

Anubhav Kumar Das

+91 88672 95324 | anubhavkumarr02@gmail.com | [linkedin](#) | [github](#)

EDUCATION

Indian Institute of Technology

B.Tech in Microelectronics and VLSI | CGPA: 8.88/10

Himachal Pradesh, India

Aug. 2024 - Present

- Relevant Coursework: Computer Architecture and Organization, Programming and Data Structures, Digital System Design, Machine Learning, Deep Learning and Applications

TECHNICAL SKILLS

Languages: SystemVerilog, Verilog, C/C++, Python

Hardware Design & Architecture: RTL Design, Microarchitecture, Computer Architecture, RISC-V (RV32I), SoC Design, AXI4 Interconnect, Systolic Arrays, Pipeline Hazard Handling, Logic Synthesis

Machine Learning & Frameworks: Deep Learning, PyTorch, Triton, Custom GPU Kernels, HW/SW Co-Design, 3D Image Registration, 3D U-Net, Parallel Computing

Tools & Platforms: Xilinx Vivado, Vivado ILA, Yosys/ABC, GTKWave, Xilinx Zynq-7000 FPGA, NVIDIA Ampere GPUs, Git, Linux

PROJECTS

Dual-Core RISC-V SoC with Integrated Output-Stationary Systolic Array AI Accelerator

[GitHub](#)

SystemVerilog, Vivado, AXI4, FPGA Implementation, RISC-V, Vitis, Vivado ILA

- Architected a dual-core RV32I cluster tightly coupled with an 8×8 Output-Stationary Systolic Array (64 INT8/INT32 MACs), utilizing private 16KB ITCMs and a shared 64KB DTCM.
- Designed a 5x2 AXI4 crossbar and 3-channel decentralized DMA with two-tier arbitration (Strict CPU Priority + Round-Robin DMA) to streamline matrix GEMM workloads.
- Developed a self-checking SystemVerilog SoC testbench featuring constrained-random stimulus, behavioral AXI4 memory models, automated golden-model verification, and hardware watchdog probes to debug and eliminate AXI transaction deadlocks.

GPU-Native 3D Medical Image Registration Framework (HW/SW Co-Design)

[GitHub](#)

Triton, C++, Python, CUDA, PyTorch, SimpleITK

- Engineered a C++/Triton B-spline kernel mapped to Ampere SMs with constexpr tile sizing, driving loop unrolling, zero-branch execution, and optimal register file utilization to deliver a 100–150x speedup over CPU solvers.
- Replaced 3D unfolding convolutions with zero-allocation tiled execution (cutting VRAM from 64x to 1x DVF); retained $4 \times 4 \times 4$ stencil data in L1 cache (128–192 KB) and eliminated backward pass write races via hardware L2 atomics (ATOM.E.ADD.F32).
- Wrapped custom kernels into torch.autograd.Function with analytical gradient verification; achieved 183.9 ms full-volume ($160 \times 224 \times 192$) warping latency on RTX A6000 (2x faster than SOTA vision transformers) with a near-zero 0.012% fold rate.

5-Stage Pipelined RV32I Processor with Dynamic Branch Prediction

[GitHub](#)

Verilog, Yosys/ABC, GTKWave, Vivado, Computer Architecture, TCL Scripting

- Architected a 5-stage in-order RV32I core featuring a dedicated Hazard Detection Unit (HDU) for 1-cycle load-use stalls, full MEM/WB forwarding, and a 1024-entry 2-bit bimodal BHT with an IF-stage target pre-decoder—delivering 0-cycle taken branch penalties and cutting effective CPI from 1.35 to 1.08.
- Executed a 4-phase verification suite covering BHT aliasing and priority inversion; synthesized RTL via Yosys/ABC (~67k NAND2 gates, 5,659 FFs) to analyze gate-level cell and wire density.

ACHIEVEMENTS

1st Place in Motion Mimicking Track and **4th overall** at the **HCLTech** DL Hackathon (approx. 600 participants).

POSITIONS OF RESPONSIBILITY

• Robotronics Club, IIT Mandi – Project X team member

Aug 2024 – Dec 2024

Contributing to the development of autonomous robotic systems for national competitions.

• Career & Placement Cell, IIT Mandi – Student Representative

May 2025 – Present

Assisting in placement activities by coordinating with recruiters, student volunteers, and the placement team.