SchedraNIC: Multi-Queue NIC Supporting Program-mable Packet Scheduling

Weitao PAN, Yichen DENG, Zekun WANG, Binhao YUE, Jiangyi SHI

Frontiers of Computer Science, DOI: 10.1007/s11704-026-51906-2

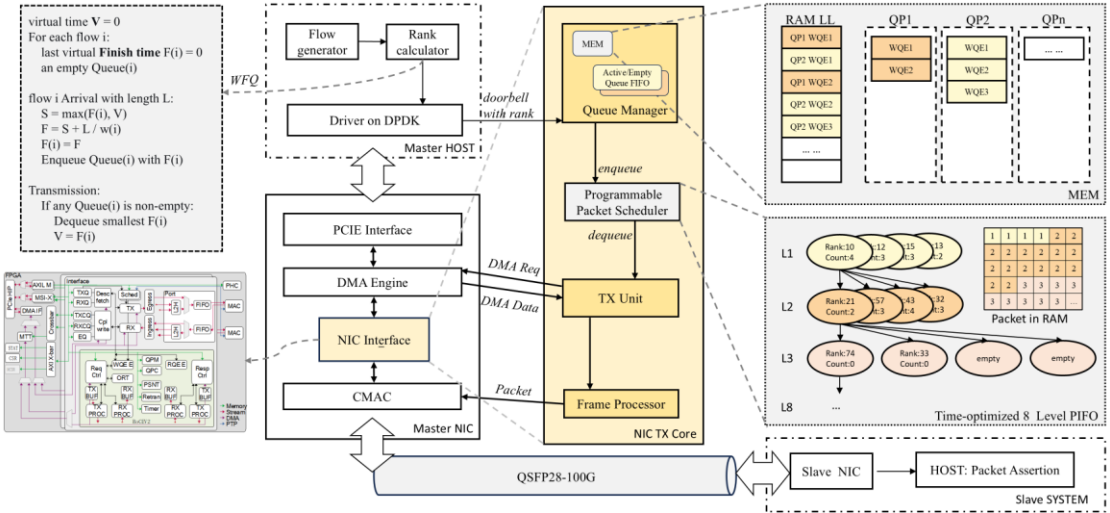
Problems & Ideas

Problems of existing packet scheduling systems:

- **Closed and non-reproducible scheduling architectures:**
  Commercial smart NICs embed scheduling logic inside proprietary hardware/software stacks, limiting flexibility and hindering research innovation.
- **Lack of hardware support in open-source platforms:**
  Existing open-source NIC frameworks (e.g., Corundum) provide programmable datapaths but lack native, efficient hardware scheduling primitives.
- **Poor scalability in multi-queue scenarios:**
  Conventional designs rely on centralized arbitration (e.g., large multiplexers), leading to severe timing and resource bottlenecks when scaling to thousands of flows.

Ideas: A programmable, hardware-efficient NIC scheduling platform

- We propose **SchedraNIC**, a multi-queue NIC architecture that integrates scalable queue management and programmable scheduling into a unified hardware platform.
- A **RAM-based linked-list multi-queue manager** eliminates arbitration bottlenecks and supports up to thousands of concurrent flows.
- A **timing-optimized BMW-PIFO scheduler** enables high-throughput rank-based packet scheduling with low hardware overhead.
- A **unified PUSH/POP abstraction** bridges scheduling algorithms and hardware execution, allowing flexible deployment of policies such as SP, WFQ, and EDF.



Main Contributions

Main Contributions

- This paper presents **SchedraNIC**, a programmable multi-queue NIC platform for efficient packet scheduling. The main contributions are summarized as follows:
- **A scalable multi-queue NIC architecture for programmable scheduling.**
  We design a hardware architecture that integrates queue management and scheduling, supporting up to thousands of concurrent flows while maintaining high throughput.
- **A RAM-based linked-list queue manager eliminating arbitration bottlenecks.**
  The proposed design replaces conventional centralized arbitration with a distributed queue structure, significantly improving scalability and reducing hardware complexity.
- **A timing-optimized BMW-PIFO scheduler for high-throughput execution.**
  By restructuring memory access into a pipelined design, the scheduler achieves efficient rank-based packet ordering with only **3–5% additional hardware overhead**.
- **A unified PUSH/POP abstraction bridging algorithms and hardware.**
  The platform enables seamless deployment of diverse scheduling policies (e.g., SP, WFQ, EDF), supporting flexible and policy-agnostic scheduling.
- **Comprehensive hardware-based evaluation on a real NIC platform.**
  Experimental results demonstrate up to **25.5% throughput improvement** and significantly improved scheduling accuracy compared to baseline systems.

