# SWETA VOODA

svooda3@gatech.edu | linkedin.com/in/sweta-vooda | swetavooda.github.io

## EDUCATION

**Georgia Institute of Technology - Atlanta, Georgia** **Aug 2023 – May 2025**

Master of Science in Computer Science, Specialization: Computing Systems – GPA: 3.8/4

- **Courses**: Advanced Operating Systems, Database Implementation, High Performance Computer Architecture, Computer Networks

**Keshav Memorial Institute of Technology** - Hyderabad, India **Aug 2018 – Aug 2022**

Bachelor of Technology in Information Technology - CGPA: 8.3/10

- **Awards**: Best Outgoing Student of Batch 2022

## WORK EXPERIENCE

**AWS, Control Tower | | Software Development Intern** **September 2024 – December 2024**

- Designed, Built, and deployed a solution to resolve a security issue in the drift detection of the service using AWS Lambda, SQS, and SNS, performing security reviews, ORRs, and collaborating with stakeholders for a seamless release.

**Sigma Computing | Software Development Intern** **June 2024 – August 2024**

- Designed and implemented Microsoft Teams as an export destination, enabling scheduled exports to SharePoint and a custom bot for automated messaging, projected to handle **35M exports annually** and contribute **$48M in ARR.**

**Amazon, Alexa Subscriptions onboarding | Software Development Engineer** **July 2022 – July 2023**

- Developed a Lambda-SQS service processing up to **460K messages/day** for **1.3M customers.**
- Innovated a metric factory to track previously undetected error metrics which led to a **72% reduction of system failures,** established automated monitoring, and enhanced code quality and operational efficiency.
- Single-handedly executed UI improvements using JSP, JSTL, Ajax, Amazon UI components as an away team project, that contributed to a **17.8% surge in Purchase Completion rates** optimizing user experience and compliance.
- Collaborated globally to **defend a DDoS attack** by identifying and blocking suspicious IPs, bolstering security, and user experience.
- Successfully **retrieved 67%** of **lost business dashboard data** utilizing strategic backfill solutions with **SQL and Redshift.**

**Adobe Systems | Product Intern** **May 2021 – August 2021**

- Designed and developed a high-throughput email campaign service by implementing Redis as a secondary cache, Addressed design aspects such as redundancy, concurrency, cache levels, eviction policies, and cache coherency.

## RESEARCH PROJECTS

**pgvector-remote a PostgreSQL extension to support multiple Vector Databases** **January 2024**

C | PostgreSQL | pgvector | VectorDB

- pgvector-remote is an extension designed to effortlessly incorporate remote vector stores, like Pinecone and Milvus, into pgvector. Implemented advanced batch insertions and query optimizations like metadata filtering and multicolumn indexing; VACUUM and LIMIT operations. Benchmarked on Big-ANN-Benchmarks and published our project to community – Article written.

**LangCache Semantic Caching Library for LLM Queries** **November 2023**

Python | Database - Query optimization | Cache | OpenAI

- Created LangCache, a semantic caching library for LLMs addressing drawbacks of **GPTCache** by implementing dynamic threshold adjustment. Improved workflow by **query optimization** through storage and indexing solutions **improving performance by 30%.**

## TECHNICAL SKILLS

- **Programming languages**: Java, C++, C, Python, Kotlin, Scala, Rust, Go; **HPC**: OpenMP, MPI, gRPC;
- **Web Development**: HTML/CSS, JavaScript, Angular, React, Spring MVC, GraphQL; **Compute Schedulers**: Slurm
- **Databases**: MySQL, PostgreSQL, SQLite, MongoDB, CockroachDB; Redis ;**VectorDBs**: pgvector, Pinecone, Milvus, ChromaDB
- **Networking**: OSPF and BGP configuration, P4; **AWS:** CDK, RedShift, Dynamodb, AWS Lambda, CloudWatch, S3; Docker

## SELECT PROJECTS

**Cache Architecture Simulator** **February 2024**

C++ | Cache | Computer Architecture

- Designed and developed a C++ cache simulator, with L1 and L2 caches, diverse cache sizes, associativity, advanced prefetching techniques and replacement algorithms. Rigorously evaluated across multiple workloads to find the best configuration.

**MapReduce Framework Development** **December 2023**

C++ | gRPC

- Implemented the MapReduce framework for processing large text files using gRPC in C++. Optimized performance through file sharding and asynchronous completion queues, ensuring dynamic workload distribution and resource efficiency.

**Implementation of Barrier Synchronization algorithms using OpenMP and MPI** **October 2023**

OpenMP | MPI | C | Slurm

- Implemented advanced barrier synchronization algorithms using OpenMP and MPI in C, drawing inspiration from Mellor-Crummey and Scott's research. Conducted comprehensive experiments analyzing algorithm behavior and performance in-depth.

**Dynamic Resource Management for Virtual Machines** **September 2023**

libvirt | C | CPU and Memory Virtualization

- Developed a vCPU scheduler and memory coordinator using C and `libvirt` toolkit to dynamically manage resources for VMs, for load balancing, memory ballooning, and efficient resource allocation, while ensuring optimal performance and stability.