



DISSERTATION DEFENSE



HOSSEIN GOLESTANI

Architectural Enhancements for Data Transport in Datacenter Systems

Thursday, July 22, 2021

12:00 pm

[Virtual \(Passcode: 000372\)](#)

ABSTRACT: Substantial efforts over recent decades have enabled us to utilize datacenters as warehouse-scale computers and benefit from “XaaS”—infrastructure, platform, software, etc. as a Service. Datacenter systems run myriad applications, which frequently communicate with each other and/or Input/Output (I/O) devices, including network adapters and storage devices. Due to the growing speed of I/O devices, hyper-tenancy, and the emergence of microservice-based programming models, the I/O software stacks have become a critical factor in end-to-end communication performance. Datacenters rely on fast, efficient Software Data Planes (SDPs), which orchestrate data transfer between applications and I/O devices. The goal of this thesis is to enhance the performance, efficiency, and scalability of SDPs through hardware-software solutions.

In this thesis, I first characterize a state-of-the-art, high-performance SDP, which relies on shared-memory queues and spin-polling cores, on a real system. I pinpoint inefficiencies of spin-polling as a notification mechanism—namely, useless work, overheads, and adverse effect on co-running hyper-threads—and demonstrate its lack of queue scalability due to cache capacity constraints. Next, I introduce HyperPlane, a notification accelerator that replaces spin-polling. HyperPlane sorts ready queues to be serviced based on a particular policy; and allows a core to not iterate on empty queues, halt when all queues are empty, and efficiently share a queue with other cores. Finally, I introduce the HyperData accelerator, which enhances data transfer in SDPs through prefetching. HyperData discovers the exact memory locations that will be referenced and prefetches the corresponding data to the target core’s L1 cache. HyperData prefetches only the necessary parts of data buffers to avoid cache pollution, and is programmable to support various queue formats.

All in all, an accelerated data plane with HyperPlane and/or HyperData has significant advantages in terms of throughput, average/tail latency, and energy efficiency over a state-of-the-art baseline, with very small power and area overheads.

CHAIR: Prof. Thomas Wenisch