ZHICHAO CAO

Address: Arizona State University, Tempe campus, Arizona
EMAIL: Zhichao.Cao@asu.edu **HOMEPAGE:** https://isearch.asu.edu/profile/4082902

RESEARCH INTEREST

Data Infrastructure: key-value stores (RocksDB, LevelDB, HBase); NoSQL databases (GraphDB, VectorDB); data deduplication; backup and archive system; file system; hierarchical storage system; distributed storage system; compute-storage disaggregation; and memory disaggregation

Storage for big data: cloud storage; object storage; storage systems for big graph; AI/ML for storage system; storage systems for AI/ML; storage system in IoT

New storage devices: Disaggregated Memory (with CXL and RDMA); Non-Volatile Memory (NVM); Shingled Magnetic Recording (SMR); Interlaced Magnetic Recording (SMR); Zoned Namespace SSDs (ZNS SSDs); DNA- and Glass-based storage

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*

EMPLOYMENTS

Assistant Professor	**Arizona State University**	Jan. 2022 - Present
Research Scientist	**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

PUBLICATIONS

26. **[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)*, 2024, Proc. ACM Manag. Data 2, 3 (SIGMOD), Article 124 (June 2024), 28 pages. [pdf]

25. **[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)*, 2024, To Appear.

24. **[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)*, 2024, To Appear.

23. [**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 (MSST)*, 2024, To Appear.

22. [**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 (MSST)*, 2024, To Appear.

21. [**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*, 2024. [pdf]

20. [**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 *(Acceptance rate: 28%).*[pdf]

19. [**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 *(Acceptance rate: 28%).*[pdf]

18. [**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]

17. [**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 Transaction on Storage*, Volume 18, Issue 2, May 2022. [pdf]

16. [**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 Transaction on Storage*, Volume 18, Issue 2, May 2022.[pdf]

15. [**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 Transaction on Storage*, Volume 18, Issue 2, May 2022. [pdf]

14. [**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 *(Acceptance rate: 64/341=23%).* [pdf]

13. [**FAST'20**] **Zhichao Cao**, Siying Dong, Sagar Vemuri, and David H.C. Du.. "Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebooke." *18th USENIX Conference on File and Storage Technologies*, 2020 *(Acceptance rate: 23/138=17% as Full Paper).* [pdf]

12. [**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 *(Acceptance rate: 26/145=18% as Full Paper).* [pdf]

11. [**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]

10. [**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]

9. [**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 *(Acceptance rate: 23/139=17% as Full Paper).* [pdf]

8. **[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]
7. **[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]
6. **[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]
5. **[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]
4. **[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 Information Systems*, 2015. [Website]
3. **[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]
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]

ACADEMIC POSTERS AND WORK-IN-PROGRESS

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.

HONORS, AWARDS & PATENT

- Indonesia-US Research Collaboration Award 2024
- Indonesia-US Research Collaboration Award 2023
- Google Cloud Research Credit 2022
- FAST Travel Grant 2019
- U.S. Patent "System and Methods for Performing Live Migrations of Software Containers", 15/261,596[P]. 2018
- FAST Travel Grant 2017
- Best Innovation Pod among all intern teams of Veritas 2016

TEACHING

CSE 511 (graduate-level) Spring 2022, Spring 2023
Data Processing at Scale *Instructor*

CSE 330 (undergraduate-level) Fall 2022, Fall 2023, Sprin 2024
Operating Systems *Instructor*

CSCI 5980 (graduate-level) Fall 2018
Big Data and Storage System *Guest lecture*

CSCI 2021 (undergraduate-level) Sprint 2014
Machine Architecture and Organization *Teaching assistant*

CURRENT STUDENTS SUPERVISED

Chang Guo	Ph.D. Student	August. 2022 - Present
Viraj Thakkar	Ph.D. Student	August. 2023 - Present

THESIS DEFENSE COMMITTEE

- Ph.D. thesis defense committee: Yiming Wei (2024)
- Master thesis defense committee: Sungho Hong (2022), Viraj Thakkar (2023), Manimozhi Sekar (2024), Vrutik Halani (2024)

ACADEMIC SERVICES

- Program Committee of USENIX FAST 2025
- Program Committee of ACM SIGMOD 2023, 2024 (demo track), 2025
- Program Committee of VLDB 2024, 2025
- Program Committee of USENIX ATC 2024
- Publicity Co-Chair of MSST 2024
- Program Committee of ACM HotStorage 2023, 2024
- Program Committee of ACM SYSTOR 2024
- Session Chair of IEEE ICCD 2023
- Program Committee of ICPP 2023
- Session Chair of ACM SIGMOD 2023
- Proceedings Co-Chair of ACM SIGMOD 2023
- Virtual Chair of ACM HotStorage 2022
- Program Committee of IEEE NAS 2022
- Program Committee of ACM APSys 2022
- Reviewer of ACM Transaction on Storage (TOS) (2022, 2023)
- Reviewer of IEEE Transactions on Computers (TC)
- Reviewer of IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)
- Reviewer of IEEE Transactions on Cloud Computing
- Reviewer of IEEE/ACM Transactions on Networking
- Reviewer of IEEE Transactions on Dependable and Secure Computing
- Reviewer of IEEE Access
- Reviewer of IEEE Intelligent Systems
- Reviewer of The International Journal for the Computer Communications
- Reviewer of The International Journal for the Future Generation Computer Systems
- Volunteer of International Conference on Parallel Processing (ICPP'14)

DEPARTMENT SERVICES

Ph.D. Admission Committee	2022 - Present
Faculty Search Committee	2022 - 2023
DSAE Graduate Program Committee (Chair)	2022 - Present

RESEARCH COOPERATION WITH INDUSTRIES

Western Digital Feb. 2022 - Present
Project: System design for Zoned Namespace SSDs ***PI***

Facebook Sep. 2018 - Sep. 2019
Project: RocksDB Workload Characterization and Benchmarking ***PI***

Hewlett-Packard (HPE) Sep. 2016 - Jun. 2018
Project: Integrating SDS with SDN *with Hao Wen*

Veritas Sep. 2016 - Jun. 2017
Project: Global data allocation and migration project ***co-PI*** *with David H.C Du*

Symantec Sep. 2015 - Jun. 2016
Project: Federated and Distributed Storage System ***co-PI*** *with David H.C Du*

NetApp Sep. 2013 - Jun. 2015
Project: Integrating Local Storage and Cloud with Access Hints ***co-PI*** *with David H.C Du*

INDUSTRIAL EXPERIENCES

Facebook Oct. 2019 - Dec. 2021
Project: Database and RocksDB Research ***Research Scientist***

- RocksDB key value store performance and data integrity research;
- RocksDB data protection and workloads exploration;
- New storage devices for key-value stores;
- Research of integrating AI/ML models with RocksDB for performance improvement.

Facebook Sep. 2018 - Sep. 2019
Project: RocksDB Workload Characterization and Benchmarking ***Research Collaborator***

- Key value store workload collecting in large scale social graph, storage system, and AI platform
- Enhance RocksDB tracing tool and analyzing tool and workload characterization;
- Propose and develop the Key-value store workload characterization methodologies.

Facebook Jun. 2018 - Aug. 2018
Project: RocksDB Query Workload Resarch ***Research Intern***

- Designed and implemented the RocksDB query level trace analyzing tool;
- Deploying the trace collecting tool in two different shadow services and delivered real-world workload analyzing and characterization;
- Proposed and implemented the RocksDB synthetic workload generator.

Veritas Jun. 2016 - Aug. 2016
Project: Docker Container Live Migration ***Research Intern***

- Designed and implemented incremental container checkpoint and restore in RunC/Docker;
- Implemented live migration local plugin and UI with automatic support;
- Designed and implemented machine learning based container live migration algorithm.

Hewlett-Packard (HPE) Jun. 2015 - Aug. 2015
Project: : Light Weight Cloud Gateway File System Development ***Research Intern***

- Designed and implemented data deduplication module for the file system;
- Implemented LRU cache with cache auto shrinking to optimize file system performance;
- Implemented multi-thread infrastructure (thread pool and thread management).

Hewlett-Packard (HPE) Jun. 2014 - Aug. 2014
Project: Source Deduplication Gateway for HP Catalyst ***Research Intern***

- Designed and implemented the light-weight gateway with source deduplication;
- Designed and developed the WSGI based RESTful request gateway for HP Catalyst to support Openstack and replace Swift.

INVITED TALKS

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, CSE Seminar***, CA, 2023.
6. "RocksDB Secondary Cache, Checksum, and Optimizations", *Nebula Graph Meetup*, 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 2020 Spring*, 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.