Kavya Shekar

skavya.com | skavya@vt.edu | LinkedIn | GitHub | Google Scholar

Education

Virginia Tech — Aug. 2024 – Present
Master of Science - Computer Science (Thesis) | **GPA:** 4/4 — Blacksburg, VA

PES University — Aug. 2018 – Jun. 2022
Bachelor of Technology - Computer Science | **GPA:** 9.48/10 — Bangalore, India

Experience

Software Development Engineer II — Jul. 2023 - Jul. 2024
Groww | *Spark, Airflow, Trino, Iceberg, GCP Dataproc* — Bangalore, India

- Worked on building an **in-house Customer Engagement Platform** - involved in finalizing requirements, design, deliverables & timelines across various teams - Product, Engineering, Design, Marketing, etc.
- Personalized campaign engine - Implemented a scalable generic spark engine that runs config-based campaigns for **triggering push notifications to a 20M user base.**
- Actively worked to migrate campaigns from a third-party engagement tool to the personalized in-house platform which **improved user app engagement CTR by 26%** and **removed manual efforts by 50%˙**

Software Development Engineer — Jun. 2022 - Jul. 2023
Groww | *RocksDB, Spark, Aerospike, Springboot, Airflow, Kafka* — Bangalore, India

- **Fine-tuned performance of RocksDB** to serve data at 4k req/s per k8s pod with a response time of <1ms.
- Built an aerospike-spark connector and a streaming spark job for **updating streaming real-time data** to aerospike at a rate of **10k records/sec**.
- Developed Spring boot based **backend APIs for serving real-time feature data** from Aerospike and RocksDB - handling loads of 2k req/s per k8s pod with a **response time of 5ms**.

Software Development Intern — Jan. 2022 - Jun. 2022
Groww | *Prometheus, Loki, Grafana, k8s, GCP Cloud* — Bangalore, India

- Built and open-sourced a custom RBAC controller to automate the updation of user roles to Grafana organizations - **completely removed the manual efforts of on-calls.** — | Blog | GitHub
- Setup **Istio service mesh** on GKE clusters, testing and setting up canary deployments for production applications
- Developed a golang-based **Cloudflare Metric Exporter** to scrape metrics from cloudflare GraphQL server.

Projects

Optimizing Spatial Join Operation | *Hadoop, Spatial Operations* — Oct. 2021

- Optimized the I/O time during the shuffle phase of spatial join using SSD storage for shuffle storage.
- Achieved 16% speed-up with uniformly distributed data upto 25 GB in size.
- Realized up to 5% speed-up with non-uniform real-world land coverage data.

Technical Skills

Frameworks: Springboot, Flask
Technologies: Apache Spark, GCP Cloud, Iceberg, Airflow, MySQL Debezium, Kubernetes, Prometheus, Loki, Grafana
Databases: RocksDB, Aerospike, MySQL
Languages: Scala, Java, Python, C++, C, Golang, GraphQL

Publications

[1] Prafullata Auradkar, **Kavya Shekar**, et al. "Short Paper : Optimized Spatial Join with Grid Sub-Partitioning". In: *2021 IEEE CCEM*. 2021, pp. 41–45. DOI: 10.1109/CCEM53267.2021.00017.

Awards & Certificates

- Best Infra Award for GCS cost optimisation : Groww Hackraft 2023
- IEEE CCEM 2021 : Best Short Research Paper Award | Link
- CNR Rao and MRD award PES University : consistently ranked in top 20% within the CSE department | Link