

Farhan Sanukri

910-444-9421 | farsan@unc.edu | [linkedin.com/in/farhan-sanukri](https://www.linkedin.com/in/farhan-sanukri) | US Citizen

EDUCATION

University of North Carolina at Chapel Hill

Exp Grad: May 2027

Bachelor of Science in Computer Science, Mathematics | Dean's List

Chapel Hill, NC

- **Teaching Assistant:** COMP455 - Models of Computation (Spring 26, Fall 26)
- **Research:** Developed computer vision models for injury classification at Duke University under McNulty Lab
- **Relevant Coursework:** Data Structures, Algorithms and Analysis, Computer Systems, Machine Learning, Artificial Intelligence, Object-Oriented Programming, Databases, Software Engineering, Operating Systems

EXPERIENCE

Amazon Web Services

June 2026 – Present

Software Development Engineer Intern

New York, NY

- Incoming Summer 2026 - Team: Elastic Container Registry (ECR)

Carolina Experience Lab

January 2026 – June 2026

Software Engineer Intern

Chapel Hill, NC

- Shipped an interactive **sandboxed agent** embedded into course readings, serving **400+** weekly students with context-aware LLM hint generation and automated submission feedback
- Engineered a per-language warm pool of **pre-booted Docker sandboxes** for student code execution, eliminating per-request cold-boot latency and sustaining **200+** concurrent runs at **p99 ≤ 500ms**

Fidelity Investments

June 2025 – August 2025

Software Engineer Intern

Durham, NC

- Developed **Java Spring Boot** backend services via the **Money Movement API** for **Accounts and Trades**, powering ACH and cross-account transfers for **50+** million accounts at **\$500M+** in daily transfer volume
- Re-architected the **Transactions API** from synchronous webhook-driven **PostgreSQL** writes to asynchronous **Kafka** event-driven pipeline with a read-optimized view, lifting throughput **42%** and cutting primary DB read load **25%**

App Team Carolina

January 2024 – December 2025

Software Engineer

Chapel Hill, NC

- Built **TypeScript/Express** API coordinating donors, drivers, and pantries for PORCH, driving **\$440K** in donations
- Implemented webhook-driven SMS notifications via **Twilio**, paging nearest available drivers on donation scheduling and notifying receiving pantries on pickup, automating **2,000+** monthly alerts; deployed on **AWS ECS** with **RDS**

PROJECTS

Reef: Parallel Agent Orchestrator | *Rust, Tokio, CRIU, gRPC, OpenTelemetry*

- Built a parallel agent runtime for **RL post-training** and eval rollouts, scaling to **500+** sandboxed agents per host
- Built a snapshot-restore warm pool checkpointing agents post-bootstrap, collapsing startup $\sim 2s \rightarrow <100ms$ (**20x**) for **>10K rollouts/hour** sustained throughput
- Designed a capability-gated execution plane streaming tokens and tool results over a unified event bus with **per-agent backpressure**, absorbing provider rate-limit cascades while holding **p99 < 50ms** tool-dispatch latency at saturation

Low-latency Predictive Bikeshare in NYC | *Go, FastAPI, XGBoost, Apache Parquet*

- Developed a real-time service for predictive bike availability across NYC, handling **500k+** GBFS vehicle events per second and serving ONNX forecasts and live locations via FastAPI with ≤ 45 P95 latency
- Cut 15-minute prediction margin to **2.7%** by XGBoost training on **Apache Parquet** history, reloading weights nightly at 200ms latency

FlashFreeze: GPU Container Snapshot Runtime | *Rust, CUDA, Linux Kernel, OCI, Tokio*

- Built a **container runtime** in **Rust** that snapshots and restores **GPU inference** mid-execution via content-addressed lazy FUSE-served chunks, eliminating the **cold-start tax** of weight loading and context initialization
- Orchestrated **runc**, **CRIO**, and **CUDA** through typed async adapters, compressing **vLLM** (large language model) cold starts from $\sim 60s$ to **under 3s** via portable manifests with GPU compatibility metadata

TECHNICAL SKILLS

Languages: TypeScript, Python, JavaScript, Java, SQL, C++, Swift, Rust, HTML/CSS, C

Frameworks/Libraries: React, Next.js, Node.js, Express.js, FastAPI, Spring Boot, Pandas, NumPy, Angular

Developer Tools: AWS (ECS, EKS, S3, Lambda), Docker, PostgreSQL, Redis, Kafka, Kubernetes, Git, Linux