Line rate 100 Gbps Ethernet Forwarding over PCI Express with Virtex-7

CESNET and INVEA-TECH have reached another milestone in their joint research of FPGA-based high-speed network data plane processing adapters. Until now, line rate forwarding of 100 Gbps Ethernet traffic in general CPU has been a huge challenge, if not impossible at all. The major obstacle is no longer the raw CPU power - the performance of modern multi-core CPUs is gigantic. The problem lies in the PCI Express interface. Its fastest variant uses 16 full duplex lines, each running at 8 GT/s. Theoretical throughput is therefore 128 Gbps. However, due to 128/130 coding and other protocol overhead, the real throughput is somewhat lower, close to the speed of the fastest 100 Gbps Ethernet standard available today.

CESNET and INVEA-TECH have been able to implement full duplex 100 Gbps Ethernet line rate packet receive and transmit with their FPGA-based hardware-accelerated 100GE network adapter COMBO-100G [1] and Intel Xeon CPU. Xilinx Virtex-7 FPGA on the card performed 100 Gbps Ethernet packet receive and DMA transfers to the host PC RAM RX buffers, together with the opposite direction: DMA transfers of packet data from the RAM TX buffers and 100 GE packet transmit. At the same time, the Intel Xeon CPU was forwarding the packets from RX to TX RAM buffers.

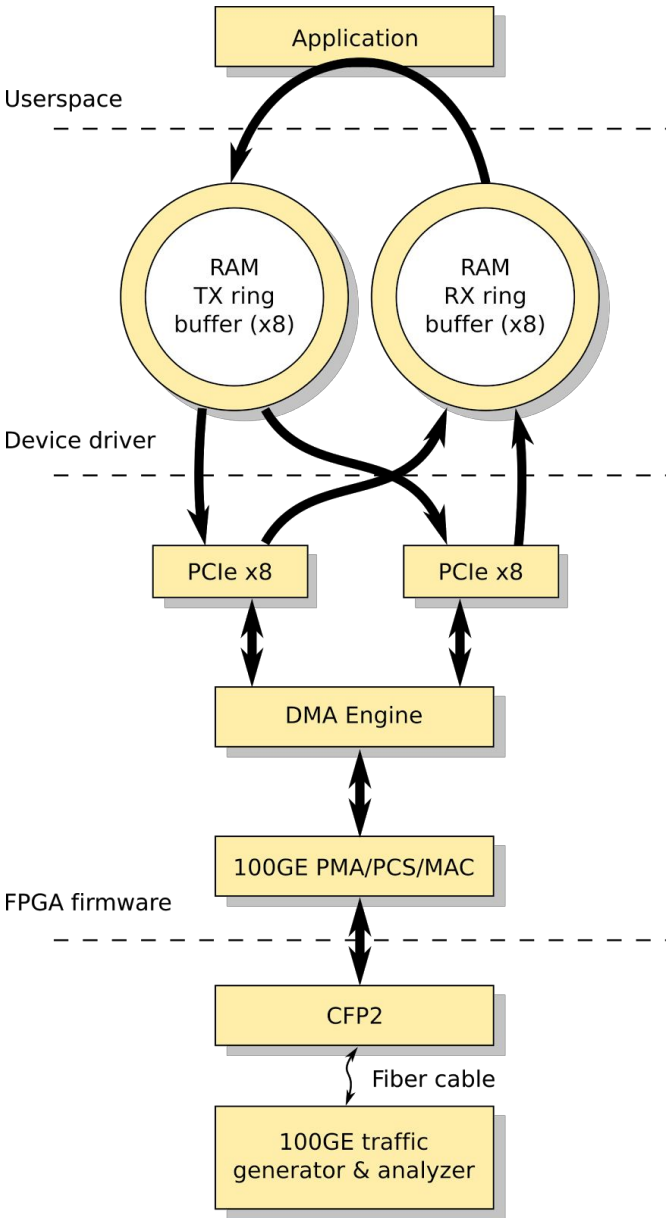

Multiple advanced techniques were employed to achieve this level of performance. Firstly, no current FPGA supports 16-lane PCI Express gen3 interface. Two 8-lane interfaces were used instead, using the so-called PCIe slot bifurcation [2]. The length and ordering of PCI Express transactions has been fine tuned using PCIe protocol analyzer (hand-optimized, actually) to achieve the optimal performance with Intel Xeon CPUs. Additional PCIe transaction buffers were added in the FPGA to extend standard Xilinx PCIe core capacity and compensate

for long latency of PCIe reads as well as limited transaction tag space. Eight independent buffers for RX and eight for TX were allocated in RAM, so that multiple CPU cores could work in parallel without the need for inter-core communication.

This achievement enables INVEA-TECH customers to create various high-speed packet processing systems including network testing and monitoring tools, security applications, high-performance computing nodes, lawful interception systems and other network appliances [3]. Furthermore, INVEA-TECH offers NetCOPE development framework that includes the above mentioned features for NFV applications and also offers selected FPGA cores to offload time-critical operations in order to further increase processing speed and reduce latency of target applications.

References

[1] http://forums.xilinx.com/t5/Xcell-Daily-Blog/PCIe-Gen3-essential-for-high-speed-FPGA-based-Ethernet-adapter/ba-p/450828

[2] http://forums.xilinx.com/t5/Xcell-Daily-Blog/Need-to-get-100G-Ethernet-data-stream-into-a-host-Intel-CPU-PCIe/ba-p/508823

[3] http://forums.xilinx.com/t5/Xcell-Daily-Blog/INVEA-TECH-s-FPGA-based-HANIC-100G-card-captures-Ethernet/ba-p/547787