SUDHARSHAN SRINIVASAN

ssriniv2@cs.uoregon.edu \diamond Website

EDUCATION

University of Oregon, Eugene PhD in Computer Science Research assistantship supported by Professor Boyana Norris and Allen Malony	Sept 2019 - Aug 2025(exp)
University of Oregon, Eugene MS in Computer Science	August 2017 - July 2019
SRM Institute of Science and Technology, Chennai BTech in Computer Sceince and Engineering	July 2013 - May 2017

SELECTED PUBLICATIONS

- Sudharshan Srinivasan, A. Khanda, Et al. "A Distributed Algorithm for Identifying Strongly Connected Components on Incremental Graphs," IEEE 35th Int'l Symposium on Computer Architecture and High-Performance Computing (SBAC-PAD), 2023
- Samuel D. Pollard, Sudharshan Srinivasan, and Boyana Norris "A performance and recommendation system for parallel graph processing implementations" in The 10th ACM/SPEC International Conference on Performance Engineering Companion, Mumbai, India, April 2019. ACM

RESEARCH EXPERIENCE

Infrastructure for dynamic graph models and GNNs

- Developing parallel algorithms and distributed infrastructure for running large-scale dynamic graphs on HPC system. Also developed frameworks for linear algebra-based graph models.
- Designing model architecture for training distributed Graph Neural Networks(GNNs).

Graph-based reinforcement learning

• Research on model architecture for graph-based MARL models. Benchmarks on Multi-particle physics and Google Research Football MARL environments

Scaling up whole genome simulations

• Developed parallel algorithms and implementations in collaboration with MesserLab to scale up SLiM, an evolutionary simulation framework to target multi-core CPU clusters.

Selection of parallel graph processing tools

 $\bullet \ \ Implemented a framework for recommending application-specific parallel graph processing toolkits and simplifying necessary build procedures. More details at https://github.com/HPCL/easy-parallel-graph \\$

Areas of focus : Distributed systems, GNNs, MARL and ML systems, parallel performance analysis and tuning

WORK EXPERIENCE

Research intern at AMD

• Focused on programming models for graph accelerators. Provided design ideas and wrote prototypes for parts of the hardware abstraction layer on the HSA runtime.

Research intern at Lawrence Livermore National Lab(LLNL)

• Implementing routing protocols for YGM, an asynchronous communication library built on top of MPI in C++. Performed experiments for large-scale asynchronous graph algorithms on the Livermore Cluster(LC).

Research intern at Argonne National Lab(ANL)

• Worked on the design of a performance portable library that lets applications using AMD's HIP API run on devices supporting Intel's level zero specification. This is part of the Exascale Computing Project aimed at performance portability of HPC applications to the Aurora supercomputer at ANL

TECHNICAL SKILLS

Programming languages: C++(Primary), Python Parallel and GPU Programming: MPI, ROCm stack, CUDA stack, OpenMP Performance analysis: Nvprof, Intel VTune Amplifier, Intel Advisor, Caliper, TAU, PAPI, Graph libraries: cuGraph, Networkx, Gunrock, Gluon, BoostGraph, HavoqtGT GNN and RL libraries: PyG, DGL, MARLlib, OpenAI GYM, Ray RLlib Aug 2024' - Current

June 2019 - Current

June 2019 - June 2020

June 2018 - June 2019

Jan'22

June'21

June'20

AWARDS AND HONORS

Graduate Teaching Fellowship

- CS425 Operating Systems: Fall 25', Winter 25'
- CS333 Applied cryptography: Spring 25'
- CS325 Software Engineering: Fall 19', Winter 19'

Graduate Research Fellowship

- Supported by HPCL lab : Fall 21' to Fall 24'
- Supported by Ecology and Evolution dept: Fall 20' to Fall 21'

Research Grants

- NSF grant for dynamic graph research: Awarded(2021), renewed(2025)
- NIH grant for scaling up whole genome simulations: Awarded(2020)

Paper Peer Reviews

- International conference on HPC: SC 25' HPC for AI track
- International Symposium on Computer Architecture: SBAC-PAD 25' Scalable Algorithms track

Proceedings volunteer

• International conference on HPC: SC21' SC22', SC23'