



## Mathematics and CS Seminar

# High Performance Data Center TCP Packet Processing

**Antoine Kaufmann (U of Washington)**

**Host: Vladimir Kolmogorov**

TCP is widely used for client-server communication in modern data centers, but TCP packet handling is notoriously CPU intensive. Our goal is to make data center TCP processing in a multi-tenant setting efficient, scalable, flexible, and secure. We propose a unique refactoring of TCP functionality that splits processing between a streamlined fast path for common operations, and an out-of-band slow path. Protocol processing executes in the kernel on dedicated cores that enforce correctness and resource isolation. Applications asynchronously communicate with the kernel through event queues, improving parallelism and cache utilization. We show that our approach can increase RPC throughput by up to 3x compared to Linux. Our fast-path can be offloaded to a programmable NIC to further improve performance and minimize CPU time for network processing. With hardware offload, data packets are delivered directly from application to application, while the NIC and kernel cooperate to enforce correctness and resource isolation. Using an emulation-based methodology, we show that our approach can increase per-core packet throughput by 8.2x compared to the Linux kernel TCP implementation.

**Tuesday, January 16, 2018 10:00am - 11:00am**

IST Austria Campus Big Seminar room Ground floor / Office Bldg West (I21.EG.101)



This invitation is valid as a ticket for the IST Shuttle from and to Heiligenstadt Station. Please find a schedule of the IST Shuttle on our webpage: <https://ist.ac.at/en/campus/how-to-get-here/> The IST Shuttle bus is marked IST Shuttle (#142) and has the Institute Logo printed on the side.