



Seminar/Talk

Redesigning Storage Systems for Future Workloads, Hardware, and Performance Requirements

Oana Balmau

University of Sidney

Host: Dan Alistarh

Cloud storage stacks are being challenged by new workloads, new hardware and new performance requirements. First, workloads evolved from following a read-heavy pattern (e.g., a static web-page) to a write-heavy profile where the read:write ratio is closer to 1:1 (e.g., as in the Internet of Things). Second, the hardware is undergoing rapid changes. The divide between fine-grained volatile memory and slow block-level storage is rapidly being bridged by the emerging byte-addressable non-volatile memory devices and the fast block-addressable NVMe SSDs (e.g., Intel Optane NVMe SSDs). Third, performance requirements in storage systems now emphasize low tail latency, in addition to high throughput. In this talk I will argue that existing storage systems have fundamental limitations that do not allow them to fully meet these challenges, and that therefore the storage stack needs to undergo radical change. In particular, using state-of-the-art key-value stores I will demonstrate that with modern workloads and hardware the bottleneck shifts from I/O to CPU, invalidating an assumption that has underpinned all past storage system design. In line with this observation I will then present a new design paradigm for key-value stores that departs from the conventional wisdom of optimizing disk usage and instead optimizes CPU usage. To do so, we keep data unsorted on disk, reduce contention for shared data structures, and do away with expensive maintenance operations. This design has been implemented in the Kvell key-value store. Kvell outperforms state-of-the-art key-value stores such as RocksDB in both read- and write-heavy workloads, running on modern NVMe SSDs. Thanks to its novel design, Kvell achieves up to 5x better throughput, and up to two orders of magnitude lower tail latency.

Monday, February 24, 2020 10:00am - 11:00am

Mondi Seminar Room 2, Central Building



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