Anshu Gupta

619-953-7358 ang037@ucsd.edu guptaanshu26 github.com/ang037

Summary

Ph.D. student in Computer Engineering with a focus on accelerating bioinformatics workloads using CPU and GPU architectures. Extensive experience in optimizing algorithms for both software (CUDA, Intel TBB) and hardware (FPGA, ASIC). 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; GPA - 4.00/4.00 *La Jolla, CA*

Indian Institute of Engineering Science and Technology, Shibpur **July 2015 – May 2019**
B.Tech. in Electronics and Telecommunication Engineering; GPA - 9.55/10.00 (Rank: 4) *West Bengal, India*

Research Experience

University of California San Diego **November 2022 – Present**
Graduate Student Researcher (Advisor: Prof. Yatish Turakhia) *La Jolla, CA*

- **Developed** ROADIES, an automated software pipeline to overcome the scalability challenges in inferring phylogenetic trees from large genomic datasets.
- Optimized algorithms and data handling, resulting in a **176x speedup** in phylogenetic analysis of genomes for projects like Bird 10K and VGP.
- Reduced computational costs and time for evolutionary research by making large-scale phylogenetic analysis more efficient and accessible.
- **Designed DP-HLS**, a generalized High-Level Synthesis framework, achieving **32x speedup** and **20x faster implementation time** for bioinformatics DP algorithms on FPGAs.
- Simplified the customization of dynamic programming algorithms, making FPGA acceleration more accessible and efficient for a wide range of bioinformatics applications.
- Extended the framework's applicability in multiple sequence alignment and protein structure prediction.

University of Bremen **May 2018 – July 2018**
DAAD Summer Research Intern *Bremen, Germany*

- Optimized fault-tolerant quantum circuits by reducing their synthesis cost using Clifford+T gate library.
- Achieved 10-50% reduction of T-gate count with same qubits, significantly optimizing synthesis costs.

Work Experience

Analog Devices, Inc. **June 2023 – September 2023**
ASIC BU Digital Design Intern *Wilmington, Massachusetta*

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

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

- Optimized radar signal processing algorithms for automotive radar-based SoC using SIMD operations, reducing DSP cycle count by 33%.
- Contributed to the digital front-end design and tape-out of low-power Battery Management Systems SoC.
- Conducted quality assurance tasks (Lint, Clock/Reset Domain Crossing (CDC/RDC), Logic Equivalence Checks (LEC)) for BMS SoC RTL design, ensuring design integrity.
- Evaluated ARM Cortex-M processor architecture and SRAM/ROM architecture's power, performance, and area trade-offs for low-power embedded SoCs.

Selected Publications [Google Scholar]

- **A. Gupta**, S. Mirarab, Y. Turakhia, "Accurate, scalable, and fully automated inference of species trees from raw genome assemblies using ROADIES", bioRxiv 2024. [paper]
- P. Niemann, **A. Gupta**, R. Drechsler, "T-depth Optimization for Fault-Tolerant Quantum Circuits", IEEE ISMVL 2019. [paper]

Technical Skills

Languages: Verilog, SystemVerilog, Python, C/C++, R, MATLAB, CUDA, Intel TBB, OpenCL
HPC & GPU Tools: CUDA, MPI, Intel TBB
FPGA Tools: AMD Xilinx Vitis HLS, Vivado Design Suite
Cloud & Workflow Management: Amazon AWS, Snakemake, Slurm, Git, Bash
Design & Verification Tools: Synopsys SpyGlass, Cadence SimVision

Projects

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

- Improved HLS4ML library for efficient Machine Learning hardware inference via High-Level Synthesis.
- Partnered with CERN's HLS4ML development team, contributing to several key feature enhancements and optimizations.

Parallelized Suffix Array Construction | *C++, CUDA, Git* **January 2023 – March 2023**

- Implemented CUDA-based Suffix Array construction, enhancing speed by 86-571x compared to CPU baseline.
- Applied various parallelization and SIMD vectorization techniques with CUDA, optimizing Suffix Array efficiency.

Parallelized Read Mapping Algorithms | *C++, CUDA, Intel TBB, Git* **January 2023 – March 2023**

- Utilized Intel Thread Building Blocks for bioinformatics read mapping, achieving a significant speed boost.
- Integrated CUDA in read mapping algorithms, streamlining genomic data processing.

Matrix Multiplication Acceleration | *CUDA, MPI, C, Intel AVX2* **October 2022 – December 2022**

- Enhanced C-based matrix multiplication speed by implementing blocking and Intel AVX2 vectorization.
- Improved CUDA matrix multiplication on K80 and T4 GPUs using blocking and shared memory techniques.
- Optimized Aliev-Panfilov solver performance with C++ and MPI on the Expanse supercomputer, achieving significant computational efficiency gains.

Branch Predictor | *C, Git* **October 2022 – December 2022**

- Developed models for Gshare and Tournament branch predictors in C, improving CPU prediction accuracy.
- Implemented Perceptron and TAGE branch predictors for enhanced prediction efficiency.
- Designed L1 Cache with FIFO replacement, optimizing memory access and CPU performance.