Joel Tony

≤ joeltony72@gmail.com

Education Birla Institute of Technology and Science Pilani, K.K. Birla Goa 2021 – 2025 | Goa, India **Campus,** B.E. Computer Science CGPA: 8.99/10 - Recipient of Institute Scholarship- 80% tuition fee waiver Relevant coursework: Distributed Computing (graduate-level), Data Storage Technologies & Networks, Computer Networks, Computer Architecture, Operating Systems, Principles of Programming Languages, Database Systems, Compiler Construction **Publications** Does Varying BeeGFS Configuration Affect the I/O Performance 2023 of HPC Workloads?, IEEE CLUSTER, 2023 ☑ Arnav Borkar[†], Joel Tony[†], Hari Vamsi K.N, Tushar Barman, Yash Bhisikar, Sreenath T.M, Arnab K. Paul **Research** Experience Scientific Computing & Imaging Institute, University of Utah, Jan 2025 – present | Salt Lake City, USA **Research** Assistant - Designing benchmarks comparing containerized vs. bare-metal HPC environments for DataSpaces 2, a distributed data framework for large-scale scientific workflows, under Prof. Manish Parashar 27. SCI Institute, University of Utah, Undergraduate Researcher Ian 2024 – Dec 2024 - Collaborating with Prof. Manish Parashar 🛛 to design a **POSIX-compliant API** for **DataSpaces** 🗹 , enhancing interoperability for large-scale scientific workflows. Data, Systems & High Performance Computing (DaSH) Lab, Oct 2022 – present | Goa, India Undergraduate Researcher - Led hybrid GPU/edge cluster deployment: Architected nodes and automated OS/node provisioning via Ansible. - Enabled secure remote access (Tailscale/Mosh) for 40+ researchers while automating cluster maintenance. - Built monitoring (Telegraf/InfluxDB/Grafana) and distributed storage (TrueNAS Scale) infrastructure. Lawrence Berkeley National Laboratory (Google Summer of May 2024 – Nov 2024 Code), Contributor - Collaborated with Dr. Jean Luca Bez 🛛 and Prof. Suren Byna 🖾 (Ohio State University) to extend the Drishti 🖾 framework, enabling I/O bottleneck analysis for non-MPI applications like distributed AI/ML workloads. - Modularized a monolithic 1,700-line codebase into Python components to enable scalable multi-trace aggregation. Embedded Systems Lab, Undergraduate Researcher Aug 2023 – Dec 2023 | Goa, India - Collaborated with Prof. K.R. Anupama I to design a **network protocol stack** for distributed cloudlet clusters, leveraging ns-2 simulations to optimize edge computing performance. **Central Electronics Engineering Research Institute** Jul 2022 – Aug 2022 | Pilani, India (CSIR-CEERI), Undergraduate Researcher - Worked under Dr. Dhiraj Sangwan (Principal Scientist) 🗹 to parallelize legacy C++/OpenCV workflows via Python migration, improving efficiency for dual-energy X-ray material classification. **Professional Experience** e6data, Performance and Research Engineering Intern 🛛 Jun 2024 – Dec 2024 | Bangalore, India - Achieved 7.5x faster serdes for string data and 5.4x avg. gains for fixed-length types. - Prevented OOM-driven query failures by designing a spill-to-disk framework for memory-intensive queries. **Coditation Systems,** Software Engineering Intern May 2023 – Jul 2023 | Pune, India

- Engineered a GPU remote development suite for graphics applications, reducing build times by 40% via Bazel's incremental artifact-based builds and cloud caching to skip redundant compilation.

Mu Sigma Inc., Data Science Intern

Jul 2018 – Aug 2018 | Bangalore, India - Applied linear regression in R Studio to analyze sales trends for anonymized Fortune 500 beverage datasets.

*i*aytau.com

jay-tau

in in/joel-tony

L +91 81057 53322

BeeGFS, IOR, Distributed File System

Projects

Data Prefetching for Edge Deep Learning Workloads,

Dynamic Striping Configuration for HPC Storage,

Apache Kafka, PyTorch, Distributed Deep Learning

- Engineered Kafka-based prefetching to mitigate I/O bottlenecks for PyTorch training on edge devices.

- Engineered a BeeGFS I/O interceptor using Darshan logs to automate file-level striping based on access

- Tailored prefetching logic to edge resource constraints, reducing idle cycles in distributed training workflows.

TCP/IP Stack Using Rust & C, Computer Networks, Rust, C

- Integrated Rust with a C-based TCP/IP stack to explore conditional memory safety guarantees and performance
- Leveraged **perf** and the **Valgrind** suite to analyze the performance gap between the integrated Rust code and original code

Autograder for Batch Evaluation of C Programs,

Software Development, Python

- Scaled grading automation for 1,000+ students, with auto-generated test cases via IO stream capture for lab evaluation.
- Secured academic integrity via symmetric encryption to authenticate submissions and flag potential plagiarism.

Personal Home Server, Docker, Linux, Networking

- Hosted Docker services (Jellyfin, Bitwarden) on Unraid, automating storage pooling and updates.
- Secured with a reverse proxy and Cloudflare Tunnels for encrypted, remote access.

Awards

ACM SIGHPC SC24 Travel Grant, ACM SIGHPC

- 1 of 6 students selected internationally-and the sole undergraduate recipient-to attend ACM/IEEE Supercomputing 2024 in Atlanta, GA, USA.

AWS AWS AI/ML Scholarship, AWS, Intel Corporation & Udacity 🛛

One of the 1000 students selected internationally for the AWS AI/ML Scholarship in collaboration with Intel and Udacity.

Teaching Experience

BITS Pilani	
Undergraduate Teaching Assistant (Fall 2024) - Principles of Programming Languages (CS F301)	Aug 2024 – Dec 2024 Goa, India
Teaching Assistant (Spring 2024) - Advanced Computer Architecture (CS G524) - Graduate Teaching As - Operating Systems (CS F372) & Computer Architecture (CS F342)	Jan 2024 – May 2024 Goa, India sistant
Undergraduate Teaching Assistant (Fall 2023) - Object-Oriented Programming (CS F213)	Aug 2023 – Dec 2023 Goa, India
Undergraduate Teaching Assistant (Spring 2023) - Computer Programming (CS F111)	Jan 2023 – Jul 2023 Goa, India
Professional Service	
ICDCS 2025, ICS 2025, IPDPS 2025, HiPC 2024, Cluster 2024, CCGrid 2024, TPDS 2023, External Reviewer	present
Supercomputing (SC) - The International Conference for High Performance Computing, Networking, Storage, and Analysis, ACM SIGHPC Volunteer	Nov 2024 Atlanta, USA
HiPC - IEEE International Conference on High Performance Computing, Data, and Analytics, <i>Student Volunteer</i>	Dec 2023 India

Aug 2023 – Dec 2023

Aug 2023 – Dec 2023

Dec 2022 – Jun 2023

2020 - present

Nov 2024

Nov 2022