


ZHELONG ZHAO

✉ zhelongzhao@hust.edu.cn  github.com/zztaki

EDUCATION

Huazhong University of Science and Technology (HUST), China 2022 – Present
Master student in Computer Science (CS), expected June 2025. Zhixing Scholarship.

Huazhong University of Science and Technology (HUST), China 2018 – 2022
B.Eng in Computer Science (CS). **National Scholarship (Top 0.2% national-wide in China)**.

WORK EXPERIENCE


Tencent Inc., Shenzhen, China 2021.7 – 2022.5
Research Intern at Cloud Architecture Platform Department. Manager: tommyhuang; Mentor: yxbillchen




- **Tencent Picture Cloud** provides services to **tens of thousands of B-side tenants**, but cache hit rate is low.
 1. Insight: Combined **cache miss rate curve** analysis and QPS on different applications, find that some applications have **high number of requests but low cache reuse rate**.
 2. Solution: Inspired by the CPU LLC shared cache partition, a **lightweight cache partition strategy** is proposed to effectively limit the amount of cache resources allocated to cache-unfriendly applications.
 3. Result: **Improve the overall cache hit rate by 9% and reduce the average access delay by 12.5ms.**
- **QQ photo album** uses SSD cache to accelerate photo download requests, but it's hit rate is low, too.
 1. Insight: After workload analysis, find that **small-size photos have a higher reuse rate**.
 2. Solution: Optimize the **insertion strategy** of the LRU cache replacement algorithm.
 3. Result: **Improve the cache hit rate by 9.6% and reduce the average access delay by 7.14ms, Tencent Technology Progress Silver Award.**

PROJECT EXPERIENCE

Ceph BlueStore Cache Policy Optimization 2023.9 - Present
Ceph BlueStore supports read cache for metadata and data to speed up data access in distributed storage.

- **Prefetch Support:** Based on Mithril, regularly mine the association of RADOS objects within the time window and generate a prefetch table. Once one object is accessed, its associated objects are prefetched.
- **Admission Support:** Based on TinyLFU, use **Counting Bloom filter** to approximate object frequency statistics. Admission or not by comparing the frequency of the eviction candidate and the accessed object.
- **New Replacement Policy Support:** Based on LRB, maintain information such as frequency and recency of objects in the data cache, use reuse distance to mark the hotness and coldness of objects, and regularly train decision tree models through the LightGBM interface to guide future cache replacement decisions.

 **libCacheSim** (Open-Source Project) 2023.8 - Present
High-performance cache simulator supports more than twenty cache algorithm simulations, supports multiple types of trace such as txt, csv, vscsi and zstd compression.

- **Prefetch Support:** Modify the critical path of the simulation process and add a cache prefetch interface.
- **Prefetch Policy Support:**  OBL  PG  Mithril. 2000+ lines were merged.

PUBLICATIONS

- Wang P, Liu Y, **Zhao Z**, et al. Smart Cache Insertion and Promotion Policy for Content Delivery Networks[C]//International Conference on Parallel Processing (ICPP). 2023: 183-192.
- Wang P, Liu Y, **Zhao Z**, et al. Adaptive Size-Aware Cache Insertion Policy for Content Delivery Networks[C]//International Conference on Computer Design (ICCD). IEEE, 2022: 195-202.
- Wang P, **Zhao Z**, Liu Y, et al. A Lightweight and Adaptive Cache Partitioning Scheme for Content Delivery Networks[C]//International Conference on Computer Design (ICCD). IEEE, 2022: 407-410.