

Ammar Ratnani

(281) 223 2900 | aratnani@stanford.edu
[linkedin.com/in/ammarratnani](https://www.linkedin.com/in/ammarratnani) | github.com/ammrat13

Education

- Stanford University** | M.S. in Computer Science June 2025
- *Coursework:* Embedded Operating Systems, Networking, Computer Architecture, VLSI Design
 - *GPA:* 4.27
- Georgia Institute of Technology** | B.S. in Computer Science May 2023
- *Coursework:* Processor Design, Secure Computer Architecture, High-Performance Computing
 - *GPA:* 4.00

Work Experience

- MINOTAUR / EE 372** | *Student Researcher* Apr. 2025 - Jun. 2025
- Optimized the performance and area of this neural network accelerator
 - Analyzed Catapult HLS schedules, Synopsys Design Compiler reports, as well as C++ and Python compiler code, to find performance bottlenecks
 - Implemented changes that improved the runtimes of MobileBERT and ResNet-18 on the accelerator by 23% and 10% respectively, while reducing the area of the design by 3%
 - Devised targeted fixes for softmax and max-pooling, reducing those particular layers' runtimes by 50% and 60% respectively
- NVIDIA** | *Software Engineering Intern* Jun. 2024 - Sept. 2024
- Worked on increasing game performance on the GeForce Now cloud gaming platform, using Windows Performance Toolkit for collection and Python for analysis
 - Created a dashboard to display CPU-side bottlenecks, including: parallelism, scheduling, inter-processor communication, hypervisor steal time, and interrupts
 - Deep-dived the causes of a performance regression on Rainbow Six: Siege when running on AMD CPUs instead of Intel
 - Improved the reliability of automated benchmarks, and even extended them with three new stress-tests and one new game
- The Aerospace Corporation** | *Software Engineering Intern* May 2023 - Aug. 2023
- Developed a fully autonomous ground station to receive images via radio transmission from the NOAA 15, 18, and 19 weather satellites
 - Integrated a GNU Radio flowgraph with Python code to demodulate and synchronize Automatic Picture Transmissions in real time on embedded hardware
 - Investigated decoding Differential Binary Phase-Shift Keyed transmissions from the NOAA GOES 16 weather satellite
 - Constructed a prototype transpiler that ingests SysMLv2 and produces HSFL COSMOS configuration files, looking to use it in a CI/CD pipeline
- Green Hills Software** | *Software Engineering Intern* Jun. 2022 - Aug. 2022
- Diagnosed performance bottlenecks on Windows, obtaining a 25% speedup by eliminating unneeded memory allocations and synchronizations
 - Deployed Windows as a guest into Green Hills's hypervisor via PXE using both wimboot and Windows Deployment Services
 - Patched OVMF to make it compatible with Windows under the hypervisor
 - Debugged xHCI USB pass-through in the hypervisor
 - Implemented VPN scalability tests, ensuring it can support up to 5000 connections

Projects and Contributions

- Zynq 7000 HDMI Peripheral:** <https://github.com/ammrat13/meta-hdmi-dev> Mar. 2024 - Jun. 2024
- Created an FPGA-based HDMI device, integrated it with the processor on this SoC, and exposed it as a framebuffer via a custom Linux kernel driver
- LLVM Cross-Compiler for the LC-3.2:** <https://github.com/lc-3-2> Feb. 2023 - Sept. 2023
- Constructed a backend to generate assembly for a variant of the LC-3 instruction set. Ported both `newlib` and `coremark` to the new architecture to verify the compiler's correctness