PRATHEEK B

+91 9746406949
website
pratheekb@iisc.ac.in
pratheekbksd@gmail.com

ABOUT

Pratheek is a Senior Software Engineer at NVIDIA, working on improving GPU memory management. He obtained his Ph.D. from the Indian Institute of Science, Bengaluru in 2026. His research revolves around improving GPU memory management, focusing on address translation and data placement for large-memory workloads. His work spans the domains of GPU micro-architecture, compiler techniques, and operating systems. Previously, he had worked on reverse-engineering Nvidia GPUs. He also maintains an active interest in CPU micro-architecture research.

EDUCATION

Ph.D. in Computer Science **2020 - 2026**
Indian Institute of Science, Bengaluru GPA : 8.0/10
Thesis : *Scaling Up GPU Memory Management*
Advised by: *Prof. Arkaprava Basu*

M.Tech. in Computer Science **2017 - 2019**
Indian Institute of Science, Bengaluru GPA : 8.4/10
Major project: *Page walk stealing in multi-tenant GPUs*

B.E. in Computer Science **2012 - 2016**
NMAM Institute of Technology, Nitte GPA : 8.9/10
Major project: *Utilities for porting TCP Evaluation Suite for NS3*

PUBLICATIONS

ISCA'26 **SuperUVM: Enhancing GPU memory oversubscription with observability**
Pratheek B., Khushit Shah, Arkaprava Basu
To appear at the 53rd International Symposium on Computer Architecture

MICRO'24 **SUV: Static analysis guided Unified Virtual Memory** [link]
Pratheek B., Guilherme Cox, Jan Vesely, Arkaprava Basu
In Proceedings of 57th ACM/IEEE International Symposium on Microarchitecture

MICRO'22 **Designing Virtual Memory System of MCM GPUs** [link]
Pratheek B.*, Neha Jawalkar*, Arkaprava Basu
In Proceedings of 55th ACM/IEEE International Symposium on Microarchitecture

HPCA'21 **Improving GPU Multi-tenancy with Page Walk Stealing** [link]
Pratheek B.*, Neha Jawalkar*, Arkaprava Basu
In Proceedings of 27th IEEE Intl. Symposium on High-Performance Computer Architecture
Best Paper Nominee

AsiaCCS'21 **(Mis)managed: A Novel TLB-based Covert Channel on GPUs** [link]
Ajay Nayak, **Pratheek B.**, Vinod Ganapathy, Arkaprava Basu
In Proceedings of 16th ACM Asia Conference on Computer and Communications Security

AWARDS AND ACHIEVEMENTS

Intel PhD Fellowship 2023-24
Best paper nominee HPCA 2021
All India Rank 10 out of ∼100,000 candidates Graduate Aptitude Test in Engineering 2017

PROFESSIONAL EXPERIENCE

Sr. Software Engineer NVIDIA
Improving GPU memory management Feb 2026 - now

Research Intern NVIDIA
Improving access counter based migration (ACBM) May 2022 - July 2022
· Analyzed different configurations of ACBM across several GPU workloads
· Implemented techniques to automatically determine the best configuration

Project Assistant IISc, Bengaluru
Page walk stealing in multi-tenant GPUs July 2019 - Dec 2019
· Studied page walker contention in multi-tenant GPUs
· Implemented and evaluated page walker sharing mechanisms optimized for throughput and fairness

Research Intern IIT Delhi
Improving parallel TSP solver June 2016 - Aug 2016
· Analyzed bottlenecks in an implementation of parallel Traveling Salesperson Problem solver
· Proposed improvements to the implementation of parallel TSP solver

SERVICE

PC member MICRO 2026
Shadow PC member Eurosys 2025
Sub-reviewer HPCA 2025
Sub-reviewer TDSC 2022

OTHER PROJECTS

Impact of multi-tenancy on UVM-based GPU workloads Course project, 2021
· Characterized behavior of unified memory (UVM) workloads under multi-tenancy
· Explored impact of varying degrees of multi-tenancy on throughput

Replacement Policies for DRAM Caches Course project, 2017
· Studied replacement policies proposed in literature for DRAM caches
· Proposed and empirically evaluated a new dynamic replacement policy for DRAM caches

Utilities for Porting TCP Evaluation Suite to NS3 Bachelors project, 2016
· Ported TMix and Delaybox network tools to the latest version of NS3 at the time
· Built prototypes for 2 of the 4 modules of the TCP Evaluation Suite

PintOS Course project, 2018
· Implemented a bare-bones OS, including scheduling, virtual memory and file system

TEACHING EXPERIENCE

Teaching assistant High Performance Computer Architecture, 2022
· Designed and evaluated assignments on GPU programming

MENTORSHIP EXPERIENCE

Improving page migration for UVM workloads IISc, Bengaluru

· Mentored M.Tech student working on improving page migration for unified memory (UVM) workloads

Memory management in GPU multi-tenancy GPUs IISc, Bengaluru

· Mentoring M.Tech. students working on analyzing the impact of multi-tenancy on UVM workloads

TALKS

Improving GPU memory management Talk at Intel, Bengaluru
Reverse engineering NVIDIA GPU TLBs Security workshop at IIT, Bombay
Improving UVM with static analysis Talk at IIT, Gandhinagar
Designing virtual memory for MCM GPUs SysTalks@India

TECHNICAL SKILLS

Programming	C, C++, Go, Python, CUDA
Tools	eBPF, LLVM, NVIDIA UVM driver, GDB, NSight, perf
Simulators	GPGPU-Sim, MGPUSim, NS-3

RELEVANT COURSES

Computer Architecture, Topics in Computer Architecture, Operating Systems, Computer Systems Security, System Virtualization, Parallel Programming