MARY WOLCOTT HALL

School of Computing ⋄ 50 S. Central Campus Drive ⋄ Salt Lake City, UT 84112
(801) 585-1039 ⋄ mhall@cs.utah.edu

EDUCATION

Rice University May 1991
Doctor of Philosophy, Computer Science
Rice University May 1989
Master of Science, Computer Science
Rice University May 1985
Bachelor of Arts, Computer Science and Mathematical Sciences Magna Cum Laude

PROFESSIONAL EXPERIENCE

2020-	Director	School of Computing, University of Utah
2012-	Professor	School of Computing, University of Utah
2008-2012	Associate Professor	School of Computing, University of Utah
2003-2008	Research Associate Professor	Computer Science, University of Southern California
1997-2003	Research Assistant Professor	Computer Science, University of Southern California
1996-2008	Project Leader	USC Information Sciences Institute
1995-1996	Visiting Assistant Professor and Senior Research Fellow	Computer Science, California Institute of Technology
1992-1994	Research Scientist	Stanford University
1990-1992	Research Scientist	Rice University

AWARDS

2020	IEEE Fellow
2010	ACM Distinguished Scientist
2015	CRA Leadership in Science Policy Institute
2010-2013	Nvidia Professor Partnership Fellowship
2003-2007	Intel Faculty Fellow
1993-1994	National Science Foundation CISE Postdoctoral Fellowship

NATIONAL LEADERSHIP

2020-	Member, Computing Research Association Board of Directors
2019	Computing and Communications Foundations Subcommittee Chair, NSF CISE Committee of Visitors
2019-	Member, DOE Computational Science Graduate Fellowships Steering Committee
2019-	Chair, Institute of Advanced Computational Science External Advisory Board, Stonybrook University; Member, 2015-2019
2015-2020	Member and ACM Representative, Computing Research Association Board of Directors
2009-2014	Chair, ACM History Committee, 2009-2013, Co-chair in 2014
2010, 2014	Chair, ACM and IEEE Computer Society Ken Kennedy Award Committee
2011, 2014	Member, IEEE Computer Society Awards Committee

PUBLICATIONS

Advisory Reports

NSF19 *NSF CISE Committee of Visitors Report For the Divisions of Computing and Communication Foundations, Computer and Network Systems, and Information and Intelligent Systems*, Azer Bestavros, Mary Hall, Sonia Fahmy and Julia Hirschberg, Nov. 2019.

CRA17 *Generation CS: Computer Science Undergraduate Enrollments Surge Since 2006* Tracy Camp (Chair), W. Richards Adrion, Betsy Bizot, Susan Davidson, Mary Hall, Susanne Hambrusch, Ellen Walker, and Stuart Zweben, Computing Research Association, Feb. 2017.

DARPA09 *ExaScale Software Study: Software Challenges in Extreme Scale Systems*, Vivek Sarkar (Editor and Lead), Saman Amarasinghe, Dan Campbell, William Carlson, Andrew Chien, William Dally, Elmootazbellah Elnohazy, Mary Hall, Robert Harrison, William Harrod, Kerry Hill, Jon Hiller, Sherman Karp, Charles Koelbel, David Koester, Peter Kogge, John Levesque, Daniel Reed, Robert Schreiber, Mark Richards, Al Scarpelli, John Shalf, Allan Snavely, Thomas Sterling, DARPA IPTO, Sept. 2009.

NSF09 *Compiler Research: The Next Fifty Years* Mary Hall, David Padua and Keshav Pingali, Communications of the ACM, Feb. 2009, Report of the NSF Future of Compiler Research and Education Workshop (2007).

Magazine and Blog Articles

CRN19 *CRA Launches BPCnet.org: A Resource Portal for Broadening Participation in Computing Efforts* Mary Hall, Shar Steed, Computing Research News 31(3), March 2019.

CACMBl18 *Broadening Participation in Computing is Easier Than You Think* Mary Hall, Richard Ladner, Diane Levitt, Manuel Prez-Quiones, Blog@CACM, Dec. 2018.

CRA17 *New Approaches to Producing High-Performance Code, Thanks to Compiler Technology* Mary Hall, CRA Board Member Highlight, Nov. 2017.

Inroads17 *Generation CS: the challenges of and responses to the enrollment surge* Tracy Camp, W. Richards Adrion, Betsy Bizot, Susan Davidson, Mary Hall, Susanne Hambrusch, Ellen Walker, and Stuart Zweben. ACM Inroads, 8(4) 59–65, October 2017.

Inroads17 *Generation CS: the mixed news on diversity and the enrollment surge* Tracy Camp, W. Richards Adrion, Betsy Bizot, Susan Davidson, Mary Hall, Susanne Hambrusch, Ellen Walker, and Stuart Zweben, ACM Inroads, 8(4) 36–42, July 2017.

Inroads17 *Generation CS: the growth of computer science* Tracy Camp, W. Richards Adrion, Betsy Bizot, Susan Davidson, Mary Hall, Susanne Hambrusch, Ellen Walker, and Stuart Zweben, *ACM Inroads*, 8(2) 44–50, May 2017.

CACM12 *Understanding ACM's Past*, Mary Hall, Communications of the ACM, 55(12) 5–5, Dec. 2012.

CACM06 *A Wiki for discussing and promoting best practices in research*, Mark D. Hill, Jean-Luc Gaudiot, Mary Hall, Joe Marks, Paolo Prinetto, Donna Baglio, Communications of the ACM, 49(9):63–64, Sept. 2006.

Journal Articles

CCPE21 *Autotuning PolyBench Benchmarks with LLVM Clang/Polly Loop Optimization Pragmas Using Bayesian Optimization* X. Wu, M. Kruse, P. Balaprakash, H. Finkel, V. Taylor, P. Hovland, M. Hall, Concurrency: Practice and Experience, Nov. 2021.

TACO19 *Data-Driven Mixed Precision Sparse Matrix Vector Multiplication for GPUs* Khalid Ahmad, Hari Sundar, Mary Hall, ACM Transactions on Architectures and Code Optimization 16(5), Dec. 2019.

IJHPCA19 *SWIRL: High-Performance Many-Core CPU Code Generation for Deep Neural Networks* Anand Venkat, Tharindu Rusira, Raj Barik, Mary Hall, Leonard Truong, International Journal of High-Performance Computing Applications, 33(6), 2019.

PIEEE18 *The Sparse Polyhedral Framework: Composing Compiler-Generated Inspector-Executor Code* M. M. Strout, M. Hall and C. Olschanowsky, Proceedings of the IEEE 106(11):1921–1934, Nov. 2018.

PIEEE18 *Autotuning in High-Performance Computing Applications* Prasanna Balaprakash, Jack Dongarra, Todd Gamblin, Mary Hall, Jeffrey K. Hollingsworth, Boyana Norris, and Richard Vuduc, Proceedings of the IEEE 106(11):2068–2083, Nov. 2018.

PARCO18 *Student Cluster Competition 2017, Team University of Utah: Reproducing Vectorization of the Tersoff Multi-Body Potential on the Intel Broadwell and Intel Skylake Platforms* J. Lake, Q. Chao, H. Eyre, E. Ford, K. Parker, K. Savoie, H. Sundar, M. Hall, Parallel Computing 79, Jul. 2018.

PARCO17 *Reproducing ParConnect for SC16* Marek Baranowski, Braden Caywood, Hannah Eyre, Janaan Lake, Kevin Parker, Kincaid Savoie, Hari Sundar, Mary Hall, Parallel Computing 70:18–21, Dec. 2017.

PARCO17 *Compiler-based code generation and autotuning for geometric multigrid on GPU-accelerated supercomputers* P. Basu, S. Williams, B. Van Straalen, L. Oliker, P. Colella, and M. Hall, Parallel Computing 64(C):50–64, May 2017.

TACO16 *Designing a Tunable Nested Data-Parallel Programming System* S. Muralidharan, M. Garland, A. Sidelnik, M. Hall, *ACM Transactions on Architecture and Code Optimization*, 13(4), December 2016.

IJHPCA13 *Towards Making Autotuning Mainstream* P. Basu, M. Hall, M. Khan, S. Maindola, S. Muralidharan, S. Ramalingam, A. Rivera, M. Shantharam, A. Venkat, International Journal of High Performance Computing Applications, 27(4), Nov. 2013.

TACO13 *A script-based autotuning compiler system to generate high-performance CUDA code* M. Khan, P. Basu, G. Rudy, M. Hall, C. Chen, and J. Chame. ACM Transactions on Architecture and Code Optimization 9(4), January 2013.

JS12 *Hierarchical parallelization and optimization of high-order stencil computations on multicore clusters* H. Dursun, M. Kunaseth, K. Nomura, J. Chame, R.F. Lucas, C. Chen, M. Hall, R.K. Kalia, A. Nakano, P. Vashishta, The Journal of Supercomputing, 62(2):946–966, Dec. 2012.

IJHPCA11 *Auto-tuning Full Applications: A Case Study* A. Tiwari, C. Chen, C. Liao, J. Chame, J. Hollingsworth, M. Hall and D. Quinlan, International Journal of High Performance Computing Applications 25(3):286-294, Aug. 2011.

TRETS11 *Domain-Specific Optimization of Signal Recognition Targeting FPGAs* M. Demertzi, P.C. Diniz, M.W. Hall, A.C. Gilbert and Y. Wang, ACM Transactions on Reconfigurable Technology and Systems 4(2), May, 2011.

CC10 *Parameterized specification, configuration and execution of data-intensive scientific workflows* V.S. Kumar, T. Kurc, V. Ratnakar, J. Kim, G. Mehta, K. Vahi, Y.L. Nelson, P. Sadayappan, E. Deelman, Y. Gil, M. Hall and J. Saltz, Cluster Computing, April 2010.

Journal Articles, cont.

IJHPCA09 *HPC and Grid Computing for Integrative Biomedical Research* T. Kurc, S. Hastings, V. Kumar, S. Langella, A. Sharma, T. Pan, S. Oster, D. Ervin, J. Permar, S. Narayanan, Y. Gil, E. Deelman, M. Hall, J. Saltz, International Journal of High Performance Computing Applications, 2009.

IJES009 *Evaluating Compiler Technology for Control-Flow Optimizations for Multimedia Extension Architectures* J. Shin, M. Hall and J. Chame. Award paper invited from MSP 7, International Journal of Embedded Systems, 2009.

JP08 *PERI Auto-Tuning* David H. Bailey, Jacqueline Chame, Chun Chen, Jack Dongarra, Mary Hall, Jeffrey K. Hollingsworth, Paul Hovland, Shirley Moore, Keith Seymour, Jaewook Shin, Ananta Tiwari, Sam Williams, Haihang You, Journal of Physics: Conference Series, Vol. 125, 2008.

PIEEE08 *Self-Configuring Applications for Heterogeneous Systems: Program Composition and Optimization Using Cognitive Techniques* M. Hall, Y. Gil and R. Lucas. Proceedings of the IEEE, Special Issue on Cutting-Edge Computing 96(5), May 2008.

JP07 *Compiler-assisted performance tuning* Chun Chen, Jacqueline Chame, Yoonju Lee Nelson, Pedro Diniz, Mary Hall and Robert Lucas, Journal of Physics: Conference Series, Vol. 78, 2007.

IJPP05 *Empirical Optimization for a Sparse Linear Solver: A Case Study* Y. Lee, P. Diniz, M. Hall and R. Lucas. International Journal of Parallel Programming, vol. 33, 2005.

TOPLAS05 *Interprocedural Parallelization Analysis in SUIF* M.W. Hall, S. Amarasinghe, B. Murphy, S. Liao and M. Lam. ACM Transactions on Programming Languages and Systems 27(4):662–731, July, 2005.

MICPRO05 *Automatic Mapping of C to FPGAs with the DEFACTO Compilation and Synthesis System* P. Diniz, M. Hall, J. Park, B. So and H. Ziegler, Elsevier MICPRO Special Issue on Field-Programmable-Gate-Arrays (FPGAs) vol. 29(2-3):51–62, April, 2005.

JILP03 *Exploiting Superword-Level Locality in Superword Register Files for Multimedia Extension Architectures* Award paper from PACT'02, Journal of Instruction-Level Parallelism 5(2003):1–28.

TPDS00 *Evaluating Automatic Parallelization in SUIF* S. Moon, B. So, and M.W. Hall. IEEE Transactions on Parallel Distributed Systems, 11(1), 36–49, Jan. 2000.

SP99 *Combining Compile-Time and Run-Time Parallelization* S. Moon, B. So and M.W. Hall. Award paper from LCR '98, Scientific Programming, Vol 7(3-4), 274–260, Aug. 1999.

CP&E98 *Adaptive Parallelism in Compiler-Parallelized Code* M.W. Hall and M. Martonosi. Concurrency: Practice and Experience, 10(14):1235–1250, Dec. 1998.

Computer96 *Maximizing Multiprocessor Performance with the SUIF Compiler* M.W. Hall, J. Anderson, S. Amarasinghe, B. Murphy, E. Bugnion and M. Lam. IEEE Computer, 29(12) Dec. 1996.

TPDS96 *Characterizing the Memory Behavior of Compiler-Parallelized Applications* E. Torrie, M. Martonosi, C. Tseng and M.W. Hall. IEEE Transactions on Parallel and Distributed Systems, 7(12):1225–1238, Dec. 1996.

JPDC96 *Interprocedural Compilation of Fortran D* M.W. Hall, S. Hiranandani, K. Kennedy and C. Tseng. Journal of Parallel and Distributed Computing, vol. 38(2):114–129, Aug. 1996.

Micro96 *Multiprocessors from a Software Perspective* S. Amarasinghe, J. Anderson, R. French, S. Liao, C. Wilson, B. Murphy, M. Lam and M.W. Hall. Invited paper selected as a Hot Chips '95 award paper, IEEE Micro Jun. 1996.

Journal Publications, cont.

IJPP96 *Memory Referencing Behavior in Compiler-Parallelized Applications* E. Torrie, M. Martonosi, C. Tseng and M.W. Hall. Invited paper selected as a PACT '95 award paper, International Journal of Parallel Programming vol. 24(4), (1996).

CPAM95 *Interprocedural Analysis and Optimization* K.D. Cooper, M.W. Hall, K. Kennedy and L. Torczon. Communications on Pure and Applied Mathematics (1995).

CL93 *A Methodology for Procedure Cloning* K.D. Cooper, M.W. Hall and K. Kennedy ICCL '92 award paper, Computer Languages, 19(2) Apr. 1993.

PIEEE93 *The ParaScope Parallel Programming Environment* K. Cooper, M.W. Hall, R. Hood, K. Kennedy, K. McKinley, J. Mellor-Crummey, L. Torczon and S. Warren. Proceedings of the IEEE (Feb. 1993).

LOPLAS92 *Efficient Call Graph Analysis* M.W. Hall and K. Kennedy. *ACM Letters on Programming Languages and Systems*, 1(3) (Sep. 1992).

LOPLAS92 *Unexpected Side Effects of Inline Substitution: A Case Study* K.D. Cooper, M.W. Hall and L. Torczon. ACM Letters on Programming Languages and Systems, 1(1) Mar. 1992.

SP&E91 *An Experiment with Inline Substitution* K. D. Cooper, M.W. Hall and L. Torczon. Software — Practice and Experience, 21(6) Jun. 1991.

TSE90 *Constructing the Call Multigraph* D. Callahan, A. Carle, M.W. Hall and K. Kennedy. IEEE Transactions on Software Engineering, 16(4) Apr. 1990.

Conference Publications

PPoPP21 *Improving communication by optimizing on-node data movement with data layout* T. Zhao, M. Hall, H. Johansen, and S. Williams. Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP '21), Mar. 21.

SC19 *Exploiting Reuse and Vectorization in Blocked Stencil Computations on CPUs and GPUs* T. Zhao, S. Williams, M. Hall, H. Johansen International Conference on Supercomputing, Networking, Storage and Analysis (SC), Nov. 2019.

PLDI19 *Sparse Computation Data Dependence Simplification for Efficient Compiler-Generated Inspectors* M. Mohammadi, K. Cheshmi, E. Davis, M. Hall, M. Dehnavi, P. Nandy, C. Olschanowsky, A. Venkat. T. Yuki, M. Strout Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), June 2019.

SC16 *Automating Wavefront Parallelization for Sparse Matrix Codes* A. Venkat, M. Mohamadi, J. Park, R. Barik, H. Rong, M. Strout, M. Hall, International Conference on Supercomputing, Networking, Storage and Analysis (SC), Nov. 2016, **Best Paper Finalist**.

IPDPS16 *Synchronization Tradeoffs in GPU Implementations of Graph Algorithms* R. Kaleem, A. Venkat, S. Pai, M. Hall, K. Pingali, Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2016.

ASPLOS16 *Architecture-Adaptive Code Variant Tuning* S. Muralidharan, A. Roy, M. Hall, M. Garland, and P. Rai, Proceedings of the ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), April 2016.

ICPP15 *Generating Efficient Tensor Contractions for GPUs* T. Nelson, A. Rivera, P. Balaprakash, M. Hall, P.D. Hovland, E. Jessup, B. Norris, Proceedings of the IEEE International Conference on Parallel Processing (ICPP), Sept. 2015.

PLDI15 *Loop and Data Transformations for Sparse Matrix Code* A. Venkat, M. Hall, M. Strout, Proceedings of the ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), June 2015.

IPDPS15 *Compiler-Directed Transformations for Higher-Order Stencils* P. Basu, S. Williams, B. van Straalen, M. Hall, L. Oliker, P. Collela, Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS), May 2015.

IPDPS14 *Nitro: A Framework for Adaptive Code Variant Tuning* S. Muralidharan , M. Shantharam, M. Hall, M. Garland, B. Catanzaro, *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS)*, May 2014.

CGO14 *Non-affine Extensions to Polyhedral Code Generation* A. Venkat, M. Shantharam, M. Hall, M. M. Strout, *Proceedings of the International Conference on Code Generation and Optimization (CGO)*, Feb. 2014.

HiPC13 *Compiler Generation and Autotuning of Communication-Avoiding Operators for Geometric Multigrid* P. Basu, S. Williams, B. Van Straalen, A. Venkat, L. Oliker, M. Hall, IEEE International Conference on High Performance Computing (HiPC), Dec. 2013.

ISWC11 *Analyzing the effect of compiler optimizations on application reliability* M. Demertzi, M. Annavaram and M. Hall, *Proceedings of the IEEE International Symposium on Workload Characterization (ISWC)*, Nov. 2011.

Conference Publications, cont.

POPL11	*EigenCFA: Accelerating Flow Analysis with GPUs* T. Prabhu, S. Ramalingam , M. Might, M. Hall, ACM SIGPLAN Principles of Programming Languages (POPL), Jan. 2011.
ICS10	*Autotuning and Specialization: Speeding up Nek5000 with Compiler Technology* J. Shin, M. Hall, J. Chame, C. Chen, P. Fischer, P.D. Hovland, International Conference on Supercomputing (ICS), June 2010.
SAAHPC09	*GPU Acceleration of the Generalized Interpolation Material Point Method* W. Chiang, M. DeLisi, T. Hummel, T. Prete, K. Tew, M. Hall, P. Wallstedt, and J. Guilkey, Symposium on Application Accelerators for High Performance Computing, July, 2009.
HPDC09	*An Integrated Framework for Parameter-based Optimization of Scientic Workflows* V. S. Kumar, P. Sadayappan, G. Mehta, K. Vahi, E. Deelman, V. Ratnakar, J. Kim, Y. Gil, M. Hall, T. Kurc, J. Saltz, Proceedings of the International Symposium on High Performance Distributed Computing, June, 2009.
HPDC09	*Model-Guided Autotuning of High-Productivity Languages for Petascale Computing* H. Zima M. Hall, C. Chen, J. Chame, Proceedings of the International Symposium on High Performance Distributed Computing, June, 2009.
IPDPS09	*A Scalable Autotuning Framework for Compiler Optimization* A. Tiwari, C. Chen, J. Chame, M. Hall and J. K. Hollingsworth, Proceedings of the International Parallel and Distributed Processing Symposium, May, 2009.
CGO05	*Combining Models and Guided Empirical Search to Optimize for Multiple Levels of the Memory Hierarchy* C. Chen, J. Chame and M. Hall. Proceedings of the Conference on Code Generation and Optimization, March, 2005.
CGO05	*Superword-Level Parallelism in the Presence of Control Flow* J. Shin, M. Hall and J. Chame. Proceedings of the Conference on Code Generation and Optimization, March, 2005.
FPGA05	*Evaluating Heuristics in Automatically Mapping Multi-Loop Applications to FPGAs* H. Ziegler and M. Hall. Proceedings of the ACM Conference on Field-Programmable Gate Arrays, February, 2005.
CC04	*Increasing the Applicability of Scalar Replacement* B. So and M. Hall, Proceedings of the Compiler Construction Conference, March, 2004.
CGO04	*Custom Data Layout for Memory Parallelism* B. So, M. Hall and H. Ziegler, Proceedings of the Code Generation and Optimization Conference, March, 2004.
DAC03	*Compiler-generated communication for pipelined FPGA applications* H. Ziegler, M. Hall and P. Diniz, In Proceedings of the ACM/IEEE Design Automation Conference, June, 2003.
DAC03	*Using estimates from behavioral synthesis tools in compiler-directed design space exploration* B. So, P. Diniz, and M. Hall, Proceedings of the ACM/IEEE Design Automation Conference, June, 2003.
PACT02	*Compiler-Controlled Caching in Superword Register Files for Multimedia Extension Architectures* J. Shin, J. Chame, and M. Hall. Proceedings of the Parallel Architectures and Compilation Techniques Conference, Sept. 2002.
ICS02	*The Architecture of the DIVA Processing-In-Memory Chip* J. Draper, J. Chame, M. Hall, C. Steele, T. Barrett, J. LaCoss, J. Granacki, J. Shin, C. Chen, C. W. Kang, I. Kim, G. Daglikoca, Proceedings of the International Conference on Supercomputing, June, 2002.

Conference Publications, cont.

FCCM02 *Coarse-Grain Pipelining for Multiple FPGA Architectures* H. Ziegler, B. So, M. Hall and P. Diniz, Proceedings of the IEEE Symposium on Field-Programmable Custom Computing Machines, April, 2002.

FCCM02 *Coarse-Grain Pipelining for Multiple FPGA Architectures* H. Ziegler, B. So, M. Hall and P. Diniz, Proceedings of the IEEE Symposium on Field-Programmable Custom Computing Machines, April, 2002.

SC99 *Mapping Irregular Computations to DIVA, a Data-Intensive Architecture* M. Hall, P. Kogge, J. Koller, P. Diniz, J. Chame, J. Draper, J. LaCoss, J. Granacki, J. Brockman, W. Athas, A. Srivastava, J. Shin, J. Park. Proceedings of SC99, Nov. 1999.

PPoPP99 *Evaluation of Predicated Array Data-Flow Analysis for Automatic Parallelization* S. Moon and M.W. Hall. Proceedings of the ACM Symposium on Principles and Practice of Parallel Programming, May, 1999.

ICS98 *Predicated Array Data-Flow Analysis for Run-Time Parallelization* S. Moon, M.W. Hall and B. Murphy. Proceedings of the ACM International Conference on Supercomputing Jul. 1998.

ICS98 *Measuring the Effectiveness of Automatic Parallelization in SUIF* B. So, S. Moon and M.W. Hall. Proceedings of the ACM International Conference on Supercomputing Jul. 1998).

SC95 *Detecting Coarse-Grain Parallelism Using an Interprocedural Parallelizing Compiler* M.W. Hall, S. Amarasinghe, B. Murphy, S. Liao and M. Lam. Proceedings of Supercomputing '95 Dec. 1995.

PACT95 *Evaluating the impact of advanced memory systems on compiler-parallelized codes* E. Torrie, M. Martonosi, C. Tseng, and M.W. Hall. Proceedings of the International Conference on Parallel Architectures and Compilation Techniques June 1995.

SIAM95 *Interprocedural Parallelization Analysis: A Case Study* M.W. Hall, B. Murphy and S. Amarasinghe. Seventh SIAM Conference on Parallel Processing Feb. 1995.

PPoPP93 *Experiences Using the ParaScope Editor: an Interactive Parallel Programming Tool* M.W. Hall, T. Harvey, K. Kennedy, N. McIntosh, K. McKinley, J. Oldham, M. Paleczny, G. Roth. Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming Apr. 1993.

SC92 *Interprocedural Compilation of Fortran D for MIMD Distributed-Memory Machines* M.W. Hall, S. Hiranandani, K. Kennedy and C. Tseng. Proceedings of Supercomputing '92, Nov. 1992.

ICCL92 *Procedure Cloning* K. D. Cooper, M.W. Hall and K. Kennedy. Proceedings of the IEEE International Conference on Computer Languages (Apr. 1992).

SC91 *Interprocedural Transformations for Parallel Code Generation* K. Kennedy and K. S. McKinley. Proceedings of Supercomputing '91, Nov. 1991.

PSDE86 *Editing and Compiling Whole Programs* K. D. Cooper, K. Kennedy, L. Torczon and A. Weingarten and M. Wolcott. Proceedings of the ACM SIGSOFT/SIGPLAN Symposium on Practical Software Development Environments Dec. 1986.

Book Chapters

LCPC20 *Optimized Code Generation for Deep Neural Networks* J. Lake, T. Pattabandi, M. Hall, Lecture Notes in Computer Science, Languages and Compilers for Parallel Computing 2020, Springer Verlag, 2022.

LCPC19 *SWIRL++: Evaluating Performance Models to Guide Code Transformation in Convolutional Neural Networks* T. Rusira Pattabandi, A. Venkat, R. Barik, and M. Hall, Lecture Notes in Computer Science, Languages and Compilers for Parallel Computing 2019, Springer Verlag, 2021.

LCPC17 *Polyhedral Compilation Support for C++ Features: A Case Study with CPPTRAJ* A. Roy, D. Roe, M. Hall, T. Cheatham, Lecture Notes in Computer Science, 2019, Volume 11403, Languages and Compilers for Parallel Computing 2017, Springer Verlag, Pages 26-35.

LCPC16 *Optimizing LOBPCG: Sparse Matrix Loop and Data Transformations in Action* K. Ahmad, A. Venkat and M. Hall, Lecture Notes in Computer Science, 2017, Volume 10136, Languages and Compilers for Parallel Computing 2016, Springer Verlag, Pages 221-231.

LCPC16 *Polyhedral Compiler Technology in Collaboration with Autotuning Important to Domain-Specific Frameworks for HPC* M. Hall and P. Basu, Lecture Notes in Computer Science, 2017, Volume 10136, Languages and Compilers for Parallel Computing, Springer Verlag.

LCPC10 *A Programming Language Interface to Describe Transformations and Code Generation* G. Rudy, M. Khan, M. Hall, C. Chen and J. Chame, Lecture Notes in Computer Science, 2011, Volume 6548, Languages and Compilers for Parallel Computing, Springer Verlag, 136–150.

PERF10 *Languages and Compilers for Autotuning* M.W. Hall and J. Chame, In Performance Tuning of Scientific Applications, edited by David Bailey, Robert F. Lucas and Sam Williams. Taylor and Francis publishers, Nov. 2010.

IWAPT10 *Autotuning and Specialization: Speeding up Matrix Multiply for Small Matrices with Compiler Technology* Jaewook Shin, Mary W. Hall, Jacqueline Chame, Chun Chen, Paul D. Hovland, Software Automatic Tuning: from concepts to state-of-the-art results, edited by Keita Teranishi, John Cavazos, Ken Naono and Reiji Suda, Springer-Verlag Publishers, 2010, 353–370.

LCPC09 *Loop Transformation Recipes for Code Generation and Auto-Tuning* Mary Hall, Jacqueline Chame, Chun Chen, Jaewook Shin and Gabe Rudy, Lecture Notes in Computer Science, 2010, Volume 5898, Languages and Compilers for Parallel Computing, Springer-Verlag, 50–64.

ICCS07 *A Combined Hardware/Software Optimization Framework for Signal Representation and Recognition* M. Demertzi, P. Diniz, M. W. Hall, A. C. Gilbert and Y. Wang, Lecture Notes in Computer Science, 2007, Volume 4487, Computational Science, ICCS 2007, 1230–1237.

LCPC05 *A Systematic Approach to Model-Guided Empirical Search for Memory Hierarchy Optimization* C. Chen, J. Chame, M. Hall and K. Lerman, Lecture Notes in Computer Science, 2006, Volume 4339, Languages and Compilers for Parallel Computing, Springer Verlag, 433–440.

LCPC04 *A Code Isolator: Isolating Code Fragments from Large Programs* Y. Lee and M. Hall, Lecture Notes in Computer Science, 2005, Volume 3602, Languages and Compilers for High Performance Computing, Springer Verlag, Page 922.

Book Chapters, cont.

LCPC03 *Search space properties for coarse-grain pipelined FPGA applications* H. Ziegler, M. Hall and B. So, Lecture Notes in Computer Science, 2004, Volume 2958, Languages and Compilers for Parallel Computing, 1–16.

LCPC02 *Bridging the Gap between Compilation and Synthesis in the DEFACTO System* P. Diniz, M. Hall, J. Park, B. So and H. Ziegler, Lecture Notes in Computer Science, 2003, Volume 2624, Languages and Compilers for Parallel Computing, Springer Verlag, Pages 52–70.

IM01 *Memory Management in a PIM-Based Architecture* M. Hall and C. Steele, Lecture Notes in Computer Science, 2001, Volume 2107, Intelligent Memory Systems, 104–121.

RAW99 *DEFACTO: A design environment for adaptive computing technology* K. Bondalapati, P. Diniz, P. Duncan, J. Granacki and M. Hall, et al. Lecture Notes in Computer Science, 1999, Volume 1586, Parallel and Distributed Processing, 570–578.

LCR97 *A Case for Combining Compile-Time and Run-Time Parallelization* S. Moon, B. So, B. Murphy and M.W. Hall, Lecture Notes in Computer Science, 1998, Volume 1511, Languages, Compilers, and Run-Time Systems for Scalable Computers, Springer Verlag, 91–106.

LCPC95 *Interprocedural analysis for parallelization* M. W. Hall, B. Murphy, S. Amarasinghe, S. Liao and M.S. Lam, Lecture Notes in Computer Science, 1996, Volume 1033, Languages and Compilers for Parallel Computing, Springer Verlag, 61–80.

LCPC93 *FIAT: A Framework for Interprocedural Analysis and Transformation* M.W. Hall, J.M. Mellor-Crummey, A. Carle and R. Rodriguez. Lecture Notes in Computer Science, 1994, Volume 768, Languages and Compilers for Parallel Computing, Springer Verlag, 522–545.

Refereed Workshop and Poster Publications

LLVM21 *Optimizing Data Layout Transformations in MLIR*, Mahesh Lakshminarasimhan, M. Hall, P. Sadayappan (2021). Seventh Annual Workshop on the LLVM Compiler Infrastructure in HPC, Nov. 2021.

MAPS21 *Predictive data locality optimization for higher-order tensor computations* T. R. Patabandi, A. Venkat, A. Kulkarni, P. Ratnalikar, M. Hall, and J. Gottschlich. 2021. Proceedings of the 5th ACM SIGPLAN International Symposium on Machine Programming (MAPS 2021).

PMBS20 *Autotuning PolyBench Benchmarks with LLVM Clang/Polly Loop Optimization Pragmas Using Bayesian Optimization* X. Wu, M. Kruse, P. Balaprakash, H. Finkel, V. Taylor, P. Hovland, M. Hall. Workshop on Performance Measurement, Benchmarking and Simulation, held in conjunction with SC'20, Nov. 2020.

IWOMP19 *A Framework for Enabling OpenMP Autotuning* V. Sreenivasan, R. Javali, M. Hall, P. Balaprakash, T. Scogland, B. de Supinski. International Workshop on OpenMP, Sept. 2019.

HPCC19 *Rigel: A Framework for OpenMP Performance. IEEE International Conference on High Performance Computing and Communication (poster paper) Tuning* P. Rameshka, P. Senanayake, T. Kannangara, P. Seneviratne, S. Jayasena, T. Rusira, M. Hall, Aug. 2019.

P3HPC18 *Delivering Performance-Portable Stencil Computations on CPUs and GPUs Using Bricks*, T. Zhao, S. Williams, M. Hall, H. Johansen, International Workshop on Performance, Portability and Productivity in HPC, Nov. 2018, held in conjunction with SC'18.

PPoPP18 *SIMD Code Generation for Stencils on Brick Decompositions*, Tuowen Zhao, Mary Hall, Protonu Basu, Samuel Williams and Hans Johansen, (poster paper), Proceedings of the 24th ACM symposium on Principles and practice of parallel programming (PPoPP '18), Feb. 2018.

IMPACT18 *Abstractions for Specifying Sparse Matrix Data Transformations*, Payal Nandy, Eddie Davis, Mahdi Soltan Mohammadi, Wei He, Mary Hall, Michelle Strout and Catherine Olschanowsky, International Workshop on Polyhedral Compilation Techniques, Jan. 2018.

IWAPT17 *Automating Compiler-Directed Autotuning for Phased Performance Behavior*, T. Rusira, M. Hall and P. Basu, in Proceedings of the International Workshop on Automatic Performance Tuning, held in conjunction with IPDPS'17, June 2017.

MERC17 *Parameterized Diamond Tiling for Parallelizing Stencil Computations* T. Dilshan, B. Senevirathne, C. Siriwardhane, V. Wijesinghe, S. Jayasena, T. Rusira, M. Hall, Moratuwa Engineering Research Conference, May 2017.

IA^{3}16 *Compiler Transformation to Generate Hybrid Sparse Computations*, H. Zhang, A. Venkat and M. Hall, Proceedings of the Sixth Workshop on Irregular Applications: Architectures and Algorithms, held in conjunction with SC16, Nov. 2016.

SC16 *A Novel Variable-Blocking Representation for Efficient Sparse Matrix-Vector Multiply on GPUs* T. Zhao, T. Rusira, K. Ahmad, and M. Hall, (poster), SC16, November, 2016.

IMPACT16 *Combining Polyhedral and AST Transformations in CHiLL*, H. Zhang, A. Venkat, P. Basu, and M. Hall, in Proceedings of the International Workshop on Polyhedral Compilation Techniques at HiPEAC, Jan. 2016.

PPoPP15 *A Collection-Oriented Programming Model for Performance Portability*, S. Muralidharan, M. Garland, B. Catanzaro, A. Sidelnik, M. Hall, (poster paper), In Proceedings of the 21st ACM symposium on principles and practice of parallel programming (PPoPP '15), Feb. 2015.

Refereed Workshop and Poster Publications, cont.

PMBS14 *Roofline Model Toolkit: A Practical Tool for Architectural and Program Analysis*, Y. J. Lo, S. Williams, B. van Straalen, T.J. Ligocki, M.J. Cordery, N.J. Wright, M.W. Hall, L. Oliker, Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems at SC2014, Nov. 2014.

WOSC14 *Converting Stencils to Accumulations for Communication-Avoiding Optimization in Geometric Multigrid* P. Basu, S. Williams, B. van Straalen, L. Oliker, M. Hall, Workshop on Optimizing Stencil Computations (WOSC) at SPLASH 2014, Oct. 2014.

HIPS12 *Improving High-Performance Sparse Libraries using Compiler-Assisted Specialization : A PETSc Case Study*, Shreyas Ramalingam, M. Hall and C. Chen, Workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS), held in conjunction with International Parallel and Distributed Processing Symposium, May 2012.

PACT11 *Understanding the Behavior of Pthread Applications on Non-Uniform Cache Architectures*, G. S. Sachdev, K. Sudan, M. W. Hall, and R. Balasubramonian, (poster paper), In Proceedings of the International Conference on Parallel Architectures and Compilation Techniques, Oct. 2011.

LCPC11 *Generating High Performance Libraries using CHiLL and Autotuning*, S. Ramalingam and M. Hall, (poster), International Workshop on Languages and Compilers for Parallel Computing, Sept. 2011. *Received Best Poster award.*

PPoPP11 *Evaluating graph coloring on GPUs," P. Grosset, P. Zhu, S. Liu, S. Venkatasubramanian, and M. Hall* In Proceedings of the 16th ACM symposium on Principles and practice of parallel programming (PPoPP '11), Feb. 2011. *Received runner-up for Best Student Poster.*

SC10 *CUDA-CHiLL: Using Compiler-Based Autotuning to Generate High-Performance GPU Libraries*, M. Khan, G. Rudy, C. Chen, M. Hall, J. Chame, (poster) SC'10, Nov. 2010.

GPU10 *Automatic High-Performance GPU code Generation using CUDA-CHiLL*, (poster) Ma lik Khan, Jacqueline Chame, Gabe Rudy, Chun Chen, Mary Hall, Mark Hall, Nvidia GPU Technology Conference, Sept. 2010.

SAAHPC10 *Takagi Factorization on GPU using CUDA*, (poster paper), Gagandeep S. Sachdev, Vishay Vanjani and Mary W. Hall, Symposium on Application Accelerators for High Performance Computing, July, 2010.

SAAHPC10 *GPU Accelerated Particle System for Triangulated Surface Meshes*, (poster paper), B. Peterson, M. Datar, M. Hall and R. Whitaker, Symposium on Application Accelerators for High Performance Computing, July, 2010.

SC09 *Autotuning and Specialization: Speeding up Nek5000 with Compiler Technology*, (poster) J. Shin, M. W. Hall, J. Chame, C. Chen, P. F. Fischer, P. D. Hovland, SC'09, Nov. 2009.

IWAPT09 Autotuning and Specialization: Speeding up Matrix Multiply for Small Matrices with Compiler Technology,"Jaewook Shin, Mary W. Hall, Jacqueline Chame, Chun Chen, Paul D. Hovland, International Workshop on Automatic Performance Tuning, October, 2009.

SAAHPC09 *Assembling Large Mosaics of Electron Microscope Images using GPU*, (poster paper) Kannan Venkataraju, Mark Kim, Dan Gerszewski, James R. Anderson, and Mary Hall,Symposium on Application Accelerators for High Performance Computing, July, 2009.

FPGA09 *Computation reuse in domain-specific optimization of signal recognition'*, (poster paper) Melina Demertzi, Pedro C. Diniz, Mary W. Hall, Anna C. Gilbert, and Yi Wang, In Proceeding of the ACM/SIGDA international symposium on Field programmable gate arrays (FPGA '09), Feb. 2009.

Refereed Workshop and Poster Publications, cont.

HPPMSE *Model-Guided Performance Tuning of Parameter Values: A Case Study with Molecular Dynamics Visualization*, Y. Nelson, B. Bansal, M. Hall, A. Nakano, and K. Lerman, Proceedings of the Workshop on High-Level Parallel Programming Models and Supportive Environments, held in conjunction with IPDPS '08, April, 2008.

POHL07 *Model-Guided Empirical Optimization for Multimedia Extension Architectures: A Case Study*, C. Chen, J. Shin, S. Kintali, J. Chame and M. Hall, Proceedings of the Workshop on Performance Optimization of High-Level Languages, held in conjunction with IPDPS '07, March, 2007.

HPEC06 *CEARCH: Cognition Enabled Architecture*, S. Crago, J. McMahon, C. Archer, K. Asanovic, R. Chaung, K. Goolsbey, M. Hall, C. Kozyrakis, K. Olukotun, U. O'Reilly, R. Pancoast, V. Prasanna, R. Rabbah, S. Ward, D. Yeung, Proceedings of the Tenth Annual High Performance Embedded Computing Workshop, September, 2006.

DaMoN06 *Processing-in-memory technology for knowledge discovery algorithms*, J. Adibi, T. Barrett, S. Bhatt, H. Chalupsky, J. Chame, M. Hall. Second International Workshop on Data Management on New Hardware (DaMoN 2006), held in conjunction with SIGMOD, June 2006.

SC05 *PIMs in Action: Delivering Memory Bandwidth* (poster), T. Barrett, S. Bhatt, J. Chame, J. Draper, M. Hall, SC'05, Nov. 2005.

PHPC05 *A Systematic Approach to Composing and Optimizing Application Workflows*, E. Deelman, M. Hall, Y. Gil, K. Lerman, and J. Saltz, In *Proceedings of the Workshop on Patterns in High Performance Computing*, May, 2005.

MSP04 *Evaluating Compiler Technology for Control-Flow Optimizations for Multimedia Extension Architectures*, J. Shin, M. Hall and J. Chame. In *Proceedings of the Media-Specific Processors Workshop*, held in conjunction with MICRO '04, December, 2004.

MSP02 *A Compiler Algorithm for Exploiting Page-Mode Memory Access in Embedded-DRAM Devices*, J. Shin, J. Chame, and M. Hall, In Proceedings of the Media-Specific Processors Workshop, held in conjunction with MICRO '02, November, 2002.

SMW00 *Compiler Transformations for Exploiting Bandwidth in PIM-Based Systems*, J. Chame, M. Hall, and J. Shin. *Proceedings of Solving the Memory Wall Workshop*, held in conjunction with the International Symposium on Computer Architecture, June 2000.

RAW99 *Parallelizing Compiler Technology for Adaptive Computing Systems*, B. So, H. Ziegler and M.W. Hall. In *Proceedings of the Workshop on Reconfigurable Computing*, held in conjunction with Conference on Parallel Architectures and Compilation Techniques, October, 1999.

SUIF97 *Adaptive Parallelism for Compiler-Parallelized Code*, M.W. Hall and M. Martonosi. In *Proceedings of 2nd SUIF Compiler Workshop* (Aug. 1997).

CPC95 *Overview of Interprocedural Parallelization Analysis*, M.W. Hall, S. Amarasinghe, B. Murphy, S. Liao and M. Lam. In *Proceedings of the Fifth International Workshop on Compilers for Parallel Computers* (June 1995).

INVITED PRESENTATIONS

Keynotes and Plenary Talks

Salishan21	*How Programming Systems Meet the Needs of New Architecture Classes: Past, Present and Future* Panel, DOE Salishan Conference on High Speed Computing, Apr. 2021.
CGO21	*Data Layout and Data Representation Optimizations to Reduce Data Movement* Keynote, ACM International Conference on Code Generation and Optimization, Mar. 2021.
WACCPD	*Achieving Performance Portability for Extreme Heterogeneity*, Keynote, Workshop on Accelerator Programming Using Directives, held in conjunction with SC20, Nov. 2021.
IA320	*Research Challenges in Compiler Technology for Sparse Tensors* Keynote, Workshop on Irregular Applications, Algorithms and Architectures, held in conjunction with SC20, Nov. 2020.
NITRD20	*Preparing the Workforce for Extreme Heterogeneity*, Plenary Talk, NITRD High-End Computing Workshop, Sept. 2020.
OSTP20	*High Performance is all about Minimizing Data Movement* Keynote, OSTP Convening, Aug. 2020.
HPDC20	*High Performance is all about Minimizing Data Movement* Keynote, ACM International Symposium on High-Performance Parallel and Distributed Computing (virtual), June 2020.
PLEMM19	*Leveraging Overlapping Goals: Compilers for High-Performance Computing and Deep Learning* Programming Languages Enthusiasts Mind Melt, Facebook Research, Sept. 2019.
PPAM19	*What is Driving Programming System Technology for Exascale and Beyond* Plenary, International Conference on Parallel Processing and Applied Mathematics, Sept. 2019.
ICPP18	*Bringing Sparse Computations into the Optimization Light* Plenary, International Conference on Parallel Processing, Aug. 2018.
LHAM15	*The Role of Compiler Optimization and Autotuning for Reducing Data Movement in High-Performance Applications* Keynote, Legacy HPC Application Migration Workshop, in conjunction with CANDAR 2015, Sapporo, Japan, Dec. 2015. Seminar, University of Tokyo, Dec. 2015.
WMCA15	*Programmer Productivity for Current and Future High-Performance Architectures* Keynote, Career Workshop for Women and Minorities in Computer Architecture, held in conjunction with ACM MICRO, Dec. 2015.
AD12	*Autotuning Compiler and Language Technology and its Role in Exascale Systems* Invited speaker, Plenary, International Conference on Automatic Differentiation, July 2012.
FCCM11	*Compiler-Based Autotuning for Productive Parallel Programming and its Relationship to Design Space Exploration* High-Level Synthesis and Parallel Computation Models Workshop held in conjunction with IEEE International Symposium on Field-Programmable Custom Computing Machines, May, 2011.

Mentoring Talks

CRA22	*Mentoring Graduate Students* co-presenter with Nancy Amato, CRA Early Career Mentoring Workshop, Feb. 2022.
CRA20	*Mentoring Graduate Students* co-presenter with Mark Hill, CRA Early Career Mentoring Workshop, Feb. 2020.
SC19	*Career Development*, Early Career Program, SC'19, Nov. 2019.
PLMW19	*How to be a Great Researcher* Programming Languages Mentoring Workshop, held in conjunction with PLDI'19, June 2019.
SC18	*A Renaissance for Domain-Specific Languages, Compilers and Code Generators in HPC and Big Data* Students@SC, SC'18, Nov. 2018.
CRAW18	*25th Anniversary of the CRA-W*, panelist, CRA-W Career Mentoring Workshop, Nov. 2018.
CRAW18	*Navigating the Job Search* co-presenter with Nancy Amato and Kathryn McKinley, CRA-W Career Mentoring Workshop, Nov. 2018.
CRA18	*Time Management and Work-Life Balance* co-presenter with Carla Brodley, CRA Early Career Mentoring Workshop, Feb. 2018.
SC17	*Writing Effective Proposals* Early Career Program, SC'17, Nov. 2017.
SC16	*Getting People to Listen* Students@SC, SC'16, Nov. 2016.
SC16	*The Role of an Autotuning Compiler in Getting to Exascale* Undergraduate Program, SC'16, Nov. 2016.
RMCWiC16	*Programming Exascale Computers* Flash Talk, Rocky Mountain Celebration of Women in Computing, October, 2016.
PLMW16	*Getting People to Listen* Programming Languages Mentoring Workshop, held in conjunction with PLDI'16, June 2016.
CRA16	*How to write a good proposal: Tips, insights, and Perspective* co-presenter with Susanne Hambrusch, CRA Early Career Mentoring Workshop, Feb. 2016.
SC15	*Making the most of grad school (in HPC)* Students@SC Education Panel, SC'15, Nov. 2015.
Rice14	*Autotuning Software, Scientific Simulation and Academic Careers* Invited Speaker for CS Lunch and Learn (undergraduates), Rice University, Oct. 2014.
SC12	*Programming Exascale Supercomputers* Broader Engagement Program, SC12, November 2012.
CRAW07	*Hot Topics and Future Directions in Programming Languages* Panelist and Moderator, CRA-W Programming Languages Summer School, May 2007.
CRAW03	*Time Management, Family, and Quality of Life Issues* Panelist, Committee on the Status of Women in Computing Research (CRA-W) Workshop, June 2003.

Invited Conference and Workshop Presentations

LCPC19	*Optimizing Data Movement and Achieving Performance Portability with Brick Data Layouts*, International Workshop on Languages and Compilers for Parallel Computing (LCPC'19), Oct. 2019.
ANL19	*Packed Data Layouts as an Abstraction for Performance Portability Across Diverse Memory Systems*, Microelectronics Workshop, Argonne National Laboratory, Oct. 2019.
Salishan19	*How do we leverage deep learning investments to pay for HPC compiler technology?* Random Access Talk, DOE Salishan Conference on High Speed Computing, April 2019.
CCDSC18	*What Will it Take to Mainstream Autotuning Compiler Technology? Crowdsourcing Autotuning (Compiler) Research*, Workshop on Clusters, Clouds, and Data for Scientific Computing (CCDSC '18) Sept. 2018.
CEED17	*Y-TUNE: Autotuning Compiler Technology for Cross-Architecture Transformation and Code Generation* CEED Co-Design Center Annual Meeting, Aug. 2017.
Dagstuhl17	*Performance Portability Using Compiler-Directed Autotuning* Dagstuhl Seminar on Performance Portability in Extreme Scale Computing: Metrics, Challenges, Solutions, Oct. 2017.
BW17	*Using CHiLL and CUDA-CHiLL PAID IME* Blue Waters Symposium, May 2017.
IWAPT16	*Compiler Optimization, Specialization and Autotuning: Achieving Productivity and High Performance for Diverse Architectures* International Workshop on Application Performance Tuning (IWAPT'16), held in conjunction with IPDPS'16, May 2016.
SIAMCSE15	*The Role of Autotuning Compiler Technology* Streamlining Application Performance Portability Minisymposium, SIAM Conference on Computational Science and Engineering, March, 2015.
CCDSC14	*Leveraging HPC Expertise and Technology in Data Analytics* Workshop on Clusters, Clouds, and Data for Scientific Computing (CCDSC '14), Sept. 2014.
USHC14	*The Creative Process in Inventing Computer Science at the University of Utah* 62nd Annual Utah State History Conference, Sept. 2014.
SIAMPP14	*Tiling Dense and Sparse Computations for Parallelism and the Memory Hierarchy of GPUs* SIAM Parallel Processing Symposium, Feb. 2014.
Salishan12	*Automating Application Mapping with Autotuning: Paving the Way to Exascale* Salishan Conference on High-Speed Computing, April 2012.
Exa11	*Autotuning Compilers: Paving the Way to Exascale* Invited Talk, Joint DOE ASCR and NNSA Exascale PI meeting, Oct. 2011.
USC11	*Compiler-Based Autotuning of Energy Applications* USC-DOE Conference on Materials for Energy Applications: Experiment, Modeling and Simulations, March, 2011.
CMU10	*A Programming Language Interface to Describe Transformations and Code Generation for Auto-Tuning* ASPLOS Program Committee Symposium, CMU, October 2010.

Invited Conference and Workshop Presentations, cont.

CScADS10	*Compiler-Based Auto-tuning for Application and Library Code* Invited talk, DOE SciDAC Center for Scalable Application Development Software Workshop on Libraries and Autotuning for Petascale Applications, August, 2010.
FF10	*Paving the Way for Programming Extreme Scale Systems* DOE Institute for Computing in Science, Future of the Field Workshop, Jul. 2010.
SIAMPP10	*Collaborative Autotuning of Scientific Applications* SIAM Parallel Processing Symposium, Feb. 2010.
CScADS09	*Big Questions in Autotuning* DOE SciDAC Center for Scalable Application Development Software Workshop on Libraries and Autotuning for Petascale Applications, August, 2009.
BMPG08	*A Compiler-Based Strategy for Performance Tuning of Scientific Applications* SC08 Workshop on Bridging Multicore's Programmability Gap, Nov. 2008.
SIAM08	*The Role of Compiler Technology in Managing the Complexity of Architectures and Applications* SIAM Annual Meeting, July, 2008.
Intel08	*TUNE: The Structure of an Autotuning Compiler* Intel Academic Software Workshop, May, 2008.
HPCSW08	*Multicore Chips and Parallel Programming* High Performance Computer Science Week, April, 2008.
SIAMPP08	*Autotuning of Scientific Applications* 13th SIAM Conference on Parallel Processing for Scientific Computing, March, 2008.
CITI007	*A Compiler-Based Strategy for Performance Tuning of Applications* invited speaker, Workshop in Memory of Ken Kennedy, Rice University, Dec. 2007.
PLHPC06	*Programming Model Building Blocks: Flexible Compilers, Tunable Components and Workflows* invited speaker, Workshop on Programming Languages for High Performance Computing, Dec. 2006.
GPGPU06	*Strategies for High-Performance Heterogeneous Applications: A Compiler Perspective* Workshop on General-Purpose GPU Computing: Practice And Experience, held in conjunction with SC'06, Nov. 2006.
LACSI06	*Model-Guided Empirical Optimization* Autotuning Workshop, Los Alamos Computer Science Institute Symposium, Santa Fe, NM, Oct. 2006.
OCASA05	*Leveraging Parallelizing Compiler Technology for System-on-a-Chip Configurable Architectures* Optimizing Compiler Assisted SoC Assembly Workshop (OCASA), held in conjunction with the Compilers, Architectures and Synthesis for Embedded Systems, Sept. 2005.
SoCal02	*Automated Design Space Exploration* SoCal Workshop, UC Irvine, March, 2002.
ISPDP95	*The Effectiveness of Interprocedural Automatic Parallelization* Seventh IEEE Symposium on Parallel and Distributed Processing, Oct. 1995.

Panels

Xpert20	*Boosting RSE productivity through advanced tools* Panel Chair, Xpert Network, Oct. 2020
LCPC19	*Is a Post Moores Law Compiler Renaissance needed?* International Workshop on Languages and Compilers for Parallel Computing, Oct. 2019.
Rice19	*Inclusion* Rice CS 35th Anniversary, Oct. 2019.
Rice19	*Systems* Rice CS 35th Anniversary, Oct. 2019.
Snowbird18	*Increasing Diversity in Computing is Easier than you think: Some Small Steps that can Make a Big Difference*, Panel organizer and speaker, CRA Snowbird'18, July 2018.
IMPACT17	*Challenges for the Polyhedral Community* Panel Organizer and Panelist, International Workshop on Polyhedral Compilation Techniques (IMPACT'17), held in conjunction with HiPEAC'17, Jan. 2017.
LCPC16	*Compilation for dynamic parallel languages* Workshop on Languages and Compilers for Parallel Computing, September, 2016.
Cluster16	*HPC vs. Big Data: Different Worlds or Common Ground?* Panel Organizer and Moderator, IEEE Cluster, Sept. 2016.
Snowbird16	*Booming Enrollments: Student Profiles/Motivations* Panel Organizer and Panelist, CRA Conference at Snowbird, July 2016.
IMPACT16	*(When) are the polyhedral techniques necessary?* International Workshop on Polyhedral Compilation Techniques (IMPACT'16), held in conjunction with HiPEAC'16, Jan. 2016.
SC13	*SC13 Silver Anniversary Panel: Retrospective in Supercomputing Technologies* Panel Organizer and Moderator, SC13, Nov. 2013.
Facebook13	*Parallel Programming: Expressiveness and Efficiency* Breakout Session Leader, Facebook Compilers Summit, Facebook, Aug. 2013.
CScADS10	*Next Generation Compiler* DOE SciDAC Center for Scalable Application Development Software Workshop on Libraries and Autotuning for Petascale Applications, Aug. 2010.
CGO08	*What Role Does Code Generation and Optimization Play for Multi-Core Enablement?* International Conference on Code Generation and Optimization, April 2008.
LACSI06	*Strategies for High-Performance Heterogeneous Applications: A Compiler Perspective* Los Alamos Computer Science Institute Symposium, Oct. 2006.
Intel03	*Development Tools for Reconfigurable Architectures* Panelist and Moderator, Intel Reconfigurable Computing and Communications Architecture Workshop, May 2003.
ASPLOS02	*Programming Models for Parallel and Distributed Programming* invited panelist at ACM Symposium on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Oct. 2002.
ICS02	*Programming Models for Future High Performance Computing Systems* International Conference on Supercomputing, June 2002.
RAW99	*Challenges for Reconfigurable Computing for the Next Decade* Workshop on Reconfigurable Computing, held in conjunction with Conference on Parallel Architectures and Compilation Techniques, Oct. 1999.
CASES99	*What are the Solutions for Embedded Systems?* anelist and Moderator, Workshop on Compilers, Architectures, and Software for Embedded Systems, Washington, DC, Oct. 1999.
CASES98	*Exploiting Parallelism in Embedded Systems* Panelist, Workshop on Compilers, Architectures, and Software for Embedded Systems, Washington, DC, Dec. 1998.

Colloquia

2019 *Bringing Sparse Computations into the Optimization Light* UC Santa Barbara, Oct. 2019.

2019 *Sparse Matrix and Tensor Representations: Access Order Determines Format* Nvidia Research, May 2019.

2018 *Bringing Sparse Computations into the Optimization Light* INRIA Ecole Normale Superiere Paris, Sept. 2018.

2017 *Y-Tune: Performance Portability Using Compiler-Directed Autotuning* Argonne National Laboratory, May 2017; Lawrence Berkeley National Laboratory, May 2017.

2016 *The Role of an Autotuning Compiler in Getting to Exascale* University of Queensland (via video), Nov. 2016.

2016 *Domain-Specific Optimization and Autotuning for Performance Portability of Supercomputing Applications* EE Department Colloquium Series, Brigham Young University, Mar. 2016.

2015 *The Role of Compiler Optimization and Autotuning for Reducing Data Movement in High-Performance Applications* University of Tokyo, Dec. 2015.

2015 *Autotuning Compiler and Library Technology for Sparse Matrix Computations* Texas A&M University, Oct. 2015.

2015 *Compiler Optimization of Computation and Communication in Stencils and Geometric Multigrid* Intel Research, Santa Clara, CA, Jul. 2015.

2014 *Making Compilers Work: Autotuning for High Performance Applications* Colloquia speaker, Rice University, Oct. 2014; University of Texas, Oct. 2014.

2011 *Compiler-Based Autotuning for Productivity and High Performance* Ohio State University, February 2011.

2009 *Autotuning Compiler Technology to Support Architectural Diversity* Dept. of Electrical and Computer Engineering, Brigham Young University, January, 2009.

2008 *The Role of Compiler Technology in Managing the Complexity of Architectures and Applications* Lawrence Berkeley National Laboratory, May, 2008.

2007 *Compiler-Assisted Performance Tuning* Argonne National Laboratory, Argonne, IL, July 2007.

2007 *Coping with Architectural Complexity using Compiler-Guided Empirical Search* Intel Corporation, May 2007.

2007 *A Compiler Strategy for Automatic Performance Tuning of Scientific Applications* Oak Ridge National Laboratory, Feb. 2007.

2006 *The Future of Compilers for Parallel Architectures, from Systems-on-a-Chip to Supercomputers* Univ. of Illinois, Urbana-Champaign Distinguished Lecture, March, 2006.

Colloquia, cont.

2001	*Compiler-Driven Synthesis through Behavioral Synthesis Estimation* Synopsys, San Jose, CA, Apr. 2001.
1995-1996	*Precise Interprocedural Analysis* University of California at Irvine, Nov. 1995; University of California at Los Angeles, Apr. 1996.
1995	*Detecting Coarse-Grain Parallelism Using an Interprocedural Parallelizing Compiler* University of California at Santa Barbara, May 1995; University of Southern California, Los Angeles, CA, Oct. 1995.
1994	*An Interprocedural Compiler for Shared-Memory Multiprocessors* Silicon Graphics Compiler Summit, San Jose, CA, July 1994; University of Illinois at Urbana-Champaign, Sept. 1994.
1994	*Interprocedural Analysis for Parallelizing Compilers* University of Michigan, Ann Arbor, MI, March 1994; University of Texas, Austin, TX, March 1994; University of Maryland, College Park, MD, April, 1994; Duke University, Durham, NC, April, 1994; University of North Carolina, Chapel Hill, NC, April, 1994; University of California at Los Angeles, April, 1994.
1993	*Interprocedural Transformations for Parallel Code Generation* Ecole des Mines, Fountainebleau, France, April 1993.
1993	*Procedure Cloning* Sun Research Labs, Mountain View, CA, Feb. 1993.
1991-1992	*Managing Interprocedural Optimization* Digital Equipment Corporation, Nashua, NH, Sept. 1991; Nasa RIACS, Mountain View, CA, Oct. 1991; University of Arizona, Tucson, AZ, Dec. 1991; North Carolina State University, Raleigh, NC, March 1992.

PROFESSIONAL SERVICE

Technical Engagement with Funding Agencies

DOE21	*Program Synthesis for Scientific Computing* Report for DOE ASCR Workshop on Program Synthesis for Scientific Computing, H. Finkel and I. Laguna, editors, Jan. 2021.
NSF20	*Report of NSF CISE Committee of Visitors*, Committee Co-Chair, NSF CISE, Mar. 2020.
DOE18	*Extreme Heterogeneity 2018: Productive Computational Science in the Era of Extreme Heterogeneity* Report for DOE ASCR Basic Needs Workshop on Extreme Heterogeneity, Jan. 2018, J.S. Vetter, R. Brightwell, M. Gokhale, P. McCormick, R. Ross, J. Shalf, K. Antypas, D. Donofrio, A. Dubey, T. Humble, C. Schuman, B. Van Essen, S. Yoo, A. Aiken, D. Bernholdt, S. Byna, K. Cameron, F. Cappello, B. Chapman, A. Chien, M. Hall, R. Hartman-Baker, Z. Lan, M. Lang, J. Leidel, S. Li, R. Lucas, J. Mellor-Crummey, P. Peltz, Jr., T. Peterka, M. Strout, and J. Wilke, Extreme Heterogeneity 2018: DOE ASCR Basic Research Needs Workshop on Extreme Heterogeneity, US Department of Energy, Office of Science, Advanced Scientific Computing Research, 2018, doi:10.2172/1473756.
DOE14	*Report of the DOE Programming Models Summit* Washington DC, Feb. 5, 2014, Michael Garland (NVIDIA), Mike Heroux (Sandia), Richard Lethin (Reservoir Labs), Sonia Sachs (DOE ASCR), Vivek Sarkar (Rice), John Shalf (LBNL), Marc Snir (UIUC), Armando Solar-Lezama (MIT), Thomas Sterling (Indiana), Kathy Yelick (LBNL).
DOE11	*Exascale Programming Challenges Report of the 2011 Workshop on Exascale Programming Challenges Marina del Rey, July 27-29, 2011* Saman Amarasinghe (MIT), Mary Hall (U. Utah), Richard Lethin (Reservoir Labs), Keshav Pingali (U. Texas-Austin), Dan Quinlan (LLNL), Vivek Sarkar (Rice U.), John Shalf (LBNL), Robert Lucas (USC/ISI), Katherine Yelick (LBNL), Pavan Balaji (ANL), Pedro C. Diniz (USC/ISI), Alice Koniges (LBNL), Marc Snir (ANL), Sonia R. Sachs (ASCR), October 2011. Released, 10/2011.
NSF10	*NSF Workshop on Archiving Experiments to Raise the Level of Scientific Research* Co-organizer and report co-author, Salt Lake City, 2010.
NSF09	Member of Organizing Committee and Host of NSF Future of Compiler Research and Education Workshop, Feb. 2007; Report co-author (CACM, Feb. 2009).

Appointed Roles in ACM and IEEE

2022	Member, IEEE von Neumann Medal Award Committee
2021	Member, IEEE von Neumann Medal Award Committee
2020	Member, ACM SIGHPC Dissertation Award
2019	Member, ACM SIGHPC Emerging Women Leader Award Committee
2018	Member, ACM SIGHPC Dissertation Award
2008-2009	Member, ACM and IEEE Computer Society Ken Kennedy Award Committee
2012-2013	Member, ACM and IEEE Computer Society Ken Kennedy Award Committee
2013	Member, ACM SIGPLAN Robin Milner Young Researcher Award Committee
2009	Member, IEEE Computer Society Sidney Fernbach Award Committee
2009	Member, IEEE Computer Society Seymour Cray Award Committee
2006	Member, ACM Health of Conferences Committee
2005-2008	Special Interest Group (SIG) Liaison to ACM History Committee

Technical Program Chairs and Editorial Positions

SC21	Technical Program Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
LCPC18	Program Co-Chair, International Workshop on Languages and Compilers for Parallel Computing
IMPACT17	Program Co-Chair, International Workshop on Polyhedral Compilation Techniques
SC14	Technical Papers Co-Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
PPoPP10	Program Chair, ACM SIGPLAN Principles and Practice of Parallel Programming
CGO09	Program Chair, ACM/IEEE Code Generation and Optimization (CGO)
PLDI05	Program Chair, ACM Programming Language Design and Implementation
PACT05	Program Co-Chair, ACM/IEEE Parallel Architecture and Compilation Techniques
SC04	Software Area Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
2005-2008	Associate Editor, ACM Computing Surveys, 2005-2008.

Conference Steering Committees

2009-	ACM SIGPLAN Principles and Practice of Parallel Programming
2010-2013	ACM SIGPLAN Programming Language Design and Implementation
2005-2007	ACM SIGPLAN Programming Language Design and Implementation
2009-2012	ACM/IEEE Code Generation and Optimization
2003-2005	ACM/IEEE Parallel Architecture and Compilation Techniques

Conference and Workshop Organization

SC21	Technical Program Chair, ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis
BPC20b	Co-Organizer, *Departmental Plans for Broadening Participation in Computing for NSF CISE Grants*, July and August 2020.
BPC20a	Co-Organizer, *Developing Department Plans for Broadening Participation in Computing for NSF Grants*, Pre-symposium event, SIGCSE'20, Mar. 2020.
BPC19	Co-Organizer, *Departmental Plans for Broadening Participation in Computer Science: A Planning Workshop for Computing Department Leadership*, UIUC
MLIR4HPC	Co-Organizer, Workshop on MLIR For HPC
FCRC19	Plenary Speaker Chair, ACM Federated Computing Research Conference
OMASE19	Workshop Co-Organizer, Workshop on Optimization, Modeling, Analysis and Space Exploration (OMASE), co-located with CGO'19
SC18	Test of Time Award Co-Chair, ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis
Cluster16	Panel Chair, IEEE Cluster Computing Conference
SC15	Test of Time Award Co-Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
SC14	Tech Papers Co-Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
SC13	Silver Anniversary Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
SC12	Exhibits Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
PLDI11	General Chair, ACM Programming Language Design & Implementation
SC09	Awards Chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
SC08	Workshop Co-chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
ATLA08	Organizer, Automatic Tuning of Libraries and Applications Workshop
SC07	Workshop Co-chair, ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
STMCS07	Organizer, ACM/IEEE Workshop on Software Tools for Multi-Core Systems
STMCS06	Originator, ACM/IEEE Workshop on Software Tools for Multi-Core Systems
SC05	Poster Chair, CM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
SPPA04	Co-organizer, Software for PIM-Based Parallel Architectures Workshop
ICS00	Workshop Chair, International Conference on Supercomputing
SoCal97	Organizer, SoCal Parallel Systems Seminar, USC/ISI
SoCal95	Originator, SoCal Parallel Systems Seminar, Caltech
PEdW91	Organizer, ParaScope Editor Workshop, Rice

Program Committees

PLDI22	ACM Programming Language Design & Implementation Program Comm.
CGO21	ACM/IEEE Code Generation and Optimization, Program Comm.
LCPC20	Workshop on Languages and Compilers for Parallel Computing
ICS20	ACM International Conference on Supercomputing, Extended Review Comm.
IPDPS20	IEEE Parallel and Distributed Processing Symposium, Extended Review Comm.
ASPLOS20	ACM Architectural Support for Programming Languages and Systems, Extended Review Comm.
LCPC19	Workshop on Languages and Compilers for Parallel Computing
FHPNC19	ACM Workshop on Functional High-Performance and Numerical Computing
EuroPar19	International European Conference on Parallel & Distributed Computing
PLDI19	ACM Programming Language Design & Implementation, External Program Comm.
PPoPP19	ACM Principles and Practice of Parallel Programming, Extended Review Comm.
IMPACT19	International Workshop on Polyhedral Compilation Techniques
CANDAR18	International Symposium on Computing and Networking
PLDI18	ACM Programming Language Design & Implementation, External Program Comm.
PPoPP18	ACM Principles and Practice of Parallel Programming, Extended Review Comm.
IMPACT18	International Workshop on Polyhedral Compilation Techniques
CANDAR17	International Symposium on Computing and Networking
LCPC17	International Workshop on Languages and Compilers for Parallel Computing
EuroPar17	International European Conference on Parallel & Distributed Computing
MAPL17	ACM Workshop on Machine Learning and Programming Languages
Chapel17	Chapel Implementer and User Workshop
PPoPP17	ACM Principles and Practice of Parallel Programming
LHAM16	International Workshop on Legacy HPC Application Migration
ISC16	International Supercomputing Conference
PPoPP16	ACM Principles and Practice of Parallel Programming
IMPACT16	International Workshop on Polyhedral Compilation Techniques
PLDI15	ACM Programming Language Design & Implementation, External Review Comm.
ASPLOS15	ACM Architectural Support for Programming Languages and Systems
MSPC14	Workshop on Memory Systems Performance and Correctness
PLDI14	ACM Programming Language Design & Implementation
IMPACT14	International Workshop on Polyhedral Compilation Techniques
SC13	ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
CC13	International Conference on Compiler Construction
PLDI13	ACM Programming Language Design & Implementation, External Review Comm.
HiPEAC13	International Conference on High-Performance and Embedded Architectures and Compilers
SC12	ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
CGO12	ACM/IEEE Code Generation and Optimization
HiPEAC12	International Conference on High-Performance and Embedded Architectures and Compilers
IMPACT12	International Workshop on Polyhedral Compilation Techniques
LCPC11	International Workshop on Languages and Compilers for Parallel Computing
ICS11	ACM/IEEE International Conference on Supercomputing
ASPLOS11	ACM Architectural Support for Programming Languages and Systems
PLDI10	ACM Programming Language Design & Implementation, External Review Comm.
LCPC09	International Workshop on Languages and Compilers for Parallel Computing)
IWAPT09	International Workshop on Application Performance Tuning
PACT09	ACM/IEEE Parallel Architecture and Compilation Techniques

Program Committees, cont.

CGO09	ACM/IEEE Code Generation and Optimization
STMCS08	ACM/IEEE Workshop on Software Tools for Multi-Core Systems
PESPMA08	ACM Workshop on Parallel Execution of Sequential Progs. on Multi-core Archs.
SC07	ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
HPPMSE07	IEEE Workshop on High-Level Parallel Programming Models and Supportive Environments
LCTES07	ACM Languages, Compilers and Tools for Embedded Computing
CGO07	ACM/IEEE Code Generation and Optimization
UbiPar06	IEEE/ACM Workshop on Programming Models for Ubiquitous Parallelism
RAW06	IEEE Reconfigurable and Adaptive Architecture Workshop
ASPLOS06	ACM Architectural Support for Programming Languages and Systems
PLDI05	ACM Programming Language Design & Implementation
CGO05	ACM/IEEE Code Generation and Optimization
SC04	ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
PLDI03	ACM Programming Language Design & Implementation
PPHEC03	ACM/IEEE Workshop on Productivity and Performance in High-End Computing
PACT03	ACM/IEEE Parallel Architecture and Compilation Techniques
SSRS03	ACM/IEEE Workshop on Software Support for Reconfigurable Systems
ICPP02	International Conference on Parallel Programming
LCTES02	ACM Languages, Compilers and Tools for Embedded Computing
LCTES01	ACM Languages, Compilers and Tools for Embedded Computing
ISCA01	ACM International Symposium on Computer Architecture
SC00	ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
ICS00	ACM/IEEE International Conference on Supercomputing
HiPC00	International Conference on High-Performance Computing
IPDPS00	IEEE International Parallel and Distributed Processing Symposium
PACT00	ACM/IEEE Parallel Architecture and Compilation Techniques
CASES00	Compilers and Architecture Support for Embedded Systems
LCR00	Workshop on Languages, Compilers, and Run-Time Systems for Scalable Computers
SC99	ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
ICPP99	International Conference on Parallel Programming
IPPS99	IEEE International Parallel Processing Symposium
PACT99	ACM/IEEE Parallel Architecture and Compilation Techniques
PPoPP99	ACM Principles and Practice of Parallel Programming
ICS98	ACM/IEEE International Conference on Supercomputing
SC97	ACM/IEEE International Conference for High Performance Computing, Networking Storage and Analysis
ICPP97	International Conference on Parallel Programming
ICS96	ACM/IEEE International Conference on Supercomputing
LCTRS95	ACM Workshop on Languages, Compilers and Tools for Real-time Systems
OOPSLA94	ACM Objected-Oriented Programming Systems, Languages and Architectures
ICCL94	International Conference on Computer Languages

Professional Service to University of Utah

2015-	School of Computing Executive Committee
2019-2020	SCI Institute Director Search Committee member
2018-2019	College of Science Dean Search Committee member
2018-	William R. Gould Lecture Series Advisory Board
2010-2011	Jay Lepreau Named Professorship Subcommittee Chair
2010-2011	University SEED Grant Committee member
2009-2010	Leadership Development Program participant

MENTORING AND OUTREACH

Mentoring Roles at University of Utah

2016-	Chair, Women in Engineering Faculty Advisory Council
2019	*Computer Science* Presentation to ACCESS Summer Program
2009-19	Chair, School of Computing Diversity Committee
2018	*Getting People to Listen* Women in Computing Brown Bag Lunch Series, Apr. 2018
2018	Panelist, Pizza and Profs, School of Computing Undergraduate Peer Mentoring Program, Feb. 2018
2017	Mentor and Coach, University of Utah SC'17 Student Cluster Competition Team
2016	Mentor and Coach, University of Utah SC'16 Student Cluster Competition Team
2014	Impostor Panel Moderator, Women in Engineering Program
2013	Organizing Committee, Women's Week, Mar. 2013.
2013-2014	Faculty Liaison, Women in Engineering Program
2012-2016	Member, Women in Engineering Faculty Advisory Council
2012-2015	Member, Presidential Commission on the Status of Women
2010	Hi-GEAR Faculty Panelist
2010	University Women in Science organizing committee

Mentoring Events and Committees

CRA20	Speed Mentor, CRA Early Career Mentoring Workshop
CRAW18	Mentor, CRA-W Early Career Mentoring Workshop
SC16	Speed Mentor, SC'16 Early Career Mentoring Workshop
PLMW16	Co-Organizer, Programming Language Mentoring Workshop at PLDI'16
SC15	Speed Mentor, SC'15 Early Career Mentoring Workshop
SC15	Mentor, SC'15 Dinner with Interesting People
GHC15	Mentor, Grace Hopper Conference Scholarship Lunch
GHC13	Scholarship Committee, Grace Hopper Conference
GHC12	Scholarship Committee, Grace Hopper Conference
GHC10	Scholarship Committee, Grace Hopper Conference
GHC07	Scholarship Committee, Grace Hopper Conference
GHC06	Scholarship Committee, Grace Hopper Conference
GHC04	Phd Forum Committee, Grace Hopper Conference
1995	Directed undergraduate in Caltech CRPC *Summer Program in Computing for Undergraduate Women*
1992-93	Mentor for Phd student in Stanford American Women in Science Mentoring Program
1991	Directed undergraduate in Rice CRPC *Spend the Summer with a Scientist* program

FIRST Robotics K-12 Outreach

2017-	Co-Creator of FLL Jr Robotics Program for 3rd Graders at Salt Lake City Community Learning Centers, reaching 44 kids so far
2019-20	FIRST Dean's List Interviewer, Dec. 2019 and Jan. 2020
2014-19	Mentor and Coach, FIRST Tech Challenge (Grade 7-12 robotics)
2019	Scorekeeper, FIRST Tech Challenge Utah State Competition, Feb. 2019
2008	FIRST Dean's List Interviewer, Dec. 2018
2018	Co-Organizer, FIRST Utah Robotics Symposium, Aug. 2018
2018	Scorekeeper, FIRST Tech Challenge Utah State Competition, Feb. 2018
2018	Scorekeeper, FIRST Tech Challenge West High Qualifying Competition, Jan. 2018, Dec. 2018
2017	Scorekeeper, FIRST Tech Challenge Park City High Qualifying Competition, Dec. 2017
2017	Scorekeeper, FIRST Tech Challenge Worlds Competition, Apr. 2017
2017	Scorekeeper, FIRST Tech Challenge Utah State Competition, Feb. 2017
2017	Scorekeeper, FIRST Tech Challenge West High Qualifying Competition, Jan. 2017
2016	Scorekeeper, FIRST Tech Challenge Park City High Qualifying Competition, Dec. 2016
2016	Field Inspector, Utah FIRST Tech Challenge State Competition, Feb. 2016
2016	Judge, Utah FIRST Lego League State Competition, Jan. 2016
2015	Judge, Utah FIRST Lego League State Competition, Jan. 2015
2010-2014	Mentor and Coach, FIRST Lego League (Age 9-14 robotics)

Other K-12 Outreach

2016	Classroom engagement: West High School CS1 Class
2014	Classroom engagement: West ELP Middle School Magnet, Seventh Grade Science
2009	Classroom engagement: West ELP Middle School Magnet, Seventh Grade Science
2011	Classroom engagement: Hawthorne Elementary ELP Magnet, Fourth Grade Class
2010	Classroom engagement: Hawthorne Elementary ELP Magnet, Third Grade Class
2009	Classroom engagement: Ensign Elementary, Second Grade Class

TEACHING

Autotuning of Dense and Sparse Matrices on GPUs, International Summer School on Parallel High Performance Computing using Accelerators, University of Minho, Braga, Portugal, June, 2014.

Compiler-based autotuning technology, International Summer School on Advanced Computer Architecture and Compilation for High-Performance and Embedded Systems, Fiugi, Italy, July, 2011.

CS 4400 (Utah) Computer Systems, undergraduate lecture and project course, 2015.

CS 4230/4961 (Utah) Parallel Programming, undergraduate lecture and project course, 2009, 2010, 2011, 2012, 2013.

CS 6235/6963 (Utah) Parallel Programming for GPUs, graduate-level lecture and project course, 2009, 2010, 2011, 2012, 2013, 2014.

CSCI 503 (USC) Parallel Programming, graduate-level lecture and project course, 2008.

ENGR 102 (USC) Freshman Engineering Academy, introduction to engineering, lecture and group project course, 2007.

CSCI 599 (USC) Parallel Programming Paradigms, graduate-level lecture and project course, 2006.

CSCI 595 (USC) Compiling for High Performance, graduate-level lecture and project course, 2002, 2004.

CSCI 565 (USC) Advanced Compiler Design, graduate-level lecture and project course on compilers and optimization, 1998, 1999.

CSCI 599 (USC), seminar on compilation techniques for high-performance computing.

Western Institute of Computer Science, a 1-week intensive course entitled *Compilers for High-Performance Uniprocessors and Multiprocessors*, with M. Lam and M. Smith, Aug. 1995.

Comp 237 (Caltech), graduate-level lecture course on compilers, optimization and parallelization, Winter, Spring, Fall, 1995.

ADVISING

Current PhD Students

Khalid Ahmad, John Jolly, Mahesh Lakshminaranan, Payal Nandy, Tharindu Rusira, Tuowen Zhao

PhD Graduates

2016	Saurav Muralidharan, *Abstractions And Autotuning Techniques For Adaptive Programming*, Nvidia Research
2016	Anand Venkat, *An Integrated Compiler And Runtime Framework For Sparse Matrix Codes*, Intel Research, Santa Clara, CA
2016	Protonu Basu, *Compiler Optimizations and Autotuning for Stencil Computations and Geometric Multigrid*, Facebook, Santa Clara, CA
2012	Malik Muhammad Murtaza Khan (USC), *Autotuning, Code Generation and Optimizing Compiler Technology for GPUs*, Norwegian Institute of Science and Technology, Trondheim, Norway
2009	Yoonju Lee Nelson (USC), *Model-Guided Performance Tuning for Application-Level Parameters*, Intel, Phoenix, AZ
2007	Chun Chen (USC), *Model-Guided Empirical Optimization for Memory Hierarchy*, deceased.
2006	Heidi Ziegler (USC), *Automatic Mapping of Coarse-Grain Pipelined Applications to FPGA Systems*, Boeing Corporation, El Segundo, CA

PhD Graduates, cont.

2005	Jaewook Shin (USC), *Compiler Optimizations for Architectures Supporting Superword-Level Parallelism*, Nvidia, Santa Clara, CA
2003	Byoungro So (USC), tAn Efficient Design Space Exploration for Balance Between Computation and Memory, Intel Research, Santa Clara, CA
2002	Sungdo Moon (USC), *Combining Compile-Time and Run-Time Parallelization*, Samsung, Seoul, South Korea

Member of PhD Committee

Brad Peterson (2019), Diptorub Deb (UNC-Chapel Hill, 2019), Alan Humphries (2019), Jie Zhao (ENS-Paris, 2018), Thomas Gilray (2017), Thomas Nelson (University of Colorado, 2016), James King (2016), Sidharth Kumar (2016), Pascal Grosset (2016), Wei-Fan Chiang (2016), Sriram Aananthakrishnan (2016), Mohammed Sourouri (University of Oslo, 2015), Yang Chen (2014), Weibin Sun (2014), Qingyu Meng (2014), Tobias Grosser (ENS-Paris, 2014), Kevin Tew (2013), Zhisong Fu (2013), Kshitij Sudan (2012), Ahn Vo (2011), Guodong Li (2010), Justin Luitjens (2010), Jeonghee Shin (USC, 2008), Ling Zhuo (EE-Systems, 2007), Wei Hong Ho (USC, 2006), Panduka Wijetunga (USC, 2006), Brian Murphy (Stanford University, 2002), Stephen Jenks (USC, 2000), Xiaogang Qui (USC, 2000), Halima Elnaga (USC, 1999), Joonho Ha (USC, 1998)

Current Masters Students

Shenruoyang Na

Masters Graduates (with thesis)

2015	Yu-Jung Lo, *Performance Modeling for Architectural and Program Analysis*
2014	Axel Rivera, *Using Autotuning for Accelerating Tensor Contraction on GPUs*
2012	Shreyas Ramalingam, *Improving High-Performance Sparse Libraries using Compiler-Assisted Specialization: a PETSc Case Study*
2010	Gabe Rudy, *CUDA-CHiLL: A Programming Language Interface for GPGPU Optimizations and Code Generation*

Masters Graduates (with project):

Vinu Sreenivasan (2019), Khalid Ahmad (2017), Amit Roy (2016), Huihui Zhang (2016), Prajakta Mane (2015), Suchit Maindola (2012), Gagandeep Sachdev (2011)

Member of Masters Committee (thesis or project)

Devin Robinson (CENG, project, 2013) Mohammed Al-Mahfoudh (project, 2013), Preethi Kothari (project, 2012), Seungkeol Choe (project, 2012), Youngsung Kim (CENG, project, 2012), Sriram Aananthakrishnan (project, 2011), Shayan Chandrashekar (project, 2011), Shreyas Subramanya (project, 2011), Wei-Fan Chiang (thesis, 2010), Tarun Prabhu (project, 2010), Chad Allen (Mechanical Engineering, thesis, 2010), Grant Watts (project, 2010)

Undergraduate advisees David Fridlander, Hong Heung, Janaan Lake, Morgan Kelley, Terry Kingston, Meyer Saminemi (BS thesis committee, 2018)

GRANTS

DOE19 *PROTEAS-TUNE* Utah PI; PI: Jeff Vetter, Oak Ridge National Laboratory, DOE Exascale Computing Project, $1.1M to Utah, 10/19-06/23.

NSF18 *EAGER: BPCnet: A Broadening Participation Resource Portal* PI; Erik Russell, co-PI (Computing Research Association) NSF 1830364, $299,966, 05/18-04/20.

DOE17b *Institute for Resource and Application Productivity through computation, Information, and Data Science (RAPIDS)*, Department of Energy Scientific Discovery through Advanced Computation (SciDAC) Institute, $300K to Utah, 09/17-08/20.

NSF17 *SHF: Medium: Hierarchical Tuning of Floating-point Computations* co-PI, PI: Ganesh Gopalakrishnan, NSF SHF-1704715, $1.2M, 08/17-07/21.

DOE17a *Autotuning Compiler Technology for Cross-Architecture Transformation and Code Generation* PI, DOE Exascale Computing Project, PI, $2.7M ($871K to Utah), 03/17-09/19.

NSF16 *SHF: Medium: Collaborative Research: An Inspector/Executor Compilation Framework for Irregular Applications* PI, NSF CCF-1564074, $400K to Utah, 08/01/16-07/31/20.

NSF16 *EAGER: Application-driven Data Precision Selection Methods* Co-PI, NSF CCF-1643056, $300K, 08/16-07/18.

BlueWaters16 *Petascale Application Improvement Discovery (PAID) Improvement Method Enabler (IME) for FDTD models of the global Earth-ionosphere waveguide*, Co-PI, Jamesina Simpson, PI. $50K, NSF via UIUC Blue Waters project, 7/16-7/17.

BlueWaters15 *Petascale Application Improvement Discovery (PAID) Improvement Method Enabler (IME) for CPPTRAJ, the Analysis code for AMBER*, Co-PI, Tom Cheatham, PI. $50K NSF via UIUC Blue Waters project, 7/15-7/17.

BlueWaters15 *Petascale Application Improvement Discovery (PAID) Subcontract to Leadership class scientific and engineering computing: breaking through the limits (Blue Waters)*, PI, $100K NSF via UIUC Blue Waters project, 2/15-7/17.

NSA14 *Using Active Harmony and CHiLL to Autotune Chapel* Jeff Hollingsworth (PI), Mary Hall (co-PI), National Security Agency, $200K to Utah, 1/14-10/17.

NSF12 *CSR: Medium: Energy-Efficient Architectures for Emerging Big-Data Workloads* Rajeev Balasubramonian (PI), Mary Hall (co-PI), Al Davis (co-PI), Feifei Li (co-PI), $875K, National Science Foundation, CNS-1302663, 7/13-6/17.

DARPA12 *Osprey: Efficient Embedded Parallel Computing Technologies* Mary Hall (Utah PI), Steve Keckler (PI, Nvidia), University of Virginia, DARPA PERFECT program, 11/12-04/14.

DOE12 *Autotuning for Exascale: Self-Tuning Software to Manage Heterogeneity* Mary Hall (PI), Lenny Oliker (LBNL), Paul Hovland (Argonne), Jacqueline Chame (USC/ISI), $525K to Utah, Department of Energy Exascale Software Stack Program, 09/12-08/15.

NSF12 *SI2-SSE: Correctness Verification Tools for Extreme Scale Hybrid Concurrency* Ganesh Gopalakrishnan (PI), Mary Hall (co-PI), Rajeev Thakur (co-PI, U. Chicago), $444K, National Science Foundation, SI2-1148127, 6/12-5/15.

DOE11 *Institute for Sustained Performance, Energy, and Resilience (SUPER)* Mary Hall (Utah PI), Ganesh Gopalakrishnan (Senior Personnel), University of Southern California (lead), Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Argonne National Laboratory, Oak Ridge National Laboratory, Massachusetts Instute of Technology, University of California San Diego, University of Maryland, University of North Carolina, University of Oregon, and University of Tennessee, Department of Energy Scientific Discovery through Advanced Computation (SciDAC) Institute, ($1.075M to Utah), 09/11-08/16.

DARPA10 *Echelon: Extreme scale Compute Hierarchies with Efficient Locality Optimized Nodes* Mary Hall (Utah PI), Nvidia Corporation (Lead); Cray, Inc., Micron Technology, Inc., Oak Ridge National Labs, Georgia Institute of Technology, Stanford University, University of California - Berkeley; Georgia Institute of Technology; University of Texas; University of Pennsylvania, Defense Advanced Research Projects Agency, ($1.256M to Utah), 8/10- 05/14.

NSF10 *SHF Small: A Compiler-Based Auto-Tuning Framework for Many-Core Code Generation* M. Hall (PI), C. Chen (co-PI), ($316K), 07/10-06/12.

DOE10 Performance Engineering Research Institute SciDAC-e Augmentation: Performance Enhancement of Simulating the Dynamics of Photoexcitation for Solar Energy Conversion," R.F. Lucas (PI), J. Chame, P. Diniz, P. Vashishta, A. Nakano, R. Kalia, M.W. Hall, J. Hollingsworth, D. Quinlan (co-PIs), Department of Energy, ($150K to Utah), 07/10-12/11.

DOE10 *Compiler-Based Autotuning* Mary Hall (PI), Co-PI(s): Jacqueline Chame (USC/ISI); Paul Hovland, Jaewook Shin (Argonne) (Utah portion moved from USC), Department of Energy ($510K), 02/2010 - 01/2012.

DOE08 *Performance Engineering Research Institute* Mary Hall (Utah PI), (Utah portion moved from USC), Department of Energy ($290K) 12/08-01/12.

Utah09 *A Novel Approach to Achieving Thread-Based Parallelism in Simulation of Transport Processes* James Sutherland (PI), Mary Hall (co-PI), Martin Berzins (co-PI). University of Utah Seed Grant ($35K), 09/2009 - 08/2010.

NSF09 *SHF:SMALL: Hardware/Software Management of Large Multi-Core Memory Hierarchies* Rajeev Balasubramonian (PI), Mary Hall (co-PI) . National Science Foundation Grant CCF-0916436 ($372K), 09/2009-08/2012.

NSF07 *CRI: CRD: Raising the Standard of Scientific Publishing Through an Experiment Archive* National Science Foundation Grant CRI-0709430 ($50K), PI with Eric Eide (co-PI), took over for Jay Lepreau, 10/2007-09/2010.

DARPA08 *Exascale Software Study* Vivek Sarkar (Rice) PI, co-PI with Bill Carlson (IDA), Andrew Chien (Intel), Bill Dally (Stanford), Mootaz Elnozahy (IBM), Robert Harrison (ORNL), Norm Jouppi (HP), Chuck Koelbel (Rice), Dave Koester (Mitre), Peter Kogge (Notre Dame), John Levesque (Cray), Dan Reed (Microsoft), John Shalf (LBNL), Allan Snavely (SDSC), Thomas Sterling (LSU), DARPA, ($70K to ISI), 6/2008-9/2008.

DOE08 *Compiler-Directed Automatic Performance Tuning* PI, co-PIs Paul Hovland and Jaewook Shin (Argonne), Department of Energy DE-FG02-08ER25834, ($1.5M, $1M to ISI), 3/2008-2/2011.

DOE08 *Gyrokinetic Particle Simulation of Turbulent Transport in Burning Plasmas* Department of Energy OFES SciDAC Center, co-PI with Pat Diamond (UCSD), Zhihong Lin (UCI), Viktor Decyk (UCLA), ($60K to ISI), 3/2008-2/2011.

Intel07 *Compiler-Assisted Performance Tuning of Libraries and Application Kernels for Multi-Core Architectures* Gift from Intel Research Council($50K), 12/2007-11/2008.

NSF07 *Petascale Hierarchical Simulations of Biopolymer Translocation through Silicon Nitride and Silica Nanopores and Nanofluidic Channels* National Science Foundation OCI-0749360 ($1M, $160K to ISI), co-PI with Priya Vashishta (PI), Rajiv Kalia and Aiichiro Nakano, 9/2007-9/2012.

NNSA07 *Compiler Requirements for enhanced FBDIMM Buffer on Board based memory technology (eBOB)* Sandia National Laboratory ($40K), PI, July-Sept. 2007.

DOE06 *Performance Engineering Research Institute* Department of Energy SciDAC Institute DE-FC02-06ER25765 ($15M, $2.25M at USC), co-PI with Robert Lucas (PI), David Bailey and Kathy Yelick (Lawrence Berkeley Laboratory), Dan Quinlan (Lawrence Livermore National Laboratory), Jeff Vetter (Oak Ridge National Laboratory), Paul Hovland (Argonne National Laboratory), Allan Snavely (UCSD), John Mellor-Crummey (Rice), Dan Reed (RENCI), Jack Dongarra (UTK), Jeff Hollingsworth (UMD), 9/2006-9/2011.

NSF06 *CSR-AES: Intelligent Optimization of Parallel and Distributed Applications* NSF Computer Systems Research Contract CSR-0615412 ($885K, $642K at USC), PI with Aiichiro Nakano, Yolanda Gil, Kristina Lerman, Ewa Deelman and Joel Saltz (Ohio State), 8/06-8/09.

NSF05 *DDDAS-SMRP: A Dynamic Data-Driven Application System for Signal and Image Processing* NSF Dynamic Data-Driven Application System Contract DDDAS-0540407($356K at USC), co-PI with Pedro Diniz and Anna Gilbert (PI) (University of Michigan), 12/05-11/08.

NSF05 *CSR-AES: Intelligent Design and Optimization of Parallel and Distributed Applications* NSF Computer Systems Research Contract CSR-0509517 ($200K, $150K at USC), PI with Aiichiro Nakano, Yolanda Gil, Kristina Lerman, Ewa Deelman and Joel Saltz (Ohio State), 7/05-6/04.

NSF04 *Extending the Limits of Large-Scale Shared Memory Multiprocessors* NSF Software Technology for High End Computing Contract CCF-0444470 ($750K, $224K at USC), co-PI with PI Kunle Olokuton (Stanford University), 11/04-10/07.

DOE04 PetaScale Application Development Analysis," DOE Contract No. DE-FG02-04ER25630 ($150K at USC), PI, 9/04-8/07.

NSA04 *Processing-in-memory Technology for Knowledge Discovery Algorithms* AFRL and NSA Grant Number FA8750-04-1-0265 ($597K), PI with co-PI Hans Chalupsky, 8/04-11/05.

Intel04 *Automated Design Space Exploration for Compiling to Heterogeneous Systems-on-a-Chip* Gift from Intel Corporation ($300K),8/04-7/07.

DARPA04 *CEARCH: Cognition Enabled ARCHitecture* DARPA ACIP Contract, ($6M, $200K at ISI-West) Senior Personnel, in collaboration with Bob Parker and Steve Crago (ISI-East), Kunle Olokuton and Sebastian Thrul (Stanford University), Anant Agarwal (MIT), Gary Bradski (Intel), Michael Witbrock (Cycorp), Rick Pancoast (Lockheed Martin), Akey (Northrop Grumman), 7/04-6/06.

Intel03 *Mapping Pipelined Computations to Reconfigurable Architectures* Gift from Intel Corporation ($36K), 9/03-2/04.

DOE06 *Exploring Multiple Shared Memory Models with FLASH* DOE Grant DE-FG02-03ER25563 ($577K at USC), PI with co-PIs Robert Lucas and Mark Horowitz at Stanford University, September 2003 through August 2006.

NSF02 *NGS: Resource-Aware Off-line and On-line Empirical Optimization* NSF Next Generation Software Contract ACI-0204040 ($666K), PI with co-PIs Pedro Diniz and Robert Lucas, 11/02-10/05.

DARPA02 *Hewlett-Packard, DARPA Industrial Research and Development* DARPA High Productivity Computing Contract ($3M, $1.5M at USC), senior personnel on subcontract with Hewlett-Packard, 8/02-2/04.

NSF02 *Compiler-Driven Design Space Exploration For Multiple FPGA Systems* NSF Computer Systems Architecture Contract CCR-0209228 ($300K), co-PI with Pedro Diniz, 8/02-7/05.

NSF98 *Predicated Analysis for Cost-Effective Run-Time Parallelization* NSF New Technology Contract ACI-9721368 ($272K), 12/98-11/02.

DARPA98 *DIVA: A Data-Intensive Architecture* DARPA Data-Intensive Systems Contract F30602-98-2-0180 ($12M, $10M at USC) PI with John Granacki, 4/98-12/02.

DARPA98 *DEFACTO: A Design Environment For Adaptive Computing Technology* DARPA Adaptive Computing Systems Contract F30602-98-2-0113 ($4M), PI with John Granacki, 10/98-6/02.

DARPA95 *Effective Parallelization via Compile-Time, Run-Time and Interactive Techniques* DARPA Scalable Software Contract DABT63-95-C-0118 ($500K), 9/95-8/98.