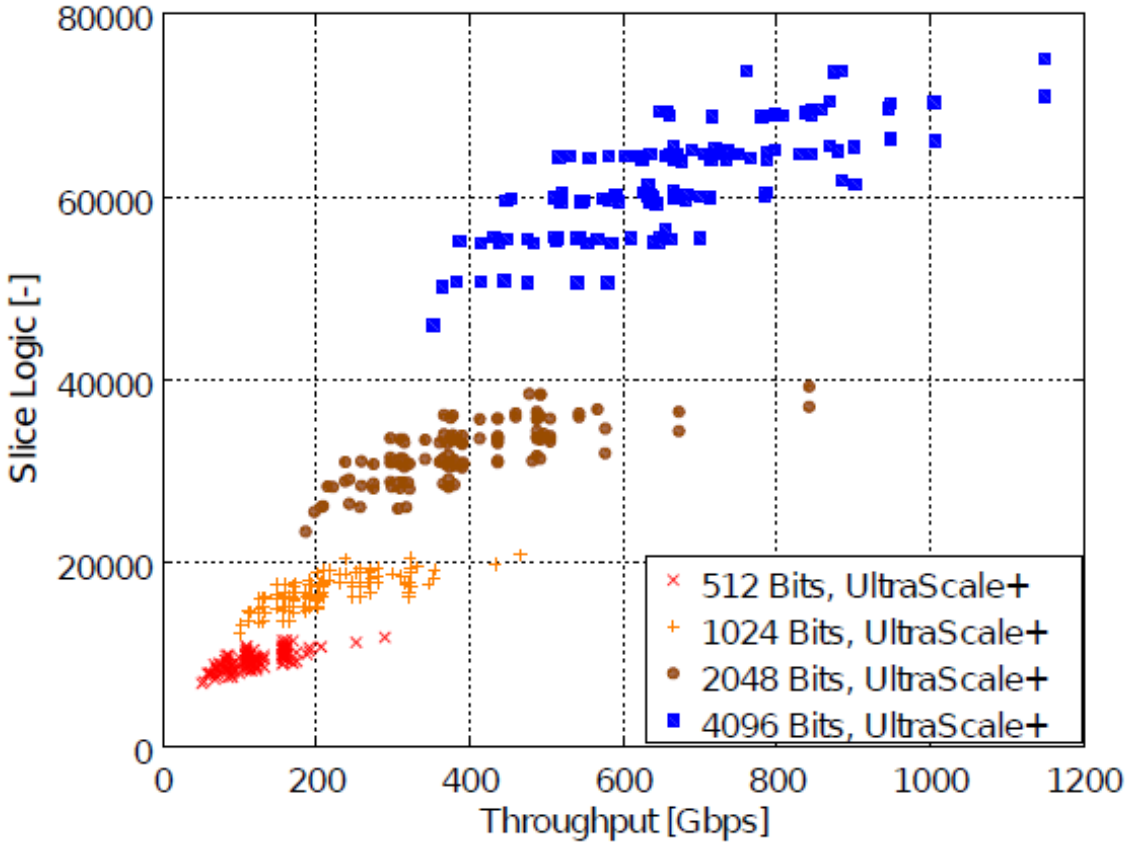

- **All solutions for simple L2 protocol stack and Xilinx UltraScale+ FPGA**
  - Minimal utilization for 1Tbps is 70k slice logic.

RESULTS
MEASURED DATA

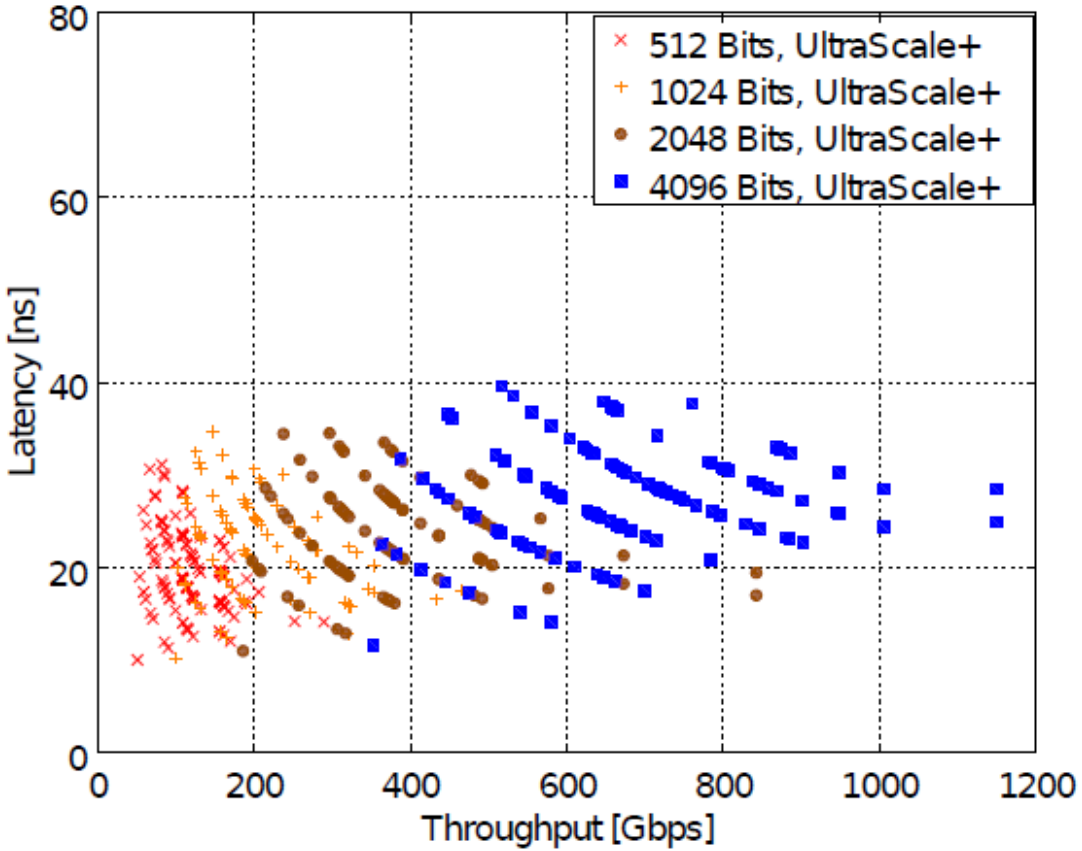

- **All solutions for simple L2 protocol stack and Xilinx UltraScale+ FPGA**
  - Minimal latency for 1Tbps is 25ns.

RESULTS
MEASURED DATA

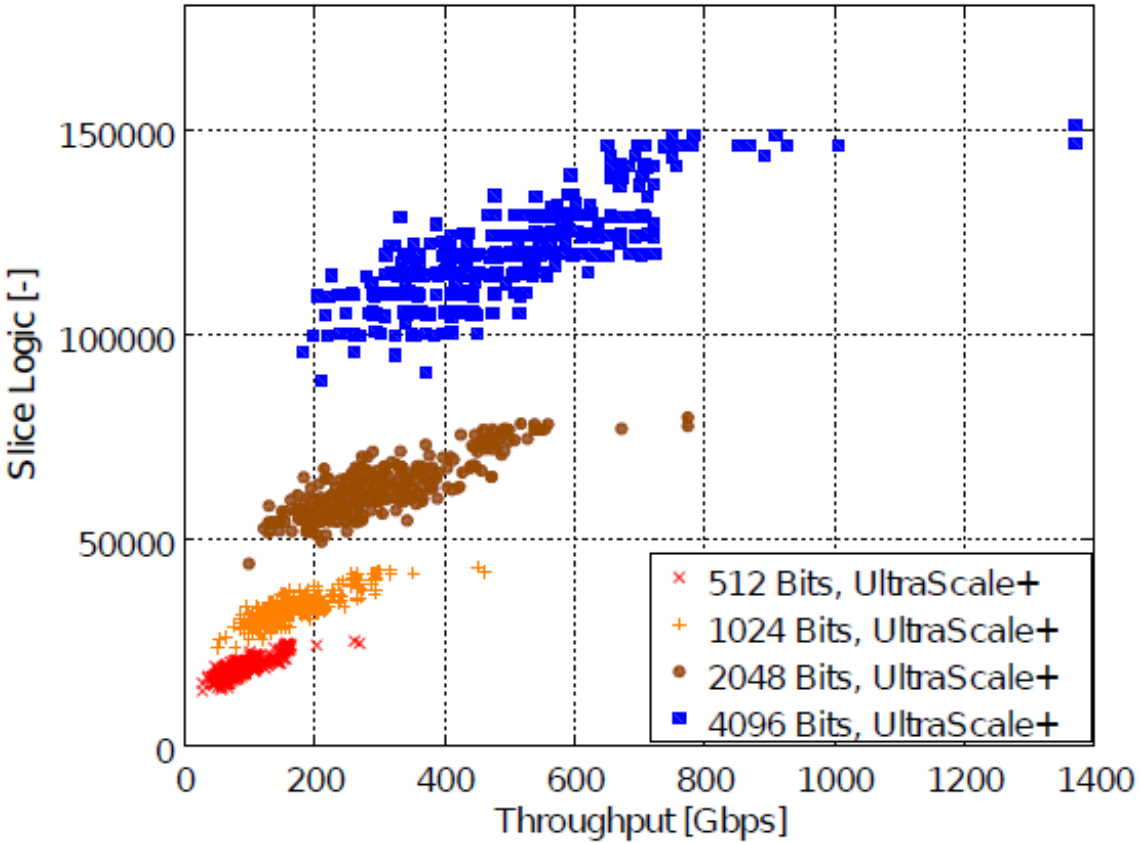

- **Solutions for Full protocol stack and UltraScale+ FPGA**
  - Minimal utilization for 1Tbps is 150k slice logic.

RESULTS
MEASURED DATA

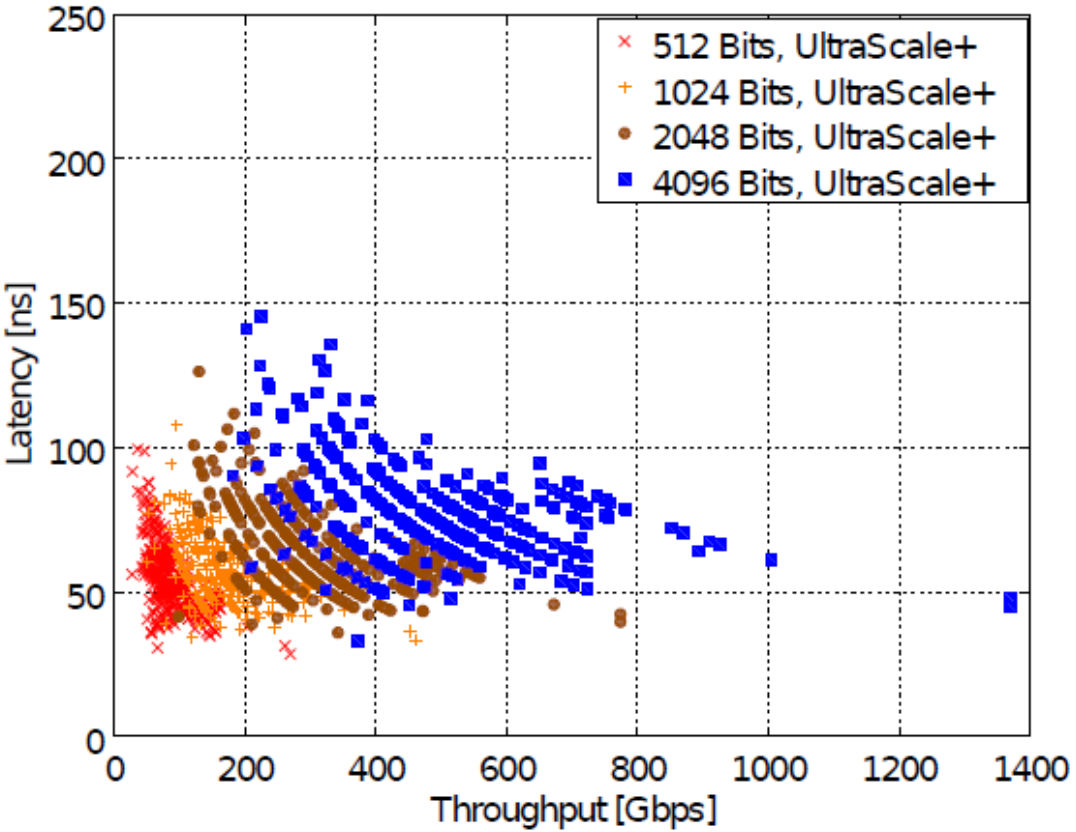

- **Solutions for Full protocol stack and UltraScale+ FPGA**
  - Minimal latency for 1Tbps is 70ns.

RESULTS
MEASURED DATA

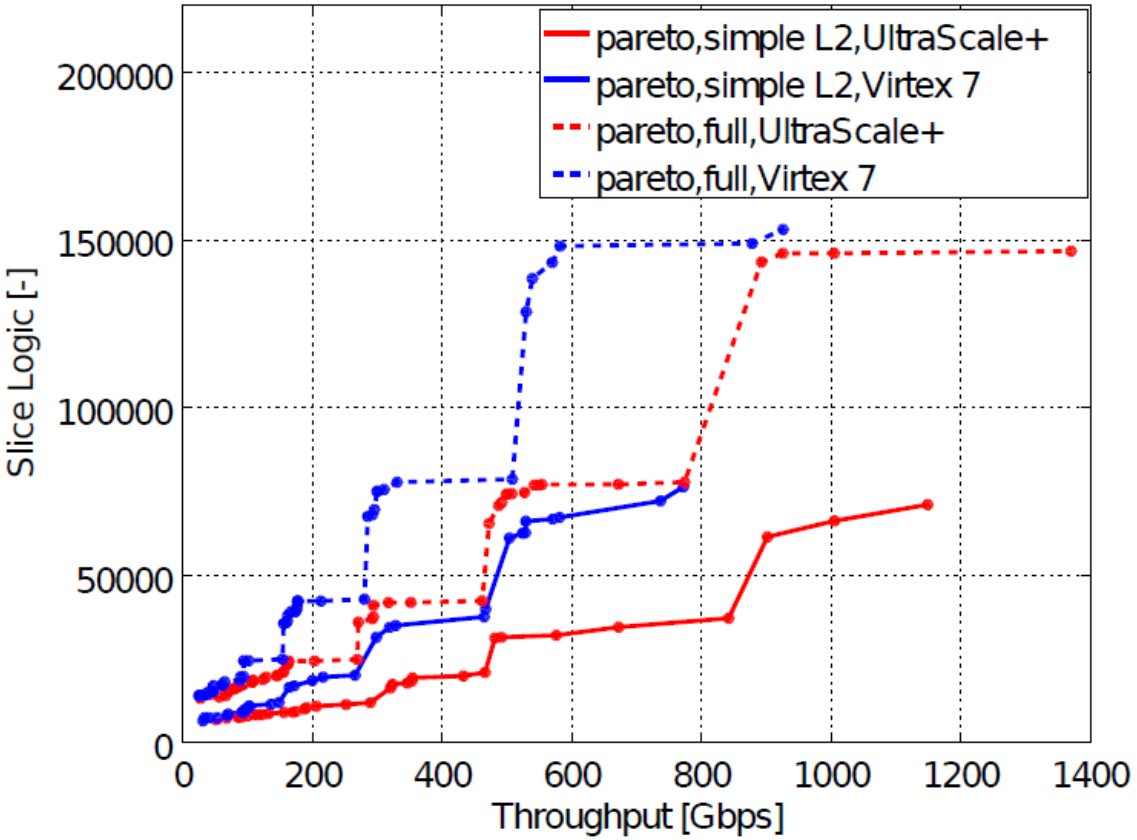

- **Pareto sets for different protocol stacks and FPGA**
  - Virtex 7 FPGA resources utilization on same throughput is twice bigger as UltraScale+ FPGA.

RESULTS
MEASURED DATA

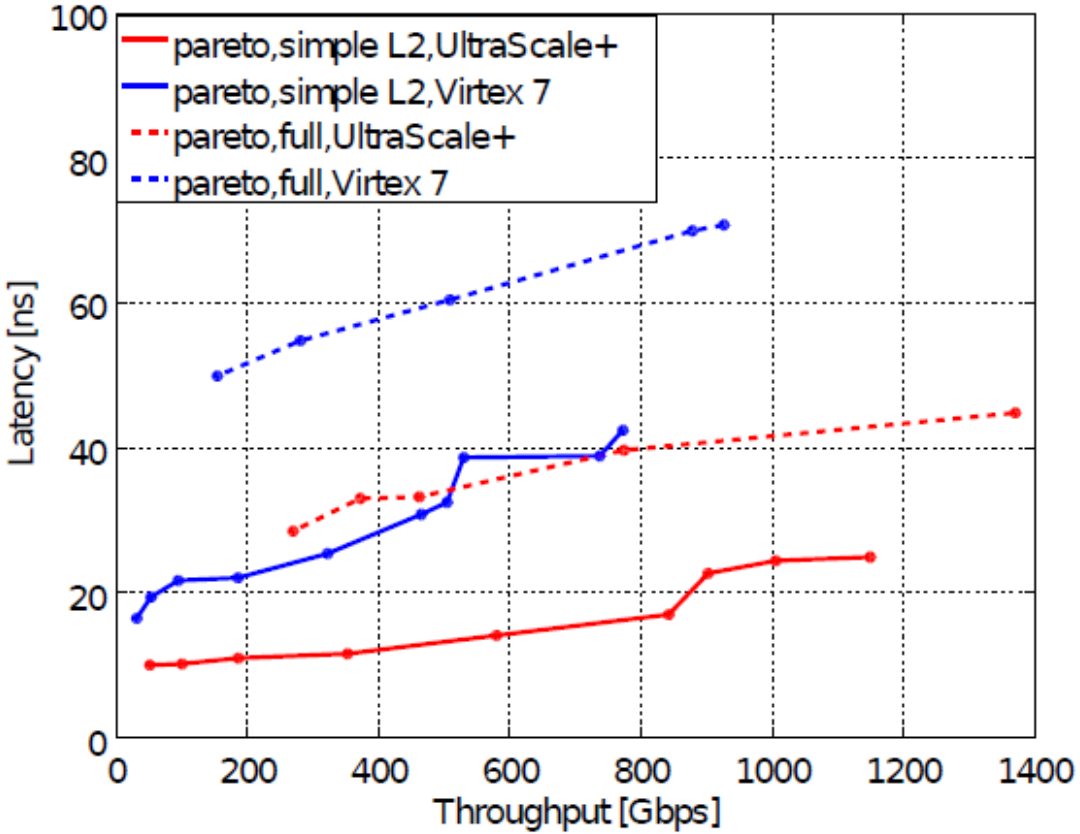

- **Pareto sets for different protocol stacks and FPGA**
  - UltraScale+ FPGA is faster and the latency is lower.

cesnet

OUR PARSER IN REAL USAGE

- **The Parser is used in several projects for our Network Acceleration Card.**
  - Xilinx UltraScale+ FPGA
  - 2x 100G Ethernet
  - 2x PCIe gen3 x16
  - low profile, half length

CONCLUSIONS AND CONTRIBUTIONS

- **Introduction of novel packet parser for FPGA:**
  - Throughput in the range from 10 to over **1000Gbps**.
  - **Wire-speed throughput on all packet lengths.**
  - Latency of 400Gbps packet parsing up to L4 is under 40ns.
  - Low resources usage: 2.4% for 200Gbps, 4.7% for 400Gbps
  - Configurable for variety of different applications.
  - **Supporting automatic generation of HDL implementation from a high-level P4 description.**

cesnet

THANK YOU FOR YOUR ATTENTION.