

XUANHAO (ERIC) ZHANG

◇ (+1) 617-447-9107 ◇ xuanhao0831@gmail.com ◇ <https://www.xh-zhang.com>

WORK EXPERIENCE

Google Inc.

Software Engineer III (Persistent Disk Team)

Kirkland, WA

Nov. 2022 - Present

- Owned the resource and permission areas of PD Podification, that decoupled 40+ backend jobs from the legacy monolithic system to the new server platform to increase release velocity.
- Led multiple projects during the pod migration, including colossus spindles capacity/permission, jobs' reshaping, and modernizing Autocap/Autopilot integration.
- Recognized the resource (GCU/RAM) efficiency problems during migration. Started from a 20% project and then collaborated with SREs to propose, plan and land a multi-quarter project to improve PD's CPU/RAM resources efficiency. It included CPU profile/scheduler analysis, memory analysis, high throughput/IOPS bursting workload analysis, vertical/horizontal auto-scale fine-tuning, as well as benchmarks.
- The PD resource efficiency project improved around 20% of the resource utilization across 300+ cloud locations with a projected 5 year savings of \$XX million without hurting SLOs. The results also got a Google Perfy Silver Award.
- Led a project to design and implement a mechanism to filter out the alerts from known device corruptions, significantly reducing the SRE on-call toil.

Google Inc.

Software Engineer II (Persistent Disk Team)

Kirkland, WA

Jul. 2021 - Nov. 2022

- Owned the PD silent data corruption (SDC) area. Identified potential SDC locations across PD and introduced new metrics, monitoring and alerts to improve observability and debuggability.
- Designed and implemented the disk metadata journaling for the PD cross-cell migration case. Utilized Flume to handle large-scale of IOs from Spanner to Colossus for both journaling and recovery pipelines. It improved the disk metadata reliability and abilities of disaster recovery.

Synopsys Inc.

Software Engineer Co-op (Methodology Core Team)

Boxborough, MA

Jan. 2020 - Aug. 2020

- Developed a Machine Learning Platform from scratch, which helps the circuit designers reserve the optimal cloud resources to run simulations.
- Built the data pipeline and ETL scripts with PySpark, and trained a MixNet model using PyTorch.
- Prototyped and benchmarked the performance of the deployed ML system, including data collection, evaluation and retraining system.
- Saved average 60% on computing resources, which prevent millions of dollar wastage on the CPU clusters. The results were presented at the Design Automation Conference 2020.

TECHNICAL SKILLS

Distributed System

C++, Python, Performance tuning and optimization, large scale deployment

ML&Robotics

PySpark, PyTorch, CUDA programming, ROS, CAD, 3D Printing

EDUCATION

Northeastern University, Boston, MA

Sep. 2018 - Dec. 2020

Master of Science in Electrical & Computer Engineering, GPA 3.84/4.0

Southeast University, Nanjing, China

Sep. 2013 - Jun. 2017

Bachelor of Engineering in Information Science and Technology, GPA 3.67/4.0.

RELATED PROJECT

Robotic Arm to Automatically Changing 3D Printer Plates

Dec. 2024 - Present

- Design and implement a robotic arm system capable of automatically removing finished 3D print plates (for Bambulab A1C) and replacing them with new ones.
- Implement ROS nodes for motion planning, sensor data processing, and error handling, ensuring robust and reliable system operation.
- Started from self-implemented inverse kinematics algorithms to leverage MoveIt! within ROS to generate smooth and collision-free trajectories for the robotic arm. Working on integrating LLM APIs to the system.