

Huanchen Zhang

Carnegie Mellon University
Department of Computer Science
Gates-Hillman Center 9023
5000 Forbes Ave, Pittsburgh, PA 15213

Email: huanche1@cs.cmu.edu
Web: <http://www.cs.cmu.edu/~huanche1/>
Cell: +1-626-688-5907
Twitter: @huanchenzhang

Research Interest

My research interests center on computer systems and databases. I have a particular interest in designing space-efficient and high-performance in-memory data structures.

Education

Carnegie Mellon University Sept. 2013 - Feb. 2020
Ph.D. in Computer Science Department *Pittsburgh, PA*
Thesis: “Memory-Efficient Search Trees for Database Management Systems”
Advisor: David G. Andersen

University of Wisconsin - Madison June 2009 - May 2013
B.S. in Computer Engineering, with *Distinctive Scholastic Achievement* *Madison, WI*
Second Major: Computer Sciences
Third Major: Mathematics
GPA: 4.0
Advisor: Remzi Arpaci-Dusseau

Research Experience

Graduate Research Assistant Sept. 2013 - Present
Carnegie Mellon University *Pittsburgh, PA*

- **Order-Preserving Search Tree Key Compression**

Designed and implemented the High-speed Order-Preserving Encoder (HOPE) for in-memory search trees. HOPE is a fast dictionary-based compressor that encodes arbitrary keys while preserving their order. HOPE allows database search trees to achieve lower query latency and better memory efficiency.

- **Succinct Range Filters (SuRF)**

Designed and implemented the first practical and general-purpose range filter called SuRF. SuRF can significantly reduce unnecessary I/Os in a database. It uses succinct encodings to achieve space that is close to the theoretic minimum while retaining high performance. SuRF won the Best Paper Award in SIGMOD’18 and is being implemented by companies such as Facebook, Alibaba, and Salesforce in their systems.

- **Hybrid Indexes**

Proposed a dual-stage index-building framework called hybrid index for in-memory OLTP databases that can substantially reduce the storage overhead of existing index structures with minimal performance penalty. We applied our semi-automatic dual-stage transformation (DST) to four different index data structures and demonstrated the effectiveness and generality of the hybrid index method.

- **OpenBw-tree**

Worked with a fellow CMU student on implementing an open-source version of Microsoft’s Bw-Tree. We clarified missing points in the original paper that are required for a correct implementation. We found that the lock-free Bw-Tree design does not perform as well as lock-based concurrent indexes for in-memory workloads.

- **Raft Consensus Algorithm Verification**

Used formal verification methods (TLA+) to prove safety properties of the Raft consensus algorithm. I found an important safety violation in Raft’s cluster membership change protocol and worked with the authors (Diego Ongaro and John Ousterhout) to fix the bug. The solution is well-received by the Raft development community and was subsequently incorporated into Diego’s thesis.

Research Intern
HP Labs

May 2016 - August 2016
Palo Alto, CA

- **Adaptive Concurrency Control**

Designed an efficient concurrency control protocol for “The Machine” architecture, where there is a large pool of non-volatile memory (NVM) at the center of the architecture shared by thousands of computing devices. The key idea is to adaptively choose between the optimistic concurrency control (OCC) and the multi-version concurrency control (MVCC) for each data slot in the NVM pool based on the execution history.

- **Metadata Storage and Indexing**

Designed a metadata storage engine that is both scalable and efficient for “The Machine”-like architecture. The key idea is to use multiple storage tiers to store metadata cost-effectively as it transitions from hot (i.e., actively accessed) to warm, to cold (i.e., kept primarily for historical and regulatory purposes). The data structures used range from lock-free concurrent radix trees to dictionary-compressed record lists.

Undergraduate Research Assistant
University of Wisconsin - Madison

Sept. 2011 - May 2013
Madison, WI

- **File System Checker**

Explored how file system checker programs behave on damaged file systems by constructing corruption models and analyzing repairing results. Helped develop fast file system checkers that support disk-order scan over dynamically allocated metadata on ext3 and XFS.

Awards & Honors

- Best Paper Award, ACM SIGMOD 2018
- 2018 ACM SIGMOD Research Highlight Award
- First Place Winner, Pittsburgh Go Tournament, 2016
- Student Travel Grant, ACM SOSP 2015, USENIX NSDI 2014
- UW-Madison Graduate with Distinctive Scholastic Achievement, 2013
- Edgar H. and Laverne R. Krainer Memorial Scholarship, 2011 & 2012
- UW-Madison Computer Sciences Summer Research Assistant Award, 2012
- The John and Elizabeth Moore Award for Excellence in General Chemistry, 2011

Publications

- [1] **Huanchen Zhang**, Lily Liu, David G. Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. “Order-Preserving Key Compression for In-Memory Search Trees” To appear in: *Proceedings of the 2020 ACM International Conference on Management of Data (SIGMOD’20)*. June 2020.
- [2] **Huanchen Zhang**, Hyeontaek Lim, Viktor Leis, David G. Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. “Succinct Range Filters” To appear in *ACM Transactions on Database Systems (TODS)*. 2020. (Invited Submission)
- [3] **Huanchen Zhang**, Hyeontaek Lim, Viktor Leis, David G. Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. “Succinct Range Filters” *ACM SIGMOD Record*, 48(1), pp.78-85.
- [4] **Huanchen Zhang**, Hyeontaek Lim, Viktor Leis, David G. Andersen, Michael Kaminsky, Kimberly Keeton, and Andrew Pavlo. “SuRF: Practical Range Query Filtering with Fast Succinct Tries” In: *Proceedings of the 2018 ACM International Conference on Management of Data (SIGMOD’18)*. June 2018, pp. 323–336.
Best Paper Award (1 out of 90 accepted papers)
- [5] Ziqi Wang, Andrew Pavlo, Hyeontaek Lim, Viktor Leis, **Huanchen Zhang**, Michael Kaminsky, and David G. Andersen. “Building a Bw-Tree Takes More Than Just Buzz Words”. In: *Proceedings of the 2018 ACM International Conference on Management of Data (SIGMOD’18)*. June 2018, pp. 473–488.

- [6] **Huanchen Zhang**, David G. Andersen, Andrew Pavlo, Michael Kaminsky, Lin Ma, and Rui Shen. “Reducing the Storage Overhead of Main-Memory OLTP Databases with Hybrid Indexes”. In: *Proceedings of the 2016 International Conference on Management of Data (SIGMOD’16)*. June 2016, pp. 1567–1581.

Non Peer-Reviewed

- [7] **Huanchen Zhang**. “The End of the x86 Dominance in Databases?” Abstract. In: *Conference on Innovative Data Systems Research (CIDR’19)*. 2019.

Patents

- [8] **Huanchen Zhang** and Kimberly Keeton. “Data Storage over Immutable and Mutable Data Stages”. Filed Sept. 2017, Granted Sept. 2019.
- [9] **Huanchen Zhang** and Kimberly Keeton. “Changing Concurrency Control Modes”. Filed May 2017. Patent Pending.

Service

- **Program Committee – SIGMOD 2020**
- Journal Reviewer – KAIS 2018

Teaching

- Head TA & Guest Lecturer – Advanced OS and Distributed Systems (CMU 15-712) – Fall 2017
- Head TA – Database Applications (CMU 15-415/615) – Fall 2016

Mentoring

Project Advisor

- Lily (Xiaoxuan) Liu (CMU M.S. 2018)
- Mengxi Chen (CMU M.S. 2017)
- Runshen Zhu (CMU M.S. 2016)
- Jiexi Lin (CMU M.S. 2016)
- Jianhong Li (CMU M.S. 2016)
- Rui Shen (CMU M.S. 2016)
- Lin Ma (Peking University B.S. 2015)

CMU CSD Ph.D. Mentoring Program

- Jalani Williams
- Han Zhang

Talks

- “Memory-Efficient Search Trees for Database Management Systems.”
 - Snowflake Inc, January 2020
 - Salesforce, January 2020
 - Tsinghua University IIIS, October 2019
 - Thesis Defense, October 2019
- “SuRF: Practical Range Query Filtering with Fast Succinct Tries.”
 - CMU Network Group Seminar, November 2018
 - Parallel Data Lab Annual Retreat, October 2018
 - SIGMOD’18 Conference Talk, June 2018
 - CMU Parallel Data Lab Seminar, May 2018
 - CMU Database Group Seminar, May 2018

- “Towards Space-Efficient High-Performance In-Memory Search Structures”
 - Thesis Proposal, April 2018
- “Succinct Trie Indexes Made Practical.”
 - CMU Database Group Seminar, February 2017
- “Distributed Metadata Store for RePO”
 - HP Labs, August 2016
- “Reducing the Storage Overhead of Main-Memory OLTP Databases with Hybrid Indexes.”
 - Parallel Data Lab Annual Retreat, October 2016
 - SIGMOD’16 Conference Talk, June 2016
 - UC-Berkeley Database Group Seminar, June 2016
 - HP Labs, May 2016
 - Parallel Data Lab Annual Retreat, October 2015

References

David G. Andersen

Professor of Computer Science Department
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
Email: dga@cs.cmu.edu

Andrew Pavlo

Assistant Professor of Computer Science Department
Carnegie Mellon University
5000 Forbes Ave, Pittsburgh, PA 15213
Email: pavlo@cs.cmu.edu

Michael Kaminsky

Senior Research Scientist
Intel Labs
4720 Forbes Ave, Pittsburgh, PA 15213
Email: michael.e.kaminsky@intel.com

Kimberly Keeton

Distinguished Technologist
Hewlett Packard Labs
1501 Page Mill Rd., MS 1123 Palo Alto, CA 94304
Email: kimberly.keeton@hpe.com