SAMEERAN JOSHI

Joshisameeran17@gmail.com | https://sameeranjoshi.github.io | Salt Lake City, UT

EDUCATION

- ***B.E in Computer Science, P.V.G's COET, Pune,*** *India* *May 2015 - April 2019*
- ***Ph.D in Computer Science, University of Utah,*** *USA* *Aug 2022 - Enrolled*
  Advisor: Dr. Mary Hall
  Focus: Compilers and Programming Models for spatial dataflow hardwares, Hardware-Software Codesign, HPC

PUBLICATIONS

- *Scheduling Language Chronology: Past, Present, and Future*, M Hall, C Oancea, AC Elster, A Rasch, **S Joshi**, AM Tavakkoli, R Schulze, ACM TACO, 2025

- *PEAK: Generating High-Performance Schedules in mlir*, Amir Mohammad Tavakkoli∗, **Sameeran Joshi**∗, Shreya Singh, Yufan Xu, P. Sadayappan, and Mary Hall. In Proceedings of the 36th International Workshop on Languages and Compilers for Parallel Computing(LCPC23). Oct. 2023

- *An NSF REU Site Based on Trust and Reproducibility of Intelligent Computation: Experience Report*, Mary Hall, Ganesh Gopalakrishnan, Eric Eide, Johanna Cohoon, Jeff M. Phillips, Mu Zhang, Shireen Y. Elhabian, Aditya Bhaskara, Harvey Dam, Artem Yadrov, Tushar Kataria, Amir Mohammad Tavakkoli, **Sameeran Joshi**, Mokshagna Sai Teja Karanam. In **EduHPC workshop** at The International Conference for High Performance Computing, Networking, Storage, and Analysis **(SC23)**

- **WIP**: ***Mapping Sparse Linear Solver on Cerebras Wafer Scale Engine***: Explored mapping of multigrid scientific application onto the kernel language of Cerebras wafer scale engine. challenges like layout the problem efficiently on a million cores, dataflow between nodes, memory mapping. Compared with Nvidia gpu state of the art solver. Preliminary results show 10-100x speedup.

EXPERIENCE

- **AMD, Campbell, CA** *May 2025 - Aug 2025*
  Compiler Intern, AIE Accelerator(NPU)
  Manager: Mahesh Ravishankar

  Mentor: James Newling

  Worked on VLIW, tiled IR compiler based on IREE, and LLVM for NPU. Supported and vectorized truncation, reduction kernels used in Deep Learning. Achieved on par speedup and faster speedups for different multidimensional matrix shapes.

- **Argonne National Lab, Lemont, IL** May *2024 – Aug 2024*
  Research Intern Technical PhD
  Supervisor: Dr. Siddhisanket Raskar

  Codegen(challenges and opportunities) for mapping HPC applications on AI testbeds, explored various accelerators including Cerebras, Sambanova, Groq, GraphCore. Added support for GraphCore backend in DaCe (Data Centric Parallel Framework)(code).

- **AMD, Bangalore, India** *June 2019 - June 2022*
  CPU Compiler Engineer, Full Time

  Worked on various parts including verification(Unit testing, finding compiler bugs using automated tools), optimization opportunities(Built internal tool using LLVM Bolt project) and frontend support in Flang(Fortran, OpenMP).

- **GCC UG project** *June 2018 - April 2019*
  Remote
  Mentor: Andi Kleen
  Extended Csmith for GCC C-Language Extensions

  Added ~15 GNU C language extensions to Csmith and found unexplored bugs (ICE's, seg faults, crashes) in GCC compiler. Found 12 critical bugs, 11 were fixed by GCC community. Increased the fuzzing code coverage of Csmith on GCC by – line coverage: 5%, function coverage: 7%, branch coverage: 4%

PROFESSIONAL SERVICE

- Volunteering CppOnSea'21, CppCon'21, SC25
- Program Committee Member LLVM Dev Meet 2021
- Artifact Evaluation Committee CGO25, ASPLOS25
- Student Travel Grant for Workshop on Sparse Tensor Computations