

Sai Vishwanath Venkatesh

saivishw@buffalo.edu · (+1) 716-275-9676

EDUCATION

UNIVERSITY AT BUFFALO

MS IN COMPUTER SCIENCE

December 2022 | Buffalo, NY

GPA: 3.42

SRM INSTITUTE OF SCIENCE AND TECHNOLOGY

BTECH IN COMPUTER SCIENCE & ENGINEERING

May 2020 | Chennai, India

Percentage : 84.2%

LINKS

 [Linkedin](#)

 [Github](#)

 [Google Scholar](#)

 [Website](#)

COURSEWORK

GRADUATE

Parallel & Distributed Programming

Intro to Parallel & Distributed Programming

Intro to Machine Learning

Distributed Systems

Algorithms for Modern Computing Systems

ML and Society

TECHNICAL SKILLS

Languages

Python • C/C++

Frameworks

OpenMP • CUDA • OpenMPI • HPX •

Tensorflow • Keras

Miscellaneous

Bash • SLURM • CMake • Apache Spark •

Docker • SQL

EXPERIENCE

SOLARILLION FOUNDATION | RESEARCHER

Jul 2018 - Feb 2021 | Chennai, India

- Lead a team to tackle otherwise computationally expensive crime detection for video by using multiple instance learning with early stopping for optical flow. Deployed and tested on a Raspberry Pi3 (ARM v8) performing 2.3x times faster than the benchmark servers (Intel Xeon) using just a third of the number of cores.
- Lead a team towards condensing one of the largest datasets (10 billion+ records) for malware classification in smartphones to 0.1% of its feature set whilst retaining threshold accuracy reported by competing methods through efficient data preprocessing and online learning.
- Developed a two-stage model to estimate the number of weeks a movie in a multiplex would successfully be in theatres based on past occupancy/crowd behaviour. The solution is currently in use by one of India's leading cinema exhibitors.

PROJECTS

SCOOL - SCALABLE COMMON OPTIMIZATION LIBRARY [GITLAB](#)

Jan 2022 - May 2022 | University at Buffalo, Buffalo, NY

- Collaborated with Dr. Jaroslaw Zola on an open-source project to design and develop a general purpose executor to solve search space optimization problems (NP-hard problems) based on the BSP(bulk synchronous parallel) paradigm.
- Achieved up to 10x relative speedup and strong scaling when tested on a 52-core machine against the state of the art (with Intel TBB) being able to tackle Bayesian network structured learning problems in a matter of minutes that would otherwise take days.

GAUSSIAN KERNEL DENSITY ESTIMATION ACCELERATION ON GPU

Nov 2021 - Dec 2021 | University at Buffalo, Buffalo, NY

- Implemented gaussian kernel density estimation in parallel for floating point numbers on a Tesla V100 using CUDA.
- Observed up to 220x relative speedup to CPU implementation using OMP.

LARGE SCALE DISTRIBUTED SORTING ON A HPC

Oct 2021 - Nov 2021 | University at Buffalo, Buffalo, NY

- Built a distributed sorting algorithm on a problem size of 64 GB distributed across a set of 256 processing elements (across 4 nodes) using Open MPI and SLURM.
- Yields up to 67% efficiency and displays weak scaling when tested on UB's High-Performance Computer at CCR.

ROOTING CONNECTED COMPONENTS AT SCALE

Oct 2021 - Nov 2021 | University at Buffalo, Buffalo, NY

- An Apache Spark implementation that finds the roots of connected components. My version provides strong scaling and good efficiency on up to 12 processors for a graph with more than 100 million edges.

FAULT TOLERANT DISTRIBUTED KEY-VALUE STORE

Jan 2022 - May 2022 | University at Buffalo, Buffalo, NY

- Constructed a fault-tolerant application(using Flask) based on the RAFT consensus protocol to maintain a distributed key-value store. App excelled on a docker network with up to 7 containers.