

Janki Bhimani, PhD

Computer Engineering

Tel: 857-991-9868

Email: bhimani@ece.neu.edu

LinkedIn: www.linkedin.com/in/jankibhimani

Webpage: <http://nucsr.coe.neu.edu/janki>



RESEARCH INTERESTS

System Performance Engineering; Flash-Based Storage Enhancement; Big Data Processing; Virtualization; Docker Container Scheduling; Datacenter Endurance and Reliability; Parallel Full-Stack Processing; High Performance Computing; Performance Modeling and Prediction; Capacity Planning; Resource Management; I/O Workload Characterization.

HIGHLIGHTS

Janki Bhimani is PhD from department of Electrical and Computer Engineering at Northeastern University, Boston. She is the recipient of best paper awards at the IEEE CLOUD in 2018 and IPCCC in 2017. Currently, she has nine publications in highly selective conferences along with ten other conference publications. She has five journal publications with two of it in IEEE Transactions. She is also main inventor of four top graded patents. She has served Northeastern University as an instructor, teaching a 4-credit course on fundamentals of engineering algorithms to undergraduate students. She received excellent feedback from her class with 4.4/5 as her instructor effectiveness mean. She has worked closely with Samsung research labs during her PhD.

EDUCATION

Doctor of Philosophy (PhD), Computer Engineering Northeastern University, Boston, MA Advisor: Ningfang Mi <i>Dissertation:</i>	<i>GPA: 3.9/4.0</i>	<i>Expected June 2019</i>
Master of Science (MS), Computer Engineering Northeastern University, Boston, MA Advisor: Ningfang Mi and Miriam Leeser <i>MS research: FiM - Fine grained Model to Predict Heterogeneous Computing Platforms Performance</i>	<i>GPA: 3.6/4.0</i>	<i>Aug. 2014</i>
Bachelor of Technology (BTech), Electrical & Electronics Engineering GITAM University, Vishakhapatnam, India <i>Major: Robotics and Programming of Embedded Systems, Minors: Circuit Design, Power Management</i>	<i>GPA: 9.26/10</i>	<i>May 2013</i>

AWARDS, HONORS, AND FELLOWSHIPS

2018	The Best Paper Award at 11th IEEE International Conference on Cloud Computing (IEEE CLOUD)
2017	The Best Paper Award at 36th IEEE International Performance Computing and Communications Conference (IPCCC)

- 2014** Double Husky Scholarship, Northeastern University
- 2012** The Best Budget Robot Award at 3rd Lunabotics International Mining Competition, NASA, FL for developing the most innovative and cost efficient Lunar Rover to operate on Synchronized Regolith at NASA Kennedy space center
- 2012** The Best Working Model Award in Junk Yard Wars at Conscientia, Indian Institute of Space Science Technology (IIST), for building a self-on-board drive machine in 1 day with limited resources of junk found in surroundings
- 2011** The Outstanding Debate Performance Award by Institute of Engineers India (IEI)
- 2010** The Impromptu Speaker Award by International Society for Technology in Education (ISTE)
- 2010 - 2013** University Merit Scholarship, GITAM University

PROFESSIONAL EXPERIENCE

Software Development Infrastructure Engineer

Samsung Semiconductors, Inc. Research Lab, San Jose, CA *Manager: Vijay Balakrishnan* *May 2018 – Aug. 2018*

- Project: KV-Kmeans
- Develop HPC Key Value API that translates file based machine learning applications to key and value based applications
- Simplify application data management by removing filesystem and offloading data storage from block based SSDs to key-value based SSDs
- Integrate OpenMP pragma with KV-kdd protocols to implement hybrid key-value based multi-threaded unsupervised clustering application with a KV SSD plug-in as primary I/O path

Performance Engineer

Samsung Semiconductors, Inc. Research Lab, San Jose, CA *Manager: Vijay Balakrishnan* *May 2017 – Aug. 2017*

- Project: PatIO – Pattern I/O Generator
- Design and develop a benchmarking framework to advance a naive synthetic I/O engine capable to produce I/Os that represent real-world workloads
- Deploy I/O generator developed to enhance and fasten the R&D phase of Autostream project for variety of workloads using Multi-stream SSDs

Engineer - Performance Architect

Samsung Semiconductors, Inc. Research Lab, San Jose, CA *Manager: Vijay Balakrishnan* *May 2016 – Aug. 2016*

- Project: Automatic Stream Assignment for Containerized Workloads
- Explore performance of an array of NVMe drives when exposed to containerized multiple parallel applications
- Innovate, design and implement a smart multi-stream assignment analytical model to estimate the improvement in write amplification factor of SSD

Graduate Research Assistant

Northeastern University, Boston, MA *Advisor: Ningfang Mi* *May 2014 – Present*

- Investigate interactions and implications of executing data intensive workloads on emerging hardware systems to achieve better performance, endurance and reliability.
- Design analytical models, prediction tools, and devised novel techniques to provide resource management for system computation, communication, memory and storage.
- Propose new ways to enhance the performance of big-data processing on large-scale enterprise cloud infrastructures existing techniques of modelling in analytical, simulation and emulation for performance prediction.
- Explore evolving flash technologies by demonstrating deployment to mitigate its proposed challenges and innovating techniques to provision endurance and reliability

Research Associate

NASA, John F. Kennedy Space Center, FL *Manager: Philip T. Metzger* *Jul. 2012 – Aug. 2012*

- Design and development responsibility towards a STEM project to build a lunar excavator which was tested on synchronized lunar regolith
- Designed an automatized actuator control system using driver IC LM293D and relay switch for voltage control
- Built a remote-controlled robot operating with a communication channel of wifi handshake and MAX232 to convert the TTL/CMOS to RS232 logic levels during serial communication of microcontroller with PC

IC Design Intern

Energy Options, Rajkot, India

Manager: Bhavesh Bhimani

Jun. 2012 – Jul. 2012

- Researched on energy efficiency and performance improvement of ‘Solar Power Harvest’
- Invented the smallest and most energy efficient IC for line follower robot, which was then used in labs of 48 regional colleges to teach students about efficient IC design and line follower robot operation
- Programmed and debugged embedded IC

TEACHING EXPERIENCE

Instructor

Northeastern University, Boston, MA

Course: EECE 2560: Fundamentals of Engineering Algorithms

Sept. 2017 – Dec. 2017

- Teach a 4-credit undergraduate level course
- Formulated coursework, assignments, projects, quizzes and exams
- Covers the design and implementation of algorithms to solve engineering problems using a high-level programming language. Reviews elementary data structures, such as arrays, stacks, queues, and lists, and introduces more advanced structures, such as trees and graphs and the use of recursion. Covers both the algorithms to manipulate these data structures as well as their use in problem solving. Emphasizes the importance of software engineering principles. Introduces algorithm complexity analysis and its application to developing efficient algorithms.

Graduate Teaching Assistant

Northeastern University, Boston, MA

Courses:

EECE 5640 High-Performance Computing

Supervisor: Mirriam Leeser

Jan. 2016 – Apr. 2016

EECE 7205 Fundamentals of Computer Engineering

Supervisor: Stefano Basagni

Sep. 2015 – Dec. 2015

EECE 5698 ST: Simulation & Performance Evaluation

Supervisor: Ningfang Mi

Jan. 2015 – Apr. 2015

CS5800 Algorithms

Supervisor: Kenneth Baclawski

Sep. 2014 – Dec. 2014

Duties:

- Conduct lectures to demonstrate practical operation of computing cluster and teach coding, debugging and profiling using recent tools corresponding to various hardware
- Mentor and evaluated students through their academic projects, assignments and class doubts
- Prepare handouts for each lecture for students to gain in-depth knowledge on topic of discussion

Student Chair

GITAM University Student Activity Center (GUSAC), India

Supervisor: Dharma Raj Cheruku

May 2011 – May 2013

- Lecturer for “Robotics Programming and Circuit”
 - Lead Speaker for workshops on “Basics and Advanced Programming”, “MATLAB”, “Algorithm and Logic Development” and “Simulation and Modeling”
 - Chief Organizer of Annual Techfest “GUSAC Carnival” for years 2012 and 2013
-

Refereed Journal Publications

1. Zhengyu Yang, Manu Awasthi, Mrinmoy Ghosh, Janki Bhimani, and Ningfang Mi, "**I/O Workload Management for All-Flash Datacenter Storage Systems Based on Total Cost of Ownership**", *IEEE Transactions on Big Data (TBDSI)*, Special Issue on the Integration of Extreme Scale Computing and Big Data Management and Analytics, 2018. DOI 10.1109/TBDATA.2018.2871114. (To Appear)
2. Janki Bhimani, Zhengyu Yang, Ningfang Mi, Jingpei Yang, Qiumin Xu, Manu Awasthi, Rajinikanth Pandurangan, and Vijay Balakrishnan, "**Docker Container Scheduler for I/O Intensive Applications running on NVMe SSDs**", *IEEE Transactions on Multi-Scale Computing Systems (TMSCS)*, 2018. (To Appear)
3. Zhengyu Yang, Janki Bhimani, Yi Yao, Cho-Hsien Lin, Jiayin Wang, Ningfang Mi, and Bo Sheng, "**AutoAdmin: Admission Control in YARN Clusters Based on Dynamic Resource Reservation**", *Scalable Computing: Practice and Experience, Special Issue on Advances in Emerging Wireless Communications and Networking (SCPE)*, 2018. Volume 19, Number 1, pp. 53–67.
4. Zhengyu Yang, Yufeng Wang, Janki Bhimani, Chiu C. Tan, and Ningfang Mi, "**EAD: Elasticity Aware Deduplication Manager for Datacenters with Multi-tier Storage Systems**", *Cluster Computing (CC)*, 2018. <https://doi.org/10.1007/s10586-018-2141-z>.
5. Zhengyu Yang, Janki Bhimani, Jiayin Wang, David Evans, and Ningfang Mi, "**Automatic and Scalable Data Replication Manager in Distributed Computation and Storage Infrastructure of Cyber-Physical Systems**", *Scalable Computing: Practice and Experience, Special Issue on Communication, Computing, and Networking in Cyber-Physical Systems (SCPE)*, 2018. Volume 18, Number 4, pp. 291–311.

Highly Selective Conference Publications

Acceptance rates $\leq 30\%$

6. Janki Bhimani, Ningfang Mi, Zhengyu Yang, Jingpei Yang, Rajinikanth Pandurangan, Changho Choi and Vijay Balakrishnan, "**FIOS: Feature Based I/O Stream Identification for Improving Endurance of Multi-Stream SSDs**", 2018 *IEEE International Conference on Cloud Computing (CLOUD'18)*, San Francisco, CA, 2018. Acceptance Rate: 15%. **(Best Paper Award)**
7. Janki Bhimani, Ningfang Mi, and Bo Sheng, "**BloomStream: Data Temperature Identification for Flash Based Memory Storage Using Bloom Filters**", 2018 *IEEE International Conference on Cloud Computing (CLOUD'18)*, San Francisco, CA, 2018. Acceptance Rate: 15%.
8. Zhengyu Yang, Morteza Hoseinzadeh, Ping Wong, John Artoux, Clay Mayers, David Thomas Evans, Rory Thomas Bolt, Janki Bhimani, Ningfang Mi, and Steven Swanson, "**H-NVMe: A Hybrid Framework of NVMe-based Storage System in Cloud Computing Environment**", *IEEE International Performance Computing and Communications Conference (IPCCC'17)*, San Diego, CA, 2017. **(Best Paper Award)**
9. Zhengyu Yang, Morteza Hoseinzadeh, Allen Andrews, Clay Mayers, David Thomas Evans, Rory Thomas Bolt, Janki Bhimani, Ningfang Mi, and Steven Swanson, "**AutoTiering: Automatic Data Placement Manager in Multi-Tier All-Flash Datacenter**", *IEEE International Performance Computing and Communications Conference (IPCCC'17)*, San Diego, CA, 2017.
10. Janki Bhimani, Ningfang Mi, Miriam Leiser, and Zhengyu Yang, "**FiM: Performance Prediction Model for Parallel Computation in Iterative Data Processing Applications**", *IEEE International Conference on Cloud Computing (CLOUD'17)*, Honolulu, HI, 2017. Acceptance Rate: 18%.
11. Han Gao, Zhengyu Yang, Janki Bhimani, Teng Wang, Jiayin Wang, Ningfang Mi, and Bo Sheng, "**AutoPath: Harnessing Parallel Execution Paths for Efficient Resource Allocation in Multi-Stage Big Data Frameworks**", *International Conference on Computer Communications and Networks (ICCCN'17)*, Vancouver, Canada, 2017. Acceptance Rate: 25%.
12. Qiumin Xu, Manu Awasthi, Krishna T. Malladi, Janki Bhimani, Jingpei Yang, and Murali Annavaram. "**Performance analysis of containerized applications on local and remote storage**" *International Conference on Massive Storage Systems and Technology (MSST'17)*, Santa Clara, CA, 2017.
13. Janki Bhimani, Jingpei Yang, Zhengyu Yang, Ningfang Mi, Qiumin Xu, Manu Awasthi, Rajinikanth Pandurangan, and Vijay Balakrishnan, "**Understanding Performance of I/O Intensive Containerized**

Applications for NVMe SSDs”, *IEEE International Performance Computing and Communications Conference (IPCCC'16)*, Las Vegas, NV, 2016. Acceptance Rate: 25.50%.

14. Zhengyu Yang, Jianzhe Tai, Janki Bhimani, Jiayin Wang, Ningfang Mi, and Bo Sheng, “**GREM: Dynamic SSD Resource Allocation in Virtualized Storage Systems with Heterogeneous VMs**”, *IEEE International Performance Computing and Communications Conference (IPCCC'16)*, Las Vegas, NV, 2016. Acceptance Rate: 25.50%.

Other Conference and Workshop Publications

Acceptance rates provided when known

15. Janki Bhimani, Jingpei Yang, Zhengyu Yang, Ningfang Mi, NHV Krishna Giri, Rajinikanth Pandurangan, Changho Choi, and Vijay Balakrishnan. “**Enhancing SSDs with multi-stream: What? why? how?**” *IEEE International Performance Computing and Communications Conference (IPCCC'17)*, San Diego, CA, 2017. (Short Paper)
16. Baiyu Chen, Zhengyu Yang, Siyu Huang, Xianzhi Du, Zhiwei Cui, Janki Bhimani, Xin Xie, and Ningfang Mi, “**Cyber-Physical System Enabled Nearby Traffic Flow Modelling for Autonomous Vehicles**”, *IEEE International Performance Computing and Communications Conference, Special Session on Cyber Physical Systems: Security, Computing, and Performance (IPCCC CPS'17)*, San Diego, CA, 2017.
17. Janki Bhimani, Zhengyu Yang, Miriam Leeser, and Ningfang Mi, “**Accelerating Big Data Applications Using Lightweight Virtualization Framework on Enterprise Cloud**”, *IEEE High Performance Extreme Computing Conference (HPEC'17)*, Waltham, MA, 2017.
18. Qiumin Xu, Manu Awasthi, Krishna T. Malladi, Janki Bhimani, Jingpei Yang, Murali Annavaram, “**Docker Characterization on High Performance SSDs**”, *IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS'17)*, Santa Rosa, CA, 2017. (Short Paper)
19. Liu Chao, Janki Bhimani, and Miriam Leeser, “**Using High Level GPU Tasks to Explore Memory and Communications Options on Heterogeneous Platforms**” *ACM Workshop on Software Engineering Methods for Parallel and High Performance Applications (SEM4HPC)*, Washington, D.C., 2017.
20. Liu Chao, Janki Bhimani, and Miriam Leeser, “**Exploring Memory Options for Data Transfer on Heterogeneous Platforms**”, *The International ACM Symposium on High-Performance Parallel and Distributed Computing (HPDC'17)*, Washington, D.C., 2017. (Short Paper)
21. Janki Bhimani, Miriam Leeser, and Ningfang Mi, “**Performance Prediction Techniques for Scalable Large Data Processing in Distributed MPI Systems**”, *IEEE International Performance Computing and Communications Conference (IPCCC'16)*, Las Vegas, NV, 2016. Acceptance Rate: 12%. (Short Paper)
22. Janki Bhimani, Miriam Leeser, and Ningfang Mi, “**Design Space Exploration of GPU Accelerated Cluster Systems for Optimal Data Transfer Using PCIe Bus**”, *IEEE High Performance Extreme Computing Conference (HPEC'16)*, Waltham, MA, 2016.
23. Janki Bhimani, Miriam Leeser, and Ningfang Mi, “**Accelerating K-Means Clustering with Parallel Implementations and GPU Computing**”, *IEEE High Performance Extreme Computing Conference (HPEC'15)*, Waltham, MA, 2015.
24. Janki Bhimani, Miriam Leeser and Ningfang Mi, “**Predicting the Performance of Machine Learning Algorithms running on Heterogeneous Computing Platforms**” *Women in Machine Learning Workshop (WiML'14)*, Montréal, Canada, 2014.

PATENTS

1. Janki Bhimani, Hingkwon Huen, Jingpei Yang, Manu Awasthi, Vijay Balakrishnan, Jason Martineau, inventors; Samsung Electronics Co Ltd, assignee. “**Intelligent controller for containerized applications**” United States patent application *US 15/379,327*. 2018 Mar 1.
2. Janki Bhimani, Jingpei Yang, Changho Choi, Jianjian Huo, inventors; Samsung Electronics Co Ltd, assignee. “**Smart I/O stream detection based on multiple attributes**” United States patent application *US 15/344,422*. 2017 Mar 16.
3. Janki Bhimani, Anand Subramanian, Vijay Balakrishnan, and Jingpei Yang, inventors; Samsung Electronics Co Ltd, assignee. “**Container workload scheduler and methods of scheduling container workloads**” United States patent application *US15/820856*.

4. Janki Bhimani, Rajinikanth Pandurangan, Vijay Balakrishnan, Changho Choi, inventors; Samsung Electronics Co Ltd, assignee. “**Representative I/O generator**” United States patent application *US 15/853419*.
5. Janki Bhimani, Rajinikanth Pandurangan, Changho Choi, Vijay Balakrishnan, inventors; Samsung Electronics Co Ltd, assignee. “**System and method for identifying hot data and stream in a solid-state drive**” *US 15/895797*.
6. Janki Bhimani, Jingpei Yang, Changho Choi, inventors; Samsung Electronics Co Ltd, assignee. “**Parallel key value based multi-thread machine learning exploiting KV-SSDs**” (Under Processing)

PROFESSIONAL ACTIVITIES

Service as Invited Technical Program Committee Member

- TCP for Work-in-Progress Track of IEEE International Conference on Cloud Computing (IEEE CLOUD) 2018
- TCP for IEEE International Performance Computing and Communications Conference (IPCCC) 2018

PEER REVIEW

Journal

- Simulation Modelling Practice and Theory (SIMPAT) Elsevier Journal, 2016, 2018
- Computers MDPI Journal, 2017
- Future Generation Computer Systems (FGCS) Elsevier Journal, 2018
- Transactions on Computers (TC), IEEE Journal, 2018
- ACM Transactions on Modeling and Performance Evaluation of Computing Systems (TOMPECS) ACM Journal, 2018

Conference

- IEEE International Conference on Cloud Computing (IEEE CLOUD), 2015, 2016, 2017, 2018
 - IEEE High Performance Extreme Computing Conference (IEEE HPEC), 2015
 - IEEE International Conference on Green Computing and Communications (GreenCom), 2015
 - International Conference on Massive Storage Systems and Technology (MSST), 2016
 - IEEE International Conference on Big Data (BigData), 2016, 2017
 - International Conference on Networking, Architecture, and Storage (NAS), 2016
 - International Conference on Parallel and Distributed Systems (ICPADS), 2016
 - Workshop on Interactions of NVM/Flash with Operating Systems and Workloads (INFLOW), 2016
 - International Conference on Performance Engineering (ICPE), 2017, 2018
 - ACM/SIGDA International Symposium on Field-Programmable Gate Arrays (FPGA), 2017
 - IEEE/IFIP International Conference on Dependable Systems and Networks (DNS), 2017
 - Big Data and Cloud Performance Workshop at INFOCOM (DCPerf), 2017
 - International Conference on Autonomic Computing (ICAC), 2017
 - International Conference on Computer Aided Design (ICCAD), 2017
 - International Conference on Cloud Computing Technologies and Applications (CloudTech), 2017
 - Field-Programmable Custom Computing Machines (FCCM), 2018
 - International Conference on Computer. Communication and Networks (ICCCN), 2018
 - IEEE International Performance Computing and Communications Conference (IPCCC), 2018
 - IEEE/ACM International Conference on Utility and Cloud Computing (UCC), 2018
-

REFERENCES

Ningfang Mi

Associate Professor

*Department of Electrical and Computer Engineering
Northeastern University*

Phone: (617) 373-3028

Email: ningfang@ece.neu.edu

Waleed Meleis

Associate Chair & Associate Professor

*Department of Electrical and Computer Engineering
Northeastern University*

Phone: (617) 373-3023

Email: meleis@ece.neu.edu

Vijay Balakrishnan

Director

*Memory Platform Lab
Samsung Semiconductor Inc.*

Phone: (408) 390-9205

Email: vijay.bala@samsung.com

Miriam E. Leeser

Professor

*Department of Electrical and Computer Engineering
Northeastern University*

Phone: (617) 373-3814

Email: mel@coe.neu.edu

David R. Kaeli

Distinguished Professor

*Department of Electrical and Computer Engineering
Northeastern University*

Phone: (617) 373-5413

Email: kaeli@ece.neu.edu