Fraunhofer FOKUS Institute for Open Communication Systems

eXpress Data Path Extensions for High-Capacity 5G User Plane Functions

SIGCOMM 2023: 1st Workshop on eBPF and Kernel Extensions

Christian Scheich, Marius Corici, Hauke Buhr, Thomas Magedanz

Outline

- Motivation
- Background
- Design
  - XDP GTP-U Extensions
  - XDP RSS
- Evaluation
  - Downlink Throughput
  - Uplink RSS
- Conclusion

Motivation

- Emerging data-intensive use-cases like Virtual Reality and high-quality video streaming challenge the throughput capacity in mobile networks.

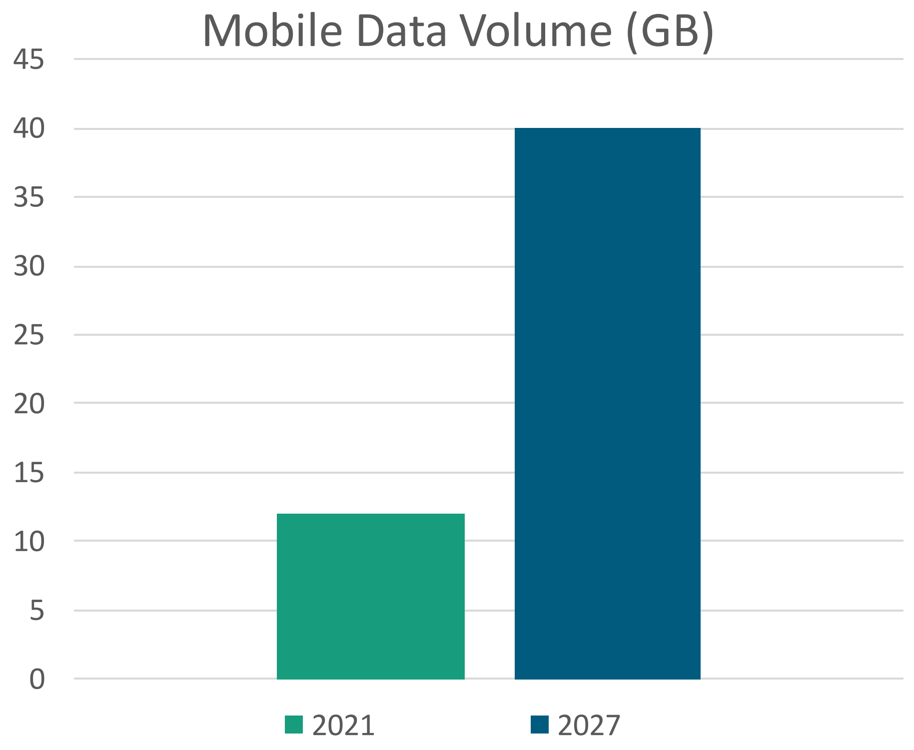

Ericsson Mobility Report June 2022

Fraunhofer FOKUS

Background: User Plane Function in 5G

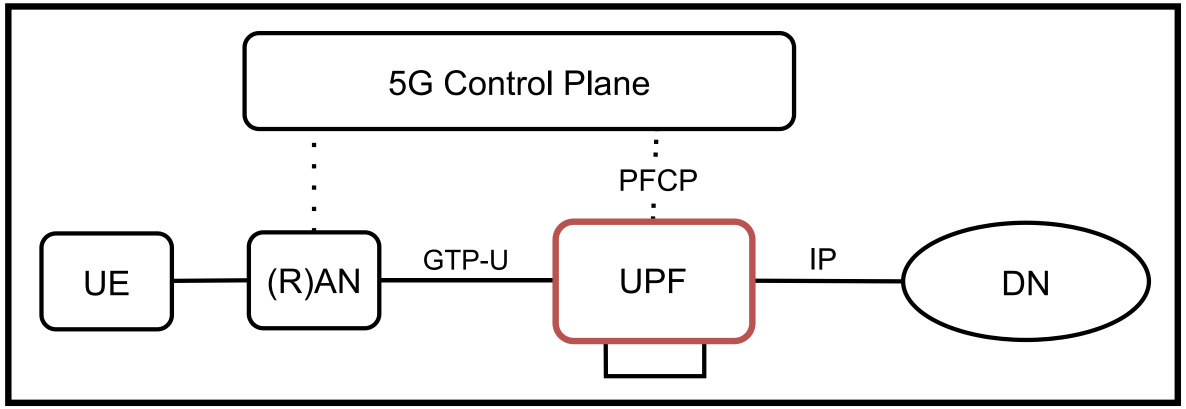

5G Architecture Example with highlighted UPF Layer (adapted from 3GPP TS 23.501 4.2.3-2)

- User Plane Function (UPF) connects Users in the RAN with the Destination Networks (DN)
- UPF is configured from the control plane via Packet Forwarding Control Protocol (PFCP)
- UPF can apply forward, buffer or apply QOS rules to the packets
- Traffic is forwarded in GTP-U Tunnels in the Data Plane

Background: Receive Side Scaling with GTP-U

- GTP header includes a unique Tunnel Endpoint Identifier (TEID)
- Network Interface Card distributes traffic with Receive Side Scaling to the available CPUs
- Sender can vary the UDP Source Port to enable RSS features, but not all Cells support this
- Receiver can be extended to load-balance with the TEID of GTP-U

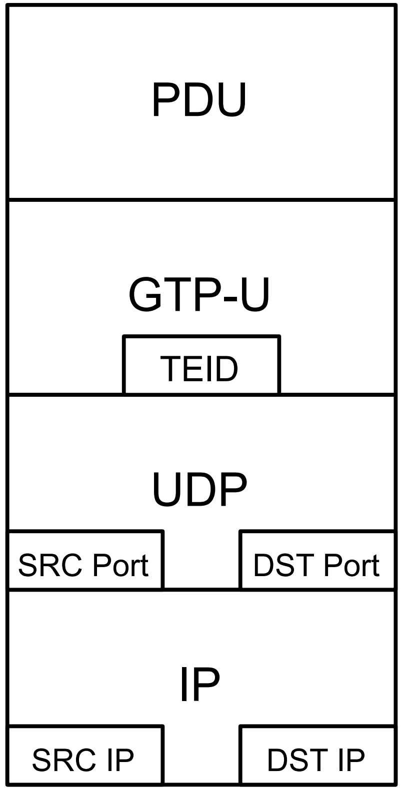

GTP-U Protocol Stack based on 3GPP TS 29.281 V17.3.0

Design: XDP GTP-U Extensions

- XDP GTP-U
  - Parses incoming packets for GTP-U header information
  - Applies actions based on the BPF-Map and the PFCP rules from user space
  - User Space UPF is fallback for buffering and extended QoS
  - Can pass to the network stack, redirect and reflect the packet
- XDP Redirect
  - Connect IP networks as router
- XDP TX:
  - Share one link on a bridged network

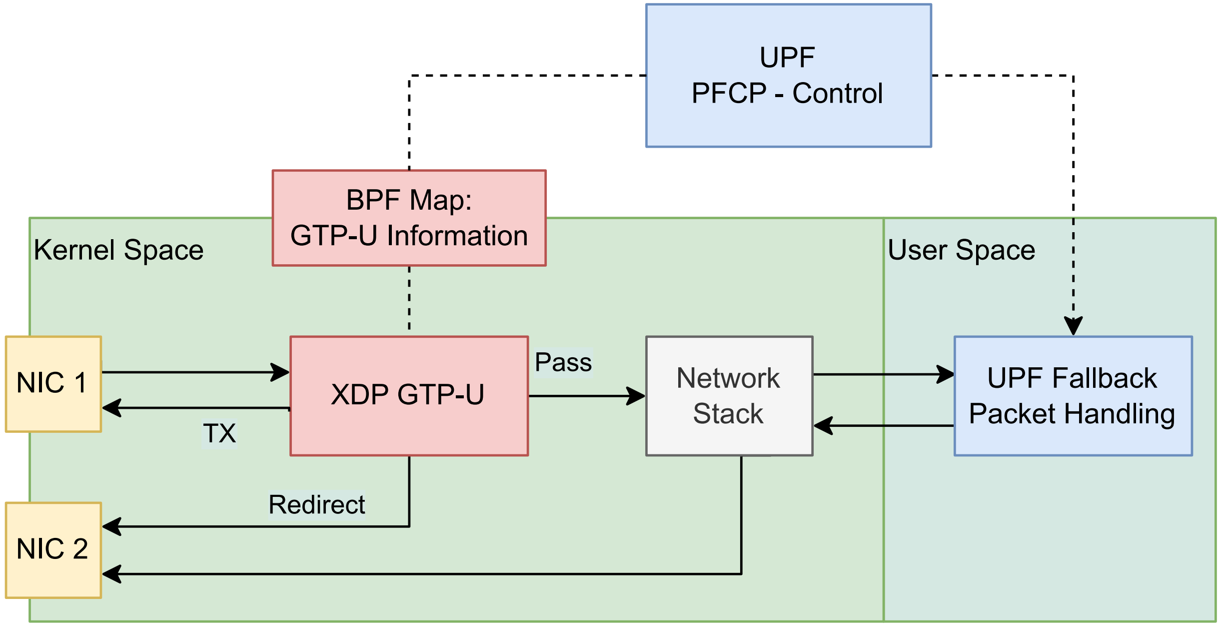

 09.10.23

Design: XDP Receive Side Scaling

- Dispatcher selects a CPU based on the GTP-U tunnel information
- XDP-GTP-U execution is balanced on multiple CPUs
- Excludes the CPU-Dispatcher from XDP-GTP-U handling

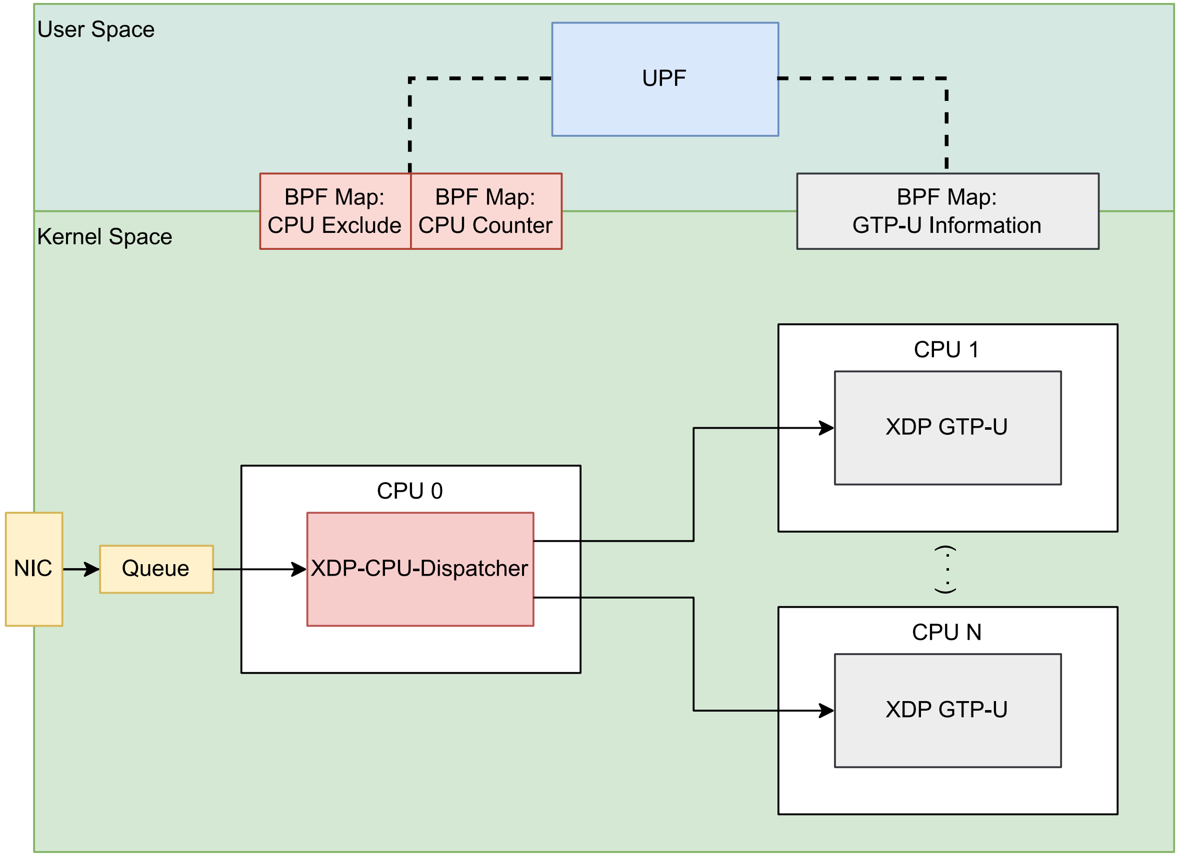

Fraunhofer FOKUS

Evaluation

- Measurement Setup for RFC2544 tests:
  - UPF is Device under test performing GTP-U encapsulation and decapsulation
  - Throughput Measurement Device:
    - sends traffic to the Device under test and counts the received packets
    - Generates traffic for 10.000 GTP-U streams

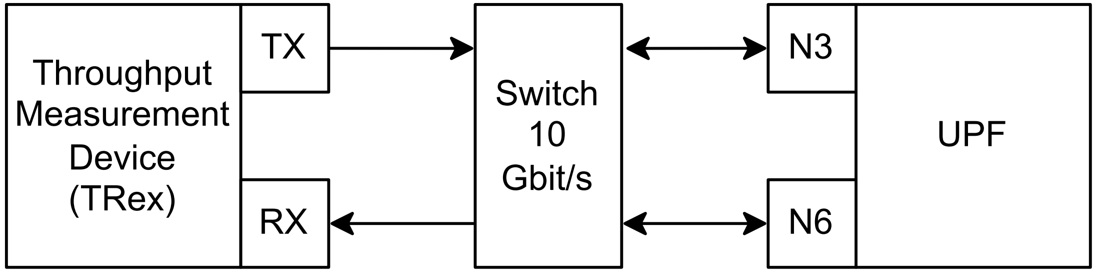

Throughput Comparison Downlink

- Nine-fold improvement of XDP-GTP with pass compared to user space
- Shortening the network stack with XDP-Redirect and XDP-TX improve further
- 15-fold throughput increase with XDP-Action TX in comparison to user space implementation

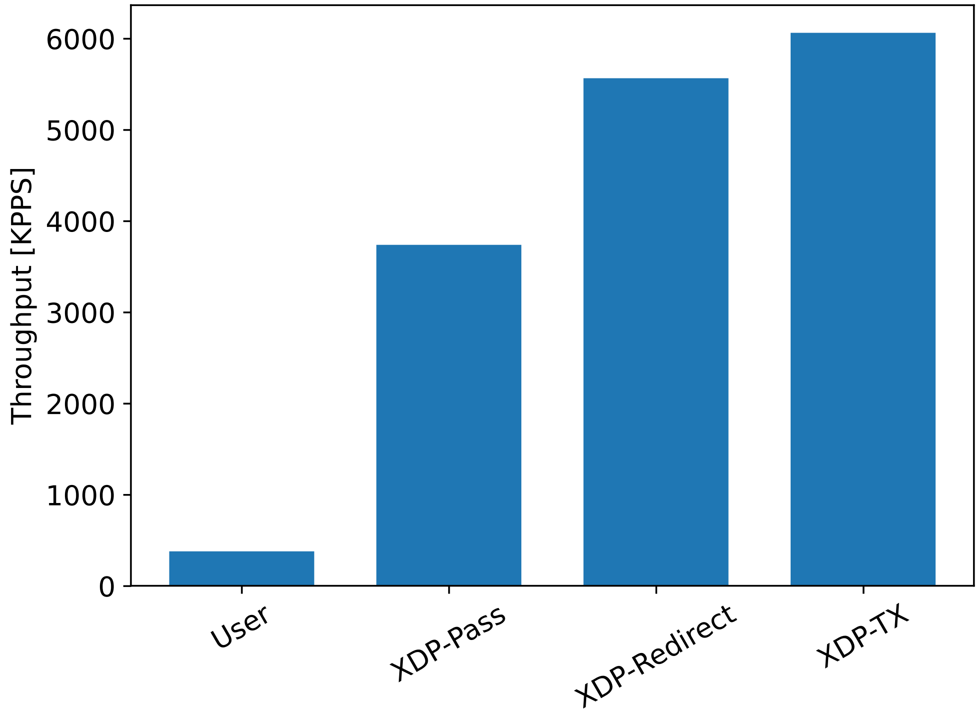

RFC 2544 Throughput Test: 10.000 Streams - 64 Byte Packets

Throughput Comparison Uplink - RSS

- XDP-Pass is lower compared to the Downlink – Limitation to one CPU
- XDP-Redirect CPU can improve about 2.5 fold

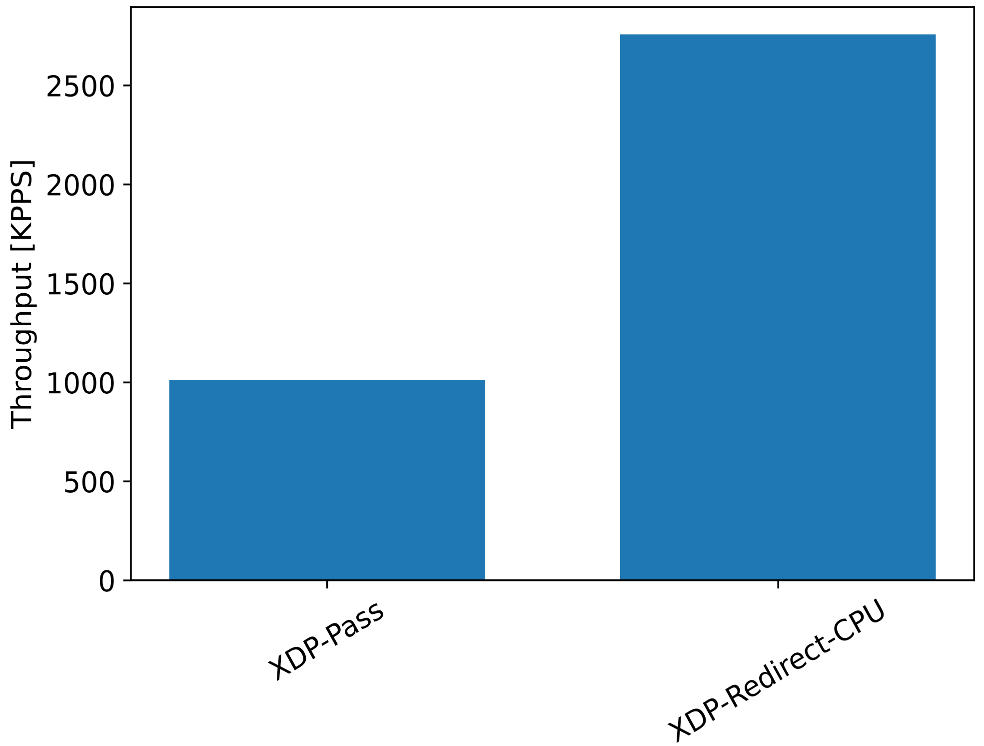

RFC 2544 Throughput Test: 10.000 Streams - 64 Byte Packets

Conclusion

- XDP increases the UPF throughput significantly in comparison to the user space implementation
- XDP CPU load-balancing on application layer for GTP-U traffic can increase the uplink capacity on the receiver side

- Further Work:
  - XDP with hardware acceleration in Smart NICs
  - Evaluate a UPF with XDP enhancements in virtualized environment

Sources

- Scott Bradner and Jim McQuaid. 1999. Benchmarking Methodology for Network Interconnect Devices. Request for Comments RFC 2544. Internet Engineering Task Force. https://doi.org/10.17487/RFC2544
- Thiago A. Navarro do Amaral, Raphael V. Rosa, David F. Cruz Moura, and Christian E. Rothenberg. 2021. An In-Kernel Solution Based on XDP for 5G UPF: Design, Prototype and Performance Evaluation. In 2021 17th International Conference on Network and Service Management (CNSM). Izmir, Turkey, 146–152. https://doi.org/10.23919/CNSM52442.2021.9615553
- Ericsson. 2022. Ericsson Mobility Report June 2022. (2022). https://www.ericsson.com/49d3a0/assets/local/reports-papers/mobility-report/documents/2022/ericsson-mobility-report-june-2022.pdf Accessed Nov. 1, 2022.
- ETSI. 2022. General Packet Radio System (GPRS) Tunnelling Protocol – User Plane (GTPv1-U). 3GPP TS 29.281 V17.3.0. (2022). https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=1699 Accessed Oct. 22, 2022.
- ETSI. 2022. Procedures for the 5G System (5GS). 3GPP TS 23.502 V16.13.0. (2022). https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3145 Accessed Oct. 11, 2022.
- ETSI. 2022. System architecture for the 5G System (5GS). 3GPP TS 23.501 V16.13.0. (2022). https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3144 Accessed Oct. 11, 2022.
- Toke Høiland-Jørgensen, Jesper Dangaard Brouer, Daniel Borkmann, John Fastabend, Tom Herbert, David Ahern, and David Miller. 2018. The EXpress Data Path: Fast Programmable Packet Processing in the Operating System Kernel. In Proceedings of the 14th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT ’18). Association for Computing Machinery, New York, NY, USA, 54–66. https://doi.org/10.1145/3281411.3281443

Sources

- Thiago A. Navarro do Amaral, Raphael V. Rosa, David F. Cruz Moura, and Christian Esteve Rothenberg. 2022. Run-Time Adaptive In-Kernel BPF/XDP Solution for 5G UPF. Electronics 11, 7 (2022), 1022. https://doi.org/10.3390/electronics11071022
- Open5GCore. 2023. Open5GCore. (2023). https://www.open5gcore.org/ Accessed Mar 1, 2023.
- Pablo Salva-Garcia, Ruben Ricart-Sanchez, Enrique Chirivella-Perez, Qi Wang, and Jose M. Alcaraz-Calero. 2022. XDP-Based SmartNIC Hardware Performance Acceleration for Next-Generation Networks. J. Netw. Syst. Manage. 30, 4 (oct 2022), 75 – 101. https://doi.org/10.1007/s10922-022-09687-z
- M Series. 2015. IMT Vision–Framework and overall objectives of the future development of IMT for 2020 and beyond. Recommendation ITU 2083, 0 (2015). https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.2083-0-201509-I!!PDF-E.pdf

Fraunhofer FOKUS

Contact

Christian Scheich
christian.scheich@fokus.fraunhofer.de

Fraunhofer FOKUS
Institute for Open Communication Systems
Kaiserin-Augusta-Allee 31
10589 Berlin, Germany
info@fokus.fraunhofer.de
www.fokus.fraunhofer.de