

Todd Gamblin

Computer Scientist, Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

P.O. Box 808, L-557
Livermore, CA, 94551
☎ 919.360.8283
✉ tgamblin@llnl.gov
<http://tgamblin.github.io>

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

- 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.
- 2014 Best Paper, International Symposium on Cluster, Cloud and Grid Computing (CCGrid).
- 2013 Best Paper, International Conference on Supercomputing (ICS).
- 2013 Invited to LLNL Computation Leadership Program.
- 2012, 2013 Two LLNL awards for building collaborative development tools for all supercomputer users.
- 2011 LLNL award for leadership on the PAVE project.
- 2010 LLNL achievement award for the Scalable Application Preparation 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.

Research Experience

- 2010-present **Lawrence Livermore National Laboratory**, *Computer Scientist*.
 - Research and develop tools to measure & analyze performance data from parallel scientific simulations.
 - Lead: deployment of the Sonar cluster: data analytics for the Livermore Computing Center.
 - Lead: Spack package manager project: combinatorial builds of over 1,300 HPC packages (spack.io).
 - Lead: HPC Developer Ecosystem (HPCDE) DevOps/collaboration team within Livermore Computing.
- 2009-2010 **Lawrence Livermore National Laboratory**, *Postdoctoral Scholar*.
Developed $O(\log(P))$ parallel clustering algorithm for on-line performance analysis.
- 2008-2009 **Lawrence Livermore National Laboratory**, *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.

2001-2002 **Williams College**, *Undergraduate Research Assistant*.
Investigated cache-conscious memory performance optimizations.

Professional Experience

2002-2003 **ValueCommerce, Co., Ltd.**, *Software Developer*, Tokyo, Japan.
Developed control panel for web hosting application; conducted performance analysis for ad servers.

Summer 2001 **EMC Corporation**, *Summer Intern*, Research Triangle Park, NC.

Summer 2000 Developed distributed performance monitoring tools and web-based user interface for NAS servers.

Research and Other Funding

2017-present **PI**, *DevRAMP: Reproducibility, Analysis, Monitoring, and Performance for Developers*, programmatic, ≈ 3 FTE/yr.
Department of Energy (DOE) Advanced Simulation and Computing (ASC) ATDM Program.

2017-2018 **PI**, *LLNL Early and Mid-Career Recognition Award*, 20% FTE/yr.
Recognized by LLNL and funded to work on expanding the Spack community.

2016-2017 **PI**, *HPC Developer Ecosystem*, programmatic, ≈ 2 FTE/yr.
Department of Energy (DOE) Advanced Simulation and Computing (ASC) ATDM Program.

2014-2019 **PI**, *Statistical Methods for Exascale Performance Modeling*, \$2,500,000.
U.S. Department of Energy (DOE) Early Career Research Program.

2014-2017 **PI**, *Sonar: An Analytics Cluster for Livermore Computing Performance Data*, \$100,000/yr.
LLNL Computation TechBase Project (3 years running)

2010-present **PI**, *Performance Analysis and Visualization at Exascale (PAVE)*, \$500,000/year.
Department of Energy (DOE) Advanced Simulation and Computing (ASC) Program.

2013-2015 **Co-PI**, *Performance Insight for Programs and Exascale Runtimes (PIPER)*, \$4,500,000.
PI Martin Schulz, LLNL. DOE Advanced Scientific Computing Research (ASCR) X-Stack 2 Program.

2012-2014 **Co-PI**, *Performance Visualization at ExaScale (PeVES)*, \$250,000.
PI Timo Bremer, LLNL. University of California Laboratory Fees Research Program.

2012-2014 **Co-PI**, *Scalable Topology-Aware Task Embedding (STATE)*, \$945,000.
PI Abhinav Bhatele, LLNL. DOE LDRD Exploratory Research.

2012 **PI**, *Data Analytics for HPC*, \$100,000.

Time Allocation Grants

2013-2015 **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. 60M hours on ANL Mira; 50M hours on ORNL Titan; 10M hours on ANL Intrepid.

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

2017 **IEEE CLUSTER**, *General Chair*.
Location to be announced at IEEE CLUSTER 2016

2014-2016 **JOWOG 34: Applied Computer Science**, *Co-Chair, Founder*.
Scalable algorithms & simulation techniques for LLNL, LANL, Sandia (US) and AWE (UK). 100+ attendees.

2015-2016 **International Workshop on HPC User Support Tools (HUST)**, *Co-Chair*.
Held at SC15-16 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.
- 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

- 2017 **ACM/IEEE Supercomputing Conference (SC)**, *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 (SC)**, *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)**, *Organizing Committee*.

Grant Reviewing

- 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 C, C++, Fortran (linking), Python, numpy, Java, Ruby, Perl, git, svn, CVS, autotools, CMake.
- 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.
- Data Analysis Sampling, sequential and parallel clustering algorithms.
- 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

see github.com/LLNL to download.

- spack** Package manager for HPC software, used widely by developers and HPC centers.
- 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

- 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

- July 16, 2017 **Request for Commits Podcast**,
Open Source and Supercomputers (Spack), Hosts: Nadia Eghbal and Mikeal Rogers,
<https://changelog.com/rfc/13>.
- June 16, 2017 **HPC Knowledge Meeting 2017**,
Managing HPC Software Complexity with Spack, San Sebastián, Spain.
- June 2, 2017 **4th International Workshop on Reproducibility in Parallel Computing (REPPAR'17)**,
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 30, 2015 **Research Computing & Engineering (RCE) Podcast**,
Spack, Hosts: Brock Palen and Jeff Squyres,
<http://www.rce-cast.com/Podcast/rce-102-spack.html>.
- March 12, 2015 **DOE Exascale Computing Initiative Runtime Systems Workshop**,
ASC's Tool Needs, Rockville, MD.
- 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] 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.
- [2] 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.
- [3] 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.
- [4] 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.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] Olga Pearce, **Todd Gamblin**, Bronis R. de Supinski, Martin Schulz, and Nancy M. Amato. MPMDFramework 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.
- [12] 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.
- [13] **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).
- [14] 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).
- [15] 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.
- [16] 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).
- [17] 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.
- [18] 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.
- [19] 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).
- [20] 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).
- [21] 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).
- [22] 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).
- [23] 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.
- [24] 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** [25] 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 2014. LLNL-CONF-645876, (19% accepted).
- [26] 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.
- [27] 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).
- [28] 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).
- [29] 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** [30] 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).
- [31] 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).
- [32] 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.

- [33] 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).
- [34] 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).
- [35] 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).
- [36] 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).
- [37] 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).
- [38] 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).
- [39] 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.
- [40] 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.
- [41] 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.
- [42] 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).
- [43] 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).
- [44] 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).
- [45] 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).
- [46] **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).
- [47] 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 Snively, 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.

- [48] 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.
- [49] **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).
- [50] 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.
- [51] **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

- [52] 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.
- [53] 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.
- [54] 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.
- [55] 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.
- [56] 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.
- [57] 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 (IWPDP'16)*, Ottawa, Canada, October 23-27 2016.
- [58] 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.
- [59] **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.
- [60] 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.
- [61] 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.
- [62] 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.
- [63] 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.
- [64] 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.

- [65] 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.
- [66] **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.
- [67] 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.
- [68] 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.
- [69] 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.
- [70] 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.
- [71] 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.
- [72] 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

- [73] **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.

Special Issues, Book Chapters, Articles

- [74] 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.
- [75] 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.
- [76] 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.
- [77] **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.
- [78] 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

- [79] **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

- [80] 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.
- [81] 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.
- [82] 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.
- [83] 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.
- [84] 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.
- [85] 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.
- [86] 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.
- [87] 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.
- [88] 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.
- [89] 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.
- [90] 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 (PPI2)*, Savannah, Georgia, USA, February 15-17 2012. LLNL-ABS-518731.
- [91] 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.
- [92] 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** [93] **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.
- [94] **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.
- [95] **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.
- [96] **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.
- [97] **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.
- [98] **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.

- [99] **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.
- [100] **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.
- [101] **Todd Gamblin**. Scalable load-balance measurement and analysis. In *SIAM Conference on Parallel Processing for Scientific Computing (PP08)*, Atlanta, GA, March 12-14 2008.
- [102] **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** [103] **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.
- [104] **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.