Hritvik Taneja

htaneja3@gatech.edu | 4709203994

Education

Georgia Institute of Technology, Atlanta, GA
Ph.D. in Computer Science, 2022 - 2027 (Expected)

Indian Institute of Technology Kanpur, India
B.Tech. in Computer Science and Engineering, 2016-2020, CGPA: 8.8/10

Publications

1. **Hritvik Taneja**, Ali Hajiabadi, Michele Marazzi, Kaveh Razavi, Moinuddin Qureshi "MIRZA: Efficiently Mitigating Rowhammer with Randomization and ALERT" *HPCA, 2026*
2. **Hritvik Taneja**, Moinuddin Qureshi "DREAM: Enabling Low-Overhead Rowhammer Mitigation via Directed Refresh Management" *ISCA, 2025*
3. **Hritvik Taneja**, Jason Kim, Jie Jeff Xu, Stephan van Schaik, Daniel Genkin, Yuval Yarom "Hot Pixels: Frequency, Power, and Temperature Attacks on GPUs and Arm SoCs" *USENIX Security, 2023*
4. Arpit Gupta, Parv Mor, **Hritvik Taneja**, Biswabandan Panda "STEVES: Pushing the Limits of Value Predictors with Sliding FCM and EVES" *Championship Value Prediction, 2019*

Research Experience

Graduate Research Assistant December 2023 - Present
Georgia Institute of Technology Atlanta, GA
Advisor: Professor Moinuddin Qureshi
Research Focus: Memory system optimizations for ML workloads

Graduate Research Assistant August 2022 - November 2023
Georgia Institute of Technology Atlanta, GA
Advisor: Professor Daniel Genkin
Research Focus: Side-channel attacks

Professional Experience

Research Intern May 2025 - Aug 2025
AMD Research Austin, TX
- Optimized data locality for LLM workloads on AMD's multi-chiplet GPUs.
- Applied tensor parallelism (TP) across chiplets to maximize locality and reduce data movement.
- Identified that synchronization latency in All-Reduce becomes a bottleneck when applying TP across high-bandwidth inter-chiplet interconnects.

Advanced Design Verification Intern May 2024 - Aug 2024
Astera Labs Santa Clara, CA
- Developed a tool for large-scale Verilog simulations on an AWS-backed cluster.
- Focused on improving the cost efficiency by reducing the unused memory and utilizing spot instances.

Infrastructure Developer Jan 2020 - Dec 2021
Plutus Research Private Limited Bengaluru, India
- Improved latency of market-data arrival to high-frequency trading (HFT) strategies – by 20% in median and by 30% in 90 percentile.
- Developed an approximate backtesting framework that uses snapshots of the price-level order book to generate trade responses — with minimal deviation in the fill ratio.

Software Engineering Intern May 2019 - Jul 2019
Nutanix, Inc. Bengaluru, India
- Improved the average turnaround time of the Nutanix Calm API by 80 fold, by caching responses.

- Use gossip protocol to ensure cache consistency, in a scaled-out setup.
- Worked on bug fixes and implemented new features in frontend and backend, as part of the Nutanix CALM team.

Key Projects

Hardware Security Lab August 2022 - February 2023
Georgia Institute of Technology Atlanta, GA
- Discovered and evaluated the data-dependent nature of frequency scaling on GPUs and Arm SoCs.
- Exploited this phenomenon to mount Pixel Stealing, History Sniffing, and Website Fingerprinting attacks on various devices.
- Apple assigned us CVE-2023-38599 for our work.

Championship Value Prediction Dec 2018 - Jun 2019
Indian Institute of Technology Kanpur Kanpur, India
- Proposed and built an FCM based value predictor, which achieves an improvement of 1.4% over state-of-the-art EVES predictor.
- Secured first (currently third) position on the unlimited track leaderboard of CVP from May'19 to November'19.

NINE LLCs in Multicore Processors Aug 2019 - Nov 2019
Indian Institute of Technology Kanpur Kanpur, India
- Designed a non-inclusive non-exclusive LLC, which marks each block either as inclusive or exclusive in a multi-core processor.
- Implemented a cache simulator equipped with MESI cache coherence protocol to analyze performance counters like messages & misses.
- Improved the performance in terms of private cache misses, LLC misses, and number of interconnect messages – **Report**.

Inter-Procedural Data-flow Analysis Aug 2019 - Nov 2019
Indian Institute of Technology Kanpur Kanpur, India
- Ported and implemented the inter-procedural analysis from [**PadhyeK13**] to LLVM (originally in Soot).
- Modified the framework [**PadhyeK13**] to account for fundamental differences between Soot and LLVM
- Designed and implemented a sign data-flow analysis using LLVM – **Code**.

Golang Compiler Jan 2019 - April 2019
Indian Institute of Technology Kanpur Kanpur, India
- A compiler for go written in go as a course project for compilers. Compiles from golang to x86 assembly.
- Supports a subset of the go language, including pointers, type checking, recursion, and some other common language features – **Code**.

Skills

Proficient: C/C++, Python, Pytorch