MARASANIGE SAMARTH MAHENDRA

Backend & Distributed Systems Engineer — Data Infrastructure

+1 (857) 707-1671 | samarth.mahendragowda@gmail.com | Boston, MA, USA | LinkedIn | GitHub | Portfolio

EDUCATION

Northeastern University — January 2024 - December 2025

Master's, Computer Science

- Courses: Programming Design, Software Engineering, Databases, Computer Systems, Algorithms, NLP/ML

Visvesvaraya Technological University — August 2018 - July 2022

Bachelor's, Computer Science - Dayananda Sagar College of Engineering

SKILLS

Languages: Python, Java, C/C++, SQL, NoSQL
Databases: PostgreSQL, Elasticsearch, MongoDB, Redis, Database Internals
Systems: Linux/Unix, Concurrency & Multithreading, Networking, gRPC
Frameworks & Libraries: Django REST Framework, Node.js
Cloud & DevOps: AWS, Docker, Kubernetes, Observability (Datadog, Prometheus), Celery, Kafka, CI/CD
Algorithms: Solved 400+ LeetCode problems (Top 1%)

PROFESSIONAL EXPERIENCE

Draup (B2B AI/SaaS, AI-Driven Sales, Talent Intelligence Platform) — Bengaluru, KA, India

Associate Software Development Engineer — *August 2022 - November 2023*

- **Owned** and evolved **core backend** and data-platform systems in a **5-engineer team**, building mission-critical, **customer-facing** production workloads using **Python, Django REST Framework, PostgreSQL**, and **Celery-based async pipelines**, supporting **high-traffic, latency-sensitive APIs and data workflows.**
- Led **PostgreSQL → Elasticsearch migration** for **real-time analytics** and applied advanced **query optimization** techniques (partitioning, indexing, materialized views), delivering **5× faster queries**, **400% execution speedups**, and **50% lower operational costs** at scale.
- Designed a **dynamic query framework** and **advanced search** and filtering engine (**Boolean operators, nested conditions)** using **PostgreSQL + Elasticsearch**, with **Redis caching**, improving query performance by **60%** and reducing new entity development time by **80%**.

Associate Software Development Engineer Intern — *April 2022 - July 2022*

- Built **Datadog dashboards**, integrated **AWS** CloudWatch alarms to **monitor platform health**, reducing issue resolution time by 30%.
- Implemented **caching** to improve the efficiency of image requests, resulting in a **70% reduction** in load times.
- Developed **self-running Jenkins jobs** for **database cleanup**, cutting manual effort and improving efficiency by **25%.**

PROJECTS & OUTSIDE EXPERIENCE

ButterDB — High-Performance Key-Value Store (C, Python) — Remote

C Programming** · B-Tree Engine · WAL · Page-Based I/O · **Concurrency — *October 2025 - November 2025*

- Built a **TCP-based key-value database** in C with a persistent, disk-backed **B-Tree storage engine** (**page-level I/O, node splitting**) supporting logarithmic-time lookups.
- Implemented fine-grained **concurrency** using node-level locks, **latch crabbing**, and **WAL synchronization** to ensure correctness, durability required for **synchronous replication** in **distributed** environments.
- **Benchmarked** and **profiled** under concurrent clients (~40K ops/sec, sub-ms p95), identifying contention hotspots and validating crash-recovery guarantees.

Orion — Dependency-Aware Parallel Task Execution Framework (C++) — Boston, MA, USA

- **Designed and implemented a minimal, Ray-style distributed task runtime in modern C++ (C++23),** modeling complex computations as directed acyclic graphs (**DAGs**) to enable seamless, horizontally scalable execution from a **single-process engine** to a **multi-node cluster.**
- **Engineered a distributed cluster scheduler utilizing gRPC and Protocol Buffers (Protobuf) for remote procedure calls (RPC),** integrating dynamic node discovery and cross-node dependency tracking to minimize scheduling latency and dispatch tasks immediately upon data readiness.
- **Built a highly concurrent, thread-safe in-memory object store and worker pool** utilizing **threads**, condition variables, and **mutex-protected queues**; implemented callback-driven unblocking to achieve deterministic **producer-consumer** coordination with zero busy-waiting.

Real-Time AI Chat + Voice Assistant & Intelligent Agent Platform — Boston, MA, USA

Tech stack: OpenAI, Twilio, Celery, FastAPI, Discord, Websockets, Render, Redis, AWS — *April 2025 - May 2025*

- **Architected an agent** system in **Python** integrating **OpenAI GPT-4** + **Google Gemini** with modular tools, **dynamic function calling**, and profile-aware responses via **MongoDB** and **Discord**.
- Built a **scalable async backend** with **FastAPI** + **WebSockets**, deployed on **Render** with a **Celery** worker handling long-running tool calls and **real-time audio** coordination.
- Deployed **live demo** via public phone number **(833) 970-3274** using **Twilio**, showcasing **job-query answering, system prompts with resume context, and cross-platform communication**. Demo @ https://www.samarthmahendra.com/

Stock Market Simulation Application - Java MVC — Boston, MA, USA

- Implemented a **modular, object-oriented trading simulation system** using **MVC and interface-driven architecture**, modeling portfolios, transactions, and **time-series market data** with strict separation of concerns.
- Applied **SOLID principles and design patterns** (Strategy, Adapter, Controller) to support **pluggable investment strategies** (dollar-cost averaging, rule-based rebalancing), **deterministic transaction replay**, and extensible analytics.
- Integrated **external financial data APIs** with caching and **rate-limit control**, implemented **cost-basis and portfolio valuation engines**, and delivered a **fully tested, production-style system** with CLI and **Swing GUI** views.

StackOverflow Full-Stack Q&A Platform — Boston, MA, USA

TypeScript, JavaScript, React.js, Node.js, MongoDB, Cypress, Jest — *February 2025 – April 2025*

- Developed Q&A web application using **React (frontend) and Node.js** with TypeScript (backend). **Architected backend controllers** and models following the **MVC pattern**, leveraging design patterns such as **Facade, Validator, Strategy, and Factory** to ensure **modularity, maintainability, and scalability.**
- Built comprehensive **end-to-end and integration test suites** using **Cypress and Jest**, automating user flows (e.g., posting questions/answers, commenting, voting) and achieving high code coverage and **reliability**.
- Integrated **CI with GitHub Actions, Cypress, and CodeQL** for PR-based linting and coverage checks.

Open Jobs Analytics Platform – Backend Infra + Monitoring — Boston, MA, USA

Puppeteer, Redis, Celery, MongoDB, Grafana, Prometheus, GPT-4o, AWS — *December 2024 - December 2024*

- Designed a **producer-consumer** architecture using **Python**, **Celery**, integrated with **Prometheus** and **Grafana**, achieving **99%** uptime. **Scraped dynamic web pages** with **Playwright** and **Puppeteer**, harvesting **1000+ data points** daily.
- **LLM-powered** CSS **selector extraction** reduced new-site onboarding by **90%.**
- Enhanced stealth capabilities with **random headers**, **user agents, referrer headers**, and **OS configurations**, reducing bot detection by up to **90%.**