Accelerating Systems with Programmable Logic Comp.

High-Speed Serial Transceivers and Busses

April 27, 2020

Sina Borrami

Uppsala University

Sina.borami@physics.uu.se

Outline

- Serial and Parallel
- SERDES
- PCIe
- CCIX, CXL, OpenCAPI

Intro to Parallel and Serial interface

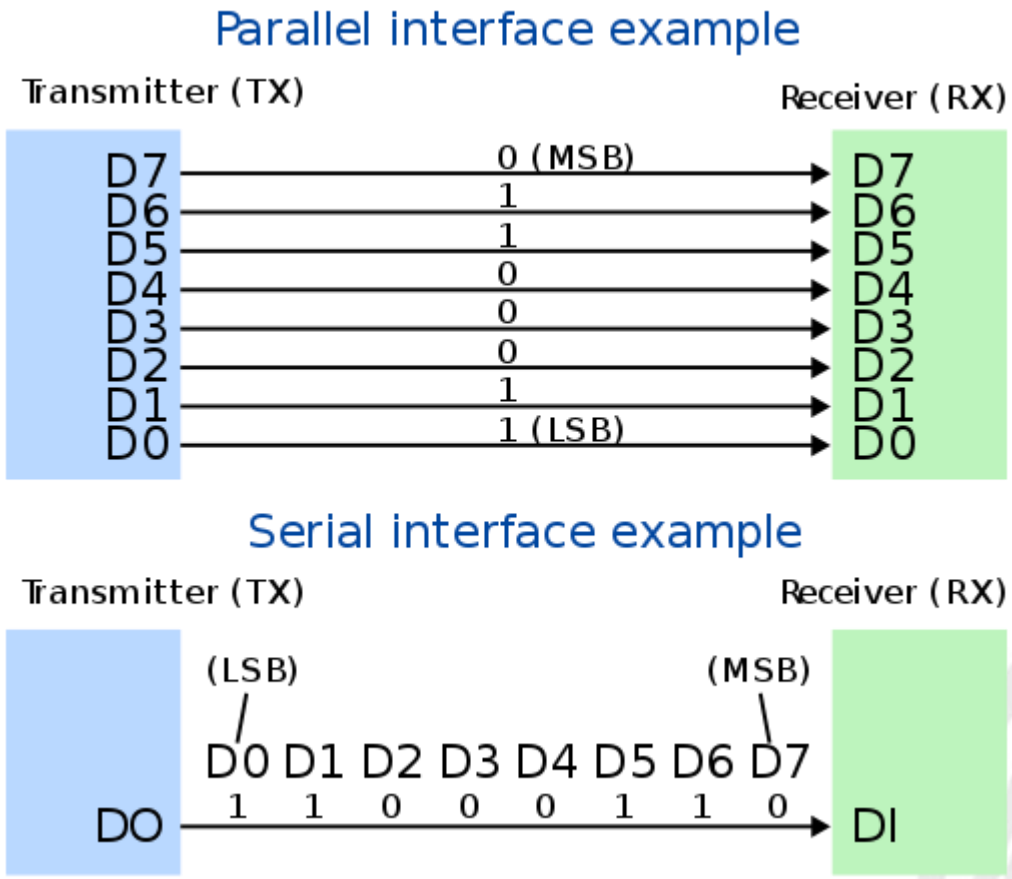

 Image source: Wikipeida.org

Which one is better

- Serial or Parallel ???

Problems with Parallel Interface

- Skew
- Cross talk
- Inefficient
- Limited BW

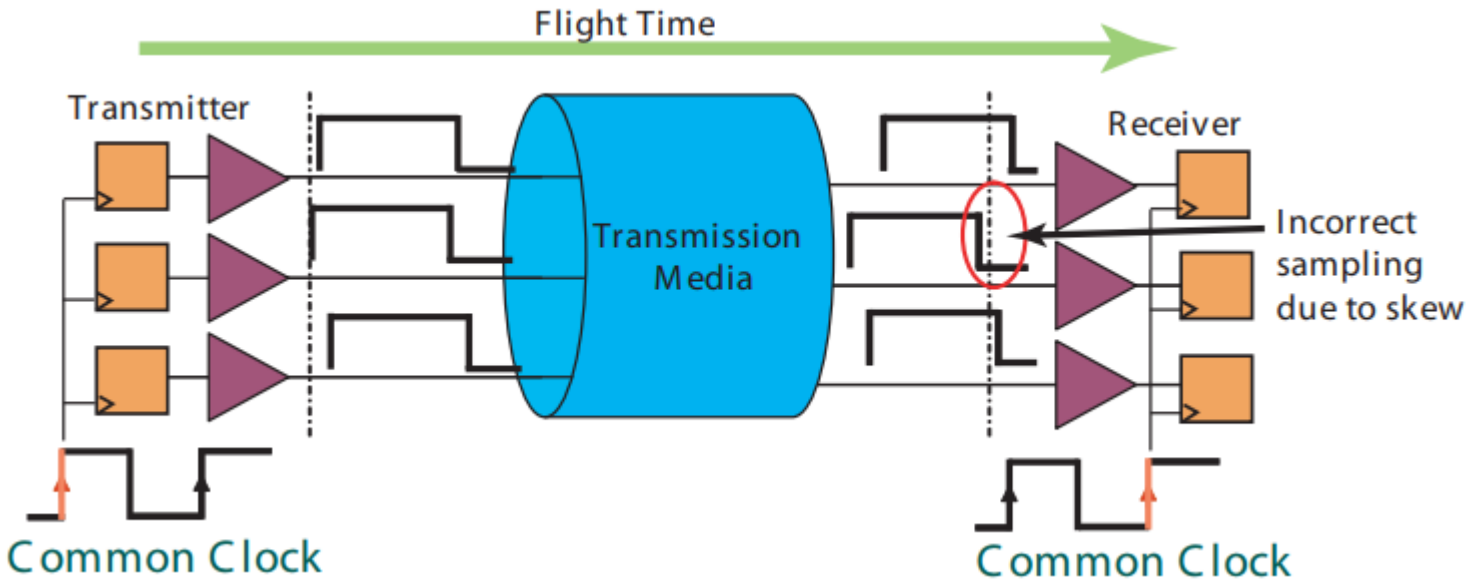

- ISA
- PCI

Solution is Serial

- USB 3.x
- Ethernet
- HDMI
- Display Port
- Aurora
- Fiber Optic
- Fiber channel
- SATA
- NVME
- Thunderbolt
- PCI-express
- Infiniband

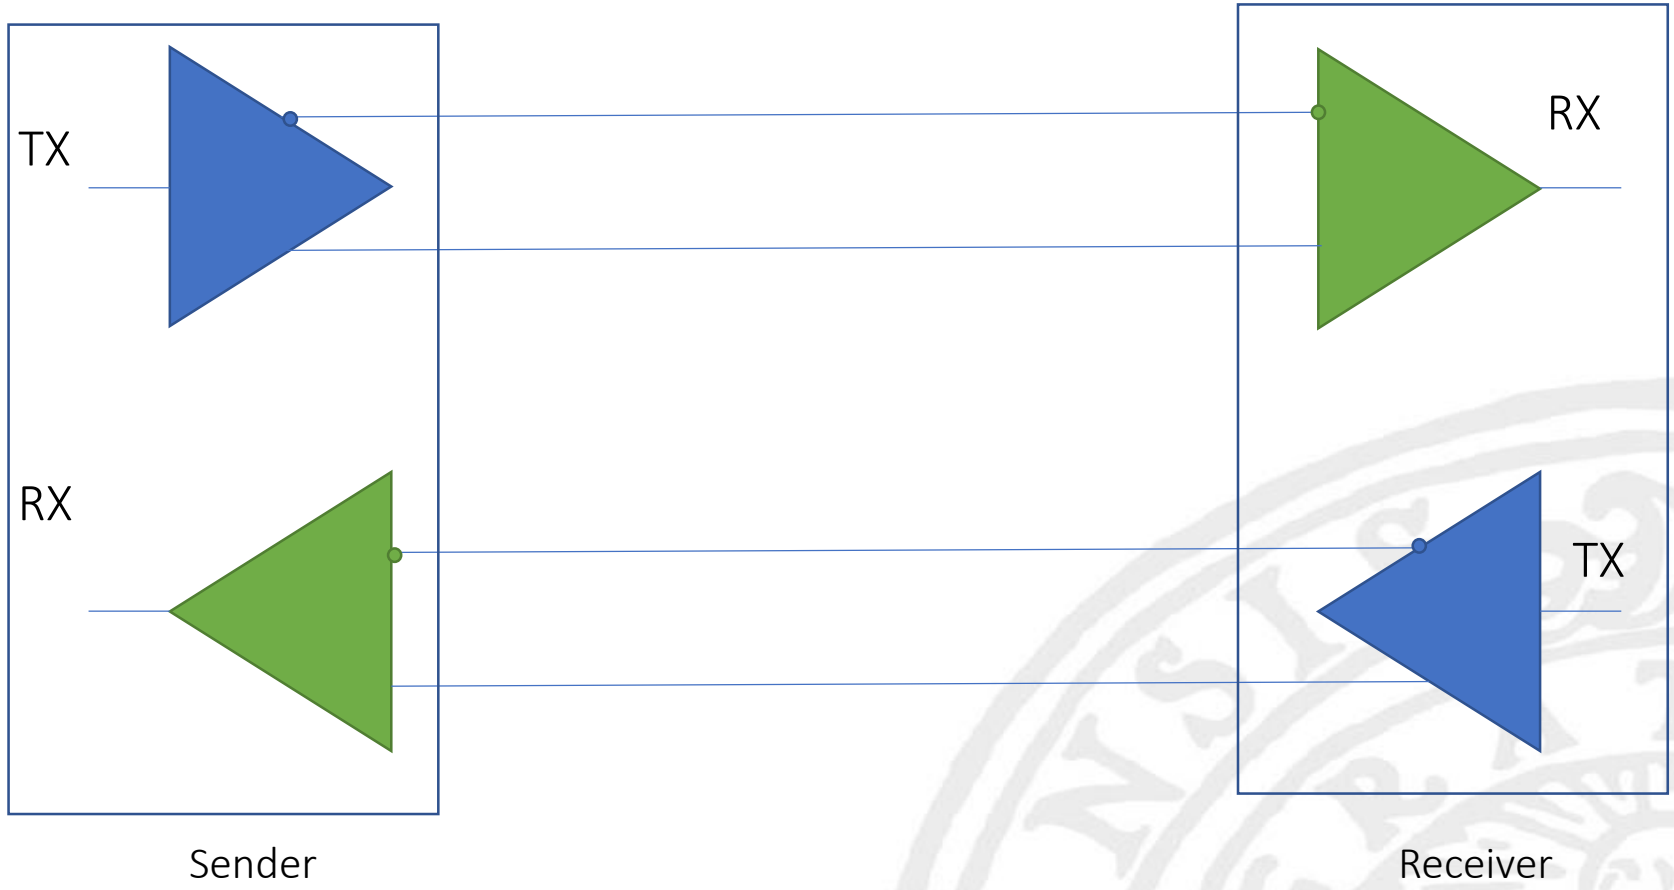

SERDES

Physical Coding Sublayer

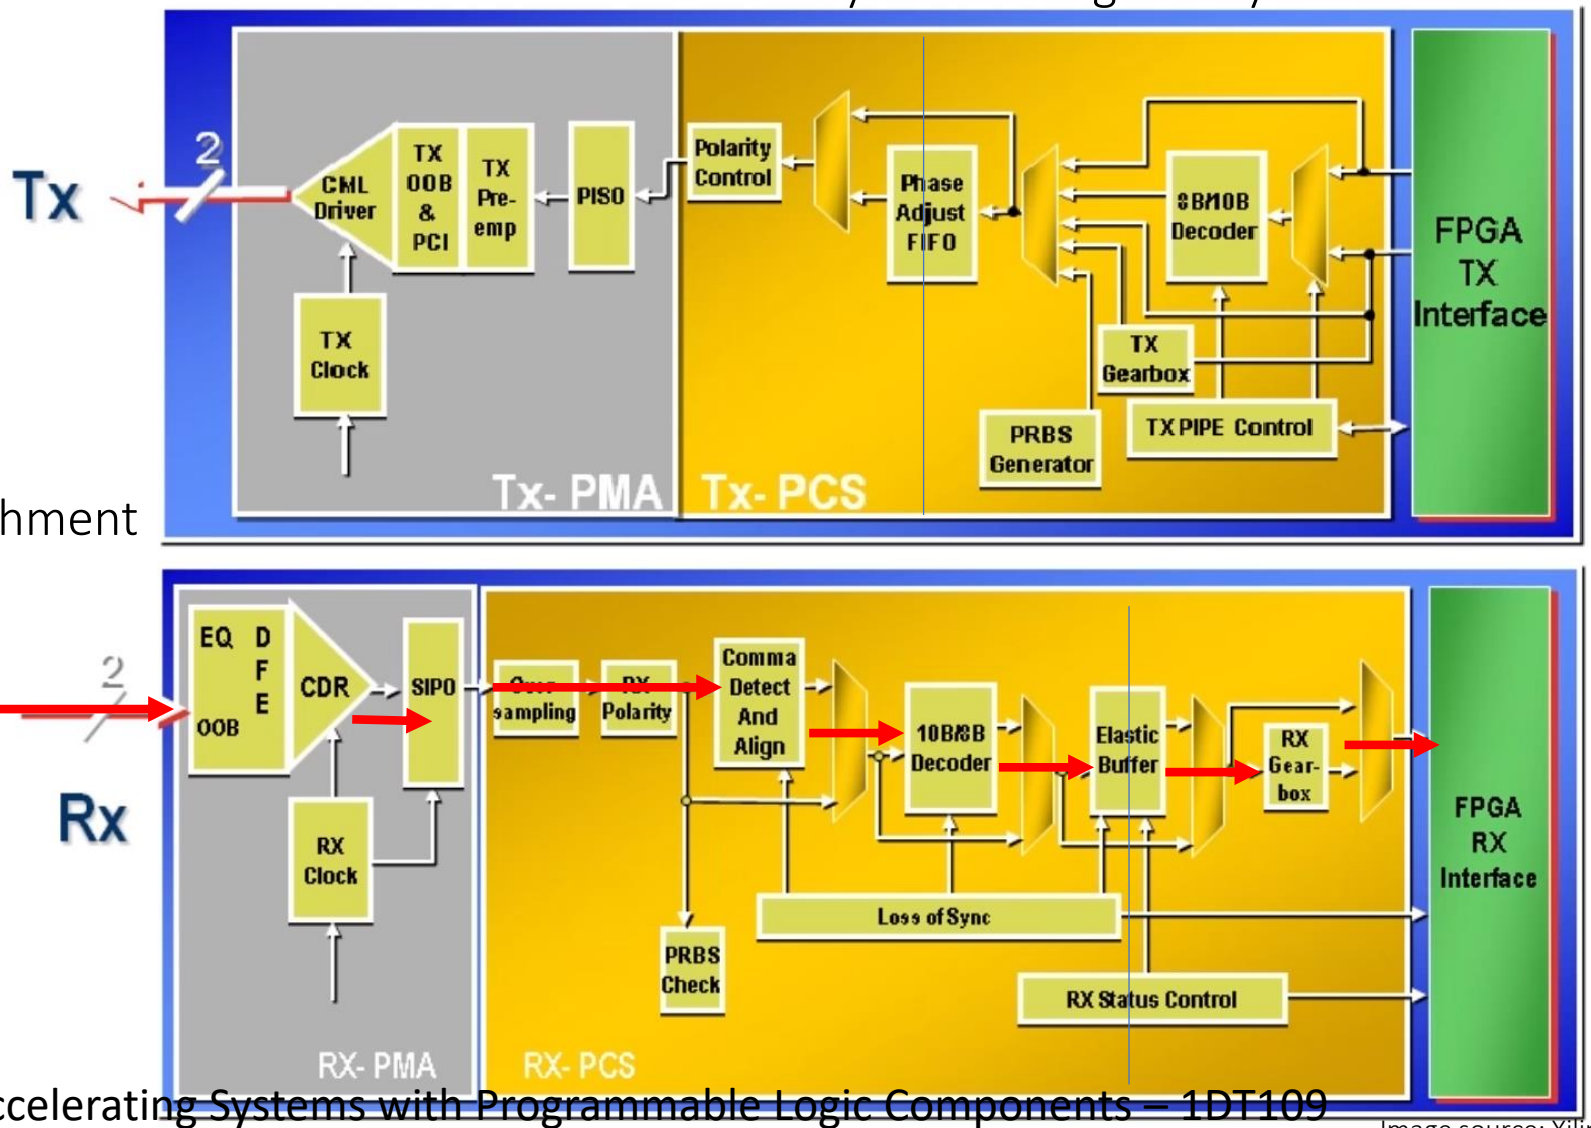

Physical Medium Attachment

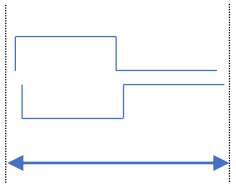

ps

Image source: Xilinx.com

Line codes

- NRZ
  - 58Gbps
- PAM4
  - 112Gbps

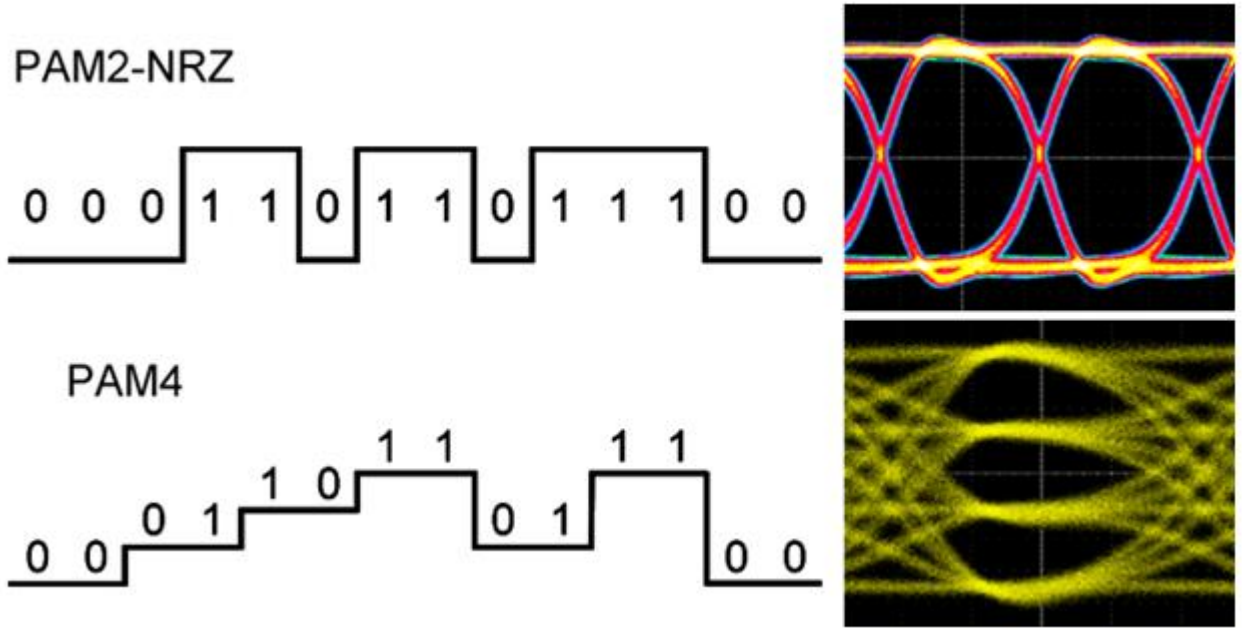

- 8b/10b encoding

 Image source: Tektronix.com

Line codes

- NRZ
  - 58Gbps
- PAM4
  - 112Gbps

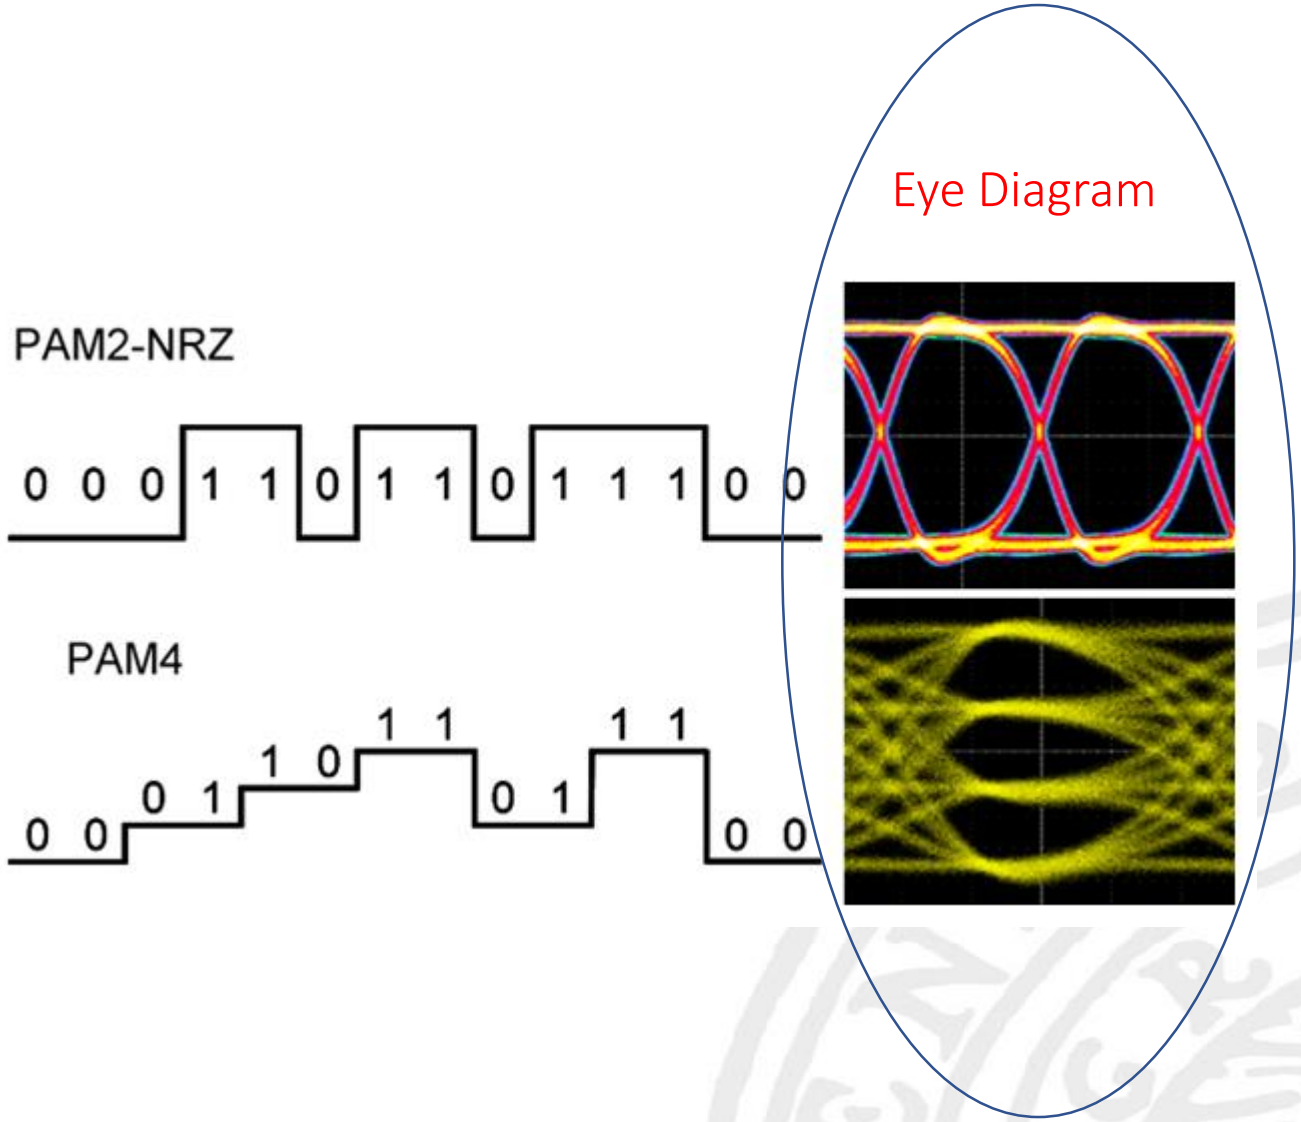

- 8b/10b encoding

Overview

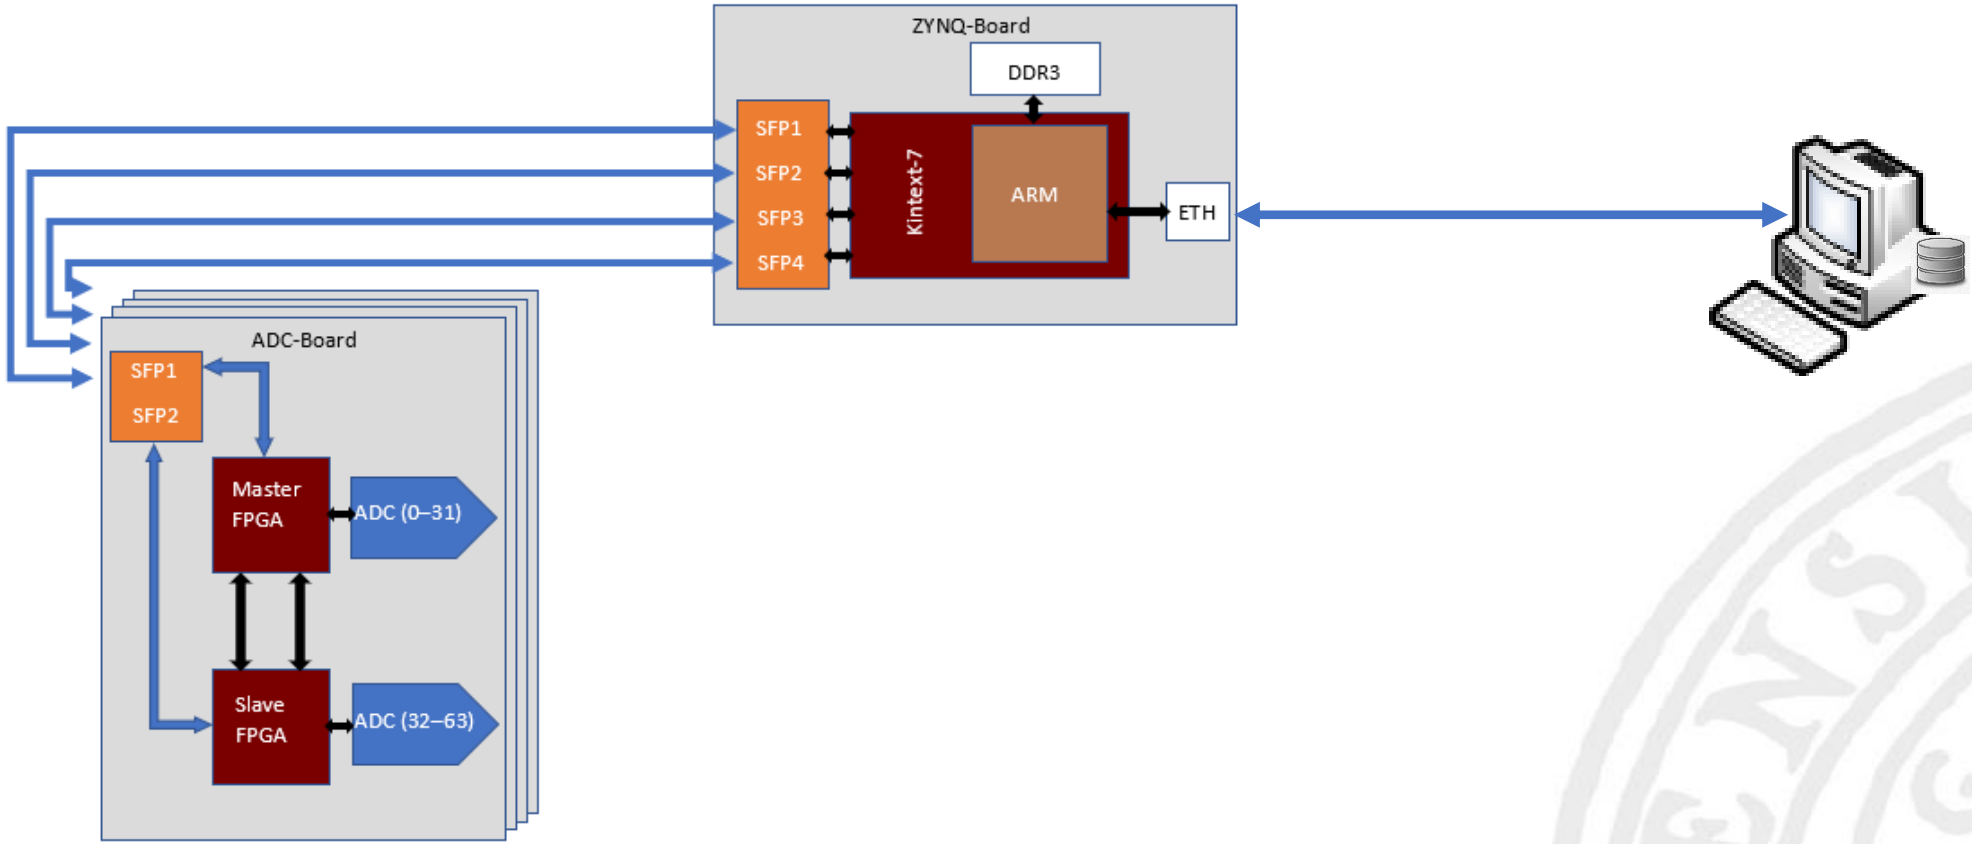
ZYNQ-Board
DDR3
SFP1
SFP2
SFP3
SFP4
Kintext-7
ARM
ETH
ADC-Board
SFP1
SFP2
Master FPGA
ADC (0–31)
Slave FPGA
ADC (32–63)

ZynqBoard

- Zynq-SOC
- 4 SFP Cage
- 1 Gigabit Ethernet
- UART USB
- 512MB DDR

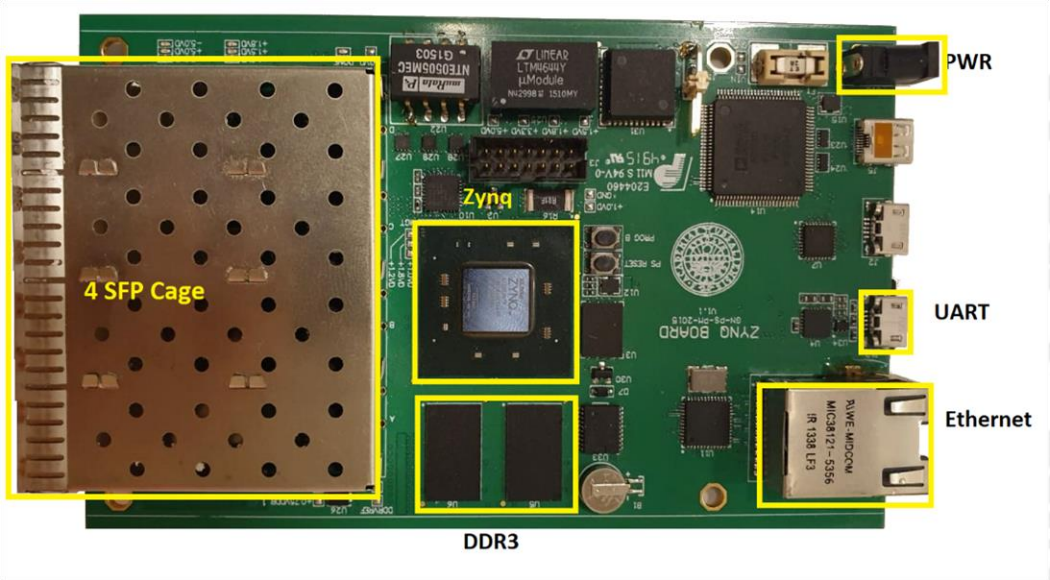

ADC Board

- 64 Channel
  - 80 MSPS
  - 14-bit resolution
- 2 FPGA
- 2 SFP Cage

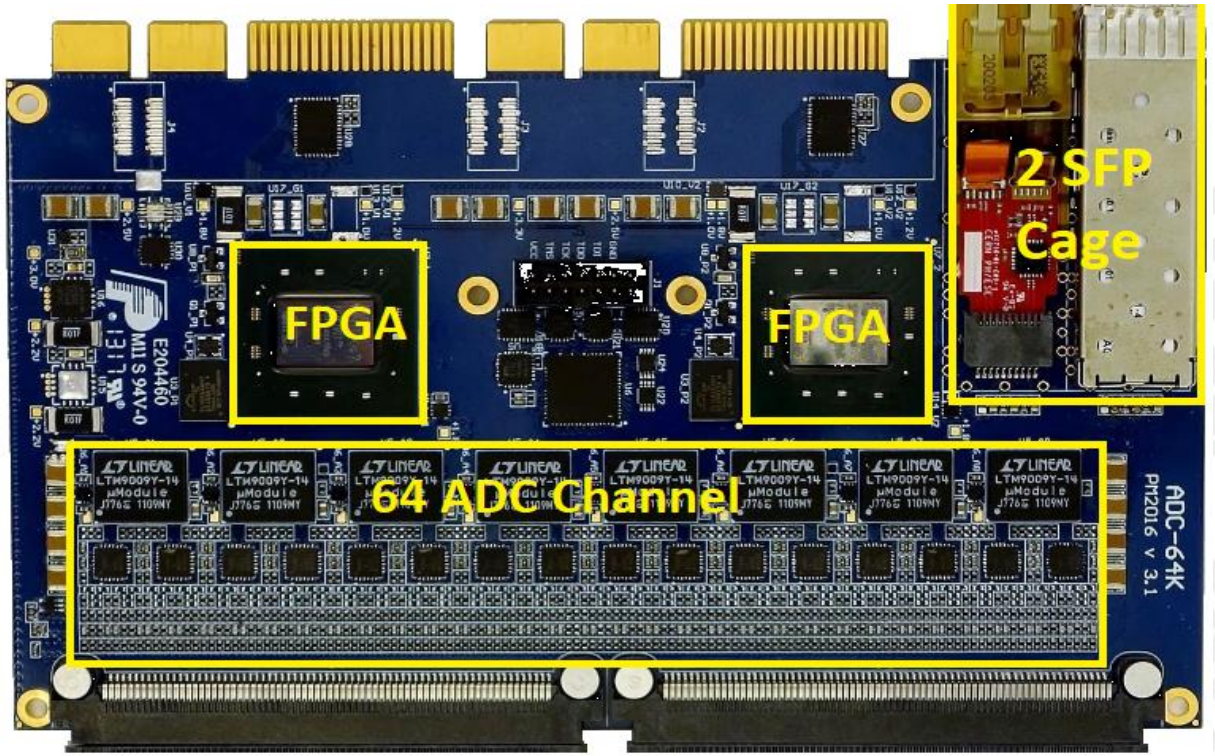

Custom Protocol

Item	Config	Rate
1	Transceiver speed	2Gbps
2	Encoding	8b/10b
3	PCS internal data width	40-bit
4	TX-RX buffers for synchronization	
5	Comma value	K 28.5
6	Comma alignment	Four Byte Boundaries
7	Four-byte Clock correction sequence	000000BC
8	K character	BC

ADC Data Format

Send

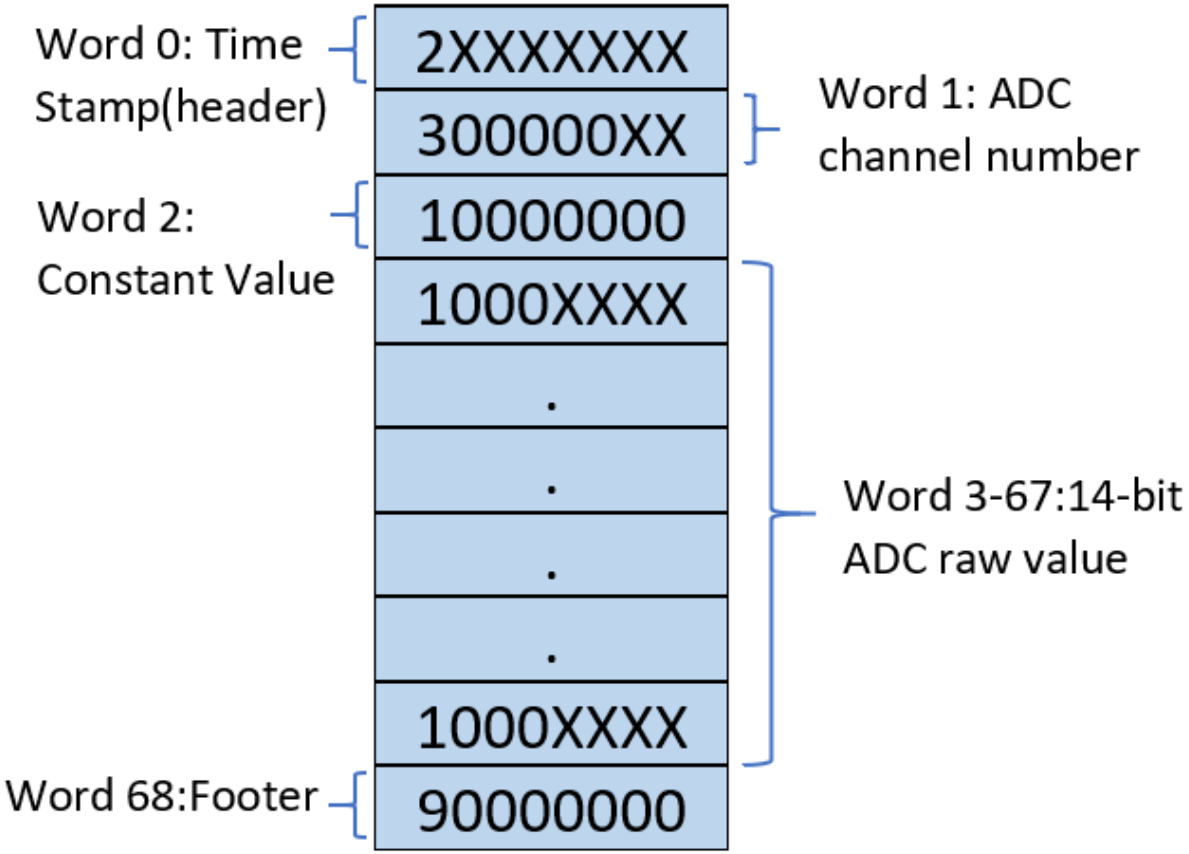

Internal Design

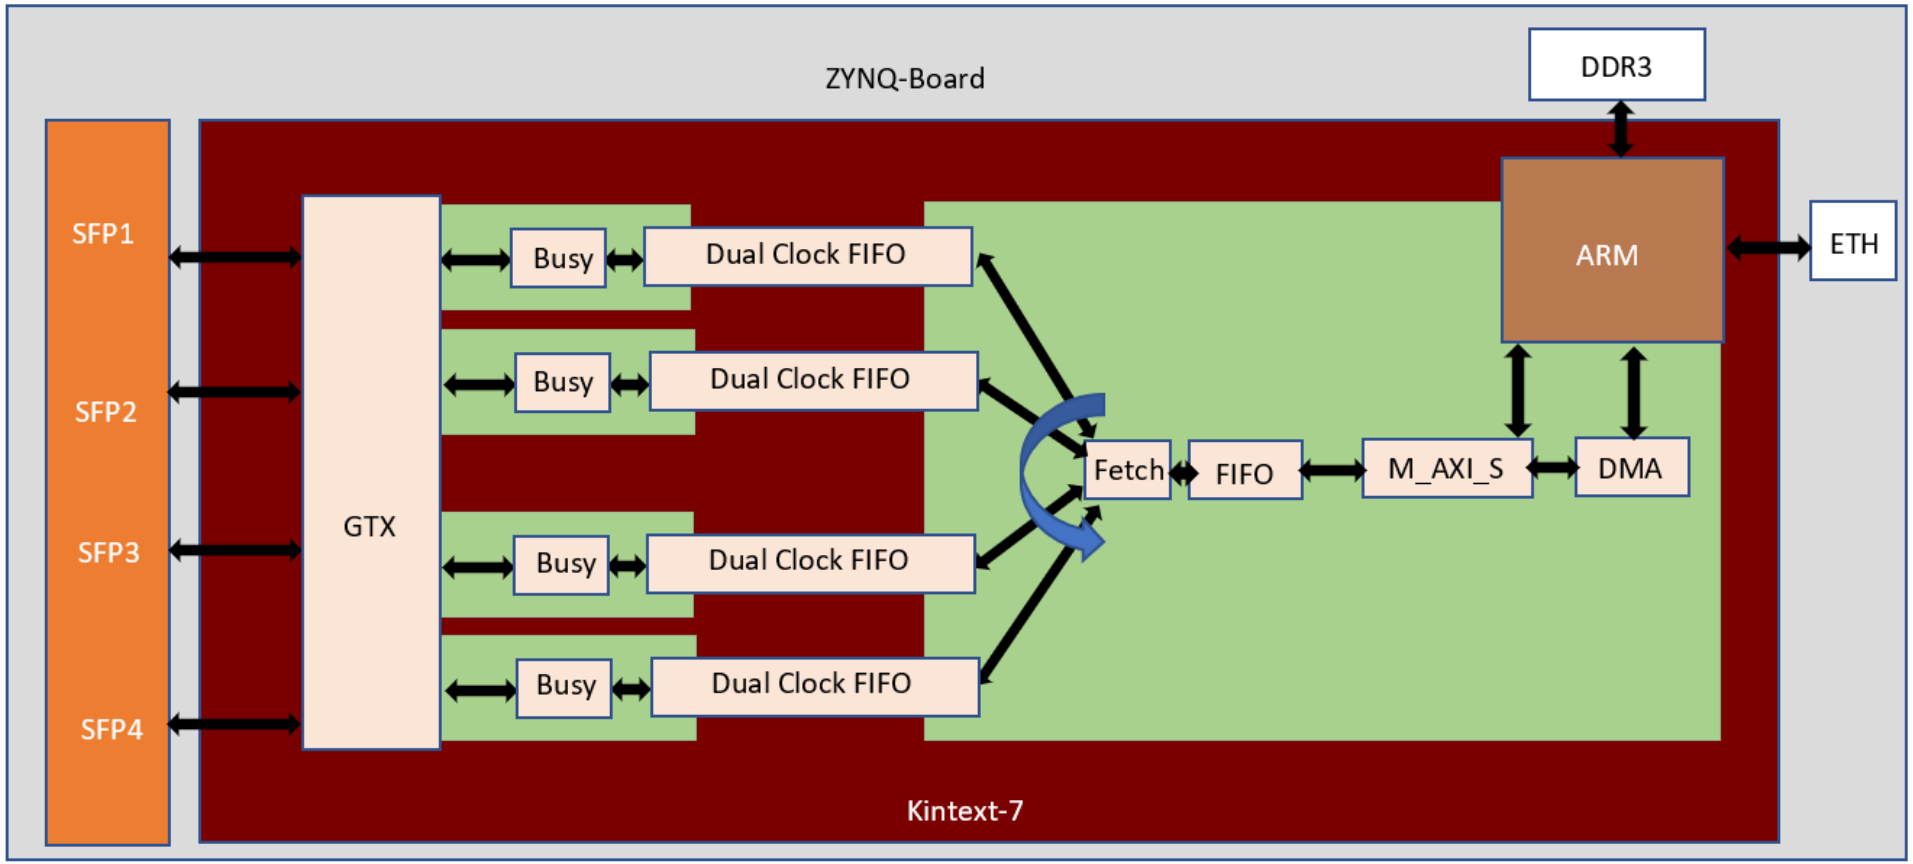

Block Design

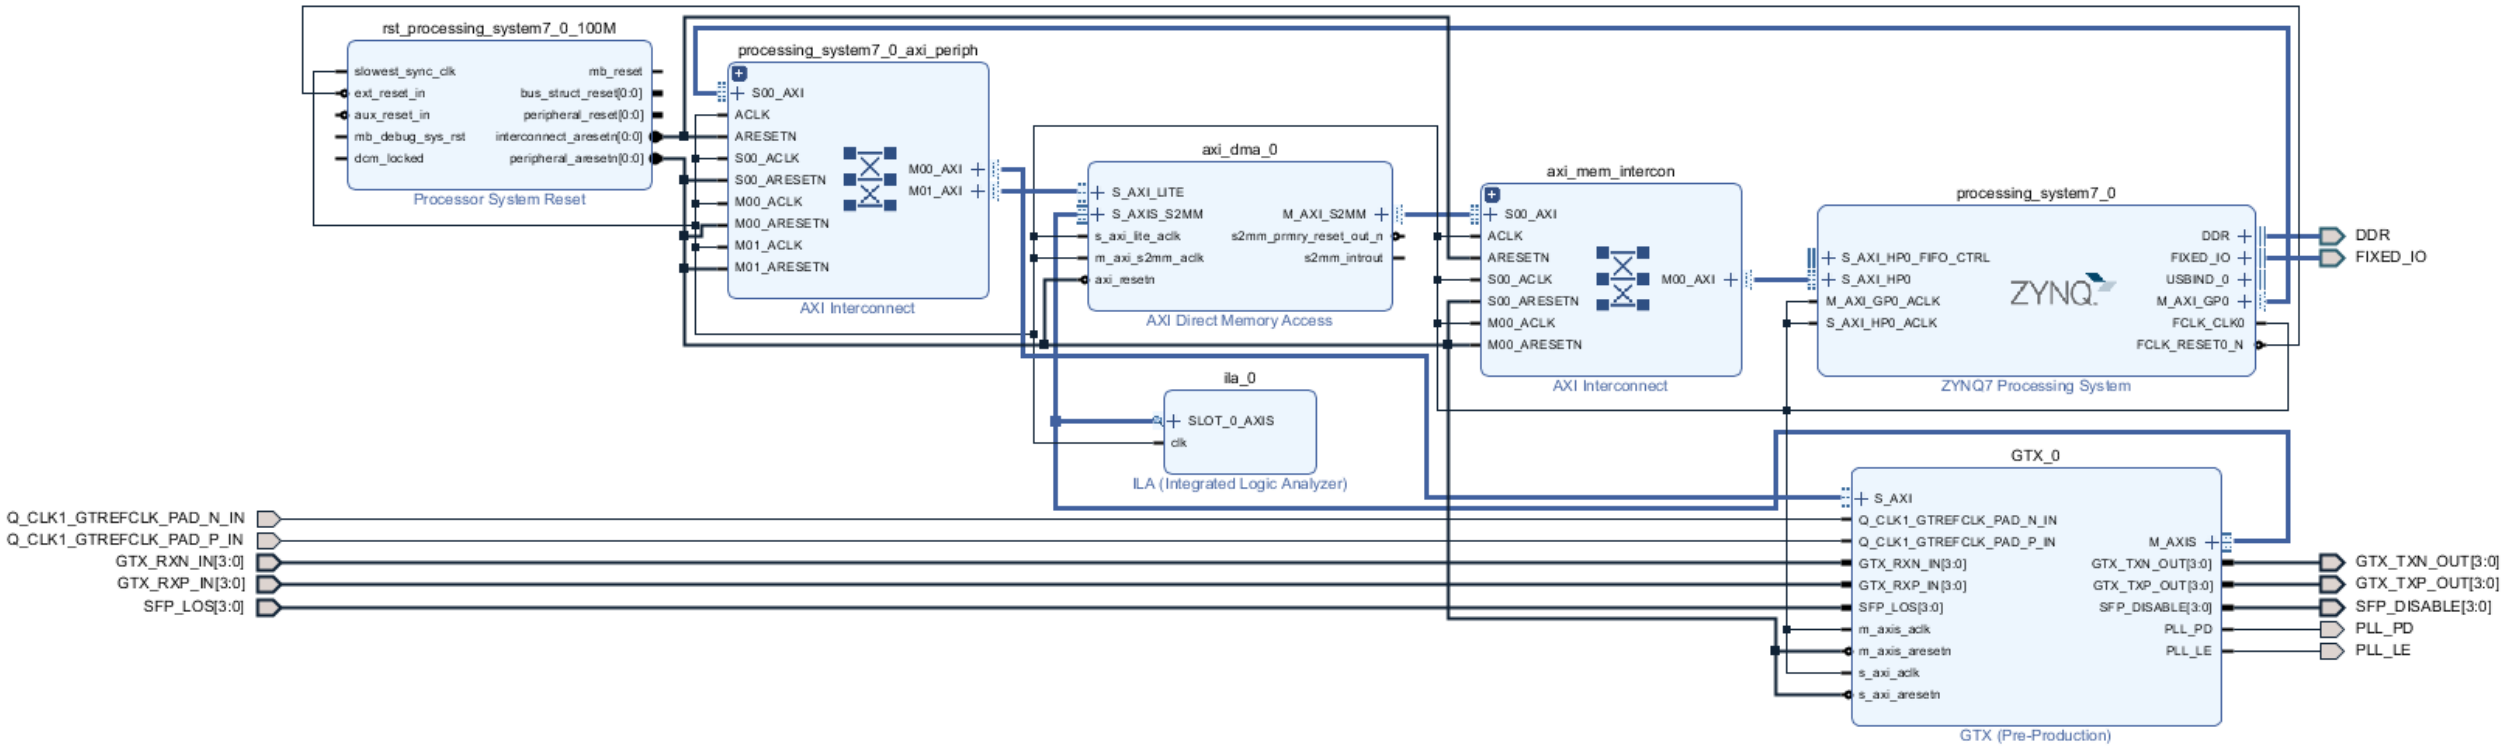

FPGA is not just logic cells

- PLL
- SERDER #4-96 (6-112Gbps)
- Buffers
- BRAM
- HBM2
- DSP engine
- AI engine
- Memory controller(DDR4/5/6)
- Security engine
- Video decoder
- Multiple CPU cores (SOC) RISC-V , ARM(R,A)

Lane

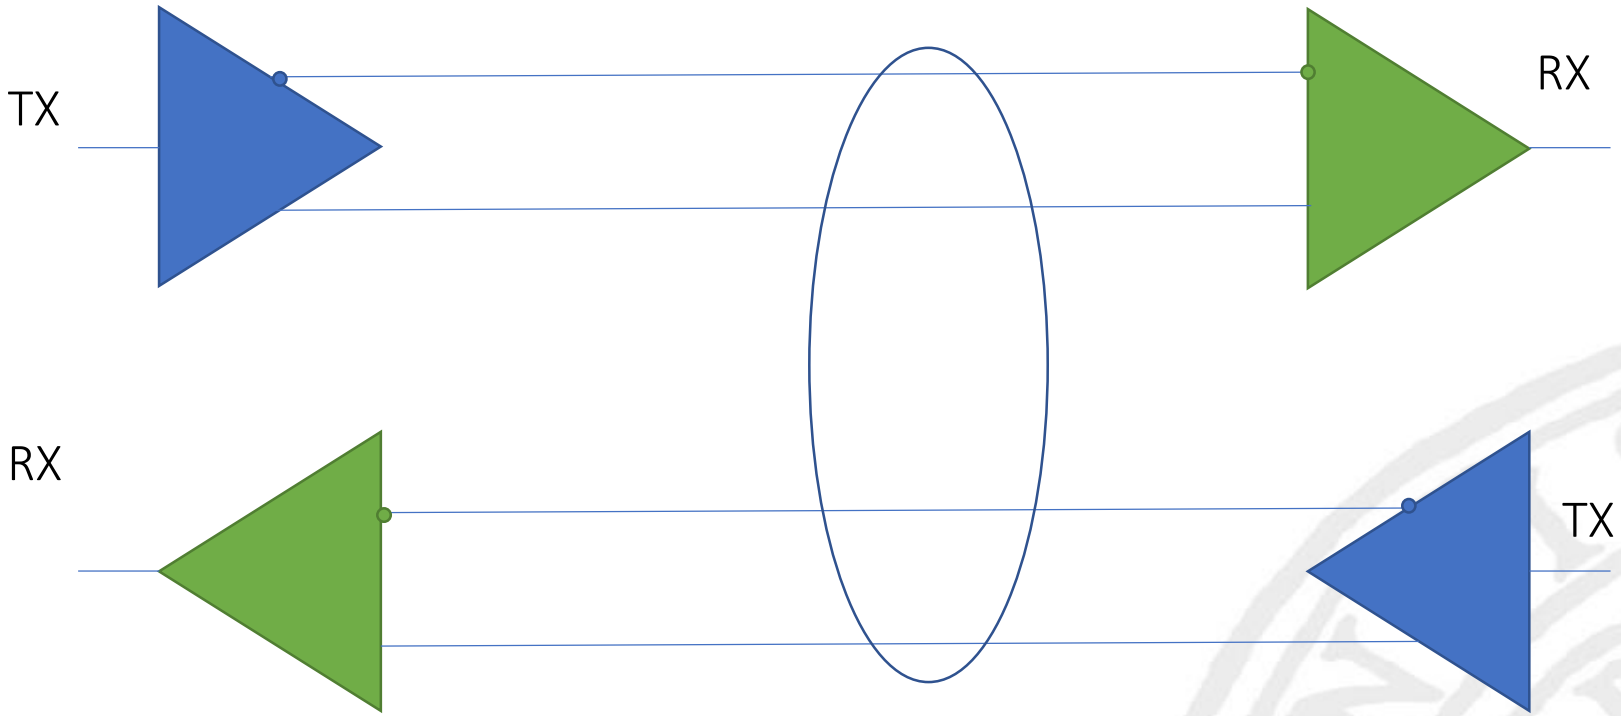

What connector is this?

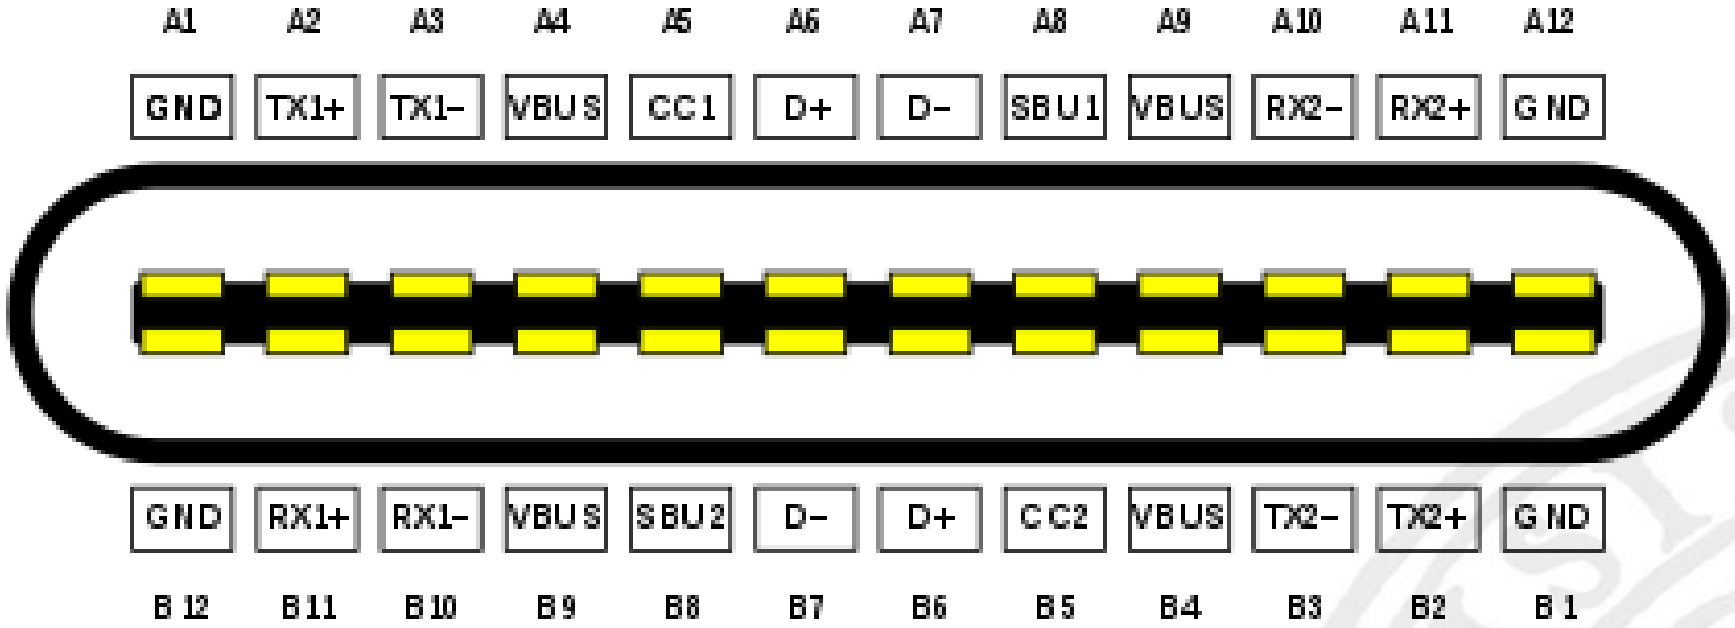

 Image source: Wikipeida.org

USB-C

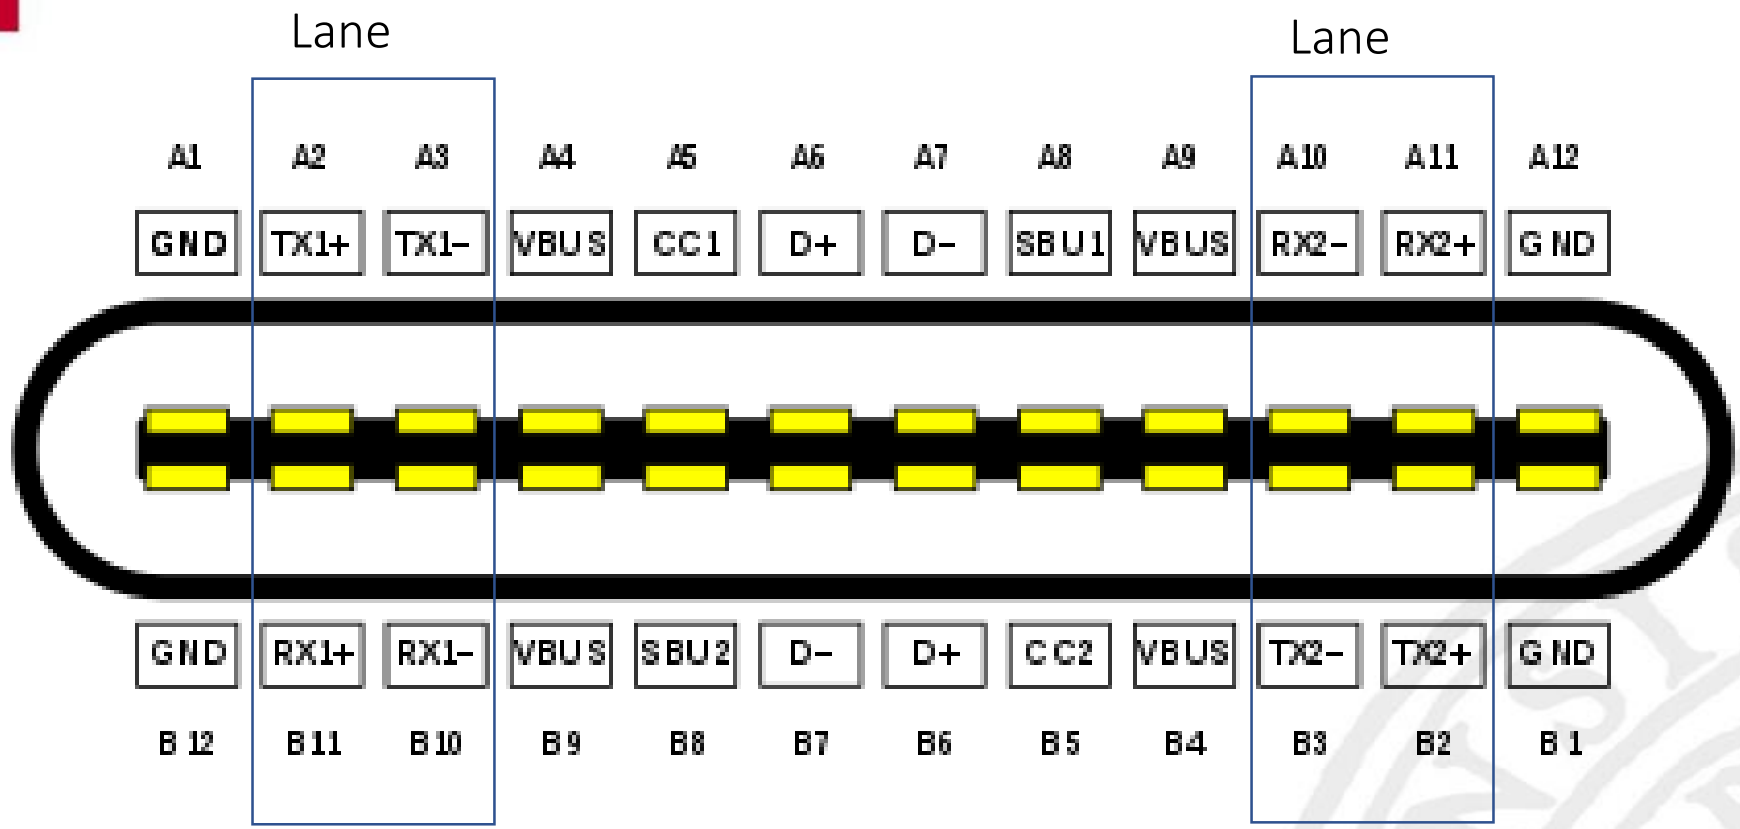

 Image source: Wikipeida.org

PCI express

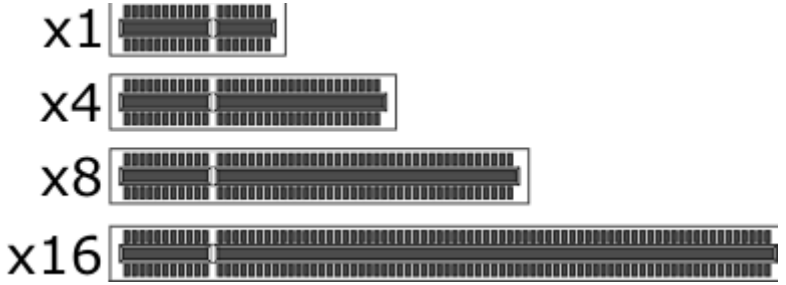

PCI Express version	Introduced	Line code	Transfer rate[i][ii]	Throughput[i][iii]				
				x1	x2	x4	x8	x16
1.0	2003	8b/10b	2.5 GT/s	250 MB/s	0.500 GB/s	1.00 GB/s	2.0 GB/s	4.0 GB/s
2.0	2007	8b/10b	5.0 GT/s	500 MB/s	1.000 GB/s	2.00 GB/s	4.0 GB/s	8.0 GB/s
3.0	2010	128b/130b	8.0 GT/s	984.6 MB/s	1.969 GB/s	3.94 GB/s	7.88 GB/s	15.75 GB/s
4.0	2017	128b/130b	16.0 GT/s	1969 MB/s	3.938 GB/s	7.88 GB/s	15.75 GB/s	31.51 GB/s
5.0	2019	128b/130b	32.0 GT/s[iv]	3938 MB/s	7.877 GB/s	15.75 GB/s	31.51 GB/s	63.02 GB/s
6.0 (planned)	2021	128b/130b	64.0 GT/s	7877 MB/s	15.754 GB/s	31.51 GB/s	63.02 GB/s	126.03 GB/s

Deskew in PCIe(channel bonding)

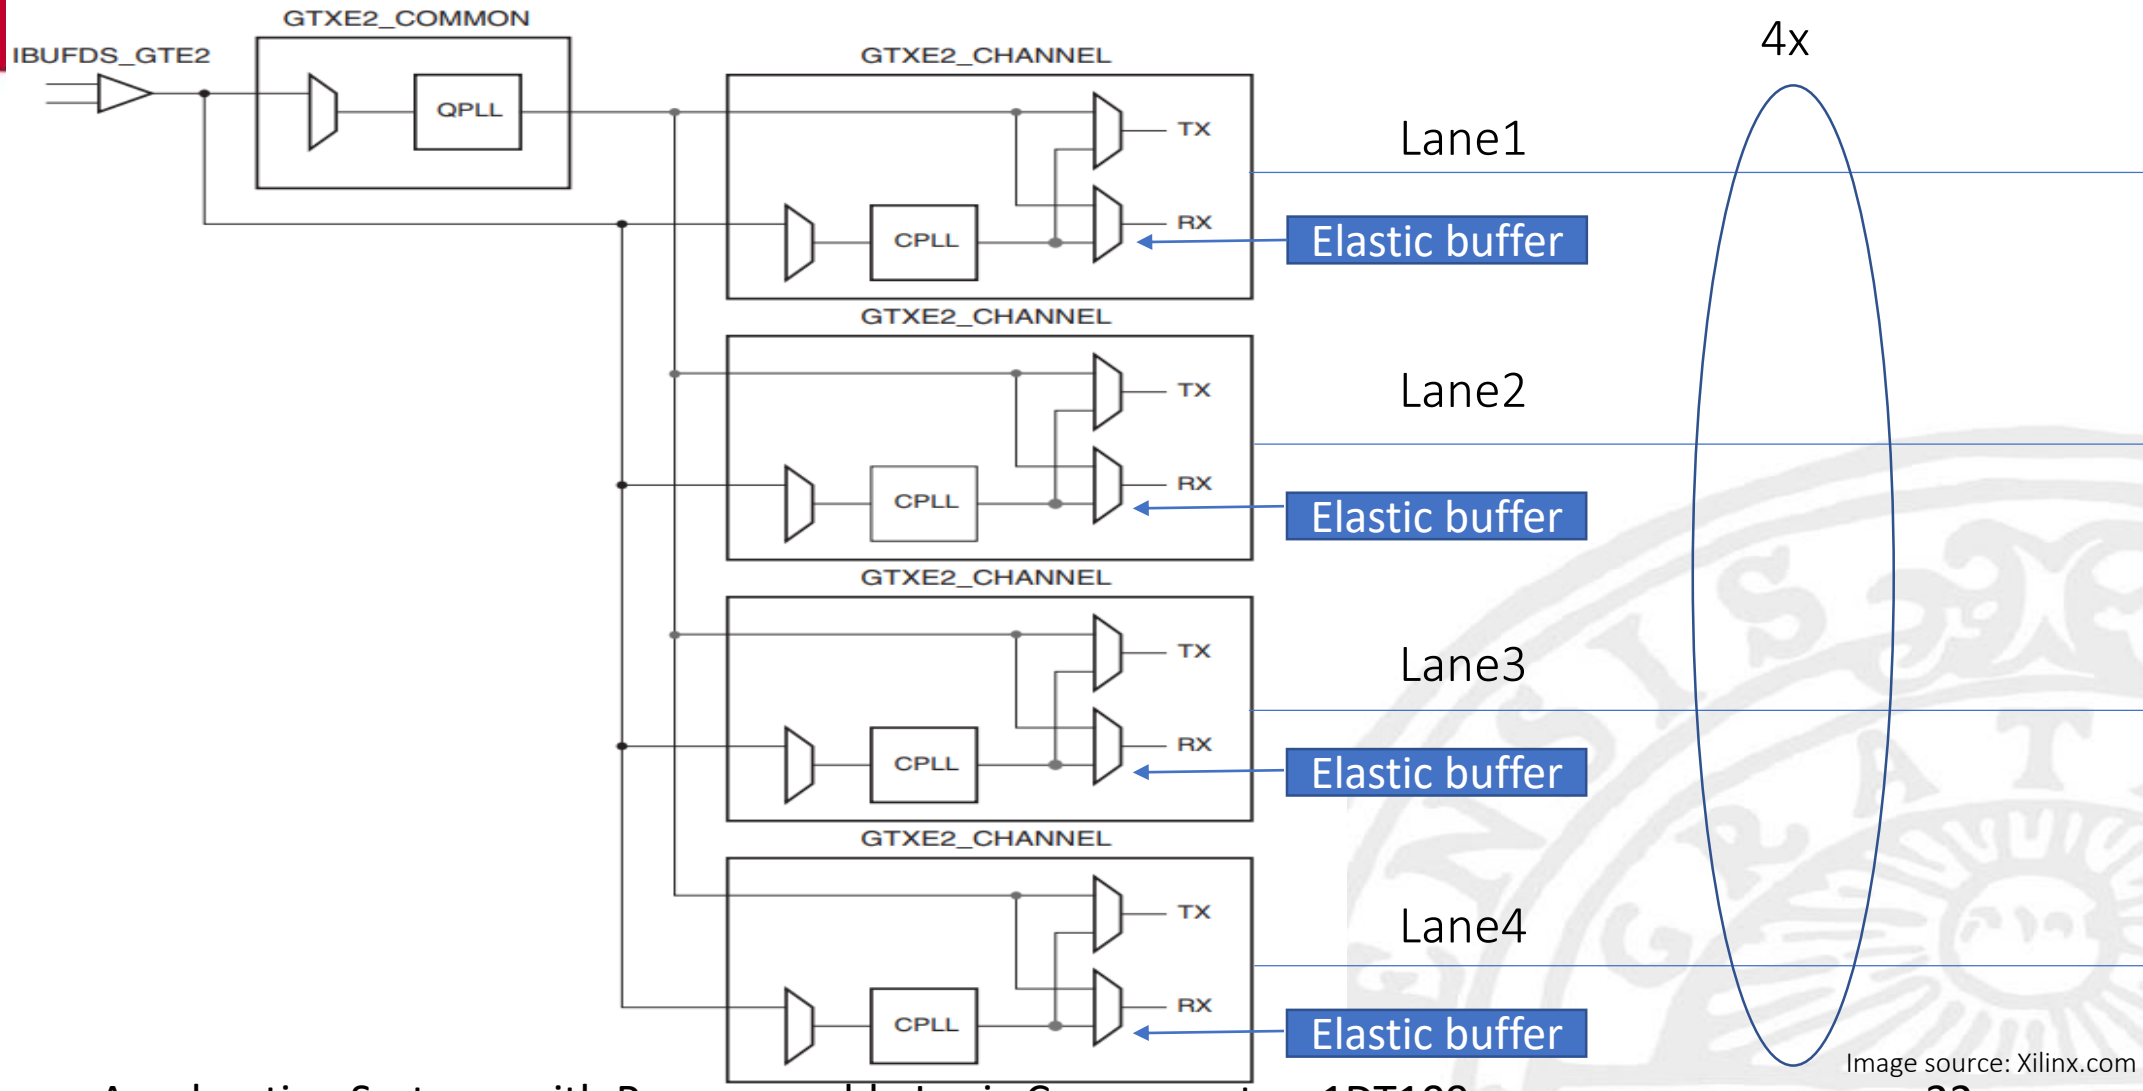

Image source: Xilinx.com

PCIe architecture

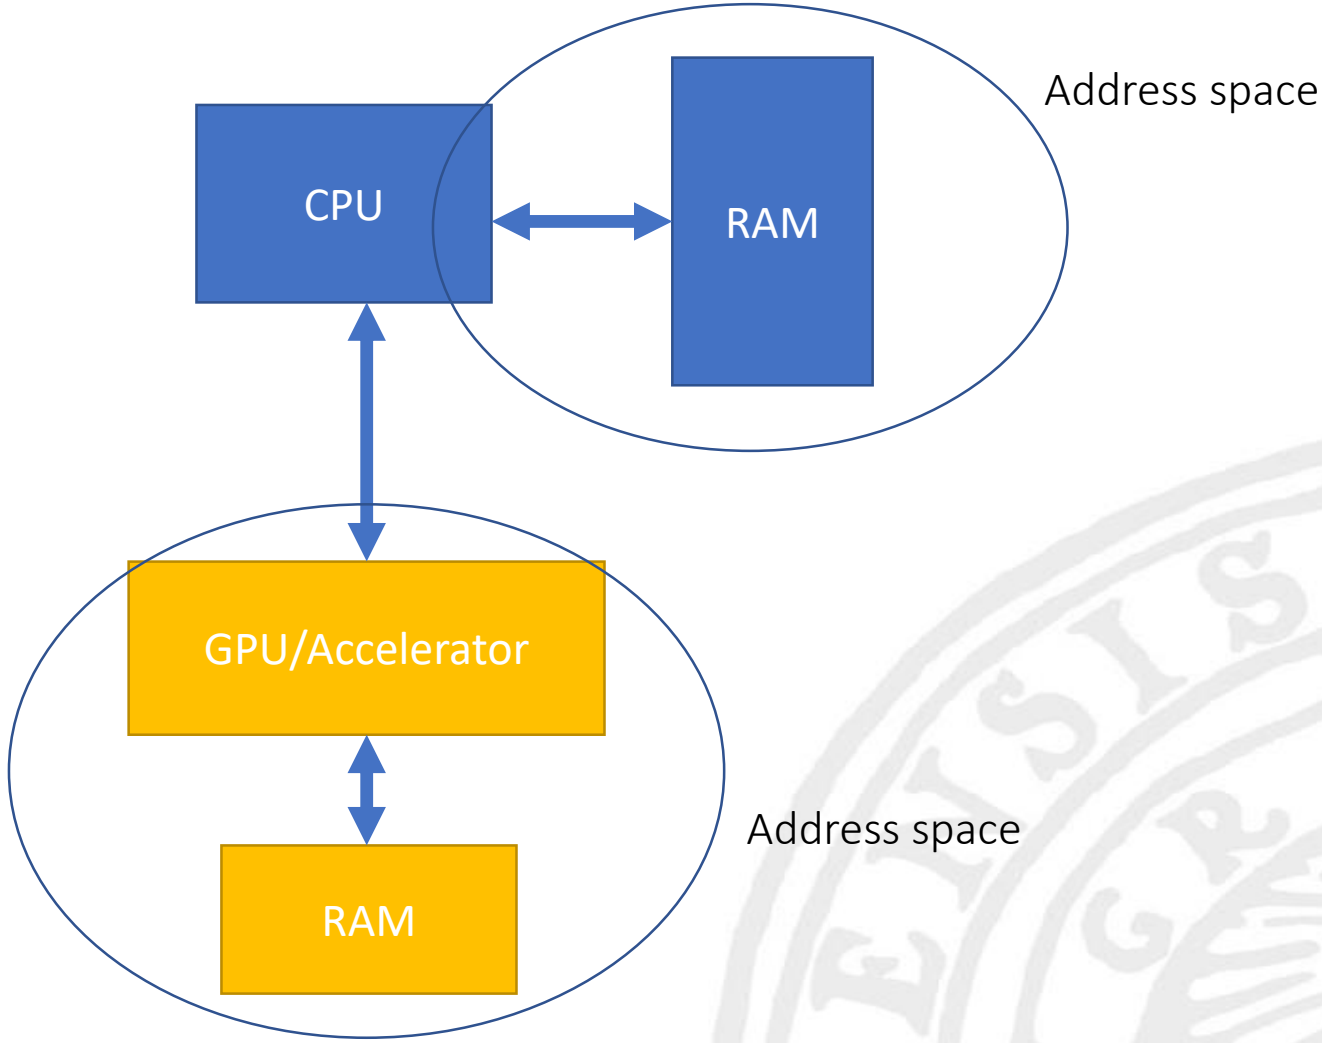

PCIe architecture

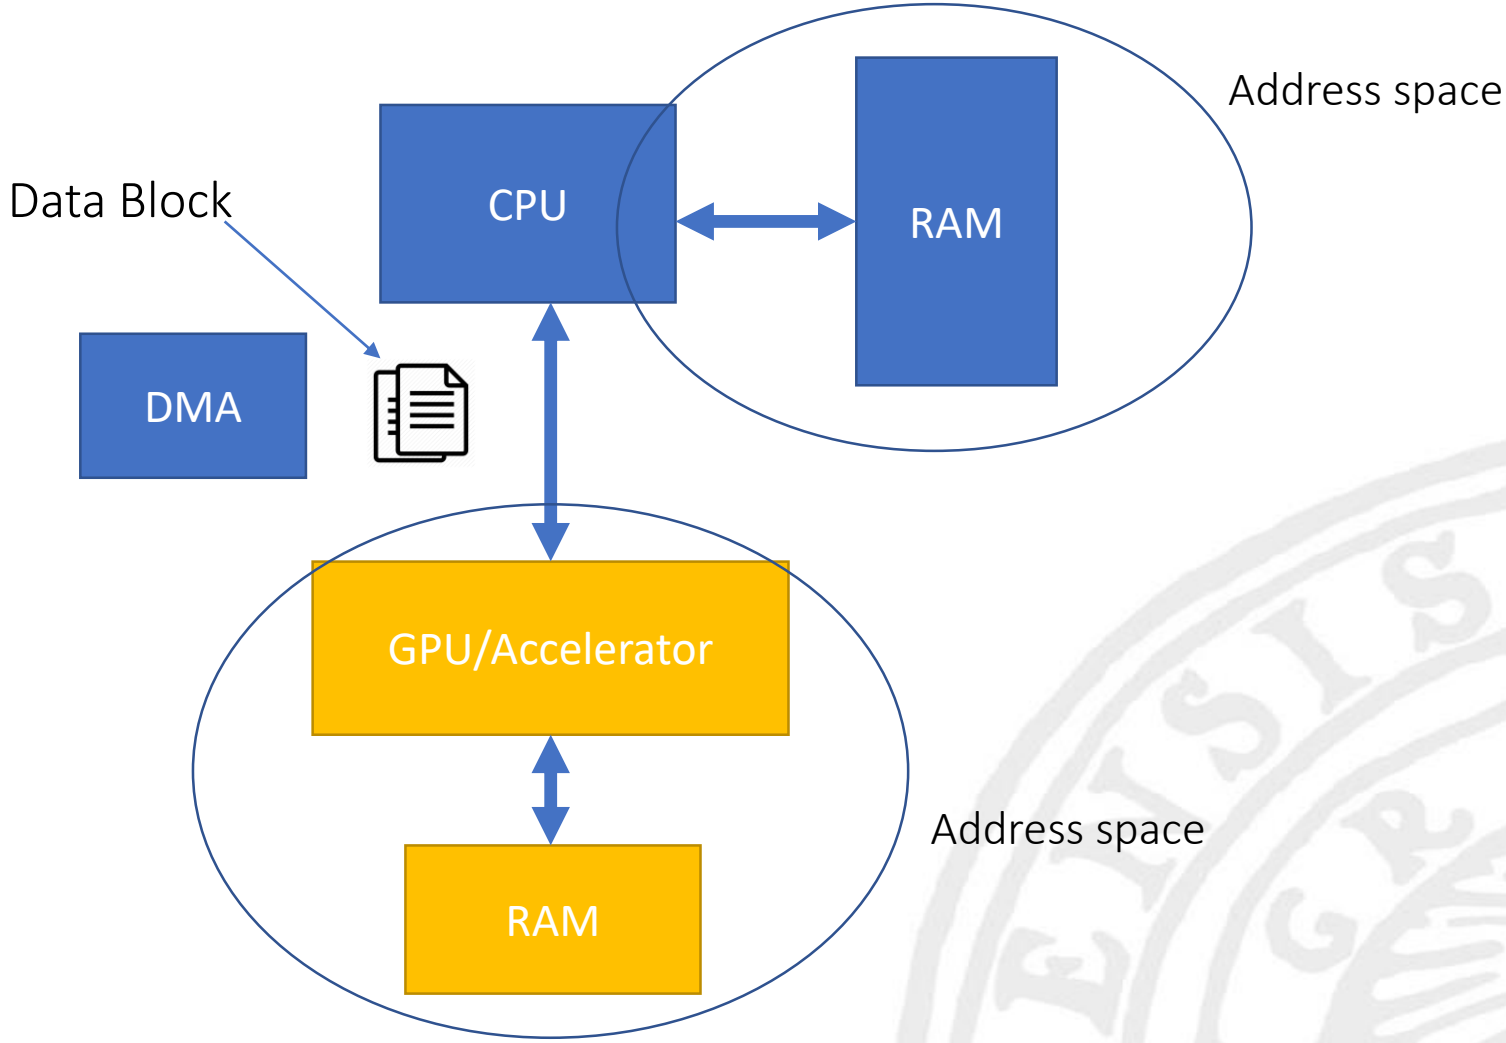

Memory Coherency

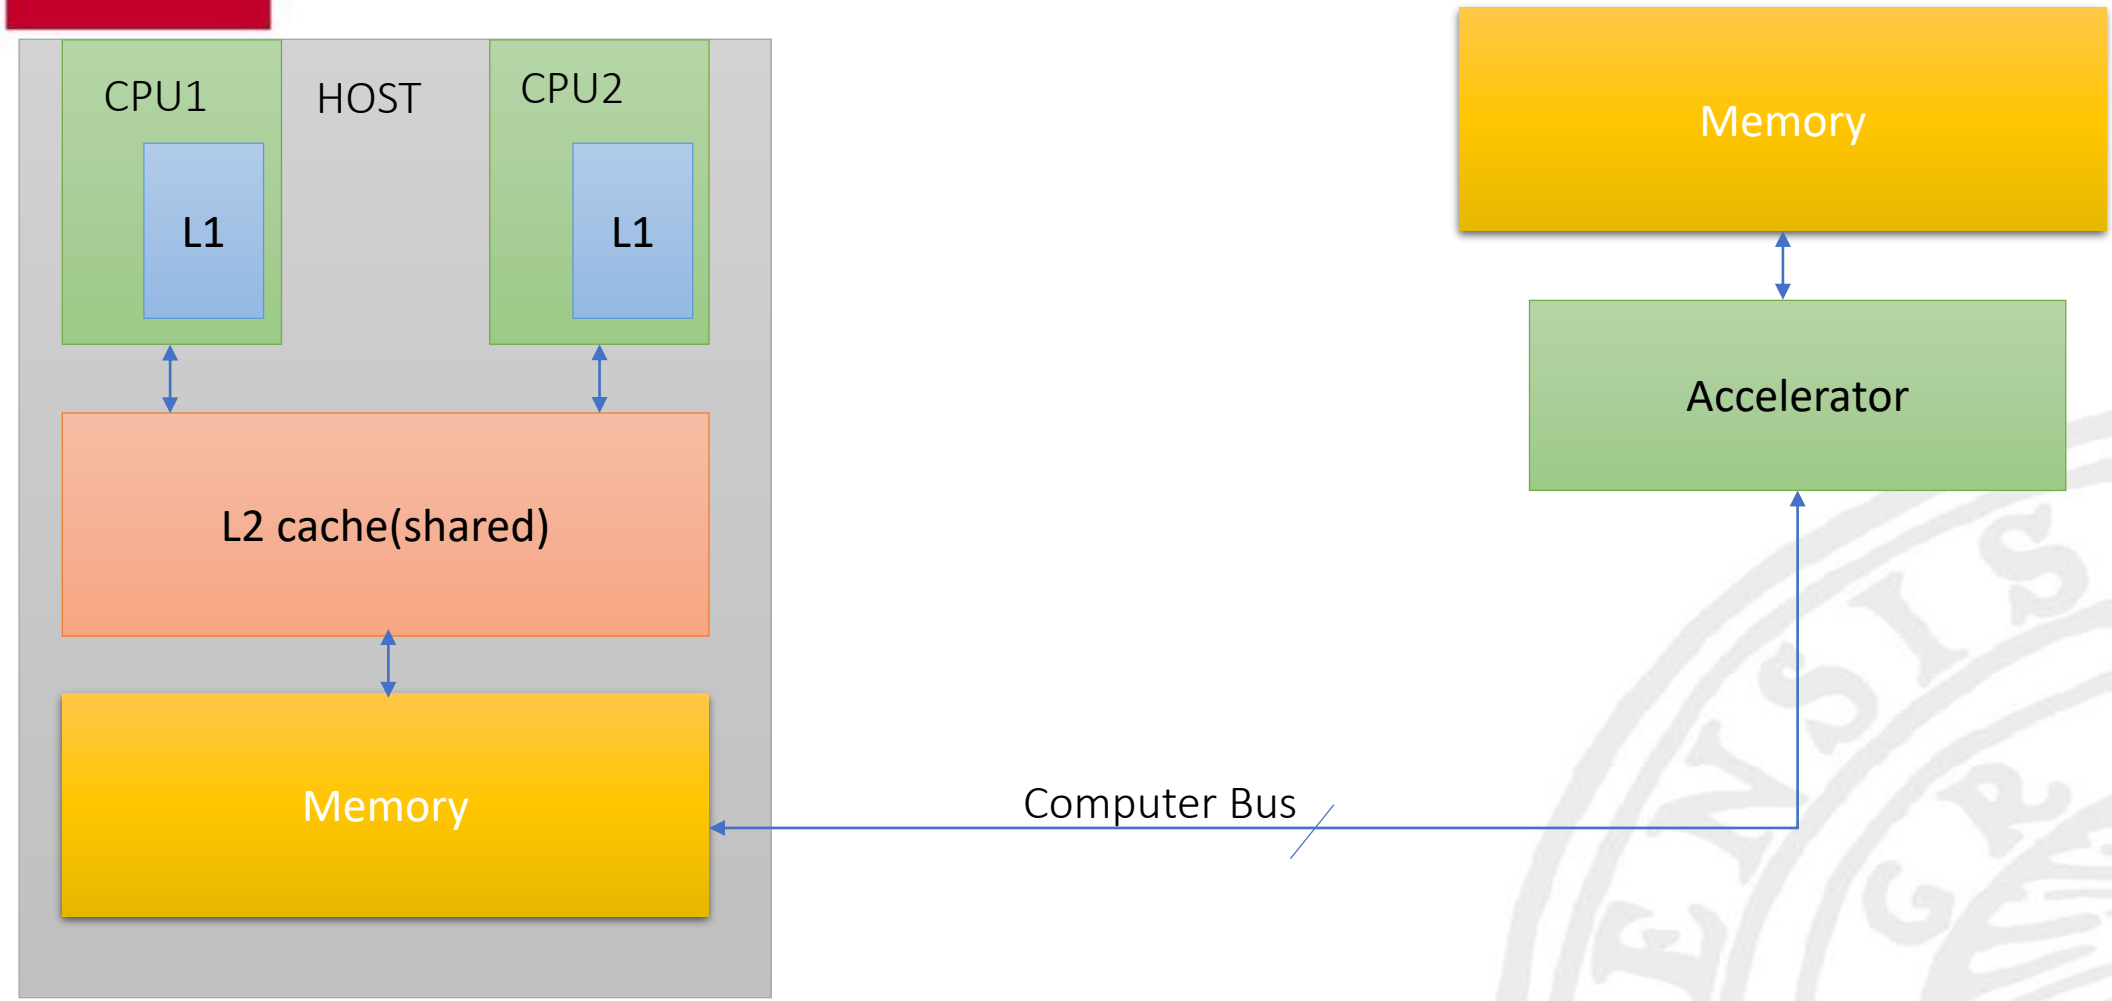

CCIX

- PCIe gen5 PHY
- NUMA
- No driver
- Complex logic
- Cache coherency

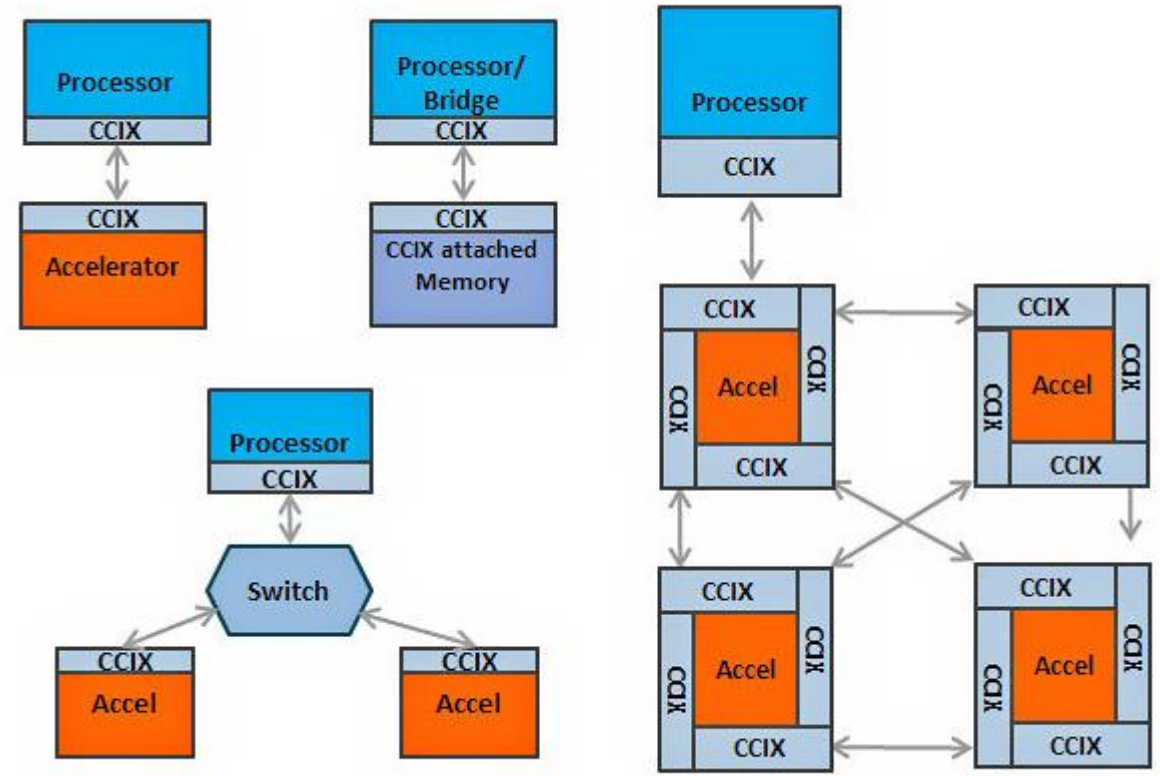

Good for autonomous driving

 Image source: Xilinx.com

CXL (2019)

- PCIe gen5 PHY
- CPU to accelerator
- Basic logic
- Root complexity is in CPU
- Server level

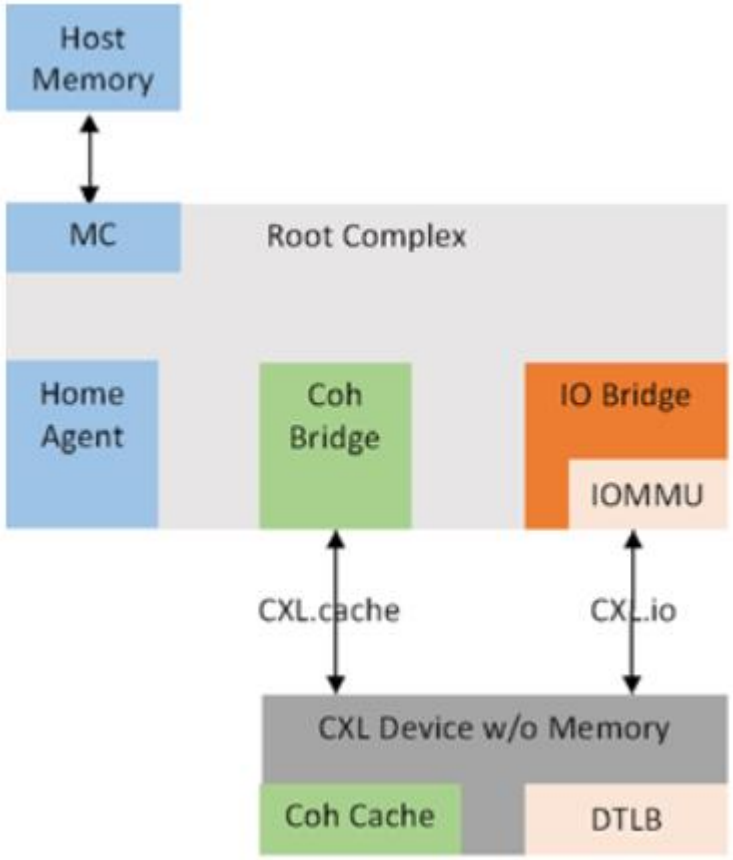

OpenCAPI

- Bluelink
- POWER9-10
- Summit super computer
- 3.0 memory coherency
- 4.0 cache coherecy
- P2P CPU to accelerator

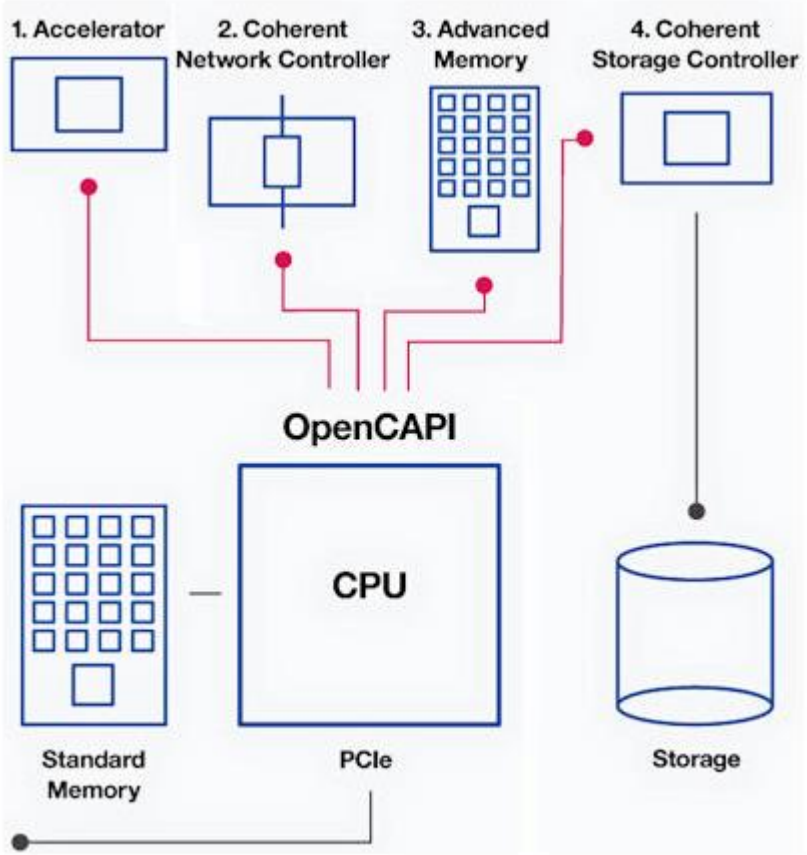

 Image source: Xilinx.com

Acronix accelerator

4 Tbps of GDDR6 Bandwidth

1× 200GbE Breakout Options:
- 2× 100GbE
- 4× 10/25/40/50 GbE

1× 400GbE Breakout Options:
- 2× 200GbE
- 4× 100GbE
- 8× 10/25/40/50 GbE

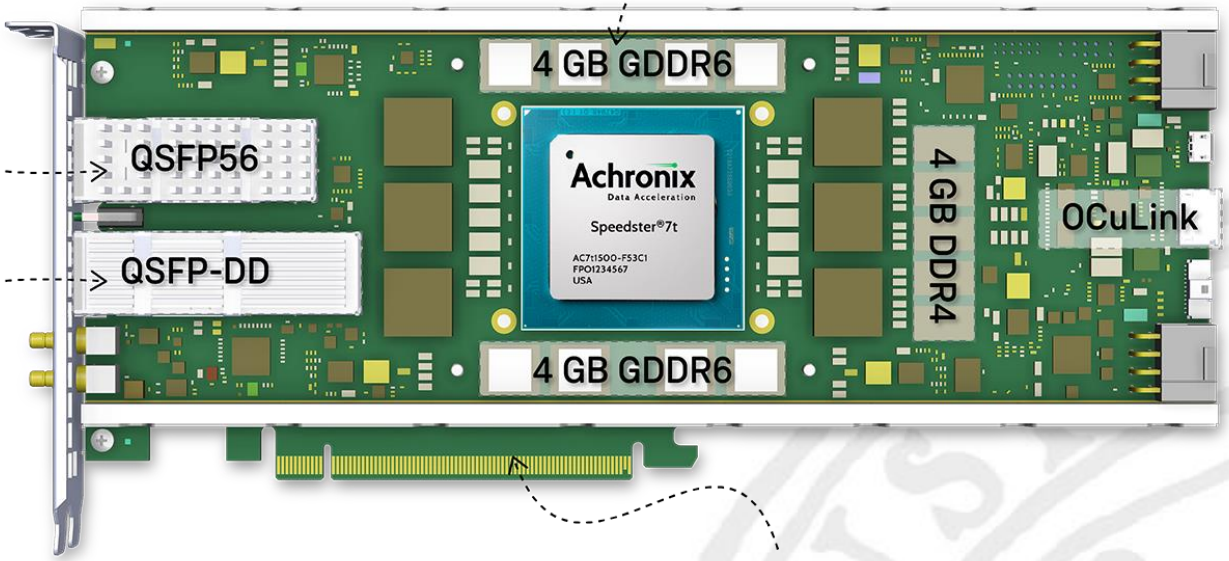

Ethernet Interfaces

Hard MAC and FEC IP for a range of industry-standard protocols and line rates

PCIe Gen3 ×16
- Speedster7t FPGAs support Gen4 and Gen5
- Card will support Gen4 with qualification
- Path to Gen5 with future versions

 Image source: Acronix.org

Xilinx accelerator

- Gen4x8 with CCIX
- 8GB HBM2
- 2x QSFP28 (100GbE)
- 2x 16GB 72b DIMM DDR4
- API: OpenCL

 Image source: Xilinx.org

V100 GPU for HPC

- 300GB/S NVLINK 2.0 cache coherent
- Summit super computer
- 7.8 DTFLOPS
- 15.7 STFLOPS
- 32GB HBM2
- API: CUDA, OpenCL,…

 Image source: Nvidia.org

Usage

- Edge computing
- Storage
- Database
- Load balancing
- Firewall
- Smart NIC
- Cryptography
- HPC
- AI