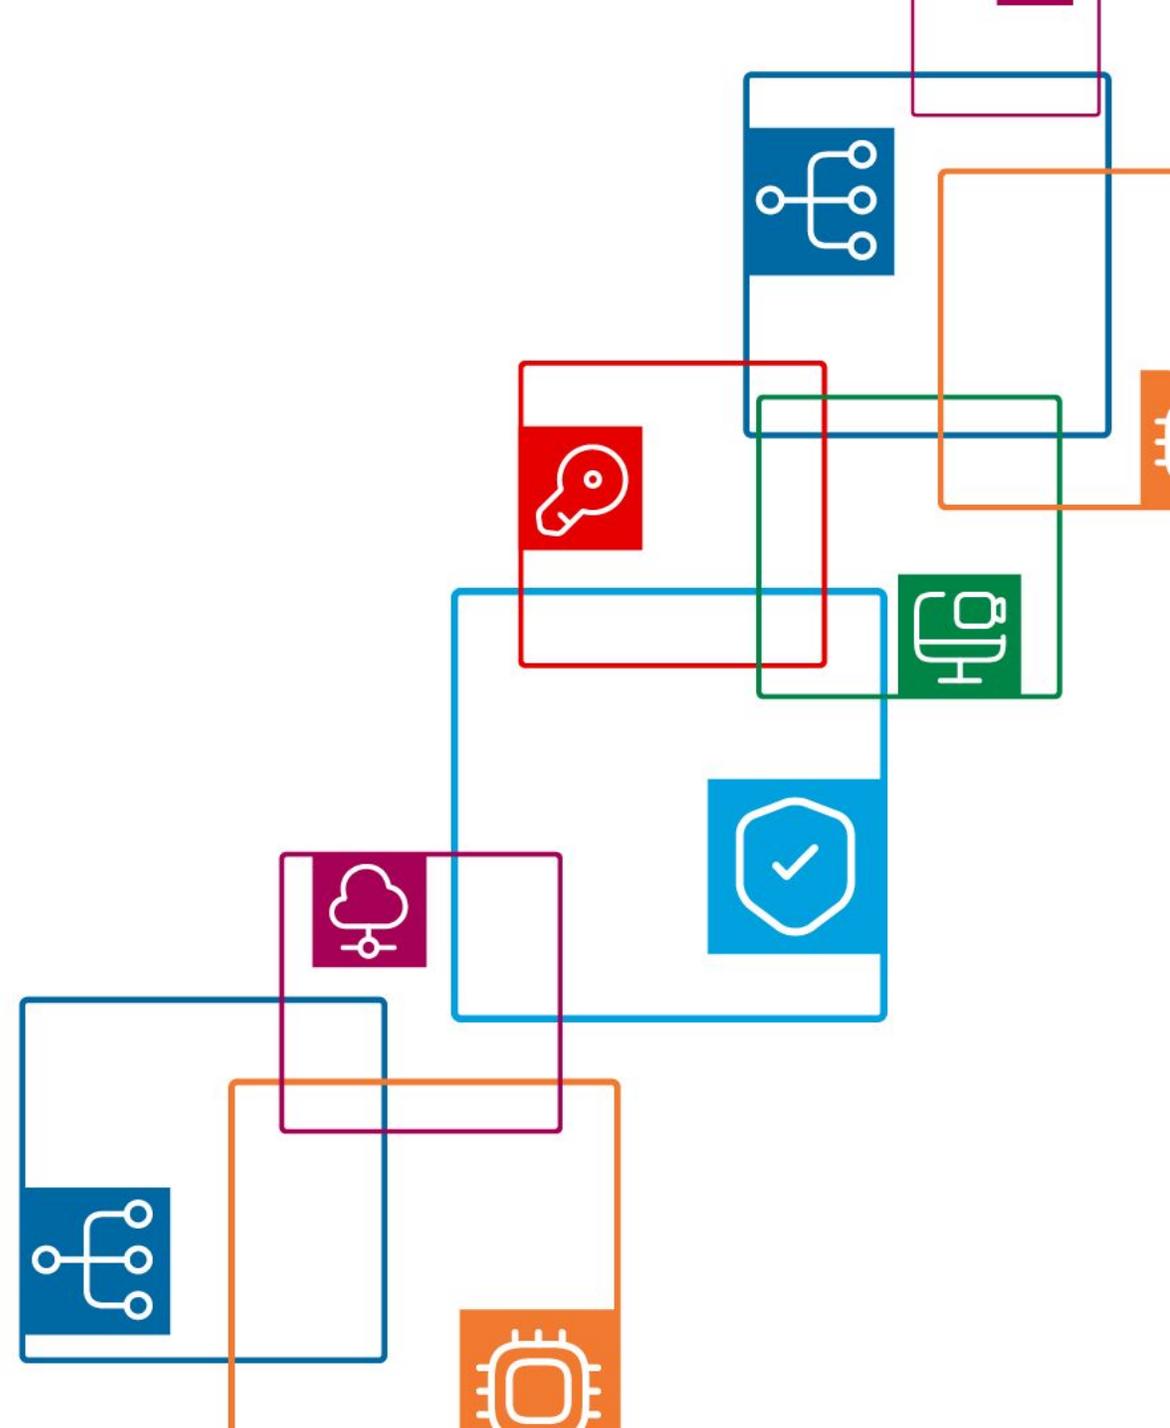




## Network Development Kit (NDK)

An open-source framework for high-speed network applications on FPGAs

Daniel Kondys  
kondys@cesnet.cz



- Introduction
- NDK on an FPGA-based SmartNIC
  - Top-level view
  - Core modules
  - Application - the user's playground
  - NDK buses
- NDK in use
- Conclusion

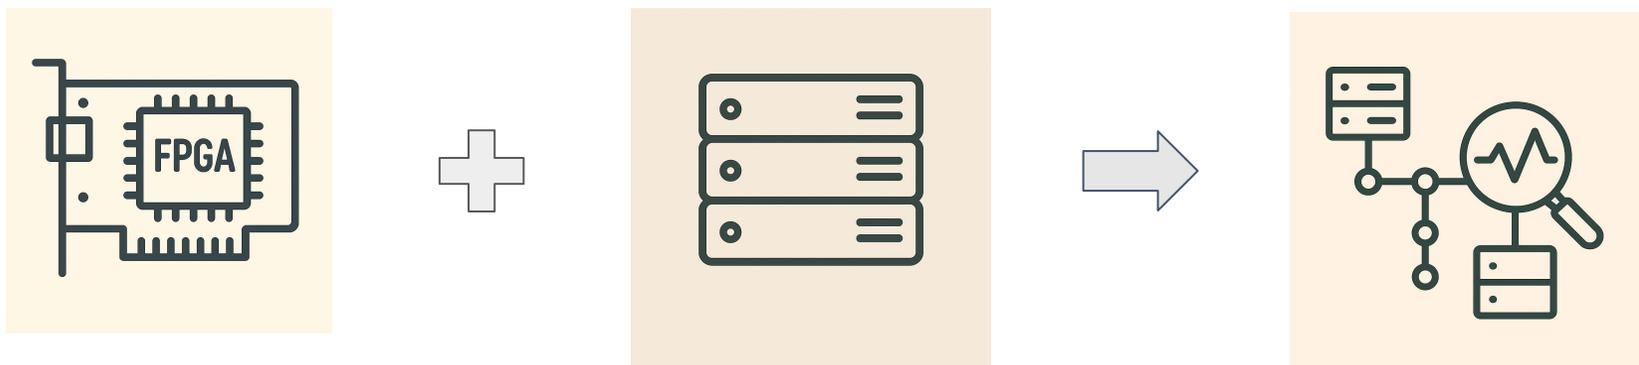


- **CESNET (Czech Education and Scientific NETwork)**
  - A collaboration of Czech public universities (founded in 1996)
  - Czech NREN (National Research and Education Network)
    - Develops and operates of national e-infrastructure for science, research, development, and education
  - Part of the GEANT association

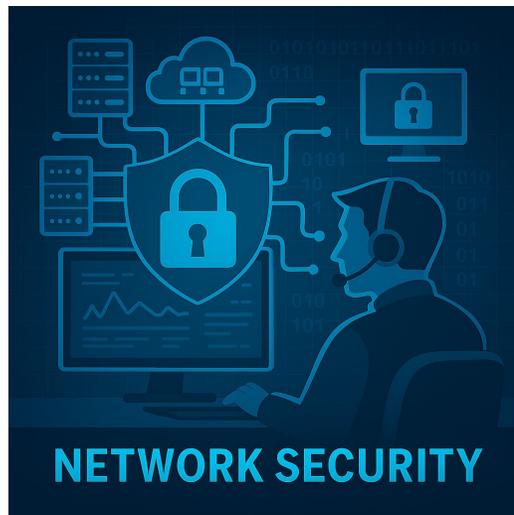




- CESNET's primary and long-term use case
- Using FPGA-based SmartNICs
  - network data pre-processing for flow monitoring
    - packet parsing
    - flow-hashing
    - filtering
    - ...
- The foundation of the NDK framework

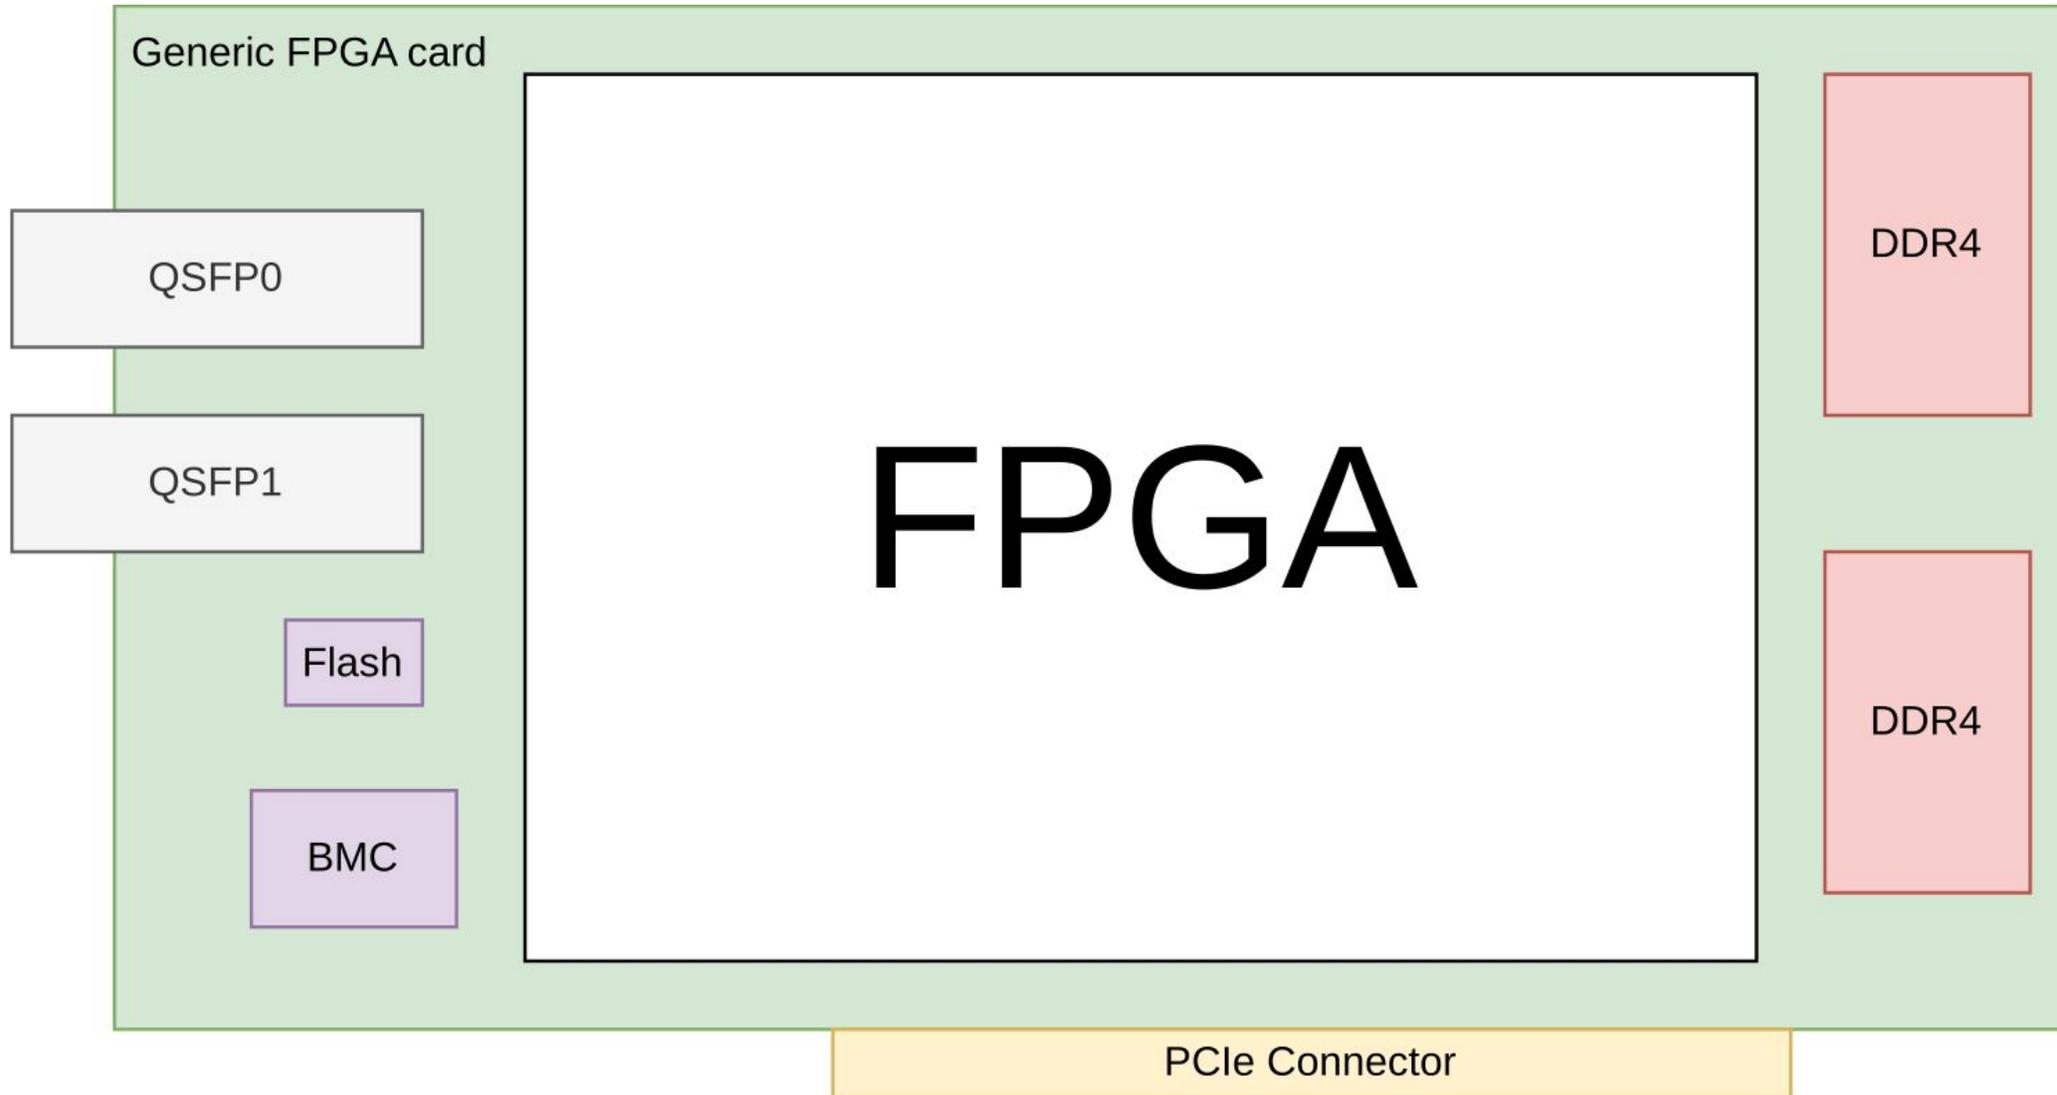


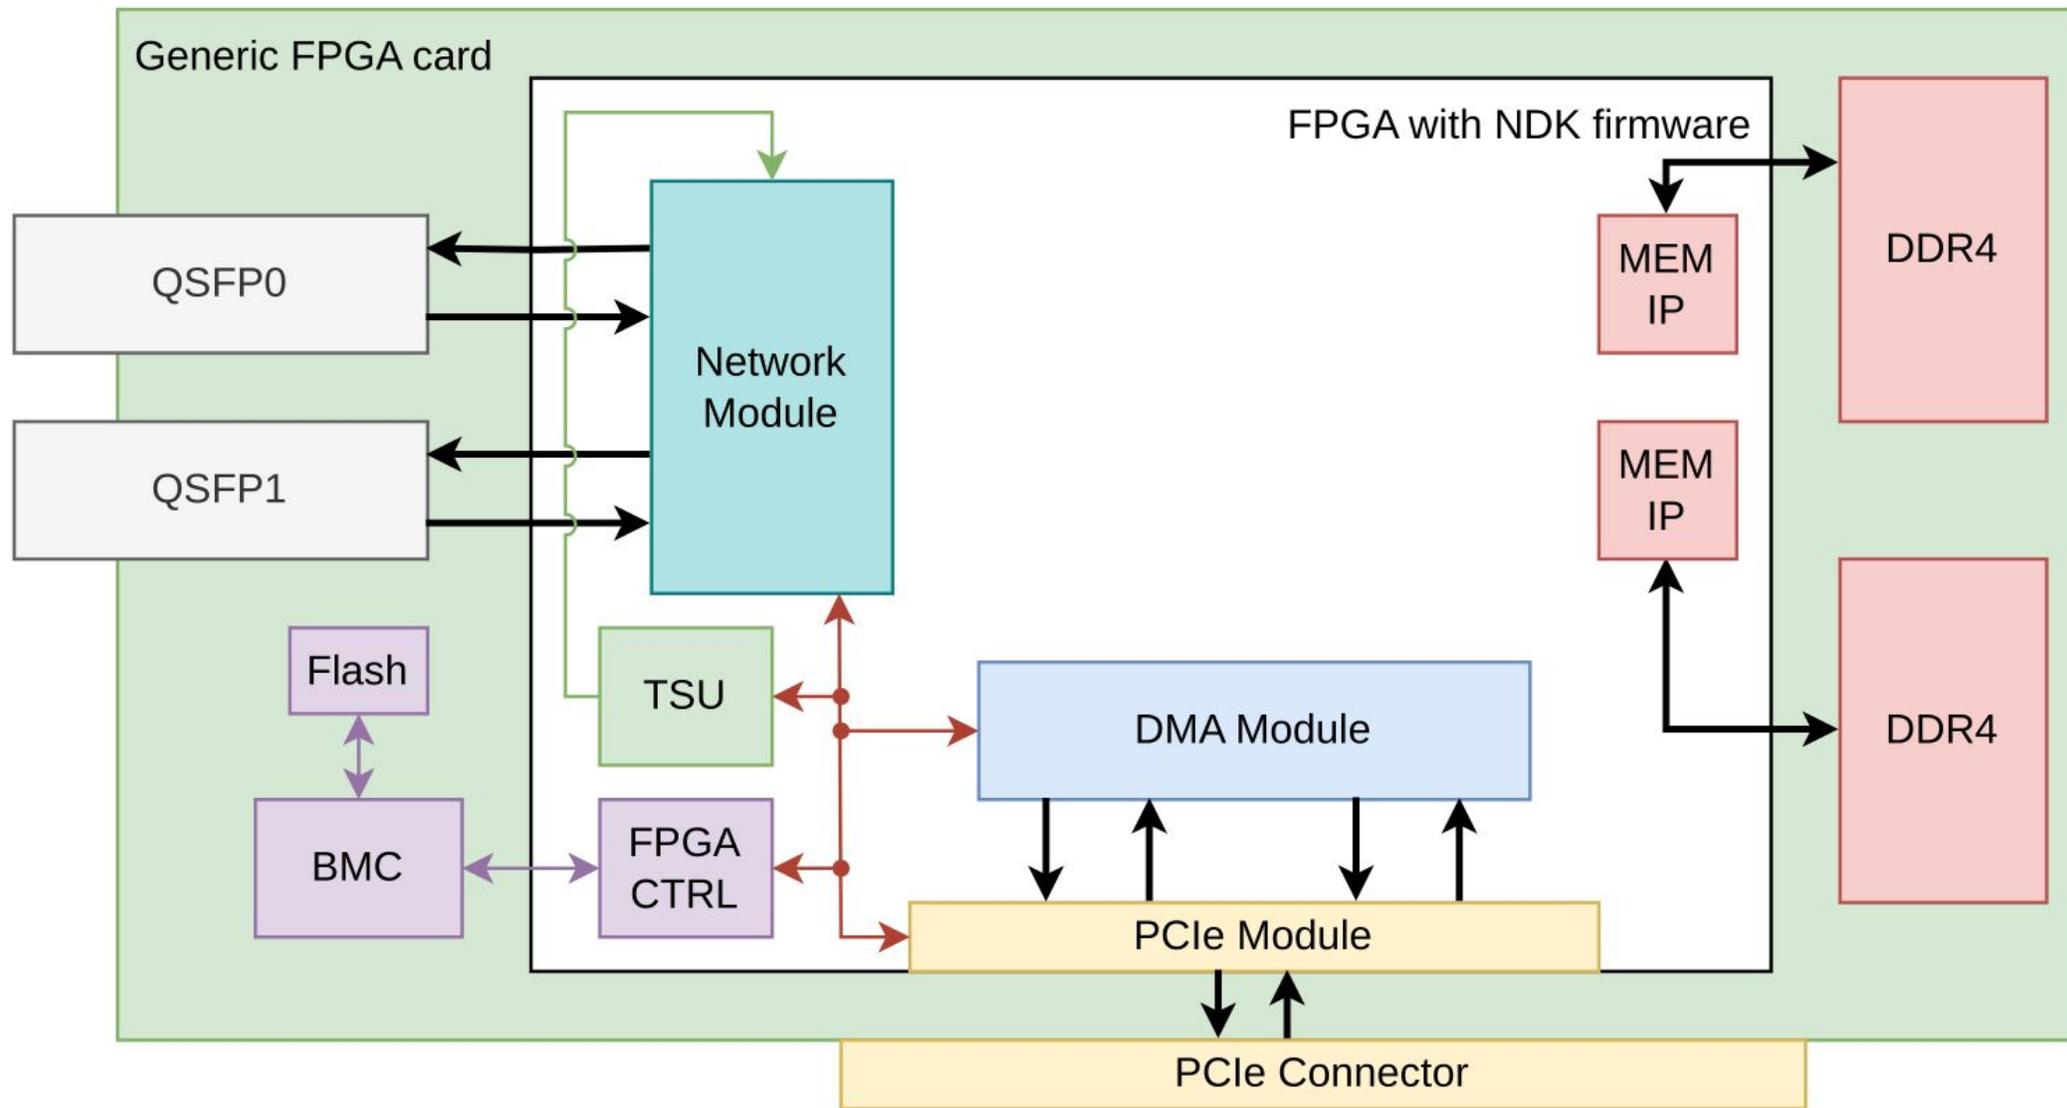
- Network security
  - IDS/IPS, DDoS protector acceleration
  - Filtering, Connection tracking, ...
- Testing network devices
  - Synthetic packet generation
  - [Realistic network traffic replay](#)
- High-Frequency Trading
  - Algorithmic trading on the stock exchange
  - Low-latency modules



# NDK on an FPGA-based SmartNIC

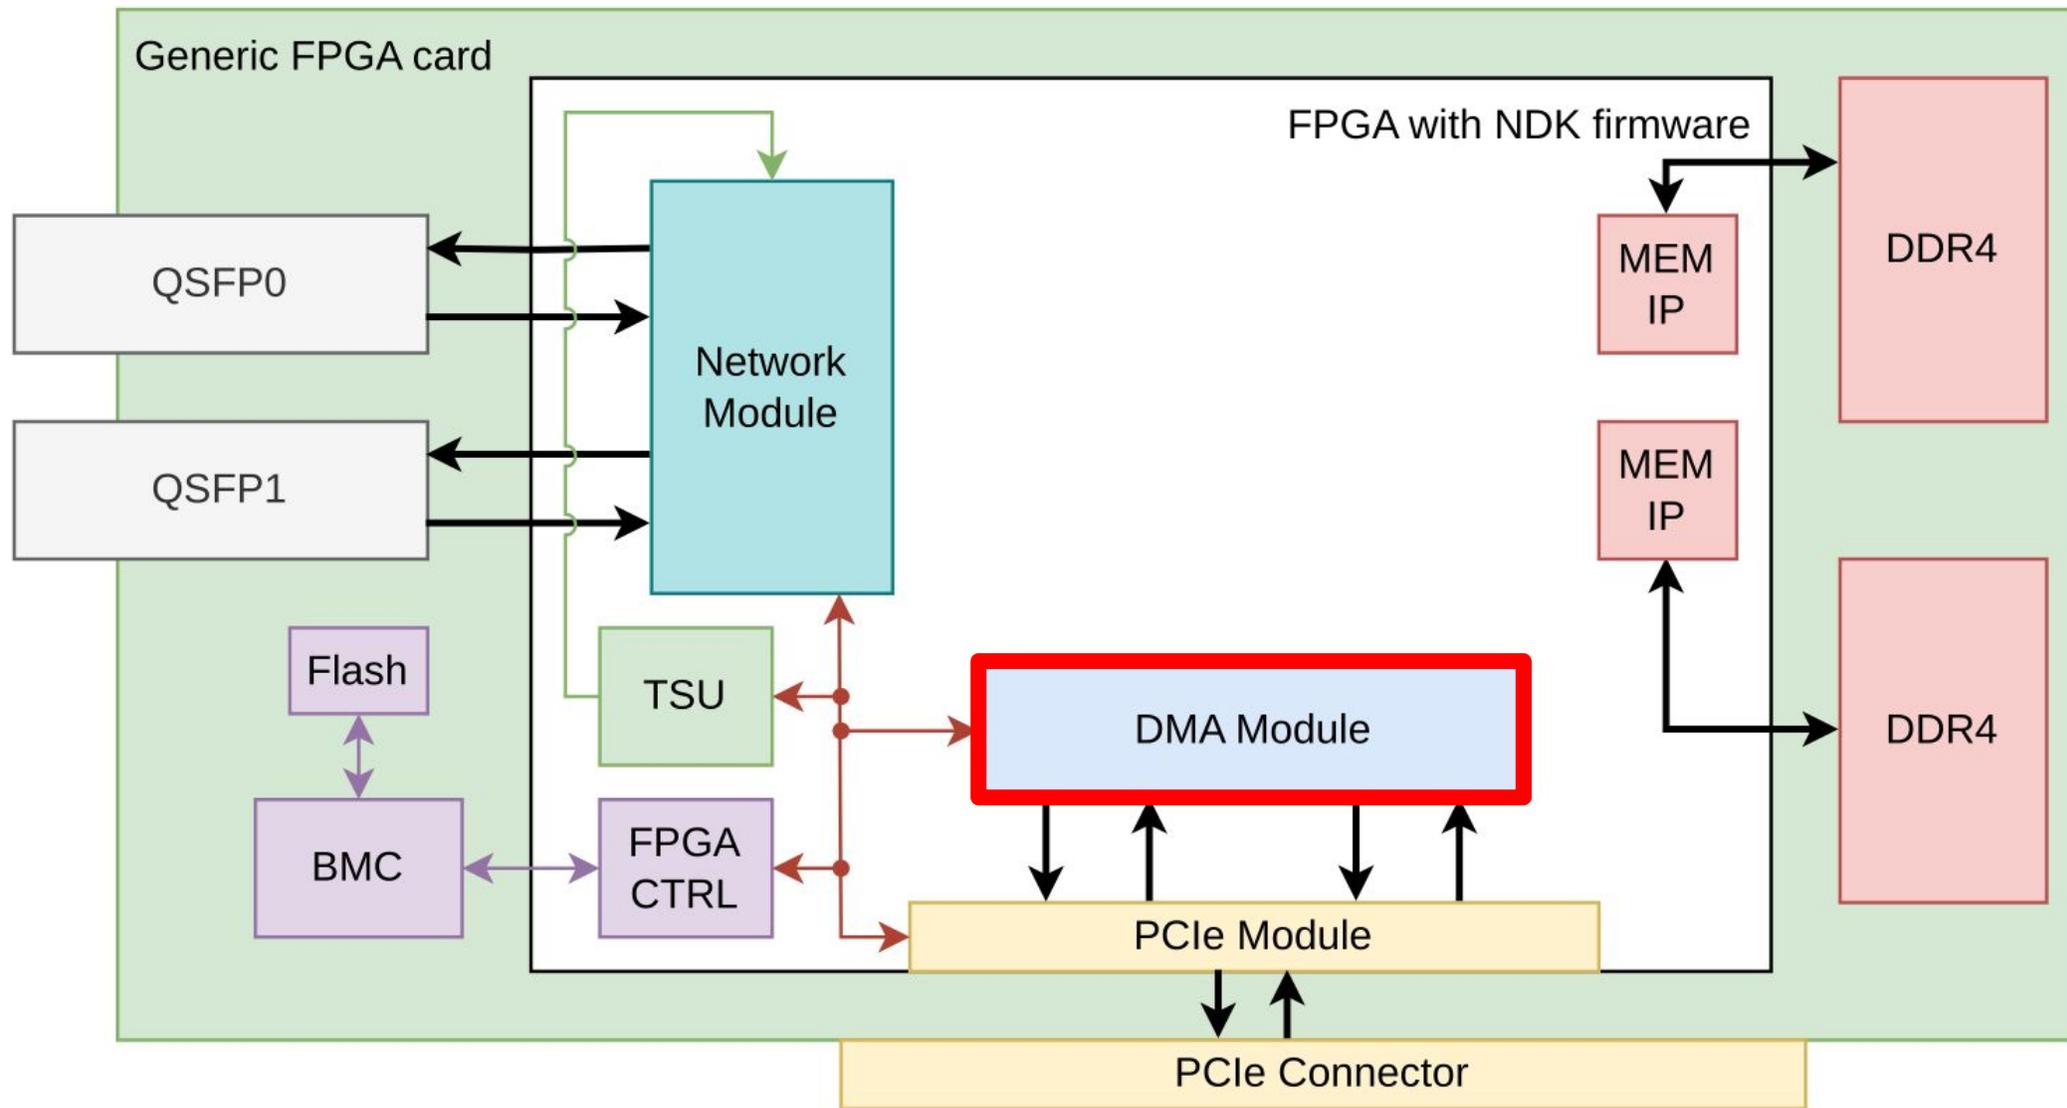






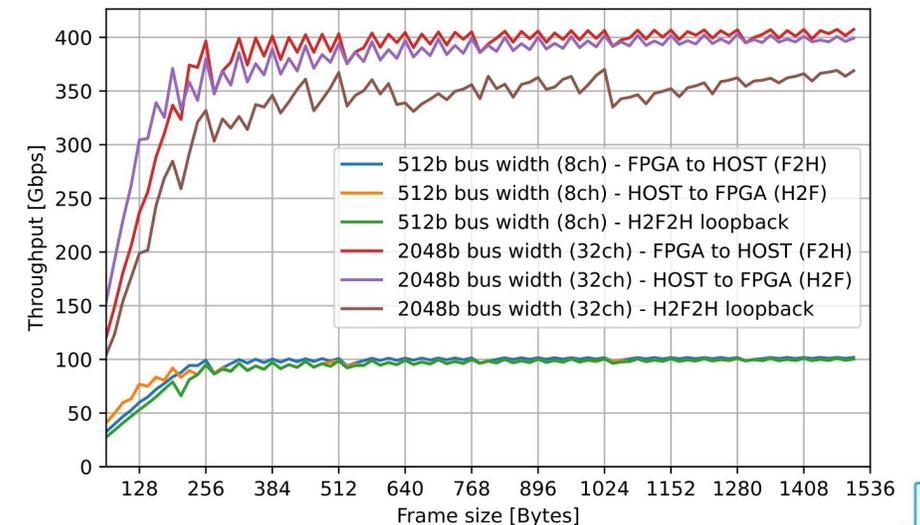
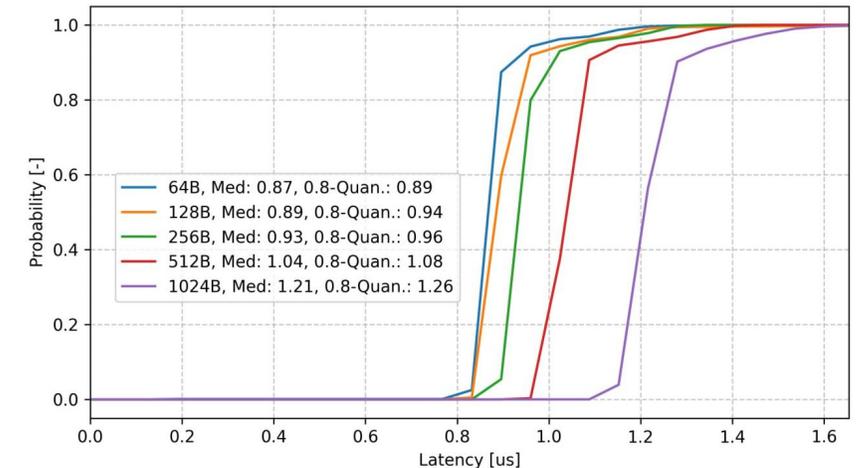
- [NDK-FPGA](#) repository on GitHub
- Flexible and easy-to-use build system
- Modules and components
  - in VHDL (2008)
  - support for Altera & AMD FPGAs
  - throughput speeds up to 400 Gbps (so far!)
- **Robust verification system**
  - in UVM (SystemVerilog) for modules and components
  - using the Cocotb framework for top-level (SW + FW) simulations
- [Wide range of supported FPGA-based NICs](#)
  - multiple speed Ethernet standards (from 10 to 400 GE)
  - up to PCIe Gen5 x16 including bifurcation (x8x8) support



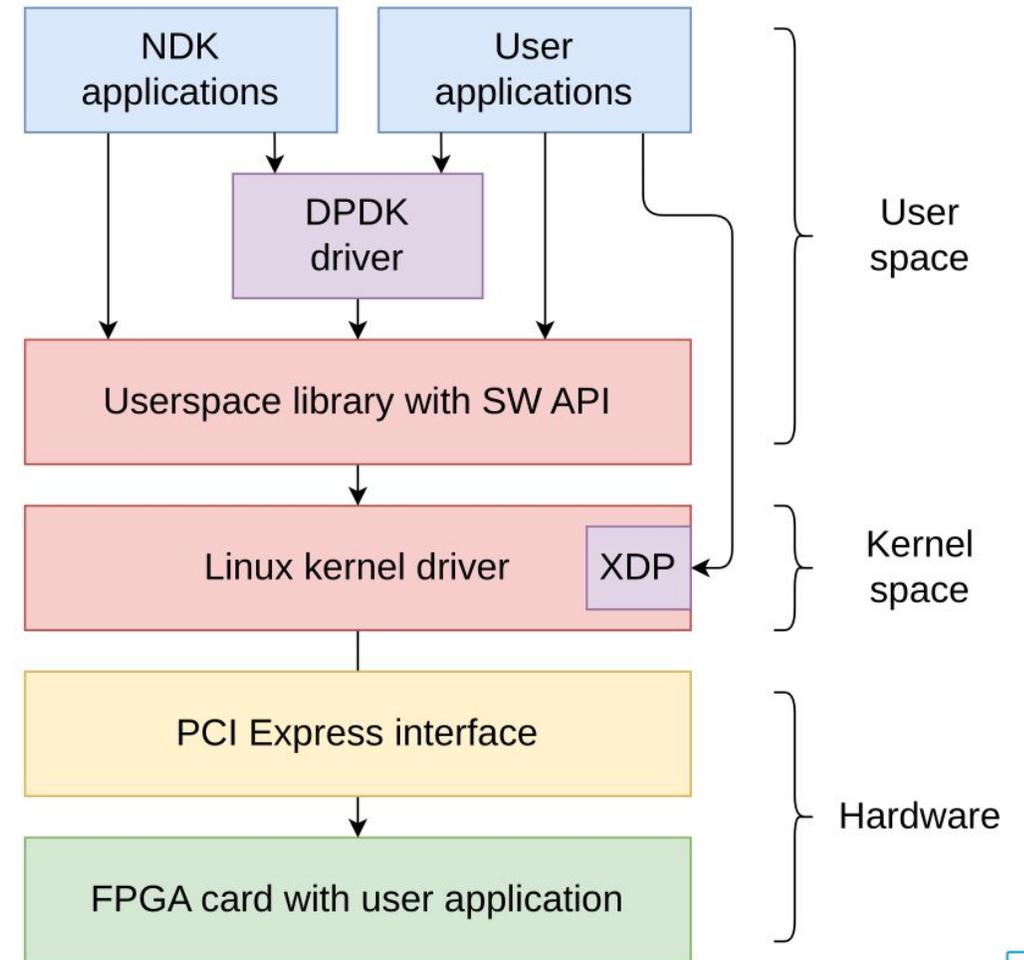


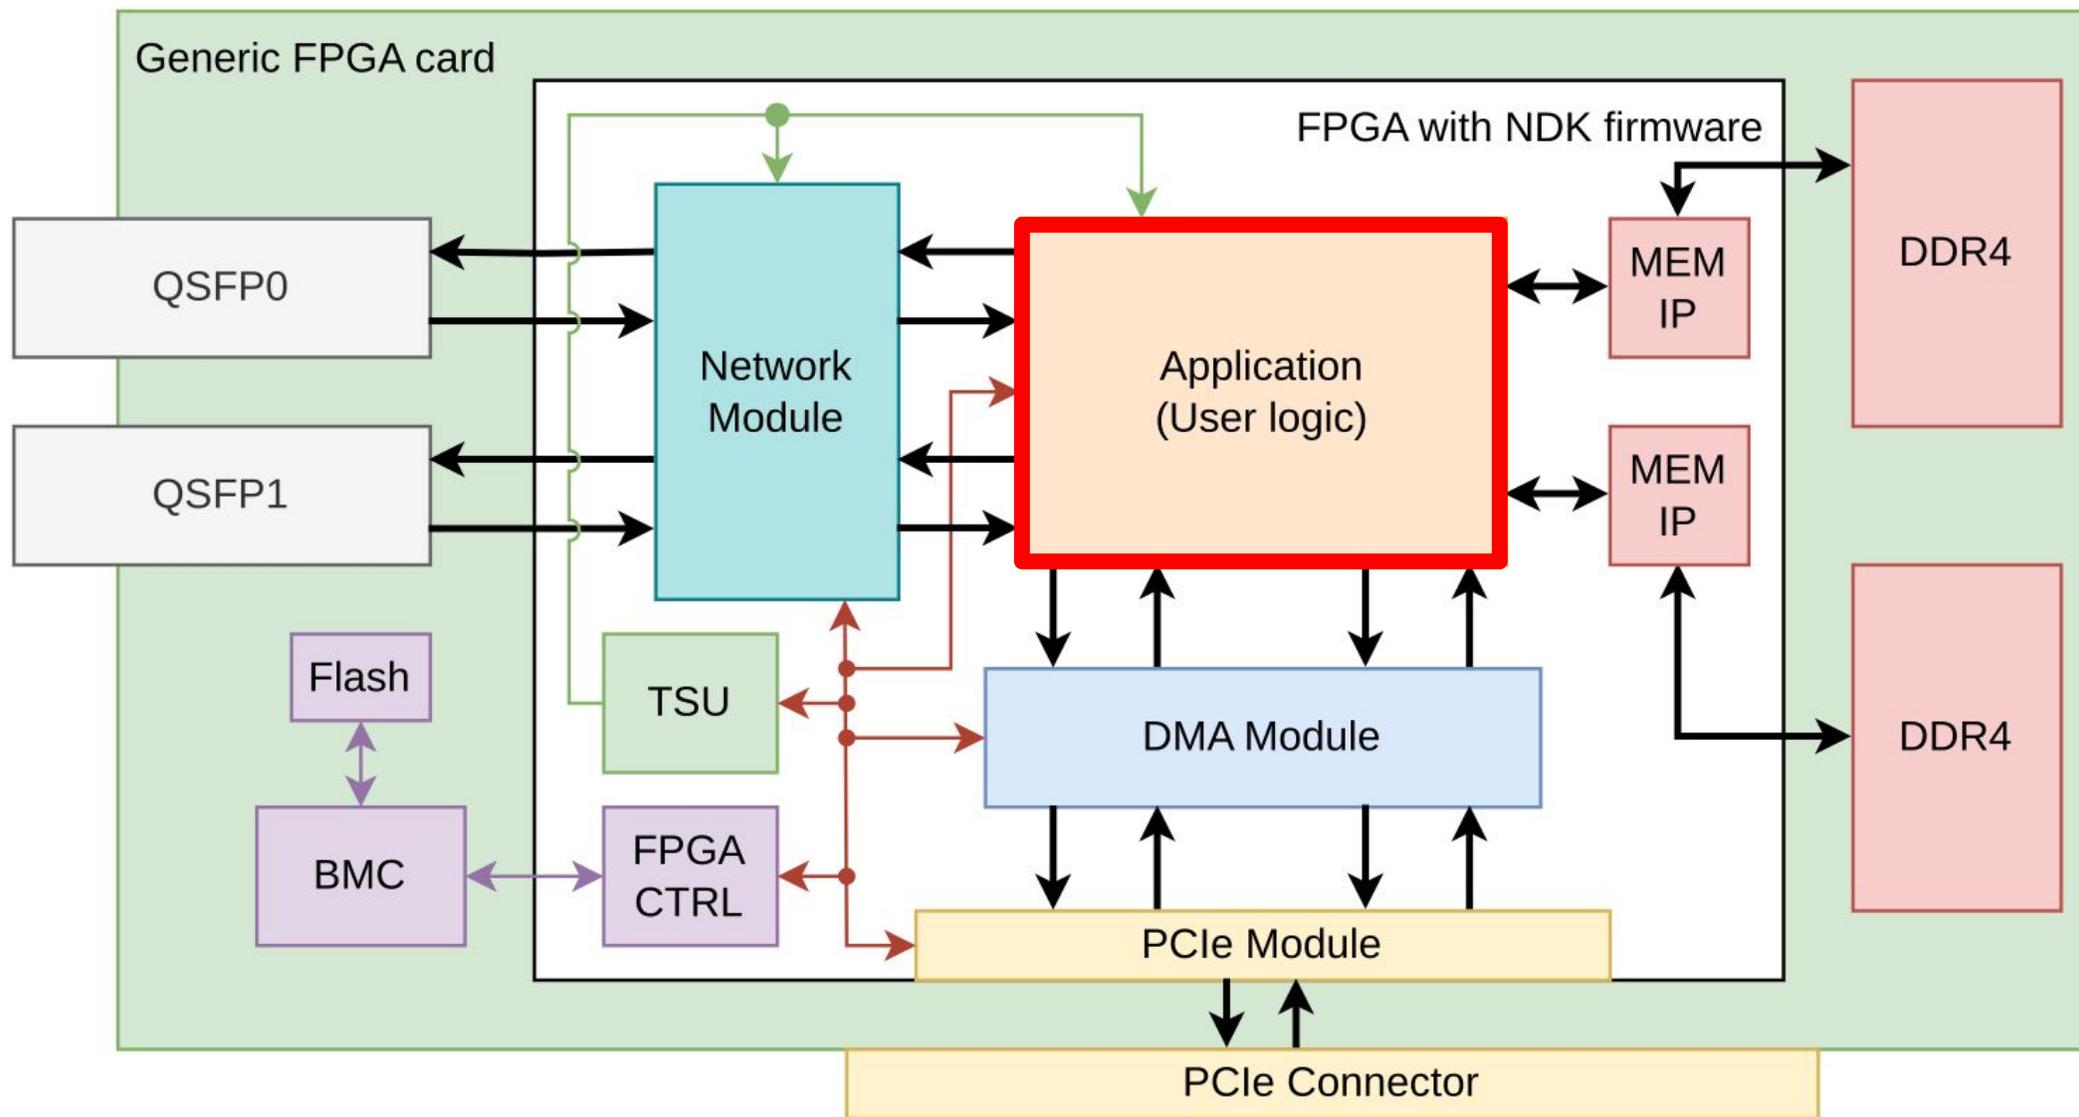
- Transfers packets to/from the host PC
- Two types
  - low-latency DMA engine
    - open-source
    - up to 100 Gbps
    - latency ~1 us
  - high-speed DMA engine
    - closed-source (available through [DynaNIC](#))
    - up to 400 Gbps
    - latency ~20 us

J. Cabal et. al., "FPL Demo: 400G FPGA Packet Capture Based on Network Development Kit," FPL 2022, Belfast, United Kingdom, 2022, pp. 474-474, doi: [10.1109/FPL57034.2022.00090](https://doi.org/10.1109/FPL57034.2022.00090).

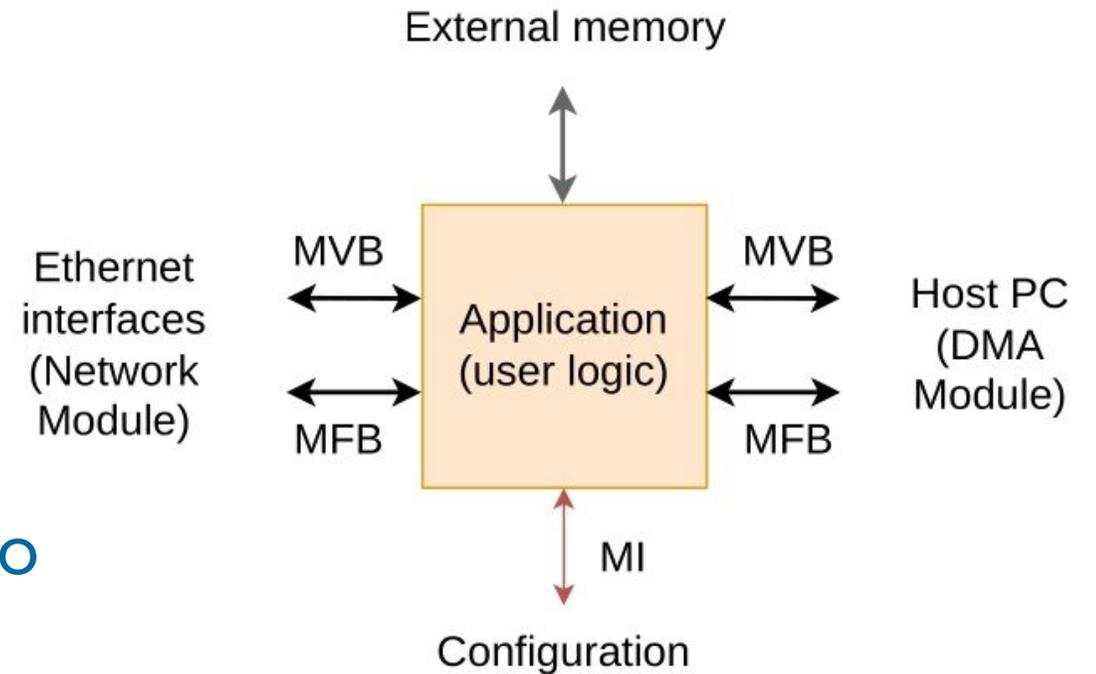


- [NDK-SW](#) repository on GitHub
- Linux kernel driver
  - nfb (NDK FPGA Board)
  - with NDP or XDP interface
- Userspace library
  - libnfb
  - with C and Python APIs
- DPDK driver
- SW applications

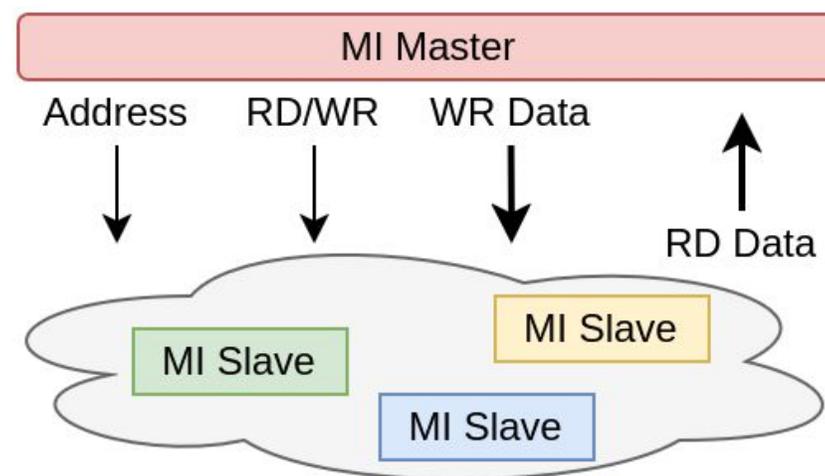




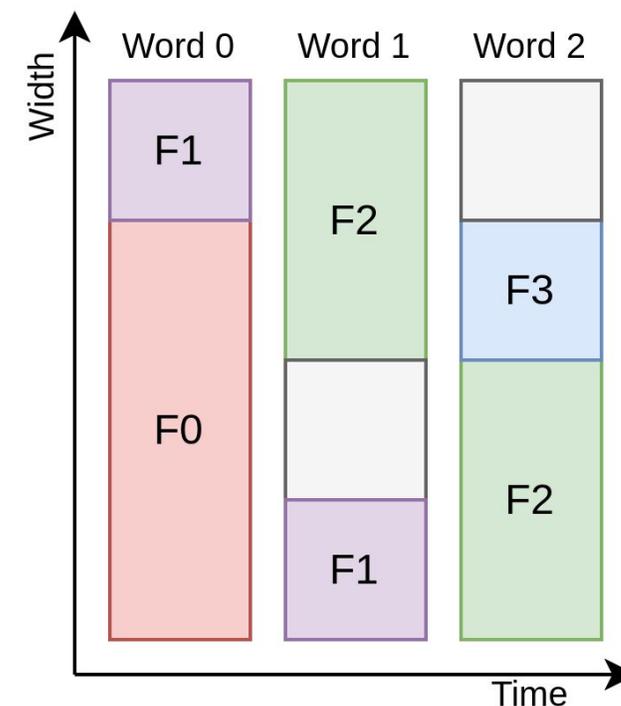
- User's playground
- Data interfaces to
  - Ethernet
  - DMA (-> host PC)
  - External memory
- Number of interfaces according to selected card
- SW access (configuration) through the MI bus



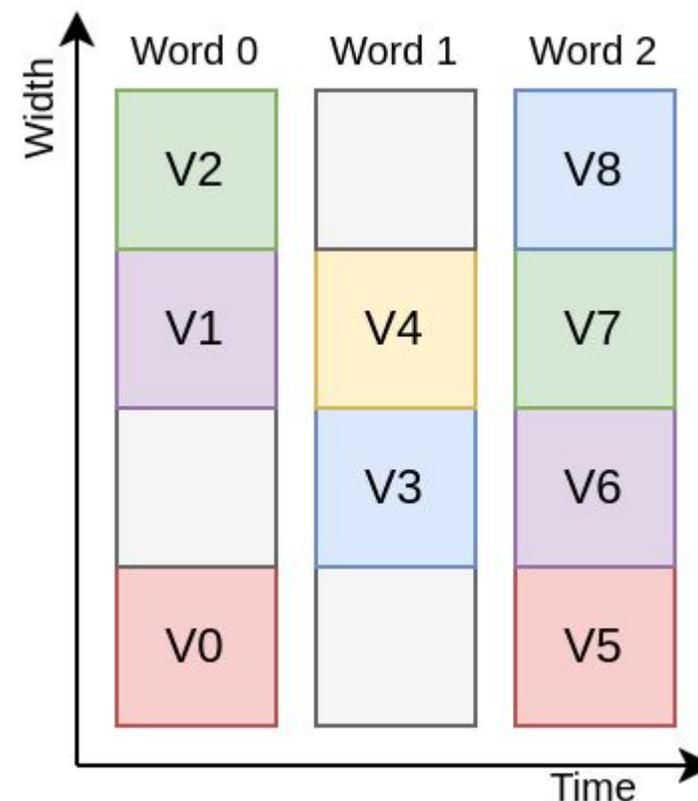
- Addressable Master-Slave bus
- SW access to Control, Status, and Statistics Registers
  - read amount of processed packets, enable packet filtration, ...
- Relatives: AXI4-Lite, Avalon Memory-Mapped
- Simple and resource-efficient
  - bursts are not supported
  - read responses in order

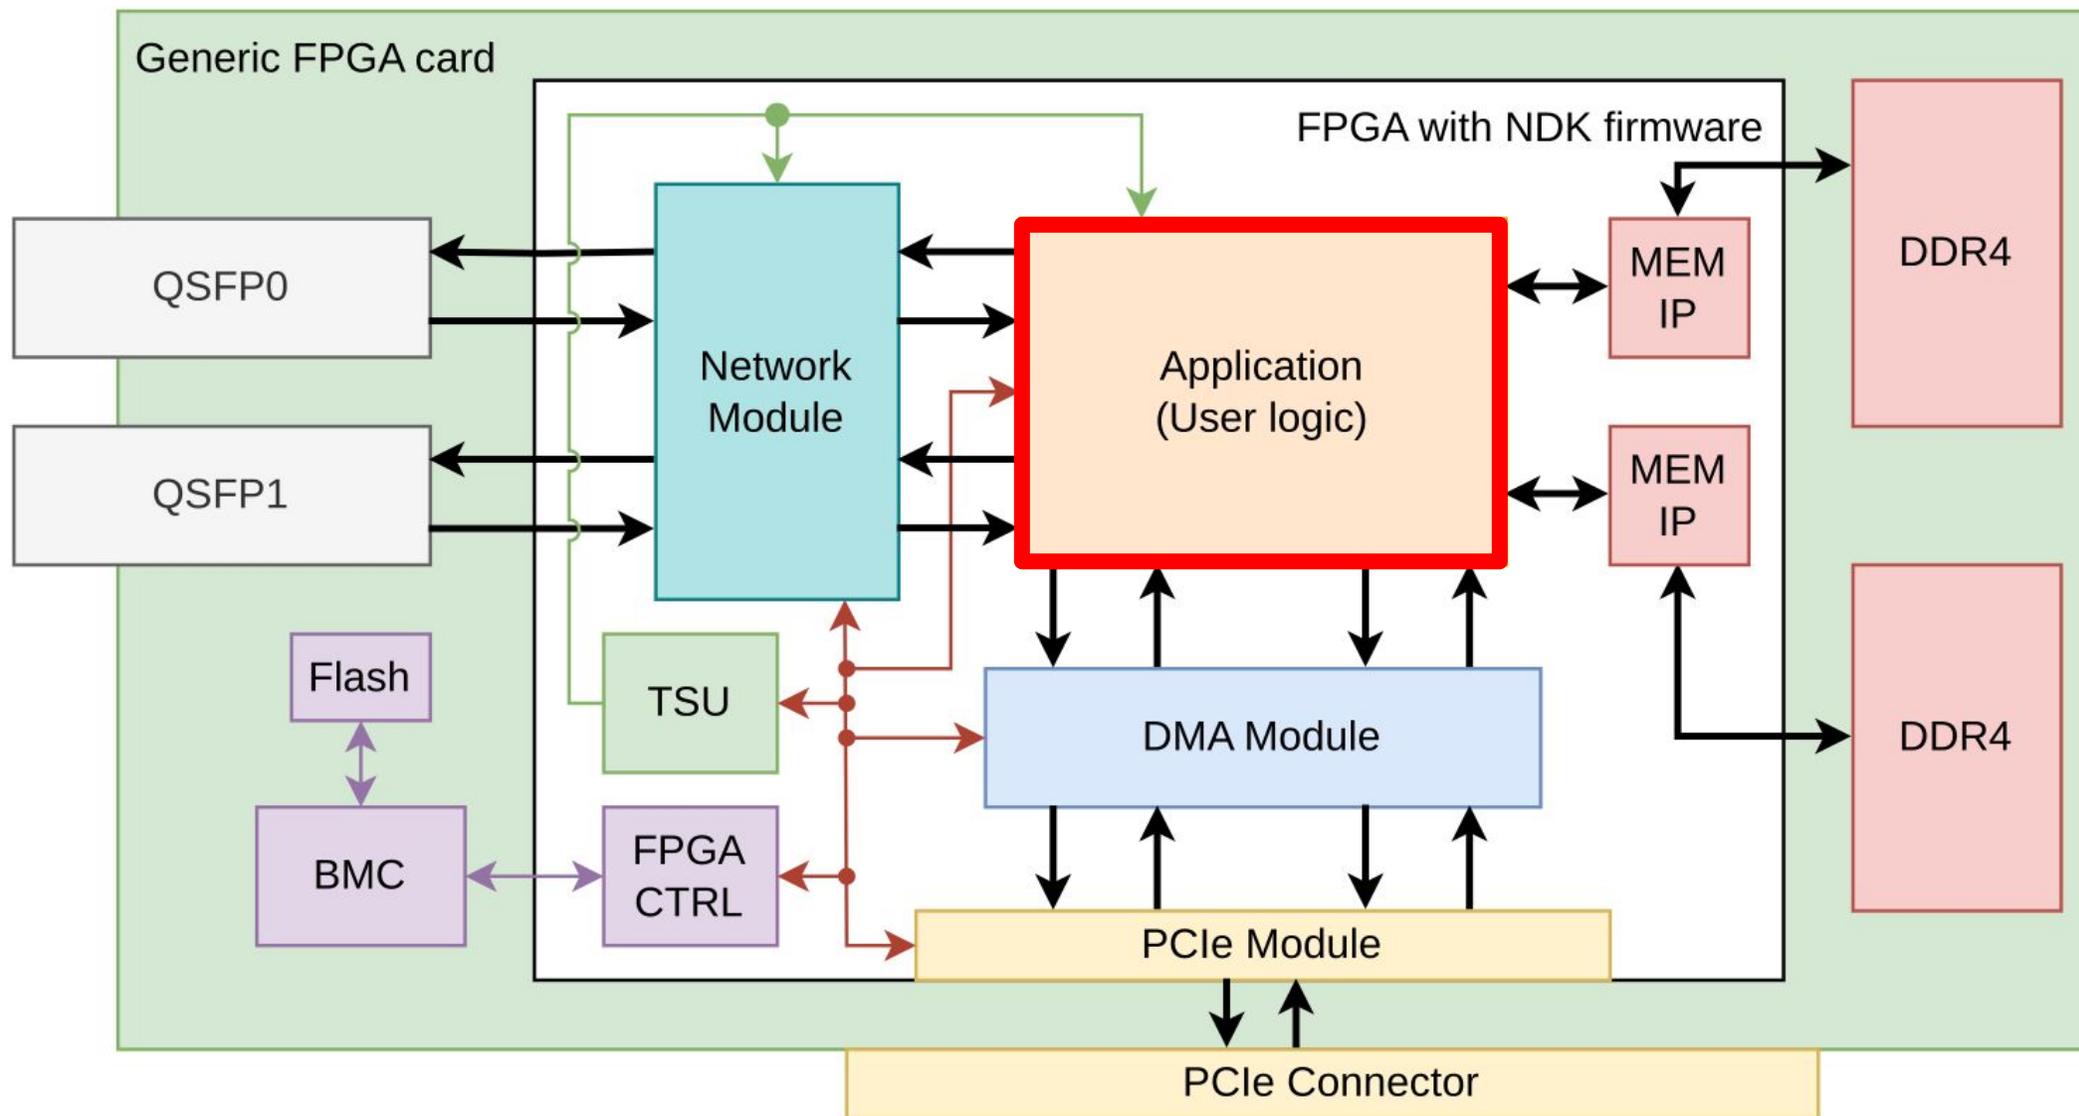


- Streaming data bus (Source -> Sink)
- Standardized relatives: AXI4 Stream, Avalon Streaming, ...
- Able to transfer multiple frames in a word
  - scales throughput by increasing width, not frequency => higher resource consumption
  - 400 Gbps  $\approx$  2048b @ 200 MHz
- Highly configurable data alignment
  - tradeoff between throughput and resource consumption

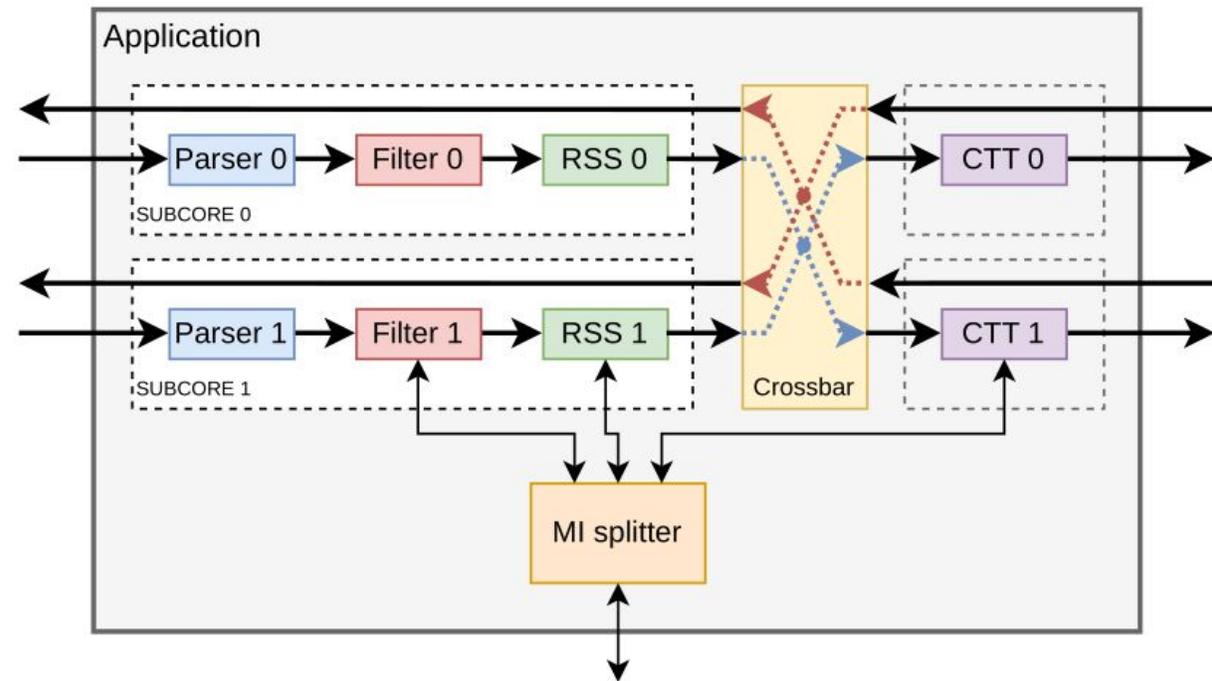


- Complementary bus for the MFB
- Multiple values per clock cycle
- Values have constant width
- Used to transfer packet metadata
  - packet length,
  - timestamp,
  - extracted packet headers (MAC, IP, ...), ...





- Independent pipelines per port
  - Parse packet headers
  - Filter out unwanted traffic
  - Receive Side Scaling (RSS) - distribution to DMA Queues
  - Connection Tracking Table (CTT)



- CESNET
  - Monitoring, IDS, ... acceleration
- [DynaNIC Semiconductors](#)
  - High-speed NIC applications
    - Network Security, Cloud Storage, Machine Learning, ...
- [Magmio](#)
  - High-Frequency Trading
- [BUT FEEC](#) (Brno University of Technology - Faculty of Electrical Engineering and Communication)
  - Post-quantum encryption



- NDK's highlights
  - Hardware acceleration of high-speed network applications
    - Support for the highest available speed standards
    - Multi-segment buses
  - Wide range of supported FPGA-based NICs
    - FPGAs from Altera and AMD
  - Flexible and easy-to-use build system
  - Standard SW APIs (DPDK, XDP)
  - Robust verification system



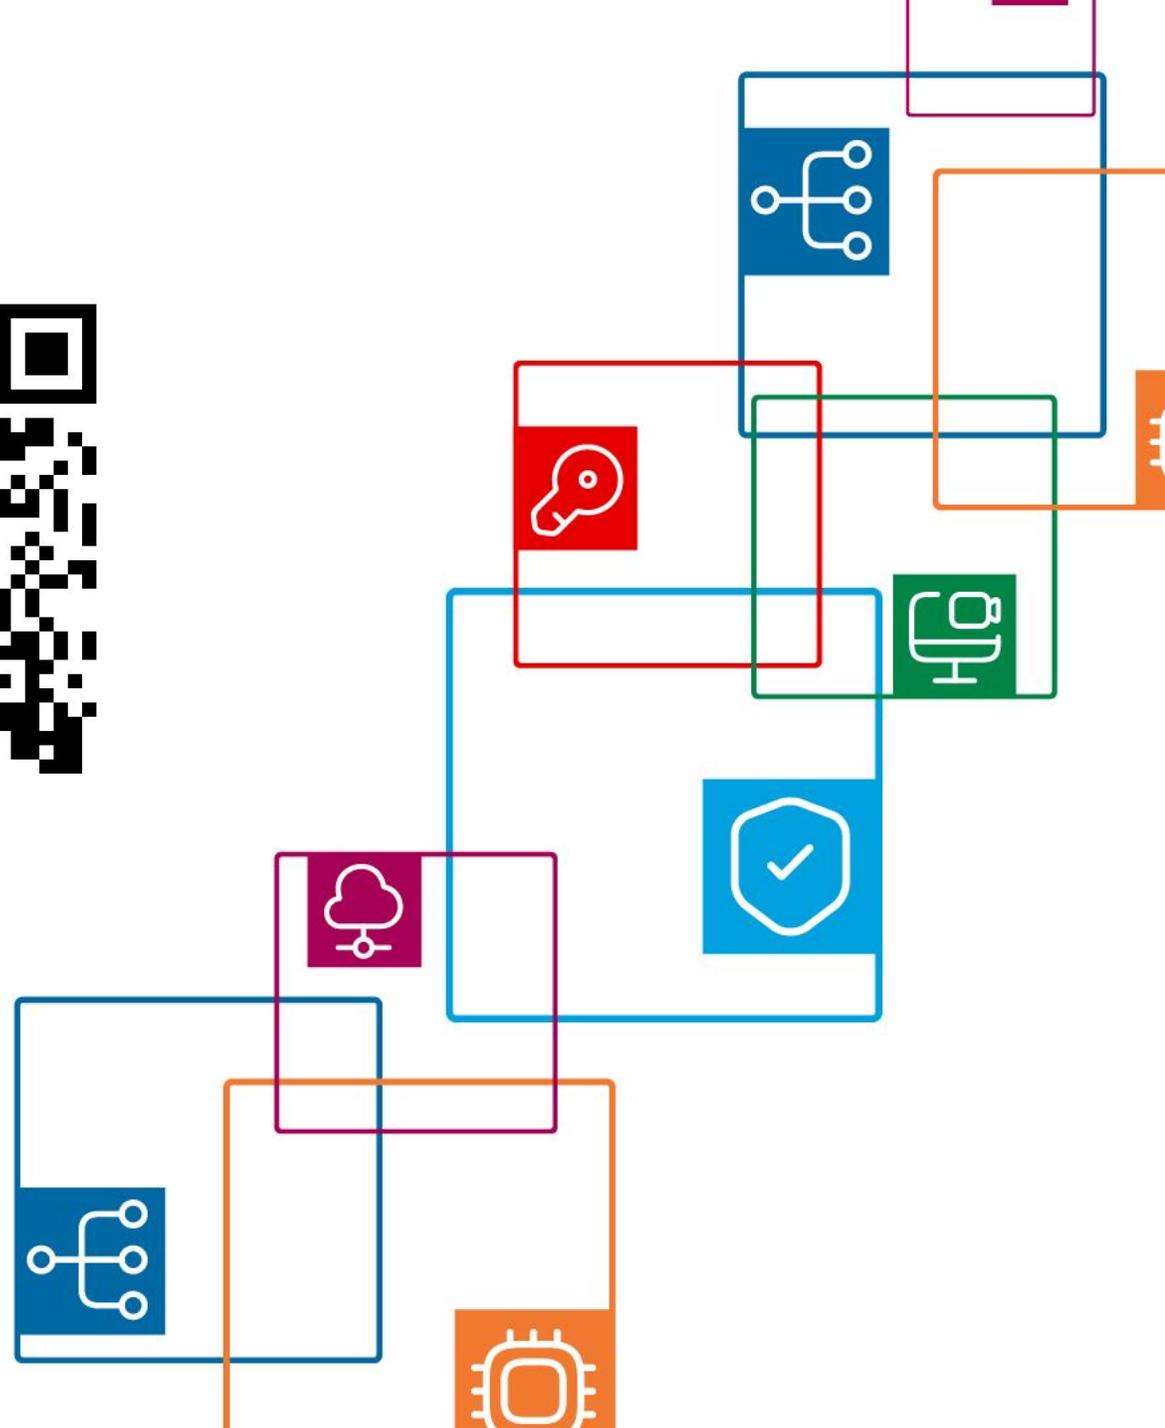
cesnet  
"...."



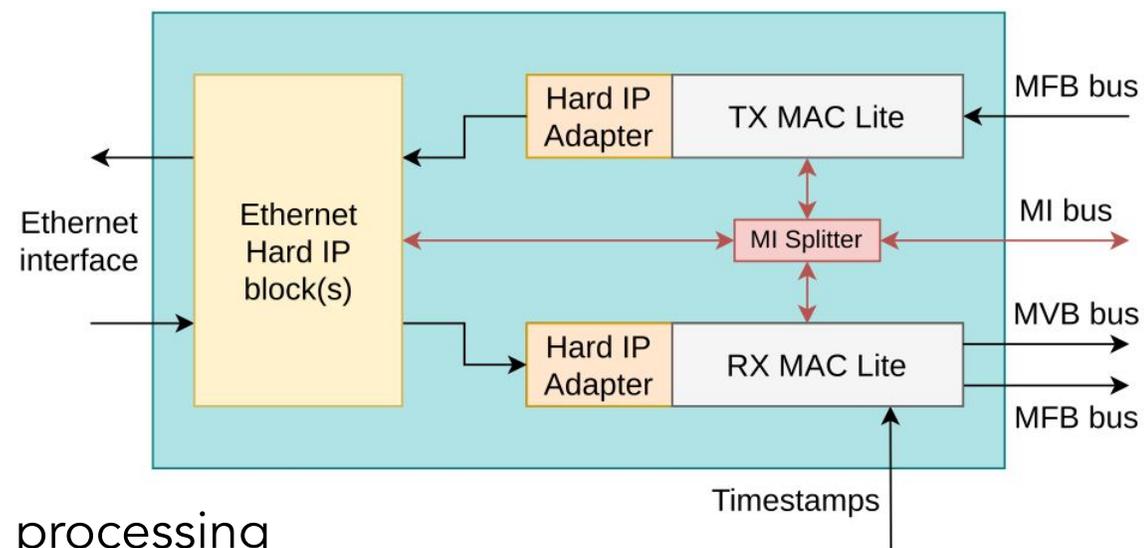
Thank you for your attention!

Checkout our [GitHub](#)

kondys@cesnet.cz



- Processes packets from/to the network
- Consists of
  - A hard IP core for Ethernet
    - IPs from different FPGA vendors
  - A bus adapter
    - converts IP's bus to a common bus
  - MAC Lite
    - Additional logic for Ethernet frame processing
- Flexible and highly configurable
  - Multiple speed standards (10 - 400 GE)
  - Multiple channel configurations



- Enables communication over the PCIe interface
- Flexible and highly configurable
  - Gen3, 4, 5
  - Bifurcation modes (1x16, 2x8)
- Consists of
  - A hard IP core for PCIe
    - IPs from different vendors
  - A bus adapter
    - converts IP's bus to a common bus
  - ROM with Device Tree of the FW
  - MI master endpoint

