

Zhichao Cao

3750 Tamayo St, APT 16
Fremont, CA, 94536
E-mail: caoxx380[at]umn.edu
Website: <https://www-users.cs.umn.edu/~caoxx380/>

Education

2013.8 — 2020.7 **Ph.D. in Computer Science**, University of Minnesota, Twin-Cities
Advisor: Prof. [David H.C. Du](#)

2009.9 — 2013.7 **Bachelor of Science in Automation**, Tsinghua University, Beijing, China
Advisor: Prof. [Qing Li](#)

Research Interests

- **Storage systems:** key-value store, NoSQL database, data deduplication, backup and archive system, file system, hierarchical storage system, distributed storage system
- **Storage for big data:** cloud storage, object storage, storage for data analytics, storage systems for big graph, AI/machine learning for storage system
- **New storage devices:** non-volatile memory (NVM), Shingled magnetic recording (SMR), Interlaced Magnetic Recording (SMR), Zone Namespace SSD (ZNS)

Publications

- [1] **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 *FAST 2020* (pp. 209-223).
- [2] **Zhichao Cao**, Shiyong Liu, et al. "Sliding Look-Back Window Assisted Data Chunk Rewriting for Improving Deduplication Restore Performance." 17th USENIX Conference on File and Storage Technologies *FAST 2019* (pp. 129-142).
- [3] **Zhichao Cao**, Hao Wen, Xiongzi Ge and David H.C. Du. "TDDFS: A Tier-aware Data Deduplication based File System." ACM *Transaction on Storage (ToS)*. February 2019.
- [4] **Zhichao Cao**, Hao Wen, Fenggang Wu, 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 *FAST 2018* (pp. 309-324).
- [5] Fenggang Wu, Bingzhe Li, Baoquan Zhang, **Zhichao Cao**, Jim Diehl, et al. "TrackLace: Data Management for Interlaced Magnetic Recording", IEEE *Transactions on Computers (TC)*. February 2020.
- [6] Fenggang Wu, Bingzhe Li, **Zhichao Cao**, Baoquan Zhang, et al. "ZoneAlloy: Elastic Data and Space Management for Hybrid SMR Drives." *HotStorage 2019*.
- [7] Fenggang Wu, Baoquan Zhang, **Zhichao Cao**, et al. "Data Management Design for Interlaced Magnetic Recording." *HotStorage 2018*.
- [8] Hao Wen, **Zhichao Cao**, Yang Zhang, et al. "JoiNS: Guaranteed QoS with Integrated Control for Networked Storage." *MASCOTS 2018*.
- [9] Shiyong Liu, **Zhichao Cao**, et al. "NVMTFS: A Non-Volatile Memory Adaptive File System for Tiered Storage System." 2018 4th International Conference on Big Data Computing and Communications *BIGCOM 2018*.

- [10] **Zhichao Cao**, Qing Li, Zeyuan Wang, et al. “A Cloud Computing Based Framework of Group Enterprise Service Integration and Sharing.” *CCIS, IEEE* 2012, 1: 371-375;
- [11] Qing Li, Zeyuan Wang, Weihua Li, **Zhichao Cao** et al. “Model-based services convergence and multi-clouds integration.” *Computers in Industry (CI)* 64, no. 7 (2013): 813-832.
- [12] Qing Li, Zeyuan Wang, **Zhichao Cao**, et al. “Process and data fragmentation-oriented enterprise network integration with collaboration modelling and collaboration agents.” *Enterprise Information Systems (EIS)* 9, no. 5-6 (2015): 468-498.
- [13] Zeyuan Wang, Qing Li, **ZhiChao Cao**, et al. “A Model-Based Deployment Framework of Integrated Public Cloud Service.” In Computer Science & Service System (CSSS), 2012 International Conference on, pp. 723-728. IEEE, 2012.
- [14] Qing Li, Dachuan Li, and **Zhichao Cao**. “Service Oriented Collaborative Simulation in Concept and Design Stages: Framework and Enabling Technologies.” International Conferences On the “Move to Meaningful Internet Systems”, pp. 5-15. Springer, Cham, 2016.

Honors and Awards

- **Patent**: system and methods for performing live migrations of software containers (U.S. Patent 15/261,596[P]. 2018-3-15.)
- **Best Innovation Pod**: “**Best Innovation Pod**” among all intern teams of **Veritas 2016**.
- **FAST 2017 Travel Grant**
- **FAST 2019 Travel Grant**

Teaching

- CSCI 2021: Machine Architecture and Organization (Spring 2014, teaching assistant)
- CSCI 5980: Big Data and Storage System (Fall 2018, guest lecture)

Industrial Experiences

Facebook

2019.10 — Present Research Scientist: Database and RocksDB

- Database performance research;
- RocksDB data protection exploration;

Facebook

2018.9 — 2019.9 Research Collaborator: RocksDB Workload Characterization

- Enhance RocksDB tracing tool and analyzing tool;
- Propose and develop the Key-value store workload characterization methodologies

Facebook

2018.6 — 2018.8 Research Intern: RocksDB Query Trace Analysing Tool Development

- 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 Lab

2016.6 — 2016.8 Research Intern: Docker Container Live Migration

- 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) Storage File System Group (IBRIX)

2015.6 — 2015.8 Research Intern: Light Weight Cloud Gateway File System Development

- 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) Storage CTO Office (3PAR)

2014.6 — 2014.8 Research Intern: Source Deduplication Gateway for HP Catalyst

- 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 (Python).

Research Cooperation with Industries

- **Facebook** (2018-2019) RocksDB project (PI)
- **Veritas** (2016-2017) Global data allocation and migration project (PI)
- **Hewlett-Packard** (2016-2018) SDS+SDN project (with Hao Wen)
- **Symantec** (2015-2016) Distributed storage system project (PI)
- **NetApp** (2013-2015) Hintstor project (with Xiongzi Ge)

Academic Talks

- [1] “Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebook.” 18th USENIX Conference on File and Storage Technologies *FAST 2020*.
- [2] “Sliding Look-Back Window Assisted Data Chunk Rewriting for Improving Deduplication Restore Performance.” 17th USENIX Conference on File and Storage Technologies *FAST 2019*.
- [3] “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 2018*.
- [4] “Optismr: Restore-Performance Optimization for Deduplication Systems Using SMR Drives.” The *WIPS Session* of 15th USENIX Conference on File and Storage Technologies *FAST 2017*.

Posters and Academic Activities

- [1] “Characterizing, Modeling, and Benchmarking RocksDB Key-Value Workloads at Facebook.” The *Poster Session* of 18th USENIX Conference on File and Storage Technologies *FAST 2020*.
- [2] “Sliding Look-Back Window Assisted Data Chunk Rewriting for Improving Deduplication Restore Performance.” The *Poster Session* of 17th USENIX Conference on File and Storage Technologies *FAST 2019*.
- [3] “ALACC: Accelerating Restore Performance of Data Deduplication Systems Using Adaptive

Look Ahead Window Assisted Chunk Caching.” The *Poster Session* of 16th USENIX Conference on File and Storage Technologies *FAST 2018*.

- [4] Fenggang Wu, **Zhichao Cao**, Baoquan Zhang, and David H.C. Du. "Wear-out Aware LSM System for QLC SSDs." *WIPS* of *FAST 2019*.
- [5] Baoquan Zhang, Fenggang Wu, **Zhichao Cao**, et al. "NVLSM-Tree: A Design of Log-Structured Merge Tree for Hybrid Volatile/Non-Volatile Memory System." *WIPS* of *FAST 2018*
- [6] "Optismr: Restore-Performance Optimization for Deduplication Systems Using SMR Drives." The *Poster Session* of 15th USENIX Conference on File and Storage Technologies *FAST 2017*.
- [7] Hao Wen, **Zhichao Cao**, Yang Zhang, et al. "Guaranteed QoS with Integrated Control for Networked Storage." The *Poster Session* of 15th USENIX Conference on File and Storage Technologies *FAST 2017*.
- [8] Xiongzi Ge, **Zhichao Cao**, and et al. "OneStore: Integrating Local and Cloud Storage with Access Hints." The *Poster Session* of *SoCC 2014*.

Professional Services

- Reviewer of *IEEE Transactions on Computers (TC)*
- Reviewer of *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD)*
- Reviewer of *International Journal of Future Generation Computer Systems*
- Reviewer of *Transactions on Cloud Computing*
- Reviewer of *Computer Communications*
- Reviewer of *IEEE Intelligent Systems*
- Reviewer of *Future Generation Computer Systems*
- Volunteer of *International Conference on Parallel Processing (ICPP'14)*

Current and Former Students Supervised

Ph.D. Students

- Yixun Wei
- Huibing Dong
- ChiaWen Hsieh
- Haoyu Gong

Visiting Scholars

- Shiyong Liu

Papers Under Review/Revision

- [1] **Zhichao Cao**, Huibing Dong, Yixun Wei, et al. "IS-HBase: A New In-Storage Computing based HBase Architecture"
- [2] **Zhichao Cao**, Fenggang Wu, Hao Wen, and David H.C. Du. "SMRTS: A Space Optimized File System for SSD-SMR based Tiered Storage"
- [3] Xiongzi Ge, **Zhichao Cao**, David H.C. Du. "HintStor: A Flexible Framework to Study I/O Access Hints in Heterogeneous Storage Systems."

- [4] Hao Wen, Bingzhe Li, **Zhichao Cao**, et al. “K8sES: Kubernetes with Enhanced Storage Service-Level Objectives.”

Research Projects

2020.2 — Present **Large Scale Archive and Backup System on DNA-based Storage System**

- Designed the deduplication framework for DNA based system;
- Propose the two-stage data partition and allocation algorithm to optimize read/write performance on DNA storage.

2020.1 — Present **Wear-leveling Optimized Key-Value Store for Zone Namespaces SSD**

- Designed the new SST file layout for zone namespaces SSD (ZNS);
- Proposed novel compaction to optimize wear-leveling on ZNS.

Implemented framework.

2019.6 — Present **High Performance Indexing for Large Scale Deduplication in Big data**

- Combine two indexing structure and designed two level adaptive in-memory index;
- Proposed high speed indexing lookup.

2019.1 — Present **Intelligent Storage for Graph Processing Optimization**

- Designed the new caching and data migration policies for graph processing platform with storage-computing-separation;
- Implemented framework.

2018.1 — 2019.8 **LSM-based Write Optimized Key-Value Store (KVS) for NVM**

- Designed the write optimized LSM tree based KVS with delayed compaction;
- Implemented the function modules based on PMDK (C++).

2018.1 — 2019.4 **Data Management Design for Interlaced Magnetic Recording**

- Designed and implemented the block swapping and remapping module (C++);
- Implemented and evaluated the data management simulator (C++).

2018.1 — 2019.2 **Storage SLO Optimization for Kubernetes (Cooperate with HPE)**

- Designed the node selection mechanism and algorithm for Kubernetes SLO requirement;
- Designed the combined resources scheduling (combing storage, network, and computing)

2017.4 — 2019.5 **NoSQL HBase Read Performance Improving**

- Implemented file level data access tracing tool for HBase to tracking I/O (C and C++);
- Implemented SCAN benchmark to evaluate read amplification and performance (Python);
- Proposed in storage nodes HFile scanning to reduce read amplification;
- Designed in storage filtering for SCAN and compaction.

2017.6 — Present **Machine Learning Assisted Selective Deduplication in Primary Storage**

- Implemented the multi-tiers file system prototype with machine learning interface;
- Designed and implemented decision tree (C4.5) based file deduplication selection.

2017.12 — 2019.1 **Self-adaptive Deduplication Engine with Chunk Rewriting**

- Proposed adaptive look-back chunk rewrite mechanism to make better tradeoffs between restore performance and deduplication ratio;
- Designed and implemented the workload adaptive chunk usage prediction model, which can achieve self-adjustment for deduplication engine;
- Implemented the deduplication engine with different deduplication configuration (capping, rewriting, caching) and restore polices.

2016.9 — 2017.9 **Self-adaptive Deduplication Restore Engine**

- Proposed adaptive look ahead window assisted chunk cache to speed up restore;
- Designed the self-managed replacement policy in chunk cache to reduce disk I/O;
- Implemented the deduplication engine with adaptive restore module (Language: C).

2015.9 —2017.8 TDDFS: Primary Data Deduplication File System for Tiered Storage

- Designed and implemented file temperature based file migration policy on SSD;
- Integrated data deduplication with data migration to speed up migration and restore;
- Implemented TDDFS in FUSE and enabled data migration between tiers and data chunk sharing in high-performance tier, reduced overhead (Language: C).

2015.10 — 2017.9 Data Management in with Machine Learning (Cooperated with Veritas)

- Designed and implemented hybrid data allocation and migration algorithm combining heuristic and machine learning algorithm (Sensitive-Hash and KNN);
- Proposed and implemented self-adaptive refining model for machine learning stage.

2013.10— 2015.8 Hints-based Heterogeneous Storage Systems (Cooperate with NetApp)

- Proposed and modified logical volume(LV) allocation module with new filter;
- Designed the integration between cloud storage and local storage via different hints.

Technical Skills

Language: C/C++, Python, Matlab, JAVA;

Open source projects: RocksDB, Deduplication, FUSE, Openstack, Hbase, Docker;