Ganeshprasad Biradar

biradarganesh25@gmail.com | 9797212892 | linkedin.com/in/biradarganesh25 | biradarganesh25.github.io

EDUCATION

Texas A&M University — College Station, Texas, USA
Master of Science in Computer Science — GPA: 4.00/4 — August 2021 – May 2023

R.V. College of Engineering — Bangalore, Karnataka, India
Bachelor of Engineering in Computer Science — GPA: 9.09/10 — August 2015 – August 2019

TECHNICAL SKILLS

- Programming Languages: Go, Python, C++, C, SQL, R, Bash
- Databases: DynamoDB, Cassandra, CockroachDB, DocDB, Redis
- Cloud & Infrastructure: Kubernetes, AWS, GCP, Azure, Docker, Terraform, Ansible, Envoy
- Frameworks & Tools: Stripe API, gRPC, Flask, PyTorch, Scikit-Learn, go-kit, go-yacc, Spark, Git, CI/CD, Jujutsu
- Core Competencies: Distributed Systems, Microservices, Attestation, Concurrency, Systems Programming, Formal Verification (TLA+, PlusCal), AuthN/AuthZ

WORK EXPERIENCE

NVIDIA — Santa Clara, California, USA

Software Engineer 3 – Financial Systems Engineering — Jan 2026 – Current

- Core contributor and maintainer for NVIDIA's centralized payment services, supporting over 10 major platforms including GeForce NOW (150K+ customers), GDN, NVStore, and Brev.
- Designed and implemented a scalable email delivery service to automate subscription lifecycle notifications (creations, renewals, and cancellations) for approximately 150K GeForce NOW users.
- Engineered robust rollback mechanisms for complex subscription payment flows for GeForce NOW (creations, upgrades, and cancellations) to ensure data consistency and system stability during unrecoverable downstream failures.
- Skills: Go, Stripe, Payment Systems, DynamoDB, Distributed Systems, Kubernetes, AWS

Software Engineer 2 – NVIDIA Attestation — Feb 2025 – Jan 2026

- Led the development of attestation-sdk, a C++ library enabling Confidential Computing and hardware-rooted trust for NVIDIA GPUs and NVSwitches; supports Hopper (SPT/PPCIE) and Blackwell (SPT/MPT) architectures and is used by major cloud providers (Azure, Alibaba, Tencent).
- Skills: C++, Confidential Computing, Trusted Execution Environments (TEE), Attestation, Systems Programming

Software Engineer 2 – Unified Access Management (UAM) — Jun 2023 – Feb 2025

- Core contributor and maintainer of UAM microservices that provide an authorization platform for all of NVIDIA's cloud products (e.g., GeForce NOW, NVIDIA Inference Microservices).
- Led a zero-downtime database migration from DocDB to Cassandra. Learned PlusCal and wrote a PlusCal spec to verify a generic zero-downtime migration algorithm, which made it to the Hacker News front page: HN article.
- Proposed and led an initiative to standardize HTTP response codes across all UAM microservices, resulting in a 90% reduction in spurious PagerDuty alerts.
- Skills: Go, go-kit, testify, Envoy, AWS, Kubernetes, Redis, Terraform, Cassandra, Distributed Systems, TLA+, Formal Verification, Systems Programming

Cockroach Labs — New York City, New York, USA

Software Engineering Intern — Sept 2022 – Dec 2022

- Engineered Apache Parquet encoding for CockroachDB changefeeds, leveraging knowledge of key-value storage, distributed transactions, and SQL layers to extend the distributed CDC pipeline and reduce downstream storage by approximately 60%.
- Designed and implemented automatic scheduling for periodic changefeeds, providing fault-tolerant incremental exports without manual jobs. Authored the accompanying SQL grammar using go-yacc.
- Skills: Distributed databases and computing, Go, go-yacc, Automata theory.

NVIDIA — Santa Clara, California, USA

Software Engineering Intern — May 2022 – Aug 2022

- Researched and implemented async rate limiting for Unified Access Management (UAM), which added no latency to the critical path. Inspired by Cloudflare's scalable rate limiting architecture designed to mitigate large-scale L7 DoS attacks.
- Tuned connection parameters and conducted load analysis, validating 0ms added latency under sustained throughput of 5000 requests/sec.
- Designed and shipped zero-downtime APIs for dynamic rate-limit configuration; achieved 95% unit-test coverage.

Citrix R&D — Bangalore, Karnataka, India

Software Development Engineer 2 — Jul 2019 – Aug 2021

- Designed and implemented cloud modules for Azure and GCP to enable high availability and back-end autoscaling features in Citrix Application Delivery Controller (ADC), increasing product adoption by approximately 20%.
- Designed a centralized cloud interface abstracting siloed, provider-specific implementations, which standardized ADC cloud interactions, improved maintainability and testability, and cut feature delivery time by 3x.
- Skills: AWS, Azure, GCP, Python, C, Ansible.