Harris Ahmad

Buffalo, NY | +1 (716) 279 5206 | harrisah@buffalo.edu | LinkedIn | GitHub

Education

University at Buffalo (UBuffalo) Jan 2025 - May 2028
PhD Computer Science and Engineering - PhD CSE Buffalo, NY
Teaching Assistant: Modern Networking Concepts (Spring'25)

Lahore University of Management Sciences (LUMS) Sep 2020 - May 2024
Bachelor's of Science in Computer Science - BS CS Lahore, PK
Teaching Assistant: Object Oriented Programming in C++ (Spring'23, Fall'23), Distributed Systems (Fall'23)

Experience

CPI-NS, UBuffalo Jan 2025 - Present
Research Software Engineer *Buffalo, NY*

- Implemented 3 concurrency and replication protocols such as **primary backup**, **chain replication** and **two phase locking** in Go to test their effectiveness in systems that are suited for collaborative AI. This work is a part of an ongoing (**NSF grant**).
- Testing the protocols using CloudLab's 4-node architecture showed these existing protocols are **unfit** for nodes that collaborate in real-time.
- Currently developing a distributed transaction **reordering** & **coordination** system for another **NSF funded project** by designing and implementing a **snapshot-based** architecture in C++ and Go that reduces **lock contention** and minimizes **transaction abort rate**, achieving a significant throughput improvement. We're submitting this to **OSDI'26**.

Linq, Asset Management (logicfarm) Jul 2024 - Dec 2024
Software Engineer *Dallas, TX (Remote)*

- Built **on-the-fly** media generation API using FastAPI and MongoDB that reduced page load time by **60%**, dynamically generating asset thumbnails as users scroll through 10K+ item inventories.
- Deployed successful SSO authentication using **OAuth 2.0** & **FusionAuth** enabling enterprise users to access systems with existing credentials.
- Created admin debugging and playground dashboards that exposed system metrics via direct OS calls, reducing troubleshooting time from hours to minutes for **50+ support tickets**.

Its IT Group May 2024 - Aug 2024
System Engineer *Lahore, PK*

- Redesigned the backend in **node.js** and **MongoDB** of an AR app for VTuber/karaoke 3D avatars leveraging WebSockets for **real-time** comments, notifications and likes across 5K+ daily active users.
- Implemented 3D model pipeline to process LiDAR scans, building USDZ to GLB converter enabling Unity integration, reducing manual measurement time by **80%**.

Interactive Media Lab, LUMS Oct 2023 - March 2024
Research Software Engineer *Lahore, PK*

- Built Awaaz-e-Sehat health platform using Python Flask framework on AWS Lambda, reducing API response time from 5s to 500ms handling 10K+ daily medical queries.
- Designed serverless data pipeline AWS (Lambda, DynamoDB, S3) to store and index **50K+** medical recordings, eliminating server costs while scaling automatically.
- Built medical transcription service using **Whisper API** and **GPT-4**, converting doctor consultations and medical history recordings to structured clinical notes with 95% accuracy, processing **1K+** recordings daily.

Networks and Systems Group, LUMS May 2022 - Dec 2023
Research Software Engineer *Lahore, PK*

- Discovered the impact of varying video bitrates across main-video resolutions: 720p causes ∼**3×** higher latent buffer loss than 360p (**10.1 MB vs 3.4 MB**), finding let to recommendations for adaptive ad-insertion/ABR changes to cut loss and cost.
- Collected data for **17600** YouTube videos and **46600** YouTube ads across 8 countries. This is the first large corpus of YouTube data, with **15+** features, that includes buffer data for each video. The dataset and codebase are publicly **available**.
- Published **findings** in Proceedings of the ACM Web Conference 2024 (**WWW'24**).
- Built distributed scraper using Selenium/PyTube to collect YouTube metrics, parallelizing across 20 machines gathering **2TB+** of performance data over a year.

Technical Skills

Programming languages: Python; C/C++; JavaScript; TypeScript; Haskell; Bash; PowerShell; Go; SQL; GraphQL.
Frameworks & libraries: Flask; Django; Express.js; React; Next.js; NumPy; SciPy; pandas; Selenium.
Developer tools: AWS; Azure; Git/GitHub; Linux; Postman; Android Studio.
Databases: MongoDB; CockroachDB; ScyllaDB; MySQL; Oracle.