

Harris Ahmad

+1 (716) 279-5206 | harrisah@buffalo.edu | Buffalo, NY, USA | linkedin.com/in/harris-ahmad1 | github.com/harris-ahmad | harris-ahmad.github.io/

EDUCATION

University at Buffalo - SUNY

PhD, Computer Science

- Teaching Assistant: Modern Networking Concepts (Spring'25)

January 2025 - May 2028

GPA: 3.7

Lahore University of Management Sciences

Bachelor's, Computer Science

- Teaching Assistant: Object Oriented Programming in C++ (Spring'23, Fall'23), Distributed Systems (Fall'23)

September 2020 - May 2024

PROFESSIONAL EXPERIENCE

University at Buffalo - SUNY

Research Software Engineer

- 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.

January 2025 - Present

Buffalo, NY

Linq Americas

Software Engineer

- 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.

July 2024 - December 2024

Dallas, TX (Remote)

Its IT Group

System Engineer

- 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%.

May 2024 - August 2024

Lahore, PK

Interactive Media Lab, LUMS

Research Software Engineer

- 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.

October 2023 - March 2024

Lahore, PK

Networks and Systems Group, LUMS

Research Software Engineer

- Discovered the impact of varying video bitrates across main-video resolutions: 720p causes ~3x 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 a distributed scraper using Selenium/PyTube to collect YouTube metrics, parallelizing across 20 machines gathering 2TB+ of performance data over a year.
- Paper link: <https://harris-ahmad.github.io/files/ytafford-www'24.pdf>

May 2022 - December 2023

Lahore, PK

SKILLS

Skills: C/C++, JavaScript, TypeScript, Go, Haskell, Bash, PowerShell, Python

Frontend: React.js, Vue.js, Quasar, AngularJS, Next.js, HTML/CSS, D3.js, jQuery, Redux.js, Electron.js

Backend: Node.js, Django, Flask, FastAPI, Express.js, Nest.js, GraphQL, gRPC, REST APIs, WebSockets

Databases: Postgres, MongoDB, Mongoose, MySQL, Redis, Elasticsearch, ScyllaDB, CockroachDB, Oracle, BigQuery, NoSQL, RDBMS

Cloud/DevOps: AWS, Google Cloud Platform, Microsoft Azure, Digital Ocean, Docker, Kubernetes, Jenkins, Heroku, Vercel, Firebase, RabbitMQ

Tools: Git, Postman, Selenium, Vim, Wireshark, Android Studio, Wordpress, Amazon Quicksight, Wireframe

ML/AI/Data: Pytorch, Tensorflow, Scikit-learn, Pandas, NumPy, SciPy, LangChain, LLM, GPT-4, Whisper API, Natural Language Processing (NLP), Data Science, Jupyter

Concepts/Other: Linux/Unix, SEO, OAuth 2.0, Solidity, Operating Systems, High Performance Computing (HPC), Microservices