

# Building Latency-Critical Datacenter Systems

Host: Avi Silberschatz



Thursday, February 18, 2021  
10:30 a.m.

Speaker: Marios Kogias

Zoom Presentation

**ABSTRACT:** Online services play a major role in our everyday life for communication, entertainment, socializing, e-commerce, etc. These services run inside the datacenter under strict tail-latency service level objectives in order to remain interactive. The emergence of new hardware for IO has enabled microsecond-scale datacenter communications that challenge the efficiency of existing operating system and network mechanisms. Also, new in-network programmable devices start being deployed in datacenters and introduce a new computing paradigm that shifts functionality traditionally performed at the end-points to the network.

In this talk I will revisit the operating systems, networking, and distributed systems infrastructure specifically targeting latency-critical datacenter systems, while drawing intuition from basic queueing theory results. In the first part of the talk, I will focus on ZygOS[SOSP 2017], a system optimized for  $\mu$ -scale, in-memory computing on multicore servers. ZygOS implements a work-conserving scheduler within a specialized operating system designed for high request rates and a large number of network connections. ZygOS revealed the challenges associated with serving remote procedure calls (RPCs) on top of a byte-stream oriented protocol, such as TCP. In the second part of the talk, I will present R2P2[ATC 2019]. R2P2 is a transport protocol specifically designed for datacenter RPCs, that exposes the RPC abstraction to the endpoints and the network, making RPCs first-class datacenter citizens. R2P2 enables pushing functionality, such as scheduling, fault-tolerance, and tail-tolerance, inside the transport protocol. I will show how using R2P2 allowed us to offload RPC scheduling to programmable switches that can schedule requests directly on individual CPU cores.

**BIO:** Marios Kogias is a researcher at MSR Cambridge. He graduated from EPFL in August 2020. His main research focus is at the intersection of operating systems and networking in the datacenter. He's working on building and understanding systems with strict tail-latency SLOs leveraging new emerging hardware. He was an IBM PhD Fellow and won the best student paper award at Eurosys2020. Before joining EPFL he got his undergrad degree from the National Technical University of Athens. He has interned at Cern, Google, and Microsoft Research.