Saurav Muralidharan

saurav_m@acm.org

RESEARCH

Neural Network Performance & Scalability, Parallel Programming Models, Compilers, GPU Computing

EDUCATION

Ph.D. (Computer Science), University of Utah, Salt Lake City, UT, USA **May 2016**

- **Advisor:** Mary Hall
- **Dissertation:** Abstractions and Strategies for Adaptive Programming

B.Tech. (Computer Science & Engineering), Kannur University, Kerala, India **June 2009**

EXPERIENCE

NVIDIA Research, Santa Clara, CA, USA **August 2016 – Present**
Research Scientist

School of Computing, University of Utah, Salt Lake City, UT, USA **August 2011 – June 2016**
Graduate Research Assistant

Designed and developed *Nitro*, a machine learning-based autotuning framework that automatically finds the best implementation (variant) of a computation for a given input and hardware architecture. It provides C++ and Python interfaces for programmers to specify variants, input dataset features, and constraints. Nitro also supports tuning across hardware architectures using multi-task learning. **[IPDPS '14] [ASPLOS '16]** [nitro-tuner.github.io]

NVIDIA Research, Santa Clara, CA, USA **May – August: 2013 & 2014**
Research Intern

Worked on *Surge*, a nested data-parallel programming model that decouples the specification of high-level computations from their hardware implementation using first-class constructs named *schedules* and *policies*. This enables the same Surge program to have multiple implementations which can then be tuned. **[PPoPP '15 Poster] [TACO '16]**

NVIDIA, Salt Lake City, UT, USA **May 2011 – August 2011**
Software Intern

Designed and implemented a PTX-to-LLVM IR translator for NVIDIA's OptiX ray tracing software.

Indian Institute of Technology (IIT) - Madras, Chennai, India **November 2009 – June 2010**
Project Associate

Explored static analysis-based approaches for CUDA code generation from sequential loop nests.

PUBLICATIONS

"Designing a Tunable Nested Data-Parallel Programming System"
S. Muralidharan, M. Garland, A. Sidelnik, M. Hall
ACM Transactions on Architecture and Code Optimization **(TACO)**, 2016.

"Architecture-Adaptive Code Variant Tuning"
S. Muralidharan, A. Roy, M. Hall, M. Garland, P. Rai
International Conference on Architectural Support for Programming Languages and Operating Systems **(ASPLOS)**, 2016.

"A Collection-Oriented Programming Model for Performance Portability"
S. Muralidharan, M. Garland, B. Catanzaro, A. Sidelnik, M. Hall
Symposium on Principles and Practice of Parallel Programming **(PPoPP)**, poster paper, 2015.

"Nitro: A Framework for Adaptive Code Variant Tuning"
S. Muralidharan, M. Shantharam, M. Hall, M. Garland, B. Catanzaro
IEEE International Parallel & Distributed Processing Symposium **(IPDPS)**, 2014.

"Towards Making Autotuning Mainstream"
P. Basu, M. Hall, M. Khan, S. Maindola, *S. Muralidharan*, S. Ramalingam, A. Rivera, M. Shantaram, A. Venkat
International Journal of High Performance Computing Applications **(IJHPCA)**, 2013.

"Galaxia: A Semi-Decentralized System for Implementing Secure-Group P2P Networks"
S. Muralidharan, S. Koroth, N. Anto, R. Pandarachalil
International Conference on Networks & Communications **(NetCoM)**, 2009.

POSTERS & TALKS

"Designing a Tunable Nested Data-Parallel Programming System", Invited Conference Talk, High Performance and Embedded Architecture and Compilation Conference **(HiPEAC '17)**, January 2017, Stockholm, Sweden

"Building High-Performance Input-Adaptive GPU Applications with Nitro", Talk, GPU Technology Conference **(GTC '15)**, March 2015, San Jose, USA

"A Framework for Input and Architecture Aware Code Variant Autotuning", Early Research Showcase Poster, The International Conference for High Performance Computing, Networking, Storage and Analysis **(SC '13)**, November 2013, Denver, USA

AWARDS

Student Travel Grant, ASPLOS 2016
Winner, University of Utah School of Computing poster competition, 2014
Academic excellence award for ranking 1st in the CS&E department, GCE Kannur, 2009

PROFESSIONAL SERVICE

Conference Program Committees: IPDPS 2018, ICS 2018
Conference Refereeing: SC 2016, PPoPP 2016, PLDI 2014, HPCC 2014, ICCS 2013
Journal Refereeing: TACO 2018, TOPC 2017

TECHNICAL SKILLS

Programming Languages: C, C++, Python
Parallel Programming: CUDA, OpenMP, MPI, Pthreads
Compiler Frameworks: LLVM, CLANG

TEACHING

School of Computing, University of Utah, Salt Lake City, UT **January 2013 – April 2013**
Graduate Teaching Assistant, "Parallel Programming for Many-Core Architectures"

GRADUATE COURSEWORK

Computer Architecture, OS, Compilers, Advanced Algorithms, Parallel Programming for GPUs, Advanced Embedded Software, Parallel Computing and HPC, Models of Computation for Massive Data

ACTIVITIES

Graduate Student Advisory Committee, 2012-2014