

CHEN ZHONG

☎ (682)256-0650 ✉ zhongch4all@gmail.com [in linkedin/zhongch4g](https://www.linkedin.com/in/zhongch4g) [github/zhongch4g.github.io](https://github.com/zhongch4g)

HIGHLIGHTS

- Built efficient caching, indexing, and key-value storage systems
- Developed a unifying framework integrating in-memory and on-disk indexes in collaboration with Tencent's database team
- Partnering with VMware vSAN team to optimize I/O interface efficiency for key-value operations

EDUCATION

University of Texas at Arlington | *Ph.D. candidate in Computer Science* Sep. 2019 – Present
Beijing University of Posts and Telecommunications | *MEng, in Software Engineering* Sep. 2016 – Jul. 2019
Jiangxi University of Science and Technology | *BEng, in Software Engineering* Sep. 2012 – Jul. 2016

WORK EXPERIENCE

Tencent America LLC. Aug. 2022 – Dec. 2022
Database Research and Development Intern | C/C++, In-memory/on-disk index, Database Bellevue, WA

- Implemented the ART tree indexing structure into production database code.
- Designed and developed high-performance extensible memory-disk index for next-generation database systems achieving 30X throughput improvement.

JD.com, Inc. Aug. 2017 – Mar. 2018
Research and Development Intern | Python, Hive, Data analysis Beijing, China

- Developed and deployed prediction model for JD.com Analytics by analyzing user behavior patterns and extracting mission-critical features for backend services.

PUBLICATIONS

- **Chen Zhong**, Q. Zhou, Y. Chen, X. Zhao, K. He, A. Pan, S. Jiang, "IndeXY: A Framework for Constructing Extensible Large Indexes for OLTP Databases" (ICDE '24).
- X. Zhao, P. Challa, **Chen Zhong**, and S. Jiang, "Developing Index Structures in Persistent Memory Using Spot-on Optimizations with DRAM" (ICPE 2024).
- S. Maharjan, S. Zhao, **Chen Zhong**, and S. Jiang, "From LeanStore to LearnedStore: Using a Learned Index to Improve Database Index Search" (HDIS 2023, *Best Paper Award*).
- X. Zhao, **Chen Zhong**, S. Jiang, "TurboHash: A Hash Table for Key-value Store on Persistent Memory", (SYSTOR '23).
- **Chen Zhong**, P. Challa, X. Zhao, S. Jiang, "Buffered Hash Table: Leveraging DRAM to Enhance Hash Indexes in the Persistent Memory" (NVMSA '22, *Best Paper Candidate*).
- **Chen Zhong**, X. Zhao, S. Jiang, "LIRS2: An Improved LIRS Replacement Algorithm" (SYSTOR '21).

RESEARCH & PROJECT EXPERIENCE

Efficient Access of Key-value Storage | *C/C++, linux kernel, SPDK, RocksDB, NVMe, KVSSD* 2023 – present

- Identified and removed performance bottleneck in on-device key-value storage systems with NVMe through detailed analysis and optimizing access patterns to reduce per-operation overhead.
- Optimized write latency and data persistence tradeoffs while developing key-value cache management strategies

A Memory-disk-spanning Index Design | *C/C++, Python, Indexing, Caching, Key-value Store* 2022 – 2023

- Developed framework to unify in-memory and on-disk indexes into single extensible system (IndeXY) with optimized policies
- Implemented performance optimizations including hot/cold data detection, granularity tuning, and intelligent data placement for efficient disk I/O management.

Improve Persistent Memory Hash Table Efficiency | *C/C++, Persistent memory* 2021 – 2022

- Identified write amplification in persistent memory hash tables due to misaligned access patterns with hardware properties.
- Designed in-DRAM buffering and write-ahead logging, achieving 2.8X throughput improvement over state-of-the-art indexes.

A Cache Replacement Algorithm | *C/C++, Block/page strategies* 2020 – 2021

- Incorporated a new data locality measure into the design of the LIRS cache replacement algorithm that improve its performance by 19.1% across various workloads, with lower time and space overheads.