

Onur Mutlu

Professor, Department of Computer Science, ETH Zürich, 8092 Zürich, Switzerland

<https://people.inf.ethz.ch/omutlu/>

onur.mutlu@inf.ethz.ch

One-Page Biography

Onur Mutlu was born in 1978 in Turkey. He earned his dual Bachelor of Science degrees in Computer Engineering and Psychology from the University of Michigan, and his Master's and PhD in Computer Engineering from the University of Texas at Austin. During his PhD, he invented the Efficient Runahead Execution technique for tolerating long memory latencies in modern processors, variants of which are implemented in commercial microprocessors (e.g., IBM POWER6, Sun/Oracle Rock). His dissertation's influence on academic and industrial research on the topic still continues. Between 2001-2005, he worked in the cutting-edge microprocessor design and research groups at Intel Corporation and Advanced Micro Devices, where he adapted his Efficient Runahead Execution [Mutlu+, HPCA'03, ISCA'05] and Feedback-Directed Data Prefetching [Srinath+, HPCA'07] mechanisms for future systems.

In 2006, he joined Microsoft Research Redmond to start the Computer Architecture Group. During his 2.5-year tenure there, he conducted seminal research into memory systems for multi-core processors. His research into memory controllers spearheaded a fresh research area in computer architecture, which continues to thrive. With his collaborator Thomas Moscibroda, he discovered that existing multi-core memory controllers were vulnerable to denial service attacks [Moscibroda+, USENIX Security'07]. In a series of works published at top venues in computer architecture since 2007, he co-devised new memory control algorithms that provide high system performance, fairness, and quality of service. These techniques turned memory controllers into a center of attention in computer architecture. In particular, his MICRO'07 paper on "Stall-Time Fair Memory Access Scheduling" and ISCA'08 paper on "Parallelism-Aware Batch Scheduling" have greatly influenced academic research on memory controllers by exposing new problems and new solutions. Variants of his Parallelism-Aware Batch Scheduler [Mutlu+, ISCA'08] are implemented in some memory controllers designed by Samsung.

In 2009, he moved to Carnegie Mellon University as an Assistant Professor. He continued attacking the "memory problem" from all angles, with impact on both major academic research directions and commercial products. He advocated and analyzed the use of Phase Change Memory technology as part of the main memory of a computing system [Lee+, ISCA'09], starting a new area in computer architecture: how to enable/exploit emerging non-volatile memory technologies. This work is considered a precursor of 3D XPoint by Intel and other similar technologies that are being designed for main memory.

His work on flash memory reliability, starting with his paper entitled "Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis" [Cai+, DATE'12] and continuing most recently to [Luo+, JSAC'16][Cai+, HPCA'17][Cai+, Proc IEEE'17][Luo+, HPCA'18], uncovered new error mechanisms, provided precious experimental data, and greatly improved flash memory lifetime and reliability. These works are regarded as prime resources for educating engineers in leading flash memory companies, Samsung, SK Hynix, Seagate, Sandisk, Micron, Toshiba, IBM, Intel.

His work on key DRAM challenges (refresh, latency, energy, variability, reliability, scaling) ignited substantial academic research and influenced industry directions. For example, his Subarray-Level Parallelism work [Kim+, ISCA'12] is advocated by Intel and Samsung to be in future DRAM standards [Kang+, Memory Forum'14]. He pioneered recent research on solving the DRAM refresh problem [Liu+, ISCA'12]. His research showed that refresh is the key scaling limiter of the DRAM technology, experimentally discovered novel problems related to it, provided precious experimental data from modern DRAM devices (in another seminal work [Liu+, ISCA'13]), and developed practical solutions to the refresh problem (e.g., [Qureshi+, DSN'15][Patel+, ISCA'17]). His work experimentally demonstrated, analyzed, and provided architectural solutions for various other critical DRAM issues (e.g., latency [Lee+, HPCA'13, HPCA'15, SIGMETRICS'17], variability [Chang+, SIGMETRICS'16], energy [David+, ICAC'11][Chang+, SIGMETRICS'17], reliability [Meza+, SIGMETRICS'15]) by analyzing modern DRAM chips using real FPGA-based experimental platforms [Hasan+, HPCA'17], providing precious data available nowhere else.

In 2014, his group discovered the RowHammer problem [Kim+, ISCA'14], a failure mechanism affecting most real DRAM chips. This work shook the fundamentals of systems security: Google and others demonstrated attacks that exploit RowHammer to take over an otherwise-secure system. RowHammer is the first example of a hardware failure mechanism that causes a practical and widespread system security vulnerability. It continues to have widespread impact on systems, security, software, and hardware communities, both academic and industrial [Mutlu, DATE'17]: e.g., it caused a new Hammer Test to be included in standard memory test programs, and Apple cited the work [Kim+, ISCA'14] in its critical security release that introduced a hardware patch to mitigate RowHammer.

For his contributions, he received the endowed Strecker Early Career Professorship at Carnegie Mellon University (2013), the first IEEE Computer Society Young Computer Architect Award (2011), the first Intel Early Career Faculty Award (2012), and the Ladd Research Award from CMU (2012). 11 of his papers were chosen as Top Picks by IEEE Micro, and more than 15 received Best Paper recognitions. He received numerous Faculty Awards from Facebook, Google, HP, IBM, Intel, Microsoft, VMware, and NSF. He was named an ACM Fellow in 2017, "for contributions to computer architecture research, especially in memory systems."

He spent significant time in industry, enabling technology transfer of his ideas. He was with Intel (Summer'01-'03 and '12) and AMD (Summer'04-'05), researching novel microprocessor designs, and with VMware (Spring-Summer'16) and Google (Summer'16), exploring new memory and system architectures. He continues to collaborate with industry on both fundamental and cutting-edge technology issues.

He has advocated open sharing of teaching and research artifacts to democratize education/research worldwide. His computer architecture course videos/materials and his group's research artifacts are freely available online. They are used by many educators, researchers and practitioners, including companies like Google and Intel.

In September 2015, he was appointed Full Professor of Computer Science at ETH Zürich, where he started in May 2016. He continues to thrive on research, education, and service for the community. More information is available on his webpage: <https://people.inf.ethz.ch/omutlu/>.

Onur Mutlu

Professor, Department of Computer Science, ETH Zürich, 8092 Zürich, Switzerland

<https://people.inf.ethz.ch/omutlu/>

onur.mutlu@inf.ethz.ch

Research, Teaching and Consulting Interests

- My main research is in computer architecture, systems, security, and bioinformatics. My work spans and stretches the boundaries between applications, systems, languages, system software, compilers, and hardware. My research tackles many issues in high performance, energy efficiency, hardware security, fault tolerance, predictable systems, dependable systems, and hardware/software cooperation. I am especially excited about novel, fundamentally-secure and fundamentally-efficient computation, communication and memory/storage paradigms, applied to emerging systems, technologies, and bioinformatics/medical applications. I am also excited about system design for bioinformatics and biologically inspired computing paradigms. Significant expertise areas include the following:
- Memory and storage systems (DRAM, flash, NVM, emerging). Memory controllers. Scalable, QoS-aware, latency-tolerant systems.
- Processor design. Multi/many-core systems. Hardware/software interaction. OS/PL/hardware interaction. Interconnects.
- Hardware security and safety. Reliability. Fault-tolerant and bug-tolerant architectures. Dependable systems.
- Bioinformatics algorithms and architectures. Genome sequence analysis and assembly techniques.

Education

University of Texas at Austin

September 2000 - August 2006

Ph.D., Computer Engineering, August 2006

Dissertation Title: *Efficient Runahead Execution Processors*

Nominated for the ACM Doctoral Dissertation Award by UT-Austin

M.S.E., Computer Engineering, May 2002

University of Michigan, Ann Arbor

September 1997 - August 2000

B.S.E., *summa cum laude*, Computer Engineering, August 2000

B.S., *with highest distinction*, Psychology, August 2000

Professional Experience

ETH Zürich, Dept. of Computer Science (D-INFK), Full Professor, *September 2015 - Present*

Carnegie Mellon University, Dept. of Electrical and Computer Engineering, Adjunct Professor, *June 2016 - Present*

Bilkent University, Dept. of Computer Engineering, Adjunct Professor, *November 2015 - Present*

Google, Visiting Research Scientist, *May 2016 - September 2016*

VMware, Scholar in Residence, *February 2016 - September 2016*

Carnegie Mellon University, Dept. of ECE, Dr. William D. and Nancy W. Strecker Endowed Early Career Professor, *January 2013 - June 2016*

Intel Corporation, Visiting Researcher, *June 2012 - August 2012*

Carnegie Mellon University, Dept. of Electrical and Computer Engineering, Assistant Professor, *January 2009 - January 2013*

Carnegie Mellon University, Dept. of Computer Science, Courtesy Professor, *January 2009 - Present*

Microsoft Research, Computer Architecture Group (Redmond, WA), Researcher, *August 2006 - January 2009*

University of Texas at Austin, Dept. of Electrical and Computer Engineering, Research Fellow, *August 2007 - January 2009*

University of Texas at Austin, Dept. of Electrical and Computer Engineering, Research and Teaching Assistant, *August 2000 - August 2006*

Advanced Micro Devices, Architecture/Performance Modeling Group (Sunnyvale, CA), Co-op Engineer, *May - August 2005*

Advanced Micro Devices, Architecture/Performance Modeling Group (Sunnyvale, CA), Co-op Engineer, *May - August 2004*

Intel Corporation, Desktop Platforms Group (Hillsboro, OR), Graduate Technical Intern, *May - August 2003*

Intel Corporation, Microprocessor Research Labs (Hillsboro, OR), Graduate Technical Intern, *May - August 2002*

Intel Corporation, Desktop Platforms Group (Hillsboro, OR), Graduate Technical Intern, *May - August 2001*

Honors and Awards

- ACM Fellow, “for contributions to computer architecture research, especially in memory systems,” 2017
- Google Faculty Research Award, 2015, 2016
- Microsoft Research Software Engineering Innovation Foundation Award, 2014
- Dr. William D. and Nancy W. Strecker Early Career Professorship (at Carnegie Mellon University), January 2013
- IBM Faculty Partnership Award, 2012, 2013
- Intel Early Career Faculty Honor Program Award, 2012
- Carnegie Mellon University College of Engineering George Tallman Ladd Research Award, 2012
- IEEE Computer Society Technical Committee on Computer Architecture Young Computer Architect Award, 2011
- Hewlett-Packard Laboratories Innovation Research Program Award, 2012
- Nvidia CUDA Center of Excellence Award, 2012, 2013 (with multiple CMU Faculty members)
- More than 25 keynote, plenary and invited talks at many conferences and workshops, including AACBB 2018, ARC 2017, MST 2017, ISMM 2016, DAC 2016, RSP 2016, ARM Research Summit 2016, FABULOUS 2016, MST 2016, NoCArc 2015, CASS 2015, SBAC-PAD 2015, SAMOS 2015, ISC 2015, ASAP 2014, DaMoN 2014, Summer Supercomputing Academy of Moscow State University 2014, IMW 2013, ISMM 2011.

- Best paper award at DFRWS-EU 2017 (Digital Forensics Conference)
Improving the Reliability of Chip-Off Forensic Analysis of NAND Flash Memory Devices
- Best paper award at RTAS 2014
Bounding Memory Interference Delay in COTS-based Multi-Core Systems
- Best paper award at ICCD 2012 (Computer Systems and Applications Track)
Row Buffer Locality Aware Caching Policies for Hybrid Memories
- Best paper award at ASPLOS 2010
Fairness via Source Throttling: A Configurable and High-Performance Fairness Substrate for Multi-Core Memory Systems
- Best paper award at VTS 2010 (awarded in 2011)
Concurrent Autonomous Self-Test for Uncore Components in System-on-Chips
- One paper selected as Honorable Mention for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2015
A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing
- One paper (of 12 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2011
Kilo-NOC: A Heterogeneous Network-on-Chip Architecture for Scalability and Service Guarantees
- Three papers (of 11 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2010
Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior
Aergia: Exploiting Packet Latency Slack in On-Chip Networks
Data Marshaling for Multi-core Architectures
- Two papers (of 13 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2009
Accelerating Critical Section Execution with Asymmetric Multi-Core Architectures
Architecting Phase Change Memory as a Scalable DRAM Alternative
- One paper (of 12 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2008
Parallelism-Aware Batch Scheduling: Enabling High-Performance and Fair Memory Controllers
- One paper (of 11 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2006
Diverge-Merge Processor (DMP): Generalized and Energy-Efficient Dynamic Predication
- Two papers (of 13 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2005
Efficient Runahead Execution: Power-efficient Memory Latency Tolerance
Wish Branches: Enabling Adaptive and Aggressive Predicated Execution
- One paper (of 15 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2003
Runahead Execution: An Effective Alternative to Large Instruction Windows
- One paper selected for CACM's "Research Highlights," 2009
Architecting Phase Change Memory as a Scalable DRAM Alternative
- Best Paper Session at MICRO 2016: *Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads*
- Best Paper Session (Runner-Up) at HPCA 2015: *Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery*
- Best Paper Session at HPCA 2014: *Improving Cache Performance by Exploiting Read-Write Disparity*
- Best Paper Session at HPCA 2010: *ATLAS: A Scalable and High-Performance Scheduling Algorithm for Multiple Memory Controllers*
- Best Paper Session at HPCA 2009: *Techniques for Bandwidth-Efficient Prefetching of Linked Data Structures in Hybrid Prefetching Systems*
- Best paper award nominations at MICRO 2016, NOCS 2015, NOCS 2012, HPCA 2015, HPCA 2014, HPCA 2010, HPCA 2009, HPCA 2007, MICRO 2006, and MICRO 2005 conferences
- Selected to the ISCA and MICRO Halls of Fame, 2009
- Selected to the HPCA Hall of Fame, 2010
- NSF CAREER Award, 2010 (*QoS-Aware, High-Performance, and Scalable Many-Core Memory Systems*)
- Microsoft Gold Star Award, 2008
- Distinguished Lecture at Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, 28 June 2012.
- PhD Dissertation nominated by UT-Austin for the ACM Doctoral Dissertation Award, 2006
- University Co-op/George H. Mitchell Award for Excellence in Graduate Research (Awarded to 6 out of 271 nominees at UT-Austin), 2005
- Intel Foundation Ph.D. Fellowship, 2004
- University of Texas Graduate School Continuing Fellowship, 2003
- University of Michigan EECS Dept. Summer Research Fellowship, 1999, 2000
- University of Michigan EECS Dept. William Harvey Seeley Award (money award given to the top undergraduate junior), 1999
- University of Michigan Branstrom Freshman Prize, 1998

Publications

(Please visit <https://people.inf.ethz.ch/omutlu/projects.htm> for electronic copies.)

(Please visit <http://scholar.google.com/citations?user=7XyGUGkAAAAJ&hl=en> for citations.)

(H-index: 71, I10-index: 197, Citation count: 17692 – according to Google Scholar, as of April 1, 2018)

Book Chapters

1. Nandita Vijaykumar, Kevin Hsieh, Gennady Pekhimenko, Samira Khan, Saugata Ghose, Ashish Shreshtha, Adwait Jog, Phillip P. Gibbons, Onur Mutlu, "Decoupling the Programming Model from Resource Management in Throughput Processors," in *Many-Core Computing: Hardware and Software*, IET, to appear in 2018.

2. Saugata Ghose, Kevin Hsieh, Amirali Boroumand, Rachata Ausavarungnirun, Onur Mutlu, “The Processing-in-Memory Paradigm: Mechanisms to Enable Adoption,” in *Beyond-CMOS Technologies for Next Generation Computer Design*, Springer, to appear in 2018.
3. Yu Cai, Saugata Ghose, Erich F. Haratsch, Yixin Luo, Onur Mutlu, “Reliability Issues in Flash-Memory-Based Solid-State Drives: Experimental Analysis, Mitigation, Recovery,” in *Inside Solid State Drives (SSDs)*, Springer, to appear in 2018.
4. Vivek Seshadri, Onur Mutlu, “Simple Operations in Memory to Reduce Data Movement,” *Advances in Computers*, Elsevier, 2017.
5. Nandita Vijaykumar, Gennady Pekhimenko, Adwait Jog, Abhishek Bhowmick, Rachata Ausavarungnirun, Chita Das, Mahmut Kandemir, Todd C. Mowry, Onur Mutlu, “A Framework for Accelerating Bottlenecks in GPU Execution with Assist Warps,” in *Advances in GPU Computing and Practice*, Elsevier, 2016.
6. Onur Mutlu, “Main Memory Scaling: Challenges and Solution Directions,” in *More Than Moore Technologies for Next Generation Computer Design*, Springer, January 2015.
7. Yoongu Kim, Onur Mutlu, “Memory Systems,” in *Computing Handbook: Computer Science and Software Engineering*, CRC Press, April 2014.
8. Chris Fallin, Greg Nazario, Xiangyao Yu, Kevin Chang, Rachata Ausavarungnirun, Onur Mutlu, “Bufferless and Minimally-Buffered Deflection Routing,” in *Routing Algorithms in Networks-on-Chip*, Springer, 2014.
9. M. Aater Suleman, Onur Mutlu, “Accelerating Critical Section Execution with Multi-Core Architectures,” in *Multicore Technology: Architecture, Reconfiguration, and Modeling*, CRC Press, July 2013.

Refereed Conference (and Major Workshop) Publications

10. Arash Tavakkol, Mohammad Sadrosadati, Saugata Ghose, Jeremie Kim, Yixin Luo, Yaohua Wang, Nika Mansouri Ghiasi, Lois Orosa, Juan G. Luna, Onur Mutlu, “FLIN: Enabling Fairness and Enhancing Performance in Modern NVMe Solid State Drives,” *Proceedings of the 45th International Symposium on Computer Architecture (ISCA)*, Los Angeles, CA, USA, June 2018.
11. Nandita Vijaykumar, Eiman Ebrahimi, Kevin Hsieh, Phillip B. Gibbons, Onur Mutlu, “The Locality Descriptor: A Holistic Abstraction to Exploit Data Locality in GPUs,” *Proceedings of the 45th International Symposium on Computer Architecture (ISCA)*, Los Angeles, CA, USA, June 2018.
12. Nandita Vijaykumar, Abhilasha Jain, Diptesh Majumdar, Kevin Hsieh, Gennady Pekhimenko, Eiman Ebrahimi, Nastaran Hajinazar, Phillip B. Gibbons, Onur Mutlu, “A Case for Richer Cross-layer Abstractions: Bridging the Semantic Gap to Enhance Memory Optimization,” *Proceedings of the 45th International Symposium on Computer Architecture (ISCA)*, Los Angeles, CA, USA, June 2018.
13. Feng Zhang, Jidong Zhai, Xipeng Shen, Onur Mutlu, Wenguang Chen, “Zwift: A Programming Framework for High Performance Text Analytics on Compressed Data,” *Proceedings of the International Conference on Supercomputing (ICS)*, Beijing, China, June 2018.
14. Anup Das, Hasan Hassan, Onur Mutlu, “VRL-DRAM: Improving DRAM Performance via Variable Refresh Latency,” *Proceedings of the 55th Design Automation Conference (DAC)*, San Francisco, CA, USA, June 2018.
15. Saba Ahmadian, Onur Mutlu, Hossein Asadi, “ECI-Cache: A High-Endurance and Cost-Efficient I/O Caching Scheme for Virtualized Platforms,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Irvine, CA, USA, June 2018.
16. Rachata Ausavarungnirun, Vance Miller, Joshua Landgraf, Saugata Ghose, Jayneel Gandhi, Adwait Jog, Christopher J. Rossbach, Onur Mutlu, “MASK: Redesigning the GPU Memory Hierarchy to Support Multi-Application Concurrency,” *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Williamsburg, VA, USA March 2018.
17. Maciej Besta, Syed Minhaj Hassan, Sudhakar Yalamanchili, Rachata Ausavarungnirun, Onur Mutlu, Torsten Hoefler, “Slim NoC: A Low-Diameter On-Chip Network Topology for High Energy Efficiency and Scalability,” *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Williamsburg, VA, USA March 2018.
18. Amirali Boroumand, Saugata Ghose, Youngsok Kim, Rachata Ausavarungnirun, Eric Shiu, Rahul Thakur, Daehyun Kim, Aki Kuusela, Allan Knies, Parthasarathy Ranganathan, Onur Mutlu, “Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks,” *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Williamsburg, VA, USA March 2018.
19. Amir M. Rahmani, Bryan Donyanavard, Tiago Muck, Kasra Moazzemi, Axel Jantsch, Onur Mutlu, Nikil Dutt, “SPECTR: Formal Supervisory Control and Coordination for Many-core Systems Resource Management,” *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Williamsburg, VA, USA March 2018.
20. Mohammad Sadrosadati, Amirhossein Mirhosseini, Seyed Borna Ehsani, Hamid Sarbazi-Azad, Mario Drumond, Babak Falsafi, Rachata Ausavarungnirun, Onur Mutlu, “LTRF: Enabling High-Capacity Register Files for GPUs via Hardware/Software Cooperative Register Prefetching,” *Proceedings of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Williamsburg, VA, USA March 2018.
21. Arash Tavakkol, Juan Gomez-Luna, Mohammad Sadrosadati, Saugata Ghose, Onur Mutlu, “MQSim: A Framework for Enabling Realistic Studies of Modern Multi-Queue SSD Devices,” *Proceedings of the 16th USENIX Conference on File and Storage Technologies (FAST)*, Oakland, CA, USA, February 2018.
22. Jeremie S. Kim, Minesh Patel, Hasan Hassan, Onur Mutlu, “The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern DRAM Devices,” *Proceedings of the 24th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Vienna, Austria, February 2018.

23. Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, Onur Mutlu, “HeatWatch: Improving 3D NAND Flash Memory Device Reliability by Exploiting Self-Recovery and Temperature-Awareness,” *Proceedings of the 24th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Vienna, Austria, February 2018.
24. Vivek Seshadri, Kevin Hsieh, Amirali Boroumand, Donghyuk Lee, Michael A. Kozuch, Onur Mutlu, Phillip B. Gibbons, Todd C. Mowry, “Ambit: In-Memory Accelerator for Bulk Bitwise Operations Using Commodity DRAM Technology,” *Proceedings of the 50th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Boston, MA, USA, October 2017.
25. Xiangyao Yu, Christopher J. Hughes, Nadathur Satish, Onur Mutlu, Srinivas Devadas, “Banshee: Bandwidth-Efficient DRAM Caching via Software/Hardware Cooperation,” *Proceedings of the 50th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Boston, MA, USA, October 2017.
26. Samira Khan, Chris Wilkerson, Zhe Wang, Alaa R. Alameldeen, Donghyuk Lee, Onur Mutlu, “Detecting and Mitigating Data-Dependent DRAM Failures by Exploiting Current Memory Content,” *Proceedings of the 50th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Boston, MA, USA, October 2017.
27. Rachata Ausavarungnirun, Joshua Landgraf, Vance Miller, Saugata Ghose, Jayneel Gandhi, Christopher J. Rossbach, Onur Mutlu, “Mosaic: A GPU Memory Manager with Application-Transparent Support for Multiple Page Sizes,” *Proceedings of the 50th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Boston, MA, USA, October 2017.
28. Yang Li, Saugata Ghose, Jongmoo Choi, Jin Sun, Hui Wang, Onur Mutlu, “Utility-Based Hybrid Memory Management,” *Proceedings of the 19th IEEE Cluster Conference (CLUSTER)*, Honolulu, Hawaii, USA, September 2017.
29. Zhiyu Liu, Irina Calciu, Maurice Herlihy, Onur Mutlu, “Concurrent Data Structures for Near-Memory Computing,” *Proceedings of the 29th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, Washington, DC, USA, July 2017.
30. Minesh Patel, Jeremie Kim, Onur Mutlu, “The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions,” *Proceedings of the 44th International Symposium on Computer Architecture (ISCA)*, Toronto, Canada, June 2017.
31. Donghyuk Lee, Samira Khan, Lavanya Subramanian, Rachata Ausavarungnirun, Gennady Pekhimenko, Vivek Seshadri, Saugata Ghose, Onur Mutlu, “Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Urbana-Champaign, IL, USA, June 2017.
32. Kevin Chang, A. Giray Yaglikci, Saugata Ghose, Abhijith Kashyap, Hasan Hassan, Aditya Agrawal, Niladrish Chatterjee, Donghyuk Lee, Mike O’Connor, Onur Mutlu, “Understanding Reduced-Voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Urbana-Champaign, IL, USA, June 2017.
33. Xiyue Xiang, Saugata Ghose, Onur Mutlu, Nian-Feng Tzeng, Lu Peng, Wentao Shi, “Carpool: A Bufferless On-Chip Network with Adaptive Multicast and Hotspot Alleviation,” *Proceedings of the International Conference on Supercomputing (ICS)*, Chicago, IL, USA, June 2017.
34. Kaan Kara, Dan Alistarh, Ce Zhang, Onur Mutlu, Gustavo Alonso, “FPGA-accelerated Dense Linear Machine Learning: A Precision-Convergence Trade-off,” *Proceedings of the 25th International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, Napa, CA, USA, April 2017.
35. Kevin Hsieh, Aaron Harlap, Nandita Vijaykumar, Dimitris Konomis, Gregory R. Ganger, Phillip B. Gibbons, Onur Mutlu, “Gaia: Geo-Distributed Machine Learning Approaching LAN Speeds,” *Proceedings of the 14th USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, Boston, MA, USA, March 2017.
36. Onur Mutlu, “The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser,” *Proceedings of the Design, Automation, and Test in Europe Conference (DATE)*, Lausanne, Switzerland, March 2017. **Invited Paper and Presentation.**
37. Aya Fukami, Saugata Ghose, Yixin Luo, Yu Cai, Onur Mutlu, “Improving the Reliability of Chip-Off Forensic Analysis of NAND Flash Memory Devices,” *Proceedings of the Digital Forensics Conference (DFRWS EU)*, Lake Constance, Germany, March 2017. **Best Paper Award.**
38. Hasan Hassan, Nandita Vijaykumar, Samira Khan, Saugata Ghose, Kevin Chang, Gennady Pekhimenko, Donghyuk Lee, Oguz Ergin, Onur Mutlu, “SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies,” *Proceedings of the 23rd IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Austin, TX, USA, February 2017.
39. Yu Cai, Saugata Ghose, Yixin Luo, Ken Mai, Onur Mutlu, Erich F. Haratsch, “Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques,” *Proceedings of the 23rd IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Austin, TX, USA, February 2017.
40. Khanh Nguyen, Lu Fang, Guoqing Xu, Brian Demsky, Shan Lu, Sanazsadat Alamian, Onur Mutlu, “Yak: A High-Performance Big-Data-Friendly Garbage Collector,” *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI)*, Savannah, GA, USA, November 2016.
41. Himanshu Chauhan, Irina Calciu, Vijay Chidambaram, Eric Schkufza, Onur Mutlu, Pratap Subrahmanyam, “NVMove: Helping Programmers Move to Byte-Based Persistence,” *4th Workshop on Interactions of NVM/Flash with Operating Systems and Workloads (INFLOW)*, Savannah, GA, USA, November 2016.

42. Nandita Vijaykumar, Kevin Hsieh, Gennady Pekhimenko, Samira Khan, Saugata Ghose, Ashish Shreshtha, Adwait Jog, Phillip P. Gibbons, Onur Mutlu, “Zorua: A Holistic Approach to Resource Virtualization in GPUs,” *Proceedings of the 49th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Taipei, Taiwan, October 2016.
43. Milad Hashemi, Onur Mutlu, Yale N. Patt, “Continuous Runahead: Transparent Hardware Acceleration for Memory Intensive Workloads,” *Proceedings of the 49th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Taipei, Taiwan, October 2016. **Best Paper Session.**
44. Kevin Hsieh, Samira Khan, Nandita Vijaykumar, Kevin K. Chang, Amirali Boroumand, Saugata Ghose, Onur Mutlu, “Accelerating Pointer Chasing in 3D-Stacked Memory: Challenges, Mechanisms, Evaluation,” *Proceedings of the 34th IEEE International Conference on Computer Design (ICCD)*, Phoenix, AZ, USA, October 2016.
45. Xiyue Xiang, Saugata Ghose, Onur Mutlu, Nian-Feng Tzeng, “A Model for Application Slowdown Estimation in On-Chip Networks and Its Use for Improving System Fairness and Performance,” *Proceedings of the 34th IEEE International Conference on Computer Design (ICCD)*, Phoenix, AZ, USA, October 2016.
46. Ashutosh Pattnaik, Xulong Tang, Adwait Jog, Onur Kayiran, Asit K. Mishra, Mahmut T. Kandemir, Onur Mutlu, Chita R. Das, “Scheduling Techniques for GPU Architectures with Processing-In-Memory Capabilities,” *Proceedings of the 25th ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Haifa, Israel, September 2016.
47. Onur Kayiran, Adwait Jog, Ashutosh Pattnaik, Rachata Ausavarungnirun, Xulong Tang, Mahmut T. Kandemir, Gabriel H. Loh, Onur Mutlu, Chita R. Das, “ μ C-States: Fine-grained GPU Datapath Power Management,” *Proceedings of the 25th ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Haifa, Israel, September 2016.
48. Samira Khan, Donghyuk Lee, Onur Mutlu, “PARBOR: An Efficient System-Level Technique to Detect Data-Dependent Failures in DRAM,” *Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Toulouse, France, June 2016.
49. Kevin Hsieh, Eiman Ebrahimi, Gwangsun Kim, Niladri Chatterjee, Mike O’Connor, Nandita Vijaykumar, Onur Mutlu, Stephen W. Keckler, “Transparent Offloading and Mapping (TOM): Enabling Programmer-Transparent Near-Data Processing in GPU Systems,” *Proceedings of the 43rd International Symposium on Computer Architecture (ISCA)*, Seoul, South Korea, June 2016.
50. Milad Hashemi, Khubaib, Eiman Ebrahimi, Onur Mutlu, Yale N. Patt, “Accelerating Dependent Cache Misses with an Enhanced Memory Controller,” *Proceedings of the 43rd International Symposium on Computer Architecture (ISCA)*, Seoul, South Korea, June 2016.
51. Kevin Chang, Abhijith Kashyap, Hasan Hassan, Samira Khan, Kevin Hsieh, Donghyuk Lee, Saugata Ghose, Gennady Pekhimenko, Tianshi Li, Onur Mutlu, “Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Antibes Juan-les-Pins, France, June 2016.
52. Adwait Jog, Onur Kayiran, Ashutosh Pattnaik, Mahmut T. Kandemir, Onur Mutlu, Ravishankar Iyer, and Chita R. Das, “Exploiting Core Criticality for Enhanced GPU Performance,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Antibes Juan-les-Pins, France, June 2016.
53. Wayne Burleson, Onur Mutlu, Mohit Tiwari, “Who Is the Major Threat to Tomorrow’s Security? You, the Hardware Designer,” *Proceedings of the 53rd Design Automation Conference (DAC)*, Austin, TX, USA, June 2016. **Invited Paper and Presentation.**
54. Yang Li, Di Wang, Saugata Ghose, Jie Liu, Sriram Govindan, Sean James, Eric Peterson, John Siegler, Rachata Ausavarungnirun, Onur Mutlu, “SizeCap: Efficiently Handling Power Surges in Fuel Cell Powered Data Centers,” *Proceedings of the 22nd IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Barcelona, Spain, March 2016.
55. Kevin K. Chang, Prashant J. Nair, Saugata Ghose, Donghyuk Lee, Moinuddin K. Qureshi, Onur Mutlu, “Low-Cost Inter-Linked Subarrays (LISA): Enabling Fast Inter-Subarray Data Movement in DRAM,” *Proceedings of the 22nd IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Barcelona, Spain, March 2016.
56. Hasan Hassan, Gennady Pekhimenko, Nandita Vijaykumar, Vivek Seshadri, Donghyuk Lee, Oguz Ergin, Onur Mutlu, “ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality,” *Proceedings of the 22nd IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Barcelona, Spain, March 2016.
57. Gennady Pekhimenko, Evgeny Bolotin, Nandita Vijaykumar, Onur Mutlu, Todd C. Mowry, Stephen W. Keckler, “A Case for Toggle-Aware Compression for GPU Systems,” *Proceedings of the 22nd IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Barcelona, Spain, March 2016.
58. Lavanya Subramanian, Vivek Seshadri, Arnab Ghosh, Samira Khan, Onur Mutlu, “The Application Slowdown Model: Quantifying and Controlling the Impact of Inter-Application Interference at Shared Caches and Main Memory,” *Proceedings of the 48th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Waikiki, Hawaii, December 2015.
59. Vivek Seshadri, Thomas Mullins, Amirali Boroumand, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, Todd C. Mowry, “Gather-Scatter DRAM: In-DRAM Address Translation to Improve the Spatial Locality of Non-unit Strided Accesses,” *Proceedings of the 48th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Waikiki, Hawaii, December 2015.
60. Jinglei Ren, Jishen Zhao, Samira Khan, Jongmoo Choi, Yongwei Wu, Onur Mutlu, “ThyNVM: Enabling Software-Transparent Crash Consistency in Persistent Memory Systems,” *Proceedings of the 48th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Waikiki, Hawaii, December 2015.

61. Donghyuk Lee, Lavanya Subramanian, Rachata Ausavarungnirun, Jongmoo Choi, Onur Mutlu, “Decoupled Direct Memory Access: Isolating CPU and I/O Traffic by Leveraging a Dual-Port DRAM,” *Proceedings of the 24th ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, San Francisco, CA, October 2015.
62. Rachata Ausavarungnirun, Onur Kayiran, Saugata Ghose, Gabriel Loh, Chita Das, Mahmut Kandemir, Onur Mutlu, “Exploiting Inter-Warp Heterogeneity to Improve GPGPU Performance,” *Proceedings of the 24th ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, San Francisco, CA, October 2015.
63. Mohammad Fattah, Antti Airola, Rachata Ausavarungnirun, Nima Mirzaei, Pasi Liljeberg, Juha Plosila, Siamak Mohammadi, Tapio Pahikkala, Onur Mutlu, Hannu Tenhunen, “A Low-Overhead, Fully-Distributed, Guaranteed-Delivery Routing Algorithm for Faulty Network-on-Chips,” *Proceedings of the 9th International Networks-On-Chip Symposium (NOCS)*, Vancouver, BC, Canada, September 2015. **One of the 3 papers nominated for the Best Paper Award by the Program Committee.**
64. Nandita Vijaykumar, Gennady Pekhimenko, Adwait Jog, Abhishek Bhowmick, Rachata Ausavarungnirun, Chita Das, Mahmut Kandemir, Todd C. Mowry, Onur Mutlu, “A Case for Core-Assisted Bottleneck Acceleration in GPUs: Enabling Efficient Data Compression with Assist Warps,” *Proceedings of the 42nd International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2015.
65. Vivek Seshadri, Gennady Pekhimenko, Olatunji Ruwase, Onur Mutlu, Phillip Gibbons, Michael Kozuch, Todd C. Mowry, Trishul Chilimbi, “Page Overlays: An Enhanced Virtual Memory Framework to Enable Fine-grained Memory Management,” *Proceedings of the 42nd International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2015.
66. Junwhan Ahn, Sungjoo Yoo, Onur Mutlu, Kiyoung Choi, “PIM-Enabled Instructions: A Low-Overhead, Locality-Aware Processing-in-Memory Architecture,” *Proceedings of the 42nd International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2015.
67. Junwhan Ahn, Sungpack Hong, Sungjoo Yoo, Onur Mutlu, Kiyoung Choi, “A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing,” *Proceedings of the 42nd International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2015. **IEEE Micro Top Picks Honorable Mention.**
68. Justin Meza, Qiang Wu, Sanjeev Kumar, Onur Mutlu, “A Large-Scale Study of Flash Memory Failures in the Field,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Portland, OR, June 2015.
69. Yu Cai, Yixin Luo, Saugata Ghose, Onur Mutlu, “Read Disturb Errors in MLC NAND Flash Memory: Characterization and Mitigation,” *Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Rio De Janeiro, Brazil, June 2015.
70. Justin Meza, Qiang Wu, Sanjeev Kumar, Onur Mutlu, “Revisiting Memory Errors in Large-Scale Production Data Centers: Analysis and Modeling of New Trends from the Field,” *Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Rio De Janeiro, Brazil, June 2015.
71. Moinuddin Qureshi, Dae Hyun Kim, Samira Khan, Prashant Nair, Onur Mutlu, Chris Wilkerson, “AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems,” *Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Rio De Janeiro, Brazil, June 2015.
72. Yixin Luo, Yu Cai, Saugata Ghose, Jongmoo Choi, Onur Mutlu, “WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management,” *Proceedings of the 31st International Conference on Massive Storage Systems and Technologies (MSST)*, Santa Clara, CA, June 2015.
73. Dongwoo Kang, Seungjae Baek, Jongmoo Choi, Donghee Lee, Sam H. Noh, Onur Mutlu, “Amnesic Cache Management for Non-Volatile Memory,” *Proceedings of the 31st International Conference on Massive Storage Systems and Technologies (MSST)*, Santa Clara, CA, June 2015.
74. Hui Wang, Canturk Isci, Lavanya Subramanian, Jongmoo Choi, Depei Qian, Onur Mutlu, “A-DRM: Architecture-aware Distributed Resource Management of Virtualized Clusters,” *Proceedings of the 11th ACM SIGPLAN/SIGOPS International Conference on Virtual Execution Environments (VEE)*, Istanbul, Turkey, March 2015.
75. Yu Cai, Ken Mai, Onur Mutlu, “Comparative Evaluation of FPGA and ASIC Implementations of Bufferless and Buffered Routing Algorithms for On-Chip Networks,” *Proceedings of the 16th International Symposium on Quality Electronic Design (ISQED)*, Santa Clara, CA, March 2015.
76. Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, Onur Mutlu, “Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery,” *Proceedings of the 21st IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Bay Area, CA, February 2015. **Best Paper Session.**
77. Donghyuk Lee, Yoongu Kim, Gennady Pekhimenko, Samira Khan, Vivek Seshadri, Kevin Chang, Onur Mutlu, “Adaptive-Latency DRAM: Optimizing DRAM Timing for the Common-Case,” *Proceedings of the 21st IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Bay Area, CA, February 2015.
78. Gennady Pekhimenko, Tyler Huberty, Rui Cai, Onur Mutlu, Phillip P. Gibbons, Michael A. Kozuch, and Todd C. Mowry, “Exploiting Compressed Block Size as an Indicator of Future Reuse,” *Proceedings of the 21st IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Bay Area, CA, February 2015.
79. Jishen Zhao, Onur Mutlu, and Yuan Xie, “FIRM: Fair and High-performance Memory Control for Persistent Memory Systems,” *Proceedings of the 47th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Cambridge, UK, December 2014.
80. Onur Kayiran, Nachiappan C. Nachiappan, Adwait Jog, Rachata Ausavarungnirun, Mahmut Kandemir, Gabriel Loh, Onur Mutlu, and Chita Das, “Managing GPU Concurrency in Heterogeneous Architectures,” *Proceedings of the 47th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Cambridge, UK December 2014.

81. Rachata Ausavarungnirun, Chris Fallin, Xiangyao Yu, Kevin Chang, Greg Nazario, Reetuparna Das, Gabriel Loh, Onur Mutlu, “Design and Evaluation of Hierarchical Rings with Deflection Routing,” *Proceedings of the 26th IEEE International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, Paris, France, October 2014.
82. Lavanya Subramanian, Donghyuk Lee, Vivek Seshadri, Harsha Rastogi, Onur Mutlu, “The Blacklisting Memory Scheduler: Achieving High Performance and Fairness at Low Cost,” *Proceedings of the 32nd IEEE International Conference on Computer Design (ICCD)*, Seoul, South Korea, October 2014.
83. Chris Fallin, Chris Wilkerson, Onur Mutlu, “The Heterogeneous Block Architecture,” *Proceedings of the 32nd IEEE International Conference on Computer Design (ICCD)*, Seoul, South Korea, October 2014.
84. Youyou Lu, Jiwu Shu, Long Sun, Onur Mutlu, “Loose-Ordering Consistency for Persistent Memory,” *Proceedings of the 32nd IEEE International Conference on Computer Design (ICCD)*, Seoul, South Korea, October 2014.
85. James A. Jablin, Thomas B. Jablin, Onur Mutlu, Maurice Herlihy, “Warp-Aware Trace Scheduling for GPUs,” *Proceedings of the 23rd ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Edmonton, Alberta, Canada, August 2014.
86. Yoongu Kim, Ross Daly, Jeremie Kim, Ji Hye Lee, Chris Fallin, Chris Wilkerson, Onur Mutlu, “Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors,” *Proceedings of the 41st International Symposium on Computer Architecture (ISCA)*, Minneapolis, MN, June 2014.
87. Vivek Seshadri, Abhishek Bhowmick, Onur Mutlu, Phillip B. Gibbons, Michael A. Kozuch, Todd C. Mowry, “The Dirty-Block Index,” *Proceedings of the 41st International Symposium on Computer Architecture (ISCA)*, Minneapolis, MN, June 2014.
88. Yixin Luo, Sriram Govindan, Bikash Sharma, Mark Santaniello, Aman Kansal, Justin Meza, Onur Mutlu, Jie Liu, Badriddine Khessib, Kushagra Vaid, “Characterizing Application Memory Error Vulnerability to Optimize Data Center Cost via Heterogeneous-Reliability Memory,” *Proceedings of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Atlanta, GA, June 2014.
89. Yu Cai, Gulay Yalcin, Onur Mutlu, Eric Haratsch, Osman Unsal, Adrian Cristal, Ken Mai, “Neighbor-Cell Assisted Error Correction for MLC NAND Flash Memories,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Austin, TX, June 2014.
90. Samira Khan, Donghyuk Lee, Yoongu Kim, Alaa Alameldeen, Chris Wilkerson, Onur Mutlu, “The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study,” *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Austin, TX, June 2014.
91. Hyoseung Kim, Dionisio de Niz, Bjorn Andersson, Mark Klein, Onur Mutlu, Ragunathan Rajkumar, “Bounding Memory Interference Delay in COTS-based Multi-Core Systems,” *Proceedings of the 20th IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, Berlin, Germany, April 2014. **Best Paper Award.**
92. Kevin Chang, Donghyuk Lee, Zeshan Chishti, Chris Wilkerson, Alaa Alameldeen, Yoongu Kim, Onur Mutlu, “Improving DRAM Performance by Parallelizing Refreshes with Accesses,” *Proceedings of the 20th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Orlando, FL, February 2014.
93. Samira Khan, Alaa Alameldeen, Chris Wilkerson, Onur Mutlu, Daniel Jimenez, “Improving Cache Performance Using Read-Write Partitioning,” *Proceedings of the 20th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Orlando, FL, February 2014. **Best Paper Session.**
94. Gennady Pekhimenko, Vivek Seshadri, Yoongu Kim, Hongyi Xin, Onur Mutlu, Michael A. Kozuch, Phillip B. Gibbons, and Todd C. Mowry, “Linearly Compressed Pages: A Main Memory Compression Framework with Low Complexity and Low Latency,” *Proceedings of the 46th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Davis, CA, December 2013.
95. Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu, Michael A. Kozuch, Phillip B. Gibbons, and Todd C. Mowry, “RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization,” *Proceedings of the 46th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Davis, CA, December 2013.
96. Yu Cai, Onur Mutlu, Erich F. Haratsch, Ken Mai, “Program Interference in MLC NAND Flash Memory: Characterization, Modeling, and Mitigation,” *Proceedings of the 31st IEEE International Conference on Computer Design (ICCD)*, Asheville, NC, October 2013.
97. Youyou Lu, Jiwu Shu, Jia Guo, Shuai Li, Onur Mutlu, “LightTx: A Lightweight Transactional Design in Flash-based SSDs to Support Flexible Transactions,” *Proceedings of the 31st IEEE International Conference on Computer Design (ICCD)*, Asheville, NC, October 2013.
98. Justin Meza, Yixin Luo, Samira Khan, Jishen Zhao, Yuan Xie, Onur Mutlu, “A Case for Efficient Hardware-Software Cooperative Management of Storage and Memory,” *Proceedings of the 5th Workshop on Energy-Efficient Design (WEED)*, Tel-Aviv, Israel, June 2013.
99. Jamie Liu, Ben Jaiyen, Yoongu Kim, Chris Wilkerson, Onur Mutlu, “An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms,” *Proceedings of the 40th International Symposium on Computer Architecture (ISCA)*, Tel-Aviv, Israel, June 2013.
100. Jose A. Joao, M. Aater Suleman, Onur Mutlu, Yale N. Patt, “Utility-Based Acceleration of Multithreaded Applications on Asymmetric CMPs,” *Proceedings of the 40th International Symposium on Computer Architecture (ISCA)*, Tel-Aviv, Israel, June 2013.
101. Adwait Jog, Onur Kayiran, Asit K. Mishra, Mahmut T. Kandemir, Onur Mutlu, Ravishankar Iyer, Chita R. Das, “Orchestrated Scheduling and Prefetching in GPGPUs,” *Proceedings of the 40th International Symposium on Computer Architecture (ISCA)*, Tel-Aviv, Israel, June 2013.

102. Asit K. Mishra, Onur Mutlu, Chita R. Das, "A Heterogeneous Multiple Network-on-Chip Design: An Application-Aware Approach," *Proceedings of the 50th Design Automation Conference (DAC)*, Austin, TX, June 2013.
103. Onur Mutlu, "Memory Scaling: A Systems Architecture Perspective," *Proceedings of the 5th International Memory Workshop (IMW)*, Monterey, CA, May 2013. **Invited Paper.**
104. Alexey Tumanov, Joshua Wise, Onur Mutlu, Gregory R. Ganger, "Asymmetry-Aware Execution Placement on Manycore Chips," *Proceedings of the 3rd Workshop on Systems for Future Multicore Architectures (SFMA)*, Prague, Czech Republic, April 2013.
105. Emre Kultursay, Mahmut Kandemir, Anand Sivasubramaniam, Onur Mutlu, "Evaluating STT-RAM Technology as an Energy Efficient Main Memory Alternative," *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Austin, TX, April 2013.
106. Yu Cai, Eric F. Haratsch, Onur Mutlu, Ken Mai, "Threshold Voltage Distribution in MLC NAND Flash: Characterization, Analysis and Modeling," *Proceedings of the Design, Automation, and Test in Europe Conference (DATE)*, Grenoble, France, March 2013.
107. Adwait Jog, Onur Kayiran, Nachiappan Chidambaram Nachiappan, Asit K. Mishra, Mahmut T. Kandemir, Onur Mutlu, Ravishankar K. Iyer, Chita R. Das, "OWL: FOCused Thread-array aWare ScheduLing for Enhanced Performance in GPGPUs," *Proceedings of the 18th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Houston, TX, March 2013.
108. Donghyuk Lee, Yoongu Kim, Vivek Seshadri, Jamie Liu, Lavanya Subramanian, Onur Mutlu, "Tiered-Latency DRAM: A Low Latency and Low Cost DRAM Architecture," *Proceedings of the 19th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Shenzhen, China, February 2013.
109. Lavanya Subramanian, Vivek Seshadri, Yoongu Kim, Ben Jaiyen, Onur Mutlu, "MISE: Providing Performance Predictability and Improving Fairness in Shared Main Memory Systems," *Proceedings of the 19th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Shenzhen, China, February 2013.
110. Reetuparna Das, Rachata Ausavarungnirun, Onur Mutlu, Akhilesh Kumar, Mani Azimi, "Application-to-Core Mapping Policies to Reduce Memory System Interference in Multi-Core Systems," *Proceedings of the 19th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, Shenzhen, China, February 2013.
111. Kevin Chang, Rachata Ausavarungnirun, Chris Fallin, Onur Mutlu, "HAT: Heterogeneous Adaptive Throttling: for On-Chip Networks," *Proceedings of the 24th IEEE International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, New York, NY, October 2012.
112. HanBin Yoon, Justin Meza, Rachata Ausavarungnirun, Rachael Harding, Onur Mutlu, "A High-Performance and Energy-Efficient Row Buffer Locality-Aware Caching Policy for Hybrid Memories," *Proceedings of the 30th IEEE International Conference on Computer Design (ICCD)*, Montreal, Quebec, Canada, September 2012. **Best Paper Award (Computer Systems and Applications Track).**
113. Yu Cai, Gulay Yalcin, Onur Mutlu, Eric Haratsch, Adrian Cristal, Osman Unsal, Ken Mai, "Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime," *Proceedings of the 30th IEEE International Conference on Computer Design (ICCD)*, Montreal, Quebec, Canada, September 2012.
114. Vivek Seshadri, Onur Mutlu, Michael A. Kozuch, Todd C. Mowry, "The Evicted-Address Filter: A Unified Mechanism to Address Both Cache Pollution and Thrashing," *Proceedings of the 21st ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Minneapolis, MN, September 2012.
115. Gennady Pekhimenko, Vivek Seshadri, Onur Mutlu, Philip B. Gibbons, Michael A. Kozuch, Todd C. Mowry, "Base-Delta-Immediate Compression: A Practical Data Compression Mechanism for On-Chip Caches," *Proceedings of the 21st ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Minneapolis, MN, September 2012.
116. George Nychis, Chris Fallin, Thomas Moscibroda, Onur Mutlu, Srinivasan Seshan, "On-Chip Networks from a Networking Perspective: Congestion and Scalability in Many-core Interconnects," *Proceedings of the 2012 ACM SIGCOMM Conference (SIGCOMM)*, Helsinki, Finland, August 2012.
117. Jamie Liu, Ben Jaiyen, Richard Veras, Onur Mutlu, "RAIDR: Retention-Aware Intelligent DRAM Refresh," *Proceedings of the 39th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2012.
118. Yoongu Kim, Vivek Seshadri, Donghyuk Lee, Jamie Liu, Onur Mutlu, "A Case for Exploiting Subarray-Level Parallelism (SALP) in DRAM," *Proceedings of the 39th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2012.
119. Rachata Ausavarungnirun, Kevin Chang, Lavanya Subramanian, Gabriel Loh, Onur Mutlu, "Staged Memory Scheduling: Achieving High Performance and Scalability in Heterogeneous Systems," *Proceedings of the 39th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, Portland, OR, June 2012.
120. Chris Fallin, Greg Nazario, Xiangyao Yu, Kevin Chang, Rachata Ausavarungnirun, Onur Mutlu, "MinBD: Minimally-Buffered Deflection Routing for Energy-Efficient Interconnect," *Proceedings of the 6th International Networks-On-Chip Symposium (NOCS)*, Lyngby, Denmark, May 2012. **One of the 5 papers nominated for the Best Paper Award by the Program Committee.**
121. Yu Cai, Eric F. Haratsch, Onur Mutlu, Ken Mai, "Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis," *Proceedings Design, Automation, and Test in Europe Conference (DATE)*, Dresden, Germany, March 2012.
122. Jose A. Joao, M. Aater Suleman, Onur Mutlu, Yale N. Patt, "Bottleneck Identification and Scheduling in Multithreaded Applications," *Proceedings of the 17th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, London, UK, March 2012.

123. Sai Prashanth Muralidhara, Lavanya Subramanian, Onur Mutlu, Mahmut Kandemir, Thomas Moscibroda, "Reducing Memory Interference in Multicore Systems via Application-Aware Memory Channel Partitioning," *Proceedings of the 44th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Porto Alegre, Brazil, December 2011.
124. Eiman Ebrahimi, Rustam Miftakhutdinov, Chang Joo Lee, Onur Mutlu, Chris Fallin, Yale N. Patt, "Parallel Application Memory Scheduling", *Proceedings of the 44th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Porto Alegre, Brazil, December 2011.
125. Veynu Narasiman, Chang Joo Lee, Michael Shebanow, Rustam Miftakhutdinov, Onur Mutlu, Yale N. Patt, "Improving GPU Performance via Large Warps and Two-Level Warp Scheduling," *Proceedings of the 44th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Porto Alegre, Brazil, December 2011.
126. Boris Grot, Joel Hestness, Steve Keckler, Onur Mutlu, "Kilo-NOC: A Heterogeneous Network-on-Chip Architecture for Scalability and Service Guarantees," *Proceedings of the 38th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, San Jose, CA, June 2011. **One of the 12 papers of 2011 selected as "Top Picks" by IEEE Micro.**
127. Eiman Ebrahimi, Chang Joo Lee, Onur Mutlu, Yale N. Patt, "Prefetch-Aware Shared-Resource Management for Multi-Core Systems," *Proceedings of the 38th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, San Jose, CA, June 2011.
128. Howard David, Chris Fallin, Eugene Gorbato, Ulf R. Hanebutte, Onur Mutlu, "Memory Power Management via Dynamic Voltage/Frequency Scaling," *Proceedings of the 8th IEEE International Conference on Autonomic Computing (ICAC)*, Karlsruhe, Germany, June 2011.
129. Michael Papamichael, James Hoe, Onur Mutlu, "FIST: A Fast, Lightweight, FPGA-Friendly Packet Latency Estimator for NoC Modeling in Full-System Simulations," *Proceedings of the 5th International Networks-On-Chip Symposium (NOCS)*, Pittsburgh, PA, May 2011.
130. Chris Fallin, Chris Craik, Onur Mutlu, "CHIPPER: A Low-complexity Bufferless Deflection Router," *Proceedings of the 17th IEEE International Symposium on High Performance Computer Architecture (HPCA)*, San Antonio, TX, February 2011.
131. Yoongu Kim, Michael Papamichael, Onur Mutlu, and Mor Harchol-Balter, "Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior," *Proceedings of the 43rd IEEE/ACM International Symposium on Microarchitecture (MICRO)*, Atlanta, GA, December 2010. **One of the 11 papers of 2010 selected as "Top Picks" by IEEE Micro.**
132. George Nychis, Chris Fallin, Thomas Moscibroda, Onur Mutlu, "Next Generation On-Chip Networks: What Kind of Congestion Control Do We Need?," *Proceedings of the 9th ACM Workshop on Hot Topics in Networks (HotNETS)*, Monterey, CA, October 2010.
133. Tanausu Ramirez, Alex Pajuelo, Oliverio Santana, Onur Mutlu, and Mateo Valero, "Efficient Runahead Threads," *Proceedings of the 19th ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Vienna, Austria, September 2010.
134. Reetuparna Das, Onur Mutlu, Thomas Moscibroda, and Chita R. Das, "Aergia: Exploiting Packet Latency Slack in On-Chip Networks," *Proceedings of the 37th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, St. Malo, France, June 2010. **One of the 11 papers of 2010 selected as "Top Picks" by IEEE Micro.**
135. M. Aater Suleman, Onur Mutlu, Jose A. Joao, Khubaib, and Yale N. Patt, "Data Marshaling for Multi-core Architectures," *Proceedings of the 37th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, St. Malo, France, June 2010. **One of the 11 papers of 2010 selected as "Top Picks" by IEEE Micro.**
136. Boris Grot, Steve Keckler, and Onur Mutlu, "Topology-aware Quality-of-Service Support in Highly Integrated Chip Multiprocessors," *Proceedings of the 6th Annual Workshop on the Interaction between Operating Systems and Computer Architecture (WIOSCA)*, in conjunction with the 37th International Symposium on Computer Architecture, St. Malo, France, June 2010.
137. Paul Bogdan, Miray Kas, Radu Marculescu, and Onur Mutlu, "QuaLe: A Quantum-Leap Inspired Model for Non-Stationary Analysis of NoC Traffic in Multi-Processor Platforms," *Proceedings of the 4th International Networks-On-Chip Symposium (NOCS)*, May 2010.
138. Yanjing Li, Onur Mutlu, Donald S. Gardner, and Subhasish Mitra, "Concurrent Autonomous Self-Test for Uncore Components in System-on-Chips" *Proceedings of the 28th IEEE VLSI Test Symposium (VTS)*, Santa Cruz, CA, April 2010. **Best Paper Award.**
139. Eiman Ebrahimi, Chang Joo Lee, Onur Mutlu, and Yale N. Patt, "Fairness via Source Throttling: A Configurable and High-Performance Fairness Substrate for Multi-Core Memory Systems," *Proceedings of the 15th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Pittsburgh, PA, March 2010. **Best Paper Award.**
140. Yoongu Kim, Dongsu Han, Onur Mutlu, and Mor Harchol-Balter, "ATLAS: A Scalable and High-Performance Scheduling Algorithm for Multiple Memory Controllers," *Proceedings of the 16th IEEE International Conference on High Performance Computer Architecture (HPCA)*, Bangalore, India, January 2010. **Best Paper Session. One of the 4 papers nominated for the Best Paper Award by the Program Committee.**
141. Reetuparna Das, Onur Mutlu, Thomas Moscibroda, and Chita R. Das, "Application-Aware Prioritization Mechanisms for On-Chip Networks," *Proceedings of the 42nd IEEE/ACM International Symposium on Microarchitecture (MICRO)*, New York, NY, pp. 280-291, December 2009.
142. Eiman Ebrahimi, Onur Mutlu, Chang Joo Lee, and Yale N. Patt, "Coordinated Control of Multiple Prefetchers in Multi-Core Systems," *Proceedings of the 42nd IEEE/ACM International Symposium on Microarchitecture (MICRO)*, New York, NY, pp. 316-326, December 2009.
143. Boris Grot, Steve Keckler, and Onur Mutlu, "Preemptive Virtual Clock: A Flexible, Efficient, and Cost-effective QOS Scheme for Networks-on-Chip," *Proceedings of the 42nd IEEE/ACM International Symposium on Microarchitecture (MICRO)*, New York, NY, pp. 268-279, December 2009.

144. Chang Joo Lee, Veynu Narasiman, Onur Mutlu, and Yale N. Patt, "Improving Memory Bank-Level Parallelism in the Presence of Prefetching," *Proceedings of the 42nd IEEE/ACM International Symposium on Microarchitecture (MICRO)*, New York, NY, pp. 327-336, December 2009.
145. Yanjing Li, Onur Mutlu, and Subhasish Mitra, "Operating System Scheduling for Efficient Online Self-Test in Robust Systems," *Proceedings of the 2009 IEEE/ACM International Conference on Computer-Aided Design (ICCAD)*, San Jose, CA, pp. 201-208, November 2009.
146. Thomas Moscibroda and Onur Mutlu, "A Case for Bufferless Routing in On-Chip Networks," *Proceedings of the 36th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 196-207, Austin, TX, June 2009.
147. José A. Joao, Onur Mutlu, and Yale N. Patt, "Flexible Reference Counting Based Hardware Acceleration for Garbage Collection," *Proceedings of the 36th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 418-428, Austin, TX, June 2009.
148. Benjamin Lee, Engin Ipek, Onur Mutlu, and Doug Burger, "Architecting Phase Change Memory as a Scalable DRAM Alternative," *Proceedings of the 36th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 2-13, Austin, TX, June 2009. **One of the 13 papers of 2009 selected as "Top Picks" by IEEE Micro. Opening Paper of the Conference. Selected for publication in CACM Research Highlights.**
149. M. Aater Suleman, Onur Mutlu, Moinuddin Qureshi, Yale N. Patt, "Accelerating Critical Section Execution with Asymmetric Multi-Core Architectures," *Proceedings of the 14th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 253-264, Washington, DC, March 2009. **One of the 13 papers of 2009 selected as "Top Picks" by IEEE Micro.**
150. Eiman Ebrahimi, Onur Mutlu, Yale N. Patt, "Techniques for Bandwidth-Efficient Prefetching of Linked Data Structures in Hybrid Prefetching Systems," *Proceedings of the 15th IEEE International Conference on High Performance Computer Architecture (HPCA)*, pp. 7-17, Raleigh, NC, February 2009. **Best Paper Session. One of the 3 papers nominated for the Best Paper Award by the Program Committee.**
151. Boris Grot, Joel Hestness, Steve Keckler, Onur Mutlu, "Express Cube Topologies for On-Chip Interconnects," *Proceedings of the 15th IEEE International Conference on High Performance Computer Architecture (HPCA)*, pp. 163-174, Raleigh, NC, February 2009.
152. Kypros Constantinides, Onur Mutlu, Todd Austin, "Online Design Bug Detection: RTL Analysis, Flexible Mechanisms, and Evaluation," *Proceedings of the 41st IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 282-293, Lake Como, Italy, November 2008.
153. Chang Joo Lee, Onur Mutlu, Veynu Narasiman, Yale N. Patt, "Prefetch Aware DRAM Controllers," *Proceedings of the 41st IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 200-209, Lake Como, Italy, November 2008.
154. Thomas Moscibroda and Onur Mutlu, "Distributed Order Scheduling and its Application to Multi-Core DRAM Controllers," *Proceedings of the 27th Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC)*, pp. 365-374, Toronto, Canada, August 2008.
155. Onur Mutlu and Thomas Moscibroda, "Parallelism-Aware Batch Scheduling: Enhancing both Performance and Fairness of Shared DRAM Systems," *Proceedings of the 35th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 63-74, Beijing, China, June 2008. **One of the 12 papers of 2008 selected as "Top Picks" by IEEE Micro.**
156. Engin Ipek, Onur Mutlu, José F. Martínez, Rich Caruana, "Self-Optimizing Memory Controllers: A Reinforcement Learning Approach," *Proceedings of the 35th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 39-50, Beijing, China, June 2008.
157. José A. Joao, Onur Mutlu, Hyesoon Kim, Rishi Agarwal, Yale N. Patt, "Improving the Performance of Object-Oriented Languages with Dynamic Predication of Indirect Jumps," *Proceedings of the 13th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 80-90, Seattle, WA, March 2008.
158. Chang Joo Lee, Hyesoon Kim, Onur Mutlu, Yale N. Patt, "Performance-aware Speculation Control using Wrong Path Usefulness Prediction," *Proceedings of the 14th IEEE International Conference on High Performance Computer Architecture (HPCA)*, pp. 39-49, Salt Lake City, UT, February 2008.
159. Onur Mutlu and Thomas Moscibroda, "Stall-Time Fair Memory Access Scheduling for Chip Multiprocessors," *Proceedings of the 40th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 146-158, Chicago, IL, December 2007.
160. Kypros Constantinides, Onur Mutlu, Todd Austin, and Valeria Bertacco, "Software-Based Online Detection of Hardware Defects: Mechanisms, Architectural Support, and Evaluation," *Proceedings of the 40th IEEE/ACM International Symposium on Microarchitecture (MICRO)*, pp. 97-108, Chicago, IL, December 2007.
161. Thomas Moscibroda and Onur Mutlu, "Memory Performance Attacks: Denial of Memory Service in Multi-Core Systems," *Proceedings of the 16th USENIX Security Symposium (USENIX SECURITY)*, pp. 257-274, Boston, MA, August 2007.
162. Hyesoon Kim, José A. Joao, Onur Mutlu, Chang Joo Lee, Yale N. Patt, Robert S. Cohn, "VPC Prediction: Reducing the Cost of Indirect Branches via Hardware-based Dynamic Devirtualization," *Proceedings of the 34th IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 424-435, San Diego, CA, June 2007.
163. Hyesoon Kim, José A. Joao, Onur Mutlu, Yale N. Patt, "Profile-assisted Compiler Support for Dynamic Predication in Diverge-Merge Processors," *Proceedings of the 5th Annual International Symposium on Code Generation and Optimization (CGO)*, pp. 367-378, San Jose, CA, March 2007.

164. Santhosh Srinath, Onur Mutlu, Hyesoon Kim, Yale N. Patt, "Feedback Directed Prefetching: Improving the Performance and Bandwidth-Efficiency of Hardware Prefetchers," *Proceedings of the 13th IEEE International Conference on High Performance Computer Architecture (HPCA)*, pp. 63-74, Phoenix, AZ, February 2007. **One of the 5 papers nominated for the Best Paper Award by the Program Committee.**
165. Hyesoon Kim, José Joao, Onur Mutlu, Yale N. Patt, "Diverge-Merge Processor (DMP): Dynamic Predicated Execution of Complex Control-Flow Graphs Based on Frequently Executed Paths," *Proceedings of the 38th IEEE/ACM Annual International Symposium on Microarchitecture (MICRO)*, pp. 53-64, Orlando, FL, December 2006. **One of the 11 papers of 2006 selected as "Top Picks" by IEEE Micro.**
166. Moinuddin K. Qureshi, Daniel N. Lynch, Onur Mutlu, Yale N. Patt, "A Case for MLP-Aware Cache Replacement," *Proceedings of the 33rd IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 167-177, Boston, MA, USA, June 2006.
167. Hyesoon Kim, M. Aater Suleman, Onur Mutlu, Yale N. Patt, "2D-Profilng: Detecting Input-Dependent Branches with a Single Input Data Set," *Proceedings of the 4th Annual International Symposium on Code Generation and Optimization (CGO)*, pp. 159-169, New York, NY, USA, March 2006.
168. Onur Mutlu, Hyesoon Kim, Yale N. Patt, "Address-Value Delta (AVD) Prediction: Increasing the Effectiveness of Runahead Execution by Exploiting Regular Memory Allocation Patterns," *Proceedings of the 38th IEEE/ACM Annual International Symposium on Microarchitecture (MICRO)*, pp. 233-244, Barcelona, Spain, November 2005. **One of the 5 papers nominated for the Best Paper Award by the Program Committee.**
169. Hyesoon Kim, Onur Mutlu, Jared Stark, Yale N. Patt, "Wish Branches: Combining Conditional Branching and Predication for Adaptive Predicated Execution," *Proceedings of the 38th IEEE/ACM Annual International Symposium on Microarchitecture (MICRO)*, pp. 43-54, Barcelona, Spain, November 2005. **One of the 13 papers of 2005 selected as "Top Picks" by IEEE Micro.**
170. Onur Mutlu, Hyesoon Kim, Yale N. Patt, "Techniques for Efficient Processing in Runahead Execution Engines," *Proceedings of the 32nd IEEE/ACM International Symposium on Computer Architecture (ISCA)*, pp. 370-381, Madison, WI, USA, June 2005. **One of the 13 papers of 2005 selected as "Top Picks" by IEEE Micro.**
171. Moinuddin K. Qureshi, Onur Mutlu, and Yale N. Patt, "Microarchitecture-Based Introspection: A Technique for Transient-Fault Tolerance in Microprocessors," *Proceedings of the 2005 IEEE International Conference on Dependable Systems and Networks (DSN)*, pp. 434-443, Yokohama, Japan, June 2005.
172. David N. Armstrong, Hyesoon Kim, Onur Mutlu, Yale N. Patt, "Wrong Path Events: Exploiting Unusual and Illegal Program Behavior for Early Misprediction Detection and Recovery," *Proceedings of the 37th IEEE/ACM Annual International Symposium on Microarchitecture (MICRO)*, pp. 119-128, Portland, OR, USA, December 2004.
173. Onur Mutlu, Hyesoon Kim, David N. Armstrong, Yale N. Patt, "Cache Filtering Techniques to Reduce the Negative Impact of Useless Speculative Memory References on Processor Performance," *Proceedings of the 16th IEEE International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, pp. 2-9, Foz do Iguacu, Brazil, October 2004. **Opening Paper of the Conference.**
174. Onur Mutlu, Hyesoon Kim, David N. Armstrong, Yale N. Patt, "Understanding The Effects of Wrong-Path Memory References on Processor Performance," *Proceedings of the 3rd ACM Workshop on Memory Performance Issues (WMPPI), in conjunction with the 31st International Symposium on Computer Architecture*, pp. 56-64, Munchen, Germany, June 2004.
175. Onur Mutlu, Jared Stark, Chris Wilkerson, Yale N. Patt, "Runahead Execution: An Alternative to Very Large Instruction Windows for Out-of-order Processors," *Proceedings of the 9th IEEE International Conference on High Performance Computer Architecture (HPCA)*, pp. 129-140, Anaheim, CA, USA, February 2003. **One of the 15 papers of 2003 selected as "Top Picks" by IEEE Micro.**

Refereed and Invited Journal Publications

176. Damla Senol, Jeremie Kim, Saugata Ghose, Can Alkan, Onur Mutlu, "Nanopore Sequencing Technology and Tools: Computational Analysis of the Current State, Bottlenecks and Future Directions," *Briefings in Bioinformatics (BIB)*, 2018.
177. Jeremie S. Kim, Damla Senol Cali, Hongyi Xin, Donghyuk Lee, Saugata Ghose, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, Onur Mutlu, "GRIM-Filter: Fast Seed Location Filtering in DNA Read Mapping Using Processing-in-Memory Technologies," *BMC Genomics*, to appear in 2018. Also appears in *Proceedings of the 16th Asia Pacific Bioinformatics Conference (APBC)*, Yokohama, Japan, January 2018.
178. Yu Cai, Saugata Ghose, Yixin Luo, Erich F. Haratsch, Onur Mutlu, "Flash Memory SSD Errors, Mitigation, and Recovery," *Proceedings of the IEEE*, September 2017.
179. Mohammed Alser, Onur Mutlu, Can Alkan, "MAGNET: Understanding and Improving the Accuracy of Genome Pre-Alignment Filtering," *IPSI Transactions on Internet Research*, [published online], July 2017.
180. Mohammed Alser, Hasan Hassan, Hongyi Xin, Oguz Ergin, Onur Mutlu, Can Alkan, "GateKeeper: A New Hardware Architecture for Accelerating Pre-Alignment in DNA Short Read Mapping," *Bioinformatics*, [published online], May 31, 2017.
181. Youyou Lu, Jiwu Shu, Long Sun, Onur Mutlu, "Improving Performance and Endurance of Persistent Memory with Loose-Ordering Consistency," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, accepted, March 2017.
182. Samira Khan, Chris Wilkerson, Donghyuk Lee, Alaa R. Alameldeen, Onur Mutlu, "A Case for Memory Content-Based Detection and Mitigation of Data-Dependent Failures in DRAM," *IEEE Computer Architecture Letters (CAL)*, November 2016.
183. Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, Onur Mutlu, "Enabling Accurate and Practical Online Flash Channel Modeling for Modern MLC NAND Flash Memory," *IEEE Journal on Selected Areas in Communications (JSAC)*, September 2016.

184. Amirali Boroumand, Saugata Ghose, Minesh Patel, Hasan Hassan, Brandon Lucia, Kevin Hsieh, Krishna T. Malladi, Hongzhong Zheng, [Onur Mutlu](#), “LazyPIM: An Efficient Cache Coherence Mechanism for Processing-in-Memory,” *IEEE Computer Architecture Letters (CAL)*, June 2016.
185. Rachata Ausavarungnirun, Chris Fallin, Xiangyao Yu, Kevin Chang, Greg Nazario, Reetuparna Das, Gabriel Loh, [Onur Mutlu](#), “A Case for Hierarchical Rings with Deflection Routing: An Energy-Efficient On-Chip Communication Substrate,” *Parallel Computing (PARCO)*, accepted, January 2016.
186. Hyoseung Kim, Dionisio de Niz, Bjorn Andersson, Mark Klein, [Onur Mutlu](#), Ragnathan Rajkumar, “Bounding and Reducing Memory Interference in COTS-based Multi-Core Systems,” *Real-Time Systems (RTS)*, accepted, January 2016.
187. Lavanya Subramanian, Donghyuk Lee, Vivek Seshadri, Harsha Rastogi, [Onur Mutlu](#), “The Blacklisting Memory Scheduler: Balancing Performance, Fairness and Complexity,” *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, accepted, January 2016.
188. Amir Yazdanbakhsh, Gennady Pekhimenko, Bradley Thwaites, Hadi Esmaeilzadeh, [Onur Mutlu](#), Todd C. Mowry, “Mitigating the Memory Bottleneck With Approximate Load Value Prediction,” *IEEE Design and Test (D&T)*, January/February 2016.
189. Hiroyuki Usui, Lavanya Subramanian, Kevin Kai-Wei Chang, [Onur Mutlu](#), “DASH: Deadline-Aware High-Performance Memory Scheduler for Heterogeneous Systems with Hardware Accelerators,” *ACM Transactions on Architecture and Code Optimization (TACO)*, January 2016.
190. Donghyuk Lee, Gennady Pekhimenko, Samira Khan, Saugata Ghose, [Onur Mutlu](#), “Simultaneous Multi-Layer Access: Improving 3D-Stacked Memory Bandwidth at Low Cost,” *ACM Transactions on Architecture and Code Optimization (TACO)*, January 2016.
191. Amir Yazdanbakhsh, Gennady Pekhimenko, Bradley Thwaites, Hadi Esmaeilzadeh, [Onur Mutlu](#), Todd C. Mowry, “RFVP: Rollback-Free Value Prediction with Safe to Approximate Loads,” *ACM Transactions on Architecture and Code Optimization (TACO)*, January 2016.
192. Hongyi Xin, Sunny Nahar, Richard Zhu, John Emmons, Gennady Pekhimenko, Carl Kingsford, Can Alkan, [Onur Mutlu](#), “Optimal Seed Solver: Optimizing Seed Selection in Read Mapping,” *Bioinformatics*, [published online], November 14, 2015.
193. Gennady Pekhimenko, Evgeny Bolotin, Mike O’Connor, [Onur Mutlu](#), Todd C. Mowry, Stephen W. Keckler, “Toggle-Aware Bandwidth Compression for GPUs,” *IEEE Computer Architecture Letters (CAL)*, May 2015.
194. Vivek Seshadri, Kevin Hsieh, Amirali Boroumand, Donghyuk Lee, Michael Kozuch, [Onur Mutlu](#), Phillip B. Gibbons, Todd C. Mowry, “Fast Bulk Bitwise AND and OR in DRAM,” *IEEE Computer Architecture Letters (CAL)*, April 2015.
195. Yoongu Kim, Weikun Yang, [Onur Mutlu](#), “Ramulator: A Fast and Extensible DRAM Simulator,” *IEEE Computer Architecture Letters (CAL)*, March 2015.
196. [Onur Mutlu](#), Justin Meza, Lavanya Subramanian, “The Main Memory System: Challenges and Opportunities,” *Communications of the Korean Institute of Information Scientists and Engineers (KIISE)*, January 2015. **Invited Article.**
197. Hongyi Xin, John Greth, John Emmons, Gennady Pekhimenko, Carl Kingsford, Can Alkan, [Onur Mutlu](#), “Shifted Hamming Distance: A Fast and Accurate SIMD-friendly Filter to Accelerate Alignment Verification in Read Mapping,” *Bioinformatics*, [published online], January 10, 2015.
198. Vivek Seshadri, Samihan Yedkar, Hongyi Xin, [Onur Mutlu](#), Phillip P. Gibbons, Michael A. Kozuch, and Todd C. Mowry, “Mitigating Prefetcher-Caused Pollution using Informed Caching Policies for Prefetched Blocks,” *ACM Transactions on Architecture and Code Optimization (TACO)*, January 2015.
199. Youyou Lu, Jiwu Shu, Jia Guo, Shuai Li, [Onur Mutlu](#), “High-Performance and Lightweight Transaction Support in Flash-Based SSDs,” *IEEE Transactions on Computers*, [published online], January 2015.
200. [Onur Mutlu](#), Lavanya Subramanian, “Research Problems and Opportunities in Memory Systems,” *Supercomputing Frontiers and Innovations*, December 2014. **Invited Article.**
201. HanBin Yoon, Justin Meza, Naveen Muralimanohar, Norman P. Jouppi, [Onur Mutlu](#), “Efficient Data Mapping and Buffering Techniques for Multi-Level Cell Phase-Change Memories,” *ACM Transactions on Architecture and Code Optimization (TACO)*, December 2014. **Best Presentation Award at HiPEAC 2015 Conference.**
202. Yu Cai, Gulay Yalcin, [Onur Mutlu](#), Erich F. Haratsch, Adrian Cristal, Osman S. Unsal, Ken Mai, “Error Analysis and Retention-Aware Error Management for NAND Flash Memory,” *Intel Technology Journal*, vol. 17 (1), May 2013. **Invited article.**
203. Hongyi Xin, Donghyuk Lee, Farhad Hormozdiari, Samihan Yedkar, [Onur Mutlu](#), Can Alkan, “Accelerating Read Mapping with FastHASH,” *BMC Genomics* 14(Suppl 1):S13, 21 January 2013. Also appears in *Proceedings of the 11th Asia Pacific Bioinformatics Conference (APBC)*, Vancouver, BC, Canada, January 2013.
204. Boris Grot, Joel Hestness, Stephen W. Keckler, and [Onur Mutlu](#), “A QoS-Enabled On-Die Interconnect Fabric for Kilo-Node Chips,” *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 32 (3), May/June 2012.
205. Eiman Ebrahimi, Chang Joo Lee, [Onur Mutlu](#), Yale N. Patt, “Fairness via Source Throttling: A Configurable and High-Performance Fairness Substrate for Multi-Core Memory Systems,” *ACM Transactions on Computer Systems*, April 2012.
206. Justin Meza, Jichuan Chang, HanBin Yoon, [Onur Mutlu](#), Parthasarathy Ranganathan, “Enabling Efficient and Scalable Hybrid Memories Using Fine-Granularity DRAM Cache Management,” *IEEE Computer Architecture Letters (CAL)*, February 2012.
207. Chang Joo Lee, [Onur Mutlu](#), Veynu Narasiman, and Yale N. Patt, “Prefetch-Aware Memory Controllers,” *IEEE Transactions on Computers*, vol. 60 (10), pp. 1406-1430, October 2011.

208. Yoongu Kim, Michael Papamichael, [Onur Mutlu](#), and Mor Harchol-Balter, "Thread Cluster Memory Scheduling," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 31 (1), January/February 2011.
209. Reetuparna Das, [Onur Mutlu](#), Thomas Moscibroda, and Chita R. Das, "Aergia: A Network-On-Chip Exploiting Packet Latency Slack," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 31 (1), January/February 2011.
210. M. Aater Suleman, [Onur Mutlu](#), Jose A. Joao, Khubaib, and Yale N. Patt, "Data Marshaling for Multi-core Systems," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 31 (1), January/February 2011.
211. Benjamin Lee, Engin Ipek, [Onur Mutlu](#), Doug Burger, "Phase Change Memory Architecture and the Quest for Scalability," *Communications of the ACM (CACM)*, Research Highlight, vol. 53 (7), pp. 99-106, July 2010.
212. M. Aater Suleman, [Onur Mutlu](#), Moinuddin Qureshi, Yale N. Patt, "Accelerating Critical Section Execution with Asymmetric Multi-Core Architectures," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 30 (1), pp. 131-141, January/February 2010.
213. Benjamin Lee, Ping Zhou, Jun Yang, Youtao Zhang, Bo Zhao, Engin Ipek, [Onur Mutlu](#), Doug Burger, "Phase Change Technology and the Future of Main Memory," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 30 (1), pp. 60-70, January/February 2010.
214. Can Alkan, Jeffrey M. Kidd, Tomas Marques-Bonet, Gozde Aksay, Francesca Antonacci, Fereydoun Hormozdiari, Jacob O. Kitzman, Carl Baker, Maika Malig, [Onur Mutlu](#), S. Cenk Sahinalp, Richard A. Gibbs, Evan E. Eichler, "Personalized Copy-Number and Segmental Duplication Maps using Next-Generation Sequencing," *Nature Genetics*, [Epub: August 30], vol. 41 (10), pp. 1061-1067, October 2009.
215. Hyesoon Kim, José A. Joao, [Onur Mutlu](#), Chang Joo Lee, Yale N. Patt, Robert S. Cohn, "Virtual Program Counter (VPC) Prediction: Very Low-Cost Indirect Branch Prediction using Conditional Branch Prediction Hardware," *IEEE Transactions on Computers (TC)*, vol. 58 (9), pp. 1153-1170, September 2009. **Online featured article.**
216. Kypros Constantinides, [Onur Mutlu](#), Todd Austin, and Valeria Bertacco, "A Flexible Software-Based Framework for Online Detection of Hardware Defects," *IEEE Transactions on Computers (TC)*, vol. 58 (8), pp. 1063-1079, August 2009.
217. [Onur Mutlu](#) and Thomas Moscibroda, "Parallelism-Aware Batch Scheduling: Enabling High-Performance and Fair Memory Controllers," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 28(1), January/February 2009.
218. José A. Joao, [Onur Mutlu](#), Hyesoon Kim, Yale N. Patt, "Dynamic Predication of Indirect Jumps," *IEEE Computer Architecture Letters (CAL)*, vol. 6, May 2007.
219. Hyesoon Kim, José A. Joao, [Onur Mutlu](#), Yale N. Patt, "Diverge-Merge Processor: Generalized and Energy-Efficient Dynamic Predication," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 27(1), pp. 94-104, January/February 2007.
220. [Onur Mutlu](#), Hyesoon Kim, Yale N. Patt, "Address-Value Delta (AVD) Prediction: A Hardware Technique for Efficiently Parallelizing Dependent Cache Misses," *IEEE Transactions on Computers (TC)*, vol. 55 (12), pp. 1491-1508, December 2006. **Featured article.**
221. [Onur Mutlu](#), Hyesoon Kim, Yale N. Patt, "Efficient Runahead Execution: Power-Efficient Memory Latency Tolerance," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 26(1), pp. 10-20, January/February 2006.
222. Hyesoon Kim, [Onur Mutlu](#), Jared Stark, Yale N. Patt, "Wish Branches: Enabling Adaptive and Aggressive Predicated Execution," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 26(1), pp. 48-58, January/February 2006.
223. [Onur Mutlu](#), Hyesoon Kim, David N. Armstrong, Yale N. Patt, "An Analysis of the Performance Impact of Wrong-Path Memory References on Out-of-Order and Runahead Execution Processors," *IEEE Transactions on Computers (TC)*, vol. 54 (12), pp. 1556-1571, December 2005.
224. [Onur Mutlu](#), Hyesoon Kim, David N. Armstrong, Yale N. Patt, "Using the First-Level Caches as Filters to Reduce the Pollution Caused by Speculative Memory References," *International Journal of Parallel Programming (IJPP)*, vol. 33 (5), pp. 529-559, October 2005.
225. [Onur Mutlu](#), Hyesoon Kim, Jared Stark, Yale N. Patt, "On Reusing the Results of Pre-Executed Instructions in a Runahead Execution Processor," *IEEE Computer Architecture Letters (CAL)*, vol. 4, January 2005.
226. [Onur Mutlu](#), Jared Stark, Chris Wilkerson, Yale N. Patt, "Runahead Execution: An Effective Alternative to Large Instruction Windows," *IEEE Micro, Special Issue: Top Picks from Computer Architecture Conferences (TOP PICKS)*, vol. 23(6), pp. 20-25, November/December 2003.

Editorial Articles

227. [Onur Mutlu](#), Scott A. Mahlke, Thomas M. Conte, Wen-mei W. Hwu, "Iterative Module Scheduling," *IEEE Micro*, January/February 2018.
228. [Onur Mutlu](#), Rich Belgard, "Common Bonds: MIPS, HPS, Two-Level Branch Prediction, and Compressed Code RISC Processor," *IEEE Micro*, July/August 2016.
229. [Onur Mutlu](#), Rich Belgard, "The 2014 MICRO Test of Time Award Winners: From 1978 to 1992," *IEEE Micro*, January/February 2016.
230. [Onur Mutlu](#), Rich Belgard, "Introducing the MICRO Test of Time Awards: Concept, Process, 2014 Winners, and the Future," *IEEE Micro*, vol. 35(2), March/April 2015.
231. Yale N. Patt, [Onur Mutlu](#), "Guest Editors' Introduction: Top Picks," *IEEE Micro, Special Issue (IEEE MICRO)*, vol. 31(1), January/February 2011.
232. Sangyeun Cho, Tao Li, [Onur Mutlu](#), "Interaction of Many-core Computer Architecture and Operating Systems: Guest Editors' Introduction," *IEEE Micro, Special Issue (IEEE MICRO)*, vol. 28(3), pp. 2-5, May/June 2008.

Significant Abstracts and Conference Presentations (Most are Refereed)

233. Onur Mutlu, “Achieving the Next Breakthroughs in NAND Flash Memory: Changing Our Fixed Mindset,” *Proceedings of the Flash Memory Summit (FMS)*, Santa Clara, CA, August 2017.
234. Damla Senol, Jeremie Kim, Saugata Ghose, Can Alkan, Onur Mutlu, “Nanopore Sequencing Technology and Tools: Computational Analysis of the Current State, Bottlenecks and Future Directions,” *Pacific Symposium on Biocomputing (PSB) Poster Session*, Hawaii, January 2017.
235. Jeremie Kim, Damla Senol, Hongyi Xin, Donghyuk Lee, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, Onur Mutlu, “Genome Read In-Memory (GRIM) Filter: Fast Location Filtering in DNA Read Mapping using Emerging Memory Technologies,” *Pacific Symposium on Biocomputing (PSB) Poster Session*, Hawaii, January 2017.
236. Onur Mutlu, “ThyNVM: Software-Transparent Crash Consistency for Persistent Memory,” *Proceedings of the Flash Memory Summit (FMS)*, Santa Clara, CA, August 2016.
237. Onur Mutlu, “Large-Scale Study of In-the-Field Flash Failures,” *Proceedings of the Flash Memory Summit (FMS)*, Santa Clara, CA, August 2016.
238. Onur Mutlu, “The Row Hammer Problem and Other Issues We May Face as Memory Becomes Denser,” *Proceedings of the 53rd Design Automation Conference (DAC)*, Austin, TX, USA, June 2016.
239. Jeremie Kim, Damla Senol, Hongyi Xin, Donghyuk Lee, Mohammed Alser, Hasan Hassan, Oguz Ergin, Can Alkan, Onur Mutlu, “Genome Read In-Memory (GRIM) Filter: Fast Location Filtering in DNA Read Mapping with Emerging Memory Technologies,” *Flash Talk and Poster at 20th Annual International Conference on Research in Computational Molecular Biology (RECOMB)*, Santa Monica, CA, USA, April 2016.
240. Onur Mutlu, “Read Disturb Errors in MLC NAND Flash Memory,” *Proceedings of the Flash Memory Summit (FMS)*, Santa Clara, CA, August 2015.
241. Hongyi Xin, Richard Zhu, Sunny Nahar, John Emmons, Gennady Pekhimenko, Carl Kingsford, Can Alkan, Onur Mutlu, “Optimal Seed Solver: Optimizing Seed Selection in Read Mapping,” *Poster at High Throughput Sequencing Algorithms and Applications (HITSEQ)*, Dublin, Ireland, July 2015.
242. Azita Nouri, Reha Oguz Selvitopi, Ozcan Ozturk, Onur Mutlu, Can Alkan, “Massively Parallel Mapping of Next Generation Sequence Reads Using GPUs,” *Short Student Research Competition paper at the 20th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Istanbul, Turkey, March 2015.
243. Gennady Pekhimenko, Todd C. Mowry, Onur Mutlu, “Energy-Efficient Data Compression for GPU Memory Systems,” *Short Student Research Competition paper at the 20th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Istanbul, Turkey, March 2015. **First place in ACM Student Research Competition.**
244. Onur Mutlu, “Error Analysis and Management for MLC NAND Flash Memory,” *Proceedings of the Flash Memory Summit (FMS)*, Santa Clara, CA, August 2014.
245. Bradley Thwaites, Gennady Pekhimenko, Amir Yazdanbakhsh, Jongse Park, Girish Mururu, Hadi Esmaeilzadeh, Onur Mutlu, Todd Mowry, “Rollback-Free Value Prediction with Approximate Loads,” *Short poster paper at the 23rd ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Edmonton, Alberta, Canada, August 2014.
246. Hongyi Xin, John Greth, John Emmons, Gennady Pekhimenko, Carl Kingsford, Can Alkan, Onur Mutlu, “Shifted Hamming Distance (SHD): A Fast and Accurate SIMD-Friendly Filter for Local Alignment in Read Mapping,” *Poster at High Throughput Sequencing Algorithms and Applications (HITSEQ)*, Boston, MA, July 2014.
247. Onur Mutlu, “Memory Scaling: A Systems Architecture Perspective,” *Proceedings of MemCon (MEMCON)*, Santa Clara, CA, August 2013.
248. Chuanjun Zhang, Glenn G. Ko, Jungwook Choi, Shang-nien Tsai, Minje Kim, Abner Guzman-Rivera, Rob A. Rutenbar, Paris Smaragdus, Mi Sun Park, Vijaykrishnan Narayanan, Hongyi Xin, Onur Mutlu, Bin Li, Li Zhao, Mei Chen, “EMERALD: Characterization of emerging applications and algorithms for low-power devices,” *Short poster paper at the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS)*, Austin, TX, April 2013.
249. Justin Meza, Jing Li, Onur Mutlu, “A Case for Small Row Buffers in Non-Volatile Main Memories,” *Short poster paper at the 30th IEEE International Conference on Computer Design (ICCD)*, Montreal, Quebec, Canada, September 2012.
250. Reetuparna Das, Rachata Ausavarungnirun, Onur Mutlu, Akhilesh Kumar, Mani Azimi, “Application-to-Core Mapping Policies to Reduce Memory Interference in Multi-Core Systems,” *Short poster paper at the 21st ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Minneapolis, MN, September 2012.
251. Nachiappan C. Nachiappan, Asit K. Mishra, Mahmut Kandemir, Anand Sivasubramaniam, Onur Mutlu, Chita R. Das, “Application-aware Prefetch Prioritization in On-chip Networks,” *Short poster paper at the 21st ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Minneapolis, MN, September 2012.
252. Gennady Pekhimenko, Todd C. Mowry, Onur Mutlu, “Linearly Compressed Pages: A Main Memory Compression Framework with Low Complexity and Low Latency,” *Short Student Research Competition paper at the 21st ACM International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Minneapolis, MN, September 2012. **Second place in ACM Student Research Competition.**
253. Hongyi Xin, Donghyuk Lee, Farhad Hormozdiari, Can Alkan, Onur Mutlu, “FastHASH: A New GPU-friendly Algorithm for Fast and Comprehensive Next-generation Sequence Mapping,” *Pacific Symposium on Biocomputing (PSB) Poster Session*, Hawaii, January 2012.

254. Onur Mutlu, “Some Opportunities and Obstacles in Cross-Layer and Cross-Component (Power) Management,” *Position paper and presentation at NSF Workshop on Cross-Layer Power Optimization and Management (NSF CPOM)*, Los Angeles, CA, February 2012.
255. Licheng Chen, Yongbing Huang, Yungang Bao, Onur Mutlu, Guangming Tan, Mingyu Chen, “Revisiting virtual channel memory for performance and fairness on multi-core architectures,” *Short poster paper at the 25th International Conference on Supercomputing (ICS)*, Tucson, AZ, May 2011.
256. Onur Mutlu, “Asymmetry Everywhere (with Automatic Resource Management),” *Position paper and presentation at CRA Workshop on Advancing Computer Architecture Research: Popular Parallel Programming (NSF ACAR)*, San Diego, CA, February 2010.

Significant Technical Reports (otherwise unpublished)

257. Yixin Luo, Saugata Ghose, Tianshi Li, Sriram Govindan, Bikash Sharma, Bryan Kelly, Amirali Boroumand, Onur Mutlu, “Using ECC DRAM to Adaptively Increase Memory Capacity,” arXiv.org, arXiv:1706.08870, June 2017.
258. Amirali Boroumand, Saugata Ghose, Minesh Patel, Hasan Hassan, Brandon Lucia, Nastaran Hajinazar, Kevin Hsieh, Krishna T. Malladi, Hongzhong Zheng, Onur Mutlu, “LazyPIM: Efficient Support for Cache Coherence in Processing-in-Memory Architectures,” arXiv.org, arXiv:1706.03162, June 2017.
259. Hongyi Xin, Jeremie Kim, Sunny Nahar, Can Alkan, Onur Mutlu, “LEAP: A Generalization Of The Landau-Vishkin Algorithm With Custom Gap Penalties,” biorxiv.org, <https://doi.org/10.1101/133157>, May 2017.
260. Vivek Seshadri, Donghyuk Lee, Thomas Mullins, Hasan Hassan, Amirali Boroumand, Jeremie Kim, Michael A. Kozuch, Onur Mutlu, Phillip B. Gibbons, Todd C. Mowry, “Buddy-RAM: Improving the Performance and Efficiency of Bulk Bitwise Operations Using DRAM,” arXiv.org, arXiv:1611.09988, November 2016.
261. Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji-Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, Onur Mutlu, “RowHammer: Reliability Analysis and Security Implications,” arXiv.org, arXiv:1603.00747, February 2016.
262. Kevin Chang, Gabriel H. Loh, Mithuna Thottethodi, Yasuko Eckert, Mike O’Connor, Lavanya Subramanian, Onur Mutlu, “Enabling Efficient Dynamic Resizing of Large DRAM Caches via A Hardware Consistent Hashing Mechanism,” SAFARI Technical Report, TR-SAFARI-2013-001, Carnegie Mellon University, April 2013.
263. Justin Meza, Jing Li, Onur Mutlu, “Evaluating Row Buffer Locality in Future Non-Volatile Main Memories,” SAFARI Technical Report, TR-SAFARI-2012-002, Carnegie Mellon University, December 2012.
264. Chris Fallin, Xiangyao Yu, Greg Nazario, Onur Mutlu, “A High-Performance Hierarchical Ring On-Chip Interconnect with Low-Cost Routers,” SAFARI Technical Report, TR-SAFARI-2011-007, Carnegie Mellon University, September 2011.
265. Chris Craik, Onur Mutlu, “Investigating the Viability of Bufferless NoCs in Modern Chip Multi-Processor Systems,” SAFARI Technical Report, TR-SAFARI-2011-004, Carnegie Mellon University, August 2011.
266. Chang Joo Lee, Eiman Ebrahimi, Veynu Narasiman, Onur Mutlu, Yale N. Patt, “DRAM-Aware Last-Level Cache Replacement,” HPS Technical Report, TR-HPS-2010-007, December 2010.
267. Chang Joo Lee, Veynu Narasiman, Eiman Ebrahimi, Onur Mutlu, Yale N. Patt, “DRAM-Aware Last-Level Cache Writeback: Reducing Write-Caused Interference in Memory Systems,” HPS Technical Report, TR-HPS-2010-002, April 2010.
268. Onur Mutlu, “Efficient Runahead Execution Processors,” Ph.D. Dissertation, HPS Technical Report, TR-HPS-2006-007, July 2006. **Nominated for the ACM Doctoral Dissertation Award by the University of Texas at Austin**

Patents and Invention Disclosures (filed while in Industry)

1. William H. Mangione Smith, Onur Mutlu, “Reliable Communications in On-chip Networks,” US Patent 8473818, June 2013.
2. Thomas Moscibroda, Onur Mutlu, “Bufferless Routing in On-Chip Interconnection Networks,” US Patent 8509078, August 2013.
3. Thomas Moscibroda, Onur Mutlu, “Prioritization of Multiple Concurrent Threads for Scheduling Requests to Shared Memory,” US Patent 8271741, September 2012.
4. Thomas Moscibroda, Onur Mutlu, “Coordination among Multiple Memory Controllers,” US Patent 8266393, September 2012.
5. Thomas Moscibroda, Onur Mutlu, “Controlling Interference in Shared DRAM Systems using Batch Scheduling,” US Patent 8180975, May 2012.
6. Onur Mutlu, José A. Joao, “Feedback Mechanism for Dynamic Predication of Indirect Jumps,” US Patent 7818551, October 2010.
7. José A. Joao, Onur Mutlu, “Target-Frequency based Indirect Jump Prediction for High-Performance Processors,” US Patent 7870371, January 2011.
8. Onur Mutlu, Thomas Moscibroda, “A Software-Configurable and Stall-Time Fair Memory Access Scheduling Mechanism for Shared Memory Systems,” US Patent 8245232, August 2012.
9. Thomas Moscibroda, Onur Mutlu, “Multi-level DRAM Controller to Manage Access to DRAM,” US Patent 8001338, August 2011.
10. Onur Mutlu, Thomas Moscibroda, “Parallelism-Aware Memory Request Scheduling in Shared Memory Controllers,” US Patent Application filed August 2007.
11. Thomas Moscibroda, Onur Mutlu, “Fairness in Memory Systems,” US Patent Application filed July 2007.
12. Jared Stark, Chris Wilkerson, Onur Mutlu, “Apparatus for Memory Communication During Runahead Execution,” US Patent Application filed December 2002.

13. Eric Sprangle, Onur Mutlu, “Method and Apparatus to Control Memory Accesses,” US Patent 6799257, September 2004.

Invited Talks and Lectures

(Please visit <https://people.inf.ethz.ch/omutlu/talks.htm> for some electronic copies and video links.)

1. “Accelerating Genome Analysis: A Primer on an Ongoing Journey”
 - **Keynote talk** at the 1st Workshop on Accelerator Architecture in Computational Biology and Bioinformatics (AACBB), Vienna, Austria, 24 February 2018.
 - ETH Seminar Talk on Hardware Architectures for Machine Learning, 8 March 2018.
2. “Rethinking the Computing Systems We Design” or “Rethinking the Computing Platforms We Design” or “Rethinking the Systems We Design” or “Key Design Challenges in Future Computing Platforms: Changing Our Fixed Mindsets”
 - Visions for the Future (Celebrating Yale@75) Workshop, UT-Austin, TX, 19 September 2014.
 - Chalmers University Future of Computer Architecture Workshop, Goteborg, Sweden, 9 May 2017.
 - Nvidia, Austin, TX, USA, 20 June 2017.
 - Bogazici University, Istanbul, Turkey, 3 August 2017.
 - VMware, Palo Alto, CA, USA, 14 August 2017.
3. “Memory Reliability, Security, and Beyond: Three Key Issues in Modern Systems”
 - DAC Design Automation Summer School, Austin, TX, USA, 18 June 2017.
4. “Future Computing Architectures” or “Future Computing and Genome Analysis Platforms” or “Future Computing Architectures: Key Challenges and Opportunities”
 - ETH Zurich Inaugural Lecture, Zurich, Switzerland, 15 May 2017.
 - ETH Zurich Industry Day, Zurich, Switzerland, 29 August 2017.
 - ETH Systems Group Industry Retreat, Engelberg, Switzerland, 22 January 2018.
 - ARM CPU Design Group Doughnuts Talk, Cambridge, UK, 20 March 2018.
 - **Keynote talk** at the 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP), Cambridge, UK, 21 March 2018.
5. “Solving the Memory Problem”
 - ETH Systems Group Industry Retreat, Engelberg, Switzerland, 18 January 2016.
 - ETH Systems Group Industry Retreat, Engelberg, Switzerland, 30 January 2017.
6. “Rethinking Memory System Design: Challenging Our Fixed Mindsets” or “Rethinking Memory System Design (and the Computing Platforms We Design Around It)” or “Rethinking Memory Systems: Challenges and Opportunities” or “Rethinking Memory System Design (for Data-Intensive Computing)” or “Memory Scaling: A Systems Architecture Perspective” or “Key Challenges and Opportunities in Memory Systems” (many variants on the same theme with varying content)
 - **Special invited talk** at the 5th IEEE International Memory Workshop, Monterey, CA, 27 May 2013.
 - MemCon, San Jose, CA, 6 August 2013.
 - Facebook, Menlo Park, CA, 7 August 2013.
 - Huawei, Sonoma, CA, 8 August 2013.
 - Samsung, San Jose, 9 September 2013.
 - Intel Science and Technology Center for Cloud Computing Retreat, Pittsburgh, PA, 8 November 2013.
 - IBM Research, Yorktown Heights, NY, 12 November 2013.
 - D.E. Shaw Research, New York, NY, 13 November 2013.
 - Intel Memory Hierarchy Workshop, Hillsboro, OR, 5 December 2013.
 - **Keynote talk** at Industry-Academia Partnership Stanford Cloud Workshop, Mountain View, CA, 6 December 2013.
 - Columbia University, New York, NY, 11 March 2014.
 - **Keynote talk** at Huawei Strategy and Technology Workshop, Santa Clara, CA, 18 March 2014.
 - **Keynote talk** at Industry-Academia Partnership Carnegie Mellon Cloud Workshop, Pittsburgh, PA, 4 April 2014.
 - New York University, Brooklyn, NY, 24 April 2014.
 - **Opening talk** at Future Memory Systems Workshop, Carnegie Mellon University, Pittsburgh, PA, 7 May 2014.
 - **Keynote talk** at Huawei Strategy and Technology Workshop, Shenzhen, China, 14 May 2014.
 - **Keynote talk** at the 9th International Conference on Green, Pervasive and Cloud Computing, Wuhan, China, 27 May 2014.
 - **Keynote talk** at the 25th IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP), Zürich, Switzerland, 20 June 2014.
 - ETH Zürich, Zürich, Switzerland, 23 June 2014.
 - **Plenary talk** at the Summer Supercomputing Academy, Moscow State University, Moscow, Russia, 2 July 2014.
 - Marmara University, Istanbul, Turkey, 9 July 2014.
 - Microsoft Research, Redmond, WA, 17 July 2014.
 - Google, Mountain View, CA, 23 July 2014.
 - Huawei, Santa Clara, CA, 6 August 2014.
 - International Summer School of AP Education Consortium, Hsinchu, Taiwan, 11 August 2014.
 - National Taiwan University, Taipei, Taiwan, 14 August 2014.
 - MediaTek, Hsinchu, Taiwan, 15 August 2014.
 - ARM, Austin, TX, 20 August 2014.
 - Broadcom, Santa Clara, CA, 21 August 2014.
 - SanDisk, San Jose, CA, 22 August 2014.
 - Micron, Boise, ID, 26 September 2014.
 - Hanyang University, Seoul, Korea, 21 October 2014.

- SK Hynix, Korea, 23 October 2014.
- Samsung, Korea, 24 October 2014.
- Carnegie Mellon University Parallel Data Laboratory Retreat Opening Talk, 27 October 2014.
- Northwestern University, Evanston, IL, 10 November 2014.
- **Keynote talk** at the 4th Workshop on Irregular Applications: Architectures and Algorithms, held with Supercomputing (SC), New Orleans, LA, 17 November 2014.
- Bogazici University, Istanbul, Turkey, 6 January 2015.
- Koc University, Istanbul, Turkey, 13 March 2015.
- Intel, Hillsboro, OR, 5 May 2015.
- ETH Zürich, Zürich, Switzerland, 11 May 2015.
- **Keynote talk** at 11th International Workshop on Data Management on Novel Hardware (DaMoN) with SIGMOD, Melbourne, Australia, 1 June 2015.
- University of Melbourne, Melbourne, Australia, 2 June 2015.
- Australian National University, Canberra, Australia, 3 June 2015.
- University of New South Wales and National Information and Communications Technology Research Centre of Excellence, Sydney, Australia, 4 June 2015.
- HP Labs, Palo Alto, CA, 25 June 2015.
- **Special invited talk** at ISC High Performance (International Supercomputing Conference), Frankfurt, Germany, 14 July 2015.
- SAP, Walldorf, Germany, 16 July 2015.
- **Keynote talk** at SAMOS XV (15th International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation), Samos, Greece, 20 July 2015.
- Bilkent University, Ankara, Turkey, 29 July 2015.
- Turku Center for Computer Science, Turku, Finland, 18 August 2015.
- KTH - Swedish Royal Institute of Technology, Stockholm, Sweden, 24 August 2015.
- Yale University, New Haven, CT, 10 September 2015.
- University of British Columbia, Vancouver, BC, Canada, 17 September 2015.
- Apple, Inc., Cupertino, CA, 22 September 2015.
- Stanford University (Systems Seminar), Palo Alto, CA, 23 September 2015.
- **Keynote talk** at the 18th International Symposium on Computer Architecture and Digital Systems (CADS), 7 October 2015.
- **Keynote talk** at the 27th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), Florianopolis, Brazil, 21 October 2015.
- **Keynote talk** at the 5th IEEE Circuits and Systems Society Workshop (CASS), Porto Alegre, Brazil, 22 October 2015.
- University of Campinas (UNICAMP), Campinas, Brazil, 23 October 2015.
- SAP, Dublin, CA, 6 November 2015.
- VMware, Palo Alto, CA, 10 November 2015.
- University of Chicago, Chicago, IL, 16 November 2015.
- Kadir Has University, Istanbul, Turkey, 23 December 2015.
- TOBB Economics and Technology University, Ankara, Turkey, 4 January 2016.
- Istanbul Technical University, Istanbul, Turkey, 6 January 2016.
- VMware, Palo Alto, CA, 3 March 2016.
- **Distinguished lecture** at Triangle Computer Science Distinguished Lecture Series, Raleigh, NC, 11 April 2016.
- University of California, Irvine, CA, 22 April 2016.
- ARM, Austin, TX, 10 June 2016.
- Intel, Austin, TX, 10 June 2016.
- **Keynote talk** at the ACM SIGPLAN International Symposium on Memory Management (ISMM), Santa Barbara, CA, 14 June 2016.
- Stony Brook University, Stony Brook, NY, USA, 13 July 2016.
- SK Telecom, Seoul, Korea, 16 August 2016.
- **Keynote talk** at the First ARM Research Summit, Cambridge, UK, 15 September 2016.
- **Keynote talk** at the 2nd Workshop on Mobile System Technologies (MST), Milan, Italy, 23 September 2016.
- **Keynote talk** at the 27th International Symposium on Rapid System Prototyping (RSP), Pittsburgh, PA, USA, 6 October 2016.
- **Keynote talk** at the 2nd EAI International Conference on Future Access Enablers of Ubiquitous and Intelligent Infrastructures (FAB-ULOUS), Belgrade, Serbia, 24 October 2016.
- IBM Research Zürich, Ruschlikon, Switzerland, 2 November 2016.
- Xilinx, San Jose, CA, 8 December 2016.
- Intel Labs, Santa Clara, CA, 9 December 2016.
- TU Dresden, Dresden, Germany, 27 February 2017.
- **Special Invited talk** at the Dagstuhl Seminar “Databases on Future Hardware”, Germany, 9 March 2017.
- **Keynote talk** at the 13th International Symposium on Applied Reconfigurable Computing (ARC), Delft, Netherlands, 4 April 2017.
- Google, Mountain View, CA, 18 May 2017.
- Technion Seiden Workshop on Beyond CMOS Technologies, Haifa, Israel, 5 June 2017.
- Intel Israel, Haifa, Israel, 8 June 2017.
- Huawei Sweden Future of Wireless Summit, Stockholm, Sweden, 13 June 2017.
- Huawei Sweden Baseband SoC Group, Stockholm, Sweden, 14 June 2017.
- IBM Research, Austin, TX, USA, 22 June 2017.
- SAP Hana Days, Tegensee, Germany, 26 August 2017.
- **Keynote talk** at the IDEA League Transiently Powered Computing Doctoral School, Delft, Netherlands, 9 November 2017.
- **Distinguished lecture** at INESC-ID Distinguished Lecture Series, University of Lisbon, Portugal, 4 December 2017.
- **Special Invited talk** at the International Workshop on Many-Core Computing: Hardware and Software, Southampton, UK, 17 January 2018.

7. “Rethinking Memory System Design (along with Interconnects)”
 - **Keynote talk** at the 8th International Workshop on Network on Chip Architectures (NoCArc), Honolulu, Hawaii, 5 December 2015.
8. Short Course on “Memory Systems” or “Memory Systems in the Multi-Core Era (An Accelerated Course)” or “Scalable Memory Systems in the Multi-Core Era” or “Rethinking Memory System Design”
 - Beihang University, 3-day Lecture Series (16 hours), Beijing, China, June 17, 25, 26, 2012.
 - Seoul National University, 2-day Lecture Series (6 hours), Seoul, Korea, June 18, 20, 2012. (Memory Scaling and Memory QoS)
 - Bogazici University, 3-day Lecture Series (15 hours), Istanbul, Turkey, June 13, 14, 17, 2013.
 - INRIA Rennes, 3-day Lecture Series (6 hours), July 4, 8, 9, 2013.
 - HiPEAC ACACES Summer School (6 hours), Fiuggi, Italy, July 15-19, 2013.
 - Turku Center for Computer Science (16 hours), Turku, Finland, August 19-21, 2015.
 - Pohang Institute of Science and Technology - POSTECH (7 hours), Pohang, Korea, 17-18 August 2016.
 - Korea Advanced Institute of Science and Technology - KAIST (7 hours), Daejeon, Korea, 19 August 2016.
 - Technion EU Cost Action Training School - Technion (5 hours), Haifa, Israel, 7 June 2017.
9. “Achieving the Next Breakthroughs in NAND Flash Memory: Changing Our Fixed Mindsets”
 - Flash Memory Summit, 10 August 2017.
10. “Some Recent Works (with a Focus on Memory)”
 - Intel, Austin, TX, USA, 21 June 2017.
 - ARM, Austin, TX, USA, 23 June 2017.
 - University of British Columbia, Vancouver, BC, Canada, 14 July 2017.
11. “Processing Data Where It Makes Sense: Enabling In-Memory Computation” or “Enabling Practical, Efficient and Large-Scale Computation Near Data”
 - Bogazici University, Istanbul, Turkey, 6 August 2015.
 - Google, Mountain View, CA, 9 November 2015.
 - Intel, Hillsboro, OR, 10 January 2017.
 - Microsoft Research Swiss Joint Research Center Workshop, Cambridge, UK, 2 February 2017.
 - **Keynote talk** at the 3rd Workshop on Mobile System Technologies (MST), Milan, Italy, 27 October 2017.
 - Semiconductor Research Corporation E-Workshop, online, 1 November 2017.
 - Institute of Neuroinformatics, ETH and University of Zurich, Zurich, Switzerland, 1 December 2017.
 - Chalmers University, Gothenburg, Sweden, 19 December 2017.
 - Microsoft Research Swiss Joint Research Center Workshop, Lausanne, Switzerland, 19 February 2018.
 - **Keynote talk** at the 3rd International Workshop on Emerging Memory Solutions and Applications (held with DATE), Dresden, Germany, 23 March 2018.
12. “Enabling Practical, Efficient, and Large-Scale Computation Near Data”
 - Microsoft Research, Cambridge, UK, 2 February 2017.
13. “Opportunities and Challenges of New Memory Technologies”
 - Samsung MRAM Global Innovation Forum, San Jose, CA, 11 November 2015.
 - ARM Research Summit, Cambridge, UK, 11 September 2017.
14. “Reliability and Security Issues of DRAM and NAND Flash Scaling”
 - Memory Reliability Forum at HPCA, Barcelona, Spain, 13 March 2016.
 - SNU International Workshop on Recent Advances in Neural Networks and Non-Volatile Memory, Seoul, Korea, 16 August 2016.
15. “The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser” or “The DRAM RowHammer Problem (and Its Reliability and Security Implications)”
 - Carnegie Mellon University CyLab Partners Conference, Pittsburgh, PA, 30 September 2015.
 - JEDEC RowHammer Task Force, 2 December 2015.
 - **Special Invited Talk** at the Design Automation Conference (DAC), 9 June 2016.
 - **Special Invited Talk** at the Design, Automation, and Test in Europe Conference (DATE), 30 March 2017.
 - **Keynote talk** at the Workshop on Resilient Systems, TU Dresden, Dresden, Germany, 21 April 2017.
 - Vrije University, Amsterdam, The Netherlands, 10 November 2017.
16. “ThyNVM: Software-Transparent Crash Consistency for Persistent Memory”
 - Flash Memory Summit, Santa Clara, CA, 8 August 2016.
 - Google, Mountain View, CA, 22 August 2016.
17. “Large-Scale Study of In-the-Field Flash Failures”
 - Flash Memory Summit, Santa Clara, CA, 10 August 2016.
18. “Read Disturb Errors in MLC NAND Flash Memory”
 - Flash Memory Summit, Santa Clara, CA, 12 August 2015.
19. “A Case for Autonomous Memory”
 - Google Systems Research Workshop, San Francisco, CA, 26 June 2015.
20. “Architectural Frameworks to Facilitate Practical, Efficient and Transparent Computation Near Data”
 - SRC (Semiconductor Research Corporation) Annual Review, Ann Arbor, MI, USA, 31 May 2017.
21. “Integrated Techniques for Scalable Management of Main Memory Performance, Energy, and QoS”
 - SRC (Semiconductor Research Corporation) Annual Review, Intel, Hillsboro, OR, 1 May 2013.

- SRC (Semiconductor Research Corporation) Annual Review, Carnegie Mellon University, Pittsburgh, PA, 6 May 2014.
 - SRC (Semiconductor Research Corporation) Annual Review, Intel, Hillsboro, OR, 6 May 2015.
 - SRC (Semiconductor Research Corporation) Annual Review, Intel, Hillsboro, OR, 29 April 2016.
 - Google, Mountain View, CA, 20 July 2016.
22. “Error Analysis and Management for MLC NAND Flash Memory”
 - Flash Memory Summit (45-minute talk), Santa Clara, CA, 7 August 2014.
 23. “Some New Ideas in Memory System Design for Data-Intensive Computing”
 - Intel Science and Technology Center for Cloud Computing Retreat, Hillsboro, OR, 4 September 2014.
 24. “A Fresh Look at DRAM Architecture: New Techniques to Improve DRAM Latency, Parallelism, and Energy Efficiency”
 - MIT, Cambridge, MA, 22 May 2013.
 - INRIA Rennes, 4 July 2013.
 - Intel Memory Hierarchy Workshop, Hillsboro, OR, 13 March 2014.
 25. “Enabling Low-Latency DRAM Architectures”
 - Intel Low Latency Architectures Workshop, Hillsboro, OR, 10 January 2018.
 - Intel Low Latency Architectures Workshop, Hillsboro, OR, 2 December 2016.
 - Intel Low Latency Architectures Workshop, Hillsboro, OR, 11 December 2015.
 - Intel Low Latency Architectures Workshop, Hillsboro, OR, 7 November 2014.
 26. “Understanding and Overcoming Challenges of DRAM Refresh”
 - Extreme Scale Scientific Computing Workshop, Moscow, Russia, 30 June 2014.
 27. “Multi-core Architectures and Shared Resource Management: Fundamentals and Recent Research”
 - Seoul National University, Lecture Series (12 hours), Seoul, Korea, July 6-9, 2010.
 - Korea Advanced Institute of Science and Technology, Global Lecture Series (15 hours), Daejeon, Korea, July 26-29, 2010.
 - Beihang University, 2-day Lecture Series (13 hours), Beijing, China, August 9, 12, 2011.
 - Bogazici University, 3-day Lecture Series (15 hours), Istanbul, Turkey, June 6, 7, 10, 2013.
 28. “Scaling the Main Memory System in the Many-Core Era” or “Main Memory Scaling: Some Ideas to Improve DRAM and Enable Hybrid Memory Systems” or “Main Memory Scaling: Some Challenges and Solution Directions”
 - IBM, Poughkeepsie, NY, 21 May 2012.
 - Invited Talk, ACM Design Automation Conference More Than Moore Technologies Workshop, San Francisco, CA, 3 June 2012.
 - Samsung Information Systems America, San Jose, CA, 4 June 2012.
 - Rambus, Sunnyvale, CA, 5 June 2012.
 - POSTECH, Pohang, Korea, 19 June 2012.
 - Samsung, Memory Division, Hwasung City, Korea, 21 June 2012.
 - SK Hynix, Seoul, Korea, 23 June 2012.
 - **Distinguished lecture** at Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, 28 June 2012.
 - Intel, Hillsboro, OR, 26 July 2012.
 - Advanced Micro Devices, Austin, TX, 21 September 2012.
 - McGill University, Montreal, Quebec, Canada, 1 October 2012.
 - Sabanci University, Istanbul, Turkey, 4 January 2013.
 - Nvidia, Austin, TX, 22 March 2013.
 - EMC, Hopkinton, MA, 20 May 2013.
 - Barcelona Supercomputing Center and Universitat Politècnica de Catalunya, Barcelona, Spain, 26 July 2013.
 - Nvidia, Santa Clara, CA, 5 August 2013.
 29. “Architecting and Exploiting Asymmetry in Multi-Core Architectures”
 - Intel, Santa Clara, CA, 9 March 2011.
 - Rambus, Sunnyvale, CA, 5 June 2012.
 - Samsung, System LSI Division, Hwasung City, Korea, 22 June 2012.
 - Intel, Hillsboro, OR, 26 July 2012.
 - Advanced Micro Devices, Bellevue, WA, 3 August 2012.
 - Intel Archfest, Hillsboro, OR, 10 August 2012.
 - Intel Science and Technology Center on Cloud Computing Board of Advisors Meeting, Pittsburgh, PA, 16 August 2012.
 - Intel Science and Technology Center on Cloud Computing Retreat, Pittsburgh, PA, 29 November 2012.
 - Bogazici University, Istanbul, Turkey, 26 December 2012.
 - Bilkent University, Ankara, Turkey, 28 December 2012.
 - EMC, Hopkinton, MA, 20 May 2013.
 - TUBITAK, Gebze, Turkey, 19 June 2013.
 - INRIA Rennes, Rennes, France, 2 July 2013.
 - Barcelona Supercomputing Center and Universitat Politècnica de Catalunya, Barcelona, Spain, 23 July 2013.
 - International Summer School of AP Education Consortium, Hsinchu, Taiwan, 11 August 2014.
 - International Workshop on Heterogeneous Computing Platforms, held with ICCAD, San Jose, CA, 6 November 2014.
 30. “Some Ideas in Designing Scalable and Efficient Multi-Core Systems”
 - Apple, Cupertino, CA, 6 June 2012.
 31. “Concurrent Autonomous Self-Test for Uncore Components in SoCs”
 - SK Hynix, Seoul, Korea, 23 June 2012.

32. “Designing QoS-Aware Memory Systems” or “Predictable Memory Systems for the Many-Core Era”
 - IBM, Poughkeepsie, NY, 21 May 2012.
 - SK Hynix, Seoul, Korea, 23 June 2012.
 - Intel, Hillsboro, OR, 1 August 2012.
 - VMware, Palo Alto, CA, 17 October 2012.
33. “Main Memory Issues (Both Volatile and Non-Volatile)”
 - Carnegie Mellon University Parallel Data Lab Retreat, Bedford, PA, 6 November 2012.
 - Carnegie Mellon University Parallel Data Lab Visit Day, Pittsburgh, PA, 11 May 2012.
34. “Some (Security-Related) Challenges in Future Computing Platforms”
 - Carnegie Mellon University College of Engineering Deans Council Meeting, Pittsburgh, PA, 26 April 2012.
35. “Some Opportunities and Obstacles in Cross-Layer and Cross-Component (Power) Management”
 - NSF Workshop on Cross-Layer Power Optimization and Management, Los Angeles, CA, 10 February 2012.
36. “Memory Systems in the Many-Core Era: Challenges, Opportunities, and Solution Directions”
 - **Joint Keynote Talk**, International Symposium on Memory Management (ISMM) and ACM Workshop on Memory System Performance and Correctness (MSPC), San Jose, CA, 5 June 2011.
 - Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, 10 August 2011.
 - Tsinghua University, Beijing, China, 11 August 2011.
 - Microsoft Research Asia, Beijing, China, 15 August 2011.
 - Massachusetts Institute of Technology, 4 November 2011.
 - University of California at Berkeley (PARLAB seminar), 15 November 2011.
37. “Issues in DRAM and NVM based Main Memory”
 - Carnegie Mellon University Parallel Data Lab Retreat, Bedford Springs, PA, 8 November 2011.
38. “Computer Performance”
 - Carnegie Mellon University ECE Department Advisory Board Presentation, 21 September 2011.
39. “Application-Aware Memory Controllers” or “Thread Cluster Memory Scheduling”
 - Xilinx Labs, San Jose, CA, 8 June 2011.
 - Gigascale Systems Research Center Mid-Year Review, Yorktown Heights, NY, 27 May 2011.
40. “Architecture and System-Level Challenges Related to Memory”
 - Focus Center Research Program Memory Cross-Cut Workshop, Cambridge, MA, 12 May 2011.
41. “Towards Practical Bufferless On-Chip Networks”
 - Intel, Santa Clara, CA, 9 March 2011.
42. “Some Ideas for Efficient and High-Performance Core Design”
 - Intel Corporation ARO Swiss Army Processor Workshop, Hillsboro, OR, 29 April 2011.
43. “PCM (NVM) as Main Memory: Opportunities and Challenges”
 - Carnegie Mellon University, Parallel Data Lab Retreat, Pittsburgh, PA, October 25, 2010.
44. “Research Challenges in Future Computing Platforms”
 - Carnegie Mellon University, ECE Department Faculty Retreat, Wheeling, WV, August 12, 2010.
 - Carnegie Mellon University, Sophomore Electrical Engineering Seminar, September 23, 2010.
 - Carnegie Mellon University, ECE/SCS Alumni Event, Austin, TX, March 24, 2013.
45. “End-to-end QoS-aware, High-Performance and Customizable Many-Core Memory Systems”
 - Intel Memory Hierarchy Meeting, Hillsboro, OR, 8 October 2010.
46. “Rethinking Core Design in the Power-Constrained Many-Core Era”
 - Intel Core Workshop, Hillsboro, OR, 27 September 2010.
47. “Some Ideas for ILP Research”
 - CRA Workshop on Advancing Computer Architecture Research, Seattle, WA, 20 September 2010.
48. “Designing High-Performance and Fair Shared Multi-core Memory Systems: Two Approaches” or “QoS-Aware Multi-Core Memory System Management”
 - Gigascale Systems Research Center E-Seminar, 23 March 2010.
 - Pennsylvania State University, CSE Colloquium, 26 March 2010.
 - ARM, Inc., Austin, TX, 8 April 2010.
 - Advanced Micro Devices, Austin, TX, 9 April 2010.
 - Microsoft Research, Redmond, WA, 27 April 2010.
 - HP Laboratories, Palo Alto, CA, 25 May 2010.
 - VMware, Palo Alto, CA, 26 May 2010.
 - Intel Corporation, Hillsboro, OR, 27 May 2010.
 - Gigascale Systems Research Center Annual Review, San Jose, CA, 29 September 2010.
 - Intel Corporation ArchFest, Hillsboro, OR, 8 October 2010.
 - ASPLOS 2011 Program Committee Symposium, Pittsburgh, PA, 22 October 2010.
 - Advanced Micro Devices, Sunnyvale, CA, 10 March 2011.

- Oracle, Redwood Shores, CA, 11 March 2011.
 - Princeton University, Princeton, NJ, 18 March 2011.
49. “Rethinking Memory System Design in the Nanoscale Many-Core Era”
 - Intel Memory Hierarchy Workshop, Hillsboro, OR, 22 January 2010.
 - ASPLOS Workshop on Architecting Memory Technologies, Pittsburgh, PA, 14 March 2010.
 50. “Asymmetry Everywhere (with Automatic Resource Management)”
 - CRA Workshop on Advancing Computer Architecture Research, San Diego, CA, 22 February 2010.
 51. “Preventing Memory Performance Attacks in Multi-Core Systems”
 - ECE Seminar, Carnegie Mellon University, 5 February 2009.
 - Massachusetts Institute of Technology, 23 April 2008.
 - Carnegie Mellon University, 15 April 2008.
 52. “Parallelism-aware Batch Scheduling: Enhancing both Performance and Fairness of Shared DRAM Systems”
 - IBM Austin Research Laboratory, Austin, TX, 19 June 2009.
 - Advanced Micro Devices Research Lab, Redmond, WA, 6 March 2009.
 - Beihang University, Beijing, China, 21 June 2008.
 53. “ATLAS: A Scalable and High-Performance Scheduling Algorithm for Multiple Memory Controllers”
 - Advanced Micro Devices Research Lab, Redmond, WA, October 2009.
 - Freescale Semiconductor, Austin, TX, 8 April 2010.
 54. “Memory Performance Attacks and Fair Memory Scheduling”
 - University of British Columbia and IEEE Computer Society, Vancouver, BC, Canada, 6 March 2008.
 - ASPLOS PC Meeting Research Seminar, Microsoft Research, 18 October 2007.
 - Multi-Core Virtual Team Meeting, Microsoft, 5 October 2007.
 55. “MSR Computer Architecture Group: Vision and Projects”
 - Presentation to Rico Malvar, MSR-Redmond Director, Microsoft Research, 12 December 2007.
 56. “Hardware-Based Devirtualization of Virtual Function Calls”
 - MSR Systems and Networking Seminar, Microsoft Research, 7 December 2006.
 57. “Runahead Execution and AVD Prediction: A Power-efficient Processing Paradigm for Tolerating Long Main Memory Latencies”
 - University of Illinois Urbana-Champaign, Computer Engineering Seminar, Urbana, IL, USA, 23 January 2007.
 - Xilinx Labs, San Jose, CA, USA, 16 June 2006.
 - Microsoft Research, Redmond, WA, USA, 12 June 2006.
 - Stanford University, Department of EE, Computer Architecture Seminar, Stanford, CA, USA, 7 June 2006.
 - MIPS Technologies, Mountain View, CA, USA, 6 June 2006.
 - IBM T.J. Watson Research Center, Yorktown Heights, NY, USA, 1 June 2006.
 - Carnegie Mellon University, Department of ECE, CALCM Seminar, Pittsburgh, PA, USA, 30 May 2006.
 - Hewlett-Packard Laboratories, Palo Alto, CA, USA, 4 May 2006.
 - University of California, San Diego, Department of Computer Science and Engineering, CA, USA, 21 April 2006.
 - University of Texas at Austin, Department of ECE, Guest Lecture for EE382N (Microarchitecture), 11-12 April 2006.
 58. “Efficient Runahead Execution”
 - Advanced Micro Devices, Sunnyvale, CA, USA, May 2005.
 - Intel Barcelona Research Center, Barcelona, Spain, November 2005.
 59. “Runahead Execution”
 - Advanced Micro Devices, Sunnyvale, CA, USA, May 2004.
 - Instituto de Informatica, Universidade Federal Rio Grande do Sul (UFRGS), Porto Alegre, Brazil, October 2004.
 60. “Runahead Execution: A Mechanism to Approximate the Performance of Large Instruction Windows”
 - Enterprise Platforms Group, Intel Corporation, Santa Clara, CA, USA, August 2002.
 - Desktop Platforms Group, Intel Corporation, Hillsboro, OR, USA, August 2002.

Conference Talks (including some Keynote, Invited, and Plenary Speeches)

61. “Processing Data Where It Makes Sense: Enabling In-Memory Computation *Keynote talk at International Workshop on Emerging Memory Solutions and Applications (held with DATE)*, Dresden, Germany, March 2018.
62. “Future Computing Architectures,” *Keynote talk at the 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*, Cambridge, UK, March 2018.
63. “Accelerating Genome Analysis: A Primer on an Ongoing Journey,” *Keynote talk at the 1st Workshop on Accelerator Architecture in Computational Biology and Bioinformatics (AACBB)*, Vienna, Austria, February 2018.
64. “Rethinking Memory System Design (and the Computing Platforms We Design Around It),” *Special invited talk at the International Workshop on Many-Core Computing: Hardware and Software*, Southampton, UK, January 2018.
65. “Processing Data Where It Makes Sense: Enabling In-Memory Computation *Keynote talk at 3rd Workshop on Mobile System Technologies (MST)*, Milan, Italy, October 2017.
66. “Opportunities and Challenges of Emerging Memory Technologies” *Invited talk at the ARM Research Summit*, Cambridge, UK, September 2017.

67. "Achieving the Next Breakthroughs in NAND Flash Memory: Changing Our Fixed Mindsets" *Flash Memory Summit*, Santa Clara, CA, August 2017.
68. "Rethinking Memory System Design (and the Computing Platforms We Design Around It)," *Keynote talk at the 13th International Symposium on Applied Reconfigurable Computing (ARC)*, Delft, Netherlands, 4 April 2017.
69. "The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser," *Special invited talk and paper at Design, Automation and Test in Europe Conference (DATE)*, Lausanne, Switzerland, 30 March 2017.
70. "Rethinking Memory System Design: Business As Usual in the Next Decade?" *Keynote talk at 2nd EAI International Conference on Future Access Enablers of Ubiquitous and Intelligent Infrastructures (FABULOUS)*, Belgrade, Serbia, October 2016.
71. "Rethinking Memory System Design" *Keynote talk at 27th International Symposium on Rapid System Prototyping (RSP)*, Pittsburgh, PA, USA, October 2016.
72. "Rethinking Memory System Design: Business As Usual in the Next Decade?" *Keynote talk at 2nd Workshop on Mobile System Technologies (MST)*, Milan, Italy, September 2016.
73. "Rethinking Memory System Design: Business As Usual in the Next Decade?" *Keynote talk at the First ARM Research Summit*, Cambridge, UK, September 2016.
74. "Large-Scale Study of In-the-Field Flash Failures," *Flash Memory Summit*, Santa Clara, CA, August 2016.
75. "ThyNVM: Software-Transparent Crash Consistency for Persistent Memory," *Flash Memory Summit*, Santa Clara, CA, August 2016.
76. "Rethinking Memory System Design," *Keynote talk at the ACM SIGPLAN International Symposium on Memory Management (ISMM)*, Santa Barbara, CA, 14 June 2016.
77. "The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser," *Special invited talk and paper at Design Automation Conference (DAC)*, 9 June 2016.
78. "Reliability and Security Issues of DRAM and NAND Flash Scaling," *Invited talk at the Memory Reliability Forum at HPCA*, Barcelona, Spain, March 2016.
79. "Rethinking Memory System Design (along with Interconnects)," *Keynote talk at the 8th International Workshop on Network on Chip Architectures (NoCArc)*, Honolulu, Hawaii, December 2015.
80. "Rethinking Memory System Design (for Data-Intensive Computing)," *Keynote talk at the 5th IEEE Circuits and Systems Society Workshop (CASS)*, Porto Alegre, Brazil, October 2015.
81. "Rethinking Memory System Design (for Data-Intensive Computing)," *Keynote talk at the 27th International Symposium on Computer Architecture and High Performance Computing (SBAC-PAD)*, Florianopolis, Brazil, October 2015.
82. "Rethinking Memory System Design (for Data-Intensive Computing)," *Keynote talk at the 18th International Symposium on Computer Architecture and Digital Systems (CADS)*, October 2015.
83. "Read Disturb Errors in MLC NAND Flash Memory," *Flash Memory Summit*, Santa Clara, CA, August 2015.
84. "Rethinking Memory System Design (for Data-Intensive Computing)," *Keynote talk at SAMOS XV (15th International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation)*, Samos, Greece, July 2015.
85. "Rethinking Memory System Design (for Data-Intensive Computing)," *Special invited talk at ISC High Performance*, Frankfurt, Germany, July 2015.
86. "Rethinking Memory System Design (for Data-Intensive Computing)," *Keynote talk at 11th International Workshop on Data Management on Novel Hardware (DaMoN), held with SIGMOD*, Melbourne, Australia, June 2015.
87. "Rethinking Memory System Design (for Data-Intensive Computing)," *Keynote talk at the 4th Workshop on Irregular Applications: Architectures and Algorithms, held with Supercomputing (SC)*, New Orleans, LA, November 2014.
88. "Architecting and Exploiting Asymmetry in Multi-Core Architectures," *Invited talk at the International Workshop on Heterogeneous Computing Platforms, held with ICCAD*, San Jose, CA, November 2014.
89. "The Heterogeneous Block Architecture," *32nd IEEE International Conference on Computer Design*, Seoul, South Korea, October 2014.
90. "Error Analysis and Management for MLC NAND Flash Memory," *Flash Memory Summit*, Santa Clara, CA, August 2014.
91. "Rethinking Memory System Design (for Data-Intensive Computing)," *Plenary Talk at the Summer Supercomputing Academy*, Moscow, Russia, July 2014.
92. "Understanding and Overcoming Challenges of DRAM Refresh," *Extreme Scale Scientific Computing Workshop*, Moscow, Russia, June 2014.
93. "Rethinking Memory System Design for Data-Intensive Computing," *Keynote talk at the 25th IEEE International Conference on Application-specific Systems, Architectures and Processors*, Zürich, Switzerland, June 2014.
94. "Rethinking Memory/Storage System Design for Data-Intensive Computing," *Keynote talk at the 9th International Conference on Green, Pervasive and Cloud Computing*, Wuhan, China, May 2014.
95. "Rethinking Memory System Design for Data-Intensive Computing," *Keynote talk at the Industry-Academia Partnership Cloud Workshop*, Pittsburgh, PA, April 2014.
96. "Rethinking Memory System Design for Data-Intensive Computing," *Keynote talk at the Industry-Academia Partnership Cloud Workshop*, Mountain View, CA, December 2013.
97. "Memory Scaling: A Systems Architecture Perspective," *MemCon 2013*, Santa Clara, CA, August 2013.

98. "An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms," *40th International Symposium on Computer Architecture*, Tel-Aviv, Israel, June 2013.
99. "A Case for Efficient Hardware-Software Cooperative Management of Storage and Memory," *5th Workshop on Energy-Efficient Design*, Tel-Aviv, Israel, June 2013.
100. "Memory Scaling: A Systems Architecture Perspective," *Invited talk at the 5th International Memory Workshop*, Monterey, CA, May 2013.
101. "Scalable Memory, Compute and Communication Architectures," *Industry-Academia Partnership Cloud Workshop*, Cambridge, MA, May 2013.
102. "Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime," *30th International Conference on Computer Design*, Montreal, Quebec, Canada, September 2012.
103. "Bottleneck Identification and Scheduling in Multithreaded Applications," *17th International Conference on Architectural Support for Programming Languages and Operating Systems*, London, UK, March 2012.
104. "Data Marshaling for Multi-core Architectures," *37th International Symposium on Computer Architecture*, St. Malo, France, June 2010.
105. "Parallelism-aware Batch Scheduling: Enhancing both Performance and Fairness of Shared DRAM Systems," *35th International Symposium on Computer Architecture*, Beijing, China, June 2008.
106. "Stall-Time Fair Memory Access Scheduling," *40th International Symposium on Microarchitecture*, Chicago, IL, USA, December 2007.
107. "Memory Performance Attacks: Denial of Memory Service in Multi-Core Systems," *16th USENIX Security Symposium*, Boston, MA, USA, August 2007.
108. "Address-Value Delta Prediction," *38th International Symposium on Microarchitecture*, Barcelona, Spain, November 2005.
109. "Techniques for Efficient Processing in Runahead Execution Engines," *32nd International Symposium on Computer Architecture*, Madison, WI, USA, June 2005.
110. "Wrong Path Events and Their Application to Early Misprediction Detection and Recovery," *37th International Symposium on Microarchitecture*, Portland, OR, USA, December 2004.
111. "Cache Filtering Techniques to Reduce the Negative Impact of Useless Speculative Memory References on Processor Performance," *16th Symposium on Computer Architecture and High Performance Computing*, Foz do Iguacu, Brazil, October 2004.
112. "Runahead Execution: An Alternative to Very Large Instruction Windows for Out-of-order Processors," *9th International Conference on High Performance Computer Architecture*, Anaheim, CA, USA, February 2003.

Panel Talks and Discussions

113. "It's time: top-tier academic computer system venues should mandate authors to make their data and code publicly available upon publication" *Panel at the 20th International Conference on Architectural Support for Programming Languages and Operating Systems*, Istanbul, Turkey, March 2015.
114. "Memory and System Balance," *Panel at 4th Workshop on Irregular Applications: Architectures and Algorithms, held with Supercomputing (SC)*, New Orleans, LA, November 2014.
115. "Meeting the Future Needs of the Data Center," *Panel at the Industry-Academia Partnership Cloud Workshop*, Cambridge, MA, May 2013.

Teaching Experience

(Please visit <https://people.inf.ethz.ch/omutlu/teaching.html> for online course materials – free for all my courses).

(Please visit <https://www.youtube.com/user/cmu18447/> for online lecture videos for my courses at Carnegie Mellon University.)

(Please visit <https://www.youtube.com/channel/UCIwQ8uOeRFgOEvBLYc3kc3g> for online lecture videos for my courses at ETH Zürich.)

ETH Zürich, Computer Science Department, Instructor for Digital Circuits and Computer Architecture (Freshman Course), *Spring 2018*

ETH Zürich, Computer Science Department, Instructor for Computer Architecture (Senior/Masters-level Course), *Fall 2017*

ETH Zürich, Computer Science Department, Co-Instructor for Digital Circuits and Computer Architecture (Freshman Course), *Spring 2017*

ETH Zürich, Computer Science Department, Co-Instructor, Graduate Seminar on Hardware Architectures for Machine Learning, *Spring 2017, 2018*

ETH Zürich, Computer Science Department, Co-Instructor for Graduate Seminar on Hardware Acceleration for Data Processing, *Fall 2016, 2017*

Carnegie Mellon University, ECE Department, Instructor for Undergraduate Computer Architecture Course ECE-447, *Spring 2012, 2013, 2014, 2015*

- Instructor rating: 4.12/5.0 (S12, 34/47), 4.33/5.0 (S13, 21/30), 4.55/5.0 (S14, 33/37), 4.54/5.0 (S15, 24/30)

- Spring 2015 online course materials, including lecture videos: <http://www.ece.cmu.edu/~ece447/s15/>

- Spring 2014 online course materials, including lecture videos: <http://www.ece.cmu.edu/~ece447/s14/>

- Spring 2013 online course materials, including lecture videos: <http://www.ece.cmu.edu/~ece447/s13/>

Carnegie Mellon University, ECE Department, Instructor for Graduate Computer Architecture Course CS/ECE-740, *Fall 2010, 2011, 2013, 2015*

- Instructor rating: 4.42/5.0 (F10, 43/53), 4.7/5.0 (F11, 19/24), 4.29/5.0 (F13, 28/33), 4.58/5.0 (F15, 12/18)

- Fall 2013 online course materials, including lecture videos: <http://www.ece.cmu.edu/~ece740/f15/>

- Fall 2013 online course materials, including lecture videos: <http://www.ece.cmu.edu/~ece740/f13/>

- Fall 2011 online course materials: <http://www.ece.cmu.edu/~ece740/f11/>

- Fall 2010 online course materials: <http://www.ece.cmu.edu/~ece740/f10/>

Carnegie Mellon University, ECE Department, Instructor for Parallel Computer Architecture Course ECE-742, *Spring 2010, 2011, Fall 2012, 2014*

- Instructor rating: 4.44/5.0 (S10, 18/20), 4.71/5.0 (S11, 7/10), 4.22/5.0 (F12, 9/14), 5.00/5.00 (F14, 5/5)
- Fall 2014 online course materials, including lecture videos: <http://www.ece.cmu.edu/~ece742/f14/>
- Fall 2012 online course materials, including lecture videos: <http://www.ece.cmu.edu/~ece742/f12/>
- Spring 2011 online course materials: <http://www.ece.cmu.edu/~ece742/2011spring/>

Carnegie Mellon University, ECE Department, Instructor for Advanced Computer Architecture Course ECE-741, *Spring 2009*

- Instructor rating based on student evaluations: 4.6/5.0 (25/35 students responded)

University of Texas at Austin, ECE Department, Teaching Assistant for Senior-level Computer Architecture Course EE360N, *Spring 2001, 2003*

- Instructor rating: 4.8/5.0 (S01, 30/50), 4.7/5.0 (S03, 59/70)

University of Texas at Austin, ECE Department, Teaching Assistant for Freshman-level Intro. to Computing Course EE306, *Fall 2000*

- Instructor rating based on student evaluations: 4.7/5.0 (35/55 students responded)

Supervised Students and Trainees

(Please visit <http://www.ece.cmu.edu/~safari/people.html> for a partial list of supervised students).

PhD Alumni

1. Lavanya Subramanian, “Providing High and Controllable Performance in Multicore Systems Through Shared Resource Management,” CMU ECE, defended April 2015. First job: Intel Labs, Santa Clara, CA, USA.
2. Yoongu Kim, “Architectural Techniques to Enhance DRAM Scaling,” CMU ECE, defended May 2015. First job: at Software Robotics, Boston, MA, USA.
3. Donghyuk Lee, “Reducing DRAM Latency at Low Cost by Exploiting Heterogeneity,” CMU ECE, defended December 2015. First job: Nvidia Research, Austin, TX, USA
4. Vivek Seshadri, “Simple DRAM and Virtual Memory Abstractions to Enable Highly Efficient Memory Subsystems,” CMU CS, defended March 2016. First job: Microsoft Research, Bangalore, India.
5. Gennady Pekhimenko, “Practical Data Compression for Modern Memory Hierarchies,” CMU CS, defended July 2016. First job: Assistant Professor at the University of Toronto CS Department, Canada.
6. Rachata Ausavarungnirun, “Techniques for Shared Resource Management in Systems with Throughput Processors,” CMU ECE, defended May 2017. First job: Postdoctoral Researcher at Carnegie Mellon University.
7. Kevin Chang, “Understanding and Improving the Latency of DRAM-Based Memory Systems,” CMU ECE, defended May 2017. First job: Postdoctoral Researcher at Carnegie Mellon University.
8. Yixin Luo, “Architectural Techniques for Improving NAND Flash Memory Reliability,” CMU CS, defended February 2018. First job: Postdoctoral Researcher at Carnegie Mellon University.

Post-doctoral Researcher Alumni

1. Kevin Chang, CMU ECE, May 2017 to January 2018. First job: Research Scientist at Facebook.
2. Saugata Ghose, CMU ECE, September 2014 to September 2016. First job: Special Faculty Member at Carnegie Mellon University.
3. Samira Khan, CMU ECE, September 2012 to September 2015. First job: Assistant Professor at the University of Virginia.

Master’s Alumni

1. Huimin Yang, CMU ECE, December 2010.
2. Chris Craik, “Investigating the Viability of Bufferless NoCs in Modern Chip Multi-Processor Systems,” CMU ECE, Master’s report, August 2011. First job: Google, Mountain View, CA, USA.
3. Jamie Liu, “RAIDR: Retention-Aware Intelligent DRAM Refresh,” CMU ECE, December 2012. First job: Google, Mountain View, CA, USA.
4. Ben Jaiyen, CMU ECE, December 2012. First job: Google, Mountain View, CA, USA.
5. HanBin Yoon, “Techniques for Data Mapping and Buffering to Exploit Asymmetry in Multi-Level Cell (Phase Change) Memory,” CMU ECE, Master’s report, May 2013. First job: Google, Mountain View, CA, USA.
6. Chris Fallin, “The Heterogeneous Block Architecture,” CMU ECE, August 2013. First job: Intel, Hillsboro, OR, USA.
7. Ameya Ambardekar, CMU ECE, August 2013. First job: Intel, Hillsboro, OR, USA.
8. Tyler Huberty, CMU ECE, December 2013. First job: Apple, Cupertino, CA, USA.
9. Abhijith Kashyap, CMU ECE, May 2016. First job: Nvidia, Santa Clara, CA, USA.

Bachelor’s Alumni

1. Xiangyao Yu, CMU ECE (Intern), August 2011. First job: PhD student at MIT.
2. Rachael Harding, CMU ECE, August 2011. First job: PhD student at MIT.

- Alexander Breuss, "Sound-based Localization," ETH Zürich D-INFK, Bachelor's Thesis, August 2017. First job: Master's student at ETH Zürich D-INFK.

Current PhD Students

- Mohammed Alser, Bilkent CS, since 2016 (co-advised)
- Amirali Boroumand, CMU ECE, since 2014
- Lukas Breitwieser, ETH Zürich D-INFK, since February 2017 (hosted at CERN)
- Hasan Hassan, ETH Zürich D-INFK, since December 2016
- Kevin Hsieh, CMU ECE, since 2014 (co-advised)
- Jeremie Kim, CMU ECE, since 2015
- Yang Li, CMU ECE, since 2013
- Justin Meza, CMU ECE, since 2010 (on leave)
- Minesh Patel, CMU ECE, since 2015
- Ivan Puddu, ETH Zürich D-INFK, since February 2017 (co-advised)
- Damla Senol, CMU ECE, since 2015 (co-advised)
- Nandita Vijaykumar, CMU ECE, since 2013 (co-advised)
- A. Giray Yaglikci, ETH Zürich D-INFK, since February 2018

Current Post-Doctoral Researchers

- Arash Tavakkol, ETH Zürich D-INFK, since August 2016.
- Yaohua Wang, ETH Zürich D-INFK, since November 2016.
- Lois Orosa, ETH Zürich D-INFK, since January 2017.
- Rachata Ausavarungnirun, CMU ECE, since May 2017.
- Juan Gomez Luna, ETH Zürich D-INFK, since August 2017.
- Yixin Luo, CMU CS/ECE, since March 2018.

Current Interns

- Mohammad Sadr, ETH Zürich D-INFK, PhD student at Sharif University, since May 2017.
- Geraldo Francisco Oliveira de Junior, ETH Zürich D-INFK, since November 2017.
- Mert Atamaner, ETH Zürich D-INFK, since November 2017.
- Can Firtina, ETH Zürich D-INFK, since February 2018.

Current Master's Research Students

- Nika Mansouri Ghiasi, ETH Zürich D-ITET, since September 2016.
- Batuhan Ceylan, ETH Zürich D-ITET, since February 2016.

Current Undergraduate Research Students

- Tobias Büchli, ETH Zürich D-INFK, since April 2018.

Past Undergraduate Research Students at CMU (incomplete)

Hitesh Arora (IIT Guwahati), Sean Buckley, Rui Cai, Albert Cho, Jared Choi, Vidushi Dadu, Parag Dixit, John Emmons (Drake), Arnab Ghosh (IIT Kanpur), John Greth, Matthias Grundmann (Karlsruhe Inst of Tech), Raghav Gupta, Arjun Hans, Yuchen Hao (Tsinghua), Rachael Harding (grad. Dec 2010), Tyler Huberty, Jeremie Kim, Ji Hye Lee, Evan Lee, Tianshi Li (Peking), Jason Lin, Connor Maines, Tom Mullins, Sunny Nahar, Naman Nandan, Gregory Nazario, Karthic Palaniappan, Harsha Rastogi (BITS Pilani), Eddie Sears, Kevin Woo, Weikun Yang (Peking), Samihan Yedkar, Xiangyao Yu (Tsinghua), Richard Zhu

Past Master's Research Students at CMU (incomplete)

Ameya Ambardekar, Abhishek Bhowmick (IIT Bombay), Chris Craik (grad. Aug 2011), Chris Fallin (Aug 2013), Mughda Satish Dharkar, Hasan Hassan (TOBB ETU), Ben Jaiyen (Dec 2012), Abhijith Kashyap (May 2016), Jamie Liu (Dec 2012), Ashish Shrestha, Huimin Yang, HanBin Yoon

Past High School Research Students at CMU (incomplete)

Aaron Cutright (Pittsburgh Allderdice High School), Steven Barash (Pittsburgh Allderdice High School), Nolan Dickey (Pittsburgh Allderdice High School)

Past Visiting Researchers at CMU (incomplete)

Jongmoo Choi (Dankook Univ.), Hiroyuki Usui (Toshiba), Rui Wang (Beihang).

Past Interns (incomplete)

Rishi Agarwal (IIT-Kanpur), Kypros Constantinides (Michigan), Alejandro Cornejo (MIT), Reetuparna Das (Penn State), Eiman Ebrahimi (UT-Austin), Mohammad Fattah (Turku), Boris Grot (UT-Austin), Engin Ipek (Cornell), Jamie Jablin (Brown), Jose Joao (UT-Austin), Yanjing Li (Stanford), Youyou Lu (Tsinghua), Gregory Nazario (CMU), Arthur Perais (INRIA), Xuehai Qian (Illinois), Jinglei Ren (Tsinghua), Hui Wang (Beihang), Ying Wang (ICT-CAS), Gulay Yalcin (BSC), Xiangyao Yu (Tsinghua), Jishen Xhao (Penn State)

Ph.D. Thesis Committee Member (incomplete)

Manuel le Gallo (ETH), Akkarit Sangpetch (CMU), Asit Mishra (Penn State), Chen-Ling Chou (CMU), Reetuparna Das (Penn State), Eiman Ebrahimi (UT-Austin), Michael Ferdman (CMU), Boris Grot (UT-Austin), Engin Ipek (Cornell), Lei Jin (Pitt), Jose Joao (UT-Austin), Adwait Jog (Penn State), Onur Kayiran (Penn State), Karthik Lakshmanan (CMU), Chang Joo Lee (UT-Austin), Michel Papamichael (CMU), Tim Rogers (UBC), Olatunji Ruwase (CMU), Marek Telgarsky (CMU), Yu Cai (CMU), Hyoseung Kim (CMU), Jishen Zhao (Penn State)

Professional Service

Technical Journal Editorships

Associate Editor, ACM Transactions on Architecture and Code Optimization (TACO), Feb. 2010 - Present.

Associate Editor, ACM Transactions on Storage (TOS), Sep. 2016 - Present.

Editorial Board Member, IEEE Micro, March 2016 - Present.

Editorial Board Member, International Journal of Supercomputing Frontiers and Innovations. Nov. 2014 - Present.

Editorial Board Member, International Journal of Big Data Intelligence (IJBDI), Sep. 2014 - Present.

Co-Guest Editor, IEEE Micro special issue (May/June 2008) on "The Interaction of Computer Architecture and Operating Systems in the Many-core Era," May/June 2008.

Co-Guest Editor, IEEE Micro special issue (Jan/Feb 2011) on "Micro's Top Picks from Computer Architecture Conferences," Jan/Feb 2011.

Technical Conference and Award Committee Chairmanships

MICRO Steering Committee Co-Chair, October 2017-Present.

MICRO Test of Time Award Committee Chair, 2015-Present.

IEEE TCCA Young Computer Architect Award Committee Chair, 2017-Present.

ICS Program Co-Chair, ACM International Conference on Supercomputing, 2016.

NVMSA Program Co-Chair, IEEE International Conference on Non-Volatile Memory Systems and Applications, 2016.

HiPEAC Program Chair, 10th International Conference on High-Performance and Embedded Architectures and Compilers, 2015.

VLSI-SoC Track Co-Chair, 21st IFIP/IEEE International Conference on Very Large Scale Integration, 2013.

MICRO Program Chair, 45th ACM/IEEE International Symposium on Microarchitecture, 2012.

MSPC Program Chair, ACM SIGPLAN Workshop on Memory System Performance and Correctness, 2012.

IEEE Micro Top Picks Selection Committee Co-Chair, IEEE Micro Special Issue on Top Picks from Computer Architecture Conferences, 2011.

IISWC Program Co-Chair, 4th IEEE International Symposium on Workload Characterization (IISWC), 2008.

Technical Conference and Award Committee Memberships

MICRO Steering Committee Member, 2013-Present.

MICRO Test of Time Award Committee Member, 2014-Present.

ICS Steering Committee Member, 2016-Present.

IEEE TCCA Young Computer Architect Award Committee Member, 2016-Present.

IEEE Micro Top Picks Selection Committee Member, IEEE Micro Special Issue on Top Picks from Computer Architecture Conferences, 2018, 2017, 2015, 2013, 2012, 2011, 2010, 2009.

MICRO Program Committee Member, International Symposium on Microarchitecture, 2017, 2015, 2014, 2012, 2008, 2007 (Did not accept invitation in 2013, 2011, 2010 due to paper submission or acceptance limits; 2018 due to time conflicts).

MICRO External Review Committee Member, International Symposium on Microarchitecture, 2016.

ISCA Program Committee Member, International Symposium on Computer Architecture, 2017, 2010, 2008 (Did not accept invitation in 2016, 2015 due to paper submission or acceptance limits).

ISCA External Review Committee Member, International Symposium on Computer Architecture, 2016, 2015.

FAST Program Committee Member, USENIX Conference on File and Storage Technologies, 2019.

PLDI Program Committee Member, ACM SIGPLAN conference on Programming Language Design and Implementation, 2017.

ASPLOS Program Committee Member, Intl. Conf. on Architectural Support for Programming Lang. and Operating Systems, 2016, 2012, 2011, 2009.

ASPLOS External Review Committee Member, Intl. Conf. on Architectural Support for Programming Lang. and Operating Systems, 2018, 2014.

HPCA Program Committee Member, International Symposium on High Performance Computer Architecture, 2018, 2017, 2015, 2012, 2010, 2009.

HPCA External Review Committee Member, International Symposium on High Performance Computer Architecture, 2016.

DSN Program Committee Member, IEEE/IFIP International Conference on Dependable Systems and Networks, 2016.
DAC Program Committee Member, Design Automation Conference, 2016, 2015, 2014.
DATE Program Committee Member, Design Automation Test in Europe Conference, 2016.
PACT Program Committee Member, International Conference on Parallel Architectures and Compilation Techniques, 2013.
PACT External Review Committee Member, International Conference on Parallel Architectures and Compilation Techniques, 2018, 2016, 2014.
ICS Program Committee Member, ACM International Conference on Supercomputing, 2018, 2017, 2016.
HotStorage Program Committee Member, USENIX Workshop on Hot Topics in Storage and File Systems, 2017.
ISMM Program Committee Member, International Symposium on Memory Management, 2013.
ISMM External Review Committee Member, International Symposium on Memory Management, 2016, 2015.
ICAC Program Committee Member, International Conference on Autonomic Computing, 2013.
ICCD Program Committee Member, IEEE International Conference on Computer Design, 2014, 2013, 2012, 2011, 2010, 2009.
IISWC Program Committee Member, IEEE International Symposium on Workload Characterization, 2014, 2011, 2009, 2008.
IISWC Steering Committee Member, IEEE International Symposium on Workload Characterization, 2008-Present.
MSPC Program Committee Member, ACM SIGPLAN Workshop on Memory Systems Performance and Correctness, 2012, 2008
IGCC Program Committee Member, International Green Computing Conference, 2013, 2012.
ASAP Program Committee Member, International Conference on Application-specific Systems, Architectures and Processors, 2013.
SBAC-PAD Program Committee Member, International Symposium on Computer Architecture and High Performance Computing, 2016, 2013.
NVMSA Program Committee Member, IEEE International Conference on Non-Volatile Memory Systems and Applications, 2017, 2016.
MEMSYS Program Committee Member, International Symposium on Memory Systems, 2017.
HPC-Asia Program Committee Member, International Conference on High Performance Computing in Asia-Pacific Region, 2018.

Technical Workshop and Tutorial Committee Memberships and Chairmanships

Workshop on Speculative Side-Channel Analysis Program Committee Member, 2018. (held with ECOOP)
The Memory Forum Co-organizer, 2014. (held with ISCA)
Future Memory Systems Workshop Organizer (held at Carnegie Mellon University), May 2014.
First ASPLOS Doctoral Workshop Program Committee Co-Chair, 2012.
FCRP Memory Cross-Cut Workshop Co-organizer, 2011.
WIOSCA Co-organizer, Workshop on the Interaction between Operating Systems and Computer Architecture, 2007-2010. (Held with ISCA).
JILP Memory Scheduling Championship Program Committee Member, 2012. (held with ISCA)
ALCHEMY Program Committee Member, Architecture, Languages, Compilation and Hardware support for Emerging ManYcore systems, 2014.
WEED Program Committee Member, Workshop on Energy-Efficient Design, 2013, 2010. (held with ISCA)
MoBS Program Committee Member, Workshop on Modeling, Benchmarking, and Simulation, 2010. (held with ISCA)
PESPMA Program Committee Member, Workshop on Parallel Execution of Sequential Programs on Multi-core Architectures, 2010, 2009. (held with ISCA)
CATARS Program Committee Member, Workshop on Compiler/Architectural Techniques for Application Reliability/Security, 2009, 2008. (held with DSN)
WDA Program Committee Member, Workshop on Dependable Architectures, 2008, 2006. (held with MICRO)
RAAW Program Committee Member, Reconfigurable and Adaptive Architecture Workshop, 2007. (held with MICRO)

Technical Reviewer for Journals (incomplete)

ACM Transactions on Architecture and Code Optimization, ACM Transactions on Design Automation of Electronic Systems, ACM Transactions on Embedded Computing Systems, Communications of the ACM, IEEE Computer, IEEE Micro, IEEE Transactions on Computers, IEEE Transactions on Computer Aided Design, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Reliability, IEEE Transactions on Nanotechnology, IEEE Transactions on VLSI Systems, IEEE Transactions on Circuits and Systems, IEEE Computer Architecture Letters, HiPEAC Journal, Journal of Parallel and Distributed Computing, ACM/IEEE Transactions on Networking, Nature Communications, Neural Networks.

Technical Reviewer for Conferences (incomplete)

ISCA, MICRO, HPCA, ASPLOS, Usenix Security Symposium, HotOS, DAC, DATE, DSN, ICS, PACT, RTSS, IPDPS, ISPASS, HiPEAC, Euro-Par, SBAC-PAD, WASP, SBCCI, IGCC, MSPC, IISWC.

Expert Reviewer for Research Proposals (incomplete)

Swiss National Science Foundation Reviewer, 2017, 2016
European Research Council Reviewer, 2017, 2016
NSF Review Panel Participant, 2007, 2008, 2009, 2010, 2011, 2012, 2013, 2014, 2015.
DOE Review Panel Participant, 2010.
Reviewer for Research Proposals in Australia, Belgium, European Union, Canada, Norway, Switzerland, 2013, 2014, 2015, 2016, 2017.
Microsoft Research Breakthrough Research Grant Proposal Reviewer, 2007.
Microsoft Research New Faculty Fellowship Candidate Reviewer, 2006, 2007.
Microsoft Research Reviewer for Miscellaneous Requests for Proposals, 2007, 2008.
Microsoft Research PhD Fellowship Candidate Reviewer, 2006, 2007, 2008.

Departmental Committee Memberships

ETH D-INFK: Faculty Hiring Committee (2018-Present)

ETH D-INFK: Doctoral Studies Committee (2016-Present), Doctoral Admissions Committee (2016-2017)

CMU ECE: Faculty Search Committee (2011-2016), Admissions Committee (2011-2014), Seminar Committee (2010-2011)

CMU CSD: ACM Doctoral Dissertation Award Committee (Fall 2012), Faculty Hiring Committee (2010)

Other Committee Memberships

Local Arrangements Chair, ASPLOS 2010.

Publicity Chair, HiPEAC 2017, HiPEAC 2016, ASPLOS 2015, HPCA 2014, HPCA 2013, HPCA 2010, IISWC 2009, ASPLOS 2008.

Publications Chair, ICS 2007.

Open Source Educational and Research Artifacts

Lecture Videos and Course Materials

Please visit <https://people.inf.ethz.ch/omutlu/teaching.html> for online course materials.

Please visit <https://www.youtube.com/user/cm18447> for online computer architecture lecture videos at Carnegie Mellon University.

Please visit <https://www.youtube.com/channel/UCIwQ8uOeRFgOEvBLYc3kc3g> for online lecture videos for my courses at ETH Zürich.

Source Code Releases

Please visit <http://www.ece.cmu.edu/~safari/tools.html> and <https://github.com/CMU-SAFARI>.

Professional Memberships

Fellow of ACM, 2017, for contributions to computer architecture research, especially in memory systems

Senior Member of IEEE

Member of ACM SIGARCH, ACM SIGMICRO, IEEE Computer Society

Member of Eta Kappa Nu Electrical Engineering Honor Society and Tau Beta Pi Engineering Honor Society, since 1998