

EDUCATION

- **Cornell University** Ithaca, NY
Ph. D. in Computer Science - GPA: 4.18 January 2021 - Present
- **Indian Institute of Technology (BHU) Varanasi** Varanasi, India
B. Tech. in Computer Science and Engineering - GPA: 9.75/10.00 - Department Rank 1 July 2014 - May 2018

RESEARCH EXPERIENCE

- **Cornell University** Ithaca, NY
PhD Student January 2022 - Present
 - Advised by Dr. Robbert van Renesse and Dr. Lorenzo Alvisi. We are preparing a draft for submission to SOSP '23.
 - **Coordination-Free Distributed Shared Log** Distributed Shared Log is a popular abstraction to build consistent, fault-tolerant distributed systems. We propose a coordination-free protocol that combines batching and pre-ordering to provide both high-throughput and low-latency. I contributed to the high-level protocol details and I am also building a high-performance RDMA implementation of the system.
- **Cornell University** Ithaca, NY
PhD Student January 2021 - June 2021
 - Advised by Dr. Rachit Agarwal. I gave a presentation on this work at NetDev 0x15. This work was also published in SIGCOMM '21.
 - **Understanding Host Network Stack Overheads** With stagnant CPU speeds and exponentially faster network bandwidth, existing network stacks fail to fully utilise the potential of modern hardware. In this project, we analysed the Linux TCP stack to uncover bottlenecks and quantify the impact of various software and hardware offloads. Our insight lead to various recommendations relevant to the design of networking systems of the future. I designed and ran several experiments to measure the impact of NIC Rx descriptor size and TCP Rx buffer size on the end-to-end throughput and it's relation to the receiver cache miss rate in the presence of DDIO and it's relation to data-copy latency. This formed the basis of an important contribution in the paper.
 - **Artifact** I was the primary contributor in preparing the final artifact for the publication, which met the highest standards of the conference.
- **Microsoft Research** Bangalore, India
Research Fellow July 2018 - December 2020
 - Advised by Dr. Muthian Sivathanu and Dr. Ramchandran Ramjee. This research led to a publication in EuroSys '20 and later was also deployed in Microsoft's Singularity.
 - **Gandiva_{fair}** is a Proportionally Fair Scheduling policy for Gandiva based on the concept of using job migration to balance the ticket-adjusted load across servers and then performing intra-server proportional share to achieve cluster-level proportional share. It transparently handles GPU heterogeneity using profiling and a novel GPU trading strategy to improve cluster efficiency while preserving fairness. I was involved in the design of the scheduling algorithm and policies and their implementation.
 - **Gandiva** is a Cluster Scheduler/Resource Manager for Deep Learning Training jobs utilizing CPU Scheduling primitives like time-slicing, migration, and profiling to schedule GPUs efficiently. I redesigned and implemented the entire cluster manager using **Kubernetes** to manage job containers and wrote the scheduling framework in **Scala** using the **Akka Actors** library for concurrency and **gRPC** for performing RPCs.
 - **GPU Proxy** is a mechanism inspired from **CRCUDA** to enable suspend/resume, and migration of Deep Learning jobs written in **PyTorch** and **TensorFlow**. I implemented **Horovod** and **NCCL** support to enable suspend/resume and migration of distributed jobs.

PUBLICATIONS

- Q. CAI, S. CHAUDHARY, M. VUPPALAPATI, J. HWANG, and R. AGARWAL. *Understanding Host Network Stack Overheads*. 2021. In Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM '21). <https://dl.acm.org/doi/abs/10.1145/3452296.3472888>
- S. CHAUDHARY, R. RAMJEE, M. SIVATHANU, N. KWATRA, and S. VISWANATHA. *Balancing Efficiency and Fairness in Heterogeneous GPU Clusters for Deep Learning*. 2020. In Proceedings of the Fifteenth European Conference on Computer Systems (EuroSys '20). <https://dl.acm.org/doi/abs/10.1145/3342195.3387555>

INTERNSHIP

- **Goldman Sachs** Bangalore, India
Summer Analyst *May 2017 - July 2017*
 - **DbBench** is a distributed Database Benchmarking system for SecDB, a proprietary object database, used to make critical decisions about the future of the database system. It can emulate the real-life workload of hundreds of clients accessing a single database endpoint and report aggregate metrics like throughput and transactions per second. I designed and implemented the entire system in Scala using Akka Actors library for concurrency and Netty for composable high-throughput, low-latency asynchronous socket IO. I also implemented a YAML based configuration system and an interpreted Turing Complete programming language to write performance tests in.

SELECTED PROJECTS

- **Profiling Linux Network Stack** <https://github.com/Terabit-Ethernet/terabit-network-stack-profiling>
This repository contains the artifacts used to perform the experiments in our SIGCOMM '21 publication.
- **LTDP for Viterbi Algorithm** <https://github.com/5hubh4m/ltdp-viterbi-algorithm> Applied the Linear Tropical Dynamic Programming model to the Viterbi algorithm to parallelise the otherwise linear algorithm using the theorem of Rank Convergence. Implemented it using OpenMPI in conjunction with OpenMP in C++. Investigated the improvement over vanilla LTDP by redistributing data to processors after each principal iteration to prevent otherwise redundant computation of converged rows.
- **Scheme Interpreter in Haskell** <https://github.com/5hubh4m/simple-scheme> Wrote an Interpreter for a subset of Lisp (the R5RS standard) in Haskell, a pure functional programming language. Implemented an REPL, a parser, variables, user-defined functions using Functional Programming concepts like maps, folds, composition, monads, do-notation, currying, pattern-matching, etc.

TEACHING ASSISTANCE

- **Cornell University** CS 6410: Advanced Systems (Fall 2022), CS 4120: Introduction to Compilers (Spring 2022), CS 4410: Operating Systems (Fall 2021), CS 4450: Introduction to Computer Networks (Spring 2021).
- **IIT (BHU) Varanasi** CSE 241N: Artificial Intelligence (Fall 2018), CSO 101: Introduction to Computer Programming (for 4 semesters).

ACHIEVEMENTS

- **Outstanding TA Award** for the course CS 4450 (Introduction to Computer Networks) taught during Spring 2021 at Cornell University.
- **IIT (BHU) Varanasi Medal** for Department Rank 1 in Computer Science and Engineering B. Tech. Class of 2014-18.
- **Summer Research Fellowship** by the Indian Academy of Sciences, Bangalore in 2016.
- **Runner Up** at Microsoft's Code.Fun.Do 2016 24-Hrs Hackathon.
- **Winner** of Parliamentary Debate Competition at AMU Literary Festival, AMU Aligarh, 2015.