

PRATHAM SAHU

Final Year Undergraduate

Department of Computer Science and Technology, Indian Institute of Technology Kanpur

☎ +91 7619678791 ✉ spratham21@iitk.ac.in 💻 [pratham-sahu](https://pratham-sahu.github.io) 🐙 [Prathamsahu52](https://Prathamsahu52.github.io) 🌐 prathamsahu52.github.io

Education

Bachelor of Technology, Computer Science and Engineering

2021 - present

Indian Institute of Technology Kanpur

CPI: 9.31/10

Scholastic Achievements

- **All India Rank 131** of 1,50,000 shortlisted candidates, **JEE Advanced 2021**, National Entrance Exam for IITs.
- **All India Rank 87** of 1.1 million candidates, **JEE Mains 2021**, Engineering entrance exam.
- Accepted into the **HPC Student Cohort**, invited to attend SC'24, IPDPS'25 by **IEEE-CS TCHPC** and **TCPP**.
- **Gold Medal**, INPhO 2021, **All India Rank 44**, qualified for National **IPhO Camp**.
- **Gold Medal**, INAO 2021, **All India Rank 19**, qualified for National **IOAA Camp**.
- **Gold Medal**, INChO 2021, **All India Rank 31**, qualified for National **ICHO Camp**.
- **Directors Scholarship**, IIT Kanpur 2022, awarded for having an exceptional **JEE Advanced** rank.
- **Fellowship Awardee**, KVPY SA 2019, secured **All India Rank 340**, conducted by **IISC, Bangalore**.
- **Fellowship Awardee**, KVPY SX 2020, secured **All India Rank 148**, conducted by **IISC, Bangalore**.
- **Academic Excellence Award**, received for exceptional performance in 3 consecutive academic sessions (2021-2024).
- **Scholarship Recipient**, National Talent Search Examination (NTSE) 2019, awarded by **NCERT**.

Research Projects and Competitions

ISC Student Cluster Competition

Aug '23 - May '24

[ISC Student Cluster Competition ↗]

Prof. Preeti Malakar & Prof. Swarnendu Biswas

- Selected as the **first Indian team** to participate, finishing **8th out of 22 teams** as the **only new team**.
- Led a team of six third-year students in the world's **largest student HPC Competition**
- Improved performance by **20x** in the μ phys package within the **ICON** application by optimizing code with **OpenACC directives** for **CUDA portability**. Submitted a pull requested to the upstream repository of the application.
- Demonstrated effective **scaling** of the **NekoCFD** application on CPU and GPU backends, visualizing results using **paraviewCLI** and profiling using **Nvidia Nsight, Tau** on the **Bridges2 supercomputer**.

Extended Dirty Page tracking in X86 Systems

July '24 - current

Submitted to ISCA'25

Prof. Debadatta Mishra

- Proposed a novel **hardware-software co-design** to extend the **dirty page tracking** mechanism in x86 systems.
- Improved dirty-page tracking granularity to **64 bytes** from the existing **4KB** granularity in x86.
- Implemented TLB modules and **Page Table Entry** modifications to track dirty pages per pmd level in a seperate page.
- Implemented baseline designs on the gem5 simulator, linux kernel and conducted extensive testing on real applications.

Controlled Inter-thread Memory Sharing in Multi-Threaded Applications

Jan '24 - May '24

[Report ↗]

Prof. Debadatta Mishra

- Designed a **novel** memory framework for **threads to achieve privilege seperation** in the same address space.
- Used **TLB entries** to boost privileges temporarily on cores with threads having higher privileges for the same page.
- Used **inter-processor interrupts** to stop other cores from accessing boosted Page Table entries and thus ensuring correctness. Optimized for efficiency using **TSX** provided by a few older x86 systems and critically analyzed the same.
- Implemented the design on the **Linux kernel using modules, core kernel code**, and provided **user-space APIs**.
- Conducted thorough testing and benchmarking, verifying correctness and assessing memory access time trade-offs.

Modelling Performance Variability in HPC Clusters

Jan '24 - current

[Report ↗]

Prof. Preeti Malakar

- Received an **A*** grade for exceptional contributions to research in **performance and variability modeling in HPC**.
- Analyzed job interference impact on performance variability in a production supercomputer(**Param Sanganak**).
- Characterized supercomputer jobs using profilers (IPMPI), I/O tracing, network tracing, and hardware counters (perf).
- Designed a novel algorithm to mitigate variability due to job interference, integrated it on the **Slurm Scheduler**.

Work Experience

Research Intern - Data-driven Systems, Insights, Experiences team

May '24 - Jul '24

Adobe Research, Big Data Experience Lab, Bengaluru

- Finetuned **LLMs like llama03** for generating **data-driven, domain-aware strategies** using customer insights.
- Modeled **customer journeys as graph walks** to identify optimal strategies for enhancing engagement.
- Used **interpretable multivariate forecasting** to evaluate the effectiveness of these strategies, driving actionability.

Intern Research Assistant

May '23 - Sep '23

Yonsei Vision and Learning Laboratory, Seoul, South Korea

Prof. Jonghyun Choi

- Developed effective deep learning algorithms for **continual learning** in autonomous driving scenarios.
- Utilized **recent coreset techniques for streaming data** to mitigate forgetting in incremental deep learning setups.
- Devised a novel technique for evolving existing episodic-replay methods in incremental classification using ensembling.

Platform Software Engineering Intern

Jul '22 - Aug '22

MSenseAI, Bengaluru

- Implemented open source IoT Platform **ThingsBoard** adding functionalities to suit product-specific requirements.
- Imparted picture annotation functionalities to the platform by integrating Annotorius, an open-source JS library.

Relevant Coursework

(*:ongoing)

- | | | |
|-------------------------------|----------------------------------|-----------------------------|
| • Linux Kernel Programming | • Networks | • Operating Systems |
| • Database Management Systems | • Advanced Algorithms | • Theory of Computation |
| • Parallel Computing (A*) | • Data Structures and Algorithms | • Probability |
| • Compilers Design | • Mathematical Logic | • Software Development |
| • Computer Organisation | • Fundamentals of Computing | • Large Data Analytics (A*) |

Selected Projects

PuppyLove2.0 - Campus Dating App

Jan '23 - Apr '23

Programming Club

[Code]

- Designed a **zero-knowledge system** for a dating application using vanilla symmetric and asymmetric key encryption.
- Developed the application using **Golang+PostgreSQL** for the backend framework and **NextJS** to build the frontend.
- Deployed the application using **Kubernetes for microservice orchestration** with security measures like re-captcha and rate-limiting to prevent large-scale DOS attacks on the server, and browser caching to release load on server.
- Achieved registration of **2,800 users** with over **600 maximum concurrent users**.

Python Compiler

Jan '24 - Apr '24

Compiler Design Course Project, Prof. Swarnendu Biswas

[Code]

- Awarded **97% score** for building an end-to-end compiler to convert a subset of **Python** language to **x86 assembly**.
- Designed lexer using flex, parser using bison and semantic analyzer to support for **classes, multilevel inheritance, arrays, and non-primitive types** in python. Included support for **type checking and implicit type conversion**.
- Used a **hierarchical symbol table** and **Abstract Syntax tree based IR** to form a **3AC IR**, which was then converted to x86 assembly code using template functions. Ensured chaining support across data structures.

CSE-Bubble

Jan '23 - Apr '23

Computer Organisation Course Project, Prof. Urbi Chatterjee

[Code]

- Implemented a **Verilog hardware description** of a **32-bit processor**, featuring ISA, ALU, and a memory unit.
- Verified correctness by running **RISC-V(MPI)** assembly code for sorting algorithms on the on a FPGA board.

Building GemOS

Aug '23 - Dec '23

Operating Systems Course Project, Prof Debadatta Mishra

[Code]

- Developed support for **system call tracing and function call tracing** on a toy OS(GemOS) with x86 architecture.
- Designed and integrated **mmap, munmap, and remap** semantics, ensuring minimal fragmentation of virtual memory.
- Implemented **regular (copy all)** and **Copy-on-Write (CoW)** semantics for fork handling within the toy OS.

CampusPay

Jan '23 - Apr '23

Software Design Course Project, Prof. Indranil Saha

[Code]

- Developed a website to handle finances and dues for the campus community and vendors in structured manner.
- Wrote the backend in **Django** and used **postgresql** for DB. Frontend was implemented in the **ReactJS framework**.
- Implemented a **client-server** model for the application and used push notification model for notification sub-system.

Teaching

Tutor

Jul '24 - Nov '24

ESO207, Data Structures and Algorithms

IIT Kanpur

- Tutored a fundamental data structures and algorithms course with 380 registered students.
- Conducted **weekly office hours** for doubt solving and managed course logistics including grading.

Technical Skills

Programming Languages: C, C++, Python, Java, Javascript, Solidity, RUST, CUDA, DPC++

Libraries: Numpy, Pandas, Matplotlib, MERN stack, NextJS, PyTorch, Django

Utilities: Linux Kernel, Kubernetes, Git, gdb, perf, Tau, Nvidia Nsight

Leadership and Service

Coordinator

Jun '23 - May '24

Programming Club

IIT Kanpur

- Managed one of the most active clubs of IIT Kanpur which delves into the multiple domains of programming
- Contributed to open-source projects targeting the campus community, such as **StudentSearch** and **PuppyLove**
- Organised **Linux install fest**, **Spring Camp** (2 weeks of workshops across 5 domains) for the community.
- Mentored 15 students to get introduced to the concepts of Operating Systems and understanding the Linux Kernel.

Secretary

Jun '22 - May '23

Programming Club

IIT Kanpur

- Assisted in conducting competitions, activities, lectures, and workshops for programming enthusiasts in IIT Kanpur.

Overall Coordinator

Jun '24 - Sep '24

Opportunity Open Source Conference

IIT Kanpur

- Organised a flagship student Open Source Conference for the first time in India with **52 speakers** from around the world and **900 registered attendees** from across Indian Institutions.
- Scheduled 70 talks on multiple domains and conducted an overnight hackathon with 46 teams during the event.

Student Guide

Sep '22- July '23

Counselling Services

IIT Kanpur

- Mentored a group of 6 freshmen academically and emotionally to get acclimatized to the new college environment.