

Zeyuan Hu

Homepage: <https://zhu45.org/>

Email: ferrishu3886@gmail.com

EDUCATION

University of Texas

Austin, TX

Sept 2017 – May 2019

- M.S. in Computer Science. (GPA: 3.88/4.00)
- Selected Coursework: Data Centers, Distributed Systems, Advanced Operating Systems, Structured Models in NLP, Machine Learning, Natural Language Processing, Automated Logic Reasoning, Algorithms

University of Wisconsin

Madison, WI

Sept 2010 – Dec 2014

- B.A. in Computer Science. (GPA: 3.74/4.00)
- B.A. in Economics with Honors. (GPA: 3.85/4.00)
- B.A. in Mathematics. (GPA: 3.81/4.00)
- Recipient of 2013 Honors Summer Sophomore Research Apprenticeship
- Recipient of 2012 Meek Bishop Scholarship in Economics, *top 2 out of 500 economics major students*

WORK EXPERIENCE

Software Engineer Internship

Schlumberger

May 2018 – August 2018

HPC infrastructure team

- Implemented a monitoring component of the in-house High-Performance Computing (HPC) engine in C++ to provide the fault tolerance and handle the “straggler” problem
- Employed SGD algorithm to dynamically learn the best timing for backup executions of the in-progress tasks based on the computation task characteristics
- Built a C++ code generator that automatically generates the application layer code based on the engine API

Software Engineer

IBM

August 2015 – August 2017

Db2 LUW federation team

- Constructed Hive and Impala wrappers with C++ and Java to support federation database between traditional RDBMS and Hadoop-based data warehouse solution
- Created automated setup tools with Shell that reduce product configuration time by 75%
- Enhanced server option optimization tools using C to reduce federation database performance tuning time by 90 % and enable the capability of tuning the product against Hive, Impala, and Spark
- Resolved over 20 defects, including a severe memory leak issue that impacted a \$1.6 million deal. *Awarded IBM Manager’s Choice Award 2016*

SELECTED PROJECTS

- **RustFS** (2018 -). Building a user-space file system that leverages NVMe SSD. Rust, SPDK
- **Strata with Lease** (2018). Extended Strata file system with Lease mechanism to support concurrent file access across processes. C.
- **HyperPebblesDB** (2018). Constructed a key-value store that is part of LevelDB family with focus on reducing write amplification. C++, CMake, Autotools
- **Distributed Key-Value Store** (2018). Built a distributed key-value store with Python that uses *eventually consistency* model with two session guarantees: *Read Your Writes* and *Monotonic Reads*.

LANGUAGES AND TECHNOLOGIES

- **Languages:** C++, C, Python, Rust, Shell, SQL, Java, Lisp, MATLAB
- **Software:** CMake, Autotools, QEMU, Docker, Tensorflow, Keras, Git, ClearCase, Hive, Impala, Maven, Hadoop