Yuvaraj Chesetti

Ph.D Student, University of Utah, Salt Lake City, UT, USA
E-Mail: chesetti@cs.utah.edu
Phone: (+1) 385-444-6659
Website: https://ykchesetti.com
Github: https://github.com/droidkid

EDUCATION

University of Utah, Salt Lake City, Utah — July 2022 — Present
Ph.D, Computer Science

Birla Institute of Technology, Mersa, India — June 2012 — April 2016
Bachelor of Science, Computer Science

INDUSTRY EXPERIENCE

Google — India

Software Engineer L4, Unified Fleet Optimization (UFO) EngProd, Bangalore — *Nov 2019 - July 2021*

- Part of a team that aimed to develop an automated missing test suggestion tool to reduce coverage gaps in large integration tests by cross-referencing production traces with test traces across the entire UFO team, which is an organization of roughly a thousand engineers.
- Delivered infrastructure to sample production traces from production and testing logs and report missing traces.
- Played a pivotal role in team bootstrapping through mentoring and delivering technical talks.

Software Engineer L4, xGA Search, Bangalore — *Nov 2018 - Nov 2019*

- Led a team of 2 in designing and implementing a curation pipeline used in delivering short videos to the Google Discover Feed.

Software Engineer L3-L4, Google Pay, Hyderabad — *July 2016 - Nov 2029*

- Member of the launch team of Tez (now rebranded as Google Pay India).
- Led the development of load testing infrastructure, allowing developers to conduct scalability tests before production launches, resulting in more stable releases. Succesfully integrated the load testing infrastructure into the release pipeline, resulting in preventing several releases that would have caused outages.
- Implemented Backend APIs for a new payment flow and performance improvements.

Media.Net, Directi — Bangalore, India

Internship — April 2015 — June 2015

Worked on implementing a Redis cache layer to pool database connections.

RESEARCH PROJECTS

Learned Indexes for Databases — *Summer 2023-Present*

- Conducting research on designing new algorithms to improve the performance of key-value stores in in external memory using learned indexes, These learned indexes employ machine learning techniques to model data distribution, and our primary focus is leveraging these models to improve core database operations, including queries, updates, joins, and merges.
- Implementing and testing these techniques specifically in Log Structured Merge (LSM) Tree-based key-value stores. Actively prototyping these methods within production-level LSM Tree key-value stores (LevelDB).
- Currently being advised by Prof. Prashant Pandey as a Ph.D. Student.

Optimizing 2D graphics using E-Graphs — *Fall 2022 - Spring 2023*

- Implemented a proof-of-concept to optimize memory usage in the 2D graphics backend library (Skia) used by Google Chrome, utilizing Egg—a high-performance rule-based rewrite system based on E-Graphs.
- Conducted the project as a Graduate Research Assistant, funded by Prof. Pavel Panchekha.

Transpiler Fuzzing — *Spring 2023*

- Prototyped applying state-of-the-art grammar-based compiler fuzzing techniques to ensure the correctness of Rust to C transpilers (C2Rust).
- Completed setting up an end-to-end project as an independent study under the guidance of Prof. Stefan Nagy.