



# Abhinav Bhatele

Brendan Iribe Center 5218  
8125 Paint Branch Drive  
College Park, MD 20742

 Associate Professor, Department of Computer Science  
University of Maryland, College Park

 [bhatele@cs.umd.edu](mailto:bhatele@cs.umd.edu)  
 [cs.umd.edu/~bhatele](http://cs.umd.edu/~bhatele)

## Education

- 2010 **Ph.D., Computer Science**, *University of Illinois at Urbana-Champaign, Urbana, Illinois, USA*
- 2007 **M.S., Computer Science**, *University of Illinois at Urbana-Champaign, Urbana, Illinois, USA*
- 2005 **B.Tech., Computer Science and Engineering**, *Indian Institute of Technology, Kanpur, India*

## Research and Professional Experience

### University of Maryland, College Park

- 2024–present Affiliate Faculty Member, *Artificial Intelligence Interdisciplinary Institute at Maryland (AIM)*
- 2022–present Associate Professor, *Department of Computer Science and Institute for Advanced Computer Studies*
- 2020–present Affiliate Faculty Member (Full), *Applied Math. & Statistics, and Scientific Computation (AMSC)*
- 2019–2022 Assistant Professor, *Department of Computer Science and Institute for Advanced Computer Studies*

### Other Appointments

- 2022–2025 Visiting Scientist, *RIKEN Center for Computational Science*
- Lawrence Livermore National Laboratory**
- 2018–2019 Principal Computer Scientist, *Center for Applied Scientific Computing*
- 2012–2018 Senior Computer Scientist, *Center for Applied Scientific Computing*
- 2011–2012 Post-doctoral Research Scholar, *Center for Applied Scientific Computing*
- University of Illinois at Urbana-Champaign**
- 2010–2011 Post-doctoral Research Associate, *Department of Computer Science*
- 2005–2010 Graduate Research Assistant, *Department of Computer Science*
- Summer 2007 **IBM Thomas J. Watson Research Center**, *Research Intern*
- Summer 2006 **IBM Thomas J. Watson Research Center**, *Research Intern*
- Summer 2004 **INRIA Nancy – Grand-Est, France**, *Research Intern*
- Summer 2003 **Indian Institute of Technology, Kanpur, India**, *Research Intern*

## Awards and Honors

- 2025 Research Excellence Honoree, University of Maryland
- 2024 Hyperion Research HPC Innovation Excellence Award: AxoNN
- 2024 **Finalist, ACM Gordon Bell Prize**, SC '24, Atlanta, GA
- 2024 Early Career Academic Achievement Alumni Award, Dept. of CS, Univ. of Illinois
- 2024 Best Paper Award, PDP '24, Dublin, Ireland
- 2024 DOE INCITE Leadership Computing Allocation Award
- 2023 **IEEE TCSC Award for Excellence in Scalable Computing** for Mid-career Researchers
- 2023 Best Research Poster Award, SC '23, Denver, CO
- 2021 NSF CAREER Award
- 2019 Technical University Munich (TUM) Institute for Advanced Study Visiting Fellow
- 2017, 2018 Outstanding Mentor Award, Lawrence Livermore National Laboratory
- 2018 **Early and Mid-Career Recognition Award**, Lawrence Livermore National Laboratory
- 2018 Best Student Paper Finalist, SC '18, Dallas, TX
- 2017 NERSC Award for Innovative Use of High Performance Computing

- 2017 Deputy Director for S&T Excellence in Publication Award, LLNL
- 2017 Finalist, 10th IEEE International Scalable Computing Challenge (SCALE 2017)
- 2016 Best Paper Award, IPDPS '16, Chicago, IL
- 2014 Best Report award, AX Division, Lawrence Livermore National Laboratory
- 2014 **IEEE TCSC Award for Excellence in Scalable Computing** for Early Career Researchers
- 2013 Best Paper Award, IPDPS '13, Boston, MA
- 2012 Best Poster Award (1st Place), Computation Postdoc Poster Symposium, LLNL
- 2011 Member of winning team for HPC Challenge Class II Award, SC '11, Seattle, WA
- 2011 **David J. Kuck Outstanding Ph.D. Thesis Award**, Dept. of Computer Science, Univ. of Illinois
- 2010 Feng Chen Memorial Best Paper Award, Dept. of Computer Science, Univ. of Illinois
- 2010 Teacher Scholar Certificate, Center of Teaching Excellence, Univ. of Illinois
- 2010 Graduate Teacher Certificate, Center of Teaching Excellence, Univ. of Illinois
- 2009 **ACM/IEEE-CS George Michael Memorial HPC Fellowship**, SC '09, Portland, OR
- 2009 Distinguished Paper Award, Euro-Par '09, Delft, The Netherlands
- 2009 David J. Kuck Outstanding M.S. Thesis Award, Dept. of Computer Science, Univ. of Illinois
- 2008 Best Graduate Poster (3rd Place) at the ACM Student Research Competition, SC '08

---

## Keynotes and Invited Panelist

- November 2025 **Panel** Students@SC Panel on *Navigating Education Systems Internationally*. Supercomputing (SC '25), St. Louis, MO
- September 2025 **Keynote** *Lost in Translation: LLMs and Whole-Repository Porting of HPC Codes*. Workshop on AI Assisted Software Development for High-performance Computing (AI4Dev '25), San Diego, CA
- March 2025 **Panel** *Explore AI-Assisted Developer Tools for Accelerated Computing Application Development*. NVIDIA GPU Technology Conference (GTC '25), San Jose, CA
- November 2024 **Panel** Early Career Program Panel on *Resilience and Time Management Tips*. Supercomputing (SC '24), Atlanta, GA
- June 2024 **Keynote** *The Future of Parallel Code Development: Will AI Lead the Way?* International Symposium on High-Performance Parallel and Distributed Computing (HPDC '24), Pisa, Italy
- March 2024 **Panel** *Exploring AI-Assisted Developer Tools for Accelerated Computing*. NVIDIA GPU Technology Conference (GTC '24), San Jose, CA
- January 2024 **Keynote** *Making Distributed Training Easier and More Communication-efficient*. International Conference on Cloud Computing, Data Science & Engineering (IEEE Confluence '24), Noida, India
- September 2022 **Keynote** *Using system monitoring data to understand and mitigate performance variability*. Workshop on Monitoring and Analysis for HPC Systems Plus Applications (HPCMASPA '22), Heidelberg, Germany
- December 2021 **Panel** ASCR Workshop on *Science of Scientific-Software Development*. Department of Energy (virtual)
- September 2011 **Panel** *Evaluation of High-Performance Computing Systems: How Well Are We Doing?* High Performance Computing and Communications (HPCC '11), Banff, Canada

---

## Invitation-only Meetings and Conferences

- February 2025 CRA Leadership Academy, Washington D.C.
- April 2021 Salishan Conference on High Speed Computing (virtual)
- July 2018 GI-Dagstuhl Seminar on Visualizing Systems and Software Performance (VSSP), Wadern, Germany
- January 2014 Dagstuhl Perspectives Workshop on "Connecting Performance Analysis and Visualization to Advance Extreme Scale Computing", Wadern, Germany
- September 2013 1st Heidelberg Laureate Forum, Heidelberg, Germany

---

## Media Coverage and Interviews

- Nov 21, 2024 Hyperion Research Announces Winners of 2024 HPC Innovation Excellence Awards. 📄 hpcwire.com
- Nov 12, 2024 Gordon Bell Prize nomination recognizes efforts to train extreme-scale large language models using Frontier. 📄 ornl.gov
- Sep 4, 2024 A day in the life of the world's fastest supercomputer. 📄 nature.com
- May 13, 2024 Bhatele Receives University of Illinois Early Career Alumni Award. 📄 umd.edu
- Dec 4, 2023 University of Maryland Team Wins DOE Award to Advance AI Using Supercomputers. 📄 hpcwire.com
- Apr 28, 2022 What's New in HPC Research: E3SM Diags, Sunway Supercomputer, AxoNN, Snellius Supercomputer & More. 📄 hpcwire.com
- Feb 25, 2021 Bhatele Receives NSF CAREER Award to Optimize Parallel Software and Systems. 📄 umd.edu
- Nov 16, 2013 SC13 Research Highlight: There Goes the Neighborhood. 📄 hpcwire.com
- Apr 16, 2013 George Michael Fellowship Helps HPC Community Thrive. 📄 hpcwire.com
- Nov 30, 2009 CS graduate student earns national honor. 📄 illinois.edu

---

## Grants

- 2025–2026 **PI**, *Modeling and Optimizing the Performance of HPC Codes and AI Frameworks*, \$271,000. DOE/LLNL
- 2024–2027 **UMD PI**, *Productive AI-Assisted HPC Software Ecosystem*, Total: \$7,050,000, UMD: \$1,050,000. Lead PI: Harshitha Menon, LLNL. DOE ASCR
- 2021–2027 **PI**, *CAREER: Self-tuning Parallel Software and Systems*, \$68,000. NSF REU
- 2021–2027 **PI**, *CAREER: Self-tuning Parallel Software and Systems*, \$550,000. NSF
- 2020–2025 **Co-I\***, *Collaborative Research: Expeditions: Global Pervasive Computational Epidemiology*, \$56,000. Lead PI: Madhav Marathe, Univ. of Virginia. NSF REU
- 2020–2025 **Co-I\***, *Collaborative Research: Expeditions: Global Pervasive Computational Epidemiology*, UMD: \$219,000. Lead PI: Madhav Marathe, Univ. of Virginia. NSF
- 2023–2025 **PI**, *Analytics and Modeling of HPC Code and Data*, \$518,000. DOE/LLNL
- 2023–2024 **PI**, *Travel: Student Support for IEEE Cluster 2023 Conference*, \$15,000. NSF
- 2022–2023 **PI**, *Automating Performance Analysis within Distributed Workflows*, \$74,000. DOE/PNNL
- 2022–2023 **PI**, *Analytics and Modeling of HPC Performance Regression Data*, \$196,000. DOE/LLNL
- 2020–2022 **PI**, *Analyzing and Optimizing I/O and Overall Performance of HPC Applications*, \$519,000. DOE/LLNL
- 2020 **PI**, *Studying the Impact of Network Traffic Classes Using Simulations*, \$53,000. DOE/LLNL
- 2019 **PI**, *Validation & Hardening of the TraceR HPC Network Simulator*, \$100,000. LLNL Organizational ISCP
- 2018–2019 **PI**, *Exploring Research Ideas in Machine Learning and Performance Engineering for HPC*, \$63,000. LLNL Early and Mid-Career Recognition Award
- 2017–2019 **Co-I\***, *ECP Proxy Applications Project*, \$2,000,000/year. PI: David Richards, LLNL. DOE Exascale Computing Project (ECP).
- 2017–2019 **Co-PI**, *Workload-driven Design Space Exploration*, \$500,000/year. PI: Maya Gokhale, LLNL. DOE ECP Hardware Integration.
- 2011–2019 **Co-I\***, *Performance Analysis and Visualization at Exascale (PAVE)*, \$500,000/year. PI: Todd Gamblin, LLNL. DOE/NNSA Advanced Simulation and Computing Program.
- 2016 **PI**, *Exploring Asynchronous Task-based Models and Runtimes in WSC Codes*, \$100,000. Linking Exploratory Application Research to Next-gen development (LEARN) program.
- 2015–2017 **Co-PI**, *Beyond the Standard Model (BSM)*. \$249,000 (no-cost extension). PI: Adolfo Hoisie, PNNL. DOE ASCR

- 2014–2017 **Co-PI**, *Planetary Scale Agent Simulations*, \$2,110,000.  
PI: Peter D. Barnes Jr., LLNL. DOE LDRD Exploratory Research.
- 2013–2016 **Co-I\***, *Performance Insight for Programs and Exascale Runtimes (PIPER)*, \$4,500,000.  
PI: Martin Schulz, LLNL. DOE ASCR X-Stack II Program.
- 2012–2015 **PI**, *Task mapping on complex network topologies for improving performance*, \$945,000.  
DOE LDRD Exploratory Research.

**\*Co-I** = Co-Investigator/Senior Personnel (significant contribution to ideas in the project)

## Time Allocation Grants

- 2026–2027 **PI**, *Performance Analysis, Modeling and Scaling of HPC and AI Applications*, 500 CPU node-hours/year, 7.5K GPU node-hours/year.  
Perlmutter (NERSC), Director Reserve Allocation Award
- 2023–2026 **PI**, *Performance Analysis and Tuning of HPC and AI Applications*, 24K node-hours/year.  
Frontier (OLCF), Director's Discretion Project Award
- 2025–2026 **Co-PI**, *Training Multi-modal Models for HPC Code and Data*, 300K node-hours on Frontier (OLCF), 200K node-hours on Aurora (ALCF), 200K node-hours on Perlmutter (NERSC).  
DOE ASCR Leadership Computing Challenge (ALCC)
- 2022–2026 **PI**, *Performance Analysis, Modeling and Scaling of HPC Applications & Tools*, 1K CPU node-hours/year, 3K GPU node-hours/year.  
Perlmutter (NERSC), Director Reserve Allocation Award
- 2025–2026 **PI**, *Training a Large Scale Model for Code Retrieval*, 100K node-hours.  
Alps (CSCS), Large Production Project
- 2025–2026 **PI**, *Democratizing AI by Training Deployable Open-source Language Models*, 50K node-hours on Aurora (ALCF).  
DOE INCITE Program
- 2024 **PI**, *Democratizing AI via Highly Scalable Parallel Deep Learning Infrastructure*, 82K GPU node-hours.  
Perlmutter (NERSC), Director Reserve Allocation Award
- 2024 **PI**, *Democratizing AI by Training Deployable Open-source Language Models*, 600K node-hours on Frontier (OLCF), 100K node-hours on Polaris (ALCF).  
DOE INCITE Program
- 2020–2023 **PI**, *Performance Analysis and Tuning of HPC and AI Applications*, 20K node-hours/year.  
Summit (OLCF), Director's Discretion Project Award
- 2015–2022 **PI**, *Performance Analysis, Modeling and Scaling of HPC Applications & Tools*, 3M core-hours/year.  
Edison & Cori (NERSC), Director Reserve Allocation Award
- 2020–2021 **PI**, *Analyzing and Optimizing Parallel I/O and Performance Tools*, 8K node-hours.  
Summit (OLCF), Director's Discretion Project Award
- 2015–2016 **PI**, *Performance Analysis, Modeling and Scaling of HPC Applications & Tools*, 20.1M core-hours at ALCF, 9.3M core-hours at OLCF.  
Mira & Titan, ASCR Leadership Computing Challenge (ALCC).
- 2012–2015 **PI**, *Scalable Topology Aware Task Embedding (STATE)*, 17M core-hours/year.  
Cab & Vulcan (LLNL LC), Multiprogrammatic & Institutional Computing Initiative.
- 2014 **PI**, *Modeling Communication Behavior on Supercomputer Networks*, 30K core-hours.  
Blacklight (PSC), Extreme Science and Engineering Discovery Environment (XSEDE).
- 2013 **PI**, *Exploring energy efficiency, memory contention and application-tools-hardware co-design on a Xeon Phi cluster*, 10K node-hours.  
Beacon (NICS), The Beacon Project.
- 2008–2011 **Co-PI**, *Charm++ and its applications*, 1M core-hours/year.  
Intrepid (ANL ALCF), Director's Discretionary Allocation.

---

## Software Projects

Software developed by me ([github.com/bhatele](https://github.com/bhatele))

- CommProxies MPI proxy codes representing communication patterns in HPC applications
- congest Interconnect CONGestion ESTimators
- hatchet Python library for analyzing hierarchical performance data
- mapping Heuristics-based task mapping library
- mol3d Molecular Dynamics proxy application in Charm++
- NoiseFinder Program to quantify OS Noise on HPC clusters
- topomgr Interconnect Topology Management/Discovery Library

Software co-developed with students, and collaborators ([github.com/hpcgroup](https://github.com/hpcgroup))

- AriesNCL Aries Network Performance Counters Monitoring Library
- AxoNN Parallel framework for training deep neural networks
- BGQNCL Blue Gene/Q Network Performance Counters Monitoring Library
- Boxfish Tool for visualizing traffic on  $n$ -dimensional torus networks
- CallFlow Tool for visualizing calling context trees
- Chatterbug Suite of communication proxies for HPC applications
- Chizu Graph partitioning based task mapping library
- Damselfly Model-based parallel network simulator
- DragonView Tool for visualizing traffic on dragonfly networks
- Graphator Serial codes to generate communication patterns for parallel programs
- Kripke-Charm Charm++ version of Kripke, a 3D Sn deterministic particle transport code
- Loimos Parallel simulator for epidemic diffusion
- Loupe Profiler for MPI programs
- LULESH-Charm Charm++ version of LULESH, a shock hydrodynamics proxy application
- Ravel Tool for visualizing parallel execution traces in logical time
- Rubik Python library for task mapping of structured codes on  $n$ -dimensional torus
- ParEval Benchmark for evaluating the ability of LLMs to write parallel code
- ParEval-Repo Repository-level Parallel Code Evaluation Benchmark
- PCCL Performant Collective Communication Library
- pipit Python library for analyzing parallel execution traces
- Plexus Parallel framework for GNN training
- TraceR Trace Replay and network simulation framework
- TreeScope Tool for visualizing traffic on fat-tree networks
- YALIS Yet Another LLM Inference System

Other open-source software I have contributed to

- Spack Package manager for HPC platforms
- Charm++ Message-passing parallel language and runtime system
- NAMD Classical Molecular Dynamics simulation software in Charm++
- OpenAtom Car-Parrinello Molecular Dynamics simulation software in Charm++

---

## Teaching Experience

University of Maryland, College Park

- CMSC416/616** Introduction to Parallel Computing, *Fall 2020, Fall 2021, Fall 2022, Fall 2023, Spring 2024, Fall 2024, Fall 2025*
- CMSC828G** Systems for Machine Learning, *Spring 2025, Spring 2026*
- AMSC663/664** Advanced Scientific Computing I and II, *Fall 2022, Spring 2023*

- CMSC714** High Performance Computing Systems, *Fall 2019, Spring 2021*
- DIT Workshop** HPC Programming Bootcamp, *Winter 2020*
- [Tutorials at Conferences](#)
- AxoNN/YALIS** Distributed Deep Learning, *ISC '23, ISC '24, SC '24, ISC '25, SC '25, NAIRR Annual Meeting '26*
- Hatchet** User-centric Automated Performance Analysis, *SC '21*
- [Other Teaching](#)
- Spring 2013 **US Hindi Foundation (USHF), Palo Alto, CA**, Hindi Level I and II
- Summer 2009 **University of Illinois**, Computer Architecture I (CS231)\*
- Summer 2008 **University of Illinois**, Computer Architecture I (CS231)\*
- Summer 2005 **Indian Institute of Technology, Kanpur**, Data Structures and Algorithms\*
- \* Full responsibility for the course, prepared and gave lectures, awarded final grades

## Advising at UMD

### Current Students

- Ph.D. Students Onur Cankur, Joshua Hoke Davis, Emir Gencer, Lannie Dalton Hough, Joy Marie Kitson, Klaudiusz Rydzy, Prajwal Singhanian, Cunyang Wei
- M.S. Students Hoffmann M. Fokum, Mazin Karjekar, Hemanth Nandakumar
- Undergraduate Students Akshita Badkundri, Aman Chaturvedi, Noah Kasica, Ishan Khillan, Nirav Koley, Ashna Nayak, Franciszek Pajak, Ishan Revankar, Sarang S, Yash Sakharkar, Hayeong Shim, Mahua Singh, Alexander Sorescu, Akarsh Srivastava, Mehul Tahlani

### Graduated Students

- Ph.D. Students Daniel Lee Nichols, Siddharth Singh
- M.S. Students Rakrish Dhakal, Shua-Hua Lin, Nikodemos Koutsoheras, Alexander Movsesyan, Srivishnu Pyda, Aditya Ranjan, Zack Sating, Pranav Sivaraman, Sathwik Yanamaddi, Yiheng Xu
- Undergraduate Students Anish Bhupalam, Anshu Chandrasekar, Ian Costello, Rakrish Dhakal, Dilan Gunawardana, Michael Ilie, Tamar Kellner, Nikodemos Koutsoheras, Carter Lewis, Thomas Li, Shu-Huai Lin, Jordan Marry, Aidan Meyer, Isaac Minn, Rohan Mishra, Alexander Movsesyan, Hasan Muhammad, Sameer Mustaqali, Rishi Keshav Pradeep, Arjun Rajaram, Aditya Ranjan, Domenic Sangiovanni, Zack Sating, Omer Sharif, Pranav Sivaraman, Aditya Tomar, John Wendlandt, Zhaojun Xie, Sathwik Yanamaddi, William Yang, Jeff Zhang
- High School Students Joshua Piety, Aditya Ranjan, Alexander Sorescu, Aditya Tomar

## Mentoring/Co-advising at LLNL

- Postdocs Christopher Earl (LLNL), Giorgis Georgakoudis (LLNL), Tanzima Islam (LLNL), Nikhil Jain (LLNL), Aniruddha Marathe (LLNL), Jae-Seung Yeom (LLNL)
- Graduate Students Bilge Acun (Illinois), Israa Alqassem (Purdue), Ashwin Bhandare (UC Davis), Saptarshi Bhowmik (Florida State), Kevin Brown (Tokyo Tech), Ronak Buch (Illinois), Emilio Castillo (BSC), Vanessa Cedeno (Virginia Tech), Jaemin Choi (Illinois), Mehmet Deveci (Ohio State), Jens Domke (TU Dresden), Venmugil Elango (Ohio State), Alfredo Giménez (UC Davis), Suraj P. Kesavan (UC Davis), Katherine E. Isaacs (UC Davis), Harsh Khetawat (NC State), Konstantinos Koiliaris (Illinois), Harshitha Menon (Illinois), Yin Yee Ng (UC Davis), Huu Tan Nguyen (UC Davis), Samuel Pollard (UOregon), Michael P. Robson (Illinois), Karthik Senthil (Illinois), Aamer Shah (RWTH Aachen), Staci Smith (Arizona), Edgar Solomonik (UC Berkeley), Ankit Srivastava (Georgia Tech), Felix Wang (Illinois), Samuel T. White (Illinois), Chunxing Yin (Georgia Tech)
- Undergraduate Students Matthew R. Kotila (UC Davis), Aditya Nigam (IIT Kanpur), Alexander Robey (Swarthmore), Kristina Sep (UC Berkeley), Palash Sharma (IIT Kanpur), Andrew R. Titus (MIT), Dylan Wang (UC Davis)

---

## University and other Academic Service

Dissertation Committee Saptarshi Bhowmik (FSU), Jaemin Choi (Illinois), Tom Cornebize (Univ. of Grenoble), Victoria Cepeda Espinoza (UMD), MG Hirsch (UMD), Yehuda Katz (UMD), Thomas Rolinger (UMD), Swati Singhal (UMD), Staci Smith (Arizona), Candace Walden (UMD)

- 2024-date AIM Compute Committee
- 2024-2027 University Senator for Computer Science
- 2021-date Faculty Advisory Committee for National Quantum Lab
- 2021-date UMD DIT Research Technology Working Group
- 2025-2027, 2022-2023 CS Department Council
- 2025-2026 CS TTK Search Committee
- 2025-2026, 2022-2023, 2020-2021 UMIACS Appointment, Promotion, and Tenure (APT) Committee
- 2022-2026 CS Merit Pay Committee
- 2024-2025 AIM TTK Search Committee
- 2023-2024 Chair, Ad-hoc Committee for BPC Plan Review & Re-approval
- 2023-2024 Member, UMD IT Council
- 2023-2024 Chair, UMD DIT Research Technology Working Group
- 2023-2024 UMIACS Steering Committee
- 2020-2024 CS Diversity and Inclusion Committee
- 2023 CS Graduate Coordinator Search Committee
- 2022-2023, 2020-2021 D&I representative, CS TTK Search Committee
- 2021-2022 CS Department Review Diversity & inclusion, outreach, and communications sub-committee
- 2019-2022 Iribe Building Committee
- 2020-2021 CS/UMIACS Student Seating Committee
- 2020-2021 CS Field Committee Chair, Computer Systems

---

## Professional Service

- 2026 Program Committee Track Chair, Scalable AI and Data Analytics Track, Euro-Par '26
- 2025-present Member, National Deep Inference Fabric (NDIF) Advisory Board
- 2025 Tutorials Chair, SC '25
- 2020-2024 Steering Committee, IEEE Cluster Conference
- 2024 Steering Committee, International Workshop on Large Language Models and HPC (LLM x HPC)
- 2024 Finance Liaison for Technical Program, SC '24
- 2024 Student Program Co-chair, IEEE Cluster '24
- 2023 Student Funding Liaison, IEEE Cluster '23
- 2021-2023 Steering Committee, ISC High Performance Conference
- 2023 Research Papers Chair, ISC, Hamburg, Germany
- 2022 General Chair, IEEE Cluster, Heidelberg, Germany
- 2022 Research Papers Deputy Chair, ISC, Hamburg, Germany
- 2021 Program Vice-Chair, Applications Track, HiPC '21
- 2021 Virtual Logistics Liaison for BOFs, SC '21
- 2021 Program Committee Track Chair, Performance Track, ICPP '21
- 2021 PhD Forum Deputy Chair, ISC, Frankfurt, Germany

2019–2021 Workshop Co-Chair, Workshop on Programming and Performance Visualization Tools (ProTools)  
 Member ACM, ACM SIGHPC, IEEE, IEEE Computer Society, IEEE TCHPC

2018, 2019 Co-organizer, Computation Summer Hackathon, LLNL

2019 Workshops and Tutorials Chair, PPOPP'19, Washington, DC

2017, 2018, 2019 Mentor for Supercomputing Mentor/Protégé Program

2018 Member, DOE CORAL2 Technical Evaluation Team (TET)

2017–2018 Chair, Tri-lab Support Team (TST), PSAAP-II Center of Excellence, University of Florida

2017, 2018 Workshop Chair, International Workshop on Visual Performance Analytics (VPA)

2016, 2018 Organizer, Mini-symposium at SIAM PP'16 (Paris, France), SIAM PP'18 (Tokyo, Japan)

2017 Posters Chair, IEEE Cluster Conference, Honolulu, HI

2017 Organizer, Mini-symposium at SIAM CSE'17, Atlanta, GA

2009, 2012, 2014 Mentor for Supercomputing Mentor/Protégé Program

2014 Program Committee Area Chair, Applications and Algorithms Track, SBAC-PAD '14

2012 Postdoc Representative, Institutional Postdoc Program Board (IPPB), LLNL

2011–2012 Member and Web Developer, Lawrence Livermore Postdoc Association (LLPA) Council

2010–2011 Member and Web Developer, Society of Postdoctoral Scholars, Univ. of Illinois

2010 CS Grad Ambassador, Dept. of Computer Science, Univ. of Illinois

2007–2011 Helped with organization of Charm++ Workshops

2009, 2010 Facilitator, Graduate Academy for College Teaching, Univ. of Illinois

2009, 2010 Mentor for the WCS Mentoring Program, Dept. of CS, Univ. of Illinois

2008, 2009 Student Volunteer, Supercomputing, Austin, TX and Portland, OR

## Peer Review Committees

Editorial Board Associate Editor, IEEE Transactions on Parallel and Distributed Systems, 2021–2024

Tech. Program Conferences: ICISTM '12, IEEE BigData (2013, 2014), IEEE Cluster (2015, 2016), Euro-Par '16,  
 Committees HotI '17, SBAC-PAD (2013–2017), HotI '19, CCGrid (2012, 2013, 2018–2020), ISC '20, HiPC (2017, 2018, 2020), VISSOFT (2020, 2021), ICPP '23, IPDPS (2014, 2015, 2017, 2019–2024), HPDC (2024–2025), ICS (2023, 2026), SC (2018, 2021, 2022, 2024, 2026)

Workshops: WHIST '12, ESCAPE (2011–2013), WRAP '15, COMHPC '16, HPCMASPA (2016–2019), PMBS (2017, 2018), SNACS (2019, 2021), WAMTA '23

Grant Review 2025, 2022, 2014 DOE/ASCR Review Panel  
 Panels 2026, 2023, 2020 NSF Review Panel

2018 DOE Predictive Science Academic Alliance Program (PSAAP) III

2016–2018 LLNL Institutional Computing Grand Challenge Awards

2014 DOE Small Business Innovation Research (SBIR)

Other 2026, 2024 SC Posters Committee  
 Committees 2023, 2022, 2021 SC Tutorials Committee

2020 CCGrid '20 Scalable Computing Challenge (SCALE) Committee

2016 SC'16 Early Career Program Committee

2014 SC'14 Workshops Committee

2012, 2013 Selection Committee, ACM/IEEE-CS George Michael Memorial HPC Fellowship


Technical Conferences: CHI '08, ICPP '09, PPOPP '11, PACT '12, SBAC-PAD '12, Euro-Par '13, ICS '14 (ERC)  
 Reviewer Journals: IEEE TPDS, ToC, IHPCA, JPDC, CPE, PARCO, FGCS, IJCS, JOSS








## Peer Reviewed Publications

### Journal Publications

- [1] Daniel Nichols, Pranav Polasam, Harshitha Menon, Aniruddha Marathe, Todd Gamblin, and **Abhinav Bhatele**. Performance-aligned LLMs for generating fast HPC code. *IEEE Transactions on Parallel & Distributed Systems*, 2026 (to appear).

- [2] Joy Kitson, Prescott C. Alexander, Joseph Tuccillo, David J. Butts, Christa Brelsford, **Abhinav Bhatele**, Sara Y. Del Valle, and Timothy C. Germann. Simulating nationwide coupled disease and fear spread in an agent-based model. *Nature Scientific Reports*, November 2025. <https://doi.org/10.1038/s41598-025-26425-y>.
- [3] Kohei Yoshida, Ryuichi Sakamoto, Kento Sato, **Abhinav Bhatele**, Hayato Yamaki, Hiroki Honda, and Shinobu Miwa. VAHRM: Variation-aware resource management in heterogeneous supercomputing systems. *IEEE Transactions on Parallel & Distributed Systems*, 36(08):1713–1727, August 2025. <https://doi.ieeecomputersociety.org/10.1109/TPDS.2025.3577252>.
- [4] Connor Scully-Allison, Ian Lumsden, Katy Williams, Jesse Bartels, Michela Taufer, Stephanie Brink, **Abhinav Bhatele**, Olga Pearce, and Katherine E. Isaacs. Design concerns for integrated scripting and interactive visualization in notebook environments. *IEEE Transactions on Visualization and Computer Graphics*, 30(9):6572–6585, September 2024. <https://doi.ieeecomputersociety.org/10.1109/TVCG.2024.3354561>.
- [5] Suraj P. Kesavan, Harsh Bhatia, **Abhinav Bhatele**, Stephanie Brink, Olga Pearce, Todd Gamblin, Peer-Timo Bremer, and Kwan-Liu Ma. Scalable comparative visualization of ensembles of call graphs. *IEEE Transactions on Visualization and Computer Graphics*, 29(3):1691–1704, March 2023. LLNL-JRNL-809459. <https://doi.ieeecomputersociety.org/10.1109/TVCG.2021.3129414>.
- [6] Huu Tan Nguyen, **Abhinav Bhatele**, Nikhil Jain, Suraj Kesavan, Harsh Bhatia, Todd Gamblin, Kwan-Liu Ma, and Peer-Timo Bremer. Visualizing hierarchical performance profiles of parallel codes using CallFlow. *IEEE Transactions on Visualization and Computer Graphics*, 27(4):2455–2468, April 2021. LLNL-JRNL-797378. <https://doi.ieeecomputersociety.org/10.1109/TVCG.2019.2953746>.
- [7] Alfredo Giménez, Todd Gamblin, Ilir Jusufi, **Abhinav Bhatele**, Martin Schulz, Peer-Timo Bremer, and Bernd Hamann. MemAxes: Visualization and analytics for characterizing complex memory performance behaviors. *IEEE Transactions on Visualization and Computer Graphics*, 24(7):2180–2193, July 2018. LLNL-JRNL-. <http://doi.ieeecomputersociety.org/10.1109/TVCG.2017.2718532>.
- [8] Harsh Bhatia, Nikhil Jain, **Abhinav Bhatele**, Yarden Livnat, Jens Domke, Valerio Pascucci, and Peer-Timo Bremer. Interactive investigation of traffic congestion on fat-tree networks using TreeScope. *Computer Graphics Forum*, 37(3):561–572, June 2018. LLNL-JRNL-. <https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13442>.
- [9] Erik W. Draeger, Xavier Andrade, John A. Gunnels, **Abhinav Bhatele**, Andre Schleife, and Alfredo A. Correa. Massively parallel first-principles simulation of electron dynamics in materials. *Journal of Parallel and Distributed Computing*, 106:205–214, February 2017. <http://www.sciencedirect.com/science/article/pii/S0743731517300734>.
- [10] Katherine E. Isaacs, Todd Gamblin, **Abhinav Bhatele**, Martin Schulz, Bernd Hamann, and Peer-Timo Bremer. Ordering traces logically to identify lateness in message-passing programs. *IEEE Transactions on Parallel and Distributed Systems*, 27(3):829–840, March 2016. LLNL-JRNL-668754. <http://doi.ieeecomputersociety.org/10.1109/TPDS.2015.2417531>.
- [11] Katherine E. Isaacs, Peer-Timo Bremer, Ilir Jusufi, Todd Gamblin, **Abhinav Bhatele**, Martin Schulz, and Bernd Hamann. Combing the communication hairball: Visualizing parallel execution traces using logical time. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):2349–2358, December 2014. LLNL-JRNL-657418. <http://doi.ieeecomputersociety.org/10.1109/TVCG.2014.2346456>.
- [12] Steven Langer, **Abhinav Bhatele**, and Charles H. Still. pF3D simulations of laser-plasma interactions in National Ignition Facility experiments. *Computing in Science and Engineering*, 16(6):42–50, November 2014. LLNL-JRNL-648736. <http://doi.ieeecomputersociety.org/10.1109/MCSE.2014.79>.
- [13] Aaditya G. Landge, Joshua A. Levine, Katherine E. Isaacs, **Abhinav Bhatele**, Todd Gamblin, Martin Schulz, Steve H. Langer, Peer-Timo Bremer, and Valerio Pascucci. Visualizing network traffic to understand the performance of massively parallel simulations. *IEEE Transactions on*

*Visualization and Computer Graphics*, 18(12):2467–2476, December 2012. LLNL-CONF-543359.  <http://doi.ieeecomputersociety.org/10.1109/TVCG.2012.286>.

- [14] Gengbin Zheng, **Abhinav Bhatele**, Esteban Meneses, and Laxmikant V. Kalé. Periodic hierarchical load balancing for large supercomputers. *Int. J. High Perform. Comput. Appl.*, 25(4):371–385, November 2011.  <http://hpc.sagepub.com/content/25/4/371>.
- [15] **Abhinav Bhatele**, Eric Bohm, and Laxmikant V. Kale. Optimizing communication for Charm++ applications by reducing network contention. *Concurrency and Computation: Practice and Experience*, 23(2):211–222, February 2011.  <https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.1637>.
- [16] **Abhinav Bhatele**, Lukasz Wesolowski, Eric Bohm, Edgar Solomonik, and Laxmikant V. Kalé. Understanding application performance via micro-benchmarks on three large supercomputers: Intrepid, Ranger and Jaguar. *Int. J. High Perform. Comput. Appl.*, 24(4):411–427, November 2010.  <http://hpc.sagepub.com/content/24/4/411>.
- [17] **Abhinav Bhatele** and Laxmikant V. Kale. Quantifying network contention on large parallel machines. *Parallel Processing Letters*, 19(04):553–572, December 2009.  <http://www.worldscientific.com/doi/abs/10.1142/S0129626409000419>.
- [18] **Abhinav Bhatele** and Laxmikant V. Kale. Benefits of topology aware mapping for mesh interconnects. *Parallel Processing Letters*, 18(04):549–566, December 2008.  <http://www.worldscientific.com/doi/abs/10.1142/S0129626408003569>.
- [19] Eric Bohm, **Abhinav Bhatele**, Laxmikant V. Kalé, Mark E. Tuckerman, Sameer Kumar, John A. Gunnels, and Glenn J. Martyna. Fine-grained parallelization of the Car-Parrinello ab initio molecular dynamics method on the IBM Blue Gene/L supercomputer. *IBM J. Res. Dev.*, 52(1/2):159–175, January 2008.  <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=5388634>.
- [20] Sameer Kumar, Chao Huang, Gengbin Zheng, Eric Bohm, **Abhinav Bhatele**, James C. Phillips, Hao Yu, and Laxmikant V. Kalé. Scalable molecular dynamics with NAMD on the IBM Blue Gene/L system. *IBM J. Res. Dev.*, 52(1/2):177–188, January 2008.  <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=5388644>.

## Conference Publications

- [21] Cunyang Wei and **Abhinav Bhatele**. Scalable all-to-all algorithms for dynamic and irregular communication patterns. In *Proceedings of the International Conference on Supercomputing, ICS '26*, June 2026 (to appear).
- [22] Onur Cankur, Brian Austin, Dhruva Kulkarni, and **Abhinav Bhatele**. Characterizing production GPU workloads using system-wide telemetry data. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '26*. IEEE Computer Society, May 2026 (to appear).
- [23] Siddharth Singh, Keshav Pradeep, Mahua Singh, Cunyang Wei, and **Abhinav Bhatele**. The big send-off: Scalable and performant collectives for deep learning. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '26*. IEEE Computer Society, May 2026 (to appear).
- [24] Cunyang Wei, Keshav Pradeep, and **Abhinav Bhatele**. The case of the elusive application performance on production GPU supercomputers. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '26*. IEEE Computer Society, May 2026 (to appear).
- [25] Jonas Geiping, Sean Michael McLeish, Neel Jain, John Kirchenbauer, Siddharth Singh, Brian R. Bartoldson, Bhavya Kailkhura, **Abhinav Bhatele**, and Tom Goldstein. Scaling up test-time compute with latent reasoning: A recurrent depth approach. In *Advances in Neural Information Processing Systems (spotlight)*, volume 38 of *NeurIPS '25 (spotlight)*. Curran Associates, Inc., December 2025.
- [26] Sean Michael McLeish, John Kirchenbauer, David Yu Miller, Siddharth Singh, **Abhinav Bhatele**, Micah Goldblum, Ashwinee Panda, and Tom Goldstein. Gemstones: A model suite for scaling laws.

In *Advances in Neural Information Processing Systems*, volume 38 of *NeurIPS '25*. Curran Associates, Inc., December 2025.

- [27] Aditya K. Ranjan, Siddharth Singh, Cunyang Wei, and **Abhinav Bhatele**. Plexus: Taming billion-edge graphs with 3D parallel full-graph GNN training. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '25*. ACM, November 2025. <https://doi.acm.org/10.1145/3712285.3759890>.
- [28] Yue Kaiyu, Vasu Singla, Menglin Jia, John Kirchenbauer, Rifaa Qadri, Zikui Cai, **Abhinav Bhatele**, Furong Huang, and Tom Goldstein. Zero-shot vision encoder grafting via LLM surrogates. In *International Conference on Computer Vision (ICCV)*, October 2025. [https://openaccess.thecvf.com/content/ICCV2025/html/Yue\\_Zero-Shot\\_Vision\\_Encoder\\_Grafting\\_via\\_LLM\\_Surrogates\\_ICCV\\_2025\\_paper.html](https://openaccess.thecvf.com/content/ICCV2025/html/Yue_Zero-Shot_Vision_Encoder_Grafting_via_LLM_Surrogates_ICCV_2025_paper.html).
- [29] Joshua H. Davis, Daniel Nichols, Ishan Khillan, and **Abhinav Bhatele**. ParEval-Repo: A benchmark suite for evaluating LLMs with repository-level HPC translation tasks. In *Proceedings of the 54th International Conference on Parallel Processing, ICPP '25*, September 2025. <http://doi.acm.org/10.1145/3754598.3754669>.
- [30] Lannie Dalton Hough and **Abhinav Bhatele**. Elias-Fano compression for space-efficient rank and select structures. In *Proceedings of the 23rd Symposium on Experimental Algorithms, SEA '25*, July 2025. <https://doi.org/10.4230/LIPIcs.SEA.2025.23>.
- [31] Aman Chaturvedi, Daniel Nichols, Siddharth Singh, and **Abhinav Bhatele**. HPC-Coder-v2: Studying code llms across low-resource parallel languages. In *Proceedings of the ISC High Performance Conference, ISC '25*, June 2025.
- [32] Joshua H. Davis, Pranav Sivaraman, Joy Kitson, Konstantinos Parasyris, Harshitha Menon, Isaac Minn, Giorgis Georgakoudis, and **Abhinav Bhatele**. Taking gpu programming models to task for performance portability. In *Proceedings of the International Conference on Supercomputing, ICS '25*, June 2025. <http://doi.acm.org/10.1145/3721145.3730423>.
- [33] Joy Kitson, Ian Costello, Jiangzhuo Chen, Diego Jiménez, Stefan Hoops, Henning Mortveit, Esteban Meneses, Jae-Seung Yeom, Madhav V. Marathe, and **Abhinav Bhatele**. Pandemics *in silico*: Scaling an agent-based simulation on realistic social contact networks. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '25*. IEEE Computer Society, June 2025. <https://doi.ieeecomputersociety.org/10.1109/IPDPS64566.2025.00050>.
- [34] Aakash Raj Dhakal, Tanzima Z. Islam, Arunavo Dey, Tapasya Patki, Daniel Nichols, **Abhinav Bhatele**, and Jae-Seung Yeom. xAMM: “Attention” to details improves cross-platform prediction accuracy. In *Proceedings of the IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGrid '25*. IEEE Computer Society, May 2025. <https://doi.ieeecomputersociety.org/10.1109/CCGRID64434.2025.00067>.
- [35] Abhimanyu Hans, John Kirchenbauer, Yuxin Wen, Neel Jain, Hamid Kazemi, Prajwal Singhanian, Siddharth Singh, Gowthami Somepalli, Jonas Geiping, **Abhinav Bhatele**, Avi Schwarzschild, and Tom Goldstein. Be like a goldfish, don't memorize! Mitigating memorization in generative LLMs. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37 of *NeurIPS '24*, pages 24022–24045. Curran Associates, Inc., December 2024. <https://doi.org/10.52202/079017-0757>.
- [36] Sean Michael McLeish, Arpit Bansal, Alex Stein, Neel Jain, John Kirchenbauer, Brian R. Bartoldson, Bhavya Kailkhura, **Abhinav Bhatele**, Jonas Geiping, Avi Schwarzschild, and Tom Goldstein. Transformers can do arithmetic with the right embeddings. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37 of *NeurIPS '24*, pages 108012–108041. Curran Associates, Inc., December 2024. <https://doi.org/10.52202/079017-3430>.
- [37] Prajwal Singhanian, Siddharth Singh, Shwai He, Soheil Feizi, and **Abhinav Bhatele**. Loki: Low-rank keys for efficient sparse attention. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet,

- J. Tomczak, and C. Zhang, editors, *Advances in Neural Information Processing Systems*, volume 37 of *NeurIPS '24*, pages 16692–16723. Curran Associates, Inc., December 2024. <https://doi.org/10.52202/079017-0532>.
- [38] Yiheng Xu, Hariharan Devarajan, Kathryn Mohror, and **Abhinav Bhatele**. ML-based modeling to predict I/O performance on different storage sub-systems. In *Proceedings of the IEEE International Conference on High Performance Computing, HiPC '24*, December 2024. <https://doi.ieeecomputersociety.org/10.1109/HiPC62374.2024.00030>.
- [39] Daniel Nichols, Harshitha Menon, Todd Gamblin, and **Abhinav Bhatele**. A probabilistic approach to selecting build configurations in package managers. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '24*. IEEE Computer Society, November 2024. <https://doi.ieeecomputersociety.org/10.1109/SC41406.2024.00090>.
- [40] Siddharth Singh, Prajwal Singhanian, Aditya Ranjan, John Kirchenbauer, Jonas Geiping, Yuxin Wen, Neel Jain, Abhimanyu Hans, Manli Shu, Aditya Tomar, Tom Goldstein, and **Abhinav Bhatele**. Democratizing AI: Open-source scalable LLM training on GPU-based supercomputers. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '24*. IEEE Computer Society, November 2024. <https://doi.ieeecomputersociety.org/10.1109/SC41406.2024.00010>.
- [41] Daniel Nichols, Joshua H. Davis, Zhaojun Xie, Arjun Rajaram, and **Abhinav Bhatele**. Can large language models write parallel code? In *Proceedings of the 33rd International Symposium on High-Performance Parallel and Distributed Computing, HPDC '24*, June 2024. <http://doi.acm.org/10.1145/3625549.3658689>.
- [42] Daniel Nichols, Aniruddha Marathe, Harshitha Menon, Todd Gamblin, and **Abhinav Bhatele**. HPC-Coder: Modeling parallel programs using large language models. In *Proceedings of the ISC High Performance Conference, ISC '24*, May 2024. <https://doi.org/10.23919/ISC.2024.10528929>.
- [43] Daniel Nichols, Alexander Mosesyan, Jae-Seung Yeom, Abhik Sarkar, Daniel Milroy, Tapasya Patki, and **Abhinav Bhatele**. Predicting cross-architecture performance of parallel programs. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '24*. IEEE Computer Society, May 2024. <https://doi.ieeecomputersociety.org/10.1109/IPDPS57955.2024.00057>.
- [44] Harshitha Menon, Daniel Nichols, **Abhinav Bhatele**, and Todd Gamblin. Learning to predict and improve build successes in package ecosystems. In *Proceedings of the 21st International Conference on Mining Software Repositories, MSR '24*. Association for Computing Machinery, April 2024. <http://doi.acm.org/10.1145/3643991.3644927>.
- [45] Harsh Khetawat, Nikhil Jain, **Abhinav Bhatele**, and Frank Mueller. Predicting GPUDirect benefits for HPC workloads. In *Proceedings of the 32nd Euromicro International Conference on Parallel, Distributed, and Network-Based Processing, PDP '24*. IEEE Computer Society, March 2024. <https://doi.ieeecomputersociety.org/10.1109/PDP62718.2024.00020>.
- [46] Siddharth Singh, Olatunji Ruwase, Ammar Ahmad Awan, Samyam Rajbhandari, Yuxiong He, and **Abhinav Bhatele**. A hybrid tensor-expert-data parallelism approach to optimize mixture-of-experts training. In *Proceedings of the International Conference on Supercomputing, ICS '23*, June 2023. <http://doi.acm.org/10.1145/3577193.3593704>.
- [47] Joshua Hoke Davis, Justin Shafner, Daniel Nichols, Nathan Grube, Pino Martin, and **Abhinav Bhatele**. Porting a computational fluid dynamics code with AMR to large-scale GPU platforms. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '23*. IEEE Computer Society, May 2023. <https://doi.ieeecomputersociety.org/10.1109/IPDPS54959.2023.00066>.
- [48] Siddharth Singh and **Abhinav Bhatele**. Exploiting sparsity in pruned neural networks to optimize large model training. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium,*

- IPDPS '23. IEEE Computer Society, May 2023. <https://doi.ieeecomputersociety.org/10.1109/IPDPS54959.2023.00033>.
- [49] Onur Cankur and **Abhinav Bhatele**. Comparative evaluation of call graph generation by profiling tools. In *Proceedings of the ISC High Performance Conference*, ISC '22, pages 213–232. Springer International Publishing, May 2022. [https://link.springer.com/chapter/10.1007/978-3-031-07312-0\\_11](https://link.springer.com/chapter/10.1007/978-3-031-07312-0_11).
- [50] Daniel Nichols, Aniruddha Marathe, Kathleen Shoga, Todd Gamblin, and **Abhinav Bhatele**. Resource utilization aware job scheduling to mitigate performance variability. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, IPDPS '22. IEEE Computer Society, May 2022. <https://doi.ieeecomputersociety.org/10.1109/IPDPS53621.2022.00040>.
- [51] Siddharth Singh and **Abhinav Bhatele**. AxoNN: An asynchronous, message-driven parallel framework for extreme-scale deep learning. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, IPDPS '22. IEEE Computer Society, May 2022. <https://doi.ieeecomputersociety.org/10.1109/IPDPS53621.2022.00065>.
- [52] Saptarshi Bhowmik, Nikhil Jain, Xin Yuan, and **Abhinav Bhatele**. A simulation study of hardware parameters for future GPU-based HPC platforms. In *Proceedings of the IEEE International Performance, Computing, and Communications Conference*, IPCCC '21. IEEE Computer Society, October 2021. <https://doi.ieeecomputersociety.org/10.1109/IPCCC51483.2021.9679359>.
- [53] Sascha Hunold, **Abhinav Bhatele**, George Bosilca, and Peter Knees. Predicting MPI collective communication performance using machine learning. In *Proceedings of the IEEE Cluster Conference*, Cluster '20, September 2020. <https://doi.ieeecomputersociety.org/10.1109/CLUSTER49012.2020.00036>.
- [54] Jaemin Choi, David Richards, Laxmikant V. Kale, and **Abhinav Bhatele**. End-to-end performance modeling of distributed GPU applications. In *Proceedings of the International Conference on Supercomputing*, ICS '20, June 2020. LLNL-CONF-809401. <http://doi.acm.org/10.1145/3392717.3392737>.
- [55] **Abhinav Bhatele**, Jayaraman J. Thiagarajan, Taylor Groves, Rushil Anirudh, Staci A. Smith, Brandon Cook, and David K. Lowenthal. The case of performance variability on dragonfly-based systems. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, IPDPS '20. IEEE Computer Society, May 2020. LLNL-CONF-772401. <https://doi.ieeecomputersociety.org/10.1109/IPDPS47924.2020.00096>.
- [56] Harshitha Menon, **Abhinav Bhatele**, and Todd Gamblin. Auto-tuning parameter choices using bayesian optimization. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, IPDPS '20. IEEE Computer Society, May 2020. LLNL-CONF-772119. <https://doi.ieeecomputersociety.org/10.1109/IPDPS47924.2020.00090>.
- [57] Giorgis Georgakoudis, Nikhil Jain, Takatsugu Ono, Koji Inoue, Shinobu Miwa, and **Abhinav Bhatele**. Evaluating the impact of energy efficient networks on hpc workloads. In *Proceedings of the IEEE International Conference on High Performance Computing*, HiPC '19, December 2019. LLNL-CONF-791976. <https://doi.ieeecomputersociety.org/10.1109/HiPC.2019.00044>.
- [58] **Abhinav Bhatele**, Stephanie Brink, and Todd Gamblin. Hatchet: Pruning the overgrowth in parallel profiles. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '19. ACM, November 2019. LLNL-CONF-772402. <http://doi.acm.org/10.1145/3295500.3356219>.
- [59] **Abhinav Bhatele**, Nikhil Jain, Misbah Mubarak, and Todd Gamblin. Analyzing cost-performance tradeoffs of HPC network designs under different constraints using simulations. In *Proceedings of the ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, PADS '19. IEEE Computer Society, June 2019. LLNL-CONF-772399. <http://doi.acm.org/10.1145/3316480.3325516>.
- [60] Emilio Castillo, Nikhil Jain, Marc Casas, Miquel Moreto, Martin Schulz, Ramon Bievide, Mateo Valero, and **Abhinav Bhatele**. Optimizing computation-communication overlap in asynchronous

- task-based programs. In *Proceedings of the International Conference on Supercomputing*, ICS '19, June 2019. LLNL-CONF-772400. <http://doi.acm.org/10.1145/3293883.3295720>.
- [61] Samuel A. Pollard, Nikhil Jain, Stephen Herbein, and **Abhinav Bhatele**. Evaluation of an interference-free node allocation policy on fat-tree clusters. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '18. IEEE Computer Society, November 2018. LLNL-CONF-745526. <https://doi.ieeecomputersociety.org/10.1109/SC.2018.00029>.
- [62] Staci A. Smith, Clara Cromey, David K. Lowenthal, Jens Domke, Nikhil Jain, Jayaraman J. Thiagarajan, and **Abhinav Bhatele**. Mitigating inter-job interference using adaptive flow-aware routing. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '18. IEEE Computer Society, November 2018. LLNL-CONF-745538. <https://doi.ieeecomputersociety.org/10.1109/SC.2018.00030>.
- [63] Kevin A. Brown, Nikhil Jain, Satoshi Matsuoka, Martin Schulz, and **Abhinav Bhatele**. Interference between I/O and MPI traffic on fat-tree networks. In *Proceedings of the International Conference on Parallel Processing*, ICPP '18, August 2018. LLNL-CONF-751958. <http://doi.acm.org/10.1145/3225058.3225144>.
- [64] Jayaraman J. Thiagarajan, Nikhil Jain, Rushil Anirudh, Alfredo Giménez, Rahul Sridhar, Aniruddha Marathe, Tao Wang, Murali Emani, **Abhinav Bhatele**, and Todd Gamblin. Bootstrapping parameter space exploration for fast tuning. In *Proceedings of the International Conference on Supercomputing*, ICS '18, June 2018. LLNL-CONF-750296. <http://doi.acm.org/10.1145/3205289.3205321>.
- [65] Jayaraman J. 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. IEEE Computer Society, May 2018. LLNL-CONF-740303. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2018.00088>.
- [66] Misbah Mubarak, Nikhil Jain, Jens Domke, Noah Wolfe, Caitlin Ross, Jianping Li, **Abhinav Bhatele**, Christopher D. Carothers, Kwan-Liu Ma, and Robert B. Ross. Toward reliable validation of HPC interconnect simulations. In *Proceedings of the Winter Simulation Conference*, WSC '17, December 2017. LLNL-CONF-733848. <https://dl.acm.org/citation.cfm?id=3242231>.
- [67] Alfredo Giménez, Todd Gamblin, **Abhinav Bhatele**, Chad Wood, Kathleen Shoga, Aniruddha Marathe, Peer-Timo Bremer, Bernd Hamann, and Martin Schulz. ScrubJay: Deriving knowledge from the disarray of hpc performance data. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '17. ACM, November 2017. LLNL-CONF-735962. <http://doi.acm.org/10.1145/3126908.3126935>.
- [68] Nikhil Jain, **Abhinav Bhatele**, Louis Howell, David Böhme, 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 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '17. ACM, November 2017. LLNL-CONF-736289. <http://doi.acm.org/10.1145/3126908.3126967>.
- [69] 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 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '17, November 2017. LLNL-CONF-736726. <http://doi.acm.org/10.1145/3126908.3126969>.
- [70] Misbah Mubarak, Philip Carns, Jonathan Jenkins, Jianping Li, Nikhil Jain, Shane Snyder, Robert B. 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 *Proceedings of the IEEE Cluster Conference*, Cluster '17, September 2017. LLNL-CONF-731482. <http://doi.ieeecomputersociety.org/10.1109/CLUSTER.2017.25>.

- [71] **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 IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGrid '17 SCALE Challenge*. IEEE Computer Society, May 2017. LLNL-CONF-690723. <http://doi.ieeecomputersociety.org/10.1109/CCGRID.2017.141>.
- [72] 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. LLNL-CONF-706801. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2017.91>.
- [73] 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. In *Proceedings of the IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, CCGrid '17*. IEEE Computer Society, May 2017. LLNL-CONF-713054. <http://doi.ieeecomputersociety.org/10.1109/CCGRID.2017.102>.
- [74] Tanzima Z. Islam, Jayaraman J. Thiagarajan, **Abhinav Bhatele**, Martin Schulz, and Todd Gamblin. A machine learning framework for performance coverage analysis of proxy applications. 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-696018. <http://doi.ieeecomputersociety.org/10.1109/SC.2016.45>.
- [75] 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. <http://doi.ieeecomputersociety.org/10.1109/SC.2016.13>.
- [76] Edgar A. Leon, Ian Karlin, **Abhinav Bhatele**, Steven H. Langer, Chris Chembreau, Louis H. Howell, Trent D'Hooge, and Matthew L. Leininger. Characterizing parallel scientific applications on commodity clusters: An empirical study of a tapered fat-tree. 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-681011. <http://doi.ieeecomputersociety.org/10.1109/SC.2016.77>.
- [77] **Abhinav Bhatele**, Nikhil Jain, Yarden Livnat, Valerio Pascucci, and Peer-Timo Bremer. Analyzing network health and congestion in dragonfly-based supercomputers. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '16*. IEEE Computer Society, May 2016. LLNL-CONF-678293. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2016.123>.
- [78] Erik Draeger, Xavier Andrade, John Gunnels, **Abhinav Bhatele**, Andre Schleife, and Alfredo Correa. Massively parallel first-principles simulation of electron dynamics in materials. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '16*. IEEE Computer Society, May 2016. LLNL-CONF-669964. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2016.46>.
- [79] Katherine E. Isaacs, **Abhinav Bhatele**, Jonathan Lifflander, David Böhme, Todd Gamblin, Martin Schulz, Bernd Hamann, and Peer-Timo Bremer. Recovering logical structure from Charm++ event traces. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '15*. ACM, November 2015. LLNL-CONF-670046. <http://doi.acm.org/10.1145/2807591.2807634>.
- [80] **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. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2015.92>.

- [81] 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. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2015.102>.
- [82] **Abhinav Bhatele**, Nikhil Jain, Katherine E. Isaacs, Ronak Buch, Todd Gamblin, Steven H. Langer, and Laxmikant V. Kale. Optimizing the performance of parallel applications on a 5D torus via task mapping. In *Proceedings of the IEEE International Conference on High Performance Computing, HiPC '14*. IEEE Computer Society, December 2014. LLNL-CONF-655465. <http://doi.ieeecomputersociety.org/10.1109/HiPC.2014.7116706>.
- [83] Ahmed Abdel-Gawad, Mithuna Thottethodi, and **Abhinav Bhatele**. RAHTM: Routing-algorithm aware hierarchical task mapping. 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-653568. <http://doi.ieeecomputersociety.org/10.1109/SC.2014.32>.
- [84] Alfredo Giménez, Todd Gamblin, Barry Rountree, **Abhinav Bhatele**, Ilir Jusufi, Peer-Timo Bremer, and Bernd Hamann. Dissecting on-node memory access performance: A semantic approach. 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-658626. <http://doi.ieeecomputersociety.org/10.1109/SC.2014.19>.
- [85] 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. <http://doi.ieeecomputersociety.org/10.1109/SC.2014.33>.
- [86] Katherine E. Isaacs, Alfredo Giménez, Ilir Jusufi, Todd Gamblin, **Abhinav Bhatele**, Martin Schulz, Bernd Hamann, and Peer-Timo Bremer. State of the art of performance visualization. In R. Borgo, R. Maciejewski, and I. Viola, editors, *Proceedings of the Eurographics Conference of Visualization (STARs), EuroVis '14*. The Eurographics Association, June 2014. LLNL-CONF-652873. <http://dx.doi.org/10.2312/eurovisstar.20141177>.
- [87] Jae-Seung Yeom, **Abhinav Bhatele**, Keith R. Bisset, Eric Bohm, Abhishek Gupta, Laxmikant V. Kale, Madhav Marathe, Dimitrios S. Nikolopoulos, Martin Schulz, and Lukasz Wesolowski. Overcoming the scalability challenges of epidemic simulations on Blue Waters. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '14*. IEEE Computer Society, May 2014. LLNL-CONF-648533. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2014.83>.
- [88] **Abhinav Bhatele**, Kathryn Mohror, Steven H. Langer, and Katherine E. Isaacs. There goes the neighborhood: performance degradation due to nearby jobs. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '13*. ACM, November 2013. LLNL-CONF-635776. <http://doi.acm.org/10.1145/2503210.2503247>.
- [89] Nikhil Jain, **Abhinav Bhatele**, Michael P. Robson, Todd Gamblin, and Laxmikant V. Kale. Predicting application performance using supervised learning on communication features. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '13*. ACM, November 2013. LLNL-CONF-635857. <http://doi.acm.org/10.1145/2503210.2503263>.
- [90] Ian Karlin, **Abhinav Bhatele**, Jeff Keasler, Bradford L. Chamberlain, Jonathan Cohen, Zachary DeVito, Riyaz Haque, Dan Laney, Edward Luke, Felix Wang, David Richards, Martin Schulz, and Charles H. Still. Exploring traditional and emerging parallel programming models using a proxy application. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '13*. IEEE Computer Society, May 2013. LLNL-CONF-586774. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2013.115>.

- [91] **Abhinav Bhatele**, Todd Gamblin, Katherine E. Isaacs, Brian T. N. Gunney, Martin Schulz, Peer-Timo Bremer, and Bernd Hamann. Novel views of performance data to analyze large-scale adaptive applications. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '12*. IEEE Computer Society, November 2012. LLNL-CONF-554552. <http://doi.ieeecomputersociety.org/10.1109/SC.2012.80>.
- [92] **Abhinav Bhatele**, Todd Gamblin, Steven H. Langer, Peer-Timo Bremer, Erik W. Draeger, Bernd Hamann, Katherine E. Isaacs, Aaditya G. Landge, Joshua A. Levine, Valerio Pascucci, Martin Schulz, and Charles H. Still. Mapping applications with collectives over sub-communicators on torus networks. In *Proceedings of the ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '12*. IEEE Computer Society, November 2012. LLNL-CONF-556491. <http://doi.ieeecomputersociety.org/10.1109/SC.2012.75>.
- [93] Laercio Pilla, Christiane Ribeiro, Daniel Cordeiro, Chao Mei, **Abhinav Bhatele**, Philippe Navaux, Francois Broquedis, Jean-Francois Mehaut, and Laxmikant V. Kale. Hierarchical approach for load balancing on parallel multi-core systems. In *Proceedings of the International Conference on Parallel Processing, ICPP '12*, September 2012. LLNL-CONF-536171. <http://doi.ieeecomputersociety.org/10.1109/ICPP.2012.9>.
- [94] Vivek Kale, **Abhinav Bhatele**, and William D. Gropp. Weighted locality-sensitive scheduling for mitigating noise on multi-core clusters. In *International Conference on High-Performance Computing, HiPC '11*. IEEE Computer Society, December 2011. LLNL-CONF-492091. <http://doi.ieeecomputersociety.org/10.1109/HiPC.2011.6152722>.
- [95] Ehsan Toton, **Abhinav Bhatele**, Eric J. Bohm, Nikhil Jain, Celso L. Mendes, Ryan M. Mokos, Gengbin Zheng, and Laxmikant V. 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, ICPADS '11*. IEEE Computer Society, December 2011. LLNL-CONF-500821. <http://doi.ieeecomputersociety.org/10.1109/ICPADS.2011.121>.
- [96] **Abhinav Bhatele**, Nikhil Jain, William D. Gropp, and Laxmikant V. Kale. Avoiding hot-spots on two-level direct networks. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '11*. ACM, November 2011. LLNL-CONF-491454. <http://doi.acm.org/10.1145/2063384.2063486>.
- [97] Edgar Solomonik, **Abhinav Bhatele**, and James Demmel. Improving communication performance in dense linear algebra via topology aware collectives. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '11*. ACM, November 2011. LLNL-CONF-491442. <http://doi.acm.org/10.1145/2063384.2063487>.
- [98] **Abhinav Bhatele**, Pritish Jetley, Hormozd Gahvari, Lukasz Wesolowski, William D. Gropp, and Laxmikant Kale. Architectural constraints to attain 1 Exaflop/s for three scientific application classes. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium, IPDPS '11*. IEEE Computer Society, May 2011. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2011.18>.
- [99] **Abhinav Bhatele**, Gagan R. Gupta, Laxmikant V. Kale, and I-Hsin Chung. Automated mapping of regular communication graphs on mesh interconnects. In *Proceedings of IEEE International Conference on High Performance Computing, HiPC '10*. IEEE Computer Society, December 2010. <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=6152722>.
- [100] **Abhinav Bhatele**, Eric Bohm, and Laxmikant V. Kalé. A case study of communication optimizations on 3D mesh interconnects. In *Proceedings of the 15th International Euro-Par Conference on Parallel Processing, Euro-Par '09*, pages 1015–1028. Springer-Verlag, August 2009. [https://link.springer.com/chapter/10.1007/978-3-642-03869-3\\_94](https://link.springer.com/chapter/10.1007/978-3-642-03869-3_94).
- [101] **Abhinav Bhatele**, Laxmikant V. Kalé, and Sameer Kumar. Dynamic topology aware load balancing algorithms for molecular dynamics applications. In *Proceedings of the 23rd international conference on Supercomputing, ICS '09*. ACM, June 2009. <http://doi.acm.org/10.1145/1542275.1542295>.

- [102] **Abhinav Bhatele**, Sameer Kumar, Chao Mei, James C. Phillips, Gengbin Zheng, and Laxmikant V. Kale. Overcoming scaling challenges in biomolecular simulations across multiple platforms. In *Proceedings of the IEEE International Symposium on Parallel and Distributed Processing, IPDPS '08*. IEEE Computer Society, April 2008. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2008.4536317>.

### Workshop Publications

- [103] Keita Teranishi, Harshita Menon, William F. Godoy, Prasanna Balaprakash, David Bau, Tal Ben-Nun, **Abhinav Bhatele**, Franz Franchetti, Michael Franusich, Todd Gamblin, Giorgis Georgakoudis, Tom Goldstein, Arjun Guha, Steven E. Hahn, Costin Iancu, Zheming Jin, Terry Jones, Tze-Meng Low, Het Mankad, Narasinga Rao Miniskar, Mohammad Alaul Haque Monil, Daniel Nichols, Konstantinos Parasyris, Swaroop Pophale, Pedro Valero-Lara, Jeffrey S. Vetter, Samuel Williams, and Aaron Young. Leveraging ai for productive and trustworthy hpc software: Challenges and research directions. In *Proceedings of the First International Workshop on Foundational Large Language Models Advances for HPC*, June 2025.
- [104] Srivishnu Pyda, Daniel Nichols, and **Abhinav Bhatele**. The shortcomings of code LLMs in modeling code properties. In *Proceedings of the IEEE/ACM International Workshop on Large Language Models for Code*, pages 193–199. IEEE Computer Society, May 2025. <https://doi.ieeecomputersociety.org/10.1109/LLM4Code66737.2025.00029>.
- [105] Arunavo Dey, Aakash Dhakal, Tanzima Z. Islam, Jae-Seung Yeom, Tapasya Patki, Daniel Nichols, Alexander Movsesyan, and **Abhinav Bhatele**. Relative performance prediction using few-shot learning. In *Proceedings of the 6th IEEE International Workshop on Deep Analysis of Data-Driven Applications*, pages 1764–1769. IEEE Computer Society, July 2024. <https://doi.ieeecomputersociety.org/10.1109/COMPASAC61105.2024.00278>.
- [106] Stephanie Brink, Ian Lumsden, Connor Scully-Allison, Katy Williams, Olga Pearce, Todd Gamblin, Michela Tauber, Katherine E. Isaacs, and **Abhinav Bhatele**. Usability and performance improvements in Hatchet. In *Proceedings of the Workshop on Programming and Performance Visualization Tools, ProTools '20*, November 2020.
- [107] Alexandre Bergel, **Abhinav Bhatele**, David Boehme, Patrick Gralka, Kevin Griffin, Marc-Andre Hermanns, Dusan Okanovic, Olga Pearce, and Tom Vierjahn. Visual analytics challenges in analyzing calling context trees. In *Programming and Performance Visualization Tools*, volume 11027 of *Lecture Notes in Computer Science*, April 2019. [https://link.springer.com/chapter/10.1007/978-3-030-17872-7\\_14](https://link.springer.com/chapter/10.1007/978-3-030-17872-7_14).
- [108] Kevin Brown, Tianqi Xu, Keita Iwabuchi, Kento Sato, Adam Moody, Kathryn Mohror, Nikhil Jain, **Abhinav Bhatele**, Martin Schulz, Roger Pearce, Maya Gokhale, and Satoshi Matsuoka. Accelerating big data infrastructure and applications (ongoing collaboration). In *Proceedings of the first US-Japan Workshop on Collaborative Global Research on Applying Information Technology*, June 2017. LLNL-CONF-727471.
- [109] Huu Tan Nguyen, **Abhinav Bhatele**, Peer-Timo Bremer, Todd Gamblin, Martin Schulz, Lai Wei, David Böhme, and Kwan-Liu Ma. VIPACT: A visualization interface for analyzing calling context trees. In *Proceedings of the 3rd Workshop on Visual Performance Analysis, VPA '16*, November 2016. LLNL-CONF-704659. <https://doi.org/10.1109/VPA.2016.9>.
- [110] Jae-Seung Yeom, Jayaraman J. Thiagarajan, **Abhinav Bhatele**, Greg Bronevetsky, and Tzanio Kolev. Data-dependent performance modeling of linear solvers for sparse matrices. In *Proceedings of the 7th International Workshop in Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems, PMBS '16*, November 2016. LLNL-CONF-704087. <http://doi.ieeecomputersociety.org/10.1109/PMBS.2016.009>.
- [111] Aniruddha Marathe, Hormozd Gahvari, Jae-Seung Yeom, and **Abhinav Bhatele**. libPowerMon: A lightweight profiling framework to profile program context and system-level metrics. In *Proceedings of the 12th Workshop on High-Performance, Power-Aware Computing, HPPAC '16*. IEEE Computer

Society, May 2016. LLNL-CONF-681427. <http://doi.ieeecomputersociety.org/10.1109/IPDPSW.2016.199>.

- [112] 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.
- [113] Martin Schulz, **Abhinav Bhatele**, David Böhme, Peer-Timo Bremer, Todd Gamblin, Alfredo Gimenez, and Kate Isaacs. A flexible data model to support multi-domain performance analysis. In Christoph Niethammer, José Gracia, Andreas Knüpfer, Michael M. Resch, and Wolfgang E. Nagel, editors, *Proceedings of the 8th International Workshop on Parallel Tools for High Performance Computing*, pages 211–229, Cham, Jan 2015. Springer International Publishing.
- [114] Collin M. McCarthy, Katherine E. Isaacs, **Abhinav Bhatele**, Peer-Timo Bremer, and Bernd Hamann. Visualizing the five-dimensional torus network of the IBM Blue Gene/Q. In *Proceedings of the 1st Workshop on Visual Performance Analysis, VPA '14*, November 2014. LLNL-CONF-661000. <http://doi.ieeecomputersociety.org/10.1109/VPA.2014.10>.
- [115] Steven Langer, **Abhinav Bhatele**, Todd Gamblin, Bert Still, Denise Hinkel, Mike Kumbera, Bruce Langdon, and Ed Williams. Simulating laser-plasma interaction in experiments at the national ignition facility on a Cray XE6. In *Cray User Group Meeting, CUG '12*, April 2012. LLNL-PROC-547711.
- [116] **Abhinav Bhatele** and Laxmikant V. Kale. Heuristic-based techniques for mapping irregular communication graphs to mesh topologies. In *Proceedings of the Workshop on Extreme Scale Computing Application Enablement - Modeling and Tools, ESCAPE '11*, September 2011. LLNL-CONF-491311.
- [117] Martin Schulz, **Abhinav Bhatele**, Peer-Timo Bremer, Todd Gamblin, Katherine Isaacs, Joshua A. Levine, and Valerio Pascucci. Creating a tool set for optimizing topology-aware node mappings. In Holger Brunst, Matthias S. Müller, Wolfgang E. Nagel, and Michael M. Resch, editors, *Tools for High Performance Computing 2011*. Springer Berlin Heidelberg, September 2011. LLNL-CONF-402937. [http://link.springer.com/chapter/10.1007/978-3-642-31476-6\\_1](http://link.springer.com/chapter/10.1007/978-3-642-31476-6_1).
- [118] Gengbin Zheng, Esteban Meneses, **Abhinav Bhatele**, and Laxmikant V. Kale. Hierarchical load balancing for charm++ applications on large supercomputers. In *Proceedings of International Workshop on Parallel Programming Models and Systems Software for High-End Computing, P2S2 '10*, September 2010. <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?tp=&arnumber=5599103>.
- [119] Eric J. Bohm, Sayantan Chakravorty, Pritish Jetley, **Abhinav Bhatele**, and Laxmikant V. Kale. Ckdirect: Unsynchronized one-sided communication in a message-driven paradigm. In *Proceedings of International Workshop on Parallel Programming Models and Systems Software for High-End Computing, P2S2 '09*. IEEE Computer Society, September 2009. <http://ieeexplore.ieee.org/xpl/articleDetails.jsp?arnumber=5365157>.
- [120] **Abhinav Bhatele**, Laxmikant V. Kale, Nicholas Chen, and Ralph E. Johnson. Pattern language for topology aware mapping. In *Proceedings of the Workshop on Parallel Programming Patterns, ParaPLOP '09*, June 2009.
- [121] **Abhinav Bhatele** and Laxmikant V. Kale. An evaluative study on the effect of contention on message latencies in large supercomputers. In *Proceedings of the Workshop on Large-Scale Parallel Processing, LSPP '09*. IEEE Computer Society, April 2009. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2009.5161094>.
- [122] **Abhinav Bhatele** and Laxmikant V. Kale. Application-specific topology-aware mapping for three dimensional topologies. In *Proceedings of the Workshop on Large-Scale Parallel Processing, LSPP '08*. IEEE Computer Society, April 2008. <http://doi.ieeecomputersociety.org/10.1109/IPDPS.2008.4536348>.
- [123] **Abhinav Bhatele** and Guojing Cong. A selective profiling tool: Towards automatic performance tuning. In *Proceedings of the International Workshop on System Management Techniques, Processes and*

Services, SMTPS '07. IEEE Computer Society, April 2007. [✉ http://doi.ieeecomputersociety.org/10.1109/IPDPS.2007.370627](http://doi.ieeecomputersociety.org/10.1109/IPDPS.2007.370627).

---

## Other Publications

### Books & Book Chapters

- [124] **Abhinav Bhatele**, Jeff Hammond, Baboulin Marc, and Carola Kruse. *High Performance Computing*, volume 13948. Springer Cham, May 2023. [✉ https://link.springer.com/book/10.1007/978-3-031-32041-5](https://link.springer.com/book/10.1007/978-3-031-32041-5).
- [125] Ana-Lucia Varbanescu, **Abhinav Bhatele**, Piotr Luszczek, and Baboulin Marc. *High Performance Computing*, volume 13289. Springer Cham, April 2022. [✉ https://link.springer.com/book/10.1007/978-3-031-07312-0](https://link.springer.com/book/10.1007/978-3-031-07312-0).
- [126] **Abhinav Bhatele**, David Boehme, Joshua A. Levine, Allen D. Malony, and Martin Schulz. *Programming and Performance Visualization Tools*, volume 11027. Springer International Publishing, April 2019. [✉ https://www.springer.com/gp/book/9783030178710](https://www.springer.com/gp/book/9783030178710).
- [127] Martin Schulz, Jim Belak, **Abhinav Bhatele**, Peer-Timo Bremer, Greg Bronevetsky, Marc Casas, Todd Gamblin, Katherine E. Isaacs, Ignacio Laguna, Joshua Levine, Valerio Pascucci, David Richards, and Barry Rountree. Performance analysis techniques for the exascale co-design process. In M. Bader, A. Bode, H.-J. Bungartz, M. Gerndt, G.R. Joubert, and F. Peters, editors, *Parallel Computing: Accelerating Computational Science and Engineering (CSE)*, pages 19–32. IOS Press, March 2014. [✉ http://www.ebooks.iospress.nl/volumearticle/35861](http://www.ebooks.iospress.nl/volumearticle/35861).
- [128] Laxmikant V. Kale and **Abhinav Bhatele**, editors. *Parallel Science and Engineering Applications: The Charm++ Approach*. CRC Press, October 2013. [✉ http://www.crcpress.com/product/isbn/9781466504127](http://www.crcpress.com/product/isbn/9781466504127).
- [129] Glenn J. Martyna, Eric J. Bohm, Ramprasad Venkataraman, Anshu Arya, Laxmikant V. Kale, and **Abhinav Bhatele**. OpenAtom: *Ab-initio* molecular dynamics for petascale platforms. In Laxmikant V. Kale and **Abhinav Bhatele**, editors, *Parallel Science and Engineering Applications: The Charm++ Approach*, pages 79–104. CRC Press, October 2013.
- [130] James C. Phillips, Klaus Schulten, **Abhinav Bhatele**, Chao Mei, Yanhua Sun, and Laxmikant V. Kale. Scalable molecular dynamics with NAMD. In Laxmikant V. Kale and **Abhinav Bhatele**, editors, *Parallel Science and Engineering Applications: The Charm++ Approach*, pages 61–77. CRC Press, October 2013.
- [131] **Abhinav Bhatele**. Topology aware task mapping. In David Padua, editor, *Encyclopedia of Parallel Computing*, pages 2057–2062. Springer US, 2011. [✉ http://www.springerlink.com/content/t301u3145x67r615](http://www.springerlink.com/content/t301u3145x67r615).
- [132] Laxmikant V. Kale, **Abhinav Bhatele**, Eric J. Bohm, and James C. Phillips. NAMD (NAnoscale Molecular Dynamics). In David Padua, editor, *Encyclopedia of Parallel Computing*, pages 1249–1254. Springer US, 2011. [✉ http://www.springerlink.com/content/t42n856361705231](http://www.springerlink.com/content/t42n856361705231).
- [133] **Abhinav Bhatele**, Benjamin Fergoso Munoz, Carolina Ana Sternberg, Hio Lam Lao, Jonathan Andrew Khu Ang, Jong-Yeon Ee, Joonwon Yoon, Joyce Wei, Kashif Altaf, Minna Yung, Mrinalini Rao, Ruqing Pan, Jong Won Han, and Zai yu Elisia Phua, editors. *International Student Guide Book 2009–2010*. Korean Cultural Center, November 2009.
- [134] Klaus Schulten, James C. Phillips, Laxmikant V. Kale, and **Abhinav Bhatele**. Biomolecular modeling in the era of petascale computing. In David Bader, editor, *Petascale Computing: Algorithms and Applications*, pages 165–181. Chapman & Hall, December 2007. [✉ http://www.crcnetbase.com/doi/abs/10.1201/9781584889106.ch9](http://www.crcnetbase.com/doi/abs/10.1201/9781584889106.ch9).

### Theses & Technical Reports

- [135] **Abhinav Bhatele**. Task mapping on complex computer network topologies for improved performance. Technical report, LDRD Final Report, Lawrence Livermore National Laboratory, October 2015. LLNL-TR-678732.
- [136] Louis Howell, Brian Gunney, and **Abhinav Bhatele**. Characterization of proxy application performance on advanced architectures: UMT2013, MCB, AMG2013. Technical report, Lawrence Livermore National Laboratory, October 2015. LLNL-TR-677974. <http://www.osti.gov/scitech/biblio/1224409>.
- [137] Janine Bennett, Robert Clay, Gavin Baker, Marc Gamell, David Hollman, Samuel Knight, Hemanth Kolla, Gregory Sjaardema, Nicole Slattengren, Keita Teranishi, Jeremiah Wilke, Matt Bettencourt, Steve Bova, Ken Franko, Paul Lin, Ryan Grant, Si Hammond, Stephen Olivier, Laxmikant Kale, Nikhil Jain, Eric Mikida, Alex Aiken, Mike Bauer, Wonchan Lee, Elliott Slaughter, Sean Treichler, Martin Berzins, Todd Harman, Alan Humphrey, John Schmidt, Dan Sunderland, Pat McCormick, Samuel Gutierrez, Martin Schulz, **Abhinav Bhatele**, David Boehme, Peer-Timo Bremer, and Todd Gamblin. ASC ATDM level 2 milestone #5325: Asynchronous many-task runtime system analysis and assessment for next generation platforms. Technical report, Sandia National Laboratories, September 2015. SAND2015-8312. <https://cfwebprod.sandia.gov/cfdocs/CompResearch/docs/ATDM-AMT-L2-Final-SAND2015-8312.pdf>.
- [138] Harshitha Menon, **Abhinav Bhatele**, Sebastien Fourestier, Laxmikant Kale, and Francois Pellegrini. Applying graph partitioning methods in measurement-based dynamic load balancing. Technical report, Dept. of Computer Science, University of Illinois, May 2015. <http://hdl.handle.net/2142/75950>.
- [139] Laxmikant V. Kale, Anshu Arya, **Abhinav Bhatele**, Abhishek Gupta, Nikhil Jain, Pritish Jetley, Jonathan Lifflander, Phil Miller, Yanhua Sun, Ramprasad Venkataraman, Lukasz Wesolowski, and Gengbin Zheng. Charm++ for productivity and performance: A submission to the 2011 HPC Class II Challenge. Technical report, Dept. of Computer Science, University of Illinois, November 2011.
- [140] **Abhinav Bhatele**. *Automating Topology Aware Mapping for Supercomputers*. PhD thesis, Dept. of Computer Science, University of Illinois, August 2010. <http://hdl.handle.net/2142/16578>.
- [141] **Abhinav Bhatele**. Application specific topology aware mapping and load balancing for three dimensional torus topologies. Master's thesis, Dept. of Computer Science, University of Illinois, December 2007. <http://charm.cs.illinois.edu/papers/BhateleMSThesis07.shtml>.

---

## Significant Presentations

### Technical Posters

- [142] Onur Cankur, Brian Austin, and **Abhinav Bhatele**. Understanding GPU utilization using LDMS data on Perlmutter. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '25, November 2025.
- [143] Siddharth Singh, Mahua Singh, Keshav Pradeep, and **Abhinav Bhatele**. Optimizing collectives with large payloads on GPU-based supercomputers. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '25, November 2025.
- [144] Prajwal Singhanian, Siddharth Singh, Lannie Dalton Hough, Ishan Revankar, Harshitha Menon, Charles Fredrick Jekel, and **Abhinav Bhatele**. Understanding communication bottlenecks in multi-node LLM inference. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '25, November 2025.
- [145] Cunyang Wei, Keshav Pradeep, and **Abhinav Bhatele**. Unmasking performance variability in GPU codes on supercomputers. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '25, November 2025.
- [146] Aman Chaturvedi, Daniel Nichols, Siddharth Singh, and **Abhinav Bhatele**. Creating code llms for hpc: It's llms all the way down. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '24, November 2024.

- [147] Aditya Tomar, Siddharth Singh, Tom Goldstein, and **Abhinav Bhatele**. Eve: Less memory, same might. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '24, November 2024.
- [148] Joshua H. Davis, Pranav Sivaraman, Isaac Minn, and **Abhinav Bhatele**. Evaluating performance portability of gpu programming models. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '23, November 2023.
- [149] Nikodemos Koutsoheras, Sayan Ghosh, Nathan Tallent, Joshua Suetterlein, and **Abhinav Bhatele**. The impact of process topology on rma programming models: A study on nersc perlmutter. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '23, November 2023.
- [150] Alexander Movsesyan, Rakrish Dhakal, Aditya Ranjan, Jordan Marry, Onur Cankur, and **Abhinav Bhatele**. Pipit: Simplifying analysis of parallel execution traces. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '23, November 2023.
- [151] Daniel Nichols, Aniruddha Marathe, Harshitha Menon, Todd Gamblin, and **Abhinav Bhatele**. Modeling parallel programs using large language models. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '23, November 2023.
- [152] Joshua H. Davis, Justin Shafner, Daniel Nichols, Nathan Grube, Pino Martin, and **Abhinav Bhatele**. Extreme-scale computational fluid dynamics with AMR on GPUs. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '22, November 2022.
- [153] Joy Kitson, Ian Costello, Diego Jiménez, Jiangzhuo Chen, Jaemin Choi, Stefan Hoops, Esteban Meneses, Tamar Kellner, Henning Mortveit, Jae-Seung Yeom, Laxmikant V. Kale, Madhav V. Marathe, and **Abhinav Bhatele**. Loimos: A large-scale epidemic simulation framework for realistic social contact networks. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '22, November 2022.
- [154] Daniel Nichols, Dilan Gunawardana, Aniruddha Marathe, Todd Gamblin, and **Abhinav Bhatele**. Noncommittal commits: Predicting performance slowdowns in version control history. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '22, November 2022.
- [155] Daniel Nichols, Jae-Seung Yeom, and **Abhinav Bhatele**. Predicting cross-platform relative performance with deep generative models. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '22, November 2022.
- [156] Siddharth Singh and **Abhinav Bhatele**. Optimizing communication in parallel deep learning via parameter pruning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '22, November 2022.
- [157] Yiheng Xu, Kathryn Mohror, Hariharan Devarajan, Cameron Stanavige, and **Abhinav Bhatele**. Chaining multiple tools and libraries using Gotcha. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '21, November 2021.
- [158] Saptarshi Bhowmik, Nikhil Jain, Xin Yuan, and **Abhinav Bhatele**. A simulation study of hardware parameters for gpu-based hpc platforms. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '20, November 2020.
- [159] Ian Costello and **Abhinav Bhatele**. Predicting the performance of jobs in the queue using machine learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '20, November 2020.
- [160] Suraj Kesavan, Harsh Bhatia, **Abhinav Bhatele**, Stephanie Brink, Olga Pearce, Todd Gamblin, Peer-Timo Bremer, and Kwan-Liu Ma. Scalable comparative visualization of ensembles of call graphs

using callflow. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '20, November 2020.

- [161] Joy Kitson, Sudheer Chunduri, and **Abhinav Bhatele**. Analyzing interconnect congestion on a production dragonfly-based system. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '20, November 2020.
- [162] Jaemin Choi, **Abhinav Bhatele**, and David Richards. Fast profiling-based performance modeling of distributed gpu applications. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '19, November 2019. LLNL-POST-. [https://sc19.supercomputing.org/proceedings/src\\_poster/src\\_poster\\_pages/spostg126.html](https://sc19.supercomputing.org/proceedings/src_poster/src_poster_pages/spostg126.html).
- [163] Emilio Castillo, Nikhil Jain, Marc Casas, Miquel Moreto, Martin Schulz, Ramon Beivide, Mateo Valero, and **Abhinav Bhatele**. Optimizing computation-communication overlap in asynchronous task-based programs. In *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP '19, New York, NY, February 2019. ACM. LLNL-POST-767691.
- [164] Rob Ross, Misbah Mubarak, Ray Bair, **Abhinav Bhatele**, Nikhil Jain, Scott Pakin, Chris Carothers, Scott Hemmert, Jeremiah Wilke, Si Hammond, and Mike Levenhagen. Simulation of exascale system interconnects. In *Exascale Computing Project Annual Meeting*, ECP '17, February 2017. LLNL-POST-719979.
- [165] Yarden Livnat, **Abhinav Bhatele**, Nikhil Jain, Peer-Timo Bremer, and Valerio Pascucci. DragonView: Toward understanding network interference in dragonfly-based supercomputers. In *Proceedings of the SCI Institute Technical Exchange*, SCIX '16, November 2016. LLNL-POST-.
- [166] Alfredo Giménez, Todd Gamblin, Peer-Timo Bremer, **Abhinav Bhatele**, and Martin Schulz. Combining disparate data sources in the HPC ecosystem. In *Proceedings of the Salishan Conference on High Speed Computing*, Salishan '16, April 2016. LLNL-POST-692697.
- [167] **Abhinav Bhatele**, Nikhil Jain, Yarden Livnat, Valerio Pascucci, and Peer-Timo Bremer. Simulating and visualizing traffic on the dragonfly network. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '15, November 2015. LLNL-POST-676008. [http://sc15.supercomputing.org/sites/all/themes/SC15images/tech\\_poster/tech\\_poster\\_pages/post109.html](http://sc15.supercomputing.org/sites/all/themes/SC15images/tech_poster/tech_poster_pages/post109.html).
- [168] 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 *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '14, November 2014. LLNL-POST-662677. [http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/tech\\_poster/tech\\_poster\\_pages/post236.html](http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/tech_poster/tech_poster_pages/post236.html).
- [169] Andrew Titus and **Abhinav Bhatele**. Supervised learning for parallel application performance prediction. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '14, November 2014. LLNL-POST-662676. [http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/src\\_poster/src\\_poster\\_pages/spost127.html](http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/src_poster/src_poster_pages/spost127.html).
- [170] Dylan Wang, **Abhinav Bhatele**, and Dipak Ghosal. Performance variability due to job placement on Edison. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '14, November 2014. LLNL-POST-662284. [http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/src\\_poster/src\\_poster\\_pages/spost120.html](http://sc14.supercomputing.org/sites/all/themes/sc14/files/archive/src_poster/src_poster_pages/spost120.html).
- [171] Katherine E. Isaacs, Todd Gamblin, **Abhinav Bhatele**, Peer-Timo Bremer, Martin Schulz, and Bernd Hamann. Extracting logical structure and identifying stragglers in parallel execution traces. In *Proceedings of the 19th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP '14, New York, NY, February 2014. ACM. LLNL-POST-649674. <http://doi.acm.org/10.1145/2555243.2555288>.

- [172] **Abhinav Bhatele**, Todd Gamblin, Steven H. Langer, Peer-Timo Bremer, and Martin Schulz. Mapping collectives over sub-communicators on torus networks. In *Current Challenges in Computing 2012: Network Science*, Napa, CA, August 2012. LLNL-POST-563791.
- [173] Aaditya Landge, Joshua A. Levine, Peer-Timo Bremer, Martin Schulz, Todd Gamblin, **Abhinav Bhatele**, Katherine Isaacs, and Valerio Pascucci. Interactive visualizations for performance analysis of heterogeneous computing clusters. In *GPU Technology Conference*, GTC '12, San Jose, CA, May 2012. LLNL-POST-518831.
- [174] **Abhinav Bhatele**, Todd Gamblin, Martin Schulz, and Peer-Timo Bremer. Intuitive visualizations for analyzing exascale workloads. In *Exascale Research Conference*, Portland, OR, April 2012. LLNL-POST-545412.
- [175] **Abhinav Bhatele**, Todd Gamblin, Brian T. N. Gunney, Martin Schulz, Peer-Timo Bremer, and Katherine Isaacs. Revealing performance artifacts in parallel codes through multi-domain visualizations. In *SIAM Conference on Parallel Processing for Scientific Computing*, SIAM PP'12, Savannah, GA, February 2012. LLNL-POST-527971. [http://meetings.siam.org/session/dsp\\_talk.cfm?p=49976](http://meetings.siam.org/session/dsp_talk.cfm?p=49976).
- [176] **Abhinav Bhatele**, Lukasz Wesolowski, Eric Bohm, Edgar Solomonik, and Laxmikant V. Kalé. Performance comparison of Intrepid, Jaguar and Ranger using scientific applications. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '09, November 2009.
- [177] **Abhinav Bhatele**, Eric Bohm, and Laxmikant V. Kalé. Topology aware task mapping techniques: an API and case study. In *Proceedings of the 14th ACM SIGPLAN symposium on Principles and Practice of Parallel Programming*, PPOPP '09. ACM, February 2009. <http://doi.acm.org/10.1145/1504176.1504225>.
- [178] **Abhinav Bhatele** and Laxmikant V. Kalé. Effects of contention on message latencies in large supercomputers. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '08, November 2008.
- [179] **Abhinav Bhatele** and Laxmikant V. Kalé. Automatic topology-aware task mapping for parallel applications running on large parallel machines. In *Proceedings of the IEEE International Parallel & Distributed Processing Symposium*, IPDPS '08, April 2008.

### Keynotes & Invited Talks

- [180] **Abhinav Bhatele**. Lost in translation: LLMs and whole-repository porting of HPC codes. In *Workshop on AI Assisted Software Development for HPC*, AI4Dev '25, San Diego, CA, September 2025.
- [181] **Abhinav Bhatele** and Siddharth Singh. Democratizing AI: Open-source scalable LLM training on GPU-based supercomputers. In *ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '24, Atlanta, GA, November 2024.
- [182] **Abhinav Bhatele**. A 4D hybrid algorithm to scale parallel training to thousands of GPUs. In *Second International Workshop on Artificial Intelligence and Big Data Analytics at Extreme Scale*, TIES '24, Providence, RI, July 2024.
- [183] **Abhinav Bhatele**. The role of generative AI in HPC code development. In *Second International Workshop on Conversational AI Interfaces for HPC*, CONVOHPC '24, Providence, RI, July 2024.
- [184] **Abhinav Bhatele**. The future of parallel code development: Will AI lead the way? In *33rd International Symposium on High-Performance Parallel and Distributed Computing*, HPDC '24, Pisa, Italy, June 2024.
- [185] **Abhinav Bhatele**. Distributed AI training and use of LLMs for parallel code development. In *Lawrence Livermore National Laboratory*, Livermore, CA, March 2024.

- [186] **Abhinav Bhatele**. Making distributed training easier and more communication-efficient. In *14th International Conference on Cloud Computing, Data Science & Engineering*, IEEE Confluence '24, Noida, India, January 2024.
- [187] **Abhinav Bhatele**. Who's winning the performance portability race on GPU platforms? In *ACM/IEEE 8th International Workshop on Extreme Scale Programming Models and Middleware*, ESPM2 '23, Denver, CO, November 2023.
- [188] **Abhinav Bhatele**. Pipit: Enabling programmatic analysis of parallel execution traces. In *Scalable Tools Workshop*, Lake Tahoe, CA, June 2023.
- [189] **Abhinav Bhatele**. HPC and AI/ML: A synergistic relationship. In *SPCL\_Bcast(COMM\_WORLD) Seminar Series, ETH Zurich*, virtual, March 2023. <https://sycl.inf.ethz.ch/Bcast>.
- [190] **Abhinav Bhatele**. Rethinking the parallelization of extreme-scale deep learning. In *CCT/LSU Colloquium on Artificial Intelligence Research and Optimization*, Baton Rouge, LA, October 2022. <https://www.cct.lsu.edu/lectures/rethinking-parallelization-extreme-scale-deep-learning>.
- [191] **Abhinav Bhatele**. Using system monitoring data to understand and mitigate performance variability. In *Workshop on Monitoring and Analysis for HPC Systems Plus Applications*, HPCMASPA '22, Heidelberg, Germany, September 2022. <https://sites.google.com/view/hpcmaspa2022>.
- [192] **Abhinav Bhatele**. Analyzing performance profiles using hatchet. In *HPC, Big Data and Data Science devroom at FOSDEM*, FOSDEM '21, virtual, February 2021. [https://fosdem.org/2021/schedule/track/hpc\\_big\\_data\\_and\\_data\\_science/](https://fosdem.org/2021/schedule/track/hpc_big_data_and_data_science/).
- [193] **Abhinav Bhatele**. Analyzing call graphs using hatchet. In *18th Annual Workshop on Charm++ and its Applications*, Charm++ Workshop '20, virtual, October 2020. <https://charm.cs.illinois.edu/workshops/charmWorkshop2020/program.shtml>.
- [194] **Abhinav Bhatele**. Modeling performance variability via machine learning. In *Workshop on Modeling & Simulation of Systems and Applications*, ModSim '20, virtual, July 2020. <https://www.bnl.gov/modsim2020/files/pdf/agenda.pdf>.
- [195] **Abhinav Bhatele**. On mitigating congestion in high performance networks. In *Technical University Munich/LRZ*, Munich, Germany, July 2019. <https://www.ias.tum.de/en/events/events-single-view/article/on-mitigating-congestion-in-high-performance-networks-talk-bei-tum-ias-visiting-fellow-pr>
- [196] **Abhinav Bhatele**. On mitigating congestion in high performance networks. In *ETH Zurich*, Zurich, Switzerland, July 2019.
- [197] **Abhinav Bhatele**, Stephanie Brink, and Todd Gamblin. Hatchet: Pruning the overgrowth in parallel profiles. In *Scalable Tools Workshop*, Lake Tahoe, CA, July 2019. [https://dyninst.github.io/scalable\\_tools\\_workshop/petascale2019/monday.html](https://dyninst.github.io/scalable_tools_workshop/petascale2019/monday.html).
- [198] **Abhinav Bhatele**. The case of performance variability on dragonfly networks. In *Tokyo Institute of Technology*, Tokyo, Japan, May 2019.
- [199] **Abhinav Bhatele**. Understanding and optimizing performance on hpc networks. In *RIKEN Center for Computational Science (CCS)*, Kobe, Japan, May 2019.
- [200] **Abhinav Bhatele**. On mitigating congestion in high performance networks. In *CS Colloquium, Rice University*, Houston, TX, February 2019. LLNL-PRES-767646.
- [201] **Abhinav Bhatele**. On mitigating congestion in high performance networks. In *CS Colloquium, University of Maryland*, College Park, MD, February 2019. LLNL-PRES-767646. <https://www.cs.umd.edu/event/2019/01/mitigating-congestion-high-performance-networks>.
- [202] Samuel A. Pollard, Nikhil Jain, Stephen Herbein, and **Abhinav Bhatele**. Interference-free node allocations on fat-tree clusters. In *ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '18, November 2018. LLNL-PRES-761657.

- [203] **Abhinav Bhatele**. Network simulation and performance analysis research at LLNL. In *Technical University Munich*, Munich, Germany, July 2018. LLNL-PRES-756255.
- [204] Jayaraman J. 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*, Vancouver, Canada, May 2018. LLNL-PRES-751869.
- [205] **Abhinav Bhatele**, Nikhil Jain, Todd Gamblin, and Laxmikant V. Kale. Network design and comparison studies using tracer simulations. In *SIAM Conference on Parallel Processing for Scientific Computing, SIAM PP'18*, Tokyo, Japan, March 2018. LLNL-PRES-742377.
- [206] Takatsugu Ono, Yuta Kakibuka, Koji Inoue, Nikhil Jain, **Abhinav Bhatele**, and Shinobu Miwa. Dynamic power management of hpc networks. In *Fifth DOE/MEXT Meeting*, Tokyo, Japan, March 2018. LLNL-PRES-747751.
- [207] **Abhinav Bhatele**. Analyzing and mitigating congestion on high performance networks. In *CS Colloquium, The University of Arizona*, Tucson, AZ, February 2018. LLNL-PRES-756280.
- [208] **Abhinav Bhatele**. Optimizing performance on hpc systems. In *CS Colloquium, University of Oregon*, Eugene, OR, February 2018. LLNL-PRES-756281. <https://cs.uoregon.edu/colloquium/analyzing-and-mitigating-congestion-high-performance-networks>.
- [209] **Abhinav Bhatele**, Jae-Seung Yeom, Nikhil Jain, Chris J. Kuhlman, Yarden Livnat, Keith R. Bisset, Laxmikant V. Kale, and Madhav V. Marathe. Scaling parallel epidemic simulations using Charm++. In *SC Birds-of-a-feather Session on Charm++ and AMPI, SC '17 BoF*, Denver, CO, November 2017. LLNL-PRES-742360.
- [210] Nikhil Jain, **Abhinav Bhatele**, Louis Howell, David Böhme, Ian Karlin, Edgar Leon, Misbah Mubarak, Noah Wolfe, Todd Gamblin, and Matthew Leininger. Predicting the performance impact of different fat-tree configurations. In *ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis, SC '17*, Denver, CO, November 2017. LLNL-PRES-742377.
- [211] **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 *NERSC Users Group meeting, NUG '17*, Berkeley, CA, September 2017. LLNL-PRES-731838.
- [212] **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 *CCGrid '17 SCALE Challenge*, Madrid, Spain, May 2017. LLNL-PRES-731838.
- [213] **Abhinav Bhatele**, Staci Smith, David K. Lowenthal, Jayaraman J. Thiagarajan, and Laxmikant V. Kale. Identifying root causes of application performance variability & anomalies. In *SIAM Conference on Computational Science & Engineering, SIAM CSE'17*, Atlanta, GA, March 2017. LLNL-PRES-725619. [http://meetings.siam.org/sess/dsp\\_talk.cfm?p=82513](http://meetings.siam.org/sess/dsp_talk.cfm?p=82513).
- [214] **Abhinav Bhatele**. Early results on exploring Charm++/ Adaptive MPI (AMPI) in WSC codes. In *JOWOG-34 Annual Computer Science Meeting*, Livermore, CA, February 2017. LLNL-PRES-721763.
- [215] **Abhinav Bhatele**. Impact of job placement and routing policies on application performance. In *11th Scheduling for Large Scale Systems Workshop*, Nashville, TN, May 2016. LLNL-PRES-692619.
- [216] **Abhinav Bhatele**. Network simulation and visualization tools: Tracer, Dragonview. In *JOWOG-34 Annual Computer Science Meeting*, Los Alamos, NM, February 2016. LLNL-PRES-675498.
- [217] **Abhinav Bhatele**. Analyzing and optimizing data movement on HPC networks. In *Workshop on Exascale Software Technologies*, Albuquerque, NM, January 2016. LLNL-PRES-681760.

- [218] **Abhinav Bhatele**. Analyzing HPC interconnects using network simulators. In *Network Dynamics and Simulation Science Laboratory (NDSSL), Virginia Tech*, Blacksburg, VA, November 2015. LLNL-PRES-679783.
- [219] **Abhinav Bhatele**. Network simulations to predict congestion and performance. In *JOWOG-34 Annual Meeting*, Livermore, CA, July 2015. LLNL-PRES-675498.
- [220] **Abhinav Bhatele**. Parallel execution models, performance prediction, and optimization. In *Indian Institute of Technology, Kanpur*, Kanpur, INDIA, May 2015. LLNL-PRES-665882.
- [221] **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 *13th Annual Workshop on Charm++ and its Applications*, Charm++ Workshop '15, Urbana, IL, May 2015. LLNL-PRES-670743. <http://charm.cs.illinois.edu/workshops/charmWorkshop2015/program.php>.
- [222] **Abhinav Bhatele**, Nikhil Jain, Xiang Ni, and Laxmikant V. Kale. Maximizing throughput on a dragonfly network. In *SIAM Conference on Computational Science & Engineering*, SIAM CSE'15, Salt Lake City, UT, March 2015. LLNL-PRES-668919. [http://meetings.siam.org/sess/dsp\\_talk.cfm?p=66707](http://meetings.siam.org/sess/dsp_talk.cfm?p=66707).
- [223] **Abhinav Bhatele**. Introduction to Git (Part I). In *Software Improvement Networking Group (SWING) Git Tutorial Series*, Livermore, CA, January 2015. LLNL-PRES-645972.
- [224] **Abhinav Bhatele**. Parallel execution models, performance prediction, and optimization. In *Indian Institute of Technology, Delhi*, Delhi, INDIA, December 2014. LLNL-PRES-665882.
- [225] **Abhinav Bhatele**. Tools for visualizing communication, network traffic, and job placement. In *8th Annual Petascale Tools Workshop*, Madison, WI, August 2014. LLNL-PRES-659275. [http://www.paradyn.org/petascale2014/monday\\_paradyn.html](http://www.paradyn.org/petascale2014/monday_paradyn.html).
- [226] **Abhinav Bhatele**. Task mapping, job placements, and routing strategies. In *12th Annual Workshop on Charm++ and its Applications*, Charm++ Workshop '14, Urbana, IL, April 2014. LLNL-PRES-654602. <http://charm.cs.illinois.edu/workshops/charmWorkshop2014/program.php>.
- [227] **Abhinav Bhatele** and Todd Gamblin. Placing communicating tasks apart to maximize bandwidth. In *SIAM Conference on Computational Science & Engineering*, SIAM CSE'13, Boston, MA, March 2013. LLNL-PRES-621732. [http://meetings.siam.org/sess/dsp\\_talk.cfm?p=53092](http://meetings.siam.org/sess/dsp_talk.cfm?p=53092).
- [228] **Abhinav Bhatele**. Exploring Charm++ for LULESH. In *SIAM Conference on Computational Science & Engineering*, SIAM CSE'13, Boston, MA, February 2013. LLNL-PRES-621033. [http://meetings.siam.org/sess/dsp\\_talk.cfm?p=53991](http://meetings.siam.org/sess/dsp_talk.cfm?p=53991).
- [229] **Abhinav Bhatele**. On maximizing bandwidth utilization on torus interconnects. In *IBM Research*, Yorktown Heights, NY, October 2012. LLNL-PRES-592213.
- [230] **Abhinav Bhatele** and Todd Gamblin. OS/Runtime challenges for dynamic topology aware mapping. In *U.S. Department of Energy Exascale Operating Systems and Runtime Research Workshop*, Washington, DC, October 2012. LLNL-PRES-587572. <https://collab.mcs.anl.gov/display/exasr/ExaOSR-2012+Workshop+Agenda>.
- [231] **Abhinav Bhatele**, Peer-Timo Bremer, Todd Gamblin, and Martin Schulz. PAVE: Intuitive visualizations for analyzing exascale workloads. In *Exascale Research Conference*, Portland, OR, April 2012. LLNL-PRES-540811.
- [232] **Abhinav Bhatele** and Laxmikant V. Kale. Topology aware resource allocation and mapping challenges at exascale. In *SIAM Conference on Parallel Processing for Scientific Computing*, SIAM PP'12, Savannah, GA, February 2012. LLNL-PRES-529376. [http://meetings.siam.org/sess/dsp\\_talk.cfm?p=47860](http://meetings.siam.org/sess/dsp_talk.cfm?p=47860).
- [233] **Abhinav Bhatele**. A mapping framework for torus networks. In *Par Lab, Computer Science Division, University of California*, Berkeley, CA, October 2011. LLNL-PRES-505691.

- [234] **Abhinav Bhatele**. Automating topology aware mapping on large supercomputers. In *Computing Sciences Seminar, Lawrence Berkeley National Laboratory*, Berkeley, CA, January 2011.
- [235] **Abhinav Bhatele**. Automating topology aware mapping on large supercomputers. In *Mathematics and Computer Science Division, Argonne National Laboratory*, Argonne, IL, January 2011. <http://www.alcf.anl.gov/events/automating-topology-aware-mapping-supercomputers>.
- [236] **Abhinav Bhatele**. Mapping your application on interconnect topologies: Effort versus benefits. In *George Michael HPC Fellow Presentation, International Conference for High Performance Computing, Networking, Storage and Analysis, SC '10*, New Orleans, LA, November 2010. [http://sc10.supercomputing.org/schedule/event\\_detail.php?evid=hpcp101](http://sc10.supercomputing.org/schedule/event_detail.php?evid=hpcp101).
- [237] **Abhinav Bhatele** and Laxmikant V. Kale. Mapping parallel applications on the machine topology: Lessons learned. In *TeraGrid '10*, Pittsburgh, PA, August 2010. <https://www.xsede.org/wwwteragrid/archive/web/events/tg10/techtrack.html#mapping>.
- [238] **Abhinav Bhatele**, Eric Lee, Ly Le, Leonardo Trabuco, Eduard Schreiner, Jen Hsin, James C. Phillips, Laxmikant V. Kale, and Klaus Schulten. Biomolecular simulations using NAMD on TeraGrid machines. In *TeraGrid '10*, Pittsburgh, PA, August 2010. <https://www.xsede.org/wwwteragrid/archive/web/events/tg10/sciencetrack.html#namd>.
- [239] **Abhinav Bhatele**. Automating topology aware mapping on large supercomputers. In *CSE Seminar, College of Computing, Georgia Tech*, Atlanta, GA, March 2010. <http://hdl.handle.net/1853/32764>.
- [240] **Abhinav Bhatele**. Automating topology aware task mapping for large supercomputers. In *Doctoral Showcase, International Conference for High Performance Computing, Networking, Storage and Analysis, SC '09*, Portland, OR, November 2009.
- [241] **Abhinav Bhatele**. Load balancing and topology aware mapping for petascale machines. In *Scaling to Petascale Summer School, NCSA, Urbana, IL*, August 2009. <http://www.greatlakesconsortium.org/events/scaling/agenda.html>.
- [242] **Abhinav Bhatele**. The Charm++ programming model and NAMD. In *Barcelona Supercomputing Center*, Barcelona, Spain, February 2009.
- [243] **Abhinav Bhatele** and Laxmikant V. Kale. IS TOPOLOGY IMPORTANT AGAIN? – Effects of contention on message latencies in large supercomputers. In *ACM Student Research Competition, International Conference for High Performance Computing, Networking, Storage and Analysis, SC '08*, Austin, TX, November 2008. <http://src.acm.org/bhatele/bhatele.html>.
- [244] **Abhinav Bhatele**. Topology aware mapping for performance optimization of science applications. In *Institute for Advanced Computing Applications and Technology (IACAT) Seminar*, Urbana, IL, October 2008. <http://charm.cs.uiuc.edu/CPC/bhatele.shtml>.
- [245] **Abhinav Bhatele**, Laxmikant V. Kale, and Sameer Kumar. Dynamic topology aware load balancing algorithms for MD applications. In *UK e-Science All Hands Meeting*, Edinburgh, UK, September 2008. <http://www.allhands.org.uk/2008/programme>.