Chenhao Ye

chenhaoy@cs.wisc.edu • https://pages.cs.wisc.edu/~chenhaoy

EDUCATION

Ph.D. student in Computer Science, University of Wisconsin–Madison — Madison, WI, Sep 2020 – Present
- Advisors: Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau
- Research Interests: Storage Systems, Distributed Systems, Databases

B.S. in Computer Science, University of Wisconsin–Madison — Madison, WI, Jan 2019 – Aug 2020
- GPA: 4.00 / 4.00, Graduated with Honors

B.S. in Electrical and Computer Engineering, Shanghai Jiao Tong University — Shanghai, China, Sep 2016 – Dec 2018
- GPA: 3.73 / 4.00

WORK EXPERIENCE

Microsoft Research, Data Systems Group — Redmond, WA, May 2024 – Aug 2024

Research Intern, supervised by Vasileios Zois and Badrish Chandramouli
- Designed, implemented, and evaluated a scalable replication protocol for a high-performance distributed key-value store.
- *Details of this project are omitted due to a Non-Disclosure Agreement (NDA).*

Snowflake, Global Platform Team — San Mateo, CA, May 2023 – Aug 2023

Software Engineer Intern, supervised by Leonidas Galanis
- Prototyped a workload replay tool that generates realistic workloads based on job statistics from the data infrastructure; this tool attracted significant interest from multiple internal customers for benchmarking and resource provisioning purposes.
- Implemented a critical multi-metadata-store feature for a new data ingestion system, which resolved the major blocker for this new system's production deployment.
- Uncovered and fixed a subtle concurrency bug in the data infrastructure that could cause data loss in production.

RESEARCH PROJECTS

Cache-Centric Multi-Resource Allocation for Storage Services, Project Leader — Jan 2021 – Present
- Present a resource allocation framework for multi-tenant storage systems that leverages the demand correlation between cache sizes and other resources (*e.g.*, I/O, network) to optimize resource utilization while maintaining fairness.
- Develop *HopperKV*, a multi-tenant Redis-based key-value store that caches data for DynamoDB; by judiciously allocating the cache sizes among tenants, HopperKV optimizes the DynamoDB utilization, achieving up to 1.9× higher throughput.
- Build *BunnyFS*, a multi-tenant local filesystem for high-performance NVMe SSDs; by optimizing page cache allocations among tenants, BunnyFS delivers up to 1.4× higher throughput.

Enabling Transaction Priority in Optimistic Concurrency Control, Project Leader — Oct 2021 – Apr 2023
- Propose a lightweight reservation mechanism for the optimistic concurrency control (OCC) protocol that protects high-priority transactions from being aborted by low-priority transactions in the case of conflicts.
- Design and implement *Polaris*, an OCC protocol that supports multiple priority levels; benchmarks show it can achieve up to 1.9× higher throughput and 17× lower latency compared to an existing OCC protocol on high-contention workloads.

MadFS: Per-File Virtualization for Userspace Persistent Memory Filesystems, Project Co-Leader — Oct 2021 – Jan 2023
- Propose a novel *per-file virtualization* technique for persistent memory filesystems, which encapsulates a set of filesystem functionalities, including metadata management, crash consistency, and concurrency control, fully in userspace; this technique significantly reduces the kernel-crossing overhead on the critical path.
- Build *MadFS*, a kernel-bypassing persistent memory filesystem based on the per-file virtualization, which achieves up to 1.5× speedup for LevelDB on YCSB workload and 1.9× for SQLite on TPC-C workload.

PUBLICATIONS

[1] Sambhav Satija, Chenhao Ye, Ranjitha Kosgi, Aditya Jain, Romit Kankaria, Yiwei Chen, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Kiran Srinivasan. Cloudscape: A Study of Storage Services in Modern Cloud Architectures. In *23rd USENIX Conference on File and Storage Technologies*. **FAST '25**

[2] Chenhao Ye, Wuh-Chwen Hwang, Keren Chen, Xiangyao Yu. Polaris: Enabling Transaction Priority in Optimistic Concurrency Control. In *Proceedings of the 2023 International Conference on Management of Data*. **SIGMOD '23**

[3] Shawn Zhong*, Chenhao Ye*, Guanzhou Hu, Suyan Qu, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Michael M. Swift. MadFS: Per-File Virtualization for Userspace Persistent Memory Filesystems. In *21st USENIX Conference on File and Storage Technologies*. (*contributed equally) **FAST '23**

[4] Yuvraj Patel, Chenhao Ye, Akshat Sinha, Abigail Matthews, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau, Michael M. Swift. Using Trāṭr to tame Adversarial Synchronization. In *31st USENIX Security Symposium*. **USENIX Security '22**

[5] Jing Liu, Anthony Rebello, Yifan Dai, Chenhao Ye, Sudarsun Kannan, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. Scale and Performance in a Filesystem Semi-Microkernel. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles*. **SOSP '21**