Jishen Zhao

jzhao@ucsd.edu

Associate Professor, Computer Science and Engineering
University of California, San Diego

http://cseweb.ucsd.edu/~jzhao

Education

2009 - 2014 Ph.D. in Computer Science and Engineering, Pennsylvania State University
2003 - 2008 M.Eng. in Electrical Engineering, Zhejiang University
1999 - 2003 B.Eng. in Electrical Engineering, Zhejiang University

Research Interests

A broad range of computer architecture and systems research that bridges system software and hardware design, with an emphasis on memory and storage systems, machine learning and system co-design, and reliability.

Professional Experience

July 2020 - present **University of California, San Diego**
Associate Professor in Computer Science and Engineering Department

Nov. 2017 - July 2020 **University of California, San Diego**
Assistant Professor in Computer Science and Engineering Department

Nov. 2017 - Dec. 2018 **University of California, Santa Cruz**
Visiting Assistant Professor in Computer Engineering Department

Jan. 2015 - Nov. 2017 **University of California, Santa Cruz**
Assistant Professor in Computer Engineering Department

Jan. 2014 - Jan. 2015 **Hewlett-Packard Labs**
Research Scientist at Systems Research Lab

Feb. 2012 - Dec. 2013 **Hewlett-Packard Labs**
Research Associate at Intelligent Infrastructure Lab

Honors & Awards

2021 One paper (of 12 total) selected for IEEE Micro Top Picks from Computer Architecture Conferences
Characterizing and Modeling Non-Volatile Memory Systems
2020 Best Paper Award at ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)
2020 Best Paper in Track Award at IEEE International Conference on Computer Design (ICCD)
2020 ACSIC (American Chinese Scholar In Computing) Rising Star Award
2019 MICRO Hall of Fame Award
2017 NSF CAREER Award
2013 Best Paper Honorable Mention Award at MICRO

Books

[1] Yuan Xie and Jishen Zhao. Die-stacking architecture. *Synthesis Lectures on Computer Architecture. Morgan & Claypool Publishers*, 2015.

Publications

[1] Zixuan Wang, Xiao Liu, Jian Yang, Theodore Michailidis, Steven Swanson, and Jishen Zhao. Characterizing and Modeling Non-Volatile Memory Systems. In *the Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2020. Acceptance Rate: 82/422=19.4%, **IEEE Micro top picks from computer architecture conferences, 2021**.

[2] Juno Kim, Yun Joon Soh, Joseph Izraelevitz, Jishen Zhao, and Steven Swanson. SubZero: Zero-copy IO for Persistent Main Memory FileSystems. In *the Proceedings of the 11th ACM SIGOPS Asia-Pacific Workshop on Systems (APSys)*, 2020. **Best Paper Award**.

[3] Cheng Fu, Huili Chen, Zhenheng Yang, Yuandong Tian Farinaz Koushanfar and, and Jishen Zhao. Enhancing Model Parallelism in Neural Architecture Search for Multi-device System. In *IEEE Micro Special Issue on Machine Learning for Systems*, 2020.

[4] Saransh Gupta, Mohsen Imani, Hengyu Zhao, Fan Wu, Jishen Zhao, and Tajana Rosing. Implementing Binary Neural Networks in Memory with Approximate Accumulation. In *the Proceedings of ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, 2020.

[5] Sabur Baidya, Yu-Jen Ku, Hengyu Zhao, Jishen Zhao, and Sujit Dey. Vehicular and Edge Computing for Emerging Connected and Autonomous Vehicle Applications. In *the Proceedings of the 57th ACM/IEEE Design Automation Conference (DAC)*, 2020. Invited paper.

[6] Hengyu Zhao, Yubo Zhang, Pingfan Meng, Hui Shi, Erran Li, Tiancheng Lou, and Jishen Zhao. Safety Score: A Quantitative Approach to Guiding Safety-Aware Autonomous Vehicle Computing System Design. In *the Proceedings of IEEE Intelligent Vehicles Symposium (IV)*, 2020.

[7] Hui Shi, Yang Zhang, Xinyun Chen, Yuandong Tian, and Jishen Zhao. Deep Symbolic Superoptimization Without Human Knowledge. In *the Proceedings of International Conference on Learning Representations (ICLR)*, 2020. Acceptance Rate: 687/2594=26.5%.

[8] Hengyu Zhao, Yubo Zhang, Pingfan Meng, Hui Shi, Li Erran Li, Tiancheng Lou, and Jishen Zhao. Driving Scenario Perception-Aware Computing System Design in Autonomous Vehicles. In *the Proceedings of the 38th International Conference on Computer Design (ICCD)*, 2020. Acceptance rate: 62/221=28.1%, **Best Paper in Track**.

[9] Cheng Fu, Huili Chen, Haolan Liu, Xinyun Chen, Yuandong Tian, Farinaz Koushanfar, and Jishen Zhao. Coda: An End-to-End Neural Program Decompiler. In *the Proceedings of Conference on Neural Information Processing System (NeurIPS)*, 2019. Acceptance Rate: 1428/6743=21.2%.

[10] Huili Chen, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. GenUnlock: An Automated Genetic Algorithm Framework for Unlocking Logic Encryption. In *International Conference on Computer Aided Design (ICCAD)*, 2019.

[11] Xiao Liu, Bhaskar Jupudi, Pankaj Mehra, and Jishen Zhao. Persistent Memory Workload Characterization: A Hardware Perspective. In *the Proceedings of IEEE International Symposium on Workload Characterization (IISWC)*, 2019.

[12] Xiao Liu, David Roberts, Rachata Ausavarungnirun, Onur Mutlu, and Jishen Zhao. Binary Star: Coordinated Reliability in Heterogeneous Memory Systems for High Performance and Scalability. In *the Proceedings of International Symposium on Microarchitecture (MICRO)*, 2019. Acceptance Rate: 79/344=23.0%.

[13] Yuanjiang Ni, Jishen Zhao, Heiner Litz, Daniel Bittman, and Ethan Miller. SSP: Eliminating Redundant Writes in Failure-Atomic NVRAMs via Shadow Sub-Paging. In *the Proceedings of International Symposium on Microarchitecture (MICRO)*, 2019. Acceptance Rate: 79/344=23.0%.

[14] Brandon Nesterenko, Xiao Liu, Qing Yi, Jishen Zhao, and Jiange Zhang. Transitioning scientific applications to using non-volatile memory for resilience. In *Proceedings of International Symposium on Memory Systems (MEMSYS)*, 2019.

[15] Huili Chen, Bita Darvish Rouhani, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. Deepmarks: A secure fingerprinting framework for digital rights management of deep learning models. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval (ICMR)*, pages 105–113, 2019.

[16] Huili Chen, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. DeepInspect: An Automated Trojan Detection Framework for Neural Networks. In *the Proceedings of International Joint Conference on Artificial Intelligence (IJCAI)*, 2019. Acceptance Rate: 850/4752=17.8%.

[17] Xiao Liu, Minxuan Zhou, Tajana Rosing, and Jishen Zhao. HR3AM: a Heat Resilient Design for RRAM based Neuromorphic Computing. In *the Proceedings of ACM/IEEE International Symposium on Low Power Electronics and Design (ISLPED)*, 2019.

[18] Kunal Korgaonkar, Joe Izraelevitz, Jishen Zhao, and Steven Swanson. Vorpal: Vector Clock-Inspired Ordering for Large Persistent Memory Systems. In *the Proceedings of the 38th ACM Symposium on Principles of Distributed Computing (PODC)*, 2019.

[19] Huili Chen, Cheng Fu, Jishen Zhao, and Farinaz Koushanfar. DeepAttest: An End-to-End Attestation Framework for Deep Neural Networks. In *the Proceedings of the 46th International Symposium on Computer Architecture (ISCA)*, 2019. Acceptance Rate: 62/365=17.0%.

[20] Shengye Wang, Xiao Liu, Jishen Zhao, and Henrik Iskov Christensen. Rorg: Service robot software management with linux containers. In *the Proceedings of International Conference on Robotics and Automation (ICRA), 2019*.

[21] Matheus A. Ogleari, Ye Yu, Chen Qian, Ethan L. Miller, and Jishen Zhao. String Figure: A scalable and elastic memory network architecture. In *the Proceedings of International Symposium on High-Performance Computer Architecture (HPCA), 2019*. Acceptance rate: 46/233=19.7%.

[22] Sihang Liu, Yizhou Wei, Jishen Zhao, Aasheesh Kolli, and Samira Khan. Pmtest: A fast and flexible testing framework for persistent memory programs. In *the Proceedings of International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2019*. Acceptance rate: 74/350=21.1%.

[23] Hengyu Zhao*, Jiawen Liu*, Matheus A. Ogleari, Dong Li, and Jishen Zhao. Processing-in-Memory for Energy-efficient Neural Network Training: A Heterogeneous Approach. In *the Proceedings of International Symposium on Microarchitecture (MICRO)*, 2018. (The first two authors contribute equally) Acceptance rate: 74/348=21.3%.

[24] Xing Hu, Matheus A. Ogleari, Jishen Zhao, Shuangchen Li, Abanti Basak, and Yuan Xie. Persistence Parallelism Optimization: A holistic approach from memory bus to RDMA network. In *the Proceedings of International Symposium on Microarchitecture (MICRO)*, 2018. Acceptance rate: 74/348=21.3%.

[25] Hengyu Zhao and Jishen Zhao. Leveraging MLC STT-RAM for energy-efficient CNN training. In *Proceedings of International Symposium on Memory Systems (MEMSYS)*, 2018.

[26] Shouyi Yin, Shibin Tang, Xinhan Lin, Peng Ouyang, Fengbin Tu, Leibo Liu, Jishen Zhao, Cong Xu, Shuangchen Li, Yuan Xie, and Shaojun Wei. Parana: A parallel neural architecture considering thermal problem of 3d stacked memory. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, 2018.

[27] Yuanjiang Ni, Jishen Zhao, Daniel Bittman, and Ethan Miller. Reducing NVM writes with optimized shadow paging. In *the 10th USENIX Workshop on Hot Topics in Storage and File Systems (HotStorage), 2018*.

[28] Guohao Dai, Tianhao Huang, Yuze Chi, Jishen Zhao, Guangyu Sun, Yongpan Liu, Yu Wang, Yuan Xie, and Huazhong Yang. GraphH: A processing-in-memory architecture for large-scale graph processing. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 2018.

[29] Matheus Ogleari, Ethan Miller, and Jishen Zhao. Steal but no force: Efficient hardware undo+redo logging for persistent memory systems. In *the Proceedings of International Symposium on High-Performance Computer Architecture (HPCA), 2018*. Acceptance rate: 54/260=20.8%.

[30] Mengjie Li, Matheus A. Ogleari, and Jishen Zhao. Logging in persistent memory: to cache, or not to cache? In *Proceedings of the International Symposium on Memory Systems*, 2017.

[31] Hengyu Zhao*, Linuo Xue*, Ping Chi, and Jishen Zhao. Approximate image storage with multi-level cell stt-mram main memory. In *Proceedings of the International Conference On Computer Aided Design (ICCAD)*, 2017. (The first two authors contribute equally) Acceptance rate: 105/399=26.3%.

[32] Liang Chang, Zhaohao Wang, Alvin Oliver Glova, Jishen Zhao, Youguang Zhang, Yuan Xie, and Weisheng Zhao. PRESCOTT: Preset-based cross-point architecture for spin-orbit-torque magnetic random access memory. In *Proceedings of the International Conference On Computer Aided Design (ICCAD)*, 2017. Acceptance rate: 105/399=26.3%.

[33] Chun-Hao Lai, Jishen Zhao, and Chia-Lin Yang. Leave the cache hierarchy operation as it is: A new persistent memory accelerating approach. In *Proceedings of of the 54th Design Automation Conference (DAC)*, 2017. Acceptance Rate: 161/676=23.8%.

[34] Jia Zhan, Itir Akgun, Jishen Zhao, Al Davis, Paolo Faraboschi, Yuangang Wang, and Yuan Xie. A unified memory network architecture for in-memory computing in commodity servers. In *Proceedings of the International Symposium on Microarchitecture (MICRO)*, 2016. Acceptance Rate: 61/283=21.6%.

[35] Yuxiong Zhu, Borui Wang, Dong Li, and Jishen Zhao. Integrated thermal analysis for processing in die-stacking memory. In *Proceedings of International Symposium on Memory Systems (MEMSYS)*, 2016.

[36] Borui Wang, Martin Torres, Dong Li, Jishen Zhao, and Florin Rusu. Performance implications of processing-in-memory designs on data-intensive applications. In *5th Workshop on Heterogeneous and Unconventional Cluster Architectures and Applications*, 2016.

[37] Ping Chi, Shuangchen Li, Cong Xu, Tao Zhang, Jishen Zhao, Yongpan Liu, Yu Wang, and Yuan Xie. A novel processing-in-memory architecture for neural network computation in reram-based main memory. In *Proceedings of the 43rd International Symposium on Computer Architecture (ISCA)*, 2016. Acceptance Rate: 54/288=18.8%.

[38] Hsiang-Yun Cheng, Jishen Zhao, Jack Sampson, Mary Jane Irwin, Aamer Jaleel, Yu Lu, and Yuan Xie. Lap: Loop-block aware inclusion properties for energy-efficient asymmetric last level caches. In *Proceedings of the 43rd International Symposium on Computer Architecture (ISCA)*, 2016. Acceptance Rate: 54/288=18.8%.

[39] Xiao Liu, Qing Yi, and Jishen Zhao. Using memory-style storage to support fault tolerance in data centers. In *USENIX Workshop on Cool Topics in Sustainable Data Centers (CoolDC)*, 2016.

[40] Shuangchen Li, Cong Xu, Jishen Zhao, Yu Lu, and Yuan Xie. Pinatubo: A processing in non-volatile memory architecture for bulk bitwise operations. In *Proceedings of the 53rd Design Automation Conference (DAC)*, 2016. Acceptance Rate: 152/876=17.4%.

[41] Jishen Zhao, Cong Xu, Tao Zhang, and Yuan Xie. BACH: A bandwidth-aware hybrid cache hierarchy design with nonvolatile memories. *Journal of Computer Science and Technology*, 31:20–35, 2016.

[42] Jinglei Ren, Jishen Zhao, Samira Khan, Jongmoo Choi, Yongwei Wu, and Onur Mutlu. ThyNVM: Enabling software-transparent crash consistency in persistent memory systems. In *International Symposium on Microarchitecture (MICRO)*, 2015. Acceptance Rate: 61/283=21.5%.

[43] Jishen Zhao, Qiaosha Zou, and Yuan Xie. Overview of 3d architecture design opportunities and techniques. *IEEE Design & Test*, 2015.

[44] Jishen Zhao, Sheng Li, Jichuan Chang, John L. Byrne, Laura L. Ramirez, Kevin Lim, Yuan Xie, and Paolo Faraboschi. Buri: Scaling big memory computing with transparent memory expansion. *ACM Transactions on Architecture and Code Optimization (TACO)*, 2015.

[45] Shuangchen Li, Ping Chi, Jishen Zhao, Kwang-Ting Cheng, and Yuan Xie. Leveraging nonvolatility for architecture design with emerging nvm. In *Proceedings of the 4th IEEE Non-Volatile Memory System and Applications Symposium (NVMSA)*, 2015.

[46] Hsiang-Yun Cheng, Jia Zhan, Jishen Zhao, Yuan Xie, Jack Sampson, and Mary Jane Irwin. Core vs. Uncore: The heart of darkness. In *Proceedings of the 52nd Design Automation Conference (DAC)*, 2015. Invited paper.

[47] Ke Chen, Sheng Li, Jung Ho Ahn, Naveen Muralimanohar, Jishen Zhao, Cong Xu, Seongil O, Yuan Xie, Jay B. Brockman, and Norman P. Jouppi. History-assisted adaptive-granularity caches (HAAG$) for high performance 3D DRAM architectures. In *Proceedings of the 29th International Conference on Supercomputing (ICS)*, 2015. Acceptance rate: 40/160=25%.

[48] Jia Zhan, Jin Ouyang, Fen Ge, Jishen Zhao, and Yuan Xie. DimNoC: A dim silicon approach towards power-efficient on-chip network. In *Proceedings of the 52nd Design Automation Conference (DAC)*, 2015. Acceptance Rate: 162/789=20.5%.

[49] Jishen Zhao, Cong Xu, Ping Chi, and Yuan Xie. Memory and storage system design with nonvolatile memory technologies. *IPSJ Transactions on System LSI Design Methodology (TSLDM)*, 8:2–11, 2015. Invited paper.

[50] Jishen Zhao, Onur Mutlu, and Yuan Xie. FIRM: Fair and high-performance memory control for persistent memory systems. In *Proceedings of the 47th International Symposium on Microarchitecture (MICRO-47)*, 2014. Acceptance rate: 53/279=19%.

[51] Jishen Zhao, Sheng Li, Doe Hyun Yoon, Yuan Xie, and Norman P. Jouppi. Kiln: Closing the performance gap between systems with and without persistence support. In *Proceedings of the 46th International Symposium on Microarchitecture (MICRO)*, pages 421–432, 2013. Acceptance rate: 39/239=16%, **Best Paper Honorable Mention**.

[52] Jishen Zhao, Guangyu Sun, Gabriel Loh, and Yuan Xie. Optimizing GPU energy efficiency with 3D die-stacking graphics memory and reconfigurable memory interface. *ACM Transactions on Architecture and Code Optimization (TACO)*, 10(4):24:1–24:25, 2013.

[53] Justin Meza, Yixin Luo, Samira Khan, Jishen Zhao, Yuan Xie, and Onur Mutlu. A case for efficient hardware/software cooperative management of storage and memory. In *Proceedings of the 5th Workshop on Energy-Efficient Design (WEED), held in conjunction with the International Symposium on Computer Architecture (ISCA-40)*, 2013.

[54] Sheng Li, Doe Hyun Yoon, Ke Chen, Jishen Zhao, Jung Ho Ahn, Jay B. Brockman, Yuan Xie, and Norman P. Jouppi. Mage: Adaptive granularity and ecc for resilient and power efficient memory systems. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, pages 33:1–33:11, 2012. Acceptance rate: 100/472=21%.

[55] Jishen Zhao and Yuan Xie. Optimizing bandwidth and power of graphics memory with hybrid memory technologies and adaptive data migration. In *Proceedings of the International Conference on Computer-Aided Design (ICCAD)*, pages 81–87, 2012. Acceptance rate: 82/338=24%.

[56] Jishen Zhao, Guangyu Sun, Gabriel H. Loh, and Yuan Xie. Energy-efficient GPU design with reconfigurable in-package graphics memory. In *Proceedings of the 18th International Symposium on Low Power Electronics and Design (ISLPED)*, pages 403–408, 2012. Acceptance rate: 34/213=16%.

[57] Jishen Zhao, Cong Xu, and Yuan Xie. Bandwidth-aware reconfigurable cache design with hybrid memory technologies. In *Proceedings of the International Conference on Computer-Aided Design (ICCAD)*, pages 48–55, 2011. Acceptance rate: 106/349=30%.

[58] Guangyu Sun, Christopher J. Hughes, Changkyu Kim, Jishen Zhao, Cong Xu, Yuan Xie, and Yen-Kuang Chen. Moguls: A model to explore the memory hierarchy for bandwidth improvements. In *Proceedings of the 38th International Symposium on Computer Architecture (ISCA)*, pages 377–388, 2011. Acceptance rate: 40/208=19%.

[59] Jishen Zhao, Xiangyu Dong, and Yuan Xie. An energy-efficient 3D CMP design with fine-grained voltage scaling. In *Proceedings of the IEEE/ACM Design, Automation, and Test in Europe Conference (DATE)*, pages 539–542, 2011.

[60] Xiangyu Dong, Jishen Zhao, and Yuan Xie. Fabrication cost analysis and cost-aware design space exploration for 3D ICs. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*, 29(12):1959–1972, December 2010.

[61] Jing Xie, Jishen Zhao, and Yuan Xie. Architectural benefits and design challenges for three-dimensional integrated circuits. In *Proceedings of the IEEE 11th biennial Asia Pacific Conference on Circuits and Systems (APCCAS)*, pages 540–543, 2010.

[62] Yibo Chen, Jishen Zhao, and Yuan Xie. 3D-nonFAR: Three-dimensional non-volatile FPGA architecture using phase change memory. In *Proceedings of the 16th International Symposium on Low Power Electronics and Design (ISLPED)*, pages 55–60, 2010. Acceptance rate: 52/210=25%.

[63] Jishen Zhao, Xiangyu Dong, and Yuan Xie. Cost-aware three-dimensional (3D) many-core multiprocessor design. In *Proceedings of the 47th Design Automation Conference (DAC)*, pages 126–131, 2010. Acceptance rate: 148/607=24%.

Patents

[P1] Rajeev Balasubramonian, Gregg B. Lesartre, Robert Schreiber, Jishen Zhao, Naveen Muralimanohar, Paolo Faraboschi. Memory device write based on mapping. *US20170220257A1*

[P2] Douglas L Voigt, Charles B Morrey, Jishen Zhao, Dhruva Chakrabarti, and Joseph E Foster. Persistent memory versioning and merging. *WO2016160035*

[P3] Stanko Novakovic, Paolo Faraboschi, Kimberly Keeton, Jishen Zhao, Robert Schreiber. Responsive server identification among multiple data servers linked to a shared memory. *WO2016137496.*

[P4] Rajeev Balasubramonian, Naveen Muralimanohar, Gregg B. Lesartre, Paolo Faraboschi, Jishen Zhao. Data write to subset of memory devices. *WO2016144291.*

[P5] Sheng Li, Jishen Zhao, Kevin T Lim, and Paolo Faraboschi. Node-based computing devices with virtual circuits. *WO2016014043.*

[P6] Sheng Li, Paolo Faraboschi, Kevin T Lim, Jishen Zhao. Node-based computing devices with protocol-based priority. *WO2016014044.*

[P7] Sheng Li, Jichuan Chang, and Jishen Zhao. Dynamic memory expansion by data compression. *WO2015142341*.

[P8] Sheng Li, Jishen Zhao, Jichuan Chang, Parthasarathy Ranganathan, Alistair Veitch, Kevin T. Lim, and Mark Lillibridge. Atomically Committing Write Requests *US Patent 20160342351*

[P9] Sheng Li, Jishen Zhao, Jichuan Chang, Mapping virtual memory pages to physical memory pages. *US Patent 20160267015*.

[P10] Doe Hyun Yoon, Sheng Li, Jishen Zhao, and Norman P. Jouppi. Multiversioned nonvolatile memory hierarchy for persistent memory. *US Patent 20160034225*.

Invited Presentations, Seminars, and Others

[1] Scalable, Energy-efficient, and Flexible Memory System Design with Emerging Technologies, *SK Hynix Inc., Jul. 2019*.

[2] Towards Safety-aware Computing System Design in Autonomous Vehicles, *Mini-Keynote at International Forum on MPSoC for Software-defined Hardware, Kanagawa, Japan, Jul. 2019*.

[3] System and Architecture Design for Safe and Reliable Autonomous Robotic Applications, *UCLA, Jun. 2019*.

[4] System and Architecture Design for Safe and Reliable Autonomous Robotic Applications, *Princeton University, 2019*.

[5] System and Architecture Design for Safe and Reliable Autonomous Robotic Applications, *University of Southern California, 2019*.

[6] System and Architecture Design for Safe and Reliable Autonomous Robotic Applications, *University of Wisconsin–Madison, 2019*.

[7] System and Architecture Design for Safe and Reliable Autonomous Robotic Applications, *University of Michigan, 2019*.

[8] System and Architecture Design for Safe and Reliable Autonomous Robotic Applications, *Georgia Institute of Technology, 2019*.

[9] Software and Hardware Co-design for Scalable and Energy-efficient Neural Network Training with Processing-in-Memory, *H.C. Torng Lecture Series in Computer Engineering, Cornell University, Nov. 2018*.

[10] Software and Hardware Co-design for Scalable and Energy-efficient Neural Network Training with Processing-in-Memory, *Hardware and Algorithms for Learning On-a-Chip (HALO), Nov. 2018*.

[11] Toward Reliable and Cost-efficient IoT Systems, *NSF Workshop on Internet-of-Things (IoT) Systems, Nov. 2018*.

[12] Software and Hardware Co-design for Scalable and Energy-efficient Neural Network Training with Processing-in-Memory, *University of Pittsburgh, Oct. 2018*.

[13] Unlocking the Full Potential of Persistent Memory Technique with Software/Hardware Coordinated Design, *Carnegie Mellon University, Oct. 2018*.

[14] Heterogeneous Processing-in-memory for Energy-efficient CNN training, *[hMini-Keynote at International Forum on MPSoC for Software-defined Hardware, Snowbird, UT, USA, Jul. 2018.*

[15] Scalable, Energy-efficient, and Flexible Memory System Design with Emerging Technologies, Western Digital Corporation, Jun., 2018.

[16] Transaction Cache: A Persistent Memory Acceleration Approach, Mini-Keynote at International Forum on MPSoC for Software-defined Hardware, Annecy, France, July 2017.

[17] Towards unlocking the full potential of persistent memory: A hardware-driven approach, University of California, San Diego, April, 2017.

[18] Workload Characterization and Hardware Support for Persistence, Ulsan National Institute of Science and Technology, Ulsan, Korea, Jul. 2016.

[19] Workload Characterization and Hardware Support for Persistence, International Forum on MPSoC for Software-defined Hardware, Nara, Japan, Jul. 2016.

[20] Persistent memory architecture research at UCSC, SK Hynix, San Jose, California, June 2016.

[21] Re-architecting the Memory-Storage Stack with NVRAMs, University of California, Merced, Aug. 2015.

[22] Memory persistence: a new dimension in memory system design, SanDisk TechCon'15, Milpitas, California, Jul. 2015.

[23] Memory Persistence: A New Dimension in Memory System Design, International Forum on MPSoC for Software-defined Hardware, Santa Barbara, California, Jul. 2015.

[24] Rethinking Memory System Organization with Emerging Technologies, Northwestern University, Mar. 2014.

[25] Rethinking Memory System Organization with Emerging Technologies, University of California Santa Cruz, Feb. 2014.

[26] Rethinking Memory System Organization with NVRAMs, HP Labs, Sep. 2013.

[27] Energy-efficient Reconfigurable Memory Hierarchy Design with Emerging Technologies, ACM Student Research Competition at ICCAD, 2013

[28] Bandwidth-Aware Reconfigurable Cache Hierarchy Design with Hybrid Memory Technologies, First CRA-W/CDC Workshop on Diversity in Design Automation and Test (WD2AT), May 2011.

[29] Reconfigurable Energy-Efficient 3D Stacked Chip-Multiprocessor Design, SRC TECHCON, Sep. 2010.

Professional Services and Activities

Technical Program Chair, Track Chair/Co-chair:

1. Program Co-chair, IEEE/ACM International Symposium on Microarchitecture (MICRO), 2021.
2. Backup Program Chair, IEEE Micro Top Picks from Compute Architecture Conferences, 2020.
3. Backup Track Program Chair, ACM/IEEE Design Automation Conference (DAC), 2019.
4. Backup Program Chair, International Symposium on Computer Architecture (ISCA), 2019.
5. Program Co-chair, International Forum on MPSoC for Software-defined Hardware, 2018.
6. Track Program Co-chair, Architecture Track, IEEE International Conference on Networking, Architecture, and Storage (NAS), 2016

Conference/Workshop Organizer:

1. Workshop and Tutorial Chair, International Symposium on High-Performance Computer Architecture (HPCA), 2020.
2. Social Media Chair, International Symposium on Computer Architecture (ISCA), 2020.
3. Student Travel Grant Chair for the IEEE International Symposium on Workload Characterization (IISWC) 2019.
4. Steering Committee Member, Persistent Programming In Real Life (PIRL), 2019.
5. Publicity Co-chair, ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2019
6. Program co-chair, International Forum on MPSoC for Software-defined Hardware (MPSOC), 2018
7. Web Co-chair, IEEE/ACM International Symposium on Microarchitecture (MICRO), 2018.
8. Program Co-chair, Non-Volatile Memories Workshop (NVMW), 2018
9. Publicity Co-chair, International Symposium on High-Performance Computer Architecture (HPCA), 2018
10. Publicity Co-chair, International Workshop of Software-Defined Data Communications and Storage (SDDCS), Co-located with ASPLOS 2017
11. DAC'16 special session organizer "Winning the Memory Challenge: Which Non-volatile Memory Technology Will Rise Above", June, 2016
12. Finance Co-chair, International Forum on MPSoC for Software-defined Hardware (MPSOC), 2015

Program Committee Member in Conferences:

1. International Symposium on Computer Architecture (ISCA), 2019, 2020
2. International Symposium on High-Performance Computer Architecture (HPCA), 2018 - 2020
3. IEEE/ACM International Symposium on Microarchitecture (MICRO), 2018 - 2020
4. International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2021
5. IEEE Micro Top Picks in Compute Architecture, 2020, 2021
6. ACM/IEEE Design Automation Conference (DAC), 2016 - 2019
7. IEEE/ACM International Conference on Computer Aided Design (ICCAD), 2018
8. International Conference on Parallel Processing (ICPP), 2020
9. USENIX Conference on File and Storage Technologies (FAST), 2020
10. International Symposium on Low Power Electronics and Design (ISLPED), 2015 - 2018
11. International Symposium on Memory Systems (MEMSYS), 2017, 2018
12. IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), 2018

13. IEEE International Symposium on Workload Characterization (IISWC), 2017, 2018, 2020
14. Architecture track, IEEE International Parallel & Distributed Processing Symposium (IPDPS), 2017
15. International Conference on Massive Storage Systems and Technology (MSST), 2016, 2017, 2020
16. ACM International Conference on Computing Frontiers (CF), 2015 - 2017
17. International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS), 2015 - 2017
18. IEEE Non-Volatile Memory System and Applications Symposium (NVMSA), 2015 - 2017
19. Great Lakes Symposium on VLSI (GLSVLSI), 2014 - 2017
20. International Conference on Networking, Architecture, and Storage (NAS), Storage track, 2017
21. International Conference for High Performance Computing, Networking, Storage and Analysis (SC), 2016
22. International Conference on Super Computing (ICS), 2016
23. IEEE International System-on-Chip Conference (SOCC), 2015
24. ESWeek Internet-of-Things Symposium (IoT), 2015
25. IFIP/IEEE International Conference on Very Large Scale Integration (VLSI-SOC), 2015 - 2016
26. Workshop on Architectures and Systems for Big Data (ASBD), 2014
27. Symposium on Computer Architecture and High Performance Computing (SBAC-PAD), 2014

Journal Editor / Guest Editor

1. Associate editor, ACM Transactions on Design Automation of Electronic Systems, 2020 -
2. Guest editor, special issue on emerging memory technologies, IEEE Micro, 2019
3. Guest editor, special issue on System-level Synthesis, IEEE Access, 2016

Proposal Panelist/Reviewer

1. Panelist for NSF CCF program, 2015, 2017, 2018, 2019
2. DOE proposal reviewer, 2019

Conference/Workshop Session Chair

1. Session "In-Memory Computing" at IEEE International Symposium on High-Performance Computer Architecture (HPCA), Vienna, Austria, February 2018.
2. Session "Next Generation System Level Technologies" at International Conference On Computer Aided Design (ICCAD), Irvine, California, November 2017.
3. Session "Managing Energy in Wearable Devices" at International Symposium on Low Power Electronics and Design (ISLPED), San Francisco, California, August 2016.
4. Session "On-the-Go Storage" at International Conference on Massive Storage Systems and Technology (MSST), Santa Clara, California, May 2016.
5. Research session "Memory that Never Forgets" at Design Automation Conference (DAC), Austin, Texas, June 2016

Reviewer/External Reviewer

1. International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2019
2. ACM Transactions on Architecture and Code Optimization (TACO), 2015 - present
3. IEEE Transactions on Very Large Scale Integration Systems (TVLSI), 2015 - present
4. Journal of Low Power Electronics and Applications (JLPEA), 2015 - present
5. Elsevier Parallel Computing (PARCO), 2015 - present
6. IEEE Transactions on Computers (TC), 2014 - present
7. ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2019
8. IEEE/ACM International Symposium on Microarchitecture (MICRO), 2017
9. IEEE International Symposium on Circuits and Systems (ISCAS), 2016
10. International Symposium on Computer Architecture (ISCA), 2015
11. ACM/IEEE Design Automation Conference (DAC), 2011, 2012, 2015
12. International Conference on Parallel Processing (ICPP), 2015

13. International Conference on Parallel Architectures and Compilation Techniques (PACT), 2013
14. International Symposium on High-performance Computer Architecture (HPCA), 2013
15. Design Automation and Test in Europe (DATE), 2012

University Services

1. UCSD CSE Website Committee, 2018-present
2. UCSD CSE Diversity Awareness Subcommittee, 2018-present
3. UCSD CSE Graduate Admissions Committee, 2018-present
4. Diversity Liaison of UCSC CE faculty recruitment, 2017
5. Advisory Board at UCSC Silicon Valley Campus, 2016

Teaching

Spring 2019	CSE 240A Principles of Computer Architecture (Graduate)
Winter 2019	CSE 291 Emerging Computer Architectures (Graduate)
Winter 2018	CSE 141 Introduction to Computer Architecture (Undergraduate)
Fall 2017	CMPE 110 Computer Architecture (Undergradute)
Spring 2017	CMPE 202 Computer Architecture (Gradute)
Winter 2017	CMPE 122/222 Intro/Advanced VLSI Digital System Design (Undergradute/Graduate)
Fall 2016	CMPE 110 Computer Architecture (Undergraduate)
Spring 2016	CMPE 290N Topics in Computer Performance (Graduate)
Winter 2016	CMPE 110 Computer Architecture (Undergraduate)
Fall 2015	CMPE 110 Computer Architecture (Undergraduate)
Spring 2015	CMPE 202 Computer Architecture (Graduate)

Supervising

Matheus Ogleari	Ph.D.	Graduated in 2019
Xiao Liu	Ph.D.	Fall 2015 - present
Mengjie Li	Ph.D.	Fall 2016 - present
Hengyu Zhao	Ph.D.	Fall 2016 - present
Cheng Fu	Ph.D.	Fall 2018 - present
Hui Shi	Ph.D.	Fall 2018 - present
Yun Joon Soh	Ph.D.	Fall 2018 - present
Zixuan Wang	Ph.D.	Fall 2018 - present
Hanxian Huang	Ph.D.	Fall 2019 - present
Theodoros Michailidis	Ph.D.	Fall 2019 - present
Haolan Liu	Master	2019 - present
Sneha Subramanian	Master	Graduated in 2019
Harshini Voruganti	Master	Graduated in 2019
Qing Zhang	Master	Graduated in 2017
Naga Venkata Sai Indubhaskar Jupudi	Master	Graduated in 2017
Borui Wang	Master	Graduated in 2016
Yiqiao Hu	Master	Graduated in 2016