



# RUTWIK JAIN

## Curriculum Vitae

+1 608-982-5056   
rnjain@wisc.edu   
cs.wisc.edu/~rnjain   
Github   
LinkedIn 

### Education

#### 2021 **Ph.D. Computer Science**

↓ University of Wisconsin–Madison

GPA – **3.82/4.0**

Present

- Ongoing Research: Variability-aware scheduling in GPU-based HPC systems
- Advisors: Prof. Matt Sinclair and Prof. Shivaram Venkataraman
- Coursework: Advanced Computer Architecture, High Performance Computing and Applications in Engineering, Computer Graphics, Operating Systems, Distributed Systems, HW/SW for Next Gen AI Systems, Fault-Tolerant Computing, Optimization, Data Visualization

#### 2015 **B.E.(Hons.) Electrical and Electronics Engineering**

↓ Birla Institute of Technology and Science, Pilani

GPA – **9.52/10.0**

2019

- Class Rank: 3/93
- Coursework: Computer Architecture, Embedded Systems Design, Digital Design, Digital Signal Processing, Microprocessors and Interfacing, Analog and Digital VLSI Design

### Publications

#### In Conference Proceedings

- 2022 Praseon Sinha, Akhil Guliani, **Rutwik Jain**, Brandon Tran, Matthew D. Sinclair, and Shivaram Venkataraman. Not all gpus are created equal: Characterizing variability in large-scale, accelerator-rich systems. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC '22*. IEEE Press, 2022.
- 2019 **Rutwik Jain** and Umang Garg. Modelling and simulation of onboard wire antennas for a 3U CubeSat (as undergraduate). In *2019 Photonics and Electromagnetics Research Symposium - Fall (PIERS - Fall 2019)*, pages 871–879, Dec 2019.
- 2018 **Rutwik Jain**, Shubham Sharma, Kushagra Aggarwal, Tushar Goyal, Abhinav Sundhar, Ujjwal Anand, Nishad Sahu, and Joy Parikh. Modes of operation for a 3u cubesat with hyperspectral imaging payload. In *Proceedings of the 69th International Astronautical Congress (IAC) 2018, Bremen, Germany*, Oct 2018.

#### Journal Articles

- 2023 Matthew D. Sinclair, Parthasarathy Ranganathan, Gaurang Upasani, Adrian Sampson, David Patterson, **Rutwik Jain**, Nidhi Parthasarathy, and Shaan Shah. Fifty years of the international symposium on computer architecture: A data-driven retrospective. *IEEE Micro*, volume 43, pages 109–124. IEEE Computer Society, nov 2023.
- 2020 Anuj K. Ojha, **Rutwik Jain**, and Praveen Kumar AV. Magnetic quadrupole mode excitation of a cylindrical dielectric resonator antenna using planar feed. *Microwave and Optical Technology Letters*, volume 62, pages 484–490, 2020.

### Work Experience

#### Jul 2019 **ASIC Design Engineer, NVIDIA**

↓ Memory Subsystem Design Team

Bangalore, India

Jul 2021

- Microarchitecture design and RTL implementation for the SoC's memory management IP.
- This included new feature implementation, modification of existing  $\mu$ Arch for PPA optimization, logic synthesis for area and timing accuracy, iterative design to incorporate DV and PD feedback.

- Jan 2019 **ASIC Intern, NVIDIA**  
↓ Memory Subsystem Design Team Bangalore, India
- Jun 2019
  - Developed a debug tool that uses a generated waveform database to identify, trace and correlate transactions as they flow through the Memory Management Unit (MMU).
  - Presented my work on developing and scaling the tool in a lightning talk at India-level internal NVIDIA conference, NTech 2019.

## Projects

- Aug 2022 **Variability-aware scheduling for large-scale, GPU-based systems**  
↓ Advisors: Prof. Matt Sinclair and Prof. Shivaram Vankataraman  
Present
  - Ongoing research to develop a variability-aware scheduler that makes application-specific GPU allocations when scheduling jobs in large-scale GPU-based systems, based on PM-related information from hardware.
- Aug 2021 **Characterizing Variability in Large-scale, Accelerator-rich Systems**  
↓ Advisors: Prof. Matt Sinclair and Prof. Shivaram Vankataraman
- Aug 2022
  - Characterized GPU variability in HPC systems and analyzed how variability is impacted by various cluster attributes like size, cooling, and GPU vendors (compared ORNL Summit, TACC Longhorn and Frontera, SNL Vortex and LLNL Corona)
  - Profiled different applications (image processing, machine translation, molecular dynamics, graph analytics) on the same cluster to evaluate application-dependence.
  - Co-authored and presented our work at Supercomputing (SC), 2022.
- Aug 2023 **Parallelized Implementation of Coordinated Animal Motion – Boid Flocking**  
↓ Course: CS 759 High Performance Computing and Applications in Engg, Fall 2023
- Dec 2023
  - Optimized a CPU-based C++ implementation of flocking behavior through OpenMP and CUDA implementations to compare frame rate performance and scalability.
  - Implemented a CUDA kernel and shared memory optimizations that allowed flocking simulations to scale gracefully as number of particles increases upto  $2^{14}$ .
- Aug 2022 **PICLE: Prefetcher for Inter-CU Locality Exploitation in GPU TLBs**  
↓ Course: CS 757 Adv Computer Architecture II, Fall 2022
- Dec 2022
  - Implemented a GPU L1-TLB prefetcher based on Valkyrie [Baruah et.al, PACT'20] which exploits inter-TLB locality to improve GPU virtual memory performance using gem5.
- Jan 2022 **Project Eagle: Open-source HW/SW Stack for AI Applications**  
↓ Course: CS 838 Next-Gen HW/SW Design for AI/ML, Spring 2022
- May 2022
  - Implemented RTL for a dual-ported, lockup-free L1 cache with prefetching for the memory system of UPCYCLE, a tiled accelerator architecture.
- Jan 2022 **Replicated Block Store**  
↓ Course: CS 739 Distributed Systems, Spring 2022
- May 2022
  - Implemented a primary-backup replica based distributed block store.
  - Analyzed and reported read and write performance, crash recovery and reintegration time, and evidence for correctness and linearizability guarantees.

## Skills

Programming: C, C++, OpenMP | Scripting: Python | RTL Design: Verilog  
GPU computing: CUDA | HPC: MPI | Optimization: Julia & JuMP  
Cycle-level Simulators: gem5 | RTL Design and Debug: ModelSim, VCS Synopsys  
Version Control: Git | Containers: Docker, Singularity | Visualization: Tableau

## Talks & Presentations

- Jan 2024 **GPU Variability: Characterization Study and Mitigation Strategies**, AMD Research.
- Nov 2022 **Characterizing GPU variability in Large-Scale, Accelerator-Rich Systems**, SC 2022
- Oct 2022 **Variability-Aware Scheduling for Efficient Exascale Computing**, UW-Madison's Architecture Affiliates 2022.

Oct 2018 **Modes of Operation of a 3U CubeSat with Hyperspectral Imaging Payload**, 69th International Astronautical Congress (IAC) 2018

## Teaching

FA 21 **CS/ECE 552: Introduction to Computer Architecture**

SP 22 Taught Verilog and HDL concepts in discussion section. Also helped students with in-class exercises and debugging projects (pipelined processor and cache implementation)

FA 22 **CS/ECE 252: Introduction to Computer Engineering**

Assisted students with homework and projects through office hours, graded homework assignments.

## Awards & Honors

- Received CS Summer Research Assistantship from UW-Madison for the 2022 summer term.
- First-year CS Departmental Scholarship from UW-Madison for academic year 2021-22.
- Merit Scholarship at BITS Pilani: 40% tuition waiver for seven semesters from 2015 to 2019.