Harshitha Menon

CONTACT
INFORMATION

Website: http://harshithamenon.com
E-mail: harshitha.menon@gmail.com
Phone: 650-741-4260

RESEARCH
INTEREST

Parallel Computing, High Performance Computing, Distributed Systems

EDUCATION

University of Illinois at Urbana–Champaign, IL, USA
PhD Candidate in Computer Science **Aug 2012 – Present**
GPA: 3.91/4.0
Advisor: Prof Laxmikant Kale

University of Illinois at Urbana–Champaign, IL, USA
Master of Science in Computer Science **Aug 2010 – May 2012**
GPA: 3.91/4.0
Advisor: Prof Laxmikant Kale

College of Engineering, Trivandrum, Kerala, India
Bachelor of Technology in Computer Science **Jul 2002 – May 2006**
Percentage: 78.4 (Ranked 5th out of 66)

WORK
EXPERIENCE

Summer Intern at LLNL **May 2013 - Aug 2013**
Lawrence Livermore National Lab

Research Assistant at UIUC **Dec 2011 - Present**
Member of PPL under the guidance of Prof. Kale

Internship at Google - Mountain View **May 2011 - Aug 2011**
Software Engineer Intern; Member of Map Reduce team

Teaching Assistant at UIUC **Aug 2010 - Dec 2011**
Course: Data Structures

Google - Hyderabad **May 2006 - Aug 2010**

PUBLICATIONS

Harshitha Menon, Lukasz Wesolowski, Gengbin Zheng, Pritish Jetley, Laxmikant Kale, Thomas Quinn, Fabio Governato
Adaptive Techniques for Clustered N-Body Cosmological Simulations [Computational Astrophysics and Cosmology 2015]

A Bastidas Fry, F Governato, A Pontzen, T Quinn, M Tremmel, L Anderson, H Menon, AM Brooks, J Wadsley
All about baryons: revisiting SIDM predictions at small halo masses [Monthly Notices of the Royal Astronomical Society 2015]

Bilge Acun, Abhishek Gupta, Nikhil Jain, Akhil Langer, Harshitha Menon, Eric Mikida, Xiang Ni, Michael Robson, Yanhua Sun, Ehsan Totoni, Lukasz Wesolowski and Laxmikant Kale *(All authors have equal contribution)*
Parallel Programming with Migratable Objects: Charm++ in Practice [SC14]

Jonathan Lifflander, Esteban Meneses, Harshitha Menon, Phil Miller, Sriram Krishnamoorthy, Laxmikant Kale *Scalable Replay with Partial-Order Dependencies for Message-Logging Fault Tolerance* [Cluster 2014]
Best Student Paper

Harshitha Menon, Laxmikant V. Kale
A Distributed Dynamic Load Balancer for Iterative Applications [SC 2013]
Best Student Paper finalist

Harshitha Menon, Bilge Acun, Simon Garcia De Gonzalo, Osman Sarood and Laxmikant Kale
Thermal Aware Automated Load Balancing for HPC Applications [Cluster 2013]

Harshitha Menon, Nikhil Jain, Gengbin Zheng and Laxmikant V. Kale
Automated Load Balancing Invocation based on Application Characteristics [Cluster 2012]

Bilge Acun, Akhil Langer, Esteban Meneses, Harshitha Menon, Osman Sarood, Ehsan Totoni, Laxmikant V. Kale
Power, Reliability, Performance: One System to Rule Them All [Under submission for IEEE Computer 2016]

Harshitha Menon, Abhinav Bhatele, Sebastien Fourestier, Franois Pellegrini and Laxmikant Kale
Applying Graph Partitioning Methods in Measurement-based Dynamic Load Balancing [Technical report]

AWARDS AND HONORS

- **Recipient of ACM/IEEE-CS George Michael Memorial HPC Fellowship** - 2014.
- **Best Paper Award** at Cluster 2014
- **Best Student Paper finalist** at SC 2013
- **Received Google Anita Borg Scholarship** - 2014.
- **Anita Borg finalist** - 2013.
- Awarded **Siebel Scholarship** Class of 2012 - The Siebel Scholars Foundation rewards excellence among the top students at the leading graduate schools.
- **Best Summer Intern Poster** at LLNL - 2013
- Member of the PPL team, the **finalist for HPC challenge** class 2 award 2012.
- Awarded **Google Fellowship for Employees**, 2010 - Competitive fellowship granted to Google employees to pursue higher education based on merit.
- **Teaching Excellence Award**, UIUC, fall 2011.
- Awarded **Founders Award** at Google in 2007 for contribution to Gmail.
- **Best Graduating Student Award**, College of Engineering Trivandrum, 2006.

PROJECTS & RESEARCH

Epidemic Algorithm for Distributed Load Balancing [SC 2013]
- A distributed algorithm for dynamic load balancing which uses gossip protocol for information propagation.
- Scales to $131K$ cores of BlueGene/Q.
- New feature in Charm++ production release.

Adaptive Techniques for Clustered N-Body Cosmological Simulations [Computational Astrophysics and Cosmology 2015]
- Addressed many challenges in scaling ChaNGa, n-body simulation code written in Charm++ which enabled it to scale to $512K$ cores of Blue Waters.
- Implemented scalable load balancing strategy and adaptive techniques to improve performance.

Automated Load Balancing based on Application Characteristics [Cluster 2012]
Master thesis with Prof.Kale
- An automated framework which takes decisions regarding load balancing for an application.
- The framework instruments multiple characteristics of an application like computation time, communication pattern; and selects an ideal load balancing strategy to be used for it.
- Decisions related to how frequently and when load redistribution needs to be done are taken by the framework; best decisions can be made with minimal efforts from end users.
- Feature in Charm++ production release.

Thermal Aware Automated Load Balancing [Cluster 2013]
- An adaptive control system that minimizes the cooling energy by using Dynamic Voltage and Frequency Scaling to control the temperature.
- The framework also detects load imbalance created by DVFS and performs automatic load balancing.

Chizu: A Framework to Enable Topology Aware Task Mapping

Summer intern project at Lawrence Livermore National Lab
- Topology aware mapping can improve the performance of a parallel application significantly.
- Developed Chizu, is a generic framework, which facilitates topology aware mapping of tasks.
- Used Scotch graph partitioner to perform mapping.
- This project is one of the LLNL's CASC Division OKRs.

Automated Distributed Configuration Manager for MapReduce
Summer intern project at Google with Dr Ken Goldman
- Implemented a distributed configuration manager for mapreduce worker formerly done manually by hit and miss trials.
- The manager adaptively configures the workers based on information shared between the workers using gossip protocol to obtain optimal performance.

Load Balancing using Graph Partitioning
- Implemented a dynamic load balancing strategy using Scotch graph partitioner in Charm++.
- Obtained performance gains of the order of 15–20% over existing algorithms.
- Suggested constructive changes which were incorporated in Scotch in the form of new partitioning strategies for the cases in which weight variance was large among vertices.

Job Execution Management system
Work project at Google
- Redesigned the existing job management system for scheduling, distributing and managing test binaries across multiple platforms.
- Enabled scalability, replication across data centers, dynamic load management and recovery on failure.
- Implemented an automated solution to monitor the health of the service.
- Used bigtable, GFS and other Google technologies for underlying infrastructure.

PROFESSIONAL & COMMUNITY
- Reviewing: TPDS, Concurrency and Computation: Practice and Experience, CiSE
- Panelist in HPC for undergraduate program at SC-13.
- Held Charm++ tutorial along with two other students at the 2014 Charm Workshop
- Guest speaker at RIT Engineering College. Gave a talk on Android platform.
- Held Robotics workshop at Google for Google Women in Engineering Award finalists.
- Mentor for nooglers; Test certified mentor - helps improve the testing practices at Google.
- Presented multiple sessions on robotics and parallel computing for high school students.