

Kan Wu

Operating Systems • Machine Learning Systems

<https://cs.wisc.edu/~kanwu> • kanwu.wisc@gmail.com • +1 (608) 960-6750 • San Francisco, CA

SUMMARY

Systems generalist interested in cross-stack hardware-model co-design. At xAI, contributed to training the text upsampler for Grok Imagine image generation, Grok speculative decoding training and inference, inference engine development, and led the Applied Inference team building large-scale (batch) inference infrastructure. At Google, co-designed models and systems for Gemma 3n on-device sparse-attention models and implemented Gemini 2.5 Flash CPU serving in gemma.cpp. And PhD in operating systems, with a focus on memory management.

PROFESSIONAL EXPERIENCE

xAI — Member of Technical Staff

MAY 2025 – PRESENT

- **[Image Generation Model Training]** Grok Imagine text upsampler training: SFT, RL training and data recipe.
- **[Speculative Decoding]** Grok models speculative decoding model design, training and inference.
- **[SGLang Inference Engine Development]** scheduling, memory management, constrained decoding, etc.
- **[Large Scale Batch Inference]** led applied inference team - large-scale inference infra (load balancing, batch inference supporting data labeling and RL, auto-scaling).

SystemsResearch@Google — Senior Software Engineer

SEP 2022 – MAY 2025

- **[Model for Edge Devices]** Gemma 3n - sparse attention (paper published in NeurIPS 2025) co-design and kernel implementations.
- **[Gemma.cpp]** Gemini 2.5 Flash on CPU inference implementation in Gemma.cpp.
- **[Operating Systems Memory Management]** OS memory efficiency - eBPF based fleet-wide memory-tiering systems (USENIX ATC 2025), profile-guided fleet hugepage efficiency (techniques upstreamed to LLVM).

EDUCATION

2016-2022

UNIVERSITY OF WISCONSIN – MADISON

MADISON, WI

*Ph.D. in Computer Science. Advisors: Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau
Areas: Operating Systems, Distributed Systems*

2012-2016

UNIVERSITY OF SCIENCE AND TECHNOLOGY OF CHINA

HEFEI, ANHUI

B.Sc. in Computer Science (with honor)

SELECTED PUBLICATIONS

[1] Spark Transformer: Reactivating Sparsity in FFN and Attention.

Chong You*, **Kan Wu***(co-first author), Zhipeng Jia*, Lin Chen*, Srinadh Bhojanapalli, Jiaxian Guo, Utku Evci, Jan Wassenberg, Praneeth Netrapalli, Jeremiah J. Willcock, Suvinay Subramanian, Felix Chern, Alek Andreev, Shreya Pathak, Felix Yu, Prateek Jain, David E. Culler, Henry M. Levy, Sanjiv Kumar.

NeurIPS'2025, techniques shipped in Google Gemma-3n models

[2] PageFlex: Flexible and Efficient User-space Delegation of Linux Paging Policies with eBPF.

Anil Yelam, **Kan Wu***(corresponding author), Zhiyuan Guo, Rajath Shashidhara, Stanko Novakovic, Suli Yang, Wei Xu, Alex C. Snoeren, Kimberly Keeton.

USENIX ATC'2025, techniques landed in Google fleet memory swapping systems

[3] The Storage Hierarchy is Not a Hierarchy: Optimizing Caching on Modern Storage Devices with Orthus.

Kan Wu, Zhihan Guo, Guanzhou Hu, Kaiwei Tu, Ramnaththan Alagappan, Rathijit Sen, Kwanghyun Park, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau.

USENIX FAST'2021

[4] **Arya: Arbitrary Graph Pattern Mining with Decomposition-based Sampling.**
Zeying Zhu*, **Kan Wu*** (co-first author), Zaoxing Liu.
USENIX NSDI'2023

[5] **NyxCache: Flexible and Efficient Multi-tenant Persistent Memory Caching.**
Kan Wu, Kaiwei Tu, Yuvraj Patel, Rathijit Sen, Kwanghyun Park, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau.
USENIX FAST'2022

PROFESSIONAL SERVICES

Program committee: OSDI'26, MLSYS'26, FAST'25, ATC'25, ATC'24, HotStorage'24

Proceedings chair: SOSP'24

OTHER PUBLICATIONS

[6] **MDK: rethinking the data center memory reclamation problem.**
Shaurya Patel, Suli Yang, Yawen Wang, **Kan Wu**, Alexandra Fedorova, Margo Seltzer, Kimberly Keeton.
USENIX OSDI'2026

[7] **OBASE: Object-Based Address-Space Engineering to Improve Memory Tiering.**
Vinay Banakar, Suli Yang, **Kan Wu**, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau, Kimberly Keeton.
USENIX OSDI'2026

[8] **Getting the MOST out of your Storage Hierarchy with Mirror-Optimized Storage Tiering.**
Kaiwei Tu, **Kan Wu**, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau.
USENIX FAST'2026

[9] **FineMem: Breaking the Allocation Overhead vs. Memory Waste Dilemma in Fine-Grained Disaggregated Memory Management.**
Xiaoyang Wang, Yongkun Li, **Kan Wu**, Wenzhe Zhu, Yuqi Li, Yinlong Xu.
USENIX OSDI'2025

[10] **SLAP: Segmented Reuse-Time-Label Based Admission Policy for Content Delivery Network Caching.**
Ke Liu, **Kan Wu**, Hua Wang, Ke Zhou, Peng Wang, Ji Zhang, Cong Li.
ACM TACO'2024

[11] **SLAP: An Adaptive, Learned Admission Policy for Content Delivery Network Caching.**
Ke Liu, **Kan Wu**, Hua Wang, Ke Zhou, Ji Zhang, Cong Li.
IPDPS'2023

[12] **WiscSort: External Sorting for Byte Addressable Storage.**
Vinay Banakar, **Kan Wu**, Yuvraj Patel, Kimberly Keeton, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau.
VLDB'2023

[13] **Cornus: Atomic Commit for Cloud DBMS with Storage.**
Zhihan Guo, Xinyu Zeng, **Kan Wu**, Wuh-Chwen Hwang, Ziwei Ren, Xiangyao Yu, Mahesh Balakrishnan, Philip A. Bernstein.
VLDB'2022

[14] **Releasing Locks As Early As You Can: Reducing Contention of Hotspots by Violating Two-Phase Locking.**
Zhihan Guo, **Kan Wu**, Cong Yan, Xiangyao Yu.
SIGMOD'2021

[15] **Read as Needed: Building WiSER, a Flash-Optimized Search Engine.**
Jun He, **Kan Wu**, Sudarsun Kannan, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau.
USENIX FAST'2020

[16] **Towards an Unwritten Contract of Intel Optane SSD.**
Kan Wu, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau.
USENIX HotStorage'2019

[17] **Exploiting Intel Optane SSD for Microsoft SQL Server.**
Kan Wu, Andrea Arpaci-Dusseau, Remzi Arpaci-Dusseau.
ACM DaMoN@SIGMOD'2019