

Jan Kubalek, Supervisor: Tomas Martinek, Faculty of Information Technology of the Brno University of Technology

Motivation

With the ongoing development of Internet and computer networks in general the amount of data processed and transferred between computers constantly increases. Backbone network nodes need to provide support for connections with speed reaching up to hundreds of gigabits per second. This support includes monitoring and analysis of Big data which requires hardware acceleration for processing and high-speed hardware-to-software communication. There are currently multiple solutions that allow network access with up to 100 Gbps and a few with 200 Gbps. This thesis deals on the development of Host-to-Card (H2C) part of a firmware DMA Module which supports various transfer speeds between 100 and 400 Gbps with the possibility of further scaling. The module is designed to be highly configurable and platform-independent.

Technologies and capabilities

- Mainly intended for **FPGA-based Smart NICs**
- Can possibly be used on different peripherals such as graphic cards or specialised data processing hardware
- High throughput for PCIe Gen3 and Gen4 x16 with future aims for PCIe Gen5
- Different FPGA platforms including **Xilinx Ultrascale+** and **Intel Stratix10**
- **Low-overhead** packet DMA transfer system compatible with **DPDK**
- Generic number of parallel DMA queues for ideal CPU core load distribution and support of virtualization

Architecture

- Architecture supports connection of a single DMA Module to various combinations of PCIe interfaces at once
- Each PCIe interface's transaction layer is provided by one PCIe Transaction Controller unit (PTC)
- This unit can also divide the PCIe data stream to multiple streams with 100 Gbps throughput based on the connected PCIe generation and number of lanes
- The DMA Module itself is divided to DMA Endpoints, each processing data with 100 Gbps throughput
- DMA Endpoint is directly connected to one of the PTC units
- On the other side, data streams from all DMA Endpoints are merged to one while sustaining the maximum data flow speed

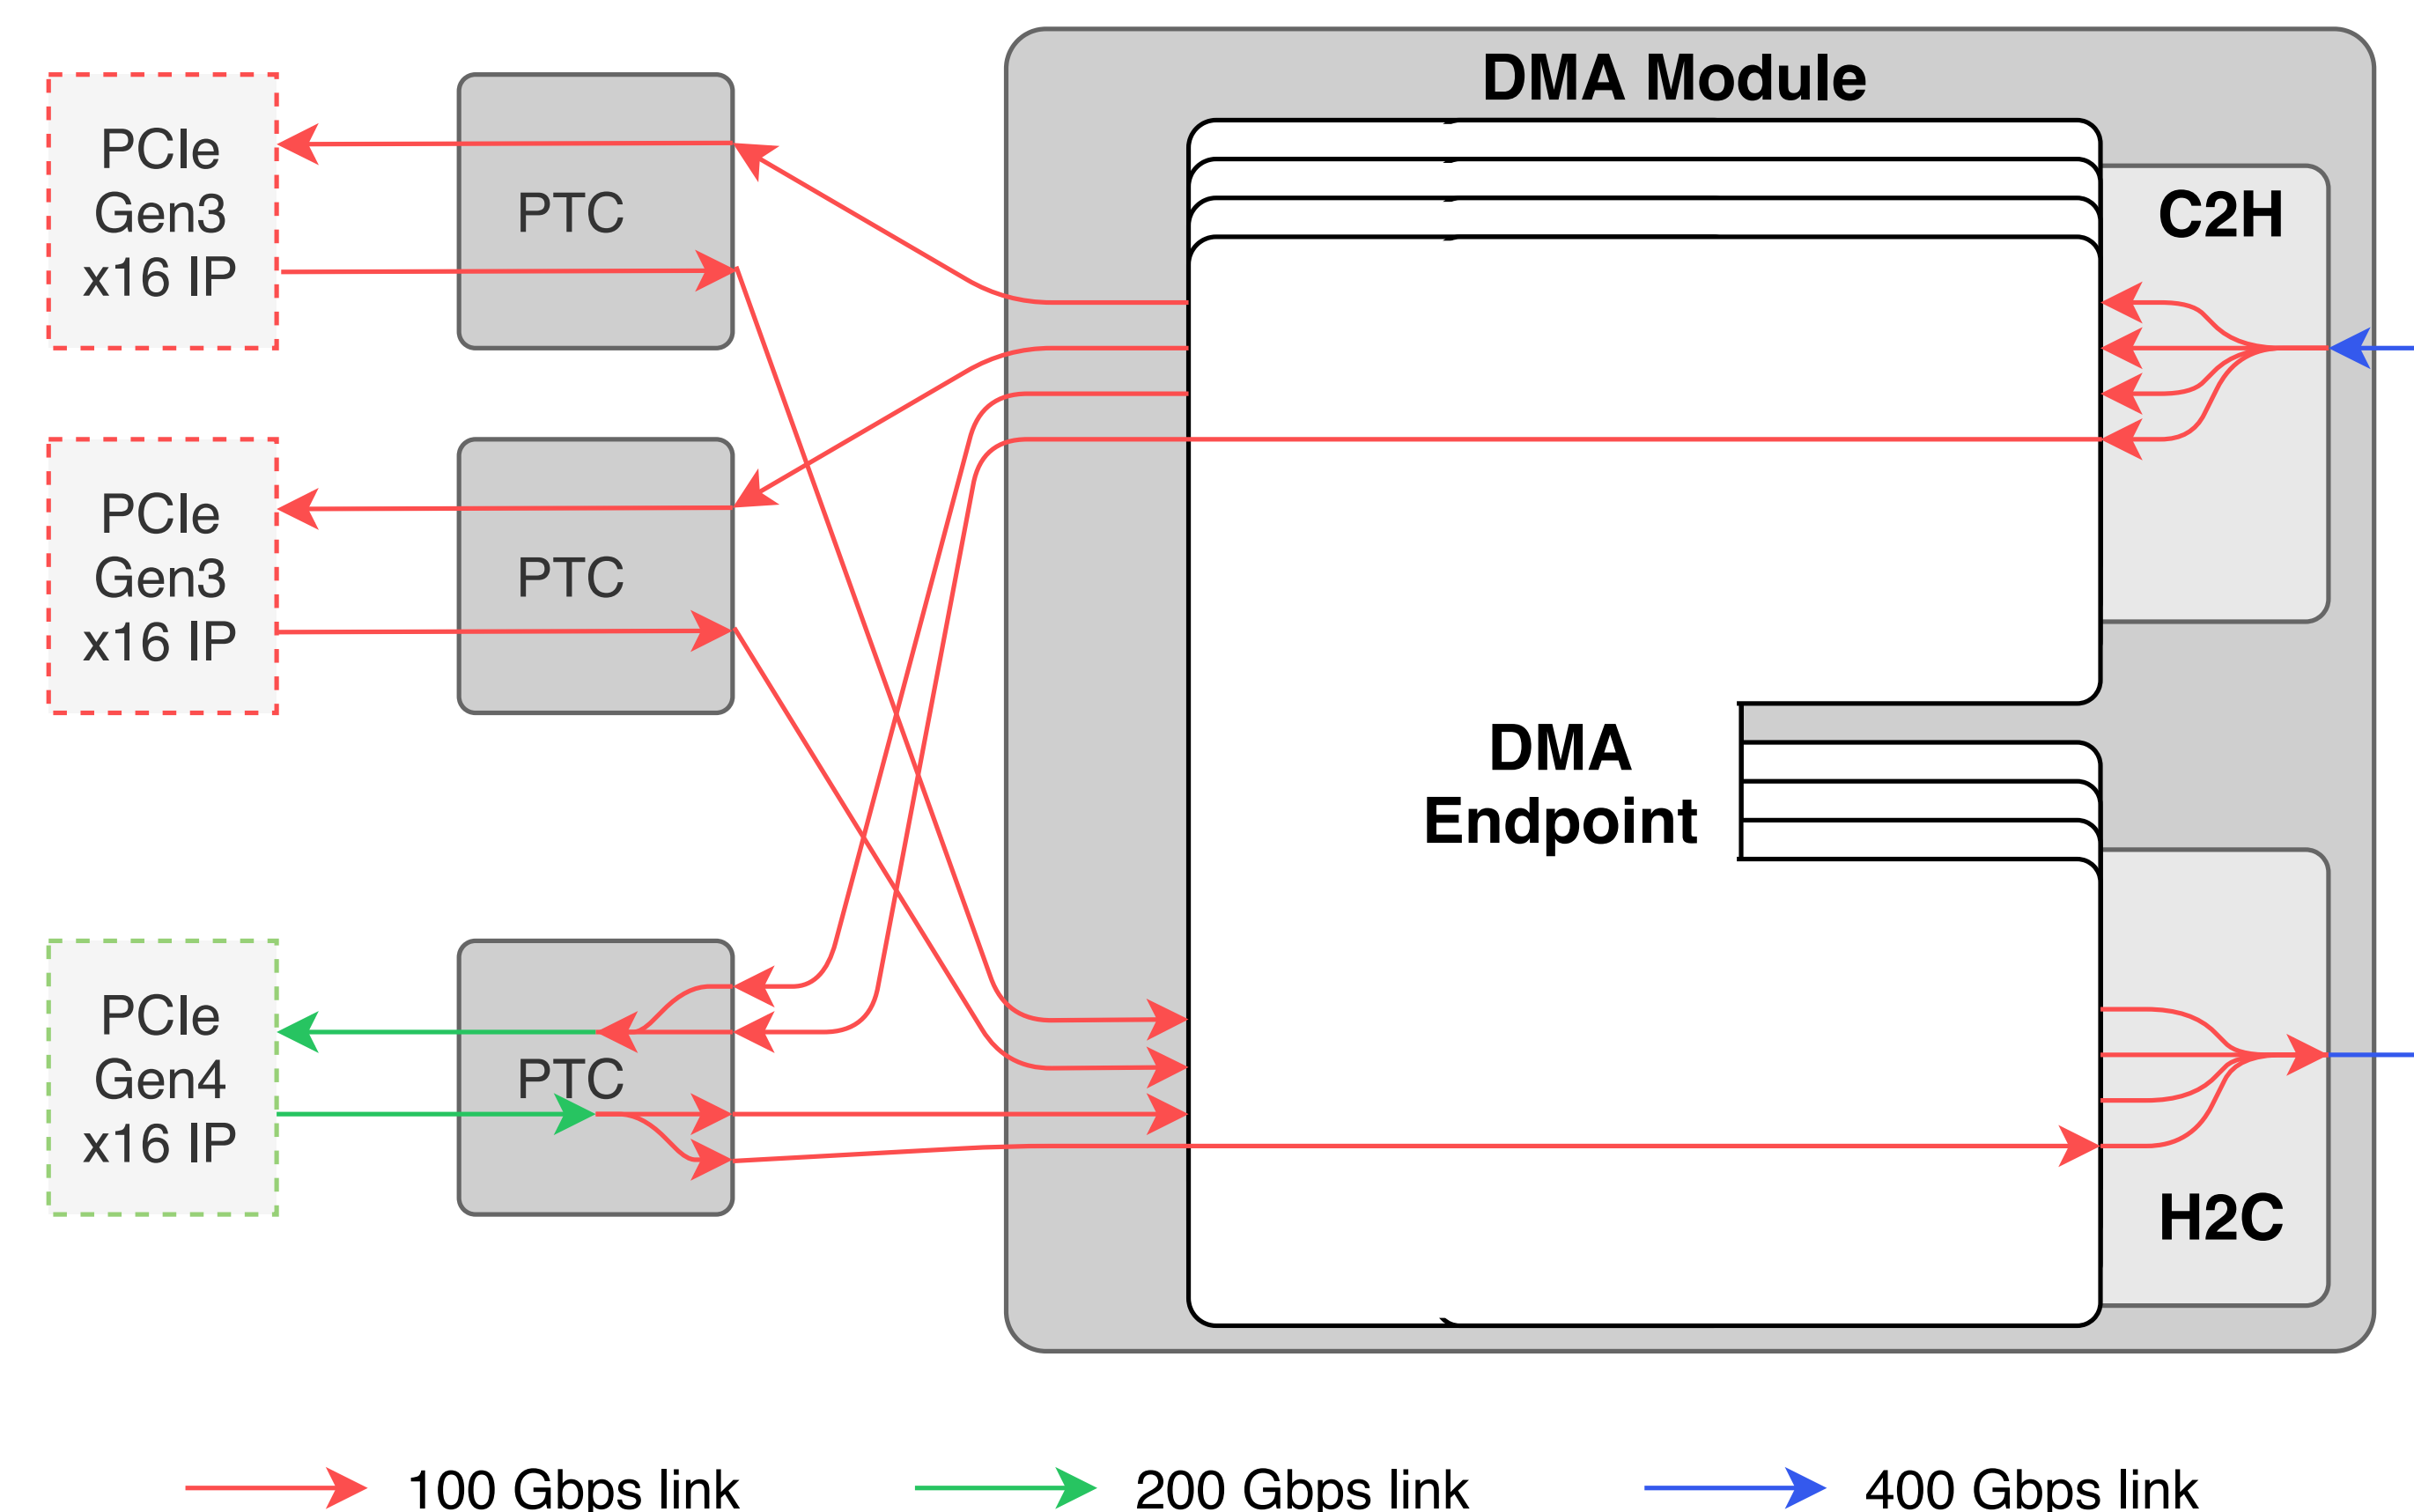


Figure 1: Example of DMA Module connection with 2x PCIe Gen3 x16 plus 1x PCIe Gen4 x16 and total throughput of 400 Gbps

Testing and measurement

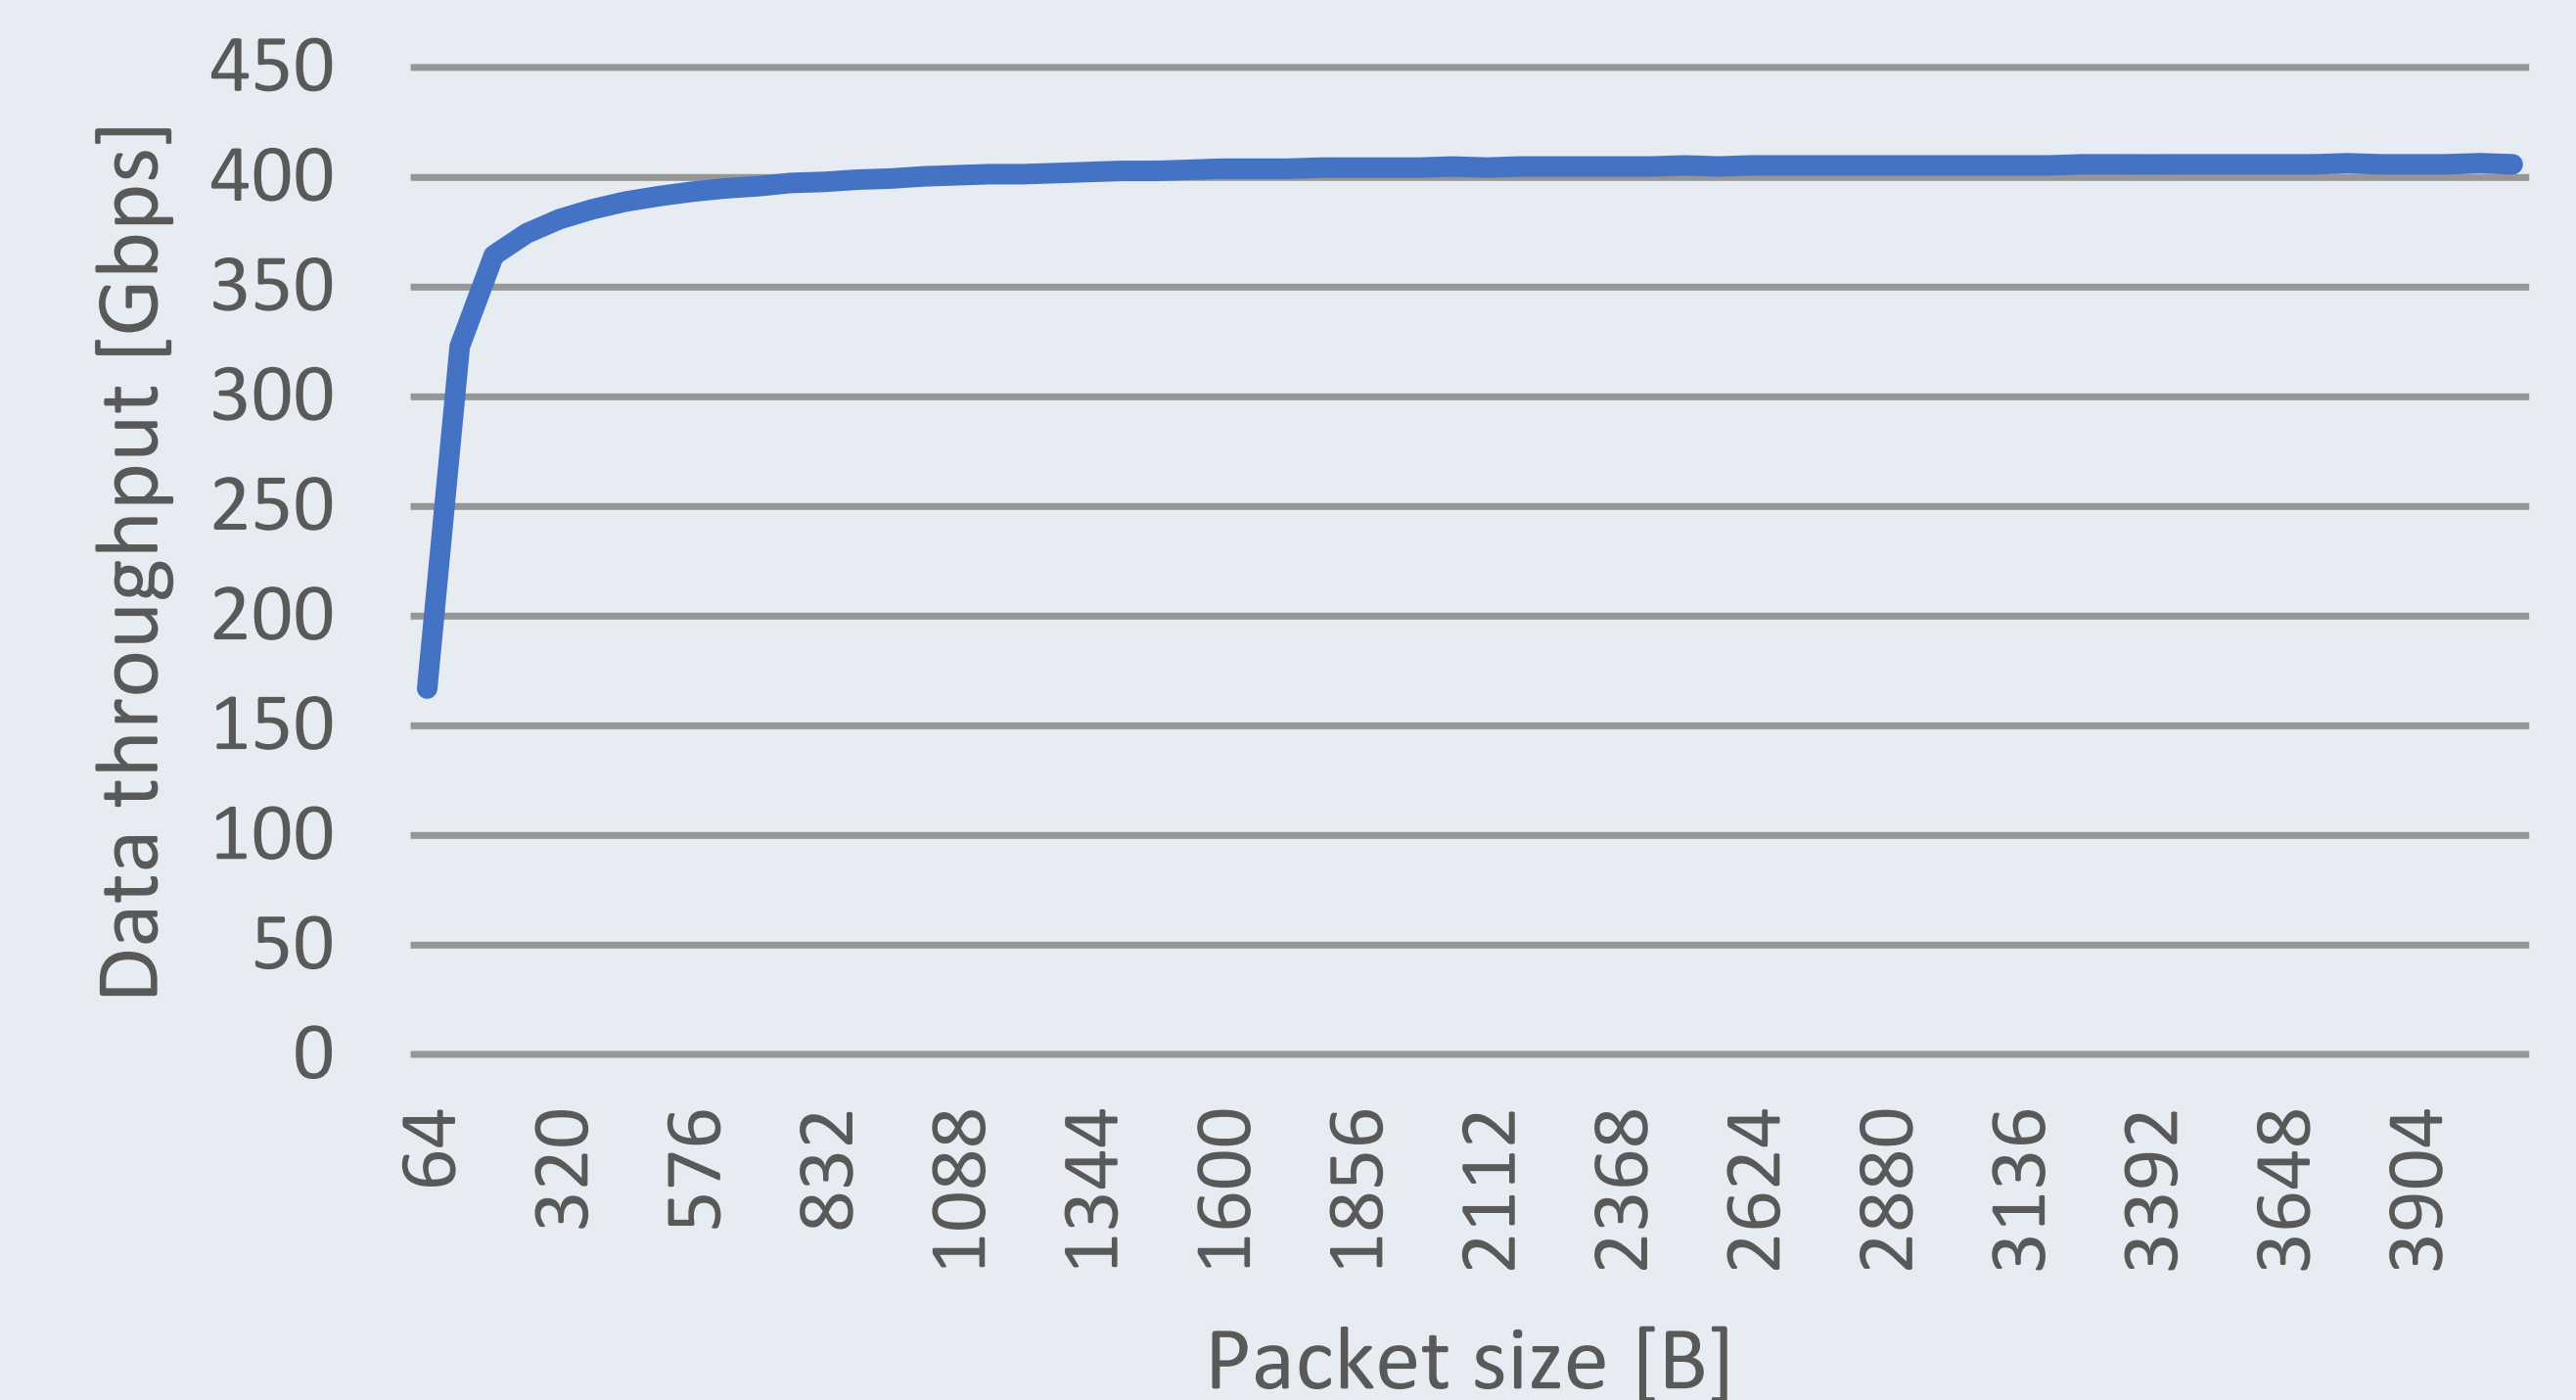


Figure 2: User data throughput in H2C direction for packet sizes 64 to 4096 B using 4x PCIe Gen4 x8 on FPGA Stratix10

- Tested on FPGA Xilinx Ultrascale+ with **1x PCIe Gen3 x16** (100 Gbps) and Intel Stratix10 with **4x PCIe Gen4 x8** (400 Gbps) and **2x PCIe Gen4 x16** (400 Gbps)
- On Stratix10 supports up to **512 DMA queues**
- FPGA resource consumption increase from 8 queues to 512 queues only **39 %**

Future work

- Expanding support for more FPGA platforms and PCIe IPs including **PCIe Gen5**
- Testing of different PCIe combinations and speeds
- Research paper describing the whole DMA Module currently being prepared