

Vishal Rao

Department of Computer Science
Georgia Institute of Technology

Email: vishal.rao@gatech.edu | Personal Website: kadle11.github.io | Phone: (+1) 470-430-6936

Research Interests

Memory Disaggregation, Vector Databases, Systems for AI and Emerging Hardware, Graph Analytics

Education

Georgia Institute of Technology

Ph.D in Computer Science | Advisor: Ada Gavrilovska 2023 – 2028 (expected)

M.S in Computer Science 2021 – 2023

PES University, Bengaluru

B.Tech in Computer Science & Engineering | Advisor: Subramaniam Kalambur 2016 – 2020

Publications

HPDC '25 Grudon: A System for Deploying Graph Workloads on Disaggregated Architectures with Near-Data Processing

Vishal Rao, Nikhil Ram Shashidhar, Suyeon Lee, and Ada Gavrilovska

The 31st International Symposium on High-Performance Parallel and Distributed Computing

RESDIS '24 Towards Disaggregated NDP Architectures for Large-scale Graph Analytics

Suyeon Lee*, Vishal Rao*, and Ada Gavrilovska (*equal contribution)

The 4th International Workshop on RESource DISaggregation in High-Performance Computing

JSSPP '21 Scheduling microservice containers on large core machines through placement and coalescing

Vishal Rao, Vishnu Singh, KS Goutham, Bharani Ujjaini Kempaiah, Ruben John Mampilli, Subramaniam Kalambur, Dinkar Sitaram

24th International Workshop on Job Scheduling Strategies for Parallel Processing

Research Experience

Research @ Kernel Research Group | GaTech

May '23 – present

Research Assistant | Advisor: Ada Gavrilovska

- Designed systems to enable smart offload to emerging platforms with near-data processing.
- Designed systems to reduce data movement overheads for graph analytics/vector search.

SoLab | SK Hynix America

Sep '26 – Dec '26

AI/HPC Systems Research Intern | Guides: Jongryool Kim

- Exploring CXL-based prefetching mechanisms to improve RAG performance.
- Designing systems for hardware platforms with disaggregated memory and compute.

Systems Architecture Lab | HPE Labs

May '25 – Dec '25

Research Associate Intern | Guides: Sekwon Lee and Harumi Kuno

- Designed a system to improve Linux kernel page migration performance.
- Designed a multi-level cache for graph-based vector embeddings.

Server Performance Group | AMD

Jun '20 – Jun '21

Research Intern | Guide: [Kanishka Lahiri](#)

- Characterized the impact of Linux kernel security mitigations on EPYC platforms.
- Investigated Linux kernel hotspots for EPYC Gen 3.

Research @ Center for Cloud Computing | PES University

Aug '17 – Aug '21

Research Assitant | Guides: [Subramaniam Kalambur](#), [Reetinder Sidhu](#)

- Designed a NUMA-aware container scheduler that reduces L3 cache contention.
- Designed chunked GEMM algorithms for custom FPGA-based accelerators.

Awards & Achievements

- SRC PRISM Annual Review '25: Best Poster Runner-up
Grudon: Deploying Graph Workloads on Disaggregated Architectures with Near-Data Processing
- 4x recipient of C. N. R. Rao Scholarship, CS Department, PES University

Service

- Artifact Evaluation Committees: SOSP '25, OSDI '25, ATC '25, OSDI '24, ATC '24
- Shadow Program Committees: EuroSys'26
- Student Volunteer: OSDI'24
- GT SCS Graduate Student Association: President (2024 - 2025)

Teaching

- CS 7210 Distributed Systems, TA Fall 2025, Fall 2022, Spring 2022
- CS 352 Cloud Computing, TA Spring 2020

Talks

- **Towards Disaggregated NDP Architectures for Large-scale Graph Analytics**
SAL Tech Talk Series, HPE Labs July 2025
SRC PRISM Liason Talk May 2025