Onur Mutlu

Professor of Computer Science
Department of Information Technology and Electrical Engineering, Gloriastrasse 35, ETH Zürich, 8092 Zürich, Switzerland
https://people.inf.ethz.ch/omutlu/
omutlu@gmail.com
https://www.youtube.com/OnurMutluLectures
Research Group Website: https://safari.ethz.ch/

Overview of Research, Teaching and Professional Interests

- My main research is in computer architecture, computing systems, hardware security, memory & storage systems, and bioinformatics.
- My work spans and stretches the boundaries between hardware and software (including applications, system software, compilers, architecture, microarchitecture, and logic circuits). 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-efficient and fundamentally-secure 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. Some of my present and past research topics include the following:
- Memory and storage systems (DRAM, flash, NVM, emerging). Memory controllers. Scalable, QoS-aware, latency-tolerant systems.
- New computing and memory paradigms, including processing in memory, biologically-inspired computation.
- Robust computing systems (including reliability, security, safety, availability). Fault-tolerant and bug-tolerant architectures.
- Multi/many-core systems. Hardware/software interaction. OS/PL/hardware interaction. Processor design. Microarchitecture.
- Architectures for machine learning and bioinformatics. System and hardware design for ML and genome analysis.
- Interconnects. Communication architectures, both small scale and large scale. Efficient router design.
- Hardware accelerators. Graphics processing units. Heterogeneous architectures. Software and hardware design issues.
- New hardware/software interfaces. Energy, programmability, performance, portability, virtualization support.
- Workload and system characterization and analysis. Profiling, simulation, benchmarking tools. Open source software and hardware.

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, Professor of Computer Science, Full Professor, ***September 2015 - Present***
Departments of Information Technology and Electrical Engineering (2020-Present) and Computer Science (2015-2020)
Stanford University, Dept. of Electrical Engineering, Visiting Professor, ***September 2023 - September 2024***
Carnegie Mellon University, Dept. of Electrical and Computer Engineering, Adjunct Professor, ***June 2016 - September 2022***
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 - September 2022***
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 Received for Published Scholarly Work

- IEEE/IFIP International Conference on Dependable Systems and Networks (DSN) Test of Time Award, 2025
  For the DSN 2015 paper "AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems"

- Distinguished artifact award at HPCA 2025 (IEEE High-Performance Computer Architecture Symposium)
  Understanding RowHammer Under Reduced Refresh Latency: Experimental Analysis of Real DRAM Chips and Implications on Future Solutions
- Best paper award at ISCA 2024 (ACM/IEEE International Symposium on Computer Architecture)
  Constable: Improving Performance and Power Efficiency by Safely Eliminating Load Instruction Execution
- IFIP Working Group 10.4 Jean-Claude Laprie Award in Dependable Computing, 2024
  For the ISCA 2014 paper "Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors"
  Citation: "The first paper to introduce, demonstrate, analyze, and propose solutions to DRAM row disturbance errors, later known as Rowhammer. The solution proposed in the paper, PARA (Probabilistic Adjacent Row Activation), was incorporated by Intel and other CPU vendors, and is still one of main ways to fix Rowhammer. It took a while for the problem to be recognized by DRAM vendors, but eventually it led to significant efforts and investment into fixing it. On the research side, the paper had a strong impact by creating new research directions into DRAM security and Rowhammer attacks and defenses, collecting an impressive number of citations over the last 10 years (1300+ citations). The paper also had an important impact on the dependability community at large by motivating others into looking at the reliability of general-purpose hardware and discovering other vulnerabilities, like Spectre and Meltdown."
- One paper (of 12 total) selected for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2023
  RowPress: Amplifying Read Disturbance in Modern DRAM Chips
- Distinguished artifact award at MICRO 2023 (ACM/IEEE International Symposium on Microarchitecture)
  Victima: Drastically Increasing Address Translation Reach by Leveraging Underutilized Cache Resources
- Five papers selected to the 50th Anniversary of ISCA (ACM/IEEE International Symposium on Computer Architecture), 2023
  Self Optimizing Memory Controllers: A Reinforcement Learning Approach (ISCA 2008)
  RAIDR: Retention-Aware Intelligent DRAM Refresh (ISCA 2012)
  An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms (ISCA 2013)
  Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors (ISCA 2014)
  A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing (ISCA 2015)
- Distinguished artifact award at ISCA 2023 (ACM/IEEE International Symposium on Computer Architecture)
  RowPress: Amplifying Read Disturbance in Modern DRAM Chips
- Best artifact award at PPoPP 2023 (ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming)
  High-Performance and Scalable Agent-Based Simulation with BioDynaMo
- Best paper award at HiPEAC 2023 (International Conference on High-Performance Embedded Architectures and Compilers)
  MetaSys: A Practical Open-source Metadata Management System to Implement and Evaluate Cross-layer Optimizations
- Best paper award at MICRO 2022 (ACM/IEEE International Symposium on Microarchitecture)
  Hermes: Accelerating Long-Latency Load Requests via Perceptron-Based Off-Chip Load Prediction
- Best paper award at SC 2022 (ACM International Conference for High Performance Computing, Networking, Storage, and Analysis)
  ProbGraph: High-Performance and High-Accuracy Graph Mining with Probabilistic Set Representations
- Intel Hardware Security Academic Award Finalist, 2022
  For the HPCA 2021 paper "BlockHammer: Preventing RowHammer at Low Cost by Blacklisting Rapidly-Accessed DRAM Rows"
- Persistent Impact Prize (of the Non-Volatile Memory Systems Workshop), 2022
  For the ISCA 2009 paper "Architecting Phase Change Memory as a Scalable DRAM Alternative"
  Citation: "in recognition of its seminal contribution to understanding how non-volatile memories might replace conventional DRAM. It also played a critical role in introducing Phase Change Memory (PCM) technology to the architecture community, spawning a wide range of follow-on work."
- IEEE High-Performance Computer Architecture Symposium (HPCA) Test of Time Award, 2021
  For the HPCA 2003 paper "Runahead Execution: An Alternative to Very Large Instruction Windows for Out-of-order Processors"
  Citation: "Runahead Execution is a pioneering paper that opened up new avenues in dynamic prefetching. The basic idea of run ahead execution effectively increases the instruction window very significantly, without having to increase physical resource size (e.g. the issue queue). This seminal paper spawned off a new area of ILP-enhancing microarchitecture research. This work has had strong industry impact as evidenced by IBM's POWER6 - Load Lookahead, NVIDIA Denver, and Sun ROCK's hardware scouting."
- Pwnie Award 2020 for Most Innovative Research, 2020.
  TRRespass: Exploiting the Many Sides of Target Row Refresh
- Best Paper Award at MICRO 2020 (ACM/IEEE International Symposium on Microarchitecture)
  Bit-Exact ECC Recovery (BEER): Determining DRAM On-Die ECC Functions by Exploiting DRAM Data Retention Characteristics
- Best Paper Award at the IEEE Security and Privacy Conference 2020
  TRRespass: Exploiting the Many Sides of Target Row Refresh
- Best Paper Award at DSN 2019 (IEEE/IFIP International Conference on Dependable Systems and Networks)
  Understanding and Modeling On-Die Error Correction in Modern DRAM: An Experimental Study Using Real Devices
- One paper (of 7 total) selected as a "Top Pick in Hardware and Embedded Security" (as one of 7 published in 2012-2017 in the area), 2019
  Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors (ISCA 2014)
- 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," 2020
TRRespass: Exploiting the Many Sides of Target Row Refresh
- One paper selected as Honorable Mention for IEEE Micro's "Top Picks from Computer Architecture Conferences," 2019
D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput
- 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 Communication of the ACM (CACM) "Research Highlights," 2009
Architecting Phase Change Memory as a Scalable DRAM Alternative
- Best Paper Session at ISPASS 2023: *Evaluating Machine Learning Workloads on Memory-Centric Computing Systems*
- 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 ISPASS 2023, FPL 2020, MICRO 2016, NOCS 2015, NOCS 2012, HPCA 2015, HPCA 2014, HPCA 2010, HPCA 2009, HPCA 2007, MICRO 2006, and MICRO 2005 conferences
- PhD Dissertation nominated by UT-Austin for the ACM Doctoral Dissertation Award, 2006
Dissertation Title: Efficient Runahead Execution Processors

Major Honors and Awards

- IEEE Computer Society Harry H. Goode Memorial Award, "for seminal contributions to computer architecture research and practice, especially in memory systems," 2025
- Huawei OlympusMons Award, "for outstanding achievements in data-centric and data-driven storage system design for high performance, efficiency and reliability," 2023
- Google Open Source Peer Bonus Award, 2023 *For contributions to the educational YouTube channel Onur Mutlu Lectures*
- Google Security and Privacy Research Award, 2022
- Intel Outstanding Researcher Award, 2022 *Citation: "Professor Onur Mutlu's emerging-workload-driven research has led to open-source benchmark suites that help hardware-software co-design and exploration for near-memory architectures. His team's research has also led to innovative near-memory architectures for deep learning and genomics applications."*
- IEEE Computer Society Edward J. McCluskey Technical Achievement Award, "for innovative and impactful contributions to computer memory systems," 2020
- ACM SIGARCH Maurice Wilkes Award, "for innovative contributions in efficient and secure DRAM systems," 2019
- Facebook AI System Hardware/Software Co-Design Research Award, 2019
- IEEE Fellow, "for contributions to computer architecture research and practice," 2018
- Elected Member of the Academy of Europe (Academia Europaea), "for outstanding achievements as a researcher," 2018
- ACM Fellow, "for contributions to computer architecture research, especially in memory systems," 2017

- Google Faculty Research Award, 2019, 2016, 2015, 2014
- 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, 2013, 2012
- 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, 2013, 2012 (with multiple CMU Faculty members)
- More than 100 keynote, plenary and special invited talks at conferences and workshops, including recently at Qualcomm Product Security Summit 2024, Google Zurich Hardware Security Summit 2024, ICECS 2024, Hardwear.io Memory Security Track 2024
- Distinguished Lecture at University of California at Irvine CS Department, Irvine, CA, 7 March 2025.
- Distinguished Lecture at University of Toronto CS Department, Toronto, ON, Canada, 7 May 2024.
- Distinguished Lecture at Arizona State University ACME Center, Tempe, AZ, USA, 28 March 2024.
- Distinguished Lecture at Duke University Athena AI Institute, Durham, NC, USA, 15 February 2024.
- Distinguished Lecture at Simon Fraser University, Burnaby, BC, Canada, 22 September 2022.
- Distinguished Lecture at Ontario Tech Engineering Research Distinguished Speaker Series, Virtual, 19 October 2021.
- Distinguished Seminar at Koc University College of Engineering, Virtual, 9 April 2021.
- Distinguished Lecture at HKUST Engineering and HKSTP Distinguished Speaker Series, Virtual, 7 October 2020.
- Distinguished Lecture at George Washington University, Washington, DC, USA, 15 February 2019.
- ECE Endowed Lectureship Series at the University of Texas at Austin, Austin, TX, USA 29 January-2 February 2019.
- Distinguished Lecture at INESC-ID Distinguished Lecture Series, University of Lisbon, Portugal, 4 December 2017.
- Distinguished Lecture at Triangle Computer Science Distinguished Lecture Series, Raleigh, NC, 11 April 2016.
- Distinguished Lecture at Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, 28 June 2012.
- Selected to the ISCA Hall of Fame, 2009
- Selected to the MICRO Hall of Fame, 2009
- Selected to the HPCA Hall of Fame, 2010
- Selected to the ASPLOS Hall of Fame, 2018
- NSF CAREER Award, 2010 (*QoS-Aware, High-Performance, and Scalable Many-Core Memory Systems*)
- Microsoft Gold Star Award, 2008
- University Co-op/George H. Mitchell Award for Excellence in Graduate Research (Awarded to 6 out of 271 nominees at UT-Austin), 2005

Major Student Research Honors and Awards

- William C. Carter PhD Dissertation Award in Dependability (for PhD Graduate A. Giray Yaglikci), 2025
  Dissertation: "Enabling Efficient and Scalable DRAM Read Disturbance Mitigation via New Experimental Insights into Modern DRAM Chips"
- IEEE VLSI Test Symposium (VTS) TTTC's E. J. McCluskey Best Doctoral Thesis Award Finalist (for PhD Graduate A. Giray Yaglikci), 2025 *Dissertation: "Enabling Efficient and Scalable DRAM Read Disturbance Mitigation via New Experimental Insights into Modern DRAM Chips"*
- IEEE International Symposium on Hardware Oriented Security and Trust (HOST) PhD Competition Finalist (for PhD Graduate A. Giray Yaglikci), 2024 *Dissertation: "Enabling Efficient and Scalable DRAM Read Disturbance Mitigation via New Experimental Insights into Modern DRAM Chips"*
- ACM PACT (Parallel Architectures and Compilation Techniques) Student Research Competition Winner (for PhD Graduate A. Giray Yaglikci), 2023 *Research: "Understanding and Leveraging the Spatial Variation in Read Disturbance Vulnerability of Real DRAM Chips"*
- European Design and Automation Association (EDAA) Outstanding Dissertation Award (for PhD Graduate Hasan Hassan), 2023
  Dissertation: Improving DRAM Performance, Reliability, and Security by Rigorously Understanding Intrinsic DRAM Operation
- National Technical University of Athens (NTUA) Best Doctoral Thesis Award (for PhD Graduate Christina Giannoula), 2023.
  Dissertation: "Accelerating Irregular Applications via Efficient Synchronization and Data Access Techniques"
- William C. Carter PhD Dissertation Award in Dependability (for PhD Graduate Minesh Patel), 2022
  Dissertation: Enabling Effective Error Mitigation in Memory Chips That Use On-Die Error-Correcting Codes
- ETH Doctoral Medal for Outstanding PhD Thesis (for PhD Graduate Minesh Patel; awarded to about 8% of all ETH PhD theses), 2022
  Dissertation: Enabling Effective Error Mitigation in Memory Chips That Use On-Die Error-Correcting Codes
- European Design and Automation Association (EDAA) Outstanding Dissertation Award (for PhD Graduate Jeremie Kim), 2021
  Dissertation: Improving DRAM Performance, Security, and Reliability by Understanding and Exploiting DRAM Timing Parameter Margins
- Semiconductor Research Corporation (SRC) Best Student Presentation Award at TECHCON 2019 (for PhD Student Damla Senol Cali), 2019
  Paper: GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis
- IEEE Turkey Ph.D. Dissertation Award (for Ph.D. Student Mohammed Alser), 2018
  Dissertation: Accelerating the Understanding of Life's Code Through Better Algorithms and Hardware Design"
- Semiconductor Research Corporation Best Student Presentation Award at TECHCON 2018 (for PhD Student Amirali Boroumand), 2018
  Paper: Google Workloads for Consumer Devices: Mitigating Data Movement Bottlenecks
- Best Poster Award at RECOMB-Seq 2018 (for work led and presented by PhD student Damla Senol Cali), 2018

Poster: Accelerating Approximate Pattern Matching with Processing-In-Memory and Single-Instruction Multiple-Data Programming

- Semiconductor Research Corporation (SRC) Best in Session Award at SRC TECHCON 2016 (for PhD Student Kevin Chang), 2016
  Paper: Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization
- Best (student) presentation award at HiPEAC 2015 (for PhD Student Justin Meza), 2015
  Paper: Efficient Data Mapping and Buffering Techniques for Multi-Level Cell Phase-Change Memories
- ACM ASPLOS (Architectural Support for Programming Languages and Operating Systems) Student Research Competition Winner (for PhD Graduate Gennady Pekhimenko), 2015
  Research: "Energy-Efficient Data Compression for GPU Memory Systems"

Publications

Please visit `https://people.inf.ethz.ch/omutlu/projects.htm` for electronic copies, slides, and source code.

DBLP record: `https://dblp.org/pid/m/OnurMutlu.html`

Google Scholar record: `https://scholar.google.com/citations?hl=en&user=7XyGUGkAAAAJ&view_op=list_works&sortby=pubdate`

Book Chapters

1. Onur Mutlu, Saugata Ghose, Juan Gomez-Luna, and Rachata Ausavarungnirun, "A Modern Primer on Processing in Memory," *Emerging Computing: From Devices to Systems - Looking Beyond Moore and Von Neumann*, Springer, 2022.
2. Rachata Ausavarungnirun and Onur Mutlu, "The Design of an Energy-Efficient Deflection-Based On-Chip Network," *Advances in Computers*, Elsevier, 2022.
3. Vivek Seshadri and Onur Mutlu, "Bulk Bitwise Execution Model in Memory: Mechanisms, Implementation, and Evaluation," *Advances in Computers*, Elsevier, 2022.
4. Bryan Donyanavard, Amir M. Rahmani, Axel Jantsch, Onur Mutlu, and Nikil Dutt, "Intelligent Management of Mobile Systems Through Computational Self-Awareness," Handbook of Research on Methodologies and Applications of Supercomputing, IGI Global, 2021.
5. 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, 2019.
6. 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, 2018.
7. 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, 2018.
8. Vivek Seshadri, Onur Mutlu, "Simple Operations in Memory to Reduce Data Movement," *Advances in Computers*, Elsevier, 2017.
9. 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.
10. Onur Mutlu, "Main Memory Scaling: Challenges and Solution Directions," in *More Than Moore Technologies for Next Generation Computer Design*, *Springer*, January 2015.
11. Yoongu Kim, Onur Mutlu, "Memory Systems," in *Computing Handbook: Computer Science and Software Engineering*, *CRC Press*, April 2014.
12. 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.
13. 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

14. Kangqi Chen, Rakesh Nadig, Andreas Kosmas Kakolyris, Manos Frouzakis, Nika Mansouri Ghiasi, Yu Liang, Haiyu Mao, Jisung Park, Mohammad Sadrosadati, and Onur Mutlu, "REIS: A High-Performance and Energy-Efficient Retrieval System with In-Storage Processing," *Proceedings of the 52nd International Symposium on Computer Architecture **(ISCA)***, Tokyo, Japan, June 2025.
15. Ismail Emir Yuksel, Akash Sood, Ataberk Olgun, Oguzhan Canpolat, Haocong Luo, Nisa Bostanci, Mohammad Sadrosadati, A. Giray Yaglikci, and Onur Mutlu, "PuDHammer: Experimental Analysis of Read Disturbance Effects of Processing-using-DRAM in Real DRAM Chips," *Proceedings of the 52nd International Symposium on Computer Architecture **(ISCA)***, Tokyo, Japan, June 2025.
16. Nisa Bostanci, Konstantinos Kanellopoulos, Ataberk Olgun, A. Giray Yaglikci, Ismail Emir Yuksel, Nika Mansouri Ghiasi, Zulal Bingol, Mohammad Sadrosadati, and Onur Mutlu, "Revisiting Main Memory-Based Covert and Side Channel Attacks in the Context of Processing-in-Memory," *Proceedings of the 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks **(DSN)***, Naples, Italy, June 2025. **Officially artifact evaluated as available, reviewed, and reproduced.**
17. Oğuzhan Canpolat, Ataberk Olgun, David Novo, Oguz Ergin, and Onur Mutlu, "EasyDRAM: An FPGA-based Infrastructure for Fast and Accurate End-to-End Evaluation of Emerging DRAM Techniques," *Proceedings of the 55th Annual IEEE/IFIP International Conference on Dependable Systems and Networks **(DSN)***, Naples, Italy, June 2025.

18. Geraldo Francisco, Mayank Kabra, Yuxin Guo, Kangqi Chen, A. Giray Yaglikci, Melina Soysal, Mohammad Sadrosadati, Joaquin Olivares, Saugata Ghose, Juan Gomez-Luna, and Onur Mutlu, "Proteus: Achieving High-Performance Processing-Using-DRAM via Dynamic Precision Bit-Serial Arithmetic," *Proceedings of the 37th ACM International Conference on Supercomputing **(ICS)***, Salt Lake City, UT, USA, June 2025.
19. Melina Soysal, Konstantina Koliogeorgi, Can Firtina, Nika MansouriGhiasi, Rakesh Nadig, Haiyu Mao, Geraldo Francisco, Yu Liang, Klea Zambaku, Mohammad Sadrosadati, and Onur Mutlu, "MARS: Processing-In-Memory Acceleration of Raw Signal Genome Analysis Inside the Storage Subsystem," *Proceedings of the 37th ACM International Conference on Supercomputing **(ICS)***, Salt Lake City, UT, USA, June 2025.
20. Christina Giannoula, Peiming Yang, Ivan Fernandez, Jiacheng Yang, Sankeerth Durvasula, Yu Xin Li, Mohammad Sadrosadati, Juan Gomez Luna, Onur Mutlu, and Gennady Pekhimenko, "PyGim: An Efficient Graph Neural Network Library for Real Processing-In-Memory Architectures," *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems **(SIGMETRICS)***, Stony Brook, NY, USA, June 2025.
21. Onur Mutlu, Ataberk Olgun, and Ismail Emir Yuksel, "Memory-Centric Computing: Solving Computing's Memory Problem," *Proceedings of the 17th International Memory Workshop **(IMW)***, Monterey, CA, May 2025. **Invited Paper and Presentation.**
22. Konstantinos Kanellopoulos, Konstantinos Sgouras, Nisa Bostanci, Andreas Kosmas Kakolyris, Berkin Kerim Konar, Rahul Bera, Rakesh Kumar, Mohammad Sadrosadati, Nandita Vijaykumar, and Onur Mutlu, "Virtuoso: Enabling Fast and Accurate Virtual Memory Research via an Imitation-based Operating System Simulation Methodology," *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Rotterdam, Netherlands, April 2025.
23. Yintao He, Haiyu Mao, Christina Giannoula, Mohammad Sadrosadati, Juan Gomez-Luna, Huawei Li, Xiaowei Li, Ying Wang, and Onur Mutlu, "PAPI: Exploiting Dynamic Parallelism in Large Language Model Decoding with a Processing-In-Memory-Enabled Computing System," *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Rotterdam, Netherlands, April 2025.
24. Yufeng Gu, Alireza Khadem, Sumanth Umesh, Ning Liang, Xavier Servot, Onur Mutlu, Ravi Iyer, and Reetuparna Das, "PIM Is All You Need: A CXL-Enabled GPU-Free System for Large Language Model Inference," *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Rotterdam, Netherlands, April 2025. **Officially artifact evaluated as available, functional, and reproduced.**
25. Mayank Kabra, Rakesh Nadig, Harshita Gupta, Rahul Bera, Manos Frouzakis, Vamanan Arulchelvan, Yu Liang, Haiyu Mao, Mohammad Sadrosadati, and Onur Mutlu, "CIPHERMATCH: Accelerating Homomorphic Encryption based String Matching via Memory-Efficient Data Packing and In-Flash Processing," *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Rotterdam, Netherlands, April 2025.
26. Haocong Luo, Ismail Emir Yuksel, Ataberk Olgun, Abdullah Giray Yaglikci, and Onur Mutlu, "Revisiting DRAM Read Disturbance: Identifying Inconsistencies Between Experimental Characterization with Device-Level Studies," *Proceedings of the 43rd IEEE VLSI Test Symposium **(VTS)***, Tempe, AZ, USA, April 2025.
27. Ataberk Olgun, Nisa Bostanci, Ismail Emir Yuksel, Abdullah Giray Yaglikci, Geraldo Francisco de Oliveira, Haocong Luo, Oguzhan Canpolat, Minesh Patel, and Onur Mutlu, "Variable Read Disturbance: An Experimental Analysis of Temporal Variation in DRAM Read Disturbance," *Proceedings of the 31st IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Las Vegas, NV, USA, March 2025.
28. Oguzhan Canpolat, Abdullah Giray Yaglikci, Geraldo Francisco de Oliveira, Ataberk Olgun, Nisa Bostanci, Ismail Emir Yuksel, Haocong Luo, Oguz Ergin, and Onur Mutlu, "Chronus: Understanding and Securing the Cutting-Edge Industry Solutions to DRAM Read Disturbance," *Proceedings of the 31st IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Las Vegas, NV, USA, March 2025. **Officially artifact evaluated as available, functional, and reproduced.**
29. Yahya Can Tugrul, Abdullah Giray Yaglikci, Ismail Emir Yuksel, Ataberk Olgun, Oguzhan Canpolat, Nisa Bostanci, Mohammad Sadrosadati, Oguz Ergin, and Onur Mutlu, "Understanding RowHammer Under Reduced Refresh Latency: Experimental Analysis of Real DRAM Chips and Implications on Future Solutions," *Proceedings of the 31st IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Las Vegas, NV, USA, March 2025. **Distinguished Artifact Award. Officially artifact evaluated as available, reusable and reproducible.**
30. Yu Liang, Aofeng Shen, Chun Jason Xue, Riwei Pan, Haiyu Mao, Nika Mansouri Ghiasi, Qingcai Jiang, Rakesh Nadig, Lei Li, Rachata Ausavarungnirun, Mohammad Sadrosadati, and Onur Mutlu, "Ariadne: Hotness-Aware and Size-Adaptive Compressed Swap Scheme for Mobile Devices," *Proceedings of the 31st IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Las Vegas, NV, USA, March 2025.
31. Onur Mutlu, Ataberk Olgun, Geraldo F. Oliveira, and İsmail Emir Yüksel, "Memory-Centric Computing: Recent Advances in Processing-in-DRAM," *Proceedings of the 70th Annual IEEE International Electron Devices Meeting **(IEDM)***, San Francisco, CA, USA, December 2024. **Invited Paper and Presentation.**
32. Hasan Hassan, Ataberk Olgun, A. Giray Yaglikci, Haocong Luo, and Onur Mutlu, "Self-Managing DRAM: A Low-Cost Framework for Enabling Autonomous and Efficient DRAM Maintenance Operations," *Proceedings of the 57th International Symposium on Microarchitecture **(MICRO)***, Austin, TX, USA, November 2024.
33. Oğuzhan Canpolat, A. Giray Yağlıkçı, Ataberk Olgun, İsmail Emir Yüksel, Yahya Can Tuğrul, Konstantinos Kanellopoulos, Oğuz Ergin, and Onur Mutlu, "BreakHammer: Enhancing RowHammer Mitigations by Carefully Throttling Suspect Threads," *Proceedings of the 57th*

*International Symposium on Microarchitecture **(MICRO)***, Austin, TX, USA, November 2024. **Officially artifact evaluated as available, functional, reusable and reproducible.**

34. Shai Bergman, Onur Mutlu, Wu Yong, Keji Huang, and Ji Zhang, "Composable Storage Servers: A Storage Paradigm for Disaggregated Systems," *Proceedings of the 17th IEEE International Conference on Networking, Architecture and Storage **(NAS)***, Guangzhou, China, November 2024.
35. Steve Rhyner, Haocong Luo, Juan Gómez-Luna, Mohammad Sadrosadati, Jiawei Jiang, Ataberk Olgun, Harshita Gupta, Ce Zhang, and Onur Mutlu, "PIM-Opt: Demystifying Distributed Optimization Algorithms on a Real-World Processing-In-Memory System," *Proceedings of the 33rd International Conference on Parallel Architectures and Compilation Techniques **(PACT)***, Long Beach, CA, USA, October 2024.
36. Ataberk Olgun, Yahya Can Tugrul, Nisa Bostanci, Ismail Emir Yuksel, Haocong Luo, Steve Rhyner, Abdullah Giray Yaglikci, Geraldo F. Oliveira, and Onur Mutlu, "ABACuS: All-Bank Activation Counters for Scalable and Low Overhead RowHammer Mitigation," *Proceedings of the 33rd USENIX Security Symposium (**USENIX SECURITY**)*, Philadelphia, PA, USA, August 2024. **Officially artifact evaluated as available, functional, and reproduced.**
37. Rahul Bera, Adithya Ranganathan, Joydeep Rakshit, Sujit Mahto, Anant V. Nori, Jayesh Gaur, Ataberk Olgun, Konstantinos Kanellopoulos, Mohammad Sadrosadati, Sreenivas Subramoney, and Onur Mutlu, "Constable: Improving Performance and Power Efficiency by Safely Eliminating Load Instruction Execution," *Proceedings of the 51st International Symposium on Computer Architecture **(ISCA)***, Buenos Aires, Argentina, July 2024. **Best Paper Award at ISCA 2024.**
38. Nika Mansouri Ghiasi, Mohammad Sadrosadati, Harun Mustafa, Arvid Gollwitzer, Can Firtina, Julien Eudine, Haiyu Mao, Joel Lindegger, Meryem Banu Cavlak, Mohammed Alser, Jisung Park, and Onur Mutlu, "MegIS: High-Performance and Low-Cost Metagenomic Analysis with In-Storage Processing," *Proceedings of the 51st International Symposium on Computer Architecture **(ISCA)***, Buenos Aires, Argentina, July 2024.
39. Julian Pavon, Ivan Vargas Valdivieso, Carlos Rojas, Cesar Hernandez, Mehmet Aslan, Roger Figueras, Yichao Yuan, Joel Lindegger, Mohammed Alser, Francesc Moll, Santiago Marco-Sola, Oguz Ergin, Nishil Talati, Onur Mutlu, Osman Unsal, Mateo Valero, and Adrian Cristal, "QUETZAL: Vector Acceleration Framework For Modern Genome Sequence Analysis," *Proceedings of the 51st International Symposium on Computer Architecture **(ISCA)***, Buenos Aires, Argentina, July 2024.
40. Oğuzhan Canpolat, A. Giray Yağlıkçı, Geraldo F. Oliveira, Ataberk Olgun, Oğuz Ergin, and Onur Mutlu, "Understanding the Security Benefits and Overheads of Emerging Industry Solutions to DRAM Read Disturbance," *4th Workshop on DRAM Security **(DRAMsec)**, held with 51st Annual International Symposium on Computer Architecture (ISCA)*, Buenos Aires, Argentina, July 2024.
41. Ismail Emir Yuksel, Yahya Can Tugrul, F. Nisa Bostanci, Geraldo F. Oliveira, A. Giray Yaglikci, Ataberk Olgun, Melina Soysal, Haocong Luo, Juan Gomez-Luna, Mohammad Sadrosadati, and Onur Mutlu, "Simultaneous Many-Row Activation in Off-the-Shelf DRAM Chips: Experimental Characterization and Analysis," *Proceedings of the 54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks **(DSN)***, Brisbane, Australia, June 2024. **Officially artifact evaluated as both code and dataset available, reviewed and reproducible.**
42. Ataberk Olgun, Majd Osseiran, Abdullah Giray Yaglikci, Yahya Can Tugrul, Haocong Luo, Steve Rhyner, Behzad Salami, Juan Gomez-Luna, and Onur Mutlu, "Read Disturbance in High Bandwidth Memory: A Detailed Experimental Study on HBM2 DRAM Chips," *Proceedings of the 54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks **(DSN)***, Brisbane, Australia, June 2024. **Officially artifact evaluated as both code and dataset available, reviewed and reproducible.**
43. Haocong Luo, Ismail Emir Yüksel, Ataberk Olgun, A. Giray Yağlıkçı, Mohammad Sadrosadati, and Onur Mutlu, "An Experimental Characterization of Combined RowHammer and RowPress Read Disturbance in Modern DRAM Chips," *Proceedings of the 54th Annual IEEE/IFIP International Conference on Dependable Systems and Networks Disrupt Track **(DSN Disrupt)***, Brisbane, Australia, June 2024.
44. Kailash Gogineni, Sai Santosh Dayapule, Juan Gomez-Luna, Karthikeya Gogineni, Peng Wei, Tian Lan, Mohammad Sadrosadati, Onur Mutlu, and Guru Venkataramani, "SwiftRL: Towards Efficient Reinforcement Learning on Real Processing-In-Memory Systems," *Proceedings of the 2024 IEEE International Symposium on Performance Analysis of Systems and Software **(ISPASS)***, Indianapolis, Indiana, May 2024.
45. Sungjun Cho, Beomjun Kim, Hyunuk Cho, Gyeongseob Seo, Onur Mutlu, Myungsuk Kim, and Jisung Park, "AERO: Adaptive Erase Operation for Improving Lifetime and Performance of Modern NAND Flash-Based SSDs" *Proceedings of the 29th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, San Diego, CA, USA, April 2024.
46. Abdullah Giray Yaglikci, Geraldo Francisco de Oliveira, Yahya Can Tugrul, Ismail Yuksel, Ataberk Olgun, Haocong Luo, and Onur Mutlu, "Spatial Variation-Aware Read Disturbance Defenses: Experimental Analysis of Real DRAM Chips and Implications on Future Solutions," *Proceedings of the 30th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Edinburgh, Scotland, UK, March 2024.
47. F. Nisa Bostanci, Ismail Emir Yuksel, Ataberk Olgun, Konstantinos Kanellopoulos, Yahya Can Tugrul, A. Giray Yaglikci, Mohammad Sadrosadati, and Onur Mutlu, "CoMeT: Count-Min-Sketch-based Row Tracking to Mitigate RowHammer at Low Cost," *Proceedings of the 30th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Edinburgh, Scotland, UK, March 2024. **Officially artifact evaluated as available, reviewed and reproducible.**
48. Ismail Emir Yuksel, Yahya Can Tugrul, Ataberk Olgun, F. Nisa Bostanci, A. Giray Yaglikci, Geraldo F. Oliveira, Haocong Luo, Juan Gomez-Luna, Mohammad Sadrosadati, and Onur Mutlu, "Functionally-Complete Boolean Logic in Real DRAM Chips: Experimental Characterization and Analysis," *Proceedings of the 30th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Edinburgh, Scotland, UK, March 2024.

49. Geraldo F. Oliveira, Ataberk Olgun, Abdullah Giray Yaglikci, F. Nisa Bostanci, Juan Gomez-Luna, Saugata Ghose, and Onur Mutlu, "MIMDRAM: An End-to-End Processing-Using-DRAM System for High-Throughput, Energy-Efficient and Programmer-Transparent Multiple-Instruction Multiple-Data Processing," *Proceedings of the 30th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Edinburgh, Scotland, UK, March 2024.
50. Konstantinos Kanellopoulos, Hong Chul Nam, Nisa Bostanci, Rahul Bera, Mohammad Sadrosadati, Rakesh Kumar, Davide-Basilio Bartolini, and Onur Mutlu, "Victima: Drastically Increasing Address Translation Reach by Leveraging Underutilized Cache Resources," *Proceedings of the 56th International Symposium on Microarchitecture (**MICRO**)*, Toronto, ON, Canada, October/November 2023. **Distinguished Artifact Award. Officially artifact evaluated as available, functional, reusable and reproducible.**
51. Konstantinos Kanellopoulos, Rahul Bera, Kosta Stojiljkovic, Nisa Bostanci, Can Firtina, Rachata Ausavarungnirun, Rakesh Kumar, Nastaran Hajinazar, Mohammad Sadrosadati, Nandita Vijaykumar, and Onur Mutlu, "Utopia: Efficient Address Translation using Hybrid Virtual-to-Physical Address Mapping," *Proceedings of the 56th International Symposium on Microarchitecture (**MICRO**)*, Toronto, ON, Canada, October/November 2023.
52. Taha Shahroodi, Gagandeep Singh, Mahdi Zahedi, Haiyu Mao, Joel Lindegger, Can Firtina, Stephan Wong, Onur Mutlu, and Said Hamdioui, "Swordfish: A Framework for Evaluating Deep Neural Network-based Basecalling using Computation-In-Memory with Non-Ideal Memristors," *Proceedings of the 56th International Symposium on Microarchitecture (**MICRO**)*, Toronto, ON, Canada, October/November 2023.
53. Jinfan Chen, Juan Gomez Luna, Izzat El Hajj, YuXin Guo, and Onur Mutlu, "SimplePIM: A Software Framework for Productive and Efficient Processing in Memory," *Proceedings of the 32nd International Conference on Parallel Architectures and Compilation Techniques (**PACT**)*, Vienna, Austria, October 2023.
54. Harshita Gupta, Mayank Kabra, Juan Gómez-Luna, Konstantinos Kanellopoulos, and Onur Mutlu, "Evaluating Homomorphic Operations on a Real-World Processing-In-Memory System," *Proceedings of the 2023 IEEE International Symposium on Workload Characterization (**IISWC**)*, Ghent, Belgium, October 2023.
55. Can Firtina, Nika Mansouri Ghiasi, Joel Lindegger, Gagandeep Singh, Meryem Banu Cavlak, Haiyu Mao, and Onur Mutlu, "RawHash: Enabling Fast and Accurate Real-Time Analysis of Raw Nanopore Signals for Large Genomes," *Proceedings of the 31st Annual Conference on Intelligent Systems for Molecular Biology and the 22nd European Conference on Computational Biology (**ISMB/ECCB**)*, Lyon, France, July 2023.
56. Onur Mutlu, "Memory-Centric Computing," *Invited Lighting Talk Paper in Proceedings of the 60th Design Automation Conference (**DAC**)*, San Francisco, CA, USA, July 2023. **Invited Paper and Presentation.**
57. Onur Mutlu and Can Firtina, "Accelerating Genome Analysis via Algorithm-Architecture Co-Design," *Invited Special Session Paper in Proceedings of the 60th Design Automation Conference (**DAC**)*, San Francisco, CA, USA, July 2023. **Invited Paper and Presentation.**
58. Haocong Luo, Ataberk Olgun, Giray Yaglikci, Yahya Can Tugrul, Steve Rhyner, M. Banu Cavlak, Joel Lindegger, Mohammad Sadrosadati, and Onur Mutlu, "RowPress: Amplifying Read Disturbance in Modern DRAM Chips," *Proceedings of the 50th International Symposium on Computer Architecture (**ISCA**)*, Orlando, FL, USA, June 2023. **Distinguished Artifact Award. Officially artifact evaluated as available, reusable and reproducible.**
59. Rakesh Nadig, Mohammad Sadrosadati, Haiyu Mao, Nika Mansouri Ghiasi, Arash Tavakkol, Jisung Park, Hamid Sarbazi-Azad, Juan Gómez Luna, and Onur Mutlu, "Venice: Improving Solid-State Drive Parallelism at Low Cost via Conflict-Free Accesses," *Proceedings of the 50th International Symposium on Computer Architecture (**ISCA**)*, Orlando, FL, USA, June 2023.
60. Gagandeep Singh, Alireza Khodamoradi, Kristof Denolf, Jack Lo, Juan Gómez-Luna, Joseph Melber, Andra Bisca, Henk Corporaal, and Onur Mutlu, "SPARTA: Spatial Acceleration for Efficient and Scalable Horizontal Diffusion Weather Stencil Computation," *Proceedings of the 37th ACM International Conference on Supercomputing (**ICS**)*, Orlando, FL, USA, June 2023.
61. Ataberk Olgun, Majd Osserian, A. Giray Yaglıkcı, Yahya Can Tugrul, Haocong Luo, Steve Rhyner, Behzad Salami, Juan Gomez-Luna, and Onur Mutlu, "An Experimental Analysis of RowHammer in HBM2 DRAM Chips," *Proceedings of the 53rd Annual IEEE/IFIP International Conference on Dependable Systems and Networks Disrupt Track (**DSN Disrupt**)*, Porto, Portugal, June 2023.
62. Juan Gómez Luna, Yuxin Guo, Sylvan Brocard, Julien Legriel, Remy Cimadomo, Geraldo F. Oliveira, Gagandeep Singh, and Onur Mutlu, "Evaluating Machine Learning Workloads on Memory-Centric Computing Systems," *Proceedings of the 2023 IEEE International Symposium on Performance Analysis of Systems and Software (**ISPASS**)*, Raleigh, NC, USA, April 2023. **Best Paper Session.**
63. Maurus Item, Juan Gómez Luna, Yuxin Guo, Geraldo F. Oliveira, Mohammad Sadrosadati, and Onur Mutlu, "TransPimLib: Efficient Transcendental Functions for Processing-in-Memory Systems," *Proceedings of the 2023 IEEE International Symposium on Performance Analysis of Systems and Software (**ISPASS**)*, Raleigh, NC, USA, April 2023.
64. M. Banu Cavlak, Gagandeep Singh, Mohammed Alser, Can Firtina, Joel Lindegger, Mohammad Sadrosadati, Nika Mansouri Ghiasi, Can Alkan, and Onur Mutlu, "TargetCall: Eliminating the Wasted Computation in Basecalling via Pre-Basecalling Filtering," *Proceedings of the 21st Asia Pacific Bioinformatics Conference (**APBC**)*, Changsha, China, April 2023.
65. Jeremie S. Kim, Can Firtina, M. Banu Cavlak, Damla Senol Cali, Nastaran Hajinazar, Mohammed Alser, Can Alkan, and Onur Mutlu, "AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes," *Proceedings of the 21st Asia Pacific Bioinformatics Conference (**APBC**)*, Changsha, China, April 2023.
66. Nour Almadhoun Alserr, Gulce Kale, Onur Mutlu, Oznur Tastan, and Erman Ayday, "Tuning Privacy-Utility Tradeoff in Genomic Studies Using Selective SNP Hiding," *Proceedings of the 21st Asia Pacific Bioinformatics Conference (**APBC**)*, Changsha, China, April 2023.

67. Lukas Breitwieser, Ahmad Hesam, Fons Rademakers, Juan Gómez Luna, and Onur Mutlu, "High-Performance and Scalable Agent-Based Simulation with BioDynaMo," *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming **(PPoPP)***, Montreal, QC, Canada, February 2023. **Best Artifact Award. Officially artifact evaluated as available, reusable and reproducible.**
68. Onur Mutlu, Ataberk Olgun, and A. Giray Yaglikci, "Fundamentally Understanding and Solving RowHammer," *Invited Special Session Paper in Proceedings of the 28th Asia and South Pacific Design Automation Conference **(ASP-DAC)***, Tokyo, Japan, January 2023. **Invited Paper and Presentation.**
69. Maciej Besta, Cesare Miglioli, Paolo Sylos Labini, Jakub Tětek, Patrick Iff, Raghavendra Kanakagiri, Saleh Ashkboos, Kacper Janda, Michal Podstawski, Grzegorz Kwasniewski, Niels Gleinig, Flavio Vella, Onur Mutlu, and Torsten Hoefler, "ProbGraph: High-Performance and High-Accuracy Graph Mining with Probabilistic Set Representations," *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis **(SC)***, Dallas, TX, USA, November 2022. **Best Paper Award. Officially artifact evaluated as available, reusable and reproducible.**
70. Rahul Bera, Konstantinos Kanellopoulos, Shankar Balachandran, David Novo, Ataberk Olgun, Mohammad Sadrosadati, and Onur Mutlu, "Hermes: Accelerating Long-Latency Load Requests via Perceptron-Based Off-Chip Load Prediction," *Proceedings of the 55th International Symposium on Microarchitecture **(MICRO)***, Chicago, IL, USA, October 2022. **Best Paper Award. Officially artifact evaluated as available, reusable and reproducible.**
71. João Dinis Ferreira, Gabriel Falcao, Juan Gómez-Luna, Mohammed Alser, Lois Orosa, Mohammad Sadrosadati, Jeremie S. Kim, Geraldo F. Oliveira, Taha Shahroodi, Anant Nori, and Onur Mutlu, "pLUTo: Enabling Massively Parallel Computation in DRAM via Lookup Tables," *Proceedings of the 55th International Symposium on Microarchitecture **(MICRO)***, Chicago, IL, USA, October 2022. **Officially artifact evaluated as available, reusable and reproducible.**
72. Jisung Park, Roknoddin Azizi, Geraldo F. Oliveira, Mohammad Sadrosadati, Rakesh Nadig, David Novo, Juan Gómez-Luna, Myungsuk Kim, and Onur Mutlu, "Flash-Cosmos: In-Flash Bulk Bitwise Operations Using Inherent Computation Capability of NAND Flash Memory," *Proceedings of the 55th International Symposium on Microarchitecture **(MICRO)***, Chicago, IL, USA, October 2022.
73. Haiyu Mao, Mohammed Alser, Mohammad Sadrosadati, Can Firtina, Akanksha Baranwal, Damla Senol Cali, Aditya Manglik, Nour Almadhoun Alserr, and Onur Mutlu, "GenPIP: In-Memory Acceleration of Genome Analysis via Tight Integration of Basecalling and Read Mapping," *Proceedings of the 55th International Symposium on Microarchitecture **(MICRO)***, Chicago, IL, USA, October 2022.
74. A. Giray Yaglıkcı, Ataberk Olgun, Minesh Patel, Haocong Luo, Hasan Hassan, Lois Orosa, Oguz Ergin, and Onur Mutlu, "HiRA: Hidden Row Activation for Reducing Refresh Latency of Off-the-Shelf DRAM Chips," *Proceedings of the 55th International Symposium on Microarchitecture **(MICRO)***, Chicago, IL, USA, October 2022.
75. Sina Darabi, Mohammad Sadrosadati, Joël Lindegger, Negar Akbarzadeh, Mohammad Hosseini, Jisung Park, Juan Gómez-Luna, Onur Mutlu, and Hamid Sarbazi-Azad, "Morpheus: Extending the Last Level Cache Capacity in GPU Systems Using Idle GPU Core Resources," *Proceedings of the 55th International Symposium on Microarchitecture **(MICRO)***, Chicago, IL, USA, October 2022.
76. Jawad Haj Yahya, Haris Volos, Davide B. Bartolini, Georgia Antoniou, Jeremie S. Kim, Zhe Wang, Kleovoulos Kalaitzidis, Tom Rollet, Zhirui Chen, Ye Geng, Onur Mutlu, and Yiannakis Sazeides, "AgileWatts: An Energy-Efficient CPU Core Idle-State Architecture for Latency-Sensitive Server Applications," *Proceedings of the 55th International Symposium on Microarchitecture **(MICRO)***, Chicago, IL, USA, October 2022.
77. A. Giray Yağlıkçı, Haocong Luo, Geraldo F. de Oliviera, Ataberk Olgun, Minesh Patel, Jisung Park, Hasan Hassan, Jeremie S. Kim, Lois Orosa, and Onur Mutlu, "Understanding RowHammer Under Reduced Wordline Voltage: An Experimental Study Using Real DRAM Devices," *Proceedings of the 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks **(DSN)***, Baltimore, MD, USA, June 2022.
78. Gagandeep Singh, Rakesh Nadig, Jisung Park, Rahul Bera, Nastaran Hajinazar, David Novo, Juan Gomez-Luna, Sander Stuijk, Henk Corporaal, and Onur Mutlu, "Sibyl: Adaptive and Extensible Data Placement in Hybrid Storage Systems Using Online Reinforcement Learning," *Proceedings of the 49th International Symposium on Computer Architecture **(ISCA)***, New York, NY, USA, June 2022.
79. Damla Senol Cali, Konstantinos Kanellopoulos, Joel Lindegger, Zulal Bingol, Gurpreet S. Kalsi, Ziyi Zuo, Can Firtina, Meryem Banu Cavlak, Jeremie Kim, Nika MansouriGhiasi, Gagandeep Singh, Juan Gomez-Luna, Nour Almadhoun Alserr, Mohammed Alser, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu, "SeGraM: A Universal Hardware Accelerator for Genomic Sequence-to-Graph and Sequence-to-Sequence Mapping," *Proceedings of the 49th International Symposium on Computer Architecture **(ISCA)***, New York, NY, USA, June 2022.
80. Christina Giannoula, Ivan Fernandez, Juan Gomez-Luna, Nectarios Koziris, Georgios Goumas, and Onur Mutlu, "SparseP: Towards Efficient Sparse Matrix Vector Multiplication on Real Processing-In-Memory Architectures," *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems **(SIGMETRICS)***, Mumbai, India, June 2022.
81. Amirali Boroumand, Saugata Ghose, Geraldo F. Oliveira, and Onur Mutlu, "Polynesia: Enabling High-Performance and Energy-Efficient Hybrid Transactional/Analytical Databases with Hardware/Software Co-Design," *Proceedings of the 38th International Conference on Data Engineering **(ICDE)***, Virtual, May 2022.
82. F. Nisa Bostanci, Ataberk Olgun, Lois Orosa, A. Giray Yaglikci, Jeremie S. Kim, Hasan Hassan, Oguz Ergin, and Onur Mutlu, "DR-STRaNGe: End-to-End System Design for DRAM-based True Random Number Generators," *Proceedings of the 28th IEEE International Symposium on High Performance Computer Architecture **(HPCA)***, Virtual, April 2022.

83. Jawad Haj Yahya, Jeremie S. Kim, A. Giray Yaglikci, Jisung Park, Efraim Rotem, Yanos Sazeides, and Onur Mutlu, "DarkGates: A Hybrid Power-Gating Architecture to Mitigate the Performance Impact of Dark-Silicon in High Performance Processors," *Proceedings of the 28th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Virtual, April 2022.
84. Mhd Ghaith Olabi, Juan Gomez Luna, Onur Mutlu, Wen-mei Hwu, and Izzat El Hajj, "A Compiler Framework for Optimizing Dynamic Parallelism on GPUs," *Proceedings of the International Symposium on Code Generation and Optimization **(CGO)***, Virtual, April 2022. **Officially artifact evaluated as available, reusable and reproducible.**
85. Nika Mansouri Ghiasi, Jisung Park, Harun Mustafa, Jeremie Kim, Ataberk Olgun, Arvid Gollwitzer, Damla Senol Cali, Can Firtina, Haiyu Mao, Nour Almadhoun Alserr, Rachata Ausavarungnirun, Nandita Vijaykumar, Mohammed Alser, and Onur Mutlu, "GenStore: A High-Performance and Energy-Efficient In-Storage Computing System for Genome Sequence Analysis," *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Lausanne, Switzerland, February-March 2022.
86. Jisung Park, Jeonggyun Kim, Yeseong Kim, Sungjin Lee, and Onur Mutlu, "DeepSketch: A New Machine Learning-Based Reference Search Technique for Post-Deduplication Delta Compression," *Proceedings of the 20th USENIX Conference on File and Storage Technologies **(FAST)***, Santa Clara, CA, USA, February 2022.
87. Juan Gomez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu, "A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses," *Invited Paper at Workshop on Computing with Unconventional Technologies **(CUT)***, Virtual, October 2021.
88. Lois Orosa, Abdullah Giray Yaglikci, Haocong Luo, Ataberk Olgun, Jisung Park, Hasan Hassan, Minesh Patel, Jeremie S. Kim, and Onur Mutlu, "A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses," *Proceedings of the 54th International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2021.
89. Hasan Hassan, Yahya Can Tugrul, Jeremie S. Kim, Victor van der Veen, Kaveh Razavi, and Onur Mutlu, "Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications," *Proceedings of the 54th International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2021.
90. Rahul Bera, Konstantinos Kanellopoulos, Anant Nori, Taha Shahroodi, Sreenivas Subramoney, and Onur Mutlu, "Pythia: A Customizable Hardware Prefetching Framework Using Online Reinforcement Learning," *Proceedings of the 54th International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2021. **Officially artifact evaluated as available, reusable and reproducible.**
91. Minesh Patel, Geraldo F. de Oliveira Jr., and Onur Mutlu, "HARP: Practically and Effectively Identifying Uncorrectable Errors in Memory Chips That Use On-Die Error-Correcting Codes," *Proceedings of the 54th International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2021. **Officially artifact evaluated as available, reusable and reproducible.**
92. Jawad Haj-Yahya, Jisung Park, Rahul Bera, Juan Gomez Luna, Taha Shahroodi, Jeremie S. Kim, Efraim Rotem, and Onur Mutlu, "BurstLink: Techniques for Energy-Efficient Conventional and Virtual Reality Video Display," *Proceedings of the 54th International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2021.
93. Maciej Besta, Raghavendra Kanakagiri, Grzegorz Kwasniewski, Rachata Ausavarungnirun, Jakub Beránek, Konstantinos Kanellopoulos, Kacper Janda, Zur Vonarburg-Shmaria, Lukas Gianinazzi, Ioana Stefan, Juan Gómez-Luna, Marcin Copik, Lukas Kapp-Schwoerer, Salvatore Di Girolamo, Nils Blach, Marek Konieczny, Onur Mutlu, and Torsten Hoefler, "SISA: Set-Centric Instruction Set Architecture for Graph Mining on Processing-in-Memory Systems," *Proceedings of the 54th International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2021.
94. Sherif A.S. Mohamed, Mohammad-Hashem Haghbayan, Antonio Miele, Onur Mutlu, and Juha Plosila, "Energy-Efficient Mobile Robot Control via Run-time Monitoring of Environmental Complexity and Computing Workload," *Proceedings of the 2021 IEEE/RSJ International Conference on Intelligent Robots and Systems **(IROS)***, Prague, Czech Republic, October 2021.
95. Amirali Boroumand, Saugata Ghose, Berkin Akin, Ravi Narayanaswami, Geraldo F. Oliveira, Xiaoyu Ma, Eric Shiu, and Onur Mutlu, "Google Neural Network Models for Edge Devices: Analyzing and Mitigating Machine Learning Inference Bottlenecks," *Proceedings of the 30th International Conference on Parallel Architectures and Compilation Techniques **(PACT)***, Virtual, September 2021.
96. Maciej Besta, Zur Vonarburg-Shmaria, Yannick Schaffner, Leonardo Schwarz, Grzegorz Kwasniewski, Lukas Gianinazzi, Jakub Beránek, Kacper Janda, Tobias Holenstein, Sebastian Leisinger, Peter Tatkowski, Esref Ozdemir, Adrian Balla, Marcin Copik, Philipp Lindenberger, Pavel Kalvoda, Marek Konieczny, Onur Mutlu, and Torsten Hoefler, "GraphMineSuite: Enabling High-Performance and Programmable Graph Mining Algorithms with Set Algebra," *Proceedings of the 45th International Conference on Very Large Databases **(VLDB)***, Copenhagen, Denmark, August 2021.
97. Jawad Haj-Yahya, Jeremie S. Kim, A. Giray Yaglikci, Ivan Puddu, Lois Orosa, Juan Gomez Luna, Mohammed Alser, and Onur Mutlu, "IChannels: Exploiting Current Management Mechanisms to Create Covert Channels in Modern Processors," *Proceedings of the 48th International Symposium on Computer Architecture **(ISCA)***, Virtual, June 2021.
98. Ataberk Olgun, Minesh Patel, A. Giray Yaglikci, Haocong Luo, Jeremie S. Kim, F. Nisa Bostanci, Nandita Vijaykumar, Oguz Ergin, and Onur Mutlu, "QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips," *Proceedings of the 48th International Symposium on Computer Architecture **(ISCA)***, Virtual, June 2021.
99. Jawad Haj-Yahya, Jeremie S. Kim, A. Giray Yaglikci, Ivan Puddu, Lois Orosa, Juan Gomez Luna, Mohammed Alser, and Onur Mutlu, "IChannels: Exploiting Current Management Mechanisms to Create Covert Channels in Modern Processors," *Proceedings of the 48th International Symposium on Computer Architecture **(ISCA)***, Virtual, June 2021.

100. Lois Orosa, Yaohua Wang, Mohammad Sadrosadati, Jeremie S. Kim, Minesh Patel, Ivan Puddu, Haocong Luo, Kaveh Razavi, Juan Gomez-Luna, Hasan Hassan, Nika Mansouri-Ghiasi, Saugata Ghose, and Onur Mutlu, "CODIC: A Low-Cost Substrate for Enabling Custom In-DRAM Functionalities and Optimizations," *Proceedings of the 48th International Symposium on Computer Architecture **(ISCA)***, Virtual, June 2021.
101. Nastaran Hajinazar, Geraldo F. Oliveira, Sven Gregorio, Joao Dinis Ferreira, Nika Mansouri Ghiasi, Minesh Patel, Mohammed Alser, Saugata Ghose, Juan Gomez-Luna, and Onur Mutlu, "SIMDRAM: An End-to-End Framework for Bit-Serial SIMD Computing in DRAM", *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Virtual, March-April 2021.
102. Jisung Park, Myungsuk Kim, Myoungjun Chun, Lois Orosa, Jihong Kim, and Onur Mutlu, "Reducing Solid-State Drive Read Latency by Optimizing Read-Retry", *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Virtual, March-April 2021.
103. Irina Calciu, M. Talha Imran, Ivan Puddu, Sanidhya Kashyap, Hasan Al Maruf, Onur Mutlu, and Aasheesh Kolli, "Rethinking Software Runtimes for Disaggregated Memory", *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Virtual, March-April 2021. **Officially artifact evaluated as available, reusable and reproducible.**
104. Feng Zhang, Zaifeng Pan, Yanling Zhou, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Xiaoyong Du, "G-TADOC: Enabling Efficient GPU-Based Text Analytics without Decompression", *Proceedings of the 37th International Conference on Data Engineering **(ICDE)***, Virtual, April 2021.
105. Christina Giannoula, Nandita Vijaykumar, Nikela Papadopoulou, Vasileios Karakostas, Ivan Fernandez, Juan Gómez-Luna, Lois Orosa, Nectarios Koziris, Georgios Goumas, and Onur Mutlu, "SynCron: Efficient Synchronization Support for Near-Data-Processing Architectures", *Proceedings of the 27th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Virtual, February-March 2021.
106. A. Giray Yaglikci, Minesh Patel, Jeremie S. Kim, Roknoddin Azizi, Ataberk Olgun, Lois Orosa, Hasan Hassan, Jisung Park, Konstantinos Kanellopoulos, Taha Shahroodi, Saugata Ghose, and Onur Mutlu, "BlockHammer: Preventing RowHammer at Low Cost by Blacklisting Rapidly-Accessed DRAM Rows," *Proceedings of the 27th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Virtual, February-March 2021. **Intel Hardware Security Academic Award Finalist (one of 4 finalists out of 34 nominations).**
107. Onur Mutlu, "Intelligent Architectures for Intelligent Computing Systems," *Invited Paper in Proceedings of the Design, Automation, and Test in Europe Conference **(DATE)***, Virtual, February 2021. **Invited Paper and Presentation.**
108. Seyed Saber Nabavi Larimi, Behzad Salami, Osman S. Unsal, Adrian Cristal Kestelman, Hamid Sarbazi-Azad, and Onur Mutlu, "Understanding Power Consumption and Reliability of High-Bandwidth Memory with Voltage Underscaling," *Proceedings of the Design, Automation, and Test in Europe Conference **(DATE)***, Virtual, February 2021.
109. Shihao Song, Anup Das, Onur Mutlu, and Nagarajan Kandasamy, "Aging-Aware Request Scheduling for Non-Volatile Main Memory," *Proceedings of the 26th Asia and South Pacific Design Automation Conference **(ASP-DAC)***, Virtual, January 2021.
110. Minesh Patel, Jeremie S. Kim, Taha Shahroodi, Hasan Hassan, and Onur Mutlu, "Bit-Exact ECC Recovery (BEER): Determining DRAM On-Die ECC Functions by Exploiting DRAM Data Retention Characteristics," *Proceedings of the 53rd International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2020. **Best Paper Award.**
111. Damla Senol Cali, Gurpreet S. Kalsi, Zulal Bingol, Can Firtina, Lavanya Subramanian, Jeremie S. Kim, Rachata Ausavarungnirun, Mohammed Alser, Juan Gomez-Luna, Amirali Boroumand, Anant Nori, Allison Scibisz, Sreenivas Subramoney, Can Alkan, Saugata Ghose, and Onur Mutlu, "GenASM: A High-Performance, Low-Power Approximate String Matching Acceleration Framework for Genome Sequence Analysis," *Proceedings of the 53rd International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2020.
112. Yaohua Wang, Lois Orosa, Xiangjun Peng, Yang Guo, Saugata Ghose, Minesh Patel, Jeremie S. Kim, Juan Gómez Luna, Mohammad Sadrosadati, Nika Mansouri Ghiasi, and Onur Mutlu, "FIGARO: Improving System Performance via Fine-Grained In-DRAM Data Relocation and Caching," *Proceedings of the 53rd International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2020.
113. Jawad Haj-Yahya, Mohammed Alser, Jeremie S. Kim, Lois Orosa, Efraim Rotem, Avi Mendelson, Anupam Chattopadhyay, and Onur Mutlu, "FlexWatts: A Power- and Workload-Aware Hybrid Power Delivery Network for Energy-Efficient Microprocessors," *Proceedings of the 53rd International Symposium on Microarchitecture **(MICRO)***, Virtual, October 2020.
114. Ivan Fernandez, Ricardo Quislant, Christina Giannoula, Mohammed Alser, Juan Gómez-Luna, Eladio Gutiérrez, Oscar Plata, and Onur Mutlu, "NATSA: A Near-Data Processing Accelerator for Time Series Analysis," *Proceedings of the 38th IEEE International Conference on Computer Design **(ICCD)***, Virtual, October 2020.
115. Leonid Yavits, Lois Orosa, João Dinis Ferreira, Mattan Erez, Ran Ginosar, and Onur Mutlu, "WoLFRaM: Enhancing Wear-Leveling and Fault Tolerance in Resistive Memories using Programmable Address Decoders," *Proceedings of the 38th IEEE International Conference on Computer Design **(ICCD)***, Virtual, October 2020.
116. Onur Mutlu, "Intelligent Architectures for Intelligent Machines," *Keynote Paper in Proceedings of the 2020 International Symposia on VLSI Design, Automation, Test **(VLSI-DAT)***, Hsinchu City, Taiwan, August 2020. **Keynote Paper and Presentation.**
117. Jorge Gonzalez, Alexander Gazman, Maarten Hattink, Mauricio G. Palma, Meisam Bahadori, Ruth Rubio-Noriega, Lois Orosa, Madeleine Glick, Onur Mutlu, Keren Bergman, and Rodolfo Azevedo, "Optically Connected Memory for Disaggregated Data Centers," *Proceedings of the 32nd International Symposium on Computer Architecture and High Performance Computing **(SBAC-PAD)***, Porto, Portugal, September 2020.

118. Gagandeep Singh, Dionysios Diamantopoulos, Christoph Hagleitner, Juan Gómez-Luna, Sander Stuijk, Onur Mutlu, and Henk Corporaal, "NERO: A Near High-Bandwidth Memory Stencil Accelerator for Weather Prediction Modeling," *Proceedings of the 30th International Conference on Field-Programmable Logic and Applications **(FPL)***, Gothenburg, Sweden, September 2020. **One of the four papers nominated for the Stamatis Vassiliadis Memorial Best Paper Award.**
119. Kevin Hsieh, Amar Phanishayee, Onur Mutlu, and Phillip B. Gibbons, "The Non-IID Data Quagmire of Decentralized Machine Learning," *Proceedings of the 37th International Conference on Machine Learning **(ICML)***, Virtual, June 2020.
120. Jeremie S. Kim, Minesh Patel, A. Giray Yaglikci, Hasan Hassan, Roknoddin Azizi, Lois Orosa, and Onur Mutlu, "Revisiting RowHammer: An Experimental Analysis of Modern Devices and Mitigation Techniques," *Proceedings of the 47th International Symposium on Computer Architecture **(ISCA)***, Valencia, Spain, July 2020.
121. Haocong Luo, Taha Shahroodi, Hasan Hassan, Minesh Patel, A. Giray Yaglikci, Lois Orosa, Jisung Park, and Onur Mutlu, "CLR-DRAM: A Low-Cost DRAM Architecture Enabling Dynamic Capacity-Latency Trade-Off," *Proceedings of the 47th International Symposium on Computer Architecture **(ISCA)***, Valencia, Spain, June 2020.
122. Nastaran Hajinazar, Pratyush Patel, Minesh Patel, Konstantinos Kanellopoulos, Saugata Ghose, Rachata Ausavarungnirun, Geraldo Francisco de Oliveira Jr., Jonathan Appavoo, Vivek Seshadri, and Onur Mutlu, "The Virtual Block Interface: A Flexible Alternative to the Conventional Virtual Memory Framework," *Proceedings of the 47th International Symposium on Computer Architecture **(ISCA)***, Valencia, Spain, June 2020.
123. Jawad Haj-Yahya, Mohammed Alser, Jeremie Kim, A. Giray Yaglikci, Nandita Vijaykumar, Efraim Rotem, and Onur Mutlu, "SysScale: Exploiting Multi-domain Dynamic Voltage and Frequency Scaling for Energy Efficient Mobile Processors," *Proceedings of the 47th International Symposium on Computer Architecture **(ISCA)***, Valencia, Spain, June 2020.
124. Behzad Salami, Erhan Baturay Onural, Ismail Emir Yuksel, Fahrettin Koc, Oguz Ergin, Adrian Cristal Kestelman, Osman S. Unsal, Hamid Sarbazi-Azad, and Onur Mutlu, "An Experimental Study of Reduced-Voltage Operation in Modern FPGAs for Neural Network Acceleration," *Proceedings of the 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks **(DSN)***, Valencia, Spain, June 2020.
125. Shihao Song, Anup Das, Onur Mutlu, and Nagarajan Kandasamy, "Improving Phase Change Memory Performance with Data Content Aware Access," *Proceedings of the ACM SIGPLAN International Symposium on Memory Management **(ISMM)***, London, UK, June 2020.
126. Lucian Cojocar, Jeremie Kim, Minesh Patel, Lillian Tsai, Stefan Saroiu, Alec Wolman, and Onur Mutlu, "Are We Susceptible to Rowhammer? An End-to-End Methodology for Cloud Providers," *Proceedings of the 41st IEEE Symposium on Security and Privacy **(S&P)***, San Francisco, CA, USA, May 2020.
127. Pietro Frigo, Emanuele Vannacci, Hasan Hassan, Victor van der Veen, Onur Mutlu, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi, "TRRespass: Exploiting the Many Sides of Target Row Refresh," *Proceedings of the 41st IEEE Symposium on Security and Privacy **(S&P)***, San Francisco, CA, USA, May 2020. **Best Paper Award. Pwnie Award 2020 for Most Innovative Research. Top Picks Honorable Mention by IEEE Micro.**
128. Feng Zhang, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Xiaoyong Du, "Enabling Efficient Random Access to Hierarchically-Compressed Data," *Proceedings of the 36th International Conference on Data Engineering **(ICDE)***, Dallas, TX, USA, April 2020.
129. Myungsuk Kim, Jisung Park, Geonhee Cho, Yoona Kim, Lois Orosa, Onur Mutlu, and Jihong Kim, "Evanesco: Architectural Support for Efficient Data Sanitization in Modern Flash-Based Storage Systems," *Proceedings of the 25th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Lausanne, Switzerland, April 2020.
130. Jiantong Jiang, Zeke Wang, Xue Liu, Juan Gómez-Luna, Nan Guan, Qingxu Deng, Wei Zhang, and Onur Mutlu, "Boyi: A Systematic Framework for Automatically Deciding the Right Execution Model of OpenCL Applications on FPGAs," *Proceedings of the 28th International Symposium on Field-Programmable Gate Arrays (**FPGA**)*, Seaside, CA USA, February 2020.
131. Jawad Haj-Yahya, Yanos Sazeides, Mohammed Alser, Efraim Rotem, and Onur Mutlu, "Techniques for Reducing the Connected-Standby Energy Consumption of Mobile Devices," *Proceedings of the 26th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, San Diego, CA USA, February 2020.
132. Konstantinos Kanellopoulos, Christina Giannoula, Roknoddin Azizi, Skanda Koppula, Nika Mansouri Ghiasi, Juan Gomez-Luna, Nandita Vijaykumar, Taha Shahroodi, and Onur Mutlu, "SMASH: Co-designing Software Compression and Hardware-Accelerated Indexing for Efficient Sparse Matrix Operations," *Proceedings of the 52nd IEEE/ACM International Symposium on Microarchitecture (**MICRO**)*, Columbus, OH, USA, October 2019.
133. Skanda Koppula, Lois Orosa, Konstantinos Kanellopoulos, Taha Shahroodi, Roknoddin Azizi, A. Giray Yaglikci, and Onur Mutlu, "EDEN: Energy-Efficient, High-Performance Neural Network Inference Using Approximate DRAM," *Proceedings of the 52nd IEEE/ACM International Symposium on Microarchitecture (**MICRO**)*, Columbus, OH, USA, October 2019.
134. Xiao Liu, David Roberts, Rachata Ausavarungnirun, Onur Mutlu, and Jishen Zhao, "Binary Star: Coordinated Reliability in Heterogeneous Memory Systems for High Performance and Scalability," *Proceedings of the 52nd IEEE/ACM International Symposium on Microarchitecture (**MICRO**)*, Columbus, OH, USA, October 2019.
135. Rahul Bera, Anant V. Nori, Onur Mutlu, and Sreenivas Subramoney, "DSPatch: Dual Spatial Pattern Prefetcher," *Proceedings of the 52nd IEEE/ACM International Symposium on Microarchitecture (**MICRO**)*, Columbus, OH, USA, October 2019.
136. Shihao Song, Anup Das, Onur Mutlu, and Nagarajan Kandasamy, "Enabling and Exploiting Partition-Level Parallelism (PALP) in Phase Change Memories," *Proceedings of the International Conference on Compilers, Architecture, and Synthesis for Embedded Systems (**CASES**)*, New York City, NY, USA, October 2019.

137. Zeke Wang, Kaan Kara, Hantian Zhang, Gustavo Alonso, Onur Mutlu, and Ce Zhang, "Accelerating Generalized Linear Models with MLWeaving: A One-Size-Fits-All System for Any-Precision Learning," *Proceedings of the 45th International Conference on Very Large Databases (**VLDB**)*, Los Angeles, CA, USA, August 2018.
138. Hasan Hassan, Minesh Patel, Jeremie S. Kim, A. Giray Yaglikci, Nandita Vijaykumar, Nika Mansourighiasi, Saugata Ghose, Onur Mutlu, "CROW: A Low-Overhead Substrate for Improving DRAM Performance and Energy-Efficiency," *Proceedings of the 46th International Symposium on Computer Architecture (**ISCA**)*, Phoenix, AZ, USA, June 2019.
139. Amirali Boroumand, Saugata Ghose, Minesh Patel, Rachata Ausavarungnirun, Hasan Hassan, Brandon Lucia, Kevin Hsieh, Nastaran Hajinazar, Krishna T. Malladi, Hongzhong Zheng, Onur Mutlu, "CoNDA: Enabling Efficient Near-Data Accelerator Communication by Optimizing Data Movement," *Proceedings of the 46th International Symposium on Computer Architecture (**ISCA**)*, Phoenix, AZ, USA, June 2019.
140. Minesh Patel, Jeremie Kim, Hasan Hassan, Onur Mutlu, "Understanding and Modeling On-Die Error Correction in Modern DRAM: An Experimental Study Using Real Devices," *Proceedings of the 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (**DSN**)*, Portland, OR, USA, June 2019. **Best Paper Award.**
141. Saugata Ghose, Tianshi Li, Nastaran Hajinazar, Damla Senol Cali, and Onur Mutlu, "Demystifying Workload–DRAM Interactions: An Experimental Study," *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**)*, Phoenix, AZ, USA, June 2019.
142. Chenxi Wang, Huimin Cui, Harry Xu, Ting Cao, John Zigman, Haris Volos, Onur Mutlu, Fang Lv, Xiaobing Feng, "Panthera: Holistic Memory Management for Big Data Processing over Hybrid Memories," *Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (**PLDI**)*, Phoenix, AZ, June 2019.
143. Gagandeep Singh, Stefano Corda, Geraldo Francisco de Oliveira, Juan Gomez-Luna, Giovanni Mariani, Sander Stujik, Onur Mutlu, Henk Corporaal, "NAPEL: Near-Memory Computing Application Performance Prediction via Ensemble Learning," *Proceedings of the 56th Design Automation Conference (**DAC**)*, San Francisco, CA, USA, June 2019.
144. Onur Mutlu, Saugata Ghose, Juan Gomez-Luna, and Rachata Ausavarungnirun, "Enabling Practical Processing in and near Memory for Data-Intensive Computing," *Proceedings of the 56th Design Automation Conference (**DAC**)*, San Francisco, CA, USA, June 2019. **Invited Paper and Presentation.**
145. Irina Calciu, Ivan Puddu, Aasheesh Kolli, Andreas Nowatzyk, Jayneel Gandhi, Onur Mutlu, Pratap Subrahmanyam, "Project PBerry: FPGA Acceleration for Remote Memory," *Proceedings of the 17th Workshop on Hot Topics in Operating Systems (**HotOS**)*, Bertinoro, Italy, May 2019.
146. Sitao Huang, Li-Wen Chang, Izzat El Hajj, Simon Garcia De Gonzalo, Juan Gomez-Luna, Sai Rahul Chalamalasetti, Mohamed El-Hadedy, Dejan Milojicic, Onur Mutlu, Deming Chen, Wen-mei Hwu, "Analysis and Modeling of Collaborative Execution Strategies for Heterogeneous CPU-FPGA Architectures," *Proceedings of the 10th ACM/SPEC International Conference on Performance Engineering **(ICPE)***, Mumbai, India, April 2019.
147. Chen Li, Rachata Ausavarungnirun, Christopher J. Rossbach, Youtao Zhang Onur Mutlu, Yang Guo, Jun Yang, "A Framework for Memory Oversubscription Management in Graphics Processing Units," *Proceedings of the 24th ACM International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)***, Providence, RI, USA, April 2019.
148. Onur Mutlu, "RowHammer and Beyond," *Proceedings of the 10th International Workshop on Constructive Side-Channel Analysis and Secure Design (**COSADE**)*, Darmstadt, Germany, April 2019. **Keynote Paper and Presentation.**
149. Jeremie S. Kim, Minesh Patel, Hasan Hassan, Lois Orosa, Onur Mutlu, "D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput," *Proceedings of the 25th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Washington, DC USA, February 2019.
150. Yang Li, Charles Lefurgy, Karthick Rajamani, Malcolm Ware, Guillermo J. Silva, Daniel D. Heimsoth, Saugata Ghose, Onur Mutlu, 'Challenges in Power Management Adoption for Public Clouds," *Proceedings of the 25th IEEE International Symposium on High Performance Computer Architecture (**HPCA**)*, Washington, DC USA, February 2019.
151. Simon Garcia de Gonzalo, Sitao Huang, Juan Gomez-Luna, Simon D. Hammond, Onur Mutlu, Wen-mei Hwu, "Automatic Generation of Warp-Level Primitives and Atomic Operations for Fast-Portable GPU Reductions," *Proceedings of the 17th Annual International Symposium on Code Generation and Optimization (**CGO**)*, Washington, DC USA, February 2019.
152. Justin Meza, Tianyin Xu, Kaushik Veeraraghavan, and Onur Mutlu, "A Large Scale Study of Data Center Network Reliability," *Proceedings of the 18th ACM Internet Measurement Conference (**IMC**)*, Boston, MA, USA, October/November 2018.
153. Kevin Hsieh, Ganesh Ananthanarayanan, Peter Bodik, Shivaram Venkataraman, Paramvir Bahl, Matthai Philipose, Phillip B. Gibbons, and Onur Mutlu, "Focus: Querying Large Video Datasets with Low Latency and Low Cost," *Proceedings of the 13th USENIX Symposium on Operating Systems Design and Implementation (**OSDI**)*, Carlsbad, CA, USA, October 2018.
154. Yaohua Wang, Arash Tavakkol, Lois Orosa, Saugata Ghose, Nika Mansouri Ghiasi, Minesh Patel, Jeremie S. Kim, Hasan Hassan, Mohammad Sadrosadati, and Onur Mutlu, "Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration"," *Proceedings of the 51st IEEE/ACM International Symposium on Microarchitecture (**MICRO**)*, Fukuoka, Japan, October 2018.
155. Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu, "Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines," *Proceedings of the 36th IEEE International Conference on Computer Design **(ICCD)***, Orlando, FL, USA, October 2018.

156. Feng Zhang, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Wenguang Chen, "Efficient Document Analytics on Compressed Data: Method, Challenges, Algorithms, Insights," *Proceedings of the 44th International Conference on Very Large Databases (**VLDB**)*, Rio de Janeiro, Brazil, August 2018.
157. 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.
158. 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.
159. 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.
160. 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**)*, Beiing, China, June 2018.
161. 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.
162. Saugata Ghose, A. Giray Yaglikci, Raghav Gupta, Donghyuk Lee, Kais Kudrolli, William X. Liu, Hasan Hassan, Kevin K. Chang, Niladrish Chatterjee, Aditya Agrawal, Mike O'Connor, and Onur Mutlu, "What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study," *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**)*, Irvine, CA, USA, June 2018.
163. Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, and Onur Mutlu, "Improving 3D NAND Flash Memory Lifetime by Tolerating Early Retention Loss and Process Variation," *Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**)*, Irvine, CA, USA, June 2018.
164. 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.
165. 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.
166. 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.
167. 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.
168. Amir M. Rahmani, Bryan Donyanavard, Tiago Muuck, 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.
169. 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.
170. 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.
171. 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.
172. 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.
173. 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.

174. 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.

175. 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.

176. 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.

177. 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.

178. 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.

179. 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.

180. 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.

181. 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.

182. 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.

183. 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.

184. 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.

185. 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.**

186. 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.**

187. 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.

188. 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.

189. 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.

190. 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.

191. 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.

192. 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.**

193. 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.
194. 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.
195. 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.
196. Onur Kayiran, Adwait Jog, Ashutosh Pattnaik, Rachata Ausavarungnirun, Xulong Tang, Mahmut T. Kandemir, Gabriel H. Loh, Onur Mutlu, Chita R. Das, "μC-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.
197. 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.
198. Kevin Hsieh, Eiman Ebrahimi, Gwangsun Kim, Niladrish 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.
199. 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.
200. 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.
201. 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.
202. 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.**
203. 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.
204. 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.
205. 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.
206. 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.
207. 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.
208. 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.
209. 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.
210. 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.
211. 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.

212. 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.**
213. 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.
214. 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.
215. 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.
216. 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. Selected to the 50th Anniversary of ISCA Special Commemorative Issue.**
217. 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.
218. 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.
219. 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.
220. 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. **DSN Test of Time Award 2021.**
221. 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.
222. 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.
223. 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.
224. 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.
225. 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.**
226. 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.
227. 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.
228. 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.
229. 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.
230. 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.
231. 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.

232. 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.

233. 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.

234. 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.

235. 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. **One of the 7 papers of 2012-2017 selected as "Top Picks in Hardware and Embedded Security" by IEEE TCAD in 2018. Selected to the 50th Anniversary of ISCA Special Commemorative Issue. Winner of the 2024 IFIP Jean-Claude Laprie Award in Dependable Computing.**

236. 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.

237. 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.

238. 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.

239. 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.

240. 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.**

241. 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.

242. 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.**

243. 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.

244. 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.

245. 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.

246. 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.

247. 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.

248. 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. **Selected to the 50th Anniversary of ISCA Special Commemorative Issue.**

249. 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.

250. 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.

251. 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.

252. Onur Mutlu, "Memory Scaling: A Systems Architecture Perspective," *Proceedings of the 5th International Memory Workshop **(IMW)***, Monterey, CA, May 2013. **Invited Paper and Presentation.**

253. 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.
254. 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.
255. 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.
256. 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.
257. 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.
258. 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.
259. 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.
260. 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.
261. 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)**.
262. 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.
263. 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.
264. 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.
265. 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.
266. 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. **Selected to the 50th Anniversary of ISCA Special Commemorative Issue.**
267. 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.
268. 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.
269. 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.**
270. 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.
271. 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.
272. 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.

273. 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.
274. 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.
275. 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.**
276. 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.
277. Howard David, Chris Fallin, Eugene Gorbatov, 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.
278. 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.
279. 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.
280. 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.**
281. 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.
282. 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.
283. 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.**
284. 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.**
285. 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.
286. 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.
287. 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.**
288. 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.**
289. 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.**
290. 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.
291. 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.
292. 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.
293. 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.

294. 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.

295. 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.

296. 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.

297. 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.**

298. 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.**

299. 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.**

300. 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.

301. 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.

302. 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.

303. 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.

304. 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.**

305. 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. **Selected to the 50th Anniversary of ISCA Special Commemorative Issue.**

306. 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.

307. 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.

308. 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.

309. 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.

310. 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.

311. 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.

312. 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.

313. 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.**

314. 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.**
315. 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.
316. Hyesoon Kim, M. Aater Suleman, Onur Mutlu, Yale N. Patt, "2D-Profiling: 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.
317. 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.**
318. 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.**
319. 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.**
320. 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.
321. 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.
322. 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.**
323. 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 **(WMPI)**, in conjunction with the 31st International Symposium on Computer Architecture*, pp. 56-64, Munchen, Germany, June 2004.
324. 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. HPCA Test of Time Award 2021.**

Refereed and Invited Journal Publications

325. William A. Simon, Irem Boybat, Riselda Kodra, Elena Ferro, Gagandeep Singh, Mohammed Alser, Shubham Jain, Hsinyu Tsai, Geoffrey W. Burr, Onur Mutlu, and Abu Sebastian, "CiMBA: Accelerating Genome Sequencing Through On-Device Basecalling via Compute-in-Memory," ***IEEE Transactions on Parallel and Distributed Systems (TPDS)***, June 2025.
326. Mohammed Alser, Julien Eudine, and Onur Mutlu, "Taming Large-Scale Genomic Analyses via Sparsified Genomics," ***Nature Communications***, 21 January 2025.
327. Shaopeng Liu, Judith S. Rodriguez, Viorel Munteanu, Cynthia Ronkowski, Nitesh Kumar Sharma, Mohammed Alser, Francesco Andreace, Ran Blekhman, Dagmara Błaszczyk, Rayan Chikhi, Keith A. Crandall, Katja Della Libera, Dallace Francis, Alina Frolova, Abigail Shahar Gancz, Naomi E. Huntley, Pooja Jaiswal, Tomasz Kosciolek, Pawel P. Łabaj, Wojciech Łabaj, Tu Luan, Christopher Mason, Ahmed M. Moustafa, Harihara Subrahmaniam Muralidharan, Onur Mutlu, Nika Mansouri Ghiasi, Ali Rahnavard, Fengzhu Sun, Shuchang Tian, Braden T. Tierney, Emily Van Syoc, Riccardo Vicedomini, Joseph P. Zackular, Alex Zelikovsky, Kinga Zielińska, Erika Ganda, Emily R. Davenport, Mihai Pop, David Koslicki, and Serghei Mangul, "Analysis of Metagenomic Data," ***Nature Reviews: Methods Primers***, 23 January 2025.
328. Joel Lindegger, Can Firtina, Nika Mansouri Ghiasi, Mohammad Sadrosadati, Mohammed Alser, and Onur Mutlu, "RawAlign: Accurate, Fast, and Scalable Raw Nanopore Signal Mapping via Combining Seeding and Alignment," ***IEEE Access***, December 2024.
329. Minesh Patel, Taha Shahroodi, Aditya Manglik, A. Giray Yaglikci, Ataberk Olgun, Haocong Luo, and Onur Mutlu, "Rethinking the Producer-Consumer Relationship in Modern DRAM-Based Systems," ***IEEE Access***, December 2024.
330. Alejandro Alonso-Marín, Ivan Fernandez, Quim Aguado-Puig, Juan Gómez-Luna, Santiago Marco-Sola, Onur Mutlu, and Miquel Moreto, "BIMSA: Accelerating Long Sequence Alignment Using Processing-In-Memory," ***Bioinformatics***, 21 October 2024.
331. Adnan Mehonic, Daniele Ielmini, Kaushik Roy, Onur Mutlu, Shahar Kvatinsky, Teresa Serrano-Gotarredona, Bernabe Linares-Barranco, Sabina Spiga, Sergey Savelev, Alexander G Balanov, Nitin Chawla, Giuseppe Desoli, Gerardo Malavena, Christian Monzio Compagnoni, Zhongrui Wang, J Joshua Yang, Ghazi Sarwat Syed, Abu Sebastian, Thomas Mikolajick, Beatriz Noheda, Stefan Slesazeck, Bernard

Dieny, Tuo-Hung (Alex) Hou, Akhil Varri, Frank Bruckerhoff-Pluckelmann, Wolfram Pernice, Xixiang Zhang, Sebastian Pazos, Mario Lanza, Stefan Wiefels, Regina Dittmann, Wing H Ng, Mark Buckwell, Horatio RJ Cox, Daniel J Mannion, Anthony J Kenyon, Yingming Lu, Yuchao Yang, Damien Querlioz, Louis Hutin, Elisa Vianello, Sayeed Shafayet Chowdhury, Piergiulio Mannocci, Yimao Cai, Zhong Sun, Giacomo Pedretti, John Paul Strachan, Dmitri Strukov, Manuel Le Gallo, Stefano Ambrogio, Ilia Valov, and Rainer Waser, "Roadmap to Neuromorphic Computing with Emerging Technologies," ***APL Materials***, 21 October 2024.

332. M. Banu Cavlak, Gagandeep Singh, Mohammed Alser, Can Firtina, Joel Lindegger, Mohammad Sadrosadati, Nika Mansouri Ghiasi, Can Alkan, and Onur Mutlu, "TargetCall: Eliminating the Wasted Computation in Basecalling via Pre-Basecalling Filtering," ***Frontiers in Genetics***, 30 September 2024.
333. Jeremie S. Kim, Can Firtina, Meryem Banu Cavlak, Damla Senol Cali, Nastaran Hajinazar, Mohammed Alser, Can Alkan, and Onur Mutlu, "AirLift: A Fast and Comprehensive Technique for Remapping Alignments between Reference Genomes," ***IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)***, August 2024.
334. Can Firtina, Melina Soysal, Joël Lindegger, and Onur Mutlu, "RawHash2: Mapping Raw Nanopore Signals Using Hash-Based Seeding and Adaptive Quantization," ***Bioinformatics***, 30 July 2024.
335. Ataberk Olgun, F. Nisa Bostanci, Geraldo F. Oliveira, Yahya Can Tugrul, Rahul Bera, A. Giray Yaglikci, Hasan Hassan, Oguz Ergin, and Onur Mutlu, "Sectored DRAM: A Practical Energy-Efficient and High-Performance Fine-Grained DRAM Architecture," ***ACM Transactions on Architecture and Code Optimization (TACO)***, June 2024.
336. Haocong Luo, Ismail Emir Yüksel, Ataberk Olgun, A. Giray Yağlıkçı, Mohammad Sadrosadati, and Onur Mutlu, "RowPress Vulnerability in Modern DRAM Chips," ***IEEE Micro***, Special Issue: Micro's Top Picks from 2023 Computer Architecture Conferences, July/August 2024.
337. Lois Orosa, Ulrich Ruhrmair, A. Giray Yaglikci, Haocong Luo, Ataberk Olgun, Patrick Jattke, Minesh Patel, Jeremie S. Kim, Kaveh Razavi, and Onur Mutlu, "SpyHammer: Understanding and Exploiting RowHammer Under Fine-Grained Temperature Variations," ***IEEE Access***, June 2024.
338. Zulal Bingol, Mohammed Alser, Onur Mutlu, Ozcan Ozturk, and Can Alkan, "GateKeeper-GPU: Fast and Accurate Pre-Alignment Filtering in Short Read Mapping," ***IEEE Transactions on Computers***, 2024.
339. Hajar Falahati, Mohammad Sadrosadati, Qiumin Xu, Juan Gómez-Luna, Banafsheh Saber Latibari, Hyeran Jeon, Shaahin Hesaabi, Hamid Sarbazi-Azad, Onur Mutlu, Murali Annavaram, and Massoud Pedram, "Cross-core Data Sharing for Energy-efficient GPUs," ***ACM Transactions on Architecture and Code Optimization (TACO)***, 2024.
340. Georgia Antoniou, Davide Bartolini, Haris Volos, Marios Kleanthous, Zhe Wang, Kleovoulos Kalaitzidis, Tom Rollet, Ziwei Li, Onur Mutlu, Yiannakis Sazeides, and Jawad Haj Yahya, "Agile C-states: A Core C-state Architecture for Latency Critical Applications Optimizing both Transition and Cold-Start Latency," ***ACM Transactions on Architecture and Code Optimization (TACO)***, 2024.
341. Mohammed Alser, Brendan Lawlor, Richard J. Abdill, Sharon Waymost, Ram Ayyala, Neha Rajkumar, Nathan LaPierre, Jaqueline Brito, Andre M. Ribeiro-dos-Santos, Nour Almadhoun, Varuni Sarwal, Can Firtina, Tomasz Osinski, Eleazar Eskin, Qiyang Hu, Derek Strong, Byoung-Do (B. D) Kim, Malak S. Abedalthagafi, Onur Mutlu, and Serghei Mangul, ""Packaging and Containerization of Computational Methods," ***Nature Protocols***, 2 April 2024.
342. Haocong Luo, Yahya Can Tugrul, F. Nisa Bostanci, Ataberk Olgun, A. Giray Yaglikci, and Onur Mutlu, "Ramulator 2.0: A Modern, Modular, and Extensible DRAM Simulator," ***IEEE Computer Architecture Letters (CAL)***, March 2024.
343. Jie Zhang, Hongjing Huang, Jie Sun, Juan Gomez Luna, Onur Mutlu, and Zeke Wang, "SparseACC: A Generalized Linear Model Accelerator for Sparse Datasets," ***IEEE Transactions on Computer Aided Design (TCAD)***, 2024.
344. Ivan Fernandez, Christina Giannoula, Aditya Manglik, Ricardo Quislant, Nika Mansouri Ghiasi, Juan Gomez Luna, Eladio Gutierrez, Oscar Plata, and Onur Mutlu, "MATSA: An MRAM-Based Energy-Efficient Accelerator for Time Series Analysis," ***IEEE Access***, March 2024.
345. Deepanjali Mishra, Konstantinos Kanellopoulos, Ashish Panwar, Akshitha Sriraman, Vivek Seshadri, Onur Mutlu, and Todd C. Mowry, "Address Scaling: Architectural Support for Fine-Grained Thread-Safe Metadata Management," ***IEEE Computer Architecture Letters (CAL)***, March 2024.
346. Can Firtina, Kamlesh Pillai, Gurpreet S. Kalsi, Bharathwaj Suresh, Damla Senol Cali, Jeremie S. Kim, Taha Shahroodi, Meryem Banu Cavlak, Joel Lindegger, Mohammed Alser, Juan Gomez Luna, Sreenivas Subramoney, and Onur Mutlu, "ApHMM: Accelerating Profile Hidden Markov Models for Fast and Energy-efficient Genome Analysis," ***ACM Transactions on Architecture and Code Optimization (TACO)***, February 2024.
347. Gagandeep Singh, Mohammed Alser, Kristof Denolf, Can Firtina, Alireza Khodamoradi, Meryem Banu Cavlak, Henk Corporaal, and Onur Mutlu, "RUBICON: A Framework for Designing Efficient Deep Learning-Based Genomic Basecallers," ***Genome Biology***, 16 February 2024.
348. Geraldo F. Oliveira, Saugata Ghose, Juan Gómez-Luna, Amirali Boroumand, Alexis Savery, Sonny Rao, Salman Qazi, Gwendal Grignou, Rahul Thakur, Eric Shiu, and Onur Mutlu, "Extending Memory Capacity in Modern Consumer Systems with Emerging Non-Volatile Memory: Experimental Analysis and Characterization using the Intel Optane SSD," ***IEEE Access***, September 2023.
349. Irina Calciu, M. Talha Imran, Ivan Puddu, Sanidhya Kashyap, Hasan Al Maruf, Onur Mutlu, and Aasheesh Kolli, "Using Local Cache Coherence for Disaggregated Memory Systems," ***ACM SIGOPS Operating Systems Review (OSR)***, 2023.

350. Ataberk Olgun, Hasan Hassan, A Giray Yağlıkçı, Yahya Can Tuğrul, Lois Orosa, Haocong Luo, Minesh Patel, Oğuz Ergin, and Onur Mutlu, "DRAM Bender: An Extensible and Versatile FPGA-based Infrastructure to Easily Test State-of-the-art DRAM Chips," ***IEEE Transactions on Computer Aided Design (TCAD)***, 2023.
351. Hashem Haghbayan, Antonio Miele, Onur Mutlu, and Juha Plosila, "Run-time Resource Management in CMPs Handling Multiple Aging Mechanisms," ***IEEE Transactions on Computers***, 2023.
352. Lois Orosa, Skanda Koppula, Yaman Umuroglu, Konstantinos Kanellopoulos, Juan Gomez-Luna, Michaela Blott, Kees Vissers, and Onur Mutlu, "EcoFlow: Efficient Convolutional Dataflows for Low-Power Neural Network Accelerators," ***IEEE Transactions on Computers***, 2023.
353. Joël Lindegger, Damla Senol Cali, Mohammed Alser, Juan Gómez-Luna, Nika Mansouri Ghiasi, and Onur Mutlu, "Scrooge: A Fast and Memory-Frugal Genomic Sequence Aligner for CPUs, GPUs, and ASICs," ***Bioinformatics***, 24 March 2023.
354. Safaa Diab, Amir Nassereldine, Mohammed Alser, Juan Gómez Luna, Onur Mutlu, and Izzat El Hajj, "A Framework for High-throughput Sequence Alignment using Real Processing-in-Memory Systems," ***Bioinformatics***, 27 March 2023.
355. Can Firtina, Jisung Park, Mohammed Alser, Jeremie S. Kim, Damla Senol Cali, Taha Shahroodi, Nika Mansouri Ghiasi, Gagandeep Singh, Konstantinos Kanellopoulos, Can Alkan, Onur Mutlu, "BLEND: A Fast, Memory-Efficient, and Accurate Mechanism to Find Fuzzy Seed Matches in Genome Analysis," ***Nucleics Acids Research: Genomics and Bioinformatics***, March 2023.
356. Alain Denzler, Rahul Bera, Nastaran Hajinazar, Gagandeep Singh, Geraldo F. Oliveira, Juan Gómez-Luna, and Onur Mutlu, "Casper: Accelerating Stencil Computation using Near-cache Processing," ***IEEE Access***, 3 March 2023.
357. Ataberk Olgun, Juan Gomez Luna, Konstantinos Kanellopoulos, Behzad Salami, Hasan Hassan, Oguz Ergin, and Onur Mutlu, "PiDRAM: A Holistic End-to-end FPGA-based Framework for Processing-in-DRAM," ***ACM Transactions on Architecture and Code Optimization (TACO)***, March 2023.
358. Geraldo F. Oliveira, Juan Gómez-Luna, Saugata Ghose, Amirali Boroumand, and Onur Mutlu, "Accelerating Neural Network Inference With Processing-in-DRAM: From the Edge to the Cloud," ***IEEE Micro***, Vol. 42, No. 6, pages 25-38, November/December 2022.
359. Nika Mansouri Ghiasi, Nandita Vijaykumar, Geraldo F. Oliveira, Lois Orosa, Ivan Fernandez, Mohammad Sadrosadati, Konstantinos Kanellopoulos, Nastaran Hajinazar, Juan Gómez Luna, and Onur Mutlu, "ALP: Alleviating CPU-Memory Data Movement Overheads in Memory-Centric Systems," ***IEEE Transactions on Emerging Topics in Computing (TETC)***, December 2022.
360. Jeremie S. Kim, Can Firtina, Meryem Banu Cavlak, Damla Senol Cali, Can Alkan, Onur Mutlu, "FastRemap: A Tool for Quickly Remapping Reads between Genome Assemblies," ***Bioinformatics***, 1 October 2022.
361. Mostafa Hadizadeh, Elham Cheshmikhani, Maysam Rahmanpour, Onur Mutlu, Hossein Asadi, "CoPA: Cold Page Awakening to Overcome Retention Failures in STT-MRAM Based I/O Buffers," ***IEEE Transactions on Parallel and Distributed Systems (TPDS)***, October 2022.
362. Mohammed Alser, Joel Lindegger, Can Firtina, Nour Almadhoun Alserr, Haiyu Mao, Gagandeep Singh, Juan Gomez-Luna, and Onur Mutlu, "From molecules to genomic variations: Accelerating genome analysis via intelligent algorithms and architectures," ***Computational and Structural Biotechnology Journal (CSBJ)***, August 2022.
363. Taha Shahroodi, Mahdi Zahedi, Can Firtina, Mohammed Alser, Stephan Wong, Onur Mutlu, and Said Hamdioui, "Demeter: A Fast and Energy-Efficient Food Profiler Using Hyperdimensional Computing in Memory," ***IEEE Access***, 1 August 2022.
364. Zaifeng Pan, Feng Zhang, Yanliang Zhou, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Xiaoyong Du, "Exploring Data Analytics Without Decompression on Embedded GPU Systems," ***IEEE Transactions on Parallel and Distributed Systems (TPDS)***, July 2022.
365. Mohammed Alser, Jeremie S. Kim, Nour Almadhoun Alserr, Stefan W. Tell, and Onur Mutlu, "COVIDHunter: COVID-19 Pandemic Wave Prediction and Mitigation via Seasonality Aware Modeling," ***Frontiers in Public Health***, June 2022.
366. Gagandeep Singh, Dionysios Diamantopoulos, Christoph Hagleitner, Juan Gómez-Luna, Sander Stuijk, Henk Corporaal, and Onur Mutlu, "Accelerating Weather Prediction Using Near-Memory Reconfigurable Fabric," ***ACM Transactions on Reconfigurable Technology and Systems (TRETS)***, June 2022.
367. Nandita Vijaykumar, Ataberk Olgun, Konstantinos Kanellopoulos, F. Nisa Bostanci, Hasan Hassan, Mehrshad Lotfi, Phillip B. Gibbons, and Onur Mutlu, "MetaSys: A Practical Open-source Metadata Management System to Implement and Evaluate Cross-layer Optimizations," ***ACM Transactions on Architecture and Code Optimization (TACO)***, June 2022.
368. Juan Gomez-Luna, Izzat El Hajj, Ivan Fernandez, Christina Giannoula, Geraldo F. Oliveira, and Onur Mutlu, "Benchmarking a New Paradigm: Experimental Analysis and Characterization of a Real Processing-in-Memory System," ***IEEE Access***, 10 May 2022.
369. Jorge Gonzalez, Mauricio G. Palma, Maarten Hattink, Ruth Rubio-Noriega, Lois Orosa, Onur Mutlu, Keren Bergman, and Rodolfo Azevedo, "Optically connected memory for disaggregated data centers," ***Journal of Parallel and Distributed Computing (JPDC)***, May 2022.
370. Feng Zhang, Jidong Zhai, Xipeng Shen, Onur Mutlu, and Xiaoyong Du, "POCLib: A High-Performance Framework for Enabling Near Orthogonal Processing on Compression," ***IEEE Transactions on Parallel and Distributed Systems (TPDS)***, February 2022.
371. Lei Chen, Jiacheng Zhao, Chenxi Wang, Ting Cao, John N. Zigman, Haris Volos, Onur Mutlu, Fang Lv, Xiaobing Feng, Guoqing Harry Xu, Huimin Cui, "Unified Holistic Memory Management Supporting Multiple Big Data Processing Frameworks over Hybrid Memories," ***ACM Transactions on Computer Systems (TOCS)***, October 2021.
372. Saba Ahmadian, Reza Salkhordeh, Onur Mutlu, and Hossein Asadi, "ETICA: Efficient Two-Level I/O Caching Architecture for Virtualized Platforms," ***IEEE Transactions on Parallel and Distributed Systems (TPDS)***, October 2021.

373. Lukas Breitwieser, Ahmad Hesam, Jean de Montigny, Vasileios Vavourakis, Alexandros Iosif, Jack Jennings, Marcus Kaiser, Marco Manca, Alberto Di Meglio, Zaid Al-Ars, Fons Rademakers, Onur Mutlu, and Roman Bauer, "BioDynaMo: A General Platform for Scalable Agent-based Simulation," ***Bioinformatics***, 16 September 2021.
374. Mohammed Alser, Jeremy Rotman, Dhrithi Deshpande, Kodi Taraszka, Huwenbo Shi, Pelin Icer Baykal, Harry Taegyun Yang, Victor Xue, Sergey Knyazev, Benjamin D. Singer, Brunilda Balliu, David Koslicki, Pavel Skums, Alex Zelikovsky, Can Alkan, Onur Mutlu, and Serghei Mangul, "Technology Dictates Algorithms: Recent Developments in Read Alignment," ***Genome Biology***, 26 August 2021.
375. Geraldo F. Oliveira, Juan Gomez-Luna, Lois Orosa, Saugata Ghose, Nandita Vijaykumar, Ivan fernandez, Mohammad Sadrosadati, and Onur Mutlu, "DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks," ***IEEE Access***, 8 September 2021.
376. Gagandeep Singh, Mohammed Alser, Damla Senol Cali, Dionysios Diamantopoulos, Juan Gómez-Luna, Henk Corporaal, and Onur Mutlu, "FPGA-based Near-Memory Acceleration of Modern Data-Intensive Applications," ***IEEE Micro***, Vol. 41, No. 4, pages 39-48, July/August 2021.
377. Mohammad Sadrosadati, Amirhossein Mirhosseini, Seyed Borna Ehsani, Hamid Sarbazi-Azad, Mario Drumond, Babak Falsafi, Rachata Ausavarungnirun, and Onur Mutlu, "Highly Concurrent Latency-Tolerant Register Files for GPUs," ***ACM Transactions on Computer Systems (TOCS)***, January 2021.
378. Mohammed Alser, Taha Shahroodi, Juan-Gomez Luna, Can Alkan, and Onur Mutlu, "SneakySnake: A Fast and Accurate Universal Genome Pre-Alignment Filter for CPUs, GPUs, and FPGAs," ***Bioinformatics***, 26 December 2020.
379. Schahram Dustdar, Onur Mutlu, and Nandita Vijaykumar, "Rethinking Divide and Conquer: Towards Holistic Interfaces of the Computing Stack," ***IEEE Internet Computing***, November/December 2020.
380. Mohammed Alser, Zulal Bingol, Damla Senol Cali, Jeremie Kim, Saugata Ghose, Can Alkan, and Onur Mutlu, "Accelerating Genome Analysis: A Primer on an Ongoing Journey," ***IEEE Micro***, Vol. 40, No. 5, pages 65-75, September/October 2020.
381. Syed M. A. H. Jafri, Hasan Hassan, Ahmed Hemani, and Onur Mutlu, "Refresh Triggered Computation: Improving the Energy Efficiency of Convolutional Neural Network Accelerators," ***ACM Transactions on Architecture and Code Optimization (TACO)***, December 2020.
382. Feng Zhang, Jidong Zhai, Xipeng Shen, Dalin Wang, Zheng Chen, Onur Mutlu, Wenguang Chen, and Xiaoyong Du, "TADOC: Text Analytics Directly on Compression," ***VLDB Journal (VLDB Journal)***, September 2020.
383. Can Firtina, Jeremie S. Kim, Mohammed Alser, Damla Senol Cali, A. Ercument Cicek, Can Alkan, and Onur Mutlu, "Apollo: A Sequencing-Technology-Independent, Scalable, and Accurate Assembly Polishing Algorithm," ***Bioinformatics***, June 2020.
384. Orestis Zachariadis, Andrea Teatini, Nitin Satpute, Juan Gomez-Luna, Onur Mutlu, Ole Jakob Elle, and Joaquin Olivares, "Accelerating B-spline Interpolation on GPUs: Application to Medical Image Registration," ***Computer Methods and Programs in Biomedicine***, September 2020.
385. Seyyed Hossein SeyyedAghaei Rezaei, Mehdi Modarressi, Rachata Ausavarungnirun, Mohammad Sadrosadati, Onur Mutlu, and Masoud Daneshtalab, "NoM: Network-on-Memory for Inter-Bank Data Transfer in Highly-Banked Memories," ***IEEE Computer Architecture Letters (CAL)***, 2020.
386. Muhammad Shafique, Mahum Naseer, Theocharis Theocharides, Christos Kyrkou, Onur Mutlu, Lois Orosa, and Jungwook Choi, "Robust Machine Learning Systems: Challenges, Current Trends, Perspectives, and the Road Ahead," ***IEEE Design & Test***, April 2020.
387. Saugata Ghose, Amirali Boroumand, Jeremie S. Kim, Juan Gomez-Luna, and Onur Mutlu, "Processing-in-Memory: A Workload-Driven Perspective," ***IBM Journal of Research and Development***, Special Issue on Hardware for Artificial Intelligence, November 2019.
388. Onur Mutlu, Jeremie S. Kim, "RowHammer: A Retrospective," ***IEEE Transactions on Computer Aided Design (TCAD)***, Special Issue on Top Picks in Hardware and Embedded Security, July 2019.
389. Onur Mutlu, Saugata Ghose, Juan Gomez-Luna, Rachata Ausavarungnirun, "Processing Data Where It Makes Sense: Enabling In-Memory Computation," ***Microprocessors and Microsystems***, [online] March 2019.
390. Mohammed Alser, Hasan Hassan, Akash Kumar Onur Mutlu, Can Alkan, "Shouji: Fast and Efficient Computation of Banded Sequence Alignment," ***Bioinformatics***, 2019.
391. Reza Salkhordeh, Onur Mutlu, Hossein Asadi, "An Analytical Model for Performance and Lifetime Estimation of Hybrid DRAM-NVM Main Memories," ***IEEE Transactions on Computers***, 2019.
392. Mohammad Sadrosadati, Seyed Borna Ehsani, Hajar Falahati, Rachata Ausavarungnirun, Arash Tavakkol, Mojtaba Abaee, Lois Orosa, Yaohua Wang, Hamid Sarbazi-Azad, and Onur Mutlu, "ITAP: Idle-Time-Aware Power Management for GPU Execution Units," ***ACM Transactions on Architecture and Code Optimization (TACO)***, January 2019.
393. Lois Orosa, Rodolfo Azevedo, and Onur Mutlu, "AVPP: Address-first Value-next Predictor with Value Prefetching for Improving the Efficiency of Load Value Prediction," ***ACM Transactions on Architecture and Code Optimization (TACO)***, December 2018.
394. Rachata Ausavarungnirun, Joshua Landgraf, Vance Miller, Saugata Ghose, Jayneel Gandhi, Christopher J. Rossbach, and Onur Mutlu, "Mosaic: Enabling Application-Transparent Support for Multiple Page Sizes in Throughput Processors" ***ACM SIGOPS Operating Systems Review (OSR)***, July 2018.
395. 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.

396. 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***, 2018. Also appears in *Proceedings of the 16th Asia Pacific Bioinformatics Conference (**APBC**)*, Yokohama, Japan, January 2018.
397. Yu Cai, Saugata Ghose, Yixin Luo, Erich F. Haratsch, Onur Mutlu, "Flash Memory SSD Errors, Mitigation, and Recovery," ***Proceedings of the IEEE***, September 2017.
398. 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.
399. 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.
400. 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.
401. 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.
402. 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.
403. 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.
404. 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.
405. Hyoseung Kim, Dionisio de Niz, Bjorn Andersson, Mark Klein, Onur Mutlu, Ragunathan Rajkumar, "Bounding and Reducing Memory Interference in COTS-based Multi-Core Systems," ***Real-Time Systems (RTS)***, accepted, January 2016.
406. 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.
407. 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.
408. 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.
409. 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.
410. 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.
411. 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.
412. 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.
413. 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.
414. Yoongu Kim, Weikun Yang, Onur Mutlu, "Ramulator: A Fast and Extensible DRAM Simulator," ***IEEE Computer Architecture Letters (CAL)***, March 2015.
415. 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.**
416. 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.
417. 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.
418. 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.
419. Onur Mutlu, Lavanya Subramanian, "Research Problems and Opportunities in Memory Systems," ***Supercomputing Frontiers and Innovations***, December 2014. **Invited Article.**

420. 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.**

421. 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.**

422. 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.

423. 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.

424. 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.

425. 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.

426. 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.

427. 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.

428. 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.

429. 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.

430. 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.

431. 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.

432. 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.

433. 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.

434. 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.**

435. 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.

436. 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.

437. José A. Joao, Onur Mutlu, Hyesoon Kim, Yale N. Patt, "Dynamic Predication of Indirect Jumps," ***IEEE Computer Architecture Letters (CAL)***, vol. 6, May 2007.

438. 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.

439. 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.**

440. 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.

441. 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.

442. 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.

443. 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.

444. 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.
445. 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

446. Hai Li, Alaa R. Alameldeen, Onur Mutlu, "Near-Memory and In-Memory Processing," ***IEEE Design and Test***, vol. 39(2), April 2022.
447. Theocharis Theocharides, Muhammad Shafique, Jungwook Choi, Onur Mutlu, "Robust Resource-Constrained Systems for Machine Learning," ***IEEE Design and Test***, vol. 37(2), April 2020.
448. Onur Mutlu, Saugata Ghose, Rachata Ausavarungnirun, "Recent Advances in DRAM and Flash Memory Architectures," ***IPSI Transactions on Internet Research, Special Issue, Guest Editors' Introduction***, vol. 14(2), July 2018.
449. Onur Mutlu, Saugata Ghose, Rachata Ausavarungnirun, "Guest Editors' Introduction: Recent Advances in Overcoming Bottlenecks in Memory Systems and Managing Memory Resources in GPU Systems," ***IPSI Transactions on Advanced Research, Special Issue, Guest Editors' Introduction***, vol. 14(2), July 2018.
450. Onur Mutlu, Scott A. Mahlke, Thomas M. Conte, Wen-mei W. Hwu, "Iterative Module Scheduling," ***IEEE Micro***, January/February 2018.
451. Onur Mutlu, Rich Belgard, "Common Bonds: MIPS, HPS, Two-Level Branch Prediction, and Compressed Code RISC Processor," ***IEEE Micro***, July/August 2016.
452. Onur Mutlu, Rich Belgard, "The 2014 MICRO Test of Time Award Winners: From 1978 to 1992," ***IEEE Micro***, January/February 2016.
453. 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.
454. Yale N. Patt, Onur Mutlu, "Guest Editors' Introduction: Top Picks," ***IEEE Micro, Special Issue (IEEE MICRO)***, vol. 31(1), January/February 2011.
455. 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)

456. Can Firtina, Maximilian Mordig, Harun Mustafa, Sayan Goswami, Nika Mansouri Ghiasi, Stefano Mercogliano, Furkan Eris, Joël Lindegger, Andre Kahles, and Onur Mutlu, "Rawsamble: Overlapping and Assembling Raw Nanopore Signals using a Hash-based Seeding Mechanism," *32nd Annual Conference on Intelligent Systems for Molecular Biology (**ISMB**)*, Montreal, QC, Canada, July 2024.
457. Meryem Banu Cavlak, Gagandeep Singh, Mohammed Alser, Can Firtina, Joël Lindegger, Mohammad Sadrosadati, Nika Mansouri Ghiasi, Can Alkan, and Onur Mutlu, "TargetCall: Eliminating the Wasted Computation in Basecalling via Pre-Basecalling Filtering," *The 13th RECOMB Satellite Conference on Biological Sequence Analysis (**RECOMB-SEQ**)*, Istanbul, Turkey, April 2023.
458. Joël Lindegger, Damla Senol Cali, Mohammed Alser, Juan Gómez-Luna, Nika Mansouri Ghiasi, and Onur Mutlu, "Scrooge: A Fast and Memory-Frugal Genomic Sequence Aligner for CPUs, GPUs, and ASICs," *The 13th RECOMB Satellite Conference on Biological Sequence Analysis (**RECOMB-SEQ**)*, Istanbul, Turkey, April 2023.
459. Gagandeep Singh, Mohammed Alser, Alireza Khodamoradi, Kristof Denolf, Can Firtina, Meryem Banu Cavlak, Henk Corporaal, and Onur Mutlu, "RUBICON: A Framework for Designing Efficient Deep Learning-Based Genomic Basecallers," *Pacific Symposium on Biocomputing (**PSB**) Poster Session*, Hawaii, January 2023.
460. Geraldo F. Oliveira, Juan Gómez-Luna, Saugata Ghose, and Onur Mutlu, "Methodologies, Workloads, and Tools for Processing-in-Memory: Enabling the Adoption of Data-Centric Architectures," *Proceedings of the 20th IEEE Computer Society Annual Symposium on VLSI **(ISVLSI)***, Paphos, Cyprus, 4 July 2022. **Invited Presentation and Paper.**
461. Ataberk Olgun, Juan Gómez-Luna, Konstantinos Kanellopoulos, Behzad Salami, Hasan Hassan, Oguz Ergin, and Onur Mutlu, "PiDRAM: An FPGA-based Framework for End-to-end Evaluation of Processing-in-DRAM Techniques," *Proceedings of the 20th IEEE Computer Society Annual Symposium on VLSI **(ISVLSI)***, Paphos, Cyprus, 4 July 2022. **Invited Presentation and Paper.**
462. Geraldo F. Oliveira, Amirali Boroumand, Saugata Ghose, Juan Gómez-Luna, and Onur Mutlu, "Heterogeneous Data-Centric Architectures for Modern Data-Intensive Applications: Case Studies in Machine Learning and Databases," *Proceedings of the 20th IEEE Computer Society Annual Symposium on VLSI **(ISVLSI)***, Paphos, Cyprus, 4 July 2022. **Invited Presentation and Paper.**
463. Ivan Fernandez, Ricardo Quislant, Christina Giannoula, Mohammed Alser, Juan Gómez-Luna, Eladio Gutiérrez, Oscar G. Plata, and Onur Mutlu, "Exploiting Near-Data Processing to Accelerate Time Series Analysis," *Proceedings of the 20th IEEE Computer Society Annual Symposium on VLSI **(ISVLSI)***, Paphos, Cyprus, 4 July 2022. **Invited Presentation and Paper.**
464. Nika Mansouri-Ghiasi, Jisung Park, Harun Mustafa, Jeremie S. Kim, Ataberk Olgun, Arvid Gollwitzer, Damla Senol Cali, Can Firtina, Haiyu Mao, Nour Almadhoun Alserr, Rachata Ausavarungnirun, Nandita Vijaykumar, Mohammed Alser, and Onur Mutlu, "GenStore: In-Storage Filtering of Genomic Data for High-Performance and Energy-Efficient Genome Analysis," *Proceedings of the 20th IEEE Computer Society Annual Symposium on VLSI **(ISVLSI)***, Paphos, Cyprus, 4 July 2022. **Invited Presentation and Paper.**
465. Christina Giannoula, Ivan Fernandez, Juan Gómez-Luna, Nectarios Koziris, Georgios I. Goumas, and Onur Mutlu, "SparseP: Efficient Sparse Matrix Vector Multiplication on Real Processing-In-Memory Architectures," *Proceedings of the 20th IEEE Computer Society Annual Symposium on VLSI **(ISVLSI)***, Paphos, Cyprus, 4 July 2022. **Invited Presentation and Paper.**

466. Juan Gómez-Luna, Yuxin Guo, Sylvan Brocard, Julien Legriel, Remy Cimadomo, Geraldo F. Oliveira, Gagandeep Singh, and Onur Mutlu, "Machine Learning Training on a Real Processing-in-Memory System," *Proceedings of the 20th IEEE Computer Society Annual Symposium on VLSI **(ISVLSI)***, Paphos, Cyprus, 4 July 2022. **Invited Presentation and Paper.**
467. Joel Lindegger, Damla Senol Cali, Mohammed Alser, Juan Gomez-Luna, and Onur Mutlu, "Algorithmic Improvement and GPU Acceleration of the GenASM Algorithm," *Proceedings of the 21st IEEE International Workshop on High Performance Computational Biology **(HiCOMB)***, Virtual, May 2022.
468. Safaa Diab, Amir Nassereldine, Mohammed Alser, Juan Gomez-Luna, Onur Mutlu, and Izzat El Hajj, "High-throughput Pairwise Alignment with the Wavefront Algorithm using Processing-in-Memory," *Proceedings of the 21st IEEE International Workshop on High Performance Computational Biology **(HiCOMB)***, Virtual, May 2022.
469. Onur Mutlu, "RowHammer and Beyond," *Invited Keynote Paper in Proceedings of the International Workshop on Constructive Side-Channel Analysis and Secure Design **(COSADE)***, Darmstadt, Germany, April 2019.
470. Onur Mutlu, "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Invited Keynote Paper in Proceedings of the 7th Mediterranean Conference on Embedded Computing **(MECO)***, Budva, Montenegro, June 2018.
471. 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.
472. 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.
473. 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.
474. Onur Mutlu, "ThyNVM: Software-Transparent Crash Consistency for Persistent Memory," *Proceedings of the Flash Memory Summit **(FMS)***, Santa Clara, CA, August 2016.
475. Onur Mutlu, "Large-Scale Study of In-the-Field Flash Failures," *Proceedings of the Flash Memory Summit **(FMS)***, Santa Clara, CA, August 2016.
476. 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.
477. 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.
478. Onur Mutlu, "Read Disturb Errors in MLC NAND Flash Memory," *Proceedings of the Flash Memory Summit **(FMS)***, Santa Clara, CA, August 2015.
479. 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.
480. 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.
481. 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.**
482. Onur Mutlu, "Error Analysis and Management for MLC NAND Flash Memory," *Proceedings of the Flash Memory Summit **(FMS)***, Santa Clara, CA, August 2014.
483. 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.
484. 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.
485. Onur Mutlu, "Memory Scaling: A Systems Architecture Perspective," *Proceedings of MemCon **(MEMCON)***, Santa Clara, CA, August 2013.
486. Chuanjun Zhang, Glenn G. Ko, Jungwook Choi, Shang-nien Tsai, Minje Kim, Abner Guzman-Rivera, Rob A. Rutenbar, Paris Smaragdis, 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.

487. 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.
488. 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.
489. 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.
490. 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.**
491. 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.
492. 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.
493. 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.
494. 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)

495. Onur Mutlu, "Retrospective: An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms," arXiV.org, arXiv:2306.16037, June 2023. **ISCA@50 25-Year Retrospective: 1996-2020.**
496. Onur Mutlu, "Retrospective: Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors," arXiV.org, arXiv:2306.16093, June 2023. **ISCA@50 25-Year Retrospective: 1996-2020.**
497. Onur Mutlu, "Retrospective: RAIDR: Retention-Aware Intelligent DRAM Refresh," arXiV.org, arXiv:2306.16024, June 2023. **ISCA@50 25-Year Retrospective: 1996-2020.**
498. Junwhan Ahn, Sungpack Hong, Sungjoo Yoo, Onur Mutlu, and Kiyoung Choi, "Retrospective: A Scalable Processing-in-Memory Accelerator for Parallel Graph Processing," arXiV.org, arXiv:2306.16024, June 2023. **ISCA@50 25-Year Retrospective: 1996-2020.**
499. Engin Ipek, Onur Mutlu, José F. Martínez, and Rich Caruana "Retrospective: Self-Optimizing Memory Controllers: A Reinforcement Learning Approach," **ISCA@50 25-Year Retrospective: 1996-2020.**
500. İsmail Emir Yüksel, Ataberk Olgun, Behzad Salami, F. Nisa Bostancı, Yahya Can Tuğrul, A. Giray Yağlıkçı, Nika Mansouri Ghiasi, Onur Mutlu, and Oğuz Ergin, "TuRaN: True Random Number Generation Using Supply Voltage Underscaling in SRAMs," arXiv.org, arXiv:2211.10894, November 2022.
501. Mohammed Alser, Julien Eudine, and Onur Mutlu "Genome-on-Diet: Taming Large-Scale Genomic Analyses via Sparsified Genomics," arXiv.org, arXiv:2211.08157, November 2022.
502. Ivan Fernandez, Aditya Manglik, Christina Giannoula, Ricardo Quislant, Nika Mansouri Ghiasi, Juan Gómez-Luna, Eladio Gutierrez, Oscar Plata, and Onur Mutlu, "Accelerating Time Series Analysis via Processing using Non-Volatile Memories," arXiv.org, arXiv:2211.04369, November 2022.
503. Gagandeep Singh, Mohammed Alser, Alireza Khodamoradi, Kristof Denolf, Can Firtina, Meryem Banu Cavlak, Henk Corporaal, and Onur Mutlu, "A Framework for Designing Efficient Deep Learning-Based Genomic Basecallers," arXiv.org, arXiv:2211.03079, November 2022.
504. Nika Mansouri Ghiasi, Mohammad Sadrosadati, Geraldo F. Oliveira, Konstantinos Kanellopoulos, Rachata Ausavarungnirun, Juan Gómez Luna, Aditya Manglik, João Ferreira, Jeremie S. Kim, Christina Giannoula, Nandita Vijaykumar, Jisung Park, and Onur Mutlu, "RevaMp3D: Architecting the Processor Core and Cache Hierarchy for Systems with Monolithically-Integrated Logic and Memory," arXiv.org, arXiv:2210.08508, October 2022.
505. Juan Gómez-Luna, Yuxin Guo, Sylvan Brocard, Julien Legriel, Remy Cimadomo, Geraldo F. Oliveira, Gagandeep Singh, and Onur Mutlu, "An Experimental Evaluation of Machine Learning Training on a Real Processing-in-Memory System," arXiv.org, arXiv:2207.07886, July 2022.
506. Hasan Hassan, Ataberk Olgun, A. Giray Yaglikci, Haocong Luo, and Onur Mutlu, "A Case for Self-Managing DRAM Chips: Improving Performance, Efficiency, Reliability, and Security via Autonomous in-DRAM Maintenance Operations," arXiv.org, arXiv:2207.13358, July 2022.
507. Minesh Patel, Taha Shahroodi, Aditya Manglik, A. Giray Yaglikci, Ataberk Olgun, Haocong Luo, and Onur Mutlu, "A Case for Transparent Reliability in DRAM Systems," arXiv.org, arXiv:2204.10378, April 2022.
508. Geraldo F. Oliveira, Saugata Ghose, Juan Gómez-Luna, Amirali Boroumand, Alexis Savery, Sonny Rao, Salman Qazi, Gwendal Grignou, Rahul Thakur, Eric Shiu, and Onur Mutlu, "Extending Memory Capacity in Consumer Devices with Emerging Non-Volatile Memory: An Experimental Study," arXiv.org, arXiv:2111.02325, November 2021.

509. A. Giray Yaglıkcı, Jeremie S. Kim, Fabrice Devaux, and Onur Mutlu, "Security Analysis of the Silver Bullet Technique for RowHammer Prevention," arXiv.org, arXiv:2106.07084, June 2021.
510. Zülal Bingöl, Mohammed Alser, Onur Mutlu, Ozcan Ozturk, and Can Alkan, "GateKeeper-GPU: Fast and Accurate Pre-Alignment Filtering in Short Read Mapping," arXiv.org, arXiv:2103.14978, March 2021.
511. Lois Orosa, Yaohua Wang, Ivan Puddu, Mohammad Sadrosadati, Kaveh Razavi, Juan Gomez-Luna, Hasan Hassan, Nika Mansouri-Ghiasi, Arash Tavakkol, Minesh Patel, Jeremie Kim, Vivek Seshadri, Uksong Kang, Saugata Ghose, Rodolfo Azevedo, Onur Mutlu, "Dataplant: In-DRAM Security Mechanisms for Low-Cost Devices," arXiv.org, arXiv:1902.07344, February 2019.
512. Arash Tavakkol, Aasheesh Kolli, Stanko Novakovic, Kaveh Razavi, Juan Gomez-Luna, Hasan Hassan, Claude Barthels, Yaohua Wang, Mohammad Sadrosadati, Saugata Ghose, Ankit Singla, Pratap Subrahmanyam, Onur Mutlu, "Enabling Efficient RDMA-based Synchronous Mirroring of Persistent Memory Transactions," arXiv.org, arXiv:1810.09360, October 2018.
513. 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.
514. 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.
515. 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.
516. 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.
517. 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.
518. 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.
519. 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.
520. 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.
521. 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.
522. 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.
523. 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. Jeremie S. Kim, Minesh H. Patel, Stephan G. Meier, Tyler J. Huberty, Onur Mutlu, "Security Techniques Based on Memory Timing Characteristics," US Patent 10,776,521, September 2020.
2. Vijay Chidambaram,, Irina Calciu, Himanshu Chauhan, Eric Schkufza, Onur Mutlu, Pratap Subrahmanyam, "Conversion Tool for Moving from Block-based Persistence to Byte-based Persistence," US Patent 10,515,029, December 2019.
3. William H. Mangione Smith, Onur Mutlu, "Reliable Communications in On-chip Networks," US Patent 8473818, June 2013.
4. Thomas Moscibroda, Onur Mutlu, "Bufferless Routing in On-Chip Interconnection Networks," US Patent 8509078, August 2013.
5. Thomas Moscibroda, Onur Mutlu, "Prioritization of Multiple Concurrent Threads for Scheduling Requests to Shared Memory," US Patent 8271741, September 2012.
6. Thomas Moscibroda, Onur Mutlu, "Coordination among Multiple Memory Controllers," US Patent 8266393, September 2012.
7. Thomas Moscibroda, Onur Mutlu, "Controlling Interference in Shared DRAM Systems using Batch Scheduling," US Patent 8180975, May 2012.
8. Onur Mutlu, José A. Joao, "Feedback Mechanism for Dynamic Predication of Indirect Jumps," US Patent 7818551, October 2010.
9. José A. Joao, Onur Mutlu, "Target-Frequency based Indirect Jump Prediction for High-Performance Processors," US Patent 7870371, January 2011.
10. Onur Mutlu, Thomas Moscibroda, "A Software-Configurable and Stall-Time Fair Memory Access Scheduling Mechanism for Shared Memory Systems," US Patent 8245232, August 2012.
11. Thomas Moscibroda, Onur Mutlu, "Multi-level DRAM Controller to Manage Access to DRAM," US Patent 8001338, August 2011.
12. Onur Mutlu, Thomas Moscibroda, "Parallelism-Aware Memory Request Scheduling in Shared Memory Controllers," US Patent Application filed August 2007.
13. Thomas Moscibroda, Onur Mutlu, "Fairness in Memory Systems," US Patent Application filed July 2007.

14. Jared Stark, Chris Wilkerson, Onur Mutlu,"Apparatus for Memory Communication During Runahead Execution," US Patent Application filed December 2002.
15. Eric Sprangle, Onur Mutlu, "Method and Apparatus to Control Memory Accesses," US Patent 6799257, September 2004.

Invited Talks, Lectures, Tutorials

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

1. "Real Processing in Memory Systems (RealPIM) Tutorial"
   - Co-organizer (with Juan Gomez-Luna) of full-day tutorial held with the HPCA Conference, Montreal, QC, Canada, 26 February 2023.
   - Co-organizer (with Juan Gomez-Luna) of full-day tutorial held with the ASPLOS Conference, Montreal, QC, Canada, 26 March 2023.
   - Co-organizer (with Juan Gomez-Luna) of full-day tutorial held with the ISCA Conference, Orlando, FL, USA, 18 June 2023.
   - Co-organizer (with Juan Gomez-Luna) of full-day tutorial held with the MICRO Conference, Toronto, ON, Canada, 29 October 2023.
2. "Memory-Centric Computing Systems Tutorial and Tutorial+Workshop"
   - Co-organizer (with Geraldo Oliveira et al.) of half-day tutorial held with the ISCA Conference, Buenos Aires, Argentina, 29 June 2024.
   - Co-organizer (with Geraldo Oliveira et al.) of half-day tutorial held with the MICRO Conference, Austin, TX, USA, 2 November 2024.
   - Co-organizer (with Geraldo Oliveira et al.) of half-day tutorial held with the PPoPP/HPCA/CGO Conference, Las Vegas, NV, USA, 1 March 2025.
   - Co-organizer (with Geraldo Oliveira et al.) of full-day workshop+tutorial held with the ASPLOS Conference, Rotterdam, Netherlands, 30 March 2025.
3. "Future of Computer Architecture and Hardware Security,"
   - **Keynote talk** at Qualcomm Product Security Summit, San Diego, CA, USA, 16 May 2024.
   - Invited talk at Microsoft Research, Redmond, WA, USA, 10 September 2024.
   - **Keynote talk** at hardwear.io Memory Security Track, Amsterdam, Netherlands, 25 October 2024.
   - Invited talk at the Hong Kong University of Science and Technology (HKUST), Hong Kong, 31 October 2024.
   - Invited talk at the University of Southern California, Los Angeles, CA, USA, 6 March 2025.
4. "ML-Assisted Memory and Storage Management,"
   - Invited talk at the AI4FACD (AI for Fully Automated Chip Design) Workshop, held with ISCA, Virtual, 29 June 2024.
5. "Storage-Centric Computing: Enabling Fundamentally Efficient Computers"
   - **Keynote talk** at the CCF Storage Conference, Guangzhou, China, 1 December 2024.
6. "Memory-Centric Computing: Recent Advances in Processing-in-DRAM,"
   - Invited talk at the 70th International Electron Devices Meeting (IEDM), San Francisco, CA, USA, 9 December 2024.
   - Invited Talk at Dagstuhl Seminar on "Disruptive Memory Technologies", Germany, 7 April 2025.
7. "Reinventing Computing and Storage" or "Storage-Centric Computing for Modern Data-Intensive Workloads"
   - **Keynote talk** at the Huawei OlympusMons Award Ceremony, Emei, Sichuan, China, 7 April 2023.
   - **Keynote talk** at the Next Generation Data Storage Architecture Innovation Forum, Luzern, Switzerland, 15 May 2023.
   - **Keynote talk** at the NCIS Conference, Guangzhou, China, 16 September 2023.
8. "Some Reflections (on DRAM)"
   - **Award speech** for ACM SIGARCH Maurice Wilkes Award Acceptance at ISCA, Phoenix, AZ, USA, 25 June 2019.
9. "Runahead Execution: A Short Retrospective,"
   - IEEE International Conference on High Performance Computer Architecture (HPCA) Test of Time Award Ceremony, Virtual, 2 March 2021.
10. "Transforming Server Architectures"
    - **Keynote talk** at the Swiss Data Center Summit (DC SUMMIT), Zurich, Switzerland, 30 June 2022.
11. "Accelerating Genome Analysis" or "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, Zurich, Switzerland, 8 March 2018.
    - Bogazici University, Istanbul, Turkey, 10 April 2018.
    - **Keynote talk** at the 17th International Workshop on High Performance Computational Biology (HiCOMB), held with IPDPS, Vancouver, BC, Canada, 21 May 2018.
    - **Keynote talk** at the 2nd Workshop on Accelerator Architecture in Computational Biology and Bioinformatics (AACBB), Washington, DC, USA, 16 February 2019.
    - Shenzhen Institute of Advanced Technology, Shenzhen, China, 17 May 2019.
    - Invited Lecture at Technion, Virtual, 26 January 2021.
    - **Keynote Talk** at the Systems for Post-Moore Architectures Workshop (SPMA), held with the EuroSys Conference (EUROSYS), Virtual, 5 April 2022.

- **Keynote Talk** at 4th Workshop on Accelerator Architecture in Computational Biology and Bioinformatics (AACBB), held with the International Symposium on Computer Architecture (ISCA), New York, NY, USA, 18 June 2022.
- Invited Talk at the Barcelona Supercomputing Center (BSC), Barcelona, Spain, 6 September 2022.
- Invited Talk at the Montenegro Academy of Sciences Conference, Montenegro, 2 November 2022.
- **Opening talk** at the 1st Workshop on Hardware Acceleration of Bioinformatics Workloads (BIO-Arch), held with Research on Computational Molecular Biology (RECOMB), Istanbul, Turkey, 14 April 2023.
- Invited Talk at the University Donja Gorica, Podgorica, Montenegro, 8 June 2023.
- Invited Talk at the University of Montenegro Faculty of Medicine, Podgorica, Montenegro, 9 June 2023.
- Invited Special Session Talk at the 60th Design Automation Conference, San Francisco, CA, 11 July 2023.

12. "RowHammer" or "The Story of RowHammer" or "RowHammer and Beyond" or 'The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser" or "The DRAM RowHammer Problem (and Its Reliability and Security Implications)" or "Security Aspects of DRAM: The Story of RowHammer" or "Securing the Memory System: The Story of RowHammer" or "Fundamentally Understanding and Solving RowHammer" or "RowHammer, RowPress & Beyond: Can We Be Free of Bitflips (Soon)?"
    - 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.
    - Vrjie University, Amsterdam, The Netherlands, 10 November 2017.
    - TU Graz Security Seminar, Graz, Austria, 28 May 2018.
    - Microsoft Research Faculty Summit, Redmond, WA, USA, 2 August 2018.
    - **Special Invited Talk** at the First Workshop on Top Picks in Hardware and Embedded Security (TopinHES), held with ICCAD, San Diego, CA, USA, 8 November 2018.
    - **Keynote talk** at the 10th International Workshop on Constructive Side-Channel Analysis and Secure Design (COSADE), Darmstadt, Germany, 4 April 2019.
    - New York University ECE Colloquium, Brooklyn, NY, USA, 25 April 2019.
    - **Keynote talk** at the 32nd IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT), Delft, Netherlands, 2 October 2019.
    - **Keynote talk** at the HiPEAC Computing Systems Week Autumn 2019 Resilience in High Performance Computing Thematic Session (RESILIENTHPC), Bilbao, Spain, 28 October 2019.
    - Invited talk at 28th IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC), Special Session on AI Frameworks, Virtual, 9 October 2020.
    - Invited talk at the Workshop on Robust and Safe Software 2.0 (RSS2), held with the 27th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Virtual, 28 February 2022.
    - **Keynote talk** at Secure Hardware, Architectures, and Operating Systems Workshop (SeHAS), held with HiPEAC 2021 Conference, Virtual, 19 January 2021.
    - Invited lecture at the New York University Hardware Security Class, Virtual, 29 April 2022.
    - Invited tutorial at 14th IEEE Electron Devices Society International Memory Workshop (IMW), Dresden, Germany, 15 May 2022.
    - Invited Talk at Qualcomm, San Diego, CA, USA, 11 May 2022.
    - Invited talk at IBM Research, Yorktown Heights, NY, USA, 16 May 2022.
    - Invited talk at Google, San Diego, CA, USA, 29 August 2022.
    - Invited talk at the University of British Columbia, Vancouver, BC, Canada, 20 September 2022.
    - Invited talk at the Huawei Strategy and Technology Workshop, Virtual, 27 September 2022.
    - Invited talk at the University of Chicago, Chicago, IL, US, 6 October 2022.
    - Invited talk at Northwestern University, Evanston, IL, US, 7 October 2022.
    - **Keynote talk** at the 6th NTUA Computing Systems Research Day, Athens, Greece, 10 January 2023.
    - Invited Talk at the University of Montenegro, Podgorica, Montenegro, 9 June 2023.
    - Invited Talk at New York University, 23 June 2023.
    - Invited Talk at the Flash Memory Summit, 10 August 2023.
    - Invited Talk at Meta, 28 September 2023.
    - Invited Talk at the Monte Verita Workshop on New Software Abstractions for Hardware Security Technology, Monte Verita, Ascona, Switzerland, 2 October 2023.
    - Invited Talk at Stanford University, 26 October 2023.
    - **Keynote talk** at the Google Zurich Hardware Security Summit, 15 November 2023.
    - Invited Talk at Dagstuhl Seminar on "Microarchitectural Attacks and Defenses", Germany, 27 November 2023.
    - Invited Talk at Apple, 28 November 2023.
    - Invited Talk at New Jersey Institute of Technology, 18 December 2023.
    - Invited Talk at North Carolina State University, 16 February 2024.
    - Invited Talk at University of California at Berkeley, Berkeley, CA, USA, 28 February 2024.
    - **Keynote talk** at Qualcomm Innovation Fellowship Europe Finals, Virtual, 31 May 2024.
    - Invited Talk at University of California at Boulder, Boulder, CO, USA, 13 June 2024.
    - **Award lecture** for the Jean-Claude Laprie Award in Dependable Computing at the DSN Conference, Brisbane, Australia, 26 June 2024.

13. "Intelligent Architectures for Intelligent Machines" or "Intelligent Architectures for Intelligent Computing Systems"
    - **Keynote talk** at SRC-Mubadala-Khalifa Forum on the Future of Artificial Intelligence Hardware Systems, Abu Dhabi, UAE, 15 April 2019.
    - Huawei Strategy and Technology Workshop (STW), Shenzhen, China, 14 May 2019.

- Invited talk at Yale: 80 in 2019: Pushing the Envelope of Computing for the Future (YALE@80), Barcelona, Spain, 1 July 2019.
- **Keynote talk** at the 19th International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS), Samos, Greece, 10 July 2019.
- Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, 17 August 2019.
- Tsinghua University, EE Department, Beijing, China, 18 August 2019.
- **Keynote talk** at the HiPEAC Computing Systems Week Autumn 2019 (HIPEAC CSWEEK), Bilbao, Spain, 28 October 2019.
- **Keynote talk** at the Khalifa University System on Chip Center Open House, Abu Dhabi, UAE, 20 November 2019.
- Invited talk at the Huawei European Research Symposium, Paris, France, 26 November 2019.
- **Keynote talk** at the 17th ChinaSys Workshop, Zhuhai, China, 20 December 2019.
- **Plenary keynote talk** at the 2020 International Symposia on VLSI (VLSI), Virtual, 11 August 2020.
- **Distinguished lecture** at HKUST Engineering and HKSTP Distinguished Speaker Series, Virtual, 7 October 2020.
- **Keynote talk** at 13th ACM International Systems and Storage Conference (SYSTOR) , Virtual, 13 October 2020.
- **Keynote talk** at National Science Foundation Workshop on Processing-In-Memory Technology (NSF-PIM) , Virtual, 26 October 2020.
- Invited talk at Texas State University Computer Science Seminar, Virtual, 13 November 2020.
- **Keynote talk** at Huawei Compute and Storage Technology Workshop, Virtual, 2 December 2020.
- Invited talk at SUSTech, Virtual, 11 December 2020.
- Seoul National University Neural Processing Research Center, Virtual, 12 January 2021.
- Invited Special Session Talk at Design, Automation and Test in Europe (DATE), Virtual, 2 February 2021.
- Invited Tech Talk at AMD, Virtual, 19 February 2021.
- **Keynote talk** at Future of Information and Communication Conference (FICC), Virtual, 29 April 2021.
- Opening talk at TU Vienna Mondays in Memory Webinar Series (MiM), Virtual, 3 May 2021.
- Closing talk at IEEE Data and Storage Symposium (DSS), Virtual, 30 June 2021.
- Invited talk at Supercomputing Frontiers Europe (SCFE), Virtual, 21 July 2021.
- **Keynote talk** at TUBA World Conference on Energy Science and Technology (WCEST), Virtual, 11 August 2021.
- Public Lecture at Frontiers of AI Accelerators: Technologies, Circuits and Applications, Croucher Advanced Study Institute Workshop (FAI-ASI), Virtual, 25 August 2021.
- Invited talk at Huawei Strategy and Technology Workshop, Virtual, 14 October 2021.
- **Distinguished lecture** at Ontario Tech Engineering Research Distinguished Speaker Series, Virtual, 19 October 2021.
- **Keynote talk** at 15th International Conference on Networking, Architecture, and Storage (NAS), Riverside, CA, USA, 25 October 2021.
- Invited webinar at IEEE SSCS and CAS Central Texas Chapter, Virtual, 2 December 2021.
- Invited talk at the Faculty Development Program of Meerut Institute of Engineering and Technology on Advanced Computing Techniques, Virtual, 6 March 2022.
- Invited talk at INSAIT Conference on Emerging Trends in AI and Computing Research, Sofia, Bulgaria, 1 October 2022.
- Invited talk at VMware, Palo Alto, CA, 13 September 2023.

14. "Future Computing Architectures" or "Future Computing and Genome Analysis Platforms" or "Future Computing Architectures: Challenges and Opportunities" or "Future Computing Platforms: 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.
    - Bogazici University, Istanbul, Turkey, 10 April 2018.
    - Apple, Cupertino, CA, USA, 17 April 2018.
    - AliBaba, Bellevue, WA, USA, 20 April 2018.
    - IBM TJ Watson Research Center PIC Talk, White Plains, NY, USA, 23 April 2018.
    - New York University (Tandon Engineering), Brooklyn, NY, USA, 24 April 2018.
    - Jacobs Foundation Visit to ETH, Zurich, Switzerland, 3 May 2018.
    - Microsoft Azure Team, Redmond, WA, USA, 3 August 2018.
    - University of Texas at Austin ECE Department Colloquium, Austin, TX, 30 January 2019.
    - Invited lecture at IEEE Computer Society Turkey Chapter, Virtual, 20 February 2021.
    - Invited lecture at Koc University ACM Student Chapter, Virtual, 22 March 2021.
    - **Keynote talk** at CSCON IEEE Turkey Computer Society Student Branch Conference (CSCON), Virtual, 2 April 2021.
    - **Distinguished seminar** at Koc University College of Engineering, Virtual, 9 April 2021.
    - Invited seminar at IEEE Egypt and IEEE Future University in Egypt Student Branch, Virtual, 16 May 2021.
    - Invited talk at Huawei Strategy and Technology Workshop, Virtual, 15 October 2021.
    - Invited lecture at the 10th EdukCircle International Convention on Engineering and Computer Technology, Virtual, 10 November 2021.
    - Closing talk at IEEE Data and Storage Symposium (DSS), Virtual, 1 December 2021.
    - Stanford University SystemX Seminar, Stanford, CA, USA, 8 February 2024.

15. "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" or "Rethinking Memory System Design: Robustness, Energy, Performance" (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 (FABULOUS), 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.
- **Keynote talk** at the 28th International Online Testing Symposium (IOLTS), Costa Brava, Spain, 3 July 2018.
- **Joint Keynote Talk**, at the 9th International Symposium on Non-Volatile Memory Systems and Applications (NVMSA) and the 24th International Symposium on Real-Time Computing Systems and Applications (RTCSA), Hakodate, Japan, 29 August 2018.
- Huawei, Santa Clara, CA, USA, 14 January 2019.

16. "Processing Data Where It Makes Sense (in Modern Computing Systems): Enabling In-Memory Computation" or "Enabling Practical, Efficient and Large-Scale Computation Near Data" or "Changing the Computing Paradigm for High Efficiency" or "Memory-Centric Computing in the Big Data Era"
    - 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.
    - Istanbul Technical University, Istanbul, Turkey, 9 April 2018.
    - **Opening lecture** at the DAC Design Automation Summer School, Austin, TX, USA, 24 June 2018.
    - **Keynote talk** at the 7th Mediterranean Conference on Embedded Computing (MECO), Budva, Montenegro, 13 June 2018.
    - **Keynote talk** at the Cordoba HiPerNav Week, Cordoba, Spain, 17 September 2018.
    - **Keynote talk** at the 2nd Workshop on Highly Efficient Neural Processing (HENP), held with ESWEEK, Turin, Italy, 3 October 2018.
    - VMware, Palo Alto, CA, USA, 5 November 2018.
    - Facebook, Menlo Park, CA, USA, 6 November 2018.
    - Qualcomm, San Diego, CA, USA, 9 November 2018.
    - Intel, Santa Clara, CA, USA, 12 November 2018.
    - Barcelona Supercomputing Center and Universitat Politecnica de Catalunya, Barcelona, Spain, 20 November 2018.
    - **Keynote talk** at the 4th Workshop on Mobile System Technologies (MST), Milan, Italy, 23 November 2018.
    - University of Pittsburgh, Pittsburgh, PA, 4 December 2018.
    - Massachussetts Institute of Technology, Cambridge, MA, 10 December 2018.
    - University of Southern California, Los Angeles, CA, 10 January 2019.
    - **Distinguished Lecture** at George Washington University, 15 February 2019.
    - **Special Invited Talk** at ISSCC Forum on "Intelligence at the Edge: How Can We Make Machine Learning More Energy Efficient?", San Francisco, CA, 20 February 2019.
    - **Special Invited Talk** at Dagstuhl Seminar on Emerging Hardware Techniques and EDA Methodologies for Neuromorphic Computing, Dagstuhl, Germany, 8 April 2019.
    - **Keynote talk** at the 29th ACM Great Lakes Symposium on VLSI (GLSVLSI), Washington, DC, USA, 10 May 2019.

- **Special Invited Tutorial** at the International Memory Workshop (IMW), Monterey, CA, USA, 12 May 2019.
- **Opening lecture** at the DAC Design Automation Summer School, Las Vegas, NV, USA, 2 June 2019.
- **Special Invited Talk** at the Design Automation Conference Special Session on In-Memory Computing, Las Vegas, NV, USA, 4 June 2019.
- TU Wien, Vienna, Austria, 18 June 2019.
- Computer Architecture Research Directions Workshop Memory Debate, with ISCA, Phoenix, AZ, USA, 23 June 2019.
- **Special Invited Talk** at the Flash Memory Summit, Santa Clara, CA, USA, 8 August 2019.
- **Keynote talk** at the International Symposium on Advanced Parallel Processing Technology (APPT), Tianjin, China, 16 August 2019.
- Peking University, Beijing, China, 19 August 2019.
- Tsinghua University, CS Department, Beijing, China, 19 August 2019.
- **Keynote talk** at the 37th IEEE International Conference on Computer Design (ICCD), 18 November 2019.

17. "Memory-Centric Computing" or "Memory-Centric Computing Systems" or "Memory-Centric Computing: Enabling Fundamentally Efficient Computers"
    - **Endowed lecture series** at the University of Texas at Austin ECE Department, Austin, TX, 29, 31 January and 1 February 2019.
    - Invited tutorial at the 19th International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS), Samos, Greece, 7 July 2019.
    - Invited tutorial at 66th International Electron Devices Meeting (IEDM), Virtual, 12 December 2020.
    - Invited talk at Apple, Cupertino, CA, USA, Virtual, 7 July 2021.
    - Invited talk at Korea Computing Industry Association PIM Software Workshop (KCIA), Virtual, 7 October 2021.
    - Education Class at Embedded Systems Week (ESWEEK), Virtual, 9 October 2021.
    - Invited lecture at IEEE Electron Devices Society Fall School (EDS FALL SCHOOL), Virtual, 26 October 2021.
    - **Plenary keynote talk** at Tsinghua Future Chips Forum (FUTURE CHIPS), Virtual, 17 December 2021.
    - **Keynote talk** at the National Technical University of Athens Research Challenges in Computer Science Symposium (NTUA-RCCS), Virtual, 18 January 2022.
    - **Keynote talk** at the Intel Interconnect and Connectivity Summit (IICS), Virtual, 9 February 2022.
    - **Keynote talk** at the Thoughtworks Engineering for Research Symposium (E4R), Virtual, 19 February 2022.
    - **Keynote talk** at the Joint International Workshop on Big Data Management on Emerging Hardware and Data Management on Virtualized Active Systems (HardBD & Active), held with the International Conference on Data Engineering (ICDE), Virtual, 9 May 2022.
    - **Keynote talk** at the 13th Annual Non-Volatile Memories Workshop (NVMW), San Diego, CA, USA, 10 May 2022.
    - Invited talk at Micron, Virtual, 3 June 2022.
    - Invited talk at EPFL CIS Edge AI Summer School, Lausanne, Switzerland, 15 June 2022.
    - **Keynote talk** at HiPChips Chiplet Workshop, held with the International Symposium on Computer Architecture (ISCA), New York, NY, 19 June 2022.
    - Invited tutorial at the 22nd International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS), Samos, Greece, 3 July 2022.
    - **Plenary keynote talk** at 20th IEEE Computer Society Annual Symposium on VLSI (ISVLSI), Paphos, Cyprus, 4 July 2022.
    - Invited talk at the Stanford Mini Symposium, Palo Alto, CA, USA, 15 July 2022.
    - Invited talk at MATEO 2022 - Multicore Architectures and Their Effective Operation (MATEO), Barcelona, Spain, 2 September 2022.
    - **Distinguished lecture** at Simon Fraser University, Burnaby, BC, Canada, 22 September 2022.
    - Invited Talk at Microsoft Azure Hardware Learning Session, Virtual, 27 September 2022.
    - Invited tutorial at the 31st International Conference on Parallel Architectures and Compilation Techniques (PACT), Chicago, IL, USA, 8 October 2022.
    - Invited tutorial at the DFG SPP Summer School on Scalable Data Management for Future Hardware, Seeheim, Germany, 12 October 2022.
    - Invited talk at the SONOVA Workshop, Zurich, Switzerland, 24 October 2022.
    - Invited talk at the LiG Seminar in Grenoble, Virtual, 29 November 2022.
    - Invited talk at the University of California, Los Angeles (UCLA), Los Angeles, CA, USA, 13 January 2023.
    - **Keynote talk** at the Memory-Centric Computing for Data-Intensive Workloads (MCC), held with the HiPEAC Conference, Toulouse, France, 18 January 2023.
    - **Keynote talk** at the 5th Workshop on Accelerated Machine Learning (AccML), held with the HiPEAC Conference, Toulouse, France, 18 January 2023.
    - **Opening talk** at the 1st Real Processing in Memory Systems Tutorial (RealPIM), held with the HPCA Conference, Montreal, QC, Canada, 26 February 2023.
    - **Opening talk** at the 2nd Real Processing in Memory Systems Tutorial (RealPIM), held with the ASPLOS Conference, Vancouver, BC, Canada, 26 March 2023.
    - **Opening talk** at the 3rd DFG Workshop on Operating Systems (WSOS), Gunzburg, Germany, 3 April 2023.
    - Educational Session talk at the 2023 IEEE Custom Integrated Circuits Conference (CICC), 23 April 2023.
    - Invited Talk at the National Technical University of Athens (NTUA), Athens, Greece, 3 May 2023.
    - **Keynote talk** at the eda Workshop, Hannover, Germany, 8 May 2023.
    - **Closing Keynote talk** at the Huawei Global Software Technology Summit, Dresden, Germany, 1 June 2023.
    - **Keynote talk** at the Montenegrin Academy of Sciences (CANU), Podgorica, Montenegro, 8 June 2023.
    - **Opening talk** at the 3rd Real Processing in Memory Systems Tutorial (RealPIM), held with the ISCA Conference, Orlando, FL, USA, 18 June 2023.
    - Invited Special Session Talk at the 23rd International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS), Samos, Greece, 5 July 2023.
    - **Keynote talk** at the In-Memory Architectures and Computing Applications (IMACAW) Workshop, held with the DAC-60 Conference, San Francisco, CA, USA, 9 July 2023.

- Invited Lightning talk at the Design Automation Conference, San Francisco, CA, USA, 13 July 2023.
- Invited talk at Groq/Maxeler, London, UK, 28 July 2023.
- Invited talk at Hunan University, China, Virtual, 17 August 2023.
- **Keynote talk** at the DOT-PIM Workshop at ESWEEK, 21 September 2023.
- Invited talk at Rivos, Sunnyvale, CA, USA, 25 October 2023.
- **Keynote talk** at the SLIP Workshop at ICCAD, 2 November 2023.
- Invited talk at the School of AI Algiers AI Summit, 14 December 2023.
- The University of Texas at Austin ECE Colloquium, Austin, TX, USA, 31 January 2024.
- **Distinguished Lecture** at Duke University Athena AI Institute, Durham, NC, USA, 15 February 2024.
- **Distinguished Lecture** at Arizona State University ACME Center, Tempe, AZ, USA, 28 March 2024.
- Stanford University SystemX Seminar, Stanford, CA, USA, 4 April 2024.
- Invited Talk at the 50th Asilomar Microcomputer Workshop (AMW), Asilomar, CA, USA, 25 April 2024.
- **Distinguished Lecture** at University of Toronto CS Department, Toronto, ON, Canada, 7 May 2024.
- **Keynote talk** IDEAS Center Industrial Advisory Board Meeting, Sunnyvale, CA, USA, 14 May 2024.
- **Keynote talk** at the First EFCL (ETH Future Computing Laboratory) Summer School, Zurich, Switzerland, 6 June 2024.
- Invited talk at the VAIL Computer Elements Workshop, Vail, CO, USA, 10 June 2024.
- Invited talk at the Samsung Forum at Samsung Electronics, San Jose, CA, USA, 18 June 2024.
- Invited talk at Rambus Labs, Sunnyvale, CA, USA, 20 August 2024.
- Invited talk at Samsung, San Jose, CA, USA, 4 September 2024.
- **Keynote talk** at the 13th International Workshop on Characterization and Modeling of Memory Devices (IWCM2), Milan, Italy, 3 October 2024.
- **Keynote talk** at the Future of Science Prize Symposium, Hong Kong, 2 November 2024.
- **Keynote talk** at the 1st International Workshop on Sustainable AI, Zurich, Switzerland, 14 November 2024.
- **Keynote talk** at the 31st IEEE International Conference on Electronics Circuits and Systems, Nancy, France, 19 November 2024.
- Invited talk at Huazhong University of Science and Technology, Wuhan, China, 28 November 2024.
- **Keynote talk** at the 10th Energy Efficient Machine Learning and Cognitive Computing Workshop (EMC2), held with HPCA, Las Vegas, NV, USA, 2 March 2025.
- **Distinguished Lecture** at University of California at Irvine CS Department, Irvine, CA, USA, 7 March 2025.
- **Computer Engineering Seminar** at Georgia Institute of Technology, Atlanta, GA, USA, 1 May 2025.

18. "Computation in Memory: An Architectural Perspective,"
    - Invited talk at the 57th Design Automation Conference Tutorial on New Era of Compute-In-Memory (DAC), Virtual, 20 July 2020.

19. "Enabling Practical, Efficient, and Large-Scale Computation Near Data"
    - Microsoft Research, Cambridge, UK, 2 February 2017.

20. "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.

21. "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.

22. Short Course on "Memory Systems" or "Memory Systems and Memory-Centric Computing Systems: Challenges and Opportunities" 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.
    - HiPEAC ACACES Summer School (6 hours), Fiuggi, Italy, July 9-13, 2018.
    - Technion Short Course (26 hours), Haifa, Israel, October 8-12, 2018.
    - University of Texas at Austin ECE Department Endowed Lecture Series (3 hours), Austin, TX, 29, 31 January and 2 February 2019.
    - TU Wien Short Course (20 hours), Vienna, Austria, 12-19 June 2019.
    - SAMOS 2019 Tutorial (4.5 hours), Samos, Greece, 7 July 2019.
    - IPM Summer School (6 hours), Tehran, Iran, 27-28 August 2019.
    - NiPS Summer School (4.5 hours), Perugia, Italy, 3 September 2019.
    - IEDM Tutorial (2 hours), Virtual, 12 December 2020.
    - SAMOS 2022 Tutorial (4.5 hours), Samos, Greece, 3 July 2022.
    - EFCL Summer School, Computer Architecture and Memory Systems Course (4 hours), Zurich, Switzerland, 3 June 2024.
    - HiPEAC ACACES Summer School (6 hours), Fiuggi, Italy, July 14-20, 2024.

23. "QoS-Aware and High-Performance Multi-Core Memory Systems"
    - Google, Mountain View, CA, USA, 18 April 2018.

24. "Venice: Improving Solid-State Drive Parallelism at Low Cost via Conflict-Free Accesses,"
    - Invited Talk at the Future of Memory and Storage Conference (FMS), Santa Clara, CA, USA, 7 August 2024.

25. "Storage-Centric Computing for Genomics and Metagenomics,"
    - Invited Talk at the Future of Memory and Storage Conference (FMS), Santa Clara, CA, USA, 6 August 2024.

26. "DRAM Bender: Open Source & Easy to Use DRAM Testing Infrastructure,"
    - Invited Talk at the Flash Memory Summit (FMS), Santa Clara, CA, USA, 10 August 2023.

27. "Sibyl: Data Placement in Hybrid Storage Systems using Reinforcement Learning,"
    - Invited Talk at the Flash Memory Summit (FMS), Santa Clara, CA, USA, 10 August 2023.

28. "In-Memory Acceleration of Genome Analysis,"
    - Invited Talk at the Flash Memory Summit (FMS), Santa Clara, CA, USA, 10 August 2023.

29. "Achieving the Next Breakthroughs in NAND Flash Memory: Changing Our Fixed Mindsets"
    - Flash Memory Summit, 10 August 2017.

30. "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.

31. "Solving the Memory Problem"
    - ETH Systems Group Industry Retreat, Engelberg, Switzerland, 18 January 2016.
    - ETH Systems Group Industry Retreat, Engelberg, Switzerland, 30 January 2017.

32. "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.

33. "Microarchitecture Research Directions,"
    - Invited Panel Talk at the Microarchitecture Research Panel at the 54th International Symposium on Microarchitecture (MICRO), Virtual, 20 October 2021.

34. "SAFARI Research Group: Introduction & Research,"
    - Invited Talk at the ETH Future Computing Laboratory Huawei Day, Virtual, 19 October 2021.
    - Talk at ETH Future Computing Laboratory Welcome Workshop (EFCL), Virtual, 6 July 2021.

35. "Applying to Graduate School & Doing Impactful Research,"
    - Invited Panel Talk at the 3rd Undergraduate Mentoring Workshop, held with the 48th International Symposium on Computer Architecture (ISCA), Virtual, 18 June 2021.

36. ""How to Build an Impactful Research Group,"
    - Invited Talk at the 56th Design Automation Conference Early Career Workshop (DAC), Las Vegas, NV, USA, June 2019.
    - Invited Talk at the 57th Design Automation Conference Early Career Workshop (DAC), Virtual, 19 July 2020.

37. "ThyNVM: Software-Transparent Crash Consistency for Persistent Memory"
    - Flash Memory Summit, Santa Clara, CA, 8 August 2016.
    - Google, Mountain View, CA, 22 August 2016.

38. "Large-Scale Study of In-the-Field Flash Failures"
    - Flash Memory Summit, Santa Clara, CA, 10 August 2016.

39. "Read Disturb Errors in MLC NAND Flash Memory"
    - Flash Memory Summit, Santa Clara, CA, 12 August 2015.

40. "A Case for Autonomous Memory"
    - Google Systems Research Workshop, San Francisco, CA, 26 June 2015.

41. "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.

42. "Using Commodity Memory Devices to Support Fundamental Security Primitives"
    - Bogazici University, Istanbul, Turkey, 25 March 2019.
    - IBM Research, Yorktown Heights, NY, USA, 26 April 2019.
    - **Visionary talk** at the Energy Secure System Architectures Workshop (ESSA), co-located with HOST, Washington, DC, USA, 10 May 2019.

43. "Memory Reliability, Security, and Beyond: Three Key Issues in Modern Systems"
    - DAC Design Automation Summer School, Austin, TX, USA, 18 June 2017.

44. "Memory System Design for AI/ML Accelerators & ML/AI Techniques for Memory System Design,"
    - Invited Talk at the Semiconductor Research Corporation (SRC) AI Hardware Program Kickoff Meeting, Virtual, 17 June 2020.
    - Invited Talk at the Semiconductor Research Corporation (SRC) AI Hardware Program Annual Review, Virtual, 29 September 2021.

- Invited Talk at the Semiconductor Research Corporation (SRC) AI Hardware Program Annual Review, San Diego, CA, USA, 30 August 2022.
- Invited Talk at the Intel-SRC Research Program Review Conference, Virtual, 10 November 2022.
- Invited Talk at the Semiconductor Research Corporation (SRC) AI Hardware Program Kickoff Meeting, Virtual (Almaden, CA, USA), 22 May 2023.
- Invited Talk at the Semiconductor Research Corporation (SRC) AI Hardware Program Annual Review, Boston, MA, USA, 23 July 2024.
- Invited Talk at Intel ArchFest, Intel, Hillsboro, OR, USA, 23 August 2024.

45. "Architectural Frameworks to Facilitate Practical, Efficient and Transparent Computation Near Data"
    - SRC (Semiconductor Research Corporation) Annual Review, Ann Arbor, MI, USA, 31 May 2017.
    - SRC (Semiconductor Research Corporation) Annual Review, Hillsboro, OR, USA, 18 April 2018.

46. "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.

47. "Error Analysis and Management for MLC NAND Flash Memory"
    - Flash Memory Summit (45-minute talk), Santa Clara, CA, 7 August 2014.

48. "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.

49. "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.

50. "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.

51. "Understanding and Overcoming Challenges of DRAM Refresh"
    - Extreme Scale Scientific Computing Workshop, Moscow, Russia, 30 June 2014.

52. "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.

53. "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 Politecnica de Catalunya, Barcelona, Spain, 26 July 2013.
    - Nvidia, Santa Clara, CA, 5 August 2013.

54. "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 Politecnica 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.

55. "Some Ideas in Designing Scalable and Efficient Multi-Core Systems"
    - Apple, Cupertino, CA, 6 June 2012.

56. "A Potpourri of Recent SAFARI Research"
    - Apple, Cupertino, CA, 13 September 2023.

57. "Machine Learning Driven Memory and Storage Hierarchy Management,"
    - VMware, Palo Alto, CA, 20 September 2023.

58. "Security of PIM Systems,"
    - Invited Talk at Dagstuhl Seminar on "Microarchitectural Attacks and Defenses", Germany, 27 November 2023.

59. "Victima: Drastically Increasing Address Translation Reach by Leveraging Underutilized Cache Resources,"
    - Invited Talk at VMware Kernel Memory Team Journal Club, Palo Alto, CA, USA, 12 April 2024.

60. "Utopia: Fast and Efficient Address Translation via Hybrid Restrictive & Flexible Virtual-to-Physical Address Mappings,"
    - Invited Talk at VMware Kernel Memory Team Journal Club, Palo Alto, CA, USA, 3 May 2024.

61. "Concurrent Autonomous Self-Test for Uncore Components in SoCs"
    - SK Hynix, Seoul, Korea, 23 June 2012.

62. "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.

63. "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.

64. "Some (Security-Related) Challenges in Future Computing Platforms"
    - Carnegie Mellon University College of Engineering Dean's Council Meeting, Pittsburgh, PA, 26 April 2012.

65. "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.

66. "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.

67. "Issues in DRAM and NVM based Main Memory"
    - Carnegie Mellon University Parallel Data Lab Retreat, Bedford Springs, PA, 8 November 2011.

68. "Computer Performance"
    - Carnegie Mellon University ECE Department Advisory Board Presentation, 21 September 2011.

69. "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.

70. "Architecture and System-Level Challenges Related to Memory"
    - Focus Center Research Program Memory Cross-Cut Workshop, Cambridge, MA, 12 May 2011.

71. "Towards Practical Bufferless On-Chip Networks"
    - Intel, Santa Clara, CA, 9 March 2011.

72. "Some Ideas for Efficient and High-Performance Core Design"
    - Intel Corporation ARO Swiss Army Processor Workshop, Hillsboro, OR, 29 April 2011.

73. "PCM (NVM) as Main Memory: Opportunities and Challenges"
    - Carnegie Mellon University, Parallel Data Lab Retreat, Pittsburgh, PA, October 25, 2010.

74. "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.

75. "End-to-end QoS-aware, High-Performance and Customizable Many-Core Memory Systems"
    - Intel Memory Hierarchy Meeting, Hillsboro, OR, 8 October 2010.

76. "Rethinking Core Design in the Power-Constrained Many-Core Era"
    - Intel Core Workshop, Hillsboro, OR, 27 September 2010.

77. "Some Ideas for ILP Research"
    - CRA Workshop on Advancing Computer Architecture Research, Seattle, WA, 20 September 2010.

78. "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.

79. "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.

80. "Asymmetry Everywhere (with Automatic Resource Management)"
    - CRA Workshop on Advancing Computer Architecture Research, San Diego, CA, 22 February 2010.

81. "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.

82. "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.

83. "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.

84. "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.

85. "MSR Computer Architecture Group: Vision and Projects"
    - Presentation to Rico Malvar, MSR-Redmond Director, Microsoft Research, 12 December 2007.

86. "Hardware-Based Devirtualization of Virtual Function Calls"
    - MSR Systems and Networking Seminar, Microsoft Research, 7 December 2006.

87. "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.

88. "Efficient Runahead Execution"

- Advanced Micro Devices, Sunnyvale, CA, USA, May 2005.
- Intel Barcelona Research Center, Barcelona, Spain, November 2005.

89. "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.
90. "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)

91. "Memory-Centric Computing: Enabling Fundamentally Efficient Computers" *Keynote talk at the 10th Energy Efficient Machine Learning and Cognitive Computing Workshop (EMC2), held with HPCA*, Las Vegas, NV, USA, 2 March 2025.
92. "Memory-Centric Computing: Recent Advances in Processing-in-DRAM," *Invited talk at the 70th International Electron Devices Meeting (IEDM)*, San Francisco, CA, USA, 9 December 2024.
93. "Storage-Centric Computing: Enabling Fundamentally Efficient Computers" *Keynote talk at the CCF Storage Conference*, Guangzhou, China, 1 December 2024.
94. "Memory-Centric Computing: Enabling Fundamentally Efficient Computers" *Keynote talk at the 31st IEEE International Conference on Electronics Circuits and Systems (ICECS)*, Nancy, France, 19 November 2024.
95. "Memory-Centric Computing: Enabling Fundamentally Efficient Computers" *Keynote talk at the 1st International Workshop on Sustainable AI*, Zurich, Switzerland, 14 November 2024.
96. "Memory-Centric Computing," *Keynote talk at the Future of Science Prize Symposium*, Hong Kong, 2 November 2024.
97. "Future of Computer Architecture and Hardware Security," *Keynote talk at hardwear.io Memory Security Track*, Amsterdam, Netherlands, 25 October 2024.
98. "Memory-Centric Computing," *Keynote talk at the 13th International Workshop on Characterization and Modeling of Memory Devices (IWCM2)*, Milan, Italy, 3 October 2024.
99. "Venice: Improving Solid-State Drive Parallelism at Low Cost via Conflict-Free Accesses," *Invited Talk at the Future of Memory and Storage Conference (FMS)*, Santa Clara, CA, USA, 7 August 2024.
100. "Storage-Centric Computing for Genomics and Metagenomics," *Invited Talk at the Future of Memory and Storage Conference (FMS)*, Santa Clara, CA, USA, 6 August 2024.
101. "ML-Assisted Memory and Storage Management," *Invited talk at the AI4FACD (AI for Fully Automated Chip Design) Workshop, held with ISCA*, Virtual, 29 June 2024.
102. "Memory-Centric Computing," *Invited talk at the VAIL Computer Elements Workshop*, Vail, CO, USA, 10 June 2024.
103. "Future of Computer Architecture and Hardware Security," *Keynote talk at Qualcomm Product Security Summit*, San Diego, CA, USA, 16 May 2024.
104. "Memory-Centric Computing," *Invited Talk at the 50th Asilomar Microcomputer Workshop (AMW)*, Asilomar, CA, USA, 25 April 2024.
105. "Memory-Centric Computing," *Keynote talk at the SLIP Workshop at ICCAD*, San Francisco, CA, USA, 2 November 2023.
106. "Memory-Centric Computing," *Keynote Talk at at the DOT-PIM Workshop at ESWEEK*, Hamburg, Germany, 21 September 2023.
107. "In-Memory Acceleration of Genome Analysis," *Invited Talk at the Flash Memory Summit (FMS)*, Santa Clara, CA, USA, 10 August 2023.
108. "DRAM Bender: Open Source & Easy to Use DRAM Testing Infrastructure," *Invited Talk at the Flash Memory Summit (FMS)*, Santa Clara, CA, USA, 10 August 2023.
109. "Sibyl: Data Placement in Hybrid Storage Systems using Reinforcement Learning," *Invited Talk at the Flash Memory Summit (FMS)*, Santa Clara, CA, USA, 10 August 2023.
110. "Fundamentally Understanding and Solving RowHammer," *Invited Talk at the Flash Memory Summit (FMS)*, Santa Clara, CA, USA, 8 August 2023.
111. "Memory-Centric Computing," *Invited Lightning Talk at the 60th Design Automation Conference (DAC)*, San Francisco, CA, USA, 13 July 2023.
112. "Accelerating Genome Analysis," *Invited Lightning Talk at the 60th Design Automation Conference (DAC)*, San Francisco, CA, USA, 11 July 2023.
113. "Memory-Centric Computing," *Keynote Talk at the IMACAW Workshop of the 60th Design Automation Conference (DAC)*, San Francisco, CA, USA, 9 July 2023.
114. "Memory-Centric Computing," *Invited Special Session Talk at the 23rd International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS)*, Samos, Greece, 5 July 2023.
115. "Memory-Centric Computing," *Invited Education Session Talk at the Custom Integrated Circuits Conference (CICC)*, San Antonio, TX, USA, 23 April 2023.
116. "Memory-Centric Computing," *Keynote Talk at the 5th Workshop on Accelerated Machine Learning (AccML), held with the HiPEAC Conference*, Toulouse, France, 18 January 2023.

117. "Memory-Centric Computing," *Keynote Talk at the Memory-Centric Computing for Data-Intensive Workloads (MCC), held with the HiPEAC Conference*, Toulouse, France, 18 January 2023.
118. "Accelerating Genome Analysis," *Invited Talk at the Montenegro Academy of Sciences Conference*, Montenegro, 2 November 2022.
119. "Memory-Centric Computing," *Invited Tutorial at the 31st International Conference on Parallel Architectures and Compilation Techniques (PACT)*, Chicago, IL, USA, 8 October 2022.
120. "Memory-Centric Computing," *Invited Talk at INSAIT Conference on Emerging Trends in AI and Computing Research*, Sofia, Bulgaria, 1 October 2022.
121. "Memory-Centric Computing," *Plenary Keynote Talk at 20th IEEE Computer Society Annual Symposium on VLSI (ISVLSI)*, Paphos, Cyprus, 4 July 2022.
122. "Memory-Centric Computing," *Invited Tutorial at the 22nd International Conference on Embedded Computer Systems: Architectures, Modeling, and Simulation (SAMOS)*, Samos, Greece, 3 July 2022.
123. "Memory-Centric Computing," *Keynote Talk at HiPChips Chiplet Workshop, held with the International Symposium on Computer Architecture (ISCA)*, New York, NY, 19 June 2022.
124. "Accelerating Genome Analysis," *Keynote Talk at 4th Workshop on Accelerator Architecture in Computational Biology and Bioinformatics (AACBB), held with the International Symposium on Computer Architecture (ISCA)*, New York, NY, USA, 18 June 2022.
125. "Security Aspects of DRAM: The Story of RowHammer," *Invited Tutorial at 14th IEEE Electron Devices Society International Memory Workshop (IMW)*, Dresden, Germany, 15 May 2022.
126. "Memory-Centric Computing," *Keynote Talk at the 13th Annual Non-Volatile Memories Workshop (NVMW)*, San Diego, CA, USA, 10 May 2022.
127. "Memory-Centric Computing," *Keynote Talk at the Joint International Workshop on Big Data Management on Emerging Hardware and Data Management on Virtualized Active Systems (HardBD & Active), held with the International Conference on Data Engineering (ICDE)*, Virtual, 9 May 2022.
128. "Accelerating Genome Analysis," *Keynote Talk at the Systems for Post-Moore Architectures Workshop (SPMA), held with the EuroSys Conference (EUROSYS)*, Virtual, 5 April 2022.
129. "The Story of RowHammer," *Invited Talk at the Workshop on Robust and Safe Software 2.0 (RSS2), held with the 27th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Virtual, 28 February 2022.
130. "Memory-Centric Computing," *Keynote Talk at the Intel Interconnect & Connectivity Summit (IICS)*, Virtual, 9 February 2022.
131. "Memory-Centric Computing," *Plenary Keynote Talk at Tsinghua Future Chips Forum (FUTURE CHIPS)*, Virtual, 17 December 2021.
132. "Future Computing Platforms: Challenges and Opportunities," *Closing Talk at IEEE Data & Storage Symposium (DSS)*, Virtual, 1 December 2021.
133. "Intelligent Architectures for Intelligent Systems," *Keynote Talk at 15th International Conference on Networking, Architecture, and Storage (NAS)*, Riverside, CA, USA, 25 October 2021.
134. "Memory-Centric Computing," *Education Class at Embedded Systems Week (ESWEEK)*, Virtual, 9 October 2021.
135. "Intelligent Architectures for Intelligent Machines," *Keynote Talk at TUBA World Conference on Energy Science and Technology (WCEST)*, Virtual, 11 August 2021.
136. "Intelligent Architectures for Intelligent Systems," *Invited Talk at Supercomputing Frontiers Europe (SCFE)*, Virtual, 21 July 2021.
137. "Intelligent Architectures for Intelligent Machines," *Closing Talk at IEEE Data & Storage Symposium (DSS)*, Virtual, 30 June 2021.
138. "Intelligent Architectures for Intelligent Machines," *Keynote Talk at Future of Information and Communication Conference (FICC)*, Virtual, 29 April 2021.
139. "Future Computing Architectures: Challenges and Opportunities," *Keynote Talk at CSCON IEEE Turkey Computer Society Student Branch Conference (CSCON)*, Virtual, 2 April 2021.
140. "Intelligent Architectures for Intelligent Computing Systems," *Invited Special Session Talk at Design, Automation and Test in Europe (DATE)*, Virtual, 2 February 2021.
141. "The Story of RowHammer," *Keynote Talk at Secure Hardware, Architectures, and Operating Systems Workshop (SeHAS)*, held with HiPEAC 2021 Conference, Virtual, 19 January 2021.
142. "Memory-Centric Computing Systems," *Invited Tutorial at 66th International Electron Devices Meeting (IEDM)*, Virtual, 12 December 2020.
143. "Intelligent Architectures for Intelligent Machines," *Keynote Talk at National Science Foundation Workshop on Processing-In-Memory Technology (NSF-PIM)*, Virtual, 26 October 2020.
144. "Intelligent Architectures for Intelligent Machines," *Keynote Talk at 13th ACM International Systems and Storage Conference (SYSTOR)*, Virtual, 13 October 2020.
145. "Revisiting RowHammer," *Invited Talk at 28th IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SoC), Special Session on AI Frameworks*, Virtual, 9 October 2020.
146. "Intelligent Architectures for Intelligent Machines," *Plenary Keynote Talk at the 2020 International Symposia on VLSI (VLSI)*, Virtual, 11 August 2020.
147. "Intelligent Architectures for Intelligent Machines," *Keynote talk at the 17th ChinaSys Workshop*, Zhuhai, China, 20 December 2019.

148. "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Keynote talk at 37th IEEE International Conference on Computer Design (ICCD)*, Abu Dhabi, UAE, 19 November 2019.
149. "Intelligent Architectures for Intelligent Machines," *Keynote talk at the HiPEAC Computing Systems Week Autumn 2019 (HIPEAC CSWEEK)*, Bilbao, Spain, 28 October 2019.
150. "RowHammer and Beyond," *Keynote talk at the HiPEAC Computing Systems Week Autumn 2019 Resilience in High Performance Computing Thematic Session (RESILIENTHPC)*, Bilbao, Spain, 28 October 2019.
151. "RowHammer and Beyond," *Keynote talk at the 32nd IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, Noordwijk, the Netherlands, 2 October 2019.
152. "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Keynote talk at the International Symposium on Advanced Parallel Processing Technology (APPT)*, Tianjin, China, 16 August 2019.
153. "Intelligent Architectures for Intelligent Machines," *Keynote talk at SAMOS XIX (19th International Conference on Embedded Computer Systems: Architectures, Modeling and Simulation)*, Samos, Greece, July 2019.
154. "Enabling Practical Processing in and near Memory for Data-Intensive Computing," *Special Invited Talk at the Design Automation Conference Special Session on In-Memory Computing (DAC)*, Las Vegas, NV, USA, 4 June 2019.
155. "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Keynote talk at the 29th ACM Great Lakes Symposium on VLSI (GLSVLSI)*, Washington, DC, USA, 10 May 2019.
156. "Intelligent Architectures for Intelligent Machines," *Keynote Talk at the Mubadala-SRC Forum on the Future of Artificial Intelligence Hardware Systems*, Abu Dhabi, UAE, April 2019.
157. "RowHammer and Beyond," *Keynote Talk at the 10th International Workshop on Constructive Side-Channel Analysis and Secure Design (COSADE)*, Darmstadt, Germany, April 2019.
158. "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Special invited talk at the International Solid State Circuits Conference (ISSCC) Forum on "Intelligence at the Edge: How Can We Make Machine Learning More Energy Efficient?"*, San Francisco, CA, USA, February 2019.
159. "Accelerating Genome Analysis: A Primer on an Ongoing Journey," *Keynote talk at the 2nd Workshop on Accelerator Architecture in Computational Biology and Bioinformatics (AACBB)*, Washington, DC, USA, February 2019.
160. "Enabling In-Memory Computation in Mobile Systems: Changing the Computing Paradigm for High Efficiency *Keynote talk at 3rd Workshop on Mobile System Technologies (MST)*, Milan, Italy, November 2018.
161. "RowHammer," *Special Invited Talk at the First Workshop on Top Picks in Hardware and Embedded Security (TopinHES), held with ICCAD*, San Diego, CA, USA, November 2018.
162. "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Keynote talk at the 2nd Workshop on Highly Efficient Neural Processing (HENP), held with ESWEEK*, Turin, Italy, October 2018.
163. "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Keynote talk at Cordoba HiPerNav Week*, Cordoba, Spain, September 2018.
164. "Rethinking Memory System Design (for Data-Intensive Computing)," *Joint Keynote talk at the 9th International Symposium on Non-Volatile Memory Systems and Applications (NVMSA) and the 24th International Symposium on Real-Time Computing Systems and Applications (RTCSA)*, Hakodate, Japan, August 2018.
165. "Rethinking Memory System Design: Robustness, Energy, Performance," *Keynote talk at the 28th International Online Testing Symposium (IOLTS)*, Costa Brava, Spain, July 2018.
166. "Processing Data Where It Makes Sense in Modern Computing Systems: Enabling In-Memory Computation," *Keynote talk at 7th Mediterranean Conference on Embedded Computing (MECO)*, Budva, Montenegro, June 2018.
167. "Accelerating Genome Analysis: A Primer on an Ongoing Journey," *Keynote talk at the 17th International Workshop on High Performance Computational Biology (HiCOMB), held with IPDPS*, Vancouver, BC, Canada, May 2018.
168. "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.
169. "Future Computing Architectures," *Keynote talk at the 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)*, Cambridge, UK, March 2018.
170. "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.
171. "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.
172. "Processing Data Where It Makes Sense: Enabling In-Memory Computation *Keynote talk at 3rd Workshop on Mobile System Technologies (MST)*, Milan, Italy, October 2017.
173. "Opportunities and Challenges of Emerging Memory Technologies" *Invited talk at the ARM Research Summit*, Cambridge, UK, September 2017.
174. "Achieving the Next Breakthroughs in NAND Flash Memory: Changing Our Fixed Mindsets" *Flash Memory Summit*, Santa Clara, CA, August 2017.

175. "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.
176. "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.
177. "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.
178. "Rethinking Memory System Design" *Keynote talk at 27th International Symposium on Rapid System Prototyping (RSP)*, Pittsburgh, PA, USA, October 2016.
179. "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.
180. "Rethinking Memory System Design: Business As Usual in the Next Decade?" *Keynote talk at the First ARM Research Summit*, Cambridge, UK, September 2016.
181. "Large-Scale Study of In-the-Field Flash Failures," *Flash Memory Summit*, Santa Clara, CA, August 2016.
182. "ThyNVM: Software-Transparent Crash Consistency for Persistent Memory," *Flash Memory Summit*, Santa Clara, CA, August 2016.
183. "Rethinking Memory System Design," *Keynote talk at the ACM SIGPLAN International Symposium on Memory Management (ISMM)*, Santa Barbara, CA, 14 June 2016.
184. "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.
185. "Reliability and Security Issues of DRAM and NAND Flash Scaling," *Invited talk at the Memory Reliability Forum at HPCA*, Barcelona, Spain, March 2016.
186. "Rethinking Memory System Design (along with Interconnects)," *Keynote talk at the 8th International Workshop on Network on Chip Architectures (NoCArc)*, Honolulu, Hawaii, December 2015.
187. "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.
188. "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.
189. "Rethinking Memory System Design (for Data-Intensive Computing)," *Keynote talk at the 18th International Symposium on Computer Architecture and Digital Systems (CADS)*, October 2015.
190. "Read Disturb Errors in MLC NAND Flash Memory," *Flash Memory Summit*, Santa Clara, CA, August 2015.
191. "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.
192. "Rethinking Memory System Design (for Data-Intensive Computing)," *Special invited talk at ISC High Performance*, Frankfurt, Germany, July 2015.
193. "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.
194. "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.
195. "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.
196. "The Heterogeneous Block Architecture," *32nd IEEE International Conference on Computer Design*, Seoul, South Korea, October 2014.
197. "Error Analysis and Management for MLC NAND Flash Memory," *Flash Memory Summit*, Santa Clara, CA, August 2014.
198. "Rethinking Memory System Design (for Data-Intensive Computing)," *Plenary Talk at the Summer Supercomputing Academy*, Moscow, Russia, July 2014.
199. "Understanding and Overcoming Challenges of DRAM Refresh," *Extreme Scale Scientific Computing Workshop*, Moscow, Russia, June 2014.
200. "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.
201. "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.
202. "Rethinking Memory System Design for Data-Intensive Computing," *Keynote talk at the Industry-Academia Partnership Cloud Workshop*, Pittsburgh, PA, April 2014.
203. "Rethinking Memory System Design for Data-Intensive Computing," *Keynote talk at the Industry-Academia Partnership Cloud Workshop*, Mountain View, CA, December 2013.
204. "Memory Scaling: A Systems Architecture Perspective," *MemCon 2013*, Santa Clara, CA, August 2013.
205. "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.

206. "A Case for Efficient Hardware-Software Cooperative Management of Storage and Memory," *5th Workshop on Energy-Efficient Design*, Tel-Aviv, Israel, June 2013.
207. "Memory Scaling: A Systems Architecture Perspective," *Invited talk at the 5th International Memory Workshop*, Monterey, CA, May 2013.
208. "Scalable Memory, Compute and Communication Architectures," *Industry-Academia Partnership Cloud Workshop*, Cambridge, MA, May 2013.
209. "Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime," *30th International Conference on Computer Design*, Montreal, Quebec, Canada, September 2012.
210. "Bottleneck Identification and Scheduling in Multithreaded Applications," *17th International Conference on Architectural Support for Programming Languages and Operating Systems*, London, UK, March 2012.
211. "Data Marshaling for Multi-core Architectures," *37th International Symposium on Computer Architecture*, St. Malo, France, June 2010.
212. "Parallelism-aware Batch Scheduling: Enhancing both Performance and Fairness of Shared DRAM Systems," *35th International Symposium on Computer Architecture*, Beijing, China, June 2008.
213. "Stall-Time Fair Memory Access Scheduling," *40th International Symposium on Microarchitecture*, Chicago, IL, USA, December 2007.
214. "Memory Performance Attacks: Denial of Memory Service in Multi-Core Systems," *16th USENIX Security Symposium*, Boston, MA, USA, August 2007.
215. "Address-Value Delta Prediction," *38th International Symposium on Microarchitecture*, Barcelona, Spain, November 2005.
216. "Techniques for Efficient Processing in Runahead Execution Engines," *32nd International Symposium on Computer Architecture*, Madison, WI, USA, June 2005.
217. "Wrong Path Events and Their Application to Early Misprediction Detection and Recovery," *37th International Symposium on Microarchitecture*, Portland, OR, USA, December 2004.
218. "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.
219. "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 (Selected)

220. "The evolution of scaling laws and what it means for sustainable AI computing," *Panel at the 10th Energy Efficient Machine Learning and Cognitive Computing Workshop (EMC2), held with HPCA*, Las Vegas, NV, USA, 2 March 2025.
221. "Are We Done?" *Panel at the 4th Workshop on DRAM Security (DRAMSec), held with ISCA-51*, Virtual, 29 June 2024.
222. "ISCA-50, a birthday panel: celebrating the past and looking to the future," *Panel at the 50th International Symposium on Computer Architecture (ISCA)*, Orlando, FL, 19 June 2023.
223. "Security Challenges in Future Memory Technologies and Confidential Compute," *Panel at the 3rd Workshop on DRAM Security (DRAMSec), held with ISCA-50*, 17 June 2023.
224. "Next-Generation Data Storage Systems," *Panel at the Next Generation Data Storage Architecture Innovation Forum*, Luzern, Switzerland, 15 May 2023.
225. "Hardware Acceleration of Bioinformatics Workloads," *Panel Moderator at RECOMB 2023*, Istanbul, Turkey, 17 April 2023.
226. "Future of Storage Systems," *Panel at the Huawei OlympusMons Award Ceremony*, Emei, Sichuan, China, 7 April 2023.
227. "Microarchitecture Research Directions," *Microarchitecture Research Panel Talk at the 54th International Symposium on Microarchitecture (MICRO)*, Virtual, 20 October 2021.
228. ""Applying to Graduate School & Doing Impactful Research," *Invited Panel Talk and Discussion at the 3rd Undergraduate Mentoring Workshop, held with the 48th International Symposium on Computer Architecture (ISCA)*, Virtual, 18 June 2021.
229. "Is Rowhammer here to stay?" *Invited Panel Talk and Discussion at the Workshop on DRAM Security (DRAMSec), held with the 48th International Symposium on Computer Architecture (ISCA)*, Virtual, 17 June 2021.
230. "How to Build an Impactful Research Group" *DAC Early Career Workshop Panel*, Las Vegas, NV, USA, 2 June 2019.
231. "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.
232. "Memory and System Balance," *Panel at 4th Workshop on Irregular Applications: Architectures and Algorithms, held with Supercomputing (SC)*, New Orleans, LA, November 2014.
233. "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/@CMUCompArch` for online lecture videos for my courses at Carnegie Mellon University.
Please visit `https://www.youtube.com/OnurMutluLectures` for online lecture videos for my courses at ETH Zürich.

ETH Zürich, CS & ECE Departments, Instructor for Digital Design and Computer Architecture (Freshman Course), ***Every Spring, 2017-2025***

ETH Zürich, ECE Department, Instructor for Computer Architecture (Senior/Masters-level Course), ***Every Fall, 2020-2024***
ETH Zürich, ECE Department, Instructor for Fundamentals of Computer Architecture (Senior-level Course), ***Spring 2025***
ETH Zürich, CS & ECE Departments, Instructor for Seminar in Computer Architecture (Bachelors/Masters course), ***Every Spring and Fall, 2018-2025***
ETH Zürich, CS Department, Instructor for Computer Architecture (Senior/Masters-level Course), ***Fall 2017'18'19***
ETH Zürich, CS Department, Co-Instructor, Graduate Seminar on Hardware Architectures for Machine Learning, ***Spring 2017'18***
ETH Zürich, CS Department, Co-Instructor for Graduate Seminar on Hardware Acceleration for Data Processing, ***Fall 2016'17***
Carnegie Mellon University, ECE Department, Instructor for Undergraduate Computer Architecture Course ECE-447, ***Spring 2012'13'14'15***
- 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 2010'11'13'15***
- 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 2010'11, Fall 2012'14***
- 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 2001'03***
- 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 `https://safari.ethz.ch/safari-alumni/` for a somewhat more complete 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. Now: Open AI.
2. Yoongu Kim, "Architectural Techniques to Enhance DRAM Scaling," CMU ECE, defended May 2015. First job: Software Robotics, Boston, MA, USA. Now: Google.
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.
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: Assistant Professor at KMUTNB, Thailand. Now: MangoBoost.
7. Kevin Chang, "Understanding and Improving the Latency of DRAM-Based Memory Systems," CMU ECE, defended May 2017. First job: Facebook.
8. Yixin Luo, "Architectural Techniques for Improving NAND Flash Memory Reliability," CMU CS, defended February 2018. First job: Google. Now: Open AI.
9. Mohammed Alser, "Accelerating the Understanding of Life's Code Through Better Algorithms and Hardware Design," Bilkent CS, defended June 2018. **2018 IEEE Turkey Doctoral Dissertation Award.** First job: Assistant Professor at Georgia State University.
10. Justin J. Meza, "Large Scale Studies of Memory, Storage, and Network Failures in a Modern Data Center," CMU ECE, defended December 2018. First job: Research Scientist at Facebook.
11. Kevin Hsieh, "Machine Learning Systems for Highly-Distributed and Rapidly-Growing Data," CMU ECE, defended September 2019. First job: Researcher at Microsoft Research.
12. Nandita Vijaykumar, "Enhancing Programmability, Portability, and Performance with Rich Cross-Layer Abstractions," CMU ECE, defended October 2019. First job: Assistant Professor at the University of Toronto.

13. Jeremie Kim, "Improving DRAM Performance, Security, and Reliability by Understanding and Exploiting DRAM Timing Parameter Margins," CMU ECE, defended July 2020. **2020 EDAA Outstanding Dissertation Award.** First job: Apple.
14. Amirali Boroumand, "Practical Mechanisms for Reducing Processor-Memory Data Movement in Modern Workloads," CMU ECE, defended November 2020. First job: Researcher at Google.
15. Gagandeep Singh, "Designing, Modeling, and Optimizing Data-Intensive Computing Systems," TU Eindhoven, defended March 2021. First job: Researcher at AMD.
16. Nastaran Hajinazar, "Data-Centric and Data-Aware Frameworks for Fundamentally Efficient Data Handling in Modern Computing Systems," Simon Fraser University, defended June 2021. First job: Research Scientist at Intel Labs.
17. Damla Senol Cali, "Accelerating Genome Sequence Analysis via Efficient Hardware/Algorithm Co-Design," CMU ECE, defended July 2021. First job: Researcher at Bionano Genomics. Now: Researcher at 10x Genomics.
18. Minesh Patel, "Enabling Effective Error Mitigation in Modern Memory Chips that Use On-Die Error-Correcting Codes," ETH Zurich, defended October 2021. **2022 William C. Carter PhD Dissertation Award in Dependability. 2022 ETH Doctoral Medal.** First job: Assistant Professor at Rutgers University.
19. Christina Giannoula, "Accelerating Irregular Applications via Efficient Synchronization and Data Access Techniques," National Technical University of Athens, defended September 2022. First job: Postdoctoral Researcher at the University of Toronto. **2023 National Technical University of Athens (NTUA) Best Doctoral Thesis Award.**
20. Hasan Hassan, "Improving DRAM Performance, Reliability, and Security by Rigorously Understanding Intrinsic DRAM Operation," ETH Zurich, defended September 2022. **2023 EDAA Outstanding Dissertation Award.** First job: Architect at Rivos, Inc.
21. Abdullah Giray Yaglikci, "Enabling Efficient and Scalable DRAM Read Disturbance Mitigation via New Experimental Insights into Modern DRAM Chips," ETH Zurich, defended April 2024. **2025 William C. Carter PhD Dissertation Award in Dependability. 2025 VTS McCluskey Award Finalist. 2024 HOST PhD Competition Finalist.** First job: Postdoctoral Researcher at ETH Zurich.
22. Lukas Breitwieser, "Design and Analysis of an Extreme-Scale, High-Performance, and Modular Agent-Based Simulation Platform," ETH Zurich, defended November 2024. First job: Researcher at CERN.
23. Can Firtina, "Enabling Fast, Accurate, and Efficient Real-Time Genome Analysis via New Algorithms and Techniques," ETH Zurich, defended November 2024. First job: Postdoctoral Researcher at ETH Zurich.
24. Geraldo de Oliveira Junior, "New Tools, Programming Models, and System Support for Processing-in-Memory Architectures," ETH Zurich, defended April 2025.

Post-doctoral Researcher Alumni

1. Haiyu Mao, ETH Zürich, September 2020 to September 2024, First job: Lecturer (Assistant Professor) at King's College, London, UK
2. Juan Gomez Luna, August 2017 to October 2023, First job: Researcher at NVIDIA Research, Zurich, Switzerland
3. Mohammed Alser, September 2018 to August 2023, First job: Assistant Professor, Georgia State University, Atlanta, GA, USA
4. Gagandeep Singh, May 2021 to March 2023, First job: Researcher, AMD Research, Zurich, Switzerland
5. Jisung Park, ETH Zurich, September 2019 to August 2022, First job: Assistant Professor, Pohang University of Science and Technology, South Korea
6. Lois Orosa, ETH Zurich, January 2019 to February 2022, First job: Managing Director at the Supercomputing Center of Galicia, CESGA
7. Nour Almadhoun Alserr, September 2020 to August 2022.
8. Jawad Haj-Yahya, ETH Zurich, 2019-2021, First job: Principal Researcher at the Huawei Zurich Research Center
9. Rachata Ausavarungnirun, CMU ECE, 2017-2019, First job: Assistant Professor, Sirindhorn International Thai-German Graduate School (TGGS)
10. Yixin Luo, CMU ECE/CS, March 2018 to June 2018. First job: Research Scientist at Google.
11. Arash Tavakkol, ETH Zürich D-INFK, August 2016 to May 2018.
12. Yaohua Wang, ETH Zürich D-INFK, 2017.
13. Kevin Chang, CMU ECE, May 2017 to January 2018. First job: Research Scientist at Facebook.
14. Saugata Ghose, CMU ECE, September 2014 to September 2016. Special Faculty Member from 2016-2021. First job: Assistant Professor at the University of Illinois, Urbana-Champaign.
15. Samira Khan, CMU ECE, September 2012 to September 2015. First job: Assistant Professor at the University of Virginia.

Master's Alumni (incomplete)

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.
10. Ashish Shrestha, CMU ECE, 2016.
11. Philipp Gamper, ETH Zurich, Fall 2017.
12. Vesna Resende Barros, "Spectral analysis of fluorescently labeled amyloid fibrils," ETH D-ITET, November 2018.
13. Peter Resutik, "Implementation of a nanopore DNA base-caller," ETH Zurich, Spring 2019.
14. Nika Mansourighiasi, "Understanding the Opportunities and Bottlenecks in Near-Data Processing Engines," ETH Zurich, Fall 2019.
15. Alain Denzler, "A Near-Cache Accelerator for Stencil Computations," ETH Zurich, Spring 2020.
16. Taha Shahroodi, "Investigating the Opportunities for Improving the Accuracy and Efficiency of Metagenomic Profiling," ETH Zurich, Fall 2020.
17. Jan Schappi, "Reducing Reference Data for Metagenomic Analysis," ETH Zurich, Spring 2021.
18. Kosta Stojiljkovic, "Alternative address translation path for frequently accessed pages," ETH Zurich, Spring 2021.
19. Joel Lindegger, "GenASM-GPU: A Practical and Efficient Implementation of GenASM on GPUs," ETH Zurich, Fall 2021.
20. Emanuele Esposito, "PIM Areas: Memory allocation and alignment for Processing-in-Memory," ETH Zurich, Spring 2021.
21. Roknoddin Azizibarzoki, "Exploring Processing-in-Storage using Inherent Computation Capability of NAND Flash Memory," ETH Zurich, Spring 2022.
22. Julien Eudine, "New Methods and Architectures for Accelerating Read Mapping on Modern CPUs, GPUs, and FPGAs.," ETH Zurich, Fall 2022.
23. João Dinis Ferreira, "Performance and Energy Modelling of a Real-World Processing-in-Memory Architecture," ETH Zurich, Fall 2022.
24. Luca Blum, "Catalyzing deep learning based metagenomic profiling," ETH Zurich, Spring 2023.
25. Alain Kohli, "Framework for Easy Use of the UPMEM Processing-in-Memory Platform using Parallel Patterns," ETH Zurich, Spring 2023.
26. Haocong Luo, "RowPress: Amplifying Read Disturbance in Modern DRAM Chips," ETH Zürich, March 2023.
27. Melina Soysal, "In-Storage and In-Memory Aceleration of Raw Signal Genome Analysis," ETH Zürich, January 2024.
28. Banu Cavlak, "OverlapCall: Improving the Accuracy and Performance of Nanopore Basecalling by Jointly Basecalling Overlapping Signals," ETH Zürich, Spring 2024.
29. Vamanan Arulchelvan, "Prometheus: Automated Framework for Near-Data Processing in Storage Systems," ETH Zürich, Spring 2024.
30. Manos Frouzakis, ETH Zürich, Summer 2024.

Bachelor's Alumni (incomplete)

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.
3. Greg Nazario, CMU ECE (Intern), 2011.
4. Harsha Rastogi, CMU ECE (Intern), Summer 2012.
5. Yuchen Hao, CMU ECE (Intern), Summer 2012.
6. Abhishek Bhowmick, CMU ECE (Intern), 2012-2013.
7. Hitesh Arora, CMU ECE (Intern), 2013.
8. Jeremie Kim, CMU ECE (Intern), 2012-2014.
9. John Emmons, CMU ECE (Intern).
10. Arnab Ghosh, CMU ECE (Intern), Summer 2015.
11. Albert Cho, CMU ECE (Intern), Summer 2015.
12. Raghav Gupta, CMU ECE, 2016.
13. Kais Kudrolli, CMU ECE, 2016.
14. Alexander Breuss, "Sound-based Localization," ETH Zürich D-INFK, Bachelor's Thesis, August 2017. First job: Master's student at ETH Zürich D-INFK.
15. Tobias Büchli, ETH Zürich D-INFK, Bachelor's Thesis, September 2018.
16. Rudolf Loretan, ETH Zürich, Bachelor's Thesis, Summer 2019.
17. Sven Gregorio, "AmbitPlus: a massively parallel DRAM based SIMD accelerator," ETH Zürich, Bachelor's Thesis, Fall 2019.
18. Julian Buchel, "Optimal spike-based signal representation on a neuromorphic chip," ETH Zürich, Bachelor's Thesis, Spring 2020.
19. Roberto Starc, "Enabling Processing-in-Memory for General Purpose Architectures," ETH Zürich, Bachelor's Thesis, Spring 2021.
20. Maurus Item, "Fast Implementation of Transcendental Functions for Processing-in-Memory," ETH Zürich, Bachelor's Thesis, Spring 2022.

21. Max Rumpf, "SequenceLab: A Comprehensive Benchmark of Computational Heuristic Methods for Comparing Genomic Sequences," ETH Zürich, Bachelor's Thesis, Fall 2022.
22. Georgijs Vilums, "Easing PIM Programmability with Parallel Patterns: A Case Study with the UPMEM Architecture," ETH Zürich, Bachelor's Thesis, Fall 2022.
23. Youn Joo Lee, "Accelerating Read Mapping using FPGAs," ETH Zürich, Fall 2022.
24. Axel Scwarzenbach, "Characterizing and Leveraging HBM-PIM," ETH Zürich, Bachelor's Thesis, Spring 2023.
25. Steve Rhyner, "Is ML Training on PiM getting real? Towards an Algorithm/Hardware Co-Design," ETH Zürich, Bachelor's Thesis, Fall 2023.
26. Arvid Gollwitzer, "Heuristic Read Mapping for Metagenomic Analyses," ETH Zürich, Bachelor's Thesis, Fall 2023.
27. Hong Chul Nam, "Victima: Drastically Increasing Address Translation Reach by Leveraging Underutilized Cache Resources," ETH Zürich, Bachelor's Thesis, Summer 2023.
28. Jeril Vathalloor, ETH Zürich, Bachelor's Thesis, Summer 2023.
29. Lukas Zink, "Towards Throughput-oriented Sparse Matrix Vector Multiplication on a Processing-in-Memory System," ETH Zürich, Bachelor's Thesis, Fall 2023.
30. Clement Rousseau, "Loop-aware Value Prediction," ETH Zürich, Bachelor's Thesis, Fall 2023.
31. Jan Mantsch, "Effectiveness of Storage Centric Computing in Accelerating Arithmetic Operations," ETH Zürich, Bachelor's Thesis, Fall 2023.
32. Caroline Hengartner, "Athena: Co-optimizing High-Performance Prefetcher and Off-chip Prediction in Multi-Core Systems," ETH Zürich, Bachelor's Thesis, Spring 2024.
33. Christina Morad, "Understanding the Power Consumption of Processing using Memory (PuM) Operations in Real DRAM Chips," ETH Zürich, Bachelor's Thesis, Spring 2024.
34. Ian Ganz, "Virtuoso: Enabling Fast and Accurate Virtual Memory Research with a Microkernel-based Simulation Methodology," ETH Zürich, Bachelor's Thesis, Spring 2024.
35. Egor Gubenko, "DRAM Characterization Toolkit: A Flexible, Easy-to-use Framework to Streamline RowHammer Experiments," ETH Zürich, Bachelor's Thesis, Spring 2024.
36. Adrian Tonica, "Enabling Massively Parallel Computation in Flash Memory via Lookup Tables," ETH Zürich, Bachelor's Thesis, Summer 2024.
37. Talu Guloglu, ETH Zürich, Bachelor's Thesis, Fall 2024.
38. Maria Makeenkova, ETH Zürich, Bachelor's Thesis, Fall 2024.

Current PhD Students

1. Rahul Bera, ETH Zürich, since September 2019
2. Nisa Bostanci, ETH Zürich, since October 2022
3. Geraldo Francisco Oliveira de Junior, ETH Zürich, since November 2018
4. Andreas Kosmas Kakolyris, ETH Zürich, since April 2025
5. Konstantinos Kanellopoulos, ETH Zürich, since June 2020
6. Haocong Luo, ETH Zürich, since May 2023
7. Nika Mansourighiasi, ETH Zürich, since April 2020
8. Rakesh Nadig, ETH Zürich, since September 2021
9. Ataberk Olgun, ETH Zürich, since January 2022
10. Ismail Emir Yuksel, ETH Zürich, since February 2024

Current Senior & Post-Doctoral Researchers

1. Can Firtina, ETH Zürich, since November 2024.
2. Konstantina Koliogeorgi, ETH Zürich, since April 2024.
3. Yu (Lenny) Liang, ETH Zürich, since September 2023.
4. Mohammad Sadrosadati, ETH Zürich, since May 2021.
5. Abdullah Giray Yaglikci, ETH Zürich, since May 2024.

Past High School Research Students at CMU and ETH (incomplete)

Aaron Cutright (Pittsburgh Allderdice High School), Steven Barash (Pittsburgh Allderdice High School), Nolan Dickey (Pittsburgh Allderdice High School), Steven Yu (Hackley School, Tarrytown, New York)

Past Visiting Researchers at CMU and ETH (incomplete)
Nanditha Rao (IIIT Bangalore), Jongmoo Choi (Dankook Univ.), Hiroyuki Usui (Toshiba), Rui Wang (Beihang), Yusen Li, Yaohua Wang (NUDT), Mohammad-Hashem Haghbayan (Turku), David Novo (LIRMM), Soramichi Akiyama (Tokyo), Behzad Salami (BSC), Chetan Kumar Vudadha (BITS Pilani)

Past Interns (incomplete)
Rishi Agarwal (IIT-Kanpur), Nisa Bostanci (TOBB), Kypros Constantinides (Michigan), Alejandro Cornejo (MIT), Reetuparna Das (Penn State), Parag Dixit (CMU), Eiman Ebrahimi (UT-Austin), Mohammad Fattah (Turku), Ivan Fernandez-Vega (Malaga), Christina Giannoula (NTUA), John Greth (CMU), Boris Grot (UT-Austin), Matthias Grundmann (KIT), Engin Ipek (Cornell), Jamie Jablin (Brown), Jose Joao (UT-Austin), Skanda Koppula (MIT), Yanjing Li (Stanford), Youyou Lu (Tsinghua), Haocong Luo (Shanghai Tech), Gregory Nazario (CMU), Ataberk Olgun (TOBB), Arthur Perais (INRIA), Xuehai Qian (Illinois), Jinglei Ren (Tsinghua), Seyyedmohammad Sadrosadati (Sharif U), Kalliopi Tzimi (Patras), Hui Wang (Beihang), Ying Wang (ICT-CAS), Gulay Yalcin (BSC), Samihan Yedkar (CMU), Xiangyao Yu (Tsinghua), Jishen Zhao (Penn State)

Ph.D. Thesis Committee Member (incomplete)
Manuel le Gallo (ETH), Kevin Loughlin (Michigan), 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), Haoran Qiu (UIUC), 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 Computing Surveys, 2020 - Present.
Associate Editor, ACM Transactions on Architecture and Code Optimization (TACO), Feb. 2010 - 2020.
Associate Editor, ACM Transactions on Storage (TOS), Sep. 2016 - Sep 2019.
Associate Editor, IEEE Transactions on Computers (TC), Jan 2019 - 2021.
Editorial Board Member, IEEE Micro, March 2016 - 2021.
Editorial Board Member, International Journal of Supercomputing Frontiers and Innovations. Nov. 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 Chairpersonships
ACM ISMM Program Co-Chair, ACM SIGPLAN International Symposium on Memory Management, 2025.
DRAMSec Workshop Program Co-Chair, 5th Workshop on DRAM Security, to be held with ISCA, 2025.
DSN Program Co-Chair, 53rd IEEE/IFIP International Conference on Dependable Systems and Networks, 2023.
ACM SIGMICRO Chair, July 2020 - July 2023.
ACM SIGMICRO Executive Committee, Past Chair, July 2023 - Present.
MICRO Steering Committee Co-Chair, October 2017-October 2018.
MICRO Test of Time Award Committee Chair, 2015-2018.
IEEE TCCA Young Computer Architect Award Committee Chair, 2017.
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 Test of Time Award Committee Member, 2014-2021.
HPCA Test of Time Award Committee Member, 2018-2020.
IEEE Computer Society IEEE Fellow Evaluation Committee Member, 2019-2020.
IEEE TCCA Young Computer Architect Award Committee Member, 2016-2018.
MICRO Steering Committee Member, 2013-Present.
ISCA Steering Committee Member, 2018.
ICS Steering Committee Member, 2016-2020.
IEEE TCCA Executive Committee Member, 2017-2018.
IISWC Steering Committee Member, IEEE International Symposium on Workload Characterization, 2008-2018.
IEEE Micro Top Picks Selection Committee Member, IEEE Micro Special Issue on Top Picks from Computer Architecture Conferences, 2023, 2022, 2021, 2020, 2018, 2017, 2015, 2013, 2012, 2011, 2010, 2009 (Did not accept invitation in 2024 due to sabbatical).

MICRO Program Committee Member, International Symposium on Microarchitecture, 2025, 2024, 2023, 2021, 2020, 2019, 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, 2023, 2021, 2020, 2017, 2010, 2008 (Did not accept invitation in 2016, 2015 due to paper submission or acceptance limits; 2019 due to time conflicts).
ISCA Industrial Session Program Committee Member, International Symposium on Computer Architecture, 2025.
ISCA External Review Committee Member, International Symposium on Computer Architecture, 2024, 2022, 2019, 2016, 2015.
ACM CCS Program Committee Member, ACM Conference on Computer and Communications Security, 2025.
RECOMB Program Committee Member, ISCB Conference on Research in Computational Molecular Biology, 2022, 2023, 2025.
USENIX ATC Program Committee Member, USENIX Annual Technical Conference, 2022.
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.
EuroSys Program Committee Member, ACM SIGOPS European Conference on Computer Systems, 2024.
ASPLOS Program Committee Member, Intl. Conf. on Architectural Support for Programming Lang. and Operating Systems, 2026, 2025, 2024, 2022, 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, 2025, 2024, 2018, 2017, 2015, 2012, 2010, 2009.
HPCA Industrial Session Program Committee Member, International Symposium on High Performance Computer Architecture, 2024.
HPCA External Review Committee Member, International Symposium on High Performance Computer Architecture, 2021, 2016.
DSN Program Committee Member, IEEE/IFIP International Conference on Dependable Systems and Networks, 2026, 2025, 2024, 2023, 2022, 2021, 2020, 2019, 2016.
DAC Program Committee Member, Design Automation Conference, 2024, 2023, 2022, 2021, 2020, 2016, 2015, 2014.
DATE Program Committee Member, Design Automation Test in Europe Conference, 2016.
IPDPS Program Committee Member, IEEE International Parallel & Distributed Processing Symposium, 2023.
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, 2025, 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.
ISPASS Program Committee Member, IEEE International Symposium on Performance Analysis of Systems and Software, 2020.
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.
MSST Program Committee Member, International Conference on Massive Storage Systems and Technology, 2020, 2019.
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, 2019, 2018, 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.
ETS Program Committee Member, European Test Symposium, 2025.

Technical Workshop and Tutorial Committee Memberships and Chairmanships

Memory-Centric Computing Tutorial/Workshop Organizer, 2025, 2024. (held with ISCA, HEART, MICRO; to be held with PPoPP, ASPLOS)
Real Processing in Memory Systems (RealPIM) Tutorial Organizer, 2023. (held with HPCA, ASPLOS, ISCA, MICRO)
Workshop on Hardware Acceleration for Bioinformatics (BIO-Arch) Organizer, 2023. (held with RECOMB)
Workshop on DRAM Security, 2025, 2024, 2023, 2022 (held with ISCA)
Workshop on Accelerator Architecture in Computational Biology and Bioinformatics, 2022, 2019, 2018. (held with HPCA)
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, ACM Computing Surveys, Proceedings of the IEEE, Nature Electronics, IEEE Journal of Solid State Circuits.

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, NVMSA.

Expert Reviewer for Research Proposals (incomplete)
Reviewer for National Research Foundations of Various European Countries, 2023, 2022, 2021, 2020, 2019, 2018, 2017.
Swiss National Science Foundation Reviewer, 2017, 2016
European Research Council Reviewer, 2017, 2016
NSF Review Panel Participant, 2007-2015.
DOE Review Panel Participant, 2010.
Reviewer for Research Proposals in Australia, Belgium, European Union, Canada, Norway, Switzerland, 2013, 2014, 2015, 2016, 2017, 2018.
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/D-ITET: Various Faculty Hiring Committees (2018-Present)
ETH D-INFK: Awards Committee Chair (2018-2020)
ETH D-INFK: Doctoral Studies Committee (2016-2020), 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, ASPLOS 2023, ASPLOS 2020, HPCA 2019, 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/@CMUCompArch` for online computer architecture lecture videos at Carnegie Mellon University.
Please visit `https://www.youtube.com/OnurMutluLectures` for online lecture videos for my courses at ETH Zürich.

Major Source Code Releases
Please visit `https://github.com/CMU-SAFARI` for all source code releases that are freely and openly available.
Below are some popular open-source tools and software.

Ramulator, `https://github.com/CMU-SAFARI/ramulator`
Ramulator2, `https://github.com/CMU-SAFARI/ramulator2`
RowHammer, `https://github.com/CMU-SAFARI/rowhammer`
RowPress, `https://github.com/CMU-SAFARI/RowPress`
MQSim, `https://github.com/CMU-SAFARI/MQSim`
SoftMC, `https://github.com/CMU-SAFARI/SoftMC`
DRAM Bender, `https://github.com/CMU-SAFARI/DRAM-Bender`

PrIM Benchmarks, https://github.com/CMU-SAFARI/prim-benchmarks
SparseP, https://github.com/CMU-SAFARI/SparseP
Ramulator-PIM, https://github.com/CMU-SAFARI/ramulator-pim
Pythia, https://github.com/CMU-SAFARI/Pythia
RawHash https://github.com/CMU-SAFARI/RawHash
SneakySnake, https://github.com/CMU-SAFARI/SneakySnake
PiDRAM, https://github.com/CMU-SAFARI/PiDRAM
Hermes, https://github.com/CMU-SAFARI/Hermes
DAMOV, https://github.com/CMU-SAFARI/DAMOV
BioDynamo, https://biodynamo.org/
MetaSys, https://github.com/CMU-SAFARI/MetaSys
Scrooge, https://github.com/CMU-SAFARI/Scrooge
IMPICA, https://github.com/CMU-SAFARI/IMPICA
GPGPUSim-Ramulator, https://github.com/CMU-SAFARI/GPGPUSim-Ramulator
Mosaic, https://github.com/CMU-SAFARI/Mosaic
VAMPIRE (DRAM Power Model), https://github.com/CMU-SAFARI/VAMPIRE
BLEND, https://github.com/CMU-SAFARI/BLEND
GenASM, https://github.com/CMU-SAFARI/GenASM

Funding Summary

- Raised more than 8 Million USD funding while at CMU from various sources including NSF, NIH, SRC/GSRC, industry donations.
- Raised more than 12.5 Million CHF funding while at ETH from various sources including SNF, EU Horizon, SRC, industry donations and projects.
- More than 3 Million USD raised in gift/donation funding at CMU from major companies including Nvidia, Microsoft, Intel, AMD, Google, Samsung, Oracle, HP, IBM, Facebook, Huawei, Qualcomm, Rambus, Seagate, VMware
- More than 5 Million USD raised in gift/donation funding at ETH from major companies including Google, Microsoft, VMware, ASML, Intel, Xilinx, Huawei, Alibaba
- NIH R01 grant as PI while at CMU (USD 1.475M): *Novel algorithms and hardware designs for ultra-fast nextgen sequence analysis*
- NSF grants as single PI while at CMU: NSF CAREER Award (USD 715K), NSF CSR Small (USD 516K), NSF CCF EAGER (USD 116K):
  CAREER: QoS-Aware, High-Performance, and Scalable Many-Core Memory Systems
  CSR: Small: High-Performance and Energy-Efficient Single-Level Stores: Efficient Coordinated Management of Storage and Memory
  EAGER: Collaborative Research: Heterogeneous Cores, Memory-Hierarchy and Communication Architectures for Future CMPs
- Collaborative NSF grants as PI while at CMU: 1 Large (USD 368K), 1 Medium (USD 456K):
  SHF: Large:Collaborative Research: Architecting the Next Generation Memory Hierarchy - A Holistic Approach
  CSR: Medium: Collaborative Research: Enabling GPUs as First-Class Computing Engines
- GSRC and SRC grants as PI while at CMU (totaling USD 570K) *GSRC: Scalable and QoS-Aware Memory Systems*
  SRC: Integrated Techniques for Scalable Management of Main Memory Performance, Energy and QoS
- NSF grants as co-PI while at CMU: 1 Medium (USD 380K), 1 Small (USD 248K):
  TC: Medium: Exploiting Multicore and Hardware Acceleration to Perform Efficient Behavior-Based Attack Detection and Repair
  CSR: Small: Effective Data Compression for Modern Memory Systems
- SNF (Swiss National Science Foundation) grant as single PI while at ETH (CHF 680K):
  Near-Data-Processing Architectures and Algorithms for Metagenomic Analysis
- EU Horizon collaborative grant as PI while at ETH (BioPIM) (CHF 500K):
  BioPIM: Processing-in-memory architectures and programming libraries for bioinformatics algorithms
- SRC grants as single PI while at ETH (totaling USD 726K):
  Architectural Frameworks to Facilitate Practical Efficient and Transparent Computation Near Data
  Memory System Design for AI/ML Accelerators & ML/AI Techniques for Memory System Design
  Memory System Design for AI/ML & ML/AI for Memory System Design
- HKUST ACCESS Center grant as single PI while at ETH (totaling USD 720K):
  Scalable Memory-Centric Heterogeneous System Design for Machine Learning

Professional Memberships

Elected Member of the Academy of Europe (Academia Europaea), 2018, elected "for outstanding achievements as a researcher"
Fellow of ACM, 2017, elected "for contributions to computer architecture research, especially in memory systems"
Fellow of IEEE, 2018, elected "for contributions to computer architecture research and practice"
Member of AAAS
Member of ACM SIGARCH, ACM SIGMICRO, IEEE Computer Society, IEEE TCCA
Member of Eta Kappa Nu Electrical Engineering Honor Society and Tau Beta Pi Engineering Honor Society, since 1998