Chen Chen | Ph.D. Candidate

⊠ mail@roychan.org • 🕆 www.roychan.org • € nerdroychan

Summary

I am a motivated Ph.D. candidate specializing in Computer Science and an enthusiast in computer systems. My expertise spans key-value stores, storage systems, data indexing and system performance. I am passionate about building fast, secure, and reliable systems software to improve the foundation of data-intensive applications.

Professional Skills

- **Data Management**: Key-Value Stores, Data Indexing, Transaction Processing, Crash Consistency, Heterogeneous Memory/Storage, Linux I/O, File Systems.
- **System Performance**: Performance Profiling, Program Analysis, Cache-aware Data Structures, Lock-free Mechanisms, Kernel Bypassing, CPU Architecture, x86 Assembly.
- **Computer Fundamentals**: Algorithms and Data Structures, Complexity Analysis, Concurrency and Parallelism, Computer Organization, Operating Systems, Networking.
- **Programming and Tools**: System Programming (C/Rust), Scripting (Python/Bash), Working with Linux distributions, GNU coreutils and virtualization/container technologies.

Education

University of Illinois at Chicago Ph.D. in Computer Science (expected 12/2024)

ShanghaiTech University *Master of Engineering in Computer Science*

ShanghaiTech University *Bachelor of Engineering in Computer Science*

Work Experience

Ph.D. Research Intern, Cockroach Labs, Inc. Storage Engine Team

05/2022 - 08/2022

01/2020 - present Shanghai, China 09/2018 - 06/2021

Chicago, IL

Shanghai, China 09/2014 - 07/2018

- Became a member of the storage team responsible for developing and maintaining the core key-value storage engine, Pebble, integral to CockroachDB.
- Participated in the research of disaggregated storage in Pebble, leveraging heterogeneous storage media/services to achieve a more cost-effective data storage in CockroachDB.
- Prototyped the initial draft of essential components for disaggregated storage in Pebble, which enables seamless key migration across multi-regional deployments and ens data compatibility.
- Summary PR of the work: O https://github.com/cockroachdb/pebble/pull/1899

Academic Experience

Ph.D. Student, University of Illinois at Chicago

01/2020 - present

Advisors: Dr. Jakob Eriksson (UIC) and Dr. Xingbo Wu (Microsoft Research)

• Decentralized Scheduling for Concurrent In-Memory Transaction Processing [1]

- Proposed a novel concurrency control protocol that that effectively reduces transaction aborts in deterministic transactions by leveraging determinism, realized by the following key features:
 - A lightweight, lock-free approach that quickly determines the global order of a transaction.
 - A decentralized scheduling algorithm that lets each individual transaction observe their globally consensual execution order without a centralized scheduler or batched execution.
 - $\cdot\,$ Integration of multiple optimization strategies to lower the protocol's space and time complexity.
- Open source: O https://github.com/decentralized-scheduling
- Efficient Data Management with Flexible Address Space [2]
 - Proposed an innovative storage abstraction, called a flexible address space, allowing structured data to move freely in an address space and be sorted efficiently by the application at a low cost.
 - Introduced the bottom-up design that realizes this concept:
 - $\cdot\,$ Designed the FlexTree index structure to support efficient data shifting operations.
 - · Implemented FlexSpace, a persistent storage engine that provides a flexible address space
 - Built FlexDB, a highly performant concurrent key-value store on top of FlexSpace.
 - Open source: **O** https://github.com/flexible-address-space
- REMIX: Efficient Range Query for LSM-Trees [3]
 - Participated in the design, implementation and evaluation of REMIX and RemixDB.

Research Assistant, ShanghaiTech University

07/2017 - 12/2019

Advisor: Dr. Shu Yin (ShanghaiTech University and Chinese Academy of Science)

- Flat Indexing in KV-embedded File Systems [4]
- Parallel User-level File System Performance Optimization [5, 6, 7]

Publications

[1] **Chen Chen**, Xingbo Wu, Wenshao Zhong, and Jakob Eriksson. Fast Abort-freedom for Deterministic Transactions. In *IEEE 38th International Parallel & Distributed Processing Symposium*

(IPDPS '24), San Francisco, USA, May 2024.

[2] **Chen Chen**, Wenshao Zhong, and Xingbo Wu. Building an efficient key-value store in a flexible address space. In *Proceedings of the Seventeenth European Conference on Computer Systems* (*EuroSys* '22), page 51–68, Rennes, France, 2022.

[3] Wenshao Zhong, **Chen Chen**, Xingbo Wu, and Song Jiang. REMIX: Efficient Range Query for LSM-trees. In *19th USENIX Conference on File and Storage Technologies (FAST '21)*, pages 51–64, February 2021.

[4] **Chen Chen**, Tongliang Deng, Jian Zhang, Yanliang Zou, Xiaomin Zhu, and Shu Yin. FILT: Optimizing KV-Embedded File Systems through Flat Indexing. In 2020 *IEEE 40th International Conference on Distributed Computing Systems (ICDCS '20)*, pages 1203–1204, November 2020.

[5] **Chen Chen**, Jianzhong Liu, Yanliang Zou, Tongliang Deng, Xiaomin Zhu, and Shu Yin. A Case Study on the Efficiency of User-Level Parallel File Systems. In *IEEE 21st International Conference on High Performance Computing and Communications (HPCC '19)*, pages 90–97, August 2019.

[6] Yanliang Zou, **Chen Chen**, Tongliang Deng, Jian Zhang, Si Chen, Xiaomin Zhu, and Shu Yin. SHC: A Method for Stackable Parallel File Systems in Userspace. In *IEEE 21st International Conference on High Performance Computing and Communications (HPCC '19)*, pages 1374–1381, August 2019.

[7] Yanliang Zou, **Chen Chen**, Tongliang Deng, Jian Zhang, Xiaomin Zhu, Si Chen, and Shu Yin. User-level parallel file system: Case studies and performance optimizations. *Concurrency and Computation: Practice and Experience*, 34(13):e6905, 2022.

[8] Noaman Ahmad, **Chen Chen**, Ben Baenen, and Jakob Eriksson. A Scalable Memory-Safe Delegation System. *In Submission*, 2024.

[9] Wenshao Zhong, Xingbo Wu, **Chen Chen**, and Jakob Eriksson. I/O-efficient Indexing for LSM-Trees. *In Submission*, 2024.

[10] **Chen Chen** and Jakob Eriksson. Enhancing In-Memory Data Structures via Adaptive Concurrency Control. *Manuscript in Preparation*, 2024.

Teaching Experience

• Graduate Teaching Assistant, University of Illinois at Chicago

- Facilitated the academic growth of more than 200 students each semester through 6 consecutive semesters as a dedicated teaching assistant.
- Hosted weekly in-person recitation sessions for Computer Organization class.
- Managed public Linux server that supports class infrastructures.
- Undergraduate Teaching Assistant, ShanghaiTech University
 - Instructed and supported students in the following classes through 2016 to 2018:
 - (1) Intro to Information Science (2) Computer Programming (3) Operating Systems

Awards

- IEEE TCPP Student Travel Award, 2024
- ACM EuroSys Student Travel Grant, 2022
- ASC Student Supercomputer Challenge Silver Medal (runner-up), ePrize Award, 2018 Media Coverage: [HPCwire 1] [HPCwire 2] [ShanghaiTech]
- ISC Student Cluster Competition Finalist Rank 4, Highest HPCG Score, 2018 Media Coverage: [HPCwire 1] [HPCwire 2]