

WAN Hu·万虎

Department of Computer Science, City University of Hong Kong
hu.wan@my.cityu.edu.hk ◊ <http://www.cs.cityu.edu.hk/~huwan2>

ABOUT ME

I am currently a Postdoctoral Fellow in the Department of Computer Science at City University of Hong Kong (CityUHK), where I received my Ph.D. in Computer Science in 2023. Prior to joining CityUHK, I received my B.Eng. in Computer Science and Technology and M.Sc. in High Reliability Embedded Systems from Capital Normal University.

RESEARCH INTERESTS

My research lies at the intersection of AI-system co-design, memory and storage systems, machine learning infrastructure, and embodied intelligence.

I have extensive experience in storage-efficient architectures and system optimization for large-scale recommendation and deep learning workloads, including non-volatile memory, in-storage computing, SSD-based embedding management, and tail-latency-aware scheduling.

My current research focuses on high-performance edge intelligence and self-evolving robotic systems. In particular, I investigate the integration of foundation models (e.g., LLMs and VLA models) with reinforcement learning to enable adaptive and self-improving robotic manipulation, while optimizing their deployment on resource-constrained edge platforms.

RECENT PUBLICATIONS

- Rabbitail: A Tail Latency-Aware Scheduler for Deep Learning Recommendation Systems with Hierarchical Embedding Storage
Hu Wan, Yun Huang, Shuhan Bai, Xuan Sun, Tei-Wei Kuo, and Chun Jason Xue
In *the 40th ACM/SIGAPP Symposium On Applied Computing (SAC)*. 2025.
- Binary Folding Compression for Efficient Software Distribution
Jinheng Li, Qiao Li, **Hu Wan**, and Chun Jason Xue
In *the 39th ACM/SIGAPP Symposium On Applied Computing (SAC)*. 2024.
Award: ACM Europe Award.
- Pipette: Efficient Fine-Grained Reads for SSDs
Shuhan Bai, **Hu Wan**, Yun Huang, Xuan Sun, Fei Wu, Changsheng Xie, Hung-chih Hsieh, Tei-wei Kuo, and Chun Jason Xue
In *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, vol. 42, no. 12, pp. 4721-4734, Dec. 2023.
- Pipette: Efficient Fine-Grained Reads for SSDs
Shuhan Bai*, **Hu Wan***, Yun Huang, Xuan Sun, Fei Wu, Changsheng Xie, Tei-wei Kuo, Jason Xue, and Hung-chih Hsieh
In *59th ACM/IEEE Design Automation Conference (DAC)*, July 10-14, 2022.
- RM-SSD: In-Storage Computing for Large-Scale Recommendation Inference.
Xuan Sun*, **Hu Wan***, Qiao Li, Chia-Lin Yang, Tei-Wei Kuo, and Chun Jason Xue
In *28th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, February 12-16, 2022.

* Asterisks in author lists denote equal contribution.

- FlashEmbedding: Storing Embedding Tables in SSD for Large-Scale Recommender Systems.
Hu Wan*, Xuan Sun*, Yufei Cui, Chia-Lin Yang, Tei-Wei Kuo, and Chun Jason Xue
 In *ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)*, August 24-25, 2021.

WORK EXPERIENCE

City University of Hong Kong (香港城市大學) Jul. 2023 - Present
 Department of Computer Science Postdoctoral Fellow
 Supervisor: Prof. Chun Jason Xue & Prof. Nan Guan

EDUCATION

City University of Hong Kong (香港城市大學) Jul. 2017 - Jun. 2023
 Department of Computer Science Ph.D.
 Supervisor: Prof. Chun Jason Xue

Capital Normal University (首都师范大学) Sep. 2014 - Jun. 2017
 College of Information Engineering M.Sc.
 Supervisor: Assoc. Prof. Yuanchao Xu

Capital Normal University (首都师范大学) Sep. 2009 - Jul. 2013
 College of Information Engineering B. Eng.

TEACHING

- CS5296: Cloud Computing: Theory and Practice, Teaching Assistant, Semester B, 2019/2020
- CS3103: Operating Systems, Teaching Assistant, Semester A, 2019/2020
- CS5296: Cloud Computing: Theory and Practice, Teaching Assistant, Semester B, 2018/2019
- CS3103: Operating Systems, External Teaching Assistant, Semester A, 2018/2019
- CS2311: Computer Programming, Teaching Assistant, Semester A, 2018/2019
- CS4465: Computational Biology and Bioinformatics, Teaching Assistant, Semester B, 2017/2018
- CS2116: Computer Systems, Teaching Assistant, Semester A, 2017/2018

SERVICES

- Organizing Committees:
 - General Chair: 1st Workshop on Frontiers in Memory and Storage Systems for LLMs, 2025
 - Registration Chair: 2nd Workshop on Frontiers in Memory and Storage Systems for LLMs, 2026
 - Web Chair: ACM/SIGAPP Symposium on Applied Computing (CPS Track), 2018–2026
- Session Chair: ACM/SIGAPP Symposium on Applied Computing (CPS Track), 2024
- Reviewer:
 - IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)
 - ACM Transactions on Embedded Computing Systems (TECS)
 - ACM Computing Surveys (CSUR)