
# The 2012 ACM SIGPLAN Workshop on Memory Systems Performance & Correctness

*Co-located with PLDI 2012*

**General Chair**
Lixin Zhang, ICT, Chinese Academy of Sciences

**Program Chair**
Onur Mutlu, Carnegie Mellon University

**Program Committee**
Rajeev Balasubramonian, Utah
Emery Berger, UMass Amherst
John Carter, IBM Research
Jichuan Chang, HP Labs
Andrew Chien, Chicago
Trishul Chilimbi, Microsoft Research
Cliff Click, Azul Systems
Eiman Ebrahimi, UT-Austin
Mattan Erez, UT-Austin
Eugene Gorbatov, Intel
Nikos Hardavellas, Northwestern
Maurice Herlihy, Brown
Brian Hirano, Oracle
Hillery Hunter, IBM Research
Jim Larus, Microsoft Research
Alvy Lebeck, Duke
Thomas Moscibroda, Microsoft Research
Naveen Muralimanohar, HP Labs
Satish Narayanasamy, Michigan
Erez Petrank, Technion
Moinuddin Qureshi, Georgia Tech
Xipeng Shen, William and Mary
Osman Unsal, BSC
Chris Wilkerson, Intel
Ben Zorn, Microsoft Research

**Steering Committee**
Emery Berger, UMass Amherst
Brad Chen, Google
Trishul Chilimbi, Microsoft Research
Chen Ding, Rochester
Madan Musuvathi, Microsoft Research
Xipeng Shen, College of William & Mary
Jeffrey S. Vetter, Oak Ridge Nat'l Lab & Georgia Tech
Ben Zorn, Microsoft Research

**Web and Submissions Chairs**
Chris Fallin, Carnegie Mellon University
Vivek Seshadri, Carnegie Mellon University

## Important Dates

Papers Due: ~~**Mar 12, 2012, 11:59pm EDT**~~
**Extended**: **Mar 22, 2012, 11:59pm EDT**
Acceptance Notification: **Apr 23, 2012**
Final Papers Due**: May 1, 2012**

Memory continues to be a major bottleneck in almost all computing systems. It is becoming more so as more cores and agents are sharing parts of the memory system and as applications that run on the cores are becoming increasingly data intensive. Continuing the tradition of six previous successful incarnations, MSPC 2012 will provide a forum for publishing and discussing *all aspects of memory performance and correctness* on a variety of systems (multi-core, desktop, embedded, server/cloud, high-performance computing, sensor, etc) and related software and hardware innovations at various levels of the technology stack. We invite new submissions that tackle issues in memory system performance, efficiency, correctness, and dependability in both hardware and software layers. Example areas of interest include but are not limited to the following:

- Hardware, software, and hybrid techniques for better memory performance, correctness, reliability, efficiency
- Memory hierarchy design for chip multiprocessors (CMPs)
- Emerging memory technologies (e.g., Phase Change Memory, MRAM)
- Characterization and analysis of memory systems performance
- Insightful experimental evaluation and analysis of memory-intensive workloads
- Static and dynamic techniques for understanding and improving memory performance and efficiency
- Managed memory and garbage collection optimizations
- Hardware and software techniques for ensuring memory safety and detecting memory-related bugs
- Hardware and software memory models and their impact on programmability and performance
- Memory system issues in accelerator-based computing (e.g., GPGPU)
- Memory system issues in embedded computers and tiny devices
- Prefetching, compression, latency tolerance techniques for memory
- Memory power and energy management techniques
- Memory reliability management techniques

Software, hardware, and hybrid approaches are encouraged. In addition, we solicit papers from practitioners describing problems and experiences with memory performance and correctness in specific application domains.

**Submission Guidelines** We encourage the submission of not-fully-polished but provocative short papers (6—8 pages; 8 pages maximum) or position abstracts (1-2 pages; 2 pages maximum). Paper submissions should use standard ACM SIGPLAN conference format (10pt), available at http://www.sigplan.org/authorInformation.htm. Copies of accepted papers will be made available at the workshop and published in the ACM digital library. Submitted papers must not be simultaneously under review for any other conference or journal, and authors should point out any substantial overlap with their previously published or currently submitted work.

**http://safari.ece.cmu.edu/MSPC2012/**