Karthikeyan Renga Rajan

Bryan, TX 77840 (Willing to Relocate) | krengara@gmail.com | +1 (979) 575-1749
Linkedin: krengara | GitHub: Karthikeyan564 | Website: karthikeyan564.github.io

Education

Texas A&M University — **College Station, TX**
Master of Science in Computer Engineering (CEEN) — August 2024 - May 2026 (Expected)

- Recipient of **departmental scholarship worth $10,747**
- **Coursework:** Computer Architecture | Deep learning | Microprocessor System Design
- **GPA:** 4.0 out of 4.0

College of Engineering, Guindy, Anna University — **Chennai, India**
Bachelor of Engineering in Electronics and Communication Engineering — August 2019 - May 2023

- CGPA: 9.41 out of 10 - 3.99 on a 4.0 equivalent

Skills

Languages: SystemVerilog | Verilog | C/C++ | SystemC | Python | Tcl | Perl | LaTeX
Frameworks: UVM | GEM5 | Scalesim | cocotb | PyTorch | CUDA | OpenGL | OpenCV | TensorFlow
Tools: Vivado | Quartus | Verilator | Zsim | Git | GitHub | MATLAB | Vim | QEMU | Linux/Unix
Protocols: PCIe | TCP/IP | DMA | USB | Ethernet | JTAG | UART

Experience

HelixLogic — **Chennai, India**
Co-founder and Digital Design Engineer — December 2022 - March 2024

- Co-founded an **ASIC design startup** focused on Ultrafast, Nanopore DNA sequencing, raising **$10,000** in startup capital in the first year.
- **Designed and Implemented an RISC-V V-extension based ASIC** in SystemVerilog to decode DNA sequences based on signals derived from a Nanopore Membrane.
- Achieved a cycle-to-cycle acceleration factor of **107 at 400 MHz** on benchmarks as compared to a Tesla V100.
- **Managed a 2-member** undergraduate student team.

ISM Laboratory, University of Toronto (In-person) — **Toronto, Canada**
Summer Research Intern — May 2022 - August 2022

- **Developed a RISC-V core in SystemVerilog** to create custom mask sequences on the fly for coded-exposure image sensors.
- Performed verification of the RISC-V core through **OpenOCD on a Nexys Video Artix-7 FPGA**.
- Improved the **Mask generation and Mask decompression modules** for generating hard-coded masks and receiving compressed masks **leading to a 24.47% gain in performance**.

Integrated Systems Lab, Anna University — **Chennai, India**
RTL Design Intern — May 2021 - January 2022

- **Designed a GPS Baseband Engine** as a team of 4 students using OpenLane and sky130 PDK.
- **Selected by SSCS PICO Contest for fabrication.**
- Implemented RTL changes to accommodate modified architecture of the ASIC.

Indian Academy of Sciences(IASc) — **Bengaluru, India**
Summer Research Fellow — June 2021 - July 2021

- Developed an **Image-Based-Rendering (IBR) based Reinforcement learning environment** for end-to-end training to avoid sim2real, domain adaptation or domain randomization etc using CUDA and OpenGL.
- **Achieved a 67% better model transfer rate than Sim2real** while achieving a **2.57x FPS speedup.**

Projects

Implementation of a Machine Learning Inference Accelerator with a Hierarchical Mesh-based NoC
https://github.com/Karthikeyan564/uEyeriss — January 2023 - April 2023

- **A SystemVerilog Implementation of the Eyeriss CNN architecture** with a Hierarchical Mesh-based Network-on-Chip(NoC) and approximate multipliers to increase speed and energy efficiency.
- Showed an **overall improvement of 77.23% over the base architecture** on golden tests in Integration Testing.
- Received the **Anna University Guindy Engineers 65 Project Award Endowment** for best bachelor's thesis among 50+ Projects.
- Additionally implemented and tested on **Xilinx Kintex-7 KC705 FPGA**.

Architectural Implementation of MERCURY: Accelerating DNN Training by Exploiting Input Similarity
https://github.com/CSCE-614-Dr-Kim-Fall-2024/MERCURY_GRP19 August 2024 - Present

- Modified PyTorch and Scalesim to implement generation of signature of input vectors using existing Processing Elements (PEs) and then perform caching using the signatures.
- Achieved a **2.14x speedup** with only 0.23% loss in accuracy on the ResNet101, VGG-19 and GoogleNet models.

H.264 Video Codec Accelerator over PCIe with DMA device drivers
https://github.com/Karthikeyan564/h264_bachelor_thesis August 2022 - December 2022

- Worked on Intra Prediction, Fractional Motion Estimation and Transform Coding modules.
- Simulated using **Verilator, SystemC-TLM, and QEMU.**
- Achieved a **VMAF of 85.46 over the H.264 White Paper's 82.19**.
- Implemented and tested on **Altera Cyclone IV FPGA**.

Metis: Programmable 256-neuron, 2048-synapse neuromorphic chip in 130nm CMOS to accelerate RSNNs
https://github.com/Karthikeyan564/Metis August 2021 - January 2022

- The chip features an analog circuit for the **leaky integrate-and-fire neurons and on-chip e-prop learning.**
- The on-chip e-prop was able to train a spiking neural network to achieve an **accuracy of 98.96% on the MNIST dataset** with a **power efficiency of 4.78pJ/SOP at 1.8V**.
- **Fabricated using the efabless Open MPW-3 shuttle**.

Hardware-Software Codesign for Verilog development over PCIe
https://github.com/Karthikeyan564/Idaten August 2022 - December 2022

- Developed a PCIe-Verilog verification setup for codesign using QEMU and SystemC-TLM enabling **development of Firmware and RTL simultaneously**.
- Uses Remote Port IPC to connect the QEMU Virtual machine and the Verilator simulation.

InfoSogCLR - A Novel loss function to improve CLIP model training on resource-constrained machines
https://github.com/Karthikeyan564/InfoSogCLR August 2024 - Present

- Implemented a novel loss function to improve the performance of CLIP model trained on a 100k subset of the Conceptual Captions 3M (CC3M) Dataset with a batch size of 128 for 30 epochs.
- Achieved a **15.09% on Image-to-Text Recall at position 1**, **17.92% on Text-to-Image Recall at position 1** on the MSCOCO dataset and **20.90% on Top 1 Accuracy** on the ImageNet dataset outperforming the SogCLR and CLIP loss functions.