Kaiwei Tu

Department of Computer Science – University of Wisconsin–Madison
Madison, WI
ktu5@wisc.edu

Education

University of Wisconsin–Madison **Madison, WI**
Ph.D. Candidate in Computer Sciences *2021–Present*
Advised by Prof. Remzi H. Arpaci-Dusseau and Prof. Andrea C. Arpaci-Dusseau

University of Wisconsin–Madison **Madison, WI**
M.S. in Computer Sciences *2019–2021*

University of Michigan **Ann Arbor, MI**
B.S. in Engineering, Computer Science, summa cum laude *2015–2019*

Shanghai Jiao Tong University **Shanghai, China**
B.S. in Electrical and Computer Engineering *2014–2019*

Research Interests

Storage Systems, Operating systems, and AI and systems co-design

Publications

[1]: **Kaiwei Tu**, Kan Wu, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. "Getting the MOST out of your Storage Hierarchy with Mirror-Optimized Storage Tiering." *USENIX FAST '26*.

[2]: Kan Wu, **Kaiwei Tu**, Yuvraj Patel, Rathijit Sen, Kwanghyun Park, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. "NyxCache: Flexible and Efficient Multi-tenant Persistent Memory Caching." *USENIX FAST '22*.

[3]: Kan Wu, Zhihan Guo, Guanzhou Hu, **Kaiwei Tu**, Ramnatthan Alagappan, Rathijit Sen, Kwanghyun Park, Andrea C. Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. "The Storage Hierarchy is Not a Hierarchy: Optimizing Caching on Modern Storage Devices with Orthus." **USENIX FAST '21**.

Research Experience

Advanced Systems Laboratory (ADSL), University of Wisconsin–Madison **Madison, WI**
Research Assistant *May 2021–Present*

- Investigating load balancing mechanisms for tiered memory systems to reduce inter-job interference.
- **MOST (FAST '26):** Developing a mirror optimized tiering scheme for heterogeneous storage hierarchies.
- **NyxCache (FAST '22):** Designed and implemented the Nyx controller for multi-tenant persistent memory caching framework; evaluated QoS, proportional sharing, and fair slowdown policies.
- **Orthus (FAST '21):** Implemented Non Hierarchy Caching (NHC) in Intel's Open CAS Linux kernel module and evaluated performance under dynamic workloads.

Professional Experience

Meta **Menlo Park, CA**
AI and Systems Co-design, Software Engineering Intern *May 2025–Aug 2025*

- Investigated hyperscale model inference workloads and developed AI workload benchmarks that contributed to DCPerf, a benchmark suite designed to reflect real world hyperscale cloud applications in data centers.

Google **Sunnyvale, CA**
Google Cloud, Research Intern *May 2022–Aug 2022*
- Collaborated with the Borglet team to evaluate the effectiveness of memory bandwidth isolation mechanisms on AMD processors.

Citrix **Fort Lauderdale, FL**
Software Development Engineer Intern *May 2018–Aug 2018*
- Enhanced cloud VM provisioning service; led development of a RESTful API and UI for deployment monitoring.

Teaching Experience

University of Wisconsin–Madison **Madison, WI**
Teaching Assistant, Advanced Operating Systems *Fall 2023*
- Led discussion sections, designed problem sets, and managed grading for a graduate level operating systems course.

University of Wisconsin–Madison **Madison, WI**
Teaching Assistant, Introduction to Computer Systems *Spring 2023*
- Held office hours, assisted students with laboratory exercises, and conducted exam review sessions.

Honors & Awards

2022, 2026: USENIX Student Grant, FAST '22 and FAST '26

2019: James B. Angell Scholar

2018: EECS Scholars

2017–2019: Dean's List; University Honors

Academic Service

2025: EuroSys 2026 Shadow Program Committee