# 100G/400G Ethernet on Agilex 7

Alberto Perro - PCIe400 Development Meeting - 2/11/2023

# Overview

- F-tile Ethernet Hard IP
- 100G Avalon Stream UDP
- 100G Segmented Stream UDP
- 400G and future developments

# F-Tile Ethernet IP

The **F-Tile Ethernet** Hard IP is the ethernet core available on Intel **Agilex 7 I and M series**.

The IP supports 10/25/40/50/**100**/200/**400** GbE and already provides 1588 Precision Time Protocol, Auto Negotiation, Link Training, and different Forward Error Correction standards (FEC and RS-FEC).

The IP supports **Avalon Stream** client interface up to 100G, then it is mandatory to switch to the **Segmented Stream** interface.

# Quartus Support

Support for the AGX7I-series has been introduced in v22.3 Pro. The IP and the example designs are in **early stage**, with many functionalities lacking or not working as expected.

The version used for the evaluation is the **22.4 Pro**, which seemed to be the most stable. Early support also means that Place and Route takes a **great amount of time**, since it is not optimized.

Current evaluation of version 23.2 Pro showed great reduction in compilation time, but there are still some functionality bugs.

# 100G UDP Avalon Stream: Firmware

A 100G UDP Transmitter Avalon Client has been designed to test the performance of the hard IP. The whole chain runs at 402.83203125 MHz.

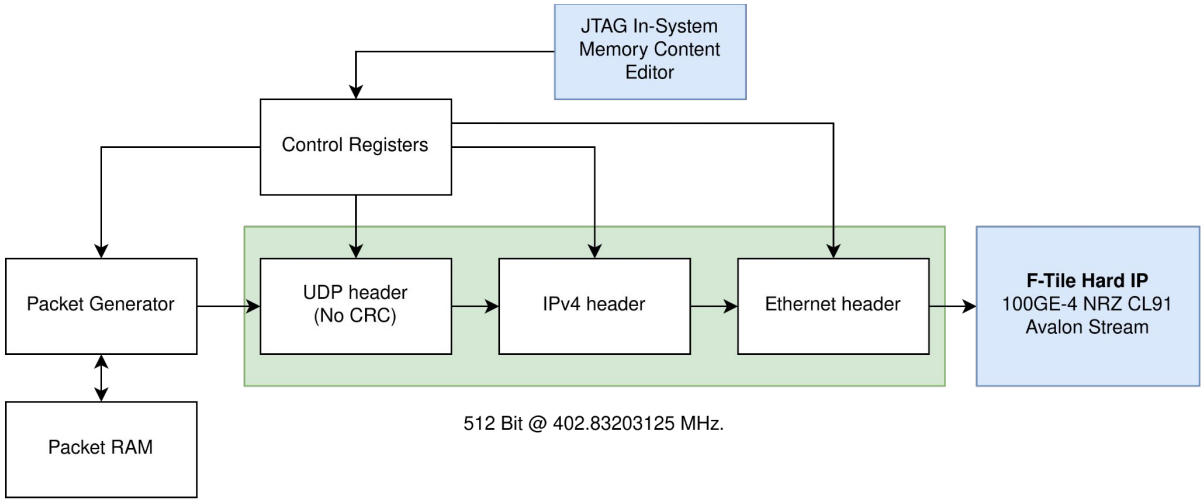

# 100G UDP Avalon Stream: Setup

The testbench uses an Intel Agilex 7 FPGA I-Series Development Board which hosts two QSFP-DD cages.

One of the cages (QSFP-DD_0) is connected via a QSFP **Direct Attach Cable** to a NVidia ConnectX-6 200GbE **Network Interface Card**.

The packet generator tags every packet to be able to measure packet loss. In a **back-to-back** configuration, **no packet loss** has been measured and **full throughput** has been achieved (1.5 Mpps at 8096 message size).

# 100G UDP Avalon Stream: Learnings

- The IP **can’t get CDR lock** with CL134 and ETC RS-FECs.
- At 100G the ready signal of the AvST **oscillates continuously**. This is by design, but raises some questions on power consumption.
- Input/Output Registers are **mandatory** to reach timing closure (< 2.5 ns clock period).
- Any CRC calculation gets tricky with a 512 bits bus at 402 MHz, so it requires special pipelining.

# 100G UDP Segmented Stream

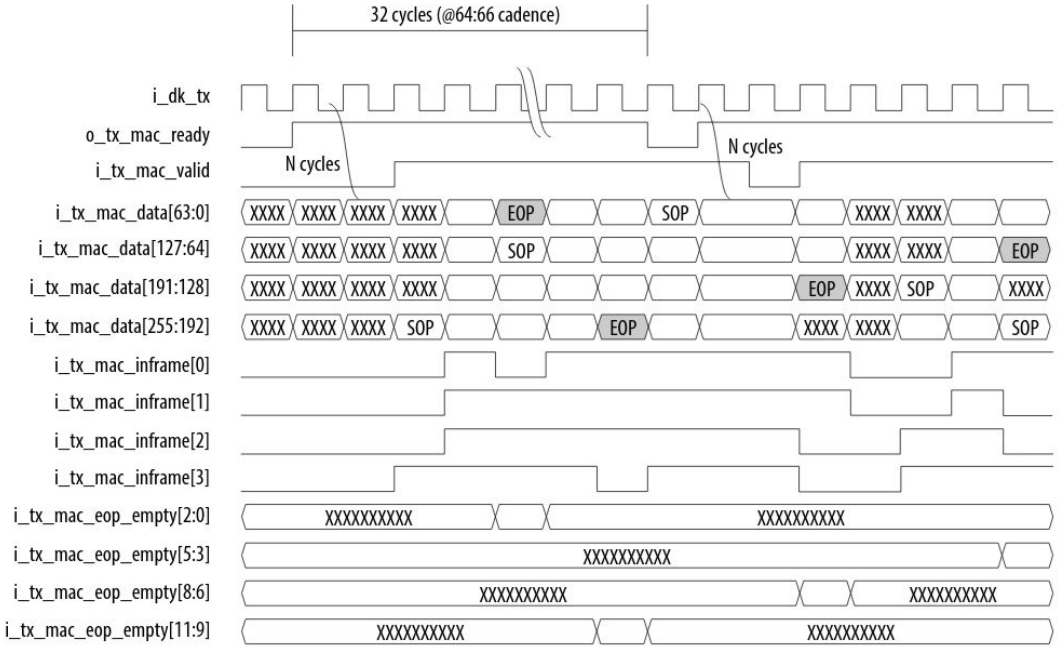


100G Segmented Interface

To go above 100G, a Segmented Stream client interface is required. This interface packs data **more efficiently** by reducing the empty bits between packets. Moreover, **bus size is reduced** as well, which helps with routing and timing.

# 100G UDP Segmented Stream

The segmented interface introduces new challenges in the design of the cores. A first test is ongoing on 100G to test the components behaviour.

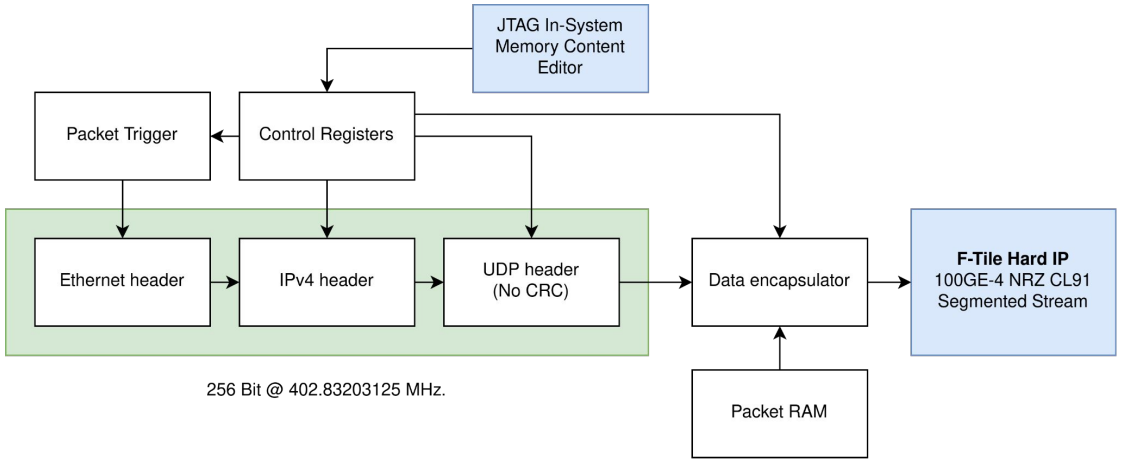

# 100G UDP Segmented Stream

The new design takes in account the necessity to add **new protocol block**s (such as RoCE).

Computation time for CRCs can be hidden in the latency of the header adders.

At 100G (256 bit bus) is **not** possible to have two full packets in the same word (the ethernet header is 20 bytes alone), so the header adders are considering **only one packet at a time**.

# 400G and future improvements

- A working avalon stream 100G UDP transmitter core has been **designed and tested**, the segmented version is under test.
- To scale to 400G (1024 bit bus) the header logic **must be updated** to handle multiple full packets in a single word.
- CRC computation will become even **more critical** at that throughput.
- Other protocols of interest have to be designed and tested (e.g. RoCE).
- The receiver chain has to be designed and tested.

# Thanks for the attention

# References

Intel F-Tile Documentation

https://www.intel.com/content/www/us/en/docs/programmable/683023/22-4/overview-16832.html

100G UDP Testbench Core

https://gitlab.cern.ch/alperro/simple_udp_vhdl/