Ralf Kundel, Leonhard Nobach, Hans-Joerg Kolbe, Tobias Meuser, Ralf Steinmetz. *FPGA-assisted Massive Packet Queueing and Traffic Shaping at the Network Edge.* To appear in the Proceedings of IEEE International Symposium on Field-Programmable Custom Computing Machines (FCCM), IEEE, 2022.

FPGA-assisted Massive Packet Queueing and Traffic Shaping at the Network Edge

Ralf Kundel*, Leonhard Nobach§, Hans-Joerg Kolbe§, Tobias Meuser*, Ralf Steinmetz*
*Multimedia Communications Lab, Technische University of Darmstadt, Germany
{*ralf.kundel, tobias.meuser, ralf.steinmetz*}*@kom.tu-darmstadt.de*
§ Deutsche Telekom Technik GmbH, Germany
{*leonhard.nobach, hans-joerg.kolbe*}*@telekom.de*

***Abstract*—Large-scale packet queueing and scheduling is the basis for today's Quality of Service (QoS) in computer access networks, especially to achieve guaranteed high throughput and low latency. While highly performant fixed-function ASICs offer sufficient functionality for most data center use-cases, as of today, they cannot support all functionality required for access networks, e.g., QoS-aware packet queueing.**

In this poster, we present an FPGA-based architecture for network packet queueing optimized for residential and mobile Internet access networks.

***Index Terms*—FPGA, Network Function Offloading, QoS, Queueing, Traffic Shaping, AQM**

I. Motivation

In our previous work, we investigated the capabilities of P4-programmable hardware for Internet service creation [1], [2]. One main finding was the lack of traffic shaping support in programmable switches. In mobile and residential Internet access creation, it is required to separate customers from each other. Therefore, for each customer (commonly up to $200,000$), a queue, sufficient packet storing memory, and advanced scheduling logic is required. One promising alternative to fixed-function ASICs are FPGAs, which have already proven their viability for network function offloading [3].

II. Proposed Design

Building on these requirements, we propose an FPGA design, realizing massive packet queueing. Concretely, we utilized a Xilinx Ultrascale+ FPGA with external DDR4 memory for packet buffering. A pipeline, described in Verilog, processes every ingressing packet: 1) Packets are classified, and a queue ID is assigned. 2) The packet is stored in the external memory. For this, a shared memory concept was chosen, *i.e.*, a physical address is allocated dynamically for every packet, and the memory is not segmented by the logical queues. 3) The packet address in the external memory and its length are stored in an FPGA-internal data structure, as shown in Figure 1. For this, a linked list data structure is realized, consisting of two memories: the first memory, *queues_mem*, holds a pointer on the first and last packet in each queue. The second memory, *buckets_mem*, consists of the list entries pointing at each other. 4) A scheduler pops packets from this queueing data structure, considering the rate limits and QoS rules to be met. 5) Finally, a transmission module reads the packet from the external memory and sends the packet on an Ethernet port. For this, the stored packet length is required, as otherwise, the actual packet length can not be determined.

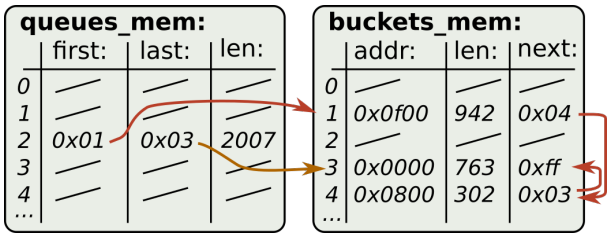

Fig. 1: FPGA-internal packet-queue data structure.

Our evaluation results prove this approach as a viable and flexible alternative to currently used ASIC-based solutions, providing the necessary performance for current and future Internet access networks. In experiments up to 100 $Gbit/s$ link speed, we observed zero unexpected packet loss, and achieved a low and deterministic latency if the configured rate limit is higher than the link speed. In the case of built up packet queues, we observed the expected scheduling behavior. Energy usage measurements have shown a base consumption of ~ 23 $Watt$, and a throughput dependant consumption of ~ 0.1 $\frac{Watt \cdot s}{Gbit}$, which is much lower compared to software-based solutions.

Acknowledgment

This work has been supported by Deutsche Telekom through the Dynamic Networks 8 project, and in parts by the German Research Foundation (DFG) as part of the project B1 and C3 within the Collaborative Research Center (CRC) 1053 - MAKI.

References

[1] R. Kundel, T. Meuser, T. Koppe, R. Hark, and R. Steinmetz, "User plane hardware acceleration in access networks: Experiences in offloading network functions in real 5g deployments," in *Proceedings of the 55th Hawaii International Conference on System Sciences*. Computer Society Press, 2022, p. 1–10.

[2] R. Kundel, L. Nobach, J. Blendin, W. Maas, A. Zimber, H.-J. Kolbe, G. Schyguda, V. Gurevich, R. Hark, B. Koldehofe, and R. Steinmetz, "Openbng: Central office network functions on programmable data plane hardware," *International Journal of Network Management*, vol. 31, no. 1, p. e2134, 2021, 25 pages.

[3] J. Naous, G. Gibb, S. Bolouki, and N. McKeown, "Netfpga: Reusable router architecture for experimental research," ser. PRESTO '08. New York, NY, USA: Association for Computing Machinery, 2008, p. 1–7.