ZHICHAO CAO

CONTACT INFORMATION

Arizona State University
BYENG 416
Tempe Campus
Tempe, AZ 85281, USA

Homepage: isearch.asu.edu/profile/4082902
E-mail: zhichao.cao@asu.edu

RESEARCH INTERESTS

Key-Value Data Systems and NoSQL Databases: LSM-based key-value store systems (RocksDB, LevelDB, HBase), caching systems (CacheLib, Kvrocks), graph databases (Nebula, Neo4j), indexing for storage and database systems (hybrid, R-tree, B-Tree, and ART), and caching optimizations for storage and databases.

Intelligent Data Systems for AI: LLM-driven optimizations for storage and database systems, similarity search indexing (in RAG), key-value data systems for direct GPU-SSD I/Os, LLM-serving optimizations (caching, offloading, and indexing), failure protection for LLM (bitflip protection, ECC, replication).

Disaggregated Data Infrastructure: Disaggregated storage systems, disaggregated memory (with CXL and RDMA), cloud storage, distributed object storage.

Storage and Data Systems for Emerging and Sustainable Devices: Non-Volatile Memory (NVM), Shingled Magnetic Recording (SMR), Interlaced Magnetic Recording (IMR), Zoned Namespace SSDs (ZNS SSDs).

EDUCATION

University of Minnesota, Twin-Cities — Aug. 2013 - Jul. 2020
Ph.D. in Computer Science — Advisor: Prof. David.H.C. Du

Thesis: High-Performance and Cost-Effective Storage Systems for Supporting Big Data Applications [pdf]

ACM Doctoral Dissertation Award Nomination by Department of Computer Science and Engineering

Tsinghua University — Sep. 2009 - Jul. 2013
B.E. in Automation (with honor) — Advisor: Prof. Qing Li

ACADEMIC AND PROFESSIONAL EXPERIENCES

Assistant Professor	**Arizona State University**	Jan. 2022 - Present
Research Scientist	**Meta (Facebook)**	Oct. 2019 - Dec. 2021
Research Collaborator	**Facebook**	Sep. 2018 - Sep. 2019
Research Intern	**Facebook**	Jun. 2018 - Aug. 2018
Research Intern	**Veritas**	Jun. 2016 - Aug. 2016
Research Intern	**Hewlett-Packard (HPE)**	Jun. 2015 - Aug. 2015
Research Intern	**Hewlett-Packard (HPE)**	Jun. 2014 - Aug. 2014
Research Assistant	**University of Minnesota, Twin-Cities**	Sep. 2013 - Sep. 2018

SELECTED HONORS AND AWARDS

- IEEE ICDE 2026 Distinguished PC Award — 2026
- NSF CAREER Award — 2025
- ACM HPDC 2025 Best Student Paper Award — 2025
- Distinguished Reviewers Board of ACM Transactions on Database Systems — 2025

- HotStorage 2024 Best Paper Award 2024
- VLDB 2024 Distinguished Reviewer Award 2024
- Arizona State University Professor of Impact Award 2024
- USENIX FAST Student Travel Grant 2019
- USENIX FAST Student Travel Grant 2017
- Best Innovation Pod among all intern teams of Veritas 2016

PUBLICATIONS

LEGEND

- (∗) Ph.D. student supervised by Dr. Cao
- (†) Other graduate or undergraduate student supervised by Dr. Cao

Peer-Reviewed Conference Publications

29. [**SIGMOD'26**] Qi Lin*, Gangqi Huang, Te Guo, Chang Guo*, Viraj Thakkar*, Zichen Zhu, Jianguo Wang, **Zhichao Cao**. "O^3-LSM: Maximizing Disaggregated LSM Write Performance via Three-Layer Offloading." *Proceedings of ACM Conference on Management of Data (SIGMOD)*, To Appear.
28. [**SeQureDB'26**] Viraj Thakkar*, Dongha Kim, Hokeun Kim, **Zhichao Cao**. "ANCHOR: A Vision for Secure Persistent Key-Value Stores in Disaggregated Data Centers." *Workshop on Secure and Private Data Management co-located with ACM SIGMOD 2026 (SeQureDB 2026)*, To Appear.
27. [**ICS'26**] Sungjin Byeon*, Kyungwook Min, Jaewan Park, Sangyun Lee, Hong-Yeon Kim, Junyoung Han, Joo-Young Hwang, **Zhichao Cao**, Youngjae Kim. "ColdMap: Compaction-Aware Cost-Benefit Zone Cleaning for ZNS-Based Key-Value Stores." *Proceedings of the ACM International Conference on Supercomputing (ICS)*, To Appear.
26. [**ICDE'26**] Fei Shao, Jia Zou, **Zhichao Cao**, Xusheng Xiao. "PROGQL: A Provenance Graph Query System for Cyber Attack Investigation." *The 42nd IEEE International Conference on Data Engineering (ICDE)*, To Appear.
25. [**EMNLP'25**] Yuhang Chen, Zhen Tan†, Ajay Kumar Jaiswal, Huaizhi Qu, Xinyu Zhao, Qi Lin*, Yu Cheng, Andrew Kwong, **Zhichao Cao**, Tianlong Chen. "Bit-Flip Error Resilience in LLMs: A Comprehensive Analysis and Defense Framework." *The 2025 Conference on Empirical Methods in Natural Language Processing (EMNLP), Main Conference*, November 2025, Pages 10425-10435. [pdf]
24. [**HPDC'25**] Chang Guo*, Ning Yan, Lipeng Wan, **Zhichao Cao**. "LegoIndex: A Scalable and Modular Indexing Framework for Efficient Analysis of Extreme-Scale Particle Data." *The 34th ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC)*, July 2025, Article No.: 25, Pages 1 - 14.[pdf] **Best Student Paper Award!**
23. [**HotStorage'25**] Chang Guo*, Norbert Podhorszki, Greg Eisenhauer, Zhiwen Xie†, Scott Klasky, **Zhichao Cao**. "Unlocking the Unusable: A Proactive Caching Framework for Reusing Partial Overlapped Data." *The 17th ACM Workshop on Hot Topics in Storage and File Systems (HotStorage)*, June 2025, Pages 129 - 136.[pdf]
22. [**SIGMOD'25**] Viraj Thakkar*, Dongha Kim, Hokeun Kim, **Zhichao Cao**. "SHIELD: Encrypting Persistent Data of LSM-KVS from Monolithic to Disaggregated Storage." *Proceedings of ACM Conference on Management of Data (SIGMOD)*, June 2025, Volume 3, Issue 3, Article No.: 217, Pages 1 - 28.[pdf]
21. [**SOSP'24**] Shushu Yi, Shaocong Sun, Li Peng, Yingbo Sun, Ming-Chang Yang, **Zhichao Cao**, Qiao Li, Myoungsoo Jung, Ke Zhou, Jie Zhang. "BIZA: Design of Self-Governing Block-Interface ZNS AFA for Endurance and Performance." *The 30th ACM Symposium on Operating Systems Principles (SOSP)*, November 2024.[pdf]

20. [**HotStorage'24**] Viraj Thakkar*,Madhumitha Sukumar†, Jiaxin Dai†, Kaushiki Singh†, **Zhichao Cao**. "Can Modern LLMs Tune and Configure LSM-based Key-Value Stores?." *16th ACM Workshop on Hot Topics in Storage and File Systems (HotStorage)*, July 2024, Pages 116 - 123. [pdf] **Best Paper Award!**

19. [**HotStorage'24**] Chongzhuo Yang†, Zhang Cao†, Chang Guo*, Ming Zhao, **Zhichao Cao**. "Can ZNS SSDs be Better Storage Devices for Persistent Cache?." *16th ACM Workshop on Hot Topics in Storage and File Systems (HotStorage)*, July 2024, Pages 55 - 62. [pdf]

18. [**SIGMOD'24**] Qiaolin Yu†, Chang Guo*, Jay Zhuang, Viraj Thakkar*,Jianguo Wang, **Zhichao Cao**. "CaaS-LSM: Compaction-as-a-Service for LSM-based Key-Value Stores in Storage-Disaggregated Infrastructure." *Proceedings of ACM Conference on Management of Data (SIGMOD)*, Volume 2, Issue 3, Article No.: 124, Pages 1 - 28. [pdf]

17. [**MSST'24**] Zhang Cao†, Chang Guo*, Ziyuan Lv†, Anand Ananthabhotla, **Zhichao Cao**. "SAS-Cache: A Semantic-Aware Secondary Cache for LSM-based Key-Value Stores." *The 38th International Conference on Massive Storage Systems and Technology*, June 2024. [pdf]

16. [**MSST'24**] Gaoji Liu†, Chongzhuo Yang†, Qiaolin Yu†, Chang Guo*, Wen Xia, **Zhichao Cao**. "Prophet: Optimizing LSM-Based Key-Value Store on ZNS SSDs with File Lifetime Prediction and Compaction Compensation." *The 38th International Conference on Massive Storage Systems and Technology* , June 2024. [pdf]

15. [**DSDE'24**] Chongzhuo Yang†, Baolin Feng, Zhang Cao†, **Zhichao Cao**. "HyzoneStore: Hybrid Storage with Flexible Logical Interface and Optimized Cache for Zoned Devices." *Proceedings of the 2024 7th International Conference on Data Storage and Data Engineering*, February 2024, Pages 71 - 77. [pdf]

14. [**ICCD'23**] **Zhichao Cao**, Hao Wen, Fenggang Wu, David H.C. Du. "SMRTS: A Performance and Cost-Effectiveness Optimized SSD-SMR Tiered File System with Data Deduplication." *The 41st IEEE International Conference on Computer Design*, 2023.[pdf]

13. [**ICCD'23**] Hao Wen, **Zhichao Cao**, Bingzhe Li, David Du, Ayman Abouelwafa, Doug Voigt, Shiyong Liu, Jim Diehl and Fenggang Wu. "K8sES: Optimizing Kubernetes with Enhanced Storage Service-Level Objectives." *The 41st IEEE International Conference on Computer Design*, 2023.[pdf]

12. [**ICCD'22**] Jingsong Yuan, Xiangyu Zou, Han Xu, **Zhichao Cao**, Shiyi Li, Wen Xia, Peng Wang and Li Chen. "A Focused Garbage Collection Approach for Primary Deduplicated Storage with Low Memory Overhead." *The 40th IEEE International Conference on Computer Design*, 2022.[pdf]

11. [**ATC'21**] Hiwot Tadese Kassa, Jason Akers, Mrinmoy Ghosh, **Zhichao Cao**, Vaibhav Gogte, Ronald Dreslinski. "Improving Performance of Flash Based Key-Value Stores Using Storage Class Memory as a Volatile Memory Extension." *2021 USENIX Annual Technical Conference*, 2021. [pdf]

10. [**FAST'20**] **Zhichao Cao**, Siying Dong, Sagar Vemuri, and David H.C. Du.. "Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebook." *18th USENIX Conference on File and Storage Technologies*, 2020. [pdf]

9. [**FAST'19**] **Zhichao Cao**, Shiyong Liu, Fenggang Wu, Guohua Wang, Bingzhe Li, and David H.C. Du. "Sliding Look-Back Window Assisted Data Chunk Rewriting for Improving Deduplication Restore Performance." *17th USENIX Conference on File and Storage Technologies*, 2019. [pdf]

8. [**HotStorage'19**] Fenggang Wu, Bingzhe Li, **Zhichao Cao**, Baoquan Zhang, Minghong Yang, Hao Wen, and David H.C. Du. "ZoneAlloy: Elastic Data and Space Management for Hybrid SMR Drives." *11th USENIX Workshop on Hot Topics in Storage and File Systems*, 2019. [pdf]

7. [**FAST'18**] **Zhichao Cao**, Hao Wen, Fenggang Wu, and David H.C. Du. "ALACC: Accelerating Restore Performance of Data Deduplication Systems Using Adaptive Look Ahead Window Assisted Chunk Caching." *16th USENIX Conference on File and Storage Technologies*, 2018. [pdf]

6. **[HotStorage'18]** Fenggang Wu, Baoquan Zhang, **Zhichao Cao**, Hao Wen, Bingzhe Li, Jim Diehl, Guohua Wang, and David H.C. Du. "Data Management Design for Interlaced Magnetic Recording." *10th USENIX Workshop on Hot Topics in Storage and File Systems*, 2018. [pdf]
5. **[MASCOTS'18]** Hao Wen, **Zhichao Cao**, Yang Zhang, Xiang Cao, Ziqi Fan, Doug Voigt, and David H.C. Du. "JoiNS: Meeting Latency SLO with Integrated Control for Networked Storage." *IEEE 26th International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems*, 2018. [pdf]
4. **[BIGCOM'18]** Shiyong Liu, **Zhichao Cao**, Zhongwen Guo , Guohua Wang , Xupeng Wang , Zhijin Qiu , and Xukun Qin. "NVMTFS: A Non-Volatile Memory Adaptive File System for Tiered Storage System." *IEEE 4th International Conference on Big Data Computing and Communications*, 2018. [Website]
3. **[OTM'16]** Qing Li, Dachuan Li, and **Zhichao Cao**. "Service Oriented Collaborative Simulation in Concept and Design Stages: Framework and Enabling Technologies." *OTM Confederated International Conferences" On the Move to Meaningful Internet Systems". Springer*, 2016.[Website]
2. **[CCIS'12]** **Zhichao Cao**, Qing Li, Zeyuan Wang, Weihua Li, Jun Li, and Ruiyang Du. "A cloud computing based framework of group-enterprise service integration and sharing." *IEEE 2nd International Conference on Cloud Computing and Intelligence Systems*, 2012. [Website]
1. **[CSSS'12]** Zeyuan Wang, Qing Li, **Zhichao Cao**, Weihua Li, Jun Li, and Ruiyang Du. "A model-based deployment framework of integrated public cloud service." *2012 International Conference on Computer Science and Service System*, 2012. [Website]

Peer-Reviewed Journal Publications

10. **[Front. Pain Res.'25]** Nick Seah, Catherine D Chong, Oana M Dumitrascu, Todd J Schwedt, **Zhichao Cao**, Teresa Wu. "Migraine is associated with a higher risk of ischemic and hemorrhagic stroke: an analysis of the All of Us database." *Frontiers in Pain Research 6 (2025): 1646142.*, 2025.[pdf]
9. **[JAS'25]** Xianyu He, Chaoshu Yang, Runyu Zhang, Huizhang Luo, **Zhichao Cao**, Jeff Zhang. "Optimizing both performance and tail latency for B+ tree on persistent memory." *Journal of Systems Architecture*, 2025.[pdf]
8. **[TC'24]** Yixun Wei†, **Zhichao Cao**, David HC Du. "CPI: A Collaborative Partial Indexing Design for Large-Scale Deduplication Systems." *IEEE Transactions on Computers*, 2024.[pdf]
7. **[TOS'22]** **Zhichao Cao**, Huibing Dong, Yixun Wei, Shiyong Liu, and David H.C. Du. "IS-HBase: An In-Storage Computing Optimized HBase with I/O Offloading and Self-Adaptive Caching in Compute-Storage Disaggregated Infrastructure." *ACM Transactions on Storage*, Volume 18, Issue 2, May 2022. [pdf]
6. **[TOS'22]** Hiwot Tadese Kassa, Jason Akers, Mrinmoy Ghosh, **Zhichao Cao**, Vaibhav Gogte, Ronald Dreslinski. "Power-optimized Deployment of Key-value Stores Using Storage Class Memory." *ACM Transactions on Storage*, Volume 18, Issue 2, May 2022.[pdf]
5. **[TOS'22]** Xiongzi Ge **Zhichao Cao**, David H.C. Du, Pradeep Ganesan, Dennis Hahn. "HintStor: A Framework to Study I/O Hints in Heterogeneous Storage." *ACM Transactions on Storage*, Volume 18, Issue 2, 2022. [pdf]
4. **[TC'20]** Fenggang Wu, Bingzhe Li, Baoquan Zhang, **Zhichao Cao**, Jim Diehl, Hao Wen, David HC Du. "Tracklace: Data management for interlaced magnetic recording." *IEEE Transactions on Computers*, Volume: 70 Issue: 3, Page(s): 347 - 358, April 2020. [pdf]
3. **[TOS'19]** **Zhichao Cao**, Hao Wen, Xiongzi Ge, and David H.C. Du. "TDDFS: A Tier-aware Data Deduplication based File System." *ACM Transaction on Storage*, 2019. [pdf]
2. **[EIS'15]** Qing Li, Zeyuan Wang, **Zhichao Cao**, Ruiyang Du, and Hao Luo. "Process and data fragmentation-oriented enterprise network integration with collaboration modeling and collaboration agents." *Enterprise In-*

formation Systems, 2015. [Website]

1. **[CI'13]** Qing Li, Zeyuan Wang, Weihua Li, **Zhichao Cao**, Ruiyang Du, and Hao Luo. "Model-based services convergence and multi-clouds integration." *Computers in Industry*, 2013.[Website]

Preprints

4. **[arXiv'25]** Qi Lin*, Zhenyu Zhang*,Viraj Thakkar*, Zhenjie Sun*, Mai Zheng, **Zhichao Cao**. "StorageXTuner: An LLM Agent-Driven Automatic Tuning Framework for Heterogeneous Storage Systems." *arXiv preprint arXiv:2510.25017*, 2025.[pdf]
3. **[arXiv'25]** Viraj Thakkar*, Qi Lin*, Kenanya Keandra Adriel Prasetyo†, Raden Haryosatyo Wisjnunandono†, Achmad Imam Kistijantoro, Reza Fuad Rachmadi, **Zhichao Cao**. "ELMo-Tune-V2: LLM-Assisted Full-Cycle Auto-Tuning to Optimize LSM-Based Key-Value Stores." *arXiv preprint arXiv:2502.17606*, 2025.[pdf]
2. **[arXiv'24]** Jie Peng†, Zhang Cao†, Huaizhi Qu, Zhengyu Zhang*, Chang Guo*, Yanyong Zhang, **Zhichao Cao**, Tianlong Chen. "Harnessing Your DRAM and SSD for Sustainable and Accessible LLM Inference with Mixed-Precision and Multi-level Caching." *arXiv preprint arXiv:2410.14740*, 2024.[pdf]
1. **[arXiv'24]** Chongzhuo Yang†, Chang Guo*, Ming Zhao, **Zhichao Cao**. "A Zoned Storage Optimized Flash Cache on ZNS SSDs." *arXiv preprint arXiv:2410.11260*, 2024.[pdf]

Peer-Reviewed Posters and Work-In-Progress Reports

18. **[FAST'26]** Zhenjie Sun*, Viraj Thakkar*, Qi Lin*, Avani Wildani, **Zhichao Cao**. "Beyond ARC - Introducing LLM-driven Adaptive Policy Replacement." *24th USENIX Conference on File and Storage Technologies*, 2026.
17. **[FAST'26]** Jun Kong*, Chang Guo*, Zhenyu Zhang*, **Zhichao Cao**. "Less Index, More Speed: Accelerating Large-Scale KVCache Index for Long-Context LLM Serving." *24th USENIX Conference on File and Storage Technologies*, 2026.
16. **[FAST'26]** Wenkai Guan, Yang Zhao, Tianlong Chen, **Zhichao Cao**, Zishen Wan. "Energy-per-Tile: Characterizing the Energy Footprint of AI Workloads on Heterogeneous, Recycled Infrastructure." *24th USENIX Conference on File and Storage Technologies*, 2026.
15. **[FAST'25]** Chang Guo*, Zhenyu Zhang*, **Zhichao Cao**. "EverCache: A Multi-Tier KVCache Engine for High-Performance and High-Efficiency LLMs Inferencing." *23rd USENIX Conference on File and Storage Technologies*, 2025.
14. **[FAST'25]** Jiajun Li†, Chang Guo*, **Zhichao Cao**. "AnyTier: An LSM-Managed Dynamic Data Tiering Framework with High Generality and Efficiency." *23rd USENIX Conference on File and Storage Technologies*, 2025.
13. **[FAST'25]** Yibo Zhao, Viraj Thakkar*, **Zhichao Cao**, Zaoxing Liu. "NetLSM: Enabling an In-Network Approach for Scheduling LSM-KVS Operations." *23rd USENIX Conference on File and Storage Technologies*, 2025.
12. **[FAST'24]** Madhumitha Sukumar†, Jiaxin Dai†, Kaushiki Singh†, Vikriti Lokegaonkar†, Viraj Thakkar*,**Zhichao Cao**. "LLM-assisted Automatic-Configuration and Tuning Framework for LSM-based Key-Value Stores." *22th USENIX Conference on File and Storage Technologies*, 2024.
11. **[MSST'24]** Kritshekhar Jha, Alexander Sutila, Ian Mcdonough, Yongfeng Wang, Lillian Seebold, **Zhichao Cao**, Ming Zhao. "ZNSCache: Zoned Namespace (ZNS) SSD based Caching." *The 38th International Conference on Massive Storage Systems and Technology*, 2024.
10. **[FAST'23]** Kritshekhar Jha, Ian Mcdonough, Alexander Sutila, **Zhichao Cao**, and Ming Zhao.. "DM-ZCache: Zoned Namespace (ZNS) SSD based Caching." *21th USENIX Conference on File and Storage Technologies*, 2023.
9. **[FAST'23]** Jinghuan Yu, Yixun Wei†, **Zhichao Cao**, David H.C. Du, and Chun Jason Xue.. "Level-based Shard Migration in Distributed LSM KV Store." *21th USENIX Conference on File and Storage Technologies*, 2023.

8. [**FAST'20**] **Zhichao Cao**, Siying Dong, Sagar Vemuri, and David H.C. Du.. "Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebook." *18th USENIX Conference on File and Storage Technologies*, 2020.
7. [**FAST'19**] **Zhichao Cao**, Shiyong Liu, Fenggang Wu, Guohua Wang, Bingzhe Li, and David H.C. Du. "Sliding Look-Back Window Assisted Data Chunk Rewriting for Improving Deduplication Restore Performance." *17th USENIX Conference on File and Storage Technologies*, 2019.
6. [**FAST'19**] Fenggang Wu, **Zhichao Cao**, Baoquan Zhang, and David H.C. Du. "Wear-out Aware LSM System for QLC SSDs." *17th USENIX Conference on File and Storage Technologies*, 2019.
5. [**FAST'19**] Fenggang Wu, Baoquan Zhang, **Zhichao Cao**, and David H.C. Du. "NVLSM-Tree: A Design of Log-Structured Merge Tree for Hybrid Volatile/Non-Volatile Memory System." *17th USENIX Conference on File and Storage Technologies*, 2019.
4. [**FAST'18**] **Zhichao Cao**, Hao Wen, Fenggang Wu, and David H.C. Du. "ALACC: Accelerating Restore Performance of Data Deduplication Systems Using Adaptive Look Ahead Window Assisted Chunk Caching.." *16th USENIX Conference on File and Storage Technologies*, 2018.
3. [**FAST'17**] **Zhichao Cao**, Fenggang Wu, Hao Wen, and David H.C. Du. "Optismr: Restore-Performance Optimization for Deduplication Systems Using SMR Drives." *16th USENIX Conference on File and Storage Technologies*, 2017.
2. [**FAST'17**] Hao Wen, **Zhichao Cao**, Yang Zhang, and David H.C. Du. "Guaranteed QoS with Integrated Control for Networked Storage." *16th USENIX Conference on File and Storage Technologies*, 2017.
1. [**SoCC'14**] Xiongzi Ge, **Zhichao Cao**, and David H.C. Du. "OneStore: Integrating Local and Cloud Storage with Access Hints." *ACM Symposium on Cloud Computing*, 2014.

INVITED TALKS

17. "The Chemical Reaction Between AI and Storage System Research", *ECE University of Minnesota, Twin Cities, Invited Talk*, MN, 2026.
16. "The Chemical Reaction Between AI and Storage System Research", *HPE Labs Palo Alto, Invited Talk*, CA, 2025.
15. "Towards Disaggregated and Intelligent Log-structured Merge-tree-based Key-Value Stores", *TikTok Headquarters San Jose, Invited Talk*, CA, 2025.
14. "Data Systems for LLMs and LLMs for Data Systems", *LinkedIn Headquarters Mountain View, Invited Talk*, CA, 2025.
13. "Z-CacheLib: Designing a High-Performance and Flash Friendly Persistent Cache using ZNS SSDs", *Cache Summit Menlo Park, Invited Talk*, CA, 2025.
12. "Decoupled-RocksDB and Its Auto-Tuning with Morden LLMs", *RocksDB End-of-Year Meetup Menlo Park, Invited Talk*, CA, 2024.
11. "Storage and Machine Learning - What can we learn from each other?", *16th ACM Workshop on Hot Topics in Storage and File Systems Santa Clara, Panelist*, CA, 2024.
10. "Can ZNS SSDs be Better Storage Devices for Persistent Cache?", *16th ACM Workshop on Hot Topics in Storage and File Systems HotStorage'24*, CA, 2024.
9. "LSM-based Key-Value Stores in AI/ML Era", *University of Chicago Chicago, Invited Talk* , IL, 2024.
8. "SMRTS: A Performance and Cost-Effectiveness Optimized SSD-SMR Tiered File System with Data Deduplication", *The 41st IEEE International Conference on Computer Design ICCD'23*, DC, 2023.
7. "Optimizing LSM-based Key-Value Stores for Disaggregated Infrastructure and New Storage Devices", *UC Santa Cruz, Invited CSE Seminar*, CA, 2023.

6. "RocksDB Secondary Cache, Checksum, and Optimizations", *Nebula Graph Meetup, Invited Talk*, CA, 2021.
5. "Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebook", *18th USENIX Conference on File and Storage Technologies [FAST'20]*, CA, 2020.
4. "RocksDB Workload Analyzing and Benchmarking", *RocksDB Community Meetup, Invited Talk*, CA, 2020.
3. "Sliding Look-Back Window Assisted Data Chunk Rewriting for Improving Deduplication Restore Performance", *17th USENIX Conference on File and Storage Technologies [FAST'19]*, MA, 2019.
2. "ALACC: Accelerating Restore Performance of Data Deduplication Systems Using Adaptive Look Ahead Window Assisted Chunk Caching", *16th USENIX Conference on File and Storage Technologies [FAST'18]*, CA, 2018.
1. "Optismr: Restore-Performance Optimization for Deduplication Systems Using SMR Drives", *15th USENIX Conference on File and Storage Technologies [FAST'17]*, CA, 2017.

PROFESSIONAL SERVICES

Editorships

- Associate Editor: ACM Transactions on Storage (2025 - Present)

Conference Organizer

- Mentorship Program Co-Chair of USENIX FAST 2027
- Mentorship Program Co-Chair of USENIX FAST 2026
- Work-in-Progress and Poster Program Co-Chair of USENIX FAST 2026
- Session Chair of USENIX FAST 2026
- Mentorship Program Co-Chair of USENIX FAST 2025
- Publicity Co-Chair of ACM HotStorage 2025
- Session Chair of USENIX ATC 2024
- Session Chair of ACM HotStorage 2024
- Publicity Co-Chair of MSST 2024
- Session Chair of IEEE ICCD 2023
- Session Chair of ACM SIGMOD 2023
- Proceedings Co-Chair of ACM SIGMOD 2023
- Virtual Chair of ACM HotStorage 2022

Technical Program Committees

- Program Committee of ACM SIGMOD 2027
- Program Committee of VLDB 2027
- Program Committee of USENIX FAST 2027
- Program Committee of ACM SIGMOD 2026
- Program Committee of VLDB 2026
- Program Committee of IEEE ICDE 2026
- Program Committee of SC 2026

- Program Committee of USENIX FAST 2026
- Program Committee of SOSP 2026 Workshop (BigMem)
- Program Committee of USENIX FAST 2025
- Program Committee of USENIX ATC 2025
- Program Committee of ACM SIGMOD 2025
- Program Committee of VLDB 2025
- Program Committee of ACM HotStorage 2025
- Program Committee of IEEE ICDCS 2025
- Program Committee of SOSP 2025 Workshop (BigMem)
- Program Committee of USENIX ATC 2024
- Program Committee of ACM SIGMOD 2024
- Program Committee of VLDB 2024
- Program Committee of ACM HotStorage 2024
- Program Committee of ACM SYSTOR 2024
- Program Committee of ACM SIGMOD 2023
- Program Committee of ACM HotStorage 2023
- Program Committee of ICPP 2023
- Program Committee of IEEE NAS 2022
- Program Committee of ACM APSys 2022

Journal Reviewing

- Reviewer of ACM Transactions on Storage (TOS) (2022, 2023, 2024, 2025, 2026)
- Reviewer of ACM Transactions on Computer Systems (TOCS) (2026)
- Reviewer of ACM Transactions on Database Systems (TODS) (2024, 2025)
- Reviewer of IEEE Micro (2024, 2025)
- Reviewer of ACM Computing Surveys (2025, 2026)
- Reviewer of ACM Transactions on Architecture and Code Optimization (TACO) (2024, 2025)
- Reviewer of IEEE Transactions on Computers (TC) (2023, 2024)
- Reviewer of IEEE/ACM Transactions on Networking (2024)
- Reviewer of IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD) (2023)
- Reviewer of IEEE Transactions on Cloud Computing (2022)
- Reviewer of IEEE Transactions on Dependable and Secure Computing (2021)
- Reviewer of IEEE Access (2021)
- Reviewer of IEEE Intelligent Systems (2021)
- Reviewer of The International Journal for the Computer Communications (2021)
- Reviewer of The International Journal for the Future Generation Computer Systems (2020)

Proposal Reviewing

- NSF TIP Directorate Panelist (2026)
- NSF CISE Directorate Panelist (2024)
- Facebook Faculty Research Award Panelist (2021)

Department/School Service

- ASU AI Ph.D. Program, Graduate Program Committee (2025 - Present)
- ASU Data Science, Analytics and Engineering Program, Graduate Program Committee Chair (2023 - 2024)
- ASU Data Science, Analytics and Engineering Program, Graduate Program Committee (2022 - Present)
- ASU Data Science, Analytics and Engineering Program, M.S. Program Committee (2023 - 2024)
- ASU SCAI Faculty Search Committee (2022 - 2024)
- ASU SCAI Ph.D. Admission Committee (2022 - 2025)

STUDENTS SUPERVISED

Ph.D. Students

Chang Guo	Ph.D. Student	August. 2022 - 2027 (expected)
Viraj Thakkar	Ph.D. Student (Fulton Fellowship)	August. 2023 - 2027 (expected)
Qi Lin	Ph.D. Student (Fulton Fellowship)	June. 2024 - 2028 (expected)
Zhenyu Zhang	Ph.D. Student	August. 2024 - 2029 (expected)
Jun Kong	Ph.D. Student	August. 2025 - 2030 (expected)
Zhenjie Sun	Ph.D. Student	August. 2025 - 2030 (expected)
Sungjin Byeon	Ph.D. Student	August. 2026 - 2030 (expected)

Other Students Mentored

- **M.S. Students:** Chongzhuo Yang (2022), Tanmesh Mishra (2023), Uma Maheshwara Swamy Desineedi (2023), Avish Khosla (2023), Jiaxin Dai (2024), Nick Seah (2024), Zhiwen Xie (2024), Sampada Nemade (2025), Srinivas Oguri (2025), Harry Samuel DeCecco (2025), Khadyothan Choudari Dasari (2025)
- **Undergraduate Students (Research):** Qiaolin Yu (2022), Ziyuan Lv (2022), Gaoji Liu (2023), Zhang Cao (2023), Vikriti Lokegaonka (2023), Madhumitha Sukumar (2024), Kaushiki Singh (2024), Bhavana Kannan (2025), Harshith Vijayan (2025), Keshava Subramanian (2025)

Thesis Defense Committee

- **Ph.D. Dissertation committee:** Yiming Wei (2024), Wangyang Ying (2025), Hong Guan (2026), Cheng (Jesse) Jing (2026), Vikram Ramaswamy (2026), Kritshekhar Jha (2026)
- **Master thesis defense committee:** Sungho Hong (2022), Viraj Thakkar (2023), Manimozhi Sekar (2024), Vrutik Halani (2024), Nick Seah (2024), Vrutik Halani (2025), Alexander Sutila (2025)
- **Master Applied Project Supervised**: Harry Samuel DeCecco (2025), Khadyothan Choudari Dasari (2025)