

Nimit Singhania

Email: singhania.nimit@gmail.com

Ph: (215) 605-0424

LinkedIn: <https://www.linkedin.com/in/nimit-singhania>

Objective

To seek a full-time opportunity in the field of Compiler and Systems Infrastructure.

Education

- Ph.D., University of Pennsylvania (August 2013 - December 2018)
Dissertation: Static Analysis for GPU Program Performance
- B.Tech., Indian Institute of Technology, Kanpur, India (July 2006 - June 2010)

Professional Experience

- **Member of Technical Staff - I, eBay**, San Jose (April 2021 - Present)
 - Working on retrieval architecture for Search for the inventory at eBay.
 - Includes analysis for the current systems and augmenting the system to improve relevance of retrieved items.
- **Consultant, Facebook**, Sunnyvale (September 2020 - April 2021):
 - Developing software kernels in C for ML Inference Hardware designed at Facebook.
 - Across-the-stack involvement from simulation models to RTL validation and Emulation.
 - Cross-functional work with design, verification and compiler teams.
- **Software Engineer, Google**, Sunnyvale (October 2018 - July 2020):
 - Piloted a CPU server platform with a new memory technology and evaluated impact on job performance.
 - Built simulators to predict the impact of landing the server into fleet in C++.
 - Helped build and maintain benchmarks for performance analysis of new and existing server platforms.
- **Research Intern, Bell Labs**, Murray Hill (June 2015 - August 2015): Developed a research prototype in LLVM to allow programmable and formally verified loop transformations in sequential code.
- **Software Engineer, IBM Research**, New Delhi, India (July 2010 - June 2013):
 - Led the work on the runtime backend for a tool (Java) to automate manual scripts for web-testing into automatic tests.
 - Innovated the backend with change-resilient representation of web page objects that helped automatic tests to be robust.

- Worked on a research prototype (Java) to identify deep bugs in large programs using modularity of code to explore code paths close to the error location.
- **Intern, NVIDIA**, Pune, India (May - June 2008):
 - Scaled the GPU benchmark suite to evenly utilize cores across multiple generations of GPU with up to 50% performance improvement on GT200 GPU.

Research Projects

- Static analysis to identify uncoalesced accesses in GPU programs: relied on a novel abstraction of threads to scale analysis with up to 70% precision. Identified 111 uncoalesced accesses as compared to a dynamic analysis with only 69 accesses, with up to 100x faster execution.
- Block-size Independence of GPU programs: developed an approach to prove independence of GPU programs from block-size in the synchronization-free setting. The analysis is used to modify block-size and improve performance of uncoalesced accesses in GPU programs.
- Programmable and Formally Verified Loop Transformations: built a semi-automatic LLVM based approach to optimize loops in sequential programs using user-specified transformation. Showed up to 10x improvement as compared to state-of-the-art optimization compilers.

Positions of Leadership

- Volunteer at Homesteaders' 4H club to promote STEM among youth (2018 - 2020).
- Former member, Penn Sargam, a South Asian fusion music group at University of Pennsylvania (2014-2018).

Skills

- Areas: Program Analysis, Computer Architecture, System Design and Infrastructure, Software Engineering, Machine Learning.
- Frameworks: LLVM, MLIR, Unix tools (Vim, Bash, sed, awk, grep), Latex, Matlab, Github.
- Languages: C/C++, Java, Matlab, Shell, Make.

Patents

- P202418US01 / 079894.6930 – Auto-Configuration of Hardware Non-linear Function Acceleration (Filed).

Peer-reviewed Publications (selected)

- Static Analysis to detect Uncoalesced Accesses in GPU Programs, R. Alur, J. Devietti, O. Navarro Leija, N. Singhanian, FMSD 2021 (Journal Paper).
- Block-Size Independence of GPU Programs, R. Alur, J. Devietti, N. Singhanian, SAS 2018.
- Loopy: Programmable and Formally Verified Loop Transformations, K. Namjoshi, N. Singhanian, SAS 2016.
- Mosaic: Precise Piecewise Affine Models from Input-Output Data R. Alur, N. Singhanian, EMSOFT 2014.