

Nikhil Jain

Fernbach Postdoctoral Fellow, Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

Livermore, CA, 94551
☎ 217.979.0918
✉ nikhil.jain@acm.org
<http://nikhil-jain.github.io/>

Education

- 2016 **Ph.D., Computer Science**, *University of Illinois at Urbana-Champaign*.
Optimization of Communication Intensive Applications on HPC Networks. Advisor: Laxmikant V. Kale
- 2009 **M.Tech., Computer Science and Engineering**, *Indian Institute of Technology*, Kanpur, India.
- 2009 **B.Tech., Computer Science and Engineering**, *Indian Institute of Technology*, Kanpur, India.

Awards & Honors

- 2017 NERSC Award for Innovative Use of HPC
- 2017 Finalist, 10th IEEE International Scalable Computing Challenge
- 2017 Outstanding Mentor Award, Lawrence Livermore National Laboratory
- 2017 Best Poster Award (2nd Place), Computation Postdoc Poster Symposium, LLNL
- 2016 Fernbach Postdoctoral Fellowship
- 2014 IBM PhD Fellowship
- 2013 Silver Award, ACM Student Research Poster Competition, Supercomputing
- 2012 Best Poster, LLNL Annual Student Poster Session
- 2012 Co-winner of HPC Challenge Class 2 (performance), Charm++ Team, Supercomputing
- 2011 Andrew and Shana Laursen Fellowship, Department of Computer Science, UIUC

Skills

- General C, C++, Fortran, Python, numpy, git, svn, CVS, autotools, CMake.
- Parallel Computing Parallel application development in MPI, OpenMP, and Charm++. Parallel runtime system development. Performance tool development in C/C++. MPI and Charm++ profiling tools. Analysis and tuning of parallel applications on large clusters (IBM Blue Gene, Cray, Linux). Experience with large science codes at UIllinois.
- Data Analytics Application of machine learning for application tuning, root cause analysis, and performance analysis of parallel codes.
- Networks and Communication HPC Network simulation tools. Communication and network modelling and analysis. Collective operations on large scale systems. Performance monitoring of networks.

Research and Work Experience

- 2016-present **Lawrence Livermore National Laboratory**, *Fernbach Postdoctoral Fellow*.
 - Research, characterize, and improve communication throughput of HPC networks.
 - Develop data analytics driven methods for performance optimization.
 - Explore new avenues for application of parallel computing.
- 2011-2016 **Dept of Computer Science, University of Illinois at Urbana-Champaign**, *Research Assistant*.
 - Created TraceR and contributed to CODES to enable scalable and accurate simulation of parallel codes on HPC networks.
 - Proposed DamselFly - a fast analytical model for understanding performance of dragonfly networks.
 - Developed new features for Charm++, a widely used parallel programming model.
 - Contributed to development and performance improvements in OpenAtom and NAMD.
- Summer 2014 **Lawrence Livermore National Laboratory**, *Research Intern*.
Analyzed prediction models to identify factors that affect performance on HPC networks.

- Summer 2012 **Lawrence Livermore National Laboratory**, *Research Intern*.
Proposed and demonstrated use of machine learning based models for predicting communication performance of HPC applications.
- Spring 2011 **Dept of Computer Science, University of Illinois at Urbana-Champaign**, *Visiting Scholar*.
Re-designed runtime-level memory management module for speed and efficiency in Charm++.
- 2009-2011 **IBM Research, India**, *Blue Scholar*.
Developed 3× faster topology-aware algorithms for collective communication on Blue Gene systems.
- Summer 2007 **Polytechnic Institute of NYU**, *Undergraduate Research Intern*.
Investigated use of obfuscation methods in logic design.

Teaching Experience

- Fall 2014 **Teaching Assistant**, *Ullinois Urbana-Champaign*, Conducted discussion hours, and designed lectures, homework assignments, and exams for Introduction to Parallel Programming.
- Spring 2009 **Teaching Assistant**, *IIT Kanpur*, Conducted discussion hours and coding sessions for Fundamental of Computing.
- Fall 2008 **Teaching Assistant**, *IIT Kanpur*, Conducted discussion hours and coding sessions for Fundamental of Computing.

Research and Other Funding

- 2017-present **Co-I**, *ECP Proxy Applications Project*, \$2,000,000/year.
PI: David Richards, LLNL. DOE ECP (website)
- 2017-present **Co-I**, *Workload-driven Design Space Exploration*, \$500,000/year.
PI: Martin Schulz, LLNL. DOE ECP Hardware Technologies.
- 2016-present **Co-I**, *Statistical Methods for Exascale Performance Modeling*, \$2,500,000.
PI: Todd Gamblin, U.S. Department of Energy (DOE) Early Career Research Program.
- 2016 **Co-I**, *Exploring Asynchronous Task-based Models and Runtimes in WSC Codes*, \$50,000.
PI: Abhinav Bhatele, Linking Exploratory Application Research to Next-gen development (LEARN) program.

Mentoring/Co-advicing

- Students Kevin Brown (Tokyo Tech), Emilio Castillo (BSC), Vanessa Cedeno (Virginia Tech), Aditya Nigam (IIT Kanpur), Samuel Pollard (UOregon), Alexander Robey (Swarthmore), Karthik Senthil (Illinois), Palash Sharma (IIT Kanpur), Staci Smith (Arizona), Rahul Sridhar (UC Irvine), Ankit Srivastava (Georgia Tech), Andrew R. Titus (MIT), Tao Wang (NC State)

Committees

- Technical Program Committees COMHPC '16, Cluster '17, HPC Asia '18, IPDPS '18, PMBS '17, SC '18, SCSC '16 '17, Supercomputing Asia '18
- Grant Review Committees DOE Small Business Innovation Research (SBIR) 2016.
- Other Reviewing TPDS, IJHPCA, JPDC, Biophysical, PACT '17, IPDPS '14 '17, Cluster '12

Software Projects

- traceR Trace-driven PDES-based network simulation of parallel applications.
- Damselfly Analytical modeling of steady-state traffic flow on dragonfly networks.
- bgnql An easy interface for network counters collection on BlueGene/Q systems.
- LeanMD A mini-application for exploring optimizations for molecular dynamics.
- topoAPI Fully-automated profiling, analysis, and mapping on HPC systems.
- Charm-ROSS Charm++ based implementation of the parallel discrete event simulator ROSS.

Open-source software I contribute to

CODES	Packet-level HPC network simulation.
Charm++	Object-based message driven parallel programming paradigm.
OpenAtom	Scalable Ab-Initio Molecular Dynamics with diverse features.
EpiSimdemics	Hybrid contagion modeling using Charm++.
ROSS	Scalable parallel discrete simulation with optimistic execution.
NAMD	Highly scalable parallel molecular dynamics.

Publications

Peer Reviewed Conference & Journal Papers

- [1] Jayaraman Thiagarajan, Rushil Anirudh, Bhavya Kailkhura, **Nikhil Jain**, Tanzima Islam, Abhinav Bhatele, Jae-Seung Yeom, and Todd Gamblin. PADDLE: Performance analysis using a data-driven learning environment. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '18* (to appear). IEEE Computer Society, May 2018.
- [2] Misbah Mubarak, **Nikhil Jain**, Jens Domke, Noah Wolfe, Caitlin Ross, Kelvin Li, Abhinav Bhatele, Christopher D. Carothers, and Rob Ross. Toward reliable validation of HPC network simulation models. *Winter Simulation Conference*, December 2017.
- [3] **Nikhil Jain**, Abhinav Bhatele, Louis Howell, David Bohme, Ian Karlin, Edgar Leon, Misbah Mubarak, Noah Wolfe, Todd Gamblin, and Matthew Leininger. Predicting the performance impact of different fat-tree configurations. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '17*, November 2017.
- [4] Aniruddha Marathe, Rushil Anirudh, **Nikhil Jain**, Abhinav Bhatele, Jayaraman Thiagarajan, Bhavya Kailkhura, Jae-Seung Yeom, Barry Rountree, and Todd Gamblin. Performance modeling under resource constraints using deep transfer learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '17*, November 2017.
- [5] Misbah Mubarak, Phil Carns, Jonathan Jenkins, Kelvin Li, **Nikhil Jain**, Shane Snyder, Rob Ross, Christopher D. Carothers, Abhinav Bhatele, and Kwan-Liu Ma. Quantifying i/o and communication traffic interference on dragonfly networks equipped with burst buffers. In *2017 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 204–215, Sept 2017.
- [6] Xiang Ni, **Nikhil Jain**, Kavitha Chandrasekar, and Laxmikant Kale. Runtime techniques for programming with fast and slow memory. In *2017 IEEE International Conference on Cluster Computing (CLUSTER)*, pages 147–151, Sept 2017.
- [7] Juan J. Galvez, **Nikhil Jain**, and Laxmikant V. Kale. Automatic topology mapping of diverse large-scale parallel applications. In *Proceedings of the International Conference on Supercomputing, ICS '17*, pages 17:1–17:10, New York, NY, USA, June 2017. ACM.
- [8] Abhinav Bhatele, Jae-Seung Yeom, **Nikhil Jain**, Chris J. Kuhlman, Yarden Livnat, Keith R. Bisset, Laxmikant V. Kale, and Madhav V. Marathe. Massively parallel simulations of spread of infectious diseases over realistic social networks. In *Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGrid '17*, pages 689–694, Piscataway, NJ, USA, May 2017. IEEE Press.
- [9] **Nikhil Jain**, Abhinav Bhatele, Xiang Ni, Todd Gamblin, and Laxmikant V. Kale. Partitioning low-diameter networks to eliminate inter-job interference. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '17*. IEEE Computer Society, May 2017.
- [10] Noah Wolfe, Misbah Mubarak, **Nikhil Jain**, Jens Domke, Abhinav Bhatele, Christopher D. Carothers, and Robert B. Ross. Preliminary performance analysis of multi-rail fat-tree networks. *Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, May 2017.
- [11] **Nikhil Jain**, Eric Bohm, Eric Mikida, Subhasish Mandal, Minjung Kim, Prateek Jindal, Qi Li, Sohrab Ismail-Beigi, Glenn Martyna, and Laxmikant Kale. Openatom: Scalable ab-initio molecular dynamics with diverse capabilities. In *International Supercomputing Conference, ISC HPC '16*, 2016.
- [12] **Nikhil Jain**, Abhinav Bhatele, Samuel T. White, Todd Gamblin, and Laxmikant V. Kale. Evaluating HPC networks via simulation of parallel workloads. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '16*. IEEE Computer Society, November 2016. LLNL-CONF-690662.
- [13] Abhinav Bhatele, **Nikhil Jain**, Yarden Livnat, Valerio Pascucci, and Peer-Timo Bremer. Analyzing network health and congestion in dragonfly-based systems. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '16*. IEEE Computer Society, May 2016. LLNL-CONF-678293.

- [14] Eric Mikida, **Nikhil Jain**, Elsa Gonsiorowski, Peter D. Barnes, Jr., David Jefferson, Christopher D. Carothers, and Laxmikant V. Kale. Towards pdes in a message-driven paradigm: A preliminary case study using charm++. In *ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS)*, SIGSIM PADS '16. ACM, May 2016.
- [15] Abhinav Bhatele, Andrew R. Titus, Jayaraman J. Thiagarajan, **Nikhil Jain**, Todd Gamblin, Peer-Timo Bremer, Martin Schulz, and Laxmikant V. Kale. Identifying the culprits behind network congestion. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, IPDPS '15. IEEE Computer Society, May 2015. LLNL-CONF-663150.
- [16] **Nikhil Jain**, Abhinav Bhatele, Jae-Seung Yeom, Mark F. Adams, Francesco Miniati, Chao Mei, and Laxmikant V. Kale. Charm++ & MPI: Combining the best of both worlds. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, IPDPS '15. IEEE Computer Society, May 2015. LLNL-CONF-663041.
- [17] Ehsan Ttoni, **Nikhil Jain**, and Laxmikant Kale. Power management of extreme-scale networks with on/off links in runtime systems. *ACM Transactions on Parallel Computing*, 2014.
- [18] Abhinav Bhatele, **Nikhil Jain**, Katherine E. Isaacs, Ronak Buch, Todd Gamblin, Steven H. Langer, and Laxmikant V. Kale. Improving application performance via task mapping on IBM Blue Gene/Q. In *Proceedings of IEEE International Conference on High Performance Computing*, HiPC '14. IEEE Computer Society, December 2014. LLNL-CONF-655465.
- [19] Bilge Acun, Abhishek Gupta, **Nikhil Jain**, Akhil Langer, Harshitha Menon, Eric Mikida, Xiang Ni, Michael Robson, Yanhua Sun, Ehsan Ttoni, Lukasz Wesolowski, and Laxmikant Kale. Parallel Programming with Migratable Objects: Charm++ in Practice. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '14, November 2014.
- [20] **Nikhil Jain**, Abhinav Bhatele, Xiang Ni, Nicholas J. Wright, and Laxmikant V. Kale. Maximizing throughput on a dragonfly network. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '14. IEEE Computer Society, November 2014. LLNL-CONF-653557.
- [21] James C. Phillips, Yanhua Sun, **Nikhil Jain**, Eric J. Bohm, and Laxmikant V. Kale. Mapping to Irregular Torus Topologies and Other Techniques for Petascale Biomolecular Simulation. In *Proceedings of ACM/IEEE SC 2014*, New Orleans, Louisiana, November 2014.
- [22] **Nikhil Jain**, Abhinav Bhatele, Michael P. Robson, Todd Gamblin, and Laxmikant V. Kale. Predicting application performance using supervised learning on communication features. In *ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '13. IEEE Computer Society, November 2013. LLNL-CONF-635857.
- [23] Xiang Ni, Esteban Meneses, **Nikhil Jain**, and Laxmikant V. Kale. Acr: Automatic checkpoint/restart for soft and hard error protection. In *ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '13. IEEE Computer Society, November 2013.
- [24] **Nikhil Jain**, JohnMark Lau, and Laxmikant V. Kalé. Collectives on two-tier direct networks. In *EuroMPI*, pages 67–77, 2012.
- [25] Anshul Mittal, **Nikhil Jain**, Thomas George, Yogish Sabharwal, and Sameer Kumar. Collective algorithms for sub-communicators. In *Proceedings of the 26th ACM International Conference on Supercomputing*, ICS '12, pages 225–234, New York, NY, USA, 2012. ACM.
- [26] Harshitha Menon, **Nikhil Jain**, Gengbin Zheng, and Laxmikant V. Kalé. Automated load balancing invocation based on application characteristics. In *IEEE Cluster 12*, Beijing, China, September 2012.
- [27] Abhinav Bhatele, **Nikhil Jain**, William D. Gropp, and Laxmikant V. Kale. Avoiding hot-spots on two-level direct networks. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '11, pages 76:1–76:11, New York, NY, USA, 2011. ACM.
- [28] Anshul Mittal, Jagabondhu Hazra, **Nikhil Jain**, Vivek Goyal, Deva P. Seetharam, and Yogish Sabharwal. Real time contingency analysis for power grids. In *Proceedings of the 17th international conference on Parallel processing - Volume Part II*, Euro-Par'11, pages 303–315, Berlin, Heidelberg, 2011. Springer-Verlag.
- [29] Ehsan Ttoni, Abhinav Bhatele, Eric Bohm, **Nikhil Jain**, Celso Mendes, Ryan Mokos, Gengbin Zheng, and Laxmikant Kale. Simulation-based performance analysis and tuning for a two-level directly connected system. In *Proceedings of the 17th IEEE International Conference on Parallel and Distributed Systems*, December 2011.

- [30] Venkatesan T. Chakaravarthy, **Nikhil Jain**, and Yogish Sabharwal. Optimizing matrix transpose on torus interconnects. In *Proceedings of the 16th international Euro-Par conference on Parallel processing: Part II*, Euro-Par'10, pages 440–451, Berlin, Heidelberg, 2010. Springer-Verlag.
- [31] **Nikhil Jain** and Yogish Sabharwal. Optimal bucket algorithms for large MPI collectives on torus interconnects. In *Proceedings of the 24th ACM International Conference on Supercomputing*, ICS '10, pages 27–36, New York, NY, USA, 2010. ACM.

Workshop Papers

- [32] Bilge Acun, **Nikhil Jain**, Abhinav Bhatele, Misbah Mubarak, Christopher D. Carothers, and Laxmikant V. Kale. Preliminary evaluation of a parallel trace replay tool for hpc network simulations. In *Proceedings of the 3rd Workshop on Parallel and Distributed Agent-Based Simulations*, PADABS '15, August 2015. LLNL-CONF-667225.
- [33] Ehsan Toton, **Nikhil Jain**, and Laxmikant V. Kale. Toward runtime power management of exascale networks by on/off control of links. In *Parallel and Distributed Processing Workshops and Phd Forum (IPDPSW), 2013 IEEE International Symposium on*, 2013.
- [34] Laxmikant Kale, **Nikhil Jain**, Akhil Langer, Esteban Meneses, Phil Miller, Osman Sarood, and Ehsan Toton. A Multi-resolution Emulation + Simulation Methodology. In *DoE Workshop on Modeling and Simulation of Exascale Systems and Applications (MODSIM)*, Seattle, Washington, September 2013.
- [35] Laxmikant Kale, Osman Sarood, Eric Bohm, **Nikhil Jain**, Akhil Langer, and Esteban Meneses. Actionable Performance Modeling for Future Supercomputers. In *DoE Workshop on Modeling and Simulation of Exascale Systems and Applications (MODSIM)*, Seattle, Washington, September 2013.
- [36] **Nikhil Jain**, Brajesh Pande, and Phalguni Gupta. Smp based solver for large binary systems. In *Proceedings of the 2009 International Conference on Parallel and Distributed Computing, Applications and Technologies*, PDCAT '09, pages 426–431, Washington, DC, USA, 2009. IEEE Computer Society.

Ph.D. Dissertation

- [37] **Nikhil Jain**. *Optimization of communication intensive applications on HPC networks*. PhD thesis, University of Illinois at Urbana-Champaign, Urbana, Illinois, February 2016.

Posters

- [38] Rahul Sridhar, Rushil Anirudh, Jayaraman J. Thiagarajan, **Nikhil Jain**, and Todd Gamblin. Model-agnostic influence analysis for performance data. In *SC '17 (to appear)*, Denver, Colorado, USA, November 2017.
- [39] **Nikhil Jain**. Improving performance of networks & applications using simulations. In *LLNL Computation Postdoctoral Poster Session*, Livermore, CA, May 2017.
- [40] Abhinav Bhatele, **Nikhil Jain**, Yarden Livnat, Valerio Pascucci, and Peer-Timo Bremer. Simulating and visualizing traffic on the dragonfly network. In *SC '15*, Austin, Texas, USA, November 2015.
- [41] **Nikhil Jain**, Abhinav Bhatele, Jae-Seung Yeom, Mark F. Adams, Francesco Miniati, Chao Mei, and Laxmikant V. Kale. Interoperating MPI and Charm++ for productivity and performance. In *SC '14*, New Orleans, Louisiana, USA, November 2014.
- [42] **Nikhil Jain** and Laxmikant Kale. Collective algorithms for sub-communicators. In *SC '13*, Denver, Colorado, USA, November 2013.
- [43] **Nikhil Jain**, Abhinav Bhatele, and Todd Gamblin. Understanding network contention on Blue Gene supercomputers. In *LLNL Computation Summer Student Poster Session*, Livermore, CA, August 2012.
- [44] Anshul Mittal, **Nikhil Jain**, Thomas George, Yogish Sabharwal, and Sameer Kumar. Collective algorithms for sub-communicators. In *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP '12, pages 315–316, New York, NY, USA, February 2012. ACM.