Kaiwei Tu

Email: ktu5@wisc.edu,Website: https://kaiweitu.github.io/

Education

University of Wisconsin - Madison — **Madison, WI**
Ph.D. candidate in Computer Science — *Sept 2019 - Present*
Advisor: Prof. Remzi H. Arpaci-Dusseau and Prof. Andrea C. Arpaci-Dusseau

University of Michigan — **Ann Arbor, MI**
Bachelor of Science in Engineering major in Computer Science, Summa Cum Laude (highest distinction) — *April 2019*

Shanghai Jiao Tong University — **Shanghai, China**
Bachelor of Science in Electrical and Computer Engineering — *August 2017*

Research Experience

The ADvanced Systems Laboratory (ADSL) — **Ann Arbor, MI**
Research Assistant — *April 2021 - present*

- I am currently exploring the development of a modern memory tiering system.
- **MOST** is a novel tiering-based approach designed for modern heterogeneous storage hierarchy.
- **Orthus** is a research project published in **FAST'21** that introduces a new caching mechanism called Non-Hierarchy Caching (NHC), designed to dynamically offload saturated caching devices when necessary. I implemented several state-of-the-art caching strategies within the kernel module of Intel's open-source caching framework, Open CAS Linux, to compare their performance against NHC. The results demonstrated that NHC outperforms existing approaches across a broad spectrum of workloads, especially in dynamic environments.
- **NyxCache** is a research project published in **FAST'22** that introduces a flexible and efficient caching framework for multi-tenant caches using persistent memory (PM). One of my key contributions was developing the Nyx-controller, a crucial component that enforces various sharing policies within a multi-tenant caching environment. I also conducted experiments that demonstrated NyxCache's ability to effectively implement policies such as quality-of-service, proportional sharing, and fair slowdown—capabilities that traditional DRAM-based methods are unable to achieve.

Work Experience

Meta — **Menlo Park, CA**
Software Engineering Intern, host: Wei Su — May 2025 - Current

- I worked with the AI & System Co-design Compute and Storage team to build a AI workload benchmark for CPU.

Google — **Sunnyvale, CA**
Research Intern, host: Carlos Villavieja and Sree Kodakara. — *May 2022 - Aug 2022*

- I worked with the Brog team, Google's large-scale cluster management system. I conducted a comprehensive evaluation of four memory bandwidth regulation mechanisms on AMD machines across the fleet. I proposed a hybrid policy that combines CPU jailing and resctrl, leading to improved memory bandwidth regulation efficiency at scale.

Citrix System — **Fort Lauderdale, FL**
Software Development Engineering Intern — *April 2018 – August 2018*

- Enhanced a provisioning service to automatically deploy Citrix VMs on the cloud. Led the design and development of a new user interface for checking provisioning information by creating a corresponding RESTful API.

Publications

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

[2] Kan Wu, **Kaiwei Tu**, Yuvraj Patel, Rathijit Sen, Kwanghyun Park, Andrea C. Arpaci-Dusseau, Remzi H. Arpaci-Dusseau. NyxCache: Flexible and Efficient Multi-tenant Persistent Memory Caching. In *Proceedings of the 20th USENIX Conference on File and Storage Technologies* (FAST '22), February 2022.

[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. In *Proceedings of the 19th USENIX Conference on File and Storage Technologies* (FAST '21), Virtual, February 2021.

Awards & Grants

- USENIX Student Grant: FAST'22.
- James B. Angell Scholar (2019 Spring); EECS Scholars (2018 Fall).
- Dean's List: 2017 Fall, 2018 Spring, 2018 Fall, 2019 Spring; University Honors: 2017 Fall, 2018 Spring, 2018 Fall.