==========

Call for Papers

MTSA'23: Workshop on Memory Technologies, Sysstems, and Applications
held in conjunction with SC23: The International Conference on High Performance Computing, Networking, Storage and Analysis

Time/Date: 1:30PM - 5:00PM, U.S. Mountain Standard Time, Monday, November 13, 20223
Location: Colorado Convention Center, Room TBD

https://passlab.github.io/mtsa/mtsa2023

Submission Deadline: August 04, 2023

==========

The growing disparity between compute and memory speed, known as the memory wall problem, has been one of the most critical and long-standing challenges in the computing industry. The prevalence of heterogeneous computing, the ongoing expansion of the memory hierarchy, and the advent of disaggregated architectures have considerably expanded the scope of this problem. Computer architecture, operating systems, storage systems, performance models, tools, and applications themselves are being enhanced or even redesigned to address the performance, programmability, and energy efficiency challenges of the increasingly complex and heterogeneous memory systems. Exploring the intersection of these research areas will enable cohesive and synergistic development and collaboration on the future of memory technologies, systems, and applications.

As the successor to the previously successful Memory-Centric High Performance Computing (MCHPC) workshop, MTSA'23: Workshop on Memory Technologies, Systems, and Applications aims to bring together researchers from industry, government labs, and academia concerned with the challenges of designing, prototyping, efficiently using, and evaluating current and emergent memory systems. The term performance for memory systems is general, including latency, bandwidth, power consumption and reliability from the aspect of hardware memory technologies to what it is manifested in the application performance. The topics of interest for the MTSA workshop include, but are not limited to:

- Software, hardware, and co-design approaches that ease the adoption and optimize the use of processing-in-memory and near-memory computing technologies.
- Evaluation, characterization, performance analysis, and use cases of emerging memory technologies, including disaggregated memory, cache-coherent interconnects e.g., CXL, heterogeneous memory, etc.
- Programming interfaces or language extensions that improve the programmability of using emerging memory technologies and systems, heterogeneous memory systems and multi-dimensional data, and unified memory systems.

- Compiler, runtime, and system techniques for optimizing data layout and placement, page migration, coherence and consistency enforcement, latency hiding and improving bandwidth utilization and energy consumption of heterogeneous memory systems.
- Enhancement or new development for operating systems, storage and file systems, and I/O systems that address challenges of existing and emerging memory technologies, heterogeneous memory systems, and the blurred boundary between memory and storage.
- Tools, modeling, evaluation, and case study of memory system behavior and application performance that reveal the limitation and characteristics of existing memory systems.
- Application development and optimization for new memory architecture and technologies.
- This year, there is a particular interest in papers that include or focus on energy and power efficiency in memory technologies.

Important Dates

- Submission Deadline **Extended** – August ~~04~~ **16**, 2023
- Notifications -- September 6, 2023
- Camera Ready Papers – September 29, 2023, including Copyright Form
- Workshop -- November 13, 2023

Submission

Authors are invited to submit manuscripts in English structured as 10 2-column pages (U.S. letter – 8.5”x11”), excluding the bibliography, using the ACM Proceedings Template (https://www.acm.org/publications/proceedings-template). Latex users, please use the “**sigconf**” option (use of the “**review**” option is recommended but not required). Word authors can use the “Interim Layout”. The workshop encourages submitters to include reproducibility information, using Reproducibility Initiative for SC'23 Technical Papers as a guideline. All manuscripts will be peer-reviewed and judged on correctness, originality, technical strength, and significance, quality of presentation, and interest and relevance to the workshop attendees. There will be one round of single-blind reviews with potential for acceptance with minor revisions before the camera ready submission.

Submitted papers must represent original unpublished research that is not currently under review for any other conference or journal. Papers not following these guidelines will be rejected without review and further action may be taken, including (but not limited to) notifications sent to the heads of the institutions of the authors and sponsors of the conference. Submissions received after the due date, exceeding the length limit, or not appropriately structured may also not be considered. At least one author of an accepted paper must register for and attend the workshop. Authors may contact the workshop organizers for more information. Papers should be submitted electronically at: https://submissions.supercomputing.org/, choose "SC23 Workshop: MTSA'23: Workshop on Memory Technologies, Systems, and Applications". The final papers will be published in the SC Workshops Proceedings volume. Copyright procedures will follow those of the SC Workshops Proceedings.

Organizers

- Tyler Allen (University of North Carolina at Charlotte, t.allen@charlotte.edu)
- Ivy B. Peng (KTH Royal Institute of Technology, ivybopeng@kth.se)
- Ron Brightwell (Sandia National Laboratories, rbbrigh@sandia.gov)
- Maya B Gokhale (Lawrence Livermore National Laboratory, gokhale2@llnl.gov)

Program Committee

- Tyler Allen, University of North Carolina at Charlotte
- Ron Brightwell, Sandia National Laboratories
- Tom Deakin, University of Bristol
- Maya Gokhale, Lawrence Livermore National Laboratory
- Kyle Hale, Illinois Institute of Technology
- Alice Koniges, University of Hawai'i, Maui High Performance Computing Center
- Seyong Lee, Oak Ridge National Laboratory
- Dong Li, University of California, Merced
- Chunhua Liao, Lawrence Livermore National Laboratory
- Stephen L. Olivier, Sandia National Laboratories
- Ivy Peng, KTH Royal Institute of Technology
- Jie Ren, College of William and Mary
- Yonghong Yan, University of North Carolina at Charlotte
- Onkar Patil, IBM Corporation

Steering Committee

- Xian-He Sun (Illinois Institute of Technology)
- Yonghong Yan (University of North Carolina at Charlotte)