

# AKSHATH RAGHAV RAVIKIRAN

765-404-8121 [araviki@purdue.edu](mailto:araviki@purdue.edu) [linkedin.com/in/akshathrr](https://linkedin.com/in/akshathrr) [github.io](https://github.com/akshathraghav.github.io)

## SUMMARY

MS ECE Candidate seeking full-time (August 2026) Digital Design roles focused on CPU microarchitecture, simulation and performance analysis.

## EDUCATION

### Purdue University, West Lafayette

Master of Science in Electrical and Computer Engineering

Bachelor of Science in Computer Engineering

**Cumulative GPA: 3.64**

August 2025 – August 2026

August 2022 – December 2025

Coursework: Computer Architecture, MOS VLSI Design, AI Hardware Design, ASIC Design, Operating Systems, Microprocessors & MCUs

Honors: ECE Senior Design Award, Eli Shay Scholarship, UIUC HDR Fellowship, Outstanding Sophomore (VIP), Dean's List & Semester Honors (7x)

Organizations: IEEE Eta Kappa Nu, Purdue SoCET, ECE Graduate Student Association, Purdue Chess Club

## TECHNICAL SKILLS

AI/ML (PyTorch/TensorFlow) | GPU Programming (CUDA/Triton) | Parallel Programming (OpenMP) | Bash, TCL Scripting. C/C++/Assembly | Modelling w/ Gem5 & GPGPU-Sim | Protocols – Serial (SPI/I<sup>2</sup>C/UART) & AMBA (AXI, AHB) | Embedded C w/ ESP-IDF. SystemVerilog, UVM | RTL simulation (QuestaSim, Verilator) | FPGA synthesis (Xilinx Vivado/Vitis) | ASIC STA & synthesis (Cadence Genus).

## EXPERIENCE

### Research Assistant

Purdue SoCET (PI: Prof. Mark Johnson, Purdue-ECE)

July 2024 – Present

West Lafayette, IN, US

- AI Hardware: Leading on-chip [Memory Subsystem](#) for Atalla Tensor Core; focusing on architecture diagramming & ISA design.
  - Built a cycle-accurate simulator of the datapath for **performance modelling** using implicit-convolution and GEMM kernels.
  - Architected a parameterizable 2MB Scratchpad with on-the-fly swizzling and a pipelined  $N \times N$  crossbar — optimized for PPA.
  - Designed FP16 datapaths between Systolic Array & Vector Core; integrating DDR4 controller for asynchronous DRAM transfers.
- GPU: Advising Hardware team in RTL Design & Python Modeling; designing custom Cardinal GPU Core for graphics workloads.
  - Implemented a [lockup-free-cache cache](#) – Achieved 100% coverage (ModelSim); Optimized to synthesize (Genus) at 700MHz.
  - Modelling per-warp divergence-mitigation heuristics in GPGPU-Sim; simulated to improve IPC [up to 13%](#) on Rodinia benchmarks.
- Enhanced the [AFTx07 RISC-V core](#) with Zicond extension for macro-fusion of conditional arithmetic/logic sequences.

### ML Engineering Researcher

Duality Lab (Contract w/ Google LLC – PI: Prof. James Davis, Purdue)

August 2023 – April 2024

West Lafayette, IN, US

- Helped re-engineer the [MaskFormer model](#) from PyTorch to TensorFlow to run on GCP TPUs & integrate into the TF Model Garden.
- Contributed to a technical [white paper](#), providing implementation guidance for TPU-focused **HW/SW co-design**.
- Integrated auxiliary losses & conducted hyperparameter tuning to increase Panoptic Quality scores by 25% on the COCO Dataset.
- Performed distributed training on GCP & debugged on RCAC Gilbreth w/ a MLOps workflow to track model improvements.

### AI Engineering Intern

BMW - Group IT

May 2025 – August 2025

Munich, Bavaria, Germany

- Built remote [AI Agents](#) w/ LangGraph to crawl enterprise apps & automatically run QA tests – eliminating Playwright automation KPIs.
- Hosted a Github Copilot Assistant for human-in-the-loop provisioning of Azure infra, helping full-stack devs deploy apps internally.
- Deployed FastMCP connectors for core DevOps platforms into GAIA (internal AI platform) w/ AWS Lambda & EventBridge.
- Engineered RAG connectors exposing searchable knowledge graphs in AI IDEs – validated to **outperform GAIA** on 10K+ LOC docs.

## SELECTED PROJECTS

### BoilerNet – Compute-Enabled Mini-NAS

KiCad 9.0, PCB Bring-Up, Edge ML, Networking Protocols

- Designed PCBs, assembled into 3D-printed enclosures, to form a Network Attached Storage w/ swappable compute blades & memory slots.
- Received Purdue ECE's [Senior Design Award](#) in Spring '25 for our **decentralized microcontroller** stack & master/slave SPI-based drivers.
- Supports INT16/FP16 quantized MobileNetV2 models with DMA-friendly data-parallel pipelines through TFLite Micro and ESP-IDF.

### RISCV Five-Stage Multicore Processor

QuestaSim, Xilinx Vivado, RTL Design, FPGA Prototyping

- [Pipelined processor](#) implementing branch prediction, forwarding/hazard detection logic and dual-core MSI snoopy cache coherency.
- Synthesized to DE2-115 FPGA at 60MHz — Performed **static timing analysis** showing  $2.77 \times$  speedup over single-cycle design.
- Memory controller arbitrates read/write memory accesses to external memory and supports variable-latency access.

### gem5 Microarchitecture Studies

gem5, Architectural Simulation, CPU Performance Analysis

- Added a Waiting Instruction Buffer into the **O3 CPU** to offload load-dependencies, [improving IPC](#) by up to 13% on SPEC benchmarks.
- Implemented a victim cache with no-allocate and mostly-exclusive policies, reducing L1 miss rates by upto 5% on high-locality workloads.

### 8-bit Wallace Tree Multiplier – Physical Design

Cadence Virtuoso ADE, GPDK45nm

- Designed a Wallace Tree Multiplier using the full-adder inversion property, saving 298 transistors and [achieving a  \$3305 \mu\text{m}^2\$  footprint](#).
- Completed **schematic-to-layout flow** with DRC/LVS and post-layout parasitic extraction, with 1.74 ns delay and 320 fJ energy (1.0 V).

### tinySpeech – Speech Recognition on Edge Devices

PyTorch, ML Quantization, MCU Programming

- Reproduced TinySpeech word-recognition models; Achieved 91% precision benchmarks w/o access to DarwinAI's proprietary code.
- Custom **quantization-aware training** pipelines allow for INT4/8 training and PerTensor/PerWeight scales, with 5% accuracy drop.
- Developed an embedded C-based [inference engine](#), optimized for INT8 precision targeting RISCV-EC architecture.