Ramilla Hrushikesh

+91-9701721767 ramillahrushikesh@gmail.com LinkedIn Github
LeetCode Codeforces CodeChef

Education

Indian Institute of Information Technology and Management **August 2023 – May 2027**
Bachelor of Technology in Mathematics and Scientific Computing *Gwalior, MP*

Relevant Coursework: Data Structures & Algorithms, Operating Systems, Database Management Systems, Computer Networks, Probability & Statistics, Object Oriented Programming.

Achievements & Problem Solving

Competitive Programming Profiles: LeetCode Guardian (**2290 peak**, Top 0.5%), CodeChef 5-Star (**2091 peak**), Codeforces Expert (**1853 peak**).
Contest Performance: Achieved **Global Rank < 30** multiple times in CodeChef contests; consistently solving complex algorithmic challenges under strict time limits.
Hackathons: Finalist in L'Oreal Sustainability Challenge and Walmart Global Hackathon.

Technical Skills

Languages: C++, Python, Java, Go, C, SQL, Bash/Shell Scripting
Systems & Low-Level: Operating Systems, Concurrency, Lock-Free Programming, Atomics, Cache Optimization, File Systems, Networking
Systems & Distributed Concepts: Storage Engines, Concurrency Control, Fault Tolerance, Crash Recovery, Log-Structured Systems
Tools & Frameworks: Git, Linux, Docker, PyTorch, TensorFlow

Projects

NullRing: Ultra-Low Latency C++20 Execution Engine | *C++20, Lock-Free, Atomics* **2025**

- Designed a high-performance execution engine for low-latency task processing in concurrent systems.
- Built a deterministic ultra-low latency execution engine achieving ~**92ns pipeline floor and ~142ns median latency**, eliminating dynamic allocation, syscalls, and virtual dispatch.
- Designed a wait-free SPSC ring buffer with **64-byte cache-line isolation** and acquire/release memory ordering, ensuring zero contention and predictable inter-core communication.
- Analyzed cache behavior to identify inter-core communication bottlenecks and optimized execution to keep data cache-resident.
- Ensured reliability and correctness under concurrent workloads with production-style constraints.

AccretionDB: WiscKey-Style LSM Storage Engine | *C++20, Storage Engines, File Systems* **2026**

- Built a storage engine for efficient large-scale key-value data management.
- Engineered a production-grade LSM storage engine using WiscKey architecture, separating keys and values to reduce write amplification to ~**2.0x**, achieving **80,000 ops/s random reads**.
- Implemented a crash-safe write path (**WAL** → **fsync** → **VLog**), guaranteeing durability and recovery correctness across power failures.
- Designed multi-level compaction with newest-write-wins semantics and log-driven garbage collection.
- Designed APIs/interfaces for system interaction and integration with external components.

OrionScheduler: Concurrent DAG Task Scheduler | *Go, Concurrency, Fault Tolerance* **2025**

- Developed a task scheduling system for reliable execution of dependent workflows.
- Built a concurrent DAG-based task scheduler supporting dependency resolution, priority-based execution, and crash-safe recovery via write-ahead logging (WAL), sustaining ~**10,000 tasks/sec**.
- Designed a single-goroutine event loop architecture to eliminate mutex contention, utilizing atomic counters for blocking metrics calls to restore throughput under concurrent load.
- Resolved critical race conditions by introducing atomic cancellation flags with dual-guard checks, and implemented failure propagation using BFS cascades.
- Designed interfaces/APIs for system interaction and integration with external workflows.

Publications, Certifications & Leadership

Publication: "Edge-AI based Real-Time Infant Monitoring using Multimodal Sensing." *Springer LNCS*, VDAT 2025. (First Author)
Certifications: IBM Data Science Professional, IBM AI Engineering Professional, Meta Back-End Developer
Leadership: Core Team Lead for Infotsav (Techno-Managerial Fest) & Aurora (Cultural Fest), managing 50+ teams, securing 10+ sponsorships, and executing multi-day event operations.