AROBS Polska

COTS FPGA technology for On-Board Switching

Dawid Linowski, Szymon Kałużyński, Tomasz Rybak (presenter)

6th SEFUW, Space FPGA Users Workshop

ESA/ESTEC

27th March 2025

Agenda

- **Introduction and objectives**
- Achronix FPGA technology
- Ethernet layer 2 switch design
- Tests in hardware
- AMD Versal NoC
- Summary

AROBS Polska – space electronics, FPGA & software

Established in:
2016

Number of employees:
20

Headquarters:
Gdańsk, Poland

Office Island in Gdansk – AROBS Polska headquarters

Areas of expertise

FLASH MEMORY MODULES

Control electronics for mechanisms and instruments

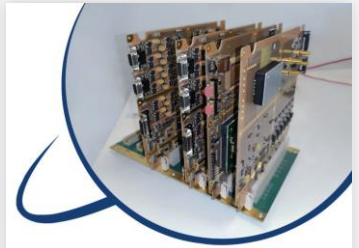

QUANTUM ENTANGLEMENT CONTROLLERS

Payload computers, data processing units

Flash Memory Modules

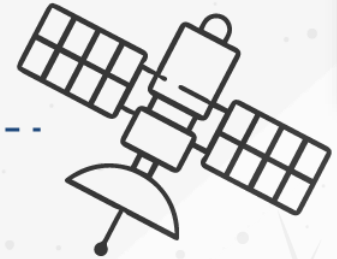

Quantum and optical communication

Electrical Ground Support Equipment

TESTING CABINET

CCSDS STANDARDISED RANGING FOR OPTICAL COMMUNICATION TERMINALS

Project’s objectives

COTS FPGA technology for On-Board Switching

- Developed under ESA contract No. 4000138940/22/NL/AF, in the scope of **ScyLight** Strategic Programme Line within **ARTES 4.0**
- Project in the frame of ESA High thRoughput Optical Network project (**HydRON**)
- Target **TRL-4**
- Project duration: 05.09.2022 – ongoing (closing Final Review).

Goals:

- Technology assessment of **Achronix FPGA**-based solution for **Ethernet layer 2 switch**
- Data rate up to **100 Gbps per single channel** (simulation and HW validation).
- Transferring achieved solution to **AMD/Xilinx Versal** technology (just simulation).
- **Comparison** of simulation performances for both technologies.

Why such a project?

HydRON:

- Part of ScyLight programme, which focuses on **optical communications**, photonics and quantum communication.
- Aims to demonstrate the world's first (all) optical multi-orbit transport **network at Terabit/sec capacity in space**
- Extends terrestial **fibre-based** networks into space -> **„Fibre in the sky".**

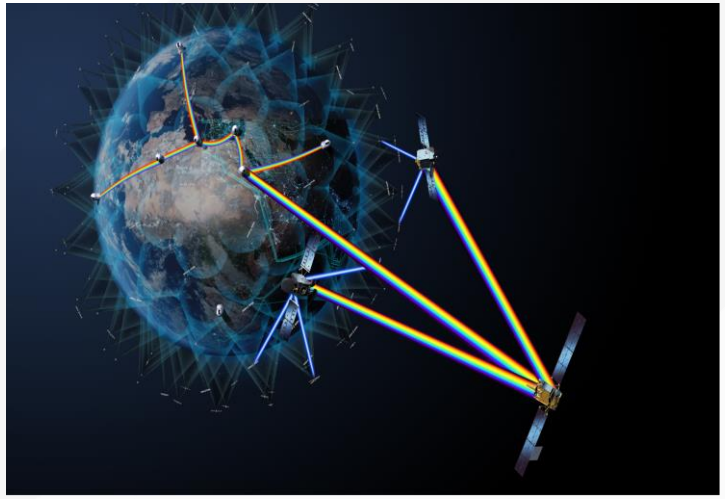

Currently used Rad-Hard By Design (RHBD) FPGA devices do not allow to implement the performance required by this programme -> **potential COTS solutions shall be investigated**.

Technology assessment

- AMD Versal FPGA:
  - built in TSMC 7 nm FinFET process technology
  - promising radiation test results -> similar behaviour expected from other FPGAs based on the same process.

- **Achronix Speedster7t** family:
  - based on the same technology process
  - incorporating **Network-on-Chip (NoC)** for high data rate connections inside the fabric
  - mainly used in data centres for high throughput, Ethernet switch-based solutions up to 400 Gbps per port.

Agenda

- Introduction and objectives
- **Achronix FPGA technology**
- Ethernet layer 2 switch design
- Tests in hardware
- AMD Versal NoC
- Summary

Achronix FPGA

Speedster7t AC7t1500 family:

- **692K 6-input LUTs** (1,522K 4-input LUT equivalents)
- **1.382M flip-flops**
- 2-dimensional **NoC** capable of > **20 Tbps**
- 32 high-speed SerDes transceivers, each up to **112 Gbps (PAM4)** and 56 Gbps (NRZ).
- Two **hard-IP Ethernet MACs**, with support up to 400 Gbps (each using 8 of 32 SerDes).
- Integrated **DDR4/5** and **GDDR6** controllers.
- Integrated **PCIe Gen5** controller (uses remaining 16 SerDes).

Speedster7t FPGA family

Part Number/Name	7t800	7t1500	7t1550
LUT count (K 6LUTs)	326	692	646
Inline Cryptography	Yes	No	Yes
MLP: Multi-fracturable MAC array	864	2,560	
LRAM (2.3k)	864	2,560	
BRAM (73.7k)	1,152	2,560	
Memory	86 Mb	195 Mb	
ML TOps: int8 or block bfloat16	20.5	61	
SerDes 112G/224G	24	32	
DDR4/5	1 DDR5 ×64 (w/ECC)	1 DDR4 ×64 (w/ECC)	
High Bandwidth Memory Channels	6 GDDR6 (1.5 Tbps, w/ECC)	16 GDDR6 (4Tbps)	
PCI Express	One ×16 PCie Gen5	One ×8, one ×16 PCie Gen5	
Ethernet	8 lanes, 2×400G/ 8×100G	16 lanes, 4×400G/ 16×100G	
2D Network-on-Chip BW – Tbps	12	20	

Achronix NoC overview

- Extremely high-speed connection between:
  - **fabric logic and interfaces** around the periphery (GDDR6, DDR4/5, Ethernet, PCIe)
  - **fabric logic within FPGA itself**.
- Cross-sectional bidirectional bandwidth of **20 Tbps**.
- Running with the main clock of **2 GHz**.
- Consists of two main parts:
  - Peripheral **ring** around the fabric with all the IP interfaces.
  - The **rows** and **columns** that run over the top of the FPGA fabric.
- **Wide (256-bit), high-speed buses without using the fabric routing resources!**

Achronix NoC overview

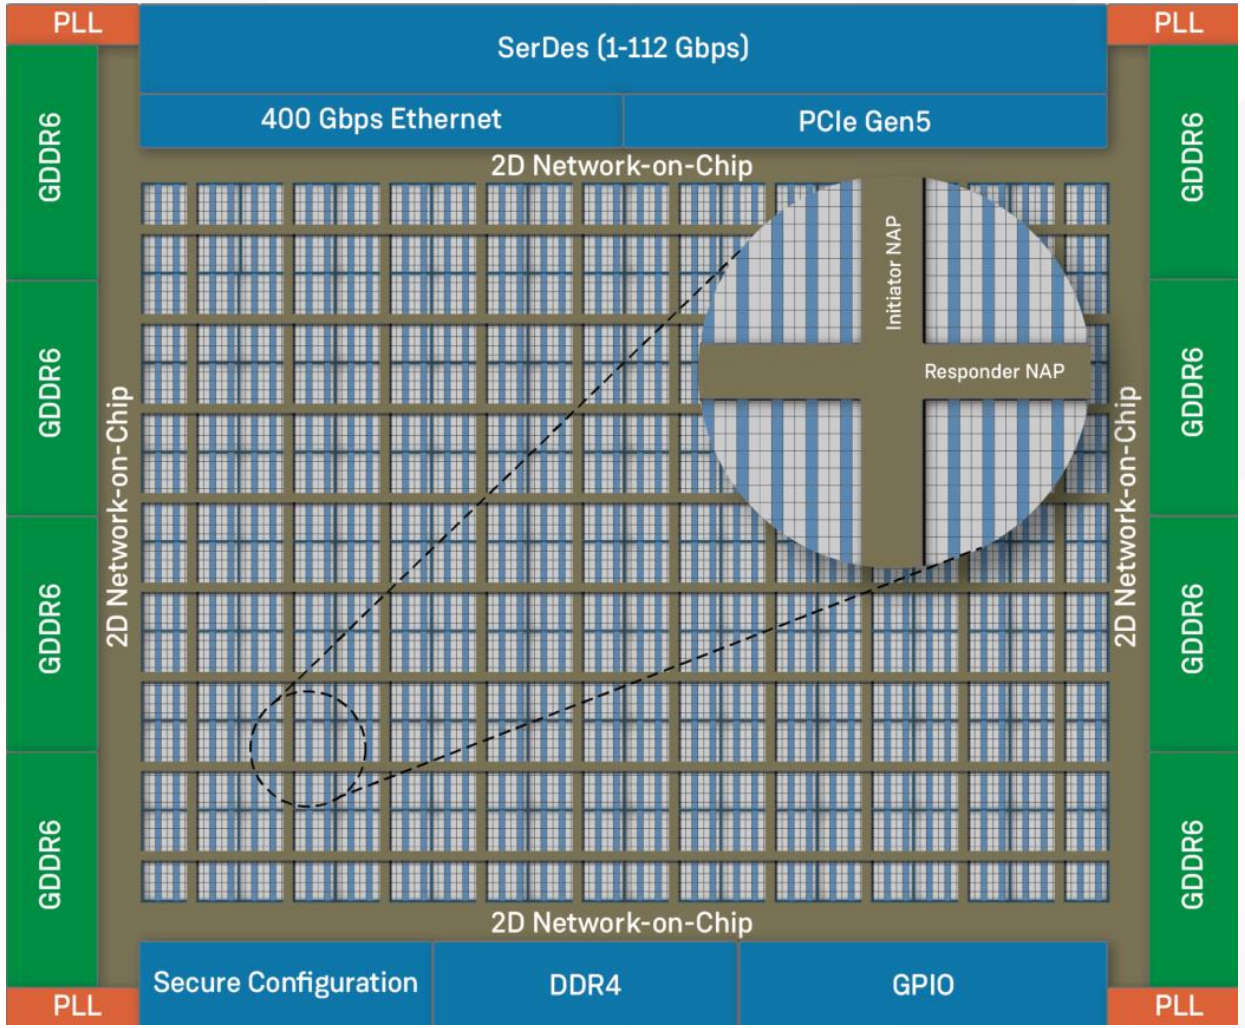

Accessing the Achronix NoC - NAP

- NoC Access Points (**NAPs**) instantiated in user logic to connect to the rows and columns of the 2D NoC.
- NAP types:
  - AXI Responder NAP – allows connection to a row of the NoC
  - AXI Initiator NAP – allows connection to a column of the NoC
  - Horizontal NAP – allows connection along one row of the NoC
  - Vertical NAP – allows data streaming along one column of the NoC
  - Ethernet NAP – allows data streaming of Eternet packets along NoC columns.
- NAP handles required Clock Domain Crossing.

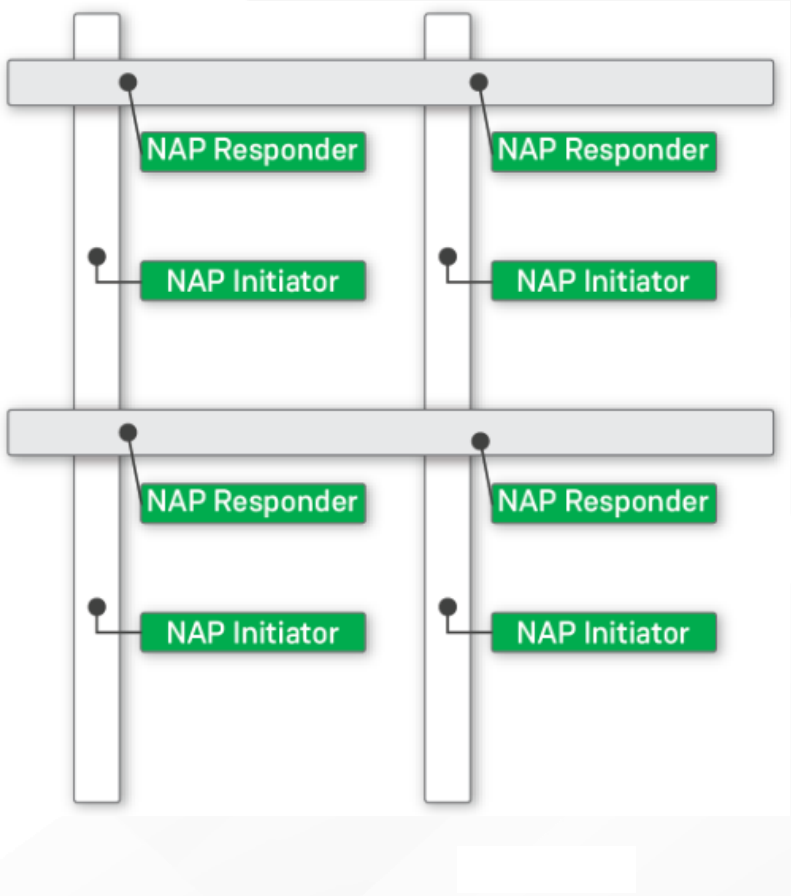

Achronix NoC limitations

- Connection to NAP follows AXI4 standard, but bursts are limited to 16 transfers.
- Both Ethernet modules are connected only to certain NoC columns (two columns per Ethernet module).
- No direct connection between rows or columns - ring usage necessary.
- NoC performance results directly from the main NoC clock frequency. 2 GHz is used only in the highest speed grade (C1).
- NoC transactions on the same column or row are point to point - no broadcast option.
- Multiple NAPs along a single column or row in the FPGA can potentially create a traffic congestion. Expected traffic patterns in the design should be evaluated beforehand and NAP locations selected to spread the transaction traffic across several rows or columns when possible.

Achronix Evaluation Board

BittWare VectorPath® S7t-VG6 accelerator card

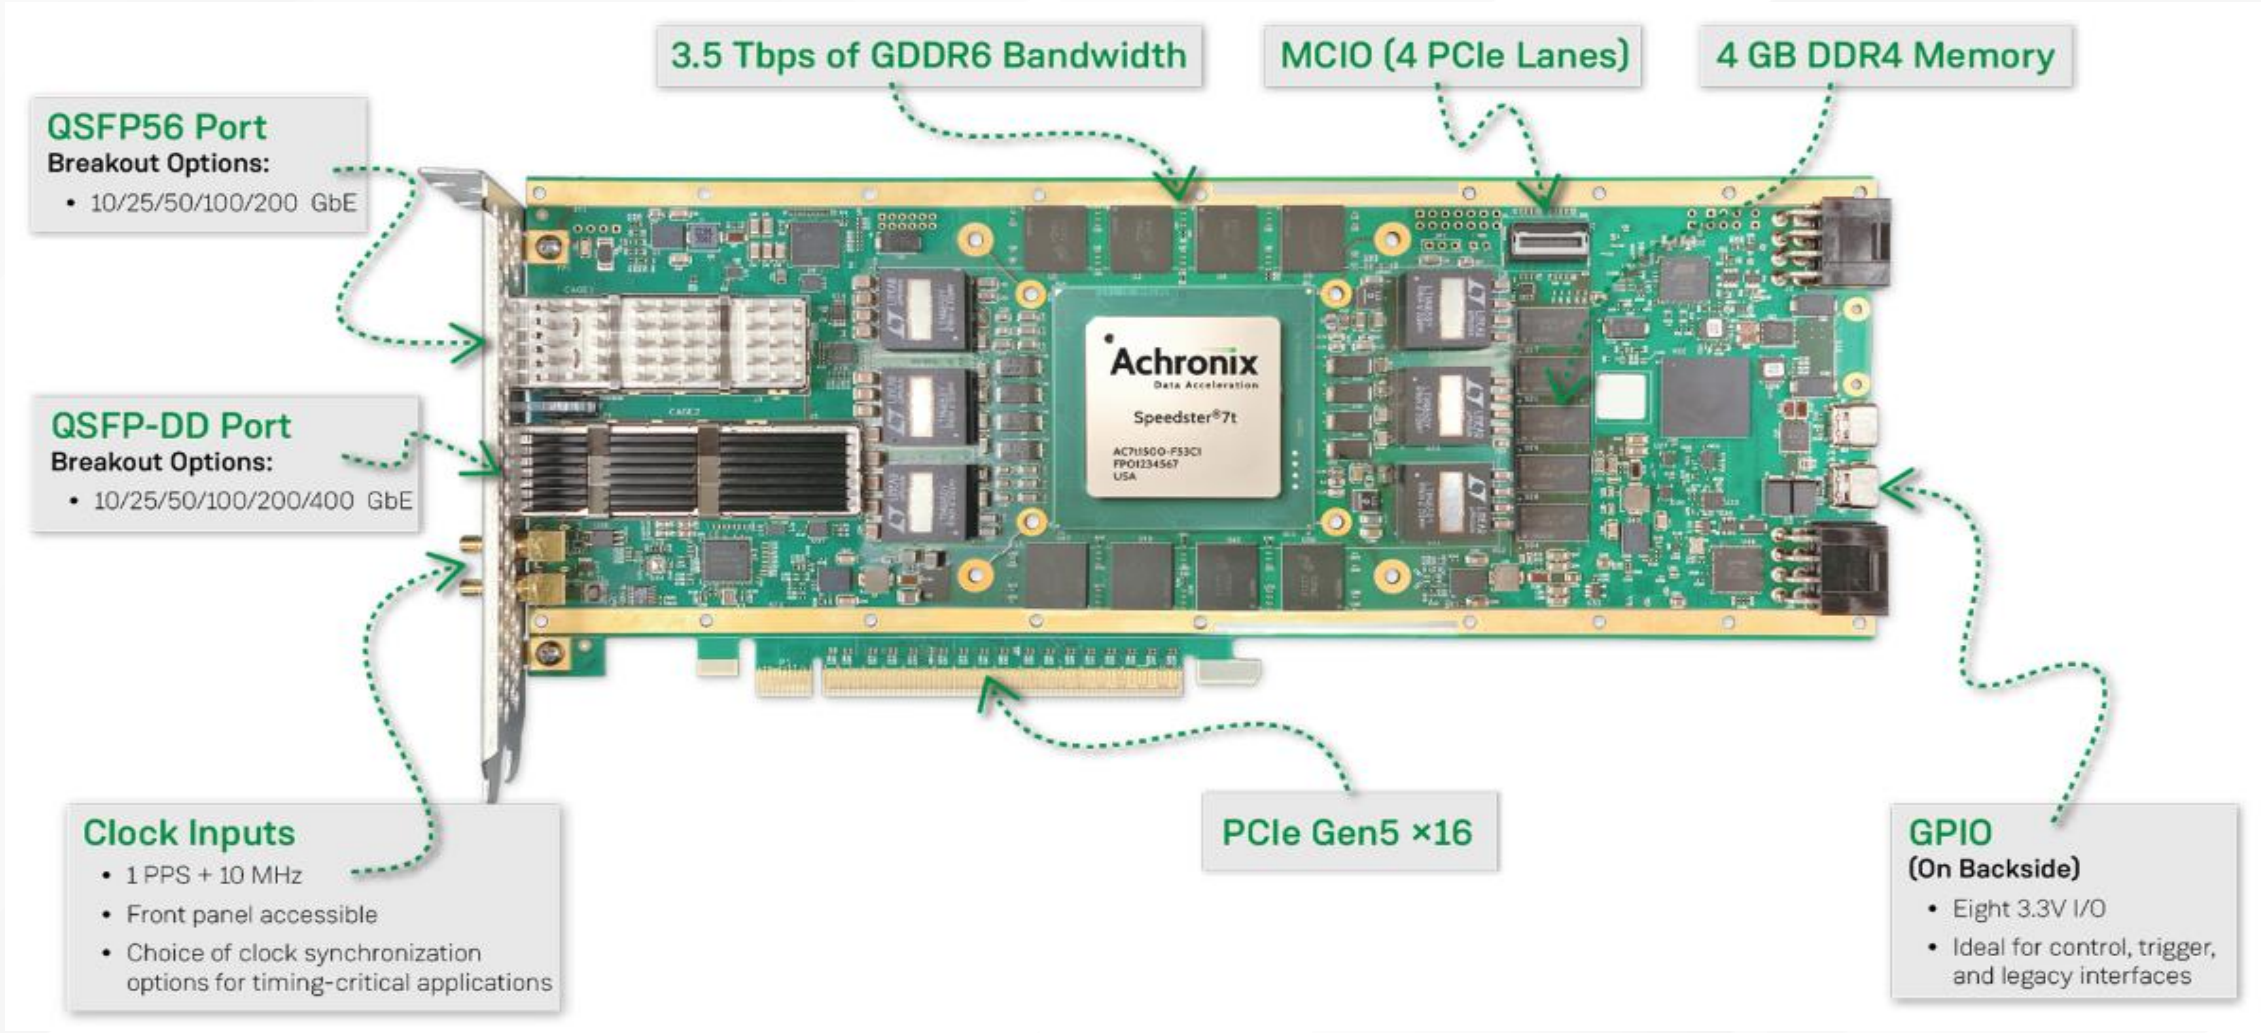

Agenda

- Introduction and objectives
- Achronix FPGA technology
- **Ethernet layer 2 switch design**
- Tests in hardware
- AMD Versal NoC
- Summary

Ethernet layer 2 switch - architecture

Composed of:

- COMMS CTRL – responsible for communication between the PC and Achronix FPGA
- SYSTEM CTRL – manages system control and status operations
- SWITCH PORT CTRL – realizes Ethernet packet switching between Ethernet ports through NoC.
- EMULATED PORT CTRL – Ethernet subsystem emulation; also Ethernet packet generation and checking.

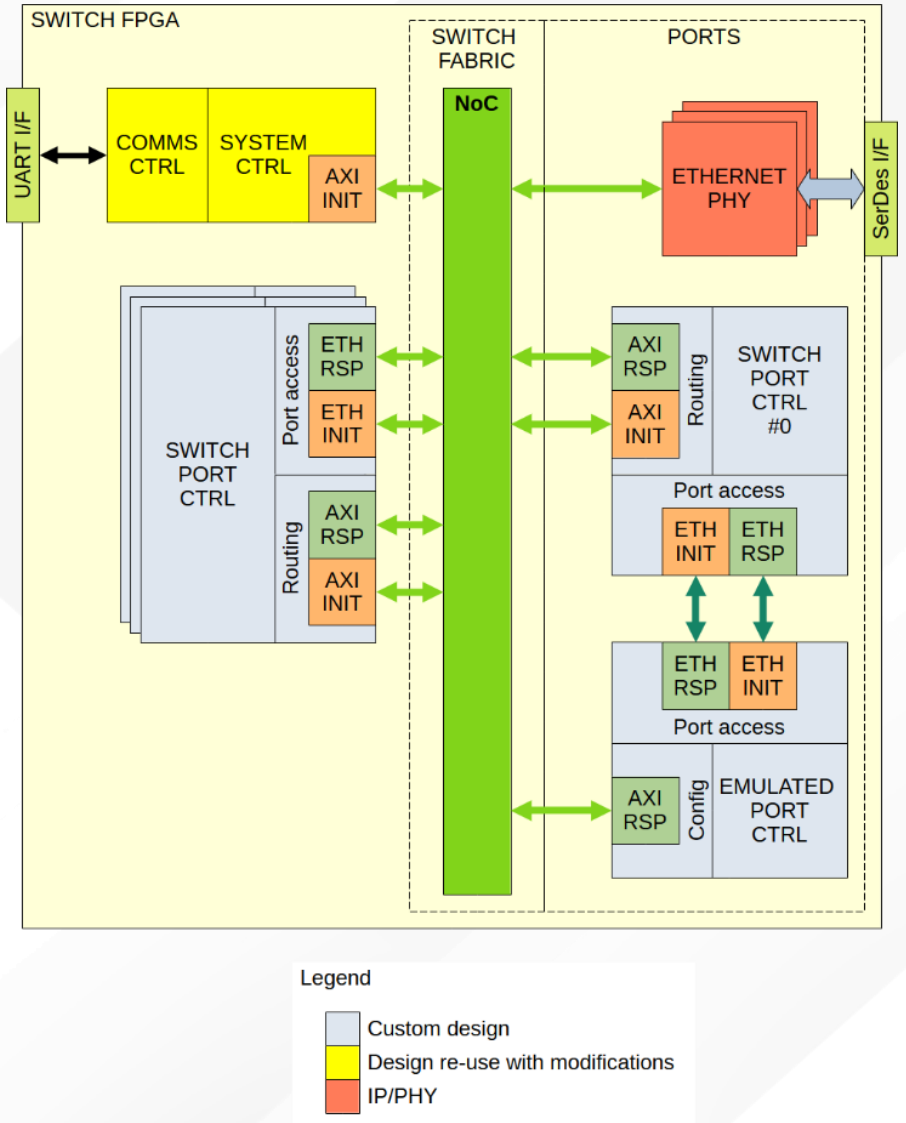

Ethernet layer 2 switch - SWITCH PORT CTRL

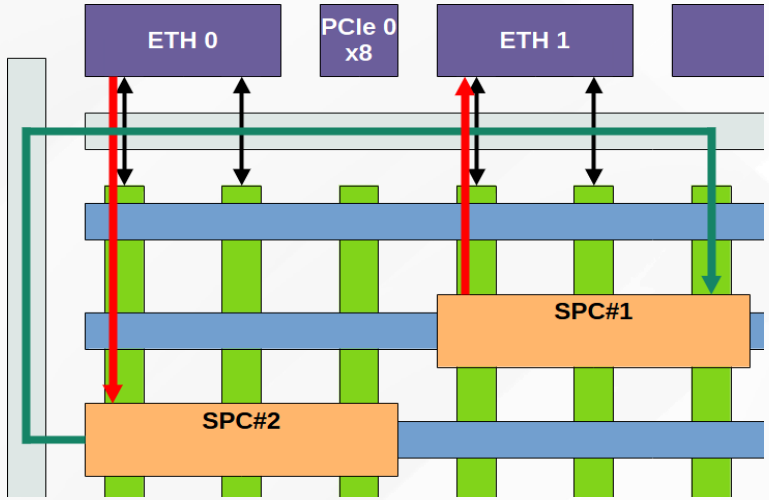

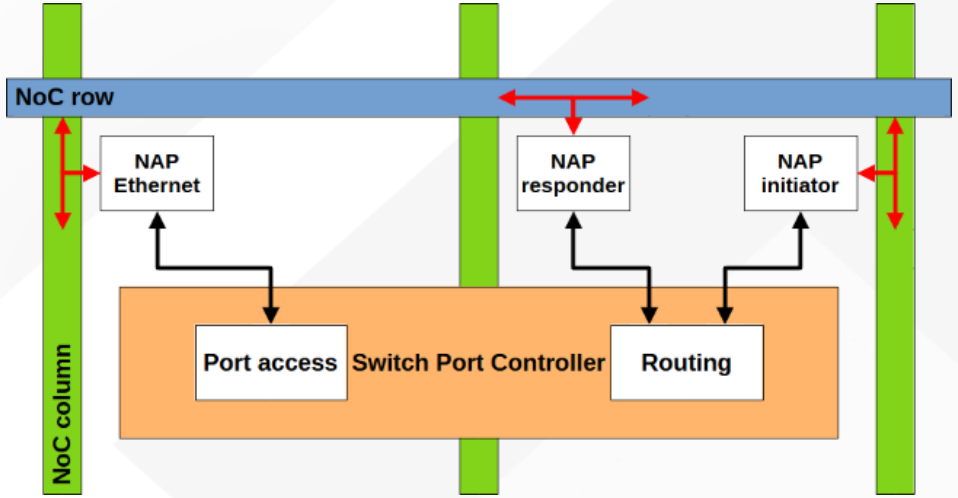

- No direct connection between two columns connected to two different Ethernet submodules – SWITCH PORT CTRL (SPC).
- Each SPC has two NoC interfaces:
  - for communication with Ethernet submodule (packets RX/TX)
  - for Ethernet packet routing to / from another SPC.

Ethernet layer 2 switch - SWITCH PORT CTRL

- SPC instances spread out in fabric to minimize congestion to the NoC ring.
- Final SPC placement:

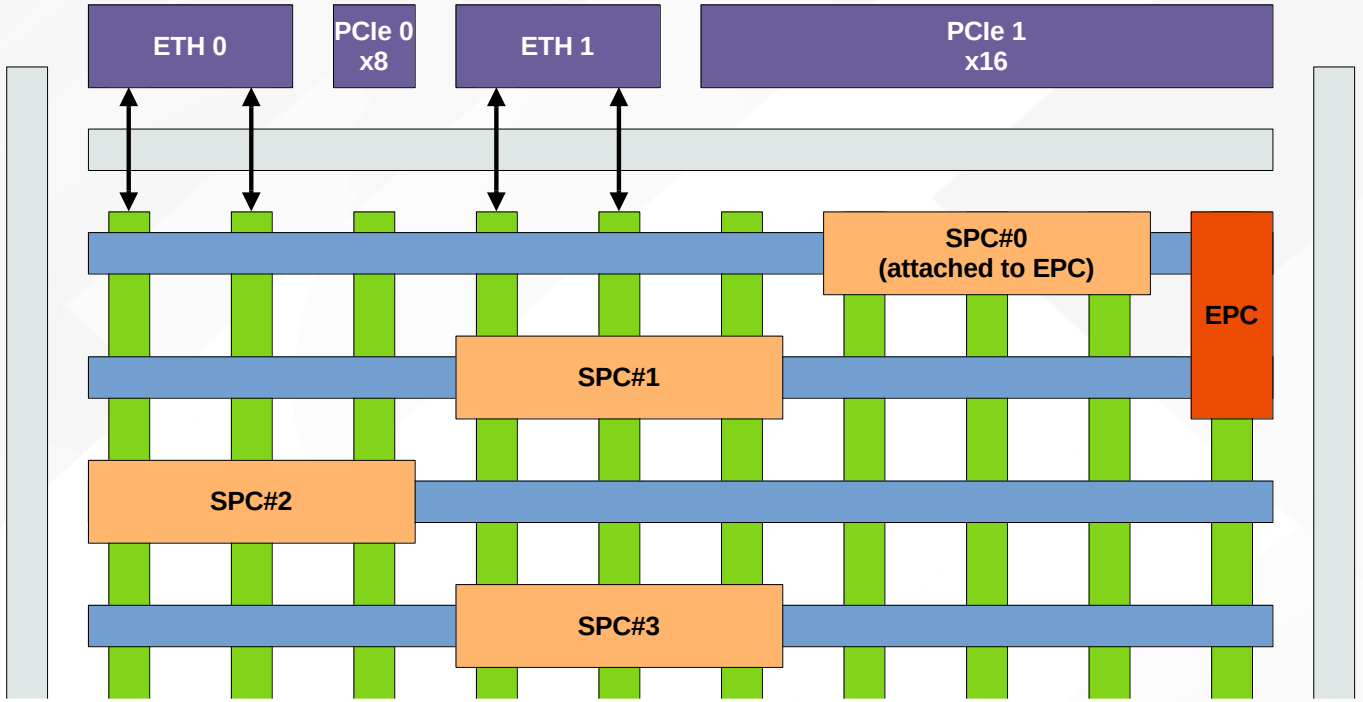

Ethernet layer 2 switch – challenges and limitations

- The BittWare **evaluation board** allows for only 56 Gbps PAM4 (28 Gbps NRZ) per transceiver lane.
- For a 100 Gbps Ethernet port the **timing closure** on AXI and Ethernet related clocks could not be reached within the project's schedule.
- Issues with **Ethernet submodule BFM** (number of Ethernet ports).

Final Ethernet ports configuration in Achronix Evaluation Kit is **3x 50GAUI-2** connected to the QSFP:

Connector	Ethernet subsystem	Number of ports	Data rate [Gbps] (per port)	SerDes lanes (per port)	SerDes rate [Gbps] (per lane)
QSFP56	ETH_1	1	50	2	25
QSFP-DD	ETH_0	2	50	2	25

Synthesis / Place&Route results

FPGA resources

Resource name	Used	Total	Usage [%]
DFF	23259	1382400	1.68
LUT	21612	691200	3.13
MLP	1	2560	less than 1
ALU8	309	172800	less than 1
LRAM	144	2560	5.62
BRAM	2	2560	less than 1

FPGA power report

Description \ Junction temp.	Power [W]		
	0°C	-40°C	125°C
Total Dynamic Power	36.2097	64.4775	64.5243
Total Static Power	19.3248	17.3300	74.8594
Core Dynamic Power	1.1351	1.1234	1.1702
Core Static Power	1.5786	0.3053	50.9195
Hard IP Dynamic Power	35.0746	63.3541	63.3541
Hard IP Static Power	17.7462	17.0247	23.9399
Dynamic Instance Power	0.8122	0.8038	0.8373
Dynamic Interconnect Power	0.3229	0.3196	0.3329
Dynamic Clock Network Power	0.1445	0.1430	0.1490
Total Power	**55.5345**	**81.8075**	**139.3837**

FPGA timing report

```
Collective Summary of All Corners
                    Slack (ns)     Frequency (MHz)
Clock / Group       setup  hold    Target  Upper Limit
i_clk_axi           0.663  0.003   210.0   243.9
i_clk_eth           0.922  0.003   260.0   342.0
i_clk_sys           3.640  0.003   100.0   157.2
```

Agenda

- Introduction and objectives
- Achronix FPGA technology
- Ethernet layer 2 switch design
- **Tests in hardware**
- AMD Versal NoC
- Summary

Test setup

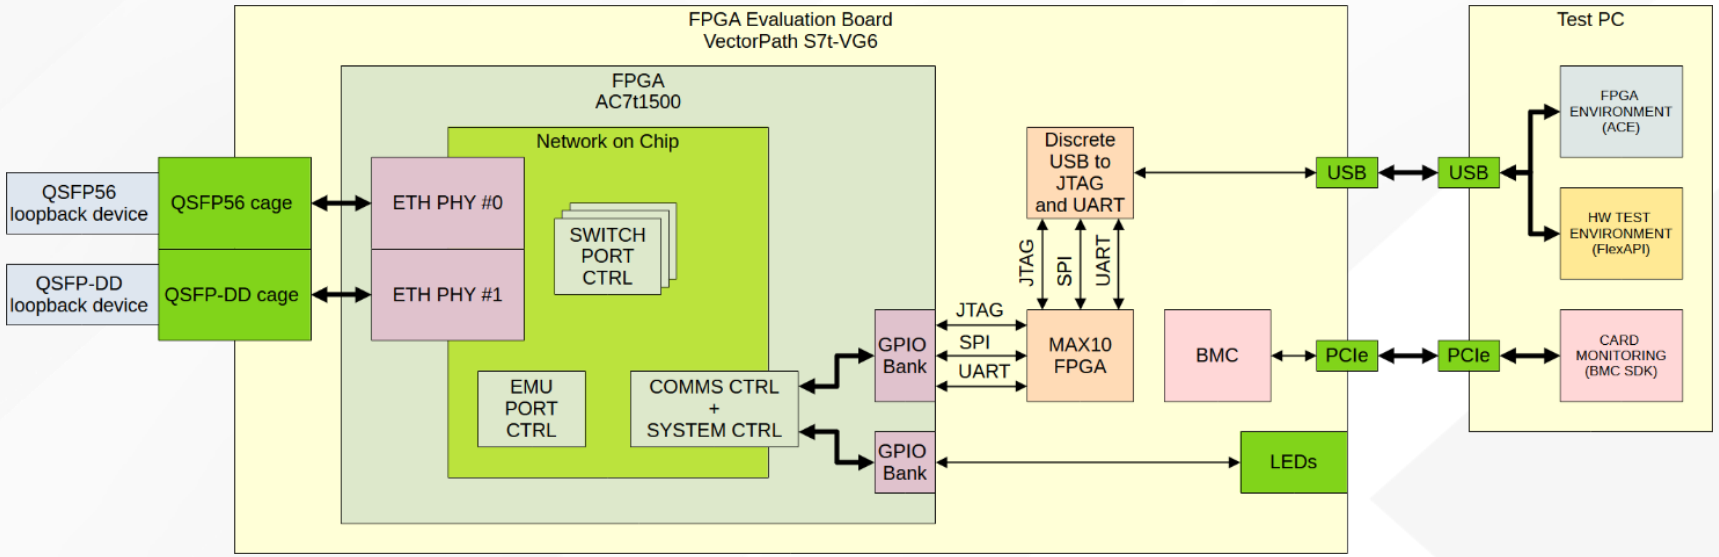

The entire Test Setup:

- The Achronix Breadboard (BB):
  - The BittWare VectorPath S7t-VG6 accelerator card with the Achronix AC7t1500 FPGA chip.
  - The QSFP56 and QSFP-DD loopback devices.
  - USB port connection to the Test PC, for MODBUS RTU communication via UART.
  - PCIe port connection to the Test PC, which provides card monitoring.
- The Test PC.

Test results

Test:

Four test cases ran, each using three different paths via all three available Ethernet ports.

Results:

- No functional errors observed
- 50 Gbps throughput per port.

Description	Throughput [Gbps]		
	Seq. #1	Seq. #2	Seq. #3
Test case 1 – 2048 packets, 256 bytes each	47.645	47.645	47.515
Test case 2 – 4096 packets, 1522 bytes each	48.555	48.555	48.555
Test case 3 – 50 packets, 9000 bytes each	49.855	49.79	49.855
Test case 4 – 32,767 packets, 9000 bytes each	49.855	49.92	49.855

Agenda

- Introduction and objectives
- Achronix FPGA technology
- Ethernet layer 2 switch design
- Tests in hardware
- **AMD Versal NoC**
- Summary

AMD Versal NoC – overview

Features:

- **VNoC** – to connect logic in the fabric
- **HNoC** – to connect the AI Engine Processing System, memory controllers and others
- **NMU** – to connect master to the NoC via AXI4
- **NSU** – to connect slave to the NoC via AXI4
- **No direct connection between Ethernet MAC and NoC**
- No ring (as in Achronix) – different VNoCs connected with HNoC
- NoC clock – max. 1080 MHz.

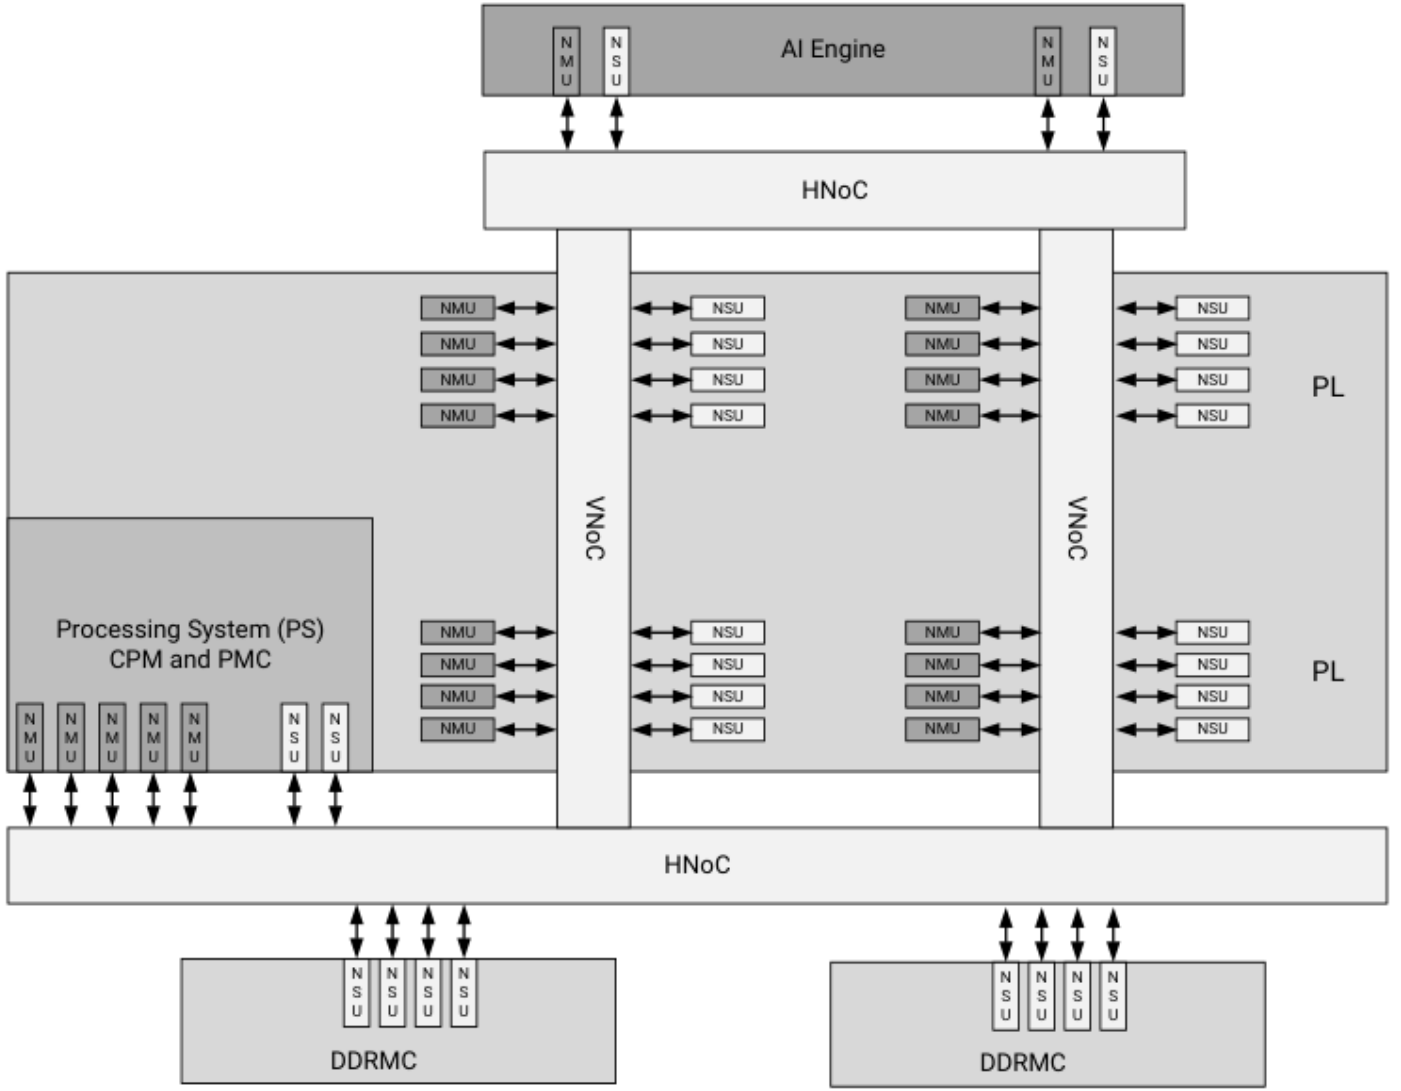

AMD Versal NoC – project approach

Initial goal: to transfer Ethernet switch solution from project with Achronix FPGA but...

- no Ethernet MAC connection to NoC
- no mixed language support in Vivado simulation when using NoC (all Verilog design necessary).

Thus, simple project to measure performance in simulation just with AXI transactions (no Ethernet packets):

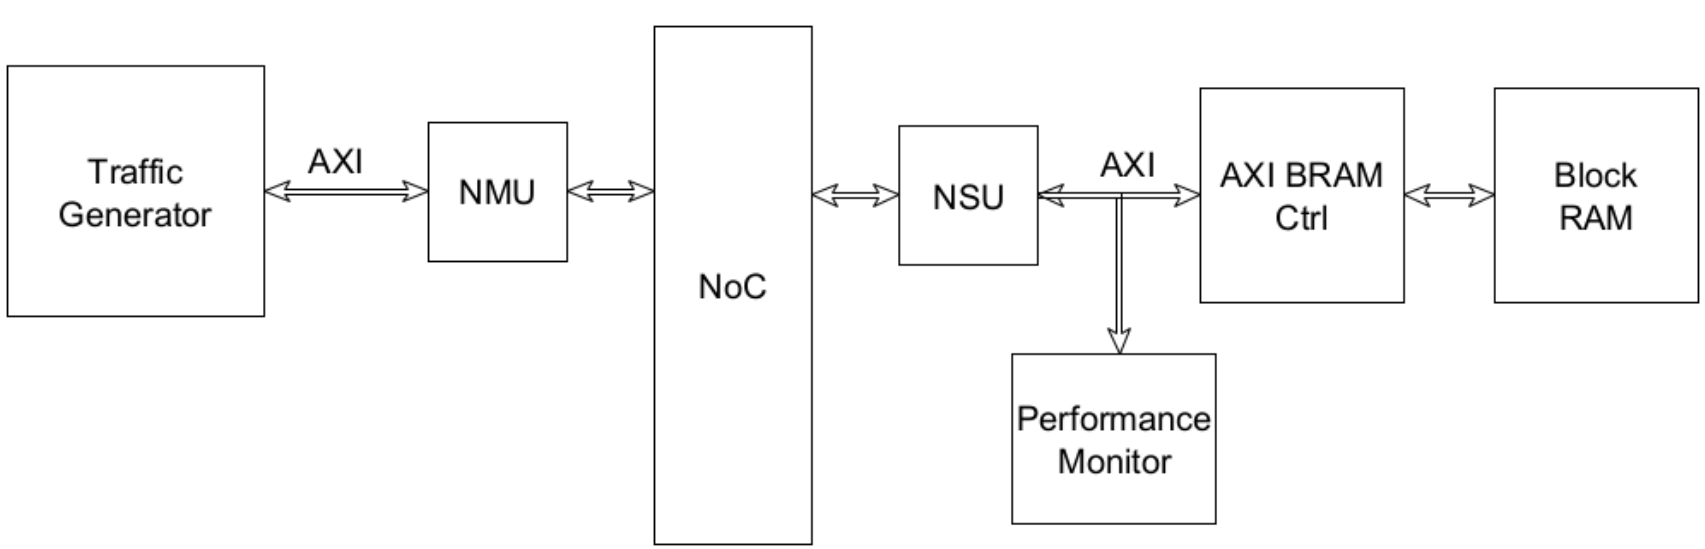

AMD Versal NoC – performance

Performance results with 210 MHz AXI clock:

Description	Throughput [Gbps]
Test case 1 – 2048 bursts, 256 bytes each	42.0033
Test case 2 – 4096 bursts, 1536 bytes each	40.5117
Test case 3 – 50 bursts, 4096 bytes each	42.0097
Test case 4 – 71,998 packets, 4096 bytes each	41.9992

Not performed:

- synthesis and layout
- timing analysis.

Agenda

- Introduction and objectives
- Achronix FPGA technology
- Ethernet layer 2 switch design
- Tests in hardware
- AMD Versal NoC
- **Summary**

Summary

For Achronix FPGA:

- 100 Gbps target bandwidth per port and transceiver lane not reached.
- 50 Gbps bandwidth per port achieved for 3 Ethernet ports in 50GAUI-2 mode.
- NoC feature tested in Ethernet layer 2 switch design.
- There are no limitations on port routability.
- Ethernet layer 2 switch solution verified in hardware:
  - no functional errors detected during transmissions
  - all Ethernet ports operated with maximum configured data rate in both directions.

For AMD Versal FPGA:

- Switch solution transfer unsuccessful due to differences in Ethernet MAC connection to NoC and issues with NoC simulation in Vivado.
- Verified NoC performance in simulation of a simple project. Lower value (42 Gbps) needs further investigation.

AROBS Polska

Thank you for your attention!

Tomasz Rybak
e-mail: tomasz.rybak@arobs.pl