

Todd Gamblin

Computer Scientist, Advanced Technology Office, Livermore Computing
Lawrence Livermore National Laboratory (LLNL)

P.O. Box 808, L-557

Livermore, CA, 94551

☎ 925-422-9319

✉ tgamblin@llnl.gov

<https://people.llnl.gov/gamblin2>

Education

- 2009 **Ph.D., Computer Science**, *University of North Carolina at Chapel Hill*.
Scalable Performance Measurement and Analysis. Advisor: Daniel A. Reed.
- 2005 **M.S., Computer Science**, *University of North Carolina at Chapel Hill*.
- 2002 **B.A., Computer Science & Japanese**, *Williams College*.

Awards & Honors

- 2021 LLNL award for critical contributions to LLNL open source software strategy
- 2020 LLNL Director's Science & Technology Award for the Spack Package Manager
- 2020 LLNL WSC Code Development award for spearheading software process improvements
- 2019 **R&D 100 Award** for Spack: A Package Manager for Supercomputers
- 2019 **R&D 100 Special Recognition** for Spack: Silver Medalist in the "Market Disruptor" category
- 2019 **Top Technology Development**: Federal Laboratory Consortium Regional Award for Spack
- 2017 LLNL Early and Mid-Career Recognition (EMCR) Award
- 2017 LLNL WCI Directorate Award
- 2015 LLNL Director's Institutional Operational Excellence (DIOE) Award for labwide collaboration tools
- 2015 LLNL award for organizing & reinvigorating the JOWOG 34 Applied CS workshop
- 2014 U.S. Department of Energy (DOE) **Early Career Research Award** (4% accepted)
- 2014 Best Paper, International Symposium on Cluster, Cloud and Grid Computing (CCGrid)
- 2013 Best Paper, International Conference on Supercomputing (ICS)
- 2012, 2013 Two awards for building collaborative development tools for all LLNL HPC users
- 2011 LLNL award for leadership on the PAVE project
- 2006-2008 Latané Interdisciplinary Fellowship
- 2007 Member, Frank Porter Graham Honor Society
- 2006 Best Poster, Los Alamos Computer Science Institute (LACSI) Symposium
- 2004 NSF East Asia and Pacific Summer Institutes (EAPSI) Fellowship

Experience

- 2019-present **LLNL, Senior Principal Computer Scientist (MTS 5)**, Livermore Computing (LC).
 - Lead: Strategic engagements with cloud providers for LC Advanced Technology Office.
 - Lead: Spack open source project: package manager for HPC, thousands of users worldwide (spack.io).
 - Lead: BUILD: LLNL Strategic research initiative on software integration ([LLNL site](https://llnl.gov/build)).
 - Lead: Packaging working group with HPE and AMD for the [El Capitan](https://www.elcapitanproject.org) procurement.
 - Lead: Packaging Technologies project within the U.S. Exascale Computing Project (exascaleproject.org).
 - Lead: DevRAMP: Reproducibility, Analysis, Monitoring, and Performance for Developers.
- 2015-2019 **LLNL, Principal Computer Scientist (MTS 4)**, Center for Applied Scientific Computing (CASC).
- 2010-2015 **LLNL, Senior Computer Scientist (MTS 3)**, Center for Applied Scientific Computing (CASC).
- 2009-2010 **LLNL, Postdoctoral Scholar**, Center for Applied Scientific Computing (CASC).
Developed $O(\log(P))$ parallel clustering algorithm for on-line performance analysis.
- 2008-2009, **LLNL, Research Intern**.
- Jun, Aug 2007 Developed scalable parallel load balance measurement tools using wavelet compression.

- 2004-2008 **Renaissance Computing Institute (RENCI), UNC Chapel Hill**, *Research Assistant*.
Researched statistical and numerical techniques for performance measurement data reduction.
- Summer 2006 **IBM Thomas J. Watson Research Center**, *Research Intern*.
Developed distributed clock synchronization for virtualized environments on IBM System z mainframes.
- Summer 2004 **Nanya-Nakamura Laboratory, University of Tokyo**, *Research Intern*.
On NSF EAPSI fellowship, investigated low-power register design using asynchronous digital logic.
- 2003-2004 **Clockless Computing Laboratory, UNC Chapel Hill**, *Research Assistant*.
Investigated modular interface designs for asynchronous (clockless) digital circuits.
- 2002-2003 **ValueCommerce, Co., Ltd.**, *Software Developer*, Tokyo, Japan.
Developed control panel for web hosting application; conducted performance analysis for ad servers.
- 2001-2002 **Williams College**, *Undergraduate Research Assistant*.
Investigated cache-conscious memory performance optimizations.
- Summer 2001 **EMC Corporation**, *Summer Intern*, Research Triangle Park, NC.
Developed distributed performance monitoring tools and web-based user interface for NAS servers.
- Summer 2000 **EMC Corporation**, *Summer Intern*, Research Triangle Park, NC.
Developed distributed, cross-platform performance monitoring tool for large networks.
- Summer 1999 **NCSU Physics Department**, *Summer Research Assistant*.
Designed imaging algorithms and an archival website for the Diffraction Enhanced Imaging (DEI) Group.

Research and Other Funding

- 2020-present **PI**, *BUILD: Binary Understanding and Integration Logic for Dependencies*, LLNL Strategic Initiative.
Laboratory Directed Research and Development (LDRD) Program, \$2.3M/year.
- 2019-present **PI**, *Packagng Technologies Project.*, Programmatic, \$1.9M/year.
Exascale Computing Project, L4 project.
- 2016-present **PI**, *DevRAMP: Reproducibility, Analysis, Monitoring, and Performance for Developers*, Programmatic, \$1.4M/year.
Department of Energy (DOE) Advanced Simulation and Computing (ASC) ATDM Program.
- 2017-2018 **PI**, *LLNL Early and Mid-Career Recognition Award*, 20% time.
Recognized by LLNL and funded to work on expanding the Spack community.
- 2017-2019 **Co-PI**, *IDEAS-ECP: Softare Productivity for the Exascale Computing Project*, PIs Mike Heroux and Lois Curfman McInnes, \$225k/yr.
U.S. Exascale Computing Project
- 2014-2019 **PI**, *Statistical Methods for Exascale Performance Modeling*, \$2,500,000.
U.S. Department of Energy (DOE) Early Career Research Program.
- 2010-2019 **PI**, *Performance Analysis and Visualization at Exascale (PAVE)*, \$500K/year.
Department of Energy (DOE) Advanced Simulation and Computing (ASC) Program.
- 2015-2016 **PI**, *Containers and Secure Services for HPC*, \$100k.
LLNL ISCP Project
- 2014-2017 **PI**, *Sonar: An Analytics Cluster for Livermore Computing Performance Data*, \$100k/yr.
LLNL Computation TechBase Project
- 2013-2015 **Co-PI**, *Performance Insight for Programs and Exascale Runtimes (PIPER)*, \$4.5M.
PI Martin Schulz, LLNL. DOE Advanced Scientific Computing Research (ASCR) X-Stack 2 Program.
- 2012-2014 **Co-PI**, *Performance Visualization at ExaScale (PeVES)*, \$250k.
PI Timo Bremer, LLNL. University of California Laboratory Fees Research Program.
- 2012-2014 **Co-PI**, *Scalable Topology-Aware Task Embedding (STATE)*, \$945k.
PI Abhinav Bhatele, LLNL. DOE LDRD Exploratory Research.
- 2012 **PI**, *Data Analytics for HPC*, \$100k.

Time Allocation Grants

- 2013-2018 **Co-PI**, *Performance Evaluation and Analysis Consortium (PEAC)*, 120M core-hours.
PI Lenny Oliker, LBNL. DOE Innovative and Novel Computational Impact on Theory and Experiment (INCITE) Program. Provided millions of compute hours for CS researchers on ANL Intrepid, Mira, and Theta; and ORNL Titan systems.
- 2013-2015 **Co-PI**, *Exploring energy efficiency, memory contention and application-tools-hardware co-design on a Xeon Phi cluster*, 1M core-hours. PI Abhinav Bhatele, LLNL. UTK Beacon Project..

Conferences & Workshops Organized

- 2021 **PackagingCon 2021**, *Co-Chair, Co-Founder*.
Virtual Event, 350+ attendees
- 2017 **IEEE CLUSTER**, *General Chair*.
Honolulu, Hawaii, 180+ attendees
- 2014-2016 **JOWOG 34: Applied Computer Science**, *Co-Chair, Co-Founder*.
Scalable algorithms & simulation techniques for LLNL, LANL, Sandia (US) and AWE (UK). 100+ attendees. Continues to be held yearly.
- 2015-2016 **International Workshop on HPC User Support Tools (HUST)**, *Co-Chair*.
Held at SC15-present in cooperation with SIGHPC.
- 2016 **Workshop on Variability in Parallel and Distributed Systems (VarSys)**, *Co-Chair*.
Held at IPDPS'16.
- 2016 **Workshop on Emerging Parallel and Distributed Runtime Systems and Middleware (IPDRM'16)**, *Co-Chair*.
Held at IPDPS'16-17 and SC'20.
- 2011-2012 **Workshop on High-performance Infrastructure for Scalable Tools (WHIST)**, *Chair*.
Held twice at ICS. Proceedings published in special issue of ParCo Journal. 30 attendees.

Conference Committees

- 2021 **ACM/IEEE Supercomputing Conference (SC21)**, *Technical Program Vice Chair*.
- 2019 **ACM/IEEE Supercomputing Conference (SC19)**, *Technical Papers Committee*.
- 2019 **ACM Symp. on Principles and Practices of Parallel Programming (PPoPP)**, *Finance Chair*.
- 2018 **ACM/IEEE Supercomputing Conference (SC18)**, *Vice Chair, Technical Papers*.
- 2017 **ACM/IEEE Supercomputing Conference (SC17)**, *Posters Committee*.
- 2017 **Parallel Architectures and Compilation Techniques (PACT)**, *Technical Papers Committee*.
- 2014-2017 **International Workshop on HPC User Support Tools (HUST)**, *Technical Papers*.
- 2014-2017 **International Workshop on Visual Performance Analytics (VPA)**, *Technical Papers*.
- 2016 **International Conference on Parallel Processing (ICPP)**, *Technical Papers Committee*.
- 2012-2016 **ACM Student Research Competition**, *Review Committee*.
- 2010-2016 **Performance Modeling, Benchmarking, & Simulation (PMBS)**, *Steering Committee*.
- 2015 **IEEE CLUSTER**, *Area Chair, Cluster Design and Architecture*.
- 2015 **PPoPP 2015**, *External Program Committee*.
- 2013,2014 **ACM/IEEE Supercomputing Conference (SC13,SC14)**, *Technical Papers Committee*.
- 2011,2014 **ACM International Conference on Supercomputing (ICS)**, *Technical Papers Committee*.
- 2014 **IEEE International Parallel and Distributed Processing Symposium (IPDPS)**, *Tech. Papers*.
- 2014 **Hardware-Software Co-Design for High Performance Computing (Co-HPC)**, *Technical Papers Committee*.
- 2013-2014 **IEEE/ACM Cluster, Cloud & Grid Computing (CCGrid)**, *Technical Papers Committee*.
- 2013-2014 **IEEE Big Data**, *Technical Papers Committee*.
- 2012-2013 **International Conference on Parallel Processing (ICPP)**, *Technical Papers Committee*.

- 2012 **IEEE CLUSTER**, *Technical Papers Committee*.
- 2012 **ACM/IEEE Supercomputing Conference (SC)**, *Posters Committee*.
- 2009 **Parallel Architectures and Compilation Techniques (PACT)**, *Web Chair*.

Grant Reviewing

- 2021 **Small Business Innovation Research (SBIR)**, *DOE Office of Science*.
- 2019 **Predictive Science Academic Alliance Program (PSAAP III)**, *DOE NNSA*.
- 2015 **Computing and Communication Foundations (CCF)**, *NSF*.
- 2012 **Predictive Science Academic Alliance Program (PSAAP II)**, *DOE NNSA*.
- 2011 **Small Business Innovation Research (SBIR)**, *DOE Office of Science*.

Memberships

- Professional [IEEE](#), [ACM](#), [SIGHPC](#)
- Teaching [Software Carpentry](#) Certified Instructor
- Open Source [OpenHPC](#) Technical Steering Committee (Developer Tools Representative)

Skills

- General Python, C, C++, Fortran (linking), Answer Set Programming (ASP), Clingo, Java, Ruby, Perl, git, svn, CVS, autotools, CMake.
- Logic Programs Answer Set Programming (ASP), Clingo grounder/solver, integration with Python.
- Packaging Dependency solving, build systems, linkers/loaders, compiler runtime systems, software modeling, binary packaging, ELF, binary relocation, large-scale continuous integration.
- High Performance Computing Parallel performance tool development in C/C++. MPI & OpenMP profiling tools, runtime systems, debugger interfaces, stack tracing. Measurement, analysis, and tuning of parallel applications on large clusters (IBM Blue Gene, Cray, Linux). Experience with large physics codes at LLNL.
- Enterprise Web Development Server-side Java: MVC, Tomcat, Apache, Struts, Velocity, Atlassian plugins, Spring Framework; Maven, Ant; JDO, JDBC, SQL with Oracle. Other: Ruby on Rails, Mongrel server, LAMP stack.
- Languages English (fluent), Japanese (conversational)

Software Projects

- [spack](#) De-facto standard for HPC package management. 1,000 contributors, 6,000+ packages. Thousands of users including HPC Centers, research labs, computational scientists, developers.
- [sbang](#) Tool that allows scripts to run with very long shebang lines.
- [cram](#) Tool for running ensembles of millions of small MPI jobs for Uncertainty Quantification.
- [rubik](#) Tool for building optimized task mappings for structured simulation codes on supercomputers.
- [callpath](#) Library for analyzing stack traces scalably in distributed applications.
- [wrap](#) Code generator for MPI profiling wrapper libraries. Used in Allinea DDT.
- [perf-dump](#) Dump per-process, per-thread PAPI counter data from MPI programs to HDF5 files.
- [muster](#) Scalable $O(\log(P))$ parallel clustering library, used in performance/debugging tools.
- [nami](#) Parallel and sequential wavelet compression library.
- [libra](#) Scalable load-balance measurement/viewer tool for MPI simulations.
- [Software developed by students and mentees](#)
- [hatchet](#) Performance data analysis tool, allows Pandas DataFrames to be indexed by call graph nodes.
- [AutomaDeD](#) Massively scalable statistical control flow analysis and performance debugging.
- [Boxfish](#) Visual tool for projecting performance data onto network topologies.
- [MemAxes](#) Visual projection of memory performance data onto processor topologies.
- [Ravel](#) Visualization of MPI traces in logical, Lamport happened-before order.

Invited Talks, Panels, and Podcasts

- December 1, 2021 **NITRD Middleware and Grid Interagency Coordination Team (MAGIC) Meeting**,
Long-term Plans for Spack, Meeting Topic: Sustainable Software.
- November 17, 2021 **Supercomputing 2021 (SC'21)**,
Panel: Performance and Correctness Tools for Extreme-Scale Computing.
- November 10, 2021 **PackagingCon 2021 (SC'21)**,
Panel: Package Manager Convergence: What Stands in the Way?.
- October 21, 2021 **Leadership Scientific Software (lssw.io)**,
Panel: Progress, impediments, priorities and gaps in leadership scientific software.
- May 28, 2021 **CppCast (Podcast)**,
Spack with Todd Gamblin and Greg Becker, Hosts: Rob Irving and Jason Turner.
- May 21, 2020 **Research Software Engineer Stories Podcast**,
Spack Attack!, Host: Vanessa Sochat.
- July 16, 2019 **The Manifest (Podcast)**,
Spack, Hosts: Andrew Nesbitt and Alex Pounds.
- October 26, 2018 **Let's Talk Exascale (Podcast)**,
Spack: The Deployment Tool for ECP's Software Stack, Host: Scott Gibson.
- August 2, 2018 **Argonne Training Program on Extreme-Scale Computing (ATPESC)**,
Dinner Talk: Building things: Spack, Software, and Sustainable Communities in HPC.
- February 7, 2018 **FOSDEM'18 Package Management Devroom**,
Panel: The Future of Package Management.
- February 6, 2018 **The Changelog (Podcast)**,
Moore's Law and High Performance Computing, Hosts: Jerod Santo and Adam Stacoviak.
- July 16, 2017 **Request for Commits (Podcast)**,
Open Source and Supercomputers (Spack), Hosts: Nadia Eghbal and Mikeal Rogers.
- June 2, 2017 **4th International Workshop on Reproducibility in Parallel Computing (REPPAR'17)**,
Keynote: Spack and Reproducibility in HPC, Orlando, FL.
Held at the International Parallel and Distributed Processing Symposium (IPDPS'17)
- February 26, 2016 **Tokyo Institute of Technology**,
Measuring and Analyzing Entire HPC Centers: The Sonar Project at LLNL, Tokyo, Japan.
- January 29, 2016 **University of California, Merced**,
Measuring and Analyzing Entire HPC Centers: The Sonar Project at LLNL, Merced, CA.
- March 26, 2015 **Lawrence Berkeley National Laboratory**,
Visualizing application semantics in performance tools, Berkeley, CA.
- October, 2014 **College of William & Mary**,
Measuring, Analyzing, and Visualizing Memory Performance, Williamsburg, VA.
- August, 2014 **ASCR Workshop on Modeling & Simulation of Systems & Applications**,
Panel: Application Representation, Seattle, WA.
- April, 2013 **University of Illinois, Charm++ Workshop**,
Intuitive Visualizations for Performance Analysis at Scale, Urbana, IL.
- June, 2010 **Tokyo Institute of Technology, Waseda University, Tokyo University**,
Preparing for Exascale: Performance Analysis Research at LLNL, Tokyo, Japan.

Publications

Peer Reviewed Conference & Journal Papers

- [1] Daniel Nichols, Aniruddha Marathe, Kathleen Shoga, **Todd Gamblin**, and Abhinav Bhatele. Resource Utilization Aware Job Scheduling to Mitigate Performance Variability. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'22)*, Virtual Event, May 30-June 3 2022.
- [2] Harshitha Menon, Konstantinos Parasyris, Tom Scogland, and **Todd Gamblin**. Using Active Learning to

- Find High-Fidelity Builds. In *Mining Software Repositories Conference (MSR'22)*, Pittsburgh, PA, May 18-24 2022.
- [3] 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. *IEEE Transactions on Visualization and Computer Graphics*, 2021.
 - [4] Huu Tan Nguyen, Abhinav Bhatele, Nikhil Jain, Suraj P. Kesavan, Harsh Bhatia, **Todd Gamblin**, Kwan-Liu Ma, and Peer-Timo Bremer. Visualizing hierarchical performance profiles of parallel codes using callflow. *IEEE Trans. Vis. Comput. Graph.*, 27(4):2455–2468, 2021.
 - [5] Chad Wood, Giorgis Georgakoudis, David Beckingsale, David Poliakoff, Alfredo Giménez, Kevin A. Huck, Allen D. Malony, and **Todd Gamblin**. Artemis: Automatic runtime tuning of parallel execution parameters using machine learning. In Bradford L. Chamberlain, Ana Lucia Varbanescu, Hatem Ltaief, and Piotr Luszczek, editors, *International Conference on High Performance Computing (ISC'21)*, volume 12728 of *Lecture Notes in Computer Science*, pages 453–472, Virtual Event, June 24 - July 2 2021. Springer.
 - [6] Tao Wang, Nikhil Jain, David Böhme, David Beckingsale, Frank Mueller, and **Todd Gamblin**. Codeseer: input-dependent code variants selection via machine learning. In Eduard Ayguadé, Wen-mei W. Hwu, Rosa M. Badia, and H. Peter Hofstee, editors, *International Conference on Supercomputing (ICS'20)*, pages 43:1–43:11, Barcelona, Spain, June 29-July 2 2020. ACM.
 - [7] Harshitha Menon, Abhinav Bhatele, and **Todd Gamblin**. Auto-Tuning Parameter Choices using Bayesian Optimization. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'20)*, May 18-22 2020.
 - [8] Michael Wyatt, Stephen Herbein, Kathleen Shoga, **Todd Gamblin**, and Michela Taufer. CanarIO: Sounding the Alarm on IO-Related Performance Degradation. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'20)*, May 18-22 2020.
 - [9] Katherine E. Isaacs and **Todd Gamblin**. Preserving Command Line Workflow for a Package Management System using ASCII DAG Visualization. *IEEE Transactions on Visualization and Computer Graphics*, 25(9), 2019.
 - [10] Abhinav Bhatele, Stephanie Brink, and **Todd Gamblin**. Hatchet: Pruning the overgrowth in parallel profiles. In *Supercomputing 2019 (SC'19)*, Denver, CO, November 17-22 2019.
 - [11] Ian Karlin, Yoonho Park, Bronis R. de Supinski, Peng Wang, Bert Still, David Beckingsale, Robert Blake, Tong Chen, Guojing Cong, Carlos H. A. Costa, Johann Dahm, Giacomo Domeniconi, Thomas Epperly, Aaron Fisher, Sara Kokkila Schumacher, Steven H. Langer, Hai Le, Eun Kyung Lee, Naoya Maruyama, Xinyu Que, David F. Richards, Björn Sjögreen, Jonathan Wong, Carol S. Woodward, Ulrike Meier Yang, Xiaohua Zhang, Bob Anderson, David Appelhans, Levi Barnes, Peter D. Barnes Jr., Sorin Bastea, David Böhme, Jamie A. Bramwell, James M. Brase, José R. Brunheroto, Barry Chen, Charway R. Cooper, Tony Degroot, Robert D. Falgout, **Todd Gamblin**, David J. Gardner, James N. Glosli, John A. Gunnels, Max Katz, Tzanio V. Kolev, I-Feng W. Kuo, Matthew P. LeGendre, Ruipeng Li, Pei-Hung Lin, Shelby Lockhart, Kathleen McCandless, Claudia Misale, Jaime H. Moreno, Rob Neely, Jarom Nelson, Rao Nimmakayala, Kathryn M. O'Brien, Kevin O'Brien, Ramesh Pankajakshan, Roger Pearce, Slaven Peles, Phil Regier, Steve Rennich, Martin Schulz, Howard Scott, James C. Sexton, Kathleen Shoga, Shiv Sundram, Guillaume Thomas-Collignon, Brian Van Essen, Alexey Voronin, Bob Walkup, Lu Wang, Chris Ward, Hui-Fang Wen, Daniel A. White, Christopher Young, Cyril Zeller, and Edward Zywickz. Preparation and optimization of a diverse workload for a large-scale heterogeneous system. In *Supercomputing 2019 (SC'19)*, Denver, CO, November 17-22 2019.
 - [12] Tao Wang, Nikhil Jain, David Böhme, Frank Mueller, and **Todd Gamblin**. Funcytuner: Auto-tuning scientific applications with per-loop compilation. In *International Conference on Parallel Processing (ICPP'19)*, Tsukuba, Japan, August 5-8 2019.
 - [13] Abhinav Bhatele, Nikhil Jain, Misbah Mubarak, and **Todd Gamblin**. Analyzing Cost-Performance Tradeoffs of HPC Network Designs under Different Constraints using Simulations. In *SIGSIM PADS 2019*, Chicago, IL, June 3-5 2019.
 - [14] Prasanna Balaprakash, Jack Dongarra, **Todd Gamblin**, Mary Hall, Jeffrey K. Hollingsworth, Boyana Norris, and Rich Vuduc. Autotuning in High-Performance Computing Applications. *Proceedings of the IEEE*, pages 1–16, 2018.
 - [15] 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, 2018.

- [16] Michael R. Wyatt II, Stephen N. Herbein, **Todd Gamblin**, Adam Moody, Dong H. Ahn, and Michela Taufer. PRIONN: Predicting Runtime and IO using Neural Networks. In *International Conference on Parallel Processing (ICPP'18)*, Eugene, OR, August 13-16 2018. LLNL-CONF-739253.
- [17] 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 *International Conference on Supercomputing (ICS'18)*, Beijing, China, June 12-15 2018. LLNL-CONF-750296.
- [18] 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 International Parallel and Distributed Processing Symposium (IPDPS'18)*, Vancouver, BC, Canada, May 21-25 2018. LLNL-CONF-740303.
- [19] 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 *Supercomputing 2017 (SC'17)*, Denver, CO, November 13-16 2017.
- [20] Nikhil Jain, Abhinav Bhatele, Louis Howell, David Böhme, Ian Karlin, Edgar Leon, Misbah Mubarak, Noah Wolfe, **Todd Gamblin**, and Matthew Leininger. Predicting the Impact of Fat-tree Configurations on Applications and Multi-job Workloads. In *Supercomputing 2017 (SC'17)*, Denver, CO, November 13-16 2017.
- [21] Aniruddha Marathe, Nikhil Jain, Rushil Anirudh, Abhinav Bhatele, Jayaraman J. Thiagarajan, Bhavya Kailkhura, Jae-Seung Yeom, Barry Rountree, and **Todd Gamblin**. Power-aware Parameter Tuning of HPC Applications using Machine Learning. In *Supercomputing 2017 (SC'17)*, Denver, CO, November 13-16 2017.
- [22] David Beckingsale, Olga Pearce, Ignacio Laguna, and **Todd Gamblin**. Apollo: Fast, Dynamic Tuning for Data-Dependent Code. In *IEEE International Parallel & Distributed Processing Symposium (IPDPS'17)*, Orlando, FL, USA, May 29-June 2 2017. LLNL-CONF-723337.
- [23] Nikhil Jain, Abhinav Bhatele, Xiang Ni, **Todd Gamblin**, and Laxmikant V. Kale. Express Mesh: A Partitionable Low-Diameter Network that Eliminates Inter-job Interference. In *IEEE International Parallel & Distributed Processing Symposium (IPDPS'17)*, Orlando, FL, USA, May 29-June 2 2017. LLNL-CONF-706801.
- [24] Hao Xu, Shasha Wen, Alfredo Gimenez, **Todd Gamblin**, and Xu Liu. DR-BW: Identifying Bandwidth Contention in NUMA Architectures with Supervised Learning. In *IEEE International Parallel & Distributed Processing Symposium (IPDPS'17)*, Orlando, FL, USA, May 29-June 2 2017.
- [25] David Böhme, **Todd Gamblin**, David Beckingsale, Peer-Timo Bremer, Alfredo Giménez, Matthew LeGendre, Olga Pearce, and Martin Schulz. Caliper: Performance Introspection for HPC Software Stacks. In *Supercomputing 2016 (SC'16)*, Salt Lake City, UT, November 13-18 2016. LLNL-CONF-699263.
- [26] 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 *Supercomputing 2016 (SC'16)*, Salt Lake City, UT, November 13-18 2016. LLNL-CONF-696018.
- [27] Nikhil Jain, Abhinav Bhatele, Sam White, **Todd Gamblin**, and Laxmikant V. Kale. Evaluating HPC Networks via Simulation of Parallel Workloads. In *Supercomputing 2016 (SC'16)*, Salt Lake City, UT, November 13-18 2016. LLNL-CONF-690662.
- [28] Ryan McKenna, Stephen N. Herbein, Adam Moody, **Todd Gamblin**, and Michela Taufer. Forecasting traffic to parallel filesystems from hpc clusters. In *IEEE Cluster 2016 (CLUSTER'16)*, Taipei, Taiwan, September 12-16 2016. LLNL-CONF-691797.
- [29] Olga Pearce, **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, and Nancy M. Amato. MPMD Framework for Offloading Load Balance Computation. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'16)*, Chicago, IL, May 23-27, 2016 2016. LLNL-CONF-665393.
- [30] Ignacio Laguna, David F. Richards, **Todd Gamblin**, Martin Schulz, Bronis R. de Supinski, Kathryn Mohror, and Howard Pritchard. Evaluating and Extending User-Level Fault Tolerance in MPI. *International Journal of High Performance Computing Applications (IJHPCA)*, January 11 2016. LLNL-JRNL-663434.
- [31] **Todd Gamblin**, Matthew P. LeGendre, Michael R. Collette, Gregory L. Lee, Adam Moody, Bronis R. de Supinski, and W. Scott Futral. **The Spack Package Manager: Bringing order to HPC software chaos.** In *Supercomputing 2015 (SC'15)*, Austin, Texas, November 15-20 2015. LLNL-CONF-669890, (22% accepted).
- [32] Katherine E. Isaacs, Abhinav Bhatele, Jonathan Lifflander, David Böhme, **Todd Gamblin**, Martin Schulz, Bernd Hamann, and Peer-Timo Bremer. Analyzing the structure of execution traces from task-based parallel

runtimes. In *Supercomputing 2015 (SC'15)*, Austin, Texas, November 15-20 2015. LLNL-CONF-670046, (22% accepted).

- [33] Ignacio Laguna, Dong H. Ahn, Bronis R. de Supinski, **Todd Gamblin**, Gregory L. Lee, Martin Schulz, Saurabh Bagchi, Milind Kulkarni, Bowen Zhou, Zhezhe Chen, and Feng Qin. Debugging high-performance computing applications at massive scales. *Communications of the ACM*, September 2015. LLNL-JRNL-652400.
- [34] 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 International Parallel and Distributed Processing Symposium (IPDPS'15)*, Hyderabad, INDIA, May 25-29 2015. LLNL-CONF-663150, (21.8% accepted).
- [35] 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 (TPDS)*, 27(3):829–840, March 2015.
- [36] Abhinav Bhatele, Nikhil Jain, Katherine 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 *International Conference on High Performance Computing (HiPC'14)*, Goa, India, December 17-20 2014. LLNL-CONF-655465.
- [37] Alfredo Giménez, **Todd Gamblin**, Barry Rountree, Abhinav Bhatele, Ilir Jusufi, Bernd Hamann, and Peer-Timo Bremer. [Dissecting On-Node Memory Access Performance: A Semantic Approach](#). In *Supercomputing 2014 (SC'14)*, New Orleans, LA, November 16-21 2014. LLNL-CONF-658626, (20.8% accepted).
- [38] 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](#). In *IEEE Symposium on Information Visualization (INFOVIS'14)*, Paris, France, November 9-14 2014. LLNL-JRNL-657418, (23% accepted).
- [39] Ignacio Laguna, David F. Richards, **Todd Gamblin**, Martin Schulz, and Bronis R. de Supinski. [Evaluating User-Level Fault Tolerance for MPI Applications](#). In *EuroMPI/Asia 2014*, Kyoto, Japan, September 9-12 2014. LLNL-CONF-656877, (46% accepted).
- [40] Olga Pearce, **Todd Gamblin**, Bronis R. de Supinski, Tom Arsenlis, and Nancy M. Amato. Load Balancing N-body Simulations with Highly Non-Uniform Density. In *International Conference on Supercomputing (ICS'14)*, Munich, Germany, June 10-13 2014. LLNL-CONF-648577, (20.9% accepted).
- [41] Katherine E. Isaacs, Abhinav Bhatele, Peer-Timo Bremer, **Todd Gamblin**, Alfredo Gimenez, Bernd Hamann, Ilir Jusufi, and Martin Schulz. [State of the Art of Performance Visualization](#). In R. Borgo, R. Maciejewski, and I. Viola, editors, *Eurographics Conference on Visualization (EuroVis'14)*, Swansea, Wales, UK, June 9-13 2014. LLNL-CONF-652873.
- [42] Subrata Mitra, Ignacio Laguna, Dong H. Ahn, Saurabh Bagchi, Martin Schulz, and **Todd Gamblin**. Accurate application progress analysis for large-scale parallel debugging. In *Programming Language Design and Implementation (PLDI'14)*, Edinburgh, UK, June 9-11 2014. LLNL-CONF-646258, (18% accepted).
- Best Paper** [43] Kento Sato, Kathryn Mohror, Adam Moody, **Todd Gamblin**, Bronis R. de Supinski, Naoya Maruyama, and Satoshi Matsuoka. A user-level infiniband-based file system and checkpoint strategy for burst buffers. In *IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'14)*, Chicago, IL, May 26-29, 2014. LLNL-CONF-645876, (19% accepted).
- [44] Ignacio Laguna, Dong H. Ahn, Bronis R. de Supinski, Saurabh Bagchi, and **Todd Gamblin**. Diagnosis of performance faults in large scale mpi applications via probabilistic progress-dependence inference. *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, May 22 2014. LLNL-JRNL-643939.
- [45] Kento Sato, Adam Moody, Kathryn Mohror, **Todd Gamblin**, Bronis R. de Supinski, Naoya Maruyama, and Satoshi Matsuoka. FMI: Fault Tolerant Messaging Interface for Fast and Transparent Recovery. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'14)*, Phoenix, AZ, May 19-23 2014. LLNL-CONF-645209, (21% accepted).
- [46] Nikhil Jain, Abhinav Bhatele, Michael P. Robson, **Todd Gamblin**, and Laxmikant V. Kale. [Predicting application performance using supervised learning on communication features](#). In *Supercomputing 2013 (SC'13)*, Denver, CO, November 17-22 2013. LLNL-CONF-556491, (20% accepted).
- [47] Lukasz G. Szafaryn, **Todd Gamblin**, Bronis R. de Supinski, and Kevin Skadron. [Trellis: Portability across architectures with a high-level framework](#). *J. Parallel Distrib. Comput.*, 73(10):1400–1413, July 18 2013.

- Best Paper** [48] Wolfgang Frings, Dong H. Ahn, Matthew LeGendre, **Todd Gamblin**, Bronis R. de Supinski, and Felix Wolf. [Massively parallel loading](#). In *International Conference on Supercomputing (ICS'13)*, Eugene, OR, June 10-14 2013. LLNL-CONF-610893, (21% accepted).
- [49] Dong H. Ahn, Michael J. Brim, Bronis R. de Supinski, **Todd Gamblin**, Gregory L. Lee, Matthew P. LeGendre, Barton P. Miller, Adam Moody, and Martin Schulz. [Efficient and Scalable Retrieval Techniques for Global File Properties](#). In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'13)*, Boston, MA, May 20-24 2013. LLNL-PROC-554055, (21% accepted).
- [50] Barry Rountree, **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, David K. Lowenthal, Guy Cobb, and Henry Tufo. [Parallelizing heavyweight debugging tools with MPIEcho](#). *Parallel Computing*, 39(3):156–166, March 2013.
- [51] 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 *Supercomputing 2012 (SC'12)*, Salt Lake City, UT, November 10-16 2012. LLNL-CONF-554552, (21% accepted).
- [52] 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 *Supercomputing 2012 (SC'12)*, Salt Lake City, UT, November 10-16 2012. LLNL-CONF-556491, (21% accepted).
- [53] Kento Sato, Adam Moody, Kathryn Mohror, **Todd Gamblin**, Bronis R. de Supinski, Naoya Maruyama, and Satoshi Matsuoka. [Design and modeling of a non-blocking checkpointing system](#). In *Supercomputing 2012 (SC'12)*, Salt Lake City, UT, November 10-16 2012. LLNL-CONF-554431, (21% accepted).
- [54] Aaditya G. Landge, Joshua A. Levine, Katherine E. Isaacs, Abhinav Bhatele, **Todd Gamblin**, Martin Schulz, Steven H. Langer, Peer-Timo Bremer, and Valerio Pascucci. [Visualizing network traffic to understand the performance of massively parallel simulations](#). In *IEEE Symposium on Information Visualization (INFOVIS'12)*, Seattle, WA, October 14-19 2012. LLNL-CONF-543359, (25% accepted).
- [55] Ignacio Laguna, Dong H. Ahn, Bronis R. de Supinski, Saurabh Bagchi, and **Todd Gamblin**. [Probabilistic diagnosis of performance faults in large scale parallel applications](#). In *International Conference on Parallel Architectures and Compilation Techniques (PACT'12)*, Minneapolis, MN, September 19-23 2012. LLNL-PROC-548642, (18% accepted).
- [56] Olga Pearce, **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, and Nancy M. Amato. [Quantifying the Effectiveness of Load Balance Algorithms](#). In *International Conference on Supercomputing (ICS'12)*, Venice, Italy, June 25-29 2012. LLNL-CONF- 523343, (22.3% accepted).
- [57] Spyros Lyberis, Polyvios Pratikakis, Dimitrios Nikolopoulos, Martin Schulz, **Todd Gamblin**, and Bronis de Supinski. [The myrmics memory allocator: Hierarchical, message-passing allocation for global address spaces](#). In *International Symposium on Memory Management (ISMM'12)*, Beijing, China, June 15-16 2012. LLNL-CONF-545875.
- [58] Sandeep Budanur, Frank Mueller, and **Todd Gamblin**. [Memory Trace Compression and Replay for SPMD Systems using Extended PRSDs](#). *Computer Journal*, 55(2):206–217, February 2012. LLNL-CONF-460598.
- [59] Allison H. Baker, Robert D. Falgout, **Todd Gamblin**, Tzanio V. Kolev, Martin Schulz, and Ulrike Meier Yang. [Scaling Algebraic Multigrid Solvers: On the Road to Exascale](#). *Competence in High Performance Computing (CiHPC)*, Dec 2011. LLNL-PROC-463941.
- [60] Ignacio Laguna, **Todd Gamblin**, Bronis R. de Supinski, Saurabh Bagchi, Greg Bronevetsky, Dong H. Ahn, Martin Schulz, and Barry Rountree. [Large scale debugging of parallel tasks with AutomaDeD](#). In *Supercomputing 2011 (SC'11)*, Seattle, WA, November 12-18 2011. LLNL-CONF-486911, (21% accepted).
- [61] Martin Schulz, Joshua A. Levine, Peer-Timo Bremer, **Todd Gamblin**, and Valerio Pascucci. [Interpreting performance data across intuitive domains](#). In *International Conference on Parallel Processing (ICPP'11)*, Taipei, Taiwan, September 13-16 2011. LLNL-CONF-476091, (22% accepted).
- [62] Allison H. Baker, **Todd Gamblin**, Martin Schulz, and Ulrike Meier Yang. [Challenges of Scaling Algebraic Multigrid across Modern Multicore Architectures](#). In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'11)*, Anchorage, Alaska, May 16-20 2011. LLNL-CONF-458074, (19.6% accepted).
- [63] Zoltan Szebenyi, **Todd Gamblin**, Martin Schulz, Bronis R. de Supinski, Felix Wolf, and Brian J. N. Wylie. [Reconciling sampling and direct instrumentation for unintrusive call-path profiling of MPI programs](#). In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'11)*, Anchorage, Alaska, May 16-20 2011. LLNL-CONF-468224, (19.6% accepted).

- [64] **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, Robert J. Fowler, and Daniel A. Reed. [Clustering performance data efficiently at massive scales](#). In *International Conference on Supercomputing (ICS'10)*, Tsukuba, Japan, June 1-4 2010. LLNL-CONF-422684, (18.5% accepted).
- [65] Bronis R. de Supinski, Sadaf Alam, David H. Bailey, Laura Carrington, Chris Daley, Anshu Dubey, **Todd Gamblin**, Dan Gunter, Paul D. Hovland, Heike Jagode, Karen Karavanic, Gabriel Marin, John Mellor-Crummey, Shirley Moore, Boyana Norris, Leonid Oliker, Catherine Olschanowsky, Philip C. Roth, Martin Schulz, Sameer Shende, Allan Snaveley, Wyatt Spear, Mustafa Tikir, Jeff Vetter, Pat Worley, and Nicholas Wright. [Modeling the office of science ten year facilities plan: The PERI architecture tiger team](#). *Journal of Physics: Conference Series, SciDAC 2009*, 2009. LLNL-CONF-413427.
- [66] Robert J. Fowler, Laksono Adhianto, Bronis R. de Supinski, Mike Fagan, **Todd Gamblin**, Mark Krentel, John Mellor-Crummey, Martin Schulz, and Nathan Tallent. [Frontiers of performance analysis on leadership class systems](#). *Journal of Physics: Conference Series, SciDAC 2009*, 2009. LLNL-CONF-413969.
- [67] **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, Robert J. Fowler, and Daniel A. Reed. [Scalable load-balance measurement for SPMD codes](#). In *Supercomputing 2008 (SC'08)*, Austin, Texas, November 15-21 2008. LLNL-CONF-406045, (21.3% accepted).
- [68] Robert J. Fowler, **Todd Gamblin**, Allan K. Porterfield, Patrick Dreher, Song Huang, and Balint Joo. [Performance engineering challenges: The view from RENCI](#). *Journal of Physics: Conference Series, SciDAC 2008*, 125:5, August 2008.
- [69] **Todd Gamblin**, Robert J. Fowler, and Daniel A. Reed. [Scalable methods for monitoring and detecting behavioral classes in scientific codes](#). In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS'08)*, Miami, FL, April 14-18 2008. (25.6% accepted).

Workshop Papers & Technical Reports

- [70] C. Liao, A. Wang, Giorgis Georgakoudis, B. de Supinski, Y. Yan, D. Beckingsale, and **Todd Gamblin**. [Machine Learning-Driven Adaptation of OpenMP Applications](#). In *Eighth Workshop on Accelerator Programming Using Directives (WACCPD'21)*, Virtual Event, November 14 2021. LLNL-CONF-826432.
- [71] 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 *IEEE/ACM International Workshop on HPC User Support Tools and Workshop on Programming and Performance Visualization Tools (HUST/ProTools'20)*, pages 49–58, Atlanta, GA, November 18 2020. IEEE.
- [72] Tal Ben-Nun, **Todd Gamblin**, Daisy S. Hollman, Hari Krishnan, and Chris J. Newburn. [Workflows are the new applications: Challenges in performance, portability, and productivity](#). In *IEEE/ACM International Workshop on Performance, Portability and Productivity in HPC (P3HPC)*, pages 57–69. IEEE, November 13 2020.
- [73] Massimiliano Culpo, Gregory Becker, Carlos Eduardo Arango Gutierrez, Kenneth Hoste, and **Todd Gamblin**. [archspec: A library for detecting, labeling, and reasoning about microarchitectures](#). In *2nd International Workshop on Containers and New Orchestration Paradigms for Isolated Environments in HPC (CANOPIE-HPC)*, pages 45–52, Atlanta, GA, November 12 2020. IEEE.
- [74] **Todd Gamblin**. [Software Integration Challenges](#). In *Collegeville Workshop 2020*, Virtual Event, July 21-23 2020.
- [75] Samuel Knight, Jeremiah Wilke, and **Todd Gamblin**. [Using Malleable Task Scheduling to Accelerate Package Manager Installations](#). In *Workshop on HPC User Support Tools (HUST-19)*, Denver, CO, November 18 2019.
- [76] Mario Melara, **Todd Gamblin**, Gregory Becker, Robert French, Matt Belhorn, Kelly Thompson, Peter Scheibel, and Rebecca Hartman-Baker. [Using Spack to Manage Software on Cray Supercomputers](#). In *Cray Users Group (CUG 2017)*, Seattle, WA, May 7-11 2017.
- [77] Gregory Becker, Peter Scheibel, Matthew P. LeGendre, and **Todd Gamblin**. [Managing Combinatorial Software Installations with Spack](#). In *Second International Workshop on HPC User Support Tools (HUST'16)*, Salt Lake City, UT, November 13 2016.
- [78] Benafsh Husain, Alfredo Giménez, Joshua A. Levine, **Todd Gamblin**, and Peer-Timo Bremer. [Relating memory performance data to application domain data using an integration API](#). In *Workshop on Visual Performance Analytics (VPA'16)*, Salt Lake City, UT, November 13 2016.
- [79] Huu Tan Nguyen, Abhinav Bhatele, Peer-Timo Bremer, **Todd Gamblin**, Martin Schulz, Lai Wei, David Boehme, and Kwan-Liu Ma. [VIPACT: A Visualization Interface for Analyzing Calling Context Trees](#). In *Workshop on Visual Performance Analytics (VPA'16)*, Salt Lake City, UT, November 13 2016.

- [80] Chad Wood, Sudhanshu Sane, Daniel Ellsworth, Alfredo Gimenez, Kevin Huck, **Todd Gamblin**, and Allen Malony. A Scalable Observation System for Introspection and In Situ Analytics. In *Workshop on Extreme-Scale Programming Tools (ESPT'16)*, Salt Lake City, UT, November 13 2016.
- [81] Subrata Mitra, Suhas Javagal, Amiya K. Maji, **Todd Gamblin**, Adam Moody, Stephen Harrell, and Saurabh Bagchi. *A Study of Failures in Community Clusters: The Case of Conte*. In *International Workshop on Program Debugging (IWPD'16)*, Ottawa, Canada, October 23-27 2016.
- [82] Ignacio Laguna, **Todd Gamblin**, Kathryn Mohror, and Martin Schulz. A global exception fault tolerance model for mpi. In *Workshop on Exascale MPI (ExaMPI14), held at Supercomputing 2014 (SC'14)*, New Orleans, LA, November 17 2014. LLNL-CONF-XXXXXX.
- [83] **Todd Gamblin**, Martin Schulz, Timo Bremer, and Abhinav Bhatele. Building models from measurement with multi-domain correlation. In *DOE Workshop on Modeling & Simulation of Exascale Systems & Applications (ModSim 2013)*, Seattle, WA, September 18-19 2013. LLNL-CONF-639760.
- [84] 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 (ExaOSR)*, Washington, DC, October 2012. LLNL-MI-563334.
- [85] Kento Sato, Adam Moody, Kathryn Mohror, **Todd Gamblin**, Bronis R. de Supinski, Naoya Maruyama, and Satoshi Matsuoka. Towards a light-weight non-blocking checkpointing system. In *HPC in Asia Workshop 2012*, Hamburg, Germany, June 17 2012.
- [86] Kento Sato, Adam Moody, Kathryn Mohror, **Todd Gamblin**, Bronis R. de Supinski, Naoya Maruyama, and Satoshi Matsuoka. Design and modeling of a non-blocking checkpoint system. In *ATIP - A*CRC Workshop on Accelerator Technologies in High Performance Computing*, May 7-10 2012.
- [87] 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 Users Group (CUG 2012)*, Stuttgart, Germany, April 29-May 3 2012. LLNL-PROC-547711.
- [88] Kento Sato, Adam Moody, Kathryn Mohror, **Todd Gamblin**, Bronis R. de Supinski, Naoya Maruyama, and Satoshi Matsuoka. Towards an asynchronous checkpointing system. In *IPSI SIG Technical Reports 2011-ARC-197 2011-HPC-132 (HOKKE-19)*, November 2011. LLNL-CONF-509152.
- [89] Martin Schulz, Abhinav Bhatele, Peer-Timo Bremer, **Todd Gamblin**, Katherine Isaacs, Joshua Levine, and Valerio Pascucci. Creating a Tool Set for Optimizing Topology-aware Node Mappings. In *5th ZIH Parallel Tools Workshop*, Dresden, Germany, September 26-27 2011. LLNL-CONF-402937.
- [90] **Todd Gamblin**, Martin Schulz, Peer-Timo Bremer, Joshua A. Levine, and Valerio Pascucci. Intuitive performance visualization techniques for topological analysis on capability machines. In *Kagoshima Summer United Workshops on Parallel, Distributed, and Cooperative Processing*, Kagoshima, Japan, July 26-29 2011. LLNL-CONF- 490044.
- [91] Barry Rountree, Guy Cobb, **Todd Gamblin**, Martin Schulz, Bronis R. de Supinski, and Henry Tufo. Parallelizing heavyweight debugging tools with MPlecho. In *Workshop on High-performance Infrastructure for Scalable Tools (WHIST'11)*, Tucson, Arizona, June 4 2011. LLNL-CONF-481092.
- [92] Lukasz G. Szafaryn, **Todd Gamblin**, Bronis R. de Supinski, and Kevin Skadron. Experiences with achieving portability across heterogeneous architectures. In *Workshop on Domain-Specific Languages and High-Level Frameworks for High Performance Computing (WOLFHPC'11)*, Tucson, Arizona, May 31 2011. LLNL-CONF-486291.
- [93] Sandeep Budanur, Frank Mueller, and **Todd Gamblin**. *Memory Trace Compression and Replay for SPMD Systems using Extended PRSDs*. In *First International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computing Systems (PMBS'10)*, New Orleans, LA, November 15 2010. LLNL-CONF-460598.
- [94] Frank Mueller, Xing Wu, Martin Schulz, **Todd Gamblin**, and Bronis R. de Supinski. *ScalaTrace: Tracing, analysis and modeling of HPC codes at scale*. In *Para 2010*, Reykjavik, Iceland, June 6-9 2010. LLNL-CONF-427005.
- [95] Bronis R. de Supinski, Robert J. Fowler, **Todd Gamblin**, Frank Mueller, Martin Schulz, and Prasun Ratn. *An open framework for scalable, reconfigurable analysis*. In *International Workshop on Scalable Tools for High-End Computing (STHEC)*, Kos, Greece, June 7-12 2008. LLNL-CONF-403954.
- [96] Francisco Chinchilla, **Todd Gamblin**, Morten Sommervoll, and Jan F. Prins. *Parallel N-body simulation using GPUs*. Technical Report TR04-032, University of North Carolina at Chapel Hill, December 2004.

Tutorials

- [97] Gregory Becker, Robert Blake, Massimiliano Culpo, Tamara Dahlgren, Adam J. Stewart, Harmen Stoppels, and **Todd Gamblin**. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Supercomputing 2021*, St. Louis, MO, November 14 2021. Full day.
- [98] Gregory Becker, Robert Blake, Tamara Dahlgren, Peter Scheibel, and **Todd Gamblin**. [Tutorial: Using Spack to Accelerate Developer Workflows](#). In *Exascale Computing Project 5th Annual Meeting*, Virtual Event, April 16 2021. Full day.
- [99] **Todd Gamblin**, Gregory Becker, Massimiliano Culpo, Mario Melara, Peter Scheibel, and Adam J. Stewart. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Supercomputing 2020*, Virtual Event, November 9-10 2020. Full day.
- [100] **Todd Gamblin**, Gregory Becker, and Sean Smith. [Spack Tutorial on AWS ParallelCluster](#). In *AWS Workshop Studio*, Virtual Event, July 28-29 2020. Over 140 attendees.
- [101] **Todd Gamblin** and Gregory Becker. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Exascale Computing Project 4th Annual Meeting*, Houston, TX, February 6 2020. Full day.
- [102] Adam J. Stewart, Massimiliano Culpo, Gregory Becker, Peter Scheibel, and **Todd Gamblin**. [Spack Community BoF](#). In *Supercomputing 2019*, Denver, CO, November 21 2019.
- [103] **Todd Gamblin**, Gregory Becker, Massimiliano Culpo, Mario Melara, Peter Scheibel, and Adam J. Stewart. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Supercomputing 2019*, Denver, CO, November 18 2019. Full day.
- [104] Gregory Becker and **Todd Gamblin**. [Tutorial: Spack for Developers](#). In *Los Alamos National Laboratory*, Los Alamos, NM, November 5 2019. Half day.
- [105] Levi Baber, Adam J. Stewart, Gregory Becker, and **Todd Gamblin**. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Practice and Experience in Advanced Research Computing (PEARC'19)*, Chicago, IL, July 31 2019. Half day.
- [106] **Todd Gamblin** and Gregory Becker. [Spack Tutorial](#). In *1st Workshop on NSF and DOE High Performance Computing Tools*, Eugene, OR, July 10-11 2019. University of Oregon.
- [107] **Todd Gamblin**, Gregory Becker, Massimiliano Culpo, and Michael Kühn. [Tutorial: Managing HPC Software Complexity with Spack](#). In *ISC High Performance*, Houston, TX, June 16 2019. Half day.
- [108] **Todd Gamblin**, Gregory Becker, Matthew P. LeGendre, and Peter Scheibel. [Spack Roundtable Discussion](#). In *Exascale Computing Project 3rd Annual Meeting*, Houston, TX, January 16 2019.
- [109] **Todd Gamblin**, Gregory Becker, Peter Scheibel, Matt Legendre, and Mario Melara. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Exascale Computing Project 3rd Annual Meeting*, Houston, TX, January 14 2019. Full day.
- [110] **Todd Gamblin**, Adam J. Stewart, Johannes Albert von der Gönna, Marc Pérache, Matt Belhorn, and Veronica Vergara. [Spack Community BoF](#). In *Supercomputing 2018*, Dallas, TX, November 13 2018.
- [111] **Todd Gamblin**, Gregory Becker, Massimiliano Culpo, Gregory L. Lee, Matt Legendre, Mario Melara, and Adam J. Stewart. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Supercomputing 2018*, Dallas, TX, November 12 2018. Full day.
- [112] **Todd Gamblin**, Gregory Becker, Peter Scheibel, Matt Legendre, and Mario Melara. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Exascale Computing Project 2nd Annual Meeting*, Knoxville, TN, February 6-8 2018. Half day.
- [113] **Todd Gamblin**, William Scullin, Matt Belhorn, Mario Melara, and Gerald Ragghianti. [Spack State of the Union](#). In *Exascale Computing Project 2nd Annual Meeting*, Knoxville, TN, February 6-8 2018.
- [114] **Todd Gamblin**, Gregory Becker, Massimiliano Culpo, Gregory L. Lee, Matt Legendre, Mario Melara, and Adam J. Stewart. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Supercomputing 2017*, Salt Lake City, Utah, November 13 2017. Full day.
- [115] **Todd Gamblin**. [Tutorial: Managing HPC Software Complexity with Spack](#). In *HPC Knowledge Meeting (HPCKP'17)*, San Sebastián, Spain, June 16 2017. 2 hours.
- [116] Gregory Becker, Matt Legendre, and **Todd Gamblin**. [Tutorial: Spack for HPC](#). Livermore Computing, Lawrence Livermore National Laboratory, Livermore, CA, April 6 2017. Half day.
- [117] **Todd Gamblin**, Massimiliano Culpo, Gregory Becker, Matt Legendre, Greg Lee, Elizabeth Fischer, and Benedikt Hegner. [Tutorial: Managing HPC Software Complexity with Spack](#). In *Supercomputing 2016*, Salt Lake City, Utah, November 13 2016. Half day.

Special Issues, Book Chapters, Other Articles

- [118] Michael A. Heroux, Lois Curfman McInnes, Rajeev Thakur, Jeffrey S. Vetter, Sherry Li, James Ahrens, Todd Munson, Kathryn Mohror, Timothy Germann, Ryan Adamson, **Todd Gamblin**, Sameer Shende, and James Willenbring. [ECP Response to the “Stewardship of Software for Scientific and High-Performance Computing” RFI](#). Technical report, U.S. Exascale Computing Project, December 13 2021.
- [119] Jeff Hittinger, Ghaleb Abdulla, Dong Ahn, Charity Follett, Jim Foraker, Scott Futral, **Todd Gamblin**, Marisol Gamboa, Michael Goldman, Kyle Halliday, Judith Hill, Tzanio Kolev, Ana Kupresanin, Ignacio Laguna, Ian Lee, Matthew LeGendre, Edgar Leon, Thomas Mendoza, Mark Miller, Kathryn Mohror, Dan Quinlan, Becky Springmeyer, and Vanessa Sochat. [LLNL Response to the DOE ASCR RFI, “Stewardship of Software for Scientific and High-Performance Computing”](#). Technical report, Lawrence Livermore National Laboratory, December 13 2021. LLNL-TR-830011.
- [120] **Todd Gamblin**. Picking an Open Source License at LLNL: Guidance and Recommendations from the Computing Directorate. Technical report, Lawrence Livermore National Laboratory, May 7 2021. LLNL-TR-822365.
- [121] [Packaging a wallop: Lawrence Livermore National Laboratory’s time-saving HPC tool eases the way for next era of scientific simulations](#). *ASCR Discovery*, August 2016.
- [122] 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 Michael Bader, Arndt Bode, Hans-Joachim Bungartz, Michael Gerndt, Gerhard R. Joubert, and Frans Peters, editors, *Parallel Computing: Accelerating Computational Science and Engineering*, volume 25 of *Advances in Parallel Computing*, pages 19–32. IOS Press, 2014. LLNL-CONF-645777.
- [123] Nathan R. Tallent and **Todd Gamblin**, editors. *Special Issue on High-performance Infrastructure for Scalable Tools*, volume 39(3) of *Parallel Computing: Systems & Applications*. March 2013.
- [124] **Todd Gamblin**. [Collaboration Leads to Advances in Power Grid Modeling](#). In *2010 Computation Annual Report*, page 32. Lawrence Livermore National Laboratory, 2010. LLNL-TR-470873.
- [125] Robert J. Fowler, **Todd Gamblin**, Gopi Kandaswamy, Anirban Mandal, Allan K. Porterfield, Lavanya Ramakrishnan, and Daniel A. Reed. Challenges of scale: When all computing becomes grid computing. In Lucio Grandinetti, editor, *High Performance Computing and Grids in Action*, volume 16 of *Advances in Parallel Computing*, pages 186–206. IOS Press, Amsterdam, 2007.

Ph.D. Dissertation

- [126] **Todd Gamblin**. [Scalable Performance Measurement and Analysis](#). PhD thesis, University of North Carolina at Chapel Hill, Chapel Hill, NC, August 2009. LLNL-TH-419482.

Posters & Presentations

- [127] **Todd Gamblin**. [Beyond version solving: implementing general package solvers with Answer Set Programming](#). In *PackagingCon 2021*, Virtual Event, November 9 2021.
- [128] **Todd Gamblin**. [How Can Package Managers Handle ABI \(In\)compatibility in C++?](#) In *Cppcon 2021*, Aurora, CO, October 29 2021.
- [129] **Todd Gamblin**. [Reusing dependencies across ecosystems: what stands in the way?](#) In *Free and Open source Software Developers’ European Meeting (FOSDEM’21)*, Brussels, Belgium, February 6 2021.
- [130] **Todd Gamblin**. [Build for your microarchitecture: experiences with Spack and archspec](#). In *Free and Open source Software Developers’ European Meeting (FOSDEM’20)*, Brussels, Belgium, February 2 2020.
- [131] **Todd Gamblin**. [Spack’s new Concretizer: Dependency solving is more than just SAT!](#) In *Free and Open source Software Developers’ European Meeting (FOSDEM’20)*, Brussels, Belgium, February 1 2020.
- [132] **Todd Gamblin**. [Decluttering HPC Software Chaos with SPACK](#). In *SIAM Conference on Parallel Processing for Scientific Computing (PP’18)*, Minisymposium on Productive Programming using Parallel Models, Tools and Scientific Workflows, Tokyo, Japan, March 7 2018.
- [133] **Todd Gamblin**. [Spack State of the Union](#). In *Exascale Computing Project 2nd Annual Meeting*, Knoxville, TN, February 6-8, 2018 2018. LLNL-PRES-746210.
- [134] **Todd Gamblin**. [Binary packaging for HPC with Spack](#). In *Free and Open source Software Developers’ European Meeting (FOSDEM’18)*, Brussels, Belgium, February 4 2018. LLNL-PRES-745747.
- [135] **Todd Gamblin**. [How compilers affect dependency resolution in Spack](#). In *Free and Open source Software Developers’ European Meeting (FOSDEM’18)*, Brussels, Belgium, February 3 2018. LLNL-PRES-745770.

- [136] **Todd Gamblin**. Recent developments in Spack. In *Easybuild User Meeting*, Amsterdam, The Netherlands, January 30 2018.
- [137] David A. Beckingsale, Olga Pearce, and **Todd Gamblin**. Lightweight, Reusable Models for Dynamically Tuning Data-Dependent Code. In *Supercomputing 2016 (SC'16)*, Salt Lake City, UT, November 13-18 2016.
- [138] Olga Pearce, **Todd Gamblin**, Bronis de Supinski, Martin Schulz, and Nancy M. Amato. Decoupled load balancing. In *Symposium on Principles and Practices of Parallel Programming (PPoPP'15)*, San Francisco, CA, February 7-11 2015.
- [139] John Gyllenhaal, **Todd Gamblin**, Adam Bertsch, and Roy Musselman. **Enabling High Job Throughput for Uncertainty Quantification on BG/Q**. In *IBM HPC Systems Scientific Computing User Group (Scicomp'14)*, Chicago, IL, 2014.
- [140] Koji Inoue, Yuki Abe, Yuichi Inadomi, Kate Isaacs, **Todd Gamblin**, and Martin Schulz. NsimPower: Large Scale Interconnection Network Simulator for Power/Performance Analysis. In *JST/CREST International Symposium on Post Petascale System Software (ISP2S2)*, Kobe, Japan, December 2-4 2014.
- [141] 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 *ACM SIGPLAN Symposium on Principles and Practices of Parallel Programming (PPoPP'14)*, Orlando, FL, February 15-19 2014. LLNL-ABS-647655.
- [142] Anshu Arya, **Todd Gamblin**, Bronis R. de Supinski, and Laxmikant V. Kale. Evaluating topology mapping via graph partitioning. In *Supercomputing 2012 (SC'12)*, Salt Lake City, UT, November 10-16 2012. LLNL-POST-491398.
- [143] Katherine E. Isaacs, Aaditya G. Landge, **Todd Gamblin**, Peer-Timo Bremer, Valerio Pascucci, and Bernd Hamann. Exploring performance data with boxfish. In *Supercomputing 2012 (SC'12)*, Salt Lake City, UT, November 10-16 2012.
- [144] Vivek Kale, **Todd Gamblin**, Torsten Hoefler, Bronis R. de Supinski, and William Gropp. Slack-conscious lightweight loop scheduling for scaling past the noise amplification problem. In *Supercomputing 2012 (SC'12)*, Salt Lake City, UT, November 10-16 2012.
- [145] Brian T. N. Gunney, Abhinav Bhatele, and **Todd Gamblin**. Tree-based load balancer in the SAMRAI framework. In *SIAM Annual Meeting*, Minneapolis, MN, July 9-14 2012. LLNL-PRES-562671.
- [146] 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 14-17 2012.
- [147] 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 (PP12)*, Savannah, Georgia, USA, February 15-17 2012. LLNL-ABS-518731.
- [148] Olga Pearce, **Todd Gamblin**, Martin Schulz, Bronis R. de Supinski, and Nancy M. Amato. Load Balance: Correlating Application-Independent Measurements with Application-Semantic Computational Models. In *Supercomputing 2010 (SC'10)*, New Orleans, LA, November 13-19 2010. LLNL-POST-442731.
- [149] Allison H. Baker, **Todd Gamblin**, Martin Schulz, and Ulrike Yang. Algebraic multigrid for multicore architectures. In *SIAM Annual Meeting*, Pittsburgh, PA, July 12-16 2010. LLNL-PRES-435936.
- Best Poster** [150] **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, Robert J. Fowler, and Daniel A. Reed. Clustering performance data efficiently at massive scales. In *LLNL Computation Postdoctoral Poster Session*, Livermore, CA, June 29 2010. LLNL-POST-433541.
- [151] **Todd Gamblin**. Preparing for Exascale: System Software and Tools Research at LLNL. In *Global Scientific Information and Computing Seminar*, Tokyo Institute of Technology, Tokyo, Japan, June 11 2010. LLNL-PRES-434295.
- [152] **Todd Gamblin**. Load Balance Analysis with Libra. In *Proceedings of the Dagstuhl Seminar on Program Development for Extreme-Scale Computing*, Dagstuhl, Germany, May 3-7 2010. LLNL-PRES-431860.
- [153] **Todd Gamblin**. Scalable In-situ Analysis Techniques. In *Proceedings of the Dagstuhl Seminar on Program Development for Extreme-Scale Computing*, Dagstuhl, Germany, May 3-7 2010. LLNL-POST-433541.
- [154] **Todd Gamblin**. Infrastructure challenges for scalable data analysis algorithms. In *ParaDyn/Condor Week 2010*, Madison, WI, April 12-16 2010. University of Wisconsin. LLNL-PRES-408863.

- [155] **Todd Gamblin**, Bronis R. de Supinski, and Martin Schulz. Performance modeling and analysis at LLNL. In *Proceedings of Joint Operations Weapons Operations Group 34 (JOWOG34)*, Sandia National Laboratories, Albuquerque, NM, July 20-24 2009. LLNL-PRES-463247.
- [156] **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, Robert J. Fowler, and Daniel A. Reed. System-wide performance equivalence class detection using clustering. In *Supercomputing 2008 (SC'08)*, Austin, TX, November 15-21 2008. LLNL-POST-408815.
- [157] **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, Robert J. Fowler, and Daniel A. Reed. Scalable Load Balance Measurement Using Wavelets. In *The Salishan Conference on High-Speed Computing*, Gleneden Beach, OR, April 21-24 2008. LLNL-POST-463395.
- [158] **Todd Gamblin**. Scalable load-balance measurement and analysis. In *SIAM Conference on Parallel Processing for Scientific Computing (PP08)*, Atlanta, GA, March 12-14 2008.
- [159] **Todd Gamblin**, Prasun Ratn, Bronis R. de Supinski, Martin Schulz, Frank Mueller, Robert J. Fowler, and Daniel A. Reed. An open framework for scalable, reconfigurable performance analysis. In *Supercomputing 2007 (SC'07)*, Reno, NV, November 10-16 2007. UCRL-POST-236200.
- Best Poster** [160] **Todd Gamblin**, Robert J. Fowler, and Daniel A. Reed. Methods for automatic behavioral stratification in scientific codes. In *Proceedings of the 7th LACSI Symposium*, Santa Fe, NM, October 17-19 2006.
- [161] **Todd Gamblin** and Daniel A. Reed. Adaptive performance monitoring and profiling in large-scale systems. In *Proceedings of the 6th LACSI Symposium*, Santa Fe, NM, October 11-13 2005.