



High-Performance and Flexible Hardware Architectures for FPGA-Based SmartNICs

Klajd Zyla, Andreas Herkersdorf

TUM School of Computation, Information and Technology
Technical University of Munich, Germany

Motivation

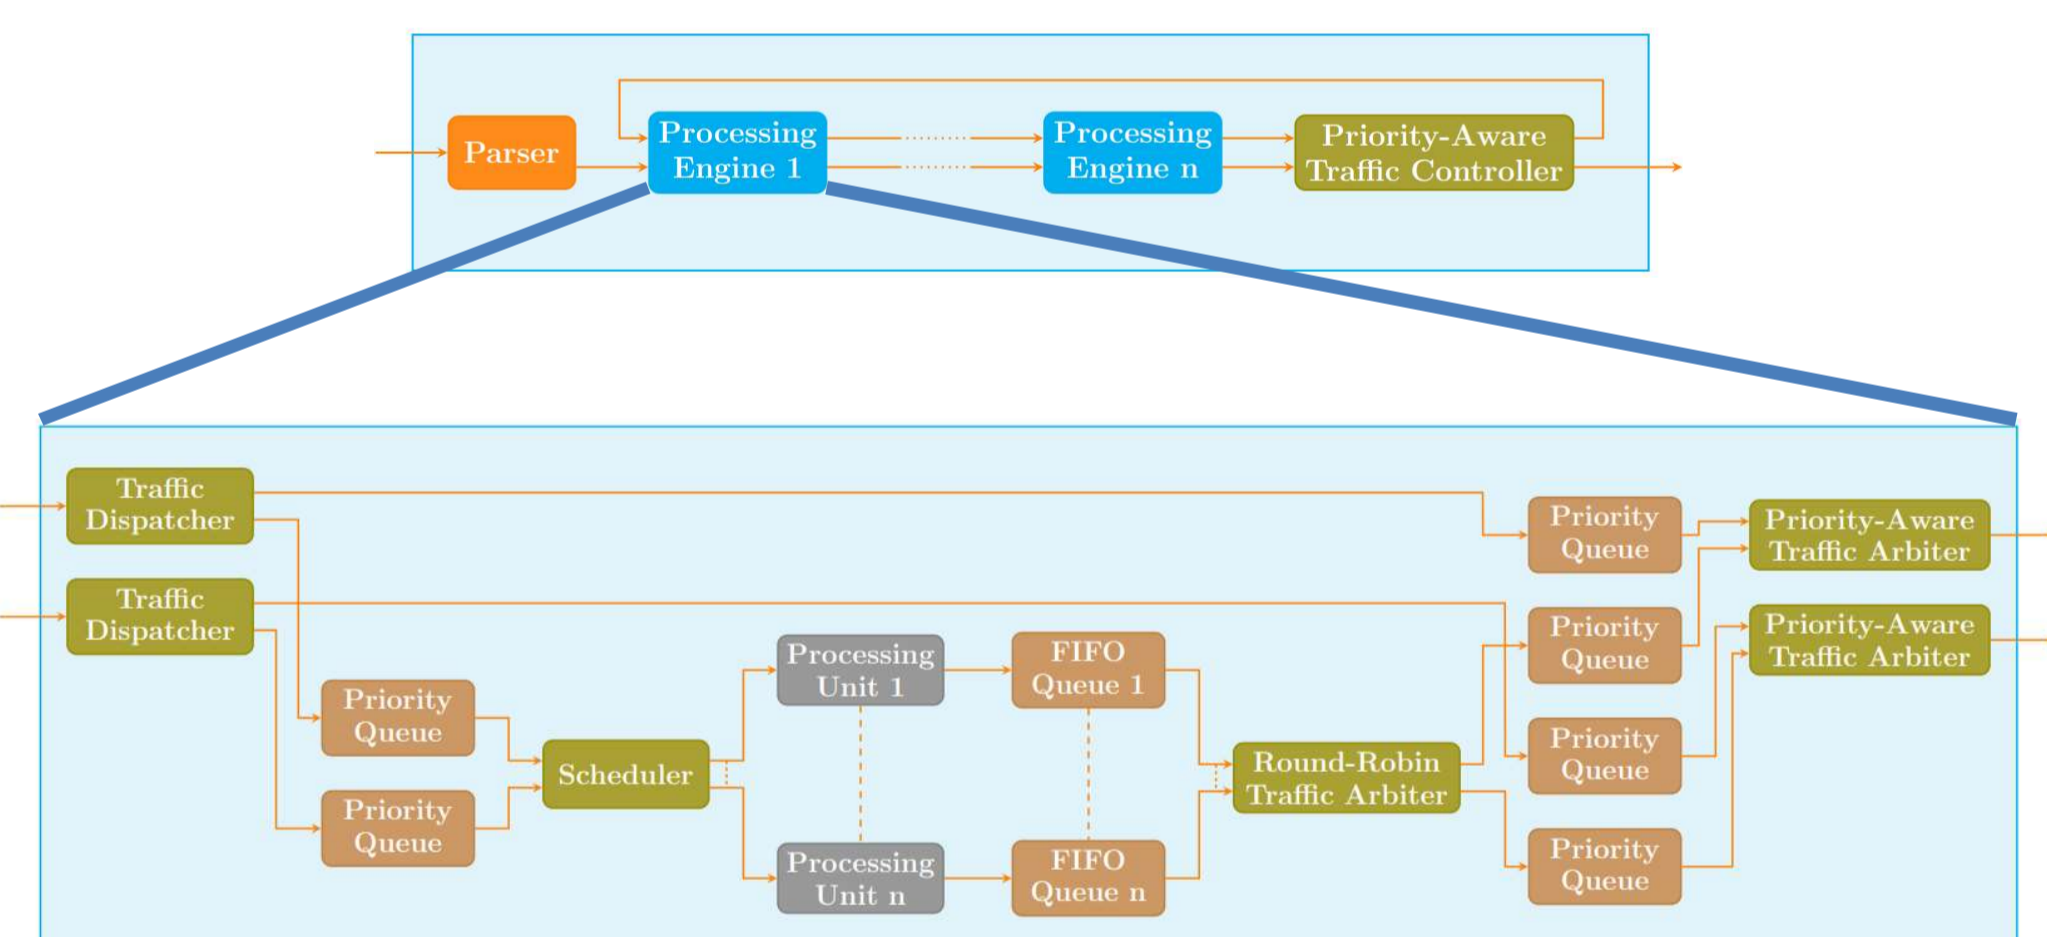
- **Challenges:**
 - Increasing processing demands on data centers.
 - High CPU power consumption and memory-access latency.
- **Solution: In-Network Computing:**
 - Shifts computational tasks to network infrastructure.
 - Benefits: Minimizes data movement, reduces latency, boosts performance.
- **Rising Number of Tasks in SmartNICs:**
 - Network traffic with increasingly diverse processing requirements.
 - Flexible packet-processing architectures are required.
- **State-of-the-Art Approaches:**
 - High-Abstraction programming frameworks – low performance
 - Hardware accelerators – limited flexibility

State of the Art

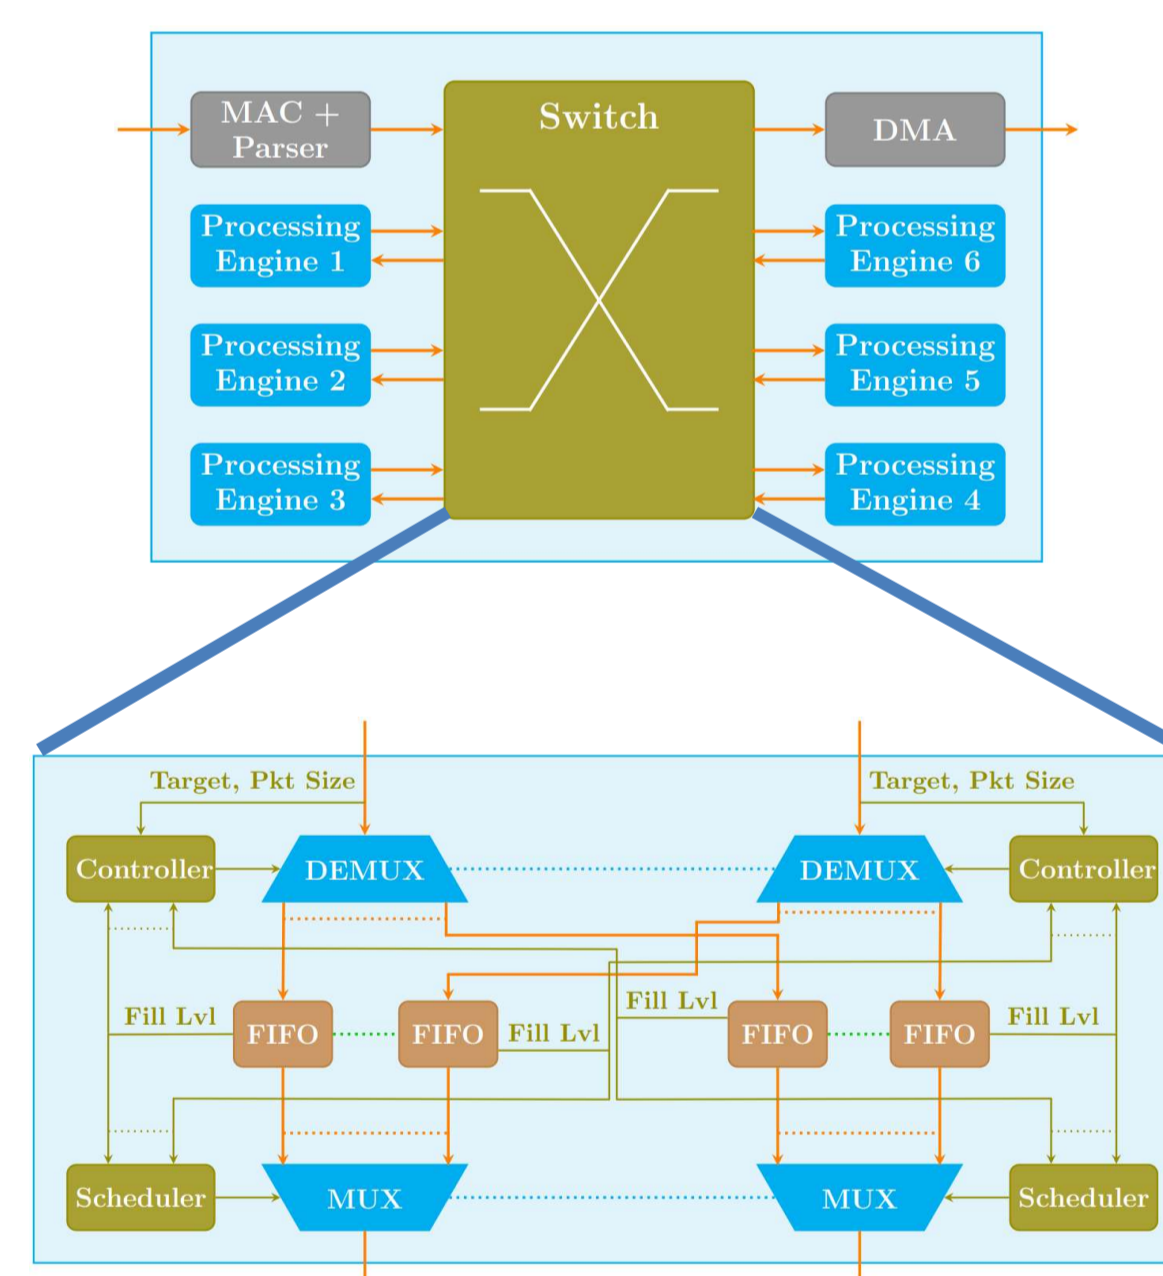
- **PANIC [1]:**
 - Input/Output-queued switch for flexible packet forwarding.
 - Centralized scheduler for priority-based and load-aware packet steering.
 - Low throughput: 48.3% for uniform random traffic.
- **CICQ switch [2]:**
 - Input/Crosspoint-queued switch avoids HoL blocking.
 - No centralized scheduler; distributed switch allocation.
 - Quadratic scalability (2^*N^2) in switch wires and queues.
- **ProNoC [3]:**
 - Open-source FPGA-optimized VC-based NoC.
 - Includes several router architecture optimizations.
 - Saturates at 28% throughput for uniform random traffic.

Proposed Hardware Architectures

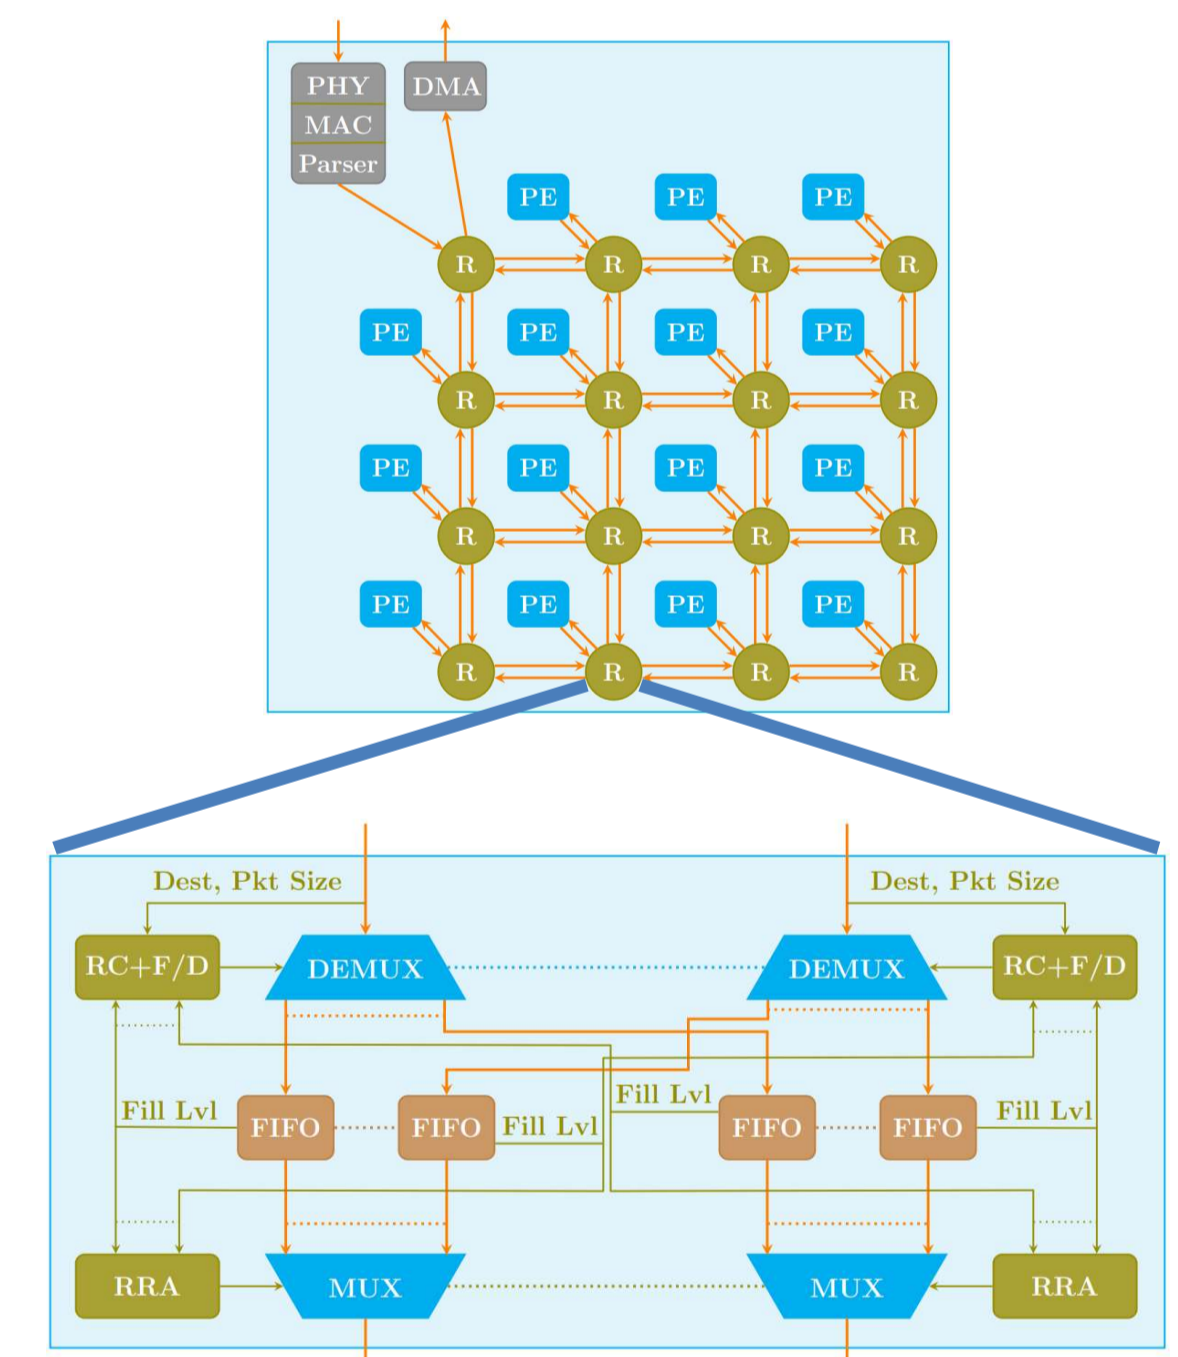
FlexRoute [4]



FlexCross [5]



HiPerNoC [6]



Experimental Evaluation

- Developed the architectures in SystemVerilog.
- Evaluated via RTL simulations in Vivado.
- Compared against state-of-the-art approaches.
- Integrated the designs into AMD OpenNIC Shell.
- Target device: Alveo U55C (Virtex UltraScale+ FPGA).
- All designs meet timing for 200 MHz (102.4 Gbit/s).
- FPGA specs: 1.3M LUTs, 601k LUTRAMs, 2.6M FFs, and 2016 BRAMs.

