# Anshu Gupta

ang037@ucsd.edu guptaanshu26 ang037.github.io

## Summary

**4th-year Ph.D. student in Computer Engineering** with expertise in **High-Performance Computing (HPC), and hardware-software co-design** for solving large-scale computational challenges in bioinformatics and AI. Highly proficient in Python, C/C++, and Verilog. Seeking research opportunities to leverage my expertise in similar fields.

## Education

**University of California San Diego** **September 2022 – Present**
Ph.D. student in Computer Science and Engineering, (M.S. awarded, GPA: 4.0) La Jolla, CA

**Indian Institute of Engineering Science and Technology, Shibpur** **July 2015 – May 2019**
B.Tech. in Electronics and Telecommunication Engineering; GPA - 9.55/10 (Rank: 4) West Bengal, India

## Relevant Publications [Google Scholar]

- **A. Gupta**, Y. Cao, J. Liang, Y. Turakhia, "DP-HLS: A High-Level Synthesis Framework for Accelerating Dynamic Programming Algorithms in Bioinformatics" - accepted at HPCA 2026. [arXiv]
- **A. Gupta**, S. Mirarab, Y. Turakhia, "Accurate, scalable, and fully automated inference of species trees from raw genome assemblies using ROADIES", Proc. Natl. Acad. Sci. U.S.A. 122 (19), e2500553122 (2025). [paper]

## Work Experience

**AMD Research and Advanced Development (RAD)** **June 2025 – Sep 2025**
Research Associate - PhD (Mentor: Stephen Neuendorffer) San Jose, California

- Emulated sub-modules of an advanced AI test chip on cutting-edge FPGA platform, contributing to the hardware and software bring-up of custom silicon.
- Conducted runtime performance characterization of AI models on the target test-chip architecture.

**Analog Devices, Inc.** **June 2023 – Sep 2023**
ASIC BU Digital Design Intern Wilmington, Massachusetts

- Conducted hardware architectural benchmarking for Cadence Tensilica Xtensa configurable processors.
- Characterized performance, power, and area metrics for low-power embedded processor configurations.

**Analog Devices India Pvt. Ltd.** **July 2019 – July 2022**
Digital Design Engineer Bengaluru, India

- Optimized embedded signal processing kernels using SIMD operations and architecture-specific vectorization, reducing DSP cycle count by 33%.
- Contributed to the digital front-end design and tape-out of a low-power Battery Management SoC.
- Evaluated ARM Cortex-M processor architecture and SRAM/ROM architecture's power, performance, and area trade-offs for low-power embedded SoCs.

## Research Experience (Mentor: Yatish Turakhia)

**ML Hardware Acceleration** **May 2025 – Ongoing**

- Designing **hardware-aware weight compression techniques** to significantly optimize and accelerate Deep Learning inference on specialized hardware.
- Research focused on maximizing computational efficiency and reducing memory footprint for deploying large-scale neural networks at the edge.

**ROADIES (Automated Species Tree Inference Tool for Genomic Datasets)** **Jan 2023 – Ongoing**

- **Led the development of** ROADIES, a reference-free, highly parallelized HPC tool for evolutionary genomics using **GPU/CPU co-processing and parallelization primitives**.
- Achieved **176x speedup** in tree generation for large-scale genomics dataset, reducing computational cost.

- Published in PNAS, featured on the Journal Cover.
- Presented at **Intelligent Systems for Molecular Biology (ISMB) 2024** conference.

**DP-HLS (FPGA Acceleration Framework)** **Nov 2022 – June 2024**

- **Developed DP-HLS**, an HLS-based framework, that achieved a **32x speedup** and **20x faster implementation time** for accelerating dynamic programming algorithms for bioinformatics workloads.
- Designed the framework for easy customization of algorithms without hardware redesign, using highly parameterizable systolic array architectures.

## Technical Skills

**GPU & HPC Acceleration**: CUDA, MPI, OpenMP, Intel TBB, Slurm, Amazon AWS, Snakemake, Conda, Bash
**AI/ML & Software**: Python, C/C++, HLS, R
**Hardware & Co-Design**: ASIC/Custom Silicon Design Flow, FPGA (AMD/Xilinx Vitis HLS/Vivado), Verilog, SystemVerilog

## Other Projects

**Parallelized Genomic and HPC Algorithms** | *C++, CUDA, Intel TBB, Git* **Jan 2023 – March 2024**

- Executed high-performance parallelization of Suffix Array construction on GPUs, achieving $86 - 571x$ speedup over optimized CPU baselines for genomic read mapping.
- Accelerated dense matrix multiplication using CUDA on NVIDIA K80/T4 GPUs and implemented Intel AVX2 vectorization for high-throughput computing
- Optimized computation on the Expanse Supercomputer by applying MPI and C++ parallelization to challenging solvers (e.g., Aliev-Panfilov)

**Optimization of HLS4ML Library** | *Vitis HLS, Vivado, Git, Python* **April 2023 – June 2023**

- Improved the HLS4ML library to achieve efficient ML hardware inference via High-Level Synthesis, focusing directly on the DL-to-FPGA flow.
- Partnered with CERN's HLS4ML team to implement feature enhancements and optimizations.

## Relevant Coursework

- Principles of Computer Architecture
- Parallel Computer Architecture
- Parallel Computation
- Parallel Computation in Bioinformatics
- Design Automation and Prototyping
- Validation and Testing in Embedded Systems
- Algorithms in Computational Biology
- Computational Evolutionary Biology