

Prayag Mohanty

Mumbai, India

[linkedin.com/in/prayag-mohanty](https://www.linkedin.com/in/prayag-mohanty) · github.com/Prayag-Mohanty · prayag-mohanty.github.io

EDUCATION

Indian Institute of Technology Bombay

M.Tech, Electrical Engineering — Integrated Circuits & Systems

Mumbai, India

2024 – 2027 (expected)

BITS Pilani, K. K. Birla Goa Campus

B.E. (Hons.), Electrical & Electronics Engineering

Goa, India

2020 – 2024

Coursework: Processor Design · Advanced Topics in Computer Architecture · VLSI Design · VLSI Testing & Verification · Embedded Systems Design · Neuromorphic Engineering · Machine Learning

RESEARCH EXPERIENCE

Graduate Researcher — Computer Architecture & Dependable Systems Lab (CADSL), IIT Bombay

Jul 2025 – Present

Advisor: Prof. Virendra Singh

- **Address Translation for Multi-Chip-Module GPUs** (M.Tech thesis, ongoing) — investigating virtual-memory translation as a first-order bottleneck when the TLB and page-walk hierarchy is partitioned across dies. Proposed coupling TLB and cache replacement by using **dead-page prediction at the L1 TLB to predict deadness of that page's blocks in the L1 cache**; prototyping a bypass-candidate configuration on the **MGvm** MCM-GPU simulator. Targeting a conference submission by mid-2027.
- **Register-File-Resident Victim Caching with Locality-Aware Warp Scheduling** — co-designed a victim cache held in underutilised register-file capacity with a warp scheduler that biases selection toward warps whose working set is still resident, so retained L1D lines are re-referenced before eviction. Implemented end to end in **GPGPU-Sim** on an NVIDIA Volta GV100 configuration; **1.042× IPC** over baseline. Presented as a poster at the Indo-Japanese Workshop on Trustable Cyber-Physical Cognitive Systems, IIT Bombay (2026).
- **Out-of-Order Scheduling for GPUs** (M.Tech seminar) — surveyed **7 OoO GPU proposals** (ISCA/HPCA/TPDS/TACO, 2016–2024) and tabulated their **performance, power and area** trade-offs: 6.9–38% speedup at 0.007–13% power and 0.05–2.4% area. Showed PTX-level simulation overstates available ILP — **LOOG's +16% on PTX becomes –16.5% on SASS** — establishing SASS-level evaluation as the representative baseline.

HARDWARE DESIGN & ARCHITECTURE PROJECTS

- **SS16 — 16-bit Superscalar Processor** (SystemVerilog, 2026): designed and verified an in-order superscalar core with a **two-pass assembler**, a **golden instruction-set simulator**, and an **automated co-simulation regression** checking RTL against the ISS every run. Extended toward out-of-order execution with register rename, ROB and issue queue (in progress).
- **RTL-to-GDSII, Sequential MAC** (OpenLane, SkyWater sky130, 2024): independently carried a 16×16 shift-and-add multiplier through synthesis, floorplan, placement, CTS, routing, parasitic extraction and multi-corner STA. Closed at **5.53 ns setup / 0.48 ns hold slack, 8,168 μm^2 , 2.59 mW** with **zero Magic/KLayout DRC violations and clean LVS**.
- **CNN Inference Accelerator on FPGA** (B.E. thesis, Vivado HLS, 2023): built an AlexNet accelerator on a Zedboard (Zynq-7020) with a pipelined core and kernel/bias/input/output caches, fixed-point widths sized via Ristretto. **36,527 LUTs, 54 BRAM, 35 DSP at 0.74 GFOPS and ~0.9 W** — a 7–64× smaller DSP footprint and ~10× lower power than comparable published accelerators.
- **ML-Driven Last-Level Cache Partitioning for CPU-GPU SoCs** (2025): proposed a hardware perceptron predicting per-core cache sensitivity from IPC, MPKI and memory-queue occupancy to set LLC way-masks every 1M cycles, replacing a static split that bought **+8.36% CPU at –7.5% GPU**. Instrumented **gem5** with probing counters for motivational data.
- **6-Stage Pipelined RISC Processor** (VHDL, 2025) and **SAT-based ATPG for the 74181 ALU** (2026): full datapath with forwarding and interlocks; Tseitin CNF + MiniSat test generation with fault collapsing, cutting ATPG effort **3.67×**.

EXPERIENCE

Research Assistant — Wadhvani Electronics Laboratory, IIT Bombay

Jul 2024 – Present

Supported Communication Lab and Electronic Design Lab; trained incoming TAs on the experiments; built and maintain the lab website (GitHub Pages, now served from the lab's internal server).

TECHNICAL SKILLS

Languages: SystemVerilog, Verilog, VHDL, C++, C, Python, CUDA

Architecture & Simulation: gem5, GPGPU-Sim, AccelSim, MGvm, Chipyard/BOOM, SpikeSim

EDA & Flow: OpenLane, Yosys, OpenSTA, Magic, KLayout, sky130 PDK, Xilinx Vivado & Vivado HLS, Quartus Prime, GTKWave

Platforms & ML: PYNQ / Zynq-7020, AXI-Lite & AXI-DMA, PCB design, GNU Radio; PyTorch, scikit-learn

SELECTED ACHIEVEMENTS

- **Badminton:** Winner, IIT Bombay PG General Championship (2026) and Institute Badminton League (2025); represented IIT Bombay at Inter-IIT (runners-up, 2024) and BITS Goa at Inter-BITS (winners, 2024).
- **Quizzing:** Winner — NTPC Electron Quiz Western Regionals (2025), Zerodha Varsity Quiz Mumbai (2024, 2026), Tata Crucible Goa (2021); Finalist, Nihilanth Inter IIT-IIM quiz festival.