

Hongbo Zhang

☎ (+1) 607-431-3346 | ✉ hongbo@cs.cornell.edu | 🏠 https://cs.cornell.edu/~hongbo | 📄 0x486F626F

Education

Cornell University

Ph.D. of Computer Science

Sept 2019 - Present

Ithaca, NY

- Teaching Assistant: CS 3410 Computer System Organization and Programming (Fall 2019)

University of Waterloo

Bachelor of Mathematics in Honours Computer Science with Digital Hardware and Honours Statistics

May 2014 - Apr 2019

Waterloo, ON

- Co-operative Program, a five-year undergraduate program that combines academic studies with internships.
- University of Waterloo President's Scholarship.

Publications

Carousel: Low-Latency Transaction Processing for Globally-Distributed Data

Xinan Yan, Linguan Yang, **Hongbo Zhang**, Xiayue Charles Lin, Bernard Wong, Kenneth Salem, Tim Brecht
In Proceedings of the ACM International Conference on Management of Data (**SIGMOD 2018**)

June 2018

Houston, TX

Research

CedrusDB

Supervisor: Robbert van Renesse, Emin Gün Sirer

Aug 2019 - Present

- Memory-mapped key-value store

Single-Level Storage System

Supervisor: Ali José Mashtizadeh

Mar 2018 - Apr 2019

- Low-overhead continuous incremental process checkpointing achieved by copy-on-write shadow objects.
- Built the prototype system by implementing a FreeBSD kernel module in C.

Carousel: Geo-Distributed Transactional Data System

Supervisor: Bernard Wong

May 2017 - Feb 2018

- Geo-distributed database system offering low-latency processing for multi-partition transactions.
- Requires At most two wide-area network roundtrips needed to commit a transaction in the absence of failures.
- Built the prototype system in Go and evaluated with Retwis and YCSB+T workload on AWS.
- Carousel reduces 95th percentile transaction latency by approximately 30% compared to TAPIR.

Optimizing Shortest Path Algorithm in Graph Databases

Supervisor: Semih Salihoglu

Sept 2016 - Dec 2016

- Optimized shortest path queries in graph database by pruning out-of-bound vertices on search tree.
- Implemented optimizations in Neo4j (Java) and ran experiments on large graph dataset.

Text Compression Algorithm

Supervisor: Alejandro López-Ortiz

May 2015 - Aug 2015

- Implemented a dictionary-based text compression algorithm in C++.
- Built a data structure that is optimized for fuzzy string matching.

Work Experience

Remind

Backend Performance Engineer

May 2016 - Aug 2016

San Francisco, CA

- Refactored the daily data analysis system and made it 5x faster.
- Built pipelines for data aggregation with the Luigi framework.
- Migrated authentication service to a new token system to reduce database I/O.

Autodesk Research

Software Developer

Sept 2015 - Dec 2015

Toronto, ON

- Worked on Project Dreamcatcher in the Computational Science Research Group.
- Evaluated several physics engines and scientific simulators for multi-body dynamics simulation.
- Developed simulator and analysis tools with Modelica and Bullet.

Avidbots

Software Developer

Jan 2015 - Apr 2015

Kitchener, ON

- Designed and implemented their first path planning algorithm for floor cleaning robots from scratch.
- Developed simulator for validations and visualization.
- Integrated the path planning algorithm into the Robot Operating System to drive real robots.

Awards

2nd Place, ACM-International Collegiate Programming Contest

Nov 2014

Central East North America Regional, Windsor Site

- In ACM-ICPC, teams of three intensively solve 10-12 algorithmic problems in 5 hours.
- Ranked 6th in UWaterloo solo qualifier before joining the UWaterloo Team.

Projects

6502 Computer in Verilog

2018

<https://github.com/0x486F626F/6502Com>

- Built a 6502 processor, RAM, and NES-like Cartridge in Verilog.
- The computer executes 6502 instructions contained in a cartridge, initialized by an iNES format ROM file

Agricultural Product Information System

2017

- Designed and built a system for agricultural manufacturers to host product information.
- Customers use the system to access user manuals and detect counterfeit goods.
- The system is currently hosting more than two million products with 5k visit per day.

Pi-Thermostat

2017

<http://static.hongbozhang.me/img/pi-thermostat.jpg>

- Built a smart, energy-efficient, and remote-controllable thermostat with Raspberry Pi.
- Deployed multiple temperature sensors in different rooms to measure temperature more accurately.

Pyusic

2017

<https://github.com/0x486F626F/pyusic-api> **Demo:** <http://music.hongbozhang.me>

- Not satisfied by the official YouTube App, I built my own web-app to playback audio tracks of YouTube music. It supports in-background, screen-off, and ad-free playing on my phone.

Image Super-Resolution

2017

<http://static.hongbozhang.me/doc/ISR.pdf>

- A study of perceptually effective image super-resolution models for single-image super resolution (SISR).
- Evaluated four SISR methods: interpolation, sparse coding, convolutional neural network (CNN), and generative adversarial network (GAN) in terms of training efficiency and perceptual quality

Quora Question Paris Challenge

2017

<http://static.hongbozhang.me/doc/Quora.pdf>

- A natural language processing (NLP) competition on identifying duplicate questions on Quora.
- Proposed a Siamese LSTM structure and developed our classifier. Ranked top 8% at submission

Looprac

2016

<https://devpost.com/software/looprac-gvui2s>

- Carpool matching app that lets drivers and passengers find each other easily.
- Worked on matching algorithm, API services, and data warehouse.

Scheduloo

2015

<https://devpost.com/software/scheduloo-web>

- Course scheduling tool for UWaterloo which provides the optimal schedule by user's preferences.

Miscellaneous

Undergraduate Research Opportunities Conference (UROC)

Oct 2016

- 50 top undergraduate students from across Canada were invited to attend this research workshop.
- Investigated counterexamples to the conjecture of compatible triangulations with Prof. Anna Lubiw.
- Report published at <https://arxiv.org/abs/1612.04861>