

# Mohil Patel

✉ mohilp1998@gmail.com | 🏠 mohilp1998.github.io | 🌐 mohilp1998 | in mohilp1998

## Education

### University of Wisconsin-Madison

Madison, WI, USA

M.S. in Computer Science | GPA: 4.0/4.0

Sep. 2021 - May 2023

**Coursework:** Distributed Systems, Big Data Systems, Advanced Operating Systems, Machine Learning, Database Management Systems

### Indian Institute of Technology Bombay

Mumbai, India

B.Tech with Honors in Electrical Engineering | Minor in Computer Science | GPA: 9.58/10.0

Jul. 2016 - May 2020

**Coursework:** Data Structures & Algorithms, Computer Architecture, Digital Image Processing, Computer Networks, Probability

## Experience

### Oracle

Redwood City, CA, USA

SOFTWARE ENGINEER | GOLDENGATE - DATABASE ORG

July 2023 - Present

- Part of Oracle's GoldenGate team, responsible for feature development, bug fixes and maintenance of **GoldenGate** product (C++)
- GoldenGate allows customers to setup **real-time database replication** pipelines across multiple databases instances & big data systems

### Samsung Semiconductor

San Jose, CA, USA

SYSTEM SOFTWARE INTERN | GPU DRIVER

May 2022 - Aug. 2022

- As intern in GPU SW team worked on the **ANGLE** project, which translates **OpenGL ES** API calls to **Vulkan** API calls at runtime in smartphones
- Developed methods to profile GPU memory usage using **Vulkan** Extension & implemented memory optimizations in **ANGLE** codebase (C++)

### Nvidia

Hyderabad, India

SOFTWARE ENGINEER | GEFORCE NOW - CLOUD GAMING SERVICE

Jul. 2020 - Jul. 2021

- Member of Nvidia's Cloud Game Streaming **QoS team**, handling **real-time algorithms** to enhance the gameplay experience dynamically
- Ran experiments & analyzed results (**python**) to understand the **impact of networks parameters on cloud game streaming**
- Implemented new algorithms & optimized existing features (C++) to improve user experience by reducing stutter, latency & packet loss

## Projects

### Low Bandwidth Google File System

COURSE PROJECT | CS739: DISTRIBUTED SYSTEMS | [REPORT](#)

Mar. 2023 - May 2023

- Improved design of GFS for Low Bandwidth conditions by **identifying duplicate blocks** on client & server to **avoid redundant data transfer**
- Implemented solution as a wrapper on top of **HDFS** & showed significant reduction in data transmitted over network for specific workloads

### Efficient Distributed Transfer Learning using Pipelined Model Parallelism

COURSE PROJECT | CS744: BIG DATA SYSTEMS | [REPORT](#)

Oct. 2022 - Dec. 2022

- Designed a system to incorporate pipelined model parallelism in transfer learning & reduced training speed by improving resource utilization
- Implemented system using **pytorch** & **async communication** library, & showed perf. improvement by training multiple **resnet50** in parallel

### Improving Performance in LSM-Tree Based Key-Value Stores using NVMe

COURSE PROJECT | CS736: ADV. OPERATING SYSTEMS | [REPORT](#)

Sept. 2022 - Dec. 2022

- Analyzed the performance of Log-Structured Merge (LSM) Tree's compaction process in **RocksDB** for NVMe SSD & HDD, using **fio** benchmark
- Improved performance by **splitting hot & cold data** between SSD & HDD, and using **SPDK** to **bypass kernel** for direct NVMe reads/writes

### Database to Graph Conversion Tool

RESEARCH PROJECT | TEAM: [MARIUS](#) | [GITHUB REPOSITORY](#)

Oct. 2021 - May 2022

- Designed a tool (**python**) which outputs a graph (as an edge list) from an input database using user-defined configuration and SQL queries
- Implemented **out-of-memory** processing to generate **billions of edges** within **few hours**. Currently supports **Postgres, MySQL & MariaDB**

### CHIP-8 Emulator

PERSONAL PROJECT | [GITHUB REPOSITORY](#)

Jan. 2021 - Mar. 2021

- Chip-8 is a **8-bit interpreted language** with **35 opcodes** & **4KB memory** used in 1970s in many microcontrollers
- Designed a chip-8 emulator with additional functionalities like timer, keyboard & graphics using **C++ and SDL2.0**

## Skills

**Languages:** (*proficient*): C, C++, Python, VHDL,  $\LaTeX$  (*familiar*): Matlab, Java, SQL, Embedded C, HTML

**Tools & Frameworks:** (*proficient*): NumPy, Pandas, Jupyter Notebook, Git (*familiar*): Pytorch, Postgres, Docker, sklearn, SciPy, Spark