
MARK D. HILL

Partner Hardware Architect, Microsoft Corp.

Gene M. Amdahl and John P. Morgridge Professor Emeritus, University of Wisconsin-Madison

April 2024

Home:

1170 Cloud Ave.

Menlo Park CA, 94025

markhill@cs.wisc.edu

<http://www.cs.wisc.edu/~markhill>

Short Biography

Mark D. Hill is the Gene M. Amdahl and John P. Morgridge Professor Emeritus of Computer Sciences at the University of Wisconsin-Madison (<http://www.cs.wisc.edu/~markhill>), following his 1988-2020 service in Computer Sciences and Electrical and Computer Engineering. His research interests include parallel-computer system design, memory system design, and computer simulation. Hill's work is highly collaborative with over 170 co-authors. He received the 2019 Eckert-Mauchly Award and is a fellow of AAAS, ACM, and IEEE. He served on the Computing Community Consortium (CCC) 2013-21 including as CCC Chair 2018-20, Computing Research Association (CRA) Board of Directors 2018-20, and Wisconsin Computer Sciences Department Chair 2014-2017. Hill was Partner Hardware Architect at Microsoft (2020-2024) where he led software-hardware pathfinding for Azure. Hill has a PhD in computer science from the University of California, Berkeley.

Long Biography

Mark D. Hill is the Gene M. Amdahl and John P. Morgridge Professor Emeritus of Computer Sciences at the University of Wisconsin-Madison (<http://www.cs.wisc.edu/~markhill>), following his 1988-2020 service in Computer Sciences and Electrical and Computer Engineering. He served on the Computing Community Consortium (CCC) 2013-21 including as CCC Chair 2018-20, Computing Research Association (CRA) Board of Directors 2018-20, and Wisconsin Computer Sciences Department Chair 2014-2017. Hill was Partner Hardware Architect at Microsoft (2020-2024) where he led software-hardware pathfinding for Azure.

Dr. Hill's research targets computer design and evaluation. He has contributed to parallel computer system design (e.g., memory consistency models and cache coherence), memory system design (caches and translation buffers), computer simulation (parallel systems and memory systems), software (e.g., page tables and cache-conscious optimizations), and deterministic replay and transactional memory. For example, he is the inventor of the widely-used *3C model* of cache behavior (*compulsory*, *capacity*, and *conflict* misses) and co-inventor of the memory consistency model *sequential consistency for data-race-free programs* that serves as a foundation for the C++ and Java multi-threaded memory specifications. Hill's work is highly collaborative with over 170 co-authors and especially with his long-time colleague David A. Wood.

Hill received the ACM - IEEE CS Eckert-Mauchly Award in 2019 for *seminal contributions to the fields of cache memories, memory consistency models, transactional memory, and simulation*. He was selected as a John P. Morgridge Endowed Chair of UW-Madison Computer Sciences in 2015. He was selected AAAS Fellow (2022) for *distinguished contributions to advanced computer architecture, particularly for memory system design*, named an ACM Fellow (2004) for *contributions to memory consistency models and memory system design*, elevated to Fellow of the IEEE (2000) for *contributions to cache memory design and analysis*, and was awarded the ACM SIGARCH Alan Berenbaum Distinguished

Service Award in 2009. He was won four important University of Wisconsin-Madison awards: WARF Named Professorship in 2013 (2nd winner from CS), Kellett in 2010 (3rd winner from CS), Vilas Associate in 2006, and Romnes Fellow in 1997. He co-wrote *A Primer on Memory Consistency and Cache Coherence* in 2011 (downloaded over 10K times) and 2nd Edition 2020, co-edited *Readings in Computer Architecture* in 2000, is co-inventor of over 40 United States patents (several of which have been co-issued in the European Union & Japan), was an ACM SIGARCH Director of ACM (1993-2007), and won an NSF Presidential Young Investigator award in 1989. He is co-author of eight papers selected by IEEE Micro Top Picks, co-won the best paper award in VLDB 2001, and has an H-index over 80 and over 30,000 citations (per Google Scholar). He has held visiting positions at Google (2018), Advanced Micro Devices (2011), University of Washington (2011), Columbia University (2010), Universidad Politécnica de Cataluña (2002-03) and Sun Microsystems (1995-96). Dr. Hill earned a Ph.D. in Computer Science from the University of California - Berkeley in 1987, an M.S. in Computer Science from Berkeley in 1983, and a B.S.E. in Computer Engineering from the University of Michigan – Ann Arbor in 1981.

Education

Ph.D. (Computer Science), University of California, Berkeley, November 1987. Thesis: *Aspects of Cache Memory and Instruction Buffer Performance* (UCB/CSD 87/381). Advisors: David A. Patterson and Alan Jay Smith.

M.S. (Computer Science), University of California, Berkeley, December 1983. Thesis: *Evaluation of On-Chip Cache Memories*. Advisor: Alan Jay Smith.

B.S.E. *summa cum laude* (Computer Engineering), University of Michigan, Ann Arbor, May 1981.

Selected Awards

Fellow of the American Association for the Advancement of Science (AAAS), 2022.

ACM - IEEE CS Eckert-Mauchly Award, 2019, *for seminal contributions to the fields of cache memories, memory consistency models, transactional memory, and simulation*. Highest computer architecture award.

John P. Morgridge Endowed Chair, 2015, in Computer Sciences at UW-Madison.

ACM Fellow, 2004, *for contributions to memory consistency models and memory system design*. Conferred by the Association for Computing Machinery (ACM) on those who have distinguished themselves by outstanding technical and professional achievements in information technology.

Fellow of the IEEE, 2000, *for contributions to cache memory design and analysis*. The highest grade of membership in the Institute of Electrical and Electronic Engineers (IEEE). Named IEEE Senior Member in 1995.

ACM SIGARCH Alan Berenbaum Distinguished Service Award, 2009. Awarded by ACM Special Interest Group on Computer Architecture (SIGARCH) *for important service to the computer architecture community*, includes travel to ISCA and \$1000. Second DSA award ever given. The award was named for Alan Berenbaum in 2016.

University of Wisconsin-Madison Awards. Recipients may come from *any* department at UW-Madison.

1. **WARF Named Professorship**, 2013. Awarded to eight professors among the over 2000 UW-Madison faculty and includes \$75,000 research funds. Second award ever in Computer Sciences. Named the Gene M. Amdahl Professor with Amdahl's blessing.

2. **Kellett Mid-Career**, 2010. Awarded to professors within twenty years of tenure and includes \$60,000 research funds. Third award ever in Computer Sciences.
3. **Vilas Associate**, 2006-2008. Awarded to professors within twenty years of tenure, includes four summer months of salary and \$25,000 research funds.
4. **Romnes Fellowship**, 1997-2002. Awarded to recently-tenured professors and include \$50,000 research funds. Second award ever in Computer Sciences.

ISCA@50 Retrospective: 1996-2020 includes A “Flight Data Recorder” for Enabling Full-system Multiprocessor Deterministic Replay.

Best Paper Awards. ACM ASPLOS Distinguished Paper 2023, IEEE Micro Top Picks 2016, IEEE Micro Top Picks 2015, IEEE Micro Top Picks 2012, Top Picks 2007 (2 papers), IEEE Micro Top Picks 2006, IEEE Micro Top Picks 2004, IEEE Micro Top Picks 2003, and VLDB 2001.

An Influential Paper from Computer’s First 50 Years. In 2017, *Amdahl’s Law in the Multicore Era* [2008] was honored as one of seven influential papers from IEEE Computer’s first 50 years (<https://www.computer.org/computer-magazine/from-the-archives-computers-legacy/>).

Selected Papers from the First 25 International Symposia on Computer Architecture, ACM Press, 1998, included *Weak Ordering – A New Definition*, co-authored with Sarita Adve for ISCA 1990.

Presidential Young Investigator Award, 1989. Five-year National Science Foundation grant and matching funds that *recognizes research and teaching accomplishments as well as academic potential*. Recipient of one of three national awards in computer architecture.

Conference Halls of Fame (for prolific publishing).

1. International Symposium on Computer Architecture (ISCA) <http://pages.cs.wisc.edu/~arch/www/iscabibhall.html>
2. International Symposium on Microarchitecture (MICRO) <http://newsletter.sigmicro.org/micro-hof.txt/view>

Erdős Number of 3, because Paul Erdős wrote a paper with Michael E. Saks, who wrote a paper with Anne E. Condon, who several papers with Hill.

Selected Professional Experience

Gene M. Amdahl & John P. Morgridge Professor Emeritus, Computer Sciences Department, University of Wisconsin, Madison, August 2020-present.

Partner Hardware Architect, Azure, Microsoft Corp., December 2020-2024.

Gene M. Amdahl & John P. Morgridge Professor, Computer Sciences Department and Electrical and Computer Engineering Department, University of Wisconsin, Madison, July 2015-August 2020.

Visiting Scientist (sabbatical), Google, Mountain View, CA, January-August 2018.

Computer Sciences Department Chair, University of Wisconsin, Madison, July 2014-June 2017.

Gene M. Amdahl Professor, Computer Sciences Department and Electrical and Computer Engineering Department, University of Wisconsin, Madison, 2013-present.

Professor, Computer Sciences Department and Electrical and Computer Engineering Department, University of Wisconsin, Madison, March 1998-2013.

Consultant (sabbatical), Advanced Micro Devices, Inc. (AMD), Bellevue, Washington, and **Visiting Scholar** (sabbatical), Computer Science & Engineering Department, University of Washington, Seattle, Washington, January-July 2011.

Visiting Senior Research Scientist (sabbatical), Computer Science Department, Columbia University, New York, New York, September-December 2010.

Profesor Visitante (sabático), Arquitectura de Computadores, Facultad de Informática, Universidad Politécnica de Cataluña, Jordi Girona 1-3, Campus Nord, Módulo D-6, 08034 Barcelona, España, Agosto 2002-Junio 2003.

Associate Professor, Computer Sciences Department and Electrical and Computer Engineering Department, University of Wisconsin, Madison, November 1993-March 1998.

Senior Visiting Scientist (sabbatical), Central Technology Office of Sun Microsystems Computer Corporation, Sun Microsystems, Inc., Menlo Park, California, September 1995-June 1996.

Assistant Professor, Computer Sciences Department and Electrical and Computer Engineering Department, University of Wisconsin, Madison, January 1988-November 1993.

Research Assistant, Electrical Engineering and Computer Science Department, University of California, Berkeley, July 1982-November 1987 (except Summer 1983).

Research Intern, Xerox PARC, Palo Alto, California, June 1983-August 1983; **Teaching Assistant**, U.C. Berkeley, September 1981-June 1982, **Research Assistant**, University of Michigan, Ann Arbor, September 1980-August 1981, **Design Engineer**, Texas Instruments, Dallas, Texas, May 1980-August 1980.

Consulting includes AMD Research (2011-2017), Microsoft Research (2009-2011), National Science Foundation (NSF) (2004-05), and Sun Microsystems (1995-2009).

Selected Grants & Donations

1. *Computing Innovation Fellows 2020 Project*, 5/2020-4/2024, \$16,737,291, National Science Foundation (CCF-2030859), Cooperative Agreement with Computing Research Association (CRA). Co-PI with PI Ellen Zegura and other Co-PIs Andrew Bernat, Liz Bradley & Ann Drobnis. Includes \$2,782,246 supplement.
2. *Supplement to Computing Community Consortium III*, 1/2020-3/2022, \$1,029,990, National Science Foundation (CCF-1734706), Cooperative Agreement with Computing Research Association (CRA). PI with Co-PIs Liz Bradley & Ann Drobnis.
3. *CSR:Small: System Support for Petabyte Memories*, 6/2018-6/2021, \$499,813, National Science Foundation (CNS-1815656), Co-PI; PI is Michael M. Swift.
4. *Computing Community Consortium III*, 4/2018-3/2022, \$8,485,854, National Science Foundation (CCF-1734706), Cooperative Agreement with Computing Research Association (CRA). PI with Co-PIs Liz Bradley & Ann Drobnis.
5. *Research Donation*, 2014, ~\$53K. Google Corp, Co-PI with Mike Swift and David Wood.
6. *XPS: CSR:Small: Architectural and Operating System Support for Non-Volatile Memory*, 9/2016-8/2018, \$499,059, National Science Foundation (CCF-1617824), Co-PI with David A. Wood; PI is Michael M. Swift.
7. *XPS: FULL: CCA: NUMB: Exploiting Non-Uniform Memory Bandwidth for Computational Science*, 9/2015-8/2018, \$749,994, National Science Foundation (CCF-1533885), Co-PI with Eftychios Sifakis & Michael M. Swift; PI is David A. Wood.
8. *XPS: FULL: CCA: An Initial Exploration of Architecture and System Support for Sparsity in Science and Engineering Applications*, 9/2014-8/2016, \$199,995, National Science Foundation (CCF-1438992), Co-PI with Eftychios Sifakis & Michael M. Swift; PI is David A. Wood.
9. *Research Donation*, 2014, ~\$90K. Google Corp, Co-PI with David Wood.
10. *CSR:Medium: WasteNot: Streamlining Virtual Memory for Modern Systems*, 9/2013-8/2016, \$749,103, National Science Foundation (CNS-1302260), Co-PI with David A. Wood; PI is Michael M. Swift.
11. *Research Donation*, 2013, ~\$90K. Google Corp, Co-PI with David Wood.
12. *SHF:Small: Energy-Optimized Memory Hierarchies*, 8/2012-7/2015, \$400,000, National Science Foundation (CCF-1218323), Co-PI with David A. Wood PI.
13. *CSR: Small: Codesign of Accelerator Interface Software and Hardware*, 8/2011-7/2014, \$400,000, National Science Foundation (CNS-1117280), Co-PI with David A. Wood PI and Michael Swift co-PI.
14. *Research Donation*, 2011, ~\$50K. AMD Corp, Co-PI with David Wood.
15. *Research Donation*, 2010, ~\$60K. Google Corp, Co-PI with David Wood.

16. *Simulation and Power Husbanding Project*, 4/2010-2/2012, \$268,800, Sandia National Laboratories, Co-PI with David A. Wood PI.
17. *SHF:Small: Power Husbanding via Architectural Techniques (PHAT)*, 8/2010-7/2013, \$500,000 National Science Foundation (CCF-1017650). Co-PI with David A. Wood PI.
18. *SHF:Small: Managing Non-Determinism in Multithreaded Software and Hardware Multithreaded Record, Replay, and Execution*, 8/2009-7/2012, \$499,998, National Science Foundation (CCF-0916725), Co-PI David A. Wood.
19. *Deterministic Multithreaded Record, Replay, and Execution*, 1/2009-9/2009, \$70,181, Sandia National Laboratories (#MSN123960 flow through from DOE 890426).
20. *CSR—AES: Deconstructing Transactional Memory: System Support for Robust Concurrent Programming*, 7/2007-6/2010, \$919,999, National Science Foundation (CNS-0720565), Co-PIs Michael M. Swift and David A. Wood.
21. *CRI: MASSIV Cluster for Designing Chip Multiprocessors*, 6/2006-5/2010, \$401,918, National Science Foundation CNS Computing Research Infrastructure (CNS-0551401), Co-PIs Gurindar S. Sohi and David A. Wood.
22. *Research Donation*, 2006, ~\$40K. Intel Corp, Co-PI with David Wood.
23. *Research Donation*, 2005, A second cluster of 32 dual-processor PCs (list ~\$200K), Sun Microsystems, Co-PI with David Wood.
24. *Research Donation*, 2005, ~\$40K. Intel Corp, Co-PI with David Wood.
25. *Research Donation*, 2004, A cluster of 32 dual-processor PCs (list \$235,054), Sun Microsystems, Co-PI with David Wood.
26. *Research Donation*, 2004, \$37,107. Intel Corp, Co-PI with David Wood.
27. *Advanced Architectures and Technologies for Chip Multiprocessors*, 9/2003-8/2008, \$1,110,370, National Science Foundation CISE ITR (CCR-0324878), Co-PI with David A. Wood.
28. *Research Donation*, 2003, \$100,000 and two 8-processor SUN FIRE V880 servers (list \$220,470), Sun Microsystems, Co-PI with David Wood.
29. *Research Donation*, 2003, \$37,107. Intel Corp, Co-PI with David Wood.
30. *SafetyNet: Synergistic Support for Availability, Designability, Programmability, & Performance*, 9/2002-8/2007, \$1,450,000, National Science Foundation CISE ITR (EIA/CNS-0205286), Co-PI with David A. Wood with Investigator is Rastislav Bodik.
31. *Exploiting the Critical Path in the Design and Performance Analysis of Modern Processors*, 2001-2004, \$422,507, National Science Foundation (CCR-0105721), Co-PI with Rastislav Bodik.
32. *Kilo-Instruction Checkpoint and Recovery for Supporting Speculation and Hardware Fault Tolerance*, 2000-2002, \$50,000, Intel Research Donation, Co-PI with David Wood.
33. *Multifacet: Exploiting Prediction and Speculation in Multiprocessor Memory Systems*, 1999-2002, \$1,081,850, National Science Foundation CISE Experimental Partnerships (EIA-9971256), Co-PI with David A. Wood with Investigators Pei Cao, Anne Condon, and Charles Fischer.
34. *Multifacet: Exploiting Prediction and Speculation in Multiprocessor Memory Systems*, 1999-2002, \$375,000, Sun Microsystems Collaborative Research Grant, Co-PI with David Wood.
35. *Lamport Clocks: Reasoning About Shared Memory Correctness*, 1998-2000, \$50,000, Intel Research Donation, Co-PI with Anne Condon.
36. *Moleware and the Molecular Computer*, 1999-2001, \$1,998,871, Defense Advanced Research Projects Agency (DARPA), Consultant, Head PI James M. Tour (Rice) with Co-PIs David L. Allara (Penn State), Mark A. Reed (Yale), Jorge Seminario (South Carolina), and Paul S. Weiss (Penn State).
37. *MIDSHIP: Managing Image Data for Scalable High Performance*, 1996-2001, \$1,598,026, National Science Foundation institutional infrastructure grant (CDA-9623632), partially matched by the University of Wisconsin Graduate School, Project co-director with Jeffery Naughton and co-PI with nine other faculty.
38. *Fine-Grain Distributed Shared Memory on SMP Clusters*, 1996-1999, \$300,000, Sun Microsystems Collaborative Research Grant, Co-PI with James Larus and David Wood.

39. *Tornado: Fine-Grain Distributed Shared Memory for SMP Clusters*, 1996-1999, \$1,134,962, National Science Foundation MIPS Experimental Systems (MIPS-9625558), Co-PI with David Wood, James Larus, and Pei Cao.
40. *Virtually Prototyping Scalable Coherence Protocols Using the Wisconsin Wind Tunnel*, 1995-1996, 2000 CM-5 service units, National Center for Supercomputing Applications (NCSA), a NSF Supercomputer Center, Co-PI with James R. Goodman and David Wood.
41. *Cooperative Shared Memory and the Wisconsin Wind Tunnel (Supplement)*, 1994-1996, \$224,896, National Science Foundation MIPS Experimental Systems (MIPS-9225097), Co-PI with James Larus and David Wood.
42. *Blizzard and Paradyn: Infrastructure and Scalable Tools for Multi-Paradigm Parallel Computers*, 1994-1997, \$2,371,525, Advanced Research Projects Agency (ARPA order no. B550) through Wright Laboratory Avionics Directorate, USAF (grant #F33615-94-1-1525), Co-PI with Barton Miller, James Larus and David Wood.
43. *Cooperative Shared Memory and the Wisconsin Wind Tunnel*, 1993-1996, \$1,483,081, National Science Foundation MIPS Experimental Systems (MIPS-9225097), Co-PI with James Larus and David Wood.
44. *PRISM: A Laboratory for Research in Future High-Performance Parallel Computing*, 1991-1995, \$2,000,000, National Science Foundation institutional infrastructure grant (CDA-9024618), partially matched by the University of Wisconsin Graduate School, Co-PI with Michael Carey, Charles Dyer, Robert Meyer, Barton Miller, and Mary Vernon (project coordinator).
45. *Presidential Young Investigator Award: Cache Memory Design*, 1989-1994, \$500,000, Five-year grant and matching funds from the National Science Foundation (MIPS-8957278). PYI matching funds have been donated by A.T.&T. Bell Laboratories, Cray Research, Digital Equipment Corporation, Sun Microsystems, and Texas Instruments.
46. *Memory Models for Large-Scale, Shared-Memory Multiprocessors*, 1991-1992, \$95,642, Digital Equipment Corporation external research grant.
47. *A High Speed Data Acquisition System for Research in Parallel Computing*, 1990-1991, \$100,642, National Science Foundation equipment grant (CDA-8920777), partially matched by the University of Wisconsin Graduate School and A.T.&T. Bell Laboratories, Co-PI with Mary Vernon.
48. *The Design of Secondary Caches*, 1989-1991, \$55,538, Two-year National Science Foundation grant (CCR-8902536).

Professional Activities

Major Consulting

- AMD Research (2011-17)
- Microsoft Research (2009-11)
- National Science Foundation (NSF) Division of Shared Cyberinfrastructure (SCI) (2004-05)
- Sun Microsystems (1995-09)

Computing Innovation Fellows 2020 (<https://cifellows2020.org/>)

- Co-Principal Investigator, Steering Committee Member, and Selection Committee Co-Chair providing 59 postdocs in computing in response to COVID-19 impact on academic job market.

National Boards

- Computing Community Consortium (CCC) Chair Emeritus 2020-21, Chair 2018-20, Vice Chair 2016-18, Executive Committee 2014-21, Council 2013-21.
- Computing Research Association (CRA) Board of Directors Member 2018-20.
- Information Science and Technology (ISAT) Study Group, 2010-2013, that reports to the Defense Advanced Research Projects Agency (DARPA) on research directions. On Steering Committee 2012-2013.

Professional organization officer

- 2015-2017, Member of Executive Committee for IEEE Technical Committee on Computer Architecture (TCCA).

- 1993-2007, Member of Board of Directors for ACM SIGARCH (Special Interest Group on Computer Architecture).
- 1995-2002, Information Director for ACM SIGARCH (Special Interest Group on Computer Architecture).

Co-Authored White Papers

- Ian Foster, Daniel Lopresti, Bill Gropp, Mark D. Hill, and Katie Schuman, *A National Discovery Cloud: Preparing the US for Global Competitiveness in the New Era of 21st Century Digital Transformation*, Computing Community Consortium (CCC) White Paper, April 14, 2021. URL: <https://cra.org/ccc/wp-content/uploads/sites/2/2021/04/CCC-Whitepaper-National-Discovery-Cloud-2021.pdf>
- Thomas M. Conte, Ian T. Foster, William Gropp, and Mark D. Hill, *Advancing Computing's Foundation of US Industry & Society*, Computing Community Consortium (CCC) Quadrennial Paper, October 29, 2020. URL: <https://cra.org/ccc/wp-content/uploads/sites/2/2020/10/Advancing-Computings-Foundation-of-US-Industry-Society-Moores-Law-2.pdf>
- Vivek Sarkar, Nancy Amato, Susan Davidson, Eric de Sturler, David Ebert, Mark D. Hill, Charles Isbell, Shwetak Patel, Chris Ramming, Divesh Srivastava, Marvin Theimer, Benjamin Zorn, *CRA Industry/Academia Committee Report*, Computing Research Association, May 29, 2020. URL: https://cra.org/wp-content/uploads/2020/07/CRA-Industry_Academia-Committee-Report.pdf
- Vasant G. Honavar, Katherine Yelick, Klara Nahrstedt, Holly Rushmeier, Jennifer Rexford, Mark D. Hill, Elizabeth Bradley, and Elizabeth Mynatt, *Advanced Cyberinfrastructure for Science, Engineering, and Public Policy*, Computing Community Consortium, Jun 30, 2017. URL: <https://arxiv.org/abs/1707.00599>
- Luis Ceze, Mark D. Hill, Karthikeyan Sankaralingam, Thomas F. Wenisch, *Democratizing Design for Future Computing Platforms*, Computing Community Consortium, Jun 26, 2017. URL: <https://arxiv.org/abs/1706.08597>
- Randy Bryant, Mark Hill, Tom Kazior, Daniel Lee, Jie Liu, Klara Nahrstedt, Vijay Narayanan, Jan Rabaey, Hava Siegelmann, Naresh Shanbhag, Naveen Verma, H.-S. Philip Wong, *Nanotechnology-inspired Information Processing Systems of the Future*, Computing Community Consortium, released February 13, 2017. URL: <http://cra.org/ccc/Nanotech-Report>
- Luis Ceze, Mark D. Hill, Thomas F. Wenisch, *Arch2030: A Vision of Computer Architecture Research over the Next 15 Years*, Computing Community Consortium, December 9, 2016. Also appears in Communications of the China Computer Federation (CCCF), Volume 13, Issue 7, July 2017, translated into Chinese by State Key Laboratory of Computer Architecture. URL: <http://arxiv.org/abs/1612.03182>
- Vasant G. Honavar, Mark D. Hill, and Katherine Yelick, *Accelerating Science: A Computing Research Agenda*, Computing Community Consortium, February, 19, 2016. URL: <http://cra.org/ccc/wp-content/uploads/sites/2/2016/02/Accelerating-Science-Whitepaper-CCC-Final2.pdf>
- Gregory D. Hager, Mark D. Hill, and Katherine Yelick, *Opportunities and Challenges for Next Generation Computing*, Computing Community Consortium, October 19, 2015. URL: <http://cra.org/ccc/wp-content/uploads/sites/2/2015/10/NextGenComputingChallenges.pdf>

Led White Paper Effort

- *21st Century Computer Architecture*, Mark D. Hill, Sarita Adve, Luis Ceze, Mary Jane Irwin, David Kaeli, Margaret Martonosi, Josep Torrellas, Thomas F. Wenisch, David Wood, Katherine Yelick, Computing Community Consortium, May 25, 2012. URL: <https://arxiv.org/abs/1609.06756> Chinese translation appeared in CACM equivalent: Computing Community Consortium (CCC), 21st Century Computer Architecture, Communication of the China Computer Federation (CCF), 2012,

Volume 8, Issue 12, p70-81. URL:

<http://www.ccf.org.cn/resources/1190201776262/2012/12/17/15.pdf>. This white paper has been annually cited by NSF in *Exploiting Parallelism and Scalability & Scalable Parallelism in the Extreme* programs that have awarded \$16M/year (\$90M total by 2019) and cited again by the 2020 *Principles and Practice of Scalable Systems* program that promises up to \$87M more in the 2020s.

- *Advancing Computer Systems without Technology Progress*, Mark D. Hill and Christos Kozyrakis, ISAT Outbrief, April 17-18, 2012, of DARPA/ISAT Workshop, March 26-27, 2012. Cited by and contributed to 2017 DARPA Electronics Resurgence Initiative, \$75M in 2017 and projected to grow to \$1.5B.

Chair

- ACM/IEEE Committee on the Health of Conferences, 2005-2006, developed WIKI (<http://wiki.acm.org/healthcc/>).

National Academy Study Member

- “Global Approaches to Advanced Computing” under the Policy and Global Affairs Division (PGA) of the National Research Council (NRC), 2011-12. Report: *The New Global Ecosystem in Advanced Computing: Implications for U.S. Competitiveness and National Security*, 2012.
- “Sustaining Growth in Computing Performance” under the Computer Science and Telecommunications Board (CSTB) of the National Research Council (NRC), 2007-2010. Report: *The Future of Computing Performance: Game Over or Next Level?*, 2011.
- “Fundamentals of Computer Science” under the Computer Science and Telecommunications Board (CSTB) of the National Research Council (NRC), 2001-2003. Report: *Computer Science: Reflections on the Field, Reflections from the Field*, 2004. Chapter 2a by Mark D. Hill. “Harnessing Moore’s Law.”

Member

- IEEE von Neuman Award Selection Committee, 2021-23.
- Computer Research Association (CRA) Selection Committee for Executive Director Finalists, 2021.
- Computer Research Association (CRA) Committee on Best Practices for Hiring, Promotion, and Scholarship, 2014-15.
http://cra.org/resources/bp-view/best_practices_memo_evaluating_scholarship_in_hiring_tenure_and_promot/
- ACM-IEEE Eckert-Mauchly Selection Committee, 2012-2014 (Chair 2014).
- ACM SIGARCH Distinguished Service Award Selection Committee, 2010-2013 (Chair 2013).
- National Science Foundation Study Panel on Computer Performance Evaluation, 2002.
- Search Committee for Editor-in-Chief of IEEE Transactions on Parallel and Distributed Systems, 2001.
- Defense Science Study Group (DSSG), 1998-1999. Directed by the Institute for Defense Analyses (IDA) and sponsored by the Defense Advanced Research Projects Agency (DARPA) (<http://dssg.ida.org>).

Founding Editor of Morgan & Claypool Synthesis Lectures on Computer Architecture (now

Springer). Edited 29 lectures (monographs) <https://www.springer.com/series/16916>, 2005-2014

- Hari Angepat, Derek Chiou, Eric S. Chung, and James C. Hoe, *FPGA-Accelerated Simulation of Computer Systems*, August 2014
- Babak Falsafi and Thomas F. Wenisch, *A Primer on Hardware Prefetching*, May 2014
- Christopher J. Nitta, Matthew K. Farrens, and Venkatesh Akella, *On-Chip Photonic Interconnects: A Computer Architect’s Perspective*, October 2013
- Tony Nowatzki, Michael Ferris, Karthikeyan Sankaralingam, Cristian Estan, Nilay Vaish, and David Wood, *Optimization and Mathematical modeling in Computer Architecture*, September 2013

-
- Ruby B. Lee, *Security Basics for Computer Architects*, September 2013
 - Luiz Andre Barroso, Jimmy Clidaras, and Urs Holzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale machines*, Second edition, July 2013
 - Michael L. Scott, *Shared-Memory Synchronization*, July 2013
 - Vijay Janapa Reddi and Meeta Sharma Gupta, *Resilient Architecture Design for Voltage Variation*, June 2013
 - Mario Nemirovsky and Dean M. Tullsen, *Multithreading Architecture*, January 2013
 - Hyesoon Kim, Richard Vuduc, Sara Baghsorkhi, Jee Choi, and Wen-mei Hwu, *Performance Analysis and Tuning for General Purpose Graphics Processing Units (GPGPU)*, November, 2012
 - Samuel P. Midkiff, *Automatic Parallelization: An Overview of Fundamental Compiler Techniques*, January 2012
 - Moinuddin K. Qureshi, Sudhanva Gurumurthi, and Bipin Rajendran, *Phase Change Memory: From Devices to Systems* November, 2011
 - Rajeev Balasubramonian, Norman P. Jouppi, and Naveen Muralimanohar, *Multi-Core Cache Hierarchies*, November 2011.
 - Daniel J. Sorin, Mark D. Hill, and David A. Wood, *A Primer on Memory Consistency and Cache Coherence*, 2011.
 - Kim Hazelwood, *Dynamic Binary Modification: Tools, Techniques, and Applications*, March 2011
 - Tzvetan S. Metodi, Arvin I. Faruque, and Frederic T. Chong, *Quantum Computing for Computer Architects*, Second Edition March 2011
 - Dennis Abts and John Kim, *High Performance Datacenter Networks: Architecture, Algorithms, and Opportunities*, March 2011
 - Antonio Gonzalez, Fernando Latorre, and Grigorios Magklis, *Processor Microarchitecture: An Implementation Perspective*, December 2011
 - Tim Harris, James Larus, and Ravi Rajwar, *Transactional Memory*, 2nd Edition, December 2010
 - Lieven Eeckhout, *Computer Architecture Performance Evaluation Methods*, December 2010
 - Marco Lanzagorta, Stephen Bique, and Robert Rosenberg, *Introduction to Reconfigurable Supercomputing*, 2009
 - Natalie Enright Jerger and Li-Shiuan Peh, *On-Chip Networks*, 2009
 - Bruce Jacob, *The Memory System: You Can't Avoid It, You Can't Ignore It, You Can't Fake it*, 2009
 - Daniel J. Sorin, *Fault Tolerant Computer Architecture*, 2009
 - Luiz Andre Barroso and Urs Holzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*, 2009
 - Stefanos Kaxiras and Margaret Martonosi, *Computer Architecture Techniques for Power-Efficiency*, 2008
 - Kunle Olukotun, Lance Hammond, and James Laudon, *Chip Multiprocessor Architecture: Techniques to Improve Throughput and Latency*, 2007
 - James R. Larus and Ravi Rajwar, *Transactional Memory*, 2006
 - Tzvetan S. Metodi, Frederic T. Chong, *Quantum Computing for Computer Architects*, 2006

Other Editorships

- Associate Editor, ACM Transactions on Parallel Computing (TOPC), 2013-2014.
 - Associate Editor, ACM Transactions on Computer Systems (TOCS), 2003-2012.
 - Associate Editor, IEEE TCCA Computer Architecture Letters, 2001-2010.
 - Associate Editor, IEEE Transactions on Parallel and Distributed Systems, 1994-1997.
-

Conference officer

- Tutorial Chair for ACM/IEEE International Symposium on Computer Architecture (ISCA), 1995.

Conference technical program committees

- Hot Chips 2022
- HPCA Industrial Track PC: 2023, 2022, 2021, 2020
- ISCA Industrial Track PC: 2022, 2021, 2020 (1st one)
- ASPLOS Extended PC: 2022, 2020, 2019
- HPCA Extended PC: 2021, 2020, 2019
- ISCA Extended PC: 2020, 2019
- MICRO Extended PC: 2020, 2019
- (Wisconsin CS department chair 2014-2017 + 2018 sabbatical)
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2014.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2012.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2011.
- IEEE International Symposium on High-Performance Computer Architecture (HPCA), 2010.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2009.
- ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 2008.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2008.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2007.
- Conference on High-Performance Embedded Architecture and Compilation (HiPEAC), 2006.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), *Chair*, 2005.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2004.
- ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI), 2002.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 2001.
- ACM Symposium on Parallel Algorithms and Architectures (SPAA), 2001.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 1999.
- ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 1998.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 1998.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 1997.
- IEEE International Symposium on High-Performance Computer Architecture (HPCA), 1997.
- ACM/IEEE International Symposium on Computer Architecture (ISCA), 1995.
- ACM Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), 1995.
- ACM Symposium on Parallel Algorithms and Architectures (SPAA), 1995.
- ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 1994.
- ACM International Conference on Supercomputing (ICS), 1993.
- ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), 1991.

Other program committees

- IEEE Micro Top Picks in Computer Architecture, 2012.
-

- Advancing Computer Architecture Research (ACAR2): What Now in ILP Research?, September 2009.
- Advancing Computer Architecture Research (ACAR1): Failure is not an Option: Popular Parallel Programming, February 2009.
- IEEE Micro Top Picks in Computer Architecture, 2009.
- IEEE Micro Top Picks in Computer Architecture, 2007.
- Workshop on Design, Architecture and Simulation of Chip Multi-Processors (dasCMP), 2007
- IEEE Micro Top Picks in Computer Architecture, 2006.
- IEEE International Symposium on High-Performance Computer Architecture (HPCA) Industrial Track, 2006.
- ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS) Workshop of Architectural Support for Improving Software Dependability (ASID), 2006.
- IEEE Micro Top Picks in Computer Architecture, 2005.
- ACM SIGPLAN Workshop on Memory System Performance (in conjunction with PLDI), 2002.
- AEA/CBEMA/CRA/CSPP/CoC/CIWT/EDUCOM Forum on the R&D Agenda for the National Information Infrastructure, 1994. Co-chair of network and system components track and principal author of *Information Appliances and Servers* section of final report (available from EDUCOM).

Presented Tutorials

- International Summer School on Advanced Computer Architecture and Compilation for Embedded Systems (ACACES), July 2006. *Transactional Memory*.
- ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, May 1994. *Implementing High-Performance Shared Memory*.
- ACM International Conference on Architectural Support for Programming Languages and Operating Systems, October 1992. *High-Performance Shared-Memory: Implementation Issues and Consistency Models* (with Sarita Adve).
- IEEE Hot Chips Symposium, August 1992. *Quantitative Approach to Microprocessor Architecture*.
- ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems, May 1992. *Cache Design and Evaluation*.
- ACM/IEEE International Symposium on Computer Architecture, May 1990. *Uniprocessor Cache Design and Evaluation*.

Guest Editor

- Special Issue of IEEE Computer on Network Interfaces, November 1998 (with Andrew Chien and Shubhendu Mukherjee).
- Special Issue of IEEE Micro on Hot Chips II, June 1991 (with David Wood).

NSF Panel Member

- Many, including Engineering Research Center (ERC) site visit team, September 2010.

Conference Panels

- *Research Methodology in the New Era of Computer Architecture*, Panel Member, International Symposium on Computer Architecture (ISCA), Virtual, June 2021.
 - *Cybersecurity: Transcending Physics, Technology, & Society*, Panel Moderator with panelists Kevin Fu (Michigan), Jon Masters (Redhat), and Zeynep Tufekci (North Carolina and New York Times), Annual Meeting of the American Association for the Advancement of Science, Washington, DC, February, 2019.
-

- *Celebration of 2017 Turing Award*, Panel Member with David Patterson (Berkeley/Google) and John Markoff (New York Times), Semicon West, San Francisco, CA, July 2018.
- *On the Implications of the Meltdown & Spectre Design Flaws*, Panel Moderator, International Symposium on Computer Architecture (ISCA), Los Angeles, CA, June 2018.
- *Is the Death of Moore's Law Making Computer Architecture Livelier Than Ever?*, Panel Member, International Symposium on Computer Architecture (ISCA), Los Angeles, CA, June 2017.
- *Whistling Past the Graveyard: Why the End of Moore's Law Matters to All of CS*, Panel Member, Computing Research Association Conference at Snowbird, Snowbird, Utah, July 2016.
- *Moore's Law Challenges Below 10nm*, Panel Member, IEEE International Solid-State Circuits Conference (ISSCC), San Francisco, CA, February 2015. The panel won the ISSCC 2015 outstanding evening session award.
- *Research Directions for 21st Century Computer Systems*, Panel Moderator, ACM/IEEE International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS), Houston, TX, March 2013.
- *La Microarchitecture est Morte. Longue Vie à la Microarchitecture*, Panel Moderator, International Symposium on Computer Architecture (ISCA), St. Malo, France, June 2010.
- *Opportunities Beyond Single-Core Microprocessors*, Panel Moderator, International Symposium on High-Performance Computer Architecture (HPCA) and ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, (PPoPP), Raleigh, North Carolina, February 2009.
- *Computing Beyond Von Neumann*, IEEE/ACM International Symposium on Microarchitecture (MICRO), December 2007.
- *How will we develop and program emerging robust, low-power, adaptive multicore computing systems?*, International Conference on Parallel and Distributed Systems, July 2006.
- *How Much Longer Will SuperScalar Microarchitectures Scale?*, International Conference on Computer Design, September 2001.
- *Big Science vs. Little Science: Do You Have to Build it?*, Panel Moderator, International Symposium on Computer Architecture (ISCA), May 1990.

Other Panels

- *Single-Threaded vs. Multithreaded: Where Should We Focus?*, Workshop on Computer Architecture Research Directions (CARD 2007) at HPCA, February 2007.
- *ManyCoreDB: Will you still need me, will you still feed me, when I'm 64?*, Workshop on Data Management on New Hardware (DaMoN 2006) at SIGMOD, June 2006.

Referee for the following journals

- ACM Computing Surveys
- ACM Transactions on Computer Systems
- Communications of the ACM
- IEEE Transactions on Computers
- IEEE Transactions on Parallel and Distributed Systems
- IEEE Computer

Referee for the following conferences

- IEEE/ACM SIGARCH International Symposium on Computer Architecture.
- ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems.
- ACM SIGARCH, SIGOPS, SIGPLAN International Conference on Architectural Support for Programming Languages and Operating Systems.

Professional Society Memberships

- IEEE, including Computer Society and Technical Committee on Computer Architecture (TCCA).
- ACM, including SIGARCH, SIGMETRICS, SIGMICRO, and SIGOPS.

Teaching

Taught approximately one course per semester at the University of Wisconsin, Madison, since January 1988.

Most frequently taught courses:

- CS/ECE 252 (Freshman/Sophomore): Introduction to Computer Engineering.
- CS/ECE 354 (Sophomore/Junior Level): Machine Organization and Programming.
- CS/ECE 552 (Senior Level): Introduction to Computer Architecture.
- CS/ECE 752 (Graduate Level): Advanced Computer Architecture I.
- CS/ECE 757 (Graduate Level): Advanced Computer Architecture II.

Ph.D.'s Supervised

Swapnil Haria, Ph.D. August 2019, [*Architecture and Software Support for Persistent and Vast Memory*](#), first employment: VMware (co-advisor with Michael Swift).

Jason Lowe-Power, Ph.D. August 2017, [*On Heterogeneous Compute and Memory Systems*](#), first employment: Univ. of California-Davis (co-advisor with David Wood)

Jayneel Gandhi, Ph.D. August 2016, [*Efficient Memory Virtualization*](#), first employment: VMware (co-advisor with Michael Swift).

Lena E. Olson, Ph.D. August 2016, [*Protecting Host Systems from Imperfect Hardware Accelerators*](#), first employment: Google.

Arkaprava (Arka) Basu, Ph.D. December 2013, [*Revisiting Virtual Memory*](#), first employment: AMD (co-advisor with Michael Swift).

Derek R. Hower, Ph.D., July 2012, [*Acoherent Shared Memory*](#), first employment: AMD.

Jayaram Bobba, Ph.D., February 2010, [*Hardware Support for Efficient Transactional and Supervised Memory Systems*](#), first employment: Intel.

Luke Yen, Ph.D., February 2009, [*Signatures in Transactional Memory Systems*](#), first employment: AMD.

Michael R. Marty, Ph.D., January 2008, [*Cache Coherence Techniques for Multicore Processors*](#), first employment: Google.

Brian Fields, Ph.D., U.C. Berkeley, December 2006, [*Using Criticality to Attack Performance Bottlenecks*](#), first employment: VMware, co-advisor with Rastislav Bodik.

Min Xu, Ph.D., May 2006, [*Race Recording for Multithreaded Deterministic Replay using Multiprocessor Hardware*](#), first employment: VMware, co-advisor with Rastislav Bodik.

Milo Martin, Ph.D., December 2003, [*Token Coherence*](#), first employment: Assistant Professor at University of Pennsylvania (Penn).

Anastassia (Natassa) Ailamaki, Ph.D., December 2000, [*Architecture-Conscious Database Systems*](#), first employment: Assistant Professor at Carnegie-Mellon University (co-advisor with David DeWitt).

Eric Schnarr, Ph.D., December 2000, [*Applying Programming Language Implementation Techniques to Processor Simulation*](#), first employment: Hypercosm Corporation (co-advisor with James R. Larus).

Trishul Chilimbi, Ph.D., August 1999, [*Cache-Conscious Data Structures--Design and Implementation*](#), first employment: Microsoft Corporation (co-advisor with James R. Larus).

Shubhendu (Shubu) Mukherjee, Ph.D., May 1998, [*Design and Evaluation of Network Interfaces for System Area Networks*](#), first employment: Digital Equipment Corporation.

Ioannis (Yannis) Schoinas, Ph.D., December 1997, [*Fine Grain Distributed Shared Memory Systems on Clusters of Workstations*](#), first employment: Intel Corporation.

Madhusudhan (Madhu) Talluri, Ph.D., August 1995, [Use of Superpages and Subblocking in the Address Translation Hierarchy](#), first employment: Sun Microsystems.

Sarita V. Adve, Ph.D., November 1993, [Designing Memory Consistency Models for Shared-Memory Multiprocessors](#), first employment: Assistant Professor at Rice University.

Richard E. Kessler, Ph.D., July 1991, [Analysis of Multi-Megabyte Secondary CPU Cache Memories](#), first employment: Cray Research.

Books

1. Vijay Nagarajan, Daniel J. Sorin, Mark D. Hill, and David A. Wood, [A Primer on Memory Consistency and Cache Coherence](#), Second Edition, Synthesis Lectures in Computer Architecture, Morgan & Claypool Publishers, doi: <https://doi.org/10.2200/S00962ED2V01Y201910CAC049>, February 2020.
2. Daniel J. Sorin, Mark D. Hill, and David A. Wood, [A Primer on Memory Consistency and Cache Coherence](#), Synthesis Lectures in Computer Architecture, Morgan & Claypool Publishers, doi:10.2200/S00346ED1V01Y201104CAC016, May 2011 (downloaded over 10,000 times).
3. Mark D. Hill, Norman P. Jouppi, and Gurindar S. Sohi, [Readings in Computer Architecture](#), Morgan Kaufmann Publishers (now Elsevier), ISBN 1-55860-539-8, 2000 (includes 85 pages of new material in two-column conference-paper format).

Patents

U.S. Patents 5,951,657, 6,883,070, 6,981,097, 8,239,633, 9,158,704, 9,954,60, 9,619,401, and 10,713,165 are assigned to the Wisconsin Alumni Research Foundation (WARF). U.S. Patents 8,868,843, 8,954,672, 9,298,615, 9,361,118, and 9,753,858 are assigned to AMD from a 2011 sabbatical and subsequent consulting. Other patents are results from Hill's sabbatical (1995-1996) and subsequent consulting at Sun Microsystems and are assigned to Sun. Many have also issued as European Union and Japanese Patents.

1. Jason Lowe-Power, David A. Wood, and Mark D. Hill, "Adaptive computer cache architecture," *United States Patent 10,713,165*, issued July 14, 2020.
2. Gabriel H. Loh and Mark D. Hill, "DRAM cache with tags and data jointly stored in physical rows," *United States Patent 9,753,858*, issued September 5, 2017.
3. Jayneel Gandhi, Mark D. Hill, Michael M. Swift, "Efficient memory management system for computers supporting virtual machines," *United States Patent 9,619,401*, issued Apr 11, 2017.
4. Arkaprava Basu, Mark D. Hill, Michael M. Swift, "I/O memory management unit providing self invalidated mapping," *United States Patent 9,954,603*, issued Jan 17, 2017.
5. Derek R. Hower, Mark D. Hill, David Wood, Steven K. Reinhardt, Benedict R. Gaster, Blake A. Hechtman, Bradford M. Beckmann, "Method for memory consistency among heterogeneous computer components," *United States Patent 9,361,118*, issued May 12, 2016.
6. Lena E. Olson, Yasuko Eckert, Vikas K. Sridharan, James M. O'Connor, Mark D. Hill, Srilatha Manne, "Methods and apparatus for soft-partitioning of a data cache for stack data," *United States Patent 9,298,615*, issued March 29, 2016, filed July 19, 2013.
7. Arkaprava Basu, Mark D. Hill, and Michael M. Swift, "Virtual memory management system with reduced latency," *United States Patent 9,158,704*, issued October 13, 2015, filed January 24, 2013.
8. Gabriel H. Loh and Mark D. Hill, "System and method for cache organization in row-based memories," *United States Patent 8,954,672*, issued February 10, 2015, filed March 12, 2012.
9. Gabriel H. Loh and Mark D. Hill, "Hardware filter for tracking block presence in large caches," *United States Patent 8,868,843 A1*, issued October 21, 2014, filed November 30, 2011.

10. David A. Wood, Mark D. Hill, Michael M. Swift, Michael R. Marty, Luke Yen, Kevin E. Moore, Jayaram Bobba, and Haris Volos, "Non-broadcast Signature-based Transactional Memory," *United States Patent* 8,239,633 B2, issued August 27, 2012, filed July 9, 2008.
 11. Robert E. Cypher, David Wood A., Mark D. Hill, and Thomas M. Wicki, "Computer system implementing synchronized broadcast using timestamps," *United States Patent* 7,366,843, issued April 29, 2008, filed June 30, 2003.
 12. David Wood, Robert C. Zak, Jr., Monica Wong-Chan, Christopher J. Jackson, Thomas P. Webber, and Mark D. Hill, "System and method for enhancing communication between devices in a computer system," *United States Patent* 7,225,383, issued May 29, 2007, filed January 19, 2000.
 13. Robert E. Cypher, Mark D. Hill, and David A. Wood, "Computer system implementing synchronized broadcast using skew control and queuing," *United States Patent* 7,136,980, issued November 14, 2006, filed June 30, 2003.
 14. Milo M. Martin, Mark D. Hill, and David A. Wood, "Token based cache-coherence protocol," *United States Patent* 6,981,097 issued December 27, 2005, filed March 14, 2003.
 15. Milo M. Martin, Daniel J. Sorin, Mark D. Hill, and David A. Wood, "Bandwidth-adaptive, hybrid, cache-coherence protocol," *United States Patent* 6,883,070 issued April 19, 2005, filed October 19, 2001.
 16. Borus Ostrovsky, Daniel R. Cassidy, John R. Fehrer, David A. Wood, Pazhani Pillai, Christopher J. Jackson, and Mark D. Hill, "Method and device for a context-based memory management system," *United States Patent* 6,826,671 issued November 30, 2004, filed October 9, 2001.
 17. Erik E. Hagersten and Mark D. Hill, "Hierarchical SMP computer system," *United States Patent* 6,826,660 issued November 30, 2004, filed February 11, 2002.
 18. Erik E. Hagersten and Mark D. Hill, "Skewed finite hashing function," *United States Patent* 6,654,866 issued November 25, 2003, filed August 27, 2001.
 19. Erik E. Hagersten and Mark D. Hill, "Repeater for use in a shared memory computing system," *United States Patent* 6,578,071 issued June 10, 2003, filed March 15, 2001.
 20. Erik E. Hagersten and Mark D. Hill, "Methods and apparatus for a directory-less memory access protocol in a distributed shared memory computer system," *United States Patent* 6,574,659, issued June 3, 2003, filed March 20, 2000.
 21. Erik E. Hagersten and Mark D. Hill, "Hybrid memory access protocol in a distributed shared memory computer system," *United States Patent* 6,496,854, issued December 17, 2002, filed February, 25, 2000.
 22. Erik E. Hagersten and Mark D. Hill, "Methods and apparatus for a directory-less memory access protocol in a distributed shared memory computer system," *United States Patent* 6,377,980, issued April 23, 2002, filed January 25, 1999.
 23. Erik E. Hagersten and Mark D. Hill, "Skewed finite hashing function," *United States Patent* 6,308,246, issued October 23, 2001, filed September 4, 1998.
 24. Erik E. Hagersten and Mark D. Hill, "Hybrid memory access protocol in a distributed shared memory computer system," *United States Patent* 6,243,742, issued June 5, 2001, filed July 30, 1998. Also *European Union Patent* 0818732 and *Japanese Patent* 10177518.
 25. Erik E. Hagersten and Mark D. Hill, "Shared memory system for symmetric multiprocessor systems," *United States Patent* 6,226,671, issued May 1, 2001, filed July 30, 1998.
 26. David A. Wood, Steven K. Reinhardt, Shubhendu S. Mukherjee, Babak Falsafi, Mark D. Hill, and Robert W. Pfile, "Cachable interface control registers for high speed data transfer," *United States Patent* 5,951,657, issued September 14, 1999, assignee Wisconsin Alumni Research Foundation (WARF), filed June 9, 1997.
 27. Erik E. Hagersten, Mark D. Hill, and David A. Wood, "Methods and apparatus for substantially memory-less coherence transformer for connecting computer node coherence domains," *United States Patent* 5,940,860, issued August 17, 1999, filed July 1, 1996. Also *European Union Patent* 0817068 and *Japanese Patent* 11003277.
 28. Erik E. Hagersten and Mark D. Hill, "Split-SMP computer system configured to operate in a protected mode having repeater which inhibits transaction to local address partition," *United States Patent* 5,923,847, issued July 13, 1999, filed July 2, 1996. Also *European Union Patent* 0817094 and *Japanese Patent* 10187646.
-

29. Erik E. Hagersten and Mark D. Hill, "Skip-level write-through in a multi-level memory of a computer system," *United States Patent* 5,903,907, issued May 11, 1999, filed July 1, 1996. Also *European Union Patent* 0817079 and *Japanese Patent* 11003280.
30. Erik E. Hagersten and Mark D. Hill, "Efficient allocation of cache memory space in a computer system," *United States Patent* 5,893,150, issued April 6, 1999, filed July 1, 1996. Also *European Union Patent* 0817078 and *Japanese Patent* 10214229.
31. Erik E. Hagersten and Mark D. Hill, "Multiprocessor system configured to detect and efficiently provide for migratory data access patterns," *United States Patent* 5,734,922, issued March 31, 1999, filed July 1, 1996.
32. Erik E. Hagersten and Mark D. Hill, "Method and apparatus for a directory-less memory access protocol in a distributed shared memory computer system," *United States Patent* 5,873,117, issued February 16, 1999, filed July 1, 1996. Also *European Union Patent* 0817067 and *Japanese Patent* 10134009.
33. Erik E. Hagersten and Mark D. Hill, "Hybrid memory access protocol for servicing memory access request by ascertaining whether the memory block is currently cached in determining which protocols to be used," *United States Patent* 5,864,671, issued January 26, 1999, filed July 1, 1996.
34. Erik E. Hagersten and Mark D. Hill, "Hierarchical SMP computer system," *United States Patent* 5,862,357, issued January 19, 1999, filed July 2, 1996. Also *European Union Patent* 0817060 and *Japanese Patent* 10187630.
35. Erik E. Hagersten, Mark D. Hill, and David A. Wood, "Methods and apparatus for a coherence transformer for connecting computer system coherence domains," *United States Patent* 5,860,109, issued January 12, 1999, filed July 1, 1996. Also *European Union Patent* 0817065 and *Japanese Patent* 10214222.
36. Erik E. Hagersten and Mark D. Hill, "Methods and apparatus for sharing stored data objects in a computer system," *United States Patent* 5,835,906, issued November 10, 1998, filed July 1, 1996. Also *European Union Patent* 0817040 and *Japanese Patent* 10187527.
37. Erik E. Hagersten, Mark D. Hill, and David A. Wood, "Methods and apparatus for a coherence transformer with limited memory for connecting computer system coherence domains," *United States Patent* 5,829,034, issued October 27, 1998, filed July 1, 1996. Also *European Union Patent* 0817069 and *Japanese Patent* 10187633.
38. Erik E. Hagersten and Mark D. Hill, "Efficient storage of data in computer systems with multiple cache levels," *United States Patent* 5,802,563, issued September 1, 1998, filed July 1, 1996. Also *European Union Patent* 0817080 and *Japanese Patent* 10214224.
39. Erik E. Hagersten and Mark D. Hill, "Extended symmetrical multiprocessor address mapping," *United States Patent* 5,796,605, issued August 18, 1998, filed July 2, 1996. (I was officially added as a co-inventor on June 27, 1998, correcting a filing error.)
40. Erik E. Hagersten and Mark D. Hill, "Extended symmetrical multiprocessor architecture," *United States Patent* 5,754,877, issued May 19, 1998, filed July 2, 1996. Also *European Union Patent* 0817092 and *Japanese Patent* 10097513.
41. Erik E. Hagersten and Mark D. Hill, "Multiprocessing system configured to detect and efficiently provide for migratory data access patterns," *United States Patent* 5,734,922, issued March 31, 1998, filed July 1, 1996. Also *European Union Patent* 0817071 and *Japanese Patent* 10143483.

Rigorously-Refereed Conference Proceedings.

Most of Mark Hill's publications are available on-line at
<http://www.cs.wisc.edu/~markhill/publications.html>. See also Google Scholar page
<http://scholar.google.com/citations?user=7IVfIWYAAAAJ>).

1. Huaicheng Li, Daniel S. Berger, Stanko Novakovic, Lisa Hsu, Dan Ernst, Pantea Zardoshti, Monish Shah, Samir Rajadnya, Scott Lee, Ishwar Agarwal, Mark D. Hill, Marcus Fontoura, Ricardo Bianchini, [Pond: CXL-Based Memory Pooling Systems for Cloud Platforms](#), *Proc. 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2023. Selected Distinguished Paper among less than 10% of papers.

2. Swapnil Haria, Mark D. Hill, Michael M. Swift, [MOD: Minimally Ordered Durable Datastructures for Persistent Memory](#), *Proc. 25th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2020.
3. Mark D. Hill and Vijay Janapa Reddi, [Gables: A Roofline Model for Mobile SoCs](#), *Proc. 25th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, February 2019.
4. Swapnil Haria, Mark D. Hill, Michael M. Swift, [Devirtualizing Memory for Heterogeneous Systems](#), *Proc. 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2018.
5. Lena E. Olson, Mark D. Hill, David A. Wood, [Crossing Guard: Mediating Host-Accelerator Coherence Interactions](#), *Proc. 22nd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 2017.
6. Sanketh Nalli, Swapnil Haria, Mark D. Hill, Michael M. Swift, Haris Volos, Kimberly Keeton, [An Analysis of Persistent Memory Use with WHISPER](#), *Proc. 22nd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, April 2017. Selected for an IEEE Micro Top Picks 2018 Honorable Mention.
7. Jayneel Gandhi, Mark D. Hill, and Michael M. Swift, [Agile Paging: Exceeding the Best of Nested and Shadow Paging](#), *43rd ACM/IEEE International Symposium on Computer Architecture (ISCA)*, June 2016. Selected as one of 2016 Top Picks in Computer Architecture for IEEE Micro, May-June 2017.
8. Vasileios Karakostas, Jayneel Gandhi, Adrian Cristal, Mark D. Hill, Kathryn S. McKinley, Mario Nemirovsky, Michael M. Swift, and Osman Unsal, [Energy-Efficient Address Translation](#), *Proc. 22nd IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, March 2016.
9. Lena E. Olson, Jason Power, Mark D. Hill, and David A. Wood, [Border Control: Sandboxing Accelerators](#), *Proc. 48th ACM International Symposium on Computer Microarchitecture (MICRO)*, December 2015.
10. Vasileios Karakostas, Jayneel Gandhi, Furkan Ayar, Adrian Cristal, Mark D. Hill, Kathryn S. McKinley, Mario Nemirovsky, Michael M. Swift, Osman Unsal, [Redundant Memory Mappings for Fast Access to Large Memories](#), *Proc. 42nd ACM/IEEE International Symposium on Computer Architecture (ISCA)*, June 2015. Selected as one of 2015 Top Picks in Computer Architecture for IEEE Micro, May-June 2016.
11. Marc S. Orr, Shuai Che, Ayse Yilmazer, Bradford M. Beckmann, Mark D. Hill, David A. Wood, [Synchronization Using Remote-Scope Promotion](#), *Proc. 20th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2015.
12. Jayneel Gandhi, Arkaprava Basu, Mark D. Hill, Michael M. Swift, [Efficient Memory Virtualization: Reducing Dimensionality of Nested Page Walks](#), *Proc. 47th ACM International Symposium on Computer Microarchitecture (MICRO)*, December 2014. IEEE Micro Top Picks Honorable Mention.
13. Derek R. Hower, Blake A. Hechtman, Bradford M. Beckmann, Benedict R. Gaster, Mark D. Hill, Steven K. Reinhardt, and David A. Wood, [Heterogeneous-race-free Memory Models](#), *Proc. 19th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, March 2014.
14. Jason Power, Mark D. Hill, and David A. Wood, [Supporting x86-64 Address Translation for 100s of GPU Lanes](#), *Proc. 20th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, February 2014.
15. Blake A. Hechtman, Shuai Che, Derek R. Hower, Yingying Tian, Bradford M. Beckmann, Mark D. Hill, Steven K. Reinhardt, and David A. Wood, [QuickRelease: A Throughput-oriented Approach to Release Consistency on GPUs](#), *Proc. 20th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, February 2014.
16. Jason Power, Arkaprava Basu, Junli Gu, Sooraj Puthoor, Bradford M Beckmann, Mark D Hill, Steven K Reinhardt, David A Wood, [Heterogeneous System Coherence for Integrated CPU-GPU Systems](#), *Proc. 46th ACM International Symposium on Computer Microarchitecture (MICRO)*, pp. TBD, December 2013.
17. Arkaprava Basu, Jayneel Gandhi, Jichuan Chang, Mark D. Hill, and Michael M. Swift, [Efficient Virtual Memory for Big Memory Servers](#), *Proc. 40th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 237-248, June 2013.
18. Arkaprava Basu, Mark D. Hill and Michael M. Swift, [Reducing Memory Reference Energy With Opportunistic Virtual Caching](#), *Proc. 39th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 297-308, June 2012.

-
19. Gabriel H. Loh and Mark D. Hill, [Efficiently Enabling Conventional Block Sizes for Very Large Die-stacked DRAM Caches](#), *Proc. 44th ACM International Symposium on Computer Microarchitecture (MICRO)*, pp. 70-78, December 2011. Selected as one of 2011 *Top Picks in Computer Architecture* for *IEEE Micro*, May/June 2012.
 20. Derek R. Hower, Polina Dudnik, David A. Wood, and Mark D. Hill, [Calvin: Deterministic or Not? Free Will to Choose](#), *Proc. 17th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 333-344, February 2011.
 21. Jayaram Bobba, Marc Lupon, Mark D. Hill, and David A. Wood, [Safe and Efficient Supervised Memory Systems](#), *Proc. 17th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 369-380, February 2011.
 22. Jayaram Bobba, Weiwei Xiong, Luke Yen, Mark D. Hill, and David A. Wood, [StealthTest: Low Overhead Online Software Testing using Transactional Memory](#), *Proc. ACM Conf. on Parallel Architectures and Compilation Techniques (PACT)*, pp. 146-155, September 2009.
 23. Luke Yen, Stark C. Draper, and Mark D. Hill, [Notary: Hardware Techniques to Enhance Signatures](#), *Proc. 41st ACM International Symposium on Computer Microarchitecture (MICRO)*, pp. 234-245, December 2008.
 24. Derek R Hower and Mark D. Hill, [Rerun: Exploiting Episodes for Lightweight Memory Race Recording](#), *Proc. 35th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 265-276, June 2008. Recommended to CACM for wider audience [CACM, pp. 93-100, June 2009].
 25. Jayaram Bobba, Neelam Goyal, Mark D. Hill, Michael M. Swift, and David A. Wood, [TokenTM: Efficient Execution of Large Transactions with Hardware Transactional Memory](#), *Proc. 35th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 127-138, June 2008.
 26. Daniel Sanchez, Luke Yen, Mark D. Hill, and Karthikeyan Sankaralingam, [Implementing Signatures for Transactional Memory](#), *Proc. 40th ACM International Symposium on Computer Microarchitecture (MICRO)*, pp. 123-133, December 2007.
 27. Jayaram Bobba, Kevin E. Moore, Haris Volos, Luke Yen, Mark D. Hill, Michael M. Swift, and David A. Wood, [Performance Pathologies in Hardware Transactional Memory](#), *Proc. 34th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 81-91, June 2007.
 28. Michael R. Marty and Mark D. Hill, [Virtual Hierarchies to Support Server Consolidation](#), *Proc. 34th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 46-56, June 2007.
 29. Luke Yen, Jayaram Bobba, Michael R. Marty, Kevin E. Moore, Haris Volos, Mark D. Hill, Michael M. Swift, and David A. Wood, [LogTM-SE: Decoupling Hardware Transactional Memory from Caches](#), *Proc. 13th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 261-272, February 2007.
 30. Michael R. Marty and Mark D. Hill, [Coherence Ordering for Ring-based Chip Multiprocessors](#), *Proc. 39th ACM International Symposium on Computer Microarchitecture (MICRO)*, pp. 309-320, December 2006.
 31. Min Xu, Rastislav Bodik, and Mark D. Hill, [A Regulated Transitive Reduction \(RTR\) for Longer Memory Race Recording](#), *Proc. 12th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 49-60, October 2006.
 32. Michelle J. Moravan, Jayaram Bobba, Kevin E. Moore, Luke Yen, Mark D. Hill, Ben Liblit, Michael M. Swift, and David A. Wood, [Supporting Nested Transactional Memory in LogTM](#), *Proc. 12th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 359-370, October 2006.
 33. Kevin E. Moore, Jayaram Bobba, Michelle J. Moravan, Mark D. Hill, and David A. Wood, [LogTM: Log-based Transactional Memory](#), *Proc. 12th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 258-269, February 2006.
 34. Min Xu, Rastislav Bodik, and Mark D. Hill, [A Serializability Violation Detector for Shared-Memory Server Programs](#), *Proc. ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pp. 1-14, June 2005.
 35. Michael R. Marty, Jesse D. Bingham, Mark D. Hill, Alan J. Hu, Milo M. K. Martin, and David A. Wood, [Improving Multiple-CMP Systems Using Token Coherence](#), *Proc. 11th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 328-339, February 2005.
-

36. Brian A. Fields, Rastislav Bodik, Mark D. Hill, and Chris J. Newburn, [Using Interaction Costs for Microarchitectural Bottleneck Analysis](#), *Proc. 36th ACM International Symposium on Computer Microarchitecture (MICRO)*, pp. 228-241, December 2003. Award paper finalist.
37. Milo M. K. Martin, Mark D. Hill, and David A. Wood, [Token Coherence: Decoupling Performance and Correctness](#), *Proc. 30th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 182-193, June 2003. Selected as one of *Top Picks in Computer Architecture* for *IEEE Micro*, November-December 2003.
38. Milo M. K. Martin, Pacia J. Harper, Daniel J. Sorin, Mark D. Hill, and David A. Wood, [Using Destination-Set Prediction to Improve the Latency/Bandwidth Tradeoff in Shared Memory Multiprocessors](#), *Proc. 30th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 206-217, June 2003.
1. Min Xu, Rastislav Bodik, and Mark D. Hill, [A “Flight Data Recorder” for Enabling Full-system Multiprocessor Deterministic Replay](#), *Proc. 30th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 122-133, June 2003. Selected for ISCA@50 Retrospective: 1996-2020. Edited by José F. Martínez and Lizy K. John, June 2023: [A Retrospective on: A “Flight Data Recorder” for Enabling Full-system Multiprocessor Deterministic Replay](#).
- 39.
40. Daniel J. Sorin, Mark D. Hill, and David A. Wood, [Dynamic Verification of End-to-End Multiprocessor Invariants](#), *IEEE International Conference of Dependable Systems and Networks (DSN, formerly FTCC)*, pp. 281-290, June 2003.
41. Carl J. Mauer, Mark D. Hill, and David A. Wood, [Full-System Timing-First Simulation](#), *Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pp. 108-116, June 2002.
42. Daniel J. Sorin, Milo M. K. Martin, Mark D. Hill, and David A. Wood, [SafetyNet: Improving the Availability of Shared Memory Multiprocessors with Global Checkpoint/Recovery](#), *Proc. 29th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 123-134, May 2002.
43. Brian Fields, Rastislav Bodik, and Mark D. Hill, [Slack: Maximizing Performance Under Technological Constraints](#), *Proc. 29th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 47-58, May 2002.
44. Milo M. K. Martin, Daniel J. Sorin, Mark D. Hill, and David A. Wood, [Bandwidth Adaptive Snooping](#), *Proc. 8th International Symposium on High Performance Computer Architecture (HPCA)*, pp. 251-262, February 2002.
45. Milo M. K. Martin, Daniel J. Sorin, Harold W. Cain, Mark D. Hill, and Mikko H. Lipasti, [Correctly Implementing Value Prediction in Microprocessors that Support Multithreading or Multiprocessing](#), *Proc. 34th ACM International Symposium on Computer Microarchitecture (MICRO)*, pp. 328-337, December 2001.
46. Anastassia G. Ailamaki, David J. DeWitt, Mark D. Hill, and Marios Skounakis, [Weaving Relations for Cache Performance](#), *Proc. 27th International Conference on Very Large Databases (VLDB)*, pp. 169-180, September 2001. Given VLDB 2001 *Best Paper Award*.
47. Eric C. Schnarr, Mark D. Hill, and James R. Larus, [Facile: A Language and Compiler for High-Performance Processor Simulators](#), *Proc. ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pp. 321-331, June 2001.
48. Milo M. K. Martin, Daniel J. Sorin, Anastassia Ailamaki, Alaa R. Alameldeen, Ross M. Dickson, Carl J. Mauer, Kevin E. Moore, Manoj Plakal, Mark D. Hill, and David A. Wood, [Timestamp Snooping: An Approach for Extending SMPs](#), *Proc. 9th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 25-36, November 2000.
49. Anastassia G. Ailamaki, David J. DeWitt, Mark D. Hill, and David A. Wood, [DBMSs on Modern Processors: Where Does Time Go?](#), *Proc. 25th International Conference on Very Large Databases (VLDB)*, pp. 266-277, September 1999.
50. Mark D. Hill, Anne E. Condon, Manoj Plakal, and Daniel J. Sorin, [A System-Level Specification Framework for I/O Architectures](#), *Proc. 11th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 138-147, June 1999.
51. E. Ender Bilir, Ross M. Dickson, Ying Hu, Manoj Plakal, Daniel J. Sorin, Mark D. Hill, and David A. Wood, [Multicast Snooping: A New Coherence Method Using a Multicast Address Network](#), *Proc. 26th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 294-304, May 1999.

52. Trishul M. Chilimbi, Mark D. Hill, and James R. Larus, [Cache-Conscious Structure Layout](#), *Proc. ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pp. 1-12, May 1999.
53. Anne E. Condon, Daniel J. Sorin, Manoj Plakal, and Mark D. Hill, [Using Lamport Clocks to Reason About Relaxed Memory Models](#), *Proc. 5th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 270-278, January 1999.
54. Ioannis Schoinas, Babak Falsafi, Mark D. Hill, James R. Larus, and David A. Wood, [Sirocco: Cost-Effective Fine-Grain Distributed Shared Memory](#), *Proc. Sixth International Conference on Parallel Architectures and Compilation Techniques (PACT 98)*, pp. 40-51, October 1998.
55. Manoj Plakal, Daniel J. Sorin, Anne E. Condon, and Mark D. Hill, [Lamport Clocks: Verifying A Directory Cache-Coherence Protocol](#), *Proc. 10th ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pp. 67-76, June 1998.
56. Shubhendu S. Mukherjee and Mark D. Hill, [Using Prediction to Accelerate Coherence Protocols](#), *Proc. 25th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 179-190, June 1998.
57. Shubhendu S. Mukherjee and Mark D. Hill, [The Impact of Data Transfer and Buffering Alternatives on Network Interface Design](#), *Proc. 4th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 207-218, February 1998.
58. Ioannis Schoinas and Mark D. Hill, [Address Translation Mechanisms in Network Interfaces](#), *Proc. 4th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pp. 219-230, February 1998.
59. Yuanyuan Zhou, Liviu Iftode, Jaswinder Pal Singh, Kai Li, Brian R. Toonen, Ioannis Schoinas, Mark D. Hill, and David A. Wood, [Relaxed Consistency and Coherence Granularity in DSM Systems: A Performance Evaluation](#), *Proc. 7th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pp. 193-205, June 1997.
60. Shubhendu S. Mukherjee, Babak Falsafi, Mark D. Hill, and David A. Wood, [Coherent Network Interfaces for Fine-Grain Communication](#), *Proc. 23rd ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 247-258, May 1996.
61. Madhusudhan Talluri, Mark D. Hill, and Yousef A. Khalidi, [A New Page Table for 64-bit Address Spaces](#), *Proc. 15th ACM Symposium on Operating Systems Principles (SOSP)*, pp. 184-200, December 1995.
62. Shubhendu Mukherjee, Shamik Sharma, Mark D. Hill, James Larus, Anne Rogers, and Joel Saltz, [Efficient Support for Irregular Applications on Distributed-Memory Machines](#), *Proc. 5th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP)*, pp. 68-79, July 1995.
63. Madhusudhan Talluri and Mark D. Hill, [Surpassing the TLB Performance of Superpages with Less Operating System Support](#), *Proc. 6th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 171-182, October 1994.
64. David A. Wood, Satish Chandra, Babak Falsafi, Mark D. Hill, James R. Larus, Alvin R. Lebeck, James C. Lewis, Shubhendu S. Mukherjee, Subbarao Palacharla, and Steven K. Reinhardt, [Mechanisms for Cooperative Shared Memory](#), *Proc. 20th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 156-167, May 1993. Reprinted in *Computer Measurement Group Transactions*, Issue 84, pp. 51-62, Spring 1994 (Special Issue on Parallel Architectures).
65. Steven K. Reinhardt, Mark D. Hill, James R. Larus, Alvin R. Lebeck, James C. Lewis, and David A. Wood, [The Wisconsin Wind Tunnel: Virtual Prototyping of Parallel Computers](#), *Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pp. 48-60, May 1993. Reprinted in *Multiprocessor Performance Measurement and Evaluation*, pp. 150-162, Laxmi Bhuyan and Xiaodong Zhang, editors, IEEE Computer Society Press, 1994.
66. Mark D. Hill, James R. Larus, Steven K. Reinhardt, and David A. Wood, [Cooperative Shared Memory: Software and Hardware for Scalable Multiprocessors](#), *Proc. 5th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pp. 262-273, October 1992. Selected as one of four papers reprinted in *ACM Transactions on Computer Systems*, Vol. 11, No. 4, pp. 300-318, November 1993.
67. Madhusudhan Talluri, Shing Kong, Mark D. Hill and David A. Patterson, [Tradeoffs in Supporting Two Page Sizes](#), *Proc. 19th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 415-424, May 1992.
68. Sarita V. Adve, Vikram S. Adve, Mark D. Hill, and Mary K. Vernon, [Comparison of Hardware and Software Cache Coherence Schemes](#), *Proc. 18th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 298-

- 308, May 1991. Reprinted in *The Cache Coherence Problem in Shared-Memory Multiprocessors: Hardware Solutions*, pp. 419-429, M. Tomasevic and V. Milutinovic, editors, IEEE Computer Society Press, 1993.
69. Sarita V. Adve, Mark D. Hill, Barton P. Miller, and Robert H. B. Netzer, [Detecting Data Races on Weak Memory Systems](#), *Proc. 18th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 234-243, May 1991.
 70. David A. Wood, Mark D. Hill, and R. E. Kessler, [A Model for Estimating Trace-Sample Miss Ratios](#), *Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pp. 79-89, May 1991.
 71. Sarita V. Adve and Mark D. Hill, [Weak Ordering - A New Definition](#), *Proc. 17th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 2-14, May 1990. Reprinted in *Selected Papers from the First 25 International Symposia on Computer Architecture*, Gurindar S. Sohi, editor, ACM Press, 1998. One of 41 papers.
 72. R. E. Kessler, Richard Jooss, Alvin Lebeck, and Mark D. Hill, [Inexpensive Implementations of Set-Associativity](#), *Proc. 16th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 131-139, June 1989.
 73. D. A. Wood, S. J. Eggers, G. Gibson, M. D. Hill, J. Pendleton, S. A. Ritchie, R. H. Katz, and D. A. Patterson, [An In-Cache Address Translation Mechanism](#), *Proc. 13th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 358-365, June 1986.
 74. Mark D. Hill and Alan J. Smith, [Experimental Evaluation of On-Chip Microprocessor Cache Memories](#), *Proc. 11th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 158-166, June 1984.
 75. D. A. Patterson, P. Garrison, M. D. Hill, D. Lioupis, C. Nyberg, T. N. Sippel, and K. S. Van Dyke, [Architecture of a VLSI Instruction Cache](#), *Proc. 10th ACM/IEEE International Symposium on Computer Architecture (ISCA)*, pp. 108-116, June 1983. Reprinted in *Tutorial on Advanced Microprocessors and High-Level Language Computer Architecture*, pp. 158-166, Veljko Milutinović, editor, IEEE Computer Society Press, 1986.

Other Refereed Conference or Workshop Proceedings.

1. Jason Power, Yinan Li, Mark D. Hill, Jignesh M. Patel, David A. Wood, [Toward GPUs being mainstream in analytic processing: An initial argument using simple scan-aggregate queries](#), *Proceedings of the Eleventh International Workshop on Data Management on New Hardware (DaMoN)*, 2015.
2. Arkaprava Basu, Derek R. Hower, Mark D. Hill, and Michael M. Swift, [FreshCache: Statically and Dynamically Exploiting Dataless Ways](#), *Proceedings of the International Conference on Computer Design (ICCD)*, 2013.
3. Derek R. Hower, Bradford M. Beckmann, Benedict R. Gaster, Blake A. Hechtman, Mark D. Hill, Steven K. Reinhardt, and David A. Wood, [Sequential Consistency for Heterogeneous-Race-Free: Programmer-centric Memory Models for Heterogeneous Platforms](#), *Workshop on Memory Systems Performance and Correctness (MSPC)*, June 2013. See revised version in ASPLOS 2014.
4. Arkaprava Basu, Jayaram Bobba, Mark D. Hill, [Karma: Scalable Deterministic Record-Replay](#), *25th International Conference on Supercomputing*, pp. 359-368, June 2011.
5. Michael M. Swift, Haris Volos, Neelam Goyal, Luke Yen, Mark D. Hill and David A Wood, [OS Support for Virtualizing Transactional Memory](#), *Third ACM SIGPLAN Workshop on Transactional Memory (TRANSACT)*, pp. 11-20, February 2008.
6. Daniel J. Sorin, Milo M. K. Martin, Mark D. Hill, and David A. Wood, [Using Speculation to Simplify Multiprocessor Design](#), *Proc. International Parallel and Distributed Processing Symposium (IPDPS)*, pp. 75-84, April 2004.
7. Alaa R. Alameldeen, Carl J. Mauer, Min Xu, Pacia J. Harper, Milo M. K. Martin, Daniel J. Sorin, Mark D. Hill, and David A. Wood, [Evaluating Non-deterministic Multi-threaded Commercial Workloads](#), *Proc. IEEE Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW-02 at HPCA)*, pp. 30-38, February 2002.
8. Sashikanth Chandrasekaran and Mark D. Hill, [Optimistic Simulation of Parallel Architectures Using Program Executables](#), *Proc. Workshop on Parallel and Distributed Simulation (PADS)*, pp. 143-150, May 1996.
9. Babak Falsafi, Alvin Lebeck, Steven Reinhardt, Ioannis Schoinas, Mark D. Hill, James Larus, Anne Rogers, and David Wood, [Application-Specific Protocols for User-Level Shared Memory](#), *Proc. Supercomputing '94*, pp. 380-389, November 1994.

10. Shubhendu S. Mukherjee and Mark D. Hill, [An Evaluation of Directory Protocols for Medium-Scale Shared-Memory Multiprocessors](#), *Proc. International Conference on Supercomputing (ICS)*, pp. 64-74, July 1994.
11. Sarita V. Adve and Mark D. Hill, [Implementing Sequential Consistency in Cache-Based Systems](#), *Proc. International Conference on Parallel Processing*, pp. I-47-50, August 1990.

Rigorously-Refereed Journal Publications

1. Mark D. Hill and Vijay Janapa Reddi, [Accelerator-Level Parallelism](#), *Communications of the ACM (CACM)*, to appear.
2. Todd Hylton, Thomas Conte, and Mark D. Hill, [A Vision to Computer Like Nature: Thermodynamically](#), *Communications of the ACM (CACM)*, Vol. 64, No. 6, pp. 35-38, June 2021.
3. Milo M. K. Martin, Mark D. Hill, and Daniel J. Sorin, [Why On-Chip Cache Coherence is Here to Stay](#), *Communications of the ACM (CACM)*, Vol. 55, No. 7, pp. 78-89, July 2012, DOI: 10.1145/2209249.2209269.
4. Brian A. Fields, Rastislav Bodik, Mark D. Hill, and Chris J. Newburn, [Interaction Cost and Shotgun Profiling](#), *ACM Transactions on Architecture and Compiler Optimizations (TACO)*, Vol. 1, No. 3, pp. 272-304, September 2004.
5. Anastassia G. Ailamaki, David J. DeWitt, and Mark D. Hill, [Data Page Layouts for Relational Databases on Deep Memory Hierarchies](#), *The VLDB Journal: The International Journal of Very Large Databases*, Springer, 11(3), pp. 198-215, 2002.
6. Daniel J. Sorin, Manoj Plakal, Anne E. Condon, Mark D. Hill, Milo M. K. Martin, David A. Wood, [Specifying and Verifying a Broadcast and a Multicast Snooping Cache Coherence Protocol](#), *IEEE Transactions on Parallel and Distributed Systems*, Vol. 13, No. 6, pp. 556-578, June 2002.
7. Frank Traenkle, Mark D. Hill, and Sangtae Kim, [Solving Microstructure Electrostatics on a Proposed Parallel Computer](#), *Computers and Chemical Engineering*, Special Issue on Parallel Computing, Vol. 19, pp. 743-757, 1995.
8. R. E. Kessler, Mark D. Hill, and David A. Wood, [A Comparison of Trace-Sampling Techniques for Multi-Megabyte Caches](#), *IEEE Transactions on Computers*, Vol. 43, No. 6, pp. 664-675, June 1994.
9. See [Cooperative Shared Memory: Software and Hardware for Scalable Multiprocessors](#), October 1992, under Rigorously-Refereed Conference Publications.
10. Sarita V. Adve and Mark D. Hill, [A Unified Formalization of Four Shared-Memory Models](#), *IEEE Transactions on Parallel and Distributed Systems*, Vol. 4, No. 6, pp. 613-624, June 1993.
11. Andreas Farid Pour and Mark D. Hill, [Performance Implications of Tolerating Cache Faults](#), *IEEE Transactions on Computers*, Vol. 42, No. 3, pp. 257-267, March 1993.
12. R. E. Kessler and Mark D. Hill, [Page Placement Algorithms for Large Real-Indexed Caches](#), *ACM Transactions on Computer Systems*, Vol. 10, No. 4, pp. 338-359, November 1992.
13. Kourosh Gharachorloo, Sarita V. Adve, Anoop Gupta, John L. Hennessy, and Mark D. Hill, [Programming for Different Memory Consistency Models](#), *Journal of Parallel and Distributed Computing*, Vol. 15, No. 4, pp. 399-407, August 1992.
14. Mark D. Hill and Alan Jay Smith, [Evaluating Associativity in CPU Caches](#), *IEEE Transactions on Computers*, Vol. 38, No. 12, pp. 1612-1630, December 1989. Reprinted in *Readings in Computer Architecture*, Mark D. Hill, Norman P. Jouppi, and Gurindar S. Sohi, editors, Morgan Kaufmann Publishers, 2000.
15. David D. Lee, Shing I. Kong, Mark D. Hill, George S. Taylor, David A. Hodges, Randy H. Katz, and David A. Patterson, [A VLSI Chip Set for a Multiprocessor Workstation - Part I: A RISC Microprocessor with Coprocessor Interface and Support for Symbolic Processing](#), *IEEE Journal of Solid State Circuits*, Vol. 24, No. 6, pp. 1688-1698, December 1989.

Other Refereed Journal Publications

1. Daniel S. Berger, Daniel Ernst, Huaicheng Li, Pantea Zardoshti, Monish Shah, Samir Rajadnya, Scott Lee, Lisa Hsu, Ishwar Agarwal, Mark D. Hill, and Ricardo Bianchini, [Design Tradeoffs in CXL-Based Memory Pools for Public Cloud Platforms](#), *IEEE Micro*, March/April 2023, DOI: [10.1109/MM.2023.3241586](#)

2. Mark D. Hill and Vijay Janapa Reddi, [Accelerator-Level Parallelism](#), *Communications of the ACM (CACM)*, Vol. 64, No. 12, December 2021, DOI: [10.1109/HPCA.2019.00047](#)
3. Mark D. Hill, [Why 'Correct' Computers Can Leak Your Information](#), *Communications of the ACM (CACM)*, Vol. 63, No. 7, July 2020, DOI: 10.1145/3399740. A one-page *Technical Perspective* introduction to a *Research Highlights* article by the original Spectre authors.
4. Jayneel Gandhi, Mark D. Hill, and Michael M. Swift, [Agile Paging for Efficient Memory Virtualization](#), *IEEE Micro Special Issue on Top Picks from Architectural Conferences*, May/June 2017. This article summarizes an ISCA2016 paper in about half as many words.
5. Mark D. Hill, Dave Christie, David Patterson, Joshua J. Yi, Derek Chiou, Resit Sendag, [Proprietary versus Open Instruction Sets](#), *IEEE Micro*, July/August 2016. Based on debate at CARD Workshop at ISCA 2015.
6. Jayneel Gandhi, Vasileios Karakostas, Furkan Ayar, Adrian Cristal, Mark D. Hill, Kathryn S. McKinley, Mario Nemirovsky, Michael M. Swift, Osman Unsal, [Range Translations for Fast Virtual Memory](#), *IEEE Micro Special Issue on Top Picks from Architectural Conferences*, May/June 2016. This article summarizes an ISCA2015 paper in about half as many words.
7. Lena E. Olson, Simha Sethumadhavan, Mark D. Hill, and David A. Wood, [Security Implications of Third-Party Accelerators](#), *IEEE Computer Architecture Letters (CAL)*, vol. 15. no. 1, Jan-Jun 2016.
8. Jason Power, Yinan Li, Mark D. Hill, Jignesh M. Patel, David A. Wood, [Supporting Very Large DRAM Caches with Compound Access Scheduling and MissMap](#), *SIGMOD Record*, Volume 44, Issue 1, April 2015.
9. Jason Power, Joel Hestness, Marc S. Orr, Mark D. Hill, and David A. Wood, [gem5-gpu: A Heterogeneous CPU-GPU Simulator](#), *IEEE Computer Architecture Letters (CAL)*, vol 13, no. 1. Jan-June 2014.
10. Gabriel H. Loh and Mark D. Hill, [Implications of Emerging 3D GPU Architecture on the Scan Primitive](#), *IEEE Micro Special Issue on Top Picks from Architectural Conferences*, May/June 2012. This article summarizes an MICRO2011 paper in about half as many words.
11. Derek R. Hower, Pablo Montesinos, Luis Ceze, Mark D. Hill, and Josep Torrellas, [Two Hardware-based Approaches for Deterministic Multiprocessor Replay](#), *Communications of the ACM (CACM)*, Vol. 52, No. 6, pp. 93-100, June 2009. Partially based on ISCA paper [ISCA 2008, pp. 265-276] recommend to CACM for wider audience.
12. Mark D. Hill and Michael R. Marty, [Amdahl's Law in the Multicore Era](#), *IEEE Computer*, Vol. 41, No. 7, pp. 33-38, July 2008. Honored as one of seven influential papers from *IEEE Computer's* first 50 years (<https://www.computer.org/computer-magazine/from-the-archives-computers-legacy/>).
13. Jayaram Bobba, Kevin E. Moore, Haris Volos, Luke Yen, Mark D. Hill, Michael M. Swift, and David A. Wood, [Performance Pathologies in Hardware Transactional Memory](#), *IEEE Micro Special Issue on Top Picks from Microarchitectural Conferences*, Vol. 28, No. 1, pp. 32-41, January-February 2008. This article summarizes an ISCA07 paper in about half as many words.
14. Michael R. Marty and Mark D. Hill, [Virtual Hierarchies](#), *IEEE Micro Special Issue on Top Picks from Microarchitectural Conferences*, Vol. 28, No. 1, pp. 99-109, January-February 2008. This article summarizes an ISCA07 paper in about half as many words.
15. Joel Emer, Mark D. Hill, Yale N. Patt, Joshua J. Yi, Derek Chiou, and Resit Sendag, [Single-Threaded vs. Multithreaded: Where Should We Focus?](#), *IEEE Micro Special Issue on Computer Architecture Debates*, Vol. 274, No. 6, pp. 14-23, November-December 2007.
16. Min Xu, Rastislav Bodik, and Mark D. Hill, [A Hardware Memory Race Recorder for Deterministic Replay](#), *IEEE Micro Special Issue on Top Picks from Microarchitectural Conferences*, Vol. 274, No. 1, pp. 48-55, January-February 2007. This article summarizes the corresponding ISCA03 and ASPLOS06 papers.
17. Brian A. Fields, Rastislav Bodik, Mark D. Hill, and Chris J. Newburn, [Interaction Cost: For when event counts just don't add up](#), *IEEE Micro Special Issue on Top Picks from Microarchitectural Conferences*, Vol. 24, No. 6, pp. 57-61, November-December 2004. This article summarizes the corresponding MICRO03 and TACO04 papers.
18. Milo M. K. Martin, Mark D. Hill, and David A. Wood, [Token Coherence: A New Framework for Shared-Memory Multiprocessors](#), *IEEE Micro Special Issue on Top Picks in Computer Architecture*, Vol. 23, No. 6, pp. 108-116, November-December 2003. This article summarizes the corresponding ISCA03 paper in about half as many words.
19. Kevin Skadron, Margaret Martonosi, David I. August, Mark D. Hill, David J. Lilja, and Vijay S. Pai, [Challenges in Computer Architecture Evaluation](#), *IEEE Computer*, Vol. 36, No. 8, pp. 30-36, August 2003.

20. Alaa R. Alameldeen, Milo M. K. Martin, Carl J. Mauer, Kevin E. Moore, Min Xu, Daniel J. Sorin, Mark D. Hill, and David A. Wood, [Simulating a \\$2M Commercial Server on a \\$2K PC](#), *IEEE Computer*, Vol. 36, No. 2, pp. 50-57, February 2003.
21. Trishul M. Chilimbi, Mark D. Hill, and James R. Larus, [Making Pointer-Based Data Structures Cache Conscious](#), *IEEE Computer*, Vol. 33, No. 12, pp. 67-74, December 2000.
22. Shubhendu S. Mukherjee, Steven K. Reinhardt, Babak Falsafi, Mike Litzkow, Steven Huss-Lederman, Mark D. Hill, James R. Larus, and David A. Wood, [Fast and Portable Parallel Architecture Simulators: Wisconsin Wind Tunnel II](#), *IEEE Concurrency*, Vol. 8, No. 4, pp. 12-20, October-December 2000.
23. Shubhendu S. Mukherjee and Mark D. Hill, [Making Network Interfaces Less Peripheral](#), *IEEE Computer*, Vol. 31, No. 10, pp. 70-76, October 1998.
24. Mark D. Hill, [Multiprocessors Should Support Simple Memory Consistency Models](#), *IEEE Computer*, Vol. 31, No. 8, pp. 28-34, August 1998.
25. David A. Wood and Mark D. Hill, [Cost-Effective Parallel Computing](#), *IEEE Computer*, Vol. 28, No. 2, pp. 69-72, February 1995.
26. Jeffrey D. Gee, Mark D. Hill, Dionisios N. Pnevmatikatos, and Alan Jay Smith, [Cache Performance of the SPEC92 Benchmark Suite](#), *IEEE Micro*, Vol. 13, No. 4, pp. 17-27, August 1993.
27. Mark D. Hill and James R. Larus, [Cache Considerations for Multiprocessor Programmers](#), *Communications of the ACM*, Vol. 33, No. 8, pp. 97-102, August 1990.
28. Mark D. Hill, [A Case for Direct-Mapped Caches](#), *IEEE Computer*, Vol. 21, No. 12, pp. 25-40, December 1988.
29. M. D. Hill, S. J. Eggers, J. R. Larus, G. S. Taylor, G. Adams, B. K. Bose, G. A. Gibson, P. M. Hansen, J. Keller, S. I. Kong, C. G. Lee, D. Lee, J. M. Pendleton, S. A. Ritchie, D. A. Wood, B. G. Zorn, P. N. Hilfinger, D. Hodges, R. H. Katz, J. Ousterhout, and D. A. Patterson, [Design Decisions in SPUR](#), *IEEE Computer*, Vol. 19, No. 11, pp. 8-22, November 1986. Reprinted in *Computers for Artificial Intelligence Processing*, pp. 273-299, B. W. Wah and C. V. Ramamoorthy, editors, John Wiley and Sons, Inc., New York, 1990. ***SPUR is now dubbed RISC IV, as the predecessor to open-source RISC V.***

In Collections

2. Min Xu, Rastislav Bodik, Mark D. Hill, [A Retrospective on: A “Flight Data Recorder” for Enabling Full-system Multiprocessor Deterministic Replay](#), In *ISCA@50 Retrospective: 1996-2020*. Edited by José F. Martínez and Lizy K. John, June 2023
3. Mark D. Hill, *Gene M. Amdahl 1922-2015*, Memorial Tributes, National Academy of Engineering, The National Academies Press, Washington DC, 2019, pp. 9-12, URL: http://pages.cs.wisc.edu/~markhill/papers/nae_memorial_tribute_amdahl_gene_2019.pdf
4. See *Evaluating Associativity in CPU Caches*, December 1989, under Rigorously-Refereed Journal Publications.
5. Sarita V. Adve and Mark D. Hill, [A Retrospective on “Weak Ordering - A New Definition”](#), In *Selected Papers from the First 25 International Symposia on Computer Architecture*, Gurindar S. Sohi, editor, ACM Press, pp. 63-66, 1998.
6. See *Weak Ordering – A New Definition*, May 1990, under Rigorously-Refereed Conference Publications.
7. See *The Wisconsin Wind Tunnel: Virtual Prototyping of Parallel Computers*, May 1993, under Rigorously-Refereed Conference Publications.
8. See *Comparison of Hardware and Software Cache Coherence Schemes*, May 1991, under Rigorously-Refereed Conference Publications.
9. Mark D. Hill, [What is Scalability?](#), pp. 89-96 in *Scalable Shared Memory Multiprocessors*, M. Dubois and S. S. Thakker, editors, Kluwer Academic Publishers, 1992. Also appeared in *ACM SIGARCH Computer Architecture News*, Vol. 18, No. 4, pp. 18-21, December 1990.
10. See *Design Decisions in SPUR*, November 1986, under Other Refereed Journal Publications.
11. See *Architecture of a VLSI Instruction Cache*, June 1983, under Rigorously-Refereed Conference Publications.

Other Publications (not appearing elsewhere)

1. Huaicheng Li, Daniel S. Berger, Stanko Novakovic, Lisa Hsu, Dan Ernst, Pantea Zardoshti, Monish Shah, Ishwar Agarwal, Mark D. Hill, Marcus Fontoura, and Ricardo Bianchini, *First-generation Memory Disaggregation for Cloud Platforms*, arXiv <https://arxiv.org/abs/2203.00241>, March 2022.
2. Mark D. Hill, *Reflections and Research Advice Upon Receiving the 2019 Eckert-Mauchly Award*, *IEEE Micro*, Vol. 39, No. 5, September/October 2019. This paper captures key aspect of EM acceptance speech at ISCA 2019.
3. Pratyush Mahapatra, Mark D. Hill, Michael M. Swift, *Don't Persist All : Efficient Persistent Data Structures*, arXiv <https://arxiv.org/abs/1905.13011>, May 2019.
4. Mark D. Hill, Jon Masters, Parthasarathy Ranganathan, Paul Turner, and John Hennessy, *On the Spectre and Meltdown Processor Security Vulnerabilities*, *IEEE Micro*, Vol. 39, No. 2, March/April 2019. The paper expands from a keynote/panel by the authors at IEEE Hot Chips 2018.
5. Mark D. Hill, *Three Other Models of Computer System Performance*, arXiv <https://arxiv.org/abs/1901.02926>, December 2018.
6. Mark D. Hill and Michael R. Marty, [Retrospective on Amdahl's Law in the Multicore Era](#), *IEEE Computer*, Vol. 50, No. 6, pp. 40-42, June 2017. Short invited paper for *Amdahl's Law in the Multicore Era* [2008] being honored as one of seven influential papers from *IEEE Computer's* first 50 years (<https://www.computer.org/computer-magazine/from-the-archives-computers-legacy/>).
7. Chris Feilbach, Adam Sperling, Eftychios Sifakis, Mark D. Hill, [Programming Heterogeneous Computers and Improving Inter-Node Communication Across Xeon Phi's](#), Univ. of Wisconsin Computer Sciences Technical Report CS-TR-2016-1834, May 2016.
8. Jason Power, Mark D. Hill, David A. Wood, [When to use 3D Die-Stacked Memory for Bandwidth-Constrained Big-Data Workloads](#), 7th Workshop on Big Data Benchmarks, Performance Optimization, and Emerging Hardware (BPOE 7) at ASPLOS, April 2016.
9. Lena E. Olson and Mark D. Hill, [Probabilistic Directed Writebacks for Exclusive Caches](#), Univ. of Wisconsin Computer Sciences Technical Report CS-TR-2016-1831, February 2016.
10. Jayneel Gandhi, Arkaprava Basu, Mark D. Hill, and Michael M. Swift, [BadgerTrap: A Tool to Instrument x86-64 TLB Misses](#), *ACM SIGARCH Computer Architecture News (CAN)*, 2014.
11. Arkaprava Basu, Bradford M. Beckmann, Steven K. Reinhardt, and Mark D. Hill, [CMP Directory Coherence: One Granularity Does Not Fit All](#), Univ. of Wisconsin Computer Sciences Technical Report CS-TR-2013-1798, June 2013.
12. Mark D. Hill, Sarita Adve, Luis Ceze, Mary Jane Irwin, David Kaeli, Margaret Martonosi, Josep Torrellas, David Wood, Katherine Yelick, et al., [21st Century Computer Architecture](#), *Computing Community Consortium (CCC) whitepaper*, May 25, 2012.
13. Nathan Binkert, Bradford Beckmann, Gabriel Black, Steven K. Reinhardt, Ali Saidi, Arkaprava Basu, Joel Hestness, Derek R. Hower, Tushar Krishna, Somayeh Sardashti, Rathijit Sen, Korey Sewell, Muhammad Shoaib, Nilay Vaish, Mark D. Hill, and David A. Wood, [The gem5 simulator](#), *ACM SIGARCH Computer Architecture News (CAN)*, 2011.
14. Josep Torrellas, Mark Oskin, et al., [Failure is not an Option: Popular Parallel Programming](#), *Report from Computing Community Consortium (CCC) Workshop on Advancing Computer Architecture Research (ACAR-1)*, released October 2010.
15. Mark D. Hill, Derek Hower, Kevin E. Moore, Michael M. Swift, Haris Volos and David A. Wood, [A Case for Deconstructing Hardware Transactional Memory Systems](#), *Dagstuhl Seminar Proceedings 07361*, editors Albert Cohen, Maria J. Garzaran, Christian Lengauer, and Samuel P. Midkiff, 2008.
16. Mark D. Hill, Jean-Luc Gaudiot, Mary Hall, Joe Marks, Paolo Prinetto, and Donna Baglio, [A Wiki for Discussing and Promoting Best Practices in Research](#), *Communications of the ACM*, Vol. 49, No. 9, pp. 63-64, September 2006.
17. Milo M.K. Martin, Daniel J. Sorin, Bradford M. Beckmann, Michael R. Marty, Min Xu, Alaa R. Alameldeen, Kevin E. Moore, Mark D. Hill, and David A. Wood, [Multifacet's General Execution-driven Multiprocessor Simulator \(GEMS\) Toolset](#), *ACM SIGARCH Computer Architecture News (CAN)*, Vol. 33, No. 4, pp. 92-99, September 2005.

18. Jason F. Cantin and Mark D. Hill, [Cache Performance for Selected SPEC CPU2000 Benchmarks](#), *ACM SIGARCH Computer Architecture News (CAN)*, Vol. 29, No. 4, pp. 13-18, September 2001.
19. Mark D. Hill, [Exploiting Market Realities to Address National Security's High-Performance Computing Needs](#), In *Defense Science Study Group 1998-1999, Volume 1: Papers 1-13*, Institute for Defense Analysis (IDA) Paper P-3531, 2000.
20. Geoffrey C. Orsak, Mark D. Hill, and Robin R. Murphy, [Formation of a National Entity for Information Infrastructure Protection](#), In *Defense Science Study Group 1998-1999, Volume 1: Papers 1-13*, Institute for Defense Analysis (IDA) Paper P-3531, 2000.
21. Andrew A. Chien, Mark D. Hill, and Shubhendu S. Mukherjee, [Guest Editor's Introduction: Design Challenges for High-Performance Network Interfaces](#), *IEEE Computer*, Vol. 31, No. 11, pp. 42-44, November 1998.
22. Shubhendu S. Mukherjee and Mark D. Hill, [A Survey of User-Level Network Interfaces for System Area Networks](#), Computer Sciences Technical Report 1340, University of Wisconsin, Madison, February 1997.
23. Mark D. Hill, [Bidirectional Technology Transfer: Sabbaticals in Industry](#), *Proc. NSF Conference on Experimental Research in Computer Systems*, pp. 201-203, June 1996. Reprinted in *Computing Research News*, Vol. 9, No. 5, November 1997.
24. Mark D. Hill, James Larus, and David Wood, [Parallel Computer Research in the Wisconsin Wind Tunnel Project](#), *Proc. NSF Conference on Experimental Research in Computer Systems*, pp. 135-145, June 1996.
25. Mark D. Hill, James R. Larus, and David A. Wood, "The Wisconsin Wind Tunnel Project," *IEEE Computer Society Technical Committee on Computer Architecture Newsletter*, Vol. 22, No. 5, pp. 12-15, August 1995.
26. Mark D. Hill, James Larus, and David Wood, [Tempest: A Substrate for Portable Parallel Programs](#), *Proc. IEEE Compcon*, pp. 327-332, March 1995. Reprinted in *Japanese Joint Symposium on Parallel Processing*, May 1995.
27. Mark D. Hill, James R. Larus, and David A. Wood, [The Wisconsin Wind Tunnel Project: An Annotated Bibliography](#), *ACM SIGARCH Computer Architecture News (CAN)*, Vol. 22, No. 5, pp. 19-26, December 1994.
28. Mark D. Hill, James R. Larus, Alvin R. Lebeck, Madhusudhan Talluri, and David A. Wood, [Wisconsin Architectural Research Tool Set \(WARTS\)](#), *ACM SIGARCH Computer Architecture News*, Vol. 21, No. 4, pp. 8-10, August 1993.
29. Mark D. Hill and David A. Wood, "Guest Editor's Introduction: Hot Chips II Symposium," *IEEE Micro*, Vol. 11, No. 3, pp. 8-9, June 1991.
30. Yul H. Kim, Mark D. Hill, and David A. Wood, [Implementing Stack Simulation for Highly-Associative Memories](#), (Extended Abstract), *Proc. ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, pp. 212-213, May 1991. Extended version appeared as Computer Sciences Technical Report 997, University of Wisconsin, Madison, February 1991.
31. Dionisios N. Pnevmatikatos and Mark D. Hill, [Cache Performance of the Integer SPEC Benchmarks on a RISC](#), *ACM SIGARCH Computer Architecture News (CAN)*, Vol. 18, No. 2, pp. 53-68, June 1990.
32. Mark D. Hill, [Test Driving Your Next Cache](#), *Magazine of Intelligent Personal Systems (MIPS)*, pp. 84-92, August 1989.

Software Distributed

1. **The gem5 Simulation Infrastructure (= GEMS + M5).** Hill co-led the development and release of the gem5 simulator (gem5.org) that merges the best aspects of the Wisconsin GEMS and Michigan M5 simulators. It provides a highly configurable full-system simulation framework, multiple ISAs (ARM, ALPHA, MIPS, Power, SPARC, and x86), diverse CPU models, and a detailed and flexible memory system, including support for multiple cache coherence protocols and interconnect models. The work seeks to make research dollars go further through cooperation, initially of universities Michigan, Princeton, MIT, Texas, and Wisconsin; and with companies AMD, ARM, HP, and MIPS. We expect gem5 componentized design and liberal BSD-like license will facilitate expanded future cooperation. over 5000 citations per Google Scholar. Time line:

2009-10: Initial GEMS and M5 merger talks and work.

2011: gem5 public release (gem5.org), boots Linux on ARM, ALPHA, and x86, CAN paper, and ISCA tutorial (slides: www.gem5.org/dist/tutorials/isca_pres_2011.pdf).

2014: gem5gpu public release for heterogeneous CPU-GPU simulator (http://www.cs.wisc.edu/multifacet/papers/cal14_gem5gpu.pdf).

2. **Multifacet GEMS (General Execution-driven Multiprocessor Simulator)** is a set of modules for Virtutech Simics that enables detailed simulation of multiprocessor systems, including Chip-Multiprocessors (CMPs). See <http://www.cs.wisc.edu/gems/> or September 2005 Computer Architecture News article. A 2005 ISCA tutorial introduced the community to GEMS, *which now has 2461 registered users, 391 users on the support mailing list, at least 86 non-UW publications, and 1865 citations*. Subsequent releases add:
 - 2006:** LogTM, making GEMS the first broadly available HTM simulation infrastructure.
 - 2007:** SMT, x86, and refined TM support.
 - 2008:** Four new multicore coherence protocols, commercial workload setup scripts, and support for Princeton’s on-chip interconnect models and Sun’s Adaptive Transactional Memory Test Platform.
 - 2009:** Default operation with Simics 3.x (but Simics 2.x still supported).
 - 2010:** Operates with Simics 4.x, but default is Simics 3.x with Simics 2.x supported.
 - 2011:** GEMS development slowed with gem5 release.
3. **WWW Computer Architecture Home Page** (<http://www.cs.wisc.edu/~arch/www>) serves as a clearinghouse for information pertaining to computer architecture. This web site was set up with Douglas Burger, enhanced by Milo Martin, and currently maintained by Min Xu. It averaged 100,000 hits per year (e.g., in 2000) with additional hits to mirrors in Europe, India, and Japan. Selected as the index for computer architecture by Yahoo (http://dir.yahoo.com/Science/Computer_Science/Architecture/) and in American Library Association’s *College & Research Libraries News* (<http://www.ala.org/acrl/resjune01.html>).
4. The **dineroIII** cache simulator has been distributed to more than one hundred universities, research labs, and companies. **Tycho**, a more sophisticated cache simulator, has been distributed to many of the same institutions. See Wisconsin Architectural Research Tool Set (WARTS) in August 1993 issue of Computer Architecture News or at URL <http://www.cs.wisc.edu/~larus/warts.html>.

Colloquia

Slides for many recent talks at <http://www.cs.wisc.edu/~markhill/includes/publications.html>

Driving Innovation in Academia and Industry, Keynote @ Computing Research Association-Industry (CRA-I) Workshop on Computing Research in Industry at ACM Federated Computing Research Conference (FCRC), June 2023.

In Computer Architecture, We Don’t Change the Questions, We Change the Answers

- Keynote at Intel Interconnect and Connectivity Summit, virtually, February 2024
- Keynote at Database Community Five-Year Meeting, MIT, Cambridge, MA, October 2023
- Keynote at Northwest Database Society Annual Meeting, Redmond, WA, May 2023
- University of Santa Cruz (UCSC), Santa Cruz, CA, March 2023
- Keynote at Data Management on New Hardware Workshop (DaMoN) @ SIGMOD, virtually, June 2022
- University of Texas, Austin, TX, March 2022
- U.C. Davis, May 2021
- Semiconductor Big Ideas Conference - New Street Research, September 2021
- Stanford SystemX Workshop, Stanford CA, June 2021.

Accelerator-level Parallelism

- U.C. Berkeley, CA, February 2022
- Computer Science & Automation Golden Jubilee Frontier Lecture, Indian Institute of Science, Bangalore, India (virtually), October 2020

- Technion (virtually), Haifa, Israel, June 2020
- Microsoft Research, Redmond, WA, February 2020
- Distinguished Lecture, Univ. of Illinois, Urbana-Champaign, IL, February 2020
- Keynote, IEEE International Symposium on Workload Characterization (IISWC), Orlando, FL, November 2019
- Harvard University, Cambridge, MA, October 2019
- Boston University, Boston, MA, October 2019
- Yale University, New Haven, CT, September 2019
- AMD, Sunnyvale, CA (850 people attended virtually), August 2019
- Keynote, FastPath Workshop on Performance Analysis of Machine Learning Systems, Madison, WI, March 2019
- University of Michigan, Ann Arbor, MI, March 2019

How Computing May Change Our World (A Talk for Non-Computer Scientists)

- Participatory Learning And Teaching Organization (PLATO) (virtually), September 2020
- Computing Research Association (CRA) (virtually), September 2020

Academic-Industrial Synergy: Stories, Pitfalls, & Advice, Keynote @ Google Computer Architecture & Deep Learning Workshop, August 2020.

Mentoring and Managing Graduate Students, CRA Mentoring Workshop (for new faculty), February 2020.

Exploiting modern microarchitectures, implications for computer architects, Hot Chips Keynote, Cupertino, CA, August 2018. Other parts of the joint keynote presented by John Hennessy, Paul Turner, and John Masters.

On the Meltdown & Spectre Design Flaws

- UPMARC Day Keynote, Uppsala University, Uppsala, Sweden, November 2018
- Wisconsin Alumni Club, Mountain View, CA, April 2018
- Google Geo, Mountain View CA, February 2018
- Google Platforms, Mountain View, CA, February 2018
- Google gChips, Mountain View, CA, January 2018

Efficient Memory Virtualization, University of Illinois, Urbana-Champaign, Illinois, November 2016.
Computer Architecture - 1975-2025, Gerald M. Masson Distinguished Lecture, Johns Hopkins University, Baltimore, MD, November 2016.

Technology, Computer Architecture, and Memory, Distinguished Lecture, Oklahoma State University, Stillwater, OK, September 2016.

Remember Memory?, Invited Talk, 40 Years of Patterson Symposium, Berkeley, CA, May 2016.

Whither Acoherent Shared Memory?, Invited Talk, Workshop on Negative Outcomes Post-mortems, and Experiences (NOPE), held in conjunction with MICRO, Honolulu, HI, December 2015.

Why On-chip Cache Coherence is Here to Stay

- AMD Research, Austin, TX, April 2015.
- Cornell University, Ithaca, New York, October 2012.

21st Century Computer Architecture

- Architecture 2030 Workshop @ ISCA 2016, Seoul, South Korea, June 2016
 - Keynote Address, International Conference on Supercomputing (ICS), Munich, Germany, June 2014.
-

- Joint Keynote Address, High Performance Computer Architecture (HPCA), ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP), and International Symposium on Code Generation and Optimization (CGO), Orlando, FL, February 2014.
- Distinguished Lecture, National Science Foundation, December 2013.
- Uppsala University, Uppsala, Sweden, September 2013.

Efficient Virtual Memory for Big Memory Servers, Advanced Micro Devices (AMD), Bellevue, WA, August 2013.

The Future of Computing Performance: Game Over or Next Level?

- Northwestern University, Distinguished Lecture, Evanston, IL, November 2011.
- University of Wisconsin, Madison, WI, June 2011.
- Advanced Micro Devices (AMD), Sunnyvale, CA, May 2011.
- University of Washington, Seattle, WA, May 2011.

Calvin: Deterministic or Not? Free Will to Choose, Dagstuhl Seminar, Dagstuhl, Germany, January 2011.

Amdahl's Law for the Multicore Era

- University of Southern California, Distinguished Lecture, Los Angeles, CA, February 2012.
- Princeton University, Princeton, New Jersey, November 2010.
- University of Maryland, College Park, Maryland, November 2010.
- Columbia University, New York, New York, November 2010.
- Epic Corporation, Verona, WI, February 2010.
- University of Washington, Seattle, WA, January 2010.
- University of Michigan, Ann Arbor, Michigan, August 2009.
- IBM Research, Yorktown Heights, New York, July 2009.
- Cornell University, Ithaca, New York, July 2009.
- Columbia University, New York, New York, July 2009.
- Salishan Conference on High-Speed Computing, Gleneden Beach, Oregon, April 2009.
- Carnegie-Mellon University, Pittsburgh, Pennsylvania, April 2009.
- Ecole Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, March 2009.
- University of California, Berkeley, California, February 2009.
- Google, Mountain View, California, February 2009.
- University of Auckland, Auckland, New Zealand, August 2008.
- University of Washington / Microsoft Research, Summer Institute on Making Parallel Programming Popular, Semiahmoo, WA, August 2008.
- IBM Research, Beijing, China, June 2008.
- Harvard University, Cambridge, MA, April 2008.
- Keynote Address, High Performance Computer Architecture (HPCA), Salt Lake City, Utah, February 2008.

Memory Consistency Models and Amdahl's Law for the Multicore Era, Uppsala Programming for Multicore Research Center (UPMARC) Summer School on Multicore Computing, Enköping, Sweden, June 2010

Is Transactional Memory an Oxymoron?

- Purdue University, West Lafayette, IN, September 2008.
- Keynote Address, Very Large Data Base Conference (VLDB), Auckland, New Zealand, August 2008.

A Case for Deconstructing Hardware Transactional Memory Systems, Dagstuhl Seminar on Parallel Programming Models, Dagstuhl, Germany, September 2007.

A Hardware Memory Race Recorder for Deterministic Replay, Intel Corp., Portland, Oregon, August 2007.

Performance Pathologies in Hardware Transactional Memory, Sun Microsystems, Mountain View, California, July 2007.

Virtualizing Hardware Transactional Memory with LogTM Signature Edition, Microsoft Research, Redmond, Washington February 2007.

LogTM: Log-based Transactional Memory

- Intel Corp., Marlborough, MA, April 2008.
- Distinguished Lecture, Texas A&M University, College Station, Texas, September 2007.
- VMware, Palo Alto, California, May 2007.
- HP Labs, Palo Alto, California, May 2007.
- Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts, November 2006.
- University of Texas at Austin, Austin, Texas, August 2006.
- Carnegie-Mellon University, Pittsburgh, Pennsylvania, May 2006.
- University of Pennsylvania, Philadelphia, Pennsylvania, January 2006.
- University of California, Berkeley, California, November 2005.
- Uppsala University, Uppsala, Sweden, November 2005.
- Virtutech Corp., Stockholm, Sweden, November 2005.
- Microsoft Research, Redmond, Washington, October 2005.

Designing Commercial Servers and the Flight Data Recorder

- Universidad Polit cnica de Catalu a, Barcelona, Spain, June 2005.
- Cornell University, Ithaca, New York, April 2005.
- Carnegie-Mellon University, Pittsburgh, Pennsylvania, January 2005.
- University of Illinois, Champaign, Illinois, October 2004.

Future Computer Advances are Between a Rock (Slow Memory) and a Hard Place (Multithreading).

- Nokia Research, Helsinki, Finland, November 2005.
- Universidad Politecnica de Valencia, Valencia, Spain, June 2005.
- National Science Foundation, Shared Cyberinfrastructure Division, Washington, DC, January 2005.
- The National Academy’s Computer Science and Telecommunication Board (CSTB) Meeting, Washington, DC, October, 2004.

Designing Commercial Servers and Token Coherence

- IBM Research, Yorktown Heights, New York, April 2005.
- University of Michigan, Ann Arbor, Michigan, September 2004.
- University of Texas, Austin, Texas, April 2004.
- Stanford University, Stanford, California, April 2004.

A Future for Parallel Computer Architectures, Keynote Talk at International Conference on Parallel Processing (ICPP), Montreal, Canada, August 2004.

Evaluating a \$2M Commercial Server on a \$2K PC, Invited Talk at Workshop on Computer Architecture Evaluation using Commercial Workloads (CAECW) at International Symposium on High-Performance Computer Architecture (HPCA), Madrid, Spain, February 2004.

Revisiting “Multiprocessors Should Use Simple Memory Consistency Models”, Dagstuhl Seminar on Consistency Models, Dagstuhl, Germany, October 2003.

Token Coherence: Enabling Faster Multiprocessors by Decoupling Performance and Correctness, Universidad Polit cnica de Catalu a, Barcelona, Spain, March 2003.

SafetyNet: Improving the Availability of Shared Memory Multiprocessors with Global Checkpoint/Recovery

- Universidad Politecnica de Valencia, Valencia, Spain, March 2003.
- University of Edinburgh, Edinburgh, Scotland, February 2003.
- IRISA / INRIA, Rennes, France, January 2003.

- Keynote address at Multithreaded Execution, Architecture, and Compilation Workshop at ACM International Symposium on Computer Microarchitecture (MICRO), Istanbul, Turkey, December 2002.
- Uppsala University, Uppsala, Sweden, December 2002.
- Universidad Politécnica de Cataluña, Barcelona, Spain, October 2002.
- Microsoft Labs, Redmond, Washington, March 2002.

Simulating a \$2M Commercial Server on a \$2K PC

- Universitat Autònoma de Barcelona, Barcelona, Spain, April 2003.
- Universidad de Murcia, Murcia, Spain, March 2003.
- Universidad de La Laguna, Tenerife, Spain, March 2003.
- Universidad Politécnica de Cataluña, Barcelona, Spain, January 2003.
- Université Paris Sud, Rennes, Orsay, January 2003.
- Virtutech Corp., Stockholm, Sweden, December 2002.
- Intel Labs, Barcelona, Spain, October 2002.

Harnessing Moore's Law

- Universidad de Las Palmas, Gran Canaria, Spain, March 2003.
- Universidad de La Laguna, Tenerife, Spain, March 2003.
- Beloit College, Beloit, WI, March 2002.

Wisconsin Multifacet: Commercial Server Design, Intel Labs, Barcelona, Spain, October 2002.

Timestamp Snooping: An Approach for Extending SMPs, Univ. of Rochester, Rochester, NY, March 2001.

Wisconsin Multifacet: Multicast Snooping and Beyond, Intel Corporation, Hillsboro, OR, July 2000.

How Computer Architecture Trends May Affect Future Distributed Systems, Keynote lecture at Principles of Distributed Computing (PODC), Portland, OR, July 2000.

Multicast Snooping: A New Coherence Method Using a Multicast Address Networks

- University of British Columbia, Vancouver, Canada, March 2000.
- University of Washington, Seattle, WA, March 2000.
- University of Toronto, February 2000.
- Distinguished lecture at University of Cincinnati, Cincinnati, OH, January 2000.
- Keynote address at ACM Symposium on Parallel Algorithms and Architectures (SPAA), St. Malo, France, July 1999.
- Sun Microsystems, Menlo Park, CA, April 1999.
- Compaq Computer Corp., Shrewsbury, MA, March 1999.
- Sun Microsystems, Burlington, MA, March 1999.
- University of Texas, Austin, TX, March 1999.

Exploiting Market Realities to Address National Security's High-Performance Computing Needs

- Parke-Davis Research Center, Ann Arbor, MI, February 2000.
- Center of Computing Sciences (CCS), Bowie, MD, January 2000.
- Institute for Defense Analysis (IDA), Alexandria, VA, November 1999.

Surfing Exponential Computer Growth, Plenary address at the Indian Science Congress, Hyderabad, India, January 1998.

Tempest: A Substrate for Portable Parallel Programs

- Northwestern University, November 1997.
- U.C. Berkeley, April 1996.
- DEC SRC, March 1996.
- U.C. San Diego, February 1996.

-
- Fujitsu, June 1995.
 - Japanese Joint Symposium on Parallel Processing (invited talk), June 1995.
 - Georgia Tech, November 1994.
- Cooperative Shared Memory and the Wisconsin Wind Tunnel*
- MIT, October 1993.
 - Univ. of Washington, August 1993.
 - Supercomputing Research Center, Bowie, MD, August 1993.
 - Argonne National Lab, July 1993.
 - Univ. of Illinois-Urbana-Champaign, April 1993.
 - U.C. Santa Barbara, February 1993.
 - DEC SRC, February 1993.
 - UCLA, January 1993.
 - USC, January 1993.
 - Workshop on Scalability of Parallel Algorithms and Architectures, Supercomputing, October 1992.
 - Stanford University, August 1992.
 - Cray Research, August 1992.
- The Wisconsin Wind Tunnel*, NSF Institutional Infrastructure Workshop, May 1992
- Memory (Consistency) Models*, Sun Microsystems Labs / SPARC International, January 1991.
- Beyond Sequential Consistency: A New Approach for Specifying Memory Models*
- IBM T.J. Watson Lab, November 1990.
 - Cray Research, October 1990.
 - University of California-Berkeley, October 1990.
 - DEC Western Research Lab, October 1990.
 - Hewlett-Packard, October 1990.
 - Sequent Computer Systems, August 1990.
- What is Scalability?*, Scalable Shared-Memory Workshop at International Symposium on Computer Architecture, May 1990.
- Weak Ordering A New Definition*
- Univ. of Illinois, February 1990.
 - Apple Computer, January 1990.
 - MIPS Co., January 1990.
 - Sun Microsystems, January 1990.
 - Scalable Interface Meeting, November 1989.
- Inexpensive Implementations of Set-Associativity*
- IBM Watson, April 1989.
 - U.C. Berkeley, January 1989.
 - DEC WRL, January 1989.
- Multicube: A Large-Scale Cache- MIMD*
- Sun Microsystems, October 1988.
 - U.C. Berkeley, May 1988.
 - DEC SRC, May 1988.
- A Case for Direct-Mapped Caches: Big and Dumb May Be Better*
- Brown University, April 1987.
-

- Cornell University, April 1987.
- UCLA, April 1987.
- CMU, April 1987.
- MIT, March 1987.
- University of Michigan, March 1987.
- University of North Carolina, March 1987.
- University of Toronto, March 1987.
- Princeton University, March 1987.
- University of Illinois-Urbana-Champaign, March 1987.
- University of Wisconsin, March 1987.
- Stanford University, March 1987.
- University of Washington, February 1987.