

Karan Kumar Gangadhar

☎ +1 929-691-9190 | ✉ kk5409@nyu.edu | 🌐 karan-kumarg | 🎮 Spielerr | 📍 New York, USA

EDUCATION

New York University (Courant Institute of Mathematical Sciences)

New York, NY

Masters of Science in Computer Science; CGPA: 4.00/4.00

Sep 2023 – May 2025

- **Relevant coursework:** Honors Analysis of Algorithms, Programming Languages, GPUs-Architecture and Programming, Operating Systems, Deep Learning, Research study on Programmable Network Switches Transpilation with **Prof. Anirudh Sivaraman**

PES University

Bengaluru, India

Bachelor of Technology in Computer Science and Engineering; CGPA: 3.98/4.00

Aug 2018 – May 2022

Specialization in Systems and Core Computing, Class Top 10%

- **Relevant coursework:** Computing Using Python, Engineering Mathematics, Problem Solving with C, Digital Design & Computer Organization, Data Structures, Introduction to Data Science, Web Technologies, Discrete Mathematics & Logic, Operation Linux, Design and Analysis of Algorithms, Database Management Systems, Microprocessor & Computer Architecture, Theory of Computation, Linear Algebra, Computer Networks, Operating Systems, Machine Intelligence, Advanced Algorithms, Big Data, Generic Programming, Design Patterns, Compiler Design, Cloud Computing, Performance Engineering, Computer Forensics, Advanced Java

SKILLS

Programming Languages: C, C++, Java, Python, Scala, P4

Databases: MySQL, MongoDB, Elasticsearch, Redis

Tools and Platforms: Git/GitHub, Linux/Unix, Docker, Kubernetes, Jenkins, GCP, AWS, Prometheus, Grafana, ELK Stack

Frameworks and Libraries: GDB, Valgrind, Bash Scripting, SpringBoot, Django, Flask, Node.js, PyTorch, Scikit-Learn

Other Technical Skills: Parallel Programming using OpenMP, CUDA and C++ threads, Compiler Optimisation using LLVM and MultiIR, Deep Learning Architecture and Programming using PyTorch, Network Switch Architecture and Programming using P4, CI/CD Pipelines, DevOps Practices, Microservices Architecture, RESTful API Development, Agile and Scrum Methodologies

Certifications: Architecting with Google Compute Engine Specialization

WORK EXPERIENCE

Goldman Sachs

Bengaluru, India

Analyst - Software Engineer

Jul 2022 – Jul 2023 (1 year)

- Developed and deployed a Java SpringBoot-based backend microservice with **RESTful APIs**, integrating MySQL and MongoDB databases for the Asset Management Dashboard. The client transaction processing time was **reduced from one week to one day** by eliminating manual spreadsheet usage.
- Implemented a **Reliability Monitoring Service** using Prometheus and Grafana, hosted on Kubernetes, to track the availability and latency of hosted services. This automation **improved application performance by over 75%** by identifying and addressing latency spikes, eliminating the need for manual failure tracking, and leveraging Agile methodologies.
- Designed and integrated new **CI/CD pipelines** using Jenkins, Docker, and Kubernetes for streamlined deployment and monitoring processes.
- **Recognised for excellence** through a selected presentation and demo of completed work to the head of the organisation.

Cisco

Bengaluru, India

Technical Undergraduate Intern

Feb 2022 – July 2022 (6 months)

- Introduced **C unit testing automation** using the Cmocka framework, significantly enhancing the ease and speed of adding new unit tests across multiple components in the Cisco IOS-XR Network OS. This initiative **reduced human effort by over 80%**.
- Integrated unit testing with the CI/CD pipeline to **automate pre-commit and nightly checks**, resulting in a **code test coverage improvement to over 95%** for selected modules.
- **Awarded "Best Intern Project"** for outstanding contributions to unit testing automation.

Goldman Sachs (Remote)

Bengaluru, India

Summer Technical Intern

Jun 2021 – Jul 2021 (2 months)

- Enhanced text-based search performance of a web application by implementing **database caching with Elasticsearch**. This optimization **accelerated search speed by 150%**, improving user experience and efficiency.
- **Redesigned backend services** to support new functionalities, enhancing code quality and introducing common **design patterns** for improved code reusability and maintainability.
- Received the **"Excellence in Engineering"** award for outstanding enhancements in search performance.

SELECTED PUBLICATIONS

Cross-Platform Transpilation of Packet-Processing Programs using Program Synthesis

Xiangyu Gao, Jiaqi Gao*, **Karan Kumar Gangadhar***, Ennan Zhai, Srinivas Narayana, and Anirudh Sivaraman

Accepted in Asia-Pacific Workshop on Networking (APNet 2024) (Co-located with SIGCOMM 2024)

RESEARCH EXPERIENCE

NYU Courant Systems Lab | Master's Thesis

New York, USA

Cross-Platform Transpilation of Packet-Processing Programs using Program Synthesis

Jan 2024 – Present (8 months)

- Developing a **first-of-its-kind transpiler** to automatically convert packet processing programs between multiple **network switch architectures** and supported programming languages such as **P4 and NPL**, using the technique of Program Synthesis.
- Employing use of **Multi-Intermediate Representation** to use a common interface for transpiling between numerous source and target platform combinations.
- **Optimisations** to fit the programs into the available **parser hardware** to successfully compile the programs and prevent failure.
- Integrating the optimisations and synthesis engine into **LLVM framework** to provide a seamless interface for end-users.

PES University Research Cell | Undergraduate Researcher | [GitHub](#)

Bengaluru, India

Automated Parallelization of Source Code using Program Comprehension

Aug 2020 – Dec 2021 (1 year 4 months)

- Explored techniques to convert any sequential program to its parallel equivalent, maximizing the CPU utilization, achieving a **500x speed-up in execution times** for larger inputs.
- Developed a **first-of-its-kind inter and intra-functional parallelization**, improving existing loop-based techniques.
- Achieved inter-functional parallelization by calculating **data and control dependencies** and developed a scheduling algorithm to perform parallel execution using **C++ threads**.
- Achieved intra-functional parallelization using techniques of **program comprehension** to identify and replace sequential algorithms with their suitable parallel equivalents.

Cloud Computing and Big Data Research Center | Summer Researcher | [GitHub](#)

Bengaluru, India

Shared Libraries on NUMA

May 2020 – Aug 2020 (3 months)

- Optimised the execution of Shared Libraries on NUMA machines by reducing the time wasted on internode communication.
- Replicated Shared Libraries executables across different nodes to eliminate latency caused by internode traffic and **improved performance by over 85%**.
- Worked with **QEMU** to emulate NUMA architecture to conduct experiments and generate results.

Microsoft Innovation Lab, PES University | Summer Researcher | [GitHub](#)

Bengaluru, India

Customer Feedback Analysis

Jun 2019 – Aug 2019 (3 months)

- Implemented a Customer Feedback Service to help any company improve its sales and services by enabling them to understand the customer views and expectations.
- Built a **sentiment analysis model using IBM Watson** to understand the severity of complaints posted on online review sites and an **AI-based ranking algorithm** to sort them by importance to highlight the most burning issues.
- Poster presentation at the IEEE CONNECCT 2019 held at IIT Bangalore.

TEACHING EXPERIENCE

Course Assistant | Prof. Mohamed Zahran, NYU Courant

Jan 2024 – May 2024 (5 months)

- Served as a Course Assistant for the course **CSCI-UA.0202-002 Operating Systems**.
- My duties included creating study material, helping resolve students' doubts about the lectures and class in general, actively monitoring and answering questions posted on an online portal for students, checking for plagiarism in lab assignments and handling of cases of plagiarism, grading midterm and final course examinations.

Grader | Prof. Mohamed Zahran, NYU Courant

Sept 2024 – Dec 2024 (4 months)

- Served as a Grader for the course **CSCI-UA.0480-051 Parallel Computing**.
- My duties included grading all homework, assignments, midterm and final course examinations and monitoring and answering questions posted on an online portal for students.

Teaching Assistant | Prof. Preet Kanwal, PES University

Jan 2022 – May 2022 (5 months)

- Served as a Teaching Assistant for the course **Compiler Design**.
- My duties included creating course slides and documentation, creating new problem statements for assignments and grading the same.

Subject Matter Expert (SME) | PES University

Jan 2020 – May 2020 (5 months)

- Served as a SME (part of PESU I/O) for the class of **NLP and Sentiment Analysis using IBM Watson**.
- My duties included mentoring a class of 30 students by myself, preparing digital content (videos and class slides), weekly quizzes and assignments, grading assignments and final project submissions, and enabling students to create their own sentiment analysis tools for different use cases.

AWARDS & ACHIEVEMENTS

CNR Rao Merit Scholarship: Awarded to Bachelor alumni placed in top 10% at PES University

Karnataka Regional Math Olympiad: Top 25 in the state and qualified for the Indian National Math Olympiad

Google Hash Code Coding Competition: India Rank: 128, Global Rank: 847 out of over 10000 participants worldwide

Best Project Award: Awarded for the best hackathon project on milestoning log data at Goldman Sachs among over 250 teams

Benefit Everyone Award: Awarded for the best Intern presentation on Unit Testing and Cmocka at Cisco among over 150 interns

HashCode Annual Hackathon: Mentored and collaborated with team working on "Customer Feedback Model" and placed in top 3

Sutherlands Innovathon - Finalist: Top 20 for Voice based Grievance System at Sutherlands Innovathon

Bhopal Smart City Hackathon - Finalist: Top 25 for Voice-based Grievance System at Bhopal Smart City Hackathon

Analytical Model for GPU Kernel Performance Prediction | [GitHub](#)

- Developed an innovative analytical model to predict kernel performance on GPUs. Despite the model's simplicity, it accurately predicts execution times for multiple GPU programs.
- Implemented several kernel functions, including matrix-vector multiplication, dot product, vector and matrix addition. Both linear and non-linear execution time functions are supported, demonstrating the model's versatility and robustness.
- The model's predictions were highly accurate, with the **percentage deviation** of predicted execution times from actual execution times falling within **6.9%**. This high level of precision underscores the model's effectiveness in real-world applications.
- Utilized **CUDA for GPU programming**, time profiling for performance measurement, and CUDA optimization analysis to enhance model accuracy. Key influencing parameters were analyzed, demonstrating expertise in GPU architecture and programming.

Future Video Frame Prediction | [GitHub](#)

- Utilized the Masked Conditional Video Diffusion (**MCVD**) model to predict future video frames with a **Jaccard index score improvement from 0.12 to 0.36** with increased sampling, demonstrating enhanced prediction accuracy.
- Employed the **U-Net** architecture for precise semantic segmentation of predicted video frames, achieving over 96% accuracy on the validation set.
- Combined MCVD for video frame prediction and U-Net for segmentation to generate and analyze high-quality video frames, showcasing significant strides in predictive capabilities and detailed scene analysis.
- Identified and worked on model limitations, such as image blurring and inference time, highlighting areas for future improvement to enhance the robustness and efficiency of video frame prediction models.

Generic Programming in C | [GitHub](#)

- A C project which implements the following Generic Programming Features in C using Pre-Processor Directives.
- **Generic Containers** - list, stack, queue, vector and hashmap that can be initialised with any type.
- **Iterators** to decouple the containers and algorithms and generic algorithms like find, count, min, max, accumulate etc, that make use of these iterators.

Implementation of Treaps | [GitHub](#)

- A C++ project which implements treaps as a data structure and added functionalities to support various operations on treaps.
- The treap is implemented with nodes built using **canonical classes**, with multiple operations supported such as insert, delete, split treap, merge two treaps, union, intersection, difference and traversal of treaps.
- A **bidirectional iterator** along with member algorithms such as find and replace are also implemented.

YACS - Yet Another Centralized Scheduler | [GitHub](#)

- A centralised scheduling framework maintaining a master node and multiple worker nodes, all simulated as different processes on a single machine using Python programming.
- 3 different scheduling algorithms implemented - **Least loaded, Round robin and Randomised.**

Image Segmentation using MSTs | [GitHub](#)

- Used graph-based representations of images to develop an efficient segmentation algorithm.
- Used Minimum Spanning Trees (MSTs) and a greedy algorithm to classify images into separate regions.
- Achieved a **linear time complexity** in the number of graph edges and is much faster compared to ML-based techniques.

Mini-compiler for C++ | [GitHub](#)

- Created a miniature version of a C++ compiler, covering basic syntax through lexical analysis, syntax analysis, intermediate code generation and code optimisation.
- Specific constructs implemented include the if, if-else and for constructs, along with the basic syntax of the language.

Voice-based Citizen Grievance System | *IBM Watson NLU*

- Collect citizen complaints, categorise and rank them based on severity using the IBM Watson NLU to help government officials tackle and respond to grievances faster and more efficiently.
- Another application of summer project done as part of the Microsoft Innovation Lab internship.