

International Workshop on Big Memory (BigMem) 2025

Co-located with ACM Symposium on Operating Systems Principles (SOSP) 2025

Seoul, Republic of Korea, October 13, 2025

<https://bigmem2025.github.io/>

Overview

The rapid evolution of memory-centric computing technologies—including Compute Express Link (CXL), persistent memory, and disaggregated architectures—is fundamentally reshaping system software design paradigms. The International Workshop on Big Memory (BigMem 2025) establishes a premier forum for researchers and practitioners to explore the challenges and opportunities in managing terabyte-to-petabyte scale memory hierarchies in modern computing systems.

BigMem 2025 aims to foster cross-disciplinary collaboration between operating system researchers, computer architects, and industry practitioners. By bringing together diverse perspectives, we seek to accelerate innovation in big memory systems and establish this workshop as the leading venue for memory-centric computing research. BigMem 2025 will be co-located with SOSP and registrations will be handled by SOSP 2025.

Topics of Interest

This workshop will focus on important research directions, including operating system support for heterogeneous memory systems integrating DRAM, CXL and/or persistent memory, memory-centric networking technologies enabling efficient remote memory access, and innovative applications leveraging massive memory capacities. Moreover, we will examine emerging challenges in memory virtualization, coherence protocols, and fault tolerance for next-generation memory architectures.

The topics include, but are not limited to: Memory-system design (hardware/software co-optimization); Disaggregated memory architectures (e.g., CXL, NVMe-over-Fabrics); Operating systems for hybrid/heterogeneous memory (NVM, DRAM, HBM); Design and operation of large-scale memory systems; Emerging memory technologies (3D XPoint, STT-MRAM, FeRAM, optical memory); Memory failure modes, reliability, and fault mitigation; Energy-efficient memory subsystems; Memory security (side-channel attacks, encryption, secure allocators); Memory safety for critical systems; Memory-centric programming models and languages; In-memory and near-memory computing architectures; Algorithmic memory optimizations (caching, prefetching, compression); Non-volatile memory

programming paradigms; In-memory databases, NoSQL stores, and analytics; Memory-driven AI/ML workloads; Embedded and autonomous systems; Memory disaggregation for cloud/edge environments; Bio-inspired memory architectures. By bridging architecture, systems, and applications, BigMem 2025 seeks to shape future research directions in operating systems.

Submission Guidelines

We invite original research contributions that have not been published previously or submitted concurrently to other venues, including any other conference or journal. Authors should prepare their work as a two-page extended abstract (references excluded) in English, formatted as a PDF document. The ACM submission template is recommended.

All submissions must be made through the workshop's online submission system and will undergo a double-blind review process by our program committee. Please ensure your submission is properly anonymized to maintain the integrity of the review. Submissions will be evaluated on their technical merit, novelty, relevance to the workshop themes, and clarity of presentation. This workshop focuses on discussion and feedback rather than archival publication, and therefore does not produce formal proceedings.

Important Dates

Submissions due: July 30, 2025 (AOE)

Notification of acceptance: August 30, 2025

Workshop presentation: October 13, 2025

Program Chairs:

Yu Hua, *Huazhong Univ. Sci. and Tech*

Ming-Chang Yang, *The Chinese Univ. of HK*

Web/Submission Chair:

Yuxuan Mo, *Huazhong Univ. Sci. and Tech.*

Program Committee Member:

Zhichao Cao, *Arizona State University*

Yue Cheng, *The University of Virginia*

Vijay Chidambaram, *The University of Texas at Austin*

Mingkai Dong, *Shanghai Jiao Tong University*

Michio Honda, *University of Edinburgh*

Sanidhya Kashyap, *EPFL*

Bingzhe Li, *The University of Texas at Dallas*

Jing Liu, *Microsoft Research*

Sam H. Noh, *Virginia Tech*

Cheng Tan, *Northeastern University*

Hong Zhang, *University of Waterloo*