

Jing Yang

ditsing@gmail.com
(+1) 360 809 1109
Sammamish, WA 98029

ditsing.com
github.com/ditsing
linkedin.com/in/ditsing

OBJECTIVE

To find the next area to deep dive into.

EXPERIENCE

Tech Lead

Persistent Disk - Reliability and release

Google Kirkland

2020 - Present

- Led 8 engineers on the reliability team. Replicated the entire Persistent Disk stack on a modern service platform and moved over all traffic. Over the decade, PD accumulated 40+ services developed by hundreds of engineers with a variety of design, architecture and guarantees. We adapted, improved and migrated all of the services in 2 years.
- Took a deep dive into each of PD's services. Created tailor-made migration mechanism individually, emphasising simplicity, consistency and service continuity. Solved challenges around configuration-parity, controlled failover, split brain, capacity planning and scaling. Paid off numerous tech debt in the process.
Scoped the project, ramped up team members, delegated the majority of the work and oversaw the project end-to-end. Migrated live I/O from millions of VMs across the globe with minimal customer-visible interruptions.
- Served as the diplomat between PD and one of its neighbours, the Google Cloud virtual machine team, owning PD's qualification and release in VMs.
- Previously worked on building an industry-leading feature that replicates disks across cloud zones with bounded staleness.

Senior Software Engineer

Blogger - Web and large scale data migration

Google Sydney

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

AppStats - Realtime timeseries database

Google Sydney

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

Evergreen - eco-friendly data centers

Facebook

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 **consensus protocols**
- 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	Harbin, China	09/2009 - 07/2013
Bachelor, Computer Science	GPA: top 5%	

PROJECT HIGHLIGHTS

Personal Projects 01/2019 - Present

- [Ruafit](#), 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.