

Atharv Rawal

+91 9611225911 • atharvrawal124@gmail.com • [in atharv-rawal](#) • [at atharvrawal](#)

Education

PES University

Bachelor's Degree in Computer Science (CGPA: 8.2/10)

Bangalore, India

2023 – 2027

Experience

Systems Developer (Summer Project)

Rust, Tokio, Python, Linux

PESU Innovation Lab [PortMux](#)

Summer 2025

- Engineered an asynchronous application-level port multiplexer in **Rust** (*Tokio*), routing multiple protocols via **non-blocking I/O** over a single exposed port.
- Developed a dynamic routing engine with **first-packet inspection** and a **YAML**-driven configuration supporting **Python**-based custom protocol handlers.
- Validated using Locust, sustaining **10,000 concurrent connections** with average latency overhead converging to **8%** above direct connection at peak load, demonstrating efficient scaling of the async architecture.
- Packaged as a **Systemd service** and published to the **Arch User Repository (AUR)**

Projects

P2P File Sharing System

Rust, AWS, TLS, TCP

- Designed and implemented a **dual-stack (IPv4/IPv6)** peer-to-peer file transfer system in **Rust** utilizing **raw TCP streams** and chunked transfer mechanisms for high-throughput streaming.
- Deployed an **AWS**-hosted signaling server to coordinate NAT traversal and peer discovery.
- Secured communication layers via **TLS encryption** and **SHA-256** checksum validation.

Automated CI/CD & Testing Pipeline

Python, GitHub Actions, Playwright

- Architected a strict CI/CD pipeline using **GitHub Actions** to enforce code quality, blocking PRs failing **>75%** test coverage, **Bandit** security audits, or **PyLint** thresholds.
- Developed a comprehensive automated testing suite encompassing unit, integration, and **Playwright**-driven E2E system tests.
- Engineered a custom fuzz testing suite to automatically validate API endpoints against Denial of Service (**DoS**) and Server-Side Request Forgery (**SSRF**) vulnerabilities.

Publication

A Survey of Hybrid Inference Systems for Large Language Models

TMLR, 2026

[OpenReview Paper](#)

- Co-authored a survey of **hybrid LLM inference**, covering speculative decoding, reasoning, model routing, multi-agent systems, and hardware-aware inference.
- Developed a **four-level taxonomy** spanning micro-, macro-, structural-, and physical-level orchestration.
- Identified the “**Orchestration Gap**”, highlighting the need for unified control to balance inference latency and reasoning quality.

Skills

Languages: Rust, C++, Python, C, JavaScript, Bash

Core CS: Data Structures & Algorithms, Operating Systems, Computer Networks, Object-Oriented Programming

Backend & Systems: Tokio, Async I/O, TCP/IP, TLS, CI/CD, Fuzz Testing, Linux Internals

AI Infrastructure: LLM Inference Optimization, Speculative Decoding, MoE Architectures

Tools: Arch Linux, Git, GitHub Actions, AWS, MongoDB, SQL