

Shashank Sati

shashanksati11@gmail.com — +91-8126067104 — github.com/shankeleven

Education

Jawaharlal Nehru University, Delhi

August 2021 – May 2026

Bachelor+Master in Technology, Electronics and Communication Engineering

Coursework: Data Structures, Design and Analysis of Algorithms, Digital Logic Design, Computer Architecture, Semiconductor Devices, Machine Learning, Deep Learning, Artificial Intelligence

Projects

SevenDB

GitHub

- Built a **distributed, reactive in-memory database** with **Raft-based deterministic subscriptions** using log-embedded emission contracts, surviving rigorous single/multi-node failure benchmarks
- Achieved **123k ops throughput** with **p50 0.11ms / p99 0.35ms** latency; Raft failover completed in **11.4ms (p99 12.7ms)**
- Validated durability via crash-recovery and reconnect benchmarks: achieved **12 ms leader failover (p99)**, **sub-1ms client resume**, **zero lost emissions**, and **at-least-once delivery** across 1k updates
- Engineered Raft-backed primitives including durable **WAL (fsync, pruning, rotation)**, **EMITRECONNECT**, for the client to resume after crashes and **per-bucket metrics**; passed **100-run determinism** tests under leader churn

Khoj

GitHub

- Implemented a **TF-IDF** based local search engine with an optimized inverted index for low-latency developer workflows
- Indexed 859 files in 3.54s (**243 files/sec, 23.1 MB/sec**); sustained **519 QPS** over 10s load tests (5200 queries) with **1.68ms avg latency**

Teach-code

GitHub

- Built a **collaborative code editor** with shared whiteboard (Excalidraw) and **real-time voice chat** for peer study of algorithms
- Implemented **WebSocket-based sync** (Next.js + Golang) achieving **sub-50ms edit propagation**
- Integrated **WebRTC for peer-to-peer voice communication** with backend signaling, achieving **sub-200ms end-to-end latency**

Open Source Contributions

Google FlatBuffers

PR #8666

- Extended **Flatbuffers Compiler (flatc)** to support **gRPC CallbackService**, providing a reactive, non-blocking execution model that reduces thread overhead compared to traditional synchronous and completion-queue architectures
- Generated **callback-based server and client stubs** (reactor-based unary + streaming APIs), with **compile-only validation** ensuring symbol correctness and full compatibility with gRPC's modern async callback model

Microsoft MSSQL

PR #19914

- Developed the **Global Object Search** feature for the VSCode-MSSQL extension, allowing users to search through **tables, views, stored procedures or functions** from the VS Code window itself

KubeVirt

PR #14177

- Implemented native **StartupProbe** support for **Virtual Machine deployment on Kubernetes**, eliminating initialization race conditions
- Built host-guest health verification via **QEMU Guest Agent** using **gRPC over Unix sockets**, accelerating failure detection with dynamic failure thresholds

Achievements

- **2× ICPC Regionalist:** Qualified for the Amritapuri and Kanpur sites for the 2024 ICPC Regionals
- Achieved state-of-the-art performance (**87.3% accuracy**) for **semantic segmentation of Indian legal judgements** using **BiLSTM-CRF**
- Authored **quic.blog**, an interactive **TCP vs QUIC visualizer**, reaching 2k+ unique visitors

Skills

Languages: Go, C++, Python, Rust, JavaScript, SQL

Systems: Distributed Systems, Databases, Concurrency, Caching

ML/DL: Classical ML, Deep Learning, NLP, LLM, BERT, LSTM

Infra/Tools: Linux, Docker, Kubernetes, gRPC, Redis, PostgreSQL, MongoDB, Kafka, AWS(Lambda, EC2, S3)