Scott Pakin

Research interests

High-performance computer systems, with an emphasis on high-speed communication networks, messaging layers, parallel programming models, and quantum computing

Experience

Los Alamos National Laboratory

Christoph Junghans 11/2022–present
Patrick McCormick 4/2010–11/2022
Adolfy Hoisie 2/2002–4/2010

- Served as team leader for the Co-design team.
- Led the Flexible Computational Science Infrastructure (FleCSI) project. FleCSI is intended to be the middleware layer for LANL's next-generation mission-critical physics applications.
- Led the Unifying Circuit-Model Quantum Computing and Quantum Annealing project, which seeks to develop new, high-level programming models that increase programmer productivity by mapping easily to both of the two dominant quantum-computing models.
- Co-led the Monte Carlo Transport Simulations on a Billion-Core Approximate-Computing Platform project. The goal of this project is to gauge the feasibility and performance of running Monte Carlo simulations on pre-production hardware designed for extreme scale with low power and physical-footprint requirements.
- Led the Taming Defects in Quantum Computers project, a multidisciplinary effort (theoretical physics and computer science) to use machine-learning techniques to generate quantum algorithms that are most robust to noise in noisy intermediate-scale quantum (NISQ) computers—primarily gate-model—without requiring costly quantum error-correcting codes.
- Led the Hardware Evaluation element of the Exascale Computing Project. Hardware Evaluation's role was to analyze and predict the performance of scientific applications under development in the context of future supercomputers.
- Co-led the Hybrid Quantum-Classical Computing (HQCC) project. HQCC drew talent from theoretical physics, experimental physics, and computer science to analyze the behavior, performance characteristics, and application of annealing-based quantum computers.
- Created and led the development of Byfl, a tool that examines ways in which application performance can be characterized in a hardware-independent manner. Byfl implements a form of "software performance counters" that are analogous

(Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

to the hardware performance counters provided by all modern CPUs but return application-centric rather than hardware-centric information.

⋆ Led the Power Aware Data Center project, an attempt to quantify and exploit "trapped capacity"—the difference between allocated power and average power draw—in supercomputer data centers. This sort of study has never been done at a comparable scale (three of the world's ten fastest supercomputers), time frame (well over a year), and with normal, production workloads (as opposed to controlled studies). The Power Aware Data Center project spans multiple organizational units at Los Alamos National Laboratory; project members include computer scientists, statisticians, facilities engineers, and supercomputer system administrators, all working together to find ways to make more efficient use of a fixed data center power budget.

⋆ Created and co-led (and towards the end, solely led) the Supercomputer Stress Test project, whose goal was to produce a modular, extensible suite of stress tests and diagnostic routines for large-scale clusters and supercomputers that may contain unique hardware components. Salient characteristics of the Supercomputer Stress Test are that it can test hardware components together or in isolation; it can be extended easily as new hardware is introduced; and it can perform fault isolation, binary searching for a minimal set of nodes and tests that fail to produce correct output. Code development was performed by a team of students at Harvey Mudd College for their senior-year Computer Science Clinic project.

⋆ Designed and implemented the Cell Messaging Layer, the world's fastest MPI-based communication layer for clusters of Cell Broadband Engine (Cell/B.E.) processors. The Cell Messaging Layer embodies a novel programming approach called the *reverse-acceleration model* that facilitates the porting of applications from conventional supercomputing platforms to the more complex—but potentially higher-performance—Cell/B.E. platform. The Cell Messaging Layer has been demonstrated on the full Roadrunner system, the first supercomputer to achieve a sustained petaflop/s of performance.

⋆ Designed and implemented coNCePTuaL, a novel, domain-specific programming language and surrounding framework designed specifically for analyzing the performance of high-performance communication networks in a formal, scientific manner. coNCePTuaL was used to help analyze proposals for Los Alamos National Laboratory's Roadrunner and Cielo supercomputer procurements.

⋆ Designed and implemented JumboMem, the first fully user-level piece of software that gives unmodified applications transparent access to memory spread

 (Continued on next page)

Mail stop B287 · Los Alamos, NM 87545 · E-mail: pakin@lanl.gov
Telephone: (505) 667-5568
URL: https://www.lanl.gov/search-capabilities/profiles/scott-pakin.shtml

across multiple computers. When a program exceeds the memory in one computer, it automatically spills over into the memory of the next computer.

⋆ Helped develop a highly accurate, quasi-analytical performance model for Krak, a 600,000-line Lagrangian hydrodynamics application from the U.S. Department of Energy's Advanced Simulation and Computing program.

⋆ Assisted in the benchmarking and analysis of various production and preproduction large-scale parallel computer systems including the Roadrunner, Lobo, Lightning, and Q supercomputers at Los Alamos National Laboratory; the Blue Gene/P, Blue Gene/L, and Purple supercomputers at Lawrence Livermore National Laboratory and the IBM T. J. Watson Research Center; the Red Storm and Thunderbird supercomputers at Sandia National Laboratory; and the Columbia supercomputer at NASA.

⋆ Supervised a number of graduate and undergraduate students on projects related to the preceding topics.

National Center for Supercomputing Applications — **Rob Pennington** — **8/2001–2/2002**

⋆ Collaborated on a project to develop NCSA's next-generation communication middleware, the Virtual Machine Interface 2.0. The intention was for VMI 2.0, when completed, to run on all of the production clusters at the National Center for Supercomputing Applications.

University of Illinois at Urbana-Champaign — **Andrew Chien** — **1/1995–8/2001**

⋆ Led development team producing Illinois Fast Messages (FM), a fast, efficient messaging layer that provides the communication infrastructure for the National Center for Supercomputing Applications's NT Supercluster, at the time one of the world's 500 fastest computing platforms. FM was also one of four academic research projects influencing Compaq/Intel/Microsoft's Virtual Interface Architecture (VIA) standard for cluster interconnects, the predecessor to InfiniBand.

⋆ Cooperated with researchers at MIT on implementing a novel coordinated process scheduling algorithm that improves the throughput and response time of parallel and distributed applications running on clusters of timeshared workstations.

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

- ⋆ Designed and implemented a new technique for improving the performance of collective-communication-intensive parallel applications running on timeshared PC clusters.

- ⋆ Supervised various undergraduate research projects.

Intel Corporation **Tom Shott** **5/1995–8/1995**

- ⋆ Helped architect and validate a commercial, programmable, cache-coherent, scalable distributed shared memory system based on Intel Pentium Pro processors.

- ⋆ Defined and implemented a tool to ensure the behavioral, functional, and register-transfer level simulators correctly modeled the hardware, software, and firmware used to implement the system's cache-coherency protocol.

- ⋆ Assisted with the development and refinement of the protocol.

Education

Ph.D. University of Illinois at Urbana-Champaign October 2001
Comp. Sci. *Thesis title:* "Unresponsiveness-tolerant collective communication"

M.S. University of Illinois at Urbana-Champaign January 1995
Comp. Sci. *Thesis title:* "The impact of message traffic on multicomputer memory hierarchy performance"

B.S. Carnegie Mellon University May 1992
Math/CS *Thesis title:* "An efficient implementation of Linda on the iWarp"

Publications

Book Chapters

Song Huang, Song Fu, **Scott Pakin**, and Michael Lang. Characterizing power and energy efficiency of Legion data-centric runtime and applications on heterogeneous high-performance computing systems. In Satyadhyan Chickerur, editor, *High Performance Parallel Computing*, pages 1–20. IntechOpen, December 10, 2018.

Scott Pakin. Myrinet. In David Padua, editor, *Encyclopedia of Parallel Computing*, pages 1239–1247. Springer, September 8, 2011.

Scott Pakin. High performance interconnects for massively parallel systems. In Ada Gavrilovska, editor, *Attaining High Performance Communications: A Vertical Approach*, chapter 1, pages 1–23. Chapman and Hall/CRC Press, September 2009.

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Journals and Magazines

J. Abhijith, Adetokunbo Adedoyin, John Ambrosiano, Petr Anisimov, William Casper, Gopinath Chennupati, Carleton Coffrin, Hristo Djidjev, David Gunter, Satish Karra, Nathan Lemons, Shizeng Lin, Alexander Malyzhenkov, David Mascarenas, Susan Mniszewski, Balu Nadiga, Daniel O'Malley, Diane Oyen, **Scott Pakin**, Lakshman Prasad, Randy Roberts, Phillip Romero, Nandakishore Santhi, Nikolai Sinitsyn, Pieter J. Swart, James G. Wendelberger, Boram Yoon, Richard Zamora, Wei Zhu, Stephan Eidenbenz, Andreas Bärtschi, Patrick J. Coles, Marc Vuffray, and Andrey Y. Lokhov. Quantum algorithm implementations for beginners. *ACM Transactions on Quantum Computing*, 3(4), July 2022.

Md Atiqul Mollah, Wenqi Wang, Peyman Faizian, Md Shafayat Rahman, Xin Yuan, **Scott Pakin**, and Michael Lang. Modeling universal globally adaptive load-balanced routing. *ACM Transactions on Parallel Computing*, 6(2):9:1–9:23, September 2019.

Scott Pakin and Steven P. Reinhardt. Programming a D-Wave annealing-based quantum computer: Tools and techniques. *Quantum Information and Computation*, 19(9–10):0721–0759, August 2019.

Peyman Faizian, Juan Francisco Alfaro, Md Shafayat Rahman, Md Atiqul Mollah, Xin Yuan, **Scott Pakin**, and Michael Lang. TPR: Traffic pattern-based adaptive routing for dragonfly networks. *IEEE Transactions on Multi-Scale Computing Systems*, 4(4):931–943, October 22, 2018.

Scott Pakin. Performing fully parallel constraint logic programming on a quantum annealer. *Theory and Practice of Logic Programming*, 18(5–6):1–22, September 2018.

Zhou Tong, **Scott Pakin**, Michael Lang, and Xin Yuan. Fast classification of MPI applications using Lamport's logical clocks. *Journal of Parallel and Distributed Computing*, 120:77–88, October 2018.

Peyman Faizian, Md Atiqul Mollah, Xin Yuan, Zaid Alzaid, **Scott Pakin**, and Michael Lang. Random regular graph and generalized De Bruijn graph with k-shortest path routing. *IEEE Transactions on Parallel and Distributed Systems*, 29(1):144–155, January 2018.

Md Atiqul Mollah, Xin Yuan, **Scott Pakin**, and Michael Lang. Rapid calculation of max-min fair rates for multi-commodity flows in fat-tree networks. *IEEE Transactions on Parallel and Distributed Systems*, 29(1):156–168, January 2018.

Ziming Zhang, Michael Lang, **Scott Pakin**, and Song Fu. TracSim: Simulating and scheduling trapped power capacity to maximize machine room throughput. *Parallel Computing*, 57:108–124, September 2016.

(Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Scott Pakin, Curtis Storlie, Michael Lang, Robert E. Fields III, Eloy E. Romero, Craig Idler, Sarah Michalak, Hugh Greenberg, Josip Loncaric, Randal Rheinheimer, Gary Grider, and Joanne Wendelberger. Power usage of production supercomputers and production workloads. *Concurrency and Computation: Practice and Experience*, 28(2):274–290, February 2016.

Xin Yuan, Santosh Mahapatra, Michael Lang, and **Scott Pakin**. Static load-balanced routing for slimmed fat-trees. *Journal of Parallel and Distributed Computing*, 74(5):2423–2432, May 2014.

Scott Pakin, Xin Yuan, and Michael Lang. Predicting the performance of extreme-scale supercomputer networks. *The Next Wave*, 20(2):7–19, November 2013.

Scott Pakin. Ten ways to fool the masses when giving performance results on GPUs. *HPCwire*, December 13, 2011.

Darren J. Kerbyson, Michael Lang, and **Scott Pakin**. Adapting wave-front algorithms to efficiently utilize systems with deep communication hierarchies. *Parallel Computing*, 37(9):550–561, September 2011.

Hikmet Dursun, Kevin J. Barker, Darren J. Kerbyson, **Scott Pakin**, Richard Seymour, Rajiv K. Kalia, Aiichiro Nakano, and Priya Vashishta. An MPI performance monitoring interface for Cell based compute nodes. *Parallel Processing Letters*, 19(4), December 2009.

Kevin J. Barker, Kei Davis, Adolfy Hoisie, Darren J. Kerbyson, Michael Lang, **Scott Pakin**, and José Carlos Sancho. Using performance modeling to design large-scale systems. *IEEE Computer*, 42(11):42–49, November 2009.

Scott Pakin, Michael Lang, and Darren J. Kerbyson. The reverse-acceleration model for programming petascale hybrid systems. *IBM Journal of Research and Development*, 53(5), September–October 2009.

Kevin Barker, Kei Davis, Adolfy Hoisie, Darren J. Kerbyson, Mike Lang, **Scott Pakin**, and José C. Sancho. A performance evaluation of the Nehalem quad-core processor for scientific computing. *Parallel Processing Letters*, 18(4):453–469, December 2008.

Scott Pakin. Good things come in little packages: An introduction to writing `.ins` and `.dtx` files. *TUGboat*, 29(2):305–314, 2008.

Scott Pakin. The design and implementation of a domain-specific language for network performance testing. *IEEE Transactions on Parallel and Distributed Systems*, 18(10):1436–1449, October 2007.

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Ram Srinivasan, Eitan Frachtenberg, Olaf Lubeck, **Scott Pakin**, and Jeanine Cook. An idealistic neuro-PPM branch predictor. *The Journal of Instruction-Level Parallelism*, 9, May 2007.

Eitan Frachtenberg, Fabrizio Petrini, Juan Fernandez, and **Scott Pakin**. STORM: Scalable resource management for large-scale parallel computers. *IEEE Transactions on Computers*, 55(12):1572–1587, December 2006.

Leon Arber and **Scott Pakin**. The impact of message-buffer alignment on communication performance. *Parallel Processing Letters*, 15(1):49–65, March 2005. ISSN 0129-6264.

Darren J. Kerbyson, Adolfy Hoisie, **Scott Pakin**, Fabrizio Petrini, and Harvey J. Wasserman. A performance evaluation of an Alpha EV7 processing node. *The International Journal of High Performance Computing Applications*, 18(2):199–209, May 1, 2004.

Scott Pakin. PerlTEX: Defining LATEX macros using Perl. *TUGboat*, 25(2):150–159, 2004. ISSN 0896-3207.

Andrew A. Chien, Mario Lauria, Rob Pennington, Mike Showerman, Giulio Iannello, Matt Buchanan, Kay Connelly, Louis Giannini, Greg Koenig, Sudha Krishnamurthy, Qian Liu, **Scott Pakin**, and Geetanjali Sampemane. Design and evaluation of an HPVM-based Windows NT supercomputer. *International Journal of High Performance Computing Applications*, 13(3):201–219, Fall 1999. Special issue on clusters and computational grids for scientific computing.

Mario Lauria, **Scott Pakin**, and Andrew Chien. Efficient layering for high speed communication: The MPI over Fast Messages (FM) experience. *Cluster Computing*, 2(2):107–116, September 1999.

Scott Pakin, Vijay Karamcheti, and Andrew A. Chien. Fast Messages: Efficient, portable communication for workstation clusters and MPPs. *IEEE Concurrency*, 5(2):60–73, April–June 1997.

Scott Pakin. Regular expressions and gender guessing. *Computer Language Magazine*, 8(12):59–68, December 1991.

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Conferences and Workshops (peer-reviewed)

Li Tang, Phil Jones, and **Scott Pakin**. Harnessing extreme heterogeneity for ocean modeling with tensors. In *The 2nd International Workshop on Extreme Heterogeneity Solutions (ExHET 2023)*, Montreal, Canada, February 25, 2023. Held in conjunction with Principles and Practice of Parallel Programming (PPoPP 2023). **Won Best Paper Award**.

Ellis Wilson, Frank Mueller, and **Scott Pakin**. Combining hard and soft constraints in quantum constraint-satisfaction systems. In *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, pages 161–174, Dallas, Texas, USA, November 13-18, 2022. IEEE Computer Society.

Li Tang and **Scott Pakin**. Cross-level characterization of program execution. In *30th International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS 2022)*, Nice, France, October 18–20, 2022.

Ellis Wilson, Frank Mueller, and **Scott Pakin**. Mapping constraint problems onto quantum gate and annealing devices. In *Second International Workshop on Quantum Computing Software*, St. Louis, Missouri, USA, November 15, 2021. IEEE/ACM.

Scott Pakin. A simple heuristic for expressing a truth table as a quadratic pseudo-Boolean function. In *IEEE International Conference on Quantum Computing and Engineering (QCE21)*. IEEE, October 17–22, 2021.

Mohamed W. Hassan, **Scott Pakin**, and Wu-chun Feng. C to D-Wave: A high-level C compilation framework for quantum annealers. In *Proceedings of the Twenty-third Annual IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, Massachusetts, USA, September 24–26, 2019. **Won Innovative Paper Award**.

Harsh Khetawat, Ashlesha Atrey, George Li, Frank Mueller, and **Scott Pakin**. Implementing NChooseK on IBM Q quantum computer systems. In Michael Kirkedal Thomsen and Mathias Soeken, editors, *Proceedings of the 11th International Conference on Reversible Computation*, volume 11497 of *Lecture Notes in Computer Science*, pages 209–223, Lausanne, Switzerland, June 24–25, 2019. Springer.

Scott Pakin. Targeting classical code to a quantum annealer. In *Proceedings of the 24th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, Providence, Rhode Island, USA, April 14–17, 2019. ACM.

Tomáš Vyskočil, **Scott Pakin**, and Hristo N. Djidjev. Embedding inequality constraints for quantum annealing optimization. In Sebastian Feld and Claudia Linnhoff-Popien, editors, *Proceedings of the First International Workshop on Quantum Technol-*

(Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

ogy and Optimization Problems (QTOP), volume 11413 of *Lecture Notes in Computer Science*, pages 11–22, Garching bei München, Germany, March 18–21, 2019. Springer.

Scott Pakin and Steven P. Reinhardt. A survey of programming tools for D-Wave quantum-annealing processors. In Rio Yokota, Michèle Weiland, David Keyes, and Carsten Trinitis, editors, *Proceedings of the 2018 International Conference on High Performance Computing (ISC-HPC)*, volume 10876 of *Lecture Notes in Computer Science*, pages 103–122, Frankfurt, Germany, June 24–28, 2018. Springer.

Zhou Tong, Xin Yuan, **Scott Pakin**, and Michael Lang. Performance and accuracy trade-offs of HPC application modeling and simulation. In *Proceedings of the 32nd IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 774–783, Vancouver, British Columbia, Canada, May 21–25, 2018.

Md Atiqul Mollah, Peyman Faizian, Md Shafayat Rahman, Xin Yuan, **Scott Pakin**, and Mike Lang. A comparative study of topology design approaches for HPC interconnects. In *Proceedings of the 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, pages 392–401, Washington, DC, USA, May 1–4, 2018.

Md Atiqul Mollah, Peyman Faizian, Md Shafayat Rahman, Xin Yuan, **Scott Pakin**, and Michael Lang. Modeling UGAL on the dragonfly topology. In Stephen Jarvis, Steven Wright, and Simon Hammond, editors, *Proceedings of the 8th International Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS)*, volume 10724 of *Lecture Notes in Computer Science*, pages 136–157, Denver, Colorado, USA, November 13, 2017. Springer.

Peyman Faizian, Md Atiqul Mollah, Md Shafayat Rahman, Xin Yuan, **Scott Pakin**, and Mike Lang. Throughput models of interconnection networks: The good, the bad, and the ugly. In *Proceedings of the 25th IEEE Annual Symposium on High-Performance Interconnects*, pages 33–40, Santa Clara, California, USA, August 28–30, 2017. IEEE.

Scott Pakin. Navigating a maze using a quantum annealer. In *Proceedings of the 2nd International Workshop on Post Moore's Era Supercomputing (PMES)*, pages 30–36, Denver, Colorado, USA, November 13, 2017. ACM.

Vignesh Adhinarayanan, Wu-chun Feng, David Rogers, James Ahrens, and **Scott Pakin**. Characterizing and modeling power and energy for extreme-scale in-situ visualization. In *Proceedings of the 31st IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 978–987, Orlando, Florida, USA, May 29–June 2, 2017. IEEE.

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Song Huang, Song Fu, **Scott Pakin**, and Michael Lang. Characterizing power and energy efficiency of Legion runtime and applications: An early experience. In *Proceedings of the Seventh International Green and Sustainable Computing Conference (IGSC)*, pages 1–8, Hangzhou, China, November 7–9, 2016. IEEE.

Scott Pakin. A quantum macro assembler. In *Proceedings of the 2016 IEEE High Performance Extreme Computing Conference (HPEC)*, Waltham, Massachusetts, USA, September 13–15, 2016.

Peyman Faizian, Md Shafayat Rahman, Md Atiqul Mollah, Xin Yuan, **Scott Pakin**, and Mike Lang. Traffic pattern-based adaptive routing for intra-group communication in dragonfly networks. In *Proceedings of the 24th IEEE Annual Symposium on High-Performance Interconnects*, pages 19–26, Santa Clara, California, USA, August 24–26, 2016. IEEE.

Zhou Tong, **Scott Pakin**, Michael Lang, and Xin Yuan. Fast classification of MPI applications using Lamport's logical clocks. In *Proceedings of the 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 618–627. IEEE, May 23–27, 2016.

Peyman Faizian, Md Atiqul Mollah, Xin Yuan, **Scott Pakin**, and Michael Lang. Random regular graph and generalized De Bruijn graph with k-shortest path routing. In *Proceedings of the 30th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 103–112. IEEE, May 23–27, 2016.

James H. Laros III, Kevin Pedretti, Stephen Olivier, Ryan E. Grant, Stephen L. Olivier, Michael Levenhagen, David DeBonis, **Scott Pakin**, Steven Martin, Matthew Kappel, and Paul Falde. ACES and Cray collaborate on advanced power managment for Trinity. In *Proceedings of the 2016 Cray User's Group Conference*, London, England, May 2016.

Song Huang, Michael Lang, **Scott Pakin**, and Song Fu. Measurement and characterization of Haswell power and energy consumption. In *Proceedings of the 3rd International Workshop on Energy Efficient Supercomputing (E2SC 2015)*, pages 7:1–7:10. ACM, 2015.

Md Atiqul Mollah, Xin Yuan, **Scott Pakin**, and Michael Lang. Fast calculation of max-min fair rates for multi-commodity flows in fat-tree networks. In *2015 IEEE International Conference on Cluster Computing (Cluster 2015)*, pages 351–360, September 2015.

Scott Pakin. Parallel post-processing with MPI-Bash. In *Proceedings of the First International Workshop on HPC User Support Tools (HUST-14)*, pages 1–12, New

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Orleans, Louisiana, November 21, 2014.

Eric Anger, Sudhakar Yalamanchili, **Scott Pakin**, and Patrick McCormick. Architecture-independent modeling of intra-node data movement. In *Proceedings of the LLVM Compiler Infrastructure in HPC Workshop (LLVM-HPC)*, pages 29–39, New Orleans, Louisiana, November 17, 2014.

Ziming Zhang, Michael Lang, **Scott Pakin**, and Song Fu. Trapped capacity: Scheduling under a power cap to maximize machine-room throughput. In *Proceedings of the 2nd International Workshop on Energy Efficient Supercomputing (E2SC'14)*, pages 41–50, New Orleans, Louisiana, November 16, 2014.

Joshua Payne, Dana Knoll, Allen McPherson, William Taitano, Luis Chacon, Guangye Chen, and **Scott Pakin**. Computational co-design of a multiscale plasma application: A process and initial results. In *Proceedings of the 28th International Parallel and Distributed Processing Symposium*, Phoenix, Arizona, May 19–23, 2014.

Xin Yuan, Santosh Mahapatra, Michael Lang, and **Scott Pakin**. LFTI: A new performance metric for assessing interconnect designs for extreme-scale HPC systems. In *Proceedings of the 28th International Parallel and Distributed Processing Symposium*, Phoenix, Arizona, May 19–23, 2014.

Mark Gamell, Ivan Rodero, Manish Parashar, Janine C. Bennett, Hemanth Kolla, Jacqueline Chen, Peer-Timo Bremer, Aaditya G. Landge, Attila Gyulassy, Patrick McCormick, **Scott Pakin**, Valerio Pascucci, and Scott Klasky. Exploring power behaviors and trade-offs of in-situ data analytics. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'13)*, Denver, Colorado, November 17–22, 2013.

Xin Yuan, Santosh Mahapatra, Wickus Nienaber, **Scott Pakin**, and Michael Lang. A new routing scheme for jellyfish and its performance with HPC workloads. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC'13)*, Denver, Colorado, November 17–22, 2013.

Scott Pakin and Patrick McCormick. Hardware-independent application characterization. In *Proceedings of the 2013 IEEE International Symposium on Workload Characterization (IISWC 2013)*, Portland, Oregon, September 22–24, 2013.

Scott Pakin and Michael Lang. Energy modeling of supercomputers and large-scale scientific applications. In *Proceedings of the 4th International Green Computing Conference (IGCC 2013), Work in Progress Workshop*, Arlington, Virginia, June 27–29, 2013.

(Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Scott Pakin and Michael Lang. Understanding the performance of two production supercomputers. In *Proceedings of the 27th IEEE International Symposium on Parallel and Distributed Processing (IPDPS 2013), 14th IEEE International Workshop on Parallel and Distributed Scientific and Engineering Computing (PDSEC)*, Boston, Massachusetts, May 24, 2013.

Xin Yuan, Santosh Mahapatra, Michael Lang, and **Scott Pakin**. RRR: A load balanced routing scheme for slimmed fat-trees. In *Proceedings of the 27th IEEE International Symposium on Parallel and Distributed Processing (IPDPS 2013), 3rd Workshop on Communication Architecture for Scalable Systems (CASS)*, Boston, Massachusetts, May 20, 2013.

Xing Wu, Frank Mueller, and **Scott Pakin**. Automatic generation of executable communication specifications from parallel applications. In *Proceedings of the 25th International Conference on Supercomputing (ICS 2011)*, Tucson, Arizona, May 31–June 4, 2011. ACM. **Nominated for Best Paper Award**.

Kevin Barker, Kei Davis, Adolfy Hoisie, Darren J. Kerbyson, Mike Lang, **Scott Pakin**, and José C. Sancho. Entering the petaflop era: The architecture and performance of Roadrunner. In *Proceedings of the 2008 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'08)*, November 15–21, 2008.

Kevin Barker, Kei Davis, Darren J. Kerbyson, Mike Lang, **Scott Pakin**, and José C. Sancho. An early performance evaluation of the SiCortex SC648. In *Proceedings of the IEEE International Symposium on Performance Analysis of Systems and Software (ISPASS), Workshop on Unique Chips and Systems (UCAS4)*, Austin, Texas, April 20, 2008.

Scott Pakin. Receiver-initiated message passing over RDMA networks. In *Proceedings of the 22nd IEEE International Parallel and Distributed Processing Symposium (IPDPS 2008)*, Miami, Florida, April 14–18, 2008.

Scott Pakin and Greg Johnson. Performance analysis of a user-level memory server. In *Proceedings of the IEEE International Conference on Cluster Computing (Cluster 2007)*, pages 249–258, Austin, Texas, September 17–20, 2007.

Kevin J. Barker, **Scott Pakin**, and Darren J. Kerbyson. A performance model of the Krak hydrodynamics application. In *Proceedings of the International Conference on Parallel Processing (ICPP 2006)*, Columbus, Ohio, August 14–18, 2006.

Adolfy Hoisie, Gregory Johnson, Darren J. Kerbyson, Mike Lang, and **Scott Pakin**. A performance comparison through benchmarking and modeling of three leading

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

supercomputers: Blue Gene/L, Red Storm, and Purple. In *Proceedings of the 2006 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'06)*, Tampa, Florida, November 11–17, 2006.

Scott Pakin. Rapid development of application-specific network performance tests. In *Proceedings of the International Conference on Computational Science (ICCS 2005), Workshop on Tools for Program Development and Analysis in Computational Science*, Atlanta, Georgia, May 22–25, 2005.

Scott Pakin. Reproducible network benchmarks with coNCePTuaL. In *Proceedings of Euro-Par 2004*, number 3149 in Lecture Notes in Computer Science, pages 64–71, Pisa, Italy, August 31–September 3, 2004.

Scott Pakin. coNCePTuaL: A network correctness and performance testing language. In *Proceedings of the International Parallel and Distributed Processing Symposium (IPDPS 2004)*, Santa Fe, New Mexico, April 28–30, 2004.

Fabrizio Petrini, Darren J. Kerbyson, and **Scott Pakin**. The case of the missing supercomputer performance: Achieving optimal performance on the 8,192 processors of ASCI Q. In *Proceedings of the 2003 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'03)*, Phoenix, Arizona, November 15–21, 2003. **Won Best Paper Award**.

Darren Kerbyson, Adolfy Hoisie, **Scott Pakin**, Fabrizio Petrini, and Harvey Wasserman. Performance testing of an EV7 AlphaServer machine. In *Proceedings of the Los Alamos Computer Science Institute (LACSI) Symposium*, Santa Fe, New Mexico, October 13–16, 2002.

Eitan Frachtenberg, Fabrizio Petrini, Juan Fernandez, **Scott Pakin**, and Salvador Coll. STORM: Lightning-fast resource management. In *Proceedings of the 2002 International Conference for High Performance Computing, Networking, Storage, and Analysis (SC'02)*, Baltimore, Maryland, November 16–22, 2002.

Scott Pakin and Avneesh Pant. VMI 2.0: A dynamically reconfigurable messaging layer for availability, usability, and management. In *The 8th International Symposium on High Performance Computer Architecture (HPCA-8), Workshop on Novel Uses of System Area Networks (SAN-1)*, Cambridge, Massachusetts, February 2, 2002.

Geetanjali Sampemane, **Scott Pakin**, and Andrew A. Chien. Performance monitoring on an HPVM cluster. In *Proceedings of the International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA 2000), International Workshop on Cluster Computing—Technologies, Environments, and Applications (CC-TEA 2000)*, Las Vegas, Nevada, June 26–29, 2000.

 (Continued on next page)

Mail stop B287 · Los Alamos, NM 87545 · E-mail: pakin@lanl.gov
Telephone: (505) 667-5568
URL: https://www.lanl.gov/search-capabilities/profiles/scott-pakin.shtml

Patrick Sobalvarro, **Scott Pakin**, Andrew Chien, and William Weihl. Dynamic coscheduling on workstation clusters. In *12th Annual International Parallel Processing Symposium & 9th Symposium on Parallel and Distributed Processing (IPPS/SPDP), 4th Workshop on Job Scheduling Strategies for Parallel Processing*, Orlando, Florida, March 1998. Published in Lecture Notes in Computer Science, vol. 1459, pp. 231–256. Springer-Verlag. ISBN 3-540-64825-9.

Mario Lauria, **Scott Pakin**, and Andrew A. Chien. Efficient layering for high speed communication: Fast Messages 2.x. In *Proceedings of the Seventh IEEE International Symposium on High Performance Distributed Computing (HPDC-7)*, pages 10–20, Chicago, Illinois, July 28–31, 1998.

Andrew Chien, **Scott Pakin**, Mario Lauria, Matt Buchanan, Kay Hane, Louis Giannini, and Jane Prusakova. High performance virtual machines (HPVM): Clusters with supercomputing APIs and performance. In Michael Heath, Virginia Torczon, Greg Astfalk, Petter E. Bjørstad, Alan H. Karp, Charles H. Koebel, Vipin Kumar, Robert F. Lucas, Layne T. Watson, and David E. Womble, editors, *Proceedings of the Eighth SIAM Conference on Parallel Processing for Scientific Computing*, Minneapolis, Minnesota, March 1997.

Scott Pakin, Mario Lauria, and Andrew Chien. High performance messaging on workstations: Illinois Fast Messages (FM) for Myrinet. In *Proceedings of the 1995 ACM/IEEE Supercomputing Conference*, volume 2, pages 1528–1557, San Diego, California, December 4–8, 1995.

Awards and Honors

Best Paper Award at EXHET 2023 for *Harnessing Extreme Heterogeneity for Ocean Modeling with Tensors*, 2023.

Los Alamos Spot Award for "Highly positive influence on co-workers or team members or extraordinary team support: Outstanding support of CCS-7 and CCS as DC", 2021.

Los Alamos Spot Award for "Courtesy or responsiveness in dealing with customers or colleagues: For his DC/DUSA presentation to the group", 2021.

Innovative Paper Award at HPEC 2019 for *C to D-Wave: A High-level C Compilation Framework for Quantum Annealers*.

Los Alamos Spot Award for "a notable accomplishment in the area of safety, security, environment, or ethics", 2019.

Los Alamos Awards Program award "in recognition of outstanding contributions to the Logical Time Trace Tool", 2014.

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Best Paper award nomination at ICS 2011 for *Automatic Generation of Communication Specifications from Parallel Applications* (3 papers nominated out of 161 submissions/35 accepted papers).

Los Alamos Awards Program achievement award "for ASC CSSE L2 milestone 'Application Enablement on Next Generation Platforms,'" February 2011.

2010 Distinguished Performance Award "in appreciation for your contribution on the Cielo Integration Team."

2010 Defense Programs Award of Excellence "for significant contributions to the stockpile stewardship program" and for "successful deployment of the Cielo petascale computer."

Los Alamos Awards Program achievement award "for your contributions and teamwork in the Roadrunner Advanced Algorithms Team," August 2008.

2007 Distinguished Performance Award "in appreciation for your contribution on the Roadrunner system integration and technical assessment team."

2007 Defense Programs Award of Excellence "for outstanding achievements in high performance computing at LANL in 2007 in support of the nuclear weapons program."

Third place in the Journal of Instruction-Level Parallelism's 2nd Championship Branch Prediction Competition (CBP-2) in the "idealistic" track, for *An Idealistic Neuro-PPM Branch Predictor*, December 2006.

Best Paper award at SC2003 for *The Case of the Missing Supercomputer Performance: Achieving Optimal Performance on the 8,192 Processors of ASCI Q* (1 winner out of 207 submissions/60 accepted papers).

2003 Defense Programs Award of Excellence "for developing new system and environment for large-scale production Linux clusters."

2002 Distinguished Performance Award "in appreciation for your contribution on the Advanced Simulation and Computing (ASCI) Q Team."

2002 Defense Programs Award of Excellence "for the completion of the December 2002 LANL ASCI Applications Milestone thru deployment and integration of the ASCI Q Machine."

Named one of the top five new hires of the 2001–2002 fiscal year by a Los

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Alamos National Laboratory Director's committee.

Honor Society of Phi Kappa Phi for superior academic performance. Inducted 1999.

Intel Foundation Graduate Fellowship for doctoral students at selected universities who are judged to be the "best of the best" in fields of study related to Intel's technology research. September 1998–May 1999.

W. J. Poppelbaum Award for the CS department's most outstanding graduate student in the areas of hardware and architecture. March 1997.

Graduate College Fellowship tuition and fee waiver for the best three CS graduate students not yet on an assistantship. September 1993–May 1994.

Professional Service

Conferences

Awards chair, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2024

General co-chair, 2023 IEEE International Conference on Cluster Computing (Cluster 2023)

Panels co-chair, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2023

Program committee member, ISC High Performance, "Performance Modeling, Evaluation, and Analysis" track, 2022–2023

Steering Committee member, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2021–2024

Steering Committee member, IEEE International Conference on Cluster Computing (Cluster), 2021–2025

Student Networking chair, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2021

Track chair, the first IEEE International Conference on Quantum Computing and Engineering (QCE), "Quantum Algorithms & Information" track, 2020

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Track vice chair, the 34th IEEE International Parallel and Distributed Processing Symposium (IPDPS), "Experiments" track, 2020

Poster committee member, the International Conference on Cluster, Cloud and Grid Computing (CCGrid), 2010

Technical Papers co-chair, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2019

Program committee member, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), "Architecture and Networks" (formerly "Networks") track, 2005, 2007–2009, 2011, 2016–2017

Workshops committee member, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2014, 2015, and 2017

Area co-chair, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), "Architecture and Networks" track, 2013

Poster committee member, the International Conference for High Performance Computing, Networking, Storage, and Analysis (SC), 2012–2013

Committee member, SC13 Committee for the Selection of the Best Paper (BP) and Best Student Paper (BSP) Awards, 2013

Program committee member, the International Conference on Parallel Processing, "Performance" track, 2016; "Performance Modeling, Analysis, and Evaluation" track, 2015; "Networking and Communications" track, 2012

Program committee member, the IEEE International Conference on Parallel and Distributed Systems (ICPADS), "Multicore Computing and Parallel/Distributed Architecture" track, 2014–2015; no track, 2006

External review committee member, the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), 2012–2014

Program committee member, the IEEE International Conference on High Performance Computing (HiPC), "System Software" track, 2014; "Architecture" track, 2009

Program committee member, the IEEE International Conference on Cluster Computing (Cluster), "Cluster Design, Configuration and Administration" track, 2014; and "Tools, System Software, and Middleware" track, 2012

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Program committee member, Euro-Par, "Performance Prediction and Evaluation" track, 2013

Program committee member, the IEEE International Conference on Computer Communication Networks (ICCCN), 2012

Program committee member, the IEEE International Parallel and Distributed Processing Symposium (IPDPS), "Architecture" track, 2011

Program committee member, the International Conference on Computer Communication Networks (ICCCN), "Network Architecture and P2P protocols" track, 2011

Program committee member, the International Conference on Cluster, Cloud and Grid Computing (CCGrid), 2011

Program committee member, the IEEE International Symposium on Parallel and Distributed Processing with Applications (ISPA), "Middleware" track, 2011; "Architecture" track, 2008

Publicity co-chair, the IEEE International Conference on Cluster Computing (Cluster), 2010

Program committee member, the 1st Symposium on Application Accelerators in High Performance Computing (SAAHPC), 2009

Program committee member, the IEEE International Symposium on High Performance Distributed Computing (HPDC), 2003

Journals

Associate editor, IEEE Transactions on Parallel and Distributed Systems (TPDS), 2020–2022

Co-editor, IEEE Transactions on Parallel and Distributed Systems (TPDS) Special Section on Non-von Neumann Technologies, 2021

Guest editor, Elsevier's Journal of Parallel and Distributed Computing (JPDC), special issue on communication architecture for scalable systems, 2012

Reviewer, ACM Transactions on Quantum Computing, Advanced Quantum Technology (Wiley), Communications of the ACM, IEEE Computer, IEEE Transactions on Computers, IEEE Transactions on Parallel and Distributed Systems, IEEE Transactions on Quantum Engineering, Information and Computation (Elsevier), Journal of Parallel and Distributed Computing (Elsevier), and Parallel Computing: Systems & Applications (Elsevier)

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Workshops

Program committee, International Workshop on Quantum Computing Software (QCS), 2020–2023

Organizing committee, Workshop on Modeling & Simulation of Systems and Applications (ModSim), 2018, 2020

Program committee member, International Workshop on Post Moore's Era Supercomputing (PMES), 2018

Program committee member, the International Workshop on Parallel Programming Models and System Software (P2S2), 2008–2021, except 2011

Program committee member, the workshop on High-Level Parallel Programming Models and Supportive Environments (HIPS), 2016

Program committee member, the International Workshop on Energy Efficiency for Supercomputing (E2SC), 2015

Program committee member, the International Workshop on Power-aware Algorithms, Systems and Architectures (PASA), 2014

Program committee member, the First International Workshop on Energy Efficiency for Supercomputing (E2SC), 2013

Co-chair, Communication Architecture for Scalable Systems (CASS) workshop, 2011–2013

Program committee member, the Computer Architecture and Operating System co-design (CAOS) workshop, 2011

Co-chair, Communication Architecture for Clusters (CAC) workshop, 2004–2010

Program committee member, the International Workshop on High-Performance Interconnects for Distributed Computing (HPI-DC), 2005, 2009

Organizing committee member, the Institute for Advanced Architectures and Algorithms (IAA) Interconnection Network Workshop, 2008

Program committee member, Communication Architecture for Clusters (CAC) workshop, 2002–2003

 (Continued on next page)

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML

Minisymposia Organizer, "Guidance from Early Applications of Quantum Computers" minisymposium at the 2019 SIAM Annual Meeting; "Lessons from Early Applications Success on Quantum Computers" at the 2018 SIAM Annual Meeting; and "Identifying Computational Methods for Early Benefit from Quantum Computing" at the 2017 SIAM Annual Meeting

Other Judge, ACM Student Research Competition (at SC13), 2013

MAIL STOP B287 · LOS ALAMOS, NM 87545 · E-MAIL: PAKIN@LANL.GOV
TELEPHONE: (505) 667-5568
URL: HTTPS://WWW.LANL.GOV/SEARCH-CAPABILITIES/PROFILES/SCOTT-PAKIN.SHTML