

TWINKLE JAIN

Boston, MA ◊ 857.707.8421 ◊ jain.t@northeastern.edu ◊ [linkedin.com/in/jaintwinkle](https://www.linkedin.com/in/jaintwinkle) ◊ www.jaintwinkle.com

EDUCATION

Northeastern University Boston, MA
Ph.D. in Computer Science, GPA: 3.7/4.0 2017 – Present
Thesis Advisor: Prof. Gene Cooperman
Thesis Title: Application-transparent strategies to optimize limited resources in HPC and Big Data

Jai Narain Vyas University (currently MBM University) Jodhpur, India
Master of Computer Applications, First Class with Distinction, equivalent GPA: 4.0/4.0 2012 – 2015

RESEARCH EXPERIENCE

Northeastern University Boston, MA
Graduate Research Assistant 2017 – Present

- Proposed a novel extensible split-process architecture to handle multiple lower halves; the software will facilitate fault-tolerance in a multi-node, multi-GPU (CUDA + MPI) HPC application via checkpoint-restart (C/R).
- Introduced a low-overhead split-process-based architecture to checkpoint CUDA applications; published a first-author conference paper named “CRAC: checkpoint-restart architecture for CUDA with streams and UVM” in SC’20.
- Designed a testing suite for MANA, an MPI-Agnostic Network Agnostic C/R software for MPI applications; presented a poster called “Checkpoint/Restart of MPI applications over the Cray GNI Network via Proxies” in MUG’18.
- Extended DMTCP, a transparent checkpoint-restart tool, to support pseudo-TTYs to save a ROS v1 application; presented the work as “DMTCP: Fixing the single point of failure of the ROS master” in ROSCon’17.

MemVerge, Inc. San Francisco, CA
Summer Engineer Intern (Remote) 2022

- Troubleshoot memory corruption-related issues in MANA to checkpoint-restart at least three scientific HPC applications contributing around 25% of the total machine hours at NERSC Supercomputing sites.

IBM TJ Watson Research Center Yorktown Heights, NY
Summer Research Intern (Remote) 2021

- Analyzed the existing resiliency support in Ray, a distributed execution framework; demonstrated a 5% improvement in execution time and averted crashes caused by memory overflow via configuration tuning in Ray.

Inria Nantes, France
Summer Research Visitor 2019 & 2020

- Evaluated and improved speculative execution’s implementation (to detect and handle stragglers) in Hadoop and Spark.
- Published findings across multiple papers: a short paper, “Stragglers’ Detection in Big Data Analytic Systems: The Impact of Heartbeat Arrival,” in CCGrid’22; a second-author paper, “On the (In)Accuracy of Stragglers Detection in Hadoop,” in FGCS’23, and a first-author paper, “Towards an effective Speculative Execution in Spark,” in VLDB’23.

Mentor Graphics Waltham, MA
Summer System Engineer Intern 2018

- Developed a C/R plugin to restore an optimized executable as a debug build; reduced the debugging time by 75%.
- Patented** the work as “Software checkpoint-restoration between distinctly compiled executables” in August 2022.

Stratus Technologies Maynard, MA
Summer Platform Engineer Intern 2017

- Assessed performance of COarse-grained LOfck-stepping (COLO) technique on QEMU for fault-tolerance in servers.

SELECTED ACADEMIC PROJECTS

Northeastern University

Compiler for a small Programming Language 2018

- Wrote a complete compiler in SML-NJ language for Andrew Appel’s Tiger programming language in a team of two.

Decrease Down-time in Live Process Migration 2017

- Decreased downtime by 80% in process migration from one host to another by prioritizing memory-pages send order.

Jai Narain Vyas University

Maze Traversing Robot 2015

- Built a wireless, camera-driven robot to find a path in a maze; Won Pixelate award in Asia’s Largest Technical Festival.

TECHNICAL KNOWLEDGE

Languages and APIs: C/C++ (CUDA/MPI/POSIX), Python (numpy/matplotlib/pandas).
Distributed Frameworks: SLURM, Hadoop, Spark, Yarn, HDFS, Kubernetes, OpenShift.