

Jing Yang

Redmond WA 98052
(+1) 360 809 1109

ditsing@gmail.com
github.com/ditsing
[LinkedIn](#)

OBJECTIVE

Proven team leader looking for team-founding opportunities.

EXPERIENCE

Senior Software Engineer Google Kirkland
Persistent Disk - Disks in every VM 2020 - Present

- Led 5 engineers on the Pod traffic migration project, which moves the entire Persistent Disk backend stack to the new service platform. PD is made by a team of 400 engineers and serves all disks on Google Cloud, from small boot disks to hyper performant data disks. Took a deep dive into 20+ services, designed tailor-made migration strategy for each service, emphasizing consistency, data safety and service continuity. Scoped the effort and design, ramped up team members, delegated the majority of the work and managed the project end-to-end.
- Owns PD's qualification and integration with the Google Cloud virtual machine team.
- Previously worked on creating an industry-leading feature that replicates disks consistently across cloud regions.

Senior Software Engineer Google Seattle
Panels - Anti-fraud with big data 2019 - 2020

Worked on an anti-fraud project that removed a large amount of false data from a valuable data set. Applied cutting edge technology in adjacent research areas on the huge amount of data collected under explicit consent and compensation. The model was later validated by manual reviews and demonstrated high accuracy.

Senior Software Engineer Google Seattle
Blogger - Web and large scale data migration 2015 - 2019

- Led several multiple-quarter projects, including large scale data migrations, new feature developments and improvements to anti-abuse techniques. Duties including oncall for the high traffic web service, technology evolutions and other general backend-related tasks. Blogger is one of the earliest and largest blog-publishing platforms.
- Led two engineers on the Facepile Data Migration project. With zero down time, we migrated 300 million data entries away from the deprecated storage infrastructure. Improved service quality and wiped all issues related to the widget from the "Top Reported Issues" list.

Software Engineer Google Sydney
AppStats - Realtime timeseries database 2013 - 2014

- Worked on a timeseries database that serves millions of queries per second from Blogger, Google Photos, Maps and other Google projects, at close-to-real-time latency. As a core contributor, made improvements end-to-end throughout the service, from building and testing infrastructure, wipeout pipelines to effectiveness of the master election algorithm.
- Designed a string prefix matching algorithm, improved the processing speed of a MapReduce pipeline by a factor of 20. The system can match millions of patterns against billions of strings.

Software Engineer Intern Facebook
Evergreen - eco-friendly data centers 2012

Developed an internal tool that predicts server power consumption based on existing mathematical models. Implemented a separate calculation engine for later reuse and a friendly web UI.

SKILLS

- Large scale data migrations
- Mentoring team members
- **Software engineering**: planning, testing, qualification and releasing
- Deep understanding of **distributed systems** and **the Raft consensus protocol**
- Proficient in **concurrent programming**: locks, atomic primitives and communication in multi-thread environments
- Proficient in **Java** (6 years), **C/C++** (6 years) and **Linux** (10+ years)
- Proficient in **SQL** and **Rust** (10000 lines)
- Proficient in **Python** and **Ruby**

EDUCATION

Harbin Institute of Technology Bachelor, Computer Science	Harbin, China GPA: top 5%	09/2009 - 07/2013
---	------------------------------	-------------------

PROJECT HIGHLIGHTS

Personal Projects 01/2019 - Present

- **Ruافت**, an implementation of the Raft consensus protocol in Rust. Carefully architected to minimize response latency while being resilient to tough network environments. Handles 20000 queries per second in an ideal environment.
- **Shadowrocks**, a lightweight stealth proxy to bypass hostile firewalls. Greatly reduced attack surface by slightly modifying the original shadowsocks specification.

MapReduce Pipeline Optimization 04/2014 - 11/2014

Designed a string prefix matching algorithm for the critical path of the pipeline. Improved the processing speed of a MapReduce pipeline by a factor of 20. Enhanced with the algorithm and other architecture improvements, the new pipeline is able to process 1.5 years' worth of legacy requests within a single day, matching millions of requests against tens of billions of keys in each run.

HONORS AND AWARDS

Gold medal in ACM-ICPC (International Collegiate Programming Contest) China Fuzhou Invitational Programming Contest, 2011.

Silver medals in ACM-ICPC Asia Regional Contest Dalian and Shanghai Site, 2011.

Silver medals in ACM-ICPC Asia Regional Contest Harbin and Chengdu Site, 2010.

Scholarship of State, 2010. Top 4.